

Orbitofrontal Cortex Value Signals Depend on Fixation Location during Free Viewing

Highlights

- In free-viewing monkeys, orbitofrontal neurons signal the distance of gaze from a cue
- The distance signal is nearly as strong as the signal representing cue value
- In some cells, value signals increase when subjects fixate on the cue
- Representation of gaze distance persists across two distinct task phases

Authors

Vincent B. McGinty, Antonio Rangel, William T. Newsome

Correspondence

vmcginty@stanford.edu

In Brief

McGinty et al. show that value signals in primate orbitofrontal cortex are driven by moment-to-moment changes in gaze location during natural free viewing. This surprisingly dynamic value code may play a key role in everyday decision-making and other motivated behaviors.



Orbitofrontal Cortex Value Signals Depend on Fixation Location during Free Viewing

Vincent B. McGinty,^{1,*} Antonio Rangel,² and William T. Newsome^{1,3}

¹Department of Neurobiology, Stanford University, Stanford, CA 94305, USA

²Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125, USA

³Howard Hughes Medical Institute, Stanford University, Stanford, CA 94305, USA

*Correspondence: vmcginty@stanford.edu

<http://dx.doi.org/10.1016/j.neuron.2016.04.045>

SUMMARY

In the natural world, monkeys and humans judge the economic value of numerous competing stimuli by moving their gaze from one object to another, in a rapid series of eye movements. This suggests that the primate brain processes value serially, and that value-coding neurons may be modulated by changes in gaze. To test this hypothesis, we presented monkeys with value-associated visual cues and took the unusual step of allowing unrestricted free viewing while we recorded neurons in the orbitofrontal cortex (OFC). By leveraging natural gaze patterns, we found that a large proportion of OFC cells encode gaze location and, that in some cells, value coding is amplified when subjects fixate near the cue. These findings provide the first cellular-level mechanism for previously documented behavioral effects of gaze on valuation and suggest a major role for gaze in neural mechanisms of valuation and decision-making under ecologically realistic conditions.

INTRODUCTION

One of the most important tasks that an organism performs is judging the economic value—the potential for reward or punishment—associated with the stimuli in its environment. This is a difficult task in natural settings, in which many stimuli need to be accurately evaluated. One way that organisms address this problem is by evaluating stimuli serially. In primates, this is done through saccadic eye movements: by shifting gaze between objects, primates can focus their perceptual and cognitive resources on one stimulus at a time (Bichot et al., 2005; DiCarlo and Maunsell, 2000; Mazer and Gallant, 2003; Sheinberg and Logothetis, 2001). A logical hypothesis, therefore, is that when primates judge the value of visual objects in natural settings, they recruit their valuation circuitry in a serial fashion, according to the location of gaze. Furthermore, this suggests that to understand ecologically realistic decisions in primates, it is critical to understand how neural valuation circuitry is influenced by gaze. While several neural mechanisms exist for encoding the value of visible objects (Kennerley et al., 2011; Padoa-Schioppa

and Assad, 2006; Platt and Glimcher, 1999; Roesch and Olson, 2004; Thorpe et al., 1983; Yasuda et al., 2012), little is known about how value-coding neurons modulate their firing when subjects move their gaze from one object to the next, as value signals in primates are usually measured in the near-absence of eye movements. In fact, most primate behavioral tasks suppress natural eye movements by requiring prolonged fixation of gaze at a single location. And in cases where gaze was not subject to strict control, there has been no analysis—or even discussion—of how value signals might relate to gaze (e.g., Bouret and Richmond, 2010; Strait et al., 2014; Thorpe et al., 1983; Tremblay and Schultz, 1999).

In contrast, several recent behavioral studies in humans have shown that simple economic decisions are influenced by fluctuations in gaze location during the choice process (Armell et al., 2008; Kraglich et al., 2012, 2010; Kraglich and Rangel, 2011; Shimojo et al., 2003; Towal et al., 2013; Vaidya and Fellows, 2015). Specifically, subjects are more likely to choose an item if they fixate on that item longer than the alternatives. While the underlying neural mechanism is unknown, computational models suggest that the effects of fixation on choice are best explained by a value signal that is modulated by the location of gaze. In these models, choices are made by comparing and sequentially integrating over time the instantaneous value of the available items. As subjects shift their fixation between items, at any given instant the value of the fixated item is amplified relative to the unfixated ones, biasing the integration process in its favor, producing a choice bias for the items fixated longer overall. Using fMRI in humans, Lim et al. (2011) tested this hypothesis by asking if changes in fixation target could modulate the decision value signals in ventromedial prefrontal cortex (vmPFC, see Basten et al., 2010; Hare et al., 2009; Kable and Glimcher, 2007). They found that vmPFC value signals were positively correlated with the value of the currently fixated object and negatively correlated with the unfixated object's value.

While these results suggest that value signals are modulated by gaze, they leave many open questions, which the current study begins to address. First, with its limited spatial and temporal resolution, fMRI cannot show whether gaze modulates value signals at the natural functional unit of the nervous system (single neurons) and at the millisecond timescale of natural free viewing. Second, the gaze studies discussed above have focused on binary choice situations, yet it is possible that gaze modulates value signals in any situation in which value is relevant, not only when facing an explicit choice. Third, an effect of gaze on

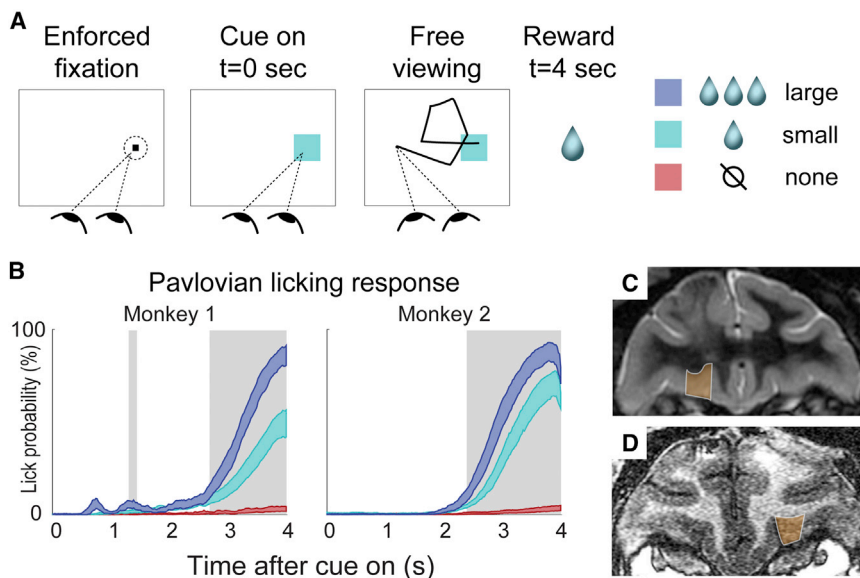


Figure 1. Task, Behavior, and Recording Sites

(A) Trial structure of the task. The FP was shown on the left or right side of the screen (randomized across trials), and the Pavlovian cue was shown at the location of the FP on each trial, after 1–1.5 s of enforced fixation. The cue colors indicated reward volume, and new cue colors were used in every session.

(B) 95% confidence intervals of the mean licking response in 50 sessions for monkey 1 and 28 sessions for monkey 2. This graph shows data from trials performed during neural data collection, after successful cue-reward learning (see [Experimental Procedures](#)). The shaded areas indicate a significant difference in licking between all three cues (Wilcoxon rank-sum tests, $p < 0.05$, corrected).

(C and D) Coronal MRI sections from monkey 1 (C) and monkey 2 (D), respectively. The orange indicates OFC recording area.

value representations has not been demonstrated in non-human primates.

To address these questions simultaneously, we use a behavioral task in which monkeys viewed reward-associated visual cues with no eye movement restrictions (free viewing), while we recorded single and multi-unit neural activity in a region known to express robust value signals for visual objects, the orbitofrontal cortex (OFC) (Abe and Lee, 2011; Morrison and Salzman, 2009; Padoa-Schioppa and Assad, 2006; Roesch and Olson, 2004; Rolls, 2015; Tremblay and Schultz, 1999; Wallis and Miller, 2003). This task explicitly manipulates the expectation of value, but, unconventionally, allows for natural gaze behavior, producing rich variation in gaze location that we then exploit to assess the effect of gaze on value coding. Importantly, reward delivery in the task did not depend on gaze behavior, meaning that any effects of gaze on value-related neural activity was not confounded by the operant demands of the task.

We found strong modulation of value coding by gaze, including cells in which value signals became amplified as the fixation drew close to the cue. Overall, the encoding of fixation location was nearly as strong as the encoding of value in the OFC population, a surprising observation given the predominance of value-coding accounts of OFC in the literature (Rolls, 2015). Taken together, these findings provide, (1) novel insight into the dynamic coding of value during free viewing, (2) evidence for a key element of computational models that account for the effects of fixation on behavioral choice, and (3) a link between the dynamics of frontal lobe value signals at the areal and cellular levels (human fMRI and monkey electrophysiology, respectively).

RESULTS

Every experimental session had two phases: an initial conditioning phase followed by neural data collection. During the conditioning phase, monkeys were trained to associate three distinct

color cues with three different volumes of juice (large = ~3 drops, small = 1 drop, and no reward = 0 drops), using a form of Pavlovian conditioning (Figure 1). New, randomly chosen colors were used in every session, and the monkeys performed the conditioning trials until they had learned the cue-reward association, indicated by different licking responses for all three cues (see [Experimental Procedures](#)). Figure 1B shows licking responses after conditioning was complete. Neurons were isolated and data collection began immediately after behavioral conditioning. The trial structure during both conditioning and data collection was identical (Figure 1A): Each trial began with a brief period of enforced gaze upon a fixation point (FP) placed at one of two locations in the screen. Then, a conditioned cue (chosen randomly) appeared at the center of the fixation window and fixation control was immediately released, allowing the monkey to move his gaze for the 4 s duration of the trial. At 4 s after cue onset, the predicted amount of juice (if any) was delivered and the trial ended. Eye position was monitored during the trial, but had no consequence on the outcome.

