



This postprint was originally published as:

Chan, S. C. Y., Schuck, N. W., Lopatina, N., Schoenbaum, G., & Niv, Y. (2021). Orbitofrontal cortex and learning predictions of state transitions. *Behavioral Neuroscience*, 135(4), 487–497.

<https://doi.org/10.1037/bne0000461>

The following copyright notice is a publisher requirement:

©American Psychological Association, 2021. This paper is not the copy of record and may not exactly replicate the authoritative document published in the APA journal. Please do not copy or cite without author's permission.

The final article is available, upon publication, at:

<https://doi.org/10.1037/bne0000461>

Provided by:

Max Planck Institute for Human Development

Library and Research Information

library@mpib-berlin.mpg.de

Orbitofrontal Cortex and Learning Predictions of State Transitions

Stephanie C. Y. Chan¹, Nicolas W. Schuck^{2, 3}, Nina Lopatina⁴, Geoffrey Schoenbaum⁵, and Yael Niv¹

¹ Department of Psychology, Princeton Neuroscience Institute, Princeton University

² Max Planck Research Group NeuroCode, Max Planck Institute for Human Development, Berlin, Germany

³ Max Planck UCL Centre for Computational Psychiatry and Ageing Research, Berlin, Germany

⁴ Lab41, IQT Labs, Menlo Park, California, United States

⁵ National Institute on Drug Abuse, Intramural Research Program, Maryland, United States

The orbitofrontal cortex (OFC) has been implicated in goal-directed planning and model-based decision-making. One key prerequisite for model-based decision-making is learning the transition structure of the environment—the probabilities of transitioning from one environmental state to another. In this work, we investigated how the OFC might be involved in learning this transition structure, by using fMRI to assess OFC activity while humans experienced probabilistic cue–outcome transitions. We found that OFC activity was indeed correlated with behavioral measures of learning about transition structure. On a trial-by-trial basis, OFC activity was associated with subsequently increased expectation of the more probable outcome; that is, with subsequently more optimal cue–outcome predictions. Interestingly, this relationship was observed no matter what outcome occurred at the time of the OFC activity, and thus is inconsistent with an interpretation of the OFC activity as representing a “state prediction error” that would facilitate learning transitions via error-correcting mechanisms. Finally, OFC activity was related to more optimal predictions only for subsequent trials involving the same cue that was observed at the time of OFC activity—this relationship was not observed for subsequent trials involving a different cue. All together, these results indicate that the OFC is involved in updating or reinforcing a learned transition model on a trial-by-trial basis, specifically for the currently observed cue–outcome associations.

Keywords: orbitofrontal cortex, state transitions, planning, model-based decision making

To flexibly plan for the future, we must be able to predict which states of the world lead to which (i.e., we need to learn a model of the “transition structure” of the world). For example, to decide whether to drink warm milk or coffee, we need to know that warm milk makes us sleepy, but coffee wakes us up. This type of planning has been termed “model-based decision-making,” in contrast to “model-free decision-making,” which does not require such a model (Daw et al., 2005).

The orbitofrontal cortex (OFC) has been shown to be particularly important for model-based decision-making (Baxter et al., 2000; De Wit et al., 2009; Izquierdo et al., 2004; McDannald et al., 2011; Rudebeck & Murray, 2011; Valentin et al., 2007; Walton et al., 2010). However, previous research has focused on showing that OFC activity relates to the expected values of future rewards (Fellows, 2007; Gottfried et al., 2003; Hampton et al., 2006; Hare et al., 2008; Monosov & Hikosaka, 2012; Padoa-Schioppa & Assad, 2006; Wallis & Kennerley, 2011). Recently, we have instead proposed that the OFC represents the current state of the task (Schuck et al., 2016) and that the OFC is especially critical for making decisions in situations where environmental stimuli do not unambiguously determine the task-relevant state (e.g., whether the state is “Thursday evening” and it is bedtime, vs. “Friday evening,” in which case I don’t want to become sleepy as I am going to a party; Bradfield et al., 2015; Chan et al., 2016; Nogueira et al., 2017; Wilson et al., 2014). However, both value and state representation are important in model-free as well as model-based decision-making, and therefore these two lines of research do not explain why the OFC is critical specifically for the latter.

Yet another line of research provides a potential explanation for the OFC’s particular prominence in model-based planning. This research suggests that the OFC is important for learning about the state-to-state “transition structure” of the world—the tendencies of

Stephanie C. Y. Chan <https://orcid.org/0000-0003-2590-7832>

Nicolas W. Schuck <https://orcid.org/0000-0002-0150-8776>

Nina Lopatina <https://orcid.org/0000-0001-6844-4941>

Geoffrey Schoenbaum <https://orcid.org/0000-0001-8180-0701>

Yael Niv <https://orcid.org/0000-0002-0259-8371>

The authors declare no competing financial interests.

This work was supported by the National Institute of Mental Health (R01MH0988861 and T32MH065214); a Sloan Research Fellowship (to Yael Niv); the National Science Foundation (Collaborative Research in Computational Neuroscience award IIS1009452); the U.S. Army Research Laboratory and the U.S. Army Research Office (W911NF1410101); the National Institute on Drug Abuse (Geoffrey Schoenbaum and Nina Lopatina), an Independent Max Planck Research Group grant awarded by the Max Planck Society (M.TN.A.BILD0004, Nicolas W. Schuck), and a Starting Grant from the European Union (ERC-2019-StG-REPLAY- 852669, Nicolas W. Schuck). The opinions expressed in this article are the authors’ own and do not reflect the view of the NIH/DHHS.

Correspondence concerning this article should be addressed to Stephanie C. Y. Chan, Department of Psychology, Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08544, United States. Email: chan.stephanie.cy@gmail.com

certain environmental states to lead to other states. One study showed that OFC-lesioned rats couldn't learn about changes in the transitions from cues to outcomes (cue–outcome associations; McDannald et al., 2011), while a study in humans linked fMRI surprise signals in lateral OFC with updates in the hippocampus of a model of transition structure (Boorman et al., 2016). Some newer studies have observed such surprise signals in the midbrain (Sharpe et al., 2017; Stalnaker et al., 2019; Takahashi et al., 2017), or in a set of areas including the midbrain and lateral OFC (Howard & Kahnt, 2018; Suarez et al., 2019). Howard and Kahnt (2018) additionally found that midbrain identity prediction errors were correlated with a cue–outcome learning and changes in outcome identity representations in the OFC. The hypothesized link between OFC and learning transition structure could also explain OFC's centrality to model-based decision-making, given that transition structure is a critical component of the “model” in such decision-making. One cannot plan and mentally simulate the future result of current actions without an accurate model of how state transitions are likely to unfold in the future.

How exactly might the OFC be involved in learning about transition structure? The OFC might itself compute or represent a prediction error at the time of unexpected outcomes, which can be used to update an internal model of the transition structure. Such “state prediction error” signals would occur upon observing state transitions that are unexpected, and could be used to guide learning so that transitions are better predicted in the future (e.g., Glascher et al., 2010). Note that these error signals are analogous to—but distinct from—reward prediction errors that are used for learning to associate states with their reward values (e.g., Montague et al., 1996; Rescorla & Wagner, 1972). However, the existing research does not make specific predictions about the role of OFC in representing or learning about transition structure, and state prediction errors are just one possible way. We, therefore, set out here to test whether the OFC might be involved in error-driven learning via signaling of state prediction errors and whether OFC activity could predict behavior related to learning transition structure. We also tested the two dominant hypotheses of OFC function—representing the current state and representing expected value.

In our experiment, black-and-white image cues led stochastically to M&M candies of different quantities and colors (outcomes). In the critical trials, the number of M&Ms was fully predictable, but their color was not, so as to generate state prediction errors in the absence of reward prediction errors. Using fMRI, we investigated activity in the human OFC at the time of these outcomes, and its relationship with participants' behavioral predictions of state transitions.

Materials and Methods

Subjects

Twenty-four volunteers from the Princeton University community participated in exchange for monetary compensation (\$20 per hour + up to \$10 performance-related bonus). All subjects were right-handed (14 female, age range 18–34 years) and stated that they liked M&Ms. Informed written consent was obtained from all subjects, and the study protocol was approved by the Institutional Review Board for Human Subjects at Princeton University.

