Supplementary Information:

The primate amygdala combines information about space and value

Christopher J. Peck*, Brian Lau*, and C. Daniel Salzman†
*These authors contributed equally to this work
†To whom correspondence should be addressed. E-mail: cds2005@columbia.edu

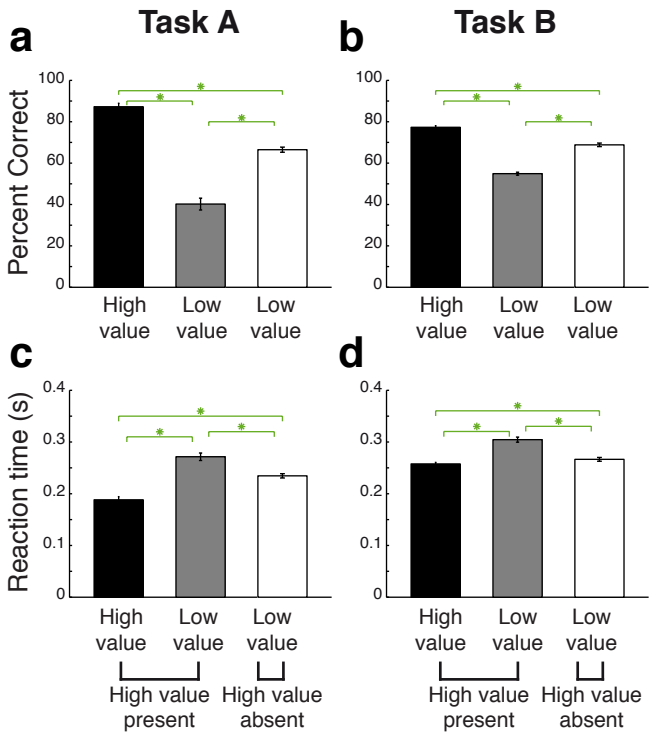

Supplementary Figure 1. Both tasks bias spatial attention as measured by performance and reaction time

Behavioral results for each task individually. Green asterisks indicate the significance of all comparisons, in both tasks (Paired Wilcoxon, $P < 10^{-6}$). (**a,b**) Percent correct on task A and task B. On high value present trials (black bars), percent correct was significantly higher when the target appeared near the high value cue location that when it appeared at the low value cue location in both tasks. Note that chance performance on Task A was approximately 23% correct, while chance performance in Task B was 50%; in each task, performance at the low value cue location was still significantly better than chance ($P < 10^{-6}$). When two low value cues appeared (high value absent trials), monkeys performed at a level that was significantly greater than when that low value cue was paired with a high value cue, and the target appeared near the low value cue location (compare grey & white bars). Thus, monkeys did not inherently perform poorly for the low value outcomes; rather, performance for low value outcomes suffered when a high value outcome biased attention away from the low value location. (**c,d**) Reaction times in the same format as (**a,b**). Reaction times were shorter when the target appeared near high value cue location than when it appeared near the low value cue location (on high value present trials). Low value cue reaction times were significantly shorter on high value absent trials than high value present trials. Thus, instances of higher hit rates were mirrored by instances of shorter reaction times for all comparisons, in both tasks.

