

# Neuronal Origins of Choice Variability in Economic Decisions

Camillo Padoa-Schioppa<sup>1,2,3,\*</sup>

<sup>1</sup>Department of Anatomy and Neurobiology

<sup>2</sup>Department of Economics

<sup>3</sup>Department of Biomedical Engineering

Washington University in St. Louis, St. Louis, MO 63110, USA

\*Correspondence: [camillo@wustl.edu](mailto:camillo@wustl.edu)

<http://dx.doi.org/10.1016/j.neuron.2013.09.013>

## SUMMARY

To investigate the mechanisms through which economic decisions are formed, I examined the activity of neurons in the orbitofrontal cortex while monkeys chose between different juice types. Different classes of cells encoded the value of individual offers (*offer value*), the value of the chosen option (*chosen value*), or the identity of the chosen juice (*chosen juice*). Choice variability was partly explained by the tendency to repeat choices (choice hysteresis). Surprisingly, near-indifference decisions did not reflect fluctuations in the activity of *offer value* cells. In contrast, near-indifference decisions correlated with fluctuations in the preoffer activity of *chosen juice* cells. After the offer, the activity of *chosen juice* cells reflected the decision difficulty but did not resemble a race-to-threshold. Finally, *chosen value* cells presented an “activity overshooting” closely related to the decision difficulty and possibly due to fluctuations in the relative value of the juices. This overshooting was independent of choice hysteresis.

## INTRODUCTION

In recent years, significant progress has been made in understanding the neural underpinnings of economic choices. In particular, much work has focused on the computation and representation of subjective values. Lesion studies have shown that value-based decisions are selectively disrupted after lesions to the orbitofrontal cortex (OFC) and/or the amygdala, but effectively spared after lesions to other brain regions (Buckley et al., 2009; Camille et al., 2011; Gallagher et al., 1999; Rudebeck and Murray, 2011; West et al., 2011). Neurophysiology experiments have found that neurons in the primate OFC encode the subjective value of different goods during economic decisions and integrate multiple dimensions on which goods can vary (Kennerley et al., 2009; Padoa-Schioppa and Assad, 2006; Roesch and Olson, 2005). Functional imaging in humans has consistently confirmed and extended these results (Kable and Glimcher, 2007; Levy et al., 2010; Peters and Büchel, 2009;

Plassmann et al., 2007). But in spite of these advances, fundamental questions remain open. Perhaps most pressingly, the precise mechanisms through which values are compared remain unclear. In this respect, OFC appears particularly noteworthy. In a computational sense, an economic decision is a process through which the values of different goods are compared and one good is eventually chosen. Studies in which monkeys chose between different juice types have shown that neurons in the OFC encode three variables: *offer value* (the value of individual goods, independent of the eventual choice), *chosen value* (the value of the chosen good, independent of its identity), and *chosen juice* (the identity of the chosen good, independent of its value) (Padoa-Schioppa and Assad, 2006, 2008). OFC thus appears to represent all the components of the decision process, suggesting that closer examination of activity in this area might shed light on key aspects of economic choice.

In the effort to unravel the neuronal mechanisms of economic decisions, it could be fruitful to establish an analogy between economic decisions and other behaviors frequently examined in neurophysiology (Sugrue et al., 2005). In particular, extensive research has focused on the decision process underlying the visual perception of motion (henceforth “perceptual decisions”). In a somewhat simplified account, two brain areas play a critical role. Neurons in the middle temporal (MT) area encode the direction of motion for the stimuli present in the visual scene at any given time. In contrast, neurons in the lateral intraparietal (LIP) area encode the binary result of the decision process. When stimuli are degraded such that the decision process stretches over longer periods of time, neurons in MT encode the instantaneous evidence from the visual stimuli, with no memory. In contrast, neurons in LIP encode the accumulated evidence in favor of one particular decision (Newsome, 1997; Shadlen et al., 1996). Tracing the analogy between economic and perceptual decisions, *offer value* cells in OFC may correspond to neurons in MT, whereas *chosen juice* may correspond to neurons in LIP. Indeed, the former seem to represent the main input to the decision process, whereas the latter seem to represent the binary outcome of the decision. In contrast, *chosen value* cells in OFC do not appear to have a clear counterpart in perceptual decisions.

The analogy with perceptual decisions highlights two fundamental and open issues in economic decision-making. First, extensive work on perceptual decisions has been devoted to understanding how fluctuations in the activity of different neuronal populations contribute to decisions near the

indifference point (threshold). In particular, the observation that near-indifference decisions are mildly, but significantly, correlated with activity fluctuations in area MT (Britten et al., 1996; Cohen and Newsome, 2009) has provided a critical link between this area and the perception of motion. In contrast, the neuronal origins of variability in economic choices have not yet been examined, and we do not yet understand what drives decisions near the indifference point. Second, the time necessary to reach either a perceptual or an economic decision depends on the decision difficulty (Padoa-Schioppa et al., 2006; Roitman and Shadlen, 2002; Soltani et al., 2012). Building on this notion, much research has focused on neuronal activity reflecting the formation of a perceptual decision over time. In particular, the activity of neurons in LIP was found to increase gradually during perceptual decisions, suggesting that these cells encode the evolving decision state of the animal (Roitman and Shadlen, 2002; Shadlen and Newsome, 2001). By comparison, less is known about how economic decisions form over time, or about how economic decisions depend on the decision difficulty. In addition to these empirical questions, considerable work on perceptual decisions has been devoted to mathematical conceptualization. Specifically, activity profiles in area LIP have been described with a variety of models, including race-to-threshold processes and dynamical systems (Bogacz et al., 2006; Gold and Shadlen, 2007; Wang, 2002). In contrast, although several proposals were recently put forth (Hunt et al., 2012; Krajbich et al., 2010; Soltani et al., 2012; Solway and Botvinick, 2012), a comprehensive model for the neuronal mechanisms of economic decisions remains elusive (see Discussion). To address these issues and gather elements that would inform future models, I examined data recorded in the OFC of monkeys engaged in economic choices.

## RESULTS

Neuronal activity in OFC was recorded in two experiments during which monkeys chose between different juice types (see [Experimental Procedures](#)). In experiment 1, animals chose between two juices labeled A and B, with A preferred (Padoa-Schioppa and Assad, 2006). Offers were represented by sets of colored squares on a computer monitor and the animals indicated their choices with an eye movement. Juice quantities varied from trial to trial and behavioral choice patterns typically presented a quality-quantity trade-off (Figures 1A and 1B). In experiment 2, the procedures were very similar except that three juices were used in each session (Padoa-Schioppa and Assad, 2008). Two of the three juices were offered in each trial, with the three juice pairs randomly interleaved.

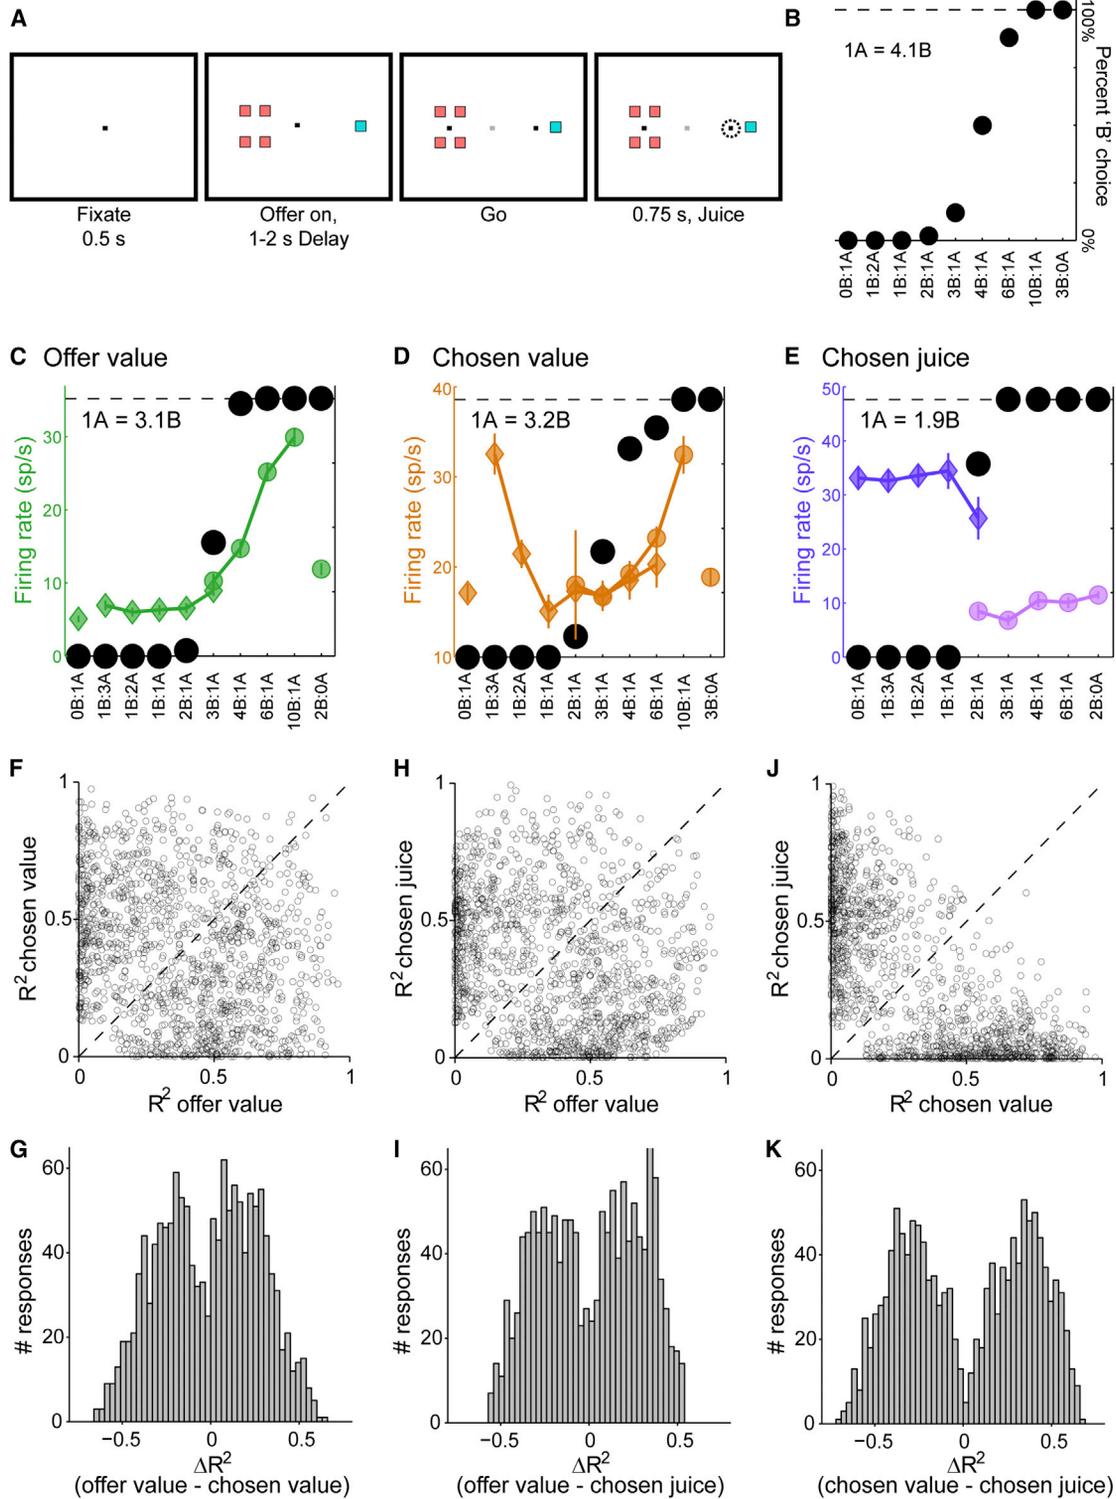
### From Neuronal Responses to Cell Classes

Previous analyses were based on neuronal responses, defined as the activity of one cell in one time window (see [Experimental Procedures](#)). It was shown that the vast majority of neuronal responses encoded one of three variables: *offer value* (Figure 1C), *chosen value* (Figure 1D), and *chosen juice* (Figure 1E). Notably, the firing rate could increase or decrease as a function of the encoded variable (positive or negative encoding). However, two important questions were not previously addressed. First,

because variables *offer value*, *chosen value* and *chosen juice* were intrinsically correlated, individual responses were often explained by more than one variable. For example, one response could have a nonzero slope when regressed onto either *offer value* or *chosen value*. In such case, the response was assigned to the variable with the highest  $R^2$ . However, this criterion did not assess whether *offer value* and *chosen value* were distinct classes of responses or, alternatively, whether the two variables represented “poles” of a continuum. Second, previous studies did not test whether *offer value*, *chosen value*, and *chosen juice* corresponded to separate groups of cells. In principle, any given neuron could encode different variables at different times. Alternatively, each cell could consistently encode a single variable. I addressed these issues as follows.

To assess whether *offer value* and *chosen value* are distinct classes of responses, I computed for each response the linear regression onto variables *offer value* and *chosen value*, from which I obtained the two  $R^2$ . I then defined  $\Delta R^2 = R^2_{\text{offer value}} - R^2_{\text{chosen value}}$ , which ranged from  $-1$  to  $+1$ . For a response perfectly explained by *offer value* (*chosen value*) and poorly explained by *chosen value* (*offer value*),  $\Delta R^2$  is close to  $+1$  ( $-1$ ). If *offer value* and *chosen value* are two poles of a continuum, the distribution of  $\Delta R^2$  should be unimodal with a peak close to zero. Conversely, if *offer value* and *chosen value* are distinct classes of responses, the distribution of  $\Delta R^2$  should be bimodal with a dip close to zero. As illustrated in Figures 1F and 1G, the distribution obtained for  $\Delta R^2$  was indeed bimodal ( $p < 0.02$ , Hartigan’s dip test). Thus, *offer value* and *chosen value* appeared to be distinct classes of responses. I repeated this analysis for the two other pairs of variables (*offer value* versus *chosen juice* and *chosen value* versus *chosen juice*). In both cases, the distribution for  $\Delta R^2$  was clearly bimodal (both  $p < 10^{-10}$ , Hartigan’s dip test; Figures 1H–1K). In conclusion, *offer value*, *chosen value*, and *chosen juice* responses are best thought of as different classes of responses, not as poles of a continuum.

To assess whether different neurons encoded different variables, I first examined data from experiment 1. Neuronal responses were classified as encoding one of four variables: *offer value A*, *offer value B*, *chosen value*, or *chosen juice*. Responses that were not task-related or that were not explained by any variable were unclassified. Given a neuron and two time windows, I defined a “classification conflict” if the neuron was classified in both time windows but it encoded different variables. A conflict was also detected if a cell encoded the same variable but with different sign. I thus sought to establish whether the incidence of classification conflicts in the population was greater, comparable, or lower than expected by chance. Chance level was estimated with a bootstrap technique (see [Supplemental Experimental Procedures](#) available online). This analysis showed that the number of classification conflicts present in the data was significantly lower than expected by chance ( $p < 10^{-10}$ , t test; Figure S1A). Conversely, for each pair of time windows, cells with consistent classification were significantly more frequent than expected by chance (Figure S1B; all  $p < 10^{-10}$ , t test). Data from experiment 2 provided very similar results both for the analysis of classification conflicts ( $p < 10^{-10}$ , t test) and for that of classification consistency (all  $p < 10^{-10}$ , t test). In other words, OFC neurons typically encoded the same variable across time windows.



**Figure 1. Categorical Encoding of Offer Value, Chosen Value, and Chosen Juice**

(A) Task design. Animals maintained center fixation and offers were represented by two sets of color squares. After a randomly variable delay, animals indicated their choice with a saccade.

(B) Typical choice pattern. The x axis represents offer types ranked by the ratio #B:#A. The y axis represents the percentage of trials in which the animal chose juice B. In this session, the animal was roughly indifferent between 1A and 4B.

(legend continued on next page)

In the light of these results, I assigned each neuron to the variable that best explained responses across all time windows (sum of  $R^2$  across time windows), taking into account the sign of the encoding. The resulting data set included 245 *offer value* cells (188/57 with positive/negative encoding), 273 *chosen value* cells (161/112 with positive/negative encoding), and 265 *chosen juice* cells. The sign of *chosen juice* cells could be assessed unequivocally only for data from experiment 2 (146 *chosen juice* cells, 96/50 with positive/negative encoding). Unless otherwise specified, all the analyses of *chosen juice* cells were performed by pooling data from the two experiments and rectifying cells with negative encoding such that the “encoded” juice elicited higher neuronal activity.

Figure 2 illustrates the average activity profile obtained for each neuronal population. Importantly, inspection of Figure 2E suggests that decisions were made within 500 ms of the offer.

### Computational Framework

Consider a session in which the animal chose between juice A and juice B. When the two offer values were sufficiently different, the animal consistently chose the same juice. However, near-indifference decisions were typically split: on some trials the animal chose juice A, in other trials it chose juice B. This phenomenon is referred to as choice variability (Figure 3A). The primary goal of this study was to shed light on the neuronal origins of choice variability.

The analyses presented here were guided by the computational framework depicted in Figure 3B (see also Padoa-Schioppa, 2011). At the outset of this study, I conceptualized the decision between two goods as a process in which two offer values are compared on the basis of a relative value. The decision outcome is represented by the identity and value of the chosen good. The three populations of neurons found in the OFC appear to match this scheme. Indeed, *offer value* cells encode the value of individual offers, whereas *chosen value* and *chosen juice* cells encode, respectively, the value and identity of the chosen good. This observation led to the working hypothesis that motivated this study, namely that each class of cells in the OFC may be identified with the corresponding computation.

