

Subpopulations of neurons in IOFC encode previous and current rewards at time of choice

David L Hocker^{1*}, Carlos D Brody^{1,2,3,4}, Cristina Savin^{1,5},
Christine M Constantinople¹

¹Center for Neural Science, New York University, New York, United States; ²Princeton Neuroscience Institute, Princeton University, Princeton, United States; ³Department of Molecular Biology, Princeton University, Princeton, United States; ⁴Howard Hughes Medical Institute, Princeton University, Princeton, United States; ⁵Center for Data Science, New York University, New York, United States

Abstract Studies of neural dynamics in lateral orbitofrontal cortex (IOFC) have shown that subsets of neurons that encode distinct aspects of behavior, such as value, may project to common downstream targets. However, it is unclear whether reward history, which may subserve IOFC's well-documented role in learning, is represented by functional subpopulations in IOFC. Previously, we analyzed neural recordings from rats performing a value-based decision-making task, and we documented trial-by-trial learning that required IOFC (Constantinople et al., 2019). Here, we characterize functional subpopulations of IOFC neurons during behavior, including their encoding of task variables. We found five distinct clusters of IOFC neurons, either based on clustering of their trial-averaged peristimulus time histograms (PSTHs), or a feature space defined by their average conditional firing rates aligned to different task variables. We observed weak encoding of reward attributes, but stronger encoding of reward history, the animal's left or right choice, and reward receipt across all clusters. Only one cluster, however, encoded the animal's reward history at the time shortly preceding the choice, suggesting a possible role in integrating previous and current trial outcomes at the time of choice. This cluster also exhibits qualitatively similar responses to identified corticostriatal projection neurons in a recent study (Hirokawa et al., 2019), and suggests a possible role for subpopulations of IOFC neurons in mediating trial-by-trial learning.

*For correspondence:
dh148@nyu.edu

Competing interest: The authors declare that no competing interests exist.

Funding: See page 23

Received: 14 May 2021

Accepted: 20 October 2021

Published: 25 October 2021

Reviewing Editor: Emilio Salinas, Wake Forest School of Medicine, United States

© Copyright Hocker et al. This article is distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use and redistribution provided that the original author and source are credited.

Editor's evaluation

This study used advanced statistical methods to classify the responses of neurons in the lateral orbito-frontal cortex (IOFC) of rats performing economic choices. The authors report that IOFC neurons can be classified into a small number of distinct types based on their characteristic temporal dynamics and preferences for task variables (e.g., choice, success, reward history). The results will be of interest to neuroscientists studying prefrontal cortex, as they demonstrate order amidst what would appear as disorganized diversity, and suggest that some neuronal types play specific functional roles; for instance, in integrating previous trial outcomes into a current choice.

Introduction

Previous experience can profoundly influence subsequent behavior and choices. In trial-based tasks, the effects of previous choices and outcomes on subsequent ones are referred to as 'sequential effects,' and while they are advantageous in dynamic environments, they produce suboptimal biases

when outcomes on each trial are independent. The orbitofrontal cortex (OFC) has been implicated in updating behavior based on previous experience, particularly when task contingencies are partially observable (*Rushworth et al., 2011; Stalnaker et al., 2014; Wallis, 2007; Wilson et al., 2014*). However, it is unclear whether behavioral flexibility in OFC is mediated by dedicated subpopulations of neurons exhibiting distinct encoding and/or connectivity. We previously trained rats on a value-based decision-making task, in which they chose between explicitly cued, guaranteed and probabilistic rewards on each trial (*Constantinople et al., 2019b; Constantinople et al., 2019a*). Despite the fact that outcomes were independent on each trial, we observed several distinct sequential effects that contributed to behavioral variability. Optogenetic perturbations of the lateral orbitofrontal cortex (lOFC) eliminated one particular sequential effect, an increased willingness to take risks following risky wins, but spared other types of trial-by-trial learning, such as spatial 'win-stay/lose-switch' biases. We interpreted this data as evidence that (1) different sequential effects may be mediated by distinct neural circuits and (2) lOFC promotes learning of abstract biases that reflect the task structure (here, biases for the risky option), but not spatial ones (*Constantinople et al., 2019a*).

Electrophysiological recordings from lOFC during this task revealed encoding of reward history and reward outcomes on each trial at the population level, which could in principle support sequential effects (*Constantinople et al., 2019a*). Recent studies of rodent OFC have suggested that despite the apparent heterogeneity of neural responses in prefrontal cortex, neurons can be grouped into distinct clusters that exhibit similar task-related responses and, in some cases, project to a common downstream target (*Hirokawa et al., 2019; Namboodiri et al., 2019*). In light of these results, we hypothesized that reward history might be encoded by a distinct cluster of neurons in lOFC. This would suggest that the lOFC-dependent sequential effect we observed (an increased willingness to take risks following risky wins) may derive from a subpopulation of neurons encoding reward history, and may potentially project to a common downstream target.

Here, we analyzed an electrophysiological dataset from lOFC during a task in which independent and variable sensory cues conveyed dissociable reward attributes on each trial (reward probability and amount; *Constantinople et al., 2019b; Constantinople et al., 2019a*). Our previous analysis investigated single unit responses for sensitivity to task parameters, and here we found that clustering of lOFC neurons based on either their trial-averaged peristimulus time histograms (PSTHs), or a feature space defined by their average conditional firing rates aligned to different task variables (*Hirokawa et al., 2019*), both revealed five clusters of neurons with different response profiles. We exploited the temporal variability of task events to fit a generalized linear model (GLM) to lOFC firing rates, generating a rich description of the encoding of various task variables over time in individual neurons. All the clusters exhibited weak encoding of reward attributes, and stronger encoding of reward history, the animal's left or right choice, and reward receipt. Encoding of these variables was broadly distributed across clusters. Only one cluster, however, encoded the animal's reward history at the time shortly preceding the choice. This distinct encoding was observable by three independent metrics (coefficient of partial determination, mutual information, and discriminability or d') and two separate clustering methods. Intriguingly, this cluster exhibited qualitatively similar responses to identified corticostriatal cells in a previous study (*Hirokawa et al., 2019*), and also exhibited the strongest encoding of reward outcomes, suggesting that these neurons are well-situated to integrate previous and current reward experience. We hypothesize that this subpopulation of neurons, which were identifiable based on their temporal response profiles alone, may mediate sequential learning effects by integrating previous and current trial outcomes at the time of choice.

Results

Rats' behavior on this task has been previously described (*Constantinople et al., 2019b; Constantinople et al., 2019a*). Briefly, rats initiated a trial by poking their nose in a central nose port (*Figure 1A–C*). They were then presented with a series of pseudo-randomly timed light flashes, the number of which conveyed information about reward probability on each trial. Simultaneously, they were presented with randomly timed auditory clicks, the rate of which conveyed the volume of water reward (6–48 μ l) baited on each side. After a cue period of $\sim 2.5 - 3$ s, rats reported their choice by poking in one of the two side ports. Reward volume and probability of reward (p) were randomly chosen on each trial. On each trial, the left or right port was randomly designated as risky ($p < 1$) or safe ($p = 1$). Well-trained rats tended to choose the side offering the greater subjective value, and trial

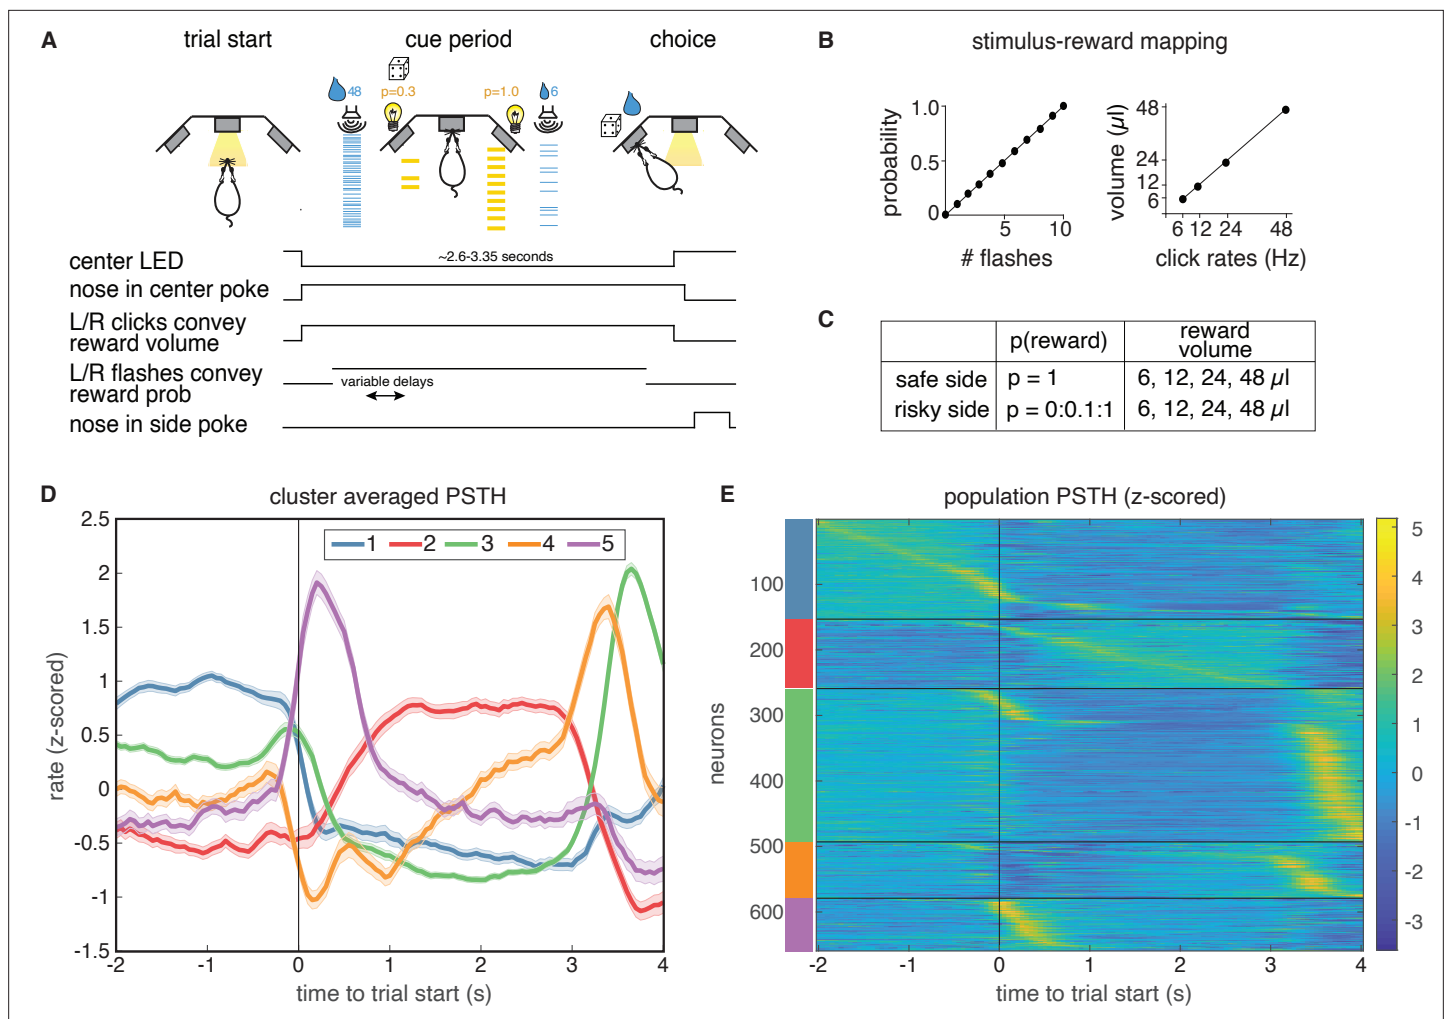


Figure 1. Task description and PSTH clustering. Single units in IOFC belong to clusters with distinct temporal response profiles. **(A)** Behavioral task: rats chose between a guaranteed and probabilistic water reward on each trial, that could each be one of four water volumes. Rats were well-trained and tended to choose the option with the greater subjective value. Reward probability and volume were cued by visual flashes and auditory click rates, respectively. **(B)** Mapping between the visual flashes/auditory clicks and reward attributes. **(C)** Range of reward attributes. **(D)** Cluster-specific, mean responses of the trial-averaged, z-scored PSTHs. Error bars denote s.e.m. **(E)** Z-scored PSTHs for neurons in each cluster, sorted by time to peak within each cluster. Colored bars on the left indicate the cluster identities from panel D. Panels A-C reproduced and modified from *Constantinople et al., 2019a*.

The online version of this article includes the following figure supplement(s) for figure 1:

Figure supplement 1. Clustering aligned to different task events.

history effects contributed to behavioral variability (*Constantinople et al., 2019b; Constantinople et al., 2019a*). We analyzed 659 well-isolated single-units in the IOFC with minimum firing rates of 1 Hz, obtained from tetrode recordings (*Constantinople et al., 2019a*).

Single units in IOFC belong to clusters with distinct temporal response profiles

To characterize the temporal dynamics of the neural responses during this task, we performed k-means clustering on trial-averaged PSTHs ('PSTH clustering'), aligned to trial initiation. We used the gap statistic (*Tibshirani et al., 2001*), which quantifies the improvement in cluster quality with additional clusters relative to a null distribution (Materials and methods), to choose a principled number of clusters, and identified five distinct clusters of responses (*Figure 1D*). Each cluster has a stereotyped period of elevated activity during the task: at trial start, during the cue period, and at or near reward delivery, with the largest cluster (cluster 3) having activity just before the rat entered the reward port

(**Figure 1—figure supplement 1**). Cluster two is the only cluster that exhibited persistent activity during the cue period. These data indicate that despite the well-documented heterogeneity of neural responses in prefrontal cortex, responses in IOFC belonged to one of a relatively small, distinct set of temporal response profiles aligned to different task events.

A recent study in rat IOFC similarly found discrete clusters of neural responses by clustering the conditional firing rates of neurons for different task-related variables ('conditional clustering') (**Hirokawa et al., 2019**). We sought to compare results from clustering on PSTHs and conditional firing rates. Therefore, we generated conditional response profiles for each neuron (**Figure 2B**; Materials and methods), and performed clustering on these conditional responses via the same procedure. This also revealed five distinct clusters that exhibited qualitatively similar temporal response profiles as the clusters that were based on the average PSTHs (**Figure 2C–E, Figure 2—figure supplement 3**). Moreover, clusters 2 and 3 were comprised of highly similar groups of neurons across both procedures, where similarity was quantified by the consistency of being labeled in a given cluster across clustering approaches (**Figure 2E**, 66 % overlap of cluster 2 neurons and 70 % of cluster three neurons). This indicates that neurons in these two clusters are robustly identified by either their temporal response profiles or their encoding of task variables (**Figure 2E**).

We evaluated the quality of clustering using two additional methods. First, we visualized clustering using t-SNE projections of data into two dimensions, and further quantified cluster separation using the Mahalanobis distance among clusters, which in this context can be interpreted as an average z-score for data in a cluster from a given cluster mean (**Figure 2—figure supplement 2**). We found that IOFC responses obtained by both procedures were qualitatively clustered in feature space, and demonstrated Mahalanobis distances with smaller intracluster distances than across-cluster distances. Second, we used the PAIRS statistic to more directly test whether IOFC responses belong to distinct clusters (**Raposo et al., 2014**). PAIRS examines the angle between neural responses in feature space, and obtains an average angle for the data points closest to it. Intuitively, a small average angle among neighboring points implies that the data is tightly packed together into clusters, whereas larger angles suggest broadly distributed, unclustered data. We found that our neural responses were more tightly packed together than would be observed for an unclustered, Gaussian reference distribution (**Figure 2—figure supplement 4**).

