Behavioral/Systems/Cognitive

Range-Adapting Representation of Economic Value in the Orbitofrontal Cortex

Camillo Padoa-Schioppa

Department of Anatomy and Neurobiology, Washington University in St. Louis, St. Louis, Missouri 63110, and Department of Neurobiology, Harvard University, Boston, Massachusetts 02115

While making economic choices, individuals assign subjective values to the available options. Values computed in different behavioral conditions, however, can vary substantially. The same person might choose some times between goods worth a few dollars, and other times between goods worth thousands of dollars, or more. How does the brain system that computes values—the "valuation system"—handle this large variability? Here we show that the representation of value in the orbitofrontal cortex (OFC), an area implicated in value assignment during economic choice, adapts to the behavioral condition of choice and, more specifically, to the range of values available in any given condition. In the experiments, monkeys chose between different juices and their choice patterns provided a measure of subjective value. Value ranges were varied from session to session and, in each session, OFC neurons encoded values in a linear way. Across the population, the neuronal sensitivity (defined as the change in neuronal activity elicited by the increase in one value unit) was inversely proportional to the value range. Conversely, the neuronal activity range did not depend on the value range. This phenomenon of range adaptation complements that of menu invariance observed in a previous study. Indeed, the activity of each neuron adapts to the range values it encodes but does not depend on other available goods. Our results thus suggest that the representation of value in the OFC is at one time instantiative of preference transitivity (menu invariance) and computationally efficient (range adaptation).

Introduction

Multiple lines of evidence implicate the orbitofrontal cortex (OFC) in value assignment underlying economic choice. Neural activity in this area is generally elevated when subjects make choices compared with when choices are made for them (Arana et al., 2003; Chaudhry et al., 2009). Moreover, blood oxygenation correlates with assigned values in a variety of choice tasks (Plassmann et al., 2007; Tom et al., 2007; Valentin et al., 2007; Hare et al., 2009; Venkatraman et al., 2009). Individual OFC neurons encode the subjective value monkeys assign to different juices when they choose between them (Padoa-Schioppa and Assad, 2006). Interestingly, neuronal activity in this area is also affected by the effort exerted by the animal to obtain the juice, by the probability with which the juice is delivered (Kennerley et al., 2009), and by the duration of a delay intervening before juice delivery (Roesch and Olson, 2005). Finally, OFC lesions specifically disrupt choice behavior (Gallagher et al., 1999; Rahman et al., 1999; Fellows and Farah, 2007; Machado and Bachevalier, 2007). Current evidence

Received Aug. 1, 2009; revised Oct. 4, 2009; accepted Oct. 7, 2009.

This work was supported by a Pathway to Independence Award from the National Institute of Mental Health (Grant R00-MH080852). I thank W. Schultz for extended discussions and for sharing unpublished results from his laboratory. I also thank J. Assad and L. Snyder for helpful comments on an earlier version of this manuscript and W. Assad for providing software for behavioral control. Analysis of the main dataset was started at Harvard University and completed at Washington University, at which I collected additional data and prepared this manuscript. The content is solely the responsibility of the author and does not necessarily represent the official views of the National Institute of Mental Health or the National Institutes of Health.

Correspondence should be addressed to Dr. Camillo Padoa-Schioppa, Department of Anatomy and Neurobiology, Washington University in St. Louis, 660 South Euclid Avenue, Campus Box 8108, St. Louis, MO 63110. E-mail: camillo@wustl.edu.

DOI:10.1523/JNEUROSCI.3751-09.2009

Copyright © 2009 Society for Neuroscience 0270-6474/09/2914004-11\$15.00/0

thus suggests that economic choices may be based on values represented in the OFC.

A broad and fundamental question is whether and how this representation of value depends on the behavioral context of choice. Notably, the behavioral context can vary in multiple ways. For example, a person in a wine store might serially compare different pairs of bottles: she might first compare a Nebbiolo to a Bordeaux, and later compare the same Nebbiolo to a Zinfandel. Changes in the alternative to a given option, which typically occur moment to moment, are referred to as changes of menu. However, the context of choice can also change in other ways. For example, while visiting northern Italy, the same person might choose between buying a 12-bottle case of that Nebbiolo and spending one more night in Venice. Broad changes such as from a wine store to a vacation, in which value ranges vary, are referred to as changes of condition. In previous work, we found that the representation of value in the OFC is invariant for changes of menu (Padoa-Schioppa and Assad, 2008). Here we investigated whether and how this representation of value depends on the behavioral condition.

Several reasons led us to hypothesize that the representation of value in the OFC might be subject to neuronal adaptation. First, an early observation of Tremblay and Schultz, considered in the light of our previous results, suggests an adapting encoding of value (Tremblay and Schultz, 1999; Padoa-Schioppa and Assad, 2008). Second, from a computational perspective, a valuation system endowed with neuronal adaptation would present clear advantages. Indeed, while values faced in different conditions can vary by orders of magnitude, any neuronal representation of value is ultimately subject to physiological limitations on neuronal firing rates. These limitations imply that choices cannot be ubiquitously accurate (i.e., consistent) unless the valuation system effectively adapts to the range of values available in any given condition. On this basis, we predicted an inverse relationship between the sensitivity of value-encoding neurons and the range of values available in any particular condition. Our results confirmed this hypothesis.

Materials and Methods

Main dataset. For this study, we pooled and reanalyzed two datasets originally collected to examine the neuronal representation of economic value [experiment 1 (Exp. 1)], 931 cells (Padoa-Schioppa and Assad, 2006)] and its independence from the menu [Exp. 2, 557 cells (Padoa-Schioppa and Assad, 2008)]. In all the experiments, monkeys chose in every trial between 2 juices labeled A and B, with A preferred. In Exp. 1, two juices (A and B) were used in each session. Offers were represented by sets of colored squares on a computer monitor, and monkeys indicated their choice through eye movements (Fig. 1a). In any given session, the amounts of the two juices varied from trial to trial, and the behavior of the monkey presented a quality/quantity trade-off. When offered equal amounts of the two juices, monkeys would consistently choose juice A. However, when juice B was offered in sufficiently large amounts, the monkeys would choose it. The "relative value" of the two juices was inferred from the indifference point-the quantity ratio for which the monkey would choose either juice equally often. For example, if the monkey was indifferent between 1A and 3B, we inferred that the value of 1A was equal to the value of 3B (relative value = 3). Trials were divided into "trial types" based on the offer type and the choice. For example, a monkey facing the offer type 4B:1A could choose either 1A or 4B, corresponding to the two trial types (4B:1A, 1A) and (4B:1A, 4B). Neuronal firing rates were examined in seven time windows: 0.5 s preoffer, 0.5 s postoffer, late delay (0.5–1.0 s after the offer), 0.5 s pre-go, reaction time (from "go" to saccade), 0.5 s prejuice, and 0.5 s postjuice. A "neuronal response" was defined as the activity of one neuron in one time window as a function of the trial type. In Exp. 2, we used identical procedures for behavioral control and neuronal recordings except that three juices (A, B, and C) were used in each session (Padoa-Schioppa and Assad, 2008). Trials with the three juice pairs (A:B, B:C, and A:C) were randomly interleaved. For the present study, we separated the three sets of trials and we relabeled, for each juice pair, the two juices as A and B (A preferred). Thus, each neuron recorded in Exp. 2 provided three separate neuronal responses for each time window (separate sets of trials). Data from Exp. 2 were then analyzed with the same procedures used for Exp. 1.

