



Supplementary Materials for

Coupled Electrophysiological, Hemodynamic, and Cerebrospinal Fluid Oscillations in Human Sleep

Nina E. Fultz, Giorgio Bonmassar, Kawin Setsompop, Robert A. Stickgold, Bruce R.
Rosen, Jonathan R. Polimeni, Laura D. Lewis.

correspondence to: ldlewis@bu.edu

This PDF file includes:

Materials and Methods
Supplementary Text
Figs. S1 to S10
References

Materials and Methods

Subject enrollment and imaging sessions:

All subject provided written informed consent and all procedures were approved by the Massachusetts General Hospital Institutional Review Board. A total of 15 subjects were studied and 13 were analyzed.

We imaged 15 subjects in sessions beginning at approximately midnight. One subject was excluded for poor EEG quality and one was excluded after stopping the scan to report he/she could not sleep comfortably. The remaining 13 subjects were aged 23–33 (11 female), and reported no neurological, psychiatric, or sleep disorders. Subjects were screened for MR contraindications and to have habitual sleep duration of no less than six hours, daily caffeine consumption of less than 300 mg, not smoke cigarettes, weigh less than 250 pounds, not be pregnant, be comfortable sleeping on their back, and not be taking medications that affect sleep. Subjects self-reported head size (ranging from small to extra-large) in a pre-screening form, and those reporting large head size were excluded due to the space constraints of the RF receive coil housing. Due to the head size constraint, a majority of the subjects were female. While we did not have an a priori prediction of gender effects and were not powered to test for these, we examined the data to evaluate whether large differences were apparent and found that CSF oscillations were detected in both male and female subjects (Fig. S8). All subjects had previously participated in a daytime MR imaging session to ensure comfort in the scanner environment. Subjects were asked to sleep only four hours the night prior to the imaging study, to increase sleep pressure during the scan. One subject failed to complete the sleep restriction protocol but was still included and still fell asleep in the scanner. During sleep acquisition sessions, subjects were asked to rest with eyes closed. In a subset of sessions ($n=6$), subjects were asked to press a button on an MR-compatible USB button box on each breath in and each breath out, to provide an additional behavioral metric of sleep (43). Instructions asked subjects to allow themselves to fall asleep but to return to performing to the task if they awoke during the scan.

EEG acquisition

EEG was acquired using MR-compatible 256-channel geodesic nets and a NA410 amplifier (Electrical Geodesics, Inc., Eugene, OR USA) at a sampling rate of 1000 Hz. EEG acquisition was synchronized to the scanner 10 MHz clock to reduce aliasing of high-frequency gradient artifacts. The scanner cryopump was temporarily shut off during EEG acquisition to reduce vibrational artifact. To acquire reference signals to be used for EEG noise removal, subjects wore a reference layer cap composed of an isolating vinyl layer and conductive satin layer on the head, with grommets inserted to allow electrodes to pass through and make contact with the scalp (44), while other electrodes remained isolated from the scalp and recorded the noise, resulting in a total of 30–36 EEG electrodes per subject.

Physiological signals were simultaneously acquired using a Physio16 device (Electrical Geodesics, Inc., Eugene, OR USA). ECG was measured through two disposable electrodes placed on the chest diagonally across the heart, with an MR-compatible lead (InVivo Corp, Philips). Respiration was measured through a piezo-electrical belt (UFI systems, Morro Bay, CA USA) around the chest.

fMRI acquisition

Subjects were scanned on a 3T Siemens Prisma scanner with the vendor-supplied 64-channel head-and-neck coil. Each session began with a 1 mm isotropic multi-echo MPRAGE to provide an anatomical reference (45). Functional runs acquired 40 interleaved BOLD-weighted EPI slices with 2.5 mm³ isotropic voxels. This acquisition volume covered most of the brain, omitting the tips of the temporal poles, which also exhibit higher signal dropout. fMRI protocols consisted of a single-shot gradient echo SMS-EPI (46) with MultiBand factor=8, matrix=92×92, blipped CAIPI shift=FOV/4, TR=367 ms, nominal echo-spacing=0.53 ms, flip angle=32-37°, and no in-plane acceleration. Two slightly different fMRI scanning acquisition parameters were used due to a scanner software upgrade partway through the study: the TE was set to 32 ms in the first seven subjects and 30 ms in the remaining subjects. In one subject the TR was set to 387 ms due to technical error. VERSE factor was set between 1 and 1.5 depending on individual subject SAR constraints. Individual runs could last up to 2 hours. If runs ended earlier, subsequent runs would be started up to a maximum total scan duration of 2.5 hours as long as subjects were still comfortable and sleeping.

