

Bistability and Non-Gaussian Fluctuations in Spontaneous Cortical Activity

Frank Freyer,^{1,2} Kevin Aquino,^{4,6} Peter A. Robinson,^{4,5,7} Petra Ritter,^{1,2,3} and Michael Breakspear^{6,8}

¹Berlin Neuroimaging Center and Department of Neurology, Charité Universitätsmedizin, and ²Bernstein Center for Computational Neuroscience, and ³Graduate School of Mind and Brain and Mind and Brain Institute, Faculty of Philosophy, Humboldt University, 10117 Berlin, Germany, ⁴School of Physics and ⁵Faculty of Medicine, University of Sydney, Sydney, New South Wales 2006, Australia, ⁶School of Psychiatry, University of New South Wales and The Black Dog Institute, Sydney, New South Wales 2031, Australia, ⁷Brain Dynamics Center, Westmead Millennium Institute, Westmead Hospital and University of Sydney, Westmead, New South Wales 2145, Australia, and ⁸Queensland Institute of Medical Research and the Royal Brisbane and Woman's Hospital, Brisbane, Queensland 4006, Australia

The brain is widely assumed to be a paradigmatic example of a complex, self-organizing system. As such, it should exhibit the classic hallmarks of nonlinearity, multistability, and “nondiffusivity” (large coherent fluctuations). Surprisingly, at least at the very large scale of neocortical dynamics, there is little empirical evidence to support this, and hence most computational and methodological frameworks for healthy brain activity have proceeded very reasonably from a purely linear and diffusive perspective. By studying the temporal fluctuations of power in human resting-state electroencephalograms, we show that, although these simple properties may hold true at some temporal scales, there is strong evidence for bistability and nondiffusivity in key brain rhythms. Bistability is manifest as nonclassic bursting between high- and low-amplitude modes in the alpha rhythm. Nondiffusivity is expressed through the irregular appearance of high amplitude “extremal” events in beta rhythm power fluctuations. The statistical robustness of these observations was confirmed through comparison with Gaussian-rendered phase-randomized surrogate data. Although there is a good conceptual framework for understanding bistability in cortical dynamics, the implications of the extremal events challenge existing frameworks for understanding large-scale brain systems.

Introduction

The characterization of spontaneous or “resting-state” data has been traditionally overshadowed by the study of task-related activity, such as the evoked response potential of psychophysiology. However, there is growing interest in resting-state data (Biswal et al., 1995), both as an influential precursor to sensory processing (Makeig et al., 2002) or motor activity (Fox et al., 2006) and as a window into the highly structured nature of spontaneous large-scale brain activity per se (Greicius et al., 2003; Achard et al., 2006). The use of resting-state electroencephalogram (EEG) data is ideal in these regards, because its high temporal resolution allows the study of activity from long to very short timescales. Lengthy acquisitions of such data permit estimation of the statistics of activity at various scales, including the mean power (or power spectra) plus higher-order moments, as reflected in long-range temporal (Linkenkaer-Hansen et al., 2001) and spatial interdependences (Breakspear and Terry, 2002). However, the only

measure traditionally reported is the mean power (i.e., the power spectrum). This leaves the nature of spontaneous fluctuations around the mean unknown. An understanding of these fluctuations may have far-reaching methodological and computational consequences.

We hence examine the likelihood distribution of power fluctuations across a range of timescales (0.5–35 Hz) in human resting-state EEG and seek low-dimensional parametric forms at each frequency. In particular, we test the null hypothesis that the distribution at all frequencies could be well described by the simple one-parameter family of exponential distributions, which contain no additional information that could be inferred from knowledge of the mean. Such distributions arise when the underlying system generates uncorrelated stochastic events and have near-maximum entropy (i.e., are close to thermodynamic equilibrium). This would accord with the view that cortical activity is well described as filtered Gaussian noise, justifying the sole use of the power spectra when reporting the properties of scalp EEG data. Rejection of this null hypothesis would imply the presence of non-Gaussian processes in the cortex and suggest the need for more sophisticated generative models.

Stochastic systems that generate Gaussian statistics are typically close to equilibrium, dominated by short-range interactions and characterized by weak global correlations (Tsallis and Brigatti, 2004). They may comprise complex local units but are of sufficiently high dimension that their large-scale statistics are nonetheless essentially uncorrelated. However, the brain has

Received Jan. 29, 2009; revised May 26, 2009; accepted May 27, 2009.

P.A.R. and K.A. were supported by the Australian Research Council (ARC). P.R. was supported by the German Federal Ministry of Education and Research (Berlin Neuroimaging Center; Bernstein Center for Computational Neuroscience, Bernstein Focus State Dependencies of Learning) and the Robert Bosch Foundation. M.B. acknowledges the support of Brain Network Recovery Group Grant JSMF22002082 and ARC Grants DP0667065 and TS0669860. We thank A. R. McIntosh, V. K. Jirsa, and K. Friston for helpful discussions.

Correspondence should be addressed to Michael Breakspear, Queensland Institute for Medical Research, Herston, QLD 4006, Australia. E-mail: Michael.Breakspear@qimr.edu.au.

DOI:10.1523/JNEUROSCI.0754-09.2009

Copyright © 2009 Society for Neuroscience 0270-6474/09/298512-13\$15.00/0

dense internal connectivity with long-range projections that scale hierarchically (Hilgetag and Kaiser, 2004; Sporns and Zwi, 2004). Through its abundant use of energy, the brain resides in a strongly non-equilibrium state. These are the hallmarks of complex systems, which exist in a wide variety of physical and biological fields and exhibit self-organization, multiple excited modes and highly correlated, nondiffusive processes (Zaslavsky, 2002). From this perspective, it might seem remarkable that, in contrast to many complex systems, the brain would exhibit only simple diffusive dynamics at the global scale. However, current computational models of neural population dynamics, such as the Fokker–Plank formulation (Renart et al., 2003; Deco et al., 2008), are premised on the “diffusion approximation,” namely that the inputs impinging on individual neurons within a population can be treated as temporally uncorrelated when computing instantaneous postsynaptic firing rates. Likewise, Bayesian formulations of human inference propose that cortical populations compute via uncorrelated Poisson processes (Ma et al., 2006). Ultimately, the spatial and temporal scales at which these assumptions are valid are empirical questions that can only be determined by examining the statistics of *in vivo* neuronal activity. The goal of the present study is to address this at the macroscopic spatial scale.

Materials and Methods

Subjects and data acquisition

Scalp EEG data were collected from a total of 16 healthy, adult volunteers (11 females; mean age, 25.3 years; range, 20–31 years). Written informed consent was obtained from each subject before their participation. Subjects were requested to rest with eyes closed while maintaining alertness. Acquisition times ranged in duration from 14 to 30 min.

All data were acquired using magnetic resonance-compatible amplifiers (hardware bandpass filter, 0.1–250 Hz; BrainAmp; Brain Products) and EEG caps (Easy-Cap; FMS) arranged according to the International 10–20 System, referenced against an electrode centered between Cz and Pz. Impedances of all electrodes were set below 5 k Ω . In seven subjects, EEG data were acquired from 60 scalp channels in a sound- and light-attenuated room. In nine subjects, EEG data were acquired from 29 scalp channels and recorded simultaneously with functional magnetic resonance imaging (fMRI) in a 1.5 T MR Vision tomograph (Siemens) using a T2*-weighted blood oxygen level-dependent (BOLD)-sensitive gradient echo planar imaging sequence (scan repetition time, 2200 ms; acquisition time, 2050 ms; echo time, 60 ms; volumes, 750; slices, 20; flip angle, 90°; voxel size, 3 × 3 × 5 mm). For additional details of the EEG–fMRI setup, see Moosmann et al. (2003) and Ritter et al. (2009a). EEG sampling rate was 5 kHz. In all cases, the electrooculogram and electrocardiogram were also recorded.

Data preprocessing

The nine datasets acquired simultaneously with fMRI contain artifacts from MR gradient switching (Ritter et al., 2009b) and ballistocardiogram effects, which are generated by heart rate-associated movements of the electrodes in the static B0 field of the MR scanner. Both types of artifacts were corrected using template estimation and subtraction (Allen et al., 1998, 2000; Ritter et al., 2007) as implemented in the Vision-Analyzer software (version 1.03; Brain Products).

Temporal independent component analyses (ICAs) were performed on all 16 datasets. Components reflecting eye movement and scalp muscle artifact were identified and subtracted from the data. Data were then visually inspected to ensure they were artifact free. For this study, we characterized temporal fluctuations from the right occipital region (channel O2). These single-channel data were low-pass filtered (70 Hz high cutoff frequency) and downsampled to 200 Hz for subsequent analysis.

Data analysis

Wavelet decompositions and derivation of empirical probability distribution functions. Wavelet-based estimators of the dynamic spectral content

of time series are a well established approach, with robust theoretical support (Torrence and Compo, 1998) and recent extensions to inference on dynamic coherence (Maraun et al., 2007). For the present study, dynamic spectrograms were derived by convolving the data with complex Morlet wavelets (center frequency, 1 Hz; bandwidth parameter, 10 s; corresponding to a kernel with a full-width at half-maximum of 0.5 s at 10 Hz). Conversion from scales to frequency was estimated as the inverse of the scaling parameter at each level of the wavelet decomposition, yielding a frequency range from 0.5 to 35 Hz in steps of 0.5 Hz. Power was estimated as the modulus squared of the corresponding wavelet coefficients. Frequency-specific probability distribution functions (PDFs) were then obtained by partitioning the fluctuations of power separately at each frequency and counting the number of observations in each bin. Bin width was determined individually at each frequency by dividing the frequency-specific power range into 200 equally sized bins.