The analysis performed here builds on previous results showing that the population of neurons in the OFC encodes three variables: *offer value* (the value of one of the two juices) (Fig. 1*b*,*c*), *chosen value* (the value of the chosen juice) (Fig. 1*d*), and *taste* (a binary variable indicating the identity of the chosen juice but not its quantity). Only neuronal responses significantly modulated by the trial type (ANOVA, p < 0.001) were included in the analysis. For each response, we performed a linear regression on each variable. A variable was said to "explain" the response if the regression slope differed significantly from zero (p < 0.05). If a given response was explained by more than one variable, the encoded variable was identified as the one with the highest R^2 . The encoding of value was linear in the sense that adding a quadratic term to the regression generally failed to significantly improve the fit (Padoa-Schioppa and Assad, 2006).

The present study focused on neuronal responses encoding the *offer value* or the *chosen value*. For each response, "slope" and "intercept" were obtained from the linear regression onto the encoded variable. The "value range" was defined as the difference between the maximum value and the minimum value available in that session, the "activity range" was defined as the product of the slope by the value range, and the "baseline activity" was defined as the activity corresponding to the minimum value available in that session. Neuronal responses recorded in different time windows likely differ for their functional significance. However, the present results held true in each time window. We thus describe our



Figure 1. Experimental design and preliminary analysis. *a*, At the beginning of each trial, the monkey fixated the center of a computer monitor. Two sets of squares appeared on opposite sides of the fixation point ("offer"). Different colors of the squares indicated different juice types and the number of squares indicated the juice amount. After a randomly variable delay (1-2s), two saccade targets appeared near the offers ("go"). The monkey indicated its choice and maintained fixation on the saccade target for 0.75 s before juice delivery. The trial was aborted if the monkey broke fixation before the go. For any juice pair, the quantities of the two juices varied randomly. For any given pair of offers (offer type), left/right positions were randomly counterbalanced. **b**, Neuronal response encoding the offer value A. In the left, black circles represent the behavioral choice pattern, i.e., the percentage of B choices (v axis) recorded for different offer types (x axis). The relative value, obtained from a sigmoid fit, is indicated in the panel. In the same panel, red symbols represent the neuronal firing rate \pm SEM. Diamonds and circles represent trials in which the monkey chose juice A and juice B, respectively. On the right, the same firing rate (y axis) is plotted against the encoded variable offer value A (x axis). The black line is obtained from the linear regression and the regression parameters (slope and intercept) are indicated. *c*, Neuronal response encoding the *offer value B* plotted against the offer type (left) and against the encoded variable (right). *d*, Neuronal response encoding the chosen value plotted against the offer type (left) and against the encoded variable (right). Regression slopes are expressed in conventional value units uV, corresponding to uA for offer value A and uB for offer value B and chosen value (see Materials and Methods). sp, Spikes.

findings pooling responses from different time windows, although we also report the main results obtained for individual time windows.

Each neuron in this dataset (main dataset) was recorded in one session with one value range. However, sizable groups of neurons were recorded with different value ranges. To examine how the encoding depended on the value range, we thus proceeded with a population-based approach. We divided the dataset in "sub-populations" of neuronal responses recorded with different value ranges. We then considered one encoding parameter (slope, activity range, etc.), analyzed it for each subpopulation, and compared the results across sub-populations (see Fig. 4*d*). This can be viewed as a repeated-measure procedure performed at the level of the neuronal population.

Unless otherwise indicated, we always express values in conventional units uV, defined as uA for *offer value A*, uB for *offer value B*, and uB for *chosen value*. Units uA (uB) are defined as equal to the value of 1 drop of juice A (juice B). Value ranges were thus "quantized" for *offer value* responses, while they could assume any numerical value for *chosen value* responses (see Fig. 4*a*–*c*, *x* axes). Subpopulations were naturally defined for *offer value* responses. For *chosen value* responses, we divided value ranges in bins of 1–2 uB width and we computed the average value range for each bin. A similar procedure was used to analyze neuronal responses in nonconventional value units (see Fig. 8).

Range adaptation for individual neurons. Population analyses of the main dataset indicated that the encoding of value in the OFC adapts to the range of values available in the behavioral condition. To confirm that adaptation takes place at the level of individual neurons, we collected an additional 111 cells from the OFC of one monkey. The experimental procedures were similar to those previously de-

scribed (Padoa-Schioppa and Assad, 2006), with minor differences. Briefly, all surgical procedures were as previously described. The behavioral task was controlled by a custom-written software (Asaad and Eskandar, 2008) based on Matlab (PC version R2007b, MathWorks). Tungsten electrodes were advanced using a motorized microdrive as previously described. Neuronal signals were amplified and filtered (Neuralynx; lowfrequency cutoff 300 Hz, high-frequency cutoff 5 kHz), recorded at 25 kHz by a dedicated processor (Power 1401, CED) and saved to disk for off-line processing. Recordings were located in the same region examined in previous studies, tentatively identified as area 13m. Clustering procedures and preliminary data analyses were all as previously described. All experimental procedures conformed to the National Institutes of Health guidelines and with the regulations of the Washington University School of Medicine.

The behavioral task was identical to that of Exp. 1 (Fig. 1*a*). However, each session included two blocks of trials, with small/large value range. For both *offer value* and *chosen value* responses, large value ranges were approximately twice as large as small value ranges. Each trial block included 200–300 trials. Data from the two trial blocks were first analyzed separately. For each neuron, time window, and trial block, we determined whether the activity was modulated by the trial type (ANOVA, p < 0.01) and, if so, whether it was explained by any variable (*offer value A, offer value B, chosen value, taste*; regression slope different from zero, p < 0.05). Each linear regression provided an R^2 and we conventionally set $R^2 = 0$ for variables that did not explain the response. For each neuron, we then identified the encoded variable as the one that best explained responses across time windows and across



Figure 2. Slope distributions. a - c, Distribution of signed slopes. The histograms represent the number of responses (*y* axis) recorded with a given regression slope (*x* axis), separately for the three encoded variables. For each variable, the two distributions obtained for the subsets of responses with positive and negative slopes are statistically indistinguishable (all p > 0.15, Wilcoxon test). d - f, Distribution of rectified slopes. The distribution obtained for *offer value A* responses is significantly broader than that obtained for *offer value B* responses ($p < 10^{-10}$, Wilcoxon test). For each variable, the median of the distribution is indicated in the corresponding panel. sp, Spikes.

trial blocks (highest total R^2). To compare across trial blocks, we imposed that neuronal responses encode the same variable in both trial blocks. These criteria identified 19 *offer value* responses and 17 *chosen value* responses.

