Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Quantile Detrending Moving Average Algorithm for Multifractal Time Series and Surfaces

Sawon Pratiher

sawon1234@gmail.com

Howrah https://orcid.org/0000-0001-8471-8148

Research Article

Keywords:

Posted Date: November 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11631864/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Quantile detrending moving average algorithm for multifractal time series and surfaces

Sawon Pratiher
Department of Electrical Engineering, Indian Institute of Technology Kharagpur, Kharagpur 721302, West Bengal, India
sawon1234@gmail.com
October 31, 2026

We introduce the multifractal quantile detrending moving average (MFQDMA) algorithm for one-dimensional (1D) time series and two-dimensional (2D) surfaces. MFQDMA replaces the moving average of the multifractal detrending moving average (MFDMA) method by a moving λ-quantile, the local-constant solution of Koenker–Bassett quantile regression, and thereby produces a one-parameter family of generalized Hurst exponents $h(q,\lambda)$ for backward, centered, and forward windows. We derive exact properties of the estimator. For a non-negative measure the quantile residual is exactly the mass of a gliding box of size $\ell = (1-\theta-\lambda)(n-1)$; consequently $h(q,\lambda)$ is asymptotically independent of λ, and the algorithm degenerates on the line $\theta+\lambda=1$, which contains the centered median. For signals, $h(q,\lambda)$ is affine invariant and satisfies $h(q,\lambda;-x) = h(q,1-\lambda;x)$. No box-sum reduction exists in 2D. On benchmarks with analytically known spectra, backward MFQDMA matches MFDMA on clean 1D cascades and fractional Gaussian noises, and lowers the error of the mass exponent $\tau(q)$ of 256×256 multiplicative cascades by 29% at $\lambda = 1/4$. When 2%–10% of a 1D cascade is replaced by outliers, Student-t, or Lévy-stable values, low-λ MFQDMA reduces the $\tau(q)$ error of MFDMA by 22%–46% and that of multifractal detrended fluctuation analysis by a factor of 2–4, whereas $\lambda \geq 3/4$ is harmful. MFQDMA does not remove the spurious multifractality of α-stable noise, does not confer invariance under monotone transforms, and gives only marginal gains for contaminated surfaces. Self-affine surfaces must be analyzed through their mixed increments; otherwise $h(2) \approx 2$ irrespective of the Hurst index. We apply the method to daily returns of 18 financial assets (12 stock indices, three exchange rates, gold, oil, Bitcoin) and 29 Dow Jones stocks. There the multifractality is mainly distributional, and low-λ MFQDMA is the estimator least sensitive to isolated bad ticks (less sensitive than MFDMA for 28 of 29 stocks, and about half as sensitive as MFDFA) and to a few crash days: removing 19 October 1987 changes the S&P 500 spectral width by 0.010, versus 0.093 for MFDMA. MFDMA remains preferable for dense microstructure noise. In 500 electroencephalogram segments and a terrain elevation model, the multifractality is dominated by linear correlations, and the λ-resolved exponents carry class information not contained in $h(q)$. A jointly submitted companion paper extends the method to cross-correlations.

PACS: 05.40.-a, 05.45.Df, 05.45.Tp, 89.75.Da

1 Introduction

Multifractal analysis characterizes the heterogeneity of singular measures, rough signals, and random surfaces through a spectrum of scaling exponents [1–4]. Since the introduction of the wavelet transform modulus maxima [5, 6] and wavelet-leader [7–9] formalisms, a second, fluctuation-based family of estimators has become the most widely used in statistical physics: detrended fluctuation analysis (DFA) [10] and its multifractal generalization MFDFA [11, 12], the detrending moving average (DMA) [13–17], and the multifractal detrending moving average (MFDMA) of Gu and Zhou [18], which they showed to outperform MFDFA on binomial measures and to extend naturally to higher dimensions [18, 19]. These ideas have been carried over to pairs of signals [20–24] and are now standard tools in, e.g., finance [25] and physiology.

All members of the DFA/DMA family estimate the local trend by least squares: a polynomial fit in DFA, an arithmetic mean in DMA. Least-squares trends are optimal for Gaussian fluctuations but have unbounded influence functions [26–29], so a small fraction of aberrant samples can bias the residuals over entire windows. Additive noise, periodic trends, short-term memory [30], heavy tails, and finite-size effects [31, 32] are known to corrupt or spuriously generate multifractality. In statistics, the canonical alternative to least squares is quantile regression [33–35], which replaces the quadratic loss by the asymmetric absolute ("check") loss and estimates conditional quantiles rather than conditional means. Its local-constant version is the moving quantile; the moving median is the classical nonlinear smoother of signal processing [36–38].

In this paper we combine the two lines of work. We define the multifractal quantile detrending moving average (MFQDMA) algorithm in one and two dimensions by replacing the moving average of MFDMA [18] with a moving λ-quantile, $0 < \lambda < 1$. (The quantile level is usually denoted τ in the quantile-regression literature [35]; we write λ to avoid confusion with the multifractal mass exponent $\tau(q)$.) We then ask, and answer quantitatively, three questions that any new multifractal estimator must face: (i) What exactly does the estimator measure, and when does it fail? (ii) Is it accurate on processes with analytically known spectra? (iii) What does it add on real data?

Our main results are as follows. First, we prove that for non-negative measures the quantile residual is *exactly* a gliding-box mass (Proposition 1). This identity explains why MFQDMA estimates the multifractal spectrum of measures, why its exponents do not depend on λ asymptotically, and why it fails catastrophically on the line $\theta+\lambda=1$, which contains the intuitively appealing centered median ($\theta=\lambda=1/2$). Second, we establish affine equivariance and a reflection symmetry $h(q,\lambda;-x) = h(q,1-\lambda;x)$ for signals (Proposition 2), and a telescoping identity that fixes the correct input convention for self-affine surfaces (Proposition 3). Third, extensive benchmarks following the protocols of Refs. [18, 22] show

that MFQDMA is a *robustness* rather than an accuracy improvement in 1D: it ties MFDMA on clean data but is 22%–46% more accurate than MFDMA when 2%–10% of the samples are contaminated. In 2D the quantile trend is genuinely different from the mean trend and yields the most accurate $\tau(q)$ for 256×256 multiplicative cascades. Fourth, we apply MFQDMA to daily returns of 18 financial assets and 29 Dow Jones stocks. We separate the influence of genuine extreme days from that of injected data errors (bad ticks, erroneous returns, and microstructure noise). Low-λ MFQDMA is the least sensitive estimator to isolated bad ticks and to the removal of a few crash days, whereas MFDMA is preferable for dense microstructure noise. A single profile-domain mechanism, transient versus persistent versus dense contamination, explains all three outcomes. Fifth, on electroencephalogram (EEG) recordings [39] and a terrain elevation model we combine MFQDMA with shuffled and phase-randomized surrogates [40, 41] and find that the λ-resolved exponents provide discriminative information beyond the standard $h(q)$. We also report the negative results that a referee would want to know about: MFQDMA does not suppress the spurious multifractality of α-stable noise, it is not invariant under monotone transformations, and the λ dependence of $h(q,\lambda)$ is *not* a measure of amplitude asymmetry.

* Joint submission. This article is the first of two companion papers submitted jointly to Physical Review E. The present paper (Paper I) contains the univariate method in 1D and 2D, its exact properties, the univariate benchmarks, and the univariate applications, including the financial study of Sec. 6.1. The companion paper [42] (Paper II) develops the bivariate extension, multifractal quantile detrending moving-average cross-correlation analysis (MF-X-QDMA), including the treatment of sign-indefinite cross-covariances, pairs of quantile levels (λ_x,λ_y), benchmarks on binomial-measure pairs, bivariate fractional Brownian motions, and two-component ARFIMA processes, and cross-correlations between the financial assets analyzed here. Paper II refers to Secs. 2 and 3 of the present paper for the univariate algorithm and its properties instead of repeating them; no figure or table is shared between the two papers.

The paper is organized as follows. Section 2 defines MFQDMA in 1D and 2D. Section 3 derives its exact properties and the analytical benchmarks. Sections 4 and 5 report the numerical experiments in 1D and 2D, Sec. 6 the applications to financial markets, EEG, and terrain, and Sec. 7 practical guidelines and limitations. Section 8 concludes.

2 The MFQDMA algorithm

2.1 Quantiles as detrending functionals

For a sample $\{u_k\}_{k=1}^n$ the arithmetic mean minimizes the quadratic loss, $\bar u = \arg\min_c \sum_k (u_k-c)^2$. Koenker and Bassett [33] observed that the λ-quantile minimizes the asymmetric absolute loss,

$$\mathscr{Q}_\lambda\{u_k\} \in \arg\min_{c\in\mathbb{R}} \sum_{k=1}^{n} \rho_\lambda(u_k-c), \qquad \rho_\lambda(v) = v\,[\lambda - \mathbb{1}(v<0)], \tag{1}$$

which reduces to the absolute loss (the median) for $\lambda=1/2$. The influence function of $\mathscr{Q}_\lambda$ is bounded, whereas that of the mean is not [27, 29]. Equation (1) can have a flat set of minimizers; we use the continuous, linearly interpolated sample quantile (definition 7 of Hyndman and Fan [43]),

$$\mathscr{Q}_\lambda\{u_k\} = (1-f)\,u_{(\lfloor\lambda(n-1)\rfloor+1)} + f\,u_{(\lfloor\lambda(n-1)\rfloor+2)}, \tag{2}$$

where $u_{(1)}\le\cdots\le u_{(n)}$ are the order statistics and $f=\lambda(n-1)-\lfloor\lambda(n-1)\rfloor$. Equation (2) is equivariant under $u\mapsto au+b$ for $a>0$ and satisfies the reflection identity $\mathscr{Q}_\lambda\{-u_k\} = -\mathscr{Q}_{1-\lambda}\{u_k\}$, two facts used in Sec. 3. A moving λ-quantile is the local-constant quantile regression of a sequence on its index, exactly as a moving average is the local-constant least-squares regression [35].

2.2 One-dimensional MFQDMA

Let $x(t)$, $t=1,\dots,N$, be a time series. We follow the notation of Gu and Zhou [18] and change only step 2.

Step 1. Construct the profile

$$y(i) = \sum_{t=1}^{i} [x(t) - \langle x\rangle], \qquad i=1,\dots,N, \tag{3}$$

where $\langle x\rangle$ is the sample mean. For a non-negative measure we use the *measure convention* $y(i)=\sum_{t\le i}x(t)$ without subtracting $\langle x\rangle$; subtracting the mean destroys the monotonicity of the cumulative measure, and in our tests it raised the error of $\tau(q)$ of a binomial measure from 0.017 to 0.92.

Step 2. For a window size n and a position parameter $\theta\in[0,1]$ define the moving λ-quantile

$$\begin{aligned}\widetilde{y}_\lambda(i) \;=\; \mathscr{Q}_\lambda\big\{y(k):\; & i-\lceil(n-1)(1-\theta)\rceil \le k \\ & \le i+\lfloor(n-1)\theta\rfloor\big\}.\end{aligned} \tag{4}$$

As in MFDMA, $\theta=0$ (backward), $\theta=1/2$ (centered), and $\theta=1$ (forward) place the window behind, around, or ahead of i. Replacing $\mathscr{Q}_\lambda$ by the arithmetic mean recovers MFDMA [18].

Step 3. Detrend the profile,

$$\varepsilon_\lambda(i) = y(i) - \widetilde{y}_\lambda(i), \tag{5}$$

for $\lceil(n-1)(1-\theta)\rceil < i \le N-\lfloor(n-1)\theta\rfloor$.

Step 4. Partition the residual sequence of length $N_\varepsilon = N-n+1$ into $N_n=\lfloor N_\varepsilon/n\rfloor$ disjoint segments of size n and compute

$$F_\nu^2(n,\lambda) = \frac{1}{n}\sum_{i=1}^{n} \varepsilon_\lambda^2\big(i+(\nu-1)n\big), \qquad \nu=1,\dots,N_n. \tag{6}$$

Step 5. The qth-order fluctuation function is

$$F_q(n,\lambda) = \left\{\frac{1}{N_n}\sum_{\nu=1}^{N_n} F_\nu^q(n,\lambda)\right\}^{1/q}, \tag{7}$$

with $\ln F_0(n,\lambda) = N_n^{-1}\sum_\nu \ln F_\nu(n,\lambda)$ for $q=0$.

Step 6. Varying n yields the scaling law

$$F_q(n,\lambda) \sim n^{h(q,\lambda)}, \tag{8}$$

which defines the λ-resolved generalized Hurst exponent $h(q,\lambda)$. The mass exponent, singularity strength, and singularity spectrum follow from the standard relations [1, 11]

$$\tau(q,\lambda) = q\,h(q,\lambda) - D_f, \quad \alpha = \frac{\partial\tau}{\partial q}, \quad f(\alpha) = q\alpha - \tau, \tag{9}$$

with $D_f = 1$ in 1D. Unless λ dependence is the object of study we write $h(q)$ and $\tau(q)$ for a fixed λ.

Figure 1 illustrates steps 2–6 on a fractional Gaussian noise (fGn) with $H = 0.7$. The family of quantile trends brackets the moving average [Fig. 1(a)], the residuals are shifted and reshaped [Fig. 1(b)], and the fluctuation functions are nearly parallel power laws [Fig. 1(c)].

2.3 Two-dimensional MFQDMA

Let $X(i_1,i_2)$, $i_1 = 1,\dots,N_1$, $i_2 = 1,\dots,N_2$, be a field: a measure on a grid, a stationary random field, or the mixed increments of a surface (see Sec. 3.3). Fields in the signal convention are demeaned; measures are not. We use isotropic windows $n_1 = n_2 = n$ and $\theta_1 = \theta_2 = \theta$, as in Ref. [18].

Step 1. For $n \le i_1 \le N_1 - \lfloor (n-1)\theta \rfloor$ and $n \le i_2 \le N_2 - \lfloor (n-1)\theta \rfloor$, compute the box sum

$$Y(i_1,i_2) = \sum_{j_1=i_1-n+1}^{i_1} \sum_{j_2=i_2-n+1}^{i_2} X(j_1,j_2). \tag{10}$$

Step 2. Extract the θ-shifted $n\times n$ submatrix $W(k_1,k_2) = X(i_1 - \lceil (n-1)(1-\theta)\rceil + k_1 - 1,\ i_2 - \lceil (n-1)(1-\theta)\rceil + k_2 - 1)$, $1 \le k_1,k_2 \le n$, and its *within-window* cumulative sums

$$\widetilde{W}(m_1,m_2) = \sum_{d_1=1}^{m_1}\sum_{d_2=1}^{m_2} W(d_1,d_2), \qquad 1 \le m_1,m_2 \le n. \tag{11}$$

The 2D quantile trend is the λ-quantile of these n^2 partial sums,

$$\widetilde{Y}_\lambda(i_1,i_2) = \mathscr{Q}_\lambda\left\{\widetilde{W}(m_1,m_2):\ 1 \le m_1,m_2 \le n\right\}. \tag{12}$$

Replacing $\mathscr{Q}_\lambda$ by the arithmetic mean $n^{-2}\sum_{m_1,m_2}\widetilde{W}(m_1,m_2)$ recovers 2D-MFDMA [Eqs. (15) and (16) of Ref. [18]].