Importantly, we only collected spike data after the cue-reward associations had been learned and licking behavior had stabilized. Thus, none of the results described here include data from the initial conditioning phase.

Fixation Patterns

Figure 2 shows the location of fixations during the free-viewing period, using data from 0.5 to 3.75 s after cue onset. We focus on this period to minimize the impact of the cue onset or the juice delivery on the analysis. Figure 2A depicts the eye position in a sample trial, which starts at the cue location, visits several locations on the screen (including the cue), and ends at the cue just before juice delivery. Figure 2B illustrates the distribution of fixations across the study. Both animals were more likely to fixate on the cue than anywhere else, regardless of cue identity (frequency at center grid square was significantly greater than at the square with next-highest frequency,

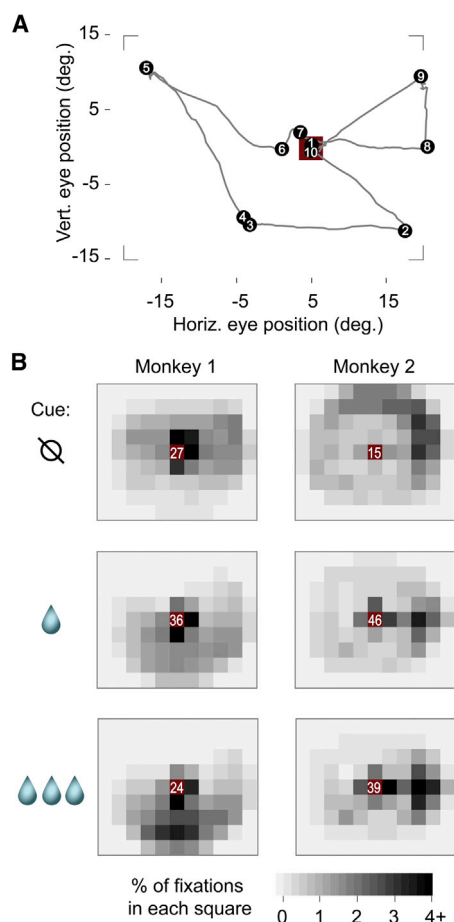


Figure 2. Fixation Locations during Free Viewing

(A) Eye position trace from a sample trial. The black dots indicate fixations, the numbers show fixation order, and the red square is the cue location (3.2×3.2 degrees). The corners show the approximate screen border.

(B) Spatial distribution of fixations averaged across all sessions, using only trials performed during neural data collection. Each small square is 5×5 degrees. The red squares at the center contain the cue, and white numbers give the percentage of fixations within that square (i.e., fixations on or very near the cue). Outside of the center, the fixation percentages are given by the gray scale.

$p < 8 \times 10^{-7}$ by Wilcoxon rank-sum test, for all six images in Figure 2B). Importantly, this is not driven by the fact that the cue is first presented at the initial FP, because the subjects almost always moved their eyes away from the cue within the first 0.5 s after cue onset (see Figure 3A for an example), and all data before 0.5 s were excluded from this analysis. Qualitatively, Figure 2B also shows that non-cue fixations were distributed widely, but tended to fall below the cue location for monkey 1 and toward the right edge of the monitor for monkey 2. Finally, the average likelihood of fixating the cue was not monotonically related to the size of the juice reward: the cue indicating a small reward was fixated more often than either the large or no reward cues by both monkey 1 ($p < 9 \times 10^{-7}$ for both small versus no reward and small versus large, Wilcoxon rank-sum test) and monkey 2 ($p < 2 \times 10^{-5}$ for both comparisons).

Together, these results show that the fixations fluctuated widely across the screen, which is necessary for the analyses below.

OFC Neurons Encode the Distance of Gaze from a Cue

We leveraged the rich, natural variability in fixation location to address the key question of our study: how fixation location influenced value signals in OFC neurons. We recorded from 176 single neurons and 107 multi-unit signals (total 283, see [Experimental Procedures](#); [Figures 1C](#) and [1D](#)). When discussing individual neural responses, we use the terms “single unit” and “multi-unit signal” as appropriate; when referring to group-level data, we use the terms “cells” or “neurons,” which encompass both single unit and multi-unit signals. Below, we present examples of the main form of gaze modulation in the OFC, which is the encoding of fixation distance from the cue. Then, we show that this distance signal is widespread at the population level and unlikely to be due to encoding of other gaze-related variables. Finally, we show that gaze distance and value signals overlap in many cells, including in a subset of cells with value signals that are greatest when subjects fixate on the cue.

Figure 3A depicts eye position and the firing of an identified single unit over one trial. In this trial, the cell fired more after fixations near the cue (gray bars in raster below x axis), but fired less following fixations away from the cue. Critically, this location-dependent modulation was strong when the “no reward” cue was shown (Figure 3B, top rows), but was weak or absent when the “small” or “large” reward was shown (Figure 3B, middle and bottom rows). Average firing rates time-locked to cue onset in each trial (Figure 3C) show that this cell fires most for the no reward cue and least for large, throughout nearly the whole trial (Figure 3C). However, trial-averaged data obscure the effect of fixation location on firing, because fixation patterns were unique in every trial. In Figure 3D, therefore, we replot these data to show firing time-locked to fixation onset, using fixations that began between 0.50 and 3.75 s after cue onset (Figure 3C, “fixations eligible for analysis,” see [Experimental Procedures](#)). Referenced to fixation onset, activity is clearly modulated by both value and fixation location: fixations near the no reward cue were followed by a burst of firing, but fixations onto the other cues elicited little or no response (Figure 3D, left). Thus, the cell’s value code—its differentiation between the cues—depends on fixation distance: it is strong following fixations onto the cue, but is weak following fixations away (Figure 3D, right).

To summarize value and fixation distance encoding in this single unit, we segmented the eye position data into saccade and fixation epochs, extracted the firing in a 200 ms window following the onset of each fixation (see [Experimental Procedures](#); [Figure S1](#)), and plotted this “fixation-evoked” firing as a function of cue value and the distance of fixation from the cue. This plot, in Figure 4A, shows the interaction between value and distance encoding exhibited by this cell: value coding is maximal when fixations land near the cue.

Figures 4B–4D show three additional examples of value and fixation encoding, with patterns distinct from the cell in Figure 4A. The multi-unit signal in Figure 4B also encodes an interaction between value and location, but with the opposite effect from the

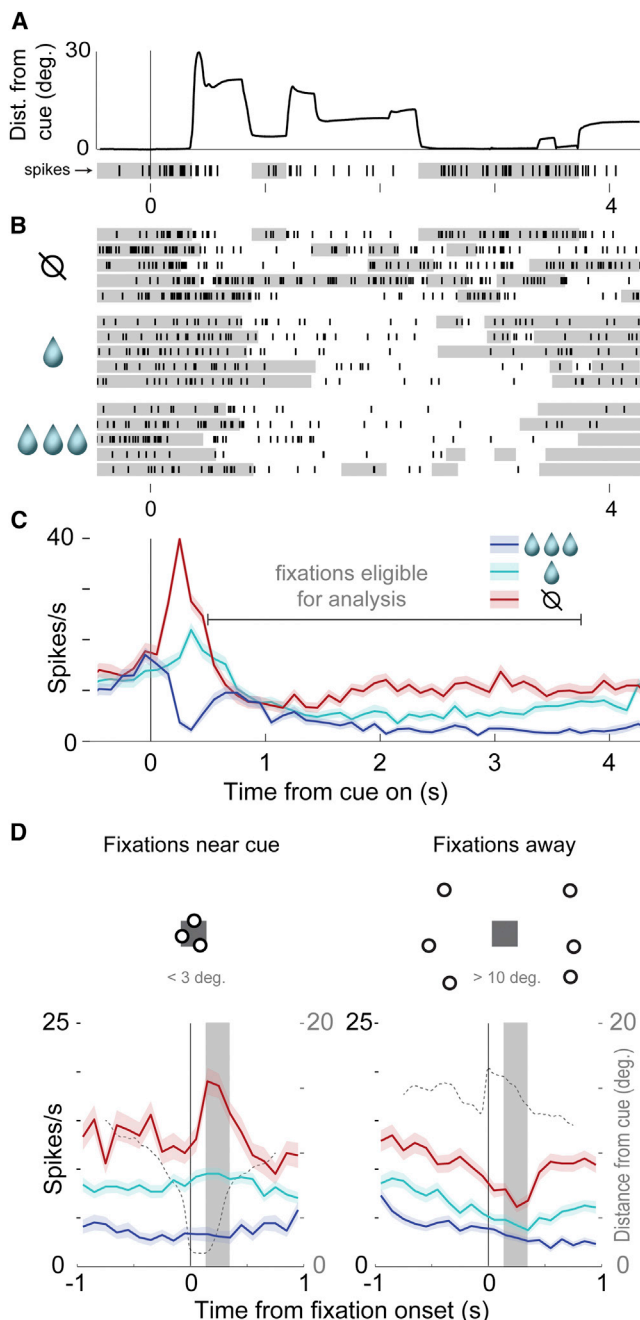


Figure 3. Value and Fixation Location Encoding in a Single Neuron
All images show the same single neuron. In (A)–(C), data are aligned to cue onset at $t = 0$ s and continue through reward delivery at $t = 4$ s. (A) Eye position and neural data in a single trial: the thick black line gives the distance of gaze from the cue, and the raster with black tick marks (below x axis) shows the spikes of a single cell. The gray shading in the raster shows when the eyes were within 5 degrees of the cue center. (B) Rasters showing spiking on 15 trials for three cues indicating different reward volumes. The top raster line is the trial in (A). The gray shading indicates eyes <5 degrees from cue center. (C) Average firing across all trials; shaded area shows SEM. The gray horizontal line shows the time range used for subsequent analyses (Experimental Procedures).

cell in Figure 4A: firing does not distinguish between the cues following on-cue fixations, but does following fixations away. The multi-unit signal in Figure 4C and single unit in Figure 4D encode both value and fixation location, but in an additive, not interactive, manner. Both distinguish between the cues, and at the same time, they modulate their overall activity level depending on gaze: one fires more overall for near-to-cue fixations (Figure 4C), while the other fires more for fixations away (Figure 4D). In contrast to these four examples with both value and fixation location effects, cells that only encode either value or distance alone yield very different firing patterns, illustrated in Figures 5H–5J.

