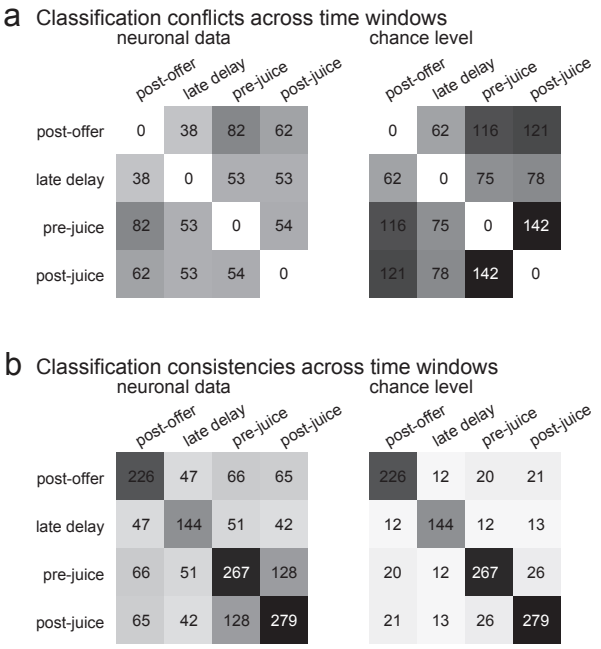

Figure S1. Analysis of classification conflict (related to Figure 1). **a.** In total, 443 cells from Exp.1 were classified in at least one time window and included in this analysis. The left panel refers to actual neuronal data. For each pair of time window, the number indicated the number of cells 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.

