



Published in final edited form as:

Nat Neurosci. ; 14(12): 1590–1597. doi:10.1038/nn.2957.

Expectancy-related changes in firing of dopamine neurons depend on orbitofrontal cortex

Yuji K. Takahashi¹, Matthew R. Roesch^{2,3}, Robert C. Wilson⁴, Kathy Toreson¹, Patricio O'Donnell^{1,5}, Yael Niv^{4,*}, and Geoffrey Schoenbaum^{1,5,6,*}

¹Department of Anatomy and Neurobiology, University of Maryland School of Medicine, 20 Penn St, HSF-2 S251, Baltimore, MD 21201

²Department of Psychology, University of Maryland College Park, College Park, MD 20742

³Program in Neuroscience and Cognitive Science, University of Maryland College Park, College Park, MD 20742

⁴Department of Psychology and Neuroscience Institute, Princeton University, Princeton, NJ 08544

⁵Department of Psychiatry, University of Maryland School of Medicine, 20 Penn St, HSF-2 S251, Baltimore, MD 21201

⁶NIDA Intramural Research Program, 251 Bayview Drive, Baltimore, MD 21224

Abstract

The orbitofrontal cortex has been hypothesized to carry information regarding the value of expected rewards. Such information is essential for associative learning, which relies on comparisons between expected and obtained reward for generating instructive error signals. These error signals are thought to be conveyed by dopamine neurons. To test whether orbitofrontal cortex contributes to these error signals, we recorded from dopamine neurons in orbitofrontal-lesioned rats performing a reward learning task. Lesions caused marked changes in dopaminergic error signaling. However, the effect of lesions was not consistent with a simple loss of information regarding expected value. Instead, without orbitofrontal input, dopaminergic error signals failed to reflect internal information about the impending response that distinguished externally similar states leading to differently valued future rewards. These results are consistent with current conceptualizations of orbitofrontal cortex as supporting model-based behavior and suggest an unexpected role for this information in dopaminergic error signaling.

Midbrain dopamine neurons signal errors in reward prediction^{1–3}. These error signals are required for learning in a variety of theoretical accounts^{4–6}. By definition, calculation of these errors requires information about the value of the rewards expected in a given circumstance or ‘state’. In temporal difference reinforcement learning (TDRL) models, such learned expectations contribute to computations of prediction errors and are modified based

Users may view, print, copy, download and 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

*These authors contributed equally to this work as described in the author contributions.

on these errors. However the neural source of this expected value signal has not been established for dopamine neurons in the ventral tegmental area (VTA). Here we tested whether one contributor might be the orbitofrontal cortex (OFC), a prefrontal area previously shown to be critical for using information about the value of expected rewards to guide behavior⁷⁻¹¹.

RESULTS

To test whether OFC contributes to reward prediction errors, we recorded single-unit activity from putative dopamine neurons in the VTA in rats with ipsilateral sham (n=6) or neurotoxic lesions (n=7) of OFC. Lesions targeted the ventral and lateral orbital and ventral and dorsal agranular insular areas in the bank of the rhinal sulcus, resulting in frank loss of neurons in 33.4% (23–40%) of this layered cortical region across the seven subjects (Fig. 1c, inset). Neurons in this region fire in anticipation of an expected reward¹² and interact with VTA to drive learning in response to prediction errors¹³. Notably, sparse direct projections from this part of OFC to VTA are largely unilateral¹⁴, and neither direct nor indirect input to VTA from contralateral OFC is sufficient to support normal learning¹³. Thus ipsilateral lesions should severely diminish any influence of OFC signaling on VTA in the lesioned hemisphere, while leaving the circuit intact in the opposite hemisphere to avoid confounding behavioral deficits.

Neurons were recorded in an odor-guided choice task used previously to characterize signaling of errors and outcome expectancies^{12,13,15}. On each trial, rats responded at one of two adjacent wells after sampling one of three different odor cues at a central port (Fig. 1a). One odor signaled sucrose reward in the right well (forced-choice right), a second odor signaled sucrose reward in the left well (forced-choice left), and a third odor signaled the availability of reward at either well (free-choice). To generate errors in the prediction of rewards, we manipulated the timing (Fig. 1b; blocks 1 and 2) or size of the reward (Fig. 1b; blocks 3 and 4) across blocks of trials. This resulted in the introduction of new and unexpected rewards, when immediate or large rewards were instituted at the start of blocks 2^{sh}, 3^{bg}, 4^{bg} (Fig. 1b), and omission of expected rewards, when delayed or small rewards were instituted at the start of blocks 2^{lo} and 4sm (Fig. 1b).

As expected, sham-lesioned rats changed their choice behavior across blocks in response to the changing rewards, choosing the higher value reward more often on free-choice trials (t-test, $t_{(100)}=18.91$, $p<0.01$; Fig. 1c, inset) and responding more accurately (t-test, $t_{(100)}=10.77$, $p<0.01$; Fig. 1d) and with shorter reaction times (t-test, $t_{(100)}=13.32$, $p<0.01$; Fig. 1e) on forced-choice trials when the high value reward was at stake. Rats with unilateral OFC lesions showed similar behavior (t-test, %choice, $t_{(84)}=14.51$, $p<0.01$: %correct, $t_{(84)}=9.88$, $p<0.01$: reaction time, $t_{(84)}=8.32$, $p<0.01$: Fig. 1c–e), and direct comparisons of all three performance measures across groups revealed no significant differences (ANOVA, sham vs lesioned: %choice, $F_{(1,184)}=0.16$, $p=0.68$: %correct, $F_{(2,183)}=2.11$, $p=0.12$: reaction time, $F_{(2,183)}=2.92$, $p=0.06$).

We identified dopamine neurons in the VTA via a cluster analysis based on spike duration and amplitude ratio. Although the use of such criteria has been questioned¹⁶, the particular

analysis used here isolates neurons whose firing is sensitive to IV apomorphine¹⁵ or quinpirole¹⁷. Additionally, neurons identified by this cluster analysis are selectively activated by optical stimulation in TH-Chr2 mutants¹⁷ and exhibit reduced bursting in TH-NR2 knockouts¹⁷. While these criteria may exclude some dopamine neurons, only neurons within this cluster signaled reward prediction errors in a significant number in our prior work¹⁵.

This analysis identified 52 of 481 recorded neurons as dopaminergic in shams (Fig. 2a) and 76 of 500 as dopaminergic in OFC-lesioned rats (Fig. 2b). These neurons had spike durations and amplitude ratios that differed significantly (> 3 standard deviations) from those of other neurons. Of these, 30 in sham and 50 in OFC-lesioned rats increased firing in response to reward (compared with baseline during the inter-trial interval; t-test, $p < 0.05$; proportions did not differ in sham vs. lesioned: X^2 -test, $X^2 = 0.86$, $df = 1$, $p = 0.35$). There were no apparent effects of OFC lesions on the waveform characteristics of these neurons (Fig. 2c, t-test, amplitude ratio: $t_{(127)} = 0.53$, $p = 0.59$; duration: $t_{(127)} = 0.78$, $p = 0.43$). The average baseline activity of reward responsive and non-responsive dopamine neurons, taken during the 500ms before light onset, was also similar in the two groups (Fig. 2f; sham vs lesioned, t-test, reward-responsive dopamine neurons, $t_{(78)} = 0.49$, $p = 0.62$; reward non-responsive dopamine neurons, $t_{(46)} = 1.57$, $p = 0.12$), as was the distribution of the baseline firing (Fig. 2d and e; sham vs lesioned: Wilcoxon, reward-responsive dopamine neurons $p = 0.86$; reward non-responsive dopamine neurons, $p = 0.09$). Interestingly, non-dopaminergic neurons fired significantly more slowly in the OFC-lesioned rats (sham vs lesioned, t-test; $t_{(851)} = 3.81$, $p < 0.01$; see Fig. S2 in Supplemental).

OFC supports dopaminergic error signals

Prior work has shown that prediction-error signaling is largely restricted to reward-responsive dopamine neurons¹⁵ (see Supplemental Figures S1 and S2 for analysis of other populations). As expected, activity in these neurons in sham-lesioned rats increased in response to unexpected reward and decreased in response to omission of an expected reward. This is illustrated by the unit examples and population responses in Figure 3a, which show that neural activity increased at the start of block 2^{sh}, when a new reward was introduced (Fig. 3a, black arrows), and decreased at the start of block 2^{lo}, when an expected reward was omitted (Fig. 3a, grey arrows). For both, the change in activity was maximal at the beginning of the block and then diminished with learning.

These patterns were significantly muted in OFC-lesioned rats. Although dopamine neurons still showed phasic firing to unexpected rewards, this response was not as pronounced at the beginning of the block, nor did it change significantly with learning (Fig. 3b, black arrows). In addition, the suppression of activity normally caused by unexpected reward omission was largely abolished (Fig. 3b, grey arrows).

