

Published in final edited form as:

Neuron. 2007 September 20; 55(6): 970–984.

Expectation modulates neural responses to pleasant and aversive stimuli in primate amygdala

Marina A. Belova^{1,*}, Joseph J. Paton^{1,*}, Sara E. Morrison¹, and C. Daniel Salzman^{1,2,3,4,5,6}

¹ Dept. of Neuroscience, Columbia University

² Dept. of Psychiatry, Columbia University

³ W.M. Keck Center on Brain Plasticity and Cognition, Columbia University

⁴ Kavli Institute for Brain Sciences, Columbia University

⁵ Mahoney Center for Brain and Behavior, Columbia University

⁶ New York State Psychiatric Institute, 1051 Riverside Drive, Unit 87, New York, NY 10032

Abstract

Animals and humans learn to approach and acquire pleasant stimuli, and to avoid or defend against aversive ones. However, both pleasant and aversive stimuli can elicit arousal and attention, and their salience or intensity increases when they occur by surprise. Thus adaptive behavior may require that neural circuits compute both stimulus valence - or value - and intensity. To explore how these computations may be implemented, we examined neural responses in the primate amygdala to unexpected reinforcement during learning. Many amygdala neurons responded differently to reinforcement depending upon whether or not it was expected. In some neurons, this modulation occurred only for rewards or aversive stimuli, but not both. In other neurons, expectation similarly modulated responses to both rewards and punishments. These different neuronal populations may subserve two sorts of processes mediated by the amygdala: those activated by surprising reinforcements of both valences - such as enhanced arousal and attention - and those that are valence-specific, such as fear or reward-seeking behavior.

Food, water, mates, and predators do not always appear predictably in a natural environment. Instead, animals and humans often encounter these stimuli by surprise, triggering a range of cognitive, physiological and behavioral responses that contribute to an emotional experience. One influential framework describes emotions in a two-dimensional space, with arousal (ranging from calm to excited) and valence, or value, (ranging from highly negative to highly positive) comprising the two axes (Russell, 1980). In support of this framework, some components of the emotional response, such as approach or defensive behaviors, depend on computations of the positive or negative value of stimuli. Other components of the emotional response, such as metabolic arousal, increased attention and the enhancement of memory formation, are common to intense stimuli of both valences. Distinct psychophysiological measures correlate with arousal and valence, respectively (Lang and Davis, 2006). Finally, in addition to triggering these different aspects of emotional responses, surprising reinforcement can also induce learning – such that past experience is used to behave adaptively in the future

Correspondence and requests for materials should be addressed to C.D.S. (cds2005@columbia.edu).

*These authors contributed equally to this work.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

– a point emphasized in many theoretical models of reinforcement learning (Mackintosh, 1975; Pearce, 1980; Rescorla and Wagner, 1972; Sutton and Barto, 1998).

Neurobiological and psychological accounts of emotion often posit opponent motivational systems, one for appetitive, or reward-related, processes and one for aversive processes (Daw et al., 2002; Dickinson and Dearing, 1979; Grossberg, 1984; Konorski, 1967; Solomon and Corbit, 1974). Within this conceptualization, affective valence and arousal are determined by the relative activation and intensity of activation, respectively, of these two systems (Lang and Davis, 2006). Arousing pleasant or aversive stimuli can therefore elicit valence-specific responses while also enhancing autonomic reactivity (e.g. blood pressure, skin conductance responses), attention, and memory formation through a common, perhaps valence-insensitive, pathway. Thus appetitive and aversive brain systems may act in an opponent manner, for valence-specific processes, as well as in a congruent manner, for processes sensitive to affective intensity but not valence.

The amygdala, a brain structure located in the medial temporal lobe that has been implicated in emotional processes, is well positioned to integrate information from appetitive and aversive systems in order to direct emotional responses. Information about primary pleasant and aversive stimuli converges in the amygdala, which receives input from sensory systems of all modalities (Amaral et al., 1992; McDonald, 1998; Stefanacci and Amaral, 2002). Furthermore, prefrontal and rhinal cortices, as well as neuromodulatory systems and the hippocampus, all project to the amygdala (Amaral et al., 1992; Ghashghaei, 2002; McDonald, 1998; Stefanacci and Amaral, 2000, 2002), potentially influencing amygdala processing in relation to expectations, memories, and other cognitive processes. Output from the amygdala to subcortical structures and to the striatum, hippocampus and cortex could help coordinate a range of physiological, cognitive, and behavioral responses (Amaral et al., 1992; Davis, 2000). Some of these responses are valence-specific, such as freezing behavior (a component of fear responses) (LeDoux, 2000), and others are insensitive to valence, such as activation of the sympathetic nervous system in relation to arousal (Lang and Davis, 2006).

We were interested in understanding whether the amygdala processes primary reinforcement in a manner that could support its role in both valence-based and arousal-based emotional responses. To examine this issue, we focused on the effect of predictability on the responses of amygdala neurons to rewards and aversive stimuli. We measured amygdala responses to rewards and aversive stimuli during two experimental procedures that manipulated monkeys' ability to anticipate reinforcement. In one procedure, a trace-conditioning task, monkeys learned that the presentation of a conditioned stimulus (CS) predicted an unconditioned stimulus (US, a reward or an aversive stimulus). In the other procedure, rewards and aversive stimuli were delivered in an un-cued and unpredictable manner. Although this procedure did not place behavioral demands on the monkeys, we refer to it as the "random task" for convenience. Since variables other than reinforcement expectation, such as attention or cognitive load, may also differ between these two tasks, we use a comparison of responses to reinforcement during the trace-conditioning and random tasks to identify cells with potential effects of expectation. To confirm that in these cells expectation modulates reinforcement responses, we then examine neural responses to reinforcement in relation to learning and behavioral signs of expectation during the trace-conditioning task.

We found that expectation indeed influences the responses of amygdala neurons to reinforcement. Some cells showed a similar effect of expectation on responses to rewards and punishments. In other cells, however, expectation modulated the responses to either rewards or aversive stimuli, but not both. Neurons with congruent effects of expectation on responses to reinforcement of both valences are well-suited to support processes such as attention, attention-based learning, arousal, and the enhancement of memory formation, all of which have

been shown to involve the amygdala (Holland and Gallagher, 1999,2006;Kensinger and Corkin, 2004;McGaugh, 2004;Oya et al., 2005;Phelps and LeDoux, 2005). On the other hand, neurons that preferentially modulate their responses to either pleasant or aversive stimuli could play a role in processes that require information about stimulus valence (i.e. positive or negative value). Many of these processes, which include learning when to exhibit defensive as opposed to approach behaviors among others, are also likely to involve the amygdala (Baxter, 2002;Everitt et al., 2003;LeDoux, 2000;Paton et al., 2006).

In the trace-conditioning task, if predictability modulates neural responses to reinforcement, then these responses should change as monkeys learn to accurately anticipate reinforcement. Indeed, in the cells that respond more to reinforcers during the random task than during the trace task, we found that responses to rewards and aversive stimuli were stronger before learning, when monkeys did not expect the reinforcement, and that these responses decreased as monkeys learned. This expectation-dependent modulation of reinforcement responses bears similarity to a “prediction error signal”, which in theoretical accounts of reinforcement learning represents the difference between expected and received reinforcement (Pearce, 1980;Rescorla and Wagner, 1972;Sutton and Barto, 1998). However, the response profile of amygdala neurons that we studied cannot be described as simply encoding “pure” prediction errors because they carried other, additional signals. For example, we have previously shown that amygdala neurons are selective for the identity and value of visual CSs (Paton et al., 2006), and here we also show that these neurons can also respond to USs selectively. Instead, the influence of predictability on amygdala responses to reinforcement suggests that amygdala neurons may combine prediction error signal input with other signals carried by the amygdala during learning. These inputs may converge onto amygdala neurons in distinct patterns so as to support both valence-specific and arousal-specific processes.

Results

Predictability modulates amygdala neural responses to pleasant and aversive stimuli

To examine the effect of reinforcement expectation on neural responses in the amygdala, we measured neural activity during performance of two behavioral tasks: one in which reinforcement was predictable (trace-conditioning task), and another in which it was not (random task). In the trace-conditioning task, monkeys learned that USs (liquid rewards or aversive air-puffs directed at the face) were predicted by visual CSs. To assess learning, we measured two behaviors: licking a spout in anticipation of reward and blinking in anticipation of an air-puff. (Fig. 1A). (In different experiments, reinforcements occurred with a probability of 100% or 80%, which we will refer to as the 100% and 80% reinforced trace-conditioning tasks.) Subsequent to learning the initial CS-US contingencies, we reversed them without warning, and monkeys learned these new contingencies. In the random task, we recorded activity while presenting rewards and punishments with an equal probability but now in an uncued and unpredictable manner (Fig. 1B). Monkeys did not lick or blink in anticipation of reinforcement during the random task, indicating that they did not predict reinforcement.

We recorded the activity of 285 amygdala neurons from 5 monkeys during both the trace-conditioning and random tasks (100% reinforced task, 116 cells: 63 and 53 from monkeys V and P, respectively; 80% reinforced task, 169 cells: 46, 41, 25, 33, and 24 cells from monkeys V, P, Lo, R, and Lu, respectively). We examined activity in the 50–600 ms after reinforcement. The responses to reinforcement often differed in the two tasks, with some neurons exhibiting differential responses to rewards only, other neurons to punishments only, and some neurons to both rewards and punishments. For example, Fig. 2A,B shows activity from an amygdala neuron that responded more strongly when reward was delivered unexpectedly (random task) than when the same reward was delivered in a cued manner (trace-conditioning task). This cell did not respond to air-puffs during either task. Other amygdala neurons showed stronger

responses to unexpected air-puffs, but did not modulate their responses to reward (e.g. Fig. 2C,D). Finally, many cells had stronger responses to both unexpected rewards and air-puffs, as exemplified by the cell shown in Fig. 2E,F.

The response profiles of the neurons depicted in Fig. 2 suggest that expectation has differential effects on different populations of amygdala neurons. To examine this on a population level, we used a receiver operating characteristic (ROC) analysis to categorize each neuron according to whether it responded significantly more to unexpected reward only, air-puff only, or both during the random task as compared to each respective reinforcer when it was expected during the trace conditioning task ($p < 0.05$, one-tailed permutation test). 47/285 cells had differential responses for reward responses alone (expectation-modulated for reward cells, rEM), 35/285 cells for air-puff responses alone (expectation-modulated for aversive stimuli cells, aEM), and 65/285 cells for both types of reinforcement (non-valenced expectation-modulated cells, nEM).

Fig. 3A–F shows normalized and averaged neural activity for the rEM, aEM, and nEM cells. These plots combine data from all five monkeys because the proportion of rEM, aEM, and nEM neurons did not vary significantly between monkeys (χ^2 test, $p > 0.05$, Supplementary Table 1). rEM, aEM, and nEM neurons were represented in our sample of neurons more often than would be expected by chance (χ^2 test, $p > 0.05$ for each response profile type). In addition, across the population, the effects of expectation on responses to rewards and air-puffs were not independent (χ^2 test, $p < 0.001$). Neurons that showed (or did not show) stronger responses to unexpected rewards tended to show (or not show) stronger responses to air-puffs, with the incidence of nEM cells occurring at a rate greater than expected by chance given the observed frequencies of each type (binomial test, $p < 10^{-5}$). Together the findings suggest that different populations of neurons in the amygdala could subservise valence specific and valence non-specific functions. We have not noticed any clear anatomical clustering of reward-related vs. punishment-related effects of expectation on reinforcement responses; these cells appeared to be intermingled within the amygdala (Fig. 1C).

