

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

Richard F. Betzel^{1–4,*}, Joshua Faskowitz^{1,3}, and Olaf Sporns^{1–4}

¹*Department of Psychological and Brain Sciences,*

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

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

(Dated: October 9, 2019)

Resting-state functional connectivity is used throughout neuroscience to study brain organization and to generate biomarkers of development, disease, and cognition. The processes that give rise to correlated activity are, however, poorly understood. Here, we decompose resting-state functional connectivity using a “temporal unwrapping” procedure to assess the contributions of moment-to-moment activity co-fluctuations to the overall connectivity pattern. This approach temporally resolves functional connectivity at a timescale of single frames, which enables us to make direct comparisons of co-fluctuations of network organization with fluctuations in the BOLD time series. We show that, surprisingly, only a small fraction of frames exhibiting the strongest co-fluctuation amplitude are required to explain a significant fraction of variance in the overall pattern of connection weights as well as the network’s modular structure. These frames coincide with frames of high BOLD activity amplitude, corresponding to activity patterns that are remarkably consistent across individuals and identify fluctuations in default mode and control network activity as the primary driver of resting-state functional connectivity. Our approach reveals fine-scale temporal structure of resting-state functional connectivity, and discloses that frame-wise contributions vary across time. These observations illuminate the relation of brain activity to functional connectivity and open a number of new directions for future research.

INTRODUCTION

Resting-state functional connectivity (rsFC) refers to the correlation structure of fMRI BOLD activity, usually estimated over the course of an entire scan session [1]. Inter-individual differences in rsFC have been linked to variation in biological age [2], cognitive state [3], and clinical status [4]. Other studies have emphasized the dynamic nature of rsFC, using sliding window techniques to generate temporally blurred estimates of rsFC across time [5, 6] and linking changes in network architecture to behavior [7] and phenotypes [8].

Despite intense interest and widespread application, the processes that underpin and shape rsFC are not fully understood. For instance, how do moment-to-moment fluctuations in connectivity contribute to the pattern of rsFC estimated over longer timescales? How are changes in connectivity supported by instantaneous fluctuations in brain activity? While approaches like innovation-driven co-activity patterns (iCAPs) [9], state-based models [10], and sliding-window analyses [11] have provided insight into the dynamics of either activity or connectivity, they generally are less well suited to provide insight into the relations between domains.

Here, we address these questions using a novel approach for modeling instantaneous co-fluctuations in rsFC. We find that at rest co-fluctuations are “bursty” and occur intermittently as part of whole-brain co-fluctuation “events” that are uncorrelated with respiration, cardiac cycle, and in-scanner motion. We then show

that rsFC estimated using only event frames is highly correlated with rsFC estimated over the entire scan session, indicating that rsFC and its system-level organization are driven by co-fluctuations during relatively few frames. Finally, we show that events are underpinned by the activation of a particular spatial mode of brain activity in which default mode and control networks anticorrelated with sensorimotor and attentional systems.

RESULTS

The strength of rsFC between two brain regions can be quantified as the Pearson correlation of their fMRI BOLD time series, which is calculated (after z-scoring) as the mean value of their element-wise product [12]. By omitting the averaging step, we can “temporally unwrap” the correlation measure, which results in a new set of time series – one for every pair of brain regions (network edges) – whose elements represent the magnitude of co-fluctuation between those regions resolved at every moment in time (Fig. 1a). These edge time series can be analyzed directly to pinpoint both the magnitude and timing of co-fluctuations between pairs of brain regions.

In the following subsections, we analyze co-fluctuation time series constructed from functional imaging data acquired as part of the Human Connectome Project [13] (see **Materials and Methods** for details). All results reported in the main text were generated using these data; we replicate these findings using a second, independently acquired dataset [14], with results reported in the **Supplementary Material**.

* [rbetzel @ indiana.edu](mailto:rbetzel@indiana.edu)

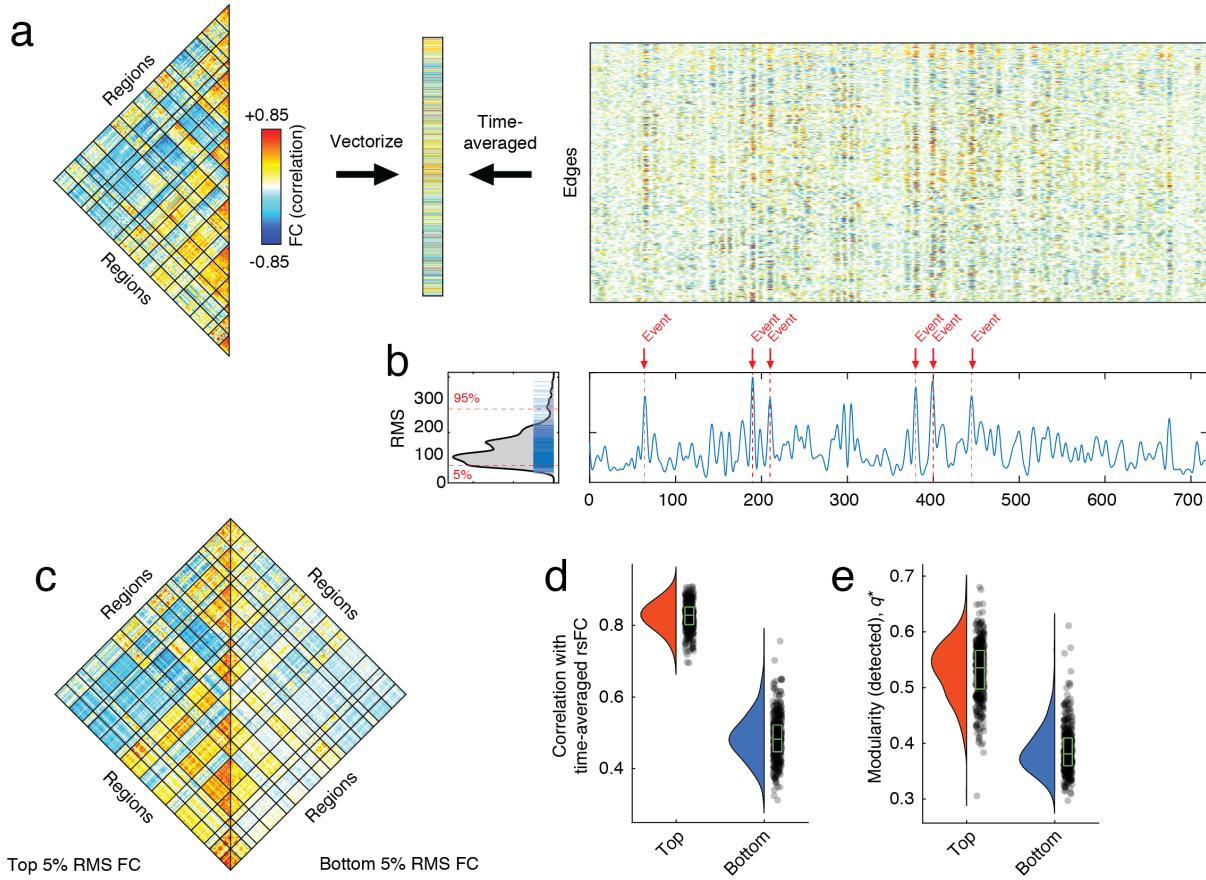


FIG. 1. Co-fluctuation time series reveal event structure of resting-state functional connectivity. (a) We use a “temporal unwrapping” of the Pearson correlation to generate co-fluctuation time series for every pair of brain regions (edges). The elements of the co-fluctuation time series are the element-wise 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 mean square across all co-fluctuation time series and plotting this value as a function of time. We consider high-amplitude values as potential “events”. In panel *b* we label moments in time corresponding to particularly high amplitude co-fluctuations that would be classified as events. The distribution of edge co-fluctuation amplitude is heavy tailed. We wanted to assess the contribution of events and non-events 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.

