



HHS Public Access

Author manuscript

Nat Neurosci. Author manuscript; available in PMC 2016 December 06.

Published in final edited form as:

Nat Neurosci. 2016 July ; 19(7): 973–980. doi:10.1038/nn.4320.

Decoding subjective decisions from orbitofrontal cortex

Erin L. Rich^{1,*} and Jonathan D. Wallis^{1,2}

¹Helen Wills Neuroscience Institute, University of California Berkeley, Berkeley, CA 94720

²Department of Psychology, University of California Berkeley, Berkeley, CA 94720

Abstract

When making a subjective choice, the brain must compute a value for each option and compare those values to make a decision. The orbitofrontal cortex (OFC) is critically involved in this process, but the neural mechanisms remain obscure, in part due to limitations in our ability to measure and control the internal deliberations that can alter the dynamics of the decision process. Here, we tracked the dynamics by recovering temporally precise neural states from multi-dimensional data in OFC. During individual choices, OFC alternated between states associated with the value of two available options, with dynamics that predicted whether a subject would decide quickly or vacillate between the two alternatives. Ensembles of value-encoding neurons contributed to these states, with individual neurons shifting activity patterns as the network evaluated each option. Thus, the mechanism of subjective decision-making involves the dynamic activation of OFC states associated with each choice alternative.

It is believed that the brain makes simple choices by computing a subjective value for each available option then comparing them to arrive at a choice^{1–3}. Behavioral evidence suggests that this comparison involves a dynamic process of rapid deliberation among options^{4, 5}. However, direct neurophysiological signatures of such a process have not been identified. Across species, the OFC plays a critical role in making value-based decisions^{6–9} and evaluating choice alternatives¹⁰, and is therefore a region likely to provide insight into the neural basis of this deliberation process.

A key challenge in measuring the dynamics of subjective decisions is that evaluation and decision-making are unobservable, cognitive processes that can vary significantly with each

Users may view, print, copy, and download text and data-mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use:http://www.nature.com/authors/editorial_policies/license.html#terms

*Correspondence to: erin.rich@berkeley.edu.

Contributions

E.L.R. and J.D.W. designed the experiment and wrote the manuscript. E.L.R. collected and analyzed the data. J.D.W. supervised the project.

Competing Financial Interests

The authors declare no competing financial interests.

Data Availability

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Code Availability

Matlab code for analyses in this study is available from the corresponding author upon request.

A Supplementary Methods Checklist is available.

iteration. For instance, when evaluating two options, A and B, one might first consider A then B. Alternatively, one could consider B then A. In addition, the time taken to evaluate each option may depend on a variety of internal factors¹¹, such as knowledge of outcomes¹², confidence¹³, attention^{4, 14} and motivation¹⁵. Because neuronal responses are inherently stochastic, studies of decision-making typically average activity across repeated trials. However, when decisions vary trial to trial, this approach can obscure critical mechanistic detail¹⁶.

Here we leverage the power of ensemble recording in OFC to analyze data from single choice trials with high temporal resolution. Single unit and local field potential (LFP) data were used to decode patterns of neural activity associated with specific choice options while monkeys made preference-based decisions. During individual choice trials, neural representations alternated between states associated with each available option, as if the network were considering them in turn. These neural states were present both at the ensemble and single neuron level, and their patterns predicted choice behavior. These results indicate that the neural basis of subjective decision-making includes the dynamic representation of choice options in OFC.

Results

Two macaques learned that specific pictures predicted specific rewards, then performed a task in which the pictures were presented either singly or in pairs, the latter of which required the subject to make a choice (Fig. 1). Up to 16 electrodes were acutely placed in OFC areas 11 and 13, with targeting based on previously obtained magnetic resonance images (Supplementary Fig. 1).

OFC represents choice options dynamically

To analyze OFC activity with single-trial resolution, we trained a linear discriminant analysis (LDA) to classify single picture trials by identifying neural patterns associated with four categories of subjective value, ranging from least to most desirable (value = 1 to 4). Single picture trials were ideal training data because they provided a clean measure of neural responses, as there was only one picture presented at a time and all pictures appeared with equal frequency. Information about value was contained in both single neuron activity and LFPs, so both signals were included in the decoder (Supplementary Figs. 3 and 4). The trained classifiers were then used to decode value representations from choice trials in overlapping 20 ms time windows.

The dynamics of the decoder output were variable across trials, consistent with the idea that subjective decisions are not stereotyped, but evolve differently with each iteration. Fig. 2a shows six representative trials, in which the posterior probability of classification in each value category (i.e. $Pr(value_x|observation_i)$ for $x = 1$ to 4) was used as a measure of representational strength of that category. For each trial, we identified the value of the picture that was ultimately chosen by the subject (*chosen*), the alternative option that was not chosen (*unchosen*), and the options that were not available on that particular trial. Note that we designated the chosen or unchosen option post-hoc, based on the subject's choice. The chosen and unchosen options had stronger representations than unavailable options, and they

tended to alternate in dominance, as if the OFC network were transiently encoding each option in turn.

These results were intriguing as they mirror what might intuitively underlie a subjective decision-making process: each option being considered in turn until a decision is made. However, they also raised a number of questions. Critically, can we relate these patterns to the underlying choice process? When a value is decoded during a choice trial, is this dependent on network-level representations? And finally, how do these decoded states relate to single neurons?

OFC representations correspond to available option values

The first evidence that representations decoded from OFC related to the choice process was that the value of pictures in the task predicted which categories were decoded from neural activity. For each 20 ms time bin, decoder classifications were regressed on the picture values on the right and left of the task screen ($n = 4739$ choice trials). Both pictures explained a similar amount of variance in the decoder classifications regardless of their location (Fig. 2b, top), consistent with earlier accounts that value encoding in OFC is not spatially dependent¹⁷. Similarly, classifications were regressed on chosen and unchosen picture values. In this case, the chosen value explained the most variance in the decoded classifications, but the unchosen value also had significant explanatory power between approximately 120 and 250 ms after the pictures appeared (multiple regression $p = 0.001$, Fig. 2b bottom). Therefore, the decoded values reflect both choice options.

Because the LDA classifies every observation even when there is only weak evidence for the categorization, we defined criteria for a stable value representation (which we refer to as a *state* of the network). The criteria required that the same value be decoded from the population for at least 4 consecutive time bins (approximately 35 ms, accounting for bin overlap) with posterior probability > 0.5 . Removing observations with low posterior probabilities preferentially removed transient states (Fig. 2c). Identical analyses using synthesized data sets confirmed that, when thresholded in this way, states were only recovered from input data with clear temporal structure and not from noisy or mixed signals (Supplementary Material and Supplementary Fig. 6). Using these criteria, the median state lasted for 10 time bins (65 ms). On choices between differently valued pictures, chosen states were more common than unchosen states (Fig. 2d). There were a mean of 5.5 stable states per trial (median = 5) that corresponded to the options available in the 800 ms following picture onset, consisting of 3.18 chosen (median = 3) and 2.64 unchosen states (median = 3. $n = 44$ sessions, Wilcoxon rank-sum $z = 14.55$, $p = 6 \times 10^{-48}$). Chosen states were also slightly longer (median = 11 bins, or 70 ms) than unchosen states (median = 9 bins, or 60 ms. Wilcoxon rank-sum $z = 5.40$, $p = 3 \times 10^{-8}$). Confirmatory analyses of the neural feature vectors found that these states correlate with the dominant dimensions of temporal variance in the population during choices (Supplementary Figs. 7 and 8).

Next we quantified how often the neural state was observed to transition from one option value to the other. Overall, there were an average of 3.5 ± 0.07 (95% CI) transitions per trial between states associated with the two available pictures in the 800 ms following choice onset. There were fewer transitions on trials in which the first identified state corresponded

to the chosen item (mean \pm 95% CI = 3.26 ± 0.09), compared to trials where the first identified state corresponded to the unchosen item (3.76 ± 0.1 ; Wilcoxon rank-sum $z = 7.21$, $p = 6 \times 10^{-13}$). Although there were not many error trials (5.8% of trials), there were more transitions between states when subjects made errors (4.0 ± 0.31) than on correct trials (3.47 ± 0.07 ; Wilcoxon rank-sum $z = 3.10$, $p = 0.002$). Finally, on correct trials, higher values of the chosen picture predicted fewer state transitions, while higher values of the unchosen picture predicted more transitions (multiple regression of number of transitions on chosen and unchosen values, significance of β coefficients $p = 7 \times 10^{-15}$ and $p = 2 \times 10^{-24}$ respectively). These results show that we consistently observe transitions between states associated with each option within a trial and the transitions map in a logical way onto the value of the available options and choice behavior.

If states represent the value of available pictures, we should recover fewer states when the two option values are the same, since the states corresponding to each option would be indistinguishable. Indeed, when the choice options had the same value ($n = 684$ trials), there were fewer states (mean = 3.3, median = 3 states) than when choice options had different values ($n = 3782$ excluding error trials, mean = 5.8, median = 6 states. Wilcoxon rank-sum $z = 27.94$, $p = 9 \times 10^{-172}$). This could arise simply because there are two value states consistent with the choice options on different-value trials as opposed to only one on same-value trials. However, the total duration of the recovered states was longer on trials where the options had the same value, compared to those where the options were of different value (mean = 451 ms per trial same value, 389 ms different value. Wilcoxon rank-sum $z = 5.27$, $p = 1 \times 10^{-7}$). This suggests that, on same-value trials, we were missing state transitions between each choice option because the states are indistinguishable to our decoder. Because decoded states could not be assigned to a single picture on trials where both options were the same value, all further analyses excluded these trials, as well as trials in which the lower value option was chosen (error trials).

In contrast to regression analysis of single neuron activity, which showed little evidence for encoding of unchosen options (Supplementary Fig. 2), the decoding analysis captured transient representations related to the value of the unchosen item. However, an alternative explanation is that the decoder returned the value of the chosen item, and noise in this classification was mistakenly attributed to a representation of the value of the unchosen item. For instance, in a choice between values 2 and 3, OFC may only represent 3 with some noise, and since 2 flanks 3, we could recover 2 with some consistency. To test this, we looked at choices where the chosen and unchosen values were not sequential (e.g. 3 versus 1, 4 versus 2). In all cases, the odds of decoding the unchosen value were higher than the odds of decoding other values that were intermediate between the chosen and unchosen values. To quantify the effect, the unchosen and intermediate trial labels were shuffled 200 times to create null distributions of odds ratio differences. P -values for differences observed in the real data were calculated from Gaussians fit to these distributions (2-sided tests). In all comparisons, the odds of decoding the unchosen value became greater than the intermediate value within approximately 200 ms after choice onset (Fig. 3). Therefore, using a decoding approach we were able to recover stable representations of both choice options.