Estimating and fitting exponential PDFs. A Gaussian process is expected when the observed events arise independently from a stochastic source. According to the central limit theorem, such events need not be drawn from a single distribution with stationary statistics but may be drawn from a large (limiting to infinite) number of independent processes that, by themselves, need not be Gaussian, provided they have finite variance. In the context of scalp EEG, which reflects summed activity across a number of large neuronal populations, this implies that the temporal fluctuations in activity may be Gaussian even if neuronal activity on smaller scales is not Gaussian (e.g., if it is a nonlinear process) as long as they are statistically independent when observed at the macroscopic scale. In the present study, this forms the null hypothesis. Rejection of this null hypothesis attributable to non-Gaussian statistics is expected when such spatially distributed macroscopic processes become sufficiently correlated.

For processes exhibiting Gaussian fluctuations in amplitude, the functional form of the corresponding power distribution (i.e., the amplitude squared) follows an exponential PDF (Balakrishnan and Basu, 1996):

$$P_x(x) = \lambda e^{-\lambda x}, \quad (1)$$

where x is the power, and λ is the shape parameter that can be estimated from an empirical distribution by taking the log of the likelihood and estimating the slope of the resulting line. For the present purposes, this was achieved by a linear regression of the PDF in linear–log coordinates. High variance at the lower end of the log-likelihood line typically arises as a result of the relatively infrequent observations. We hence iteratively excluded such bins until a goodness-of-fit threshold ($R^2 > 0.95$) was achieved. It is important to note that this procedure was only pursued to provide a best estimate of the shape parameter λ when Equation 1 was at least approximately satisfied, not to make any inferences regarding the best functional form for the distribution. To gain a better insight into the functional form of the PDF, particularly the asymptotic scaling behavior of both tails, the fitted PDFs were formally evaluated in log–linear and log–log coordinates. For the data, this was achieved by partitioning the range of the log of the power into equally sized bins and counting the number of observations in each bin. For the fitted distributions, this required a change of variables to $y = \log(x)$ and, hence,

$$P_y(y) = \left| \frac{dx}{dy} \right| P_x[x(y)] = \exp(y - \lambda e^y). \quad (2)$$

Independent component decomposition. It is possible that deviations from an exponential fit could reflect multiple independent spatial sources (i.e., cortical regions) contributing in a complex, heterogeneous manner to scalp channels. We hence also applied the estimation and fitting procedure to time series data obtained from ICA of the whole-head EEG (Makeig et al., 1996). This algorithm produces the maximally temporally independent signals available in the channel data, hence unmixing spatially discrete sources of the scalp channel data.

Phase-randomized surrogate data. Non-negligible residual errors and biases in an empirical fit to an exponential PDF may arise from trivial causes such as a finite sample set, measurement bias, or methodological processes, such as wavelet-based estimates of moment-to-moment power in a colored noise dataset. To test the null hypothesis that observed

deviations are attributable to such trivial effects, we repeated the analysis on surrogate data constructed by Fourier-based phase randomization of the original EEG data (Theiler et al., 1992). Such a process preserves the mean power spectrum of the original data while rendering its statistics Gaussian. Two hundred such surrogate datasets were computed for each subject and used to represent the null distribution.

Results

Representative distributions and parametric fits at different temporal scales are presented first. A systematic survey across the frequency domain is then provided, together with a formal model-based comparison of Gaussian and candidate non-Gaussian fits on empirical and phase-randomized surrogate data.

Exemplar results: Gaussian, bimodal, and super-Gaussian

In this section, we present examples of empirical PDFs of power across a range of different temporal frequencies.

Example of a simple unimodal exponential distribution

At some temporal scales, empirical distributions of the log power were observed to exhibit a unimodal form with featureless tails above and below the mean. Exponential PDFs fitted to these distributions capture the observed variability across the entire power domain. Figure 1 shows such an example from a single subject at 34 Hz, of the power-likelihood distributions in log-linear (Fig. 1*a*) and log-log (Fig. 1*b*) coordinates. In this example, the exponential PDF is an excellent quantitative estimate of the data, as evident in both log-linear and log-log space. Importantly, there are no obvious systematic deviations (biases) evident above or below the mean. Residual error variance increases near the lower tail (left-hand), but this is spread across both sides of the tail and hence does not show a systematic bias. Increased variance is expected at the ends of the tails because of the relatively few observations.

The observed and fitted distributions derived from the same data, after phase randomization, are plotted in Figure 1, *c* and *d*. It is hence evident that the exponential PDF closely fits a time series that has been deliberately rendered Gaussian. Once again, there are some unbiased residual errors at the lower left tail. Although fitting phase-randomized data with an exponential PDF is an expected result, it is nonetheless important to verify empirically because it confirms that combined use of the wavelet decomposition and curve-fitting procedures used have *prima facie* validity, that is, they do capture the nature of a distribution when it is known to be Gaussian by construction. Comparison of Figure 1, *a* and *b*, with Figure 1, *c* and *d*, also shows that, in this example, phase randomization had no discernable effect on the statistics of the original empirical distribution. This accords with a view that power fluctuations in this subject, at this scale, likely arise from uncorrelated stochastic processes.

Examples of bimodal exponential distributions

In most subjects, the appearance of two distinct modes in the alpha rhythm (8–12 Hz) marks a striking deviation from the unimodal exponential distribution shown above. An example of

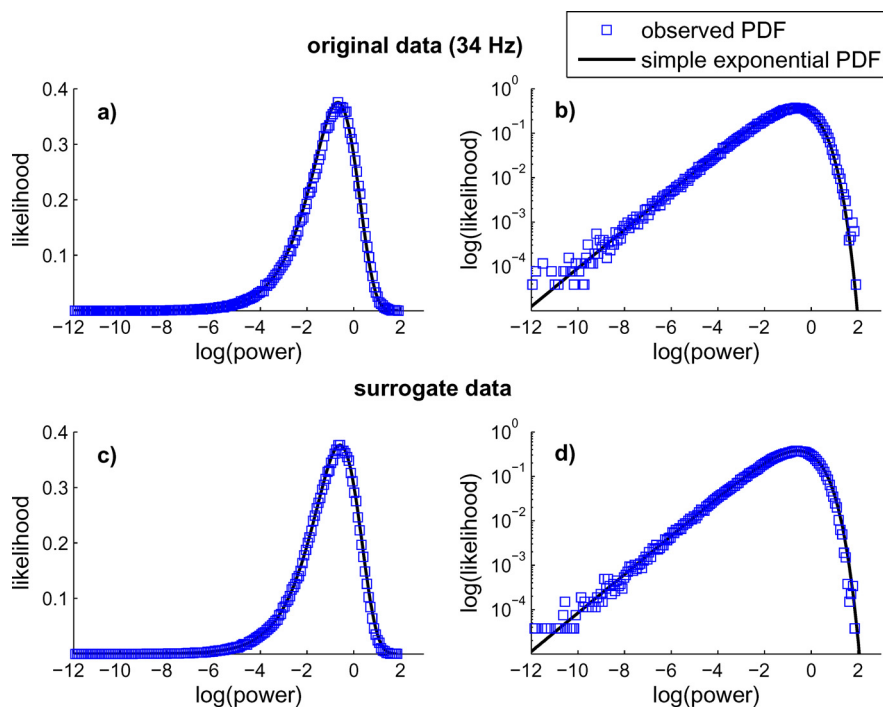


Figure 1. Example of a simple unimodal distribution of power at 34 Hz in a single subject (S10). *a, b*, Observed PDF in log-linear (*a*) and log-log (*b*) coordinates obtained from original data. Blue squares show the observed distribution, and the black line shows the fitted exponential function (Eq. 1). *c, d*, The observed and fitted distributions obtained from surrogate data in log-linear (*c*) and log-log (*d*) coordinates.

an observed bimodal distribution in a single subject at 10 Hz in log-linear and log-log coordinates is given in Figure 2. To model this distribution, a single-exponential PDF was regressed to fit the dominant mode (black line), which in this case is the mode with greatest power. A second exponential distribution was then estimated from the residuals of this fit (gray line). It can be seen that the two exponentials provide a remarkably good fit, as evident by comparing their sum (red line) with the original distribution. Figure 2, *c* and *d*, shows the same data rendered Gaussian through phase randomization. The second mode has been abolished and the surrogate data hence affords a close fit with a single-exponential PDF. Forcing a second exponential fit to the residual errors in the phase-randomized case actually renders a worse fit than one mode alone, as evident in the large errors of the summed exponentials (red line) in Figure 2, *c* and *d*.

To ensure that the abolishment of the second mode is not an artifact of using the Fourier-based phase-randomization method of producing surrogate data, we also inspected PDFs derived from a wavelet-based resampling technique (Breakspear et al., 2003). Briefly, these surrogate data are created by resampling the coefficients of a discrete wavelet decomposition, using a low-order discrete wavelet basis set (we used the Daubechies wavelet of order 3). The results accorded with those derived from the Fourier approach. Supplemental Figure S1 (available at www.jneurosci.org as supplemental material) shows the outcome of applying this alternative procedure to the same data as shown in Figure 2, showing that the outcome is not dependent on the particular surrogate algorithm used.