rsFC is driven by short-lived and high-amplitude co-fluctuation events

When estimated using Pearson correlation, rsFC is expressed as a normalized and time-averaged (over the entire scan session) measure of how strongly the activity of two brain regions co-fluctuates. While past studies have used sliding window methods to generate estimates of moment-to-moment fluctuations in rsFC [5, 6], the use of a windowing procedure results in a temporally blurred estimate of rsFC. This restricts the time scale of obser-

vations regarding dynamic changes in functional connectivity to the width of the time window, generally on the order of dozens of frames (1 minute of real time). Here, we address this limitation using co-fluctuation times series, which allow us to accurately assess contributions made to rsFC by single frames without the necessity of a sliding window.

When analyzed across the whole brain, we find that edge time series exhibit “bursty” behavior, such that the amplitude of co-fluctuations (quantified by computing the root mean square; RMS) moves around a mean value, but is punctuated by brief, intermittent, and dis-

proportionately large fluctuations, which we refer to as “events” (Fig. 1b). These events are not directly related to cardiac and respiratory cycles, in-scanner head motion (Fig. S1), and spectral properties of fMRI BOLD time series (Fig. S2), and appear aperiodic with heavy-tailed distributions of event size, event durations, and inter-event intervals (Fig. S3).

To better understand how instantaneous co-fluctuations contribute to whole-brain rsFC, we isolated high-amplitude “events” and compared them with low-amplitude episodes (top and bottom 5% in terms of co-fluctuation amplitude; 60 frames for HCP; see Fig. S4 for comparisons at other percentiles). We then estimated rsFC separately for each category, using only fMRI BOLD data corresponding to those time points and compared the resulting networks. First, we found that connection weights were significantly stronger during events than non-events (within-sample *t*-test; $p < 10^{-15}$; Fig. 1c). Next, we calculated the similarity of rsFC estimated during events and low-amplitude episodes with respect to time-averaged rsFC estimated using the full time series. We found that the event networks were highly correlated with rsFC ($r = 0.81 \pm 0.05$) while the non-event networks were much less correlated ($r = 0.54 \pm 0.07$) and that these differences were highly significant (*t*-test, $p < 10^{-15}$; Fig. 1d). We also performed an analogous comparison of network modularity [15], an index that can be interpreted as a measure of how segregated a network’s systems are from one another. As before, we found that modularity was greater in the event networks ($q = 0.51 \pm 0.06$) compared to the non-event networks ($q = 0.37 \pm 0.05$) (*t*-test, $p < 10^{-15}$; Fig. 1e).

In the supplement we show similar results in a second dataset (Fig. S5). We also demonstrate that these effects persist with highly conservative motion censoring (Fig. S6), when using an alternative strategy for estimating networks from the top and bottom 5% time points (Fig. S7), and when comparing against a null model that preserves the temporal structure of events while sampling frames randomly from the entire time series (Fig. S8).

Collectively, these results suggest that rsFC, estimated over long time scales, is driven by a small number of brief, intermittent, and high-amplitude co-fluctuations. The network structure over these points in time contributes disproportionately to the overall modularity and system-level organization of cerebral cortex, as estimated from long-time averages of rsFC. In contrast, low-amplitude co-fluctuations are only weakly related to the overall pattern of rsFC and correspond to less modular architectures.

FC events are driven by fluctuations of task-positive/task-negative mode of brain activity

In the previous section we demonstrated that time-averaged rsFC can be explained by high-amplitude co-

fluctuations that occur during a relatively small number of frames. It remains unclear, however, whether co-fluctuation events are underpinned by a specific pattern of brain activity or whether they reflect contributions from multiple distinct patterns. Here, we address this question directly, using co-fluctuation time series, which enable us to temporally localize network-level events at the resolution of individual frames, and to compare co-fluctuations directly with brain activity.

As a first point of comparison, we calculated the RMS of both the co-fluctuation time series as well as the z-scored fMRI BOLD time series. We found that, across subjects, these time series were highly correlated ($r = 0.97$), indicating that co-fluctuation events have an almost one-to-one correspondence with high-amplitude BOLD fluctuations (Fig. 2a). This relationship is expected; because co-fluctuations are calculated as products of z-scored regional activity, their amplitudes will necessarily be correlated with one another.

Given that fluctuations in BOLD activity are greater during events than non-events, we asked whether they formed a consistent and recognizable pattern of activity. To address this question, we calculated the mean activity pattern for each subject during their events and non-events and computed between-subject and between-scan similarity (Fig. 2b). In general, activity during events was more correlated across subjects compared to the activity patterns during non-events (*t*-test, $p < 10^{-15}$). To better understand what was driving these correlations, we performed a principal components analysis of the activity patterns during events and non-events, aggregated over all subjects and scans. We focused on the first principal component (PC1), which explained 26% of total variance. The coefficients for PC1 were, on average, much greater for events than non-events (*t*-test, $p < 10^{-15}$; Fig. 2c), indicating that PC1 was descriptive of activity patterns during events but less so for non-events. We then mapped component scores for PC1 onto the cortical surface and found that PC1 corresponded to a mode of activity that delineates regions in default mode and control networks from sensorimotor and attentional networks (Fig. 2d,e). We replicated these results in a second dataset (see Fig. S9).

These results suggest that underlying co-fluctuation events is a mode of brain activity whose spatial pattern resembles the traditional task-positive/task-negative division of the brain [16]. This pattern of activity is similar across individuals, suggesting a conserved mechanism by which rsFC emerges from brain activity. These observations suggest a fundamental link between distinct patterns of brain activity and connectivity while further clarifying the origins of co-fluctuation “events.”

