Published in final edited form as:
Neuron. 2011 September 22; 71(6): 1127–1140. doi:10.1016/j.neuron.2011.07.016.

Different time courses for learning-related changes in amygdala and orbitofrontal cortex

Sara E. Morrison[1], **Alexandre Saez**[1], **Brian Lau**[1], and **C. Daniel Salzman**[1,2,3,4,5,6]

[1]Dept. of Neuroscience, Columbia University

[2]Dept. of Psychiatry, Columbia University

[3]Kavli Institute for Brain Sciences, Columbia University

[4]Mahoney Center for Brain and Behavior, Columbia University

[5]W.M. Keck Center on Brain Plasticity and Cognition, Columbia University

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

Abstract

The orbitofrontal cortex (OFC) and amygdala are thought to participate in reversal learning, a process in which cue-outcome associations are switched. However, current theories disagree on whether OFC directs reversal learning in the amygdala. Here, we show that during reversal of cues' associations with rewarding and aversive outcomes, neurons that respond preferentially to stimuli predicting aversive events update more quickly in amygdala than OFC; meanwhile, OFC neurons that respond preferentially to reward-predicting stimuli update more quickly than those in the amygdala. After learning, however, OFC consistently differentiates between impending reinforcements with a shorter latency than the amygdala. Finally, analysis of local field potentials (LFPs) reveals a disproportionate influence of OFC on amygdala that emerges after learning. We propose that reversal learning is supported by complex interactions between neural circuits spanning the amygdala and OFC, rather than directed by any single structure.

The prefrontal cortex (PFC) is important for the ability to respond flexibly and adaptively to changing environmental contingencies (Miller and Cohen, 2001). This process has been studied using reversal learning, a behavioral paradigm in which reinforcement contingencies are switched without warning. Lesions of the orbitofrontal cortex (OFC) – an area that comprises much of the ventral surface of the PFC – impair reversal learning in primates and rodents (Chudasama and Robbins, 2003; Fellows and Farah, 2003; Iversen and Mishkin, 1970; Izquierdo et al., 2004; Schoenbaum et al., 2003), although the impairment may be mitigated by limiting lesions to specific sub-regions or by using techniques that spare fibers of passage (Kazama and Bachevalier, 2009). The OFC has extensive bidirectional connections with the amygdala (Carmichael and Price, 1995; Ghashghaei et al., 2007; Stefanacci and Amaral, 2000, 2002), a subcortical brain structure that is important for a wide range of behaviors with an appetitive or aversive affective component, such as learning and

Author contributions: S.E.M. collected the data and analyzed the single unit data with advice from B.L. and C.D.S.; B.L. wrote the software to perform Granger analysis; A.S. analyzed the LFP data with assistance from B.L.; S.E.M. and C.D.S. designed the experiment and wrote the manuscript with input from B.L. and A.S.

updating cue-outcome associations (Murray and Izquierdo, 2007). Based on recent evidence, OFC has joined the amygdala as a key brain area in which both appetitive and aversive information are processed (Baxter and Murray, 2002; Belova et al., 2007; Belova et al., 2008; Hosokawa et al., 2007; Morrison and Salzman, 2009; Paton et al., 2006; Phelps and LeDoux, 2005; Salzman and Fusi, 2010; Salzman et al., 2007), and information about both valences often converges at the level of single neurons (Belova et al., 2008; Morrison and Salzman, 2009). Both the OFC and the amygdala are important for a variety of behavioral tasks that require the flexible reassignment of values to stimuli (Murray and Izquierdo, 2007; Murray and Wise, 2010).

Despite these findings, there is no consensus on the specific roles of the OFC and amygdala in the neural circuits underlying flexible behavior. Some authors have posited that OFC is specialized for supporting flexible behavior because it is better or faster than other brain areas, such as the amygdala, at signaling new cue-outcome associations – i.e., associations between a conditioned stimulus (CS) and an unconditioned stimulus (US) (Rolls et al., 1996; Rolls and Grabenhorst, 2008). Under this framework, during reversal learning the OFC is thought to rapidly detect the new CS-US associations and emit a “reversal signal” that facilitates the updating of CS-US contingencies in the amygdala. Other authors have suggested that the OFC plays a different role in reversal learning: maintaining the pre-reversal CS-outcome associations after reversal (Schoenbaum et al., 2009). In this model, the persistent representation of pre-reversal CS-US contingencies in OFC is thought to provide a basis for comparison with ongoing events, facilitating error-based updating in the amygdala and other areas.

We sought to test these hypotheses by simultaneously recording in amygdala and OFC in order to compare the onset and time course of neural changes during reversal learning. We reasoned that if OFC directs the reversal of associations in the amygdala – perhaps via a reversal signal – then the encoding of new CS-US associations should emerge more rapidly in the OFC than the amygdala during reversal learning. Alternately, if OFC maintains the previous CS-US associations during reversal learning, then the encoding of new associations should appear slowly in OFC and more rapidly in other brain areas such as the amygdala. Previous studies have identified neural activity that encodes the reinforcement associations of stimuli in primate OFC or amygdala separately (Belova et al., 2007; Belova et al., 2008; Bermudez and Schultz, 2010; Hosokawa et al., 2007; Morrison and Salzman, 2009; Nishijo et al., 1988; Padoa-Schioppa and Assad, 2006; Paton et al., 2006; Roesch and Olson, 2004; Rolls, 1992; Thorpe et al., 1983; Tremblay and Schultz, 1999). By recording from OFC and amygdala simultaneously, we were able to examine the time course of changing neural responses during and after reversal learning in both areas for two populations of neurons: those that respond more strongly to stimuli that predict reward (“positive” value-coding neurons) and neurons that respond more strongly to stimuli that predict aversive events (“negative” value-coding neurons).

Surprisingly, we found marked differences between positive and negative cell populations in the relative dynamics of their changing signals: negative value-coding cells “learned” faster in amygdala, while positive value-coding cells learned faster in OFC. Only after completion of reversal learning was there evidence consistent with the idea that one brain area (OFC) may drive processing in the other (amygdala). Thus, the debate concerning which area directs learning in the other area must be expanded to account for valence-dependent differences in dynamics. Instead of being driven by a particular brain area, our data indicate that reversal learning is more likely driven by complex interactions between neural circuits spanning both brain areas.

Results

Task and behavior

We recorded single neuron activity and LFPs simultaneously from the OFC and amygdala of two monkeys performing a Pavlovian trace-conditioning task with a reversal learning component (Fig. 1a). In each session, monkeys learned the associations of two novel, abstract visual CSs; one CS, the “positive” image, was followed by a rewarding US (liquid reward), while the other CS, the “negative” image, was associated with an aversive US (an air-puff to the face). We monitored monkeys' learning by tracking the amount of licking at the reward spout in expectation of reward and eye closure (“blinking”) in expectation of air-puff. After monkeys learned the initial reinforcement contingencies, we reversed the associations of the positive and negative CSs without warning, and monkeys learned the new contingencies, as indicated by changes in licking and blinking after reversal.

We determined the onset of monkeys' learning-related behavioral changes using a change point test (Gallistel et al., 2004; Paton et al., 2006). In the example shown in Figure 1b,c, anticipatory licking and blinking rates begin to change quickly after the reversal of reinforcement contingencies, although the monkey did not switch to the appropriate behavior until it had experienced at least one pairing of each image with its new reinforcement outcome. Across experiments, monkeys were no more likely to lick on the first positive trial after reversal, or to blink on the first negative trial, after first experiencing a trial of the other type (Fig. 1d,e; Wilcoxon, $p > 0.5$ for both), and this did not change with experience (comparison between first and second half of recording sessions; chi-square test, $p > 0.05$). Thus, monkeys do not appear to develop a working concept of reversal to guide their behavior on this task (a higher-level strategy); rather, they learn reversals by experiencing each cue paired with its associated outcome.

Amygdala and OFC neural responses to reversals of reinforcement contingencies

We recorded from 217 neurons targeting area 13 of the right OFC (Ongur and Price, 2000), and 222 neurons in the right amygdala (Fig. 2). We used a two-way ANOVA with factors for CS value (positive or negative) and CS identity to detect neurons that have activity reflecting the association of CSs with reward or air-puff. Many cells in each brain area showed a significant main effect of CS value on neural firing in the CS and/or trace intervals ($n = 86$ in each area, $p < 0.01$). We further categorized these 172 cells by their preference for CS valence: neurons that fired more strongly in response to the positive or negative CS were designated “positive” or “negative” value-coding cells, respectively. We identified substantial populations of positive and negative value-coding cells in each brain area (41 positive cells and 45 negative cells in OFC; 27 positive cells and 59 negative cells in amygdala). These value-coding cells responded more strongly either to CSs predicting reward (positive cells) or air-puff (negative cells) both before and after reversal (Fig. 3).

We applied the change point test to neural activity to determine whether the activity of these value-coding cells could underlie the behavioral changes seen following reversal (Fig. 1b,c). Figure 4a-d illustrates the responses of a positive value-coding cell from OFC recorded during the behavioral session depicted in Fig. 1b,c. The neural response to Image 1 decreased as its associated outcome changed from positive to negative (Fig. 4a,c), and the response to Image 2 increased as the image changed from negative to positive (Fig. 4b,d). For each image, the change in neural response starts to occur at the same time as one or both shifts in licking and blinking behavior. Using this procedure, we identified a trial number corresponding to the onset of the change in activity of each value-coding neuron, and we compared it to when licking and blinking behavior began to change upon reversal for the same image. For each group of value-coding cells, neural change points either were not

different from behavioral change points (Fig. 4f-h; sign rank test, $p > 0.05$) or were slightly earlier than behavioral changes (Fig. 4e; sign rank test, $p < 0.05$). The change point differences did not differ between groups (Wilcoxon, $p > 0.05$ for all comparisons). Thus, neural activity in OFC – as well as in amygdala (Paton et al., 2006) – could contribute to reversal learning.

Relative dynamics of value-coding cells during reversal learning

We next examined the differences in the *time course*, as opposed to onset, of the neural changes among positive and negative value-coding cells in OFC and amygdala. An unexpected pattern of differences emerged: among positive value-coding cells (Fig. 4i), OFC neurons exhibited a larger change in activity from the 12 trials before to the 12 trials after the change point (significant for positive trials; Wilcoxon, $p < 0.05$). However, among negative value-coding cells, amygdala neurons exhibited a larger change in activity than OFC neurons (Fig. 4j; Wilcoxon, $p < 0.05$ for both trial types). Thus, positive and negative value-coding neurons in amygdala and OFC appear to "learn" at different rates relative to each other.

