

Resting-State fMRI Dynamics and Null Models: Perspectives, Sampling Variability, and Simulations

Robyn Miller¹, Tulay Adalı³, Yuri Levin-Schwarz³, Vince Calhoun^{1,2}

¹*The Mind Research Network*

²*The University of New Mexico, Department of Electrical and Computer Engineering*

³*The University of Maryland Baltimore County, Department of Computer Science and Electrical Engineering*

Abstract: Studies of resting state functional MRI (rs-fMRI) are increasingly focused on “dynamics”, or on those properties of brain activation that manifest and vary on timescales shorter than the scan’s full duration. This shift in focus has led to a flurry of interest in developing hypothesis testing frameworks and null models applicable to the dynamical setting. Thus far however, efforts in this direction have been weakened by a number of crucial shortcomings, several of which are outlined and discussed in this short paper. The discussion below is important because the weaknesses exhibited by many of the recently proposed models are not merely quantitative or parametric. Rather they are fundamental, falling into categories that ultimately neutralize the ability of the resulting models to play a scientifically clarifying role, while simultaneously injecting distracting new vectors of murkiness and confusion into a young but promising research enterprise these efforts are surely intended to support. One key point emphasized in this paper is that since functionally relevant temporal variations in brain activation are undeniably continual in living human rs-fMRI subjects (e.g., the period of time during which all neural activity throughout a human brain carries no functional relevance whatsoever is something whose existence, even during stroke or seizure, seems highly improbable), “absence of brain dynamics” as a null hypothesis for fMRI imaging studies is neither scientifically interesting nor well-defined. As we learn more about the phenomenon of functionally relevant brain dynamics (FRBD) and its imaging correlates, scientifically meaningful null hypotheses and well-tuned null models will naturally emerge. In the course of our discussion, we revisit the important concept of stationarity, what it is, how to assess it from multiple realizations of a process and finally attempt to parse through the benefits and limitations of applying this property in models of rs-fMRI data. We hope that the discussions herein are useful, and promote thoughtful consideration of these important issues.

1. Introduction

Studies of resting state functional magnetic resonance imaging (rs-fMRI) have been increasingly focused on properties of functional activation that manifest and vary on timescales shorter than the full duration of the scan¹. This approach is a more natural way to analyze such data as we know the brain is a highly dynamic organ. In resting-state studies, we do not have the benefit of external indicators of either behavior or functional brain processes, other than the narrow case of subject wakefulness which can be studied using simultaneous electroencephalogram (EEG) recordings²⁻⁴. In the case of resting data, we are seeking evidence in the scan itself of shifts in a subject's cognitive focus, emotional state, attention or consciousness level: e.g., we are attempting to locate the temporal and correlative signatures of complex internal ("ecologically authentic") sequences of mental tasks.

The growing focus on examining "dynamics" in rs-fMRI has spawned a number of efforts^{1,5-7} to formulate hypothesis testing frameworks aimed at statistically distinguishing variations (univariate, multivariate and relational/correlative) in network behavior that arise from reconfigurations of the subject's *cognitive, attentional, sensory, emotional* (CASE) state from those that have a high likelihood of occurring when a subject's brain is not undergoing or adapting to CASE shifts. The work on null models and hypothesis testing (e.g., ⁸⁻¹¹) frameworks for shorter timescale analysis of network behavior and dynamics has produced some preliminary insights, but remains rather muddled with regard to:

- 1) Creating an appropriate null space of network timecourses lacking the variations consistent with actual shifts in mental functioning, e.g., that
 - a. there is no task, experimental or ecological condition whose signature presents as a type of epochal variation¹ generically observable in that null space, and
 - b. the timeseries features that occur most rarely in this space are exactly those that are most strongly consistent with a brain undergoing shifts in CASE state.
- 2) Developing a metric that is sensitive to dynamic connectivity, i.e., that with some degree of specificity rises in the presence of those univariate or multivariate timeseries variations that are most plausibly connected with actual shifts in mental functioning
- 3) Avoiding the scenario of self-contradictory null models, i.e., distributional tails of a non-dynamic null model are not easily replicated in explicitly dynamic null models.

In this short position paper, we would like to address these points in particular as they relate to a recent paper ⁹. Our concerns are applicable well beyond the scope of ⁹ but this

¹ The term "epoch" has been used in this paper to refer to a duration, time period, and not necessarily of repetitive nature.

recent paper motivates us to discuss some of the inherent challenges of dynamic fMRI analyses, and helps illustrate the key pitfalls of employing naive null models in efforts to detect the presence of ubiquitous, complex and poorly understood phenomena such as functionally-relevant brain dynamics. We hope our discussion clarifies the implications of⁹, while also conveying the challenges of the project its authors undertook.

2. Modeling Functionally Relevant Brain Dynamics vs. Modeling Their Absence (Data Models vs. Null Models)

The relationship between well-constructed models of observed data and null models for a given phenomenon depends largely on the relationship between observed/observable data and the phenomenon being investigated. In situations such as those presented by the study of resting-state functionally relevant brain dynamics (rs-FRBD) using human subject fMRI data, where the null hypothesis is that a certain phenomenon is not present but the empirical data being modeled (in this case real fs-fMRI data) happens to be data in which this phenomenon is continually present, then the goals of building a null model for the phenomenon and accurately modeling the data can diverge significantly. Though there are many cases where highly accurate models of the data are also appropriate models of a setting in which the null hypothesis applies, e.g., they are also good null models. In the fMRI setting, since we tend to correct for motion artifacts, then a good model of this data that might, for example, be useful for identifying new scans that should be examined for motion contamination.

One of the challenges in the field of brain dynamics is that it is difficult to generate null models as the phenomena of interest are rather poorly defined and are also continual in the material being recorded from. Thus, modeling this data with high fidelity is not going to be the same as producing a null model for brain dynamics. In fact, a valid null simulation model of multivariate signals lacking features corresponding to rs-FRBD would by necessity diverge from actual scans observed in living people. Depending on what turn out to be the most reliable timeseries indicators of CASE-driven brain dynamics, it is possible that some valid null model might exhibit some similarities with the observed data (or output from good models of that data). However, a model built on the null hypothesis of no brain dynamics would by necessity only rarely create realizations that look similar to real rs-fMRI multivariate timeseries.

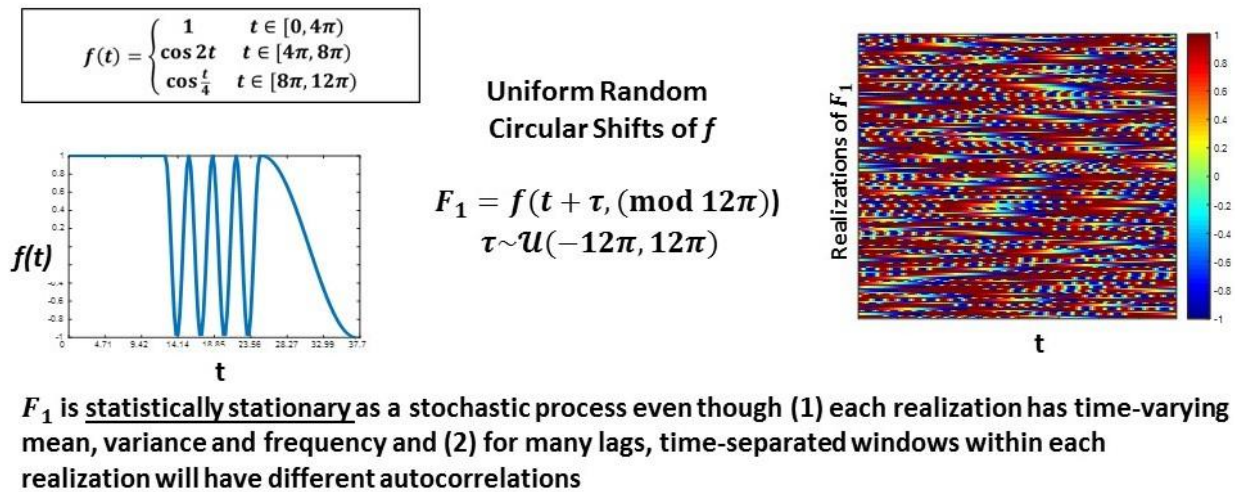
2.1. Statistical Stationarity, Gaussianity and rs-Brain Dynamics

Statistical stationarity is defined through the invariance of its joint probabilistic distribution across any number of samples, to any time shift, and as such, with real data, it can be only inferred using multiple realizations of a given stochastic process. A process is stationary if the statistical characterization computed over N realizations of any k -tuple of timepoints of length j $\{t_{j_1}, \dots, t_{j_k}\}$ and their τ -translates $\{t_{j_1} + \tau, \dots, t_{j_k} + \tau\}$ converge to identical values as $N \rightarrow \infty$. A practical way to infer stationarity is by estimating a finite set of moments as they are easier to compute than a full probability distribution. Of course, it is important to note that any definition of stationarity depends on the interval over which

it is evaluated. For a given interval over which stationarity holds it is quite possible for strong nonstationarities to manifest over shorter intervals as we demonstrate next. Matching white Gaussian noise to a template spectrum actually produces every possible signal with the given time-averaged spectral content given by the target. Statistical stationarity without other explicit constraints on the process does not imply that individual realizations of the univariate or multivariate timeseries (e.g., for fMRI these are individual subjects) are not featuring pronounced temporal epochs (see **Figure 1**). More specifically, even white Gaussian noise spectrally matched to a narrow-band spectral template can be markedly epochal on an average individual realization (see **Figure 2**).

The simulated null model of ⁹ is built from spectrally and covariance constrained multivariate Gaussian processes (SCC Gaussians). The use of an SCC Gaussian timeseries as a null model for functionally-relevant brain dynamics (FRBD) rests on an implicit assumption that (in the case of a 95% confidence level) fewer than 5% of the signals in this space exhibit characteristics consistent with FRBD. The use of fMRI as an imaging modality also rests on an implicit assumption that some significant proportion of the data being recorded reflects functionally relevant brain dynamics. It would be hard, for example, to justify studying even scan-length (static) spectral and relational/connectivity characteristics of timeseries whose fluctuations are believed to be largely artifactual. The functioning human brain during any state of wakefulness is (hopefully inarguably) continually engaged in myriad temporally-varying combinations of cognitive, sensorimotor, attentional, emotional, planning, imagination and memory-related tasks. Many of these functions are in use even during sleep. Thus, any null model of multivariate timeseries whose characteristics are highly consistent with empirically observed fMRI-based brain measurements has little utility, since the phenomenon that it is testing for is ubiquitous rather than rare. The space of SCC multivariate Gaussians replicates real fMRI network timeseries with sufficient fidelity to induce broad consistency in measurable characteristics between the simulated data and the empirical data it was modeled upon ¹². Moreover, there is no *a priori* reason to believe that aberrant or “tail” phenomena in this space should be more strongly associated with functionally-relevant brain dynamics than with measurement noise, motion or other artifacts, e.g., the sort of features that might warrant examining a scan for possible removal rather than positioning it as an exemplar of functionally-relevant resting state brain dynamics.

Statistical Stationarity Does not Ensure Absence of Evident Temporal Epochs



$$E(F_1(t)) = E(F_1(t')) = 0.3354 \text{ for all } t, t' \quad E(F_1(t)F_1(t')) = c(\Delta) \text{ for all } \{t, t': |t - t'| = \Delta\}$$

Figure 1 The function f is a highly stylized example of a signal with distinguishable temporal epochs. The first, second, and final third have different means, variances and characteristic frequencies. The stochastic process F_1 , however, whose realizations are obtained through uniform random circular shifts of f , is statistically stationary in that that the statistical summaries assessed at distinct timepoints over large numbers of realizations are the same. All realizations have spectral and epochally clear variations, which would be reflective of FRDB, but as a stochastic process the collection of phase-shifted versions of f are statistically stationary.