DISCUSSION

Here, we presented a general approach for “temporally unwrapping” Pearson correlations to generate time se-

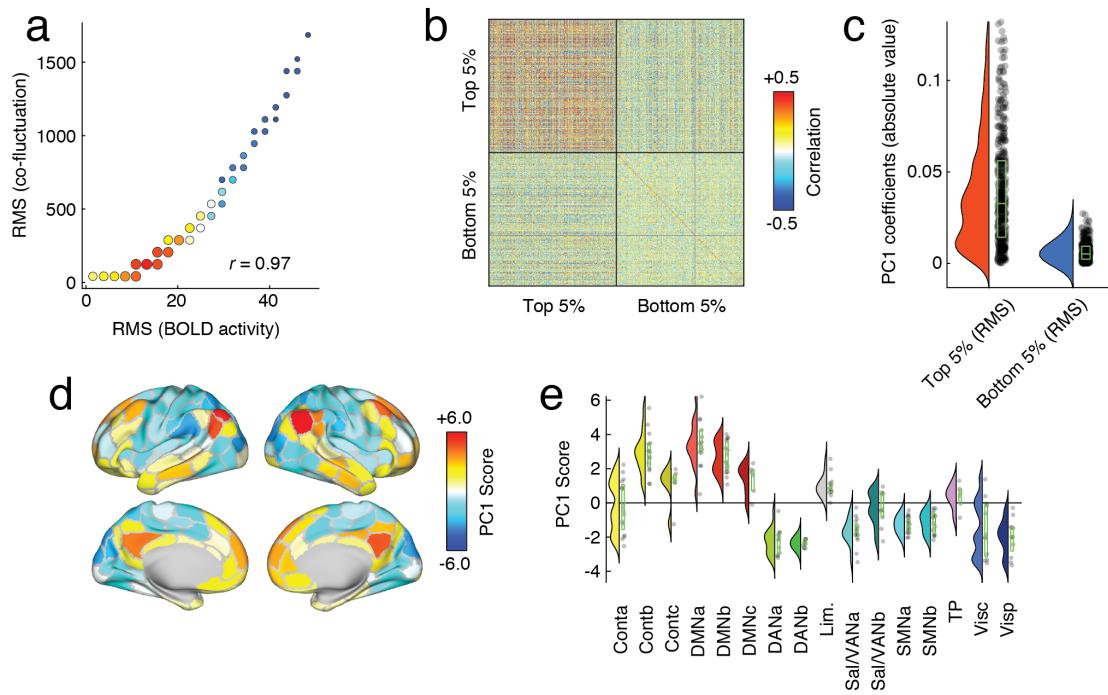


FIG. 2. Relationship of network co-fluctuations with BOLD fluctuations. In the previous section we demonstrated that resting-state functional connectivity could be explained on the basis of relatively few frames during which high-amplitude co-fluctuations occurred. Here, we relate those co-fluctuation frames to BOLD activity fluctuations. We first the root mean 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 co-fluctuation events, which drive resting-state functional connectivity, are underpinned by instantaneous activation and deactivation of default mode and control network areas.

ries of inter-regional co-fluctuations along network edges. This simple procedure enables us to parse the contributions made by individual frames to rsFC. We find that, in general, we can accurately estimate whole-brain rsFC and its system-level organization using data from a relatively small number of frames. Importantly, we link these frames to a high-modularity brain state and to a specific mode of brain activity, in which default mode and control networks fluctuate in opposition to sensorimotor and attention systems.

Origins of co-fluctuation events

Altogether, these findings speak to the origins of whole-brain rsFC and modularity in terms of co-fluctuation events that are, themselves, driven by instantaneous and opposed fluctuations of association and sensorimotor cortex. These observations clarify our understanding of, as well as challenge core assumptions underpinning the origins of rsFC. Perhaps the most pressing

question concerns the neurobiological processes underlying both the co-fluctuation events and their concomitant BOLD signal patterns. Our findings suggest that these events are unrelated to head motion, cardiac cycle, or respiration. However, they make disproportionately large contributions to the observed patterns of rsFC [17], a measure that is used widely throughout neuroscience and neuroimaging communities. The analysis presented here suggests that high co-fluctuation events may correspond to previously observed “network states” associated with time-varying changes in modular structure [18, 19]. Determining the causes of such events will help to better contextualize findings related to rsFC and suggest avenues for future research.

For example, the categorization of time points into “events” and “non-events” may have important implications for “functional fingerprinting” [20]. We find that events are associated with activation patterns that are more similar across subjects than activation patterns during non-events – How repeatable are those patterns within an individual? Are “functional fingerprints” more

strongly expressed during events or non-events [21]? Are event or non-event patterns more useful for accurately identifying a subject [22]? Do the statistics of event and non-event patterns differ across genetic and phenotypic markers as well as clinical, developmental, or cognitive state? How does underlying white-matter connectivity relate to rsFC during events and non-events? Past studies have reported that rsFC decouples from anatomy during modular brain states [23, 24]. Our approaches allows us to investigate this relationship with greatly improved temporal resolution.

Contributions of brain activity to the system-level organization of rsFC

Lastly, our findings hint at a crucial link between instantaneous fluctuations in activity and the organization of rsFC [25, 26]. Many studies have found that the community structure of rsFC resembles known co-activation patterns, including task-evoked activity [27, 28]. Here, we proposed a strategy that enables us to tease apart the precise contribution of instantaneous BOLD fluctuations (and their topography) to rsFC.

We demonstrated that a particular pattern of activity involving default mode and control regions is primarily responsible for driving co-fluctuation events and, in turn, whole-brain rsFC. While this mode made the greatest contribution, it is likely that other modes make non-trivial contributions as well. By extending the definition of an event to include lower-amplitude fluctuations, we expect to find patterns of activity that correspond to other, well-known brain systems [9]. Moreover, we speculate that these patterns likely recombine in different proportions as a function of task complexity and domain [27, 29] and across individuals [30]. In future work, the proportion of variance explained by different patterns and other statistics related to events, including the frequency with which they occur, may serve as potent correlates of cognitive and disease state. Because events appear to drive the overall configuration of rsFC, we further speculate that their statistics may serve as important complements to traditional measures of rsFC.

Conclusion

In conclusion, our study discloses a novel a link between cortical activity and rsFC, facilitating a statistical explanation of the brain's system-level architecture in terms of intermittent, short-lived, high-amplitude fluctuations in activity and co-activity. Our methodological framework is readily applicable to other imaging datasets and recording modalities, including observations at neuronal scales, enabling the study of neural co-activity at unprecedented temporal resolution.

MATERIALS AND METHODS

Datasets

The Human Connectome Project (HCP) dataset [13] included resting state functional data (rsfMRI) from 100 unrelated adult subjects (54% female, mean age = 29.11 \pm 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 [31]. 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 [14] included rsfMRI from 10 adults (50% female, mean age = 29.1 \pm 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

HCP Functional Preprocessing

Functional images in the HCP dataset were minimally preprocessed according to the description provided in [31]. 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 [32]. The resultant CIFTI file for each HCP subject used in this study followed the file naming pattern: *_REST{1,2}_{L,R}_Atlas_MSKAll.dtseries.nii.

