Sankhyā : The Indian Journal of Statistics
2024, Volume 86-B, Part 1, pp. 191–223

Rough Volatility: Fact or Artefact?

Rama Cont
University of Oxford, Oxford, UK
Purba Das
King's College London, London, UK

Abstract

We investigate the statistical evidence for the use of 'rough' fractional processes with Hurst exponent $H < 0.5$ for modeling the volatility of financial assets, using a model-free approach. We introduce a non-parametric method for estimating the roughness of a function based on discrete sample, using the concept of normalized p-th variation along a sequence of partitions. Detailed numerical experiments based on sample paths of fractional Brownian motion and other fractional processes reveal good finite sample performance of our estimator for measuring the roughness of sample paths of stochastic processes. We then apply this method to estimate the roughness of realized volatility signals based on high-frequency observations. Detailed numerical experiments based on stochastic volatility models show that, even when the instantaneous volatility has diffusive dynamics with the same roughness as Brownian motion, the realized volatility exhibits rough behaviour corresponding to a Hurst exponent significantly smaller than 0.5. Comparison of roughness estimates for realized and instantaneous volatility in fractional volatility models with different values of Hurst exponent shows that, irrespective of the roughness of the spot volatility process, realized volatility always exhibits 'rough' behaviour with an apparent Hurst index $\widehat{H} < 0.5$. These results suggest that the origin of the roughness observed in realized volatility time series lies in the estimation error rather than the volatility process itself.

AMS (2000) subject classification. Primary 60L90, 62P05; Secondary 60G22, 62G10, 62M07.
Keywords and phrases. Roughness, variation index, p-th variation, realized volatility, instantaneous volatility, fractional Brownian motion, hurst exponent, high-frequency data.

1 Introduction

1.1 Fractional Processes in Finance: from Long-Range Dependence to 'Rough Volatility' Beginning with Mandelbrot and Van Ness (1968), fractional Brownian motion and fractional Gaussian noise have been used as

building blocks of stochastic models of various phenomena in physics, engineering (Lévy-Véhel et al. 2005) and finance (Baillie 1996, Bollerslev and Ole Mikkelsen 1996, Comte and Renault 1998, Cont 2007, Gatheral et al. 2018, Rogers 1997, Willinger et al. 1999). Fractional Brownian motion has two remarkable properties which have contributed towards its adoption as a building block in stochastic models: first, its ability to model long-range dependence, as measured by the slow decay $\sim T^{2H-2}$ of auto-correlation functions of increments, where $0 < H < 1$ is the Hurst exponent; second, its ability to generate trajectories which have varying levels of Hölder regularity ('roughness'). The former is a property that manifests itself over long time scales while the latter manifests itself over short time scales and, in general, these two properties are unrelated. But in the case of fractional Brownian motion, the two properties are linked through self-similarity and governed by the Hurst exponent $0 < H < 1$: for $H > 1/2$ one obtains long-range dependence in increments and trajectories smoother than Brownian motion while for $H < 1/2$ one obtains 'anti-correlated' increments and trajectories rougher than Brownian motion[1]. Processes driven by fractional Gaussian noise with $H < 1/2$ are thus sometimes referred to as 'rough processes'.

In early applications to financial data (Baillie 1996, Bollerslev and Ole Mikkelsen 1996, Comte and Renault 1998, Willinger et al. 1999), fractional processes were adopted in order to model *long-range dependence* effects in financial time series (Cont 2005). More specifically, statistical evidence of *volatility clustering* (Cont 2007) - positive dependence of the amplitude of returns over long time scales - led to the development of stochastic volatility models driven by fractional Brownian motion. A well-known example of such a fractional stochastic volatility model is the one proposed by Comte and Renault (1998) who modelled the dynamics of the (instantaneous) volatility $\sigma(t)$ of an asset as:

$$Y(t) = \ln \sigma(t) \qquad dY(t) = -\gamma Y(t)dt + \theta dB^H(t) \tag{1}$$

where B^H is a fractional Brownian motion (fBM) with Hurst exponent H. This long-range dependence in volatility is modelled by choosing fractional models which correspond to $1 > H > 1/2$ (Bollerslev and Ole Mikkelsen 1996, Breidt et al. 1998, Comte and Renault 1998, Fleming and Kirby 2011, Lahiri and Sen 2020).

A recent strand of literature, starting with Gatheral et al. (2018), has suggested the use of fractional Brownian models with $H < 1/2$ for mod-

[1] Incidentally, Hurst and Hölder happen to have the same initials, adding to the confusion...

elling volatility. Unlike previous studies based on auto-correlations of various volatility estimators over long time scales (Baillie 1996, Bollerslev and Ole Mikkelsen 1996, Cont 2005, Fleming and Kirby 2011), Gatheral et al. (2018) rely on the analysis of the behaviour of volatility estimators over short intraday time scales in order to assess the 'roughness' of these signals and concluded that volatility is 'rough' i.e. has paths with a Hölder regularity which is strictly less than 1/2, suggest to use stochastic models with sample paths rougher than Brownian motion.

However, it has not been lost on experts working in this area that previous estimation results for fractional models in the literature on long-range dependence in volatility, pointed towards Hurst exponents $H > 0.5$ (close to 0.55) (Baillie 1996, Comte and Renault 1998, Lahiri and Sen 2020) while the recent 'rough volatility' literature indicates Hurst exponents close to 0.1. Together with the well-known statistical issues plaguing the estimation of Hurst exponents (Beran 1994, Rogers 2019), these conflicting results call for a critical examination of the empirical evidence for 'rough volatility'.