Step 3. The residual matrix is $\varepsilon_\lambda(i_1,i_2) = Y(i_1,i_2) - \widetilde{Y}_\lambda(i_1,i_2)$.

Steps 4–6. Partition ε_λ into $N_n^{(1)}\times N_n^{(2)}$ disjoint $n\times n$ boxes, compute $F^2_{\nu_1\nu_2}(n,\lambda) = n^{-2}\sum \varepsilon_\lambda^2$ in each box, average the qth moments as in Eq. (7), and obtain $h(q,\lambda)$ from $F_q(n,\lambda) \sim n^{h(q,\lambda)}$. Equation (9) applies with $D_f = 2$.

Two implementation remarks matter. (a) The cumulative sums in Eq. (11) must be computed *within* each window. In 1D, a global cumulative sum differs from the within-window one by a per-window constant that cancels in Eq. (5); in 2D the inclusion–exclusion correction varies across the window and does not cancel. Using a global 2D cumulative sum raised the $\tau(q)$ error of a 256×256 cascade from 0.14 to 1.42 in our tests. (b) The position parameter enters only through the shift of the trend window relative to the box sum of Eq. (10). The mean trend is separable, $n^{-2}\sum_m \widetilde{W}(m) = n^{-2}\sum_d W(d)(n-d_1+1)(n-d_2+1)$, and is evaluated with two 1D convolutions. The quantile trend is not separable, so all n^2 partial sums per window are materialized from an integral image (Appendix A).

3 Exact properties and analytical benchmarks

3.1 Measures: the gliding-box representation

Proposition 1 (Gliding-box representation)**.** *Let $x(t) \ge 0$ and let y be the measure-convention profile, extended to non-integer arguments by linear interpolation. Then, for every admissible i in Eq. (5),*

$$\begin{aligned} \varepsilon_\lambda(i) &= y(i) - y(i-\ell), \\ \ell &\equiv \ell(n,\theta,\lambda) = \lceil (n-1)(1-\theta)\rceil - \lambda(n-1). \end{aligned} \tag{13}$$

For $\ell > 0$, $\varepsilon_\lambda(i) = \mu\big((i-\ell,i]\big)$ is the mass of a gliding box of length ℓ ending at i. For $\ell < 0$ it is minus the mass of the box $(i,i+|\ell|]$.

Proof. Because $x \ge 0$, y is nondecreasing, so the order statistics of the window in Eq. (4) are its elements in index order: $y_{(k+1)} = y(i - \lceil (n-1)(1-\theta)\rceil + k)$, $k = 0,\dots,n-1$. Inserting this into Eq. (2) gives $\widetilde{y}_\lambda(i) = y\big(i - \lceil (n-1)(1-\theta)\rceil + \lambda(n-1)\big)$, with linear interpolation between neighbors. Equation (5) then gives Eq. (13). □

We verified Eq. (13) numerically for binomial measures and several (n,θ,λ); the maximal deviation is 10^{-15} (machine precision). For comparison, the MFDMA residual of a measure is a smoothed box sum with triangular weights, $y(i) - \langle y\rangle_{\rm win} = \sum_k w_k\, x(i-k)$. MFQDMA replaces this kernel by a sharp box, which is why the two methods behave similarly on 1D measures.

Corollary 1 (Degeneracy line)**.** *If $|\ell| < 1$, the residual $\varepsilon_\lambda(i) = \ell\, x(i)$ (for $\ell \ge 0$, and analogously for $\ell < 0$) does not grow with n. Hence $F_q(n,\lambda)$ does not scale and the estimate of $h(q)$ is spurious. Since $\ell = (1-\theta-\lambda)(n-1) + O(1)$, this happens on the line $\theta + \lambda = 1$. In particular, the centered median $\theta = \lambda = 1/2$ gives $\ell = 0$ for odd n (so $\varepsilon_\lambda \equiv 0$) and $\ell = 1/2$ for even n.*

Corollary 2 (Asymptotic λ independence)**.** *Away from the degeneracy line, $|\ell| = c\,n + O(1)$ with $c = |1-\theta-\lambda| > 0$. If the box masses of size s have partition function $\sum \mu(s)^q \sim (s/N)^{\tau(q)}$ with N/s boxes, then $F_\nu(n,\lambda)$ is the root-mean-square gliding-box mass at scale cn inside segment ν and*

$$F_q^q(n,\lambda) \approx \frac{n}{N}\sum_\nu \mu_\nu(cn)^q \propto (cn)^{1+\tau(q)}. \tag{14}$$

Therefore $h(q,\lambda) = [\tau(q)+1]/q$ independently of λ and θ. Changing λ only rescales the effective box size, which shifts $\ln F_q$ by a constant and leaves the slope unchanged.

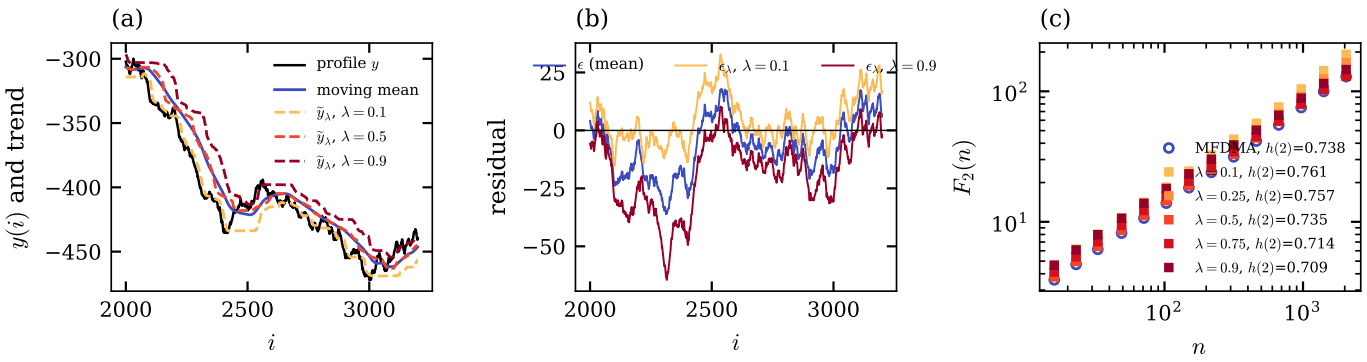


Figure 1: The 1D MFQDMA algorithm on an fGn with $H = 0.7$ ($N = 2^{14}$, $\theta = 0$, $n = 128$). (a) Profile $y(i)$ with the moving mean of MFDMA and the moving quantiles $\widetilde{y}_\lambda$ for $\lambda = 0.1, 0.5, 0.9$. (b) Residuals $\varepsilon(i)$ of MFDMA and $\varepsilon_\lambda(i)$ for $\lambda = 0.1$ and 0.9. (c) $F_2(n,\lambda)$ for MFDMA and five quantile levels; the slopes $h(2,\lambda)$ are given in the legend.

Corollary 2 has an important negative consequence. For measures, any λ dependence of $h(q,\lambda)$ is a finite-size effect and cannot encode, e.g., an asymmetry of the multipliers. We confirm this numerically in Sec. 4.6. Corollary 1 explains the numerical observation that the centered MFQDMA fails on measures by two orders of magnitude (Sec. 4.3), whereas centered MFDMA merely loses accuracy [18].

3.2 Signals: equivariance and reflection

Proposition 2 (Affine equivariance and reflection)**.** *In the signal convention [Eq. (3)], let $x'(t) = a\,x(t) + b$ with $a \neq 0$. Then, in 1D and in 2D,*

$$h(q,\lambda;x') = h(q,\lambda;x), \qquad a > 0, \tag{15}$$
$$h(q,\lambda;x') = h(q,1-\lambda;x), \qquad a < 0. \tag{16}$$

Proof. Demeaning removes b, so $y' = a\,y$ (or $\widetilde{W}' = a\widetilde{W}$ in 2D). By the equivariance and reflection identities of Eq. (2), $\varepsilon'_\lambda = a\,\varepsilon_\lambda$ for $a > 0$ and $\varepsilon'_\lambda = a\,\varepsilon_{1-\lambda}$ for $a < 0$. Hence $F'_q(n,\lambda) = |a|F_q(n,\lambda)$ or $|a|F_q(n,1-\lambda)$, and the slopes are unchanged or exchanged. □

Numerically, Eqs. (15) and (16) hold to 10^{-15} in 1D and to 5×10^{-12} in 2D. Proposition 2 guarantees that amplifier gains and offsets cannot affect MFQDMA exponents. It also implies that for a sign-symmetric process ($x \overset{d}{=} -x$), $h(q,\lambda)$ and $h(q,1-\lambda)$ are identically distributed. An antisymmetric index $h(q,\lambda) - h(q,1-\lambda)$ therefore vanishes on average for all Gaussian processes and cannot detect anything else in them.

3.3 Surfaces: the input convention

Proposition 3 (Telescoping of mixed increments)**.** *Let $G(i_1,i_2) = X(i_1,i_2) - X(i_1-1,i_2) - X(i_1,i_2-1) + X(i_1-1,i_2-1)$ be the mixed increment field of a surface X. The box sum of Eq. (10) applied to G equals the rectangular increment of X over the corners of the box, and the same holds for every within-window partial sum $\widetilde{W}$. If X is self-similar with Hurst index H, e.g., an isotropic fractional Brownian surface [44], all these quantities scale as n^H in distribution. Then 2D-MF(Q)DMA applied to G yields $h(q) = H$ and $\tau(q) = qH - 2$.*

The proof is immediate from the telescoping of double sums and the self-similarity $X(c\,\mathbf{r}) \overset{d}{=} c^H X(\mathbf{r})$. Applied to X itself, by contrast, the box sums are dominated by n^2 times the local level of the surface, and $h(q)$ saturates near $D_f = 2$ for every H (Sec. 5.2). An uncorrelated field gives $h = D_f/2 = 1$ in this convention, the 2D analog of $h = 1/2$ for white noise in 1D. We therefore use three conventions: a measure is analyzed directly without demeaning, a stationary field directly after demeaning, and a self-affine surface through its mixed increments G.

3.4 Analytical benchmarks

* 1D binomial measure. The p-model on $N = 2^k$ points [2, 18] has

$$\tau_{\rm th}(q) = -\log_2\left[p^q + (1-p)^q\right], \tag{17}$$
$$h_{\rm th}(q) = \frac{\tau_{\rm th}(q)+1}{q}, \qquad q \neq 0,$$
$$h_{\rm th}(0) = \tau'_{\rm th}(0) = -\frac{\ln p + \ln(1-p)}{2\ln 2}. \tag{18}$$

The value at $q = 0$ follows from L'Hôpital's rule applied to $\tau_{\rm th}(q) + 1 \to 0$. The entropy-weighted expression $-[p\ln p + (1-p)\ln(1-p)]/\ln 2$ that is sometimes used is incorrect here.

* 2D four-multiplier cascade. At each of k iterations a square is divided into quadrants that receive fractions $p_1,\ldots,p_4$ ($\sum_i p_i = 1$) of its mass [18]. Then

$$\tau_{\rm th}(q) = -\log_2 \sum_{i=1}^{4} p_i^q, \quad h_{\rm th}(q) = \frac{\tau_{\rm th}(q)+2}{q}, \tag{19}$$
$$\alpha_{\rm th}(q) = -\frac{\sum_i p_i^q \ln p_i}{\ln 2 \sum_i p_i^q}, \qquad f_{\rm th} = q\alpha_{\rm th} - \tau_{\rm th}, \tag{20}$$
$$h_{\rm th}(0) = \alpha_{\rm th}(0) = -\frac{1}{4\ln 2}\sum_{i=1}^{4} \ln p_i. \tag{21}$$

The identities $\tau_{\rm th}(0) = -2$, $\tau_{\rm th}(1) = 0$, and $\tau_{\rm th}(q) = 2(q-1)$ for $p_i \equiv 1/4$ serve as consistency checks of any implementation.

* Monofractals. An fGn with Hurst index H has $h(q) = H$. A fractional Brownian surface analyzed through G has $h(q) = H$ and $\tau(q) = qH - 2$ (Proposition 3).

4 Numerical experiments in one dimension

4.1 Protocol

The benchmark corpus is summarized in Table 1. All exact Gaussian processes were generated by circulant embedding [45, 46] with Hermitian-symmetric spectra, and α-stable variables by the Chambers–Mallows–Stuck method [47]. Unless stated otherwise, we use $q \in [-4,4]$ in steps of 0.5. MFDMA and MFQDMA use 14 logarithmically spaced window sizes $16 \le n \le 2048$, and MFDFA (linear detrending) uses $64 \le s \le 8192$, the ranges recommended in Ref. [18]. Exponents are least-squares slopes of $\ln F_q$ versus $\ln n$, and α is obtained by central differences. Accuracy is the root-mean-square error over the 17 moments,

$$\mathrm{RMSE}_\tau = \left\{ \frac{1}{17} \sum_q \left[\tau(q) - \tau_{\mathrm{th}}(q) \right]^2 \right\}^{1/2}, \qquad (22)$$

and analogously RMSE_h for $h(q)$ versus H for monofractals.

4.2 Binomial measures

Figure 2 shows the analysis of a binomial measure with $p = 0.3$ and $N = 2^{16}$. The fluctuation functions of backward MFQDMA are excellent power laws over more than two decades [Fig. 2(a)], and the estimated $\tau(q)$ and $f(\alpha)$ agree with Eqs. (17) and (18) [Figs. 2(b) and 2(d)]. The deviations $\Delta\tau(q) = \tau(q) - \tau_{\mathrm{th}}(q)$ of all backward MFQDMA variants stay within about 0.07 of zero. This is comparable to backward MFDMA and much smaller than for MFDFA, whose error grows to 0.2 at $q = 4$ [Fig. 2(c)].

Table 2 summarizes eight cascades. We reproduce the ranking of Ref. [18]: backward MFDMA ($\mathrm{RMSE}_\tau = 0.025$) beats forward MFDMA, MFDFA, and centered MFDMA. All backward MFQDMA variants come within 0.004–0.007 of backward MFDMA, which is less than one standard deviation across realizations. On clean 1D measures MFQDMA is therefore as accurate as the best existing method, but not more accurate. Figure 4(a) illustrates Corollary 2: the five curves $h(q,\lambda)$ coincide with each other and with the theory.

4.3 Dependence on the window position and the degeneracy line

Figure 3 tests Proposition 1 and Corollary 1. Panel (a) shows the geometry: on a monotone profile the λ-quantile of the window is the profile value at lag ℓ, so the residual is the mass of the box $(i-\ell, i]$. Panel (b) shows the residual standard deviation as a function of θ for $n = 65$. For each λ it collapses to zero exactly at $\theta = 1-\lambda$, while the MFDMA residual only has a shallow minimum at $\theta = 1/2$. Panel (c) shows that RMSE_τ of MFQDMA peaks at $\theta \approx 1-\lambda$ for every λ, reaching 2.3 for the centered median. Backward MFQDMA ($\theta = 0$) is safe for every $\lambda < 1$, so we recommend $\theta = 0$.