These effects are quantified in Figures 3c and 3d, which plot the average activity across all reward-responsive dopamine neurons in each group, on each of the first and last 10 trials in all blocks in which we delivered a new, unexpected reward (blocks 2^{sh}, 3^{bg}, 4^{bg}) or omitted an expected reward (blocks 2^{lo} and 4sm). In sham-lesioned rats, reward-responsive dopamine neurons increased firing upon introduction of a new reward and suppressed firing on

omission of an expected reward. In each case, the change in firing was maximal on the first trial and diminished significantly thereafter (Fig. 3c). Two-factor ANOVA's comparing firing to unexpected reward (or reward omission) to background firing (average firing during inter-trial intervals) revealed significant interactions between trial period and trial number in each case (reward vs background: $F_{19,532} = 4.37$, $p < 0.0001$; omission vs background: $F_{19,532} = 3.57$, $p < 0.0001$). Post hoc comparisons showed that activity on the first five trials differed significantly from background, and from activity on later trials, for both unexpected reward and reward omission (p 's < 0.01). Furthermore, the distribution of difference scores comparing each neuron's firing early and late in the block was shifted significantly above zero for unexpected reward (Fig. 3e; Wilcoxon signed-rank test, $p < 0.01$) and below zero for reward omission (Fig. 3f; Wilcoxon signed-rank test, $p < 0.01$), and there was a significant inverse correlation between changes in firing in response to unexpected reward and reward omission (Fig. 3g; $r = 0.43$, $p < 0.05$). These results are consistent with bidirectional prediction error signaling in the individual neurons at the time of reward in the sham-lesioned rats.

By contrast, the activity of reward-responsive dopamine neurons in OFC-lesioned rats did not change substantially across trials or in response to reward omission (Fig. 3d). Two-factor ANOVA's comparing these data to background firing revealed a main effect of reward ($F_{1,48} = 46.3$, $p < 0.0001$) but no effect of omission nor any interactions with trial number (F 's < 1.29 , p 's > 0.17), and post hoc comparisons showed that the reward-evoked response was significantly higher than background on every trial in Fig. 3d (p 's < 0.01), while the omission-evoked response did not differ on any trial (p 's > 0.05).

Examination of the difference scores across individual neurons in OFC-lesioned rats showed similar effects. For example, while the distribution of these scores was shifted significantly above zero for unexpected reward (Fig. 3h; Wilcoxon signed-rank test, $p < 0.01$), the shift was significantly less than that in shams (Fig. 3e vs 3h; Mann-Whitney U test, $p < 0.001$), as was the actual number of individual neurons in which reward-evoked activity declined significantly with learning (black bars in Fig. 3e vs 3h; X^2 -test, $X^2 = 5.12$, $df = 1$, $p = 0.02$). Furthermore, not a single neuron in the lesioned rats suppressed firing significantly in response to reward omission (Fig. 3i), and the distribution of these scores was less shifted than in shams (Fig. 3f vs 3i; Mann-Whitney U test, $p < 0.001$) and did not differ from zero (Fig. 3i; Wilcoxon signed-rank test, $p = 0.12$).

Thus, ipsilateral lesions of OFC substantially diminished the normal effect of learning on firing in response to unexpected reward and reward omission in VTA dopamine neurons. This effect was observed despite the fact that the rats' behavior indicated that they learned to expect reward at the same rate as controls (see Fig. 1). These results, along with a parallel analysis of activity at the time of delivery of the titrated delayed reward in blocks 1 and 2 (see Supplemental Figure S3), all point to a critical contribution of OFC to the prediction errors signaled by VTA dopamine neurons at the time of reward.

According to prevailing accounts such as temporal difference reinforcement learning (TDRL), prediction error signals should also be evident in response to cues. Consistent with this, reward-responsive dopamine neurons in sham-lesioned rats responded phasically

during and immediately after sampling of the odors, and this phasic response differed according to the expected value of the trial. Thus, on forced-choice trials, the average firing rate was higher during (and immediately after) sampling of the high value cue than during sampling of the low value cue (Fig. 4a). This difference was not present in the initial trials of a block but rather developed with learning (Fig. 4a, dashed vs solid lines). A two-factor ANOVA comparing firing to the odor cues across all neurons revealed a significant main effect of value ($F_{1,28} = 12.2$, $p < 0.01$) and a significant interaction between value and learning ($F_{1,28} = 18.0$, $p < 0.001$). We also quantified the effect of value by calculating the difference in firing to the high and low value cues for each neuron before and after learning; the distribution of this score was shifted significantly above zero after (Fig. 4a, late distribution; Wilcoxon signed-rank test, $p < 0.01$) but not before learning (Fig. 4a, early distribution; Wilcoxon signed-rank test, $p = 0.68$).

This pattern was also evident on free-choice trials, in which a single odor cue was presented but either of the two rewards could be selected by responding to the appropriate well. Dopaminergic activity in sham-lesioned rats increased during sampling of the single cue, and then diverged in accordance with the *future* choice of the animal, increasing more prior to selection of the high value well than the low value well. ANOVA comparing firing between odor offset (when the rat was still in the odor port) and earliest well response confirmed this effect (Fig. 4b; $F_{1,28} = 8.33$, $p < 0.01$), as did the distribution of the difference scores comparing firing during this period on high minus low value trials for each neuron (Fig. 4b, distribution; Wilcoxon signed-rank test, $p < 0.01$).

These cue-evoked effects were also altered in OFC-lesioned rats. On forced-choice trials, reward-responsive dopamine neurons fired differentially based on cue value (Fig. 4c; a two-factor ANOVA demonstrated a significant main effect of value ($F_{1,48} = 10.4$, $p < 0.01$) and a significant interaction between value and learning ($F_{1,48} = 6.36$, $p < 0.05$)), and the distribution of the difference scores comparing firing to the high and low value cues after learning was shifted significantly above zero (Fig. 4c, late distribution; Wilcoxon signed-rank test, $p < 0.01$). However, the differential firing in OFC-lesioned rats on forced-choice trials was weaker than in shams, and the number of neurons in which firing showed a significant effect of cue value was significantly lower in OFC-lesioned than sham-lesioned rats (black bars in Fig. 4a vs 4c; X^2 -test, $X^2 = 5.19$, $df = 1$, $p = 0.02$).

In addition, on free-choice trials, the difference in firing that emerged after cue-sampling in sham-lesioned rats was wholly absent in OFC-lesioned rats (Fig. 4d); a two-factor ANOVA comparing firing during this post-cue-sampling period in OFC-lesioned rats with that in shams demonstrated a significant interaction between group and value (Fig. 4b versus 4d; $F_{1,78} = 4.05$, $p < 0.05$), and post-hoc testing showed that the significant difference present in shams was not present in lesioned rats (Fig. 4d; $F_{1,48} = 1.71$, $p = 0.2$).

OFC modulates dopaminergic activity in vivo

The data above suggest that OFC modulates the firing of VTA dopamine neurons. To test this directly, we recorded juxtacellularly from VTA neurons in anesthetized rats. We identified 15 neurons with amplitude ratios and spike durations similar to those of the putative dopamine neurons recorded in the awake, behaving rats. These neurons exhibited

low baseline firing rates (3.54 ± 1.35 Hz) as well as bursting patterns characteristic of dopaminergic neurons¹⁸. Six stained with Neurobiotin and co-localized tyrosine hydroxylase (Fig 5a).

Eleven (73.3%) exhibited a statistically significant suppression of firing during and immediately following electrical stimulation of the OFC (five pulse, 20 Hz trains; Fig. 5b and 5c), including all six TH+ neurons (Fig. 5c; black lines). In each case, inhibition began during the 200 ms period of OFC stimulation and lasted for several hundred ms, averaging 393.3 ± 184.9 ms (range 220 – 740 ms). Inhibition was sometimes followed by a rebound excitation. Inhibition was not observed during the inter-stimulation interval in these neurons and thus was a specific effect of OFC stimulation.

Of the 4 neurons that did not show a significant suppression, two showed a significant increase in firing in response to OFC stimulation, suggesting that OFC can excite as well as inhibit firing in dopamine neurons, while two showed significant suppression epochs only after the end of stimulation; these were considered non-responsive since suppression began well after OFC stimulation. The average latency of onset of the OFC-dependent responses was 93.9 ± 106.9 ms (range 0 – 980 ms).

OFC does not convey value to VTA

Our results show that the OFC contributes to intact error signaling by dopamine neurons in VTA. To understand the nature of this contribution, we utilized computational modeling. In all models, we used the TDRL framework⁴ that has been used extensively to describe animal learning in reward driven tasks and the generation of phasic dopaminergic firing patterns¹⁹. In this framework a prediction error signal δ_t at time t is computed as $\delta_t = r_t + V(S_t) - V(S_{t-1})$, where r_t is the currently available reward (if any), S_t is the current state of the task (the available stimuli etc.) and $V(S_t)$ is the value of the current state, that is, the expected amount of future rewards, and $V(S_{t-1})$ is the value of the previous state, that is, the total predicted rewards prior to this timepoint. The prediction error is used to learn the state values through experience with the task, by increasing $V(S_{t-1})$ if the prediction error δ_t is positive (indicating that obtained and future expected rewards exceed the initial expectation), and decreasing $V(S_{t-1})$ if the prediction error is negative (indicating over-optimistic initial expectations that must be reduced). These prediction errors are the signals thought to be reported by dopamine neurons¹⁹ and were indeed well-matched to the neural data from VTA dopamine neurons recorded in sham-lesioned rats (Fig. 6a).

Due to the involvement of OFC in signaling reward expectancies⁷⁻¹⁰, we initially hypothesized that OFC might convey to dopamine neurons the value of states in terms of the expected future reward $V(S_t)$ at each point in time. However, modeling the OFC lesion by removing expected values from the calculation of prediction errors failed to replicate the experimental results (Fig. 6b, Model 1; for details of this and subsequent models see Methods). Specifically, while removal of (learned) values accurately predicted that firing in dopamine neurons at the time of unexpected reward or reward omission would remain unchanged with learning, this model could not account for the reduced initial response to unexpected rewards in OFC- versus sham-lesioned rats (Fig. 3c vs. 3d), nor could it generate differential firing to the odor cues on forced-choice trials (Fig. 4c). Thus, a complete loss of

value input to dopamine neurons did not reproduce the effects of OFC lesions on error signaling.