Compounding this issue is the fact that (spot) volatility is not directly observed but estimated from price series, with an inherent estimation error which has been the subject of many studies (Barndorff-Nielsen and Shephard 2002, Jacod and Protter 2012, Lahiri and Sen 2020). This estimation error is far from i.i.d.: it is known to possess path-dependent features (see Jacod and Protter, 2012). As a result, measures of roughness for realized volatility indicators may be quite different from those of the underlying 'spot volatility'. This is simply because the convergence of high-frequency volatility estimators in L^p norms does not imply in any way their functional convergence in Hölder norms or other norms related to roughness.

As already pointed out by Rogers (1997), these two properties, namely the short-range behaviour which determines the roughness of the path, and the long-range dependence property, can (and should) be modeled through different mechanisms. Bennedsen et al. (2022) discuss several such approaches.

The focus of the literature on parametric models based on fractional Brownian motion or fractional Gaussian noise concentrates these two, very different, properties in a single parameter: the Hurst exponent H (Bolko et al. 2022, Fukasawa et al. 2022, Gatheral et al. 2018). Such parametric approaches proceed as follows: one estimates a parametric model for volatility dynamics based on some fractional Gaussian driving noise with Hurst exponent $0 < H < 1$ using a MLE (Fukasawa et al. 2022) or method of moments (Bolko et al. 2022). Then, based on the estimated value of this parameter H, one concludes that "volatility is rough" if $\hat{H} < 0.5$.

The validity of such approaches hinges on the assumption that the class of models used is well-specified. As pointed out by Bennedsen et al. (2022), this is unlikely to be the case for SDEs driven by fractional Gaussian noise if one wants to accommodate both (long-range) dependence properties and (short-range) roughness properties.

To avoid this caveat, we propose a model-free *non-parametric* method which focuses solely on the roughness properties of sample paths. Although less ambitious in its scope -we only focus on roughness properties rather than developing a full model for volatility dynamics- our approach is robust to the specification errors and estimation biases which plague parametric methods.

1.2 Contribution In this work we investigate whether the assertion that 'volatility is rough' (i.e. rougher than typical paths of Brownian motion) is supported by empirical evidence. We address this question in detail by re-examining the statistical evidence from high-frequency financial data. We investigate the statistical evidence for the use of 'rough' fractional processes with Hurst exponent $H < 0.5$ for the modelling of volatility of financial assets, using a non-parametric, model-free approach.

We introduce a non-parametric method for estimating the roughness of a function/path based on a (high-frequency) discrete sample, using the concept of normalized p-th variation along a sequence of partitions, and discuss the consistency of our estimator in a pathwise setting. We investigate the finite sample performance of our estimator for measuring the roughness of sample paths of stochastic processes using detailed numerical experiments based on sample paths of fractional Brownian motion and other fractional processes. We then apply this method to estimate the roughness of realized volatility signals based on high-frequency observations. We did a detailed numerical experiment based on stochastic volatility model. Our experiment shows that even when the instantaneous (spot) volatility has diffusive dynamics driven by Brownian motion, the realized volatility exhibits rough behaviour corresponding to a Hurst exponent significantly smaller than 0.5. In particular, as shown in Section 4.1, under the null hypothesis that (instantaneous) volatility follows an Ornstein-Uhlenbeck process with $H = 0.5$ one typically observes an apparent Hurst exponent $\hat{H} \simeq 0.1$ for realized volatility, similar to values reported in empirical studies as evidence of 'rough volatility'. This suggests that the origin of the roughness observed in realized volatility time series lie in the estimation error rather than the volatility process itself. Comparison of roughness estimates for realized and instantaneous volatility in fractional volatility models for different values of Hurst parameter H

shows that whatever the value of H for the (spot) volatility process, realized volatility always exhibits 'rough' behaviour.

Our results are broadly consistent with the observations by Rogers (2019). We pinpoint more precisely the origin of the apparent 'rough' behaviour of volatility as being the estimation error inherent in the estimation of realized volatility from discrete observations. In particular, our results question whether the empirical evidence presented from high-frequency volatility estimates may be used to infer that 'volatility is rough'. In fact, as shown in Section 4.1, we show that estimated values of the Hurst exponent close to 0.1 for *realized volatility* are typically obtained in stochastic volatility model with diffusion dynamics for which $H = 0.5$.

2 Measuring the Roughness of a Path

Determining the roughness of realized volatility from a sample path plays a crucial role in model specification. In practice, we observe only a single price path sampled on a discrete grid of observation times, so one is faced with the problem of determining the roughness of a process from a single path sampled at high frequency. We present various concepts for measuring the roughness of a path and discuss how they may be used to design estimators from high-frequency observations.

2.1 p-th Variation and Roughness Index of a Path Consider a sequence of partitions $\pi = (\pi^n)_{n\geq 1}$ of $[0, T]$ where

$$\pi^n = \left(0 = t_0^n < t_1^n < \cdots < t_{N(\pi^n)}^n = T\right) \tag{2}$$

represents observation times 'at frequency n'. We denote $N(\pi^n)$ to be the number of intervals in the partition π^n. Denote respectively by $|\pi^n| = \sup_{i=1,\cdots,N(\pi^n)} |t_i^n - t_{i-1}^n|$, and $\underline{\pi^n} = \inf_{i=1,\cdots,N(\pi^n)} |t_i^n - t_{i-1}^n|$, the size of the largest and the smallest interval of π^n. In this paper, we will always assume

$$|\pi^n| = \sup_{i=1,\cdots,N(\pi^n)} |t_i^n - t_{i-1}^n| \overset{n\to\infty}{\to} 0.$$

