Report

Current Biology

Novelty, Salience, and Surprise Timing Are Signaled by Neurons in the Basal Forebrain

Highlights

- The basal forebrain (BF) predicts, and signals the statistics of, many salient events
- Some BF cells ramp in anticipation of novel objects and surprising reinforcements
- Their ramps are sensitive to internal estimates of, and confidence in, event timing
- Other BF cells burst to external events and convey their high-order statistics

Authors

Kaining Zhang, Charles D. Chen, Ilya E. Monosov

Correspondence

ilya.monosov@gmail.com

In Brief

Zhang et al. show that cells in the primate basal forebrain (BF), an area that mediates activity of the neocortex, predict the timing of events that capture attention, such as surprising reinforcements and novel objects. After a salient event, other BF cells' burst activations rapidly signal higher-order statistical information about motivational salience, novelty, and surprise.

Zhang et al., 2019, Current Biology *29*, 134–142
January 7, 2019
https://doi.org/10.1016/j.cub.2018.11.012

CellPress

CellPress

Current Biology
Report

Novelty, Salience, and Surprise Timing Are Signaled by Neurons in the Basal Forebrain

Kaining Zhang,[1,2] Charles D. Chen,[1] and Ilya E. Monosov[1,2,3,*]
[1]Department of Neuroscience, Washington University in St. Louis, St. Louis, MO 63110
[2]Department of Biomedical Engineering, Washington University in St. Louis, St. Louis, MO 63110, USA
[3]Lead Contact
*Correspondence: ilya.monosov@gmail.com
https://doi.org/10.1016/j.cub.2018.11.012

SUMMARY

The basal forebrain (BF) is a principal source of modulation of the neocortex [1–6] and is thought to regulate cognitive functions such as attention, motivation, and learning by broadcasting information about salience [2, 3, 5, 7–19]. However, events can be salient for multiple reasons—such as novelty, surprise, or reward prediction errors [20–24]—and to date, precisely which salience-related information the BF broadcasts is unclear. Here, we report that the primate BF contains at least two types of neurons that often process salient events in distinct manners: one with phasic burst responses to cues predicting salient events and one with ramping activity anticipating such events. Bursting neurons respond to cues that convey predictions about the magnitude, probability, and timing of primary reinforcements. They also burst to the reinforcement itself, particularly when it is unexpected. However, they do not have a selective response to reinforcement omission (the unexpected absence of an event). Thus, bursting neurons do not convey value-prediction errors but do signal surprise associated with external events. Indeed, they are not limited to processing primary reinforcement: they discriminate fully expected novel visual objects from familiar objects and respond to object-sequence violations. In contrast, ramping neurons predict the timing of many salient, novel, and surprising events. Their ramping activity is highly sensitive to the subjects' confidence in event timing and on average encodes the subjects' surprise after unexpected events occur. These data suggest that the primate BF contains mechanisms to anticipate the timing of a diverse set of important external events (via ramping activity) and to rapidly deploy cognitive resources when these events occur (via short latency bursting).

RESULTS

Previous work suggests that two prominent neuronal activation patterns in the BF support its mediation of cognitive functions in response to salient events: phasic bursting [8, 25], which has been identified in the brains of rodents, and tonic activations [4, 8], which in monkeys are often seen in neurons that also ramp to the time of delivery of uncertain or noxious outcomes [4]. To date, it remains unclear how these neuronal activations signal surprise and/or novelty and how their surprise-related responses relate to errors in estimates of state values, referred to as reward prediction errors (RPEs). Therefore, how bursting and ramping BF activations contribute to cognitive functions remains poorly understood. Here, we assessed whether prediction-related phasic bursting and ramping activity occur in distinct groups of neurons and tested whether and how the BF represents prediction errors, surprise, value, novelty, and timing.

CS-Related Phasic and Ramping Activity Are Observed in Mostly Distinct BF Cell Groups that Differentially Signal Reinforcement Statistics

We recorded BF neurons in 5 monkeys that participated in a Pavlovian procedure in which they experienced reward predictions that varied in magnitude and probability [4, 26, 27]. A reward-probability block contained five conditioned stimuli (CSs) associated with five probabilistic reward predictions (0, 25, 50, 75, and 100% of 0.25 mL of juice). A reward-amount block contained five other CSs associated with certain reward predictions of varying reward amounts (0.25, 0.1875, 0.125, 0.065, and 0 mL). During neuronal recording, any neuron that displayed ramping and/or phasic burst responses in the CS epoch of this Pavlovian procedure was recorded (n = 70; monkey H = 15, monkey P = 16, monkey B = 10, monkey R = 12, and monkey Z = 17).

Example neurons are shown in Figure 1A. The first neuron (Figure 1A, top) displayed short latency bursting after the presentation of the probability and amount CSs. This phasic activation was greatest following the presentation of the CS associated with the highest expected value in either the reward-probability or the reward-amount block and least following the presentation of the CSs associated with the lowest expected value (no reward). In either block, the bursting activity was strongly correlated with the expected value (Spearman's rank correlation; probability block, $\rho = 0.84$, $p < 0.0001$; amount block, $\rho = 0.86$, $p < 0.0001$). The second neuron (Figure 1A, bottom) had a very different response. Shortly after the CSs were presented, it displayed a consistent CS-onset-related inhibition that was greatest in the low-value trials and less apparent during high-value trials, on average roughly scaling with the expected value. In the reward-probability block, this initial change was followed by ramping activity to the time of the uncertain

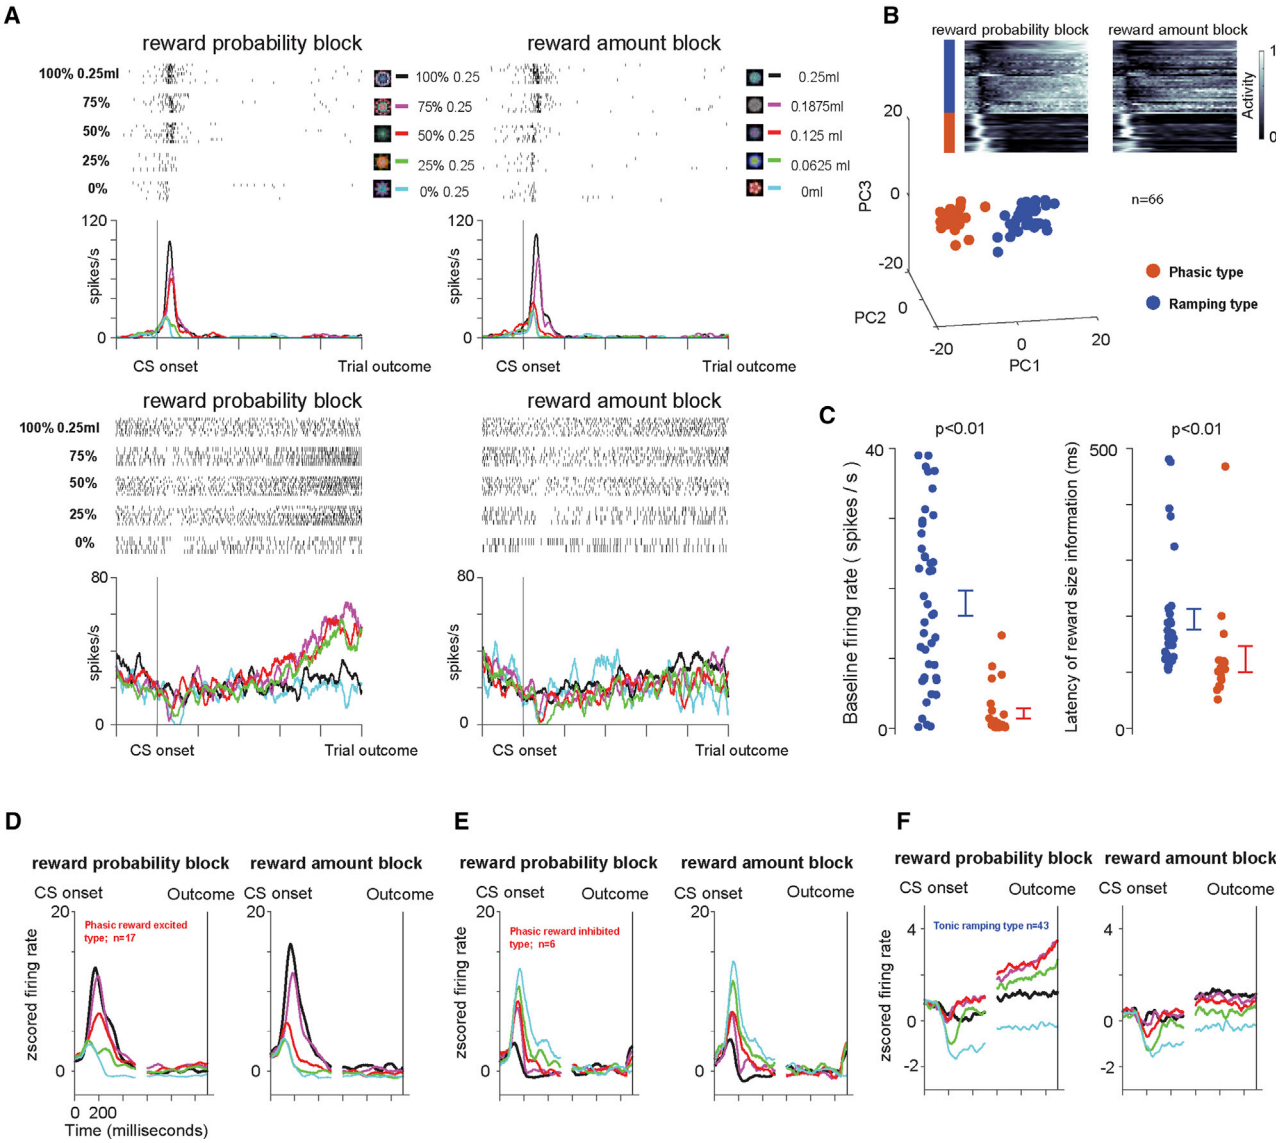

Figure 1. Two Groups of BF Neurons Encode the Magnitude and Probability of Reinforcement in Distinct Manners

