Published in final edited form as:
Neuron. 2013 December 4; 80(5): . doi:10.1016/j.neuron.2013.09.013.

Neuronal origins of choice variability in economic decisions

Camillo Padoa-Schioppa[1,2,3]

[1]Department of Anatomy and Neurobiology, Washington University in St Louis, St Louis, MO 63110

[2]Department of Economics, Washington University in St Louis, St Louis, MO 63110

[3]Department of Biomedical Engineering, Washington University in St Louis, St Louis, MO 63110

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 pre-offer 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 Buchel, 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

Correspondence: camillo@wustl.edu.

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, while *chosen juice* may correspond to neurons in LIP. Indeed, the former seem to represent the main input to the decision process, while 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 Exp.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 (Fig. 1ab). In Exp. 2, the procedures were very similar except that 3 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* (Fig. 1c), *chosen value* (Fig. 1d) and *chosen juice* (Fig. 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 non-zero 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 $\Delta R^2 = R^2_{offer\ value} - R^2_{chosen\ value}$, which ranged from −1 to +1. For a response perfectly explained by *offer value* (*chosen value*) and poorly explained by *chosen value* (*offer value*), ΔR^2 is close to +1 (−1). If *offer value* and *chosen value* are two poles of a continuum, the distribution of ΔR^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 ΔR^2 should be bimodal with a dip close to zero. As illustrated in Fig. 1fg, the distribution obtained for ΔR^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* vs. *chosen juice* and *chosen value* vs. *chosen juice*). In both cases, the distribution for ΔR^2 was clearly bimodal (both $p<10^{-10}$, Hartigan's dip test; Fig. 1h–k). 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.

Second, to assess whether different neurons encoded different variables, I first examined data from Exp.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 Fig. S1). 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; Fig. S1a). Conversely, for each pair of time windows, cells with consistent classification were significantly more frequent than expected by chance (Fig. S1b; all $p<10^{-10}$, t-test). Data from Exp.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.

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 Exp.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.

Fig. 2 illustrates the average activity profile obtained for each neuronal population. Importantly, inspection of Fig. 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 (Fig. 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 Fig. 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 hysteresis”. One example session is illustrated in Fig. 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 (Fig. 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 Exp.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 (Fig. 4c).

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

$$\begin{aligned} \text{choice B} &= 1/(1+e^{-X}) \\ X &= a_0 + a_1 \log(\#B/\#A) + a_2(\delta_{n-1,B} - \delta_{n-1,A}) \end{aligned} \quad (1)$$

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 (for example with receipt of the third juice in Exp.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, while 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 (Fig. 4d). The median of the distribution m = 0.124 was significantly greater than zero ($p<10^{-10}$, 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.

Fluctuations in the activity of offer value cells do not explain choice variability

In the framework of Fig. 3b, *offer value* cells (Fig. 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. 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 Fig. 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 (Fig.

5b) and in several different variants of this analysis (Supplemental Procedures p.S1, Fig. 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 et al (2007). This analysis focused on the 500 ms following the offer. I constructed the following logistic model:

$$\text{choice E} = 1/(1+e^{-X})$$
$$X = a_0 + a_1 \log(\#E/\#O) + a_2(\delta_{n-1,E} - \delta_{n-1,O}) + a_3 \phi_{\text{residual}} \quad (2)$$

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 (φ) onto the variable *offer value E*. Other notations were as in Eq.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 Exp.2 contributed each with two data points). Fig. 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 (Fig. 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 the decision was 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 (Fig. 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 Fig. 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 Fig. 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 roughly 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).

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 Fig. 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 hard (light blue and light red lines in Fig. 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 Fig. 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. 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, the indifference point typically corresponded to lower juice quantities (see Fig. 1b). Thus the difference in neural activity between easy decisions (dark blue in Fig. 6a) and split decisions (light blue in Fig. 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 (Fig. 6b). 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 Fig. 6b 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 pre-offer activity of *chosen juice* cells and near-indifference decisions was also tested with a logistic analysis. I constructed the following model:

$$\begin{aligned} \text{choice}\,\mathrm{E} &= 1/(1+\mathrm{e}^{-\mathrm{X}}) \\ \mathrm{X} &= \mathrm{a}_0+\mathrm{a}_1\log(\#\mathrm{E}/\#\mathrm{O})+\mathrm{a}_2\,\phi \end{aligned} \quad (3)$$

For each *chosen juice* cell, φ was the firing rate in the 500 ms preceding the offer. Fig. S5c illustrates the distribution for a_2/a_1 obtained across the population. The median of the distribution m = 0.005 was significantly greater than zero ($p<10^{-8}$, Wilcoxon signed-rank test). In essence, this means that when the pre-offer activity of *chosen juice* cells increased by 1 spike, the animal made its choice as though the quantity of the encoded juice was multiplied by a factor of ≈ 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 pre-offer 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 Fig. 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 (Fig. 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 reflecting 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 controlled for, the presence of residual predictive activity (Fig. 6d–g) provided evidence in favor of H1. For a quantitative assessment of residual predictive activity, I constructed the following logistic model:

$$\begin{gathered}\text{choice E}=1/(1+e^{-X})\\ X=a_0+a_1\log(\#E/\#O)+a_2(\delta_{n-1,E}-\delta_{n-1,O})+a_3\,\phi_{\text{residual}}\end{gathered}\quad(4)$$

For each *chosen juice* cell, $\varphi_{\text{residual}}$ was the residual firing rate remaining after the linear regression of the raw firing rate φ onto the variable $(\delta_{n-1,E} - \delta_{n-1,O})$. The null hypothesis corresponded to $a_3/a_1 = 0$. Fig. 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 greater than zero ($p<0.02$, Wilcoxon signed-rank test). In other words, trial-by-trial fluctuations in the pre-offer activity of *chosen juice* cells were significantly correlated with the decision of the animal, even when the outcome of the previous trials is controlled for.