The concept of p-th variation along a sequence of partitions $\pi = (\pi^n)_{n\geq 1}$ with $0 = t_0^n < \ldots < t_k^n < \ldots < t_{N(\pi^n)}^n = T$ is defined following Cont and Perkowski (2019):

Definition 1 *(p-th variation along a sequence of partitions)* $x \in C^0([0, T], \mathbb{R})$ *has finite p-th variation along the sequence of partitions* $\pi = (\pi^n, n \geq 1)$ *if*

there exists a continuous increasing function $[x]_\pi^{(p)} : [0,T] \to \mathbb{R}_+$ *such that*

$$\forall t \in [0,T], \qquad \sum_{\substack{[t_j^n, t_{j+1}^n] \in \pi^n : \\ t_j^n \leq t}} \left| x(t_{j+1}^n) - x(t_j^n) \right|^p \overset{n\to\infty}{\longrightarrow} [x]_\pi^{(p)}(t). \tag{3}$$

We call $[x]_\pi^{(p)}$ *the* p-th variation *of* x *along the sequence of partitions* π. *We denote* $V_\pi^p([0,T],\mathbb{R})$ *the set of all continuous paths with finite* p*-th variation along* π.

We may formalize the concept of roughness using the notion of variation index and roughness index:

Definition 2 (Variation index). *The variation index of a path* x *along a partition sequence* π *is defined as the smallest* $p \geq 1$ *for which* x *has finite* p*-th variation along* π*:*

$$p^\pi(x) = \inf \left\{ p \geq 1 \ : \ x \in V_\pi^p([0,T],\mathbb{R}) \right\}.$$

Definition 3 (Roughness index). *The roughness index of a path* x *(along* π*) is defined as*

$$H^\pi(x) = \frac{1}{p^\pi(x)}.$$

When the underlying sequence of partitions is clear, we will omit π and denote these indices as $p(x)$ and $H(x)$. A similar roughness index was introduced by Han and Schied (2022).

For a (real-valued) stochastic process $X : [0,T] \times \Omega \to \mathbb{R}$ the roughness index $p^\pi(X(.,\omega))$ of each sample path $X(.,\omega)$ may be different in principle. Nevertheless there are many important classes of stochastic processes which have an *almost-sure* roughness index. For example, the roughness index of fractional Brownian motion (fBM) matches with the corresponding Hurst parameter/ Hölder exponent:

Example 1 *Brownian motion* B *has variation index* $p^\pi(B) = 2$ *and roughness index* $H^\pi(B) = \frac{1}{2}$ *along any refining partition sequence* π *or any partition* π *with* $|\pi^n| \log n \to 0$ (Dudley 1973).

Fractional Brownian motion B^H *has variation index* $p^{\mathbb{T}}(B^H) = \frac{1}{H}$ *and roughness index* $H^{\mathbb{T}} = H$ *along the dyadic partition sequence* $\mathbb{T}$.

In general, the existence of a variation index is not obvious. For further details see (Cont and Das 2024, Han and Schied 2022).

2.2 Normalized p-th Variation It is not easy to use p-th variation directly on empirical data for estimating roughness based on discrete observations, as this involves checking convergence to an unknown limit. We introduce a normalized version of p-th variation which has better asymptotic properties (Cont and Das 2024):

Definition 4 (Normalized p-th variation along a sequence of partitions) *Let π be a sequence of partitions of $[0,T]$ with mesh $|\pi^n| \to 0$ and $\pi^n = \left(0 = t_1^n < t_2^n < \cdots < t_{N(\pi^n)}^n = T\right)$. $x \in V_\pi^p([0,T],\mathbb{R})$ is said to have normalized p-th variation along π if there exists a continuous function $w(x,p,\pi) : [0,T] \to \mathbb{R}$ such that:*

$$\forall t \in [0,T], \sum_{\pi^n \cap [0,t]} \frac{\left|x\left(t_{i+1}^n\right) - x\left(t_i^n\right)\right|^p}{[x]_\pi^{(p)}(t_{i+1}^n) - [x]_\pi^{(p)}(t_i^n)} \times \left(t_{i+1}^n - t_i^n\right) \xrightarrow{n\to\infty} w(x,p,\pi)(t). \tag{4}$$

We denote $N_\pi^p([0,T],\mathbb{R})$ the class of all continuous functions for which the normalized p-th variation[2] *exists.*

The terminology is justified by the following result from (Cont and Das 2024) which shows that, for a large class of functions with $p-$th variation, the normalized p-th variation exists and is linear:

Theorem 2.1 *Let $x \in V_\pi^p([0,T],\mathbb{R})$ for some $p > 1$ where π be a sequence of partitions of $[0,T]$ with vanishing mesh $|\pi^n| \to 0$. Furthermore, if the p-th variation is absolutely continuous then:*

$$x \in N_\pi^p([0,T],\mathbb{R}) \qquad \text{and} \qquad \forall t \in [0,T],\ w(x,p,\pi)(t) = t.$$

The following result shows that normalized p-th variation is a 'sharp' statistic: if a function has finite p-th variation along a sequence of partitions π then for all $q \neq p$ the normalized p-th variation is either infinite or zero.

Theorem 2.2 *Let π be a sequence of partitions of $[0,T]$ with mesh $|\pi^n| \to 0$. Let $x \in V_\pi^p([0,T],\mathbb{R})$ with $[x]_\pi^{(p)} \in (0,\infty)$ for some $p > 1$.*

(i) For all $t \in (0,T]$ and for all $q > p$; $w(x,q,\pi)(t) = \infty$.