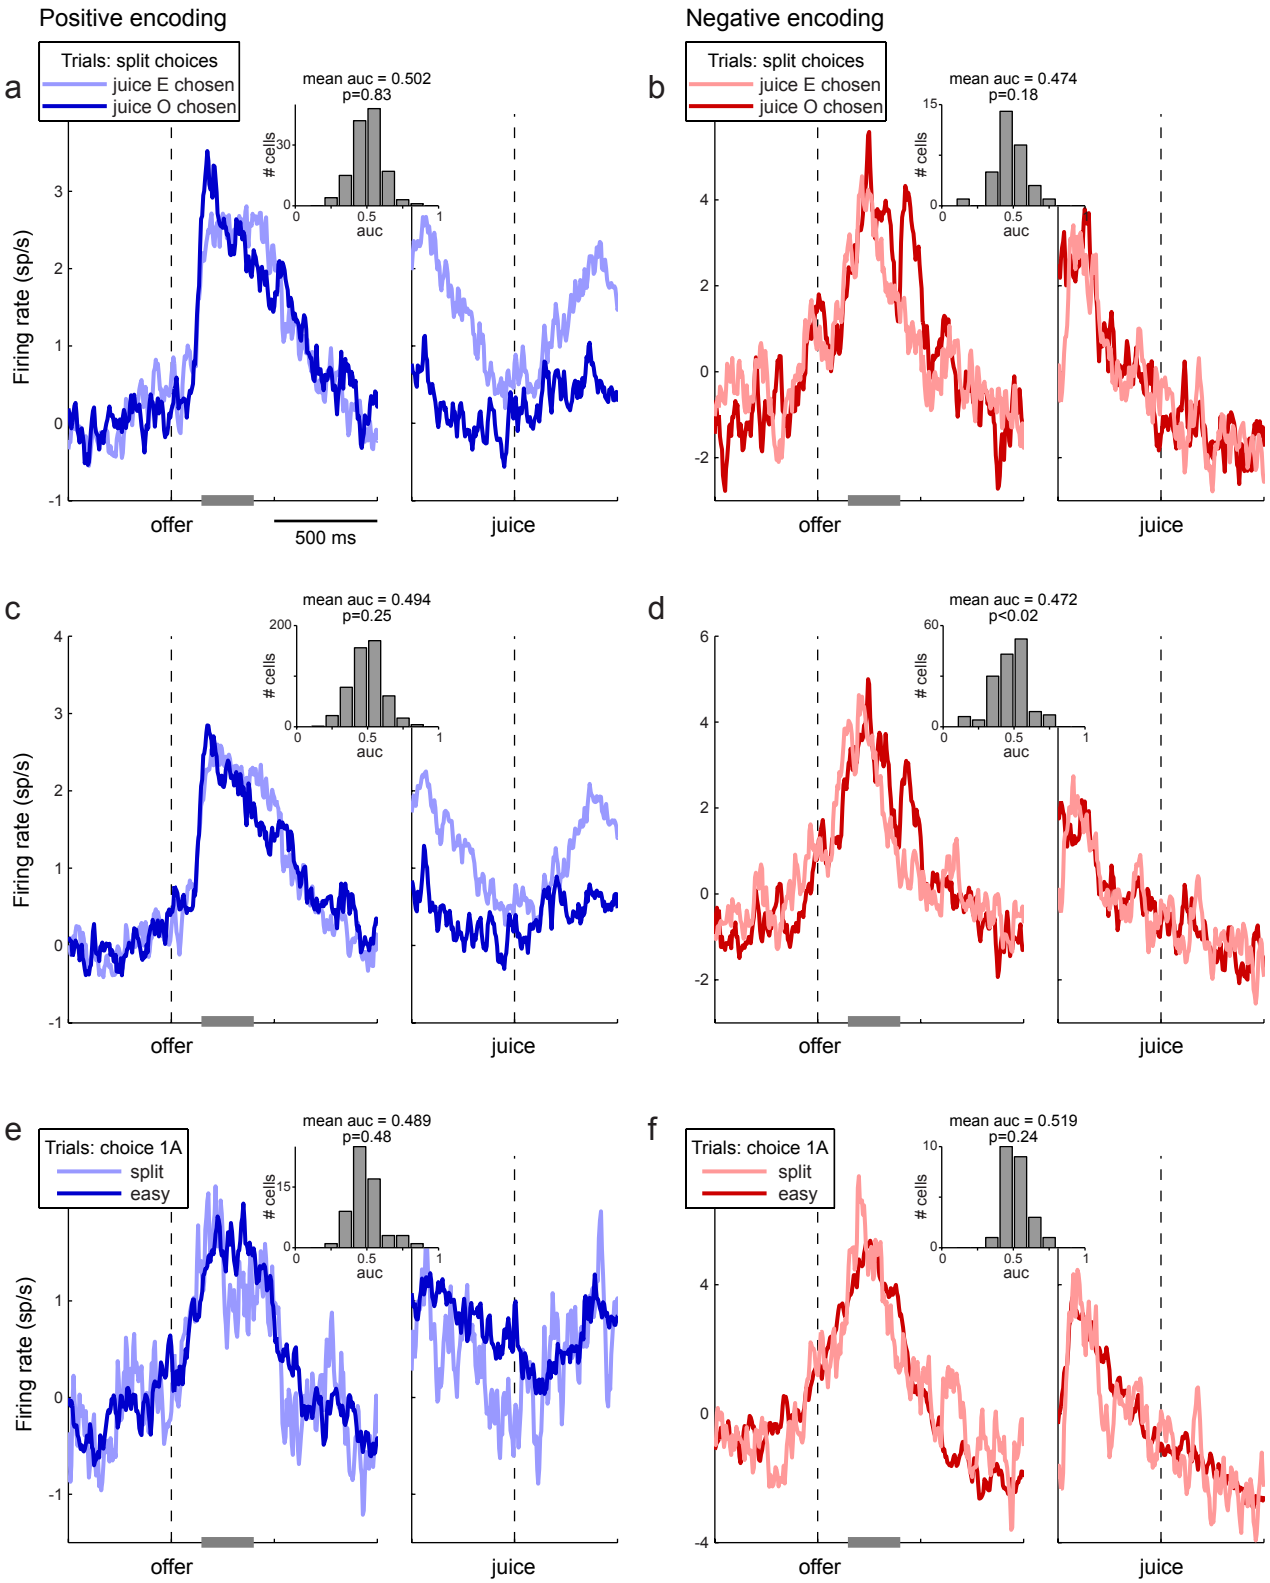

Figure S2. Control analyses for offer value cells (related to Figure 5). **ab.** Same analysis as in Fig.5ab including only neurons that were tuned in the 150-400 ms after the offer. Traces for positive and negative encoding are from 130 cells and 33 cells, respectively. **cd.** In this analysis, trials 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 (≥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.

