



Published in final edited form as:

Cell Rep. 2024 October 22; 43(10): 114772. doi:10.1016/j.celrep.2024.114772.

The representation of decision variables in orbitofrontal cortex is longitudinally stable

Manning Zhang (张曼宁)^{1,2,4}, Alessandro Livi^{1,4}, Mary Carter¹, Heide Schoknecht¹, Andreas Burkhalter¹, Timothy E. Holy^{1,2}, Camillo Padoa-Schioppa^{1,2,3,5,*}

¹Department of Neuroscience, Washington University in St. Louis, St. Louis, MO 63110, USA

²Department of Biomedical Engineering, Washington University in St. Louis, St. Louis, MO 63110, USA

³Department of Economics, Washington University in St. Louis, St. Louis, MO 63110, USA

⁴These authors contributed equally

⁵Lead contact

SUMMARY

The computation and comparison of subjective values underlying economic choices rely on the orbitofrontal cortex (OFC). In this area, distinct groups of neurons encode the value of individual options, the binary choice outcome, and the chosen value. These variables capture both the choice input and the choice output, suggesting that the cell groups found in the OFC constitute the building blocks of a decision circuit. Here, we show that this neural circuit is longitudinally stable. Using two-photon calcium imaging, we record from the OFC of mice engaged in a juice-choice task. Imaging of individual cells continues for up to 40 weeks. For each cell and each session pair, we compare activity profiles using cosine similarity, and we assess whether the neuron encodes the same variable in both sessions. We find a high degree of stability and a modest representational drift. Quantitative estimates indicate that this drift would not randomize the circuit within the animal's lifetime.

Graphical Abstract

This is an open access article under the CC BY-NC license (<http://creativecommons.org/licenses/by-nc/4.0/>).

*Correspondence: camillo@wustl.edu.

AUTHOR CONTRIBUTIONS

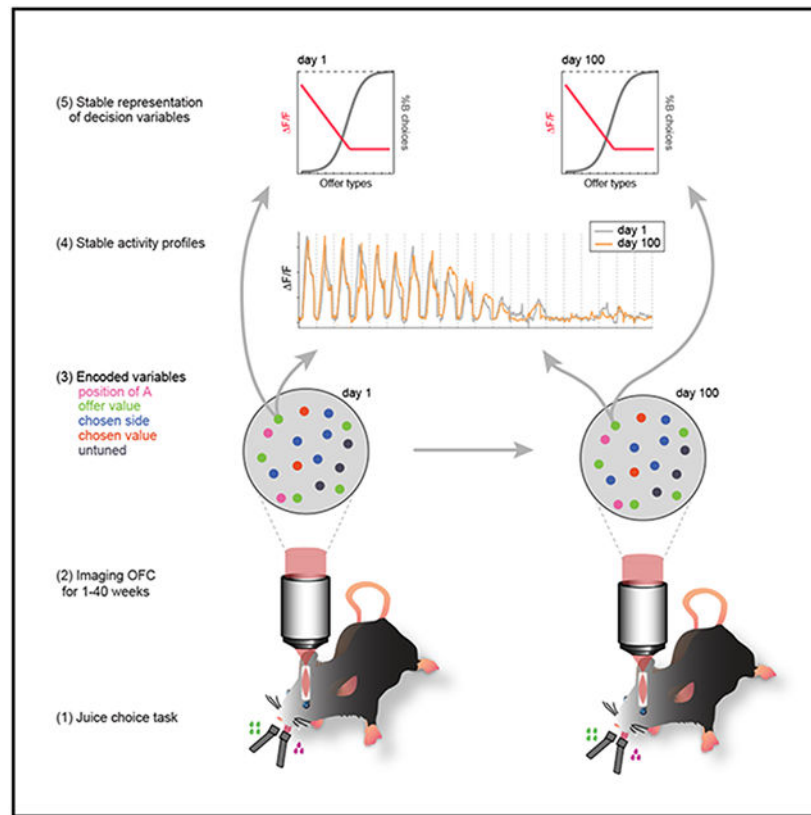
M.Z., A.L., T.E.H. and C.P.-S. designed the study. M.Z. and A.L. performed the experiments with assistance from M.C. and H.S. M.Z. and A.L. processed and analyzed the data. A.B. helped with anatomy and histology. C.P.-S. and T.E.H. supervised the project. M.Z. wrote the initial draft. All authors edited the manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.celrep.2024.114772>.



In brief

Economic choices depend on the orbitofrontal cortex (OFC). In mice choosing between juice flavors, Zhang et al. record from individual neurons for up to 40 weeks. Activity profiles and the representation of decision variables are longitudinally stable. Hence, different cell groups in the OFC form a stable decision circuit.

INTRODUCTION

Economic choices entail computing the values of different options; a decision is then made by comparing values. Work in primates and rodents has shown that these mental processes rely on the orbitofrontal cortex (OFC). Lesion or inactivation of this area severely disrupts choices.¹⁻⁴ Furthermore, electrical stimulation in monkeys showed that offer values represented in OFC are causal to choices⁵ and that neuronal activity in OFC directly contributes to value comparison.⁶ Neurophysiology studies examined neuronal activity in monkeys and rodents choosing between different juice types. Neurons in the OFC were found to encode different decision variables, including the value of individual offers, the binary choice outcome, and the chosen value.^{2,3,7,8} These variables capture both the choice input (offer value) and the choice output (choice outcome, chosen value). Moreover, population dynamics in OFC are consistent with the formation of a decision.^{9,10} Taken together, these lines of evidence suggest that different cell groups identified in OFC constitute the building blocks of a neural circuit in which economic

decisions are formed. The organization and the mechanisms governing this circuit are poorly understood. A fundamental question is whether this circuit is stable over extended periods of time; alternatively, the representation of decision variables might undergo substantial reorganization.¹¹

Longitudinal studies of other neural circuits have reported a spectrum of results. In the parietal cortex of mice engaged in a virtual navigation task, variables encoded by individual cells underwent major reorganization over weeks.¹² Interestingly, this drift was constrained by, but not strictly confined to, the null space.¹³ Representational drift was also found in visual¹⁴ and olfactory¹⁵ areas. In particular, longitudinal recordings in the primary visual cortex found relatively stable responses to artificial stimuli (gratings) but substantial drift in the representation of natural scenes.¹⁶⁻¹⁸ Place fields in the mouse hippocampus were found to be quite unstable.¹⁹ However, the neuronal propensity to spike was preserved over time,²⁰ and changes in spatial tuning were associated primarily with increased experience in the arena, as opposed to the mere passage of time.²¹ In the premotor area higher vocal center (HVC) of songbirds, the number and timing of spikes of individual neurons in relation to courtship songs remained remarkably stable over weeks.²² Similarly, in mice performing an olfaction-guided spatial memory task, responses in the medial prefrontal cortex remained stable for 3–4 weeks.²³ With respect to orbital regions, neurons in the ventromedial OFC of mice undergoing Pavlovian conditioning responded stably for 6 days.²⁴ Most relevant here, in monkeys performing economic choices, neurons in central OFC (homologous to the region examined here) presented a stable representation of decision variables across sessions within a day.²⁵ However, it remains unclear whether this representation is also stable over longer periods of time. Importantly, behavioral choice patterns—i.e., the relative values (ρ) of different goods and the choice accuracy (η)—can vary significantly from day to day.^{2,7} One hypothesis is that this behavioral variability reflects instability in the decision circuit. Alternatively, choice variability could result from a stable decision circuit and simply reflect variability in the input to that circuit or fluctuations in the motivational state of the animal.

To shed light on these issues, we used two-photon (2P) calcium (Ca^{2+}) imaging in combination with gradient-index (GRIN) lenses to record from the OFC of mice performing an economic choice task.² Recordings in each animal continued for many weeks, and we could reliably image the same cells in many sessions. Neurons were analyzed separately in each session, and the results were compared across sessions. Activity profiles and the decision variables encoded by individual neurons remained remarkably stable over many weeks. We observed some neuronal changes associated with the degree of experience in the task. In addition, we found a limited representational drift associated with the passage of time. This phenomenon was statistically significant but quantitatively modest. Indeed, the time necessary for a complete reorganization of the neuronal circuit was estimated at >3.5 years—longer than the animal's lifetime. Our results suggest that day-to-day variability in choices is not due to instability in the decision circuit. The presence of a stable circuit within OFC provides the opportunity to examine the mechanisms underlying economic choices.

RESULTS

Economic choices under the microscope

Mice were accustomed to head fixation under the 2P objective and then trained to make binary choices. The choice task closely resembled that used for neurophysiology studies in mice and monkeys.^{2,7} In each session, the mouse chose between two types of juice labeled A and B (A preferred) and offered in variable quantities. Offers were represented by olfactory stimuli. The odor identity represented the juice type, and the odor concentration represented the juice quantity. For each juice type, we used 4 or 5 quantity/concentration levels. The two odors were presented on the left and on the right, and the animal indicated its choice by licking one of two spouts placed near its mouth (Figure 1A). Offered quantities and their spatial configuration varied pseudo-randomly from trial to trial. Figure 1B illustrates the trial structure (see STAR Methods).

Mice choices reliably presented a trade-off between juice type and juice quantity (Figure 1C). For each session, we analyzed choices using logistic regressions (STAR Methods). This analysis provided measures for the ρ , the η , and the side bias (ϵ). Indicating with q_A and q_B the juice quantities offered in each trial, ρ was the quantity ratio q_B / q_A that made the animal indifferent between the two offers, η was proportional to the sigmoid steepness, and ϵ was a bias favoring one side over the other. Our whole dataset included 343 sessions from 13 animals. Figures 1D-1F illustrate the distribution for ρ , η , and ϵ across sessions. Notably, the ϵ was generally modest ($|\epsilon| < \rho$), with a distribution roughly centered on zero.

We examined whether and how choices changed as a function of the animals' experience in the task. For each mouse, we counted days starting with the first day it chose between different juice types. For each behavioral session, we performed a logistic analysis and derived measures for ρ , η , and ϵ . For each mouse and each parameter, we studied how values varied longitudinally by fitting them with an exponential function (STAR Methods). Figure S1 illustrates the results obtained for each mouse in the dataset. For some animals, one or more parameters showed a clear longitudinal trend. We found significant effects for 6 of the 12 mice in which we could conduct this analysis, and 4 other mice showed similar trends that did not reach significance level (Table S1).

Importantly, ρ and ϵ ultimately capture aspects of the animal's subjective preferences, and longitudinal trends of these parameters do not necessarily indicate a learning process or improved performance. In contrast, in a normative sense, η should be as large as possible (i.e., sigmoids should be as steep as possible).²⁶ Thus, for a population analysis we focused on η . Pooling sessions from all animals, we fitted the whole dataset with a single exponential function. The longitudinal trend was statistically significant (Figure 2). In other words, performance improved over days of experience in the task. The time constant derived from the exponential fit was $\tau = 31$ days.

Longitudinal imaging of OFC during choice behavior

We conducted 2P imaging while mice performed the choice task. Mice expressing GCaMP6f under a synapsin promoter were implanted with a GRIN lens targeting OFC (Figure 3A;

STAR Methods). In each animal, we recorded from multiple fields of view (FOVs) (in different sessions). Images from each session were segmented using the CaImAn package,²⁷ and the results were checked manually. In any FOV, we recorded from 12–225 neurons (median = 77 neurons) (Figures 3B and 3C). Across sessions, we recorded from 2–12 FOVs in each mouse. Thus, we could record from 140–850 individual neurons in each mouse.

Data for the present study were collected in 13 mice and 77 FOVs. For longitudinal analysis, we recorded from each FOV 2–16 times over a maximum span of 12–41 weeks. For each pair of sessions in which we visited the same FOV, we manually aligned the two sets of images using blood vessels and neurons with distinctive features as landmarks. Figure 3D illustrates one example FOV visited 5 times over the course of 83 days. Images from different days were remarkably similar. To match individual neurons across pairs of sessions, we used a Bayesian procedure (see STAR Methods). Given a pair of sessions, we typically found that 60%–80% of cells identified in any session could be matched with cells identified in the other session. Figure 3E illustrates the results of this analysis for two sessions 11 days apart. Notably, the majority of cells in this FOV could be matched across sessions. Considering all session pairs, the fraction of cells that could be matched across sessions decreased as the number of days between the two sessions increased but remained >50% even when recordings were 100 days apart (Figure 3F).