To examine this apparent difference in time course, we calculated a "difference index," – the difference in average normalized neural response to the two CSs – over the trials following reversal, using a six-trial moving window (Fig. 5a,b). We quantified the time course of the difference indices for each neural population by calculating a scale-adjusted latency or "threshold" for a fitted sigmoid curve, representing the trial number when the curve reached a specific percentage of its maximum value (see Experimental Procedures). The curves reached this threshold at significantly different times for amygdala and OFC (F-test, $p < 0.001$), and this difference had an opposite sign for positive and negative value-coding cells. The scale-adjusted latency was 15.15±0.68 trials for positive OFC and 18.24±1.07 trials for positive amygdala cells; and 16.97±2.12 trials for negative OFC and 10.03±0.65 trials for negative amygdala cells (margins of error are based on 95% prediction intervals). Thus, for positive cells, the OFC group changed more rapidly than the amygdala group, but for negative cells, the amygdala group changed more rapidly than its counterpart in OFC.

The "difference index" provides a straightforward way to analyze the time course of changing neural responses, but it does not take into account the possible contributions of other factors, such as the sensory characteristics of images. Therefore, we used a sliding ANOVA analysis to examine how the unique contributions to neural activity of image identity and image value change after reversal. For each value-coding cell, we calculated the average proportion of explainable variance in neural activity that was due to image value – a "contribution-of-value index" – using data from 6 trials of each type before and after reversal (24 total trials), and thereafter sliding the post-reversal 6-trial window in 1 trial steps (i.e., using trials 2-7, then 3-8, etc.). As before, we fit sigmoid functions to the index and tested for differences in latency between the curves. This analysis, shown in Figure 5c,d, confirmed the findings of the difference index analysis. The contribution of image value to the activity of positive OFC cells increased more rapidly and reached a plateau 6.4 trials earlier than that of positive amygdala cells (Fig. 5c); conversely, the contribution of value to the activity of negative amygdala cells reached a plateau 13.7 trials sooner than that of negative OFC cells (Fig. 5d; F-test, $p < 0.001$ in both cases).

Finally, we found that the average onset of changes in neural activity and behavior was similar, consistent with the change-point analysis (see Fig. 4). These data indicate that although neurons in both brain areas begin to update their signaling fast enough to drive the onset of behavioral learning, the dynamics of learning differ. The appetitive system (comprised of positive value-coding neurons) changes more rapidly in OFC, but the aversive system (comprised of negative value-coding neurons) updates more rapidly in amygdala.

Multi-dimensional regression analysis of neural dynamics across reversal

We next examined how the time course of value-related signals *within* trials changes during learning (Fig. 6; Fig. S1). Here, as in Figure 5c,d, we calculated a contribution-of-value index in 6-trial windows stepped by 1 trial over the reversal learning period; but now we applied the analysis to neural activity in 200 ms bins advanced in 20 ms increments across the trial. For positive OFC cells and negative amygdala cells, the contribution-of-value index achieves significance ($p < 0.01$ for at least 3 consecutive bins) shortly after CS onset during the first trial after reversal (Fig. 6a,d, top rows in each panel); in contrast, the value index for negative OFC cells and positive amygdala cells reaches significance later in the trial (or, for negative OFC cells, not until the time of reinforcement). For the first data point after reversal, negative amygdala cells reach this threshold significantly earlier in the trial than negative OFC cells (permutation test; $p = 0.01$), while positive OFC cells reach threshold earlier than positive amygdala cells, with the difference approaching significance ($p = 0.056$).

Thus, the population responses of positive OFC and negative amygdala neurons accurately encode the associations of images with reward or air-puff before the population responses of positive amygdala neurons and negative OFC neurons have significantly adapted to the new reinforcement contingencies. These timing differences are unlikely to be accounted for by differences in the contribution-of-value index magnitude: the maximum index magnitudes are similar for negative cells in OFC and amygdala, while the peak magnitude is actually higher for positive amygdala cells than for positive OFC cells. This analysis also reveals that significant contribution-of-value indices shift to earlier times on later trials compared to the earliest trials after reversal, resembling the back-propagation of value signals predicted by temporal difference models of reinforcement learning (Sutton and Barto, 1998).

In the multidimensional sliding ANOVA, lingering pre-reversal reinforcement associations are absorbed by the image identity factor. This is illustrated in Figure 7a-d (see also Fig. S2), in which the putative contribution of image identity – a contribution-of-image index – is plotted as a function of time and trial number. The image identity term captures a large amount of the variance in neural activity immediately after reversal, and this effect is especially salient in the two slower-changing groups. Consistent with our previous findings (Morrison and Salzman, 2009; Paton et al., 2006), some image identity encoding remains in all populations even after learning has taken place. Finally, the interaction term made a relatively small contribution that did not differ systematically across groups (Fig. 7e-h), and therefore cannot explain group differences in learning. The neuronal groups were not different with regard to the proportion of cells with a significant interaction effect (chi-square test, $p > 0.1$), nor with regard to the average magnitude of the interaction effect (t-test, $p > 0.05$ for all comparisons).

After learning, OFC encodes value with a shorter latency than amygdala

Positive and negative value-coding cells in OFC and amygdala have strongly divergent time courses of activity over the trials just following reversal; but how does the time course of value-coding activity compare in the two areas *after* learning? Further examination of when the value index reaches significance (asterisks in Fig. 6), focusing on trial 6 and later after reversal, reveals that both subgroups in OFC seem to signal outcome expectations earlier during the trial than their counterparts in the amygdala (latencies to significance as established by Fisher p-value < 0.01 were significantly earlier in OFC than amygdala; Wilcoxon, $p < 0.01$).

Focusing on post-learning trials, we examined the contribution of image value to each cell's activity throughout the trial. Figure 8 illustrates that OFC neurons as a population are

quicker to encode image value, regardless of their positive or negative CS value preference. Compared to amygdala, we found relatively more OFC neurons with the earliest significant value contributions – less than 150 ms following cue onset (χ^2 test, $p < 0.05$). Moreover, the average contribution-of-value signal reached significance for the OFC earlier than amygdala by about 40-60 ms for both positive and negative cells (Fig. 8e,f; F-test, $p < 0.01$). We fit sigmoid curves to the early portion (first 500 ms after image onset) of the average contribution-of-value signal for each group; in both cases, the time to reach the scale-adjusted threshold for the OFC group was significantly shorter than that for the amygdala group (F-test, $p < 0.01$). Thus, in contrast to the robust differences between positive and negative neurons in the timing of the value signal *during* learning, OFC neurons encoded image value more rapidly during the trial than amygdala neurons *after* learning.

Granger causality analysis of local field potentials

The post-learning timing differences in the single unit data suggest that OFC might preferentially influence signaling in amygdala after learning. We looked for evidence to support this notion by examining LFPs recorded simultaneously in OFC and amygdala. We recorded LFPs from 853 sites in two monkeys, yielding 1282 simultaneously recorded OFC-amygdala pairs. We estimated the directed influences between OFC and amygdala using Granger causality, which measures the degree to which the past values of one LFP predict the current values of another (see Experimental Procedures). Looking at a broad range of frequencies (5-100 Hz), we computed Granger causality in sliding windows across the trial for all post-reversal trials. We found that the average influence in both directions – OFC-to-amygdala and amygdala-to-OFC – was significantly elevated during the image presentation and trace intervals (Wilcoxon, $p < 0.01$; Fig. 9a), indicative of a task-related increase in the exchange of information between these areas. Granger causality was generally significantly greater in the OFC-to-amygdala direction (Fig. 9a, blue line) than in the amygdala-to-OFC direction (Fig. 9a, green line) throughout much of the trial (asterisks; permutation test, $p < 0.05$).

We also examined whether Granger causality changes as a function of learning. For each time window across the trial, we subtracted the Granger causality in the amygdala-to-OFC direction from the causality in the OFC-to-amygdala direction, yielding a measure of the relative strength of directed influence between the LFPs from each brain area. We calculated this directional difference separately for 6 trials of each type at the beginning of the reversal block and the end of the block (Fig. 9b). We found that, averaged over a wide frequency band, Granger causality *during* the learning process is actually stronger in the amygdala-to-OFC direction; causality in the OFC-to-amygdala direction becomes predominant only *after* learning has taken place, consistent with the single unit findings. Moreover, this effect appears to be related to task engagement, as it emerges most prominently after CS onset. Finally, averaging across time during the trial, we found that this learning-related directional effect is robust across frequencies ranging from the beta band (12-25 Hz) through the lower gamma band (25-40 Hz; Fig. 9c).

Discussion

We compared the dynamics of simultaneously recorded neural signals in amygdala and OFC while monkeys performed a reversal learning task with both appetitive and aversive reinforcement contingencies. We found that neurons in amygdala and OFC exhibited different relative time courses when updating representations of impending reinforcement. Both amygdala and OFC neurons began to update their representations rapidly after a reversal of reinforcement contingencies; but the rates of change in amygdala and OFC depended upon the preferred valence of neurons. Positive value-coding cells in OFC adapted to the reversed reinforcement contingencies significantly more quickly than positive value-

coding cells in the amygdala; conversely, negative value-coding cells in OFC adapted more slowly than their counterparts in the amygdala. These data suggest that distinct sequences of neural processing lead to the updating of activity in the appetitive and aversive neuronal subpopulations.

It has long been theorized that the amygdala is specialized, at least in part, for responding to aversive events and generating the associated responses of withdrawal, avoidance, or defense (Morrison and Salzman, 2010; Phelps and LeDoux, 2005). Consistent with this idea, optical stimulation of pyramidal neurons in the lateral amygdala can act as a US to produce fear conditioning (Johansen et al., 2010). Thus, the fast adaptation of neural activity in the subcortical aspect of the aversive network – i.e., negative value-coding cells in the amygdala – might reflect the evolutionary preservation of a rapid-detection system for possible threats. Note, however, that negative value-coding cells do not exclusively encode aversive events, nor do positive value-coding cells respond only to rewarding events; rather, information about both rewarding and aversive cues and outcomes often converges in both positive and negative cells in the amygdala (Belova et al., 2008) and OFC (Morrison and Salzman, 2009).

Despite this convergence, most neurophysiological studies in non-human primates have focused on reward processing. These studies have shown that neural activity in a number of brain areas evolves in concert with learning motivated by reward: from subcortical structures like the striatum (Brasted and Wise, 2004; Pasupathy and Miller, 2005; Tremblay et al., 1998) and hippocampus (Wirth et al., 2003), to cortical areas ranging from the frontal eye fields, supplementary eye fields and premotor cortex (Brasted and Wise, 2004; Chen and Wise, 1995a; Chen and Wise, 1995b; Chen and Wise, 1996; Mitz et al., 1991) to various sub-regions of prefrontal cortex (e.g., Pasupathy and Miller, 2005), including OFC (Tremblay and Schultz, 2000). The current work builds on these prior studies in several ways, including the addition of aversive stimuli and the use of a Pavlovian rather than instrumental task.