EEG preprocessing

Gradient artifacts were removed through average artifact subtraction (47), using a moving average of the previous 20 TRs. This gradient artifact removal acts as a high-pass filter and EEG analysis therefore focused on >0.2 Hz signals. Electrodes were then re-referenced to the common average, computing this separately for electrodes contacting the head, and those placed on the reference layer. Channels on the cheeks and borders of the reference cap were excluded from the common average. Ballistocardiogram artifacts were removed using regression of reference signals from the isolated EEG electrodes (44, 48). Since there was a larger number of noise electrodes than signal electrodes, the regression was performed after subsampling the noise electrodes, using only every fourth isolated electrode. Because the position of and physiological noise influences on the electrodes can vary over the long recording times used here, we implemented a dynamic time-varying regression of the reference signals. Beta coefficients for the best-fit regression within 30 s sliding time windows were fit using least-squares; these beta values were then linearly interpolated over the nonoverlapping windows. The resulting interpolated beta value at every time point was then used for a local subtraction of the reference signals from the modeled EEG recording. This regression was performed individually for each EEG channel.

BOLD fMRI data preprocessing

Processing streams are shown in Fig. S9. CSF inflow analysis used the raw acquired fMRI data with slice-timing correction but without any motion correction, as motion correction corrupts the voxel slice position information needed for inflow analysis, and motion correction cannot be accurately performed on edge slices where tissue moves in and out of the imaging volume. All analysis therefore was performed after selecting low-motion epochs. For BOLD analysis, fMRI data were slice-timing corrected using FSL (<https://fsl.fmrib.ox.ac.uk/fsl>) and motion corrected with AFNI (<https://afni.nimh.nih.gov>). Physiological noise removal was performed using dynamic

regression based on the concept of RETROICOR (49) and adapted for fast fMRI as follows. The respiratory trace was bandpass filtered between 0.16–0.4 Hz using a finite impulse response filter and the instantaneous phase was computed as the angle of the Hilbert transform. The cardiac peaks were detected automatically using the FASST toolbox (<http://www.montefiore.ulg.ac.be/~phillips/FASST.html>) and the phase was modeled as varying linearly between each identified peak. Sine and cosine basis functions using the phase of the signal and its second harmonic were generated as regressors for physiological noise. This regression was performed over 1000 s windows sliding every 400 s to enable high-quality physiological noise removal as the heart rate and respiratory rate varied throughout the scan. No spatial smoothing was applied.

Sleep/wake segment extraction

Because our analysis of CSF signals required non-motion-corrected data (to measure signal at the edge slices) and long continuous epochs (to analyze continuous low-frequency dynamics), we restricted our analysis to periods of stable wake or sleep with low motion. Our analysis focused on long segments of stable and unambiguous continuous NREM sleep or wake (REM epochs were not seen in our data, likely due to the long time needed to reach REM sleep). We therefore based our sleep and wake segment identification by examining ongoing dynamics in the EEG spectrograms rather than performing conventional sleep scoring in discrete windows. To select continuous sleep and wake segments, we plotted occipital EEG spectrograms (from the channel nearest to OZ with good recording quality), behavior (in the subset of subjects performing the task), and translational and rotational motion estimated automatically from the fMRI timeseries using AFNI. EEG signatures of sleep included loss of occipital alpha (8–12 Hz) rhythms and increased delta (0.5–4 Hz) and theta (4–8 Hz) power. The occipital EEG channel was selected both to provide the ability to identify disappearance of occipital alpha rhythms at sleep onset, allowing for clear segmentation of wake and sleep, and because occipital EEG has the highest signal quality in the MR environment. Periods of at least 90 s of low motion and either stable wake or NREM were manually identified and extracted for further analysis. Since we required stable low-motion sleep for >90 s, whereas sleep substages are typically defined in 30 s windows (50), and since N2 is the majority of human NREM sleep, these longer >90 s NREM segments were predominantly N2 sleep, but could nevertheless also include some N1 or N3 (Fig. S2).

ROI definition

All fMRI signal analyses were performed within individual subjects in the original spatial frame of the fMRI acquisition space, without transforming to a common average, to avoid spatial blurring of these signals across space in the registration process.

The ROI for the fourth ventricle and aqueduct was defined anatomically on the functional images. An initial registration matrix between the functional and anatomical images was calculated using boundary-based registration (51). The registered MPRAGE was overlaid onto the functional image to identify the approximate position of the ventricle/aqueduct, and then the brightest voxels on the functional image were selected to identify the CSF region. ROIs for the BOLD signal in cortical gray matter were defined using the automated segmentation generated by Freesurfer (<https://surfer.nmr.mgh.harvard.edu>) (52) on the MPRAGE, and then registered to the

functional volume, to include all cortical gray matter present in the imaging volume (i.e. most of the cortex).