Activity profiles are longitudinally stable

After cell segmentation, we normalized the fluorescent trace for each cell and obtained the $\Delta F / F$ signal. All analyses of neural activity were based on this signal. We first examined neuronal activity profiles. A trial type was defined by two offered quantities (q_A and q_B), their spatial configuration, and a choice. For each session, we examined $\Delta F / F$ traces recorded in each trial. We joined two time windows aligned with the offer onset (from –600 to +1,600 ms) and the first lick (from –600 to +600 ms). We then averaged traces across trials for each trial type. Finally, we concatenated the resulting mean traces across trial types and obtained the activity profile (see STAR Methods). We repeated this operation for each cell and each pair of sessions, including only trial types present in both sessions. Figures 4A–4C illustrate the activity profiles recorded for 3 example cells in two sessions ~7 days apart. Notably, activity profiles differ substantially from cell to cell but appear quite stable for each cell. We quantified the likeness between pairs of activity profiles using cosine similarity (CS; see STAR Methods). Intuitively, activity profiles were vectors where each component corresponded to a time interval; CS was the cosine of the angle between two such vectors. In principle, CS could vary between –1 and 1, and CS = 1 if the two traces were identical. For the three example cells illustrated in Figures 4A–4C, we measured CS = 0.96, CS = 0.94, and CS = 0.90, respectively.

To assess how CS varied across the population and as a function of time, we proceeded as follows. Counting days starting from the first day of recording, for any two sessions, T1 and T2 indicated the recording days (T1 – T2), and $\Delta T = T2 - T1$ was the time distance. We examined how CS varied as a function of both T1 (i.e., experience in the task) and ΔT (i.e., time passage). First, we considered the plane defined by T1 and ΔT , and we partitioned it in two-dimensional time bins of 2×2 weeks. Both T1 and ΔT ranged between 0 and 20

weeks (100 time bins total). Second, for each time bin, we considered neurons that had been recorded in two sessions with corresponding T1 and ΔT . Because of the large number of FOVs included in this study and the multitude of recording sessions obtained for each FOV, our dataset included a considerable number of cells for most bins (Figure 4E). Third, taking advantage of this rich dataset, we examined the distribution of CS obtained in each time bin.

Figure 4D illustrates the results for one representative time bin ($n = 747$ cells). Notably, the CS distribution was broad and almost entirely >0.5 (median CS = 0.677). Since our measures of $\Delta F / F$ were nearly always >0 , profile vectors were essentially confined to a single orthant of a high-dimensional space, which implies that CS was necessarily large. To gauge whether measures of CS were higher than one would expect by chance, we examined two control treatments. First, for any FOV recorded in two sessions, we arbitrarily mismatched the cell pairings by randomly coupling neurons from session 1 to neurons from session 2 (treatment: mismatched cells). We then proceeded with the computation of CSs and generated the resulting distribution. As expected, CSs were generally >0.5 . However, the median of this distribution (median CS = 0.617) was significantly lower than that measured when neurons were matched ($p = 2.26 \times 10^{-38}$; Kruskal-Wallis test). Second, for any cell recorded in two sessions, we arbitrarily shuffled the activity profile by randomly permuting the components of the profile vector (treatment: shuffled times). Again, CSs obtained for this treatment were generally >0.5 . However, the median of this distribution (median CS = 0.616) was significantly lower than that measured for actual activity profiles ($p = 1.37 \times 10^{-29}$; Kruskal-Wallis test).

We repeated these analyses for each time bin. Thus for every pair of T1 and ΔT , we obtained three measures of median(CS) corresponding to the three treatments (same cells, mismatched cells, and shuffled times). Heatmaps in Figures 4F-4H illustrate our results. Several points are noteworthy.

First, the median(CS) measured for the same cells was higher than that computed for either of the two control treatments. This was true for every value of T1 and ΔT , and the difference was statistically significant for most values of T1 and ΔT (Figures 4I and 4J; Kruskal-Wallis test). In other words, the activity profiles recorded for individual cells in different sessions were significantly more similar than expected by chance.

Second, a direct comparison of the two control treatments indicated that, for most time bins, the median(CS) values were statistically indistinguishable (Figure 4K). Importantly, the time shuffling procedure is non-biological and only preserves the fact that $\Delta F / F$ is generally >0 . Thus, one might expect that the median(CS) obtained for the time shuffled control would be substantially lower than that obtained for mismatched neurons. Conversely, the result of Figure 4K indicates a high degree of variability in the activity profiles of different cells.

Third, visual inspection of Figure 4F suggested that median(CS) increased as a function of T1 and decreased as a function of ΔT . To confirm this observation, we defined larger time bins of 6×6 weeks (Figure 5A). We refer to these as Bins. We then compared the CS distributions measured for different Bins. (1) To ascertain whether activity profiles varied as a function of T1, we compared CSs measured in Bin 1 (T1 = weeks 0–6; ΔT = weeks 0–6)

and Bin 2 (T1 = weeks 15–20; ΔT = weeks 0 – 6). Confirming our impression, median(CS) was significantly higher in Bin 2 than in Bin 1. In other words, the activity profiles of individual neurons became more reproducible as the animals became more experienced in the task. (2) To ascertain whether activity profiles varied as a function of ΔT , we compared CSs measured in Bin 1 (T1 = weeks 0–6; ΔT = weeks 0 – 6) and Bin 3 (T1 = weeks 0–6; ΔT = weeks 15 – 20). Indeed, median(CS) was significantly lower in Bin 3 than in Bin 1. In other words, the activity profiles of individual neurons became more dissimilar as a function of the time distance between sessions (representational drift).

Representational drift and longitudinal stability

To precisely quantify longitudinal drifts, we examined how CS varied as a function of ΔT (time passage) and T1 (experience in the task) with a 1-day resolution. In drift-diffusion processes, distance from the initial position increases as the square root of time. Thus, to model longitudinal changes in CS in terms of a drift process, we performed the regression $CS(\Delta T) = a_0 + a_1 \Delta T^{1/2}$ (Figure 5B). We also computed $b = \text{median}(CS)$ for the “mismatched cells” treatment. In this context, b can be regarded as the chance level for CS—i.e., the expected value for CS if the population underwent full reorganization. From the fitted parameters, we estimated the time necessary for full reorganization of the neural circuit given the drift rate, namely ΔT_{FR} , as follows:

$$\Delta T_{FR} = ((b - a_0) / a_1)^2 = 1,278 \text{ days}.$$

In other words, the time necessary for a full reorganization was >3.5 years—longer than the animal’s lifetime. In conclusion, the representational drift observed in this neuronal population was statistically significant but quantitatively modest.

Control analyses based on an alternate definition of the activity profile (Figure S2) or restricted to task-related cells yielded similar results (Figure S3).

A similar analysis confirmed that CS increased as a function of T1 (Figure 5C).

Task performance and reproducibility of the activity profiles in individual animals

We next investigated whether the increase of CS as a function of the animal’s experience in the task was related with the trend observed in behavioral performance (Figure 2) and whether these phenomena and their possible correlation could be observed in individual animals.

First, pooling sessions from all animals, we found that CS and η were indeed significantly correlated ($p = 0.04$; Figure 5D).

Second, to assess whether the relation between CS and η was driven by the animal’s experience in the task, we considered for each mouse and each FOV the last recording session (behavioral plateau). Any other session preceded the last session, and we defined T3 as the time difference between the session of interest and the last session $T3 = T - T_{\text{last session}}$ ($T3 < 0$).

Third, we regressed CS on η separately for each mouse (Figure S4A). Since the two measures were noisy, the linear term did not differ significantly from zero in most animals. That aside, we aimed to assess whether the resulting linear function captured a real relationship between CS on η and whether these measures and their relationship were related to T3. To do so, one would ideally define a single variable X incorporating CS, η , and their relationship and examine how X varies as a function of T3. Thus, for each animal and for each session, we computed the estimated CS' based on η and the fitted linear function (Figure S4A), and we defined X as the geometric mean between CS and CS'. We predicted that X and T3 would be positively correlated, and this prediction held true in the majority of individual mice (Figures 5E and S4B). The correlation between X and T3 reached statistical significance ($p < 0.01$) in 5/12 animals.

These results suggest that neuronal activity profiles became more reproducible as mice gained experience in the task and that this trend induced higher behavioral accuracy.

Neuronal encoding of decision variables

We next examined the variables encoded by individual neurons in OFC. This analysis was similar to that previously conducted for spiking activity.^{2,7} We defined five time windows aligned with different behavioral events (Figure 1B; STAR Methods). Again, a trial type was defined by two offered quantities, their spatial configuration, and a choice. For each cell, time window, and trial type, we averaged $\Delta F / F$ over trials. A neuronal response was defined as the activity of one neuron in one time window as a function of the trial type. Each FOV was analyzed based on a single session (i.e., each cell contributed only once to this analysis). Thus our dataset included 6,620 neurons from 13 animals and 73 FOVs (see STAR Methods).

The analysis proceeded in steps. First, we examined the activity of each cell in each time window with an ANOVA (factor: trial type; $p < 0.01$). Neurons that satisfied the criterion in 1 time windows ($n = 932$ cells) were identified as task-related and included in subsequent analyses. Second, we examined how neuronal responses varied as a function of the trial type. Qualitatively, different responses seemed to encode different variables, including the value of individual offers, the spatial configuration of the offers, the chosen value, and the chosen side (Figures 6A-6D). For a quantitative analysis, we defined 12 candidate variables possibly encoded by neurons in OFC (STAR Methods; Table S2). We then regressed each task-related response separately against each variable. If the regression slope differed significantly from zero ($p < 0.05$), then the variable was said to explain the response, and we noted the corresponding R^2 . If a response was explained by more than one variable, we identified the variable providing the best explanation (maximum R^2). Thus, we generated a population table for the number of responses best explained by each variable in each time window (Figure 6E). We aimed to identify a small number of variables that would explain most of the neuronal dataset. Following previous studies,^{2,7} we selected variables using two procedures, namely stepwise and best subset (see STAR Methods). Both procedures selected the same 7 variables, namely position of A, offer value A, offer value B, offer value ipsi, offer value contra, chosen side, and chosen value (Figures 6F-6H; see discussion).

On the basis of these results, we classified each neuron by assigning it to the variable providing the best explanation across time windows (maximum total R^2). Task-related cells that were not explained by any variable were classified as untuned.

The encoding of decision variables is longitudinally stable

We aimed to assess whether individual neurons encoded the same variable over extended periods of time. For each FOV, we selected two sessions based on the behavioral performance and the number of cells (see STAR Methods). We labeled them as session 1 (earlier) and session 2 (later), and we identified neurons that had been recorded and were task-related in both sessions. Repeating this operation for all FOVs provided a dataset of 649 cells. We then examined the classification obtained for each neuron in each session.

The responses recorded for any given neuron in the two sessions often appeared qualitatively similar (Figure S5). For a population analysis, we constructed a contingency table where rows and columns represented the variables encoded in sessions 1 and 2, respectively, and entries were cell counts (Figure 7A). Cells maintaining their functional role would populate the main diagonal. Conversely, neurons encoding different variables in the two sessions would be off diagonal. Neurons that encoded one variable in one session and were untuned in the other session would populate the rightmost column or the bottom row.

Because different variables were encoded with different frequencies, the numbers in Figure 7A are not immediately interpretable. Thus, for each entry in the table, we also computed the corresponding odds ratio (OR; STAR Methods). For each entry in Figure 7B, $OR = 1$ was the chance level, and $OR > 1$ or $OR < 1$ indicated that the cell count was higher or lower than expected by chance, respectively. Two aspects of Figure 7B are most remarkable. First, all diagonal entries were significantly >1 . In other words, the neuronal representation of each variable was significantly more stable than expected by chance. Second, a reasonable hypothesis might be that, over the course of many weeks, some neurons become tuned to the task while other neurons leave the pool. To the contrary, all entries in the rightmost column and bottom row of Figure 7B were at or significantly below chance level. In other words, cells were typically tuned in both sessions or not tuned at all. In summary, the encoding of decision variables in the OFC was very stable (see also Figure S6).

Figure 7 included only two sessions for each FOV. However, our dataset was substantially richer because we recorded from the same FOV 2–16 times over the span of 12–41 weeks. Thus, we conducted a longitudinal assessment of the encoding stability. For each FOV and each pair of sessions, we generated a contingency table and an OR table similar to those displayed in Figure 7. Encoding consistency was ultimately captured by the diagonal nature of the OR table. To quantify this trait, we used the diagonalization index (DI; STAR Methods, Equation 9). We examined how the DI varied as a function of the time distance (ΔT). Our analyses indicated that (1) for any ΔT , the consistency in classification was significantly higher than that expected by chance, and (2) DI did not significantly decrease as a function of ΔT (Figure S7).

DISCUSSION

Neurophysiology studies in primates and rodents suggested that different groups of neurons in OFC constitute the building blocks of a decision circuit. To investigate the structure and, thus, the decision mechanisms, we developed a protocol for 2P Ca²⁺ imaging in mice performing an economic choice task. Imaging of OFC was performed through GRIN lenses. The protocol provided very stable signals in the sense that we were able to record from the same FOV and the same neurons repeatedly for up to 40 weeks. Taking advantage of this preparation, we examined the degree of longitudinal stability in neuronal activity profiles and the representation of decision variables. Several results are noteworthy.