4.4 Monofractal signals

Table 3 and Figs. 4(b) and 4(c) summarize the fGn benchmark. All methods recover $h(2)$ within 0.06 of H. The centered MFDMA is the most accurate ($\mathrm{RMSE}_h = 0.006$), in agreement with Ref. [22]. Backward MFQDMA (0.021–0.034) is comparable to backward MFDMA (0.028) and to MFDFA (0.023). For signals the profile is not monotone, so Corollary 1 does not apply, and the centered median is merely inaccurate (0.104) rather than degenerate. The spurious width $\max_q h - \min_q h$ of MFQDMA grows with H and λ [Fig. 4(c)] and stays below 0.10; for $H \ge 0.8$, $\lambda = 0.1$ gives the narrowest spectra among the backward DMA-type methods.

4.5 Robustness to contamination

We contaminate the binomial measure ($p = 0.3$, $N = 2^{16}$) by replacing a randomly chosen fraction $\varepsilon \in \{1,2,5,10\}\%$ of its values with (i) outliers $\pm 20\sigma$ with random sign, (ii) Student-t variables with two degrees of freedom scaled by 4σ, or (iii) symmetric α-stable variables with $\alpha_L = 1.2$ scaled by σ. Here σ is the standard deviation of the clean measure. Each configuration is repeated five times. The analysis uses the measure convention and $\tau_{\mathrm{th}}(q)$ of the clean cascade as reference. The results are given in Table 4 and Fig. 5.

Three conclusions follow. (i) For $\varepsilon \ge 2\%$, MFQDMA with $\lambda = 0.1$ lowers the error of backward MFDMA by 23%–37% for outliers, 25%–46% for Student-t, and 22%–32% for Lévy contamination. It lowers the error of MFDFA by a factor of 1.9–4.2. (ii) At $\varepsilon = 1\%$ of $\pm 20\sigma$ outliers, low-λ MFQDMA and MFDMA are tied (0.105–0.111 versus 0.110). (iii) Upper quantiles are harmful. With $\lambda = 0.9$ the error at $\varepsilon = 1\%$ is five times that of MFDMA, because the upper quantile follows the envelope inflated by positive outliers. The paired ratios in Fig. 5(d) show that the ranking $\lambda = 0.1 \lesssim 0.25 < 1 < 0.9$ holds for every contamination type up to $\varepsilon = 5\%$; at 10% even $\lambda = 0.9$ is slightly better than MFDMA for outliers.

The improvement is real but moderate, and a clarification is necessary. Robust-statistics arguments based on the breakdown point of the median do not transfer to profile-based estimators. A single outlier in x produces a step in y that persists for all later indices and enters n consecutive windows. No order statistic of the profile can remove it. The quantile trend reduces the influence of the step on the local trend but does not eliminate it. MFQDMA therefore does not have a 50% breakdown point, and claims of this kind for median-based detrending should be treated with caution. The argument does, however, single out the case in which order statistics help most. A *transient* profile outlier, which displaces a single value of y, changes a moving λ-quantile by at most one rank but shifts the moving mean in n windows. In financial data this is exactly the signature of an isolated misprinted price (Sec. 6.1).

Table 1: Benchmark corpus with analytically known references. N is the series length (1D) or the side length (2D).

Process	Parameters	N, realizations
Binomial measure	$p \in \{0.25, 0.30, 0.35, 0.40\}$	$2^{14}, 2^{16}$; 8
fGn	$H \in \{0.2, 0.3, 0.5, 0.7, 0.8, 0.9\}$	2^{16}; 18
Contaminated binomial	$p = 0.3$, 3 types, $\varepsilon \leq 10\%$	2^{16}; 61
α-stable noise	$\alpha_L \in \{1.2, 1.5, 1.8, 2\}$	2^{15}; 20
Monotone maps	log-normal cascade, 6 maps	2^{15}; 18
Random-orientation cascade	4 multiplier pairs	2^{15}; 20
2D cascade	3 multiplier sets	256, 512; 6
Contaminated 2D cascade	2 types, $\varepsilon \leq 10\%$	256; 17
Fractional Brownian surface	$H \in \{0.2, \ldots, 0.9\}$	256; 19

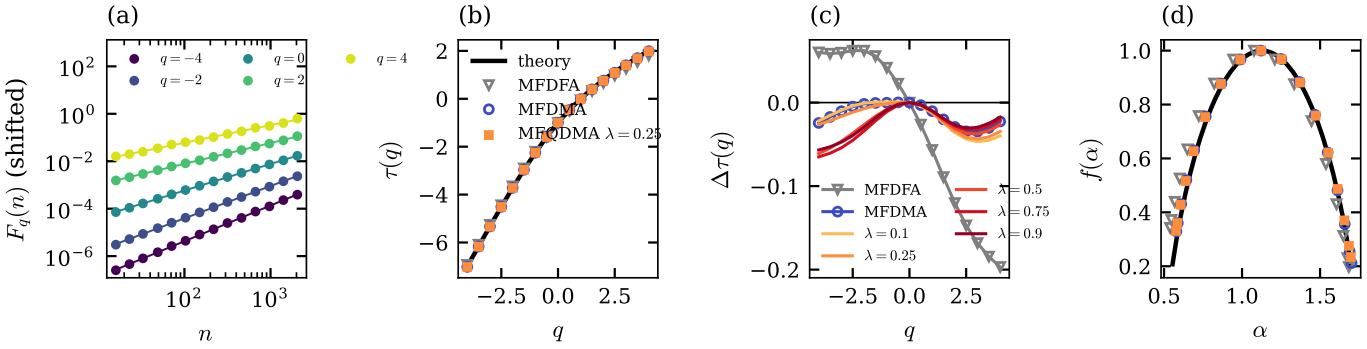


Figure 2: Binomial measure, $p = 0.3$, $N = 2^{16}$. (a) $F_q(n)$ of backward MFQDMA ($\lambda = 0.25$) for $q = -4, -2, 0, 2, 4$, shifted vertically for clarity; lines are least-squares fits. (b) $\tau(q)$ from MFDFA, backward MFDMA, and backward MFQDMA versus Eq. (17). (c) $\Delta\tau(q)$ for all methods and five quantile levels. (d) Singularity spectra versus theory.

Table 2: RMSE_τ (mean $\pm$ s.d.) over the eight binomial measures of Table 1. The centered MFQDMA lies on the degeneracy line of Corollary 1.

Method	θ	RMSE_τ
MFDFA	—	0.155 ± 0.054
MFDMA	0	$\mathbf{0.025 \pm 0.016}$
MFDMA	1/2	0.142 ± 0.034
MFDMA	1	0.041 ± 0.026
MFQDMA, $\lambda = 0.10$	0	0.029 ± 0.020
MFQDMA, $\lambda = 0.25$	0	0.029 ± 0.018
MFQDMA, $\lambda = 0.50$	0	0.032 ± 0.015
MFQDMA, $\lambda = 0.75$	0	0.031 ± 0.012
MFQDMA, $\lambda = 0.90$	0	0.029 ± 0.009
MFQDMA, $\lambda = 0.50$	1	0.050 ± 0.030
MFQDMA, $\lambda = 0.50$	1/2	2.278 ± 0.116

4.6 Heavy tails, monotone transforms, and amplitude asymmetry

* α-stable noise. Independent α-stable variables are monofractal in their correlations but generate spurious finite-size multifractality through their heavy tails [11, 32]. Figure 6(a) shows that the spurious width of MFQDMA, between 0.73 and 0.77 for $\alpha_L = 1.2$ and $N = 2^{15}$, is the same as that of MFDMA (0.73) and MFDFA (0.70) within 0.07. A quantile trend cannot help here, because the heavy tails enter through the qth moments of the residuals in Eq. (7), not through the trend.

* Monotone transforms. Rank-based statistics are invariant under strictly increasing maps, and one might expect MFQDMA to inherit this property. It does not, because the quantile acts on the profile, which is a sum of transformed values. For a log-normal cascade ($\sigma = 0.4$), the mean absolute change of $h(q)$ under the maps $\sqrt{\cdot}$, $\sqrt[3]{\cdot}$, log, Gaussianization, and $(\cdot)^3$ is as large for MFQDMA as for MFDMA [Fig. 6(b)]; MFDFA is the least sensitive (1.35 versus 1.60–1.74 under cubing).

Table 3: RMSE_h (mean $\pm$ s.d.) of $h(q)$ versus H over 18 fGn series ($N = 2^{16}$).

Method	θ	RMSE_h
MFDFA	—	0.023 ± 0.022
MFDMA	0	0.028 ± 0.020
MFDMA	1/2	$\mathbf{0.006 \pm 0.003}$
MFDMA	1	0.028 ± 0.021
MFQDMA, $\lambda = 0.10$	0	0.028 ± 0.017
MFQDMA, $\lambda = 0.25$	0	0.034 ± 0.021
MFQDMA, $\lambda = 0.50$	0	0.031 ± 0.021
MFQDMA, $\lambda = 0.75$	0	0.025 ± 0.022
MFQDMA, $\lambda = 0.90$	0	0.021 ± 0.014
MFQDMA, $\lambda = 0.50$	1/2	0.104 ± 0.126

* Amplitude asymmetry. Random-orientation binomial cascades, in which each split assigns the multipliers $(p_a, 1 - p_a)$ or $(p_b, 1 - p_b)$ with probability $1/2$, give a λ-spread $\max_\lambda h(q, \lambda) - \min_\lambda h(q, \lambda)$ (averaged over q) of 0.0085, 0.0116, and 0.0225 for $(p_a, p_b) = (0.35, 0.65)$, $(0.30, 0.50)$, and $(0.25, 0.75)$, and zero for the uniform case $(0.5, 0.5)$. The spread is of order 10^{-2} and increases with the strength of multifractality, not with any asymmetry of the multipliers. This

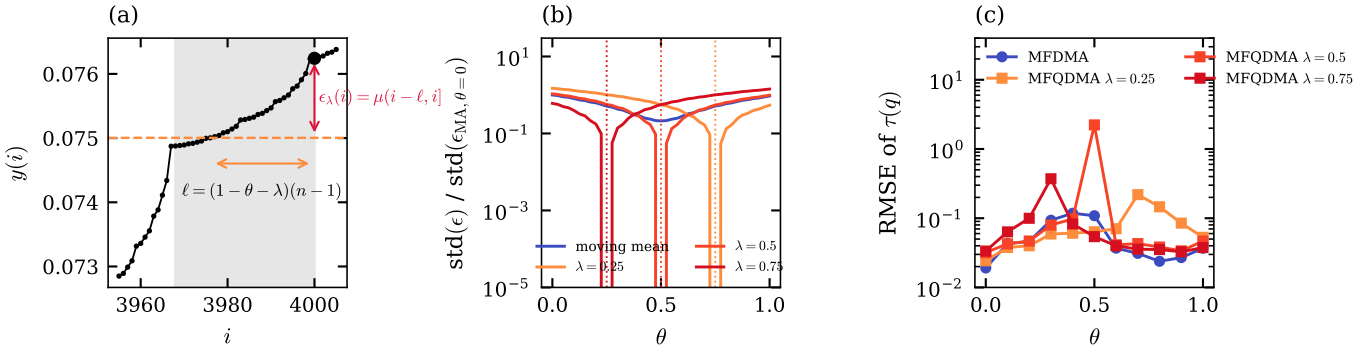


Figure 3: Proposition 1 and the degeneracy line. (a) A monotone profile of a binomial measure ($p = 0.3$) with a backward window (shaded, $n = 33$). The $\lambda = 0.25$ quantile (dashed) equals $y(i-\ell)$, so $\varepsilon_\lambda(i)$ is the box mass $\mu(i-\ell, i]$. (b) Standard deviation of the residual versus θ ($n = 65$), normalized by that of backward MFDMA. Dotted lines mark $\theta = 1-\lambda$. (c) RMSE$_\tau$ versus θ for $N = 2^{16}$.

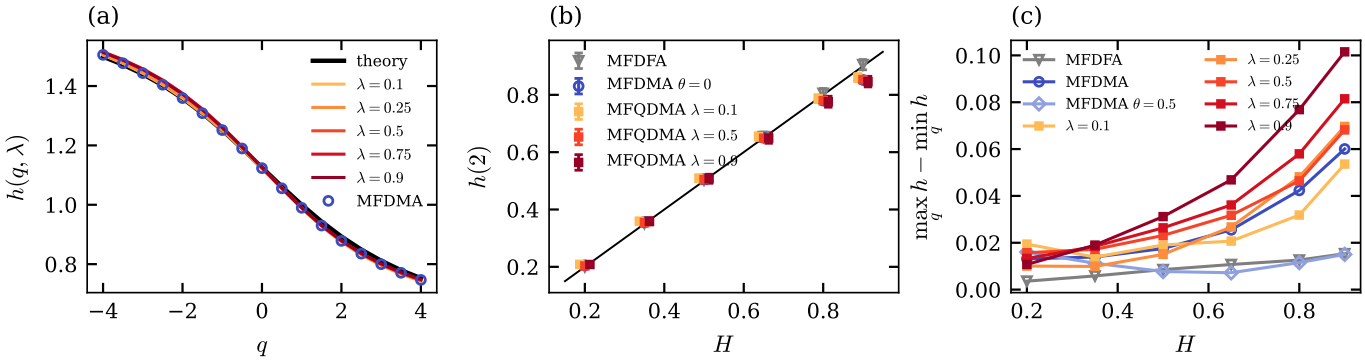


Figure 4: (a) Corollary 2: $h(q,\lambda)$ of a binomial measure ($p = 0.3$, $N = 2^{16}$) for five quantile levels, compared with MFDMA and with Eq. (17). (b) $h(2)$ versus H for fGn (mean $\pm$ s.d. of three realizations, $N = 2^{16}$); the line is $h = H$. (c) Spurious width $\max_q h(q) - \min_q h(q)$ for fGn.

is what Corollary 2 predicts for a finite-size effect. The λ-dependence of $h(q,\lambda)$ for measures is therefore not an asymmetry index.

5 Numerical experiments in two dimensions

5.1 Multiplicative cascades

We generate four-multiplier cascades with $(p_1, p_2, p_3, p_4) = (0.05, 0.15, 0.25, 0.55)$ (set A), $(0.10, 0.20, 0.30, 0.40)$ (set B, the parameters of Ref. [18]), and $(0.02, 0.18, 0.30, 0.50)$ (set C), on 256^2 and 512^2 grids. We use eight logarithmically spaced window sizes $6 \le n \le N/8$ and the measure convention. Figure 7 shows set A on a 512^2 grid. The fluctuation functions are power laws, and the singularity spectra of 2D-MFDMA and 2D-MFQDMA agree with Eqs. (20) and (21) over the full range $0.9 \lesssim \alpha \lesssim 4.3$.

