High-amplitude cofluctuations in cortical activity drive functional connectivity

Farnaz Zamani Esfahlani[a,1], Youngheun Jo[a,1], Joshua Faskowitz[a,b], Lisa Byrge[a], Daniel P. Kennedy[a,b,c], Olaf Sporns[a,b,c,d], and Richard F. Betzel[a,b,c,d,2]

[a]Department of Psychological and Brain Sciences, Indiana University, Bloomington, IN 47405; [b]Program in Neuroscience, Indiana University, Bloomington, IN 47405; [c]Cognitive Science Program, Indiana University, Bloomington, IN 47405; and [d]Network Science Institute, Indiana University, Bloomington, IN 47405

Edited by Marcus E. Raichle, Washington University in St. Louis, St. Louis, MO, and approved September 11, 2020 (received for review March 24, 2020)

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 cofluctuations 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 cofluctuations of network organization with fluctuations in the blood oxygen level-dependent (BOLD) time series. We show that surprisingly, only a small fraction of frames exhibiting the strongest cofluctuation 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. Finally, we demonstrate that cofluctuation amplitude synchronizes across subjects during movie watching and that high-amplitude frames carry detailed information about individual subjects (whereas low-amplitude frames carry little). 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 directions for future research.

functional connectivity | dynamics | time-varying connectivity | naturalistic stimuli

Resting-state functional connectivity (rsFC) refers to the correlation structure of functional magnetic resonance imaging (fMRI) blood oxygen level-dependent (BOLD) activity, usually estimated over the course of an entire scan session (1, 2). Interindividual differences in rsFC have been linked to variation in biological age (3, 4), cognitive state (5), and clinical status (6). Other studies have emphasized the dynamic nature of rsFC, using sliding window techniques to generate temporally blurred estimates of rsFC across time (7–9) and linking changes in network architecture to behavior (10, 11) and phenotypes (12, 13).

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?

In principle, these questions can be addressed using sliding-window approaches to track fluctuations in rsFC across time. However, the windowing procedure induces a blurring effect, making it impossible to localize time-varying connectivity in time and assess the contributions made by individual frames (14). On the other hand, methods like coactivation patterns (CAPs) allow for brain dynamics to be characterized at the resolution of single frames (15–19). Using these types of methods, previous studies have shown that coactivity patterns fluctuate across time, with the brain's system-level organization expressed only during a select set of time points (20). However, these approaches generally require the specification of a seed region or a threshold for determining what constitutes high-amplitude activity. Consequently, a comprehensive and mathematically precise explanation for how these coactivity patterns combine to give rise to rsFC over longer timescales remains elusive (21).

Here we address these questions using a mathematically exact decomposition of rsFC into its frame-wise contributions, explicitly linking instantaneous patterns of cofluctuation to rsFC over longer timescales (22, 23). We find that at rest, cofluctuations are bursty and occur intermittently as part of whole-brain cofluctuation events that are uncorrelated with respiration, cardiac cycle, and in-scanner motion. We then show that rsFC estimated using only high-amplitude frames is highly correlated with rsFC estimated over the entire scan session, indicating that rsFC and its system-level organization are driven by cofluctuations during relatively few frames. We then show that high-amplitude cofluctuations are underpinned by the activation of a particular spatial mode of brain activity in which default mode and control networks are anticorrelated with sensorimotor and attentional systems. We then present two careful examinations

Significance

Despite widespread applications, the origins of functional connectivity remain elusive. Here we analyze human functional neuroimaging data. We decompose resting-state functional connectivity across time to assess the contributions of moment-to-moment activity cofluctuations to the overall connectivity pattern. We show that functional connectivity is driven by a small number of high-amplitude frames. We show that these frames are underpinned by a specific mode of brain activity; that the topography of this mode gets modulated during in-scanner tasks; and that high-amplitude frames encode personalized, subject-specific information. In summary, our parameter-free method provides an exact mathematical link between functional connectivity and frame-wise cofluctuations, creating opportunities for studying both static and time-varying functional brain networks.

Author contributions: F.Z.E., Y.J., and R.F.B. designed research; F.Z.E., Y.J., and R.F.B. performed research; F.Z.E., Y.J., J.F., L.B., D.P.K., and O.S. contributed new reagents/analytic tools; F.Z.E., Y.J., and R.F.B. analyzed data; and F.Z.E., Y.J., J.F., L.B., D.P.K., O.S., and R.F.B. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1]F.Z.E. and Y.J. contributed equally to this work.

[2]To whom correspondence may be addressed. Email: rbetzel@indiana.edu.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2005531117/-/DCSupplemental.

First published October 22, 2020.

of high-amplitude cofluctuations. First, we demonstrate that time series of cofluctuation amplitude synchronize across subjects during movie watching, and second, we show that subjects' "fingerprints" are enhanced during high-amplitude frames compared to low-amplitude frames.

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 (24). 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 cofluctuation between those regions resolved at every moment in time (Fig. 1*A*). These edge time series can be analyzed directly to pinpoint both the magnitude and timing of cofluctuations between pairs of brain regions.

In *rsFC Is Driven by Short-Lived and High-Amplitude Cofluctuation Events* and in *High-Amplitude Frames Are Driven by Fluctuations of Task-Positive/Task-Negative Mode of Brain Activity*, we analyze cofluctuation time series constructed from functional imaging data acquired as part of the Human Connnectome Project (25, 26) (see *Materials and Methods* for details). All results reported in those sections generated using these data; we replicate these findings using a second dataset (27), with results reported in *SI Appendix*. In *Intersubject Synchrony of Whole-Brain Cofluctuation Amplitude during Passive Movie Watching* and in *High-Amplitude Cofluctuations Enhance Identifiability*, we analyze an independently acquired movie-watching dataset (28, 29) and data from the Midnight Scan Club (MSC) (27, 30), respectively.

rsFC Is Driven by Short-Lived and High-Amplitude Cofluctuation Events. While past studies have used sliding window methods to generate estimates of moment-to-moment fluctuations in rsFC (8, 9), the use of a windowing procedure results in a temporally blurred estimate of rsFC. This restricts the timescale of observations regarding dynamic changes in functional connectivity to the width of the time window, generally on the order of dozens of frames (approximately 1 min of real time, although we note that recently, several groups have explored methods for estimating time-varying FC at a frame-wise timescale, e.g., refs. 31 and 32). Here we address this limitation using cofluctuation times series.

When analyzed across the whole brain, we find that edge time series exhibit bursty behavior, such that the amplitude of cofluctuations (quantified by computing the root sum square [RSS]) moves around a mean value but is punctuated by brief, intermittent, and disproportionately large fluctuations (Fig. 1*B*). These high-amplitude frames are not directly related to cardiac and respiratory cycles, in-scanner head motion (*SI Appendix*, Fig. S1), and spectral properties of fMRI BOLD time series (*SI Appendix*, Fig. S2) and appear aperiodic with heavy-tailed distributions of size, duration, and intervals (*SI Appendix*, Fig. S3).

