Multifractal Analysis of Stock Price Volatility Using Wavelet Transform

Hideomi Totsuka *

1 Introduction

In recent years, numerous empirical studies have documented the roughness of volatility. Volatility, which measures the magnitude of price fluctuations, is a key statistic in financial markets. Using high-frequency equity-market data, Gatheral *et al.* (2018) reported roughness and monofractality[1] of realized volatility, with a Hurst exponent around $H \approx 0.1$ inferred from its scaling law. Fukasawa (2021) proposed an adaptive Whittle estimator based on the error distribution of realized volatility and the autocovariance of an fBm-driven stochastic volatility model, and showed that for the realized volatility of major stock indices $H \leq 0.06$, suggesting even stronger roughness than in Gatheral *et al.* In general, the Hurst exponent H, an index of roughness, is defined via scaling relations such as self-similarity $X_{\lambda t} \overset{d}{\approx} \lambda^H X_t$ and $\mathrm{Var}(X_{t+\Delta} - X_t) \propto \Delta^{2H}$.

However, a single exponent does not necessarily describe an entire time series. In multifractal signals, singularities of differing strengths coexist, and scaling varies with the moment order q. A statistical framework is therefore required to characterize such multifractality. Using MultiFractal Detrended Fluctuation Analysis (MF-DFA), Takaishi (2020) reported that Bitcoin volatility exhibits multifractality, with a generalized Hurst exponent $h(q) \leq 0.21$, and found a similar tendency for the S&P 500.

Beyond MF-DFA, there exist wavelet-based approaches to multifractal analysis[2]. Among these, the wavelet transform modulus maxima (WTMM) method emphasizes coefficients near singularities through mother wavelets with vanishing moments and, based on the scaling law $Z(q, s) \sim s^{\tau(q)}$ of the partition function $Z(q, s)$ constructed from modulus-maxima lines connected across scales, enables the estimation of the scaling function $\tau(q)$, the generalized Hurst exponent $h(q)$, and the singularity spectrum $f(\alpha)$[3]. This property affords relative robustness to nonstationary components and low-order trends, and, by allowing the localization of singularities, makes the method well suited to financial data.

The purpose of this study is to apply MF-WTMM to realized volatility estimated from high-frequency data on Japanese and U.S. stock indices, and to volatility indices, in order to clarify and comparatively assess roughness

* College of Economics, Nihon University, e-mail: totsuka.hideomi@nihon-u.ac.jp

1) A fractal is a structure exhibiting self-similarity; monofractality refers to self-similarity characterized by a single singularity. For an overview, see Mandelbrot (1998).

2) See, for example, Chhabra and Jensen (1989) and Muzy *et al.* (1993).

3) In this paper, we refer to this approach as multifractal WTMM (MF-WTMM).

and multifractality. Empirically, we analyze the log-differences of the realized volatility of the Nikkei 225 and the S&P 500, as well as the log-differences of closing prices for the Nikkei 225 VI and the CBOE VIX, comparing the q-dependence of the generalized Hurst exponent and the width and peak location of the singularity spectrum $f(\alpha)$. Our main findings are as follows. First, for the realized volatility of stock indices, the width of $f(\alpha)$ is relatively narrow, indicating behavior close to monofractal scaling. Second, for volatility indices, the width of $f(\alpha)$ is broad with a clear q-dependence, revealing pronounced multifractality. In particular, for the CBOE VIX, the Hölder exponent α extends markedly into the negative range, suggesting that sharp, spike-like fluctuations are dominant.

The remainder of the paper is organized as follows. Section 2 describes the multifractal methods employed in the analysis. Section 3 presents the data — two realized-volatility series for stock indices and two volatility indices — and the empirical results. Section 4 discusses the findings in light of the evidence. Section 5 concludes and outlines directions for future research.

2 Multifractal Analysis

2.1 Continuous Wavelet Transform

For a square-integrable function $f(t) \in L^2(\mathbb{R})$, the continuous wavelet transform (CWT) is defined by

$$W_\psi[f](u,s) = \int_{-\infty}^{\infty} f(t)\psi^*_{u,s}(t)dt. \tag{2.1}$$

Here,

$$\psi_{u,s}(t) = \frac{1}{\sqrt{s}}\psi\left(\frac{t-u}{s}\right), \tag{2.2}$$

and $\psi \in L^2(\mathbb{R})$ is called the mother wavelet. With this normalization, $\|\psi_{u,s}\|_2 = \|\psi\|_2$ holds[4]. Note that some references use L^1 normalization (with factor $1/s$). In that case, the L^1 norm (area) of the wavelet is preserved, which is convenient for probabilistic/measure-theoretic interpretations; in this article we adopt the standard L^2 normalization.

The scale s dilates or compresses the time axis by a factor of s ($t \mapsto (t-u)/s$) and, to preserve energy under this dilation, the amplitude is normalized by $(1/\sqrt{s}$. The effective (central) frequency then varies approximately in proportion to $1/s$. The shift u translates the time location of the wavelet by u.

2.2 Mother Wavelet

The mother wavelet serves as the basis function of the continuous wavelet transform (CWT), and its properties — shape, localization, and number of vanishing moments — directly affect estimation accuracy and detection sensitivity. This subsection outlines the definition of a mother wavelet and the design requirements.