(A) Responses of two example BF neurons (top and bottom) to the presentation of 10 fractal objects associated with certain and uncertain predictions of juice rewards in the reward-probability block (left) and reward-amount block (right).

(B) Clustering of BF neurons based on average activity in the probability block. The inset heatmap shows the activity of 66 BF neurons (normalized from 0 to 1 to the minimum and maximum in the reward-probability block) from the time of the CS onset to the time of the trial outcome (reward or no reward) in the reward-probability and reward-amount blocks. Each line represents the average activity across all 5 trial types in the block for each neuron. Below are the results of principal-component analyses performed on those normalized CS response functions. *K*-means clustering (STAR Methods) was used to separate the neurons into two groups: red group (n = 23) and blue group (n = 43). The group identities of the neurons are also indicated by a color bar to the left of the heatmap.

(C) The two clusters of neurons (red and blue) display distinct baseline firing rates (left) and latencies of value coding (right) in the reward-amount block. Each dot represents data from a single neuron. Error bars around the mean show the SEM.

(D) Average responses of the neurons in the blue group in the reward-probability block (left) and reward-amount block (right).

(E and F) Average responses of the neurons in the red group in the reward-probability block (left) and reward-amount block (right). (E) shows neurons that displayed greater activation for reward versus no-reward trials, while (F) shows neurons that displayed greater activation for no-reward trials.

See also Figure S1, STAR Methods, and the associated Figure S2 for details and anatomical locations of neuronal recordings.

(or risky) reward delivery (following 75%, 50%, and 25% CSs). The neuron's activity was significantly fit by a model of uncertainty ($\rho = 0.77$, $p = 0.0001$; measured in the last 500 ms) but not expected value ($\rho = -0.01$, $p = 0.94$). In the reward-amount block, in which all trials were certain, the neuron represented the expected value until the time of the reinforcement in its tonic activity ($\rho = 0.70$, $p < 0.0001$; measured in the last 500 ms). These example neurons suggest that the BF may contain functionally distinct classes of neurons: phasic bursting neurons that co-vary with the magnitude and probability of reinforcements

and tonic neurons that ramp, predicting the timing of uncertain outcomes.

To test this, we clustered BF neurons based on their average responses. Only neurons that had been recorded in every condition in both blocks were included (n = 66/70). Importantly, their response vectors were obtained by averaging the neuronal activity across all five CSs in the reward-probability block and were subsequently normalized from 0 to 1 (Figure 1B, inset). Therefore, neuronal tuning (e.g., representation of reward probability) and baseline firing rates were not considered in the clustering analysis.

This analysis revealed two clusters (Figure 1B). The first cluster (red; n = 23) showed clear bursting after the CS onset (see the neurons' response vectors in Figure 1B, inset). In contrast, the second cluster (blue; n = 43) showed an initial suppression following the CS onset and a slow ramp-like increase in activity as the trial's outcome neared.

The two clusters had different baseline firing rates (Figure 1C): one had relatively high firing rates (blue cluster; average frequency = 18 Hz; SD = 12 Hz) and the other low (red cluster; average frequency = 2.1 Hz; SD = 3.5 Hz). Both clusters' initial CS responses co-varied with the magnitude of the predicted reward, initially coding expected value, but the latency of this information was different among the two clusters. The expected value was conveyed earlier by the neurons in the phasic bursting red cluster (Figure 1C, right; rank-sum test; $p < 0.01$; blue cluster, average = 195 ms, median = 159 ms, SD = 104 ms; red cluster, average = 123 ms, median = 100 ms, SD = 96 ms).

These clusters differed in how they represented both the probability and amount of reinforcement (Figures 1D–1F and S1). Phasic bursting neurons (red cluster) signaled the expected value of the CSs in their bursting activations. The bursting activity was correlated with the probability in the probability block ($\rho = 0.60$, $p < 0.0001$) and with the reward amount in the amount block ($\rho = 0.71$, $p < 0.0001$). Tonic ramping neurons' initial suppression co-varied with the expected value ($\rho = 0.47$, $p < 0.0001$ in the probability block; $\rho = 0.46$, $p < 0.0001$ in the amount block). However, in trials in which reward was uncertain, they displayed additional ramping activity toward the trial outcome [4]. The activity during these 75%, 50%, and 25% CS trials was correlated with the probability of reinforcement delivery (Spearman's rank correlation; $\rho = 0.25$, $p = 0.0043$; pre-outcome analysis window −0.5 s before outcome was delivered). And, on average, the pre-outcome activity in the reward-probability block (across all 5 trial types) was correlated with uncertainty ($\rho = 0.71$, $p < 0.0001$; same analysis window as above).

Locations of phasic bursting and tonic ramping neurons were reconstructed using *in vivo* MRI STAR Methods [28] (Figure S2). Both phasic bursting and tonic ramping neurons were found within the BF, in the diagonal band of Broca and the nucleus basalis of Meynert [1, 4, 29].

Phasic and Ramping Neurons Signal Early Versus Late Rewards under Temporal Uncertainty

Does ramping of BF neurons encode the estimated timing of uncertain rewards? If so, then if rewards were certain but their timing was uncertain, the neurons should display ramping activity to the time of the earliest possible reward. Second, phasic bursting neurons' bursts seemed to scale with the expected values of the CSs, regardless of whether the value was manipulated by probability or amount. Might these neurons also encode the value of early versus late rewards?

To answer these questions, we designed a reward-timing procedure (Figure S3). Here, five distinct visual-fractal objects served as CSs that predicted either (1) a probabilistic delay before a reward with deterministic delivery (delays = 1.5 or 4.5 s; reward-timing-uncertain CSs) or (2) a deterministic delay before a reward with 0.5 probability of delivery (reward-probability CS). To test how phasic bursting neurons and tonic ramping neurons encode temporally uncertain reward predictions, we first identified them using the task in Figure 1 and then recorded them in this reward-timing procedure (n = 52; monkey W = 21, monkey B = 6, monkey R = 15, and monkey Z = 10).

Tonic ramping neurons displayed ramping activity in the 0.75, 0.5, and 0.25 reward-timing-uncertain conditions (Figure S3). The magnitude of this activation was correlated with the probability of reinforcement delivery at 1.5 s (Spearman's rank correlation; $\rho = 0.48$, $p < 0.0001$; analysis window, 1 s to 1.5 s). Interestingly, significant ramping was also observed to certain late reward at 4.5 s (Figure S3). Therefore, BF ramping tracks reward delivery during temporal-reward uncertainty (before 1.5 s) and during relatively longer epochs in which there is temporal uncertainty due to noise in interval timing.

Phasic bursting neurons' activity scaled with reward timing such that highest activity was evoked by CSs predicting the earliest reward (Figure S3). Their average activity was correlated with reward probability at 1.5 s (Spearman's rank correlation; $\rho = 0.41$, $p = 0.0012$; analysis window, 0 s to 0.5 s). Unlike the tonic ramping neurons, the phasic bursting neurons did not anticipate the late reward at 4.5 s (Figure S3).

BF Phasic and Ramping Neurons Signal Reinforcement Surprise in Distinct Manners

A long-standing question is whether the BF signals errors in state values, or RPEs—a key signal for updating reward values and mediating economic choice [30, 31]. An alternative is that BF neurons signal a rectified (unsigned) prediction error [32, 33] rather than a value (signed) prediction error, which is better suited to control attention and mediate memory of salient events. We tested which type of prediction error is signaled by the BF by analyzing responses to reward deliveries and reward omissions after 25%, 50%, and 75% predictions (Figure 2).