To better understand how instantaneous cofluctuations contribute to whole-brain rsFC, we isolated high-amplitude frames and compared them with low-amplitude episodes (top and bottom 5% in terms of cofluctuation amplitude; 60 frames for HCP; see *SI Appendix*, 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 high-amplitude frames than low-amplitude (within-sample t test, $p < 10^{-15}$; Fig. 1*C*). Next, we calculated the similarity of rsFC estimated during high-

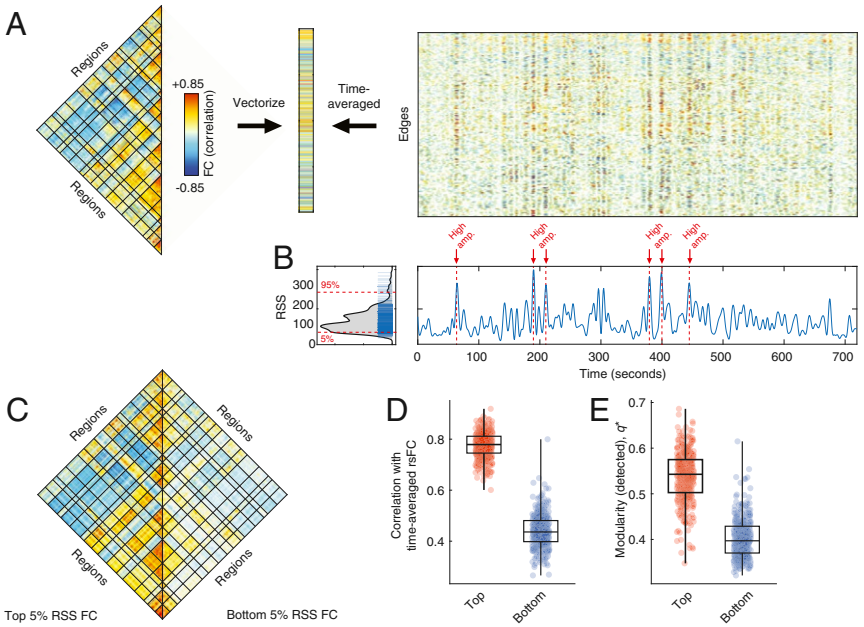

Fig. 1. Cofluctuation time series reveal bursty structure of resting-state functional connectivity. (*A*) We use a temporal unwrapping of the Pearson correlation to generate cofluctuation time series for every pair of brain regions (edges). The elements of the cofluctuation time series are the element-wise 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 a resting-state functional connectivity matrix. (*B*) We find that the cofluctuation time series contains moments in time where many edges collectively cofluctuate. We can identify these moments by calculating the RSS across all cofluctuation time series and plotting this value as a function of time. In *B* we label high- and low-amplitude frames. The distribution of edge cofluctuation 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 cofluctuation 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.

and low-amplitude episodes with respect to time-averaged rsFC estimated using the full time series. We found that the high-amplitude networks were highly correlated with rsFC ($r = 0.81 \pm 0.05$) while the low-amplitude networks were much less correlated ($r = 0.54 \pm 0.07$) and that these differences were highly significant (t test, $p < 10^{-15}$; Fig. 1*D*). We also performed an analogous comparison of network modularity (33), 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 high-amplitude networks ($q = 0.51 \pm 0.06$) compared to the networks estimated from low-amplitude frames ($q = 0.37 \pm 0.05$) (t test, $p < 10^{-15}$; Fig. 1*E*).

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

Collectively, these results suggest that rsFC, estimated over long timescales, is driven by a small number of brief, intermittent, and high-amplitude cofluctuations. 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 cofluctuations are only weakly related to the overall pattern of rsFC and correspond to less modular architectures.

High-Amplitude Frames 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 cofluctuations that occur during a relatively small number of frames. It remains unclear, however, whether high-amplitude frames are underpinned by a specific pattern of brain activity or whether they reflect contributions from multiple distinct patterns. Here we address this question directly, by investigating the patterns of brain activity that occur at the same time as high-amplitude frames.

As a first point of comparison, we calculated the RSS of both the cofluctuation 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 high-amplitude frames have an almost one-to-one correspondence with high-amplitude BOLD fluctuations (Fig. 2*A*). This relationship is expected; because cofluctuations 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 high-amplitude frames compared to low-amplitude frames, 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 high- and low-amplitude frames and computed between-subject and between-scan similarity (Fig. 2*B*). In general, activity during high-amplitude frames was more correlated across subjects compared to the activity patterns during low-amplitude frames (t test, $p < 10^{-15}$). To better understand what was driving these correlations, we performed a principal components analysis of the activity patterns during high- and low-amplitude frames, 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 high-amplitude frames compared to low-amplitude (t test, $p < 10^{-15}$; Fig. 2*C*), indicating that PC1 was descriptive of activity patterns during high-amplitude frames but less so for low-amplitude frames. 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. 2 *D* and *E*). We replicated these results in a second dataset (*SI Appendix*, Fig. S9).

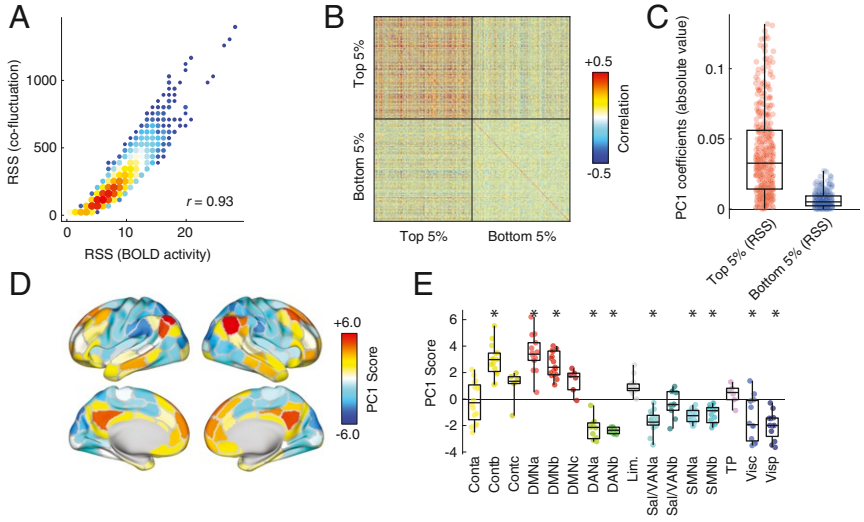

Fig. 2. Relationship of network cofluctuations with BOLD fluctuations. In *rsFC Is Driven by Short-Lived and High-Amplitude Cofluctuation Events* we demonstrated that resting-state functional connectivity could be explained on the basis of relatively few frames during which high-amplitude cofluctuations occurred. Here we relate those cofluctuation frames to BOLD activity fluctuations. We first calculate the RSS amplitude of BOLD activity at each time point and compare that to the amplitude of cofluctuations. (*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 cofluctuation 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* and *E*) the PC1 score corresponded to activity patterns that emphasized correlated fluctuations of default mode and control networks that were weakly or anticorrelated with fluctuations elsewhere in the brain. Asterisks indicate systems whose mean PC1 score was significantly greater (more positive or negative) than expected by chance (permutation test; FDR fixed at 5%; $p_{adjusted} = 0.018$). These observations suggest that high-amplitude cofluctuations, which drive resting-state functional connectivity, are underpinned by instantaneous activation and deactivation of default mode and control network areas.

These results suggest that underlying high-amplitude frames is a mode of brain activity whose spatial pattern resembles the traditional task-positive/task-negative division of the brain (34). 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 high-amplitude frames.