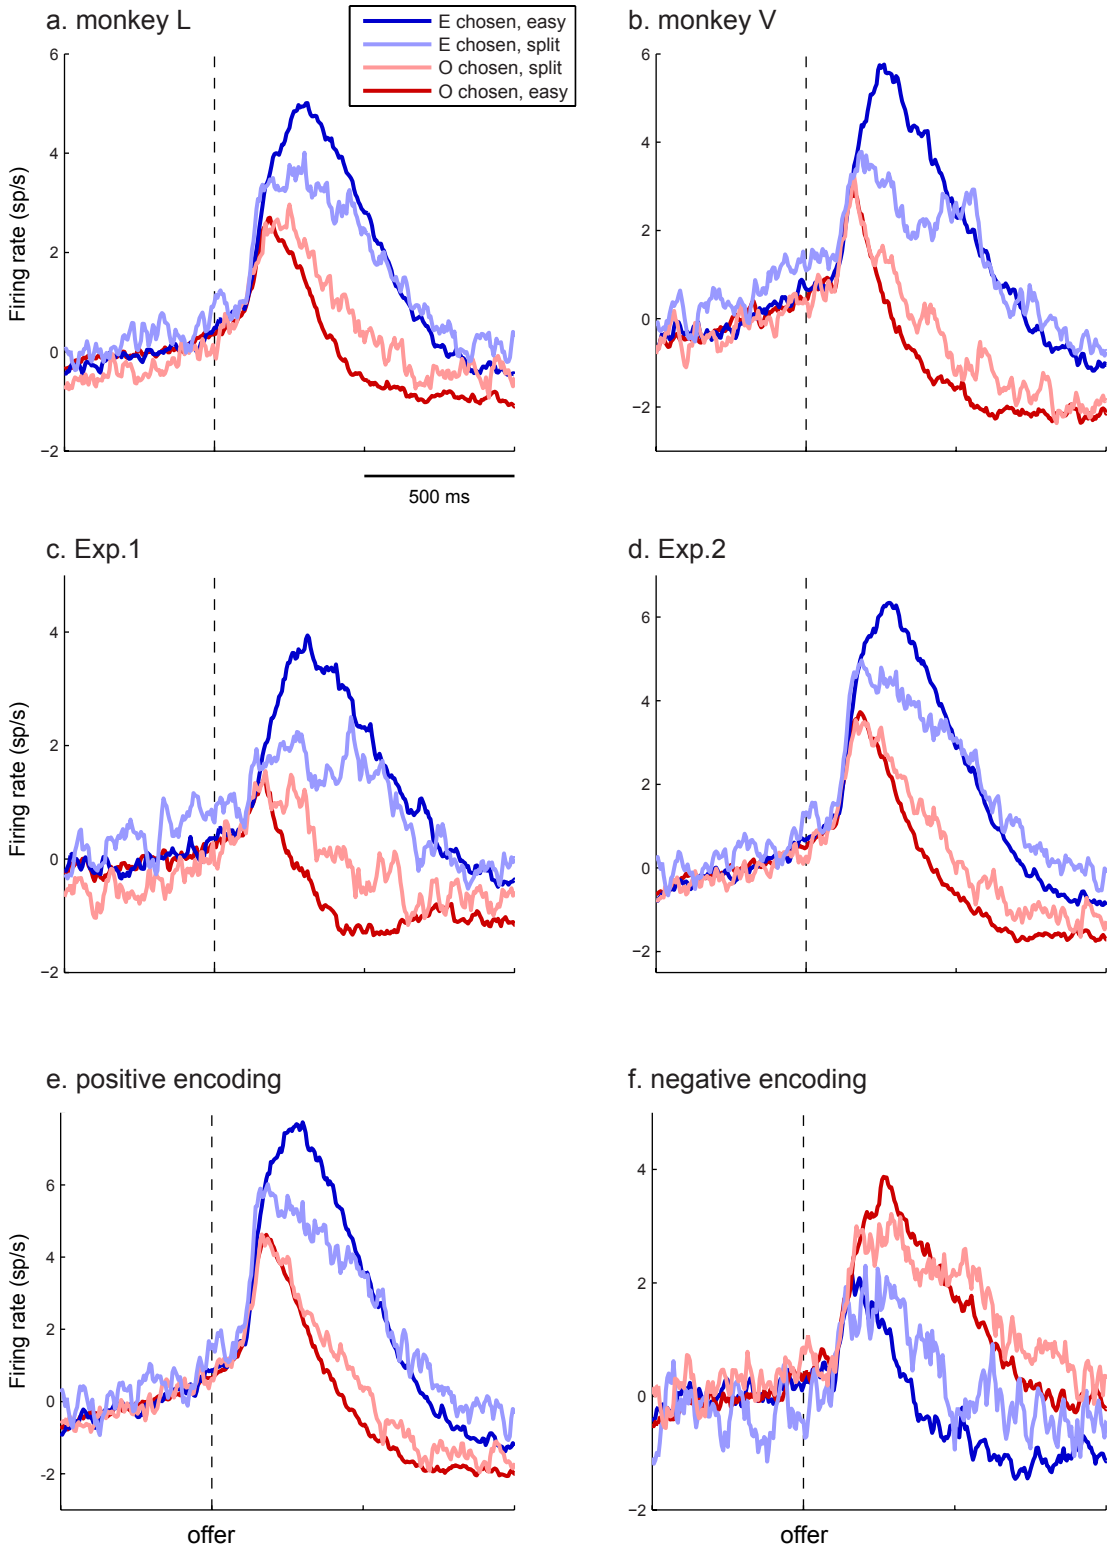

Figure S3. Control analyses for chosen juice cells (related to Figure 6). To verify the robustness of the results obtained for chosen juice cells, I repeated the analysis for different subsets of neurons: cells from monkey L (**a**, 169 cells), cells from monkey V (**b**, 96 cells), cells recorded in 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.

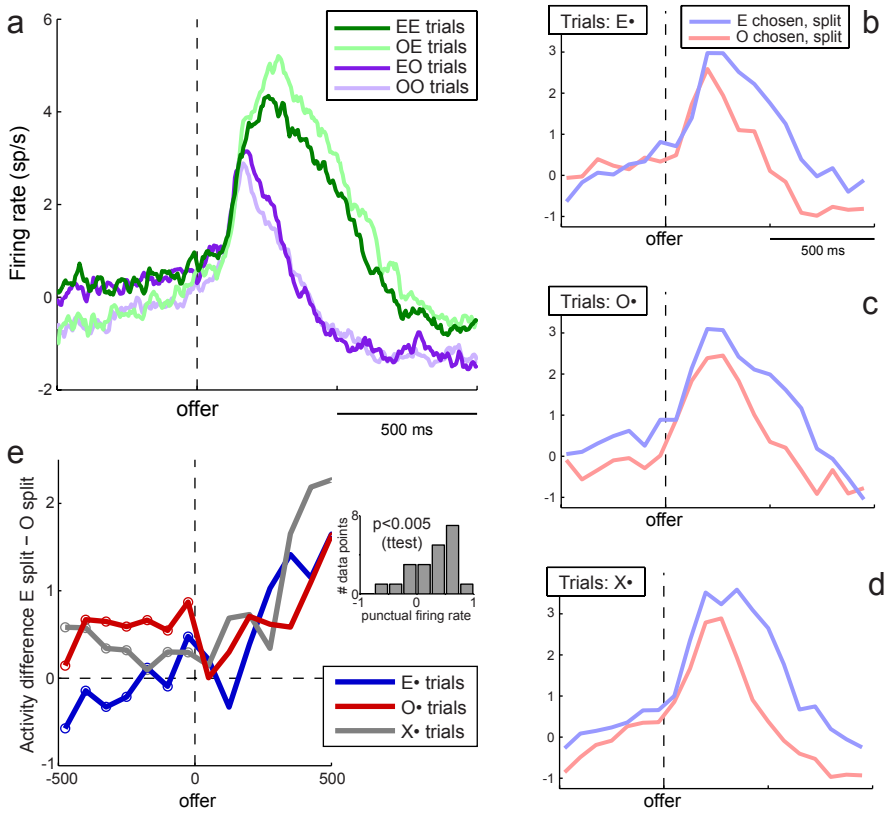

Figure S4. Chosen juice cells, activity in relation to the previous trial (related to Figure 6). **a.** Trials were divided depending on both the previous outcome and the current choice (see legend). The activity of chosen juice cells mainly depends on the current choice (green traces above purple traces after the offer). However, there is a tail activity from the previous trial (dark traces above purple 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.

