

High-amplitude co-fluctuations in cortical activity drive functional connectivity

Farnaz Zamani Esfahlani^{1,*}, Youngheun Jo^{1,*}, Joshua Faskowitz^{1,2}, Lisa Byrge¹, Daniel P. Kennedy¹⁻³, Olaf Sporns¹⁻⁴, and Richard F. Betzel^{1-4†}

¹*Department of Psychological and Brain Sciences,*

²*Program in Neuroscience,* ³*Cognitive Science Program,*

⁴*Network Science Institute, Indiana University, Bloomington, IN 47405*

MATERIALS AND METHODS

Datasets

We analyzed three separate datasets. Specifically, we focused on resting-state data from both The Human Connectome Project and Midnight Scan Club. These data were processed similarly, the details of which are described in this section. The third dataset, which has been analyzed elsewhere [1], includes both resting-state and movie-watching data from a cohort of 29 individuals. This dataset was processed separately using a different procedure and is described in its own section.

The Human Connectome Project (HCP) dataset [2] included resting state functional data (rsfMRI) from 100 unrelated adult subjects (54% female, mean age = 29.11 ± 3.67, age range = 22-36). The study was approved by the Washington University Institutional Review Board and informed consent was obtained from all subjects. Subjects underwent four 15 minute rsfMRI scans over a two day span. A full description of the imaging parameters and image preprocessing can be found in [3]. The rsfMRI data was acquired with a gradient-echo EPI sequence (run duration = 14:33 min, TR = 720 ms, TE = 33.1 ms, flip angle = 52°, 2 mm isotropic voxel resolution, multiband factor = 8) with eyes open and instructions to fixate on a cross. Images were collected on a 3T Siemens Connectome Skyra with a 32-channel head coil.

The Midnight Scan Club (MSC) dataset [4] included rsfMRI from 10 adults (50% female, mean age = 29.1 ± 3.3, age range = 24-34). The study was approved by the Washington University School of Medicine Human Studies Committee and Institutional Review Board and informed consent was obtained from all subjects. Subjects underwent 12 scanning sessions on separate days, each session beginning at midnight. 10 rsfMRI scans per subject were collected with a gradient-echo EPI sequence (run duration = 30 min, TR = 2200 ms, TE = 27 ms, flip angle = 90°, 4 mm isotropic voxel resolution) with eyes open and with eye tracking recording to monitor for prolonged eye closure (to assess drowsiness). Images were collected on a 3T Siemens Trio.

Image Preprocessing of HCP and MSC datasets

HCP Functional Preprocessing

Functional images in the HCP dataset were minimally preprocessed according to the description provided in [3]. Briefly, these data were corrected for gradient distortion, susceptibility distortion, and motion, and then aligned to a corresponding T1-weighted (T1w) image with one spline interpolation step. This volume was further corrected for intensity bias and normalized to a mean of 10000. This volume was then projected to the *32k_fs_LR* mesh, excluding outliers, and aligned to a common space using a multi-modal surface registration [5]. The resultant CIFTI file for each HCP subject used in this study followed the file naming pattern: `*_REST{1,2}_{L,R}_Atlas_MSMA11.dtseries.nii`.

MSC Functional Preprocessing

Functional images in the MSC dataset were preprocessed using *fMRIPrep* 1.3.2 [6], which is based on Nipype 1.1.9 [7]. The following description of *fMRIPrep*'s preprocessing is based on boilerplate distributed with the software covered by a “no rights reserved” (CC0) license. Internal operations of *fMRIPrep* use Nilearn 0.5.0 [8], ANTs 2.2.0, FreeSurfer 6.0.1, FSL 5.0.9, and AFNI v16.2.07. For more details about the pipeline, see the section corresponding to workflows in *fMRIPrep*'s documentation.

The T1-weighted (T1w) image was corrected for intensity non-uniformity with `N4BiasFieldCorrection` [9, 10], distributed with ANTs, and used as T1w-reference throughout the workflow. The T1w-reference was then skull-stripped with a Nipype implementation of the `antsBrainExtraction.sh` workflow, using NKI as the target template. Brain surfaces were reconstructed using `recon-all` [11], and the brain mask estimated previously was refined with a custom variation of the method to reconcile ANTs-derived and FreeSurfer-derived segmentations of the cortical gray-matter using `Mindboggle` [12]. Spatial normalization to the *ICBM 152 Nonlinear Asymmetrical template version 2009c* [13] was performed through nonlinear registration with `antsRegistration`, using brain-extracted versions of both T1w volume and template. Brain tissue segmentation of cerebrospinal

* These authors contributed equally

† rbetzel @ indiana.edu

fluid (CSF), white-matter (WM) and gray-matter (GM) was performed on the brain-extracted T1w using FSL’s `fast` [14].

Functional data was slice time corrected using AFNI’s `3dTshift` and motion corrected using FSL’s `mcflirt` [15]. *Fieldmap-less* distortion correction was performed by co-registering the functional image to the same-subject T1w image with intensity inverted [16] constrained with an average fieldmap template [17], implemented with `antsRegistration`. This was followed by co-registration to the corresponding T1w using boundary-based registration [18] with 9 degrees of freedom. Motion correcting transformations, field distortion correcting warp, BOLD-to-T1w transformation and T1w-to-template (MNI) warp were concatenated and applied in a single step using `antsApplyTransforms` using Lanczos interpolation.

Several confounding time-series were calculated based on this preprocessed BOLD: framewise displacement (FD), DVARS and three region-wise global signals. FD and DVARS are calculated for each functional run, both using their implementations in Nipype [19]. The three global signals are extracted within the CSF, the WM, and the whole-brain masks.

The resultant NIFTI file for each MSC subject used in this study followed the file naming pattern `*_space-T1w_desc-preproc_bold.nii.gz`.

Image Quality Control

All functional images in the HCP dataset were retained. The quality of functional images in the MSC were assessed using *fMRIPrep*’s visual reports and *MRIQC* 0.15.1 [20]. Data was visually inspected for whole brain field of view coverage, signal artifacts, and proper alignment to the corresponding anatomical image. Information about these image quality metrics can be found within *MRIQC*’s documentation [21].

Functional Networks Preprocessing

Parcellation Preprocessing

A functional parcellation designed to optimize both local gradient and global similarity measures of the fMRI signal [22] (*Schaefer200*) was used to define 200 areas on the cerebral cortex. These nodes are also mapped to the *Yeo* canonical functional networks [23]. For the HCP dataset, the *Schaefer200* is openly available in *32k_fs_LR* space as a CIFTI file. For the MSC and HBM datasets, a *Schaefer200* parcellation was obtained for each subject using a Gaussian classifier surface atlas [24] (trained on 100 unrelated HCP subjects) and FreeSurfer’s `mrisc_label` function. These tools utilize the surface registrations computed in the `recon-all` pipeline to transfer a group average atlas to subject space based

on individual surface curvature and sulcal patterns. This method rendered a T1w space volume for each subject. For use with functional data, the parcellation was resampled to 2mm T1w space.

Functional Network Preprocessing

The mean BOLD signal for each cortical node data was linearly detrended, band-pass filtered (0.008-0.08 Hz) [25], confound regressed and standardized using Nilearn’s `signal.clean`, which removes confounds orthogonally to the temporal filters [26]. The confound regression employed [27] included 6 motion estimates, time series of the mean CSF, mean WM, and mean global signal, the derivatives of these nine regressors, and the squares these 18 terms. Furthermore, a spike regressor was added for each fMRI frame exceeding a motion threshold (HCP = 0.25 mm root mean squared displacement, MSC = 0.5 mm framewise displacement). This confound strategy has been shown to be relatively effective option for reducing motion-related artifacts [25]. Following this preprocessing and nuisance regression, residual mean BOLD time series at each node was recovered.

We also analyzed a version of these data that did not include global signal regression as a pre-processing step. This procedure included 6 motion parameters and their temporal derivatives, the squares of these 12 terms, and 10 `aCompCor` components [25, 28].

Image Preprocessing of Indiana University Dataset

Demographics

We analyzed MRI data collected from $N_s = 29$ subjects (5 female, 24 male; 25 were right-handed). This cohort was male-dominant, as subjects were intended to serve as controls for a study in autism spectrum disorder, which is more common in men than women. At the time of their first scan, the average subject age was 24.9 ± 4.7 years [29].

MRI acquisition and processing

MRI images were acquired using a 3T whole-body MRI system (Magnetom Tim Trio, Siemens Medical Solutions, Natick, MA) with a 32-channel head receive array. Both raw and prescan-normalized images were acquired; raw images were used at all preprocessing stages and in all analyses unless specifically noted. During functional scans, T2*-weighted multiband echo planar imaging (EPI) data were acquired using the following parameters: TR/TE = 813/28 ms; 1200 vol; flip angle = 60°; 3.4 mm isotropic voxels; 42 slices acquired with interleaved order covering the whole brain; multi-band acceleration factor of 3. Preceding the first functional

scan, gradient-echo EPI images were acquired in opposite phase-encoding directions (10 images each with P-A and A-P phase encoding) with identical geometry to the EPI data (TR/TE = 1175/39.2 ms, flip angle = 60°) to be used to generate a fieldmap to correct EPI distortions, similar to the approach used by the Human Connectome Project [30]. High-resolution T1-weighted images of the whole brain (MPRAGE, 0.7 mm isotropic voxel size; TR/TE/TI = 2499/2.3/1000 ms) were acquired as anatomical references.

All functional data were processed according to an in-house pipeline using FEAT (v6.00) and MELODIC (v3.14) within FSL (v. 5.0.9; FMRIB’s Software Library, www.fmrib.ox.ac.uk/fsl), Advanced Normalization Tools (ANTs; v2.1.0) [31], and Matlab_R2014b. This pipeline was identical to the **GLM + MGTR** procedure described in [32].