First, we explicitly examined longitudinal changes in behavioral choice patterns. An earlier study noted that capuchin monkeys can effectively choose between pairs of foods never encountered together before,²⁸ suggesting that economic choice is a natural behavior. Work in other species²⁹⁻³¹ and anecdotal observations in macaques support this view. For example, when we train monkeys in a standard juice choice task,⁷ we familiarize animals with the trial structure using only one juice type; subsequently, we introduce choices between two different juice types. Typically, animals present non-trivial trade-offs from the very first session in which two juice types are offered against each other. Importantly, the notion that economic choice is a natural behavior does not exclude that performance in the choice task might improve with experience. Here, we conducted a longitudinal analysis to that effect. Ultimately, task performance is quantified by η .^{26,32} For given ρ , a larger η (i.e., a steeper sigmoid) means that there are fewer trials in which the animal chooses the lower value. We found that η increased over the course of a few weeks. In principle, this trend could be due to multiple factors. (1) While choosing between two juices, mice must decode the offer stimuli — i.e., two odors presented in variable concentrations on the two sides of the nose. Although mice have a sophisticated olfactory system, it seems likely that the perceptual component of the task was somewhat challenging for them and that, over the course of many sessions, animals improved the ability to decode olfactory stimuli. (2) Performance in the choice task certainly required some degree of mental focus. It is quite possible that, over sessions, mice became increasingly capable of focusing on the choice task for many trials. This point is particularly salient if one considers that, in binary choices, there is nominally no error (animals always receive some juice quantity). Supporting the notion of increasing mental focus, the number of trials any animal could perform in any one sitting typically increased with experience. (3) Finally, the decision mechanisms might become more efficacious over time. Longitudinal training improves performance in a variety of natural functions, including perception and motor control; by the same token, longitudinal training could improve performance in decision making. Future research should gauge the relative weights of these factors.

Second, we identified a set of variables encoded by individual neurons in OFC. These include variables associated with individual offers (offer value A, offer value B, offer value ipsi, offer value contra), their spatial configuration (position of A), and the choice outcome (chosen side, chosen value). This finding strongly resonates with the results of neurophysiology studies in primates⁷ and mice.² One difference is that the set of variables identified here includes both juice-based and spatial variables, while previous work in mice

found a prevalence of spatial variables.² We will address this apparent discrepancy in a separate report.

Third, longitudinal analyses provided evidence for high stability in the activity profiles of individual cells and the representation of decision variables. Specifically, activity profiles became more reproducible as the animals became more experienced in the task. To some extent, activity profiles also drifted as a function of time passage. However, this drift was quantitatively modest, and the estimated time necessary for full reorganization exceeded the lifetime of the animal. Concurrently, the encoding of decision variables did not reveal any systematic drift. Considering that only a fraction of OFC neurons appear task-related in any given session, a reasonable hypothesis was that, over the course of many weeks, some neurons would leave the active pool while other neurons would become tuned to the task. Another reasonable hypothesis was that the functional role of any given cell would evolve over time. Our results argue against both hypotheses. Notwithstanding the fact that classification procedures were somewhat noisy (see below), neurons in OFC were typically tuned consistently across sessions or not tuned at all. Furthermore, neurons typically encoded the same variables across sessions. These findings strongly support the notion of a stable decision circuit. Thus, our results set the stage to examine the connectivity of this neural circuit and, ultimately, the winner-takes-all mechanisms underlying economic decisions.

It is interesting to consider our findings in relation to the literature on representational drift. Previous studies reported a broad spectrum of findings—substantial drift was found in sensory and parietal areas,^{12,14-18} and a stable representation was observed in a premotor region of zebra finches,²² while observations in the hippocampus were mixed.¹⁹⁻²¹ Our results place the representation of economic decision variables in OFC on the stable side of this spectrum. Under the assumption that brain regions controlling high cognitive functions are more flexible and/or less rigidly organized,³³ our findings might appear somewhat counterintuitive. At the same time, the fact that decision variables in OFC are encoded categorically³⁴⁻³⁶ likely contributes to the stability of this neuronal representation. Indeed, a neuron drifting from one decision variable to another would have to undertake a quantum leap, qualitatively different from what might take place, for example, in primary visual cortex. In a computational perspective, some degree of stability seems necessary for a circuit performing a winner-takes-all process.

Limitations of the study

Measures of mean(DI) (Figure S7) indicate that our classification procedures are noisy. Indeed, our success rate can be estimated at $\text{mean(DI)}^{1/2} \approx 55\%$. A previous study in monkeys²⁵ using a similar approach provided an estimated success rate of $\approx 80\%$. The difference might be due to two factors, namely that (1) the mouse OFC represents more variables (making the classification more challenging) and (2) mice completed fewer trials each day. Future studies shall address this issue, for example, by classifying each cell based on multiple sessions.

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the lead contact, Camillo Padoa-Schioppa (camillo@wustl.edu).

Materials availability

This study did not generate new unique reagents.

Data and code availability

- The data described in this paper and the Matlab code used for the analysis are available at Mendeley Data <https://data.mendeley.com/datasets/7cg46th7vz/1> (<https://doi.org/10.17632/7cg46th7vz.1>).
- Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

STAR★METHODS

EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS

All experimental procedures were approved by the Washington University Institutional Animal Care and Use Committee (IACUC).

The study was conducted on $N = 13$ mice (6 males, 7 females) expressing the genetically encoded Ca^{2+} indicator GCaMP6f under a pan-neuronal promoter. $N = 10$ C57BL/6J animals (B6; Jackson Laboratory, stock #000664) were injected the pAAV.Syn.GCaMP6f.WPRE.SV40 (Addgene, stock #100837) virus.³⁷ $N = 2$ Rasgrf2-2A-dCre animals (Jackson Laboratory, stock #022864) and $N = 1$ PV-Cre animal (Jackson Laboratory, stock #008069) were injected a 1:1 mixture of pAAV.Syn.GCaMP6f.WPRE.SV40 (Addgene, stock #100837) virus³⁷ and pAAV.CAG.LSL.tdTomato (Addgene, stock #100048) virus.³⁸ In every mouse strain, we aimed to express GCaMP6f under human synapsin (hSYN) promoter. During the imaging experiments, animals were 2.5–7 months old.

METHOD DETAILS

Surgical procedures—All recordings focused on the left hemisphere. In a single surgery, we injected the virus(es), implanted a gradient-index (GRIN) lens, and fixed a head bar to the skull. During surgery, anesthesia was provided by 0.8–1.5% isoflurane in 2 L/min medical air. A heating pad maintained the body temperature at 37°C. Buprenorphine Sustained-Release (SR) (1 mg/kg, subcutaneous (s.c.) injection) was injected for long-term (2 days) pain relief, and dexamethasone (2 mg/kg, intramuscular (i.m.) injection) was injected to limit inflammation. We also injected 2% lidocaine hydrochloride locally at the incision site to reduce pain. Sterilized saline (37°C, 0.50 mL, intraperitoneal (i.p.) injection) was injected at the end of the surgery to avoid dehydration. We first retracted the skin overlying the left and right hemisphere of the mouse to expose the skull overlying frontal

and parietal cortices. We then performed a 2–3 mm diameter craniotomy spanning OFC. Table S3 details the target locations in OFC for each mouse. After removing the skull, we injected the GCaMP6f virus (or a mix of GCaMP6f and LSL-tdTomato viruses) into the brain at two locations (AP: +0.15 mm, ML: –0.15 mm; reference point: target location) and (AP: –0.15 mm, ML: +0.15 mm, reference point: target location) with a Hamilton syringe (Hamilton stock #65458-02). At each location, we injected 100 nL in the tissue at DV depth ranging –0.1 to +0.1 mm with 0.05 mm step size (reference point: target location). The injection speed was 20–40 nL/min and we waited 3 min between steps. Because this type of syringe has a bevel at the needle tip, we corrected the injection depth by adding 0.7 mm (= distance between the tip of the needle and the opening's center). After withdrawing the needle, we slowly inserted a 1 mm diameter GRIN lens (Inscopix ProView lens probe, stock #1050–004605, magnification factor = 1.2, numerical aperture (NA) = 0.5 in air) into the brain and down to OFC (AP: 0 mm, ML: 0 mm, DV: –0.15 to –0.3 mm; reference point: target location). The 0.15–0.3 mm (see Table S3) were subtracted to correct for the fact that this GRIN lens has a working distance ranging 100–300 μ m. This procedure was performed with a GRIN lens holder (Inscopix) at the speed of 0.6 mm/min. During lens insertion, to give some opportunity for the brain to rebound, we retracted 0.2 mm upon first reaching depths of 1, 1.5, and 2 mm. Once in place, the GRIN lens was secured to the skull using C&B Metabond dental cement (Parkell stock #S380). Finally, we attached a custom head bar to the skull using dental cement. After surgery, mice were given dexamethasone (2 mg/kg, i.m.) for inflammation and carprofen (5 mg/kg, s.c.) for pain, if needed. Animals were given 7–14 days to recover before we started training or imaging sessions.

Our procedures for GRIN lens insertion is standard for studies that image deep tissue including OFC.^{24,40} One concern might be whether the presence of a GRIN lens affected the behavior and/or neuronal activity in OFC. Two observations mitigate this possible issue. First, the behavioral performance of animals implanted with a GRIN lens was very similar to that we previously measured in non-implanted mice² and to that recorded by other authors in rats.^{3,41} Second, neuronal responses recorded here using calcium imaging were generally similar to those previously recorded in mice² and rats³ using neurophysiology. For these reasons, we are confident that the lens implantation did not significantly alter OFC function.

Histology—To verify the locations of GRIN lens implantation and GCaMP6f expression, after completing all the *in vivo* imaging experiments, we performed *trans*-cardiac perfusion⁴² to harvest the mouse brains. Mice were anesthetized deeply with ketamine/xylazine cocktail (10 mg/kg, i.p. injection) and then perfused with phosphate-buffered saline (PBS) (mixed with heparin with ratio of 100:1) and 4% paraformaldehyde (PFA). We removed the head bar and GRIN lens. Brains were then extracted from the skull and embedded into 15% then 30% sucrose/PBS solution at 4°C for several days, until they sunk into the bottom of the solution. We then covered the brains with O.C.T. compound (Tissue-Tek) and froze them at –80°C. Brains were sliced in 40 μ m thick coronal sections using cryostats (Leica Biosystems, CM1950) and cover-slipped with antifade mounting medium with DAPI (Vector Laboratories, VECTASHIELD). Brain slices were imaged through inverted microscopy (Leica DMI6000B) with 2.5X, 5X and 10X objectives (Leica Plan Apo). Slice imaging results were compared to the Allen Mouse Brain Atlas (Coronal

Atlas) from AP: +2.60 mm to +2.80 mm. We used the lesion gap caused by the GRIN lens as a landmark to check whether recorded FOVs were within OFC, accounting for the GRIN lens range of working distances.

Choice task and training protocol—The choice task closely resembled that used for non-human primates⁷ and was nearly identical to that described in a previous mouse neurophysiology study.² During training or imaging sessions, mice were head-fixed and placed in a customized tube under the microscope. In each session, the animal chose between two different juices labeled A and B (juice A preferred) and offered invariable quantities. In each trial, we delivered two odor stimuli (octanal and octanol) simultaneously from the left and right of the mouse nose (Figure 1A). The odor identities represented the juice type, and the odor concentration represented the juice quantity. For juice A, we used 5 different concentrations to represent 5 quantities (1–5 drops). For juice B, we used 4 or 5 different concentrations to represent 4 (2,3,4,7 drops) or 5 quantities (2–6 drops or 1,2,3,5,6 drops), respectively. The association between odors and juices was balanced across mice (in 8 animals, juices A/B were paired with octanal/octanol, respectively; in 5 animals, the association was reversed). The offered quantities were linearly related to the odor concentration (e.g., odor levels 1 to maximum number of ppm linearly represented quantity 1 to maximum number of juice). Offered quantities q_A and q_B (and the corresponding odor concentrations) varied pseudo-randomly from trial to trial. For any two quantities, the spatial configuration of the offers (left/right) varied pseudo-randomly from trial to trial.

Odors were presented for 2.4 s (offer period). Immediately thereafter, a 0.2 s tone was played (go signal). In each trial, the animal indicated its choice by licking one of two spouts. The animal had to indicate its choice within 5 s. The lick was detected by a battery-operated touch circuit,⁴³ and the chosen juice was immediately delivered. Juice delivery typically took 150–1050 ms (~150 ms/per drop). If the animal did not respond within 5 s, the trial was aborted. Licks preceding the onset of the go signal were ignored and did not result in any juice delivery. In forced choice trials (~35% of trials), only one odor was presented on one side (one correct response). If the animal licked the wrong spout, a 2–5 s white noise sound was played. The exact length of the white noise sound varied across sessions but was fixed within each session. Between two trials, we activated a vacuum system that removed the odors remaining from the previous trial. The inter-trial interval (ITI) varies in the range 2.5–5 s.