In conclusion, predictive activity reflected additional sources of stochasticity in addition to the tail activity from the previous trial.

Activity overshooting of chosen value cells

I now turn to *chosen value* cells (Fig. 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▶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 (Fig. 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 (insert in Fig. 7a). 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 Fig. 7a was very robust (Fig. 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 (Fig. 7b). Restricting the analysis to cells that were significantly tuned yielded similar results (Fig. 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 (Fig. 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; Fig. 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 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 Fig. 7a and Fig. 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 Fig. S2e and as a chosen value in Fig. 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, while *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 Fig. 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 (ρ). Given two goods, ρ can change over relatively long periods of time, for example due to changes in internal motivation (Padoa-Schioppa and Assad, 2006). More subtly, ρ could fluctuate on a trial-by-trial basis. Interestingly, stochastic fluctuations of ρ would induce overshooting in the activity of *chosen value* cells similar to that shown in Fig. 7a. To appreciate this point, consider trials in which the animal chose between juices A and B. If value functions are linear, ρ is the quantity ratio that makes the animal indifferent between the two juices:

$$V(A)=\rho V(B) \quad (5)$$

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

The activity overshooting of *chosen value* cells can thus be explained by fluctuations of ρ. 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 ρ was significantly higher than that of *total value* ($p<0.01$, Kruskal-Wallis test; p.S3; Fig. S7).

In summary, evidence suggested that the activity overshooting observed in *chosen value* cells reflected trial-by-trial fluctuations in ρ. 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 (Fig. 8, Supplement p.S5). 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.

I. Variables *offer value*, *chosen value* and *chosen juice* were encoded by different classes 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.

II. 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 −0.5 (CPs) 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.

III. 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).

IV. 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 micro-stimulation.

V. 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 ρ. Under this interpretation, an important question relates to how ρ is instantiated at the neuronal level. In the framework of Fig. 3b, ρ 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 Fig. 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 non-spatial. 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 towards. 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 (Fig. 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 Exp.1 (931 cells), animals chose in each session between two juices. In Exp.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 4 primary time windows: post-offer (0.5 s after the offer), late delay (0.5 – 1 s after the offer), pre-juice (0.5 s before juice delivery) and post-juice (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 Exp.1; 340 from Exp. 2).

Statistical analyses

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. The label “easy” refers to 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▶qB), 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 Fig. 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 Fig. 5b. The results reported in Fig. 5ab 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 Fig. 6a, I identified offer types in which decisions were split and 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 Fig. 6c and Fig. 7a, I focused on trials in which the animal chose 1 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 Fig. 7ce and Fig. S2ab 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 Fig. 2a–d), 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 Exp.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”.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

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

References

Andersen RA, Cui H. Intention, action planning, and decision making in parietal-frontal circuits. Neuron. 2009; 63:568–583. [PubMed: 19755101]

Bennur S, Gold JI. Distinct representations of a perceptual decision and the associated oculomotor plan in the monkey lateral intraparietal area. J Neurosci. 2011; 31:913–921. [PubMed: 21248116]

Bisley JW, Goldberg ME. Attention, intention, and priority in the parietal lobe. Annu Rev Neurosci. 2010; 33:1–21. [PubMed: 20192813]

Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD. The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. Psychological Review. 2006; 113:700–765. [PubMed: 17014301]

Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA. A relationship between behavioral choice and the visual responses of neurons in macaque MT. Vis Neurosci. 1996; 13:87–100. [PubMed: 8730992]

Buckley MJ, Mansouri FA, Hoda H, Mahboubi M, Browning PG, Kwok SC, Phillips A, Tanaka K. Dissociable components of rule-guided behavior depend on distinct medial and prefrontal regions. Science. 2009; 325:52–58. [PubMed: 19574382]

Camille N, Griffiths CA, Vo K, Fellows LK, Kable JW. Ventromedial frontal lobe damage disrupts value maximization in humans. J Neurosci. 2011; 31:7527–7532. [PubMed: 21593337]

Cohen MR, Newsome WT. Estimates of the contribution of single neurons to perception depend on timescale and noise correlation. J Neurosci. 2009; 29:6635–6648. [PubMed: 19458234]

Gallagher M, McMahan RW, Schoenbaum G. Orbitofrontal cortex and representation of incentive value in associative learning. J Neurosci. 1999; 19:6610–6614. [PubMed: 10414988]

Gold JI, Shadlen MN. The neural basis of decision making. Annu Rev Neurosci. 2007; 30:535–574. [PubMed: 17600525]

Haefner RM, Gerwinn S, Macke JH, Bethge M. Inferring decoding strategies from choice probabilities in the presence of correlated variability. Nat Neurosci. 2013; 16:235–242. [PubMed: 23313912]

Hare TA, Schultz W, Camerer CF, O’Doherty JP, Rangel A. Transformation of stimulus value signals into motor commands during simple choice. P Natl Acad Sci USA. 2011; 108:18120–18125.

Hunt LT, Kolling N, Soltani A, Woolrich MW, Rushworth MFS, Behrens TEJ. Mechanisms underlying cortical activity during value-guided choice. Nature Neuroscience. 2012; 15:470–U169.