Experimental Design

Each trial began with .5–8 s of fixation (truncated exponential distribution, mean 2.4 s). Then one of four black-and-white image cues depicting outdoor scenes appeared for 1.2 s (see Figure 1a). On 75% of the trials, this was followed by the opening of a box around the image (.2 s). Then, a set of M&Ms appeared below the image and fell into a bowl, over the course of .9 s. As the M&Ms fell into the bowl, one clinking sound was emitted for each M&M in the set. A tally at the bottom of the screen (not shown in Figure 1a) indicated the total number of M&Ms received so far, for each of the four possible colors.

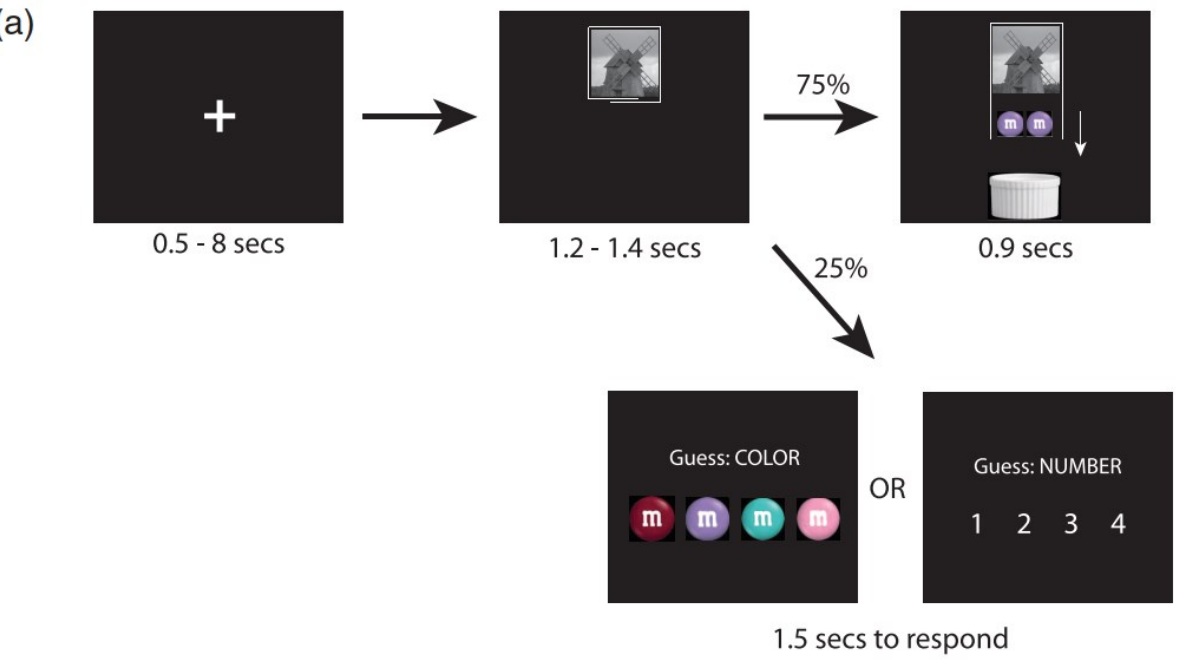
Each of the four image cues was associated with different numbers and colors of M&Ms according to a predetermined schedule of reinforcement (Figure 1b). Cue A and Cue B were designed to elicit state prediction errors throughout the experiment due to a probabilistic schedule of M&M color, but not reward prediction errors, because they always dropped exactly 2 M&Ms. Cue C, in contrast, was associated with 2 M&Ms of a fixed color, thus eliciting no prediction errors once the contingencies had been learned. Finally, Cue D was designed to elicit only reward prediction errors—~~it dropped either 1 or 4 M&Ms of a fixed color (as with the other image cues, Cue D led to 2 M&Ms on average, such that all 4 cues were equated for average reward value).~~ For each subject, the images and M&M colors were assigned randomly from a pool of 20 images and 5 M&M colors.

Subjects earned one real M&M of a given color for every 17 “virtual” M&Ms that they received in the task. Subjects were requested to refrain from eating or drinking (except water) for at least 3 hr prior to the experiment, so that the M&Ms would be especially rewarding. Nonstandard M&M colors were chosen to circumvent preexisting preferences for specific M&M colors and to achieve perceptually distinct outcomes that are of equal value. Note also that our analyses of state prediction error always combine Cue A and Cue B trials, so that any potential value differences between the two colors cancel out. In a postexperiment questionnaire, subjects rated the appeal of the M&Ms on a scale from 1 (*not appealing at all*) to 5 (*very appealing*). The mean rating was $3.8 \pm .2$.

Twenty-five percent of all trials (pseudorandomly distributed) were “guess trials.” On these trials, the appearance of the black-and-white image cue was followed by a prompt reading “Guess: COLOR” or “Guess: NUMBER.” At the appearance of the prompt, the image cue disappeared. Subjects were given 1.5 s to guess what color/number of M&Ms *would have* fallen on that trial. Subjects received 10¢ for every question correctly answered. The purpose of the guess trials was to encourage subjects to pay attention to the image cue and to actively make a prediction of the upcoming M&M outcome on *every* trial—because the allowed response time was so short, subjects had to prepare an answer upon viewing the image cue in case a guess prompt followed.

Subjects performed 72 training trials outside of the scanner, to familiarize themselves with the task and to learn the stimulus–outcome contingencies. During training, subjects received and ate the M&Ms they earned (approximately 7 M&M candies). They were then informed that future M&Ms they earned would be given to them after the ensuing scanning session, and they performed another 420 trials in the MRI scanner. At the end of the experiment, subjects received all M&Ms earned while in the scanner. The 420 trials were evenly distributed between the 4 image cues, with trial order pseudorandomized so that the total number of

Figure 1
Experimental Design



(b)

End state

2 x Color1

2 x Color2

2 x Color3

1 x Color4

4 x Color4

Start state		Image A	2/3	1/3	--	--	--	State PE
		Image B	1/3	2/3	--	--	--	
		Image C	--	--	1	--	--	No PE
		Image D	--	--	--	2/3	1/3	State PE & Reward PE

Note. Trials began with fixation, then, one of four image cues (four different “start states”) appeared. On most trials, the box around the image opened, and a number of colored M&M’s dropped from the image, clinking as they fell into a bowl (five different “end states”). On the randomly interspersed “guess” trials, the image cue was instead followed by a prompt to guess (within 1.5 s) either the color or number of M&M’s that would have fallen on that trial. (b) Cue–outcome contingencies for each of the four images (transition matrix for the experiment). Numbers in table indicate probability of each end state (M&M outcome) given each start state (image cue). Larger state prediction errors are expected for rarer outcomes (smaller transition probabilities). Images and M&M colors were assigned randomly for each subject. our analyses focused on cue A and cue B trials, which were designed to elicit state prediction errors in the absence of reward prediction errors PE = prediction error.

M&M’s collected increased at the same rate for every color. The experiment was divided into five scan sessions of approximately 10 min each.

Behavioral Measures

We evaluated three types of behavioral measures, computed separately for each subject and for each prediction trial type (image cue type × number/color prediction): (a) overall performance over the course of the experiment, (b) change in performance over the course of the experiment, and (c) sensitivity to the most recent outcome (a proxy for learning rate).

To assess overall performance, we computed the fraction of responses that were optimal (i.e., for which the subject selected the common outcome), across all scan sessions. To measure the change in performance, we computed the difference in performance from the beginning to the end of the experiment as the fraction of optimal responses in the last scan session minus the fraction of optimal responses in the training session. To assess sensitivity to the previous outcome, we computed the probability of predicting the

common outcome after observing the common outcome on the previous trial with the same image cue, compared to the probability of predicting the common outcome after observing the uncommon outcome on the previous trial with the same image cue. The difference between these two quantities served as a proxy for learning rate—subjects with high learning rate would be more sensitive to the most recent outcome and would show a larger difference between the two quantities.

fMRI Acquisition

Functional brain images were acquired using a 3T MRI scanner (Skyra; Siemens Erlangen, Germany), and were preprocessed using FSL (<http://fsl.fmrib.ox.ac.uk/fsl/>). An echoplanar imaging sequence was used to acquire 40 slices of 2-mm thickness with a 1-mm gap (repetition time (TR) = 2.4s, echo time (TE) = 27 ms, flip angle = 71°, the field of view = 196 mm, phase encoding direction = anterior to posterior). We optimized our fMRI sequence for OFC signal acquisition by including a gap between slices, using shimming and fieldmap unwarping, and tilting the slices by approximately 30° from the axial plane toward a coronal orientation (Deichmann et al., 2003). Fieldmaps consisted of forty 3-mm slices, centered at the centers of the echoplanar slices, with TR = 500 ms, TE1 = 3.99 ms, TE2 = 6.45 ms, field of view = 196 mm. At the end of the five functional scanning sessions, an MPRAGE anatomical scan was acquired, consisting of 176 1-mm axial slices, TR = 2.3s, TE = 3.08 ms, flip angle = 9°, and field of view = 256 mm.

