Neuron
Article

Distinct Tonic and Phasic Anticipatory Activity in Lateral Habenula and Dopamine Neurons

Ethan S. Bromberg-Martin,[1,*] Masayuki Matsumoto,[1,2] and Okihide Hikosaka[1]
[1]Laboratory of Sensorimotor Research, National Eye Institute, National Institutes of Health, Bethesda, MD 20892, USA
[2]Primate Research Institute, Kyoto University, Inuyama, Aichi 484-8506, Japan
*Correspondence: bromberge@mail.nih.gov
DOI 10.1016/j.neuron.2010.06.016

SUMMARY

Dopamine has a crucial role in anticipation of motivational events. To investigate the underlying mechanisms of this process, we analyzed the activity of dopamine neurons and one of their major sources of input, neurons in the lateral habenula, while animals anticipated upcoming behavioral tasks. We found that lateral habenula and dopamine neurons anticipated tasks in two distinct manners. First, neurons encoded the timing distribution of upcoming tasks through gradual changes in their tonic activity. This tonic signal encoded rewarding tasks in preference to punishing tasks and was correlated with classic phasic coding of motivational value. Second, neurons transmitted a phasic signal marking the time when a task began. This phasic signal encoded rewarding and punishing tasks in similar manners, as though reflecting motivational salience. Our data suggest that the habenula-dopamine pathway motivates anticipation through a combination of tonic reward-related and phasic salience-related signals.

INTRODUCTION

Our ability to anticipate motivational events is a common thread that runs throughout our everyday lives, known well to the child who eagerly counts the minutes before she can unwrap her Christmas presents or who fidgets anxiously in the waiting room while awaiting her appointment with the dentist. This ability is thought to be critically dependent on dopamine release within the basal ganglia. Interval timing and self-paced behavior are distorted by dopaminergic agonists and antagonists and are impaired by dopaminergic lesions and degeneration in Parkinson's disease (Buhusi and Meck, 2006). Yet, current knowledge about dopamine neurons predominantly focuses on their phasic reactions triggered by external sensory stimuli, notably their responses to rewards and punishments which are thought to cause motivational learning (Schultz et al., 1997; Joshua et al., 2008; Matsumoto and Hikosaka, 2009b). Few studies have investigated whether and how dopamine neurons change their activity anticipatorily in the seconds leading up to an event expected in the future (Romo and Schultz, 1990; Fiorillo et al., 2003, 2008; Schultz, 2007). In addition, few studies have examined how dopamine neurons represent stimuli that act as purely temporal cues, allowing the timing of future events to be anticipated without providing new information about their motivational value (Satoh et al., 2003).

Here, we report that dopamine neurons convey two distinct signals that tonically anticipate the time of future behavioral tasks and phasically mark their time of occurrence. We were able to discover several of the principles underlying these signals by using an experimental design to dissociate neural activity related to rewards and punishments. Tonic anticipatory activity preferentially encoded rewards rather than punishments and was correlated with classic phasic encoding of motivational value. These tonic signals were therefore broadly consistent with current theories that dopamine neurons report "reward prediction errors" (Schultz et al., 1997), with the modification that dopamine neurons do so in a tonic as well as a phasic manner. In contrast, phasic responses to the start of a behavioral task followed a distinct coding principle. These signals encoded rewarding and punishing tasks in very similar manners, as if representing the task's motivational importance or "salience" (Horvitz, 2000; Redgrave and Gurney, 2006; Lin and Nicolelis, 2008; Joshua et al., 2009). Thus, the same neurons transmitted tonic activity and classic phasic signals related to motivational value, as well as an additional phasic signal at the start of the task more closely resembling motivational salience.

In addition, we report a candidate neural pathway to provide dopamine neurons with these tonic and phasic signals. We recorded neural activity in the lateral habenula, a nucleus located in the epithalamus that exerts control over multiple neuromodulatory systems including dopamine, serotonin, and norepinephrine (Lecourtier and Kelly, 2007). We considered the lateral habenula to be a good candidate for causing dopamine anticipatory signals because habenula lesions disrupt the timing of reward-oriented behavior (Macdougall et al., 1969; Thornton and Evans, 1984; Thornton et al., 1990), lateral habenula activation causes dopamine neurons to be powerfully inhibited (Christoph et al., 1986; Ji and Shepard, 2007), and lateral habenula neurons transmit reward and punishment signals that resemble a sign-reversed version of classic dopamine phasic responses (Matsumoto and Hikosaka, 2007, 2009a). Here, we extend these results by showing that lateral habenula neurons also contain sign-reversed versions of the tonic and phasic task anticipation signals found in dopamine neurons, in a

manner consistent with a habenula → dopamine direction of transmission. These data implicate the habenula-dopamine pathway in anticipation of rewarding and salient events.

RESULTS

Tonic Encoding of Variable Intervals before Rewarding Tasks

We analyzed a database of neurons recorded while monkeys performed three behavioral tasks—a reward-biased saccade task (lateral habenula n = 65, dopamine n = 64; Matsumoto and Hikosaka, 2007), an information choice task (dopamine n = 47; Bromberg-Martin and Hikosaka, 2009), and a Pavlovian conditioning procedure (lateral habenula n = 74; dopamine n = 103; Matsumoto and Hikosaka, 2009a, 2009b).

We first analyzed neural activity recorded during the reward-biased saccade task while animals waited for the next trial to begin (Figure 1). The stimulus display was very simple: during the intertrial interval, the animal viewed a blank screen, after which a white spot of light appeared that signaled the start of the next trial (Figure 1A). After the trial start cue appeared, the animal performed a visually guided saccade task to gain probabilistic juice rewards (Experimental Procedures). The trial start cue was primarily a temporal cue, indicating the timing of the upcoming trial without providing new information about its expected reward value. Nonetheless, because the trial start cue marked the start of a new opportunity to gain rewards, animals are thought to assign it motivational value which triggers an excitatory dopamine response (Satoh et al., 2003; Takikawa et al., 2004). Consistent with these studies, the trial start cue triggered strong phasic inhibition in lateral habenula neurons and strong phasic excitation in dopamine neurons (Figures 1B and 1C, right), similar to the responses of these neurons to cues that explicitly indicate future rewards (Matsumoto and Hikosaka, 2007). Neural response latencies were consistent with a habenula → dopamine direction of transmission, as lateral habenula neurons were inhibited at a latency of 119 ± 3 ms after which dopamine neurons were excited at a longer latency of 127 ± 2 ms ($p < 0.05$, bootstrap test; ± indicates SE; see Figure S1 available online).

A close examination of neural activity revealed a second task-anticipatory signal in tonic spike activity during the intertrial interval (ITI). In these experiments the ITI was randomized from 2.2–3.2 s (Figure 1A, lateral habenula n = 43, dopamine n = 42). During the first 2.2 s of the ITI neural activity was sustained at a tonic, baseline level. Then, at the first moment when the trial start cue could potentially appear, neurons began to change their level of tonic activity in a linear, ramp-like manner opposite to the direction of their phasic responses (Figures 1B and 1C, left). This tonic change in activity continued until the ITI ended with the onset of the trial start cue. Thus, lateral habenula neurons produced a gradual tonic excitation during the ITI followed by a later phasic inhibition; dopamine neurons produced gradual tonic inhibition during the ITI followed by a later phasic excitation. This pattern of tonic activity resembles a form of negative "reward prediction error," occurring at each moment when the trial start cue failed to appear and growing in magnitude as time elapsed during the ITI and the trial start

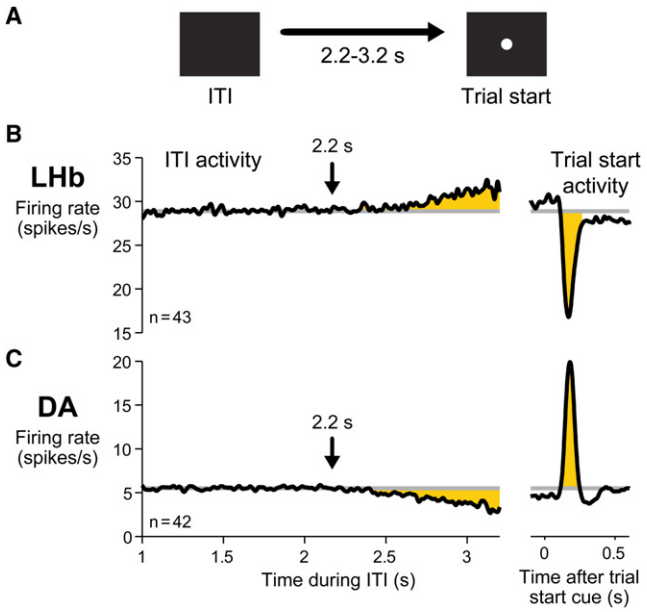

Figure 1. Lateral Habenula and Dopamine Neurons Tonically Encode Variable Intervals before Rewarding Tasks

(A) Events during the intertrial interval (ITI) of the reward-biased saccade task. During the ITI the screen was blank. After a randomized 2.2–3.2 s delay, a cue appeared marking the start of the next trial.

