Network computations underlying learning from symbolic gains and losses

Hua Tang[1*], Ramon Bartolo[1, 2], Bruno B. Averbeck[1*]

1. Laboratory of Neuropsychology, National Institute of Mental Health, National Institutes of Health, Bethesda, MD, USA.

2. Laboratory of Sensorimotor Research, National Eye Institute, National Institutes of Health, Bethesda, MD, USA.

*Correspondence: hua.tang@nih.gov (H.T.), bruno.averbeck@nih.gov (B. B. A.)

Abstract

Reinforcement learning (RL) engages a network of areas, including the orbitofrontal cortex (OFC), ventral striatum (VS), amygdala (AMY), and mediodorsal thalamus (MDt). This study examined RL mediated by gains and losses of symbolic reinforcers across this network. Monkeys learned to select options that led to gaining tokens and avoid options that led to losing tokens. Tokens were cashed out for juice rewards periodically. OFC played a dominant role in coding information about token updates, suggesting that the cortex is more important than subcortical structures when learning from symbolic outcomes. We also found that VS showed increased responses specific to appetitive outcomes, and AMY responded to the salience of outcomes. In addition, analysis of network activity showed that symbolic reinforcement was calculated by temporal differentiation of accumulated tokens. This process was mediated by dynamics within the OFC-MDt-VS circuit. Thus, we provide a neurocomputational account of learning from symbolic gains and losses.

Introduction

Neurophysiology and anatomy have identified large-scale forebrain networks that underlie specific behavioral processes [1, 2]. Most of the work in physiology has focused on establishing the function of single cortical or subcortical areas, by looking for correlates of hypothesized functions in the responses of single neurons or populations within each area. The recent advent of circuit work has extended the study of single areas to pairs of areas, by looking at specific inputs to areas or projection targets of areas [3]. However, this work has not extended the study to networks of areas. Much of this work has shown that anatomically connected areas tend to process similar information once one moves beyond primary sensory or motor areas [4], and lesions or manipulations to multiple nodes across a connected network can lead to similar deficits in tasks [5, 6]. An exception to this is fMRI, which has shifted to studying the contribution of identified networks to behavior [7, 8]. Thus, examining connected areas within networks is important to understand how the network orchestrates behavior.

Although early studies of reinforcement learning (RL) focused on the striatum and the dopamine innervation of the striatum [9], subsequent work showed that additional structures, including a network composed of the OFC [10, 11], AMY [12, 13], and MDt [5] were also involved in RL [2]. Although causal manipulations and neurophysiology have shown that each of these areas plays a role in RL, the results across tasks and areas are not always consistent. For example, lesions to the VS cause deficits in bandit tasks that use probabilistic delivery of primary reinforcers [13]. However, these deficits are limited to learning to associate rewards with objects and not actions [14]. This is not consistent with actor-critic models of RL that suggest that VS stores a general state value representation for policy learning [15, 16]. Similarly, AMY and VS appear to play stronger roles in probabilistic bandit tasks but not in token-based RL tasks, in which tokens are used as reinforcers [17-19]. In token-based RL, VS is only involved in learning to discriminate between relative gains and plays no role in learning to discriminate between gains and losses [17]. AMY, on the other hand, appears to play almost no role in token-based RL [18]. Previous neurophysiology studies have shown that token-based learning may engage cortical networks more than subcortical networks [20, 21], whereas learning based directly on primary rewards may preferentially engage subcortical networks [22, 23]. Beyond these examples, some learning-related behaviors may rely more on inference processes than incremental value updates that characterize RL and may, therefore, tap into different networks [24, 25]. Thus, the way in which the ventral cortico-striatal-thalamo-cortical network, including the OFC, VS, AMY, and MDt, orchestrates RL, particularly with symbolic reinforcers, is unclear.

In the present study, we have carried out physiology recordings across the ventral cortico-striatal-thalamo-cortical network using a token-based RL task. We examined single neuron, population, and network

computations underlying the performance of the task. We began by examining single neuron coding of task variables across areas, and found, consistent with previous work, that the task-relevant variables were encoded broadly across areas. We found a small enhancement of token encoding in the OFC, chosen stimulus identity in the AMY, and choice direction (action) encoding in the MDt. However, other variables, including chosen value and primary reward, were encoded similarly across areas. Thus, finding a neural correlate of a task variable within a single area provides minimal evidence that the area is important for that variable. More differences emerged when we examined how neurons were tuned to task variables across areas. Specifically, although both OFC and VS encoded chosen value and token outcomes, VS strongly encoded appetitive choices and outcomes (i.e., gains) with increased firing rates, while OFC was more balanced with a bias towards encoding appetitive choices and outcomes with decreased firing rates. AMY uniquely showed enhanced salience coding and no correlation between the encoding of chosen value and token outcomes. Furthermore, AMY showed the strongest correlations between population coding of primary and symbolic reinforcers, whereas OFC strongly encoded both reinforcers but discriminated between them at the population level. Finally, when we examined network computations underlying token reinforcement, we found that OFC and VS both computed token outcomes as time derivatives of accumulated tokens and relayed this information reciprocally through the MDt. Taken together, this shows unique contributions of each area, revealed by analysis of large-scale recordings across the ventral cortico-striatal-thalamo-cortical network, as well as evidence for why lesions across the nodes can affect RL differentially with different paradigms.

Results

Two rhesus monkeys were trained on a two-armed bandit task. In this task, the monkeys collected tokens, which were periodically exchanged for juice rewards (Fig. 1a-b). In every block of 108 trials, we introduced four novel images. In each trial, two of the four images were presented on the screen. The choice of an image led stochastically to one of -2, -1, +1, or +2 tokens (Fig. 1b). The image-outcome relationships were unknown to the monkeys at the start of the block. Monkeys had to learn the values of the images by choosing one of them and observing the outcome. Token outcomes were stochastic such that in 75% of the trials, the monkeys received the number of tokens associated with the chosen image, and in 25% of the trials, the number of tokens did not change. Tokens were accumulated across trials and were cashed out every four to six trials, with one drop of juice for each token.

The primary reward was apple juice (#juice), but it was only delivered every four to six trials. In each trial, the monkeys had a chance to gain or lose tokens by choosing one of the images. The chosen value information was first carried by the images (cValue), which predicted the change of tokens (Δtoken). The number of accumulated tokens (#token) was always presented on the screen. So, value information appeared in different forms, including cValue, Δtoken, #token, and #juice (Extended Data Fig. 1a). The following analyses focus on the neural representation and interactions of these forms of value.

Choice behavior was influenced by gaining/losing tokens and token numbers

We quantified choice behavior by measuring the fraction of times the monkeys chose the image associated with a higher value in each condition (e.g., choose +1 when presented with -1 and +1 images). The monkeys learned to distinguish the values of two images shown on the screen within about 10 trials for the Gain/Loss and Gain/Gain conditions but learned minimally in the Loss/Loss condition (Fig. 1g, and Extended Data Fig. 1b). We also fit a Rescorla-Wagner (RW) reinforcement learning model to the choice behavior (Fig. 1g), which has been explored previously [17]. The model is used below to examine token reward prediction error (RPE). To quantify how the choice behavior (i.e., whether they chose the better option) was affected by task variables, we fit a multi-way ANOVA model to it (Fig. 1h-i). The choice behavior was significantly modulated by the number of observations (Trial, $F_{17,\ 44950} = 109.98$, $p < 0.001$), Gain/Loss condition (Gain/Loss, $F_{1,\ 44950} = 286.91$, $p < 0.001$), value difference of two options (Δvalue, $F_{3,\ 44950} = 73.37$, $p < 0.001$), and the number of tokens before choice (#token, $F_{10,\ 44950} = 8.68$, $p < 0.001$), but not by the number of trials since last token cashout (CashID, $F_{5,\ 44950} = 2.07$, $p = 0.066$). The number of observations indicates the learning process, both the Gain/Loss condition and Δvalue indicate the updating of tokens, #token

indicates the accumulated value, and the CashID indicates the probability of receiving a primary reward. The result indicates that the monkeys understood the meaning of tokens and adjusted their behavior to get more tokens as learning progressed.

Neural encoding of choices, primary and symbolic reinforcers

We performed simultaneous recordings of population neural activity across four regions of the ventral cortico-striatal-thalamo-cortical network (Fig. 1c and Extended Data Fig. 1c-e) from two macaque monkeys using multi-site linear probes. We collected 606 neurons in the 13L region of the orbitofrontal cortex (OFC), 829 neurons in the core region of the ventral striatum (VS), 1607 neurons in the basolateral amygdala (AMY), and 1035 neurons in the medial portion of the mediodorsal thalamus (MDt) using a semi-chronic recording procedure (Fig. 1d-f). Neuronal activity differed across areas in different task epochs (Extended Data Fig. 1f-g).

To examine how choices and rewards were represented in each area, we fit the responses of single neurons with a sliding window ANOVA model. The model included multiple task-relevant factors, which were the number of tokens (#token), the change of token numbers (Δtoken), the number of juice drops delivered on cashout trials (#juice), the image pair presented (condition), the stimulus identity (cStim), *a priori* value (cValue, i.e., +2, +1, -1, -2), and direction (cDir) of chosen images. The factors #token, Δtoken, and #juice (Fig. 2a-c) were the reinforcement signals that drove the monkeys' behavior. Neurons in all areas (> 10%), but more in the OFC (> 20%), showed a strong representation of #token (Fig. 2a, h). This is consistent with the tokens that were always present on the screen during the trial. There was also a phasic increase in all areas when tokens were updated. The Δtoken was also encoded by neurons in all areas but more strongly in the OFC and VS (Fig. 2b). OFC also led the encoding of Δtoken in time (Fig. 2b, inset). The #juice activated more than 40% of neurons in every area (Fig. 2c, i). OFC played a significant role in encoding the reinforcement signals, having the highest proportion of neurons encoding these task variables. A substantial proportion of neurons in these areas also showed responses to choices, including the identity (Fig. 2d), *a priori* value (Fig. 2e, g), and direction (Fig. 2f) of chosen images. Learning-related values derived from the RW model were encoded consistently across areas (Extended Data Fig. 4b). The stimulus identities, chosen values, and directions were most strongly represented in the AMY, OFC, and MDt, but differences were often subtle. This result shows that task variables related to the choices and rewards were encoded across the ventral network, with different regions encoding specific types of information at different phases.

Diverse encoding of value update information

The updating of value information was represented as gaining or losing tokens. In the first analysis, we used only linear encoding (Fig. 2b). However, neurons may encode value information in more diverse ways. To quantitatively characterize the encoding patterns across the areas, we examined the activity of each neuron using a series of linear regression models (Extended Data Fig. 2a). We identified the best-fitting model for each neuron and classified them into different functional categories according to how they were tuned to outcomes.

Here, we classified neurons encoding Δtoken into five types: neurons encoding (1) value linearly across both Gain and Loss (Fig. 3a, f), (2) value salience (Fig. 3b, g), (3) categorial Gain/Loss (Fig. 3c, h), (4) value only for Gain (Fig. 3d, i), and (5) value only for Loss (Fig. 3e, j). Note that each encoding type could be positively or negatively related to the neural activity. The first category encoded Δtoken linearly (Fig. 3a). These neurons encoded the gains and losses on a linear value axis, in other words, processing the gaining and losing of tokens in the same internal value system. Many OFC neurons, but few in the AMY, were of this type. However, more AMY neurons encoded the salience of Δtoken. These neurons represent both gains and losses but with inverse correlations of neural activity and value (Fig. 3b). The categorical Gain/Loss signal groups each option as gain or loss, regardless of the magnitude. These neurons were found throughout the ventral network, especially in the OFC and VS (Fig. 3c). Many more neurons in these areas encoded gains rather than losses. VS and OFC neurons showed robust and phasic responses to gains (Fig. 3d). Only a small proportion of neurons encoded losses, mostly in the OFC (Fig. 3e). Before the token update, the chosen images predicted the value update. We also classified the neurons encoding cValue into the same five types (Extended Data Fig. 3a). They showed similar encoding patterns, with activity locked to the onset of the images. This analysis shows that these areas played unique roles in encoding the chosen value and token outcome information.

An alternative way to examine the encoding of value update information is to measure the representations of gains and losses separately by running two regressions, either using Gain trials (i.e., $\Delta token \in [0, 1, 2]$) or Loss trials (i.e., $\Delta token \in [0, -1, -2]$), then calculate the correlation of the Gain and Loss regression coefficients (Fig. 3k-o). The OFC, VS, and MDt populations showed high co-encoding of gains and losses, which indicates the encoding of gains and losses on similar value axes (Fig. 3k). Especially for OFC, the effect was consistent for two seconds after the token update (Fig. 4l). This was consistent with the single-cell result (Fig. 3a), indicating that OFC processes the gaining and losing of tokens closer to objective outcomes than the other areas (Fig. 4k; Fisher's z-transformation, $p < 0.05$). Conversely, AMY showed

negative correlations between Gain and Loss regression coefficients (Fig. 4k, n), indicating populational encoding of value salience (Fig. 3b).

Co-encoding of the value update information

In each trial, the monkeys may gain or lose tokens. The predicted value update was carried by the images (cValue) before the token outcome was delivered (Δtoken). Did neurons in these areas encode the two forms of value information in a similar way? To address this question, we first compared the number of neurons that encoded both forms of value linearly. Neurons encoding cValue in a window 500 ms before the token update and Δtoken in a window 500 ms after the token update were classified as encoding both cValue and Δtoken (blue dots in Fig. 4a-d). Fewer neurons in the AMY than in the other areas linearly encoded both factors (Fig. 4a-d, f; chi-square test, $p < 0.01$). To address how similar AMY neurons encoded cValue and Δtoken, we then calculated the correlation of the cValue and Δtoken regression coefficients in each area (Extended Data Fig. 2b-c). AMY neurons showed a significantly lower correlation between cValue and Δtoken than the other areas (Fig. 4a-d, g; Fisher's z-transformation, $p < 0.05$ for significant neurons, $p < 0.001$ for all neurons). This result shows that neurons in the AMY encoded the value carried by the images and token outcome less consistently, which is also consistent with AMY primarily encoding salience.

We also found asymmetric encoding of gains and losses in the VS, OFC and MDt (Fig. 4e; chi-square test; VS, $x^2 = 95.74$, $p < 0.001$; OFC, $x^2 = 19.96$, $p < 0.001$; MDt, $x^2 = 10.87$, $p < 0.05$). More VS neurons had positive cValue and Δtoken regression coefficients. Together with the results in Fig. 3, this indicates that more neurons in the VS fired more when choosing higher-valued images and getting more tokens in the Gain conditions (Fig. 4b, e, quadrant 4). OFC neurons showed more balanced encoding in increased and decreased activities (Fig. 4e, quadrant 3 vs. 4; chi-square test, $x^2 = 3.11$, $p = 0.078$). Thus, neurons in the VS tended to respond with increased firing rates to the choice of good options and outcomes, whereas neurons in the OFC had both positive and negative tuning to the same variables.

Co-encoding of primary and symbolic reinforcers

Although a growing number of studies have adopted symbolic reinforcers, whether they are encoded in the same way as primary reinforcers remains an open question. To address this, we compared the encoding of #token and #juice. We first calculated the number of neurons encoding each variable. Neurons encoding #token in a window 500 ms after the token update and encoding #juice in a window 500 ms after juice delivery were classified as encoding both #token and #juice (Extended Data Fig. 3b). More neurons in the OFC than in the other areas encoded both #token and #juice (Fig. 4h; chi-square test, $p < 0.001$). We then

calculated the correlation of the #token and #juice regression coefficients in each area. Surprisingly, neuronal populations in the AMY and VS showed a higher similarity in encoding #token and #juice (Fig. 4i and Extended Data Fig. 3c; Fisher's z-transformation, $p < 0.05$). This result shows that the symbolic and primary reinforcers were more similarly encoded in the AMY and VS. Although OFC encoded both at the highest level, it did so with different population responses and, therefore, can discriminate these reinforcers best among the areas.

Value updates in the ventral striatum and orbitofrontal cortex

We used a stochastic reward schedule such that tokens were only updated in 75% of the trials, and token losses only occurred when tokens could be lost. The encoding of Δtoken (Fig. 3), therefore, may also include the encoding of RPE (Fig. 5a). To address this, we further fit the behaviors with a reinforcement learning model (Fig. 1g) and classified each neuron into Δtoken or RPE categories, depending on which variable best described the neuron's responses, using linear regression models (Extended Data Fig. 4a-b). RPE was calculated with the RW model. Neurons encoding Δtoken responded to the token outcome independently of the learning, whereas neurons encoding RPE encoded the difference between the token outcome and the predicted token outcome, with the prediction estimated by the RW model. Overall, more neurons in the OFC and VS encoded Δtoken or RPE than in the other two areas (Extended Data Fig. 4c-f; chi-square test, $p < 0.001$). We then split neurons based on whether their regression coefficients were positive or negative (e.g., whether they increased or decreased their firing rate for RPE; Fig. 5b-e). For example, a neuron was called a +RPE neuron when classified in the RPE category and with a positive regression coefficient.

Although a similar proportion of neurons in the OFC and VS encoded token outcome information, they showed different patterns. Most neurons in the OFC at the time of token outcome encoded -Δtoken and +RPE (Fig. 5b). In other words, they fired more when losing more tokens (-Δtoken) or when they did not lose tokens after choosing a loss option (+RPE). Specifically, this was aligned with the value gradient carried by the choices (Fig. 5b inset), even though it was a signal related to the outcome. On the other hand, most neurons in the VS encoded +Δtoken and +RPE (Fig. 5c). Neurons in the VS, therefore, fired more for all good outcomes when gaining or not losing more tokens. This was well aligned with the value gradient carried by the outcomes but not the choices (Fig. 5c inset). This suggests that value updates were referenced to outcomes in the VS and to choices in the OFC.

