

Orbitofrontal Cortex Uses Distinct Codes for Different Choice Attributes in Decisions Motivated by Curiosity

Highlights

- Monkeys sacrifice water reward for advance information about gamble outcomes
- OFC neurons signal both water amount and informativeness of a gamble
- OFC neurons do not tend to integrate these variables to code subjective value
- OFC neurons respond differently to water rewards and cues predicting those rewards

Authors

Tommy C. Blanchard, Benjamin Y. Hayden, Ethan S. Bromberg-Martin

Correspondence

benhayden@gmail.com

In Brief

Blanchard et al. find that monkeys are curious, willing to sacrifice some water for advance information about gamble outcomes. OFC neurons signal water amount and informativeness of gambles, but do not tend to integrate these variables into a common currency.



Orbitofrontal Cortex Uses Distinct Codes for Different Choice Attributes in Decisions Motivated by Curiosity

Tommy C. Blanchard,¹ Benjamin Y. Hayden,^{1,3,*} and Ethan S. Bromberg-Martin^{2,3}

¹Department of Brain and Cognitive Sciences and Center for Visual Science, University of Rochester, Rochester, NY 14627, USA

²Department of Neuroscience and Kavli Institute for Brain Science, Columbia University, New York, NY 10032, USA

³Co-senior author

*Correspondence: benhayden@gmail.com

<http://dx.doi.org/10.1016/j.neuron.2014.12.050>

SUMMARY

Decision makers are curious and consequently value advance information about future events. We made use of this fact to test competing theories of value representation in area 13 of orbitofrontal cortex (OFC). In a new task, we found that monkeys reliably sacrificed primary reward (water) to view advance information about gamble outcomes. While monkeys integrated information value with primary reward value to make their decisions, OFC neurons had no systematic tendency to integrate these variables, instead encoding them in orthogonal manners. These results suggest that the predominant role of the OFC is to encode variables relevant for learning, attention, and decision making, rather than integrating them into a single scale of value. They also suggest that OFC may be placed at a relatively early stage in the hierarchy of information-seeking decisions, before evaluation is complete. Thus, our results delineate a circuit for information-seeking decisions and suggest a neural basis for curiosity.

INTRODUCTION

Decision makers are often confronted with the opportunity to make choices that provide information about the world (Gottlieb et al., 2013). This information generally comes at a cost, even if it's just the opportunity cost associated with foregoing other possible options. Nonetheless, information is so useful that we may be endowed with a basic drive to seek it out, even when it serves no obvious immediate purpose (Loewenstein, 1994). This drive for information is poorly understood, but is relevant for understanding learning, decision making, and social interactions (Gottlieb et al., 2013). Recent studies have begun to identify the structures associated with curiosity and motivated information seeking more generally (Bromberg-Martin and Hikosaka, 2009, 2011; Gruber et al., 2014; Kang et al., 2009; Phillips et al., 2012).

Like any other good, information can enter into decision-making processes and influence our reward-based (i.e., economic)

decisions. For example, monkeys performing an *information-seeking* choice task will preferentially choose to have the outcomes of risky gambles revealed immediately, rather than to remain in a state of uncertainty while waiting for the outcome to be delivered (Bromberg-Martin and Hikosaka, 2009, 2011). This behavior can be modeled in a standard economic framework in which monkeys integrate two dimensions of an option (here, its information content and the volume of its primary reward, e.g., water or juice) into a single scale to create a single dimension of subjective value. The monkey's subjective value then serves as the basis for its choices (Padoa-Schioppa, 2011).

The fact that information and primary reward are integrated to produce behavior suggests that they are also integrated neurally. Because information and primary reward such as food and water are distinct in many respects (visual versus gustatory, abstract versus appetitive, etc.), they are presumably first detected by different neural systems, then combined to create a common scalar value signal (Levy and Glimcher, 2012; Montague and Berns, 2002; Padoa-Schioppa, 2011; Padoa-Schioppa and Assad, 2006; Raghuraman and Padoa-Schioppa, 2014). Data from monkeys performing the information-seeking task suggest that one neural instantiation of this value scale may be the firing patterns of neurons that encode reward prediction errors (RPEs) (Bromberg-Martin and Hikosaka, 2009, 2011). Specifically, during this task, RPE-coding cells generate similar signals for both primary reward and informational reward, a pattern found in both midbrain dopamine neurons (DA neurons) and one of their major inputs, the lateral habenula (LHb) (Bromberg-Martin and Hikosaka, 2009, 2011). These data suggest that integration of different value types onto a single scale occurs prior to the neural circuitry that computes RPEs (see also Lak et al., 2014).

We hypothesize that this integration process involves outputs from the OFC, a reward area that is anatomically early in the reward hierarchy and that serves as an indirect input to the dopamine system (Takahashi et al., 2011). The OFC is important for signaling information about reward, reward learning, and regulation of reward-related cognition (Rushworth et al., 2011; Wallis, 2007; Wilson et al., 2014; Padoa-Schioppa, 2011). OFC may be involved in economic choice at least two ways. First, it could be a stage where all choice-relevant features are maintained in separate buffers, constituting a complete representation of task state. This would then be used as raw material from which downstream areas could compute an integrated value signal

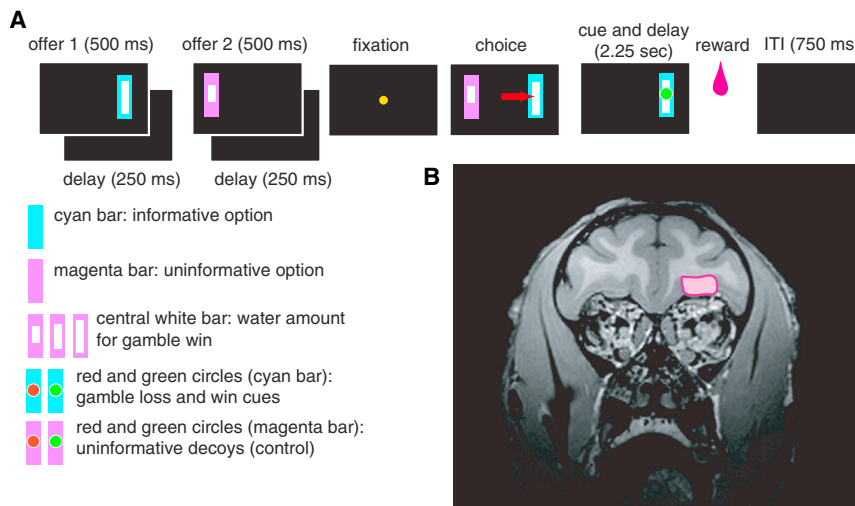


Figure 1. Task Design and Recording Location

(A) Basic task design. Two offers were presented in sequence, followed by a blank period. The monkey then had to fixate a central target. The two options then reappeared and the monkey chose one with a gaze shift. Then a cue appeared which was either informative (indicating whether the trial would be rewarded) or uninformative (leaving the monkey in a state of uncertainty). Following a 2.25-s delay, the monkey obtained the outcome. Cyan and magenta bars indicated informative and uninformative options, respectively. An inscribed white rectangle indicated gamble stakes. An inscribed red or green circle was the cue.

(B) MRI indicating position of 13 m (see [Figure S1](#) for a more detailed figure).

(Wilson et al., 2014). In this case, OFC neurons would code the presence of information or the presence of appetitive reward, but would not code their combined value or utility. Alternatively, the OFC could implement the next stage of evaluation where features are combined to create the value that guides decisions (Padoa-Schioppa, 2011; Padoa-Schioppa and Assad, 2006; Raghuraman and Padoa-Schioppa, 2014). In this case, activity in OFC would depend on both appetitive reward and information in a correlated manner, and precisely to the extent that the two variables influence decisions.

To test between these hypotheses, we recorded activity of OFC neurons using a novel “curiosity tradeoff task.” This task is a variant of the information-seeking task developed by Bromberg-Martin and Hikosaka (Bromberg-Martin and Hikosaka, 2009, 2011). On each trial of our task, a monkey chose between gambles that differed on two dimensions: (1) *water amount* associated with winning the gamble and (2) *informativeness*, i.e., whether a cue revealed the gamble outcome in advance of its delivery. Importantly, the information allowed monkeys to fully predict the chosen gamble’s outcome, but could not be used to influence the outcome in any way. Thus, any value the monkeys assigned to information was due to its intrinsic worth, rather than any objective benefit for gathering water reward.

We find that monkeys reliably choose to sacrifice water to obtain immediate information about the outcome of the gamble. Furthermore, our task allows us to measure the precise manner in which animals integrate water amount and informativeness into their judgments of subjective value. We could then test whether OFC neurons integrate these variables in the same manner as the animals do in their choice behavior (if the OFC represents subjective value) or whether OFC neurons encode these variables independently (if the OFC represents an abstract task state).

We find that OFC neurons encode both water amount and informativeness, consistent with a role for the OFC in curiosity-guided choices. Furthermore, much as OFC primary reward signals reflect the value subjects assign to those rewards (Critchley and Rolls, 1996; Padoa-Schioppa and Assad, 2006; Tremblay and Schultz, 1999), OFC information signals are correlated with

the value that monkeys assign to information. However, OFC neurons had no systematic tendency to integrate the values of water and information in an appropriate manner to code the overall subjective value that guides decisions. Instead, we find that OFC neurons coded these variables in orthogonal manners, consistent with a representation of abstract task state. For example, if a given neuron was positively tuned for water amount, it was no more likely than chance to be either positively or negatively tuned for informativeness. These results are consistent with the idea that OFC precedes the value computation that guides decisions, and suggest a role in motivating curiosity-guided choices.