Results

Distribution of regression slopes and model of neuronal adaptation

Our main dataset included 937 responses encoding the *offer value* (Fig. 1*b*,*c*) and 817 responses encoding the *chosen value* (Fig. 1*d*). We previously found that the encoding of value is linear in our experiments (Padoa-Schioppa and Assad, 2006). In other words, there is a linear relationship between the firing rate ϕ of a neuron encoding the *offer value* or the *chosen value* and the encoded value *V*,

$$\phi = c_0 + c_1 V, \tag{1}$$

where parameters c_0 and c_1 represent, respectively, the intercept and the slope of the encoding.

To study how the representation of value depends on the behavioral condition, we first examined the regression slope. We considered separately neuronal responses encoding *offer value A*, *offer value B*, and *chosen value*. For each response we computed the numerical value of the regression slope, and we plotted the results obtained for the entire population in a histogram (Fig.



Figure 3. Model of range adaptation. *a*, Firing rate. The model describes the activity of a value-encoding neuron recorded in different behavioral conditions. The generic variable *value of X* corresponds either to the variable *offer value X* (if X is one particular juice) or to the variable *chosen value* (if X is the chosen juice). The relationship between the neuronal firing rate (*y* axis) and the value variable *(x* axis) is assumed to be linear. Different behavioral conditions are characterized by different value ranges. The fundamental assumption of the model is that the neuronal activity range remains constant and independent of the value range. In other words, the relationship between the neuronal firing rate and the encoded value depends on the value range: the neuron adapts in such a way that its activity range describes, in any behavioral condition, the entire range of available values. *b*, *c*, Slope. According to the model, the slope of the encoding (corresponding to the neuronal sensitivity) is proportional to the inverse value range.



Figure 4. Regression slope versus value range. a-c, Slope distributions. The three panels refer to neuronal responses encoding offer value A (a), offer value B (b), and chosen value (c). In each panel, the regression slope (y axis) is plotted against the value range (x axis) and each black cross represents one neuronal response. Value ranges are measured in conventional units. Color diamonds represent the mean slopes computed for the corresponding subpopulation. In the case of chosen value responses, means are computed binning value ranges (see Materials and Methods). For all encoded variables, the mean regression slope decreases as a function of the value range. d, e, Mean slopes. Mean slopes are plotted against the value range (ΔV , d) and against the inverse value range ($1/\Delta V$, e). The three colors represent the three encoded variables. The relationship slope $\propto 1/\Delta V$ can be observed in both panels and is confirmed by a linear fit (see Results).

2a-c). This analysis revealed several points. First, for each encoded variable, neuronal responses can have either positive or negative slopes. Second, for each variable, the two distributions obtained for the subsets of responses with positive slopes and

rectified negative slopes were statistically indistinguishable (all p > 0.15, Wilcoxon test). We thus combined for each variable the two subsets of responses by computing the slope's absolute value (Fig. 2d-f). Interestingly, the distributions obtained for the three variables were different. In particular, the slope distribution for offer value A responses was significantly broader than that for offer value B responses ($p < 10^{-10}$, Wilcoxon test). This result may seem somewhat puzzling. In principle, neurons encoding offer value A and offer value B would seem to be "of the same type" as their responses encode the value of one of the juices available in the session. Labels A and B indicate only which one of the two juices is preferred, and it is not clear why this should affect the slope of the encoding. How can we explain the difference in slope distribution?

One possible explanation follows from the observation that, in the experiments, the range of juice B offered to the monkeys was typically larger than the range of juice A. Indeed, in all the sessions, the minimum quantity offered for each juice was always 0 drops, whereas the maximum quantity varied. Since monkeys preferred juice A to juice B, we generally offered larger quantities of juice B to ensure that the animals would choose juice B for some offer types. For example, in one representative session (Fig. 1b), the quantity of juice A offered varied between 0 drops and 3 drops (value range = 3 units), while the quantity of juice B offered varied between 0 drops and 6 drops (value range = 6units). (also see Fig. 1*c*,*d*).

To illustrate how the difference in value range may relate to the difference in regression slope, we developed a qualitative model of neuronal adaptation (Fig. 3a). The model describes, for example, the activity of one neuron encoding the variable offer value X recorded in different sessions, in which juice X is offered in different value ranges. The fundamental assumption of the model is that the activity range of the neuron remains unchanged in different sessions. In any given session, the neuron adapts to the range of values of X offered to the monkey in such a way that the activity range of the neuron $(\Delta \phi, y \text{ axis})$ corresponds to the value range (ΔV , x axis) offered in that session. The key prediction of this adaptation model is that the slope of the encoding

should be proportional to the inverse value range (Fig. 3*b*,*c*). In other words, according to this model, the following relationship should hold true: slope $\propto 1/\Delta V$ (where \propto stands for "is proportional to"). In principle, the same model of neuronal adaptation

might also describe the activity of neurons encoding the *chosen value*, leading to the same predicted relationship between slope and value range.

Regression slope and inverse value range

Each neuron in our main dataset was recorded in one session with one value range. To test the predictions of the adaptation model, we thus performed a population analysis. We studied the distribution of regression slopes as a function of the value range (Fig. 4a-c). For each of the three encoded variables (offer value A, offer *value B*, and *chosen value*), we observe that regression slopes decrease as a function of the value range. This phenomenon is particularly clear when the mean slope is plotted against the value range (Fig. 4d). As predicted by the adaptation model, the relationship between the mean slope and the inverse value range appears approximately linear (Fig. 4e). For a statistical test, we performed a linear fit of the mean slope onto the inverse value range as follows: $slope = a_0 + a_1 \cdot 1/\Delta V$. The result obtained for the coefficient a_0 does not differ from zero (95% confidence interval). In other words, the measured relationship between the mean slope and the inverse value range is statistically indistinguishable from the predicted relationship *slope* $\propto 1/\Delta V$. Interestingly, this result does not depend on the particular units used to express values (see below, Scale invariance).

Figure 4, *d* and *e*, suggests that *offer value A* and *offer value B* responses are indeed of the same type and that the difference in slope distribution (Fig. 2) was due to the difference in value range. To confirm this point, we restricted the comparison of slope distributions to responses recorded with the same value range. In

this case, we did not find any significant difference between *offer* value A and offer value B (analysis performed for $\Delta V = 3$ and $\Delta V = 4$, both p > 0.5, Wilcoxon test). We thus pooled all offer value responses in subsequent analyses.