In general, the modulation of reinforcement responses in the random task compared to the trace-conditioning task was greater when reinforcement during trace conditioning occurred with a probability of 100% (Fig. 3G–I). These data are consistent with the notion that CSs do not fully predict reinforcement when delivery occurs with a probability of 80%; consequently, the difference in the predictability of rewards and air-puffs between the two tasks is smaller.

The interpretation that expectation modulates neural responses to reinforcement was also supported by an analysis of data from within the trace-conditioning task, using monkeys' anticipatory behavior as a proxy for their expectation. Recall that on the trace-conditioning task, monkeys learned to lick in anticipation of reward and blink in anticipation of an aversive air-puff. Assuming that licking but not blinking indicates expectation of reward, and vice-versa for expectation of air-puff, we asked whether reinforcement responses were higher when monkeys incorrectly predicted the upcoming reinforcement. Combining all the rEM, aEM, and nEM cells, we divided trials into two categories: one containing trials in which the monkey blinked but did not lick, and the other containing trials in which the monkey licked but did not blink. For each reinforcement type, we normalized responses within experiments in the same manner as for population PSTHs, and then combined data across experiments to compare activity in the two categories. For both reward and air-puff, we found that reinforcement responses were significantly greater when the monkey predicted the wrong type of reinforcement (Fig. 4A,B; analysis performed in 100 ms bins, stepped in 20 ms steps). These effects were weakened but not eliminated if the first four trials of initial learning and the first four trials after reversal were removed (most "incorrect" predictions occurred during these learning trials). In addition, the effects were smaller on the task with 80% reinforcement probability (Supplementary Figure 1). Overall, this analysis further supports the notion that

responses to reinforcement in the amygdala are modulated by expectation, and that violations of expectation have a greater impact on neural responses on the 100% reinforcement probability task.

Amygdala neurons weakly represent omitted rewards and punishments

So far we have focused on comparisons of responses to expected and unexpected reinforcement. Next, we examined whether the omission of an expected reinforcement also modulates amygdala neural activity by considering data from the 169 neurons studied during the 80% reinforced trace-conditioning task. In this task, reinforcement was omitted on 20% of trials. We sought to determine whether amygdala neural activity changed at the time of the omitted reinforcement. To detect reinforcement omission, monkeys must internally time when reinforcement should occur in relation to a CS presentation. The inherent variability of internal timing, combined with inherent variability of neural response latencies to reinforcement, led us to use an analysis of response latency to find neurons that responded to omitted rewards and punishments. This analysis identified whether and when a neuron significantly changed its response around the time of expected reinforcement when such reinforcement was omitted ($p < 0.05$, see methods). Of note, many amygdala neurons have low baseline firing rates (56% of neurons had baseline firing rates below 10 Hz), limiting the sensitivity of statistical approaches for detecting significant decreases in firing rate.

Fig. 5A,B shows a cell identified by the analysis of latency that fires somewhat more strongly to unexpected as compared to expected reward, and that decreases its firing rate to omitted rewards. In addition, the same cell increases its firing rate when an air-puff is omitted, and it decreases its firing rate to both predicted and unpredicted air-puffs. Thus this cell's response profile reveals opponent effects of rewards, punishments, and their omission in relation to expectation. Fig. 5C,D shows a similar opponent effects for a cell that increases its response to unexpected air-puff but not reward. Strikingly, an omitted reward also elicits an increase in response, and an omitted punishment elicits a modest decrease in response.

Theories of reinforcement learning often posit that valence-specific signals that represent the receipt or omission of reinforcement in relation to expectation play an instructive role in learning (Rescorla and Wagner, 1972; Sutton and Barto, 1998). Assuming that appetitive and aversive systems act as mutual opponents in the brain, as has been proposed previously (Daw et al., 2002; Grossberg, 1984; Solomon and Corbit, 1974), we grouped neurons according to whether or not their responses to omitted reinforcement were consistent with their providing a valence-specific signal. 32/169 (19%) neurons responded with significantly less firing to omitted rewards and/or significantly greater firing to omitted air-puffs (Fig. 5E,F; see Table 1), as determined by the analysis of response latency. These neurons encode outcomes of omitted reinforcement that are better than expected by increasing firing rate and/or of omitted reinforcement worse than expected by decreasing firing rate. Fig. 5G,H shows the response profile for 26/169 (15%) of neurons identified by the latency analysis as decreasing firing rate significantly after omitted air-puffs and/or increasing firing rate significantly after omitted reward. The neurons in Fig. 5E,F and Fig. 5G,H are non-overlapping populations (Table 1). Neurons from each of the 5 monkeys contributed to each panel depicted in Figs. 5E–H. From 50–600 ms after reinforcement, all curves within panels 5E–H are significantly different from each other (t-tests, $p < 0.05$ in each case), and on average omitted reinforcement resulted in a significant decrease in firing in Figs. 5E,H ($p < 0.01$, t-tests).

These data show that a small number of neurons in the amygdala modulate their firing rate when reinforcement is omitted. Such neurons often appear to integrate information about appetitive and aversive stimuli in relation to expectation in an opponent fashion, with omitted reinforcement having an opposite influence on firing rate for rewards and aversive stimuli. Across all 169 neurons, however, on average omitted reinforcement did not result in a

significant change in firing rate (Fig. 5I,J, t-tests comparing activity from the last 500 ms of the trace interval to activity 50–600 ms after the time of reinforcement, $p > 0.05$ for both omitted reward and air-puff responses). In addition, the responses after omitted reinforcement differed from the responses to expected rewards and air-puffs (t-test, $p < 0.05$). Thus omitted reinforcement has both a smaller and less frequent effect on firing rate than the increase in firing rate observed when reinforcement occurred unexpectedly.

Reinforcement responses dissipate in relation to learning about value at the behavioral and neural levels

During the trace-conditioning task, monkeys learned to predict reinforcement, as indicated by their anticipatory licking and blinking (Paton et al., 2006). However, immediately after we reversed the value assignments of CSs, which occurred without warning, monkeys' predictions about reinforcement were violated. We noticed that many neurons had stronger responses to rewards and air-puffs immediately after such image value reversals. Figs. 6A,B shows the responses of two cells with stronger responses to reinforcement after reversal for reward and air-puff respectively. Both neurons also encoded CS value, and neural activity reflecting CS value changed as reinforcement responses dissipated (see also Supplementary Figure 2 for additional examples).

For cells recorded during the 100% reinforcement trace-conditioning task, 65 cells had stronger responses to air-puffs and/or rewards during the random task (Supplementary Table 1); on average, these cells had elevated responses to reward and air-puff immediately after reversal, with responses dissipating in 5 or fewer trials (Fig. 6C,D). The stronger responses to rewards and air-puffs after reversals in reinforcement contingencies were likely due to the surprising reinforcement received, rather than to the expected reinforcement that was omitted. As shown in Fig. 5, on average omitted reinforcement has only a small effect on firing rate. Furthermore, neurons which did not respond more strongly to reinforcers in the random task also did not show a stronger response to reinforcement after reversal in the trace-conditioning task (Supplementary Figure 3). Finally, in the 65 cells with stronger responses to reinforcement on the random task, the mean difference between reinforcement responses on the trace-conditioning and random tasks was correlated with the change in response to reinforcement observed immediately after reversals during the trace-conditioning task (Supplementary Figure 4). Overall, therefore, comparing responses across the trace-conditioning and random tasks, and comparing responses with respect to reversals in reinforcement contingencies, reveal similar effects on neural response properties. The consistency in these data further supports the conclusion that expectation modulates responses to rewards and aversive stimuli in the amygdala.

To examine the relationship between changing reinforcement responses and behavioral learning, we constructed behavioral “learning curves” for positive-to-negative and negative-to-positive reversals. To construct these curves, we scored every trial according to whether monkeys licked and/or blinked in anticipation of reinforcement, and then normalized and averaged behavioral responses across experiments (see Methods). The dissipation in reinforcement responses shown in Figs. 6C,D was correlated with behavioral learning, indicating that reinforcement responses decline as monkeys learn to predict reinforcement (Fig. 6E,F; reward correlation: $r = -0.89$, $p < 10^{-7}$; air-puff correlation: $r = -0.74$, $p < 10^{-4}$). A similar relationship between the dissipation in reinforcement responses after reversal and learning was not found when we performed the same analysis on data from the 80% reinforced trace-conditioning task. This appeared to be largely due to a much smaller effect of reversal on reinforcement responses in that task (Supplementary Fig. 5).

We next considered whether the dissipation in reinforcement responses during learning was correlated with the time course of the evolving neural representation of CS value. To test this

idea, we derived an estimate of value encoded by the amygdala as a function of trial number by assuming that appetitive and aversive systems act as mutual opponents in the brain (Daw et al., 2002; Grossberg, 1984; Solomon and Corbit, 1974). Accepting this assumption, we take the difference between a given cell's activity on rewarded and punished trials as the "value signal" provided by that cell. For each neuron we computed the difference between the normalized (in the same manner as for population PSTHs) level of activity on rewarded and punished trials in 200 ms bins (shifted in 50 ms steps) in each of the 20 trials after reversal. We then averaged these normalized "value signals" across neurons that encoded positive and negative CS value, respectively, during the 100% trace-conditioning task (Fig. 7A,B). We defined a cell as encoding value if neural responses to both images (during either or both the visual stimulus or trace intervals) changed significantly and in opposite directions when reinforcement contingencies reversed (change point test, $p < 0.05$ after Bonferroni correction), or if a change point test detected a change in activity for only one image, but a 2-way ANOVA, with image value and image identity as main factors, confirmed a significant effect of image value ($p < 0.05$) (Paton et al., 2006). Positive value coding cells fired more in response to images associated with reward, and negative value-coding cells fired more in response to images associated with air-puff.

CS value-coding, estimated as the average normalized value signal from CS onset until the end of the trace interval, evolved rapidly after reversal in cells encoding positive and negative value (Fig. 7C,D). The evolving representation of CS value was correlated with the reduction over successive trials of neural responses to USs (Fig. 7E,F; reward responses: $r = -0.83$, $p < 10^{-5}$; air-puff responses $r = -0.58$, $p < 0.01$). Thus, the dynamics of reinforcement responses are consistent with the notion that this changing activity may be related to the process of updating representations of value in the amygdala during learning.

Multiple signals carried by individual amygdala neurons

Many amygdala neurons that modulated their response to reinforcement in an expectation-dependent manner also encoded other information. Here we characterize and distinguish among 3 different signals often carried by amygdala neurons: (1) the modulation of neural responses to reinforcement by expectation; (2) the encoding of CS value; and (3) neural responses selective for rewards or aversive air-puffs. In addition, we have previously shown that amygdala neurons often responded selectively to the visual CSs themselves during the trace-conditioning task (Paton et al., 2006).

