


Interhemispheric transfer of working memories

Graphical Abstract



Authors

Scott L. Brincat, Jacob A. Donoghue, Meredith K. Mahnke, Simon Kornblith, Mikael Lundqvist, Earl K. Miller

Correspondence

ekmiller@mit.edu

In Brief

Brincat et al. use bilateral recording to show working memories transferring between the right and left prefrontal cortex. Transferred memories engage different ensembles than feedforward-induced memory traces. Trace transfer is accompanied by directed interhemispheric theta/beta synchrony.

Highlights

- A gaze shift swapped a working memory's location between visual hemifields
- This induced transfer of the working memory trace between the right and left PFC
- Trace transfer was accompanied by directed interhemispheric theta/beta synchrony
- Transferred traces activated different ensembles than feedforward-induced traces

Article

Interhemispheric transfer of working memories

Scott L. Brincat,^{1,3} Jacob A. Donoghue,^{1,3} Meredith K. Mahnke,¹ Simon Kornblith,¹ Mikael Lundqvist,^{1,2} and Earl K. Miller^{1,4,*}

¹The Picower Institute for Learning and Memory and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139 USA

²Department of Psychology, Stockholm University, 106 91 Stockholm, Sweden

³These authors contributed equally

⁴Lead contact

*Correspondence: ekmiller@mit.edu

<https://doi.org/10.1016/j.neuron.2021.01.016>

SUMMARY

Visual working memory (WM) storage is largely independent between the left and right visual hemifields/cerebral hemispheres, yet somehow WM feels seamless. We studied how WM is integrated across hemifields by recording neural activity bilaterally from lateral prefrontal cortex. An instructed saccade during the WM delay shifted the remembered location from one hemifield to the other. Before the shift, spike rates and oscillatory power showed clear signatures of memory laterality. After the shift, the lateralization inverted, consistent with transfer of the memory trace from one hemisphere to the other. Transferred traces initially used different neural ensembles from feedforward-induced ones, but they converged at the end of the delay. Around the time of transfer, synchrony between the two prefrontal hemispheres peaked in theta and beta frequencies, with a directionality consistent with memory trace transfer. This illustrates how dynamics between the two cortical hemispheres can stitch together WM traces across visual hemifields.

INTRODUCTION

Imagine driving on the freeway. A car passes you and holds your attention as you wait to see if it will cut in front of you. Even if you briefly close your eyes or shift your gaze elsewhere, you are able to maintain the car's location in mind and are surprised if it changes unexpectedly. This relies on visual working memory (WM), the ability to maintain images in mind in their absence. Decades of evidence points toward prefrontal cortex (PFC) as a key node in the cortical network underlying WM (D'Esposito and Postle, 2015; Funahashi et al., 1989; Fuster and Alexander, 1971; Miller et al., 1996; Romo et al., 1999; Ungerleider et al., 1998; Voytek and Knight, 2010).

In both humans and monkeys, visual WM seems largely independent between the left and right visual hemifields, which project to the right and left cerebral hemispheres, respectively. WM has a very limited storage capacity (Luck and Vogel, 1997, 2013). But the capacity within one visual hemifield is largely unaffected by the number of objects in the other hemifield (Buschman et al., 2011; Delvenne, 2005; Umemoto et al., 2010). Correspondingly, the neural correlates of WM storage and WM load (how many items are held in memory) primarily reflect items within the contralateral hemifield (Funahashi et al., 1990; Kastner et al., 2007; Kornblith et al., 2015; Luria et al., 2016; Rainer et al., 1998).

Nevertheless, visual cognition seems seamless across the visual field, even when eye movements switch the remembered location of objects between visual hemifields. In such situations,

are memory representations transferred from one cerebral hemisphere to the other, or are they “bound” to the hemisphere where they were initially stored? If transferred, do transferred WMs use the same neural representation as ones induced by feedforward visual inputs, and what are the mechanisms underlying transfer?

We used a novel variant of the delayed nonmatch-to-sample task. A midline-crossing saccade during the memory delay switched the hemifield of a remembered item. We found that WM traces are transferred from one prefrontal hemisphere to the other, that transferred traces recruit neural ensembles distinct from those induced by feedforward inputs, and that this transfer is facilitated by rhythmic coupling between the cerebral hemispheres.

RESULTS

Monkeys were trained to perform a modified version of a non-match-to-sample visual WM task (Figure 1A). They fixated on a point on the left or right (50% of trials randomly) of a computer screen. An object briefly appeared as a sample in the center of the screen, thus in the right or left visual hemifield, respectively (Figure 1B, left). The sample could be one of two different objects, at one of two different locations slightly above or below the center of the screen. The monkeys were required to

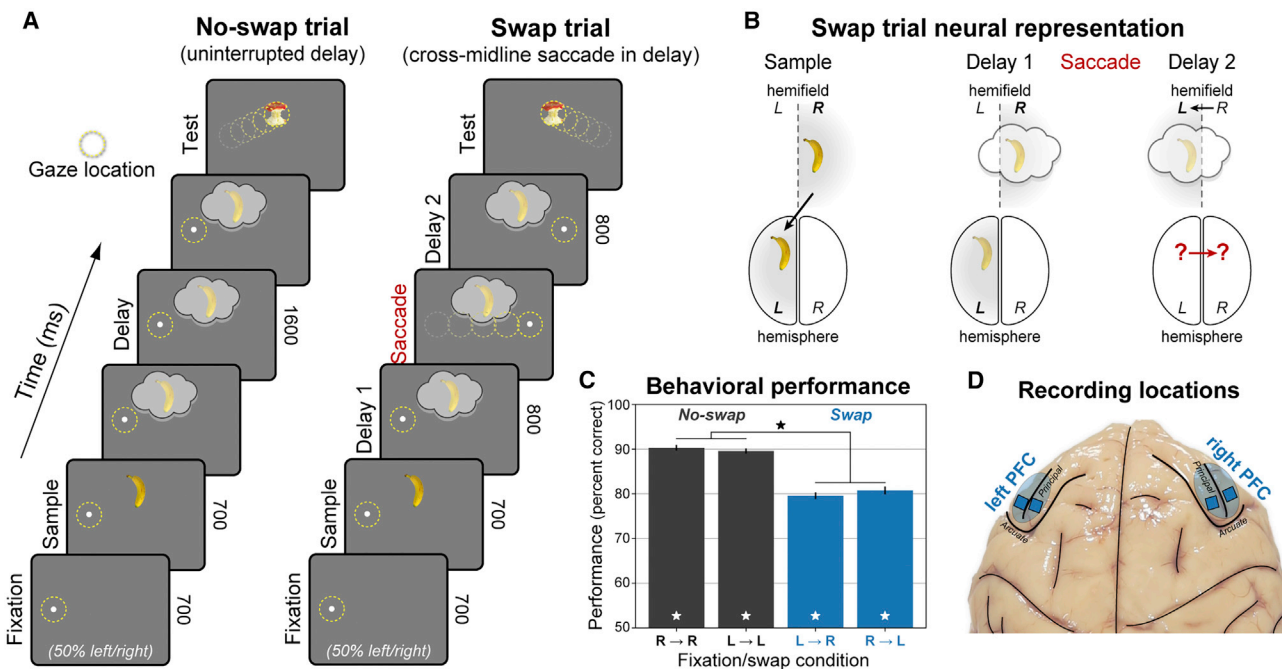


Figure 1. Behavioral and electrophysiological methods

(A) Hemifield-swap working memory (WM) task. Subjects fixated to the left or right while a sample object was presented in the center, placing it in the right or left visual hemifield, respectively. Samples could be one of two objects, presented in one of two locations (above or below center). Both the object and its location needed to be remembered over a blank 1.6 s delay. After the delay, a series of two test objects was displayed, and subjects responded to the one that did not match the sample in object identity or upper/lower location (response to first object shown for brevity). In “no-swap” trials (left), the WM delay was uninterrupted. In “swap” trials (right), subjects were instructed to saccade to the opposite side mid-delay, switching the visual hemifield of the remembered location relative to gaze.

(B) Neural representation of WM in swap trials. The WM trace is initially encoded in the prefrontal hemisphere contralateral to the sample. We tested whether the change in gaze on swap trials caused the WM traces to transfer to the other hemisphere.

(C) Mean performance (\pm SEM across 56 sessions) for each swap condition and visual hemifield. Monkeys performed the task well (white stars, significant versus chance), with a small but significant decrease (black star) on swap trials.

(D) Electrophysiological signals were recorded bilaterally from 256 electrodes in lateral prefrontal cortex (PFC).

remember both the object identity and location over a blank delay and then compare it with a test object. If it did not match the sample in either identity or upper/lower location, they were trained to saccade directly to it. Otherwise, they held fixation until a second, always non-matching test object appeared.

A random 50% of trials had an uninterrupted delay, with no change in the remembered sample location relative to gaze (“no-swap” trials; Figure 1A, left). On the other trials, halfway through the delay the fixation point jumped across the midline to the opposite side, instructing an immediate saccade and refixation on it for the remainder of the delay. This shifted the remembered sample’s retinotopic location to the opposite visual hemifield (“swap” trials; Figure 1A, right). Performance was good for all conditions ($p \leq 1 \times 10^{-4}$ for all, randomized sign test across 56 sessions), albeit somewhat worse on swap trials ($p \leq 1 \times 10^{-4}$, permutation paired t test; Figure 1C). There were no significant differences in performance when the sample ($p = 0.71$) or test object ($p = 0.10$) appeared in the left versus right hemifield, so all results were pooled across them. We recorded multi-unit activity (MUA) and local field potentials (LFPs) from 256 electrodes in four chronic arrays implanted bilaterally in both hemi-

spheres of lateral PFC during task performance (Figure 1D). All presented statistics were based on nonparametric randomization tests across sessions, corrected for multiple comparisons across frequencies and time points (Benjamini and Yekutieli, 2001; see STAR methods for details).

Laterality of WM-related activity

We first examined no-swap trials to establish a reliable signature of the laterality of the WM trace. Data from each prefrontal hemisphere and sample object hemifield were analyzed separately, and then results were pooled on the basis of whether the sample was contralateral or ipsilateral to the recorded hemisphere.

