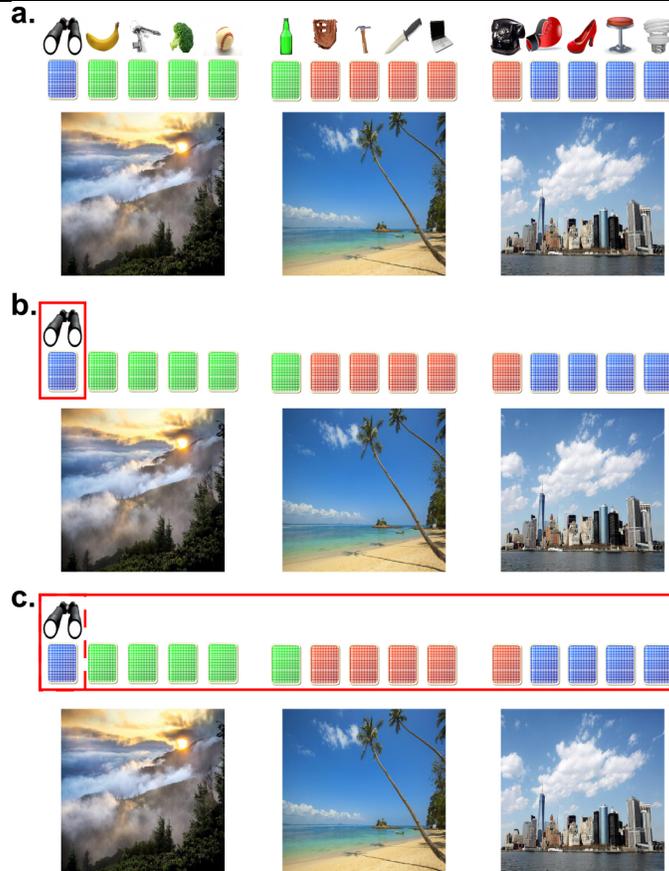


Supplementary Figure 1

**Context guides which episodes are sampled next.**

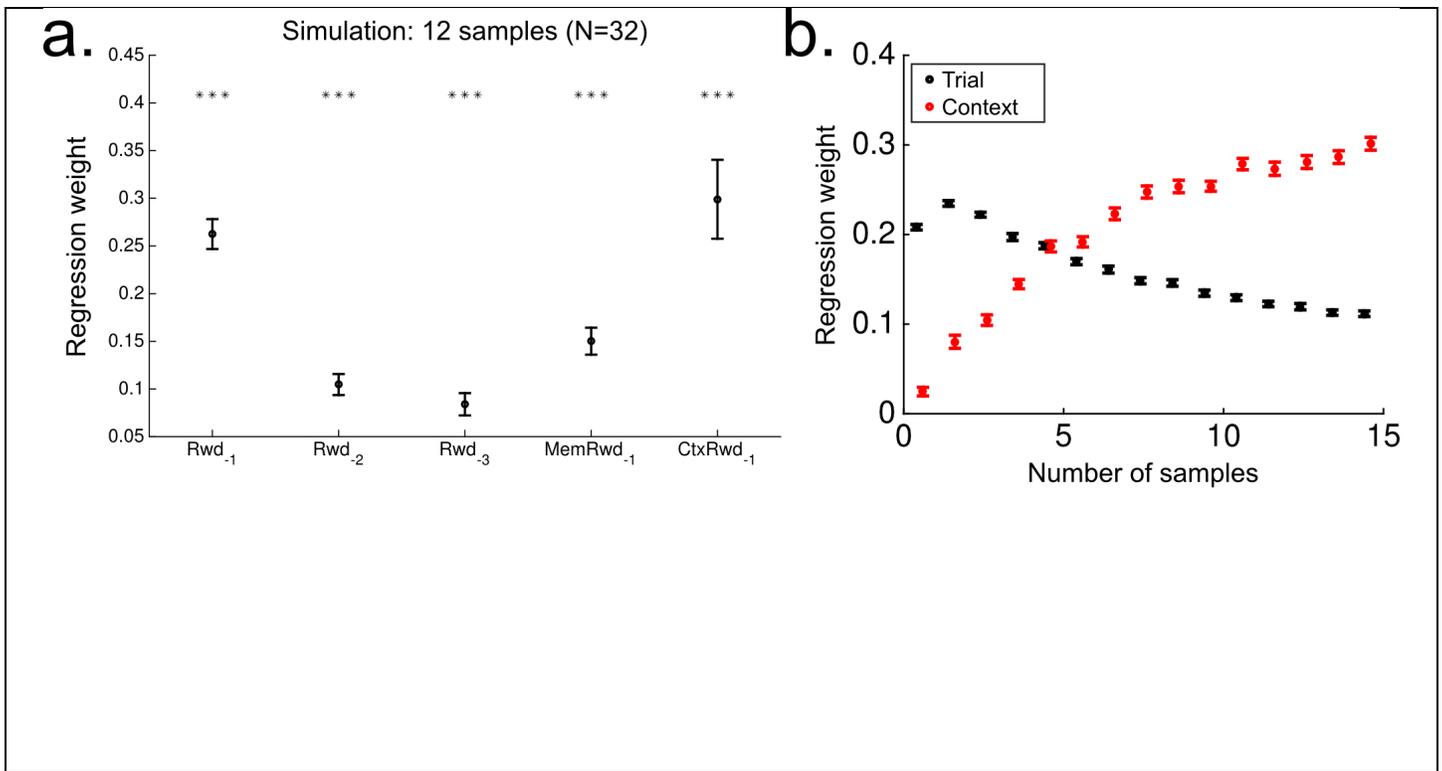
In the proposed mechanism, object images presented as memory probes cause participants to reinstate the trial episode on which the probed image was first encountered (first item in thought bubble). The reinstated episode carries with it the option chosen and reward received on that trial (\$10 bill or phase-scrambled bill reflecting \$0 reward). The episode also carries with it information about the context, or “casino room”, in which that trial took place. Reinstated context can lead to the subsequent reinstatement of other trials from that same context (ensuing items in thought bubble). Therefore, we predicted that memory probes will cause subsequent choices to be biased by both the rewards received on the reminded trial as well as those received on other trials in the same context.



