Behavioral/Systems/Cognitive

Neural Correlates of Variations in Event Processing during Learning in Basolateral Amygdala

Matthew R. Roesch,[5,6] **Donna J. Calu,**[2] **Guillem R. Esber,**[4] **and Geoffrey Schoenbaum**[1,3]

[1]Department of Anatomy and Neurobiology, [2]Program in Neuroscience, and [3]Department of Psychiatry, University of Maryland School of Medicine, Baltimore, Maryland 21201, [4]Department of Psychological and Brain Sciences, Johns Hopkins University, Baltimore, Maryland 21209, and [5]Department of Psychology and [6]Program in Neuroscience and Cognitive Science, University of Maryland College Park, College Park, Maryland 20742

The discovery that dopamine neurons signal errors in reward prediction has demonstrated that concepts empirically derived from the study of animal behavior can be used to understand the neural implementation of reward learning. Yet the learning theory models linked to phasic dopamine activity treat attention to events such as cues and rewards as static quantities; other models, such as Pearce–Hall, propose that learning might be influenced by variations in processing of these events. A key feature of these accounts is that event processing is modulated by unsigned rather than signed reward prediction errors. Here we tested whether neural activity in rat basolateral amygdala conforms to this pattern by recording single units in a behavioral task in which rewards were unexpectedly delivered or omitted. We report that neural activity at the time of reward is providing an unsigned error signal with characteristics consistent with those postulated by these models. This neural signal increased immediately after a change in reward, and stronger firing was evident whether the value of the reward increased or decreased. Further, as predicted by these models, the change in firing developed over several trials as expectations for reward were repeatedly violated. This neural signal was correlated with faster orienting to predictive cues after changes in reward, and abolition of the signal by inactivation of basolateral amygdala disrupted this change in orienting and retarded learning in response to changes in reward. These results suggest that basolateral amygdala serves a critical function in attention for learning.

Introduction

A common assumption of associative learning theories is that the prediction error computed on a trial—the difference between the reward expected and that which actually occurs—serves to adjust future predictions and ultimately guide behavior. But whereas some models (Rescorla and Wagner, 1972; Sutton and Barto, 1998) use prediction errors to drive learning directly, others (Mackintosh, 1975; Pearce and Hall, 1980) use them to drive variations in event processing (attention) which can then influence the rate of learning (Schultz and Dickinson, 2000). For models in the former class, the sign of the error is relevant because it specifies the direction of the change in associative strength. For instance, according to Rescorla and Wagner (1972), the error is positive when reward is better than expected and, accordingly, cues should gain positive associative strength. Conversely, the error is negative when reward is worse than expected, and cues should acquire negative associative strength. In contrast, models in the latter class focus on the unsigned or absolute value of the error to modulate the processing of events in the environment—

and consequently how much is learned about them. For example, according to the original Pearce–Hall (1980) model, a cue should be most thoroughly processed (and learned about) when it is a poor predictor of reward. As the cue becomes a more reliable predictor (i.e., the error associated with it approaches zero), processing (and learning) should decline.

Models based on signed errors (Rescorla and Wagner, 1972; Sutton and Barto, 1998) have taken on special significance with the discovery that activity in midbrain dopamine neurons correlates with these signed errors in relatively simple learning paradigms. Thus firing in these neurons increases in the face of unexpected reward (positive error) and is suppressed when reward is unexpectedly omitted (negative error). The wealth of evidence from various labs confirming and expanding these results (Mirenowicz and Schultz, 1994; Montague et al., 1996; Schultz et al., 1997; Hollerman and Schultz, 1998; Waelti et al., 2001; Bayer and Glimcher, 2005; Pan et al., 2005; Bayer et al., 2007; D'Ardenne et al., 2008; Matsumoto and Hikosaka, 2009) and identifying similar correlates in other brain areas (Hong and Hikosaka, 2008; Matsumoto and Hikosaka, 2009) contrasts with the lack of evidence for neural correlates of unsigned prediction errors—e.g., increased firing when reward is either better or worse than expected. Here we report such correlates in firing at the time of reward in the basolateral amygdala (ABL); reward-selective neurons fired more strongly immediately after a change in reward than later after learning, and stronger firing was evident whether the value of the reward increased or decreased. As predicted by an amended version of the Pearce–Hall model (1982),

the change in firing developed over several trials as expectations for reward were repeatedly violated. Further, consistent with the hypothesis that this neural signal contributes to attention and learning, the change in firing was correlated with faster orienting behavior on subsequent trials, and inactivation of ABL disrupted this change in orienting and retarded learning in the task.

Materials and Methods

Subjects. Male Long–Evans rats were obtained at 175–200 g from Charles River Labs. Rats were tested at the University of Maryland School of Medicine in accordance with School of Medicine and National Institutes of Health guidelines.

Surgical procedures, inactivation, and histology. Surgical procedures followed guidelines for aseptic technique. Electrodes were manufactured and implanted as in prior recording experiments. Rats had a drivable bundle of ten 25 μm diameter FeNiCr wires (Stablohm 675, California Fine Wire) chronically implanted in the left hemisphere dorsal to either ABL ($n = 7$; 3.0 mm posterior to bregma, 5.0 mm laterally, and 7.5 mm ventral to the brain surface) or VTA ($n = 5$; 5.2 mm posterior to bregma, 0.7 mm laterally, and 7.0 mm ventral to the brain surface). Data from VTA have been previously reported (Roesch et al., 2007). Immediately before implantation, these wires were freshly cut with surgical scissors to extend ~1 mm beyond the cannula and electroplated with platinum (H_2PtCl_6, Aldrich) to an impedance of ~300 kΩ. Cephalexin (15 mg/kg, p.o.) was administered twice daily for 2 weeks postoperatively to prevent infection.

For validation of our procedures for identifying dopamine neurons in VTA, some rats also received sterilized SILASTIC catheters (Dow Corning) to allow intravenous apomorphine infusions using published procedures. An incision was made lateral to the midline to expose the jugular vein. The catheter was inserted into the right jugular vein and secured using sterile silk sutures. The catheter then passed subcutaneously to the top of the skull where it was connected to the modified 22 gauge cannula (Plastics One) head mount, which was anchored on the skull using 5 jeweler's screws and grip cement. A plastic blocker was placed over the open end of the cannula connector. Analgesic buprenorphine (0.1 mg/kg, s.c.) was administered postoperatively, and the catheters were flushed every 24–48 h with an antibiotic gentamicin/saline solution (0.1 ml at 0.08 mg/50 ml) to maintain patency of the catheter and to reduce chance of infection.