RESULTS

Monkeys Value Advance Information about Gamble Outcomes

On each trial, monkeys chose between two gambles represented by visual stimuli on the left and right sides of a screen. Each gamble yielded either a water reward or no reward with equal probability. The water amount for each gamble was drawn randomly from the range 75–375 μl in 15- μl steps, and was indicated to the monkey by the height of a white inset bar (Figure 1). Gambles also varied in their *informativeness*, which was indicated to the monkey by their color. Choosing the informative gamble (Figure 1, cyan bar) always led to the presentation of a visual stimulus that cued the gamble’s outcome. Choosing the uninformative gamble (Figure 1, magenta bar) led to the presentation of a visual stimulus that provided no new information. Monkeys could not make use of the information to influence the outcome of the gamble; the information merely gave them 2.25 s of advance notice.

Monkeys preferred both greater water amounts and informative cues. They chose the option with the greater water amount on 81% of trials, and the option with greater informativeness on 67% of trials (both $p < 0.0001$, binomial test). Monkeys exhibited only small choice biases favoring offers based on non-reward features such as their location on the screen (choice of rightmost offer: subject B 54%, subject H 53%) or presentation

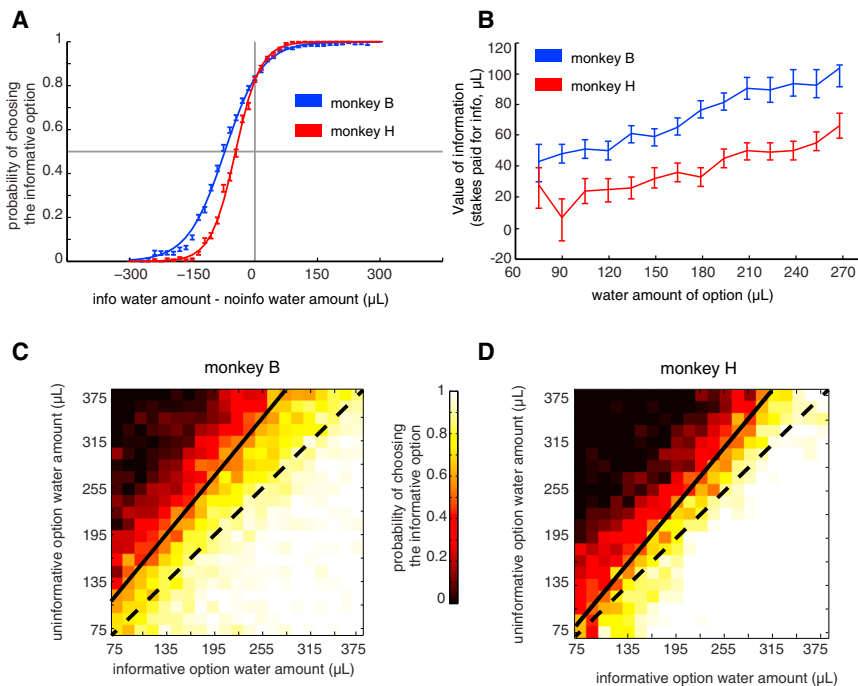


Figure 2. Monkeys Pay for Information about Future Reward

(A) Monkey preference for the informative option as a function of the water amount difference between the informative and uninformative options. Error bars indicate standard error.

(B) The subjective value of information (i.e., the amount of offered stakes the monkey paid to gain the information) as a function of offered water amount. Error bars indicate 95% confidence intervals.

(C and D) Heatmap, showing preference for the informative option as a function of the water amount of the informative and uninformative options. Black line indicates the indifference curve, indicating the indifference points of the animal (i.e., when preference for the two options is equivalent)

order (choice of first presented offer: subject B 51%, subject H 51%).

We calculated each monkey's probability of choosing the informative option as a function of the difference in water amount between the two options (Figure 2A). When the two options had equal water amounts, both monkeys strongly preferred information (subject B: 83% choice; subject H: 83% choice). Monkeys were only indifferent when the informative option offered a considerably smaller amount of water (subject B: on average 74- μL difference; subject H: 44- μL difference). This indifference point identifies the monkey's willingness to pay for information, and hence serves as a measure of the subjective value that the monkey assigns to that information (cf. [Deaner et al., 2005](#)). Once we account for the 50% probability of gambles, our data indicate that the monkeys would give up 37 μL (subject B) or 22 μL (subject H) of water to gain information. This translates to a substantial fraction of the water they were offered. The expected reward size averaged over all offers was 112.5 μL of water. Thus, monkeys paid an average of 33% of offered water (subject B) or 20% of offered water (subject H) in exchange for just a few seconds of advance information.

Do monkeys adjust their willingness to pay for information about a gamble's outcome based on the water amount at stake? To test this, we estimated the monkey's probability of choosing info as a function of the water amounts of *both* the informative and uninformative options (Figures 2C and 2D). We then plotted an indifference curve, tracing through all combinations of water amounts for the two options for which the monkeys chose the two options with equal probability (Figures 2C and 2D, black line). If monkeys assigned a fixed value to information regardless of the stakes, then the indifference curve would be a straight line with a slope of 1. Instead, the indifference curve had a slope steeper than 1 for both monkeys ($m = 1.24$ for subject B, 1.23

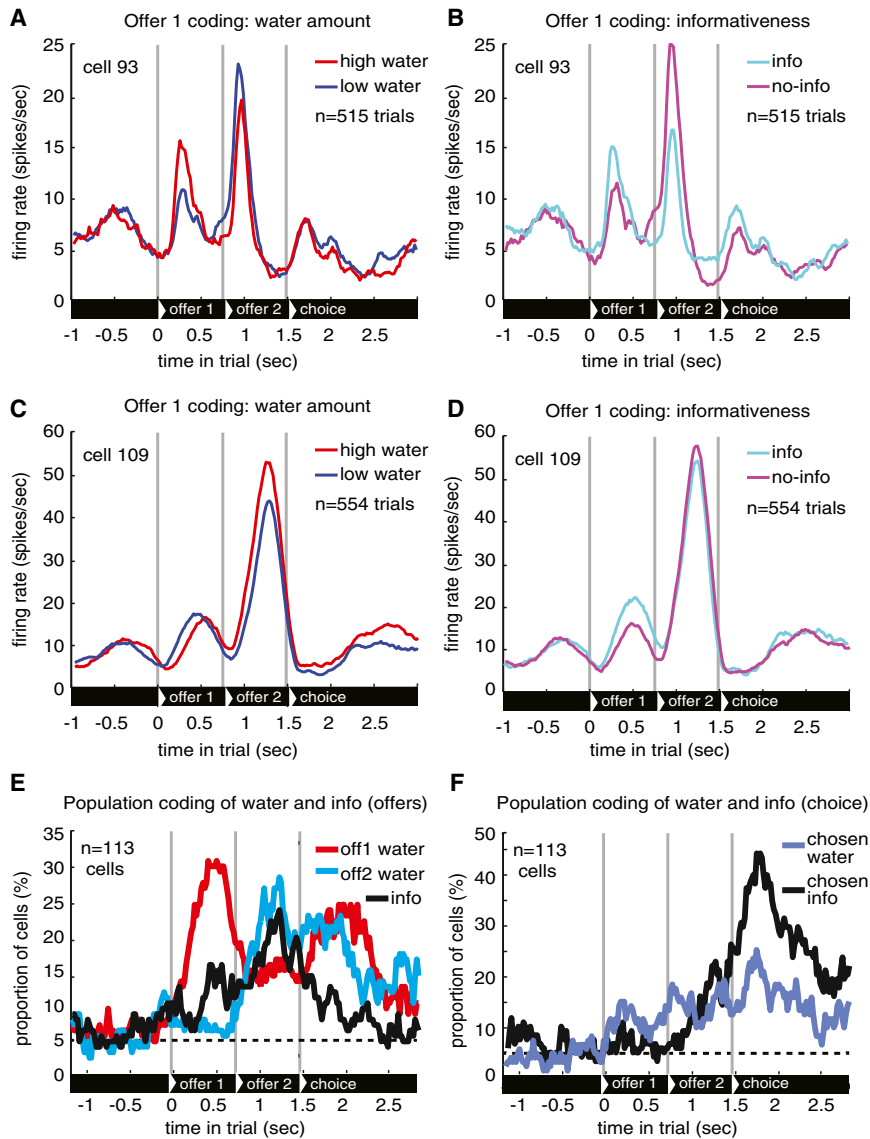
for subject H). The subjective value of information was an essentially linear function of the stakes in both monkeys (Figure 2B). These data suggest that the value of information may have a multiplicative effect on the value of water amount, just as probability does in a conventional gambling task, time does in a discounting task, or effort does in an effort task. Other analysis confirmed that the increasing value of information was not due to decreasing marginal utility of water (Figure S2). Thus, monkeys integrated both the availability of information and the amount of water at stake in order to arrive at their decisions.

OFC Neurons Code Offered Water Amount and Informativeness

We collected responses of 113 OFC neurons ($n = 72$ in subject B and $n = 41$ in subject H). We obtained an average of 522 trials per neuron (range: 396–818 trials). We first examined neural coding of water amount and informativeness. These two features of each offer were chosen independently, making it straightforward to separately measure their influences on neural activity. Furthermore, we presented the offers to the monkey one by one, which allowed us to separately measure neural coding of the two offers.