A mother wavelet is a basis function used in the CWT to expand a time series signal across multiple scales. It is temporally localized, has zero mean, and possesses a prescribed number of vanishing moments, thereby enabling efficient extraction of local variations and singularities in the signal. In general, a mother wavelet $\psi(t)$ is defined to satisfy the following conditions:

4） The L^2 norm is defined by

$$\|g\|2 = \left(\int_{-\infty}^{\infty} |g(t)|^2 dt\right)^{1/2}$$

and, in signal analysis, corresponds to the total energy of the signal.

1. **Zero-mean property**

$$\int_{-\infty}^{\infty} \psi(t)dt = 0. \tag{2.3}$$

This eliminates sensitivity to low-frequency components and emphasizes local fluctuations.

2. **Regularity and localization**

$\psi(t)$ should be localized in time and sufficiently smooth. These properties stabilize the transform coefficients across scales and allow accurate estimation of the location and strength of singularities.

3. **Vanishing moments**

The order n of vanishing moments is the largest integer such that

$$\int_{-\infty}^{\infty} t^k \psi(t)dt = 0, \quad (k = 0, 1, \ldots, n-1). \tag{2.4}$$

A larger n removes polynomial trends of higher degree and increases sensitivity to sharp singularities.

4. **Admissibility condition**

For the CWT to be invertible, the following must hold in the frequency domain:

$$C_\psi = \int_0^{\infty} \frac{\left|\hat{\psi}(\omega)\right|^2}{\omega} d\omega < \infty, \tag{2.5}$$

where $\hat{\psi}(\omega)$ denotes the Fourier transform of $\psi(t)$. This condition implies the zero-mean property and guarantees suppression of low-frequency components.

The choice of mother wavelet depends on the analytical objective. For example, the Morlet wavelet combines a sinusoidal carrier with a Gaussian window, offering high frequency resolution and suitability for oscillatory signals. By contrast, the Mexican hat (Ricker) wavelet is the second derivative of a Gaussian and is effective for detecting singularities and peaks. Because the selection directly influences accuracy and sensitivity, it should be made with careful consideration of the signal characteristics and the goals of the analysis.

Mother wavelets are broadly classified into *real-valued* and *complex-valued* types. Real-valued wavelets are primarily used to extract local singularities, whereas complex-valued wavelets provide both amplitude and phase information and are useful for time–frequency analysis. Below, we present representative examples of each class.

2.2.1 Real-Valued Mother Wavelets

- **Mexican hat**

The Mexican hat (also known as the Ricker wavelet) is obtained as the second derivative of a Gaussian function. Its shape features a negative trough at the center surrounded by positive sidelobes — resembling a sombrero — hence the name. It is given by

$$\psi(t) = -\sigma^2 \frac{d^2}{dt^2}\left(e^{-\frac{t^2}{2\sigma^2}}\right) = \left(1 - \frac{t^2}{\sigma^2}\right) e^{-\frac{t^2}{2\sigma^2}}. \tag{2.6}$$

The Mexican hat has two vanishing moments and thus removes constant and linear trends. It is well suited to detecting abrupt changes and singularities in signals and is frequently employed in multifractal analysis.

- **Derivative-of-Gaussian (DOG) wavelets**

  While the Mexican hat corresponds to the second derivative of a Gaussian, a broader family of wavelets is obtained by taking the *n*th derivative of a Gaussian, known as *DOG (Derivative of Gaussian)* wavelets. The general form is

$$\psi^{(n)}(t) = \frac{d^n}{dt^n}\left(e^{-\frac{t^2}{2}}\right). \tag{2.7}$$

  Larger n yields higher-order vanishing moments, enabling the removal of polynomial components of higher degree. In particular, $n = 1$ is effective for edge detection, $n = 2$ for singularity detection, and $n \geq 3$ for separating smoother components. In multifractal analysis, DOG wavelets play a foundational role in probing scaling behavior.

- **Real Morlet wavelet**

  The Morlet wavelet is the product of a complex sinusoid and a Gaussian window; when only its real part is used, it is referred to as the *real Morlet wavelet*. Its general form is

$$\psi(t) = \Re\left\{\frac{1}{\pi^{1/4}} e^{i\omega t} e^{-\frac{t^2}{2}}\right\} = \frac{1}{\pi^{1/4}} \cos(\omega t) e^{-\frac{t^2}{2}}. \tag{2.8}$$

  Because the Gaussian window provides temporal localization and the cosine term supplies periodic structure, the real Morlet is effective for emphasizing periodicity and oscillatory components. It is therefore used when oscillatory features must be captured or frequency selectivity is required.

2.2.2 Complex-Valued Mother Wavelets

- **Complex Morlet wavelet**

  The complex Morlet wavelet treats the Morlet form explicitly as a complex exponential, i.e., a complex sinusoid modulated by a Gaussian envelope. A commonly used form is

$$\psi(t) = \frac{1}{\pi^{1/4}}\left(e^{i\omega_0 t} - e^{-\omega_0^2/2}\right)\exp\left(-\frac{t^2}{2}\right), \tag{2.9}$$

  where ω_0 denotes the carrier frequency. Having both real and imaginary parts, the complex Morlet simultaneously captures amplitude and phase, exhibiting excellent localization for time–frequency analysis of oscillatory signals. It is one of the most widely used complex wavelets.