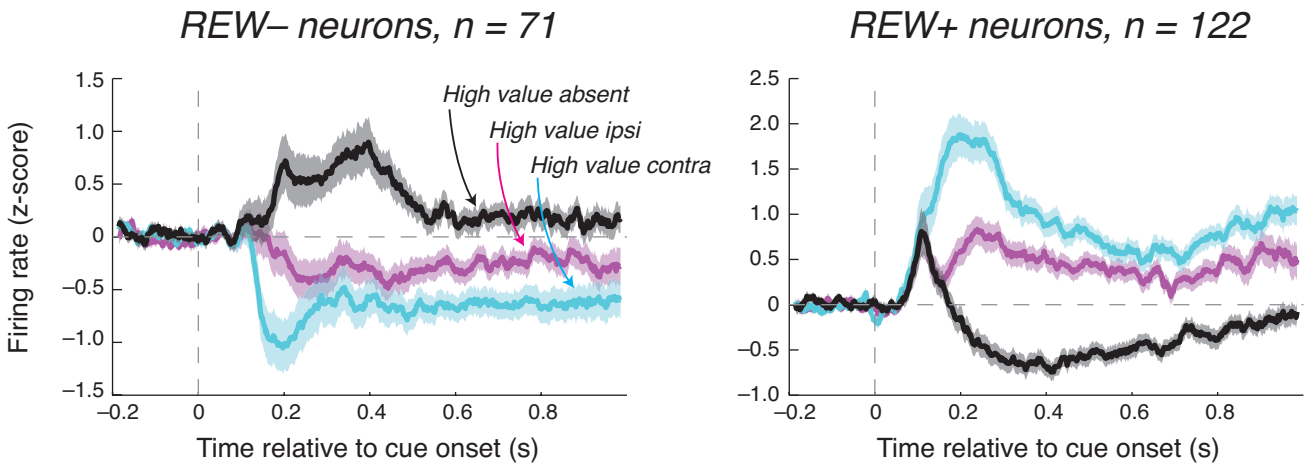

Supplementary Figure 2. Population PSTHs for REW– and REW+ neurons

Firing rates averaged across the population are plotted as a function of time relative to cue onset for each onset. Z-scores were calculated by subtracting the mean baseline firing rate over all trial types (200 ms before cue onset) and dividing by the within group standard deviation of firing rates (100–1000 ms after cue onset) according to:

$$\frac{z_i = fr_i - \mu}{s}, where\ fr_i\ is\ the\ firing\ rate, \mu\ is\ the\ mean\ baseline\ firing\ rate$$

$$s = \sqrt{\frac{\sum_g^3 \sum_{i \epsilon g} (x_i - \bar{x}_{i \epsilon g})^2}{N - 1}},\ \ where\ g\ is\ the\ trial\ type, N\ is\ the\ number\ of\ trials$$

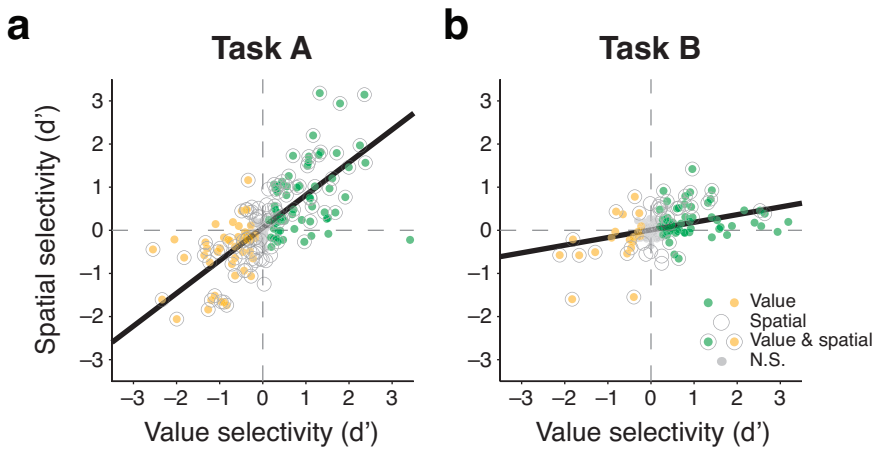

Supplementary Figure 3. Value and spatial selectivity are systematically related in both tasks

Scatterplot of value selectivity vs. spatial selectivity indices for each individual neuron in Task A (**a**) and Task B (**b**). The line represents the weighted least squares regression fit in each case (slope is significantly greater than 0 in both tasks ($P < 10^{-4}$). Moreover, in both tasks the vast majority of data points fall in the upper right and lower left quadrants of the figure, indicating that spatial and value selectivity are systematically related in both tasks. REW+ neurons respond strongest when the high-valued cue appears on the contralateral side, and REW– neurons respond strongest when the low-valued cue appears on the contralateral side. Plot style indicates the significance of selectivity for each neuron (see legend), as in **Figure 3b**. We note that the neurons were recorded in the left hemisphere in task A and the right hemisphere in task B, suggesting that this correspondence between value and spatial selectivity is true for both the left and right amygdalae.