Preprocessing

All functional images were preprocessed using high-pass filtering (filter at 1/100 Hz), motion correction (six-parameter rigid body transformation), correction for B0 magnetic inhomogeneities (fieldmap unwarping), spatial smoothing (Gaussian kernel with a full width at half maximum of 5 mm), and coregistration of functional and structural scans. For General Linear Model (GLM) results, we additionally performed spatial normalization of subject-level results to match a template in Montreal Neurological Index (MNI) space (12-parameter affine transformation).

Functional Parcellation of Orbitofrontal Cortex

Regions of interest for the OFC were obtained from Kahnt et al. (2012), who used *k*-means clustering of functional connectivity patterns to parcellate OFC into subregions. We used the parcellation of OFC into two clusters, which correspond with medial–lateral subdivisions of OFC found in studies of cytoarchitectonic structure and of intraregional anatomical connectivity (Carmichael & Price, 1996; Öngür & Price, 2000).

Obtaining Mean Percent Signal Change at M&M Outcomes

Using the FSL toolbox (<http://fsl.fmrib.ox.ac.uk/fsl/>), we performed a GLM analysis with the following regressors: One regressor for the onsets of each type of image cue (A, B, C, D); one regressor for the onsets of the M&M outcomes for Cue C; one regressor for the onsets of the uncommon outcomes for each of the image cues A, B, and D (three regressors total); one regressor for the onsets of the common outcomes for each of the image cues A, B, and D (three regressors total); and one parametric regressor for the clinks of the M&Ms into the bowl (one, two, or four clinks). These 12 regressors were convolved with a standard hemodynamic response function. In addition, the design matrix included 6 motion regressors and an intercept (constant) term.

Regressor weights for each voxel and each scan session were converted to percent signal change (PSC) by multiplying by the appropriate scale factor for events of length 0.1 s convolved with the standard double-gamma hemodynamic response function, and then dividing by the mean of the voxel's timecourse for that scan session. These per-scan numbers were averaged across scans for each subject. To obtain the PSC for a region of interest, the PSC was averaged across all voxels in the region of interest.

Obtaining Trial-by-Trial Estimates of Percent Signal Change at M&M Outcomes

To obtain trial-by-trial estimates of PSC in a region of interest (ROI) at each M&M outcome, we fit a separate GLM for each trial. This GLM was identical to the one used for estimating mean PSC (above), except that the regressor for the condition of the trial of interest was split into two—one regressor modeled the onset for the trial of interest only, and a second regressor modeled the onsets of all other trials in that condition (Mumford et al., 2012). These GLMs were fitted to data that were preprocessed in FSL, but the GLMs themselves were fitted using in-lab code written in MATLAB, for computational reasons.

Multivariate Pattern Analysis Classification

The purpose of our multivariate pattern analysis (MVPA) analyses was to test whether activity in OFC at the time of the M&M outcomes contained information about the start state and end state (stimulus and outcome) for each transition. We analyzed the trials that were designed to elicit state prediction errors (Cue A and Cue B trials).

Given our rapid event-related design, we first used a GLM to deconvolve neighboring events, regress out motion artifacts, and to de-noise examples through averaging (Mumford et al., 2012). To maximize power, we divided each scan in two (1st and 2nd half). The GLM then included, for each half of each scan session, regressors modeling the appearance of the M&Ms for each of four trial types of interest (Cue A followed by M&M Color 1, Cue A followed by M&M Color 2, Cue B followed by M&M Color 1, Cue B followed by M&M Color 2), totaling 8 regressors per scan. The regressors were convolved with a canonical hemodynamic response function. In addition, for each scan session, we modeled head motion using six motion regressors and the mean activity using an intercept regressor. We estimated this GLM on each subject's smoothed, motion-corrected fMRI data using the FSL toolbox (<http://fsl.fmrib.ox.ac.uk/fsl/>).

We used the resulting patterns of voxel-wise regressor weights for the four trial types (two regressor weights per run and trial type; z-scored within voxels) as training and testing examples for a support vector machine (SVM) classification algorithm with a linear kernel (nu-SVM, as implemented in LIBSVM; Chang & Lin, 2011), under a leave-one-out cross-validation scheme, using the Princeton MVPA Toolbox (<https://code.google.com/p/princeton-mvpa-toolbox/>).

We used a standard cost (nu) parameter of 1 for the SVM (results did not depend strongly on this parameter).

To classify the start state, we classified training and testing examples according to the image cue (Cue A or Cue B). To classify end state, we classified training and testing examples according to the M&M color (Color 1 or Color 2).