To compare tuning properties of neurons, we quantified each neuron's tuning with the regression coefficients from a linear regression of firing rate against offered water amount and informativeness. To compare between regression coefficients from neurons with different firing properties we first normalized (Z scored) neural firing rates and regressors (see [Supplemental Experimental Procedures](#)). However, we obtained the same qualitative results from performing regression on the raw data (data not shown). Neurons with significant regression coefficients for water amount or informativeness ($p < 0.05$ for this and all further statistical tests) were deemed to *code* these variables.

We observed significant coding of water amount for Offer 1 in 30% of neurons ($n = 34/113$; Figures 3A, 3C, and 3E) and informativeness in 15% of neurons ($n = 17/113$; Figures 3B, 3D, and 3E). The number of informativeness-coding neurons was greater than expected by chance ($p = 0.0005$, one-sided



binomial test against the 5% expected by chance). When the second offer was presented, we observed significant coding of water amount in 29% of neurons ($n = 33/113$) and informativeness in 25% of neurons ($n = 28/113$). Both proportions were greater than chance ($p < 0.001$). Latency analysis suggested that OFC water amount and information signals are present simultaneously in OFC (Figure S3).

We find a strong positive correlation between the signals our population of neurons used to encode the water amounts of the two offers ($r = +0.68$, $p < 0.001$; Figures 4A and 4C). In other words, a neuron that is excited by the presentation of a large water amount for Offer 2 will also tend to be excited by large water amounts for Offer 1. We find a similar positive correlation for informativeness signals between the two offers ($r = +0.33$, $p < 0.001$; Figures 4B and 4C). Thus, neural tuning to the individual aspects of the two offers was similar regardless of the order of presentation. Further analysis confirmed that neu-

Figure 3. OFC Neurons Signal Offered Water Amount and Informativeness

(A–D) PSTHs of two example neurons, showing (A and C) responses to first offers with different water amounts and (B and D) responses to first offers with different informativeness. (E) Percentage of cells showing significant correlation between firing rate and the water amounts of the two offers as well as the informativeness of the offers (note that there is only one “informativeness” variable here because if the first offer was informative the second offer was always non-informative, and vice versa). Dashed line indicates the percent of significant cells expected by chance. (F) Percentage of cells showing significant correlation between firing rate and the chosen offer’s water amount and informativeness.

rons consistently coded features of the currently presented offer, regardless of other variables. For instance, neurons used similar codes for the water amounts of the informative and non-informative offers, and did not have a predominant tendency to encode the water amount of the second offer relative to the previously presented first offer (Figure S4) consistent with previous studies (Rudebeck et al., 2013).

OFC Information Signals Grow with the Value of Information

OFC signals for primary reward are known to be sensitive to the subjective value of those rewards. We therefore asked whether the same was true for OFC information signals. Did OFC neurons simply encode a binary distinction, information versus no information? Or did OFC neurons signal the value that

the information has to the animal? Our data allow us to test between these hypotheses because animals assigned greater value to information when a greater amount of water was at stake. Thus, if OFC neurons signal the value of information, their information signals should grow with the offered water amount.

Indeed, OFC information signals were enhanced during high-stakes offers. The cell in Figure 5A, for instance, had activity negatively related to offer informativeness, and this negative informativeness signal was stronger on trials when the offered water amount was high (Figure 5A). To quantify this phenomenon, we examined the neuron’s regression coefficient for the term representing the interaction between informativeness and water amount (“Info \times Water,” Figure 5B). This cell had a significant negative coefficient for informativeness, indicating that it was inhibited by informative offers, and a significant negative coefficient for the Info \times Water interaction, indicating that its inhibition was stronger on trials when the offered water amount was high.

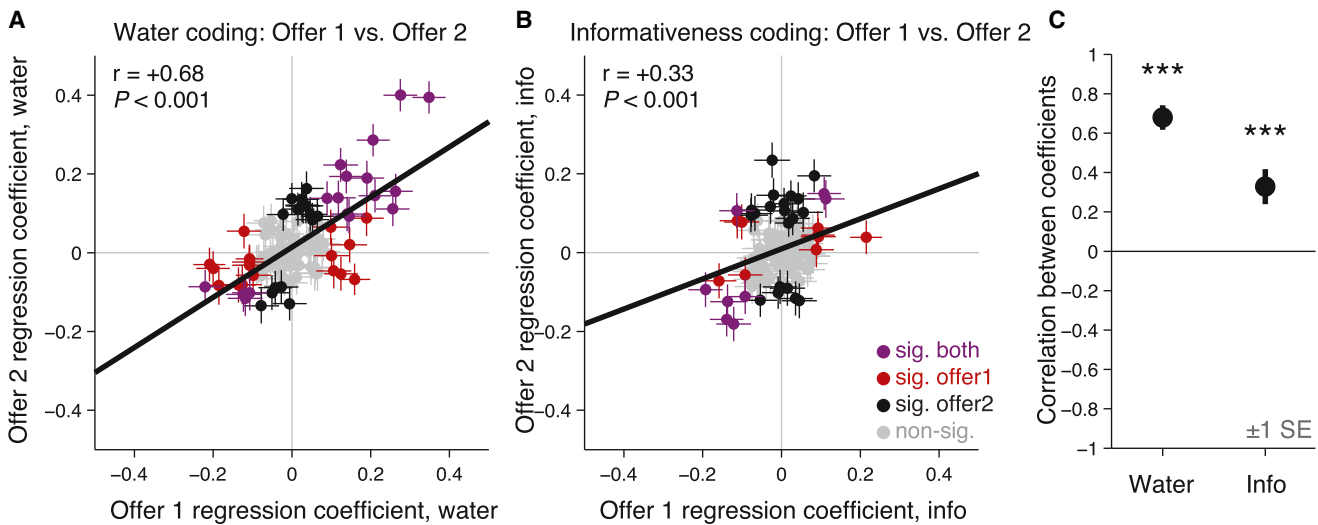


Figure 4. OFC Neurons Signal Both Offers in Similar Manners

(A and B) The correlation between each neuron's regression coefficients for (A) water amount and (B) informativeness for the two offers. Black lines indicate the line of best fit (linear regression). Red points are neurons that significantly encode the variable for Offer 1, black significantly encode the variable for Offer 2, purple significantly encode both, and gray fail to reach significance for either. Error bars indicate standard error of estimated regression coefficients.

(C) Correlation between the regression coefficients (± 1 SE).

The OFC population as a whole had a similar response pattern. Neurons were generally modulated by the Info \times Water interaction in consistent manners for both offers ($r = +0.31$, $p < 0.001$), similar to their consistent coding of main effects (Figure 4). Hence, for this analysis we pooled data by averaging each neuron's regression coefficients from the two offers. We then asked whether neural responses to information, measured as the main effect of informativeness, were consistently modulated by water amount, measured by the Info \times Water interaction. Indeed, these regression coefficients had a clear positive correlation ($r = +0.41$, $p < 0.001$; Figures 5C and 5E; similar results were found from analyzing individual offers, Offer 1: $r = +0.20$, $p = 0.032$; Offer 2: $r = +0.31$, $p = 0.001$). In other words, cells that were responsive to information were more responsive during high-stakes offers, when animals assigned the information greater value.

Our data also allow us to test between two hypotheses about the detailed mechanisms that generate information seeking. One hypothesis is that subjects value information because it allows them to physically or mentally prepare for reward delivery, thus increasing the amount of subjective value they can extract from the primary reward (Perkins, 1955). If this were the case then OFC water signals should be enhanced for informative offers, because informed water reward would have higher value than uninformed water reward. Alternately, the brain could assign a distinct value to information in its own right. If that were the case then OFC water signals should have no net enhancement by information, because the presence of information would have no effect on the value of water.

Our data support the latter view: there was no systematic tendency for signals coding water reward to be enhanced by the promise of information about those rewards. There was no significant correlation between neural modulation by the Info \times Water interaction and neural coding of water amount ($r =$

-0.13 , $p = 0.175$; Figures 5D and 5E). Furthermore, the neural interaction effect was significantly more correlated with information signals than water amount signals (difference of correlations = 0.54, bootstrap 99.5% CI excludes zero; Figure 5E). This suggests that information is assigned value in its own right, rather than merely enhancing the value of water reward, at least at the level of the OFC.

Orthogonal, Not Integrated, Coding of Offered Water Amount and Informativeness

OFC neurons appear to signal the value of information to the animal. However, monkeys prefer both water and information. This raises the question: do OFC neurons integrate the values of both water and information, and thus encode the overall subjective value of the option that guides decisions? If so, they should respond to water amount and informativeness with the same sign, and with strength proportional to their influence on choice. In contrast, if neurons code multiple task variables independently, then they ought to use unrelated signals to encode water amount and informativeness.