Clear visually evident bimodal distributions of alpha activity (8–12 Hz) were observed in 13 of the total 16 subjects whose data was acquired and inspected. These distributions are shown in supplemental Figure S2 (available at www.jneurosci.org as supplemental material). Data from two other subjects (S05 and S06)

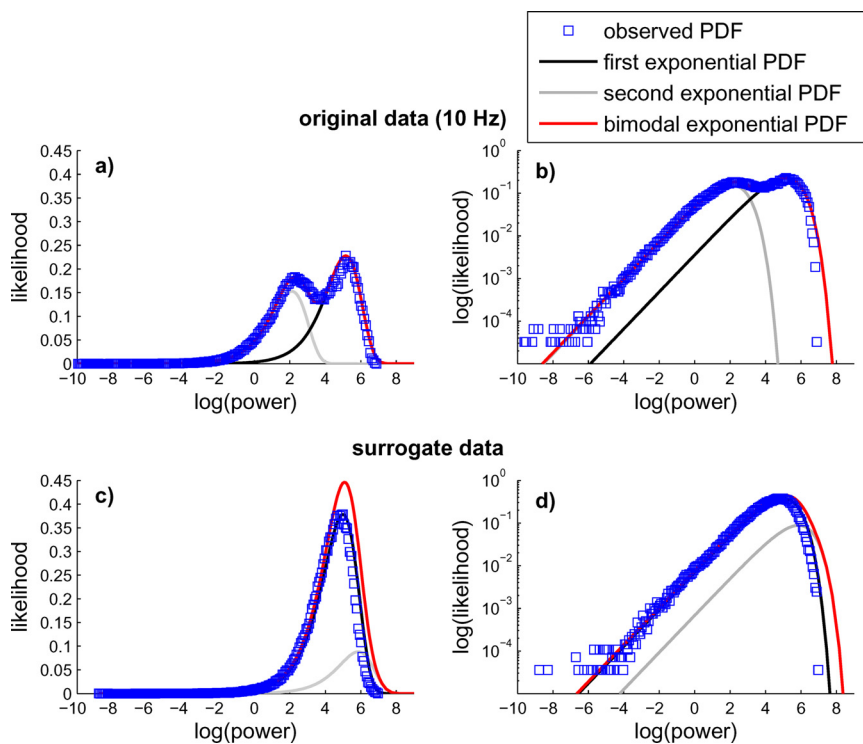


Figure 2. Example of strongly bimodal distribution in a single subject (S11) at 10 Hz. Panels as per Figure 1. The red line shows the sum of the two unimodal forms.

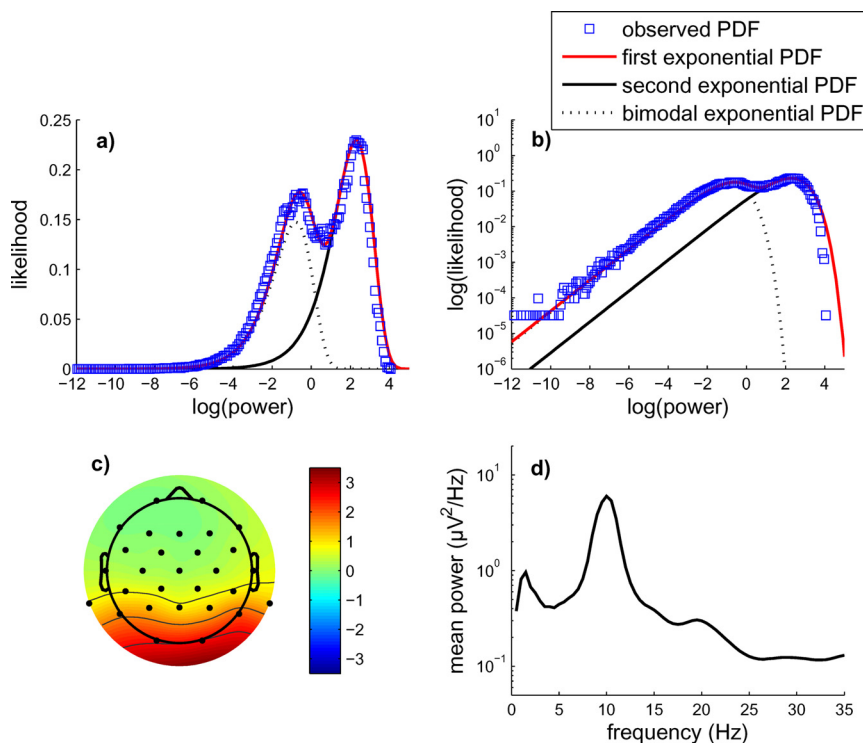


Figure 3. *a, b*, Example of a bimodal distribution in a single independent component in log–log (*a*) and log–linear (*b*) coordinates at 10.5 Hz. *c*, Spatial distribution of this component. Color bar shows weighting of component over scalp electrodes (dimensionless units). *d*, Power spectrum estimated from wavelet coefficients.

exhibited somewhat less distinct bimodal distributions at 8 and 23 Hz. For completeness, the PDFs for these two subjects at the respective (non-alpha) frequencies where they were observed are shown in supplemental Figure S3 (available at www.jneurosci.org

as supplemental material). One additional subject (S07) had no clear bimodal distribution within any of the frequencies studied.

To address the possibility that the bimodal distributions arise from two spatially distinct sources, each with different amplitude statistics, we analyzed time series derived from independent component decompositions of the full multichannel EEG, identifying the single dominant posterior alpha component. In subjects who showed a bimodal alpha distribution, the corresponding posterior alpha mode was invariably bimodal, typically with even greater separation of the modes. Figure 3 shows an example of a bimodal ICA distribution, again in log–log (*a*) and log–linear (*b*) coordinates taken from a single subject at 10.5 Hz. The spatial topology of this ICA mode (Fig. 3*c*) shows a single posterior source, consistent with the principles that ICA unmixes single (linearly mixed) brain sources of scalp activity (Makeig et al., 1996). The power spectrum of this component (Fig. 3*d*) evidences a strong peak in the alpha rhythm, albeit obscuring the bimodal nature of this peak.

These observations suggest that the two modes arise through a temporally heterogeneous process, whereby two distinct dynamical processes arising from a single cortical region are alternately expressed. Figure 4 shows examples of the temporal expression of bimodal distributions in two representative subjects at their peak alpha frequency. In the first subject (Fig. 4*a–c*), the second (higher-power) mode is expressed erratically throughout the entire 500 s example (Fig. 4*a*). The time series exhibits a burst-like behavior from the lower-power first mode (Fig. 4*b*), evident as sudden increases in the magnitude of the first temporal derivative of power (Fig. 4*c*). Each burst is weakly visible in the raw electrode time series. In this subject, this bursting behavior is present throughout the entire dataset. In contrast, data from the second subject (Fig. 4*d–f*) shows two distinct periods of time ($t < 180$ s and $t > 1000$ s) when the second (higher-power) mode is constantly present. In the interim period (180 s $< t < 1000$ s), there are sporadic instances when the second mode appears to be expressed, although these are typically lower in amplitude. Greater irregularity is also visible in the raw electrode time series (Fig. 4*a,d*), underlining the gross, macroscopic impact of alpha rhythm bimodality.

Although the bimodal distribution clearly affords a far better fit to the data, it does include a second parameter. To formally compare the fits, we used the Bayesian information criterion (BIC), which includes a penalty term for model complexity:




Figure 4. Examples of mode switching in two subjects (top, S14; bottom, S11). *a–f*, Time series for the original O2 electrode data (*a, d*), the power (amplitude squared) of the wavelet coefficients for this frequency (*b*, 10 Hz; *e*, 10.5 Hz), and its temporal derivative (*c, f*). Panels have gray shading whenever the system is in the high-amplitude mode. The red line in *b* and *e* show the power at which the modes have equal likelihoods. This threshold was derived from the crossing of the two PDFs (see Fig. 2).

$$\Lambda = n \ln \left(\frac{\text{RSS}}{n} \right) + k \ln(n), \quad (3)$$

where RSS is the sum of the squared residuals, n is the number of observations (discrete power bins), and k is the number of free parameters ($k = 1$ for the unimodal fit and $k = 2$ for the bimodal fit). Given two or more candidate models, the “best” model will yield relatively low values, reflecting small residual variance after penalization for the number of free parameters. Figure 5 shows a (single-subject) example of the difference between the BIC values for a unimodal versus a bimodal fit across all frequencies studied. Hence, a strongly positive value favors the bimodal fit, whereas a strongly negative value favors a unimodal fit. Results for the empirical data are shown in black squares. The mean and 95% confidence intervals, estimated parametrically from 200 surrogate datasets, are depicted in red circles.

The stand-out feature in the original data is a strong upward deflection in the BIC difference between 8 and 11.5 Hz. This supports the results clearly visible in Figure 2 (the same subject)

while additionally providing a quantitative measure of model preference and indicating the approximate frequency domain over which the second mode is expressed. In stark contrast, results from the surrogate data are consistently negative, indicating that, after penalizing for the second parameter, the bimodal fit is statistically inferior to a single-exponential fit for the unimodal distribution that is invariably generated from surrogate data. The original data also show a second peak at ~ 14.5 – 15.5 Hz. Visual inspection of the PDF at these frequencies reveals only weak evidence for a second mode, evident as a point of inflection in the right-hand tail. The significance of the broader positive region of the curve in the beta range (13–30 Hz) is considered below.

In systems with multistability, the distribution of times for which the system “dwells” in each mode can also be informative of the possible nature of the underlying process (Nakamura et al., 2007). Figure 6 shows the cumulative distribution of dwell times derived from the two same subjects whose data was shown in Figure 4. It can be seen that dwell-time distributions of each

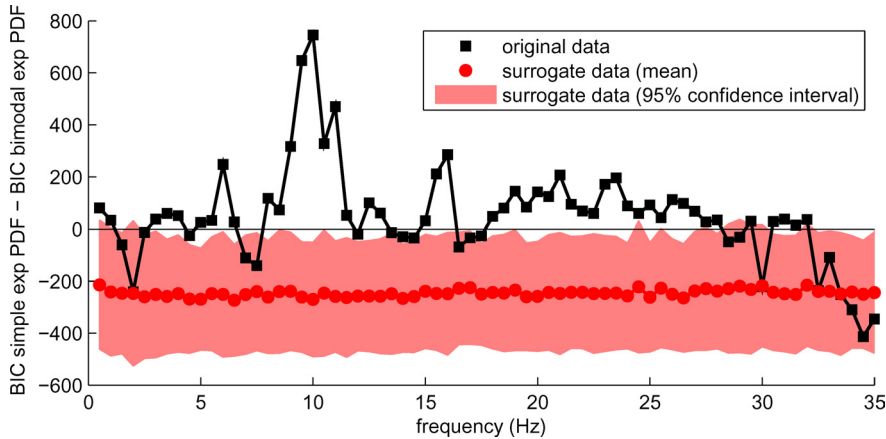


Figure 5. Difference between Bayesian information criterion for unimodal and bimodal fits in a single subject (S11). Positive (negative) values indicate preference for bimodal (unimodal) exponential PDF. Results from original data are shown in black. Mean and 95% confidence intervals for surrogate data are shown in red.

