Supplemental data items

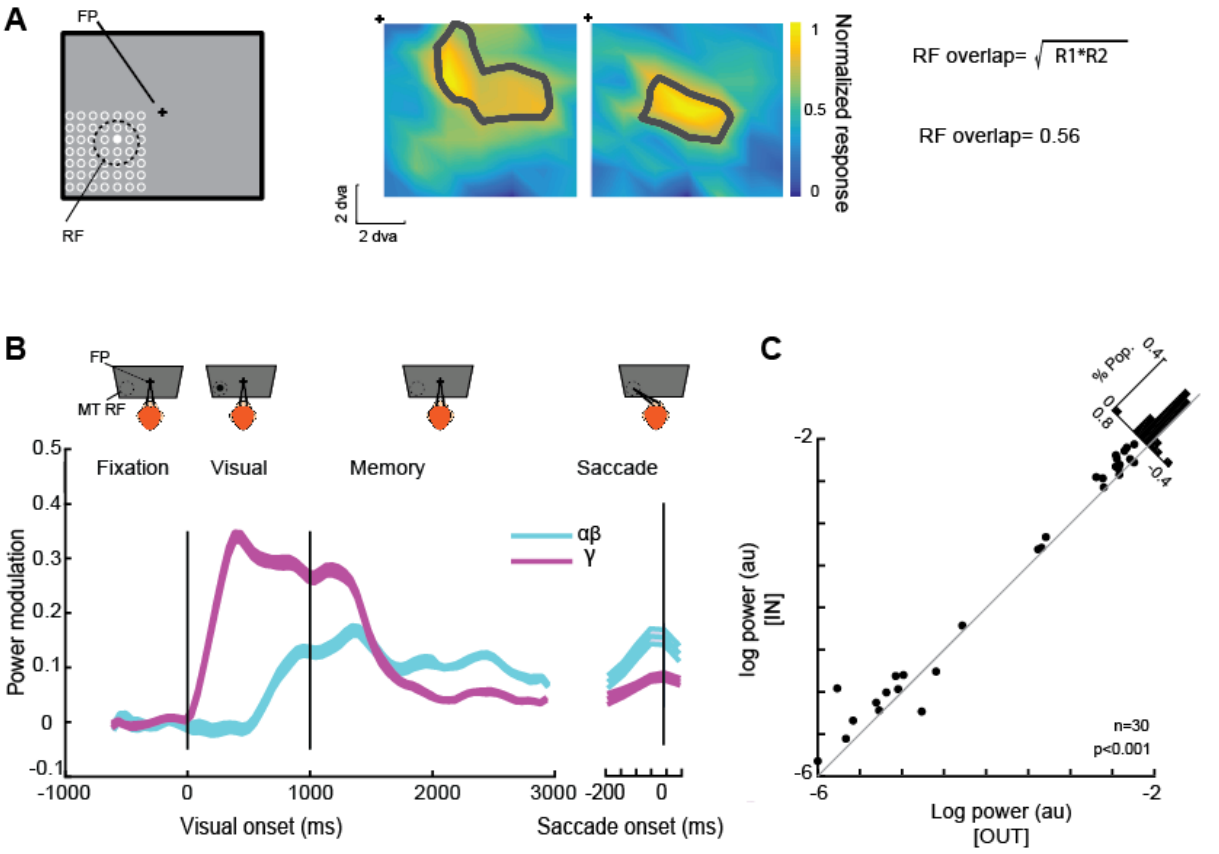

Figure S1. Receptive field (RF) mapping and changes in LFP power across time and sessions. Related to Figures 1& 7. A) In the RF mapping task (left), the monkey fixates on the fixation point (FP) in the center of the screen, while visual probes (filled white circle) flash in a grid covering the estimated RF location; potential probe locations are indicated by open white circles and not actually shown onscreen. For each neuron, an RF profile is constructed based on locations in which the visually-evoked activity is at least 50% of the maximum probe-evoked response; plots on the right show the normalized responses and RF profiles of two example neurons. The RF overlap between two neurons is measured based on these RF profiles using the formulas on the right. B) The time course of LFP power modulation (LFP power of each condition is normalized to the pre-stimulus baseline) in the alpha-beta (αβ; blue) and gamma (γ; purple) bands over the course of a trial on the MGS task, aligned to the time of stimulus onset (left) or saccade onset (right). γ-band power increased immediately after cue presentation in the memory IN condition compared to OUT, remained elevated throughout the visual period (n = 480, Δ γ power = 0.299 ± 0.025, $p < 10^{-3}$, Wilcoxon signed-rank test, then dropped following removal of the cue, but remained slightly elevated compared to memory OUT trials during the memory period (n = 480, Δ γ power = 0.056 ± 0.013, $p < 10^{-3}$, Wilcoxon signed-rank test, memory period is considered from 500-1000 ms after stimulus offset for fixed-delay sessions and 1000-1500 ms after stimulus offset for variable delay sessions (see STAR Methods)). The αβ power was significantly elevated for memory IN trials beginning a few hundred milliseconds before