### Choice Hysteresis

Choice patterns in the experiments were generally saturated, indicating that the animals had strict preferences. However, when the two offers had similar values, monkeys were more likely to choose the same juice that they had chosen in the previous trial. I refer to this behavioral phenomenon as “choice hys-

teresis.” One example session is illustrated in Figure 4A, where I separated trials into two groups depending on the outcome of the previous trial. The choice pattern obtained when the outcome of the previous trial was juice A (trials A•) was displaced to the right (higher indifference point) compared to the choice pattern obtained when the outcome of the previous trial was juice B (trials B•). Choice hysteresis was consistent across sessions (Figure 4B). The indifference point measured in A• trials was typically higher than that measured in B• trials ( $p < 10^{-10}$ , sign test). In some cases, the outcome of the previous trial was neither juice A nor juice B. These trials (X• trials) followed incomplete trials or, in experiment 2, trials in which the animal chose the third juice offered in the session. The indifference point measured in X• trials was typically between those obtained for A• trials and B• trials. Importantly, choice hysteresis largely dissipated after one trial (Figure 4C).

To quantify choice hysteresis more precisely, I used a logistic analysis. I constructed the following logistic model:

$$\text{choice B} = 1 / (1 + e^{-X}) \quad (\text{Equation 1})$$

$$X = a_0 + a_1 \log(\#B / \#A) + a_2(\delta_{n-1,B} - \delta_{n-1,A}).$$

The variable choice B was equal to 1 if the animal chose juice B and 0 otherwise. #A and #B were, respectively, the quantities of juices A and B offered to the animal in any given trial. The current trial was referred to as trial  $n$  and the variable  $\delta_{n-1,j}$  was equal to 1 if in the previous trial the animal received juice  $j$  and 0 otherwise. Note that the difference  $(\delta_{n-1,B} - \delta_{n-1,A})$  was equal to 1,  $-1$ , or 0 depending on whether the previous trial ended with receipt of juice B, juice A, or otherwise (e.g., with receipt of the third juice in experiment 2). The logistic regression provided an estimate for parameters  $a_0$ ,  $a_1$ , and  $a_2$ . By construction,  $a_1 > 0$ . In the simplified model with  $a_2 = 0$ ,  $a_1$  was the inverse temperature and a measure of choice variability, whereas the indifference point was provided by  $\exp(-a_0/a_1)$ . Choice hysteresis corresponded to  $a_2 > 0$ . However, it was useful to quantify the effect of choice hysteresis with the normalized coefficient  $a_2/a_1$ . This logistic regression was performed for each session in the data set (304 sessions total). I thus obtained a distribution for  $a_2/a_1$  across sessions (Figure 4D). The median of the distribution  $m = 0.124$  was significantly  $>0$  ( $p < 10^{-10}$ , Wilcoxon sign test). Behaviorally, this means that the effect of obtaining juice B in the previous trial was equivalent to multiplying the quantity of juice B by a factor of  $\exp(m) \approx 1.13$ .

In subsequent analyses, I examined the contributions of different neuronal populations to choice variability while controlling for choice hysteresis.

(C) Response encoding the *offer value*. Black symbols represent the behavioral choice pattern and green symbols represent the firing rate recorded in the 500 ms after the offer. Each data point represents one trial type and diamonds and circles represent, respectively, trials in which the animal chose juice A and juice B. Error bars represent SE.

(D) Response encoding the *chosen value*. Color symbols represent the firing rate recorded in the 500 ms after the offer.

(E) Response encoding the *chosen juice*. Color symbols represent the firing rate recorded in the 500 ms before juice delivery.

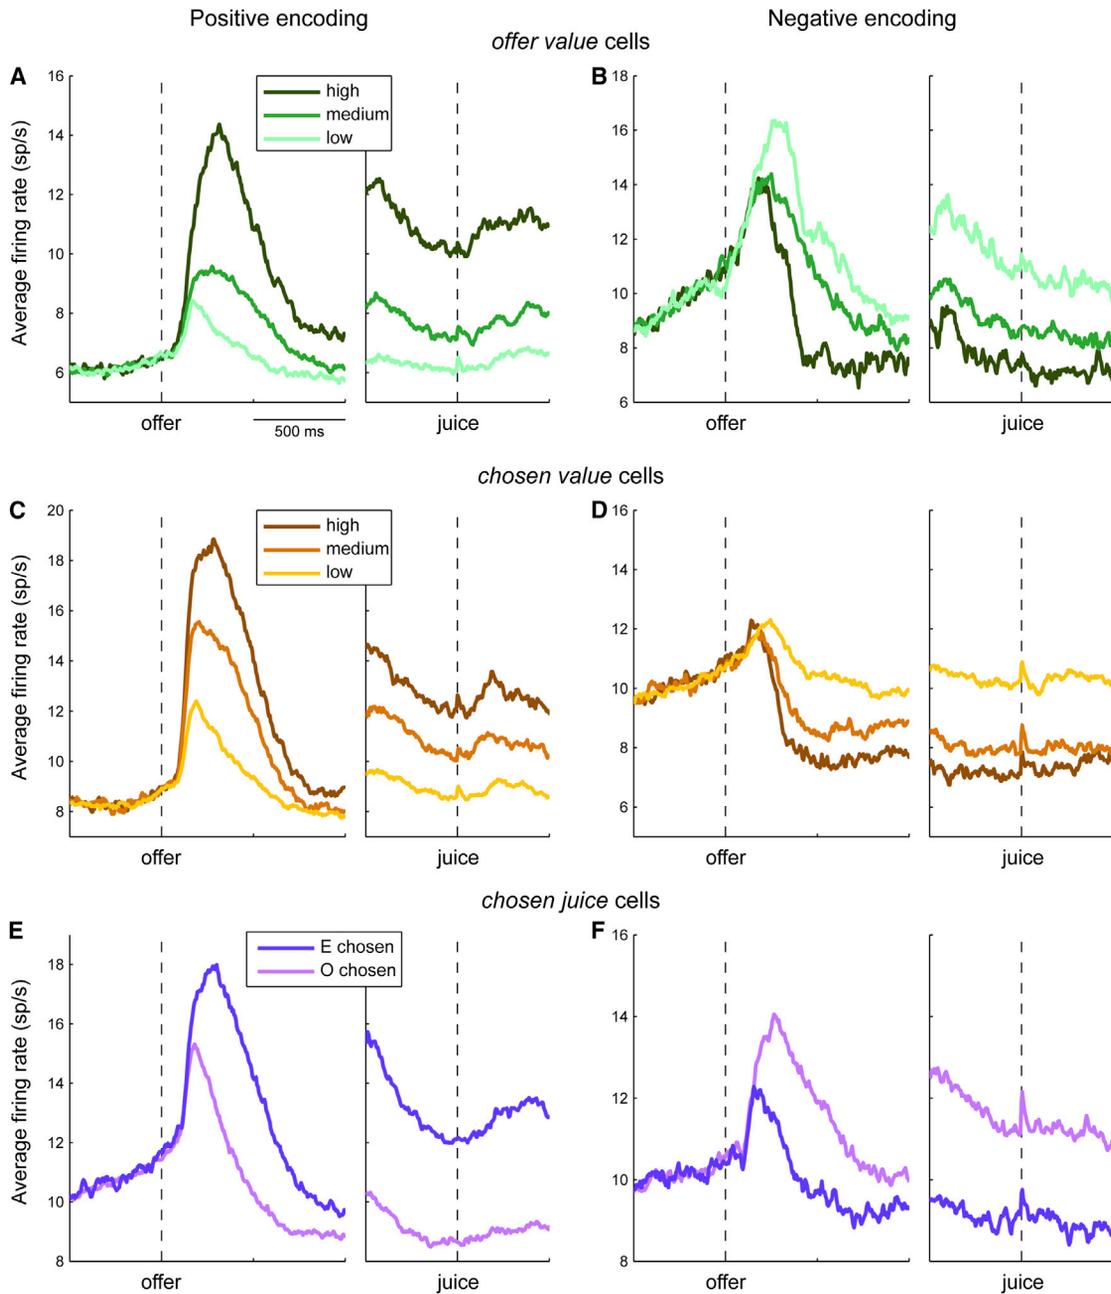
(F and G) Categorical encoding of *offer value* versus *chosen value*. The scatter plot and the histogram include all the responses classified as encoding either the *offer value* or the *chosen value*. For each response, I considered the two  $R^2$ s obtained from linear regressions against variables *offer value* and *chosen value*.

(F) The two  $R^2$ s are plotted against each other. (G) Illustrates the distribution obtained for  $\Delta R^2$ .

(H and I) *Offer value* versus *chosen juice*.

(J and K) *Chosen value* versus *chosen juice*.

See also Figure S1.



**Figure 2. Average Activity Profiles**

(A and B) Average population activity for *offer value* cells. For each cell, trials were divided into three tertiles based on the value of the encoded juice (high, medium, and low). The activity of each tertile was averaged across the population. (A) and (B) show the activity for neurons with positive/negative encoding (188/57 cells).

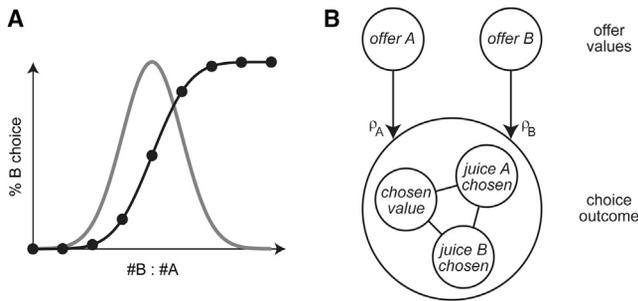
(C and D) Average population activity for *chosen value* cells. For each cell, trials were divided into three tertiles based on the chosen value (high, medium, and low). (C) and (D) show the activity for neurons with positive/negative encoding (161/112 cells).

(E and F) Average population activity for *chosen juice* cells. The figure includes only data from experiment 2, for which positive/negative encoding could be established (see main text). For each chosen juice cell, trials were divided depending on whether the animal chose the juice encoded by the cell (E chosen) or the other juice (O chosen). (E) and (F) show the average activity for neurons with positive/negative encoding (96/50 cells).

**Fluctuations in the Activity of Offer Value Cells Do Not Explain Choice Variability**

In the framework of Figure 3B, *offer value* cells (Figure 1C) represent the primary input to the decision process. Intuitively and by

analogy with results in perceptual decisions, it is reasonable to wonder whether choice variability reflects fluctuations in the activity of *offer value* cells. To examine this issue, I analyzed all *offer value* cells focusing on offer types for which decisions were split.



**Figure 3. Computational Framework**

(A) Choice variability (cartoon). Consider a session in which the animal chose between juice A and juice B. The experimental design and the analysis assumed that, for each offer type, the percent of B choices (black dots) depended only on the ratio #B:#A. Choice variability corresponds to the fact that the normal distribution derived from the sigmoid fit has nonzero variance. The mean of the distribution is the indifference point.

(B) Computational framework. The decision model proposed here assumes that there is an input layer represented by *offer value* cells. The input feeds into a circuit that includes *chosen juice* neurons and *chosen value* neurons, which collectively represent the choice outcome. This computational framework does not specify the architecture of the network, and is thus compatible with a variety of possible architectures (Bogacz et al., 2006). The relative value between the two goods ( $\rho$ ) can generally be thought of as a ratio of synaptic efficacies. For example, in a mutual-inhibition model or in a pooled-inhibition model, the input units (*offer value* cells) feed into response units (*chosen juice* cells) with synaptic efficacies  $\rho_A$  and  $\rho_B$ . In this scenario and under reasonable assumptions, the relative value  $\rho$  equals the ratio  $\rho_A/\rho_B$ . Alternatively, in a pooled-inhibition model,  $\rho \neq 1$  could also emerge from an imbalance of the other synaptic efficacies defined in the network. In principle, synaptic efficacies (and their ratios) can fluctuate stochastically on a trial-by-trial basis.

Thus, in some trials the animal chose the juice encoded by the neuron under consideration (juice E), whereas in other trials the animal chose the other juice (juice O). For each offer type, the firing rate was averaged separately for the two groups of trials. The resulting traces were averaged across offer types to obtain two traces for each *offer value* cell: one for trials in which the animal chose the encoded juice (E chosen) and another for trials in which the animal chose the other juice (O chosen). These traces were baseline-subtracted and averaged across neurons. As illustrated in Figure 5A (positive encoding), the resulting population traces appeared indistinguishable throughout the 1 s following the offer. A receiver operating characteristic (ROC) (see Experimental Procedures) analysis focused on the 150–400 ms after the offer confirmed this impression. Specifically, the area under the curve (AUC; also referred to as choice probability) did not consistently differ from the null hypothesis of 0.5 (mean AUC = 0.504;  $p = 0.6$ , t test). Thus there was no evidence that the activity of *offer value* cells was elevated on trials in which the animal chose the juice they encoded. Similar results were obtained for negative encoding cells (Figure 5B) and in several different variants of this analysis (Supplemental Experimental Procedures; Figure S2).

To further test the possible relationship between fluctuations in the activity of *offer value* cells and near-indifference decisions, I ran a logistic analysis using an approach similar to that of Yang and Shadlen (2007). This analysis focused on the 500 ms following the offer. I constructed the following logistic model:

$$\text{choice } E = 1/(1 + e^{-X})$$

$$X = a_0 + a_1 \log(\#E/\#O) + a_2(\delta_{n-1,E} - \delta_{n-1,O}) + a_3 \varphi_{\text{residual}}$$

(Equation 2)

For each *offer value* cell, E was the juice encoded by the cell, O was the other juice, and  $\varphi_{\text{residual}}$  was the residual firing rate remaining after the linear regression of the raw firing rate ( $\varphi$ ) onto the variable *offer value* E. Other notations were as in Equation 1. The null hypothesis corresponded to  $a_3/a_1 = 0$ . The logistic regression was performed for each *offer value* cell in the data set (cells from experiment 2 contributed each with two data points). Figure S5B illustrates the distribution for  $a_3/a_1$  obtained across the population. In this histogram, cells with positive and negative encoding were pooled after inverting the sign of  $a_3$  for cells with negative encoding. The median of the distribution  $m = 0.001$  was in the expected direction but did not reach statistical significance ( $p = 0.12$ , Wilcoxon signed-rank test).

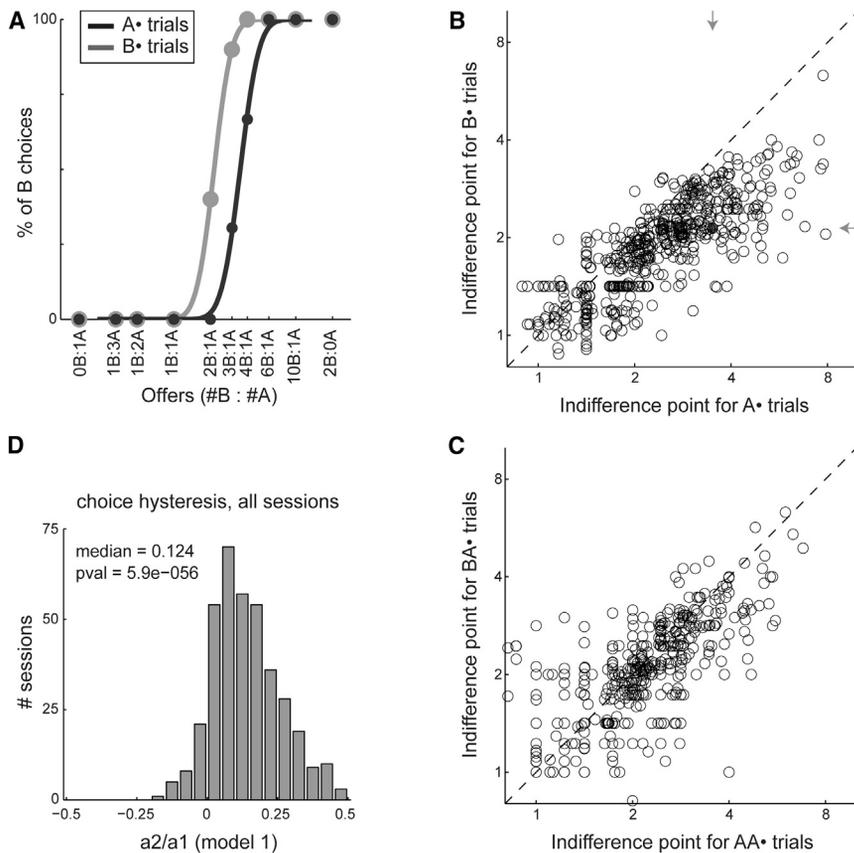
In summary, I did not find consistent evidence that near-indifference decisions correlate with stochastic fluctuations in the activity of *offer value* cells. This result is somewhat surprising and qualitatively different from observations on perceptual decisions (see Discussion).

### Chosen Juice Cells, Decision Difficulty, and Predictive Activity

I next examined *chosen juice* cells (Figure 1E). By definition, the activity of these neurons depended on the type of juice the animal chose, but not on its value. One important question was whether and how their activity depended on the decision difficulty. To address this issue, I pooled cells from the two experiments and rectified neurons such that the encoded juice was defined as that which elicited higher activity. For each cell, I divided trials depending on whether the monkey chose the encoded juice (juice E) or the other juice (juice O) and on whether decisions for that offer type were easy or split (see Experimental Procedures). I thus obtained four groups of trials: “E chosen easy,” “E chosen split,” “O chosen split,” and “O chosen easy.” For each group, I averaged the activity profiles across trials and across cells (Figure 6A). Several aspects of the results are noteworthy.

First, even though the encoding was basically binary (high or low depending on the chosen juice), the activity profile clearly depended on the decision difficulty. In particular, consider trials in which the monkey chose the encoded juice (blue lines in Figure 6A). In the time window 200–450 ms following the offer, the activity was significantly higher for easy decisions than for split decisions (ROC analysis: across the population, mean AUC = 0.556;  $p < 10^{-10}$ , t test). Conversely, for trials in which the animal chose the other juice (red lines in Figure 6A), the drop of activity in the same time window was significantly more pronounced when decisions were easy than when they were split (mean AUC = 0.477;  $p < 10^{-4}$ , t test).