Inflow analysis:

For analysis of inflow dynamics, the CSF ROI was split into separate subROIs within individual slices, consisting of its projection onto the bottom four slices of the functional acquisition volume. The mean signal from the subROI of the CSF on each slice was then extracted. Only subjects for whom the CSF region could be identified on all four bottom slices were included in this analysis ($n=11$ subjects). The critical velocity for inflow signals in slice 2 was calculated (53) as the slice width (2.5 mm) divided by the temporal gap between slice 1 and slice 2 (220.2 ms, due to interleaved acquisition). To capture the range of signal fluctuations between low- and high-flow conditions, the signal magnitude was calculated as the relative ratio of the 95th percentile and 5th percentile of the signal in each ROI over time.

Spectral power analysis:

The power of fMRI and EEG signals was calculated using multitaper spectral estimation (Chronux, (54)). Power in the BOLD signal was estimated in the 90 s segments identified as described above, and the mean power in each patient was calculated across each segment. Pairwise comparisons across sleep and wake segments were computed within the patients who exhibited both sleep and wake data ($n=11$; the remaining two subjects only had sleep segments) using the Wilcoxon signed-rank test, and confidence intervals were estimated by inverting the test. When comparisons were performed across segments (Fig. S2), because the segments were not paired, differences and confidence intervals were calculated using the Wilcoxon rank-sum test. The BOLD and CSF analyses used 5 tapers and the EEG analysis used 59 tapers. EEG analyses were performed on the occipital EEG channel closest to OZ that was also identified as having good data quality in order to minimize ballistocardiogram artifact induced by motion in the magnetic field (as occipital channels are anchored by the head, whereas frontal channels exhibit larger artifacts) and to allow analysis of occipital alpha to track sleep onset. 95% confidence intervals for the power spectra were estimated by bootstrapping across subjects with 1000 times with replacement.

Cross-correlation analysis:

The CSF and BOLD timeseries were respectively extracted from the ventricle and cortical gray matter ROIs (Fig. S9). For comparing the relative flow, the BOLD and CSF signals were lowpass filtered below 0.1 Hz. This filtering was performed to reduce noise but results were also consistent when repeating the analysis without this filtering step (Fig. S10). Next, the time derivative of the BOLD signal was computed. For the analysis in Fig. 3, this derivative was multiplied by -1 and all negative values were set to zero. This zero-thresholding was applied to restrict the analysis to inflow but not outflow signals, because our imaging technique only measures inflow of CSF and not outflow. Cross-correlation was computed after detrending data and was normalized such that the autocorrelation of each signal was 1. The EEG was filtered into the 0.2–4 Hz band using a finite impulse response filter, and its instantaneous amplitude was estimated as the magnitude of the Hilbert transform, which was then smoothed with a moving average of

4 s. To create a threshold for statistical testing and estimate p-values, the segment indices of the sleep data were shuffled such that the two variables being correlated were drawn from different data segments, and the maximal cross-correlation over lags (-20 to 20 s) was computed. This procedure was repeated 1000 times and the p-value was calculated from percentiles of this shuffled distribution.

CSF peak-locked analysis:

Peaks in the CSF signal during sleep were identified by filtering the CSF signal between 0-0.1 Hz and identifying local maxima that surpassed a threshold of 20% amplitude. The 0-0.1 Hz filtering was used only to identify the peaks of the CSF waves, and the unfiltered CSF signals were then analyzed, to minimize processing of the data. These peaks were then used to extract a peak-locked signal in the EEG 0.2-4 Hz amplitude envelope, the time derivative of the BOLD signal, and the CSF signal (Fig. 4a,b,c). To test for statistical significance of the peak in the EEG signal, the relative timing of each peak-locked signal was shuffled randomly, and the mean EEG amplitude signal was recalculated on the shuffled data. This process was repeated 1000 times to create a null distribution for peak EEG amplitudes.

Model simulations:

We first simulated the relationship between EEG power, BOLD, and CSF signals by calculating the flow of CSF as the negative of CBF. Physiological parameters were first drawn from previous studies (33, 35) to test whether these dynamics could be explained without requiring any additional parameter fitting. Then, we numerically optimized model parameters to estimate the best-fit coupling functions between EEG, hemodynamics, and CSF. The equations used in the model are reported in the Supplementary Text.