the start of the memory period, and this increase was maintained during the memory period. αβ power also increased just prior to initiating a saccade toward (vs. away from) the RF. The shaded areas show the standard error across all recorded LFP signals. C) Because of the similarity between channels recorded in a single session, we examined the significance of the αβ power increase across sessions by averaging the power magnitude of all recorded channels in each session. We observed an increase in αβ power when comparing session averages for memory locations inside vs. outside the MT RF (the values are converted to log space for better illustration, n = 30, Δ αβ power = 0.116 ± 0.035, $p < 10^{-3}$, Wilcoxon signed-rank test).

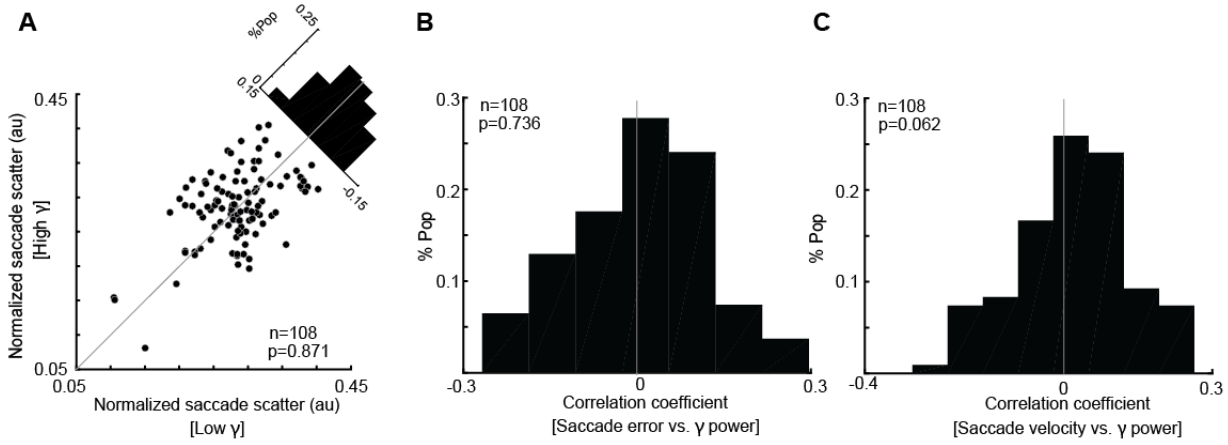

Figure S2. γ-band power during the memory period is not correlated with behavior. Related to Figure 2. A) A comparison of saccade scatter for high γ power trials vs. low γ power trials, across all γ-memory-selective LFP signals (n = 108), shows no significant difference in saccade scatter based on the magnitude of γ power (Δ saccade scatter = 0.001 ± 0.005, p = 0.871). B) The distribution of correlation coefficients between saccade error and γ power, across all γ-memory-selective LFP signals, shows a mean correlation which is not significantly different from zero (n = 108, $r_{error\&\gamma power} = 0.001 \pm 0.012$, p = 0.736). C) The distribution of correlation coefficients between saccade velocity and γ power, across all γ-memory-selective LFP signals, shows a mean correlation not significantly different from zero (n = 108, $r_{velocity\&\gamma power} = 0.020 \pm 0.011$, p = 0.062 There were no memory-selective LFP signals for frequencies less than 8 Hz, and so we did not attempt to measure behavioral correlations in these lower frequencies.