MSC Functional Preprocessing

Functional images in the MSC dataset were preprocessed using *fMRIprep* 1.3.2 [33], which is based on Nipype 1.1.9 [34]. 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 [35], 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* [36, 37], 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* [38], 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 [39]. Spatial normalization to the *ICBM 152 Nonlinear Asymmetrical template version 2009c* [40] was performed through nonlinear registration with *antsRegistration*, using brain-extracted versions of both T1w volume and template. Brain tissue segmentation of cerebrospinal fluid (CSF), white-matter (WM) and gray-matter (GM) was performed on the brain-extracted T1w using FSL's *fast* [41].

Functional data was slice time corrected using AFNI's *3dTshift* and motion corrected using FSL's *mcflirt* [42]. *Fieldmap-less* distortion correction was performed by co-registering the functional image to the same-subject T1w image with intensity inverted [43] constrained with an average fieldmap template [44], implemented with *antsRegistration*. This was followed by co-registration to the corresponding T1w using boundary-based registration [45] 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 [46]. 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 [47]. Data was visually inspected for whole brain field of view coverage, signal artifacts, and proper alignment to the corresponding anatomical image. Functional data were excluded if greater than 25% of the frames exceeded 0.2 mm framewise displacement [48]. Furthermore, functional data were excluded if marked as an outlier (exceeding 1.5x inter-quartile range in the adverse direction) in more than half of the following image quality metrics (calculated within-dataset, across all functional acquisitions): *dvars*, *tsnr*, *fd_mean*, *aor*, *aqi*, *snr*, and *etc*. Information about these image quality metrics can be found within *MRIQC*'s documentation [49].

Functional and Structural Networks Preprocessing

Parcellation Preprocessing

A functional parcellation designed to optimize both local gradient and global similarity measures of the fMRI signal [50] (*Schaefer200*) was used to define 200 areas on the cerebral cortex. These nodes are also mapped to the *Yeo* canonical functional networks [25]. 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 [51] (trained on 100 unrelated HCP subjects) and FreeSurfer's *mrfs_ca_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) [48], confound regressed and standardized using Nilearn's *signal.clean*, which removes confounds orthogonally to the temporal filters [52]. The confound regression employed [53] 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 [48]. Following this pre-

processing and nuisance regression, residual mean BOLD time series at each node was recovered.

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 parcels 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 with 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 $\sigma_i = \sqrt{\frac{1}{T-1} \sum_t [x_i(t) - \mu_i]^2}$ 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 [54]. 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 [15]:

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

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 $v^\pm = \sum_i k_i^\pm$.

AUTHOR CONTRIBUTIONS

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

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/>.

CODE AVAILABILITY

All processing and analysis code is available upon reasonable request.

[1] R. C. Craddock, S. Jbabdi, C.-G. Yan, J. T. Vogelstein, F. X. Castellanos, A. Di Martino, C. Kelly, K. Heberlein, S. Colcombe, and M. P. Milham, *Nature methods* **10**, 524 (2013).

[2] S. Gu, T. D. Satterthwaite, J. D. Medaglia, M. Yang, R. E. Gur, R. C. Gur, and D. S. Bassett, *Proceedings of the National Academy of Sciences* **112**, 13681 (2015).

[3] M. W. Cole, D. S. Bassett, J. D. Power, T. S. Braver, and S. E. Petersen, *Neuron* **83**, 238 (2014).

[4] A. Fornito, A. Zalesky, and M. Breakspear, *Nature Reviews Neuroscience* **16**, 159 (2015).

[5] E. A. Allen, E. Damaraju, S. M. Plis, E. B. Erhardt, T. Eichele, and V. D. Calhoun, *Cerebral cortex* **24**, 663 (2014).

[6] A. Zalesky, A. Fornito, L. Cocchi, L. L. Gollo, and M. Breakspear, *Proceedings of the National Academy of Sciences* **111**, 10341 (2014).

[7] A. Kucyi and K. D. Davis, *Neuroimage* **100**, 471 (2014).

[8] E. Damaraju, E. A. Allen, A. Belger, J. M. Ford, S. McEwen, D. Mathalon, B. Mueller, G. Pearlson, S. Potkin, A. Preda, *et al.*, *NeuroImage: Clinical* **5**, 298 (2014).

[9] F. I. Karahanoglu and D. Van De Ville, *Nature communications* **6**, 7751 (2015).

[10] D. Vidaurre, S. M. Smith, and M. W. Woolrich, *Proceedings of the National Academy of Sciences* **114**, 12827 (2017).

[11] R. Hindriks, M. H. Adhikari, Y. Murayama, M. Ganzetti, D. Mantini, N. K. Logothetis, and G. Deco, *Neuroimage* **127**, 242 (2016).

[12] S. M. Smith, K. L. Miller, G. Salimi-Khorshidi, M. Webster, C. F. Beckmann, T. E. Nichols, J. D. Ramsey, and M. W. Woolrich, *Neuroimage* **54**, 875 (2011).

[13] 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).

[14] 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).

[15] M. Rubinov and O. Sporns, *Neuroimage* **56**, 2068 (2011).

[16] M. Mennes, C. Kelly, X.-N. Zuo, A. Di Martino, B. B. Biswal, F. X. Castellanos, and M. P. Milham, *Neuroimage* **50**, 1690 (2010).

[17] E. Tagliazucchi, P. Balenzuela, D. Fraiman, and D. R. Chialvo, *Frontiers in physiology* **3**, 15 (2012).

[18] R. F. Betzel, M. Fukushima, Y. He, X.-N. Zuo, and O. Sporns, *NeuroImage* **127**, 287 (2016).