Unlike in 1D, no gliding-box reduction exists in 2D. The partial sums $\widetilde{W}(m_1, m_2)$ of a measure increase in each index separately, but their λ-quantile over the $n \times n$ index square is not a fixed lattice element: which element it is depends on the local geometry of the measure. The 2D quantile trend is therefore a genuinely different detrending functional. Table 5 and Fig. 8(a) show that it is also more accurate. On 256^2 grids, backward 2D-MFQDMA with $\lambda = 0.25$ has the lowest error for all three multiplier sets, with improvements of 22% (A), 34% (B), and 36% (C) over backward 2D-MFDMA, or 29% averaged over the three sets. On 512^2 grids all errors are smaller and the differences shrink. $\lambda = 0.5$ is marginally better than MFDMA for all three sets (by 0.3%–22%), whereas $\lambda = 0.25$ is better only for set B, and $\lambda = 0.75$ is marginally best for set C. Averaged over both sizes and all sets, RMSE$_\tau$ = 0.135 ($\lambda = 0.25$), 0.147 ($\lambda = 0.5$), and 0.167 (MFDMA). For every trend functional the backward window is the most accurate [Fig. 8(a)], in agreement with Ref. [18].

5.2 Monofractal surfaces and the input convention

Figure 9 and Table 6 test Proposition 3 on isotropic fractional Brownian surfaces (256^2, spectral synthesis). Applied to the surface X (one realization for each $H \in \{0.3, 0.5, 0.7, 0.9\}$), every method gives $h(2)$ between 1.93 and 2.16 for every H (dashed lines), with fit qualities $R^2 > 0.997$. These fits look excellent but measure only the n^2 growth of the box sums. Applied to the mixed increments G (three realizations for each

Table 4: RMSE_τ of the contaminated binomial measure ($p = 0.3$, $N = 2^{16}$, mean of five replicates). The clean values ($\varepsilon = 0$) are MFDFA 0.101, MFDMA 0.019, and MFQDMA 0.025, 0.024, 0.032, 0.034, 0.030 for $\lambda = 0.1, 0.25, 0.5, 0.75, 0.9$. Bold: best per column.

	Outliers $\pm 20\sigma$				Student-t_2				Lévy, $\alpha_L = 1.2$			
Method	1%	2%	5%	10%	1%	2%	5%	10%	1%	2%	5%	10%
MFDFA	0.237	0.415	0.938	1.570	0.302	0.476	0.862	1.219	0.424	0.581	0.858	1.102
MFDMA ($\theta = 0$)	0.110	0.166	0.418	0.810	0.122	0.210	0.460	0.728	0.293	0.395	0.618	0.764
MFQDMA $\lambda = 0.10$	0.111	**0.127**	**0.264**	0.545	**0.085**	**0.113**	**0.288**	**0.548**	**0.194**	**0.270**	**0.430**	**0.593**
MFQDMA $\lambda = 0.25$	**0.105**	0.128	0.275	0.534	0.094	0.147	0.359	0.641	0.253	0.340	0.522	0.677
MFQDMA $\lambda = 0.50$	0.126	0.149	0.276	**0.509**	0.097	0.199	0.410	0.633	0.306	0.392	0.633	0.734
MFQDMA $\lambda = 0.75$	0.269	0.316	0.400	0.561	0.135	0.312	0.525	0.696	0.346	0.453	0.733	0.833
MFQDMA $\lambda = 0.90$	0.586	0.694	0.713	0.679	0.369	0.576	0.769	0.800	0.470	0.566	0.818	0.905

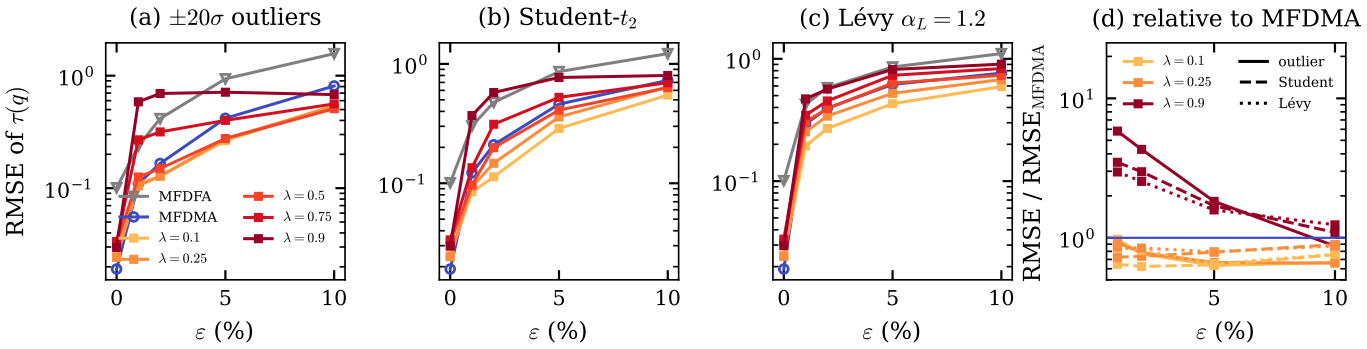


Figure 5: Robustness of $\tau(q)$ estimates of a contaminated binomial measure ($p = 0.3$, $N = 2^{16}$). (a)–(c) RMSE_τ versus the contaminated fraction ε for three contamination types (mean of five replicates). (d) Ratio $\mathrm{RMSE}_\tau(\mathrm{MFQDMA})/\mathrm{RMSE}_\tau(\mathrm{MFDMA})$ for $\lambda = 0.1, 0.25, 0.9$, paired by replicate; values below the horizontal line favor MFQDMA.

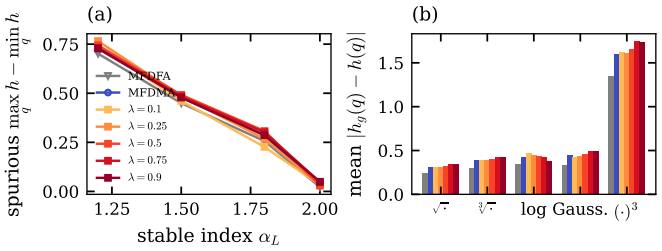


Figure 6: Negative results. (a) Spurious width of $h(q)$ for independent symmetric α-stable noise ($N = 2^{15}$, five realizations) versus the stable index. (b) Mean absolute change of $h(q)$ of a log-normal cascade under strictly increasing maps. Bar order in each group: MFDFA, MFDMA, and MFQDMA with $\lambda = 0.1, 0.25, 0.5, 0.75, 0.9$.

$H \in \{0.2, 0.3, 0.5, 0.7, 0.9\}$), all methods track H: $h(2)$ increases from 0.32 to 0.95, the spectra are flat within 0.03–0.09, and $\mathrm{RMSE}_h = 0.083$–0.092. The residual positive bias at small H is the same for all methods and is caused by the finite-size spectral synthesis of the surfaces, not by the estimators. 2D-MFQDMA and 2D-MFDMA are equivalent on monofractal surfaces.

5.3 Contaminated surfaces

Replacing even 1% of the pixels of the set-A cascade (256^2) with $\pm 20\sigma$ outliers or Student-t_2 values raises RMSE_τ of the backward methods from 0.28–0.42 to 0.9–1.6 [Figs. 8(b) and 8(c)]. 2D-MFQDMA with $\lambda = 0.25$ is better than 2D-MFDMA in six of the eight configurations, by 2%–21% (e.g., 3.16 versus 3.72 at 5% outliers), but the estimates of all methods are unusable. Contaminated surfaces should be cleaned before multifractal analysis, whichever detrending functional is used.

Table 5: RMSE_τ of 2D cascades with backward windows ($\theta = 0$): moving-mean trend (2D-MFDMA) and quantile trends (2D-MFQDMA). Bold: best per column.

	256×256			512×512		
Trend	A	B	C	A	B	C
Moving mean	0.356	0.196	0.197	0.087	0.050	0.115
$\lambda = 0.25$	**0.279**	**0.128**	**0.125**	0.099	0.040	0.137
$\lambda = 0.50$	0.316	0.173	0.157	**0.084**	**0.039**	0.115
$\lambda = 0.75$	0.422	0.282	0.252	0.109	0.103	**0.112**

Table 6: $h(2)$ of the mixed increments of fractional Brownian surfaces, backward windows: mean of three realizations, with the standard deviation in units of the last digit in parentheses.

H	2D-MFDMA	$\lambda = 0.25$	$\lambda = 0.5$	$\lambda = 0.75$
0.2	0.327(13)	0.321(17)	0.315(13)	0.320(7)
0.3	0.402(16)	0.394(21)	0.391(17)	0.394(10)
0.5	0.572(22)	0.559(28)	0.563(24)	0.562(16)
0.7	0.756(29)	0.741(35)	0.749(31)	0.748(23)
0.9	0.946(35)	0.929(41)	0.940(37)	0.940(29)

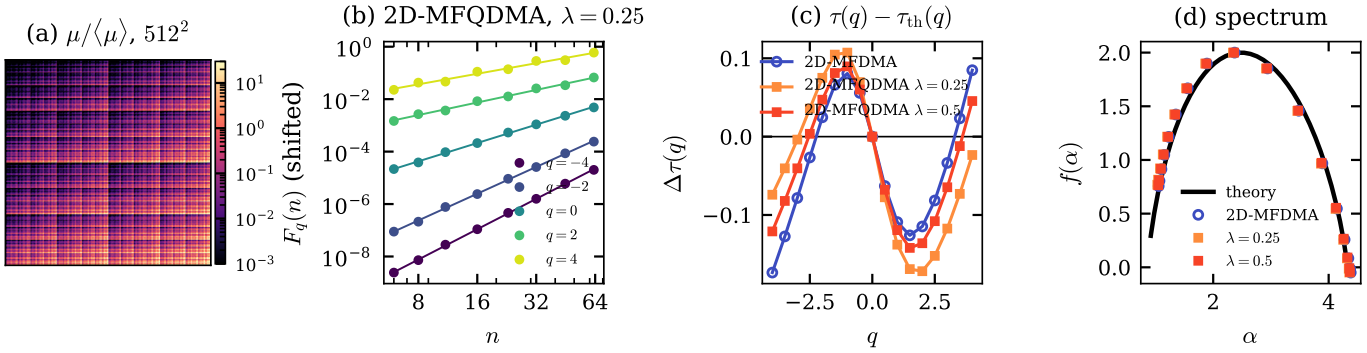


Figure 7: 2D four-multiplier cascade, set A, 512×512. (a) Measure normalized by its mean (logarithmic color scale). (b) $F_q(n)$ of backward 2D-MFQDMA ($\lambda = 0.25$), shifted vertically. (c) $\Delta\tau(q)$ for backward 2D-MFDMA and 2D-MFQDMA. (d) $f(\alpha)$ versus Eqs. (20). For this realization $\mathrm{RMSE}_\tau = 0.087$ (MFDMA), 0.099 ($\lambda = 0.25$), and 0.084 ($\lambda = 0.5$).

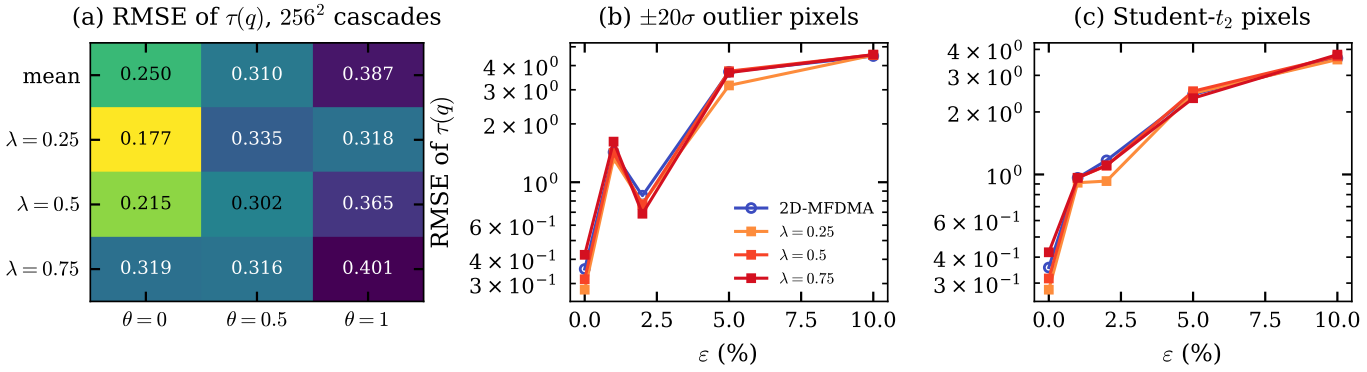


Figure 8: (a) RMSE_τ of 256^2 cascades (mean over sets A–C) for all combinations of trend functional and window position. (b), (c) RMSE_τ of the set-A cascade (256^2) when a fraction ε of pixels is replaced by $\pm 20\sigma$ outliers or Student-t_2 values (mean of two replicates).

6 Applications

6.1 Financial markets

Financial returns are the most widely studied multifractal signals [25, 48–54]. They are also the clearest case in which a multifractal estimator must cope with outliers that are not part of the process: erroneous prints ("bad ticks"), mis-keyed prices, and microstructure noise are routinely removed before econometric analysis [55–57]. Filters are never perfect, so the sensitivity of an estimator to the errors that remain is a practical question.

* Data. We analyze daily adjusted closing prices $P(t)$ of 18 assets from the `qrmdata` collection of Hofert, Hornik, and McNeil [58, 59] (Yahoo Finance data, read with a dependency-free parser of the R serialization format). The assets are 12 stock indices: S&P 500, Dow Jones Industrial Average (DJIA), NASDAQ-100, FTSE 100, DAX, CAC 40, SMI, Euro Stoxx 50, Nikkei 225, Hang Seng, Shanghai Stock Exchange (SSE) Composite, and CSI 300. They also include three exchange rates against the U.S. dollar, gold, Brent crude oil, and Bitcoin. The SSE Composite, DJIA, and NASDAQ-100 cover the markets studied with MFDMA and MF-X-DMA in Refs. [18, 22]. In addition, we analyze the 29 DJIA constituents with complete price histories from 2000 to 2015 ($N = 4024$ each). Weekend records of the non-crypto assets are removed, and we analyze the log returns $r(t) = \ln P(t) - \ln P(t-1)$ in the signal convention with backward windows, $10 \le n \le \min(500, N/10)$ days (MFDFA: $20 \le s \le \min(1000, N/5)$), and $q \in [-4, 4]$. The series contain between 2662 and 16 606 returns, with excess kurtosis from 1.8 (EUR/USD) to 159 (gold) (Table 7).