Because the magnitude of the relationship between the neural activity, EEG, and flow is not known, and BOLD does not provide quantitative information, this model did not employ physical units but instead aimed to capture the time-varying dynamics of these signals as quantified by the correlation coefficient. We calculated the mean correlation coefficient between the predicted CSF flow and the measured CSF timeseries across all segments. We compared this metric to a shuffled distribution in which the prediction from one time segment was tested against the data from a different time segment. This shuffling was repeated 1000 times and the threshold was set as the 2.5th and 97.5th percentile of the shuffled distribution. When parameter fitting was used (Fig. 4d), the R-value was further assessed through tenfold cross-validation, and statistics report the mean and standard deviation of R-values measured in cross-validated test data. Model performance was also evaluated across segments after sorting segments based on CSF wave amplitude, calculated as the range of the 0-0.1 Hz filtered CSF signal.

Supplementary Text

Model with fixed parameters

The blood flow response to neural activity was calculated as:

$$f(t) = -n * h(t) \quad (1)$$

where $f(t)$ is the relative cerebral blood flow (CBF), which is always positive and is normalized to a value of 1; n is the power envelope of the EEG signal between 0.2 Hz and 4 Hz; and $h(t)$ is the flow impulse response to neural activity, modeled as the gamma distribution:

$$h(t) = \frac{(t/\tau_f)^{(z-1)} \exp(-t/\tau_f)}{\tau_f(z-1)!} \quad (2)$$

The value for τ_f was set at 2.1 and z was set at 3 based on the results in Simon and Buxton, 2015.

We added a term for CSF such that a decrease in blood volume would elicit an increase in CSF volume. Because the exact link between volume and CSF flow rates is not known, we approximated CSF flow as the opposite of the cerebral blood flow:

$$CSF = -f(t) + 1; \quad (3)$$

This model simplifies the relationship between blood and CSF by assuming that blood flow and CSF flow changes are exactly coupled, and assuming that net CSF flow is zero. The CSF term thus includes an offset of 1, as CSF flow can be negative or positive, and is centered at zero. In contrast, the CBF $f(t)$ term is always positive, representing inflow of fully oxygenated blood, and is normalized to 1. The cross-correlation between this CSF prediction and the CSF signal was then calculated.

Model with varying parameters

We used numerical optimization to examine the best-fit impulse response between EEG and CSF. We first fit the shape and scale parameters of a gamma distribution, using as the cost function the root-mean-squared error between the CSF prediction and the true CSF signal. This process was used to generate the EEG-CSF impulse response in Fig. 4d.


To compare the best-fit impulse response for EEG-CSF to the model predictions of CBF and cerebral blood volume (CBV), we used the CBF equations as above and calculated the predicted CBV change using the balloon model as in Buxton et al., 2004:

$$\frac{dv}{dt} = \frac{1}{\tau_{MTT}} (f(t) - f_{out}(v,t)) \quad (4)$$

$$f_{out} = v^{\frac{1}{\alpha}} + \tau_v \frac{dv}{dt} \quad (5)$$

The physiological parameters were drawn from Buxton et al., 2004, and fixed at $\tau_{MTT}=4$, $E_0=0.4$, $\alpha=0.4$. To find the best fit delay between flow and volume, we varied either the CBF parameters (τ_f and z) or varied the viscoelastic time constant τ_v in the previously published range between 0 and 30. The model fitting minimized the difference between the derivative of CBV and the inverse of CSF flow (e.g., positive CSF flow when d/dt CBV is negative).

Schematic of signal reaching steady state



As successive RF pulses are applied to an imaging volume, the signal intensity decays to a steady state which remains for the rest of the imaging period. This causes an initial transient in the images at the first few timepoints.

Schematic of flow-related enhancement




Fig. S1: Schematic of inflow effects in rapid acquisition paradigms. The stationary tissue inside the imaging volume is at steady state, having already experienced a large number of RF pulses. When fresh fluid begins to flow into the volume, the fluid will be bright, as it has not yet experienced an RF pulse. After the first image is taken, the fluid will now have experienced 1 RF pulse, and will have reduced signal as it moves upward into the second slice, while the continuing flow of fresh fluid will lead to continued bright signals in the bottom slice. As fluid continues to flow into upper slices, it will further decrease its signal (receiving additional RF excitation as each image is taken), until it too reaches steady state. This signal corresponds to the well-established principle of MR flow-related enhancement, which appears in fast fMRI acquisition (53, 55).