One critical question to resolve regarding the different signals carried by amygdala neurons concerns whether signals encoding CS value may be instead described as encoding pure prediction errors, such as those formulated by temporal difference (TD) learning algorithms (Schultz et al., 1997; Sutton and Barto, 1998). In TD learning, responses to unexpected reinforcement transfer with learning so that a response to the CS that predicts reinforcement develops as the response to a previously unexpected US disappears (Schultz et al., 1997). Thus it is possible that what appears to be a signal encoding CS value could be better explained as simply encoding a TD error specific for one valence of reinforcement. In this scenario, a neuron encoding a pure TD error could have neural responses to reinforcement modulated by expectation and neural responses to CSs that reflect value. However, the response profile of amygdala neurons encoding CS value does not match the prototypical phasic, short latency response profile of neurons encoding TD errors. Across the population of amygdala neurons, the representation of CS value extends temporally from shortly after CS onset and continues through the trace interval (Fig. 8A–D; see also (Paton et al., 2006)). Very few neurons appear to be candidates for encoding pure TD errors. Among neurons encoding positive CS value within our sample, only 5 had short latency, phasic value-related signals while also modulating their responses to only rewards depending upon expectation. No neurons that we recorded

encoded negative value with an equivalent criterion; 1 neuron showed a short latency but longer duration signal encoding negative CS value.

Although these 6 cells are candidates for encoding a relatively pure prediction error, closer inspection reveals that they differ from pure TD errors in a number of ways (Supplementary Fig. 6). First, these cells exhibit reinforcement selectivity that is maintained after learning, whereas a prediction error should lose its US selectivity when an animal learns to predict reinforcement. Second, these cells exhibit a robust excitatory response to the CS conditioned with the opposite value to the preferred one. Overall, within the population of amygdala cells we recorded, we did not identify a single example of a cell whose response profile was consistent with its encoding a pure TD-style prediction error signal. Since cells that have reinforcement responses modulated by expectation are not well described as carrying pure TD errors, neural responses reflecting CS value and neural responses to reinforcement modulated by expectation therefore appear to be separate signals carried by amygdala neurons.

Amygdala neurons also often carry a third type of signal: selectivity for rewards and air-puffs. We used an ROC analysis to construct a reinforcement selectivity index for each cell, with ROC values > 0.5 indicating that a cell responded more strongly to rewards than air-puffs. 178/285 cells recorded had a selective response to the type of reinforcement (ROC different from 0.5, permutation test, $p < 0.05$, two-tailed), with 103 neurons preferring rewards, and 75 neurons preferring air-puffs. Of the 103 selective for reward, 20 showed no significant response to air-puff, 45 showed inhibition, and 38 showed excitation. Of the 75 cells selective for air-puff, 41 showed no response to reward, 20 showed inhibition, and 14 showed excitation (Supplementary Table 2). The rates at which we find cells that exhibit reinforcement selectivity did not vary significantly across the five monkeys (χ^2 test, $p > 0.05$, Supplementary Table 3). In addition, we did not find any significant relationship between reinforcement selectivity and the modulation of reinforcement responses by expectation (Supplementary Figure 7).

On average, neurons encoding positive and negative CS value maintained higher levels of activity after rewards and air-puffs, respectively (Fig. 8A–D; see also Fig. 7A,B). Indeed, the majority of neurons encoding positive CS value had reinforcement selectivity indices > 0.5 , and the majority of neurons encoding negative CS value had reinforcement selectivity indices < 0.5 (inset histograms, Fig. 8A–D). US selectivity was similar for the trace-conditioning and random tasks, as revealed by a significant correlation between indices representing US selectivity in each task ($r = 0.66$, $p < 10^{-4}$, and $r = 0.79$, $p < 10^{-3}$ for comparisons between the random task and the trace-conditioning tasks with 100% and 80% reinforcement probabilities, respectively). These observations are consistent with the idea that amygdala neurons may continually encode the value of sensory events in the environment, whether those events predict reinforcement (CSs) or are reinforcements themselves (USs).

The modulation of reinforcement responses depending upon expectation is often superimposed on signals representing CS and US value (Fig. 8E,F); in fact, 34/116 (29%) and 27/169 (16%) of neurons carried all three signals in the 100% and 80% reinforcement probability versions of the trace-conditioning task, respectively. Moreover, for the neurons studied when reinforcement probability was 100%, the presence of an expectation-modulated reinforcement response was associated with a higher probability that the neuron encoded CS value (χ^2 test, $p < 0.005$). Forty-five of the 65 neurons showing stronger responses to unexpected compared to expected reinforcement also encoded CS value, whereas only 21 out of the 51 remaining neurons encoded value. These data suggest that these different signals may be linked during the process of learning about value. However, we did not find that this relationship between signals also applied to the neurons studied during the 80% reinforced trace-conditioning task (χ^2 test, $p = 0.4$), perhaps because CSs do not fully predict reinforcement when delivery occurs with a probability of 80%.

Discussion

We have examined the responses of amygdala neurons to rewards and aversive air-puffs during two tasks: a trace-conditioning procedure in which monkeys learned to predict reinforcement based on CSs, and a “random” task, in which reinforcement occurred unpredictably. Our central finding is that expectation frequently modulated amygdala neurons’ responses to rewards and aversive stimuli. During the trace-conditioning task, on average, neural responses were larger when monkeys incorrectly predicted rewards or air-puffs (Fig. 4). Furthermore, by testing every neuron with both expected and unexpected rewards and punishments, we identified two different populations of neurons in the amygdala. The activity of some amygdala neurons reflects a valence-specific effect of expectation on responses to pleasant or aversive stimuli, but not both; the activity of other amygdala neurons reflects similar effects of expectation on responses to reinforcements of both valences (Fig. 3).

The two groups of neurons we have described could underlie the role of the amygdala in integrating a wide array of sensory information in order to drive physiological, cognitive, and behavioral responses related to emotion (Balleine and Killcross, 2006;Baxter, 2002;Everitt et al., 2003;Holland and Gallagher, 1999;Lang and Davis, 2006;LeDoux, 2000;McGaugh, 2004;Paz et al., 2006;Phelps, 2006;Phelps and LeDoux, 2005). Some processes, such as fear or reward-seeking behavior, are specific to stimulus valence (positive or negative value). Other processes, such as arousal, attention, and memory formation enhancement, are instead sensitive to stimulus intensity regardless of valence. The different types of expectation modulation we have described in amygdala neurons may play specific roles in each of these two sorts of processes.

Although valence non-specific effects of expectation on reinforcement responses could derive from or contribute to processes like generalized attention and arousal, valence-specific effects of expectation cannot be easily explained by such processes. For example, one might hypothesize that generalized attention or arousal could modulate neural activity in a valence-specific manner if 1) a neuron responds to only reward or punishment, and 2) attention or arousal changes the gain of the neuron’s responses. However, the effects of expectation on neural responses were unrelated to neurons’ inherent selectivity for rewards or air-puffs (Supplementary Fig. 7). Moreover, as shown in Fig. 3A,B, on average rEM cells actually responded more strongly to expected air-puffs than to expected rewards. Therefore, neurons with valence-specific effects of expectation likely reflect their specific connectivity with neural circuits that process rewards or aversive stimuli in relation to expectation, but not both. Consistent with this notion, in 21 experiments, multiple neurons with different properties – valence specific and valence non-specific effects of expectation - were recorded simultaneously. Generalized attention or arousal is unlikely to account for the differential response properties of these neurons. Instead, one would have to invoke multiple types of attention and/or arousal mechanisms, including some that are valence-specific, to explain how neurons with different properties were recorded simultaneously.

The modulation of reinforcement responses by expectation that we describe in amygdala neurons bears some similarity to prediction error signals. Prediction error signals have been proposed to play an important role in associative learning, albeit in somewhat different ways depending upon the theoretical approach. For example, in the trial-based Rescorla-Wagner formulation (Rescorla and Wagner, 1972), prediction error signals indicate the sign and magnitude of the adjustment required to predict reinforcement accurately in the future, and thereby could act as teaching signals that update neural representations of value. TD learning algorithms generalize the trial-based Rescorla-Wagner formulation to characterize prediction errors in real-time, and posit a similar role for prediction error signals (Sutton and Barto, 1998). In TD learning algorithms, error signals are continuously computed by taking the

difference in the value of two successive “states”, where a state is roughly defined as the overall situation. In the context of our study, we may conceptualize a trial as containing two states: one initiated by CS presentation and the other by US presentation. In this scenario, both a CS that is known to predict a particular US and a surprising US could cause a phasic prediction error signal in a cell. Midbrain dopamine neurons appear to encode reward prediction error signals as defined by TD models (Schultz et al., 1997), but punishment prediction error signals, to our knowledge, have not been described at the cellular level. Error signals also play a role in attention-based models of reinforcement learning; in this formulation, these signals serve to enhance the associability of the CS during learning by modulating attention to the CS (Mackintosh, 1975; Pearce, 1980).

Although the fact that expectation can modulate amygdala responses to reinforcement is reminiscent of TD prediction error signals, it is clear that in general amygdala neurons do not simply encode “pure” TD errors. Amygdala neurons carry a number of other signals, including signals selective for visual stimulus identity (Paton et al., 2006), temporally extended signals that are selective for CS value (Paton et al., 2006), and signals selective for USs (Fig. 8). Moreover, the vast majority of amygdala neurons (including those that display expectation-dependent modulation of reinforcement responses) do not have phasic CS responses after learning that are consistent with TD error signals. Even those neurons with phasic CS value signals and with expectation-modulated reinforcement responses exhibit other response properties (Supplementary Figure 6). Thus the response profile of amygdala neurons differs from what has been ascribed to dopamine neurons for learning about reward (Schultz et al., 1997). In addition, few amygdala neurons appear to have activity that reflects “negative” reward or punishment prediction errors (Fig. 5). Negative prediction errors occur when the reinforcement received is smaller than that expected or omitted altogether, and our data revealed a weak and infrequent representation of omitted reinforcement of both valences in the amygdala.

The convergence of two types of signals – those related to value (for both CSs and USs), and those related to the predictability of reinforcement – suggests that changes in neural responses to CSs during learning may arise from the reciprocal interactions between these signals within the amygdala. Consistent with this notion, we found that reinforcement responses dissipated during learning, as monkeys learned to predict reinforcement (Fig. 6). The dynamics of this dissipation were correlated with behavioral learning about value (Fig. 6), and with the evolving representation of value in the amygdala (Fig. 7). Together, these data link neural responses to reinforcement that are modulated by expectation to the dynamic process of learning about value at the neural and behavioral levels. However, since different populations of amygdala neurons represent expectation in either a valence-specific (Fig. 3A–D) or valence non-specific (Fig. 3E,F) manner, the mechanism by which expectation-modulated reinforcement signals may help induce learning about value remain unclear. Neurons in which expectation modulates reinforcement responses in a valence-specific manner may not be the only type that participate in learning about value. Neurons that have enhanced responses for both unexpected rewards and punishments still could play a role in driving learning about value according to attention-based theories of reinforcement learning (Mackintosh, 1975; Pearce, 1980). In fact, the amygdala has been proposed to play a critical role in processing how surprising rewards influence learning (Holland and Gallagher, 2006); and surprising punishments could have the same effect on attention-based learning mechanisms.