Kable JW, Glimcher PW. The neural correlates of subjective value during intertemporal choice. Nat Neurosci. 2007; 10:1625–1633. [PubMed: 17982449]

Kang I, Maunsell JH. Potential confounds in estimating trial-to-trial correlations between neuronal response and behavior using choice probabilities. J Neurophysiol. 2012; 108:3403–3415. [PubMed: 22993262]

Kennerley SW, Dahmubed AF, Lara AH, Wallis JD. Neurons in the frontal lobe encode the value of multiple decision variables. J Cogn Neurosci. 2009; 21:1162–1178. [PubMed: 18752411]

Kohn A. Visual adaptation: physiology, mechanisms, and functional benefits. J Neurophysiol. 2007; 97:3155–3164. [PubMed: 17344377]

Krajbich I, Armel C, Rangel A. Visual fixations and the computation and comparison of value in simple choice. Nat Neurosci. 2010; 13:1292–1298. [PubMed: 20835253]

Levy I, Snell J, Nelson AJ, Rustichini A, Glimcher PW. Neural representation of subjective value under risk and ambiguity. J Neurophysiol. 2010; 103:1036–1047. [PubMed: 20032238]

Lim SL, O'Doherty JP, Rangel A. The decision value computations in the vmPFC and striatum use a relative value code that is guided by visual attention. J Neurosci. 2011; 31:13214–13223. [PubMed: 21917804]

Newsome WT. The King Solomon lectures in neuroethology. Deciding about motion: linking perception to action. J Comp Physiol [A]. 1997; 181:5–12.

Nienborg H, Cohen MR, Cumming BG. Decision-related activity in sensory neurons: correlations among neurons and with behavior. Annu Rev Neurosci. 2012; 35:463–483. [PubMed: 22483043]

Padoa-Schioppa C. Range-adapting representation of economic value in the orbitofrontal cortex. J Neurosci. 2009; 29:14004–14014. [PubMed: 19890010]

Padoa-Schioppa C. Neurobiology of economic choice: a good-based model. Annu Rev Neurosci. 2011; 34:333–359. [PubMed: 21456961]

Padoa-Schioppa C, Assad JA. Neurons in orbitofrontal cortex encode economic value. Nature. 2006; 441:223–226. [PubMed: 16633341]

Padoa-Schioppa C, Assad JA. The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. Nat Neurosci. 2008; 11:95–102. [PubMed: 18066060]

Padoa-Schioppa C, Jandolo L, Visalberghi E. Multi-stage mental process for economic choice in capuchins. Cognition. 2006; 99:B1–B13. [PubMed: 16043168]

Peters J, Buchel C. Overlapping and distinct neural systems code for subjective value during intertemporal and risky decision making. J Neurosci. 2009; 29:15727–15734. [PubMed: 20016088]

Plassmann H, O'Doherty J, Rangel A. Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. J Neurosci. 2007; 27:9984–9988. [PubMed: 17855612]

Roesch MR, Olson CR. Neuronal activity in primate orbitofrontal cortex reflects the value of time. J Neurophysiol. 2005; 94:2457–2471. [PubMed: 15958600]

Roitman JD, Shadlen MN. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J Neurosci. 2002; 22:9475–9489. [PubMed: 12417672]

Rudebeck PH, Murray EA. Dissociable effects of subtotal lesions within the macaque orbital prefrontal cortex on reward-guided behavior. J Neurosci. 2011; 31:10569–10578. [PubMed: 21775601]

Shadlen MN, Britten KH, Newsome WT, Movshon JA. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. Journal of Neuroscience. 1996; 16:1486–1510. [PubMed: 8778300]

Shadlen, MN.; Kiani, R.; Hanks, TD.; Churchland, AK. Neurobiology of decision making: an intentional framework. In: Engel, C.; Singer, W., editors. Better than conscious? Decision making, the human mind, and implications for institutions. Cambridge, MA: MIT Press; 2008.

Shadlen MN, Newsome WT. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. J Neurophysiol. 2001; 86:1916–1936. [PubMed: 11600651]

Simmons JM, Richmond BJ. Dynamic changes in representations of preceding and upcoming reward in monkey orbitofrontal cortex. Cereb Cortex. 2008; 18:93–103. [PubMed: 17434918]

Soltani A, De Martino B, Camerer C. A range-normalization model of context-dependent choice: a new model and evidence. PLoS Comput Biol. 2012; 8:e1002607. [PubMed: 22829761]

Solway A, Botvinick MM. Goal-directed decision making as probabilistic inference: a computational framework and potential neural correlates. Psychological Review. 2012; 119:120–154. [PubMed: 22229491]

Sugrue LP, Corrado GS, Newsome WT. Choosing the greater of two goods: neural currencies for valuation and decision making. Nat Rev Neurosci. 2005; 6:363–375. [PubMed: 15832198]

Summerfield C, Tsetsos K. Building bridges between perceptual and economic decision-making: neural and computational mechanisms. Front Neurosci. 2012; 6:70. [PubMed: 22654730]

Wang XJ. Probabilistic decision making by slow reverberation in cortical circuits. Neuron. 2002; 36:955–968. [PubMed: 12467598]

Wang XJ. Decision making in recurrent neuronal circuits. Neuron. 2008; 60:215–234. [PubMed: 18957215]

West EA, DesJardin JT, Gale K, Malkova L. Transient inactivation of orbitofrontal cortex blocks reinforcer devaluation in macaques. J Neurosci. 2011; 31:15128–15135. [PubMed: 22016546]

Williams ZM, Elfar JC, Eskandar EN, Toth LJ, Assad JA. Parietal activity and the perceived direction of ambiguous apparent motion. Nat Neurosci. 2003; 6:616–623. [PubMed: 12730699]

Wyart V, Tallon-Baudry C. How ongoing fluctuations in human visual cortex predict perceptual awareness: baseline shift versus decision bias. J Neurosci. 2009; 29:8715–8725. [PubMed: 19587278]

Yang T, Shadlen MN. Probabilistic reasoning by neurons. Nature. 2007; 447:1075–1080. [PubMed: 17546027]