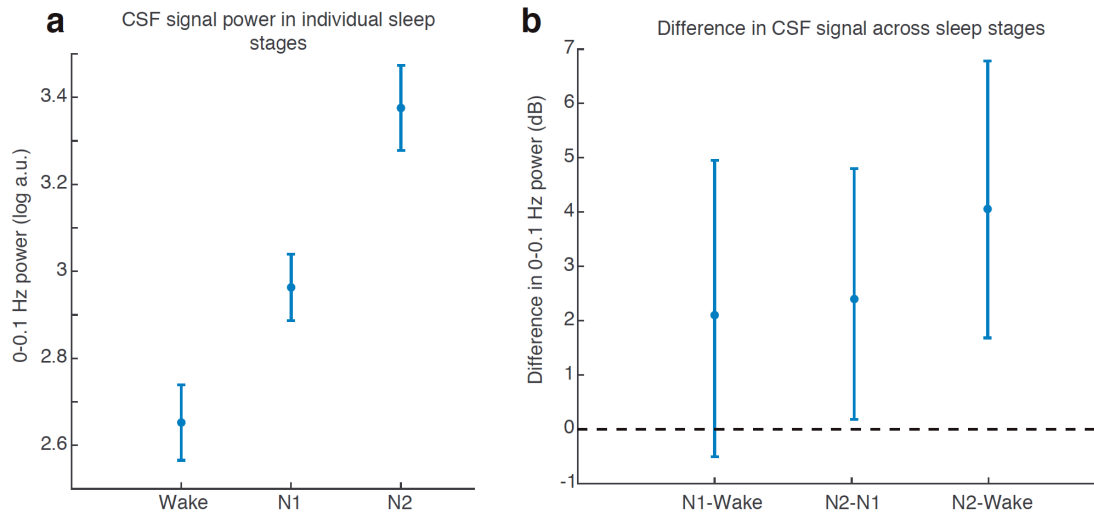


Fig. S2: Sleep scoring of sleep segments. Our analysis used the spectral dynamics of the EEG and metrics of behavior and motion to identify low-noise, stable, sleep segments lasting at least 90 seconds in order to examine continuous low-frequency signals. When applying conventional sleep scoring in 30 s windows, N2 predominates in these segments, in part because it is the stage in which people spend the majority of their NREM sleep, but also because it is most likely to be maintained for the long stable segments as subjects sleep in the scanner. A) CSF signal power during wake, N1-only, and N2-only segments (wake segments $n=45$; N1-only segments $n=23$; N2-only segments $n=73$). Error bars are standard error. B) Effect size for the difference between conditions. Error bars are 95% confidence intervals calculated by Wilcoxon ranksum; black dashed line indicates zero; n is as in panel A. N2 sleep exhibits significantly more CSF signal power than N1 sleep ($p=0.039$, Wilcoxon ranksum). We did not detect a statistically significant increase in the CSF signal in N1 relative to wake ($p=0.19$, Wilcoxon ranksum), but we note that there is reduced statistical power for this analysis (as compared to N2), as few of the segments were exclusively N1. It therefore remains possible that N1 contains a smaller CSF effect that did not reach statistical significance. While some N3 scores were present in our data, sufficient contiguous N3 scores for this analysis of segments containing exclusively one sleep stage were not present.

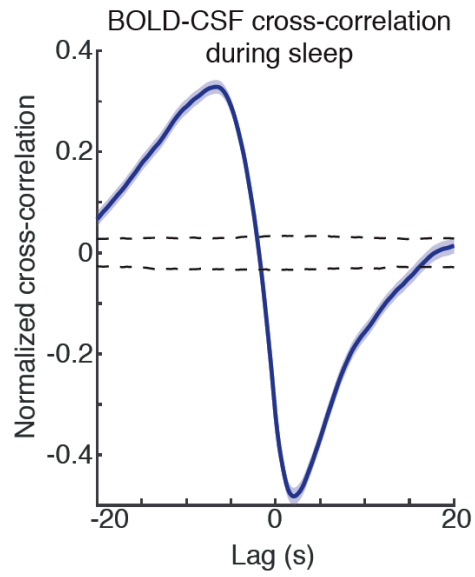


Fig. S3: Correlation between cortical BOLD and CSF signals. The BOLD signal is significantly anticorrelated with the CSF signal. Shaded region is standard error across segments and black dashed lines are 95% intervals of shuffled data.