On any given day, mice typically performed the choice task for 30 min (1 session; 170–270 trials) and received ~0.7–1.3 mL liquid. Juice pairs were varied across mice and included: (A) 15% sucrose vs. (B) water; (A) apple juice vs. (B) blueberry juice; (A) apple juice vs. (B) water; (A) 15% sucrose vs. (B) cranberry juice; (A) 15% sucrose vs. (B) grape juice; (A) water vs. (B) apple juice; (A) pomegranate juice vs. (B) water; (A) apple juice vs. (B) white grape juice; (A) white grape juice vs. (B) water; (A) grape juice vs. (B) cranberry juice; (A) 15% sucrose vs. (B) cranberry juice; (A) 15% sucrose vs. (B) peppermint tea; (A) elderberry juice vs. (B) white grape juice; (A) elderberry juice vs. (B) water; (A) 15% sucrose vs. (B) elderberry juice.

Training in this task typically took 2.5–4 months and proceeded as described by Kuwabara et al.² Briefly, in the first phase of training, animals were presented only one odor in each trial— either octanal (associated with juice A) or octanol (associated with juice B) – on the left or on the right. To obtain the juice, they had to lick the corresponding spout (forced choices). Over sessions, we introduced different odor concentrations corresponding to different juice quantities. Once performance reached 80% correct, we moved to the second phase of training, in which we delivered both odors simultaneously and the animal had to choose between the two options (binary choice). Each training phase typically took 1–1.5 months. Throughout the training period, the animal weight was maintained at 75%–85% of baseline.⁴⁴ Neural recordings started from the beginning of the second phase.

Procedures for two-photon calcium imaging—2P Ca²⁺ imaging was performed with an Ultima system (Prairie Technologies) built on an Olympus BX61W1 microscope. A mode-locked laser (Mai Tai DeepSee Ti:Sapphire laser, Spectra-Physics) tuned to 920 nm was raster scanned at 7–11 Hz for excitation while emitted GCaMP6f fluorescence was collected through a green filter (525/70 nm). Laser power at the sample was 20–60 mW. Dwell time was set to 3.2–4 μ s. To increase the imaging speed, resolution along the y axis was reduced by a factor of 2. The final pixel size was $2.29 \times 4.58 \mu$ m. We used a 20×0.45 NA objective (Olympus LCPLN20 \times IR objective). Thus in each animal we could image a cylinder of 850 μ m diameter and 350 μ m depth. By varying the focal depth, we could image from multiple fields of view (FOVs). For data collection, we restricted the imaging to depths between 120 and 370 μ m, with 90% of FOVs placed between 180 and 330 μ m. Within a FOV, we often selected a smaller region of interest, to achieve an imaging frequency of 7 Hz. Throughout this paper, we refer to this region of interest as the FOV.

QUANTIFICATION AND STATISTICAL ANALYSIS

Logistic analysis of choice data—All the analyses of behavioral and neuronal data were performed in MATLAB (MathWorks). Choice patterns were analyzed using logistic regressions.³² For each session, we focused on binary choice trials (i.e., we excluded forced choices) and we built the following logistic model:

$$P_B = 1 / (1 + \exp(-X))$$

$$X = a_0 + a_1 \log(q_B / q_A) + a_2(\delta_{A, \text{right}} - \delta_{B, \text{right}})$$

(Equation 1)

In Equation 1, P_B is the probability of choosing juice B; q_A and q_B are the offered quantities; $\delta_{A, \text{right}} = 1$ if juice A was offered on the right and 0 otherwise, and $\delta_{B, \text{right}} = 1 - \delta_{A, \text{right}}$. From the fitted parameters a_0 and a_1 , we computed the relative value of the two juices $\rho = \exp(-a_0 / a_1)$, the choice accuracy $\eta = a_1$, and the side bias $\varepsilon = -a_2 / a_1$. Intuitively, ρ is the quantity ratio q_B / q_A that makes the animal indifferent between the two juices (Figure 1C). The parameter η captures the choice accuracy (or choice consistency), also referred to as inverse temperature; it is proportional to the sigmoid steepness and inversely related to choice variability. The parameter ε quantifies a bias favoring one of the two sides; specifically, $\varepsilon < 0$ and $\varepsilon > 0$ indicate a choice bias favoring the left and right offers, respectively.

In general, the side bias measured in our experiments was relatively modest ($|\varepsilon| < \rho$). Thus in the analysis of neuronal activity, we defined candidate variables based on the relative value, disregarding the side bias. To increase our statistical power in computing these variables, we built a reduced logistic model:

$$\begin{aligned} P_B &= 1 / (1 + \exp(-X)) \\ X &= a_0 + a_1 \log(q_B / q_A) \end{aligned} \quad (\text{Equation 2})$$

From parameters fitted with Equation 2, we computed the relative value $\rho = \exp(-a_0 / a_1)$ and the sigmoid steepness $\eta = a_1$. However, for some of the analyses, we imposed restrictive criteria, including only sessions with $|\varepsilon| \leq 1$ (cosine similarity analysis and odd ratios analysis; Fig. 4; Fig. 7) or $|\varepsilon| \leq 0.5$ (population analysis; Fig. 6).

Longitudinal analysis of choice behavior—We conducted a longitudinal analysis of choice behavior. For each mouse, we counted days starting with the first in which the animal chose between two different juices ($T = 0$). For each session, we derived the parameters ρ , η , and ε from the logistic fit (Equation 1). For each animal and each parameter, we removed outliers (i.e., data points that differed >3 standard deviations from the mean, see Figure 2A), and we fitted the remaining values with an exponential function:

$$X(T) = a + b \exp(-c) \quad (\text{Equation 3})$$

In Equation 3, T indicates the day of the session and $X = \rho$, η , or ε . For each fitted parameter a , b , and c , we obtained a 95% confidence interval. Fitted parameter a captured the plateau value (for $T \rightarrow \infty$, $X \rightarrow a$). If the 95% confidence interval for either parameter b or c did not include 0, there was evidence for a significant longitudinal trend; in this case, c was the inverse time constant of this trend.

We performed this analysis for each animal and each behavioral parameter ρ , η , and ε . One animal (M1062) was excluded from the analysis because we only had data from 4 sessions. We also conducted a population analysis focusing on the choice accuracy η and pooling sessions from all 13 mice. We removed outliers (data points that differed >3 standard deviations from the mean) and we fitted the whole dataset with a single exponential function (Equation 3).

Cell segmentation and session registration—Cell segmentation was conducted using the CaImAn package.²⁷ First, we performed a frame-to-frame non-rigid motion correction to remove potential movement of individual neural signals. Then we extracted and estimated the neuronal fluorescent signal with constrained non-negative matrix factorization (CNMF). We set the criteria to identify a neuron: the signal-to-noise ratio should be > 1 . After that, we manually checked individual neurons and excluded those without a donut-shaped cell morphology. This manual checking process is blind of cell's functional property (i.e., the encoded variable), but only relied on the shape of its spatial footprint. For each cell

identified with these procedures, we normalized the fluorescence trace F at each time point to a fixed baseline generated from the sum of the 20th percentile of the average F signal and background signal (generated from CNMF) within each cell region and within a 200 frames time window. This normalization resulted in the $\Delta F / F$ signal, which was used for all subsequent analyses.

Data for this study came from 13 mice and 77 FOVs. For each mouse, recording started on the first day of the second training phase – i.e., when the animal first chose between different juice flavors. Each FOV was recorded in 2–16 sessions (mean = 5.95, std = 2.64). For each FOV and each pair of sessions, we registered images and matched neurons using a Bayesian procedure.³⁹ First, we aligned the images of the two FOVs by projecting their neuronal footprints onto the same image and finding the rotations and translations that yield the highest cross-correlation between their projections. Then we modeled the distributions of centroid distances and spatial correlations as a weighted sum of the distributions of two subpopulations of cell-pairs, representing same cells and different cells. Using Bayes' rule, we obtained the probability (P_{same}) for any pair of neighboring cells from different sessions to be the same cell, given their spatial correlation and centroid distance. Finally, we imposed a threshold of $P_{\text{same}} > 0.5$ (following previous work³⁹). Cell pairs that satisfied this criterion were identified as matching. Using these procedures, we typically found that 60%–80% of cells from any given session could be matched with cells recorded in the same FOV in another session.

Cosine similarity analysis of activity profiles—Referring to the choice task, an offer type was defined by two offered quantities (q_A, q_B); a trial type was defined by two offered quantities, their spatial configuration, and a choice. Because neural activity was recorded at different imaging frequency in different sessions, for the analysis of activity profiles, we first interpolated each cell's time course and resampled it at 20 Hz. We focused on the $\Delta F / F$ signal recorded in two time windows aligned with the offer onset (–600 ms, +1600 ms) and the first lick (–600 ms, +600 ms). For each trial, we joined these two time windows and obtained a single trace. We then averaged traces across trials for each trial type. The activity profile was defined as the concatenation of mean traces obtained for different trial types. Importantly, when comparing the activity profiles recorded for one cell in two sessions, we only included trial types present in both sessions.

The likeness of activity profiles was quantified using cosine similarity (CS). For any neuron i recorded in sessions 1 and 2, we indicated the two activity profiles as $AP_{i,1}$ and $AP_{i,2}$, respectively. The CS was defined as:

$$CS_{\text{same cells}} = \frac{AP_{i,1} \cdot AP_{i,2}}{|AP_{i,1}| \cdot |AP_{i,2}|} \quad (\text{Equation 4})$$

where “ \cdot ” indicates a vector product and “ $|x|$ ” indicates the Euclidean norm of the vector x .

Because our measures of $\Delta F / F$ were mostly >0 , AP vectors were effectively confined to a single orthant of a high-dimensional space. Consequently, CS was expected to be high. To obtain chance level estimates for CS, we generated two benchmark treatments. First, given two sessions 1 and 2, we randomly paired each neuron of session 1 with a neuron of session 2 (treatment: mismatched neurons). Indicating with i and j two randomly paired neurons in the FOV, we computed:

$$CS_{\text{mismatched cells}} = \frac{AP_{i,1} \cdot AP_{j,2}}{|AP_{i,1}| \cdot |AP_{j,2}|} \quad (\text{Equation 5})$$

We repeated this operation 10,000 times and we obtained a distribution for $CS_{\text{mismatched cells}}$. Second, for each (matched) neuron i , we shuffled the activity profile in session 2 by randomly permuting the vector components (treatment: shuffled times). We then computed:

$$CS_{\text{shuffled times}} = \frac{AP_{i,1} \cdot AP_{i,2,\text{shuffled}}}{|AP_{i,1}| \cdot |AP_{i,2,\text{shuffled}}|} \quad (\text{Equation 6})$$

We repeated this operation 10,000 times for each cell, and we obtained a distribution for $CS_{\text{shuffled times}}$.

Longitudinal analysis of cosine similarity—For longitudinal analyses, we counted days for each animal starting the first day in which they chose between different juices (see above). Recordings continued for 12–41 week, depending on the animal. For any two sessions 1 and 2, we indicate with T1 and T2 the days in which they took place (always $T1 \leq T2$). We also defined $\Delta T = T2 - T1$. On some occasions, we ran two sessions back-to-back in the same day ($\Delta T = 0$). For a longitudinal analysis of activity profiles, we examined CS as a function of the experience in the task (T1) and the time difference between two sessions (ΔT). We partitioned the (T1, ΔT) parameter space in 2-week x 2-week bins. For given T1 and ΔT , we considered the distribution of CSs obtained across animals and FOVs, and we computed the median(CS). Thus we obtained values of median(CS) as a function of T1 and ΔT . These operations were repeated for the three treatments (same neurons, mismatched neurons, shuffled times). Statistical comparisons between distributions of CSs were conducted using the Kruskal-Wallis test. To increase the statistical power, for some analyses we defined larger time intervals (6-week x 6-week), referred to as Bins (see results).

To precisely quantify longitudinal drifts in neuronal activity profiles, we examined how CS varied as a function of T1 and ΔT with a 1 day temporal resolution. In drift diffusion processes, the drifting variable is stochastic and the expected distance from the initial position increases with the square root of time. Thus we model $CS(T)$ with a square root function:

$$CS(T) = a_0 + a_1 T^{1/2}$$

(Equation 7)

where $T = T1$ or ΔT . This analysis indicated that the time to full reorganization would exceed the animal's typical life time. For a control, we also modeled $CS(T)$ with an exponential function and obtained very similar results.

Cosine similarity, choice accuracy, and task experience in individual animals

—We examined the relationship between cosine similarity (CS), choice accuracy (η), and experience in the task in individual animals. For each mouse, we performed a linear regression of CS onto η . Then, for each session, we computed CS' , which was defined as the cosine similarity estimated based on η and on the fitted linear relationship (Figure S4). For each session we computed the variable X defined as the geometric mean of CS (measured) and CS' (estimated) – in formulas, $X = (CS \cdot CS')^{1/2}$. Finally, we studied the relation between the variable X and the time to the last session ($T3$). This relation was examined using linear regressions for individual animals and an ANCOVA (independent lines). Importantly, sessions with very few neurons (<10th percentile) and sessions for which the measure of CS' was an outlier were excluded from this analysis. Outliers were identified using the interquartile criterion.⁴⁵ Defining IQR as the interquartile range, values below the first quartile minus 1.5*IQR or above the third quartile plus 1.5*IQR were identified as outliers and excluded.