Intersubject Synchrony of Whole-Brain Cofluctuation Amplitude during Passive Movie Watching. In the previous sections we showed that rsFC can be viewed as an average of time-varying cofluctuations. We also showed that time-averaged rsFC is disproportionately impacted by high-amplitude frames that are, themselves, underpinned by a specific mode of brain activity and were not clearly related to motion or physiological artifacts. What, then, is the purpose of high-amplitude frames? Are they random cofluctuations, or are they related to fluctuations in an individual's brain/cognitive state? To address these questions, we explored the cofluctuation time series for a cohort of 29 subjects that were scanned multiple times at rest and while passively viewing complex, naturalistic stimuli (movies) (28).

Specifically, we computed edge time series for all subjects and scans for both conditions. From these edge time series, we estimated the cofluctuation amplitude across all node pairs. For a given scan, this procedure results in 29 time series (1 per subject) of identical length. We found that cofluctuation time series were correlated across subjects during movie watching (Fig. 3 *A* and *B*) but uncorrelated during rest (Fig. 3 *C* and *D*). We directly compared the distributions of intersubject correlations between conditions (all scans from the same condition pooled together), discovering that as expected, the mean intersubject correlation was greater during movie watching than at rest (permutation test, $p < 0.05$; Fig. 3*E*). Importantly, we found no difference between conditions for the overall amplitude of RSS values (permutation test, $p = 0.07$; Fig. 3*F*).

Next, we explored differences between movie watching and rest in terms of brain activity patterns during high- and low-amplitude frames. Our exploration consisted of two analyses. First, and as in the previous section, we extracted activity patterns during high- and low-amplitude frames (top and bottom 5% by cofluctuation amplitude) separately for the movie-watching and resting conditions. We then performed principal component analysis (PCA) on these matrices and retained the top PC score for each condition. Interestingly, these PC scores exhibited distinct topography; the movie-watching PC (Fig. 3*G*) emphasized activity in visual and dorsal attention networks, whereas the resting PC (Fig. 3*H*) recapitulated the pattern shown in the previous section, emphasizing a task-positive/task-negative mode of activity. To directly compare these two patterns, we computed their element-wise (region-wise) difference and grouped these differences by system (Fig. 3 *I* and *J*). As expected, we found statistically significant differences in the dorsal attention and visual systems (movie > rest; false discovery rate fixed at $q = 0.05$) and salience/ventral attention system (movie < rest). These differences are further evident when we plot the PCs against one another, revealing that these systems deviate from the identity line (Fig. 3*K*).

We note that another strategy for comparing movie-watching and resting-state conditions is to analyze them simultaneously by concatenating high-amplitude activity patterns from both conditions into a single matrix and jointly decomposing that matrix using PCA. This procedure results in modes of activity that are shared across both conditions. Here we retain both the first and second PCs (*SI Appendix*, Fig. S10 *A* and *C*), whose spatial

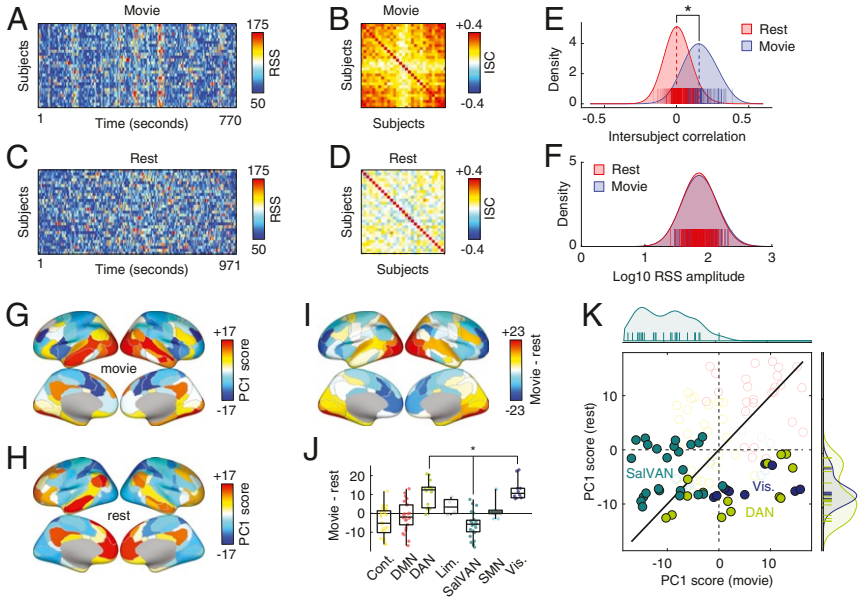

Fig. 3. Whole-brain cofluctuation amplitude synchronizes during passive movie watching. We compared cofluctuation time series during resting state and movie watching. For both conditions, we computed cofluctuation time series for 29 subjects. We show those time series in *A* (movie) and *C* (rest). We find that when subjects watch movies, their cofluctuation time series are synchronous, presumably due to the shared audiovisual stimulus. At rest, cofluctuation time series are asynchronous. We demonstrate this synchrony by computing the intersubject correlation matrix of subjects' cofluctuation time series. We show matrices for movie watching and rest in *B* and *D*, respectively. By comparing the elements of these matrices, we demonstrate statistically that movie watching leads to increased intersubject correlations. We show the distributions in *E*. We find, however, that the overall amplitude of fluctuations (RSS) is not statistically different from one condition to the other (*F*). To further contrast these two conditions, we repeated the analysis from *High-Amplitude Frames Are Driven by Fluctuations of Task-Positive/Task-Negative Mode of Brain Activity* to identify modes of brain activity that underpin high-amplitude frames. We find that the resting mode recapitulates the topographic distribution reported in the previous section (*H*), emphasizing a task-positive/task-negative division. During movie watching, however, the mode of activity emphasizes contributions of visual and dorsal attention networks (*G*). In *I*–*K*, we compare rest and movie-watching modes of activity more directly. *I* depicts the region-wise differences in modes, *J* groups those differences by system, and *K* presents them as a scatterplot, highlighting differences associated with visual, dorsal attention, and salience/ventral attention networks.

topography is similar to what we show in Fig. 3 *G* and *H*. As expected, we find differences between the two maps in terms of their PC coefficients, with movie-watching frames loading more strongly onto the first map (permutation test, $p < 0.05$; *SI Appendix*, Fig. S10*B*) and resting frames loading more strongly onto the second (permutation test, $p < 0.05$; *SI Appendix*, Fig. S10*D*).

Viewed collectively, these results complement our previous findings that cofluctuation time series are not clearly related to motion or physiological artifacts. Importantly, we demonstrate that subjects' cofluctuation time series synchronize when jointly presented with complex, time-varying, and naturalistic stimuli. This observation, combined with the topographic differences between movie-watching and resting-state activity during high-amplitude frames, strongly suggests that cofluctuation amplitude is at least in part modulated by subjects' cognitive states.

High-Amplitude Cofluctuations Enhance Identifiability. In the previous sections, we showed that high-amplitude frames contribute disproportionately to the brain's static FC patterns, shape its modular structure, are underpinned by a population-level mode of brain activity, and synchronize during viewing of naturalistic stimuli. Do they also enhance the identifiability of individual subjects? That is, does FC estimated during high-amplitude frames bear a stronger signature of a subject than FC estimated during low-amplitude frames?