Population firing rate

A close match between the activity of neurons in the OFC and the predictions of the adaptation model can also be observed examining directly the firing rate of individual neurons as a function of the encoded value. We first describe the results obtained for the population of 937 responses encoding the *offer value* (Fig. 5). For each response, we subtracted the baseline activity corresponding to the minimum value available in that session. We then rectified neuronal responses with negative regression slope and plotted the resulting firing rate as a function of the encoded value (Fig. 5*a*). Different colors in the figure label subpopulations of neuronal responses recorded with different value ranges. A qualitative inspection suggests that the various subpopulations have similar



Figure 5. Population firing rate for *offer value* responses. *a*, Individual responses. The entire population of 937 neuronal responses encoding the *offer value* is shown (*offer value A* and *offer value B* responses are combined). Neuronal responses were baseline-subtracted, rectified and plotted here (*y* axis) against the *offer value B* responses are combined). Neuronal responses were baseline-subtracted, rectified and plotted here (*y* axis) against the *offer value B* responses are combined). Neuronal responses were baseline-subtracted, rectified and plotted here (*y* axis) against the *offer value (x* axis). Different colors highlight different value ranges. Qualitatively, we observe that for each value range neuronal activities are broadly distributed. However, the distributions recorded for different value ranges appear rather similar. *b*, Average neuronal responses. Each line represents the average neuronal response obtained for given value range (see color legend). Neuronal adaptation can be observed for any value, as average neuronal responses recorded with different value ranges are well separated throughout the value spectrum (e.g., compare the activity recorded at V = 2 for various ranges ΔV). *c*, Distribution of activity ranges. Each histogram illustrates the distribution of activity ranges obtained for the subpopulation of responses recorded with the corresponding value range (color codes as in *b*). Small triangles indicate the medians. Several statistical tests failed to find any significant correlation between activity range and value range (p > 0.13, Kruskal–Wallis test; p > 0.7, correlation analysis). Analyses in *b* and *c* were performed only for subpopulations of at least 40 responses. sp, Spikes.

distribution of maximum firing rates. This point is especially clear when we average firing rates separately for each subpopulation (Fig. 5b). The emerging picture well matches that of the adaptation model (Fig. 3a).

To further test the relationship between the neuronal firing rate and the range of values available in any behavioral condition, we defined for each neuronal response the activity range $\Delta\phi$ as the product between the regression slope and the value range. Thus, Equation 1 can be rewritten as follows:

$$\phi = \phi_0 + \Delta \phi \cdot \frac{V - V_0}{\Delta V}, \qquad (2)$$

where $\Delta \phi = c_1 \cdot \Delta V$ is the activity range, $\phi_0 = c_0 + c_1 V_0$ is the baseline activity, and V_0 is the minimum value available in that session. We then analyzed the distribution of activity ranges at the population level (Fig. 5c). Multiple statistical analyses found that the distributions of activity ranges obtained for different sub-



Figure 6. Population firing rate for *chosen value* responses. *a*, Individual responses. The entire population of 817 neuronal responses is shown. Neuronal responses were baseline subtracted and rectified, and they are plotted here (*y* axis) against the baseline-subtracted *chosen value* (*x* axis). Different colors highlight different value ranges. *b*, Average neuronal responses. We binned value ranges and divided neuronal responses in subpopulations recorded with different value ranges. Each line in the plot represents the average neuronal response obtained for a given value range (see color legend). Average neuronal responses recorded with different value ranges are separated throughout the value spectrum (e.g., compare the activity recorded at *V* = 2 for various ranges ΔV). *c*, Distribution of activity ranges. Each histogram illustrates the distribution of activity ranges obtained for the subpopulation of responses recorded with the corresponding value range (color codes as in *b*). Small triangles indicate the medians. Several tests failed to find any significant correlation between activity range and value range (p > 0.05, Kruskal–Wallis test; p > 0.4, correlation analysis). Analyses in *b* and *c* were performed only for subpopulations of at least 40 responses. sp, Spikes.

populations of responses did not depend on the value range (p > 0.13, Kruskal–Wallis test; p > 0.7, correlation analysis).

Thus, for *offer value* responses, the distribution of neuronal activity ranges remains constant and does not depend on the range of values available in the behavioral condition. One concern might be whether this result is due to saturating neuronal activity or ceiling effects. The analysis of the average firing rate (Fig. 5*b*), however, rules out this possibility. Indeed, the average firing rates recorded for different value ranges (Fig. 5*b*, indicated with different colors) are well separated throughout the value spectrum, and the measure obtained for any particular value (e.g., V = 2) varies orderly depending on the value range. This indicates that the linear relationship between regression slope and inverse value range or, equivalently, the fact that the neuronal activity range does not depend on the value range, is genuinely due to neuronal adaptation.

We repeated these same analyses for the population of 817 responses encoding the *chosen value* and we obtained very similar results (Fig. 6). A qualitative inspection of individual responses suggests that different subpopulations recorded with different value ranges have similar distributions of maximum firing rates (Fig. 6a). Computing the average firing rates for different sub-

populations (Fig. 6*b*) provides a picture resembling that of the adaptation model (Fig. 3*a*). Importantly, average firing rates are separated throughout the value spectrum. Finally, statistical tests indicate that the activity range $\Delta \phi$ is independent of the value range ΔV (p > 0.05, Kruskal–Wallis test; p > 0.4, correlation analysis) (Fig. 6*c*).

These results were obtained pooling responses from different time windows. However, analyses of data from individual time windows provided a very similar picture (supplemental Fig. S1, available at www.jneurosci.org as supplemental material). In particular, we confirmed for each time window and for each encoded variable (offer value and chosen value) that the activity range $\Delta \phi$ was independent of the value range ΔV (all p > 0.05, Kruskal– Wallis test; all p > 0.1, correlation analysis). In another control, we restricted the analysis to the neuronal population previously shown to be menu invariant (Padoa-Schioppa and Assad, 2008). The results of this analysis were very similar to those obtained for the entire population, indicating that menu invariance and range adaptation are complementary phenomena (supplemental Fig. S2, available at www.jneurosci. org as supplemental material) (see Discussion).

In conclusion, the activity of valueencoding neurons in the OFC appears to adapt to the range of values available in any behavioral condition—a result that holds true both for *offer value* neurons and for *chosen value* responses. Thus, the encoding of value in this area seems well described by Equation 2, in which the activity range ($\Delta \phi$) is independent of the

value range (ΔV). Additional recordings confirmed this result at the level of individual neurons (see Range adaptation for individual neurons).

Testing additional forms of neuronal adaptation

One important question is whether the encoding of value in the OFC presents other forms of neuronal adaptation. Referring to Equation 2, the question is whether the baseline activity (ϕ_0) depends on the value range and whether the baseline activity and/or the activity range ($\Delta \phi$) depend on other variables that describe the behavioral condition. We examined several hypotheses.

First, we tested whether the activity range of neurons encoding the *offer value* of one particular juice depended on whether that juice was preferred or nonpreferred. In our experiments, any given juice (e.g., cranberry juice) could be the preferred juice in some sessions (e.g., when it was offered against peppermint tea) and the nonpreferred juice in other sessions (e.g., when it was offered against grape juice). We thus analyzed the distribution of activity ranges recorded for neurons encoding the *offer value* of cranberry juice in sessions in which cranberry juice was preferred (juice A) or nonpreferred (juice B). As illustrated (Fig. 7*a*), the two distributions were very similar and statistically indistinguishable (p > 0.2, Wilcoxon test). We conducted this same analysis for the six different juices for which we had at least five responses for each preference status and we obtained consistent results. The distribution of activity ranges recorded for *offer value* neurons did not depend on whether the encoded juice was preferred or nonpreferred (all p > 0.05, Wilcoxon test).