Ramping neurons' outcome-related activity on average was correlated with unsigned prediction errors (Figures 2A and 2B). After the trial outcome, the magnitude of their activity was greatest during reward-delivered trials following 25% reward predictions and greatest during reward-omission trials following 75% reward predictions. Reward-omission and reward-delivery outcome responses were significantly correlated with expectancy (Figure 2B), albeit in opposite manners.

Phasic bursting neurons' outcome-related activity also signaled prediction errors following reward deliveries. Their delivery responses were correlated with expectancy (Figure 2D, red), displaying highest activations following reward deliveries in 25% reward trials. However, unlike the ramping neurons, these neurons did not discriminate reward omissions following different uncertain reward predictions (Figures 2C and 2D). To verify that the lack of relationship between reward-omission-related activity and reward probability was not due to firing

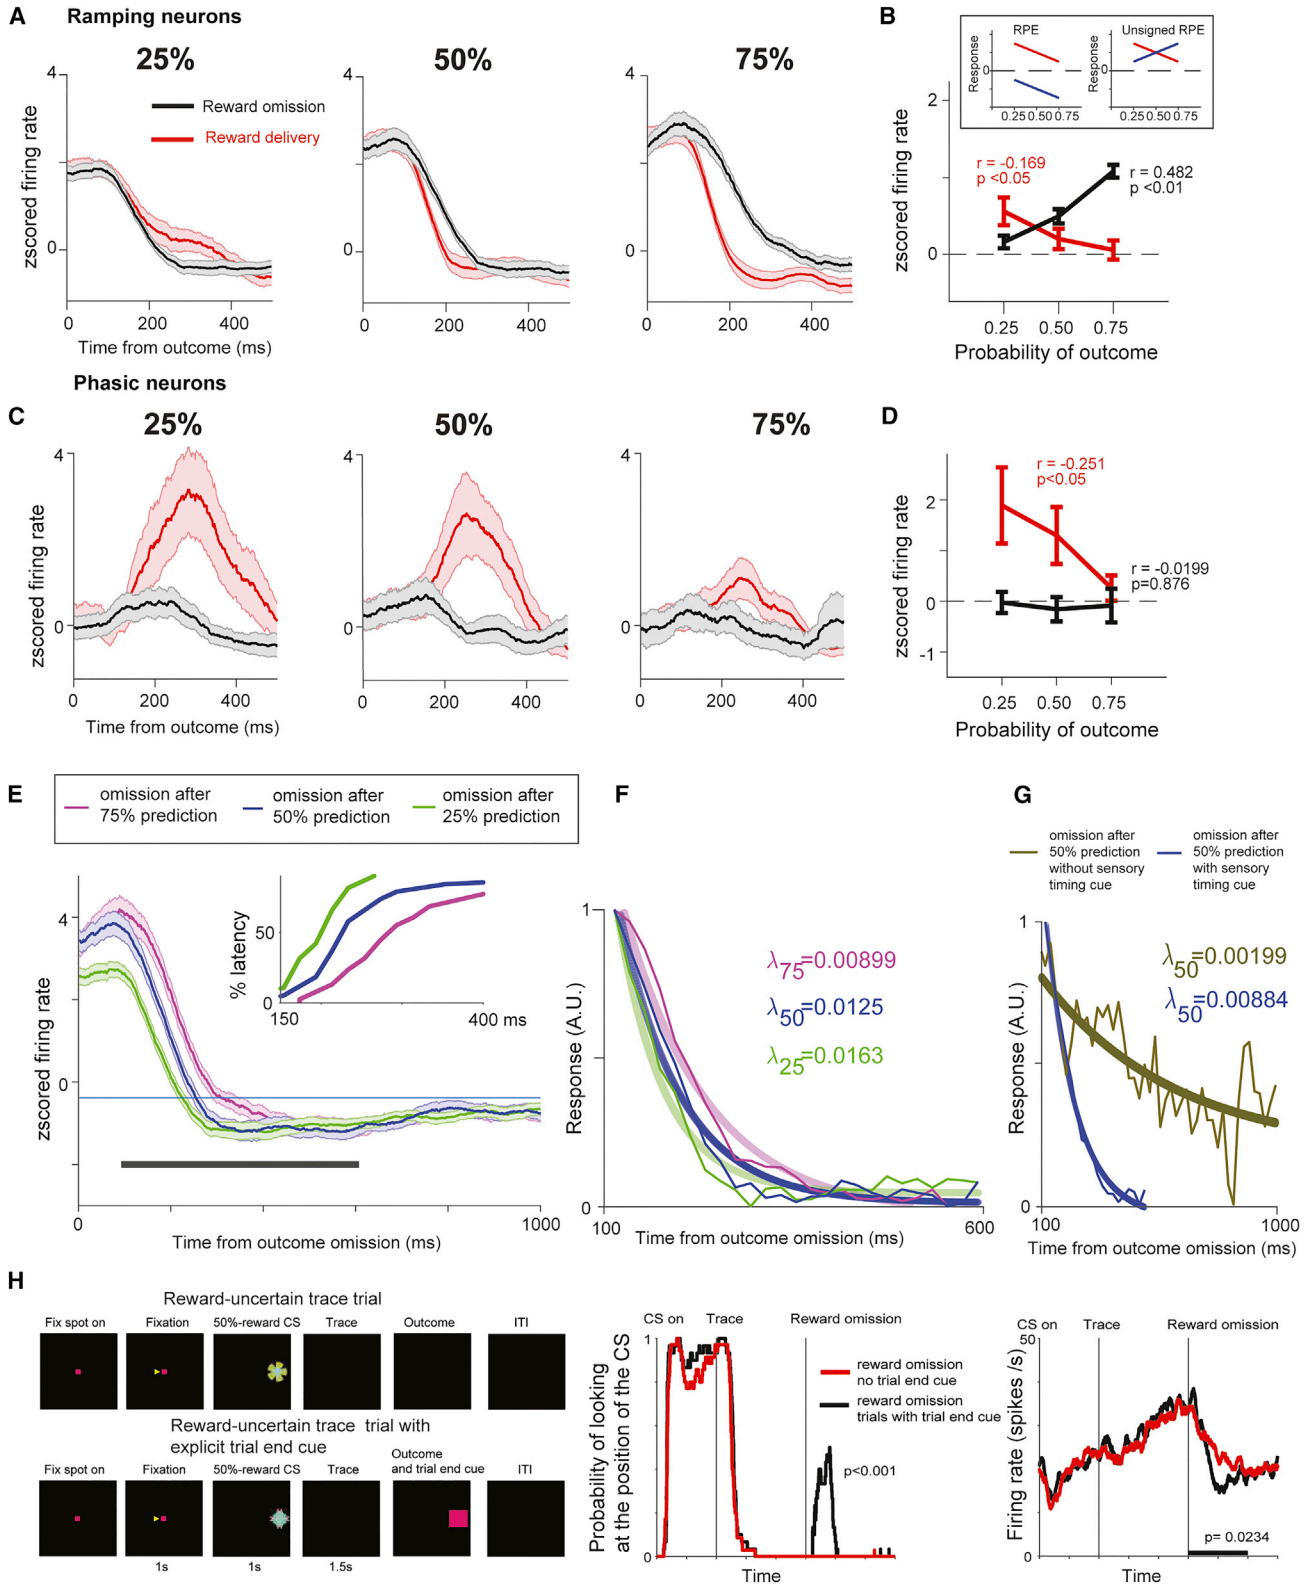

Figure 2. Differential Coding of Surprise in Ramping and Bursting BF Neurons

(A) Ramping neurons' average outcome activity in 25%, 50%, and 75% conditions. Red shows reward-delivered trials; black shows no-reward trials.

(B) Ramping neurons' average responses for reward-delivery and no-reward trials. Linear correlations of responses with reward expectancy are indicated (time window: 100 ms to 400 ms; p values were obtained with 10,000 permutations; STAR Methods). The results of the correlations suggest that the activity resembles

(legend continued on next page)

rate normalization, we repeated the correlation analyses in Figure 2D on raw omission-related spike counts and observed the same results ($p = 0.89$). Hence, a key feature of the value RPE—a reward-omission-related suppression—was missing from the phasic bursting neurons.

BF bursting can be elicited by rewarding and aversive, noxious events [4, 8, 25]. Therefore, why was bursting not apparent in response to unexpected reward omissions? The most parsimonious explanation is that the lack of omission responses was due to a lack of external salient events cueing reward omissions and a lack of sensitivity of phasic bursting neurons to internally generated errors in subjective value.

Surprise has a temporal dimension, and ramping neurons clearly display ramping signals to the timing of uncertain or salient reinforcements (Figures 1 and 2) [4], the magnitude of which is correlated with monkeys' confidence in reward delivery (Figures 1 and S3; ramping responses: $0.25 < 0.5 < 0.75$). Might BF ramping activity encode estimates of outcome timing under uncertainty?

To test this, we took advantage of the fact that in our tasks, CSs co-terminated with outcomes. During omission trials, no external cues indicated that the reward was omitted. If ramping reflects information about the animals' internal temporal estimates, then we should have seen different ramping-down responses following omissions in 25%, 50%, and 75% trials.