(ii) For all $t \in [0,T]$ and for all $q < p$; $w(x,q,\pi)(t) = 0$

In particular, Brownian motion almost surely has linear normalized quadratic variation.

[2] For $p = 2$ we will call this quantity as 'normalized quadratic variation'.

Example 2 (Normalized quadratic variation for Brownian motion) *Let B be a Wiener process on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and $(\pi^n)_{n\geq 1}$ be a sequence of partitions of $[0,T]$ with $|\pi^n| \log n \to 0$. Then for all $t \in (0,T]$:*

$$\mathbb{P}\left(w(B,2,\pi)(t) = t\right) = 1.$$

Remark 1 *In the above example, instead of taking any partition sequence with $|\pi^n| \log n \to 0$, we can also take any sequence of refining partitions. The proof is then different and uses the martingale convergence theorem.*

Example 3 (Stochastic integrals) *Let $X(t) = \int_0^t \sigma(u) dB_u$ where σ is an adapted process with $\int_0^T \sigma^2(u)du < \infty$. Then, for any refining partition sequence π with vanishing mesh and for all $t \in (0,T]$,*

$$\mathbb{P}\left(w(X,2,\pi)(t) = t\right) = 1.$$

Remark 2 *In the statement of Example 3, we can replace the assumption of refining partitions with partitions satisfying $|\pi^n| \log(n) \to 0$.*

The following example follows from (Viitasaari, 2019, Proposition 4.1):

Example 4 (Normalized p-th variation for fractional Brownian motion) *Let B^H be a fractional Brownian motion with Hurst index H, on a probability space $(\omega, \mathcal{F}, \mathbb{P})$. B^H has normalized p-th variation along the dyadic partition $\mathbb{T} = (\mathbb{T}^n)_{n\geq 1}$ for $p = 1/H$ almost-surely:*

$$\mathbb{P}\left(w\left(B^H, \frac{1}{H}, \mathbb{T}\right)(t) = t\right) = 1 \quad \forall t \in (0,T].$$

2.3 Estimating Roughness from Discrete Observations Given observations on a refining time partition π^L, we define the 'normalized p-th variation statistic' which is the discrete counterpart of the normalized p-th variation:

$$W(L,K,\pi,p,t,X) := \sum_{\pi^K \cap [0,t]} \frac{\left|X(t_{i+1}^K) - X(t_i^K)\right|^p}{\sum_{\pi^L \cap [t_i^K, t_{i+1}^K]} \left|X(t_{j+1}^L) - X(t_j^L)\right|^p} \times \left(t_{i+1}^K - t_i^K\right). \tag{5}$$

The definition of the statistic (5) involves two frequencies: a 'block' size K and a sample size $L \gg K$. As the partition is refining, π^K is a subpartition of π^L. The denominator is estimated by grouping the sample of size L into K many groups, where each group contains $\frac{L}{K}$ consecutive data points.

The statistic (5) converges to the normalized p-th variation (4) as the sample size L and block size increase:

$$\lim_{K\to\infty}\lim_{L\to\infty} W(L,K,\pi,p,t,x) = w(x,p,\pi)(t).$$

It is thus natural to define roughness estimators for a discretely sampled signal in terms of (5).

The *variation index estimator* $\widehat{p}_{L,K}(X)$ associated with the signal sampled on π^L is then obtained by computing $W(L,K,\pi,p,t,X)$ for different values of p and solving the following equation for $p^{\pi}_{L,K}(X)$:

$$W(L,K,\pi,p^{\pi}_{L,K}(X),T,X) = T. \tag{6}$$

One can either fix a window length T or solve (6) in a least squares sense across several values of T. An estimator for the roughness index is subsequently defined as:

$$\widehat{H}^{\pi}_{L,K}(X) = \frac{1}{p^{\pi}_{L,K}(X)}. \tag{7}$$

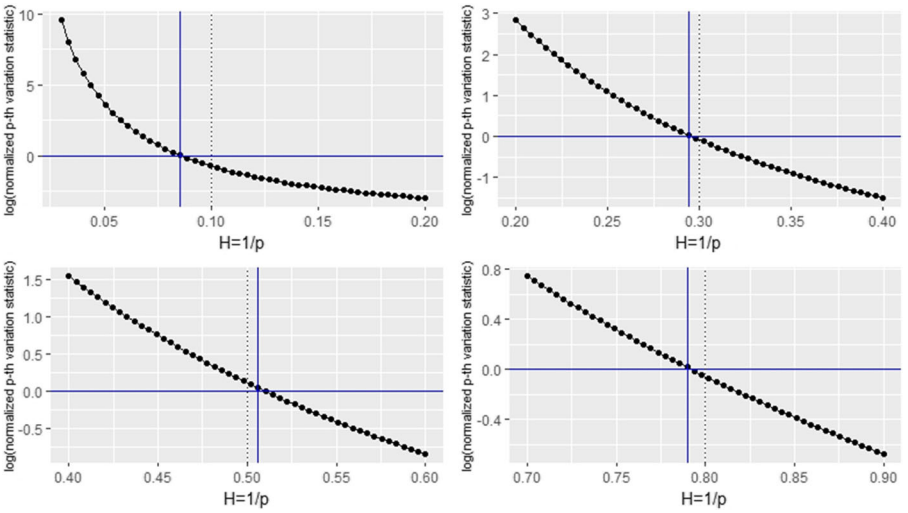

Figure 1: Normalized p-variation statistic for fractional Brownian motion as a function of Hurst parameter $H \in \{0.1, 0.3, 0.5, 0.8\}$. Blue vertical line represents $\widehat{H}_{L,K}$ with $L = 300 \times 300, K = 300$). The horizontal dotted line represents the true value H

We will denote the roughness estimator (7) as $\widehat{H}_{L,K}$ when the underlying dataset ands the corresponding partition sequence are clear. Asymptotic properties of these estimators under high-frequency sampling are studied in Cont and Das (2024).