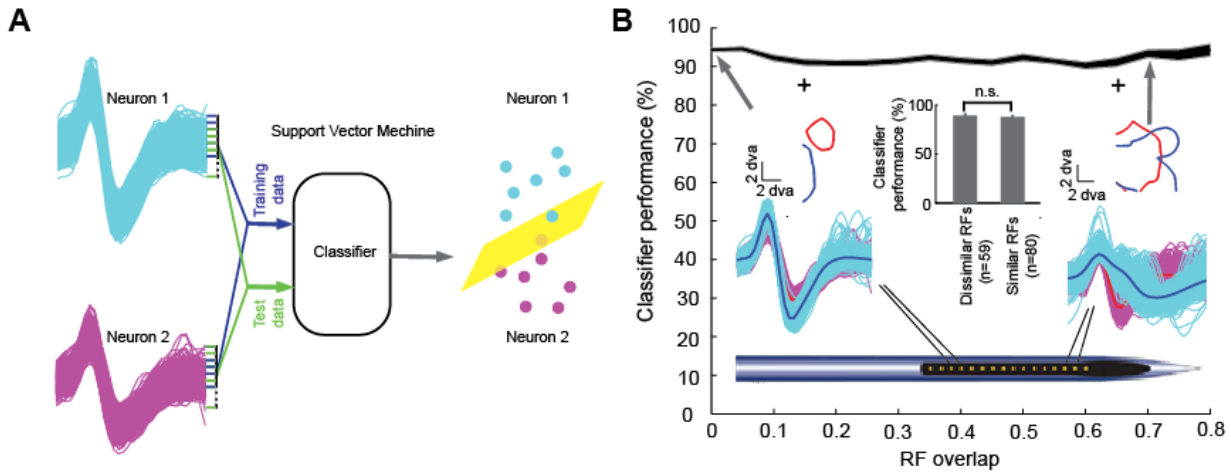

Figure S3. SVM control for the quality of neuronal isolations. Related to Figure 7. A) Schematic of the SVM classifier used to test the isolation quality of two spike waveforms recorded during the same session (n = 1092 pairs). In order to make sure that the manually sorted neurons are well isolated, we used Support Vector Machine (SVM) classifiers with linear kernels to categorize the spikes of simultaneously recorded pairs of neurons. To do that, we applied the SVM to a window of 200 spikes from each neuron in a pair, and slid that 200-spike window through the recording session. Within each window, we randomly selected 100 spikes from each neuron to train the classifier (training data, blue) and used the remaining 100 spikes to test the classifier performance (test data, green). The SVM classified each waveform as belonging to one of the two neurons, and classifier performance was compared to the assignment based on manual spike sorting. We averaged the classifier performance across different windows over the session to quantify how well the spikes from each neuronal pair are isolated. B) Classifier performance is high and independent of RF overlap between the neurons. The main plot shows the average classifier performance in categorizing the spike waveforms of simultaneously recorded neurons as a function of the overlap between the neurons' RFs. The plot was obtained using a 0.1 sliding bin on the x-axis. The overall high performance of the classifier indicates that the spikes of neuronal pairs are well isolated (the mean classifier performance across all the pairs = 92% ± 0.27, n = 1092, $p < 10^{-3}$ compared to chance level, i.e. 50%). Classifier performance is high across a range of shared RF ratios, and the dependence of classifier performance on the overlap between neurons' RFs is very slight (Pearson correlation; correlation coefficient = -0.06, p = 0.03, n = 1092). The RF contours and spike waveforms of two sample pairs which were recorded on adjacent contacts of the array electrode are shown. The RFs of the sample pair shown on the left (recorded by contacts #13 and #14) have no overlap, while the RFs of the sample pair shown in right (recorded by contacts #2 and #3) are highly overlapping (gray arrows indicate the average classifier performance at the RF overlap values of these pairs). For further analysis, we focused on all the neuronal pairs which were recorded on recording contacts in close proximity to one another (distance < 200 µm; n = 139 pairs), and grouped them into similar (RF overlap > 0.5) and dissimilar RFs based on their RF overlaps (see Method). The bar plot shows the mean classifier performance in categorizing the spike waveforms of pairs of neurons with similar and dissimilar RFs. The result shows that there is no significant difference between the classifier performances for these two groups of neuronal pairs (the mean classifier performance for pairs with similar RFs = 88% ± 1.36, n = 80, the mean classifier performance for pairs with dissimilar RFs = 89% ± 1.4, n = 59, p = 0.757, Wilcoxon rank-sum test).