Value and Fixation Distance Encoding Is Mixed in the Population

We now look beyond individual examples to the population of recorded neurons ($n = 283$) to ask how often OFC cells encode fixation location, especially in comparison to the value signals for which this region is known, and to ask how often both value and location are encoded by individual cells. We fit for every cell (single unit or multi-unit signal) a linear model that explains firing as a function of three variables: cue value, distance of fixation from the cue, and the value-by-distance interaction (here, cue value is its associated reward volume, because the cue-reward association as well as reward timing and probability remain constant throughout the session). We then ask how many cells show statistically significant effects of value or gaze alone, and, critically, how many show both effects, or show an interaction between value and gaze (as in Figures 3, 4A, and 4B). We also address the encoding of gaze-related variables other than fixation distance.

Distance Encoding Is Abundant

Table 1 shows the percentage of OFC neurons with significant effects of cue value, fixation distance from the cue, and the value-by-distance interaction in the generalized linear models (GLM). Many cells had significant effects of cue value, consistent with prior observations (Morrison and Salzman, 2009; Rolls, 2015). Critically, nearly as many were significantly modulated by fixation distance or the value-by-distance interaction. The number of cells encoding these fixation variables reached its maximum around 200 ms after fixation onset, consistent with the typical visual response latency in OFC (Figure S2). For all variables, the number of neurons with significant effects far exceeded chance levels ($p < 0.001$), established by a permutation test (Table 1, right column). In Table S1, we further explore these results, showing: (1) results were similar for the two subjects, with the exception that monkey 2 had fewer significant effects of value and gaze distance. (2) Results were similar for single and multi-unit signals. And (3) results were similar when the GLM was performed using the same post-fixation firing window

(D) The left peri-stimulus time histogram (PSTH) shows average firing time-locked to fixations near the cues (<3 degrees), and the right shows firing time-locked to fixations away (>10 degrees); shaded areas show SEM. The dotted line in each PSTH indicates average eye position over all trials (right axis scale), and the solid gray box indicates the post-fixation analysis window used for this neuron to generate Figure 4A. See also Figure S1. The dots and squares above the PSTHs are illustrations, not the actual fixation data.

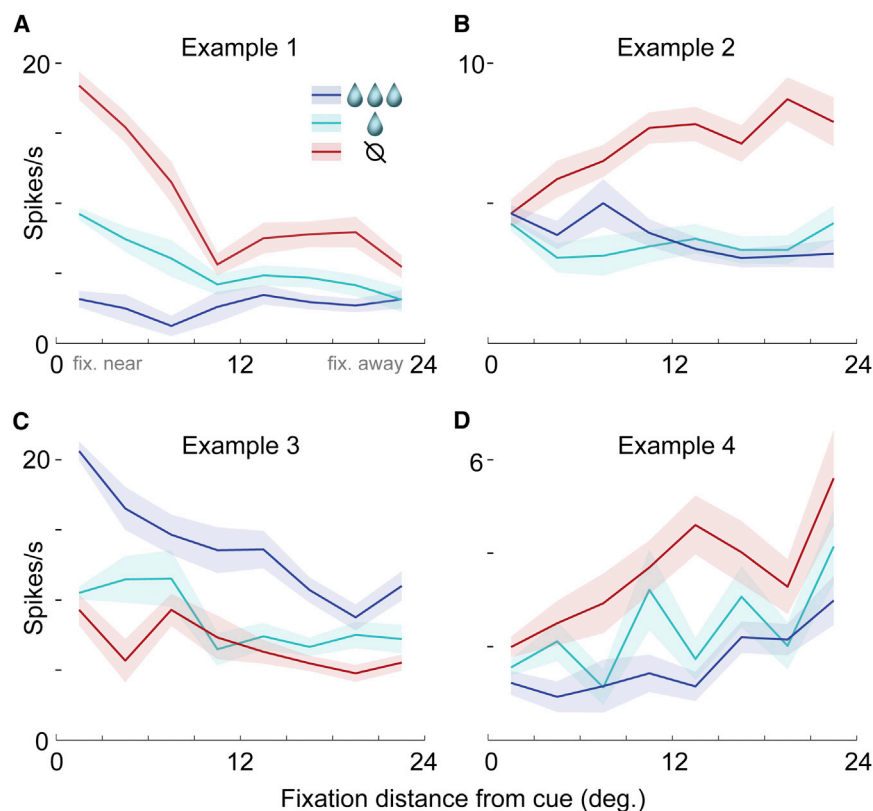


Figure 4. Four Examples of Value and Fixation Location Encoding

Each image shows data from a different single unit (A and D) or multi-unit signal (B and C).

(A) Firing as a function of cue value (colors) and the distance of fixation from the cue (x axis, 3 degree bins). The lines indicate mean and the shaded areas indicate SEM. The firing was measured in a 200 ms window following the onset of each fixation; this single unit is the same as in Figure 3, and its 200 ms window is shown by the gray boxes in Figure 3D.

(B–D) Three additional examples of firing modulated by both cue value and fixation distance. The 200 ms firing windows were defined individually for each example (see Experimental Procedures).

in all cells (rather than the cell-specific windows used in the main analysis, see Experimental Procedures).

We noted that some cells began encoding cue value within ~100 ms of cue onset (Figure 3C). This early value encoding is perhaps the best characterized response in prior OFC studies (Morrison and Salzman, 2009; Padoa-Schioppa and Assad, 2006; Roesch and Olson, 2004; Thorpe et al., 1983; Tremblay and Schultz, 1999; Wallis and Miller, 2003). We asked how similar this “early” value signal was to the value signal measured in the main GLM, which used only data 0.5 s post-cue. First, we measured firing 50–500 ms after cue onset and fit a linear model with cue value as the only variable; 32% of neurons had a significant effect of value ($p < 0.05$, corrected). We then compared the beta co-efficients for value from this early GLM to the co-efficients for value obtained during the cue viewing period (0.5–3.75 s), using Spearman’s ρ , a rank-based statistic that is resistant to the effects of outliers. The correlation was $\rho = 0.400$ ($p < 10^{-10}$), suggesting a population value code that is similar, but not identical, at these two time points. This may reflect an effect of novelty in the cue onset response: cue onset entails the arrival of a new visual stimulus and updated reward expectation, whereas these two factors do not change during the time over which the main GLM is estimated.

To determine how cue color— independent of cue value— influenced firing, in some sessions, we dissociated color from value by abruptly reversing the color-value associations of the large and no reward cues ($n = 109$ neurons). Consistent with other OFC recordings (Morrison and Salzman, 2009; Thorpe et al., 1983; Wallis and Miller, 2003), only a small minority of neu-

rons encoded only cue color and no other value-related variable: 2.8% with significant effects in a GLM, at $p < 0.05$, corrected. See Figure S3 for details.

Importantly, we tested several alternative hypotheses, none of which could explain the abundant encoding of gaze distance revealed in the main GLM. First, we determined that the GLM results in Table 1 were not attributable to oculomotor variables other than gaze distance, such as saccade velocity (Figure S4) or saccade direction (as in Bruce and Goldberg, 1985; see legend of Figure S1).

Next, we asked whether a different form of eye position encoding could better explain the effects of gaze distance. To do so, we fit two additional models, a “gaze angle” model (Equation 2) that describes eye position in terms of the absolute angle of gaze in head-centered coordinates, which differs from the cue distance because the cue appears randomly on the left or right side of the screen center in each trial; and a “gradient model” (Equation 3) that describes eye position in terms of horizontal and vertical coordinates. We then identified all cells that had significant effects ($p < 0.05$, corrected) of the eye position variables in any of the three models (Equations 1–3). Grouping these cells together (34.5% of the population), we then asked which of the three models provided the best fit for each cell, indicated by the Akaike’s information criterion (AIC). A large majority of cells in this group, 66.7%, were better fit by the cue distance model, whereas only 23.1% and 10.3% were better fit by gaze angle and the gradient model, respectively. This provides strong support for cue distance—our primary hypothesis—as being the main mode of gaze modulation in OFC neurons and suggests that only a small minority of cells may incorporate gaze angle or gradient coding signals.