At the end of recording, the final electrode position was marked by passing a 15 μA current through each electrode. The rats were then perfused, and their brains were removed and processed for histology using standard techniques.

For inactivation, rats that had been previously trained on the task received guide cannula bilaterally 2 mm above ABL (AP -2.7, ML ±5.0, relative to bregma, and DV -6.5 relative to skull surface). The infusion cannulae were fitted with plastic dummies (Plastics One). After recovery and retraining, they were tested in a series of sessions in which they received infusions of NBQX, a competitive AMPA receptor antagonist. Eight rats received bilateral infusions of either PBS or the inactivating agent, NBQX (1,2,3,4-tetrahydro-6-nitro-2,3-dioxo-benzo[f]quinoxaline-7-sulfonamide disodium salt hydrate from Sigma Aldrich). The infusion procedure began with the removal of the plastic dummy followed by the insertion and securing of the infusion needle to the guide cannula. The infusion needle extended 2 mm below the guide cannula. Either 0.2 μl of NBQX (20 mg/ml dissolved in 0.1 M PBS) or PBS (0.1 M) alone was infused over 1 min using a 5 μl Hamilton syringe (Hamilton) controlled by a microinfusion pump. The injection needle was left in place for 2 min following infusion of the agent before the plastic dummies were replaced in the guide cannulae. Approximately 10 min after removal of the injection needle, the rats performed the choice task.

Dopamine cell identification in VTA recordings. Neurons were screened for wide waveform and amplitude characteristics (Roesch et al., 2007), and then tested with a nonspecific dopamine agonist, apomorphine (0.60–1.0 mg/kg i.v.). The apomorphine test consisted of ~30 min of baseline recording, apomorphine infusion, and ~30 min postinfusion

recording. Rats were not engaged in any task activities during the apomorphine test and were allowed to move freely in the recording chamber.

Behavioral task. Recording was conducted in aluminum chambers ~18″ on each side with sloping walls narrowing to an area of 12″ × 12″ at the bottom. A central odor port was located above and two adjacent fluid wells on a panel in the right wall of each chamber. Two lights were located above the panel. The odor port was connected to an air flow dilution olfactometer to allow the rapid delivery of olfactory cues. Task control was implemented via computer. Port entry and licking was monitored by disruption of photobeams. Odors were chosen from compounds obtained from International Flavors and Fragrances.

The basic design of a trial is illustrated in Figure 1. Trials were signaled by illumination of the panel lights inside the box. When these lights were on, nosepoke into the odor port resulted in delivery of the odor cue to a small hemicylinder located behind this opening. One of three different odors was delivered to the port on each trial, in a pseudorandom order. At odor offset, the rat had 3 s to make a response at one of the two fluid wells located below the port. One odor instructed the rat to go to the left to get reward, a second odor instructed the rat to go to the right to get reward, and a third odor indicated that the rat could obtain reward at either well. Odors were presented in a pseudorandom sequence such that the free-choice odor was presented on 7/20 trials and the left/right odors were presented in equal numbers (±1 over 250 trials). In addition, the same odor could be presented on no more than 3 consecutive trials.

Once the rats were shaped to perform this basic task, we introduced blocks in which we independently manipulated the size of the reward delivered at a given side or the length of the delay preceding reward delivery. Once the rats were able to maintain accurate responding through these manipulations, we began recording sessions. For recording, one well was randomly designated as short (500 ms) and the other long (1–7 s) at the start of the session (Fig. 1a, block 1). In the second block of trials these contingencies were switched (Fig. 1a, block 2). The length of the delay under long conditions abided the following algorithm. The side designated as long started off as 1 s and increased by 1 s every time that side was chosen until it became 3 s. If the rat continued to choose that side, the length of the delay increased by 1 s up to a maximum of 7 s. If the rat chose the side designated as long <8 out of the last 10 choice trials, then the delay was reduced by 1 s to a minimum of 3 s. The reward delay for long forced-choice trials was yoked to the delay in freechoice trials during these blocks. In later blocks, we held the delay preceding reward delivery constant (500 ms) while manipulating the size of the expected reward (Fig. 1a). The reward was a 0.05 ml bolus of 10% sucrose solution. For big reward, an additional bolus was delivered after 500 ms. At least 60 trials per block were collected for each neuron. Rats were mildly water deprived (~30 min of *ad libitum* water per day) with *ad libitum* access on weekends.

Single-unit recording. Procedures were the same as described previously (Roesch et al., 2007). Wires were screened for activity daily; if no activity was detected, the rat was removed, and the electrode assembly was advanced 40 or 80 μm. Otherwise active wires were selected to be recorded, a session was conducted, and the electrode was advanced at the end of the session. Neural activity was recorded using two identical Plexon Multichannel Acquisition Processor systems, interfaced with odor discrimination training chambers. Signals from the electrode wires were amplified 20× by an op-amp headstage (Plexon, HST/8o50-G20-GR), located on the electrode array. Immediately outside the training chamber, the signals were passed through a differential preamplifier (Plexon, PBX2/16sp-r-G50/16fp-G50), where the single unit signals were amplified 50× and filtered at 150–9000 Hz. The single unit signals were then sent to the Multichannel Acquisition Processor box, where they were further filtered at 250–8000 Hz, digitized at 40 kHz and amplified at 1–32×. Waveforms (>2.5:1 signal-to-noise ratio) were extracted from active channels and recorded to disk by an associated workstation with event timestamps from the behavior computer. Waveforms were not inverted before data analysis.