- [19] J. M. Shine, P. G. Bissett, P. T. Bell, O. Koyejo, J. H. Balsters, K. J. Gorgolewski, C. A. Moodie, and R. A. Poldrack, *Neuron* **92**, 544 (2016).
- [20] E. S. Finn, X. Shen, D. Scheinost, M. D. Rosenberg, J. Huang, M. M. Chun, X. Papademetris, and R. T. Constable, *Nature neuroscience* **18**, 1664 (2015).
- [21] C. Peña-Gómez, A. Avena-Koenigsberger, J. Sepulcre, and O. Sporns, *Cerebral Cortex* **28**, 2922 (2017).
- [22] E. Amico and J. Goñi, *Scientific reports* **8**, 8254 (2018).
- [23] M. Fukushima, R. F. Betzel, Y. He, M. P. van den Heuvel, X.-N. Zuo, and O. Sporns, *Brain Structure and Function* **223**, 1091 (2018).
- [24] P. Barttfeld, L. Uhrig, J. D. Sitt, M. Sigman, B. Jarraya, and S. Dehaene, *Proceedings of the National Academy of Sciences* **112**, 887 (2015).
- [25] 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).
- [26] M. Pedersen, A. Omidvarnia, A. Zalesky, and G. D. Jackson, *Neuroimage* **181**, 85 (2018).
- [27] S. M. Smith, P. T. Fox, K. L. Miller, D. C. Glahn, P. M. Fox, C. E. Mackay, N. Filippini, K. E. Watkins, R. Toro, A. R. Laird, *et al.*, *Proceedings of the National Academy of Sciences* **106**, 13040 (2009).
- [28] N. A. Crossley, A. Mechelli, P. E. Vértes, T. T. Winton-Brown, A. X. Patel, C. E. Ginestet, P. McGuire, and E. T. Bullmore, *Proceedings of the National Academy of Sciences* **110**, 11583 (2013).
- [29] T. Yarkoni, R. A. Poldrack, T. E. Nichols, D. C. Van Essen, and T. D. Wager, *Nature methods* **8**, 665 (2011).
- [30] C. Gratton, T. O. Laumann, A. N. Nielsen, D. J. Greene, E. M. Gordon, A. W. Gilmore, S. M. Nelson, R. S. Coalson, A. Z. Snyder, B. L. Schlaggar, *et al.*, *Neuron* **98**, 439 (2018).
- [31] 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).
- [32] 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).
- [33] 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.
- [34] K. Gorgolewski, C. D. Burns, C. Madison, D. Clark, Y. O. Halchenko, M. L. Waskom, and S. Ghosh, *Frontiers in Neuroinformatics* **5**, 13 (2011).
- [35] 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.
- [36] 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).
- [37] B. Avants, C. Epstein, M. Grossman, and J. Gee, *Medical Image Analysis* **12**, 26 (2008).
- [38] A. M. Dale, B. Fischl, and M. I. Sereno, *NeuroImage* **9**, 179 (1999).
- [39] 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).
- [40] V. Fonov, A. Evans, R. McKinstry, C. Almlie, and D. Collins, *NeuroImage* **47**, *Supplement 1*, S102 (2009).
- [41] Y. Zhang, M. Brady, and S. Smith, *IEEE Transactions on Medical Imaging* **20**, 45 (2001).
- [42] M. Jenkinson, P. Bannister, M. Brady, and S. Smith, *NeuroImage* **17**, 825 (2002).
- [43] S. Wang, D. J. Peterson, J. C. Gatenby, W. Li, T. J. Grabowski, and T. M. Madhyastha, *Frontiers in Neuroinformatics* **11** (2017), 10.3389/fninf.2017.00017.
- [44] 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).
- [45] D. N. Greve and B. Fischl, *NeuroImage* **48**, 63 (2009).
- [46] J. D. Power, A. Mitra, T. O. Laumann, A. Z. Snyder, B. L. Schlaggar, and S. E. Petersen, *NeuroImage* **84**, 320 (2014).
- [47] O. Esteban, D. Birman, M. Schaer, O. O. Koyejo, R. A. Poldrack, and K. J. Gorgolewski, *PloS one* **12**, e0184661 (2017).
- [48] L. Parkes, B. Fulcher, M. Yücel, and A. Fornito, *Neuroimage* **171**, 415 (2018).
- [49] “Mriqc: Advancing the automatic prediction of image quality in mri from unseen sites,” (2019).
- [50] 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).
- [51] 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).
- [52] M. A. Lindquist, S. Geuter, T. D. Wager, and B. S. Caffo, *Human brain mapping* **40**, 2358 (2019).
- [53] T. D. Satterthwaite, M. A. Elliott, R. T. Gerraty, K. Ruparel, J. Loughead, M. E. Calkins, S. B. Eickhoff, H. Hakonarson, R. C. Gur, R. E. Gur, *et al.*, *Neuroimage* **64**, 240 (2013).
- [54] M. E. Newman and M. Girvan, *Physical review E* **69**, 026113 (2004).
- [55] D. Prichard and J. Theiler, *Physical review letters* **73**, 951 (1994).
- [56] J. D. Power, K. A. Barnes, A. Z. Snyder, B. L. Schlaggar, and S. E. Petersen, *Neuroimage* **59**, 2142 (2012).

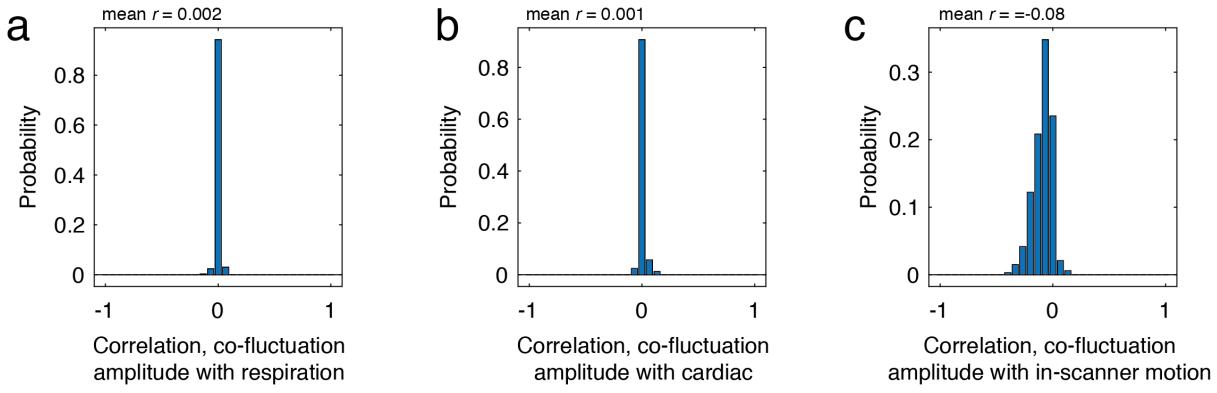


FIG. S1. Comparison of co-fluctuation amplitude with confounding variables. In the main text we calculated the magnitude of co-fluctuation at every frame. A concern is that variation in this measure could be attributed to physiological and motion-related variables of non-neural origins. To address this concern, we calculated the correlation of co-fluctuation amplitude with three variables: respiration and cardiac 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 calculated the mean value within a frame. We calculated these variables for every subject and scan session in the HCP dataset and computed their correlation with the co-fluctuation amplitude. The distributions of correlation coefficients, shown here in panels *a*, *b*, and *c*, were tightly centered on zero, suggesting that co-fluctuation amplitude is not obviously related to standard physiological or motion-related variables.