Neuronal encoding of decision variables—We conducted a series of analyses to identify the variables encoded by individual neurons in OFC using the same procedures previously adopted for the analysis of spike counts.^{2,7} Previous work in monkeys and rodents found that neurons in OFC present categorical encoding,³⁴⁻³⁶ and that different groups of cells in this area encode different variables.^{2,3,7,8,35,46-49} Importantly, these variables are often highly correlated (e.g., *offer value* and *chosen value*). In this situation, to identify the variables encoded by the population, it is necessary to construct an analysis that (a) includes many candidate variables and (b) incorporates a principled algorithm for variable selection. The analysis adopted here was designed exactly for this purpose (for details, see [7]) and has been used in numerous studies.^{2,3,7,8,35,46,47} For these analyses, we selected one session per FOV based on two criteria: (1) we only considered sessions with $|\epsilon| \leq 0.5$ and (2) among these sessions, we selected the one with the largest number of recorded cells. Neural activity ($\Delta F / F$) was examined in five time windows: pre-offer (0.6–0 s before the offer onset), post-offer (0.4–1 s after the offer onset), late delay (1–1.6 s after the offer onset), pre-juice (0.6–0 s before juice delivery onset), post-juice (0–0.6 s after juice delivery onset) (Figure 1B). Notably, these time windows were almost identical to those previously defined in neurophysiology studies in monkeys and mice^{2,7} (here we used 600 ms windows instead of 500 ms windows because GCaMP6 signals are slow compared to neurophysiology traces⁵⁰). Again, arial type was defined by two offered quantities, their spatial configuration, and a choice. For each cell, each time window, and each trial type, we averaged the $\Delta F / F$ activity the time window and over trials. A neuronal response was

defined as the average activity of one neuron in one time window as a function of the trial type.

Task-related responses were identified with an ANOVA (factor: trial type; $p < 0.01$). Neurons that passed this criterion in at least one time window were identified as task-related and included in subsequent analyses. Following previous studies, we defined twelve candidate variables that neurons in OFC might encode (Table S2). We performed a linear regression of each response passing the ANOVA criterion on each variable, and obtained the corresponding R^2 . If the regression slope differed significantly from zero ($p < 0.05$), the variable was said to explain the response. If the regression slope was indistinguishable from zero, we arbitrarily set the corresponding R^2 to zero. Neuronal responses were often explained by more than one variable; the variable providing the best fit was that with the highest R^2 .

To identify a limited number of variables that could explain most of the neuronal population, we proceeded with a variable selection analysis. As in previous studies, we used a stepwise procedure and a best-subset procedure.^{7,51} The stepwise procedure selects at each iteration the variable providing the maximum explanatory power, considering time windows separately. Once a variable is selected, all the responses explained by that variable (including those that would be explained by some other variable with a higher R^2) are removed from the dataset. The residual dataset is then examined, and a new variable is selected. The marginal explanatory power of a variable is defined as the fraction of responses that are explained by that variable and that are not explained by any other selected variable. The algorithm continues until the marginal explanatory power of newly selected variables is $< 2\%$ of the total. Importantly, at each iteration we verify that the marginal explanatory power of each selected variable satisfies this criterion. Previously selected variables whose marginal explanatory power drops below 2% in the presence of a newly selected variable would be eliminated from the selected subset. In practice, however, this situation did not occur for the present data.

Importantly, the stepwise procedure does not guarantee optimality. In contrast, the best-subset procedure is exhaustive. For $k = 2, 3, \dots$, the procedure examines all the possible subsets of k variables, computes the total explanatory power (i.e., the number of responses collectively explained), and then selects the subset providing the maximum explanatory power. In this study, the two procedures yielded the same results.

Based on the results of the variable selection analysis, we classified each neuron by assigning it to the variable providing the largest total R^2 across time windows. This sum was restricted to the time windows for which the cell passed the ANOVA criterion. Task-related cells that were not explained by any variable were classified as *untuned*. Subsequent analyses on the longitudinal stability of neuronal functional roles were based on this classification.

Stability in the representation of decision variables: Analysis of odds ratios—

We conducted a series of analyses to assess whether the variable encoded by any given cell remained stable over extended periods of time. Intuitively, our approach was as follows. For

each FOV and any two sessions, we focused on cells identified (matching) and task-related in both sessions. For each cell, we considered variables encoded in the two sessions; if in one session the cell was task-related but did not encode any variable, the cell was classified as *untuned* for that session. On this basis, we constructed for each FOV and for any two sessions a contingency table where rows and columns corresponded to 8 possible classes (7 variables plus *untuned*) recorded in the two sessions, and entries indicated cell counts. In this table, neurons with consistent encoding populated the main diagonal.

Based on this general idea, we conducted a population analysis as follows. First, for each FOV in our dataset, we selected two sessions based on two criteria: (1) we imposed $|\epsilon| \leq 1$, and (2) we selected the two sessions with the largest number of matching, task-related cells. Second, we pooled data across FOVs and examined the resulting table of cell counts (Figure 7A). Third, to assess whether cell counts measured for any pair of variables was above or below chance level, we computed the corresponding table of odds ratios (OR).⁵² Indicating with C the table of cell counts, for any entry (p, q) , OR was defined as follows:

$$\begin{aligned} a_{1,1} &= C_{p,q} \\ a_{1,2} &= \sum_{j \neq q} C_{p,j} \\ a_{2,1} &= \sum_{i \neq p} C_{i,q} \\ a_{2,2} &= \sum_{i \neq p, j \neq q} C_{i,j} \\ \text{OR} &= \frac{a_{1,1} a_{2,2}}{a_{1,2} a_{2,1}} \end{aligned}$$

(Equation 8)

We computed OR for each entry and obtained the full table (Figure 7B). Importantly, for each entry, the chance level was $\text{OR} = 1$. Conversely, $\text{OR} > 1$ ($\text{OR} < 1$) indicated that the cell count for that entry was above (below) chance. For each entry, we used Fisher's exact test (two tails)⁵³ to assess whether the departure from chance was statistically significant.

Analysis of diagonalization index—The degree of stability in the neuronal representation of decision variables was captured by the relative weight of the diagonal in the OR table. To quantify this weight, we used the diagonalization index (DI), defined as follows. For a $d \times d$ nonnegative matrix A ,

$$\text{DI} = \frac{\sum_{i=1 \dots d} A_{i,i}}{\left(\sum_{i=1 \dots d} \sum_{j=1 \dots d} A_{i,j} \right)}$$

(Equation 9)

Of note, DI ranges between 0 and 1, and $\text{DI} = 1$ corresponds to a perfectly diagonal matrix – i.e., faultless stability. More generally, the closer DI is to 1, the more stable is the representation.

If all variables were encoded by a similar fraction of cells, the chance level for DI would be $1/n$, where n is the number of variables (including *untuned*). However, different variables were encoded with different frequencies. Thus to obtain chance level estimates for DI, we generated a control measurement. Given two sessions 1 and 2, we randomly paired each neuron of session 1 with a neuron of session 2 (treatment: mismatched cells). We then generated the contingency table and calculated DI. We repeated this procedure 10,000 times and we averaged the DIs. For longitudinal analyses, we repeated this procedure for each pair of sessions.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGMENTS

We thank J. Assad and E. Han for comments on the manuscript and members of the Padoa-Schioppa and Holy laboratories for helpful discussions. This work was supported by the National Institutes of Health (grants R01-DA055709, R01-DA032758, and R21-DA042882 to C.P.-S. and R01-DC020034 to T.E.H.).

REFERENCES

- Gardner MPH, Sanchez D, Conroy JC, Wikenheiser AM, Zhou J, and Schoenbaum G (2020). Processing in lateral orbitofrontal cortex is required to estimate subjective preference during initial, but not established, economic choice. *Neuron* 108, 526–537.e4. 10.1016/j.neuron.2020.08.010. [PubMed: 32888408]
- Kuwabara M, Kang N, Holy TE, and Padoa-Schioppa C (2020). Neural mechanisms of economic choices in mice. *Elife* 9, e49669. 10.7554/eLife.49669. [PubMed: 32096761]
- Gore F, Hernandez M, Ramakrishnan C, Crow AK, Malenka RC, and Deisseroth K (2023). Orbitofrontal cortex control of striatum leads economic decision-making. *Nat. Neurosci* 26, 1566–1574. 10.1038/s41593-023-01409-1. [PubMed: 37592039]
- Yu LQ, Dana J, and Kable JW (2022). Individuals with ventromedial frontal damage display unstable but transitive preferences during decision making. *Nat. Commun* 13, 4758. 10.1038/s41467-022-32511-w. [PubMed: 35963856]
- Ballesta S, Shi W, Conen KE, and Padoa-Schioppa C (2020). Values encoded in orbitofrontal cortex are causally related to economic choices. *Nature* 588, 450–453. 10.1038/s41586-020-2880-x. [PubMed: 33139951]
- Ballesta S, Shi W, and Padoa-Schioppa C (2022). Orbitofrontal cortex contributes to the comparison of values underlying economic choices. *Nat. Commun* 13, 4405. 10.1038/s41467-022-32199-y. [PubMed: 35906242]
- Padoa-Schioppa C, and Assad JA (2006). Neurons in orbitofrontal cortex encode economic value. *Nature* 441, 223–226. [PubMed: 16633341]
- Pastor-Bernier A, Stasiak A, and Schultz W (2019). Orbitofrontal signals for two-component choice options comply with indifference curves of revealed preference theory. *Nat. Commun* 10, 4885. 10.1038/s41467-019-12792-4. [PubMed: 31653852]
- Rich EL, and Wallis JD (2016). Decoding subjective decisions from orbitofrontal cortex. *Nat. Neurosci* 19, 973–980. 10.1038/nn.4320. [PubMed: 27273768]
- Balewski ZZ, Elston TW, Knudsen EB, and Wallis JD (2023). Value dynamics affect choice preparation during decision-making. *Nat. Neurosci* 26, 1575–1583. 10.1038/s41593-023-01407-3. [PubMed: 37563295]
- Driscoll LN, Duncker L, and Harvey CD (2022). Representational drift: Emerging theories for continual learning and experimental future directions. *Curr. Opin. Neurobiol* 76, 102609. 10.1016/j.conb.2022.102609. [PubMed: 35939861]