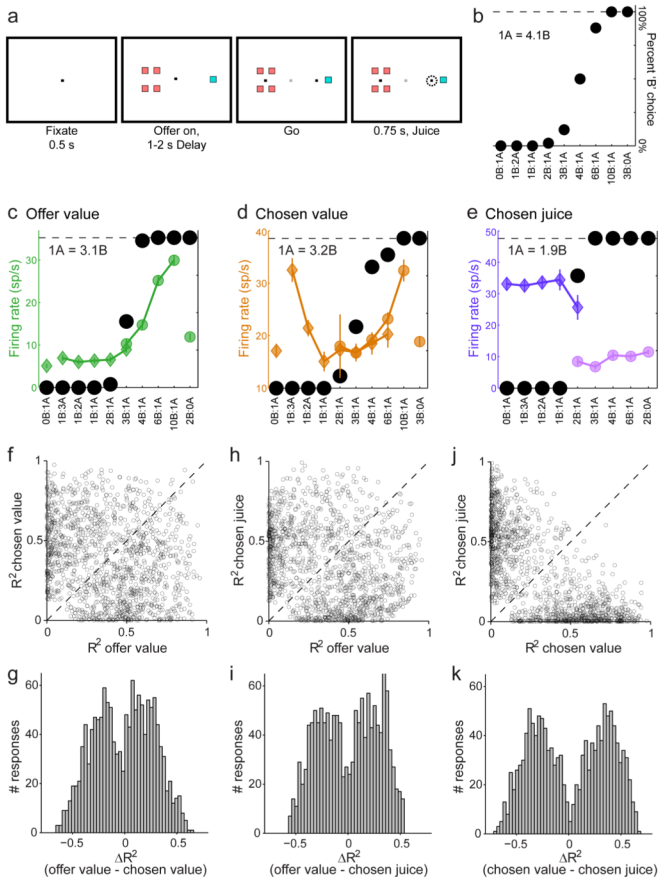

Figure 1.

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. **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. **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. **fg.** 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 obtained from linear regressions against variables *offer value* and *chosen value*. In panel (f), the two R^2 are plotted against each other. Panel (g) illustrates the distribution obtained for ΔR^2. **hi.** *Offer value* versus *chosen juice*. **jk.** *Chosen value* versus *chosen juice*. (See also Fig. S1.)

