Neuron
Article

Reward Value Comparison via Mutual Inhibition in Ventromedial Prefrontal Cortex

Caleb E. Strait,[1,*] Tommy C. Blanchard,[1] and Benjamin Y. Hayden[1]
[1]Department of Brain and Cognitive Sciences, Center for Visual Science, University of Rochester, Rochester, NY 14627, USA
*Correspondence: cstrait@bcs.rochester.edu
http://dx.doi.org/10.1016/j.neuron.2014.04.032

SUMMARY

Recent theories suggest that reward-based choice reflects competition between value signals in the ventromedial prefrontal cortex (vmPFC). We tested this idea by recording vmPFC neurons while macaques performed a gambling task with asynchronous offer presentation. We found that neuronal activity shows four patterns consistent with selection via mutual inhibition: (1) correlated tuning for probability and reward size, suggesting that vmPFC carries an integrated value signal; (2) anti-correlated tuning curves for the two options, suggesting mutual inhibition; (3) neurons rapidly come to signal the value of the chosen offer, suggesting the circuit serves to produce a choice; and (4) after regressing out the effects of option values, firing rates still could predict choice—a choice probability signal. In addition, neurons signaled gamble outcomes, suggesting that vmPFC contributes to both monitoring and choice processes. These data suggest a possible mechanism for reward-based choice and endorse the centrality of vmPFC in that process.

INTRODUCTION

In reward-based (i.e., economic) choice, decision makers select options based on the values of the outcomes they yield (Padoa-Schioppa, 2011; Rangel et al., 2008). Elucidating the mechanisms of reward-based choice is a fundamental problem in economics, psychology, cognitive science, and evolutionary biology (Glimcher, 2003; Rangel et al., 2008; Rushworth et al., 2011). Recent scholarship suggests that reward value comparisons can be efficiently implemented by mutual inhibition between representations of the values of the options (Hunt et al., 2012, 2013; Jocham et al., 2012). This mutual inhibition hypothesis is analogous to one closely associated with memory-guided perceptual comparisons (Hussar and Pasternak, 2012; Machens et al., 2005; Romo et al., 2002; Wang, 2008). This theory is also supported by neuroimaging results consistent with its general predictions (Basten et al., 2010; Boorman et al., 2009; FitzGerald et al., 2009). However, support is greatly limited by the lack of single-unit evidence for what is ultimately a neuronal hypothesis.

We chose to record in area 14 of the ventromedial prefrontal cortex (vmPFC), a central region of the monkey ventromedial reward network that is analogous to human vmPFC (Ongür and Price, 2000). We chose vmPFC for five reasons. First, a large number of neuroimaging and lesion studies have identified the vmPFC as the most likely locus for reward value comparison (Levy and Glimcher, 2012; Rangel and Clithero, 2012; Rushworth et al., 2011). Second, lesions to vmPFC are associated with deficits in choices between similarly valued items, possibly leading to inconsistent choices and shifts in choice strategy (Camille et al., 2011; Fellows, 2006; Noonan et al., 2010; Walton et al., 2010). Third, activity in this area correlates with the difference between offered values, suggesting that it may implement a value comparison process (Boorman et al., 2013; FitzGerald et al., 2009; Philiastides et al., 2010). Some recent neuroimaging specifically suggests that vmPFC is the site of a competitive inhibition process that implements reward-based choice. Blood oxygen levels in vmPFC track the relative value between the chosen option and the next-best alternative (Boorman et al., 2009, 2013). Fourth, the vmPFC BOLD signal shifts from signaling value to signaling value difference in a manner consistent with competitive inhibition (Hunt et al., 2012). Fifth, relative GABAergic and glutamatergic concentrations—chemical signatures of inhibition/excitation balance—in vmPFC are correlated with choice accuracy (Jocham et al., 2012).

Some previous studies have identified correlates of choice processes in a closely related (and adjacent) structure, the lateral orbitofrontal cortex (lOFC) (Padoa-Schioppa, 2009, 2013; Padoa-Schioppa and Assad, 2006). A key prediction of choice models is that representations of value in lOFC are stored in a common currency format and compared locally within lOFC (Padoa-Schioppa, 2011). We chose to record in the vmPFC rather than the lOFC because some evidence suggests the function of lOFC may be more aptly characterized as credit assignment, salience, reward history, or flexible control of choice (Feierstein et al., 2006; Hosokawa et al., 2013; Kennerley et al., 2011; Noonan et al., 2010; O'Neill and Schultz, 2010; Ogawa et al., 2013; Roesch et al., 2006; Schoenbaum et al., 2009; Walton et al., 2010; Watson and Platt, 2012; Wilson et al., 2014).

We used a modified version of a two-option risky choice task we have used in the past (Hayden et al., 2010, 2011a). To temporally dissociate offered value signals from comparison and selection signals, we presented each of the two offers asynchronously before allowing overt choice. We found that four patterns that are consistent with the idea that vmPFC contributes to choice through mutual inhibition of value representations: (1) in response to the presentation of the first offer, neurons carried

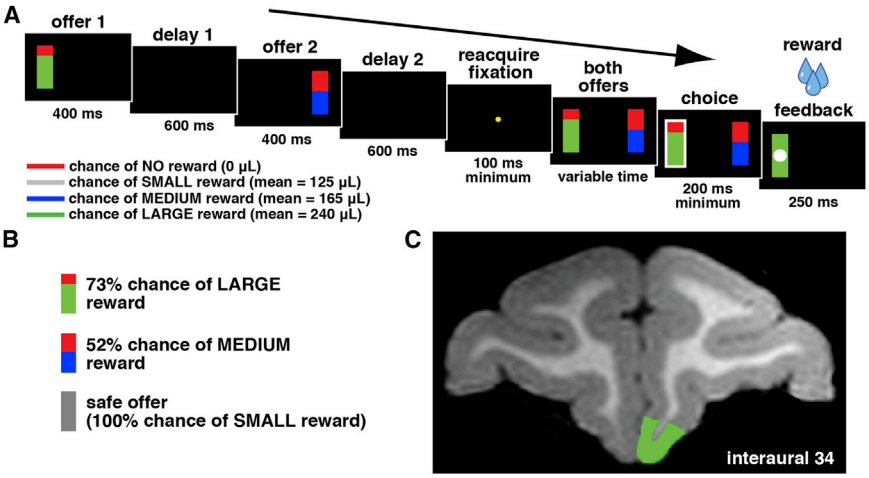

Figure 1. Task and Recording Location

(A) Timeline of gambling task. Two options were presented, each offering a gamble for water reward. Each gamble was represented by a rectangle, some proportion of which was gray, blue, or green, signifying a small, medium, or large reward, respectively. The size of this colored region indicated the probability that choosing that offer would yield the corresponding reward. Offers appeared in sequence, offset by 1 s and in a random order for 400 ms each. Then, after fixation, both offers reappeared during a decision phase. Outcomes that yielded reward were accompanied by a visual cue: a white circle in the center of the chosen offer.

(B) Example offers. Probabilities for blue and green offers were drawn from a uniform distribution between 0% and 100% by 1% increments. Gray (safe) offers were always associated with a 100% chance for reward.