Finally, we also compared the estimated number of clusters in the data to those obtained using other clustering metrics, in particular the silhouette score and adjusted Rand Index (**Figure 2—figure supplement 5**). The silhouette score quantifies cluster quality with two components: it provides high scores when points within a cluster are close in feature space, and penalizes small inter-cluster distances that arise when clusters are nearby in feature space. The adjusted Rand Index (ARI) does not assess cluster quality through distances among datapoints in feature space, but instead is a measure of reproducibility of a clustering result. Specifically, it quantifies the probability that two labels from separate clustering results would arise by chance. We generated labels for ARI by subsampling our data (90 % sampling without replacement) and calculated ARI on the subsampled dataset and the full dataset. ARI results demonstrated a broad range of possible numbers of clusters with reproducible cluster labels (**Figure 2—figure supplement 5A**). In contrast, the silhouette score gave inconclusive results on our IOFC data (**Figure 2—figure supplement 5B**). The dimensionality and covariance of our IOFC responses lies in a regime in which clusters are tightly packed in feature space, a regime in which silhouette score is known to underestimate the number of clusters in other contexts (**Garcia-Dias et al., 2020; Garcia-Dias et al., 2018**). To assess this more directly, we compared the results from different clustering metrics on artificial data with known statistics (deriving from five clusters). We found that in the data-like regime, the silhouette score substantially underestimated the optimal number of clusters (two clusters, **Figure 2—figure supplement 6**, see Discussion and Materials and methods). The gap statistic underestimated cluster number on synthetic data, but to a lesser degree. Taken together, these analyses suggest that the gap statistic is a conservative but principled metric for determining the number of clusters in the statistical regime of the IOFC data.

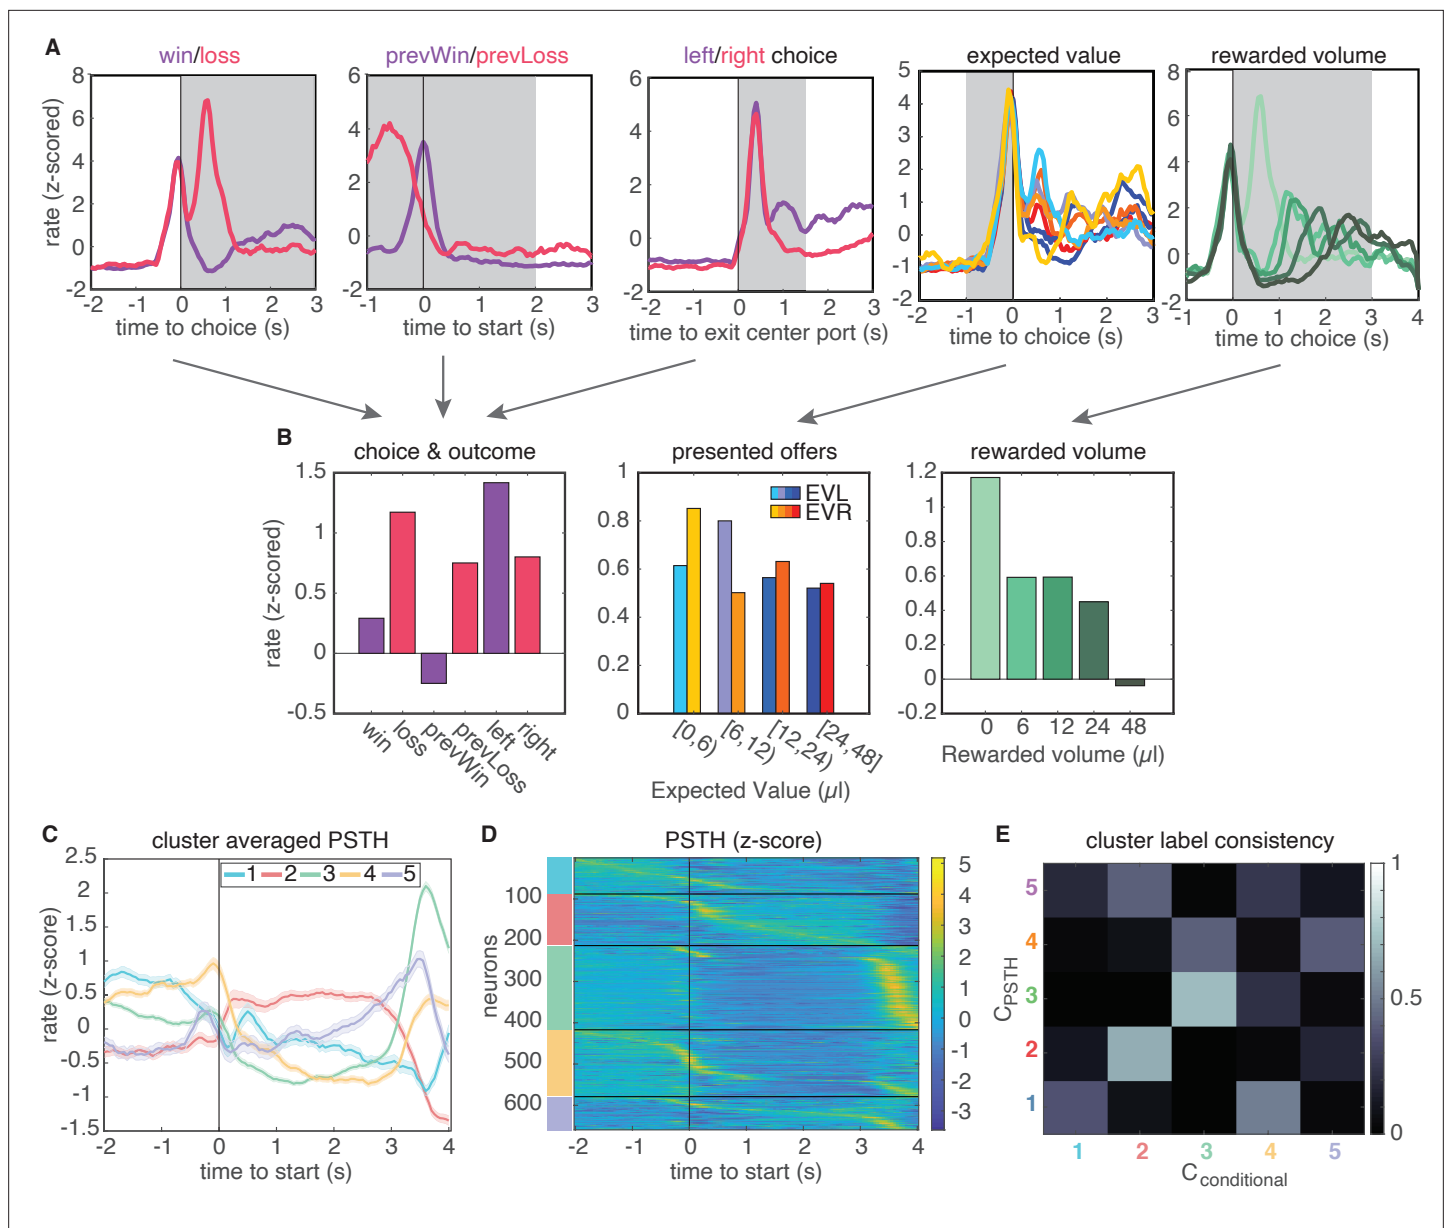


Figure 2. Conditional feature space clustering. **(A)** Feature-conditioned PSTHs of an example neuron. The average z-scored firing rate in each gray time window comprises an element in the feature space. **(B)** Feature space for conditional clustering for the sample neuron from A. The average firing rate for each condition is concatenated to yield a 19-dimensional feature space. Here, features are grouped into three qualitative types (for presentation purposes only): responses for choice and trial outcome (left); responses to expected value of left and right offers (middle); and responses to received reward volume (right). **(C)** Cluster-averaged PSTHs, aligned to trial initiation, using this conditional feature space. Error bars denote s.e.m. **(D)** Z-scored PSTHs for all neurons, sorted by time to peak within each cluster. **(E)** Consistency of cluster labeling, calculated as the conditional probability $P(C_{\text{conditional}}|C_{\text{PSTH}})$ of belonging in any 'conditional' cluster ($C_{\text{conditional}}$), given that a unit belongs to a certain PSTH-defined cluster (C_{PSTH}).

The online version of this article includes the following figure supplement(s) for figure 2:

Figure supplement 1. Gap statistic.

Figure supplement 2. Similarity across clustering methods.

Figure supplement 3. Clustering aligned to different task events.

Figure supplement 4. PAIRS analysis.

Figure supplement 5. Silhouette score and adjusted Rand index evaluation of cluster number.

Figure supplement 6. Synthetic data study of silhouette score performance.

Neural responses in IOFC are well-captured by a generalized linear model that includes attributes of offered rewards, choice, and reward history

We next sought to characterize the encoding properties of neurons in each cluster. To that end, we modeled each neuron's trial-by-trial spiking activity with a generalized linear model with Poisson noise (**Figure 3**). This approach has been previously used in sensory areas to model the effects of external stimuli upon neural firing (*Pillow et al., 2005; Pillow et al., 2008*), and more recently has been extended to higher order areas of cortex during decision-making tasks (*Park et al., 2014*). In this model, the time-dependent firing rate (λ_t) of a neuron is the exponentiated sum of a set of linear terms.

$$\lambda_t = \exp \left[\sum_{s=1}^S \left(X_{t-\tau:t}^{(s)} * k_s(\tau) \right) \right], \quad (1)$$

where the variables $X_{t-\tau:t}^{(s)}$ are the stimuli and behavioral events against which neural activity is regressed; $k_s(\tau)$ denote the response kernels that are convolved with the task variables to model time-dependent responses to each variable (**Figure 3C**). The probability that a given number of spikes would occur in a discrete time bin of $\Delta t = 50\text{ms}$ is given by a homogeneous Poisson process with mean $\lambda_t \Delta t$.

The task variables and example kernels for our model are shown for a sample neuron in **Figure 3C**. Variables were binarized such that 1 (0) denoted the occurrence (absence) of an event. The task variables included reward (win or loss on the current or previous trial), choice (left or right), and the timing of cues indicating reward attributes for the left and right offers (reward volume conveyed by auditory clicks and reward probability conveyed by visual flashes). We also included the average reward rate prior to starting each trial and the location of the trial within the session (not shown). Given the large number of model parameters, we used a smooth, log-raised cosine temporal basis and L2 regularization to prevent overfitting (Materials and methods). Based on model comparison, we found that including a spike history term as in other GLM approaches did not improve our model, presumably due to the fact that we are modeling longer timescale responses.

The chosen model parameters were selected by model comparison against several alternative models using cross-validation. Model comparison favored a simpler binary win/loss representation of rewarded outcomes over richer representations of reward volume on the current or previous trial (**Figure 3—figure supplement 1**). The model reproduced the trial-averaged PSTHs of individual neurons (**Figure 3E**). To quantify model performance, we calculated the proportion of variance explained (R^2) for held-out testing data (**Figure 3F**). The model captured a high percentage of variance for most of the neurons in our dataset. A small fraction of neurons exhibited negative R^2 values (69 units, **Figure 3—figure supplement 2**), indicating that the model produced a worse fit of the test data than the data average. Our liberal inclusion criteria did not require neurons to exhibit task-modulation of their firing rates, so these neurons were likely not task-modulated, and were excluded from subsequent analyses.

Clustering reveals distributed encoding of most task variables across subpopulations of IOFC neurons

We next sought to characterize the extent to which neurons that belonged to different clusters might be 'functionally distinct,' and encode different task-related variables (*Hirokawa et al., 2019; Namboodiri et al., 2019*). To address this question, we computed two complementary but independent metrics, both based on the GLM fit: the coefficient of partial determination and mutual information. The coefficient of partial determination (CPD) quantifies the contribution of a single covariate (e.g. left/right choice) to a neuron's response by comparing the goodness-of-fit of the encoding model with and without that covariate. In other words, the CPD quantifies the amount of variance that is explained by kernels representing different aspects of the model. We computed the CPD in a sliding time window throughout the trial, and compared CPD values to a shuffle test, in which trial indices were randomly shuffled relative to the design matrix before generating model predictions. CPD values that were within the 95 % confidence intervals of the shuffled distribution were set to 0 before averaging CPD values over neurons in a cluster. The average CPD values for different task variables is shown for clusters based on each clustering procedure in **Figure 4A and C**. Note that CPD plots are aligned to different events, depending on the covariate.