In contrast to nearly all primate studies, studies in rodents have examined reversal learning in both the reward and aversive domains, and suggest that complex interactions between amygdala and OFC occur during reinforcement learning (Saddoris et al., 2005; Schoenbaum et al., 1999; Schoenbaum et al., 2009; Stalnaker et al., 2007). For example, Schoenbaum and colleagues have shown that amygdala lesions impair the development of cue-selective activity in OFC that normally develops as rats learn about reversed reinforcement contingencies (Stalnaker et al., 2007). In a complementary study, the authors reported that OFC lesions impede the ability of the amygdala to adjust its firing to a CS after a reversal (Saddoris et al., 2005). These and other experiments have led the authors to suggest that OFC plays a prominent role in representing reinforcement expectations, even when those expectations are no longer correct (Schoenbaum et al., 2009). By retaining a representation of the pre-reversal outcome expectancies, OFC activity could provide inputs essential for the generation of prediction error signals in other brain areas – such as the ventral tegmental area – which could in turn direct flexible neural encoding in the amygdala and elsewhere.

Our findings do not support the idea that OFC neurons, as a whole, encode pre-reversal outcome expectation for a longer period than their counterparts in the amygdala, as has been proposed (Schoenbaum et al., 2009). We showed that negative value-coding neurons – those that respond preferentially to stimuli that are linked with aversive events – are indeed slower to shift their representation of stimulus-outcome contingencies in OFC than in the amygdala. On the other hand, positive value-coding neurons fully reverse their encoding more rapidly in OFC than in the amygdala. Thus, the question of which brain area is “in charge” during reversal learning is almost certainly the wrong question. Instead of a simple feed-forward

process – one brain area learning about the reversal and sending instructive signals to another – these data suggest a more complex neural circuit, in which appetitive and aversive neural networks participate in a multipart interchange of information during learning.

There are a number of reasons why the current findings paint a more complex picture than prior studies of appetitive and aversive reversal learning in rodents (Saddoris et al., 2005; Schoenbaum et al., 1999; Schoenbaum et al., 2009; Stalnaker et al., 2007). Prior studies have not separately analyzed the dynamics of neuronal subpopulations that prefer positive or negative valence, which we propose might participate in distinct appetitive and aversive networks. Moreover, the current study is the first, to our knowledge, to utilize simultaneous recording of individual neurons in the amygdala and OFC. Because simultaneous recordings are performed in the same subjects under the same behavioral conditions, the technique is advantageous for analyzing timing differences between neural signals in two different brain areas. Finally, the anatomical areas referred to as OFC in rodents may not directly correspond to OFC as it has been studied in primates. We and other primate neurophysiologists have typically investigated area 13 and other granular and dysgranular parts of OFC (Padoa-Schioppa and Assad, 2006; Roesch and Olson, 2004; Tremblay and Schultz, 1999); however, a direct homologue to rodent OFC is more likely to be found in the agranular areas located posterior to typical recording sites in monkeys (Murray and Wise, 2010; Wise, 2008). A distinctive feature of primate neuroanatomy is an expansion of prefrontal areas such as OFC, involving the emergence of dysgranular and granular cortex that are absent in rodents, and concomitant elaboration of interconnectivity with the amygdala (Ghashghaei et al., 2007; Ongur and Price, 2000; Wise, 2008). This elaboration of PFC may support enhanced cognitive flexibility, contributing to the more complex social, cognitive, and behavioral repertoire of primates (Wise, 2008).

Other authors have argued that OFC is specialized for supporting flexible behavior because it is better or faster than other brain areas, such as the amygdala, at rapidly signaling new stimulus-outcome associations (Rolls and Grabenhorst, 2008). Early work by Rolls and colleagues seemed to show that a larger percentage of neurons in OFC, compared to amygdala, shift their cue selectivity upon reversal, and that they do so almost immediately, whereas amygdala neurons change their selectivity far more slowly if at all (Sanghera et al., 1979; Thorpe et al., 1983). Under this schema, OFC would first detect reversal, and then send a “reversal signal” to other brain areas, directing them to adjust their representations. However, this model is not supported by recent work showing rapidly changing signals in the amygdala during reversal learning, nor by the current work, which points to more complex interactions underlying reversal learning. Overall, the current findings challenge existing theories of OFC-amygdala interaction during reversal learning: whether the idea that OFC simply transmits updated associative information to the amygdala and other brain areas (Rolls and Grabenhorst, 2008), or that OFC holds on to pre-reversal expectations to form a basis for prediction error signal computations (Schoenbaum et al., 2009). Instead, OFC and amygdala may be best understood as comprising at least two neural subsystems – an appetitive system and an aversive system – which exhibit different temporal dynamics. These dynamics may have arisen out of evolutionary pressure to rapidly detect and respond to threats, or to approach potential new rewards with caution.

One distinctive role for OFC may come into play only *after* learning about reinforcement contingencies. After learning, we found that OFC neurons consistently signal impending reinforcement more rapidly than amygdala. This may reflect the primary role of PFC in executive functions and emotional regulation with regard to both rewarding and aversive experiences. Consistent with this idea, Granger causality analysis of LFP signals suggests a greater influence of OFC processing on amygdala than the opposite, an effect that emerges with learning. This effect was especially prominent in the beta frequency band, which has

been suggested to be well-suited for long-range interactions between brain areas (Kopell et al., 2000). Importantly, despite the directional effect, the analysis of LFP data suggests continuous and dynamic reciprocal interactions between OFC and amygdala during task engagement. We note that this finding does not exclude the possibility that a third brain area – such as another area of PFC – could influence both OFC and amygdala in a consistently asymmetric manner.

Likewise, these findings do not preclude the participation of other brain areas in reversal learning. The striatum is a major output target for both OFC and amygdala (Carmichael and Price, 1995; Fudge et al., 2002; Haber et al., 1995), and thus a likely site of interaction and integration of signals from the two brain areas. The striatum itself contains neurons that signal changing reinforcement contingencies during instrumental tasks (Brasted and Wise, 2004; Pasupathy and Miller, 2005; Tremblay et al., 1998), and one study reported that signals update even more rapidly in striatum than in PFC upon repeated reversals of the visuomotor reinforcement contingencies associated with familiar stimuli (Pasupathy and Miller, 2005). This raises the possibility that as stimulus sets and their associated possible reinforcements upon reversals become increasingly familiar, the striatum may assume a more prominent role in directing adaptive behavioral responses to the changing environment.

Our results indicate that reversal learning likely involves complex interactions between anatomically intermingled neural circuits spanning the amygdala and OFC, and perhaps other brain structures. Fully testing the predictions of the current work, however, may require the development of techniques that can specifically manipulate activity in appetitive or aversive neurons in targeted structures – in contrast to, e.g., inactivating the entire structure – during task performance. Finally, cognitive flexibility refers to more than simply reversing the associations of outcomes with stimuli. Often, there is no one-to-one mapping between stimuli and reinforcement; rather, organisms often must learn about behavioral, environmental and social contexts in order to accurately predict reinforcement based on sensory cues. Future experiments must determine whether the same relationship holds between the appetitive and aversive networks in amygdala in OFC when organisms must link together complex combinations of information to adapt flexibly to the environment.

Experimental procedures

Our methods for electrophysiological recording have been described previously (Morrison and Salzman, 2009). All 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.

Behavioral task

We used a trace conditioning procedure to induce learning about the associations with reinforcement of three novel abstract images (fractal patterns) in every experiment (Fig 1a). In each trial, monkeys foveated a central fixation point for 1 s, and then, while maintaining fixation, viewed an image for 300 ms (monkey R) or 350 ms (monkey L). During fixation, we required the monkey to maintain its gaze within 3.5° of the fixation spot, as measured with an infrared eye tracker (ASL, Applied Science Laboratories). Images occupied an 8° square at fixation. After image viewing, a 1.5 s trace interval with no fixation requirement ensued. After the trace interval, we delivered, with 80% probability, a large liquid reward after the positive image (1.0/1.8 mL of water for monkeys L/R respectively), or a 100 ms 40-60 psi air-puff directed at the monkey's face after the negative image. Air-puffs were directed at one of two possible locations on the monkey's face, chosen randomly on every trial. A third “weak positive” image was followed on 80% of trials by a smaller reward. All

three trial types were presented in pseudo-random order, separated by a 3 s intertrial interval. We waited for a variable number of trials after monkeys learned the initial reinforcement contingencies before, without warning, reversing the images paired with large reward and air-puff. There was one reversal per session, which occurred after 30–60 presentations of each stimulus. The image associated with small reward kept the same reinforcement contingencies.

Data collection

Behavioral measures—We assessed monkeys' anticipatory licking and blinking to determine whether they had learned the associations between CSs and USs (Morrison and Salzman, 2009). To measure licking, we placed the reward delivery tube ~1 cm from the monkey's mouth, and measured when the tongue interrupted an infrared beam passing between the mouth and the reward delivery tube. We measured anticipatory blinking using an infrared eye tracker, which outputs a characteristic voltage when eye position is lost. The loss of the eye position signal corresponded to eye closures as visualized by an infrared camera.

Electrophysiological recordings—We recorded neural activity from 217 neurons in the right OFC and 222 neurons in the right amygdala of two rhesus monkeys (*Macaca mulatta*): 141 OFC cells and 136 amygdala cells from a 6 kg female (monkey L); 76 OFC cells and 86 amygdala cells from a 13 kg male (monkey R). In each recording session, we individually advanced up to four tungsten microelectrodes (impedance: ~2 MΩ; FHC Instruments) into each brain area using a motorized multielectrode drive (NAN). 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 analyzed all well isolated neurons; monkeys performed a fixation task or no task during the search for well isolated neurons.

Local field potentials (LFPs) were recorded from every electrode, regardless of whether a single unit was present, resulting in 853 LFP sites recorded. The LFP signal was extracted from the raw signal by band-pass filtering from 0.7 to 170 Hz and sampled at 1 kHz. It was then notch-filtered to exclude 60 Hz line noise.

Reconstruction of recording sites—The neuronal sample was taken from overlapping regions of OFC and amygdala in the two monkeys. Based on comparison of MR images with a monkey brain atlas (Paxinos, 2000), we tentatively assign our recording sites primarily to areas 13m and 13a of OFC; for a small number of neurons, recording sites may have extended to area 14o (using the subdivision of OFC by Ongür and Price, 2000). Recording sites in amygdala were probably located primarily in the lateral and basal nuclei, with fewer recordings likely in the central and accessory basal nuclei.

Data Analysis

Identification of neural responses to conditioned stimuli—Classification of cells was performed using spike data from two time intervals: the CS interval (90–440 ms and 90-390 ms after CS onset for monkeys L and R respectively) and the trace interval (90–1500 ms after CS offset). We selected 90 ms after CS onset as the beginning of the CS interval because > 90% of the latencies in each brain area were > 90 ms (see Morrison and Salzman, 2009).