Decoded representations predict choice times

With clear evidence that both options are represented during choice evaluation, we examined whether those representations had consequences for the subjects' choice behavior. We used a multiple regression to determine whether variance in the strength of the chosen and unchosen value representations, as measured by LDA posterior probabilities, predicted the amount of time it took the animal to make a decision (*choice time*). Choice times were defined as the interval between picture onset and the beginning of the fixation that would ultimately be the subject's choice, and were log-transformed to correct for skewed distributions. Faster choice times were predicted by stronger representations of the chosen item, while slower choice times were predicted by stronger representations of the unchosen item (Fig.4a and 4c). This effect was significant from approximately 150 ms to 450 ms after pictures were presented to the subject (multiple regression $p = 0.005$), a time window consistent with previous reports of OFC value encoding^{18, 19} that also overlapped with decision times in the task. Median choice times for subjects M and N were 223 ms and 224 ms, with 74.9% and 79.6% of all choices executed in under 450 ms. The representation of chosen and unchosen items also predicted variability in choice times between trials in which subject chose between the same options (e.g. all choices between value 1 and 3), although effects were smaller, likely due to fewer trials in the comparison and less choice time variability within a condition (Supplementary Fig. 9).

The positive association between the unchosen option and choice times could be interpreted in two ways. It could be an effect specific to representations of the alternative value on offer, so that it competes more with the chosen value. Or it could be that when OFC represents anything other than the chosen option, choice times are slower. To resolve this, an identical analysis was performed that substituted unavailable options for the unchosen option. The chosen option still negatively predicted choice times, but the unavailable options had no predictive value (Fig.4b and 4d). Thus, longer choice times were specifically predicted by representations of the unchosen option.

OFC representations predict deliberation

To gain insight into the subjective difficulty of a choice, we evaluated how long the subject viewed each option before choosing one. In most cases, subjects' eyes moved directly from the central starting point to fixate their chosen item (*quick decisions*). On other trials there was evidence of deliberation; subjects viewed both options a number of times before making a choice (Fig. 5a). Consistent with results in humans⁴, option viewing increased when decisions were more difficult (Fig. 5b). Given this, we analyzed the variability in subjective decision difficulty, as reflected in option viewing behavior. To control for objective decision difficulty, we established two matched groups of trials based on the animals' eye movements ($n = 438$ trials per group). In the first trial group, subjects made quick decisions, and in the second they deliberated. Groups were matched for picture identity and position, as well as the option that was ultimately chosen. The only difference was how much the subject deliberated over the decision.

For quick decisions, representations of the chosen item grew stronger than those of the unchosen item approximately 150 ms after the appearance of the pictures (Fig. 5c,e and

Supplementary Fig. 10a). However in deliberative decisions there was a much smaller difference between the strength of the two representations (Fig. 5d,f). Therefore, differences in the strength of option representations predicted whether subjects made quick decisions or deliberated over their choice. These results could arise from weak representations of both options during deliberation. However, this was not the case. In all conditions, the stimuli on the screen had stronger representations than unavailable options (Fig. 5c–f and Supplementary Fig. 10a). Furthermore, the differences between quick and deliberative decisions were found on identical trials and cannot be accounted for by objective decision difficulty. Instead, they must reflect the internal process of making a subjective preference decision.

We considered two additional hypotheses of how the representational strength of each option may change as a decision evolves. First, we asked whether, on deliberation trials, chosen representations slowly strengthen until the choice is made. That is, given that the choice occurs at a later and variable time when the subject deliberates, would we see the chosen representations dominating if trials were aligned to the time of the choice rather than picture onset? Contrary to expectations, chosen representations were not strong at the time of choice on deliberative decisions (Supplementary Fig. 10b). Thus, while the relative strength of each representation in OFC predicted how decisively the subject chose, it was not necessary for OFC to be encoding the value of the chosen option at the time of the decision.

Second, we examined how the strength of OFC representations related to eye movements. We focused on deliberative trials, defined as those where there was more than one saccade to each option before the choice was made. We looked at the average representational strength of chosen and unchosen options in a window of time ± 200 ms from the time the subject foveated a picture for the first two image fixations. There were no clear differences between chosen and unchosen representations when the subject viewed either image (Fig. 6), although there was a slight increase in strength of both representations as viewing began. Next, we focused on those trials where the subject looked at one option but chose the other. The chosen representation steadily gained in strength as the trial progressed and diverged from the unchosen representation. However, there was again no evidence that the option being foveated was reflected in OFC either before or after viewing began (Fig. 6b).

Decoded states do not depend on any single neuron

Network activity in OFC alternates between states corresponding to each choice option and we have shown that the strength of these states predicts choice behavior. A critical question is whether these states are a property of the neural features in the LDA or a result of a single neuron dominating the decoder classifications. We addressed this in two ways, both of which indicated that states do not depend on any single neuron.

First, if decoded states depended on one or a few neurons, removing one key feature from the decoder should drastically change the time series of decoded states. To assess this, one neuron was removed from both the training (single picture) trials and the choice trials, the LDA was retrained, and observations from the choice trials were re-classified. This procedure was repeated for each single neuron in each session. On individual trials, removal of any one neuron resulted in very little change to the decoded states (Fig. 7a).

To quantify across neurons and sessions, classifications from a decoder trained on the full neuronal ensemble were compared to those from a decoder trained with one neuron left out. Time series of decoder outputs were concatenated across trials and compared with Pearson correlations. There was a strong correlation between outputs derived from the full decoder, and the outputs from the decoder with one neuron removed (Fig. 7b, median $r^2 = 0.89$). Only 2 neurons reduced the correlation to 0.3 or less (Fig. 7b). Note that when we correlated the output of the full decoder with a shuffled decoder, whereby all neural features were shuffled independently with respect to behavioral conditions, r^2 values were always less than 0.1 (1000 shuffles: median $r^2 = 0.0014$, maximum = 0.089). Thus, there was no instance where the removal of one neuron reduced the correlation to the level of shuffled data. Therefore, across 44 behavioral sessions, decoded representations were unaffected by the removal of any single neuron.

Given that removing a single neuron did not markedly alter decoded states, we next asked how many neurons could be removed, on average, before reducing the similarity of classifications to chance levels. To do this, we started by removing each neuron in turn and calculating the correlations as above. Then we removed every possible pair of neurons, then triplets, and so on. If the number of possible combinations of neurons to remove exceeded 200, 200 combinations were selected at random to compute that data point. As expected, the classification similarity decreased with the exclusion of more neurons (Fig. 7c). After removing all single neurons and classifying solely on LFP features, correlations continued to be above chance levels, as defined by the trial-shuffled data above. This emphasizes that the choice-related signals we are detecting are robust and highly distributed in the OFC.

Single neurons dynamically represent choice options

Next, we asked how individual neurons encoded the states associated with different choice options. We considered three hypotheses. First, the recovered states could be exclusively a distributed network property, not discernable in single neurons. Second, the states could manifest as a unitary population of value-encoding neurons that changed firing rates as the network shifted from one state to another, thereby encoding both options. In this case, we would expect single neurons to encode the value of option A when the decoded state is A, and the value of option B when the decoded state is B. Finally, there may be separate populations of neurons encoding each option. For instance, when the network is in state A, a group of neurons may encode the value of option A and become inactive or non-selective when the network is in state B. Similarly, a different group of neurons may encode option B but not A. To address these hypotheses, firing rates from a held-out neuron, that did not contribute to the decoder, were aligned to the onset of states defined by the decoder, and separately analyzed for value encoding when the current state corresponded to the value of the picture on the right or left of the task screen.

An example neuron is shown in Fig. 8a. This neuron encoded the value of the left picture, but only when the value decoded from the network corresponded to the left picture. When the value of the right picture was decoded, the neuron switched to encode the value of the right picture and not the left. Note that, in the second and third frames of this plot, the neuron showed no encoding during the first 100 ms of the target state, but did encode value

prior to state onset. This reflects encoding from the prior state. For example, the neuron does not encode the value of the left picture when the network transitions to the right state (second panel), but immediately prior to that transition, the network will often be in the left state and so encoding of the value of the left picture is evident prior to the transition. Therefore, this neuron encoded the value of both right and left pictures, but only when the state decoded from the rest of OFC corresponded to that picture.

We observed this phenomenon in a large proportion of OFC neurons. To quantify the effect, firing rates from the held-out neuron were aligned to the onset of decoded states. For each neuron in turn, we regressed firing rates (bins of 50 ms, stepped forward by 10 ms) on two predictors: the value of the picture associated with the current state (e.g. the value of the left picture when the network state corresponds to the left picture) and the value of the alternate picture (e.g. the value of the right picture when the network corresponds to the left picture). After the onset of a state, more neurons encoded the value of that state than the alternate option (Fig. 8b). We next regressed firing rates against four different regressors, corresponding to each combination of state and picture value. There was a strong correlation between the magnitude of each neuron's encoding of left picture values during left states and right picture values during right state, as measured by the β coefficient of each regressor ($r^2 = 0.55$, $p = 6 \times 10^{-81}$, Fig. 8c). These results show that there is a single population of value-encoding neurons in OFC that dynamically shifts from encoding the value of one choice option to the other.

Discussion

Using a novel decoding approach, we showed that representations in OFC alternate between distinct neural states corresponding to the value of different choice options. This coding, instantiated in ensembles of neural features as well as single neurons, was transient and temporally variable from one trial to the next. Because of this variability, these states were not observed with standard regression analyses, and could only be decoded from simultaneously acquired neural signals. This approach enabled us to see the process of an internal subjective decision unfold during single choice trials. The relationship between recovered neural states and behavioral variability in the task support the notion that these dynamic representations of choice alternatives are a critical feature of OFC's role making in value-based decisions.

Decoding approaches recover unobservable neural processes

The neuropsychological evidence strongly supports a critical role for OFC in value-based decision-making^{7, 9}. The most prominent feature of neural coding in OFC during decision tasks, including the current one, tends to be the value of the chosen option^{20, 21}. This is often used as evidence to support the role of OFC in value-based decision-making, yet such representations can only be defined post-hoc, by the choice the animal made. If other options are not represented it is unclear how comparison between them could take place in OFC. A recent study²² used an artificial neural network to model decision-making and found that inhibitory interneurons, included in the model for biological plausibility, encoded chosen value, and choices involved a dynamical competition between the excitatory drive from offer

value neurons, and inhibition from these chosen value neurons. Chosen value was not simply the output of the network, but was integral to the process of offer comparison. However, there was no role for neurons encoding the unchosen value. In contrast to this model, we were able to decode representations of both chosen and unchosen values, and show independent effects of each on choice behavior.

Other theoretical models of value-based choice have argued that it involves a competition between neural representations of different options^{23, 24}. Such ideas evolved from models of perceptual decision-making, in which separate populations of neurons are excited by different options and compete with each other through mutual inhibition^{25, 26}. Here we demonstrate that, rather than competing groups of OFC neurons, subjective decision-making involves the OFC network transitioning through multiple states, dynamically representing the value of both chosen and unchosen options.