Second, the activity of *chosen juice* cells did not resemble a race-to-threshold. Indeed, although the traces for easy and split decisions converged, they did so ~500 ms after the offer for E trials and, most strikingly, well into the descent phase that followed the activity peak. In this respect, there appears to be a difference between *chosen juice* cells in OFC and neurons in LIP (but see Discussion).



**Figure 4. Choice Hysteresis**

(A) Choice hysteresis in one example session. Trials were separated into two groups depending on the outcome of the previous trial. The choice pattern obtained when the previous trial was juice A (trials A\*, dark gray) was displaced to the right (higher indifference point) compared to the choice pattern obtained when the outcome of the previous trial was juice B (trials B\*, light gray).

(B) Choice hysteresis across sessions. Each data point corresponds to one juice pair in one session, and the two axes indicate the indifference point measured in A\* trials (x axis) and B\* trials (y axis). Arrows point to the session shown in (A) (gray circle).

(C) Choice hysteresis largely dissipated within one trial. The panel compares AA\* trials and BA\* trials.

(D) Logistic analysis. The x axis represents the ratio  $a_2/a_1$  defined in Equation 1, the y axis represents the number of session (304 total).

Comparing the two distributions of firing rates, I measured the AUC. Across the population, the mean AUC significantly exceeded the null hypothesis of 0.5 (mean AUC = 0.527,  $p < 10^{-6}$ , t test), indicating that predictive activity was typically present in individual *chosen juice* cells.

One possible concern was whether the activity of *chosen juice* cells was genuinely binary. Indeed, in the experiments,

Third, an interesting phenomenon can be observed in the 500 ms preceding the offer. The activity profiles recorded when decisions were easy (dark blue and dark red in Figure 6A) were essentially indistinguishable, consistent with the intuition that the animal could not have made a decision before the offer. However, the activity profiles recorded when decisions were split (light blue and light red lines in Figure 6A) seemed to defy this intuition. Indeed, the activity preceding choices of the encoded juice was clearly higher than that preceding choices of the other juice. By analogy with effects observed in other behavioral tasks (Shadlen and Newsome, 2001; Williams et al., 2003; Wyart and Tallon-Baudry, 2009), I refer to this as “predictive activity.” In the framework of Figure 3B, a possible interpretation of the predictive activity is that trial-by-trial fluctuations in the initial state of the neuronal assembly, reflected in the activity of *chosen juice* cells, contributed to the decision of the animal. In this view, when one of the two offer values clearly dominated, the initial state was irrelevant: animals always chose the dominant offer. However, near the indifference point, when there was no clearly dominant offer, relatively small fluctuations in the initial state effectively biased the decision. (More specific hypotheses are discussed below.)

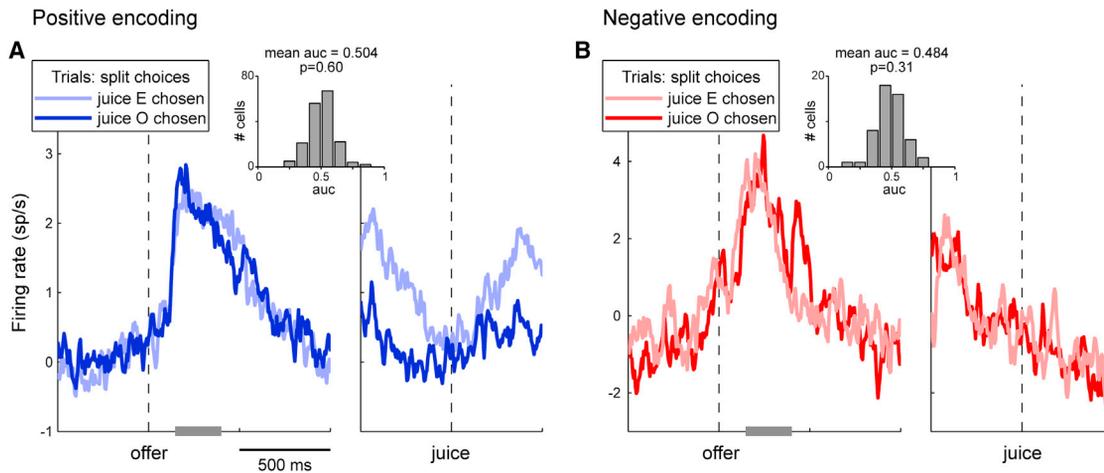
It was important to assess whether predictive activity was generally present in individual cells. To examine this issue, I performed an ROC analysis focused on the 500 ms before the offer. For each cell, I identified offer types in which decisions were split, and I divided trials into two groups depending on the chosen juice.

the indifference point typically corresponded to lower juice quantities (see Figure 1B). Thus the difference in neural activity between easy decisions (dark blue in Figure 6A) and split decisions (light blue in Figure 6A) could be explained if the activity of *chosen juice* cells depended to some extent on the chosen juice quantity. To address this issue, I isolated trials in which the animal chose one drop of the preferred juice (1A), and I identified neurons encoding the *chosen juice* A. I then divided offer types into easy and split and repeated the analysis (Figure 6C). The results confirmed those based on all the trials: (1) the activity recorded after the offer was significantly higher when decisions were easy, (2) the traces did not seem to reach a specific threshold, and (3) the activity recorded prior to the offer was elevated in split-decision trials. Note that in Figure 6C, the chosen option was identical for both traces, so differences in the activity of *chosen juice* cells cannot be explained by quantity-dependent encoding. Rather, all the differences between the two traces seem genuinely related to the decision difficulty.

The relationship between the preoffer activity of *chosen juice* cells and near-indifference decisions was also tested with a logistic analysis. I constructed the following model:

$$\begin{aligned} \text{choice } E &= 1/(1 + e^{-X}) \\ X &= a_0 + a_1 \log(\#E/\#O) + a_2 \phi. \end{aligned} \quad (\text{Equation 3})$$

For each *chosen juice* cell,  $\phi$  was the firing rate in the 500 ms preceding the offer (in sp/s). Figure S5C illustrates the



**Figure 5. Choice Variability Is Not Explained by Fluctuations of Offer Value Cells**

(A) Population with positive encoding. The analysis focused on offer types where choices were split. For each offer type, trials were divided depending on the animal's choice (juice E or juice O) and the activity was averaged separately for the two groups of trials ( $\geq 2$  trials per trace). The resulting traces were averaged across offer types for each cell and then across cells. The eventual choice of the animal does not correlate with fluctuations in the activity of *offer value* cells. Average traces shown here are from 177 cells. The gray bar highlights the time window on which the ROC analysis was conducted (150–400 ms after the offer). Inset: the histogram shows the distribution of AUC obtained across the population. The mean AUC was statistically indistinguishable from 0.5.

(B) Population with negative encoding. Same procedures as in (A). Average traces shown here are from 52 cells. Across the population, the AUC was indistinguishable from 0.5.

See also Figure S2.

distribution for  $a_2/a_1$  obtained across the population. The median of the distribution  $m = 0.005$  was significantly  $>0$  ( $p < 10^{-8}$ , Wilcoxon signed-rank test). In essence, this means that when the preoffer activity of *chosen juice* cells increased by one spike per second, the animal made its choice as though the quantity of the encoded juice was multiplied by a factor of  $\approx 1.005$ .

In summary, these results suggest that near-indifference decisions are partly driven by the initial state of the neuronal assembly, which fluctuates on a trial-by-trial basis and is reflected in the preoffer activity of *chosen juice* cells.

### Residual Predictive Activity of Chosen Juice Cells

While discussing the predictive activity, one important caveat relates to the presence of choice hysteresis. Indeed, previous work has found that reward-related activity in the OFC can outlast the trial end (Simmons and Richmond, 2008). Thus on any given trial, *chosen juice* cells might present some tail activity from the previous trial. Because of choice hysteresis, such tail activity would appear as predictive activity for hard decisions. Indeed, referring to Figure 6A, more “E chosen split” trials follow trials in which the animal chose juice E, and more “O chosen split” trials follow trials in which the animal chose juice O. To assess the relation between choice hysteresis and predictive activity, I examined whether the outcome of the previous trial affected the activity of *chosen juice* cells (Figure S4). Consistent with previous results, the activity of *chosen juice* cells early in the trial was slightly elevated after trials in which the animal chose juice E and slightly depressed after trials in which the animal chose juice O. This tail activity was in the same direction as, and thus confounded with, the predictive activity.

Importantly, the two interpretations for the predictive activity (tail activity from the previous trial or baseline fluctuation reflect-

ing a bias in the current choice) are not mutually exclusive. Indeed, predictive activity could in principle provide a neuronal mechanism for choice hysteresis. In this respect, it is interesting to assess whether predictive activity was entirely explained as tail activity from the previous trial (H0) or, alternatively, whether predictive activity also reflected additional sources of stochasticity (H1). To examine this issue, I separated trials into three groups depending on whether in the previous trial the animal chose the juice encoded by the cell (E• trials), the other juice offered (O• trials), or neither juice (X• trials). Because the outcome of the previous trial was fixed, the presence of the residual predictive activity (Figures S4B–S4E) provided evidence in favor of H1. For a quantitative assessment of residual predictive activity, I constructed the following logistic model:

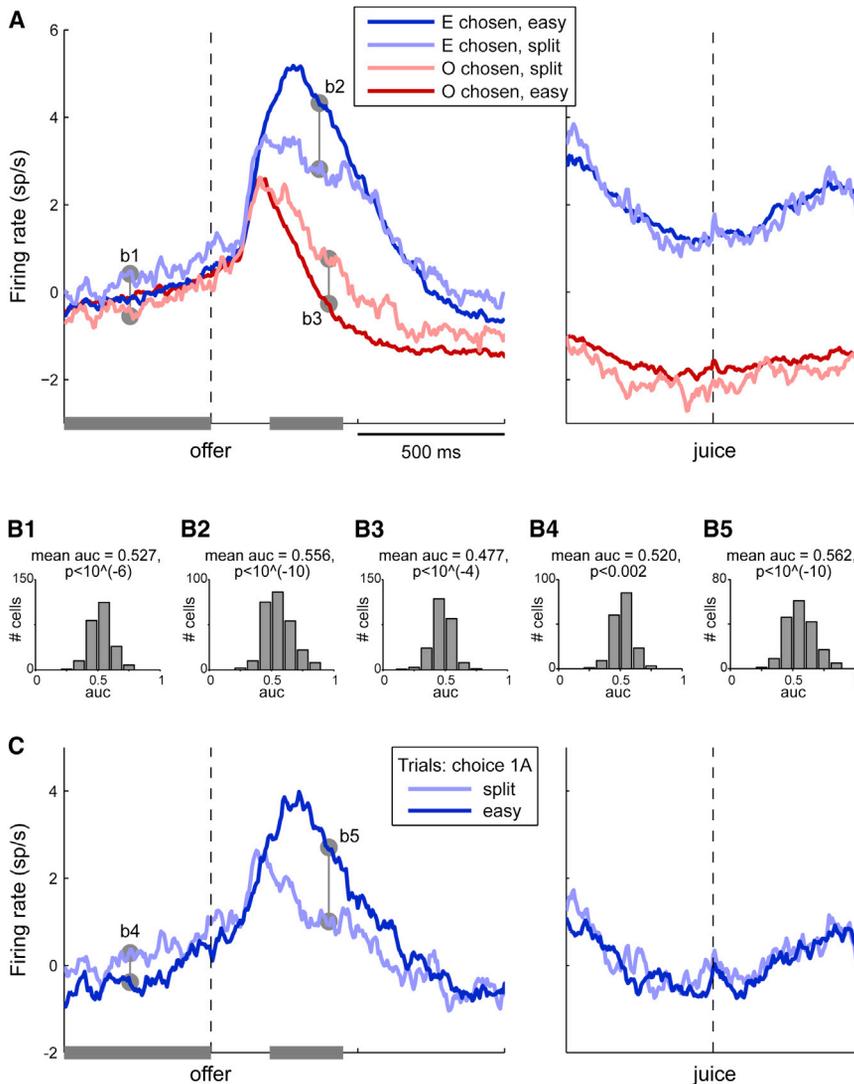
$$\text{choice E} = 1/(1 + e^{-X})$$

$$X = a_0 + a_1 \log(\#E/\#O) + a_2(\delta_{n-1,E} - \delta_{n-1,O}) + a_3 \varphi_{\text{residual}}$$

(Equation 4)

For each *chosen juice* cell,  $\varphi_{\text{residual}}$  was the residual firing rate remaining after the linear regression of the raw firing rate  $\varphi$  onto the variable  $(\delta_{n-1,E} - \delta_{n-1,O})$ . The null hypothesis corresponded to  $a_3/a_1 = 0$ . Figure S5D illustrates the distribution for  $a_3/a_1$  obtained across the population. The median of the distribution  $m = 0.002$  was small but significantly  $>0$  ( $p < 0.02$ , Wilcoxon signed-rank test). In other words, trial-by-trial fluctuations in the preoffer activity of *chosen juice* cells were significantly correlated with the decision of the animal, even when the outcome of the previous trials was controlled for.

In conclusion, predictive activity reflected additional sources of stochasticity above and beyond the tail activity from the previous trial.



**Figure 6. Activity Profiles of Chosen Juice Cells**

(A) All trials. Neurons from the two experiments were rectified (see main text) and pooled. Trials were divided depending on whether the animal chose the juice encoded by the cell (juice E) or the other juice (juice O) and on whether the decisions were easy or split. Average traces shown here are from the 257 cells for which I could compute all four traces ( $\geq 2$  trials per trace). The activity after the offer depended on the decision difficulty but did not resemble a race-to-threshold. In the 500 ms before the offer, the activity for “E chosen split” trials was elevated compared to that for “O chosen split” trials (predictive activity).

(B) ROC analyses. Histograms show the results obtained for the five comparisons indicated in (A) and (C).

(C) Control for juice quantity. This analysis focused on trials in which the animal chose one drop of the preferred juice (1A). Trials were divided into easy and split and average traces shown here are from the 181 cells for which I could compute both traces ( $\geq 2$  trials per trace). All the effects described in (A) were also observed when the quantity of the chosen juice was fixed. See also Figures S3–S5.

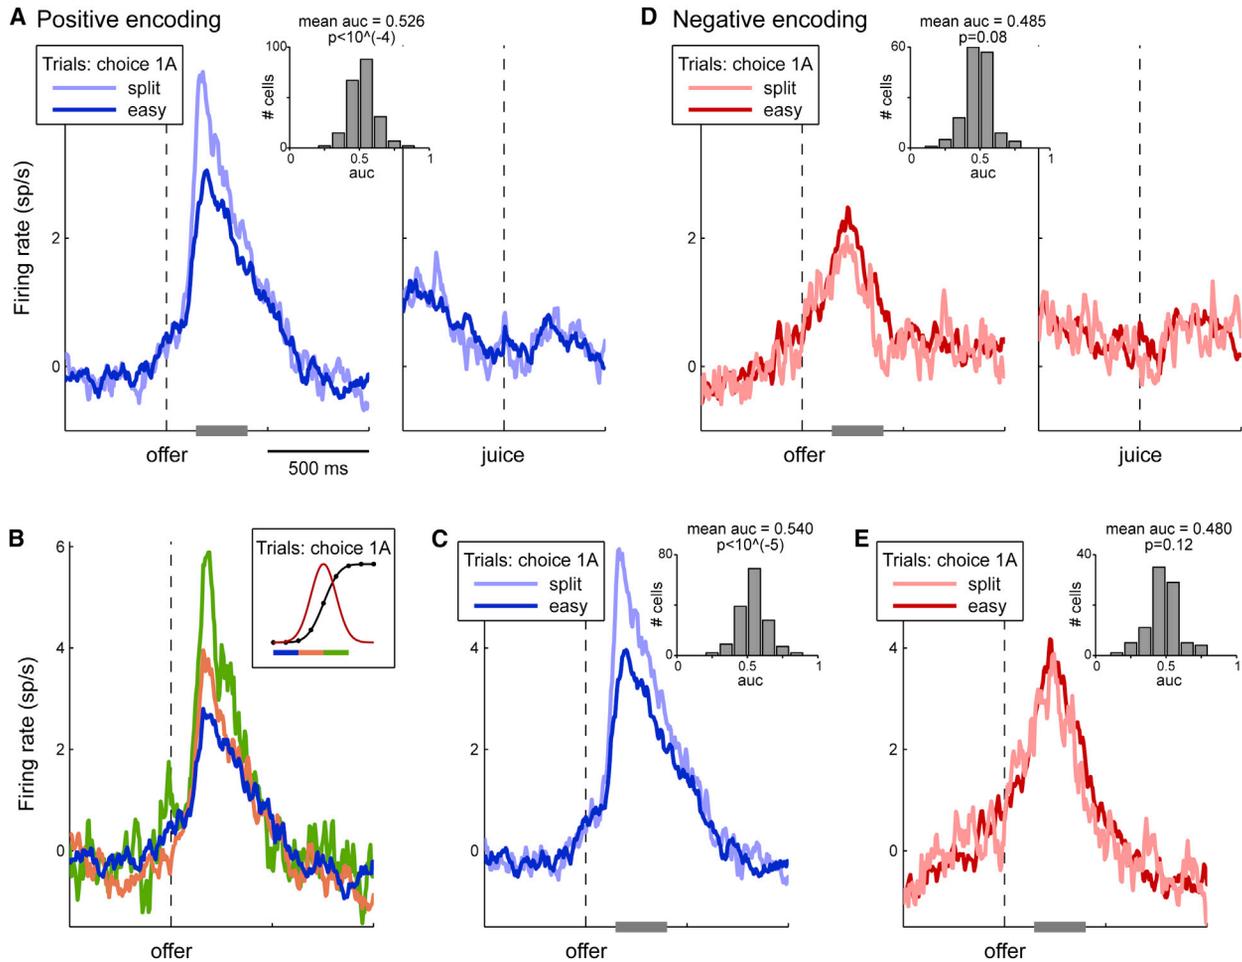
### Activity Overshooting of Chosen Value Cells

I now turn to *chosen value* cells (Figure 1D). To examine their activity in relation to choice variability, I focused on trials in which the animals chose one drop of the preferred juice over various amounts of the other juice (trials 1A  $\blacktriangleright$  qB, where q is the quantity of juice B offered). The motivation for this analysis was as follows. In principle, choice variability could ensue if the value of any particular good fluctuated from trial to trial. If so, one would expect that the activity of *chosen value* cells, conditioned on the animal choosing 1A, would be enhanced when the alternative offer is more desirable. To test this prediction, I divided offer types into easy and split. Consistent with the prediction, the activity of *chosen value* cells with positive encoding was clearly higher for split decisions compared to easy decisions (Figure 7A). This effect, termed “activity overshooting,” was evident in the time window 150–400 ms after the offer, which corresponds roughly to the time period in which the decision was made.

