Article

Neuron

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.

McGinty et al., 2016, Neuron *90*, 1299–1311
June 15, 2016
http://dx.doi.org/10.1016/j.neuron.2016.04.045

CellPress

Neuron
Article

Orbitofrontal Cortex Value Signals Depend on Fixation Location during Free Viewing

Vincent B. McGinty,[1,*] Antonio Rangel,[2] and William T. Newsome[1,3]
[1]Department of Neurobiology, Stanford University, Stanford, CA 94305, USA
[2]Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125, USA
[3]Howard 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 (Armel et al., 2008; Krajbich et al., 2012, 2010; Krajbich 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

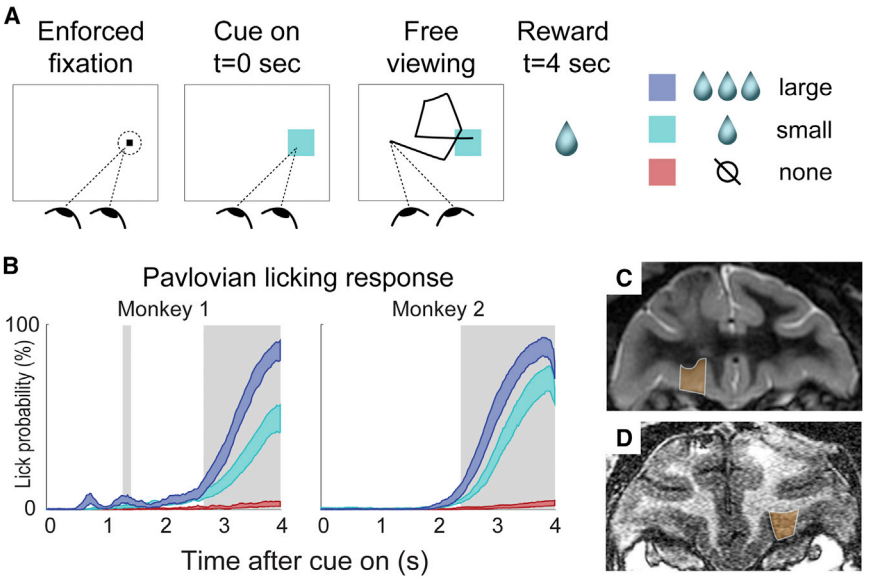

Figure 1. Task, Behavior, and Recording Sites

(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,

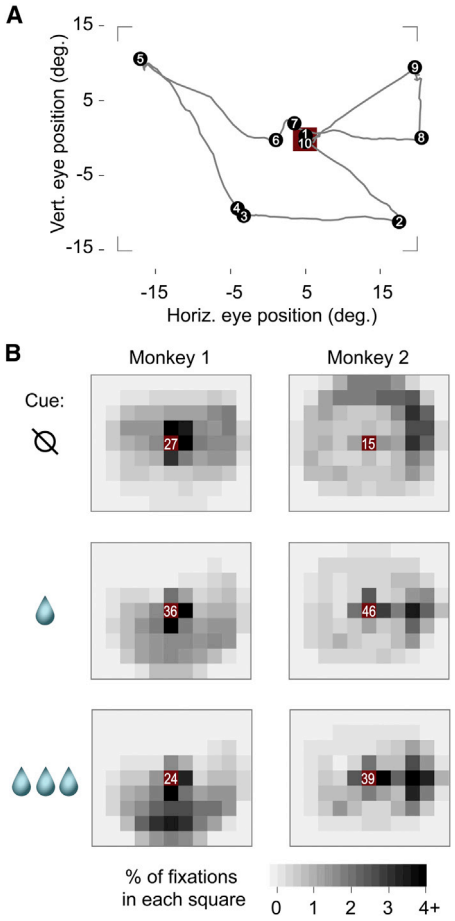

Figure 2. Fixation Locations during Free Viewing

(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

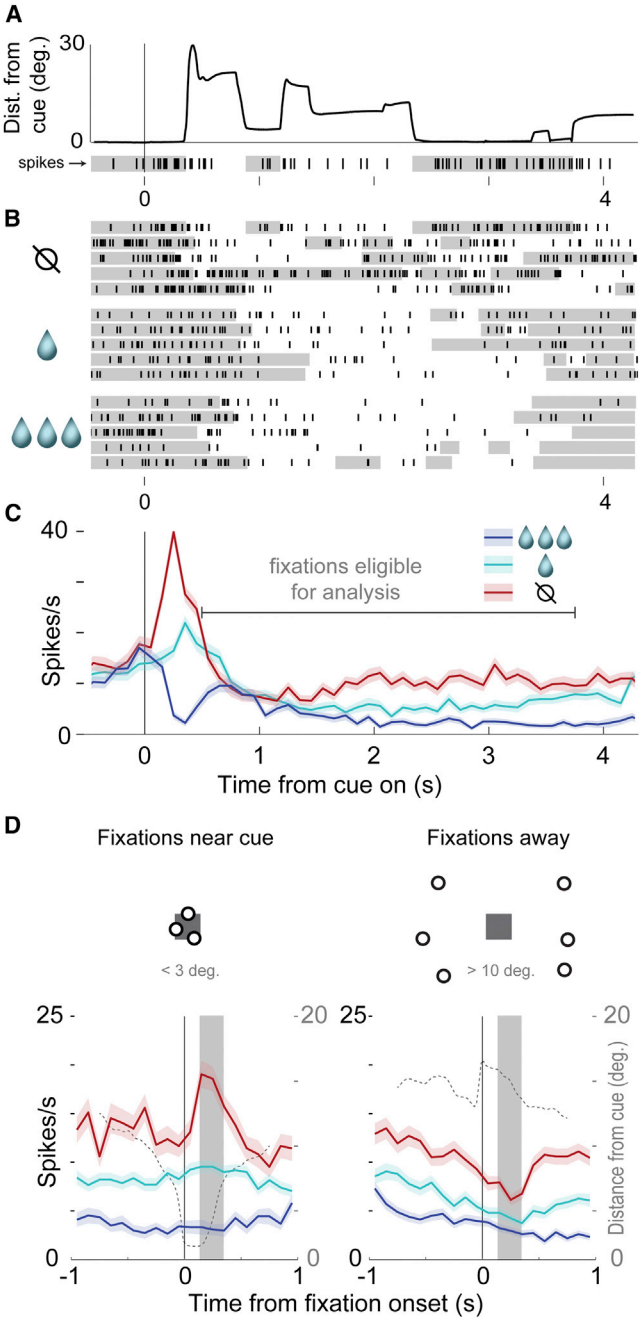

Figure 3. Value and Fixation Location Encoding in a Single Neuron

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).

(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.

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

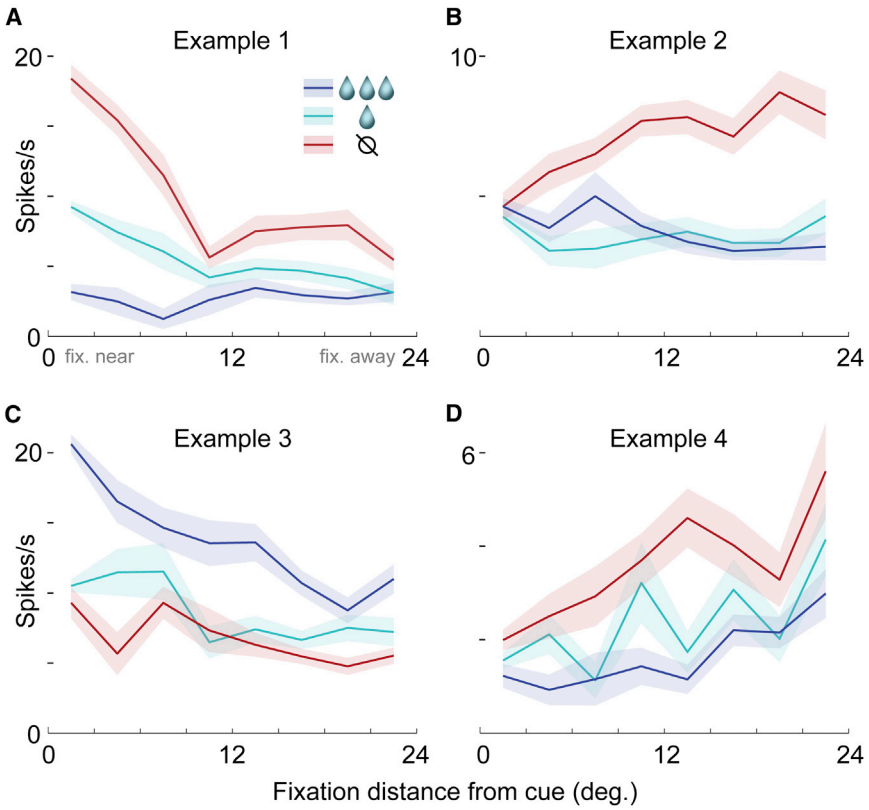

Figure 4. Four Examples of Value and Fixation Location Encoding

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 *rho*, 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 neurons 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.