To test this question, we calculated the cofluctuation amplitude at each time point (Fig. 4*A*, *i*). We then isolated the frames with the highest amplitude and estimated FC using only these frames (Fig. 4*A*, *ii*). Repeating this procedure for all subjects and scans resulted in a set of 100 feature vectors (10 subjects $\times$ 10 scans) that encode a subjects' FC patterns (Fig. 4*A*, *iii*). We then compute the $[100 \times 100]$ similarity matrix—otherwise known as the identifiability matrix—and compare the mean within-subject similarity to the mean between-subject similarity (Fig. 4*A*, *iv*). This measure—the differential identifiability—indicates how much more similar subjects' FC patterns are to themselves than to those of other subjects. We repeat this entire procedure using only low-amplitude frames (Fig. 4*A*, *v–vii*) and compare low- and high-amplitude differential identifiability.

In general, we find that intersubject similarity is greater during high-amplitude frames than compared to low-amplitude frames and over a range of thresholds (Fig. 4*B*). Importantly, we also find that the gap between within- and between-subject similarity, differential identifiability, is also greater when estimated using high-amplitude frames (Fig. 4 *C–E*). This observation suggests that high-amplitude frames may carry more individualized and distinguishable information about subjects than low-amplitude frames.

Discussion

Here we presented a general approach for temporally unwrapping Pearson correlations to generate time series of interregional cofluctuations 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. Our results also suggest that the cofluctuation patterns, at a coarse scale, capture cognitively relevant fluctuations in brain state and that high-amplitude frames encode signatures of an individual.

Decomposing Static FC into Cofluctuation Snapshots. Central to this paper is the observation that static FC can be nonparametrically decomposed into a series of time-varying snapshots, each of which expresses an instantaneous pattern of interregional cofluctuation. Critically, the average of these patterns across time is exactly equal to whole-brain static FC. This mathematical truism allows us to neatly assess the contributions of momentary

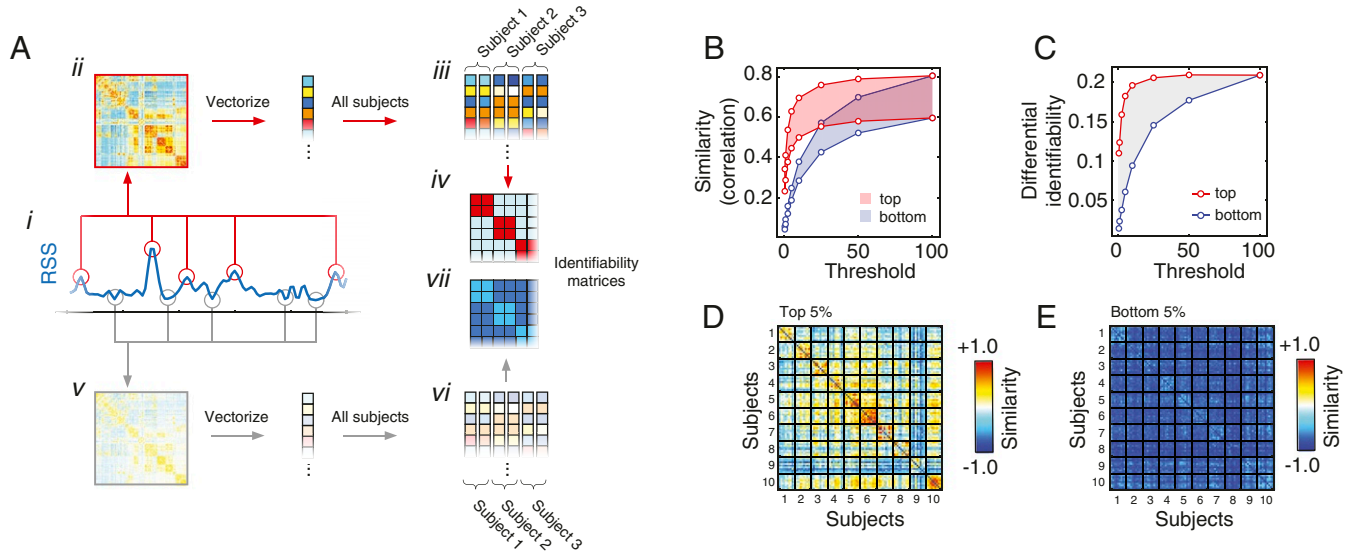

Fig. 4. Connectome fingerprints are strong during high-amplitude frames and weaker during low-amplitude frames. We investigated whether subject-specific features of rsFC were more prevalent during high- or low-amplitude cofluctuations. To address this question, we identified frames (time points) with the highest and lowest cofluctuation amplitude and estimated subjects FC using these data only. We then calculated the intersubject similarity matrix, i.e., the identifiability matrix. (*A*) Illustration of this general procedure, beginning by isolating high-amplitude time points (*i*), estimation of FC (*ii*), repeating this procedure for all subjects (*iii*), and estimating the intersubject similarity matrix (*iv*). An identical procedure was carried out for low-amplitude frames and is illustrated in *v–vii*. (*B*) We calculated the mean within- and between-subject similarity using both the top (red) and bottom (blue) frames, ordered by cofluctuation amplitude. For each set of frames, we produce two separate curves, one for the within-subject similarity and another for the between-subject similarity. The area between the curves is the differential identifiability, or the extent to which subjects' FC patterns are more similar to themselves than to FC estimated from other subjects. (*C*) We found that differential identifiability was always greater when FC was estimated using the top frames, ordered by amplitude. For the sake of visualization, we show identifiability matrices estimated using (*D*) high- and (*E*) low-amplitude frames.

cofluctuations to the overall FC pattern and to establish a clear link between FC and fluctuations in brain activity.

Our findings, which complement previous work (15–18, 20, 32), leverage a mathematically exact decomposition of FC into its frame-wise contributions to suggest that static FC is driven by contributions from relatively few time points, namely, those with the highest levels of cofluctuation amplitude. Frames with low levels of cofluctuation, on the other hand, contribute little. Because the cofluctuation time series are estimated at a temporal resolution of single frames, we directly compared high-amplitude frames with coincident patterns of brain activity. We established that at rest, high-amplitude cofluctuations occur in tandem with a specific mode of brain activity that emphasizes oppositional activation of sensorimotor and association cortex. Notably, similar patterns of activity have been reported in other state-based analyses of brain dynamics (39, 40), linking this pattern of activity to mentation during wakeful rest and working memory performance.

The observations reported here both clarify and challenge core assumptions concerning static FC and brain network dynamics (41). Specifically, our findings suggest that whole-brain FC follows a bursty trajectory through a high-dimensional state space, with extended periods of quietude punctuated by brief and intermittent events, whose timing is not clearly related to motion or physiological artifacts (42). This observation leads to several questions, the most important of which concerns the origins of high-amplitude frames. Are they spontaneous occurrences? Are they relevant, in any way, to ongoing cognitive processes? How individualized are high-amplitude frames?