Results

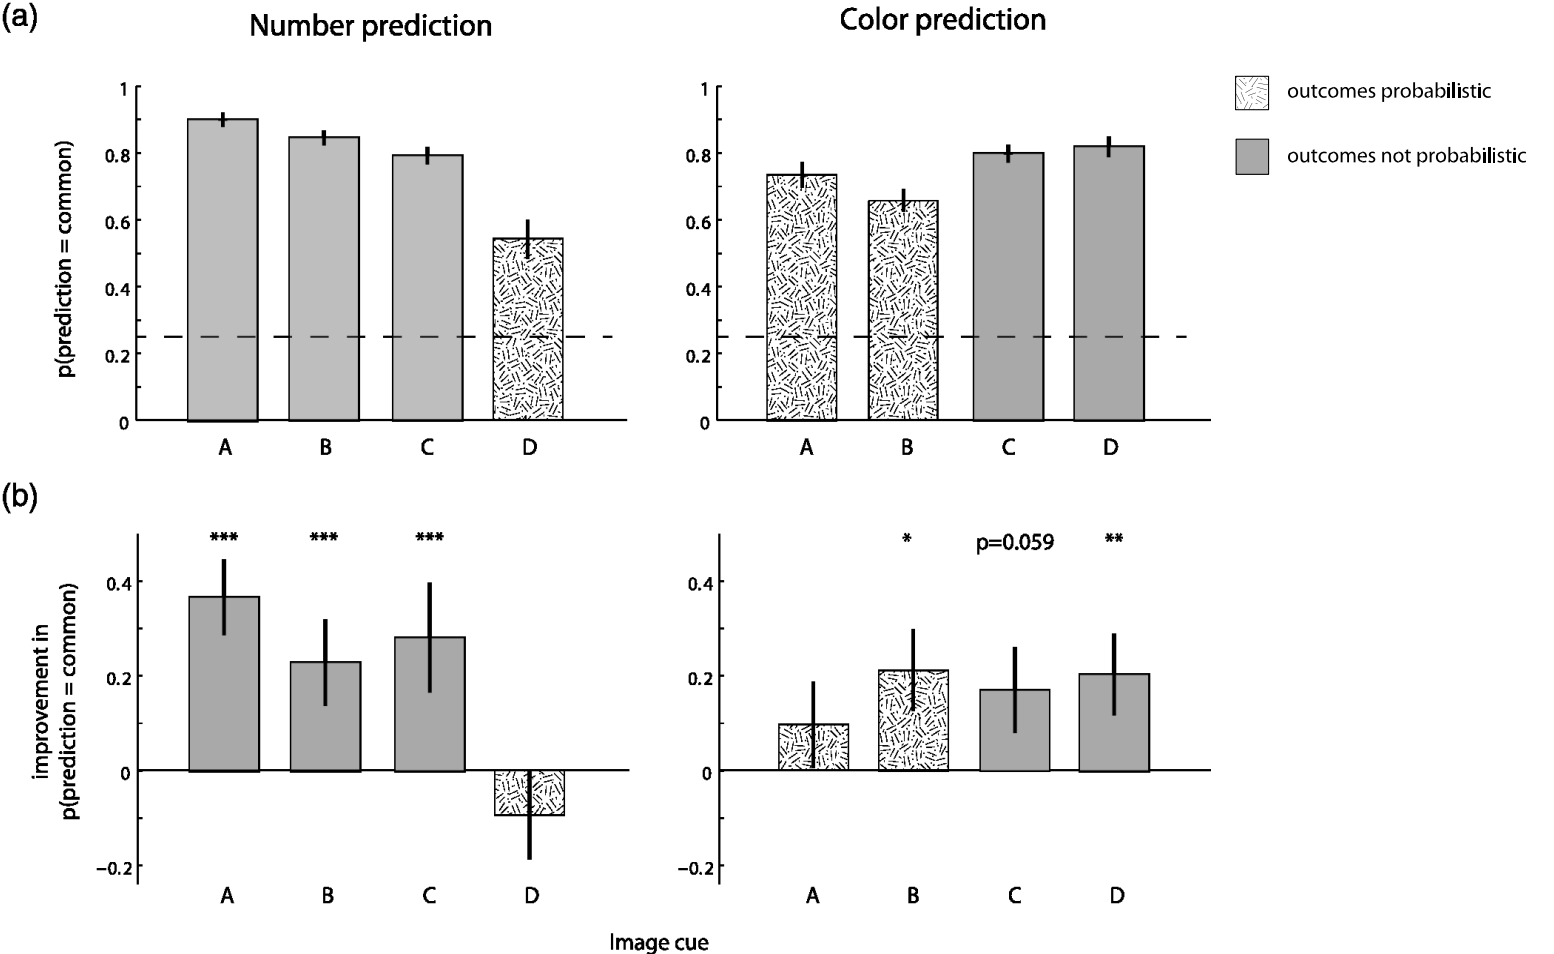
Behavioral Performance

For the prediction task, the optimal strategy was to predict the most common outcome on every trial, given that this would maximize subjects' payout. Overall, subjects predicted the most common outcome $77 \pm 2\%$ of the time. The 23% nonoptimal guesses may have resulted from a combination of probability matching (for probabilistic transitions, Erev & Barron, 2005; Vulkan, 2000), imperfect knowledge of transition probabilities, and noise. Figure 2 a shows subjects' performance on each trial type. Subjects performed significantly above chance for all trial types ($p < 10^{-6}$; one-sided bootstrap test).

Overall Learning Across the Experiment

Subjects became predicted the more common outcome more often as the experiment progressed, as measured by the difference between performance on the last scan session compared to performance during the training session (before entering the scanner) (Figure 2b). The only exception was in predicting the number of M&Ms for Cue D. Here, the optimal prediction was 1 M&M; however, participants predicted this amount on only around half the prediction trials and predicted the rare 4 M&Ms otherwise, possibly because of the high salience and appeal of the 4 M&Ms outcome. That is, although the 4 M&M outcome was delivered on only 1/3 of the trials involving Cue D, participants may have been confused regarding its frequency, or they may have predicted 4 M&Ms as a form of "wishful thinking." Over the course of the

Figure 2
Behavioral Performance Evidenced Learning Throughout the Task



Note. Hatched bars indicate that the outcomes were probabilistic for that cue and dimension (i.e., cue D for number, and cues A and B for color). Error bars indicate standard error of the mean. (a) Subjects performed above chance in all prediction tasks. plotted is the probability of choosing the more common outcome (the optimal prediction) for number prediction trials and color prediction trials across the whole experiment. Dashed line: chance. (b) Subjects' performance improved significantly across the experiment, except for color prediction for cue C (which only improved numerically) and number prediction for cue D. Plotted is the difference in probability of choosing the more common outcome in the last session compared to the training session. positive differences indicate learning during the task.
* $p < .05$. ** $p < .01$. *** $p < .0001$.

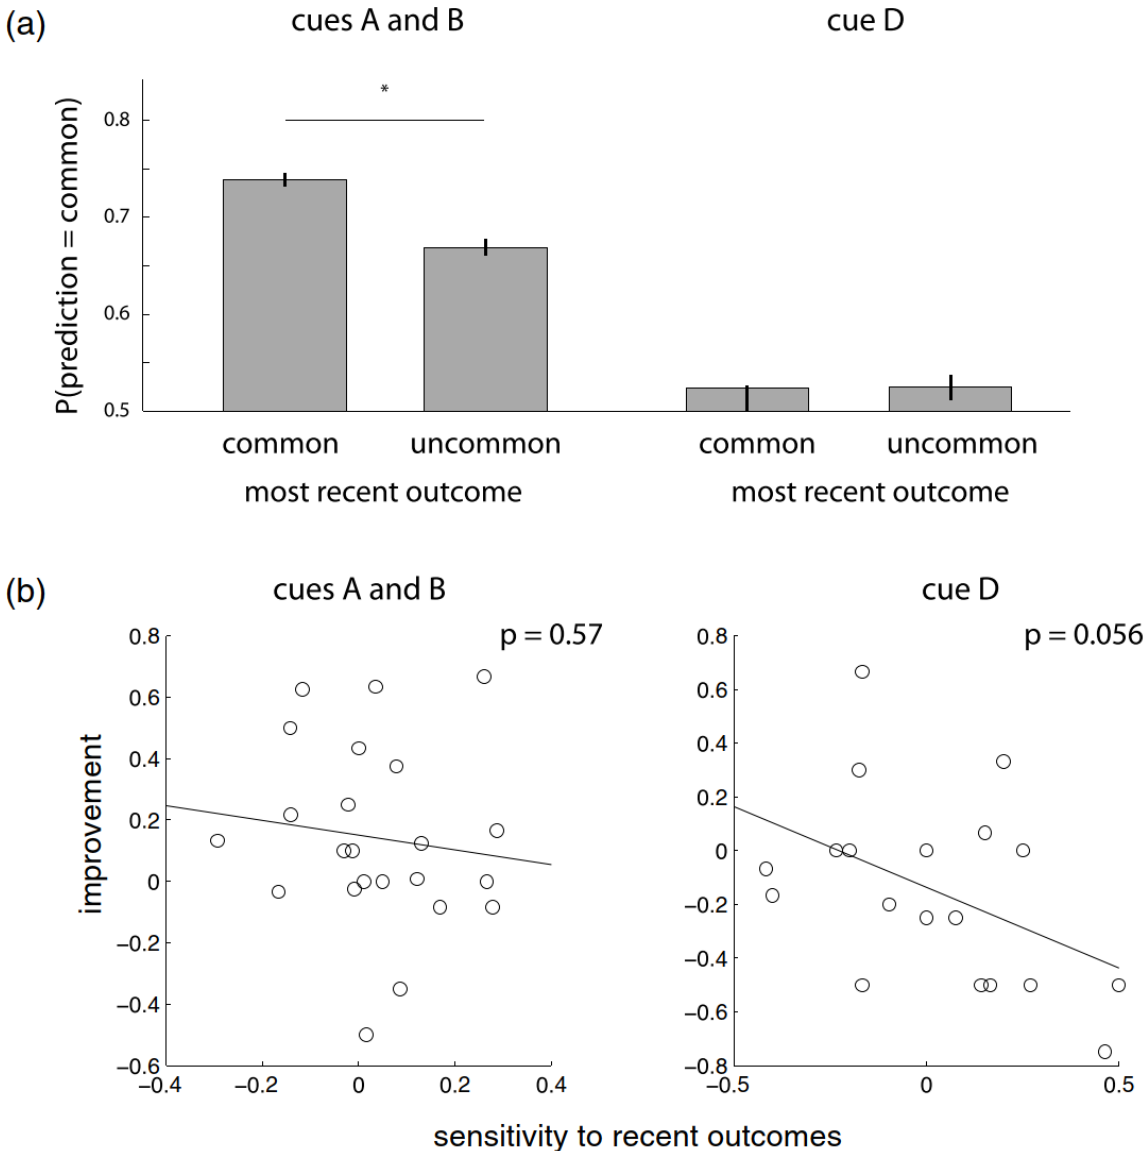
task, predictions of the outcome of this cue did not improve and even got worse numerically (Figure 2b).

Learning From Recent Outcomes

We evaluated each subject's sensitivity to the most recent outcome as a behavioral proxy for learning rate—a subject with a high learning rate should be relatively more likely to expect an outcome that she recently experienced, while a subject with a low learning rate should be less affected by recent experience. To measure this, we compared the probability of the subject predicting the common outcome for a specific cue after most recently experiencing the common outcome for that cue, versus after most recently experiencing the uncommon outcome. Stronger sensitivity to the most recent outcome, that is, higher learning rates, should manifest as larger differences between the two quantities. We evaluated learning for the scan sessions, as these were the sessions for which we could correlate learning with brain activity.