2.4 Finite Sample Behaviour of the Roughness Estimator We will now study the finite sample behaviour of the roughness estimator $\widehat{H}^{\pi}_{L,K}(X)$ using high-frequency simulations of fractional Brownian motions. In the simulation examples unless mentioned otherwise we will use a uniform partition sequence of $[0, 1]$ with:

$$\pi^n = \left(0 < \frac{1}{n} < \frac{2}{n} < \cdots < 1\right).$$

To assess the finite sample accuracy of the estimator we compare the roughness index estimator $\widehat{H}^{\pi}_{L,K}$ with the underlying Hurst exponent $H \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8\}$. For every simulated path we compute cor-

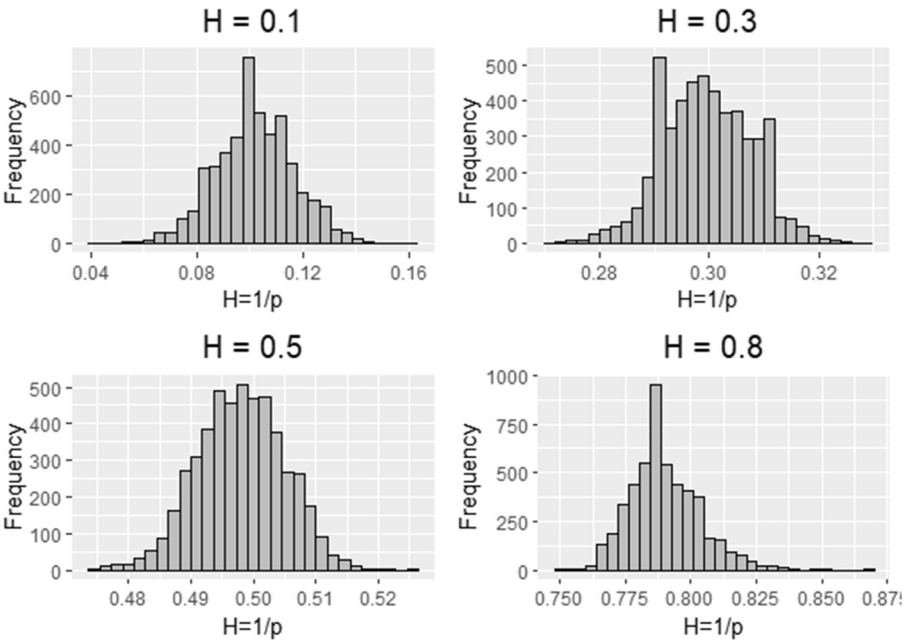

Figure 2: Histogram for estimated roughness index $\widehat{H}_{L,K}$ with $L = 300 \times 300, K = 300$ generated from 5000 independent simulations of fractional Brownian motion with Hurst parameter $H \in \{0.1, 0.3, 0.5, 0.8\}$

responding statistics $W(L = 300 \times 300, K = 300, \pi, q, t = 1, X = B^H)$ for different values of q, in order to estimate $\widehat{H}_{L,K}$. In Fig. 1, the black line represents the value of $\log\left(W(L = 300 \times 300, K = 300, \pi, q = p, t = 1, X = B^H)\right)$ plotted against roughness index $1/p$ in log-scale. The blue horizontal line represents the estimated roughness index $\widehat{H}_{L,K}$ whereas the dotted horizontal line represents the Hurst parameter. The blue horizontal line is computed by taking the intersection of the vertical line at 0 and black line Fig. 2 shows the histograms of roughness estimator $\widehat{H}^{\pi}_{L,K}$ generated from 5000 independent paths. We observe that for datasets with length, $L = 300 \times 300$ our roughness estimator $\widehat{H}_{L=300\times300,K=300}$ has satisfactory accuracy. Table 1 provides summary statistics for roughness index $\widehat{H}$ of 5000 simulated fractional Brownian motions.

Table 2 provides the variance, bias and mean square error (MSE) from 5000 independent estimated roughness index $\hat{H}_{L,K}$ with $L = 300 \times 300, K = 300$ for corresponding fractional Brownian motion.

Figure 3 represents a similar plot for simulated fractional Brownian motion with Hurst parameter $H = 0.1$. In Fig. 3, in left, similar to Fig. 1, $\log(W(L = 2000 \times 2000, K = 2000, \pi, p, t = 1, X = B^H))$ is plotted against $H = 1/p$ and the right plot represents the histograms of the estimator $\widehat{H}_{L=2000\times2000,K=2000}$ for 3000 independent simulations. The summary statistics for this estimator is provided in Table 3.