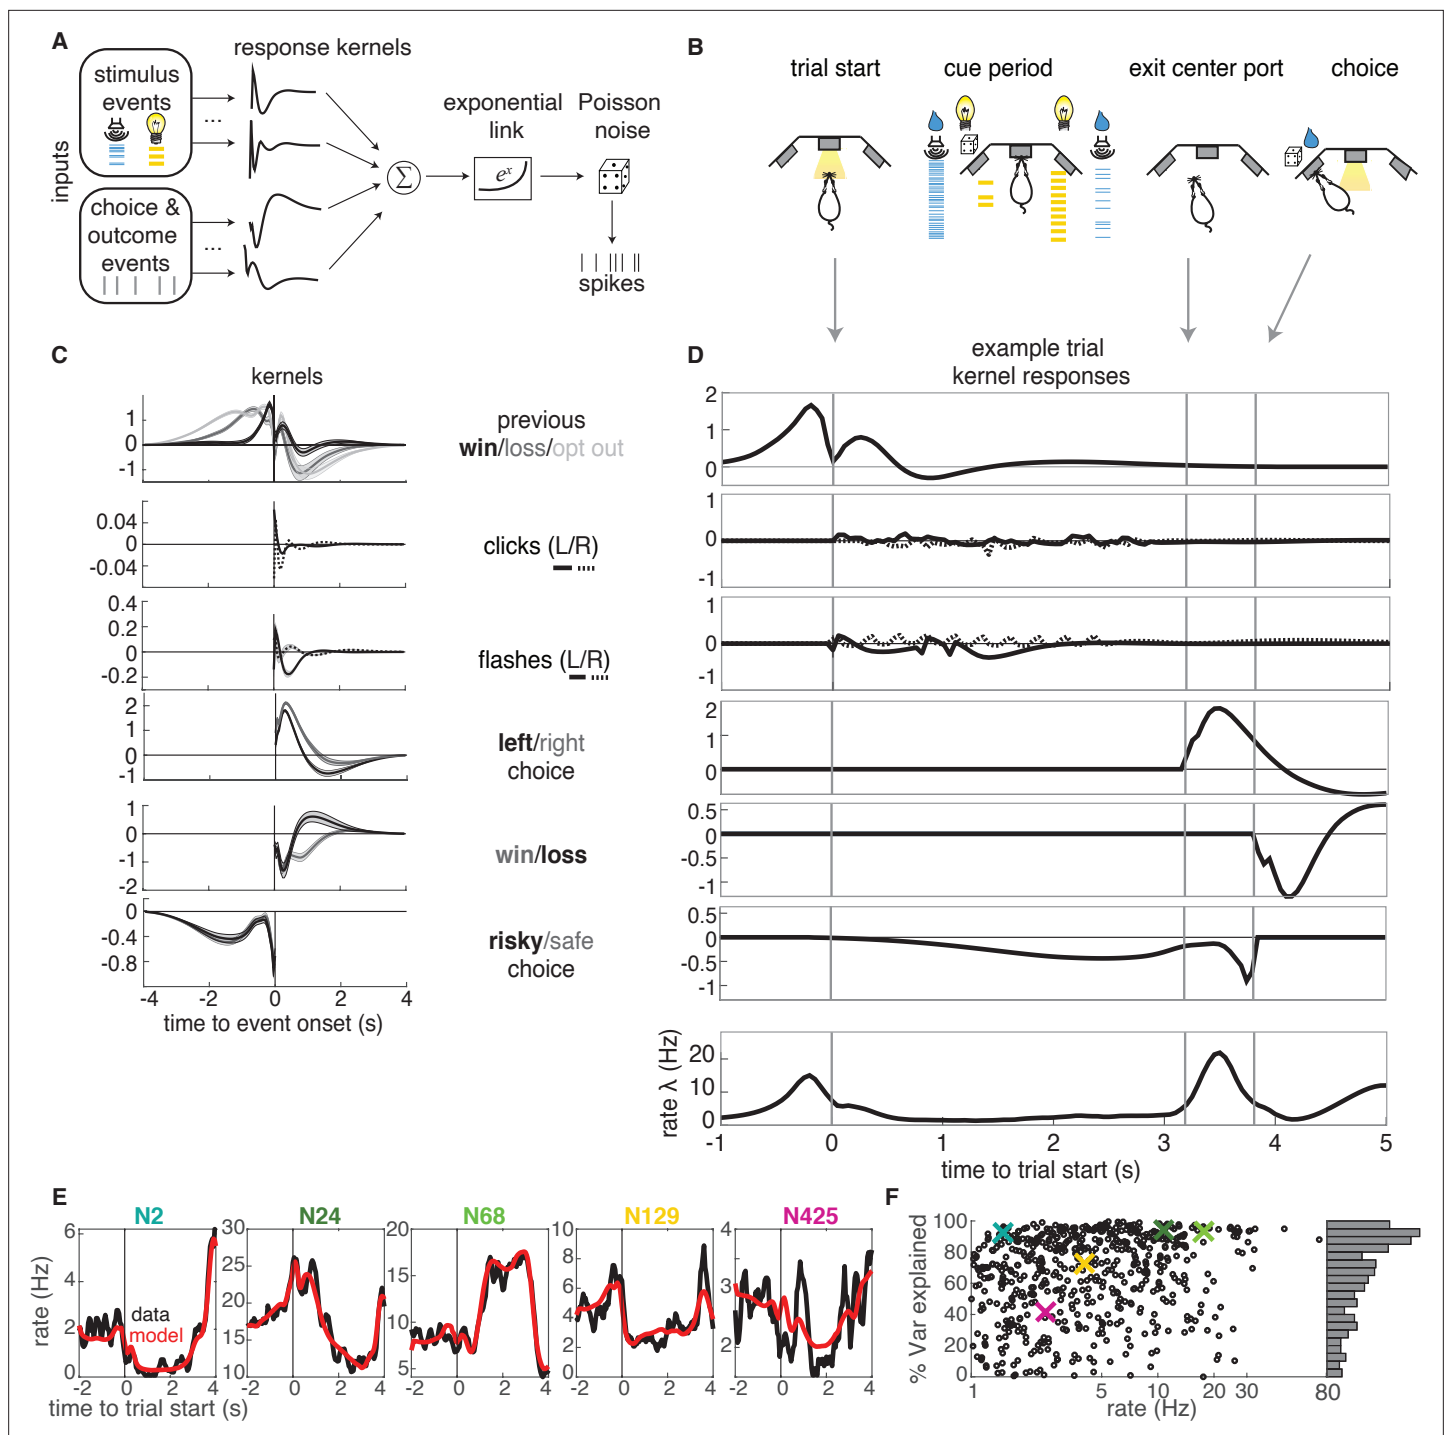


Figure 3. GLM analysis. **(A)** Model schematic: the timing of external stimuli conveying reward attributes, as well as timing of choices and outcomes on each trial are model inputs. Nonlinear response kernels convolved with each input generate time dependent responses that are summed and exponentiated to give a mean firing rate, λ_t , in each time bin. Spikes are generated from a Poisson process with mean firing rate λ_t . **(B)** Task schematic illustrating the key choice and outcome inputs to the model. Error bars denote estimated kernel standard deviation (Materials and Methods). **(C)** Kernel fits for a sample neuron. Kernels are grouped by the aspects of the task that they model. The kernels in bold in panel C are the kernels that are active in this trial. The resulting model-predicted firing rate is shown in the bottom row. **(E)** Representative PSTHs to held-out testing data from five different neurons (black) and model prediction (red). **(F)** Variance explained for each neuron, with sample neurons from E denoted by correspondingly colored crosses. The distribution of R^2 values is presented along edge of the panel.

The online version of this article includes the following figure supplement(s) for figure 3:

Figure 3 continued on next page

Figure 3 continued

Figure supplement 1. Model comparison.

Figure supplement 2. Variance explained.

We also computed the mutual information (MI) between neural spike trains and different model covariates (Materials and methods). Our approach relates the statistics of trial-level events to firing rates, which allows us to assess the information content for each stimulus throughout the entire time course of a trial. Such an approach does not easily generalize to the statistics of within-trial events such as the information represented in stimulus clicks and flashes, so we restricted the MI analysis to the other covariates: reward history, choice, and reward outcome. The average MI for these variables is shown for clusters based on the trial-averaged PSTHs in **Figure 4B**.

In general, regardless of the metric (CPD or MI) or clustering procedure (PSTH clustering or conditional clustering), task variables appear to be broadly encoded across neural subpopulations, with similar temporal dynamics and average CPD/MI values across clusters. All clusters encoded cues conveying reward volume and probability during the cue period (**Figure 4C**), although it is worth noting that the magnitude of CPD for clicks and flashes was an order of magnitude lower than for the other task variables. Therefore, encoding of cues representing reward attributes was substantially weaker than encoding of reward history, choice, and outcome. It may be surprising that neurons with strikingly different PSTHs appear to encode task variables with similar time courses. However, PSTHs marginalize out all conditional information, so a PSTH carries no information about encoding of task or behavioral variables, per se. Additionally, to assess whether different clusters may have a different composition across neuron types, we classified single units based on their waveforms into regular spiking units with broad action potential (AP) and after-hyperpolarization (AHP) widths, or fast-spiking units with narrow AP and AHP widths. We found the distribution of these two unit types to be similar across clusters (**Figure 4—figure supplement 1**).

Reward history was most strongly encoded at the time of trial initiation and decayed over the course of the trial, consistent with previous analyses (**Constantinople et al., 2019a**). Neurons belonging to cluster 2, which strongly overlap in both clustering procedures (**Figure 2E**), seem to exhibit slightly more pronounced encoding of reward history during the trial, compared to the other clusters, although these neurons were also persistently active during this time period. Encoding of choice and reward outcomes were phasic, peaking as the animal made his choice and received outcome feedback, respectively, and were broadly distributed across clusters (**Figure 4A–B**).

Reward history re-emerges before choice in a distinct subset of IOFC neurons

The CPD and MI analysis revealed one aspect of encoding that is unique to cluster 3. Both clustering methods identified largely overlapping populations of neurons in this cluster (**Figure 2E**), indicating that these neurons exhibited similar temporal response profiles as well as encoding. Like neurons in the other clusters, neurons in cluster 3 encoded reward history at trial initiation, and that encoding decreased through the cue period. However, for neurons in this cluster, encoding of reward history re-emerged at the time preceding the animal's choice on the current trial (**Figure 5A–B**, green lines). This 'bump' of reward history encoding late in the trial was unique to cluster 3 regardless of the clustering method, and was observable in both CPD and MI (**Figure 4A–B**). This result was further corroborated by computing the average discriminability index (d') for reward history, which is a model-agnostic metric that quantifies the difference in mean firing rate in each condition, accounting for response variance over trials. Cluster 3 was unique from the other clusters for having a subset of neurons with high d' values for reward history at the time preceding the animal's choice (**Figure 5C**). We additionally found that encoding of previous wins and losses during this time period only extended to the previous trial, and did not encode the outcome of additional past trials (**Figure 5—figure supplement 1**). Clustering of PSTHs aligned to later events in the trial showed more detailed structure of this late-in-trial encoding of reward history, with some neurons encoding the previous trial just before the choice, and others at the time of choice (**Figure 5—figure supplement 2C**). We note that clustering on PSTHs aligned to later events produced similar results, in terms of the number of clusters and the within-cluster responses (**Figure 5—figure supplement 2B**).

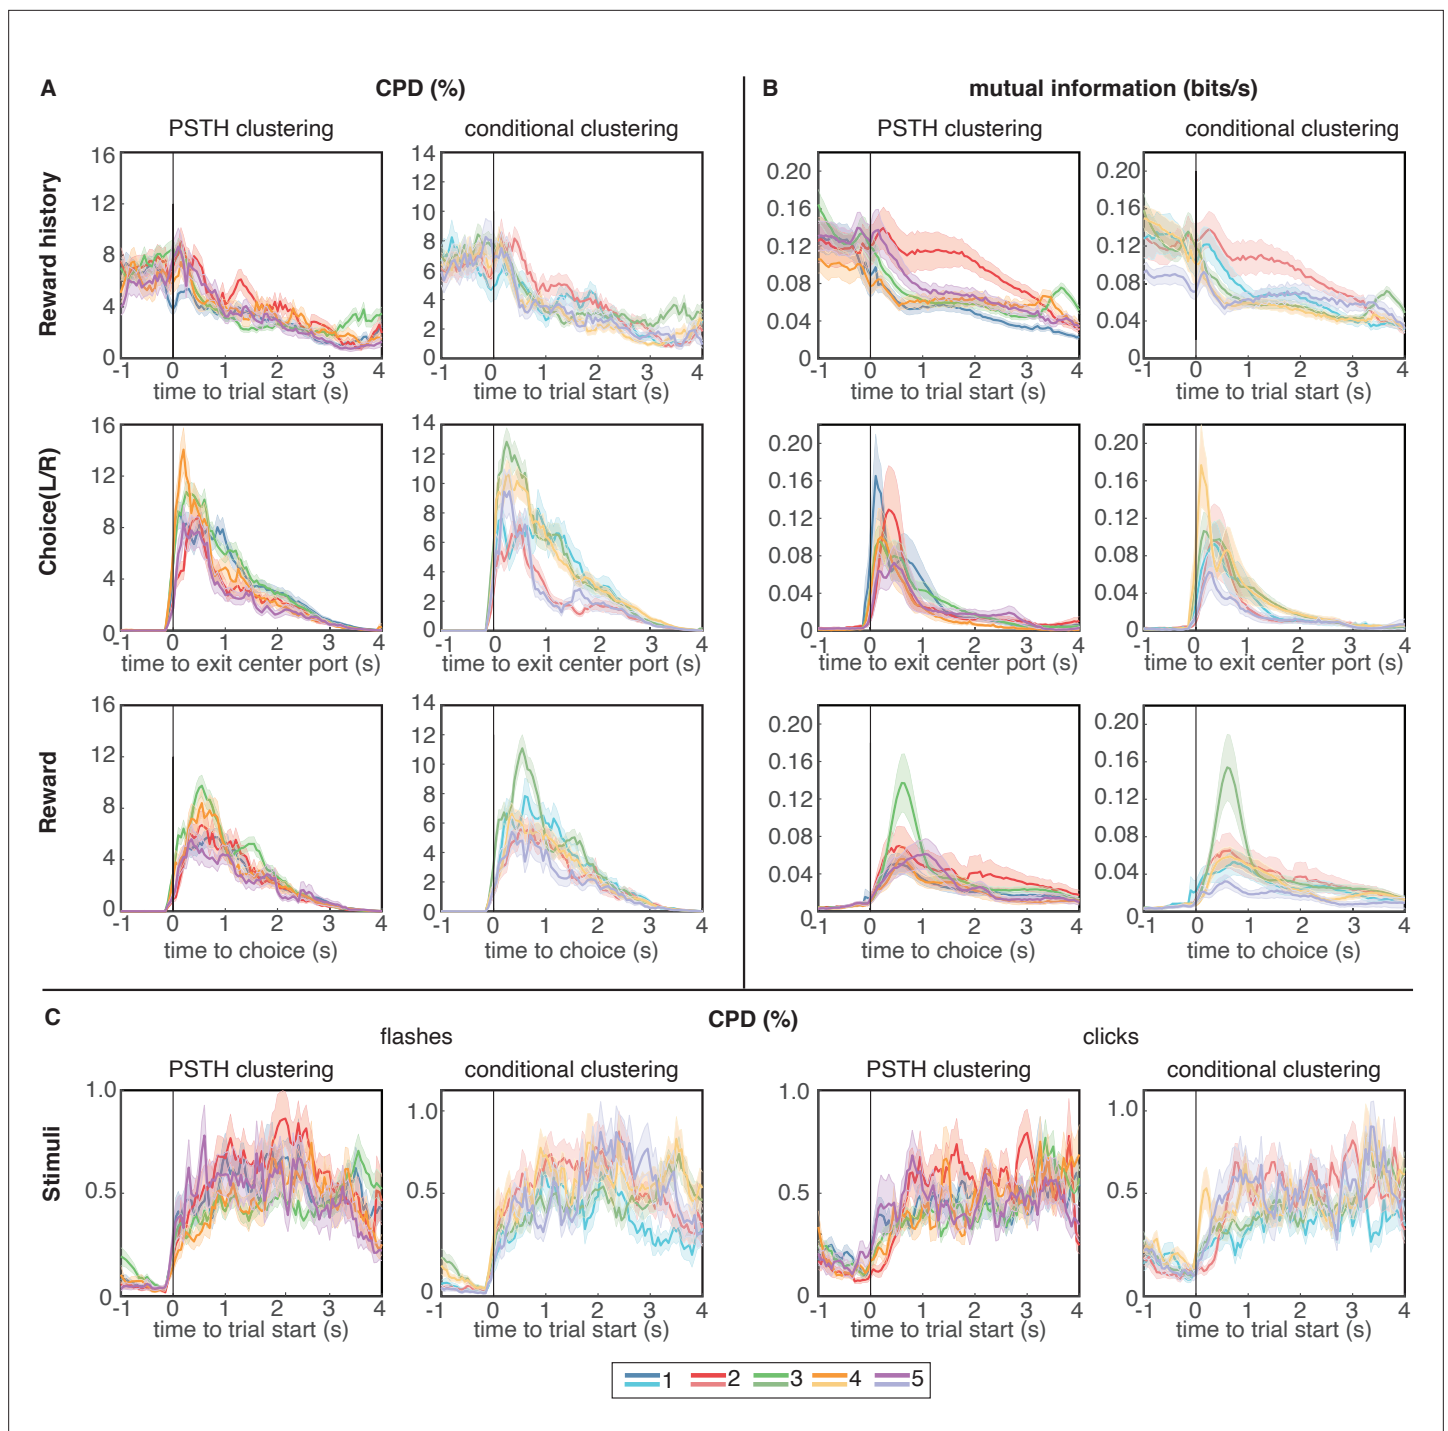


Figure 4. Coefficient of partial determination (CPD) and mutual information (MI) reveal broad encoding of all task parameters in each cluster. **(A)** CPD for choice and reward outcome encoding: Reward history panels quantify encoding of outcome on the previous trial (previous win/loss/opt out), choice panels convey encoding of choices on the current trial (left/right), and reward panels quantify encoding of outcomes on the current trial (win/loss). **(B)** Mutual information for same task parameters. **(C)** CPD of flashes (left panels) and clicks (right panels) that encode reward probability and reward volume, respectively. CPD and MI values are averaged over neurons in each cluster; error bars show s.e.m. Results for PSTH clustering are in the left columns, and results for conditional clustering are in the right columns.

The online version of this article includes the following figure supplement(s) for figure 4:

Figure supplement 1. Waveform analysis.