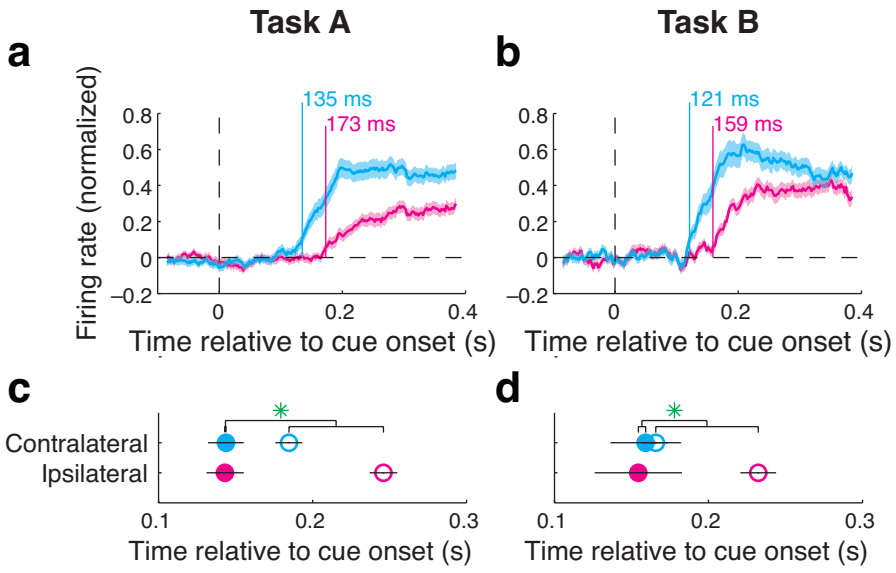

Supplementary Figure 4. Value information about contralateral stimuli, as compared to value information about ipsilateral stimuli, emerges more rapidly in the amygdala during both tasks

(**a,b**) Timing of value discrimination for the population of value selective neurons in Task A and Task B presented in the same manner as **Figure 4b** in the paper. Average firing rate differences between high value cue present trials and high value cue absent trials (high value cue on the contralateral side, cyan; high value cue on the ipsilateral side, red; shading indicates standard error) are plotted as a function of time from cue onset. As in the manuscript, firing rates differences for REW– neurons were sign-adjusted before averaging. Firing rate differences were normalized such as to maintain relative differences in the magnitude of the two signals. For both tasks, differential firing begins to occur more rapidly when the high value cue appears on the contralateral side (bootstrap, $P < 0.005$). (**c,d**) Mean latencies for the set of cells where the contralateral and ipsilateral visual onset latencies and/or the contralateral and ipsilateral value latencies could be estimated, separated by task (filled symbols correspond to visual onset latencies estimated from the fixation task, unfilled symbols correspond to value latencies estimated from the attention tasks, as in **Figure 5d** in the paper). Horizontal bars indicate the standard error for the distribution of single cell latencies. In both tasks, ipsilateral value latencies greatly lag contralateral value latencies, and this difference was significantly greater than the difference in visual onset latencies across spatial location (green asterisks; Wilcoxon, $P < 0.05$). Note that the visual onset latencies were characterized using the fixation task described in the manuscript.