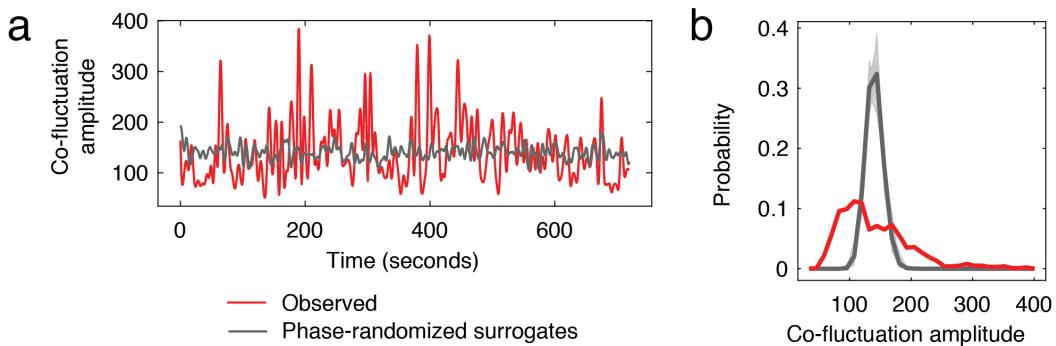


FIG. S2. Comparison of co-fluctuation amplitude from observed data and from phase-randomized surrogates. In the main text we calculate co-fluctuation amplitude as root mean square of edge co-fluctuations at each moment in time. Here, we compare these observed amplitudes with those estimated from phase-randomized surrogate time series. The phase randomization procedure has been described in detail elsewhere [55]. 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 and subsequently calculate their co-fluctuation time series and overall amplitude (panel *a*). We repeated this procedure 100 times and found that the distribution of co-fluctuation amplitude for the observed data was broad and included a heavy tail that was not present in the surrogate data (panel *b*). This observation suggests that the observed co-fluctuations (in particular the high-amplitude ‘events’) cannot be explained by spectral properties of the fMRI BOLD time series alone.

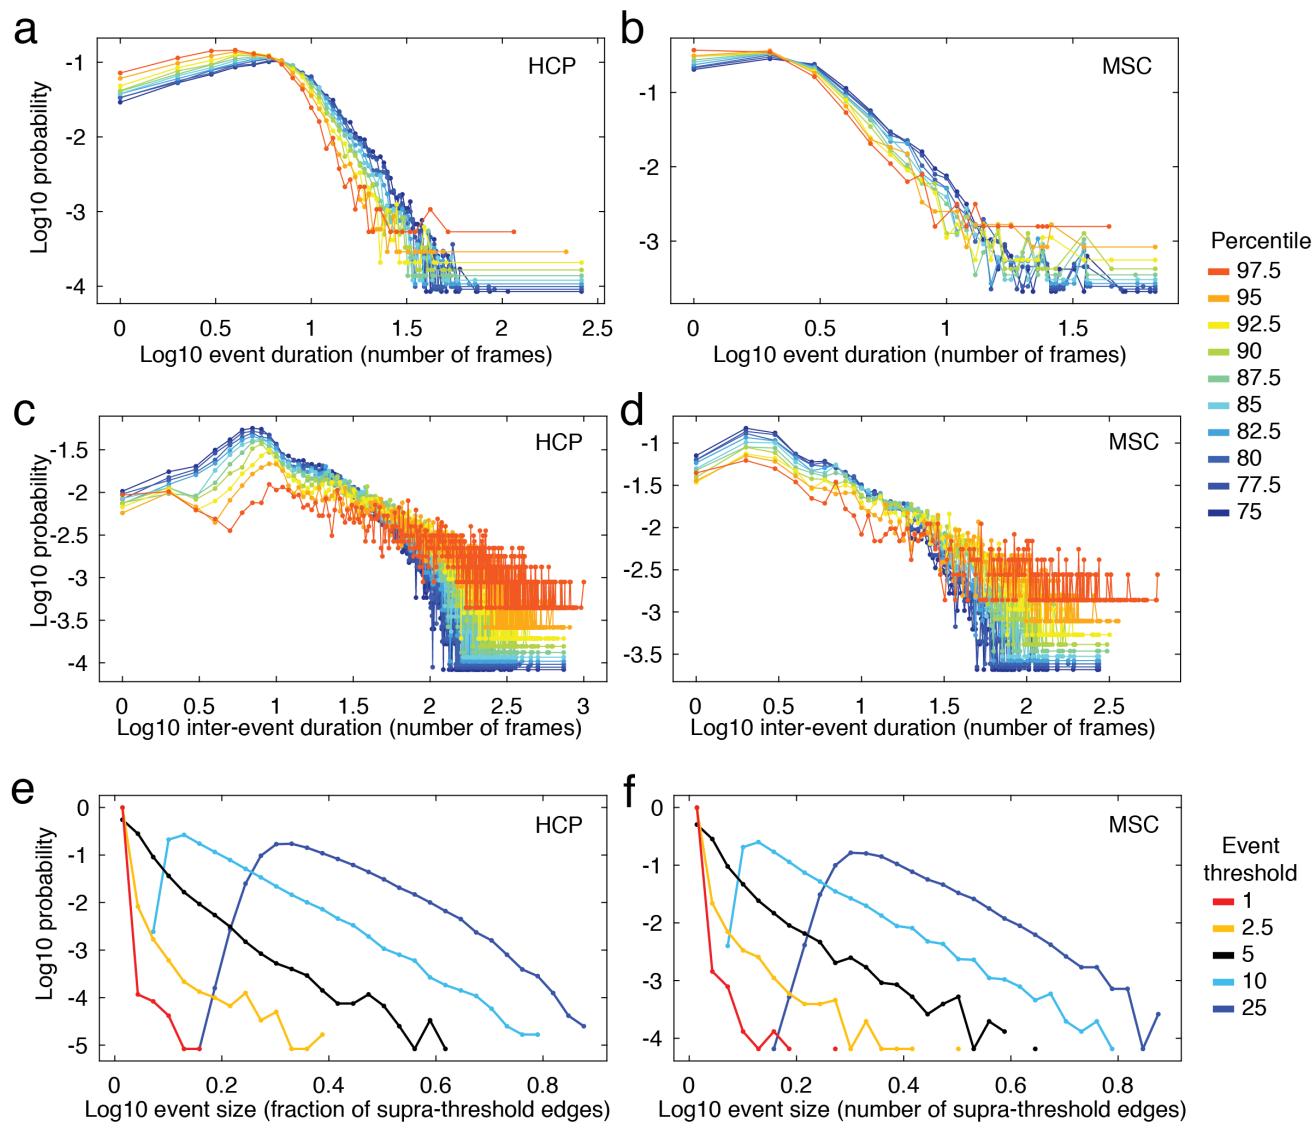


FIG. S3. Event duration and inter-event interval distributions. For each scan session, we calculated the co-fluctuation amplitude at every frame. We imposed percentile-based thresholds on these data (percentiles calculated based on pooled data from all subjects and all scan sessions). Thresholding the co-fluctuation amplitude time series results in a binary classification of time points as either “events” or “non-events”. From these observations, we calculated two quantities: “event duration” as the number of consecutive frames classified as events and “inter-event duration” as the number of frames between successive events. We repeated this analysis for both the HCP and MSC datasets. Panels *a* and *b* show event 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 inter-event durations for HCP and MSC datasets. Additionally, we assessed the size of events, as measured by the fraction of all edges whose co-fluctuation amplitude at a given frame exceeded some threshold. Here, we identified events as time points at which the co-fluctuation amplitude was in the top 1%, 2.5%, 5%, 10%, and 25% (event thresholds are indicated by different colors in each plot). Then, for each time point classified as an event, 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 event sizes follow a broad and heavy-tailed distribution, suggesting that they follow no characteristic scale of description.