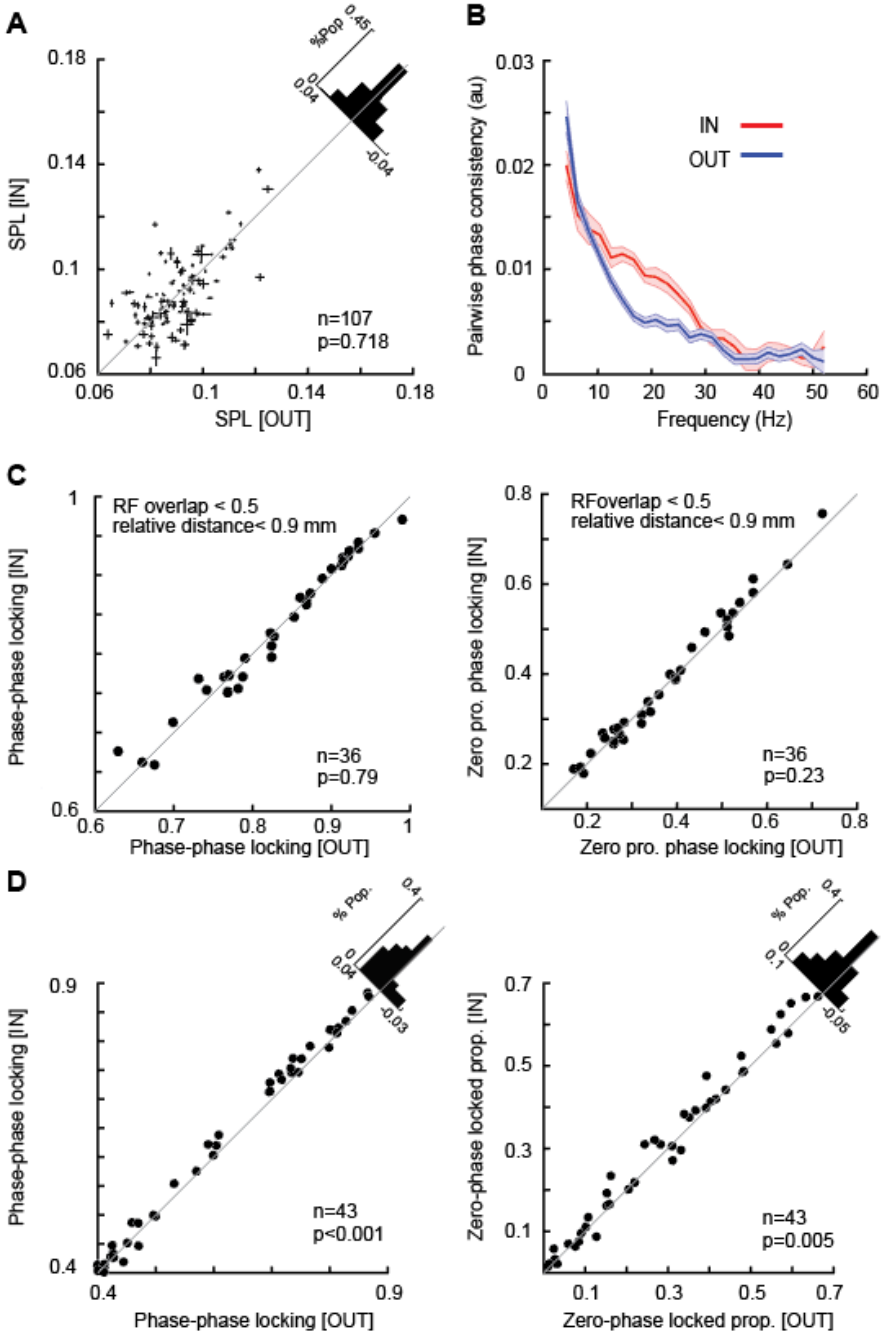

Figure S4. Controls for SPL, PPL and ZPL. Related to Figures 3 & 7. A) The SPL, computed after spike and LFP data from different trials of the same condition were shuffled, for memory IN vs. OUT. This shuffled measure of SPL was not different for memory IN vs. memory OUT (n = 107, Δ shuffled SPL = 0.001 ± 0.0001, p = 0.718). This shows that the temporal relationship between the spike time and LFP phase in a single trial is necessary to observe the WM-induced change in SPL (Figure 3C). Additionally, when we quantified the SPL using an artificial set of spikes based on a Poisson process (matching the firing rate of the neuron in each condition, but with a random timing, Dayan & Abbott, 2005) and an artificial set of LFPs (matching the power of αβ frequency for each condition, Prichard & Theiler, 1994) we did not find a difference between the IN and OUT conditions (n = 107, Δartificial SPL = -0.453 ± 2.347 x10^{-4}, p = 0.810). This analysis shows that the experimentally observed change in SPL was not a mere consequence of the changes in LFP power in the αβ band. B) The pairwise phase consistency is a bias-free method for measuring the spike

phase locking (Vinck, van Wingerden, Womelsdorf, Fries, & Pennartz, 2010). We calculated the consistency using the FieldTrip toolbox (Oostenveld, Fries, Maris, & Schoffelen, 2011). There is a significant increase in the spike-phase locking in the αβ frequency range for memory IN compared to OUT measured using this method (Δ pairwise phase consistency = 0.004 ± 0.0006, $p < 10^{-3}$). This analysis shows that the observed SPL changes were not due to nonsignificant differences in the firing rate or number of trials between IN and OUT conditions. C) Memory-induced changes in PPL and ZPL depended on RF overlap, not cortical