[a]The 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 $\beta_{VAL \times DIST}$ was also greater than chance (expected 3.2%, actual 6.4%, $p = 0.0002$), as was the co-occurrence of β_{DIST} and $\beta_{VAL \times DIST}$ (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 $\beta_{VAL \times DIST}$ (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 $\beta_{VAL \times DIST}$ 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 and 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"

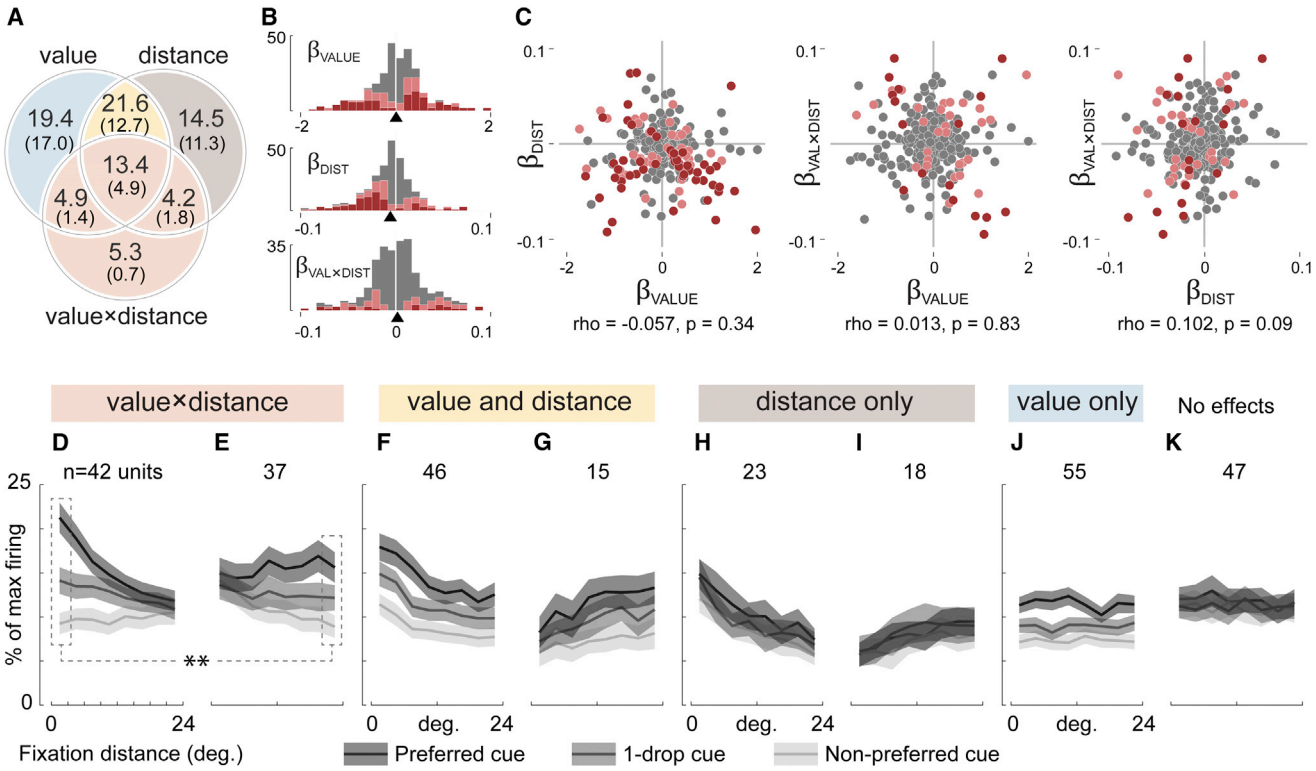

Figure 5. Value and Fixation Distance Encoding in the Population

(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 $p < 0.05$) (D). 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 ($\beta_{VALUE} > 0$) 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 ($\beta_{VALUE} < 0$) 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

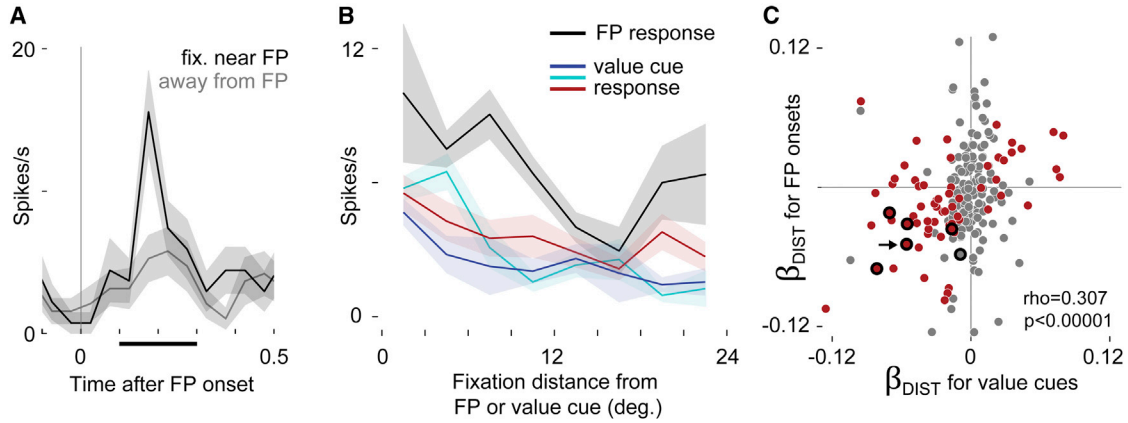

Figure 6. Gaze Distance Modulates Firing Evoked by Fixation Point Onset

(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 $\beta_{DIST\text{-}FP}$, 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 $\beta_{DIST\text{-}FP}$. Because fewer observations were available in the FP data (see text), only 228 neurons had sufficient data to fit the model and calculate $\beta_{DIST\text{-}FP}$ 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, $\beta_{DIST\text{-}FP}$. On average, $\beta_{DIST\text{-}FP}$ 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 $\beta_{DIST\text{-}FP}$ 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). $\beta_{DIST\text{-}FP}$ 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 *rho* below the natural limit of 1.0. The first factor is the inherent noise entailed in estimating β_{DIST} and $\beta_{DIST\text{-}FP}$, which will propagate into the calculation of *rho*. As we show below, this noise was substantial for $\beta_{DIST\text{-}FP}$, 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 $\beta_{DIST\text{-}FP}$.

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

To determine how differences in spatial sampling influence *rho*, 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 timescale) of natural primate gaze behavior.

Second, studies in human subjects show that longer fixation on a given item increases its likelihood of being chosen (Armel et al., 2008; Krajbich et al., 2012, 2010; Krajbich 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 (Krajbich 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:

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

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* × *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:

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

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 *rho*, 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 $\beta_{DIST\text{-}FP}$ 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 $\beta_{DIST\text{-}FP}$ 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.

Armel, 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 inferotemporal 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., Mammo, D., and Newsome, W.T. (2012). Tracking the eye noninvasively: simultaneous comparison of the scleral search coil and optical tracking techniques in the macaque monkey. Front. Behav. Neurosci. *6*, 49.

Krajbich, 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.

Krajbich, I., Armel, C., and Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. Nat. Neurosci. *13*, 1292–1298.