In principle, expectation could modulate reinforcement responses in the amygdala through a local, *de novo* computation, or these responses could reflect inputs that already represent unexpected reinforcement. The amygdala is reciprocally connected with numerous brain structures – such as midbrain dopamine neurons and prefrontal cortex – where expectation-modulated reinforcement signals may also be represented (Schultz and Dickinson, 2000). The

physiological responses we observe may therefore derive from these inputs. We tend to favor this interpretation because many aspects of amygdala neural responses, reviewed above, are inconsistent with the idea that these neurons simply compute a TD-style prediction error. Instead, the observed effects of expectation on reinforcement responses may reflect a physiological signature of error signal input from other structures, and therefore are consistent with the notion that the amygdala may be a site of convergence and interaction for value and error signals.

In addition to the representation of positive and negative value in the amygdala (Paton et al., 2006), a number of brain structures – including orbitofrontal, cingulate, and parietal cortices, and the basal ganglia – have been shown to contain single neurons that encode the reward value of stimuli and/or of potential motor actions (Dorris and Glimcher, 2004; McCoy and Platt, 2005; Padoa-Schioppa and Assad, 2006; Roesch and Olson, 2004; Salzman et al., 2005; Samejima et al., 2005; Sugrue et al., 2004, 2005). Neurophysiological recordings have shown evidence of reward responses modified by expectation in many of these areas (Schultz and Dickinson, 2000), including the amygdala (Sugase-Miyamoto and Richmond, 2005). However, unlike the present study, previous studies have not tested whether neurons represent both pleasant and aversive stimuli in relation to expectation, thus revealing how neurons differentially integrate information from appetitive and aversive systems. Moreover, previous work typically has not described the time course of these signals in relation to learning on a trial-by-trial basis, and these studies have not simultaneously examined these signals in relation to an evolving neural representation of value. Recently, experiments using fMRI have suggested that BOLD responses reflect both reward (Haruno and Kawato, 2006; O'Doherty, 2003b; Seymour et al., 2005) and punishment (Ploghaus et al., 2000; Seymour et al., 2004; Seymour et al., 2005) prediction error in the cerebellum, orbitofrontal, parietal, right insula and cingulate cortices, ventral striatum and the hippocampus. Despite its long history of being implicated in aversive learning (LeDoux, 2000; Paton et al., 2006), only one recent study has identified the amygdala as having a BOLD signal reflecting punishment error signals by using monetary loss as an aversive stimulus (Yacubian et al., 2006). In this study, however, reward prediction error signals were not identified in the amygdala, perhaps because of the limited spatial and temporal resolution of fMRI.

Compared to investigations of how expectation modulates reward responses, much less information is available concerning how expectation may modulate responses to aversive stimuli. Although dopamine neurons are frequently described as representing reward error signals, conflicting data exist concerning whether dopamine neurons also provide a phasic response to unexpected aversive stimuli (Guarraci et al., 1999; Horvitz, 2000; Mirenowicz and Schultz, 1996; Ungless, 2004). Dopamine neurons have been found to be active tonically during stress (Anstrom and Woodward, 2005), and recent evidence indicates that dopamine may play a role in fear conditioning and in modulating amygdala processing (Grace and Rosenkranz, 2002; Marowsky et al., 2005; Nader and LeDoux, 1999). Serotonergic neurons, which also project to the amygdala (Amaral et al., 1992), are another candidate source for signals about unexpected aversive stimuli. Opponent appetitive and aversive systems have long been proposed as a basis for critical aspects of reinforcement learning and affective behavior (Grossberg, 1984; Solomon and Corbit, 1974). A recent proposal suggests that, for aversive processing, the serotonergic system may mirror the role of the dopaminergic system in reward (Daw et al., 2002), but neurophysiological evidence demonstrating this is lacking. Thus, establishing the direction and nature of interactions between brain areas involved in appetitive and aversive processing remains an important goal for understanding how neural circuits implement reinforcement learning.

Motivated by a conceptual framework in which stimulus valence and intensity both contribute to our emotional experience, we have examined the activity of single neurons in the amygdala

to both predictable and unpredictable pleasant and aversive stimuli. We discovered that expectation modulates responses to reinforcement in a valence-specific manner in some neurons, and in a valence non-specific manner in other neurons. These different types of response properties may underlie the role of the amygdala in multiple processes related to emotion, including reinforcement learning, attention, and arousal. Future work must develop experimental approaches for unraveling the complex anatomical circuitry and mechanisms by which amygdala neurons influence learning and the many emotional processes related to the valence and intensity of reinforcing stimuli.

Methods

General Methods

During experiments, monkeys sat in a Plexiglas primate chair (Crist Instruments) with their eyes 57 cm in front of a 21" Sony CRT monitor. Monkeys were monitored using infrared video. Our general methods have been described previously (Paton et al., 2006). All animal procedures conformed to NIH guidelines and were approved by the Institutional Animal Care and use Committees at New York State Psychiatric institute and Columbia University.

Behavior

Trace-conditioning task—We used a trace-conditioning procedure to induce learning about the reinforcement associated with three visual CSs (novel fractal patterns) in every experiment. To begin a trial, monkeys foveated a fixation point (FP). After 1 s, an 8 degree square image (CS) appeared over the fovea for 300 or 350 ms. During fixation, monkeys maintained gaze within 3.5 degrees of the FP, as measured with an Applied Science Laboratories infrared eye tracker sampling at 240 Hz. 1.5 s (trace interval) after image and FP disappearance, US delivery occurred with a 100% or 80% probability (for the 100% and 80% reinforced versions of the task, respectively). USs were liquid rewards (0.1–0.9 ml of water), a 50–100 ms, 40–60 psi aversive air-puff or nothing. Air-puff delivery occurred at one of two possible locations on the monkey's face chosen randomly on every trial. All 3 trial types were presented in blockwise randomized order (2 trials of each type randomized within a block), with a 3.5 s intertrial interval. We reversed CS value assignments without warning, with an initially rewarded CS now followed by air-puff, and vice-versa for the punished CS. We assessed monkeys' learning by measuring licking and blinking responses (see below). Other than differing in reinforcement probability, the 80% and 100% reinforced versions of the task differed in that one CS on the 100% task was non-reinforced and never reversed, whereas the equivalent CS on the 80% task was associated with a small reward and never reversed.

Random task—This procedure typically followed the trace-conditioning task. With no behavioral requirements enforced, monkeys were presented randomly with rewards and air-puffs of the exact same duration and magnitude as during the trace-conditioning task. Times between the administrations of reinforcement were generated from a truncated exponential distribution with a mean of 5–6 seconds, minimum of 3.5 seconds and maximum of 15–20 seconds. Within an experiment, air-puffs and rewards occurred randomly and with equal frequency.

Behavioral measures—The same behavioral measures were used in both tasks. To measure licking, we measured every ms whether the monkey's tongue interrupted an infrared beam passed between the monkey's mouth and the reward delivery tube positioned 1–2 cm away. Our eye tracker registered anticipatory blinking by outputting a characteristic voltage when the eye closed, which we verified as corresponding to eye closures as visualized from an infrared camera that monitored the monkey during the experiment. For quantitative analysis

of behavioral learning, we scored each trial according to whether the monkey licked and/or closed its eyes in the 500 ms preceding reinforcement.

Data collection

We recorded neural activity from 365 neurons in the right amygdala of 5 rhesus monkeys (*Macaca mulatta*) weighing 4–14 kg. 196 neurons were recorded during the 100% reinforced trace-conditioning procedure. We also studied 116 of these neurons during the random task. In addition, we recorded the activity of 169 neurons on the 80% reinforced trace-conditioning task and the random task. We positioned recording chambers directly over the amygdala based on MRI. In each experiment, we individually advanced 1–4 metal microelectrodes (FHC Instruments) into the brain through guide tubes using either a hydraulic microdrive (Narishige) or a motorized multi-electrode drive (NAN). Guide tubes were supported within a grid with holes spaced 1 mm or 1.3 mm apart. We used the Plexon system for signal amplification, filtering, digitizing of spike waveforms, and spike sorting using a principal component analysis platform (on-line with off-line verification). We included all well-isolated neurons in this study; monkeys either performed a fixation task or no task during the search for well-isolated neurons. MRI images were used to reconstruct the borders of the amygdala and recording sites (Fig 1C). The neuronal sample was taken from an overlapping region, predominantly the basolateral region, of the amygdala in each monkey.

Data analysis

We analyzed responses to USs during an epoch spanning from 50–600 ms after US onset, which we selected on the basis of an analysis of response latency and duration. Briefly, we compared activity from a baseline period, the last 500 ms of the trace interval, to activity in the 800 ms following US onset. We first constructed distributions of the number of spikes in 20 ms bins, shifted by 1 ms, across the baseline interval. Next, we determined which 20 ms bins, slid in 1 ms steps, met a criterion response during the 800ms after US onset. Criterion for an excitatory response was met if at least 20 consecutive overlapping bins exceeded 99% of the baseline activity, for an inhibitory response if at least 20 consecutive bins contained fewer spikes than 95% of baseline bins. Latency was defined as the beginning of the first of 20 consecutive significant bins. 90% of latencies were greater than 50 ms. We used 600 ms as the end of the reinforcement epoch because a similar analysis found that 90% of cells had a response duration of 568 ms or less. Using customized time windows for air-puff and reward did not change the central findings reported here, so we use the 50–600 ms window for both reinforcements. We also used this type of latency analysis to determine whether a neuron increased or decreased its response when rewards and punishments are omitted, during the 80% reinforced trace-conditioning task. For trials with omitted reinforcement, we designated the earliest possible onset of a response as the normal time of US delivery had it not been omitted.

ROC analyses—We used a ROC analysis adapted from signal detection theory (Green and Swets, 1966) for 3 analyses. In general, the outcome of these analyses is not different when using other non-parametric statistical techniques. We chose to use ROC analysis because it provides an index that quantifies the degree and sign of the separation between two distributions without making assumptions about their shapes. For all ROC analyses, we assess statistical significance with permutation tests, in which we shuffle data 1000 times in order to evaluate whether an ROC value computed on the data was significantly greater than 95% of the shuffled values ($p < 0.05$, one-tailed test), or greater than 97.5% or less than 2.5% of the shuffled values ($p < 0.05$, two-tailed test). The basic results we report in this paper remain consistent if we used more stringent statistical criteria, such as $p < 0.01$.

The first ROC analysis determined whether unexpected compared to expected reinforcement modulated amygdala neural activity. We compared responses to USs in the trace-conditioning task (expected USs) to responses from the random task (unexpected USs) to identify three categories of neurons: rEM, aEM, and nEM neurons. For this analysis, we considered trials from both the 100%-reinforced and 80% reinforced trace-conditioning tasks to be “expected” reinforcement trials. Unexpected responses were the responses to the same reinforcement presented during the random task. Of course, when reinforcement is actually received after a CS predicts reinforcement with 80% probability, there may be a small positive prediction error (since the monkey is not expecting reward with 100% probability). However, the presence of this small prediction error would actually make it more difficult to detect a difference between these trials and the unexpected reinforcement trials from the random task.

The second ROC analysis compared responses to expected rewards and air-puffs, determining whether a neuron fired more strongly to rewards or air-puffs. For these analyses, ROC values greater than 0.5 indicated cells that fired more strongly when reinforcement was unexpected or cells that fired more strongly to rewards.