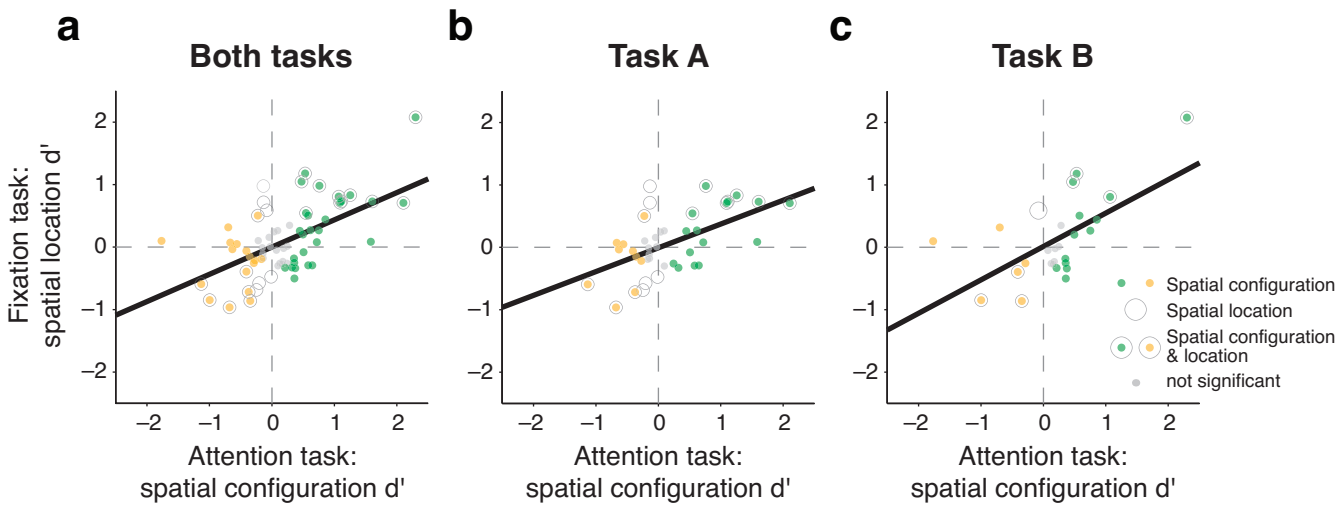

Supplementary Figure 5. Spatial configuration preference in the attention task is predictive of spatial location preference in the fixation task

(**a**) Relationship between spatial configuration selectivity in the attention task and spatial location selectivity in the fixation task for neurons recorded in both tasks (n = 68). We calculated d' values based on firing rates 100 to 500 ms after cue onset in each task. For the attention task, we compared responses on high value cue contralateral trials with responses on high value cue ipsilateral trials (as in **Fig. 3b**). For the fixation task, we compared responses when the stimulus appeared contralateral to those when the stimulus appeared ipsilateral. We used a weighted least squares regression to assess the relationship between the two sets of d' values, using the inverse standard error of d' (determined by bootstrapping) in the fixation task as the weights. The black line indicates the significant relationship between spatial location d' and spatial configuration d' (β = 0.44, C.I. 0.28 – 0.59; $P < 10^{-3}$, bootstrap analysis). The fact that the regression slope is less than 1 suggests that spatial encoding is weaker in the fixation task than in the attention tasks. The plot style of the data points indicates the significance of d' values for individual neurons (as indicated by the legend). (**b,c**) Data from each task presented separately in the same manner as (**a**). In each case, the regression slope is significant (bootstrap analysis, $P < 0.01$).

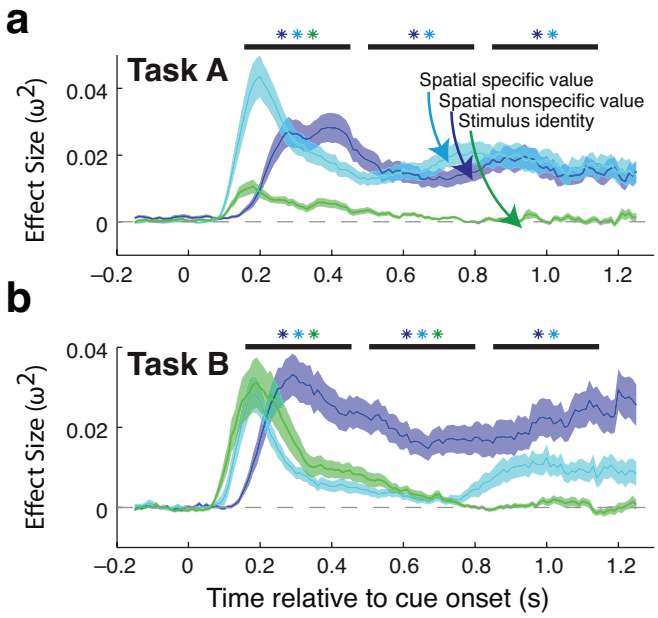

Supplementary Figure 6. Time course of spatially specific value, spatially non-specific value, and stimulus identity signals in each task