Data analysis. Units were sorted using Offline Sorter software from Plexon, using a template-matching algorithm. Sorted files were then processed in NeuroExplorer to extract unit timestamps and relevant event markers. These data were subsequently analyzed in Matlab. We focused
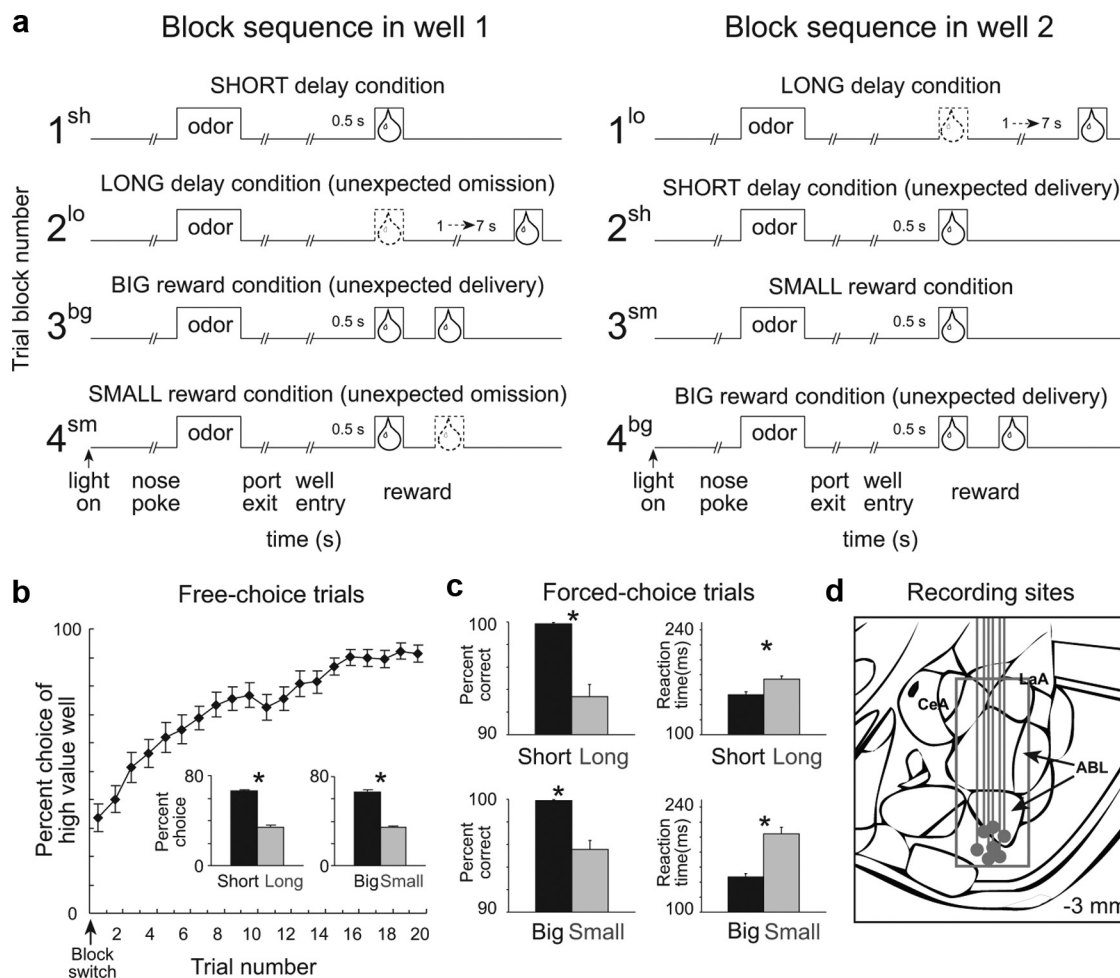
Figure 1. Task, behavior, and recording sites. *a*, Line deflections indicate the time course of stimuli (odors and rewards) presented to the animal on each trial. Other trial events are listed below. At the start of each recording session, one well was randomly designated as short (a 0.5 s delay before reward) and the other long (a 1–7 s delay before reward) (block 1). In the second block of trials these contingencies were switched (block 2). In blocks 3–4, we held the delay constant while manipulating the number of the rewards delivered. *b*, The impact of delay length and reward size choice behavior on free-choice trials. Bar graphs show average percentage choice for short versus long or big versus small across all free-choice trials. *c*, The impact of delay length and reward size on behavior on forced-choice trials. Bar graphs show percentage correct (left) and reaction time (right) across all recording sessions for different delays (top) and sizes (bottom). *d*, Location of recording sites. Gray dots represent final electrode position. Gray box marks extent of recording sites. Asterisks indicate planned comparisons revealing statistically significant differences (t test; $p < 0.05$). CeA, Central nucleus of amygdala; LaA, lateral amygdala; ABL, basolateral amygdala. Error bars indicate SEMs.

our analysis on neural activity from 0 to 1000 ms after reward delivery. Wilcoxon tests were used to measure significant shifts from zero in distribution plots ($p < 0.05$). *t* tests or ANOVAs were used to measure within cell differences in firing rate ($p < 0.05$). Pearson χ^2 tests ($p < 0.05$) were used to compare the proportions of neurons.

Modeling. The Pearce–Hall simulation was based on the extension of the original model (Pearce et al., 1982). The values for the remaining parameters were $\gamma = 0.6$ (Fig. 3) or 0.4, and 0.8 and $S = 0.05$. The Rescorla–Wagner simulation was based on the original proposal (Rescorla and Wagner, 1972). The product of α and β parameters was set as 0.05. All simulations were conducted across 20 training trials, and parameters were chosen to approximate the scale of the neural data; note the critical features of the shape of the curves are not dependent on these parameters.

Results

ABL neurons were recorded in a choice task, illustrated in Figure 1*a*. On each trial, rats responded to one of two adjacent wells after sampling an odor at a central port. Rats were trained to respond to three different odor cues: one that signaled reward in the right well (forced-choice), a second that signaled reward in the left well (forced-choice), and a third that signaled reward in either well

(free-choice). At the start of different blocks of trials, we manipulated the timing or size of the reward, thereby increasing (Fig. 1*a*, blocks 2 sh, 3 bg, and 4 bg) or decreasing (Fig. 1*a*, blocks 2 lo and 4 sm) its value unexpectedly. As illustrated in Figure 1, *b* and *c*, the rats changed their behavior in response to these value manipulations, choosing the more valuable reward significantly more often on free-choice trials and responding significantly faster and with greater accuracy for it on forced-choice trials (t test; df = 69; t values >6; p values <0.05). Thus the rats perceived the differently delayed and sized rewards as having different values and rapidly learned to change their behavior to reflect this within each trial block.