BF ramping returned to baseline earliest during 25% reward trials and latest during 75% trials (Figure 2E). Decay of the ramping also roughly scaled with reward expectation: it was greatest following omissions during 25% and least during 75% trials (Figure 2F; bootstrapping; the 95% confidence intervals of 25%, 50%, and 75% decay rates exclude each other). Note that different firing rates across different trial types could not explain these results because before obtaining the decay rates, we first normalized each trial type from 0 to 1.

Next, we studied the activity of BF ramping neurons in the reward-timing procedure because it contained two distinct 50% reward predictions: one in which the CS co-terminated with the outcome and one in which the CS remained on the screen (Figure S3). In the first condition, the animals obtained a signal about the timing of the trial, while in the other, they did not. The decay rate of BF ramping neurons was again sensitive to temporal predictions: it was greater when the animals did not receive an explicit temporal cue (Figure 2G; bootstrapping; the 95% confidence intervals of the decay rates exclude each other). Finally, we analyzed another task that contained two types of 50% reward CS trials with identical timing and reward statistics. The two trials differed in one way—one of them contained an external trial-end cue that indicated when the trial was over. Consistent with the results of Figure 2G, when an explicit cue was given during reward-omission trials, the ramping-down activity displayed a relatively rapid drop-off (Figure 2H, right). In sum, Figures 2E–2H show that BF ramping activity is strongly influenced by evidence about and confidence in the timing of reinforcements.

Object Novelty and Sensory Surprise Are Signaled by the BF

The data thus far show that BF neurons are sensitive to surprise. However, surprises arise due to violations in belief states following a probabilistic prediction, when there is a deviation of the outcome from the mean of expected-outcomes [20], or as a result of novelty due to a comparison of a sensory events with representations of past experiences. To test how the BF represents novelty, we designed an object-sequence task in which novel objects were fully expected.

Monkeys experienced four sequences of object presentations (S1, S2, S3, and S4). Each sequence contained 3 familiar objects and 1 novel object. The novel object was always in the second position in the sequence. If a neuron has a selective novelty response, it should respond more strongly and consistently to the novel object than to the familiar objects in the sequence. To assess whether novelty responses were dominantly due to task relevance or reward prediction, following S2 and S4, monkeys performed a reaction-time delayed non-matching-to-sample (DNMS) task (Figure 3A, right). During the DNMS task, an object that was novel during the presentation of S2 (or S4 if the DNMS trial followed S4) was presented along with a novel object that had never been experienced. The trial continued until the monkeys fixated this novel object for 0.5 s to get a reward (Figure 3A, right). The monkeys' behaviors indicated that they understood the task and utilized previous experiences to increase their reward

the toy model of unsigned RPEs (or surprise). The inset shows cartoon models of theoretical outcome responses coding RPEs (left) and unsigned RPEs (right). If neurons signal unsigned RPEs, then they should display greatest responses to reward deliveries following 25% reward predictions and smallest responses following a 75% reward prediction. The same neurons should display greatest responses to reward omissions following 75% reward predictions and smallest responses following 25% reward predictions. Alternatively, if neurons encode signed RPEs, then they will display inhibitions following omissions whose magnitudes ought to be inversely related to the probability of a reward.

(C) Outcome activity of phasic bursting neurons. Conventions are the same as in (A).

(D) Phasic bursting neurons' responses resembled RPE coding only in reward-delivery trials (red; time window: 200 ms to 500 ms).

(E) Activity of BF ramping neurons during 25%, 50%, and 75% reward-probability trials in which the reward was omitted. The ramping activity returned to the inter-trial baseline level (thin blue line) at different latencies across these three types of trials: earliest during 25% trials and latest during 75% trials. Cumulative distributions of these latencies are shown in the inset. The black bar below the activity indicates the time window for the analyses in (F).

(F) Exponential fits (thick lines) to the population's binned activity (thin lines; STAR Methods). Fits and decay rates (right) were calculated for the population after the activity for each trial type was normalized from 0 to 1, such that for each of the three conditions, the starting point is 1. A.U., arbitrary units.

(G) Same as (F), except here we compared the fit and decay rate during 50% trials in which an explicit cue indicated the end of the trial (dark blue) with the fit and decay rate during 50% trials in which no explicit cue was given (and the CS remained on the screen; Methods).

(H) Left: trace conditioning with and without explicit visual cues that signaled the end of the trial. Middle: the monkey's gaze behavior indicated that it attended to the trial-end cue (presented at the same location as the CS; rank-sum test; $p < 0.001$). Right: explicit knowledge of trial timing reduced the reward-omission-related ramping activity (monkey W; 8 neurons; $p = 0.0234$; signed-rank test). The analysis window used to study gaze behavior and neuronal activity is indicated by the black bar.

The shaded regions throughout this figure represent the SEM. See also Figure S3 for activity in the temporal-uncertainty procedure separately.

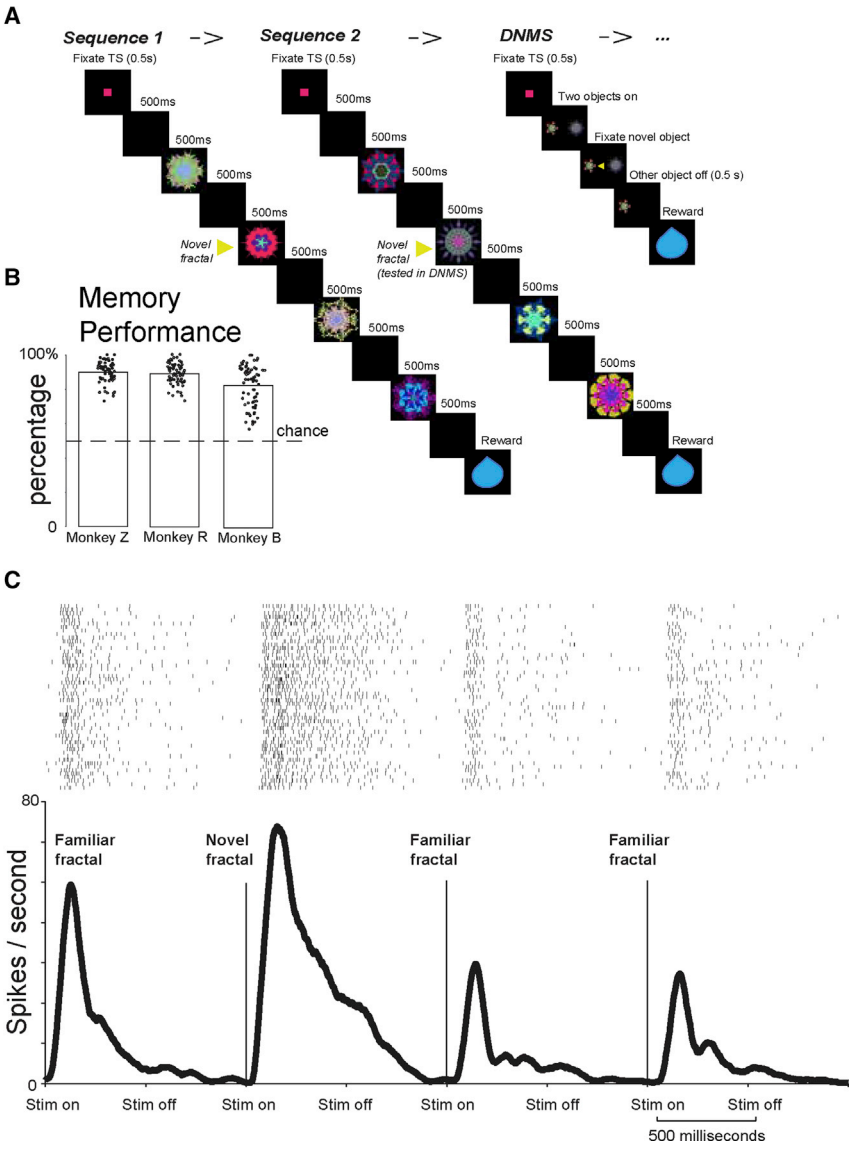

Figure 3. Object-Sequence Task

(A) The monkey was first shown sequences of fractals. Each sequence contained 4 fractals, in which the 1st, 3rd, and 4th fractals were fixed familiar objects and the 2nd fractal was always novel. After the two sequences, the monkeys performed a DNMS task in which one object was novel and the other was the object that was previously novel in sequence 2. Monkeys fixated the novel object for reward.

(B) Behavioral performance for three monkeys. y axis shows the percentage of first saccades to the novel object in DNMS. The percentages are significantly different from 0.5 for all three monkeys ($p < 0.01$; signed-rank test).

(C) Example BF phasic bursting neuron's responses to the four objects in a sequence. The response was highest for the second (novel) fractal (rank-sum test; $p < 0.05$).

rate. Their first saccade following the presentation of the two fractals most often landed on the novel object, where their gazes remained until the non-selected stimulus disappeared and a reward was delivered.