(B and C) Average firing rate of lateral habenula neurons (B) and dopamine neurons (C) aligned on the start of the ITI (left) and the onset of the trial start cue (right). The light gray line indicates baseline firing rate. The yellow shaded area indicates the deviation from baseline firing rate.

See also Figure S1.

cue became increasingly expected (Fiorillo et al., 2008). At first glance these tonic changes in firing rate appeared quite modest in size, reaching a peak excitation or inhibition of ~2.5 spikes/s. However, the prolonged nature of this activity meant that it caused a change in spike count comparable to classical phasic responses (yellow shaded area in Figures 1B and 1C).

In some neurons the tonic effects were strong enough to be seen on single trials (Figure 2). The habenula neuron in Figures 2A and 2B stayed close to a baseline of 50 spikes/s and then at the 2.2 s mark began a linear increase of activity up to 75 spikes/s by the end of the ITI. In a plot of its single trial spike activity, a white band appears indicating its phasic inhibition in response to the trial start cue. A dark area can also be seen in the last few hundred milliseconds of the longest ITIs, indicating an increased tonic spike rate in anticipation of the next trial (bottom raster plot, Figure 2A). Similarly, the dopamine neuron in Figures 2D and 2E stayed close to a baseline of 8 spikes/s and then decreased its rate during the variable portion of the ITI to a low of ~1 spike/s. It emitted a burst of spikes in response to the trial start cue, and during the longest ITIs this burst was preceded by several hundred milliseconds during which its tonic spike rate was visibly decreased (bottom raster plot, Figure 2D).

To test the prevalence of tonic activity in single neurons, we fit each neuron's activity during the variable portion of the ITI using a linear ramp-like function with two parameters, a starting firing rate and an ending firing rate (Figures 2B and 2E; Experimental

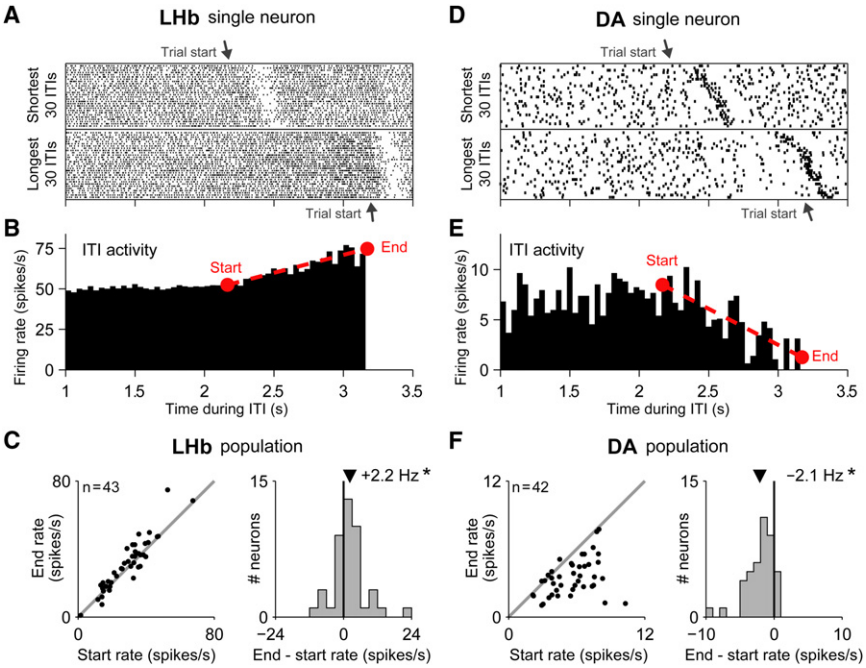

Figure 2. Prevalence of Tonic Activity in Single Neurons

(A) Spike activity of a lateral habenula neuron with strong tonic activity during the ITI. Each row is a trial and each dot is a spike. Trials are sorted by the time that the trial start cue appeared. Top: shortest 30 ITIs; bottom: longest 30 ITIs. Arrows mark the earliest and latest trial start times (gray arrow).

(B) Average firing rate of the lateral habenula neuron during the ITI. Black histogram: ITI firing rate in 40 ms bins. Red dots: fitted start and end firing rates.

(C) Left: fitted start firing rate (x axis) and end firing rate (y axis) for each lateral habenula neuron. Right: histogram of changes in firing rate, (end rate – start rate). Text indicates the mean change in firing rate across the population. Asterisk indicates statistical significance ($p < 0.05$, t test).

(D–F) Same as (A)–(C), for dopamine neurons. See also Figure S2.

Procedures). The change in tonic activity was then measured as the difference, in rates (end rate—start rate; Figures 2C and 2F). The majority of lateral habenula neurons had positive changes in activity indicating tonic excitation, and nearly all dopamine neurons had negative changes in activity indicating tonic inhibition (habenula $p = 0.01$, dopamine $p < 10^{-6}$, Wilcoxon signed-rank test of median tonic effect; Figures 2C and 2F). We fit a similar ramp-like function to phasic neural responses to the trial start cue to test whether they were modulated by elapsed time in the same manner (Fiorillo et al., 2008; Figure S2). Neurons had a trend for stronger phasic trial start responses after long ITIs, but this trend did not reach significance at the population level (habenula $p = 0.07$, dopamine $p = 0.14$).

We next asked whether the ramping changes in trial-averaged firing rates were caused by consistent ramping changes in tonic activity, or by averaging of brief, phasic changes in activity occurring at different times on each trial (Niv et al., 2005). To test between these hypotheses, we fit each neuron's spiking activity using two probabilistic models, a *tonic model* and a *phasic model* (Figure S3). The two models made identical predictions about the neuron's trial-averaged firing rate, but made distinct predictions about the pattern of activity on single trials—predicting either a gradual, ramp-like change in activity (for the tonic model) or a stable baseline interrupted by occasional brief changes in firing rate (for the phasic model). The tonic model provided a better fit than the phasic model for the majority of neurons ($p \leq 0.001$, binomial test).

Tonic Signals Track the Temporal Distribution of Rewarded Trials

The data shown so far suggested that lateral habenula and dopamine neurons tonically encoded variable temporal intervals before upcoming behavioral tasks. To test the generality of this phenomenon, we analyzed neural activity recorded in several tasks that used a wide variety of ITI distributions (Figure 3; Experimental Procedures). The results were replicated in every case: neurons had a linear ramp-like change in activity that started around the first moment in time when the next trial could begin. The same pattern was found in both lateral habenula and dopamine neurons (Figures 3B–3D and 3F–3H) and in every animal tested in every task (Figure S4).

Note that this pattern is distinct from ramping anticipatory tonic activity found in other cortical and subcortical areas. In many areas, activity ramps up in anticipation of the time of predictable events, reaching its maximal level at the event's expected time of occurrence. Yet when lateral habenula and dopamine neurons were tested using a fully predictable constant ITI of 2.2 s, their tonic ramping activity was much less prominent (Figures 3A and 3E). Thus, tonic activity did not appear to build up to a peak at the time of the next trial but instead occurred only when the trial start cue was omitted or delayed past its potential time of occurrence (Fiorillo et al., 2008). Neural activity was generally faithful to the true distribution of ITIs, although the neural timing of events was not perfectly precise. In Figure 3G, the trial could first begin at the 4.1 s mark, but dopamine neurons were inhibited a few hundred milliseconds in advance (signed-rank test on activity 300 ms before the earliest trial start time, median −0.29 spikes/s, $p = 0.01$). Similarly, neurons tested with a constant ITI had a trend for slight changes in tonic activity just before the trial start cue appeared (Figures 3A and 3E; habenula median +0.42 spikes/s, $p = 0.27$; dopamine median −0.27 spikes/s, $p = 0.046$).

Tonic Signals Encode Rewards in Preference to Punishments

Thus far, we examined tonic activity as animals anticipated trials that led to rewards. We next asked whether this activity was specific to anticipation of rewards or whether it occurred similarly in anticipation of punishments. To answer this question

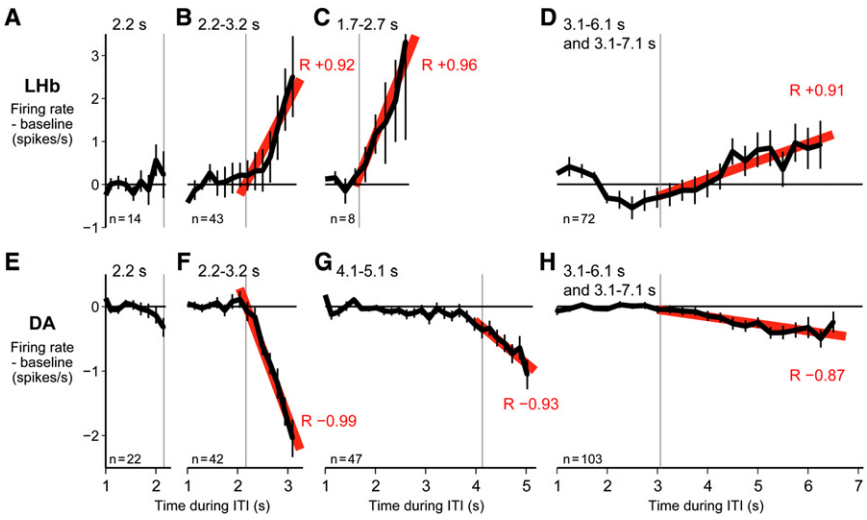

Figure 3. Lateral Habenula and Dopamine Neurons Encode the Temporal Distribution of Rewarding Tasks

(A–D) Average activity of lateral habenula neurons during tasks with a constant ITI (A) and variable ITIs (B–D). Text indicates the range of ITIs. Gray vertical lines mark the first possible time the trial could start. Black lines are mean baseline-subtracted firing rate in non-overlapping bins. The bin width for each plot was adjusted based on the range of trial start times during the ITI and the number of recorded neurons. The bin widths for (A)–(D) were 150 ms, 150 ms, 200 ms, and 250 ms. Error bars are ± 1 SEM. Data in (D) are combined from ITI distributions of 3.1–6.1 s and 3.1–7.1 s, and the last error bar in (C) is cropped above. Red lines indicate a linear least-squares fit to the plotted data points and red text indicates the linear correlation. All correlations were significant ($p < 0.005$, permutation test).