distance. We identified pairs of LFP recordings with RF overlap < 0.5, for which the mean distance of electrodes are the same as the mean distance of electrodes with similar RFs. The amplitude of phase-phase locking (PPL) between αβ oscillations on these nearby channels with dissimilar RFs was not significantly different when the memory location was inside the MT RF vs. outside MT RF (n = 36, Δ PPL= 0.003 ± 0.002, p = 0.790). The proportions of phase-phase locked trials with zero phase lag was also not significantly different based on memory location for the memory IN condition compared to OUT (n = 36, Δ ZPL= 0.004 ± 0.003, p = 0.236). This indicates that RF overlap, rather than cortical distance, was the critical factor in determining whether WM would induce greater synchrony and coherence between the αβ oscillations at two sites. D) Volume conduction control. After reducing the shared signal between channels from each channel using spatial filtering (Srinath & Ray, 2014), the amplitude of PPL between αβ oscillations on channels with similar RFs was still significantly different when the memory location was inside the MT RF vs. outside MT RF (n = 43, Δ PPL = 0.0092 ± 0.0021, $p < 10^{-3}$). The proportions of phase-phase locked trials with zero phase lag was also significantly different based on memory location for the memory IN condition compared to OUT (n = 43, Δ ZPL = 0.0133 ± 0.0043, p = 0.0055). Again, we found a robust relationship between how coherently two channels oscillate (PPL) and the degree to which the RFs of neurons from these sites overlap in visual space (correlation between controlled PPL and RF overlap in two conditions $r_{OUT} = 0.394$, $p < 10^{-3}$, $r_{IN} = 0.405$, $p < 10^{-3}$ & $r_{\text{Overlap \& }\Delta\text{PPL}} = 0.143$, $p < 10^{-3}$, Pearson correlation).