* Spectra and the origin of multifractality. Figure 10 shows the analysis of the S&P 500 (1950–2015). The fluctuation functions of MFQDMA are power laws from two weeks to two years [Fig. 10(b)], and all estimators give $h(2) \approx 0.50$–0.53, i.e., uncorrelated returns, as expected for an efficient market. The generalized Hurst exponents decrease with q, but the estimators differ at negative q [Fig. 10(c)]. Across the 18 assets the moving-average methods often produce a nonmonotone $h(q)$ at $q < 0$, which yields a hooked, non-concave $f(\alpha)$ [Fig. 10(d)]: $h(q)$ is monotone within 0.005 for 16 of 18 assets with MFDFA, 15 with MFQDMA at $\lambda = 0.1$, and 11 with MFDMA. Spectral widths therefore depend on the estimator (Table 7, Fig. 11), and differences in $\Delta\alpha$ of order 0.1 between markets should not be overinterpreted for any single method. The fit quality is $R^2 \ge 0.96$ for every asset and every estimator except gold, whose record contains 4.8% zero returns from

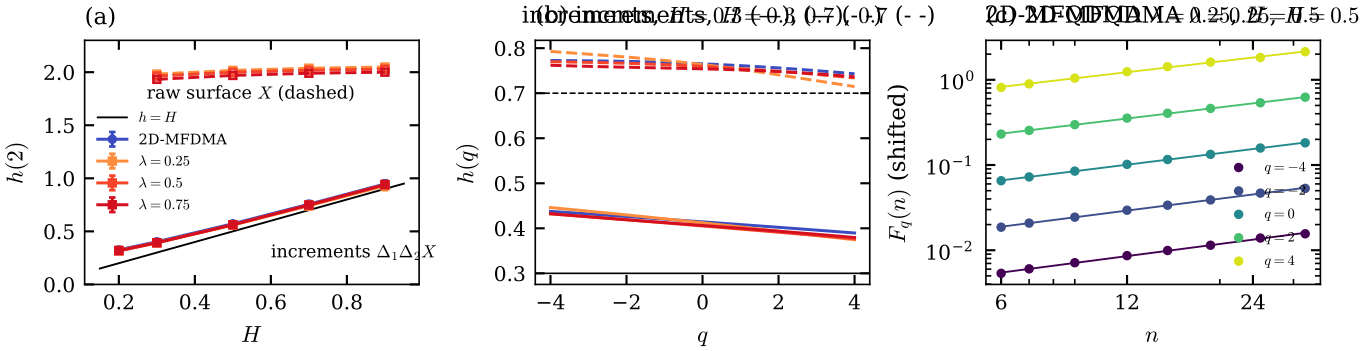


Figure 9: Fractional Brownian surfaces (256^2). (a) $h(2)$ versus H from the raw surface X (dashed) and from its mixed increments $\Delta_1\Delta_2 X$ (solid, mean $\pm$ s.d.); the black line is $h = H$. (b) $h(q)$ of the increments for $H = 0.3$ (solid) and $H = 0.7$ (dashed); horizontal lines mark H. (c) $F_q(n)$ of 2D-MFQDMA ($\lambda = 0.25$) on the increments of a surface with $H = 0.5$.

Table 7: Daily log returns of 18 assets: period, length N, excess kurtosis κ, spectral width $\Delta\alpha$ for five estimators, and the percentage of $\Delta\alpha$ retained by shuffled and Fourier phase-randomized surrogates (mean of five surrogates each) for MFDMA and MFQDMA with $\lambda = 0.1$. Backward windows; data from Ref. [58].

				$\Delta\alpha$					Shuffled (%)		Phase-rand. (%)	
Asset	Period	N	κ	MFDFA	MFDMA	$\lambda = 0.1$	$\lambda = 0.25$	$\lambda = 0.5$	MFDMA	$\lambda = 0.1$	MFDMA	$\lambda = 0.1$
S&P 500	1950–15	16606	27.3	0.28	0.18	0.27	0.25	0.21	123	73	34	24
DJIA	1985–15	7796	41.9	0.38	0.23	0.17	0.19	0.26	124	80	24	52
NASDAQ-100	1985–15	7627	7.5	0.16	0.09	0.16	0.06	0.12	201	146	66	56
FTSE 100	1984–15	8332	9.7	0.22	0.38	0.38	0.25	0.42	43	61	22	31
DAX	1990–15	6354	4.6	0.27	0.26	0.49	0.29	0.29	57	47	25	18
CAC 40	1990–15	6548	4.4	0.22	0.25	0.14	0.15	0.26	57	97	27	80
SMI	1990–15	6349	6.4	0.32	0.18	0.26	0.17	0.19	88	79	34	31
Euro Stoxx 50	1987–15	7444	5.7	0.29	0.27	0.41	0.30	0.28	49	51	15	16
Nikkei 225	1984–15	7879	8.0	0.33	0.32	0.29	0.33	0.37	43	78	22	40
Hang Seng	1987–15	7213	57.7	0.34	0.45	0.24	0.26	0.50	74	102	31	73
SSE Composite	1990–15	6391	151.9	0.72	0.72	0.47	0.58	0.73	77	135	15	33
CSI 300	2005–15	2662	3.2	0.12	0.19	0.22	0.16	0.25	81	120	49	37
EUR/USD	2000–15	4173	1.8	0.24	0.25	0.54	0.39	0.27	62	29	35	24
GBP/USD	2000–15	4173	4.0	0.19	0.08	0.39	0.19	0.09	246	62	124	16
JPY/USD	2000–15	4173	3.1	0.29	0.24	0.29	0.31	0.27	71	79	27	36
Gold	1970–15	9690	159.4	0.58	0.96	1.12	1.10	0.97	51	53	6	9
Brent oil	1987–15	7257	14.6	0.20	0.18	0.26	0.20	0.17	98	63	59	53
Bitcoin	2010–18	2873	88.0	0.66	0.58	0.60	0.53	0.58	113	111	21	20

unchanged daily fixings ($R^2 = 0.90$–0.99).

The surrogate decomposition [11, 32, 60] gives a consistent answer for all methods (Fig. 11(b)). Shuffled returns retain a median 76% of $\Delta\alpha$ (interquartile range 60%–112%) over all assets and estimators, whereas phase-randomized returns, which keep the linear correlations but Gaussianize the amplitudes, retain only 31% (21%–43%). The shuffled surrogate retains more than the phase-randomized surrogate in 99% of the 126 asset–estimator pairs. The multifractality of daily returns is therefore carried mainly by the fat-tailed return distribution and finite-size effects, in agreement with Refs. [32, 60] and in contrast with the correlation-driven multifractality of the EEG (Sec. 6.2). MFQDMA reaches the same conclusion as MFDMA and MFDFA (medians 79%/32% for $\lambda = 0.1$, 75%/27% for MFDMA, and 59%/15% for MFDFA). This is consistent with Sec. 4.6: changing the trend functional does not remove distribution-driven width.

* Influence of genuine extreme days. A spectral estimate that depends on a handful of days is fragile in any finite sample. We remove the k largest absolute returns of each asset ($k = 1$, 5, 10, 20, 50; at most 1.9% of the sample) and measure the change of the exponents by

$$D_h = \left\{ \frac{1}{17} \sum_q \left[h_{\mathrm{mod}}(q) - h(q) \right]^2 \right\}^{1/2}, \qquad (23)$$

where $h_{\mathrm{mod}}(q)$ is computed from the modified series with the same estimator. For $k \leq 20$, MFQDMA at low quantile levels is the least sensitive estimator [Fig. 12(a), Table 8]. For $k = 5$ the median D_h over the 18 assets is 0.012 for $\lambda = 0.1$, versus 0.024 for MFDMA and 0.017 for MFDFA, and $\lambda = 0.1$ is less sensitive than MFDMA for 14 of 18 assets (Wilcoxon signed-rank $p = 0.03$). The median change of the spectral width is 0.024, versus 0.058 for MFDMA and 0.067 for MFDFA. A single day illustrates the point. Removing the crash of 19 October 1987 ($r = -0.229$, a 34σ event in robust units) from 16 606 S&P 500 returns lowers $\Delta\alpha$ by 0.093 for MFDMA

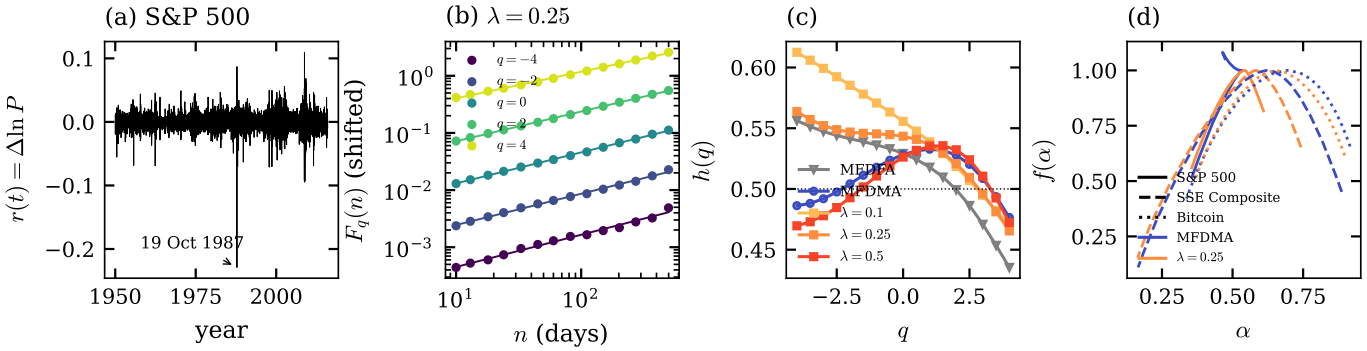


Figure 10: Daily returns. (a) Log returns of the S&P 500, 1950–2015. (b) $F_q(n)$ of backward MFQDMA ($\lambda = 0.25$), shifted vertically. (c) $h(q)$ of the S&P 500 for MFDFA, MFDMA, and MFQDMA ($\lambda = 0.1$, 0.25, 0.5); the dotted line marks $h = 1/2$. (d) $f(\alpha)$ of the S&P 500, the SSE Composite, and Bitcoin for MFDMA and MFQDMA ($\lambda = 0.25$).

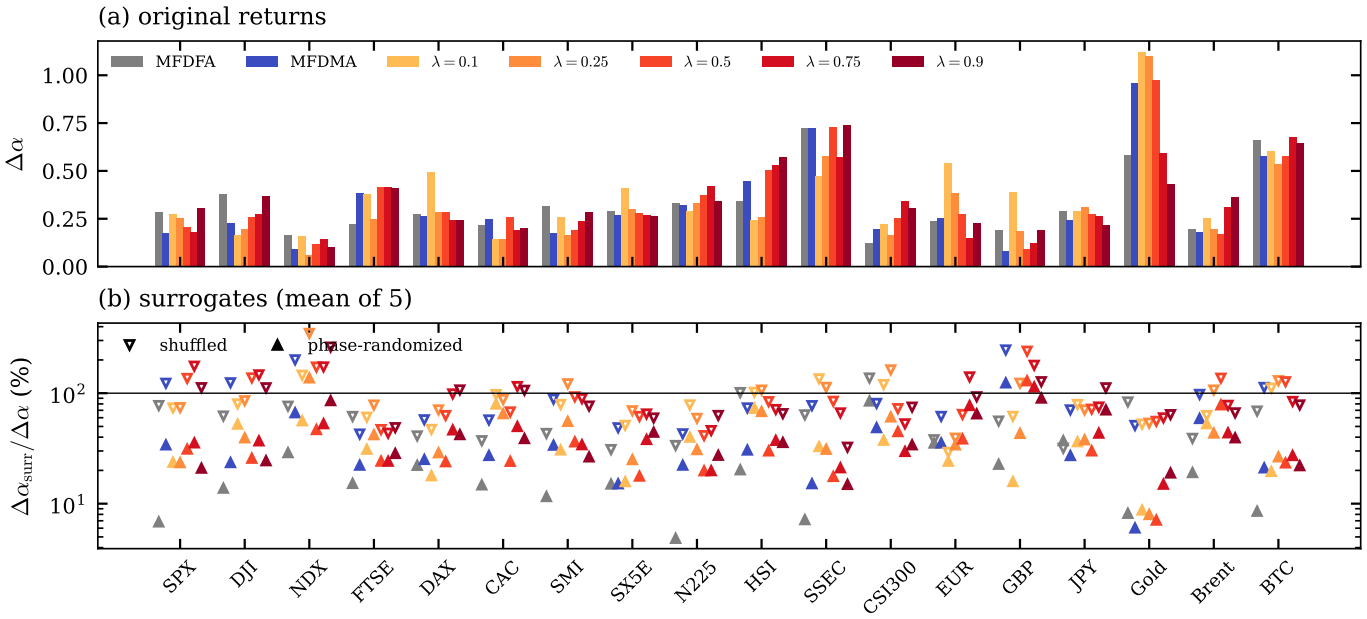


Figure 11: Eighteen assets. (a) Spectral width $\Delta\alpha$ of daily returns for MFDFA, MFDMA, and MFQDMA at five quantile levels. (b) Percentage of $\Delta\alpha$ retained by shuffled (open down-triangles) and phase-randomized (filled up-triangles) surrogates, mean of five each; colors as in (a).

and 0.085 for MFDFA, but only by 0.010 for MFQDMA with $\lambda = 0.1$. For $k = 50$ the advantage disappears, because removing the 50 largest returns changes the return distribution itself. The median ($\lambda = 0.5$) is *more* sensitive than MFDMA at every k. Low quantile levels are required.

* Injected data errors. To quantify the utility of MFQDMA in the presence of noise and outliers, we contaminate the real returns with three kinds of data errors and measure D_h relative to the estimate from the uncontaminated series. All three are expressed in units of the return standard deviation σ_r. (i) *Bad ticks*: an isolated misprinted price $\ln P(t) \to \ln P(t) \pm 10\sigma_r$ at a fraction ε of randomly chosen days. Each bad tick creates a return pair $(+s, -s)$. (ii) *Return outliers*: an erroneous return $r(t) \to r(t) \pm 10\sigma_r$. This is a permanent level shift of the price. (iii) *Microstructure noise*: i.i.d. Gaussian noise of standard deviation $\eta\sigma_r$ added to every log price [57]. We use $\varepsilon \in \{0.1, 0.25, 0.5, 1\}\%$, $\eta \in \{0.1, 0.25, 0.5, 1\}$, five replicates per asset (18 assets), and three per stock (29 DJIA stocks). The results are shown in Figs. 12 and 13 and Table 8.

The three error types give three different answers, each of which follows from the mechanism discussed in Sec. 4.5.

Bad ticks are where MFQDMA is most useful. Every MFQDMA variant is less sensitive than MFDMA at every contamination level, except $\lambda = 0.5$ at $\varepsilon = 0.1\%$, which ties [Fig. 12(b)]. At $\varepsilon = 0.5\%$ the median ratio $D_h(\mathrm{MFQDMA})/D_h(\mathrm{MFDMA})$ is 0.82–0.95 over the 18 assets and 0.79–0.92 over the 29 stocks. With $\lambda = 0.25$, MFQDMA is less sensitive for 18 of 18 assets ($p = 8 \times 10^{-6}$) and 28 of 29 stocks ($p = 10^{-8}$). MFDFA is about twice as sensitive as either moving-window method (ratio 2.0 and 1.9). In the rolling-window analysis [Fig. 13(c)], 0.5% bad ticks change the ten-year $\Delta\alpha$ of the S&P 500 by 0.017 on average for $\lambda = 0.1$,

Table 8: Sensitivity of $h(q)$ of daily returns to genuine extremes and to injected data errors: median D_h [Eq. (23), $\times 10^3$] over 18 assets ("Mkts") or 29 DJIA stocks ("Stocks"). "Top 5": removal of the five largest absolute returns, with the median $|\delta\Delta\alpha|$ in the next column. Errors are injected at $\varepsilon = 0.5\%$ (bad ticks, return outliers) or $\eta = 0.5$ (microstructure noise). "Rolling": mean $|\delta\Delta\alpha|$ caused by 0.5% bad ticks in 113 ten-year windows of the S&P 500. Bold: least sensitive per column.