Krajbich, 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., Sleezer, 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.

Neuron, Volume *90*

Supplemental Information

Orbitofrontal Cortex Value Signals Depend on Fixation Location during Free Viewing

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

SUPPLEMENTAL FIGURES

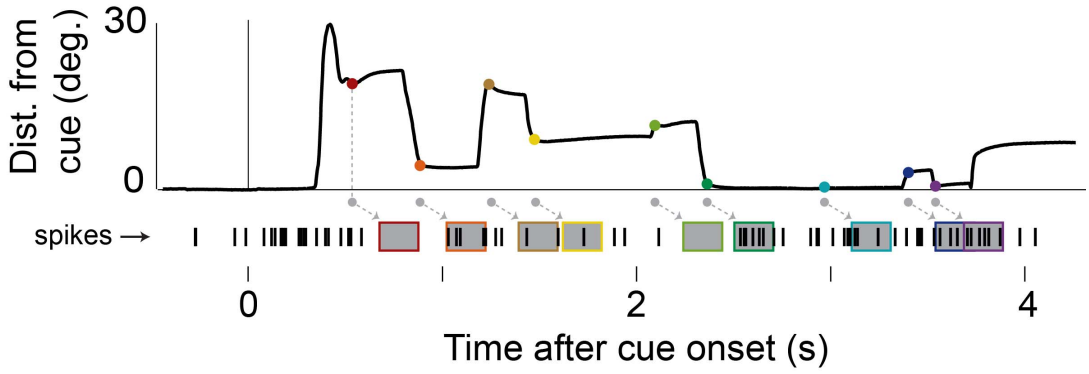

FIGURE S1, related to Figure 3: Extracting fixation-evoked firing from eye position and spiking data: Eye position and neural data in a single trial: the thick black line gives the distance of gaze from the cue, and the raster with black tick marks (below x-axis) shows the spikes of an identified single unit. This is the same data as in Figure 3A. Fixations were detected on the basis of eye velocity (Kimmel et al. 2012). The onset time of each fixation in this trial is marked with a colored dot on the eye trace, and with a corresponding gray dot below the trace. (Note that fixations outside of the 0.50-3.75s eligibility window were not analyzed (Experimental Procedures), and therefore are not marked.) The firing associated with each fixation was calculated as follows: For each fixation, we measured the spike count in a 200ms window that was offset from the beginning of the fixation by a fixed lag. The windows for each fixation are indicated by gray rectangles, with colored borders that correspond to the dots on the eye trace. The lag was determined individually for each neuron based on its activity, as described in the Experimental Procedures; the lag in this example is 140ms, and the median lag across all neurons was 160ms.

Saccades sometimes occurred during the 200ms windows over which spiking was measured (e.g. during the light green window at ~2.3s above). To determine whether this ongoing saccadic activity could contribute to variability in fixation-evoked responses, we asked whether peri-saccadic firing exhibited classical encoding of saccadic vectors (particular combinations of direction and amplitude, (Bruce and Goldberg, 1985)) using a factorial ANOVA. Fewer than 1% of neurons had peri-saccadic encoding of saccade vectors (significant effect at $p<0.05$, corrected), indicating that ongoing saccadic behavior contributed negligibly to the fixation-evoked response.

In rare instances, two firing windows overlapped with one another (e.g. the last two fixations above), so that a single spike could be counted as being evoked by both fixations. However, fewer than 4% of fixations had the potential to produce twice-counted spikes, due to the sparse firing of many neurons and to the 100ms minimum fixation dwell time that we imposed on the data. Separate analyses that excluded these fixations yielded results that were not different from those shown here (not shown).

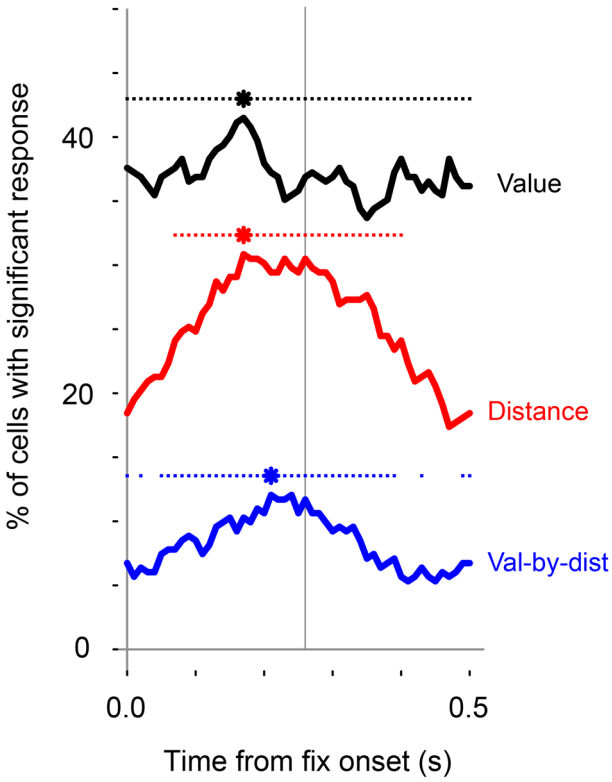

FIGURE S2, related to Figure 5: Time course of value and gaze distance encoding relative to fixation onset: The GLM given by Equation 1 was fit to firing time-locked to fixation onset, using 200ms windows moved in 10ms increments, with window centers ranging from 0 to 500ms relative to fixation onset. For each window, we found the percentage of cells with significant effects ($p<0.05$, corrected within time point, by Holm's method), and these percentages are plotted in the solid black, red, and blue lines. Each asterisk shows the time of the peak percentage for a given variable, and the adjacent dotted lines show times points where the percentages are *not different* from peak levels, by a chi-squared test for proportions (p *greater than* 0.05 after correction within variable, by Holm's method). The gray vertical line shows the center of the typical window used in the main analysis (results in Figures 2-5); this was found by taking the median of all the cell-specific post-fixation windows, which were calculated as described in Experimental Procedures.

For the distance and interaction term variables (red and blue) the maximum percentage occurs approximately 0.2ms after fixation onset, and falls below peak levels in earlier and later time windows (near 0 and 0.5s). The peak at 0.2s is consistent with the typical visual response latency of OFC neurons. The below-peak percentages in the early and late windows is consistent with the fact that gaze distance is changing constantly throughout a trial, meaning that the firing in these time windows reflects the (uncontrolled) gaze distance of fixations occurring before and after the reference fixation. In contrast, encoding of the value variable does not significantly differ from peak levels across the time windows tested. This is expected due to the structure of the task and of the model: within a trial, cue value is constant (unlike gaze distance), and using the GLM it is possible to extract the contribution that value makes to firing *independent* of the contribution of gaze distance, whose effect on firing is captured by the other two regressors.