We next considered whether a partial loss of the value signal might explain the observed effects of OFC lesions. A partial loss might occur if the lesions were incomplete (as they were) or if another brain region, such as contralateral OFC and/or ventral striatum, were also providing value signals. While this produced slower learning, it did not prevent asymptotically correct values from being learned for the low and high reward port as well as for the two choices in the free-choice trials. This occurs because the remaining value-learning structures still updated their estimates of values based on ongoing prediction errors and were thus able to compensate for the loss of some of the value-learning neurons. Thus, according to this model, prediction errors to reward should still decline with training, and prediction errors to cues/choices should still increase with training, predictions at odds with the empirical data (Fig. 6b, Model 2).

A second way a partial loss of the value signal might occur is if only some of the rats had lesions sufficient to prevent value learning, while others had enough intact neural tissue to support value learning. This would amount to a partial loss of values *between* subjects (Fig. 6b, Model 3). This model did accurately predict some features of the population data, such as diminished (but still significant) differential firing to the odor cues on forced-choice trials (Fig. 4c). However, it too failed to explain the reduced initial response to unexpected rewards in OFC- versus sham-lesioned rats (Fig. 3c vs. 3d), and it could not explain the absence of differential firing as a result of future expectation of low or high rewards on free-choice trials (Fig. 4d). Moreover, this between-subjects account was at odds with the observation that none of the individual neurons in OFC-lesioned rats showed intact error signaling (Figs. 3h–j).

OFC signals state information to VTA

Our models thus failed to support the hypothesis that the OFC conveys some or all information about expected reward value to VTA neurons. Models in which this value signal was completely or even partially removed had particular difficulty accounting for both the residual differential firing based on the learned value of the odor cues on forced-choice trials (Fig. 4c) and the loss of differential firing based on the value of the impending reward on free-choice trials (Fig. 4d). In each model, these were either both present or both absent.

The fundamental difference between forced-choice and free-choice trials is that in the former, two different cues were available to signal the two different expected values, while in the latter, signaling of the different values depended entirely on internal information regarding the rats' impending decision to respond in one direction or the other. Based on this distinction, we hypothesized that OFC might provide not state values *per se*, but rather more complex information about internal choices and their likely outcomes necessary to define precisely what state the task is currently in, particularly for states that are otherwise externally ambiguous (as is the case on free-choice trials). The provision of this information would allow other brain areas to derive more accurate value expectations for such states, and subsequently signal this information to VTA. Thus, in our fourth model, we hypothesized that the OFC provides input regarding state identity to the module that computes and learns

the values, rather than directly to dopamine neurons. Removing the OFC would therefore leave an intact value learning system, albeit one forced to operate without some essential state information.

Consistent with this hypothesis, removing choice-based disambiguating input and leaving value learning to operate with more rudimentary, stimulus-bound states (Fig. S4a) produced effects that closely matched empirical results from OFC-lesioned rats (Fig. 6b, Model 4). Specifically, this model reproduced the patterns of cue selectivity evident in dopamine neurons in lesioned rats: on forced-choice trials, learning using only stimulus-bound states resulted in weaker, but still significant, differential prediction errors to the two odor cues; however, on free-choice trials, the lesioned model did not show divergent signaling at the decision-point since it lacked the ability to use internal information about the impending choice to distinguish between the two decision states. Values for the two choices in the free-choice trials could not be learned no matter the size of the learning rate parameter or the duration of learning.

Importantly, however, the consistency of this model with the neural data went beyond the effects on free-choice trials that motivated the model. In particular, the lesioned model exhibited firing to unexpected rewards and to reward omission that changed only very mildly through learning, which is similar to the neural data. Additionally, firing to an unexpected reward early in a block was significantly lower than in the unlesioned model, again closely matching the neural results (Fig. 6b, data). Overall, this fourth model best captured the contribution of OFC to learning and prediction error signaling in our task (please see Supplemental for additional discussion of modeling results).

DISCUSSION

Here we showed that OFC is necessary for normal error signaling by VTA dopamine neurons. Dopamine neurons recorded in OFC-lesioned rats showed a muted increase in firing to an unexpected reward, and this firing failed to decline with learning as was the case in sham-lesioned rats. These same neurons also failed to suppress firing when an expected reward was omitted, exhibited weaker differential firing to differently-valued cues, and failed to show differential firing based on future expected rewards on free-choice trials. Computational modeling showed that while several of these features could be approximated by postulating that OFC provides predictive value information to support the computation of reward prediction errors, they were much better explained by an alternative model in which OFC was responsible for conveying information about impending actions in order to disambiguate externally-similar states leading to different outcomes. This suggests that rather than signaling expected values *per se*, the OFC might signal state information, thereby facilitating the derivation of more accurate values, particularly for states that are primarily distinguishable based on internal rather than external information.

These results have important implications for understanding OFC and the role of VTA dopamine neurons in learning. Regarding OFC, these results provide a mechanism whereby information relevant to predicting outcomes, signaled by OFC neurons and critical to outcome-guided behavior, might influence learning²⁰. While the involvement of OFC as a

critical source of state representations is different from the role previously ascribed to OFC in learning (that of directly signaling expected reward values²⁰ or even prediction errors²¹), it would explain more clearly why this area is important for learning in some situations but not others, inasmuch as the situations requiring OFC, such as over-expectation and rapid reversal learning^{8,13,22,23}, are ones likely to benefit from disambiguation of similar states that lead to different outcomes. In these behavioral settings in particular (reversal learning, extinction, over-expectation), optimal performance would be facilitated by the ability to create new states based on internally-available information (ie recognition that contingencies have changed)²⁴. Recent models suggest that state representations of tasks are, themselves, learned^{25,26}. Whether OFC is necessary for this learning process is not clear, but our results show that OFC is key for representing the resulting states. This idea is consistent with findings that OFC neurons encode all aspects of a task in a distributed and complex manner^{27–29} and with data showing that the OFC is particularly important for accurately attributing rewards to preceding actions^{30,31}, since this depends critically on representation of previous choices. In this regard, it is worth noting that OFC neurons have been shown to signal outcome expectancies in a response-dependent fashion in the current and other behavioral settings^{12,32–36}.

The proposed contribution of OFC is also complimentary to current proposals that other brain regions, especially the ventral striatum, are important for value learning in TDRL models³⁷. OFC has strong projections to ventral striatum³⁸. Thus, information from OFC may facilitate accurate value signals in the ventral striatum, which might then be transmitted to midbrain dopaminergic neurons via inhibitory projections to contribute to prediction error signaling. Such a relay would seem essential in order to explain how glutamatergic output from the OFC acts to inhibit activity in VTA dopamine neurons, as demonstrated here and elsewhere³⁹. Other potential relays might include rostromedial tegmental nucleus, lateral habenula, or even GABAergic interneurons within VTA, all of which receive input from OFC and can act to inhibit VTA dopamine neurons. Notably non-dopaminergic neurons in VTA, many of which are likely to be GABAergic, did exhibit significantly lower baseline firing rates in OFC-lesioned rats than in controls (Fig. S2). These different pathways are not mutually exclusive, and each would be consistent with the long latency, primarily inhibitory effects of OFC stimulation on dopamine activity shown here *in vivo*.

Finally, these results expand the potential role of VTA dopamine neurons in learning, by showing that the teaching signals encoded by these neurons are based, in part, on prefrontal representations. These prefrontal representations are critical for goal-directed or model-based behaviors⁴⁰; OFC in particular is necessary for changes in conditioned responding after reinforcer devaluation and other behaviors^{7,8,41,42} that require knowledge of how different states (cues, responses and decisions, rewards, and internal consequences) are linked together in a task. However, with the exception of two recent reports^{43,44}, this knowledge has not been thought to contribute to the so-called cached values underlying dopaminergic errors. Our results show that these prefrontal representations do contribute to the value signal used by dopamine neurons to calculate errors. Correlates with action sequences, inferred values, and impending actions evident in recent dopamine recording studies could derive from access to these states and the transitions between them thought to

reside in orbital and other prefrontal areas^{45,46}. Full access to model-based task representations the states, transition functions, and derived values - would dramatically expand the types of learning that might involve dopaminergic error signals to complex associative settings⁴⁷⁻⁴⁹ more likely to reflect situations in the real world.

METHODS

Behavioral and Single-Unit Recording Methods

Subjects—Thirteen male Long-Evans rats (Charles Rivers, 4–6 mo) were tested at the University of Maryland School of Medicine in accordance with the University and NIH guidelines.

Surgical procedures—Recording electrodes were surgically implanted under stereotaxic guidance in the one hemisphere of VTA (5.2 mm posterior to bregma, 0.7 mm lateral and 7.0 mm ventral, angled 5° toward the midline from vertical). Some rats (n = 7) also received neurotoxic lesions of ipsilateral OFC by infusing N-methyl-D-aspartic acid (NMDA) (12.5 mg/ml) at 4 sites in each hemisphere: at 4.0 mm anterior to bregma, 3.8 mm ventral to the skull surface, at 2.2 mm (0.1 ul) and 3.7 mm (0.1 ul) lateral to the midline and at 3.0 mm anterior to bregma, 5.2 mm ventral to the skull surface, at 3.2 mm (0.05 ul) and 4.2 mm lateral to the midline (0.1 ul). Controls (n = 6) received sham lesions in which burr holes were drilled and the pipette tip lowered into the brain but no solution was delivered.

Behavioral task, single-unit recording, statistical analyses—Unit recording and behavioral procedures were identical to those described previously¹⁵. Statistical analyses are described in the main text.