Linking High-Amplitude Cofluctuations to Cognition and Individual Differences. We performed two separate analyses in order to help clarify the origins of high-amplitude frames. First, we performed a comparison of their structure at rest and during movie watching (28). While the amplitude of cofluctuations was not statistically different between conditions, we found that cofluctuations were correlated during movie watching, suggesting that high-amplitude frames may be driven by audiovisual features in movies. This finding supports the hypothesis that the timing of high-amplitude frames is linked to perception and processing of sensory information and further suggests that high-amplitude frames are not simply spontaneous occurrences. These observations open up possibilities for future studies that leverage the temporal structure of cofluctuation amplitudes to track changes in an individual's cognitive state across time.

Importantly, we also discover differences in the mode of brain activity underpinning high-amplitude cofluctuations in movie watching compared to rest. In particular, we found stronger expression of visual and dorsal attention networks, brain systems that one might hypothesize to play an important role in processing visual information and redirecting attentional resources while viewing complex naturalistic stimuli (43, 44). This finding also demonstrates that although high-amplitude cofluctuations can occur spontaneously in correlated, modular systems, their character and timing are modulated by time-varying sensory input, presenting an opportunity for future studies to comprehensively map the task-evoked topography of high-amplitude cofluctuations (45).

In our second analysis, we asked whether high-amplitude cofluctuations were personalized and idiosyncratic (27, 46, 47). To address this question, we estimated subjects' FC separately using high- and low-amplitude frames and compared these networks in terms of their differential identifiability—the extent to which the similarity of FC patterns was stronger within subjects than between (48). Surprisingly, we found that identifiability was significantly stronger during high-amplitude cofluctuations, suggesting that subject-specific information is expressed more strongly during those frames.

Collectively, these findings suggest that the structure of high-amplitude cofluctuations is highly organized. It tracks time-varying fluctuations in cognitive state and is deeply personalized. These are key observations with clear implications for the study of brain–behavior associations, clinical neuroscience, and phenotype discovery, where the ability to make inferences is limited by the amount of data available. Our results suggest that by capitalizing on the fact that high-amplitude cofluctuations carry more subject-specific information than low-amplitude cofluctuations, it may be possible to generate robust network-level biomarkers using a relatively small number of frames and reducing the amount of required data (49) [we note this concept is being explored with other imaging modalities (50)]. This approach may be especially useful in clinical and developmental neuroscience, which study populations with characteristics that generally prohibit the extended scan durations necessary for stably estimating FC (51). For instance, because high-amplitude cofluctuations carry more information about static FC than low-amplitude cofluctuations and because they encode identifiable features of subjects, implementing experimental paradigms that elicit large cofluctuations at greater frequency may obviate the need for long scan sessions and large quantities of data and yield superior estimates of FC.

System-Level Organization Emergences from the Structure of High-Amplitude Cofluctuations. Last, our findings hint at a crucial link between instantaneous fluctuations in activity and the organization of rsFC (31, 52). Many studies have found that the community structure of rsFC resembles known coactivation patterns, including task-evoked activity (53, 54). Here we proposed a strategy that enabled 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 high-amplitude frames and, in turn, whole-brain rsFC. While this mode made the greatest contribution, it is likely that other modes make nontrivial contributions as well. By extending the definition to include lower-amplitude fluctuations, we expect to find patterns of activity that correspond to other, well-known brain systems (16). Moreover, we speculate that these patterns likely recombine in different proportions as a function of task complexity and domain (53, 55) and across individuals (47).

We note that other studies have shown, using a wide range of approaches, that time-varying and static rsFC are related to one another and that this relationship depends, to a large extent, on high-amplitude network states. For instance, ref. 18 identified time points of high-amplitude activity using sliding-window methods and seed-based analyses to show that activity coincides with increased correlation (rsFC). Indeed, comparable results have been reported using similar methods (15, 17, 20, 36, 38, 42, 56–58). The principal finding of these studies is that high-amplitude activity is somehow related to stronger FC or the expression of particular brain systems.

While illuminating, these papers have some limitations. Notably, they describe relationships between FC and high-amplitude states but lack a mathematical mechanism for why this relationship exists. In other cases, the observations require user-defined parameters to determine what constitutes high-amplitude activity, to specify a seed voxel or region for computing FC, or the width of a sliding window. In contrast, our approach addresses these limitations directly. Our first finding is in agreement with the above-referenced papers but is precipitated from a mathematically precise decomposition of static FC into its exact frame-wise contributions. With this decomposition, we can determine how individual time points combine to give rise to patterns of time-averaged, static FC. This recombination requires no additional parameterization and no sliding window and is naturally compared for all pairwise

connections (not just seeds). In summary, the framework we present here serves to unify these previous observations by providing a mathematical framework to explicitly link instantaneous patterns of cofluctuation to static FC.

In future work, the proportion of variance explained by different patterns and other statistics related to high-amplitude frames, including the frequency with which they occur, may serve as potent correlates of cognitive and disease state. Because high-amplitude frames appear to drive the overall configuration of rsFC, we further speculate that their statistics may serve as important complements to traditional measures of rsFC.

Future Work. The approach developed here presents several exciting opportunities for future studies. These include investigating time-varying FC using cofluctuation patterns, which provide frame-wise estimates of network structure and circumvent limitations of sliding window approaches (7, 14, 59). Other possibilities include mapping the relationship of structural connectivity to regional fluctuations or interregional cofluctuations during high- and low-amplitude frames (60) and studying individual differences in cognitive, development, and disease state based on features extracted from high-amplitude frames, which we show provide more reliable estimates of subject-level networks.

A remaining open question concerns the neurobiology that shapes high-amplitude cofluctuations. On one hand, their infrequent occurrence could reflect a dynamic strategy for limiting the consumption of metabolic resources (9, 61). This theory is supported by previous reports demonstrating that metabolic activity is elevated within the default mode network (62–64), a system that overlaps closely with the dominant mode of activity we find underpins high-amplitude frames. On the other hand, high-amplitude frames are, to some extent, a mathematical necessity emerging in correlated, modular systems. For instance, a group of brain regions form a mutually correlated functional module, then by definition their activity (and coactivity) will follow similar time courses, with a tendency to fluctuate together at the same instants in time. However, the observation that high-amplitude cofluctuations synchronize during movie watching suggests that they are also related to some underlying psychological process (with a presumed neurobiological correlate). Future work should investigate the neurobiological underpinnings of high-amplitude cofluctuations in greater detail.

Importantly, the entire cofluctuation time series enterprise could be extended in several important ways, including by applying it to other imaging modalities, e.g., electrophysiological recordings (65–68) or fluorescence imaging data (69, 70). Additionally, it would be straightforward to calculate cofluctuation time series after partialing out the effects of activity from other regions in the brain (24) or to investigate temporal dependencies and lags between brain regions (71, 72).

Finally, the approach developed here decomposes functional connections into their exact frame-wise contributions. We speculate that this decomposition might offer a selective means of addressing in-scanner motion, e.g., by identifying and censoring time points impacted by motion on an edge-wise basis (73, 74).