Some neurons coded both variables in similar manners—for instance, neurons that fired more for informative offers and fired more for large water amounts (Figures 3A and 3B). However, other neurons coded them in opposite manners—for instance, neurons that fired more for informative offers, but fired more for small water amounts (Figures 3C and 3D). If neurons tend to be tuned in the same way for the two variables then their regression coefficients should be positively correlated; if neurons code the variables independently, their correlation should be zero. We found that the regression coefficients for water amount and informativeness had no significant correlation, with an R value close to zero. The same result occurred consistently for neural responses to the first offer ($r = -0.08$, 95% CI

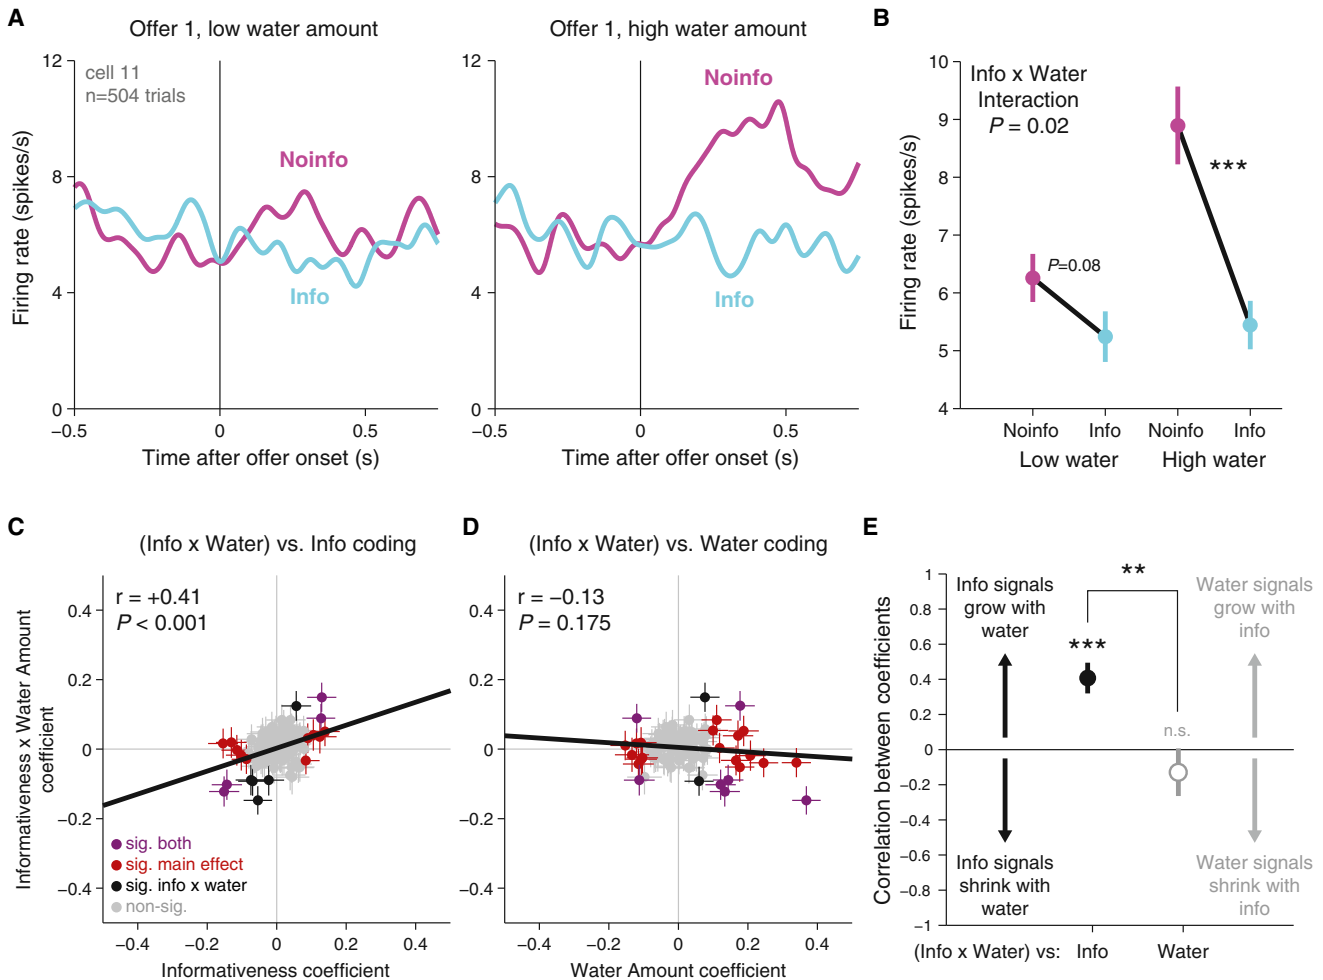


Figure 5. OFC Information Signals Grow with the Value of Information

(A and B) An example neuron that responded to Offer 1 with activity related to informativeness. The neural response (A) and mean firing rate (B) are plotted separately based on the offer's informativeness and water amount. In parallel with behavior, this neuron's information-related activity was strongest for high-stakes offers. Thus, this neuron had a negative main effect of Informativeness and a negative Informativeness \times Water Amount interaction.

(C) Neural modulations by the (Info \times Water) interaction (y axis) were strongly correlated with coding of Info (x axis). Each data point is a single neuron. Each neuron's coding of these variables was measured using the average of its regression coefficients from independent analyses of Offer1 and Offer2; analyses of each individual offer gave similar results. Same format as Figures 4A and 4B.

(D) Neural modulations by the (Info \times Water) interaction (y axis) were not significantly correlated with coding of Water Amount (x axis).

(E) Summary of results from (C) and (D). The positive correlation between Interaction coding and Info coding indicates that neural information signals were larger for offers with high water amounts (black dot, *** indicates $p < 0.001$). However, water amount signals had no significant tendency to be larger for informative or non-informative offers (gray dot). The former correlation was significantly greater than the latter (** indicates that the difference between correlations had a bootstrap 99% confidence interval that excluded zero). All error bars indicate standard error.

$[-0.26, +0.11]$, $p = 0.414$; Figures 6A and 6D) and for responses to the second offer ($r = -0.06$, 95% CI $[-0.24, +0.13]$, $p = 0.514$; Figures 6B and 6D). Furthermore, there was no significant correlation between the regression coefficients for water amount and for the interaction that modulated the value of information (Figure 5E).

It is possible we did not detect a correlation between water and information coding because we had too few trials to detect these signals. This seems unlikely, however, as we were able to detect strong and significant correlations between the same regression coefficients when comparing *within*-attribute, e.g.,

water coding of Offer 1 versus water coding of Offer 2 (Figure 4). Furthermore, we also detected clear correlations between regression coefficients when comparing *across*-attribute, e.g., informativeness coding versus interaction effects (Figure 5). As an additional test, we used a cross-validation procedure to test whether our analysis could reliably detect correlations between neural signals. We separated our data for each neuron into two halves, consisting of odd-numbered trials and even-numbered trials, and repeated the same regression procedure as above on each half of the data. Then, we compared the water regression coefficient calculated from the odd trials to the water

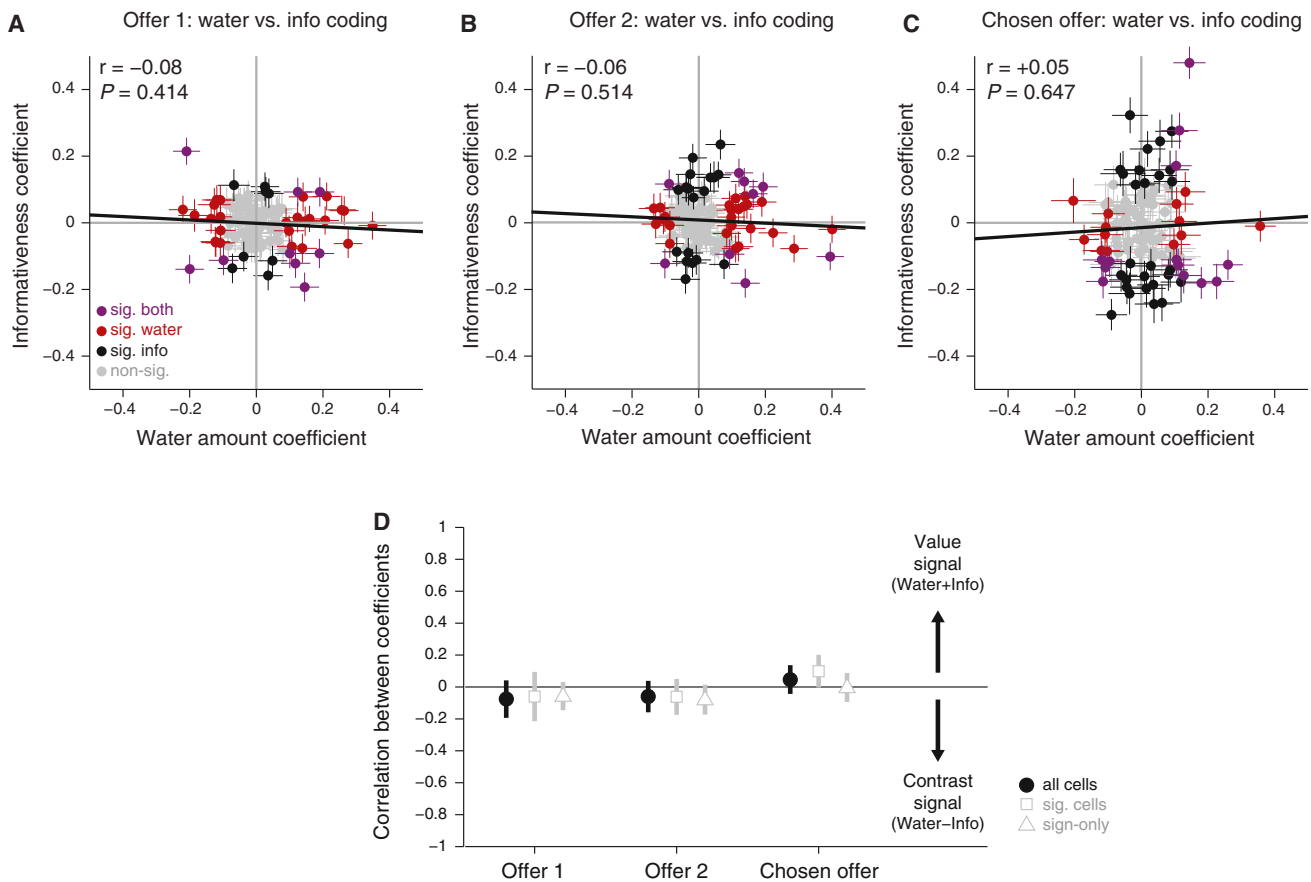


Figure 6. OFC Neurons Code Water Amount and Information in Uncorrelated Manners

(A–C) The correlation between each neuron’s regression coefficient for water amount and informativeness for (A) the first offer, (B) the second offer, and (C) the chosen offer. Panels follow a similar layout as Figures 4A and 4B. Error bars indicate standard error of estimated regression coefficients.