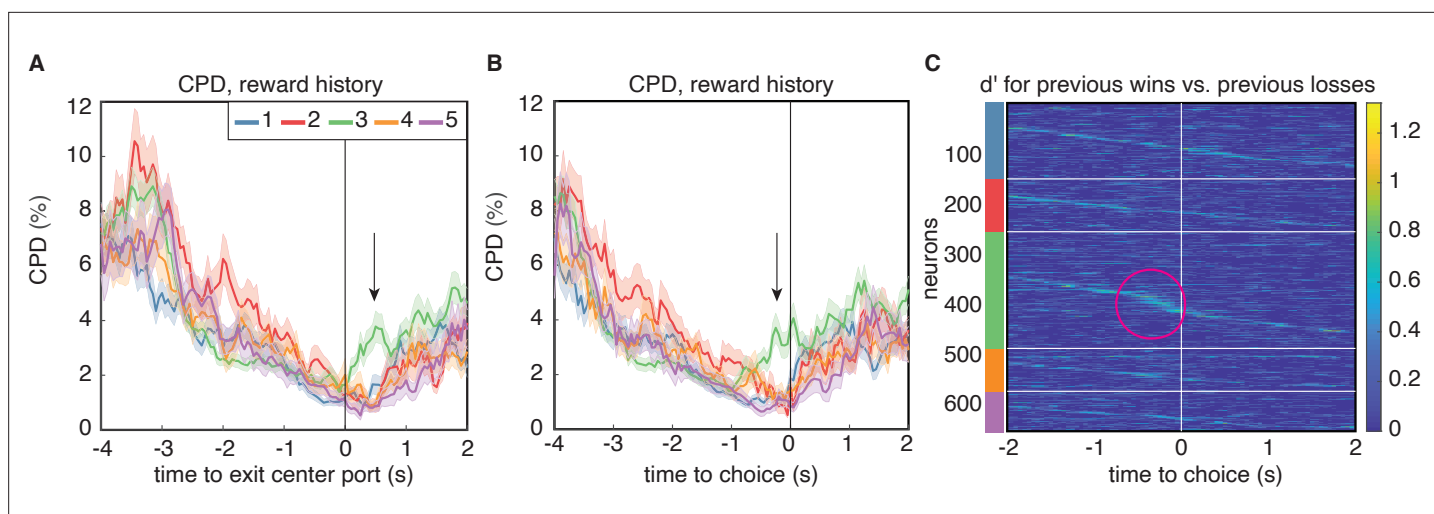


Figure 5. Encoding of reward history emerges late in the trial for cluster 3. (A) CPD for reward history, aligned to leaving the center poke after the cue period. Note the peak in encoding that is isolated to cluster 3 (black arrow). (B) CPD similarly aligned to choice, demonstrates that this encoding occurs before reward delivery on the current trial. (C) Sensitivity, d' , across time and neurons for reward history has activity isolated to cluster 3 (magenta circle). d' is sorted within the PSTH-based clusters by time-to-peak.

The online version of this article includes the following figure supplement(s) for figure 5:

Figure supplement 1. Linear regression on recent trial history.

Figure supplement 2. Clustering on later trial events.

Notably, cluster 3 also exhibited the most prominent encoding of reward outcome compared to the other clusters (**Figure 4A–B**, green, bottom row). This suggests that this subset of neurons may be specialized for representing or even integrating information about reward outcomes on previous and current trials. We wondered if this might reflect adaptive value coding, in which value representations in OFC have been observed to dynamically reflect the statistics of offered rewards (*Kobayashi et al., 2010; Webb et al., 2014; Yamada et al., 2018; Padoa-Schioppa, 2009; Rustichini et al., 2017; Conen and Padoa-Schioppa, 2019*). Adaptive value coding, which is thought to reflect a divisive normalization or range adaptation mechanism, allows the circuit to efficiently represent values in diverse contexts in which their magnitudes may differ substantially. As such, it provides key computational advantages, such as efficient coding, or the maximization of mutual information between neural signaling and the statistics of the environment (*Webb et al., 2014; Yamada et al., 2018; Padoa-Schioppa, 2009; Rustichini et al., 2017; Simoncelli and Olshausen, 2001; Carandini and Heeger, 2011; Tymula and Glimcher, 2020; Louie and Glimcher, 2012*).

According to divisive normalization models of subjective value, the value of an option or outcome is divided by a recency-weighted average of previous rewards (*Tymula and Glimcher, 2020; Louie and Glimcher, 2012; Webb et al., 2014*). Therefore, if neurons in OFC exhibit adaptive value coding, we would predict that they would exhibit stronger reward responses following unrewarded trials, and weaker responses following rewarded trials (**Figure 6C**). Put another way, regressing the firing rate against current and previous rewards should reveal coefficients with opposite signs (*Kennerley et al., 2011*). Neurons with regression coefficients for current and previous rewards with the same sign would be modulated by current and previous rewards, but not in a way that is consistent with the adaptive value coding hypothesis (**Figure 6D**).

To test this hypothesis, we regressed the firing rates of neurons in a 1 s window after reward receipt against reward win/loss outcomes on the current trial, as well as the win/loss outcome from the previous trial (**Figure 6**, Materials and methods). We found that 180 neurons (27 % of population) exhibited significant coefficients for both current reward and reward history regressors, and 45 % of these units demonstrated adaptive value coding indicated by the opposite signs of regression coefficients for current and previous reward outcomes (81 units, **Figure 6A**, blue). To determine if these adaptive neurons preferentially resided in a particular cluster, we calculated the cluster-specific probability of a neuron demonstrating either significant coefficients for rewarded volume and reward

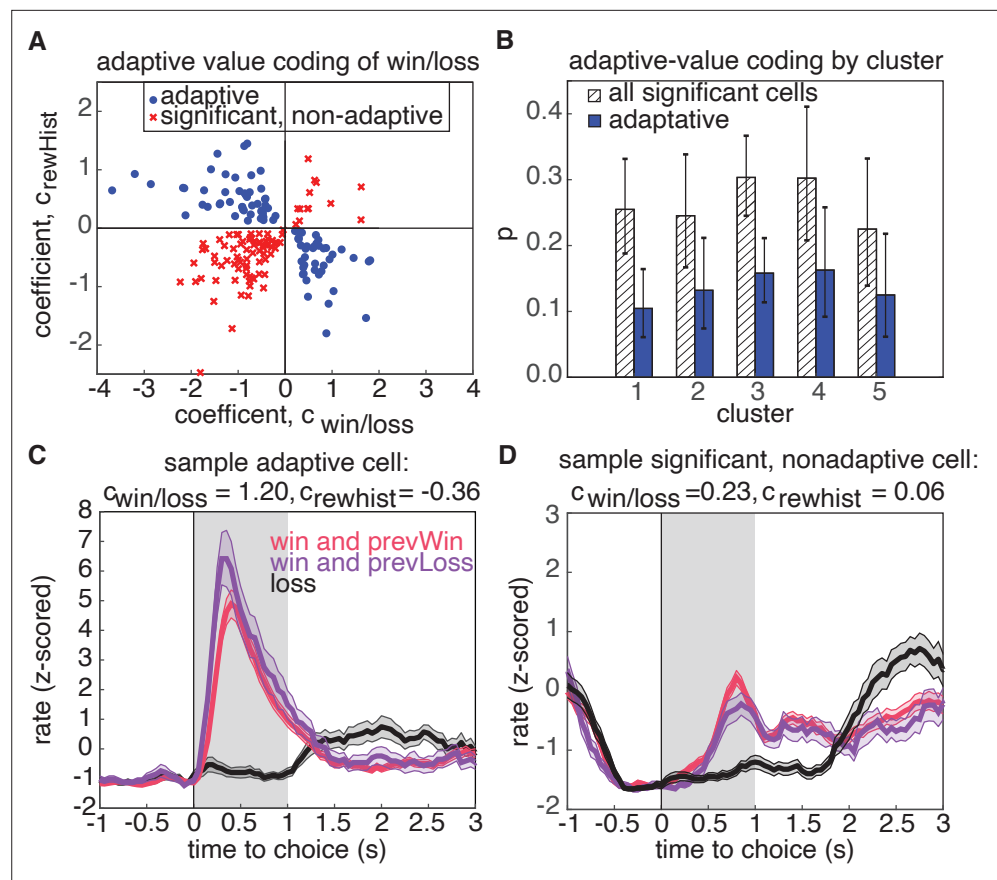


Figure 6. Adaptive value coding analysis using linear regression of firing rates for the 1 s period following reward delivery against current and previous reward outcomes. **(A)** Regression coefficients for the current reward outcome, $c_{win/loss}$ (parameterized as win = 1, and loss = -1), and previous trial outcome, $c_{rewHist}$. 180 neurons had significant coefficients for both regressors ($p < 0.05$, t-test), and 81 neurons had coefficients with opposite signs, consistent with adaptive value coding (blue dots). The remaining neurons have differential responses due to reward history, but inconsistent with adaptive value coding (red crosses). **(B)** Probability that a model with significant regressors for both current and past reward outcome would come from a given cluster. Shaded regions denote all models from panel A, and blue bars show the probability for adaptive neurons only. Error bars are the 95 % confidence interval of the mean for a binomial distribution with observed counts from each cluster. **(C)** Example cell demonstrating adaptive value coding. Shaded gray region denotes time window used to compute mean firing rate for the regression. **(D)** Sample cell demonstrating significant modulation due to reward history, but with a relationship inconsistent with adaptive value coding.

The online version of this article includes the following figure supplement(s) for figure 6:

Figure supplement 1. Linear regression of RPE.

Figure supplement 2. Adaptive value coding of rewarded volume.

history, or having coefficients with opposite signs, consistent with adaptive value coding (**Figure 6B**). We found that no cluster preferentially contained adaptive units with higher probability (comparison of 95 % binomial confidence intervals). We also investigated whether adaptive value coding was present for rewarded volume representations, and similarly found that neurons did not preferentially reside in any one cluster (30 adaptive units out of 79 significant units, **Figure 6—figure supplement 2**). Notably, activity during the reward epoch did not appear to encode a reward prediction error (**Figure 6—figure supplement 1**), defined as the difference between the expected value and the outcome of the chosen option, as has been previously reported in rat and primate OFC (**McDannald et al., 2011; Takahashi et al., 2009; Takahashi et al., 2013; Kennerley et al., 2011; Miller et al., 2018**).

Neurons in Cluster 3 have similar responses to previously reported striatum projection neurons

Studies of OFC indicate that long-range projections to different downstream circuits exhibit distinct encoding and/or make distinct contributions to behavior (*Groman et al., 2019; Hirokawa et al., 2019; Namboodiri et al., 2019*). For instance, previous work from *Hirokawa et al., 2019* used optogenetic tagging to record from OFC neurons that project to the ventral striatum, and found that these cells exhibited responses that seemed to correspond to a distinct cluster with stereotyped responses. This cluster encoded the trial outcome just after reward delivery, with larger responses following unrewarded trials, and also encoded the negative integrated value of the chosen option during the inter-trial interval, until the start of the next trial (**Figure 7A**). We examined the average PSTHs of each cluster aligned to trial outcome and the start of the next trial, to see if any cluster resembled the corticostriatal responses described in *Hirokawa et al., 2019*. Specifically, we plotted the cluster-averaged firing rates for trials classified as a loss, a low-reward trial (6 or 12 μ l), or a high-reward trial (24 or 48 μ l) (**Figure 7B**). We found that several clusters prominently encoded the reward outcome after reward (clusters 1, 3, 5); however, only cluster 3 encoded the negative (i.e., inverted) value of the reward during the inter-trial interval until the start of the next trial. Therefore, one of the clusters that we identified in our data exhibits qualitatively similar responses to the corticostriatal projection neurons identified by clustering in *Hirokawa et al., 2019*.

Discussion

We analyzed neural responses in IOFC from rats performing a value-based decision-making task, in which they chose between left and right offers with explicitly cued reward probabilities and volumes. Despite the apparent response heterogeneity in our dataset, two independent clustering methods revealed that neurons belonged to one of a small number of distinct clusters. We clustered based on the marginalized PSTHs over all trials, and found that subpopulations of neurons exhibited characteristic temporal response profiles. To our knowledge, using temporal responses that do not explicitly contain conditional information is a novel approach for identifying distinct neural response profiles in prefrontal cortex. We also clustered based on a conditional feature space for each neuron, consistent with previous work (*Namboodiri et al., 2019; Hirokawa et al., 2019*). The feature space for each neuron corresponded to its average firing rate on different trial types, in select time windows that most closely corresponded to the differentiating covariate on each trial (e.g. wins vs. losses at the time of reward feedback). Notably, these independent clustering methods identified robustly identifiable groups of neurons for two of the clusters, indicating that neurons in these clusters can be identified by either their temporal response profiles or their task variable encoding, as defined by the feature space.

Previous studies that clustered using a conditional feature space identified a larger number of clusters than we did here (*Namboodiri et al., 2019; Hirokawa et al., 2019*). This could reflect the different metrics used to select the number of clusters (the gap statistic in this study; compared to PCA and silhouette scores [*Namboodiri et al., 2019*], or adjusted Rand index [*Hirokawa et al., 2019*]). Another difference is the time window that was used to generate the conditional feature space. In *Hirokawa et al., 2019*, for instance, the authors used the same time window for each feature, which was after the rat made a choice but before it received feedback about the outcome of that choice. There was no such epoch in our behavioral paradigm, precluding a more direct comparison. However, a principled choice of clusters using the gap statistic still provided a useful tool for investigating the encoding of task variables in IOFC.

Rat IOFC weakly encodes reward attributes

Our task design allowed us to isolate neural responses to sensory cues that conveyed information about distinct reward attributes, because these cues were presented independently and variably in time. However, our GLM revealed weak encoding of these reward attributes – reward probability and volume – across all clusters, regardless of the clustering method. This was observable by examination of the relative magnitude of the flash and click kernels in individual neurons, and also by the CPD metric. The average flash and click CPD values were an order of magnitude smaller than for the other covariates, indicating that flashes and clicks did not contribute substantially to neural

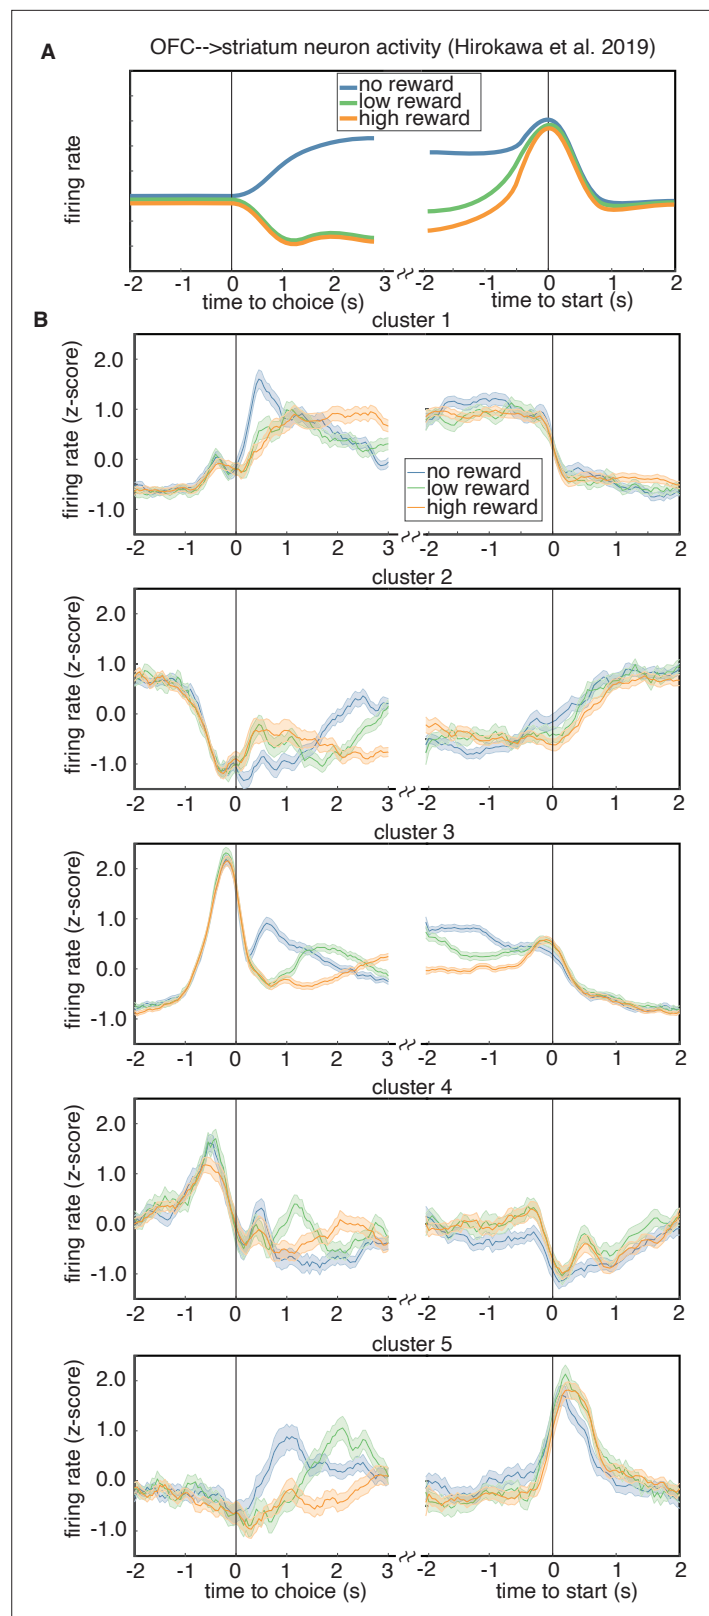


Figure 7. Rewarded volume encoding by cluster. **(A)** Schematic of the encoding properties of a cluster of putative corticostriatal projection neurons from *Hirokawa et al., 2019*. Neurons from this cluster encode the outcome of a trial just after choice, as well as negative (i.e. inverted) integrated value during the intertrial interval, until the start of the next trial. **(B)** Cluster-averaged response to rewarded volume magnitude. Each sub-panel shows the average response to either a trial loss (blue), a low reward (6 or 12 μ l, green), or a high reward (24 or 48 μ l, orange).