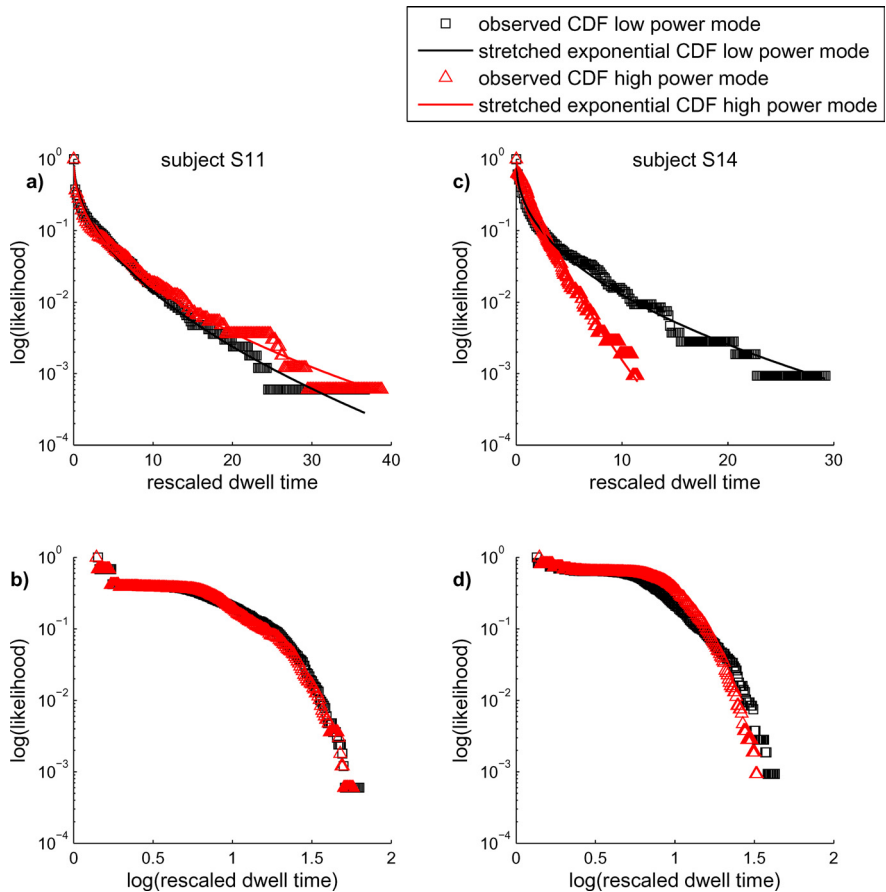


Figure 6. Representative CDFs for the bimodal dwell times in two subjects, rescaled to their mean. **a** and **b** show data from subject S11, and **c** and **d** show data from subject S14. Empirical distributions for low-power and high-power modes are given as black squares and red triangles, respectively, in linear–log (**a, c**) and log–log (**b, d**) coordinates. Stretched exponential CDFs are plotted as solid lines in the linear–log coordinates with $\alpha = 1.45$, $\beta = 0.45$ for the low-power (black) mode and $\alpha = 1.35$, $\beta = 0.5$ for the high-power (red) mode in the first subject (**a**). The parameters in the second subject (**b**) are $\alpha = 1.55$, $\beta = 0.45$ for the low-power (black) mode and $\alpha = 1.15$, $\beta = 0.75$ for the high-power (red) mode.

mode (which have been rescaled to their mean) appear to follow similar functional forms. When viewed in log–log coordinates (Fig. 6*b,d*), it is clear that the dwell-time distributions in these subjects do not follow a power-law relationship. When viewed in linear–log coordinates (Fig. 6*a,c*), the distributions instead suggest a gamma-type stretched exponential function:

$$P(x) = (-\alpha x^\beta). \quad (4)$$

Solid lines in each panel show fits to the respective data. As evident in the form of Equation 4, the parameter β determines the curvature of the curves (in linear–log space). When β is close to 1, Equation 4 approaches a basic exponential form $e^{-\alpha x}$, whereas if β approaches 0, the curve becomes highly concave (in linear–log coordinates), trending toward a power-law effect. In the first subject (Fig. 6*a,b*), $\beta = 0.45$ for the low-power mode and $\beta = 0.5$ for the high-power mode. For the second subject, the corresponding values are $\beta = 0.45$ (low-power mode) and $\beta = 0.75$ (high-power mode). Hence, the dwell times for the high-power mode in this subject are closer to a basic exponential form. Similar fits were obtained in all 13 subjects showing a bimodal alpha pattern. These are shown in supplemental Figure S4 (available at www.jneurosci.org as supplemental material). For the lower-power mode, β was generally close to 0.5 (mean, 0.49; range, 0.40–0.60), whereas for the higher mode, β was higher in all subjects except one (mean, 0.68; range, 0.50–0.90).

Examples of unimodal non-Gaussian distributions

In all subjects and across a broad range of frequencies, typically in the beta range (~13–30 Hz), the empirical PDFs frequently show a consistent upward bias of the right-side tail (i.e., above the mean) from exponential toward power-law scaling. Figure 7 provides examples of this effect in two subjects. It is critical to note that the empirical distributions (Fig. 7*a,b*) show a systematic bias and not simply large unbiased residuals over the right-hand tails. In clear contrast, this bias is absent in PDFs derived from the surrogate data (Fig. 7*c,d*).

Research into complex, highly correlated systems motivates the consideration of an alternative stochastic model (Bramwell et al., 1998). Above the mean, fluctuations in these systems scale according to the following functional form:

$$P(x) = [\exp(x - \lambda e^x)]^a, \quad (5)$$

where the parameter λ is estimated in the same manner as for the exponential PDF. Setting $a = 0.5$ yields the heavy right-hand tail that has been empirically observed in data from numerous complex systems (Bramwell et al., 2000). This double-exponential form, also known as the “Fisher–Tippett distribution,” clearly captures the trend in the empirical PDF, correcting the bias resulting from a simple-exponential fit above the mean. Phase randomization returns the tails back to simple-exponential scaling. Similar results

Phase randomization returns the tails back to simple-exponential scaling. Similar results

were seen at many frequencies in all subjects and are presented in supplemental Figure S5 (available at www.jneurosci.org as supplemental material).

By fixing $a = 0.5$, Equation 5 becomes a one-parameter double exponential and can hence be compared directly against the simple exponential without need for parameter penalization. Figure 8 shows the ratio of the RSS for the exponential and double-exponential fits in a single subject. In this subject, the ratio is relatively small (<10) for frequencies below 16 Hz and above 30 Hz. However, within the beta band, between 16 and 30 Hz, there is a very strong increase in the ratio, reflecting a decrease in the residual error for the Fisher–Tippett fit in this range and hence a corresponding increase in the relative log evidence, peaking at 21 Hz.

In comparison, the ratios derived from the surrogate data (shown in red) are uniformly small, typically <1 , with low variability between surrogate realizations. However, in this subject, as in all others, the ratio of the residuals in the empirical data is outside the surrogate distribution at most frequencies, although they are relatively low (<10). Visual inspection of the data shows that this is typically attributable to a small number of outlier observations, typically near the very edges of the right-hand tail, and not to a consistent bias as evident in Figure 7 and supplemental Figure S3 (available at www.jneurosci.org as supplemental material). Such outlier observations increase the residual error for both the simple- and double-exponential PDFs. However, the relative ratio of the RSS is invariably only large (>10) when there is a consistent upward bias of the right-hand tail associated with very small errors for the Fisher–Tippett fit.

The preferential performance of the Fisher–Tippett fit over the exponential PDF for the right-hand tail implies that large-amplitude (“extremal”) events occur far more frequently than expected from the (null) simple-exponential form. Such events can be clearly seen in time series of power at the frequencies exhibiting Fisher–Tippett statistics versus the same data after phase randomization (Fig. 9*a,b*). Notably, compared with amplitude dynamics at frequencies showing bimodal distributions (Fig. 4), power fluctuations here do not exhibit obvious irregular switching because they are still drawn from a unimodal distribution. Hence, fluctuations are sampled from a single underlying mode. The critical difference is that their log likelihood falls less rapidly than an exponential functional form, and hence large-amplitude “events” are sporadically observed. Hence, the kurtosis—a measure of the width of the tails—is increased approximately twofold.

What are the characteristics of extremal events in the time–frequency plane? By calculating the ratio of the simple-

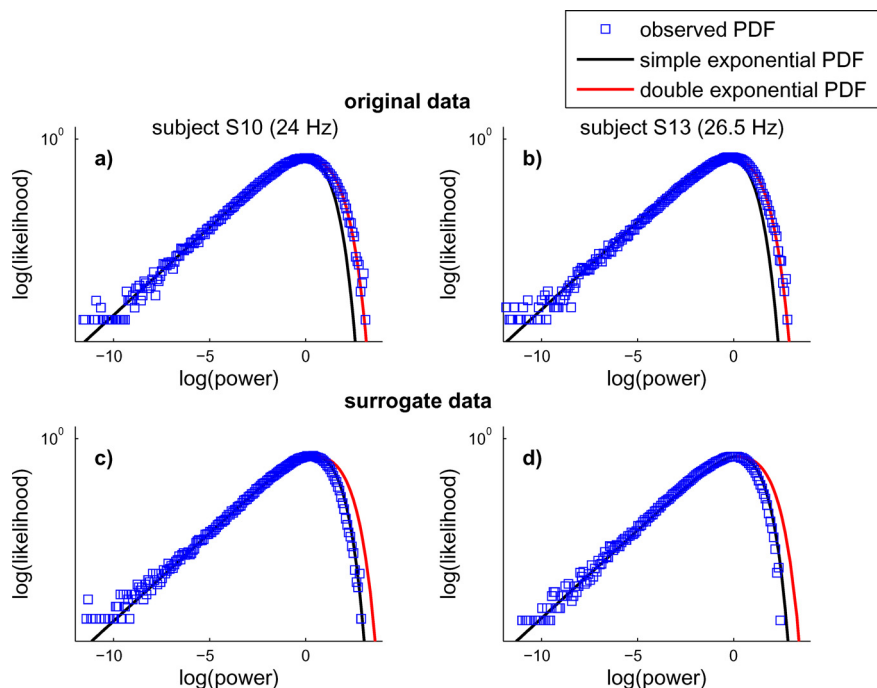


Figure 7. Examples of non-Gaussian unimodal distributions in two subjects. *a* and *b* show observed PDFs. Fitted simple-exponential PDFs (Eq. 1) are depicted as black solid lines, and fitted double-exponential PDFs (Eq. 5) are shown in red above the mean. *c* and *d* show corresponding fits to phase-randomized surrogate data from same subjects.