Juxtacellular Recording Methods

Subjects—Nine male Long-Evans rats (Charles Rivers) were tested on post-natal day 60 at the University of Maryland School of Medicine in accordance with the University and NIH guidelines.

Surgical and recording procedures—Rats were anesthetized with chloral hydrate (400 mg/kg, i.p.) and placed in a stereotaxic apparatus. A bipolar concentric stimulating electrode was placed in the OFC (3.2 mm anterior and 3.0 mm lateral to bregma, and 5.2 mm ventral from brain surface) connected to an Isoflex stimulus isolation unit and driven by a Master-8 stimulator. Electrical stimulation of the OFC consisted of a 5 pulse 20 Hz train delivered every 10 sec (pulse duration 0.5 ms, pulse amplitude 500 μ A). Recording electrodes (resistance of 10–25 M Ω) were filled with a 0.5 M NaCl 2% Neurobiotin (Vector Laboratories, Burlingame, CA) solution then lowered in the VTA (5.0–5.4 mm posterior to bregma, 0.5–1.0 mm lateral and 7.8–8.5 mm ventral). Signals were amplified 10x (intracellular recording amplifier Neurodata IR-283), filtered (cutoff 1KHz, amplification 10x, Cygnus Technologies Inc.), digitized (amplification 10x, Axon Instruments Digidata 1322A), and acquired with Axoscope software (lowpass filter 5KHz, highpass filter 300 Hz, sampling 20 KHz). Baseline activity measurements were taken from the initial 5 min recording of the neuron, including mean firing rate, burst analysis, and duration of action

potential and amplitude ratio. Neurons with a mean baseline firing rate of < 6 Hz and a long duration action potential (> 1.5 ms) were considered to be putative dopamine neurons, and subjected to burst firing analysis based on established criteria¹⁸. Using these criteria, the majority of neurons recorded exhibited bursting activity (5/7, 71.42%). To assess the response to OFC stimulation, the mean value and standard deviation of baseline activity was calculated using the 2000 ms prior to the stimulation. Onset of inhibition (or excitation) was considered to be two consecutive bins after stimulation began in which the spike count was two standard deviations or more below (or above) the mean bin value ($p < 0.001$). Offset of the response was considered to be two consecutive bins in which the bin values were no longer two standard deviations away from the mean bin value. When the value of two standard deviations below the mean fell below zero, the number of consecutive bins required to signify the onset of inhibition was increased to maintain the same criterion for significance ($p < 0.001$).

Histology—Cells were labeled with Neurobiotin by passing positive current pulses (1.0–4.0 nA; 250 ms on/off square pulses; 2 Hz) and constant positive current (0.5–5.0 nA) through the recording electrode. For Neurobiotin/Tyrosine hydroxylase immunohistochemistry, tissue was sectioned at 40 μm on a freezing microtome and collected in 0.1 M PBS. Following a one hour pretreatment with 0.3% triton and 3% normal goat serum in PBS, the sections were incubated overnight with Alexa 568-conjugated streptavidin (1:800, Molecular Probes, USA) and a monoclonal mouse anti-TH antibody (1:5,000, Swant, Switzerland). The sections were then rinsed in PBS several times and incubated with a FITC-conjugated goat anti-mouse antibody for 90 minutes (1:400, Jackson Laboratories, USA). After rinsing in PBS the sections were mounted on glass slides and coverslipped in Vectashield, then examined under fluorescence on an Olympus FluoView 500 confocal microscope. Confocal images were captured in 2 μm optical steps. To confirm stimulating electrode placements, the OFC was sectioned at 50 μm and Nissl stained.

Computational Modeling Methods

Task representation—Figure S4a shows the state-action diagram of the task in the intact model. Although simplified, this state-action sequence captures all of the key aspects of a trial. To account for errors, and in line with the behavioral data (figure 1d), we also included a 20% probability that the rat would make a mistake on a forced-choice trial, e.g. going to the right reward port after a left signal. These erroneous transitions are denoted by the grey arrows in figure S4a.

Note that having the “enter left port” state be the same on both forced and free trials allows our model to generalize between rewards received on free trials and those received on forced trials, i.e. if a rat receives a long-delay reward by turning left on a forced trial, this architecture allows it to expect that the same long-delay reward will be delivered after turning left on a free trial.

After the animal moves to the reward port, it experiences state transitions according to one of two wait-for-reward state sequences (“left rew 1” to “left rew 3” or “right rew 1” to “right rew 3”). These states indicate all the possible times that the reward could be delivered by the

experimenter. Specifically, “left rew 1” is the state at which a reward drop is delivered to the left port on the “small” and “short” trials, and at which the first drop of reward is delivered on “big” trials. “left rew 2” is the time of the second reward drop on “big” trials. Rewards are never delivered at “wait left”, but this state, and the probabilistic self transition specified by the arrow returning to “wait left”, implements a variable delay between the time of the early rewards and the “long” reward delivered at “left rew 3”. Finally, after the “left rew 3” reward state the task transitions into the end state which signifies the end of the trial. State transitions are similar for the “right reward” sequence.

Update equations—We modeled the task using an actor/critic architecture and a decaying eligibility trace ⁴. We used this framework because of its common use in modeling reinforcement learning in the brain, however, the same pattern of results was obtained when modeling state-action values (rather than state values) and the state-action-response-state-action (SARSA) temporal difference learning algorithm ^{4,46,50}. Thus our results do not depend strongly on the particular implementation of TDRL.

At the start of the experiment, we set the critic’s initial values for each state $V(s)$, and the actor’s initial action preferences for action $M(s,a)$ to zero. At the beginning of each trial all eligibility traces $e(s)$ were set to zero, except that of the initial state s_I that was set as $e(s_I) = 1$. The model dynamics were then as follows:

At each time step t of a trial, the model rat is in state s_t . First, the actor chooses action a_t with probability $\pi(s_t, a_t)$ given by

$$\pi(s_t, a_t) = \frac{\exp(M(s_t, a_t))}{\sum_b \exp(M(s_t, b))} \quad (1)$$

where $M(s_t, x)$ are the actor’s preference weights for action x in state s_t , and b enumerates all possible actions at state s_t . After taking action a_t the model rat transitions to state s_{t+1} (whose value is $V(s_{t+1})$) and observes reward r_{t+1} . The prediction error, δ , is then computed according to

$$\delta = r_{t+1} + \gamma V(s_{t+1}) - V(s_t) \quad (2)$$

where $0 < \gamma < 1$ is the discount factor. At each time step, the prediction error signal is used to update all state values in the critic according to

$$V(s) \leftarrow V(s) + e(s)\eta\delta \quad (3)$$

where $0 < \eta < 1$ is the critic’s learning rate, and $e(s)$ is the eligibility of state s for updating (see below). The actor’s preference for the chosen action a_t at state s_t is also updated according to the same prediction error:

$$M(s_t, a_t) \leftarrow M(s_t, a_t) + \beta\delta \quad (4)$$

where $0 < \beta < 1$ is the actor's learning rate. Finally the eligibility trace of the last visited state, $e(s_{t+1})$, is set to one and all non-zero eligibility traces decay to zero with rate $0 < \lambda < 1$ according to

$$\begin{aligned} e(s) &\leftarrow \lambda e(s) \quad \forall s \\ e(s_t) &= 1 \end{aligned} \quad (5)$$

The free parameters of the model (γ, η, β) were set manually to accord qualitatively with the data from the sham lesioned rats, and were not altered when modeling the OFC-lesioned group.

Value lesion (Model 1)—In this model, we removed all values from the critic, i.e. we replaced all $V(s) = 0$ in the above equations. This resulted in prediction errors given by

$$\delta = r_{t+1} \quad (6)$$

which are only non-zero at the time of reward presentation.

Partial lesion within subjects (Model 2)—In this model, we assumed that f is the fraction of the critic that has been lesioned in each subject, leaving only $(1-f)$ of the value intact. The prediction error was thus:

$$\delta = r_{t+1} + (1-f)\gamma(V_{M2}(s_{t+1}) - V_{M2}(s_t)) \quad (7)$$

and this prediction error was used as a teaching signal for the intact parts of the critic. Crucially, this led to learning of values $V_{M2}(s)$ in the intact parts of the critic, albeit at a slower rate than the sham-lesion model. As a result of the error-correcting nature of TDRL, ultimately the learned values were:

$$V_{M2}(s_t) = \frac{V(s_t)}{1-f} \quad (8)$$

and the prediction errors at the end of training were similar to those of the non-lesioned model, showing that the intact parts of the critic were able to fully compensate for the loss of parts of the critic, at least if training continued for a large enough number of trials.

Partial population lesion (Model 3)—In this model, we assumed that some fraction f of the population had been successfully lesioned according to the value lesion model (Model 1), while the other rats were unaffected. The unlesioned fraction thus learn according to the prediction error of the unlesioned model, equation (2), while the lesioned fraction experience the prediction error of Model 1, equation (6). Thus, averaged over the population, the prediction error is

$$\delta = r_{t+1} + (1-f)\gamma(V_{M3}(s_{t+1}) - V_{M3}(s_t)) \quad (9)$$

Superficially, this prediction error resembles that of Model 2. Unlike Model 2, however, this prediction error is not directly used as a teaching signal, rather, each part of the population learns from a different prediction error signal (see above) and the learned values in the intact part of the population are simply $V_{M3}(s_t) = V(s_t)$.