Populational representation of gains and losses

Next, we examined how neural populations in each area encoded gains and losses at the time of choice and token outcome by measuring population response trajectories[26]. This analysis used pseudo-populations composed of 500 neurons recorded across sessions. We focused on responses in a specific low-dimensional subspace that captured variance due to the eight combinations of cValue and Δtoken (four cValues × Δtoken or 0 for each cValue). We first defined the axes of the task-related subspace using linear regression coefficients. Then, the condition-averaged population response was projected on each axis to estimate the representation of the corresponding task variables across time (Extended Data Fig. 2d).

The one-dimensional trajectories reflected the dominant tuning at the single-cell level, and correlations between coding of cValue and Δtoken at the population level (Figs. 3-5). Trajectories diverged into groups aligned with the value gradient following cue onset for cValue (Fig. 6a-d) and following token update for Δtoken (Fig. 6e-h). Because all neurons were z-transformed before these analyses, the axes reflect the strength of encoding the corresponding variable. OFC showed the strongest and most balanced representation of positive and negative choices and outcomes among the four areas, with larger divergence among trajectories in both the cValue (Fig. 6a) and Δtoken (Fig. 6e) axes. The largest deviations for OFC, however, were for Loss conditions. OFC trajectories also showed the population representation of RPEs, with zero token outcomes for Loss options merged with Gian outcomes and zero token outcomes for Gian options merged with Loss outcomes (Fig. 6e). Because cValue and Δtoken were correlated in OFC, the cValue trajectories also crossed following token update (Fig. 6a). VS population activity also discriminated gains and losses well, but with a bias to the Gian conditions. The Gain trajectories had steeper peaks (Fig. 6b) and were more isolated (Fig. 6b, f) than the Loss trajectories. The VS also showed the crossover in cValue trajectories after outcomes (Fig. 6b). AMY trajectories grouped into Gain and Loss groups for cValue (Fig. 6c), indicating overall populational coding of Gain vs. Loss. However, the Δtoken trajectories, since about 250 ms after the token update, aligned with absolute values in both Gain and Loss conditions (Fig. 6g), such that all zero token outcomes had the weakest values, and both gains and losses were more positive, indicating the overall populational coding of the salience of Δtoken. MDt trajectories showed similar patterns to those of OFC, including the crossover following zero token outcomes. This suggests that they shared similar information. Overall, these results are consistent with the results shown in Figs. 3-5, which confirm the unique contributions of each area in encoding gains and losses.

Flow of token information in the ventral network

Brain areas function as part of big networks but not as individual isolated areas. The task-relevant information also is not isolated in each area. To understand the flow of token information within the network, we measured the linear dynamics of the trajectories within a trial. This analysis used only simultaneously recorded ensembles and was carried out trial-by-trial. First, we projected the population neural activity from single trials into a 3-D subspace to generate the population state-space representation of #token, Δtoken, and cDir (Extended Data Fig. 2d). Then, we put these task-variable-specific latent variables together into a matrix X. The rows included all the latent variables from all the areas, and the columns included each trial (Fig. 7a). We stacked all the sessions (N = 16) with more than ten neurons recorded from each area simultaneously, then estimated the loading matrix A (Fig. 7b). The matrix A characterized the flow of information among variables and areas (Extended Data Fig. 2e), with its columns representing the source areas and variables, and the rows representing the target areas and variables (Fig. 7b). The matrices were estimated separately at each point in time, as a local linear approximation to the dynamics, which were likely nonlinear.

The eigenvalues of the A matrices captured the time constant of the dynamics of information flow among the areas in the network. The top two eigenvalues, which show two peaks, indicate robust information flow during the cue and token update epochs, and the third eigenvalue captured information flow during the token update epoch (Fig. 7d). The left and right eigenvectors of the matrix A capture the temporal dynamics of output (Extended Data Fig. 5a-c) and input (Extended Data Fig. 5d-f) information. The first eigenvector also reflected the temporal structure of the task, showing phasic activity that was locked to the associated task variables. The values in the A matrices indicate the strength of information flow between a specific pair of areas and task variables. The temporal dynamics of the information flow varied across the time course of the trial (Fig. 7c). We examined the flow of token information across areas and task variables (Extended Data Fig. 6) and found the strongest information flow existed within the same area for most conditions. This is consistent with other work showing that within-region dynamics are higher dimensional than across-region dynamics [27].

Our dynamics analysis showed good specificity between task variables. For example, despite the fact that cDir, #token, and Δtoken were represented across areas, there were minimal interactions between cDir and the value signals (Extended Data Fig. 6f-h). Aiming to address the flow of token information, we focused on the interactions within and between the Δtoken and #token. For Δtoken (Fig. 7e), information flows showed peaks phase-locked to the token update. There was strong reciprocal information flow between the

OFC and MDt, reflecting the underlying anatomy [2]. We also found flow of information from AMY to the other areas, especially VS. For the #token (Fig. 7f), the information flow was relatively continuous along the time course of the trial, which is consistent with the tokens being present on the screen across trials.

More interestingly, there was a temporal derivative in the flow of information from #token to Δtoken in the OFC and VS (Fig. 7g, 100 ms before vs. after derivative; Fisher's z-transformation; VS to VS, $p < 0.001$; OFC to OFC, $p < 0.05$). This shows the mechanism by which Δtoken was calculated in the network. Specifically, OFC and VS computed the difference in #token following and prior to the choice (i.e., $\Delta token = \#token_{after} - \#token_{before}$). Also, we found a stronger flow of information from the #token to Δtoken than from the Δtoken to #token (Extended Data Fig. 6b, d; Fisher's z-transformation, $p < 0.05$). This suggests that the updating of values in this four-area network originated predominantly in the VS and OFC, consistent with the single-cell results (Fig. 5b-c). Because both areas showed time-derivative dynamics, it seems likely that they may both be calculating the update. However, the derivative in the OFC appeared about 100 ms earlier than in the VS, similar to the time of single neurons in the OFC leading Δtoken encoding relative to the VS (Fig. 2b inset). During this calculation, phasic reciprocal information flow also existed between OFC-MDt and MDt-VS (Fig. 7g), indicating that MDt bridged the information flow from the OFC to VS.

297 # Discussion

298 We examined the single-cell, population, and network representation of token-based RL across the ventral
299 cortico-striatal-thalamo-cortical network [2, 19, 28]. We found that although a similar fraction of the population
300 across these areas encoded the task variables, they encoded the information in different ways. VS showed
301 a strong bias towards encoding appetitive choices and outcomes, including positive RPEs, with increased
302 firing rates. On the other hand, OFC showed a more balanced encoding of positive and negative choices
303 and outcomes with increased and decreased firing rates. OFC also strongly encoded primary and symbolic
304 reinforcers but distinguished them at the population level. AMY, on the other hand, encoded both
305 reinforcers similarly at the population level. AMY also showed enhanced salience coding and did not show
306 correlations between choice and outcome value coding. Finally, MDt showed enhanced coding of actions
307 and mediated information flow between the OFC and VS during token update calculations.

308 ## Primary and symbolic reinforcers

309 Human studies typically use symbolic reinforcers, while animal studies use primary reinforcers. However,
310 similarities and differences between the population coding of primary and symbolic reinforcers have not
311 been examined. AMY and, to some extent, VS populations showed stronger correlations between primary
312 and symbolic reinforcers and, therefore, encoded tokens and primary reinforcers similarly. This may
313 explain why earlier lesion work found AMY linked conditioned stimuli to the specific reward properties of
314 the unconditioned stimuli they predicted [29]. OFC showed the most substantial coding of both primary and
315 symbolic reinforcers but encoded them with the lowest similarity, therefore discriminating them well. Thus,
316 OFC represents value information with high fidelity, using a code that preserves detailed information about
317 the type and valence of reinforcement [30].

318 ## Gains and losses

319 We found that VS showed a strong bias towards monotonically encoding gains, specifically coding
320 rewarding choices and outcomes, and better than the expected outcomes, with increased firing rates. Thus,
321 VS consistently signaled token gains with increased firing rates. This is consistent with human imaging
322 studies [31] and the longstanding suggestion that VS is important for motivation [32, 33]. AMY neurons showed
323 enhanced tuning for salience. The AMY population, uniquely among the areas, did not show a correlation
324 between linear choice value and outcome value encoding, and showed negative correlations, at the
325 population level, between gain and loss encoding, both of which are also consistent with salience coding.
326 Previous studies found AMY responses to appetitive and aversive stimuli but could not assess salience

because they only used single outcomes for each valence [34, 35]. OFC showed monotonic encoding of value across both positive and negative outcomes. Unlike VS and AMY, OFC neurons encoded values with both positive and negative slopes with a slight bias towards neurons responding more for negative value choices and loss outcomes. Recordings in the pregenual anterior cingulate cortex (ACC) [36] and the insula [21] similarly found neuronal populations coding appetitive and aversive choices with both positive and negative slopes. Thus, cortical areas show a more balanced coding of gains and losses than the VS and AMY.

OFC contributes to the encoding of symbolic reinforcers

We found that OFC encoded accumulated tokens and changes in tokens at both the single-cell and population levels. These findings are consistent with previous work showing that OFC codes state information about the environment [37-40], as tokens and token updates define states and state transitions in this task [41]. State representation is also referred to as a cognitive map, particularly when states have to be inferred [42, 43]. It is a map because states are nodes on graphs, and one has to know not only the current state, but also subsequent states to which one can transition. Within reinforcement learning, states are the environmental variables relevant to the learning and action selection process [44]. We also found that OFC encoded choice value, which is consistent with previous work showing that OFC codes economic value [45]. However, we have found that choice value was broadly encoded across our network, and OFC did not encode value more strongly than other areas. We did find differences in the way in which OFC encoded values relative to the other areas, as discussed above. OFC has a high-fidelity representation of the state that includes negative outcomes and the strongest encoding of accumulated tokens.

The calculation of value updates

We also examined network computations within targeted information dimensions. Analysis of single-trial communication using neurons recorded from multiple regions simultaneously can provide evidence of dynamic network processes that underlie behaviors [46]. Most studies that have examined interactions between areas with neurophysiology data have focused on pairwise interactions between cortical regions and have also not identified dynamic signatures of computations. Instead, they have shown that interactions occur within specific dimensions and often during particular periods within a task. For example, an early study found that choice-relevant signals were relayed from the prefrontal to the parietal cortex to guide behavior [47]. A more recent study found that stronger value coding in the OFC led to accelerated ramping of signals in the ACC [48]. And we recently identified object-to-direction information flow among lateral prefrontal cortex subregions in a task in which the location of a valuable stimulus had to be identified before a saccade could be directed toward it to make a choice [49].

358 The current study measured the flow of value information using simultaneous population recordings from
359 four areas in the ventral network. This allowed us to control for multiple, although not all, inputs to each
360 area. We found that token update information was calculated in the OFC and VS as a time-derivative of the
361 information about accumulated tokens (i.e., $\Delta token_{t+1} = \#token_{t+1} - \#token_t$). The signal was earlier
362 in the OFC but stronger in the VS. Thus, the change in accumulated tokens drove information about token
363 outcomes. This is consistent with the finding that OFC and VS strongly encoded Δtoken and token
364 prediction errors at the single-cell level. Note that in previous work with probabilistic reward outcomes, we
365 did not find strong encoding of RPE in these structures [12, 42]. We found minimal support for the alternative
366 hypothesis that token outcomes were integrated to generate information about accumulated tokens (i.e.,
367 $\#token_{t+1} = \#token_t + \Delta token_t$). Thus, analysis of neural dynamics identified an explicit computational
368 correlate of a behavioral process. We also found that MDt mediated interactions between the OFC and VS
369 with respect to token updates, by mediating reciprocal information flow between OFC-MDt and MDt-VS.
370 Interestingly, we did not identify a direct interaction between the OFC and the VS, even though they are
371 connected directly [50, 51]. This suggests that the token updates were mediated within the cortical-thalamic-
372 striatal circuit but not the corticostriatal circuit.

373 Effects of loss of specific nodes on reinforcement learning

374 These results also provide insight into the task-dependent effects of lesions in previous work. For example,
375 tasks that use probabilistic delivery of primary reinforcers have shown deficits following lesions to all nodes
376 of the ventral network [5, 6, 10, 11, 13, 52]. However, lesions of the AMY have almost no effect on learning in the
377 tokens task, and lesions of the VS only affect learning to discriminate gain magnitude but not gains vs.
378 losses [17, 18]. Here, we found that AMY showed enhanced salience coding. Salience coding can be used as a
379 learning signal in the context of probabilistic reward outcomes because salience and gains relative to no
380 outcome are equivalent. However, salience coding cannot be used for learning in token-based reinforcement
381 because large gains and large losses are represented similarly. This may also explain why we found no
382 evidence for the effects of AMY lesions on unsigned prediction error variables in Pearce-Hall models of
383 RL that had been reported previously [53] when examined in probabilistic reward tasks [13]. We also found that
384 VS was strongly biased toward encoding appetitive outcomes. Therefore, lesions to the VS would be
385 expected to have a larger effect on discriminating among gains, as opposed to between gains and losses [17].
386 Although we did not record in the insula, recent work has shown that the insula has substantial coding of
387 losses in token-based decision-making [21], consistent with early fMRI results [54]. Therefore, it is possible that
388 manipulations of the insula would show effects specific to loss outcomes. Finally, we found that OFC may

preferentially mediate the calculation of token reinforcement, and therefore, manipulations of OFC may have large effects on learning in the tokens task.

Conclusion

Information about task variables was represented across the ventral network. Although all areas represented task variables, they did so differently. AMY encoded choice and outcome salience, and also encoded the primary and symbolic reinforcers similarly. VS and OFC encoded value information, with VS strongly biased towards coding positive outcomes with increased firing rates and OFC more balanced towards coding positive and negative outcomes. OFC and VS calculated Δtoken as the time-derivative of accumulated tokens, with a shorter latency in the OFC. MDt mediated the interactions about token updates between the OFC and VS during this process. Importantly, the computations relevant to symbolic reinforcement appear to be more strongly mediated by cortical structures, in this case, OFC, than subcortical structures, which may differ from primary reinforcement.

REFERENCES

1. Alexander, G.E., DeLong, M.R. & Strick, P.L. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu. Rev. Neurosci.* **9**, 357-381 (1986).
2. Averbeck, B.B. & Murray, E.A. Hypothalamic Interactions with Large-Scale Neural Circuits Underlying Reinforcement Learning and Motivated Behavior. *Trends Neurosci.* **43**, 681-694 (2020).
3. Kim, C.K., Adhikari, A. & Deisseroth, K. Integration of optogenetics with complementary methodologies in systems neuroscience. *Nat. Rev. Neurosci.* **18**, 222-235 (2017).
4. Rich, E.L. & Averbeck, B.B. Neurophysiology: Differentiating Functional Contributions across Prefrontal Cortex. in *The Frontal Cortex: Organization, Networks, and Function* (ed. M.T. Banich, S.N. Haber & T.W. Robbins) (MIT Press, Cambridge, MA, 2024).
5. Chakraborty, S., Kolling, N., Walton, M.E. & Mitchell, A.S. Critical role for the mediodorsal thalamus in permitting rapid reward-guided updating in stochastic reward environments. *Elife* **5** (2016).
6. Walton, M.E., Behrens, T.E., Buckley, M.J., Rudebeck, P.H. & Rushworth, M.F. Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. *Neuron* **65**, 927-939 (2010).
7. Betzel, R.F. & Bassett, D.S. Multi-scale brain networks. *NeuroImage* **160**, 73-83 (2017).
8. Chau, B.K.*, et al.* Contrasting Roles for Orbitofrontal Cortex and Amygdala in Credit Assignment and Learning in Macaques. *Neuron* **87**, 1106-1118 (2015).
9. Houk, J.C., Adams, J.L. & Barto, A.G. A model of how the basal ganglia generate and use neural signals that predict reinforcement. in *Models of information processing in the basal ganglia.* 249-270 (The MIT Press, Cambridge, MA, US, 1995).
10. Rudebeck, P.H., Saunders, R.C., Lundgren, D.A. & Murray, E.A. Specialized Representations of Value in the Orbital and Ventrolateral Prefrontal Cortex: Desirability versus Availability of Outcomes. *Neuron* **95**, 1208-1220 e1205 (2017).
11. Rudebeck, P.H., Saunders, R.C., Prescott, A.T., Chau, L.S. & Murray, E.A. Prefrontal mechanisms of behavioral flexibility, emotion regulation and value updating. *Nat. Neurosci.* **16**, 1140-1145 (2013).
12. Costa, V.D., Mitz, A.R. & Averbeck, B.B. Subcortical Substrates of Explore-Exploit Decisions in Primates. *Neuron* **103**, 533-545 e535 (2019).
13. Costa, V.D., Dal Monte, O., Lucas, D.R., Murray, E.A. & Averbeck, B.B. Amygdala and Ventral Striatum Make Distinct Contributions to Reinforcement Learning. *Neuron* **92**, 505-517 (2016).
14. Rothenhoefer, K.M.*, et al.* Effects of Ventral Striatum Lesions on Stimulus-Based versus Action-Based Reinforcement Learning. *J. Neurosci.* **37**, 6902-6914 (2017).
15. Joel, D., Niv, Y. & Ruppin, E. Actor-critic models of the basal ganglia: new anatomical and computational perspectives. *Neural Netw.* **15**, 535-547 (2002).
16. Takahashi, Y., Schoenbaum, G. & Niv, Y. Silencing the critics: understanding the effects of cocaine sensitization on dorsolateral and ventral striatum in the context of an actor/critic model. *Front. Neurosci.* **2**, 86-99 (2008).
17. Taswell, C.A., Costa, V.D., Murray, E.A. & Averbeck, B.B. Ventral striatum's role in learning from gains and losses. *Proc Natl Acad Sci U S A* **115**, E12398-E12406 (2018).
18. Taswell, C.A., Janssen, M., Murray, E.A. & Averbeck, B.B. The motivational role of the ventral striatum and amygdala in learning from gains and losses. *Behav. Neurosci.* **137**, 268-280 (2023).
19. Averbeck, B. & O'Doherty, J.P. Reinforcement-learning in fronto-striatal circuits. *Neuropsychopharmacology* **47**, 147-162 (2022).
20. Seo, H. & Lee, D. Behavioral and neural changes after gains and losses of conditioned reinforcers. *J. Neurosci.* **29**, 3627-3641 (2009).