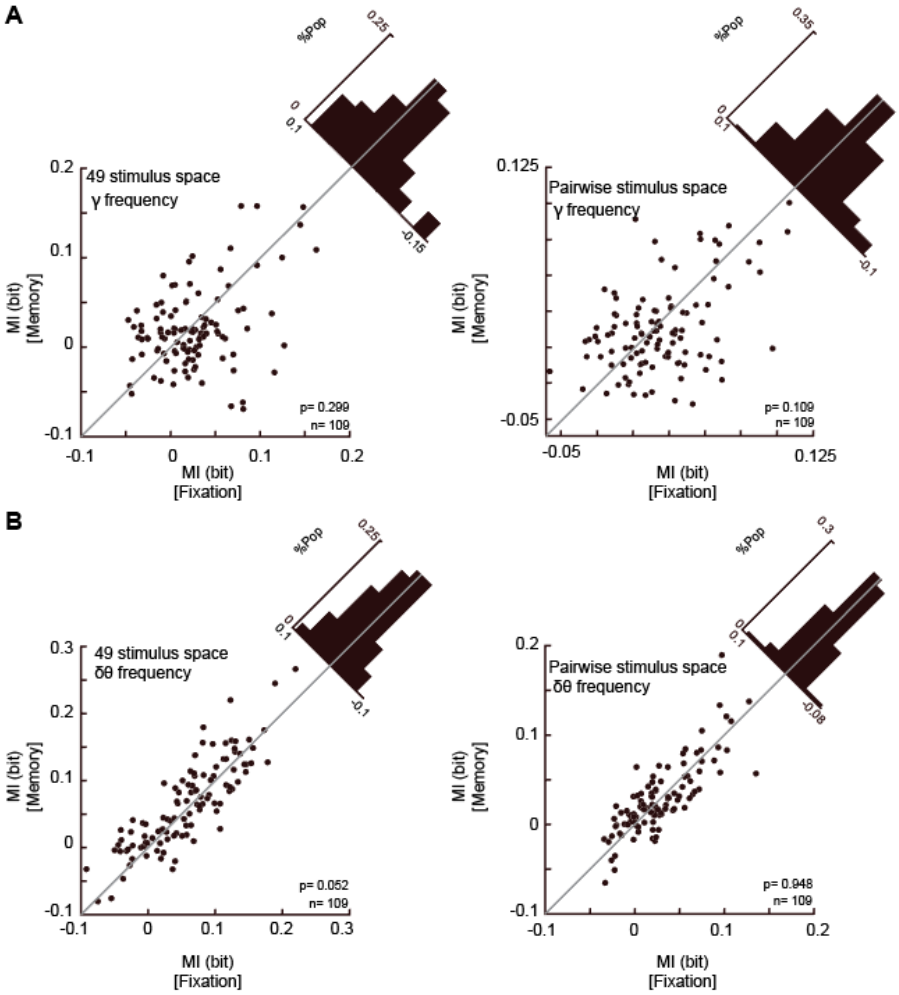

Figure S5. WM-induced enhancement of mutual information (MI) does not occur in the γ or delta-theta (δθ) frequencies. Related to Figure 5. A) There was no significant increase in the MI between probe location and the phase of ongoing γ oscillations at the time of spikes, either when calculated using all 49 probe locations (n = 109, MI_{memory} = 1.976 ± 0.0004, $MI_{fixation}$ = 2.649 ± 0.0004, p = 0.299, 8% of sites showed an individually significant increase in MI, and 8% significantly decreased), or for combinations of probe pairs (n = 109, MI_{memory} = 1.395 ± 0.0003, $MI_{fixation}$ = 1.932 ± 0.0003, p = 0.109, 20% of sites showed an individually significant increase in MI, and 27.5% significantly decreased). B) There was no increase in the MI between probe location and the phase of ongoing δθ oscillations at the time of spikes, calculated using all 49 probe locations (upper right, n = 109, MI_{memory} = 0.068 ± 0.001, $MI_{fixation}$ = 0.059 ± 0.001, p = 0.052, 11% of sites showed an individually significant increase in MI, and 8% significantly decreased) or based on all combinations of probe pairs (lower right, n = 109, MI_{memory} = 2.754 ± 0.0003, $MI_{fixation}$ = 2.705 ± 0.0003, p = 0.948, 24% of sites showed an individually significant increase in MI, and 23% significantly decreased).