Time course of amygdala signals for the population of neurons recorded in Task A (**a**) and Task B (**b**). Signals averaged over the population of neurons for each task. Curves depict the mean and standard error (shaded region) of effect size (ω^2) measured across the population. Black bars indicate the time bins used for statistical analysis, and asterisks indicate that the distribution of ω^2 for that time bin was significantly greater than during the baseline period ($P < 0.05$). Stimulus identity coding is stronger in Task B and remains significant in the second time bin (500–800 ms) unlike in Task A; this might be due to differences in neuronal samples and/or monkey differences, among other reasons. However, this difference does not affect the main point of this figure, which is that stimulus identity selectivity is extinguished in the time period where the target is likely to appear (850–1150 ms), whereas value encoding is sustained during this same time.

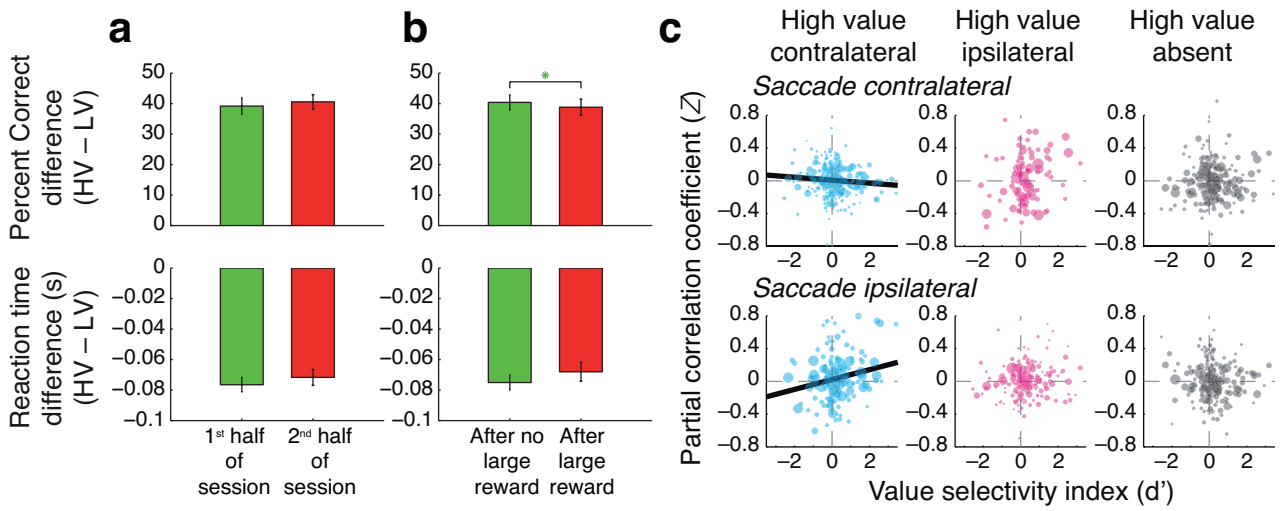

Supplementary Figure 7. Satiation and recent reinforcement outcome history do not explain correlations between firing rates and reaction times

In principle, correlations between firing rates and reaction times (**Fig. 7**) could have resulted from a drift in satiation level (which would cause performance to suffer as the session progresses) or a dependence of behavior on outcome history (where rewards on previous trials may cause a systematic changes in performance). Here, we present behavioral and neural analyses showing that neither factor could explain our results.