Relationship with Existing Approaches. We note that the analysis of cofluctuation time series is conceptually similar to several existing methods (20, 36–38) or in some cases even builds upon shared mathematical machinery (75). For instance, multiplication of temporal derivatives (MTDs) (76) calculates the element-wise products using differenced activity time series for all pairs of nodes. These time series are then convolved with a kernel to generate smooth estimates of time-varying nFC. Although similar, our approach relies on untransformed activity to estimate edge time series, thereby preserving the relationship between static nFC and the mean value of each edge time series. Furthermore, our approach omits the smoothing step, making it, in principle, capable of detecting fluctuations in network structure over shorter timescales compared to MTDs. Another related method is CAPs (15, 16), which extracts and clusters voxel- or vertex-level activity during high-activity frames. Because a voxel can be coactive under different contexts, the cluster centroids spatially overlap with one another.

Although these approaches arrive at similar conclusions, they possess distinct advantages and disadvantages that make some methods uniquely well suited for testing specific hypotheses and research questions. For instance, MTDs and the analysis of cofluctuation time series presented here are appropriate for tracking patterns of connectivity across time. In the case of cofluctuation time series, which are mathematically related to the static FC pattern, our approach is especially well suited for assessing the contributions of frame-wise cofluctuation patterns to the brain's overall FC (we note that this relationship, to our knowledge, has not been previously discussed in the extant literature). CAPS and innovation-driven CAPS on the other hand, are better suited for studying activity patterns and tracking their cooccurrences across time. In principle, a systematic and careful comparison of these methods could be carried out in future work.

Finally, we note that our approach incorporates elements of both CAPs and MTDs. Like MTDs and other sliding-window methods, our approach yields a time series of node-by-node matrices, each of which encodes pairwise relationships among brain regions across time. In the case of sliding-window methods, the elements in each matrix represent estimates of instantaneous correlations, which requires first estimating variances (a second moment) from a limited set of observations, which can lead to inaccurate and noisy FC inference, especially when the number of observations is small (narrow windows) (14, 77). Instead, the matrices generated using our approach encode instantaneous cofluctuation magnitudes, like CAPs. That is, each element indicates the magnitude with which the activity of two brain regions is instantaneously deflecting in the same direction. Importantly, deflections are calculated with respect to means and variances estimated using the full set of observations (the entire scan). Consequently, the time-varying estimates of cofluctuations reported here may be less sensitive to noise than sliding-window methods.

Additionally, it is important to note that while our approach and CAPs are similar in some ways, e.g., they both operate on single-frame timescales, our method is nonetheless distinct and has unique advantages. Most importantly, our method is built upon a mathematically exact decomposition of static rsFC into its frame-wise contributions. This decomposition enables us to quantify, precisely, how individual time points impact static rsFC. Here we use this property to demonstrate that a relatively small number of frames are necessary to explain rsFC [a finding that has been reported elsewhere (15, 17, 36)], to show that cofluctuations become synchronized during movie watching, and that high-amplitude frames can be used to enhance subjects' connectome-based fingerprints. Second, our method is parameter-free; the decomposition does not depend on a specification of ad hoc thresholds for high-amplitude activity nor does it require that we select a seed region or brain system in advance. Rather, our decomposition method considers all activity levels and the entire network simultaneously. As noted earlier, we envision this approach being useful for future studies of time-varying FC, for generating more sensitive and subject-specific biomarkers, and in conjunction with non-fMRI imaging modalities.

Limitations. One of the most key limitations concerns the calculation and interpretation of cofluctuation time series. The procedure for calculating edge time series begins by z-scoring each brain region's activity time series. This procedure, however, is only appropriate if the sample mean and SD are temporally invariant (78). If there is a sustained increase or decrease in activity, e.g., the effect of a blocked task, then the z-scoring procedure can result in a biased mean and SD resulting in poor estimates of fluctuations in activity. To minimize the likelihood of this

occurring, we focused on resting-state and movie-watching data rather than blocked tasks. In future work, investigation of task-evoked cofluctuations could be investigated by employing already common preprocessing steps, e.g., constructing task regressors to remove the first-order effect of tasks on activity (79).

Conclusion. In conclusion, our study discloses 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 coactivity. Our methodological framework is readily applicable to other imaging datasets and recording modalities, including observations at neuronal scales, enabling the study of neural coactivity at unprecedented temporal resolution.

Materials and Methods

Datasets. We analyzed three separate datasets. Specifically, we focused on resting-state data from both The Human Connectome Project (HCP) and MSC. These data were processed similarly, the details of which are described in this section. The third dataset, which has been analyzed elsewhere (28, 49, 80), 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 processing pipelines resulted in parcellations of cortex into $N = 200$ parcels (nodes) (81) in the case of the HCP and MSC datasets and $N = 114$ parcels (52) in the case of the Indiana University dataset. The processing pipelines are described in detail in *SI Appendix*.

Cofluctuation 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 or parcels are functionally connected to each other. By far the most common measure of statistic dependence is the Pearson correlation coefficient. Let $\mathbf{x}_i = [x_i(1), \ldots, x_i(T)]$ and $\mathbf{x}_j = [x_j(1), \ldots, x_j(T)]$ 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 SD. 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 of 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 cofluctuation 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 tth 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 cofluctuations between i and j. An analogous vector can easily be calculated for every pair of parcels (network nodes), resulting in a set of cofluctuation (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 (82). Intuitively, it attempts to decompose a network into nonoverlapping subnetworks such that the observed density of connections within subnetworks 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 $\delta(g_i, g_j)$ 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 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 (33):

$$q^* = q^+ + \frac{v^-}{v^+ v^-} q^-, \quad [2]$$

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

Differential Identifiability. Let A_1 be an $N \times N$ FC matrix. We can represent this matrix as an $M = N \times (N-1)/2$-dimensional vector by extracting its upper triangle elements. We can assess the similarity of two matrices, A_1 and A_2, by computing the similarity of their vector representations. Suppose we had multiple scans from multiple individuals. Let I_{self} and I_{others} be the average within- and between-subject similarity. Differential identifiability, then, is simply $I_{diff} = I_{self} - I_{diff}$ (48). Intuitively, the larger the value of I_{diff}, the stronger the population-level fingerprint (46).

Code Availability. Code and example data for generating and analyzing edge time series are available in Github (https://github.com/brain-networks/edge-ts).

Data Availability. Two of the imaging datasets come from publicly available, open-access repositories. HCP data can be accessed freely via https://db.humanconnectome.org/ after signing a data use agreement. MSC data can be accessed via OpenfMRI at https://openneuro.org/datasets/ds000224/versions/1.0.3. Postprocessed data from the Indiana University study data and code have been deposited in Figshare and GitHub (postprocessed data are available at https://doi.org/10.6084/m9.figshare.12971162).

ACKNOWLEDGMENTS. R.F.B. and F.Z.E. 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 (grants R01MH110630 and R00MH094409 to D.P.K. and T32HD007475 Postdoctoral Traineeship to L.B.).