(E–H) Same format as (A)–(D) for dopamine neurons. The bin widths for (E)–(H) were 150 ms, 150 ms, 150 ms, and 250 ms.

See also Figure S3.

we analyzed 74 lateral habenula neurons and 103 dopamine neurons recorded during a Pavlovian conditioning procedure in which rewards and punishments were presented in separate blocks (Matsumoto and Hikosaka, 2009a, 2009b; Figure 4A). In the *appetitive block*, each trial yielded either a reward (fruit juice) or no reward. In the *aversive block* the task design was identical except rewards were replaced with punishments (aversive air puffs). Animals understood the task because they discriminated between the two blocks both neurally and behaviorally (Matsumoto and Hikosaka, 2009a).

During the appetitive block, neural ITI activity reflected clear tonic excitation in lateral habenula neurons and tonic inhibition in dopamine neurons (red lines, Figures 4B and 4C; same as Figures 3D and 3H). During the aversive block neural activity showed a similar pattern, again producing excitation in lateral habenula neurons and inhibition in dopamine neurons (blue lines, Figures 4B and 4C). Lateral habenula neurons also had an overall tendency for lower firing rates during the ITI of the aversive block (Figure 4B).

At first sight, this pattern might be taken to imply that neurons anticipated rewards and punishments in similar manners. However, there was a marked tendency for changes in tonic activity during the ITI to be larger during the appetitive block than during the aversive block (dopamine appetitive mean −0.57 spikes/s, aversive −0.22 spikes/s, $p = 0.0017$; habenula appetitive +1.71 spikes/s, aversive +0.69 spikes/s, $p = 0.096$; signed-rank test). This raised the possibility that neurons preferentially anticipated rewards rather than punishments. In particular, the task was designed so that the blocks were presented in an alternating order: each aversive block was immediately followed by an appetitive block, and vice versa. Thus, tonic activity during the aversive block might have actually reflected the degree of proximity to the upcoming appetitive block.

To test this possibility, we analyzed the manner in which tonic ramping activity changed over the course of the two blocks (Figures 4D and 4E). This revealed a distinct pattern: tonic ramping activity was close to zero at the start of the aversive block, became significant near the end of the aversive block, and then maintained a high level during the appetitive block. The same pattern was found in both lateral habenula and dopamine neurons. Notably, in both populations tonic activity was stronger during the last half of the appetitive block than during the first half of the aversive block (habenula $p = 0.03$; dopamine $p < 10^{-3}$; signed-rank test), indicating that a transition from the appetitive block to the aversive block caused an abrupt decrease in tonic anticipatory activity. This pattern suggests that tonic activity preferentially encoded proximity to future rewards rather than future punishments.

We next asked whether tonic activity occurred in the same neurons as classic phasic responses to reward delivery. We calculated the correlation between a neuron's tonic activity and its phasic differential responses to reward cues, reward outcomes, punishment cues, and punishment outcomes (Experimental Procedures; Figure S4). This analysis revealed that appetitive block tonic activity was strongest in neurons that phasically signaled rewards and punishments in opposite directions, as though encoding motivational value. Specifically, dopamine neurons with strong tonic inhibition during the ITI also had strong positive responses to reward cues and outcomes, and had negative responses to punishment cues and outcomes (each $p < 0.05$, permutation test). In an analogous manner, lateral habenula neurons with strong tonic excitation during the ITI also had strong negative responses to reward cues and outcomes and had positive responses to punishment cues (each $p < 0.05$, permutation test). Thus, tonic encoding of rewards was linked to phasic encoding of motivational value.

As a further test of this conclusion, we focused on dopamine neurons. Whereas lateral habenula neurons primarily encode punishments in terms of motivational value, dopamine neurons can be divided into multiple types with distinct motivational signals: some are excited by punishments, while others are

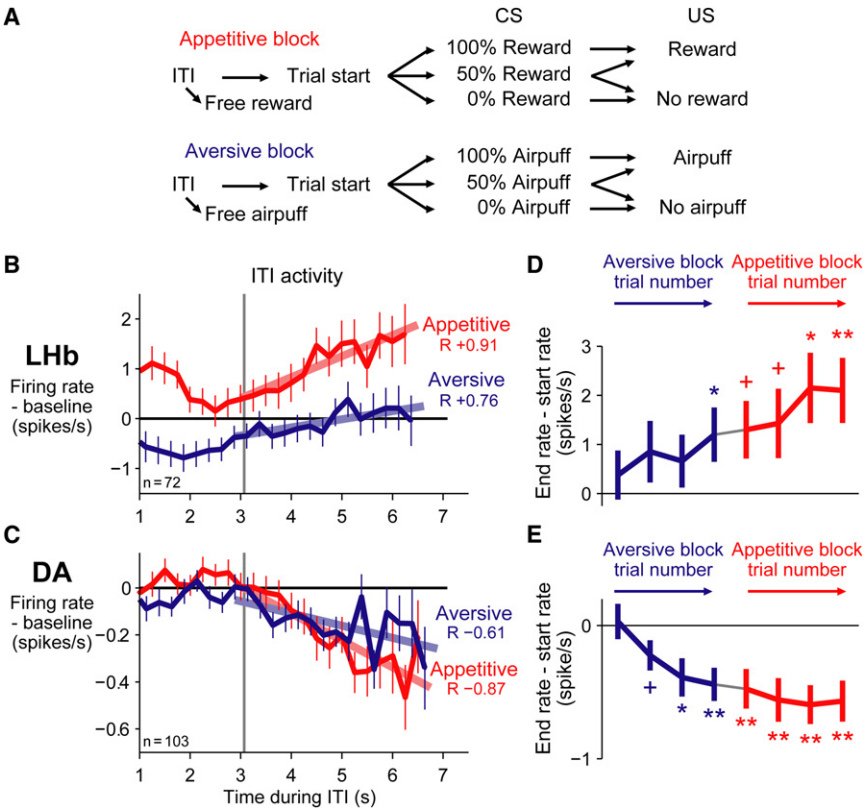

Figure 4. Tonic Activity Preferentially Encodes Rewarding Tasks

(A) Pavlovian conditioning procedure. In the appetitive block, a visual conditioned stimulus (CS) predicted juice rewards (US). In the aversive block, CSs predicted air puffs. Different CSs indicated different outcome probabilities. On a small number of ITIs an uncued "free" outcome was delivered, either a reward during the appetitive block or an air puff during the aversive block.

(B and C) Average neural ITI activity during the Pavlovian procedure, plotted separately for the appetitive block (red) and aversive block (blue). Same format as Figure 3. All correlations are significant ($p < 0.02$, permutation test).

(D and E) Changing intensity of tonic ITI ramping activity over the course of the aversive block (blue) and appetitive block (red). Activity is plotted for the 1st–4th quarters of each block, defined as trial numbers 2–12, 13–22, 23–32, and 33–42. For each quarter of each block we fit each neuron's ITI activity using a ramp function (as in Figure 2) and calculated the neuron's tonic effect as the difference (end firing rate) − (start firing rate). Each data point is the mean of the single-neuron tonic effects. Error bars are ± 1 SEM. Symbols indicate statistical significance and trends (+ indicates $p < 0.10$, * indicates $p < 0.05$, ** indicates $p < 0.01$).

See also Figure S4.

inhibited (Matsumoto and Hikosaka, 2009a, 2009b). To test whether these cells had distinct forms of tonic activity, we sorted dopamine neurons into types based on their phasic responses to aversive events (Figure 5C; Experimental Procedures). Dopamine neurons fell into four major groups: "aversive-inhibited type" that were inhibited by aversive cues and outcomes, "aversive-excited type" that were excited by aversive cues and outcomes, "aversive-mixed type" that were excited by aversive cues but inhibited by aversive outcomes, and "aversive-nonsignificant" type that did not react to aversive events with a consistent response (Figures 5A and 5B, right; Figure 5C).

This analysis revealed two distinct patterns of tonic activity. The first pattern resembled the time course seen in the population average of dopamine neuron activity: tonic inhibition that was close to zero at the start of the aversive block, then grew progressively stronger and reached a maximal level during the appetitive block (Figure 5A). This pattern occurred in the aversive-inhibited, aversive-mixed, and aversive-nonsignificant type neurons—in other words, neurons that were inhibited or non-responsive to the delivery of aversive outcomes (Figure 5A). For each of these types tonic activity was stronger in the appetitive block than the aversive block, indicating preferential encoding of rewards rather than punishments (each $p < 0.05$, signed-rank test).