To compute the estimator $\widehat{H}_{L,K}$ we have different possible choices of $K \ll L$. Figure 4 shows how the estimator $\widehat{H}^{\pi}_{L,K}$ varies with K for fractional Brownian motion with Hurst parameter $H = 0.1$. The black line represents the $\widehat{H}_{L,K}$ plotted against different values of K whereas the blue vertical line represents the value for $L = 300 \times 300, K = 300$. We observe that when we

Table 1: Summary statistics for estimated roughness index $\widehat{H}_{L,K}$ for fractional Brownian motion B^H with $L = 300 \times 300, K = 300$

H	Min.	Lower quartile	Median	Mean	Upper quartile	Max.
0.1	0.0427	0.0918	0.1009	0.1018	0.1118	0.1627
0.2	0.1627	0.1936	0.2009	0.2001	0.2064	0.2391
0.3	0.2700	0.2936	0.2991	0.2995	0.3045	0.3282
0.4	0.3755	0.3936	0.3991	0.3991	0.4045	0.4245
0.5	0.4736	0.4936	0.4991	0.4989	0.5045	0.5264
0.6	0.5682	0.5936	0.5973	0.5982	0.6027	0.6282
0.7	0.6627	0.6900	0.6955	0.6955	0.7009	0.7282
0.8	0.7500	0.7809	0.7882	0.7900	0.7973	0.8682

Table 2: Variance, MSE and bias for $\widehat{H}_{L,K}$ calculated from 5000 independent sample paths for fractional Brownian motion B^H with $L = 300 \times 300, K = 300$

H	Variance	Bias	MSE
0.1	0.0002	0.0018	3.3124×10^{-06}
0.2	0.0001	8.0727×10^{-05}	6.5169×10^{-09}
0.3	6.5883×10^{-05}	−0.0005	2.1665×10^{-07}
0.4	5.4964×10^{-05}	−0.0009	7.2528×10^{-07}
0.5	4.8770×10^{-05}	−0.0011	1.2220×10^{-06}
0.6	5.3925×10^{-05}	−0.0018	3.3137×10^{-06}
0.7	7.7617×10^{-05}	−0.0045	2.0325×10^{-05}
0.8	0.0002	−0.0100	0.0001

vary K in the neighbourhood $K \approx \sqrt{L}$ the estimator performs well and is not very sensitive to the choice of K in this range.

In summary, we observe that for realistic sample sizes and frequencies encountered in high-frequency financial data, the estimator is quite accurate and not sensitive to the block size K in the range $K \approx \sqrt{L}$.

3 Spot Volatility and Realized Volatility

Unlike asset prices which may be observed and sampled directly from market data, *(spot) volatility* is not directly observable and as a consequence must be estimated from prices. The estimation methods frequently used are based on the concept of *realized volatility* (Andersen et al. 2003, Barndorff-Nielsen and Shephard 2002). Let us consider price dynamics given by a stochastic volatility model driven by a Brownian motion B:

$$dS_t = \sigma_t S_t dB_t + \mu_t S_t dt \tag{8}$$

Table 3: Summary statistics for estimated roughness index $\widehat{H}_{L,K}$ for fractional Brownian motion B^H with $H = 0.1, L = 2000 \times 2000, K = 2000$

Hurst index H	Min.	Lower quartile	Median	Mean	Upper quartile	Max.
0.1	0.086	0.096	0.099	0.099	0.103	0.117

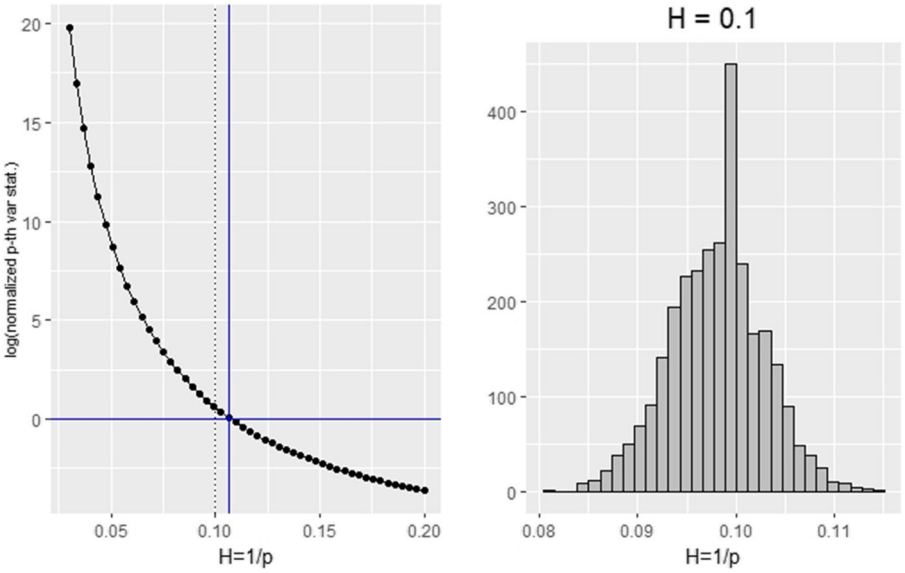

Figure 3: Simulation results for fractional Brownian motion with $H = 0.1$. **Left:** The log of normalized p-th variation statistic is plotted against $H = 1/p$ in black. The blue vertical line represents the estimated roughness index $\widehat{H}_{L,K}$ (with $L = 2000 \times 2000, K = 2000$). Dotted line: true value $H = 0.1$. **Right:** Histogram of estimated roughness index $\widehat{H}_{L,K}$ generated by simulating $n = 3000$ independent fractional Brownian paths with Hurst parameter $H = 0.1$

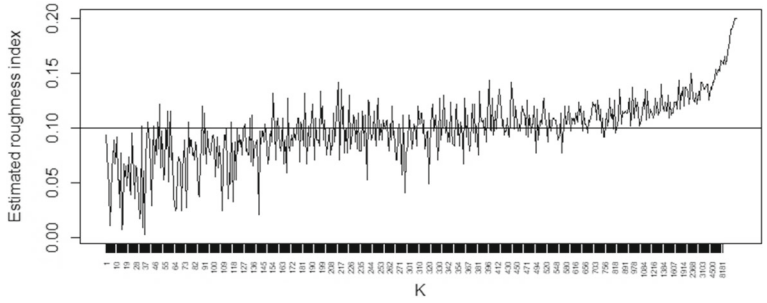

Figure 4: Estimated roughness index $\widehat{H}_{L=300\times300,K}$ as a function of block size K for a fractional Brownian motion with Hurst parameter $H = 0.1$

where σ_t represents the instantaneous or 'spot' volatility. In general, σ_t is represented as a random process itself driven by either (another) Brownian motion (Lewis 2000) or a fractional process (Comte and Renault 1998).