Second, we tested whether the activity range of value-encoding neurons depended on the relative value the monkey assigned to the juices at the time of recording. Given two juices (e.g., apple juice and peppermint tea), preferences were usually consistent across sessions (monkeys preferred apple juice). However, the relative value of the two juices could vary from session to session depending, for example, on the thirst of the animal. For this analysis we thus considered all the neuronal responses encoding the *chosen value* recorded in sessions in which apple juice was offered against peppermint tea (Fig.



Figure 7. Invariance of the neuronal activity range. *a*, The activity range does not depend on the juice preference. Of the 100 neuronal responses encoding the *offer value* of cranberry juice in our dataset, 26 responses were recorded in sessions in which cranberry juice was preferred (juice A) and 74 responses were recorded in sessions where cranberry juice was nonpreferred (juice B). These two groups of responses are plotted separately here (black, preferred; gray, nonpreferred). The histogram shows the percentage of responses (*y* axis) recorded with different activity ranges (*x* axis). A comparison revealed that the distributions measured for the two groups of responses are very similar (p > 0.2, Wilcoxon test). *b*, The activity range does not depend on the value. Our dataset includes 47 *chosen value* responses recorded in session to session. In the scatter plot, each circle represents one neuronal response, and the activity range (*y* axis) is plotted against the relative value of the two juices (*x* axis). No systematic relation between the two measures was found (95% confidence interval, linear fit). sp, Spikes.

7*b*). A linear fit of the activity range onto the relative value did not reveal any significant correlation between the two variables (95% confidence interval). We repeated this analysis across the population as follows. For *chosen value* responses, we considered in each instance the set of neuronal responses recorded with one particular pair of juices. For *offer value* responses, we considered in each instance the set of responses encoding the value of one particular juice recorded in sessions in which that juice was offered against another specific juice. We obtained a total of 32 instances (we restricted this analysis to instances with at least 20 data points). The results were rather consistent. In 29 of 32 instances (91%) a linear fit failed to reveal any significant correlation between the activity range and the relative value (95% confidence interval).

Third, we examined whether the baseline activity (ϕ_0 in Eq. 2) is also subject to adaptation. We performed a series of analyses similar to those described above for the activity range. In particular, we investigated whether the baseline activity depends on the value range, on the juice preference, or on the relative value. In summary, we did not find any evidence of adaptation. These results are detailed in supplemental Figures S3–S5, available at www.jneurosci.org as supplemental material.

Finally, we examined whether either the activity range or the baseline activity depend on the "other" value range (e.g., whether the activity of neurons encoding the *offer value A* depends on the range of values of B). We did not find any systematic relationship (supplemental Fig. S6, available at www.jneurosci.org as supplemental material).

Scale invariance

One interesting implication of Equation 2 is that the encoding of value in the OFC is "scale invariant." In other words, the relationship between the neuronal firing rate ϕ and the value V should not depend on the particular units used to measure value (if the same units are used for V and ΔV). Consequently, the results presented above should also not depend on the value units. We tested this specific prediction as follows. Usually, we express val-



Figure 8. Scale invariance. For each response in the main dataset, we computed the inverse value range and the regression slope expressing values either in uA or in uB. We then binned the inverse value range (bin size 0.05–0.2, see Materials and Methods). In the scatter plot, mean regression slopes (*y* axis) are plotted against the inverse value range (*x* axis), separately for different encoded variables (*offer value* or *chosen value*) and different value units (uA or uB, see color legend). Thus, each neuronal response contributes to this plot twice (once for each value unit). Regression lines are obtained from an analysis of covariance of the mean regression slope using the inverse value range as predictor and the dividing data in four groups (2 encoded variables \times 2 value units). Computing the full statistical model amounts to testing whether the regression lines differ significantly from each other for their intercepts (main factor group) and/or for their slopes (interaction). The results did not find any such significant effect (both p > 0.1).

ues in conventional units uV corresponding to uA for *offer value A*, uB for *offer value B*, and uB for *chosen value*. However, the behavioral choice pattern establishes in each session a relationship between uA and uB—the relative value equals uA/uB—such that any quantity of either juice can in fact be measured in either units. Thus, we can test, for example, whether the relationship



Figure 9. Range adaptation for individual neurons. *a*, One response. The figure illustrates the activity of one *chosen value* response recorded with a small value range (left, $\Delta V = 4.3$ uB) and with a large value range (center, $\Delta V = 8.1$ uB). In both panels, filled black circles represent the behavioral choice pattern and empty color symbols represent the firing rate (all conventions as in Fig. 1*b*). Right, The same neuronal response is plotted against the *chosen value* separately for the two trial blocks. The emerging picture closely resembles that of the adaptation model (Fig. 3). *b*, Change in regression slope. In the scatter plot, slopes recorded with small ΔV (*x* axis) are plotted against slopes recorded with large ΔV (*y* axis). Each circle represents one neuronal response, and different colors identify *offer value* and *chosen value* responses (see legend). A square represents one outlier (*y* = 3.58). Consistent with adaptation, the vast majority of responses lie below the identity line. *c*, Change in activity range. Axes represent the activity ranges recorded with small ΔV (*x* axis) and large ΔV (*y* axis). Each circle represents one outlier colors identify *offer value* and *chosen value* represents one outlier (*y* = 21.45). Activity ranges were approximately stable across trial blocks. sp, Spikes.

between the regression slope and the value range (*slope* $\propto 1/\Delta V$) is indeed independent of the value units.

For *offer value* and *chosen value* responses, we plotted the mean slope against the inverse value range, expressing all values in uA or, separately, all values in uB (Fig. 8). The two sets of data points lie essentially on the same line. Indeed, an analysis of covariance did not find any significant difference between the four fitted lines (p > 0.1). Note that each data point in Fig. 8 represents the mean regression slope (y axis) measured for the subpopulation of responses recorded with a given value range (x axis). By changing value units, we essentially reshuffled responses on the x axis. That the relationship between mean slope and value range was unaffected by this procedure implies that the relationship $slope \propto 1/\Delta V$ indeed holds true, and that the results obtained in conventional units (Fig. 4*e*) were not accidental.

In addition to providing a useful control, scale invariance is also interesting from a theoretical point of view. It implies that there is no intrinsic unit for economic value at the neuronal level, at least in the OFC.

Range adaptation for individual neurons

The results presented in previous sections were based on population analyses. To confirm that individual cells undergo range adaptation, we collected an additional 19 *offer value* responses and 17 chosen value responses in sessions divided in two trial blocks, with small and large value ranges (small ΔV block and large ΔV block). The activity of one neuron encoding the chosen value is illustrated (Fig. 9a). Consistent with adaptation, the slope of the encoding was shallower in the large ΔV block compared with the small ΔV block. This was typically the case for these neurons. Indeed, for 34 of 36 responses the regression slope was shallower in the large ΔV block compared with the small ΔV block, a very unlikely result if regression slopes varied randomly (Fig. 9b) $(p < 10^{-9}, binomial)$ test). In contrast, the activity range did not systematically depend on the value range (Fig. 9*c*) (p > 0.1, binomial test).