In more detail, individual anatomical images were bias-corrected and skull-stripped using ANTs, and segmented into gray matter, white matter, and CSF partial volume estimates using FSL FAST. A midspace template was constructed using ANTs’ *buildtemplateparallel* and subsequently skull-stripped. Composite (affine and diffeomorphic) transforms warping each individual anatomical image to this midspace template, and warping the midspace template to the Montreal Neurological Institute MNI152 1mm reference template, were obtained using ANTs.

For each functional run, the first five volumes (≈ 4 seconds) were discarded to minimize magnetization equilibration effects. Framewise displacement traces for this raw (trimmed) data were computed using *fsl_motion_outliers*. Following [32, 33], we performed FIX followed by mean cortical signal regression. This procedure included rigid-body motion correction, fieldmap-based geometric distortion correction, and non-brain removal (but not slice-timing correction due to fast TR [30]). Preprocessing included weak highpass temporal filtering (>2000 s FWHM) to remove slow drifts [30] and no spatial smoothing. Off-resonance geometric distortions in EPI data were corrected using a fieldmap derived from two gradient-echo EPI images collected in opposite phase-encoding directions (posterior-anterior and anterior-posterior) using FSL *topup*.

We then used FSL-FIX [34] to regress out independent components classified as noise using a classifier trained on independent but similar data and validated on hand-classified functional runs. The residuals were regarded as “cleaned” data. Finally, we regressed out the mean cortical signal (mean BOLD signal across gray matter partial volume estimate obtained from FSL FAST). All analyses were carried out on these data, which were registered to subjects’ skull-stripped T1-weighted anatomical imaging using Boundary-Based Registration (BBR) with *epi_reg* within FSL. Subjects’ functional images were then transformed to the MNI152 reference in a single step, using ANTs to apply a concatenation of the affine transformation matrix with the composite (affine + diffeomor-

phic) transforms between a subject’s anatomical image, the midspace template, and the MNI152 reference. Prior to network analysis, we extracted mean regional time series from regions of interest defined as sub-divisions of the 17-system parcellation reported in [23] and used previously [35–37]. Wakefulness during movie and rest scans was monitored in real-time using an eye tracking camera (Eyelink 1000).

Naturalistic stimuli

All movies were obtained from Vimeo (<https://vimeo.com>). They were selected based on multiple criteria. First, to ensure that movies represented novel stimuli, we excluded any movie that had a wide theatrical release. Secondly, we excluded movies with potentially objectionable content including nudity, swearing, drug use, etc. Lastly, we excluded movies with intentionally startling events that could lead to excessive in-scanner movement.

Each movie lasted approximately 1 to 5 minutes. Each movie scan comprised between four and six movies with genres that included documentaries, dramas, comedies, sports, mystery, and adventure. See Table. S1 for more details.

Co-fluctuation time series

Constructing networks from fMRI data (or any neural time series data) requires estimating the statistical dependency between every pair of time series. The magnitude of that dependency is usually interpreted as a measure of how strongly (or weakly) those voxels are functionally connected to each other. By far the most common measure of statistic dependence is the Pearson correlation coefficient. Let $\mathbf{x}_i = [x_i(1), \dots, x_i(T)]$ and $\mathbf{x}_j = [x_j(1), \dots, x_j(T)]$ be the time series recorded from voxels or parcels i and j , respectively. We can calculate the correlation of i and j by first z-scoring each time series, such that $\mathbf{z}_i = \frac{\mathbf{x}_i - \mu_i}{\sigma_i}$, where $\mu_i = \frac{1}{T} \sum_t x_i(t)$ and $\sigma_i = \frac{1}{T-1} \sum_t [x_i(t) - \mu_i]$ are the time-averaged mean and standard deviation. Then, the correlation of i with j can be calculated as: $r_{ij} = \frac{1}{T-1} \sum_t [z_i(t) \cdot z_j(t)]$. Repeating this procedure for all pairs of parcels results in a node-by-node correlation matrix, i.e. an estimate of FC. If there are N nodes, this matrix has dimensions $[N \times N]$.

To estimate *edge*-centric networks, we need to modify the above approach in one small but crucial way. Suppose we have two z-scored parcel time series, \mathbf{z}_i and \mathbf{z}_j . To estimate their correlation we calculate the mean their element-wise product (not exactly the average, because we divide by $T-1$ rather than T). Suppose, instead, that we never calculate the mean and simply stop after calculating the element-wise product. This operation would result in a vector of length T whose elements encode the moment-by-moment co-fluctuations magnitude of parcels

i and j . For instance, suppose at time t , parcels i and j simultaneously increased their activity relative to baseline. These increases are encoded in \mathbf{z}_i and \mathbf{z}_j as positive entries in the t th position, so their product is also positive. The same would be true if i and j *decreased* their activity simultaneously (because the product of negatives is a positive). On the other hand, if i increased while j decreased (or *vice versa*), this would manifest as a negative entry. Similarly, if either i or j increased or decreased while the activity of the other was close to baseline, the corresponding entry would be close to zero.

Accordingly, the vector resulting from the element-wise product of \mathbf{z}_i and \mathbf{z}_j can be viewed as encoding the magnitude of moment-to-moment co-fluctuations between i and j . An analogous vector can easily be calculated for every pair of parcels (network nodes), resulting in a set of co-fluctuation (edge) time series. With N parcels, this results in $\frac{N(N-1)}{2}$ pairs, each of length T .

Modularity maximization

Modularity maximization is a heuristic for detecting communities in networks [38]. Intuitively, it attempts to decompose a network into non-overlapping sub-networks such that the observed density of connections within sub-networks maximally exceeds what would be expected by chance, where chance is determined by the user. The actual process of detecting communities is accomplished by choosing community assignments that maximize a modularity quality function, Q , defined as:

$$Q = \sum_{ij} B_{ij} \delta(g_i, g_j) \quad (1)$$

where $B_{ij} = A_{ij} - P_{ij}$ is the $\{i, j\}$ element of the modularity matrix, which represents the observed weight of the connection between nodes i and j minus the expected weight. The variable g_i is the community assignment of node i and $\delta(x, y)$ is the Kronecker delta function, whose value is 1 when $g_i = g_j$ and 0 otherwise. The modularity, Q , is effectively a sum over all edges that fall within communities and is optimized when the the observed weights of connections is maximally greater than the expected. In general, larger values of Q are thought to reflect superior community partitions.

Signed and correlation matrices

In this manuscript, we used the following variant of modularity, q^* , which has been shown to be especially well-suited for use with correlation matrices [39]:

$$q^* = q^+ + \frac{v^-}{v^+v^-} q^- \quad (2)$$

where $q^\pm = \frac{1}{v^\pm} \sum_{ij} (r_{ij}^\pm - \frac{k_i^\pm k_j^\pm}{v^\pm}) \delta(g_i, g_j)$. In this expression, r_{ij}^\pm represents either the positive or negative elements of the correlation matrix, $k_i^\pm = \sum_j r_{ij}^\pm$, and $v^\pm = \sum_i k_i^\pm$.

Brain-behavior analysis using co-fluctuation data

In the main text, we speculated that, because high-amplitude co-fluctuations encode subject-specific information (see Fig. 4 in the main text), we might refine brain-behavior relationships by focusing on FC estimated using those frames only. Here, we tested this hypothesis directly, reconstructing FC using the top and bottom 5% frames (ordered by co-fluctuation amplitude). For each subject, we combined reconstructed FC across all four scan sessions (REST1_LR, REST1_RL, REST2_LR, and REST2_RL). Then, following [40], we regressed out of motion (average framewise displacement), number of censored frames, subject height, subject weight, systolic and diastolic blood pressure, total brain volume, and total intracranial volume (as well as the demeaned squares of these measures).

Again, following [40], we extracted 158 behavioral, demographic, and trait variables. The total number of variables was reduced to 131 after removing variables with missing entries for any subject. We then standardized (z-scored) each variable across subjects and performed a principal component analysis (PCA). The result was a series of latent variables that explained “modes” or “patterns” of behavioral and demographic variation.

Here, we focus on PC1 (it explains 20.3% of behavioral variance). The items with the strongest positive and negative loadings suggest that this component indexes general psychiatric and life function. The largest positive loadings are composite scales from the Achenbach Adult Self-Report (ASR; [41]) and sum problem counts over different clinical syndrome scales. The largest negative loadings, on the other hand, index positive affect, social, and emotional support. Accordingly, we refer to PC1 as a “life and psychiatric well-being” index (Fig. S15a).

For PC1 (and the next nine components) we computed its correlation with connection weights. We performed this analysis separately for FC reconstructed from high- and low-amplitude frames, resulting in different correlation patterns across all edges (Fig. S15b,c). We evaluated correlation maps statistically at the level of brain systems Fig. S15d) – computing the average within- and between-system correlation and comparing the average values against a null distribution generated by randomly permuting nodes’ system labels.

In the case of FC reconstructed from high-amplitude frames, we found significant system-level correlations involving many brain systems (Fig. S15e; multiple comparisons corrected for by fixing the false discovery rate 5%; $p_{adjusted} = 2.01 \times 10^{-4}$). To visualize these differences in anatomical space, we masked the full correlation map

to retain only those connections whose endpoints that were significant at a system level and were significant at a connection-level ($p < 0.05$; uncorrected). We note that this analysis is for the sake of visualization, only. For PC1, we found a constellation of strong, positive correlations of PC1 with edges linking striate cortex (visual system) to regions within somatomotor cortex (Fig. S15f). Similarly, we found strong negative correlations of PC1 with edges linking parahippocampal cortex (DMN) with somatomotor cortex (Fig. S15f).