**Supplementary Figure 2**

**Memory probes can provide multiple sources of value information.**

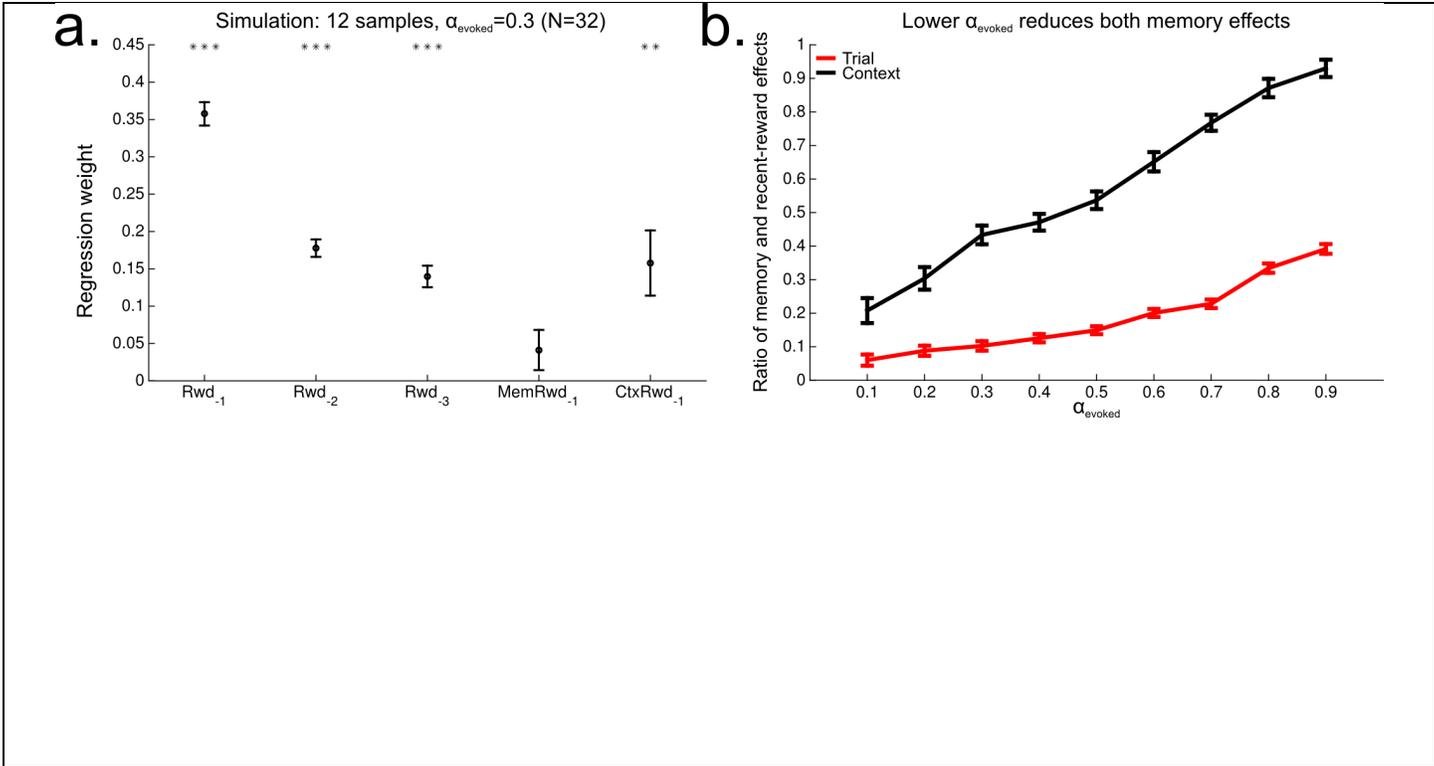
As shown in Figure S1, probes can trigger the recall of both the trial on which the probed image was received, and other trials in the same context. The payoff probabilities are designed such that during the initial 10 trials in each room, one deck is more likely to pay out a reward than the others, but that deck is not likely to pay out well in the last twenty trials of that room (Figure 2). Thus, the participant is likely to choose a different deck for the final twenty trials of each room than she did for the first 10 trials. This design feature allows us to distinguish effects of the reminded context from that of the reminded trial. **a.** First, each object image uniquely identifies one choice trial. **b.** This object image is episodically associated with the choice made, and reward received, on the given trial. **c.** However, the value information carried by subsequent retrievals of contextually-related trials is more likely to reflect trials on which a different option was chosen, and rewarded. We hypothesized that choices after a reminder would be influenced both by the rewards received on the reminded trial and the rewards received across the reminded context room.