We studied 39 BF neurons identified using experiment 1 (monkey B = 6, monkey R = 11, and monkey Z = 22). Phasic bursting neurons robustly discriminated the novel object from the familiar objects. An example phasic bursting neuron is shown in Figure 3C. This neuron responded selectively to the novel object ($p < 0.01$; rank-sum test). This selective response could not be explained by priming or reward proximity because the novel objects always appeared in the second position in the sequence (Figure 3A) rather than the first or the last. Like the example neuron, the population of phasic bursting neurons (Figure 4A) and the single neurons (Figure 4B) selectively discriminated the novel object versus familiar objects.

Phasic bursting neurons' strong and selective novelty responses in the object-sequence task were present when the novel object was relevant or irrelevant for subsequent memory behaviors (Figure 4C). That is, during both S1 and S3, BF phasic neurons displayed stronger responses to novel objects than to familiar objects (signed-rank tests; $p < 0.01$). Their novelty responses were also consistently enhanced by task relevance (Figures 4C and 4D).

An important consideration for the interpretation of novelty responses is that novelty, in primates, is thought to exert a strong influence on behavior [20, 34–36], especially on gaze behavior. However, the type of influence (attentional, motivational, or both) that is exerted has been unclear. We designed a novel behavioral procedure that revealed that object novelty indeed has a motivational value (Figure S4). This finding necessitates that future studies assess the role of BF activity in mediating the motivational effects of object novelty on behavior.

We previously showed that CSs predicting uncertain (surprising) rewards attract overt attention more than CSs predicting certain rewards [37]. Here and in a previous report [4], we showed that BF ramping neurons anticipate uncertain reward delivery (Figures S1 and S3). So, might these neurons also anticipate other attention-capturing stimuli such as novel objects? While in contrast to the phasic neurons, the ramping neurons had a weaker novelty-selective response (rank-sum test comparing single neurons' area under ROC curve values; $p = 0.035$), they indeed displayed ramping that anticipated at least two critical events in the object sequence task: the presentation of novel objects and rewards occurring after a long interval (Figure S4).

The temporal cortex, a major target of BF projections [1], is sensitive to sequence violations [38]. To test whether the BF is sensitive to unexpected violations in object sequences, we replaced an object in S2 with an object from S1 or an object from S4 with an object from S3 in ~11% of trials. These replacements avoided RPEs because the proximity to the reward was

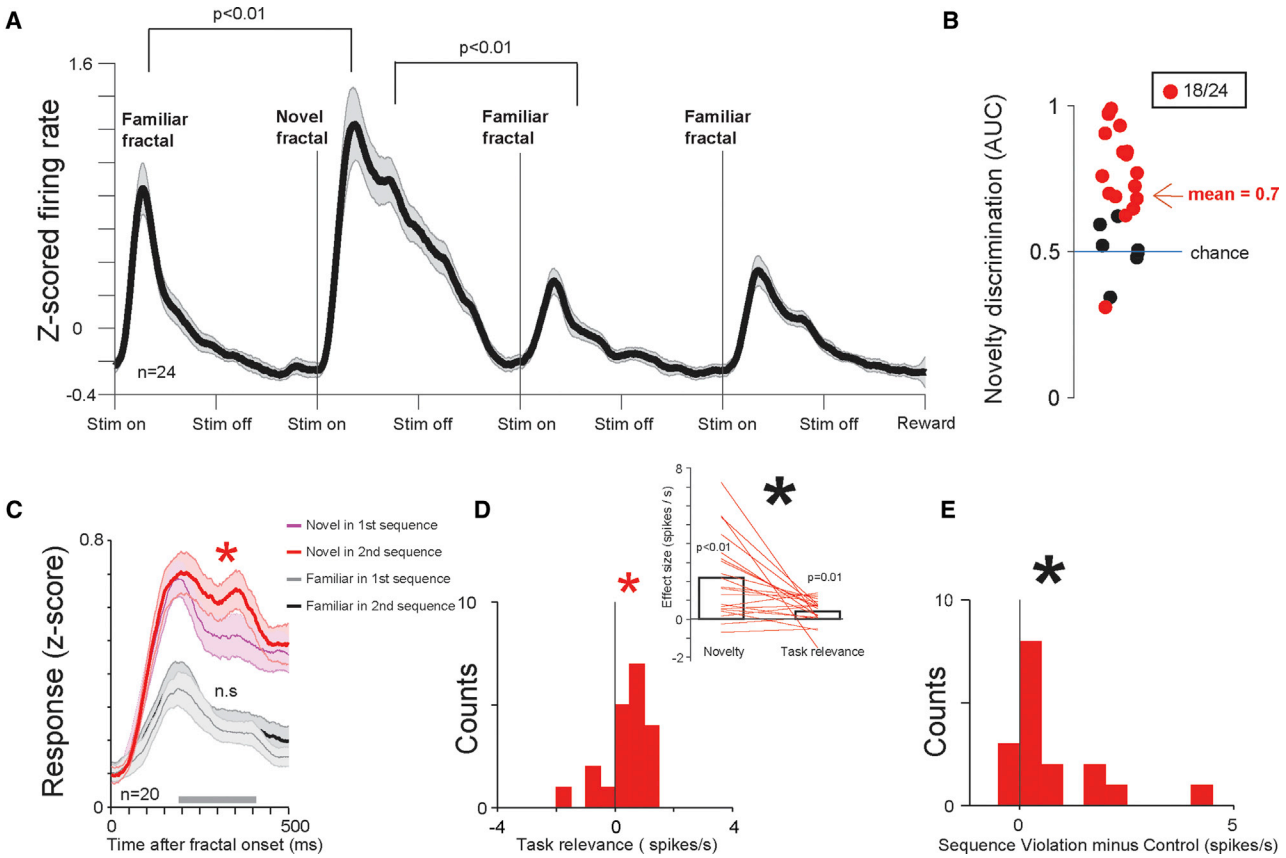

Figure 4. Phasic Bursting Neurons Signal Novelty and Surprise Not Directly Related to Reward

(A) Average activity of phasic bursting neurons in the object-sequence task. The shaded region represents the SEM.

(B) Area under the ROC curve (AUC) for each phasic neuron that assessed the ability of the neuron to discriminate novel versus familiar objects. Red dots are neurons that can significantly discriminate novel versus familiar objects (time window: 200 ms to 400 ms).

(C) Phasic neurons group average responses to novel fractals in sequence 1 (thin blue line), sequence 2 (thick red line), and then to the last 2 familiar fractals in sequence 1 (thin gray line) and sequence 2 (thick black line). The shaded region represents the SEM. The asterisk indicates significant difference ($p < 0.05$) between novel fractal responses in sequences 1 and 2. n.s., not significant.

(D) Lower left: histogram of single neurons' response differences for novel fractals in sequence 2 (or 4) and sequence 1 (or 3). The red asterisk indicates a significant difference from 0 ($p < 0.05$). Upper right: for each neuron, the data from the histogram (right) was compared with the strength of novelty discrimination (left). Both novelty-discrimination and task-relevance effects are significant, but the novelty effect is stronger ($p < 0.05$).

(E) At low probability (11%), one of the familiar fractals in sequence 2 (or 4) was substituted with another familiar fractal from sequence 1 (or 3) (STAR Methods). Phasic neurons' responses were enhanced ($p < 0.01$) by this object-sequence violation.

See also Figure S4 for the activity of tonic ramping neurons.

not changed. Sequence violations produced small but significant increases in the population responses of phasic and tonic neurons (Figures 4 and S4). So, the BF can broadcast information about novel and surprising sensory events that are not directly associated with primary reinforcements.

Conclusions

We report that the primate BF contains at least two types of neurons that process a diverse set of salient events in distinct manners—with phasic burst responses when they are occur and with ramping activity—in anticipation of their occurrence.

Ramping neurons signaled internal variables closely tied to confidence in the timing of surprises and novel events. Their activity may represent (or provide a readout of [39]) an internal clock that is well-suited to guide anticipatory temporal attention, particularly in uncertain or novel contexts. Phasic bursting neurons rapidly and precisely conveyed statistical information about the timing, magnitude, and probability of reinforcement predictions and about the surprise of reinforcement deliveries. They were highly sensitive to sensory novelty and to errors in the subjects' beliefs about the sequences of sensory events. These neurons' short latency bursting could rapidly coordinate many regions of the neocortex that receive BF projections to mediate the processing of a wide range of external salient events and orchestrate appropriate responses to them [8, 12, 29, 40, 41].