Classification of cells as value-coding—We performed a two-way ANOVA with image value and image identity as main factors. The ANOVA was performed separately on spike counts from the CS and trace intervals for each cell, excluding 5 trials of each type

from the start of the initial and reversal blocks. If there was a significant effect of image value in either or both intervals ($p < 0.01$), the cell was classified as value-coding. We found very few cells (4 OFC, 2 amygdala) that had opposite value preferences in the CS and trace intervals; we excluded these cells from further analysis. A relatively small number of cells showed a significant interaction effect ($p < 0.01$) without a significant main effect of value or image identity (21 OFC cells, 30 amygdala cells). These cells were not categorized as value-coding.

Change point analysis—We used a change point test (Gallistel et al., 2004) to determine the trial number at which licking, blinking, and neural activity began to change in relation to reversals. We applied the test to activity from either the CS or trace interval, depending on which had the largest contribution of value by ANOVA. Change point analysis was performed on scored behavior: trials in which licking (or blinking) occurred in the last 500 ms of the trace interval were scored as 1; other trials were scored as 0. We identified change points using a threshold of $p < 0.05$, correcting for multiple comparisons. If multiple change points were identified, we used the change point closest to reversal. In addition, we calculated the normalized activity (Z-scored with reference to baseline firing rate) for each value-coding cell before and after the identified change point (using 12 trials of each type before and after the change point) and averaged it together with all cells encoding the same valence in the same brain area (Fig. 4i,j).

Sliding difference analysis—We computed a “difference index” comparing each neuron's response on each trial to the two images that reverse reinforcement contingencies (Fig. 5a,b). We examined firing rates in the 90-590 ms after CS onset, normalizing the firing rates by subtracting the baseline firing rate and dividing by its standard deviation. For each value-coding cell, starting 10 trials of each type before reversal, we calculated the difference in the average normalized response to the two CSs in windows of 6 trials, stepped by 1 trial. For positive value-coding cells, we subtracted the response to Image 1 (which changes from positive to negative) from the response to Image 2 (which changes from negative to positive); for negative value-coding cells, we subtracted the response to Image 2 from the response to Image 1, so that all difference indices change in the same direction across reversal. We then averaged the difference indices for each trial across all cells in each group, and fit the average difference indices with a Weibull function (Eq. 1). Finally, for display, we normalized the fit functions and the data points by subtracting the lower asymptote of the Weibull function and dividing by the upper asymptote. Results were significant and went in the same direction for both monkeys, so the data were combined.

Sliding regression analysis—To quantify the time course of neural changes after reversal, we applied a sliding two-way ANOVA with main factors of image value and image identity on spike counts from a time window 90-590 ms after CS onset (Fig. 5c,d). This window exhibited the strongest divergence of activity among neuronal subgroups, but other time windows, including the entire CS and trace interval, produced similar results. For each value-coding cell, we performed the sliding ANOVA using data from the last 6 trials of each type before reversal, and a group of 6 trials of each type from after reversal, “slid” in 1 trial steps. For example, the first ANOVA would be computed using trials 1-6 of each type after reversal; the next, using trials 2-7 of each type, etc.

The total variance obtained from each iteration of the ANOVA (SS_{total}) was partitioned into image value (SS_{val}), image identity (SS_{id}), interaction (SS_{int}) and error (SS_{err}) terms. The strength of the contribution of image value (R^2_{val}) was quantified as (SS_{val}/SS_{total}). For each trial after reversal, we then averaged the contribution of image value (R^2_{val}) obtained for each cell over all cells in a subgroup. We normalized the population average by dividing by the maximum average R^2_{val}. We then fitted the neural data from each subgroup with a

Weibull function (Eq. 1). Results were similar and statistically significant for both monkeys, so the data were combined.

Quantitative analysis of neural dynamics—In several instances, we fit neural data with a sigmoid curve using a Weibull function:

$$f(x)=u+(l-u)\exp(-x/\alpha)^{\beta} \quad (1)$$

which modeled the data as a function of trial number after reversal (Fig. 5) or time during the trial (Fig. 8e,f). The *u* and *l* parameters adjust the upper and lower asymptotes of the fit curve, respectively, and the β parameter adjusts the shape of the curve. The α parameter can be considered to be a scale-adjusted rise latency: it is equivalent to the value of *x* (trial number) for which *f(x)* reaches a certain percentage of its maximum value. This value depends upon the upper and lower limits of the fit. Specifically, when *x* is equal to α, the function reaches a level defined by:

$$f(x)\approx 0.63*u+0.37*l \quad (2)$$

We could determine whether the α parameter was significantly different for two data sets by fitting the data twice: once with all parameters free, and once with the α parameter constrained to be the same for the two data sets. An F-test was used to determine whether separate α parameters explained the data better than a single α parameter; i.e., whether one fit reaches its scale-adjusted threshold significantly earlier than the other. We also examined whether the 95% confidence bounds for the α parameters overlap. In some cases, we report the difference between the α parameters to quantify the separation between two curves.

Multi-dimensional sliding regression analysis—To evaluate the detailed time course of changes in neural activity and behavior after reversal, we performed a sliding ANOVA analysis. For every value-coding cell we divided each trial into 200 ms bins that were slid across the trial in 20 ms steps, and obtained the spike count for each bin. Then, for the data from each bin, we calculated the two-way ANOVA using the last 6 trials of each type before reversal, and a group of 6 trials of each type after reversal, slid in 1 trial steps.

Again, the total variance obtained from each iteration of the ANOVA (SS_{total}) was partitioned into image value (SS_{val}), image identity (SS_{id}), interaction (SS_{int}) and error (SS_{err}) terms. The strength of the contribution of image value (R^2_{val}) was quantified as (SS_{val}/SS). The proportion of total explainable variance (SS_{exp}) was calculated as ($1 - SS_{err}/SS_{total}$). To compare the time course of neural changes in the different groups of neurons, we calculated the average R^2_{val} for each bin across all cells in the subgroup, and divided by the average SS_{exp} for that bin, to obtain the value index for that bin. (For the previous sliding regression analysis, which used data from a larger time window, we did not divide by the total explainable variance because the error term was generally small.)

We also obtained an estimate of the contribution of image identity, the “identity index,” for each bin, which can be quantified as (SS_{id}/SS_{total})/(SS_{exp}) (Fig. 7a-d). We similarly calculated an “interaction index” for each bin, which can be quantified as (SS_{int}/SS_{total})/(SS_{exp}), representing the contribution of the image value/identity interaction factor to neural activity (Fig. 7e-h).

To calculate a "rise-time" for the neural data – i.e., when the value index becomes significant on each trial – we used the Fisher method (Fisher, 1925, 1948) to combine the image value factor p-values obtained for each bin across all cells in a group. The rise-time was defined as the beginning of the first three consecutive bins for which the Fisher p-value was < 0.01. To determine whether rise-times were significantly different across groups, we used a permutation test with 1000 shuffles. For each shuffle, we randomly assigned each cell to one of the groups being compared (e.g., positive value-coding cells in OFC vs. positive value-coding cells in amygdala), and calculated rise-times for each group. We then calculated a rise-time difference for each trial, and finally compared the actual rise-time difference from each trial with the population of differences derived from the shuffle.

Post-learning sliding regression—To visualize the latency and timing of value-related activity, we applied a sliding ANOVA to neural data from post-learning trials (the last 20 trials of each type from the initial and reversal blocks). For each value-coding cell, we divided the trial into 200 ms bins, slid by 20 ms, and did a two-way ANOVA with factors of image value and image identity on the spike count from each bin. The SS_{val} obtained for each bin is the contribution-of-value signal. To construct Fig. 8a-d, we determined for each cell the first bin in which the value signal reached statistical significance ($p < 0.01$). We then averaged the contribution-of-value signal in each bin across all cells in each group, and normalized the results by the maximum average signal (Fig. 8e,f). To compare the time course of the average signal, we fit Weibull curves (Eq. 1) to the average data from the first 500 ms after CS onset. We used an F-test to determine that the α parameters were different for the curves fit to OFC and amygdala data.

Granger causality analysis of LFP signals—We assessed directional influences between OFC and amygdala using Granger causality analysis (Granger, 1969). One signal X(t) Granger-causes another signal Y(t) if the linear prediction of future values of Y is improved by taking into account the past values of X. We fit a bivariate autoregressive (AR) model to the two-dimensional vector (X(t),Y(t)), where X and Y are segments of LFP recorded in OFC and amygdala, respectively. We then estimated Granger causalities for each direction of influence (OFC-to-amygdala and amygdala-to-OFC) in the frequency domain (Geweke, 1982) from the AR parameters (Brovelli et al., 2004).

We examined the evolution of Granger causality by analyzing brief segments of LFP signal starting 0.5 s before CS onset until US onset (200 ms window, stepped by 50 ms, yielding 43 steps). The short window ensured that the LFPs within it could be considered stationary. For each step, the 200-ms LFP segments from trials of the same type (positive or negative) were concatenated separately and the parameters of the AR model for the resulting time series were estimated using the Nutall-Strand method (Schlögl, 2006). We fixed the AR model order to 50, and assessed model fit by testing for lack of residual correlations (Li and Mcleod, 1981). We determined the statistical significance for Granger causality at each time-frequency bin using the frequency-domain test described by Breitung and Candelon (2006).

To average Granger causality values, we first normalized these values for each pair to the value estimated for the first time window (-0.5 to -0.3 s relative to CS onset). This was performed separately for each frequency bin between 0-100 Hz. For each pair and trial type, we only averaged data from pairs that yielded Granger causality values with 4 consecutive significant time bins ($p < 0.01$, spanning 350 ms). To compare the Granger causality in the two different directions of influence (Fig. 9a), all trials of the reversal block were combined as described above. At each time bin, the Granger causality values for all frequencies from 5-100 Hz were averaged together. We determined the statistical significance of the

difference between the two directions (OFC-to-amygdala and amygdala-to-OFC) using a permutation test (10,000 shuffles).

To assess the effect of learning on the influence between the amygdala and OFC (Fig. 9b,c), only the 6 first trials of each type (12 total) after reversal and the last 6 trials of each type in the experiment were used. Granger causality was computed for these two sets of trials, and we compared its relative magnitude in both directions during and after reversal learning. For each set of trials, the Granger causality values were averaged across pairs and trial types as described above. The difference between the mean Granger causality in the two directions was then compared for the during-learning and post-learning sets in the time domain (Fig. 9b) by averaging across frequencies from 5-100 Hz; Figure 9c does this in the frequency domain by averaging across times from CS onset until the end of the trace interval. The significance of the difference between during-learning and post-learning was assessed by permutation test (10,000 shuffles).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank members of the Salzman lab and the Mahoney Center at Columbia for helpful comments and discussions; S. Dashnaw and J. Hirsch for MRI support; M. Belova and J. Paton for advice; and K. Marmon for invaluable technical assistance. This research was supported by grants from NIMH, NIDA, NEI, and the James S. McDonnell and Gatsby foundations. S.E.M. received support from an NSF graduate fellowship and from an individual NIMH NRSA. B.L. received support from NIMH and the Helen Hay Whitney Foundation. A.S. was supported by the Kavli Foundation.