**Supplementary Figure 3**

**Simulations with a context-aware sampling model.**

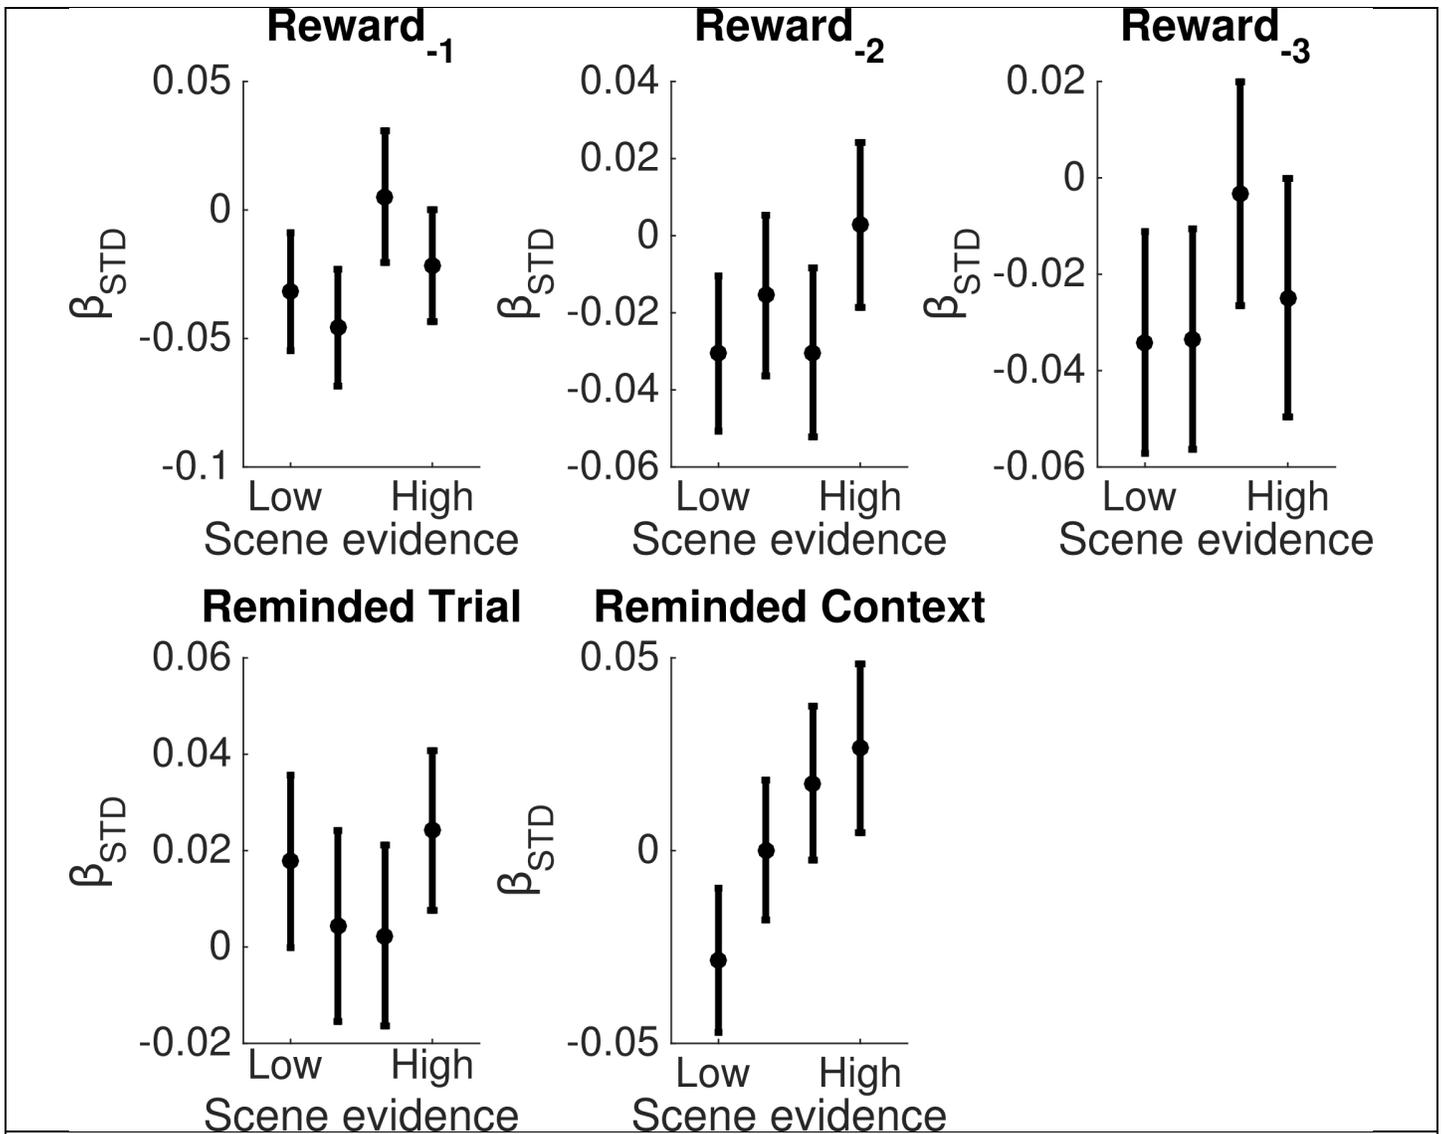
We ran simulations using a context-aware sampling model. The model reinstates a first episode with a probability related to its temporal recency, then some number of other episodes in turn. Each new reinstated episode is, with some probability, drawn from the same context as the previous reinstated episode. (See *Methods* for details of the model's implementation.) In this task, when more than a few samples are drawn, the context-aware sampling model predicts that the effect of reminded context should be greater than the effect of reminded trials. **a.** Shown here is the regression model from Figure 2, fit to a simulated population of 32 subjects that sampled 12 episodes between each probe and the ensuing choice. Error bars are +/- 1 SEM. **b.** Simulations show that the influence of context reward should increase with greater numbers of episodic samples. We simulated the context-aware sampling model to generate populations of 1,600 subjects (50 groups of 32), holding fixed all parameters except the number of samples drawn in support of each decision. We then performed the regression analysis shown in Figure 2 of the main text on this simulated data, and plot here the regression weights for single-trial (item) rewards and context rewards. As the number of samples increased, so did the ratio of the effect of context reward to that of the reminded trial (correlation between context and single-trial effects:  $R(13)=-0.9436$ ,  $P=1.3125e-07$ ). Error bars are +/- 1 SEM.