- **Cauchy wavelet**

  The Cauchy wavelet is based on a rational-analytic function on the complex plane and can be written in the time domain as

$$\psi(t) = \frac{\Gamma(p+1)}{2\pi}\left(\frac{i}{t+i}\right)^{p+1}, \qquad p > 0, \tag{2.10}$$

where p is an integer corresponding to the order of vanishing moments. The Cauchy wavelet is analytic and, in the frequency domain, is one-sided (localized on positive frequencies), thus emphasizing phase information. It is particularly effective for edge detection and for analyzing asymmetric singularities, and is a standard example alongside the Morlet among complex-valued wavelets.

2.3 Vanishing Moments and Singularity Detection

The number of vanishing moments is a key factor that determines how effectively a wavelet suppresses smooth components of a signal and how sensitively it detects singularities. This subsection presents the mathematical relationship between vanishing moments and singularity detection.

As discussed in Section 2.2, one of the defining conditions of a mother wavelet is the presence of *vanishing moments*, which, in a wavelet transform, remove smooth components and accentuate local discontinuities and singular structures. When vanishing moments are present, low–order polynomial terms are annihilated, thereby attenuating global trends and smooth variations and enhancing the contribution of singularities.

In particular, if a function $f(t)$ has Hölder exponent α at a point t_0, and if the number of vanishing moments n of the wavelet exceeds α, then the CWT coefficients $W_\psi[f](u,s)$ obey the scaling law

$$W_\psi[f](u,s) \sim s^{\alpha+1/2} \quad (s \to 0), \tag{2.11}$$

which enables the estimation of α. Conversely, if the number of vanishing moments is too small, smooth components are insufficiently removed and singularity detection becomes difficult. On the other hand, increasing the number of vanishing moments beyond what is necessary weakens localization and degrades time resolution. Therefore, an appropriate choice of the number of vanishing moments — balancing suppression of smooth components against localization — must be made in accordance with the characteristics of the signal under analysis.

2.4 Wavelet Transform Modulus Maxima method

2.4.1 Analysis of Singularities

When a signal $f(t)$ has a singularity (a discontinuity or nondifferentiability) at $t = t_0$, the continuous wavelet transform coefficients $W_\psi[f](u,s)$ attain large values in a neighborhood of t_0, and their absolute values often form local maxima. This occurs because the mother wavelet ψ, by virtue of its vanishing moments, cancels smooth components and thereby accentuates singular behavior. Owing to this property, one can quantitatively extract scale–dependent structures and fractal characteristics of the signal. The WTMM method exploits this by tracking such maxima along the scale direction and analyzing their power–law behavior, with the aim of systematically characterizing the distribution of singularities present in the signal.

2.4.2 On Modulus Maxima of Wavelet Coefficients

The wavelet transform modulus maxima (WTMM) serve as key indicators reflecting singularities of a signal. Let $W_\psi[f](u,s)$ denote the continuous wavelet transform of a signal $f(t)$. Define

$$M(u,s) = \left|W_\psi[f](u,s)\right|. \tag{2.12}$$

A WTMM at location $u = u_0$ is defined as a local maximum of $M(u, s)$ with respect to u, i.e., a point satisfying

- it is a stationary point,

$$\left.\frac{\partial M(u,s)}{\partial u}\right|_{u=u_0} = 0, \tag{2.13}$$

- and it is a strict local maximum,

$$\left.\frac{\partial^2 M(u,s)}{\partial u^2}\right|_{u=u_0} < 0. \tag{2.14}$$

Points meeting these conditions reflect the locations of local abrupt changes or singularities and thus play a role in emphasizing singular behavior. Moreover, by tracking the chain of modulus maxima as the scale s varies, one can systematically assess the strength and persistence of singularities.

2.4.3 Maxima Lines of Wavelet Coefficients

Modulus maxima of the wavelet transform identified at a single scale can be linked across scales to form *maxima lines*. A maxima line is a chain of modulus–maximum points that starts from a maximum at some scale s_0 and is continuously tracked toward smaller scales; it can be interpreted as a curve corresponding to the location of a singularity in the signal. For finite-length data, maxima lines are typically confined to a finite range of scales.

Let l denote a maxima line of the CWT coefficients $W_\psi[f](u, s)$ at the reference scale s_0. Ideally, l satisfies the following:

- If a point (u, s) lies on l, then $s \le s_0$ and (u, s) is a modulus maximum of the wavelet transform coefficients.

- For every scale $s \le s_0$ there exists a point on l at that scale (in practice, this holds approximately on a discretized set of scales).

2.5 Partition Function and Scaling Exponent

In MF-WTMM, a scale-dependent partition function is constructed from modulus maxima, and the multifractal spectrum is derived from its scaling law[5)]. This section organizes, in a systematic manner, the development from the definition of the partition function to the estimation of scaling exponents.

2.5.1 Why Introduce a Partition Function?

In local analyses based on WTMM lines, one can estimate the local Hölder exponent $\alpha(t_0)$ associated with each singularity. However, an entire signal is not necessarily characterized by a single exponent. In multifractal signals, singularities of differing strengths coexist, and a statistical framework is required to describe them collectively.

To this end, introducing the moment order q, the partition function is defined by

$$Z(q, s) = \sum_{l \in \mathcal{L}(s)} \left(\sup_{(u,s') \in l, s' \le s} \left| W_\psi[f](u, s') \right| \right)^q . \tag{2.15}$$

5）For details, see Bacry *et al.* (1993) and Jaffard (1997).