Spiking was stronger and more informative for contralateral than ipsilateral objects. Average MUA was significantly higher for contralateral than ipsilateral samples throughout the trial (Figure 2A, stars; $p < 0.01$, corrected, permutation paired t test). In fact, for ipsilateral samples, MUA was above baseline on average only during the response to the sample presentation and a brief “ramp-up” at the end of the delay (Figure 2A, dots; $p < 0.01$, corrected, randomized sign test). Ipsilateral samples had only a weak effect on spiking during the delay. Unlike

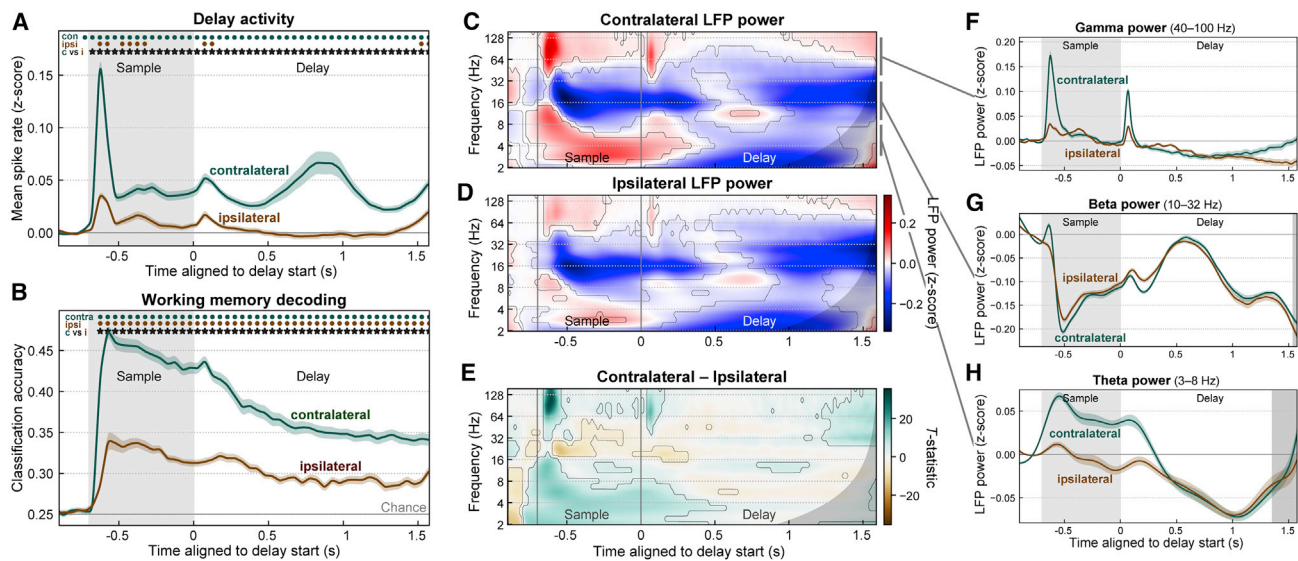


Figure 2. Contralateral bias in prefrontal cortex

(A) Population mean spike rates (multi-unit activity Z-scored to baseline, \pm SEM across 56 sessions) for sample objects contralateral (green) and ipsilateral (brown) to the recorded prefrontal hemisphere (pooled across left and right). Activity for contralateral samples was greater than baseline (green dots) and greater than activity for ipsilateral samples (stars). See also [Figure S1](#).
 (B) Mean (\pm SEM) accuracy for decoding the item held in WM (object identity and upper/lower location) from prefrontal population spike rates, for samples in the contralateral (green) and ipsilateral (brown) visual hemifield. Contralateral decoding accuracy was greater than ipsilateral (stars).
 (C and D) Mean time-frequency LFP power (Z-scored to baseline) for contralateral (C) and ipsilateral (D) samples. Contours indicate significant change from baseline. Gray regions to right of (C)–(H) indicate time points with possible temporal smearing of test-period effects. Gamma (\sim 40–100 Hz) and theta (\sim 3–8 Hz) power increased from baseline (red), while beta (\sim 10–32 Hz) power was suppressed from baseline (blue).
 (E) Contrast (paired-observation t-statistic map) between contralateral and ipsilateral power. Contours indicate significant difference.
 (F–H) Summary of LFP power for contralateral (green) and ipsilateral (brown) sample objects, pooled within frequency bands: gamma (F), beta (G), and theta (H). Because of off-axis structure in the time–frequency data, these one-dimensional (1D) plots cannot fully capture it and are thus included only to aid visualization. All modulations from baseline were stronger for contralateral samples, but only gamma showed a difference during the delay period. See also [Figure S2](#).

contralateral samples, which increased spiking, ipsilateral samples elicited a balance between increased and decreased spiking relative to baseline ([Figure S1](#)).

We computed the accuracy of decoding the contents of WM as a measure of information conveyed in the neural population about WMs. We used a linear discriminant classifier to decode object identity and upper/lower location from the pattern of spiking activity across all MUAs within each hemisphere at each time point. Cross-validated decoding accuracy was significantly above chance for both contralateral and ipsilateral sample objects ([Figure 2B](#), dots; $p < 0.01$, sign test), but it was significantly higher for contralateral ([Figure 2B](#), stars; $p < 0.01$, paired t test). Thus, prefrontal sensory and memory-related spiking showed a clear contralateral bias in strength and information content, as previously reported (see [Discussion](#)).

LFP power also exhibited a contralateral bias, especially for higher frequencies. Gamma power (\sim 40–100 Hz) was significantly elevated relative to baseline during sample object presentation and a “ramp-up” at the end of the delay for both contralateral ([Figure 2C](#)) and ipsilateral ([Figure 2D](#)) samples ($p < 0.01$, sign test; summarized in [Figure 2F](#)). Gamma power induced by contralateral sample objects was significantly higher than for ipsilateral objects during the sample onset and offset transients and the pre-test “ramp-up” ([Figure 2E](#); $p < 0.01$, paired t test).

Theta power (\sim 3–8 Hz) also showed a contralateral bias during the sample response but not the memory delay ([Figures 2E](#) and [2H](#)). In contrast, beta power (\sim 10–32 Hz) showed effects in the opposite direction overall: significant decreases in power from baseline during the sample and late delay periods for both contralateral and ipsilateral sample objects ([Figure 2G](#)). Like gamma and theta enhancement, beta suppression was significantly stronger for contralateral than ipsilateral samples ([Figure 2E](#)) but, like theta power, only during the sample object response. Beta oscillatory burst rates exhibited a stronger and more sustained laterality than beta power, but analysis of LFP bursting produced otherwise similar results ([Figures S2A](#) and [S2B](#)). These results indicate that prefrontal LFP power, like spiking activity, exhibits a clear contralateral bias.

Transfer of WMs between cerebral hemispheres

We then leveraged the neural signature of WM laterality to examine what happened when the remembered location was shifted by a saccade to the opposite visual hemifield relative to the center of gaze. We propose two alternative hypotheses. On the one hand, a WM trace might be bound to the initial representation induced by visual inputs and might simply remain in the hemisphere where it was originally encoded ([Figure 3A](#)). This *stable trace model* predicts no change in neural signatures after the mid-delay saccade

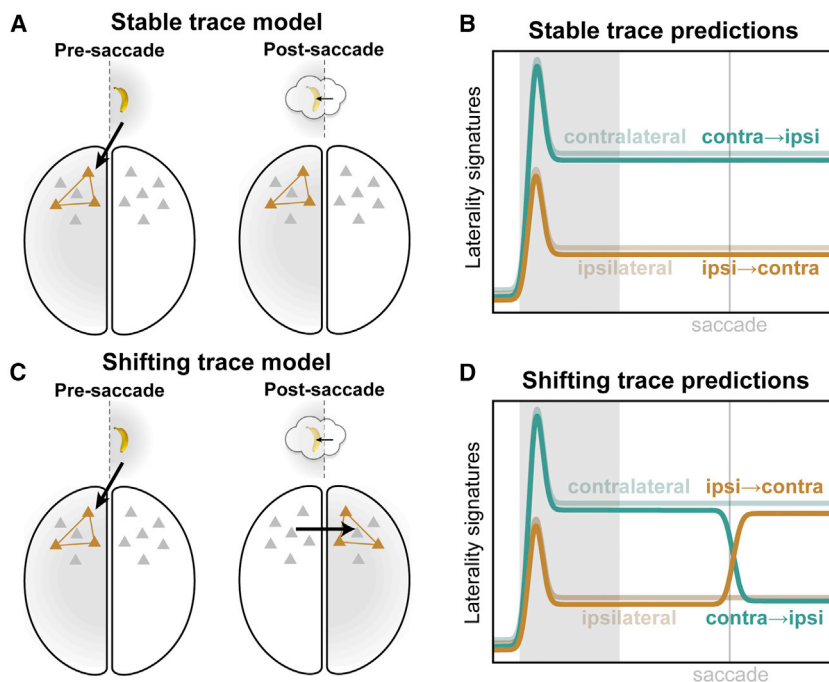


Figure 3. Competing models of swap effects

(A) The stable trace model posits that once a working memory is encoded in a given cortical hemisphere (left), it will remain there (right), despite the remembered location shifting from one hemifield to the other (inset).

(B) This model predicts that neural signatures of memory trace laterality will be unaltered by the mid-delay saccade in our task.

(C) The shifting trace model assumes that when the hemifield of the remembered location is swapped, the memory trace will be transferred from one cortical hemisphere to the other.

(D) This model predicts a post-saccadic inversion of the neural signatures of laterality: shifting the remembered location into the contralateral hemifield (orange) should come to approximate the constant contralateral location (desaturated green), while shifting it ipsilateral (green) should come to look like constant ipsilateral trials (desaturated brown).

(Figure 3B). Alternatively, when the hemifield switches, the neural trace itself might also move from the hemisphere it was originally encoded in to the opposite hemisphere, now contralateral to the remembered location (Figure 3C). This *shifting trace model* predicts neural signatures of laterality in the swap trials will invert after the midline-crossing saccade (Figure 3D).

We found evidence for the shifting trace model, an inversion of neural laterality signatures after the midline-crossing saccade. This was apparent in average MUA rate (Figure 4A) and in information carried by MUA (Figure 4B). As in the no-swap trials (Figures 2A and 2B), prefrontal MUA starts the delay with a bias toward the contralateral hemifield (Figures 4A and 4B, left; H symbols: $p < 0.01$, corrected, sample hemifield main effect in hemifield \times shift condition permutation two-way ANOVA). But after the saccade (Figures 4A and 4B, right), both MUA and decoding accuracy increased for remembered locations that shifted from the ipsilateral to the contralateral hemifield (orange) relative to that for ipsilateral locations on no-swap trials (desaturated brown; brown stars: $p < 0.01$, paired t test). By contrast, on swap trials when the saccade shifted the remembered location from the contralateral to ipsilateral hemifield (green), there was a decrease compared with the contralateral location on no-swap trials (desaturated green; green stars: $p < 0.01$). For average MUA (Figure 4A), there was a complete inversion. The spike rates after a saccade that shifted the remembered location into a given hemifield almost exactly matched those for a static memory in the same hemifield. Around the time of the saccade, there was also increased spiking for both swap conditions, relative to the no-swap trials (S symbols: $p < 0.01$, shift condition main effect). Nevertheless, later in the delay, the predicted inversion effect was dominant (X symbols: $p < 0.01$, interaction effect). For decoding accuracy (Figure 4B), ipsilateral-shifting trials (green) were near the value of constant ipsilateral trials

(desaturated brown). Contralateral-shifting trials (orange) exhibited a bump of increased accuracy after the saccade (brown stars: $p < 0.01$, paired t test) but subsequently declined and never attained the level of constant contralateral trials (desaturated green). This imperfect transfer of information may explain why behavioral performance was significantly decreased in swap trials.