Phasically bursting neurons did not discriminate among expected and unexpected reinforcement omissions that monkeys had to detect internally (e.g., omissions were not cued). Thus, in contrast to many dopamine neurons, they did not convey phasic RPEs [30, 42, 43]. Notably, a set of recent studies showed that not all dopamine-phasic responses signal RPEs wholly or purely. Instead, some dopamine neurons convey an alerting

signal complementary to BF bursting [44–48]. Future studies must assess how the BF phasic bursting and dopamine neurons work together to mediate behavior. One possibility is that BF phasic bursting (conveyed to the neocortex in response to a salient event) is followed by the release of dopamine in the basal ganglia. This dopaminergic release would then support striatal value (or motivational-salience) assignments to events being processed by the cortex (under the mediation of the BF). How dopamine would do so may ultimately depend on when and where it is released [44–48].

The BF contains prominent groups of cholinergic, GABAergic, and glutamatergic projection neurons. Previous work in rodents has identified putative GABAergic CS-related phasic bursting neurons, reinforcement-salience-related bursting cholinergic neurons, and other tonically active neurons in the rodent BF [8, 11, 25]. It will now be particularly important to identify which neurotransmitters are released (or co-released) by phasic bursting and ramping neurons in primates.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- CONTACT FOR REAGENT OR RESOURCE SHARING
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
  - Probability Amount procedure
  - Temporal Uncertainty Procedure
  - Object Sequence Procedure
  - Reward and Novelty Motivated Gaze Task
- QUANTIFICATION AND STATISTICAL ANALYSIS

SUPPLEMENTAL INFORMATION

Supplemental Information includes four figures and can be found with this article online at https://doi.org/10.1016/j.cub.2018.11.012.

ACKNOWLEDGMENTS

This work was supported by the National Institute of Mental Health under award number R01MH110594, the Defense Advanced Research Projects Agency (DARPA) Biological Technologies Office (BTO) ElectRx program through the CMO grant/contract no. HR0011-16-2-0022, Edward Mallinckrodt, Jr. Foundation, and the McDonnell Center for Systems Neuroscience. We are grateful to Ms. Julia Pai, Dr. Noah Ledbetter, and Mr. J. Kael White for assisting in data acquisition, to Ms. Kim Kocher for fantastic animal care and training, and to Dr. Okihide Hikosaka and the National Eye Institute Intramural Research Program for supporting the recording experiments in monkeys H and P for a previous study [4]. We thank Ms. Julia Pai and Drs. Ethan Bromberg-Martin and Timothy Holy for helpful discussions and Dr. Hiroyuki Nakahara and Ms. Jamie Moffa for reading earlier versions of this manuscript.

AUTHOR CONTRIBUTIONS

I.E.M. designed research; K.Z., C.D.C., and I.E.M. performed research; K.Z. analyzed the data; K.Z. and I.E.M. wrote the paper.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: September 3, 2018
Revised: October 12, 2018
Accepted: November 2, 2018
Published: December 20, 2018

REFERENCES

1. Mesulam, M.M., Mufson, E.J., Levey, A.I., and Wainer, B.H. (1983). Cholinergic innervation of cortex by the basal forebrain: cytochemistry and cortical connections of the septal area, diagonal band nuclei, nucleus basalis (substantia innominata), and hypothalamus in the rhesus monkey. J. Comp. Neurol. *214*, 170–197.
2. Everitt, B.J., and Robbins, T.W. (1997). Central cholinergic systems and cognition. Annu. Rev. Psychol. *48*, 649–684.
3. Baxter, M.G., and Chiba, A.A. (1999). Cognitive functions of the basal forebrain. Curr. Opin. Neurobiol. *9*, 178–183.
4. Monosov, I.E., Leopold, D.A., and Hikosaka, O. (2015). Neurons in the primate medial basal forebrain signal combined information about reward uncertainty, value, and punishment anticipation. J. Neurosci. *35*, 7443–7459.
5. Zaborszky, L., Csordas, A., Mosca, K., Kim, J., Gielow, M.R., Vadasz, C., and Nadasdy, Z. (2013). Neurons in the basal forebrain project to the cortex in a complex topographic organization that reflects corticocortical connectivity patterns: an experimental study based on retrograde tracing and 3D reconstruction. Cereb. Cortex *1*, 118–137.
6. Turchi, J., Saunders, R.C., and Mishkin, M. (2005). Effects of cholinergic deafferentation of the rhinal cortex on visual recognition memory in monkeys. Proc. Natl. Acad. Sci. USA *102*, 2158–2161.
7. Lin, S.C., Brown, R.E., Hussain Shuler, M.G., Petersen, C.C., and Kepecs, A. (2015). Optogenetic dissection of the basal forebrain neuromodulatory control of cortical activation, plasticity, and cognition. J. Neurosci. *35*, 13896–13903.
8. Hangya, B., Ranade, S.P., Lorenc, M., and Kepecs, A. (2015). Central cholinergic neurons are rapidly recruited by reinforcement feedback. Cell *162*, 1155–1168.
9. Voytko, M.L. (1996). Cognitive functions of the basal forebrain cholinergic system in monkeys: memory or attention? Behav. Brain Res. *75*, 13–25.
10. Chudasama, Y., Dalley, J.W., Nathwani, F., Bouger, P., and Robbins, T.W. (2004). Cholinergic modulation of visual attention and working memory: dissociable effects of basal forebrain 192-IgG-saporin lesions and intraprefrontal infusions of scopolamine. Learn. Mem. *11*, 78–86.
11. Avila, I., and Lin, S.C. (2014). Distinct neuronal populations in the basal forebrain encode motivational salience and movement. Front. Behav. Neurosci. *8*, 421.
12. Raver, S.M., and Lin, S.-C. (2015). Basal forebrain motivational salience signal enhances cortical processing and decision speed. Front. Behav. Neurosci. *9*, 277.
13. Pinto, L., Goard, M.J., Estandian, D., Xu, M., Kwan, A.C., Lee, S.H., Harrison, T.C., Feng, G., and Dan, Y. (2013). Fast modulation of visual perception by basal forebrain cholinergic neurons. Nat. Neurosci. *16*, 1857–1863.
14. Peck, C.J., and Salzman, C.D. (2014). The amygdala and basal forebrain as a pathway for motivationally guided attention. J. Neurosci. *34*, 13757–13767.
15. Richardson, R.T., and DeLong, M.R. (1990). Context-dependent responses of primate nucleus basalis neurons in a go/no-go task. J. Neurosci. *10*, 2528–2540.
16. Wilson, F.A., and Rolls, E.T. (1990). Neuronal responses related to reinforcement in the primate basal forebrain. Brain Res. *509*, 213–231.
17. Wilson, F.A., and Ma, Y.Y. (2004). Reinforcement-related neurons in the primate basal forebrain respond to the learned significance of task events rather than to the hedonic attributes of reward. Brain Res. Cogn. Brain Res. *19*, 74–81.