More generally, we found that correlation maps estimated from the high-amplitude frames resulted system-level correlations of greater magnitude compared to low-amplitude frames (Fig. S15g). Notably, this effect was consistent across and statistically significant the first ten PCs (paired-sample t-tests; $p < 0.05$).

These observations suggest that FC computed using high-amplitude frames leads to stronger brain-behavior correlations at the system level, broadly supporting our hypothesis. If nothing else, our observations suggest that FC computed from high- and low-amplitude frames can yield dissimilar and distinct patterns of brain-behavior correlations, motivating further exploration.

AUTHOR CONTRIBUTIONS

RFB, JF, and OS conceived of study. JF and LB processed data. FZE, YJ, JF, LB, DPK, OS, and RB carried out all analyses, wrote, edited, and revised the submitted manuscript.

ACKNOWLEDGMENTS

RFB and FZE acknowledge support from Indiana University Office of the Vice President for Research Emerging Area of Research Initiative, Learning: Brains, Machines and Children. This work was supported by the NIH (R01MH110630 and R00MH094409 to DPK and T32HD007475 Postdoctoral Traineeship to LB).

DATA AVAILABILITY

All imaging data come from publicly-available, open-access repositories. Human connectome project data can be accessed *via* <https://db.humanconnectome.org/app/template/Login.vm> after signing a data use agreement. Midnight scan club data can be accessed *via* OpenfMRI at <https://www.openfmri.org/dataset/ds000224/>. The Indiana University dataset is available upon reasonable request.

CODE AVAILABILITY

All processing and analysis code is available upon reasonable request.

-
- [1] R. F. Betzel, L. Byrge, F. Z. Esfahlani, and D. P. Kennedy, *NeuroImage*, 116687 (2020).
- [2] D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. Behrens, E. Yacoub, K. Ugurbil, W.-M. H. Consortium, *et al.*, *Neuroimage* **80**, 62 (2013).
- [3] M. F. Glasser, S. N. Sotiropoulos, J. A. Wilson, T. S. Coalson, B. Fischl, J. L. Andersson, J. Xu, S. Jbabdi, M. Webster, J. R. Polimeni, *et al.*, *Neuroimage* **80**, 105 (2013).
- [4] E. M. Gordon, T. O. Laumann, A. W. Gilmore, D. J. Newbold, D. J. Greene, J. J. Berg, M. Ortega, C. Hoyt-Drazen, C. Gratton, H. Sun, *et al.*, *Neuron* **95**, 791 (2017).
- [5] E. C. Robinson, S. Jbabdi, M. F. Glasser, J. Andersson, G. C. Burgess, M. P. Harms, S. M. Smith, D. C. Van Essen, and M. Jenkinson, *Neuroimage* **100**, 414 (2014).
- [6] O. Esteban, C. Markiewicz, R. W. Blair, C. Moodie, A. I. Isik, A. Erramuzpe Aliaga, J. Kent, M. Goncalves, E. DuPre, M. Snyder, H. Oya, S. Ghosh, J. Wright, J. Durnez, R. Poldrack, and K. J. Gorgolewski, *Nature Methods* (2018), 10.1038/s41592-018-0235-4.
- [7] K. Gorgolewski, C. D. Burns, C. Madison, D. Clark, Y. O. Halchenko, M. L. Waskom, and S. Ghosh, *Frontiers in Neuroinformatics* **5**, 13 (2011).
- [8] A. Abraham, F. Pedregosa, M. Eickenberg, P. Gervais, A. Mueller, J. Kossaifi, A. Gramfort, B. Thirion, and G. Varoquaux, *Frontiers in Neuroinformatics* **8** (2014), 10.3389/fninf.2014.00014.
- [9] N. J. Tustison, B. B. Avants, P. A. Cook, Y. Zheng, A. Egan, P. A. Yushkevich, and J. C. Gee, *IEEE Transactions on Medical Imaging* **29**, 1310 (2010).
- [10] B. Avants, C. Epstein, M. Grossman, and J. Gee, *Medical Image Analysis* **12**, 26 (2008).
- [11] A. M. Dale, B. Fischl, and M. I. Sereno, *NeuroImage* **9**, 179 (1999).
- [12] A. Klein, S. S. Ghosh, F. S. Bao, J. Giard, Y. Häme, E. Stavsky, N. Lee, B. Rossa, M. Reuter, E. C. Neto, and A. Keshavan, *PLOS Computational Biology* **13**, e1005350 (2017).
- [13] V. Fonov, A. Evans, R. McKinstry, C. Almlí, and D. Collins, *NeuroImage* **47**, **Supplement 1**, S102 (2009).
- [14] Y. Zhang, M. Brady, and S. Smith, *IEEE Transactions on Medical Imaging* **20**, 45 (2001).
- [15] M. Jenkinson, P. Bannister, M. Brady, and S. Smith, *NeuroImage* **17**, 825 (2002).
- [16] S. Wang, D. J. Peterson, J. C. Gatenby, W. Li, T. J. Grabowski, and T. M. Madhyastha, *Frontiers in Neu-*

- roinformatics **11** (2017), 10.3389/fninf.2017.00017.
- [17] J. M. Treiber, N. S. White, T. C. Steed, H. Bartsch, D. Holland, N. Farid, C. R. McDonald, B. S. Carter, A. M. Dale, and C. C. Chen, PLOS ONE **11**, e0152472 (2016).
- [18] D. N. Greve and B. Fischl, NeuroImage **48**, 63 (2009).
- [19] J. D. Power, A. Mitra, T. O. Laumann, A. Z. Snyder, B. L. Schlaggar, and S. E. Petersen, NeuroImage **84**, 320 (2014).
- [20] O. Esteban, D. Birman, M. Schaer, O. O. Koyejo, R. A. Poldrack, and K. J. Gorgolewski, PloS one **12**, e0184661 (2017).
- [21] “Mriqc: Advancing the automatic prediction of image quality in mri from unseen sites,” (2019).
- [22] A. Schaefer, R. Kong, E. M. Gordon, T. O. Laumann, X.-N. Zuo, A. J. Holmes, S. B. Eickhoff, and B. T. Yeo, Cerebral Cortex , 1 (2017).
- [23] B. Thomas Yeo, F. M. Krienen, J. Sepulcre, M. R. Sabuncu, D. Lashkari, M. Hollinshead, J. L. Roffman, J. W. Smoller, L. Zöllei, J. R. Polimeni, *et al.*, Journal of neurophysiology **106**, 1125 (2011).
- [24] B. Fischl, A. Van Der Kouwe, C. Destrieux, E. Halgren, F. Ségonne, D. H. Salat, E. Busa, L. J. Seidman, J. Goldstein, D. Kennedy, *et al.*, Cerebral cortex **14**, 11 (2004).
- [25] L. Parkes, B. Fulcher, M. Yücel, and A. Fornito, Neuroimage **171**, 415 (2018).
- [26] M. A. Lindquist, S. Geuter, T. D. Wager, and B. S. Caffo, Human brain mapping **40**, 2358 (2019).
- [27] T. D. Satterthwaite, M. A. Elliott, R. T. Gerraty, K. Ruparel, J. Loughhead, M. E. Calkins, S. B. Eickhoff, H. Hakonarson, R. C. Gur, R. E. Gur, *et al.*, Neuroimage **64**, 240 (2013).
- [28] Y. Behzadi, K. Restom, J. Liau, and T. T. Liu, Neuroimage **37**, 90 (2007).
- [29] L. Byrge and D. P. Kennedy, Human Brain Mapping (2020).
- [30] S. M. Smith, C. F. Beckmann, J. Andersson, E. J. Auerbach, J. Bijsterbosch, G. Douaud, E. Duff, D. A. Feinberg, L. Griffanti, M. P. Harms, *et al.*, Neuroimage **80**, 144 (2013).
- [31] B. B. Avants, N. J. Tustison, G. Song, P. A. Cook, A. Klein, and J. C. Gee, Neuroimage **54**, 2033 (2011).
- [32] L. Byrge and D. P. Kennedy, NeuroImage **171**, 376 (2018).
- [33] G. C. Burgess, S. Kandala, D. Nolan, T. O. Laumann, J. D. Power, B. Adeyemo, M. P. Harms, S. E. Petersen, and D. M. Barch, Brain connectivity **6**, 669 (2016).
- [34] G. Salimi-Khorshidi, G. Douaud, C. F. Beckmann, M. F. Glasser, L. Griffanti, and S. M. Smith, Neuroimage **90**, 449 (2014).
- [35] R. F. Betzel, L. Byrge, Y. He, J. Goñi, X.-N. Zuo, and O. Sporns, Neuroimage **102**, 345 (2014).
- [36] L. Byrge, J. Dubois, J. M. Tyszka, R. Adolphs, and D. P. Kennedy, Journal of Neuroscience **35**, 5837 (2015).
- [37] R. F. Betzel, B. Mišić, Y. He, J. Rumschlag, X.-N. Zuo, and O. Sporns, arXiv preprint arXiv:1510.08045 (2015).
- [38] M. E. Newman and M. Girvan, Physical review E **69**, 026113 (2004).
- [39] M. Rubinov and O. Sporns, Neuroimage **56**, 2068 (2011).
- [40] S. M. Smith, T. E. Nichols, D. Vidaurre, A. M. Winkler, T. E. Behrens, M. F. Glasser, K. Ugurbil, D. M. Barch, D. C. Van Essen, and K. L. Miller, Nature neuroscience **18**, 1565 (2015).
- [41] T. M. Achenbach, *The Achenbach system of empirically based assessment (ASEBA): Development, findings, theory, and applications* (University of Vermont, Research Center for Children, Youth, & Families, 2009).
- [42] D. Prichard and J. Theiler, Physical review letters **73**, 951 (1994).
- [43] J. D. Power, K. A. Barnes, A. Z. Snyder, B. L. Schlaggar, and S. E. Petersen, Neuroimage **59**, 2142 (2012).