A second pattern of tonic activity was found in aversive-excited type dopamine neurons (Figure 5B). Their tonic activity was quite weak; when combined over both blocks it was not significantly different from zero (mean −0.17 spikes/s, $p = 0.49$, signed-rank test). Furthermore, their tonic activity did not reach its maximal level during the appetitive block. On the contrary, it was significant when measured over the entire aversive block (mean −0.36 spikes/s, $p = 0.02$) and was close to zero when measured over the entire appetitive block (mean −0.02 spikes/s, $p = 0.93$; although the difference between blocks did not reach significance, $p = 0.23$). Thus, aversive-excited type dopamine neurons did not appear to tonically encode rewards and overall had weak or nonexistent tonic activity.

Phasic Task Anticipation Signals Encode Rewards and Punishments in Similar Manners

In addition to tonic activity anticipating the next trial, neurons had phasic responses marking the time when the trial start cue appeared and the task began. Did these tonic and phasic signals follow the same coding principles? In particular, did phasic signals also encode rewards in preference to punishments?

To test this, we measured neural responses to the trial start cue during the appetitive and aversive blocks (Figure 6). We found that neurons responded in similar manners during both blocks: lateral habenula neurons had similar inhibitions, and dopamine neurons had similar excitations (Figures 6A and 6B). Their response strength was not significantly different between the two blocks (habenula $p = 0.22$, dopamine $p = 0.21$; signed-rank test). Furthermore, neural responses during the two blocks were tightly correlated (Figures 6C and 6D). Most neurons clustered around the identity line indicating identical responses during both blocks (Figures 6C and 6D).

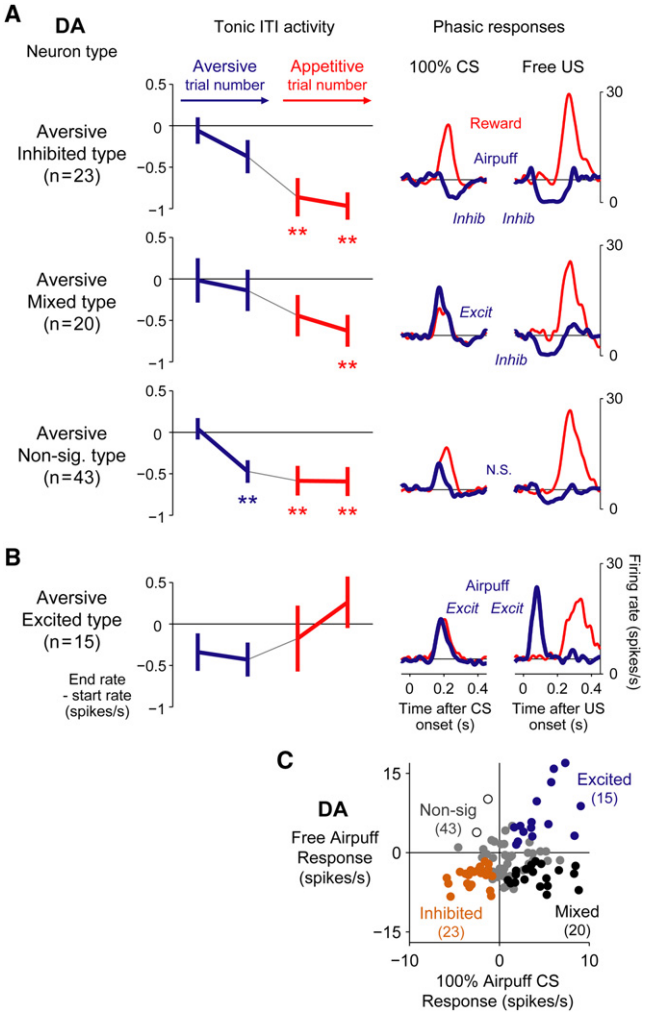

Figure 5. Time Course of Tonic Activity in Multiple Types of Dopamine Neurons

(A and B) Left: tonic ITI activity in multiple types of dopamine neurons with different responses to aversive events. Neurons were sorted into types based on their excitatory and inhibitory responses to air puff CSs and USs. Same format as Figure 4E, except that due to the relatively small number of neurons for each type, activity in each block was analyzed using two bins representing the first and second halves of each block (trials 2–22 and 23–42). Right: phasic responses of each neuron type to reward and air puff CSs and USs.

(C) Classification of dopamine neuron types based on responses to aversive cues (x axis, response to 100% airpuff CS) and aversive outcomes (y axis, response to free airpuff US). Responses were defined as the firing rate in a window after event onset minus the rate in a window before event onset (Experimental Procedures). Dots represent neurons and colors represent types of neurons: cells inhibited by the CS and US ("Inhibited," orange), excited by the CS and US ("Excited," blue), excited by the CS and inhibited by the US ("Mixed," black), and nonsignificantly responsive ("Non-sig," gray). Open circles indicate two neurons that had a rare mixed pattern of inhibition by the CS and excitation by the US.

See also Figure S5.

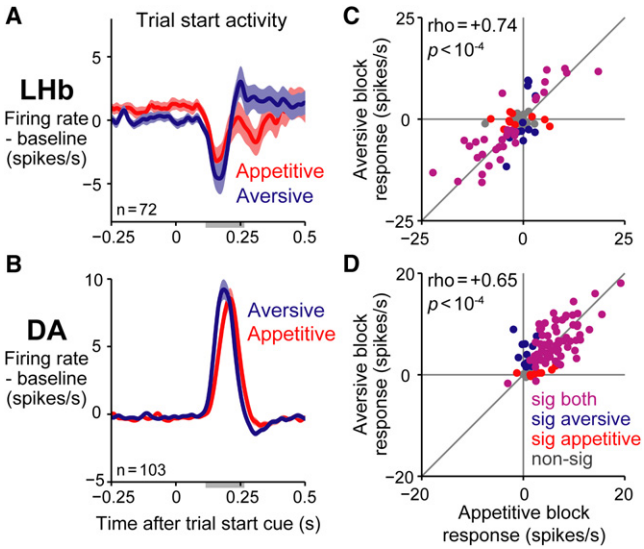

Figure 6. Phasic Trial Start Activity Encodes Both Rewarding and Aversive Tasks

(A and B) Population average activity in response to the trial start cue during the appetitive block (red) and aversive block (blue), separately for lateral habenula neurons (A) and dopamine neurons (B). Activity is baseline subtracted. Shaded region indicates ± 1 SEM. Neurons had similar responses in both appetitive and aversive blocks.

(C and D) Comparison between response to the trial start cue during the appetitive block (x axis) and aversive block (y axis). The response is the rate difference between a postcue window (gray bar below x axis) and a precue window (250 ms before the cue). Each dot is a single neuron. Colors indicate neurons with responses significantly different from zero during the appetitive block (red), aversive block (blue), or both (purple) ($p < 0.05$, signed-rank test). Text indicates the rank correlation and its p value (permutation test).

See also Figure S8.

This response pattern was found in all types of neurons and occurred at all times during the appetitive and aversive blocks (Figure 7). When dopamine neurons were classified into types based on their responses to aversive events, all types had positive responses to the trial start cue that were sustained throughout the aversive and appetitive blocks (Figure 7B). Furthermore, all types had similar response magnitudes during the two blocks ($p > 0.14$, signed-rank test) or else had slightly stronger responses during the aversive block (aversive-excited type, $p = 0.02$). This pattern was particularly striking in aversive-inhibited type dopamine neurons. These neurons showed the strongest possible evidence for phasic coding of motivational value: they were excited by reward cues and outcomes and were inhibited by aversive cues and outcomes (Figure 5A). In addition, these neurons detected the difference in value between the two blocks: upon a transition to the appetitive block they gained a stronger tonic inhibition during the ITI (Figure 5A) and gained an inhibitory response to the "0% outcome CS" which then cued omission of reward (Figure S5). Even so, these neurons treated the two blocks very similarly in their phasic responses to the trial start cue. They were strongly excited during both blocks with equal response magnitudes (mean response: aversive block +4.6 spikes/s, appetitive block +4.1 spikes/s, $p = 0.67$, signed-rank test).