**Supplementary Figure 4**

**Simulations with imperfect memory reinstatement.**

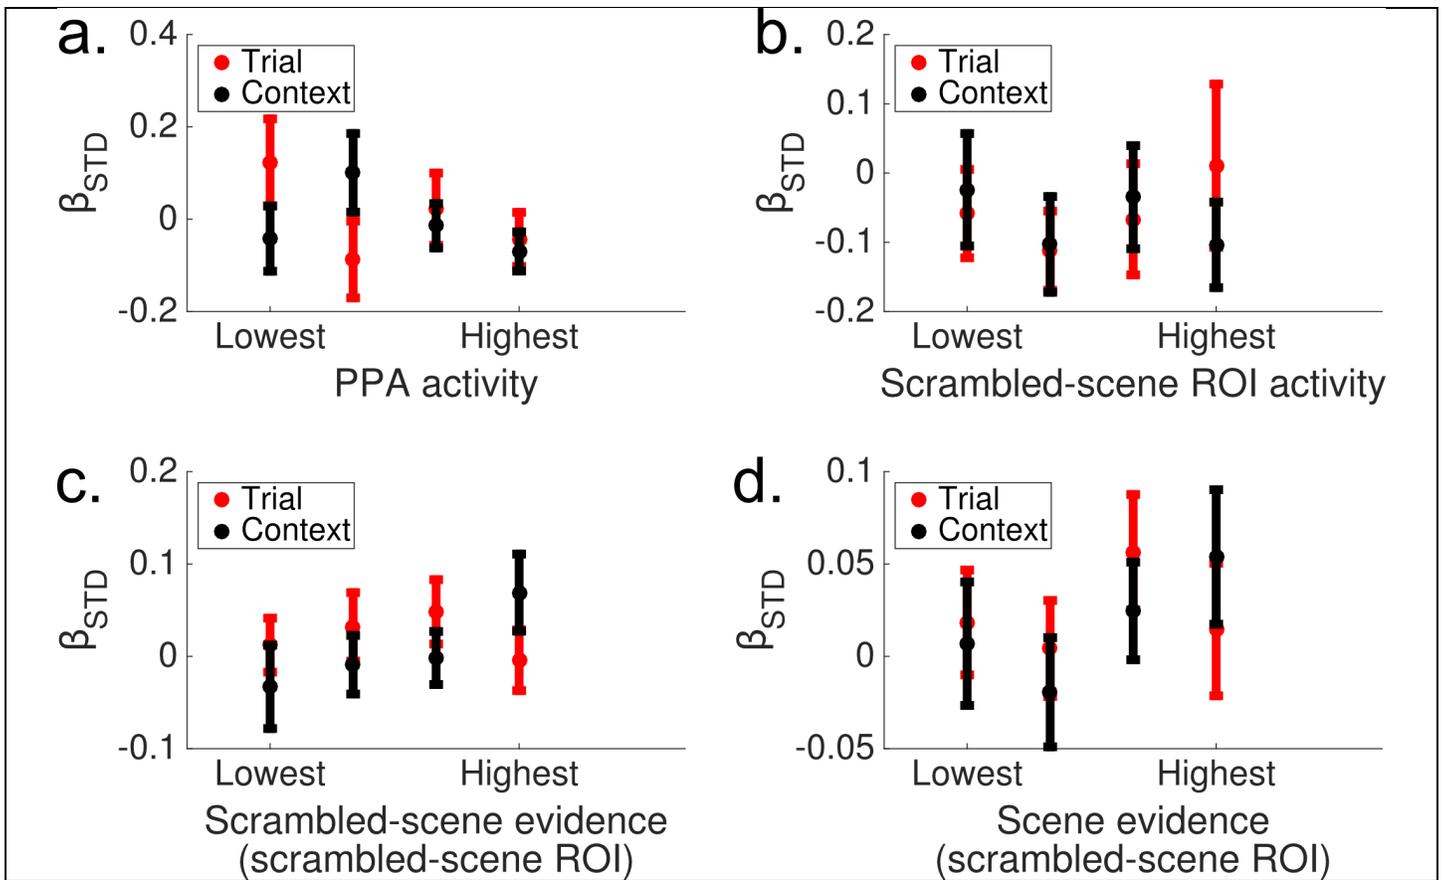
We used the simulation to illustrate how our mechanism can eliminate the effect of reminded trials while sparing an effect of recent rewards and context. We followed the simulation procedure of Figure S3, again generating populations of 1,600 subjects, this time with each population using a different value of  $\alpha_{\text{evoked}}$ , the probability that a reminder probe would result in reward information being reinstated from the reminded trial and context. All other model parameters were fixed. **a.** Shown here is the full regression model fit to a simulated population of 32 subjects that sampled 12 episodes before each choice, with  $\alpha_{\text{evoked}}$  set to 0.3. Error bars are  $\pm 1$  SEM. **b.** As the probability of reinstating reward information decreases, so does the contribution of both item and context memory to decisions. We ran the experiment for 9 populations of 1,600 simulated subjects each. Each population used a different value of the  $\alpha_{\text{evoked}}$  parameter. We then measured the regression weights for the reminded trial and the reminded context, plotted here by their ratio to the regression weight for reward experienced one trial ago. As  $\alpha_{\text{evoked}}$  decreased, so did the ratio of each type of memory-based effect to the effect of recent rewards. While the effect of reminded context is preserved at even very small values for  $\alpha_{\text{evoked}}$ , the effect of the reminded trial drops to near zero. This pattern matches the reduced effect of both reminded trial and context in Experiment 2. Error bars are  $\pm 1$  SEM.