Spectrally Matching Gaussian Noise to a Fixed Narrow-Band Target Spectrum Does not Ensure Absence of Evident Temporal Epochs

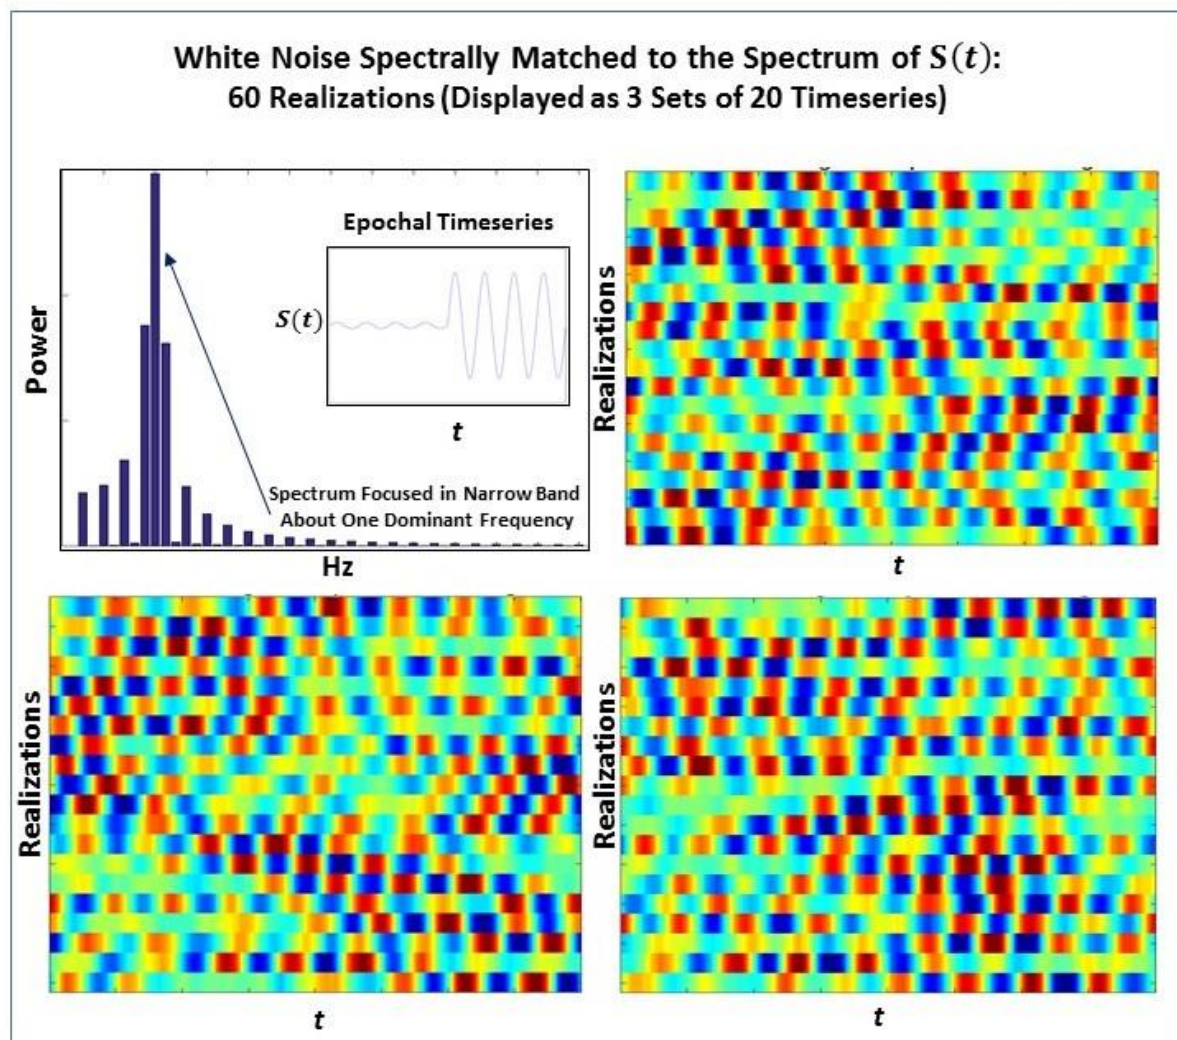


Figure 2 The signal $S(t)$ consisting of a low-amplitude 0.08Hz segment followed by a high-amplitude 0.08 Hz segment is one manifestation of a signal with a narrow-band spectrum (shown top left) focused at 0.08 Hz. Matching 60 timeseries of Gaussian white noise to this spectrum yields (in sets of 20) the timeseries shown in the other three panels of this figure. It is evident that spectral-matching of Gaussian noise to a generic template spectrum can contain discernible temporal epochs with any given realization. Thus, such a model cannot be a good null model for dynamic connectivity as it will contain the very dynamics we are interested in studying.

2.2. Multivariate Signals Lacking Plausible Markers of rs-Brain Dynamics (Null Models)

The development of valid null models for rs-FRBD is substantially hindered by a dearth of fMRI recordings from living subjects under conditions that all but preclude the ongoing CASE shifts unavoidably present even in sleeping or mentally impaired subjects. Resting state fMRI data is generally recorded under conditions in which functionally relevant brain dynamics ought to be continually present. Thus, the measurable features of empirically

observed multivariate fMRI network timeseries are intractably “contaminated” from the standpoint of parameterizing a null space in which signal properties reflecting rs-FRBD are ensured to be rare.

Another challenge for hypothesis testing of rs-FRBD resides in identifying quantifiable signal features for which every upward increment of the associated measure unambiguously yields stronger evidence for the presence of rs-FRBD. Without this property, observations from the distributional tails of the measure are simply improbable, but not necessarily in ways that relate to FRBD. Kurtosis, for example, rises with the number and magnitude of observations in a sample that would be outliers if the underlying process was stationary Gaussian, i.e., a Gaussian with constant mean and variance. Modestly elevated kurtosis might well reflect some unusually strong or active brain dynamics – this would have to be demonstrated, but is not implausible. However, extremely large kurtosis values occur when a sample contains numerous wildly large (in magnitude) observations. A brain recording with these characteristics is more likely to suggest noise, motion or a subject undergoing repeated electrical shocks than anything connected with actual rs-FRBD occurring during the scan (see **Figure 4**). More generally though, the effort to statistically validate the presence of a phenomenon that is almost axiomatically continual in any valid recording seems misguided. The more interesting hypotheses about fMRI dynamics would not focus on *whether* they exist, but rather on how they might manifest differentially over different timescales, spatial scales and functional scales.

3. Measuring Functionally Relevant Brain Dynamics vs. Identifying Outlying Observations

A valid metric of brain dynamics should rise monotonically with the strength of the signal features that are, at our current level of scientific knowledge, widely believed to have associations with CASE or task-driven shifts in brain function. The metric should also be as blind as possible to signal features believed to represent nuisance factors. The problem of quantifying signal features that have a high likelihood of representing evolving CASE states is admittedly very difficult. Every procedure will be biased by assumptions whose validity cannot be ensured based the current state of knowledge. The best we can hope for right now is for researchers to be clear about what signal properties they are seeking to quantify, and why they believe that a strong presence of these properties should be taken as evidence of FRBD (or of nuisance factors that are not easily separable from FRBD at current levels of measurement resolution).

We show in this section that a measure based on kurtosis, while sensitive to outliers, is not an ideal metric to capture brain dynamics and it is quite easy to show that kurtosis can be more sensitive to very rare outliers than it will be to more prevalent FRBD. We show in both stylized examples but also in real data that kurtosis preferentially captures signal features likely to arise from measurement disruptions (e.g., motion), while suppressing evidence of more extended spectral epochs within network timeseries. We also propose a new metric, Φ , which we believe shows some of the desired properties.

3.1. Univariate and Multivariate Kurtosis Under Stationary Gaussianity

Univariate excess or normalized kurtosis, the fourth statistical moment rescaled by squared variance and centered by subtraction of 3 has a well-understood distribution under the Gaussianity assumption that is applicable to ⁹ and here. Under this assumption, there is a closed-form transformation (dependent on the sample size, n) that converts observed excess kurtosis $b = \frac{\mathbb{E}((X-\mu_X)^4)}{\mathbb{E}((X-\mu_X)^2)^2} - 3$ into an unbiased estimator of true kurtosis,

ℓ_n (see Technical Supplement) that is distributed as a standard normal $\mathcal{N}(0,1)$ ¹³ and hence now has altered limits. The use of ℓ_n enables statistical evaluation of departures from stationary Gaussianity. When we refer to values of univariate kurtosis, these are values of the unbiased estimator ℓ_n . Samples that yield elevated values of ℓ_n (say, $\ell_n > 1.96$, so that $p < 0.05$) contain high-magnitude observations that are too numerous and/or too extreme for the sample to have even a 5% chance of having been generated by a stationary Gaussian process.

There is a similar transform for Mardia's multivariate kurtosis (m.v. kurtosis), with a similar interpretation. The unbiased estimator $B_{n,p}$ (dependent on sample size, n , and vector length, p , see Technical Supplement) for Mardia's multivariate kurtosis ¹⁴ is:

$$B_{n,p} = \frac{1}{n} \sum_{i=1}^n (x_i - \mu_i)^T C^{-1} (x_i - \mu_i)$$

where C^{-1} is the $n \times n$ inverse covariance matrix of the time-indexed p -vectors $\{x_1, x_2, \dots, x_n\}$.

3.2. Wavelet-Based Metric of Spectrally-Distinguishable Temporal Epochs

We briefly introduce a novel metric $\Phi: \mathbb{R}^T \rightarrow \mathbb{R}$ that explicitly captures the presence of spectrally distinguishable temporal epochs in a timeseries (see **Figure 3** and the Technical Supplement). The metric has a natural multivariate extension $\tilde{\Phi}: \mathbb{R}^{N \times T} \rightarrow \mathbb{R}$ that evaluates spectrally distinguishable temporal epochs in multivariate timeseries. Φ is not a primary focus of this short paper, but it plays a role in the discussion that follows because it provides a more targeted measurement than, for example, kurtosis, of timeseries characteristics that could from FRBD.