18. Masuda, R., Fukuda, M., Ono, T., and Endo, S. (1997). Neuronal responses at the sight of objects in monkey basal forebrain subregions during operant visual tasks. Neurobiol. Learn. Mem. *67*, 181–196.
19. Fukuda, M., Masuda, R., Ono, T., and Tabuch, E. (1993). Responses of monkey basal forebrain neurons during visual discrimination task. In Progress in Brain Research, T.P. Hicks, S. Molotchnikoff, and T. Ono, eds. (Elsevier), pp. 359–369.
20. Barto, A., Mirolli, M., and Baldassarre, G. (2013). Novelty or surprise? Front. Psychol. *4*, 907.
21. Wallis, J.D., and Rich, E.L. (2011). Challenges of interpreting frontal neurons during value-based decision-making. Front. Neurosci. *5*, 124.
22. Wang, T., and Mitchell, C.J. (2011). Attention and relative novelty in human perceptual learning. J. Exp. Psychol. Anim. Behav. Process. *37*, 436–445.
23. Hayden, B.Y., Heilbronner, S.R., Pearson, J.M., and Platt, M.L. (2011). Surprise signals in anterior cingulate cortex: neuronal encoding of unsigned reward prediction errors driving adjustment in behavior. J. Neurosci. *31*, 4178–4187.
24. Preuschoff, K., 't Hart, B.M., and Einhäuser, W. (2011). Pupil dilation signals surprise: evidence for noradrenaline's role in decision making. Front. Neurosci. *5*, 115.
25. Lin, S.C., and Nicolelis, M.A. (2008). Neuronal ensemble bursting in the basal forebrain encodes salience irrespective of valence. Neuron *59*, 138–149.
26. Monosov, I.E., and Hikosaka, O. (2013). Selective and graded coding of reward uncertainty by neurons in the primate anterodorsal septal region. Nat. Neurosci. *16*, 756–762.
27. White, J.K., and Monosov, I.E. (2016). Neurons in the primate dorsal striatum signal the uncertainty of object-reward associations. Nat. Commun. *7*, 12735.
28. Daye, P.M., Monosov, I.E., Hikosaka, O., Leopold, D.A., and Optican, L.M. (2013). pyElectrode: an open-source tool using structural MRI for electrode positioning and neuron mapping. J. Neurosci. Methods *213*, 123–131.
29. Turchi, J., Chang, C., Frank, Q.Y., Russ, B.E., David, K.Y., Cortes, C.R., Monosov, I.E., Duyn, J.H., and Leopold, D.A. (2018). The basal forebrain regulates global resting-state fMRI fluctuations. Neuron *97*, 940–952.e944.
30. Lak, A., Stauffer, W.R., and Schultz, W. (2014). Dopamine prediction error responses integrate subjective value from different reward dimensions. Proc. Natl. Acad. Sci. USA *111*, 2343–2348.
31. Schultz, W. (2002). Getting formal with dopamine and reward. Neuron *36*, 241–263.
32. Roesch, M.R., Calu, D.J., Esber, G.R., and Schoenbaum, G. (2010). All that glitters ... dissociating attention and outcome expectancy from prediction errors signals. J. Neurophysiol. *104*, 587–595.
33. Pearce, J.M., and Hall, G. (1980). A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. Psychol. Rev. *87*, 532–552.
34. Berlyne, D.E. (1970). Novelty, complexity, and hedonic value. Percept. Psychophys. *8*, 279–286.
35. Foley, N.C., Jangraw, D.C., Peck, C., and Gottlieb, J. (2014). Novelty enhances visual salience independently of reward in the parietal lobe. J. Neurosci. *34*, 7947–7957.
36. Tiitinen, H., May, P., Reinikainen, K., and Näätänen, R. (1994). Attentive novelty detection in humans is governed by pre-attentive sensory memory. Nature *372*, 90–92.
37. Monosov, I.E. (2017). Anterior cingulate is a source of valence-specific information about value and uncertainty. Nat. Commun. *8*, 134.
38. Meyer, T., and Olson, C.R. (2011). Statistical learning of visual transitions in monkey inferotemporal cortex. Proc. Natl. Acad. Sci. *108*, 19401–19406.
39. Paton, J.J., and Buonomano, D.V. (2018). The neural basis of timing: distributed mechanisms for diverse functions. Neuron *98*, 687–705.
40. Liu, R., Crawford, J., Callahan, P.M., Terry, A.V., Jr., Constantinidis, C., and Blake, D.T. (2017). Intermittent stimulation of the nucleus basalis of Meynert improves working memory in adult monkeys. Curr. Biol. *27*, 2640–2646.e2644.
41. Shuler, M.G., and Bear, M.F. (2006). Reward timing in the primary visual cortex. Science *311*, 1606–1609.
42. Matsumoto, M., and Hikosaka, O. (2009). Two types of dopamine neuron distinctly convey positive and negative motivational signals. Nature *459*, 837–841.
43. 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.
44. Takahashi, Y.K., Langdon, A.J., Niv, Y., and Schoenbaum, G. (2016). Temporal specificity of reward prediction errors signaled by putative dopamine neurons in rat VTA depends on ventral striatum. Neuron *91*, 182–193.
45. Bromberg-Martin, E.S., Matsumoto, M., and Hikosaka, O. (2010). Dopamine in motivational control: rewarding, aversive, and alerting. Neuron *68*, 815–834.
46. Matsumoto, M., and Takada, M. (2013). Distinct representations of cognitive and motivational signals in midbrain dopamine neurons. Neuron *79*, 1011–1024.
47. Takahashi, Y.K., Batchelor, H.M., Liu, B., Khanna, A., Morales, M., and Schoenbaum, G. (2017). Dopamine neurons respond to errors in the prediction of sensory features of expected rewards. Neuron *95*, 1395–1405.e3.
48. Babayan, B.M., Uchida, N., and Gershman, S.J. (2018). Belief state representation in the dopamine system. Nat. Commun. *9*, 1891.
49. Apicella, P., Ravel, S., Deffains, M., and Legallet, E. (2011). The role of striatal tonically active neurons in reward prediction error signaling during instrumental task performance. J. Neurosci. *31*, 1507–1515.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental Models: Organisms/Strains		
Rhesus macaques	PrimGen, NIH Animal Center at Poolesville	Macaca mulatta
Software and Algorithms		
MATLAB	Mathworks	https://www.mathworks.com/

CONTACT FOR REAGENT OR RESOURCE SHARING

Further information and reasonable requests for resources, data, and code should be directed to and will be fulfilled by the Lead Contact, Dr. Ilya E. Monosov (ilya.monosov@gmail.com).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Six adult sexually mature male rhesus monkeys (monkeys B, R, Z, W, H, and P; ages: 7-10 years old) were used for recording experiments. All procedures conform to the Guide for the Care and Use of Laboratory Animals and were approved by the Institutional Animal Care and Use Committee at Washington University (monkeys B, R, W, and Z) and the National Eye Institute (monkeys P and H).

METHOD DETAILS

All monkeys underwent surgery under general anesthesia. For each monkey, a plastic head holder and recording chamber were fixed to the skull under general anesthesia and sterile conditions. Chambers were tilted laterally from midline by 35 degrees and aimed at the basal forebrain and anterior portion of striatum. After the monkeys recovered from surgery, they participated in behavioral and neurophysiological experiments.

While the monkeys participated in behavioral procedures we recorded single neurons in the basal forebrain. The recording sites were determined with 1 mm-spacing grid system and with the aid of MR images (3T) obtained along the direction of the recording chamber. This MRI-based estimation of neuron recording locations was aided by custom-built software. Single-unit recording was performed using glass-coated electrodes (Alpha Omega). During each recording session, an electrode was inserted into the brain through a stainless-steel guide tube and advanced by an oil-driven micromanipulator (MO-97A, Narishige). Signal acquisition (including amplification and filtering) was performed using Alpha Omega 44 kHz SNR system. Action potential waveforms were identified online by multiple time-amplitude windows with an additional template matching algorithm (Alpha-Omega).

Neuronal recordings were restricted to single well-isolated neurons in the basal forebrain that displayed task related ramping or phasic-bursting activity following the presentation of the task conditioned stimuli in the Probability Amount procedure (described below). The ventral pallidum (defined using anatomical criteria and previous electrophysiological criteria, such as high and irregular firing rate) was not part of this study. The locations of the BF recordings are detailed in Figure S2. Reconstruction procedures were detailed previously [28].

Eye position was obtained with an infrared video camera (Eyelink, SR Research). Behavioral events and visual stimuli were controlled by MATLAB (Mathworks, Natick, MA) with Psychophysics Toolbox extensions. Juice, used as reward, was delivered with a solenoid delivery reward system (CRIST Instruments). Juice-related licking was measured and quantified using previously described methods. Airpuffs were delivered through a narrow tube placed ~6-8cm from the monkey's face.

Probability Amount procedure

To study (1) neuronal representations of reward probability and amount, and (2) delivery-related responses following uncertain predictions, we trained monkeys on a Pavlovian conditioning procedure. Pavlovian conditioning was used to avoid fluctuations in reward rate across trials or fluctuations in outcome timing within single trials (related to action performance) which theoretically may affect outcome prediction error signals [49].

The Pavlovian conditioning procedure contained two blocks of trials: a reward-probability block and a reward-amount block. Each trial started with the presentation of a green trial-start cue at the center. The monkeys had to maintain fixation on this trial-start cue for 1 s; then the trial start cue disappeared and one of the CSs was presented pseudo randomly. After 2.5 s (for monkeys B, Z, and R) or 1.5 s (monkeys H and P), the CS disappeared, and juice (if scheduled for that trial) was delivered. The longer duration was introduced for monkey B, Z, and R to verify that the ramping activity in the BF reaches maximum at the time of the outcome across different CS durations. The reward-probability block contained five visual fractal object CSs associated with five probabilistic reward predictions

(0, 25, 50, 75 and 100% of 0.25 mL of juice). The reward-amount block contained five objects associated with certain reward predictions of varying reward amounts (0.25, 0.1875, 0.125, 0.065 and 0ml). Each block consisted of 20 trials (monkeys B, Z, and R) and 40 trials (monkeys P and H) with fixed proportions of trial types (each of the five CSs appears four times in each block or 8 times in each block, depending on block length). The expected values of the five CSs in the probability block matched the expected values of the five CSs in the amount block. This two-block design removed confounds introduced by risk seeking-related changes in subjective values of the CSs [26, 27].

Before neuronal recordings began, the monkeys' knowledge of the CSs was confirmed by a choice procedure that was detailed previously [4, 26]. Briefly, in separate experimental sessions, the monkeys' choice preference was tested for the CSs. Each trial started with the presentation of the trial-start cue at the center, and the monkeys had to fixate it. Then two CSs appeared 10 degrees to the left and right. The monkeys had to make a saccade to one of the two CSs within 5 s and fixate it for at least 750 ms. Then, the unchosen CS disappeared, and after a brief delay the outcome (associated with the chosen CS) was delivered, and the chosen CS disappeared. If the monkey failed to fixate one of the CSs, the trial was aborted and all stimuli disappeared. The trials were presented pseudo randomly, so that a block of 180 trials contained all possible combinations of the 10 CSs four times. To verify that the monkeys' knowledge is stable during recording, we also monitored licking behavior and confirmed that it, like the choices, scaled with the expected values of the probability CSs and amount CSs (two separate Spearman's correlations, threshold: $p < 0.05$). The CS epoch responses of the 31 neurons recorded in monkeys H and P were previously analyzed in [4].