Here, $\mathcal{L}(s)$ denotes the set of WTMM lines that are detected outside the cone of influence (COI)[6] and that, being linked from finer to coarser scales, reach scale s. Positive q emphasizes large moduli (strong singularities), whereas negative q emphasizes small moduli (weak singularities). Over an appropriate scaling range,

$$Z(q, s) \sim s^{\tau(q)} \tag{2.16}$$

holds.

2.5.2 Definition and Scaling Law of $Z(q, s)$

From the preceding discussion, the modulus satisfies $\left|W_\psi[f](u, s)\right| \leq s^{\alpha+1/2}$. Moreover, the number $N_\alpha(s)$ of maxima at scale s associated with Hölder exponent α admits the estimate

$$N_\alpha(s) \leq s^{-f(\alpha)}, \tag{2.17}$$

where $f(\alpha)$ is the singularity spectrum, which quantifies — via fractal dimension — the abundance of points having local regularity α. Consequently, the scaling law of the partition function defined in (2.15) is

$$Z(q, s) \leq \sum_\alpha N_\alpha(s) s^{q(\alpha+\frac{1}{2})} \approx \int s^{q(\alpha+\frac{1}{2})-f(\alpha)} d\alpha. \tag{2.18}$$

In the limit $s \downarrow 0$, the smallest exponent governs the behavior, yielding

$$Z(q, s) \leq s^{\tau(q)}, \qquad \tau(q) = \inf_\alpha \left\{ q\left(\alpha + \frac{1}{2}\right) - f(\alpha) \right\}. \tag{2.19}$$

Here, because we adopted the L^2 normalization $1/\sqrt{s}$, the exponent $\tau(q)$ is shifted by $q/2$. In general, one therefore sets

$$\tilde{Z}(q, s) = s^{-q/2} Z(q, s), \tag{2.20}$$

so that

$$\tilde{Z}(q, s) \leq s^{\tilde{\tau}(q)}, \qquad \tilde{\tau}(q) = \inf_\alpha \{q\alpha - f(\alpha)\}. \tag{2.21}$$

In this form, the Legendre transform holds in its standard form,

$$f(\alpha) = \inf_q \{q\alpha - \tilde{\tau}(q)\}. \tag{2.22}$$

2.6 Generalized Hurst Exponent and Singularity Spectrum

2.6.1 Generalized Hurst Exponent

The generalized Hurst exponent $h(q)$ is a statistic that quantifies the extent to which a time series or signal exhibits long–range dependence (long memory) or persistence.

- $h < 0.5$: indicates anti-persistence; increases are likely to be followed by decreases, implying a strong short-term reversal tendency.

6) When the transform is computed over a finite interval $[0, T]$, parts of the wavelet extend beyond the data near the endpoints $t = 0$ and $t = T$, leading to inaccurate coefficients. The region affected by this endpoint influence — determined by s and $\psi(t)$ — is called the cone of influence.

- $h = 0.5$: corresponds to Brownian motion or i.i.d. white noise, with no long-range correlation.

- $0.5 < h \leq 1$: indicates persistence; past increases (decreases) tend to induce future increases (decreases), reflecting a long-memory structure.

Many real-world data, including those from natural phenomena and financial markets, are too complex to be characterized by a single h. In multifractal signals that contain locally varying scaling behavior and singularities, fluctuations are nonuniform, and the scaling behavior changes with the moment order q. To capture this complexity, the generalized Hurst exponent $h(q)$ is introduced.

The generalized Hurst exponent is defined in terms of the q-dependent scaling exponent $\tau(q)$ as

$$h(q) = \frac{\tilde{\tau}(q) - \tilde{\tau}(0)}{q}, \tag{2.23}$$

where $\tilde{\tau}(q)$ is the scaling exponent estimated from (2.21) and serves as a fundamental descriptor of multifractality. With this definition, $h(q)$ has the following features:

- For $q > 0$: it places greater weight on large-amplitude (high-intensity) fluctuations, emphasizing the scaling of strongly varying regions.

- For $q < 0$: it is sensitive to small-amplitude (low-intensity) fluctuations, emphasizing the scaling of quiescent regions.

By continuity at $q = 0$, we define the generalized Hurst exponent as $h(0) := \lim_{q\to 0}(\tilde{\tau}(q) - \tilde{\tau}(0))/q = \tilde{\tau}'(0) = \alpha(q)|_{q=0}$, which is distinct from the peak location of the singularity spectrum α_{peak}.

Accordingly, the generalized Hurst exponent $h(q)$ serves to distinguish and evaluate the self-similarity of large versus small fluctuations in a signal. If $h(q)$ is independent of q, the signal is monofractal; if $h(q)$ varies with q, the signal exhibits multifractal characteristics. Thus, via the scaling exponent $\tau(q)$, the generalized Hurst exponent provides a core quantitative measure of a signal's diverse self-similarities and is indispensable for understanding multifractality that cannot be captured by a single Hurst exponent.