To assess whether the activity overshooting was generally measurable for individual cells, I performed an ROC analysis

focusing on the 150–400 ms after the offer. For each cell, I identified trials in which the animal chose 1A, and I divided them into easy and split decisions. Comparing the two distributions of firing rates, I obtained a measure for the AUC (Figure 7A, inset). In general, the AUC varied substantially across cells. However, the mean AUC for the population was significantly above the null hypothesis of 0.5 (mean AUC = 0.526,  $p < 10^{-4}$ , t test). In other words, individual cells typically presented an activity overshooting.

The result illustrated in Figure 7A was very robust (Figure S6). In a variant of this analysis, I divided the amount of juice B offered into three segments. Confirming the first observation, the activity of *chosen value* cells gradually varied as a function of the quantity of juice B (Figure 7B). Restricting the analysis to cells that were significantly tuned yielded similar results (Figure 7C). With respect to *chosen value* cells with negative encoding, one would expect a higher firing rate for easy decisions compared to split decisions. Focusing again on the 150–400 ms after the offer, this prediction was qualitatively met (Figure 7D), although the difference in signal was rather small and did not reach significance threshold (mean AUC was 0.485;  $p = 0.08$ , t test). Restricting the analysis to significantly tuned cells yielded similar results (mean AUC = 0.480;  $p = 0.12$ , t test; Figure 7E). Hence it was not clear whether *chosen value* cells with positive and negative encoding differed qualitatively or, alternatively, whether the measure obtained for cells with



**Figure 7. Activity Overshooting in Chosen Value Cells**

(A) Population with positive encoding. The analysis included only trials in which the animal chose one drop of juice (e.g., juice A). Trials were divided into two groups depending on whether the offer type was easy or split (see legend). Each trace represents the average activity profiles for positive encoding *chosen value* cells. Averages were calculated including only cells for which I could compute both traces ( $\geq 2$  trials per trace). Cells from experiment 1 contributed to each average with at most one trace (choices of 1A). However, some cells from experiment 2 contributed with two traces (choices of 1A or 1B). In total, each population trace shown here is the average of 212 individual traces from 151 cells. The insert illustrates the results of the ROC analysis.

(B) Same analysis as in (A), splitting trials into three groups. In all cases, the animal chose 1A over qB, with variable q. The three groups of trials correspond to easy decisions (dark blue), split decisions with  $q < \text{mean}(p)$  (orange) and split decisions with  $q \geq \text{mean}(p)$  (green). Each population trace is the average of 112 individual traces from 90 cells.

(C) Same analysis as in (A) including only cells that were significantly tuned in the 150–400 ms after the offer. Each population trace is the average of 156 individual traces from 112 cells.

(D) Population with negative encoding. Each population trace is the average of 155 individual traces from 106 cells. Consistent with the hypothesis that the *chosen value* fluctuated from trial to trial, the dark red line was slightly above the light red line in the 150–400 ms after the offer. However, the effect did not reach significance threshold.

(E) Same analysis as in (D) including only cells that were significantly tuned in the 150–400 ms after the offer. Each population trace is the average of 91 individual traces from 62 cells.

See also Figures S6 and S7.

negative encoding was, for some reason, noisier. Thus, subsequent analyses of *chosen value* cells focused on the population with positive encoding.

### Interpreting the Activity Overshooting: the Relative Value as a Stochastic Variable

Comparing the results for *chosen value* cells with those for *offer value* cells may seem to present a puzzle. Consider the

analyses illustrated in Figures 7A and S2E, respectively. Both analyses focused on trials in which the animal chose 1A. In both cases, the activity of neurons encoding the value of 1A (as an offer value in Figure S2E and as a chosen value in Figure 7A) was analyzed as a function of the quantity of juice B. The rationale for the two analyses was similar. However, *chosen value* cells presented a robust overshooting, whereas *offer value* cells showed no such effect. In other words, it

seemed that the overshooting was not driven by fluctuations in the activity of *offer value* cells. So how can the overshooting be explained?

In the framework of Figure 3B, decisions depend on the subjective value of each juice and on the “exchange rate” between the two juices, referred to as the relative value ( $\rho$ ). Given two goods,  $\rho$  can change over relatively long periods of time, for example due to changes in internal motivation (Padoa-Schioppa and Assad, 2006). More subtly,  $\rho$  could fluctuate on a trial-by-trial basis. Interestingly, stochastic fluctuations of  $\rho$  would induce overshooting in the activity of *chosen value* cells similar to that shown in Figure 7A. To appreciate this point, consider trials in which the animal chose between juices A and B. If value functions are linear,  $\rho$  is the quantity ratio that makes the animal indifferent between the two juices:

$$V(A) = \rho V(B). \quad (\text{Equation 5})$$

Assume now that  $\rho$  is a stochastic variable with given distribution. The choice of the animal in any particular trial imposes a constraint on the possible realizations of  $\rho$  in that trial. Consider, for example, trials in which the monkey chose one 1A over qB (trials 1A  $\blacktriangleright$  qB). Disregarding other sources of choice variability, Equation 1 implies that  $\rho \geq q$ . Thus the average  $\rho$  in trials 1A  $\blacktriangleright$  qB increases as a function of q. This variability will also be reflected in the activity of *chosen value* cells. Furthermore, considering trials in which the animal chose 1A, Equation 5 implies  $\rho = \text{chosen value}$  (in units of juice B). In conclusion, if  $\rho$  fluctuates stochastically, the activity of *chosen value* cells in trials 1A  $\blacktriangleright$  qB increases as a function of q.

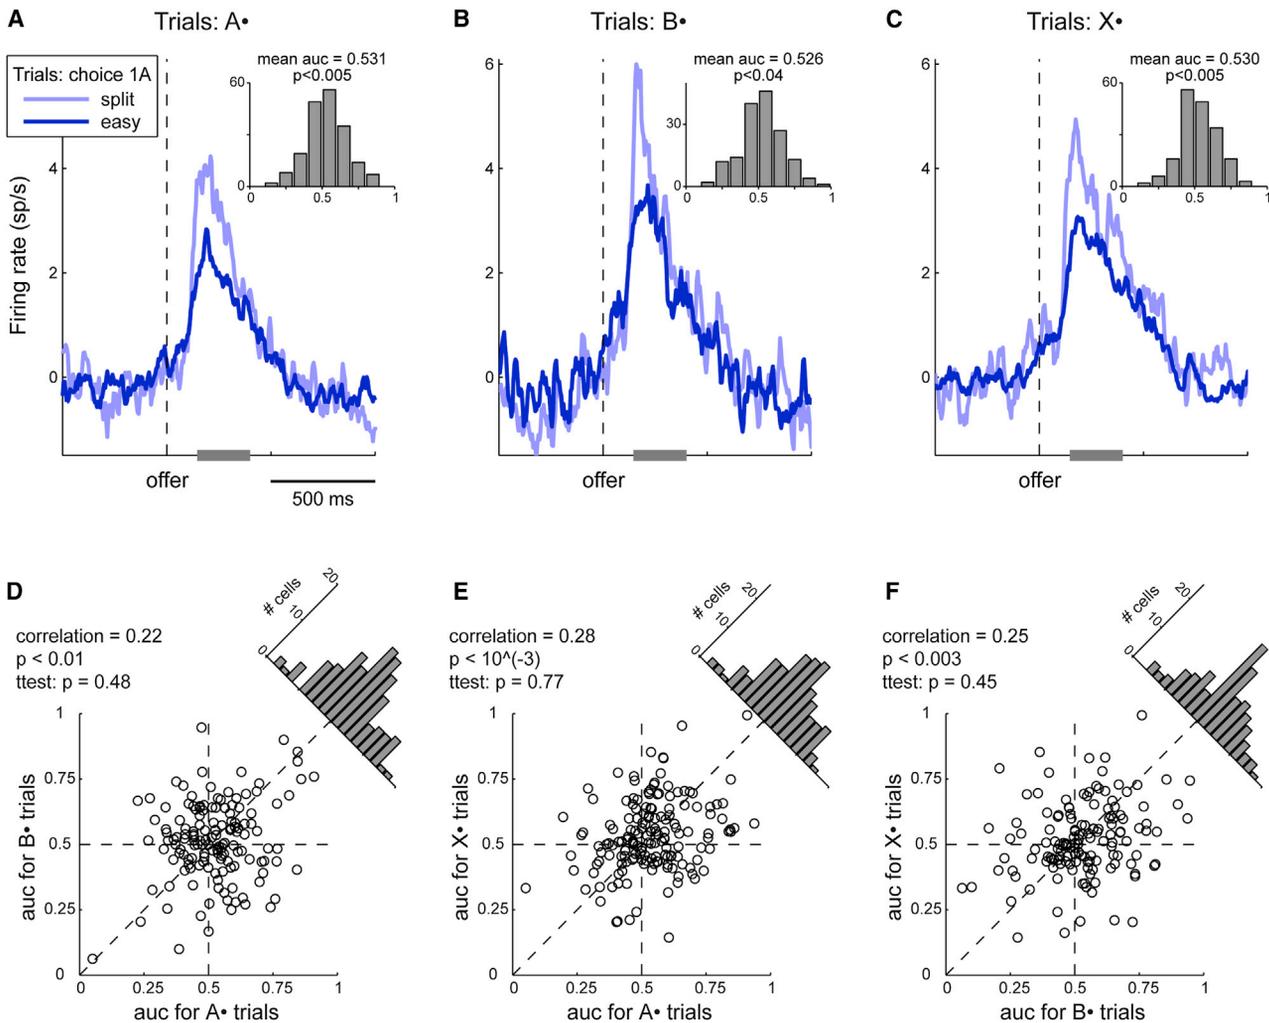
The activity overshooting of *chosen value* cells can thus be explained by fluctuations of  $\rho$ . An alternative hypothesis is that *chosen value* cells actually encode the *total value*. However, a quantitative analysis found that the explanatory power of *chosen value* corrected for fluctuations of  $\rho$  was significantly higher than that of *total value* ( $p < 0.01$ , Kruskal-Wallis test; Supplemental Experimental Procedures; Figure S7).

In summary, evidence suggested that the activity overshooting observed in *chosen value* cells reflected trial-by-trial fluctuations in  $\rho$ . If this is true, two neuronal phenomena described here appear related to choice variability: predictive activity of *chosen juice* cells and activity overshooting of *chosen value* cells. One important question was whether these phenomena were different manifestations of the same underlying source of variability or, alternatively, whether activity overshooting and predictive activity were mutually independent. As a first step to examine this issue, I took advantage of the fact that predictive activity was largely accounted for by choice hysteresis, and I repeated the analysis of *chosen value* cells while controlling for the outcome of the previous trial. The results provided strong evidence that the activity overshooting was independent of choice hysteresis (Figure 8; Supplemental Experimental Procedures). In contrast, the relation between the activity overshooting of *chosen value* cells and the residual predictive activity of *chosen juice* cells remains to be examined.

## DISCUSSION

I presented five primary results, each of which bears comments.

- (1) Variables *offer value*, *chosen value*, and *chosen juice* were encoded by different groups of neurons. The fact that these variables are encoded categorically and by different neurons appears rather significant and opens numerous questions regarding, for example, the possible correspondence between the three cell classes identified here and morphologically defined cell types. Addressing this and related issues is a primary goal for future research.
- (2) Trial-by-trial fluctuations in the activity of *offer value* cells did not explain choice variability in near-indifference decisions. Future work might revisit this issue with higher statistical power (e.g., collecting a larger data set or perhaps asking the animals to fixate individual offers). But taking the current findings at face value, how might this negative result be explained? Recent theoretical work demonstrates that choice probabilities (CPs)  $-0.5$  are proportional to the matrix product of noise correlations by read-out weights (Haefner et al., 2013). In this perspective, a distribution of CPs might be centered on zero due to several possible reasons. First, read-out weights could equal zero (in this case, *offer value* cells do not contribute to the decision). However, CPs would also be close to zero if noise correlations were very small, or if noise correlations within and across groups of cells encoding the *offer value* of different juices were similar, or if different neurons had positive and negative read-out weights. Starting from these considerations, future research shall examine noise correlations in the OFC.
- (3) The activity of *chosen juice* cells after the offer depended on the decision difficulty but did not resemble a race-to-threshold. At the outset of the study, I traced an analogy between *chosen juice* cells and neurons in LIP. The predictive activity found for *chosen juice* cells resembles that observed in LIP (Shadlen and Newsome, 2001) and thus supports this analogy. At the same time, the activity profile of *chosen juice* cells after the offer differs qualitatively from that described for LIP (Roitman and Shadlen, 2002; Shadlen and Newsome, 2001). Two observations seem relevant to this discrepancy. First, the concept of accumulation of evidence over time, which is central to perceptual decisions (Gold and Shadlen, 2007), does not equally apply to economic decisions. Indeed the “evidence” in economic decisions (i.e., offer values) is immediately available, not delivered gradually over time. Second, the steady-state activity of neurons in LIP during standard perceptual decisions may partly encode a motor plan (Andersen and Cui, 2009; Bisley and Goldberg, 2010) as distinguished from the decision outcome. In contrast, when decision outcomes and motor plans are dissociated, decision signals in LIP are transient and qualitatively similar to those illustrated here for *chosen juice* cells (Bennur and Gold, 2011).



**Figure 8. The Overshooting of Chosen Value Cells Is Independent of Choice Hysteresis**

(A) Analysis of *chosen value* cells restricted to A• trials. The insert illustrates the result of the ROC analysis performed in the 150–400 ms after the offer. All conventions are as in Figure 7A. The activity overshooting observed in *chosen value* cells is independent of the outcome of the previous trial.

(B) Analysis of B• trials.

(C) Analysis of X• trials.

(D) Comparing the AUC obtained for A• trials and B• trials. Each data point represents one neuron. Across the population, the two measures of AUC were significantly correlated (correlation = 0.22,  $p < 0.01$ ). In other words, the AUC for any given cell was reproducible. However, the difference between the two measures of AUC was statistically indistinguishable from zero (slanted histogram;  $p = 0.48$ , t test).

(E) Comparing the AUC obtained for A• trials and X• trials.

(F) Comparing the AUC obtained for B• trials and X• trials.

(4) Prior to the offer, *chosen juice* cells presented predictive activity correlated with the upcoming decision. Previous studies observed similar phenomena in other decision tasks (Shadlen and Newsome, 2001; Williams et al., 2003; Wyart and Tallon-Baudry, 2009). One notable difference is that the predictive activity found here is largely (but not entirely) related to the outcome of the previous trial. Predictive activity might suggest that *chosen juice* cells actively participate in the decision process. However, current results do not necessary imply this scenario. Indeed, an equally valid hypothesis is that other neurons, not yet identified, participate in or even determine the

decision, and separately inform the activity of *chosen juice* cells. In this latter scenario, the relation between the predictive activity documented here and the decision would be correlational, not causal. Disambiguating between these hypotheses will likely require different technical approaches such as selective microstimulation.

(5) In a limited time window shortly after the offer, *chosen value* cells presented an activity overshooting related to the decision difficulty. The present analyses suggest that the activity overshooting reflected stochastic fluctuations in the relative value  $\rho$ . Under this interpretation, an important question relates to how  $\rho$  is instantiated at the

neuronal level. In the framework of Figure 3B,  $\rho$  can be thought of as akin to a ratio of synaptic efficacies. Future work should examine this hypothesis in detail. At the same time, the framework schematized in Figure 3B is very general and compatible with a variety of possible architectures (Bogacz et al., 2006). More specific hypotheses with respect to the architecture might conceivably provide additional or alternative interpretations for the overshooting of *chosen value* cells.

### Comparing Mechanisms for Economic and Perceptual Decisions

It has often been hypothesized that the neural systems governing economic and perceptual decisions share fundamental principles and core mechanisms (Gold and Shadlen, 2007; Shadlen et al., 2008; Summerfield and Tsetsos, 2012; Wang, 2008). Upon a closer examination, the two neuronal systems do present important similarities, but also clear differences. First, *offer value* cells do not show consistent choice probabilities, unlike MT cells. Second, the activity of *chosen juice* cells does not resemble a race-to-threshold, unlike that of LIP cells (but see above). Third, *chosen value* cells do not have an obvious analog in perceptual decisions. Consequently, there is no known counterpart for the activity overshooting. Fourth, the encoding of value in the OFC undergoes range adaptation and, more generally, depends on the behavioral context in ways that differ from those found in MT (Kohn, 2007; Padoa-Schioppa, 2009; Padoa-Schioppa and Assad, 2008). Last but not least, neuronal activity in the OFC is nonspatial. In summary, economic decisions appear to involve distinct neuronal mechanisms that cannot be simply equated to those underlying perceptual decisions.