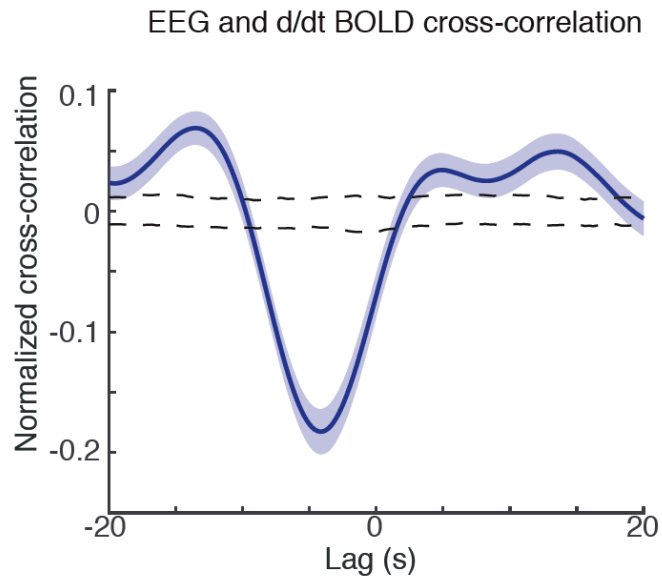


Fig. S4: EEG cross-correlation with the derivative of BOLD. The EEG slow-delta amplitude envelope is significantly correlated with the derivative of the BOLD signal, with a max R of -0.18 at a delay of -4.2 s. While a cross-correlation is not able to capture all aspects of the link between these two signals due to their nonlinear relationship, it nonetheless identifies that they are significantly coupled, with EEG preceding the BOLD derivative. Shaded region is standard error across segments and black dashed lines are 95% intervals of shuffled data.




Fig. S5: Best-fit impulse response prediction of CSF using the EEG data. A) The normalized cross-correlation of the CSF signal vs. the prediction of CSF using the EEG data demonstrates successful model prediction of CSF dynamics ($p < 0.001$, $n = 176$ segments, shuffling test). B) When considering only the sleep segments with the largest CSF waves (range $> 50\%$, $n = 19$ segments), the model prediction performance is increased, demonstrating that the EEG is more predictive of CSF when larger CSF waves are present ($p < 0.001$, shuffling test). Shaded region is standard error across segments. C) Cross-correlation of CSF data and model-based CSF prediction (using only EEG data) for each individual sleep segment. Segments are sorted by the amplitude of CSF peaks present in each segment (smaller index is smaller CSF waves). Sorting segments by CSF wave amplitude demonstrates that the slow-delta EEG model prediction performs well in individual segments, and that it explains the most variance in the segments with the largest CSF waves.

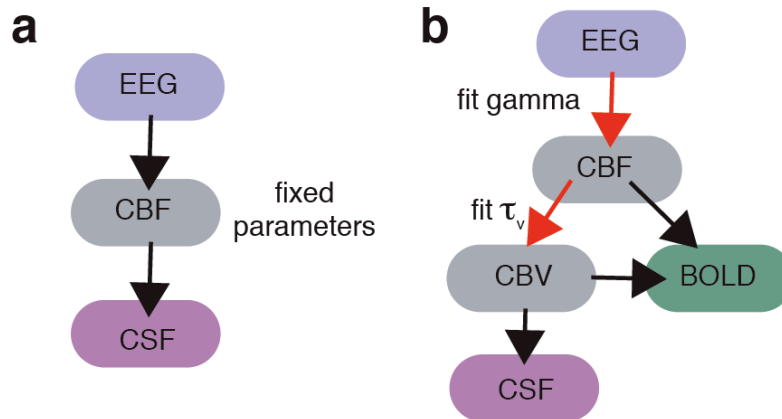


Fig. S6: Comparison of models implementing links between neural activity and CSF. Model variables include cerebral blood flow (CBF) and cerebral blood volume (CBV). A) We first implemented a model using previously established parameters without any additional fitting, and a simple assumption for CSF flow: we model the CSF signal as compensating changes in cerebral blood volume, and calculate it as the opposite of the change in CBF. We therefore tested whether the CSF dynamics were well predicted by the EEG-locked changes in CBF, using fixed parameters published in previous studies. This model performed as well as the best-fit impulse response, indicating that these biophysical mechanisms would produce the time-lagged dynamics we observe. B) We then allowed parameter settings to vary, finding two parameter settings that predicted the best-fit impulse response mapping EEG to the CSF signal, displayed in Fig. S7. In this case, we optimized model parameters to produce a CSF signal that was opposite to the d/dt of the CBV signal (i.e., positive CSF signal corresponding to negative d/dt CBV).