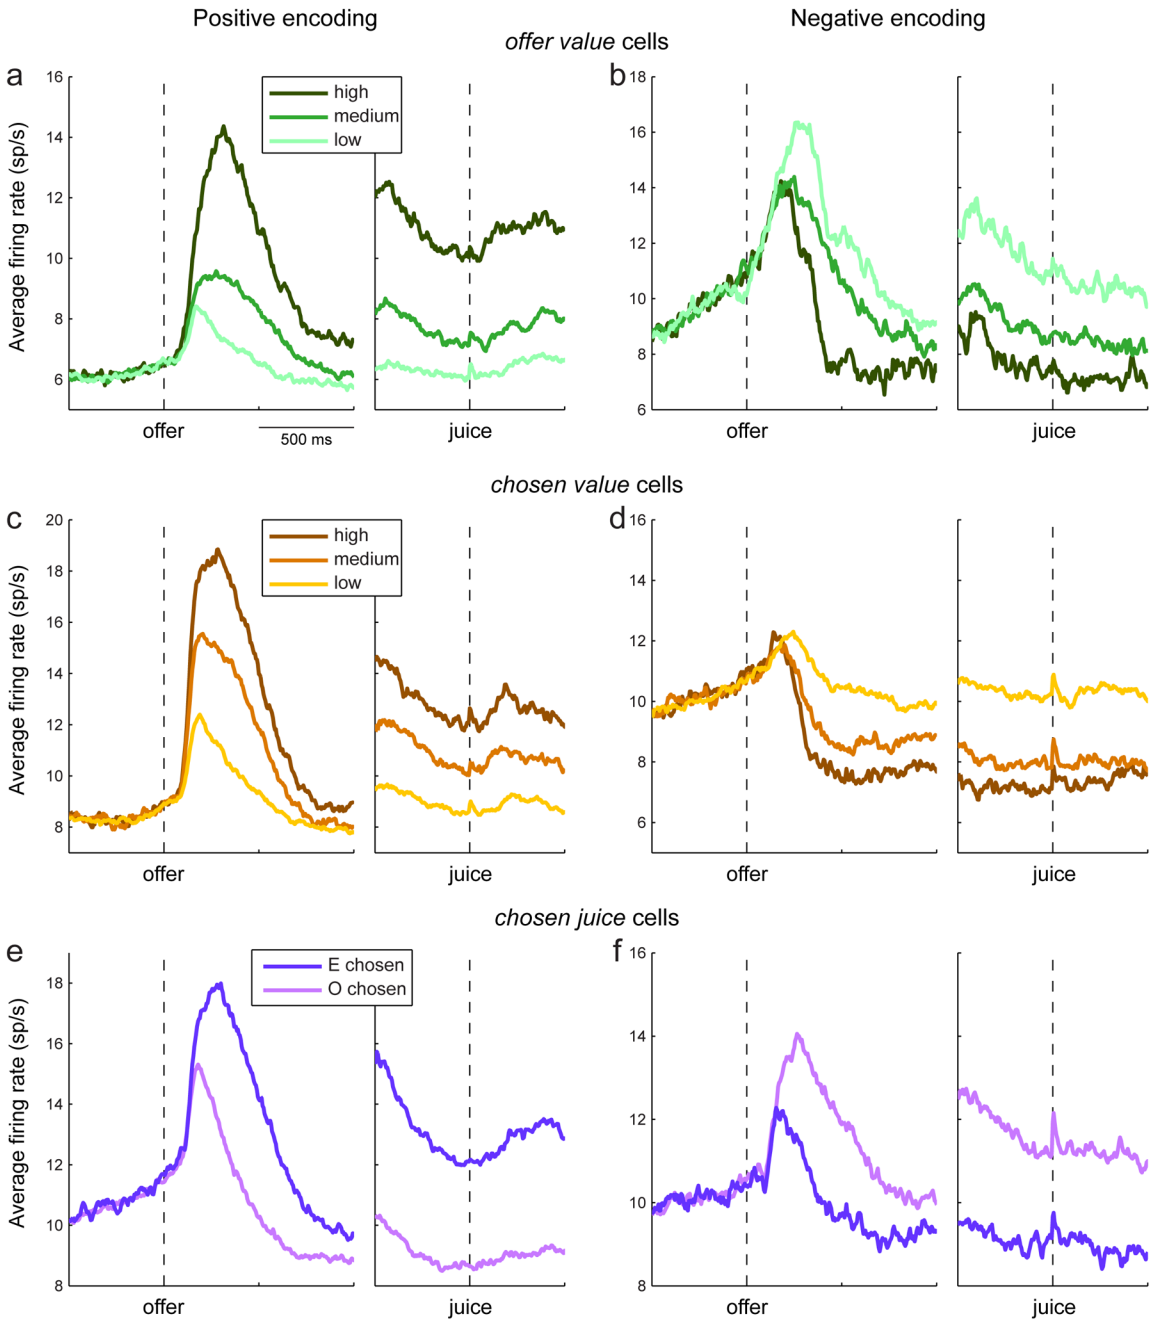

Figure 2.

Average activity profiles. **ab.** 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. Panels a/b show the activity for neurons with positive/negative encoding (188/57 cells). **cd.** Average population activity for *chosen value* cells. Panels c/d show the activity for neurons with positive/negative encoding (161/112 cells). For each cell, trials were divided into three tertiles based on the chosen value (high, medium and low). **ef.** Average population activity for *chosen juice* cells. The figure includes only data from Exp.2, for which positive/negative encoding could be established (see main text). Panels e/f show the average activity for

neurons with positive/negative encoding (96/50 cells). 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).

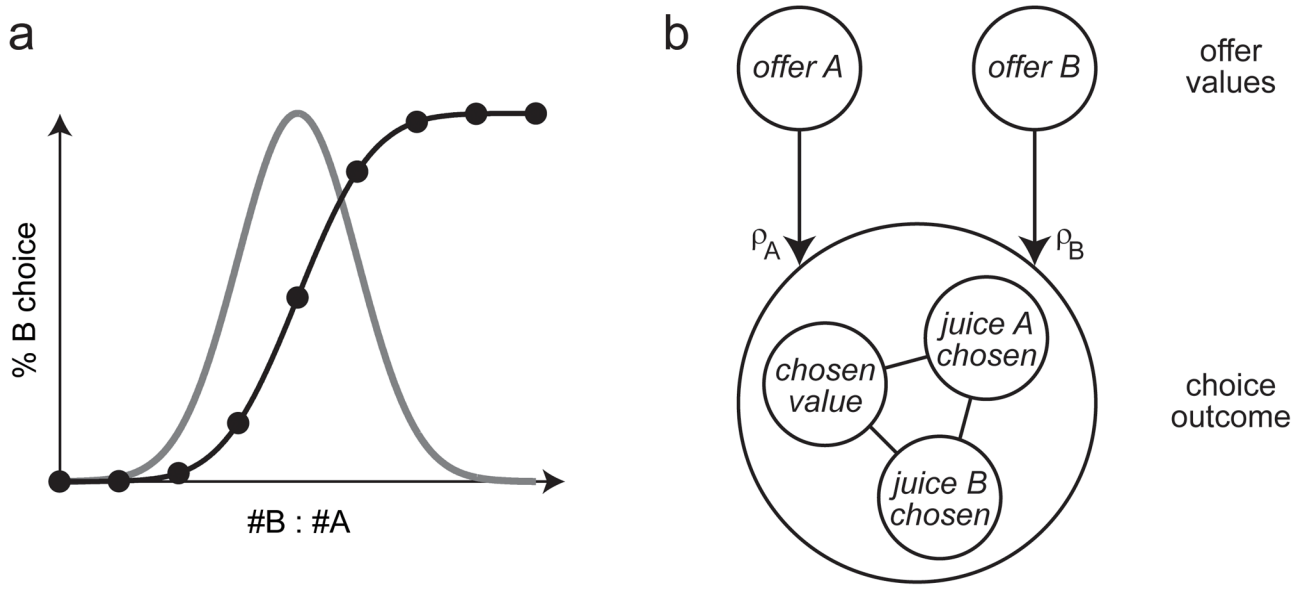

Figure 3.

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 non-zero 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 (ρ) 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 ρ_A and ρ_B. In this scenario and under reasonable assumptions, the relative value ρ equals the ratio ρ_A/ρ_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.