2.6.2 Singularity Spectrum

The singularity spectrum is a key measure in multifractal analysis that describes the statistical distribution of local singularities (discontinuities and sharp variations) in a signal. For signals exhibiting multifractality, the overall behavior cannot be captured by a single power–law exponent; instead, the degree of local irregularity varies from place to place. To quantify this local roughness, the Hölder exponent α is introduced. In practice, however, direct estimation of the Hölder exponent α is difficult, so one first estimates the scaling exponent $\tau(q)$ from the partition function. Using the resulting scaling exponent $\tilde{\tau}(q)$, the singularity spectrum $f(\alpha)$ is obtained via the Legendre transform,

$$\alpha(q) = \frac{d\tilde{\tau}(q)}{dq}, \quad f(\alpha) = q\alpha - \tilde{\tau}(q). \tag{2.24}$$

Thus, the singularity spectrum $f(\alpha)$ provides the correspondence between the strength of local singularities, measured by α, and the fractal dimension of the set on which they occur. It plays a central role in visualizing the multifaceted scaling behavior and complexity of a signal. A broad $f(\alpha)$ indicates strong multifractality, whereas a narrow spectrum implies behavior close to monofractal (single–scaling) dynamics.

3 Data and Empirical Results

3.1 Data

This study uses the realized volatilities of the Nikkei 225 and the S&P 500 obtained from the Oxford-Man Institute's Realized Library[7] and the closing values of the Nikkei 225 VI and the CBOE VIX obtained from Bloomberg. Realized volatility is constructed as the square root of the sum of squared high-frequency returns and used as an estimator of volatility[8].

The sample period for the realized volatilities of the Nikkei 225 and the S&P 500 and for the Nikkei 225 VI and the CBOE VIX spans from January 4, 2002 to December 30, 2021. Following prior studies[9], the empirical analysis is conducted on the log-differenced series of realized volatility and of the volatility indices, $LV_t = \log \sigma_t - \log \sigma_{t-1}$, where σ_t denotes realized volatility or the volatility index at time $t = 1, ..., T$.

3.2 Empirical Analysis

Figure 1 presents the generalized Hurst exponents and the singularity spectra for the Nikkei 225, the S&P 500, the Nikkei 225 VI, and the CBOE VIX, estimated via MF-WTMM.

Table 1 reports the maximum value $f_{\max}$ of the singularity spectrum together with the corresponding Hölder exponent α and the minimum and maximum values of α.

We next discuss the empirical findings for the Nikkei 225, the S&P 500, the Nikkei 225 VI, and the CBOE VIX.

3.2.1 Empirical Findings for the Nikkei 225

From the generalized Hurst exponent in Figure 1(a), the following observations emerge. The region $q < 0$ emphasizes weak fluctuations and small-amplitude variations. On the negative-q side, $h(q)$ takes relatively higher values, suggesting that small-amplitude fluctuations exhibit weaker mean reversion (or longer-range dependence) than large-amplitude fluctuations. By contrast, in the $q > 0$ range, $h(q)$ declines, indicating stronger mean reversion in large volatility movements.As a function of q, $h(q)$ decreases and exhibits a concave (downward) shape typical of multifractal structure. This implies that segments with different scaling properties coexist within the series, reflecting complex variation that cannot be summarized by a single Hurst exponent.

A pronounced q-dependence of $h(q)$ suggests that market volatility is not governed by a single scaling law; rather, large and small fluctuations coexist with distinct temporal correlations. Given that these results are based on high-frequency data, they may also capture effects of market microstructure noise and volatility clustering. Overall, the log-differenced realized volatility of the Nikkei 225 exhibits clear multifractality: scaling differs

7） See Heber et al. (2009).
8） See McAleer and Medeiros (2008).
9） See Gatheral *et al.* (2018) and Takaishi (2020).

Figure 1: Generalized Hurst exponents and singularity spectra:
(a), (b) Nikkei 225; (c), (d) S&P 500; (e), (f) Nikkei 225 VI; (g), (h) CBOE VIX.