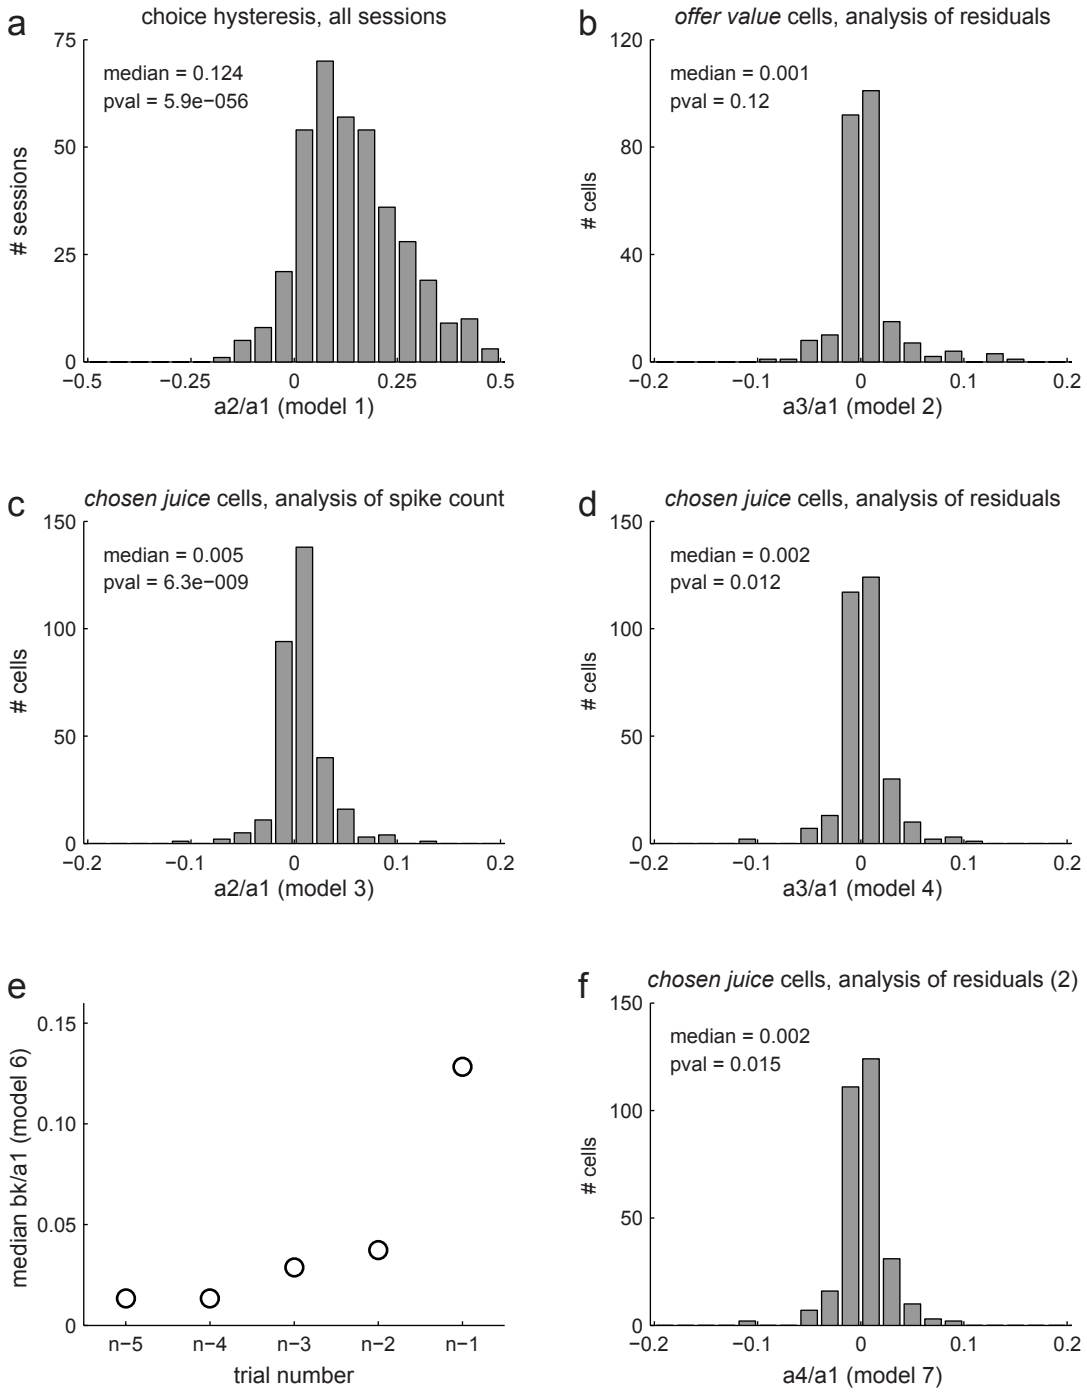

Figure S5. Results of logistic analyses (related to Figures 4, 5, 6). **a.** Choice hysteresis (same as Fig.4a). The x-axis represents the ratio a2/a1 defined in Eq.1, the y-axis represents the number of sessions (304 total). **b.** Analysis of offer value cells. The x-axis represents the ratio a3/a1 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).