21. Yang, Y.P., Li, X. & Stuphorn, V. Primate anterior insular cortex represents economic decision variables proposed by prospect theory. *Nat Commun* **13**, 717 (2022).
22. Tang, H., Costa, V.D., Bartolo, R. & Averbeck, B.B. Differential coding of goals and actions in ventral and dorsal corticostriatal circuits during goal-directed behavior. *Cell Rep.* **38**, 110198 (2022).
23. Parker, N.F.*, et al.* Choice-selective sequences dominate in cortical relative to thalamic inputs to NAc to support reinforcement learning. *Cell Rep.* **39**, 110756 (2022).
24. Bartolo, R. & Averbeck, B.B. Prefrontal Cortex Predicts State Switches during Reversal Learning. *Neuron* **106**, 1044-1054 e1044 (2020).
25. Blanco-Pozo, M., Akam, T. & Walton, M.E. Dopamine-independent effect of rewards on choices through hidden-state inference. *Nat. Neurosci.* **27**, 286-297 (2024).
26. Mante, V., Sussillo, D., Shenoy, K.V. & Newsome, W.T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78-84 (2013).
27. Semedo, J.D., Zandvakili, A., Machens, C.K., Yu, B.M. & Kohn, A. Cortical Areas Interact through a Communication Subspace. *Neuron* **102**, 249-259 e244 (2019).
28. Averbeck, B.B. & Costa, V.D. Motivational neural circuits underlying reinforcement learning. *Nat. Neurosci.* **20**, 505-512 (2017).
29. Cardinal, R.N., Parkinson, J.A., Hall, J. & Everitt, B.J. Emotion and motivation: the role of the amygdala, ventral striatum, and prefrontal cortex. *Neurosci Biobehav Rev* **26**, 321-352 (2002).
30. McDannald, M.A.*, et al.* Orbitofrontal neurons acquire responses to 'valueless' Pavlovian cues during unblocking. *Elife* **3**, e02653 (2014).
31. Yacubian, J.*, et al.* Dissociable systems for gain- and loss-related value predictions and errors of prediction in the human brain. *J. Neurosci.* **26**, 9530-9537 (2006).
32. Mogenson, G.J., Jones, D.L. & Yim, C.Y. From motivation to action: functional interface between the limbic system and the motor system. *Prog. Neurobiol.* **14**, 69-97 (1980).
33. Shiflett, M.W. & Balleine, B.W. At the limbic-motor interface: disconnection of basolateral amygdala from nucleus accumbens core and shell reveals dissociable components of incentive motivation. *Eur. J. Neurosci.* **32**, 1735-1743 (2010).
34. Belova, M.A., Paton, J.J., Morrison, S.E. & Salzman, C.D. Expectation modulates neural responses to pleasant and aversive stimuli in primate amygdala. *Neuron* **55**, 970-984 (2007).
35. Paton, J.J., Belova, M.A., Morrison, S.E. & Salzman, C.D. The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature* **439**, 865-870 (2006).
36. Amemori, K. & Graybiel, A.M. Localized microstimulation of primate pregenual cingulate cortex induces negative decision-making. *Nat. Neurosci.* **15**, 776-785 (2012).
37. Schuck, N.W., Cai, M.B., Wilson, R.C. & Niv, Y. Human Orbitofrontal Cortex Represents a Cognitive Map of State Space. *Neuron* **91**, 1402-1412 (2016).
38. Basu, R.*, et al.* The orbitofrontal cortex maps future navigational goals. *Nature* **599**, 449-452 (2021).
39. Costa, K.M.*, et al.* The role of the lateral orbitofrontal cortex in creating cognitive maps. *Nat. Neurosci.* **26**, 107-115 (2023).
40. Gardner, M.P.H. & Schoenbaum, G. The orbitofrontal cartographer. *Behav. Neurosci.* **135**, 267-276 (2021).
41. Burk, D.C., Taswell, C., Tang, H. & Averbeck, B.B. Computational mechanisms underlying motivation to earn symbolic reinforcers. *bioRxiv*, 2023.2010. 2011.561900 (2023).
42. Costa, V.D. & Averbeck, B.B. Primate Orbitofrontal Cortex Codes Information Relevant for Managing Explore-Exploit Tradeoffs. *J. Neurosci.* **40**, 2553-2561 (2020).
43. Wilson, R.C., Takahashi, Y.K., Schoenbaum, G. & Niv, Y. Orbitofrontal cortex as a cognitive map of task space. *Neuron* **81**, 267-279 (2014).

44. Averbeck, B.B. Theory of choice in bandit, information sampling and foraging tasks. *PLoS Comput. Biol.* **11**, e1004164 (2015).
45. Padoa-Schioppa, C. & Cai, X. The orbitofrontal cortex and the computation of subjective value: consolidated concepts and new perspectives. *Ann. N. Y. Acad. Sci.* **1239**, 130-137 (2011).
46. Vinck, M.*, et al.* Principles of large-scale neural interactions. *Neuron* **111**, 987-1002 (2023).
47. Crowe, D.A.*, et al.* Prefrontal neurons transmit signals to parietal neurons that reflect executive control of cognition. *Nat. Neurosci.* **16**, 1484-1491 (2013).
48. Balewski, Z.Z., Elston, T.W., Knudsen, E.B. & Wallis, J.D. Value dynamics affect choice preparation during decision-making. *Nat. Neurosci.* **26**, 1575-1583 (2023).
49. Tang, H., Bartolo, R. & Averbeck, B.B. Reward-related choices determine information timing and flow across macaque lateral prefrontal cortex. *Nat Commun* **12**, 894 (2021).
50. Haber, S.N. & Knutson, B. The reward circuit: linking primate anatomy and human imaging. *Neuropsychopharmacology* **35**, 4-26 (2010).
51. Averbeck, B.B., Lehman, J., Jacobson, M. & Haber, S.N. Estimates of projection overlap and zones of convergence within frontal-striatal circuits. *J. Neurosci.* **34**, 9497-9505 (2014).
52. Rudebeck, P.H., Ripple, J.A., Mitz, A.R., Averbeck, B.B. & Murray, E.A. Amygdala Contributions to Stimulus-Reward Encoding in the Macaque Medial and Orbital Frontal Cortex during Learning. *J. Neurosci.* **37**, 2186-2202 (2017).
53. Li, J., Schiller, D., Schoenbaum, G., Phelps, E.A. & Daw, N.D. Differential roles of human striatum and amygdala in associative learning. *Nat. Neurosci.* **14**, 1250-1252 (2011).
54. Pessiglione, M., Seymour, B., Flandin, G., Dolan, R.J. & Frith, C.D. Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature* **442**, 1042-1045 (2006).
55. Hwang, J., Mitz, A.R. & Murray, E.A. NIMH MonkeyLogic: Behavioral control and data acquisition in MATLAB. *J. Neurosci. Methods* **323**, 13-21 (2019).
56. Chaure, F.J., Rey, H.G. & Quian Quiroga, R. A novel and fully automatic spike-sorting implementation with variable number of features. *J. Neurophysiol.* **120**, 1859-1871 (2018).
57. Olejnik, S. & Algina, J. Measures of Effect Size for Comparative Studies: Applications, Interpretations, and Limitations. *Contemp. Educ. Psychol.* **25**, 241-286 (2000).

Methods

Subjects

The experiments were performed on one adult male (10 kg) and one adult female (7.5 kg) rhesus macaque (*Macca mulatta*). They were 8-10 years old. The monkeys were pair-housed when possible and had access to food 24 hours per day. On testing days, the monkeys were placed on water control and earned their juice through performing the task. On non-testing days, the monkeys were given *ad libitum* access to water.

Ethics

Experimental procedures for all monkeys followed *the Guide for the Care and Use of Laboratory Animals* and were approved by the National Institute of Mental Health Animal Care and Use Committee.

Experimental Setup

Monkeys were trained to perform a saccade-based two-armed bandit task. Stimuli were presented on a 19-inch LCD monitor situated 40 cm from the monkeys' eyes. During training and testing, the monkeys sat in a primate chair with their heads restrained. Stimulus presentation and behavioral monitoring were controlled by MonkeyLogic [55]. The eye movements were monitored at 400 fps using a Viewpoint eye tracker (Arrington Research, Scottsdale, AZ) and sampled at 1 kHz. A fixed amount of apple juice was delivered through a pressurized plastic tube gated by a solenoid valve on rewarded trials.

Task Design

The task was developed and first used in our previous study [17]. Each session had nine blocks. Each block used four novel images associated with different values, including +2, +1, −1, and −2. The monkeys obtained more tokens by choosing images associated with larger values. Choosing one of the images led to gaining or losing a corresponding number of tokens. Also, token numbers could not be negative, so choosing a loss image when there were no accumulated tokens had no effect. We used a stochastic reward schedule. The number of tokens updated 75% of the time and did not change 25% of the time. To complete a trial successfully, the monkey first acquired and held central fixation for 400-600 ms. Then, two of the four images were randomly selected by the computer and displayed on the screen. The animal made its selection by saccading to one of them. The unchosen image disappeared when the monkey reached the target. Saccade fixation was maintained on the chosen image for 500 ms. After that, the chosen image also disappeared, and the token number was updated according to the chosen image. Tokens accumulated across trials and were cashed out for juice every four to six trials, with the interval randomly selected. At cash-

551 out, the animals were given one drop of juice for each token. When each drop of juice was delivered, one
552 token was removed from the screen. Up to 12 tokens were accumulated and were displayed on the screen
553 across trials.

554 The task had six individual conditions defined by the possible values of the image pairs. The conditions
555 within a block of 108 trials (6 conditions × 2 counterbalanced for left and right sides × 9 repetitions) were
556 presented pseudo-randomly. The animals saw each condition twice, once on the left and once on the right,
557 every 12 trials before seeing any condition a third time. We introduced four novel images at the beginning
558 of each block. Images provided as choice options were normalized for luminance and spatial frequency
559 using the SHINE toolbox for MATLAB, described previously [17].

560 Surgical procedures

561 Each monkey was surgically implanted with a titanium headpost, and a 25 × 35 mm recording chamber to
562 allow vertical grid access to the OFC, VS, AMY, and MDt (Fig. 1d-f). Grid holes for the MDt had 16°
563 angles to allow better access to the target area (Extended Data Fig. 1d). Chamber placements were planned
564 and verified with T1 and T2 magnetic resonance imaging (MRI, 3.0 T). Small burr holes were drilled above
565 each target area. A grid was installed, and one guide tube was inserted through each burr hole. The guide
566 tubes were lowered to about 1-4 mm (depending on the target areas) above the target areas and were glued
567 to the grid. We removed the guide tubes and placed new ones in the adjacent locations after 3-5 recording
568 days (Extended Data Fig. 1d-e). The locations of the guide tubes were verified with MRI after guide tube
569 replacement. All sterile surgeries and MRI scans were performed under anesthesia.

570 Neurophysiological Recordings

571 Neurophysiology recordings (Extended Data Fig. 1c-d) began after the monkeys had recovered from the
572 surgery. We lowered one linear electrode array (V-probe, Plexon Inc, Dallas, TX) into each guide tube on
573 every recording day. Thirty-two channel electrodes with 150 μm inter-contact spacing probes were used in
574 the VS and MDt, and 64-channel electrodes with 150 μm inter-contact spacing probes were used in the
575 OFC and AMY. The probes were advanced to their target location by a four-channel micromanipulator
576 (NAN Instruments, Nazareth, Israel) attached to the recording chamber. The depths of the neurons were
577 estimated by their recording locations relative to the tip of the guide tubes (verified with MRI).
578 Electrophysiological data were acquired with a 512-channel Grapevine System (Ripple, Salt Lake City,
579 UT). The spike acquisition threshold was set at a 4.0 × root mean square (RMS) of the baseline signal for
580 each electrode. Behavioral event markers from MonkeyLogic and eye-tracking signals from Viewpoint

were sent to the Ripple acquisition system. The extracellular signals were high-pass filtered (1 kHz cutoff) and digitized at 30 kHz to acquire the single-cell activity. Spikes were sorted offline via Wave_clus 3 [56].

Choice behavior

Each block had six conditions. The conditions pseudo-randomly appeared in the task. The number of observations in each condition increased as a function of learning. We quantified choice behavior during the task by measuring the fraction of choosing the image associated with a higher value in each condition. We then measured how different task variables affected the choice behavior by fitting a multi-way ANOVA model. Factors, including the number of observations (Trial), Gain/Loss condition (Loss/Loss, Loss/Gian, or Gain/Gain), value difference of the options (Δvalue, e.g., Δvalue = 4 in the condition of -2 vs. +2), the number of tokens before choice (#token), and the number of trials since last token cashout (CashID) were used in the model. The model was run session by session to measure the contribution of each variable in each session, using all the sessions to acquire the statistics for each variable. All trials in which monkeys chose one of the two stimuli were analyzed. Trials in which the monkey broke fixation, failed to make a choice, or attempted to saccade to more than one target were excluded.

Effect size

To quantify the contribution of each task variable to the monkey's choice. We computed each factor's effect size, ω^2, from the ANOVA model output. $\omega^2$is an unbiased estimator of the amount of variance in neural activity explained by each task variable, and ranges between -1 and 1 [57]. It is given by:

$$\omega^2 = \frac{df_{effect} \times (MS_{effect} - MS_{error})}{SS_{total} + MS_{error}}$$

where df_{effect} refers to the degrees of freedom associated with the factor, MS_{effect} refers to the mean squares, MS_{error} is the mean squared error, SS_{total} is the sum of squares of all factors.

Responsive neurons

To identify neural responses to different task components, we fit a sliding window multi-way ANOVA model to spike counts computed in 200 ms bins, advanced in 50 ms increments, and time-locked to token update. Factors, including the number of tokens on the screen (#token, may change after the choices), the change of token numbers (Δtoken), the drops of juice delivered (#juice), image pairs with different value combinations presented (condition, e.g., +2 vs. -1), the order of blocks (blockID), the identity (cStim), *a priori* value (cValue), and direction (cDir) of chosen images were used in the model. The stimulus identity

was the specific image used in each block to represent each outcome, which was the interaction of cValue and blockID. The factor blockID was used to remove non-stationarity due to drift.

Statistic test for the proportion of neurons

The binomial test was applied to test whether the proportion of responsive neurons was significantly above the chance level (5% most of the time). The chi-square test was used to compare the proportions of responsive neurons between different pairs of brain areas or among four areas. Significant encoding for each factor at each time bin was evaluated at $p < 0.05$. A neuron that showed a significant response to a factor in no less than three contiguous bins in the statistics was considered to be responsive to that factor.

Neuronal coding regression analysis

To quantitatively characterize how the updating of values (including cValue and Δtoken) was encoded in different areas, we examined the activity of each neuron using a series of multivariate linear regression models [21]. The dependent variable, neural activity, was first z-scored by subtracting the mean response from the firing rate at each time and in each trial and dividing the result by the standard deviation of the responses. Both the mean and the standard deviation were computed by combining the neurons' responses across all trials and times. We then described the z-scored responses of neuron i at time t as a linear combination of several task variables. The independent variables were the same as the factors used in the ANOVA model:

$$r_{i,t}(k) = \beta_{i,t,\mathrm{cValue}} * cValue(k) + \beta_{i,t,\Delta token} * \Delta token(k) + \cdots + \beta_{i,t,cDir} * cDir(k) + \varepsilon$$

where $r_{i,t}(k)$ is the z-scored response of neuron i at time t and on trial k, $\Delta token(k)$ is the change of tokens on trial k. The regression coefficients, $\beta_{i,t,f}$, describe how much the trial-by-trial firing rate of neuron i, at a given time t during the trial, depends on the corresponding task variable f.

We tested all potential combinations of tuning for Δtoken (six forms: General value, Salience, Gain/Loss, Gain, Loss, nan) and cValue (six forms: General value, Salience, Gain/Loss, Gain, Loss, nan). Nan means the variable was removed from the model. A particular variable could only be represented by one specific form but not by combinations of more than one form (e.g., a model containing the Δtoken variable could include Gain or Salience but not both). In total, we tested $6 \times 6 = 36$ models for each neuron. This method is also illustrated in Extended Data Fig. 2a. We determined the best-fitting model for each neuron using the Akaike information criterion (AIC). Neurons were classified into different functional categories (up to one category per variable) according to the combination of the forms of these two variables included in the best-fitting model. We computed the AIC:

$$AIC = 2p - 2\ln(\hat{L})$$

where p is the number of free parameters in the model, and $\hat{L}$ is the maximized value of the likelihood function.