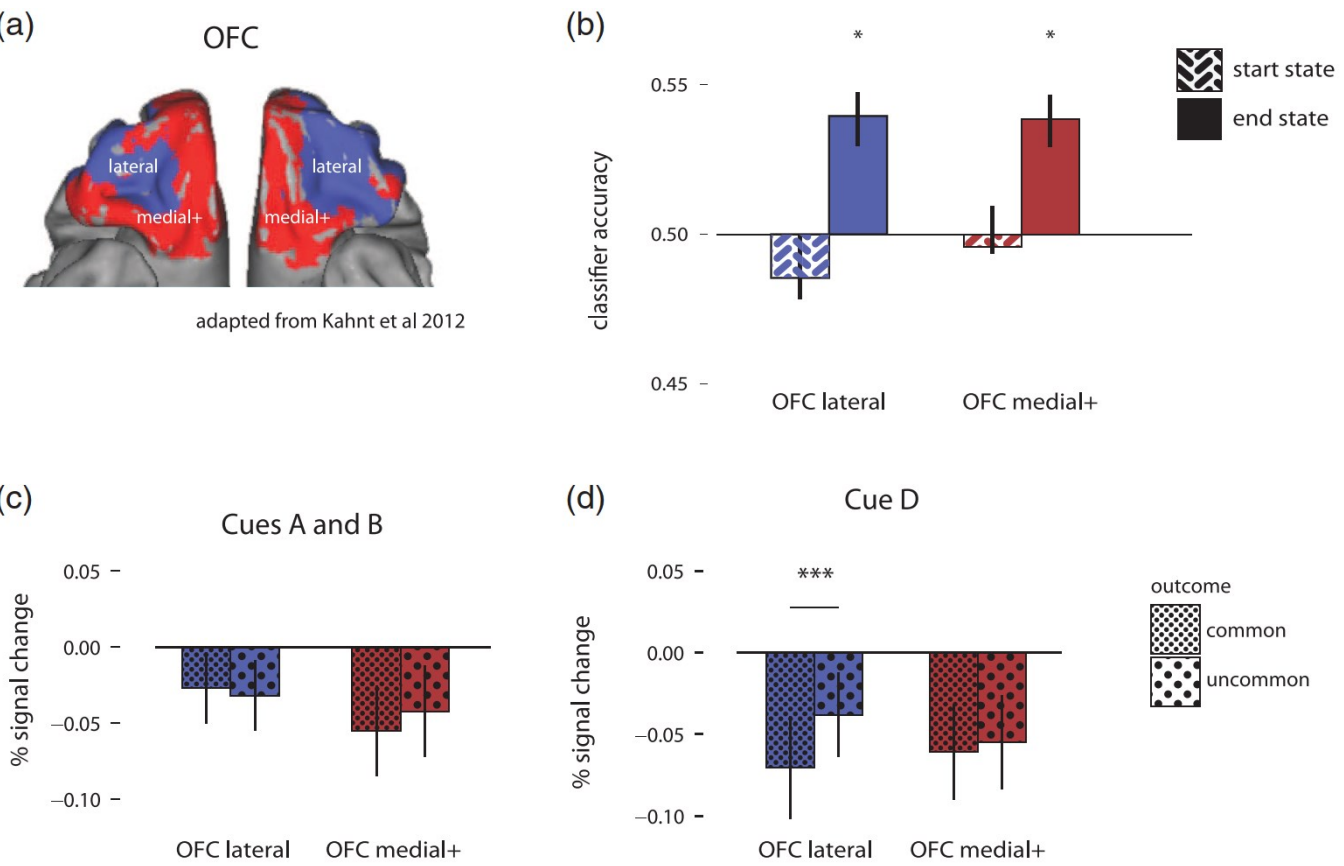
For color prediction on Cue A and Cue B trials, subjects showed significantly greater probability of choosing the common outcome if the most recent outcome was common, suggesting that subjects were learning about Cue A and B outcomes from experience during the scan sessions (Figure 3a, left). This pattern of learning was not apparent for Cue D number prediction trials, consistent with the low overall accuracy and low improvement across the experiment for predicting the number of M&Ms for Cue D (Figure 3a, right).

Figure 3
Subjects' Predictions Were Sensitive to the Most Recent Outcomes for Cues A and B, Evidencing Trial-by-Trial Learning



Note. (a) Subjects were more likely to predict the common color (for cue A and B trials, where color of the M&M outcome varied; left) after observing the common outcome (as opposed to the uncommon outcome) on the most recent trial with the same image cue. This pattern did not hold for predictions of number in the condition where number of M&Ms varied (number prediction on cue D trials; right). Means \pm SEM. (b) There was no correlation between sensitivity to recent outcomes (calculated as the difference between the probability of predicting the common outcome after recently observing the common outcome versus the uncommon outcome for the same cue; see panel a) and performance improvement across the experiment (calculated as the difference in proportion of optimal predictions between the last session and the training session; see Figure 2b).

Figure 4
M&M Outcomes, But Not Image Cues, Were Classifiable From Multivariate OFC Patterns of Activity



Note. (a) Subregions of OFC, displayed on the orbital surface of the brain. These regions of interest were obtained on a different dataset by Kahnt et al. (2012), who parcellated the OFC using k-means clustering of functional connectivity. (b) cross-validated classification performance (across subjects) for start state (image cue) and end state (M&M color) for cue A and B trials, using multivariate linear classifiers on OFC activity. error bars indicate SEM. * $p < .05$. (c and d) No significant differences in univariate BOLD responses for common versus uncommon outcomes were observed, except for cue D (for which uncommon outcomes also corresponded to larger rewards) in lateral OFC. Plotted is percent signal change in subregions of OFC at the time of the common outcomes and the uncommon outcomes. *** $p < .005$.

Note that higher sensitivity to recent outcomes does not necessarily imply a greater improvement in performance across the experiment, because high learning rates can, in fact, lead to more highly fluctuating responses. Indeed, as shown in Figure 3b, we did not find evidence for a correlation between sensitivity to recent outcomes and improvement across the experiment in Cue A and B color prediction, and we found a marginally negative correlation with improvement in Cue D number prediction.

Identity of Outcomes (But Not of Image Cues) was Decodable From Multivariate Orbitofrontal Cortex Activity—Orbitofrontal Cortex Does Not Simply Represent Perceptual Input

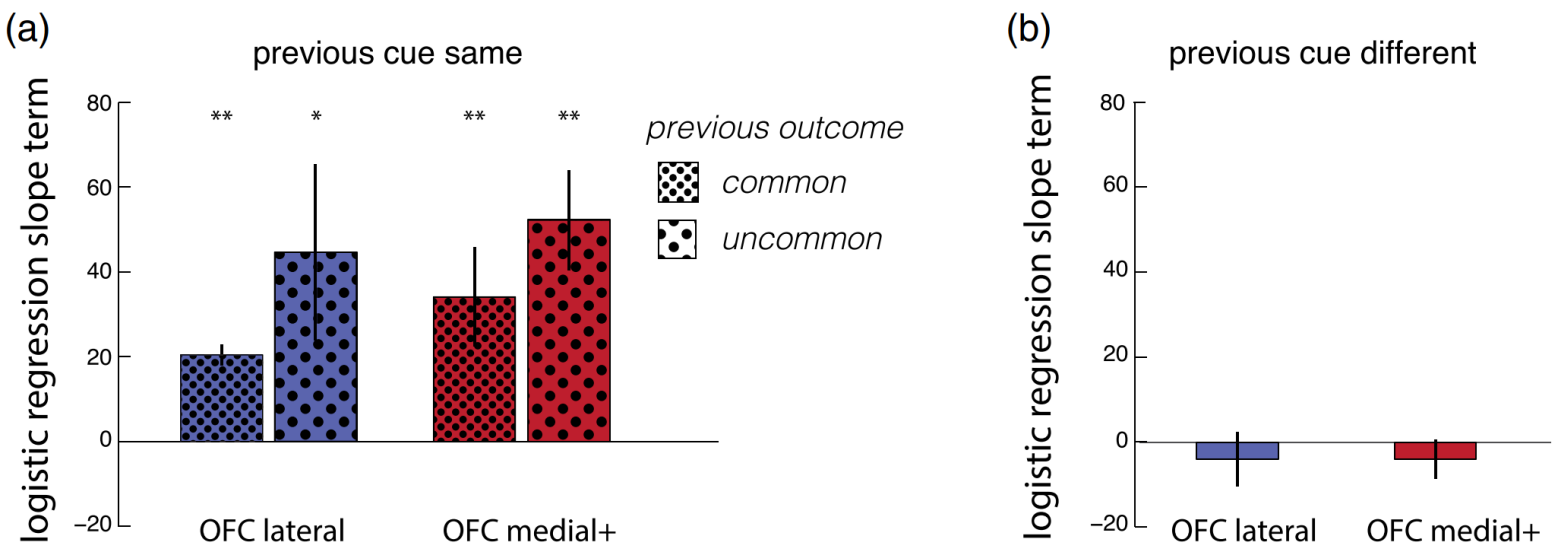
To evaluate OFC representations of the current state, we used multivariate classification methods to classify the outcome states (Color 1 vs. Color 2) at the time of the M&M outcome for Cue A and Cue B trials. We analyzed a predefined OFC ROI (Figure 4a). Cross-validated classifier performance was significantly above chance (50%) for classifying M&M outcome (classification accuracy 53.9% and 54.0%, $p = .013$ and $.012$, for medial and lateral OFC respectively; one-sided bootstrap test), indicating reliable representations of outcome state in both medial and lateral OFC (Figure 4b). In contrast, we did not find above-chance classifier performance for the image cue (Cue A vs. Cue B) at the time of the outcome (classification accuracy 49.6% and 48.5%, $p = .61$ and $.75$, for medial and lateral OFC respectively; one-sided bootstrap test). This is despite the fact that, on each trial, the image cue was still on the screen at the time that the M&M outcome appeared, and in fact occupied a much larger area of the screen than the M&Ms, indicating that OFC representations of the current state do not simply reflect perceptual input.