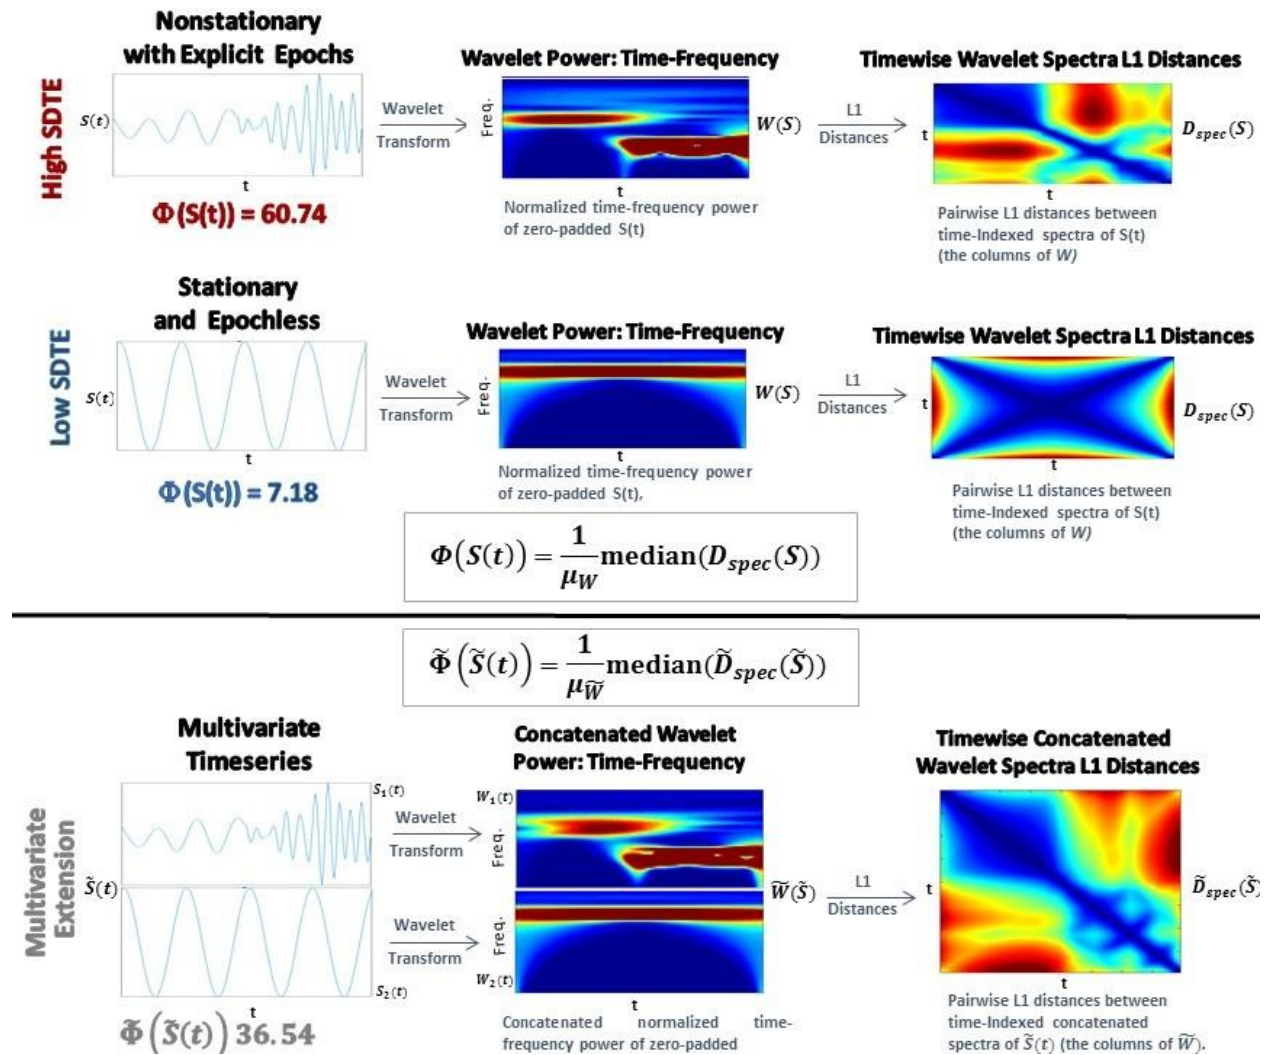


Figure 3 The metric Φ (details given in the Technical Supplement) is intended to capture spectral nonstationarities, or spectrally distinguishable temporal epochs (SDTEs), in univariate (top) and multivariate (bottom) timeseries. The univariate version uses wavelets to capture temporally-localized spectral information, yielding a set of time-indexed spectra (top middle), which we normalize to have total power equal to the product of the number of frequencies and the number of timepoints. We then compute pairwise L1 distances between the time-indexed spectra (top right) and compute the median off-diagonal value of the resulting $T \times T$ matrix. The multivariate extension $\tilde{\Phi}$ of Φ concatenates the time-frequency spectra of all univariate constituents along the frequency dimension (bottom middle), leaving the time dimension unaltered. In this case, we compute the median off-diagonal elements of the $T \times T$ matrix of pairwise differences between time-indexed concatenated spectra (bottom middle). The multivariate metric is higher when constituent univariate time-frequency spectra exhibit their largest within-signal spectral differences in mutually distinct temporal intervals. The case shown here does not illustrate the role of differential intervals of SDTEs among constituent timeseries.

3.3. Epochal Stationarity and Kurtosis

The presence of spectrally distinguishable temporal epochs across realizations, i.e., multiple subjects' connectivity characteristics in individual or multivariate network TCs is one reasonable potential form of evidence for rs-FRBD. Although it is also possible that this type of phenomenon could arise from nuisance factors, epochal behavior has structure that makes it less likely to be sourced dominantly in nuisance factors such as

motion, measurement noise or physiological rhythms. Kurtosis, which can help identify the presence of outliers in Gaussian data, has been proposed as a metric to detect FRBD, however kurtosis is highly susceptible to unstructured amplitude variations. Moreover, as an outlier metric, kurtosis has greater sensitivity to sharp, transient, high-amplitude anomalous intervals than to signals with amplitude and frequency variation on more functionally relevant timescales (see **Figure 4**). In fact, the properties leading to elevated kurtosis are sometimes more present in an epochally stationary signal than in an epochally nonstationary signal (see **Figure 5**), i.e., one that is stationary except within a given duration. This is not to say that measuring epochal nonstationarity is straightforward. There are many features and timescales on which the nonstationarity might be exhibited, and most prospective metrics will present some combination of over-sensitivity to irrelevant features and blindness to important features and/or timescales. We are currently working on a flexible, tunable approach to capturing the kind of epochally structured frequency domain variation that promises to provide valid evidence for brain dynamics after careful evaluation of sensitivity to nuisance factors.

Univariate Kurtosis and ϕ

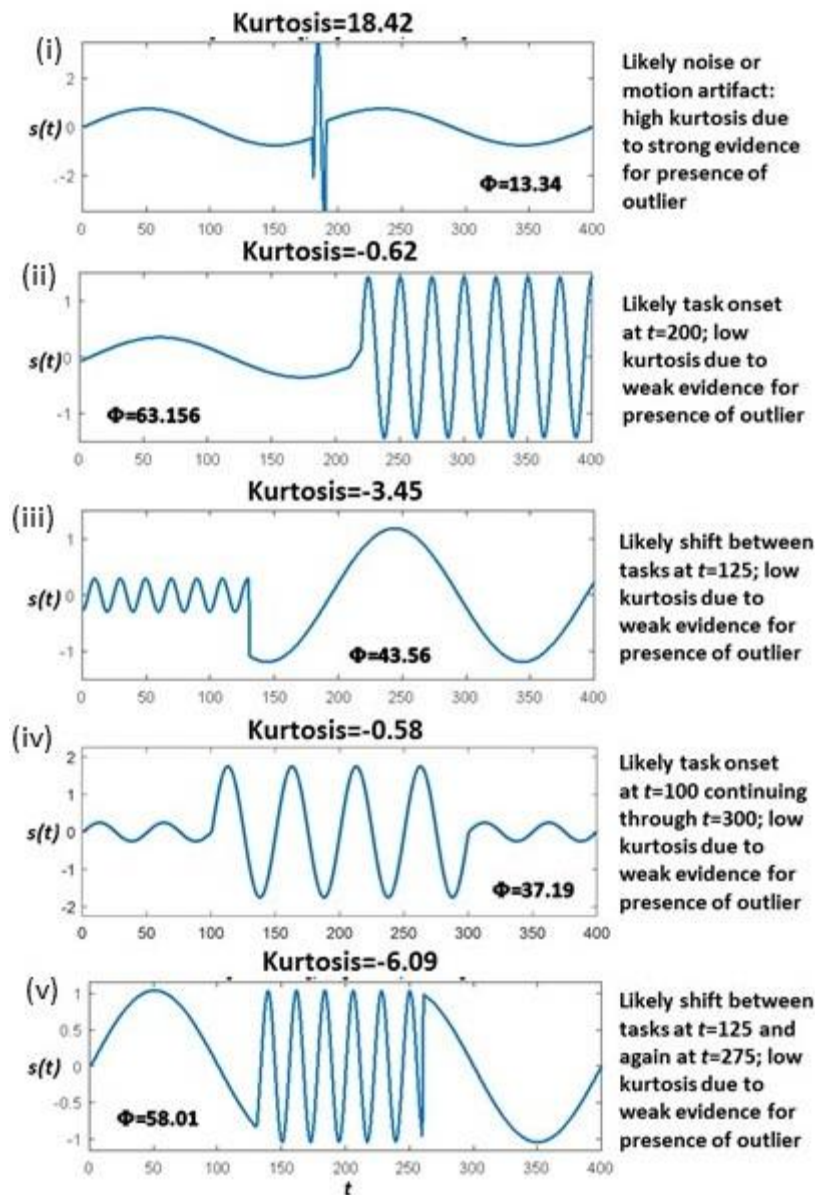




Figure 4 (i) Highest kurtosis applies in a signal with a transient high-amplitude high-frequency interval, more consistent with a motion or noise than with shifting CASE states; Φ is extremely low for this seeming artifact; (ii-iv) Much lower kurtosis in stylized signals with highly distinguishable temporal epochs are more consistent with FRBD; much higher Φ in these strong FRBD candidates; (v) Lowest kurtosis in stylized signal with very strong and distinguishable temporal epochs; much higher Φ for this strong FRBD candidate. Positive univariate excess kurtosis indicates super-Gaussianity (and is not particularly useful for indicating functionally relevant dynamics). As such, kurtosis is high when there are a larger number of high-amplitude observations than should arise under the assumption of Gaussianity. Univariate kurtosis (transformed via equation (1) in the Technical Supplement) to distribute, assuming Gaussianity, as a standard normal random variable) is negative on each of the stylized examples (ii)-(v) that exhibit distinguishable temporal epochs consistent with functionally relevant brain dynamics. It is very large and positive only in the example containing a single high-amplitude, high-frequency “spike” (i). The behavior that appears in the upper tail of the kurtosis distribution is more consistent with motion artifacts or measurement error than anything previous imaging or EEG studies have found to be associated with experimental tasks. The metric, Φ , introduced in this work, is at least three times larger for

examples (ii)-(v) that exhibit distinguishable temporal epochs consistent with functionally relevant brain dynamics than for the case (i) that features a single high-frequency high-amplitude “spike” embedded in an otherwise spectrally epochless signal. As such Φ exhibits the behavior we would expect, whereas kurtosis is not particularly useful for detecting behavior consistent with relevant brain dynamics for the examples shown above.

Multivariate Kurtosis and Φ

Two Timeseries Containing Spectrally Distinguishable Temporal Epochs with Multiple Connectivity Regimes



Two Spectrally Static Timeseries with One Relatively Stable Connectivity Regime





Figure 5 Stylized signal pairs; (Left) two signals, each with highly distinguishable spectral epochs and whose connectivity switches sign on one epoch vs. (Right) two perfectly stationary pure sinusoids whose connectivity does not change at all; m.v. kurtosis is higher for the stationary pairs with unchanging connectivity than for the pairs with shifting connectivity. **Note that Φ is higher for the pair the includes epochs.** Excess multivariate kurtosis indicates multivariate super-Gaussianity (and is not particularly useful for indicating functionally relevant dynamics). Multivariate kurtosis (transformed via equation (4)) from the Technical Supplement) is assumed to be distributed, assuming multivariate Gaussianity on the part of the random vector, as a multivariate standard normal random variable) on both stylized multivariate examples. The first (left) features two signals that each exhibit highly distinguishable spectral epochs and whose correlative behavior is also dynamic: they are perfectly correlated, then perfectly anti-correlated, then again perfectly correlated. This is a very dynamic context but not only presents negative multivariate kurtosis, its kurtosis value is even more negative than the second example (right) that features two spectrally unchanging signals whose mutual correlations are consistently zero. The multivariate measure Φ is twice as large in the dynamic example (left) compared to the static example (right).