We recorded 284 ABL neurons across 107 sessions in 7 rats performing this task. Recording locations are illustrated in Figure 1*d*. Seventy of these neurons significantly increased firing after reward delivery (1 s) relative to baseline (1 s before trial start) across all trial types (t test; $p < 0.05$). Consistent with reports that ABL encodes value (Nishijo et al., 1988; Schoenbaum et al., 1998, 1999; Sugase-Miyamoto and Richmond, 2005; Paton et al., 2006; Belova et al., 2007, 2008; Tye et al., 2008; Fontanini et al., 2009),

a

b
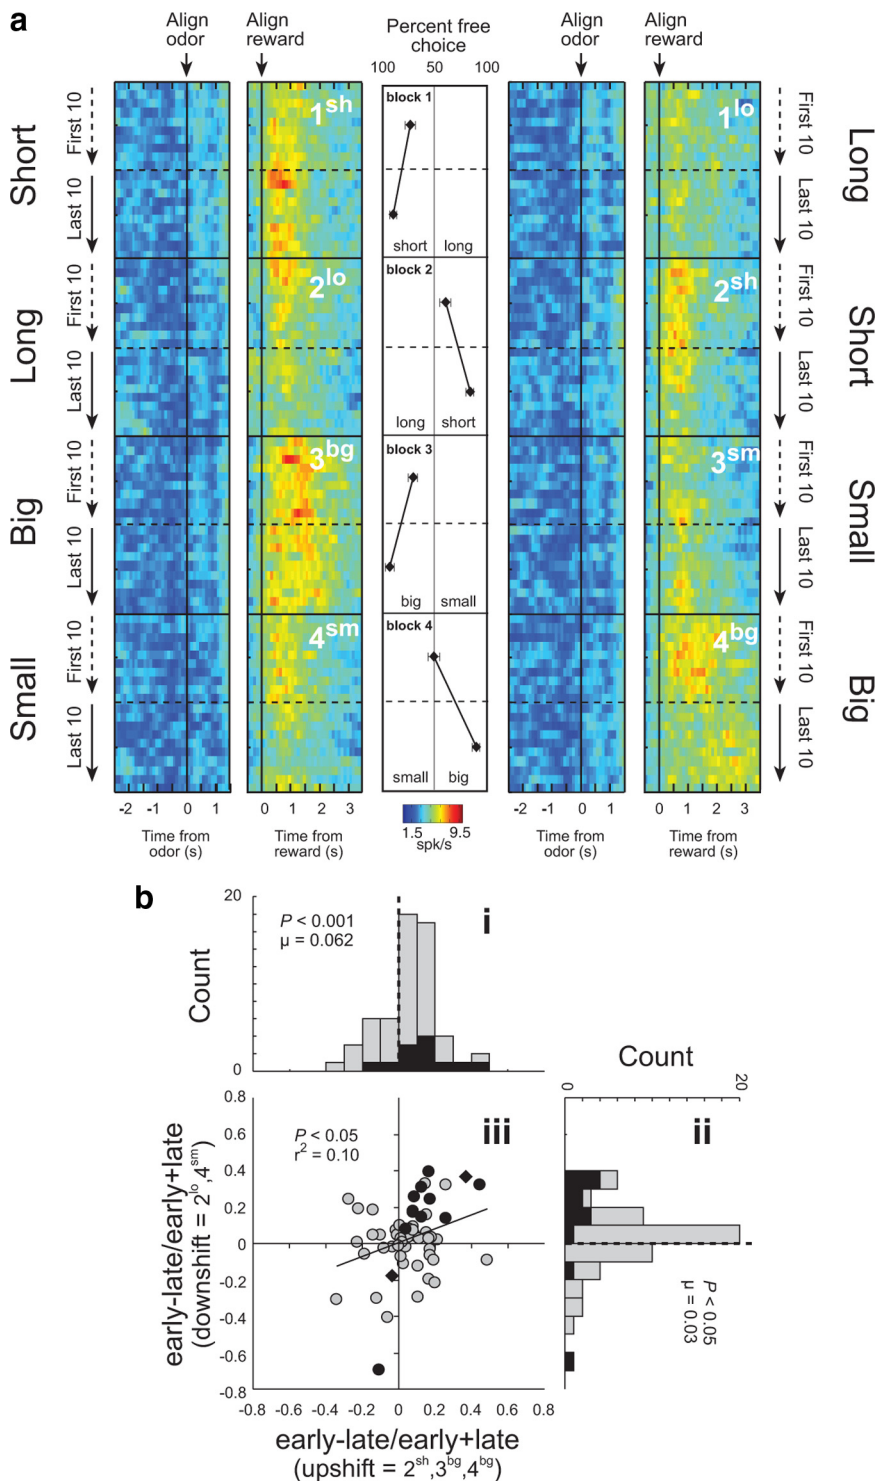
Figure 2. Neural activity in ABL is increased in response to unexpected reward delivery and omission. ***a***, Heat plots showing average activity over all ABL neurons (*n* = 58) that showed a significant effect of delay or size in an ANOVA during the 1 s after reward delivery. Activity over the course of the trials is plotted during the first and last 20 (10 per direction) trials in each training block (Fig. 1*a*; blocks 1– 4). Activity is shown, aligned on odor onset ("align odor") and reward delivery ("align reward"). Blocks 1– 4 are shown in the order performed (top to bottom). Thus, during block 1, rats responded after a "long" delay or a "short" delay to receive reward (actual starting direction—left/right—was counterbalanced in each block and is collapsed here). In block 2, the locations of the "short" delay and "long" delay were reversed. In blocks 3– 4, delays were held constant but the size of the reward ("big" or "small") varied. Line display between heat plots shows the rats' behavior on free-choice trials that were interleaved within the forced-choice trials. Value of 50% means that rats responded the same to both wells. ***b***, Distribution of indices [(early – late)/(early + late)] representing the difference in firing to reward delivery (1 s) and omission (1 s) during trials 3–10 (early) and during the last 10 trials (late) after upshifts (***i***) (2 ^{sh}, 3 ^{bg}, and 4 ^{bg}) and downshifts (***ii***) (2 ^{lo} and 4 ^{sm}). Note the first 2 trials were excluded in calculating the contrasts based on data in Figure 3*b*, showing that firing did not increase on these initial trials, consistent with predictions of the Pearce–Hall model. Filled bars in distribution plots (***i***, ***ii***) indicate the number of cells that showed

58 of the 70 reward-responsive neurons exhibited differential firing (*t* test; *p* < 0.05) based on either the timing or size of the reward after learning.

These 58 outcome-selective ABL neurons also exhibited changes in reward-related firing between the beginning and end of each block of trials. Differences were particularly apparent when reward was delivered unexpectedly (upshift), as at the start of blocks 2 ^{sh}, 4 ^{bg}, and 3 ^{bg}. Average activity in the outcome-selective ABL neurons, illustrated in Figure 2*a*, appears higher at the start of these blocks, when choice performance was poor, than at the end, when performance had improved, even though the actual reward being delivered was the same. At the same time, activity also appears higher in Figure 2*a* at the start of blocks 2 ^{lo} and 4 ^{sm}, when the value of the reward declined unexpectedly (downshift).