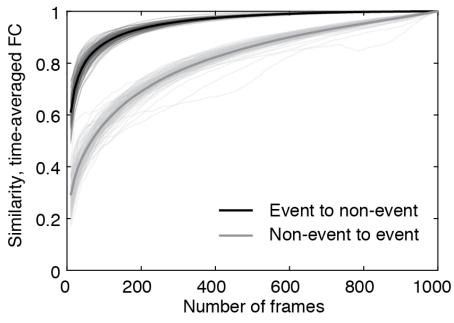


FIG. S4. Similarity of time-averaged FC with FC estimated using fewer frames. In the main text, we showed that rsFC, when estimated using the top 5% of frames (ordered by co-fluctuation magnitude) resulted in a connectivity matrix that was much more similar to time-averaged FC than the matrix generated 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 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.

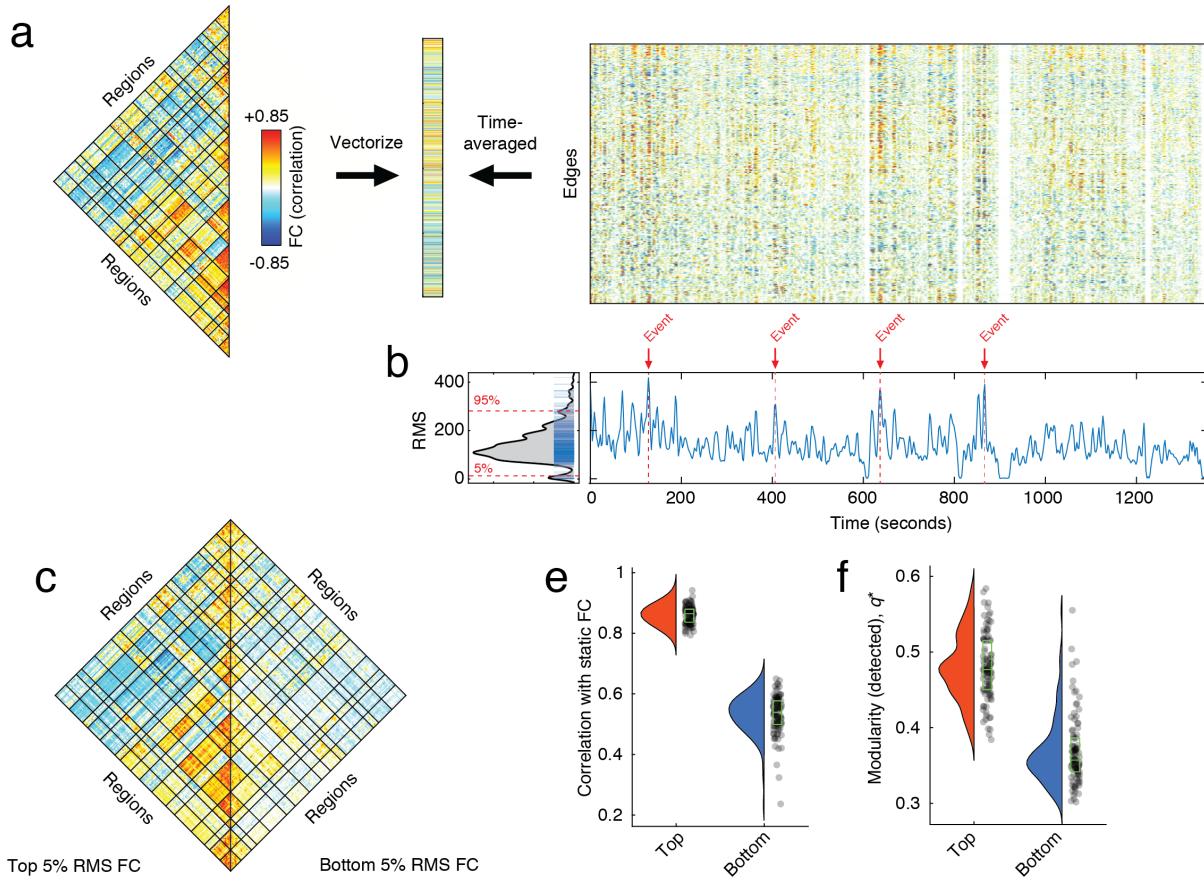


FIG. S5. Co-fluctuation time series reveal event structure of resting-state functional connectivity for MSC dataset. (a) We use a “temporal unwrapping” of the Pearson correlation to generate co-fluctuation time series for every pair of brain regions (edges). The elements of the co-fluctuation time series 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 mean square across all co-fluctuation time series and plotting this value as a function of time. We consider high-amplitude values as potential “events”. The distribution of edge co-fluctuation amplitude is heavy tailed. We wanted to assess the contribution of events and non-events 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.