In contrast to our main results, regression analyses found no significant encoding of unchosen option values, consistent with previous reports^{20, 21}. This raises the question of how we are able to decode information about both options. This is likely due to two advantages afforded our approach relative to more traditional approaches that analyze one neuron at a time. First, decoding is based on multi-dimensional data and can recover information that is not strongly present in any single dimension. Second, our recovered value representations are dynamic, with the chosen item being more common and the unchosen being relatively fleeting and not time-aligned to any externally observable event. In this case, the unchosen representations could be lost when averaging over time or across trials.

Decoding information from multidimensional neural features has been used for some time to understand the neural basis of motor behavior²⁷ and sensory processing²⁸, and for controlling neural prosthetics²⁹. It is a less common approach to understanding cognitive function, but has provided important insights into processes such as memory³⁰ and attention³¹. For example, decoding spike trains from rat hippocampal neurons found transient activity at decision points in a spatial maze task that mapped onto potential future trajectories³². As in the current results, neural activity represented choice options sequentially, rather than simultaneously. Similar flickering back and forth of competing hippocampal representations occurs as an animal transitions between different environments³³. These investigations emphasize that decoding is able to reveal neural computations that are transient and temporally variable, and in some cases impossible to control or observe externally.

Internal processing involved in preference decisions

Decoded states not only corresponded to available options, but they also predicted the time it took to make a choice and the amount of deliberation involved. Despite this, the relationship between overt behavior and OFC activity was only indirect. There was no evidence that OFC representations needed to build to some threshold in order for a decision to occur, and aligning neural activity with eye movements did not reveal any systematic changes in the neural representation as gaze shifted. In psychophysical models, gaze direction can provide insights into the choice process. For example the longer one looks at an alternative item, the longer one must look at the chosen item before it is selected, and subjects generally choose

the item they looked at last⁴. However, the extent to which one can infer the neural mechanisms of decision processes from eye movements is unclear because eye movements can only reflect computations that reach the level of motor output. There are likely additional neural processes not included in psychophysical models based solely on eye movements, such as working memory. A subject could be looking at option A, but considering the properties of option B, held in working memory. The fact that there is not a simple relationship between OFC activity and eye movements underscores the importance of decoding the decision-making process from neural activity in those areas that are putatively involved in covert cognitive processing, rather than inferring it from overt behavior³⁴. Activity in OFC may contribute to the comparison process while downstream areas are responsible for using this to guide motor output and determine the final choice response³⁵. Indeed, recent analysis of the dynamical interactions between prefrontal regions suggest a role for dorsolateral prefrontal cortex in this process³⁶.

Role of OFC representations in decision-making

Current hypotheses regarding the contribution of OFC to decision-making suggest a role in model-based learning³⁷, predicting imagined outcomes^{38, 39}, or most generally, forming a “cognitive map” that relates aspects of the current context or environment to the task at hand⁴⁰. All of these views emphasize a forward-looking, or predictive, function of OFC. As such, we could interpret the neural states corresponding to different choice options as transient activations that anticipate, or ‘imagine’^{38, 39}, outcomes that are likely to ensue. In the present task, these map onto subjective values, whereas in other tasks other features may be critical^{20, 41}.

Recent findings suggest that a small number of OFC neurons are modulated by the spatial location of an offered reward^{42, 43}, while others report almost no spatial tuning¹⁷. Here, we found that the location of the picture did not affect our ability to decode its value. While it may be that a small percentage of neurons with spatial biases were present, it did not impact population level value decoding in this task. Overall, it is likely that OFC is not inherently biased to encode either spatial or object information, but rather the extent to which it encodes either will depend on how relevant it is for optimal decision-making²⁰. We can contrast this with recent results from dorsolateral prefrontal cortex, which seems biased to encode the spatiotemporal organization of behavior, even when that information is not directly relevant to the task⁴⁴. Interestingly, a similar viewpoint has been reached by contrasting OFC with the hippocampus. OFC tends to parse information according to task relevance, while the hippocampus does so according to spatiotemporal organization of behavior^{45, 46}.

OFC has strong connectivity with subcortical areas involved in reward processing as well as sensory processing cortical regions^{47, 48}, so the transient value representations are potentially linked to widely distributed networks involved in outcome prediction. Cortical microcircuit models have recently been developed that produce network transitions on a similar timescale to those we observed^{49, 50}. These models use Markov chain Monte Carlo sampling to implement a process of probabilistic inference. Extending these ideas to OFC,

would imply that OFC networks are engaged in a process of attempting to infer the optimal choice given the available offers.

Methods

Subjects and task

All procedures were in accord with the National Institute of Health guidelines and recommendations of the University of California at Berkeley Animal Care and Use Committee. Subjects were two male rhesus macaques (*Macaca mulatta*), aged 7 and 9 years, weighing 14 and 9 kg at the time of recording. One additional subject began training on this task but was excluded prior to completion for poor behavioral performance and inability to complete sufficient trials per behavioral session. Subjects sat in a primate chair, viewed a computer screen and manipulated a bidirectional joystick fixed to the front of the chair. Stimulus presentation and behavioral contingencies were controlled using MonkeyLogic software⁵¹. Eye movements were tracked with an infra-red camera (ISCAN, Woburn, MA).

Reward-predicting stimuli consisted of 8 unique images of natural scenes, sized approximately $2^\circ \times 3^\circ$ of visual angle. Stimuli were selected randomly for presentation. For reasons unrelated to the current report, rewards consisted of either fruit juice or a conditioned reinforcer, represented by a blue reward bar visible on the task screen. Prior to task training, subjects were conditioned to associate the length of the reward bar with a proportional amount of fruit juice obtained at the end of a block of four completed trials. Subjects learned this association well and valued pictures predicting larger gains in reward bar over those predicting smaller gains. Subjects M and N chose pictures that predicted larger over smaller gains on 91% and 97% of choices respectively.

Four stimuli predicted juice reward of different amounts (0.05, 0.10, 0.18 and 0.30 ml), and four stimuli predicted an increase in the length of the reward bar (i.e. secondary reinforcement) of different amounts. For simplicity, we refer to these as ordinal values (1 = smallest reward; 4 = largest reward). Actual outcomes associated with each picture were probabilistic, in that on average, 4/7 trials (~57%) delivered the chosen reward type and quantity. On 1/7 trials (~14%), the chosen reward type was delivered, but the value was one of the other three values. On 1/7 trials, the chosen reward value was delivered, but the reward type was the opposite of the chosen type, and on 1/7 trials, both reward value and reward type were inconsistent with the chosen picture.

Amounts of juice and secondary reinforcer were titrated so that subjects had approximately equal preferences between reward types of the same ordinal value. For example, subjects consistently chose pictures of ordinal value 3 over 1 or 2, regardless of whether each of the pictures was associated with juice or reward bar as an outcome. Subject M and N chose rewards with higher ordinal value on 92% and 95% of choices respectively, and on 91% and 95% of choices in which they were offered one picture that predicted primary reward and one that predicted secondary reward.

Neurophysiological recording

After initial task training, subjects were implanted with head positioners and titanium chambers oriented over bilateral frontal lobes. Up to 16 electrodes were acutely lowered to OFC on each recording day through bilateral craniotomies following methods described in detail elsewhere⁵². Electrode placement was in OFC areas 11 and 13, with targeting based on previously obtained MR images of each subject and mapping of gray and white matter boundaries during electrode placement. The recorded neurons represent a random selection from the target regions, in that neurons were not pre-screened for selective responses. All well-isolated neurons in the target region were recorded and included in the analyses. Neural signals were acquired with a Plexon MAP system (Plexon, Dallas, TX). Because trial conditions were randomized during recording, data collection was not explicitly blinded.

Sessions in which subjects performed < 300 trials were not included (3 subject M, 1 subject N), as sufficient sampling of single-picture trials was required for decoding analyses. Analyses included a total of 451 neurons (259 subject M, 192 subject N) and 455 LFP channels (251 subject M, 204 subject N) recorded over the course of 44 sessions (24 subject M, 20 subject N). On average, simultaneously recorded ensembles included 10 neurons (minimum = 4, maximum = 21) and 10 LFP channels (minimum = 3, maximum = 14).

Preprocessing of neural signals

LDAs were trained and tested on multidimensional data, consisting of action potentials from isolated single units and simultaneously recorded LFPs, acquired within single day recording sessions. Spiking data was transformed to a time series with 1 kHz resolution, where the presence of a spike was indicated by 1, and absence 0. For LFPs, raw field potentials acquired at 1 kHz were first evaluated visually and any channels in which the signal voltage was clipped by the acquisition system were removed from further analysis. The remaining channels were notch filtered at 60, 120 and 180 Hz and band-passed using a finite impulse response filter in 6 frequency bands: delta (2–4 Hz), theta (4–8 Hz), alpha (8–12 Hz), beta (12–30 Hz), gamma (30–60 Hz), high-gamma (70–200 Hz). Analytic amplitudes were obtained from Hilbert transforms of the pass bands and z-scored. All time series (spikes and pass band amplitudes) were smoothed by a 50 ms boxcar, and aligned to the appearance of the picture(s) on each trial.

Statistics

All statistical analyses were performed with Matlab (MathWorks, Natick, MA). Details of the specific analyses are provided below. Except in the case of reaction and choice times, which were log-transformed, data were assumed to be normal, though this was not formally tested. Non-parametric tests and corrections for multiple comparisons were performed where appropriate, and all statistical tests were two-sided. No statistical methods were used to predetermine our sample sizes.

Analysis of choice information encoded by OFC neurons

To determine what task information was encoded by OFC neurons, individual neuron firing rates were analyzed with a multiple regression. Firing rate was averaged across two epochs: the 500 ms immediately preceding (fixation) or following (choice) picture onset. Six

regressors were used to predict firing rate: value of the chosen (Ch) and unchosen (UnCh) items, reward type for chosen and unchosen items, the trial number within blocks of 4 (defined by cash-in of the secondary reinforcer), and the size of the reward bar representing the earned secondary reinforcer (Supplementary Fig. 2).

Decoding value from OFC neural activity

LDA (*classify* function in Matlab Statistical Toolbox) was used to classify patterns of neural activity associated with each of the four value levels. LDA assumes that different categories of data consist of Gaussian distributions, and attempts to find the optimal weightings for all features in the data (in this case, all neural signals) so that a hyperplane optimally separates each pair of categories. Our model included no priors since each category was equally likely. The classification of a new data point was then based on which side of the hyperplane that point occurs. LDA was chosen because it is robust, commonly used and computationally efficient. Further, it has been shown to perform well in decoding information from neural data^{53–56}.