Population coding similarity

The regression coefficient, $\beta_{i,t,f}$, from the multivariate linear regression model reflects the weight of one task variable, f, in explaining the variation of the neuron's activity at time t. The regression coefficients, $\overrightarrow{\beta_f}$, from a neural population, represent the weights of each neuron in the population encoding one task variable. Hence, computing the correlation between two regression coefficient vectors from the same neural population tells us how similar a neural population encodes two task variables (Extended Data Fig. 2b-c):

$$\mathrm{r} = \mathrm{corr}(\overrightarrow{\beta_{f1}}, \overrightarrow{\beta_{f2}})$$

where $\overrightarrow{\beta_f}$ is a vector that consists of regression coefficients of task variable f in a neural population. To assess the significance of the correlation between two regression coefficients, we transformed the correlation coefficients into normally distributed z-scores using Fisher's z-transformation.

Reward prediction error

The RPE was defined as the difference between the change of tokens $\Delta token(t)$ and the estimated value of the chosen image, which is given by:

$$RPE(t) = \Delta token(t) - v_i(t)$$

where t means trial order, and i means the chosen option among two options. The updating of value v_i was estimated using the Rescorla–Wagner equation:

$$v_i(t+1) = v_i(t) + \rho_h RPE(t)$$

where ρ is the learning rate. The updated value estimate $v_i(t+1)$ equals the previous value estimate $v_i(t)$ plus the RPE scaled by the cue-dependent learning rate ρ_h for images associated with different *a priori* values h. These values were then passed through a soft-max function to give choice probabilities for the image pairs presented in each trial:

$$d_j(t) = (1 + e^{\beta(v_i(t) - v_j(t))})^{-1}, d_i(t) = 1 - d_j(t)$$

where β is the choice consistency or inverse temperature parameter, fit across all six cue conditions, and i and j are the two choice options. We then maximized the likelihood of the animal's choices, D, given the parameters, using the cost function:

$$f(D|\alpha_i, \beta) = \prod_t [\, d_1(t)c_1(t) + d_2(t)c_2(t)$$

where $d_1(t)$ is the choice probability value for one option on trial t, and $c_1(t)$ and $c_2(t)$ are indicator variables that take on a value of 1 if the corresponding option was chosen, and 0 otherwise. This model was fit across blocks in each session for each monkey to give one set of free parameters for each session.

Targeted dimensionality reduction

We used the regression coefficients described above to identify dimensions in state space representing each task variable [26]. For each variable, we first build a set of coefficient vectors $\overrightarrow{\beta_{f,t}}$ whose entries $\beta_{f,t}(i)$ correspond to the regression coefficient for task variable f, time t, and neuron i. The vectors $\overrightarrow{\beta_{f,t}}$ (of length N_{unit}) are obtained by simply rearranging the entries of the vectors $\overrightarrow{\beta_{\iota,t}}$ (of length N_f) computed above. Each vector, $\overrightarrow{\beta_{f,t}}$, thus corresponds to a direction in state space that accounts for variance in the population response at time t, due to variation in task variable f.

We used time-dependent regression vectors, B_t, which is a matrix with each column corresponding to the one regression vector $\overrightarrow{\beta_{f,t}}$. We referred to them as the 'task-related axes'. These axes span a 'regression subspace' representing the task-related information the neural population coded. We then projected the single-trial population responses onto these axes to study the representation of the task-related variables in each trial:

$$X_t = {B_t}^T R_t$$

where R_t is single-trial population response at time t, which is of dimension $N_{unit} \times N_{trial}$. And X_t is the set of time-series vectors over all task variables and trials, which is of dimension $N_f \times N_{trial}$. It represents the population coding of task-related information at time t of every single trial. This method is also illustrated in Extended Data Fig. 2d.

Linear dynamics

We fit an autoregressive, linear model to the X_t,

$$X(t+1) = AX(t) + \varepsilon(t)$$

where $X(t) = [X_{OFC}(t), X_{VS}(t), X_{AMY}(t), X_{MDt}(t)]^T$ is a matrix of dimension $(N_{area} \cdot N_f) \times N_{trial}$, and $X_{area}(t) = [x_{k,\#token}(t), x_{k,\Delta token}(t), x_{k,cDir}(t)]$ is a matrix of dimension $N_{trial} \times N_f$. $\varepsilon(t)$ is the noise term. A is a 12 × 12 dynamics matrix, where the 12 = 4 × 3 dimensions correspond to the combinations of brain areas and task variables. The model was fit using data points in a non-overlapping 25 ms moving window. This resulted in a time-dependent estimate of the matrix A. Eigenvalues of A closer to 0 indicate a faster decay, and eigenvalues near 1 would correspond to a slower decay. The method is also illustrated in Fig. 7a-c and Extended Data Fig. 2e.

A z-test was applied to test for significant differences between the actual and shuffled results for the elements of the A matrices. The shuffled data were generated using the same matrix X but with randomized trial indexes across rows ($areas \times variables$). The actual results with a mean outside the 99% confidence interval of the shuffled results showed a significant response.

One-dimensional task-variable-specific trajectory

We also carried out the targeted dimensionality reduction using condition-averaged population response to study the populational representation of gains and losses. We projected the condition-averaged population responses onto these axes:

$$X_t = {B_t}^T R_t$$

where R_t is condition-averaged population response at time t, which is of dimension $N_{unit} \times N_{condition}$. And X_t is the set of time-series vectors over all task variables and conditions, which is of dimension $N_f \times N_{condition}$.

The variability of trajectories across different conditions was estimated using a bootstrap procedure. According to the null hypothesis, we generated data with no differences among the conditions by sampling with replacement from the combined set of all conditions. We then calculated the sum of the standard error of each condition from the mean of all conditions for each bootstrap condition set. We did this 1000 times. It gave us 1000 different sampled standard errors from the null distribution. We then compared the standard error of the actual data to the standard errors in the null distribution. If the actual standard error is in the 99% confidence interval of the null distribution, the trajectories significantly differ among the conditions (i.e., $p < 0.01$).

Data availability

The datasets supporting the current study will be publicly available as of the date of publication. DOIs are listed in the key resources table.

Code availability

All original codes will be publicly available as of the date of publication. DOIs are listed in the key resources table.

Acknowledgments

We thank Christos Constantinidis, Vincent Costa and Diana Burk for their valuable comments, as well as Andy Mitz, Craig Taswell, Miriam Janssen and Sarah Falkovic for their technical help. This work was supported by the Intramural Research Program of the National Institute of Mental Health (ZIA MH002928) to B. B. A. and BBRF NARSAD Young Investigator Award (30892) to H. T. Anatomical MRI scanning was carried out in the Neurophysiology Imaging Facility Core (NIMH, NINDS, NEI).

Author contributions

R. B., H. T., and B. B. A. designed research; H. T. and R. B. performed research; H. T. and B. B. A. analyzed data; H. T. and B. B. A. wrote the paper.

Competing interests

The authors declare no competing interests.

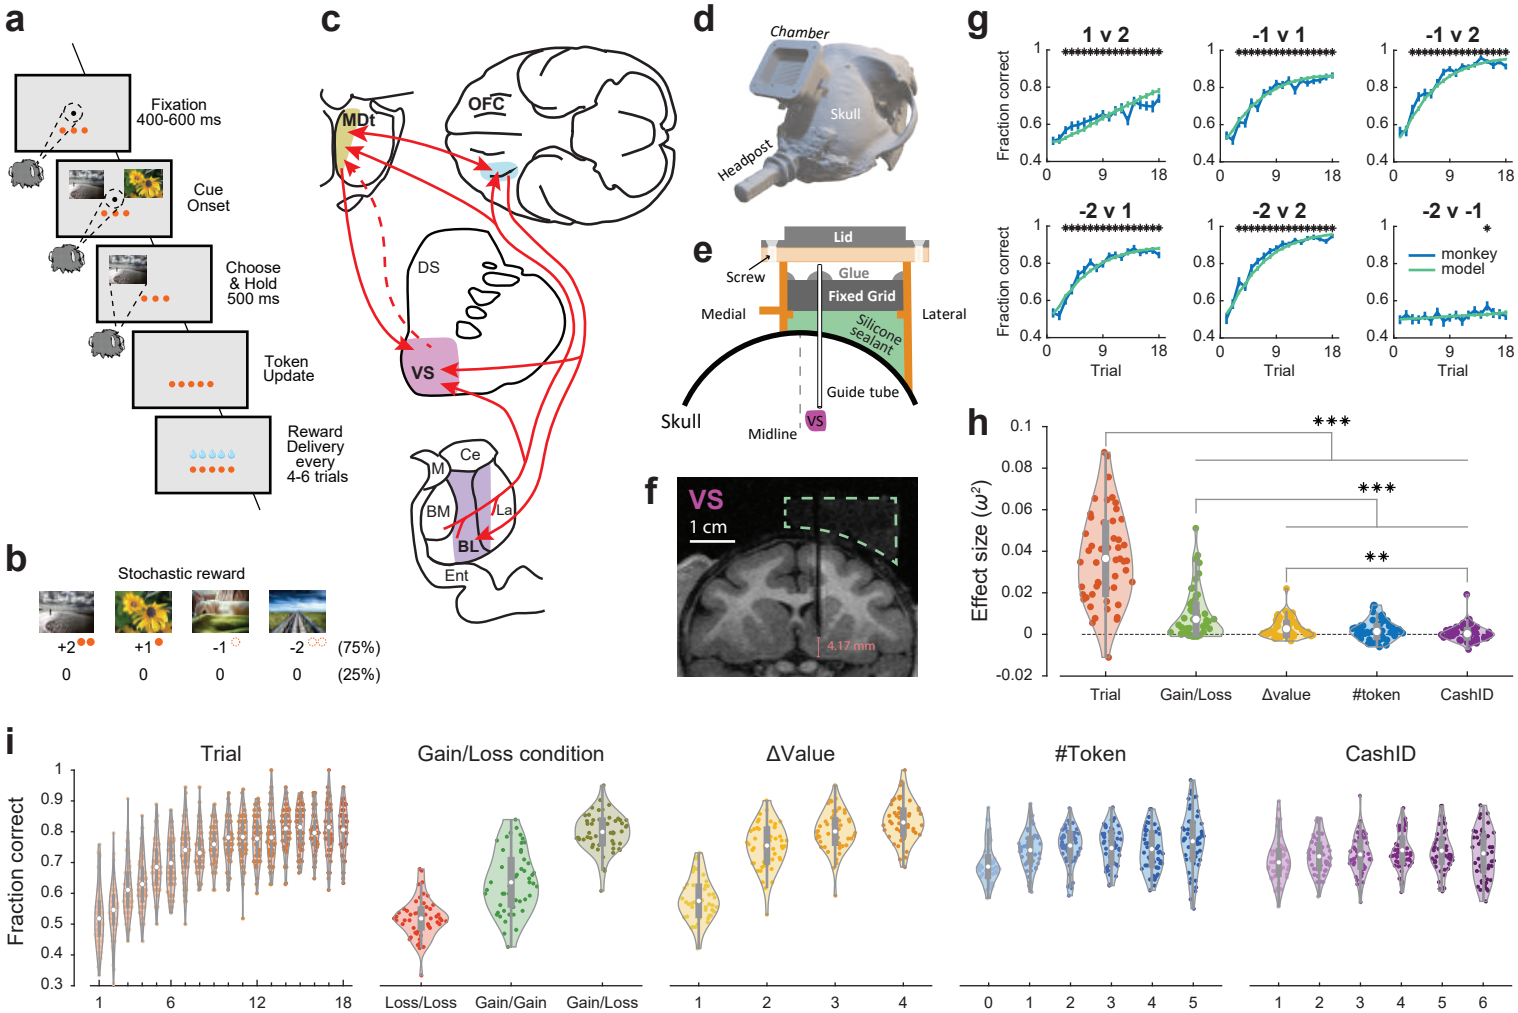
a
Fixation
400-600 ms
Cue
Onset
Choose
& Hold
500 ms
Token
Update
Reward
Delivery
every
4-6 trials
b
Stochastic reward
+2
+1
-1
-2
(75%)
0
0
0
0
(25%)
c
MDt
OFC
DS
VS
Ce
M
BM
La
BL
Ent
d
Chamber
Skull
Headpost
e
Lid
Screw
Glue
Fixed Grid
Medial
Lateral
Silicone sealant
Guide tube
Skull
Midline
VS
f
VS
1 cm
4.17 mm
g
1 v 2
-1 v 1
-1 v 2
-2 v 1
-2 v 2
-2 v -1
Fraction correct
Trial
monkey
model
h
Effect size (ω^2)
Trial
Gain/Loss
Δvalue
#token
CashID
i
Trial
Gain/Loss condition
ΔValue
#Token
CashID
Fraction correct
Loss/Loss
Gain/Gain
Gain/Loss

Figure 1

Figure legends

Fig. 1 | Task, behavior, recording method, and recorded areas.

a-b, Behavioral task. **a**, Structure of an individual trial. Successive frames illustrate the sequence of events. In each trial, only two of the four images were presented. Monkeys chose between them and gained or lost (stochastic: 75% change, 25% no change) the corresponding number of tokens. Choices could be made as soon as the images were shown. Accumulated tokens were cashed out every four to six trials, with one drop of juice for each token. **b**, In each block, the monkeys learned the values (+2, +1, −1, −2) of four novel images.

c, Schematic recording areas, including the orbitofrontal cortex (OFC), ventral striatum (VS), basolateral amygdala (AMY), and mediodorsal thalamus (MDt). Red arrows indicate the anatomical connections between these areas. The dashed line indicates an indirect connection.

d-f, Semi-chronic recording. **d,** A chamber was implanted. **e**, Then, guide tubes were inserted into each target area through a grid. The guide tubes stopped above the target areas and stayed in the brain across recording days. The empty space in the chamber was filled with silicone sealant to prevent potential infections. Multi-site linear probes were lowered into target areas through the guide tubes every recording day. **f**, A coronal section view of the chamber shows the shadow of a guide tube and the target area under MRI. The dashed green polygon indicates the silicone sealant.

g-i, Choice behavior. **g**, Fraction of choosing the image with a higher value in each condition. Blue lines represent monkey behavior, and green lines represent the prediction of the reinforcement learning model. Error bars represent mean ± SEM. The black asterisks indicate a significant difference in monkey behavior from the change level (one-sample t-test, $p < 0.01$). **h**, The effect size of each task variable. Paired t-test, $^{**}p < 0.01$, $^{***}p < 0.001$. Non-significant pairs are not indicated. **i**, Choice behavior as a function of the task variables. Results were averaged from two monkeys, n = 50 sessions. Violin plots showing the range of values calculated across all sessions (center dot, mean; box, 25th to 75th percentiles; whiskers, ± 1.5 × the interquartile range; dots, each session; shade, density curve).

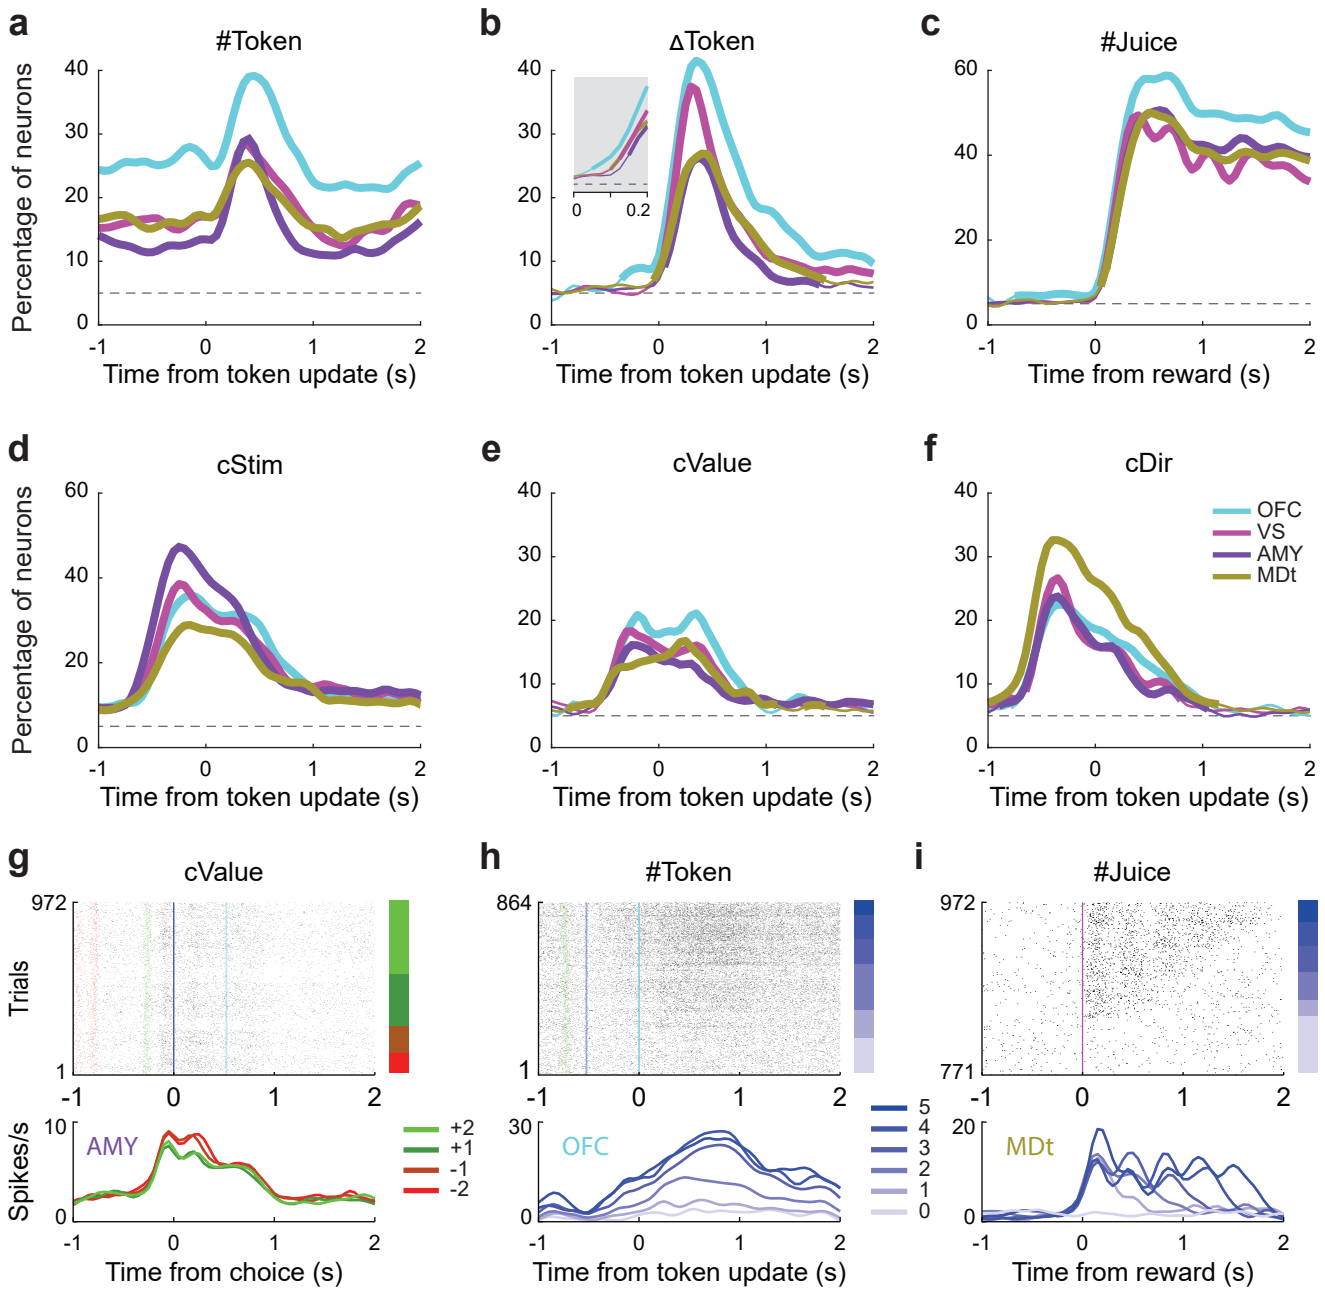
a
#Token
Percentage of neurons
Time from token update (s)
b
ΔToken
Time from token update (s)
c
#Juice
Time from reward (s)
d
cStim
Percentage of neurons
Time from token update (s)
e
cValue
Time from token update (s)
f
cDir
OFC
VS
AMY
MDt
Time from token update (s)
g
cValue
Trials
Spikes/s
AMY
+2
+1
-1
-2
Time from choice (s)
h
#Token
OFC
5
4
3
2
1
0
Time from token update (s)
i
#Juice
MDt
Time from reward (s)