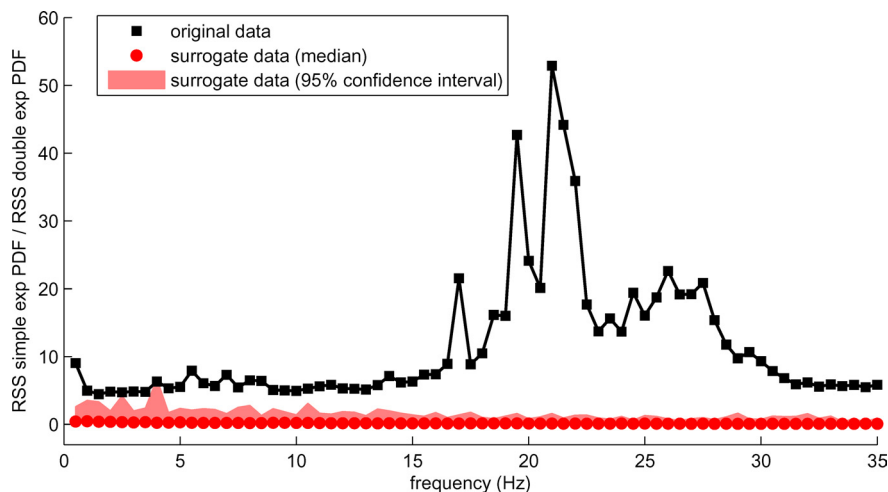


Figure 8. Ratio of the RSS of the single-exponential PDF (Eq. 1) to the double-exponential PDF (Eq. 5) in a single subject (S12). Mean and 95% confidence intervals (± 1.96 SDs) were estimated from surrogate data using nonparametric rank ordering.

exponential and Fisher–Tippett curves, it is possible to define the threshold above which amplitude events are 95% more likely to be drawn from a Fisher–Tippett rather than a simple-exponential PDF. Examples of extremal events in the time–frequency plane, after imposing this threshold in a single subject, are shown in Figure 9, *c* and *d*. For example, the double peak at 22 Hz in *a* is seen to span several adjacent frequencies in *d* at ~ 1328 s. Their appearance in this case, as in all subjects, suggests that, after taking into account the smoothing kernel of the complex Morlet wavelet, extremal events are relatively isolated in both temporal and frequency domains. Detailed inspection of the data failed to reveal more complex phenomena, such as chirped signals tracking diagonally across the time–frequency plane.

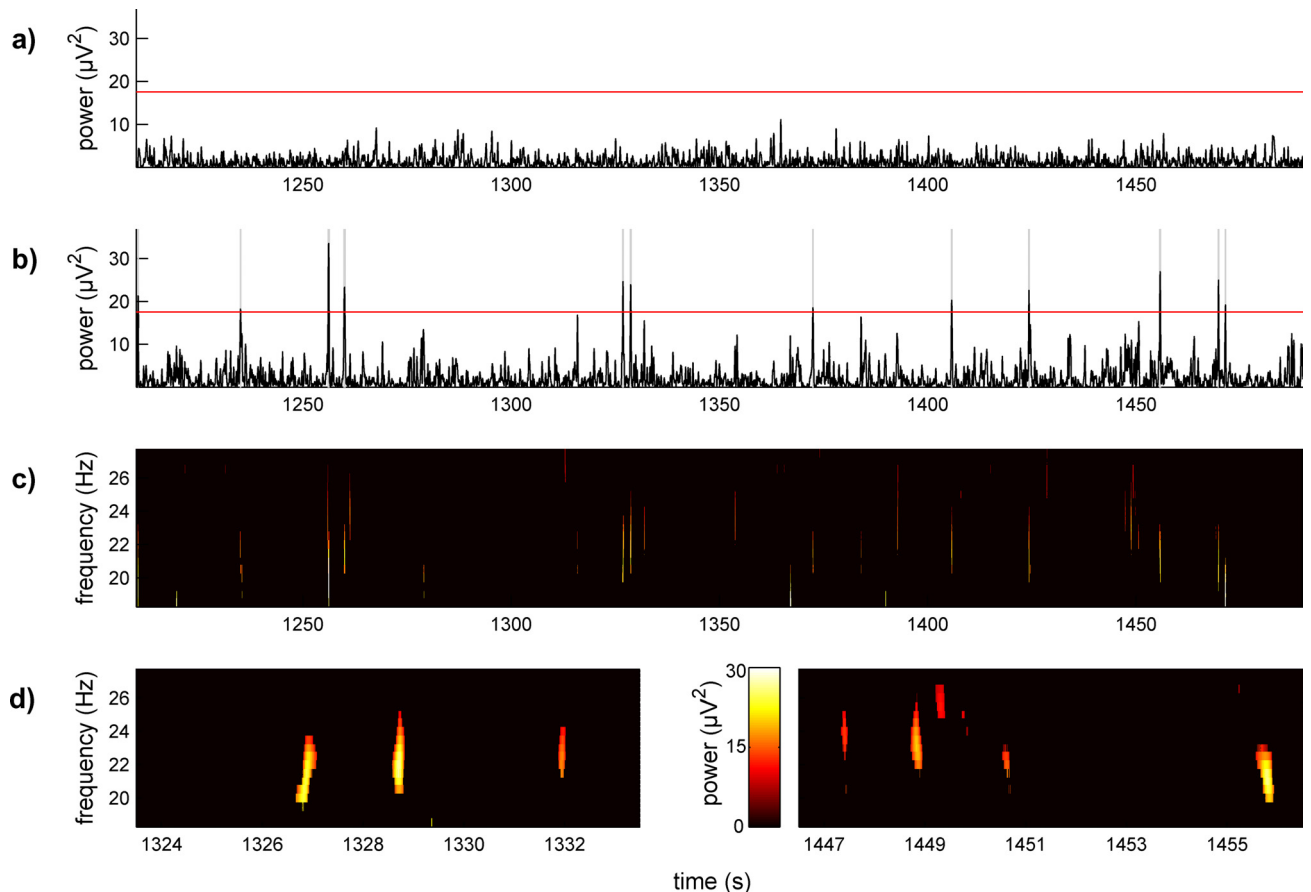


Figure 9. Exemplar extremal events in a single subject (S12). *a, b*, Power time series at 22 Hz in surrogate (*a*) versus real data (*b*), both scaled to have unit mean. Red line shows the 95% relative likelihood for the double-exponential compared with single-exponential fit. *c* and *d* show two complementary views of the time–frequency plane of the empirical data after thresholding at the 95% relative likelihood. Colors depict superthreshold power.

Group-level findings

In this section, we present the intersubject consistency and variability of the observed bimodal and non-Gaussian statistics.

Bimodal distribution

The summary statistics of bimodal PDFs observed at the single-subject level in all 16 subjects are presented in Table 1. These results were compiled by inspecting the BIC differences for each subject and verifying the results against visual inspection of the frequency-specific PDFs to exclude spurious results attributable to outlier observations or poor performance of the fitting routine. They evidence remarkable consistency in 13 of the subjects showing a robust center frequency in the alpha range (mean \pm SD, 10.65 ± 0.72 Hz; range, 9.5–12 Hz).

One subject (S06) had bimodal activity that spanned a wide range of frequencies (19.0–27.0 Hz) in the beta range. In one subject (S05), the bimodal activity was somewhat less distinct and evident just below the defined alpha range. For completeness, the PDFs for these two subjects at the respective (non-alpha) frequencies where they were observed are shown in supplemental Figure S3 (available at www.jneurosci.org as supplemental material). The single case of bimodal beta activity in the beta activity (S06) is notable for exhibiting the same classic bimodal form as typically evident at approximately half that frequency in most other subjects. Only one of the 16 subjects (S07) had no evidence of bimodal activity in the frequencies studied. Interestingly, this subject had consistently high-amplitude alpha activity and hence a high-power peak in

the alpha range, suggesting that the data were acquired during a period of time when the alpha rhythm was consistently in the high-power mode. For thoroughness, the empirical PDFs for these three subjects at 10.5 Hz (the peak frequency for bimodal activity in the grand average) are also shown in supplemental Figure S6 (available at www.jneurosci.org as supplemental material). The lack of any suggestion of a second mode is quite clear in these PDFs. These results hence argue that the existence of bimodal alpha activity is not consistently observed in all subjects.

To further characterize the nature of the bimodal alpha activity in those 13 subjects in whom it was apparent, we calculated the grand average difference between the unimodal and bimodal BIC across this subgroup. Robust evidence for a bimodal fit, when compared with the surrogate data, is evident across the conventional alpha range (8–12 Hz). Above 12 Hz, there exists a broad range of frequencies for which the difference is close to 0, yet nonetheless well above the surrogate distribution, a phenomenon seen at the single-subject level (as in Fig. 5). However, detailed inspection of the PDFs at the single-subject level shows that this finding is not attributable to weak bimodality (such as may be evident as a point of inflection). Rather, it is typically because of the existence of a Fisher–Tippett distribution. When compared with a single-exponential PDF, forcing a second fitted exponential PDF, usually with a very low maximum likelihood, is typically able to decrease the residual error term at least sufficiently to offset the parameter penalty term in Equation 3, hence yielding

near equivalent BIC values for bimodal and unimodal exponential fits between ~12 and 28 Hz.