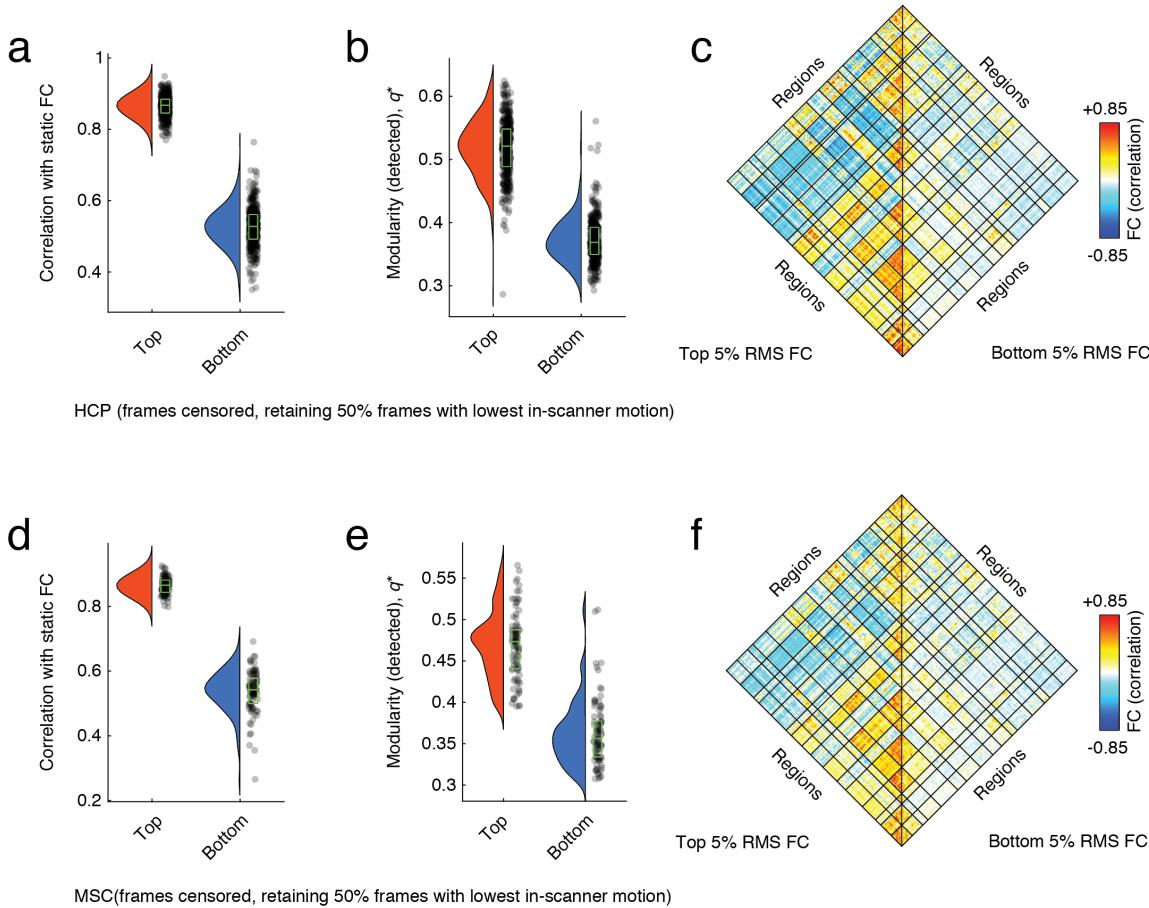


FIG. S6. Effect of frame-censoring. In the main text, we demonstrated that FC estimated using frames corresponding to high-amplitude co-fluctuations was more similar to time-averaged FC than the FC estimated using low-amplitude co-fluctuation frames. Here, we perform an identical analysis using only the bottom 50% frames in terms of in-scanner motion [56]. 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.

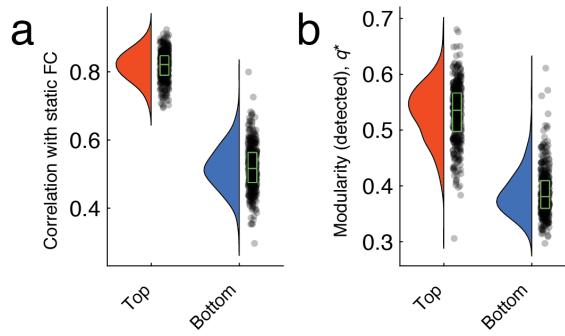


FIG. S7. Alternative strategy for estimating FC from a limited number of frames. In the main text, we estimated FC from the top and bottom 5% of frames by extracting fMRI BOLD activity from computing the correlation structure. An alternative strategy for estimating FC is to simply average co-fluctuation 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.

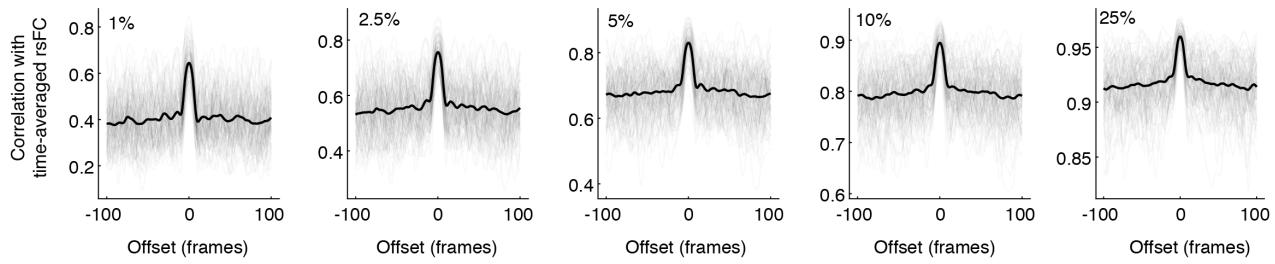


FIG. S8. Effect of ‘jittering’ on correspondence between time-averaged rsFC and rsFC estimated using reduced number of frames. In the main text, we demonstrated that the correspondence between time-averaged rsFC and rsFC estimated using high-amplitude co-fluctuation frames was significantly greater than 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 ‘events’ as frames whose co-fluctuation amplitude exceeded some label, 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 event frames (100 frames forward and backward in time). We repeat this analysis with different event thresholds (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 event 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 event frames, themselves.

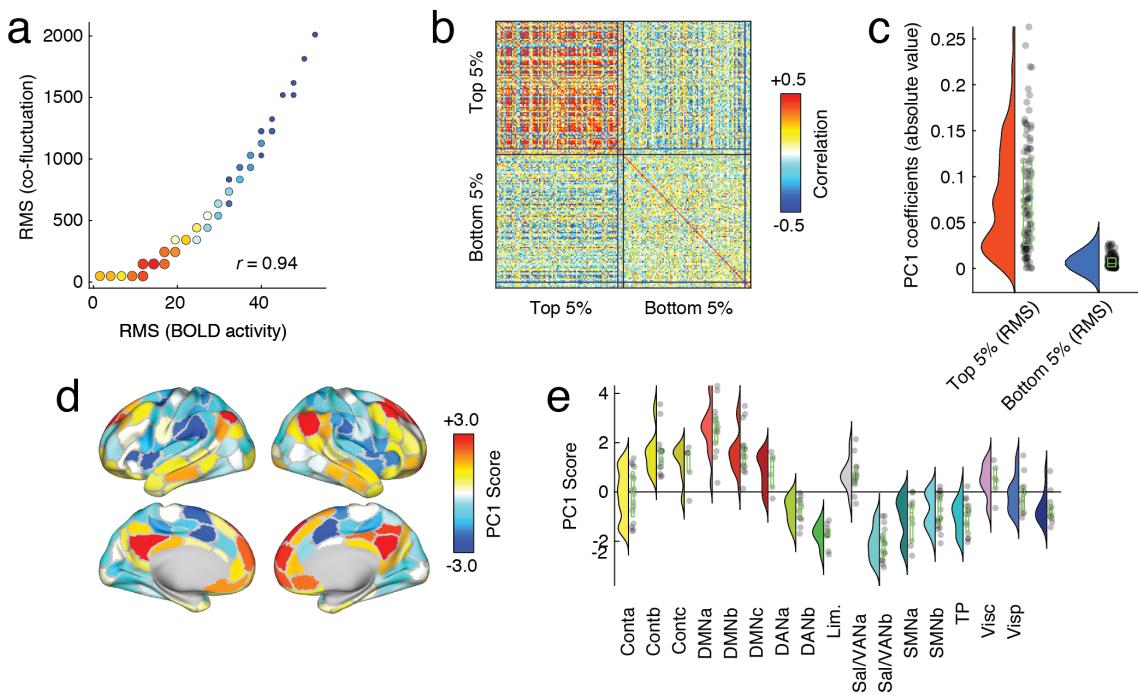


FIG. S9. Relationship of network co-fluctuations with BOLD fluctuations for MSC dataset. Here, we replicate results from the main text using the MSC data. Specifically, we relate high-amplitude co-fluctuations to fluctuations in fMRI BOLD activity. We subsequently demonstrate that the high-amplitude 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 mean 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 co-fluctuation events, which drive resting-state functional connectivity, are underpinned by instantaneous activation and deactivation of default mode and control network areas.