(C) Magnetic resonance image of monkey B. Recordings were made in area 14 of vmPFC (highlighted in green).

a signal that correlated with both its reward probability and reward size; these signals were positively correlated. This suggests that vmPFC neurons carry integrated value representations. (2) After presentation of the second offer, but before choice, neural responses were correlated with values of both options, but with anti-correlated tuning for the two options, suggesting the two values serve to mutually inhibit neuronal responding. (3) Neurons rapidly came to signal the value of the chosen offer but not the unchosen one, suggesting that the processes we are observing generate a choice. (4) After accounting for option values, variability in firing rates after presentation of the offers predicted choices. This fourth finding is analogous to the idea of choice probability in perceptual decision making and provides a strong link between neural activity in vmPFC and control of choices (Britten et al., 1996; Nienborg and Cumming, 2009). Collectively, these patterns are consistent with the idea that vmPFC stores values and compares them through a mutual inhibition process (Hunt et al., 2012; Jocham et al., 2012; Machens et al., 2005; Wang, 2008).

We made an additional observation that fleshes out our understanding of the mechanisms of reward value comparison in vmPFC. We found that vmPFC neurons tracked gamble outcomes; these monitoring signals were even stronger than choice-related signals. Unlike similar signals observed in the posterior cingulate cortex (PCC) and dorsal anterior cortex (dACC), these responses did not predict strategic adjustments (Hayden et al., 2008, 2011a). We infer that monitoring functions of vmPFC are subject to downstream gating before influencing behavior (cf. Blanchard and Hayden, 2014).

RESULTS

Preferences Patterns for Risky Gambles

Two monkeys performed a two-option gambling task (see Experimental Procedures; Figures 1A and 1B). Options differed on two dimensions: probability (0%–100% by 0.1% increments) and reward size (either medium, 165 µl, or large, 240 µl) (see Experimental Procedures). On 12.5% of trials, one option was a small safe choice (100% chance of 125 µl). Subjects chose the offer with the higher expected value 85% of the time, suggesting that they generally understood the task and sought to maximize reward (n = 70,350 trials for all preference pattern analyses).

Both monkeys were risk seeking, meaning that they preferred risky to safe offers with matched expected values (Figure 2A). We quantified risk preferences by computing points of subjective equivalence (PSE) between safe offers and gambles (Hayden et al., 2007). The PSE for large reward (green) gambles (0.39 of the value of the safe offer) was lower than for medium (blue) gambles (0.52). This difference, and also the fact that both large- and medium-reward PSEs were lower than 1, indicates strong risk-seeking tendencies (cf. McCoy and Platt, 2005). This risk-seeking pattern is consistent with what we and others have observed in rhesus monkeys (Hayden et al., 2011a; Heilbronner and Hayden, 2013; Monosov and Hikosaka, 2013; O'Neill and Schultz, 2010; Seo and Lee, 2009; So and Stuphorn, 2012) and is inconsistent with one recent study showing risk aversion in rhesus monkeys (Yamada et al., 2013).

To delineate the factors that influence monkeys' choices, we implemented a logistic general linear model with choice (offer 1 versus offer 2) as a function of seven regressors: both reward sizes, both reward probabilities, outcome of previous trial (reward versus no reward), choice of previous trial (offer 1 versus offer 2), and side of offer 1 (left versus right). Choice was significantly affected by both reward sizes (offer 1: $t = 115.89$; offer 2: $t = -114.77$; $p < 0.0001$ in both cases) and both probabilities (offer 1 probability: $t = 107.31$; offer 2 probability: $t = -109.65$; $p < 0.0001$ in both cases) (Figure 2B). Choice was not affected by outcome of previous trial ($t = 0.73$, $p = 0.47$), by chosen offer order on previous trial ($t = 1.37$, $p = 0.17$), or by side of offer 1 ($t = 1.60$, $p = 0.11$). Moreover, previous outcomes did not affect choice coded by side (left offer versus right offer; $X^2 = 1.17$, $p = 0.28$), same order offer as previous trial ($X^2 = 1.03$, $p = 0.31$), same side offer as previous trial ($X^2 = 0.91$, $p = 0.34$), or previous offer expected value (high versus low; $X^2 = 1.70$, $p = 0.19$). The lack of an observed trial-to-trial dependence is

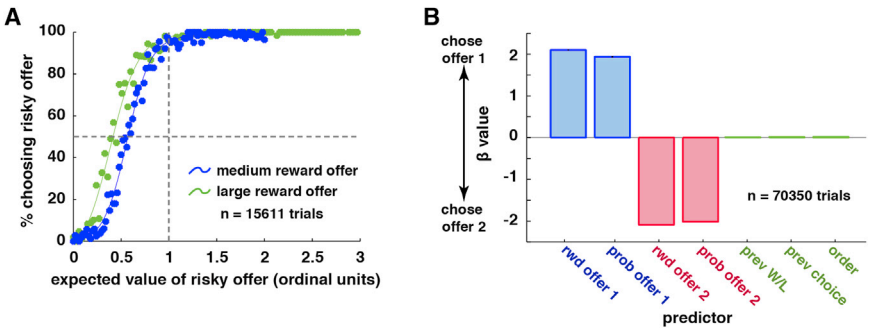

Figure 2. Behavioral Results

(A) Likelihood of choosing risky offer instead of a safe one as a function of risky offer expected value. Data are separated for high-value (green) and medium-value (blue) gambles. Fits are made with a Lowess smoothing function. Expected values are calculated in units of ordinal expected value (see Experimental Procedures).

(B) Effects of seven trial variables on choice (offer 1 versus 2) using a logistic GLM. Tested variables are as follows: (1) the reward and (2) probability for offer 1, the (3) reward and (4) probability for offer 2, (5) the outcome of the most recent trial (win or choose safe = 1, loss = 0), (6) the previous choice (first = 1, second = 0), and (7) the order of presentation of offers (left first = 1, right first = 0). Error bars in all cases are smaller than the border of the bar and are therefore not shown.

inconsistent with an earlier study using a similar task where we observed a weak trial-to-trial dependence (Hayden et al., 2011a). We suspect the difference in preferences is due to the small changes in task design between the earlier studies and the present one.

Single Unit Responses

We recorded the activity of 156 vmPFC neurons while monkeys performed our gambling task (106 neurons in monkey B; 50 neurons in monkey H). To maximize our sensitivity to potentially weak neuronal signals, we deliberately recorded large numbers of trials for each cell (mean of 1,036 trials per neuron; minimum of 500 trials). Neurons were localized to area 14 (for precise demarcation, see Figure S1 available online). For purposes of analysis, we defined three task epochs. Epochs 1, 2, and 3 began with the presentation of offer 1, the presentation of offer 2, and the reward, respectively, and each lasted 500 ms. We found that 46.15% of neurons (n = 72/156) showed some sensitivity to task events, as indicated by individual cell ANOVAs of firing rate against epoch for the three task epochs and a 500 ms intertrial epoch ($p < 0.0001$, binomial test). All proportions presented below refer to all neurons, not just the ones that produced a significant response modulation.

Neurons Represent Value in a Common Currency-Like Format

Monkeys clearly attend to both probability and reward size in evaluating offers (Figure 2B). We found that the firing rates of a small but significant number of neurons significantly encoded reward size (n = 18/156, $p < 0.05$, linear regression) and probability (n = 12/156) in epoch 1. These proportions are both greater than would be expected by chance (binomial test, $\alpha = 0.05$, $p = 0.0003$ for reward size and $p = 0.025$ for probability.) Safe offers, which occurred on 12.5% of trials, introduce a negative correlation between reward size and probability, so trials with safe offers are excluded from this analysis. Therefore, reward size and probability were strictly uncorrelated in the design of the task.