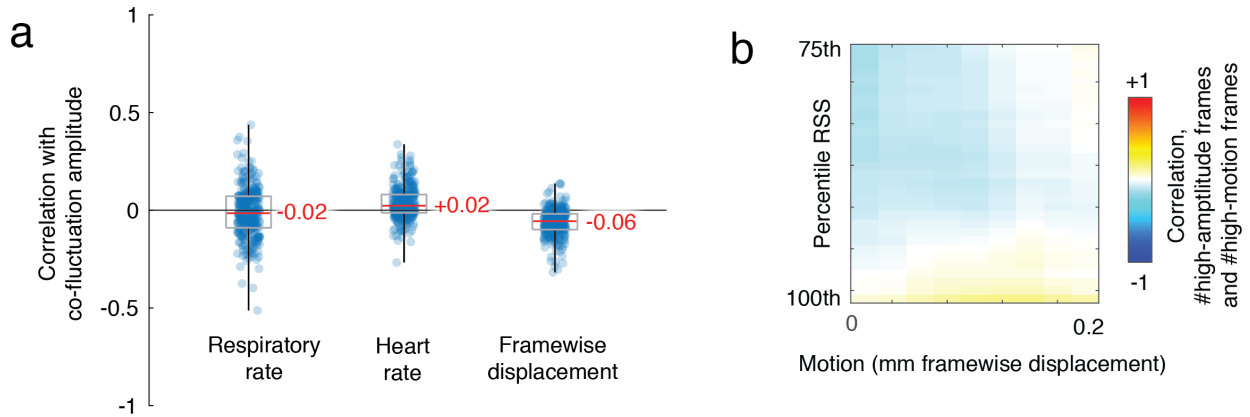


FIG. S1. **Comparison of co-fluctuation amplitude with confounding variables.** (a) In the main text we calculated the magnitude of co-fluctuation at every frame. A concern is that variation in this measure could be attributed to physiological and motion-related variables of non-neural origins. To address this concern, we calculated the correlation of co-fluctuation amplitude with three variables: respiratory and heart rate data, as well as in-scanner head motion (relative root mean square error framewise displacement). In-scanner motion is already sampled at the same frequency as the BOLD acquisition; for the two physiological variables, we estimated their instantaneous rates by smoothing (low-pass filter), and computing the time between successive peaks of the periodic signal. Using linear interpolation, we then resampled these values to be at each TR. We performed this procedure for every subject and scan session in the HCP dataset and computed their correlation with the co-fluctuation amplitude. The distributions of correlation coefficients tended to be weak and close to zero (mean values shown in red text), suggesting that co-fluctuation amplitude is not obviously related to physiological or motion-related variables. (b) We also explored the relationship of motion and RSS (co-fluctuation amplitude) across individuals. We defined two thresholds: a motion threshold for classifying frames as “high motion” or not, and an RSS threshold for classifying frames as “high-amplitude” or not. We systematically varied these thresholds (motion from 0 - 0.2 mm; RSS from the 75th to the 100th percentile, defined using RSS data pooled from all subjects’ scans). For a given motion and co-fluctuation threshold, we calculated the number of high-motion and high-amplitude frames for each subject and scan and computed the correlation between these two variables. In general, we found weak correlations ($r = -0.03 \pm 0.06$ on average; $r = 0.16$ maximum absolute).

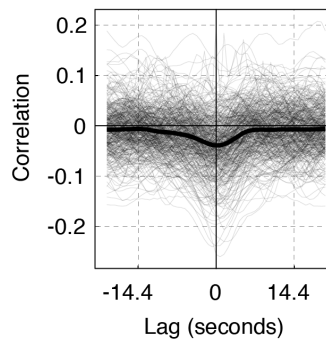


FIG. S2. **Lagged relationships between motion and RSS time series.** We generated, for each subject and scan, the RSS time series (root sum square of inter-regional co-fluctuations). We then computed the lagged cross correlation between RSS and framewise displacement (max lag of ± 20 seconds). In general, we found the correlation was weak, negative, and peaked at a lag of 0. These analyses suggest that at a single-subject level, framewise displacement (a common measure of in-scanner motion) and RSS are not strongly associated with one another. If anything, these observations suggest that high-amplitude co-fluctuations tend to occur when motion is low (albeit the relationship is weak).

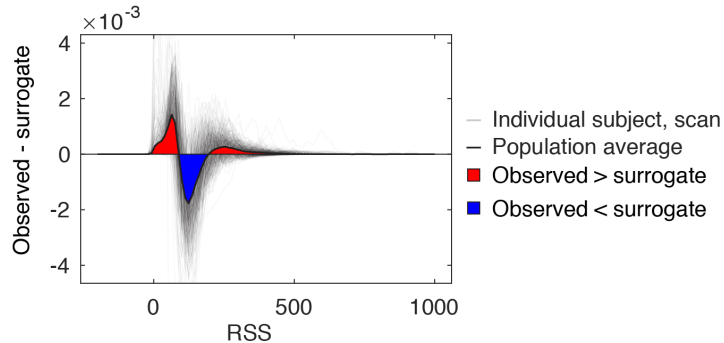


FIG. S3. **Comparison of co-fluctuation amplitude from observed data and from phase-randomized surrogates.**

In the main text we calculate co-fluctuation amplitude as root sum square of edge co-fluctuations at each moment in time. Here, we compare these observed amplitudes with those estimated from phase-randomized surrogate time series. The phase randomization procedure has been described in detail elsewhere [42]. Briefly, this procedure entails taking the discrete Fourier transform of each regional BOLD time series, adding random phase at each frequency bin, and taking the inverse Fourier transform, generating a surrogate time series for that region with same power spectrum but random phase properties. We repeat this procedure for all $N = 200$ regions, uniformly adding the same random phase to each frequency bin, thus ensuring that the surrogate time series preserve the correlation structure of the original data. We repeated this procedure 100 times per scan and subject and found that, as expected, the distribution of co-fluctuation amplitude for both the observed and synthetic data was broad and included a heavy tail. However, upon more detailed examination, we found that the observed data exhibited a greater proportion of very high- and low-amplitude frames. These observations suggest that truly high-amplitude frames are rare, even when compared directly against surrogate data that preserved the spectral properties of regional time series and the multivariate correlation structure of observed fMRI BOLD time series.

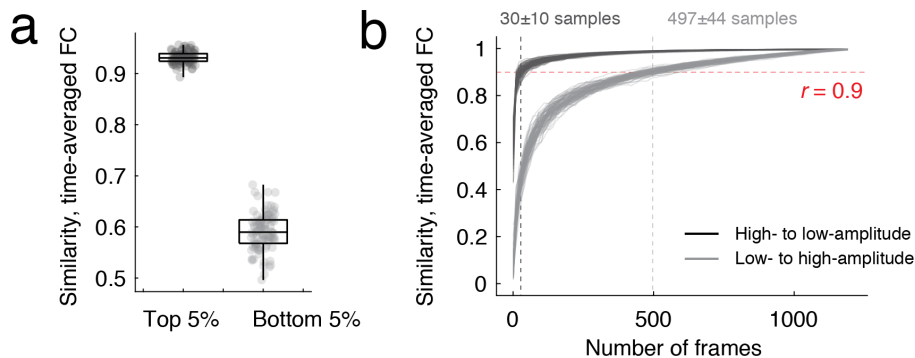


FIG. S4. **Reconstructing the correlation structure of synthetic data using high- and low-amplitude frames.**

Here, we generate 100 sets of synthetic time series from a fixed covariance structure (defined as the sample correlation matrix estimated from recorded fMRI data). We repeat the primary analysis from the main text using these data. Namely, we estimate co-fluctuation (edge) time series for each sample and use high- and low-amplitude frames to reconstruct the ground-truth covariance matrix used to generate the time series data. (a) As in the main text, we exhibit disparities in the similarity of reconstructed matrices using the top and bottom 5% frames (in descending order of RSS amplitude). Specifically, we find that high-amplitude frames yield better estimates than low-amplitude frames ($p < 0.05$; t-test). (b) We generalized this analysis, sampling different proportions of frames beginning with either the highest- or lowest-amplitude and reconstructing the correlation matrix using only those frames. We found that the similarity of the matrix reconstructed from high-amplitude frames increased quickly, requiring $\approx 30 \pm 10$ frames to achieve a similarity of $r = 0.9$. Beginning with, low-amplitude frames, we found that $\approx 497 \pm 44$ samples were required to achieve a comparable level of similarity.