Figure 2

Fig. 2 | Neural encoding of choices, primary and symbolic reinforcers.

a-f, Percentage of neurons in each area encoding #token (**a**), Δtoken (**b**), #juice (**c**), the identity (**d**), *a priori* value (**e**), and direction (**f**) of chosen images. Data were aligned to either the token update or the onset of juice delivery. Inset in **b**: ANOVA with 25 ms bin showing response latencies to Δtoken different among areas. Dashed horizontal lines represent the chance level. Thick lines indicate a significant difference between the corresponding area and chance level (binomial test, $p < 0.01$).

g-i, Example neurons encoded the *a priori* value of the chosen images (**g**), #token (**h**), and #juice (**i**). Each row in the raster plot represents the spikes during a trial. Red, green, blue, cyan, and magenta lines represent fixation, cue onset, choice, token update, and juice delivery. The bars on the right side represent the trial groups under different conditions. The bottom curves indicate mean activities, split by trial groups defined by the bars.

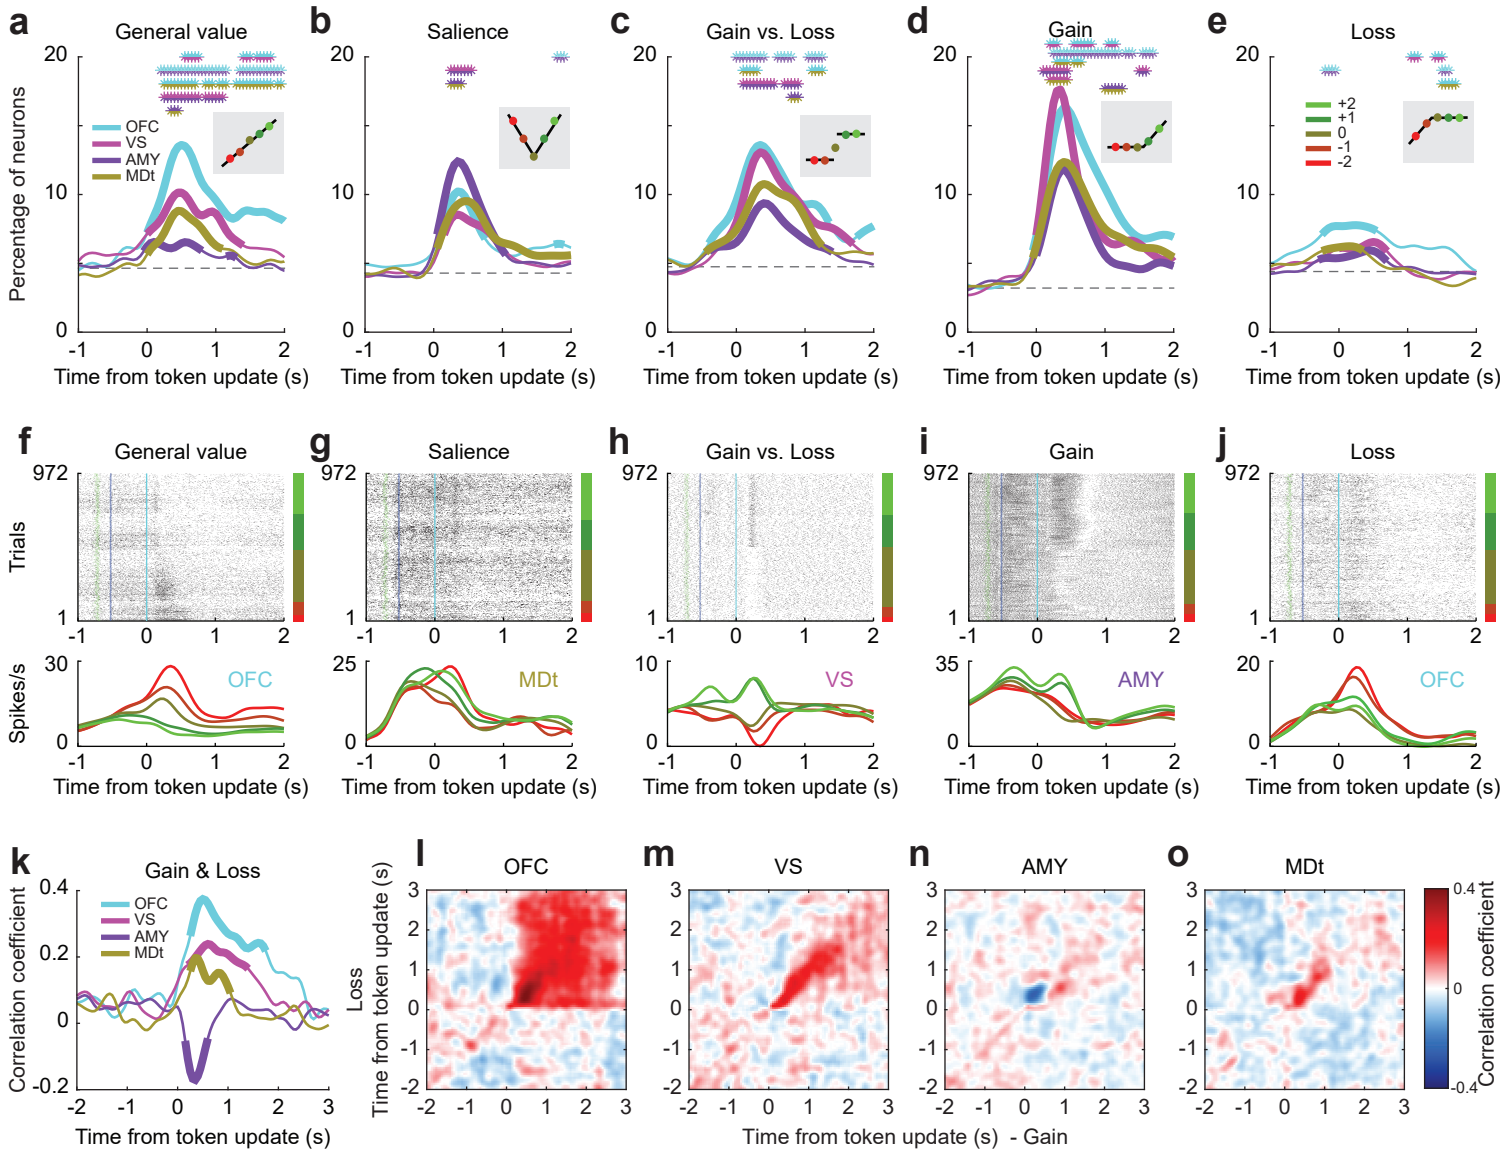
a
General value
b
Salience
c
Gain vs. Loss
d
Gain
e
Loss
Percentage of neurons
OFC
VS
AMY
MDt
+2
+1
0
-1
-2
Time from token update (s)
f
g
h
i
j
Trials
Spikes/s
k
Gain & Loss
Correlation coefficient
l
m
n
o
Loss
Time from token update (s) - Gain

Figure 3

Fig. 3 | Diverse representation of gains and losses.

a-e, Percentage of neurons in each area encoding Δtoken. The insets indicate the corresponding tuning curves. The coefficients could be positive or negative, meaning the tuning curves could be flipped vertically. **a**, General value: monotonic encoding values across gain and loss contexts. **b**, Salience: monotonic encoding values across gain and loss contexts but with opposite directions. **c**, Gain/Loss: categorical encoding values in gain and loss contexts. **d**, Gain: encoding values only in the gain context. **e**, Loss: encoding values only in the loss context. Dashed horizontal lines represent chance levels, defined by the mean of all areas during the fixation epoch. Thick lines indicate a significant difference between the corresponding area and chance level (binomial test, $p < 0.05$). The double-colored asterisks indicate a significant difference between pairs of areas indicated by the colors (chi-square test, $p < 0.05$).

f-j, Example neurons for each category.

k-o, Co-encoding of gains and losses. Correlations of the Gain and Loss regression coefficients at the same (**k**) or cross (**l-o**) time points for each area. Thick lines in **k** indicate a significant difference between the corresponding area and baseline level during the fixation period (binomial test, $p < 0.05$).

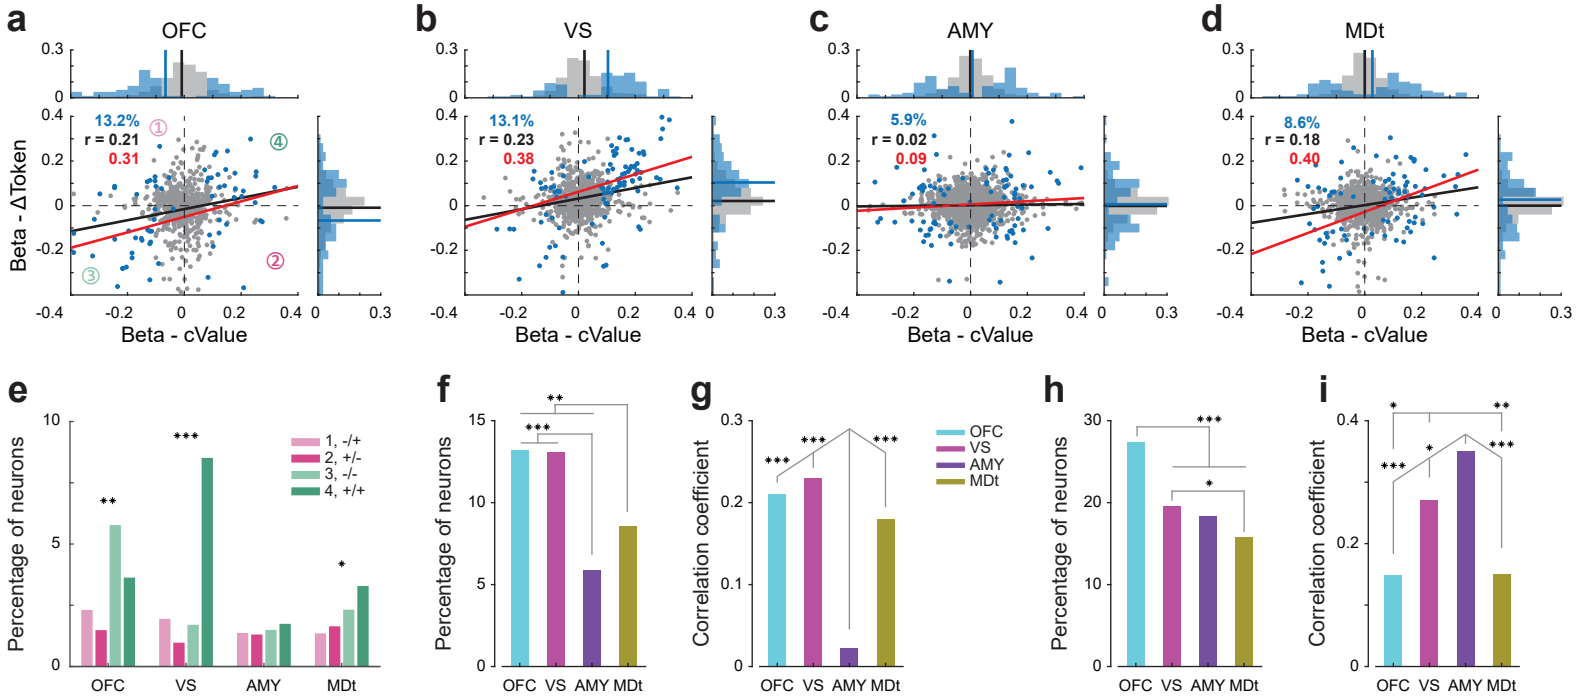
a
OFC
13.2%
r = 0.21
0.31
Beta - ΔToken
Beta - cValue
b
VS
13.1%
r = 0.23
0.38
c
AMY
5.9%
r = 0.02
0.09
d
MDt
8.6%
r = 0.18
0.40
e
Percentage of neurons
1, -/+
2, +/-
3, -/-
4, +/+
OFC
VS
AMY
MDt
f
Percentage of neurons
g
Correlation coefficient
h
Percentage of neurons
i
Correlation coefficient

Figure 4

Fig. 4 | Population encoding of value update and reinforcers.

a-g, Co-encoding of cValue and Δtoken. **a-d**, The x- and y-axes represent the regression coefficients of cValue and Δtoken. The blue dots represent the neurons encoding both of them. The gray dots represent the other neurons. Percentages on the left corner indicate the proportions of neurons encoding the cValue and Δtoken. Black and red lines are the lines of best fit for all the neurons and significant neurons, and r indicates the corresponding correlation coefficient. The bars summarize the distributions of blue and gray dots, with blue and black lines representing the means. **e**, Summary of the distributions of blue dots in each quadrant (also indicated in **a**) for each area. Chi-square test, *p < 0.05, **p < 0.01, ***p < 0.001. **f**, The proportion of neurons encoding cValue and Δtoken in each area. Summary of the percentages in **a-d**. Chi-square test, **p < 0.01, ***p < 0.001. **g**, Encoding similarity between cValue and Δtoken in each area. Summary of the black r in **a-d**. Fisher's z-transformation, ***p < 0.001. Non-significant pairs are not indicated.

h-i, Co-encoding of #token and #juice. **h**, The proportion of neurons encoding #token and #juice in each area. Chi-square test, *p < 0.05, ***p < 0.001. **i**, Encoding similarity between #token and #juice in each area. Fisher's z-transformation, *p < 0.05, **p < 0.01, ***p < 0.001.