The third ROC analysis was performed to determine the latency and duration of signals encoding CS value. We compared neural activity from trials when a CS was associated with a reward to trials when the same CS was associated with an air-puff. We computed ROC values in 100 ms bins, stepping in 20 ms increments, from CS onset until US onset, again using a permutation test to assess statistical significance. We defined value signal latency as the time of the first significant bin of at least four consecutive significant bins with the same valence of value coding and duration as the spanning from the first to the last significant bin.

Population PSTHs—Before averaging neural activity across cells in 10 ms non-overlapping bins, we normalized data from each experiment by subtracting the mean and dividing by the standard deviation of baseline activity. We used the last 1000 ms before FP onset or US onset as the baseline for the trace-conditioning and random tasks, respectively. Distributions of baseline activity combined data from all trials from both tasks so that only one computed mean and standard deviation was used for normalization of each neuron. We smoothed each PSTH by convolving it with a 50 ms Gaussian at half-width.

Behavioral learning curves—To quantify learning, we normalized and averaged behavioral responses across all experiments in which unexpected reward or air-puff significantly increased reinforcement responses. We first scored each of the 20 trials before and after reversal as a response (1) or no response (0). We then divided each trial’s response by the mean behavioral response across trials, and subtracted the mean of the normalized values from each individual normalized value. We multiplied by -1 normalized responses that decreased after reversal, so that the slope of learning curves was always positive. We did this separately for learning about positive-to-negative and negative-to-positive reversals, so that we could compare each type of learning to the responses to rewards and punishments during learning (see Fig. 6E,F).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

We thank E. Kandel, S. Fusi, P. Dayan, Y. Niv, J. Dudman, G. Corrado, P. Glimcher, S. Siegelbaum, B. Lau and members of the Salzman lab and the Mahoney Center at Columbia for helpful comments and discussions; S. Dashnaw and J. Hirsch for MRI support; D. Schneider for help with data collection, K. Marmon for invaluable technical assistance. This research was supported by the NIMH, and the Klingenstein, Keck, Sloan, James S. McDonnell and

NARSAD foundations, and by a Charles E. Culpeper Scholarship to C.D.S. J.J.P. received support from N.I.C.H.D. and N.E.I. institutional training grants. S.E.M. received support from the N.S.F.

References

- Amaral, D.; Price, J.; Pitkanen, A.; Carmichael, S. Anatomical Organization of the Primate Amygdaloid Complex. In: Aggleton, J., editor. *The Amygdala: Neurobiological Aspects of Emotion, Memory, and Mental Dysfunction*. New York: Wiley-Liss, Inc; 1992. p. 1-66.
- Anstrom KK, Woodward DJ. Restraint increases dopaminergic burst firing in awake rats. *Neuropsychopharmacology* 2005;30:1832–1840. [PubMed: 15886724]
- Balleine BW, Killcross S. Parallel incentive processing: an integrated view of amygdala function. *Trends in Neurosciences* 2006;29:272–279. [PubMed: 16545468]
- Baxter M, Murray EA. The amygdala and reward. *Nature Reviews Neuroscience* 2002;3:563–573.
- Davis, M. The role of the amygdala in conditioned and unconditioned fear and anxiety. In: Aggleton, J., editor. *The Amygdala: A Functional Analysis*. New York: Oxford University Press; 2000. p. 213-287.
- Daw ND, Kakade S, Dayan P. Opponent interactions between serotonin and dopamine. *Neural Networks* 2002;15:603–616. [PubMed: 12371515]
- Dickinson, A.; Dearing, MF. Appetitive-aversive interactions and inhibitory processes. In: Dickinson, A.; Boakes, RA., editors. *Mechanisms of learning and motivation*. Hillsdale, N.J.: Erlbaum; 1979. p. 203-231.
- Dorris MC, Glimcher PW. Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron* 2004;44:365–378. [PubMed: 15473973]
- Everitt BJ, Cardinal RN, Parkinson JA, Robbins TW. Appetitive behavior: impact of amygdala-dependent mechanisms of emotional learning. *Annals of the New York Academy of Sciences* 2003;985:233–250. [PubMed: 12724162]
- Ghashghaei H, Barbas H. Pathways for emotion: interactions of prefrontal and anterior temporal pathways in the amygdala of the rhesus monkey. *Neuroscience* 2002;115:1261–1279. [PubMed: 12453496]
- Grace AA, Rosenkranz JA. Regulation of conditioned responses of basolateral amygdala neurons. *Physiology & Behavior* 2002;77:489–493. [PubMed: 12526988]
- Green, DM.; Swets, JA. *Signal Detection Theory and Psychophysics*. New York: John Wiley and Sons, Inc.; 1966.
- Grossberg S. Some normal and abnormal behavioral syndromes due to transmitter gating of opponent processes. *Biological Psychiatry* 1984;19:1075–1118. [PubMed: 6148110]
- Guarraci FA, Frohardt RJ, Kapp BS. Amygdaloid D1 dopamine receptor involvement in Pavlovian fear conditioning. *Brain Research* 1999;827:28–40. [PubMed: 10320690]
- Haruno M, Kawato M. Different neural correlates of reward expectation and reward expectation error in the putamen and caudate nucleus during stimulus-action-reward association learning. *Journal of neurophysiology* 2006;95:948–959. [PubMed: 16192338]
- Holland PC, Gallagher M. Amygdala circuitry in attentional and representational processes. *Trends in Cognitive Sciences* 1999;3:65–73. [PubMed: 10234229]
- Holland PC, Gallagher M. Different Roles for Amygdala Central Nucleus and Substantia Innominata in the Surprise-Induced Enhancement of Learning. *Journal of Neuroscience* 2006;26:3791–3797. [PubMed: 16597732]
- Horvitz JC. Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. *Neuroscience* 2000;96:651–656. [PubMed: 10727783]
- Kensinger EA, Corkin S. Two routes to emotional memory: distinct neural processes for valence and arousal. *Proceedings of the National Academy of Sciences of the United States of America* 2004;101:3310–3315. [PubMed: 14981255]
- Konorski, J. *Integrative activity of the brain: An interdisciplinary approach*. Chicago: University of Chicago Press; 1967.
- Lang PJ, Davis M. Emotion, motivation, and the brain: reflex foundations in animal and human research. *Prog Brain Res* 2006;156:3–29. [PubMed: 17015072]
- LeDoux JE. Emotion circuits in the brain. *Annual Review of Neuroscience* 2000;23:155–184.

- Mackintosh NJ. Blocking of conditioned suppression: role of the first compound trial. *Journal of Experimental Psychology: Animal Behavior Processes* 1975;1:335–345. [PubMed: 1202140]
- Marowsky A, Yanagawa Y, Obata K, Vogt KE. A specialized subclass of interneurons mediates dopaminergic facilitation of amygdala function. *Neuron* 2005;48:1025–1037. [PubMed: 16364905] [see comment]
- McCoy AN, Platt ML. Risk-sensitive neurons in macaque posterior cingulate cortex. *Nature neuroscience* 2005;8:1220–1227.
- McDonald AJ. Cortical pathways to the mammalian amygdala. *Progress in Neurobiology* 1998;55:257–332. [PubMed: 9643556]
- McGaugh JL. The amygdala modulates the consolidation of memories of emotionally arousing experiences. *Annual Review of Neuroscience* 2004;27:1–28.
- Mirenowicz J, Schultz W. Preferential activation of midbrain dopamine neurons by appetitive rather than aversive stimuli. *Nature* 1996;379:449–451. [PubMed: 8559249]
- Nader K, LeDoux JE. Inhibition of the mesoamygdala dopaminergic pathway impairs the retrieval of conditioned fear associations. *Behavioral Neuroscience* 1999;113:891–901. [PubMed: 10571473]
- O’Doherty J, Dayan P, Friston K, Critchley H, Dolan RJ. Temporal difference models and reward-related learning in the human brain. *Neuron* 2003b;38:329–337.
- Oya H, Adolphs R, Kawasaki H, Bechara A, Damasio A, Howard MA 3rd . Electrophysiological correlates of reward prediction error recorded in the human prefrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America* 2005;102:8351–8356. [PubMed: 15928095]
- Padoa-Schioppa C, Assad JA. Neurons in the orbitofrontal cortex encode economic value. *Nature* 2006;441:223–226. [PubMed: 16633341]
- Paton J, Belova M, Morrison S, Salzman C. The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature* 2006;439:865–870. [PubMed: 16482160]
- Paz R, Pelletier JG, Bauer EP, Pare D. Emotional enhancement of memory via amygdala-driven facilitation of rhinal interactions. *Nature neuroscience* 2006;9:1321–1329.
- Pearce J, Hall G. A model for Pavlovian conditioning: variations in the effectiveness of conditioned but not unconditioned stimuli. *Psychol Rev* 1980;87:532–552. [PubMed: 7443916]
- Pelphs EA. Emotion and cognition: insights from studies of the human amygdala. *Annu Rev Psychol* 2006;57:27–53. [PubMed: 16318588]
- Pelphs EA, LeDoux JE. Contributions of the amygdala to emotion processing: from animal models to human behavior. *Neuron* 2005;48:175–187. [PubMed: 16242399]
- Ploghaus A, Tracey I, Clare S, Gati JS, Rawlins JN, Matthews PM. Learning about pain: the neural substrate of the prediction error for aversive events. *Proceedings of the National Academy of Sciences of the United States of America* 2000;97:9281–9286. [PubMed: 10908676]
- Rescorla, RA.; Wagner, AR. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In: Black, AH.; Prokasy, WF., editors. *Classical Conditioning II: Current Research and Theory*. New York: Appleton Century Crofts; 1972. p. 64-99.
- Roesch MR, Olson CR. Neuronal activity related to reward value and motivation in primate frontal cortex. *Science* 2004;304:307–310. [PubMed: 15073380]
- Russell JA. A circumplex model of affect. *Journal of Personality and Social Psychology* 1980;39:1161–1178.
- Salzman CD, Belova MA, Paton JJ. Beetles, boxes and brain cells: neural mechanisms underlying valuation and learning. *Current Opinion in Neurobiology* 2005;15:721–729. [PubMed: 16271457]
- Samejima K, Ueda Y, Doya K, Kimura M. Representation of action-specific reward values in the striatum. *Science* 2005;310:1337–1340. [PubMed: 16311337]
- Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science* 1997;275:1593–1599. [PubMed: 9054347]
- Schultz W, Dickinson A. Neuronal coding of prediction errors. *Annual Review of Neuroscience* 2000;23:473–500.