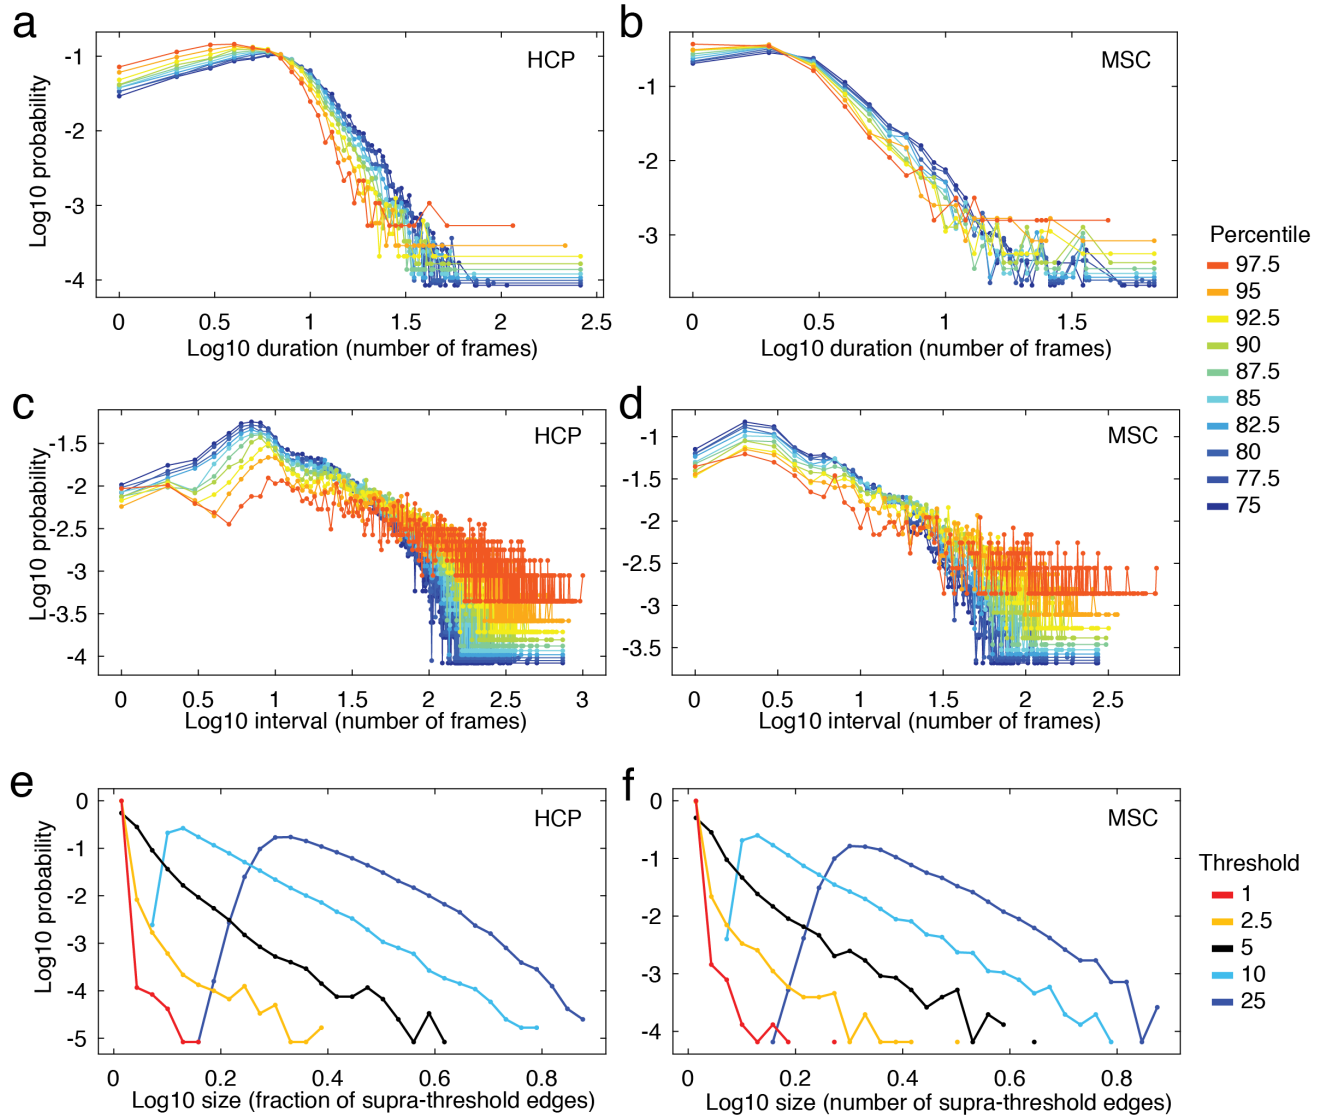


FIG. S5. **Duration and intervals of high-amplitude co-fluctuations.** For each scan session, we calculated the co-fluctuation amplitude at every frame. We imposed percentile-based thresholds on these data (percentiles calculated based on pooled data from all subjects and all scan sessions). Thresholding the co-fluctuation amplitude time series results in a binary classification of time points as either “high-” or “low-amplitude”. From these observations, we calculated two quantities: “duration” as the number of consecutive frames classified as high-amplitude and “intervals” as the number of frames between successive high-amplitude frames. We repeated this analysis for both the HCP and MSC datasets. Panels *a* and *b* show durations for HCP and MSC datasets, respectively. Note that the distribution is broad and includes a heavy tail, indicating a lack of periodicity. Panels *c* and *d* depict intervals for HCP and MSC datasets. Additionally, we assessed the size of high-amplitude co-fluctuations, as measured by the fraction of all edges whose co-fluctuation amplitude at a given frame exceeded some threshold. Here, we identified high-amplitude frames as time points at which the co-fluctuation amplitude was in the top 1%, 2.5%, 5%, 10%, and 25% (thresholds are indicated by different colors in each plot). Then, for each time point classified as high-amplitude, we calculated the fraction of all edges whose absolute co-fluctuation amplitude exceed the 75th percentile. We performed this procedure using both HCP (*e*) and MSC (*f*) data and found that size follow a broad and heavy-tailed distribution, suggesting that high-amplitude frames follow no characteristic scale of description.




FIG. S6. **Similarity of time-averaged FC with FC estimated using fewer frames.** In the main text, we showed that rsFC, when estimated using the top 5% of frames (ordered by co-fluctuation magnitude) resulted in a connectivity matrix that much more similar to time-averaged FC than the matrix generated using the bottom 5% of frames. Here, we show that this relationship persists irrespective of percentile. To do this for a given subject and scan session, we ordered frames according to co-fluctuation magnitude from greatest to least. Then, we extracted the top and bottom k frames (varying k from 3 to to T , where T is the total number of frames in the scan session), estimating FC using those k frames, and calculating the similarity with time-averaged FC. This procedure results in a similarity value at every k for both the top and bottom frames. We repeated this analysis for all 100 subjects in the HCP dataset. We find that across the full range of k , FC estimated using frames corresponding to high-amplitude co-fluctuations was always more similar to the time-averaged FC than those estimated using low-amplitude frames.

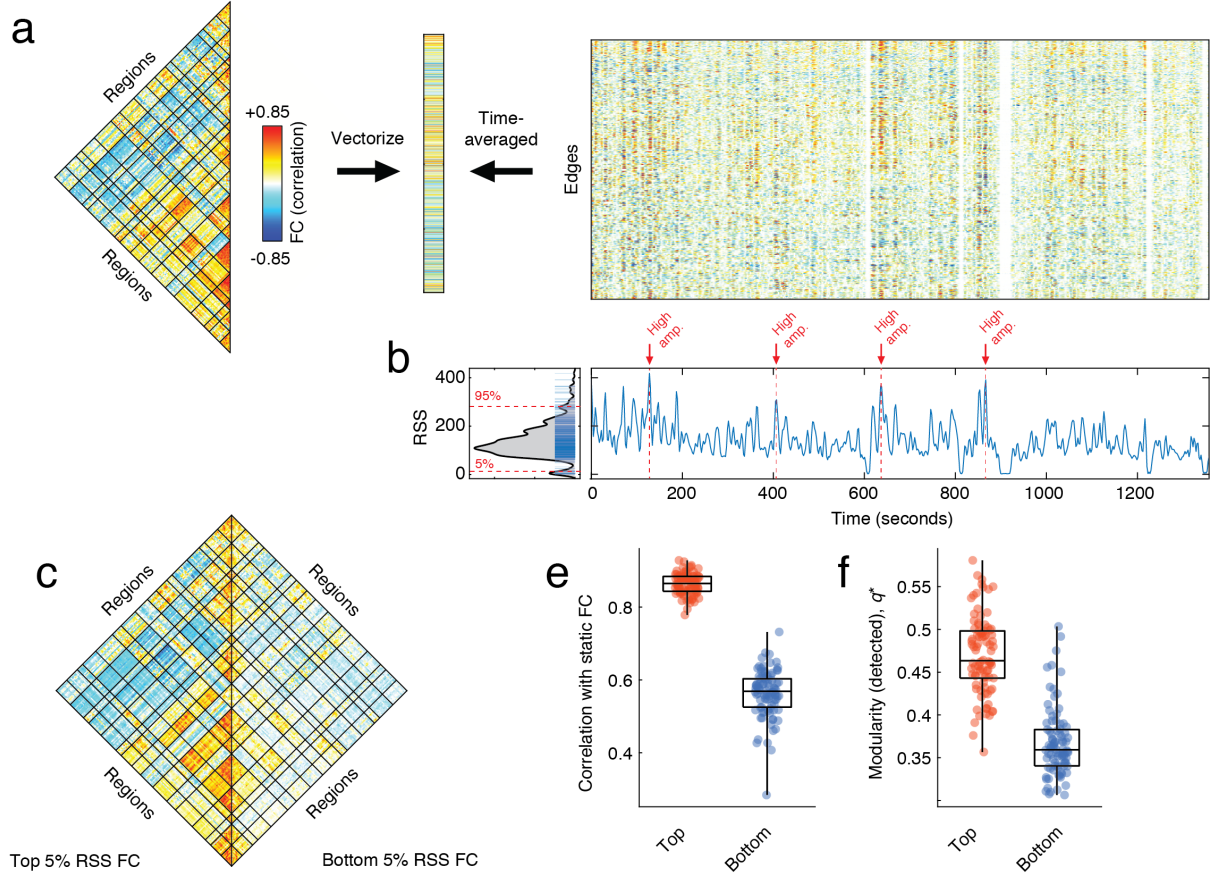


FIG. S7. **Co-fluctuation time series reveal burst structure of resting-state functional connectivity for MSC dataset.** (a) We use a “temporal unwrapping” of the Pearson correlation to generate co-fluctuation time series for every pair of brain regions (edges). The elements of the co-fluctuation time series are the element-wise products of z-scored regional BOLD time series that, when averaged across time, yield vectors that are exactly equal to the Pearson correlation coefficient and can be rearranged to create an resting-state functional connectivity matrix. (b) We find that the co-fluctuation time series contains moments in time where many edges collectively co-fluctuate. We can identify these moments by calculating the root sum square across all co-fluctuation time series and plotting this value as a function of time. In panel *b* we label high- and low-amplitude frames. The distribution of edge co-fluctuation amplitude is heavy tailed. We wanted to assess the contribution of high- and low-amplitude frames to the overall pattern of functional connectivity. To do this, we extracted the top and bottom 5% of all time points (ordered by co-fluctuation amplitude) and estimated functional connectivity from those points alone. (c) Average functional connectivity across 100 subjects using top 5% (left) and bottom 5% (right). (d) In general, the networks estimated using the top 5% of time points were much more similar to traditional functional connectivity than those estimated using the bottom 5% of time points. (e) We performed a similar comparison of network modularity using networks reconstructed using top and bottom 5% frames.