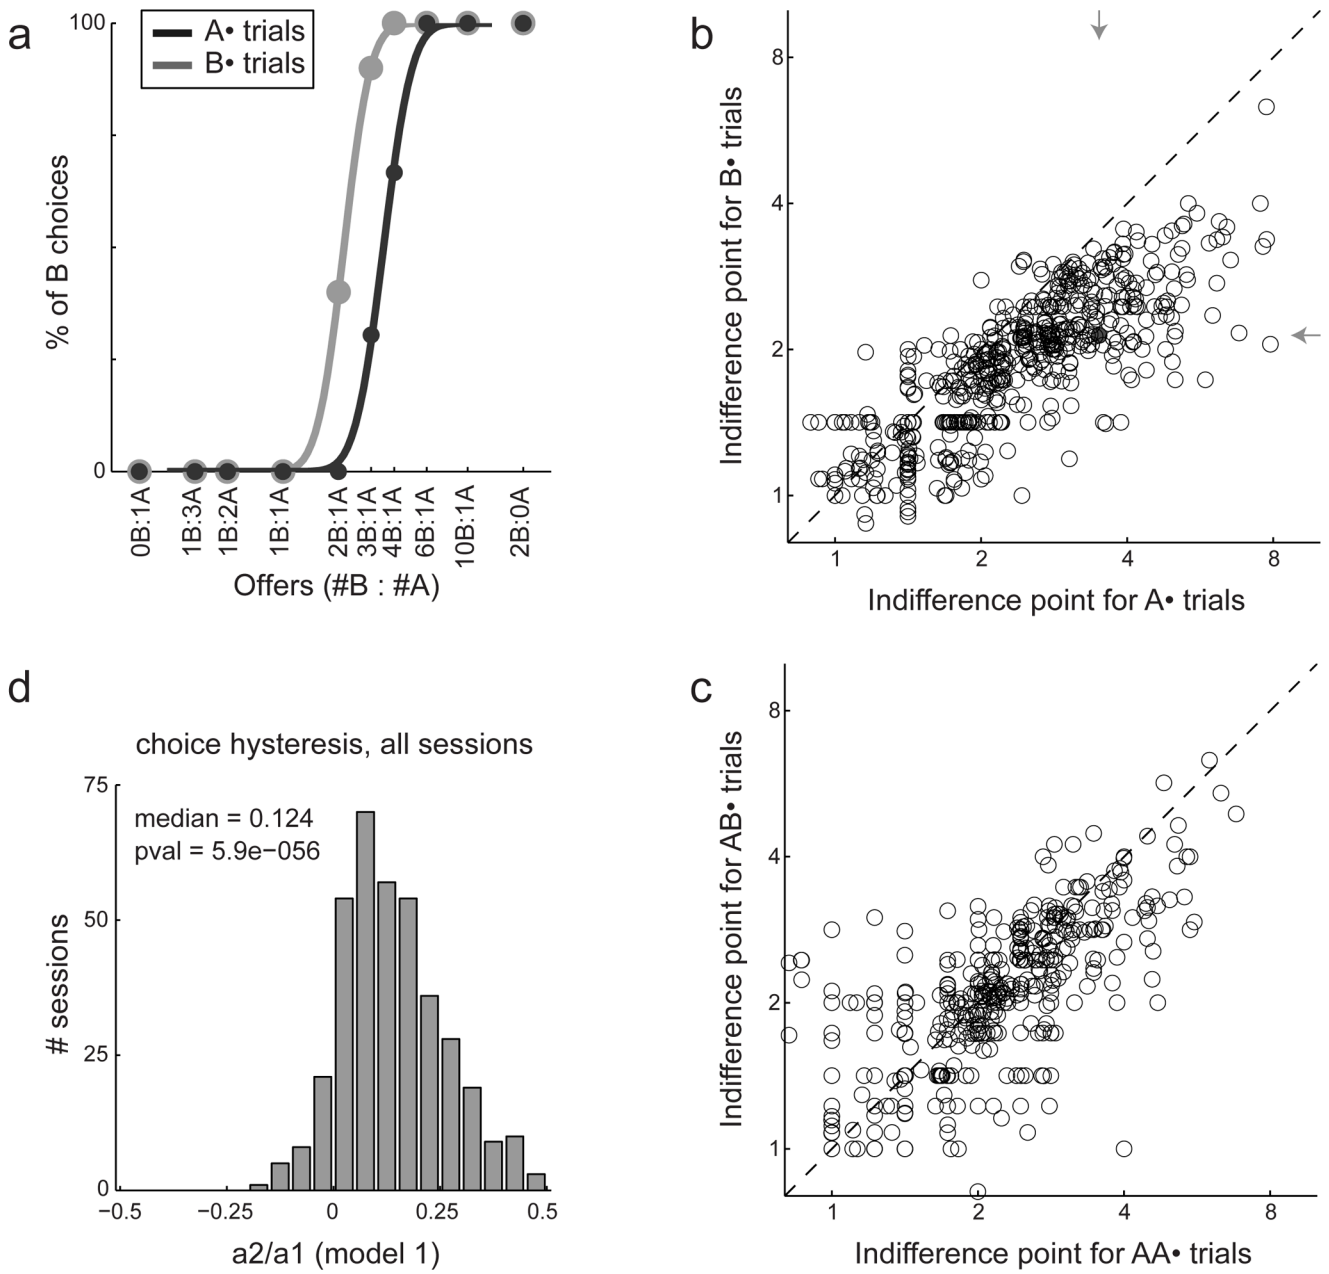

Figure 4.

Choice hysteresis. **a.** Behavioral data from one example session. Trials were separated 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•, 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 Eq.1, the y-axis represents the number of session (304 total).

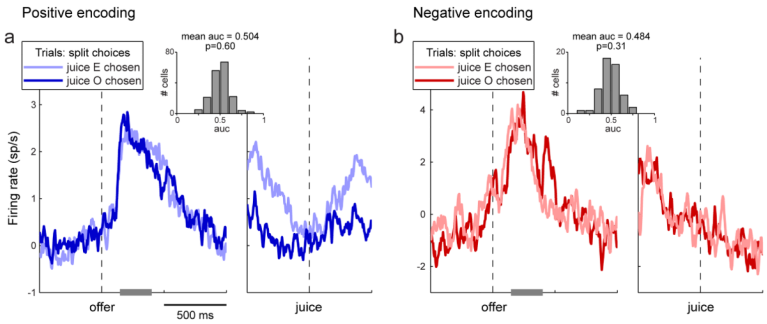

Figure 5.