Similar effects were seen in LFP power (Figure 5). During and just after the saccade, gamma (Figures 5A, 5B, and 5D) and theta (Figures 5A, 5B, and 5F) power were stronger overall for swap than no-swap trials. Later in the delay, however, gamma power inverted and became stronger for remembered locations moving from the ipsilateral to the preferred contralateral hemifield (Figures 5C and 5D; $p < 0.01$, hemifield \times swap interaction effect) as predicted by the shifting trace model. Theta power showed a similar inversion at the very end of the delay, though this is likely due to temporal smearing of test period effects (Figures 5C and 5F; gray regions indicate extent of possible test effects).

Beta power (Figures 5A–5C and 5E) exhibited complex multi-phasic dynamics on swap trials (Figure 5F). It was suppressed initially after the saccade but later became enhanced, relative to power on no-swap trials (Figure 5F). On top of these overall dynamics, however, beta power on the swap conditions inverted: it became stronger for remembered locations shifting into the ipsilateral hemifield than for those shifting into the contralateral hemifield (Figures 5C and 5F). Thus, as for spiking, prefrontal LFP power signatures of WM laterality also exhibited the inversion predicted by the shifting trace model (Figure 3D). These results support the hypothesis that the memory trace is transferred from one cortical hemisphere to the other.

Interhemispheric transfer activates novel neural ensembles

A model consistent with the data thus far is that WMs transferred between hemispheres recruit the same neural ensembles as memory traces activated by feedforward visual inputs into the

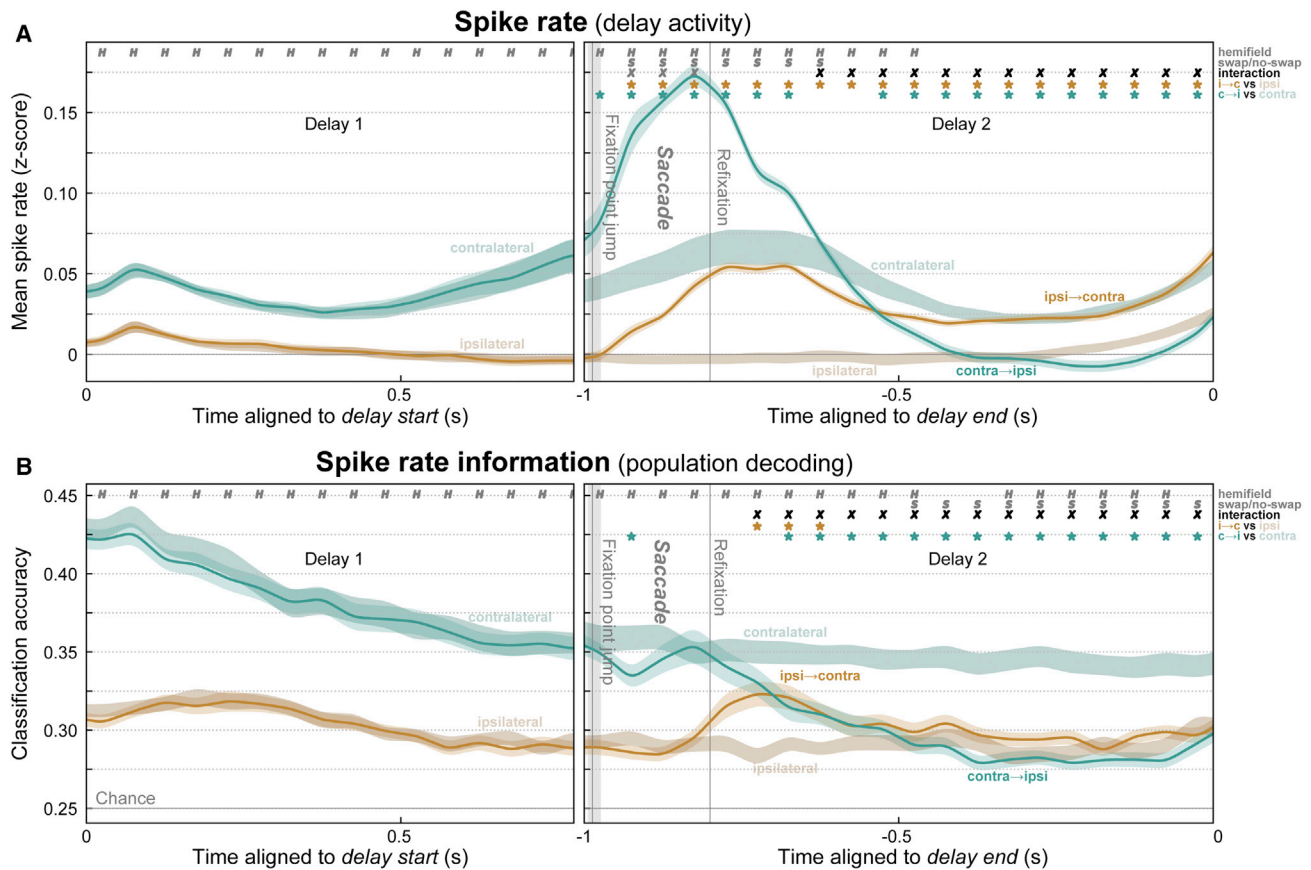


Figure 4. Evidence for interhemispheric transfer of working memory traces in prefrontal spiking activity

(A) Mean (\pm SEM) multi-unit spike rates for all trials where the remembered location was constant in the contralateral (desaturated green) or ipsilateral (desaturated brown) hemifield or where it swapped from ipsilateral to contralateral (orange) or from contralateral to ipsilateral (green). Before the mid-delay saccade, there was only a significant effect of the sample hemifield (H symbols). Around the saccade, activity was greater overall for swap than no-swap trials (S symbols). Later, the swap trials inverted and approximated activity in the corresponding no-swap trials (X symbols, significant hemifield \times swap condition interaction). Stars indicate significant difference of swap conditions from their respective no-swap baseline.

(B) Mean (\pm SEM) accuracy for decoding the item held in WM from spike rates. As predicted, post-saccade accuracy decreased in contralateral-to-ipsilateral trials. Decoding on ipsilateral-to-contralateral trials also significantly improved above baseline, but only transiently, and it never reached the level of constant contralateral trials.

same cortical hemisphere. Under this *generic ensemble model* (Figure 6A), when a given object—say, a banana in the upper location—is held in WM within a cortical hemisphere, it uses the same neural ensemble (i.e., the same pattern of spiking across our electrode arrays) whether it arrived there via feedforward inputs from visual cortex (Figure 6A, left) or via interhemispheric inputs from the contralateral hemisphere (Figure 6A, right). We could test this because the saccade on swap trials brought the remembered sample location to the same retinotopic coordinates where it appeared on the no-swap trials. An alternative model is motivated by the fact that unique prefrontal ensembles are activated by different combinations of input features and task contexts (Rigotti et al., 2013). Perhaps the same information arriving via different circuits—feedforward (Figure 6C, left) versus interhemispheric (Figure 6C, right)—also activates different ensembles. We call this alternative the *novel ensemble model*.

We used a cross-classification method to adjudicate between the models. We trained a classifier to decode the object identity

and its upper/lower location from population spiking at each time point on contralateral no-swap trials. We tested whether these classifiers could predict the same information on ipsilateral-to-contralateral swap trials, which brought that same object to the same location as the contralateral no-swap trials. Note that training and testing were both performed on the same cortical hemisphere—separately for each hemisphere, then results were pooled across them—meaning the cross-classification is across task conditions, not cortical hemispheres. Thus, we tested whether the same information was reflected in the same neural pattern in a given hemisphere, regardless of how it arrived there. If both conditions activate the same ensembles, as assumed by the generic trace model (Figure 6A), then this cross-classification (Figure 6B, orange) should result in high decoding accuracy, similar to that obtained from both training and testing on constant contralateral trials (Figure 6B, desaturated green). If these conditions activate different neural ensembles, as suggested by the novel ensemble model (Figure 6C), then

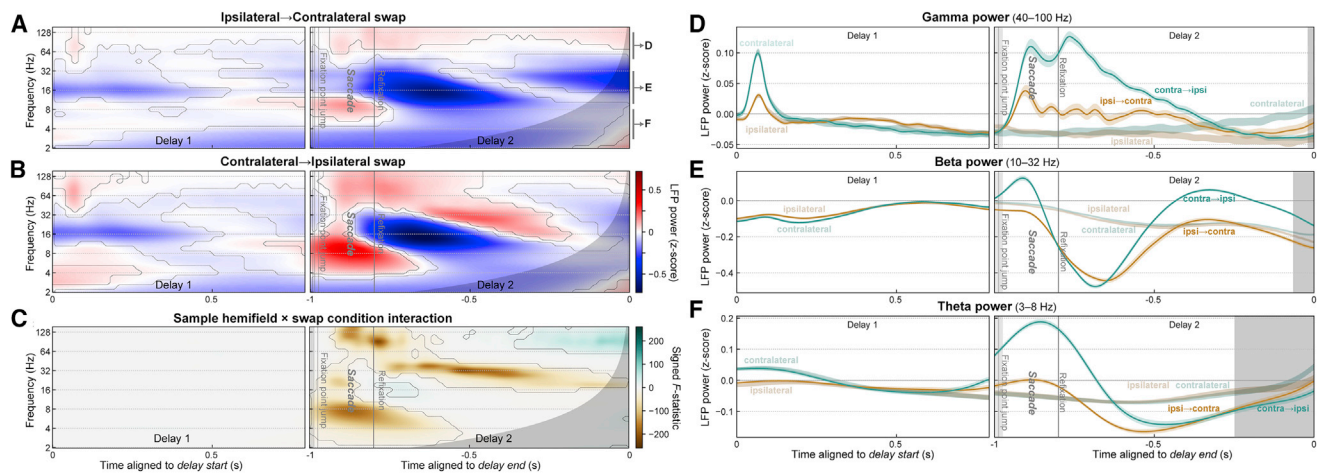


Figure 5. Evidence for interhemispheric transfer of working memory traces in prefrontal LFP power

(A and B) Mean time-frequency LFP power for trials where the remembered location shifted from the ipsilateral to contralateral (A) or from the contralateral to ipsilateral (B) hemifield. Contours indicate significant difference from pre-sample baseline. Gray regions at right of all panels indicate time points with possible temporal smearing of test-period effects.