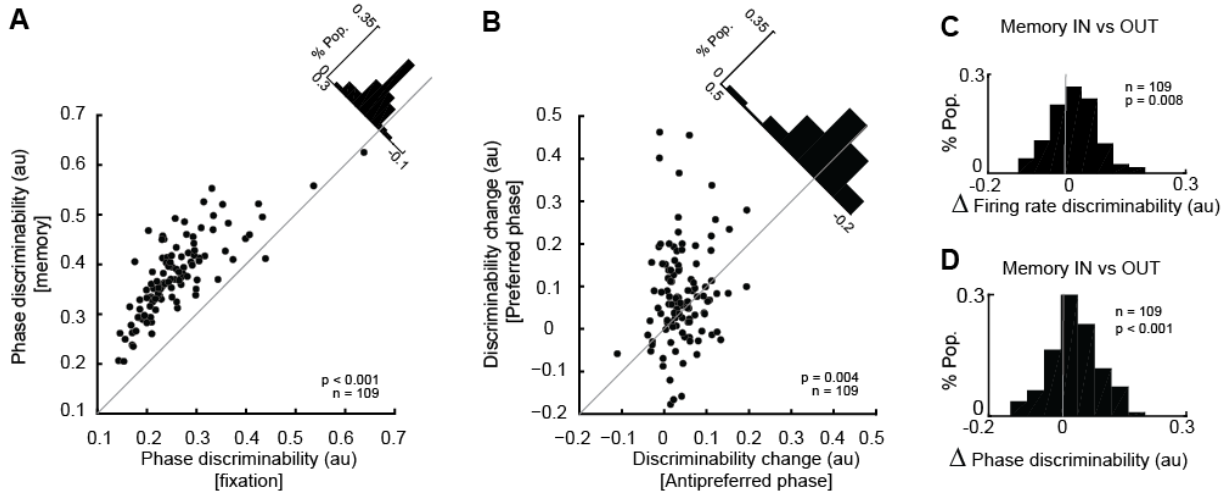

Figure S6. Control for different number of samples, spike rate & arousal for measuring discriminability. Related to Figure 6. A) Sample size control for WM-induced changes in discriminability based on phase. For measuring discriminability based on firing rate, we have a sample of the spike count for each trial, and the total number of samples is equal to the number of trials. For measuring discriminability based on the phase at spike times, we have a sample at the time of each spike, so the total number of samples for this dataset will be larger. Here we randomly chose a subset of samples from the phase dataset, equal to the number of trials for the firing rate method, and then measured the discriminability on this subset of samples. We repeated this procedure 1000 times and averaged the discriminability values over all repetitions. The discriminability calculated based on phase at spike times is still significantly greater for memory IN vs. fixation ($n = 109$, $d'_{fixation} = 0.259 \pm 0.007$, $d'_{memory} = 0.376 \pm 0.007$, $p < 10^{-3}$, Wilcoxon signed-rank test). B) A control for the difference in spike counts between the preferred and anti-preferred phases when measuring the change in discriminability (Figure 6D). We randomly deleted spikes from the preferred phase data set to equalize the mean firing rate across probes between the preferred and anti-preferred phase data, and then measured the discriminability. We repeated this procedure 100 times and averaged the discriminability values over all repetitions. The discriminability change is still significantly greater for preferred phase vs. anti-preferred phase ($n = 109$, $\Delta d'_{preferred\text{-}phase} = 0.080 \pm 0.011$, $\Delta d'_{anti\text{-}preferred\ phase} = 0.043 \pm 0.005$, $p = 0.004$, Wilcoxon signed-rank test). C) To control for differences in arousal between the memory and fixation periods, we compared discriminability during the memory period between the two memory conditions. We found that when the d' index is measured using the firing rate, the discriminability of probe locations is greater during memory IN than OUT ($n = 109$, $\Delta d' = 0.042 \pm 0.006$, $p = 0.008$, Wilcoxon signed-rank test). D) Quantifying the d' index using the spike phases also revealed significant greater discriminability during memory IN than OUT ($n = 109$, Δ discriminability $= 0.063 \pm 0.005$, $p < 10^{-3}$, Wilcoxon signed-rank test).