The accuracy of the LDA was tested on single picture trials using different input data and leave-one-trial-out cross-validation. For neuron spiking data, the activity of each neuron was averaged in time windows of 80 ms, stepped forward by 20ms. For LFPs, analytic amplitudes were averaged in the same time windows for each of the 6 frequency bands defined above. Neurons and LFPs recorded on each testing day were analyzed separately as unique ensembles, and reported results are the average across 44 sessions. In each time window, neural data were used to decode picture identity, where pictures 1 to 4 were those predicting primary reward of smallest to largest amounts, and pictures 5 to 8 were those predicting secondary reward of smallest to largest amounts. Confusion matrices of decoder classifications were constructed for the time of peak decoding. Since these analyses found that both spiking and LFP decoders identified reward types unreliably, we focused the remaining analyses on decoding reward value during choice trials.

To decode value information during choice trials, we used the same LDA to classify the 8 pictures into 4 value categories. The classifier was trained on data from single picture trials, averaged in a time window that coincided with peak decoding of single picture values (200 ms epoch beginning 100 ms after picture onset). We used these classifiers to decode value representations from choice trials in time windows of 20 ms, stepped forward by 5 ms over an epoch from 600 ms before to 800 ms after the appearance of the choice options.

Posterior probabilities, taken as an output of the MATLAB function, estimated the probability of each value category in the trained decoder, given an observation of neural data. We normalized these probabilities so that the sum over all posterior probabilities was 1. The benefit of using the posterior probability measure, instead of the classification of the observation, is that it includes more quantitative information about how well the observation was classified. For example, a point sitting far from the hyperplane and near the center of the distribution of a class in the training data would have a high posterior probability associated with that category. In contrast, another observation with the same classification, but sitting close to the hyperplane, and farther from the center of the distribution would have a lower posterior probability associated with that category.

Comparison between categorical and linear decoding

We trained a general linear model to predict picture values (1 to 4) from all neural features on single picture trials. The resulting weights were then applied to the same features during individual trials, and combining the weighted features gave us a continuous prediction of value. To compare the LDA and linear model, we assessed their ability to decode data from single picture trials. Probability density functions were estimated from the mean and standard deviation of each distribution of decoded values, and these were used to obtain the probability that each trial belonged to each category. The trial was assigned to the category with the highest probability.

Analysis of synthetic data sets

To determine whether the LDA would return state fluctuations when they do not exist, we created synthetic data sets with known temporal structures, derived from real neural data recorded on single picture trials. Each synthetic data set consisted a matrix of a trials \times time \times neural features. Single picture trials from each session were divided into four distributions, corresponding to the four picture values in the task, for each neural feature. For each trial in a synthetic data set, a time course of neural activity was constructed by randomly sampling from the relevant distribution. For example, if the choice was between value 4 and value 1, at each time point when the synthetic data was determined to represent 4, samples would be drawn from the 4 distribution. This procedure was performed independently for each neural feature and each time point. Thus, we ensured that the synthesized data matched the actual neural data in terms of the number of selective and non-selective neural features, the degree of noise in the observations and the overlap between different categories. However, the time course over which the values were represented was artificially constructed (see Supplementary Fig. 6). Trials from each recorded session were used, so that the synthetic sets also reflected the across-session variability. The number of trials in each set was the same as the number of choice trials, and the option values were the same as the real choice options, excluding choices between pictures of the same value. The time dimension was 800 data points, corresponding to the 800 ms after choice onset. The input data consisted of single neuron firing rates and amplitudes of the frequency-decomposed LFP, as in the real data set. Synthesized data were submitted to the same decoding analyses described above.

Population vector analysis

To confirm the existence of discrete neural states, we examined population vectors that comprise the inputs to the LDA to determine whether they show evidence of similar neural states. Given the nature of OFC responses to value, averaged population vectors are not informative because some neurons have higher firing rates for high value items, while others have high firing rates for low value items, and averaging them together cancels out these effects. Therefore, we examined multi-dimensional population vectors, using principal component analysis (PCA) for dimensionality reduction. For each session, multi-dimensional population vectors were created by concatenating choice trials, where each neuron or LFP input was a feature measured repeatedly over time. PCA identified orthogonal dimensions accounting for the most temporal (i.e. across-trial) variance. The data were then projected onto these principal components (PCs) for analysis.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank M. Rangel-Gomez for comments on the manuscript. This work was funded by NIDA R01 DA19028 and NIMH R01 MH097990 J. D. Wallis, and by the Hilda and Preston Davis Foundation and NIDA K08 DA039051 to E. L. Rich. This research was partially funded by the Defense Advanced Research Projects Agency (DARPA) under Cooperative Agreement Number W911NF-14-2-0043, issued by the Army Research Office contracting office in support of DARPA'S SUBNETS program. The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

References

1. Padoa-Schioppa C. Neurobiology of economic choice: a good-based model. *Annu Rev Neurosci.* 2011; 34:333–359. [PubMed: 21456961]
2. Rangel A, Camerer C, Montague PR. A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience.* 2008; 9:545–556. [PubMed: 18545266]
3. Rushworth MF, Mars RB, Summerfield C. General mechanisms for making decisions? *Curr Opin Neurobiol.* 2009; 19:75–83. [PubMed: 19349160]
4. Krajbich I, Armel C, Rangel A. Visual fixations and the computation and comparison of value in simple choice. *Nat Neurosci.* 2010; 13:1292–1298. [PubMed: 20835253]
5. Dai J, Busemeyer JR. A probabilistic, dynamic, and attribute-wise model of intertemporal choice. *J Exp Psychol Gen.* 2014; 143:1489–1514. [PubMed: 24635188]
6. Wallis JD. Cross-species studies of orbitofrontal cortex and value-based decision-making. *Nat Neurosci.* 2012; 15:13–19. [PubMed: 22101646]
7. Rudebeck PH, Murray EA. The orbitofrontal oracle: cortical mechanisms for the prediction and evaluation of specific behavioral outcomes. *Neuron.* 2014; 84:1143–1156. [PubMed: 25521376]
8. Jones JL, et al. Orbitofrontal cortex supports behavior and learning using inferred but not cached values. *Science.* 2012; 338:953–956. [PubMed: 23162000]
9. Fellows LK. Orbitofrontal contributions to value-based decision making: evidence from humans with frontal lobe damage. *Ann N Y Acad Sci.* 2011; 1239:51–58. [PubMed: 22145875]
10. Padoa-Schioppa C. Neuronal origins of choice variability in economic decisions. *Neuron.* 2013; 80:1322–1336. [PubMed: 24314733]
11. Doya K. Modulators of decision making. *Nat Neurosci.* 2008; 11:410–416. [PubMed: 18368048]
12. Sugrue LP, Corrado GS, Newsome WT. Matching behavior and the representation of value in the parietal cortex. *Science.* 2004; 304:1782–1787. [PubMed: 15205529]
13. Lak A, et al. Orbitofrontal cortex is required for optimal waiting based on decision confidence. *Neuron.* 2014; 84:190–201. [PubMed: 25242219]
14. Shimojo S, Simion C, Shimojo E, Scheier C. Gaze bias both reflects and influences preference. *Nat Neurosci.* 2003; 6:1317–1322. [PubMed: 14608360]
15. Bouret S, Richmond BJ. Ventromedial and orbital prefrontal neurons differentially encode internally and externally driven motivational values in monkeys. *J Neurosci.* 2010; 30:8591–8601. [PubMed: 20573905]
16. Churchland MM, Yu BM, Sahani M, Shenoy KV. Techniques for extracting single-trial activity patterns from large-scale neural recordings. *Curr Opin Neurobiol.* 2007; 17:609–618. [PubMed: 18093826]
17. Grattan LE, Glimcher PW. Absence of spatial tuning in the orbitofrontal cortex. *PLoS One.* 2014; 9:e112750. [PubMed: 25386837]
18. Rich EL, Wallis JD. Medial-lateral organization of the orbitofrontal cortex. *J Cogn Neurosci.* 2014; 26:1347–1362. [PubMed: 24405106]

19. Wallis JD, Miller EK. Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *Eur J Neurosci.* 2003; 18:2069–2081. [PubMed: 14622240]
20. Hosokawa T, Kennerley SW, Sloan J, Wallis JD. Single-neuron mechanisms underlying cost-benefit analysis in frontal cortex. *J Neurosci.* 2013; 33:17385–17397. [PubMed: 24174671]
21. Padoa-Schioppa C, Assad JA. Neurons in the orbitofrontal cortex encode economic value. *Nature.* 2006; 441:223–226. [PubMed: 16633341]
22. Rustichini A, Padoa-Schioppa C. A neuro-computational model of economic decisions. *J Neurophysiol.* 2015; 114:1382–1398. [PubMed: 26063776]
23. Hunt LT, et al. Mechanisms underlying cortical activity during value-guided choice. *Nat Neurosci.* 2012; 15:470–476. S471–S473. [PubMed: 22231429]
24. Jocham G, Hunt LT, Near J, Behrens TE. A mechanism for value-guided choice based on the excitation-inhibition balance in prefrontal cortex. *Nat Neurosci.* 2012; 15:960–961. [PubMed: 22706268]
25. Wang XJ. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron.* 2002; 36:955–968. [PubMed: 12467598]
26. Usher M, McClelland JL. The time course of perceptual choice: the leaky, competing accumulator model. *Psychol Rev.* 2001; 108:550–592. [PubMed: 11488378]
27. Georgopoulos AP, Schwartz AB, Kettner RE. Neuronal population coding of movement direction. *Science.* 1986; 233:1416–1419. [PubMed: 3749885]
28. Fukushima M, Saunders RC, Leopold DA, Mishkin M, Averbach BB. Differential coding of conspecific vocalizations in the ventral auditory cortical stream. *J Neurosci.* 2014; 34:4665–4676. [PubMed: 24672012]
29. Willett FR, Suminski AJ, Fagg AH, Hatsopoulos NG. Improving brain-machine interface performance by decoding intended future movements. *J Neural Eng.* 2013; 10:026011. [PubMed: 23428966]
30. Lee AK, Wilson MA. Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron.* 2002; 36:1183–1194. [PubMed: 12495631]
31. Tremblay S, Pieper F, Sachs A, Martinez-Trujillo J. Attentional filtering of visual information by neuronal ensembles in the primate lateral prefrontal cortex. *Neuron.* 2015; 85:202–215. [PubMed: 25500502]
32. Johnson A, Redish AD. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J Neurosci.* 2007; 27:12176–12189. [PubMed: 17989284]
33. Jezek K, Henriksen EJ, Treves A, Moser EI, Moser MB. Theta-paced flickering between place-cell maps in the hippocampus. *Nature.* 2011; 478:246–249. [PubMed: 21964339]
34. Seidemann E, Meilijson I, Abeles M, Bergman H, Vaadia E. Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *J Neurosci.* 1996; 16:752–768. [PubMed: 8551358]
35. Cai X, Padoa-Schioppa C. Contributions of orbitofrontal and lateral prefrontal cortices to economic choice and the good-to-action transformation. *Neuron.* 2014; 81:1140–1151. [PubMed: 24529981]
36. Hunt LT, Behrens TE, Hosokawa T, Wallis JD, Kennerley SW. Capturing the temporal evolution of choice across prefrontal cortex. *Elife.* 2015; 4
37. McDannald MA, et al. Model-based learning and the contribution of the orbitofrontal cortex to the model-free world. *Eur J Neurosci.* 2012; 35:991–996. [PubMed: 22487030]
38. Barron HC, Dolan RJ, Behrens TE. Online evaluation of novel choices by simultaneous representation of multiple memories. *Nat Neurosci.* 2013; 16:1492–1498. [PubMed: 24013592]
39. Takahashi YK, et al. Neural estimates of imagined outcomes in the orbitofrontal cortex drive behavior and learning. *Neuron.* 2013; 80:507–518. [PubMed: 24139047]
40. Wilson RC, Takahashi YK, Schoenbaum G, Niv Y. Orbitofrontal cortex as a cognitive map of task space. *Neuron.* 2014; 81:267–279. [PubMed: 24462094]
41. Schoenbaum G, Takahashi Y, Liu TL, McDannald MA. Does the orbitofrontal cortex signal value? *Ann N Y Acad Sci.* 2011; 1239:87–99. [PubMed: 22145878]