In practice the price S_t is observed on a (non-uniform) time grid

$$\pi^n = \left(0 = t_0^n < t_1^n < \cdots < t_{N(\pi^n)}^n = T\right) \tag{9}$$

where n represents a 'sampling frequency'. In order to study high-frequency asymptotics of roughness estimators, we assume the mesh $|\pi^n| \to 0$ as n increases. An example to keep in mind is the dyadic partition sequence: $\pi^n = \left(t_i^n = \frac{iT}{2^n}, i = 0, \cdots, 2^n\right)$ but none of the results below requires a uniform grid.

The spot volatility process σ_t may then be recovered from the *quadratic variation* of the log-price $X = \log S$ along this particular grid:

$$\sigma_t^2 = \frac{d}{dt}[\log S]_\pi(t) \tag{10}$$

where the quadratic variation $[\log S]_\pi$ of the log-price

$$[\log S]_\pi(t) = \lim_{n\to\infty} \sum_{\pi^n \cap [0,t]} \left(\log \frac{S(t_{i+1}^n)}{S(t_i^n)}\right)^2 = \lim_{n\to\infty} RV_t(\pi^n)^2$$

is computed as a high-frequency limit of the *realized variance* (Andersen et al. 2003, Barndorff-Nielsen and Shephard 2002) along the sampling grid π^n, defined as

$$RV_t(\pi^n)^2 = \sum_{\pi^n \cap [0,t]} \left(\log \frac{S(t_{i+1}^n)}{S(t_i^n)}\right)^2 = \sum_{\pi^n \cap [0,t]} \left(X(t_{i+1}^n) - X(t_i^n)\right)^2. \tag{11}$$

Realized volatility is defined as the square root of the realized variance.

Definition 5 (Realized volatility). *The realized volatility of a price process S over $[t, t+\Delta]$ sampled along the observation grid π^n is defined as:*

$$RV_{t,t+\Delta}(\pi^n) = \sqrt{\sum_{\pi^n \cap [t,t+\Delta]} \left(X(t_{i+1}^n) - X(t_i^n)\right)^2} = \sqrt{[X]_{\pi^n}(t+\Delta) - [\mathrm{X}]_{\pi^n}(\mathrm{t})} \tag{12}$$

where $X = \log S$.

If the price S_t follows a stochastic volatility model (8) with **instantaneous volatility** σ_t then realized variance converges to the quadratic variation of $\log S$ (also called '*integrated variance*') as sampling frequency increases (Barndorff-Nielsen and Shephard 2002, Jacod and Protter 2012):

$$RV_t(\pi^n)^2 \underset{n\to\infty}{\overset{\mathbb{P}}{\to}} IV_t := \int_0^t \sigma_u^2 du, \qquad RV_{t,t+\Delta}(\pi^n) \underset{n\to\infty}{\overset{\mathbb{P}}{\to}} \sqrt{IV_{t,t+\Delta}} = \sqrt{\int_t^{t+\Delta} \sigma_u^2 du}. \quad (13)$$

Along a single price path observed at high-frequency, we can compute the realized variance (12) and the realized volatility $RV_{t,t+\Delta}(\pi^n)$ in (13) may be used as a practical indicator of volatility:

$$RV_{t,t+\Delta}(\pi^n) \simeq \sqrt{\Delta}\, \sigma_t.$$

Several empirical studies have attempted to estimate the roughness of 'realized volatility' signals using high-frequency data, i.e. Andersen et al. (2003), Barndorff-Nielsen and Shephard (2002), Cont and Mancini (2011), Jacod and Protter (2012), Podolskij and Vetter (2009). A well-known reference is the study of Gatheral et al. (2018) where the authors estimate the roughness index of S&P500 realized volatility by performing the following logarithmic regression:

$$m(q,\Delta) := \frac{1}{n}[\log RV]_{\pi^n}^{(q)} = \frac{1}{\mathrm{n}} \sum_{\mathrm{t}=1}^{\mathrm{n}} |\log(\mathrm{RV_{t+\Delta}}) - \log(\mathrm{RV_t})|^{\mathrm{q}} \approx \mathrm{C_q} \Delta^{\xi_{\mathrm{q}}}. \quad (14)$$

The coefficients ξ_q are observed to behave linearly in q:

$$\xi_q \approx q\ \widehat{H}^R.$$

Regression of ξ_q on q yields an estimate $\widehat{H}^R$ of Hurst/Hölder index, for which Gatheral et al. (2018) report the value $\widehat{H}^R = 0.13$. Based on these observations, they propose a fractional SDE for (spot) volatility:

$$d \log \sigma_t^2 = \mu_t dt + \eta dB_t^H.$$

As we see from equation (14), the method used in Gatheral et al. (2018) actually uses p-th variation of the $\log(RV)$ to calculate the roughness of the underlying volatility process. Figure 5 is a replication of the log-regression model described above to estimate the roughness index of the volatility of 5-min S&P 500. However, in an interesting simulation study using paths

