Supplementary Figure 1

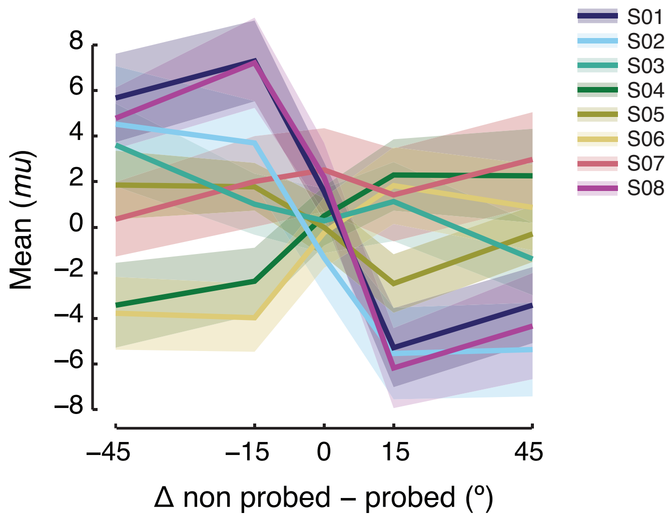

Supplementary Figure 1. When reporting the second of two to-be-remembered orientations, different participants show different biases relative to the orientation of the item that was memorized first. For participants 1, 2, 3 and 8 memory of the second orientation is shifted in a direction opposite to the first orientation. For participants 4 and 6 the reverse happens, and reports of the second orientation are attracted towards the first orientation. For participants 5 and 7 there is no significant bias in either direction. Here we used a Bayesian approach to fit individual subject data (described in more detail in Suchow et al., 2013), which constructs a full probability distribution over the model parameters, and uses a non-informative prior. This allows us to obtain 95% confidence intervals for individual participants (shaded regions) around the Maximum of the posterior distribution (MAP estimate) that serves as a point estimate for each parameter value (solid colors).

Supplementary Figure 2

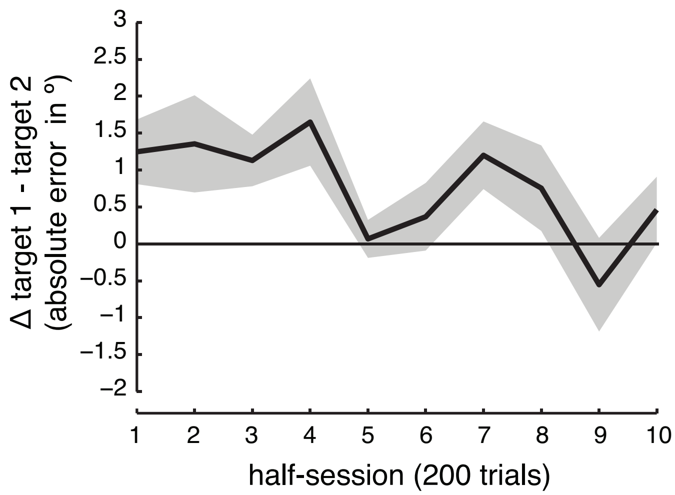

Supplementary Figure 2. Relative performance for the two memory targets over time. As soon as the first grating is shown, there is information available about the possible orientation of the second grating because the second grating can only have a limited number of orientations relative to the first. For example, if the first grating happens to be oriented at 36º (a truly random value between 1° and 180º), the second grating can only be 171°, 21°, 36º, 52º, or 81°. Note that, except for the two authors, participants had no idea about the systematic relationship between the two orientations, although learning (explicit or implicit) might have taken place during the thousand experimental trials participants were exposed to. We examined whether this might explain the overall performance benefit for the second grating compared to the first. For each 200 trials (half of a block) we calculated how much better performance for the second grating was compared to the first (absolute error target 1 – absolute error target 2), where values > 0 indicate an advantage for the second grating. The figure shows that even if participants learned the relationship between the two gratings over time, this did not increase the advantage enjoyed by the second grating. Instead, performance becomes more similar for the two gratings as participants progress through the experiment, with the initial performance benefit for the second target dwindling over time ($F_{(1,7)} = 6.507$; $p = 0.038$). Thus, it is unlikely that memory for the second grating was better due to participants having knowledge about its possible orientation prior to its presentation. The black line plots the group-averaged data with shaded grey regions representing ± 1 SEM.