Univariate Orbitofrontal Cortex Responses at the Time of Outcome Did Not Signal State Prediction Errors

In general, we did not observe significant differences in univariate BOLD responses for common versus uncommon outcomes (corresponding to hypothesized small vs. large prediction errors). The exception was in lateral OFC for Cue D, where the BOLD response was more negative for the common (1 M&M) outcome as compared to the uncommon 4 M&M outcome (Figure 4c, d; $p < .005$, one-sided bootstrap test), suggesting possible sensitivity to reward value or salience in lateral OFC in particular.

Figure 5

Greater OFC Activity at the Time of Both Common and Uncommon Outcomes Was Related to Higher Probability of Predicting the Common Outcome on Ensuing Trials With the Same Cue



Note. (a) Within-subject, trial-by-trial correlations of OFC activity with learning from recent outcomes, for cue A and B trials. Plotted is the mean slope term from logistic regression of [probability of predicting the common outcome] on [percent signal change in OFC subregion at the time of outcome, for the most recent trial with the same image cue], fit for each subject separately, and also separately for trials where the previous outcome was the common outcome or where the previous outcome was the uncommon outcome. * $p < .05$. ** $p < .01$. (b) OFC activity at outcomes was not related to the probability of predicting the common outcome on ensuing trials with a different cue. Plotted is the mean slope term from logistic regression of [probability of predicting the common outcome] on [percent signal change in OFC subregion at previous outcome, for the most recent trial with a different image cue], fitted for each subject separately bars indicate mean slope terms across subjects \pm SEM.

Trial-by-Trial Correlations of Orbitofrontal Cortex Activity With Learning From the Most Recent outcome

Given that subjects' behavior demonstrated learning from the most recent outcome for Cues A and B during the scan sessions (Figure 3a, Section "Learning From Recent Outcomes"), we evaluated whether OFC activity could predict this learning, on a trial-by-trial basis. For each subject, we used logistic regression on OFC activity at the time of an outcome to predict whether the subject would choose the common outcome in the subsequent "guess" trial involving the same cue. The results of this analysis indicated an involvement of OFC in learning about transitions, but they were not consistent with a straightforward interpretation of OFC activity as reflecting a state prediction error.

Based on a prediction-error account of OFC, we would expect that the slope term of the logistic regression would be positive for the common outcomes and negative for the uncommon outcomes—that is, if greater OFC activity at the time of an outcome indicates a larger prediction error (and therefore more learning to predict that outcome), then it should lead to a greater probability of the subject predicting the *same* outcome on the next trial (common after observing a common outcome, and uncommon after observing an uncommon outcome). Instead, we found that the fit slope terms were positive for both trials where the most recent outcome was the common outcome, as well as for trials where the most recent outcome was the uncommon outcome. In other words, no matter the outcome (common or uncommon), greater BOLD activity in OFC at the time of an outcome was correlated with a higher probability of subjects predicting the *common* outcome on the next trial with the same cue (Figure 5a). These results suggest the OFC's involvement on a trial-by-trial basis in learning task contingencies.

To investigate the possibility that the result in Figure 5a was driven by OFC activity simply indicating an increased level of task engagement or task structure representation, we performed a second, analogous analysis in which we regressed subjects' predictions on OFC activity at the most recent outcome involving a *different* cue. Note that this regression predicts choice from trials that were, on average, temporally closer than the trials used in the above analysis, since different-cue trials are more common ($15.0 \pm .4$ s between subsequent different-cue trials, 38.6 ± 1.0 s between subsequent same-cue trials). As a result, we might expect the relationship between OFC activity and choice to be even stronger in this analysis. Contrary to this prediction, we found that the regression slope terms were not significantly different from zero, as shown in Figure 5b. Thus OFC activity at the time of an outcome only seemed to affect behavioral predictions on ensuing trials when those trials involved the *same* cue. Together, these analyses suggest that the OFC is involved in updating or reinforcing a learned transition model on a trial-by-trial basis, specifically for the currently observed cue–outcome associations.

Discussion

Lesion and inactivation studies have shown that the OFC is particularly important for planning and model-based decisionmaking. However, prior work implicating OFC in the representation of expected values did not necessarily explain why this area should be important for model-based decision-making in particular, since expected values are generated in both model-based and model-free learning. Here, we have shown that OFC activity at the time of an outcome is related to learning about the transition structure of a task (the tendencies of certain states to lead to other states), which

is necessary for accurate planning, shedding new light on the question of why the OFC is critical for model-based decisions.

Using an experimental design that permits constant updating of (probabilistic) transitions between states, we showed that activity in the OFC is correlated with behavioral measures of learning about transition structure. Greater OFC activity at the time of an outcome was positively correlated with the likelihood of optimally predicting the outcome on the next trial with the same image cue, supporting the hypothesized involvement of the OFC in learning of transition structure.

State-transition learning in our experiment was distinct from value-based learning that is thought to be implemented in the dopaminergic system (Jocham et al., 2011; Kravitz et al., 2012), because the trials of interest always led to a predictable number of 2 M&Ms. Our analyses also combined conditions (Cue A and Cue B trials) in which the identities (M&M colors) of the common and uncommon outcomes were reversed, so that any potential differences in value for different M&M colors would cancel out. Therefore, our results positively identify a role for the OFC in learning a nonvalue-related quantity.

In particular, our results suggest a role for OFC in the activation or reinforcement of an already-learned transition model (and thus the reinforcement of the more optimal prediction), which may be in line with previous work indicating OFC representation of a cognitive map of task space (Chan et al., 2016; Schuck et al., 2016; Schuck & Niv, 2019; Wilson et al., 2014). At the same time, our results are agnostic as to where exactly this model is stored or accessed from—they show only the involvement of OFC in *updating* or *reinforcing* an internal model of the transition structure.

What kind of learning might underlie the observed relationships between OFC activity and behavior? Both Hebbian and error-based learning might lead to the pattern of results we observed, though these algorithms would require additional signals to drive the learning process (e.g., from a learned model), in order for increased OFC activity to strengthen the relationship between a cue and the common outcome, rather than the currently observed outcome. Other, yet unknown algorithms, may of course also be possible—in contrast to our increasingly detailed understanding of how world models are used for deploying goal-directed behavior (Daw, 2018), the mechanism for learning world models has not been well described in the literature as of yet (Niv, 2019; O'Doherty et al., 2017).

Previous work has proposed a state prediction-error algorithm for learning state transitions, analogous to learning about state values from reward prediction errors observed in dopaminergic neurons. A state prediction error would signal surprise-related information at the time of an unexpected state (regardless of the state's value) and could be used to adjust internal estimates of transition probabilities by increasing the estimated probability of transitioning to the just-observed state. However, our results did not straightforwardly affirm the hypothesis that the OFC supports state-transition learning by representing a state prediction error signal. At the least, we did not find evidence for the encoding of this error signal in the OFC's univariate response to outcomes, as we did not observe differences in OFC activity for common versus uncommon outcomes (corresponding to small vs. large state prediction errors). Further, univariate OFC activity at the time of an outcome was not correlated with greater subsequent expectations of that particular outcome. Instead, OFC activity was related to the greater subsequent expectation of the *more common* outcome (i.e., more optimal prediction by subjects), regardless of whether the OFC activity was measured at the time of a common or uncommon outcome.