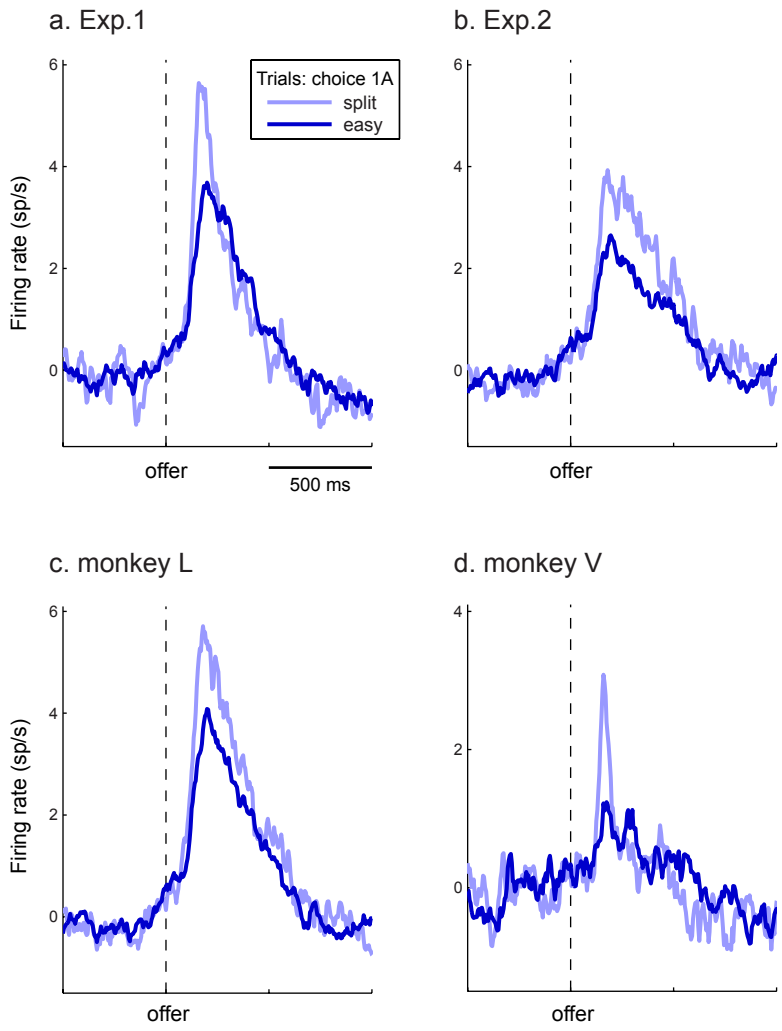

Figure S6. Control analyses for chosen value cells (related to Figure 7). Same analysis and all conventions as in Fig.7a. **a.** Data from Exp.1 (88 traces from 87 cells). **b.** Data from Exp.2 (124 traces from 64 cells). **c.** Data from monkey L (144 traces from 104 cells). **d.** Data from monkey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