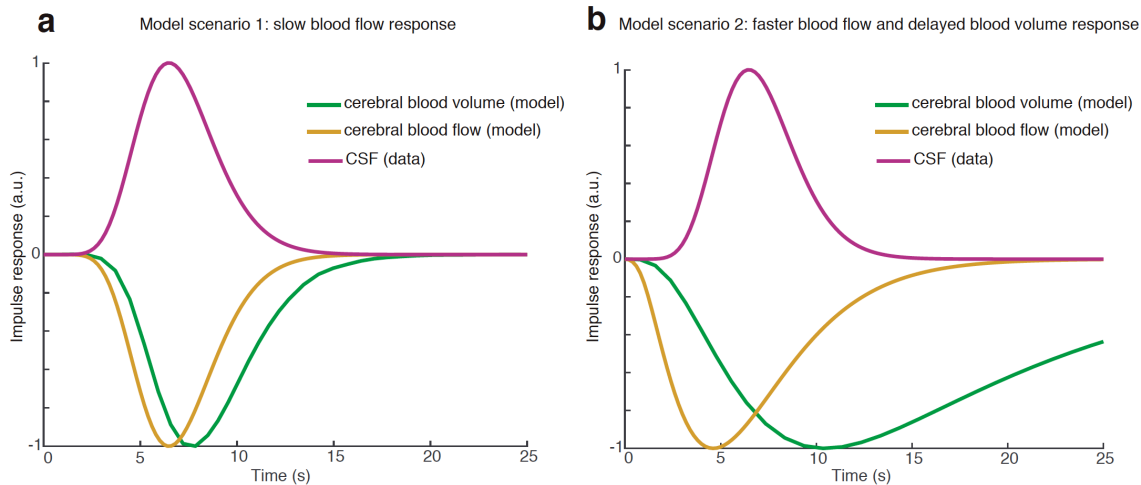


Fig. S7: The best fit impulse response for the CSF dynamics is consistent with a biophysical model of cerebral blood volume dynamics. The numerically fit CSF impulse response shows a similar, but slightly slower waveform, as compared to the fixed-parameter CBF impulse response. In this figure we demonstrate two model scenarios that are consistent with this CSF impulse response. First, this impulse response timing is within the established physiological range for CBF responses, so this result would be consistent with a slightly slower CBF coupling to spontaneous slow-delta EEG in sleep (as compared to task-induced fMRI measurements). Alternatively, it could reflect delayed changes in blood volume relative to blood flow (56). A) Implementation of delayed CBF scenario: the CBF impulse response timing matches the empirical CSF impulse response, corresponding to a slightly slower but still physiological blood flow response, as compared to the fixed-parameter model. B) Implementation of the delayed CBV scenario: plotting impulse response of the best-fit CBV impulse response (green) when holding the CBF impulse response constant at the fixed parameters, using a viscoelastic time constant of 30 s. This time constant provided the best fit within the physiological range of [0 30] s. The responses of modeled blood flow (yellow) and CSF data (purple) are also shown for comparison.

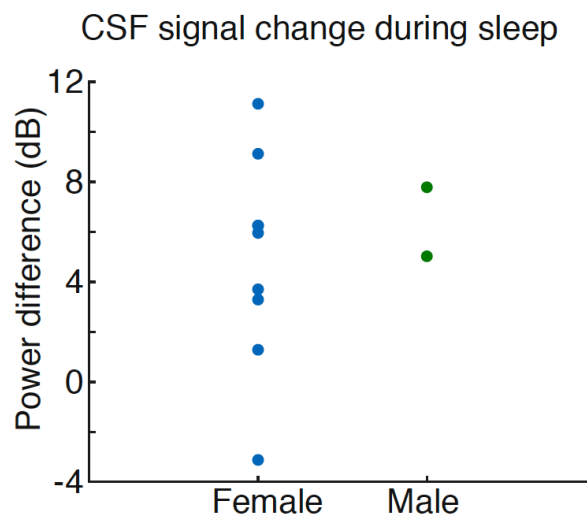


Fig. S8: CSF effects in male and female subjects. Our study enrolled more female subjects than male subjects, so we also examined whether the CSF oscillation increase during sleep was detected in both sexes. We split the CSF low-frequency (0-0.1 Hz) power change during sleep, reported in Fig. 1i, into male and female subjects. We observed that the two male subjects in the study exhibited CSF effects within the range of the female subjects. While this study was not designed to test for sex differences in CSF dynamics during sleep, future work could examine whether such differences exist.