42. Luk CH, Wallis JD. Choice coding in frontal cortex during stimulus-guided or action-guided decision-making. *J Neurosci*. 2013; 33:1864–1871. [PubMed: 23365226]
43. Strait CE, et al. Neuronal selectivity for spatial position of offers and choices in five reward regions. *J Neurophysiol*. 2015 jn.00325.02015.
44. Lara AH, Wallis JD. Executive control processes underlying multi-item working memory. *Nat Neurosci*. 2014; 17:876–883. [PubMed: 24747574]
45. Farovik A, et al. Orbitofrontal cortex encodes memories within value-based schemas and represents contexts that guide memory retrieval. *J Neurosci*. 2015; 35:8333–8344. [PubMed: 26019346]
46. McKenzie S, et al. Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. *Neuron*. 2014; 83:202–215. [PubMed: 24910078]
47. Carmichael ST, Price JL. Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. *J Comp Neurol*. 1995; 363:615–641. [PubMed: 8847421]
48. Carmichael ST, Price JL. Sensory and premotor connections of the orbital and medial prefrontal cortex of macaque monkeys. *J Comp Neurol*. 1995; 363:642–664. [PubMed: 8847422]
49. Buesing L, Bill J, Nessler B, Maass W. Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons. *PLoS Comput Biol*. 2011; 7:e1002211. [PubMed: 22096452]
50. Habenschuss S, Jonke Z, Maass W. Stochastic computations in cortical microcircuit models. *PLoS Comput Biol*. 2013; 9:e1003311. [PubMed: 24244126]
51. Asaad WF, Eskandar EN. A flexible software tool for temporally-precise behavioral control in Matlab. *J Neurosci Methods*. 2008; 174:245–258. [PubMed: 18706928]
52. Lara AH, Kennerley SW, Wallis JD. Encoding of gustatory working memory by orbitofrontal neurons. *J Neurosci*. 2009; 29:765–774. [PubMed: 19158302]
53. Boulay CB, Pieper F, Leavitt M, Martinez-Trujillo J, Sachs AJ. Single-trial decoding of intended eye movement goals from lateral prefrontal cortex neural ensembles. *J Neurophysiol*. 2015 jn.00788.02015.
54. Averbeck BB, Crowe DA, Chafee MV, Georgopoulos AP. Neural activity in prefrontal cortex during copying geometrical shapes. II. Decoding shape segments from neural ensembles. *Exp Brain Res*. 2003; 150:142–153. [PubMed: 12669171]
55. Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA. Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci*. 2002; 5:805–811. [PubMed: 12134152]
56. Crowe DA, et al. Prefrontal neurons transmit signals to parietal neurons that reflect executive control of cognition. *Nat Neurosci*. 2013; 16:1484–1491. [PubMed: 23995071]

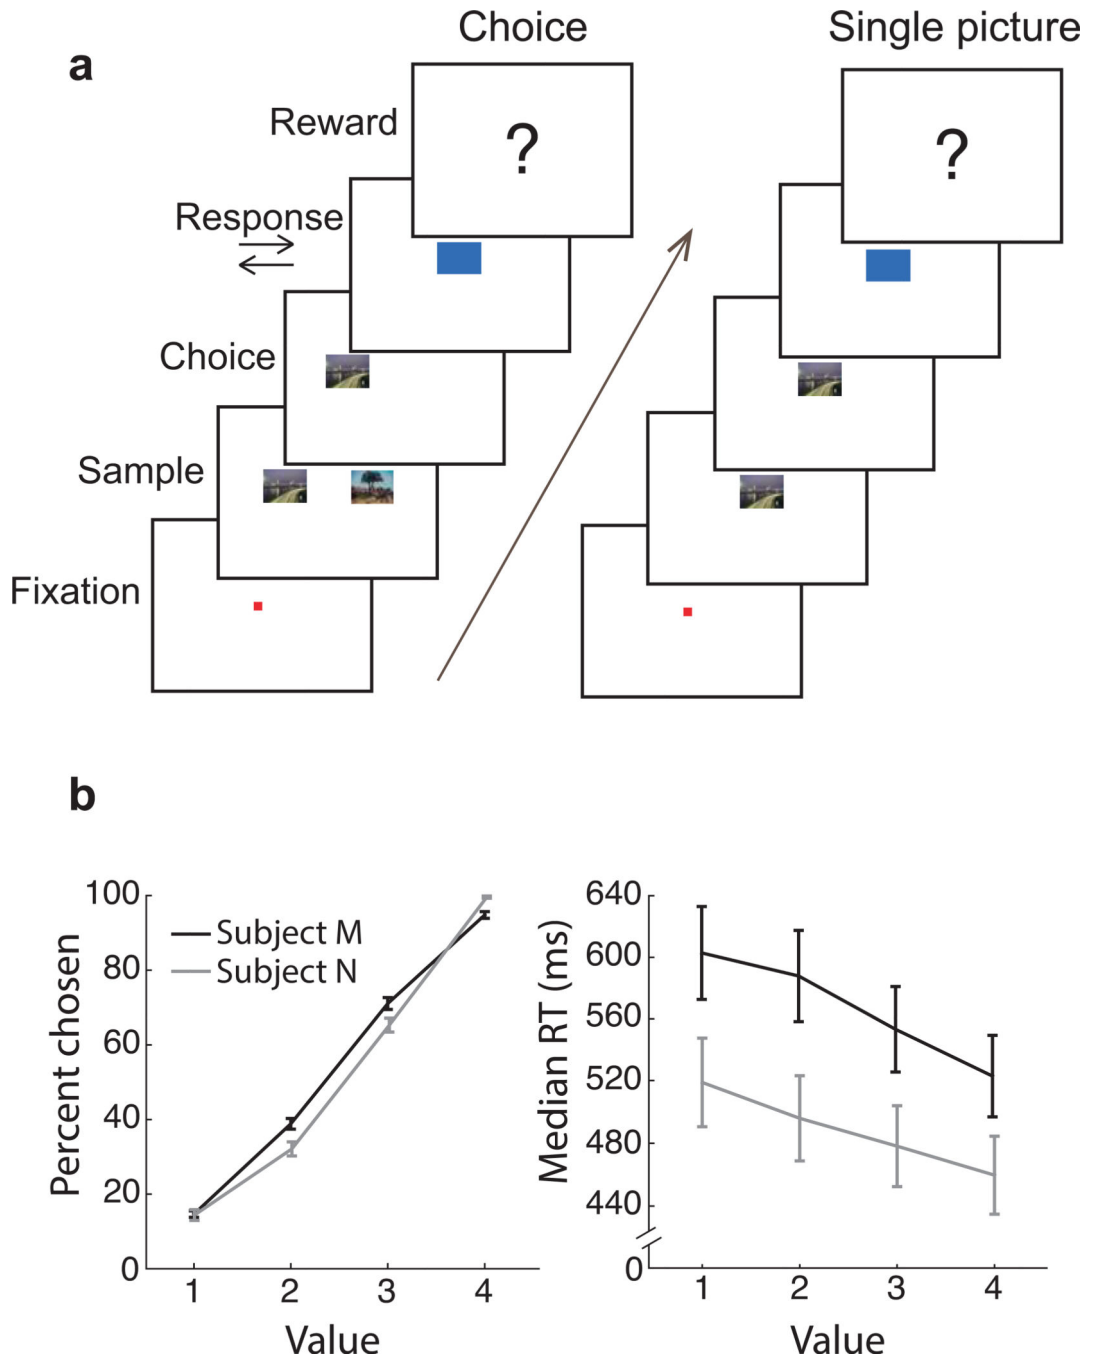


Fig. 1. Behavioral task and performance

(a) To begin a trial, subjects fixated a central point for 450 ms. On choice trials, two pictures at $\pm 5^\circ$ visual angle predicted different reward amounts. Subjects freely viewed both images and chose one by fixating it for 450 ms. After a choice, another cue appeared instructing a right or left joystick response, which, if executed correctly, resulted in the reward associated with the picture chosen at the beginning of the trial. Single picture trials were identical to choice trials, except only one randomly selected picture was shown. Subjects had to fixate the picture for 450 ms and make the subsequently instructed joystick response to obtain

reward. **(b)** Both subjects learned eight reward-predicting pictures well, choosing more valuable pictures on choice trials (left; regression of percent chosen per session on picture value. Subject M: $n = 96$ (24 sessions \times 4 values) $r^2 = 0.96$ $p = 2 \times 10^{-67}$, Subject N: $n = 80$ $r^2 = 0.94$ $p = 3 \times 10^{-50}$), and making faster joystick responses for higher value pictures on single picture trials (right; RT = Reaction Time; Regression of $\log(\text{RT})$ on picture value: Subject M $r^2 = 0.34$ $p = 6 \times 10^{-10}$, Subject N $r^2 = 0.25$ $p = 3 \times 10^{-6}$). Error bars are \pm SEM.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