OFC encodes states (Model 4)—In this model we hypothesized that the effect of lesioning OFC is to change the state representation of the task. Specifically, we suggest that the OFC allows the rat to disambiguate states that require internal knowledge of impending actions but are otherwise externally similar, such as the state of “I have chosen the large reward in a free-choice trial” versus “I have chosen a small reward in a free-choice trial”. We thus simulated the OFC lesion by assuming that the lesioned model no longer distinguishes between states based on the chosen action and cannot track correctly which type of reward is associated with which wait-for-reward state. Note that this ambiguity is caused not because the left and right reward ports are physically indistinguishable, but because, without knowledge of the mapping between the physical location of the ports and the abstract reward schedules at each port (in itself a type of internal “expectancy” knowledge), the rat cannot tell which of the two possible wait-for-reward sequences it is in.

Mapping between prediction errors and VTA neural response—To facilitate comparison between the model results and those of the experiments we transformed the prediction errors (which can be positive or negative) into predicted firing rates using a simple linear transformation,

$$\text{neural firing} = \text{baseline} + \text{scale factor} \times \text{prediction error}$$

with 5 spikes/second as the value for the baseline and negative prediction errors having a lower scale factor of 0.8 and positive prediction errors having a scale factor of 4.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

References

1. Hollerman JR, Schultz W. Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neuroscience*. 1998; 1:304–309. [PubMed: 10195164]
2. Pan WX, Schmidt R, Wickens JR, Hyland BI. Dopamine cells respond to predicted events during classical conditioning: evidence for eligibility traces in the reward-learning network. *Journal of Neuroscience*. 2005; 25:6235–6242. [PubMed: 15987953]
3. Bayer HM, Glimcher P. Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*. 2005; 47:129–141. [PubMed: 15996553]
4. Sutton, RS.; Barto, AG. *Reinforcement Learning: An introduction*. MIT Press; 1998.
5. Pearce JM, Hall G. A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. 1980; 87:532–552.
6. Rescorla, RA.; Wagner, AR. *Classical Conditioning II: Current Research and Theory*. Black, AH.; Prokasy, WF., editors. Appleton-Century-Crofts; 1972. p. 64-99.

7. Pickens CL, et al. Different roles for orbitofrontal cortex and basolateral amygdala in a reinforcer devaluation task. *Journal of Neuroscience*. 2003; 23:11078–11084. [PubMed: 14657165]
8. Izquierdo AD, Suda RK, Murray EA. Bilateral orbital prefrontal cortex lesions in rhesus monkeys disrupt choices guided by both reward value and reward contingency. *Journal of Neuroscience*. 2004; 24:7540–7548. [PubMed: 15329401]
9. O’Doherty J, Deichmann R, Critchley HD, Dolan RJ. Neural responses during anticipation of a primary taste reward. *Neuron*. 2002; 33:815–826. [PubMed: 11879657]
10. Gottfried JA, O’Doherty J, Dolan RJ. Encoding predictive reward value in human amygdala and orbitofrontal cortex. *Science*. 2003; 301:1104–1107. [PubMed: 12934011]
11. Padoa-Schioppa C, Assad JA. Neurons in orbitofrontal cortex encode economic value. *Nature*. 2006; 441:223–226. [PubMed: 16633341]
12. Roesch MR, Taylor AR, Schoenbaum G. Encoding of time-discounted rewards in orbitofrontal cortex is independent of value representation. *Neuron*. 2006; 51:509–520. [PubMed: 16908415]
13. Takahashi Y, et al. The orbitofrontal cortex and ventral tegmental area are necessary for learning from unexpected outcomes. *Neuron*. 2009; 62:269–280. [PubMed: 19409271]
14. Vazquez-Borsetti P, Cortes R, Artigas F. Pyramidal neurons in rat prefrontal cortex projecting to ventral tegmental area and dorsal raphe nucleus express 5-HT_{2A} receptors. *Cerebral Cortex*. 2009; 19:1678–1686. [PubMed: 19029064]
15. Roesch MR, Calu DJ, Schoenbaum G. Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards. *Nature Neuroscience*. 2007; 10:1615–1624. [PubMed: 18026098]
16. Margolis EB, Lock H, Hjelmstad GO, Fields HL. The ventral tegmental area revisited: Is there an electrophysiological marker for dopaminergic neurons? *Journal of Physiology*. 2006; 577:907–924. [PubMed: 16959856]
17. Xin J, Costa RM. Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature*. 2010; 466:457–462. [PubMed: 20651684]
18. Grace AA, Bunney BS. The control of firing pattern in nigral dopamine neurons: burst firing. *Journal of Neuroscience*. 1984; 4:2877–2890. [PubMed: 6150071]
19. Schultz W, Dayan P, Montague PR. A neural substrate for prediction and reward. *Science*. 1997; 275:1593–1599. [PubMed: 9054347]
20. Schoenbaum G, Roesch MR, Stalnaker TA, Takahashi YK. A new perspective on the role of the orbitofrontal cortex in adaptive behaviour. *Nature Reviews Neuroscience*. 2009; 10:885–892. [PubMed: 19904278]
21. Sul JH, Kim H, Huh N, Lee D, Jung MW. Distinct roles of rodent orbitofrontal and medial prefrontal cortex in decision making. *Neuron*. 2010; 66:449–460. [PubMed: 20471357]
22. Chudasama Y, Robbins TW. Dissociable contributions of the orbitofrontal and infralimbic cortex to pavlovian autoshaping and discrimination reversal learning: further evidence for the functional heterogeneity of the rodent frontal cortex. *Journal of Neuroscience*. 2003; 23:8771–8780. [PubMed: 14507977]
23. Fellows LK, Farah MJ. Ventromedial frontal cortex mediates affective shifting in humans: evidence from a reversal learning paradigm. *Brain*. 2003; 126:1830–1837. [PubMed: 12821528]
24. Gershman SJ, Niv Y. Learning latent structure: carving nature at its joints. *Current Opinion in Neurobiology*. 2010; 20:251–256. [PubMed: 20227271]
25. Redish AD, Jensen S, Johnson A, Kurth-Nelson Z. Reconciling reinforcement learning models with behavioral extinction and renewal: implications for addiction, relapse, and problem gambling. *Psychological Review*. 2007; 14:784–805. [PubMed: 17638506]
26. Gershman SJ, Blei DM, Niv Y. Time, context and extinction. *Psychological Review*. 2010; 117:197–209. [PubMed: 20063968]
27. Ramus SJ, Eichenbaum H. Neural correlates of olfactory recognition memory in the rat orbitofrontal cortex. *Journal of Neuroscience*. 2000; 20:8199–8208. [PubMed: 11050143]
28. van Duuren E, Lankelma J, Pennartz CMA. Population coding of reward magnitude in the orbitofrontal cortex of the rat. *Journal of Neuroscience*. 2008; 28:8590–8603. [PubMed: 18716218]

29. van Duuren E, et al. Single-cell and population coding of expected reward probability in the orbitofrontal cortex of the rat. *Journal of Neuroscience*. 2009; 29:8965–8976. [PubMed: 19605634]
30. Walton ME, Behrens TEJ, Buckley MJ, Rudebeck PH, Rushworth MFS. Separable learning systems in the macaque brain and the role of the orbitofrontal cortex in contingent learning. *Neuron*. 2010; 65:927–939. [PubMed: 20346766]
31. Tsuchida A, Doll BB, Fellows LK. Beyond reversal: a critical role for human orbitofrontal cortex in flexible learning from probabilistic feedback. *Journal of Neuroscience*. 2010; 30:16868–16875. [PubMed: 21159958]
32. Tsujimoto S, Genovesio A, Wise SP. Monkey orbitofrontal cortex encodes response choices near feedback time. *Journal of Neuroscience*. 2009; 29:2569–2574. [PubMed: 19244532]
33. Feierstein CE, Quirk MC, Uchida N, Sosulski DL, Mainen ZF. Representation of spatial goals in rat orbitofrontal cortex. *Neuron*. 2006; 51:495–507. [PubMed: 16908414]
34. Furuyashiki T, Holland PC, Gallagher M. Rat orbitofrontal cortex separately encodes response and outcome information during performance of goal-directed behavior. *Journal of Neuroscience*. 2008; 28:5127–5138. [PubMed: 18463266]
35. Abe H, Lee D. Distributed coding of actual and hypothetical outcomes in the orbital and dorsolateral prefrontal cortex. *Neuron*. 2011; 70:731–741. [PubMed: 21609828]
36. Young JJ, Shapiro ML. Dynamic coding of goal-directed paths by orbital prefrontal cortex. *Journal of Neuroscience*. 2011; 31:5989–6000. [PubMed: 21508224]
37. O'Doherty J, et al. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science*. 2004; 304:452–454. [PubMed: 15087550]
38. Voorn P, Vanderschuren LJMJ, Groenewegen HJ, Robbins TW, Pennartz CMA. Putting a spin on the dorsal-ventral divide of the striatum. *Trends in Neurosciences*. 2004; 27:468–474. [PubMed: 15271494]
39. Lodge DJ. The medial prefrontal and orbitofrontal cortices differentially regulate dopamine system function. *Neuropsychopharmacology*. 2011; 36:1227–1236. [PubMed: 21307842]
40. Daw ND, Niv Y, Dayan P. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*. 2005; 8:1704–1711. [PubMed: 16286932]
41. Burke KA, Franz TM, Miller DN, Schoenbaum G. The role of the orbitofrontal cortex in the pursuit of happiness and more specific rewards. *Nature*. 2008; 454:340–344. [PubMed: 18563088]
42. Ostlund SB, Balleine BW. Orbitofrontal cortex mediates outcome encoding in Pavlovian but not instrumental learning. *Journal of Neuroscience*. 2007; 27:4819–4825. [PubMed: 17475789]
43. Daw ND, Gershman SJ, Seymour B, Dayan P, Dolan RJ. Model-based influences on humans' choices and striatal prediction errors. *Neuron*. (in press).
44. Simon DA, Daw ND. Neural correlates of forward planning in a spatial decision task in humans. *Journal of Neuroscience*. (in press).
45. Bromberg-Martin ES, Matsumoto M, Hong S, Hikosaka O. A pallidus-habenula-dopamine pathway signals inferred stimulus values. *Journal of Neurophysiology*. 2010; 104:1068–1076. [PubMed: 20538770]
46. Morris G, Nevet A, Arkadir D, Vaadia E, Bergman H. Midbrain dopamine neurons encode decisions for future action. *Nature Neuroscience*. 2006; 9:1057–1063. [PubMed: 16862149]
47. Hampton AN, Bossaerts P, O'Doherty JP. The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. *Journal of Neuroscience*. 2006; 26:8360–8367. [PubMed: 16899731]
48. Glascher J, Daw N, Dayan P, O'Doherty JP. Prediction error signals underlying model-based and model-free reinforcement learning. *Neuron*. 2010; 66:585–595. [PubMed: 20510862]
49. McDannald MA, Lucantonio F, Burke KA, Niv Y, Schoenbaum G. Ventral striatum and orbitofrontal cortex are both required for model-based, but not model-free, reinforcement learning. *Journal of Neuroscience*. 2011; 31:2700–2705. [PubMed: 21325538]
50. Niv Y, Daw ND, Dayan P. Choice values. *Nature Neuroscience*. 2006; 9:987–988. [PubMed: 16871163]