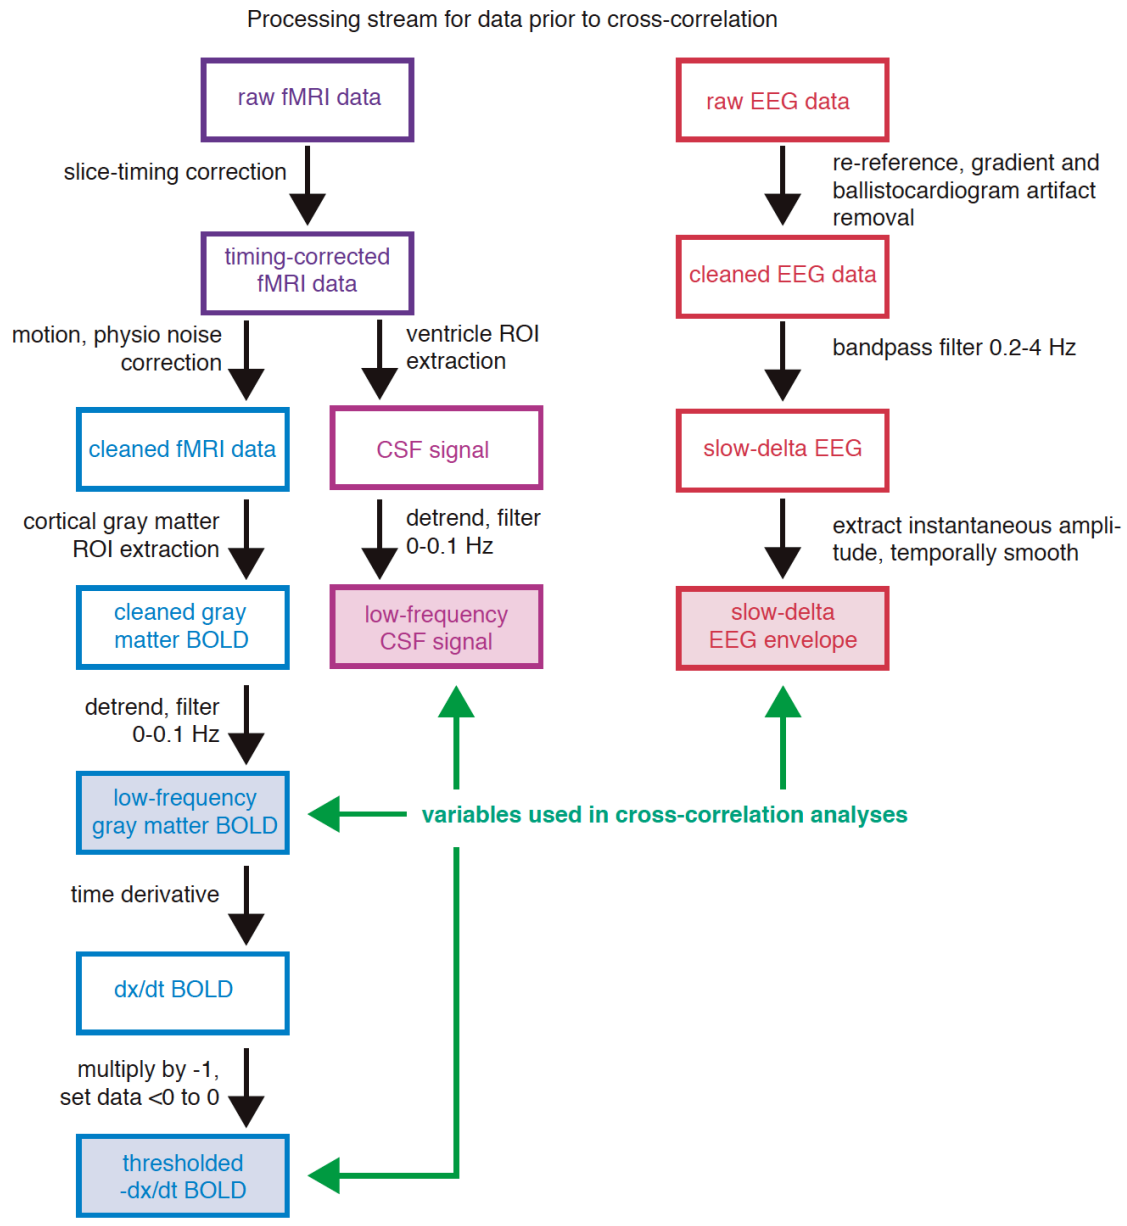


Fig. S9: Flowchart of preprocessing steps on each data type used in cross-correlation analyses. After preprocessing and filtering, pairs of signals were compared using cross-correlation analyses (Fig. 3d, Fig. S3, Fig. S4). For comparison, the cross-correlation without filtering is displayed in Fig. S10.

Cross-correlations without filtering fMRI data




Fig. S10: Cross-correlation results when omitting the 0-0.1 Hz bandpass filtering of BOLD and CSF signals. Repeating the cross-correlation without the bandpass filter steps shown in Fig. S9 yields qualitatively similar results. The correlations with the time derivative of cortical gray matter BOLD are slightly noisier when the filtering step is omitted, as the derivative calculation increases the contribution of high-frequency noise. A) Cross-correlation of CSF and derivative of BOLD has a max $R = -0.40$ at delay = -0.4 s. B) Cross-correlation of CSF and BOLD (original signal, not derivative) has a max $R = -0.48$ at delay = -2.0 s. C) Cross-correlation of EEG and derivative BOLD has a max $R = -0.10$ at delay = -3.8 s.