3.4. Empirical Data and Simulation Regimes

A set of network timecourses from a clinical rs-fMRI study on which dynamic functional network connectivity (FNC) results have already been published¹⁵, and five simulation regimes modeled on that data are employed to explore and illustrate the role of

Gaussianity and statistical stationarity as well as, spectral and covariance stationarity in modeling rs-FRBD (and/or its absence). It is important to note that we refer to stationarity in the true sense as statistical (non)stationarity in what follows. We use the terms “spectral (non)stationarity” and “covariance (non)stationarity” to refer to other definitions used including ⁹ that invoke the concept of (non)stationarity through the analysis of a single realization either in the spectral domain (for the former definition) or using covariance function (for the latter definition). The approaches are lightly outlined here, with more details provided the Technical Supplement.

Real Data: We used previously published ¹⁵ network timecourse data from a large multisite clinical resting-state fMRI study. These timecourses (314 subjects, 47 networks, 158 timepoints), subsequently filtered for frequencies at most 0.08 Hz and z-scored, are referred to below as “Real Data”¹⁵ (see **Figure 6**; top left). The average power in each frequency bin in [0.003,0.08] Hz for all network TCs for all subjects is denoted \bar{P}_{data} . The average cross-network covariance matrix for all subjects is denoted \bar{C}_{data} . Every simulation regime described below consists of 1000 simulated subjects, each characterized by a set of 47, length-158 timeseries.

SCC Gaussians: statistically stationarity without constraint on SDTEs: Following ⁹, each simulated subject in the SCC Gaussian regime is a multivariate timeseries resulting from the projection of a 47×158 matrix of low-pass filtered white noise spectrally matched to \bar{P}_{data} onto the eigenspace of \bar{C}_{data} (see **Figure 6**; top middle).

Covariance-Dynamic SCC Gaussians: piecewise stationary, two distinct covariance regimes, no constraint on SDTEs (CD-SCC Gaussian): This regime introduces explicit covariance nonstationarity. Each CD-SCC Gaussian subject starts as a 47×158 matrix consisting of low-pass filtered white noise spectrally matched to \bar{P}_{data} , which is then divided into three windows, determined by a middle window of randomly chosen length between 40 and 60 TRs. The middle window is projected onto the eigenspace of \bar{C}_{data} , which has structure and is strongly connected, while the first and final windows are projected onto the eigenspace of C_{weak} , a covariance matrix reflecting very weak unstructured network connectivity (see **Figure 6**; bottom left).

SCC Gaussians with Noise: statistically stationary with a single spike randomly inserted into a small proportion of network timeseries (“Noisy SCC Gaussian”): This regime introduces extremely sparse, high-amplitude noise to the SCC Gaussian setting. Each Noisy SCC Gaussian subject starts as an SCC Gaussian subject, i.e., as a 47×158 matrix consisting of low-pass filtered white noise spectrally matched to \bar{P}_{data} then covariance matched to \bar{C}_{data} . Of the 47 timeseries in this matrix, between 3 and 15 are selected at random to carry a single high frequency spike centered at some randomly selected timepoint. The entire multivariate timeseries contains between 3 and 15 of these noise artifacts, with at most one in any given univariate timeseries (see **Figure 6**; top right).

Spectrally and Statistically Nonstationary: explicitly nonstationary both statistically and epochally (“SS Nonstationary”): This regime introduces simulated task-responsiveness to the SCC Gaussian setting. Each simulated SS Nonstationary subject starts as an SCC Gaussian subject (a 47×158 multivariate timeseries produced by subjecting white noise to spectral and covariance constraints exhibited by the real data). The 47 networks for each subject are divided into at most 17 task-positive networks (TPNs), at most 11 task-negative networks (TNNs), with the remaining 19 – 29 networks designated as nonresponders (NRs). Following the same procedure employed for the CD-SCC Gaussian regime, multivariate timeseries in the SS Nonstationary regime are divided into three windows determined by a middle window of randomly-chosen length between 40 and 60 TRs. The hypothetical task takes place during the middle window, in which (relative to the first and final window) the selected TPNs exhibit faster, higher amplitude behavior, the selected TNNs exhibit slower, lower amplitude behavior and the NRs exhibit no change in behavior (see **Figure 6**; bottom middle).

Covariance-Dynamic Spectrally and Statistically Nonstationary: explicitly nonstationary both statistically and epochally with two distinct covariance regimes (“CD-SS Nonstationary”): This regime introduces explicit covariance nonstationarity to the SS Gaussian setting. Each simulated CD-SS Nonstationary subject starts as an SS Nonstationary subject (see immediately above). In this regime, however, the temporally task-responsive middle window is additionally subjected to explicitly different covariance constraints than the task-free first and final windows. Following the procedure from the CD-SCC Gaussian, the middle window of CD-SS Nonstationary subjects is projected onto the eigenspace of $\bar{\mathbf{C}}_{data}$, while the first and final windows are projected onto the eigenspace of \mathbf{C}_{weak} (see **Figure 6**; bottom right).

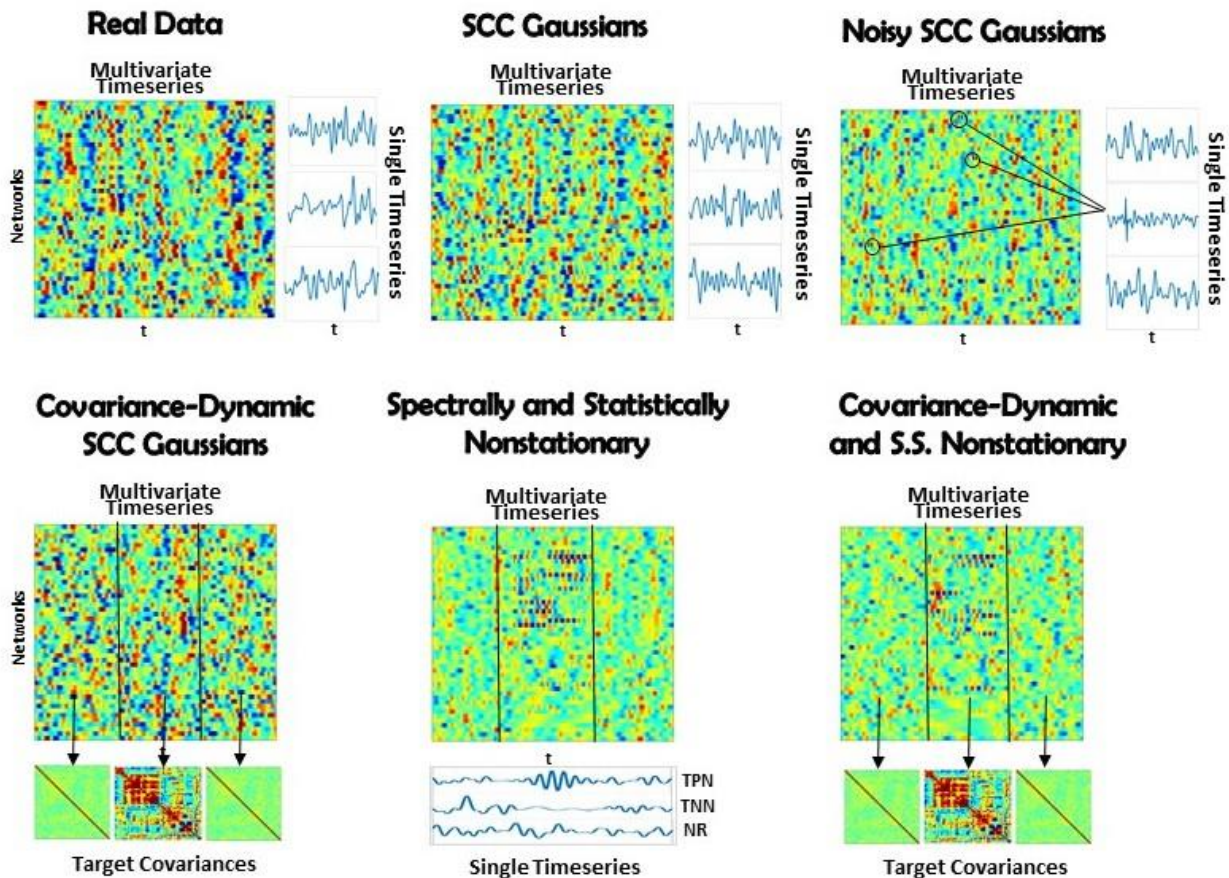


Figure 6 (Top Left) Real Data: Multivariate timeseries consisting of network timecourses from an actual rs-fMRI subject¹⁵ (3 single network examples displayed immediately to the right); (Top Middle) SCC Gaussians: Multivariate timeseries of SCC Gaussians modeled on mean spectral content and mean mutual covariance of empirical network timecourses from the study shown at the top left.; (Top Right) Noisy SCC Gaussians: The same as SCC Gaussians but with between 3 and 15 high-amplitude, high-frequency “spikes” interspersed, examples circled. No more than one spike is inserted in any individual timecourse and at most 15 of the 47 timecourses contain spikes. (Bottom Left) Covariance-Dynamic SCC Gaussians: This regime starts like the SCC Gaussians, in that the timeseries are spectrally matched to the mean spectral content of the target empirical dataset. However, the next stage involves covariance matching the middle interval to the mean mutual covariance of the empirical networks, while the first and final intervals are matched to a very weakly connected covariance structure. (Bottom Middle) Spectrally and Statistically Nonstationary: This regime also starts like the SCC Gaussians, in that the timeseries are spectrally and covariance matched to the mean spectral content and covariance structure of the target empirical dataset. Here though there is a middle window in which a subset of TPNs exhibits high-amplitude, high frequency behavior, a subset of TNNs exhibits low-amplitude, low-frequency behavior and most networks are NR to the stimulus. (Bottom Right) Covariance Dynamic and SS Nonstationary: This regime starts like the Covariance Dynamic SCC Gaussians, then a subset of TPNs and TNNs are chosen to respond in the same way as in the Spectrally and Statistically Nonstationary regime.

3.5. Kurtosis is Overly Sensitive to Noise Artifacts and Very Rarely Identifies Explicit Spectral, Statistical and Covariance Nonstationarities

As mentioned earlier, the utility of kurtosis as a measure of functionally relevant brain dynamics is mitigated by its highly-tuned sensitivity to spikes and outliers in the data. We saw this earlier in stylized univariate timeseries (see **Figure 4**). The issue is equally evident (see **Figure 7**) in more complex multivariate simulations involving different degrees of nonstationarity, exhibited in different ways. Employing m.v. kurtosis as a

measure of FRBD implies that timeseries on which it assumes values large enough to provide significant evidence ($p < 0.05$) against the null hypothesis of stationary multivariate Gaussianity are those in which the features associated with FRBD are most markedly present. More specifically, these features increase monotonically with m.v. kurtosis. In **Figure 7** however, we see that that the only the simulation regime in which a non-negligible percentage (93%) of realizations is identified as exhibiting FRBD using m.v. kurtosis is the regime featuring a handful of spikes in an otherwise stationary 47×158 multivariate Gaussian. Only 5.3% of realizations from the explicitly covariance-dynamic regime are identified using the m.v. kurtosis metric as exhibiting FRBD, and for the other nonstationary regimes the percentage of realizations identified as exhibiting FRBD is less than a tenth of a percent. Similarly, the actual data recorded from subjects undergoing continual CASE-shifts exhibits no evidence by the m.v. kurtosis criterion of having arisen from a source in which functionally relevant brain dynamics are present. This clearly highlights the limitations of kurtosis as an indicator of FRBD. Elevated kurtosis indicates the sample contains too many points that are too extreme in magnitude under an assumption of stationary Gaussianity; it is less effective at capturing the dynamic changes in temporal behavior (including covariance) that are more plausible markers of FRBD and are richly present in most realizations of even stationary multivariate Gaussian processes.