(D) Correlation between the regression coefficients (± 1 SE), calculated using all cells (black), cells with at least one significant effect (gray squares), or the signs of regression coefficients regardless of their magnitudes (gray triangles).

regression coefficient calculated from the even trials, and the informativeness coefficient calculated from the odd trials to the informativeness coefficient calculated from the even trials. If unreliable estimation of regression coefficients was the major contributor to our null effect, then the coefficients estimated from the two halves of the data would likewise show little or no correlation. Instead, the regression coefficients were strongly and significantly correlated. This was true for both water amount and informativeness coding, and for both the first and second offers (water amount: first offer $r = +0.72$, $p < 0.001$; second offer $r = 0.67$, $p < 0.001$. Informativeness: first offer $r = +0.31$, $p = 0.002$; second offer $r = +0.51$, $p < 0.001$). Thus, we were able to consistently detect neural signals, even when we calculated them using only half of our data set. It therefore seems unlikely that our finding of near-zero correlations was a consequence of insufficient data or some other cause of poor signal-to-noise.

Furthermore, the lack of detectable correlation between water and informativeness regression coefficients was not due to the presence of non-responsive neurons. The result persisted even if we restricted our analysis to neurons that had significant coding of at least one of the two variables (first offer: $n = 42/113$, $r =$

-0.06 , $p = 0.698$; second offer: $n = 52/113$, $r = -0.06$, $p = 0.655$; Figure 6D, squares). Nor does this result appear to be an artifact of our normalization procedure. Performing the analysis using non-normalized firing rates and regressors produced either no significant correlation (first offer: $r = -0.04$, $p = 0.671$) or, if anything, a weak tendency for negative correlation due to two outlier neurons with large firing rate modulations (second offer: $r = -0.25$, $p = 0.010$; reduced to $r = -0.02$, $p = 0.824$ after removal of outliers). Even a correlation of vectors containing only the sign of coding directions for each neuron (that is, either +1 or -1) produced no significant correlation (first offer: $r = -0.06$, $p = 0.431$; second offer: $r = -0.08$, $p = 0.311$; Figure 6D, triangles). In other words, cells that were excited or inhibited by water amount were equally likely to be excited or inhibited by informativeness. Thus, although a large fraction of OFC neurons encode water amount and informativeness of the offers, we conclude that they do so using orthogonal codes, rather than integrating them into a single scalar value signal.

It is important to note that our findings do not rule out the possibility that a subset of OFC neurons have value-like signals. After all, if the OFC uses an orthogonal code, there should be

subpopulations of neurons that signal water and information in all possible combinations. We were indeed able to find a small subpopulation of neurons with trends for *value-like* integration. These neurons coded information and water with the same sign and even appeared to assign higher value to information as the stakes increased (Figure S5B), as seen in behavior (Figure 2B). However, there was a similarly large subpopulation of neurons with exactly the opposite response pattern, *anti-value-like* integration (Figure S5C). These neurons coded the *absence* of information with the same sign as water, and signaled the *absence* of information more strongly as the stakes increased. Thus, OFC activity was significantly different from the pattern expected under the null hypothesis that cells predominantly coded subjective value (Figure S6). However, this finding is exactly what one would expect if OFC neurons carry all possible combinations of water and information signals, and by chance, some neurons happened to carry a combination that resembled the way monkeys computed subjective value during our task. Thus, OFC value-like signals in our task appear to be due to orthogonal coding of offer features, rather than due to value coding having a privileged status in OFC.

Orthogonal Coding during Choice Period

So far, our results indicate that OFC neurons encode features of valued offers in uncorrelated manners. It is possible that OFC neurons do predominantly signal the subjective value that guides choices, but only for the chosen offer at the time of choice. We therefore calculated water and informativeness coding indices, as we did above, but this time for the chosen offer, and in a time window encompassing the time just before and after the choice was made. To ensure that our analysis was not biased by the animals' tendency to choose offers with specific water and information parameters, we used a trial-matching procedure (see Supplemental Experimental Procedures). In essence, we performed our analysis on a subset of trials such that each trial where the informative option was chosen was paired with a trial where the non-informative option was chosen and had a similar water amount.

We found that a considerable fraction of neurons significantly coded the chosen offer's water amount (23% of neurons, $n = 26/113$) and informativeness (35% of neurons, $n = 40/113$; Figure 3F). Both of these proportions are much higher than expected by chance ($p < 0.001$, one-sided binomial tests). Furthermore, as in the offer epochs, neural information signals grew with the stakes of the chosen offer (correlation between regression coefficients for informativeness and Info \times Water interaction: $r = +0.21$, $p = 0.025$).

However, as we found in the offer epochs, OFC had no systematic tendency to integrate these features into a single value scale. There was no significant correlation between regression coefficients for water amount and informativeness ($r = +0.05$, 95% CI $[-0.14, +0.23]$, $p = 0.622$; Figures 6C and 6D). Once again, the same result held even if the analysis was performed on the subset of cells with significant coding of at least one feature ($n = 53/113$, $r = +0.10$, $p = 0.470$; Figure 6D, squares), was performed using non-normalized firing rates and regressors ($r = +0.07$, $p = 0.471$), was performed using only the sign of the regression coefficients ($r = 0.00$, $p = 0.957$; Figure 6D, triangles),

or was performed between coefficients for water amount and the Info \times Water interaction ($r = 0.05$, $p = 0.626$). And once again, our cross-validation analysis was able to reliably detect neural signals, to the extent that regression coefficients calculated from one half of the data were correlated with the same coefficients calculated from the other half of the data (water amount: $r = +0.39$, $p < 0.001$; informativeness: $r = +0.63$, $p < 0.001$). Thus, OFC neurons encode the features of the chosen offer, but have little systematic tendency to integrate them into a single value scale, even around the time of choice.

Neurons Respond Differently to Outcome-Related Cues and Outcomes Themselves

Although neurons generally did not integrate the water amount and informativeness of offers, it remained possible that they might respond consistently to water-related events throughout a trial. For instance, OFC neurons might respond similarly to water-predictive cues and to unpredicted water outcomes themselves, as dopamine and LHB neurons do (Bromberg-Martin and Hikosaka, 2009, 2011). We therefore analyzed neural responses to cues and outcomes. Following informative cue onset on informative trials, neurons quickly encoded whether the cue indicated a gamble win or loss, i.e., whether water would be delivered or omitted (Figure 7A, black line). Importantly, these signals were not simply encoding the red/green color of the visual cue, because neurons had little more than chance discrimination between the same red/green cues on uninformative trials when they were irrelevant to the task (Figure 7A, dashed gray line). Instead, on uninformative trials many neurons were strongly responsive when the gamble was resolved by water reward delivery or omission (Figure 7A).

We next asked whether neurons used similar signals to encode the offered water amount and the resolution of the gamble. There was indeed a significant positive correlation between water coding in response to offers and win/loss coding in response to informative cues ($r = +0.31$, $p = 0.001$; Figure 7D). Thus OFC neurons tended to have consistent water-amount tuning for stimuli in the same sensory modality (visual offer versus visual cue). However, this correlation was not absolute, and some neurons signaled offered and cued water in different directions (Figure 7D). For instance, the cell in Figure 7B was more activated by offers of high rather than low water amounts, but was more activated by cue and outcome feedback indicating that water would be omitted rather than delivered.

Furthermore, OFC neurons did not appear to use a consistent code to signal feedback about water conveyed through different sensory modalities (e.g., visual cue versus water outcome). Those water signals had weak negative correlation (offer versus outcome, $r = -0.20$, $p = 0.037$; Figure 7E) or no significant correlation (cue versus outcome, $r = +0.09$, $p = 0.331$, Figure 7F). The weak correlation between cue and outcome coding was especially striking. Both informative cues and uninformed outcomes conveyed very similar feedback to the animal: they were the first stimulus during the trial that told the animals whether they would receive a water reward. Yet in our task, OFC neurons were similarly likely to encode the cues and outcomes in the same direction (e.g., Figure 7B) or in opposite directions (e.g., Figure 7C).