A second potential distinction between neuron types occurred in the lateral habenula. Whereas dopamine neurons responded to the trial start cue with exclusive excitation (91 cells excited, 1 cell inhibited; each $p < 0.05$, signed-rank test), lateral habenula

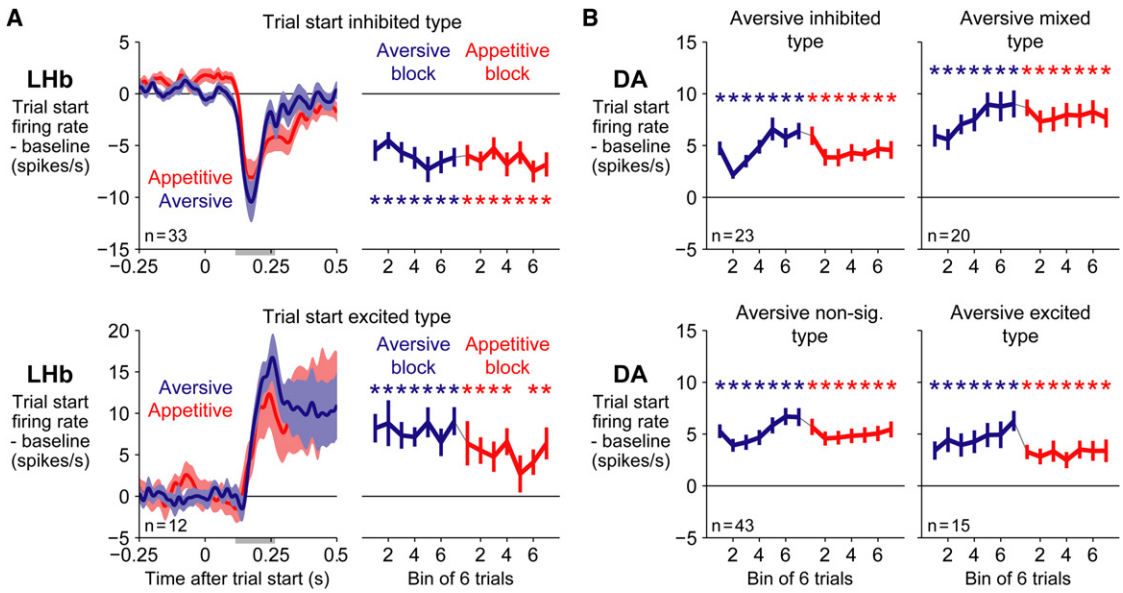

Figure 7. Trial Start Activity Encodes Rewarding and Aversive Tasks in All Neuron Types

(A) Left: population average activity in response to the trial start cue, shown separately for lateral habenula neurons that responded with inhibition (top, "Trial start inhibited") or excitation (bottom, "Trial start excited"; signed-rank test, $p < 0.05$). Same format as Figure 6A. Right: time course of trial start responses during the aversive block (blue) and appetitive block (red). Each data point is the mean response to the trial start cue during a selected group of trials within each block; blocks were divided into seven bins each containing six trials. Error bars are ±1 SEM. Asterisks indicate statistical significance ($p < 0.05$, signed-rank test). To prevent selection bias, this plot displays cross-validated data: the data displayed for each bin only include neurons whose inhibitory (top) or excitatory (bottom) responses were statistically significant when that bin's data were excluded from the analysis.

(B) Time course of trial start responses for the four types of dopamine neurons, using the classification in Figure 5. Each data point is the mean response to the trial start cue during a selected group of trials within each block; blocks were divided into seven bins each containing six trials. Error bars are ±1 SEM. Asterisks indicate statistical significance ($p < 0.05$, signed-rank test).

See also Figure S6.

neurons had two response types: some cells were excited and other cells were inhibited (12 excited, 33 inhibited; Figure 6C). Nonetheless, when these two types of lateral habenula neurons were analyzed separately, each type had phasic responses that were sustained throughout both aversive and appetitive blocks (Figure 7A). Again, response magnitudes were similar during the two blocks (trial start inhibited type, $p = 0.89$, signed-rank test) or were slightly stronger during the aversive block (trial start excited type, $p = 0.04$). Lateral habenula inhibitions occurred at a shorter latency than dopamine excitations, consistent with a habenula → dopamine direction of transmission (habenula inhibition 124 ± 5 ms, dopamine excitation 133 ± 2 ms, $p < 0.05$, bootstrap test; Figure S1). Lateral habenula excitations occurred at longer latencies (159 ± 9 ms) and were tonically sustained, suggesting that they may be generated by a different neural source; even so, these cells behaved similarly to other habenula neurons in their responses to other task events (Figure S6).

These data indicate that the phasic trial start response was distinct from other signals in lateral habenula and dopamine neurons. Unlike tonic ITI activity which preferentially encoded rewards, and unlike phasic responses to cues and outcomes which often encoded rewards and punishments in opposite manners, phasic responses to the trial start cue encoded rewarding and punishing tasks in similar manners.

Phasic Task Anticipation Signals Are Correlated with Orienting Reactions

Rewards and punishments are both motivationally salient events that have a potent ability to capture attention and eye movements (Lang and Davis, 2006; Matsumoto and Hikosaka, 2009b). Indeed, we observed that the trial start cue during the Pavlovian procedure typically evoked a rapid saccadic eye movement that shifted the animal's gaze to its location, even though no eye movement was required. We reasoned that if the neural response to the trial start cue reflected its motivational salience, then neurons might track trial-to-trial variations in its ability to attract saccadic eye movements. To test this, we analyzed the relationship between neural activity and behavior. For each neuron and for each block condition of the Pavlovian procedure we divided trials into two groups based on whether the trial start cue evoked a saccade with a fast reaction time or a slow reaction time (Figure 8). On trials when the cue evoked a fast saccadic reaction time, neural responses to the cue were enhanced. This effect occurred in both lateral habenula and dopamine neurons during both appetitive and aversive blocks (each $p < 0.05$, signed-rank test; Figures 8B and 8C).

Note that the neural response to the trial start cue was not simply coding for an eye movement command. First, the neural response was time-locked to stimulus onset, not to saccade

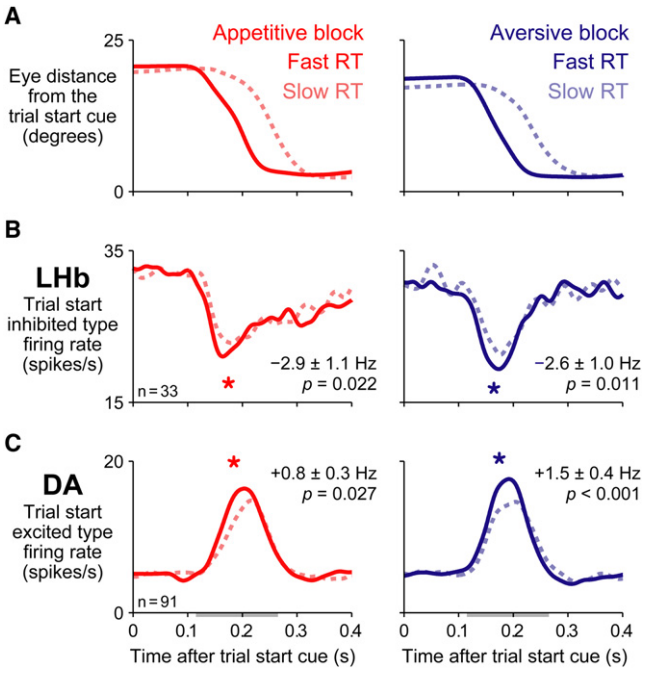

Figure 8. Trial Start Activity Is Correlated with Orienting Reactions

(A) Mean distance between the eye and the center of the trial start cue, plotted separately for the appetitive block (red, left) and aversive block (blue, right) and for the half of saccades when the animal's saccadic reaction time was fastest (solid lines, "Fast RT") or slowest (dashed lines, "Slow RT").

(B) Same format as (A), plotting the mean firing rate of the lateral habenula neurons that were inhibited by the trial start cue. The firing rate was quantified using the trial start analysis window (gray bar below the x axis in C); text indicates the mean difference in firing rate between fast and slow trials, its standard error, and the p value (signed-rank test; asterisks indicate $p < 0.05$).

(C) Same as (B), for dopamine neurons that were excited by the trial start cue. See also Figure S7.

onset (Figure S7). On trials with fast reaction times, most of the neural response occurred after the saccade had already been initiated (Figure 8). Second, neural responses did not simply reflect the magnitude of motor activation, because fast and slow saccadic reactions had similar amplitudes and durations (Figure S7). Third, neurons were not activated by spontaneous saccades or by other behavior during the intertrial interval (Figure S7). Thus, lateral habenula and dopamine neurons did not directly encode a movement plan but did track trial-by-trial variations in the cue's ability to attract saccadic eye movements.

DISCUSSION

Our data show that lateral habenula and dopamine neurons carry tonic and phasic signals that anticipate upcoming behavioral tasks in distinct manners. Tonic signals preferentially encoded rewarding tasks, while phasic signals encoded rewarding and punishing tasks in similar manners.

Tonic Anticipation of Rewards

It has long been theorized that the dopamine system is not limited to phasic responses and can also encode motivational events in its level of tonic activity (Goto et al., 2007; Niv et al., 2007; Grace, 1991). However, single-neuron evidence for tonic motivational signals has been mixed. Early reports suggested that some dopamine neurons have small changes in tonic activity in anticipation of arm movements to obtain food rewards (Romo and Schultz, 1990) or delivery of probabilistic rewards (Fiorillo et al., 2003), but these changes have not always been reported in later studies including our present datasets (Mirenowicz and Schultz, 1996; Hollerman and Schultz, 1998; Satoh et al., 2003; Morris et al., 2004; Matsumoto and Hikosaka, 2007, 2009b; Joshua et al., 2008; Bromberg-Martin and Hikosaka, 2009; Figure S5). A more recent study showed that dopamine neurons can decrease their spiking activity before the delivery of variably timed rewards (Fiorillo et al., 2008). Our data is consistent with this proposal, showing that dopamine neurons decrease their activity before the start of variably timed behavioral tasks. This result was replicated in five animals during three distinct tasks, and was complemented by reciprocal tonic activity in the lateral habenula consistent with its inhibitory influence over dopamine neurons. These data provide strong support for the presence of tonic anticipatory activity in the dopamine system.