firing. Therefore, while behavioral analyses have shown that rats used the flashes and clicks to guide their choices (*Constantinople et al., 2019b*), these cues were not strongly represented in IOFC. This weak encoding of reward attributes, whose combination would specify the subjective value of each offer (*Constantinople et al., 2019b*), is potentially consistent with a recent study of rat IOFC during a multi-step decision-making task that enabled dissociation of choice and outcome values (*Miller et al., 2017; Miller et al., 2018*). Recordings in that study revealed weak encoding of choice values, but strong encoding of outcome values, and optogenetic perturbations suggested a critical role for IOFC in guiding learning but not choice (*Miller et al., 2018*). Other studies in rat IOFC have reported strong encoding of reward and outcome values specifically following action selection, but not preceding choice (*Sul et al., 2010; Miller et al., 2018; Steiner and Redish, 2014; Steiner and Redish, 2012; Mainen and Kepecs, 2009; Hirokawa et al., 2019*). However, recordings from rat IOFC in sensory preconditioning paradigms have revealed cue-evoked responses that may reflect inferred value (*Sadacca et al., 2018; Takahashi et al., 2013*). Studies in non-human primates have also reported strong encoding of offer values in OFC after presentation of a stimulus, and before choice (*Padoa-Schioppa and Assad, 2006; Kennerley et al., 2011; Lin et al., 2020; Rich and Wallis, 2016*). A recent study in mouse OFC used olfactory cues to convey reward attributes (juice identity and volume) and reported encoding of offer values before choice (*Kuwabara et al., 2020*). It is unclear what factors may account for the pronounced encoding of offer value before choice in *Kuwabara et al., 2020*, but minimal encoding of offer value before choice in our data. One important difference is that the cues that conveyed reward attributes in the present study needed to be integrated over time. The integration process itself likely does not occur in OFC (*Lin et al., 2020*). It is possible that OFC would only represent the offer value once the animal had accumulated enough sensory evidence to infer it. If that inference occurred at a different time on each trial, depending on the strength of the sensory evidence and stochasticity of the accumulation process, then representations of offer value would not be observable in the trial-averaged neural response.

Adaptive coding in rat IOFC

Previous studies in primate medial (*Yamada et al., 2018*) and central-lateral (*Kennerley et al., 2011; Tremblay and Schultz, 1999; Padoa-Schioppa, 2009; Rustichini et al., 2017*) OFC have reported subsets of neurons that adjust the gain of their firing rates to reflect the range of offered rewards, a phenomenon referred to as adaptive value coding. This type of coding is efficient because it would allow OFC to accurately encode values in diverse contexts that may vary substantially in reward statistics (*Webb et al., 2014*), analogous to divisive normalization in sensory systems (*Carandini and Heeger, 2011; Simoncelli and Olshausen, 2001*).

According to divisive normalization models of subjective value, the value of an option is divided by a recency-weighted average of previous rewards (*Louie and Glimcher, 2012; Tymula and Glimcher, 2020; Webb et al., 2014*). Therefore, we would predict that neurons implementing adaptive value coding would exhibit stronger responses following unrewarded trials, and weaker responses following rewarded trials. Consistent with this hypothesis, we identified a subset of neurons that had significant coefficients for rewards on current and previous trials, with opposite signs, and while this fraction was somewhat modest (~15%), it was comparable to the proportion of adaptive value coding neurons observed in central-lateral primate OFC (*Kennerley et al., 2011*). Notably, we did not find any evidence of encoding of a reward prediction error during this epoch, or the difference between the expected value of the chosen lottery and the outcome. Therefore, the differential encoding of reward outcomes depending on reward history reflected a discrepancy between reward outcomes and recent experience, not a discrepancy between reward outcomes and expectations (as in a reward prediction error). These were dissociable in our task, due to the sensory cues that explicitly indicated expected value on each trial. We emphasize that adaptive value coding is likely not a specialization of IOFC, but probably occurs broadly in brain areas that represent subjective value.

Mixed selectivity in OFC

The complexity and diversity of responses found in the prefrontal cortex, including the OFC, has led to questions about whether or not neurons with mixed selectivity or specialized responses represent behavioral and cognitive variables (*Rigotti et al., 2013; Hirokawa et al., 2019; Raposo et al., 2014; Mante et al., 2013; Namboodiri et al., 2019*). We tested if OFC neural responses belonged

to clusters in our task using two separate methods: the gap statistic, and the PAIRS statistic. Both methods confirm that there is statistically significant clustering in our data. We emphasize, however, that mixed selectivity is fundamentally a property of individual neurons. We did not directly investigate whether individual IOFC neurons in our dataset exhibited mixed selectivity, but instead focused on encoding at the level of clusters. The broad encoding of task variables in each cluster suggests that IOFC neurons probably exhibit mixed selectivity, but a formal assessment was beyond the scope of our study.

Representations of reward history in OFC

In studies that require animals to learn and update the value of actions and outcomes from experience (i.e. in accordance with reinforcement learning), values are often manipulated or changed over the course of the experiment to assess behavioral flexibility, value-updating, and goal-directed behavior. In rodents and primates, lesion and perturbation studies have shown that the OFC is critical for inferring value in these contexts, suggesting an important role in learning and dynamically updating value estimates for model-based reinforcement learning (*Gallagher et al., 1999; Izquierdo et al., 2004; Pickens et al., 2005; Noonan et al., 2010; West et al., 2011; Gremel and Costa, 2013; Miller et al., 2017; Miller et al., 2018*). Neural recordings in these dynamic paradigms have revealed activity in OFC that reflects reward history and outcome values, which could subserve evaluative processing and learning (*Nogueira et al., 2017; Sul et al., 2010; Miller et al., 2018*).

Other studies, including this one, have used sensory cues (e.g. static images, odors) to convey information about reward attributes. Once learned, the mapping between these cues and reward attributes is fixed, and the subject must choose between options with explicitly cued values. Neurons in the central-lateral primate OFC and rat IOFC have been shown to represent the values associated with these sensory cues in their firing rates, as well as reward outcomes and outcome values (*Thorpe et al., 1983; Schoenbaum et al., 1998; Tremblay and Schultz, 1999; Wallis and Miller, 2003; Padoa-Schioppa and Assad, 2006; Sul et al., 2010; Padoa-Schioppa, 2011; Levy and Glimcher, 2012; Rudebeck and Murray, 2014; Sadacca et al., 2018; Hirokawa et al., 2019; Lin et al., 2020*). However, the extent to which OFC is causally required for these tasks is a point of contention, and may differ across species (*Ballesta et al., 2020; Gardner et al., 2020; Kuwabara et al., 2020*). Notably, even when reward contingencies are fixed over trials, animals often show sequential learning effects (*Lak et al., 2020; Constantinople et al., 2019b*), and reward history representations in OFC have been reported in tasks with stable reward contingencies (*Constantinople et al., 2019a; Padoa-Schioppa, 2013; Kennerley et al., 2011*).

In this study, we have described dynamic trial-by-trial changes in firing rates that reflected reward history just preceding the choice. These reward history representations might influence ongoing neural dynamics supporting the upcoming choice (*Mochol et al., 2021*), and/or mediate learning. In contrast to broadly distributed adaptive value coding, this activity was restricted to a particular subset of neurons that were identifiable by two independent clustering methods, and that exhibited the strongest encoding of reward outcomes. Intriguingly, the responses of neurons in this cluster (cluster 3) are qualitatively similar to the responses of identified corticostriatal cells in *Hirokawa et al., 2019*, raising the possibility that cluster 3 neurons also project to the striatum. IOFC neurons that project to the striatum likely mediate learning: ablating OFC projections to the ventral striatum during a reversal learning task demonstrated that this projection specifically supports learning from negative outcomes (*Groman et al., 2019*).

In reinforcement learning accounts of basal ganglia function, it is thought that cortical inputs to the striatum encode information about the state the animal is in, and corticostriatal synapses store the learned values of performing actions in particular states in their synaptic weights (i.e. Q-values), which can be updated and modulated by dopamine-dependent plasticity (*Averbeck and Costa, 2017; Doya, 2000*). Coincident activation of cortical inputs and striatal spiking can tag corticostriatal synapses for plasticity in the presence of dopamine (*Yagishita et al., 2014*). One reason that IOFC neurons might encode reward history at the time of choice is so that contextual inputs reflecting the animal's state would include its recent reward history. We have previously shown that optogenetic perturbations of IOFC in this task, triggered when rats exited the center port, disrupted sequential trial-by-trial learning effects (*Constantinople et al., 2019a*). These sequential learning effects are not optimal for decision making in this task, as trials were independent and contained explicitly

cued reward contingencies. Suboptimal sequential biases are observed across species and domains (Blanchard et al., 2014; Croson and Sundali, 2005; Neiman and Loewenstein, 2011), and reflect an innate difficulty in judging true randomness in an environment (Nickerson, 2002). The neural circuits supporting these sequential biases could potentially derive from the subpopulation of IOFC neurons in cluster 3, which may mediate the effect of previous trials on the animal's behavior either by influencing IOFC dynamics supporting the upcoming choice (Mochol et al., 2021), or via projections to the striatum that mediate learning.

Materials and methods

Animal subjects and behavior

The details for the animal subjects, behavioral task, and electrophysiological recordings have been described elsewhere (Constantinople et al., 2019a; Constantinople et al., 2019b). Briefly, neural recordings from three male Long-Evans rats were used in this work. Animal use procedures were approved by the Princeton University Institutional Animal Care and Use Committee and carried out in accordance with National Institutes of Health standards.

Rats were trained in a high-throughput facility using a computerized training protocol. The task was performed in operant training boxes with three nose ports, each containing an LED. When the LED from the center port was illuminated, the animal was free to initiate a trial by poking his nose in that port (trial start epoch). While in the center port, rats were continuously presented with a train of randomly timed auditory clicks from a left and right speaker. The click trains were generated by Poisson processes with different underlying rates (Hanks et al., 2015); the rates from each speaker conveyed the water volume baited at each side port. Following a variable pre-flash interval ranging from 0 to 350 ms, rats were also presented with light flashes from the left and right side ports, where the number of flashes conveyed reward probability at each port. Each flash was 20 ms in duration, presented in fixed bins, and spaced every 250 ms to avoid perceptual fusion of consecutive flashes. After a variable post-flash delay period from 0 to 500 ms, there was an auditory 'go' cue and the center LED turned back on. The animal was then free to leave the center port (exit center port epoch) and choose the left or right port to potentially collect reward (choice epoch).

Electrophysiology and data pre-processing for spike train analyses and model inputs

Tetrodes were constructed from twisted wires that were either PtIr (18 μm , California Fine Wire) or NiCr (25 μm , Sandvik). Tetrode tips were platinum- or gold-plated to reduce impedances to 100–250 k Ω at 1 kHz using a nanoZ (White Matter LLC). Microdrive assemblies were custom-made as described previously (Aronov and Tank, 2014). Each drive contained eight independently movable tetrodes, plus an immobile PtIr reference electrode. Each animal was implanted over the right OFC. On the day of implantation, electrodes were lowered to 4.1 mm DV. Animals were allowed to recover for 2–3 weeks before recording. Shuttles were lowered 30–60 μm approximately every 2–4 days.

Data were acquired using a Neuralynx data acquisition system. Spikes were manually sorted using MClust software. Units with fewer than 1 % inter-spike intervals less than 2 ms were deemed single units. For clustering and model fitting, we restricted our analysis to single units that had a mean firing rate greater than 1 Hz (659 units). To convert spikes to firing rates, spike counts were binned in 50 ms bins and smoothed using Matlab's smooth.m function with a 250 ms moving window. Similarly, our neural response model fit spike counts in discretized bins of 50 ms. When parsing data into cross-validated sets of trials balanced across conditions, a single trial was considered as the window of [–2,6]s around trial start. In all other analyses, conditional responses were calculated on trials with a window [–2,4]s for data aligned to trial start, and [–4,4]s for choice-aligned or exit center port-aligned responses.

Clustering of IOFC responses

Feature space parametrization

To analyze the heterogeneity in time-dependent IOFC responses, our first clustering procedure utilized trial-averaged PSTHs from each neuron to construct the feature space for clustering. Specifically, for a set of N neurons and trials of T timepoints, we z-scored the PSTH of each neuron and combined

all responses into a matrix $Z \in \mathbb{R}^{N \times T}$, then performed PCA to obtain principal components, W , and score, M , as $Z = MW^T$. We found that $k = 18$ components explained >95 % of the covariance in Z and used the first k columns of M , the PSTH projected onto the top k PC, as our feature space on which to perform clustering.

Our second clustering procedure used time-averaged, conditional neural responses to construct the feature space for k-means clustering. This is similar in form to the approach in **Hirokawa et al., 2019**. The feature space consisted of its z-scored firing rate conditioned on choice, reward outcome, reward history, presented offer value (EV of left and right offers), and rewarded volume, in time bins that most often corresponded to differential encoding of each variable (**Figure 2A and B**). Specifically, we used conditional PSTHs that depend on a single condition, and marginalized away all other conditions. The conditions, $X^{(M)} = x_j$, are grouped into three categories for our task. Choice and outcome information is $\mathbf{X}^{(\text{reward})} \in \{\text{win,loss}\}$, $\mathbf{X}^{(\text{rewardHistory})} \in \{\text{previous Win,previous loss}\}$, $\mathbf{X}^{(\text{choice})} \in \{\text{left,right}\}$. Presented offer attributes on left and right ports are the expected value of reward as $EV = pV$, where p is reward probability conveyed through flashes, and V is the volume offer conveyed through Poisson clicks. Values were binned on a log-2 scale: $\mathbf{X}^{(\text{EVL})}, \mathbf{X}^{(\text{EVR})} \in \{[0, 6], [6, 12], [12, 24], [24, 48]\mu\text{l}\}$. Rewarded value is $\mathbf{X}^{(\text{rewVol})} \in \{0, 6, 12, 24, 48\mu\text{l}\}$.