Percentage of Multivariate Observations Found to be Nonstationary or Dynamic by the Multivariate Kurtosis Metric

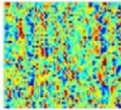
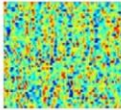
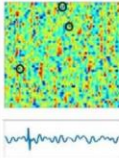
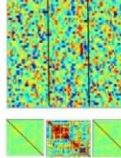
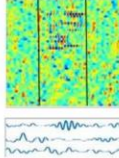
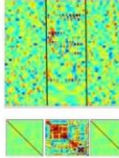
Multivariate Kurtosis (Percentage of Observations Whose Z-Transformed Multivariate Kurtosis is Elevated Enough to Offer Significant ($\alpha = 0.05$) Evidence Against the Observation Being a Stationary Multivariate Gaussian)	Real Data (N=314)	Statistically Stationary SCC Gaussians (N=1000)	Statistically Stationary SCC Gaussians (with a few spikes) (N=1000)	Covariance-Dynamic SCC Gaussians (N=1000)	Spectrally and Statistically Nonstationary (N=1000)	Covariance-Dynamic and SS Nonstationary (N=1000)
						
	0.0%	0.0%	93.4%	5.3%	0.0%	0.1%

Figure 7 The percentage of multivariate timeseries from each indicated simulation regime (and the real rs-fMRI study on which the simulated data was modeled) that present significant evidence ($p < 0.05$) of having been generated by some process that is not a stationary multivariate Gaussian. The explicitly implemented spectral and statistical nonstationarities of the two SS Nonstationary regimes (columns 5 and 6) are effectively never found to exhibit significant evidence against being generated by stationary multivariate Gaussians. In the two explicitly covariance-nonstationary regimes (columns 4 and 6), one of which is also spectrally and statistically nonstationary variations (column 6), at most 5% of the 1000 simulated subjects – each of which exhibits the explicit nonstationarity – are identified as unlikely to

have arisen from a stationary multivariate Gaussian process. Like the SCC Gaussian simulations (column 2), the SS Nonstationary simulations (columns 5 and 6) and the covariance-nonstationary simulations (columns 4 and 6), empirical observations from real subjects (column 1) in whom the phenomenon of interest (FRBD) is ubiquitous are not statistically distinguishable from realizations of a stationary multivariate Gaussian process. The only regime that multivariate kurtosis reliably distinguishes from realizations of a stationary multivariate Gaussian process is the case of SCC Gaussians in which a single high-amplitude, high-frequency spike is inserted into between 3 and 15 of the 47 univariate timeseries from the multivariate observation (column 3). This regime is basically just a lightly contaminated version of the SCC Gaussian regime (column 1) and of all of the simulation regimes exhibits the least evidence of functionally relevant brain dynamics. The behavior underlying upper-tail observations of multivariate kurtosis looks more like scan contamination than anything task-paradigm fMRI studies suggest would be strongly associated with FRBD.

Conversely, as shown in **Figure 8**, the proposed metric (defined in **Figure 3** and the Technical Supplement) responds in a more reasonable way to those features of real and simulated multivariate timeseries that have strong likelihood of reflecting FRBD vs. those that are simply aberrant in some other way. Results show that, in contrast to multivariate kurtosis, the lightly contaminated SCC Gaussians in terms of $\tilde{\Phi}$ exhibit significantly less evidence of reflecting FRBD than either of the SS Nonstationary regimes (**Figure 8**, row 2, columns 5 and 6) and real data is statistically indistinguishable from the covariance-dynamic SCC Gaussian regime (**Figure 8**, row 1, column 4). So $\tilde{\Phi}$ is putting regimes exhibiting different kinds of temporally epochal behavior—including the epochal behavior seen in statistically stationary Gaussian processes—in what seems a plausible ordering with respect to dynamism: SCC Gaussians \preceq Real Data \approx Covariance-Dynamic SCC Gaussians \preceq Noise-Contaminated SCC Gaussians \preceq SS Nonstationary \preceq Covariance-Dynamic SS Nonstationary (where curly binary relations indicate ordinal evidence of potentially relevant multivariate epochal behavior as measured by $\tilde{\Phi}$ (see legend for more details).

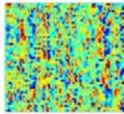
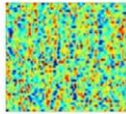
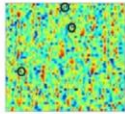
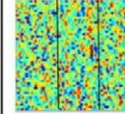
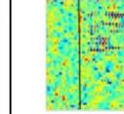
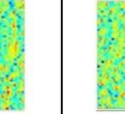
<i>T-Statistic for Multivariate Metric $\tilde{\Phi}$ of SDTEs</i> <i>(T-Test of $\tilde{\Phi}$ on Row Regime vs. Column Regime; shown where significant ($p < 0.05$))</i>	Real Data (N=314) 	Statistically Stationary SCC Gaussians (N=1000) 	Statistically Stationary SCC Gaussians (with a few spikes) (N=1000) 	Covariance-Dynamic SCC Gaussians (N=1000) 	Spectrally and Statistically Nonstationary (N=1000) 	Covariance-Dynamic and SS Nonstationary (N=1000) 
Real Data	*	11.51	-10.63	*	-81.63	-81.74
SCC Gaussians	-11.51	*	-36.45	-25.76	-184.95	-184.67
SCC Gaussians (with a few spikes)	10.63	36.45	*	19.02	-106.75	-106.82
Covariance-Dynamic SCC Gaussians	*	25.76	-19.02	*	-168.06	-167.82
Spectrally and Statistically Nonstationary	81.63	184.95	106.75	168.06	*	*
Covariance-Dynamic and SS Nonstationary	81.74	184.67	106.82	167.82	*	*

Figure 8 This table presents the *t*-statistic (where significant at the $p < 0.05$ level) for pairwise *T*-tests of $\tilde{\Phi}$ on the row regime vs. the column regime. From this standpoint, we see that the real network timecourses (row 1) exhibit significantly greater presence of SDTEs than the SCC Gaussians (column 2) modeled on them, significantly less evidence of SDTEs than the lightly contaminated “noisy” SCC Gaussians (column 3) and the explicitly SS Nonstationary Regimes (columns 5 and 6) and are statistically indistinguishable from the explicitly covariance-dynamic SCC Gaussians (column 4). Unlike what was found using multivariate kurtosis, the covariance-static SCC Gaussians (row 2) are in terms of $\tilde{\Phi}$ significantly less dynamic than the covariance-dynamic SCC Gaussians (column 4) and both SS Nonstationary regimes (columns 5 and 6). Again, in contrast to multivariate kurtosis, the lightly contaminated SCC Gaussians (row 3) in terms of $\tilde{\Phi}$ are significantly less dynamic than both SS Nonstationary regimes (columns 5 and 6). $\tilde{\Phi}$ is putting regimes exhibiting different kinds of temporally epochal behavior – including the epochal behavior seen in statistically stationary Gaussian processes – in what seems a plausible ordering with respect to dynamism: SCC Gaussians \leq Real Data \approx Covariance-Dynamic SCC Gaussians \leq Noise-Contaminated SCC Gaussians \leq SS Nonstationary \leq Covariance-Dynamic SS Nonstationary (where curly binary relations indicate ordinal evidence of potentially relevant multivariate epochal behavior as measured by $\tilde{\Phi}$).

4. Ensuring that the Null and Alternative Hypotheses Yield Different Distributional Tails

One key challenge in creating null models is to ensure that the null distribution is truly a null, *i.e.*, the behavior of interest is not generically present within the null itself. The field of dynamic connectivity in particular has struggled with this as there are multiple threads of research proposing different models of dynamic behavior including covariance dynamics^{1,15} and oscillatory dynamics¹⁶. In this context, it is very important to keep any

proposed null model within the narrow context within which it is able to reject a particular hypothesis about brain dynamics. For example, a null model of oscillatory dynamics may well contain within it dynamically changing covariances, and, as such, it is not a particularly good null model for that particular scenario¹⁰. In addition to that, even models of covariance dynamics have been shown to be limited to a rather particular set of parameters, which we do not yet fully understand how to set for the human brain⁸. There is also evidence that the brain functions as a nonlinear dynamical system¹⁷. The result is that many of the null models that have been proposed are making strong assumptions about brain dynamics which, while having some justification, are not able to rule out dynamics of a different sort. We provide some additional discussion on this important point below.

4.1. Whole-Brain Windowed Connectivity States and Occupancies from Statistically Stationary and Explicitly Nonstationary Null Models Strongly Resemble Each Other

As we have discussed previously, creating a valid null model is difficult without explicitly understanding the properties that multivariate network timeseries might exhibit in response to complex CASE variations. One approach, suggested recently by ⁹, employs a space of low-pass filtered multivariate white noise, spectrally matched to the average spectrum of empirical timecourse data and then projected onto the eigenspace of empirically observed mean network covariance. As discussed in Section 3, this is a space of timeseries explicitly modeled on real data recorded from a material in which the phenomenon of interest (i.e., functionally relevant brain dynamics) is continually present. It is therefore not a space in which univariate or multivariate temporal behavior plausibly sourced in CASE variations is vanishingly rare. Which is to say, it is not a useful null space for the identification of rs-FRBD. Moreover, it is a space of signals whose time-varying behavior is jointly determined by all simulation parameters and assumptions: the auxiliary spectral and covariance constraints as well as the Gaussianity assumption and the primary assumption of statistical stationarity. Above and beyond the problematic success of this simulation model in replicating empirical data recorded under circumstances in which the phenomenon of interest is uninterrupted and continual, this simulation regime's value as a null model is further undermined by the unexamined role of auxiliary parameters in shaping key measures and distributions. This can be seen in **Figure 9**, rows 2 and 3, where we break the model's core assumption of statistical stationarity, a property, arguably incorrectly, associated by the authors with an absence of FRBD, without discernibly disrupting either the clusters formed by short-timescale FNCs or the average cluster occupancy rates (see the Technical Supplement for a brief background on this approach). The distributional tails of occupancy rates for each connectivity state in the SCC Gaussian and the SS Nonstationary regimes have significant overlap (see **Figure 10**) and as such this model can quite easily rule both for and against dynamic connectivity at the same time, an obviously flaw in the approach. Setting aside for a moment whether stationary Gaussianity is a criterion that characterizes or excludes signal features reflecting rs-FRBD, the model as specified is not useful for finding evidence

against this assumption, that is not simultaneously providing evidence against its complement, i.e., the distributional tails of occupancy rate as a test statistic are not hypothesis-separating. This example illustrates the difficulties of building hypothesis-testing frameworks for phenomena whose distinguishing quantifiable characteristics are not well understood. If, in contrast to ⁹, one realizes that the space within which one is working contains the very dynamics that one is trying to rule out (a point subsequently made by ¹²) the conclusions that can be made are unconvincing and uninteresting.