Do single neurons represent both reward size and probability, or do neurons specialize for one or the other component variable, as lOFC neurons appear to (O'Neill and Schultz, 2010; Roesch et al., 2006)? To address this question, we compared regression coefficients for firing rate versus probability to coefficients from the regression of firing rate versus reward size (in epoch 1). We found a significant positive correlation between these coefficients ($r = 0.25$, $p = 0.0023$) (Figure 3A). We confirmed that this correlation is significant using a bootstrap (and thus, nonparametric) correlation test ($p = 0.0155$; see Experimental Procedures). These effects were even stronger using a 500 ms epoch beginning 100 ms later, suggesting that value responses in vmPFC may be sluggish ($r = 0.34$, $p < 0.0001$). These data are consistent with the idea that vmPFC represents value in a common currency-like format and suggest the possibility that these values may be compared here as well (Montague and Berns, 2002; Padoa-Schioppa, 2011).

If we assume that neurons represent offer values, defined here as an offer's reward size multiplied by its probability, we can assess the frequency of tuning for offer value in our sample. We find that responses of 10.9% (n = 17/156, $p = 0.0009$, binomial test) of neurons correlated with the value of offer 1 in epoch 1. This percentage rose to 16.66% (n = 26/156) using a 500 ms epoch that begins 100 ms later. Of these 26 neurons, 34.62% (n = 9/26) showed positive tuning for offer value in epoch 1 while the remainder showed negative tuning (this bias toward negative tuning is significant; binomial test, $p < 0.0001$). See Supplemental Information for neuronal response characteristics separated by offer 1 reward size.

Neurons Code Offer Values Simultaneously and Antagonistically

Figures 3B and 3C show value-related responses of an example neuron. Its firing rates signal the value of offer 1 in epoch 1 ($r = 0.18$, $p < 0.0001$, linear regression) and in epoch 2, although the direction is reversed, and the effect is weaker for the second epoch ($r = -0.09$, $p = 0.0025$). This neuron also showed tuning for offer 2 in epoch 2 ($r = 0.21$, $p < 0.0001$), meaning it coded both values simultaneously. Population data are shown in Figure 3D. In epoch 2, 10.26% of neurons (n = 16/156, this proportion is significant by a binomial test $p = 0.0022$,) encoded offer value 1 and 15.38% of neurons (n = 24/156, $p < 0.0001$) encoded offer value 2. The number of neurons signaling offer value 2 rose to 16.03% (n = 25/156, $p < 0.0001$. binomial test) 100 ms later.

The observation that tuning direction for offer values 1 and 2 are anticorrelated in our example neuron suggests that these values interact competitively to influence its firing when information about both options is available (Figure 4A). At the population

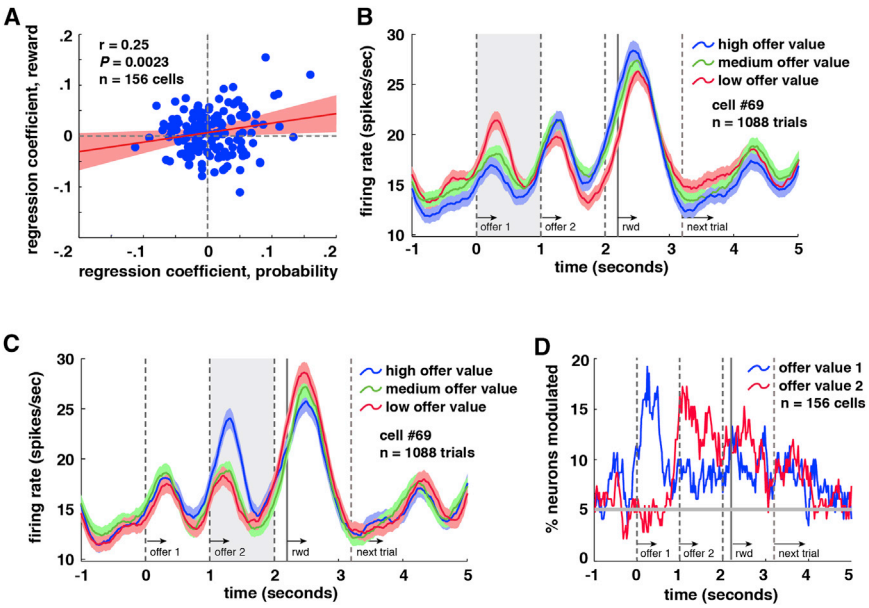

Figure 3. Coding of Offer Values in vmPFC Neurons

(A) Scatter plot of coefficients for tuning for probability (x axis) and reward size (y axis). Coefficients are significantly correlated, suggesting a common currency coding scheme. Each point corresponds to one neuron in our sample. Data are shown with a least-squares regression line and confidence intervals in red.

(B) Average responses (±1 SE in firing rate) of an example neuron to task events, separated by binned expected value of offer 1. This neuron showed tuning for offer value 1 during epoch 1 (shaded region).

(C) Responses of the same neuron (±1 SE in firing rate) separated by binned expected value of offer 2. The neuron showed tuning for offer value 2 during epoch 2 (shaded region).

(D) Plot of proportion of neurons (%) with responses significantly tuned to offer value 1 (blue) and offer value 2 (red) with a 500 ms sliding boxcar. Horizontal line indicates 5%; significance bar at alpha = 0.05.

level, regression coefficients for offer value 1 in epoch 2 are anticorrelated with coefficients for offer value 2 in the same epoch ($r = -0.218$, $p = 0.006$) (Figure 4B). We confirmed the significance of this correlation using a bootstrap correlation test ($p = 0.0061$; see Experimental Procedures). To match the criteria used above, these analyses do not include trials with safe options; however, if we repeat the analysis but include the safe offer trials as well, we still find an anticorrelation ($r = -0.162$, $p = 0.044$).

We have shown that neurons encode the value of offer 1 in epochs 1 and 2. But does vmPFC use a similar format to represent offers 1 and 2 as they initially appear, or does it use opposed ones? Our results support the former idea. We found a significant positive correlation between the regression coefficients for offer 1 in epoch 1 and those for offer 2 in epoch 2 ($r = 0.453$, $p < 0.0001$) (see Figure 4C). We confirmed the significance of this correlation using a bootstrap correlation test ($p < 0.0001$; see Experimental Procedures). Thus, whatever effect a larger offer 1 had on firing rates during epoch 1 in each neuron—whether excitatory or suppressive—the same effect was observed for those neurons to a larger offer 2 in epoch 2. This indicates that vmPFC neurons code the currently offered option in a common framework (cf. Lim et al., 2011).

Neurons Signal Chosen Offer Value, Not Unchosen Offer Value

Neurons in vmPFC represent the values of both offers simultaneously, but do they participate in selecting a preferred one? If they participate in choice, we may expect to see the gradual formation of a representation of the value of the chosen option and the dissolution of the value of the unchosen one. Figure 4D shows the proportion of neurons whose activity is significantly modulated by chosen offer values (blue) and by unchosen offer values (red). (Note that this figure shows a peak during epoch 3 that is even larger than the peak in epoch 2; this is because the value of the chosen offer was highly correlated with the value of the outcome, and outcome coding was stronger than other effects; see below.)