	Top 5 removed		Bad ticks		Return outliers		Microstructure noise		Rolling
Method	D_h	$\lvert\delta\Delta\alpha\rvert$	Mkts	Stocks	Mkts	Stocks	Mkts	Stocks	$\lvert\delta\Delta\alpha\rvert$
MFDFA	16.9	0.067	25.1	25.6	65.1	71.8	49.4	51.9	0.048
MFDMA ($\theta = 0$)	23.5	0.058	12.6	13.5	59.4	65.8	**35.4**	**39.7**	0.026
MFQDMA $\lambda = 0.1$	**12.2**	**0.024**	**10.5**	11.3	**47.1**	**59.4**	42.2	60.8	**0.017**
MFQDMA $\lambda = 0.25$	15.7	0.037	11.4	12.1	56.6	60.1	36.0	46.9	0.020
MFQDMA $\lambda = 0.5$	27.1	0.071	12.2	12.5	65.1	68.9	37.1	41.8	0.025
MFQDMA $\lambda = 0.75$	27.3	0.057	10.6	11.0	66.1	79.7	43.1	50.5	—
MFQDMA $\lambda = 0.9$	25.6	0.039	10.6	**10.8**	58.8	86.5	64.4	72.6	—

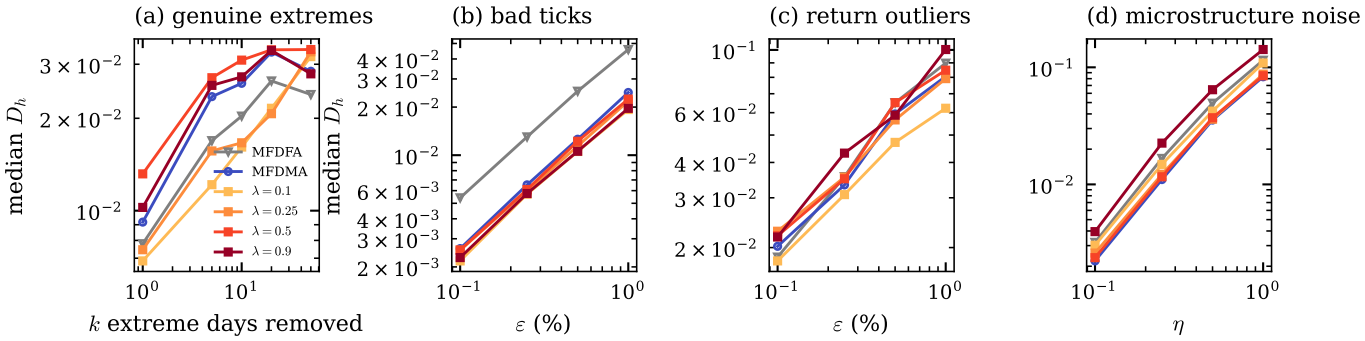


Figure 12: Sensitivity of $h(q)$ of daily returns (median D_h over 18 assets). (a) Removal of the k largest absolute returns. (b) Isolated bad ticks at a fraction ε of days. (c) Erroneous returns (price level shifts) at a fraction ε of days. (d) Microstructure noise of relative amplitude η on every log price. Colors and markers as in panel (a).

versus 0.026 for MFDMA and 0.048 for MFDFA. The mechanism is that a bad tick is a *transient* profile outlier: since $y(i) = \ln P(i) - \ln P(0) - i\langle r\rangle$, it displaces a single profile value. A single value changes a moving λ-quantile of n profile values by at most one rank, whereas it shifts the moving mean by s/n in all n windows that contain it.

Return outliers are persistent profile steps. As argued in Sec. 4.5, no order statistic can remove them. Low quantile levels still help ($\lambda = 0.25$: less sensitive for 25 of 29 stocks, $p = 7 \times 10^{-4}$, median ratio 0.84; $\lambda = 0.1$: ratio 0.90 over the 18 assets, not significant). The median and upper quantiles are significantly *more* sensitive than MFDMA for the stocks.

Microstructure noise adds white noise to every profile value. For dense Gaussian perturbations the mean is the efficient location estimator, and a quantile is less efficient (the median has an asymptotic efficiency of $2/\pi$ for Gaussian data [29]). Accordingly, MFDMA is the least sensitive method at every noise level [Fig. 12(d)]. MFQDMA with $\lambda = 0.25$ or 0.5 is close to it (median ratios 1.03–1.12), whereas $\lambda = 0.1$ and $\lambda = 0.9$ are 1.3–1.9 times more sensitive.

For the financial analyst, the conclusion is specific. Backward MFQDMA with $\lambda \approx 0.1$–0.25 is the estimator of choice when the price record may contain isolated erroneous prints or when estimates should not hinge on a few crash days. MFDMA remains preferable for dense microstructure noise, and MFDFA is the most sensitive of the three to sparse errors. Because D_h measures sensitivity rather than accuracy, the remaining errors in all cases ($D_h \approx 0.01$–0.06 at $\varepsilon = 0.5\%$) are small compared with typical between-market differences in $h(q)$. They are, however, comparable to the differences in $\Delta\alpha$ that are often interpreted as differences in market efficiency [25].

6.2 Electroencephalograms

∗ Data and preprocessing. We analyze the Bonn EEG database [39]: five sets of 100 single-channel segments of 4097 samples (23.6 s at 173.61 Hz) band-limited to 0.53–40 Hz by the original authors. Sets Z and O are surface recordings from healthy volunteers with eyes open and closed. Sets N and F are intracranial interictal recordings from the hemisphere opposite to and within the epileptogenic zone. Set S contains ictal recordings. The data are already segmented and filtered, so we apply no further filtering. Table 9 shows why: refiltering changes $h(q, 1/2)$ by more than the between-class effect (0.055 for F versus N), whereas affine operations change nothing (Proposition 2). Local-slope analysis shows no single power law at large scales. The slopes collapse beyond $n \approx 200$ samples, close to the period of the 0.53 Hz high-pass corner (327 samples). We therefore fit $12 \le n \le 128$ samples (69–737 ms) in the signal convention, with $\theta = 0$, $q \in [-5, 5]$, and $\lambda \in \{0.1, 0.25, 0.5, 0.75, 0.9\}$.

∗ Origin of the multifractality. Following Refs. [11, 18, 22], we compare 20 segments per set with shuffled surrogates,

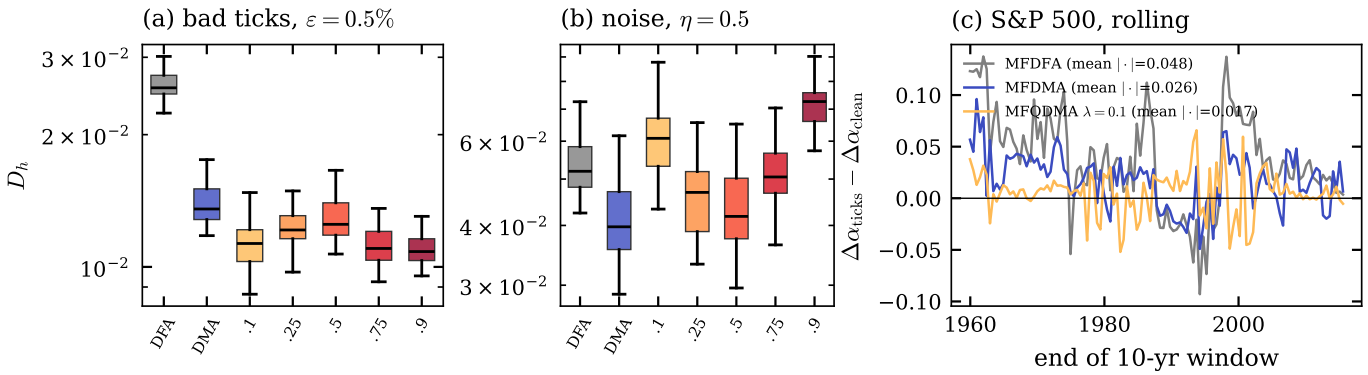


Figure 13: (a), (b) Distribution of D_h over 29 DJIA stocks (2000–2015) for 0.5% bad ticks and for microstructure noise with $\eta = 0.5$; boxes show quartiles, whiskers the 5%–95% range; labels give the method or λ. (c) Change of $\Delta\alpha$ caused by 0.5% bad ticks in rolling ten-year windows of the S&P 500 (step: half a year).

Table 9: Sensitivity of the EEG exponents to preprocessing: mean $|\Delta h(q,1/2)|$ relative to the unprocessed data. For comparison, the median F-versus-N difference is 0.055.

Operation	mean $\lvert\Delta h(q,1/2)\rvert$
z scoring	0.0000
50 Hz notch	0.0004
Linear detrending	0.0068
Downsampling by 2	0.0344
FIR high-pass, 0.5 Hz	0.1024
FIR band-pass, 1–40 Hz	0.2988

which destroy all correlations but keep the distribution, and with Fourier phase-randomized surrogates [40, 41], which keep the linear correlations but destroy nonlinear ones. Table 10 and Fig. 14 show the results. Shuffling retains only 12%–42% of the spectral width $\Delta\alpha$ for every method and set, and the width of the originals exceeds that of the shuffled data at $p < 4 \times 10^{-6}$ (Mann–Whitney) in all 35 set–method combinations. Phase randomization retains 64%–122%. The multifractality of the EEG is therefore correlation-driven and predominantly linear; the amplitude distribution accounts for a minority (12%–42%) of it. This conclusion is the same for all seven estimators. The ictal set S is where the methods differ: MFQDMA with $\lambda = 0.9$ gives the best mean fit quality ($R^2 = 0.910$, versus 0.841 for MFDMA and 0.870 for MFDFA) [Fig. 14(c)], whereas $\lambda = 0.5$ gives the worst (0.812). Scaling exponents of the ictal recordings should be interpreted with care for all methods.

* The λ-resolved exponents. For signals, Corollary 2 does not apply, and $h(q,\lambda)$ can depend on λ. Figure 15 shows the median $h(q,\lambda)$ of all 500 segments. In every set, h is largest at the extreme quantiles for $q < 0$. The ictal set has a distinct pattern, with the smallest exponents at $\lambda = 0.25$–0.75 for $q > 0$. We summarize the λ dependence by the spread $W(q) = \max_\lambda h(q,\lambda) - \min_\lambda h(q,\lambda)$. At $q = -5$ its median increases monotonically from the healthy sets (Z: 0.222; O: 0.232) through the interictal sets (N: 0.293; F: 0.334) to the ictal set (S: 0.524) [Fig. 16(a)]. This ordering is significant (Kruskal–Wallis $H = 186.6$, $p = 3 \times 10^{-39}$; Spearman correlation with a clinical severity rank $\rho = 0.60$, $p = 10^{-50}$), and all nine pathology contrasts remain significant after Benjamini–Hochberg correction. The specificity control Z versus O, a nonpathological change of state, is null ($p = 0.84$, Cliff's $\delta = 0.017$). The spread is correlated with $\Delta\alpha$ only moderately (Spearman $\rho = 0.40$). It is not redundant with it: $\Delta\alpha$ peaks in set N, whereas W increases monotonically up to set S. We stress that $q = -5$ was selected after inspecting the data. Figure 16(b) therefore reports $W(q)$ for all q; the class ordering holds for $q \le 0$.

* Discrimination. To test whether $h(q,\lambda)$ carries information beyond $h(q)$, we trained ℓ_2-regularized logistic classifiers with standardized features and evaluated them by 10×10 repeated stratified cross-validation (Table 11). At equal feature dimension (21 exponents), the best single λ improves the AUC of MFDMA for all four contrasts. The largest gain is for the clinically hardest contrast, F versus N (0.789 versus 0.713), which separates interictal recordings from the epileptogenic zone and the opposite hemisphere. Using all five quantile levels improves the AUC significantly in all four tasks (paired Wilcoxon test over the 100 folds, $p = 2 \times 10^{-7}$ to 2×10^{-5}). A dimension-matched null model, in which MFDMA is complemented by four noisy copies of its own exponents, gives no gain (-0.033 to $+0.003$), so the improvement reflects information rather than dimensionality. Univariate effect sizes are largest at low λ and positive q [Fig. 16(c)], where $|\delta|$ reaches 0.43 compared with at most 0.32 for MFDMA; in the multivariate classifier the best single level for F versus N is nevertheless $\lambda = 0.75$, which indicates that the levels carry complementary information. Consistent with Proposition 2, the antisymmetric index $h(q,\lambda) - h(q,1-\lambda)$ performed poorly (AUC 0.51–0.68) and reduced the AUC when added to $h(q)$ in three of the four tasks.

These EEG results come from one public database with 100 segments per class drawn from a small number of patients. They show that the λ dimension carries information, but they do not establish a clinical biomarker. Replication on multichannel databases is required.

Table 10: Bonn EEG, 20 segments per set. Spectral width $\Delta\alpha$ of the originals, and percentage of $\Delta\alpha$ retained by shuffled and phase-randomized surrogates (one surrogate per segment).