The hypothesis examined in the present study, namely that good-based decisions take place within the OFC, differs from a recently-proposed “attentional drift-diffusion model” (ADDM) (Krajbich et al., 2010). According to the ADDM, subjects switch their attention back and forth between the options and, at any given time, a comparator increments a decision variable in favor of the attended option. The comparator is thought to reside in the dorso-medial prefrontal cortex (dmPFC). Few considerations are in order. First, the relation between fixation patterns and choices (Krajbich et al., 2010) may, at least in part, reflect a causal relation opposite to that assumed in the ADDM. In other words, subjects might tend to look longer at offers they are leaning toward. Second, the evidence implicating dmPFC (Hare et al., 2011) is based on analyses of aggregate data and builds on assumptions that may not hold when neuronal responses are heterogeneous. Third, according to the ADDM, neurons in the OFC would encode not the *chosen value* per se, but rather the variable *chosen value* – *other value* (Lim et al., 2011). However, vanishingly few neurons were found to encode this variable (Padoa-Schioppa and Assad, 2006; 2008). In summary, current support for the ADDM is not conclusive. These considerations, together with an established literature showing that lesions to the OFC selectively impair value-based decisions, justify the hypothesis examined in this study.

To conclude, I showed that three variables intimately related to economic decisions—*offer value*, *chosen value*, and *chosen juice*—are encoded by three distinct groups of neurons in the

OFC. My analyses suggest that choice variability may be driven partly by the initial state of the neuronal assembly (revealed by the predictive activity of *chosen juice* cells) and partly by stochastic fluctuations in the relative value of the juices (revealed by the activity overshooting of *chosen value* cells). Finally, this study highlighted important analogies but also significant differences between the neuronal mechanisms of economic and perceptual decisions.

### EXPERIMENTAL PROCEDURES

#### Task Design and Preliminary Analyses

Data analyzed in this study are from two experiments (Padoa-Schioppa and Assad, 2006, 2008). Procedures for behavioral control, neuronal recording, and preliminary analyses have been described in detail. In both experiments, trials started with the animal fixating the center of a computer monitor (Figure 1A). After 0.5 s, two sets of colored squares, representing the two offers, appeared on the two sides of the fixation point. For each offer, the color represented the juice type and the number of squares represented the juice amount. The animal maintained central fixation for a randomly variable delay (1–2 s), after which the fixation point was extinguished and two saccade targets appeared by the offers (go signal). The animal indicated its choice with a saccade and maintained peripheral fixation for 0.75 s before juice delivery. Two animals, L and V, participated in each experiment. In experiment 1 (931 cells), animals chose in each session between two juices. In experiment 2 (557 cells), animals chose between three juices offered pairwise, and trials with the three juice pairs were randomly interleaved.

An offer type was defined by two offers (e.g., [1A:3B]). Juice pairs and offer types varied from session to session. Within a session, different offer types were pseudo-randomly interleaved. Their frequency varied, but each offer type was typically presented at least 20 times in each session. A trial type was defined by an offer and a choice (e.g., [1A:3B,A]). The analysis presented in the section, *From Neuronal Responses to Cell Classes*, and the subsequent cell classification was based on four primary time windows: postoffer (0.5 s after the offer), late delay (0.5–1 s after the offer), prejuice (0.5 s before juice delivery), and postjuice (0.5 s after juice delivery). A neuronal response was defined as the activity of one cell in one time window as a function of the trial type. Task-related responses were identified with an ANOVA (factor trial type,  $p < 0.001$ ). Across experiments, 843/1,488 (57%) neurons were task-related in at least one time window. Previous studies showed that variables *offer value*, *chosen value*, and *chosen juice* explain the vast majority of task-related responses. To classify responses, I performed a linear regression of each task-related response on each variable. A variable was said to explain a response if the regression slope differed significantly from zero ( $p < 0.05$ ). If a variable did not explain a response,  $R^2$  was set equal to zero. If more than one variable explained one response, the response was assigned to the variable with the highest  $R^2$ . Across experiments, neurons encoding one of the three variables in at least one time window were 783/1,488 (53%); 443 from experiment 1; 340 from experiment 2).

#### Statistical Analyses

Several analyses presented in this Article were conducted by dividing trials into two groups—easy and split. In all cases, split refers to offer types in which the animal split its decisions between the two offers, conditioned on the fact that the animal chose either option at least twice; easy refers to offer types in which the animal consistently chose the same option. The label “easy” captures the fact that these decisions were presumably easier. Restricting the analysis to trials in which the animal chose one drop of juice A against variable quantities  $q$  of juice B (1A  $\blacktriangleright$   $q$ B), easy/split also corresponds to low/high values of  $q$ .

All ROC analyses were done on row spike counts, without time averaging or baseline correction. The details of the analysis, however, differed to some extent depending on the neuronal population. For Figure 5A, I identified offer types in which decisions were split. For each offer type, I divided trials into two groups depending on the chosen juice (E or O). The two groups were compared with an ROC, from which I measured the area under the curve

(AUC). This AUC is equivalent to the measure of choice probability defined for perceptual decisions (Britten et al., 1996; Nienborg et al., 2012). To obtain a single AUC for each neuron, I averaged the AUC across offer types (Kang and Maunsell, 2012). The same procedure was used for Figure 5B. The results reported in Figures 5A and 5B were obtained with an arithmetic average. However, the results obtained weighing the AUC obtained for each offer type with the geometric mean of the two trial numbers (corresponding to choices of E and O) were essentially identical. For Figure 6A, I identified offer types in which decisions were split, and I divided trials into two groups depending on the chosen juice (E or O). In this case, trials from different offer types were pooled and compared with the ROC analysis. For Figures 6C and 7A, I focused on trials in which the animal chose one drop of the preferred juice (1A). These trials were divided into two groups depending on whether decisions with the corresponding offer type were easy or split. The two groups of trials were compared directly with the ROC analysis. The analyses illustrated in Figures 7C–7E, S2A, and S2B were restricted to cells that were significantly tuned in the 150–400 ms following the offer. For each cell, tuning was established with an ROC analysis of all trials, dividing them into tertiles of *chosen value* (as in Figures 2A–2D), comparing the activity obtained for the “high” and “low” tertiles and requiring that the AUC differ significantly from 0.5 ( $p < 0.05$ ).

For logistic analyses, data from experiment 2 were divided into three groups corresponding to the three juice pairs. For simplicity, I refer to each of these groups of trials as a “session.”

#### SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures and seven figures and can be found with this article online at <http://dx.doi.org/10.1016/j.neuron.2013.09.013>.

#### ACKNOWLEDGMENTS

This research was supported by the National Institute on Mental Health (R03-MH093330) and the National Institute on Drug Addiction (R01-DA032758). I thank J. Assad, K. Conen, A. Rustichini, L. Snyder, and members of my laboratory for helpful discussions and comments on the manuscript.

Accepted: September 12, 2013

Published: December 4, 2013

#### REFERENCES

- Andersen, R.A., and Cui, H. (2009). Intention, action planning, and decision making in parietal-frontal circuits. *Neuron* 63, 568–583.
- Bennur, S., and Gold, J.I. (2011). Distinct representations of a perceptual decision and the associated oculomotor plan in the monkey lateral intraparietal area. *J. Neurosci.* 31, 913–921.
- Bisley, J.W., and Goldberg, M.E. (2010). Attention, intention, and priority in the parietal lobe. *Annu. Rev. Neurosci.* 33, 1–21.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J.D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* 113, 700–765.
- 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.
- Buckley, M.J., Mansouri, F.A., Hoda, H., Mahboubi, M., Browning, P.G., Kwok, S.C., Phillips, A., and Tanaka, K. (2009). Dissociable components of rule-guided behavior depend on distinct medial and prefrontal regions. *Science* 325, 52–58.
- 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.
- Cohen, M.R., and Newsome, W.T. (2009). Estimates of the contribution of single neurons to perception depend on timescale and noise correlation. *J. Neurosci.* 29, 6635–6648.
- Gallagher, M., McMahan, R.W., and Schoenbaum, G. (1999). Orbitofrontal cortex and representation of incentive value in associative learning. *J. Neurosci.* 19, 6610–6614.
- Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.* 30, 535–574.
- Haefner, R.M., Gerwinn, S., Macke, J.H., and Bethge, M. (2013). Inferring decoding strategies from choice probabilities in the presence of correlated variability. *Nat. Neurosci.* 16, 235–242.
- Hare, T.A., Schultz, W., Camerer, C.F., O’Doherty, J.P., and Rangel, A. (2011). Transformation of stimulus value signals into motor commands during simple choice. *Proc. Natl. Acad. Sci. USA* 108, 18120–18125.
- Hunt, L.T., Kolling, N., Soltani, A., Woolrich, M.W., Rushworth, M.F.S., and Behrens, T.E.J. (2012). Mechanisms underlying cortical activity during value-guided choice. *Nat. Neurosci.* 15, 470–476, S1–S3.
- Kable, J.W., and Glimcher, P.W. (2007). The neural correlates of subjective value during intertemporal choice. *Nat. Neurosci.* 10, 1625–1633.
- Kang, I., and Maunsell, J.H. (2012). Potential confounds in estimating trial-to-trial correlations between neuronal response and behavior using choice probabilities. *J. Neurophysiol.* 108, 3403–3415.
- Kennerley, S.W., Dahmubed, A.F., Lara, A.H., and Wallis, J.D. (2009). Neurons in the frontal lobe encode the value of multiple decision variables. *J. Cogn. Neurosci.* 21, 1162–1178.
- Kohn, A. (2007). Visual adaptation: physiology, mechanisms, and functional benefits. *J. Neurophysiol.* 97, 3155–3164.
- Krajbich, I., Armel, C., and Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nat. Neurosci.* 13, 1292–1298.
- Levy, I., Snell, J., Nelson, A.J., Rustichini, A., and Glimcher, P.W. (2010). Neural representation of subjective value under risk and ambiguity. *J. Neurophysiol.* 103, 1036–1047.
- 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.
- Newsome, W.T. (1997). The King Solomon lectures in neuroethology. Deciding about motion: linking perception to action. *J. Comp. Physiol. A Neuroethol. Sens. Neural Behav. Physiol.* 181, 5–12.
- Nienborg, H., Cohen, M.R., and Cumming, B.G. (2012). Decision-related activity in sensory neurons: correlations among neurons and with behavior. *Annu. Rev. Neurosci.* 35, 463–483.
- 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., and Assad, J.A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature* 441, 223–226.
- Padoa-Schioppa, C., and Assad, J.A. (2008). The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nat. Neurosci.* 11, 95–102.
- Padoa-Schioppa, C., Jandolo, L., and Visalberghi, E. (2006). Multi-stage mental process for economic choice in capuchins. *Cognition* 99, B1–B13.
- Peters, J., and Büchel, C. (2009). Overlapping and distinct neural systems code for subjective value during intertemporal and risky decision making. *J. Neurosci.* 29, 15727–15734.
- Plassmann, H., O’Doherty, J., and Rangel, A. (2007). Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *J. Neurosci.* 27, 9984–9988.
- Roesch, M.R., and Olson, C.R. (2005). Neuronal activity in primate orbitofrontal cortex reflects the value of time. *J. Neurophysiol.* 94, 2457–2471.
- Roitman, J.D., and Shadlen, M.N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* 22, 9475–9489.

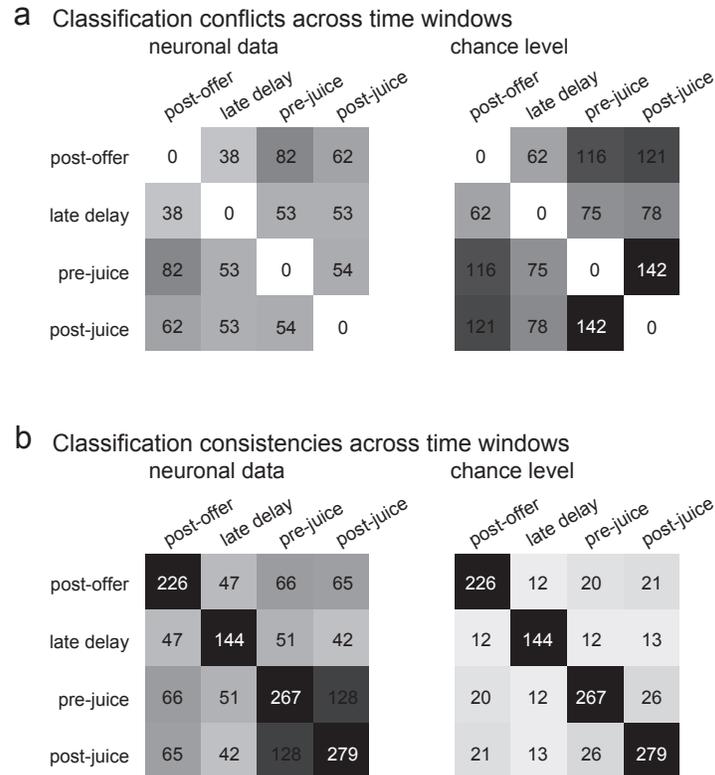
- Rudebeck, P.H., and Murray, E.A. (2011). Dissociable effects of subtotal lesions within the macaque orbital prefrontal cortex on reward-guided behavior. *J. Neurosci.* *31*, 10569–10578.
- Shadlen, M.N., and Newsome, W.T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* *86*, 1916–1936.
- Shadlen, M.N., Britten, K.H., Newsome, W.T., and Movshon, J.A. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J. Neurosci.* *16*, 1486–1510.
- Shadlen, M.N., Kiani, R., Hanks, T.D., and Churchland, A.K. (2008). Neurobiology of decision making: an intentional framework. In *Better Than Conscious? Decision Making, the Human Mind, and Implications for Institutions*, C. Engel and W. Singer, eds. (Cambridge: MIT Press).
- Simmons, J.M., and Richmond, B.J. (2008). Dynamic changes in representations of preceding and upcoming reward in monkey orbitofrontal cortex. *Cereb. Cortex* *18*, 93–103.
- Soltani, A., De Martino, B., and Camerer, C. (2012). A range-normalization model of context-dependent choice: a new model and evidence. *PLoS Comput. Biol.* *8*, e1002607.
- Solway, A., and Botvinick, M.M. (2012). Goal-directed decision making as probabilistic inference: a computational framework and potential neural correlates. *Psychol. Rev.* *119*, 120–154.
- 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.
- Summerfield, C., and Tsetsos, K. (2012). Building bridges between perceptual and economic decision-making: neural and computational mechanisms. *Front. Neurosci.* *6*, 70.
- Wang, X.J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* *36*, 955–968.
- Wang, X.J. (2008). Decision making in recurrent neuronal circuits. *Neuron* *60*, 215–234.
- West, E.A., DesJardin, J.T., Gale, K., and Malkova, L. (2011). Transient inactivation of orbitofrontal cortex blocks reinforcer devaluation in macaques. *J. Neurosci.* *31*, 15128–15135.
- Williams, Z.M., Elfar, J.C., Eskandar, E.N., Toth, L.J., and Assad, J.A. (2003). Parietal activity and the perceived direction of ambiguous apparent motion. *Nat. Neurosci.* *6*, 616–623.
- Wyart, V., and Tallon-Baudry, C. (2009). How ongoing fluctuations in human visual cortex predict perceptual awareness: baseline shift versus decision bias. *J. Neurosci.* *29*, 8715–8725.
- Yang, T., and Shadlen, M.N. (2007). Probabilistic reasoning by neurons. *Nature* *447*, 1075–1080.