12. Driscoll LN, Pettit NL, Minderer M, Chetih SN, and Harvey CD (2017). Dynamic reorganization of neuronal activity patterns in parietal cortex. *Cell* 170, 986–999.e16. 10.1016/j.cell.2017.07.021. [PubMed: 28823559]
13. Rule ME, Loback AR, Raman DV, Driscoll LN, Harvey CD, and O’Leary T (2020). Stable task information from an unstable neural population. *Elife* 9, e51121. 10.7554/eLife.51121. [PubMed: 32660692]
14. Deitch D, Rubin A, and Ziv Y (2021). Representational drift in the mouse visual cortex. *Curr. Biol* 31, 4327–4339.e6. 10.1016/j.cub.2021.07.062. [PubMed: 34433077]
15. Schoonover CE, Ohashi SN, Axel R, and Fink AJP (2021). Representational drift in primary olfactory cortex. *Nature* 594, 541–546. 10.1038/s41586-021-03628-7. [PubMed: 34108681]
16. Li M, Liu F, Jiang H, Lee TS, and Tang S (2017). Long-term two-photon imaging in awake macaque monkey. *Neuron* 93, 1049–1057.e3. 10.1016/j.neuron.2017.01.027. [PubMed: 28215557]
17. Marks TD, and Goard MJ (2021). Stimulus-dependent representational drift in primary visual cortex. *Nat. Commun* 12, 5169. 10.1038/s41467-021-25436-3. [PubMed: 34453051]
18. Xia J, Marks TD, Goard MJ, and Wessel R (2021). Stable representation of a naturalistic movie emerges from episodic activity with gain variability. *Nat. Commun* 12, 5170. 10.1038/s41467-021-25437-2. [PubMed: 34453045]
19. Ziv Y, Burns LD, Cocker ED, Hamel EO, Ghosh KK, Kitch LJ, El Gamal A, and Schnitzer MJ (2013). Long-term dynamics of CA1 hippocampal place codes. *Nat. Neurosci* 16, 264–266. 10.1038/nn.3329. [PubMed: 23396101]
20. Lee JS, Briguglio JJ, Cohen JD, Romani S, and Lee AK (2020). The statistical structure of the hippocampal code for space as a function of time, context, and value. *Cell* 183, 620–635.e22. 10.1016/j.cell.2020.09.024. [PubMed: 33035454]
21. Geva N, Deitch D, Rubin A, and Ziv Y (2023). Time and experience differentially affect distinct aspects of hippocampal representational drift. *Neuron* 111, 2357–2366.e5. 10.1016/j.neuron.2023.05.005. [PubMed: 37315556]
22. Katlowitz KA, Picardo MA, and Long MA (2018). Stable sequential activity underlying the maintenance of a precisely executed skilled behavior. *Neuron* 98, 1133–1140.e3. 10.1016/j.neuron.2018.05.017. [PubMed: 29861283]
23. Muysers H, Chen HL, Hahn J, Folschweiller S, Sigurdsson T, Sauer JF, and Bartos M (2024). A persistent prefrontal reference frame across time and task rules. *Nat. Commun* 15, 2115. 10.1038/s41467-024-46350-4. [PubMed: 38459033]
24. Namboodiri VMK, Otis JM, van Heeswijk K, Voets ES, Alghorazi RA, Rodriguez-Romaguera J, Mihalas S, and Stuber GD (2019). Single-cell activity tracking reveals that orbitofrontal neurons acquire and maintain a long-term memory to guide behavioral adaptation. *Nat. Neurosci* 22, 1110–1121. 10.1038/s41593-019-0408-1. [PubMed: 31160741]
25. Xie J, and Padoa-Schioppa C (2016). Neuronal remapping and circuit persistence in economic decisions. *Nat. Neurosci* 19, 855–861. 10.1038/nn.4300. [PubMed: 27159800]
26. Rustichini A, Conen KE, Cai X, and Padoa-Schioppa C (2017). Optimal coding and neuronal adaptation in economic decisions. *Nat. Commun* 8, 1208. 10.1038/s41467-017-01373-y. [PubMed: 29084949]
27. Giovannucci A, Friedrich J, Gunn P, Kalfon J, Brown BL, Koay SA, Taxidis J, Najafi F, Gauthier JL, Zhou P, et al. (2019). CaImAn an open source tool for scalable calcium imaging data analysis. *Elife* 8, e38173. 10.7554/eLife.38173. [PubMed: 30652683]
28. Padoa-Schioppa C, Jandolo L, and Visalberghi E (2006). Multi-stage mental process for economic choice in capuchins. *Cognition* 99, B1–B13. [PubMed: 16043168]
29. Noë R, Van Hooff JA, and Hammerstein P (2001). *Economics in Nature : Social Dilemmas, Mate Choice, and Biological Markets* (New York: NY Cambridge University Press).
30. Hurley SL, and Nudds M (2006). *Rational Animals?* (Oxford University Press).
31. Bourgeois-Gironde S, Addeksi E, and Boraud T (2021). Economic behaviours among non-human primates. *Philos. Trans. R. Soc. Lond. B Biol. Sci* 376, 20190676. 10.1098/rstb.2019.0676. [PubMed: 33423625]
32. Padoa-Schioppa C. (2022). Logistic analysis of choice data: A primer. *Neuron* 110, 1615–1630. 10.1016/j.neuron.2022.03.002. [PubMed: 35334232]

33. Miller EK, and Cohen JD (2001). An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci* 24, 167–202. 10.1146/annurev.neuro.24.1.16724/1/167 [pii]. [PubMed: 11283309]
34. Onken A, Xie J, Panzeri S, and Padoa-Schioppa C (2019). Categorical encoding of decision variables in orbitofrontal cortex. *PLoS Comput. Biol* 15, e1006667. 10.1371/journal.pcbi.1006667. [PubMed: 31609973]
35. Hirokawa J, Vaughan A, Masset P, Ott T, and Kepecs A (2019). Frontal cortex neuron types categorically encode single decision variables. *Nature* 576, 446–451. 10.1038/s41586-019-1816-9. [PubMed: 31801999]
36. Padoa-Schioppa C. (2013). Neuronal origins of choice variability in economic decisions. *Neuron* 80, 1322–1336. 10.1016/j.neuron.2013.09.013. [PubMed: 24314733]
37. Chen TW, Wardill TJ, Sun Y, Pulver SR, Renninger SL, Baohan A, Schreiter ER, Kerr RA, Orger MB, Jayaraman V, et al. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* 499, 295–300. 10.1038/nature12354. [PubMed: 23868258]
38. Oh SW, Harris JA, Ng L, Winslow B, Cain N, Mihalas S, Wang Q, Lau C, Kuan L, Henry AM, et al. (2014). A mesoscale connectome of the mouse brain. *Nature* 508, 207–214. 10.1038/nature13186. [PubMed: 24695228]
39. Sheintuch L, Rubin A, Brande-Eilat N, Geva N, Sadeh N, Pinchasof O, and Ziv Y (2017). Tracking the same neurons across multiple days in Ca(2+) imaging data. *Cell Rep.* 21, 1102–1115. 10.1016/j.celrep.2017.10.013. [PubMed: 29069591]
40. Jennings JH, Kim CK, Marshel JH, Raffiee M, Ye L, Quirin S, Pak S, Ramakrishnan C, and Deisseroth K (2019). Interacting neural ensembles in orbitofrontal cortex for social and feeding behaviour. *Nature* 565, 645–649. 10.1038/s41586-018-0866-8. [PubMed: 30651638]
41. Gardner MPH, Conroy JS, Shaham MH, Styer CV, and Schoenbaum G (2017). Lateral orbitofrontal inactivation dissociates devaluation-sensitive behavior and economic choice. *Neuron* 96, 1192–1203.e4. 10.1016/j.neuron.2017.10.026. [PubMed: 29154127]
42. Wu J, Cai Y, Wu X, Ying Y, Tai Y, and He M (2021). Transcardiac Perfusion of the Mouse for Brain Tissue Dissection and Fixation. *Bio. Protoc* 11, e3988. 10.21769/BioProtoc.3988.
43. Slotnick B (2009). A simple 2-transistor touch or lick detector circuit. *J. Exp. Anal. Behav* 91, 253–255. 10.1901/jeab.2009.91-253. [PubMed: 19794837]
44. Guo ZV, Hires SA, Li N, O'Connor DH, Komiyama T, Ophir E, Huber D, Bonardi C, Morandell K, Gutnisky D, et al. (2014). Procedures for behavioral experiments in head-fixed mice. *PLoS One* 9, e88678. 10.1371/journal.pone.0088678. [PubMed: 24520413]
45. Tukey JW (1977). *Exploratory Data Analysis* (Addison-Wesley Pub. Co.).
46. Ballesta S, and Padoa-Schioppa C (2019). Economic decisions through circuit inhibition. *Curr. Biol* 29, 3814–3824.e5. 10.1016/j.cub.2019.09.027. [PubMed: 31679936]
47. Hosokawa T, Kennerley SW, Sloan J, and Wallis JD (2013). Single-neuron mechanisms underlying cost-benefit analysis in frontal cortex. *J. Neurosci* 33, 17385–17397. 10.1523/JNEUROSCI.2221-13.2013. [PubMed: 24174671]
48. Yun M, Kawai T, Nejime M, Yamada H, and Matsumoto M (2020). Signal dynamics of midbrain dopamine neurons during economic decision-making in monkeys. *Sci. Adv* 6, eaba4962. 10.1126/sciadv.aba4962. [PubMed: 32937434]
49. Masset P, Ott T, Lak A, Hirokawa J, and Kepecs A (2020). Behavior- and modality-general representation of confidence in orbitofrontal cortex. *Cell* 182, 112–126.e18. 10.1016/j.cell.2020.05.022. [PubMed: 32504542]
50. Zhang Y, Rózsa M, Liang Y, Bushey D, Wei Z, Zheng J, Reep D, Broussard GJ, Tsang A, Tsegaye G, et al. (2023). Fast and sensitive GCaMP calcium indicators for imaging neural populations. *Nature* 615, 884–891. 10.1038/s41586-023-05828-9. [PubMed: 36922596]
51. Glantz SA, and Slinker BK (2001). *Primer of Applied Regression & Analysis of Variance*, 2nd Edition (McGraw-Hill, Medical Pub. Division).
52. Szumilas M. (2010). Explaining odds ratios. *J Can Acad Child Adolesc Psychiatry* 19, 227–229. [PubMed: 20842279]
53. Agresti A. (2019). *An Introduction to Categorical Data Analysis*, Third edition (Edition: Wiley).

Highlights

- Mice choose between different juice flavors based on subjective values
- Using two-photon calcium imaging, neurons in OFC are recorded for up to 40 weeks
- The activity profiles of individual cells are longitudinally stable
- The decision variables encoded by individual cells are longitudinally stable

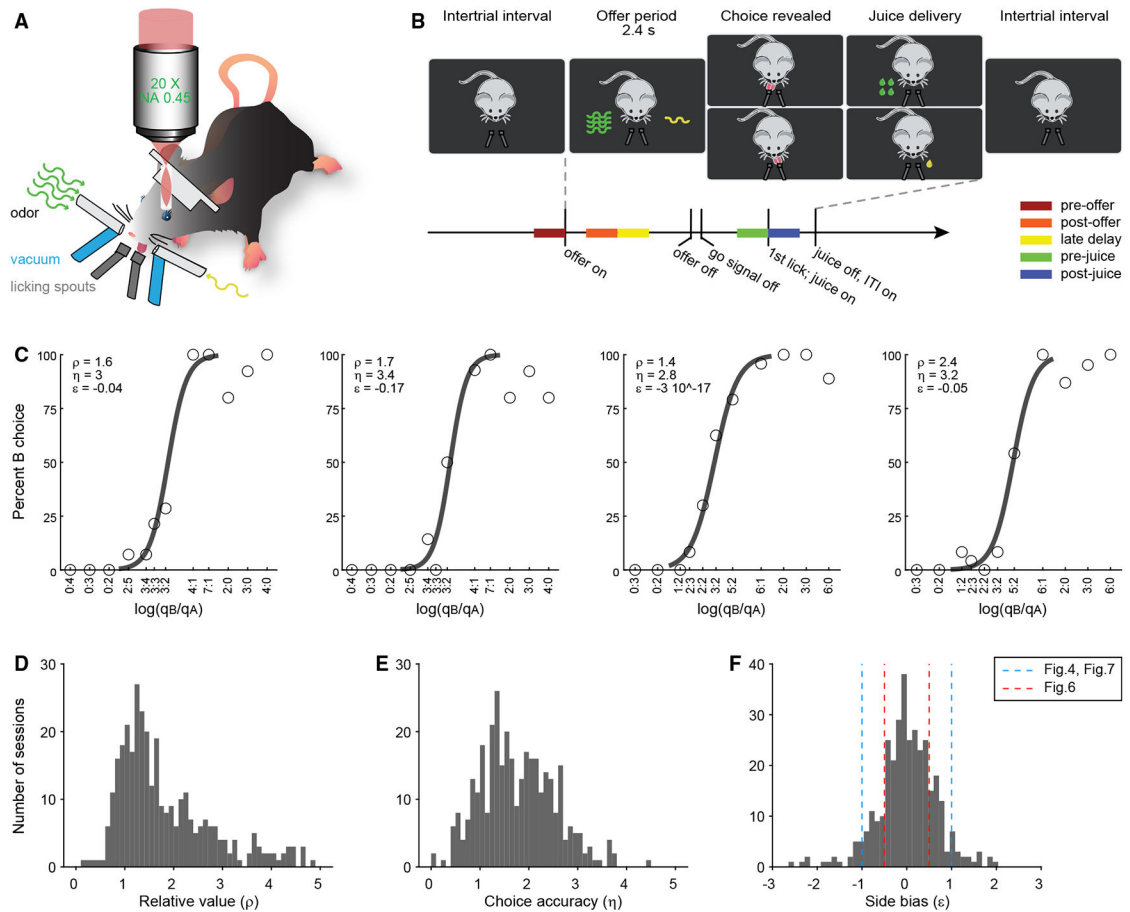


Figure 1. Choice task and behavioral performance

(A) Experimental apparatus. The animal was head fixed under the microscope objective. Two odors, each representing one of two offered juices, were presented simultaneously from the left and right. The mouse indicated its choice by licking one of two spouts.

(B) Trial structure. Following an inter-trial interval (ITI; 2.5–5 s), odors were delivered from the two odor ports. The offer period lasted 2.4 s and ended with an acoustic tone (0.2 s; go signal). The animal indicated its choice within 5 s. Immediately after the first lick following the go signal, the corresponding juice was delivered. During the subsequent ITI, a vacuum system removed all remaining odors from the field.

(C) Example sessions and choice patterns. In each plot, offer types are represented on the x axis. The y axis indicates the percentage of trials in which the animal chose juice B. Sigmoid curves were obtained from logistic regressions. Relative value (ρ), choice accuracy (η), and side bias (ϵ) are indicated for each session. (D–F) Population histograms for behavioral parameters. In (F), the dashed lines highlight $|\epsilon| = 0.5$ and $|\epsilon| = 1$ (see STAR Methods).