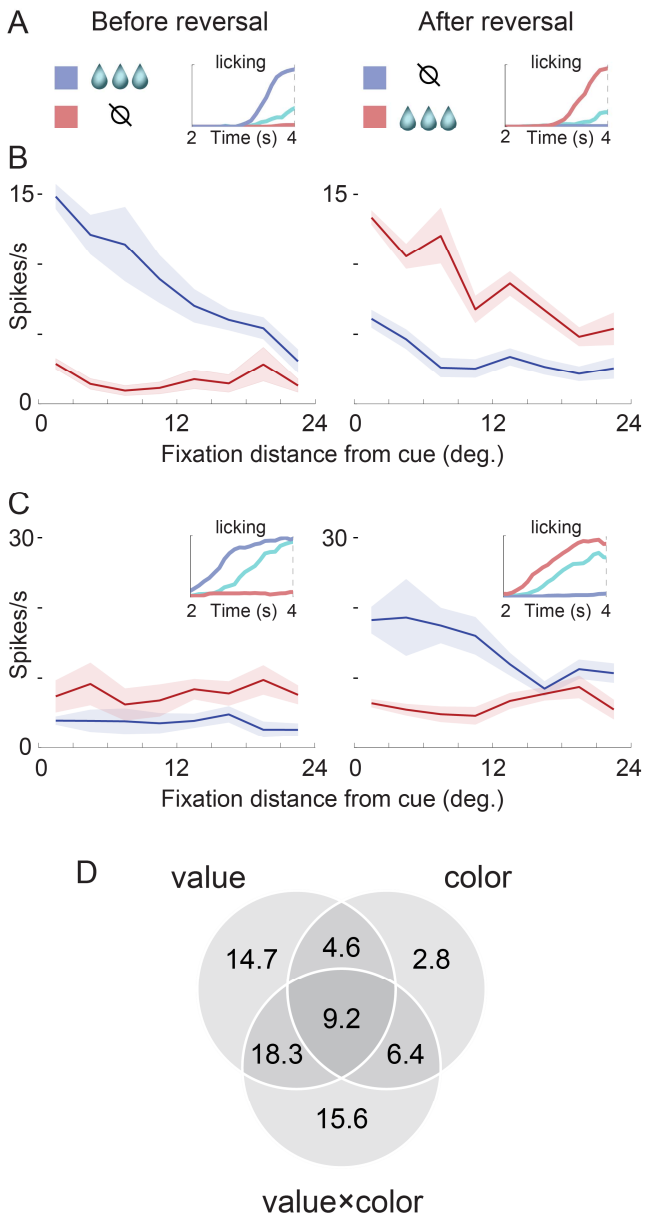

FIGURE S3, related to Figure 5: Color-value reversal test: (A) Reversal of color-reward associations, and the licking responses before and after reversal. (B) Fixation-evoked firing of an identified single unit as a function of fixation location, before and after reversal (same session as in A). Firing for small reward cue is omitted for clarity. (C) A second example of single unit firing before (left) and after reversal (right), with licking responses shown in the inset. (D) Percent of neurons (n = 109 tested) with firing significantly modulated ($p<0.05$ corrected) by cue value, cue color, and the value-by-color interaction, as determined by a GLM. 46.8% of cells were significantly modulated by cue value, 49.5% by the value-by-color interaction, and 22.9% by cue color.

A

% of neurons with effects at p<0.05	Variables of interest			Oculomtor co-regressors		
	value	distance	value-by-dist	saccade ampl.	saccade velocity	dwell time
uncorrected	59.4	48.8	27.9	17.3	12.7	19.8
corrected	35.0	28.6	9.2	2.1	0.7	3.9

B

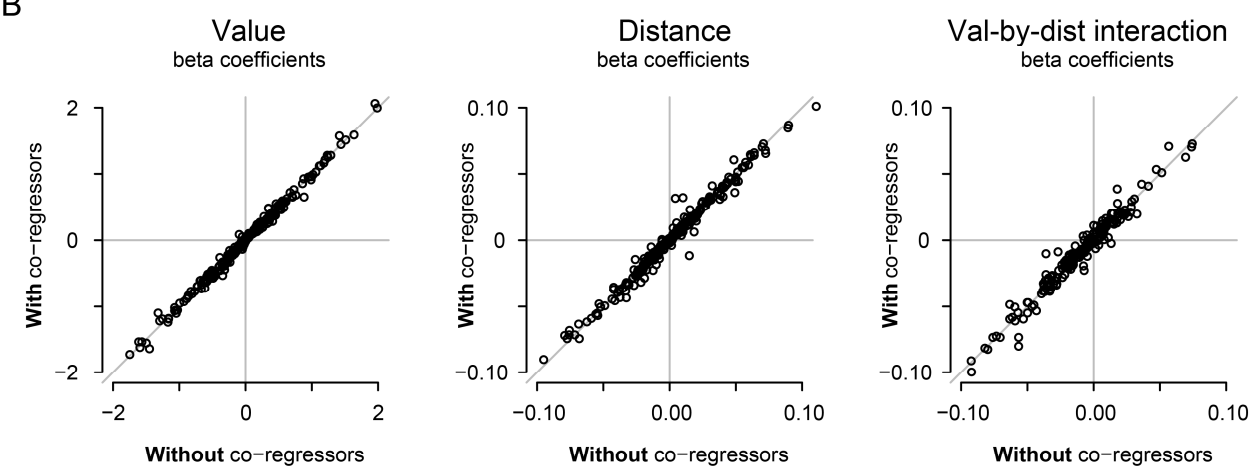

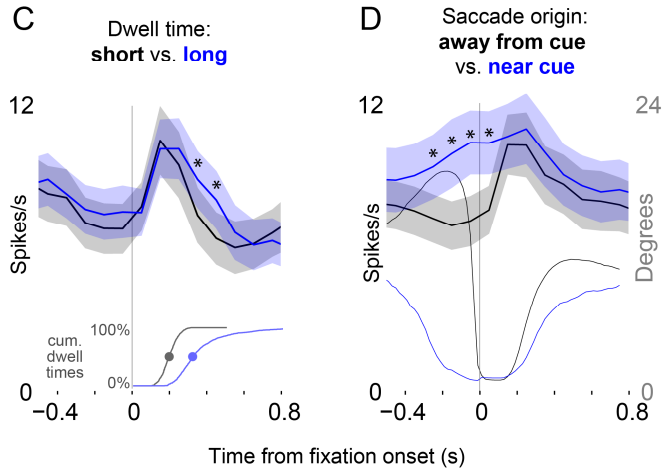

FIGURE S4, related to Table 1: Effects of other oculomotor variables: (A) Percent of neurons with significant effects in a GLM that was similar to the one given by Equation 1, but included three co-regressors that describe oculomotor variables *other* than fixation distance. Including the co-regressors does not substantially change the results for the variables of interest (compare to Table 1). (B) In each plot, the x-axis shows the beta coefficients resulting from the GLM without oculomotor co-regressors (Equation 1, same data as in Figure 5), and the y-axis shows the coefficients obtained with these co-regressors included in the model. All correlations are above $r = 0.98$. (C,

D) To confirm the GLM results in A, in a separate analysis we considered a subset of the data that isolates *excitatory* responses evoked by fixations on or near the cues, and that also holds cue value constant (see Supplemental Experimental Procedures). The PSTHs here show the responses evoked by these on-cue fixations, averaged across cells (mean and SEM). The stars indicate a significant difference between the blue and black curves ($p<0.05$ uncorrected, by Wilcoxon rank sum test). In Panel C, fixations for each cell were median split according to the dwell time of the fixation. Note the overall similar response regardless of dwell time. The thin lines below show the cumulative dwell times across all fixations used in this graph; dots indicate medians. 59 neurons contributed to this analysis. In Panel D, fixations within each cell were divided according to whether the prior fixation was near or away from the cue, effectively dividing them according to the amplitude of the prior saccade. Note the similar response after t = 0, indicating that saccade amplitude did not modulate the response. The large difference in firing prior to t=0 reflects the fact that the fixation immediately prior to t=0 was on or near the cue for the blue data, but was away from the cue for the black data. The thin lines show the eye position relative to the cue, averaged across all fixations within a given condition (right y-axis scale). 69 neurons contributed to this analysis.