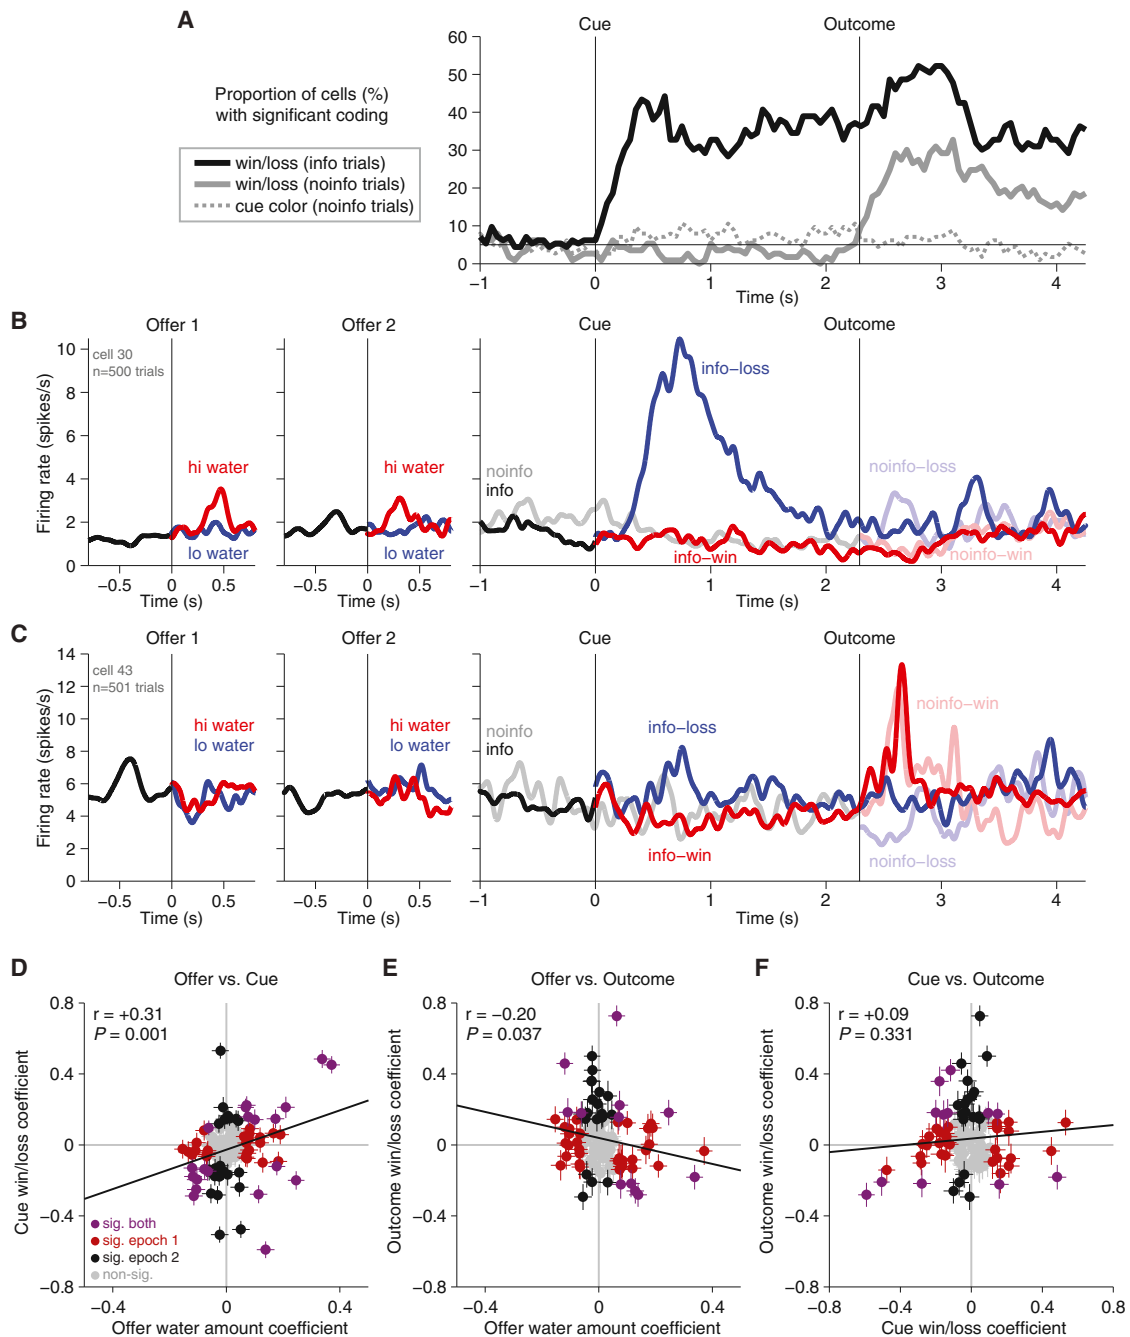


Figure 7. OFC Neurons Respond Differently to Outcome-Related Cues and Outcomes Themselves

(A) Percentage of cells showing significant correlation between firing rate and the win/loss outcome of the gamble on informative trials (black) and uninformative trials (gray solid line), as well as the cue color of the non-informative cues (gray dashed line). Horizontal black line indicates chance levels. Neurons respond to informative cues but not uninformative cues.

(B) Example neuron that was excited by high-water offers but was also excited by cues and outcomes indicating a gamble loss.

(C) Example neuron that was excited by cues indicating a gamble loss but outcomes indicating a gamble win.

(D–F) The correlation between each neuron's regression coefficients for (D) water amount and informative cue, (E) water amount and uninformed gamble outcome, and (F) informative cue and uninformed gamble outcome. Panels follow a similar layout as Figures 4A and 4B. Error bars indicate standard error of estimated regression coefficients.

This response pattern very different from cells such as LHB and DA neurons, which generally signal cue and outcome feedback in the same direction. To test this explicitly, we applied the same analysis to 95 LHB neurons previously recorded in a similar information seeking task (Bromberg-Martin and Hikosaka, 2011; DA neurons were also recorded in this task, but could not be fairly analyzed for this purpose because they were selected for recording on the basis of their cue and outcome responses; Bromberg-Martin and Hikosaka, 2009). Indeed, LHB neurons had a very strong correlation between win/loss coding in response to informative cues and uninformed outcomes ($r = +0.77$), and this was significantly greater than the correlation in OFC neurons (difference of correlations = 0.68, bootstrap 99.9% CI excludes zero).

DISCUSSION

We recorded responses of single neurons in area 13 of the OFC of two monkeys performing a curiosity tradeoff task. The prospect of immediate, rather than delayed resolution of the gamble increased its subjective value. We made use of this fact to study the representation of value in OFC neurons. We find that individual OFC neurons encode the two variables that influence value—water amount and informativeness of the gamble. However, they do not appear to integrate these variables, and instead use orthogonal codes. They also do not respond consistently to predictive cues and the receipt of the outcome. Thus, although these dimensions are integrated in DA neurons, they are largely uncorrelated in OFC. These results are consistent with the idea that OFC activity precedes and influences dopamine responses, and that OFC can be situated prior to the computations that instantiate reward-based decisions (McDannald et al., 2012; Noonan et al., 2010; Rushworth et al., 2011; Takahashi et al., 2011). Moreover, they support the idea that OFC represents task state rather than integrated value (Wilson et al., 2014).

Implications for OFC Function

Value representation in OFC is important for understanding the neural bases of economic choice. The *goods-based model* holds that the primary function of OFC neurons is to represent the values of offers and choices in a single value scale (Padoa-Schioppa, 2011; Padoa-Schioppa and Assad, 2006; Raghuraman and Padoa-Schioppa, 2014). Our results suggest that this integration does not extend to information, even when it is assigned subjective value. Other models of OFC function highlight its role as a structure that regulates learning, task-switching, executive control, and even metacognition (Kepecs et al., 2008; Ogawa et al., 2013; Roesch et al., 2006; Rushworth et al., 2011; Schoenbaum et al., 2011; Tsujimoto et al., 2009). Our results are broadly consistent with the predictions of the “task state” theory of OFC (Wilson et al., 2014). According to this theory, the function of OFC is to represent the current task state for use in guiding both choice and reinforcement learning. OFC is thus an input to and first stage of choice. Task state includes (but is not limited to) variables that influence value. Because task variables may influence choice and learning in different ways depending on the context, they may not be integrated into a single value variable. Our results support this prediction.

OFC neurons did encode some task variables in systematic ways. For instance, neurons carried similar water amount signals in response to the first and second offers, and for offers and cues. Previous work has shown that OFC responses to stimuli predicting a primary reward are proportional to the amount of value that reward has to the subject (Critchley and Rolls, 1996; Padoa-Schioppa and Assad, 2006; Roesch and Olson, 2005; Tremblay and Schultz, 1999) and are correlated with their evoked behavioral responses (Morrison and Salzman, 2009). We find a similar result here, suggesting that OFC information signals may reflect the value of information to the animal.

This suggests that OFC is at an early intermediate stage of computations, where abstract features of the task have begun to be combined into meaningful signals suitable to guide learning and decision-making, but have not yet been integrated into decision variables such as subjective value. We might call this the “aspects of value” hypothesis. In this view, OFC neurons might encode the amount of water associated with a cue (an abstract feature) or the subjective value of that water (an intermediate computation), but few neurons would integrate water with all other forms of reward to compute the subjective value of the option.

Our data also have implications for the role of the OFC in processing feedback about the outcomes of choice. In our task, the first feedback about whether the choice would yield a reward was conveyed by either informative cues or by outcome delivery. In contrast with DA neurons, OFC neurons did not respond to water-predicting cues the same way that they respond to the water outcome. We do find, however, that OFC neurons had related water-coding responses to visually presented offers and visual reward cues. Thus, OFC responses to reward and reward-related stimuli may depend on the sensory modality of the stimuli (in our task, visual versus tactile) rather than coding reward feedback per se.