- Seymour B, O'Doherty JP, Dayan P, Koltzenburg M, Jones AK, Dolan RJ, Friston KJ, Frackowiak RS. Temporal difference models describe higher-order learning in humans. *Nature* 2004;429:664–667. [PubMed: 15190354]
- Seymour B, O'Doherty JP, Koltzenburg M, Wiech K, Frackowiak R, Friston K, Dolan R. Opponent appetitive-aversive neural processes underlie predictive learning of pain relief. *Nature neuroscience* 2005;8:1234–1240.
- Solomon RL, Corbit JD. An opponent-process theory of motivation. I. Temporal dynamics of affect. *Psychological Review* 1974;81:119–145. [PubMed: 4817611]
- Stefanacci L, Amaral DG. Topographic organization of cortical inputs to the lateral nucleus of the macaque monkey amygdala: a retrograde tracing study. *Journal of Comparative Neurology* 2000;421:52–79. [PubMed: 10813772]
- Stefanacci L, Amaral DG. Some observations on cortical inputs to the macaque monkey amygdala: an anterograde tracing study. *Journal of Comparative Neurology* 2002;451:301–323. [PubMed: 12210126]
- Sugase-Miyamoto Y, Richmond BJ. Neuronal signals in the monkey basolateral amygdala during reward schedules. *Journal of Neuroscience* 2005;25:11071–11083. [PubMed: 16319307]
- Sugrue LP, Corrado GS, Newsome WT. Matching behavior and the representation of value in the parietal cortex. *Science* 2004;304:1782–1787. [PubMed: 15205529][see comment]
- Sugrue LP, Corrado GS, Newsome WT. Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat Rev Neurosci* 2005;6:363–375. [PubMed: 15832198]
- Sutton, R.; Barto, A. *Reinforcement Learning*. Cambridge, Massachusetts: MIT Press; 1998.
- Ungless MA. Dopamine: the salient issue. *Trends Neurosci* 2004;27:702–706. [PubMed: 15541509]
- Yacubian J, Glascher J, Schroeder K, Sommer T, Braus DF, Buchel C. Dissociable systems for gain- and loss-related value predictions and errors of prediction in the human brain. *Journal of Neuroscience* 2006;26:9530–9537. [PubMed: 16971537]

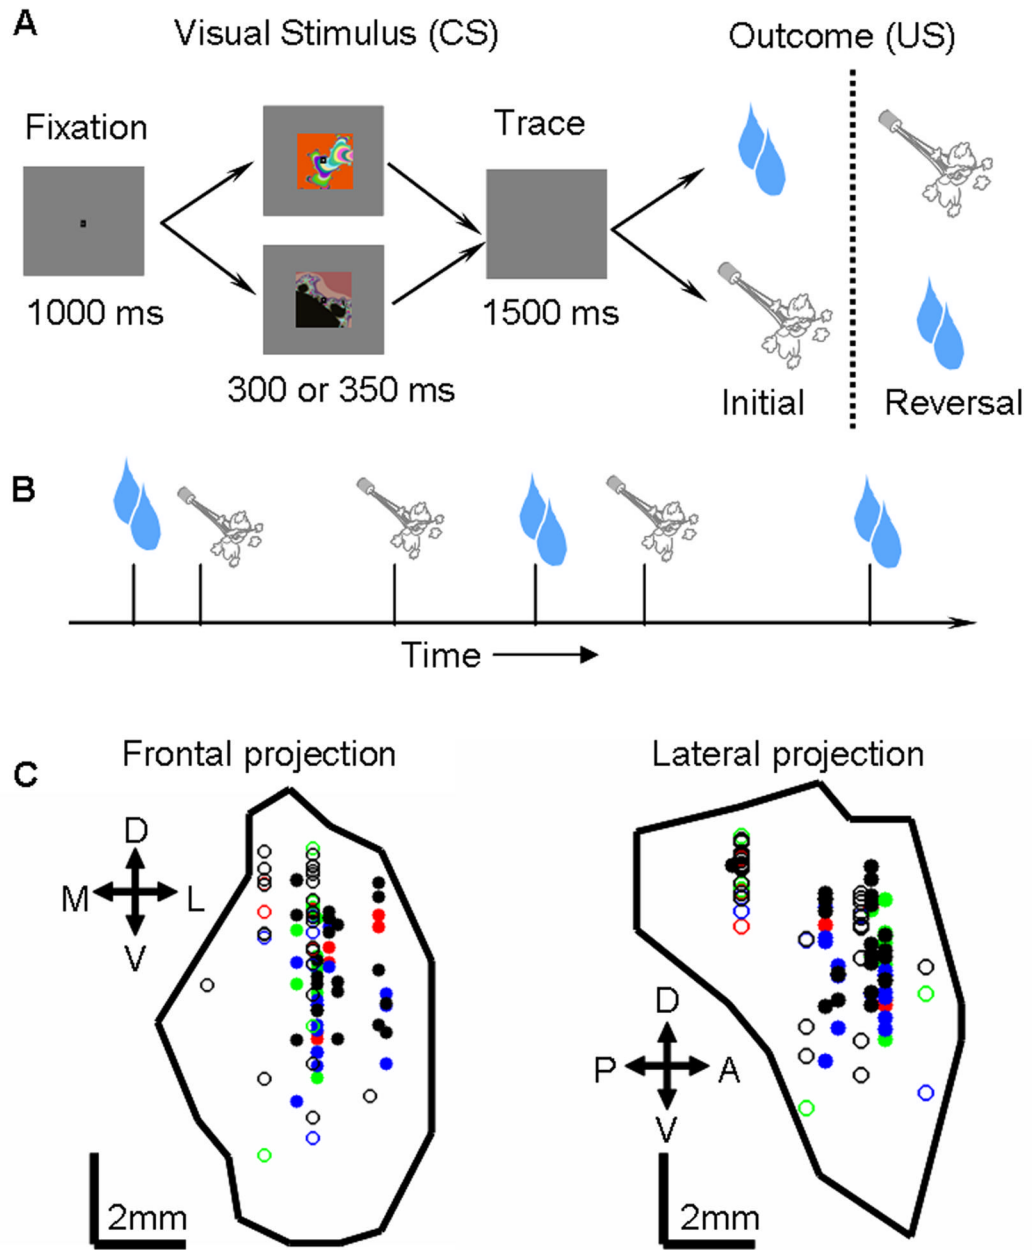


Figure 1. Behavioral tasks and recording site reconstruction

A. Trace-conditioning task. Sequence of events for trials. For CSs followed by large rewards or punishments, reinforcement contingencies reverse without warning after initial learning. Not depicted: a third trial type, in which non-reinforcement or a small reward followed a CS. B. Random task. Rewards and air-puffs were presented with equal probability in a random order. C. Anatomical reconstruction of recording sites in monkey L, with amygdala extent and site locations estimated by MRI. Left, coronal slice. Right, sagittal slice. Symbols indicate properties of recorded cells. Green, rEM cells; red, aEM cells; black, nEM cells; blue, no effect of expectation; solid and open circles, 100% and 80% probability trace-conditioning tasks, respectively.

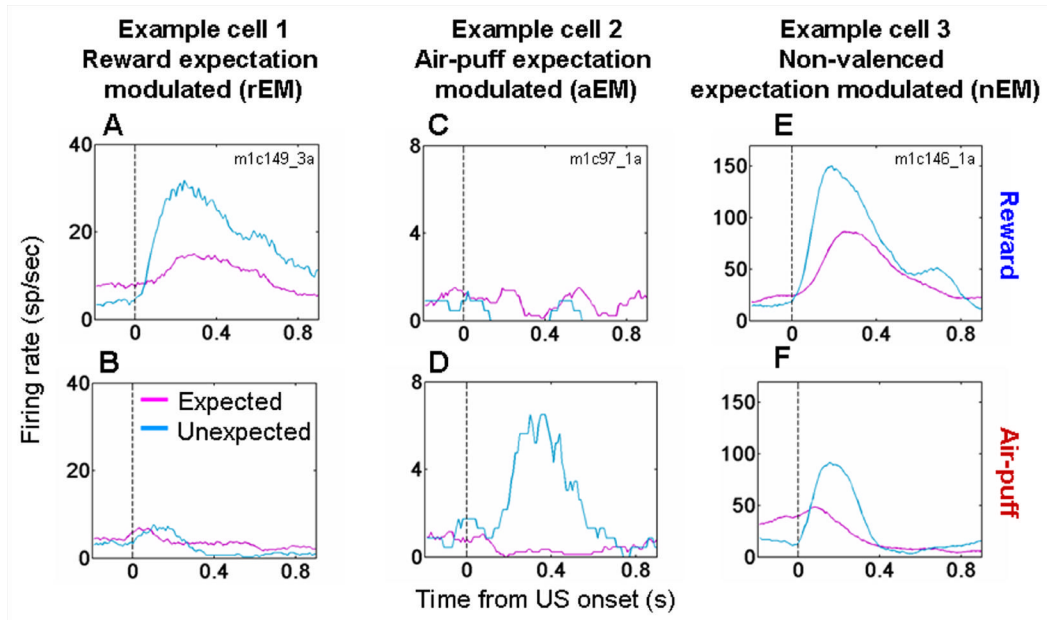


Figure 2. Expectation modulates neural responses to reinforcement in the amygdala

A,B. Peri-Stimulus Time Histograms (PSTHs) (which average neural responses as a function of time across trials) from one amygdala cell exhibiting a stronger response to reward (A) but not air-puff (B) when reinforcement was unexpected (random task, blue) compared to expected (trace-conditioning task, magenta). C,D. PSTHs from a neuron with stronger responses to unexpected air-puff (D) but not to unexpected reward (C). E,F. PSTH from a cell with stronger responses to both valences of unexpected reinforcement. All PSTHs smoothed with a 10 ms moving average.

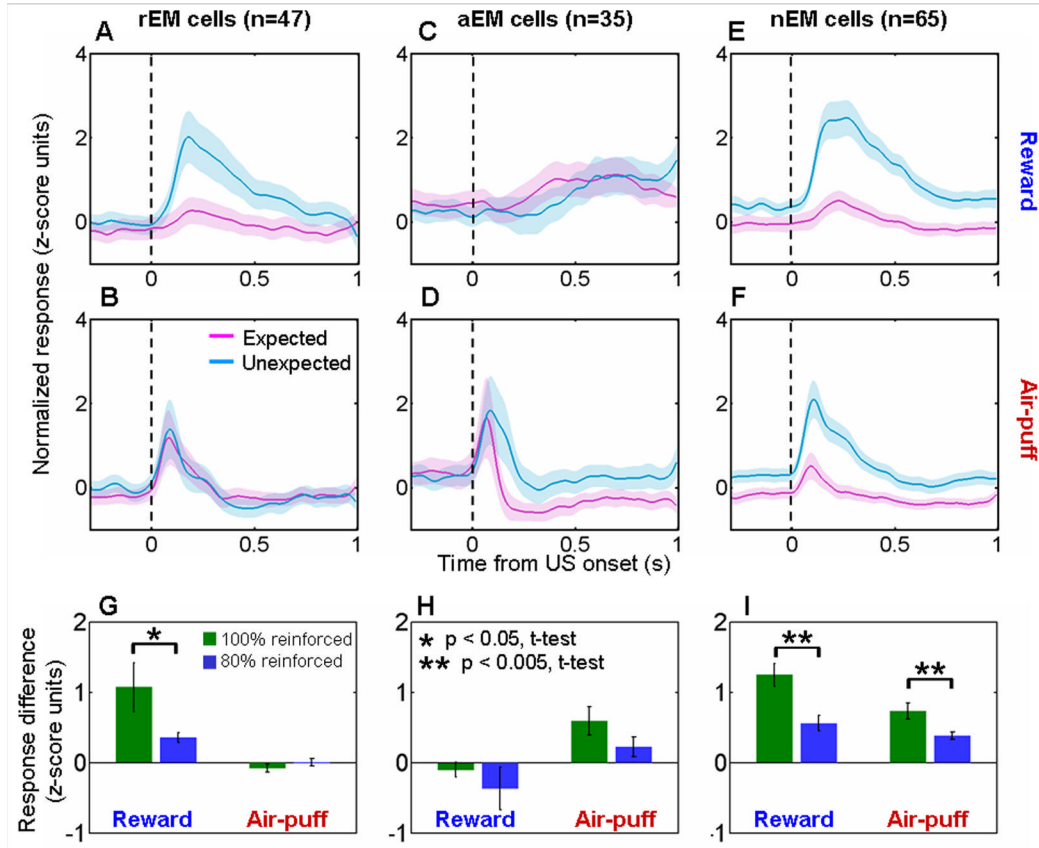


Figure 3. Valence specific and valence non-specific modulation of reinforcement responses by expectation

A–F. Normalized and averaged population PSTHs showing responses to expected (magenta, trace-conditioning task) and unexpected (blue, random task) rewards (A,C,E) and punishments (B,D,F). A,B, rEM neurons (n = 47). C,D, aEM neurons (n = 35). E,F, nEM neurons (n = 65). Shading, s.e.m. G–I. Difference in normalized activity for unexpected compared to expected rewards and air-puffs, shown separately for data from the 100% and 80% reinforcement probability trace-conditioning task. rEM cells (G), aEM cells (H), nEM cells (I). The difference in air-puff response seen in (H) does not achieve statistical significance (p = 0.14, t-test).