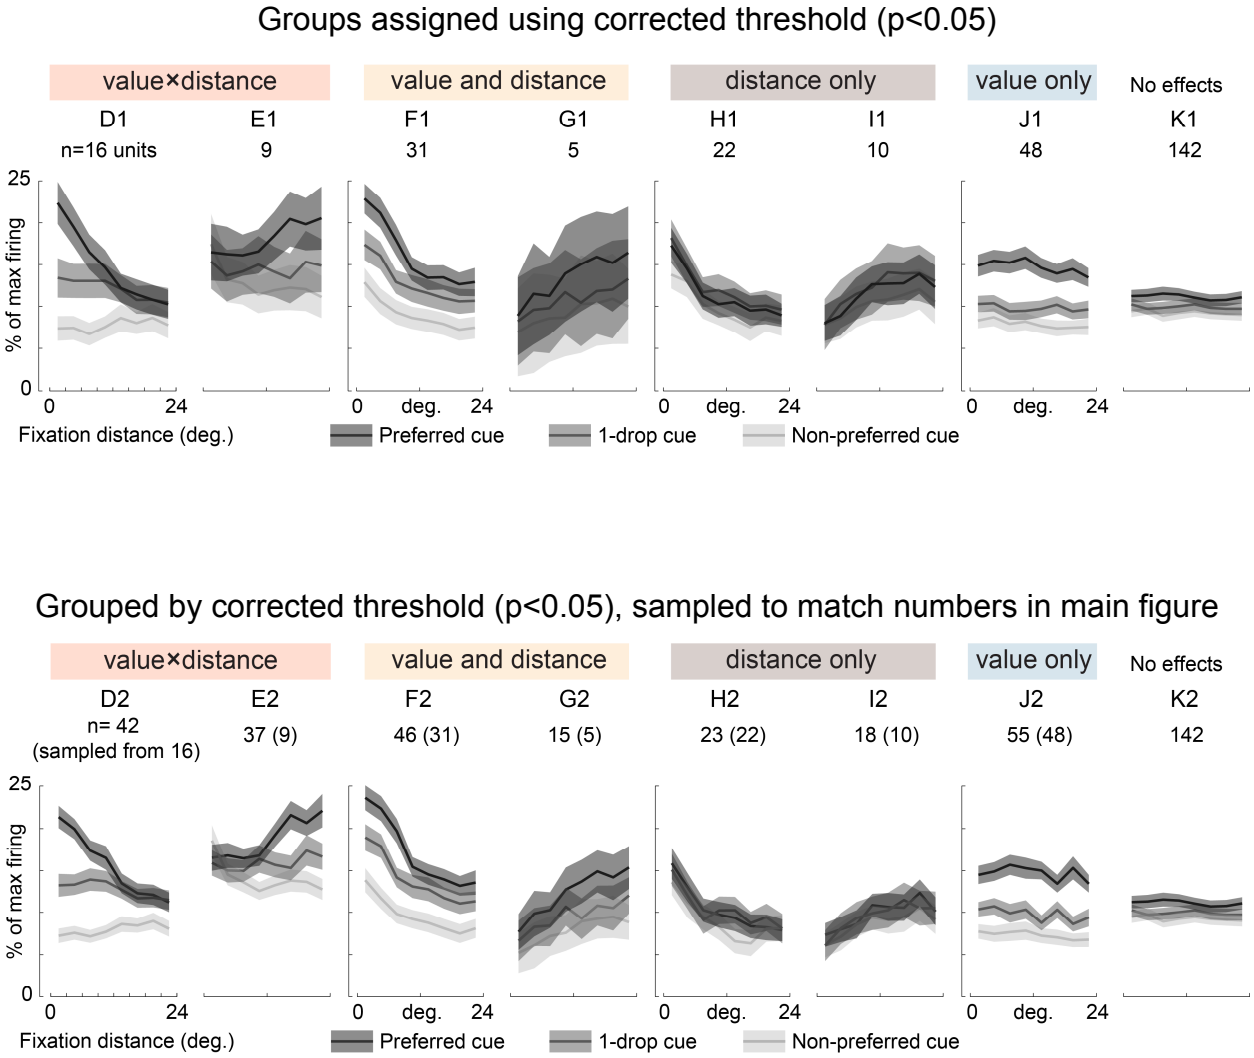

FIGURE S5, related to Figure 5, population firing using cells grouped according to a corrected threshold. In panels D1-K1, the conventions and data are the same as in Figure 5D-K, except that cells are sorted into groups according to the corrected significance threshold, corresponding to the percentages shown in parentheses in Figure 5A. The firing patterns here are very similar to those in the main figure. However, two features bear notice: First, the individual cells in panel K1 are classified as having “no effect”, but as a group they show what appears to be stratification of firing according to cue value. This suggests that one or more cells with nontrivial value effects are misclassified here (due to the use of a corrected threshold), ultimately masking these important population-level effects. To minimize this kind of misclassification and masking of responses, we use an uncorrected threshold in the main figure (Figure 5D-K).

Second, the error bars (shaded regions) are larger in D1-K1 than in the main figure, which could be due to either the reduced number of cells contributing to the graphs (numbers above the graphs) or to the response variability among those cells that do contribute. To determine the effects of cell number here, in panels D2-K2 we re-plot the data in a way that preserves the overall response patterns of panels D1-K1, yet adjusts for the reduced number of neurons. This was done by drawing samples, with replacement, from the cells used in panels D1-K1, until the cell numbers matched those in the main figure. Qualitatively, the error bar sizes in D2-K2 are very similar to those in the main figure, suggesting that the larger error bars in D1-K1 are largely due to the reduced cell number.

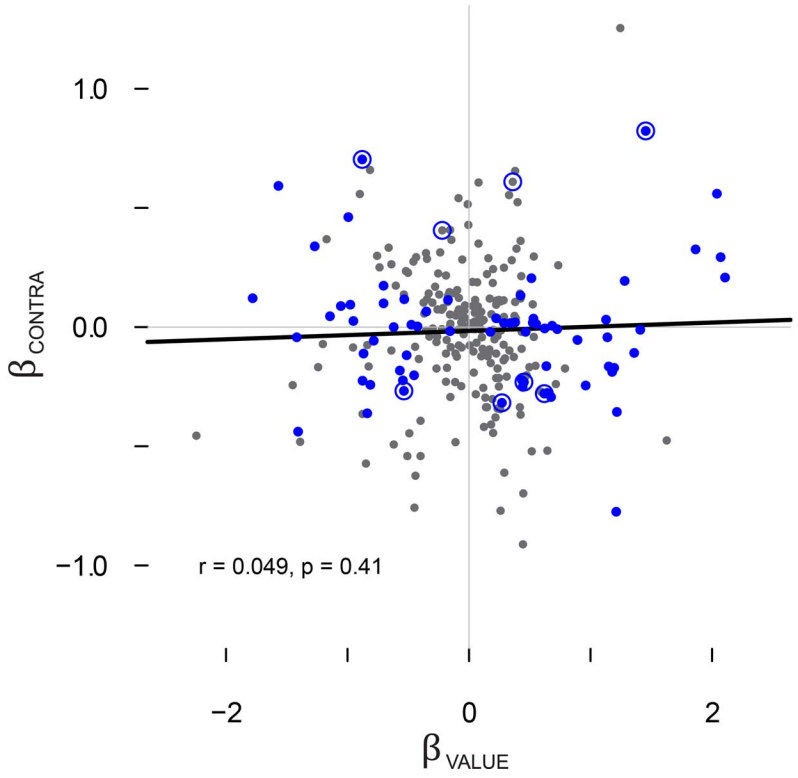

Figure S6, related to Figure 5: value preference compared to preference to contralateral cues. We asked whether a cell's preference for higher- or lower-valued cues (corresponding to positive or negative values of β_{VALUE}, respectively) was correlated with a tendency to fire more or less for fixations that placed the cue in the visual hemifield contralateral to the recording site. See Peck at al. (2013), and compare to their Figure 3B. We fit a GLM that explained fixation-evoked firing as a function of two variables: cue value, and a dummy variable indicating whether, for a given fixation, the cue was placed in the hemifield contralateral to the recording site. To maximize the possibility of detecting neurons with contra/ipsilateral firing preferences, we only used fixations located greater than 5 degrees to the left or right of the cue (which excluded on-cue fixations from this analysis). The resulting regression coefficients (β_{VALUE} and β_{CONTRA}) are compared in the scatter plot above. Each cell is indicated by a point. Solid blue points indicated a significant effect of β_{VALUE} ($p<0.05$, corrected), and blue ring, a significant effect of β_{CONTRA} ($p<0.05$, corrected). Unlike in the amygdala recordings reported by Peck et al. (2013), there is no correlation between these effects. The solid black fit line and statistics refer to Pearson's correlation coefficient. Using Spearman's coefficient, the correlation is rho = -0.080 and p = 0.18).

TABLE S1, ADDITIONAL GLM RESULTS

% of neurons with effects at p<0.05 corrected	Regressors		
	value	distance	value-by-dist
Monkey 1 (n=144)	41.7	43.0	9.7
Monkey 2 (n=139)	30.2*	18.0*	7.9
single unit (n=176)	36.4	31.2	8.5
multi-unit (n=107)	35.5	29.9	9.3
using same firing window for all cells	36.9	30.5	11.7

Table S1, related to Table 1: The percentage of neurons with significant effects in the primary GLM (Equation 1). * indicates significant difference from Monkey 1 ($p<0.05$ by chi-squared test for proportions). To generate the results in the bottom row, spiking was measured using the same 200ms post-fixation firing window in all cells (160-360ms, center at 260ms), which was the median window across all the cell-specific windows used in the main analysis.