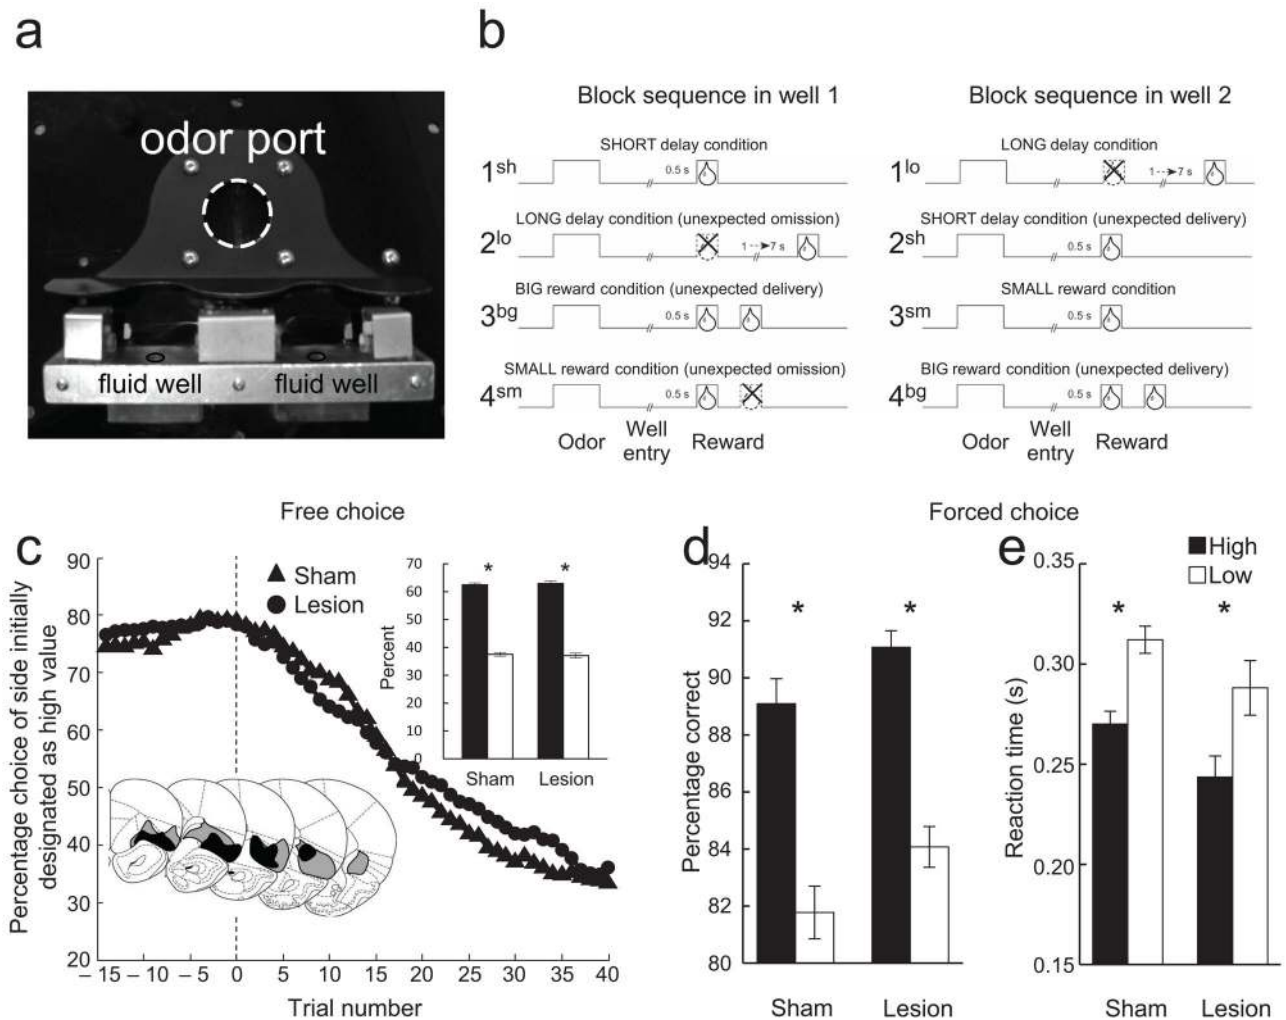


Figure 1. Apparatus and behavioral results

(a) Odor port and fluid wells. (b) Time course of stimuli (odors and rewards) presented to the animal on each trial. At the start of each recording session one well was randomly designated as short (a 0.5 s delay before reward) and the other, long (a 1–7 s delay before reward) (block 1). In the second block of trials, these contingencies were switched (block 2). In blocks 3 and 4, the delay was held constant while the number of the rewards delivered was manipulated. Expected rewards were thus omitted on long and small trials at the start of blocks 2 (2^{lo}) and 4 (4sm), respectively, and rewards were delivered unexpectedly on short and big trials at the start of blocks 2 (2^{sh}) and 3 and 4 (3^{bg} and 4^{bg}), respectively. (c) Choice behavior in trials before and after the switch from high-valued outcome (averaged across short and big) to a low-valued outcome (averaged across long and small). Inset bar graphs show average percentage choice for high- versus low-value outcomes across all free-choice trials. Inset brain sections illustrate the extent of the maximum (gray) and minimum (black) lesion at each level within OFC in the lesioned rats. (d and e) Behavior on forced-choice trials. Bar graphs show percent correct (d) and reaction times (e) in response to the high and low value cues across all recording sessions. *, t-test or other comparison, $p < 0.05$ or better; ‘ns’ denotes non-significant. Error bars: SEM.