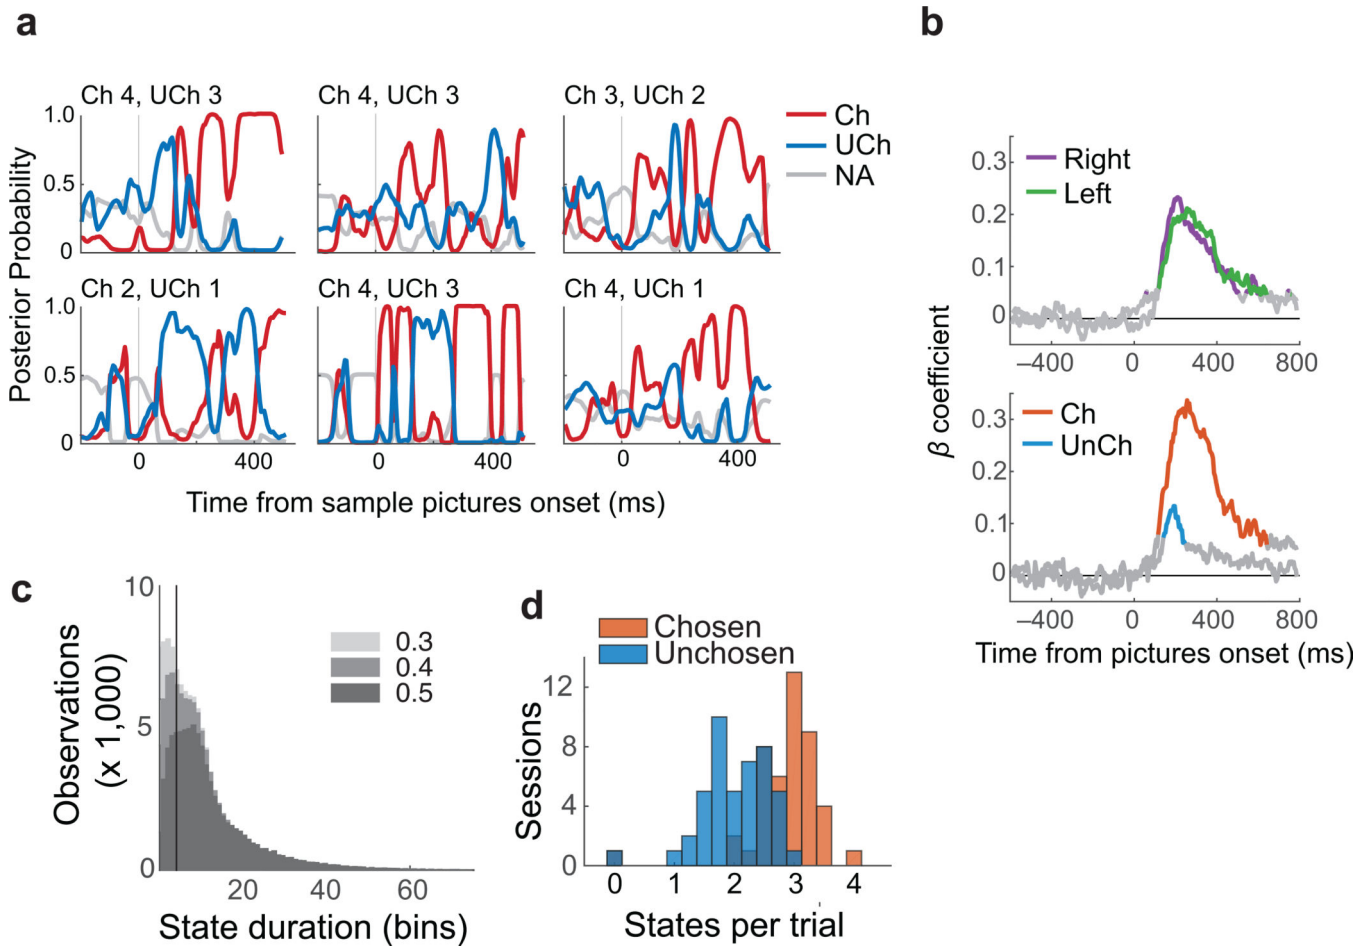


Fig. 2. Decoded choice dynamics

(a) Posterior probabilities derived from LDAs for chosen (red), unchosen (blue) and unavailable options (gray, average of both unavailable options), shown for six typical trials. (b) Pictures on the right and left of the screen (Top) or that were chosen (Ch) and not chosen (UnCh) (Bottom) explained significant variance in decoder classifications. Colored lines show times with significant β coefficients from the multiple regression ($p < 0.001$ to account for multiple comparisons). (c) Histograms of putative states decoded from choice trials, according to the number of consecutive time bins in which the same value was decoded (duration). All data are plotted three times, each with a different threshold (gray shades). Observations with posterior probabilities below the designated threshold were removed. The vertical line indicates a 4 bin duration, which was the cut-off used to define a stable state. (d) The number of stable states per trial, averaged for each session. Chosen states were more prevalent than unchosen states.

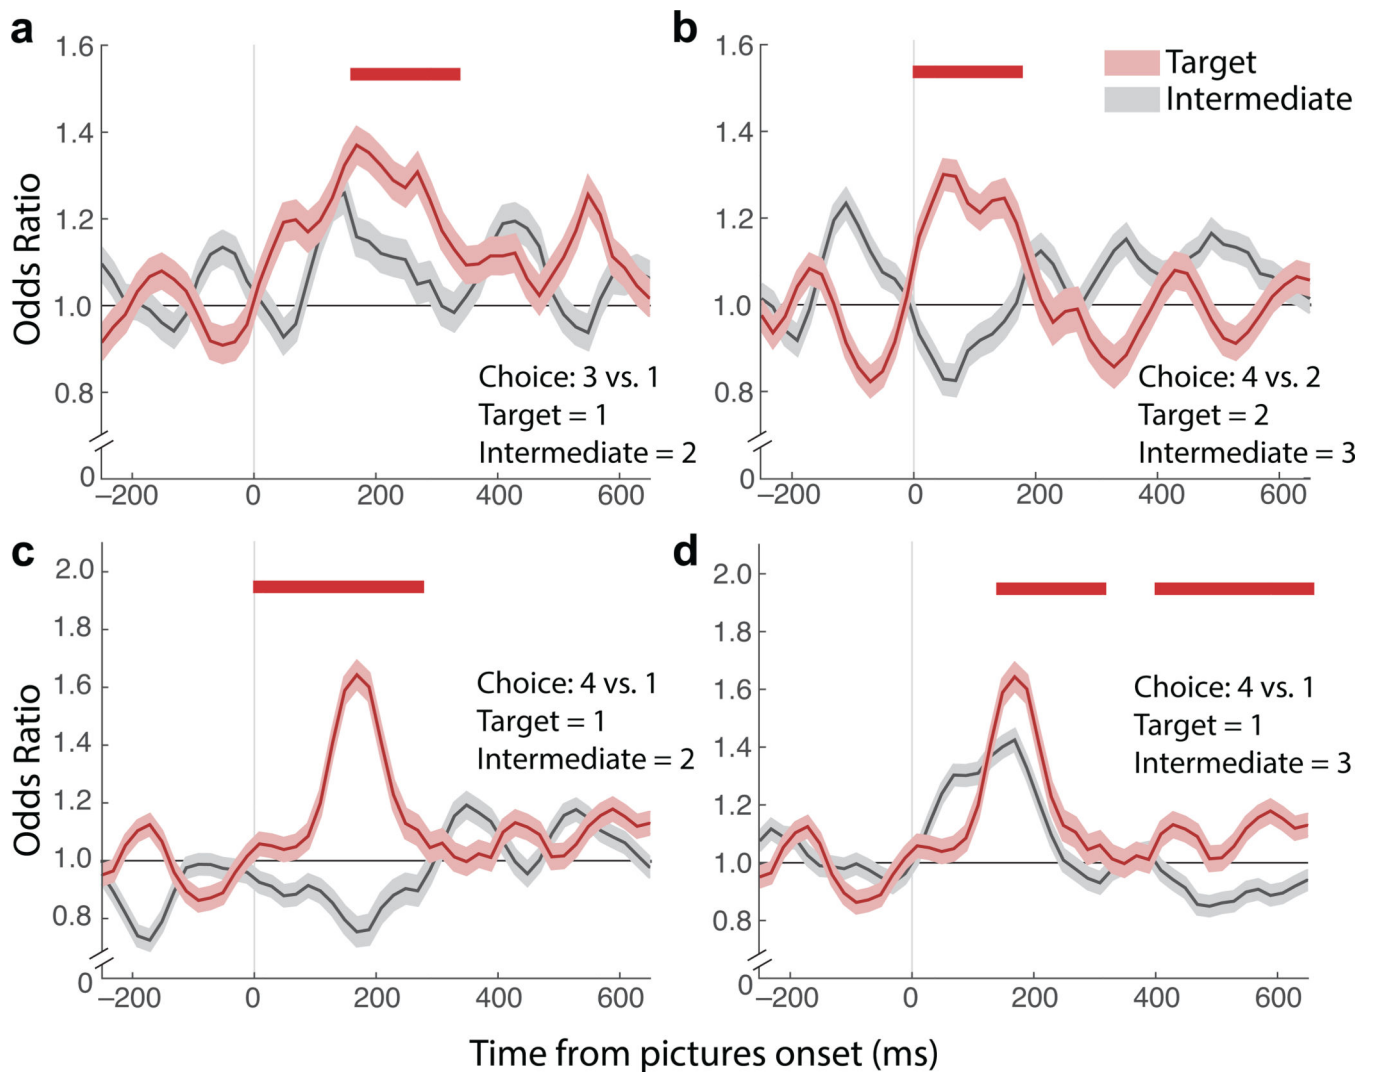


Fig. 3. Odds of decoding unchosen options

The odds of decoding the unchosen value (“target”, red) were calculated as an odds ratio among all trials with the same chosen value, and compared to the odds of decoding a value between the chosen and unchosen option (“intermediate”, gray) among the same trials. **(a)** For correct choices in which 3 was chosen, the odds of decoding 1 given that 1 was present were higher than the odds of decoding 2 given that 1 was present. **(b)** For correct choices in which 4 was chosen, the odds of decoding 2 given that 2 was present were higher than the odds of decoding 3 given that 2 was present. **(c)** For correct choices in which 4 was chosen, the odds of decoding 1 given that 1 was present were higher than the odds of decoding 2 given that 1 was present. **(d)** The odds of decoding 1 or 3, given that 1 was present, for correct choices in which 4 was chosen. Here, decoder noise raises the odds of decoding 3, however decoding 1 was still more likely. Shading = \pm SEM. Odds ratios were calculated in 70 ms epochs, stepped by 15 ms. Red bar shows differences greater than shuffled trials at $p < 0.01$ for at least six time bins (approximately 100 ms). This significance level was established by finding the threshold that reduced pre-stimulus false discoveries to 1%.