**Neuron, Volume 80**

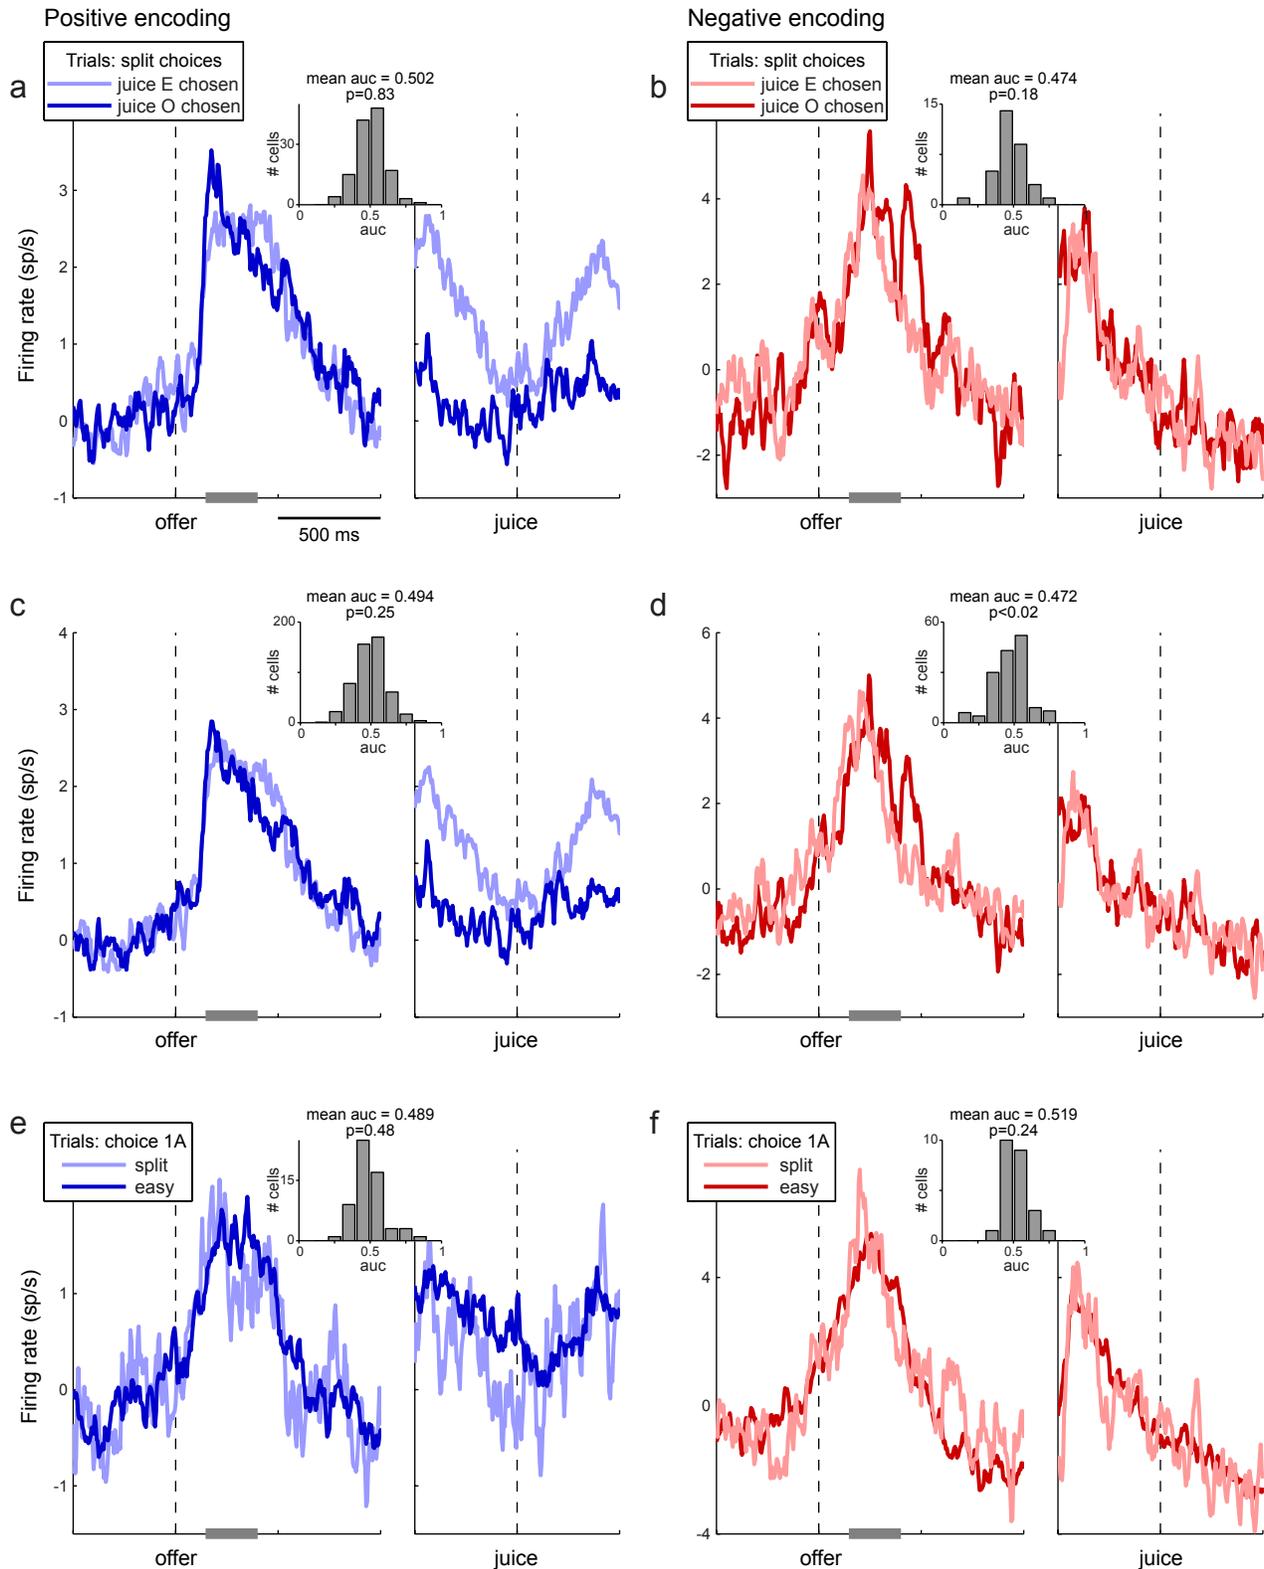
**Supplemental Information**

**Neuronal Origins of Choice Variability  
in Economic Decisions**

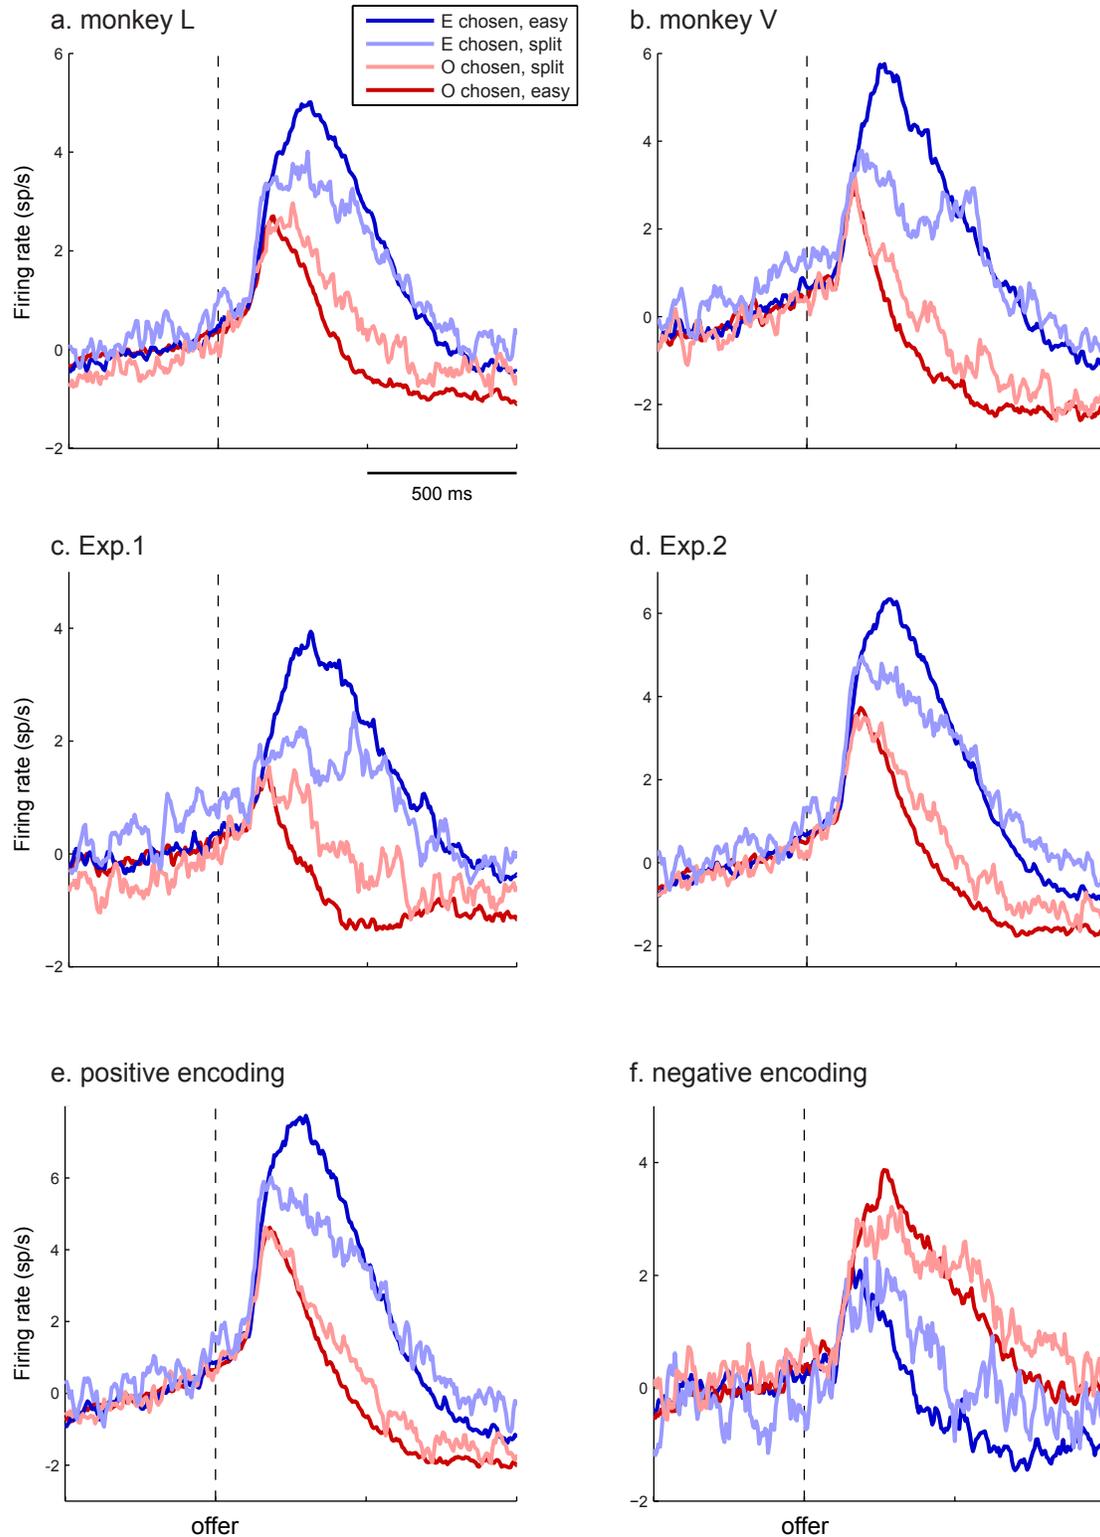
**Camillo Padoa-Schioppa**



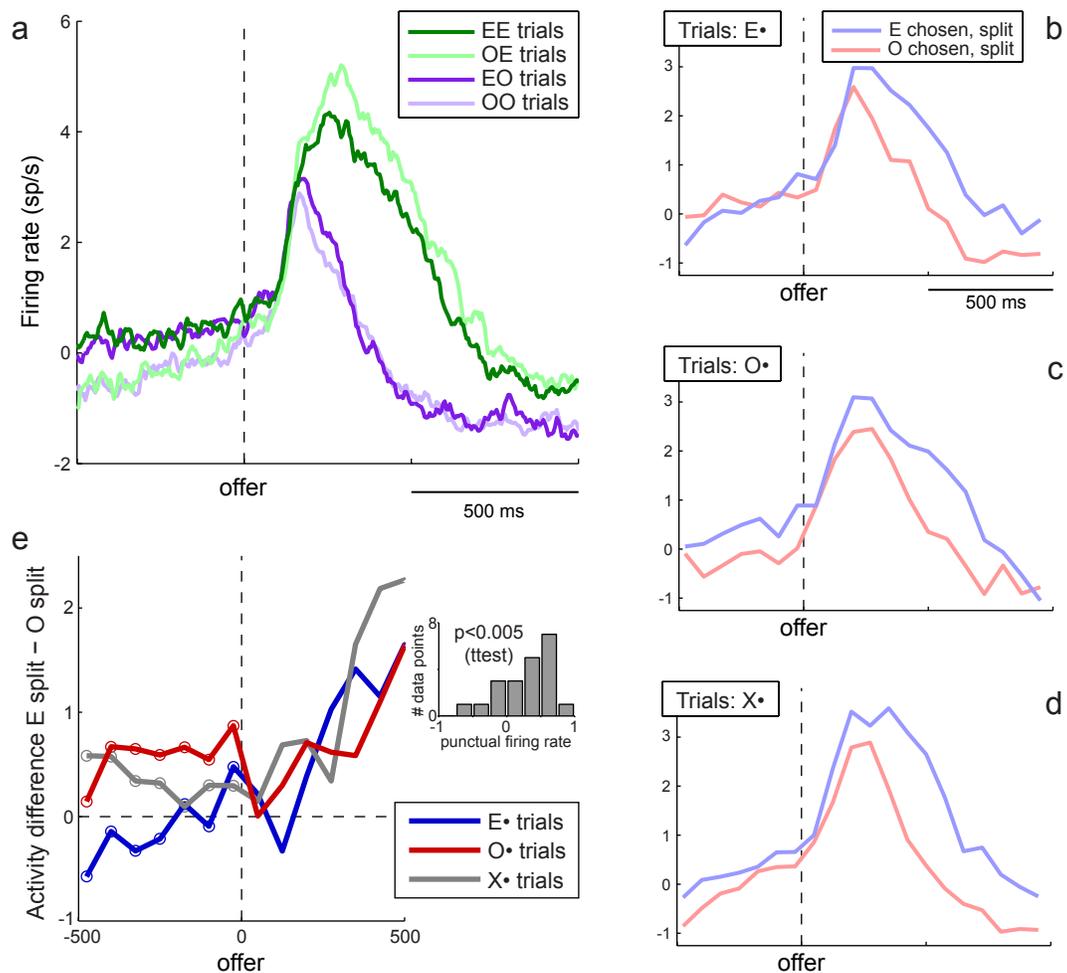
**Figure S1.** Analysis of classification conflict (related to Figure 1). **a.** In total, 443 cells from Exp.1 were classified in at least one time window and included in this analysis. The left panel refers to actual neuronal data. For each pair of time window, the number indicated the number of cells that showed classification conflict. The right panel refers to chance level and indicates, for each pair of time windows, the mean of the distribution obtained from the bootstrap (rounded). Shades of gray illustrate the same numbers graphically. For every pair of time windows, actual classification conflicts were significantly fewer than expected by chance. **b.** Analysis of classification consistency. For each pair of time windows, the number indicated the number of cells that presented consistent classification. The right panel refers to chance level. In every pair of time windows (except on the diagonal), the consistency of classification was much more frequent than expected by chance.



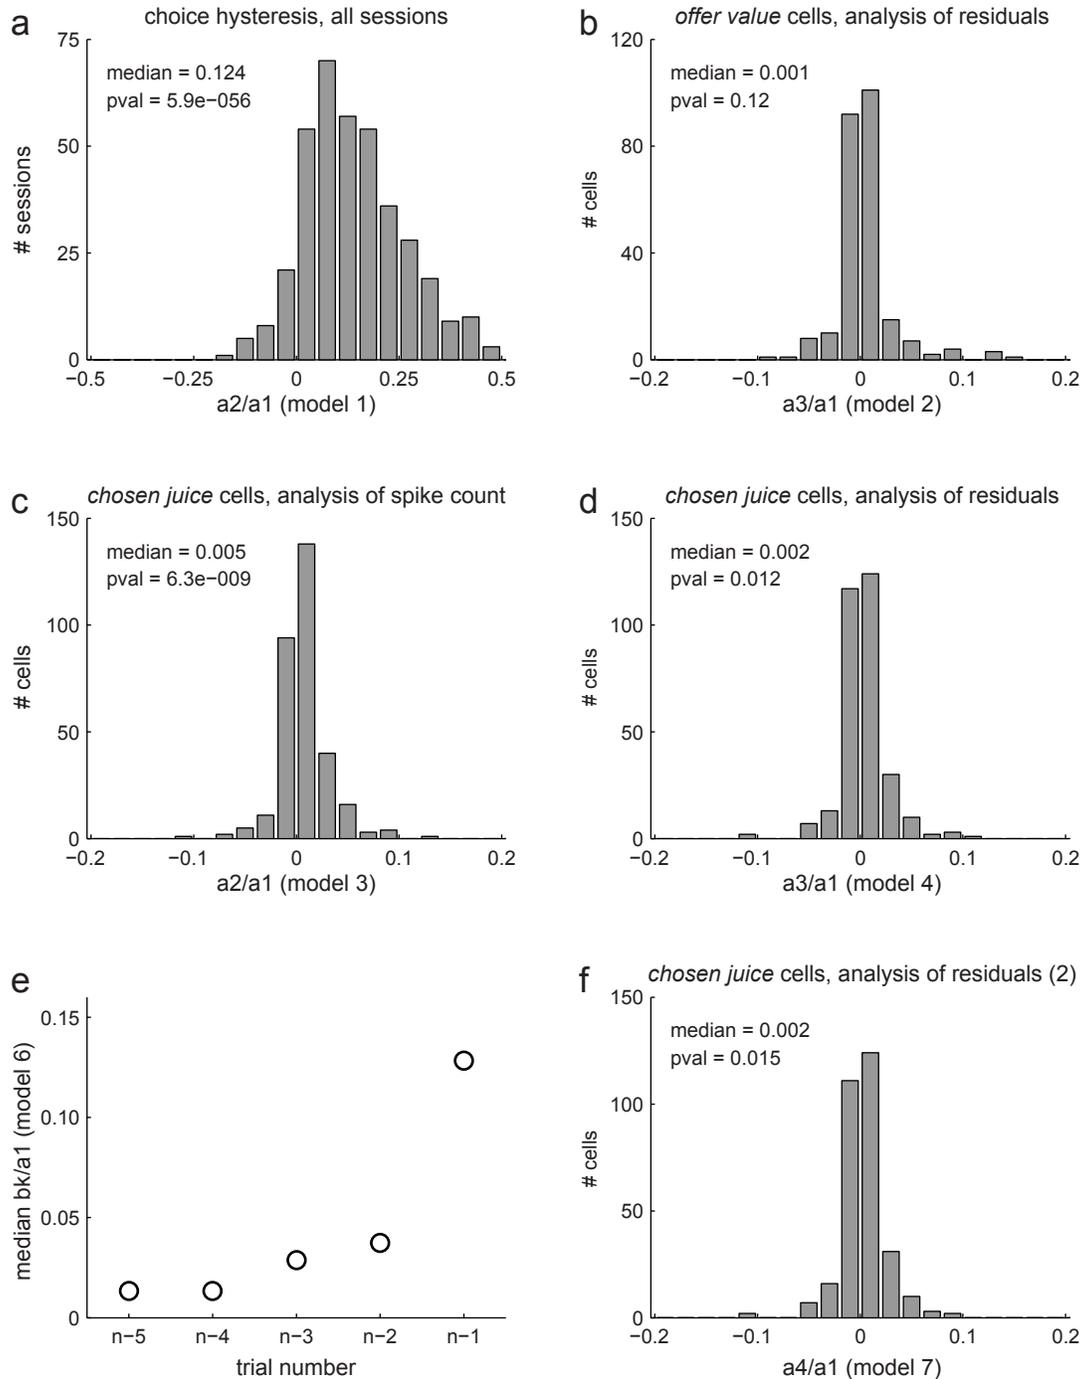
**Figure S2.** Control analyses for offer value cells (related to Figure 5). **ab.** Same analysis as in Fig.5ab including only neurons that were tuned in the 150-400 ms after the offer. Traces for positive and negative encoding are from 130 cells and 33 cells, respectively. **cd.** In this analysis, trials were split depending on the outcome of the previous trial (trials E•, O•, and X•). Each neuron thus contributed up to three traces. Population traces for positive and negative encoding are the average of 509 traces and 154 traces, respectively. **e.** Activity in relation to the other value (positive encoding). This analysis focused on offer value cells and on trials in which the animal chose one drop of the preferred juice (1A). Trials were divided into two groups depending on whether the offer type was easy (dark blue) or split (light blue) (see Experimental Procedures). Average traces shown here are from the 59 cells for which I could compute both traces ( $\geq 2$  trials per trace). The results fail to support the hypothesis that near-indifference decisions were driven by fluctuations in the activity of offer value cells. **f.** Activity in relation to the other value (negative encoding cells). Average traces shown here are from 24 cells.