(**a,b**) Behavioral effects are similar between the first and second half of experimental sessions and between trials following low value and high value outcomes. For these analyses, we plot the difference in hit rate/reaction time when the target appeared near the high value cue location minus when it appeared near the low value cue location (high value present trials only). We were primarily interested in this behavioral difference since changes in performance that applied to both target locations could not explain the opposite signs of correlation in **Figure 7** of the main manuscript. For example, if satiation caused a non-spatial increase in reaction times over the course of the session (with a corresponding drift in firing rates), then the correlations for both saccade directions would have the same sign; a similar argument applies to a non-spatial change in behavior due to outcome history. Instead, we see opposite-signed correlations for contralateral and ipsilateral saccades suggesting that these results cannot be explained by a change in how the monkey allocates spatial attention due to satiation and/or outcome history. (**a**) Comparison of behavioral effects of between the first and second half of experimental sessions. The differences in hit rate (top) and reaction times (bottom) were not different between the first and second half of the session ($P > 0.18$). (**b**) Same as in (**a**) but comparing behavior on trials following either a low value or high value outcome. Here we found that the hit rate effect was greater following trials without a big reward ($P = 0.03$), although there was no difference in the reaction time effect ($P = 0.54$). Because the reaction time effect was not different, it is unlikely that differential behavior following large rewards could explain the reaction time correlations we observed (**Fig. 7**). (**c**) Replication of the analysis in **Figure 7a** using partial correlation to remove the influence of trial number and outcome history. This was accomplished using the matlab function 'partialcorr', where the controlled variables were the number of each trial within the session and whether the previous trial had resulted in a large reward (1) or not (0). The relationship between value selectivity and partial correlation coefficients was significant only for high value contralateral trials ($P < 0.05$ for contralateral and ipsilateral saccades).

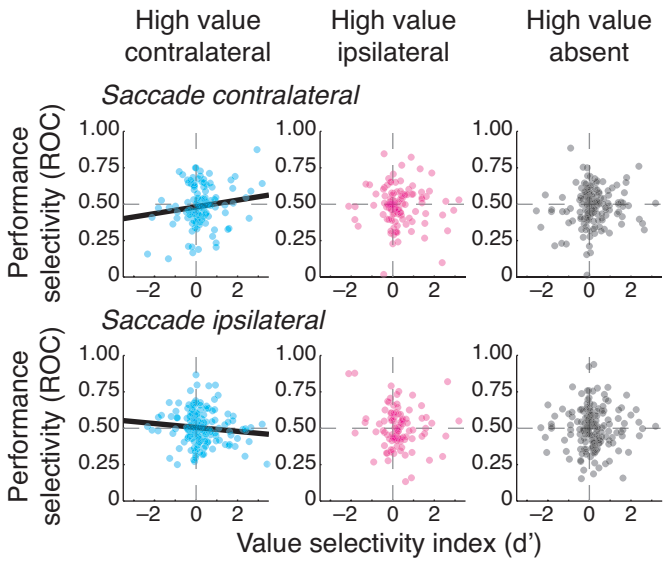

Supplementary Figure 8. Firing rate fluctuations predict correct performance

Here, we perform the same analysis as in **Figure 7a** with the exception that we compare firing rates between correct and error trials, as opposed to correlations between firing rates and reaction times. On the y-axis, we use an ROC analysis to compare firing rate distribution on correct and error trials (performance selectivity index), as opposed to the correlation coefficients used to assess reaction time correlations (as in **Fig. 7**). The relationship between value selectivity and performance selectivity indices was significant only on high value contralateral trials for trials where the target appeared contralaterally (P = 0.04) and ipsilaterally (P = 0.01); no significant relationships were observed on other trial types (P > 0.16). Note that here the signs of effects are opposite those of **Figure 7a** in the manuscript indicating that increases in firing rate are associated with improved performance. To test whether these results differed between tasks, we repeated the analysis including a task factor (1 for task A, 0 for task B) and did not observe a significant interaction effect of task by reward selectivity for any condition (ANOVA, P > 0.14).