We found weak coding for the value of the chosen option even during epoch 1 (7.69% of cells, n = 12/156, binomial test; this proportion just barely achieves statistical significance, $p = 0.05$). This activity is not "precognitive" because monkeys can sometimes guess their chosen option if the first offer is good enough. We found coding of chosen value during the first 200 ms of the presentation of offer 2 (11.54% of cells, n = 18/156, $p = 0.0003$). We used this short epoch (200 ms instead of the 500 ms we used in other analyses) because it allows us to more closely inspect the time course of this signal. By a 200 ms epoch 200 ms later into the second epoch, chosen value coding was observed in 17.31% of cells ($n = 27/156$, $p < 0.0001$). In contrast, 7.69% of cells encoded the value of the unchosen offer during the first epoch (binomial test; again, this proportion is right at the significance threshold, $p = 0.05$), and only 6.4% (n = 10/156) of neurons encoded unchosen values at the beginning of the second epoch and 200 ms into it (not significant, $p = 0.159$). These results indicate that neurons in vmPFC preferentially encode the value of the chosen offer—and do so rapidly once both offers appear.

Variability in Firing Rates Predicts Choice

To explore the connection between neural activity in vmPFC and offer selection, we made a calculation similar to choice probability (Britten et al., 1996). For each neuron, we regressed firing rate in epoch 1 onto offer value, probability, and reward size. We then examined whether the sign of the residuals from this regression predicted choice (offer 1 versus offer 2) for each neuron. This analysis provides a measure of residual variance in firing rate after accounting for the three factors that influence value. We found a significant correlation between residual firing rate variance and choice in 11.53% (n = 18/156, $p = 0.0003$, binomial test) of cells, which is more than is expected by chance.

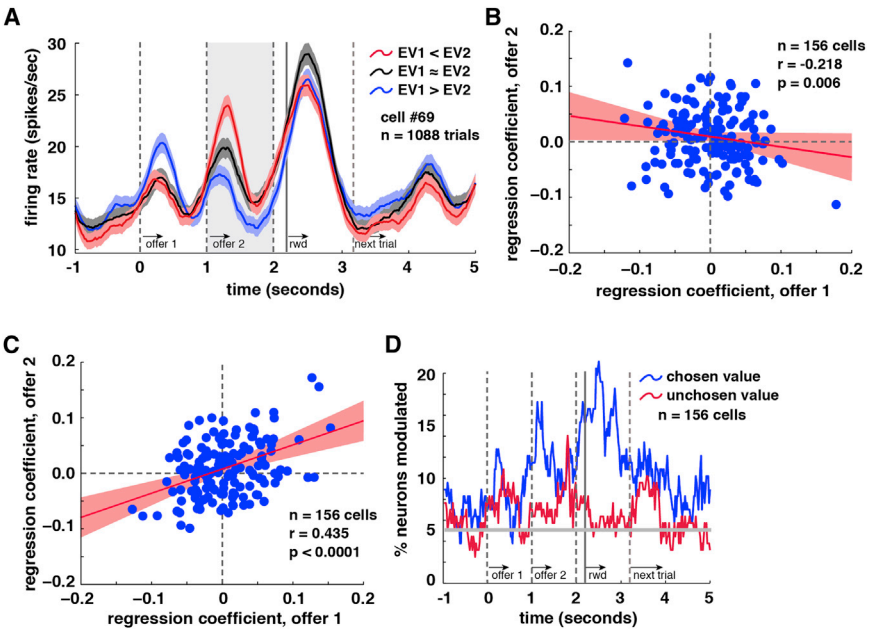

Figure 4. vmPFC Neuron Activity Related to Comparison and Choice

(A) Average responses of example neuron (±1 SE in firing rate) separated by binned expected value difference between offer values (offer value 1 minus offer value 2). During epoch 2, this neuron showed higher firing rates when offer value 2 was greater than offer value 1 (red) and lower firing when offer value 1 was greater than offer value 2 (blue).

(B) Scatter plot of coefficients for tuning for offer value 1 during epoch 2 (x axis) and for offer value 2 during epoch 2 (y axis). Least-squares regression line and confidence intervals are shown in red.

(C) Scatter plot of coefficients for tuning for offer value 1 during epoch 1 (x axis) and for offer value 2 during epoch 2 (y axis). Least-squares regression line and confidence intervals are shown in red.

(D) Plot of proportion of neurons that show a significant correlation between neural activity and the value of the chosen (blue) and unchosen (red) offers (500 ms sliding boxcar).

Similarly, residual variation in firing rate in response to offer value 2 during epoch 2 predicted choice in 12.18% of cells (n = 19/156, p = 0.0001, binomial test). This link between firing rates and choice is consistent with the fourth key prediction of the competitive inhibition hypothesis.

Neurons in vmPFC Strongly Encode Outcome Values

Outcome-monitoring signals were particularly strong during our task. Figure 5A shows responses of an example neuron with trials separated by gamble outcome. This neuron signaled received reward size in epoch 3 ($r = -0.11$, $p = 0.0047$, linear regression). We observed a significant relationship between firing rate and gamble outcome in 18.59% of cells (n = 29/156; $p < 0.0001$, binomial test) (Figure 5B). In an epoch beginning 400 ms later, this proportion rose to 25% of cells (n = 39/156; $p < 0.0001$). Of these cells, 56.41% (n = 22/39) showed negative tuning (no significant bias, $p = 0.55$, binomial test). Interestingly, outcome coding persisted across the delay between trials. Specifically, previous trial outcome was a major influence on firing rates during both epochs 1 (14.74% of cells, n = 23/156, $p < 0.0001$, binomial test) and 2 (16.03% of cells, n = 25/156, $p < 0.0001$) (Figure 5C).

Is the vmPFC coding format for outcome related to its coding format for offer values? We next compared tuning profiles for outcome and offer value 1 (we found that coding in epochs 1 and 2 is shared; see above). Specifically, we asked whether, in our population of cells, regression coefficients for offer value 1 in epoch 1 are correlated with regression coefficients for received reward size in epoch 3. We found a significant correlation between regression coefficients for offer value 1 in epoch 1 and regression coefficients for received reward size in epoch 3 ($r = 0.22$, $p = 0.0054$). This suggests that vmPFC neurons use a single coding scheme to represent offer values and represent outcomes.

Do vmPFC neurons signal outcomes or the difference between expected outcome and received outcome? To investigate this issue, we performed a stepwise regression to determine whether postoutcome responses in vmPFC are related to reward size (first) and to the probability of that reward (second). Specifically, we performed a stepwise regression on average neural firing rates in epoch 3 onto gamble outcome and the probability that the chosen option would yield a reward. To deal with the problem that many neurons have negative tuning, we flipped the values for neurons that had negative individual tuning profiles.

We first examined all risky trials together (medium reward size, blue/red bars, and high reward size, green/red bars). With these trials, gamble outcome regressor met the criteria for model inclusion ($\beta = 0.1058$, $p < 0.0001$), but the reward probability of the chosen option did not ($\beta = -0.0034$, $p = 0.8077$). We then repeated these analyses for the medium- and high-reward size trials together, in case there was an interaction with reward size. We find similar results when examining only trials where a blue option was chosen (gamble outcome: $\beta = 0.1224$, $p < 0.0001$; chosen option reward probability: $\beta = 0.0188$, $p = 0.4093$) and when examining only trials where a green option was chosen (gamble outcome: $\beta = 0.1211$, $p < 0.0001$; chosen option reward probability: $\beta = 0.0244$, $p = 0.1602$). This indicates that vmPFC neurons signal pure outcome, not the deviation of outcomes from expectation.