The conditional PSTH responses were z-scored, and then each conditional PSTH was averaged over the time window in which the behavioral variable was maximally encoded (dictated by peak location in CPD, see Materials and methods below) to obtain a conditional firing rate as a single feature for clustering. Specifically, the time windows for reward and reward volume information were averaged over $[0, 3]s$ after reward delivery, $[-1, 2]s$ after trial start for reward history, $[0, 1.5]s$ after exiting the center port for left/right choice, and $[-1, 0]s$ before the animal's choice (i.e. entering the side port) for expected value of presented offers. These 19 features were combined and pre-processed using PCA in the same way as the PSTH-based clustering to yield 11 features.

Evaluation of k-means cluster quality: gap statistic

For each clustering procedure, we utilized k-means clustering to locate groups of functionally distinct responses in IOFC, and used the gap statistic criterion to determine a principled choice of the best number of clusters (evalclusters.m in Matlab) (**Tibshirani et al., 2001**). Specifically, we locate the largest cluster size K for which there was a significant jump in gap score $Gap(K)$,

$$Gap(K) \geq Gap(K - 1) + 2SE(K - 1). \tag{2}$$

This is similar to a standard option in evalclusters.m ('SearchMethod' = 'firstMaxSE'), which finds the smallest instance in which a non-significant jump in cluster size is located. The two methods often agree. Finally, we used 5000 samples for the reference distribution to ensure convergence of results in the gap statistic.

Alternative cluster evaluation metrics: silhouette score and adjusted Rand index (ARI)

The silhouette score calculates a weighted difference between inter- and intra-cluster distances. It yields high values for data that are tightly packed in the feature space and also far away from neighboring clusters. Intra-cluster distances for a data point i quantify the similarity of data within a cluster as,

$$a(i) = \frac{1}{N_{C_i} - 1} \sum_{j \in C_i, i \neq j} d(i, j), \tag{3}$$

where $d(i, j)$ is the Euclidean distance between points i and j , and N_{C_i} is the number of data points in cluster C_i . Similarly, the inter-cluster distance of data point i to data points of the closest neighboring cluster is given by

$$b(i) = \min_{k \neq i} \frac{1}{N_{C_k}} \sum d(i, j). \tag{4}$$

The silhouette score is given by

$$S = \frac{1}{N} \sum_i \frac{b(i) - a(i)}{\max[a(i), b(i)]}, \tag{5}$$

where N is the number of data points. Silhouette scores were calculated in python using sklearn.metrics.silhouette_score with the default options (i.e, metric='euclidean').

The adjusted Rand index is a measure of reproducibility of a dataset. In this context, it calculates the probability that two cluster labelings on related data would arise by chance. We calculated ARI in python using sklearn.metrics.adjusted_rand_score. To estimate statistics, we calculated ARI between a set of labels generated from our full dataset and a sub-sampled dataset for each cluster number K . We sampled 90 % of the population without replacement, and generated 100 such datasets to get a distribution of ARI values for each value of K .

To study the behavior of the silhouette score in the same data regime as the IOFC neural data, we generated ground-truth datasets of $N = 500$ data points with different ratios of within-cluster and total data covariance. In particular, we generated datasets with a $K = 5$ clusters (100 data points per cluster) that had the same dimensionality ($D = 121$) and the same total data covariance as the neural data. These clusters had within-cluster covariance that matched the average within-cluster covariance of the IOFC data. We then preprocessed this surrogate data in the same manner as our true data by projecting onto the PC components capturing $> 95\%$ variance (total-data covariance and average within-cluster covariance shown in **Figure 2—figure supplement 6A,B**), and performed silhouette score analysis in this dimensionality-reduced subspace. We investigated different ratios of within-cluster and across-cluster by scaling the total-data covariance by a constant factor (0.1, 0.5, 1, 2, 3, 5) to determine to what extent the silhouette score can detect clusters that are tightly packed in the feature space. To gather statistics for each ratio of variances, we generated 100 datasets with random cluster means (drawn from a multivariate normal distribution) and plotted the mean \pm s.e.m silhouette scores (**Figure 2—figure supplement 6B**). As a comparison, we also performed the gap statistic analysis on the same datasets (**Figure 2—figure supplement 6C**). Finally, we visualized the groupings of these clusters by nonlinear projections (tSNE) in a two-dimensional feature space (**Figure 2—figure supplement 6D**).

Cluster labeling consistency and cluster similarity

To compare the consistency of results between the PSTH clustering and conditional clustering (**Figure 2E**), we calculated $P(C_{\text{conditional}}|C_{\text{PSTH}})$, the conditional probability of a neuron being assigned to cluster $C_{\text{conditional}}$ from the conditional clustering procedure, given that it was assigned to cluster C_{PSTH} in the PSTH cluster procedure. We evaluated the similarity of clusters within a given clustering procedure in two ways. First, we performed TSNE embedding of the features space to visualize cluster similarity in two dimensions (sklearn.manifold.TSNE in python, perplexity = 50, n_iter = 5000) (**Hinton and Roweis, 2002**). We then colored each sample in this 2-D space based on cluster identity (**Figure 2—figure supplement 2A, B**). We quantified the distance amongst clusters by calculating the cluster averaged Mahalanobis distance to the other clusters (**Mahalanobis, 1936**). The Mahalanobis distance $D_M(x)$ calculates the distance of a sample x_A to a distribution B with a known mean, μ_B , and covariance, S_B :

$$D_M(x_A, B) = \sqrt{(x_A - \mu_B)^T S_B^{-1} (x_A - \mu_B)}. \quad (6)$$

The cluster-averaged distances in **Figure 2—figure supplement 2C, D** average D_M over all samples from a given cluster A as

$$D_M(A, B) = \mathbb{E}_{x_A \in A} [D_M(x_A, B)]. \quad (7)$$

PAIRS statistic

To test the null hypothesis that the data is not clustered at all, we followed the approach provided in **Raposo et al., 2014**. We used our pre-processed PSTH and conditional feature spaces that were dimensionality reduced using PCA, and further processed them with a whitening transform so that the data had zero mean and unit covariance. For each data point, we estimated the average angle with k of its closest neighbors, θ_{data} . We then compared the data to independent draws from a reference Gaussian distribution ($N(0, \mathbb{I})$), with the same number of data points and the same dimensionality as our data. To gather statistics, we generated $N = 10,000$ of these null model datasets, and aggregated the estimated angles $\theta_{\text{ref}}^{(n)}$ into a grand reference distribution. We chose our number of nearest neighbors

k for averaging by locating the k such that our reference distribution had a median angle greater than $\pi/4$, following the approach in [Raposo et al., 2014](#) ($k=3$ for PSTH clustering, $k=8$ for conditional clustering). Distributions of angles between neighboring datapoints are presented in [Figure 2—figure supplement 4](#). We further quantified the similarity of these two distributions through a Kolmogorov Smirnov test.

PAIRS is a summary statistic of the entire population, using the median from the data distribution, $\tilde{\theta}_{\text{data}}$, and the median of the grand reference distributions $\tilde{\theta}_{\text{ref}}$,

$$\text{PAIRS} = \frac{\tilde{\theta}_{\text{ref}} - \tilde{\theta}_{\text{data}}}{\tilde{\theta}_{\text{ref}}}. \quad (8)$$

Reference PAIRS statistics were generated for each of the reference data sets, and statistical significance of the PAIRS statistic was assessed by calculating the two-sided p value for the data PAIRS compared to the distribution of reference PAIRS values. We assumed a normal distribution for the PAIRS reference values.

Generalized linear model of neural responses in IOFC

Our neural response model is a generalized linear model with an exponential link function that estimates the probability that spiking from a neuron will occur in time bin t with a rate λ_t , and with Poisson noise, given a set of time-dependent task parameters. Specifically, for a design matrix \mathbf{X} with S columns of task variables $X_t^{(s)}$, the probability of observing y_t spikes from a neuron in time bin $t + \Delta t$ is modeled as,

$$p(y_t|\mathbf{X}) \sim \text{Poisson}(\lambda_t) = \frac{1}{y_t!} (\lambda_t \Delta t)^{y_t} e^{-\lambda_t \Delta t} \quad (9)$$

$$\lambda_t = \exp \left[\sum_{s=1}^S \left(X_{t-\tau:t}^{(s)} * k_s(\tau) \right) + \theta_0 \right], \quad (10)$$

where λ_t is the rate parameter of the Poisson process. Task variables are linearly convolved with a set of response kernels $k_s(\tau)$ that allow for a time-dependent response from a neuron that may occur after (causal) or before (acausal) the event in the trial. The kernels are composed of a set of N_s basis functions, $\phi_s^{(k)}$, linearly weighted by model parameters $\theta_s^{(k)}$:

$$k_s(\tau) = \sum_{j=1}^{N_s} \theta_s^{(j)} \phi_s^{(j)}(\tau). \quad (11)$$

Additionally, we include a parameter θ_0 that captures the background firing rate of each neuron.

We used 15 time-dependent task variables in our model that indicate both the timing of an event in the task, as well as behavioral or conditional information. Parameterization of the model in this manner with a one-hot encoding of each condition per variable allows for asymmetric responses to different outcomes (e.g. wins vs. losses), and also captures the variable timing of each event in the task. The task variables were the following: (4) stimulus variables of left and right clicks and flashes (aligned to trial start). (4) Choice variables for choosing either the left or the right port (aligned to exit center port). Similarly, we included an alternate parameterization of choosing either the safe port ($p = 1$) or the risky port ($p < 1$). (5) Outcome variables were wins or losses on the current trial (aligned to choice); and reward history variables of previous wins, losses, or a previous opt-opt (aligned to trial start). We included a previous reward rate task variable that was calculated as the average rewarded volume from all previous trials in the session (aligned to trial start). We also included a 'session progress' variable as the normalized $[0,1]$ trial number in the session (aligned to trial start). This covariate captures motivational or satiety effects on firing rate over the course of a session. Finally, model comparison of cross-validated log-likelihood indicated that an autoregressive spike-history kernel was not necessary for our model.

Parameter learning, hyperparameter choice, and model validation

The set of parameters θ of the model are the kernel weights $\theta_s^{(j)}$ and the background firing rate θ_0 , and were fit by minimizing \mathcal{L} , the negative-log likelihood $-\log[p(y|\theta)]$ with an additional L_2 penalty acting as a prior over model parameters (Park et al., 2014):

$$\mathcal{L} = -\log p(y|\theta) - \log p(\theta) = \sum_t \lambda_t - y_t \log \lambda_t + \frac{\xi}{2} \|\theta\|^2. \quad (12)$$

Model parameters were chosen through four-fold cross validation, and ξ was found through a grid search optimization on a held-out test set. Specifically, we split the data from each neuron into five equal parts that were balanced among trial contingencies (i.e. equal amounts of wins/losses, previous wins/losses, and left/right choices per partition). One partition (test set) was not utilized in fitting θ , and was held out for later model comparison and hyperparameter choice. The remaining four partitions were used in cross validation to fit four models, with each model fit on three partitions and assessed on the fourth ‘validation’ partition. The model with the lowest negative log-likelihood on the validation set was chosen for further analysis. This procedure was repeated iteratively on an increasingly smaller grid of initial hyperparameter values $\xi \in [10^{-5}, 10]$, and the hyperparameter yielding the lowest negative log-likelihood on the test partition was chosen. We chose this approach for hyperparameter optimization in lieu of approximations such as calculating evidence (Bishop, 2006), as we found the underlying approximations to be limiting for our data. In general, when building our model we assessed the aspects of other hyperparameter choices (i.e. kernel length, symmetric vs. asymmetric conditional kernels, number of task variables) with cross-validated negative log-likelihood on held-out test data. See the model comparison section below for further details. Finally, kernel covariances and standard deviations were estimated using the inverse of the Hessian of \mathcal{L} .

Basis functions

We utilized a log-scaled, raised cosine basis function set for our model (Park et al., 2014). These functions offer the ability to generate impulse-like responses shortly after stimulus presentation, as well as broader, longer time-scale effects. The form of the basis function is given as

$$\phi_j(\tau) = \frac{1}{2} \cos(a \log[\tau] - b_j) + \frac{1}{2}, \quad (13)$$

where a is parameter that controls breadth of support of each basis function, and b_j controls the location of its maxima. The parameters a, b_j were chosen to spread the set of basis functions $\{\phi_j\}$ in roughly a log-linear placement across their range of support. This gives a better coverage than linear spacing, as the basis functions increase in breadth as b_j increases. We used 7 such functions, and additionally augmented our basis set with two decaying exponential basis functions placed in the first two time bins to capture any impulse-like behavior at the onset of the task variable. This gives $M = 9$ basis functions for all kernels in the study. After a cursory model comparison of different kernel lengths we took a conservative approach and utilized a range of support of 4 s for each kernel.

Model comparison

To choose the best-fit model to our data, we performed model comparison on held-out testing data. For comparison, we considered models with additional task variables such as current and previously rewarded volume; as well as reduced models that omitted the variables relating to previous opt out trials, the previous reward rate, and the session progress. We also considered a reduced model that omitted the reward history contribution entirely. In each case, we assessed the population level change in model performance via a Wilcoxon signed rank test (Figure 3—figure supplement 1). Our chosen model demonstrated a significantly lower population median in its negative log-likelihood than other models. We note that while our chosen model performed better than the reduced model at the population level, the median changes in neural responses were relatively slight (Figure 3—figure supplement 1C, left panel). However, some neurons showed relatively strong improvement from introducing previousOptOut, previousRewardRate, and sessionProgress; which motivated us to keep them for the entire population (i.e., Figure 3—figure supplement 1C, middle panel). Additionally, we investigated a symmetrized form of our model that used a binary (± 1) encoding of task variables, as opposed to a one-hot encoding. This model was rejected at the population level via model comparison (not shown).

Coefficient of partial determination

The coefficient of partial determination (CPD) is a measure of the relative amount of % variance explained by a given covariate, and here we use it to quantify the encoding of behavioral information in single units. The CPD is defined as *Miller et al., 2018*:

$$CPD(X^{(i)})_t = \frac{SSE(\mathbf{X}_{-i})_t - SSE(\mathbf{X}_{all})_t}{SSE(\mathbf{X}_{-i})_t} \times 100, \quad (14)$$

where *SSE* is the trial-summed, squared error between data and model. \mathbf{X}_{all} refers to the full model with all covariates. \mathbf{X}_{-i} implies a reduced model with the effect of $X^{(i)}$ omitted. Since our model utilized a one-hot encoding for each condition, we calculated CPD by omitting the following groups of covariates: $\mathbf{X}^{(reward)} \in \{x^{(win)}, x^{(loss)}\}$, $\mathbf{X}^{(rewardHistory)} \in \{x^{(prevWin)}, x^{(prevLoss)}, x^{(prevOptOut)}\}$, $\mathbf{X}^{(choice)} \in \{x^{(left)}, x^{(right)}\}$, $\mathbf{X}^{(clicks)} \in \{x^{(Lclick)}, x^{(Rclick)}\}$, $\mathbf{X}^{(flashes)} \in \{x^{(Lflash)}, x^{(Rflash)}\}$.

Our measure of CPD assessed the encoding of the covariate of interest (e.g., previous win or loss) separately from encoding of the general event-aligned response (e.g., trial start). As such, our reduced model averaged the kernels that corresponded to the event-aligned response to create behaviorally irrelevant task variables. For example, for CPD of reward encoding: $X_t^{(win)}, X_t^{(loss)} \rightarrow \frac{1}{2}(X_t^{(win)} + X_t^{(loss)})$. For reward, reward history, and choice CPD calculations, we additionally weighted each trial type in *SSE* such that each condition contributed equally to CPD, and omitted any bias due to an imbalance in trial statistics. Due to data limitations, we utilized the full set of training, testing, and validation data to calculate CPD. Units with a model fit of $R^2 > 0$ (590/659 units) were used in CPD calculation. Additionally, we found that CPD measures for neurons can sometimes be noisy, and we further excluded neurons as outliers in the cluster-averaged CPD calculation if they had a CPD value > 0.5 . This excluded a further 10 neurons.