For the response illustrated in Figure 9a, it can be noted that the firing rate corresponding to chosen value 6 uB was lower in the large ΔV block than in the small ΔV block, suggesting that the difference in slope was indeed due to neuronal adaptation as opposed to ceiling effects. To verify that this was true across the population, we considered for each response the firing rate recorded for the maximum value available in both trial blocks. We rectified responses with negative encoding and we compared the firing rate across trial blocks. For a significant majority of responses, the firing rate was lower in the large ΔV block compared with the small ΔV block (27/36 responses; p < 0.001, binomial test). This result is consistent with adaptation but not with ceiling effects.

In conclusion, this additional dataset shows that range adaptation occurs at the

level of individual neurons, as depicted in Figure 3 and summarized in Equation 2.

Partial adaptation on the time scale of individual trials

One interesting question concerns the timing of neuronal adaptation. In our experiments, the value range was not explicitly indicated, but it remained constant for a large number of trials. Since different trial types were randomly interleaved, monkeys learned the value range by performing multiple trials at the beginning of each session. Hence, neuronal adaptation presumably took place gradually over the course of multiple trials. However, recent results from other laboratories show that if the value range is explicitly updated on a trial-by-trial basis, at least partial neuronal adaptation can be observed on a short time scale (Elliott et al., 2008; Kobayashi et al., 2009). We thus examined whether trial-by-trial adaptation also occurred in our experiments. Specifically, we examined whether, in any given trial, neuronal responses encoding the value available in that trial were partly modulated by the value available in the previous trial.

Neuronal adaptation leads to a specific prediction. Consider a neuron encoding the *chosen value* (Fig. 10*a*, red). There is a linear relationship between the firing rate and the value chosen by the monkey in the current trial. Now consider trials such that the

chosen value = 4. If the cell undergoes adaptation, then the activity in the current trial should be slightly elevated if the chosen value in the previous trial was lower than 4. Conversely, the activity in the current trial should be slightly depressed if the chosen value in the previous trial was higher than 4. To test this prediction, we divided trials in two groups, depending on whether the chosen value in the current trial (trial *n*) was higher or lower than the chosen value in the previous trial (trial n -1). Indicating with V(k) the value chosen in trial k, the two groups of trials were thus defined by V(n) > V(n-1) and V(n) <V(n - 1), respectively. Consistent with neuronal adaptation, the activity for trials V(n) > V(n-1) was slightly higher than the activity for trials V(n) < V(n-1) (Fig. 10a).

To quantify this effect, we computed for each trial type the difference in firing rate between trials V(n) > V(n-1) and trials V(n) < V(n-1), normalized by the firing rate obtained pooling all trials. We then defined δ as the average normalized difference across trial types. In essence, δ represents the percentage modulation of trial n - 1 on the activity measured on trial n. For the response illustrated (Fig. 10*a*), we obtained $\delta = 14\%$. We thus computed δ for each response in the main dataset. For responses with negative encoding, we inverted the sign of δ . Thus, for every response, trial-by-trial adaptation would result in $\delta > 0$. For both *offer value* and chosen value responses, δ varied substantially across the population (Fig. 10b). However, for both variables there was a significant majority of responses with $\delta >$ 0 (both $p < 10^{-10}$, binomial test). Moreover, mean(δ) was significantly greater than zero (both $p < 10^{-9}$, *t* test). Quantitatively, we found that mean(δ) was $\sim 6\%$ for both variables. To test the reliability of this result, we repeated the analysis for trial n - 2, trial n - 3, etc. (i.e., we divided

trials depending on whether the current value was higher/lower than the value two trials before, three trials before, etc.). As expected, the effect faded rapidly with the trial distance: mean(δ) was ~6% for trial n - 1, ~2% for trial n - 2, and indistinguishable from zero for earlier trials (Fig. 10*c*). For a control, we repeated the analysis for trial n + 1 (i.e., we divided trials depending on whether the current value was higher/lower than the value in the next trial). We did not expect any such dependence because neurons cannot adapt to future events. Indeed, mean(δ) was statistically indistinguishable from zero (Fig. 10*c*).

In conclusion, partial neuronal adaptation occurred on the time scale of individual trials, as the value available in any one trial modulated neuronal responses in the following trial. Notably, this effect is in the direction opposite to that expected if the value available in any one trial elicited long neuronal responses that outlast the trial end (Simmons and Richmond, 2008).



Figure 10. Partial adaptation on the time scale of individual trials. *a*, One response. Red symbols represent the response computed pooling all trials. Trials were then separated in two groups, with V(n) > V(n - 1) (blue symbols) and V(n) < V(n - 1) (green symbols). For each trial type, trials with V(n) = V(n - 1) were assigned to the group with fewer trials. For most trial types, the "blue firing rate" [trials V(n) > V(n - 1)] was slightly higher than the "green firing rate" [trials V(n) < V(n - 1)], consistent with neuronal adaptation. For each trial type, we computed the difference between the blue firing rate and the green firing rate, and we normalized it by the red firing rate. This normalized difference averaged across trial types defined δ . Thus, δ represents the mean percentage modulation of trial n - 1 on the activity recorded on trial n. For this particular response, $\delta = 0.14$. *b*, Population analysis, distribution of δ for trial n - 1. For both *offer value* (top) and *chosen value* (bottom) responses, the *x* axis represents δ and the *y* axis represents the number of responses. Both variables present a large variability. However, δ was greater than zero in a significant majority of cases [binomial test (bino)], and mean(δ) (\pm SEM) (y axis) is plotted against the trial number (*x* axis) separately for *offer value* and *chosen value* responses. Filled squares indicate data points statistically different from zero (binomial test, p < 0.01). Mean(δ) is $\sim 6\%$ for trial n - 1, $\sim 2\%$ for trial n - 2, and statistically indistinguishable from zero for earlier trials. As expected, mean(δ) is also indistinguishable from zero for trial n - 1, sp, Spikes.

Discussion

Values processed in different behavioral conditions can vary by orders of magnitude. This variability is analogous to that faced by the visual system, in which various mechanisms of neuronal adaptation ensure that contrast sensitivity remains high independently of the range of luminance (Laughlin, 1989; Clifford et al., 2007). Our present results show that an analogous process of neuronal adaptation takes place in the valuation system. In particular, we found that OFC neurons encode value in a linear way and that their neuronal sensitivity (i.e., the slope of the encoding) is inversely proportional to the range of values available in any given condition. Conversely, the distribution of activity ranges does not depend on the value range. Interestingly, similar phenomena of neuronal adaptation were also observed in dopamine neurons (Tobler et al., 2005) and in the anterior cingulate cortex (Sallet et al., 2007).