Tonic activity preferentially encoded rewarding tasks and was correlated with phasic coding of motivational value. This suggests that tonic activity was also related to motivational value. Another potential explanation is that tonic activity was related to general arousal and that arousal levels were higher during the rewarding task than the punishing task. Note, however, that phasic neural responses to the trial start cue signaled these tasks with equal strength. Thus, whatever the underlying cause of tonic activity during the intertrial interval, it represented the tasks in a different manner than phasic responses to the trial start cue.

Tonic activity occurred at the time during the ITI when the trial start cue could be predicted to appear but was omitted or delayed, and occurred in the same direction as phasic neural responses when a reward itself was omitted or delayed (Figures S5 and S6). In this sense tonic activity was consistent with encoding of "reward prediction errors," albeit expressed in a tonic fashion rather than as a classic phasic pulse (Fiorillo et al., 2008). Note, however, that tonic signals did not map exactly onto conventional notions of reward prediction errors or motivational value. Tonic signals had a similar pattern in aversive-inhibited and aversive-mixed type dopamine neurons, even though these types responded in opposite directions when presented with aversive visual cues. Conversely, tonic signals were very different in aversive-inhibited and aversive-excited type dopamine neurons, even though these types were both excited by rewards.

How might these tonic anticipatory signals be used by the brain? Existing theories of dopamine function suggest several possibilities. First, dopamine release is thought to act as a reinforcement signal that causes adjustments in future behavior (Wise, 2004; Schultz et al., 1997). In this view, tonic inhibition could act as a "teaching signal" indicating that the trial began later than expected, causing the estimated ITI duration to be lengthened on future trials. Second, dopamine release is thought to generate motivation to persist in the pursuit of future rewards

(Berridge and Robinson, 1998; Salamone et al., 2007). Tonic inhibition could decrease the motivation to continue the upcoming task, promoting a switch to alternative activities. Third, dopamine release in the dorsal striatum has been proposed to set the "clock speed" of an internal timing mechanism (Buhusi and Meck, 2006). Tonic inhibition could slow down the clock, promoting greater patience when an upcoming task is delayed. To test between these possibilities, it could be necessary to record dopamine neuron activity during behavioral tasks that require rapid learning of timing distributions (Frank et al., 2009) and tradeoffs between patience and switching (Balci et al., 2009).

Tonic and phasic reward signals were both present in single neurons in the lateral habenula, suggesting that these signals are combined upstream of dopamine neurons and may not be restricted to the dopaminergic system. Lateral habenula activity has a potent influence on other neuromodulatory systems as well, including serotonin and norepinephrine (Lecourtier and Kelly, 2007). Further study will be needed to find the precise relationship between lateral habenula activity and neurons in downstream structures. Notably, lateral habenula and dopamine neurons tended to have similar changes in firing rate during their tonic ramping activity (~1–2 spikes/s) and phasic trial start responses (~10–15 spikes/s), but these were superimposed on very different baseline firing rates (~30 spikes/s for lateral habenula neurons versus ~5 spikes/s for dopamine neurons; see also Figure S3). This indicates that lateral habenula and dopamine firing rates are not simply scaled versions of each other but have a more complex relationship, possibly influenced by additional input from other brain areas.

Phasic Anticipation of Rewards and Punishments

A major goal of recent research has been to discover the neural basis of two distinct motivational signals, "valence" and "salience" (Anderson et al., 2003; Small et al., 2003; Jensen et al., 2007; Litt et al., 2010). Neurons encoding valence signal rewards and punishments in opposite manners, as if representing desire (Roitman et al., 2008). Neurons encoding salience signal rewards and punishments in similar manners, as if representing motivational importance or arousal (Lin and Nicolelis, 2008). This distinction between valence and salience is of special relevance to the lateral habenula-dopamine pathway. Several influential theories propose that dopamine primarily encodes a form of valence (or "value" or "wanting") for the purpose of learning and motivating reward-seeking behavior (Schultz et al., 1997; Berridge and Robinson, 1998; Wise, 2004). Other theories propose that dopamine primarily encodes a form of salience (or "alerting" or "timing") for the purpose of shifting attention to unpredicted events (Horvitz, 2000; Redgrave et al., 1999; Schultz, 1998) and marking their time of occurrence (Redgrave and Gurney, 2006).

Our data shows that lateral habenula and dopamine neurons are not bound to follow this simple dichotomy and can fulfill both roles at different times during a single task. Lateral habenula neurons and aversive-inhibited type dopamine neurons had differential responses to reward and punishment cues and outcomes, as though encoding valence. Yet the same neurons had similar responses to the start of rewarding and punishing tasks, as though encoding salience. The salience-like response was remarkably consistent, occurring in the great majority of dopamine neurons regardless of their various cue and outcome response types. This data is partially consistent with a proposal that dopamine neuron activity at the start of a behavioral task is related to motivational impact rather than expected reward value (Satoh et al., 2003). These salience-like signals may be sent to the habenula-dopamine pathway by neurons that specifically encode motivational salience, such as neurons of the basal forebrain (Lin and Nicolelis, 2008). On the other hand, given that salience-like and reward-related signals coexisted in single neurons, it is possible that the salience-like signals are sent to the habenula-dopamine pathway by the same brain areas that send them reward-related signals, such as the globus pallidus (Hong and Hikosaka, 2008).

How do lateral habenula and dopamine neurons decide when to encode valence and when to encode salience? Our data do not provide a conclusive answer, but do suggest a hypothesis. In the Pavlovian procedure, salience-like signals occurred in response to the trial start cue which marked the timing of an upcoming sequence of events but did not reveal new information about their value. In contrast, valence-related signals occurred for cues and outcomes that provided new information about motivational value but revealed little or no new information about the future event timing.

A close examination of dopamine neuron data from previous studies provides further evidence that their activity has a separable component related to the salience of timing cues. In experiments that use a trial start cue, the trial start cue evokes phasic excitation ("timing") and the later presentation of a negative stimulus that predicts lower than expected reward value evokes clear phasic inhibition ("value") (Satoh et al., 2003; Nakahara et al., 2004; Takikawa et al., 2004; Matsumoto and Hikosaka, 2007, 2009b; Bromberg-Martin and Hikosaka, 2009). Other experiments did not use a trial start cue. According to our hypothesis, the timing function would then be transferred to the first stimulus on each trial, which would gain an additional excitatory component to its response. Indeed, in these experiments when a negative stimulus is the first event of a trial, dopamine neurons are no longer primarily inhibited but instead are nonresponsive or even weakly excited (Mirenowicz and Schultz, 1996; Fiorillo et al., 2003; Morris et al., 2004; Day et al., 2007; Fiorillo et al., 2008; Joshua et al., 2008). Often these responses are strikingly biphasic (Schultz and Romo, 1990; Mirenowicz and Schultz, 1996; Waelti et al., 2001), as though a fast excitatory salience response was superimposed on a longer-latency inhibitory value response (Joshua et al., 2009). As a further test of this phenomenon, we analyzed data from a small number of experiments using a modified Pavlovian procedure in which the trial start cue was removed. According to our hypothesis, the trial start cue's timing function should be transferred to the first conditioned stimulus to appear on each trial. As predicted, the response to those stimuli gained an additional excitatory component in dopamine neurons and an additional inhibitory component in lateral habenula neurons and did so in a similar manner during both appetitive and aversive blocks (Figure S8).

An important goal for future experiments will be to discover whether this salience-like activity is related to abstract functions

of salience and timing or is a product of timing cues triggering motivational processes supported by the habenula-dopamine pathway such as orienting (Han et al., 1997) and information-seeking (Bromberg-Martin and Hikosaka, 2009). In particular, neural responses were correlated with the speed of orienting reactions to the start of a new task trial. Humans and animals orient to salient cues that indicate the timing of upcoming rewards and punishments (Hayhoe and Ballard, 2005; Van Damme et al., 2006; Peck et al., 2009; Matsumoto and Hikosaka, 2009b) and often treat them as incentives, actively seeking environments where informative cues are available (Badia et al., 1979; Miller, 1987; Bromberg-Martin and Hikosaka, 2009). By supporting these processes, salience-like activity may allow humans and animals to anticipate rewards and punishments with greater reliability and precision.