It is, of course, possible to include the extra three subjects in such a group-level approach. Their statistics have only a minor effect on the mean but naturally increase the variance. However, the present goal is not to make a population-level inference regarding the existence of a single frequency at which a bimodal PDF occurs, because some subjects clearly do not evidence this in a single resting-state recording. Intersubject variability in the existence and character of bimodal distributions is an interesting subject that would arguably require a larger database.

To further investigate the deviation from a unimodal exponential process within the alpha range using a conventional nonparametric statistic, we applied the Kolmogorov–Smirnov (K–S) test on the empirical distributions of both original and surrogate data at each timescale and compared the resulting statistics. The K–S test is based on comparing the empirical distribution function (eCDF) to a theoretical distribution function (tCDF). The resulting statistics reflect the maximum distance between eCDF and tCDF and can be read as a goodness-of-fit measure (Chakravarti et al., 1967; Stephens, 1974). For the tCDF, we used the CDF corresponding to the simple-exponential PDF (Eq. 1), representing the null hypothesis of an underlying Gaussian process. However, it must be noted that the conversion from a K–S statistic to an associated *p* value rests on the assumption that the data are serially independent, which is certainly not the case for raw EEG data or coefficients derived from a continuous wavelet decomposition. We hence used the K–S test simply as a measure of goodness of fit and used the same Fourier-resampling approach for statistical inference.

The grand-averaged frequency-specific K–S statistics are plotted for original and surrogate data in supplemental Figure S7 (available at www.jneurosci.org as supplemental material). Essentially the K–S test replicates the finding observed when using log RSS (compare this figure with Fig. 10). In particular, the K–S statistic diverges for the empirical (black), but not surrogate (red), distributions specifically within the alpha range (8–12 Hz). It is also worth noting that the K–S statistic is lower for the surrogate compared with the empirical data throughout the measured frequency range; this is likely attributable to a variety of ad hoc reasons, such as outlier observations and non-Gaussian (Fisher–Tippett) distributions within the beta range. Also of note is that the K–S statistic gradually decreases with increasing frequency for the surrogate distributions. This is likely attributable to the decreasing length of serial correlations at higher frequencies and hence the increas-

Table 1. Bimodal PDF properties in all subjects

Subject	Acquisition method	Peak BIC difference (unimodal — bimodal)	Frequency for peak BIC difference (Hz)	Approximate range of bimodal distribution (Hz)
S01	A	434	9.5	8.5–11.5
S02	A	389	11.0	8.5–13.5
S03	A	252	12.0	8.5–13.0
S04	A	374	11.5	8.5–12.5
S05	A	181	8.0	7.5–8.0
S06	A	280	23.0	19.0–27.0
S07	A			
S08	C	532	10.5	8.5–12.0
S09	C	421	11.0	8–12.0
S10	C	268	10.0	8.5–14.0
S11	C	745	10.0	8.5–11.5
S12	C	328	10.0	8.5–11.5
S13	C	467	10.5	9.0–12.5
S14	C	351	10.5	8.0–12.0
S15	C	478	11.5	8.5–12.0
S16	C	399	10.5	10.5–11.5

Subject order as per supplemental Figure S2 (available at www.jneurosci.org as supplemental material). Acquisition method: C, combined EEG–fMRI; A, EEG alone. The end points of the bimodal range were defined according to the following criteria: clear point of inflexion in PDF causing increased residual error in single-exponential fit; strongly decreased residual error with second exponential PDF; and adjacent to clearly bimodal PDF (i.e., with 2 distinct maximums). S07 had no discernable bimodal distribution.

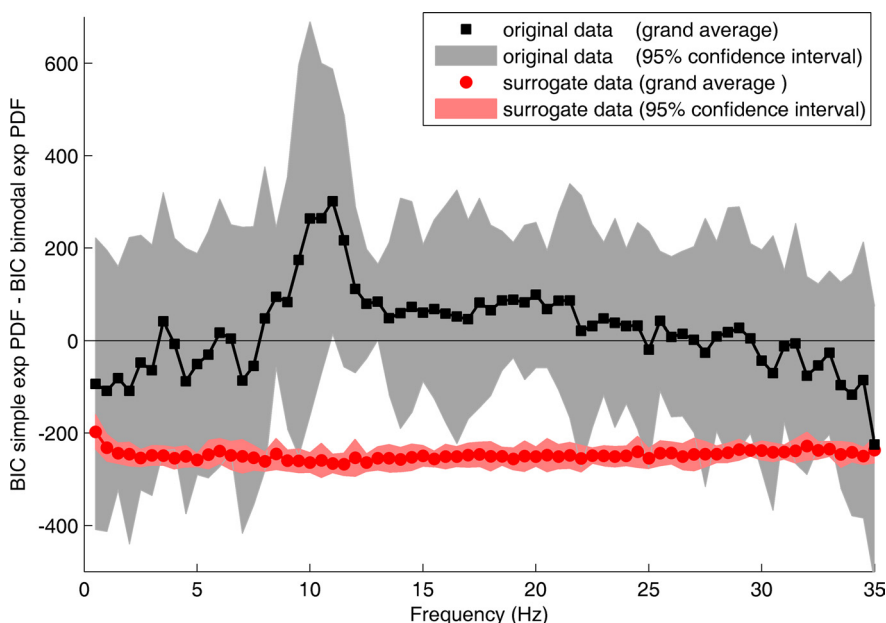


Figure 10. Grand average BIC difference across 13 of the 16 subjects whose data showed distinctly bimodal PDFs, as shown in supplemental Figure S2 (available at www.jneurosci.org as supplemental material). Mean and 95% confidence intervals (± 1.96 SDs) for surrogate data are shown in red.

ing effective degrees of freedom in the data at higher frequencies (to which the K–S test is sensitive).

Non-Gaussian unimodal distributions

Table 2 summarizes the statistics of non-Gaussian PDFs in all 16 subjects. It hence documents the existence of non-Gaussian “fat right-hand tails” within the beta range in all subjects. Evidence of non-Gaussian scaling is also often present at frequencies in the theta and low alpha range (4–8.5 Hz). In contrast to the bimodal PDFs, non-Gaussian distributions showed far greater intersubject variability. It is not uncommon for the relative log evidence to dip close to unity for narrow-frequency intervals separating broader intervals when the relative log evidence is consistently high (>10). Because of this high intersubject frequency evidence,

Table 2. Properties of non-Gaussian distributions in all subjects

Subject	Acquisition method	Maximum relative log evidence	Frequency for maximum ratio (Hz)	Approximate range of Fisher–Tippett distribution (Hz)
S01	A	26	8.0	7.0–8.5
		22	22.5	21.5–25.0
S02	A	50	27.5	26.0–35.0
		46	17.5	17.5
S03	A	35	23.5	20.5–24.0
		31	15.0	14.0–16.5
S04	A	47	34.5	30.5–34.5
S05	A	202	29.0	29.0–29.5
		42	13.5	11.0–13.5
S06	A	76	28.0	27.5–29.5
S07	A	50	8.5	8.0–9.0
		20	19.0	17.0–20.0
S08	C	49	18.0	17.0–20.5
		42	30.5	27.5–32.5
S09	C	82	32.5	22.5–31.5
		26	5.5	2.5–7.5
S10	C	75	24.0	22.5–27.0
		29	4.5	4.5–5.5
S11	C	59	22.0	21.5–25.5
		53	17.0	15.5–20.5
S12	C	53	21.0	18.5–27.5
		138	26.5	26.5–27.0
S13	C	60	17.5	18.0–21.5
		91	30.0	26.5–33.0
S14	C	50	15.0	13.0–15.5
		21	19.5	18.5–22.5
S15	C	105	31.0	30.5–31.5
		83	27.5	22.0–28.0

Relative log evidence denotes the maximum ratio of log residual error (Fisher–Tippett fit to simple exponential fit). The range was defined to include all adjacent frequencies, which scaled closely to the Fisher–Tippett distribution above the mean. In the case in which there are two distinct intervals in the frequency domain showing non-Gaussian PDFs, the interval showing the second highest relative log evidence is given in a second row for each subject.

a second-level subgroup statistic (equivalent to Fig. 10) was not performed.

Notably, the only subject without bimodal activity (S07) had low RSS for a Fisher–Tippett distribution in the alpha range.

At a number of frequencies in all subjects, there existed outlier observations in a small number of power subintervals. The resulting PDFs did not show a clear scaling behavior suggestive of either a single or double exponential, producing quite large errors for both unimodal fits. Because these outlier observations were inevitably above the background likelihood trend, they invariably yielded smaller errors for the Fisher–Tippett distribution, although the relative log ratio was typically <10 . These frequencies were not included in Table 2 as evidence for Fisher–Tippett scaling.

Discussion

After an extensive survey of neuronal oscillations across a hierarchy of scales and the $1/f$ scaling laws of their power spectrum, Buzsáki (2006, p 128) surmised that “rare but extremely large events are inevitable, because at one point $1/f$ dynamics become supersensitive to either external perturbations or its internal processes, responding with very large synchronized events.” However, he then concluded that “such large events never occur.” In addressing this issue, we have shown that the problem does not lie in observing such events but rather in forming a proper representation of the null distribution for their non-occurrence. By estimating PDFs across a spectrum of timescales in spontaneous EEG data and comparing the empirical PDFs with null distributions generated from phase-randomized surrogate data, we observe the robust occurrence of two classes of large-amplitude events in

resting-state EEG. First, in the alpha range (8–12 Hz), 13 of 16 subjects showed a distinctly bimodal distribution with a high-power mode expressed irregularly in discrete time windows. Second, in all subjects and at many temporal scales, the empirical PDFs exhibited a consistent upward bias above the mean from exponential toward power-law scaling. This translates into sporadically occurring high-amplitude events in the time–frequency domain. We therefore argue that the extremal events foreseen by Buzsáki are a robust and routine property of spontaneous cortical activity.