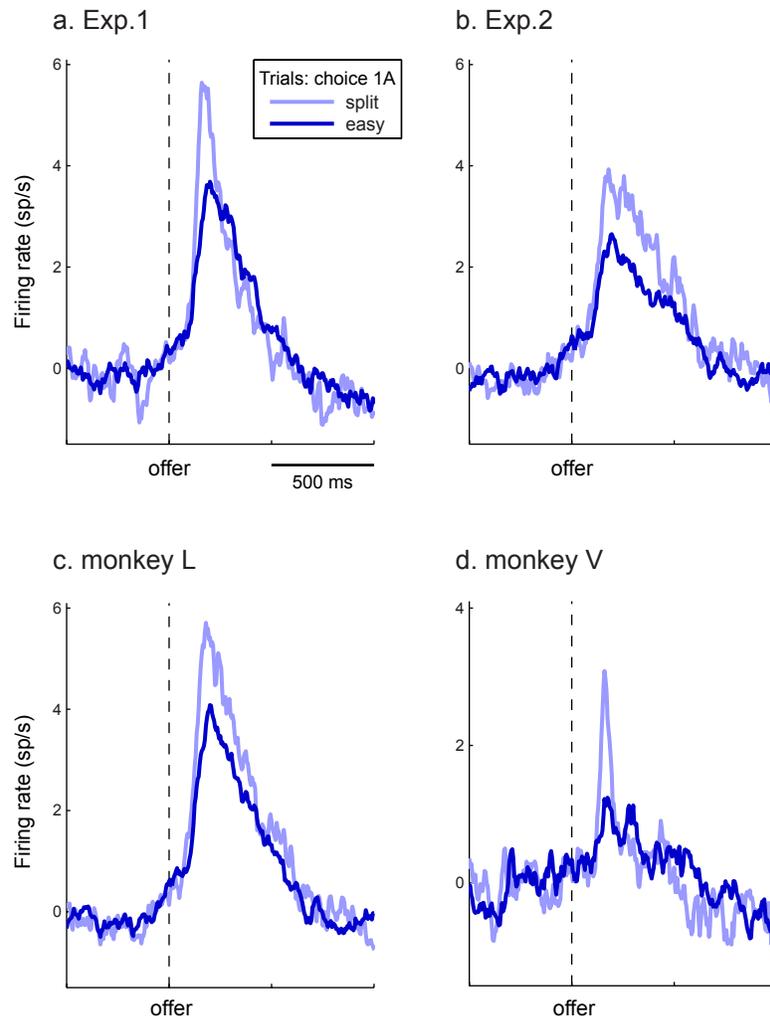
**Figure S3.** Control analyses for chosen juice cells (related to Figure 6). To verify the robustness of the results obtained for chosen juice cells, I repeated the analysis for different subsets of neurons: cells from monkey L (a, 169 cells), cells from monkey V (b, 96 cells), cells recorded in Exp.1 (c, 119 cells), cells recorded in Exp.2 (d, 146 cells). In all those cases, neurons with positive and negative encoding were pooled together. Data from Exp.2 were further broken down into positive encoding (e, 96 cells) and negative encoding (f, 50 cells). Both phenomena described for Fig.5a – namely the dependence on the decision difficulty and the predictive activity – can be observed for each subset of cells.



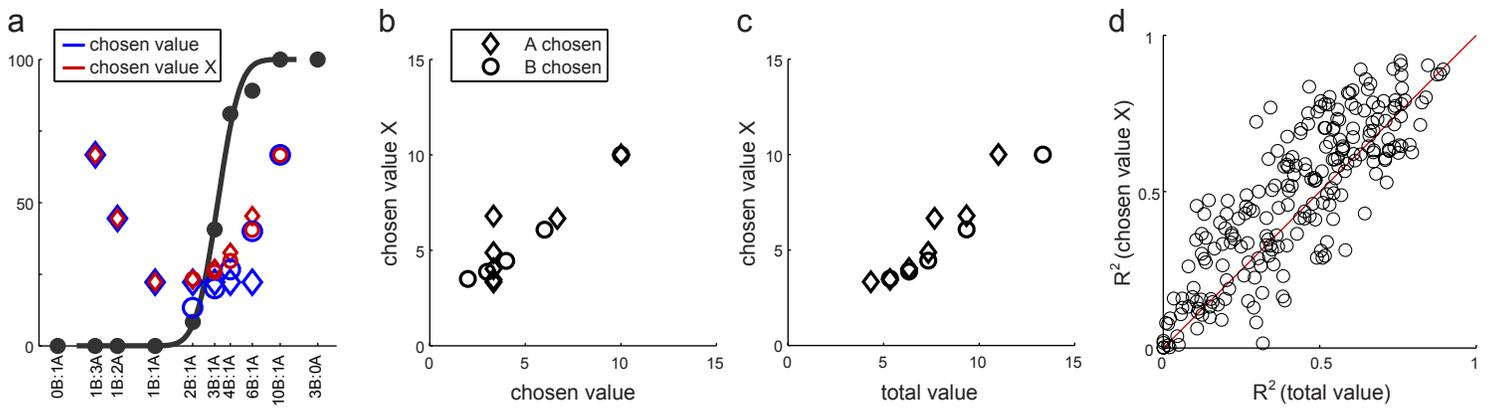
**Figure S4.** Chosen juice cells, activity in relation to the previous trial (related to Figure 6). **a.** Trials were divided depending on both the previous outcome and the current choice (see legend). The activity of chosen juice cells mainly depends on the current choice (green traces above purple traces after the offer). However, there is a tail activity from the previous trial (dark traces above light traces before the offer). **b-d.** Residual predictive activity. These plots focus only on split decisions. For these plots, the activity traces were coarse-grained by averaging firing rates in 75 ms bins (non-overlapping). **e.** Residual predictive activity, combined. Each line represents the difference between the two traces shown in (b-d). The combined distribution was displaced above zero.



**Figure S5.** Results of logistic analyses (related to Figures 4, 5, 6). **a.** Choice hysteresis (same as Fig.4a). The x-axis represents the ratio  $a2/a1$  defined in Eq.1, the y-axis represents the number of sessions (304 total). **b.** Analysis of offer value cells. The x-axis represents the ratio  $a3/a1$  defined in Eq.2, the y-axis represents the number of cells (324 total). Note that offer value cells from Exp.2 contributed to the histogram with 2 data points. **c.** Chosen juice cells, predictive activity. The x-axis represents the ratio  $a2/a1$  defined in Eq.3, the y-axis represents the number of cells (411 total). **d.** Chosen juice cells, residual predictive activity. The x-axis represents the ratio  $a3/a1$  defined in Eq.4. **e.** Time course of choice hysteresis. The x-axis represents trial number and the y-axis represents the median regression coefficient ( $bk/a1$ ) across the population (see model 6). The data point for trial  $n-1$  (roughly) corresponds to the median of the distribution in panel (a). The effect of choice hysteresis per se was essentially confined to trial  $n-1$ . In addition, there was a smaller effect that could be measured over several trials likely due to small drifts of relative value within the course of a session. **f.** Chosen juice cells, residual predictive activity accounting for the previous 2 trials. The results obtained here (model 7) are almost identical to those obtained accounting for the previous 1 trial (panel (d), model 4). In each analysis, I removed data points for which the logistic regression did not converge. The analysis of offer value cells focused on the 500 ms after the offer. The analysis of chosen juice cells focused on the 500 ms before the offer. Histograms in b, c, d and f include neurons with positive and negative encoding (the sign of the x-axis was reversed for negative encoding cells).



**Figure S6.** Control analyses for chosen value cells (related to Figure 7). Same analysis and all conventions as in Fig.7a. **a.** Data from Exp.1 (88 traces from 87 cells). **b.** Data from Exp.2 (124 traces from 64 cells). **c.** Data from monkey L (144 traces from 104 cells). **d.** Data from monkey V (68 traces from 47 cells).



**Figure S7.** Overshooting of chosen value cells: contrasting variable chosen value X and total value (related to Figure 7).  
**a.** Comparing chosen value and chosen value X, example session. The choice pattern is the same as in Fig.1d. Blue and red symbols refer to variables chosen value and chosen value X, respectively. Away from the indifference point, the two variables are essentially identical. However, near the indifference point, the two variables differ. Specifically, chosen value X is higher than chosen value for trials in which the animal seemingly chooses the "lesser" option.  
**b.** Chosen value X (y-axis) versus chosen value (x-axis). Same data as in (a). **c.** Chosen value X (y-axis) versus total value (x-axis). **d.** Contrasting the explanatory power of chosen value X and total value. Each symbol represents one cells and one trial group (A•, B• and X•). The y-axis (x-axis) represents the  $R^2$  obtained from the linear regression of the neuronal firing rate onto the variable chosen value X (total value). It can be noted that most of the data points lie above the diagonal line, indicating that chosen value X generally provided a better fit for the data.

## Supplemental Experimental Procedures

### From neuronal responses to cell classes

Previous analyses showed that individual responses in the OFC encode individual variables. Indeed, adding a second variable or a quadratic term to the linear regression usually failed to significantly improve the linear fit (Padoa-Schioppa and Assad, 2006). To examine whether *offer value* and *chosen value* were different classes of responses or, alternatively, two poles of a continuum, I first focused on Exp.1. I considered all the responses encoding either *offer value* or *chosen value*. Notably, *offer value* was a collapsed variable and responses could in fact encode either *offer value A* or *offer value B*. For each response, I considered each of the  $R^2$  obtained from the linear regressions onto the encoded variable and the other, non-encoded variable (independently of whether the latter explained the response). I then computed the difference  $\Delta R^2 = R^2_{offer\ value} - R^2_{chosen\ value}$ . This was done in one of two ways. For *offer value* responses,  $R^2_{offer\ value}$  was always the higher of the two  $R^2$  provided by *offer value A* and *offer value B*. For *chosen value* responses,  $R^2_{offer\ value}$  was either the higher of the two  $R^2$  provided by *offer value A* and *offer value B* or, alternatively, one of the two  $R^2$  randomly selected. The results reported here refer to the latter procedure. The former procedure provided very similar results (a bimodal distribution for  $\Delta R^2$ ;  $p < 0.01$ , Hartigan's dip test), except that the distribution was displaced towards higher values of  $\Delta R^2$  (as expected). Responses from Exp.2 (were *offer value* responses could encode *offer value A*, *offer value B* or *offer value C*) were treated similarly. Analogous procedures were used to compare variables *chosen value* and *chosen juice* and variables *offer value* and *chosen juice*. Data from the two experiments are pooled in Fig.1f-k.

Next, I sought to establish whether the incidence of classification conflicts actually found in the population was greater, comparable or lower than the incidence expected if conflicts occurred by chance. To estimate chance level, I used a bootstrap technique. For each time window, each cell was reassigned to a new variable with a random permutation of the variables recorded across the population in that time window. The permutation was done separately for each time window and the procedure was repeated for 1,000 times. This procedure thus provided, for each pair of time windows, two distributions for the number of classification conflicts and for the number of classification consistencies expected by chance. The procedure also provided a distribution for the total number of conflicts expected across the population. The results of this analysis are shown in Fig.S1.

### Activity profiles

Several analyses presented in the paper were conducted by dividing trials into two groups – easy and split. In all cases, split refers to offer types in which the animal split its decisions between the two offers, conditioned on the fact that the animal chose either option at least twice; easy refers to offer types in which the animal consistently chose the same option. To calculate the activity profile (i.e., the spike density function), trials were aligned at the time of the offer and separately at the time of juice delivery. For each alignment and each trial, the spike train was smoothed using the method of So and Stuphorn (2010). Spike times, expressed in 1 ms resolution, were convolved with the kernel:

$$\begin{aligned} k(t) &= (1 - \exp(-t/\tau_g))^* \exp(-t/\tau_d) && \text{for } t \geq 0 \\ k(t) &= 0 && \text{for } t < 0 \end{aligned}$$

This kernel mimics a post-synaptic potential and ensures that each spike only exerts its influence forward in time. Following previous work (Sayer et al., 1990; So and Stuphorn, 2010), I used  $\tau_g = 1$  ms and  $\tau_d = 20$  ms. For each cell, I then averaged the spike trains across all relevant trials and obtained a smoothed activity profile. Finally, I coarse-grained the signal by averaging the activity in non-adjacent 5 ms bins. This binning was performed only for display purposes; all statistical analyses were based on spike counts.

With the exception of Fig.2, all activity profiles are displayed after baseline subtraction. To calculate them, I subtracted from the activity of each cell the average activity in the 0.5 s preceding the offer. I then averaged the activity profiles across the relevant population.

### Analysis of activity profiles by quantile

For each *offer value* and *chosen value* cell, I divided trials into three tertiles depending on the value of the encoded variable (high, medium, low). I then averaged the activity of each tertile across the population. This was done separately for cells with positive and negative encoding (Fig.2a-d). Several aspects emerged from this analysis. First, the overall baseline activity ranged between 6 and 10 Hz for the various populations. The overall modulation (activity difference between the first and last tertile) ranged from roughly 2 to 6 Hz. Second, neurons with negative encoding did not simply decrease their activity compared to baseline. Rather, they often showed an increased activity for lower values of the encoded variable (this was most clear for *offer value* cells). Third, different groups of cells (e.g., *offer value* cells with negative encoding) presented robust preparatory activity preceding the offer.

Before conducting a similar analysis for *chosen juice* cells, I examined the sign of the encoding for this neuronal population. Indeed, previous work described the sign of the encoding for *offer value* and *chosen value* cells (Padoa-Schioppa, 2009), but it did not establish whether negative encoding also exists for *chosen juice* cells. In fact, this issue cannot be addressed based on data from Exp.1, where only two juices A and B were included in each session, because one cannot disambiguate between higher firing rate for one juice (positive encoding) and lower firing rate for the other juice (negative encoding). However, the sign of the encoding can be examined in data from Exp.2, where three juices (A, B and C) were included in each session. In this case, a neuron encoding, for example, *chosen juice A* with a positive sign would have high activity when the animal chooses juice A and low activity when the animal chooses either juice B or juice C. In contrast, a neuron encoding *chosen juice B* with a negative sign would have high activity when the animal chooses either juice A or juice C and low activity when the animal chooses juice B. In total, 146 *chosen juice* cells were recorded in Exp.2. Across this population, the sign of the encoding was positive for 96 (66%) cells and negative for 50 (34%) cells.

Based on this classification, I analyzed the average neuronal signal for *chosen juice* cells with positive and negative encoding. I divided trials depending on whether the animal chose the juice encoded by the cell (E) or the other juice (O; Fig.2e). For both groups of cells, the overall baseline activity and overall modulation during the delay were roughly equal to 10 Hz and 2 Hz, respectively. Both groups of cells presented preparatory activity preceding the offer. After the offer, *chosen juice* cells did not simply decrease their activity when the animal chose the other juice. Interestingly, the activity of this population clearly discriminated between the two juices starting <200 ms after the offer.

### Control analyses for *offer value* cells

A general concern is whether the negative results of Fig.5ab were veridical or due to spurious factors in the analysis. I considered several possible factors.

First, the ROC analysis in Fig.5ab focused on the time window 150-400 ms after the offer. This window was chosen based on inspection of Fig.2e and for consistency with the analysis of *chosen value* cells (Fig.7). At the same time, it is reasonable to question whether it would be more appropriate to run the ROC analysis on a later time window. However, the analysis repeated on the time window 300-500 ms after the offer yielded very similar results for both positive encoding cells (mean AUC = 0.509;  $p = 0.28$ ; t-test) and negative encoding cells (mean AUC = 0.480;  $p = 0.19$ ; t-test).

Second, the analysis in Fig.5ab pooled all *offer value* cells, including those that were not tuned in the time period immediately following the offer (these cells became tuned later in the trial). One concern might be that these cells effectively added noise and thus obfuscated the signal of interest here. Thus I repeated this analysis including in the pool only *offer value* cells that were significantly tuned in the 150–400 ms following the offer (see Experimental Procedures). The results (Fig.S2ab) confirmed those illustrated in Fig.5ab for both positive encoding cells (mean AUC = 0.502;  $p = 0.83$ ; t-test) and negative encoding cells (mean AUC = 0.474;  $p = 0.18$ ; t-test).