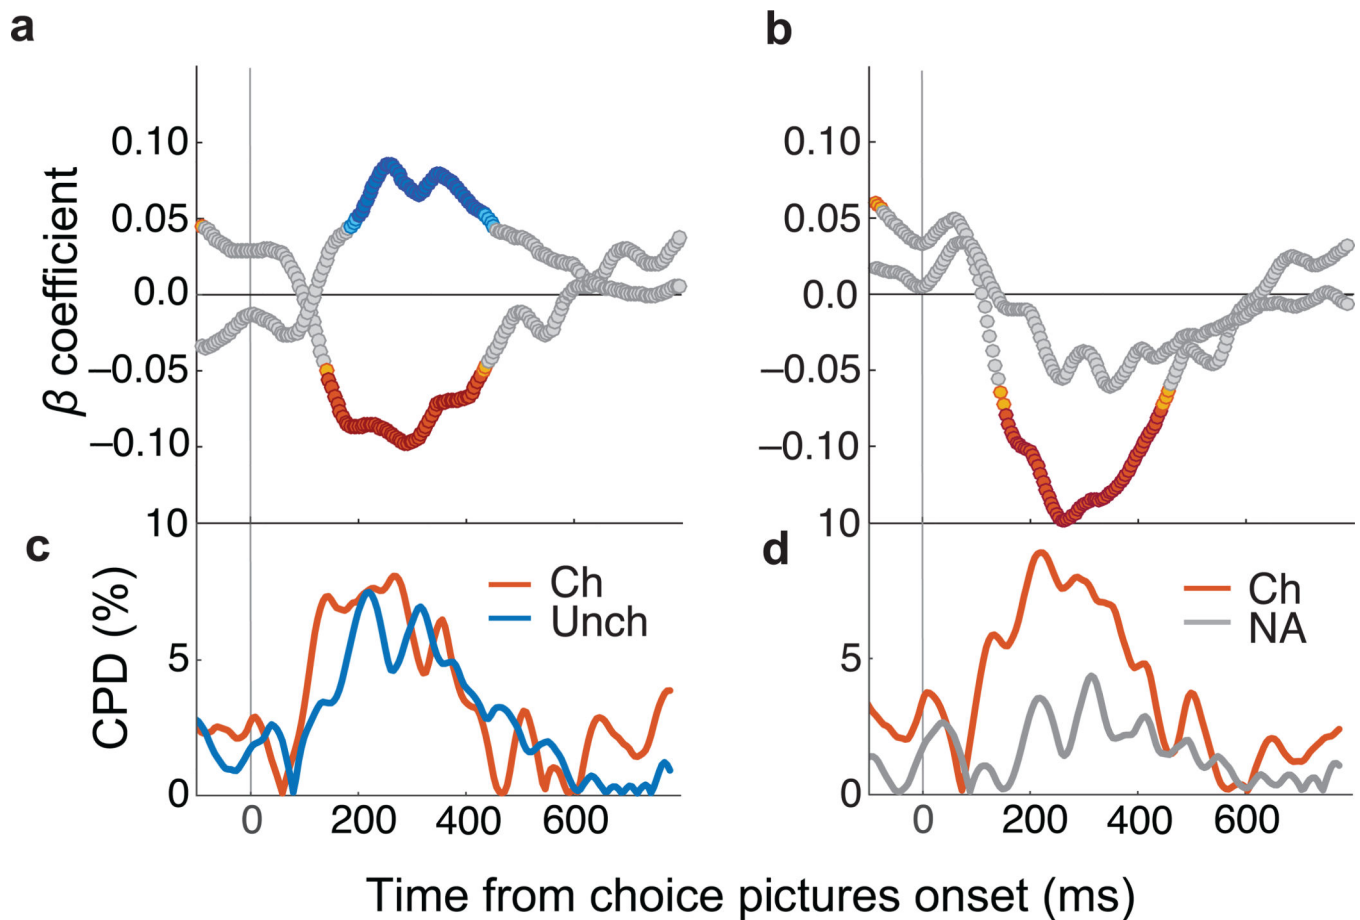


Fig. 4. Decoded representations predict choice times

(a) Posterior probabilities (averaged in 20 ms bins stepped by 5 ms) predicted choice times. β coefficients from multiple regressions of choice times on two factors: probability of chosen and unchosen states (maximum variance inflation factor (VIF) = 1.05). (b) β coefficients from multiple regressions of choice times on two factors: probability of the chosen and unavailable states (maximum VIF = 2.22). Choice times were log-transformed and probabilities z-scored so β coefficients could be compared. Orange = chosen $p < 0.005$, Red = chosen $p < 0.001$, Teal = unchosen $p < 0.005$, Blue = unchosen $p < 0.001$, gray = not significant. (c–d) The percent of choice time variance accounted for by each factor, quantified with coefficients of partial determination (CPD). Ch = chosen, Unch = unchosen, NA = not available.

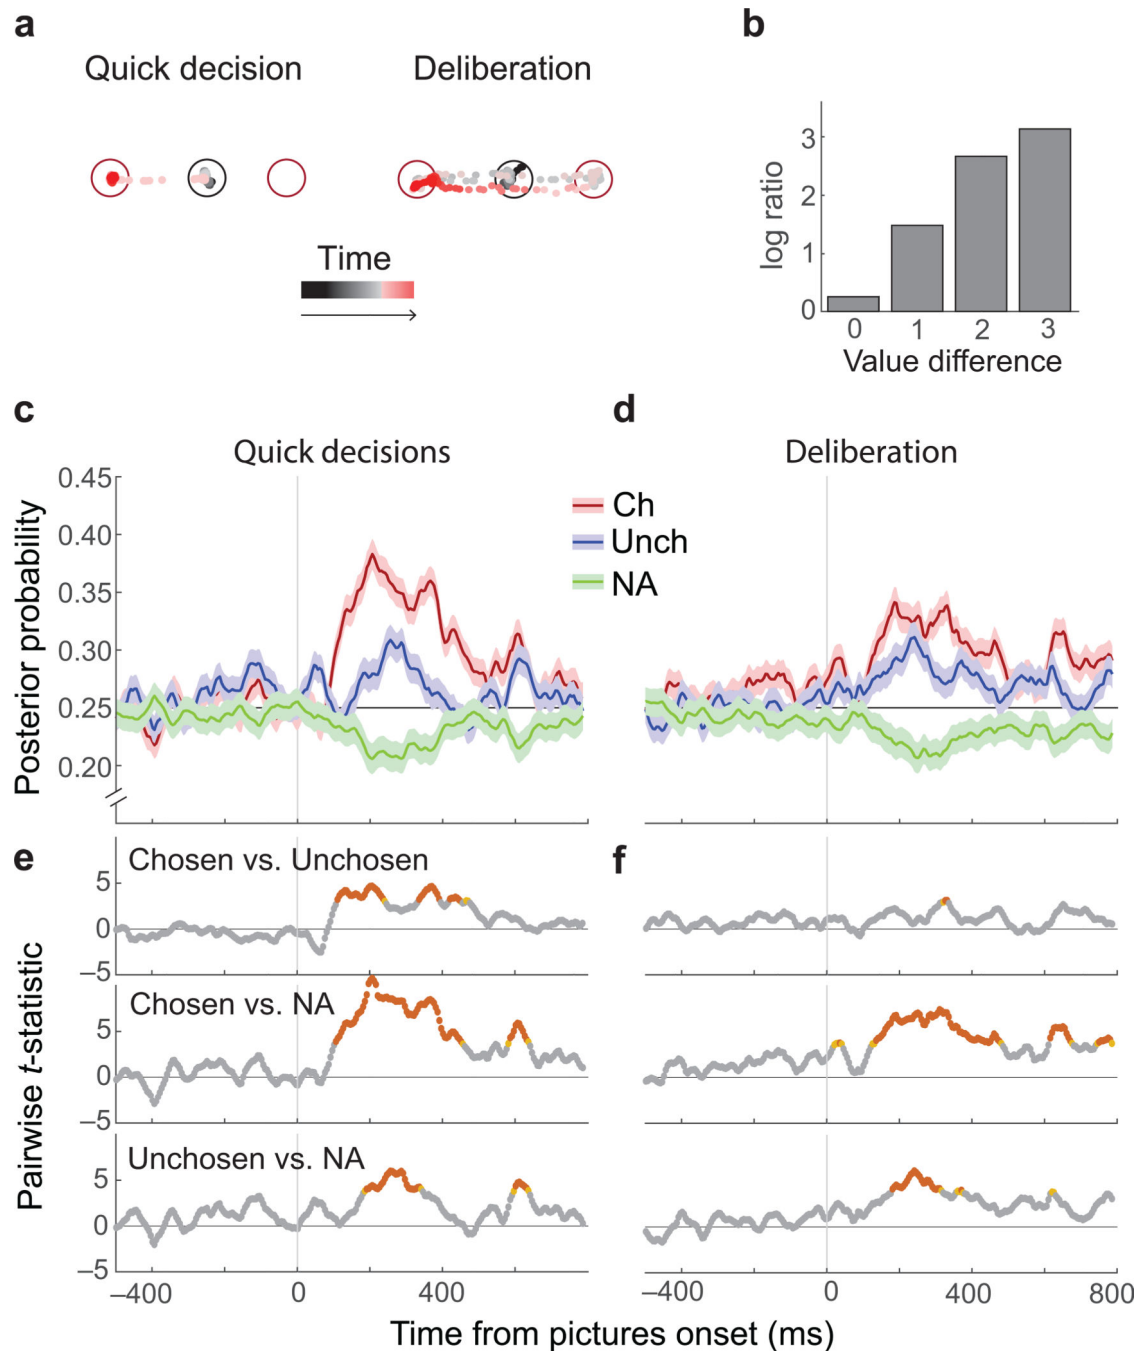


Fig. 5. Quick versus deliberative decisions

(a) Eye positions relative to the fixation point (black circle) and two choice options (red circles) are shown as points varying from black (start of trial) to red (option selection) for two example trials. In quick decisions the subjects' eyes moved from the fixation point to one picture. In deliberative decisions, eyes fell within $\pm 2.5^\circ$ of the center of one picture at least twice before selecting an option. (b) Quick decisions increased with increasing difference in option values. The height of each bar is the overall ratio (log-scale) of quick to deliberative decisions across all trials with a given value difference. (c–d) The average (\pm

SEM) probability that neural activity represented the chosen picture (red), the unchosen picture (blue) or an unavailable option (averaged across two unavailable options; green) for quick or deliberative decisions. (e–f) For quick decisions, there was a larger difference in the relative strength of the chosen and unchosen representation. At each time point, a 3×2 ANOVA with factors of representation type (chosen/ unchosen/ unavailable) and decision type (quick/ deliberative) was performed. The interaction term reached significance ($p < 0.05$) in multiple bins after picture onset, indicating that neural representations varied by decision type (not shown). Tukey’s HSD assessed pairwise contrasts. Red: $p < 0.005$ orange $p < 0.01$ gray $p > 0.01$.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