(C) Swap inversion effect. F-statistic map for sample hemifield × swap condition interaction, signed to indicate if power was greater when remembered location ends up in contralateral (green) or ipsilateral (brown) hemifield. Contours indicate significant interaction.

(D–F) Summary of LFP power for ipsilateral-to-contralateral (green) and contralateral-to-ipsilateral (orange) trials, pooled within frequency bands labeled in (A): gamma (D), beta (E), and theta (F). Around the time of the saccade, LFP power in all bands showed strong effects of the swap condition (saccade → versus no saccade). Later in the post-saccade delay, signatures in all bands inverted, as predicted by the shifting trace model.

See also [Figure S2](#).

cross-classification accuracy ([Figure 6D](#), orange) should be poor, similar to that obtained from cross-classification testing on constant ipsilateral trials ([Figure 6D](#), desaturated brown). We find evidence for both models at different time points during the post-saccade delay.

For reference, we replot (from [Figure 4B](#)) classification accuracy when both training and cross-validated testing were performed on constant contralateral trials ([Figure 6E](#), desaturated green). As a control, we computed cross-classification accuracy when the same no-swap contralateral trained classifiers were tested on no-swap ipsilateral trials ([Figure 6E](#), desaturated brown). This reflects baseline cross-classification generalization due solely to any bilaterality in the prefrontal neural code. This control analysis resulted in poor accuracy, further evidence for independence between the contralateral and ipsilateral hemifields. For most of the post-saccade delay, cross-classification of the contralateral-shift trials ([Figure 6E](#), orange) was also poor and not significantly different from the control ($p \geq 0.01$, paired t test). Here, classifier training and testing were performed at the same time points relative to the end of the delay. Similar results were obtained with training and testing at all possible relative times ([Figure S3](#)), indicating results are not dependent on the specific timing scheme used.

These results mostly support the novel ensemble model. For much of the post-saccade delay, decoders generalized poorly from the constant contralateral trials to those in which it was transferred from the opposite hemisphere ([Figure 6E](#)). The initial bump in decoding accuracy seen when the saccade shifted the remembered location from ipsilateral to contralateral ([Figure 4B](#)) was not detected using the classifiers from no-swap trials. Thus, different neural ensembles were activated by the same sensory

information in WM, depending on whether it arrived via ipsilateral visual inputs or via the contralateral hemisphere. However, near the end of the delay, the cross-classification of contralateral-shift trials increased relative to the control ([Figure 6E](#), $p < 0.01$, paired t test). This suggests that, in anticipation of using the WM, interhemispherically transferred memory traces converged somewhat toward an ensemble representation similar to feedforward-induced traces.

Interhemispheric synchrony during memory transfer

Our results thus far suggest transfer of WM traces between prefrontal hemispheres. If so, we should expect to see evidence of communication between hemispheres around the time of putative transfer. We would further expect signals to flow causally from the hemisphere contralateral to the initial sample hemifield (the “sender”) toward the hemisphere contralateral to the post-saccade hemifield (the “receiver”). Evidence suggests that phase synchrony between cortical areas helps regulate the flow of information ([Fries, 2015](#)). Thus, we measured oscillatory synchrony between LFPs in the two prefrontal hemispheres using pairwise phase consistency (PPC), an unbiased measure of phase synchrony ([Figure 7](#)).

Around the saccade on swap trials—when WM trace transfer putatively occurs—interhemispheric theta (~4–10 Hz) and high-beta (~18–40 Hz) synchrony both exhibited a transient peak ([Figure 7A](#); $p < 0.01$, paired t test versus pre-sample baseline). No such peaks were observed at analogous time points on no-swap trials ([Figure 7B](#)). These differences were confirmed by examining the contrast between swap and no-swap trials ([Figure 7C](#); $p < 0.01$, paired t test). In contrast, during this period there was a suppression of interhemispheric synchrony relative

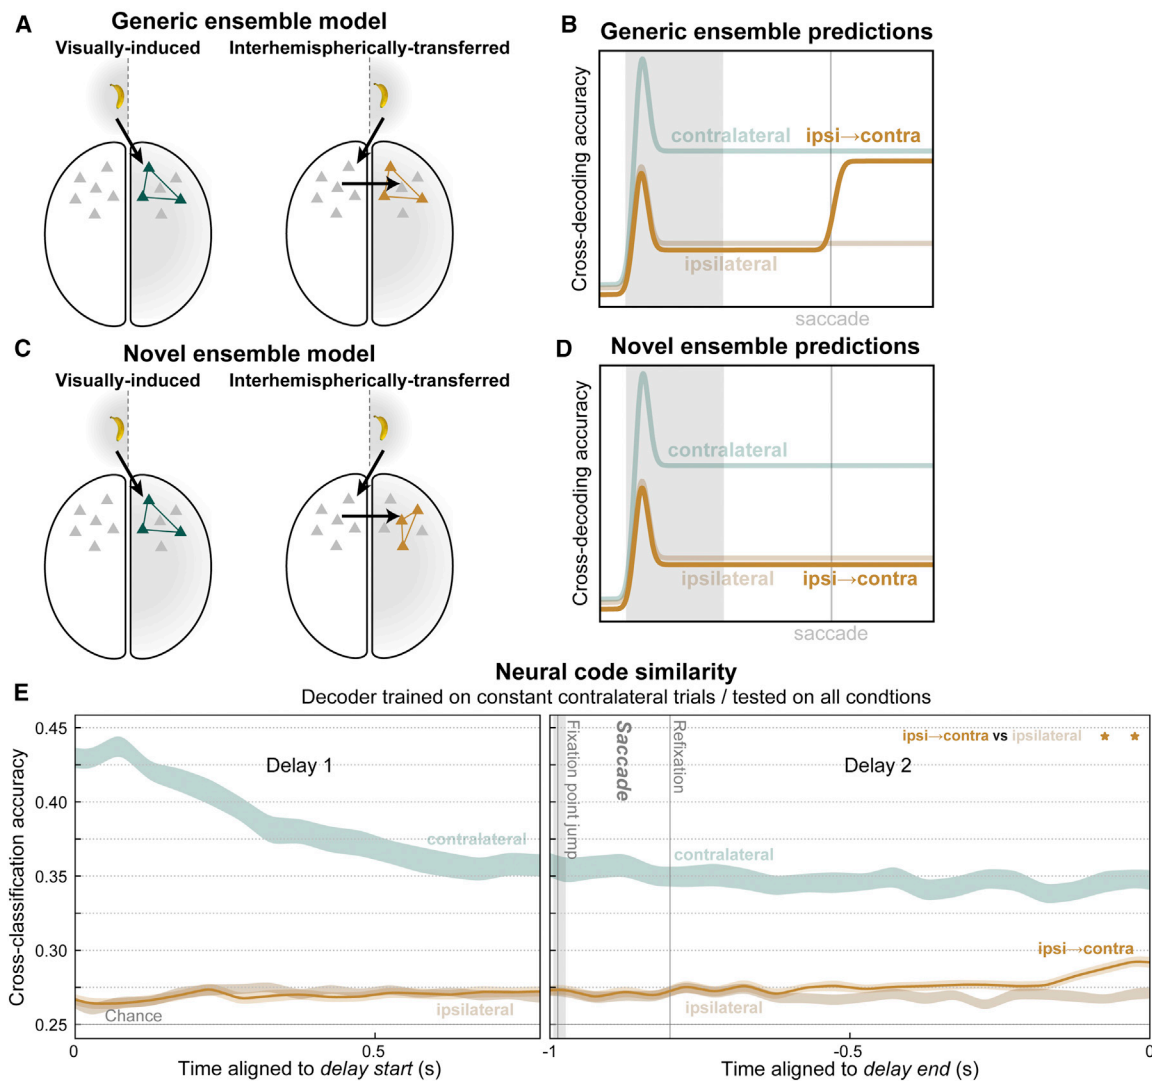


Figure 6. Transferred working memory traces used novel ensembles but converged toward visually induced ensembles at delay end

(A) The generic ensemble model assumes that a given memory trace will activate the same neural ensemble (colored neurons) whether it arrives in prefrontal cortex via feedforward inputs from visual cortex (left) or via interhemispheric inputs from the opposite cortical hemisphere (right). (B) It predicts that a classifier trained to decode working memory traces on constant contralateral trials will also be able to decode contralateral-shifting swap trials (orange). (C) The novel ensemble model posits that interhemispheric inputs activate a distinct ensemble (right) from visual inputs (left), even for the same memory trace. (D) It predicts failure of contralateral-trained decoders to generalize to contralateral-shifting trials. (E) For most of the post-saccadic delay, cross-decoding accuracy for ipsilateral-to-contralateral swap trials (orange) did not significantly differ from constant ipsilateral trials, as predicted by the novel ensemble model. Near the end of the delay, a significant difference emerged (stars), indicating contralateral-shifting trials became more similar to constant contralateral trials, as predicted by the generic ensemble model. See also [Figure S3](#).

to baseline within the alpha/low-beta band (~11–17 Hz), which extended into the post-saccade delay period ([Figure 7C](#)). Similar results, but with higher frequency positive effects extending into the gamma band (~40–100 Hz), were observed for synchrony between LFPs at distinct sites within each hemisphere ([Figures 7D–7I](#)). Theta and gamma increases were stronger in the sender hemisphere ([Figures 7D, 7F, and 7H](#)). Alpha/beta suppression from baseline, like beta power, was more similar in strength between hemispheres. These results suggest evidence for inter-

hemispheric signal communication underlying memory trace transfer and that communication occurs via theta and high-beta but not alpha/beta synchrony.

To test whether signals flow from the sender to the receiver hemisphere, we measured spectral Granger causality between LFPs in the two prefrontal hemispheres. Granger causality is a measure of how much variance in one signal can be explained by the recent history of another (here, LFPs at distinct sites), beyond what can be explained by the signal's own dynamics.