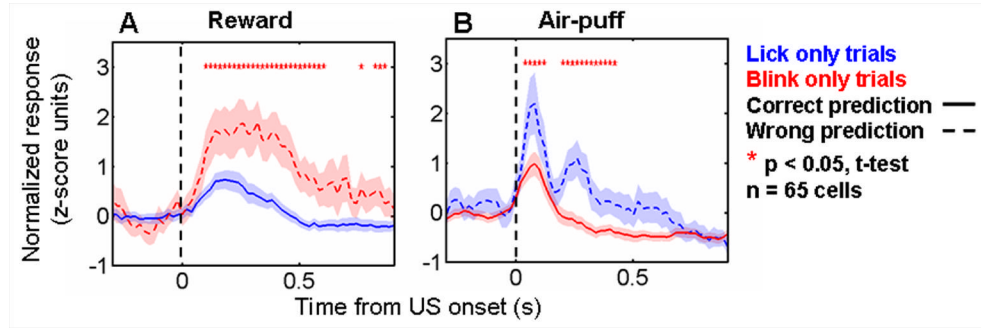


Figure 4. Expectation modulates responses to reinforcement within the trace-conditioning task
 A,B. Normalized and averaged PSTHs showing responses to rewards (A) and air-puffs (B) for the rEM, aEM, and nEM cells studied with the 100% reinforced trace-conditioning task. Trials are sorted according to whether monkeys expected air-puff (anticipatory blinking but not licking, red curves) or reward (anticipatory licking but not blinking, blue curves). On average, the responses to reinforcement were greater when monkeys incorrectly predicted the upcoming reinforcement. Red asterisks, activity significantly different in the two types trials in a 100 ms bin, $p < 0.05$, t-test.

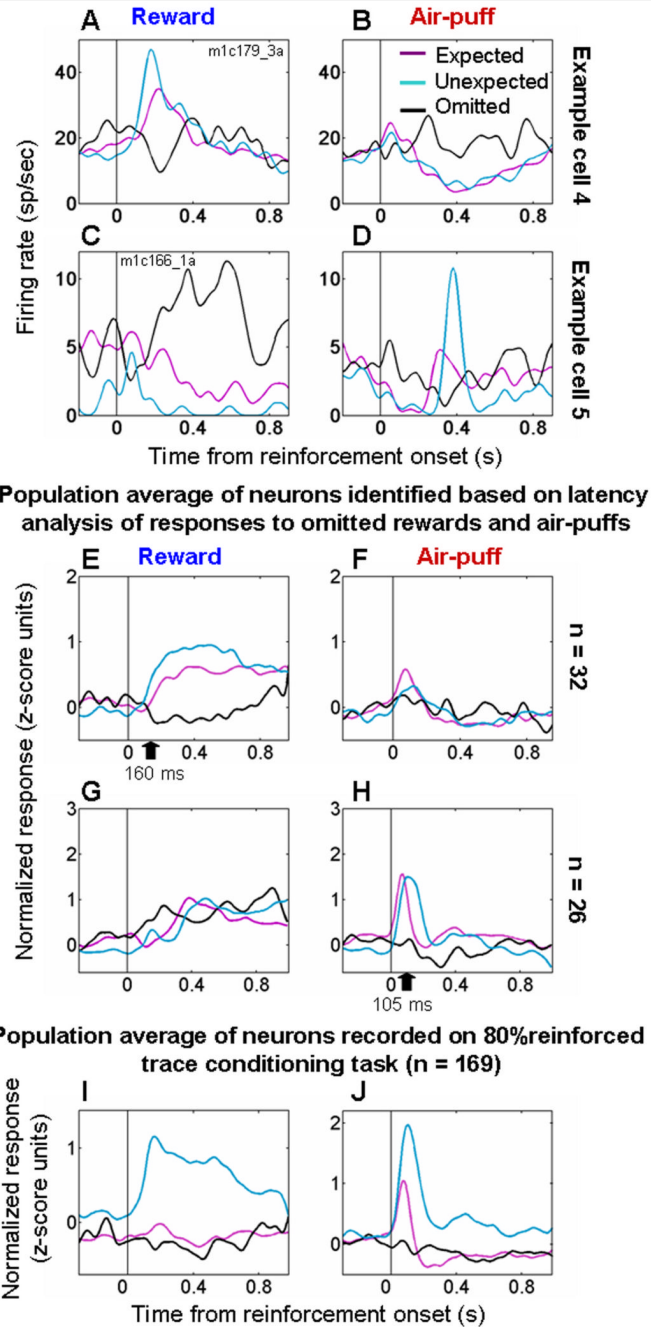


Figure 5. Neural responses to omitted reinforcement

A,B. PSTHs from an amygdala cell showing increased firing to unexpected rewards (A) and omitted air-puffs (B), and decreased firing to air-puffs (B) and omitted rewards (A). C,D. PSTH from an amygdala cell showing increased firing to unexpected air-puffs (D) and omitted rewards (C), and decreased firing to rewards (C) and omitted air-puffs (D). E,F. Normalized and averaged population PSTHs for neurons showing evidence of decreased firing to omitted rewards and/or increased firing to omitted air-puffs (n = 32, Table 1). G,H. Normalized and averaged population PSTHs for neurons showing evidence of decreased firing to omitted air-puffs and/or increased firing to omitted rewards (n = 26, Table 1). I,J. Normalized and averaged PSTHs showing responses to rewards (I) and air-puffs (J) from all cells recorded during both

the trace-conditioning task with 80% reinforcement probability and the random task ($n = 169$). For all plots: magenta, black, and blue curves, responses to expected, omitted, and random reinforcement, respectively. Black arrows in E and H, mean latency of decreased responses to omitted rewards and air-puffs. PSTHs smoothed with a 50 ms Gaussian.

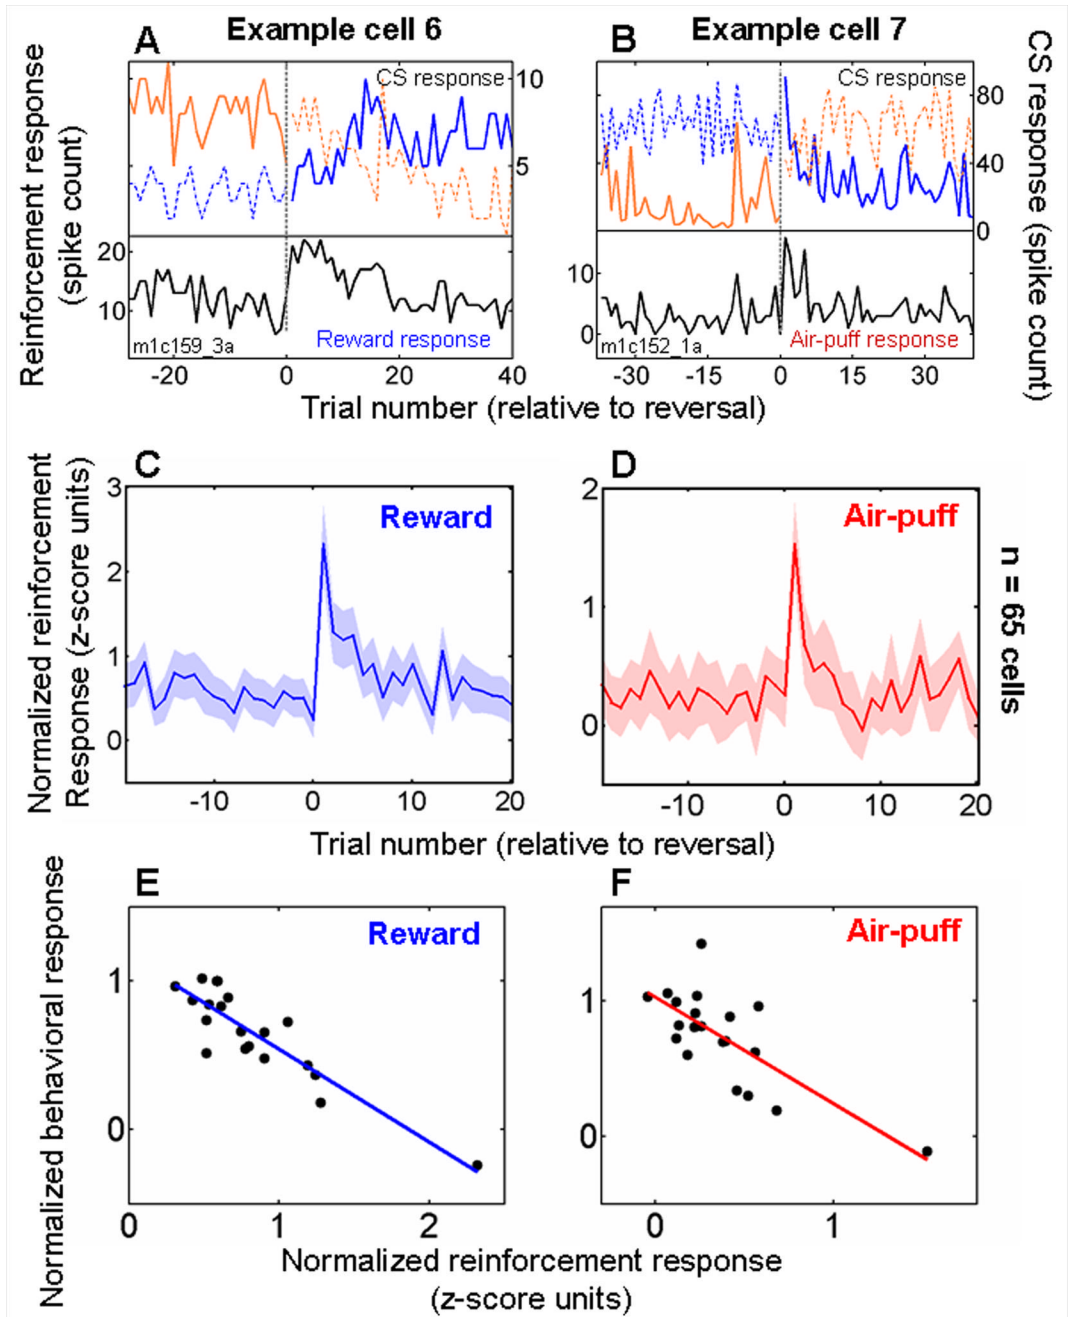


Figure 6. Responses to reinforcement decrease during learning

A. Data from one experiment in which a neuron responded more strongly to reward immediately after CS-US contingency reversal, and it rapidly changed its firing rate to the two CSs as reinforcement responses decreased. Black curve, response to reward. Orange and blue curves, visual stimulus interval responses to each of the two images, respectively. Solid and dashed lines, CS followed by reward or air-puff, respectively. B. Data from another neuron that responded more strongly to air-puff after image value reversal. Orange and blue curves, trace-interval responses to the two images, respectively. Labeling conventions the same as in A. C,D. Normalized and averaged neural responses to rewards (C) and air-puffs (D) plotted as a function of trial number relative to reversal for neurons recorded during the 100% reinforced

trace-conditioning task. C, rEM and nEM cells combined. D, aEM and nEM cells combined. Shaded regions, s.e.m. E,F. Correlation between mean responses to reward (E) and air-puff (F) on the 20 trials after reversal, taken from the data points in panels C,D, and normalized and averaged behavioral performance on the same trials (see Methods). In all panels, trial 1 is the first trial after reversal in image value.