DISCUSSION

We recorded responses of neurons in area 14 of vmPFC while rhesus monkeys performed a gambling task with staggered presentation of offers. We observed four major effects. First, neurons carried an abstract value signal that depended on both probability and reward size. Second, when information about both options was available, responses were antagonistically modulated by values of the two options. Third, neurons rapidly came to signal the value of the chosen offer but not the unchosen

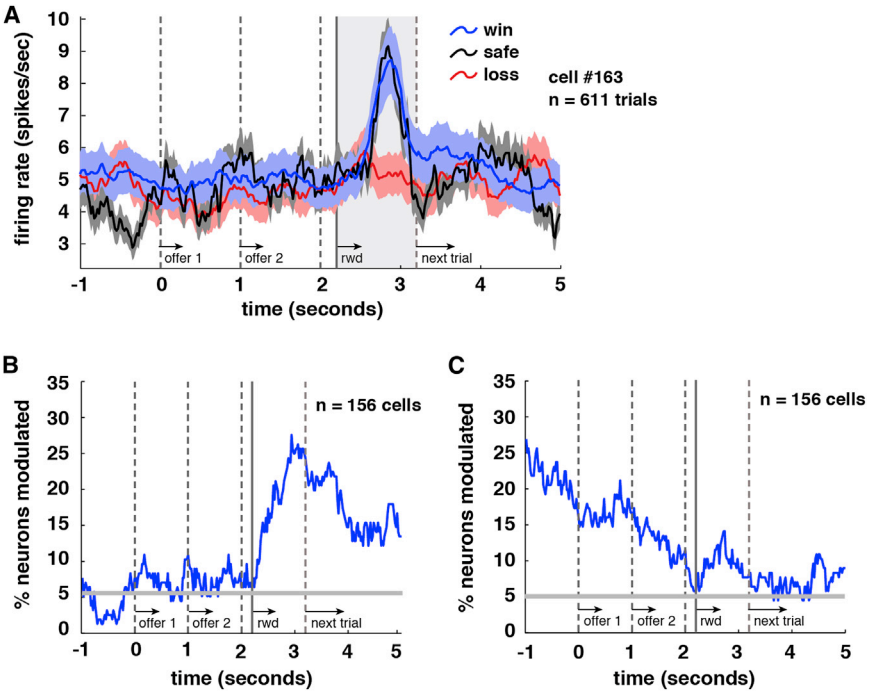

Figure 5. Coding of Outcomes in vmPFC Neurons

(A) Average responses (±1 SE in firing rate) of an example neuron to task events separated by outcome. This neuron showed a positive tuning for outcome during epoch 3 (shaded area).

(B) Plot of proportion of neurons significantly tuned for outcomes as a function of time in task using a 500 ms sliding window.

(C) Same data as in (B), but sorted for outcome on previous trial instead of on current trial. Influence of outcome on previous trial was strong and lasted throughout the current trial.

one. Fourth, after accounting for option values, residual variability in firing rates around the time of choice predicted choice. While we do not show directly that vmPFC neurons engage in mutual inhibition, these results are consistent with the theory that value comparison reflects a competition for control of vmPFC responses through mutual inhibition (Cisek, 2012; Hunt et al., 2012; Jocham et al., 2012; Wang, 2008).

Although reward correlates are observed in many brain areas, we suspect that vmPFC may be specialized for reward value comparisons. A great deal of neuroimaging evidence supports this hypothesis (Levy and Glimcher, 2012; Rushworth et al., 2011). The lOFC does not appear to integrate different dimensions of risky choices into a single value, suggesting that it may be predecisional. Moreover, value-coding neurons there do not show choice probability correlates, suggesting they may be only peripherally involved in choice (Padoa-Schioppa, 2013). Finally, human and monkey lesions in lOFC do not produce choice deficits but learning deficits. Indeed, recent comprehensive theories of lOFC function suggest that it carries multiple different values useful for controlling choice but does not itself implement choice (Rushworth et al., 2011; Wilson et al., 2014). In a similar vein, while the anterior cingulate cortex codes reward values, its signals appear to be postdecisional (Blanchard and Hayden, 2014; Cai and Padoa-Schioppa, 2012). These findings are consistent with the idea that dACC is a controller but not a decider (Shenhav et al., 2013). Finally, the lateral intraparietal cortex (LIP) is associated with choice processes, but it does not appear to represent values (Leathers and Olson, 2012) and does not show value comparison signals (Louie et al., 2011). These results suggest that choice occurs elsewhere; neuroimaging and anatomical evidence suggest that vmPFC is the site; our results endorse this idea.

Nonetheless, these results do not suggest that vmPFC is the only area in which value comparison occurs. Value comparison may, in some circumstances, occur in the lOFC, the ventral striatum (Cai et al., 2011), and the premotor cortex (Hunt et al., 2013). Indeed, it is not certain that value comparison occurs exclusively in one region instead of multiple regions acting in parallel (Cisek, 2012). However, in any of these cases, our results provide direct evidence for a specific mechanism by which value comparison occurs.

One limitation of the present study is that monkeys were overtrained on the task, which may change choice behavior or how reward information is represented in the brain. This is a limitation of all single-unit behavioral studies in monkeys. It is possible that large-scale recording grids combined with innovative recording techniques might help with this problem in the future.

Four recent reports describe response properties of vmPFC neurons. Bouret and Richmond (2010) demonstrated that neurons in area 14 preferentially encode internal sources of reward information, such as satiety, over external sources of reward information, such as visually offered reward or gamble offers. While we did not compare vmPFC to lOFC as they did, our results demonstrate that strong and significant external value and comparison signals can be readily observed in area 14 with a sufficiently demanding task. Monosov and Hikosaka (2012) showed that in a Pavlovian task, separate populations of area 14 neurons preferentially encode reward size and probability. Our recordings suggest that at least some neurons in area 14 can integrate probability and reward size into a combined signal. One possible explanation for the difference in the two sets of findings is that, unlike Monosov and Hikosaka, we used a choice task, which demands active consideration of both aspects of reward. Watson and Platt (2012) found that social information is prioritized in vmPFC (and in lOFC), even relative to its influence on preferences. In combination with our findings, these results suggest that social influences may be treated as qualitatively different than other factors that influence value (but see Smith et al., 2010). Rich and Wallis (2014) found generally weak and inconsistent responses in area 14 (which they call mOFC), suggesting that their task, which did not require value comparison, did not strongly selectively drive these neurons.

Relative to our recordings in a similar task in another medial prefrontal structure, dACC, we find that neuronal responses in vmPFC are weaker and have less consistent tuning directions (Hayden and Platt, 2010). This difference may reflect that we have not yet identified the ideal driving stimuli for vmPFC. Another possibility is a bias in recorded cell types. Unlike dACC, vmPFC lacks a prominent layer 5 (Vogt, 2009), which means that our sample of neurons may contain fewer output cells and more interneurons (Hayden et al., 2011a, 2011b). These responses may also simply be representative of vmPFC. The vmPFC responses we report here are generally small and long lasting, making them reminiscent of those observed in PCC (Hayden et al., 2008, 2009; Heilbronner et al., 2011). Intriguingly, PCC shows strong anatomical and functional connections with vmPFC (Andrews-Hanna et al., 2010; Vogt and Pandya, 1987) and, like it, is part of the poorly understood default mode network (Raichle and Gusnard, 2005). Integrating our understanding of default mode function with choice is an important goal for future studies.

Finally, we were surprised that the largest and most robust responses in vmPFC were outcome monitoring signals. Outcome monitoring signals are common in both ACC and PCC, and in these areas, they predict adjustments in behavior that follow specific outcomes (Hayden et al., 2008, 2011a). In contrast, the outcome signals we observed in vmPFC did not predict changes in behavior. This lack of an effect suggests that value monitoring signals in vmPFC may be somewhat automatic (that is, not contingent on the outcome having a specific effect) and are subject to a downstream gating process (that is, they do not affect behavior directly). Thus, these signals may be considered monitoring signals while those in cingulate may be more helpfully classified as control signals. Given the anatomy, we suspect that vmPFC may be one input for the control signals generated by cingulate cortex. Interestingly, a recent report suggests that monitoring signals that do not affect behavior are also observed on the dorsolateral surface of the prefrontal cortex (Genovesio et al., 2014).