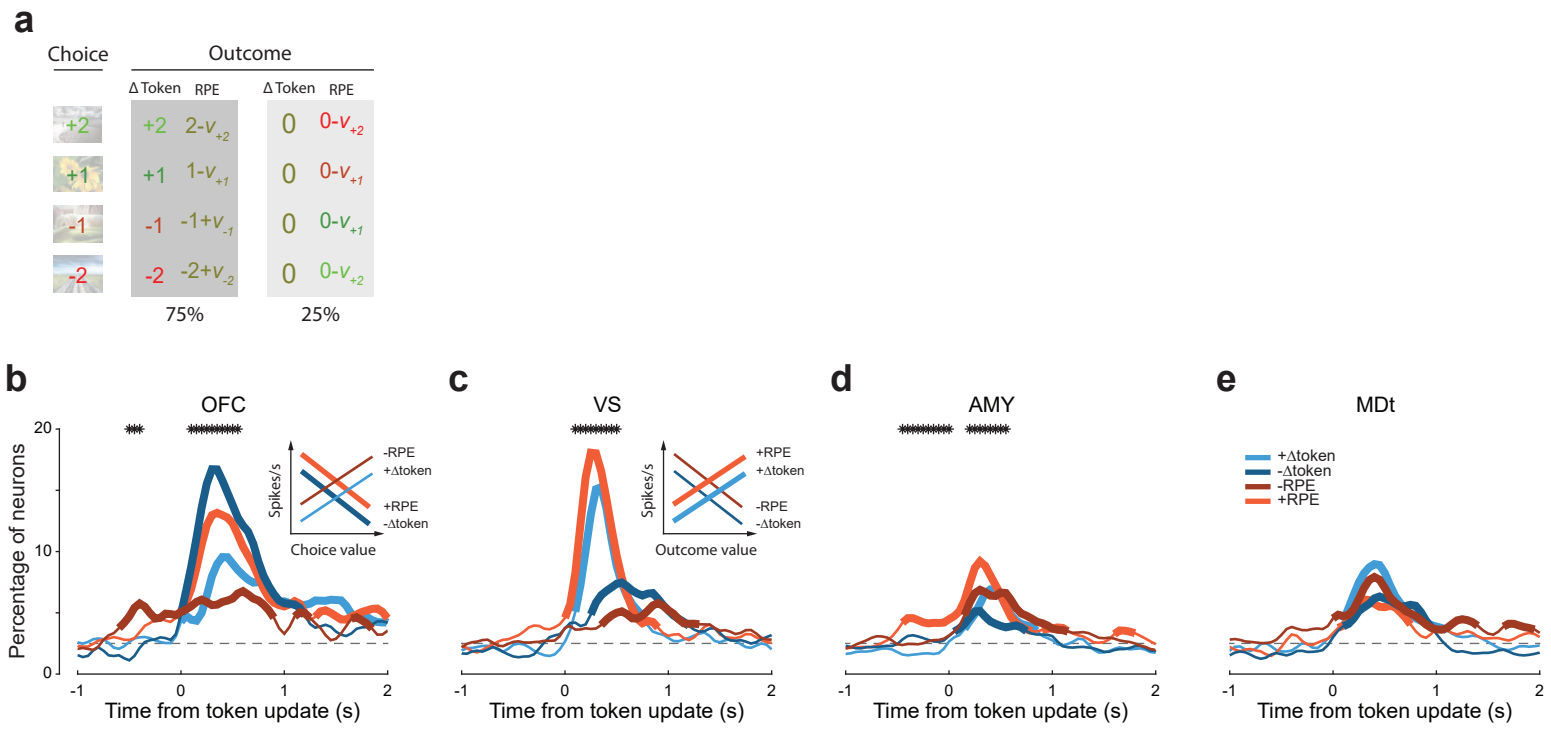
a
Choice
Outcome
Δ Token
RPE
Δ Token
RPE
+2
+2
2-v+2
0
0-v+2
+1
+1
1-v+1
0
0-v+1
-1
-1
-1+v-1
0
0-v+1
-2
-2
-2+v-2
0
0-v+2
75%
25%
b
OFC
Percentage of neurons
Time from token update (s)
Spikes/s
Choice value
-RPE
+Δtoken
+RPE
-Δtoken
c
VS
Time from token update (s)
Spikes/s
Outcome value
+RPE
+Δtoken
-RPE
-Δtoken
d
AMY
Time from token update (s)
e
MDt
Time from token update (s)
+Δtoken
-Δtoken
-RPE
+RPE

Figure 5

Fig. 5 | Value updates in the ventral striatum and orbitofrontal cortex.

a, Relationship of *a priori* value, Δtoken, and RPE. v_h indicates the RW model-estimated value of the image with *a priori* value h, which got closer to h as learning progressed.

b-e, Percentage of neurons in each area encoding +Δtoken, -Δtoken, +RPE, and -RPE. Dashed horizontal lines represent chance levels. Thick lines indicate a significant difference between the corresponding area and chance level (binomial test, $p < 0.025$). The black asterisks indicate a significant difference among the four categories (chi-square test, $p < 0.01$). Insets: OFC neurons encoded Δtoken and RPE in a manner relevant to the value of choices. VS neurons encoded Δtoken and RPE correlate with the value of outcomes. Thick lines represent more neurons.

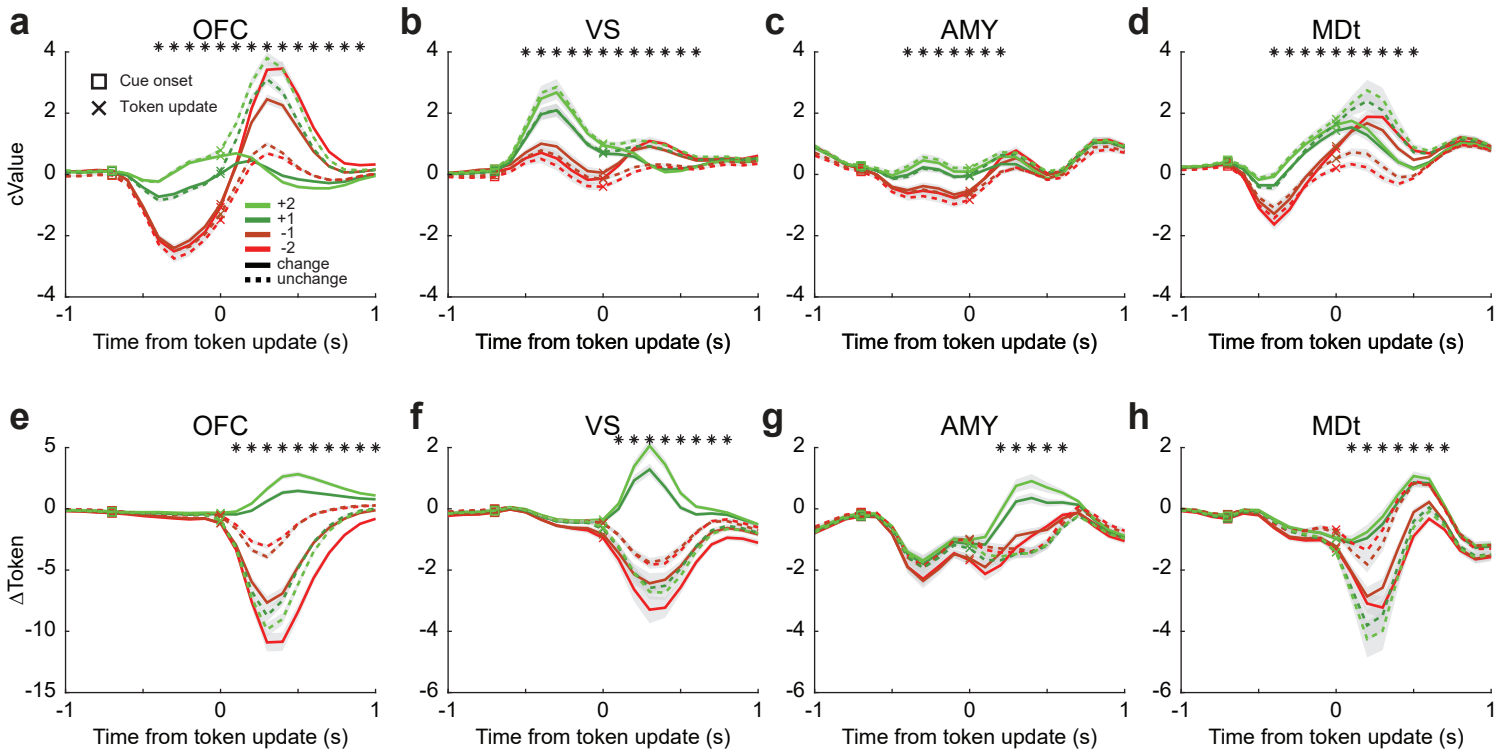
a
OFC
Cue onset
Token update
cValue
+2
+1
-1
-2
change
unchange
Time from token update (s)
b
VS
c
AMY
d
MDt
e
OFC
ΔToken
f
VS
g
AMY
h
MDt

Figure 6

Fig. 6 | Populational representation of gains and losses

a-d, Projections of neural population activity on the cValue axis in the task-variable-specific one-dimensional subspace. The y-axis indicates the value of the projection, divided by the number of neurons (always 500). All neural activity was also z-scored. Therefore, the y-axis represents the average z-scored population deviation per cell. Trials were grouped by cValue and outcome. Solid lines indicate the conditions when tokens were delivered, and dashed lines indicate the conditions in which tokens were not delivered. The analysis was repeated 1000 times by randomly choosing 500 neurons from each population without replacement. Shaded zones show the mean ± SD across iterations. The squares and crosses represent the cue onset and token update. The black asterisks indicate a significant difference among the eight conditions (bootstrap test, $p < 0.01$).

e-h, Projections of neural population activity on the Δtoken axis in the task-variable-specific one-dimensional subspace.

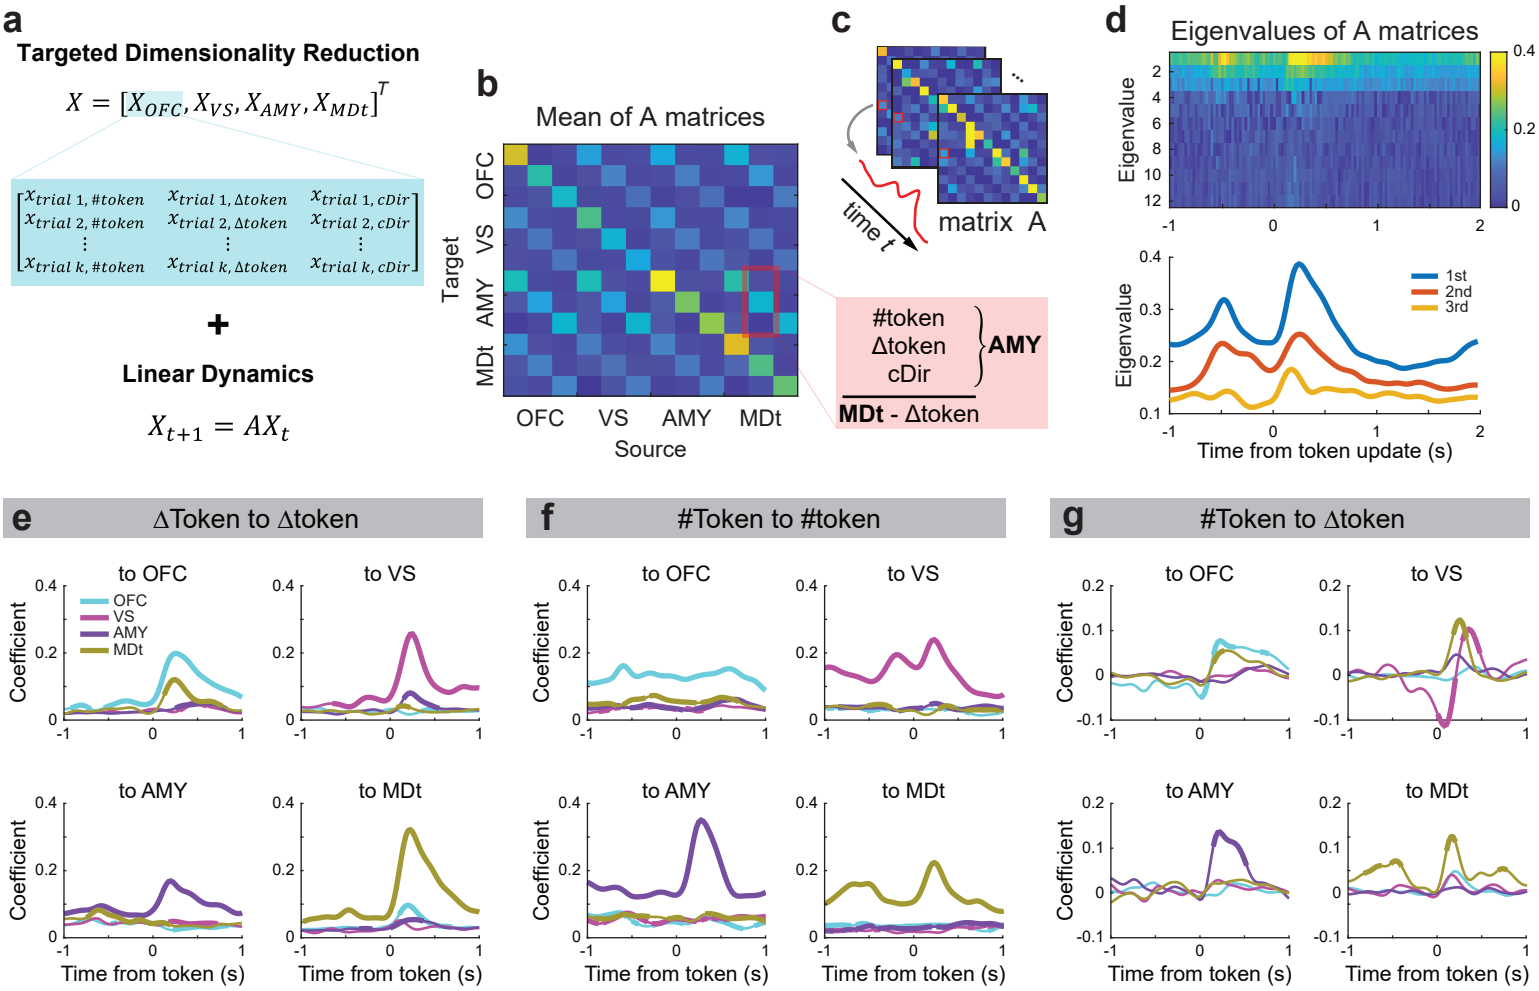
a
Targeted Dimensionality Reduction
$X = [X_{OFC}, X_{VS}, X_{AMY}, X_{MDt}]^T$
$\begin{bmatrix} x_{trial\ 1,\#token} & x_{trial\ 1,\Delta token} & x_{trial\ 1,cDir} \\ x_{trial\ 2,\#token} & x_{trial\ 2,\Delta token} & x_{trial\ 2,cDir} \\ \vdots & \vdots & \vdots \\ x_{trial\ k,\#token} & x_{trial\ k,\Delta token} & x_{trial\ k,cDir} \end{bmatrix}$
+
Linear Dynamics
$X_{t+1} = AX_t$
b
Mean of A matrices
Target
OFC
VS
AMY
MDt
Source
#token
Δtoken
cDir
AMY
MDt - Δtoken
c
time t
matrix A
d
Eigenvalues of A matrices
Eigenvalue
1st
2nd
3rd
Time from token update (s)
e
ΔToken to Δtoken
f
#Token to #token
g
#Token to Δtoken
to OFC
to VS
to AMY
to MDt
Coefficient
Time from token (s)

Figure 7

Fig. 7 | Flow of token information in the ventral network.

a-d, Detecting network communications with low dimensional information dynamics. **a**, The task-related population response matrix consisted of single-trial-level population projections. Linear dynamic regression measured time-dependent flows of information across task variables and brain areas. **b**, The loading matrix A maps the relationship between rows in matrix X. The x-axis represents the source areas and variables, and the y-axis represents the target areas and variables. For example, the values indicated by the red square represent the flow of Δtoken information from the MDt to #token, Δtoken, and cDir information in the AMY. **c**, The continuous values in A matrices represent the strength of information flow across task variables and brain areas along the time course of the trial. **d**, The eigenvalues of A matrices capture the strength of the overall information flow among the areas in the network.

e-g, The flow of token information across areas, including within the Δtoken (**e**), #token (**f**), and between them (**g**). The titles indicate the source and target task variables. Each panel indicates one target area, and lines indicate source areas. Thick lines indicate a significant difference between the corresponding area and shuffled data (two sides z-test, $p < 0.01$). Shuffling was done by keeping the same conditions but shuffling the trial order of matrix X.