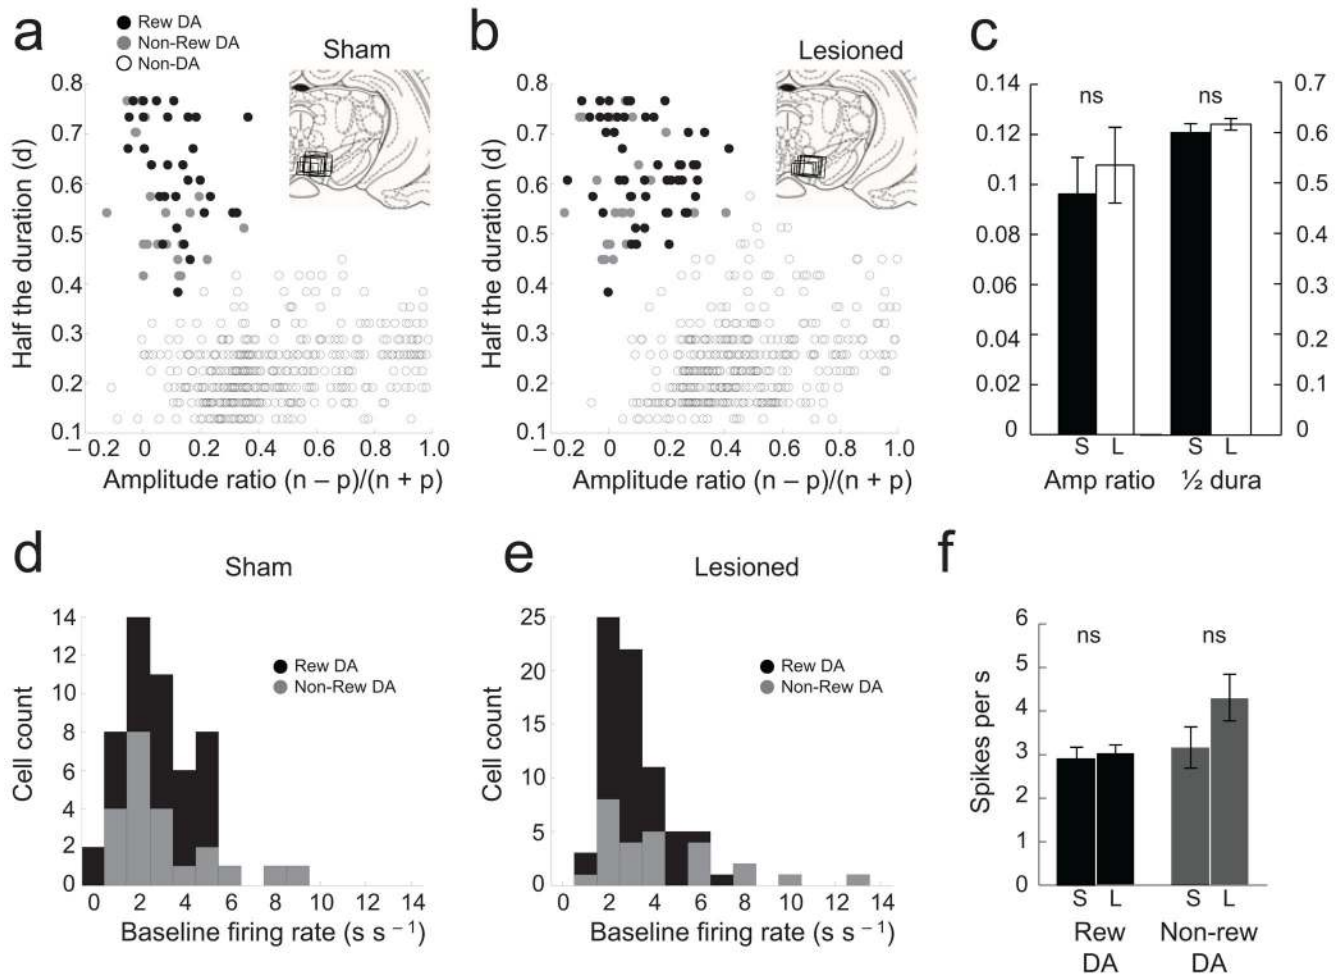


Figure 2. Identification, waveform features, and firing rates of putative dopamine and non-dopamine neurons

(a and b) Results of cluster analysis based on the half time of the spike duration (d) and the ratio comparing the amplitude of the first positive and negative waveform segments ($(n-p)/(n+p)$). The center and variance of each cluster was computed without data from the neuron of interest, and then that neuron was assigned to a cluster if it was within 3 deviations of the cluster's center. Neurons that met this criterion for more than one cluster were not classified. This process was repeated for each neuron. Reward-responsive dopamine neurons are shown in black; reward non-responsive dopamine neurons are shown in grey; neurons that classified with other clusters, no clusters or more than one cluster are shown in open circles. Insets in each panel indicate location of the electrode tracks in sham (a) and OFC-lesioned rats (b). (c) Bar graphs indicate average amplitude ratio and half duration of putative dopamine neurons in sham (black) and OFC-lesioned rats (open). (d-f) Distribution and average baseline firing rates for reward-responsive (black) and non-responsive (grey) dopamine neurons in sham (d and f) and OFC-lesioned rats (e and f). S, sham; L, lesioned. *, t-test or other comparison, $p < 0.05$ or better; 'ns' denotes non-significant. Error bars: SEM.

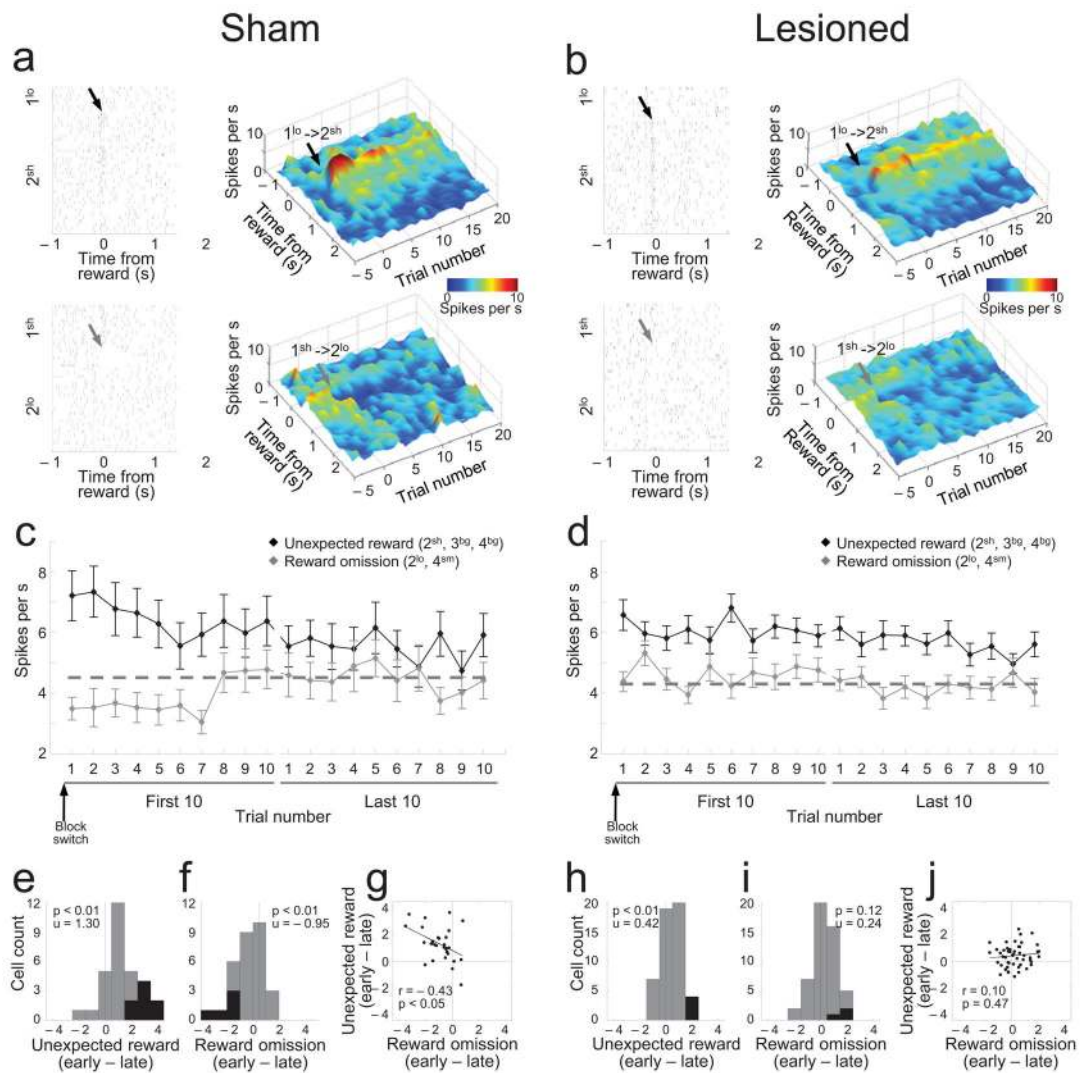


Figure 3. Changes in activity of reward-responsive dopamine neurons in response to unexpected reward delivery and omission

(a and b) Activity in a representative neuron (raster) or averaged across all reward-responsive dopamine neurons (heat plot) in sham (a) and OFC-lesioned rats (b) in response to introduction of unexpected reward in block 2^{sh} (upper plots, black arrows) and omission of expected reward in block 2^{lo} (lower plots, grey arrows). (c and d) Average firing during the period 500ms after reward delivery or omission in reward-responsive dopamine neurons in sham (c) and OFC-lesioned rats (d) in blocks in which an unexpected reward was instituted (blocks 2^{sh}, 3^{bg}, and 4^{bg}, black lines) or an expected reward omitted (blocks 2^{lo} and 4sm, grey lines). Dotted lines indicate background firing. (e–j) Distribution of difference scores and scatter plots comparing firing to unexpected reward and reward omission early versus late in relevant trial blocks in sham (e–g) and OFC-lesioned rats (h–j). Difference scores were computed from the average firing rate of each neuron in the first 5 minus the last 15 trials in relevant trial blocks. Black bars represent neurons in which the difference in firing was statistically significant (t-test; $p < 0.05$). The numbers in upper left of each panel indicate results of Wilcoxon signed-rank test (p) and the average index (u).