The present approach was informed by the characterization of correlated, non-Gaussian processes in a wide variety of other complex systems, whereby a simple functional form is sought for the probability distribution of a measure of system activity (Bramwell et al., 1998, 2000). As shown by Chapman et al. (2002), it is the presence of non-vanishing higher-order moments that distinguishes systems with nontrivial, strongly correlated fluctuations from those dominated by uncorrelated, diffusive ones. Intriguing applications of this approach have revealed universal power-law distributions in human locomotor activity (Nakamura et al., 2007) and non-Gaussian statistics in heart rate variability during daily activity (Kiyono et al., 2005). The present report also mirrors the

observation of non-Poisson statistics in a variety of human activities ranging from communication to work patterns (Barabási, 2005), as can be seen by comparing Figure 1 of Barabási’s report with Figure 9 of the present one. To our knowledge, this is the first observation of such statistics in spontaneous cortical electrical activity. It accords with an innovative methodological advance by Anemüller et al. (2003) who generalized ICA to the complex domain for cortical EEG data. Their approach enables disambiguation of nonstationary, correlated sources of cortical activity that each exhibit super-Gaussian statistics. The detailed characterization of super-Gaussian empirical PDFs in the present study, particularly in the beta range, arguably provides strong support for this approach over standard ICA models. Another related study is that of Suckling et al. (2008) who estimated the Hölder exponent, a measure of correlated fluctuations, from time series of the resting-state cortical BOLD signal. As with the present approach, they also analyzed timescale-specific activity via a multiscale wavelet decomposition. They hence reported scale invariance of this exponent that was, moreover, sensitive to age and pharmacological perturbation. However, the deviation from a Gaussian process in their data was marginal, possibly attributable to the fact that the Hölder exponent is primarily intended as a singularity analysis and may be less sensitive to non-Gaussian amplitude fluctuations or attributable to the low-pass temporal filtering property of the hemodynamic response function that converts neuronal activity into an fMRI BOLD signal.

An important study of nontrivial temporal correlations in cortical activity is that of Linkenkaer-Hansen et al. (2001) who studied modulations in alpha (8–13 Hz) and beta (15–25 Hz)

power in both EEG and magnetoencephalographic data. They observed power-law scaling of temporal correlations at time-scales ranging from a few seconds to several minutes. These long-range correlations were also present, although attenuated, during sensory input (Linkenkaer-Hansen et al., 2004). It has been proposed that this observation reflects the propagation of “avalanches” of critical neuronal activity in a network close to an instability (Poil et al., 2008), as detailed previously in neuronal cultures by Beggs and Plenz (2003). Long-tailed, possibly scale-free, distributions of synchronous activity in multichannel EEG (Breakspear et al., 2004; Stam and de Bruin, 2004), MEG and fMRI (Kitzbichler et al., 2009) have also been reported across a broad spectrum of timescales, including those of the present study. By revealing the typical presence of two distinct modes of alpha activity and formally estimating parameters for each underlying mode, our study advances this research by providing a more principled (parametric) estimate for switching between high- and low-power modes and hence for estimating the scaling behavior of their respective dwell times. Although we did find long-tailed dwell times for both modes, these appeared to be better captured by stretched exponential forms rather than the power-law scaling functions used by Poil et al. (2008).

A number of methodological points require careful consideration. The present goal of describing power distributions in scalp EEG arose from an objective to characterize the covariance structure of simultaneously acquired EEG and fMRI data. We hence observed the bimodal and non-Gaussian PDFs in the nine simultaneous acquisitions reported in Tables 1 and 2. The classic appearance of eyes closed, resting-state alpha fluctuations underlying the bimodal distributions in these data, and the intersubject variability seen in the non-Gaussian statistics provided *prima facie* support for the neuronal origin of these phenomena. In addition, an exploratory regression of both bimodal and extremal EEG events onto the BOLD data showed a variety of regional signal increases and decreases in cortical regions and did not suggest a classic artifactual effect, such as a ventricular or pericranial pattern. A full characterization of these BOLD correlates is to be the subject of a future report. Additional empirical support arose from the replication of both bimodal and non-Gaussian high-power tails in the additional seven EEG recordings acquired without fMRI in a sound- and light-attenuated room, as also reported in the present study.

It is also critical to note that we estimated the exponential shape parameter λ from a linear regression in linear (power)–log (likelihood) coordinates. However, noting that the estimation of residual errors using a linear regression may yield varied results (Armstrong, 1985; Clauset et al., 2009), we studied the nonlinear functional forms of the PDFs in log (power)–log (likelihood) coordinates. This enables a more accurate visualization of the scaling behavior of the tails of the PDFs and, crucially, quantitative analysis of any systematic scaling bias. We also avoided direct inference on the goodness of fit of the original regression, focusing instead on relative model selection of several candidate models using the BIC approach. The findings were also consistent with those derived from the K–S statistic (Clauset et al., 2009) and validated using surrogate data. Our goal was hence to find the best unbiased model, through a formal model comparison, for these functional forms. Moreover, we also studied the behavior of the candidate models on Gaussian-rendered data to ensure that no systematic methodological bias could lead to spurious non-Gaussian distributions. Finally, we also repeated the groupwise analysis using the K–S test, replicating the finding using BIC.

We chose to analyze the original channel data, rather than

source reconstructed data, because source reconstruction techniques inevitably embody a number of assumptions about the noise structure of the underlying data sources. However, finding non-Gaussian distributions in channel data, which arise from large regions of cortex, arguably has a positive meaning. According to the central limit theorem, the linear superposition of uncorrelated sources is expected to converge toward a Gaussian distribution even if the sources themselves are not Gaussian. Hence, the finding of a non-Gaussian distribution in scalp channel EEG data argues, conversely, that the underlying sources must be correlated and non-Gaussian. Recent theoretical work suggests a non-Gaussian central limit for correlated stochastic processes (Tsallis, 2006). In other words, we submit that both the bimodal and super-Gaussian distributions that we observe in channel data reflect the presence of the same nontrivial distributions in cortical sources and indeed at all scales from the macroscopic down.

The findings we report here have significant implications. The observation of a bimodal PDF in the alpha range supports previous arguments for the presence of a weakly stable, noise-perturbed steady-state attractor in the underlying corticothalamic system (Stam et al., 1999; Valdes et al., 1999) and argues for cortical field models that allow for a nearby bifurcation (Robinson et al., 1997, 2001, 2002; Lopes da Silva et al., 2003). However, the frequent occurrence of regular switching between two distinct modes in our data adds critical details. First, the bifurcation should arguably be of a subcritical Hopf nature to permit distinct jumps between two distinct modes, because a noise-perturbed supercritical bifurcation would be expected to lead to constant (and therefore unimodal) mixing of subthreshold and suprathreshold dynamics. Second suprathreshold oscillations hence occur in healthy states and not only during pathological conditions such as epilepsy, as argued previously (Robinson et al., 1997, 2002; Breakspear et al., 2006). Theoretical work on Hopf bifurcations in time-delayed stochastic systems (Longtin, 1991) provides a potential for understanding the stretched exponential forms that describe the dwelling distributions in our data.

The observation of super-Gaussian (double-exponential) PDFs primarily in the beta range is also of importance, particularly because this frequency range underlies critical cognitive and behavioral processes, such as motor learning and coordination (Boonstra et al., 2007). This finding suggests the presence of nontrivial spatial and temporal correlations among macroscopic neuronal populations. Although these may be expected to occur in small-scale neural systems, as supported by an abundance of theoretical arguments (Amari et al., 2003; Kuhn et al., 2003), current computational models of large-scale neural systems are premised on the diffusion approximation, namely that inputs impinging on individual neurons within a population can be treated as being temporally uncorrelated when computing firing rates (Renart et al., 2003; Deco et al., 2008).

Such models do, however, predict spatial and temporal correlations across multiple scales, including $1/f$ behavior near critical points where instabilities occur (Robinson et al., 1997, 2001, 2002; Robinson, 2003). Extending such models to more explicitly incorporate strong correlations and nontrivial higher-order moments would link them to the rich theoretical developments that currently underpin the study of other complex correlated systems (Zaslavsky, 2002; Tsallis, 2006). These developments explain how systems with both dynamical and stochastic processes, such as a densely connected physiological system under the influence of noise, can intermittently be dominated by the dynamical forces, causing the system to be transiently “trapped” near low-

dimensional, highly synchronous manifolds (Shlesinger et al., 1993).

Such processes arise (although not exclusively so) in systems of coupled oscillators that unfold on an underlying structure with scale-free (i.e., power-law) connectivity (Vasily et al., 2006), a possible basic structural principle of the cortex (Freeman and Breakspear, 2006). Characterizing the temporal statistics of the sporadic expression of extremal events, for example, if they are clustered in time or display formal intermittency (Platt et al., 1993), may help to formally investigate these possibilities. Future work is also required to understand their computational role and metabolic correlates.