Supplementary Figure 3

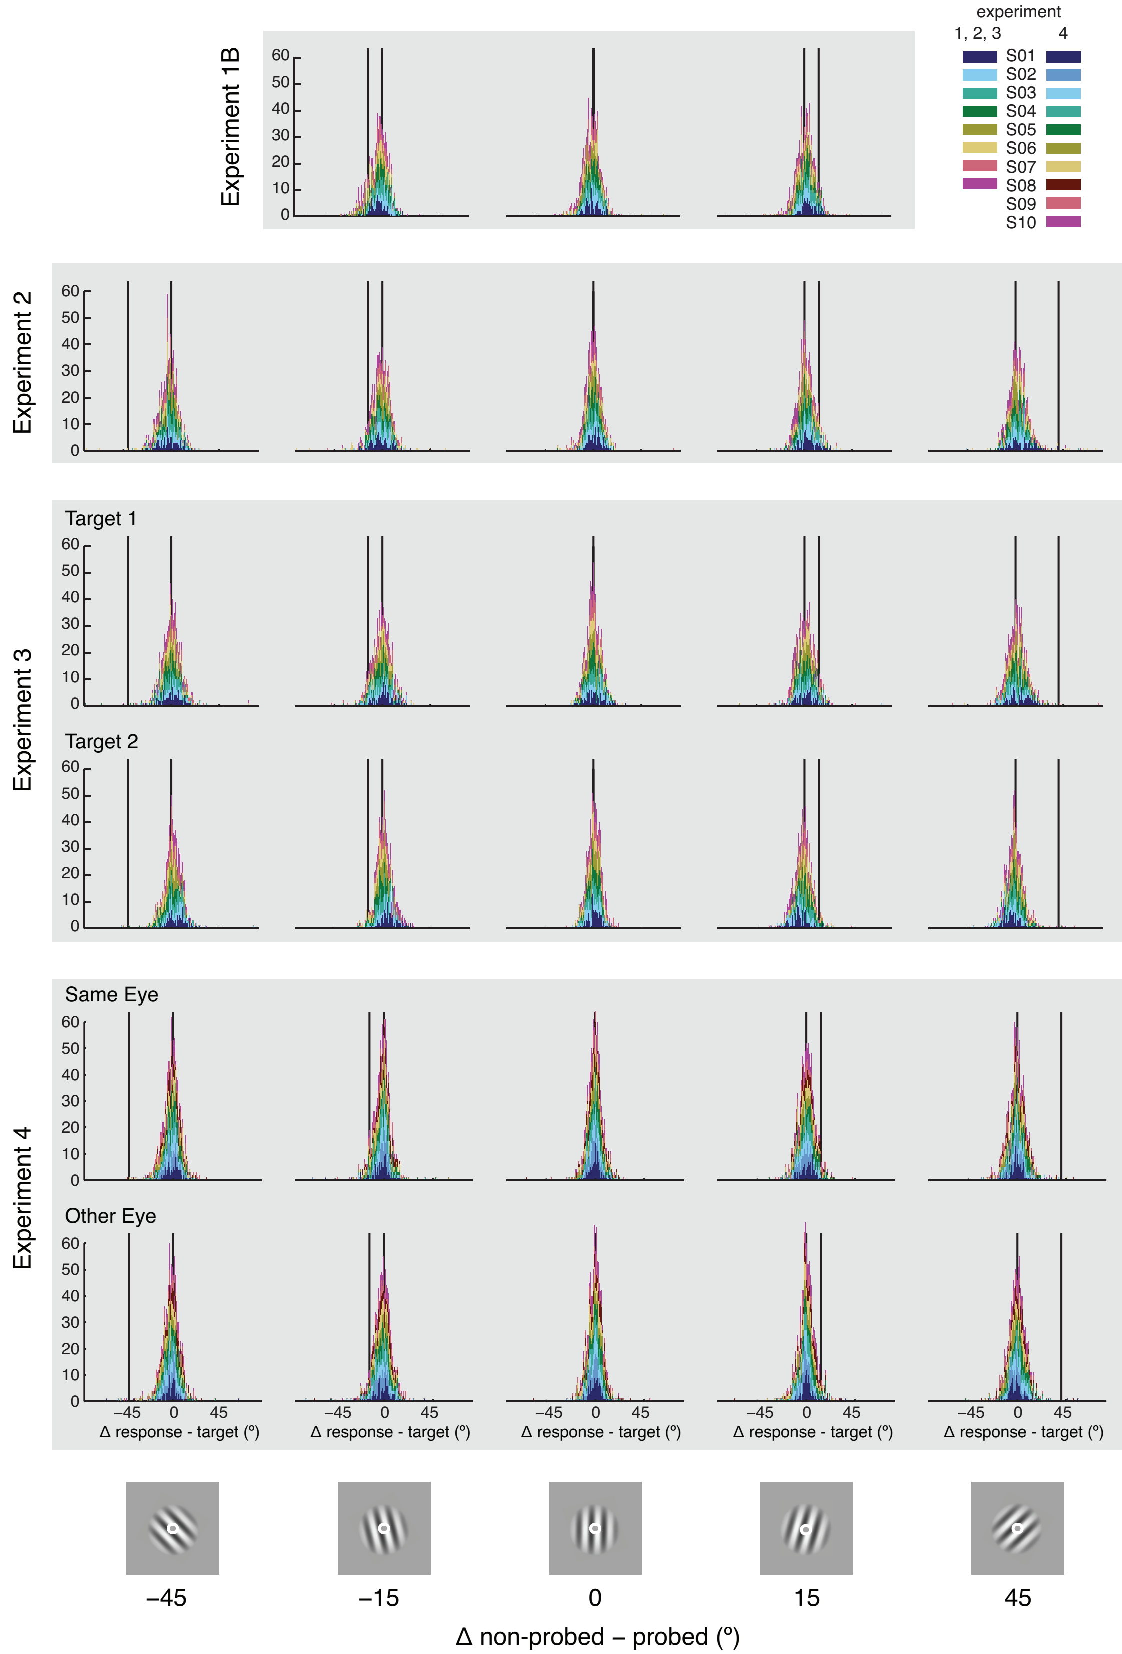

Supplementary Figure 3. Error histograms for all experiments (selected conditions) and all subjects. The memory biases (in the mean response and memory noise) found throughout our study do not originate from participants mistakenly reporting the wrong orientation. Each histogram plots response frequency against the response errors (reported orientation - target orientation). Grey panels represent the different experiments in this study, showing histograms for Experiments 1B, 2, 3 (both targets), and 4 (only visibly presented irrelevant gratings where memory biases were detected). Per grey panel, the histograms display the distribution of responses for conditions where the relative orientation differences were -45º, -15º, 0º, 15º, and 45º. Only these relative orientation differences are shown here to ease comparison between experiments, and for the sake of conciseness. Different colors represent individual participants, while the black vertical lines in each histogram represent the target orientation (at 0º) and the orientation of the second grating presented in that particular condition.