Simulation Regimes and Windowed Connectivity

Windowed Connectivity States and Occupancy Rates from Real Data and Under Stationary and Nonstationary Simulation Regimes

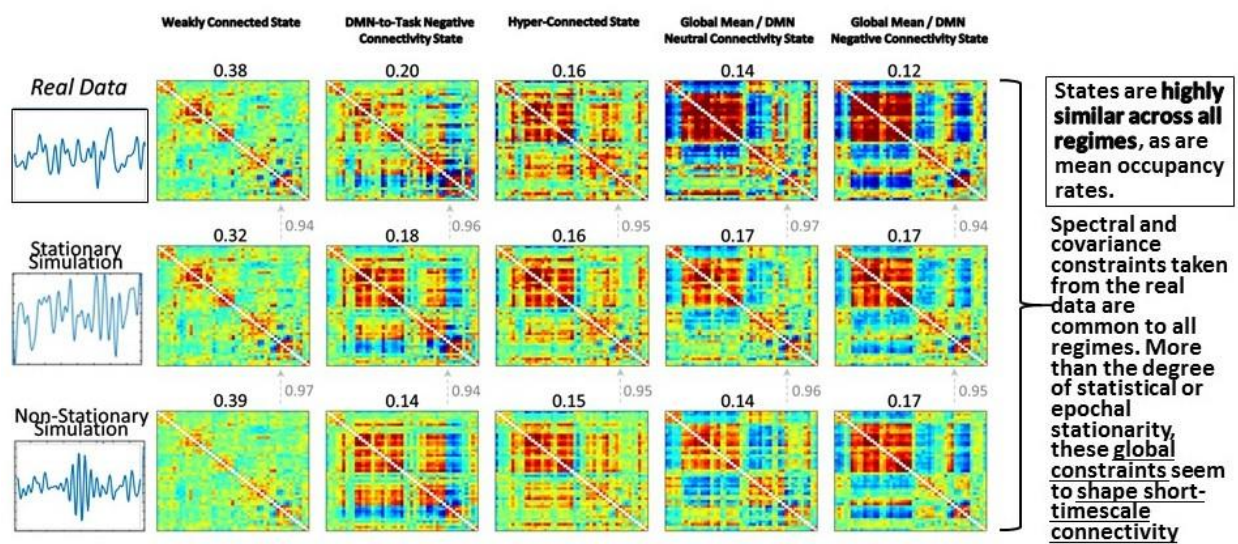


Figure 9 As we have pointed out previously, statistical stationarity, even multivariate stationary Gaussianity, can be richly inclusive of the types of time-varying multivariate behavior consistent with known brain responses to experimental tasks. Thus, the windowed connectivity states that a statistically stationary multivariate process moves through have every reason to resemble those of the real data upon which the process was tightly modeled (rows 1 and 2). However, it is also the case that short-timescale connectivity measurements from explicitly nonstationary processes (see “Statistically and Spectrally Nonstationary” regime in Section 3.4) subjected to spectral and global covariance constraints drawn from the real data cluster in the same way (row 3) as both the real data and the stationary Gaussian simulation modeled upon it. This suggests that the short-timescale connectivity states and occupancies are driven more by auxiliary constraints on mean spectrum and mean covariance than by whether the underlying process is statistically stationary. As we have shown, statistical stationarity does not preclude a multivariate signal from passing through connectivity states resembling those potentially arising from FRBD in real data. But more importantly it passes through the same connectivity states in the same way as explicitly nonstationary processes subjected to the same auxiliary constraints and as such is not particularly useful as a null condition for detection relevant brain dynamics.

Simulation Regimes and Lack of Hypothesis-Separating Distributional Tails

Overlapping Tails in the Occupancy Distributions of Windowed Connectivity States Under Stationary and Nonstationary Simulation Regimes

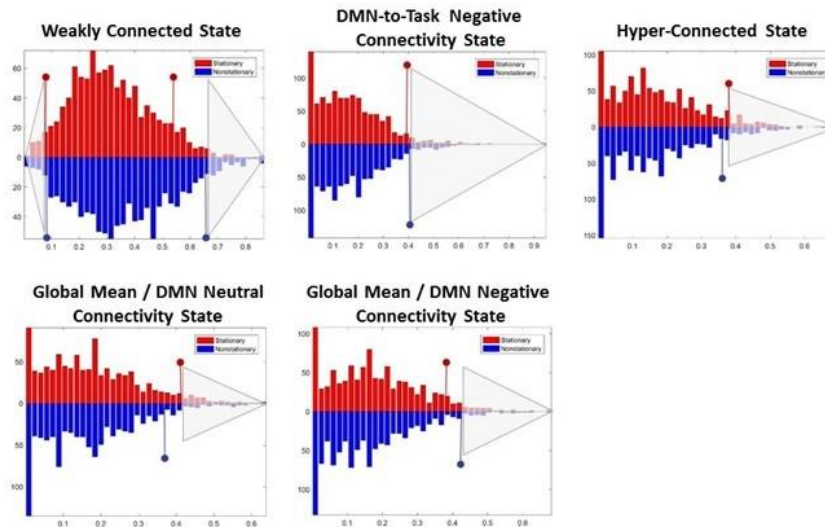


Figure 10 For a null model to be useful, it should exhibit evidence of the phenomenon being investigated extremely rarely. SCC Gaussians arguably present evidence of FRBD as ubiquitously as the real fMRI data does. It is also not clear a priori that tail phenomena in SCC Gaussians should have an especially strong association to FRBD rather than, for example, to motion artifacts or measurement noise. Moreover, the measures evaluated on a null model are only helpful in identifying behavior of interest if the distributional tails of the measure do not overlap with those of the same measure evaluated on a model that violates assumptions of the original null hypothesis. The upper tails of the distributions of occupancy rates on the windowed connectivity states from SCC Gaussians (red) and the SS Nonstationary simulation regime (blue) have substantial overlap (semi-transparent grey triangles indicate the overlapping part of the tails for each state). Insufficient clarity on the unique, distinguishing features of the phenomenon being studied and of the models/measures being employed neutralizes the measure's role in hypothesis testing as the same observation can then present significant evidence against multiple, mutually contradictory null hypotheses. In terms of occupancy rates of windowed connectivity states, we see here that a newly scanned subject who spends 65% of their time in state 2 exhibits significant evidence against one statistically stationary null hypothesis and also against an explicitly nonstationary null hypothesis.

5. Discussion

In this position paper, we have attempted to clarify and develop some of the important issues related to dynamic connectivity within the resting brain. Our focus has mostly centered around the use of metrics to detect possible dynamic behavior and also the creation of appropriate null models of functionally relevant brain dynamics, including but not limited to dynamic connectivity. We have discussed crucial limitations hindering some existing proposed models, proposed a possible metric for detection of functionally relevant brain dynamics, and have, we hope, provided the context for a rich ongoing discussion of where the field should head from here. In the remainder of this paper, we highlight a few high-level questions that we hope make this point clear.

5.1. Are Functional Brain Dynamics Rare?

One of the key points we emphasized relates to the creation of a null simulation model. A useful null simulation model should be built on assumptions complementary to those that would apply if the phenomenon of interest were present, and then combined with a test statistic that is sufficiently sensitive to the null assumptions that it very rarely achieves extreme values when they are absent. One example of this latter concern is ¹⁰, in which a particular and carefully constructed null model was proposed. The test statistic employed however was variance, whose distributional tails on their static null simulation model matched the tails of variance computed on a wide range of explicitly dynamic variants of the model in which parameters were meaningfully changing with time. The study of functional brain dynamics presents challenges to this entire type of modeling. We know that functional brain dynamics are constantly present in conscious human beings, so the phenomenon is anything but rare empirically.

5.2. Are Functional Brain Dynamics Separable from Nuisance Factors and Background Brain Rhythms?

The null model for a phenomenon necessarily produces, with very low probability, values of some test statistic that are consistent with the phenomenon which is being tested. This kind of model makes sense primarily when there is a very clear understanding of the range of values a particular test statistic will assume in the presence of the hypothesized phenomenon. So although in this short paper we are focusing primarily on the insufficiency of null models built on many realizations of some statistically stationary multivariate process modeled on empirical timecourse spectra and covariance, the larger problem is really that our present understanding cannot rule out the possibility of nontrivial intersection between signal features foreseeably connected with brain dynamics and those arising from nuisance factors and background brain rhythms. Even a spectrally pure signal, e.g., a single-frequency sinusoid, features amplitude changes that could be consistent with the ebb and flow of a network's contributions to temporally varying CASE demands. Certainly pure sinusoids and other epochally stationary signals (whose phase-randomized stochastic analogues are statistically stationary) could provide evidence of network responsiveness in certain experimental task paradigms, e.g., those involving repetitive motor or sensory tasks. Since most temporal changes in a signal, including raw amplitude changes, could plausibly be correlated with some complex sequence of CASE conditions, the present state of knowledge makes it difficult to construct null models that can claim to yield, almost-exclusively, timeseries (multivariate or univariate) lacking features prospectively associated with brain dynamics.

This is true even when the null model is narrow. Indeed, it can be difficult to ensure that the test statistic being assessed does not have distributional tails roughly matched to the distributional tails of that statistic on a similarly narrow but explicitly "dynamic" model. In that case, we can easily conceive of examples where we are in the position of having an empirical measurement of the test statistic that simultaneously leads to rejection of one static null hypothesis and various related dynamic null hypotheses. In such a case, the desired test has been rendered essentially useless. Absence of functionally relevant brain

dynamics in fMRI is the wrong null hypothesis for future research: the phenomenon is extremely high-dimensional and ill-defined and the question in this form is without scientific interest. As we have argued in this paper, the use of an “absence of brain dynamics” null model is not a particularly useful one for studying living humans as we expect functionally relevant temporal variations in brain activation to be constantly occurring throughout the experiment. As the phenomenon of FRBD is better understood, improved null models will naturally emerge and we hope that the discussions herein are useful and promote thoughtful consideration of these important issues.

6. Conclusions

To summarize, serious and continuing investigation of dynamic multivariate brain activation patterns (including dynamic connectivity) is scientifically important and central to many core open questions in brain science. The time-varying measurements provided by BOLD fMRI currently play a vital supporting role in this overall project. We have discussed some of the limitations of existing null models and metrics for capturing dynamics, and proposed a new wavelet-based metric which appears to provide sensible results in a number of simulated scenarios. Finally, we urge caution in the development of null models in the context of dynamic connectivity. Especially for studies in which subjects are not engaging in a common, narrow experimental task, the relevant features, temporal and spatial/functional scales are not yet well understood and moreover the question of *whether* functionally relevant brain dynamics “exist” in a living human subject seems both misguided and impossible to formulate in scientifically well-posed terms. Specific well-defined questions about how particular signal features evolve on a range of spatial and temporal scales could produce more useful and testable hypotheses about how the brain signals we measure relate to high-level processes by which the brain organizes, directs and rotates through some of its central tasks: e.g., cognition, sense-making, generative thinking, memory-formation, memory-retrieval and emotion regulation, among others.

7. References

- [1] R. M. Hutchison, T. Womelsdorf, E. A. Allen, P. Bandettini, V. D. Calhoun, M. Corbetta, S. D. Penna, J. Duyn, G. Glover, J. Gonzalez-Castillo, D. A. Handwerker, S. D. Keilholz, V. Kiviniemi, D. A. Leopold, F. de Pasquale, O. Sporns, M. Walter, and C. Chang, "Dynamic functional connectivity: promises, issues, and interpretations," *NeuroImage*, vol. 80, pp. 360-378, 2013, PMC3807588.
- [2] E. Tagliazucchi and H. Laufs, "Multimodal imaging of dynamic functional connectivity," *Front Neurol*, vol. 6, p. 10, 2015, 4329798.
- [3] L. Wu, T. Eichele, and V. D. Calhoun, "Reactivity of hemodynamic responses and functional connectivity to different states of alpha synchrony: a concurrent EEG-fMRI study," *NeuroImage*, vol. 52, pp. 1252-1260, 2010, PMC pending #248905.
- [4] E. Allen, E. Damaraju, L. Wu, T. Eichele, and V. D. Calhoun, "EEG signatures of dynamic functional network connectivity states," *Brain Topogr*, in press, PMC Journal - in progress.
- [5] R. M. Hutchison, T. Womelsdorf, J. S. Gati, S. Everling, and R. S. Menon, "Resting-state networks show dynamic functional connectivity in awake humans and anesthetized macaques," *Hum Brain Mapp*, vol. 34, pp. 2154-2177, Mar 22 2012.