While OFC cells can encode the location of visible targets (Abe and Lee, 2011; Feierstein et al., 2006; Roesch et al., 2006; Strait et al., 2015; Tsujimoto et al., 2009), in our study the location of the cues did not influence the results. To show this, we fit a model that had the same terms as the main GLM (Equation 1), plus an additional term for the side of screen where the cue appeared. Only 4.4% of neurons showed a significant effect of side

Table 1. Significant Effects in GLM

Percent of Neurons with Effects at $p < 0.05^a$	Regressors			Max Expected by Chance (All Variables)
	Value	Distance	Value by Distance	
Uncorrected	59.4	53.7	27.9	9.8
Corrected	36.0	30.7	8.8	2.5

See also Table S1 and Figure S4.

^aThe percentage of neurons (out of 283) significantly modulated by variables in a GLM (Equation 1). The maximum percentages expected by chance (right column) were determined by finding the maximum percentage of significant effects of any single variable within 1,000 randomly permuted data sets (see Experimental Procedures). Corrected p values were obtained with Holm's variant of the Bonferroni correction.

($p < 0.05$, corrected), and the beta co-efficients for value, distance, value-by-distance were virtually identical to those from the main model ($r > 0.998$ for all three terms). The negligible effect of target location in our study may be rooted in experimental design: the studies cited above required operant responses to targets at particular locations in visual space, whereas no such response was required in our task.

Taken together, these additional analyses show that the gaze distance effects in Table 1 and Figures 3, 4, and 5 are not attributable to other oculomotor or eye position variables that could be represented in the OFC.

Value and Distance Encoding Are Mixed

We next asked how often OFC cells showed more than one significant effect in the GLM, e.g., $p < 0.05$ for both cue value and fixation distance. If significant effects are randomly and independently distributed across cells, then the chance that two or three effects occur in any one cell is product of the individual proportions of significant effects of each variable, shown in Table 1. Such a pattern would be consistent with the “mixed” selectivity of variables observed in other frontal lobe structures (Machens et al., 2010; Mante et al., 2013; Miller and Cohen, 2001; Rigotti et al., 2013). In contrast, if the co-occurrence rate were less than expected, it would suggest that cells tend to encode only one variable at a time, consistent with “discrete” encoding.

The Venn diagram in Figure 5A shows the proportion of cells with either one or multiple significant effects, at both uncorrected and corrected thresholds of $p < 0.05$. The sum of the proportions in each of the three circles gives the same proportions as in Table 1. When using a corrected threshold (bottom row of Table 1; numbers in parentheses in Figure 5A), the proportion of cells with effects of β_{VALUE} was 36.0%, and for β_{DIST} was 30.7%, and so the expected proportion of neurons with both effects was their product, 11.0%. As shown in Figure 5A, the actual proportion was higher: 12.7% had effects of both β_{VALUE} and β_{DIST} , and an additional 4.9% had all three effects, for a total of 17.7%, which was greater than expected by chance ($p = 1 \times 10^{-6}$ by chi-square test). By this same procedure, co-occurrence of β_{VALUE} and $\beta_{\text{VAL} \times \text{DIST}}$ was also greater than chance (expected 3.2%, actual 6.4%, $p = 0.0002$), as was the co-occurrence of β_{DIST} and $\beta_{\text{VAL} \times \text{DIST}}$ (expected 2.7%, actual 6.7%, $p = 1 \times 10^{-6}$). Finally, 4.9% showed all three significant effects, greater than the expected rate of 1.0% ($p = 1 \times 10^{-10}$). Together, these four comparisons show that cells tend to encode multiple

variables slightly more often than expected by chance, arguing strongly for mixed encoding of these variables.

To further show the mixed encoding of value and gaze variables, we now address the individual (Figure 5B) and joint distributions of β_{VALUE} , β_{DIST} , and $\beta_{\text{VAL} \times \text{DIST}}$ (Figure 5C). Four features of these distributions bear notice: First, the median of β_{DIST} was negative (-0.0057 , $p = 2.4 \times 10^{-7}$ by Wilcoxon rank-sum test), meaning that overall, cells fired more for fixations near the cue and less for fixations away. Thus, the responses in Figures 4A and 4C illustrate the most common distance effects in the population. Second, the median of β_{VALUE} was not different from zero (0.007 , $p = 0.81$); this means that cells were equally likely to increase or decrease firing as a function of value. Third, all three variables had continuous, unimodal distributions (Figure 5B). Fourth, the joint distributions were essentially featureless clouds, lacking distinct clusters, with no correlation among the variables (Figure 5C, see caption for statistics). This argues against discrete selectivity, which would result in scatter plots with points clustered around the vertical and horizontal axis lines (i.e., each point would be non-zero for one variable, but nearly zero for the others). Unlike in the amygdala (Peck et al., 2013), there was no correlation between β_{VALUE} and a cell's preference for contralaterally located cues (Figure S6).

As a final illustration of the mixture of value and gaze signals, we now use the model results to divide cells into groups with quantitatively different response patterns and show the average firing rates of these groups in Figures 5D–5K. The colors in Figure 5A show how cells were grouped, and these colors map onto the headings in Figures 5D–5K. Grouping was done with uncorrected thresholds ($p < 0.05$) to capture the influence of cells with weakly significant effects that may contribute to firing modulation at the population level (see also Figure S5). The first group consists of all cells with significant effects of $\beta_{\text{VAL} \times \text{DIST}}$ in the GLM, corresponding to the light red circle in Figure 5A and the two plots in Figures 5D and 5E. Two firing rate plots are needed for this group to show the two possible interaction patterns: some cells have more value coding for fixations near the cue (Figure 5D; $n = 42$ cells), and others for fixations away (Figure 5E; $n = 37$). A second group consisted of neurons with both value and distance effects, but with no interaction effect (yellow in Figures 5A, 5F, and 5G). This group was also split into two plots, according to whether on-cue fixations produced maximum or minimum firing. The gray and blue groups illustrate neurons with only distance (Figures 5H and 5I) or only value effects (Figure 5J), respectively.

For the cells that mix value and gaze distance signals, a critical question is how the overall, population-level value signal changes as gaze moves from place to place, given that human fMRI results indicate that the frontal lobe preferentially encodes (with increased BOLD signal) the value of fixated items (Lim et al., 2011). Consistent with this, some cells have stronger value coding when fixating near the cues (Figure 5D). However, others show the opposite response pattern, stronger value signals when fixating away (Figure 5E). At first glance, these opposing population level responses may appear incompatible with the fMRI results; however, closer inspection reveals otherwise: First, at the two extremes of fixation distance, value encoding was stronger within the “near”