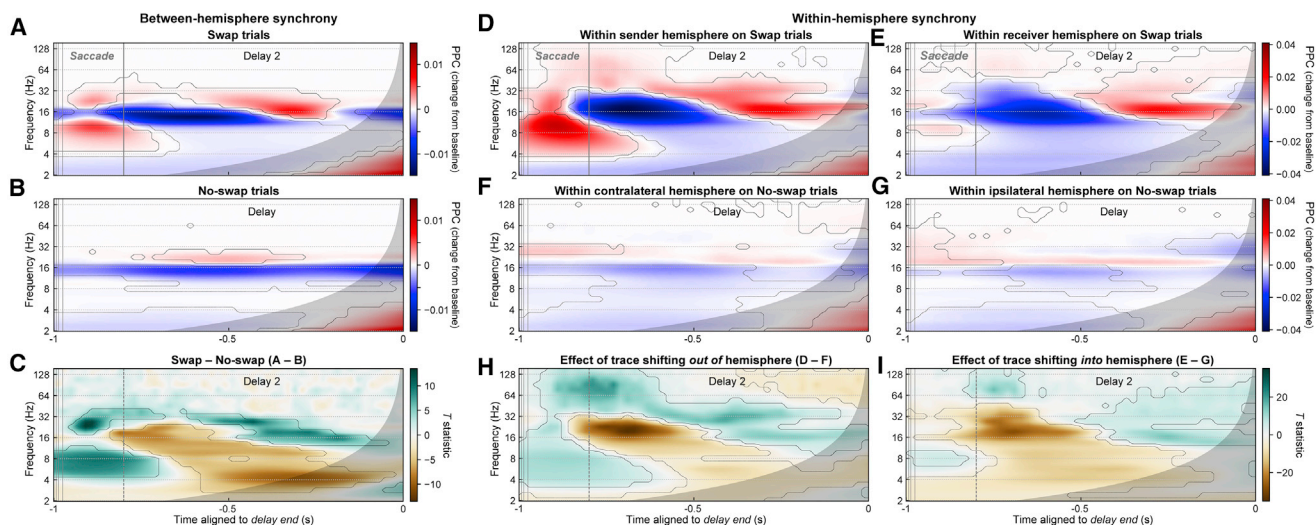


Figure 7. Interhemispheric beta/theta synchrony may mediate working memory trace transfer

(A and B) Mean phase synchrony (pairwise phase consistency [PPC]) between all pairs of LFPs in the two prefrontal hemispheres, for swap (A) and no-swap (B) trials, expressed as the change in PPC from the pre-sample fixation-period baseline. Contours indicate significant differences from baseline. Gray regions indicate time points with possible influence of test-period effects.

(C) Contrast (paired t-statistic map) between swap and no-swap conditions. Contours indicate significant between-condition difference. During the time period of putative interhemispheric memory trace transfer (–1 to –0.8 s), there was a significant increase (green) in interhemispheric synchrony in the theta (~4 to 10 Hz) and beta (~18 to 40 Hz) bands and a decrease (brown) in the alpha/low-beta band (~11 to 17 Hz).

(D and E) PPC between LFP pairs within the sender (D) and receiver (E) hemisphere on swap trials.

(F and G) PPC between LFP pairs within the contralateral (F) and ipsilateral (G) hemispheres on no-swap trials.

(H and I) t-Statistic maps for contrasts between swap and no-swap results for each hemisphere. These contrasts represent the effect of a WM trace shifting out of (H) and into (I) a hemisphere.

Spectral Granger causality expresses these causal interactions in the frequency domain (Dhamala et al., 2008). As predicted, causality was greater in the sender-to-receiver direction (Figure 8A, green) than in the opposite, receiver-to-sender direction (orange). Although this difference was significant for all frequencies from approximately 10–40 Hz (Figure 8B, gray stars; $p < 0.01$, paired t test), it was greatest at the same high-beta frequencies that synchrony peaked at (~20–40 Hz; Figure 8B, gray curve). This asymmetric directionality was not observed in no-swap trials between sites contralateral and ipsilateral to the sample location (Figures 8C and 8D). Between-region differences in signal-to-noise ratio (SNR) can result in false-positive Granger causality. As our results were pooled across left-to-right and right-to-left hemispheres, any between-hemisphere SNR differences should have averaged out. To further rule out this possibility, we recomputed Granger causality on time-reversed data, which preserves SNR differences but should reverse the direction of true causal signal flow (Bastos and Schoffelen, 2016; Haufe et al., 2013). This is exactly what we observed (Figure S4). These results indicate that during the time period of putative WM trace transfer, signals flow in the predicted direction, from the sender to the receiver hemisphere.

DISCUSSION

Our results suggest that WM traces can be transferred from one prefrontal hemisphere to the other. Previous studies have only provided indirect evidence of this. In WM tasks, when sample

and test stimuli appear in opposite visual hemifields, masking during the delay is more effective in the expected test hemifield (Zaksas et al., 2001). Associative cues presented to one visual hemifield can elicit associated representations in the opposite-hemisphere visual cortex (Tomita et al., 1999). Here, we provided direct neurophysiological evidence of WM transfer between hemispheres.

A possible anatomical substrate is direct connections between hemispheres via the corpus callosum (Barbas and Pandya, 1984). Cortical feedforward processing has been shown to be mediated by gamma and theta oscillations while feedback processing is mediated by alpha/beta oscillations (Bastos et al., 2015; Buschman and Miller, 2007; van Kerkoerle et al., 2014). Our results showed increases in interhemispheric theta and high-beta synchrony, and decreases in alpha/low-beta, during the presumed time of memory trace transfer. One possibility is that interhemispheric communication occurs both in frequency bands used by feedforward (theta) and feedback (beta) processing. Perhaps this allows flexible interaction of interhemispheric signals with either feedforward or feedback signals. Another possibility is that higher frequency cortical communication—typically in the traditional gamma band—is shifted slightly downward in frequency (to 40 Hz and lower) to maintain reliable transmission despite the time delays imposed by long callosal axons. Of course, there are other possible routes for interhemispheric transfer other than direct prefrontal callosal connections. Signals might be transferred at some higher level region and then conveyed via feedback to PFC, or conceivably transfer could

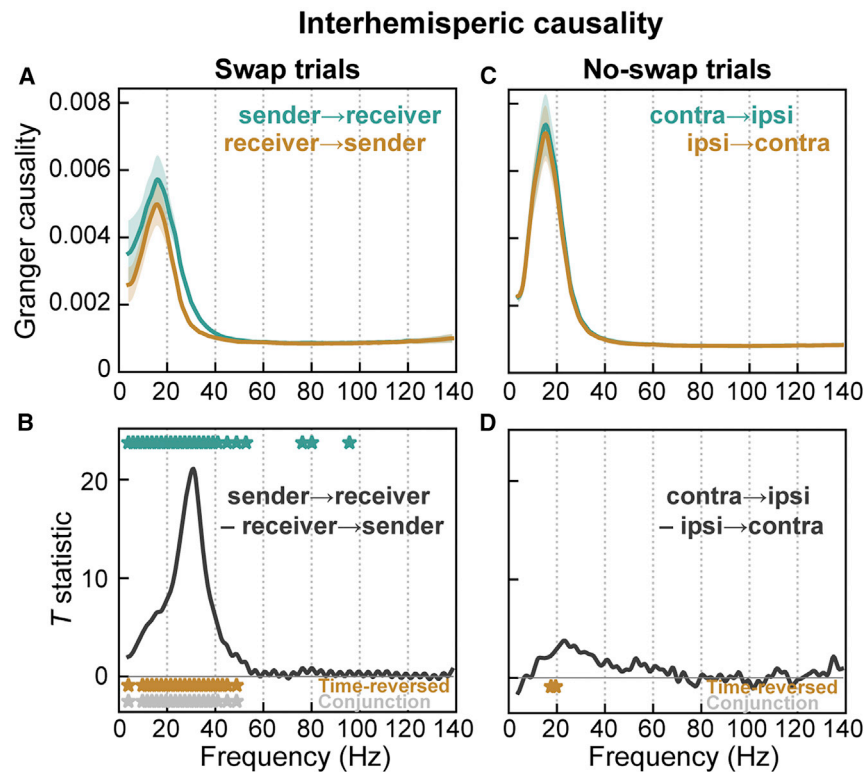


Figure 8. Granger causality indicates flow between prefrontal hemispheres in same direction as putative memory trace transfer

(A and B) Spectral Granger causality in swap trials during time period of putative memory trace transfer from prefrontal hemisphere contralateral to initial sample location (“sender”) to hemisphere contralateral to post-saccade location (“receiver”). (A) Mean (\pm SEM) causality in the sender-to-receiver direction (green) and in the receiver-to-sender direction (orange). (B) t-Statistic for contrast between causal directions. Causality was significantly greater in the sender-to-receiver direction across all frequencies ~10–40 Hz (green stars). This directional asymmetry reversed, as expected, for time-reversed data (orange stars; full results in Figure S4). Gray stars indicate frequencies significant for both forward and time-reversed data.

(C and D) No asymmetry in interhemispheric causality was observed during analogous time points (relative to delay end) in no-swap trials between contralateral and ipsilateral hemispheres.

120 ms after the rise in the receiver. Interhemispheric transfer in this task thus involves a “soft handoff,” in which information overlaps in time in both hemispheres before being cleared out of the sending hemisphere. This is similar to what has

happen at lower levels and feed forward to PFC. Ultimately, resolving this issue will likely require causal manipulation of connectivity.

Gamma and alpha/beta oscillations have characteristic WM-related temporal dynamics: gamma is elevated and beta suppressed during encoding (sample period) and readout (end-of-delay “ramp-up”) of WM, while the opposite is true during WM maintenance (delay period; Lundqvist et al., 2016, 2018). We replicated these anti-correlated gamma/beta dynamics both when a WM is statically maintained, and when it is transferred between hemispheres. We found that theta dynamics also broadly correlate with gamma and anti-correlate with alpha/beta. We also found similar dynamics in these frequency bands in the synchrony between distinct sites both within a prefrontal hemisphere and between hemispheres, with the exception that synchronization between hemispheres at traditional gamma frequencies was absent.

To better understand the temporal dynamics of interhemispheric WM transfer, we compared the latencies of all studied neural signals (Figure S5). The earliest signals after the saccade instruction included beta synchrony within the sender hemisphere and between hemispheres, consistent with the idea that it may be involved in establishing a communication channel for interhemispheric transfer. Beta signals then exhibited a bilateral decrease around the time the receiver hemisphere begins to convey information about the WM trace, consistent with the idea that suppressing beta may disinhibit cortex and allow WMs to be expressed (Lundqvist et al., 2016). Finally, WM information decreased in the sender hemisphere about

been found for tracking moving objects across the midline (Drew et al., 2014).