OFC Reward Signals: Integrated versus Independent Coding

Multiple groups have reported that OFC neurons do not necessarily integrate multiple task variables into a single value signal (Schoenbaum et al., 2009; Wilson et al., 2014). For example, OFC neurons rarely integrate reward size, probability, and effort costs (Kennerley et al., 2009), reward size and risk (O’Neill and Schultz, 2010), and reward size and delay (Roesch et al., 2006). These studies required subjects to choose between options that varied along a single dimension at a time, and hence did not require subjects to integrate multiple attributes to make their decisions. Thus, it remained possible that the OFC neurons *would have* encoded integrated value if it had been required. One study addressed this issue using a choice task in which monkeys integrated social and liquid reward (Watson and Platt, 2012). They reported that largely separate populations of OFC neurons encode the receipt of social and liquid reward; however, they did not report whether these neurons integrated these reward at the time when the options were presented and the decision was being made, leaving open the possibility that OFC neurons do encode integrated value at the time of decision making. Our work addresses these limitations directly, by using a task in which monkeys traded off multiple attributes of reward, and by

examining OFC neural activity at the time when monkeys made their decisions.

Our results may appear to paint a different picture of OFC than careful work by Padoa-Schioppa and colleagues. In their experiments, a clear majority of OFC neurons encoded integrated values, and did so in a manner matching behavioral preferences (Padoa-Schioppa and Assad, 2006; Raghuraman and Padoa-Schioppa, 2014). However, careful examination shows that these findings are fully compatible and paint a nuanced picture of OFC function. A critical point is that the multiple attributes in their experiments were all related to a single event, an upcoming *liquid reward*, such as its taste, quantity, and probability (Padoa-Schioppa and Assad, 2006; Raghuraman and Padoa-Schioppa, 2014). In contrast, the two attributes in our experiment were related to different events: the stakes were related to the *liquid reward*, while informativeness was related to an upcoming *visual cue*. This meant that the two attributes in our task were linked to distinct future events that differed in their sensory modality (visual versus gustatory), timing (immediate presentation of the cue versus delayed delivery of water), and reason for being valued (curiosity versus thirst). Thus, even if OFC neurons are capable of integrating multiple attributes of a liquid reward, they may use different signals to code informational reward. Notably, OFC neurons in our task *did* integrate a pair of features that were both related to the same form of reward, the informational reward (Figure 5).

Implications for Curiosity-Guided Behavior

Our results show that desire for information (i.e., curiosity) can be captured and quantified in the laboratory. This makes it possible to use standard economic models to estimate the value of information to the subject as well as the combined value of offers that differ in water amount, risk, and informativeness. In particular, we found that the subjective value of information about future reward increased strongly and linearly with the stakes of the gamble.

What is not clear is exactly what causes information to have value in our task, as it does not lead to any benefit in terms of earning a greater amount of primary reward. Our data do place a constraint on models of information seeking by suggesting that, at least at the level of the OFC, information is assigned a true value of its own rather than merely modulating the value of primary reward (Figure 5E). One viable mechanism would be for the informative option to have greater salience because it is followed by cues with variable values (those that predict either reward or no reward), thus causing the informative option to be reinforced more strongly than the uninformative option (Esber and Haselgrove, 2011). This explanation fits with our neural results as well, as OFC neurons have previously been shown to respond to risk/uncertainty and salience (Kepecs et al., 2008; Ogawa et al., 2013; O'Neill and Schultz, 2010). Future work should explore the possible link between the value of information and salience.

The information coding signals we observed in OFC neurons are most directly comparable to those found in LHB and DA neurons (Bromberg-Martin and Hikosaka, 2009, 2011). However, LHB and DA neurons used a common code for water amount and informativeness: cells that responded to water-pre-

dictive cues responded in the same direction to information-predictive cues, consistent with a “common currency” representation of subjective value (similar integration has been found in DA neurons for other attributes of reward, such as juice type and risk; Lak et al., 2014). Thus, LHB and DA neurons appear to reflect the *output* of value computations.

Our results therefore suggest a potential circuit for curiosity-based decisions, in which informational and primary reward are represented independently in OFC and then combined into a single value scale in downstream areas. Integrated value does appear to be represented in vmPFC and in the dopamine system, as well as in areas even further downstream, like dACC and dlPFC. Previous work suggests an involvement of vmPFC in choice (Strait et al., 2014) and places dACC post-decisionally (Blanchard and Hayden, 2014; Cai and Padoa-Schioppa, 2012). Thus, the OFC may have an important role in curiosity-guided behavior, and in decision-making more generally, as a cortical area where task-relevant choice features can be highlighted and then sent to areas that perform value computations, decision making, and learning.

What is the precise role of the OFC in these value computations? Our neural data raise one intriguing possibility, that the “hunger for information” may be more than just a metaphor. We found that OFC information signals were greater when monkeys assigned higher value to information. Previous studies of our targeted region of OFC (13) found similar results for food reward. OFC responses to the sight, smell, and taste of food are greater with hunger (Critchley and Rolls, 1996; Pritchard et al., 2007; but see Bouret and Richmond, 2010). Furthermore, this region of OFC is critical for updating the value of food-associated objects when hunger gives way to satiety (Rudebeck and Murray, 2011; West et al., 2011). We therefore hypothesize that, just as the OFC regulates seeking of appetitive reward in response to internal states like hunger and thirst, the OFC may regulate information seeking in response to internal states like uncertainty and curiosity.

More generally, our results show that the chance to get information is not simply assigned a fixed value and immediately integrated into other value representations. Instead, its value must be constructed by a neural computation process that is sensitive to the statistics of predicted future reward. Our work provides a basis for future studies to delineate the circuits that perform these computations and generate curiosity-guided behavior.

EXPERIMENTAL PROCEDURES

All animal procedures were approved by the University Committee on Animal Resources at the University of Rochester and were designed and conducted in compliance with the Public Health Service's Guide for the Care and Use of Animals.

Information Tradeoff Task

Monkeys performed a two-option gambling task (Figure 1). Two offers were presented in sequence on each trial. The first offer appeared for 500 ms, followed by a 250 ms blank period; a second option appeared for 500 ms followed by a 250 ms blank period. Every trial had one informative and one uninformative option. The order of presentation (informative versus uninformative) and location of presentation (info-on-left versus info-on-right) varied randomly by trial. The offered water amount varied randomly for each option (75 to 375 μ l water in 15 μ l increments).

Each offer was represented by a rectangle 300 pixels tall and 80 pixels wide (11.35 degrees tall and 4.08 degrees wide). All options offered a 50% probability of gamble win, to be delivered 2.25 s after the choice. Gamble offers were defined by two parameters, informativeness and water amount. Informative gambles (cyan rectangle) indicated that the subject would see a 100% valid cue immediately after choice indicating whether the gamble was won or lost (although receipt always occurred 2.25 s after choice). Uninformative gambles (magenta rectangle) indicated that a random cue would appear immediately after choice, and thus the animal had to wait the full 2.25-s delay to discover whether the gamble was won or lost. Valid and invalid cues were physically identical (green and red circles inscribed on the chosen rectangle). Each offer contained an inner white rectangle. The height of this rectangle linearly scaled with the water amount to be gained in the case of a gamble win.

Offers were separated from the fixation point by 550 pixels (27.53 degrees). Monkeys were free to fixate upon the offers when they appeared (and in our observations almost always did so). After the offers were presented, a central fixation spot appeared. Following 100 ms fixation, both offers reappeared simultaneously and the animal chose one by shifting gaze to it for 200 ms. Failure to maintain gaze returned the monkey to a choice state; thus monkeys were free to change their mind within 200 ms (although they seldom did so). Then the 2.25-s delay began, and the cue was immediately displayed. After the delay, if the gamble was won, a reward was delivered. If it was lost, no reward was delivered. All trials were followed by a 750 ms inter-trial interval (ITI) with a blank screen.

Statistical Methods

To calculate the subjective value of information for each water amount (Figure 2B), we first determined the subjective value of informative and uninformative options for each possible reward amount. We fit a separate logistic regressions for each water amount w . This regression model regressed choice of the informative option (1 or 0) against the water amount offered by the uninformative option, using only trials where the informative option offered w . We then calculated the subjective value of the informative option, in terms of μl of water offered by the uninformative option, using the point of subjective equality (where the logistic regression curve crossed $y = 0.5$). We only included points from 75–270 μl , because above this range the animal's preference for information was near ceiling, which prevents accurate estimation of the value of information. We calculated the indifference point for the highest and lowest values of the 95% confidence interval for our logistic regression estimate (error bars in Figure 2B). To calculate indifference lines in the heat maps (Figures 2C and 2D), we used the same calculation. Because subjective values are in terms of a water amount for an uninformative option, the subjective value of each informative option corresponds to a unique point on our heatmap. We fit a linear function through these points to create the curve.

PSTHs of neural activity were constructed by aligning spike rasters to task events and averaging rates across trials. Single-unit PSTHs were smoothed with a 200 ms time bin (Figures 3A–3D) or a Gaussian filter with $\text{SD} = 30$ ms (Figures 6B and 6C). The analysis of the percent of cells with significant signals, in Figures 3E, 3F, and 6A, were performed using a running 500 ms boxcar. The time windows for the scatterplots were as follows: offers 1 and 2, 480 ms windows starting 260 ms after offer onset; chosen offer, a 500 ms window centered at the time of choice; cues and outcomes, 800 ms windows starting 200 ms after event onset.

Some statistical tests of neuron activity were only appropriate when applied to single neurons because of variations in response properties across the population. In such cases, a binomial test was used to determine whether a significant portion of single neurons reached significance on their own, thereby allowing conclusions about the neural population as a whole.