- [6] M. G. Preti, T. A. Bolton, and D. Van De Ville, "The dynamic functional connectome: State-of-the-art and perspectives," *Neuroimage*, Dec 26 2016.
- [7] V. D. Calhoun, R. Miller, G. Pearlson, and T. Adali, "The chronnectome: time-varying connectivity networks as the next frontier in fMRI data discovery," *Neuron*, vol. 84, pp. 262-274, Oct 22 2014, PMC4372723.
- [8] S. Shakil, C. H. Lee, and S. D. Keilholz, "Evaluation of sliding window correlation performance for characterizing dynamic functional connectivity and brain states," *Neuroimage*, vol. 133, pp. 111-128, Jun 2016.
- [9] T. O. Laumann, A. Z. Snyder, A. Mitra, E. M. Gordon, C. Gratton, B. Adeyemo, A. W. Gilmore, S. M. Nelson, J. J. Berg, D. J. Greene, J. E. McCarthy, E. Tagliazucchi, H. Laufs, B. L. Schlaggar, N. U. Dosenbach, and S. E. Petersen, "On the Stability of BOLD fMRI Correlations," *Cereb Cortex*, Sep 02 2016.
- [10] R. Hindriks, M. H. Adhikari, Y. Murayama, M. Ganzetti, D. Mantini, N. K. Logothetis, and G. Deco, "Can sliding-window correlations reveal dynamic functional connectivity in resting-state fMRI?," *Neuroimage*, vol. 127, pp. 242-256, Feb 15 2016, PMC4758830.
- [11] A. Zalesky and M. Breakspear, "Towards a statistical test for functional connectivity dynamics," *Neuroimage*, vol. 114, pp. 466-470, Jul 1 2015.
- [12] R. Liegeois, T. O. Laumann, A. Z. Snyder, H. J. Zhou, and B. T. T. Yeo, "Interpreting Temporal Fluctuations In Resting-State Functional Connectivity MRI," *bioRxiv*, 2017.
- [13] M. K. Cain, Z. Zhang, and K. H. Yuan, "Univariate and multivariate skewness and kurtosis for measuring nonnormality: Prevalence, influence and estimation," *Behav Res Methods*, Oct 17 2016.
- [14] K. V. Mardia, "Measures of Multivariate Skewness and Kurtosis with Applications," *Biometrika*, vol. 57, pp. 519-8, 1970.
- [15] E. Allen, E. Damaraju, S. M. Plis, E. Erhardt, T. Eichele, and V. D. Calhoun, "Tracking whole-brain connectivity dynamics in the resting state," *Cereb Cortex*, vol. 24, pp. 663-676, 2014, PMC3920766.
- [16] M. Golos, V. Jirsa, and E. Dauce, "Multistability in Large Scale Models of Brain Activity," *Plos Computational Biology*, vol. 11, Dec 2015.
- [17] B. J. He, "Spontaneous and task-evoked brain activity negatively interact," *J Neurosci*, vol. 33, pp. 4672-4682, Mar 13 2013, PMC3637953.

Technical Supplement

SCC Multivariate Gaussian Processes

Let $w_{m,T}^{(i)}$ be one $m \times T$ matrix of white Gaussian noise, $\bar{\mathbf{P}}$ the average power spectrum for all subjects and all m network TCs in an actual resting state fMRI study, and \mathbf{C}_m the $m \times m$ population mean cross-network covariance. The set of timeseries one gets by inverse Fourier transforming a set of T normalized random complex coefficients $\theta_i = \frac{\omega_i}{\|\omega_i\|}$, $i = 1, 2, \dots, T$ weighted according to a fixed template spectrum $\bar{\mathbf{P}} = \{p_1, p_2, \dots, p_T\}$, ie, the set:

$$S(\bar{\mathbf{P}}, T) = \{s(t): s(t) = iFFT(\{p_1\theta_1, p_2\theta_2, \dots, p_T\theta_T\}, \text{random } \bar{\boldsymbol{\theta}}, \text{fixed } \bar{\mathbf{P}})\}$$

includes the discrete analogue (as we are working with computers on finite samples) of all length- T piecewise continuous functions with spectrum $\bar{\mathbf{P}}$. In particular, these timeseries exhaustively exhibit the full range of epochal behavior possible under the global spectral constraint $\bar{\mathbf{P}}$. Since $\bar{\mathbf{P}}$ is the average spectrum of network TCs from a real rs-fMRI study, this ensures that the epochal signatures of CASE-driven rs-FRBD are replicated in $S(\bar{\mathbf{P}}, T)$. The set of $m \times T$ multivariate timeseries

$$S_m(\bar{\mathbf{P}}, T) = \{s(t) = [s_1(t); s_2(t); \dots; s_m(t)]: s_i(t) \in S(\bar{\mathbf{P}}, T), i = 1, 2, \dots, m\}$$

similarly exhaust the range of m -fold jointly realizable within-timeseries epochs subject to the duration T and spectral constraint $\bar{\mathbf{P}}$. Finally, the set

$$S(\bar{\mathbf{P}}, T, \mathbf{C}_m) = \{s(t, \mathbf{C}_m) \equiv [s(t)]^T (\mathbf{E}(\mathbf{C}_m) \sqrt{\mathbf{V}(\mathbf{C}_m)}) : s(t) \in S_m(\bar{\mathbf{P}}, T)\}$$

of projections

$$s(t, \mathbf{C}_m) \equiv [s(t)]^T (\mathbf{E}(\mathbf{C}_m) \sqrt{\mathbf{V}(\mathbf{C}_m)})$$

of $s(t) \in S_m(\bar{\mathbf{P}}, T)$ onto the eigenspace of \mathbf{C}_m is also a collection of Gaussian multivariate process with average spectrum $\bar{\mathbf{P}}$ (where $\mathbf{E}(\mathbf{C}_m)$, $\mathbf{V}(\mathbf{C}_m)$ are respectively the eigenvectors and diagonal matrix of eigenvalues of \mathbf{C}_m). The multivariate Gaussian timeseries in $S(\bar{\mathbf{P}}, T, \mathbf{C}_m)$ exhibit the joint-spectral and covariation epochs that can arise in $m \times T$ multivariate timeseries with average spectrum $\bar{\mathbf{P}}$ and whose full-duration covariance is brought into close approximation to \mathbf{C}_m by projection onto its eigenspace. No explicit constraints are imposed on shorter-timescale epochs of covariation within elements of $S(\bar{\mathbf{P}}, T, \mathbf{C}_m)$; the set of possibilities is shaped primarily by ancillary constraints such as $\bar{\mathbf{P}}, T, \mathbf{C}_m$ and the linear projection by which $s(t)$ is aligned toward \mathbf{C}_m .

Wavelet-Based Metric of Spectrally-Distinguishable Temporal Epochs

The metric $\Phi: \mathbb{R}^T \rightarrow \mathbb{R}$ (still under active development) of within-timeseries spectrally distinguishable temporal epochs is computed on a univariate timeseries $S(t)$, $t \in \{t_0, t_1, t_2, \dots, T\}$ as follows (see Figure 5 in the main text):

1. Apply Matlab's continuous wavelet transform (with the Morse wavelet) to $S(t)$. This gives an $F \times T$ matrix $W(S)$ of wavelet coefficient magnitudes. The rows of $W(S)$ are timeseries of the power in

each of the F frequencies. First rescale $W(S)$ by the inverse of its mean, making its elements sum to FT .

2. The $T \times T$ symmetric matrix $D_{Spec}(S)$ of pairwise L1 distances between the columns of $W(S)$ contains evidence of temporal variation in the core spectrum of S . $D_{Spec}(S)$ contains evidence of *epochal* spectral variation in S .
3. Finally, we set $\Phi(S) \equiv \frac{1}{FT} \text{median}(\text{uppertri}(D_{Spec}(S)))$ to be the median off-diagonal values of $D_{Spec}(S)$ rescaled by the average time-frequency power in $W(S)$. Rescaling by the inverse summed power in $W(S)$ keeps the value of Φ strictly bounded in $[0, T]$.

The multivariate extension $\tilde{\Phi}$ of Φ assumes that $W(S)$ of size $F \times T$ has already been computed for some collection \mathbf{S} of N length- T timeseries S_i , $i = 1, 2, \dots, N$. Let $\tilde{W}(\mathbf{S})$ be the $NF \times T$ matrix of vertically concatenated $W(S_i)$'s. Now $\tilde{D}_{Spec}(\mathbf{S})$ is the matrix of pairwise L1 distance between columns of $\tilde{W}(\mathbf{S})$ the concatenated spectra of the S_i 's, and $\tilde{\Phi}(\mathbf{S}) \equiv \frac{1}{NFT} \text{median}(\text{uppertri}(\tilde{D}_{Spec}(\mathbf{S})))$ is the appropriately rescaled median off-diagonal value of L1 distances between the time-indexed concatenated spectra of the timeseries S_i , $i = 1, 2, \dots, N$.

Empirical Data and Simulation Regimes

Real Data

We used previously published [1] network timecourse data from a large multisite clinical resting-state fMRI study. Preprocessing and network identification followed protocols detailed in [1] that we simply outline here. Resting state functional magnetic resonance imaging data (160 volumes of echo planar imaging BOLD fMRI, TR = 2 sec.) was collected from 163 healthy controls (117 males, 46 females; mean age 36.9) and 151 age and gender matched patients with schizophrenia (114 males, 37 females; mean age 37.8) during eyes closed condition at 7 different sites across United States. After standard preprocessing, the fMRI data from all subjects was decomposed using group ICA into 100 maximally spatially independent spatial maps (<http://mialab.mrn.org/software>) of which 47 were identified as functionally meaningful networks. The networks fell into seven broad categories: sub-cortical (SC), auditory (AUD), visual (VIS), sensorimotor (SM), cognitive control (CC), default mode network (DMN) and cerebellar (CB). Subject specific spatial maps and timecourses were obtained from the group level spatial maps via spatio-temporal regression. The timecourses were detrended, despiked and subjected to additional postprocessing steps detailed in [1].

These timecourses (314 subjects, 47 networks, 158 timepoints), further filtered for frequencies at most 0.08 Hz and then z-scored, are referred to below as "Real Data". The average power at each frequency bin in $[0.003, 0.08]$ Hz for all network TCs for all subjects is denoted \bar{P}_{data} . The average cross-network covariance matrix for all subjects is denoted \bar{C}_{data} .

Every simulation regime described below consists of 1000 subjects, each characterized by a set of 47, length-158 timeseries.

SCC Gaussians: statistically stationary without constraint on SDTEs

Following [2], each simulated subject in the SCC Gaussian regime is a multivariate timeseries resulting from the projection of a 47×158 matrix of low-pass filtered white noise spectrally matched to \bar{P}_{data} onto the eigenspace of \bar{C}_{data} (Figure 6 from the main text (top middle)).

Covariance-Dynamic SCC Gaussians: piecewise stationary, two distinct covariance regimes, no constraint on SDTEs (CD-SCC Gaussian)

Each simulated subject in the CD-SCC Gaussian regime starts as a set of 47 length-158 low-pass filtered white noise timeseries (organized in a 47×158 matrix) spectrally matched to $\bar{\mathbf{P}}_{data}$. The multivariate timeseries is divided into three windows, determined by a middle window of randomly-chosen length between 40 and 60 TRs initially centered about the temporal midpoint $t_m = 79$ then translated a random distance in $[0,8]$ either forward or backward. This gives a middle window $([t_1, t_2], 40 \leq |t_2 - t_1| \leq 60, 41 \leq t_1 \leq t_2 \leq 117)$ that is roughly central but differs in extent and in degree of centrality between subjects. This simulation regime employs two target covariance matrices: one, as usual, is the mean covariance $\bar{\mathbf{C}}_{data}$ of the real data; the other \mathbf{C}_{weak} is modeled on $\bar{\mathbf{C}}_{data}$ but with much weaker connectivity (except for preserved variances along the diagonal) and an additive layer of very low-magnitude Gaussian noise. For each subject, the random length middle window described above is projected onto the eigenspace of $\bar{\mathbf{C}}_{data}$, a modularly structured covariance matrix, while the first and final windows are each projected onto the eigenspace of \mathbf{C}_{weak} . The resulting multivariate timeseries are explicitly undergoing a covariance shift as they enter and leave a 40-60TR window spanning the temporal midpoint of the scan (Figure 6 from the main text (bottom left)).