Previous works have found mixed evidence of state prediction errors being represented in the OFC. Howard and Kahnt (2018) and Suarez et al. (2019) discovered evidence of state prediction errors in the midbrain, OFC, piriform cortex, amygdala, lateral prefrontal cortex, and posterior parietal cortex. Boorman et al. (2016) found correlations of state prediction errors with univariate activity in lateral OFC, but only when signed positively or negatively according to whether the update increased or decreased the odds of a preferred outcome. However, Glascher et al. (2010) tested for univariate correlations with the (unsigned) magnitude of an inferred state prediction error signal and implicated the dorsolateral prefrontal cortex and intraparietal sulcus (but not the OFC) in their analysis.

It is possible that the discrepancy between our findings and previous work showing state prediction errors in OFC is due to the specific differences between our experiment and previous work—differences that perhaps highlight the specific circumstances required for observation of state prediction errors in the OFC. First, state transition probabilities in the current experiment (and in Glascher et al., 2010) were stationary (though stochastic), whereas the aforementioned previous works used nonstationary distributions with sudden reversals in cue–outcome contingencies. Recent work has indicated that humans and animals interpret large changes in the environment differently from small surprises, with large surprises taken as indicating a new underlying “latent state” of the environment (Collins & Frank, 2013; Dunsmoor et al., 2015; Gershman et al., 2010, 2017). The process of latent state inference and representation has been associated with both the hippocampus and the OFC (Niv, 2019; Schuck & Niv, 2019).

A second difference is that our main analyses averaged over any differences in value for different outcomes by aggregating across Cue A and Cue B, in order to disentangle reward prediction errors and state prediction errors. Glascher et al. also analyzed state prediction errors only in the absence of rewards and did not find evidence for state prediction errors in OFC. Indeed, for the trial type where we purposefully did not take these measures to disentangle state and reward prediction errors (Cue D trials, where the image cues led probabilistically to varying numbers—rather than colors—of M&Ms), we did find evidence of univariate differences in activation of lateral OFC for uncommon vs. common (i.e., high vs. low value) outcomes. This is reminiscent of the results from Boorman et al. (2016), who found correlations with state prediction errors in lateral OFC only when the inferred errors were signed positively or negatively according to the change in the value of a state transition. Howard and Kahnt (2018) also showed that responses to (unsigned) identity and positive value prediction errors were correlated across participants, at least in the midbrain. These results suggest some kind of shared neural response between state and reward prediction errors in these studies.

Previous work in rats that implicated the OFC in learning about transition structure, concentrated on the lateral OFC (McDannald et al., 2011), and fMRI work in humans also specifically implicated the lateral OFC in this type of process (Boorman et al., 2016; Howard & Kahnt, 2018; Suarez et al., 2019). We tested our hypotheses in the entirety of the OFC, using a previously determined functional connectivity-based parcellation of OFC into medial and lateral subregions (Kahnt et al., 2012). Medial and lateral OFC

showed very similar results across all our analyses. Of course, this does not rule out the possibility that there may exist a different parcellation of OFC that would lead to differing results across subregions. We note also that the homology of OFC between rodents and humans is currently unclear, and OFC subdivisions are particularly complex given observed considerable anatomical variability within individuals (Chiavaras & Petrides, 2000; Wallis, 2012). We should also take care in interpreting the negative BOLD response in OFC—this negative BOLD response has been previously observed (e.g., Boorman et al., 2009), but is not yet fully understood.

It is important to note that our secondary analyses did provide further support for two other mainstream theories of OFC function, in addition to the theory of a role in learning transition structure. First, we found that we could decode representation of outcome identity in OFC using MVPA of BOLD activity at the times of the outcomes, consistent with a recent theory implicating OFC in the representation of the current state (Wilson et al., 2014), for which evidence is increasingly amassing (e.g., Bradfield et al., 2015; Chan et al., 2016; Howard et al., 2020; Klein-Flügge et al., 2013; Nogueira et al., 2017; Schuck et al., 2016; Zhou et al., 2020). Second, we found evidence for value sensitivity in univariate BOLD responses in lateral OFC in a separate task condition (Cue D), in which the number (but not color) of M&Ms was unpredictable, consistent with previous work demonstrating that OFC represents the value of rewards (Fellows, 2007; Gottfried et al., 2003; Hampton et al., 2006; Hare et al., 2008; Monosov & Hikosaka, 2012; Padoa-Schioppa & Assad, 2006; Wallis & Kennerley, 2011; though note that value might be construed as just one feature in the representation of the current state; Lopatina et al., 2015).

In conclusion, the present results provide support for an emerging understanding of the relationship between the OFC and acquisition of state-to-state transition structure. Our findings may suggest a role for OFC in reactivation and reinforcement of an already learned state-transition model, relating to proposals that the OFC stores such a model (Wilson et al., 2014). Our findings also build upon previous work showing that rats with OFC lesions show impaired learning about changes in state transitions (McDannald et al., 2011) and that surprise signals in human OFC are related to changes in hippocampal representations of state transitions (Boorman et al., 2016). Importantly, while our results are not straightforwardly aligned with a state prediction error hypothesis, at least in this learning regime, they do still indicate OFC's involvement in learning about transition structure. Our findings may therefore serve to constrain future models of the particular learning algorithms that may underlie learning about transition structure, facilitating a fuller understanding of the involvement of OFC in learning and model-based decision-making.

References

- Baxter, M. G., Parker, A., Lindner, C. C. C., Izquierdo, A. D., & Murray, E. A. (2000). Control of response selection by reinforcer value requires interaction of amygdala and orbital prefrontal cortex. *The Journal of Neuroscience*, *20*, 4311–4319.
- Boorman, E. D., Behrens, T. E. J., Woolrich, M. W., & Rushworth, M. F. S. (2009). How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. *Neuron*, *62*, 733–743.
- Boorman, E. D., Rajendran, V. G., O'Reilly, J. X., & Behrens, T. E. (2016). Two anatomically and computationally distinct learning signals predict changes to stimulus-outcome associations in hippocampus. *Neuron*, *89*, 1343–1354.
- Bradfield, L. A., Dezfouli, A., van Holstein, M., Chieng, B., & Balleine, B. W. (2015). Medial orbitofrontal cortex mediates outcome retrieval in partially observable task situations. *Neuron*, *88*, 1268–1280.
- Carmichael, S. T., Price, J. L. (1996) Connectional networks within the orbital and medial prefrontal cortex of macaque monkeys. *Journal of Comparative Neurology* *371*, 179–207.
- Chan, S. C. Y., Niv, Y., & Norman, K. A. (2016). A probability distribution over latent causes, in the orbitofrontal cortex. *The Journal of Neuroscience*, *36*(30), 7817–7828. <https://doi.org/10.1523/JNEUROSCI.0659-16.2016>
- Chang, C.-C., & Lin, C.-J. (2011) LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology* *2*(3), Article 27.
- Chiavaras, M. M., & Petrides, M. (2000). Orbitofrontal sulci of the human and macaque monkey brain. *Journal of Comparative Neurology*, *422*(1), 35–54.
- Collins, A. G. E., & Frank, M. J. (2013). Cognitive control over learning: Creating, clustering, and generalizing task-set structure. *Psychological Review*, *120*(1), 190–229. <https://doi.org/10.1037/a0030852>
- Daw, N. D. (2018). Are we of two minds? *Nature Neuroscience*, *21*(11), 1497–1499.
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, *8*, 1704–1711.
- De Wit, S., Corlett, P. R., Aitken, M. R., Dickinson, A., & Fletcher, P. C. (2009). Differential Engagement of the ventromedial prefrontal cortex by goal-directed and habitual behavior toward food pictures in humans. *The Journal of Neuroscience*, *29*, 11330–11338.
- Deichmann, R., Gottfried, J., Hutton, C., & Turner, R. (2003). Optimized EPI for fMRI studies of the orbitofrontal cortex. *NeuroImage*, *19*, 430–441.
- Dunsmoor, J. E., Niv, Y., Daw, N., & Phelps, E. A. (2015). Rethinking extinction. *Neuron*, *88*(1), 47–63. <https://doi.org/10.1016/j.neuron.2015.09.028>
- Erev, I., & Barron, G. (2005). On Adaptation, Maximization, and Reinforcement Learning Among Cognitive Strategies. *Psychological Review*, *112*(4), 912–931. <https://doi.org/10.1037/0033-295X.112.4.912>
- Fellows, L. K. (2007). The role of orbitofrontal cortex in decision making. *Annals of the New York Academy of Sciences*, *1121*, 421–430.
- Gershman, S. J., Blei, D. M., & Niv, Y. (2010). Context, learning, and extinction. *Psychological Review*, *117*(1), 197–209. <https://doi.org/10.1037/a0017808>
- Gershman, S. J., Monfils, M. H., Norman, K. A., & Niv, Y. (2017). The computational nature of memory modification. *eLife*, *6*, Article e23763.
- Glascher, J., Daw, N., Dayan, P., & O'Doherty, J. P. (2010). States versus rewards: Dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron*, *66*, 585–595.
- Gottfried, J. A., O'Doherty, J., & Dolan, R. J. (2003). Encoding predictive reward value in human amygdala and orbitofrontal cortex. *Science*, *301*, 1104–1107.
- Hampton, A. N., Bossaerts, P., & O'Doherty, J. P. (2006). The Role of the Ventromedial Prefrontal Cortex in Abstract State-Based Inference during Decision Making in Humans. *The Journal of Neuroscience*, *26*(32), 8360–8367. <https://doi.org/10.1523/JNEUROSCI.1010-06.2006>
- Hare, T. A., O'Doherty, J., Camerer, C. F., Schultz, W., & Rangel, A. (2008). Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. *The Journal of Neuroscience*, *28*, 5623–5630.
- Howard, J. D., & Kahnt, T. (2018). Identity prediction errors in the human midbrain update reward-identity expectations in the orbitofrontal cortex. *Nature Communications*, *9*(1), Article 1611.