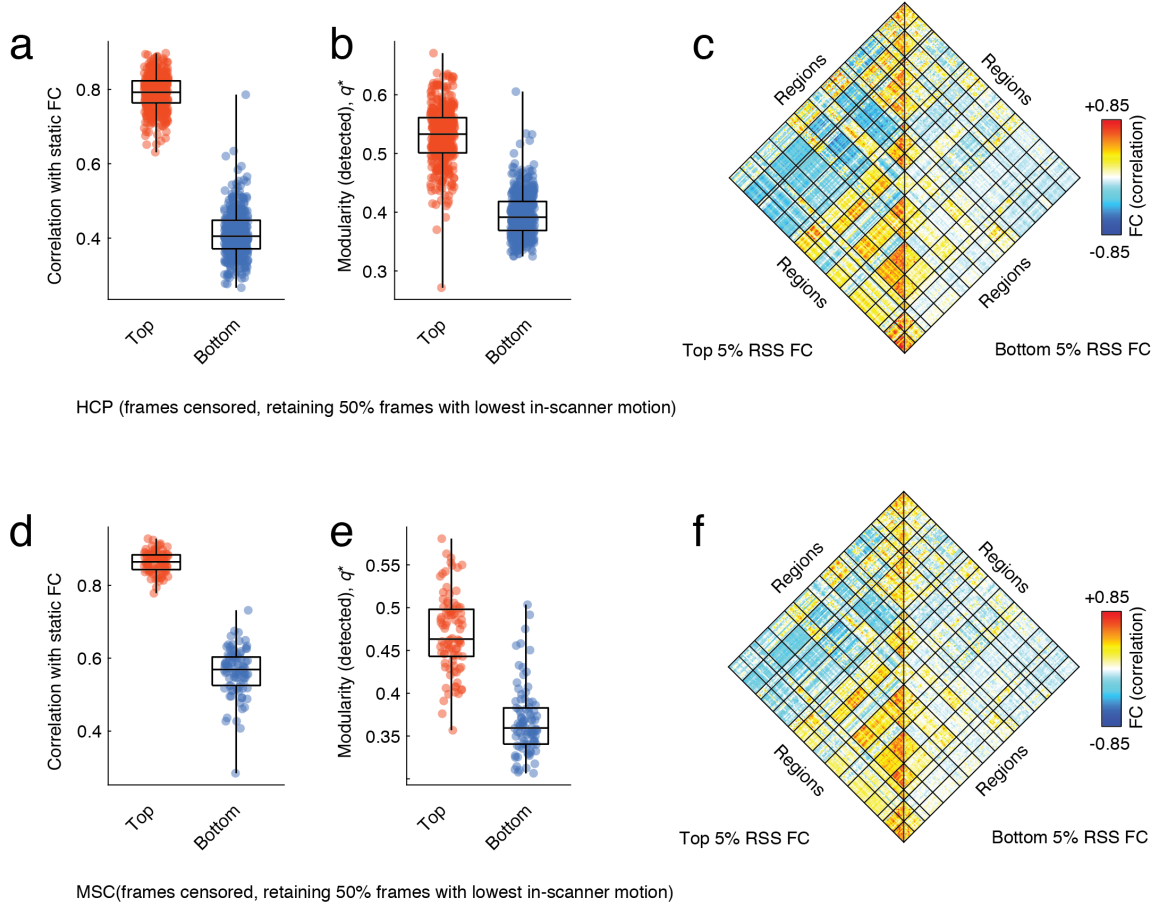


FIG. S8. **Effect of frame-censoring.** In the main text, we demonstrated that FC estimated using frames corresponding to high-amplitude co-fluctuations was more similar to time-averaged FC than the FC estimated using low-amplitude co-fluctuation frames. Here, we perform an identical analysis using only the bottom 50% frames in terms of in-scanner motion [43]. Note that this procedure results in time series that include exactly half of the original frames. (a) Correlation of FC estimated using top and bottom 5% of frames, ordered by co-fluctuation amplitude. As in the main text, the top 5% are more correlated with time-averaged FC than the bottom 5%. (b) Modularity of FC estimated using only the top and bottom 5% of frames. As in the main text, the top 5% are more modular than the bottom. (c) Group-averaged FC matrices estimated using the top 5% of frames (left) and the bottom 5% of frames (right). Panels a, b, and c depict results using HCP data, while d, e, and f depict analogous results using data from the MSC.




FIG. S9. **Alternative strategy for estimating FC from a limited number of frames.** In the main text, we estimated FC from the top and bottom 5% of frames by extracting fMRI BOLD activity from computing the correlation structure. An alternative strategy for estimating FC is to simply average co-fluctuation time series over the top/bottom frames, ordered by co-fluctuation amplitude. Here, we perform this analysis on HCP data and show that (a) FC from the top 5% of frames in terms of co-fluctuation amplitude is more similar to time-averaged FC than FC from the bottom 5% of frames and that (b) FC from the top 5% results in more modular networks than FC from the bottom 5%. These results are in exact agreement with what was reported in the main text.

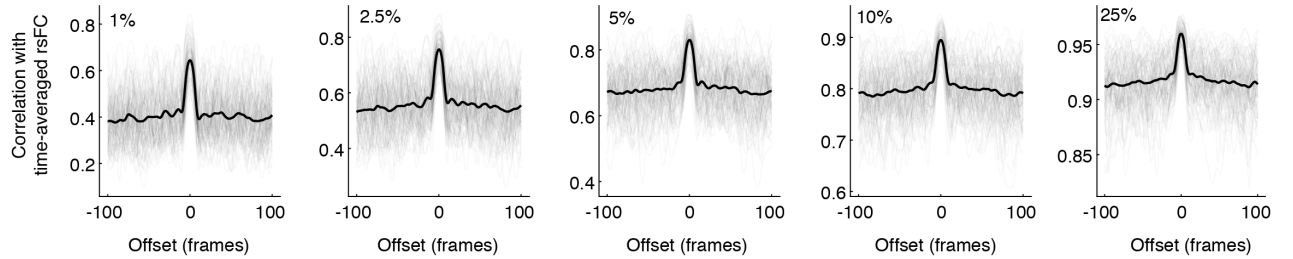


FIG. S10. **Effect of “jittering” on correspondence between time-averaged rsFC and rsFC estimated using reduced number of frames.** In the main text, we demonstrated that the correspondence between time-averaged rsFC and rsFC estimated using high-amplitude co-fluctuation frames was significantly greater than the correspondence using low-amplitude frames. This comparison of the highest- and lowest-amplitude frames can be viewed as a comparison of extremes. A more general test would be to compare the correspondence of rsFC from high-amplitude frames with rsFC from randomly-sampled frames. A truly random sample, however, may destroy any temporal autocorrelation in the time series data. Instead, we identified high-amplitude frames, and used the circular shift operator to move these frames forward and backward in time, approximately preserving their temporal structure. Here, we show the correlation of time-averaged rsFC with rsFC estimated using the offset frames (100 frames forward and backward in time). We repeat this analysis with different thresholds for considering co-fluctuations to be high-amplitude (from left to right, the top 1%, 2.5%, 5%, 10%, and 25%). In general, we find that the correlation with time-averaged rsFC is peaked exactly at an offset of 0 and rapidly decays to a baseline level. This observation holds for all thresholds, and suggests that random samples with temporal structure that preserves autocorrelative properties of the co-fluctuation amplitude time series will, in general, result in estimates of rsFC with poorer correspondence to time-averaged rsFC than the high-amplitude frames, themselves.