Third, the analysis in Fig.5ab averaged traces across offer types and then across cells. For positive (negative) encoding cells, this procedure could overweight high-value (low-value) offer types within each neuron, or could over-emphasize cells with higher firing rates, effectively reducing the statistical power of the analysis. I controlled for this issue as follows. In any time window, the encoding of value in OFC is linear and range adapting (Padoa-Schioppa, 2009). In formulas,  $\varphi = \varphi_0 + \Delta\varphi * V / \Delta V$ , where  $\varphi$  is the firing rate,  $\varphi_0$  is the baseline activity,  $\Delta\varphi$  is the activity range,  $V$  is the encoded value and  $\Delta V$  is the value range. (Note that for offer value cells in the experiment the minimum value  $V_0$  was always zero, so that  $\Delta V = V_{max}$ .) Here I am interested in small fluctuations on  $\varphi$  related to an endogenous factor (i.e., whether the encoded juice was eventually chosen). It is reasonable to assume that, if they exist, such fluctuations are proportional to the firing rate. Thus the formula can be re-written as follows  $\varphi * (1 + \varepsilon) = \varphi_0 + \Delta\varphi * V / \Delta V$ , where  $\varepsilon$  is the fluctuation. This makes it clear that  $\varepsilon$  depends on both  $V / \Delta V$  and  $\Delta\varphi$ . In essence, the analysis of Fig.5ab aims at studying  $\varepsilon$  by averaging neuronal traces across offer types (i.e., across values) and across cells. However, by simply averaging the firing rates, the analysis overweighs offers with large values (because  $\varepsilon$  increases with  $V / \Delta V$ ) and cells with large activity range (because  $\varepsilon$  increases with  $\Delta\varphi$ ). Thus to increase the resolution of the analysis one would like to rescale the firing rate (and  $\varepsilon$ ). I did so in two steps. First, I rescaled  $\varphi \rightarrow \varphi' = (\varphi - \varphi_0) * \Delta V / V$  (value range rescaling). Second, I rescaled  $\varphi' \rightarrow \varphi'' = \varphi' / \Delta\varphi$  (activity range rescaling). In these transformations, I used for  $\varphi_0$  the average activity in the 500 ms before the offer and for  $\Delta\varphi$  the activity range in the 500 ms after the offer (post-offer time window). None of these variants of the analysis affected the results for the mean AUC (all  $p > 0.3$ ; t-test).

Fourth, because of choice hysteresis, the signal of interest here might be examined with higher resolution by separating trials depending on the outcome of the previous trial. I thus conducted a variant of the analysis as follows. For each *offer value* cell, I separated trials into three groups depending on the outcome of the previous trial (trials E•, O• and X•). For each group of trials, I identified offer types for which offers were split, and I calculated the two neuronal traces and the AUC as in Fig.5ab. The results are shown in Fig.S2cd. For positive encoding neurons, the results confirmed those of Fig.5a (mean AUC = 0.494;  $p = 0.25$ , t-test). For negative encoding cells, the population AUC was significantly below 0.5 (mean AUC = 0.472;  $p < 0.02$ , t-test). Note that this departure is in the direction predicted by the hypothesis that fluctuations of *offer value* cells drive near-indifference decisions. However, when I examined data from individual animals, the effect was significant only for one monkey (L, mean AUC = 0.44;  $p < 0.005$ , t-test) and not for the other (V, mean AUC = 0.494;  $p = 0.64$ , t-test). In conclusion, the evidence that choices are driven by fluctuations of *offer value* cells was at best tentative.

In another analysis, I specifically examined trials in which the animal chose one drop of the preferred juice (1A). I focused on neurons encoding *offer value A* (the preferred juice) and I divided trials depending on the quantity of juice B offered in alternative to 1A. The rationale for this analysis was as follows. In principle, one can hypothesize that choice variability reflects stochastic fluctuations in the subjective value of any particular juice. In particular, the subjective value of 1A, represented by the activity of *offer value A* cells, might randomly fluctuate from trial to trial. All other things equal, one would expect that positive fluctuations in the activity of *offer value A* cells would facilitate choices of juice A. By the same token, one would expect that the activity of *offer value A* cells, conditional on the animal choosing 1A, would be enhanced (by chance) when the alternative offer is more desirable.

This argument would predict that the activity of positive encoding *offer value A* cells would be higher when 1A is chosen against large amounts of juice B compared to when 1A is chosen against small amounts of juice B. To test the prediction, I divided trials in easy (offer types for which the animal always chose the same option) and split (offer types for which the animal split its decisions between two options). Contrary to the prediction, the activity recorded for the two groups of trials was indistinguishable throughout the 1 s following the offer (Fig.S2e). Similar results were found for negative encoding cells (Fig.S2f).

### Time course of choice hysteresis and its relation to predictive activity

As noted in Fig.4c, choice hysteresis largely dissipated after one trial. To quantify its time course more precisely, I constructed a logistic model taking into consideration the five trials preceding the current one, as follows:

$$\text{choice B} = 1 / (1 + e^{-X})$$

$$X = a_0 + a_1 \log (\#B / \#A) + \sum_{k=1:5} b_k (\delta_{n-k, B} - \delta_{n-k, A}) \quad (6)$$

Variable  $\delta_{n-k, J} = 1$  if the animal chose and received juice J in trial (n-k), and 0 otherwise. Across the population, I found the following values: median ( $b_{n-1}/a_1$ ) = 0.128,  $p < 10^{-10}$ ; median ( $b_{n-2}/a_1$ ) = 0.037,  $p < 10^{-10}$ ; median ( $b_{n-3}/a_1$ ) = 0.029,  $p < 10^{-10}$ ; median ( $b_{n-4}/a_1$ ) = 0.013,  $p < 10^{-5}$ ; median ( $b_{n-5}/a_1$ ) = 0.013,  $p < 10^{-4}$  (all Wilcoxon sign test; Fig.S5e). These results confirm that choice hysteresis was predominantly related to the previous trial. At the same time, there was also an effect that persisted for several trials and reached a plateau level of  $\sim 0.013$ . This plateau might be due to the fact that the animals' preferences often drifted toward the preferred juice during the course of the session, probably due to reduced thirst.

In light of this result, one concern might be whether the residual predictive activity of *chosen juice* cells observed in Fig.S5d is in fact related to the persistence of choice hysteresis past the previous trial. To examine this issue, I repeated the analysis of firing rates residuals taking into consideration the preceding two trials. Specifically, I constructed the following logistic model:

$$\text{choice E} = 1 / (1 + e^{-X})$$

$$X = a_0 + a_1 \log (\#E / \#O) + a_2 (\delta_{n-1, E} - \delta_{n-1, O}) + a_3 (\delta_{n-2, E} - \delta_{n-2, O}) + a_4 \varphi_{\text{residual 2}} \quad (7)$$

For each *chosen juice* cell,  $\varphi_{\text{residual 2}}$  is the residual firing rate remaining after the bilinear regression of the raw firing rate  $\varphi$  onto variables  $(\delta_{n-1, E} - \delta_{n-1, O})$  and  $(\delta_{n-2, E} - \delta_{n-2, O})$ . The null hypothesis corresponds to  $a_4/a_1 = 0$ . Across the population, the median of the distribution was  $m = 0.002$  ( $p < 0.02$ , all Wilcoxon signed-rank test; Fig.S5f). Notably, this measure is almost identical to that obtained for  $a_3/a_1$  in model 4, which considered only the previous trial (n-1). This observation suggests that the residual predictive activity of *chosen juice* cells is not due to the persistent plateau effect or to drifting preferences. In other words, baseline fluctuations in the activity of *chosen juice* cells appear to explain a portion of choice variability above and beyond that explained by behavioral analyses alone.

### Overshooting of *chosen value* cells: control for variable *total value*

Consider offers [aA:bB], where a and b are quantities of juices A and B, respectively. The experimental design and all the analyses were based on two assumptions. First, it was assumed that the choice pattern (i.e., the percent of trials in which the animal chose juice B) depended only on the quantity ratio b/a. Second, it was assumed that value functions were linear. In other words, indicating with  $V(qX)$  the value assigned to a quantity q of juice X, it was assumed that  $V(qX) = qV(X)$ . If this is

the case, choice patterns can be described in one dimension as a function of the quantity ratio  $b/a$ . Then the relative value ( $\rho$ ) is defined as the quantity ratio that makes the animal indifferent between the two juices:  $V(A) = \rho V(B)$ .

The activity overshooting of *chosen value* cells can essentially be described as follows. Restricting the analysis to trials in which the animal chose 1A over  $qB$  ( $1A \blacktriangleright qB$ ), the activity of *chosen value* cells recorded in the time window 150-400 ms after the offer increased as a function of  $q$ . As discussed in the main text, the overshooting can be explained qualitatively if one assumes that the relative value of the two juices fluctuated from trial to trial. In the following, I refer to this hypothesis as *chosen value* cells encoding the variable *chosen value*  $X$ , which is the same as the variable *chosen value* corrected for fluctuations of  $\rho$  (see below). However, an alternative explanation is that *chosen value* cells actually encode the variable *total value* (defined as the sum of the two offer values, which increases as a function of  $q$ ). These two hypotheses were contrasted as follows.

To compute the variable *chosen value*  $X$ , it is necessary to specify the probability distribution for the relative value  $\rho$ . In the following analyses, I assumed that, once controlled for choice hysteresis, choice variability was entirely due to fluctuations of  $\rho$ . If this is true, then the probability distribution for  $\rho$  can be derived from the choice pattern. Choice patterns in the experiments were well fitted with a normal sigmoid (probit function) in log space (typical  $R^2 > 0.95$ ). If choice variability is entirely due to fluctuations of  $\rho$ , the underlying normal distribution can be viewed as a probability distribution for the variable  $x = \log \rho$ . Thus the probability distribution for  $\rho$  is  $N(x(\rho), \mu, \sigma) dx/d\rho = N(\log \rho, \mu, \sigma) 1/\rho$ . On this basis, one can compute the variable *chosen value*  $X$  in each trial, as follows.

First consider one trial in which the animal chose 1A over  $qB$ . As noted in the main text, if values are expressed in units of juice B, Eq.5 implies *chosen value*  $X = \rho \geq q$ . Now consider many trials in which the animal chose 1A over  $qB$ . On average, the variable *chosen value*  $X_{1A \blacktriangleright qB} = \langle \rho \rangle_{\rho \geq q}$  is equal to:

$$\langle \rho \rangle_{\rho \geq q} = \frac{\int_q^{\infty} N(\log \rho, \mu, \sigma) \cdot 1/\rho \cdot \rho d\rho}{\int_q^{\infty} N(\log \rho, \mu, \sigma) \cdot 1/\rho d\rho} = \frac{\int_{\log(q)}^{\infty} N(x, \mu, \sigma) \cdot e^x dx}{\int_{\log(q)}^{\infty} N(x, \mu, \sigma) dx} \quad (8)$$

More generally, when the animal chose  $a$  drops of juice A over  $b$  drops of juice B ( $aA \blacktriangleright bB$ ), *chosen value*  $X_{aA \blacktriangleright bB} = a$  *chosen value*  $X_{1A \blacktriangleright b/aB}$ . This can be calculated substituting  $b/a$  for  $q$  in Eq.8.

Now consider trials in which the animal chose  $qB$  over 1A. To proceed formally as when the animal chose juice A, I define  $\xi$  such that  $B = \xi A$  and  $x = \log \xi$ . In this case, the probability distribution for  $x$  is  $N(x, -\mu, \sigma)$  and the probability distribution for  $\xi$  is  $N(\log \xi, -\mu, \sigma) 1/\xi$ . Thus the variable *chosen value*  $X_{1B \blacktriangleright 1/q A} = \langle \xi \rangle_{\xi \geq 1/q}$  is equal to:

$$\langle \xi \rangle_{\xi \geq 1/q} = \frac{\int_{1/q}^{\infty} N(\log \xi, -\mu, \sigma) \cdot 1/\xi \cdot \xi d\xi}{\int_{1/q}^{\infty} N(\log \xi, -\mu, \sigma) \cdot 1/\xi d\xi} = \frac{\int_{-\log(q)}^{\infty} N(x, -\mu, \sigma) \cdot e^x dx}{\int_{-\log(q)}^{\infty} N(x, -\mu, \sigma) dx} \quad (9)$$

Importantly, Eq.9 expresses the *chosen value X* in units of juice A. To express all chosen values in units of juice B, I multiply for the average conversion factor  $\langle \rho \rangle$ . In conclusion, one obtains for each trial type a measure of the variable *chosen value X*.

The results of this calculation are illustrated for one representative session in Fig.S7a. Away from the indifference point, *chosen value X* is nearly identical to *chosen value*. However, close to the indifference point, *chosen value X* is generally higher than *chosen value*. Fig.S7bc also illustrate the fact that although variables *chosen value*, *chosen value X*, and *total value* are highly correlated, they are distinguishable. In particular, it can be noted that *total value* and *chosen value X* are most correlated near the indifference point, but appreciably different away from the indifference point.

To contrast the explanatory power of variables *chosen value X* and *total value*, I specifically examined the 150-400 ms after the offer and I restricted the analysis to neurons from Exp.1 that were significantly tuned in this time window (positive encoding). For an accurate measure of *chosen value X*, I removed the variability due to choice hysteresis by dividing trials into three groups depending on the outcome of the previous trial. The three groups of trials A•, B• and X• were analyzed separately, with all the trials included in the analysis. For each cell, for each group of trials and for each trial type, I computed the variables *chosen value X* and *total value* and I averaged the activity across trials. Then I performed a linear regression of the neuronal firing rate onto each variable, from which I obtained the two  $R^2$ . (Note that these procedures are essentially the same as used in previous studies (Padoa-Schioppa and Assad, 2006).) As illustrated in Fig.S7d, the  $R^2$  obtained for *chosen value X* was generally higher than that obtained for *total value* ( $p < 0.01$ , Kruskal-Wallis test). This result indicates that the explanatory power of *chosen value X*, corresponding to the hypothesis that the overshooting of *chosen value* cells is due to fluctuations of relative value  $\rho$ , is significantly higher than that of *total value*.

#### The overshooting of *chosen value* cells is independent of choice hysteresis

This study describes two neuronal phenomena seemingly related to choice variability: predictive activity of *chosen juice* cells and activity overshooting of *chosen value* cells. One important question is whether these phenomena are different manifestations of the same underlying source of variability or, alternatively, whether activity overshooting and predictive activity are mutually independent. To examine this issue, I took advantage of the fact that predictive activity was largely accounted for by the outcome of the previous trial (choice hysteresis). To assess whether the choice variability related to the activity overshooting added to, or was redundant with, that related to the choice hysteresis, I repeated the analyses of *chosen value* cells described in Fig.7a while controlling for the outcome of the previous trial. The analysis included only trials in which the animal chose one drop of the preferred juice (1A) against various amounts of the other juice ( $qB$ ). These trials were divided into three groups depending on the outcome of the previous trial (trials A•, B• and X•). Each group of trials was further divided depending on whether the offer type was easy or split (see Experimental Procedures). As illustrated in Fig.8a-c, the activity of *chosen value* cells presented the overshooting even when the previous trial's outcome was controlled for.

For each *chosen value* cell and for each group of trials (A•, B• and X•), I also performed the ROC analysis and computed the AUC. The results obtained pooling trials (insert in Fig.7a) held true separately for each group of trials (all  $p < 0.05$ , t-test; inserts in Fig.8a-c). I also noted that the mean AUC obtained for each group of trials was quantitatively similar to that obtained pooling all trials (pooling trials, mean AUC = 0.526; for A•, B• and X• trials, mean AUC = 0.531, 0.526 and 0.530, respectively). These measures suggest that the activity overshooting is independent of the outcome of the previous trial.

To further test the relation between activity overshooting and choice hysteresis, I compared for each *chosen value* cell the AUC obtained for A• trials and that obtained for B• trials (Fig.8d). Two important results emerged from this analysis. First, although each measure was rather noisy, the two measures were significantly correlated across the population (correlation coefficient = 0.22,  $p < 0.01$ ). This correlation is important because it indicates that the AUC is a reproducible measure for any given *chosen value* cell (Britten et al., 1996). Second, the difference between the two AUC obtained for the two groups of trials, examined at the population level, was statistically indistinguishable from zero ( $p = 0.48$ , t-test; insert in Fig.8d). This result stands as strong evidence that activity overshooting was independent of choice hysteresis. Indeed, if even a portion of the activity overshooting had been redundant with choice hysteresis, the AUC measured in A• trials would be overall smaller than that measured in B• trials – contrary to the observation. I repeated this analysis comparing A• trials and X• trials (Fig.8e) and, separately, B• trials and X• trials (Fig.8f). The results reinforced the conclusions already drawn. First, in both cases there was a significant correlation between the AUC measured for any given cell in different groups of trials (both  $p < 0.003$ ). Second, in both cases the difference between the two measures of AUC obtained for the two groups of trials was statistically indistinguishable from zero (inserts in Fig.8ef). In conclusion, the activity overshooting of *chosen value* cells is independent of choice hysteresis.

### Supplemental References

- 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.
- Padoa-Schioppa, C. (2009). Range-adapting representation of economic value in the orbitofrontal cortex. *J Neurosci* 29, 14004-14014.
- Padoa-Schioppa, C., and Assad, J.A. (2006). Neurons in orbitofrontal cortex encode economic value. *Nature* 441, 223-226.
- Sayer, R.J., Friedlander, M.J., and Redman, S.J. (1990). The time course and amplitude of EPSPs evoked at synapses between pairs of CA3/CA1 neurons in the hippocampal slice. *J Neurosci* 10, 826-836.
- So, N.Y., and Stuphorn, V. (2010). Supplementary eye field encodes option and action value for saccades with variable reward. *J Neurophysiol* 104, 2634-2653.