Neural coding was quantified using the fitted coefficients from a linear regression model in which a neuron's single-trial firing rates were modeled as a constant factor plus a weighted linear combination of multiple variables. The main analysis used the offer's water amount (in μl) and informativeness (0 if non-informative, 1 if informative). Analyses involving interaction effects used a model with an additional term representing the interaction between water amount and informativeness ((water amount – mean water amount) \times (informativeness – mean informativeness)). Unless otherwise stated, both the firing rates and regressors for each neuron were z-scored (i.e., they were shifted and

scaled to have mean = 0 and standard deviation = 1), to allow comparison between cells with different firing rates, and comparison of the effects of regressors with different units (μl for water amount versus a binary variable for informativeness). The regression coefficients, their standard errors, and their p values were calculated using the MATLAB function "glmfit." For some analyses, data was pooled from Offer 1 and Offer 2 by averaging their regression coefficients (Figure 5). The standard errors of correlations between the regression coefficients (Figures 4C, 5E, and 6D) were calculated using bootstrapping, as the standard deviation of the correlations calculated from 200 bootstrap data sets in which the neurons were resampled with replacement.

Analysis of previously recorded LHb neurons (Bromberg-Martin and Hikosaka, 2011) was done using the same procedure, using the analysis time windows from the previously published paper (cue response: 100–350 ms after cue onset; outcome response: 200–450 ms after outcome onset). That task was similar to the present task except that options varied only in informativeness, not water amount. Trials were gambles for reward that were equally likely to end in a win (big reward, 880 μl water) or a loss (small reward, 40 μl water), and these outcomes were cued by either informative or non-informative visual cues. For details, see (Bromberg-Martin and Hikosaka, 2011).

SUPPLEMENTAL INFORMATION

Supplemental Information includes six figures and Supplemental Experimental Procedures and can be found with this article online at <http://dx.doi.org/10.1016/j.neuron.2014.12.050>.

AUTHOR CONTRIBUTIONS

T.C.B. and B.Y.H. designed research. T.C.B. performed research. T.C.B., E.S.B.-M., and B.Y.H. analyzed data and wrote the paper.

ACKNOWLEDGMENTS

We thank Caleb Strait and Alex Thomé for useful discussions and Marc Mancarella and Meghan Castagno for assistance in data collection. This research was supported by NIH R01(DA038106) to B.Y.H.

Received: October 18, 2014

Revised: November 10, 2014

Accepted: December 11, 2014

Published: January 22, 2015

REFERENCES

- Blanchard, T.C., and Hayden, B.Y. (2014). Neurons in dorsal anterior cingulate cortex signal postdecisional variables in a foraging task. *J. Neurosci.* 34, 646–655.
- 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.
- Bromberg-Martin, E.S., and Hikosaka, O. (2009). Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron* 63, 119–126.
- Bromberg-Martin, E.S., and Hikosaka, O. (2011). Lateral habenula neurons signal errors in the prediction of reward information. *Nat. Neurosci.* 14, 1209–1216.
- Cai, X., and Padoa-Schioppa, C. (2012). Neuronal encoding of subjective value in dorsal and ventral anterior cingulate cortex. *J. Neurosci.* 32, 3791–3808.
- Critchley, H.D., and Rolls, E.T. (1996). Hunger and satiety modify the responses of olfactory and visual neurons in the primate orbitofrontal cortex. *J. Neurophysiol.* 75, 1673–1686.
- Deaner, R.O., Khara, A.V., and Platt, M.L. (2005). Monkeys pay per view: adaptive valuation of social images by rhesus macaques. *Curr. Biol.* 15, 543–548.
- Esber, G.R., and Haselgrove, M. (2011). Reconciling the influence of predictiveness and uncertainty on stimulus salience: a model of attention in associative learning. *Proc. Biol. Sci.* 278, 2553–2561.

- Gottlieb, J., Oudeyer, P.-Y., Lopes, M., and Baranes, A. (2013). Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends Cogn. Sci.* *17*, 585–593.
- Gruber, M.J., Gelman, B.D., and Ranganath, C. (2014). States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron* *84*, 486–496.
- Kang, M.J., Hsu, M., Krajbich, I.M., Loewenstein, G., McClure, S.M., Wang, J.T., and Camerer, C.F. (2009). The wick in the candle of learning: epistemic curiosity activates reward circuitry and enhances memory. *Psychol. Sci.* *20*, 963–973.
- Kennerley, S.W., Dahmubed, A.F., Lara, A.H., and Wallis, J.D. (2009). Neurons in the frontal lobe encode the value of multiple decision variables. *J. Cogn. Neurosci.* *21*, 1162–1178.
- Kepecs, A., Uchida, N., Zariwala, H.A., and Mainen, Z.F. (2008). Neural correlates, computation and behavioural impact of decision confidence. *Nature* *455*, 227–231.
- Lak, A., Stauffer, W.R., and Schultz, W. (2014). Dopamine prediction error responses integrate subjective value from different reward dimensions. *Proc. Natl. Acad. Sci. USA* *111*, 2343–2348.
- Levy, D.J., and Glimcher, P.W. (2012). The root of all value: a neural common currency for choice. *Curr. Opin. Neurobiol.* *22*, 1027–1038.
- Loewenstein, G. (1994). The psychology of curiosity: a review and reinterpretation. *Psychol. Bull.* *116*, 75–98.
- McDannald, M.A., Takahashi, Y.K., Lopatina, N., Pietras, B.W., Jones, J.L., and Schoenbaum, G. (2012). Model-based learning and the contribution of the orbitofrontal cortex to the model-free world. *Eur. J. Neurosci.* *35*, 991–996.
- Montague, P.R., and Berns, G.S. (2002). Neural economics and the biological substrates of valuation. *Neuron* *36*, 265–284.
- 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.
- Noonan, M.P., Walton, M.E., Behrens, T.E.J., Sallet, J., Buckley, M.J., and Rushworth, M.F.S. (2010). Separate value comparison and learning mechanisms in macaque medial and lateral orbitofrontal cortex. *Proc. Natl. Acad. Sci. USA* *107*, 20547–20552.
- O'Neill, M., and Schultz, W. (2010). Coding of reward risk by orbitofrontal neurons is mostly distinct from coding of reward value. *Neuron* *68*, 789–800.
- Ogawa, M., van der Meer, M.A.A., Esber, G.R., Cerri, D.H., Stalnaker, T.A., and Schoenbaum, G. (2013). Risk-responsive orbitofrontal neurons track acquired salience. *Neuron* *77*, 251–258.
- Padoa-Schioppa, C. (2011). Neurobiology of economic choice: a good-based model. *Annu. Rev. Neurosci.* *34*, 333–359.
- Padoa-Schioppa, C., and Assad, J.A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature* *441*, 223–226.
- Perkins, C.C., Jr. (1955). The stimulus conditions which follow learned responses. *Psychol. Rev.* *62*, 341–348.
- Phillips, K.A., Subiaul, F., and Sherwood, C.C. (2012). Curious monkeys have increased gray matter density in the precuneus. *Neurosci. Lett.* *518*, 172–175.
- Pritchard, T.C., Schwartz, G.J., and Scott, T.R. (2007). Taste in the medial orbitofrontal cortex of the macaque. *Ann. N Y Acad. Sci.* *1121*, 121–135.
- Raghuraman, A.P., and Padoa-Schioppa, C. (2014). Integration of multiple determinants in the neuronal computation of economic values. *J. Neurosci.* *34*, 11583–11603.
- Roesch, M.R., and Olson, C.R. (2005). Neuronal activity in primate orbitofrontal cortex reflects the value of time. *J. Neurophysiol.* *94*, 2457–2471.
- 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.
- Rudebeck, P.H., and Murray, E.A. (2011). Dissociable effects of subtotal lesions within the macaque orbital prefrontal cortex on reward-guided behavior. *J. Neurosci.* *31*, 10569–10578.
- 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.
- Rushworth, M.F.S., Noonan, M.P., Boorman, E.D., Walton, M.E., and Behrens, T.E. (2011). Frontal cortex and reward-guided learning and decision-making. *Neuron* *70*, 1054–1069.
- Schoenbaum, G., Roesch, M.R., Stalnaker, T.A., and Takahashi, Y.K. (2009). A new perspective on the role of the orbitofrontal cortex in adaptive behaviour. *Nat. Rev. Neurosci.* *10*, 885–892.
- Schoenbaum, G., Takahashi, Y., Liu, T.-L., and McDannald, M.A. (2011). Does the orbitofrontal cortex signal value? *Ann. N Y Acad. Sci.* *1239*, 87–99.
- 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.
- Takahashi, Y.K., Roesch, M.R., Wilson, R.C., Toreson, K., O'Donnell, P., Niv, Y., and Schoenbaum, G. (2011). Expectancy-related changes in firing of dopamine neurons depend on orbitofrontal cortex. *Nat. Neurosci.* *14*, 1590–1597.
- 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.
- Wallis, J.D. (2007). Orbitofrontal cortex and its contribution to decision-making. *Annu. Rev. Neurosci.* *30*, 31–56.
- Watson, K.K., and Platt, M.L. (2012). Social signals in primate orbitofrontal cortex. *Curr. Biol.* *22*, 2268–2273.
- West, E.A., DesJardin, J.T., Gale, K., and Malkova, L. (2011). Transient inactivation of orbitofrontal cortex blocks reinforcer devaluation in macaques. *J. Neurosci.* *31*, 15128–15135.
- 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.