1. R. C. Craddock *et al.*, Imaging human connectomes at the macroscale. *Nat. Methods* **10**, 524–539 (2013).
2. B. P. Rogers, V. L. Morgan, A. T. Newton, J. C. Gore, Assessing functional connectivity in the human brain by fMRI. *Magn. Reson. Imaging* **25**, 1347–1357 (2007).
3. S. Gu *et al.*, Emergence of system roles in normative neurodevelopment. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 13681–13686 (2015).
4. R. F. Betzel *et al.*, Changes in structural and functional connectivity among resting-state networks across the human lifespan. *Neuroimage* **102**, 345–357 (2014).
5. M. W. Cole, D. S. Bassett, J. D. Power, T. S. Braver, S. E. Petersen, Intrinsic and task-evoked network architectures of the human brain. *Neuron* **83**, 238–251 (2014).
6. A. Fornito, A. Zalesky, M. Breakspear, The connectomics of brain disorders. *Nat. Rev. Neurosci.* **16**, 159–172 (2015).
7. R. M. Hutchison *et al.*, Dynamic functional connectivity: Promise, issues, and interpretations. *Neuroimage* **80**, 360–378 (2013).
8. E. A. Allen *et al.*, Tracking whole-brain connectivity dynamics in the resting state. *Cereb. Cortex* **24**, 663–676 (2014).
9. A. Zalesky, A. Fornito, L. Cocchi, L. L. Gollo, M. Breakspear, Time-resolved resting-state brain networks. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 10341–10346 (2014).
10. A. Kucyi, K. D. Davis, Dynamic functional connectivity of the default mode network tracks daydreaming. *Neuroimage* **100**, 471–480 (2014).
11. P. Barttfeld *et al.*, Signature of consciousness in the dynamics of resting-state brain activity. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 887–892 (2015).
12. E. Damaraju *et al.*, Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia. *Neuroimage Clin.* **5**, 298–308 (2014).