SUPPLEMENTAL 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 (13.5-15.0 kg) designated Monkey 1 and Monkey 2. Using aseptic surgical techniques, they were implanted with an MR-compatible head holder; later, a craniotomy was performed allowing access to the right OFC in Monkey 1 and left OFC in Monkey 2, and a recording chamber (Crist Instruments, Hagerstown, MD) was implanted over the craniotomy. The monkeys performed the behavioral task while head-restrained and seated ~57 cm from a fronto-parallel CRT monitor used to display the task stimuli (background luminance ~4 cd/m^2). Eye position was monitored at 400Hz using a scleral search coil system (C-N-C Engineering) in Monkey 1, and a non-invasive optical system in Monkey 2 (Eyelink, SR Research). These different eye tracking methods yield similar data (Kimmel et al., 2012).

A tube for fluid rewards was placed ~2-4mm outside the mouth, and the monkeys could only obtain all of a given reward by touching their tongue to the end of the tube during delivery. Both monkeys quickly learned to retrieve the juice in this way, and typically consumed all of the juice delivered on every trial.

Contact between the tongue and juice tube (the “licking response”) was detected by connecting the input lead of a single channel amplifier (A-M Systems, 400Hz sampling) to the fluid reservoir, and the ground to the seat of the primate chair. Tongue contact abruptly reduced the amplitude of ambient (e.g. 60Hz) noise on the channel, so that the presence/absence of contact could be determined by setting an appropriate noise threshold for each session. The licking-vs.-time plots in Figure 1 and Figure S3 show the percentage of trials in which contact was present at a given time point.

Task flow and stimulus presentation were controlled using the REX software suite (Laboratory of Sensorimotor Research, National Eye Institute) and dedicated graphics display hardware (Cambridge Research Systems). Neural signals were measured from single tungsten electrodes (FHC Inc., Bowdoin, ME) placed at the target

locations using a motorized drive (NAN Instruments, Nazareth, Israel). Neural activity, eye position, and task event data were acquired and stored using a Plexon MAP system (Plexon, Inc., Dallas, TX).

Behavioral task

In every session the subjects performed a behavioral task in two separate parts: an initial conditioning phase, in which subjects learned arbitrary cue-reward associations, and subsequent neural recording phase, in which we measured OFC activity during task performance.

In this task, we used a form of Pavlovian conditioning (Morrison and Salzman, 2009) to train the monkeys to associate three different color cues with the delivery of three juice volumes: ~3 drops (large reward), 1 drop (small reward), and 0 drops (no reward). We used a ~3:1:0 ratio to elicit three distinct levels of behavioral response in the subjects (licking, see Figure 1B). The juice volumes were constant within a session, but the ratio of large to small reward volume varied across sessions (2.5-2.7 for Monkey 1, and 2.5-3.0 for Monkey 2), to compensate for small changes in the subjects' fluid sensitivity across the duration of the study.

The trial structure in both phases was identical, and illustrated in Figure 1A. Trials began with the presentation a gray square fixation point (FP, 0.5 degrees per side), placed randomly 5 degrees to the left or right of the screen center. After the FP was fixated for a randomly chosen interval of 1-1.5 seconds, one of the three color cues was selected at random and presented at the location of the FP. The cues were square color patches, 3.2 degrees per side, and mutually equiluminant at ~22 cd/m^2. Once the cue appeared, the monkey was free to move his eyes for the duration of the trial. Gaze position was monitored while the cue was shown, but fixations had no consequence for trial outcome, which was perfectly predicted by the cue. On large and small reward trials, juice delivery began (fluid solenoid opened) exactly 4 seconds after cue onset. In all trials, the cue remained visible until about 4.3s after onset, corresponding to the end of reward delivery (solenoid closing) for the large reward cue. Trials were separated by a random 2-4 second inter-trial-interval (ITI), which lasted from the previous cue offset to the appearance of the next FP.

New cue colors were chosen at the beginning of each session (by randomly sampling equidistant points on a color wheel within the CIELUV color space), requiring the animals to learn the new color-reward associations prior to data collection. Learning was assessed each session by measuring the total duration of tongue contact with the juice tube (licking response) in the 4 second period *prior* to juice delivery (Fiorillo et al., 2008; Morrison and Salzman, 2009). This anticipatory licking was indiscriminate at the start of a new session when reward associations were unknown (not shown), and then with subsequent trials became commensurate with the reward size: large > small > no reward, with no reward ≈ 0 seconds. The initial conditioning phase was terminated when licking durations over the prior 60-100 trials were significantly different for all three trial types (rank sum test, $p<0.01$ uncorrected). See Figure 1B for average licking after initial conditioning was complete.

After the termination of the conditioning phase, the subjects continued to perform the task while we isolated and recorded suitable neural activity (see below). Thus, neural data was collected only *after* the cue-reward pairings had stabilized.

During some neural recordings, we executed a “reversal test”, in which the cue-reward associations of the no-reward and large reward cues were switched abruptly and without warning (Morrison and Salzman, 2009; Thorpe et al., 1983). The monkeys typically learned the new associations within 5-10 presentations of each cue, which we confirmed by assessing post-reversal licking behavior with a rank sum test, as was done for the initial cue-reward conditioning. Reversals occurred only during the neural recording phase, typically after data had been obtained for 300-600 trials under the initial cue-reward associations. By recording neural responses both before and after reversal, we were able to assess the effects of cue color (independent of value) on OFC neural activity.

With the exception of the reversal test, we took several steps to ensure that subjective cue value remained stationary over the period during which neural data were collected and analyzed. First, when a reversal was performed, the first 40 trials after reversal were not used for neural data analysis, which was sufficient time for the new cue-reward associations to be learned, and for OFC responses to adapt (Morrison et al., 2011). Second, 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. Third, a session was discarded if the licking responses did not maintain selectivity after data collection began; these sessions were rare. Thus, the main analyses of this paper use data in which cue-reward associations were well learned and did not change within the experimental session.

Recording and data collection

We recorded from 283 neural unit signals in OFC, 144 from Monkey 1 and 139 from Monkey 2. Single electrodes were introduced into the brain through a sharpened guide tube whose tip was inserted 1-3mm below the dura. OFC was identified on the basis of gray/white matter transitions, and by consulting a high-resolution MRI acquired from each animal after chamber implantation. We targeted the fundus and lateral bank of the medial orbital sulcus and the laterally adjacent gyrus (Figure 1C,D), a region targeted in several other studies (Morrison and Salzman, 2009; Padoa-Schioppa and Assad, 2006; Tremblay and Schultz, 1999; Wallis and Miller, 2003), corresponding approximately to Walker's area 13 (Öngür and Price, 2000).

We typically used two electrodes simultaneously at different sites within the OFC. After initial conditioning (and while subjects continued to perform the task) we searched for neural signals using the following protocol: First, we slowly advanced the electrodes through the OFC until we isolated putative "single units", indicated by large and well-isolated waveforms, on both electrodes. Next, we collected data for 40-60 pilot trials, and used them to generate peristimulus firing rate histograms (PSTHs) of firing time-locked to cue onset, averaged across all trial types (cue values). Based on visual inspection of the PSTHs, we determined whether any of the single units showed "task-related firing", broadly defined as any apparent increase, decrease, or other change in firing rate at any point during the 0-4s cue display period. Because the PSTHs were averaged across trial types, they did not show whether cells encoded cue value, and so value encoding was not a criterion used to assess task-related firing. Likewise, fixation data were not used or referenced when constructing these PSTHs, so gaze encoding was not a criterion used to assess task-related firing. If any single unit on either electrode showed apparent task-related firing, data collection began for both electrodes.

(Thus, not all single units in the data set showed task related firing.) Otherwise, the units were abandoned, and the search continued. The process was repeated until at least one task-responsive single unit was found.