Fundamental differences between the tasks make us cautious in drawing strong conclusions for the performance results. In task A, the majority of error trials were those in which the monkey failed to make any saccade after the target's appearance (67% of all errors); the other error trials consisted of saccades to the opposite hemifield as the target (28%) and saccades to the same hemifield that were not directed at the target (5%). On the other hand, correct trials always consisted of a saccade towards the same hemifield as the target. Thus, a difference in activity between correct and error trials may in part reflect a difference in saccade behavior on these trials related to motor preparation or the cognitive processes that accompany purposeful saccades. Limiting the data set to the 5% of error trials where saccades were directed to the same hemifield as in correct trials would at least control for the direction of the saccade but does not yield sufficient data for analysis. In task B, the monkey always made a saccade after the appearance of the choice targets; error trials were those in which he saccaded to the wrong choice target, indicating the wrong direction of orientation change. Although it is encouraging that we found similar results in the two tasks, we present these results with caveats given the differences in behavior on error trials. These issues do not affect the reaction time correlations (**Fig. 7a**), since all data used for that analysis included a movement.

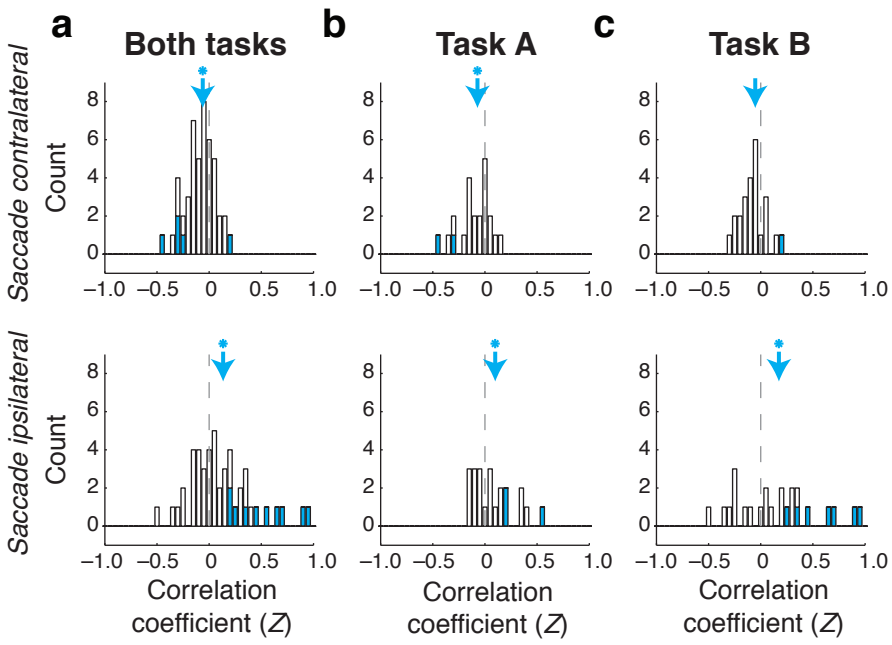

Supplementary Figure 9. Amygdala neural activity is correlated with spatial attention in both tasks

The use of two related, but distinct, tasks helps establish our results more firmly and makes clear their applicability as a robust phenomenon mediating behavior. To gain sufficient statistical power needed to demonstrate the presence of these effects on both tasks, we flipped the sign of the correlation coefficient for negative value neurons so that we could combine REW+ and REW– neurons. This procedure is analogous to what other investigators do when they combine data from neurons having different receptive field properties. (**a**) For the data combined across the two tasks we employed, the distributions of sign-corrected correlation coefficients were significantly less than zero for contralateral saccades (top) and significantly greater than zero for ipsilateral saccades (bottom; bootstrap analysis, $P < 0.01$); this figure is just a presentation of the same data presented in **Figure 7b** of the manuscript, but with the REW– and REW+ neurons combined after flipping the sign of correlation coefficients for REW– neurons. (**b,c**) shows the same analysis now applied to the 25% most value-selective neurons in each task in the same format as (**a**). Three of the distributions have means significantly shifted from 0 (bootstrap analysis, $P < 0.05$), and the last distribution has a mean nearly-significantly different from 0 ($P = 0.063$). Thus the basic effects we report are present in both tasks. Finally, we found no evidence that the magnitude of correlation coefficients was different for the two saccade directions; here, we compared the distribution of contralateral correlation coefficients in (**a**, top) with the flipped (multiplied by negative one) distribution of ipsilateral correlation coefficients in (**a**, bottom) and found no significant difference (bootstrap analysis, p = 0.57). Finally, the mean of the correlation coefficients did not differ between the positive and negative value neurons that went into (**a**) (bootstrap analysis, $P > 0.26$ for each saccade direction).