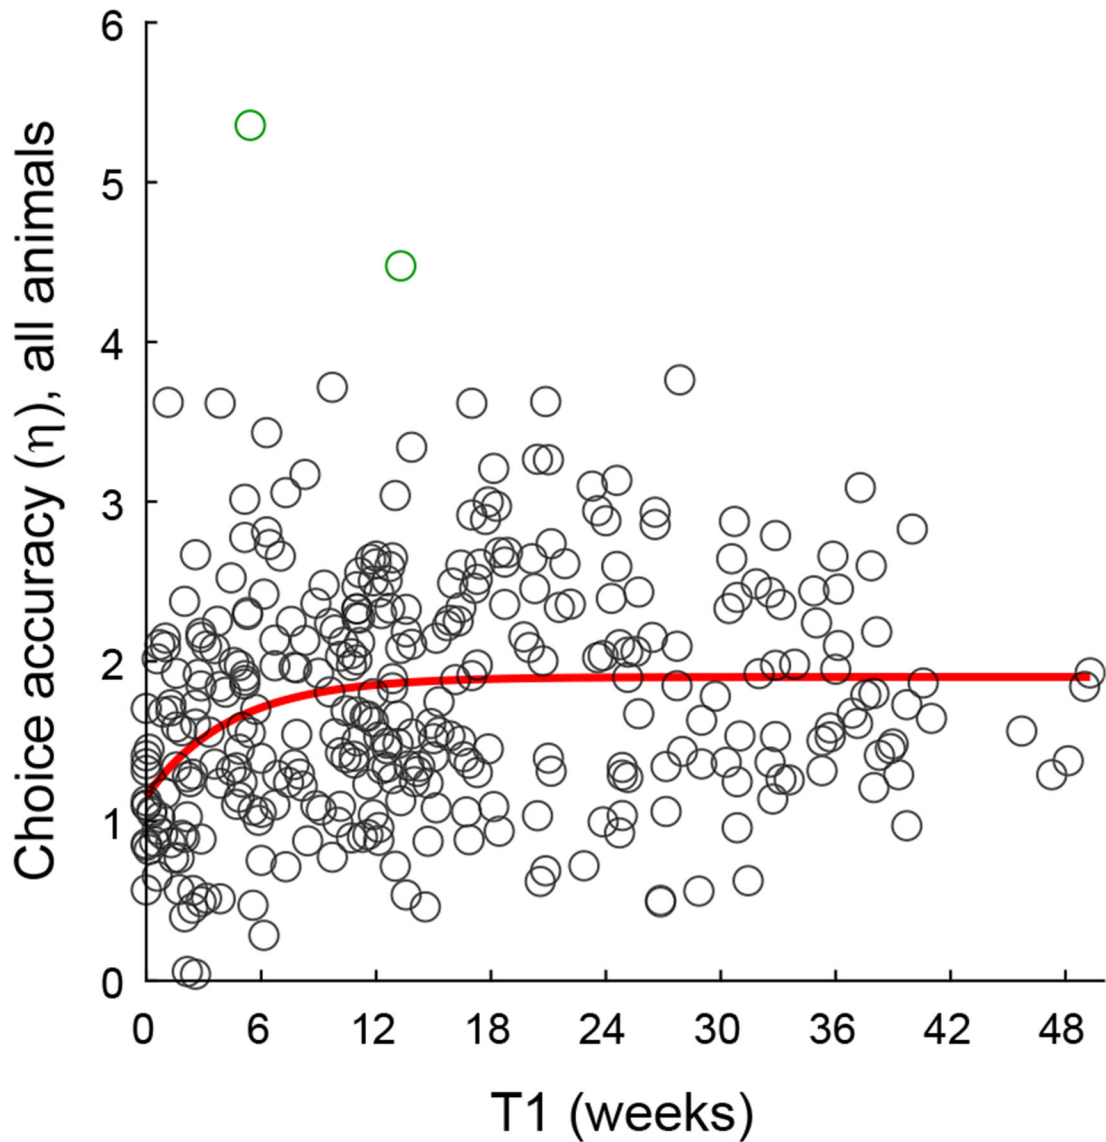


Figure 2. Longitudinal analysis of choice accuracy

We pooled 343 sessions from 13 mice (entire dataset). Each circle represents one session, and the red line was obtained from an exponential fit. Two outliers (green circles) were excluded from the fit. Referring to Equation 3, the fitted parameters (95% confidence interval) were $a = -0.75 (-1.03, -0.46)$, $b = 0.03(0.005, 0.06)$, and $c = 1.90(1.78, 2.02)$. Thus, the behavioral performance, quantified by choice accuracy η , improved significantly as mice became more expert in the task.




Figure 3. Longitudinal two-photon imaging of OFC

(A) Viral expression and lens location. Left: coronal slice (anterior-posterior = +2.65 mm) obtained from a mouse 2 months after a surgery in which we injected AAV-Syn-GCaMP6f in the OFC and implanted a 1 mm diameter GRIN lens. Cell nuclei were stained by DAPI (white), and cells expressing GCaMP6f are contained in the green injection site. The location of the GRIN lens is clearly visible. The borders of OFC are highlighted (white dotted line) by matching them with a coronal section obtained from the Allen Brain Atlas (2011; right). The focal plane (red dotted line) was calculated based on the lens working distance.

(B) Example FOV. Imaging segmentation revealed the presence of $n = 169$ cells. White contours show neuronal locations. Red contours show 12 example cells illustrated in (C). A, anterior; M, medial. Scale bar, 100 μm .

(C) $\Delta F/F$ traces (one trial) of 12 example neurons (numbered top to bottom). Vertical lines indicate different behavioral events. Red/blue dots indicate left/right licks.

(D1–D5) Example FOV recorded in five sessions over 83 days. Images remained remarkably stable throughout the recordings. Scale bar, 100 μm .

(E1 and E2) Images for days 1 and 11 obtained after cell segmentation and cell matching. Orange contours indicate matching cells ($n = 98$ cells). White contours indicate unshared cells ($n = 63$ cells for day 1; $n = 33$ cells for day 11).

(F) Longitudinal stability. We recorded 18 sessions from one single FOV over 100 days. The image illustrates the fraction of shared cells (number of shared cells/mean of cell numbers between FOVs) for every pair of sessions as a function of the time distance.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

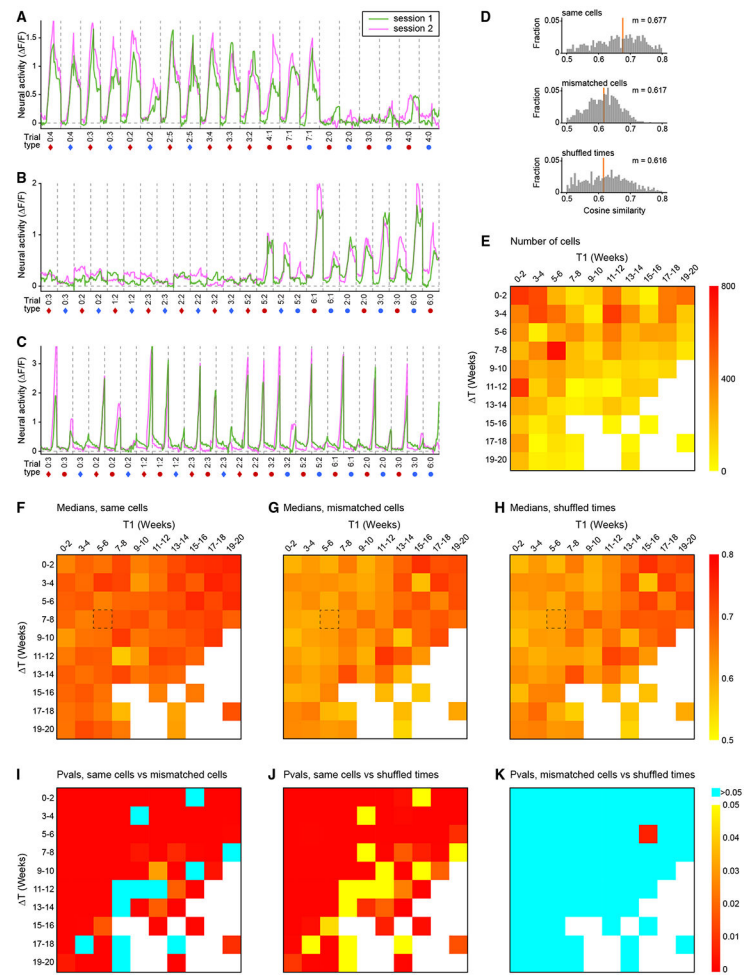


Figure 4. CS analysis

(A) Example cell. We examined the cell activity in two sessions (days 31 and 75). Single-trial traces were obtained by joining two time windows aligned with the offer onset and the time of first lick. In each session, we averaged traces over trials for each trial type, and we concatenated the mean traces obtained for different trial types. The resulting signal is referred to as the activity profile. We repeated this operation for both sessions, and we computed the cosine similarity (CS) between the two activity profiles. Here, activity profiles recorded in sessions 1 and 2 are shown in green and pink, respectively. Gray dashed lines separate trial types. Details about the trial types are indicated below the x axis: numbers indicate offered quantities ($q_B:q_A$); diamonds and circles indicate that the animal chose juice A and juice B, respectively; and red and blue indicate that the animal chose left and right, respectively. For this neuron, $CS = 0.96$.

(B) Example cell recorded on days 122 and 129 ($CS = 0.94$).

(C) Example cell recorded on days 1 and 63 ($CS = 0.90$).

(D) Distribution of CS for the three treatments (same cells, mismatched cells, shuffled times). Here, $T1 = \text{weeks } 5-6$, $\Delta T = \text{weeks } 7-8$, and $n = 747$ cells. For each distribution, the orange vertical line indicates the median. We measured $\text{median}(CS_{\text{same}}) = 0.677$, $\text{median}(CS_{\text{mismatched}}) = 0.617$, and $\text{median}(CS_{\text{shuffled}}) = 0.616$. The median measured for the CS_{same}

distribution is significantly higher than that measured for the $CS_{\text{mismatched}}$ distribution ($p = 2.26 \times 10^{-38}$, Kruskal-Wallis test) and that measured for the CS_{shuffled} distribution ($p = 1.37 \times 10^{-29}$, Kruskal-Wallis test).

(E) Number of cells matched across sessions as a function of T1 and ΔT . Session pairs were divided in bins of 2×2 weeks. Here, the number of cells available in each bin are illustrated with a heatmap (see color legend). Each entry in the table indicates the number of cells recorded with the corresponding T1 (x axis) and ΔT (y axis). Bins with ≤ 10 cells were removed from this figure.

(F–H) Measures of median(CS) as a function of T1 and ΔT same cells, mismatched cells, and shuffled times. In each image, median(CS) is illustrated with a heatmap. The black dashed rectangle highlights the bin illustrated in (D).

(I–K) Statistical comparison of different treatments. (I) illustrates the p values obtained comparing the distributions of CSs measured for “same cells” and “mismatched cells” treatments. (J) and (K) are for the other two comparisons. All p values (see color bar) are from a Kruskal-Wallis test.

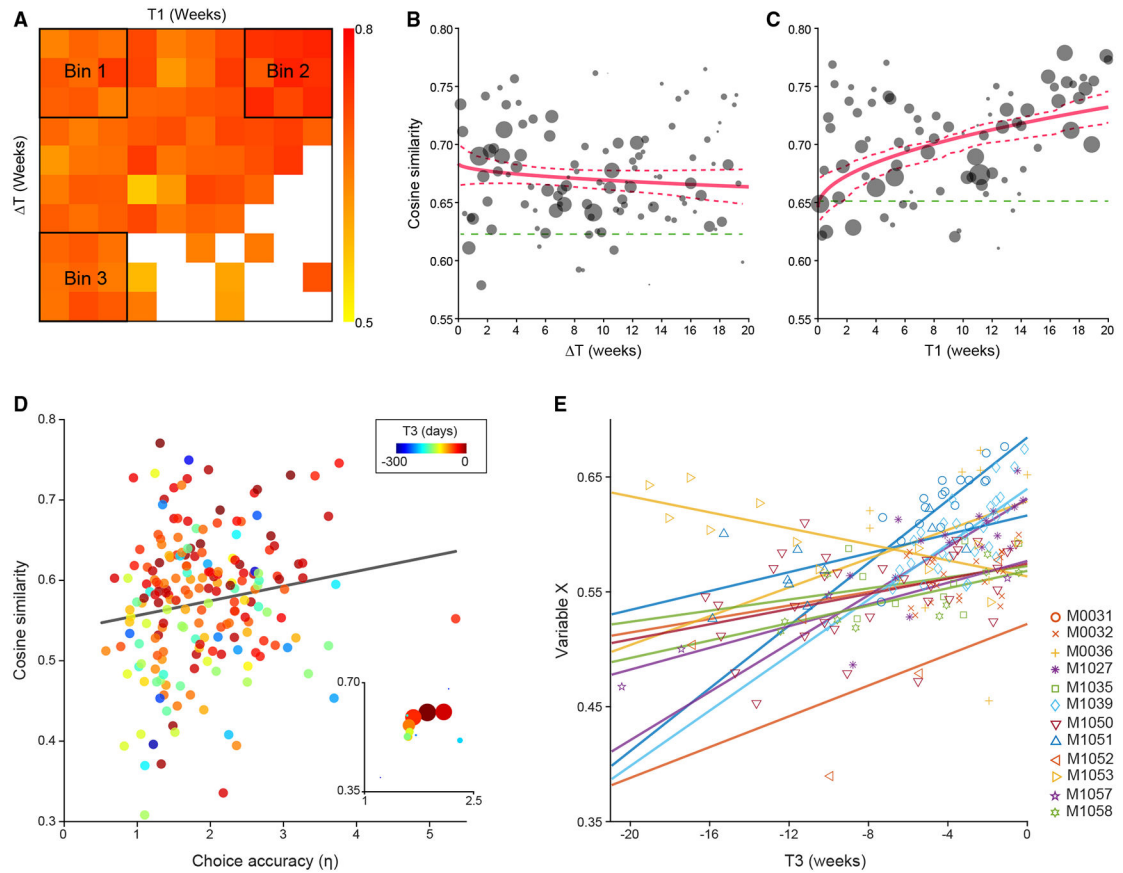


Figure 5. Longitudinal trends of cosine similarity

(A) Median(CS) as a function of T1 and ΔT (same as Figure 4F); 6×6 week Bins used for longitudinal comparisons are highlighted. Median(CS) was significantly higher in Bin 2 than in Bin 1 ($p = 1.55 \times 10^{-115}$) and significantly lower in Bin 3 than Bin 1 ($p = 0.05$, Kruskal-Wallis test).