	$\Delta\alpha$ (originals)					Shuffled (%)					Phase-randomized (%)				
Method	Z	O	N	F	S	Z	O	N	F	S	Z	O	N	F	S
MFDFA	0.360	0.535	0.681	0.870	1.050	30	18	15	19	12	103	87	94	74	64
MFDMA	0.471	0.401	0.882	0.830	0.724	29	33	16	20	21	97	80	83	88	99
MFQDMA $\lambda = 0.10$	0.526	0.568	0.934	1.095	1.080	33	37	21	21	16	111	85	101	91	94
MFQDMA $\lambda = 0.25$	0.461	0.510	0.847	0.968	0.871	37	42	17	21	21	122	83	97	82	103
MFQDMA $\lambda = 0.50$	0.496	0.444	0.886	0.847	0.817	29	33	19	22	22	101	89	86	89	96
MFQDMA $\lambda = 0.75$	0.586	0.406	0.978	0.863	0.895	29	35	21	22	20	81	93	79	96	96
MFQDMA $\lambda = 0.90$	0.664	0.497	1.097	1.022	1.170	25	34	19	20	15	79	98	83	97	86

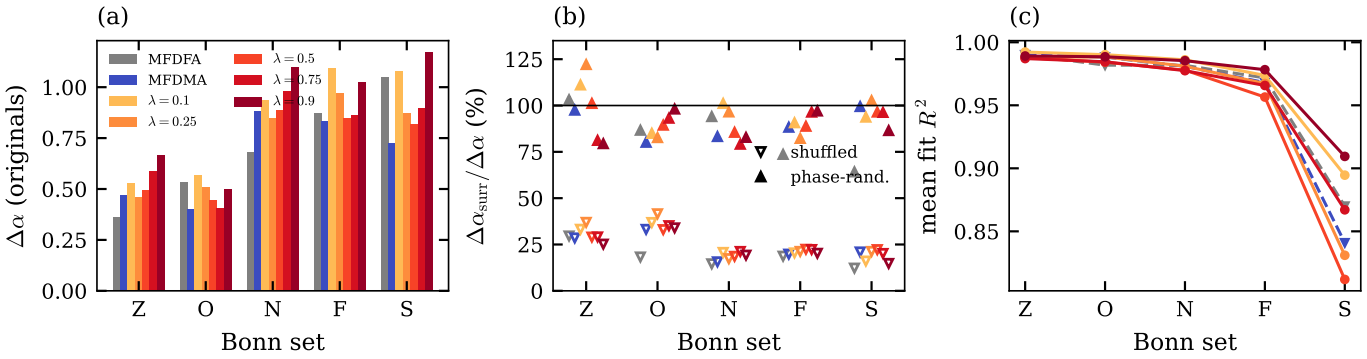


Figure 14: Bonn EEG, 20 segments per set. (a) Mean spectral width $\Delta\alpha$ of the originals. (b) Fraction of $\Delta\alpha$ retained by shuffled (open down-triangles) and phase-randomized (filled up-triangles) surrogates. (c) Mean fit quality R^2. Colors as in (a).

Table 11: Cross-validated AUC (10×10 folds) of classifiers based on 21 exponents.

Contrast	MFDMA $h(q)$	MFQDMA $h(q, \lambda^*)$	λ^*
F vs N	0.713	**0.789**	0.75
S vs F	0.938	**0.951**	0.90
Z vs S	0.991	**0.998**	0.10
Z vs O	0.896	**0.920**	0.10/0.90

6.3 Terrain elevation

As a 2D application we analyze a digital elevation model of the Jacksboro fault region (344×403 grid; distributed with Matplotlib [61]). Elevation is a self-affine surface, so by Proposition 3 we analyze its mixed increments (Fig. 17). Applied to the raw elevation, all methods give $h(2) = 1.93$–1.96, the saturation value. Applied to the increments, they give $h(2) = 0.59$–0.66, an estimate of the roughness exponent of the terrain, with $R^2 = 0.96$–0.97 (Table 12). Shuffling the increments yields $h(2) \approx 1.04$–1.07, the uncorrelated value $D_f/2$, and removes 80%–97% of $\Delta\alpha$. A 2D phase-randomized surrogate retains 57%–80% of $\Delta\alpha$. As for the EEG, the multifractality of the terrain is mainly carried by linear correlations, with a nonlinear contribution of 20%–43%. 2D-MFQDMA with $\lambda = 0.25$ gives the best fit ($R^2 = 0.973$), but the differences between methods are small.

We also analyzed a 256×256 magnetic-resonance image. Neither method produced acceptable scaling: $R^2 = 0.58$–0.61, except 0.80 for $\lambda = 0.5$. The resulting spectra violated $f(\alpha) \leq D_f$, so we do not report them. A large fraction of the image is constant background, and the increments vanish there.

Table 12: Terrain elevation model: $h(2)$ of the raw surface and of its mixed increments; $\Delta\alpha$ of the increments and of their surrogates; mean R^2 for the increments. Rows labeled by λ are 2D-MFQDMA; backward windows except where noted.

	$h(2)$		$\Delta\alpha$			
Trend	raw	incr.	orig.	shuf.	phase	R^2
MFDMA	1.960	0.622	0.545	0.037	0.379	0.970
MFDMA$_{\theta=1/2}$	1.927	0.661	0.524	—	—	0.963
$\lambda = 0.25$	1.953	0.634	0.559	0.017	0.449	0.973
$\lambda = 0.50$	1.955	0.589	0.588	0.045	0.403	0.965
$\lambda = 0.75$	1.953	0.631	0.696	0.135	0.398	0.972

7 Discussion

7.1 Practical guidelines

Table 13 collects the rules established in Secs. 3–6. The two most important are negative ones. Never use the centered median on a measure (Corollary 1). Never analyze a self-affine surface directly (Proposition 3). Both mistakes produce excellent-looking power laws ($R^2 > 0.99$) with meaningless exponents and are therefore easy to miss.

7.2 Computational cost

The moving mean costs $O(N)$ per scale by cumulative sums, whereas a moving quantile computed by selection in each window costs $O(Nn)$. For $N = 2^{16}$ and 14 scales, our vectorized implementation needs 0.13 s for MFDMA, 4.3 s for MFQDMA at one λ, and 7.9 s for five λ values in a single

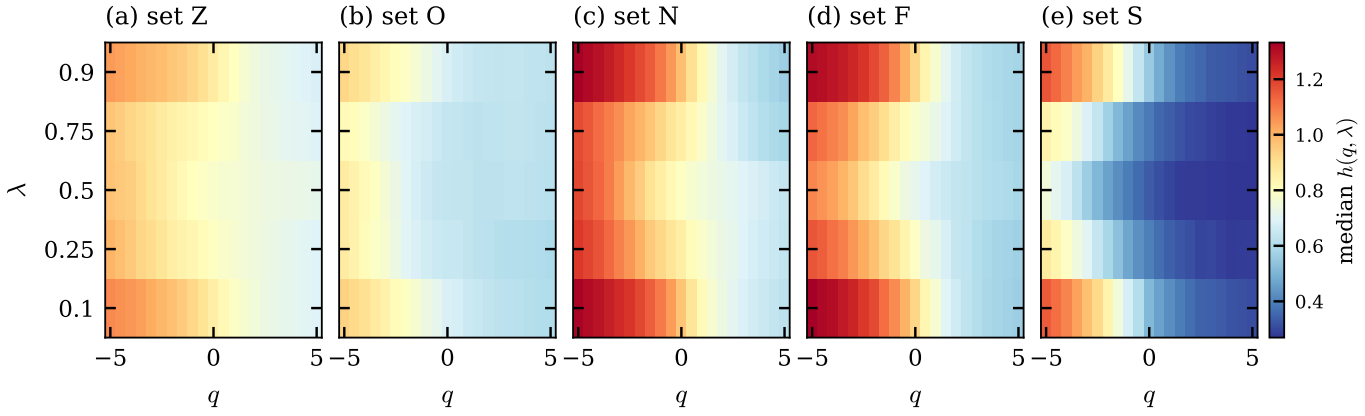


Figure 15: Median λ-resolved generalized Hurst exponents $h(q,\lambda)$ of the Bonn EEG sets (100 segments each, $\theta = 0$, $12 \le n \le 128$). (a) Z, healthy, eyes open; (b) O, healthy, eyes closed; (c) N, interictal, contralateral; (d) F, interictal, epileptogenic zone; (e) S, ictal.

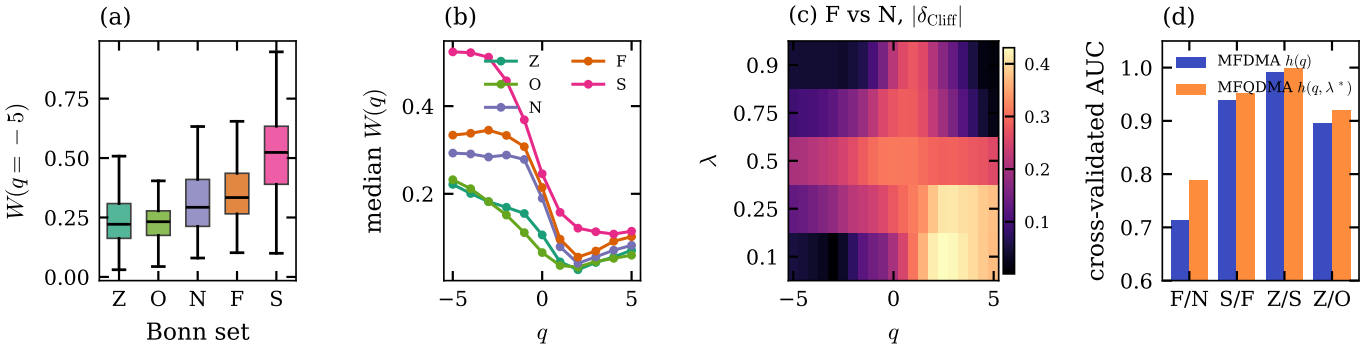


Figure 16: Information carried by the quantile level (500 EEG segments). (a) λ-spread $W(q=-5)$ per set. (b) Median $W(q)$ versus q. (c) Effect size $|\delta|$ (Cliff) of the F-versus-N contrast for each (q,λ). (d) Cross-validated AUC of logistic classifiers using the 21 exponents $h(q)$ of MFDMA or the 21 exponents $h(q,\lambda^*)$ of MFQDMA at the best single λ^* (equal feature dimension).

selection pass (MFDFA: 0.02 s). Running-quantile algorithms based on balanced trees or double heaps reduce the cost to $O(N \log n)$ [38]. In 2D the quantile trend requires all n^2 partial sums per window, at cost $O(N_1 N_2 n^2)$ per scale: 4.1 s per λ for a 255×255 field, versus 0.05 s for the separable mean trend. The cost is moderate for single images and is the main practical limitation for large image collections.

7.3 Relation to other estimators and to the companion paper

MFQDMA belongs to the DMA family and inherits its interpretation [13, 16, 18]; for measures it reduces to a sharp gliding-box estimator (Proposition 1). It is complementary to wavelet-leader methods [7–9], which achieve robustness through supremum-based leaders rather than through the trend. The λ dimension is a new, trend-side degree of freedom. For measures it is redundant (Corollary 2); for signals it carries information (Sec. 6.2).

The bivariate generalization raises questions that do not arise here, and we treat them in Paper II [42]. The cross-covariance of two residual series can be negative, so the qth-order cross moments require a sign convention [21–23]. The two series may be detrended at different quantile levels (λ_x, λ_y). The theory of Proposition 1 extends to products of box masses of two measures. It gives exact cross exponents of cascade pairs and a Cauchy–Schwarz bound, which show that the relation $h_{xy} = (h_{xx} + h_{yy})/2$ used in Refs. [21, 22] holds only for identical multipliers. Quantile levels $\lambda \neq 1/2$ also leave a residual offset that is harmless in the univariate fluctuation function but creates spurious cross-correlation; Paper II quantifies this effect and gives remedies for it. Paper II also reports benchmarks on binomial-measure pairs, bivariate fractional Brownian motions, and two-component ARFIMA processes, following Ref. [22]. It also covers the bivariate counterpart of Sec. 6.1: cross-correlations between pairs of the assets analyzed here, and their sensitivity to errors in one or both price records. There, a median trend combined with first-order moments gives the scale-resolved correlation estimate that is least sensitive to bad ticks.

7.4 Limitations

(i) On clean 1D data MFQDMA is not more accurate than backward MFDMA (Tables 2 and 3). (ii) Its robustness is partial and does not extend to heavy-tailed or monotonically

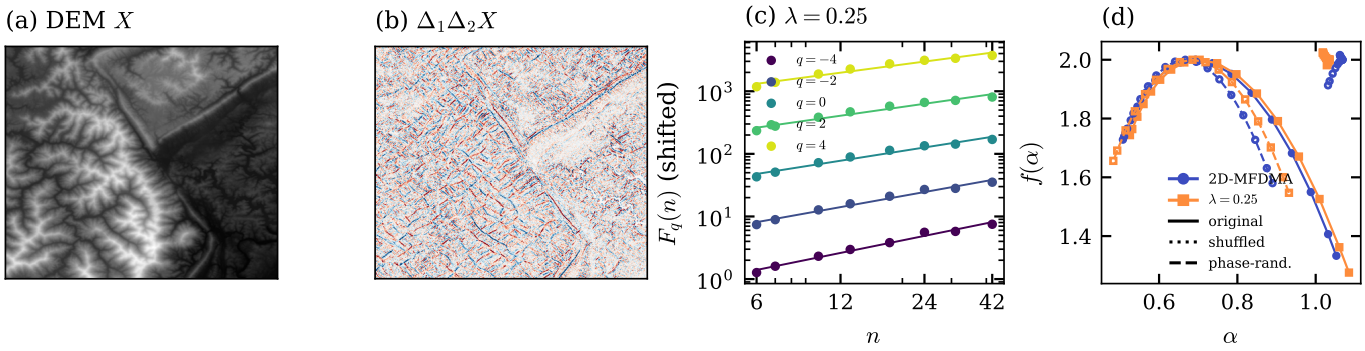


Figure 17: Terrain elevation model. (a) Elevation X. (b) Mixed increments $\Delta_1\Delta_2 X$. (c) $F_q(n)$ of 2D-MFQDMA ($\lambda = 0.25$) on the increments, shifted vertically. (d) $f(\alpha)$ of the increments (solid), of shuffled increments (dotted, clustered near $\alpha \approx 1$), and of a 2D phase-randomized surrogate (dashed), for 2D-MFDMA and 2D-MFQDMA.

Table 13: Recommended use of MFQDMA.

Choice	Recommendation
Window position	$\theta = 0$; avoid $\theta = 1-\lambda$ for measures (degenerate), and $\theta = 1/2$ in general.
Quantile level	$\lambda \in [0.1, 0.25]$ for contaminated data; never $\lambda \geq 0.75$ if outliers are suspected. Use several λ for λ-resolved features of signals.
1D input	Measure: profile without demeaning. Signal: demeaned profile.
2D input	Measure: direct. Stationary field: demeaned. Self-affine surface: mixed increments.
Scales	As for MFDMA [18]; check local slopes for crossovers such as filter corners.
Financial returns	$\theta = 0$, $\lambda \approx 0.1$–0.25 if isolated bad prints or crash days may dominate; MFDMA for dense microstructure noise; avoid $\lambda \geq 0.5$ when outliers are suspected.
Diagnostics	Shuffled and phase-randomized surrogates; check $f(\alpha) \leq D_f$, monotonicity of $h(q)$, and R^2.

transformed signals (Sec. 4.6). (iii) In 2D the accuracy gains are largest for small fields and shrink at 512^2 (Table 5). (iv) Our financial analysis uses daily data; the intraday (5-min) SSE Composite returns of Ref. [18] and tick data, where bad ticks and microstructure noise are most prevalent [55, 56], remain to be analyzed. The financial data errors were injected synthetically into real returns, so D_h measures sensitivity, not accuracy. The EEG and terrain results each rest on a single database. (v) The quantile level λ and the moment order q for the EEG spread W were not preregistered.