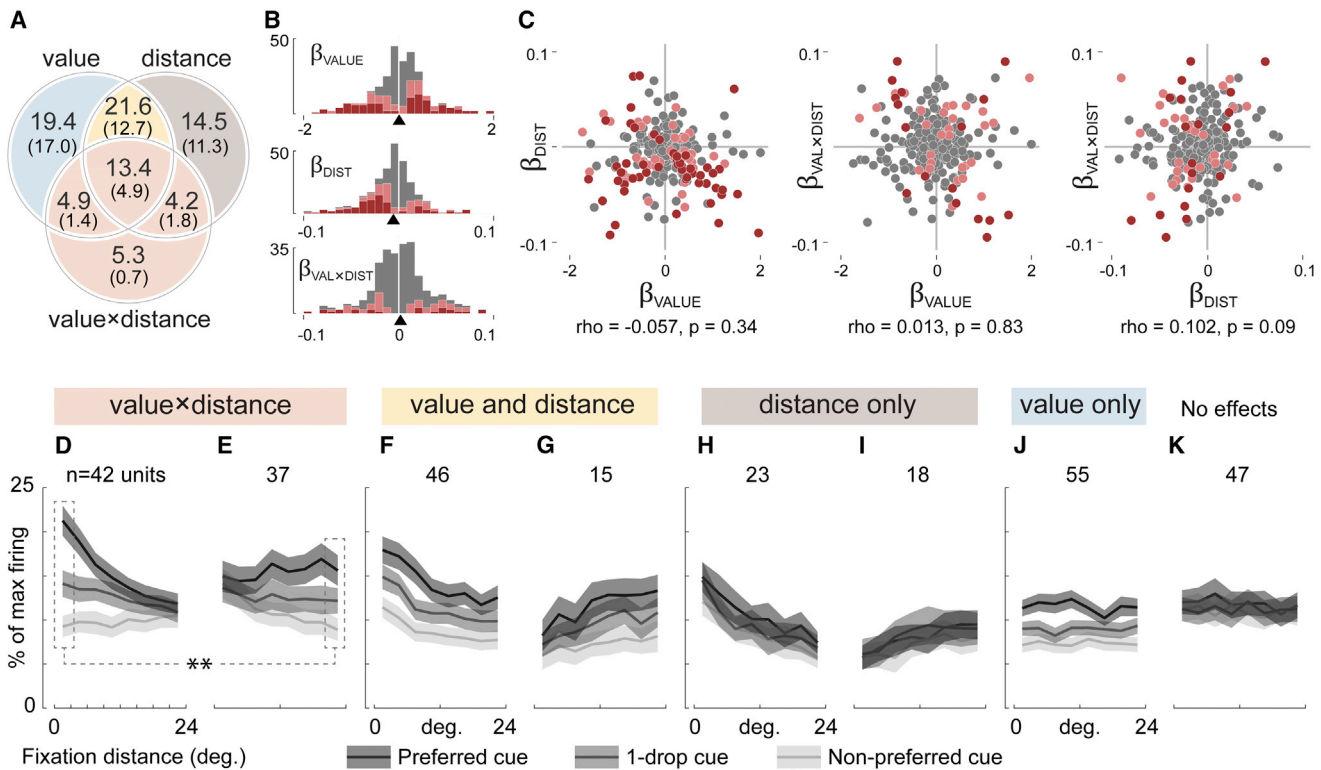


Figure 5. Value and Fixation Distance Encoding in the Population

(A) Percent of neurons modulated by value, fixation distance, and the value-by-distance interaction determined by a GLM (Equation 1). The large numerals give significant effects at an uncorrected threshold ($p < 0.05$), and the numerals in parentheses give effects with p value correction applied (Holm's method, $p < 0.05$). The colors indicate groups of neurons averaged together for plotting in (D)–(J): light red for (D) and (E), yellow for (F) and (G), gray for (H) and (I), and blue for (J). (B) Distributions of beta co-efficients from the GLM. The pink indicates significance at $p < 0.05$ uncorrected, and the red indicates $p < 0.05$ corrected. The arrowheads on the x axis indicate medians. (C) Joint beta co-efficient distributions. The pink and red indicate a significant effect on both axes at $p < 0.05$ uncorrected and $p < 0.05$ corrected, respectively. The statistics indicate Spearman's correlation co-efficient. (D and E) Average firing of cells placed in groups according to the GLM results shown in (A) (large numerals, uncorrected threshold of $p < 0.05$). The firing was normalized to the maximum within each cell. The x axis gives the distance of fixation from the cue, and the gray scale gives cue identity as follows: cells with greater firing for high value cues ($\beta_{\text{VALUE}} > 0$) had the large reward cue as the "preferred" cue and the no-reward cue as "non-preferred". The cells with greater firing for low value ($\beta_{\text{VALUE}} < 0$) had the opposite assignment. The heading colors show membership in the Venn diagram in (A): (D) and (E) are all the cells with significant effects of the value-by-distance interaction ($p < 0.05$, uncorrected). In (D), cells have stronger value coding for near cue fixations, whereas in (E), cells have stronger value coding for fixations away. The dotted lines and stars indicate stronger value coding in (D) compared to (E) at the selected distance bins (see main text). (F and G) Shows cells that additively combine value and distance effects with no interaction, with cells in (F) firing more overall for near-cue fixations and cells in (G) firing more for fixations away. (H and I) Have only distance effects, firing more (H) or less (I) as gaze approaches the cue. (J) Shows a value-only effect. (K) The average of all neurons not in (D)–(J) (i.e., all effects $p > 0.05$, uncorrected). See also [Figures S2, S3, S5, and S6](#).

subgroup (Figure 5D) compared to the "away" subgroup (Figure 5E). Specifically, the firing difference between large and no reward cues (coded as preferred and non-preferred in Figures 5D and 5E) was 12.0% (SEM 1.3%) for on-cue fixations in Figure 5D, but was only 6.7% (SEM 0.8%) for away-from-cue fixations in Figure 5E ($p < 0.002$, Wilcoxon rank-sum test, see dotted lines in Figures 5D and 5E). Thus on-cue fixations produce a stronger, population-level value code than fixations away.

A related comparison can be made for cells that mix value and gaze in an additive fashion (yellow in Figures 5A, 5F, and 5G).

Cells that fire more overall for on-cue fixations (Figure 5F) outnumber those that do the opposite (Figure 5G): $n = 46$ versus 15, $p = 3 \times 10^{-5}$ by Fisher's test for proportions. This reflects the overall negative trend in the regression co-efficients for fixation distance (Figure 5B). Thus, the overall spike output of cells that mix value and distance signals is greater when looking near cues than when looking away.

Together, these data demonstrate a mixture of value and fixation location encoding in our population. That so many cells are modulated by both variables, including significant interactions, means that the overall value signal expressed by OFC is highly

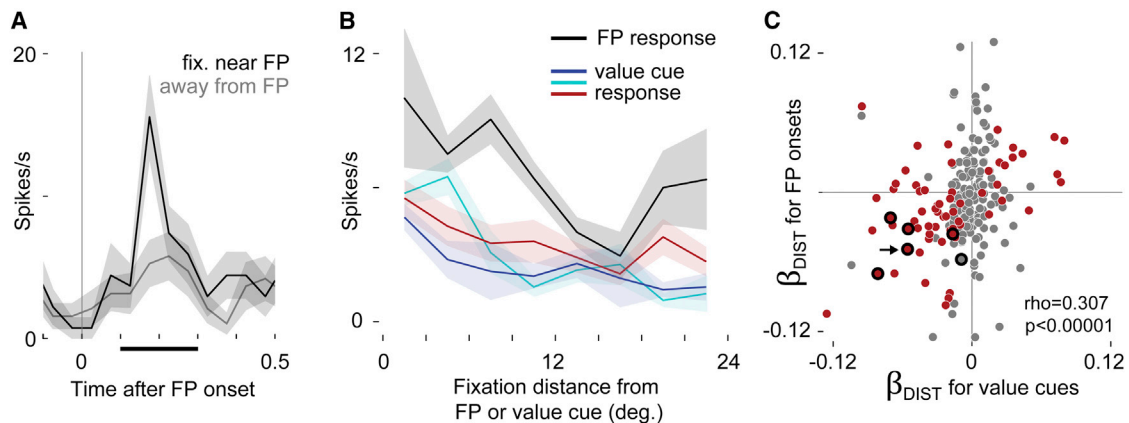


Figure 6. Gaze Distance Modulates Firing Evoked by Fixation Point Onset

(A) Firing in an identified single unit time-locked to the onset of the FP. The black shows when gaze was near the FP at onset (<5 degrees), and the gray shows when gaze was away (>15 degrees). The firing was measured in a 200 ms window (black bar on x axis) to generate the black line in (B).

(B) Same cell as (A). The firing as a function of distance of gaze from the FP (black) or from the value cues (colors, as in Figures 1, 2, 3, and 4). In (A) and (B), the lines show means, and the shaded areas show SEM.

(C) Comparison of gaze distance effects in the value cue data and FP-evoked responses. The x axis gives β_{DIST} from the value cue data (same as β_{DIST} in Figure 5), and the y axis gives $\beta_{\text{DIST-FP}}$, calculated in a separate GLM using firing evoked by FP onsets. Each dot indicates a neuron, and the arrow shows the cell in (A) and (B). A red point indicates a significant effect in the value cue data, and a thick black ring, the FP data ($p < 0.05$ corrected for both). The overlapping red points within black rings indicate cells with significant effects of both β_{DIST} and $\beta_{\text{DIST-FP}}$. Because fewer observations were available in the FP data (see text), only 228 neurons had sufficient data to fit the model and calculate $\beta_{\text{DIST-FP}}$ and very few neurons showed significant effects. The correlation statistic reflects 228 cells, but (C) shows only 218 due to axis limits; the correlation for the visible data points alone ($n = 218$) is $\rho = 0.310$, $p < 1 \times 10^{-5}$.

dynamic, varying with fixation location as it changes moment-to-moment during free viewing.

Gaze Distance Encoding Persists in a Separate Behavioral Context

Recent theories suggest a role for the OFC in representing task context (Rudebeck and Murray, 2014; Wilson et al., 2014), and that OFC neurons, like those other frontal lobe regions, may be highly sensitive to situational demands. We therefore asked whether gaze modulation differed for two distinct behavioral contexts within our task. Similar gaze encoding across contexts would suggest that gaze signals are inherent in OFC, whereas context-dependent gaze encoding would suggest a mechanism that is recruited selectively according to task structure.

The first context is the free viewing of the value cues (analyzed above). To measure gaze effects in a second context, we took advantage of a salient visual stimulus that was distinct from the value cues: the onset of the FP at trial start. Unlike the value cues, FP onset is uncertain, occurring 2–4 s after the last trial, randomly on the left or right of the screen. The FP itself has no explicit value, other than signaling the potential for an uncertain reward, contingent on an action which the animal has not yet planned (saccade to FP and hold fixation). Thus, the moment of FP onset is a distinct behavioral context from the value cue viewing period.

We focused on instances where the eyes were stationary at FP onset and measured post-FP firing over the same cell-specific window used for fixation-evoked firing in the value cue data (as in Figure 3D). Because the subjects were free viewing before FP onset, their gaze was sometimes near and sometimes distant from the FP when it appeared (3.6% of onsets with gaze <3

degrees from FP). Figure 6A shows data from an identified single unit with strong FP-evoked excitation when the eyes were near the FP, but only weak excitation when the eyes were away. Figure 6B shows data from this same single cell, with both FP-evoked firing and fixation-evoked firing during value cue viewing plotted as a function of gaze distance from the FP/cue. In both the FP and value cue contexts, the cell fired more for gaze distances near the stimulus and less when gaze was far away.

At the population level, we compared the distance encoding for FP onsets to the distance encoding for value cues. First, we fit a GLM that explained FP-evoked firing as a function of fixation distance from the FP, yielding beta co-efficients for every cell, $\beta_{\text{DIST-FP}}$. On average, $\beta_{\text{DIST-FP}}$ was slightly negative (median -0.006 , $p < 0.016$ by rank-sum test, and $n = 228$ neurons with sufficient data to fit the model), indicating that most cells fired more when gaze was near the FP. Figure 6C compares $\beta_{\text{DIST-FP}}$ to the β_{DIST} estimates derived from the value cue data (same data as Figure 5); the arrow shows data from the single cell in Figures 6A and 6B, which has significant gaze distance effects in both GLMs ($p < 0.05$, corrected). $\beta_{\text{DIST-FP}}$ was positively correlated ($\rho = 0.307$, $p < 1 \times 10^{-5}$) with β_{DIST} , meaning that across the population, cells firing more for FP onsets near the center of fixation tended to also fire more following fixations near to the value cues.