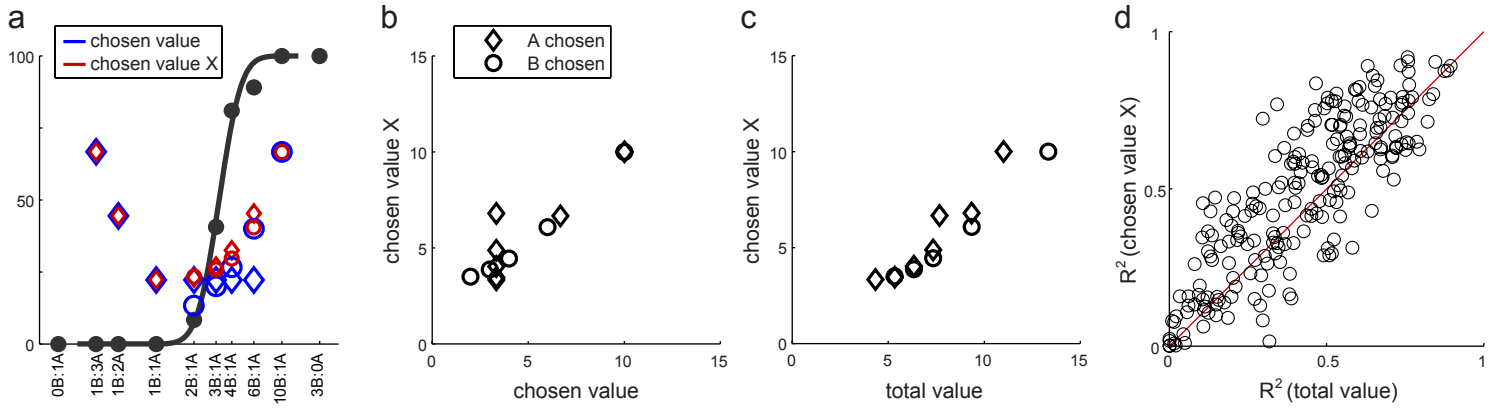

Figure S7. Overshooting of chosen value cells: contrasting variable chosen value X and total value (related to Figure 7). **a.** Comparing chosen value and chosen value X, example session. The choice pattern is the same as in Fig.1d. Blue and red symbols refer to variables chosen value and chosen value 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 R2 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 ΔR^2; $p<0.01$, Hartigan's dip test), except that the distribution was displaced towards higher values of ΔR^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:

$$k(t) = (1 - \exp(-t/\tau_g)) * \exp(-t/\tau_d) \quad \text{for } t \geq 0$$
$$k(t) = 0 \quad \text{for } t < 0$$

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 τ_g = 1 ms and τ_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 φ is the firing rate, φ_0 is the baseline activity, $\Delta\varphi$ is the activity range, V is the encoded value and ΔV 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 φ 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 ε is the fluctuation. This makes it clear that ε depends on both $V / \Delta V$ and $\Delta\varphi$. In essence, the analysis of Fig.5ab aims at studying ε 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 ε increases with $V / \Delta V$) and cells with large activity range (because ε increases with $\Delta\varphi$). Thus to increase the resolution of the analysis one would like to rescale the firing rate (and ε). 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 φ_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:

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

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

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 ~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:

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

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

For each *chosen juice* cell, $\varphi_{\text{residual 2}}$ is the residual firing rate remaining after the bilinear regression of the raw firing rate φ 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 (ρ) is defined as the quantity ratio that makes the animal indifferent between the two juices: V(A) = ρ 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►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 ρ (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 ρ. In the following analyses, I assumed that, once controlled for choice hysteresis, choice variability was entirely due to fluctuations of ρ. If this is true, then the probability distribution for ρ 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 ρ, the underlying normal distribution can be viewed as a probability distribution for the variable x = log ρ. Thus the probability distribution for ρ is N(x(ρ), μ, σ) dx/dρ = N(log ρ, μ, σ) 1/ρ. 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* = ρ ≥ q. Now consider many trials in which the animal chose 1A over qB. On average, the variable *chosen value* $X_{1A\blacktriangleright qB}$ = < ρ >$_{\rho \geq q}$ is equal to:

$$< \rho >_{\rho \geq q} = \frac{\int_{q}^{\infty} N(\log \rho, \mu, \sigma) \cdot 1/\rho \cdot \rho \, d\rho}{\int_{q}^{\infty} N(\log \rho, \mu, \sigma) \cdot 1/\rho \, d\rho} = \frac{\int_{\log(q)}^{\infty} N(x, \mu, \sigma) \cdot e^{x} \, dx}{\int_{\log(q)}^{\infty} N(x, \mu, \sigma) \, dx} \qquad (8)$$

More generally, when the animal chose a drops of juice A over b drops of juice B (aA►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 ξ such that B = ξ A and x = log ξ. In this case, the probability distribution for x is N(x, -μ, σ) and the probability distribution for ξ is N(log ξ, -μ, σ) 1/ξ. Thus the variable *chosen value* $X_{1B\blacktriangleright 1/q\ A}$ = < ξ >$_{\xi \geq 1/q}$ is equal to:

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

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 <ρ>. 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 ρ, is significantly higher than that of *total value*.

<u>The overshooting of *chosen value* cells is independent of choice hysteresis</u>

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 the 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.