Menu invariance and range adaptation

The present results complement our previous findings. In another study, we showed that the encoding of value in the OFC is menu invariant. In other words, the response of a neuron encoding the value of one good (a particular juice in the case of offer value neurons, or the chosen juice in the case of chosen value neurons) does not depend on the type and value of the other good available at the same time (i.e., the nonencoded good) (Padoa-Schioppa and Assad, 2008). Here we showed that the encoding of value in the OFC is range adapting. In other words, the response of a neuron encoding the value of one good (a particular juice in the case of offer value neurons, or the chosen juice in the case of chosen value neurons) adapts to the range of values available for that good in the behavioral condition (i.e., the distribution of encoded values). In essence, both offer value and chosen value neurons adapt to the range of values they encode, while they do not depend on the goods or values that they do not encode. Menu invariance and range adaptation are thus complementary phenomena (indeed, they were observed in the same neuronal population). Together, they indicate that the encoding of value in the OFC is at one time instantiative of preference transitivity (menu invariance) and computationally efficient (range adaptation).

One important issue relates to the time scale of range adaptation. As in the study of Tremblay and Schultz (1999) (see below), in our experiments value ranges remained constant for a large number of trials. In contrast, menu invariance was demonstrated by randomly interleaving trials with different juice pairs. This might suggest that range adaptation is an intrinsically slow process, whereas menu invariance is a property of the short time scale (Padoa-Schioppa and Assad, 2008). However, our current analysis and other reports (Elliott et al., 2008; Kobayashi et al., 2009) indicate that at least partial adaptation can take place on the time scale of individual trials. In fact, menu invariance is fully consistent with range adaptation occurring on a short time scale. Consider for example a neuron encoding the offer value of apple juice. The encoding undergoes adaptation in the sense that the neuronal activity elicited by a particular quantity of apple juice depends on the range of apple juice available in that session. At the same time, the encoding is menu invariant in the sense that the neuronal activity elicited by a particular quantity of apple juice does not depend on the type and quantity of the other juice available. Thus, not only is the encoding of value both menu invariant and range adapting. Menu invariance holds true independently of the time course of range adaptation. Importantly, this understanding captures the essence of menu invariance and its fundamental link with preference transitivity (Grace, 1993; Tversky and Simonson, 1993; Shafir, 1994; Padoa-Schioppa and Assad, 2008).

The present results shed new light on the early observation of Tremblay and Schultz (1999). In their study, monkeys were delivered one of three possible juices (A, B, or C, in decreasing order of preference). Trials were blocked, and only two juices were used in each block. A subset of OFC neurons responded to juice A but not to juice B in "A:B" blocks, and to juice B but not to juice C in "B:C" blocks. Our results suggest that these might have been *chosen value* neurons whose activity encoded the value received by the monkey in any given trial, and whose encoding adapted to the range of values received within each block. In A:B blocks, any such neuron would respond with maximum firing rate when the monkey received juice B. In B:C blocks, the same neuron would respond with maximum firing rate when the monkey received juice B and with minimum firing rate when the monkey received juice C. In this view, neurons in the OFC encode value (not relative preference), but in a range-adapting way. One interesting question is whether range adaptation also explains similar observations in other brain regions (Hosokawa et al., 2007).

Open questions and perspectives

While the evidence for neuronal adaptation appears robust, important questions remain. First, the exact mathematical definition of "value range" remains to some extent undetermined. In this study, we defined the value range as the difference between the maximum value and the minimum value available in the behavioral condition. However, for offer value responses, the minimum value was always equal to zero in our experiments. Thus, we cannot rule out that offer value neurons might adapt to the maximum encoded value per se-an alternative definition of value range. This distinction was theoretically possible for chosen *value* neurons, but in practice our data did not disambiguate between the two hypotheses. Thus, more experiments are necessary to conclusively define the value range. Second, our results demonstrate that the activity range of OFC neurons represents different value ranges in different conditions, a "quantitative adaptation." In addition, it is possible that OFC neurons undergo a "qualitative adaptation" when choices are made in different behavioral settings. For example, a neuron encoding the offer value of apple juice in sessions in which apple juice is available might encode the offer value of another juice in sessions in which apple juice is not available. Given the relatively large number (10-12)of juices used in our experiments, and the potentially much larger number of goods available in nature, this putative mechanism of flexible encoding seems a credible hypothesis. Future work will investigate this intriguing issue.

An increasing body of work links the OFC to the computation of subjective value underlying economic choice. Imaging studies show that neural activity in the OFC and ventromedial prefrontal cortex correlates with value in choices between food and money (Plassmann et al., 2007; Hare et al., 2008), in risky choices (Rolls et al., 2008; Venkatraman et al., 2009), in intertemporal choices (Kable and Glimcher, 2007; Hare et al., 2009), in choices that involve gains and losses (Tom et al., 2007; De Martino et al., 2009), and in food devaluation paradigms (Valentin et al., 2007). At the single-cell level, our work shows that individual neurons in the OFC encode the offer value and the chosen value, two variables seemingly necessary for economic choice (Padoa-Schioppa and Assad, 2006). Moreover, neuronal activity in this area is affected by a variety of variables relevant to choice behavior, including food-specific satiety (Rolls et al., 1989; Pritchard et al., 2008), exerted effort, probability of juice delivery (Kennerley and Wallis, 2009; Kennerley et al., 2009), and time delays (Roesch and Olson, 2005). Single-cell recordings in rats provide corroborating evidence (Roesch et al., 2006; van Duuren et al., 2007; Kepecs et al., 2008), although there are notable differences between rodents and primates (Feierstein et al., 2006; Zald, 2006; Wise, 2008). Last but not least, numerous studies demonstrate that OFC lesions specifically impair choice behavior (Bechara et al., 1996; Pasquier and Petit, 1997; Gallagher et al., 1999; Rahman et al., 1999; Fellows and Farah, 2007; Koenigs and Tranel, 2007; Machado and Bachevalier, 2007; Baxter et al., 2009). This accumulated evidence suggests that economic choices may be based on values represented in the OFC (Fellows, 2007; Padoa-Schioppa, 2007; Wallis, 2007). The present study sheds further light onto the nature of this computation, indicating that OFC neurons represent value in an adapting and thus computationally efficient way. More generally, it supports the understanding that fundamental principles first documented in sensory and motor systems, such as neuronal adaptation, also govern neurobiological processes underlying more complex behaviors, such as economic choice.