The natural scale of positive correlations is 0–1.0, and on this scale, the correlation in Figure 6C ($\rho = 0.307$) appears modest. However, two factors in our data could constrain the upper limit of ρ below the natural limit of 1.0. The first factor is the inherent noise entailed in estimating β_{DIST} and $\beta_{\text{DIST-FP}}$, which will propagate into the calculation of ρ . As we show below, this noise was substantial for $\beta_{\text{DIST-FP}}$, which was estimated with many fewer

observations than β_{DIST} (142 FP onsets versus 1,425 fixations per cell, on average). The second factor was the very different sampling of visual space in the FP and cue viewing contexts: the monkeys often looked directly at the value cues (31.7% grand average frequency of on-cue fixations in Figure 2B), whereas their gaze was near the FP for only a fraction of FP onsets (3.6% of onsets with gaze <3 degrees from FP). Even if a cell had identical underlying gaze effects in both contexts, the different spatial sampling could produce different estimates of this effect, and therefore different values of β_{DIST} and $\beta_{\text{DIST-FP}}$.

We therefore asked how estimation noise and sampling patterns influenced the calculation of the correlation statistic, by estimating the theoretical upper limit of ρ . To determine how estimation noise influences ρ , we first calculated the reliability (a measure of self-correlation) of β_{DIST} and $\beta_{\text{DIST-FP}}$. For β_{DIST} , the reliability was 0.87, whereas for $\beta_{\text{DIST-FP}}$ it was only 0.45, consistent with the fact that fewer observations were used to estimate $\beta_{\text{DIST-FP}}$. Next, assuming that the underlying gaze effect in the two contexts is identical, the theoretical upper limit on the correlation we could observe between β_{DIST} and $\beta_{\text{DIST-FP}}$ is the square root of the product of their reliabilities: 0.63 (see Supplemental Information).

To determine how differences in spatial sampling influence ρ , we used a resampling method. Within each neuron, we resampled the value cue data in a way that matched the sampling conditions of the FP data (fewer observations, and, critically, fewer fixations onto targets) and then recalculated β_{DIST} . We then compared this resampled β_{DIST} to the original β_{DIST} across the population with Spearman's correlation co-efficient, yielding an estimate of how similar β_{DIST} in the value cue data was to itself, but under the sampling constraints of the FP data. In 500 repetitions of this process, the median correlation was 0.47 and the maximum was 0.59.

Together, these two methods suggest that the correlation tested in Figure 6C would be approximately 0.6 (not 1.0) if the underlying effects in the two data sets were identical. On this scale, the observed correlation ($\rho = 0.307$) is more convincing and gives firm grounds to conclude that neurons with fixation distance effects during value cue viewing also have similar effects at the moment of FP onset.

DISCUSSION

This study is motivated by two fundamental observations. First, to survive, animals must judge the economic value of the stimuli in their world. Second, primates frequently process visual stimuli one at a time, by shifting the location of gaze among objects in visual space (Bichot et al., 2005; DiCarlo and Maunsell, 2000; Mazer and Gallant, 2003; Motter and Belky, 1998; Sheinberg and Logothetis, 2001). While many studies have addressed the neural correlates of value in primates, very few have asked how these value signals are influenced by changes of gaze, which all primates perform constantly throughout waking life. Here, by leveraging natural gaze behavior in awake monkeys, we identified OFC neurons that simultaneously signal object value and gaze location, providing new insight into cellular-level dynamics in the valuation circuitry of the primate frontal lobe.

In recordings from two free-viewing monkeys, we identified OFC neurons (both single and multi-unit signals) encoding cue value—as expected from previous studies—but also found many encoding the distance of fixation relative to the cue. Neurons encoding gaze distance were almost as abundant as those encoding value and many cells encoded both simultaneously. Indeed, most value coding neurons also had some form of gaze modulation (Figure 5A).

This particular form of gaze encoding—distance of gaze from the cue—could not be explained by other oculomotor variables, nor by other forms of eye position modulation such as the head-centered angle of gaze or planar gain fields. Moreover, the location of the cue (left or right side of the screen) had virtually no impact on our results, suggesting that the gaze effects we report are unrelated to spatial representation of targets in OFC reported previously (Abe and Lee, 2011; Feierstein et al., 2006; Roesch et al., 2006; Strait et al., 2015; Tsujimoto et al., 2009). Finally, cells with gaze-dependent firing in reference to Pavlovian cues also showed similar gaze modulation in their responses to onset of the FP, a stimulus whose form and behavioral significance were distinct from the value cues. This cross-context gaze encoding suggests that gaze modulation is not exclusive to the particular structure of the task used here.

Perhaps the most striking feature of our data is the cells encoding the interaction between value and gaze, i.e., a value code that is amplified or attenuated according to where the subjects look (Figures 3, 4A, 4B, 5D, and 5E). The existence of these cells suggests that gaze can act as a filter that shapes and constrains the overall value signal expressed by the OFC, according to moment-by-moment changes in gaze location.

The Role of Gaze Modulated Value Signals in Decision-Making

Although we employed a Pavlovian task that required no decision from the subjects, our findings bear on recent studies of visual fixations, choice preference, and value representations in humans.

First, when humans are asked to shift gaze between two objects while choosing between them, BOLD signals in the ventral striatum and medial prefrontal cortex are positively correlated with the value of the object fixated at a given moment and negatively correlated with the value of the other object (Lim et al., 2011). In that study, fixation was controlled by the experimenter and fixations were prolonged to the scale of the hemodynamic response (1–4 s). Our results are consistent with the measurements of Lim et al. (2011) and give a potential cellular-level basis for these effects by demonstrating fixation-modulated value signaling within single neurons in the context (and on the time-scale) of natural primate gaze behavior.

Second, studies in human subjects show that longer fixation on a given item increases its likelihood of being chosen (Armell et al., 2008; Kraglich et al., 2012, 2010; Kraglich and Rangel, 2011; Towal et al., 2013; Vaidya and Fellows, 2015). The computational model that best explains this behavior is one in which fixation actively biases choices by amplifying the value signals of fixated items (Kraglich et al., 2010). A critical component of this mechanism is an input signal that reflects the value of the current object of gaze. Our study identifies for the first time a

neuronal-level signal that is consistent with such an input: a large subset of OFC neurons whose activity is modulated by ongoing changes in fixation location (Table 1; Figure 5). In particular, we identified cells that amplify their coding of value when gaze is focused on a cue (Figure 5D). Experiments by our group and others (Malalasekera et al., 2014) are currently underway to assess the role of fixation-based value signaling in single neurons during economic choice.

One outstanding question is whether gaze-driven value signals occur in many different tasks or in only a few behavioral contexts. As outlined by Wilson et al. (2014) and by Rudebeck and Murray (2014), OFC activity appears to be highly dependent on task context, e.g., current goals, task rules, event history, and internal states of hunger or thirst. Moreover, task context could itself be explicitly encoded by OFC (Saez et al., 2015). In this study, we found gaze encoding in two distinct contexts in the same subset of cells, providing some evidence in favor of persistent encoding of gaze.

Visual Attention: a Potential Mechanism Underlying Gaze-Based Value Coding

This study focuses on the role of gaze in value coding for two reasons. First, shifts of gaze plays a critical role in natural settings, in which animals must quickly process the many stimuli they encounter. Second, evidence from the human literature discussed above suggests that natural gaze behavior might play an important role in value-based choice.

However, our results give rise to a natural question: Are the gaze-driven effects in OFC specific to overt shifts in fixation, or are these effects in fact the result of a more general process: visual attention? Visual attention is the selective processing of particular objects or locations in visual space and is typically studied in two forms: overt attention involves actively moving the gaze onto objects of interest, whereas covert attention involves attending to objects away from the center of gaze, usually while holding the eyes still (Findlay and Gilchrist, 2003). Attentional shifts influence behavior as well as neural signals in many brain regions (Desimone and Duncan, 1995; Kastner and Ungerleider, 2000), and recent theories suggest that executing overt and covert shifts of attention involve similar neural circuits (Moore et al., 2003).

Our findings clearly show that OFC value signals are modulated by overt attentional shifts (equivalent to changes of gaze location), an insight made possible by recording during unrestricted free viewing. However, this approach does not show whether OFC value signals are modulated by covert attentional shifts. Thus, it is an open question whether the value modulation we observed reflects a general attentional process encompassing both overt and covert mechanisms.

While there is no direct evidence for covert attentional effects in OFC, incidental findings in prior studies suggest it is possible. For example, new stimuli that appear in the visual field can draw, or “capture”, covert visual attention. Rudebeck et al. (2013) show that in monkeys maintaining central fixation, the addition of a new peripheral stimulus shifts OFC responses to reflect the value of this stimulus and decreases encoding of stimuli that were already present, a potential signature of attentional capture. A second example can be found in Padoa-Schioppa

and Assad (2006); they show OFC cells that are sensitive to the value of one item, but insensitive to the value of other items shown simultaneously, an effect that might be expected if only that one item, but not the others, were covertly attended. Third, recent work from Peck et al. (2013) demonstrates covert attentional modulation in the primate amygdala, a region with connectivity and neural function similar to the OFC (Carmichael and Price, 1995; Kravitz et al., 2013; Rolls, 2015). Finally, the OFC receives a large input from the anterior inferotemporal cortex (aIT) (Kravitz et al., 2013; Saleem et al., 2008). Neurons in aIT selectively encode of visual color and form, and their responses often reflect attended objects to the exclusion of other objects, whether attention is overt or covert (DiCarlo and Maunsell, 2000; Moore et al., 2003; Moran and Desimone, 1985; Richmond et al., 1983; Sheinberg and Logothetis, 2001). Thus, the OFC could in theory inherit covert attentional modulation from aIT.