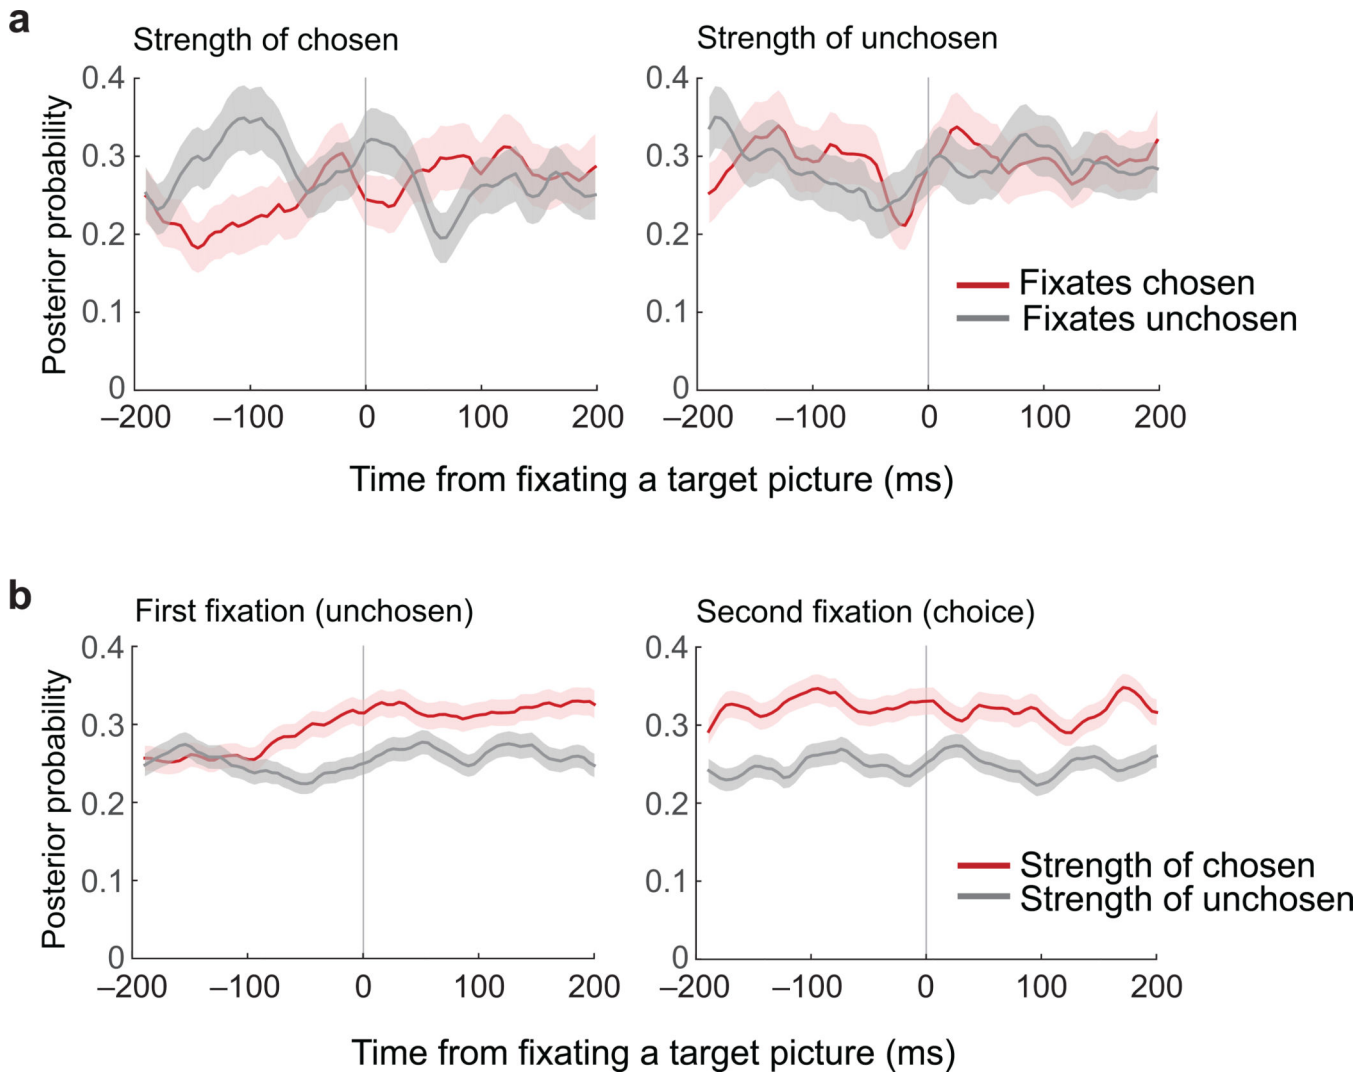


Fig. 6. Targets of gaze do not affect OFC representations

(a) Strength of chosen and unchosen representations during deliberation trials, aligned to the times the subject fixated one of the two choice options. The first and second fixations were included in this analysis, and only if they consisted of one saccade to each picture ($n = 1058$ trials). (b) Strength of chosen and unchosen representations on trials with one saccade to the unchosen item, followed by a saccade to the chosen item ($n = 550$ trials). Plots are aligned to the first and second fixation in each sequence. Lines are mean posterior probabilities \pm SEM. The chosen representation becomes progressively stronger than the unchosen representation, but this process was not affected by the fixations.




Fig. 7. Removing single neurons does not disrupt decoded states

(a) An example trial, with choice options of values 3 and 4. Colors indicate the value decoded at each point in time. The top row shows the values decoded from the full ensemble, and the rows beneath show the values decoded when each of eight neurons was held out. (b) For every neuron, a reduced ensemble was created by holding it out. Correlations were calculated between the time series of values decoded from the reduced ensemble and the corresponding full ensemble, and r^2 values are shown. (c) The average effects of holding different numbers of neurons out from the full ensemble. Each point is the average r^2 from one session, in which the same number of neurons were held out in different combinations. The black line is an exponential curve fit to the distribution. When 100% of neurons were removed, values were decoded from LFP data alone, and the minimum r^2 observed was 0.049.

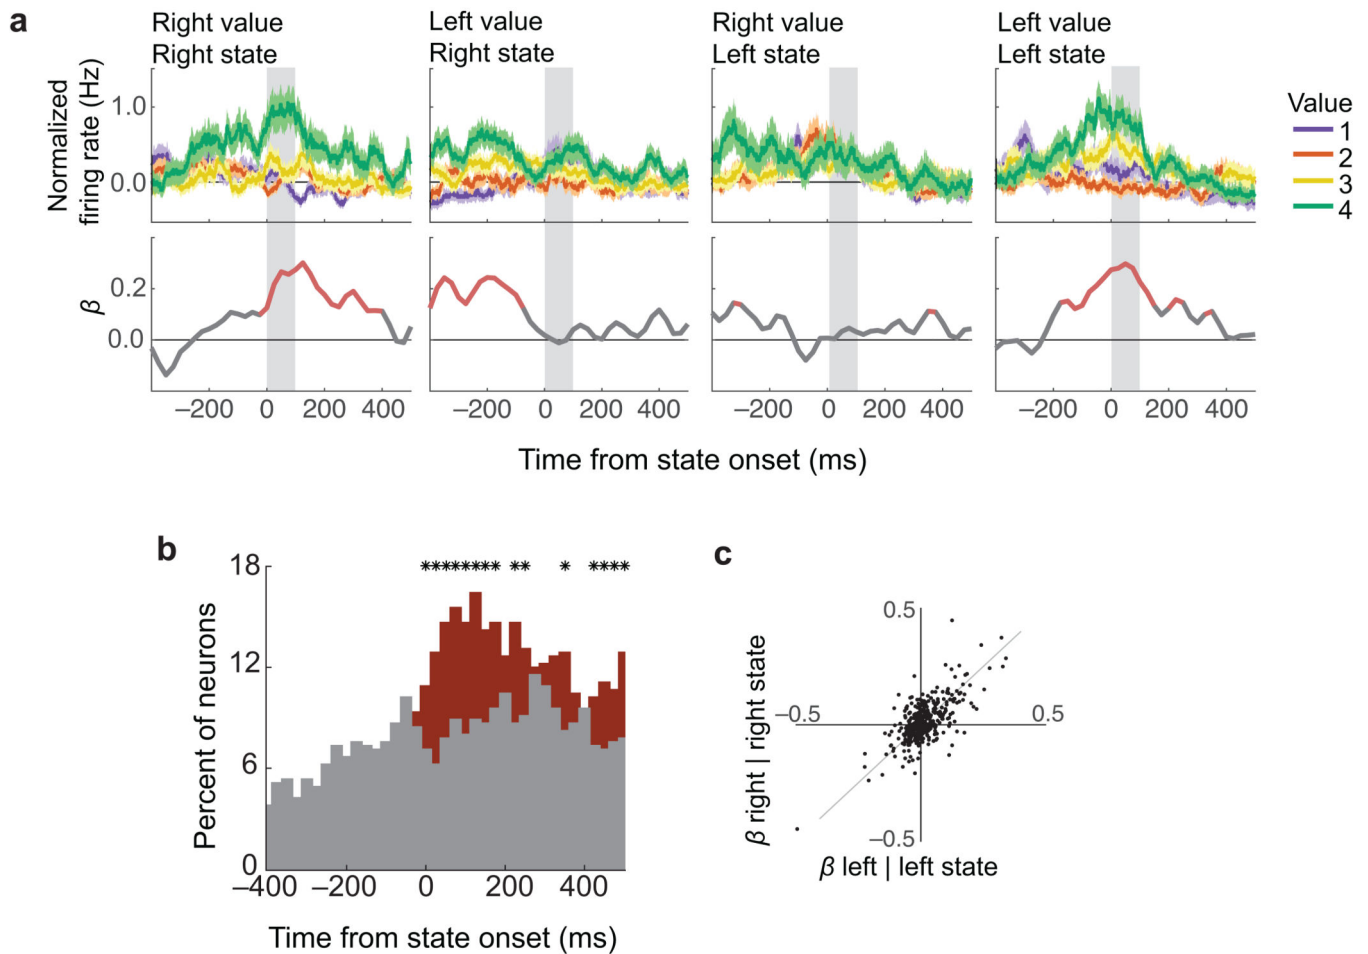


Fig. 8. Single neurons encode current states

Normalized firing rates of single neurons, re-aligned to the onsets of decoded states. **(a)** A neuron that encoded current states. In each subplot (average firing rates \pm SEM), neuron activity was aligned to states in which the right or left picture was decoded from the rest of the population (Right, Left states respectively). During 'Right states', this neuron encoded the value of pictures on the right (Right value), but not left (Left value; first two subplots). During 'Left states', it encoded 'Left values', but not 'Right values'. Bottom panels show β coefficients from a multiple regression of firing rate (averaged over 50 ms, stepped by 10 ms) on four factors corresponding to titles on the subplots (e.g. 'Right value' during 'Right states', etc.). Red: $p < 0.01$. **(b)** Percent of neurons encoding current state value (i.e. 'Right values' during 'Right states' and 'Left values' during 'Left states') (red) and those encoding the value of the alternate picture (i.e. 'Right values' during 'Left states' and 'Left values' during 'Right states') (gray). Significance was defined by non-zero β coefficients in a multiple regression ($p < 0.01$). * = red and gray proportions differ (χ^2 test, $p < 0.01$). **(c)** Across all neurons, β coefficients in the first 100 ms after state onset for 'Right values' during 'Right states' (right | right) were correlated with 'Left values' during 'Left states' (left | left) ($r^2 = 0.55$, $p = 6 \times 10^{-81}$; gray line = unity).