In contrast to perceptual decision making, very little work has looked at the mechanisms of reward-based decisions. Kacelnik and colleagues (2011) have investigated this problem and have specifically compared two hypotheses: (1) the tug-of-war hypothesis, in which there is a mutual inhibition between value representations, and (2) the race-to-threshold hypothesis, in which value representations compete, noninteractively, and the first one to achieve some threshold is chosen. While Kacelnik's work provides strong support for the race-to-threshold model, ours would seem to support the tug-of-war hypothesis. In particular, the finding that vmPFC neurons gradually come to represent the value of the chosen option at the expense of the unchosen would appear difficult to reconcile with a pure race-to-threshold model. Instead, our finding of value difference signals is consistent with a version of the race-to-threshold model that involves competition between racing value representations. Nonetheless, these results do not endorse a single model of reward-based choice. Unfortunately, by presenting options asynchronously, we were unable to measure reaction times in our task, meaning a direct comparison is impossible. It seems that further work will be needed to more fully compare these two hypotheses.

One of the most interesting aspects of these postreward signals is that vmPFC appeared to use a similar coding framework to encode outcomes and offers. One speculative explanation for this finding is that offer signals are essentially reactivations of reward representations (Kahnt et al., 2011). Monkeys might consider offers by predicting the activation they would generate if they received that reward. If so, then choice may work through competition between mental simulations of outcomes. While this hypothesis is speculative, it is at least tenuously supported by the existence of direct anatomical projections to vmPFC from hippocampus and amygdala, structures associated with associative learning (Carmichael and Price, 1995), and by evidence of co-occurring outcome and value signals throughout the medial frontal lobe (Luk and Wallis, 2009). Future studies will be needed to more fully test this hypothesis.

EXPERIMENTAL PROCEDURES

Surgical 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. Two male rhesus macaques (*Macaca mulatta*) served as subjects. A small prosthesis for holding the head was used. Animals were habituated to laboratory conditions and then trained to perform oculomotor tasks for liquid reward. A Cilux recording chamber (Crist Instruments) was placed over the vmPFC. Position was verified by magnetic resonance imaging with the aid of a Brainsight system (Rogue Research Inc.). Animals received appropriate analgesics and antibiotics after all procedures. Throughout both behavioral and physiological recording sessions, the chamber was kept sterile with regular antibiotic washes and sealed with sterile caps.

Recording Site

We approached vmPFC through a standard recording grid (Crist Instruments). We defined vmPFC as the coronal planes situated between 29 and 44 mm rostral to the interaural plane, the horizontal planes situated between 0 and 9 mm from the ventral surface of vmPFC, and the sagittal planes between 0 and 8 mm from the medial wall (Figures 1C and S1). These coordinates correspond to area 14 (Ongür and Price, 2000). Our recordings were made from a central region within this zone. We confirmed recording location before each recording session using our Brainsight system with structural magnetic resonance images taken before the experiment. Neuroimaging was performed at the Rochester Center for Brain Imaging, on a Siemens 3T MAGNETOM Trio Tim using 0.5 mm voxels. We confirmed recording locations by listening for characteristic sounds of white and gray matter during recording, which in all cases matched the loci indicated by the Brainsight system with an error of <1 mm in the horizontal plane and <2 mm in the z direction.

Electrophysiological Techniques

Single electrodes (Frederick Haer & Co., impedance range 0.8 to 4 MΩ) were lowered using a microdrive (NAN Instruments) until waveforms of between one and three neuron(s) were isolated. Individual action potentials were isolated on a Plexon system (Plexon). Neurons were selected for study solely on the basis of the quality of isolation; we never preselected based on task-related response properties. All collected neurons for which we managed to obtain at least 500 trials were analyzed; no neurons that surpassed our isolation criteria were excluded from analysis.

Eye Tracking and Reward Delivery

Eye position was sampled at 1,000 Hz by an infrared eye-monitoring camera system (SR Research). Stimuli were controlled by a computer running Matlab (Mathworks) with Psychtoolbox (Brainard, 1997) and Eyelink Toolbox (Cornelissen et al., 2002). Visual stimuli were colored rectangles on a computer monitor placed 57 cm from the animal and centered on its eyes (Figure 1A).

A standard solenoid valve controlled the duration of juice delivery. The relationship between solenoid open time and juice volume was established and confirmed before, during, and after recording.

Behavioral Task

Monkeys performed a two-option gambling task (Figures 1A and 1B). The task was similar to one we have used previously (Hayden et al., 2010, 2011a) with two major differences: (1) offers were presented asynchronously, and (2) two different winning reward sizes (medium and large) offers were available, depending on the gamble.

Two offers were presented on each trial. Each offer was represented by a rectangle 300 pixels tall and 80 pixels wide (11.35° of visual angle tall and 4.08° of visual angle wide). Options offered either a gamble or a safe (100% probability) bet for liquid reward. Gamble offers were defined by two parameters, reward size and probability. Each gamble rectangle was divided into two portions, one red and the other either blue or green. The size of the blue or green portions signified the probability of winning a medium (mean 165 μl) or large reward (mean 240 μl), respectively. These probabilities were drawn from a uniform distribution between 0% and 100%. The rest of the bar was colored red; the size of the red portion indicated the probability of no reward. Safe offers were entirely gray and always carried a 100% probability of a small reward (125 μl).

On each trial, one offer appeared on the left side of the screen and the other appeared on the right. Offers were separated from the fixation point by 550 pixels (27.53° of visual angle). The side of the first and second offer (left and right) were randomized by trial. Each offer appeared for 400 ms and was followed by a 600 ms blank period. Monkeys were free to fixate upon the offers when they appeared (and in our casual observations almost always did so). After the offers were presented separately, a central fixation spot appeared and the monkey fixated on it for 100 ms. Following this, both offers appeared simultaneously and the animal indicated its choice by shifting gaze to its preferred offer and maintaining fixation on it for 200 ms. Failure to maintain gaze for 200 ms did not lead to the end of the trial, but instead returned the monkey to a choice state; thus, monkeys were free to change their mind if they did so within 200 ms (although in our observations, they seldom did so). Following a successful 200 ms fixation, the gamble was immediately resolved and reward delivered. Trials that took more than 7 s were considered inattentive trials and were not included in analysis (this removed <1% of trials). Outcomes that yielded reward were accompanied by a visual cue: a white circle in the center of the chosen offer (see Figure 1A). All trials were followed by an 800 ms intertrial interval with a blank screen.

Probabilities were drawn from uniform distributions with a resolution only limited by the size of the computer screen's pixels. This let us present hundreds of unique gambles. Offer types were selected at random with a 43.75% probability of blue gamble, a 43.75% probability of green gambles, and 12.5% probability of safe offers.

Statistical Methods

PSTHs were constructed by aligning spike rasters to the presentation of the first offer and averaging firing rates across multiple trials. Firing rates were calculated in 20 ms bins but were generally analyzed in longer (500 ms) epochs. For display, PSTHs were smoothed using a 200 ms running boxcar.