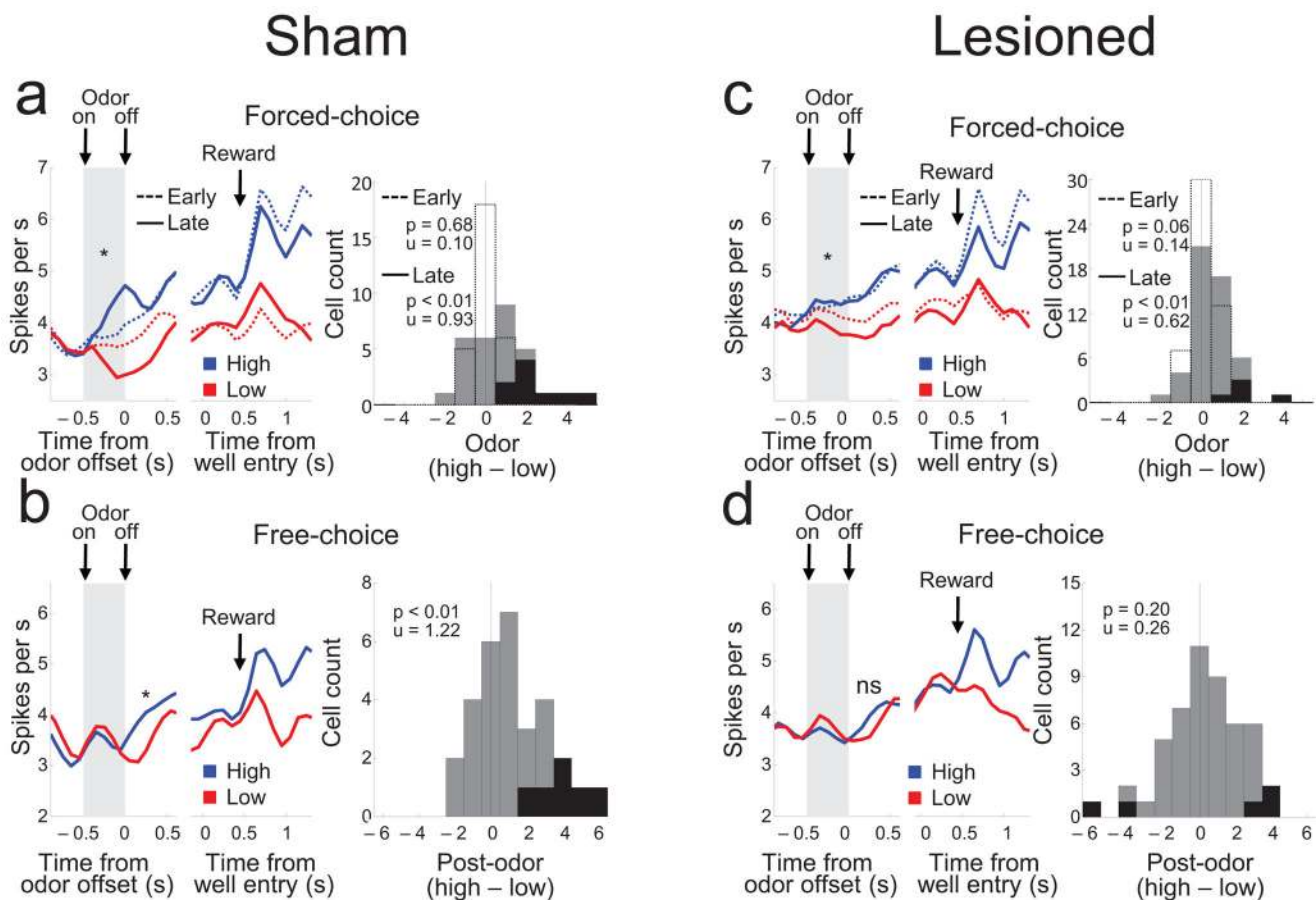


Figure 4. Changes in activity of reward-responsive dopamine neurons during and after odor cue sampling on free- and forced-choice trials

(a and c) Neural activity during forced-choice trials in shams **(a)** and OFC-lesioned rats **(c)**. Line plots show average activity synchronized to odor offset or well entry across all blocks on trials involving the high (blue lines) and low (red lines) value cues. Activity is shown separately for the first 15 (early, dotted lines) and the last 5 trials (late, solid lines) in each block, corresponding to the time during and after learning in response to a change in the size or timing of reward (Fig. 1c). Histograms show the distribution of difference scores comparing firing during sampling of the high minus the low value cues, early (dotted line bars) and late (solid bars) in the blocks. **(b and d)** Neural activity on free-choice trials in shams **(b)** and OFC-lesioned rats **(d)**. Line plots show average activity synchronized to odor offset or well entry across all blocks on high (blue lines) and low (red lines) value trials. Histograms show the distribution of difference scores comparing firing between odor offset and earliest well response on high minus low value trials. Black bars represent neurons in which the difference in firing was statistically significant late (t-test; $p < 0.05$). The numbers indicate results of Wilcoxon signed-rank test (p) and the average index (u). *, $p < 0.05$ or better on post-hoc testing; ns, non-significant.

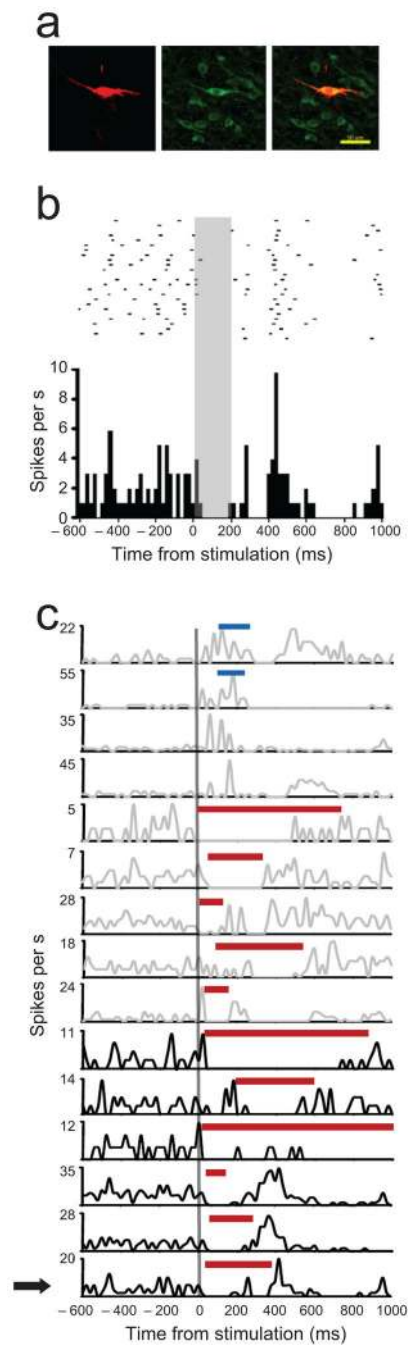


Figure 5. Changes in dopamine neuron activity in response to OFC stimulation

(a) An example of putative dopamine neuron labeled with Neurobiotin (left, red) and TH (middle, green). This neuron had morphological characteristics of dopamine neurons (bipolar dendritic orientations) and showed co-localization of Neurobiotin and TH (right, merged). Scale bar=50 μ m. (b) Raster plot and peri-stimulus histogram showing activity in the TH+ neuron in (a) before, during and after OFC stimulation. Top: each line represents a trial and dots indicate time of action potential firing; grey box indicates period of OFC stimulation. Bottom: cumulative histogram depicting firing across all trials and revealing a

pause during the stimulation. **(c)** Firing rate plots showing activity in each of the 15 recorded neurons before, during and after OFC stimulation. Arrow indicates raster of neuron shown in **(a, b)**. Each line shows the average firing rate per stimulation trial for a given neuron. Activity is aligned to onset of OFC stimulation. The vertical grey line indicates the onset of OFC stimulation. Thirteen neurons showed periods of significant inhibition (red bars) or excitation (blue bars) that began during the stimulation. Excluding secondary or rebound excitation or inhibition evident in the figure, these neurons did not exhibit significant epochs elsewhere in the inter-stimulus interval (not shown). Grey rate plots, putative dopamine neurons; black rate plots, TH immunopositive neurons.

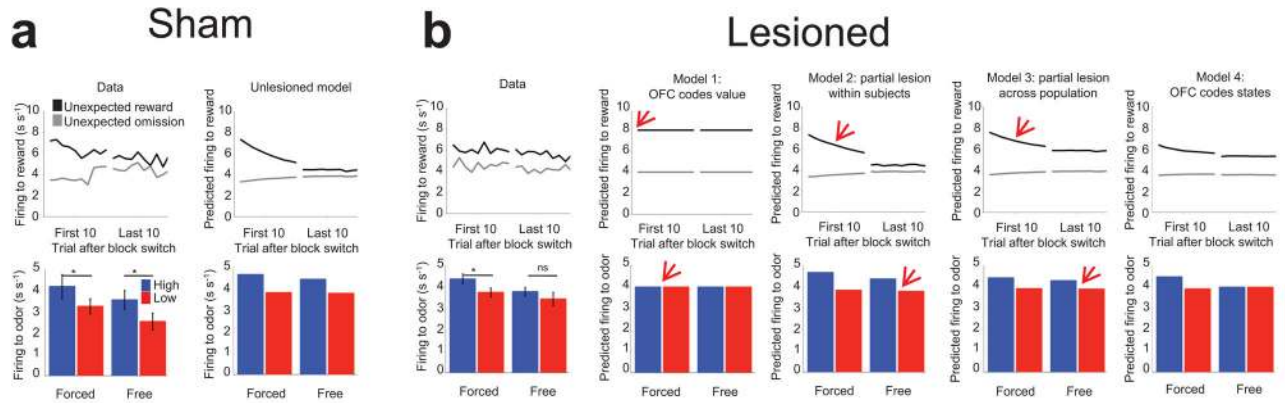


Figure 6. Comparison of model simulations and experimental data

(a) The unlesioned TDRL model and experimental data from the sham-lesioned rats. Top: At the time of unexpected reward delivery or omission, the model predicts positive (black) and negative (grey) prediction errors whose magnitude diminishes as trials proceed. Bottom: At the time of the odor cue, the model reproduces the increased responding to high value (blue) relative to low value (red) cues on forced trials. Likewise, the model predicts differential firing at the time of decision on free-choice trials. **(b)** Lesioning the hypothesized OFC component in each model produced qualitatively different effects. Red arrows highlight discrepancies between the models and the experimental data, where these exist. Model 1, which postulates that OFC conveys expected values to dopamine neurons, cannot explain the reduced firing to unexpected rewards at the beginning of a block, nor can it reproduce the differential responding to the two cues on forced-choice trials. Models 2 and 3, which assume a partial lesion of value encoding, cannot account for the lack of significant difference between high and low value choices on free-choice trials in the recorded data, and incorrectly predict significantly diminished responses at the time of reward after learning. Only Model 4, in which OFC encoding enriches the state representation of the task by distinguishing between states based on impending actions, was able to fully account for the results at the time of unexpected reward delivery or omission and at the odor on free- and forced-choice trials.