References

- Achard S, Salvador R, Whitcher B, Suckling J, Bullmore E (2006) A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs. *J Neurosci* 26:63–72.
- Allen PJ, Polizzi G, Krakow K, Fish DR, Lemieux L (1998) Identification of EEG events in the MR scanner: the problem of pulse artifact and a method for subtraction. *Neuroimage* 8:229–239.
- Allen PJ, Josephs O, Turner R (2000) A method for removing imaging artifact from continuous EEG recorded during functional MRI. *Neuroimage* 12:230–239.
- Amari S, Nakahara H, Wu S, Sakai Y (2003) Synchronous firing and higher-order interactions in neuron pool. *Neural Comput* 15:127–142.
- Anemüller J, Sejnowski TJ, Makeig S (2003) Complex independent component analysis of frequency-domain electroencephalographic data. *Neural Netw* 16:1311–1323.
- Armstrong JS (1985) Long-range forecasting. New York: Wiley.
- Balakrishnan N, Basu AP (1996) The exponential distribution: theory, methods, and applications. New York: Gordon and Breach.
- Barabási AL (2005) The origin of bursts and heavy tails in human dynamics. *Nature* 435:207–211.
- Beggs JM, Plenz D (2003) Neuronal avalanches in neocortical circuits. *J Neurosci* 23:11167–11177.
- Biswal B, Yetkin FZ, Haughton VM, Hyde JS (1995) Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med* 34:537–541.
- Boonstra TW, Daffertshofer A, Breakspear M, Beek PJ (2007) Multivariate time-frequency analysis of electromagnetic brain activity during bimanual motor learning. *Neuroimage* 36:370–377.
- Bramwell ST, Holdsworth PC, Pinton JF (1998) Universality of rare fluctuations in turbulence and critical phenomena. *Nature* 396:552–554.
- Bramwell ST, Christensen K, Fortin J, Holdsworth PC, Jensen HJ, Lise S, Lopez JM, Nicodemi M, Pinton J, Sellitto M (2000) Universal fluctuations in correlated systems. *Phys Rev Lett* 84:3744–3747.
- Breakspear M, Terry JR (2002) Topographic organisation of nonlinear interdependence in multichannel human EEG. *Neuroimage* 16:822–835.
- Breakspear M, Brammer M, Robinson PA (2003) Construction of multivariate surrogates for nonlinear data using the wavelet transform. *Physica D* 182:1–22.
- Breakspear M, Williams LM, Stam CJ (2004) Topographic analysis of phase dynamics in neural systems reveals formation and dissolution of “dynamic cell assemblies”. *J Comput Neurosci* 16:49–68.
- Breakspear M, Roberts JA, Terry JR, Rodrigues S, Mahant N, Robinson PA (2006) A unifying explanation of primary generalized seizures through nonlinear brain modeling and bifurcation analysis. *Cereb Cortex* 16:1296–1313.
- Buzsáki G (2006) Rhythms of the brain. New York: Oxford UP.
- Chakravarti IM, Laha RG, Roy J (1967) Handbook of methods of applied statistics, Vol 1. New York: Wiley.
- Chapman S, Rowlands G, Watkins N (2002) Extremum statistics: a framework for data analysis. *Nonlinear Proc Geophys* 9:409–418.
- Clauset A, Shalizi CR, Newman ME (2009) Power-law distributions in empirical data. *SIAM Rev*. Advance online publication. Retrieved June 16, 2009. arXiv:0706.1062v2.
- Deco G, Jirsa VK, Robinson PA, Breakspear M, Friston K (2008) The dynamic brain: from spiking neurons to neural masses and cortical fields. *PLoS Comput Biol* 4:e1000092.
- Fox MD, Snyder AZ, Zacks JM, Raichle ME (2006) Coherent spontaneous activity accounts for trial-to-trial variability in human evoked brain responses. *Nat Neurosci* 9:23–25.
- Freeman W, Breakspear M (2007) Scale-free neocortical dynamics. *Scholarpedia* 2:1357.
- Grecius MD, Krasnow B, Reiss AL, Menon V (2003) Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proc Natl Acad Sci U S A* 100:253–258.
- Hilgetag CC, Kaiser M (2004) Clustered organization of cortical connectivity. *Neuroinformatics* 2:353–360.
- Kitzbichler MG, Smith ML, Christensen SR, Bullmore E (2009) Broadband criticality of human brain network synchronization. *PLoS Comput Biol* 5:e1000314.
- Kiyono K, Struzik ZR, Aoyagi N, Togo F, Yamamoto Y (2005) Phase transition in a healthy human heart rate. *Phys Rev Lett* 95:058101.
- Kuhn A, Aertsen A, Rotter S (2003) Higher-order statistics of input ensembles and the response of simple model neurons. *Neural Comput* 15:67–101.
- Linkenkaer-Hansen K, Nikouline VV, Palva JM, Ilmoniemi RJ (2001) Long-range temporal correlations and scaling behavior in human brain oscillations. *J Neurosci* 21:1370–1377.
- Linkenkaer-Hansen K, Nikulin VV, Palva JM, Kaila K, Ilmoniemi RJ (2004) Stimulus-induced change in long-range temporal correlations and scaling behaviour of sensorimotor oscillations. *Eur J Neurosci* 19:203–211.
- Longtin A (1991) Noise-induced transitions at a Hopf bifurcation in a first-order delay-differential equation. *Phys Rev A* 44:4801–4813.
- Lopes da Silva F, Blanes W, Kalitzin SN, Parra J, Suffczynski P, Velis DN (2003) Epilepsies as dynamical diseases of brain systems: basic models of the transition between normal and epileptic activity. *Epilepsia* 44 [Suppl 12]:72–83.
- Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. *Nat Neurosci* 9:1432–1438.
- Makeig S, Bell AJ, Jung TP, Sejnowski TJ (1996) Independent component analysis of electroencephalographic data. *Adv Neural Inf Process Syst* 8:145–151.
- Makeig S, Westerfield M, Jung TP, Enghoff S, Townsend J, Courchesne E, Sejnowski TJ (2002) Dynamic brain sources of visual evoked responses. *Science* 295:690–694.
- Maraun D, Kurths J, Holschneider M (2007) Nonstationary Gaussian processes in wavelet domain: synthesis, estimation, and significance testing. *Phys Rev E Stat Nonlin Soft Matter Phys* 75:016707.
- Moosmann M, Ritter P, Krastel I, Brink A, Thees S, Blankenburg F, Taskin B, Obrig H, Villringer A (2003) Correlates of alpha rhythm in functional magnetic resonance imaging and near infrared spectroscopy. *Neuroimage* 20:145–158.
- Nakamura T, Kiyono K, Yoshiuchi K, Nakahara R, Struzik ZR, Yamamoto Y (2007) Universal scaling law in human behavioral organization. *Phys Rev Lett* 99:138103.
- Platt N, Spiegel EA, Tresser C (1993) On-off intermittency: a mechanism for bursting. *Phys Rev Lett* 70:279–282.
- Poil SS, van Ooyen A, Linkenkaer-Hansen K (2008) Avalanche dynamics of human brain oscillations: relation to critical branching processes and temporal correlations. *Hum Brain Mapp* 29:770–777.
- Renart A, Brunel N, Wang XJ (2003) Mean-field theory of irregularly spiking neuronal populations and working memory in recurrent cortical networks. In: *Computational neuroscience: a comprehensive approach* (Feng JF, ed). Boca Raton, FL: CRC.
- Ritter P, Becker R, Graefe C, Villringer A (2007) Evaluating gradient artifact correction of EEG data acquired simultaneously with fMRI. *Magn Reson Imaging* 25:923–932.
- Ritter P, Moosmann M, Villringer A (2009a) Rolandic alpha and beta EEG rhythms’ strengths are inversely related to fMRI-BOLD signal in primary somatosensory and motor cortex. *Hum Brain Mapp* 30:1168–1187.
- Ritter P, Becker R, Freyer F, Villringer A (2009b) EEG quality: the image acquisition artifact. In: *EEG-fMRI: physiological basis, technique, and applications* (Mulert C, Lemieux L, eds). Heidelberg: Springer, in press.
- Robinson PA (2003) Neurophysical theory of coherence and correlations of electroencephalographic and electrocorticographic signals. *J Theor Biol* 222:163–175.
- Robinson PA, Rennie CJ, Wright JJ (1997) Propagation and stability of waves of electrical activity in the cerebral cortex. *Phys Rev E Stat Nonlin Soft Matter Phys* 56:826–840.

- Robinson PA, Rennie CJ, Wright JJ, Bahramali H, Gordon E, Rowe DL (2001) Prediction of electroencephalographic spectra from neurophysiology. *Phys Rev E Stat Nonlin Soft Matter Phys* 63:1–18.
- Robinson PA, Rennie CJ, Rowe DL (2002) Dynamics of large-scale brain activity in normal arousal states and epileptic seizures. *Phys Rev E Stat Nonlin Soft Matter Phys* 65:1–9.
- Shlesinger MF, Zaslavsky GM, Klafter J (1993) Strange kinetics. *Nature* 363:31–37.
- Sporns O, Zwi JD (2004) The small world of the cerebral cortex. *Neuroinformatics* 2:145–162.
- Stam CJ, de Bruin EA (2004) Scale-free dynamics of global functional connectivity in the human brain. *Hum Brain Mapp* 22:97–109.
- Stam CJ, Pijn J, Suffczynski P, Lopes da Silva F (1999) Dynamics of the alpha rhythm: evidence for non-linearity? *Clin Neurophysiol* 110:1801–1813.
- Stephens MA (1974) EDF statistics for goodness of fit and some comparisons. *J Am Stat Assoc* 69:730–737.
- Suckling J, Wink AM, Bernard FA, Barnes A, Bullmore ET (2008) Endogenous multifractal brain dynamics are modulated by age, cholinergic blockade and cognitive performance. *J Neurosci Methods* 174:292–300.
- Theiler J, Eubank S, Longtin A, Galdrikian B, Farmer J (1992) Testing for nonlinearity: the method of surrogate data. *Physica D* 58:77–94.
- Torrence C, Compo GP (1998) A practical guide to wavelet analysis. *Bull Am Meteorolog Soc* 79:61–78.
- Tsallis C (2006) Occupancy of phase space, extensivity of S_q , and q -generalized central limit theorem. *Physica A* 365:7–16.
- Tsallis C, Brigatti E (2004) Nonextensive statistical mechanics: a brief introduction. *Continuum Mech Therm* 16:223–235.
- Valdes PA, Jimenez JC, Riera J, Biscay R, Ozaki T (1999) Nonlinear EEG analysis based on a neural mass model. *Biol Cybern* 81:415–424.
- Vasily E, Tarasov VE, Zaslavsky GM (2006) Fractional dynamics of coupled oscillators with long-range interaction. *Chaos* 16:023110.
- Zaslavsky GM (2002) Chaos, fractional kinetics and anomalous transport. *Phys Rep* 371:461–580.