Another way to think about our results is in terms of reference frames. In a retinotopic (gaze-centered) reference frame, locations are described relative to gaze (Figure S6A, top). Virtually all studied areas of visual cortex code in a retinotopic reference frame (Figure S6B, top; Cohen and Andersen, 2002; Golomb and Kanwisher, 2012). In contrast, a spatiotopic (world-centered) reference frame represents locations in the real world, independent of gaze (Figures S6A and S6B, bottom). In addition to other more established non-retinotopic spatial coding schemes (Chafee et al., 2007; Graziano and Gross, 1998; Olson, 2003), explicit spatiotopic reference frames may exist in higher level cortex (Dean and Platt, 2006; Duhamel et al., 1997).

In our study, whether the eyes were fixated to the left or right, the sample item was always at the same central spatiotopic (real-world) location. Thus, if PFC encoded locations in a spatiotopic reference frame (Figure S6B, bottom), its activity should be relatively invariant to the location of the sample object relative to the eyes (Figure S6C, bottom row). Arguing against this possibility is the differentiation between contralateral and ipsilateral sample locations in PFC activity (Figure 2). This suggests a retinotopic reference frame (Figure S6C, top row). On the other hand, if WM, at the cognitive level, also maintained locations in a retinotopic reference frame, remembered locations would be anchored to the location of gaze and would simply shift around with a saccade (Figure S6A, top). This makes predictions (Figure S6C, left column) identical

to the stable trace model (Figures 2A and 2C) (i.e., that the memory trace remains in the original hemisphere). Instead, we found that neural signatures of laterality invert after the saccade (Figures 4 and 5), ruling out a retinotopic reference frame for WM (Figure S6C, left column). Our results are consistent with the remembered location's being maintained within a spatiotopic reference frame at the cognitive level (Ong et al., 2009) but within a retinotopic reference frame at the neural level (Figure S6C, upper right quadrant).

How can we reconcile this? One possibility is that cognition reflects some putative higher level area with a spatiotopic reference frame that updates the retinotopic representation after each saccade. Arguing against this is the relative paucity of evidence for any explicit spatiotopic representation in the visual cortical hierarchy (Dean and Platt, 2006; Duhamel et al., 1997; Golomb and Kanwisher, 2012). Alternatively, spatiotopic cognition may not rely on an explicit spatiotopic neural representation. It could instead be coded implicitly by a retinotopic representation that is locally updated after a saccade. This would shift a remembered location in the opposite direction of each saccade vector to its new coordinates on the retinotopic map (Pouget and Snyder, 2000). For a midline-crossing saccade, this would entail transfer of the memory trace between the left and right visual hemifields/hemispheres.

In many areas of visual and visuomotor cortex, receptive fields (RFs) show anticipatory shifts to their future, post-saccadic location even before the onset of a saccade (Colby et al., 1995). Other studies suggest that during a saccade, RFs contract toward the location of the saccade target then later expand out to their final post-saccadic location (Chen et al., 2018; Neupane et al., 2016; Zirnsak et al., 2014). This implies a change in the population neural code for location around the time of a saccade because the ensemble responsive to a given location during the RF contraction will be different from those responsive to the same location before and after the saccade. This could explain the lack of cross-decoding between static and shifted memory traces (Figure 6C). These dynamics have yet to be demonstrated in PFC, but their properties in other areas argue against a role. RF contraction effects extend to only ~300 ms after a saccade in area V4 (Neupane et al., 2016) and the frontal eye field (Chen et al., 2018), whereas poor cross-decoding extends to over 500 ms in our results. Contraction effects in V4 occur mainly for neurons with RFs in the same hemifield as the saccade endpoint (Neupane et al., 2016), whereas in the key ipsilateral-to-contralateral condition in our results, RFs would be in the opposite hemifield to the saccade.

An alternative account is motivated by findings of “nonlinear mixed selectivity” in PFC (Rigotti et al., 2013). It suggests that PFC is best understood as a random network, in which unique ensembles are activated by different combinations of input features and task contexts (Bouchacourt and Buschman, 2019). This predicts distinct ensembles are activated depending on the route by which information arrives, feedforward versus interhemispheric. Regardless of the specific mechanism, our results indicate that just before the memory trace is to be read out for comparison with a test object, its neural

code shifts to become more like the ensemble used for static, feedforward-induced memory traces in the same location. This convergence likely facilitates downstream comparison and decision-making processes by allowing similar mechanisms and readout weights for reading the same information out from WM stores.

Processing in most visual cortical areas, including PFC, is strongly biased toward the contralateral visual hemifield (Funahashi et al., 1990; Hagler and Sereno, 2006; Kastner et al., 2007; Medendorp et al., 2007; Pasternak et al., 2015; Rainer et al., 1998; Voytek and Knight, 2010; Wimmer et al., 2016). Prefrontal spiking activity (Buschman et al., 2011) and gamma power (Kornblith et al., 2015) increase with WM load—the number of items held in memory at one time—but only for items in the contralateral visual hemifield. In contrast, beta power shows increasing suppression for increasing numbers of items in either visual hemifield (Kornblith et al., 2015; Medendorp et al., 2007). Our results confirm these findings. This distinction might reflect the fact that beta oscillations are thought to correlate with broadly selective inhibitory processes (Engel and Fries, 2010; Jensen and Mazaheri, 2010; Lundqvist et al., 2016). It might also reflect beta's having a stronger influence from more bilateral top-down or recurrent signals, while spiking and theta/gamma oscillations are dominated by feedforward signals from strongly lateralized visual cortex (Bastos et al., 2015; van Kerkoerle et al., 2014).

A somewhat surprising result of this lateralization is that WM capacity is largely independent between the two visual hemifields. WM has a very limited capacity for holding multiple items at one time (Luck and Vogel, 1997, 2013). However, in both monkeys (Buschman et al., 2011) and humans (Delvenne, 2005; Umemoto et al., 2010), even when capacity is saturated in one visual hemifield, additional items can be stored in WM if they appear in the opposite hemifield. Similar effects of hemifield independence have been observed with spatial attention (Alvarez et al., 2012) and attentional tracking of moving objects (Alvarez and Cavanagh, 2005). This strong hemifield independence, however, seems inconsistent with the apparently seamless nature of visual WM. Our results provide a possible resolution to this paradox. They suggest that, in such situations, the two prefrontal hemispheres briefly sync up using theta and beta oscillations in order to physically transfer a WM trace from one cortical hemisphere to its new location on the retinotopic map in the opposite hemisphere.

We found neural dynamics that could support the seamless transfer of information between the cerebral hemispheres. We expect some of these same mechanisms are also used when interhemispheric communication is required for other cognitive processes, such as tracking moving objects across the midline or comparing visual information between the left and right hemifields. Fast, reliable interhemispheric communication is critical for many real-world behaviors, including sports, driving, and air traffic control. Interhemispheric communication is also thought to be disrupted in some disorders, such as dyslexia (Dhar et al., 2010). We hope that an understanding of the neural mechanisms of interhemispheric communication may lead to new ways to repair and optimize it.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- **KEY RESOURCES TABLE**
- **RESOURCE AVAILABILITY**
 - Lead contact
 - Materials availability
 - Data and code availability
- **EXPERIMENTAL MODEL AND SUBJECT DETAILS**
- **METHOD DETAILS**
 - Behavioral paradigm
 - Electrophysiological data collection
- **QUANTIFICATION AND STATISTICAL ANALYSIS**
 - Preprocessing
 - Mean activity analysis
 - Population decoding analysis
 - Synchrony analysis
 - Causality analysis
 - Oscillatory burst analysis
 - Latency analysis
 - Postprocessing and plotting
 - Hypothesis testing

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.neuron.2021.01.016>.

ACKNOWLEDGMENTS

We thank Alexa D'Ambra for assistance and Jesus Ballesteros, Andre Bastos, Sayak Bhattacharya, Alex Major, Morteza Moazami, Dimitris Pinotsis, and Jefferson Roy for helpful comments. This work was supported by NIMH R37MH087027, ONR MURI N00014-16-1-2832, the JPB Foundation (E.K.M.), and NIGMS T32GM007753 (J.A.D.).

AUTHOR CONTRIBUTIONS

J.A.D., S.K., and E.K.M. designed the experiments. J.A.D., M.K.M., and S.K. performed the experiments and recorded the data. S.L.B. curated the data and conceived and performed the analyses. S.L.B., E.K.M., and M.L. wrote the manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: August 13, 2020
Revised: December 17, 2020
Accepted: January 14, 2021
Published: February 8, 2021

REFERENCES

Alvarez, G.A., and Cavanagh, P. (2005). Independent resources for attentional tracking in the left and right visual hemifields. *Psychol. Sci.* *16*, 637–643.

Alvarez, G.A., Gill, J., and Cavanagh, P. (2012). Anatomical constraints on attention: hemifield independence is a signature of multifocal spatial selection. *J. Vis.* *12*, 9.

Barbas, H., and Pandya, D.N. (1984). Topography of commissural fibers of the prefrontal cortex in the rhesus monkey. *Exp. Brain Res.* *55*, 187–191.

Bastos, A.M., and Schoffelen, J.-M. (2016). A tutorial review of functional connectivity analysis methods and their interpretational pitfalls. *Front. Syst. Neurosci.* *9*, 175.

Bastos, A.M., Vezoli, J., Bosman, C.A., Schoffelen, J.-M., Oostenveld, R., Dowdall, J.R., De Weerd, P., Kennedy, H., and Fries, P. (2015). Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron* *85*, 390–401.

Benjamini, Y., and Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Ann. Stat.* *29*, 1165–1188.

Bouchacourt, F., and Buschman, T.J. (2019). A flexible model of working memory. *Neuron* *103*, 147–160.e8.

Buschman, T.J., and Miller, E.K. (2007). Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science* *315*, 1860–1862.

Buschman, T.J., Siegel, M., Roy, J.E., and Miller, E.K. (2011). Neural substrates of cognitive capacity limitations. *Proc. Natl. Acad. Sci. U S A* *108*, 11252–11255.

Chafee, M.V., Averbeck, B.B., and Crowe, D.A. (2007). Representing spatial relationships in posterior parietal cortex: single neurons code object-referenced position. *Cereb. Cortex* *17*, 2914–2932.

Chen, X., Zirnsak, M., and Moore, T. (2018). Dissonant representations of visual space in prefrontal cortex during eye movements. *Cell Rep.* *22*, 2039–2052.

Cohen, Y.E., and Andersen, R.A. (2002). A common reference frame for movement plans in the posterior parietal cortex. *Nat. Rev. Neurosci.* *3*, 553–562.

Colby, C.L., Duhamel, J.-R., and Goldberg, M.E. (1995). Oculocentric spatial representation in parietal cortex. *Cereb. Cortex* *5*, 470–481.

D'Esposito, M., and Postle, B.R. (2015). The cognitive neuroscience of working memory. *Annu. Rev. Psychol.* *66*, 115–142.