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

Throughout data collection, reward for gray, blue, and green offers were associated with a few different sets of reward sizes due in part to the use of two different juicer solenoids. Despite this, reward sizes maintained the same sizes relative to each other. To account for overall variations in reward size, our analyses consistently make use of an ordinal coding of reward size, with gray, blue, and green offers offering 1, 2, and 3 juice units, respectively.

To test if certain signals tend to occur within the same neurons, we used the following bootstrap method. For each neuron, we calculated regression coefficients for those signals. We then calculated the correlation between those two sets of regression coefficients. We repeated this process 10,000 times using randomly reshuffled firing rates. We used the percentile at which the original data correlation coefficient fell in this distribution of randomized correlation coefficients as the p value for a single-tailed test, which we multiplied by two to calculate the p value for a two-tailed test. For example, if the correlation coefficient from the original data was greater than 90% of the randomized correlation coefficients, we considered the tuning significant at $p = 0.05$.

We performed one analysis to investigate how variance in firing related to variance in preference. First, we determined the best-fit curve for firing rate in epoch 1 as a function of the expected value of the first offer. In one analysis we fit to a line; in a second one we fit to the best-fit second-order polynomial. (We tested third and fourth order polynomials as well and found similar results; data not reported.) We next classified each trial based on whether the observed firing rate in epoch 1 was greater or lower than a value predicted by the best-fit functions. Finally, we correlated choice (coded as 1 or 0, indicating choice of offer 1 or 2) with whether firing rate was higher or lower than expected, on a trial-by-trial basis. We tested for a significant relation within each individual neuron using Pearson's correlation test of these two sets of variables with trial as the unit of analysis. We then repeated this analysis for epoch 2.

In this paper we made a deliberate decision to use expected values rather than subjective values in correlating neural activity with value. The primary reason for this is that it is the most agnostic approach one can take with regard to the causes of risk seeking. While it may be standard practice to transform values into utilities, behavioral economics has demonstrated that utility curve shape cannot explain risk attitudes in general (Kahneman and Tversky, 2000; Rabin, 2000). Our research has demonstrated that these arguments apply to monkeys as well (Hayden et al., 2010; Heilbronner and Hayden, 2013; Strait and Hayden, 2013). Moreover, using expected values bypasses the troubling question of what timescale to use to determine value functions, a decision that can have great consequences on data interpretation (Sugrue et al., 2005). Fortunately, the question of whether we use expected value or subjective value is unlikely to have more than a marginal effect on our numbers and no effect on the qualitative findings we report. This is most directly demonstrated by the fact that our findings all reproduce if we restrict our analyses to high- and medium-value gambles alone. Because these gambles have only two outcomes, utility transformations have no effect. In any case, because the mapping function between firing rate and value is nonlinear and quite noisy, the subtle changes caused by using subjective value are almost certain to produce effects that are around the level of statistical noise.

SUPPLEMENTAL INFORMATION

Supplemental Information includes two figures and can be found with this article online at http://dx.doi.org/10.1016/j.neuron.2014.04.032.

ACKNOWLEDGMENTS

This research was supported by a R00 (DA027718), a NARSAD Young Investigator Reward from the Brain and Behavior Research Foundation, and a Sloan Foundation fellowship to B.Y.H. We thank Tim Behrens, Sarah Heilbronner, and John Pearson for useful discussions and Aaron Roth and Marc Mancarella for assistance in data collection.

Accepted: April 21, 2014
Published: May 29, 2014

REFERENCES

Andrews-Hanna, J.R., Reidler, J.S., Sepulcre, J., Poulin, R., and Buckner, R.L. (2010). Functional-anatomic fractionation of the brain's default network. Neuron *65*, 550–562.

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.

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.

Boorman, E.D., Behrens, T.E., Woolrich, M.W., and Rushworth, M.F. (2009). How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. Neuron *62*, 733–743.

Boorman, E.D., Rushworth, M.F., and Behrens, T.E. (2013). Ventromedial prefrontal and anterior cingulate cortex adopt choice and default reference frames during sequential multi-alternative choice. J. Neurosci. *33*, 2242–2253.

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.

Brainard, D.H. (1997). The Psychophysics Toolbox. Spat. Vis. *10*, 433–436.

Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., and Movshon, J.A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. Vis. Neurosci. *13*, 87–100.

Cai, X., and Padoa-Schioppa, C. (2012). Neuronal encoding of subjective value in dorsal and ventral anterior cingulate cortex. J. Neurosci. *32*, 3791–3808.

Cai, X., Kim, S., and Lee, D. (2011). Heterogeneous coding of temporally discounted values in the dorsal and ventral striatum during intertemporal choice. Neuron *69*, 170–182.

Camille, N., Griffiths, C.A., Vo, K., Fellows, L.K., and Kable, J.W. (2011). Ventromedial frontal lobe damage disrupts value maximization in humans. J. Neurosci. *31*, 7527–7532.

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.

Cisek, P. (2012). Making decisions through a distributed consensus. Curr. Opin. Neurobiol. *22*, 927–936.

Cornelissen, F.W., Peters, E.M., and Palmer, J. (2002). The Eyelink Toolbox: eye tracking with MATLAB and the Psychophysics Toolbox. Behav. Res. Methods Instrum. Comput. *34*, 613–617.

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.

Fellows, L.K. (2006). Deciding how to decide: ventromedial frontal lobe damage affects information acquisition in multi-attribute decision making. Brain. *129*, 944–952.

FitzGerald, T.H., Seymour, B., and Dolan, R.J. (2009). The role of human orbitofrontal cortex in value comparison for incommensurable objects. J. Neurosci. *29*, 8388–8395.

Genovesio, A., Tsujimoto, S., Navarra, G., Falcone, R., and Wise, S.P. (2014). Autonomous encoding of irrelevant goals and outcomes by prefrontal cortex neurons. J. Neurosci. *34*, 1970–1978.

Glimcher, P.W. (2003). Decisions, Uncertainty, and the Brain: the Science of Neuroeconomics. (Cambridge, Mass.: MIT Press).

Hayden, B.Y., and Platt, M.L. (2010). Neurons in anterior cingulate cortex multiplex information about reward and action. J. Neurosci. *30*, 3339–3346.

Hayden, B.Y., Parikh, P.C., Deaner, R.O., and Platt, M.L. (2007). Economic principles motivating social attention in humans. Proc. Biol. Sci. *274*, 1751–1756.

Hayden, B.Y., Nair, A.C., McCoy, A.N., and Platt, M.L. (2008). Posterior cingulate cortex mediates outcome-contingent allocation of behavior. Neuron *60*, 19–25.

Hayden, B.Y., Pearson, J.M., and Platt, M.L. (2009). Fictive reward signals in the anterior cingulate cortex. Science *324*, 948–950.

Hayden, B.Y., Heilbronner, S.R., and Platt, M.L. (2010). Ambiguity aversion in rhesus macaques. Front. Neurosci. *4*, 166.

Hayden, B.Y., Heilbronner, S.R., Pearson, J.M., and Platt, M.L. (2011a). Surprise signals in anterior cingulate cortex: neuronal encoding of unsigned reward prediction errors driving adjustment in behavior. J. Neurosci. *31*, 4178–4187.

Hayden, B.Y., Pearson, J.M., and Platt, M.L. (2011b). Neuronal basis of sequential foraging decisions in a patchy environment. Nat. Neurosci. *14*, 933–939.

Heilbronner, S.R., and Hayden, B.Y. (2013). Contextual factors explain risk-seeking preferences in rhesus monkeys. Front. Neurosci. *7*, 7.

Heilbronner, S.R., Hayden, B.Y., and Platt, M.L. (2011). Decision salience signals in posterior cingulate cortex. Front. Neurosci. *5*, 55.