References

Baxter MG, Murray EA. The amygdala and reward. Nat Rev Neurosci. 2002; 3:563–573. [PubMed: 12094212]

Belova MA, Paton JJ, Morrison SE, Salzman CD. Expectation modulates neural responses to pleasant and aversive stimuli in primate amygdala. Neuron. 2007; 55:970–984. [PubMed: 17880899]

Belova MA, Paton JJ, Salzman CD. Moment-to-moment tracking of state value in the amygdala. J Neurosci. 2008; 28:10023–10030. [PubMed: 18829960]

Bermudez MA, Schultz W. Responses of amygdala neurons to positive reward-predicting stimuli depend on background reward (contingency) rather than stimulus-reward pairing (contiguity). J Neurophysiol. 2010; 103:1158–1170. [PubMed: 20032233]

Brasted PJ, Wise SP. Comparison of learning-related neuronal activity in the dorsal premotor cortex and striatum. Eur J Neurosci. 2004; 19:721–740. [PubMed: 14984423]

Breitung J, Candelon B. Testing for short- and long-run causality: a frequency domain approach. J Econometrics. 2006; 132:363–378.

Brovelli A, Ding M, Ledberg A, Chen Y, Nakamura R, Bressler SL. Beta oscillations in a large-scale sensorimotor cortical network: directional influences revealed by Granger causality. Proc Natl Acad Sci U S A. 2004; 101:9849–9854. [PubMed: 15210971]

Carmichael ST, Price JL. Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. J Comp Neurol. 1995; 363:615–641. [PubMed: 8847421]

Chen L, Wise S. Neuronal activity in the supplementary eye field during acquisition of conditional oculomotor associations. J Neurophysiol. 1995a; 73:1101–1121. [PubMed: 7608758]

Chen LL, Wise SP. Supplementary eye field contrasted with the frontal eye field during acquisition of conditional oculomotor associations. J Neurophysiol. 1995b; 73:1122–1134. [PubMed: 7608759]

Chen LL, Wise SP. Evolution of directional preferences in the supplementary eye field during acquisition of conditional oculomotor associations. J Neurosci. 1996; 16:3067–3081. [PubMed: 8622136]

Chudasama Y, Robbins TW. Dissociable contributions of the orbitofrontal and infralimbic cortex to pavlovian autoshaping and discrimination reversal learning: further evidence for the functional heterogeneity of the rodent frontal cortex. J Neurosci. 2003; 23:8771–8780. [PubMed: 14507977]

Fellows LK, Farah MJ. Ventromedial frontal cortex mediates affective shifting in humans: evidence from a reversal learning paradigm. Brain. 2003; 126:1830–1837. [PubMed: 12821528]

Fisher, RA. Statistical methods for research workers. Edinburgh: Oliver and Boyd; 1925.

Fisher RA. Combining independent tests of significance. American Statistician. 1948; 2:30.

Fudge JL, Kunishio K, Walsh P, Richard C, Haber SN. Amygdaloid projections to ventromedial striatal subterritories in the primate. Neuroscience. 2002; 110:257–275. [PubMed: 11958868]

Gallistel CR, Fairhurst S, Balsam P. The learning curve: implications of a quantitative analysis. Proc Natl Acad Sci U S A. 2004; 101:13124–13131. [PubMed: 15331782]

Geweke J. Measurement of linear dependence and feedback between multiple time series. J Am Stat Assoc. 1982; 77:304–324.

Ghashghaei HT, Hilgetag CC, Barbas H. Sequence of information processing for emotions based on the anatomic dialogue between prefrontal cortex and amygdala. Neuroimage. 2007; 34:905–923. [PubMed: 17126037]

Granger CWJ. Investigating causal relationships by econometric models and cross-spectral methods. Econometrica. 1969; 37:424–438.

Haber SN, Kunishio K, Mizobuchi M, Lynd-Balta E. The orbital and medial prefrontal circuit through the primate basal ganglia. J Neurosci. 1995; 15:4851–4867. [PubMed: 7623116]

Hosokawa T, Kato K, Inoue M, Mikami A. Neurons in the macaque orbitofrontal cortex code relative preference of both rewarding and aversive outcomes. Neurosci Res. 2007; 57:434–445. [PubMed: 17239463]

Iversen SD, Mishkin M. Perseverative interference in monkeys following selective lesions of the inferior prefrontal convexity. Exp Brain Res. 1970; 11:376–386. [PubMed: 4993199]

Izquierdo A, Suda RK, Murray EA. Bilateral orbital prefrontal cortex lesions in rhesus monkeys disrupt choices guided by both reward value and reward contingency. J Neurosci. 2004; 24:7540–7548. [PubMed: 15329401]

Johansen JP, Hamanaka H, Monfils MH, Behnia R, Deisseroth K, Blair HT, LeDoux JE. Optical activation of lateral amygdala pyramidal cells instructs associative fear learning. Proc Natl Acad Sci U S A. 2010; 107:12692–12697. [PubMed: 20615999]

Kazama A, Bachevalier J. Selective aspiration or neurotoxic lesions of orbital frontal areas 11 and 13 spared monkeys' performance on the object discrimination reversal task. J Neurosci. 2009; 29:2794–2804. [PubMed: 19261875]

Kopell N, Ermentrout GB, Whittington MA, Traub RD. Gamma rhythms and beta rhythms have different synchronization properties. Proc Natl Acad Sci U S A. 2000; 97:1867–1872. [PubMed: 10677548]

Li WK, Mcleod AI. Distribution of the residual autocorrelation in multivariate ARMA time series models. J Roy Statist Soc B. 1981; 43:231–239.

Miller EK, Cohen JD. An integrative theory of prefrontal cortex function. Annu Rev Neurosci. 2001; 24:167–202. [PubMed: 11283309]

Mitz AR, Godschalk M, Wise SP. Learning-dependent neuronal activity in the premotor cortex: activity during the acquisition of conditional motor associations. J Neurosci. 1991; 11:1855–1872. [PubMed: 2045890]

Morrison SE, Salzman CD. The convergence of information about rewarding and aversive stimuli in single neurons. J Neurosci. 2009; 29:11471–11483. [PubMed: 19759296]

Morrison SE, Salzman CD. Re-valuing the amygdala. Curr Opin Neurobiol. 2010; 20:221–230. [PubMed: 20299204]

Murray EA, Izquierdo A. Orbitofrontal cortex and amygdala contributions to affect and action in primates. Ann N Y Acad Sci. 2007; 1121:273–296. [PubMed: 17846154]

Murray EA, Wise SP. Interactions between orbital prefrontal cortex and amygdala: advanced cognition, learned responses and instinctive behaviors. Curr Opin Neurobiol. 2010; 20:212–220. [PubMed: 20181474]

Nishijo H, Ono T, Nishino H. Single neuron responses in amygdala of alert monkey during complex sensory stimulation with affective significance. J Neurosci. 1988; 8:3570–3583. [PubMed: 3193171]

Ongur D, Price JL. The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. Cereb Cortex. 2000; 10:206–219. [PubMed: 10731217]

Padoa-Schioppa C, Assad JA. Neurons in the orbitofrontal cortex encode economic value. Nature. 2006; 441:223–226. [PubMed: 16633341]

Pasupathy A, Miller EK. Different time courses of learning-related activity in the prefrontal cortex and striatum. Nature. 2005; 433:873–876. [PubMed: 15729344]

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]

Paxinos, G.; Huang, X.; Toga, AW. The rhesus monkey brain in stereotaxic coordinates. San Diego: Academic Press; 2000.

Phelps EA, LeDoux JE. Contributions of the amygdala to emotion processing: from animal models to human behavior. Neuron. 2005; 48:175–187. [PubMed: 16242399]

Roesch MR, Olson CR. Neuronal activity related to reward value and motivation in primate frontal cortex. Science. 2004; 304:307–310. [PubMed: 15073380]

Rolls, E. Neurophysiology and functions of the primate amygdala. In: Aggleton, J., editor. The amygdala: Neurobiological aspects of emotion, memory, and mental dysfunction. New York: Wiley-Liss, Inc.; 1992. p. 143-166.

Rolls ET, Critchley HD, Mason R, Wakeman EA. Orbitofrontal cortex neurons: role in olfactory and visual association learning. J Neurophysiol. 1996; 75:1970–1981. [PubMed: 8734596]

Rolls ET, Grabenhorst F. The orbitofrontal cortex and beyond: from affect to decision-making. Prog Neurobiol. 2008; 86:216–244. [PubMed: 18824074]

Saddoris MP, Gallagher M, Schoenbaum G. Rapid associative encoding in basolateral amygdala depends on connections with orbitofrontal cortex. Neuron. 2005; 46:321–331. [PubMed: 15848809]

Salzman CD, Fusi S. Emotion, cognition, and mental state representation in amygdala and prefrontal cortex. Annu Rev Neurosci. 2010; 33:173–202. [PubMed: 20331363]

Salzman CD, Paton JJ, Belova MA, Morrison SE. Flexible neural representations of value in the primate brain. Ann N Y Acad Sci. 2007

Sanghera MK, Rolls ET, Roper-Hall A. Visual responses of neurons in the dorsolateral amygdala of the alert monkey. Exp Neurol. 1979; 63:610–626. [PubMed: 428486]

Schlögl A. Comparison of multivariate autoregressive estimators. Signal Processing. 2006; 86:2426–2429.

Schoenbaum G, Chiba AA, Gallagher M. Neural encoding in orbitofrontal cortex and basolateral amygdala during olfactory discrimination learning. J Neurosci. 1999; 19:1876–1884. [PubMed: 10024371]

Schoenbaum G, Roesch MR, Stalnaker TA, Takahashi YK. A new perspective on the role of the orbitofrontal cortex in adaptive behaviour. Nat Rev Neurosci. 2009; 10:885–892. [PubMed: 19904278]

Schoenbaum G, Setlow B, Nugent SL, Saddoris MP, Gallagher M. Lesions of orbitofrontal cortex and basolateral amygdala complex disrupt acquisition of odor-guided discriminations and reversals. Learn Mem. 2003; 10:129–140. [PubMed: 12663751]

Stalnaker TA, Franz TM, Singh T, Schoenbaum G. Basolateral amygdala lesions abolish orbitofrontal-dependent reversal impairments. Neuron. 2007; 54:51–58. [PubMed: 17408577]

Stefanacci L, Amaral DG. Topographic organization of cortical inputs to the lateral nucleus of the macaque monkey amygdala: a retrograde tracing study. J Comp Neurol. 2000; 421:52–79. [PubMed: 10813772]

Stefanacci L, Amaral DG. Some observations on cortical inputs to the macaque monkey amygdala: an anterograde tracing study. J Comp Neurol. 2002; 451:301–323. [PubMed: 12210126]

Sutton, R.; Barto, A. Reinforcement Learning. Cambridge, Massachusetts: MIT Press; 1998.