(B) CS decreases with time passage (ΔT). To have comparable numbers of cells for different values of ΔT , we restricted this analysis to pairs of sessions with $T1 \leq 6$ weeks. For each day, we computed median(CS). Here, the radius of each circle is proportional to the number of cells. We fitted data with a drift model $CS(\Delta T) = a_0 + a_1 \Delta T^{1/2}$ (weighing data points by the number of cells) and obtained $a_0 = 0.6835$ and $a_1 = -1.697 \cdot 10^{-3} \text{ day}^{-1/2}$. Red dashed lines indicate the 95% confidence interval. For the same population, we also computed $b = \text{median}(CS_{\text{mismatched cells}}) = 0.622821$ (green dashed line). Thus, we estimated the time necessary for full reorganization, $\Delta T_{FR} = 1,278$ days.

(C) CS increases with task experience (T1). This analysis was restricted to pairs of sessions with $\Delta T \leq 6$ weeks. The red line was obtained from the fit $CS(T1) = a_0 + a_1 T1^{1/2}$. The green dashed line indicates the median(CS) obtained for mismatching cells.

(D) CS and choice accuracy (population). Each dot represents one session, and the line was derived from a linear regression ($p = 0.04$). The color of the dot represents T3 (see color bar). In the inset, dots represent 25-day bins, and colors indicate T3.

(E) CS, choice accuracy, and experience in the task (individual animals). Axes represent T3 (x axis) and variable X (y axis), and each data point represents one session. Different colors indicate different animals, and solid lines were derived from an ANCOVA (independent lines). For 11 of 12 animals, the two measures were positively correlated.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

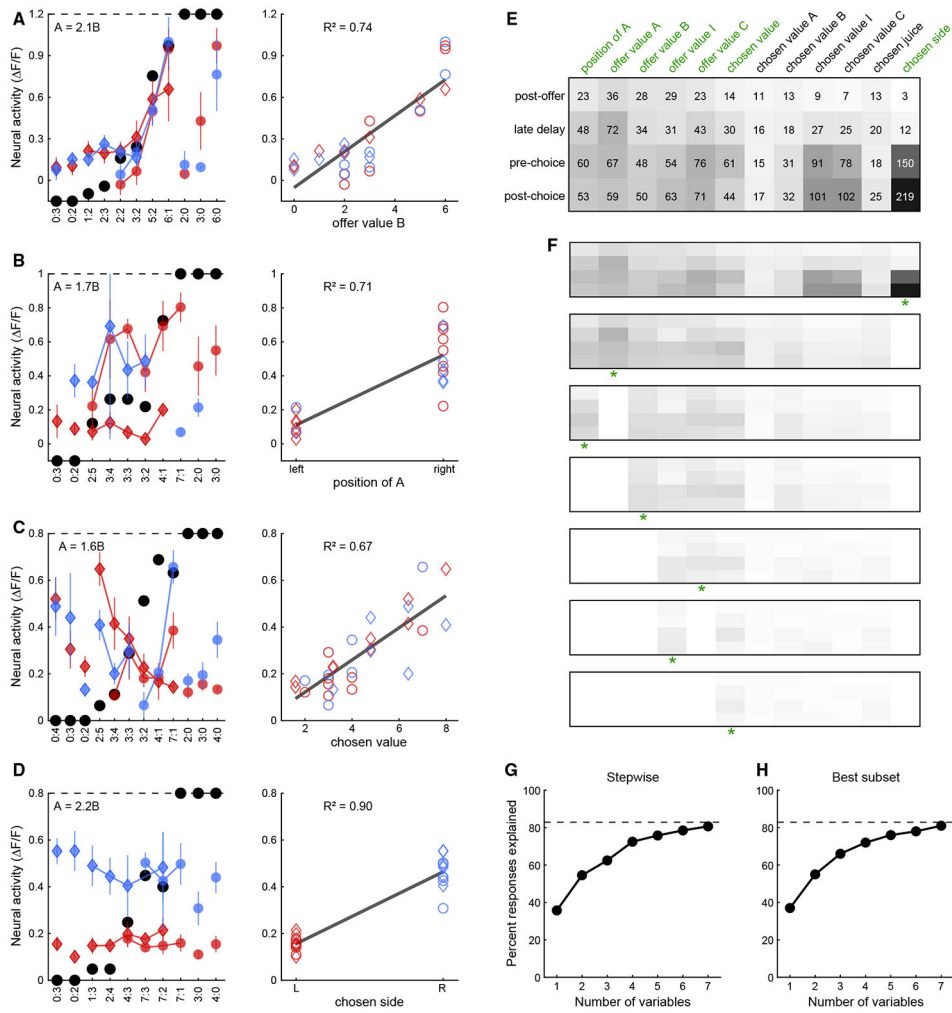


Figure 6. Decision variables encoded in OFC

(A) Example cell encoding the offer value B. Left: offer types are ordered by the quantity ratio q_B / q_A (x axis). Black dots represent the choice pattern. Color symbols represent the neuronal activity ($\Delta F / F$). Diamonds and circles are for trials in which the animal chose juice A and juice B, respectively. Red and blue are for trials in which the animal chose the offer on the left and right, respectively. Right: the same neuronal response is plotted against the binary variable offer value B. The cell activity increased as a function of q_B . The dark gray line was obtained from a linear regression.

(B) Example cell encoding the position of A. The activity of this cell was roughly binary—low and high when juice A was offered on the left and right, respectively.

(C and D) Example cells encoding the chosen value and the chosen side. Neuronal responses illustrated here are from (A) late delay, (B) late delay, (C) post-juice, and (D) pre-juice time windows. All conventions are as in (A).

(E) Population analysis. Each response passing the ANOVA criterion was regressed against each of the 12 variables. The images illustrates for each time window (row) the number of responses best explained by each variable (column). Here, numbers and shades of gray indicate cell counts.

(F) Variable selection and stepwise procedure. The top image is the same as in (E). In the first iteration, we identified the variable that explained the highest number of responses (chosen side). We removed from the pool all the responses explained by that variable (35.73% of the total). The residual pool is illustrated in the second image. The procedure was then repeated. At each iteration, we identified the variable that explained the largest number of responses in the residual pool (highlighted by a green asterisk). If the marginal explanatory power was $\geq 2\%$, then we retained the variable and removed from the pool all the responses explained by that variable. The procedure continued until all newly selected variables failed the 2% criterion.

(G) Percentage of explained responses and stepwise procedure. The x axis indicates the number of variables (iterations). On the y axis, 100% corresponds to the number of responses passing the ANOVA criterion; the dotted line corresponds to the number of responses collectively explained by the 12 variables included in the analysis.

(H) Percentage of explained response and best-subset procedure. The same format as in (G) is used. In (E), variables selected by both the stepwise and best-subset procedures are highlighted in green.

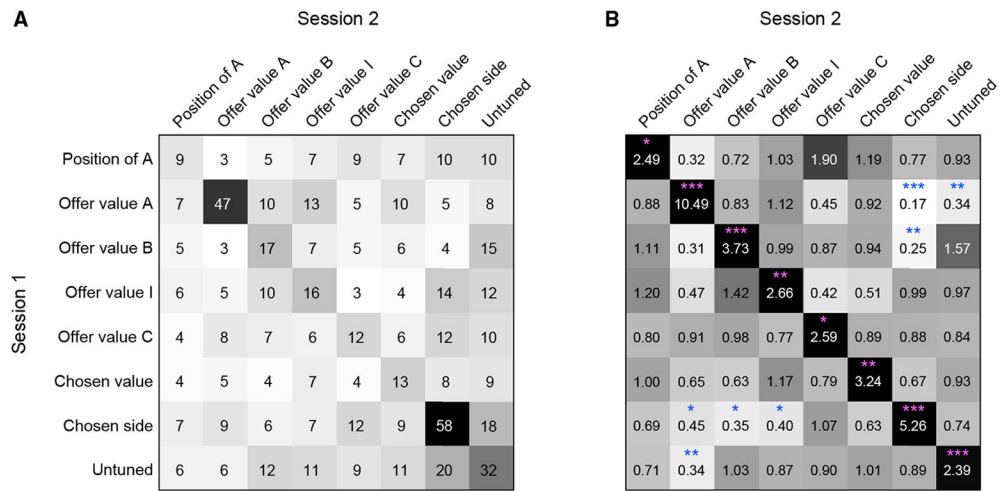


Figure 7. The encoding of decision variables is longitudinally stable

(A) Contingency table ($n = 649$ cells). Rows and columns represent the variables encoded in sessions 1 and 2, respectively; entries and gray shades are cell counts.

(B) Odds ratios (ORs). For each entry, Fisher's exact test assessed whether the departure from chance ($OR = 1$) was statistically significant. Pink and blue asterisks indicate that the cell count was significantly above and below chance, respectively ($*p < 0.05$, $**p < 0.01$, and $***p < 0.001$).

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Bacterial and virus strains		
pAAV.Syn.GCaMP6f.WPRE.SV40	Chen et al. ³⁷	Addgene AAV1; 100837-AAV1
pAAV.CAG.LSL.tdTomato	Oh et al. ³⁸	Addgene; 100048
Chemicals, peptides, and recombinant proteins		
Isoflurane	Covetrus	029405
Buprenorphine Sustained-Release (SR)	Wedgewood pharmacy	79926-058-17
Dexamethasone	Mylan Pharmaceuticals	67457-422-00
Lidocaine hydrochloride	Aspen veterinary resources	16479689
Sterilized saline (0.9%)	Pfizer	00409-4888-50
Carprofen	Rimadyl	Rimadyl injectable carprofen
DAPI mounting solution	Vector Laboratories	VECTASHIELD H-2000
C&B Metabond dental cement	Parkell	S380
Ketamine	Dechra	https://www.dechra-us.com/our-products/us/companion-animal/cat/prescription/ketamine-hydrochloride-injection
Xylazine	Dechra	https://www.dechra-us.com/our-products/us/equine/horse/prescription/rompunxylazine-injection-100-mg-ml
Paraformaldehyde (PFA)	Thermo Fisher scientific	28908
Octanal	Sigma	05608
Octanol	Sigma	W280011
Critical commercial assays		
Hamilton Neuros syringe	Hamilton	Hamilton: 65458-02
Inscopix ProView lens probe	Inscopix	1050-004605
Cryostats	Leica	Leica Biosystems, CM1950
LabVIEW	National Instruments	https://www.ni.com/en/support/downloads/software-products/download.labview.html#544096
Inverted microscopy	Leica	DMI6000B
Mai Tai DeepSee Ti:Sapphire laser	Spectra-Physics	https://www.spectra-physics.com/en/f/mai-tai-deepsee-ultrafast-laser
20 ×0.45 NA objective	Olympus	LCPLN20×IR objective
BX61W1 microscope	Olympus	BX61W1
Prairie Ultima Two-photon <i>in vivo</i> rig	Bruker	Ultima
Deposited data		
Experimental models: Organisms/strains		
Mouse: C57BL/6J	The Jackson Laboratory	JAX: 000664
Mouse: B6; 129S-Rasgrf2tm1(cre/foxA)Hze/J	The Jackson Laboratory	JAX: 022864
Mouse: B6; 129P2-Pvalbtm1(cre)Arbr/J	The Jackson Laboratory	JAX: 008069
Software and algorithms		
MATLAB	Mathworks	https://www.mathworks.com/products/matlab.html
CaImAn package	Giovannucci et al. ²⁷	https://github.com/flatironinstitute/CaImAn-MATLAB
ImageJ (Fiji)	ImageJ	https://imagej.net/software/fiji/
CellReg algorithm for matching cells across days	Sheintuch et al. ³⁹	https://github.com/zivlab/CellReg