Choice variability is not explained by fluctuations of *offer value* cells. **a.** 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 (≥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). In the insert, the histogram shows the distribution of AUC obtained across the population. The mean AUC was statistically indistinguishable from 0.5. **b.** 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 Fig. S2.)

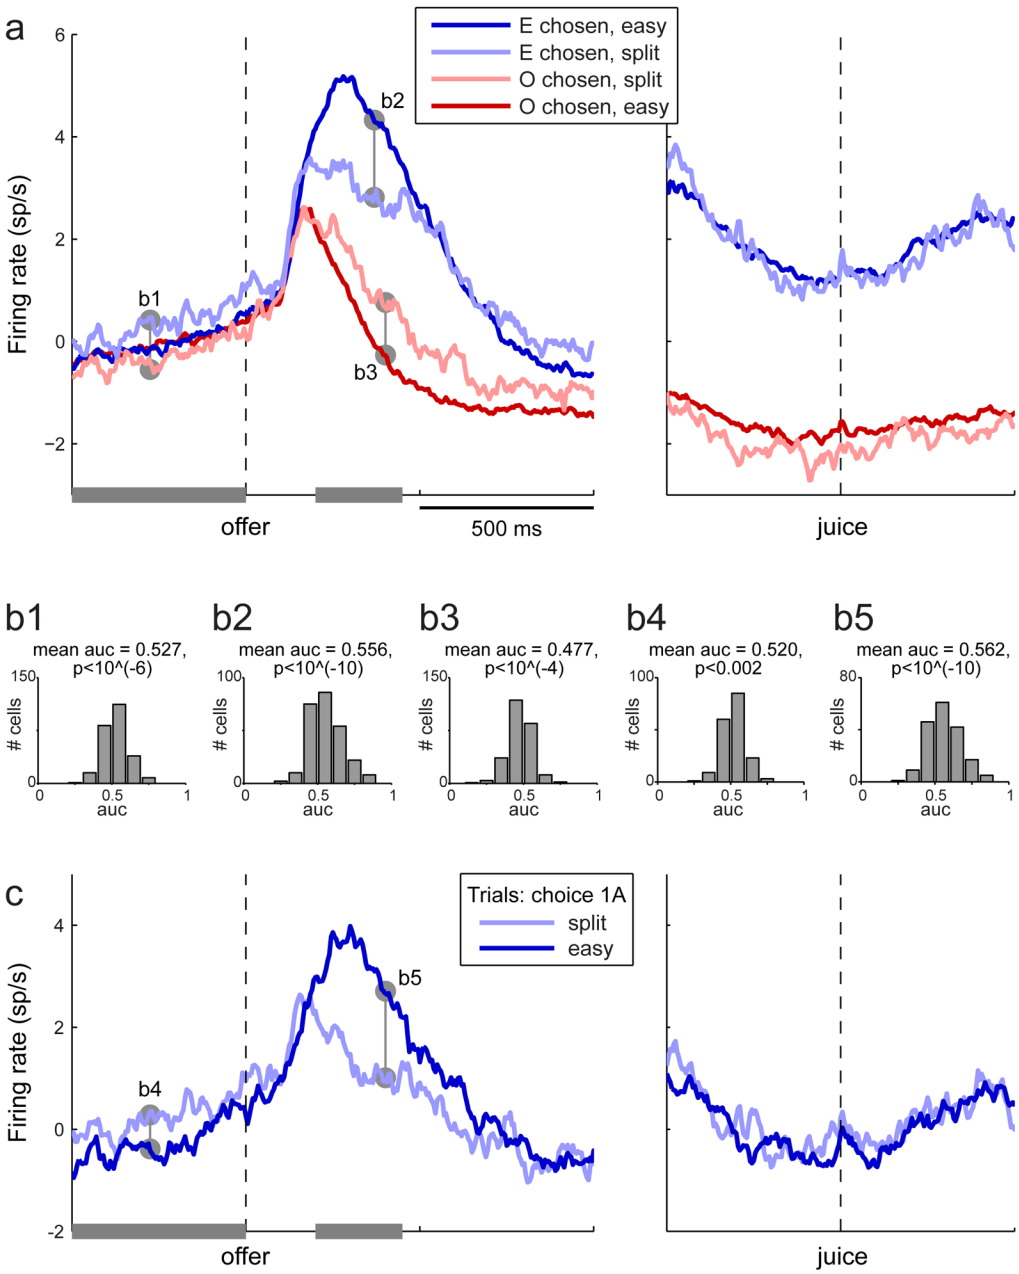

Figure 6.

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 (E) or the other 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 (≥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 panels (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 (≥2 trials per trace). All the effects described in (a) were also observed when the quantity of the chosen juice was fixed. (See also Figs. S3–S5.)