Dean, H.L., and Platt, M.L. (2006). Allocentric spatial referencing of neuronal activity in macaque posterior cingulate cortex. *J. Neurosci.* *26*, 1117–1127.

Delvenne, J.-F. (2005). The capacity of visual short-term memory within and between hemifields. *Cognition* *96*, B79–B88.

Dhamala, M., Rangarajan, G., and Ding, M. (2008). Analyzing information flow in brain networks with nonparametric Granger causality. *Neuroimage* *41*, 354–362.

Dhar, M., Been, P.H., Minderaa, R.B., and Althaus, M. (2010). Reduced inter-hemispheric coherence in dyslexic adults. *Cortex* *46*, 794–798.

Drew, T., Mance, I., Horowitz, T.S., Wolfe, J.M., and Vogel, E.K. (2014). A soft handoff of attention between cerebral hemispheres. *Curr. Biol.* *24*, 1133–1137.

Duhamel, J.-R., Bremmer, F., Ben Hamed, S., and Graf, W. (1997). Spatial invariance of visual receptive fields in parietal cortex neurons. *Nature* *389*, 845–848.

Engel, A.K., and Fries, P. (2010). Beta-band oscillations—signalling the status quo? *Curr. Opin. Neurobiol.* *20*, 156–165.

Fries, P. (2015). Rhythms for cognition: communication through coherence. *Neuron* *88*, 220–235.

Funahashi, S., Bruce, C.J., and Goldman-Rakic, P.S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiol.* *61*, 331–349.

Funahashi, S., Bruce, C.J., and Goldman-Rakic, P.S. (1990). Visuospatial coding in primate prefrontal neurons revealed by oculomotor paradigms. *J. Neurophysiol.* *63*, 814–831.

Fuster, J.M., and Alexander, G.E. (1971). Neuron activity related to short-term memory. *Science* *173*, 652–654.

Golomb, J.D., and Kanwisher, N. (2012). Higher level visual cortex represents retinotopic, not spatiotopic, object location. *Cereb. Cortex* *22*, 2794–2810.

Graziano, M.S., and Gross, C.G. (1998). Spatial maps for the control of movement. *Curr. Opin. Neurobiol.* *8*, 195–201.

Hagler, D.J., Jr., and Sereno, M.I. (2006). Spatial maps in frontal and prefrontal cortex. *Neuroimage* *29*, 567–577.

- Haufe, S., Nikulin, V.V., Müller, K.-R., and Nolte, G. (2013). A critical assessment of connectivity measures for EEG data: a simulation study. *Neuroimage* **64**, 120–133.
- Jensen, O., and Mazaheri, A. (2010). Shaping functional architecture by oscillatory alpha activity: gating by inhibition. *Front. Hum. Neurosci.* **4**, 186.
- Kastner, S., DeSimone, K., Konen, C.S., Szczepanski, S.M., Weiner, K.S., and Schneider, K.A. (2007). Topographic maps in human frontal cortex revealed in memory-guided saccade and spatial working-memory tasks. *J. Neurophysiol.* **97**, 3494–3507.
- Kornblith, S., Buschman, T.J., and Miller, E.K. (2015). Stimulus load and oscillatory activity in higher cortex. *Cereb. Cortex* **26**, 3772–3784.
- Ledoit, O., and Wolf, M. (2003). Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *J. Empir. Finance* **10**, 603–621.
- Luck, S.J., and Vogel, E.K. (1997). The capacity of visual working memory for features and conjunctions. *Nature* **390**, 279–281.
- Luck, S.J., and Vogel, E.K. (2013). Visual working memory capacity: from psychophysics and neurobiology to individual differences. *Trends Cogn. Sci.* **17**, 391–400.
- Lundqvist, M., Rose, J., Herman, P., Brincat, S.L., Buschman, T.J., and Miller, E.K. (2016). Gamma and beta bursts underlie working memory. *Neuron* **90**, 152–164.
- Lundqvist, M., Herman, P., Warden, M.R., Brincat, S.L., and Miller, E.K. (2018). Gamma and beta bursts during working memory readout suggest roles in its volitional control. *Nat. Commun.* **9**, 394.
- Luria, R., Balaban, H., Awh, E., and Vogel, E.K. (2016). The contralateral delay activity as a neural measure of visual working memory. *Neurosci. Biobehav. Rev.* **62**, 100–108.
- Manly, B.F.J. (2007). *Randomization, Bootstrap, and Monte Carlo Methods in Biology* (Chapman & Hall/CRC).
- Medendorp, W.P., Kramer, G.F.I., Jensen, O., Oostenveld, R., Schoffelen, J.-M., and Fries, P. (2007). Oscillatory activity in human parietal and occipital cortex shows hemispheric lateralization and memory effects in a delayed double-step saccade task. *Cereb. Cortex* **17**, 2364–2374.
- Miller, E.K., Erickson, C.A., and Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J. Neurosci.* **16**, 5154–5167.
- Neupane, S., Guitton, D., and Pack, C.C. (2016). Two distinct types of remapping in primate cortical area V4. *Nat. Commun.* **7**, 10402.
- Olson, C.R. (2003). Brain representation of object-centered space in monkeys and humans. *Annu. Rev. Neurosci.* **26**, 331–354.
- Ong, W.S., Hooshvar, N., Zhang, M., and Bisley, J.W. (2009). Psychophysical evidence for spatiotopic processing in area MT in a short-term memory for motion task. *J. Neurophysiol.* **102**, 2435–2440.
- Oostenveld, R., Fries, P., Maris, E., and Schoffelen, J.-M. (2011). FieldTrip: open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Comput. Intell. Neurosci.* **2011**, 156869.
- Pasternak, T., Lui, L.L., and Spinelli, P.M. (2015). Unilateral prefrontal lesions impair memory-guided comparisons of contralateral visual motion. *J. Neurosci.* **35**, 7095–7105.
- Pouget, A., and Snyder, L.H. (2000). Computational approaches to sensorimotor transformations. *Nat. Neurosci.* **3** (Suppl), 1192–1198.
- Rainer, G., Asaad, W.F., and Miller, E.K. (1998). Memory fields of neurons in the primate prefrontal cortex. *Proc. Natl. Acad. Sci. U S A* **95**, 15008–15013.
- Rigotti, M., Barak, O., Warden, M.R., Wang, X.-J., Daw, N.D., Miller, E.K., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585–590.
- Romo, R., Brody, C.D., Hernández, A., and Lemus, L. (1999). Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature* **399**, 470–473.
- Tomita, H., Ohbayashi, M., Nakahara, K., Hasegawa, I., and Miyashita, Y. (1999). Top-down signal from prefrontal cortex in executive control of memory retrieval. *Nature* **401**, 699–703.
- Torrence, C., and Compo, G. (1998). A practical guide to wavelet analysis. *Bull. Am. Meteorol. Soc.* **79**, 61–78.
- Umemoto, A., Drew, T., Ester, E.F., and Awh, E. (2010). A bilateral advantage for storage in visual working memory. *Cognition* **117**, 69–79.
- Ungerleider, L.G., Courtney, S.M., and Haxby, J.V. (1998). A neural system for human visual working memory. *Proc. Natl. Acad. Sci. U S A* **95**, 883–890.
- van Kerkoerle, T., Self, M.W., Dagnino, B., Gariel-Mathis, M.-A., Poort, J., van der Togt, C., and Roelfsema, P.R. (2014). Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. *Proc. Natl. Acad. Sci. U S A* **111**, 14332–14341.
- Vinck, M., van Wingerden, M., Womelsdorf, T., Fries, P., and Pennartz, C.M.A. (2010). The pairwise phase consistency: a bias-free measure of rhythmic neuronal synchronization. *Neuroimage* **51**, 112–122.
- Voytek, B., and Knight, R.T. (2010). Prefrontal cortex and basal ganglia contributions to visual working memory. *Proc. Natl. Acad. Sci. U S A* **107**, 18167–18172.
- Wimmer, K., Spinelli, P., and Pasternak, T. (2016). Prefrontal neurons represent motion signals from across the visual field but for memory-guided comparisons depend on neurons providing these signals. *J. Neurosci.* **36**, 9351–9364.
- Zaksas, D., Bisley, J.W., and Pasternak, T. (2001). Motion information is spatially localized in a visual working-memory task. *J. Neurophysiol.* **86**, 912–921.
- Zirnsak, M., Steinmetz, N.A., Noudoost, B., Xu, K.Z., and Moore, T. (2014). Visual space is compressed in prefrontal cortex before eye movements. *Nature* **507**, 504–507.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental models: organisms/strains		
<i>Macaca mulatta</i>	California National Primate Research Center	https://cnprc.ucdavis.edu/
Software and algorithms		
Python 3	Anaconda	3.6; RRID: SCR_008394; https://www.anaconda.com/
scikit-learn	Anaconda	0.22.1; RRID: SCR_002577; https://www.anaconda.com/
statsmodels	Anaconda	0.11.1; RRID: SCR_016074; https://www.anaconda.com/
MATLAB	The Mathworks Inc.	R2019b; RRID: SCR_001622; https://www.mathworks.com/
FieldTrip	FieldTrip developers	RRID: SCR_004849; https://www.fieldtriptoolbox.org/

RESOURCE AVAILABILITY

Lead contact

Requests for information should be directed to the lead contact, Earl Miller (emkiller@mit.edu).

Materials availability

This study did not generate new unique reagents.

Data and code availability

The neural data and code used in this study will be made available upon reasonable request to the lead contact.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

The nonhuman primate subjects in our experiments were two adult (ages 16 and 9) rhesus macaques (*Macaca mulatta*), one male and one female. All procedures followed the guidelines of the Massachusetts Institute of Technology Committee on Animal Care and the National Institutes of Health.

METHOD DETAILS

Behavioral paradigm

Subjects performed a delayed nonmatch-to-sample working memory (WM) task (Figure 1A). They began task trials by holding gaze for 700 ms on a fixation point randomly displayed at 4.5° left or right of the center of a computer screen. A sample object was then shown for 700 ms in the center of the screen, thus in the right or left visual hemifield, respectively. Two sample objects were used each session, chosen from a commercial photo library (Hemera Photo-Objects). Sample objects were displayed in one of two positions, 3.4° above or below the screen center. After a 1.6 s delay, a test object was displayed for 400 ms. The monkeys were required to saccade to it if it did not match the remembered sample in either object identity or upper/lower location. If the test was identical to the sample, they withheld response. Then, after a 100 ms blank period, a nonmatching test object was always shown, which required a saccade. Response to the non-match was rewarded with juice, followed by a 3.2 s inter-trial interval.