This impression was confirmed by a 2 factor ANOVA comparing activity at the time of reward and reward omission (1 s) in each neuron across learning (early vs late) and shift type (upshift vs downshift). Since these neurons were outcome-selective, many of them exhibited a significant effect of shift type. However, according to this analysis, 10 of the 58 neurons (17%) also fired significantly more early in a block, after a change in reward, than later, after learning. Only two neurons showed the opposite firing pattern, a proportion which was not different from chance (χ^2; *p* < 0.05).

Furthermore, only 1 of these 10 neurons exhibited a significant interaction between learning and shift type, and even that neuron fired significantly more immediately after either shift type (*t* test; *p* < 0.05). Thus the effect of learning on firing in these neurons occurred for both increases and decreases in expected reward. This is illustrated in Figure 2*b*, which shows the contrast in activity (early vs late) for each neuron, plotted separately for blocks involving upshifts and downshifts in reward. Both distributions were shifted significantly above zero, indicating higher firing early, after a change in reward, than later, after learning (Fig. 2*bi*,

←

a main effect (*p* < 0.05) of learning (early vs late). ***iii***, Correlation between contrast indices shown in ***i*** and ***ii***. Filled points in scatter plot (***iii***) indicate the number of cells that showed a main effect (*p* < 0.05) of learning (early vs late). Black diamonds indicate those neurons that also showed an interaction with shift type (upshift vs downshift). Analysis and figures shown here include both free- and forced-choice trials. *p* values for distributions are derived from Wilcoxon test.
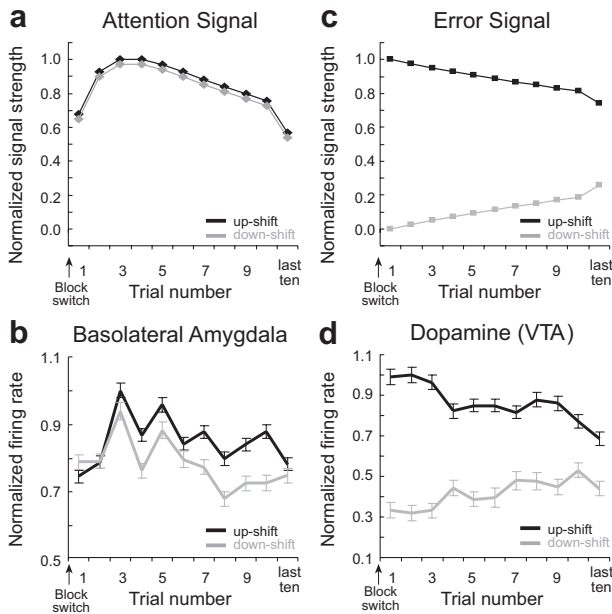
Figure 3. Neural activity in ABL increases gradually in response to unexpected reward and omission consistent with Pearce–Hall attention. *a*, *c*, Signals predicted by the Pearce–Hall (*a*) and Rescorla–Wagner (*c*) models after unexpected delivery (black) and omission (gray) of reward. *b*, *d*, Average firing (500 ms after reward delivery) in ABL (*b*) and for dopamine (*d*) neurons in VTA during the first 10 trials in blocks 2 ^{sh}, 3 ^{bg}, and 4 ^{bg} (upshifts = black) and in blocks 2 ^{lo} and 4 ^{sm} (downshifts = gray) normalized to the maximum. As reported previously (Roesch et al., 2007), dopamine neurons exhibit very short phasic changes in firing to unexpected reward and omission. The same time epoch was used to analyze both datasets here, to make the analysis equivalent. Asterisks indicate a significant difference between the first and third trial. See Results for ANOVA. Analysis and figures shown here include both free- and forced-choice trials. Error bars indicate SEMs.

Figure 4. Activity in ABL was correlated with odor-port orienting. *a*, Speed at which rats initiated trials after house light illumination during the first and last 10 trials in blocks 2– 4. These data were normalized to the maximum and inverted for ease of comparison with the model and neural data in Figure 3. *b*, Correlation between changes in firing in ABL on trial *n* and the orienting response on trial *n* + 2. Analysis and figures shown here include both free- and forced-choice trials. Error bars indicate SEMs.

Wilcoxon; $p < 0.005$; $u = 0.062$; Fig. 2*bii*, Wilcoxon; $p = 0.05$; $u = 0.030$). In addition, the two distributions did not differ (Wilcoxon; $p = 0.371$), and changes in firing to upshifts and downshifts in reward were positively correlated across the outcome-selective neurons ($p < 0.02$; $r^2 = 0.096$) (Fig. 2*biii*). Thus activity in the outcome-selective ABL neurons was higher at the start of a new training block, when reward was better or worse than expected, and declined as the rats learned to predict the value of reward. This pattern of firing is generally consistent with the notion of an unsigned error term found in stimulus processing models such as that of Pearce and Hall (1980).

Another distinctive feature of the activity in these ABL neurons is that firing did not immediately increase at the start of a new block, in response to a change in reward, but rather appeared to gather momentum and peak a few trials into the block. This is illustrated in Figure 3*b*. This gradual increase occurred in response to both upshifts and downshifts in reward value. Accordingly a 2 factor ANOVA of these data revealed main effects of trial ($F_{(19,58)} = 1.81$; $p < 0.05$) and shift type ($F_{(1,58)} = 5.6$; p values <0.05) but no significant interaction ($p = 0.98$). *Post hoc* comparisons showed that activity in ABL increased significantly after the first trial in response to upshifts and downshifts in reward, peaking on the third trial (p values <0.005), before returning to baseline. Although this property of the signal appears at odds with the original Pearce–Hall model (1980), this is precisely what is anticipated by the amended version of the model published two years later (Pearce et al., 1982). According to this version, an attentional signal should change gradually, increasing over several trials following a shift in reward contingencies, before declin-

ing back to baseline. This assumption was prompted by the observation that attentional changes seem somewhat anchored by previous levels of attention, essentially lagging prediction errors. A trial by trial simulation of the amended Pearce–Hall model, using the equations provided by the authors, is shown in Figure 3*a* to illustrate side-by-side theoretical changes in prediction error and attention at the outset of a new training block.