13. Ü. Sakoğlu *et al.*, A method for evaluating dynamic functional network connectivity and task-modulation: Application to schizophrenia. *MAGMA.* **23**, 351–366 (2010).
14. R. Hindriks *et al.*, Can sliding-window correlations reveal dynamic functional connectivity in resting-state fMRI?. *Neuroimage* **127**, 242–256 (2016).
15. X. Liu, J. H. Duyn, Time-varying functional network information extracted from brief instances of spontaneous brain activity. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 4392–4397 (2013).
16. F. I. Karahanoğlu, D. Van De Ville, Transient brain activity disentangles fMRI resting-state dynamics in terms of spatially and temporally overlapping networks. *Nat. Commun.* **6**, 7751 (2015).
17. E. Tagliazucchi, P. Balenzuela, D. Fraiman, D. R. Chialvo, Criticality in large-scale brain fMRI dynamics unveiled by a novel point process analysis. *Front. Physiol.* **3**, 15 (2012).
18. N. Petridou, C. C. Gaudes, I. L. Dryden, S. T. Francis, P. A. Gowland, Periods of rest in fMRI contain individual spontaneous events which are related to slowly fluctuating spontaneous activity. *Hum. Brain Mapp.* **34**, 1319–1329 (2013).
19. C. C. Gaudes, N. Petridou, S. T. Francis, I. L. Dryden, P. A. Gowland, Paradigm free mapping with sparse regression automatically detects single-trial functional magnetic resonance imaging blood oxygenation level dependent responses. *Hum. Brain Mapp.* **34**, 501–518 (2013).
20. T. W. Allan *et al.*, Functional connectivity in MRI is driven by spontaneous bold events. *PloS One* **10**, e0124577 (2015).
21. M. G. Preti, T. A. W. Bolton, D. Van De Ville, The dynamic functional connectome: State-of-the-art and perspectives. *Neuroimage* **160**, 41–54 (2017).
22. J. Faskowitz, F. Zamani Esfahlani, Y. Jo, O. Sporns, R. F. Betzel, Edge-centric functional network representations of human cerebral cortex reveal overlapping system-level architecture. bioRxiv:799924 (13 October 2019).
23. Y. Jo *et al.*, The diversity and multiplexity of edge communities within and between brain systems. https://doi.org/10.1101/2020.05.05.067777 (6 May 2020).
24. S. M. Smith *et al.*, Network modeling methods for fMRI. *Neuroimage* **54**, 875–891 (2011).
25. D. C. Van Essen *et al.*, The Wu-minn human connectome project: An overview. *Neuroimage* **80**, 62–79 (2013).
26. D. C. Van Essen *et al.*; Wu-Minn HCP Consortium, Data from "Human connectome project resting-state data." ConnectomeDB. https://db.humanconnectome.org/. Accessed 1 September 2019.
27. E. M. Gordon *et al.*, Precision functional mapping of individual human brains. *Neuron*, **95**, 791–807 (2017).
28. R. F. Betzel, L. Byrge, F. Z. Esfahlani, D. P. Kennedy. Temporal fluctuations in the brain's modular architecture during movie-watching. *Neuroimage* **213**, 116687 (2020).
29. R. Betzel, D. Kennedy, L. Byrge, Data from "Resting-state and movie-watching data." Figshare. https://doi.org/10.6084/m9.figshare.12971162. Deposited 17 September 2020.
30. E. M. Gordon *et al.*, Data from "The midnight scan club (MSC) dataset." OpenNeuro. https://openneuro.org/datasets/ds000224/versions/1.0.3. Accessed 1 September 2019.
31. M. Pedersen, A. Omidvarnia, A. Zalesky, G. D. Jackson. On the relationship between instantaneous phase synchrony and correlation-based sliding windows for time-resolved fMRI connectivity analysis. *Neuroimage* **181**, 85–94 (2018).
32. J. Cabral *et al.*, Exploring mechanisms of spontaneous functional connectivity in MEG: How delayed network interactions lead to structured amplitude envelopes of band-pass filtered oscillations. *Neuroimage* **90**, 423–435 (2014).
33. M. Rubinov, O. Sporns, Weight-conserving characterization of complex functional brain networks. *Neuroimage* **56**, 2068–2079 (2011).
34. M. Mennes *et al.*, Inter-individual differences in resting-state functional connectivity predict task-induced bold activity. *Neuroimage* **50**, 1690–1701 (2010).
35. E. Tagliazucchi, M. Siniatchkin, H. Laufs, D. R. Chialvo, The voxel-wise functional connectome can be efficiently derived from co-activations in a sparse spatio-temporal point-process. *Front. Neurosci.* **10**, 381 (2016).
36. I. Cifre, M. Zarepour, S. G. Horovitz, S. Cannas, D. R. Chialvo, On why a few points suffice to describe spatiotemporal large-scale brain dynamics. arXiv:1707.00759 (3 July 2017).
37. D. Gutierrez-Barragan, M. A. Basson, S. Panzeri, A. Gozzi, Infraslow state fluctuations govern spontaneous fMRI network dynamics. *Curr. Biol.* **29**, 2295–2306 (2019).
38. X. Zhang, W.-J. Pan, S. D. Keilholz, The relationship between bold and neural activity arises from temporally sparse events. *Neuroimage* **207**, 116390 (2020).
39. E. J. Cornblath *et al.*, Temporal sequences of brain activity at rest are constrained by white matter structure and modulated by cognitive demands. *Commun Biol.* **3**, 261 (2020).
40. T. Karapanagiotidis *et al.*, The psychological correlates of distinct neural states occurring during wakeful rest. bioRxiv:2019.12 (12 March 2020).
41. G. Deco, V. K. Jirsa, A. R. McIntosh, Emerging concepts for the dynamical organization of resting-state activity in the brain. *Nat. Rev. Neurosci.* **12**, 43–56 (2011).
42. W. H. Thompson, P. Fransson, Bursty properties revealed in large-scale brain networks with a point-based method for dynamic functional connectivity. *Sci. Rep.* **6**, 39156 (2016).
43. S. Sonkusare, M. Breakspear, C. Guo, Naturalistic stimuli in neuroscience: Critically acclaimed. *Trends Cogn. Sci.* **23**, 699–714 (2019).
44. U. Hasson, Y. Nir, I. Levy, G. Fuhrmann, R. Malach, Intersubject synchronization of cortical activity during natural vision. *Science* **303**, 1634–1640 (2004).
45. M. W. Cole *et al.*, Multi-task connectivity reveals flexible hubs for adaptive task control. *Nat. Neurosci.* **16**, 1348–1355 (2013).
46. E. S. Finn *et al.*, Functional connectome fingerprinting: Identifying individuals using patterns of brain connectivity. *Nat. Neurosci.* **18**, 1664–1671 (2015).
47. C. Gratton *et al.*, Functional brain networks are dominated by stable group and individual factors, not cognitive or daily variation. *Neuron* **98**, 439–452 (2018).
48. E. Amico, J. Goñi, The quest for identifiability in human functional connectomes. *Sci. Rep.* **8**, 8254 (2018).
49. L. Byrge, D. P. Kennedy, High-accuracy individual identification using a "thin slice" of the functional connectome. *Netw. Neurosci.* **3**, 363–383 (2019).
50. J. A. Nielsen, R. W. Mair, J. T. Baker, R. L. Buckner, Precision brain morphometry: Feasibility and opportunities of extreme rapid scans. bioRxiv:530436 (26 January 2019).
51. T. O. Laumann, *et al.*, Functional system and areal organization of a highly sampled individual human brain. *Neuron* **87**, 657–670 (2015).
52. B. T. Yeo *et al.* The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.* **106**, 1125–1165 (2011).
53. S. M. Smith *et al.*, Correspondence of the brain's functional architecture during activation and rest. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 13040–13045 (2009).
54. N. A. Crossley *et al.*. Cognitive relevance of the community structure of the human brain functional coactivation network. *Proc. Natl. Acad. Sci. U.S.A.* **110**. 11583–11588 (2013).
55. T. Yarkoni, R. A. Poldrack, D. C. Van. Essen, T. E. Nichols, D. W. Tor, Large-scale automated synthesis of human functional neuroimaging data. *Nat. Methods.* **8**, 665–670 (2011).
56. R. F. Betzel, M. Fukushima, H. Ye, X.-N. Zuo, O. Sporns, Dynamic fluctuations coincide with periods of high and low modularity in resting-state functional brain networks. *Neuroimage* **127**, 287–297 (2016).
57. W. H. Thompson, F. Peter, The mean–variance relationship reveals two possible strategies for dynamic brain connectivity analysis in fMRI. *Front. Hum. Neurosci.* **9**, 398 (2015).
58. J. Cabral *et al.*, Cognitive performance in healthy older adults relates to spontaneous switching between states of functional connectivity during rest. *Sci. Rep.* **7**, 5135 (2017).
59. D. J. Lurie *et al.*, Questions and controversies in the study of time-varying functional connectivity in resting fMRI. *Netw. Neurosci.* **4**, 30–69 (2020).
60. L. E. Suárez, D. R. Markello, R. F. Betzel, B. Misic, Linking structure and function in macroscale brain networks. *Trends Cogn. Sci.* **24**, 302–315 (2020).
61. S. B. Laughlin, J. T. Sejnowski, Communication in neuronal networks. *Science* **301**, 1870–1874 (2003).
62. M. E. Raichle *et al.*, A default mode of brain function. *Proc. Natl. Acad. Sci. U.S.A.* **98**, 676–682 (2001).
63. M. E. Raichle, A. Z. Snyder, A default mode of brain function: A brief history of an evolving idea. *Neuroimage* **37**, 1083–1090 (2007).
64. S. Passow *et al.*, Default-mode network functional connectivity is closely related to metabolic activity. *Hum. Brain Mapp.* **36**, 2027–2038 (2015).
65. R. F. Betzel *et al.*, Structural, geometric and genetic factors predict interregional brain connectivity patterns probed by electrocorticography. *Nat. Biomed. Eng.***3**, 902–916 (2019).
66. A. N. Khambhati, K. A. Davis, T. H. Lucas, B. Litt, D. S. Bassett. Virtual cortical resection reveals push-pull network control preceding seizure evolution. *Neuron* **91**, 1170–1182 (2016).
67. S. P. Burns *et al.*, Network dynamics of the brain and influence of the epileptic seizure onset zone. *Proc. Natl. Acad. Sci. U.S.A.*, **111**, E5321–E5330 (2014).
68. B. J. He, A. Z. Snyder, J. M. Zempel, M. D. Smyth, M. E. Raichle. Electrophysiological correlates of the brain's intrinsic large-scale functional architecture. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 16039–16044 (2008).
69. K. Mann, C. L. Gallen, T. R. Clandinin, Whole-brain calcium imaging reveals an intrinsic functional network in drosophila. *Curr. Biol.* **27**, 2389–2396 (2017).
70. R. Betzel, Organizing principles of whole-brain functional connectivity in zebrafish larvae. bioRxiv:496414 (15 December 2018).
71. A. Mitra, M. E. Raichle, How networks communicate: Propagation patterns in spontaneous brain activity. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **371**, 20150546 (2016).
72. A. Mitra, A. Z. Snyder, B. Tyler, M. E. Raichle, Lag threads organize the brain's intrinsic activity. *Proc. Natl. Acad. Sci. U.S.A.* **112**, E2235–E2244 (2015).
73. J. D. Power, K. A. Barnes, A. Z. Snyder, B. L. Schlaggar, S. E. Petersen, Spurious but systematic correlations in functional connectivity mri networks arise from subject motion. *Neuroimage* **59**, 2142–2154 (2012).
74. J. D. Power *et al.*, Methods to detect, characterize, and remove motion artifact in resting state fMRI. *Neuroimage*, **84**, 320–341 (2014).
75. E. S. B. van Oort *et al.*, Functional parcellation using time courses of instantaneous connectivity. *Neuroimage*, **170**, 31–40 (2018).
76. J. M. Shine *et al.*, Estimation of dynamic functional connectivity using multiplication of temporal derivatives. *Neuroimage* **122**, 399–407 (2015).
77. N. Leonardi, D. Van De Ville, On spurious and real fluctuations of dynamic functional connectivity during rest. *Neuroimage* **104**, 430–436 (2015).
78. R. Liegeois, T. O. Laumann, A. Z. Snyder, J. Zhou, B. T. T. Yeo, Interpreting temporal fluctuations in resting-state functional connectivity MRI. *Neuroimage* **163**, 437–455 (2017).
79. MI. W. Cole *et al.*, Task activations produce spurious but systematic inflation of task functional connectivity estimates. *Neuroimage* **189**, 1–18 (2019).
80. L. Byrge, D. P. Kennedy, Accurate prediction of individual subject identity and task, but not autism diagnosis, from functional connectomes. *Hum Brain Mapp.* **41**, 2249–2262 (2020).
81. A. Schaefer *et al.*, Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cereb. Cortex* **28**, 3095–3114 (2017)
82. M. E. J. Newman, M. Girvan, Finding and evaluating community structure in networks. *Phys. Rev. E* **69**, 026113 (2004).