We assessed the significance of the CPD result by comparing it to a null distribution of 500 CPD values that were generated by shuffling the trial labels among the relevant covariates (e.g. shuffle win and loss labels across trials, keeping timing of event and all other covariates fixed). CPD was deemed significant if it fell outside of the one-sided 95 % confidence interval of the shuffle distribution, and plotted values in *Figure 4* subtract off the mean of the shuffle distribution from the CPD.

Mutual information

Mutual information (MI) was used to calculate how much information about task variables is contained in LOFC firing rates, in different time windows throughout a trial. We calculated MI between spikes \mathbf{Y}_t and a group of covariates $\mathbf{X}^{(m)} \in \{x_i\}$ (detailed in the above section) as:

$$\begin{aligned} MI_t &= H(\mathbf{X}^{(m)}) - H(\mathbf{Y}_t | \mathbf{X}^{(m)}) \\ &= - \sum_i p(x_i) \log(x_i) + \sum_{j,k} p(x_j) p(y_t | X^{(m)} = x_k) \log[p(y_t | X^{(m)} = x_k)]. \end{aligned} \quad (15)$$

The first term is the entropy of the stimulus, which is calculated from the empirical distribution. The second term is the conditional entropy of spiking, and requires calculation of the conditional distribution $p(y_t | X^{(m)} = x_k)$ by marginalizing over the fully conditional distribution that contains all other covariates from the model.

We modeled the fully conditional distribution $p(y_t | X^{(1)}, X^{(2)}, \dots, X^{(S)}, \dots, X^{(m)} = x_k)$ via sampling of a doubly stochastic process, in which normal distributions of model parameters were propagated through our GLM and Poisson spiking. Specifically, for each time bin within each trial we sampled 500 parameter values from the normal distribution for each θ , where the covariance of θ was estimated as the inverse of the Hessian of \mathcal{L} . These samples were then passed through the exponential nonlinearity to generate a log-normal distribution of λ values, which were used as the rate parameter for a Poisson process that generated spikes. This spiking distribution of 500 samples was truncated at 10 spikes/50 ms bin (a 200 Hz cutoff), and the conditional distribution $p(y_t | X^{(m)} = x_k)$ was then taken as the average over trials in which $X^{(m)} = x_k$. Units with a model fit of $R^2 > 0$ were used in MI calculations. We assessed significance of MI in each time bin by comparing to an analogously created distribution of 500 MI values in which trial labels were shuffled. Significant MI fell outside of the 95 % confidence interval of this distribution.

Waveform analysis

To investigate if different types of neurons may underlie different clusters, we extracted the action potential (AP) peak and after-hyperpolarization (AHP) peak from the waveform associated with each single unit. Waveforms were recorded on tetrodes, and the channel with the highest amplitude signal was used for waveform analysis. An initial manual curation of the waveforms excluded 18/659 units with poor isolation of the largest-amplitude channel, due to damage of the electrode. The remaining waveforms were standardized by upsampling the waveform by a factor of 20, and then mean-subtracting and normalizing by their maximum range. AP and AHP peaks were identified, then a threshold around zero ($\alpha = 0.05$) was set to identify the beginning and end of the peak. We used the AP half-width rather than full width, as we found this to be a more robust measure of AP activity. Clustering of the two putative RSU and FSU units was performed using k-means.

Discriminability for reward history

The discriminability between previous wins and previous losses in **Figure 5C** was calculated in its unsigned form

$$d'_t = \frac{|\mu_{t,(\text{prevWin})} - \mu_{t,(\text{prevLoss})}|}{\sqrt{\frac{1}{2}(\sigma_{t,(\text{prevWin})}^2 + \sigma_{t,(\text{prevLoss})}^2)}}$$

To account for an inflated range of nonsignificant d' values around zero due to the unsigned form of d' , we subtracted from this quantity the mean d' of a trial-shuffled distribution. The shuffled distribution was generated by shuffling trial labels 1,000 times. Significant d' values were identified as being outside of the one-sided 95 % confidence interval of the shuffled distribution. Trial types were balanced when generating the shuffle distributions.

Linear regression models of pre-choice and post-choice epochs

For the linear regression of pre-choice epoch in **Figure 5—figure supplement 1**, the model used the trial history of wins and losses on the previous five trials as regressors, $x_{\text{win/loss}}^{(i)}$, the choice on that trial, $x_{\text{choice}}^{(n)}$, and an offset term, c_0 :

$$r_n = c_0 + x_{\text{choice}}^{(n)} + \sum_{i=n-1}^{n-5} c_{\text{win/loss}}^{(i)} x_{\text{win/loss}}^{(i)} \quad (16)$$

The regressors were binary +1/−1 variables for win(+ 1)/loss(−1) outcomes and left(+ 1)/right(−1) choice. Model coefficients c were fit with Matlab's fitlm function, and significance was assessed via an F-test comparing to the baseline model containing only c_0 . The significance of each coefficient was assessed via a t-test with a cutoff of $p < 0.05$ for significance.

Similarly, the results in **Figure 6** and **Figure 6—figure supplement 2** investigating the adaptation of reward representations regressed r_n , the average firing rate in the 1 s interval after reward onset on each trial n (post-choice epoch), against previous and current trial outcomes. The first model investigated adaptation of the rewarded outcome representation by including a binary win(+ 1)/loss(−1) regressor for current trial reward outcome, a binary win(+ 1)/loss(−1) regressor for previous reward outcome, the left(+ 1)/right(−1) choice on that trial, and an offset term:

$$r_n = c_{\text{win/loss}} x_{\text{win/loss}}^{(n)} + c_{\text{rewhist}} x_{\text{win/loss}}^{(n-1)} + c_{\text{choice}} x_{\text{choice}}^{(n)} + c_0 \quad (17)$$

Similarly, the other model investigating adaptation of reward volume representations instead used a regressor for rewarded volume, and a binary loss(1/0) regressor for outcome on the current trial:

$$r_n = c_{\text{vol}} x_{\text{vol}}^{(n)} + c_{\text{loss}} x_{\text{loss}}^{(n)} + c_{\text{rewhist}} x_{\text{win/loss}}^{(n-1)} + c_{\text{choice}} x_{\text{choice}}^{(n)} + c_0 \quad (18)$$

Finally, the regression in **Figure 6—figure supplement 1** modeled the post-choice epoch using a reward prediction error as a regressor:

$$r_n = c_{\text{rpe}} x_{\text{rpe}}^{(n)} + c_{\text{choice}} x_{\text{choice}}^{(n)} + c_0 \quad (19)$$

where the RPE regressor was the difference in rewarded volume and the expected value of the chosen option on that trial: $x_{\text{rpe}}^{(n)} = V_{\text{reward}}^{(n)} - p^{(n)}V_{\text{choose}}^{(n)}$.

Acknowledgements

We thank Carla Golden, Alex Piet, Abby Muthusamy, and Anne Churchland for helpful discussions and comments on the manuscript.

Additional information

Funding

Funder	Grant reference number	Author
National Institute of Mental Health	R00MH111926	Christine M Constantinople
National Institute of Mental Health	R01MH125571	Cristina Savin Christine M Constantinople
Alfred P. Sloan Foundation		Christine M Constantinople
Klingenstein-Simons Fellowship Award in Neuroscience		Christine M Constantinople
National Science Foundation	1922658	Cristina Savin

The funders had no role in study design, data collection and interpretation, or the decision to submit the work for publication.

Author contributions

David L Hocker, Data curation, Formal analysis, Methodology, Software, Visualization, Writing - original draft, Writing - review and editing; Carlos D Brody, Senior author on associated article, Writing - review and editing; Cristina Savin, Conceptualization, Formal analysis, Funding acquisition, Methodology, Project administration, Supervision, Writing - review and editing; Christine M Constantinople, Conceptualization, Data curation, Funding acquisition, Investigation, Methodology, Project administration, Supervision, Writing - original draft, Writing - review and editing

Author ORCIDs

David L Hocker  <http://orcid.org/0000-0003-3091-421X>

Carlos D Brody  <http://orcid.org/0000-0002-4201-561X>

Christine M Constantinople  <http://orcid.org/0000-0003-4435-4460>

Ethics

This study was performed in accordance with the National Institutes of Health standards. Animal use procedures were approved by the Princeton University Institutional Animal Care and Use Committee (protocol #1853).

Decision letter and Author response

Decision letter <https://doi.org/10.7554/eLife.70129.sa1>

Author response <https://doi.org/10.7554/eLife.70129.sa2>

Additional files

Supplementary files

- Transparent reporting form

Data availability

Physiology, behavior, GLM statistical modeling, and clustering results source data are publicly available on Zenodo (<https://doi.org/10.5281/zenodo.5592702>). Code for the GLM model fit, clustering,

and all subsequent analysis will also be provided publicly via Github at <https://github.com/constantinoplelab/published> (copy archived at <https://archive.softwareheritage.org/swh:1:rev:cd86802e0efc64e9d6b78408f260bd0982000b8c>).

The following dataset was generated:

Author(s)	Year	Dataset title	Dataset URL	Database and Identifier
Hokcer D, Savin C, Constantinople C	2021	Subpopulations of neurons in IOFC encode previous and current rewards at time of choice	https://doi.org/10.5281/zenodo.5592702	Zenodo, 10.5281/zenodo.5592702

References

- Aronov D**, Tank DW. 2014. Engagement of neural circuits underlying 2d spatial navigation in a rodent virtual reality system. *Neuron* **84**: 442–456. DOI: <https://doi.org/10.1016/j.neuron.2014.08.042>, PMID: 25374363
- Averbeck BB**, Costa VD. 2017. Motivational neural circuits underlying reinforcement learning. *Nature Neuroscience* **20**: 505–512. DOI: <https://doi.org/10.1038/nn.4506>, PMID: 28352111
- Ballesta S**, Shi W, Conen KE, Padoa-Schioppa C. 2020. Values encoded in orbitofrontal cortex are causally related to economic choices. *Nature* **588**: 450–453. DOI: <https://doi.org/10.1038/s41586-020-2880-x>, PMID: 33139951
- Bishop CM**. 2006. Pattern Recognition and Machine Learning. Springer.
- Blanchard TC**, Wilke A, Hayden BY. 2014. Hot-hand bias in rhesus monkeys. *Journal of Experimental Psychology* **40**: 280–286. DOI: <https://doi.org/10.1037/xan0000033>
- Carandini M**, Heeger DJ. 2011. Normalization as a canonical neural computation. *Nature Reviews. Neuroscience* **13**: 51–62. DOI: <https://doi.org/10.1038/nrn3136>, PMID: 22108672
- Conen KE**, Padoa-Schioppa C. 2019. Partial adaptation to the value range in the macaque orbitofrontal cortex. *The Journal of Neuroscience* **39**: 3498–3513. DOI: <https://doi.org/10.1523/JNEUROSCI.2279-18.2019>, PMID: 30833513
- Constantinople CM**, Piet AT, Bibawi P, Akrami A, Kopec C, Brody CD. 2019a. Lateral orbitofrontal cortex promotes trial-by-trial learning of risky, but not spatial, biases. *eLife* **8**: e49744. DOI: <https://doi.org/10.7554/eLife.49744>, PMID: 31692447
- Constantinople CM**, Piet AT, Brody CD. 2019b. An analysis of decision under risk in rats. *Current Biology* **29**: 2066–2074. DOI: <https://doi.org/10.1016/j.cub.2019.05.013>, PMID: 31155352
- Croson R**, Sundali J. 2005. The gambler's fallacy and the hot hand: Empirical data from casinos. *Journal of Risk and Uncertainty* **30**: 195–209. DOI: <https://doi.org/10.1007/s11166-005-1153-2>
- Doya K**. 2000. Complementary roles of basal ganglia and cerebellum in learning and motor control. *Current Opinion in Neurobiology* **10**: 732–739. DOI: [https://doi.org/10.1016/S0959-4388\(00\)00153-7](https://doi.org/10.1016/S0959-4388(00)00153-7), PMID: 11240282
- Gallagher M**, McMahan RW, Schoenbaum G. 1999. Orbitofrontal cortex and representation of incentive value in associative learning. *The Journal of Neuroscience* **19**: 6610–6614. DOI: <https://doi.org/10.1523/JNEUROSCI.19-15-06610.1999>, PMID: 10414988
- Garcia-Dias R**, Prieto CA, Almeida JS, Ordovás-Pascual I. 2018. Machine learning in APOGEE: Unsupervised spectral classification with K-means. *Astronomy & Astrophysics* **612**:A98. DOI: <https://doi.org/10.1051/0004-6361/201732134>
- Garcia-Dias R**, Vieira S, Pinaya WHL, Mechelli A. 2020. Clustering analysis. Garcia-Dias R (Ed). *Machine Learning*. Elsevier. p. 227–247.
- Gardner MPH**, Sanchez D, Conroy JC, Wikenheiser AM, Zhou J, 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. DOI: <https://doi.org/10.1016/j.neuron.2020.08.010>, PMID: 32888408
- Gremel CM**, Costa RM. 2013. Orbitofrontal and striatal circuits dynamically encode the shift between goal-directed and habitual actions. *Nature Communications* **4**: 1–12. DOI: <https://doi.org/10.1038/ncomms3264>, PMID: 23921250
- Groman SM**, Keistler C, Keip AJ, Hammarlund E, DiLeone RJ, Pittenger C, Lee D, Taylor JR. 2019. Orbitofrontal circuits control multiple reinforcement-learning processes. *Neuron* **103**: 734–746. DOI: <https://doi.org/10.1016/j.neuron.2019.05.042>, PMID: 31253468
- Hanks TD**, Kopec CD, Brunton BW, Duan CA, Erlich JC, Brody CD. 2015. Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* **520**: 220–223. DOI: <https://doi.org/10.1038/nature14066>, PMID: 25600270
- Hinton G**, Roweis ST. 2002. Stochastic neighbor embedding. NIPS Citeseer. 833–840.
- Hirokawa J**, Vaughan A, Masset P, Ott T, Kepecs A. 2019. Frontal cortex neuron types categorically encode single decision variables. *Nature* **576**: 446–451. DOI: <https://doi.org/10.1038/s41586-019-1816-9>, PMID: 31801999
- Izquierdo A**, Suda RK, Murray EA. 2004. Bilateral orbital prefrontal cortex lesions in rhesus monkeys disrupt choices guided by both reward value and reward contingency. *The Journal of Neuroscience* **24**: 7540–7548. DOI: <https://doi.org/10.1523/JNEUROSCI.1921-04.2004>, PMID: 15329401