SCC Gaussians with Noise: statistically stationary with a single spike randomly inserted into a small proportion of network timeseries ("Noisy SCC Gaussian")

Each simulated subject in the Noisy SCC Gaussian regime starts as an SCC Gaussian subject (a 47×158 multivariate timeseries produced by subjecting white noise to spectral and covariance constraints exhibited by the real data). Of the 47 timeseries in this matrix, between 3 and 15 are selected at random to carry a single high frequency spike. This spike is a high-frequency, high-amplitude artifact centered at some randomly chosen timepoint in each of the selected networks. The entire multivariate timeseries contains between 3 and 15 of these noise artifacts, with at most one appearing in any given univariate timeseries (Figure 6 from the main text (top right)).

Spectrally and Statistically Nonstationary: explicitly nonstationary both statistically and epochally ("SS Nonstationary")

Each simulated subject in the Nonstationary regime starts as a SCC Gaussian subject (a 47×158 multivariate timeseries produced by subjecting white noise to spectral and covariance constraints exhibited by the real data). We designate 22 of the 47 simulated networks, (with row-indices corresponding to those of auditory, visual, sensorimotor and select cognitive control networks in \mathbf{C}_{dat}) as task-positive (TPNs). Another of 15 of the remaining 25 networks (with row-indices corresponding to default mode networks and select subcortical networks in \mathbf{C}_{dat}) as task-negative (TNNs). For each subject, a randomly selected 50%-75% of the TPNs and 50-75% of the TNNs are selected to exhibit stylized responsiveness to a hypothetical stimulus, leading each simulated subject to have at least 19 and most 29 responders among their 47 networks. Following the same procedure employed for the CD-SCC Gaussian regime, multivariate timeseries in the Nonstationary regime are divided into three windows whose endpoints are determined by a middle window of randomly-chosen length between 40 and 60 TRs initially centered about the temporal midpoint $t_m = 79$, then translated a random distance in $[0,8]$ either forward or backward. As mentioned above, this gives a middle window $([t_1, t_2], 40 \leq |t_2 - t_1| \leq 60, 41 \leq t_1 \leq t_2 \leq 117)$ that is roughly central but differs in extent and in degree of centrality between subjects. In those networks selected as responding TPNs, the

middle window is filtered for frequency content in [0.06,0.08] Hz, then rescaled to have amplitude slightly higher than the first and last windows. In responding TNNs, the middle window is filtered for frequency content in [0.006,0.05] Hz and rescaled to have amplitude slightly lower than the first and last windows. So TPNs get faster and stronger in response to the hypothetical task or stimulus, while TNNs go into a slower shallower activation regime. The 19-29 nonresponding networks remain as they were following the initial spectral and covariance-matching steps. This yields a 47×158 matrix in which 11-17 rows contain TPN timeseries, each with a fast high-amplitude middle window, 7-11 rows contain TNN timeseries, each with slow shallow middle window and 19-29 low-pass filtered (non-windowed) SCC Gaussians unchanged after the initial spectral and covariance-matching step. This set of multivariate timeseries is explicitly spectrally and statistically nonstationary (Figure 6 from the main text (bottom middle)).

Covariance-Dynamic Spectrally and Statistically Nonstationary: explicitly nonstationary both statistically and epochally with two distinct covariance regimes ("CD-SS Nonstationary")

Each simulated subject in the CD-Nonstationary regime starts as a Nonstationary subject as defined immediately above. However, following the CD-SCC Gaussian regime, this simulation regime employs the two, distinct target covariance matrices detailed above in the description of the CD-SCC Gaussian regime. For each subject, the random length middle window in which a subset of networks is spectrally perturbed (as detailed in the section immediately preceding) is projected onto the eigenspace of the modularly structured covariance matrix $\bar{\mathbf{C}}_{data}$ while the first and final windows are each projected onto the eigenspace of \mathbf{C}_{weak} . The resulting multivariate timeseries are explicitly undergoing spectral, amplitude *and* covariance shifts as they enter and leave a 40-60TR window spanning the temporal midpoint of the scan (Figure 6 from the main text (bottom right)).

Univariate and Multivariate Kurtosis

Univariate kurtosis, is the fourth statistical moment, $m_4 = \mathbb{E}((X - \mu_X)^4)$ of a random variable X , rescaled by the variance $m_2 = \mathbb{E}((X - \mu_X)^2)$ squared. For a normal random variable $\mathbb{E}\left(\frac{m_4}{m_2^2}\right) = 3$, so *excess univariate kurtosis* indicative of super-Gaussianity (unusually heavy tails) is given by $b = \frac{m_4}{m_2^2} - 3$. An unbiased estimator of excess univariate kurtosis [3] is:

$$\ell_n = \frac{b_n}{SE(b_n)} \quad (1)$$

where

$$b_n = \frac{n-1}{(n-2)(n-3)}((n+1)b + 6) \quad (2)$$

and

$$SE(b_n) = 2 \sqrt{\frac{6n(n-1)(n^2-1)}{(n-2)(n+1)(n-3)(n+3)(n+5)}} \quad (3)$$

Multivariate kurtosis is a generalization of univariate kurtosis introduced by [4]. For a given length- n multivariate process consisting of p univariate timeseries, Maria's multivariate kurtosis is defined as

$$B_{n,p} = \frac{1}{n} \sum_{i=1}^n (x_i - \mu)^T C^{-1} (x_i - \mu) \quad (4)$$

where

$$C_{i,j} = \frac{1}{n} \sum_{i=1}^n (x_i - \mu)(x_j - \mu)^T. \quad (5)$$

For a multivariate Gaussian process consisting of p univariate timeseries of length $n \rightarrow \infty$, $\mathbb{E}(B_{n,p}) = p(p+2)$. Excess multivariate kurtosis, $B_{n,p} - p(p+2)$ is evidence of multivariate super-Gaussianity. An unbiased estimator [5] of excess $B_{n,p}$ is:

$$\mathcal{B}_{n,p} = \frac{B_{n,p} - \frac{p(p+2)(n-1)}{(n+1)}}{\sqrt{\frac{8p(p+2)}{n}}}. \quad (6)$$

$\mathcal{B}_{n,p}$ has a standard normal distribution $\mathcal{N}(0,1)$ when the underlying process is actually multivariate Gaussian, so the upper tail represents strong evidence against the null hypothesis that all observations all arose from fixed Gaussian distributions in each dimension.

Sliding Window Dynamic Connectivity

Short-timescale network or region of interest (ROI) connectivity estimates evaluated on successive sliding windows through activation timecourses are a commonly employed [1, 6-9] vehicle through which to investigate so-called dynamic connectivity in resting-state fMRI. The general idea is straightforward: starting from a set of R length- T network or ROI timecourses that emerges from a standard pre-preprocessing pipeline for rs-fMRI, slide a window of fixed length L vertically through each subject's $T \times R$ multivariate timeseries, advancing stepwise by some increment l until the whole timeseries is exhausted. The window can be rectangular or have tapered edges; the pipeline that selects networks or ROIs and processes the resulting timeseries is generally unaltered relative to the non-windowed setting, and the metric of connectivity – often but not always correlation or covariance – is typically not unaltered from the non-windowed setting, in part to allow for comparisons between static (i.e., scan-length) findings and windowed (short time-scale) findings. There is considerable debate [2, 10, 11] regarding appropriate window-length subject for example, to the spectral content of the signal and other considerations. Due to the number of connectivity measurements this approach generates per subject, it is common to attempt to summarize the short-timescale connectivity patterns in the study by clustering, using k -means, the entire set of windowed observations, leading to some collection of k summary connectivity states. These k connectivity states summarizing transient connectivity patterns in the entire population yield easily computable information (e.g., occupancy rates, mean dwell times and transition probabilities) about subject-level time-varying connectivity.

As indicated above, in the section about empirical data and simulation regimes, we used previously published [1] network timecourse data from a large multisite clinical resting-state fMRI study as the empirical basis for simulation models. The original study filtered timecourses for spectral content under $[0.003, 0.125]$ Hz, but in keeping with [2], in this paper we filtered for content $[0.003, 0.08]$ Hz. Otherwise our initial pipeline leading up to windowing was identical to that published in [1] (and outlined above in the section on empirical data and simulation regimes). Again following the published study [1], our windows had length 22 TR and were advanced by 1 TR at each step, leading to a total of 136 windows per subject.

Although [1] used Gaussian tapering, here we employed rectangular windows. There was no discernible difference between the two approaches, so the simpler approach was utilized. Connectivity between networks on each window is measured as pairwise correlation between the windowed network timecourses. In the original study, the elbow criterion suggested $k = 5$ clusters for this data, a choice that we retain here.

For consistency and comparability between empirical data and the various simulation regimes, all simulation regimes were windowed and clustered (using $k = 5$) with the same protocols and parameters as the empirical data

- [1] E. Damaraju *et al.*, "Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia," *Neuroimage: Clinical*, vol. 5, pp. 298-308, 2014.
- [2] T. O. Laumann *et al.*, "On the Stability of BOLD fMRI Correlations," *Cereb Cortex*, Sep 02 2016.
- [3] M. K. Cain, Z. Zhang, and K. H. Yuan, "Univariate and multivariate skewness and kurtosis for measuring nonnormality: Prevalence, influence and estimation," *Behav Res Methods*, Oct 17 2016.
- [4] K. V. Mardia, "Measures of Multivariate Skewness and Kurtosis with Applications," (in English), *Biometrika*, vol. 57, no. 3, pp. 519-&, 1970.
- [5] T. Sumikawa, K. Koizumi, and T. Seo, "Measures of multivariate skewness and kurtosis in high-dimensional framework," *Hiroshima Statistical research Group Technical Reports*, Technical Report vol. 13, no. 3, p. 22, 2013.
- [6] B. M. Baczowski, T. Johnstone, H. Walter, S. Erk, and I. M. Veer, "Sliding-window analysis tracks fluctuations in amygdala functional connectivity associated with physiological arousal and vigilance during fear conditioning," *Neuroimage*, vol. 153, pp. 168-178, Jun 2017.
- [7] E. Premi *et al.*, "Dynamic connectivity state oscillations in presymptomatic GRN disease: from Connectome to Chronnectome," (in English), *Journal of Neurochemistry*, vol. 138, pp. 381-381, Aug 2016.
- [8] L. Douw, D. G. Wakeman, N. Tanaka, H. S. Liu, and S. M. Stuffebeam, "State-Dependent Variability of Dynamic Functional Connectivity between Frontoparietal and Default Networks Relates to Cognitive Flexibility," (in English), *Neuroscience*, vol. 339, pp. 12-21, Dec 17 2016.
- [9] E. Allen, E. Damaraju, S. M. Plis, E. Erhardt, T. Eichele, and V. D. Calhoun, "Tracking whole-brain connectivity dynamics in the resting state," *Cereb Cortex*, vol. 24, no. 3, pp. 663-676, 2014.
- [10] R. Hindriks *et al.*, "Can sliding-window correlations reveal dynamic functional connectivity in resting-state fMRI?," (in English), *Neuroimage*, vol. 127, pp. 242-256, Feb 15 2016.
- [11] S. Shakil, C. H. Lee, and S. D. Keilholz, "Evaluation of sliding window correlation performance for characterizing dynamic functional connectivity and brain states," (in English), *Neuroimage*, vol. 133, pp. 111-128, Jun 2016.