Importantly, in addition to single unit signals, we also collected any "multi-unit" signals (low amplitude, poorly isolated waveforms) that were present at the same time. Note that multiunit signals were *not* subject to the online screening described above: if they were present at the time data collection began, they were recorded regardless of their task-related activity, and then analyzed alongside the single units. As a result, this data set contains single units that were screened for task-related activity according to the broad criteria above, as well as unscreened multi-unit signals.

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 data collection, spikes were assigned offline to individual units based upon the principal component features of the waveforms (Plexon Offline Sorter 2.0). On rare occasions, neurons initially designated as single units were re-categorized as multi-unit signals, if they showed an abundance of short inter-spike intervals (more than 0.05% of intervals below 2ms). 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. After unit sorting, the data were imported into MATLAB and the R software environment for analysis.

Data Analysis

Overview 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, from the eye position data we detected individual fixations (periods of stationary gaze), and counted the spikes in a 200ms window following fixation onset (see below). We call these spike counts the "fixation-evoked firing" or "fixation-evoked response" (Figure S1), and they served as the fundamental unit of observation for most

analyses. We then used generalized linear models (GLM's) to quantify the effect of cue value and gaze location on each cell's firing. For some analyses, we fit GLM's to data in which the fixation-evoked firing rates were randomly permuted across the observations, in order to test null hypothesis that firing was random with respect to the independent variables.

Unless otherwise specified, 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 *rho*, 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 $p<0.05$.

Fixation detection, data selection, and extraction of fixation-evoked firing. The basic unit of analysis was an individual fixation, given by a period of stable gaze bracketed by saccadic eye movements. Saccades and fixation epochs were detected using methods described by Kimmel et al. (2012), based upon the work of Engbert and Kliegl (2003). Briefly, in each trial we calculated the variance of the horizontal and vertical components of eye velocity within that trial; we then established a velocity threshold that was defined as the ellipse whose radii were 6 times the horizontal and vertical variances. Saccades were defined as epochs >5ms in which velocity exceeded the threshold ellipse, and fixations were defined as the epochs of stable gaze between saccades. Fixation onset was defined as the first time-stamp at which the velocity decreased to within the threshold ellipse, unless the saccade appeared to slightly overshoot the cue location (occurring in 51.1% of fixations used for analysis), in which case fixation onset was the point of maximal overshoot – i.e. the beginning of the first "ring" phase, described in detail in Kimmel et al. (2012).

This procedure allowed us to compute a key variable: the *time of onset of each fixation*, which we used to determine which fixations were eligible for analysis, and as a reference time for measuring fixation-associated firing (details below, see also Figure S1). For a fixation to be eligible for analysis, it had to meet three criteria. First, its *onset* had to occur between t=0.5 and 3.75 seconds after cue onset, to exclude from analysis any firing related to cue onset or reward delivery (at t=0 and 4 seconds, respectively).

Note that this explicitly *excludes* fixations that begin before t=0.5s but continue beyond this time. Second, the fixation location had to be within 24 degrees of the cue center, and the preceding saccade amplitude had to be < 35 degrees, to stay within the calibrated range of the eye tracker. Third, fixation duration had to exceed 100ms. This, combined with the time necessary to perform a saccade (mean ~60ms), allowed for consecutive fixation onsets to be separated in time, so that spiking associated with two consecutive fixations could be distinguished from one another.

To simplify the analyses, we focused on firing *after* the onset of each fixation. This is justified by the fact that OFC neurons typically respond to visual stimuli at a latency of 100ms or more (presumably due to obligatory visual processing (Kravitz et al., 2013)), and is similar to the approach used by others to assess ventral visual stream neural activity during free viewing (DiCarlo and Maunsell, 2000; Sheinberg and Logothetis, 2001).

In particular, for each fixation we computed the fixation-evoked firing, illustrated in Figure S1, which was the spike count within a 200ms window following the onset of each fixation. 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. The time window was defined as follows: We collapsed across all trial types (cue values), and identified fixations onto the cue (< 3 degrees from center) that were immediately preceded by fixations away from the cue (> 3.5 degrees). These instances were chosen in order to capture at a coarse level how a cell responds to an abrupt change in visual input – i.e. moving gaze from a non-cue location onto the cue. We then constructed a PSTH aligned to the onset of these on-cue fixations (10ms resolution), and measured average spiking from 0 to 600ms after fixation onset. Within this range, we then searched for the 200ms window with the greatest change in firing rate (*increase or decrease*) relative to the average firing over the 1500ms preceding fixation onset – a long "baseline" interval that by design averages over many prior fixations. Thus, the earliest possible window began at 0ms and ended at 200ms after fixation, and the latest possible window began at 400ms and ended at 600ms. The median analysis window across all cells began at 160ms and ended at 360ms.

The primary analyses in this paper use the cell-specific firing windows as described above. However, we also performed some analyses using a fixed post-fixation window for all cells (Table S1), or using all time windows from 0 to 600ms (Figure S2).

Because the firing window was 200ms in duration, it was possible for a single spike to be counted as being evoked by *two* fixations, if they began within 200ms of one another. However, fewer than 4% of fixations had the potential to produce twice-counted spikes, due, primarily, to the sparse firing of many neurons, and the fact that fixation onsets were separated in time by a 100ms minimum dwell duration (above) plus the time needed for the intervening saccade (~60ms on average). Separate analyses that excluded these fixations (not shown) yielded results that were virtually the same as those shown here.

Spiking data was not normalized or smoothed, except in Figure 5D-K, where the spike counts were scaled within each neuron to between 0 and 100%, measured across all fixations.

Main General Linear Model. Our main results are based on the estimation of the following GLMs, which assumed that fixation-evoked spike counts follow a negative binomial distribution. This is a count-based distribution for which the variance is equal to or greater than the mean (McGinty et al., 2013; Venables and Ripley, 2002), which is often the case in cortical neurons (Ardid et al., 2015; Churchland et al., 2010). Although the neural responses to value and gaze distance were not linear in all cells, preliminary inspection of the data indicated that treating them as linear was a good approximation for most neurons.

The specification of the main GLM is given by

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

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 × Dist* is the interaction of the Value and Distance variables (computed after centering them).

Note that this model codes fixation location as a single variable: the angular distance of fixation from the cue. As described below, we carried out additional analyses to test for other spatial representation schemes (see below).

The GLMs were estimated for each neuron separately. We then carried out population-level comparisons using a variety of tests. Unless otherwise specified, 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 *rho*, 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 $p<0.05$.

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

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

The second used horizontal and vertical distances to the cue, which leads to the following specification:

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

The relative fit of the models was evaluated by comparing the goodness of fit for each alternative model to the one for the main model, using Akaike's information criterion (AIC).

Comparison of initial cue-evoked value signal with value signal during free viewing

Results in main text. Firing was measured 50-500ms after cue onset, and a GLM was fit in which this firing was explained by a single regressor: cue value. To assess value

coding during the cue viewing period (0.5-3.75s), a GLM was fit explaining fixation-evoked firing as a function of cue value, using *only* the fixations <3 degrees from the cue. These fixations were used, because, at onset, the cues were always presented at the center of the fixation window and therefore near the center of gaze; we wished to select from the cue viewing data a subset of fixations that likewise placed the cue near the gaze center. The beta coefficients for cue-evoked and fixation-evoked value encoding were compared using Spearman's correlation coefficient.

GLM for color/value reversal test Results in Figure S3. For the GLM applied to the reversal test data, the three regressors were cue *Value*, the cue *Color* (coded as a two-level factor), and the *Value-by-Color* interaction. Significance was assessed at the $p<0.05$ level, corrected.

GLM with oculomotor co-regressors Results in Figure S4A-B. We asked whether the GLM results shown in Table 1 and Figure 5 could be explained by oculomotor variables other than the distance of fixation from the cue. For each cell we fit a GLM using the same terms as Equation 1, but with three additional regressors that describe other oculomotor features: the amplitude of the saccade prior to the fixation, the velocity of that saccade, and the dwell time of the fixation. The percentage of significant effects were calculated in the same manner as for Table 1.