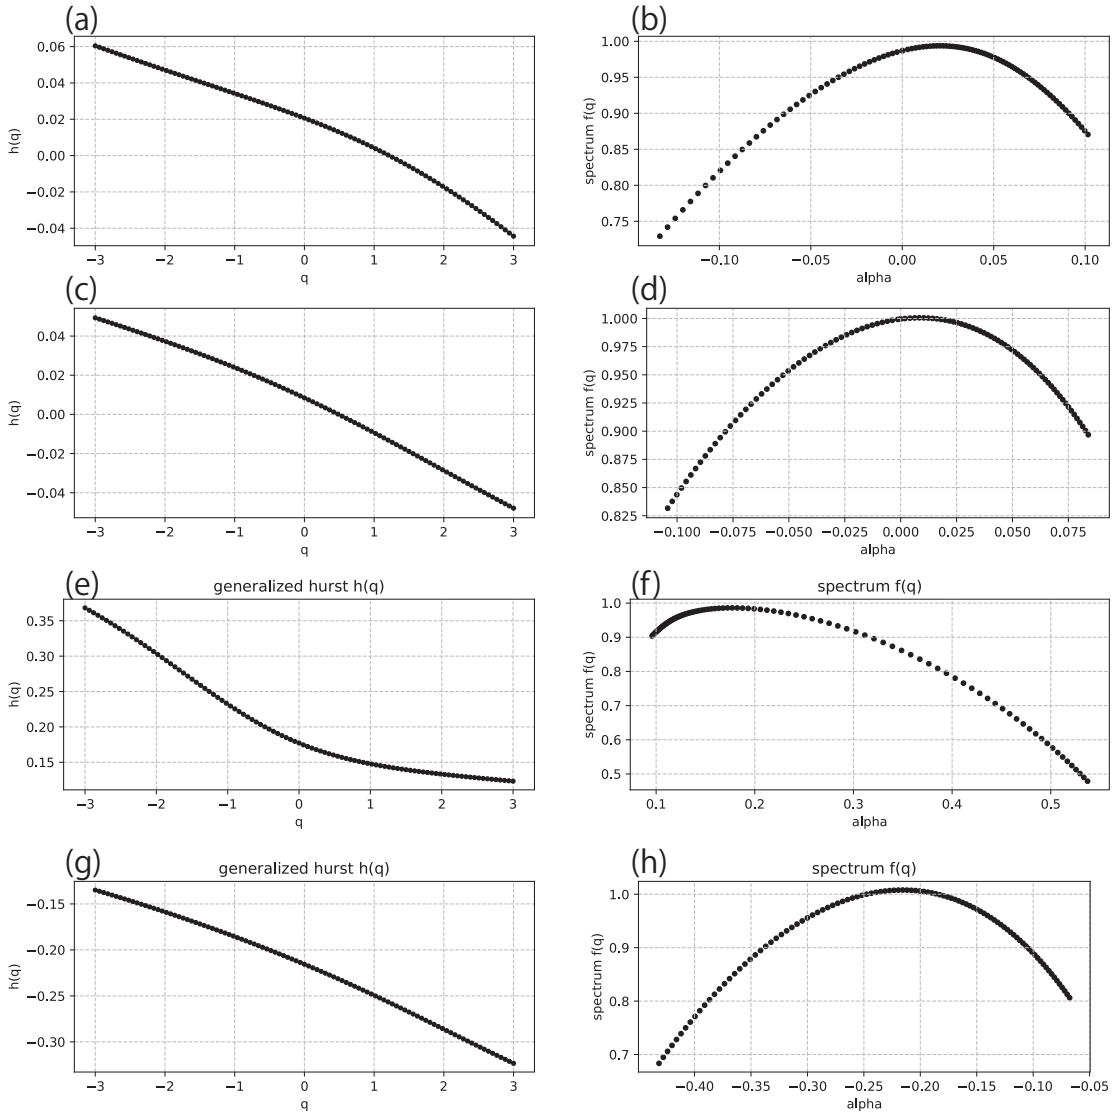

Table 1: Summary statistics of the singularity spectrum $f(\alpha)$

	Maximum of $f(\alpha)$	at maximum	Minimum α	Maximum α
Nikkei 225	0.9935	0.0205	−0.1325	0.1015
S&P 500	1.0001	0.0083	−0.1042	0.0841
Nikkei 225 VI	0.9860	0.1770	0.0959	0.5372
CBOE VIX	1.0078	−0.2156	−0.4315	−0.0677

systematically between the small-amplitude regime ($q < 0$) and the large-amplitude regime ($q > 0$), implying that a single self-similar process is insufficient to describe the data.

Turning to the singularity spectrum in Figure 1(b) and Table 1, the maximum value is $f_{\max} = 0.994$, indicating that the fractal dimension of the dominant set of singularities is close to 1, i.e., there exists a component that occupies the time axis almost uniformly. This implies that the support of the dominant singularities is effectively one-dimensional (i.e., close to dimension 1). The peak location at $\alpha = 0.021$ shows that the average Hölder regularity is low ($\alpha \approx 0.02$), indicating highly irregular, near–white-noise-like behavior.

The Hölder range $[-0.133, 0.102]$ is very narrow, implying weak multifractality and a structure close to single scaling; the inclusion of negative values at the edges nevertheless reflects the presence of locally sharp changes and spike-like behavior. The spectrum is sharply peaked with short, narrow tails on both sides — a slender inverted-U shape — characteristic of a series largely governed by a single scaling law. Although the distribution of α is narrow, indicating only modest variation in local scaling, minor departures from monofractality are detectable, while complex multifractal structure is not dominant.

3.2.2 Empirical Findings for the S&P 500

From the generalized Hurst exponent in Figure 1(c), the following observations emerge. The fact that the curve $h(q)$ has a nonzero slope with respect to q indicates a multifractal structure not captured by a single scaling exponent. Negative q is sensitive to small fluctuations (quiescent regimes); a gentle rise to the left suggests that minute variations may possess longer-term correlation. Positive q is sensitive to large fluctuations (spiky volatility); a slight decline to the right implies that large movements exhibit low persistence (anti-persistent behavior). Overall, the figure provides visual and quantitative evidence that the log-differenced realized volatility of the S&P 500 displays a clear multifractal structure.

Turning to the singularity spectrum in Figure 1(d) and Table 1, the peak value is essentially 1, close to the topological dimension (dimension 1) of an ideal multifractal set, indicating that the dominant set of singularities is effectively one-dimensional. The Hölder exponent at the maximum is very close to 0, implying that the prevailing local regularity is not very smooth and is characterized by nearly neutral scaling behavior. The Hölder range $[-0.1042, 0.0839]$ is extremely narrow, suggesting limited diversity in local scaling exponents and hence weak multifractality. In sum, while the log-differenced realized volatility of the S&P 500 behaves close to monofractal dynamics, it retains slight multifractal features.