A random 50% of trials had an uninterrupted 1.6 s WM delay (Figure 1A, left). In the other 50%, at 800 ms into the delay, the fixation point jumped to the opposite location on the screen (Figure 1A, right). The monkeys were trained to immediately saccade to it and reacquire fixation. Once fixation was acquired again, the WM delay was continued for another 800 ms, equating the full time of fixated delay period with the no-swap condition.

All stimuli were displayed on an LCD monitor. An infrared-based eye-tracking system (Eyelink 1000 Plus, SR-Research, Ontario, CA) continuously monitored eye position at 1 kHz.

Electrophysiological data collection

The subjects were chronically implanted bilaterally in the lateral prefrontal cortex (PFC) with four 8x8 iridium-oxide “Utah” microelectrode arrays (1.0 mm length, 400 μm spacing; Blackrock Microsystems, Salt Lake City, UT), for a total of 256 electrodes (Figure 1C).

Arrays were implanted bilaterally, one array in each ventrolateral and dorsolateral PFC. Electrodes in each hemisphere were grounded and referenced to a separate subdural reference wire. LFPs were amplified, low-pass filtered (250 Hz), and recorded at 30 kHz. Spiking activity was amplified, filtered (250–5,000 Hz), and manually thresholded to extract spike waveforms. All threshold-crossing spikes on each electrode were pooled together and analyzed as multi-unit activity (MUA).

QUANTIFICATION AND STATISTICAL ANALYSIS

All correctly performed trials were included in analyses. All analyses of individual MUA and LFP channels were averaged across all electrodes in each hemisphere, and analyses of channel pairs were averaged across all between-hemisphere or all within-hemisphere pairs. Analysis was initially performed separately for each prefrontal hemisphere and sample object hemifield. Results were then pooled based on whether the sample was contralateral or ipsilateral to a given hemisphere by averaging across appropriate hemisphere/hemifield combinations. This resulted in a set of observations for each experimental session. All plots depict means and standard errors across all 56 sessions in this dataset, and all statistics were performed with sessions treated as observations ($n = 56$). All preprocessing and analysis was performed in Python 3.6 or MATLAB R2019b (The Mathworks, Inc, Natick, MA).

Preprocessing

Spike rates were computed by binning spike timestamps in non-overlapping 50 ms windows. Spike rates were square-root transformed prior to analysis to convert their Poisson-like distributions to approximately normal. LFPs were re-referenced offline to remove any common-source noise by subtracting off the mean across all electrodes in each array. Evoked potentials were removed by subtracting off the mean signal across trials within each condition (object, upper/lower location, and visual hemifield). Thus, all of our analysis is on the remaining induced component. For most analyses, the resulting signals were convolved with a set of complex Morlet wavelets (wavenumber 6). LFP power was log-transformed to render its distribution approximately normal.

Mean activity analysis

To normalize out any overall differences in activity between neurons and task conditions, we z-scored spike rates to the fixation baseline. Rates were mean-pooled across the 200 ms before sample object onset separately for each neuron and condition, and the mean and standard deviation across all within-condition trials was computed. These were used to z-score rates across all time points and trials for that neuron and condition. The same transformation was used for analysis of LFP power, except it was also computed separately for each frequency, in order to also normalize out the typical $1/f$ distribution of power across frequency.

To assay for spiking activity changes from baseline in either direction (increases or decreases), we took the absolute value of the measured z-scores for each multi-unit (Figure S1A). Significance was evaluated using paired-observation tests relative to baseline. We also measured delay activity separately for subpopulations of multi-units with responses above or below baseline (Figure S1B). Since PFC units often have complex dynamics, this segregation into subpopulations was performed independently for each unit and time point. To avoid circularity in the analysis, segregation into subpopulations and computation of mean activity within each subpopulation were performed on odd and even-numbered trials, respectively.

Population decoding analysis

Spike rates for all multi-units within a given hemisphere were used as independent features in a linear classifier that decoded which of four task conditions—two objects \times two upper/lower locations—was present in each trial. Classification was performed independently on spike rate data from each time point (50 ms window). All reported classification accuracies were obtained via 5-fold cross-validation, in which trials were randomly split into five non-overlapping subsets and each classifier was trained on four of these, while its accuracy was evaluated on the final, untrained one. This process was repeated five times with each subset acting as the test set once, and the final results were averaged across the five folds. The same procedure was used for the cross-classification analysis (Figure 6), except that training and testing trials were selected from different task conditions. For the cross-temporal analysis (Figure S3), cross-classification was also performed using all possible combinations of pairs of time points for training and testing. All decoding analysis was performed with a linear discriminant classifier with optimal covariance shrinkage (Ledoit and Wolf, 2003), using the Python scikit-learn library.

As an alternative measure of difference in neural activity elicited by task conditions, we also computed the mean rate for each multi-unit's preferred, intermediate, and nonpreferred WM items (objects \times upper/lower locations; Figures S1C and S1D). This was computed separately for each multi-unit and time point, since PFC units often have complex dynamics with different preferences at different time points. To avoid circularity, calculation of item preference and of mean rates for the preference-sorted items was performed on odd and even-numbered trials, respectively.

Synchrony analysis

LFP-LFP phase synchrony was computed from the LFP of the complex wavelet transform, using the pairwise phase consistency (PPC). PPC is a measure of how consistent across trials the relative phase angles between two signals are, independent of their absolute phase. It is an unbiased estimator of the square of the mean vector resultant length (Kornblith et al., 2015; Vinck et al., 2010).

Presented results represent averages across all electrode pairs between hemispheres (Figures 7A–7C), all electrode pairs within the hemisphere contralateral to the sample hemifield (the “sender” on swap trials; Figures 7D, 7F, and 7H), and all electrode pairs within the hemisphere ipsilateral to the sample hemifield (“receiver” on swap trials: Figures 7E, 7G, and 7I) within each session. These values were then averaged across all sessions.

Causality analysis

Directional causal influences between LFPs in the two prefrontal hemispheres were measured using bivariate nonparametric spectral Granger causality (Dhamala et al., 2008). Traditional Granger causality quantifies how much variance in one signal can be explained by the recent history of another signal, beyond what can be explained by the history of the signal itself. Spectral Granger expresses these causal interactions in the frequency domain. Unlike traditional parametric causality measures, it is estimated directly via factorization of the cross-spectral density matrix, without relying on estimation of a specific autoregressive model. For this analysis only, LFP spectra were computed around the putative trace transfer period (–1 to –0.5 s relative to delay end) via the multitaper method (4 Hz frequency bandwidth, 3 dpss tapers). Trials were balanced across swap and no-swap conditions in each session by sampling a random subset of trials from the larger group. Presented results represent averages across all between-hemisphere electrode pairs within each session, and across all sessions. As a control, this analysis was also performed on data reversed in time (Figure S4). This is expected to reverse the direction of true causality, but preserve the direction of any apparent causality due to confounding between-area differences in power or signal-to-noise (Bastos and Schoffelen, 2016; Haufe et al., 2013). Inference was based on the conjunction of significance in the standard analysis and significance in the opposite causal direction in the time-reversed analysis. This analysis was performed using the FieldTrip toolbox for MATLAB (Oostenveld et al., 2011).

Oscillatory burst analysis

Methods followed those of our previous work on oscillatory bursts (Lundqvist et al., 2016). Power was computed from a wavelet transform, as described above, and pooled within several frequency “sub-bands”: low-beta (10–20 Hz), high-beta (20–32 Hz), low-gamma (40–65 Hz), gamma (55–90 Hz), and high-gamma (70–100 Hz). Within each sub-band, the mean and SD of power across all trials and within-trial time points was computed and used to threshold the data at 2 SDs above the mean. Bursts were defined as periods where power remained above threshold for at least 3 oscillatory cycles (defined at the center frequency of each sub-band). This resulted in a binary representation of time points within versus outside bursts. These were pooled separately within all beta and gamma sub-bands by labeling as bursts time points where a burst occurred within any sub-band (logical OR across sub-bands). Burst rate for each task condition, band, and time point was computed as the proportion of trials containing within-burst time points. Finally, the difference in burst rate at each time point from the fixation baseline was computed and plotted in Figure S2.

Latency analysis

Latencies were measured for all reliable temporal features—points halfway from a baseline to a peak or trough—of neural signals surrounding the epoch of WM trace transfer (Figure S5). These were measured on across-session mean neural signals, and their 99% confidence intervals were computed by bootstrapping across sessions 10,000 times. Each bootstrap iteration resampled sessions with replacement, computed the same across-session mean neural signals, and estimated resampled latencies from them.

Postprocessing and plotting

To clarify trends in the data, all plotted results were smoothed with 1D or 2D Gaussian kernels for plotting purposes only. The Gaussian standard deviations used were 10 ms for all time axes, and 0.05 octaves for all frequency axes.

The temporal extent over which time-frequency effects due to the test object could possibly extend into the delay period was computed as the time for wavelet power at each frequency to drop by a factor of e^{-2} relative to the test stimulus onset (Torrence and Compo, 1998). This is a conservative overestimate of the backward extent of test-object effects since it does not take any latency of test-evoked neural signals into account. This period is plotted as a transparent gray region at the right of all time-frequency and frequency-band summary plots.

Due to off-axis diagonal structure in some of the time-frequency data (time-frequency inseparability), it cannot be fully captured in any 1D summary. Thus, all statistical analyses and inference were performed on the full time-frequency data. However, we also include 1D plots of the time course within frequency bands simply to aid in understanding the results. The cutoffs for each frequency band were chosen to approximate values typically used to delineate standard frequency bands in the literature: theta (3–8 Hz), beta (10–32 Hz), and gamma (40–100 Hz).

Hypothesis testing

All hypothesis tests used non-parametric randomization methods that do not rely on specific assumptions about data distributions (Manly, 2007). Each session was treated as an observation, and randomizations were performed across sessions. All randomization statistics were resampled 10,000 times and evaluated in a two-tailed fashion.

To test whether a mean value differed significantly from baseline, we used a randomized sign test in which a t -statistic was computed on both the observed data, and on data where the sign of each baseline-centered observation was randomly flipped.

To test whether the means of paired observations were significantly different, we used a permutation paired t test, in which a paired-sample t -statistic was computed on the observed data, and on data with the labels of each pair of observations randomly swapped. To test significance of multiple main effects and their interaction, we used a permutation 2-way ANOVA in which an F -statistic was computed on the observed data, and on data where the multi-factor labels were randomly shuffled as a group across trials. All tests were corrected for multiple comparisons across time points and/or frequencies using a procedure that controls the false discovery rate under arbitrary dependence assumptions (Benjamini and Yekutieli, 2001) using the Python statsmodels module.