Determining whether covert attention modulates value signals in the OFC—and elsewhere—is a critical open question for future studies. If covert attention modulates value signals, it would complicate the interpretation of OFC neural data obtained from subjects maintaining fixation at a single location; under those conditions the neural value signals could in theory be subject to a covert attentional filter, shifting uncontrolled and unmonitored between different objects in the periphery. To resolve this question will require experiments that explicitly control, or at least monitor, covert attentional deployment onto visual objects of differing value.

Conclusions

Value signaling is a fundamental function of the brain. Here, we have shown value signals in the OFC that are modulated by moment-to-moment changes in gaze during natural free viewing. The abundance of this gaze encoding, its cell-level mixture with value signals, and its persistence across task phases suggests it is a major modulator of OFC function. Critical open questions are whether these value signals show these same dynamics in other instances of motivated behavior (such as decision-making), and, if so, how they influence these behaviors. Given that primates are free viewing throughout waking life, the answers to these questions have the potential to transform our understanding of the OFC and to bring us closer to understanding the neural basis of motivated behavior in the real world.

EXPERIMENTAL PROCEDURES

Subjects and Apparatus

All procedures were performed in accordance with the NIH Guide for the Care and Use of Laboratory Animals and were approved by the Animal Care and Use Committee of Stanford University. The subjects were two adult male rhesus monkeys with recording chambers allowing access to the OFC. They performed the task while head-restrained and seated before a monitor showing the task stimuli. Eye position was monitored at 400 Hz. Juice reward was delivered via a tube placed ~3 mm outside the mouth, and the monkeys retrieved reward by touching their tongue to the end of the tube during delivery. Contact between the tongue and juice tube (the “licking response”, see below) was monitored at 400 Hz, as described previously (Fiorillo et al., 2008). The licking-versus-time plots in Figures 1 and S3 show the percentage of trials in which contact was detected at a given time point. Task flow and stimulus presentation were controlled using the REX software suite (Laboratory of Sensorimotor Research, NEI).

Behavioral Task

Figure 1 illustrates the Pavlovian conditioning task. Licking responses were assessed by measuring the total duration of tongue contact with the juice tube in the 4 s period prior to juice delivery (Fiorillo et al., 2008; Morrison and Salzman, 2009). Because new color cues were used in every session, the licking response was initially indiscriminate (data not shown) and then with subsequent trials became commensurate with the reward: large > small > no reward. The initial conditioning phase was terminated (data collection began) when licking durations over the prior 60–100 trials were different for all three trial types (rank-sum test, $p < 0.01$, uncorrected). A session was discarded if the licking responses did not maintain selectivity after data collection began.

During some neural recordings, we abruptly switched, or “reversed”, the cue-reward associations of the no-reward and large reward cues (Morrison and Salzman, 2009; Thorpe et al., 1983). By comparing neural activity before and after reversal, we assessed the encoding of cue value independent of cue color. The first 40 trials after reversal were not analyzed to provide sufficient time for the new cue-reward associations to be learned and for OFC responses to adapt (Morrison et al., 2011). Reversal data were discarded if licking responses did not update to reflect the new cue-reward associations. All analyses except those for Figure S3 (which focuses on the effects of reversal) use only either pre- or post-reversal data for a given cell, but never both, based on which block contained the most trials.

Recording and Data Collection

We recorded from 283 neural unit signals in OFC, 144 from monkey 1 and 139 from monkey 2, using single tungsten electrodes (FHC). OFC was identified on the basis of gray/white matter transitions and by consulting MRIs acquired after chamber implantation (Figures 1C and 1D). We recorded both putative single neurons (single unit), and signals consisting of the mixed activity of multiple neurons (multi-unit). To avoid confusion, we use the terms single unit or multi-unit when referring to individual responses (as appropriate) and the terms cells and neurons when referring to group data that encompasses both single and multi-unit responses. After offline spike sorting, 176 neurons were designated as single units and 107 as multi-unit. Our findings do not differ between single- and multi-unit signals (Table S1) and so they are presented together. The data set contains cells that were lightly screened for task-related activity (broadly defined) as well as unscreened cells (see Supplemental Information).

Data Analysis

The objective was to determine how neural activity was modulated by both cue value and by the location of fixation. Because the subjects were free viewing, fixation timing and location were highly variable across trials; thus, the fundamental units of analysis were individual fixations not individual trials.

First, we detected individual fixations (periods of stationary gaze) by calculating a velocity threshold based on the velocity variance within a given trial (see Kimmel et al. 2012 and Supplemental Information). For a fixation to be eligible for analysis, its onset had to occur between $t = 0.5$ and 3.75 s after cue onset, to exclude from analysis firing related to cue onset or reward delivery (at $t = 0$ and 4 s, respectively). Fixations also had to be located within the calibrated range of the eye tracker and had to be at least 100 ms in duration (see Supplemental Information).

Next, for each fixation we computed the fixation-evoked firing, which was the spike count within a 200 ms window following fixation onset (illustrated in Figure S1). Importantly, the start and end of the post-fixation time window was defined uniquely for each neuron to account for cells that have different response latencies to changes in visual input (see Supplemental Information). The primary analyses in this paper uses this cell-specific firing window; however, we performed additional analysis using a fixed post-fixation window for all cells (Table S1) or using a range of time windows from 0 to 600 ms (Figure S2). In Figures 5D–5K, the spike count data were scaled within each neuron to between 0% and 100%, measured across all fixations.

We then used GLM's to quantify the effect of cue value and gaze location on firing. Our main results are based on the estimation of the following

GLM, which assumed that fixation-evoked spike counts follow a negative binomial distribution:

$$\log(Y) = \beta_0 + \beta_{VAL} * Value + \beta_{DIST} * Distance + \beta_{VAL \times DIST} * Val \times Distance, \quad (\text{Equation 1})$$

where each observation is a fixation (as defined above), Y is the fixation-evoked firing for that fixation, *Value* refers to the volume of juice associated with the cue in each trial (scaled so that 0 corresponds to the no-reward cue and 1 corresponds to the large cue), *Distance* refers to the distance of gaze from the cue center for each fixation (coded in degrees; range 0 to 24), and $Val \times Dist$ is the interaction of the *Value* and *Distance* variables (computed after centering them).

We estimated two additional GLMs to assess alternative schemes for encoding fixation location. The first one used the absolute angle of gaze in head-centered coordinates:

$$\log(Y) = \beta_0 + \beta_{VAL} * Value + \beta_{ANG} * Angle + \beta_{VAL \times ANG} * Val \times Angle. \quad (\text{Equation 2})$$

The second used horizontal and vertical distances to the cue:

$$\log(Y) = \beta_0 + \beta_{VAL} * Value + \beta_{HOR} * Horizontal + \beta_{VER} * Vertical. \quad (\text{Equation 3})$$

The relative fits of the models were evaluated by comparing the goodness of fit for each alternative model to the one for the main model, using AIC.

The GLMs were estimated for each neuron separately. We then carried out population-level comparisons using a variety of tests. To test for differences in means, we used Wilcoxon rank-sum tests, which are robust to outliers and non-normally distributed data. Correlations were assessed using Spearman's ρ , an outlier-resistant measure of association. When p value corrections were applied, Holm's modification of the Bonferroni correction was used with a threshold of $p < 0.05$.

For some GLMs, we address the issue of multiple comparisons by fitting the GLMs to data for which the spike counts were permuted. Permutation was performed by randomly shuffling the fixation-evoked spike counts among the fixations within a given cell, eliminating any systematic relationship between spiking and the regressor variables, leaving only chance correlations (see Supplemental Information).

Responses to the FP

We also examined firing evoked by the onset of the FP at the beginning of each trial to test whether these responses were also modulated by the distance of fixation from the stimulus. To do so, we measured firing time-locked to each FP onset by counting the spikes within the same cell-specific 200 ms windows described above. FP onsets were subjected to eligibility criteria similar to those imposed on fixations (see Supplemental Information). We then estimated a GLM that contained only a single regressor: the distance of gaze from the FP at the time of onset. The resulting estimates of $\beta_{DIST-FP}$ were then compared to the estimates of β_{DIST} described above (Main GLM, Equation 1), on a cell-by-cell basis.

To interpret this result, we asked whether the observed correlation between β_{DIST} and $\beta_{DIST-FP}$ is subject to an upper bound due to noisiness inherent in their estimation or to the fact that fixation locations in the value cue data differed from fixation locations at FP onset. This upper bound was estimated with two complementary methods: one that uses the reliability of the estimates in each data set, and the other that uses a random resampling procedure. The two methods produced similar results (see Supplemental Information).

SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures, six figures, and one table and can be found with this article online at <http://dx.doi.org/10.1016/j.neuron.2016.04.045>.

AUTHOR CONTRIBUTIONS

Conceptualization, V.B.M., A.R., and W.T.N.; Methodology, V.B.M. and W.T.N.; Investigation, V.B.M.; Formal Analysis, V.B.M.; Writing – Original Draft,

V.B.M.; Writing – Review and Editing, V.B.M., A.R., and W.T.N.; Visualization, V.B.M.; Supervision, W.T.N. and A.R.; and Funding Acquisition, V.B.M. and W.T.N.

ACKNOWLEDGMENTS

We thank J. Brown, S. Fong, J. Powell, J. Sanders, and E. Carson for technical assistance. C. Chandrasekaran, R.B. Ebitz, and S. Morrison commented on the manuscript. This work was supported by the Howard Hughes Medical Institute (W.T.N.), United States Air Force grant FA9550-07-1-0537 (W.T.N.), and by NIH grants K01 DA036659-01 and T32 EY20485-03 (V.B.M.).