Given the remarkable fit provided by the amended Pearce–Hall model (1982) and the role attributed to unsigned errors within this theoretical context, it seems natural to speculate that this signal may be related to variations in event processing. This possibility is even more tantalizing in view of the striking similarity between changes in the ABL signal and changes in the rats' latency to approach the odor port at the start of each trial. Like the ABL signal, the speed to approach the odor port increased gradually over several trials following a block change (Fig. 4*a*), and firing in ABL was directly correlated with faster approach to the odor port on subsequent trials (Fig. 4*b*). Latency to approach the odor port precedes knowledge of the upcoming reward, thus this measure cannot reflect the value of the upcoming reward. However, faster odor-port approach latencies may reflect error-driven increases in the processing of trial events (e.g., cues and/or reward), as rats accelerate the reception of those events when shifted contingencies need to be worked out. In this sense, approaching the odor port faster can be looked upon as similar to conditioned orienting responses made to more traditional cues such as tones and lights—in this case directed at the source of the odors, the first event in the chain leading to reward. The parallel is significant because conditioned orienting responses, also known as investigatory reflexes, have been shown to recover from habituation when learned contingencies are shifted. Indeed, by and large experiments have confirmed that changes in the vigor of these responses mirror theoretical changes in Pearce–Hall attention (Kaye and Pearce, 1984; Pearce et al., 1988; Swan and Pearce, 1988), much as changes in odor port approach latencies do here.

To investigate the relationship between the ABL signal and odor-port approach latency, we conducted an additional experiment in which we inactivated ABL during performance of the recording task. Eight rats with bilateral guide cannulae targeting ABL were trained on the task used in recording. Thirty-eight sessions were conducted after bilateral infusions of either NBQX or vehicle. The order of infusions was counterbalanced such
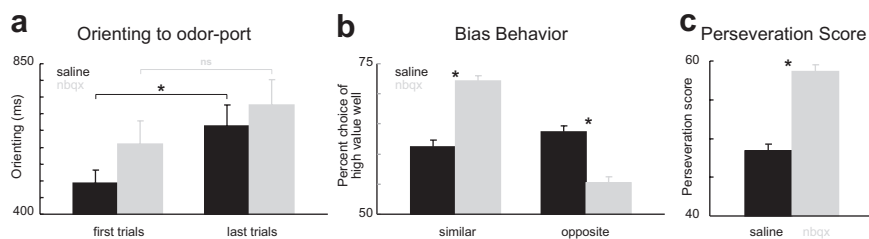
Figure 5. Inactivation of ABL disrupted odor-port orienting and impaired learning in the task. *a*, Speed at which rats oriented to the odor port after house-light illumination. Shown is the average latency in the first two trials after a change in reward versus the last two trials before a change in reward for vehicle versus NBQX sessions. For this comparison, the very first trial in a block was counted as the last trial before a change in reward, since orienting to the odor port on that trial preceded knowledge of the change in reward value. *b*, Choice performance in vehicle versus NBQX sessions, plotted according to whether the well values in a particular trial block were similar to or opposite from those learned at the end of the prior session. *c*, The height of the bars represents a perseveration score in vehicle versus NBQX sessions, computed by averaging choice performance in "similar" trial blocks with 100 minus choice performance in "opposite" blocks. See Results for statistics. Error bars indicate SEMs.

that each NBQX session had a corresponding vehicle session for comparison; rats also received reminder training between inactivation and vehicle sessions. ANOVA of orienting revealed a main effect of trial ($F_{(19,684)} = 3.79$, $p < 0.0001$) and a significant interaction with inactivation ($F_{(19,684)} = 1.74$, $p = 0.026$). Consistent with the correlation between neural activity and behavior shown in Figure 4*b*, rats responded significantly faster on the first trials after a change in reward during vehicle ($p = 0.012$) but not NBQX sessions ($p = 0.23$); there were no effects involving shift type (p values >0.05) (Fig. 5*a*).

Furthermore, consistent with the proposal that the ABL signal might influence learning via modulation of stimulus processing, the effect of ABL inactivation on orienting was accompanied by a subtle but significant learning impairment. Specifically, in the NBQX sessions, rats showed an increased bias to respond based on the value of the wells learned at the end of the prior day's session. Thus if the rats had completed the prior session with the high value reward on the right, they responded more to the right throughout the NBQX session. This effect is evident in Figure 5*b*, which plots the choice performance in each trial block in vehicle versus NBQX sessions according to whether the well values in the block were similar to or opposite from those learned at the end of the prior session. There is no difference in performance between "similar" and "opposite" blocks in vehicle sessions; however, in NBQX sessions, the rats performed significantly better when the location of the high value reward was similar to what had been learned in the prior session. ANOVA of these data showed a main effect of similarity ($F_{(1,37)} = 9.34$, $p < 0.004$) and a significant interaction between treatment and similarity ($F_{(1,37)} = 5.50$, $p < 0.05$). To further quantify this effect, we also computed a perseveration score by averaging choice performance in "similar" trial blocks with 100 minus choice performance in "opposite" blocks. As illustrated in Figure 5*c*, perseveration scores from NBQX sessions were significantly higher than those from vehicle sessions ($F_{(1,37)} = 5.50$, $p < 0.05$).

Discussion
In this paper, we have documented that firing to reward in ABL neurons generally conforms to the theoretical notion of an unsigned prediction error. In light of the amygdala's role in learning, and in particular in Pearce–Hall attentional phenomena, we have interpreted this signal as being related to variations in event processing. We have supported our argument by demonstrating the close relationship between the neural signal at the time of reward in ABL and a behavioral measure that likely reveals

(and certainly correlates with) error-induced variations in event processing.

The close correspondence between the neural signal and the criteria and predictions of the Pearce–Hall model is important because this model has proven highly effective at describing a number of important learning phenomena (Hall and Pearce, 1979; Kaye and Pearce, 1984; Swan and Pearce, 1988; Wilson et al., 1992) that lie outside the scope of theories based on simple reward prediction errors (Rescorla and Wagner, 1972; Sutton and Barto, 1998). The central theme running through these phenomena is that events will be better attended to, and hence learned about faster, when their consequences are surprising or unexpected. As a result of its success, Pearce–Hall model has been influential in guiding research into the neuroscience of attention, particularly in connection with the amygdala circuitry (Holland and Gallagher, 1993, 1999). However, there have been only a few reports of neural correlates consistent with the Pearce–Hall model. For example, several groups (Belova et al., 2007; Lin and Nicolelis, 2008; Matsumoto and Hikosaka, 2009) have argued for signaling of "motivational salience" in amygdala and a number of other areas based on similar firing to unexpected appetitive and aversive predictors or outcomes, and Tye et al. (2010) have recently reported that ~1 in 10 reward-responsive neurons in ABL exhibit increased firing during extinction, when the expected reward is omitted. This latter result is remarkably similar to the increased activity we find in outcome-selective neurons when a reward unexpectedly declines in value. Our results strengthen the case that this increase reflects a Pearce–Hall-like mechanism, since the same neurons also increased firing when reward value increases unexpectedly, and in both cases the change in firing occurred over several trials rather than all at once, as predicted by the amended version of the Pearce–Hall model (1982). Interestingly such a gradual change was not observed by Belova et al. (2007). One possible explanation for this discrepancy is the use of an aversive outcome in this study. The possibility of a punishment may have caused less weight to be placed on prior trials, resulting in a sharper, more immediate increase in firing in response to an unexpected outcome.