**Supplementary Figure 4**

**Each regressor of interest plotted as a function of scene reinstatement evidence.**

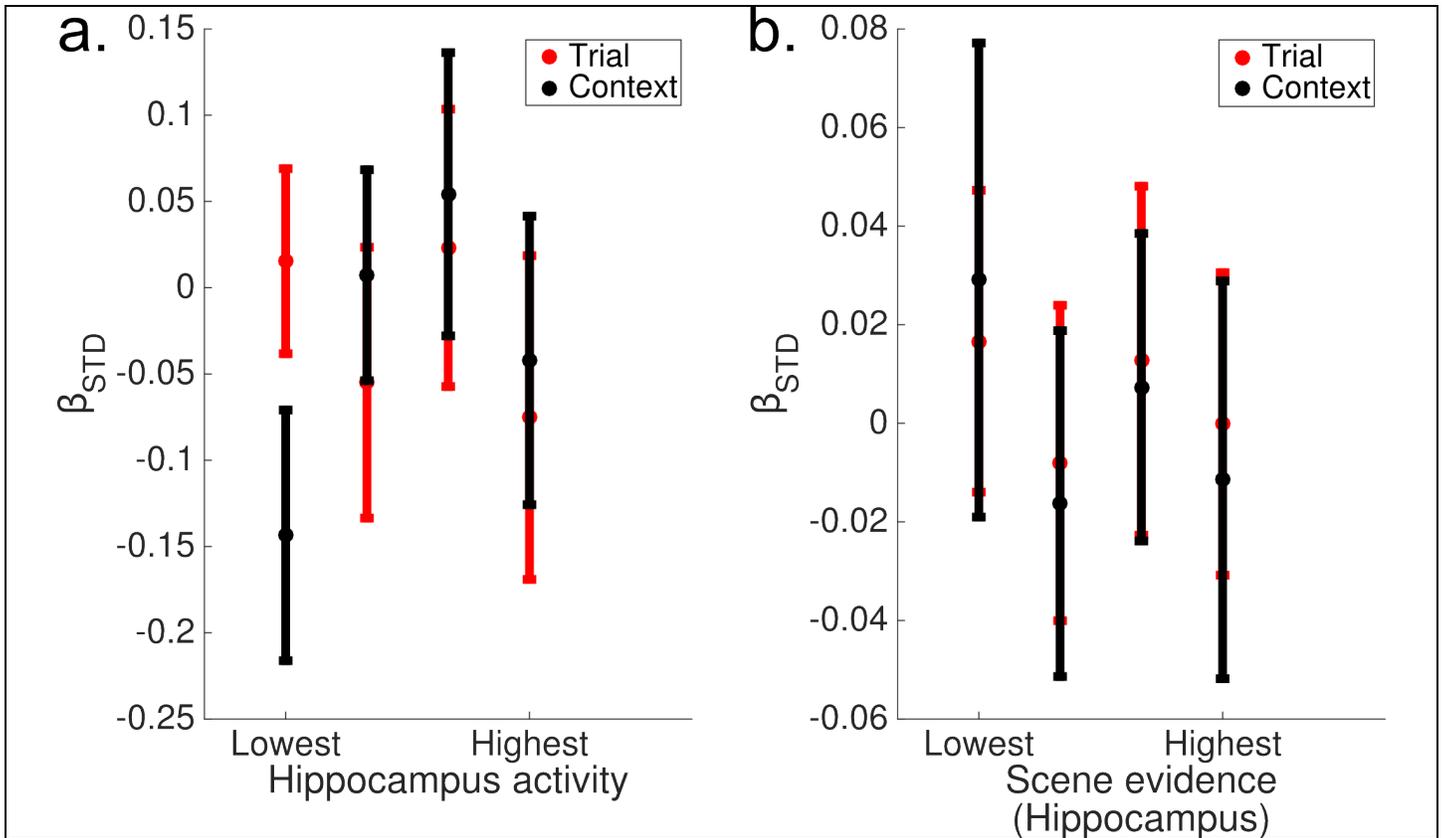
As described in the main text, we also tested whether regressors of interest other than context (recent reward, cued single-trial reward) were modulated by scene evidence, by repeating the quartiles analysis of Figure 3c for each regressor. We found that scene evidence did not reliably modulate any of the regressors of interest other than context ( $R_{-1}$ :  $t(31)=0.3188$ ,  $P=0.7520$ ;  $R_{-2}$ :  $t(31)=0.9716$ ,  $P=0.3388$ ;  $R_{-3}$ :  $t(31)=0.8135$ ,  $P=0.4221$ ; Reminded trial:  $t(31)=1.5052$ ,  $P=0.1424$ ).