Received: June 17, 2015

Revised: January 30, 2016

Accepted: April 22, 2016

Published: June 2, 2016

REFERENCES

- Abe, H., and Lee, D. (2011). Distributed coding of actual and hypothetical outcomes in the orbital and dorsolateral prefrontal cortex. *Neuron* 70, 731–741.
- Armell, K.C., Beaumel, A., and Rangel, A. (2008). Biasing simple choices by manipulating relative visual attention. *Judgm. Decis. Mak.* 3, 396.
- Basten, U., Biele, G., Heekeren, H.R., and Fiebach, C.J. (2010). How the brain integrates costs and benefits during decision making. *Proc. Natl. Acad. Sci. USA* 107, 21767–21772.
- Bichot, N.P., Rossi, A.F., and Desimone, R. (2005). Parallel and serial neural mechanisms for visual search in macaque area V4. *Science* 308, 529–534.
- Bouret, S., and Richmond, B.J. (2010). Ventromedial and orbital prefrontal neurons differentially encode internally and externally driven motivational values in monkeys. *J. Neurosci.* 30, 8591–8601.
- Bruce, C.J., and Goldberg, M.E. (1985). Primate frontal eye fields. I. Single neurons discharging before saccades. *J. Neurophysiol.* 53, 603–635.
- Carmichael, S.T., and Price, J.L. (1995). Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. *J. Comp. Neurol.* 363, 615–641.
- Desimone, R., and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* 18, 193–222.
- DiCarlo, J.J., and Maunsell, J.H.R. (2000). Form representation in monkey inferior temporal cortex is virtually unaltered by free viewing. *Nat. Neurosci.* 3, 814–821.
- Feierstein, C.E., Quirk, M.C., Uchida, N., Sosulski, D.L., and Mainen, Z.F. (2006). Representation of spatial goals in rat orbitofrontal cortex. *Neuron* 51, 495–507.
- Findlay, J.M., and Gilchrist, I.D. (2003). *Active Vision* (Oxford University Press).
- Fiorillo, C.D., Newsome, W.T., and Schultz, W. (2008). The temporal precision of reward prediction in dopamine neurons. *Nat. Neurosci.* 11, 966–973.
- Hare, T.A., Camerer, C.F., and Rangel, A. (2009). Self-control in decision-making involves modulation of the vmPFC valuation system. *Science* 324, 646–648.
- Kable, J.W., and Glimcher, P.W. (2007). The neural correlates of subjective value during intertemporal choice. *Nat. Neurosci.* 10, 1625–1633.
- Kastner, S., and Ungerleider, L.G. (2000). Mechanisms of visual attention in the human cortex. *Annu. Rev. Neurosci.* 23, 315–341.
- Kennerley, S.W., Behrens, T.E.J., and Wallis, J.D. (2011). Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nat. Neurosci.* 14, 1581–1589.
- Kimmel, D.L., Mamm, D., and Newsome, W.T. (2012). Tracking the eye non-invasively: simultaneous comparison of the scleral search coil and optical tracking techniques in the macaque monkey. *Front. Behav. Neurosci.* 6, 49.
- Krajibich, I., and Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proc. Natl. Acad. Sci. USA* 108, 13852–13857.
- Krajibich, I., Armell, C., and Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nat. Neurosci.* 13, 1292–1298.
- Krajibich, I., Lu, D., Camerer, C., and Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Front. Psychol.* 3, 193.
- Kravitz, D.J., Saleem, K.S., Baker, C.I., Ungerleider, L.G., and Mishkin, M. (2013). The ventral visual pathway: an expanded neural framework for the processing of object quality. *Trends Cogn. Sci.* 17, 26–49.
- Lim, S.-L., O'Doherty, J.P., and Rangel, A. (2011). The decision value computations in the vmPFC and striatum use a relative value code that is guided by visual attention. *J. Neurosci.* 31, 13214–13223.
- Machens, C.K., Romo, R., and Brody, C.D. (2010). Functional, but not anatomical, separation of “what” and “when” in prefrontal cortex. *J. Neurosci.* 30, 350–360.
- Malalasekera, N., Hunt, L.T., Miranda, B., Behrens, T.E.J., and Kennerley, S.W. (2014). Neuronal correlates of value depend on information gathering and comparison strategies. In: *Neuroscience Meeting Planner*. Presented at the Society for Neuroscience Annual Meeting, Society for Neuroscience, Washington, DC, p. 206.09.
- Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78–84.
- Mazer, J.A., and Gallant, J.L. (2003). Goal-related activity in V4 during free viewing visual search. Evidence for a ventral stream visual salience map. *Neuron* 40, 1241–1250.
- Miller, E.K., and Cohen, J.D. (2001). An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* 24, 167–202.
- Moore, T., Armstrong, K.M., and Fallah, M. (2003). Visuomotor origins of covert spatial attention. *Neuron* 40, 671–683.
- Moran, J., and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science* 229, 782–784.
- Morrison, S.E., and Salzman, C.D. (2009). The convergence of information about rewarding and aversive stimuli in single neurons. *J. Neurosci.* 29, 11471–11483.
- Morrison, S.E., Saez, A., Lau, B., and Salzman, C.D. (2011). Different time courses for learning-related changes in amygdala and orbitofrontal cortex. *Neuron* 71, 1127–1140.
- Motter, B.C., and Belky, E.J. (1998). The zone of focal attention during active visual search. *Vision Res.* 38, 1007–1022.
- Padoa-Schioppa, C., and Assad, J.A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature* 441, 223–226.
- Peck, C.J., Lau, B., and Salzman, C.D. (2013). The primate amygdala combines information about space and value. *Nat. Neurosci.* 16, 340–348.
- Platt, M.L., and Glimcher, P.W. (1999). Neural correlates of decision variables in parietal cortex. *Nature* 400, 233–238.
- Richmond, B.J., Wurtz, R.H., and Sato, T. (1983). Visual responses of inferior temporal neurons in awake rhesus monkey. *J. Neurophysiol.* 50, 1415–1432.
- Rigotti, M., Barak, O., Warden, M.R., Wang, X.-J., Daw, N.D., Miller, E.K., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature* 497, 585–590.
- Roesch, M.R., and Olson, C.R. (2004). Neuronal activity related to reward value and motivation in primate frontal cortex. *Science* 304, 307–310.
- Roesch, M.R., Taylor, A.R., and Schoenbaum, G. (2006). Encoding of time-discounted rewards in orbitofrontal cortex is independent of value representation. *Neuron* 51, 509–520.
- Rolls, E.T. (2015). Taste, olfactory, and food reward value processing in the brain. *Prog. Neurobiol.* 127–128, 64–90.
- Rudebeck, P.H., and Murray, E.A. (2014). The orbitofrontal oracle: cortical mechanisms for the prediction and evaluation of specific behavioral outcomes. *Neuron* 84, 1143–1156.

- Rudebeck, P.H., Mitz, A.R., Chacko, R.V., and Murray, E.A. (2013). Effects of amygdala lesions on reward-value coding in orbital and medial prefrontal cortex. *Neuron* 80, 1519–1531.
- Saez, A., Rigotti, M., Ostojic, S., Fusi, S., and Salzman, C.D. (2015). Abstract context representations in primate amygdala and prefrontal cortex. *Neuron* 87, 869–881.
- Saleem, K.S., Kondo, H., and Price, J.L. (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.
- Sheinberg, D.L., and Logothetis, N.K. (2001). Noticing familiar objects in real world scenes: the role of temporal cortical neurons in natural vision. *J. Neurosci.* 21, 1340–1350.
- Shimojo, S., Simion, C., Shimojo, E., and Scheier, C. (2003). Gaze bias both reflects and influences preference. *Nat. Neurosci.* 6, 1317–1322.
- Strait, C.E., Blanchard, T.C., and Hayden, B.Y. (2014). Reward value comparison via mutual inhibition in ventromedial prefrontal cortex. *Neuron* 82, 1357–1366.
- Strait, C.E., Slezzer, B.J., Blanchard, T.C., Azab, H., Castagno, M.D., and Hayden, B.Y. (2015). Neuronal selectivity for spatial position of offers and choices in five reward regions. *J. Neurophysiol.* 115, 1098–1111.
- Thorpe, S.J., Rolls, E.T., and Maddison, S. (1983). The orbitofrontal cortex: neuronal activity in the behaving monkey. *Exp. Brain Res.* 49, 93–115.
- Towal, R.B., Mormann, M., and Koch, C. (2013). Simultaneous modeling of visual saliency and value computation improves predictions of economic choice. *Proc. Natl. Acad. Sci. USA* 110, E3858–E3867.
- Tremblay, L., and Schultz, W. (1999). Relative reward preference in primate orbitofrontal cortex. *Nature* 398, 704–708.
- Tsujimoto, S., Genovesio, A., and Wise, S.P. (2009). Monkey orbitofrontal cortex encodes response choices near feedback time. *J. Neurosci.* 29, 2569–2574.
- Vaidya, A.R., and Fellows, L.K. (2015). Testing necessary regional frontal contributions to value assessment and fixation-based updating. *Nat. Commun.* 6, 10120.
- Wallis, J.D., and Miller, E.K. (2003). Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *Eur. J. Neurosci.* 18, 2069–2081.
- Wilson, R.C., Takahashi, Y.K., Schoenbaum, G., and Niv, Y. (2014). Orbitofrontal cortex as a cognitive map of task space. *Neuron* 81, 267–279.
- Yasuda, M., Yamamoto, S., and Hikosaka, O. (2012). Robust representation of stable object values in the oculomotor Basal Ganglia. *J. Neurosci.* 32, 16917–16932.