8 Conclusions

We have introduced MFQDMA, which replaces the moving average of MFDMA by a moving quantile in one and two dimensions, and characterized it exactly and numerically. For non-negative measures the quantile residual is a gliding-box mass. This identity explains the correctness of the method on measures, the independence of its exponents from λ, and the degeneracy line $\theta + \lambda = 1$ that contains the centered median. For signals the exponents are affine invariant and obey the reflection symmetry $h(q,\lambda;-x) = h(q,1-\lambda;x)$. Surfaces must be analyzed through their mixed increments. Numerically, backward MFQDMA ties MFDMA on clean 1D benchmarks, is 22%–46% more accurate under 2%–10% contamination at low λ, and gives the most accurate $\tau(q)$ of 256^2 cascades. It does not cure heavy-tail artifacts or confer rank invariance. In daily returns of 18 financial assets the multifractality is mainly distributional. Low-λ MFQDMA is the estimator least affected by isolated bad ticks and by the removal of a few crash days, while MFDMA is preferable for dense microstructure noise. On EEG and terrain data, surrogate analysis shows that the multifractality is correlation-driven and largely linear. On the EEG, the λ-resolved exponents improve the classification of brain states beyond $h(q)$ at equal feature dimension. The cross-correlation extension is presented in the jointly submitted companion paper [42].

Acknowledgments

The author thanks R. G. Andrzejak and co-workers for making the Bonn EEG database publicly available.

* Data availability. The Bonn EEG database is publicly available from the University of Bonn (see Ref. [39]). The elevation model is distributed with Matplotlib [61]. The financial price series are from the `qrmdata` collection [58] (CRAN). The code, the generators of the benchmark corpus, and all numerical results underlying the tables and figures will be deposited in a public repository upon publication.

A Implementation of the 2D quantile trend

Let $P(i_1,i_2) = \sum_{j_1 \leq i_1, j_2 \leq i_2} X(j_1,j_2)$ be the integral image with $P(0,\cdot) = P(\cdot,0) = 0$. The box sum of Eq. (10) is $Y = P(i_1,i_2) - P(i_1-n,i_2) - P(i_1,i_2-n) + P(i_1-n,i_2-n)$, as in Appendix A of Ref. [22]. For a window whose upper-left corner is (r_1,r_2)

the within-window partial sums of Eq. (11) are

$$\begin{aligned}\widetilde{W}(m_1,m_2) &= P(r_1+m_1,r_2+m_2)-P(r_1,r_2+m_2)\\ &\quad -P(r_1+m_1,r_2)+P(r_1,r_2),\end{aligned} \tag{24}$$

so all n^2 partial sums of every window are obtained from $(n+1)\times(n+1)$ sliding views of P without recomputing sums. The λ-quantile of each window is obtained by a partial sort (selection) of the two order statistics needed in Eq. (2). Several λ values are computed in one selection pass. Windows are processed in row chunks to bound memory. The mean trend uses the separable ramp weights given in Sec. 2.3. With these choices the quantile and mean versions differ only in the trend functional, and the 2D-MFDMA implementation reproduces the ranking of window positions reported in Ref. [18].

B Generators and consistency checks

All analytical references were checked against the identities of Sec. 3.4 before use. For the 2D cascades, $\tau_{\mathrm{th}}(0) = -2$, $\tau_{\mathrm{th}}(1) = 0$, $h_{\mathrm{th}}(0) = \alpha_{\mathrm{th}}(0)$, and a numerical derivative of τ_{th} at $q = 0$ agree to 10^{-6} or better. For fGn, the empirical autocovariance of the circulant-embedding generator agrees with $\gamma(k) = \frac{1}{2}(|k+1|^{2H} - 2|k|^{2H} + |k-1|^{2H})$ to within 3×10^{-3}. Hermitian symmetry of the random spectrum is essential: taking the real part of an unsymmetrized complex sample raised this error to 6×10^{-2}.

References

[1] Thomas C. Halsey, Mogens H. Jensen, Leo P. Kadanoff, Itamar Procaccia, and Boris I. Shraiman. Fractal measures and their singularities: The characterization of strange sets. *Phys. Rev. A*, 33:1141, 1986.

[2] Charles Meneveau and Katepalli R. Sreenivasan. Simple multifractal cascade model for fully developed turbulence. *Phys. Rev. Lett.*, 59:1424, 1987.

[3] Benoit B. Mandelbrot. Intermittent turbulence in self-similar cascades: divergence of high moments and dimension of the carrier. *J. Fluid Mech.*, 62:331, 1974.

[4] Daniel Schertzer and Shaun Lovejoy. Physical modeling and analysis of rain and clouds by anisotropic scaling multiplicative processes. *J. Geophys. Res.*, 92:9693, 1987.

[5] Jean-François Muzy, Emmanuel Bacry, and Alain Arneodo. Wavelets and multifractal formalism for singular signals: Application to turbulence data. *Phys. Rev. Lett.*, 67:3515, 1991.

[6] Alain Arneodo, Nicolas Decoster, and Stéphane G. Roux. A wavelet-based method for multifractal image analysis. I. Methodology and test applications on isotropic and anisotropic random rough surfaces. *Eur. Phys. J. B*, 15:567, 2000.

[7] Stéphane Jaffard. Wavelet techniques in multifractal analysis. In *Fractal Geometry and Applications: A Jubilee of Benoît Mandelbrot*, volume 72 of *Proc. Symp. Pure Math.*, page 91. American Mathematical Society, Providence, RI, 2004.

[8] Herwig Wendt, Patrice Abry, and Stéphane Jaffard. Bootstrap for empirical multifractal analysis. *IEEE Signal Process. Mag.*, 24(4):38, 2007.

[9] Herwig Wendt, Stéphane G. Roux, Stéphane Jaffard, and Patrice Abry. Wavelet leaders and bootstrap for multifractal analysis of images. *Signal Process.*, 89:1100, 2009.

[10] C.-K. Peng, Sergey V. Buldyrev, Shlomo Havlin, Michael Simons, H. Eugene Stanley, and Ary L. Goldberger. Mosaic organization of DNA nucleotides. *Phys. Rev. E*, 49:1685, 1994.

[11] Jan W. Kantelhardt, Stephan A. Zschiegner, Eva Koscielny-Bunde, Shlomo Havlin, Armin Bunde, and H. Eugene Stanley. Multifractal detrended fluctuation analysis of nonstationary time series. *Physica A*, 316: 87, 2002.

[12] Espen A. F. Ihlen. Introduction to multifractal detrended fluctuation analysis in Matlab. *Front. Physiol.*, 3:141, 2012.

[13] Enzo Alessio, Anna Carbone, Gianni Castelli, and Vincenzo Frappietro. Second-order moving average and scaling of stochastic time series. *Eur. Phys. J. B*, 27:197, 2002.

[14] Anna Carbone, Gianni Castelli, and H. Eugene Stanley. Analysis of clusters formed by the moving average of a long-range correlated time series. *Phys. Rev. E*, 69: 026105, 2004.

[15] Limei Xu, Plamen Ch. Ivanov, Kun Hu, Zhi Chen, Anna Carbone, and H. Eugene Stanley. Quantifying signals with power-law correlations: A comparative study of detrended fluctuation analysis and detrended moving average techniques. *Phys. Rev. E*, 71:051101, 2005.

[16] Sergio Arianos and Anna Carbone. Detrending moving average algorithm: A closed-form approximation of the scaling law. *Physica A*, 382:9, 2007.

[17] Anna Carbone. Algorithm to estimate the Hurst exponent of high-dimensional fractals. *Phys. Rev. E*, 76:056703, 2007.

[18] Gao-Feng Gu and Wei-Xing Zhou. Detrending moving average algorithm for multifractals. *Phys. Rev. E*, 82: 011136, 2010.

[19] Gao-Feng Gu and Wei-Xing Zhou. Detrended fluctuation analysis for fractals and multifractals in higher dimensions. *Phys. Rev. E*, 74:061104, 2006.

[20] Boris Podobnik and H. Eugene Stanley. Detrended cross-correlation analysis: A new method for analyzing two nonstationary time series. *Phys. Rev. Lett.*, 100:084102, 2008.

[21] Wei-Xing Zhou. Multifractal detrended cross-correlation analysis for two nonstationary signals. *Phys. Rev. E*, 77: 066211, 2008.

[22] Zhi-Qiang Jiang and Wei-Xing Zhou. Multifractal detrending moving-average cross-correlation analysis. *Phys. Rev. E*, 84:016106, 2011.

[23] Paweł Oświęcimka, Stanisław Drożdż, Marcin Forczek, Stanisław Jadach, and Jarosław Kwapień. Detrended cross-correlation analysis consistently extended to multifractality. *Phys. Rev. E*, 89:023305, 2014.

[24] Xi-Yuan Qian, Ya-Min Liu, Zhi-Qiang Jiang, Boris Podobnik, Wei-Xing Zhou, and H. Eugene Stanley. Detrended partial cross-correlation analysis of two nonstationary time series influenced by common external forces. *Phys. Rev. E*, 91:062816, 2015.

[25] Zhi-Qiang Jiang, Wen-Jie Xie, Wei-Xing Zhou, and Didier Sornette. Multifractal analysis of financial markets: a review. *Rep. Prog. Phys.*, 82:125901, 2019.

[26] Peter J. Huber. Robust estimation of a location parameter. *Ann. Math. Stat.*, 35:73, 1964.

[27] Frank R. Hampel. The influence curve and its role in robust estimation. *J. Am. Stat. Assoc.*, 69:383, 1974.

[28] Frank R. Hampel, Elvezio M. Ronchetti, Peter J. Rousseeuw, and Werner A. Stahel. *Robust Statistics: The Approach Based on Influence Functions*. Wiley, New York, 1986.

[29] Peter J. Huber and Elvezio M. Ronchetti. *Robust Statistics*. Wiley, Hoboken, NJ, 2nd edition, 2009.

[30] Josef Ludescher, Mikhail I. Bogachev, Jan W. Kantelhardt, Aicko Y. Schumann, and Armin Bunde. On spurious and corrupted multifractality: The effects of additive noise, short-term memory and periodic trends. *Physica A*, 390:2480, 2011.

[31] Stanisław Drożdż, Jarosław Kwapień, Paweł Oświęcimka, and Rafał Rak. Quantitative features of multifractal subtleties in time series. *Europhys. Lett.*, 88:60003, 2009.

[32] Wei-Xing Zhou. The components of empirical multifractality in financial returns. *Europhys. Lett.*, 88:28004, 2009.

[33] Roger Koenker and Gilbert Bassett. Regression quantiles. *Econometrica*, 46:33, 1978.

[34] Roger Koenker and Kevin F. Hallock. Quantile regression. *J. Econ. Perspect.*, 15(4):143, 2001.

[35] Roger Koenker. *Quantile Regression*. Cambridge University Press, Cambridge, 2005.

[36] John W. Tukey. *Exploratory Data Analysis*. Addison-Wesley, Reading, MA, 1977.

[37] Neal C. Gallagher and Gary L. Wise. A theoretical analysis of the properties of median filters. *IEEE Trans. Acoust., Speech, Signal Process.*, 29:1136, 1981.

[38] Wolfgang Härdle and William Steiger. Optimal median smoothing. *J. R. Stat. Soc. Ser. C (Appl. Stat.)*, 44:258, 1995.

[39] Ralph G. Andrzejak, Klaus Lehnertz, Florian Mormann, Christoph Rieke, Peter David, and Christian E. Elger. Indications of nonlinear deterministic and finite-dimensional structures in time series of brain electrical activity: Dependence on recording region and brain state. *Phys. Rev. E*, 64:061907, 2001.

[40] James Theiler, Stephen Eubank, André Longtin, Bryan Galdrikian, and J. Doyne Farmer. Testing for nonlinearity in time series: the method of surrogate data. *Physica D*, 58:77, 1992.

[41] Thomas Schreiber and Andreas Schmitz. Surrogate time series. *Physica D*, 142:346, 2000.

[42] Sawon Pratiher. Multifractal quantile detrending moving-average cross-correlation analysis, 2026. companion paper, submitted jointly to Phys. Rev. E.

[43] Rob J. Hyndman and Yanan Fan. Sample quantiles in statistical packages. *Am. Stat.*, 50:361, 1996.

[44] Benoit B. Mandelbrot and John W. Van Ness. Fractional Brownian motions, fractional noises and applications. *SIAM Rev.*, 10:422, 1968.

[45] Robert B. Davies and David S. Harte. Tests for Hurst effect. *Biometrika*, 74:95, 1987.

[46] C. R. Dietrich and Garry N. Newsam. Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix. *SIAM J. Sci. Comput.*, 18:1088, 1997.

[47] John M. Chambers, Colin L. Mallows, and B. W. Stuck. A method for simulating stable random variables. *J. Am. Stat. Assoc.*, 71:340, 1976.

[48] Rosario N. Mantegna and H. Eugene Stanley. Scaling behaviour in the dynamics of an economic index. *Nature*, 376:46, 1995.

[49] Parameswaran Gopikrishnan, Vasiliki Plerou, Luís A. Nunes Amaral, Martin Meyer, and H. Eugene Stanley. Scaling of the distribution of fluctuations of financial market indices. *Phys. Rev. E*, 60:5305, 1999.

[50] Rama Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quant. Finance*, 1:223, 2001.

[51] Emmanuel Bacry, Jean Delour, and Jean-François Muzy. Multifractal random walk. *Phys. Rev. E*, 64:026103, 2001.

[52] Laurent Calvet and Adlai Fisher. Multifractality in asset returns: Theory and evidence. *Rev. Econ. Stat.*, 84:381, 2002.

[53] Kaushik Matia, Yosef Ashkenazy, and H. Eugene Stanley. Multifractal properties of price fluctuations of stocks and commodities. *Europhys. Lett.*, 61:422, 2003.

[54] Paweł Oświęcimka, Jarosław Kwapień, and Stanisław Drożdż. Multifractality in the stock market: price increments versus waiting times. *Physica A*, 347:626, 2005.

[55] Christian T. Brownlees and Giampiero M. Gallo. Financial econometric analysis at ultra-high frequency: Data handling concerns. *Comput. Stat. Data Anal.*, 51:2232, 2006.

[56] Ole E. Barndorff-Nielsen, Peter Reinhard Hansen, Asger Lunde, and Neil Shephard. Realized kernels in practice: trades and quotes. *Econom. J.*, 12:C1, 2009.

[57] Yacine Aït-Sahalia, Per A. Mykland, and Lan Zhang. How often to sample a continuous-time process in the presence of market microstructure noise. *Rev. Financ. Stud.*, 18:351, 2005.

[58] Marius Hofert, Kurt Hornik, and Alexander J. McNeil. qrmdata: Data sets for quantitative risk management practice, 2025. R package version 2025-07-24-3, https://CRAN.R-project.org/package=qrmdata.

[59] Alexander J. McNeil, Rüdiger Frey, and Paul Embrechts. *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press, Princeton, NJ, revised edition, 2015.

[60] Jozef Baruník, Tomaso Aste, T. Di Matteo, and Ruipeng Liu. Understanding the source of multifractality in financial markets. *Physica A*, 391:4234, 2012.

[61] John D. Hunter. Matplotlib: A 2D graphics environment. *Comput. Sci. Eng.*, 9:90, 2007.