Thorpe SJ, Rolls ET, Maddison S. The orbitofrontal cortex: neuronal activity in the behaving monkey. Exp Brain Res. 1983; 49:93–115. [PubMed: 6861938]

Tremblay L, Hollerman JR, Schultz W. Modifications of reward expectation-related neuronal activity during learning in primate striatum. J Neurophysiol. 1998; 80:964–977. [PubMed: 9705482]

Tremblay L, Schultz W. Relative reward preference in primate orbitofrontal cortex. Nature. 1999; 398:704–708. [PubMed: 10227292]

Tremblay L, Schultz W. Modifications of reward expectation-related neuronal activity during learning in primate orbitofrontal cortex. J Neurophysiol. 2000; 83:1877–1885. [PubMed: 10758099]

Wirth S, Yanike M, Frank LM, Smith AC, Brown EN, Suzuki WA. Single neurons in the monkey hippocampus and learning of new associations. Science. 2003; 300:1578–1581. [PubMed: 12791995]

Wise SP. Forward frontal fields: phylogeny and fundamental function. Trends Neurosci. 2008; 31:599–608. [PubMed: 18835649]

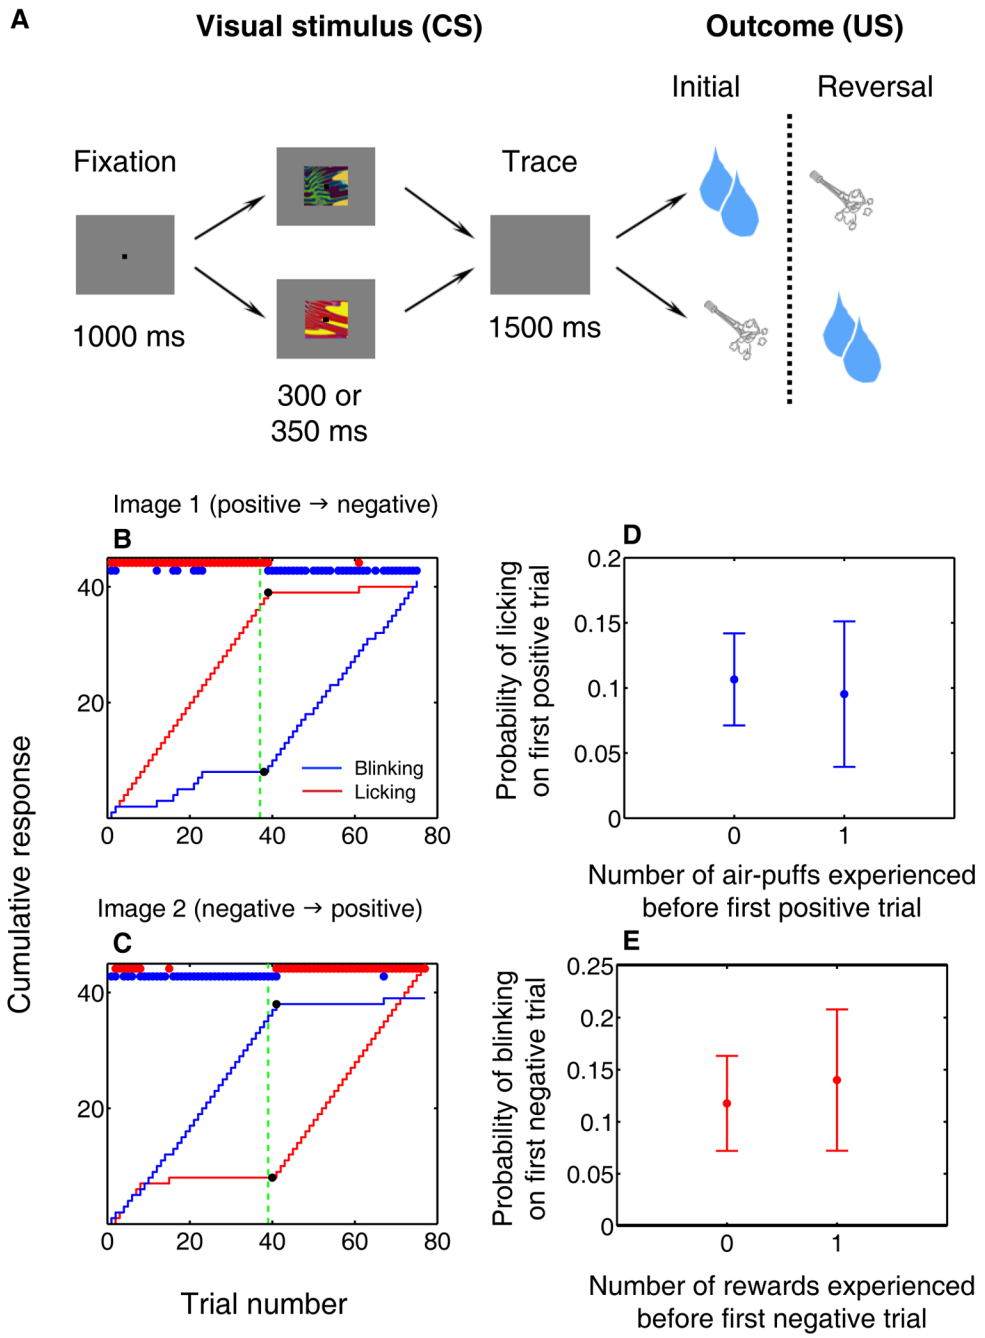

Figure 1. Task and behavioral responses during reversal learning

(A) Trace conditioning task. Images associated with large reward and air-puff reverse contingencies without warning after initial learning. Not shown: a third trial type, in which small reward follows a CS with no reversal. (**B,C**) Cumulative (curves) and trial-by-trial (red and blue dots) indicators of licking (red) and blinking (blue). Black dots, change points. Green dotted line, image value reversal. (**D,E**) Mean probability (±SEM) of monkeys' exhibiting the "correct" behavior – licking on positive trials (**D**), blinking on negative trials (**E**) – on the first trial of each type after reversal, after having experienced zero or one trials of the opposite type.

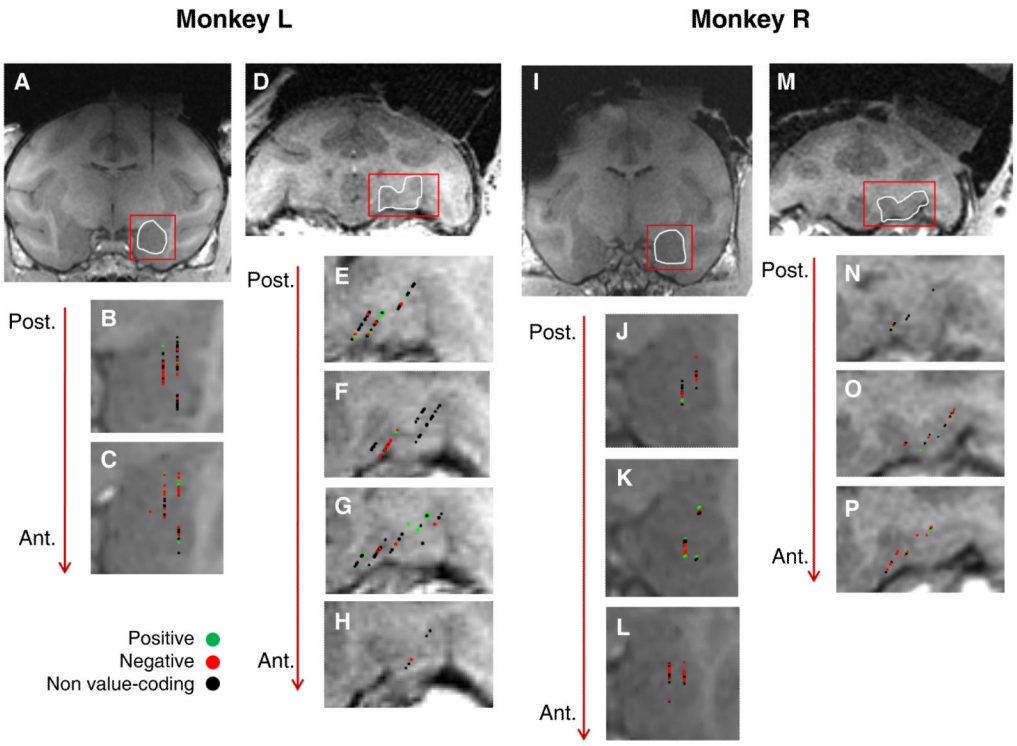

Figure 2. Recording sites in amygdala and orbitofrontal cortex

(A-P) Coronal T1-weighted MRI in Monkey L (**A-H**) and Monkey R (**I-P**). **(A)** MR image showing the artifact from an electrode inserted dorsal to the amygdala. The images in **A-C** and **I-L** have been reoriented to trace the path of the electrode. **(B,C)** Magnified images of the amygdala in two consecutive slices (1 mm apart) from the MRI shown in **A**. Symbols indicate properties of cells. **(D)** Coronal MRI showing a typical area of OFC from which we recorded. **E-H,** Magnified images of OFC in four consecutive slices (1 mm apart) from the MRI shown in **D. (I)** MR image, reoriented as described above, showing the amygdala. **(J-L)** Magnified images of the amygdala in three consecutive slices (1 mm apart) from the MRI shown in **I. (M)** Coronal MRI showing a typical area of OFC from which we recorded. **(N-P)** Magnified images of OFC in three consecutive slices (1 mm apart) from the MRI shown in **M**. Conventions as in **B,C** applied for all panels.