Supplementary Figure 4

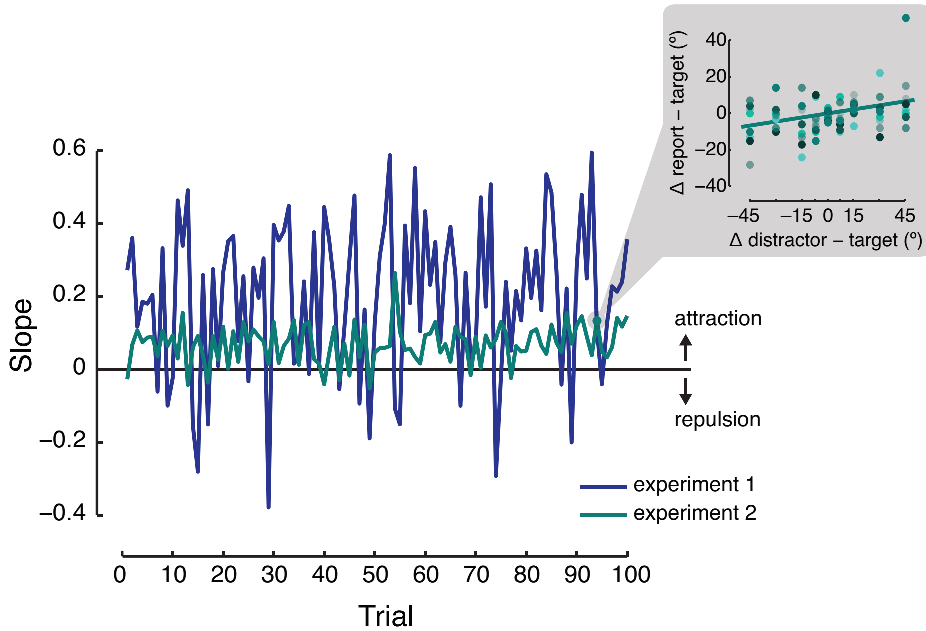

Supplementary Figure 4. We performed a trial-by-trial analysis to observe how shifts of the distribution mean evolved over time. For every n^{th} trial of a given target-distractor condition we calculated the 'bias' on that trial (or simply put: the difference between target and response), and paired this with the corresponding n^{th} trials 'biases' from the full range of target-distractor differences. We plotted these 'biases' for all subjects against target-distractor conditions (see insert, different colors represent different participants), and determined the slope of a line fitted to these points. The slope is taken as a measure of the response shift on the n^{th} trials. The main panel depicts how the response shift evolves over time for Experiments 1 (dark blue) and 2 (green-blue) and demonstrates that memory was shifted (i.e. biased towards the irrelevant distractor) from the beginning onwards, and this shift did not increase over the course of the experimental trials.