Temporal Uncertainty Procedure

To assess how monkeys' BF neurons encoded uncertain predictions about reward timing, monkeys B, R, Z were trained on an additional Pavlovian procedure (Figure S3). Following a trial start cue fixation period (same as above), one of five CSs were presented. These CSs predicted either (1) a probabilistic delay before a reward with deterministic delivery (reward-timing-uncertain CSs); or (2) a deterministic delay before a reward with 0.5 probability of delivery (reward-probability CS). In trials with one of the four reward-timing-uncertain CSs, reward was always delivered either 1.5 s after CS onset or 4.5 s after CS onset. Depending on the reward-timing uncertain CS, the reward was delivered at 1.5 s with 0.25, 0.50, 0.75, or 1 probability. In trials with the reward-probability CS, reward was delivered with a delay of 1.5 s after CS onset with 0.50 probability. During, the 0.25, 0.50, and 0.75 CS trials, when reward was not delivered at 1.5 s, the CS remained on the screen until reward was delivered at 4.5 s. During the 0.50 reward probability CS, the CS turned off at the time of the outcome (when reward was either delivered or omitted). The inter-trial-interval ranged from 2 to 6.5 seconds.

Training was verified by monkeys' reward anticipatory licking behavior. The data suggested that they understood the meanings of the CSs and were highly sensitive to the timing and probability of reward (Figure S3B). First, during the four reward-timing-uncertain CSs, monkeys displayed increased licking behavior before 1.5 s, then a decrease in licking behavior after 1.5 s if reward was not delivered, then finally an increase in licking behavior to the time of reward at 4.5 s. During reward omissions, in 75% reward trials licking behavior remained higher than 25% and 50% trials, even 0.5 s after the reward was omitted at 1.5 s ($p < 0.01$, rank-sum test, time window 2 s to 2.5 s after the onset of fractal). Also, the mean magnitude of anticipatory licking behavior before possible reward delivery at 1.5 s across all trials increased with the probability of reward delivery at 1.5 s (Spearman's rank correlation, $\rho = 0.38$, $p = < 0.0001$; Figure S3). These behavioral results indicate that the magnitude and persistence of the monkeys' anticipatory behavior were strongly influenced by reward timing conveyed by the CSs.

Object Sequence Procedure

An object sequence task was used to study how BF neurons encode sensory predictions and object novelty. Monkeys B, R, and Z experienced four distinct sequences of object presentations (S1, S2, S3, S4). The object sequences began following a 0.5 s period of fixation on the trial start cue that appeared in the center of the screen. Each sequence contained 3 familiar objects and 1 novel object. These objects were presented in the center of the screen and occupied ~3 degrees visual angle. The novel object was always presented in second position in the sequence. Therefore, the novel object was surprising because it was never experienced by the monkeys, but its presentation did not deviate from the animals' expectations. Monkeys performed more than 10,000 trials before recordings began. Following sequences S2 and S4, the monkeys performed a reaction-time Delayed Non-matching-to-Sample task (DNMS). During DNMS, an object that was novel during the presentation of S2 (or S4 if the DNMS trial followed S4) was presented with a novel object that has never been experienced. The objects were presented 10 degrees from the center, to the left and the right of the fixation point. The trial continued until the monkeys fixated the novel object for 0.5 ms to get a reward. The monkeys were never penalized for looking at the previously experienced object. Therefore, the novel objects in S2 or S4 did not have an explicit reward association, but aided the monkey in subsequent DNMS trials. On ~11% of S2 or S4 presentations, the first or the third fractal was replaced by a corresponding fractal from sequences S1 and S3 (in S2 from S1; and in S4 from S3). For example, if the first fractal in S2 was replaced, the first fractal from S1 was always displayed instead. In this way, sequence violations did not alter the relationship of the individual fractals to the timing of reward delivery. We used the probability-amount procedure to identify phasically bursting BF neurons and uncertainty ramping neurons and studied them in the object sequence procedure. All phasic bursting neurons included in Figure 4 had greatest responses for 100% reward CSs.

Reward and Novelty Motivated Gaze Task

To test if monkeys are motivated by novelty we trained Monkeys R and Z on a novel saccadic task (Figure S4) that measured their eagerness to observe a novel visual object. First, a fixation dot appeared in the center of the screen. 0.5 s after the onset of the fixation dot, a visual object fractal appeared 10 degrees to the right or the left of the fixation dot. The monkey was required to continue fixating the dot in the center. After 0.35 s the fixation spot disappeared and the monkey was free to make saccades. Reward was always delivered 3 s after the fractal onset. Therefore, the monkeys' saccadic behavior after the fixation spot disappeared did not affect reward delivery. In this task, the monkeys experienced four different trial types. The first two types of trials contained a novel (type 1) or 1 of 2 familiar (type 2) visual fractal fractal objects. Two additional trial types (3-4) tested whether the monkeys were motivated by the possibility of viewing a novel fractal. In trial type 3, 1 of 2 distinct familiar fractal objects appeared. After fixation spot disappeared, if the monkey fixated the familiar object, it was immediately replaced by a novel object. In trial type 4, 1 of 2 other distinct familiar objects appeared. If the monkey fixated this object, it was replaced by 1 of 2 other familiar objects. If novelty is salient, we ought to observe faster target acquisition times (duration between the time when the stimulus was presented and when the monkey saccades to its location) in trial type 1 than 2. Also, if novelty exerts motivational effects on saccadic behavior, then we ought to see faster target acquisition times in trial type 3 than 4.

QUANTIFICATION AND STATISTICAL ANALYSIS

To generate spike density functions, spike times were convolved with a Gaussian kernel (σ = 100 ms). Statistical tests were two-tailed. All permutation tests used 10000 shuffles. For all analyses and figures that included deliveries and omissions of rewards, unless explicitly stated in the text, a neuron was included if it had at least 2 trials for reward delivery and omission.

To cluster the single neurons' average responses in the probability block (Figure 1), first we performed principal component analysis (PCA). We then applied Silhouette and Calinski-Harabasz tests to confirm the optimal number of clusters (n = 2). K-means clustering was used to cluster the data based on PCs into 2 clusters (for this, using the first 3 PCs and up to 10 PCs resulted in very similar group membership).

To calculate the latency of reward size coding information (Figure 1C) we performed a correlation of firing rate and value in time (in 100 ms bins moving 1 ms steps) for each neuron. For each time bin we calculated the p value of the Spearman's rank correlation of neuron's activity with reward amount in the reward amount block. Reward size coding latency was defined as the first time p was lower than 0.01 (but similar results were obtained at $p < 0.05$). These statistical-latency analyses do not determine the actual latency of information coding per se because they utilize an arbitrary threshold. Instead, they are useful for demonstrating relative latencies across two groups of neurons.

To calculate the baseline rate that was used to derive the latency with which ramping neurons returned to baseline (Figure 2E), we picked the time window from 1000 ms to 500 ms before trial start cue appeared and used the average firing rate in this time window as the baseline.

To fit the outcome related activity with exponential functions (Figure 2), we first derived spike density functions using overlapping bins of 50 ms (in 20 ms steps). Then we used a least-squares method to fit the data by the function: $A^{*}e^{-\lambda t}+C$, in which λ is the decay rate, representing how fast the firing rate decreases. λ is restrained by the interval (0,0.06). To determine if the decay rates were significantly different across the different reward-omission conditions we used bootstrapping to calculate the confidence interval of the difference between two decay rates and tested if the 95% confidence interval excluded a difference of zero. Bootstrapping was done by randomly resampling the neurons with replacement (500 times). Each time resampling was done, we obtained a set of decay rates by fitting the neurons' average activity to the function shown above. For Figures 2A–2D, data from probability-amount and reward timing procedures were pooled (see outcome responses separately in Figures S1 and S3).

In the DNMS object sequence task, reward was delivered as long as the monkey fixated on the novel object for 0.5 s, regardless if he had looked at the other object. To evaluate the monkey's performance, we focused on the primary choice the monkey made, i.e., the first object he fixated for 0.5 s. To calculate performance, we obtained the percentage of trials in which the monkeys' primary choices were the novel objects.

For single neuron analyses (Figure 4B–4E) of novelty, task-relevance, and sequence-violations in the object sequence task, we subtracted the activity 100 ms before the object presentation from the activity measured after the object was presented (the time window was 200 ms to 400 ms unless otherwise stated). In this way, changes in firing rate that were unrelated to the objects were not considered in the analyses.

Neuronal discrimination of object novelty was assessed by calculating area under the receiver operating characteristic (ROC) curve. ROC areas of 0 and 1 are equivalent statistically; both indicate that two distributions are completely separated. The analysis was structured so that ROC area values greater than 0.5 indicate that the activity during novel object presentation was greater than familiar.