In this regard, it is worth noting that the signal reported here stands in contrast with that observed in midbrain dopamine neurons previously recorded in this same behavioral setting (Roesch et al., 2007). Data from this study were reanalyzed in Figure 3*d* for the sake of comparison. Activity in dopamine neurons was higher when reward was better than expected but showed suppression when reward was worse than expected. Moreover, changes in activity in these neurons were maximal the first time expectations were violated and then returned to baseline in the course of learning. A two-factor ANOVA of these data revealed a main effect of shift type and a significant interaction (p values <0.05). *Post hoc* comparisons showed that the change in activity away from baseline was never greater than on the first trial (p values >0.68). As illustrated by the trial-by-trial simulation in Figure 3*c*, this is precisely the pattern expected for a signed prediction error according to the Rescorla–Wagner (or TDRL) model.

The similar yet divergent signal carried by ABL and the midbrain dopamine neurons illustrates that signed and unsigned errors can be dissociated in the brain. This dissociation highlights

that these two kinds of signals are not mutually exclusive and may well operate in parallel, as suggested by recent hybrid theories which integrate the two (Dayan et al., 2000; Le Pelley, 2004). Support for this notion comes from behavioral work showing that damage within amygdala can affect learning dependent on attentional processes, while leaving unaffected learning that could be mediated by signed prediction errors (Holland and Gallagher, 1993). Such a formulation appears to most closely reflect the neurophysiology of reward systems in the brain.

To be sure, further work is needed to verify the connection between the neural signal in ABL and the theoretical framework in which we have chosen to couch it. One outstanding question is what specific trial events are influenced by the signal we have identified. Since changes in firing were confined to the time of reward delivery, an obvious possibility is that the signal may drive variations in the processing (i.e., salience, associability, etc) of the reward itself. Although this interpretation lies outside the scope of even the amended Pearce–Hall model, which is concerned with variations in the processing of predictors rather than outcomes, more recent evidence shows learning can also be driven by enhanced processing of an outcome when expectations about its occurrence are violated (Hall et al., 2005; Holland and Kenmuir, 2005).

At the same time, the signal we have identified may also provide an unsigned error to increase processing of cues, in keeping with the traditional Pearce–Hall view. Although we did not observe changes in cue-evoked activity in ABL at the time of shifts in reward, such effects could occur downstream. This would be consistent with reports of cue-evoked activity related to salience or attention in other areas (Rajkowski et al., 1994; Lin and Nicolelis, 2008; Matsumoto and Hikosaka, 2009). Alternatively, increased processing of the cues could be implemented as subthreshold changes in synaptic plasticity in neurons that receive information about antecedent cues. This mechanism would facilitate changes in cue-evoked firing that have been reported to occur with learning in amygdala (Nishijo et al., 1988; Schoenbaum et al., 1999; Paton et al., 2006; Tye et al., 2008), but would not necessarily change cue-evoked spiking activity immediately after a shift in reward. Of course, the development of cue-evoked activity with learning has been reported throughout the brain, including prominently in ABL and in many regions that receive input directly from ABL. Notably, in at least two of these regions, such associative encoding is disrupted by lesions or inactivation of ABL (Schoenbaum et al., 2003; Ambroggi et al., 2008).

Regardless of whether signaling of unsigned errors by ABL neurons is related to variations in processing of the predictors or the outcomes, these data suggest that the role of the ABL in a variety of learning processes may need to be reconceptualized or at least modified to include contributions to attentional processes identified in Pearce–Hall and related models. For example, ABL is clearly critical for encoding associative information properly (Davis, 2000; Gallagher, 2000; LeDoux, 2000; Murray, 2007), and associative encoding in downstream areas requires input from ABL (Schoenbaum et al., 2003; Ambroggi et al., 2008). This has been interpreted as reflecting an initial role in acquiring the simple associative representations (Pickens et al., 2003). However, an alternative account—not mutually exclusive—is that ABL may also augment the allocation of attentional resources to directly drive acquisition of the same associative information in other areas. This would support accounts of amygdala function that have emphasized vigilance (Davis and Whalen, 2001) and reports of neural activity reflecting uncertainty (i.e., prediction error) has been reported in ABL and in prefrontal regions that receive input

from ABL (Behrens et al., 2007; Herry et al., 2007; Kepecs et al., 2008).

References

Ambroggi F, Ishikawa A, Fields HL, Nicola SM (2008) Basolateral amygdala neurons facilitate reward-seeking behavior by exciting nucleus accumbens neurons. Neuron 59:648–661.

Bayer HM, Glimcher PW (2005) Midbrain dopamine neurons encode a quantitative reward prediction error signal. Neuron 47:129–141.

Bayer HM, Lau B, Glimcher PW (2007) Statistics of midbrain dopamine neuron spike trains in the awake primate. J Neurophysiol 98:1428–1439.

Behrens TE, Woolrich MW, Walton ME, Rushworth MF (2007) Learning the value of information in an uncertain world. Nat Neurosci 10:1214–1221.

Belova MA, Paton JJ, Morrison SE, Salzman CD (2007) Expectation modulates neural responses to pleasant and aversive stimuli in primate amygdala. Neuron 55:970–984.

Belova MA, Paton JJ, Salzman CD (2008) Moment-to-moment tracking of state value in the amygdala. J Neurosci 28:10023–10030.

D'Ardenne K, McClure SM, Nystrom LE, Cohen JD (2008) BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. Science 319:1264–1267.

Davis M (2000) The role of the amygdala in conditioned and unconditioned fear and anxiety. In: The amygdala: a functional analysis (Aggleton JP, ed), pp 213–287. Oxford: Oxford UP.