EXPERIMENTAL PROCEDURES

Database

We analyzed data collected in four previous studies. All experimental procedures and recording techniques can be found in our previous studies (Matsumoto and Hikosaka, 2007, 2009a, 2009b; Bromberg-Martin and Hikosaka, 2009). In brief, subjects were five rhesus macaque monkeys (*Macaca mulatta*), D, E, L, N, and Z. All procedures for animal care and experimentation were approved by the Institute Animal Care and Use Committee and complied with the Public Health Service Policy on the humane care and use of laboratory animals. A plastic head holder, scleral search coils, and plastic recording chambers were implanted under general anesthesia and sterile surgical conditions. Monkeys sat in a primate chair, facing a screen onto which visual stimuli were projected. Lateral habenula neurons were selected based on responsiveness to the experimental task. Midbrain dopamine neurons were recorded in and around the substantia nigra (Matsumoto and Hikosaka, 2007) or both substantia nigra and ventral tegmental area (Matsumoto and Hikosaka, 2009b; Bromberg-Martin and Hikosaka, 2009). Neurons were presumed to be dopaminergic based on their irregular tonic firing at 0.5–10 Hz and broad spike waveforms. Dopamine neurons were selected based on excitation by free reward (Matsumoto and Hikosaka, 2007, 2009b) or based on positive discrimination for both reward-predictive cues and unexpected reward outcomes or positive discrimination for one of those task events and no discrimination for the other task event (Bromberg-Martin and Hikosaka, 2009).

Behavioral Tasks

In this study, we analyzed task data recorded during intertrial intervals and during the response to the trial start cue. During the intertrial interval the animal faced a black screen. The trial start cue was a small white dot of light that appeared at the center of the screen. Neurons were recorded after animals had extensive experience with the tasks and intertrial intervals being tested. The ITI duration on each trial was randomly generated at 1 ms resolution. For descriptions of each task see below and Supplemental Experimental Procedures.

In the reward-biased saccade task (Matsumoto and Hikosaka, 2007) the intertrial intervals were 2.2 s (animal L, 14 lateral habenula and 22 dopamine neurons; Figures 3A and 3E), 1.7–2.7 s (animal L, 8 lateral habenula neurons; Figure 3C), or 2.2–3.2 s (animal L, 15 lateral habenula and 20 dopamine neurons; animal E, 28 lateral habenula and 20 dopamine neurons; Figures 3B and 3F). The trial start cue acted as a fixation point which animals were required to fixate to begin the trial. Animals were then required to saccade to visual targets indicating future reward or no-reward outcomes. Half of trials ended in a reward (0.3 ml of apple juice) and the other half were unrewarded.

In the information choice task (Bromberg-Martin and Hikosaka, 2009) the intertrial interval was 4.1–5.1 s (animal E, 20 dopamine neurons; animal Z, 27 dopamine neurons; Figure 3G). The trial start cue acted as a fixation point which animals were required to fixate to begin the trial. Animals then performed a saccadic decision task to choose whether to view visual cues that provided information about future reward outcomes. Half of trials ended in a big reward (~1.0 ml of water) and the other half ended in a small reward (0.04 ml of water).

In the Pavlovian procedure (Matsumoto and Hikosaka, 2009a), the intertrial intervals were 3.1–6.1 s (animal D, 11 lateral habenula and 34 dopamine neurons) or 3.1–7.1 s (animal D, 16 lateral habenula and 1 dopamine neuron; animal N, 45 lateral habenula and 68 dopamine neurons; Figures 3D and 3H). The animal was not required to make any behavioral response. On each trial the trial start cue was presented for 1 s, followed by a visual conditioned stimulus (CS) for 1.5 s, followed by the outcome (unconditioned stimulus [US] or US omission). The task alternated between appetitive blocks in which the US was a juice reward, and aversive blocks in which the US was an air puff (20–30 psi, delivered near the face through a narrow tube). Each block had three CSs: 100% CS (followed by the US on all trials), 50% CS (followed by either US or no US with equal probability), and 0% CS (followed by no US on all trials). The 0% CS was identical in both appetitive and aversive blocks; the other CSs were different images in the two blocks. On a small number of trials no trial start cue or CS was presented and the US was delivered without any signal ("free reward" or "free air puff"). Each block consisted of 42 trials (12 100% CS, 12 50% CS, 12 0% CS, 6 free outcome). The two blocks alternated without any external signal indicating the block transition, and each neuron was recorded for at least four blocks. Animals reliably detected block transitions and reversed their behavior and classic phasic responses within 1–3 trials (Matsumoto and Hikosaka, 2009a).

Data Analysis

The start of the ITI was defined as outcome onset for the information choice task and outcome offset for the other tasks. The end of the ITI was defined as the time of the trial start cue plus 40 ms. The firing rate in response to the trial start cue was defined as the firing rate in a window 115–265 ms after cue onset, which was chosen to include the major part of the excitatory and inhibitory neural response in all three tasks. Each neuron's activity during the ITI was fitted with a linear ramp-like function with two parameters, a starting rate and an ending rate. The function can be considered to represent the spiking probability at each millisecond because it was fitted to binary spike data (0 for no spike or 1 for a spike, at millisecond resolution). The parameters providing the maximum-likelihood fit were found using the MATLAB function "fminunc." The amount of data for measuring tonic activity was linearly decreasing during the ITI because only a small fraction of ITIs lasted for the maximal duration; nonetheless, our analysis procedures recovered an accurate estimate of the time course of ITI activity, indicated by its correct behavior on simulated datasets (data not shown).

Correlations were calculated with rank correlation (rho) except in Figures 3 and 4 where linear correlation (r) was used to evaluate whether firing rate changes were close to linear with time. Significance was determined with permutation tests (20,000 permutations). Each neuron's baseline firing rate was calculated using a window 1000–1600 ms after the start of the ITI for the reward-biased saccade task, and 1600–2100 ms after the start of the ITI for the information choice task and Pavlovian procedure to avoid contamination from phasic responses to the previous trial's outcome. In Figures 4–7, baseline activity was calculated using data from both appetitive and aversive blocks.

For plots of smoothed activity, firing rates were smoothed with a Gaussian kernel (σ = 10 ms). For the plots of single-neuron or population average trial start responses (Figures 6 and 7), the response to the trial start cue was defined as the firing rate during a 115–265 ms window after cue onset minus the firing rate in a window 250 ms before cue onset.

In the Pavlovian procedure, all analyses were restricted to trials when the current block's appetitive or aversive identity could be known. Thus, trials were excluded if they occurred at the start of a block before the first reward or air puff outcome was delivered or 100% or 50% CS was presented. Phasic neural firing rates were defined as the firing rate in a time window after an event chosen to contain the main component of neural responses, which were as follows: trial start cue, 115–265 ms; CS, lateral habenula 150–400 ms, dopamine 150–325 ms; reward US, lateral habenula 200–500 ms, dopamine 200–400 ms; air puff US, lateral habenula 50–150 ms, dopamine 50–200 ms; reward US omission, lateral habenula 200–500 ms, dopamine 200–500 ms;

air puff US omission, lateral habenula 50–150 ms, dopamine 150–350 ms (Matsumoto and Hikosaka, 2009a, 2009b).

Dopamine neurons were classified into four types based on their responses to both the aversive CSs and USs. Similar results were obtained if neurons were separately classified based on CS responses alone or US responses alone as in Matsumoto and Hikosaka (2009b). The responses to the 100% air puff CS and the free air puff US were defined as the phasic neural firing rates for those events minus the firing rate in a window 250 ms before event onset. Then for each neuron, both responses were tested for being significantly different from zero (Wilcoxon signed-rank test; significance was determined at a level of $p < 0.05^{1/2}$, so that the false positive rate detecting a neuron with two significant responses was controlled at the level of $\alpha = (0.05^{1/2})^2 = 0.05$). Neurons were classified as aversive responsive if both responses reached significance: as aversive-inhibited type if both responses were negative, aversive-excited type if both responses were positive, and aversive-mixed type if the response was positive to the 100% air puff CS and negative to the free air puff US. Only two neurons had the opposite mixed type (negative to the 100% air puff CS and positive to the free air puff US) which were excluded from this analysis. The remaining neurons were classified as aversive-nonsignificant type.

A full description of the analysis of neural response latencies, saccadic reaction times, and tonic and phasic models of spiking activity is in Supplemental Experimental Procedures.

SUPPLEMENTAL INFORMATION

Supplemental Information includes eight figures and Supplemental Experimental Procedures and can be found with this article online at doi:10.1016/j.neuron.2010.06.016.

ACKNOWLEDGMENTS

We thank H. Nakahara, S. Kaveri, T. Yang, L. Paninski, S. Hong, M. Yasuda, S. Yamamoto, Y. Tachibana, and H. Kim for valuable discussions. This work was supported by the intramural research program at the National Eye Institute.

Accepted: June 1, 2010
Published: July 14, 2010

REFERENCES

Anderson, A.K., Christoff, K., Stappen, I., Panitz, D., Ghahremani, D.G., Glover, G., Gabrieli, J.D.E., and Sobel, N. (2003). Dissociated neural representations of intensity and valence in human olfaction. Nat. Neurosci. *6*, 196–202.

Badia, P., Harsh, J., and Abbott, B. (1979). Choosing between predictable and unpredictable shock conditions: data and theory. Psychol. Bull. *86*, 1107–1131.

Balci, F., Freestone, D., and Gallistel, C.R. (2009). Risk assessment in man and mouse. Proc. Natl. Acad. Sci. USA *106*, 2459–2463.

Berridge, K.C., and Robinson, T.E. (1998). What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? Brain Res. *28*, 309–369.

Bromberg-Martin, E.S., and Hikosaka, O. (2009). Midbrain dopamine neurons signal preference for advance information about upcoming rewards. Neuron *63*, 119–126.