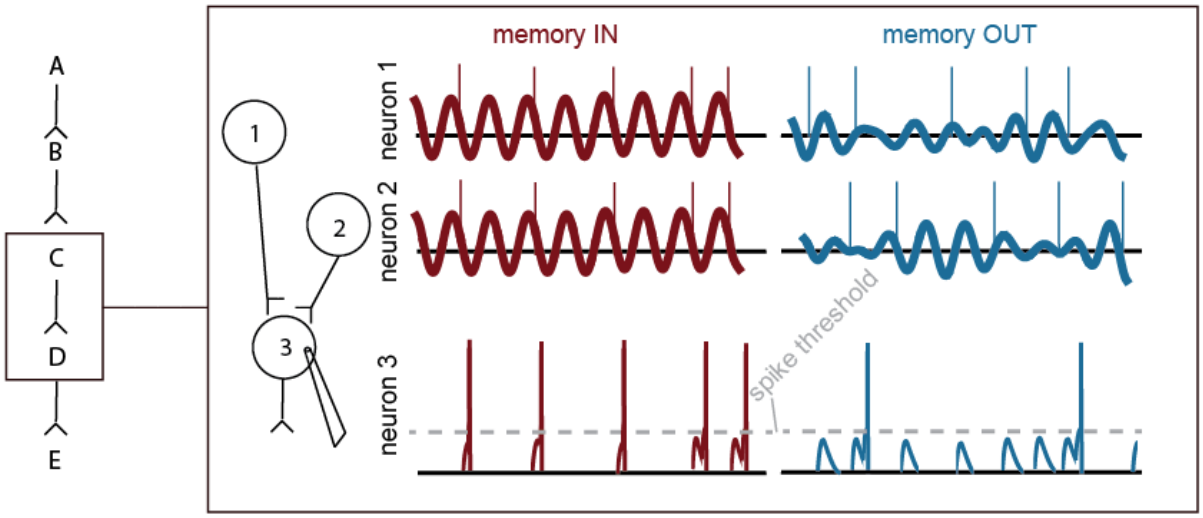

Figure S7. Conceptual model for how changes in phase-locking at each stage of neuronal processing can create firing rate changes in downstream neurons. Related to Figure 7. For one stage of neuronal processing (C synapsing onto D, left), we illustrate the effect of phase-locking changes in two presynaptic neurons with similar RFs (neurons 1 and 2), both providing input to a postsynaptic neuron (neuron 3). The spiking activity and LFP oscillations for the presynaptic neurons are shown in the top two rows, for conditions when a memory is being maintained inside (memory IN, left, red) vs. outside (memory OUT, right, blue) their RFs. These illustrations show several of the effects observed in experimental data for αβ oscillations: 1) During memory IN, oscillations are stronger. 2) During memory IN, oscillations are phase-locked across sites with similar RFs. 3) During memory IN, spikes at these sites are more phase-locked to the oscillations. The bottom row shows the evoked excitatory postsynaptic potentials (EPSPs) and spiking activity for the postsynaptic neuron. In this schematic, neuron 3 must receive two incoming spikes in close temporal proximity in order for the EPSPs that these spikes produce to combine and exceed the threshold for generating a spike in neuron 3 (spike threshold, dashed grey line). During the memory IN condition, spikes from the two input neurons are more likely to arrive close together in time, due to the phase locking of these spikes with the oscillation. This produces a higher firing rate in neuron 3 during the memory IN condition, despite firing rates in the two conditions being the same for the two presynaptic neurons. Note that the spikes evoked in the postsynaptic neurons are themselves phase locked to the same oscillation as the input neurons. Similar changes in spike timing could occur at every synaptic relay (left), including those where firing rate changes also occur, and have an additive effect across multiple levels of the cortical hierarchy. It is also possible that such changes in spike timing would have subthreshold effects on the downstream neuron for low firing rates (for example, in the absence of visual input), and that the effects of these spike timing changes would only become visible in the presence of a stronger input signal.