Supplementary Figure 5

**Other activity does not predict behavioral effects.**

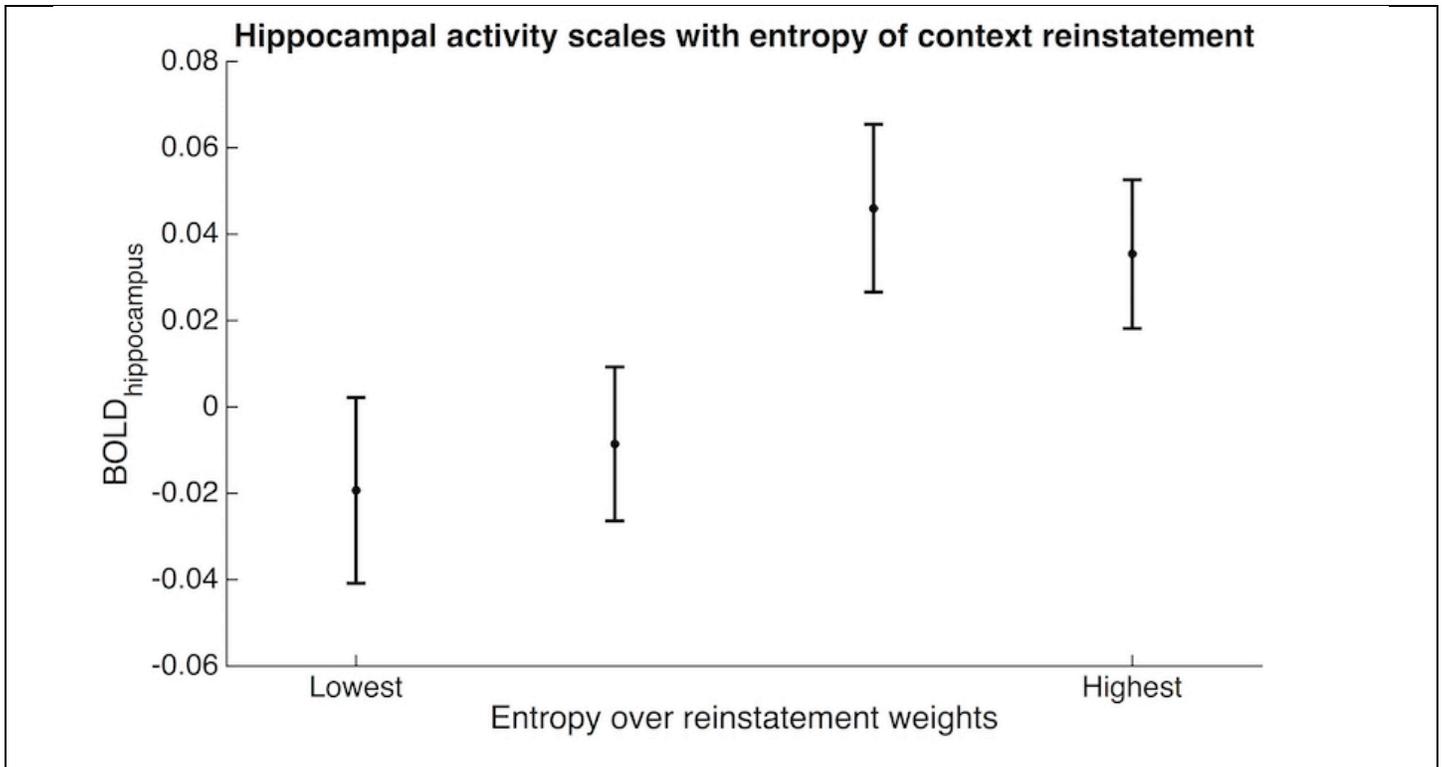
To confirm that context reward was specifically modulated by scene reinstatement in the PPA, we also repeated the quartiles analysis to look for an effect of activity or classifier evidence in several other regions of interest. **a.** Univariate activity in PPA. Across quartiles and subjects, neither reminded trial (mean slope=-0.0422, SEM 0.0397,  $t(31)=1.0205$ ,  $P=0.3154$ ), nor reminded context (mean slope=-0.0162, SEM 0.0184,  $t(31)=0.8765$ ,  $P=0.3875$ ) showed a relationship with PPA activity, nor did we observe an interaction between these relationships ( $t(31)=0.3136$ ,  $P=0.7559$ ). **b.** The same analysis was repeated for a region of interest that was differentially responsive to the “scrambled” scenes that were used as a control in our localizer task. Across quartiles and subjects, neither reminded trial (mean slope=0.0194, SEM 0.0479,  $t(31)=0.3949$ ,  $P=0.6956$ ), nor reminded context (mean slope=-0.0214, SEM 0.0212,  $t(31)=1.0085$ ,  $P=0.3210$ ) showed a relationship with univariate activity levels in this ROI, nor did we observe an interaction between these relationships ( $t(31)=0.6750$ ,  $P=0.5047$ ). **c.** We trained classifiers to decode “scrambled scene” evidence and scene evidence from this ROI. Neither reminded trial (mean slope=-0.0033, SEM 0.0160,  $t(31)=0.2048$ ,  $P=0.8391$ ), nor reminded context (mean slope=0.0314, SEM 0.0200,  $t(31)=1.5707$ ,  $P=0.1264$ ) showed a relationship with “scrambled scene” evidence, nor did we observe an interaction between these effects ( $t(31)=1.4886$ ,  $P=0.1467$ ). **d.** Neither reminded trial (mean slope=0.0041, SEM 0.0132,  $t(31)=0.3079$ ,  $P=0.7602$ ), nor reminded context (mean slope=0.0185, SEM 0.0158,  $t(31)=1.1678$ ,  $P=0.2518$ ) showed a relationship with scene evidence in this ROI, nor did we observe an interaction between these effects ( $t(31)=0.6751$ ,  $P=0.5046$ ). Error bars are +/- 1 SEM.