References

- Arana FS, Parkinson JA, Hinton E, Holland AJ, Owen AM, Roberts AC (2003) Dissociable contributions of the human amygdala and orbitofrontal cortex to incentive motivation and goal selection. J Neurosci 23:9632–9638.
- Asaad WF, Eskandar EN (2008) A flexible software tool for temporallyprecise behavioral control in Matlab. J Neurosci Methods 174:245–258.
- Baxter MG, Gaffan D, Kyriazis DA, Mitchell AS (2009) Ventrolateral prefrontal cortex is required for performance of a strategy implementation task but not reinforcer devaluation effects in rhesus monkeys. Eur J Neurosci 29:2049–2059.
- Bechara A, Tranel D, Damasio H, Damasio AR (1996) Failure to respond autonomically to anticipated future outcomes following damage to prefrontal cortex. Cereb Cortex 6:215–225.
- Chaudhry AM, Parkinson JA, Hinton EC, Owen AM, Roberts AC (2009) Preference judgements involve a network of structures within frontal, cingulate and insula cortices. Eur J Neurosci 29:1047–1055.
- Clifford CW, Webster MA, Stanley GB, Stocker AA, Kohn A, Sharpee TO, Schwartz O (2007) Visual adaptation: neural, psychological and computational aspects. Vision Res 47:3125–3131.
- De Martino B, Kumaran D, Holt B, Dolan RJ (2009) The neurobiology of reference-dependent value computation. J Neurosci 29:3833–3842.
- Elliott R, Agnew Z, Deakin JF (2008) Medial orbitofrontal cortex codes relative rather than absolute value of financial rewards in humans. Eur J Neurosci 27:2213–2218.
- Feierstein CE, Quirk MC, Uchida N, Sosulski DL, Mainen ZF (2006) Representation of spatial goals in rat orbitofrontal cortex. Neuron 51:495–507.
- Fellows LK (2007) Advances in understanding ventromedial prefrontal function: the accountant joins the executive. Neurology 68:991–995.
- Fellows LK, Farah MJ (2007) The role of ventromedial prefrontal cortex in decision making: judgment under uncertainty or judgment per se? Cereb Cortex 17:2669–2674.
- Gallagher M, McMahan RW, Schoenbaum G (1999) Orbitofrontal cortex and representation of incentive value in associative learning. J Neurosci 19:6610–6614.
- Grace RC (1993) Violations of transitivity: implications for a theory of contextual choice. J Exp Anal Behav 60:185–201.
- Hare TA, O'Doherty J, Camerer CF, Schultz W, Rangel A (2008) Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. J Neurosci 28:5623–5630.
- Hare TA, Camerer CF, Rangel A (2009) Self-control in decision-making involves modulation of the vmPFC valuation system. Science 324:646–648.
- Hosokawa T, Kato K, Inoue M, Mikami A (2007) Neurons in the macaque orbitofrontal cortex code relative preference of both rewarding and aversive outcomes. Neurosci Res 57:434–445.
- Kable JW, Glimcher PW (2007) The neural correlates of subjective value during intertemporal choice. Nat Neurosci 10:1625–1633.
- Kennerley SW, Wallis JD (2009) Evaluating choices by single neurons in the frontal lobe: outcome value encoded across multiple decision variables. Eur J Neurosci 29:2061–2073.
- Kennerley SW, Dahmubed AF, Lara AH, Wallis JD (2009) Neurons in the frontal lobe encode the value of multiple decision variables. J Cogn Neurosci 21:1162–1178.
- Kepecs A, Uchida N, Zariwala HA, Mainen ZF (2008) Neural correlates, computation and behavioural impact of decision confidence. Nature 455:227–231.
- Kobayashi S, De Carvalho OP, Schultz W (2009) Adaptation of reward sensitivity in the orbitofrontal cortex. Soc Neurosci Abstr 35:784.8.
- Koenigs M, Tranel D (2007) Irrational economic decision-making after ventromedial prefrontal damage: evidence from the Ultimatum Game. J Neurosci 27:951–956.

- Laughlin SB (1989) The role of sensory adaptation in the retina. J Exp Biol 146:39–62.
- Machado CJ, Bachevalier J (2007) The effects of selective amygdala, orbital frontal cortex or hippocampal formation lesions on reward assessment in nonhuman primates. Eur J Neurosci 25:2885–2904.
- Padoa-Schioppa C (2007) Orbitofrontal cortex and the computation of economic value. Ann N Y Acad Sci 1121:232–253.
- Padoa-Schioppa C, Assad JA (2006) Neurons in orbitofrontal cortex encode economic value. Nature 441:223–226.
- Padoa-Schioppa C, Assad JA (2008) The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. Nat Neurosci 11:95–102.
- Pasquier F, Petit H (1997) Frontotemporal dementia: its rediscovery. Eur Neurol 38:1–6.
- Plassmann H, O'Doherty J, Rangel A (2007) Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. J Neurosci 27:9984–9988.
- Pritchard TC, Nedderman EN, Edwards EM, Petticoffer AC, Schwartz GJ, Scott TR (2008) Satiety-responsive neurons in the medial orbitofrontal cortex of the macaque. Behav Neurosci 122:174–182.
- Rahman S, Sahakian BJ, Hodges JR, Rogers RD, Robbins TW (1999) Specific cognitive deficits in mild frontal variant of frontotemporal dementia. Brain 122:1469–1493.
- Roesch MR, Olson CR (2005) Neuronal activity in primate orbitofrontal cortex reflects the value of time. J Neurophysiol 94:2457–2471.
- Roesch MR, Taylor AR, Schoenbaum G (2006) Encoding of timediscounted rewards in orbitofrontal cortex is independent of value representation. Neuron 51:509–520.
- Rolls ET, Sienkiewicz ZJ, Yaxley S (1989) Hunger modulates the responses to gustatory stimuli of single neurons in the caudolateral orbitofrontal cortex of the macaque monkey. Eur J Neurosci 1:53–60.
- Rolls ET, McCabe C, Redoute J (2008) Expected value, reward outcome, and temporal difference error representations in a probabilistic decision task. Cereb Cortex 18:652–663.
- Sallet J, Quilodran R, Rothé M, Vezoli J, Joseph JP, Procyk E (2007) Expectations, gains, and losses in the anterior cingulate cortex. Cogn Affect Behav Neurosci 7:327–336.
- Shafir S (1994) Intransitivity of preferences in honey bees: support for 'comparative' evaluation of foraging options. Anim Behav 48:55–67.
- Simmons JM, Richmond BJ (2008) Dynamic changes in representations of preceding and upcoming reward in monkey orbitofrontal cortex. Cereb Cortex 18:93–103.
- Tobler PN, Fiorillo CD, Schultz W (2005) Adaptive coding of reward value by dopamine neurons. Science 307:1642–1645.
- Tom SM, Fox CR, Trepel C, Poldrack RA (2007) The neural basis of loss aversion in decision-making under risk. Science 315:515–518.
- Tremblay L, Schultz W (1999) Relative reward preference in primate orbitofrontal cortex. Nature 398:704–708.
- Tversky A, Simonson I (1993) Context-dependent preferences. Manage Sci 39:117–185.
- Valentin VV, Dickinson A, O'Doherty JP (2007) Determining the neural substrates of goal-directed learning in the human brain. J Neurosci 27:4019–4026.
- van Duuren E, Escámez FA, Joosten RN, Visser R, Mulder AB, Pennartz CM (2007) Neural coding of reward magnitude in the orbitofrontal cortex of the rat during a five-odor olfactory discrimination task. Learn Mem 14:446–456.
- Venkatraman V, Payne JW, Bettman JR, Luce MF, Huettel SA (2009) Separate neural mechanisms underlie choices and strategic preferences in risky decision making, Neuron 62:593–602.
- Wallis JD (2007) Orbitofrontal cortex and its contribution to decisionmaking. Annu Rev Neurosci 30:31–56.
- Wise SP (2008) Forward frontal fields: phylogeny and fundamental function. Trends Neurosci 31:599–608.
- Zald DH (2006) The rodent orbitofrontal cortex gets time and direction. Neuron 51:395–397.