Hosokawa, T., Kennerley, S.W., Sloan, J., and Wallis, J.D. (2013). Single-neuron mechanisms underlying cost-benefit analysis in frontal cortex. J. Neurosci. *33*, 17385–17397.

Hunt, L.T., Kolling, N., Soltani, A., Woolrich, M.W., Rushworth, M.F., and Behrens, T.E. (2012). Mechanisms underlying cortical activity during value-guided choice. Nature Neurosci. *15*, 470–476.

Hunt, L.T., Woolrich, M.W., Rushworth, M.F., and Behrens, T.E. (2013). Trial-type dependent frames of reference for value comparison. PLoS Comput. Biol. *9*, e1003225.

Hussar, C.R., and Pasternak, T. (2012). Memory-guided sensory comparisons in the prefrontal cortex: contribution of putative pyramidal cells and interneurons. J. Neurosci. *32*, 2747–2761.

Jocham, G., Hunt, L.T., Near, J., and Behrens, T.E. (2012). A mechanism for value-guided choice based on the excitation-inhibition balance in prefrontal cortex. Nat. Neurosci. *15*, 960–961.

Kacelnik, A., Vasconcelos, M., Monteiro, T., and Aw, J. (2011). Darwin's "tug-of-war" vs. starlings' "horse-racing": how adaptations for sequential encounters drive simultaneous choice. Behav. Ecol. Sociobiol. *65*, 547–558.

Kahneman, D., and Tversky, A. (2000). Choices, Values, and Frames. (Cambridge: Cambridge University Press).

Kahnt, T., Heinzle, J., Park, S.Q., and Haynes, J.D. (2011). Decoding the formation of reward predictions across learning. J. Neurosci. *31*, 14624–14630.

Kennerley, S.W., Behrens, T.E., and Wallis, J.D. (2011). Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. Nat. Neurosci. *14*, 1581–1589.

Leathers, M.L., and Olson, C.R. (2012). In monkeys making value-based decisions, LIP neurons encode cue salience and not action value. Science *338*, 132–135.

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.

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.

Louie, K., Grattan, L.E., and Glimcher, P.W. (2011). Reward value-based gain control: divisive normalization in parietal cortex. J. Neurosci. *31*, 10627–10639.

Luk, C.H., and Wallis, J.D. (2009). Dynamic encoding of responses and outcomes by neurons in medial prefrontal cortex. J. Neurosci. *29*, 7526–7539.

Machens, C.K., Romo, R., and Brody, C.D. (2005). Flexible control of mutual inhibition: a neural model of two-interval discrimination. Science *307*, 1121–1124.

McCoy, A.N., and Platt, M.L. (2005). Risk-sensitive neurons in macaque posterior cingulate cortex. Nat. Neurosci. *8*, 1220–1227.

Monosov, I.E., and Hikosaka, O. (2012). Regionally distinct processing of rewards and punishments by the primate ventromedial prefrontal cortex. J. Neurosci. *32*, 10318–10330.

Monosov, I.E., and Hikosaka, O. (2013). Selective and graded coding of reward uncertainty by neurons in the primate anterodorsal septal region. Nat. Neurosci. *16*, 756–762.

Montague, P.R., and Berns, G.S. (2002). Neural economics and the biological substrates of valuation. Neuron *36*, 265–284.

Nienborg, H., and Cumming, B.G. (2009). Decision-related activity in sensory neurons reflects more than a neuron's causal effect. Nature *459*, 89–92.

Noonan, M.P., Walton, M.E., Behrens, T.E., Sallet, J., Buckley, M.J., and Rushworth, M.F. (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., Esber, G.R., Cerri, D.H., Stalnaker, T.A., and Schoenbaum, G. (2013). Risk-responsive orbitofrontal neurons track acquired salience. Neuron *77*, 251–258.

Ongür, D., and Price, J.L. (2000). The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. Cereb. Cortex *10*, 206–219.

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

Padoa-Schioppa, C. (2011). Neurobiology of economic choice: a good-based model. Annu. Rev. Neurosci. *34*, 333–359.

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

Padoa-Schioppa, C., and Assad, J.A. (2006). Neurons in the orbitofrontal cortex encode economic value. Nature *441*, 223–226.

Philiastides, M.G., Biele, G., and Heekeren, H.R. (2010). A mechanistic account of value computation in the human brain. Proc. Natl. Acad. Sci. USA *107*, 9430–9435.

Rabin, M. (2000). Risk aversion and expected-utility theory: A calibration theorem. Econometrica *68*, 1281–1292.

Raichle, M.E., and Gusnard, D.A. (2005). Intrinsic brain activity sets the stage for expression of motivated behavior. J. Comp. Neurol. *493*, 167–176.

Rangel, A., and Clithero, J.A. (2012). Value normalization in decision making: theory and evidence. Curr. Opin. Neurobiol. *22*, 970–981.

Rangel, A., Camerer, C., and Montague, P.R. (2008). A framework for studying the neurobiology of value-based decision making. Nat. Rev. Neurosci. *9*, 545–556.

Rich, E.L., and Wallis, J.D. (2014). Medial-lateral organization of the orbitofrontal cortex. J. Cogn. Neurosci. 1–16.

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.

Romo, R., Hernández, A., Zainos, A., Lemus, L., and Brody, C.D. (2002). Neuronal correlates of decision-making in secondary somatosensory cortex. Nat. Neurosci. *5*, 1217–1225.

Rushworth, M.F., 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.

Seo, H., and Lee, D. (2009). Behavioral and neural changes after gains and losses of conditioned reinforcers. J. Neurosci *29*, 3627–3641.

Shenhav, A., Botvinick, M.M., and Cohen, J.D. (2013). The expected value of control: an integrative theory of anterior cingulate cortex function. Neuron *79*, 217–240.

Smith, D.V., Hayden, B.Y., Truong, T.K., Song, A.W., Platt, M.L., and Huettel, S.A. (2010). Distinct value signals in anterior and posterior ventromedial prefrontal cortex. J. Neurosci. *30*, 2490–2495.

So, N., and Stuphorn, V. (2012). Supplementary eye field encodes reward prediction error. J. Neurosci. *32*, 2950–2963.

Strait, C.E., and Hayden, B.Y. (2013). Preference patterns for skewed gambles in rhesus monkeys. Biol. Lett. *9*, 20130902.

Sugrue, L.P., Corrado, G.S., and Newsome, W.T. (2005). Choosing the greater of two goods: neural currencies for valuation and decision making. Nat. Rev. Neurosci. *6*, 363–375.

Vogt, B.A. (2009). Cingulate neurobiology and disease. (Oxford, New York: Oxford University Press).

Vogt, B.A., and Pandya, D.N. (1987). Cingulate cortex of the rhesus monkey: II. Cortical afferents. J. Comp. Neurol. *262*, 271–289.

Walton, M.E., Behrens, T.E., Buckley, M.J., Rudebeck, P.H., and Rushworth, M.F. (2010). Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. Neuron *65*, 927–939.

Wang, X.J. (2008). Decision making in recurrent neuronal circuits. Neuron *60*, 215–234.

Watson, K.K., and Platt, M.L. (2012). Social signals in primate orbitofrontal cortex. Curr. Biol. *22*, 2268–2273.

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.

Yamada, H., Tymula, A., Louie, K., and Glimcher, P.W. (2013). Thirst-dependent risk preferences in monkeys identify a primitive form of wealth. Proceedings of the National Academy of Sciences.