Effects of fixation dwell time and saccade amplitude on fixation-evoked excitation Results in Figure S4C-D. The GLM with oculomotor co-regressors described above indicated that very few neurons significantly encoded fixation dwell time or amplitude of the preceding saccade, suggesting that these two oculomotor variables contribute only negligibly to fixation-evoked firing. To confirm this, we selected a subset of the data designed to hold the cue value and fixation location constant, and asked whether fixation-evoked firing was modulated by either dwell time or saccade amplitude. The data subset was constructed as follows: First, only fixations onto the cues were used (distance < 3 degrees); these were chosen because they are the most frequently visited single location (Figure 2). Second, the data were restricted to isolate only neurons with

excitatory responses to on-cue fixations (effect of β_{DIST} $p<0.01$, uncorrected), because they were more prevalent than neurons inhibited by on-cue fixations (Figure 5B, D). Third, within each cell, only the trial type (cue value) that produced the largest fixation-evoked excitation was used, in order to keep cue value constant. This resulted in a subset of 69 neurons with sufficient data for analysis.

Thus, by design, the average fixation-evoked response in this data subset was excitatory; the question we asked is how this excitation is modulated by dwell time and saccade amplitude. To assess saccade amplitude effects (results in Figure S4D), the fixations for each were cell were divided according to whether the prior fixation was near the cue (< 3 deg., blue) or away from the cue (> 10 deg., black, n=69); because all fixations in the subset were < 3 degrees from the cue, dividing by prior location in this way effectively divides them according to saccade amplitude. To assess the dwell time effects (results in Figure S4C), the data subset was further refined to include only fixations that were immediately preceded *and followed* by fixations away from the target (yielding 59 neurons with sufficient data), so that consecutive fixations onto the cue were excluded. Then, for each cell, the fixations were median split into those with "short" and "long" dwell times, and the average responses were plotted in PSTHs.

Permutation tests. 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. We thereby establish the "chance levels" of the statistics shown in Table 1, with the following procedure: First, we created 1000 permuted data sets by randomly shuffling the spikes in every cell in the population 1000 times. We then fit GLMs to the permuted cells. Then, for each variable of interest, we calculated the percentage of significant effects (which are spurious by due to the permutation) in each of the 1000 sets; the maximum percentage across all sets was taken to be the maximum expected by chance. This provides a very conservative estimate of the different statistics that could be observed by chance, and leads to a confidence level equivalent to $p<0.001$.

Responses to the FP We also examined firing evoked by the onset of the fixation point (FP) at the beginning of each trial, to test whether these responses were also modulated by the distance of fixation from the stimulus. For this data, we 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 $\beta_{DIST\text{-}FP}$ were then compared to the estimates of β_{DIST} described above (Main GLM, Equation 1), on a cell-by-cell basis.

We measured firing time-locked to each FP onset by counting the spikes within the same cell-specific 200ms windows described above. Analysis was restricted to include FP onsets for which the gaze was < 24 degrees from the FP (matching the range for fixations in the value cue data), was stationary at FP onset, and remained stationary for at least 100ms thereafter (matching the minimum fixation dwell time in the value cue data). We required at least 50 such acceptable FP onsets for each cell, the minimum necessary to fit the GLM. This yielded 228/283 cells eligible for this analysis. The resulting estimates of $\beta_{DIST\text{-}FP}$ were compared to the β_{DIST} using Spearman's correlation coefficient.

To interpret the resulting correlation, we asked whether the observed correlation between β_{DIST} and $\beta_{DIST\text{-}FP}$ 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. To estimate this upper bound, we calculated the reliability of each data set (R_{DIST}, R_{FP}) using a split halves procedure and Spearman-Brown correction. Assuming that the underlying effects are identical in the two contexts, the upper limit for the correlation is given by the square root of R_{DIST}*R_{FP} (See Supplemental Experimental Procedures and Nunnally (1970)). As a complement to this calculation, we also used a resampling procedure to assess how the β_{DIST} for the value cue data would correlate with *itself,* but under sampling conditions that that closely approximated the ones used to estimate $\beta_{DIST\text{-}FP}$, in terms of number of observations and the distribution of gaze distances (see Supplemental Experimental Procedures).

Finding the upper limit on the correlation between gaze effects in Figure 6C Results in main text. We used two methods to estimate the theoretical upper limit on the

correlation between β_{DIST} and $\beta_{DIST\text{-}FP}$ (reported in Figure 6C), a limit that results from the inherent noisiness of the two sets of coefficients, as well as differences in the sampling of visual space in the FP onset and cue viewing contexts. Both methods used the value cue data in Figure 5, and the FP onset data set described in the main text.

Method 1, reliability statistics**:** To estimate the reliability of β_{DIST}, we use the "split-halves" procedure: the data in each neuron are split into even and odd trials, and the main GLM (Equation 1) is run on each half of the data (Nunnally, 1970). This yields two sets of beta values, and we find the Spearman's correlation (*rho*) between these two halves. We then apply the Spearman-Brown correction to this statistic:

$$\frac{n * rho}{(1 + (n - 1) * rho)}$$

using $n = 2$ to scale to the full size of the data (Nunnally, 1970). This gives the expected *reliability* (R_{DIST}) of the full data set, i.e. the correlation that could be expected between the existing data, and a new data set of the same size (same number of trials per cell) collected from the same cells. Using the same procedure, we find the reliability of the distance beta coefficients in the FP onset data (R_{FP}).

Given the reliability statistics from two data sets *A* and *B*, one can find the theoretical maximum correlation that could be expected between *A* and *B.* This is given by:

$$\bar{r}_{AB} * \sqrt{R_A * R_B}$$

where R_A and R_B are the respective reliabilities of *A* and *B*, and $\bar{r}_{AB}$ is the true correlation of the *underlying effects* that generate *A* and *B* (Nunnally, 1970). If we assume that the effects in A and B are identical, then $\bar{r}_{AB}$ is equal to 1, and the maximum observed correlation in the data is the square root of the product of the reliabilities.

Method 2, resampling: Resampling proceeded as follows: for a given cell, the data in both sets were divided according to the fixation distance from the cue/FP, into 0.5 degree bins. Then, for each bin, samples were drawn *with replacement* from the

value cue data, until the number of samples was equal to the number of observations in that particular distance bin in the *FP data*. The proportion of no reward, small reward, and large reward cues approximated the proportion in the original value cue data.

A GLM (Equation 1) was then fit to these resampled value cue data, and the resulting β_{DIST} was compared to the original β_{DIST} (from Figure 5) at the population level using Spearman's correlation coefficient. This resampling, estimation, and correlation was performed 500 times to produce a distribution of correlations between the original and resampled data. (Resampled data sets differ from one another because sampling is performed with replacement.) We then interpreted the resulting correlations as the practical upper bound for the correlation that could possibly be observed between the β_{DIST} estimates obtained from the value cue and FP data sets.

Because the resampling was performed with replacement using discrete gaze distance bins, even a data set that is resampled to match the original (at a precision of 0.5 degrees) would likely have a correlation with the original data of < 1.0, implying that this method itself imposes an upper limit to the self-similarity that can be measured in our data. To estimate this limit, we performed the above procedure (500 replications) on value cue data resampled to match *its own* number of observations and gaze distance distribution. The median self-correlation was *rho* = 0.928, with 99% of observations between 0.895 and 0.953. The fact that this correlation is < 1.0 means that the resampling procedure itself introduces an artefactual "loss" of self-correlation, even when the resampling attempts to match the original. However, this degree of artefactual loss is small compared to the decrease in self-correlation that results when the value cue data were sampled to match the FP data (see text).

SUPPLEMENTAL REFERENCES

Bruce, C.J., Goldberg, M.E., 1985. Primate frontal eye fields. I. Single neurons discharging before saccades. J. Neurophysiol. 53, 603–635.

Engbert, R., Kliegl, R., 2003. Microsaccades uncover the orientation of covert attention. Vision Res. 43, 1035–1045. doi:10.1016/S0042-6989(03)00084-1

Kimmel, D.L., Mammo, D., 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. doi:10.3389/fnbeh.2012.00049

Nunnally, J.C., 1970. Introduction to Psychological Measurement. McGraw-Hill Book Company.

Sheinberg, D.L., 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.