Buhusi, C.V., and Meck, W.H. (2006). What makes us tick? Functional and neural mechanisms of interval timing. Nat. Rev. Neurosci. *6*, 755–765.

Christoph, G.R., Leonzio, R.J., and Wilcox, K.S. (1986). Stimulation of the lateral habenula inhibits dopamine-containing neurons in the substantia nigra and ventral tegmental area of the rat. J. Neurosci. *6*, 613–619.

Day, J.J., Roitman, M.F., Wightman, R.M., and Carelli, R.M. (2007). Associative learning mediates dynamic shifts in dopamine signaling in the nucleus accumbens. Nat. Neurosci. *10*, 1020–1028.

Fiorillo, C.D., Tobler, P.N., and Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. Science *299*, 1898–1902.

Fiorillo, C.D., Newsome, W.T., and Schultz, W. (2008). The temporal precision of reward prediction in dopamine neurons. Nat. Neurosci. *11*, 966–973.

Frank, M.J., Doll, B.B., Oas-Terpstra, J., and Moreno, F. (2009). Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation. Nat. Neurosci. *12*, 1062–1068.

Goto, Y., Otani, S., and Grace, A.A. (2007). The Yin and Yang of dopamine release: a new perspective. Neuropharmacology *53*, 583–587.

Grace, A.A. (1991). Phasic versus tonic dopamine release and the modulation of dopamine system responsivity: a hypothesis for the etiology of schizophrenia. Neuroscience *41*, 1–24.

Han, J.S., McMahan, R.W., Holland, P., and Gallagher, M. (1997). The role of an amygdalo-nigrostriatal pathway in associative learning. J. Neurosci. *17*, 3913–3919.

Hayhoe, M., and Ballard, D. (2005). Eye movements in natural behavior. Trends Cogn. Sci. *9*, 188–194.

Hollerman, J.R., and Schultz, W. (1998). Dopamine neurons report an error in the temporal prediction of reward during learning. Nat. Neurosci. *1*, 304–309.

Hong, S., and Hikosaka, O. (2008). The globus pallidus sends reward-related signals to the lateral habenula. Neuron *60*, 720–729.

Horvitz, J.C. (2000). Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. Neuroscience *96*, 651–656.

Jensen, J., Smith, A.J., Willeit, M., Crawley, A.P., Mikulis, D.J., Vitcu, I., and Kapur, S. (2007). Separate brain regions code for salience vs. valence during reward prediction in humans. Hum. Brain Mapp. *28*, 294–302.

Ji, H., and Shepard, P.D. (2007). Lateral habenula stimulation inhibits rat midbrain dopamine neurons through a GABA(A) receptor-mediated mechanism. J. Neurosci. *27*, 6923–6930.

Joshua, M., Adler, A., Mitelman, R., Vaadia, E., and Bergman, H. (2008). Midbrain dopaminergic neurons and striatal cholinergic interneurons encode the difference between reward and aversive events at different epochs of probabilistic classical conditioning trials. J. Neurosci. *28*, 11673–11684.

Joshua, M., Adler, A., and Bergman, H. (2009). The dynamics of dopamine in control of motor behavior. Curr. Opin. Neurobiol. *19*, 615–620.

Lang, P.J., and Davis, M. (2006). Emotion, motivation, and the brain: reflex foundations in animal and human research. Prog. Brain Res. *156*, 3–29.

Lecourtier, L., and Kelly, P.H. (2007). A conductor hidden in the orchestra? Role of the habenular complex in monoamine transmission and cognition. Neurosci. Biobehav. Rev. *31*, 658–672.

Lin, S.-C., and Nicolelis, M.A. (2008). Neuronal ensemble bursting in the basal forebrain encodes salience irrespective of valence. Neuron *59*, 138–149.

Litt, A., Plassmann, H., Shiv, B., and Rangel, A. (2010). Dissociating valuation and saliency signals during decision-making. Cereb. Cortex, in press. Published online May 5, 2010. 10.1093/cercor/bhq065.

Macdougall, J.M., Van Hoesen, G.W., and Mitchell, J.C. (1969). Anatomical organization of septal projections in maintenance of DRL behavior in rats. J. Comp. Physiol. Psychol. *68*, 568–575.

Matsumoto, M., and Hikosaka, O. (2007). Lateral habenula as a source of negative reward signals in dopamine neurons. Nature *447*, 1111–1115.

Matsumoto, M., and Hikosaka, O. (2009a). Representation of negative motivational value in the primate lateral habenula. Nat. Neurosci. *12*, 77–84.

Matsumoto, M., and Hikosaka, O. (2009b). Two types of dopamine neuron distinctly convey positive and negative motivational signals. Nature *459*, 837–841.

Miller, S.M. (1987). Monitoring and blunting: validation of a questionnaire to assess styles of information seeking under threat. J. Pers. Soc. Psychol. *52*, 345–353.

Mirenowicz, J., and Schultz, W. (1996). Preferential activation of midbrain dopamine neurons by appetitive rather than aversive stimuli. Nature *379*, 449–451.

Morris, G., Arkadir, D., Nevet, A., Vaadia, E., and Bergman, H. (2004). Coincident but distinct messages of midbrain dopamine and striatal tonically active neurons. Neuron *43*, 133–143.

Nakahara, H., Itoh, H., Kawagoe, R., Takikawa, Y., and Hikosaka, O. (2004). Dopamine neurons can represent context-dependent prediction error. Neuron *41*, 269–280.

Niv, Y., Duff, M.O., and Dayan, P. (2005). Dopamine, uncertainty and TD learning. Behav. Brain Funct. *1*, 6.

Niv, Y., Daw, N.D., Joel, D., and Dayan, P. (2007). Tonic dopamine: opportunity costs and the control of response vigor. Psychopharmacology (Berl.) *191*, 507–520.

Peck, C.J., Jangraw, D.C., Suzuki, M., Efem, R., and Gottlieb, J. (2009). Reward modulates attention independently of action value in posterior parietal cortex. J. Neurosci. *29*, 11182–11191.

Redgrave, P., and Gurney, K. (2006). The short-latency dopamine signal: a role in discovering novel actions? Nat. Rev. Neurosci. *7*, 967–975.

Redgrave, P., Prescott, T.J., and Gurney, K. (1999). Is the short-latency dopamine response too short to signal reward error? Trends Neurosci. *12*, 146–151.

Roitman, M.F., Wheeler, R.A., Wightman, R.M., and Carelli, R.M. (2008). Real-time chemical responses in the nucleus accumbens differentiate rewarding and aversive stimuli. Nat. Neurosci. *11*, 1376–1377.

Romo, R., and Schultz, W. (1990). Dopamine neurons of the monkey midbrain: contingencies of responses to active touch during self-initiated arm movements. J. Neurophysiol. *63*, 592–606.

Salamone, J.D., Correa, M., Farrar, A., and Mingote, S.M. (2007). Effort-related functions of nucleus accumbens dopamine and associated forebrain circuits. Psychopharmacology (Berl.) *191*, 461–482.

Satoh, T., Nakai, S., Sato, T., and Kimura, M. (2003). Correlated coding of motivation and outcome of decision by dopamine neurons. J. Neurosci. *23*, 9913–9923.

Schultz, W. (1998). Predictive reward signal of dopamine neurons. J. Neurophysiol. *80*, 1–27.

Schultz, W. (2007). Multiple dopamine functions at different time courses. Annu. Rev. Neurosci. *30*, 259–288.

Schultz, W., and Romo, R. (1990). Dopamine neurons of the monkey midbrain: contingencies of responses to stimuli eliciting immediate behavioral reactions. J. Neurophysiol. *63*, 607–624.

Schultz, W., Dayan, P., and Montague, P.R. (1997). A neural substrate of prediction and reward. Science *275*, 1593–1599.

Small, D.M., Gregory, M.D., Mak, Y.E., Gitelman, D., Mesulam, M.M., and Parrish, T. (2003). Dissociation of neural representation of intensity and affective valuation in human gustation. Neuron *39*, 701–711.

Takikawa, Y., Kawagoe, R., and Hikosaka, O. (2004). A possible role of midbrain dopamine neurons in short- and long-term adaptation of saccades to position-reward mapping. J. Neurophysiol. *92*, 2520–2529.

Thornton, E.W., and Evans, J.A. (1984). The effects of lesions of the habenula nucleus on lever press behaviour during a tandem operant schedule with contrasting response requirements. Behav. Brain Res. *12*, 327–334.

Thornton, E.W., Bradbury, G.E., Wickens, A.P., Mottram, D.R., and McClelland, R. (1990). Intra-habenular injection of 6-hydroxydopamine produces impaired acquisition of DRL operant behavior. Behav. Neural Biol. *53*, 291–297.

Van Damme, S., Crombez, G., Hermans, D., Koster, E.H., and Eccleston, C. (2006). The role of extinction and reinstatement in attentional bias to threat: a conditioning approach. Behav. Res. Ther. *44*, 1555–1563.

Waelti, P., Dickinson, A., and Schultz, W. (2001). Dopamine responses comply with basic assumptions of formal learning theory. Nature *412*, 43–48.

Wise, R.A. (2004). Dopamine, learning and motivation. Nat. Rev. Neurosci. *5*, 483–494.