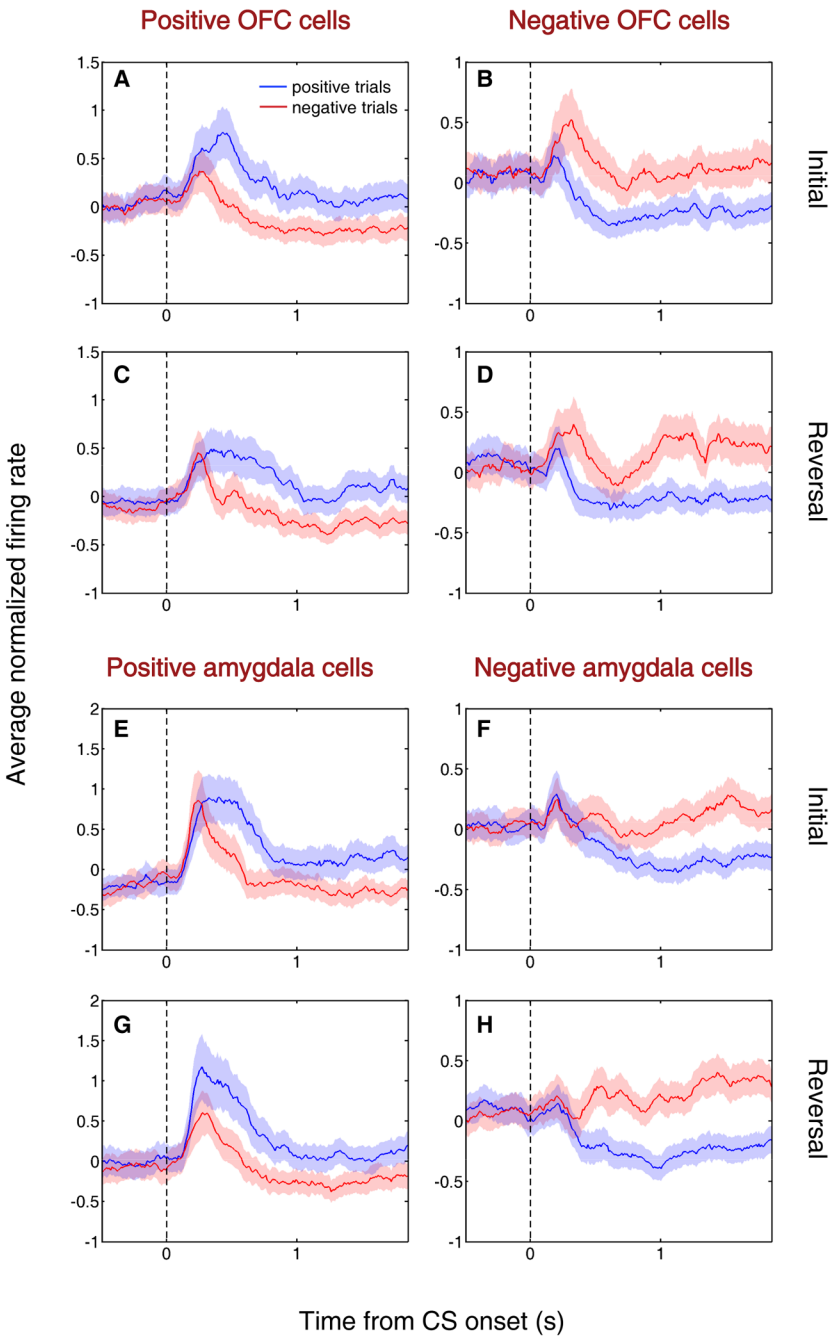

Figure 3. Value-coding cells encode reinforcement contingencies before and after reversal

(A-H) Average normalized neural activity (±SEM) of positive value-coding cells (**A,C,E,G**) and negative value-coding cells (**B,D,F,H**) in OFC (**A-D**) and amygdala (**E-H**), shown separately for initial (**A,B,E,F**) and reversal (**C,D,G,H**) blocks. The first five trials of each type in each block are omitted because activity was often changing during these trials. Peristimulus time histograms (PSTHs) are constructed by building individual PSTHs using 10 ms non-overlapping bins, Z-scoring (subtracting the mean of all bins and dividing by the standard deviation), and then averaging all cells in each bin. The resulting population PSTH is smoothed using a 10-bin moving average. PSTHs are truncated at US onset. Blue line,

average activity during positive trials; red, average activity during negative trials. Vertical dotted line, CS onset.

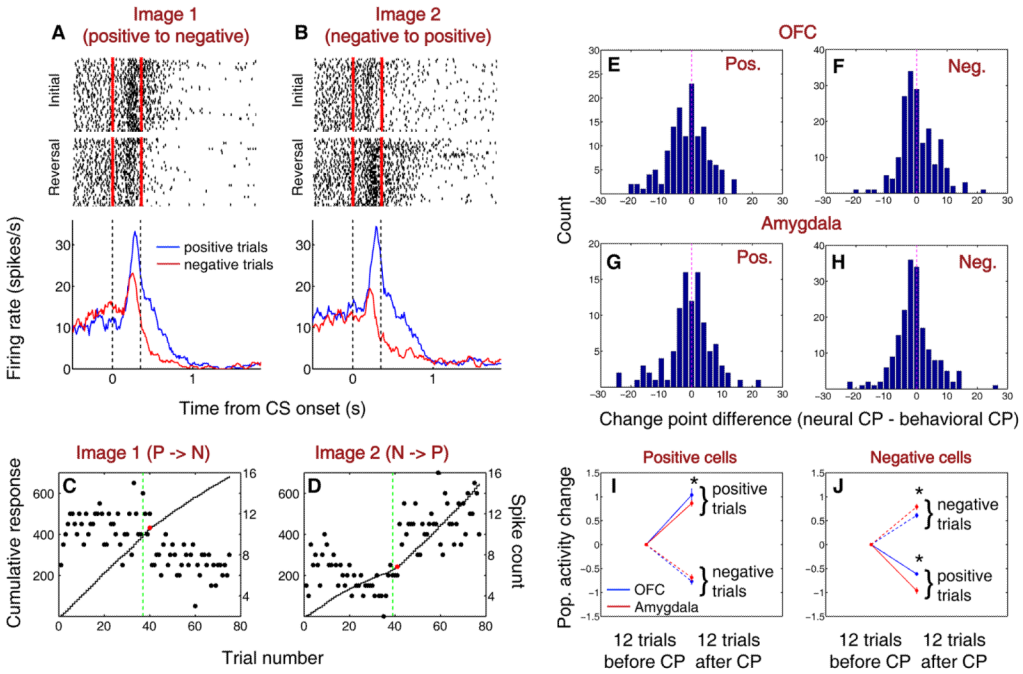

Figure 4. Positive and negative value-coding cells in amygdala and OFC show learning-related changes in activity with similar onsets but different time courses

(A,B) Rasters and peristimulus time histograms (PSTHs) illustrating the activity of a positive value-coding cell recorded in OFC during the behavior shown in Fig. 1b,c. Rasters and PSTHs are truncated at US onset. Red ticks and vertical dashed lines indicate CS onset and offset. **(C,D)** Change point analysis of the neural activity shown in **A,B**. Spike count (dots) and cumulative spike count (curves) during the CS interval, plotted as a function of trial number for image 1 (**C**) and image 2 (**D**). Red dots, change points. Green dotted line, image value reversal. **(E-H)** Distribution of differences between neural and behavioral change points (CPs; calculated as neural CP – behavioral CP) from positive value-coding neurons in OFC (**E**), negative value-coding neurons in OFC (**F**), positive value-coding neurons in amygdala (**G**), and negative value coding neurons in amygdala (**H**). **(I,J)** Normalized average neural activity in the 12 trials before and after the CP. Asterisks, $p < 0.05$, Wilcoxon).

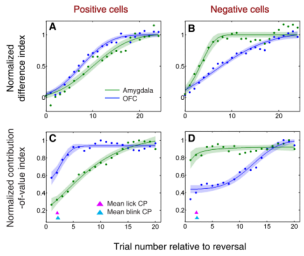

Figure 5. Positive and negative value-coding neurons in amygdala and OFC exhibit different time courses of learning-related activity

(A,B) Normalized average difference index plotted as a function of trial number relative to reversal for positive value-coding cells (**A**) and negative value-coding cells (**B**) in OFC (blue) and amygdala (green). Data is smoothed using a 6-trial average, stepped by 1 trial. **(C,D)** Normalized average contribution of image value to neural activity, derived from ANOVA, plotted as a function of trial number after reversal for positive value-coding neurons (**C**) and negative value-coding neurons (**D**). Red and cyan arrowheads indicate mean licking and blinking change points, respectively; the width of each arrowhead's base indicates SEM. In **A-D**, neural activity is derived from 90-590 ms after CS onset. Curves are best-fit sigmoids (± 95% prediction intervals).

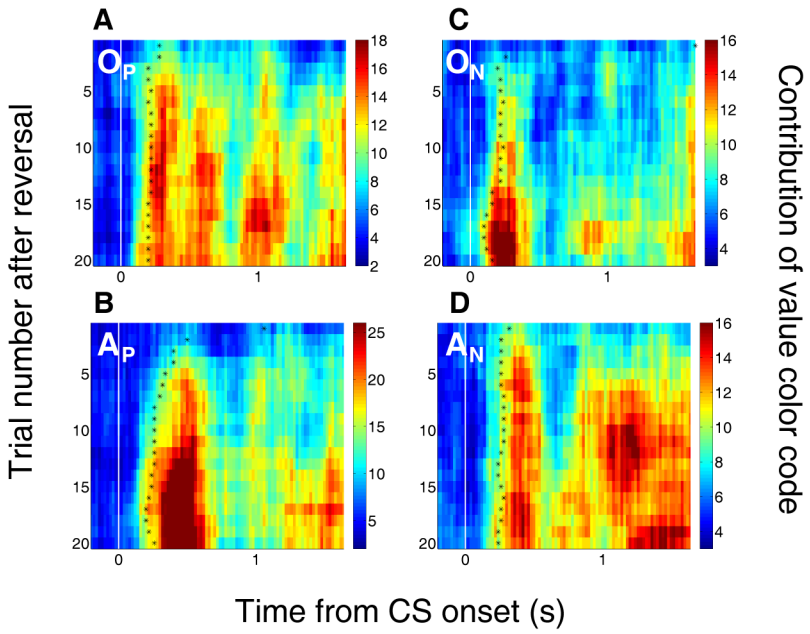

Figure 6. Multidimensional analysis of the contribution of image value to neural responses during reversal learning

(A-D) Average contribution of image value to neural activity from positive value-coding neurons in OFC (**A**) and amygdala (**B**), and from negative value-coding cells in OFC (**C**) and amygdala (**D**), plotted as a function of trials after reversal and time from image onset (white line). Plots end at US onset. Black asterisks, time when the contribution-of-value index becomes significant (asterisks placed in the center of the first of at least 3 consecutive significant bins, Fisher $p < 0.01$). Bin size, 200 ms. Bin steps, 20 ms. A grayscale version is provided (Fig. S1).