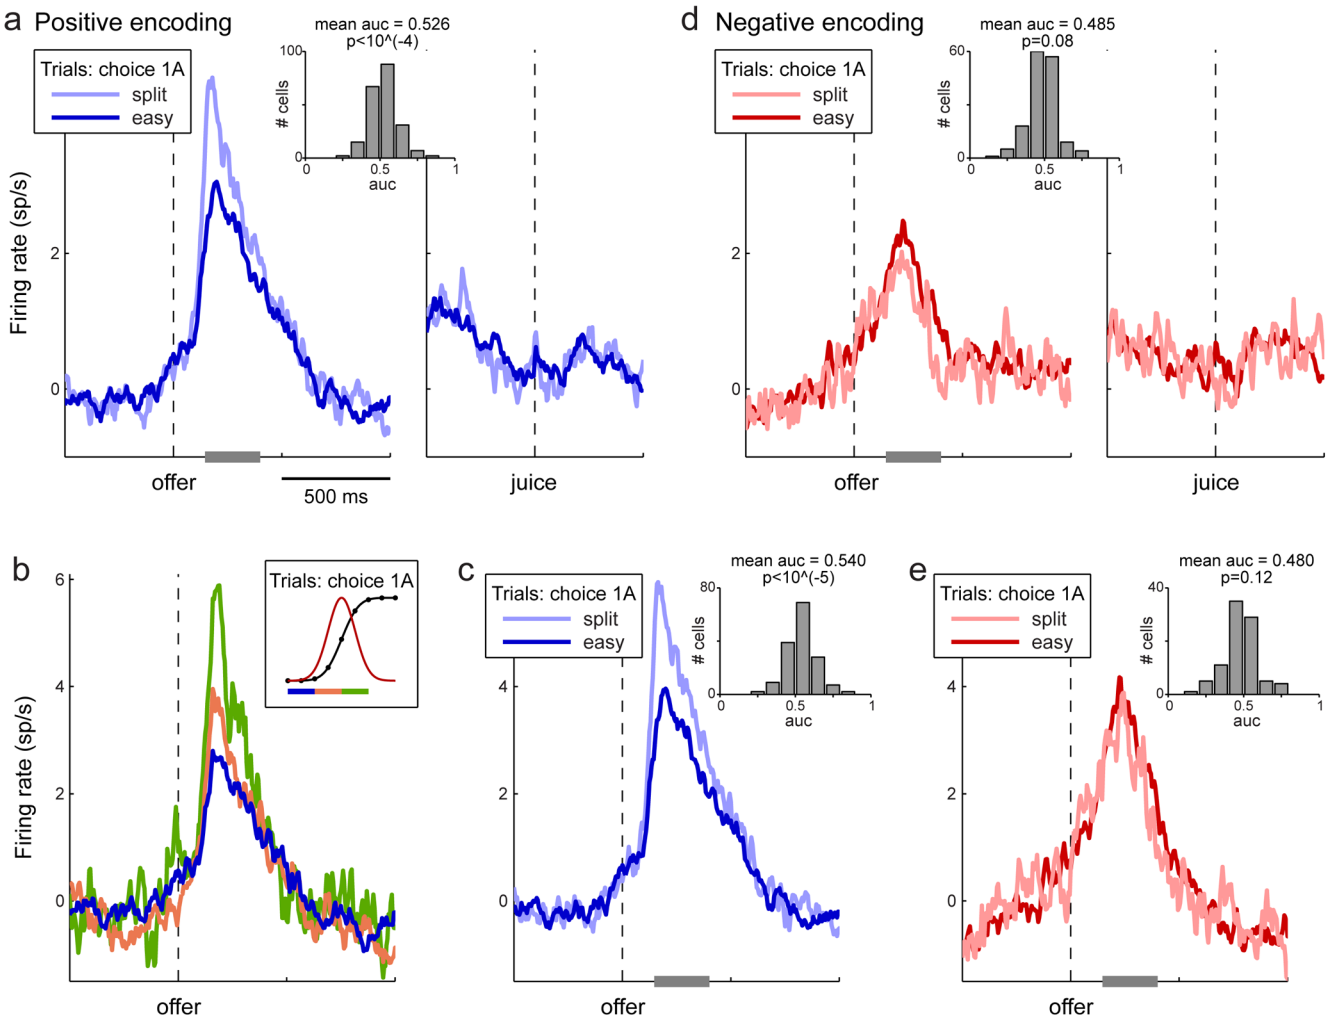

Figure 7.

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 (≥2 trials per trace). Cells from Exp.1 contributed to each average with at most one trace (choices of 1A). However, some cells from Exp.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 (b), 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 < mean(ρ) (yellow) and split decisions with q ≥ mean(ρ) (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 was not significant. **e.** Same analysis as in (e) 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 Figs. S6–S7.)

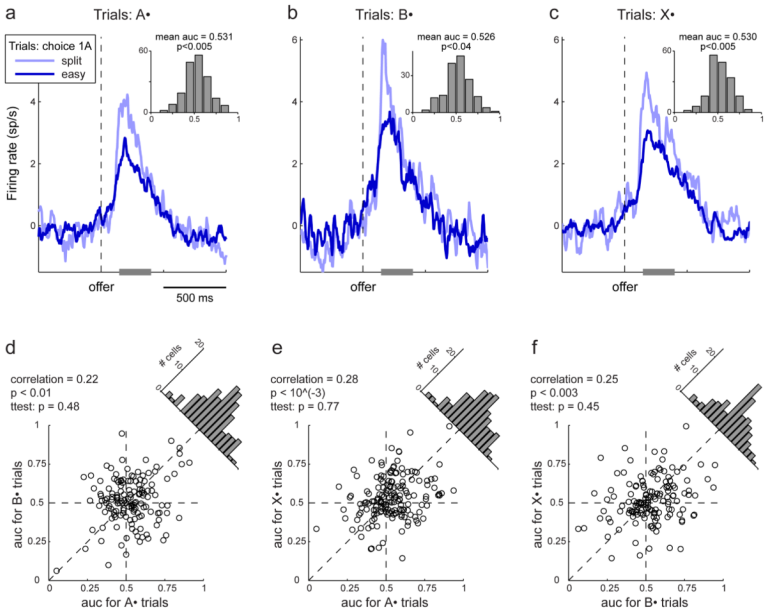

Figure 8.

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 Fig. 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 were significantly correlated (i.e., the AUC for any given cell was reproducible). However, the difference between the two measures 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.