**Supplementary Figure 7**

Neither hippocampal activity nor classifier evidence predict trial or context effect.

The hippocampus has long been understood to be critical to episodic memory encoding and retrieval. We repeated the analysis of Figure 3c using an anatomically defined mask of bilateral hippocampus (two subjects' data excluded due to badly warped anatomical masks). **a.** Univariate hippocampal activity did not scale with either the single-trial reward effect (mean slope=-0.0231 SEM 0.0288,  $t(29)=-0.7754$ ,  $P=0.4444$ ), nor the context reward effect (mean slope=0.0277 SEM 0.0285,  $t(29)=0.9657$ ,  $P=0.3422$ ), nor was there an interaction between the effects ( $t(29)=1.0715$ ,  $P=0.2928$ ). **b.** We trained a classifier to discriminate scenes in this region, and did not observe a reliable relationship between hippocampal scene evidence and either reminded trial (mean slope=-0.0069 SEM 0.0140,  $t(29)=-0.4898$ ,  $P=0.6280$ ) or reminded context reward effects (mean slope=0.0045 SEM 0.0148,  $t(29)=0.3010$ ,  $P=0.7656$ ). Error bars are  $\pm 1$  SEM.



**Supplementary Figure 8**

**Activity in bilateral hippocampus scales with entropy over scene reinstatement evidence.**

We investigated the hypothesis that hippocampal activity could reflect retrieval of memories in support of decisions [1,2]. We used the scene-specific reinstatement weights to investigate activity related to memory retrieval. In previous studies, we observed that hippocampal activity increases along with uncertainty about an action's outcome, both for simple sequential responses and goal-directed planning decisions [3,4]. We interpreted those findings as consistent with hippocampus' known role in memory retrieval. In the context of action evaluation, memory retrievals could constitute evidence about the outcome of the actions under consideration (similar to how forward trajectories are "replayed" as rodents make navigation decisions [5]). In this task, greater uncertainty about the associated context should lead to a wider range of next-step outcomes to evaluate. We therefore reasoned that activity in hippocampus at choice might scale with uncertainty about the probed item's context. We computed the evidence that participants reinstated each context image on each trial, by taking the correlation between per-scene template patterns and PPA activity on that trial. The resulting six numbers were then normalized to create a probability distribution for that trial, reflecting the relative likelihood that each scene was being remembered. The entropy over this distribution can thus be considered the uncertainty over the reinstated context. Consistent with a role for hippocampus in retrieving memories that are used to evaluate outcomes, this entropy value is reliably correlated with hippocampal activity at the time of choice (mean  $R=0.0670$ , SEM 0.0261,  $t(29)=2.5733$ ,  $P=0.0155$ ). Error bars are +/- 1 SEM.