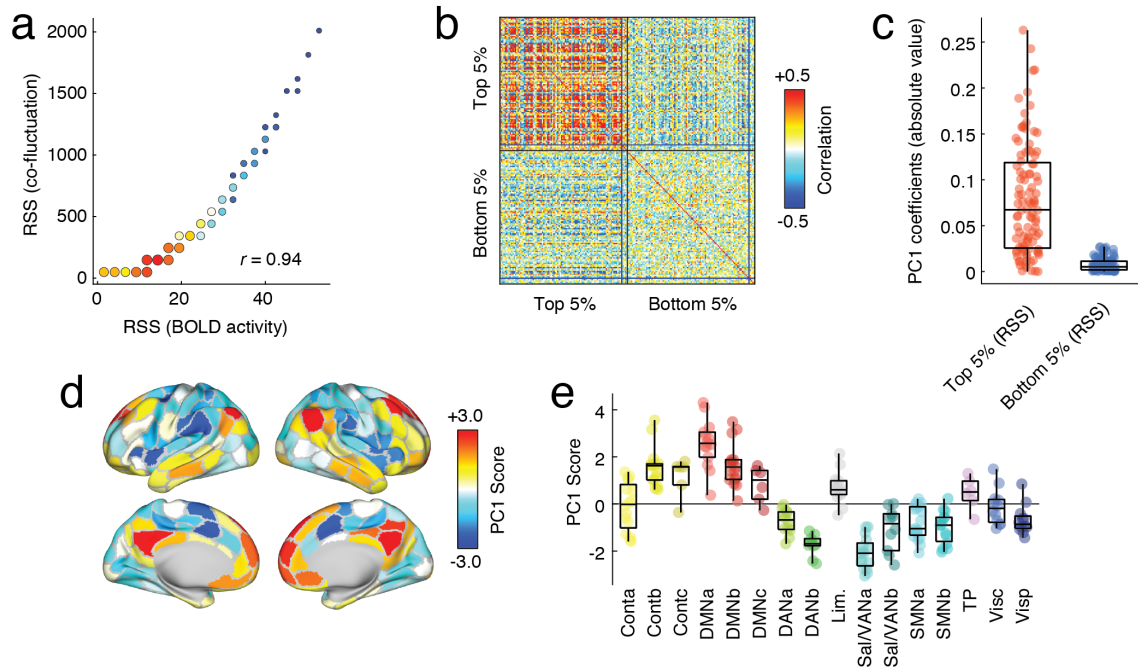


FIG. S11. **Relationship of network co-fluctuations with BOLD fluctuations for MSC dataset.** Here, we replicate results from the main text using the MSC data. Specifically, we relate high-amplitude co-fluctuations to fluctuations in fMRI BOLD activity. We subsequently demonstrate that the high-amplitude fluctuations are driven by activity patterns involving control and default mode networks, and that these patterns are expressed similarly across individuals. As in the main text, we first calculate the root sum square amplitude of BOLD activity at each time point and compare that to the amplitude of co-fluctuations. (a) Pooling data from across subjects, we find that these two variables are highly correlated. (b) To investigate this relationship further, we extract mean activity patterns for each subject and for each scan during the top and bottom 5% time points, indexed according to co-fluctuation amplitude. Here, we show the correlation matrix of those activity vectors. (c) We then performed a principal component analysis of this correlation matrix and found that absolute value of coefficients for the first component (PC1) were greater for the top 5% than the bottom 5%, and (d, e) the PC1 score corresponded to activity patterns that emphasized correlated fluctuations of default mode and control networks that were weakly or anti-correlated with fluctuations elsewhere in the brain. These observations suggest that high-amplitude co-fluctuations, which drive resting-state functional connectivity, are underpinned by instantaneous activation and deactivation of default mode and control network areas.

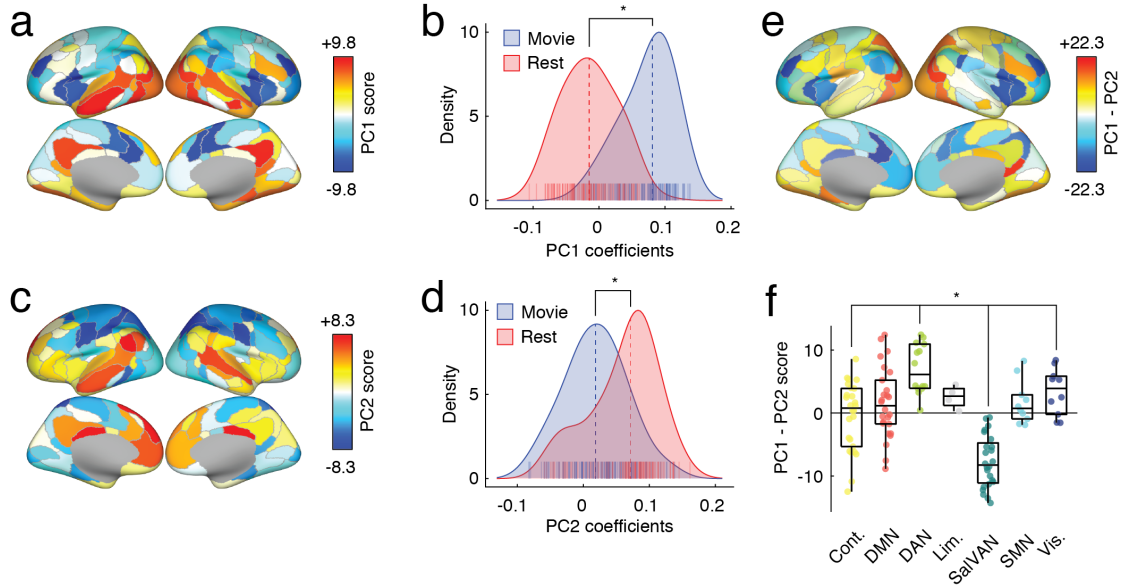


FIG. S12. **Joint PCA analysis of movie-watching and resting-state data.** In the main text we described a procedure for comparing modes of brain activity during high-amplitude frames while subjects were either at rest or watching movies. Here, we present an alternative analysis strategy. The procedure in the main text involved identifying activity patterns during high- and low-amplitude frames, concatenating these patterns across subjects, and performing a PCA on the resulting matrix. Importantly, this procedure was carried out separately for resting-state and movie scans. We then compared the first principal component for each condition. Here, we extract high-amplitude frames from resting-state and movie-watching scans, concatenate them in the same matrix, and perform a joint decomposition using PCA. Our analysis focuses on the first two PCs (PC1 and PC2), whose topographic distributions are shown in panels *a* and *c*. We note that the pattern of PC1 is highly similar to the component obtained from our analysis of the movie-watching data described in the main text. Similarly, PC2 is similar to the component obtained from our analysis of the resting-state data described in the main text. Interestingly, we find that the coefficients loading onto PC1 are stronger for movie-watching scans than for resting-state scans (panel *b*), while the opposite is true for PC2 (panel *d*). We highlight regional and system-level differences in panels *e* and *f*. We note that, as in the main text, we find stronger engagement of dorsal attention and visual networks during movie-watching and stronger engagement of control and salience/ventral attention networks at rest.

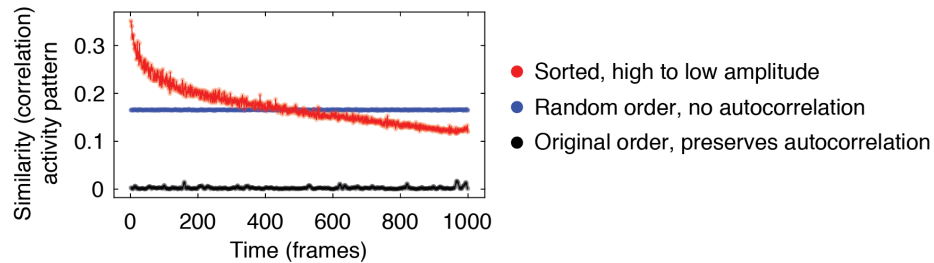


FIG. S13. **Inter-subject correlations of brain activity during rest.** In the main text we demonstrated that ‘high-amplitude frames were characterized by a shared mode of brain activity. The procedure for doing so involved extracting high- and low-amplitude frames using an arbitrary threshold of 5%. Here, we show similar results by rank-ordering frames. Specifically, we calculate the mean inter-subject similarity (absolute Pearson correlation) of brainwide activity patterns. We do this, first, without reordering the resting-state time series (black curve). Because subjects are not locked to a specific stimulus, we find that, on average, subjects’ activity patterns are uncorrelated. The original ordering of the data preserves, at a single-subject level, autocorrelative properties of the fMRI BOLD time series. Next, we destroy this autocorrelation by randomly reordering each subject’s time series and recomputing mean intersubject similarity. While the overall similarity is slightly greater it is still modest. However, when we reorder frames by their overall amplitude (descending order), we find that high-amplitude frames tend to be more strongly correlated with one another, complementing results in the main text and suggesting that high-amplitude co-fluctuations are underpinned by a shared pattern of brain activity.