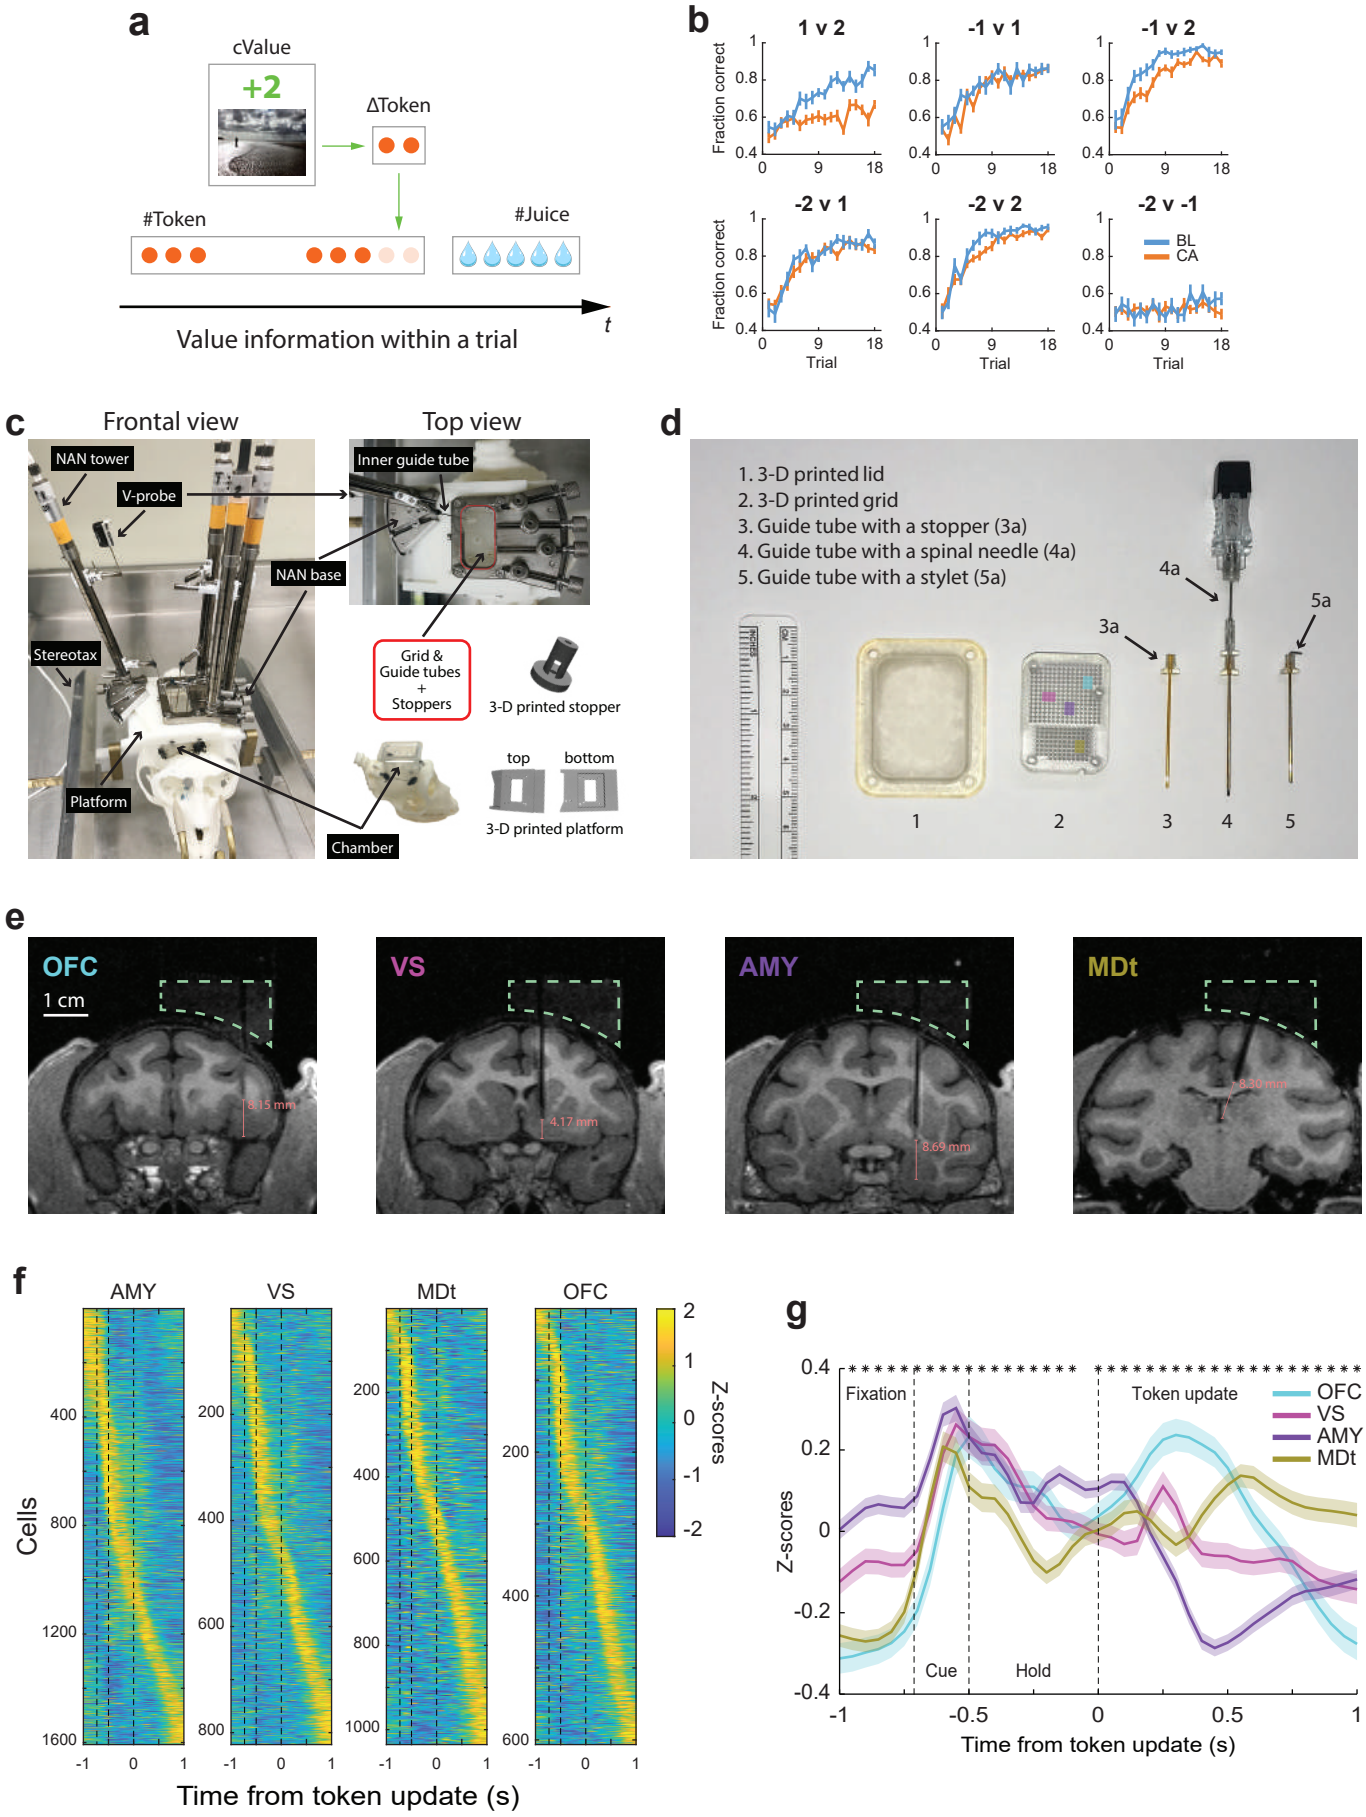

Extended Data Figure 1

Extended figure legends

Extended Data Fig. 1 | Behavior, recorded areas, and neural activity.

a, Different forms of value in the task. The primary reward was 1-12 drops of apple juice (#juice). The accumulated tokens were cashed out for juice every four to six trials. So, the reward signal was stored in the number of tokens (#token) for most trials. In each trial, the value update information was first carried by the images (cValue) and then transferred to the change of tokens (Δtoken).

b, The choice behaviors of monkeys BL and CA.

c, Recoding setup on a 3D-printed skull in a stereotaxic frame. Only one probe is shown.

d, Components of the recording chamber. 1) 3-D printed lid. 2) 3-D printed grid, with vertical holes for the OFC, VS, and AMY (indicated by the colored squares), and holes with a 16° angel for the MDt; 3) Guide tube with a 3-D printed stopper glued at the top. 4) Guide tube with a spinal needle inside. This setup was used when inserting the guide tubes into the brain. The spinal needle penetrated the dura. 5) Guide tube with a stylet inside. This setup was used after the surgery. The stylets were about 0.5 mm longer than the guide tube protrusion to avoid clogging. They were removed during recording to allow probes to enter the guide tubes. All 3-D printed components were printed using biocompatible materials. All processes regarding guide tubes were carried out using sterile procedures.

e, A coronal section view of the chamber and the shadow of a guide tube for each area under MRI. Orange lines indicate the longest distance the probes could advance from the tip of the guide tubes.

f-g, Z-scored neural activity. **f**, Heatmaps of the activity of all the neurons recorded from the AMY (n = 1607), VS (n = 829), MDt (n = 1035), and OFC (n = 606). Firing rates were z-scored based on activity during the presented period. Warmer colors indicate higher firing rates. Neurons were aligned to the token update and rank-ordered by the latency of the response peak. One row represents one neuron. **g**, Mean z-scored activities of each area. Shaded zones represent mean ± SEM. The black asterisks indicate a significant difference among the four areas (one-way ANOVA, $p < 0.01$). Vertical black dashed lines represent the onset of images, the start of holding on the chosen image, and token update, which split the time course into four epochs: fixation, cue, hold, and token update.

Single unit analysis

a

$$r_i = \beta_{\#token,i} * \#token + \beta_{\Delta token,i} * \Delta token + \cdots + \beta_{cValue,i} * cValue + \varepsilon$$

Neural population analysis

b

$$r_1 = \beta_{\#token,1}$$
$$r_2 = \beta_{\#token,2}$$
$$r_3 = \beta_{\#token,3}$$
$$r_4 = \beta_{\#token,4} * \#token + \beta_{\#juice,4} * \#juice + \cdots + \beta_{cDir,4} * cDir$$
$$r_5 = \beta_{\#token,5}$$
$$\vdots$$
$$r_i = \beta_{\#token,i}$$

$\beta_{\#juice,1}$, $\beta_{\#juice,2}$, $\beta_{\#juice,3}$, $\beta_{\#juice,4}$, $\beta_{\#juice,5}$, $\vdots$, $\beta_{\#juice,i}$

$\beta_{cDir,1}$, $\beta_{cDir,2}$, $\beta_{cDir,3}$, $\beta_{cDir,4}$, $\beta_{cDir,5}$, $\vdots$, $\beta_{cDir,i}$

c

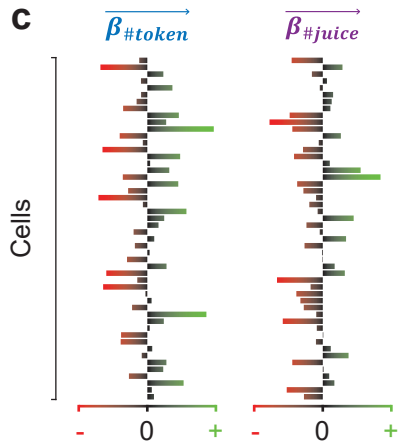

Network dynamics

d

Targeted Dimensionality Reduction

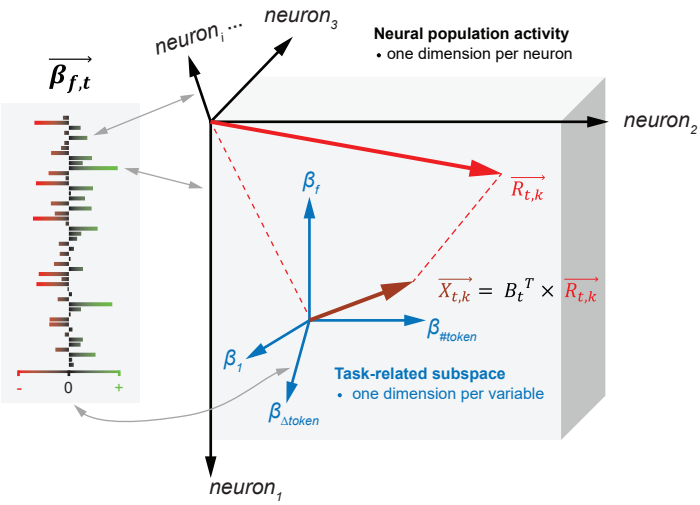

e

Linear Dynamics: $X_{t+1} = AX_t$

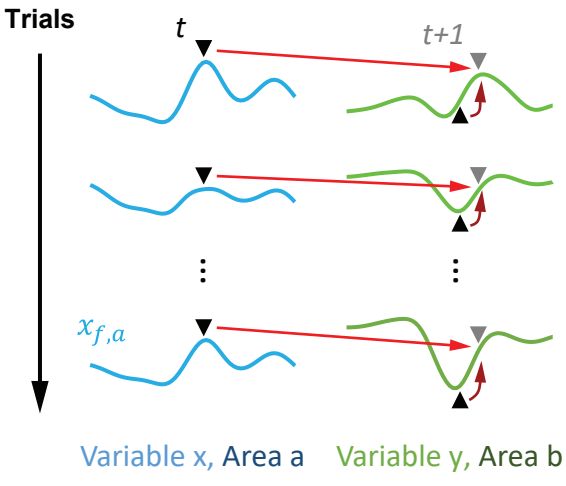

Extended Data Figure 2

Extended Data Fig. 2 | Systematic procedures for analyzing neural population data at single-cell, population, and network levels.

a, Classifying neurons by replacing independent variables (in this case, Δtoken and cValue) with several potential combinations in multivariate linear regression models. A neuron was classified into the categories as the combination included in the best-fitting model. Nan means removing the variable from the model.

b-c, Approach to calculate population coding similarity. **b**, The regression coefficient vector, $\overrightarrow{\beta_f}$, of a neural population represents the weights of each neuron in the population encoding one task variable f. It indicates the relationship between neural population activity and the task variable. **c**, The cross-correlation between two regression coefficient vectors (simulation of 50 neurons) indicates how similar a neural population encoded the two task variables.

d, Targeted dimensionality reduction. Task-variable-specific latent variable at time t on trial k, $\overrightarrow{X_{t,k}}$ (one vector in the matrix X_t), was calculated by projecting single-trial population activity $\overrightarrow{R_{t,k}}$ (one vector in the matrix R_t) into the subspace generated by axes B_t, which consisted of $\overrightarrow{\beta_{f,t}}$. Each coefficient vector $\overrightarrow{\beta_{f,t}}$ maps the relationship between all the axes in the high-dimensional space (one β for one axis) and the task variable f axis at time t in the low-dimensional subspace. We also carried out the same approach using condition-averaged population response, but not at the single trial level, analyzing the populational representation of gains and losses. See also Methods.

e, Linear dynamics. Matrix A measures the within-area/variable interactions (dark red arrows) and across-area/variable interactions (light red arrows) at each point in time. The one-dimensional single-trial task-variable-specific latent variable $x_{f,a}$ consisted of the projections of $\overrightarrow{X_{t,k}}$ from area a on the axis β_f across the time course of the trial.

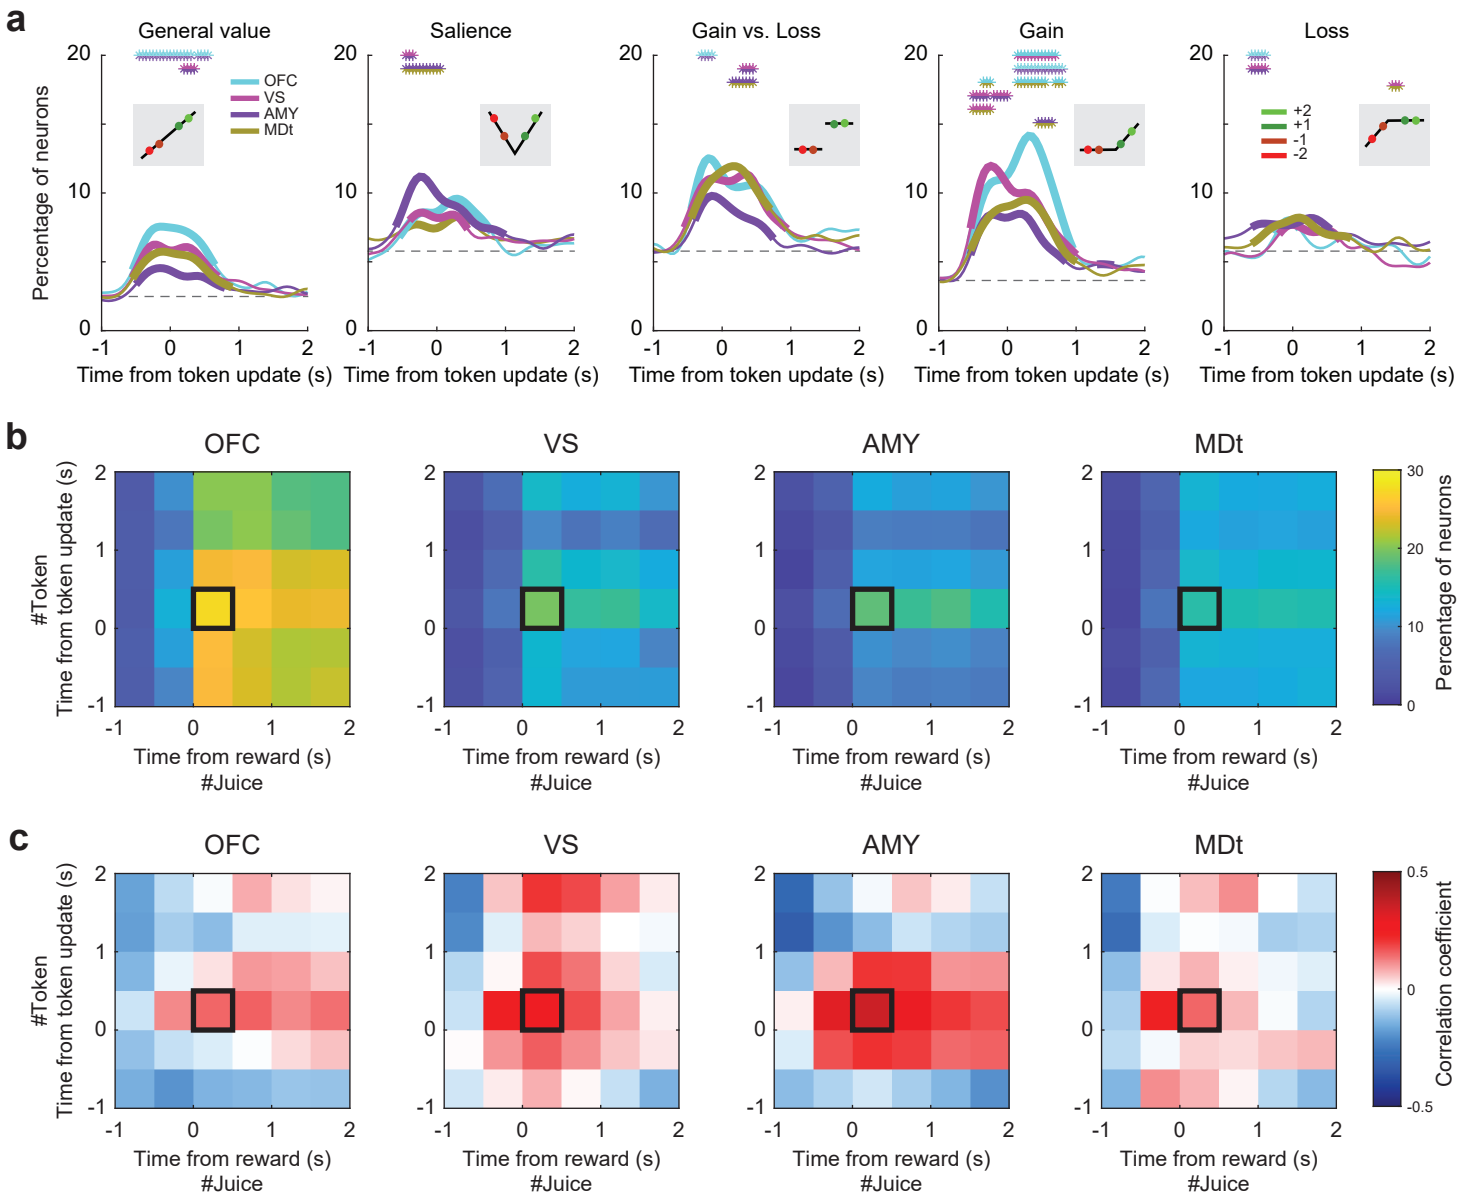

Extended Data Figure 3

Extended Data Fig. 3 | Coding of chosen value, tokens, and the number of juice rewards.

a, Percentage of neurons in each area encoding the *a priori* value of chosen images (cValue).

b, Percentage of neurons encoding #token and #juice. The x-axis represents #juice, and the y-axis represents #token. Black squares indicate the time windows used for the bar plots in Fig. 4h-i.

c, Encoding similarity between #token and #juice in each area.