- Howard, J. D., Reynolds, R., Smith, D. E., Voss, J. L., Schoenbaum, G., & Kahnt, T. (2020). Targeted stimulation of human orbitofrontal networks disrupts outcome-guided behavior. *Current Biology*, 30(3), 490–498.e4. <https://doi.org/10.1016/j.cub.2019.12.007>
- Izquierdo, A., Suda, R. K., & Murray, E. A. (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.
- Jocham, G., Klein, T. A., & Ullsperger, M. (2011). Dopamine-mediated reinforcement learning signals in the striatum and ventromedial prefrontal cortex underlie value-based choices. *The Journal of Neuroscience*, 31, 1606–1613.
- Kahnt, T., Chang, L. J., Park, S. Q., Heinzle, J., & Haynes, J.-D. (2012). Connectivity-based parcellation of the human orbitofrontal cortex. *The Journal of Neuroscience*, 32, 6240–6250.
- Klein-Flügge, M. C., Barron, H. C., Brodersen, K. H., Dolan, R. J., & Behrens, T. E. J. (2013). Segregated encoding of reward–identity and stimulus–reward associations in human orbitofrontal cortex. *The Journal of Neuroscience*, 33, 3202–3211.
- Kravitz, A. V., Tye, L. D., & Kreitzer, A. C. (2012). Distinct roles for direct and indirect pathway striatal neurons in reinforcement. *Nature Neuroscience*, 15, 816–818.
- Lopatina, N., McDannald, M. A., Styer, C. V., Sadacca, B. F., Cheer, J. F., & Schoenbaum, G. (2015). Lateral orbitofrontal neurons acquire responses to upshifted, downshifted, or blocked cues during unblocking. *eLife*, 4, Article e11299.
- McDannald, M. A., Lucantonio, F., Burke, K. A., 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.
- Monosov, I. E., & Hikosaka, O. (2012). Regionally distinct processing of rewards and punishments by the primate ventromedial prefrontal cortex. *The Journal of Neuroscience*, 32, 10318–10330.
- Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *The Journal of Neuroscience*, 16, 1936–1947.
- Mumford, J. A., Turner, B. O., Ashby, F. G., & Poldrack, R. A. (2012). Deconvolving BOLD activation in event-related designs for multivoxel pattern classification analyses. *NeuroImage*, 59, 2636–2643
- Niv, Y. (2019). Learning task-state representations. *Nature Neuroscience*, 22(10), 1544–1553.
- Nogueira, R., Abolafia, J. M., Drugowitsch, J., Balaguer-Ballester, E., Sanchez-Vives, M. V., & Moreno-Bote, R. (2017). Lateral orbitofrontal cortex anticipates choices and integrates prior with current information. *Nature Communications*, 8, Article 14823.
- O’Doherty, J. P., Cockburn, J., & Pauli, W. M. (2017). Learning, reward, and decision making. *Annual Review of Psychology*, 68(1), 73–100. <https://doi.org/10.1146/annurev-psych-010416-044216>
- Öngür, D., & Price, J. L. (2000). The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. *Cerebral Cortex*, 10, 206–219.
- Padoa-Schioppa, C., & Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature*, 441, 223–226.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). Appleton-Century-Crofts.
- Rudebeck, P. H., & Murray, E. A. (2011). Dissociable effects of subtotal lesions within the macaque orbital prefrontal cortex on reward-guided behavior. *The Journal of Neuroscience*, 31, 10569–10578.
- Schuck, N. W., Cai, M. B., Wilson, R. C., & Niv, Y. (2016). Human orbitofrontal cortex represents a cognitive map of state space. *Neuron*, 91, 1402–1412.
- Schuck, N. W., & Niv, Y. (2019). Sequential replay of nonspatial task states in the human hippocampus. *Science*, 364(6447), Article eaaw5181. <https://doi.org/10.1126/science.aaw5181>
- Sharpe, M. J., Chang, C. Y., Liu, M. A., Batchelor, H. M., Mueller, L. E., Jones, J. L., Niv, Y., & Schoenbaum, G. (2017). Dopamine transients are sufficient and necessary for acquisition of model-based associations. *Nature Neuroscience*, 20(5), 735–742. <https://doi.org/10.1038/nn.4538>
- Stalnaker, T.A., Howard, J.D., Takahashi, Y.K., Gershman, S. J., Kahnt, T., & Schoenbaum, G. (2019). Dopamine neuron ensembles signal the content of sensory prediction errors. *eLife*, 8, Article e49315. <https://doi.org/10.7554/eLife.49315>
- Suarez, J. A., Howard, J. D., Schoenbaum, G., & Kahnt, T. (2019). Sensory prediction errors in the human midbrain signal identity violations independent of perceptual distance. *eLife*, 8, Article e43962. <https://doi.org/10.7554/eLife.43962>
- Takahashi, Y. K., Batchelor, H. M., Liu, B., Khanna, A., Morales, M., & Schoenbaum, G. (2017). Dopamine neurons respond to errors in the prediction of sensory features of expected rewards. *Neuron*, 95(6), 1395–1405.e3. <https://doi.org/10.1016/j.neuron.2017.08.025>
- Valentin, V. V., Dickinson, A., & O’Doherty, J. P. (2007). Determining the neural substrates of goal-directed learning in the human brain. *The Journal of Neuroscience*, 27, 4019–4026.
- Vulkan, N. (2000). An economist’s perspective on probability matching. *Journal of Economic Surveys*, 14(1), 101–118.
- Wallis, J. D. (2012). Cross-Species studies of orbitofrontal cortex and valuebased decision-making. *Nature Neuroscience*, 15(1), 13–19. <https://doi.org/10.1038/nn.2956>
- Wallis, J. D., & Kennerley, S. W. (2011). Contrasting reward signals in the orbitofrontal cortex and anterior cingulate cortex. *Annals of the New York Academy of Sciences*, 1239, 33–42.
- Walton, M. E., Behrens, T. E. J., Buckley, M. J., Rudebeck, P. H., & Rushworth, M. F. S. (2010). Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. *Neuron*, 65, 927–939.
- Wilson, R. C., Takahashi, Y. K., Schoenbaum, G., & Niv, Y. (2014). Orbitofrontal cortex as a cognitive map of task space. *Neuron*, 81, 267–279.
- Zhou, J., Zong, W., Jia, C., Gardner, M. P. H., & Schoenbaum, G. (2020). Prospective representations in rat orbitofrontal ensembles. *Neuroscience*. Advance online publication. <https://doi.org/10.1101/2020.08.27.268391>