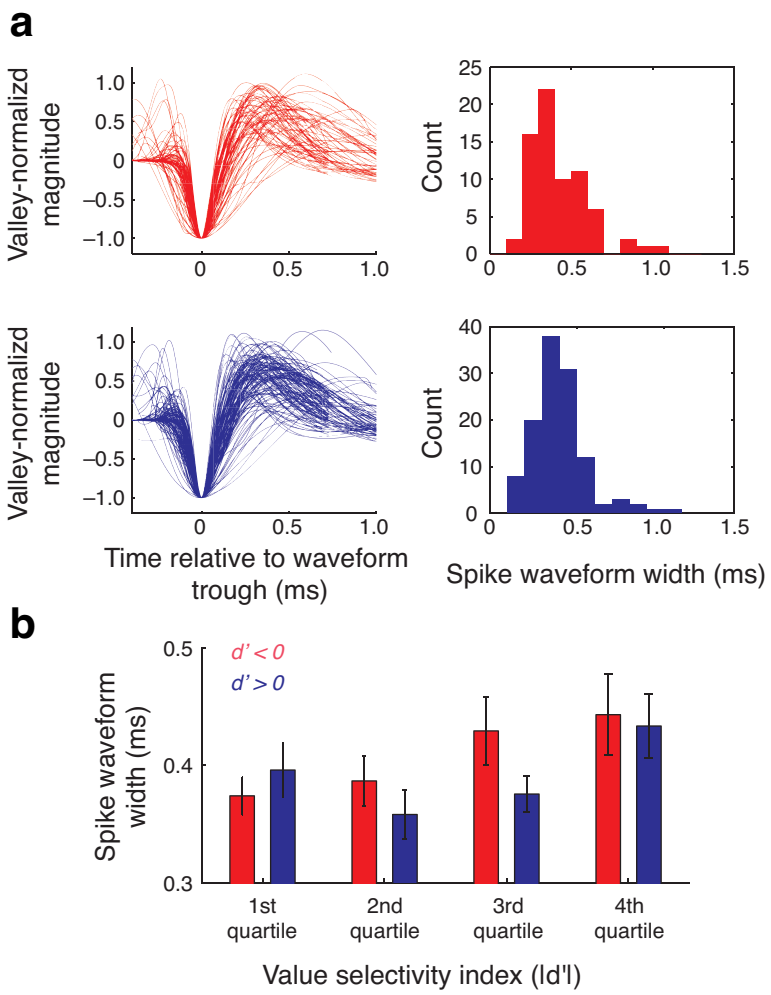

Supplementary Figure 10. Neural waveforms do not differ for REW+ and REW– neurons

(**a**) Average spike waveforms for individual neurons (left) and the distributions of spike waveform widths (right). Data are plotted for the populations of REW– neurons (top, red) and REW+ neurons (bottom, blue). Average waveforms were spline interpolated yielding a timing precision of 2.5 μs and the spike waveform width was defined as the trough to peak duration. A total of 10 cells were excluded from this analysis because either the trough or peak time could not be determined given the small time window around the spike saved during data acquisition. The distributions of spike waveform widths did not differ between REW– and REW+ neurons (Wilcoxon, $P = 0.66$). (**b**) Spike waveform width as a function of value selectivity magnitude. Neurons were split into quartiles based on the absolute value of the value selectivity d'; this was done separately for neurons with $d' < 0$ and with $d' > 0$. Vertical bars indicate the standard error of spike waveform widths for each group. Using a 2-way ANOVA, we verified that there was a significant (positive) relationship between the magnitude of value selectivity and waveform width ($P = 0.0381$); no significant effect was observed for the sign of selectivity ($P = 0.3020$) or an interaction between sign and magnitude ($P = 0.4464$).