3.2.3 Empirical Findings for the Nikkei 225 VI

From the generalized Hurst exponent in Figure 1(e), the following observations emerge. As the moment order q increases, $h(q)$ decreases monotonically, reflecting a concave (downward) scaling function $\tilde{\tau}(q)$ and indicating a clear multifractal structure in the log-differenced Nikkei 225 VI. The values of $h(q)$ lie roughly in the range 0.15–0.35, centered well below Brownian motion ($h = 0.5$), implying that anti-persistence dominates and price variations tend to revert over short horizons. For $q < 0$, $h(q)$ is somewhat higher, whereas for $q > 0$ it declines, suggesting that small fluctuations exhibit relatively longer-range correlation, while large fluctuations tend to self-reverse over short periods.

From the singularity spectrum in Figure 1(f) and Table 1, we find the following. The peak value $f_{\max} = 0.986$ is very close to 1, indicating that the dominant set of singularities is effectively one-dimensional. The Hölder

exponent at the maximum is $\alpha \approx 0.177$, which represents the typical local roughness; on average, the log-differenced series follows this scaling index. The Hölder range is $[0.0959, 0.5372]$, with a width of about 0.44, evidencing the coexistence of segments with distinct scaling properties and thus pronounced multifractality. The spectrum exhibits a concave profile with a slight rightward extension, suggesting that contributions from small fluctuations (the larger-α side) are relatively abundant compared with those from large fluctuations (the smaller-α side). The peak near $\alpha \sim 0.18$ captures the characteristic local roughness, around which singularities of varying strength are distributed.

3.2.4 Empirical Findings for the CBOE VIX

From the generalized Hurst exponent in Figure 1(g), we observe a clear dependence of $h(q)$ on the moment order q, which is consistent with multifractality. The curve decreases as q moves to positive values, indicating that the scaling behavior of large-amplitude fluctuations differs from that of small-amplitude fluctuations. Overall, the pronounced q-dependence of $h(q)$ provides evidence that the CBOE VIX cannot be characterized by a single scaling exponent and exhibits strong intermittency across fluctuation regimes.

From the singularity spectrum in Figure 1(h) and Table 1, we find the following. The peak value $f_{\max} = 1.008$ indicates that the dominant fractal set has dimension close to 1. The peak occurs at a Hölder exponent of approximately $\alpha \approx -0.216$; the negativity of α suggests that sharp, spike-like changes and steep variations are prevalent in the series. The spectrum spans the interval $[-0.431, -0.068]$, revealing a relatively broad multifractal structure: the wider the support, the richer the mixture of singularities (segments with differing roughness) contained in the time series. The graph exhibits a characteristic concave, roughly parabolic profile, pointing to self-similarity and asymmetry typical of multifractal systems.

4 Discussion

First, we compare the generalized Hurst exponents of the Nikkei 225, the S&P 500, the Nikkei 225 VI, and the CBOE VIX. There are three commonalities across the four series. In all cases, the generalized Hurst exponent $h(q)$ varies with q, confirming multifractality — i.e., the coexistence of distinct local scalings that cannot be captured by a single Hurst exponent. This points to features typical of financial markets and volatility indicators, such as nonlinear dependence and heteroskedasticity. For the Nikkei 225, the S&P 500, and the CBOE VIX, $h(q)$ lies well below 0.5 over wide ranges, indicating strong mean-reverting (anti-persistent) tendencies in the scaling sense. This tendency is especially pronounced for the volatility indices (the Nikkei 225 VI and the CBOE VIX), where abrupt, shock-type fluctuations dominate the statistical properties. The curve $h(q)$ is asymmetric between $q < 0$ and $q > 0$, indicating that different scaling laws govern small versus large fluctuations. Comparing stock indices with volatility indices yields the following. For the stock indices, the central tendency of $h(q)$ is well below 0.5, indicating anti-persistence rather than Brownian behavior, although multifractality is present, their fluctuations are relatively homogeneous and moderate. By contrast, for the volatility indices, the peak location α-peak shifts to either side of zero, indicating differences in typical local regularity (smoother segments vs spike-like fluctuations), and the broader dispersion of $h(q)$ suggests stronger multifractality than in the stock indices. Comparing Japanese and U.S. indicators, we find the following. For the Nikkei 225 and the S&P 500, the center of α_{peak} lies near zero in both cases, but the Nikkei 225 exhibits a somewhat broader spread, suggesting a greater influence of short-term

abrupt movements. For the singularity spectra, the peak Hölder exponent differs in sign: the Nikkei 225 VI peaks at $\alpha_{peak} \approx +0.18$ whereas the CBOE VIX peaks at $\alpha_{peak} \approx -0.22$. This likely reflects differences in market structure and suggests that the U.S. VIX experiences more frequent spikes and stronger local irregularity at short horizons. Taken together, these results indicate that the realized volatilities of the Nikkei 225 and the S&P 500 are comparatively homogeneous and stable, whereas the Nikkei 225 VI and the CBOE VIX, which are sensitive to market anxiety, exhibit pronounced multifractality.