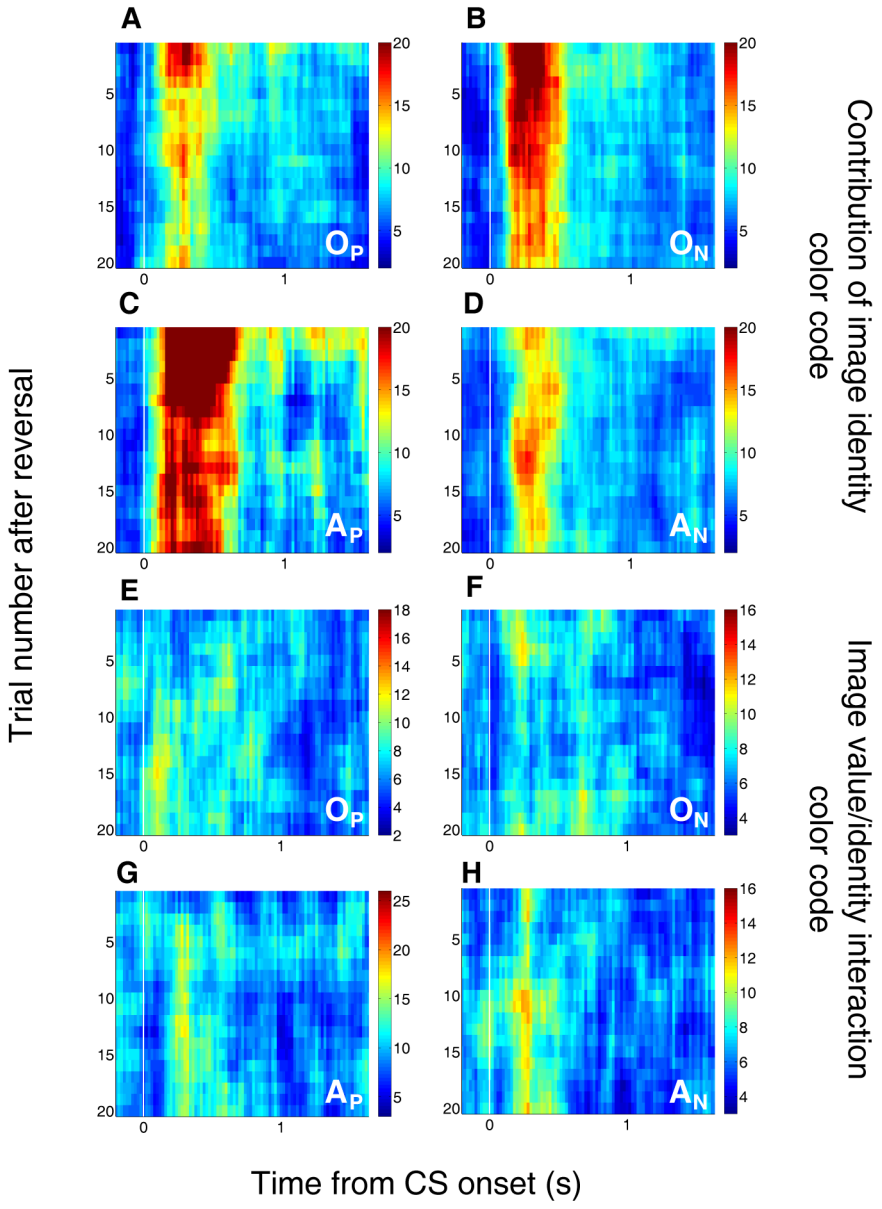

Figure 7. Encoding of image identity and interaction between image identity and image value

(A-D) Average contribution of image identity to neural activity from positive value-coding neurons in OFC (**A**) and amygdala (**B**) and from negative value-coding neurons in OFC (**C**) and amygdala (**D**). Just after reversal, when some neurons' activity still reflects the previous image-value contingencies, the variance that is actually due to coding the old image values becomes assigned to image identity. Therefore, the contribution of image identity is higher immediately after reversal, and then tends to decrease as neural responses to CSs change during learning. This decrease is slower in the slower-changing neuronal subgroups (positive amygdala and negative OFC cells). **(E-H)** Average contribution of image value/identity interaction to neural activity from positive value-coding neurons in OFC (**E**) and amygdala (**F**) and from negative value-coding neurons in OFC (**G**) and amygdala (**H**). Note the relatively low magnitude of the interaction factor in comparison to the contribution of the main value factor (Fig. 6) at equivalent times and number of trials after reversal. Other conventions as in Figure 6. A grayscale version is provided (Fig. S2).

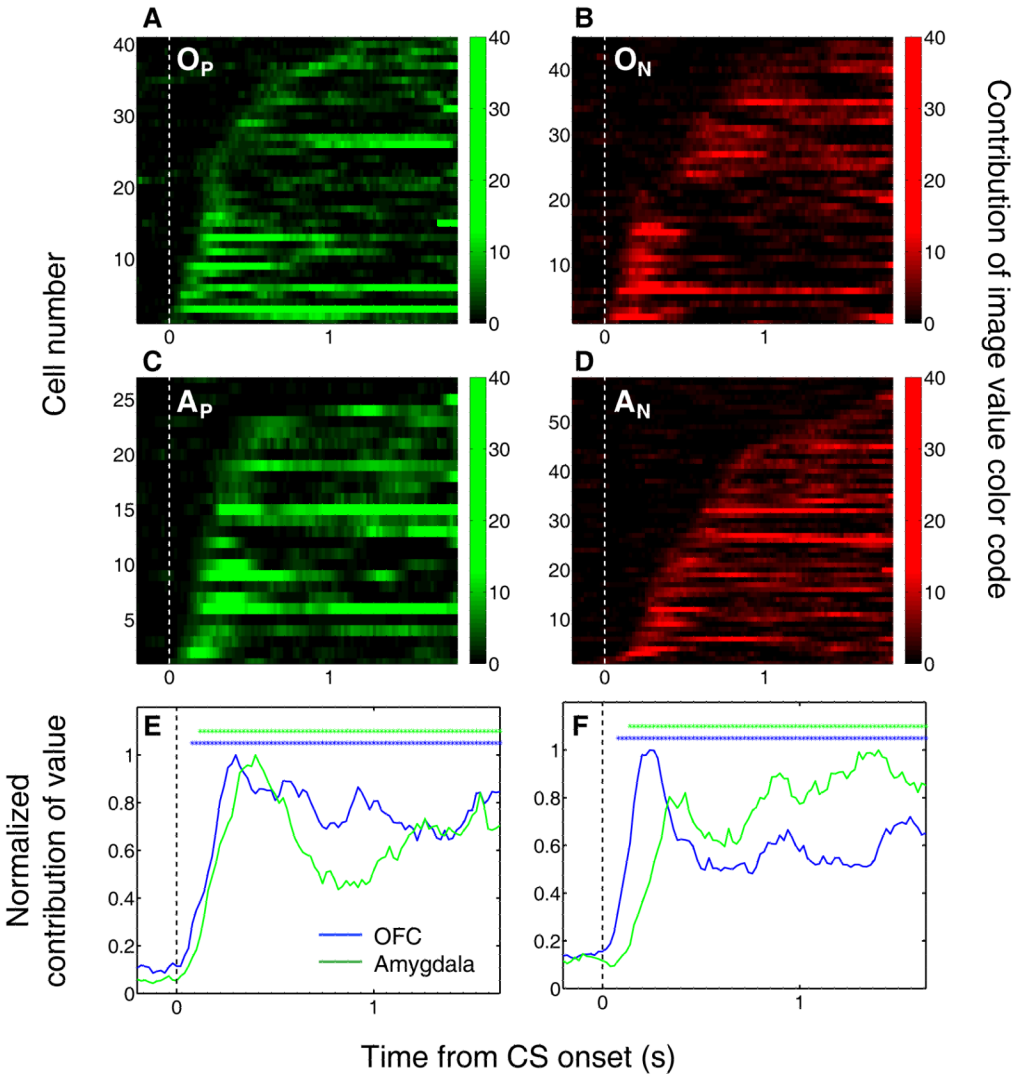

Figure 8. After learning, OFC activity encodes image value more rapidly than amygdala activity

(A-D) Contribution of image value as a function of time for all positive value-coding cells in OFC (**A**) and amygdala (**C**), and all negative value-coding cells in OFC (**B**) and amygdala (**D**). Each row represents one neuron, aligned to CS onset. The analysis was applied only to data from the last 20 trials of each type from initial and reversal learning. Neurons are sorted in order of the time after CS onset when the contribution of value reaches significance ($p < 0.05$). Bin size, 200 ms; bin steps, 20 ms. Plots are truncated at US onset. **(E,F)** Normalized average contribution of image value as a function of time for positive value-coding cells (**E**) and negative value-coding cells (**F**). Asterisks, time points at which the average contribution of value is significant (Fisher $p < 0.0001$) for OFC (blue) or amygdala (green).

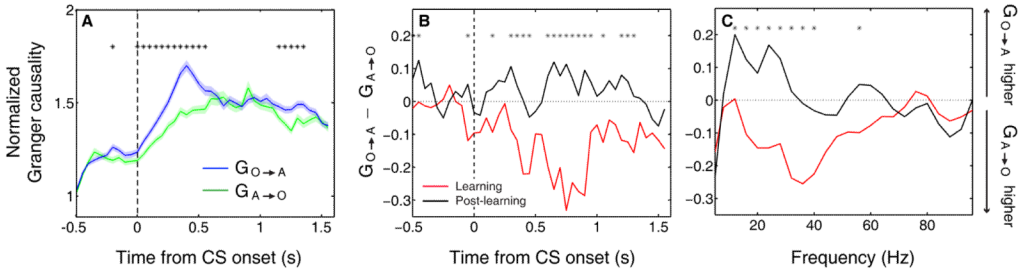

Figure 9. Granger causality in the OFC-to-amygdala direction increases with learning

(A) Average normalized Granger causality (±SEM) for the OFC-to-amygdala direction (blue) and the amygdala-to-OFC direction (green). For each pair of OFC-amygdala LFP recordings, Granger causality was computed for all trials after reversal, then averaged across pairs. Only pairs with significant Granger causality at some point during the trial were included in the average, which combines frequencies from 5-100 Hz. Granger causality in both directions is significantly enhanced after CS onset (Wilcoxon, $p < 0.01$), and causality in the OFC-to-amygdala direction is stronger than causality in the amygdala-to-OFC direction over the trial as a whole. Asterisks, bins with significantly different causality for the two directions (permutation test, $p < 0.05$). (**B-C**) Granger causality changes with learning. The difference between the mean Granger causality in the two directions (subtracting amygdala-to-OFC from OFC-to-amygdala) was separately calculated for early ('During-learning', red) and late ('Post-learning, black') trials after reversal. This comparison is shown for all frequencies 5-100 Hz as a function of time within the trial (**B**) and for the CS and trace intervals combined as a function of frequency (**C**). Asterisks, bins where the difference between during- and post-learning values was significant (permutation test, $p < 0.05$).