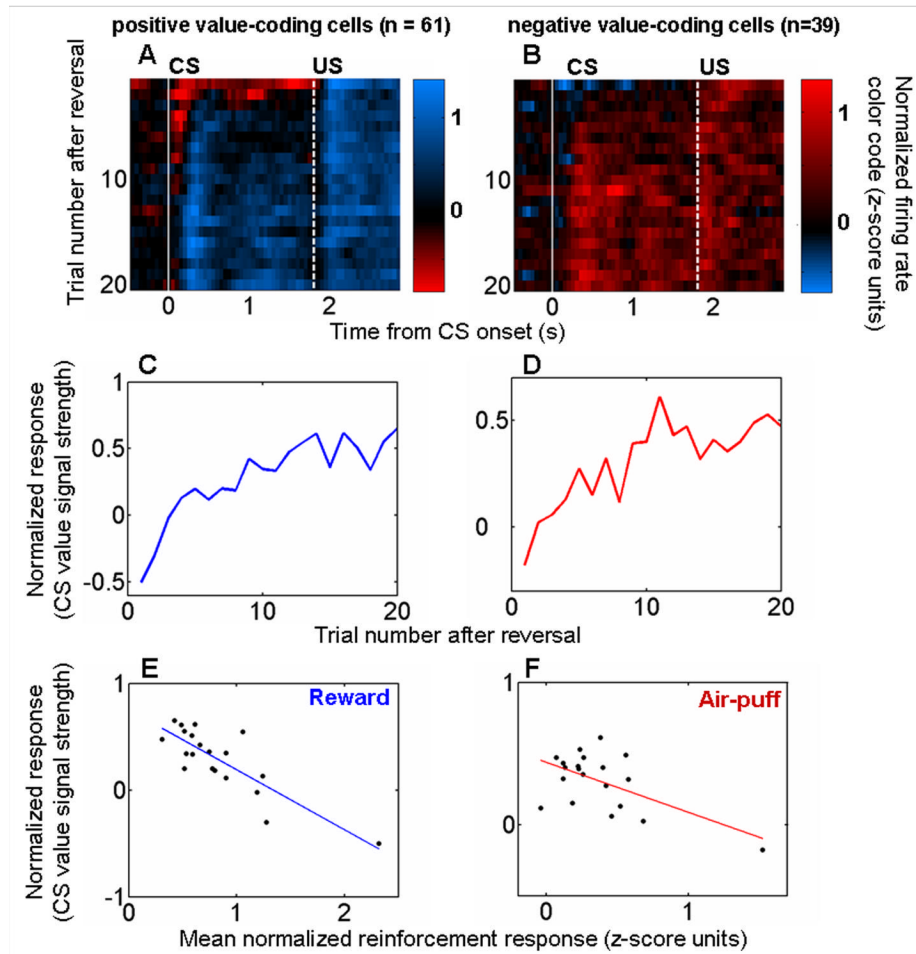


Figure 7. The decrease in responses to reinforcement is correlated with the evolving representation of value in the amygdala

A,B. Color maps showing the representation of value during reversal learning as a function of trial number and of time within trials. Positive (A) and negative (B) CS value-coding cells from the 100% reinforcement probability trace-conditioning task shown separately. Vertical white lines, CS onset (solid line) and US onset (dashed line). Each 200 ms bin shows the difference in firing rate between rewarded and punished trials; bins were advanced in 50 ms steps. Trial 1, first trial after reversal. Bin starting at time 0, interval from 0–200 ms after CS onset. C,D. CS value-coding extracted from the color map by taking the mean value in the interval between the white lines for neurons encoding positive (C) and negative (D) CS value. E,F. Recorded reinforcement responses (data points from Figs. 6C,D) plotted against the evolution of CS value coding (data points from Figs. 7C,D) during trials 1–20 after reversal. Regression lines are shown in blue and red.

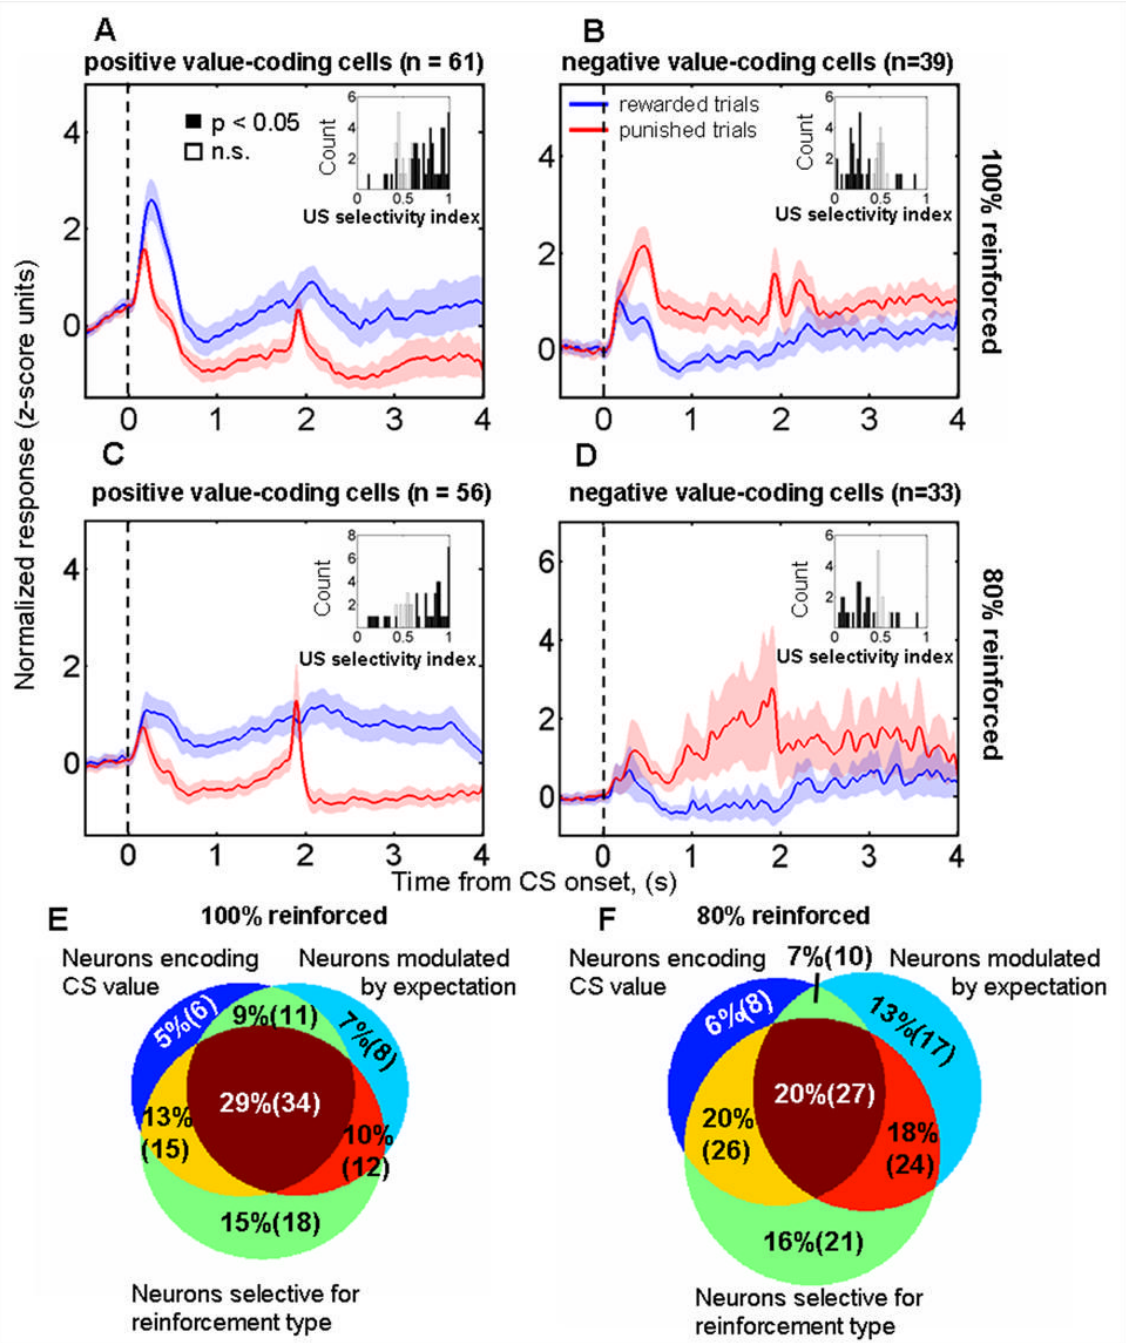


Figure 8. Amygdala neurons carry multiple signals

A–D. Normalized and averaged population PSTHs from neurons encoding positive (A,C) and negative (B,D) CS value during the trace-conditioning task with 100% (A,B) and 80% (C,D) reinforcement probability, respectively. Data aligned at CS onset; reinforcement occurred at either 1.8 or 1.85 secs depending upon the experiment. Inset histograms, ROC reinforcement selectivity indices for each of the positive (A,C) and negative (B,D) value coding neurons. Neurons encoding positive CS value tend to have reinforcement selectivity indices > 0.5, and neurons encoding negative CS value tend to have reinforcement selectivity indices < 0.5. E,F. Venn diagram showing the number of neurons classified as encoding CS value, as having reinforcement selective responses, and as having significantly stronger responses to

unexpected compared to expected rewards and/or punishments for the trace-conditioning task with 100% (E) and 80% (F) reinforcement probabilities, respectively.

Table 1

Categorization of cells according to their response to omitted reinforcement

Cells were categorized as increasing, decreasing, or not changing by comparing activity after an omitted reinforcement to activity in the last 1 sec of the trace interval.

Response to omitted reward	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	
Response to omitted air-puff	↑	↑	↑	↑	↑	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓
Number of cells	3	20	9	5	9	12	9	5	9	12	6	6	7	7	7	7	7	7	7	7	7



Excitatory response



Inhibitory response



No response

Blue – cells depicted on Fig. 5E,F

Red – cells depicted on Fig. 5G,H