We next compare the singularity spectra of the Nikkei 225, the S&P 500, the Nikkei 225 VI, and the CBOE VIX. There are three common features across the four spectra. First, in all series the singularity spectrum $f(\alpha)$ has a finite width rather than a single sharp peak, indicating the presence of multifractal structure that cannot be described by a single scaling law. Second, each spectrum exhibits a concave, approximately parabolic shape, which is consistent with a canonical multifractal distribution and suggests that the MF-WTMM procedure and the underlying scaling assumptions are theoretically coherent. Third, the widths extending to the left and right of the peak are asymmetric, reflecting distinct scaling properties for small versus large fluctuations — an attribute consistent with leverage effects and volatility clustering in financial markets. Comparing stock indices with volatility indices yields the following. For the realized volatilities of the Nikkei 225 and the S&P 500, the spectra are narrow and peak near $\alpha \approx 0$, indicating dynamics close to single-scaling behavior (near-monofractality) with limited dispersion in local regularity. By contrast, for the Nikkei 225 VI and the CBOE VIX the spectra are wide, revealing pronounced multifractality — consistent with the view that heterogeneous fluctuations and abrupt market shocks dominate the statistical properties of volatility indices. Comparing Japanese and U.S. indicators also reveals differences. The admissible range of the Hölder exponent is wider for the Nikkei 225 than for the S&P 500, suggesting that the Japanese market may be more exposed to short-run abrupt movements. Moreover, the Nikkei 225 VI spectrum peaks on the positive side ($\alpha \approx 0.18$), suggesting comparatively smooth local behavior, whereas the CBOE VIX spectrum peaks on the negative side ($\alpha \approx -0.22$), indicating spike-like fluctuations and stronger local irregularity. Taken together, these results show that the realized volatilities of the Nikkei 225 and the S&P 500 are relatively homogeneous and stable, while the Nikkei 225 VI and the CBOE VIX exhibit strong multifractality and pronounced sensitivity to market sentiment.

5 Conclusion and Future Directions

Using the MF-WTMM method, this study has shown that the realized volatilities of Japanese and U.S. stock indices are, on the whole, narrow multifractals[10], whereas the corresponding volatility indices exhibit wide multifractality. The generalized Hurst exponent $h(q)$ varies with q for all four series, indicating the coexistence of distinct self-similarities between the large-amplitude regime ($q > 0$) and the small-amplitude regime ($q < 0$). The peaks of the singularity spectra $f(\alpha)$ are generally close to 1[11], implying the presence of a dominant set broadly distributed along the time axis with (approximately) one-dimensional support, while their widths differ markedly across series, which maps neatly into differences in multifractality strength.

As a next step, it is necessary to consider multifractal analysis based on the Wavelet Leaders (WL) method using

10）They may in fact be close to single scaling; this warrants further investigation.

11）For the CBOE VIX, the estimate slightly exceeds 1, but this lies within numerical error.

the discrete wavelet transform. Since the WL method is often reported to achieve higher accuracy in detecting singularities than MF-WTMM, we expect it to enable a more detailed characterization of the multifractality of realized volatilities and volatility indices.

References

［1］ Bacry, E., Muzy, J. -F., and Arnéodo, A., (1993). Singularity spectrum of fractal signals from wavelet analysis: Exact results, *Journal of Statistical Physics*, **70**(3-4), 635–674.

［2］ Bennedsen, M., Lunde, A., and Pakkanen, M., S., (2022). Decoupling the Short- and Long-Term Behavior of Stochastic Volatility, *Journal of Financial Econometrics*, **20**(5), 961–1006.

［3］ Chhabra, A., and Jensen, R., V., (1989). Direct determination of the f(α) singularity spectrum, *Physical Review Letters*, **62**(12), 1327–1330.

［4］ Fukasawa, M., (2021). Volatility has to be rough, *Quantitative Finance*, **21**(1), 1–8.

［5］ Fukasawa, M., Takabatake, T., and Westphal, R., (2019). Is Volatility Rough?, https://arxiv.org/abs/1905.04852.

［6］ Gatheral, J., Jaisson, T., and Rosenbaum, M., (2018). Volatility is rough, *Quantitative Finance*, **18**(6), 933–949.

［7］ Heber, G., Lunde, A., Shephard, N., and Sheppard, K., (2009). Oxford-Man Institute's Realized Library, version 0.1. Oxford-Man Institute, University of Oxford.

［8］ Jaffard, S., (1997). Multifractal formalism for functions. part I: Results valid for all functions, *SIAM Journal on Mathematical Analysis*, **28**, 944–970.

［9］ Kantelhardt, J., W., Zschiegner, S., A., Koscielny-Bunde, E., Havlin, S., Bunde, A., and Stanley, S., E., (2002). Multifractal detrended fluctuation analysis of nonstationary time series, *Physica A*, **316**(1-4), 87–114.

［10］ Mandelbrot, M., M., (1998). Multifractals and 1/f Noise, *Wild Self-Affinity in Physics*, Springer.

［11］ McAleer, M., Medeiros, M.C., (2008). Realized volatility: A review, *Econometric Reviews*, **27**, 10–45.

［12］ Mpanda, M., M., and Gorjão, L., R., (2025). Fluctuation Analysis of Volatility in South African Stock Market Indices, *Computational Economics*, https://doi.org/10.1007/s10614-025-11027-7.

［13］ Muzy, J. -F., Bacry, E., and Arneodo, A., (1993). Multifractal formalism for fractal signals: The structure-function approach versus the wavelet-transform modulus-maxima method, *Physical Review E*, **47**(2), 875–884.

［14］ Takaishi, T., (2020). Rough volatility of Bitcoin, *Finance Research Letters*, **32**, 101379.

［15］ Torrence, C., and Compo, G., P., (1998). A practical guide to wavelet analysis, *Bulletin of the American Meteorological Society*, **79**(1), 61–78.