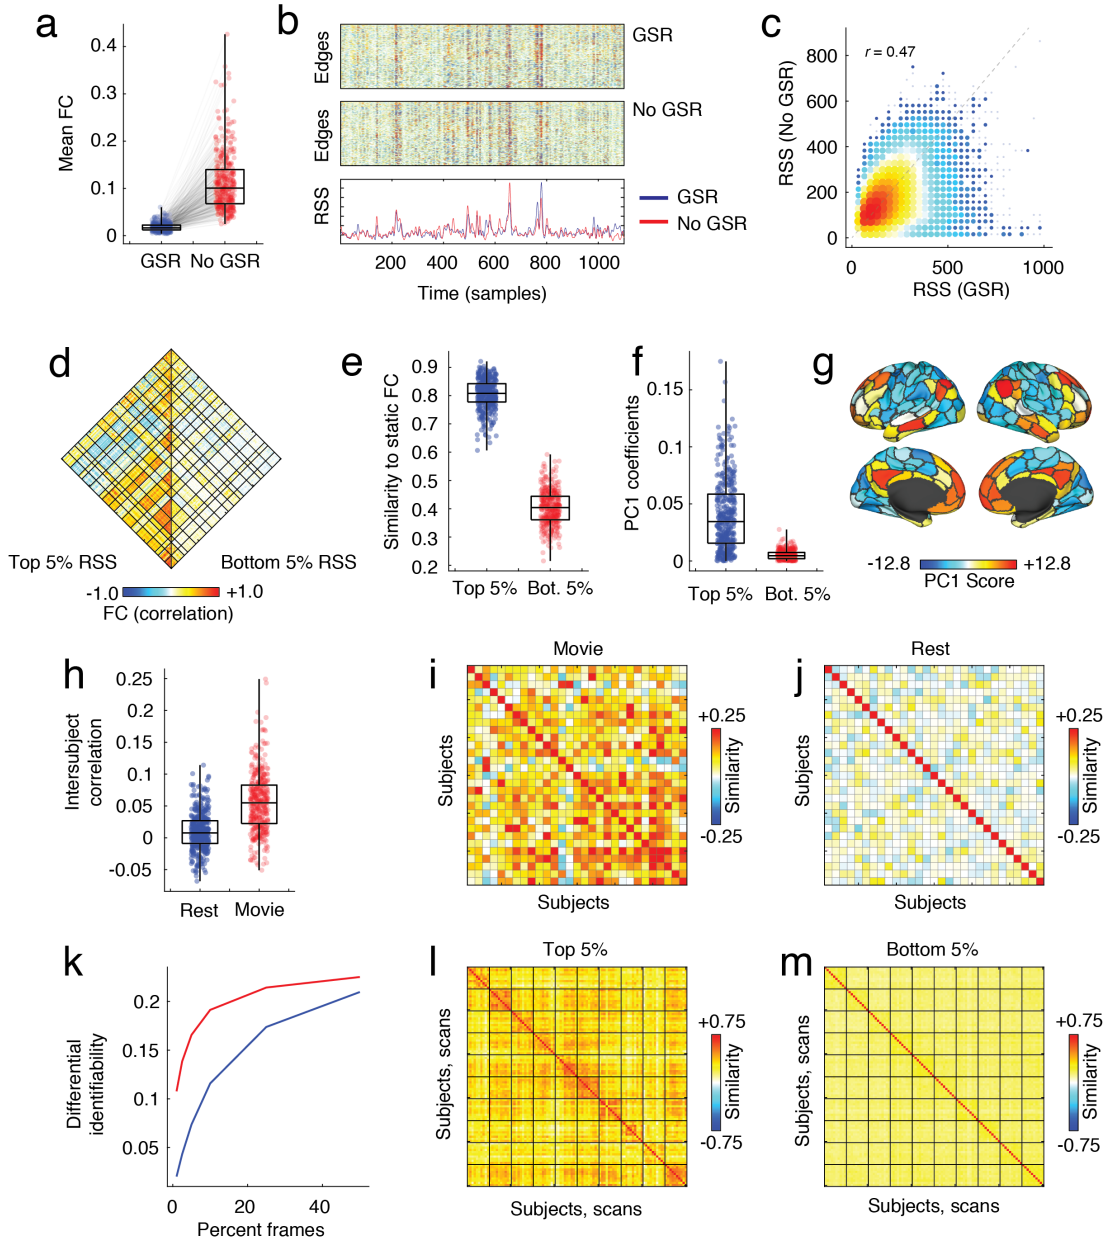


FIG. S14. **Effect of global signal regression on main results.** In the main text, we analyzed resting-state and movie-watching data that had been pre-processed using pipelines that included global signal regression (GSR). Here, we show the effect of omitting GSR from pre-processing pipelines. (a) Comparison of mean static FC (average over all pairs of nodes). As expected, omitting GSR results in an increase in mean FC. (b) Side-by-side comparison of edge (co-fluctuation) time series calculated with (*top*) and without GSR (*middle*) for representative scan. Note that the co-fluctuation amplitude time series remains correlated (*bottom*). (c) Two-dimensional histogram comparing the amplitude of co-fluctuation time series for data processed with/without GSR. As in the main text, we find that FC estimated from the top and bottom 5% of frames (ordered by RSS) is dissimilar (panel d), and that FC estimated from the top 5% of frames was more similar to the static FC estimated using all frames (panel e). We then extracted activity patterns during high-/low-amplitude frames from all subjects and scans and performed PCA on the complete set. We found that the first component indexed high-amplitude frames (panel f) and corresponded to a mode of activity that emphasized the division between DMN/CONT networks and the rest of the brain (panel g). Next, we reanalyzed movie-watching data and found that, as before, subjects' co-fluctuation time series are more correlated during movie-watching than at rest (panel h). In panels i and j we show inter-subject similarity matrices for movie-watching and rest, respectively. Lastly, we reanalyzed MSC data. Replicating results from the main text, we found that without GSR, subjects were still more identifiable based on their high-amplitude frames (panel k). In panels l and m we show example identifiability matrices estimated using FC reconstructed from top/bottom 5% frames (ordered by RSS).

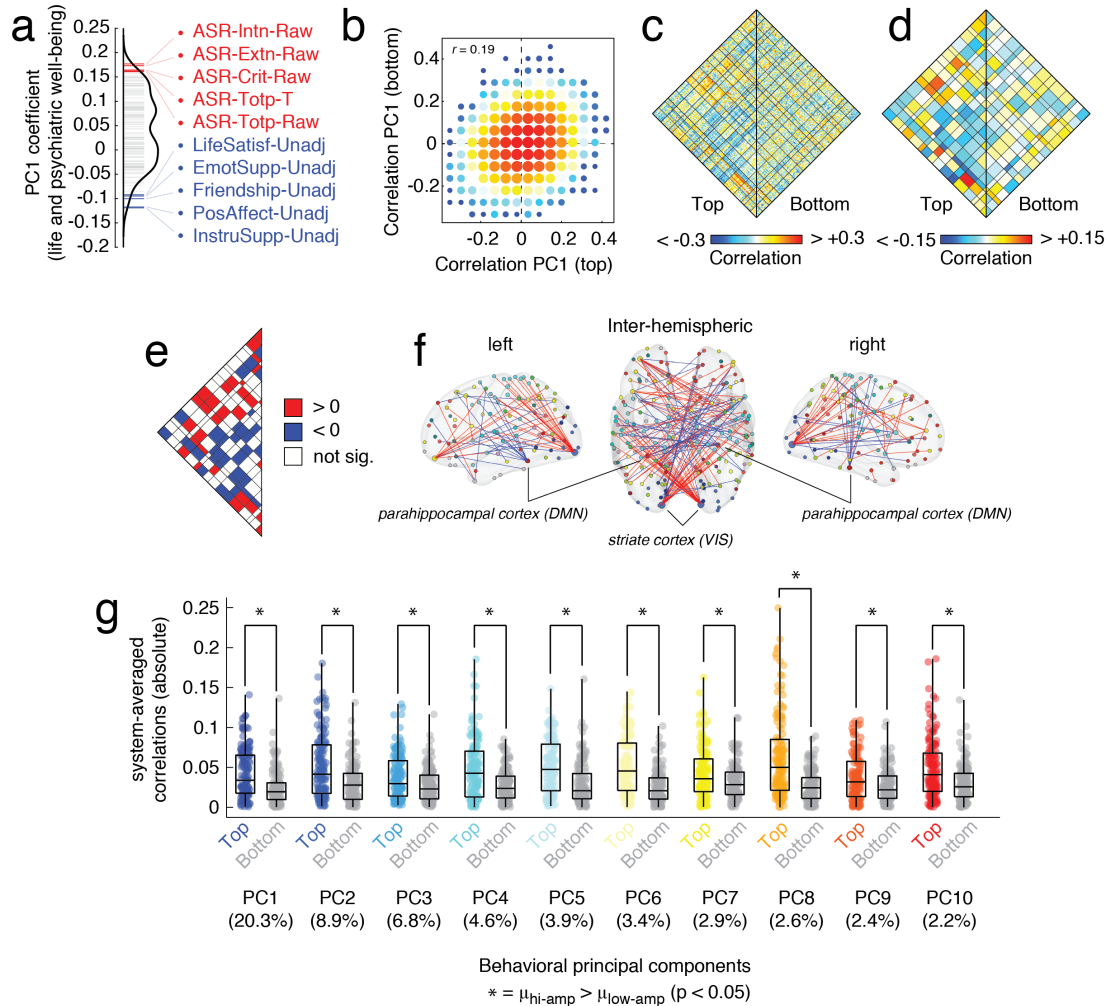


FIG. S15. **Associating co-fluctuations with behavior, traits, and demographics.** (a) First principal component (PC1) obtained from decomposition of standardized behavioral, trait, and demographic data from HCP U100 dataset. We computed the correlation of PC1 with every edge using FC reconstructed from top/bottom 5% frames (ordered by RSS) (b) Scatterplot of correlation patterns. Note that, in general, correlation magnitude is weak. (c) Correlation patterns from top (*left*) and (*bottom*) frames. (d) System-averaged correlation magnitude. (e) Pairs of systems whose average correlation magnitude survived statistical testing (random permutation of regions' system labels). (f) Individual connections that pass system-level significance testing (corrected for multiple comparisons; false discovery rate fixed at 5%; $p_{\text{adjusted}} = 2.01 \times 10^{-4}$) and $p < 0.05$ at the level of connections (uncorrected). We note that this anatomical representation of brain-behavior correlations is intended for the purposes of visualization only. The primary statistical analysis was carried out at the level of brain systems, where we controlled for multiple comparisons. (g) Comparison of absolute system-level correlations for first ten PCs.

Scan	Title	Genre	Runtime
1	Man Up and Go	documentary/emotional	4m20s
1	The First 70	documentary	3m
1	Fixation	documentary/adventure	1m42s
1	The Living	drama	2m
1	SAMSARA	documentary/“unparalleled sensory experience”	1m35s
1	Blood Brother	documentary	2m20s
2	Birdmen	documentary/adventure	3m59s
2	Groomed	drama	1m30s
2	Cold	outdoor/sports	2m
2	Sleepwalkers	drama	2m
2	A Kind of Show	comedy	1m
3	Geofish	documentary/adventure	4m40s
3	The Debut	outdoor/sports	3m23s
3	Dreams of a Life	documentary/mystery	2m10s
3	The Front Man	documentary	2m30s
3	This Is Vanity	drama	1m
4	Planetary	documentary	4m30s
4	Sign Painters	documentary	2m50s
4	Florida Man	documentary/drama	2m
4	The Sleeping Bear	drama	3m40

TABLE S1. Movies included in each movie scan.