- Kennerley SW**, Behrens TEJ, Wallis JD. 2011. Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nature Neuroscience* **14**: 1581–1589. DOI: <https://doi.org/10.1038/nn.2961>, PMID: [22037498](https://pubmed.ncbi.nlm.nih.gov/22037498/)
- Kobayashi S**, Pinto de Carvalho O, Schultz W. 2010. Adaptation of reward sensitivity in orbitofrontal neurons. *The Journal of Neuroscience* **30**: 534–544. DOI: <https://doi.org/10.1523/JNEUROSCI.4009-09.2010>, PMID: [20071516](https://pubmed.ncbi.nlm.nih.gov/20071516/)
- Kuwabara M**, Kang N, Holy TE, Padoa-Schioppa C. 2020. Neural mechanisms of economic choices in mice. *eLife* **9**: e49669. DOI: <https://doi.org/10.7554/eLife.49669>, PMID: [32096761](https://pubmed.ncbi.nlm.nih.gov/32096761/)
- Lak A**, Hueske E, Hirokawa J, Masset P, Ott T, Urai AE, Donner TH, Carandini M, Tonegawa S, Uchida N, Kepecs A. 2020. Reinforcement biases subsequent perceptual decisions when confidence is low, a widespread behavioral phenomenon. *eLife* **9**: e49834. DOI: <https://doi.org/10.7554/eLife.49834>, PMID: [32286227](https://pubmed.ncbi.nlm.nih.gov/32286227/)
- Levy DJ**, Glimcher PW. 2012. The root of all value: a neural common currency for choice. *Current Opinion in Neurobiology* **22**: 1027–1038. DOI: <https://doi.org/10.1016/j.conb.2012.06.001>, PMID: [22766486](https://pubmed.ncbi.nlm.nih.gov/22766486/)
- Lin Z**, Nie C, Zhang Y, Chen Y, Yang T. 2020. Evidence accumulation for value computation in the prefrontal cortex during decision making. *PNAS* **117**: 30728–30737. DOI: <https://doi.org/10.1073/pnas.2019077117>
- Louie K**, Glimcher PW. 2012. Efficient coding and the neural representation of value. *Annals of the New York Academy of Sciences* **1251**: 13–32. DOI: <https://doi.org/10.1111/j.1749-6632.2012.06496.x>, PMID: [22694213](https://pubmed.ncbi.nlm.nih.gov/22694213/)
- Mahalanobis PC**. 1936. On the generalized distance in statistics. Proceedings of the National Institute of Sciences (Calcutta) National Institute of Science of India; 1936. .
- Mainen ZF**, Kepecs A. 2009. Neural representation of behavioral outcomes in the orbitofrontal cortex. *Current Opinion in Neurobiology* **19**: 84–91. DOI: <https://doi.org/10.1016/j.conb.2009.03.010>, PMID: [19427193](https://pubmed.ncbi.nlm.nih.gov/19427193/)
- Mante V**, Sussillo D, Shenoy KV, Newsome WT. 2013. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**: 78–84. DOI: <https://doi.org/10.1038/nature12742>, PMID: [24201281](https://pubmed.ncbi.nlm.nih.gov/24201281/)
- McDannald MA**, Lucantonio F, Burke KA, Niv Y, Schoenbaum G. 2011. Ventral striatum and orbitofrontal cortex are both required for model-based, but not model-free, reinforcement learning. *The Journal of Neuroscience* **31**: 2700–2705. DOI: <https://doi.org/10.1523/JNEUROSCI.5499-10.2011>, PMID: [21325538](https://pubmed.ncbi.nlm.nih.gov/21325538/)
- Miller KJ**, Botvinick MM, Brody CD. 2017. Dorsal hippocampus contributes to model-based planning. *Nature Neuroscience* **20**: 1269–1276. DOI: <https://doi.org/10.1038/nn.4613>, PMID: [28758995](https://pubmed.ncbi.nlm.nih.gov/28758995/)
- Miller KJ**, Botvinick MM, Brody CD. 2018. Value Representations in Orbitofrontal Cortex Drive Learning, but Not Choice. [bioRxiv]. DOI: <https://doi.org/10.1101/245720>
- Mochol G**, Kiani R, Moreno-Bote R. 2021. Prefrontal cortex represents heuristics that shape choice bias and its integration into future behavior. *Current Biology* **31**: 1234–1244. DOI: <https://doi.org/10.1016/j.cub.2021.01.068>, PMID: [33639107](https://pubmed.ncbi.nlm.nih.gov/33639107/)
- Namboodiri VMK**, Otis JM, van Heeswijk K, Voets ES, Alghorazi RA, Rodriguez-Romaguera J, Mihalas S, Stuber GD. 2019. Single-cell activity tracking reveals that orbitofrontal neurons acquire and maintain a long-term memory to guide behavioral adaptation. *Nature Neuroscience* **22**: 1110–1121. DOI: <https://doi.org/10.1038/s41593-019-0408-1>, PMID: [31160741](https://pubmed.ncbi.nlm.nih.gov/31160741/)
- Neiman T**, Loewenstein Y. 2011. Reinforcement learning in professional basketball players. *Nature Communications* **2**: 1–8. DOI: <https://doi.org/10.1038/ncomms1580>, PMID: [22146388](https://pubmed.ncbi.nlm.nih.gov/22146388/)
- Nickerson RS**. 2002. The production and perception of randomness. *Psychological Review* **109**: 330–357. DOI: <https://doi.org/10.1037/0033-295X.109.2.330>
- Nogueira R**, Abolafia JM, Drugowitsch J, Balaguer-Ballester E, Sanchez-Vives MV, Moreno-Bote R. 2017. Lateral orbitofrontal cortex anticipates choices and integrates prior with current information. *Nature Communications* **8**: 14823. DOI: <https://doi.org/10.1038/ncomms14823>, PMID: [28337990](https://pubmed.ncbi.nlm.nih.gov/28337990/)
- Noonan MP**, Walton ME, Behrens TEJ, Sallet J, Buckley MJ, Rushworth MFS. 2010. Separate value comparison and learning mechanisms in macaque medial and lateral orbitofrontal cortex. *PNAS* **107**: 20547–20552. DOI: <https://doi.org/10.1073/pnas.1012246107>
- Padoa-Schioppa C**, Assad JA. 2006. Neurons in the orbitofrontal cortex encode economic value. *Nature* **441**: 223–226. DOI: <https://doi.org/10.1038/nature04676>, PMID: [16633341](https://pubmed.ncbi.nlm.nih.gov/16633341/)
- Padoa-Schioppa C**. 2009. Range-adapting representation of economic value in the orbitofrontal cortex. *The Journal of Neuroscience* **29**: 14004–14014. DOI: <https://doi.org/10.1523/JNEUROSCI.3751-09.2009>, PMID: [19890010](https://pubmed.ncbi.nlm.nih.gov/19890010/)
- Padoa-Schioppa C**. 2011. Neurobiology of economic choice: a good-based model. *Annual Review of Neuroscience* **34**: 333–359. DOI: <https://doi.org/10.1146/annurev-neuro-061010-113648>, PMID: [21456961](https://pubmed.ncbi.nlm.nih.gov/21456961/)
- Padoa-Schioppa C**. 2013. Neuronal origins of choice variability in economic decisions. *Neuron* **80**: 1322–1336. DOI: <https://doi.org/10.1016/j.neuron.2013.09.013>, PMID: [24314733](https://pubmed.ncbi.nlm.nih.gov/24314733/)
- Park IM**, Meister MLR, Huk AC, Pillow JW. 2014. Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nature Neuroscience* **17**: 1395–1403. DOI: <https://doi.org/10.1038/nn.3800>, PMID: [25174005](https://pubmed.ncbi.nlm.nih.gov/25174005/)
- Pickens CL**, Saddoris MP, Gallagher M, Holland PC. 2005. Orbitofrontal lesions impair use of cue-outcome associations in a devaluation task. *Behavioral Neuroscience* **119**: 317–322. DOI: <https://doi.org/10.1037/0735-7044.119.1.317>
- Pillow JW**, Paninski L, Uzzell VJ, Simoncelli EP, Chichilnisky EJ. 2005. Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. *The Journal of Neuroscience* **25**: 11003–11013. DOI: <https://doi.org/10.1523/JNEUROSCI.3305-05.2005>, PMID: [16306413](https://pubmed.ncbi.nlm.nih.gov/16306413/)
- Pillow JW**, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP. 2008. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* **454**: 995–999. DOI: <https://doi.org/10.1038/nature07140>, PMID: [18650810](https://pubmed.ncbi.nlm.nih.gov/18650810/)

- Raposo D**, Kaufman MT, Churchland AK. 2014. A category-free neural population supports evolving demands during decision-making. *Nature Neuroscience* **17**: 1784–1792. DOI: <https://doi.org/10.1038/nn.3865>, PMID: [25383902](https://pubmed.ncbi.nlm.nih.gov/25383902/)
- Rich EL**, Wallis JD. 2016. Decoding subjective decisions from orbitofrontal cortex. *Nature Neuroscience* **19**: 973–980. DOI: <https://doi.org/10.1038/nn.4320>, PMID: [27273768](https://pubmed.ncbi.nlm.nih.gov/27273768/)
- Rigotti M**, Barak O, Warden MR, Wang XJ, Daw ND, Miller EK, Fusi S. 2013. The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**: 585–590. DOI: <https://doi.org/10.1038/nature12160>, PMID: [23685452](https://pubmed.ncbi.nlm.nih.gov/23685452/)
- Rudebeck PH**, Murray EA. 2014. The orbitofrontal oracle: cortical mechanisms for the prediction and evaluation of specific behavioral outcomes. *Neuron* **84**: 1143–1156. DOI: <https://doi.org/10.1016/j.neuron.2014.10.049>, PMID: [25521376](https://pubmed.ncbi.nlm.nih.gov/25521376/)
- Rushworth MFS**, Noonan MP, Boorman ED, Walton ME, Behrens TE. 2011. Frontal cortex and reward-guided learning and decision-making. *Neuron* **70**: 1054–1069. DOI: <https://doi.org/10.1016/j.neuron.2011.05.014>, PMID: [21689594](https://pubmed.ncbi.nlm.nih.gov/21689594/)
- Rustichini A**, Conen KE, Cai X, Padoa-Schioppa C. 2017. Optimal coding and neuronal adaptation in economic decisions. *Nature Communications* **8**: 1–14. DOI: <https://doi.org/10.1038/s41467-017-01373-y>
- Sadacca BF**, Wied HM, Lopatina N, Saini GK, Nemirovsky D, Schoenbaum G. 2018. Orbitofrontal neurons signal sensory associations underlying model-based inference in a sensory preconditioning task. *eLife* **7**: e30373. DOI: <https://doi.org/10.7554/eLife.30373>
- Schoenbaum G**, Chiba AA, Gallagher M. 1998. Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. *Nature Neuroscience* **1**: 155–159. DOI: <https://doi.org/10.1038/407>, PMID: [10195132](https://pubmed.ncbi.nlm.nih.gov/10195132/)
- Simoncelli EP**, Olshausen BA. 2001. Natural image statistics and neural representation. *Annual Review of Neuroscience* **24**: 1193–1216. DOI: <https://doi.org/10.1146/annurev.neuro.24.1.1193>, PMID: [11520932](https://pubmed.ncbi.nlm.nih.gov/11520932/)
- Stalnaker TA**, Cooch NK, McDannald MA, Liu TL, Wied H, Schoenbaum G. 2014. Orbitofrontal neurons infer the value and identity of predicted outcomes. *Nature Communications* **5**: 1–13. DOI: <https://doi.org/10.1038/ncomms4926>
- Steiner AP**, Redish AD. 2012. The road not taken: neural correlates of decision making in orbitofrontal cortex. *Frontiers in Neuroscience* **6**: 131. DOI: <https://doi.org/10.3389/fnins.2012.00131>
- Steiner AP**, Redish AD. 2014. Behavioral and neurophysiological correlates of regret in rat decision-making on a neuroeconomic task. *Nature Neuroscience* **17**: 995–1002. DOI: <https://doi.org/10.1038/nn.3740>, PMID: [24908102](https://pubmed.ncbi.nlm.nih.gov/24908102/)
- Sul JH**, Kim H, Huh N, Lee D, Jung MW. 2010. Distinct roles of rodent orbitofrontal and medial prefrontal cortex in decision making. *Neuron* **66**: 449–460. DOI: <https://doi.org/10.1016/j.neuron.2010.03.033>, PMID: [20471357](https://pubmed.ncbi.nlm.nih.gov/20471357/)
- Takahashi YK**, Roesch MR, Stalnaker TA, Haney RZ, Calu DJ, Taylor AR, Burke KA, Schoenbaum G. 2009. The orbitofrontal cortex and ventral tegmental area are necessary for learning from unexpected outcomes. *Neuron* **62**: 269–280. DOI: <https://doi.org/10.1016/j.neuron.2009.03.005>, PMID: [19409271](https://pubmed.ncbi.nlm.nih.gov/19409271/)
- Takahashi YK**, Chang CY, Lucantonio F, Haney RZ, Berg BA, Yau HJ, Bonci A, Schoenbaum G. 2013. Neural estimates of imagined outcomes in the orbitofrontal cortex drive behavior and learning. *Neuron* **80**: 507–518. DOI: <https://doi.org/10.1016/j.neuron.2013.08.008>, PMID: [24139047](https://pubmed.ncbi.nlm.nih.gov/24139047/)
- Thorpe SJ**, Rolls ET, Maddison S. 1983. The orbitofrontal cortex: neuronal activity in the behaving monkey. *Experimental Brain Research* **49**: 93–115. DOI: <https://doi.org/10.1007/BF00235545>, PMID: [6861938](https://pubmed.ncbi.nlm.nih.gov/6861938/)
- Tibshirani R**, Walther G, Hastie T. 2001. Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society* **63**: 411–423. DOI: <https://doi.org/10.1111/1467-9868.00293>
- Tremblay L**, Schultz W. 1999. Relative reward preference in primate orbitofrontal cortex. *Nature* **398**: 704–708. DOI: <https://doi.org/10.1038/19525>, PMID: [10227292](https://pubmed.ncbi.nlm.nih.gov/10227292/)
- Tymula A**, Glimcher P. 2020. Expected subjective value theory (esvt): A representation of decision under risk and certainty. *SSRN Electronic Journal* **1**: 2783638. DOI: <https://doi.org/10.2139/ssrn.2783638>
- Wallis JD**, Miller EK. 2003. Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *The European Journal of Neuroscience* **18**: 2069–2081. DOI: <https://doi.org/10.1046/j.1460-9568.2003.02922.x>, PMID: [14622240](https://pubmed.ncbi.nlm.nih.gov/14622240/)
- Wallis JD**. 2007. Orbitofrontal cortex and its contribution to decision-making. *Annual Review of Neuroscience* **30**: 31–56. DOI: <https://doi.org/10.1146/annurev.neuro.30.051606.094334>, PMID: [17417936](https://pubmed.ncbi.nlm.nih.gov/17417936/)
- Webb R**, Glimcher PW, Louie K. 2014. Rationalizing context-dependent preferences: divisive normalization and neurobiological constraints on choice. *SSRN Electronic Journal* **10**: 2462895. DOI: <https://doi.org/10.2139/ssrn.2462895>
- West EA**, DesJardin JT, Gale K, Malkova L. 2011. Transient inactivation of orbitofrontal cortex blocks reinforcer devaluation in macaques. *The Journal of Neuroscience* **31**: 15128–15135. DOI: <https://doi.org/10.1523/JNEUROSCI.3295-11.2011>, PMID: [22016546](https://pubmed.ncbi.nlm.nih.gov/22016546/)
- Wilson RC**, Takahashi YK, Schoenbaum G, Niv Y. 2014. Orbitofrontal cortex as a cognitive map of task space. *Neuron* **81**: 267–279. DOI: <https://doi.org/10.1016/j.neuron.2013.11.005>, PMID: [24462094](https://pubmed.ncbi.nlm.nih.gov/24462094/)
- Yagishita S**, Hayashi-Takagi A, Ellis-Davies GCR, Urakubo H, Ishii S, Kasai H. 2014. A critical time window for dopamine actions on the structural plasticity of dendritic spines. *Science* **345**: 1616–1620. DOI: <https://doi.org/10.1126/science.1255514>, PMID: [25258080](https://pubmed.ncbi.nlm.nih.gov/25258080/)
- Yamada H**, Louie K, Tymula A, Glimcher PW. 2018. Free choice shapes normalized value signals in medial orbitofrontal cortex. *Nature Communications* **9**: 1–11. DOI: <https://doi.org/10.1038/s41467-017-02614-w>