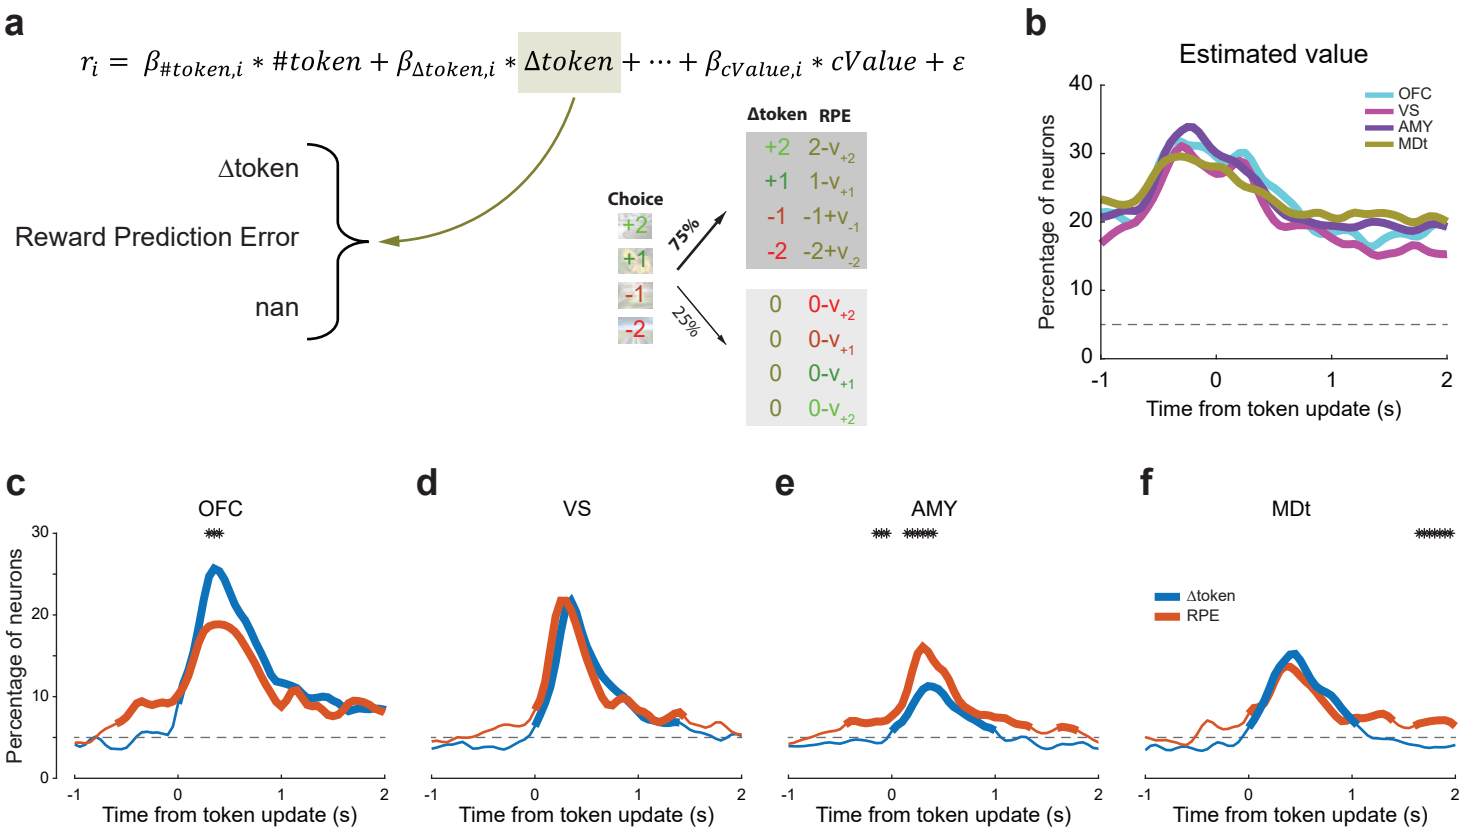

Extended Data Figure 4

Extended Data Fig. 4 | Value updates in the ventral striatum and orbitofrontal cortex.

a, Replacing the independent variable Δtoken with potential combinations (Δtoken, RPE, or nan). A neuron was classified into the category that the best-fitting model indicated.

b, Percentage of neurons in each area encoding the value of chosen images, v_h, estimated by the RW model. Dashed horizontal lines represent the chance level. Thick lines indicate a significant difference between the corresponding area and chance level (binomial test, $p < 0.01$).

c-f, Percentage of neurons in each area encoding Δtoken and RPE. Dashed horizontal lines represent chance levels. Thick lines indicate a significant difference between the corresponding area and chance level (binomial test, $p < 0.05$). The black asterisks indicate a significant difference between the two categories (chi-square test, $p < 0.01$).

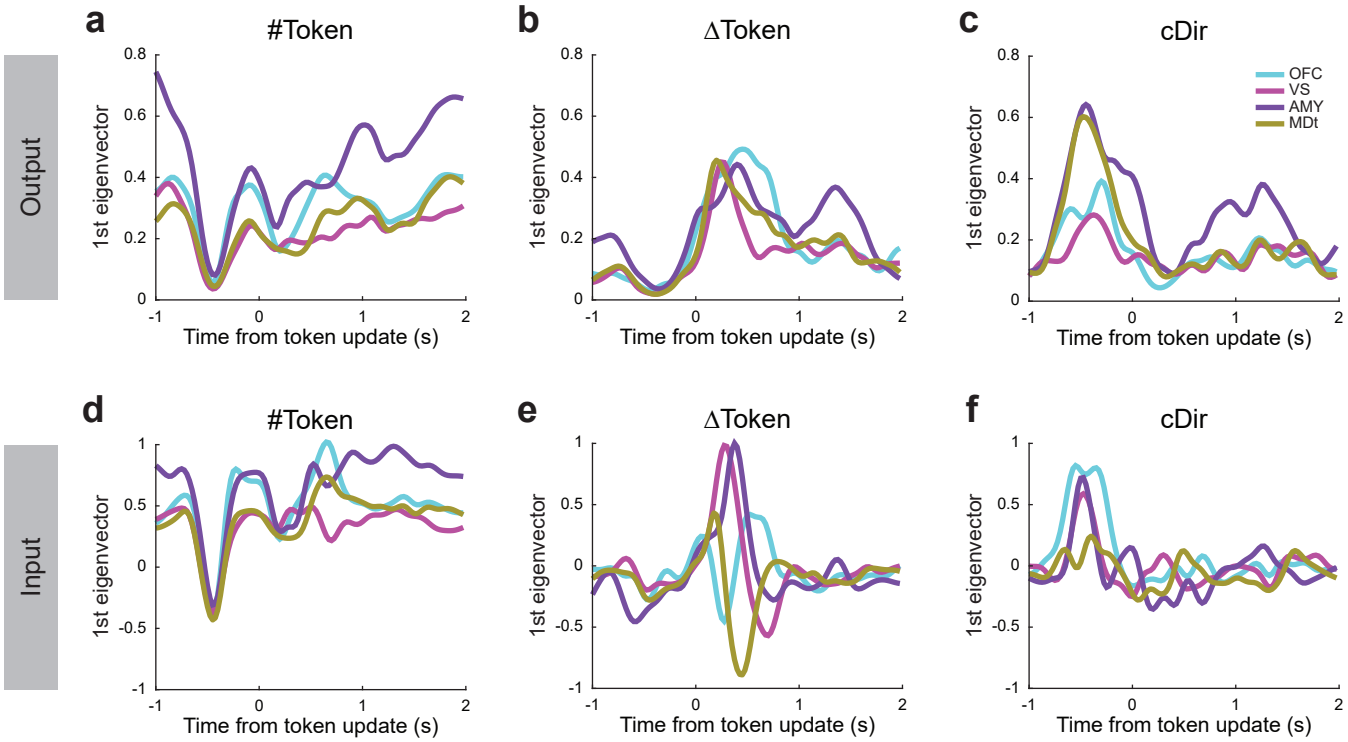

Extended Data Figure 5

Extended Data Fig. 5 | Eigenvectors of the matrix A.

a-c, The columns of eigenvectors capture the information flowing from the source areas and task variables.

d-f, The rows of inversed eigenvectors capture the information flowing into the source areas and task variables.

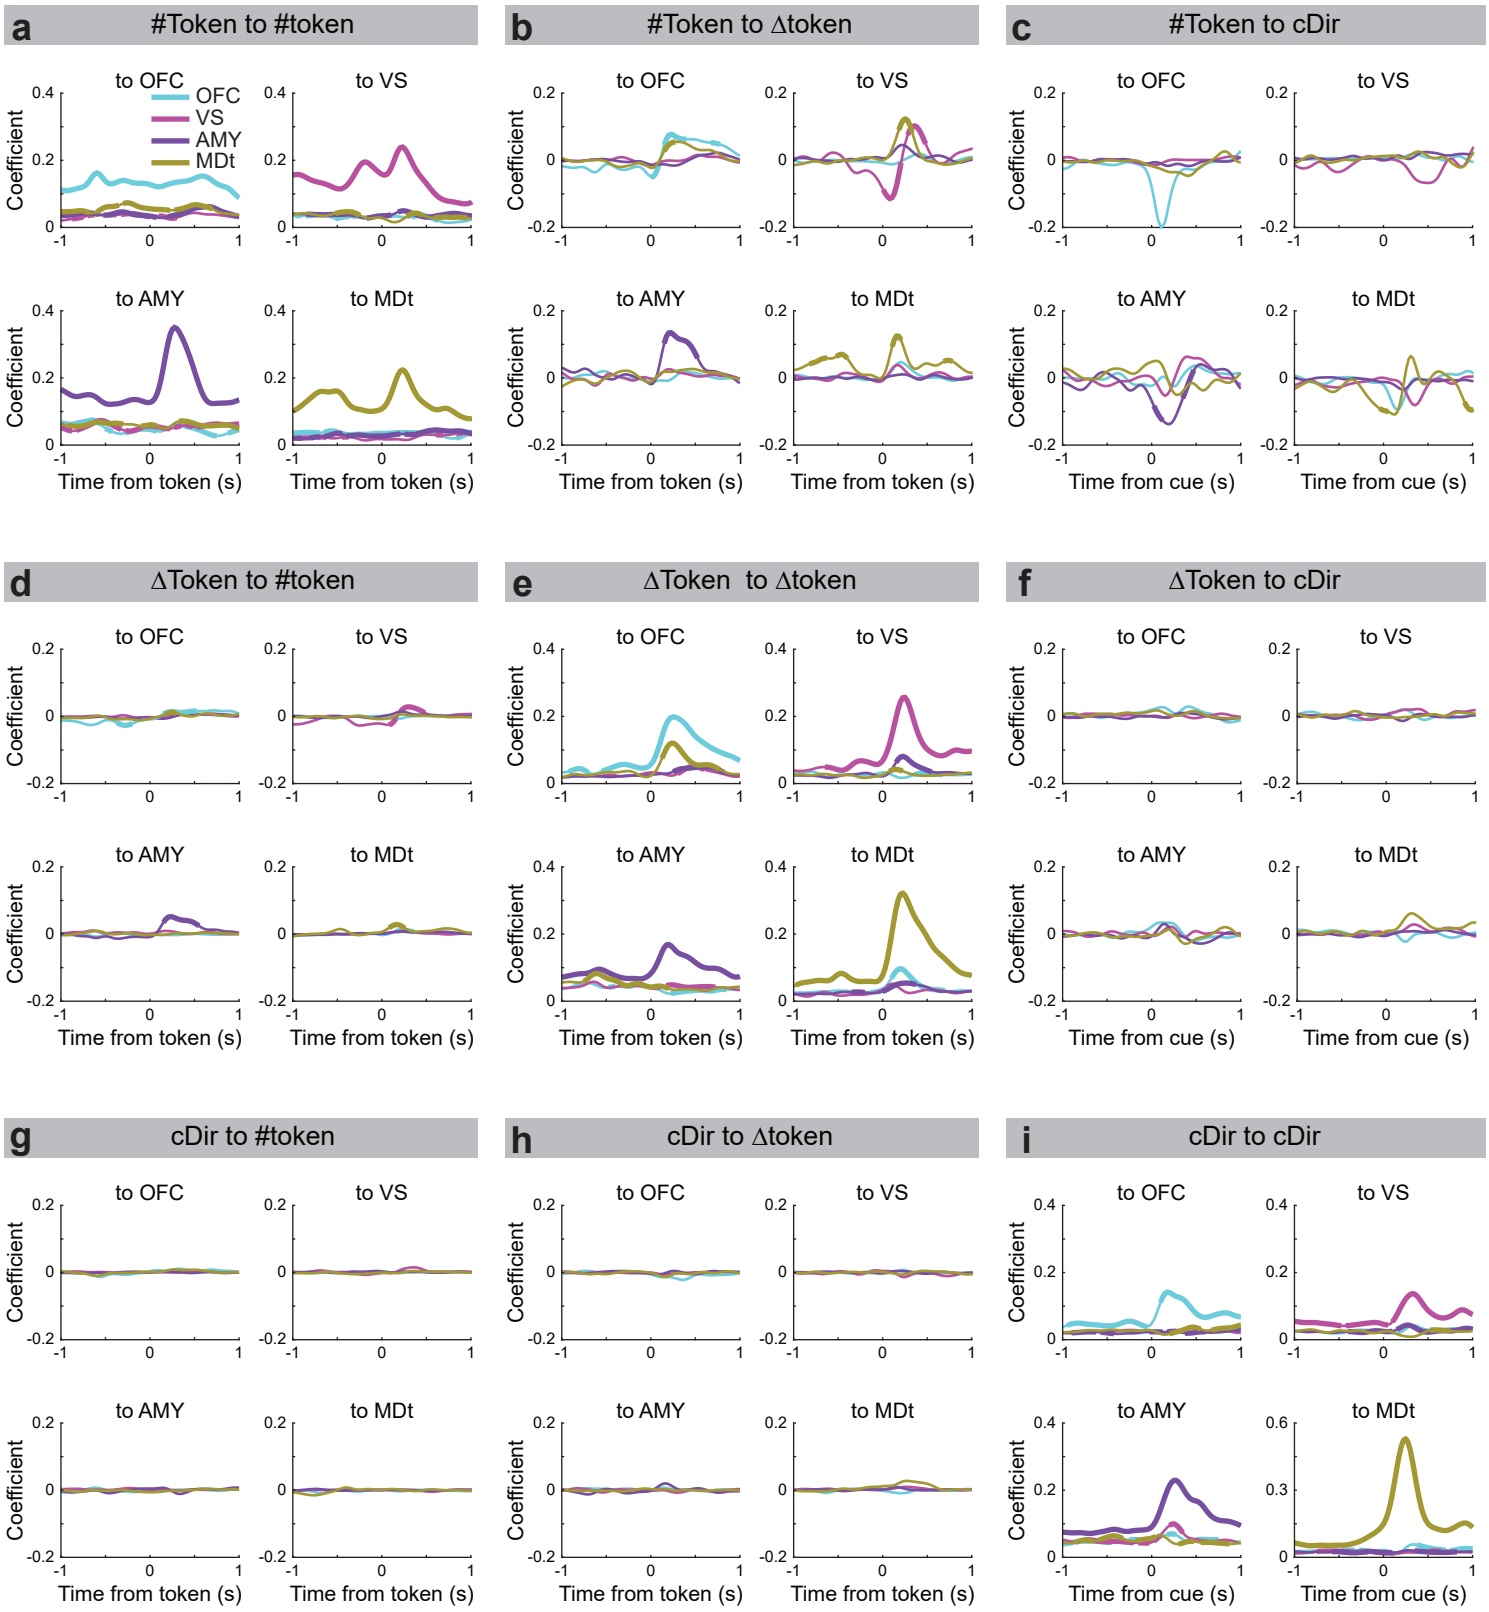

Extended Data Figure 6

Extended Data Fig. 6 | Flow of token and direction information in the ventral network.

a-i, The flow of token and chosen direction information across areas and task variables. All combinations among the three task variables are shown.