Davis M, Whalen PJ (2001) The amygdala: vigilance and emotion. Molecular Psychiatry 6:13–34.

Dayan P, Kakade S, Montague PR (2000) Learning and selective attention. Nat Neurosci 3:1218–1223.

Fontanini A, Grossman SE, Figueroa JA, Katz DB (2009) Distinct subtypes of basolateral amygdala neurons reflect palatability and reward. J Neurosci 29:2486–2495.

Gallagher M (2000) The amygdala and associative learning. In: The amygdala: a functional analysis (Aggleton JP, ed), pp 311–330. Oxford: Oxford UP.

Hall G, Pearce JM (1979) Latent inhibition of a CS during CS-US pairings. J Exp Psychol Anim Behav Process 5:31–42.

Hall G, Prados J, Sansa J (2005) Modulation of the effective salience of a stimulus by direct and associative activation of its representation. J Exp Psychol Anim Behav Process 31:267–276.

Herry C, Bach DR, Esposito F, Di Salle F, Perrig WJ, Scheffler K, Lüthi A, Seifritz E (2007) Processing of temporal unpredictability in human and animal amygdala. J Neurosci 27:5958–5966.

Holland PC, Gallagher M (1993) Amygdala central nucleus lesions disrupt increments, but not decrements, in conditioned stimulus processing. Behav Neurosci 107:246–253.

Holland PC, Gallagher M (1999) Amygdala circuitry in attentional and representational processes. Trends Cogn Sci 3:65–73.

Holland PC, Kenmuir C (2005) Variations in unconditioned stimulus processing in unblocking. J Exp Psychol Anim Behav Process 31:155–171.

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

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

Kaye H, Pearce JM (1984) The strength of the orienting response during Pavlovian conditioning. J Exp Psychol Anim Behav Process 10:90–109.

Kepecs A, Uchida N, Zariwala HA, Mainen ZF (2008) Neural correlates, computation and behavioural impact of decision confidence. Nature 455:227–231.

LeDoux JE (2000) The amygdala and emotion: a view through fear. In: The amygdala: a functional analysis (Aggleton JP, ed), pp 289–310. New York: Oxford UP.

Le Pelley ME (2004) The role of associative history in models of associative learning: a selective review and a hybrid model. Q J Exp Psychol B 57:193–243.

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

Mackintosh NJ (1975) A theory of attention: variations in the associability of stimuli with reinforcement. Psychol Rev 82:276–298.

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

Mirenowicz J, Schultz W (1994) Importance of unpredictability for reward responses in primate dopamine neurons. J Neurophysiol 72:1024–1027.

Montague PR, Dayan P, Sejnowski TJ (1996) A framework for mesencephalic dopamine systems based on predictive Hebbian learning. J Neurosci 16:1936–1947.

Murray EA (2007) The amygdala, reward and emotion. Trends Cogn Sci 11:489–497.

Nishijo H, Ono T, Nishino H (1988) Single neuron responses in alert monkey during complex sensory stimulation with affective significance. J Neurosci 8:3570–3583.

Pan W-X, Schmidt R, Wickens JR, Hyland BI (2005) Dopamine cells respond to predicted events during classical conditioning: evidence for eligibility traces in the reward-learning network. J Neurosci 25:6235–6242.

Paton JJ, Belova MA, Morrison SE, Salzman CD (2006) The primate amygdala represents the positive and negative value of visual stimuli during learning. Nature 439:865–870.

Pearce JM, Hall G (1980) A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. Psychol Rev 87:532–552.

Pearce JM, Kaye H, Hall G (1982) Predictive accuracy and stimulus associability: development of a model for Pavlovian learning. In: Quantitative analyses of behavior (Commons ML, Herrnstein RJ, Wagner AR, eds), pp 241–255. Cambridge, MA: Ballinger.

Pearce JM, Wilson PN, Kaye H (1988) The influence of predictive accuracy on serial conditioning in the rat. Q J Exp Psychol B 40:181–198.

Pickens CL, Saddoris MP, Setlow B, Gallagher M, Holland PC, Schoenbaum G (2003) Different roles for orbitofrontal cortex and basolateral amygdala in a reinforcer devaluation task. J Neurosci 23:11078–11084.

Rajkowski J, Kubiak P, Aston-Jones G (1994) Locus coeruleus activity in monkey: phasic and tonic changes are associated with altered vigilance. Brain Res Bull 35:607–616.

Rescorla RA, Wagner AR (1972) A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In: Classical conditioning II: current research and theory (Black AH, Prokasy WF, eds), pp 64–99. New York: Appleton-Century-Crofts.

Roesch MR, Calu DJ, Schoenbaum G (2007) Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards. Nat Neurosci 10:1615–1624.

Schoenbaum G, Chiba AA, Gallagher M (1998) Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. Nat Neurosci 1:155–159.

Schoenbaum G, Chiba AA, Gallagher M (1999) Neural encoding in orbitofrontal cortex and basolateral amygdala during olfactory discrimination learning. J Neurosci 19:1876–1884.

Schoenbaum G, Setlow B, Saddoris MP, Gallagher M (2003) Encoding predicted outcome and acquired value in orbitofrontal cortex during cue sampling depends upon input from basolateral amygdala. Neuron 39:855–867.

Schultz W, Dickinson A (2000) Neuronal coding of prediction errors. Annu Rev Neurosci 23:473–500.

Schultz W, Dayan P, Montague PR (1997) A neural substrate for prediction and reward. Science 275:1593–1599.

Sugase-Miyamoto Y, Richmond BJ (2005) Neuronal signals in the monkey basolateral amygdala during reward schedules. J Neurosci 25:11071–11083.

Sutton RS, Barto AG (1998) Reinforcement learning: an introduction. Cambridge MA: MIT.

Swan JA, Pearce JM (1988) The orienting response as an index of stimulus associability in rats. J Exp Psychol Anim Behav Process 14:292–301.

Tye KM, Stuber GD, de Ridder B, Bonci A, Janak PH (2008) Rapid strengthening of thalamo-amygdala synapses mediates cue-reward learning. Nature 453:1253–1257.

Tye KM, Cone JJ, Schairer WW, Janak PH (2010) Amygdala neural encoding of the absence of reward during extinction. J Neurosci 30:116–125.

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

Wilson PN, Boumphrey P, Pearce JM (1992) Restoration of the orienting response to a light by a change in its predictive accuracy. Q J Exp Psychol B 44:17–36.