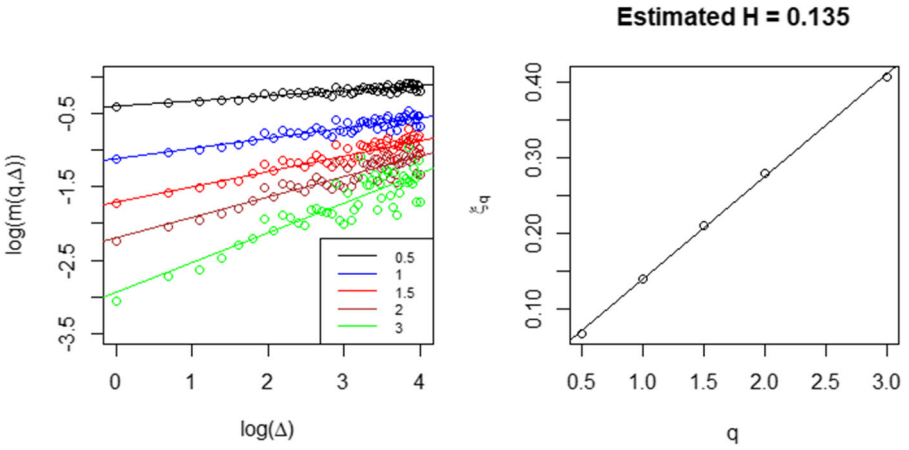

Figure 5: **Left:** A reproduction of the log-regression method introduced by Gatheral et al. (2018) using SPX 5-minute realized volatility from the Oxford-Man Institute's Realized Library. **Right:** Regression of ξ_q vs q: the estimated slope/ roughness index is $\widehat{H}^R = 0.135$

simulated from a Brownian OU volatility process, Rogers (2019) showed that the scaling behavior claimed as evidence for 'rough volatility' is also observed in a Brownian OU model over a range of time scales. He also showed that estimators of the roughness index based on log-regression of empirical p-th variation have poor accuracy.

Similar evidence for the lack of accuracy of such estimators based on log-regression of p-th variation is shown by Fukasawa et al. (2022).

4 Numerical Experiments

We now compare various roughness estimators for instantaneous volatility σ_t with those obtained from realized volatility $RV_{t,t+\Delta}(\pi^n)$ using price trajectories simulated from stochastic volatility models with varying degrees of "roughness".

4.1 Stochastic Volatility Diffusion Models Let us first consider the following stochastic volatility where volatility is simply (the modulus of) a Brownian motion:

Example 5

$$dS_t = \sigma_t S_t dB_t, \quad with \quad \sigma_t = \sigma_0 |W_t|, \tag{15}$$

where B, W are independent Brownian motions.
Figure 6 represents a path of the instantaneous volatility σ_t and the realized volatility $RV_t(\pi^L)$ computed as in equation (12) by taking 300 consecutive data-points, which corresponds to a 5-minute moving window. The center plot of Fig. 6 represents the estimation error, which is defined as the difference of instantaneous and realized volatility. The ACF of the estimation error shows a strongly time-dependent pattern which rules out IID behaviour and indicates a complex dependence structure. The right plot of Fig. 6 represents the log-estimation error plot, which is defined as the difference of log-instantaneous and log-realized volatility.

The estimated roughness index of instantaneous and realized volatility are observed to be very different. In the left graph of Fig. 7 we plot $\log(W(K = 500, L = 500 \times 500, \pi, p, t = 1, X = RV))$ against $H = 1/p$ for the realized volatility. On the other hand, the right graph is for instantaneous volatility. The estimated roughness index for realized volatility ($\widehat{H}_{L=500\times500,K=500}(RV) \approx 0.27$) is significantly smaller than the roughness index of the instantaneous volatility ($\widehat{H}_{L=500\times500,K=500}(\sigma) \approx 0.49$) suggest-

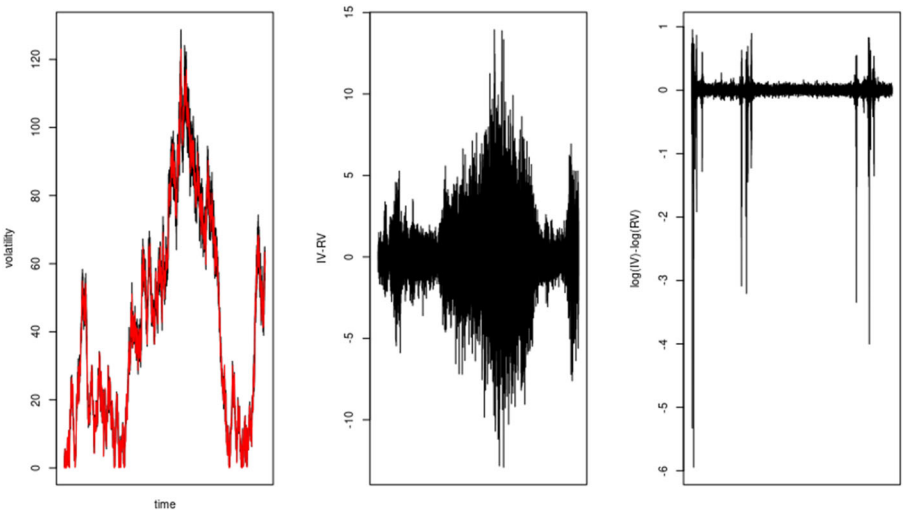

Figure 6: Simulations based on Brownian stochastic volatility model (15): $dS_t = S_t\sigma_t dB_t, \sigma_t = |W_t|$, where B and W are independent Brownian motions. **Left:** Spot volatility σ_t (red) vs realized volatility RV_t (black). **Center:** Estimation error $RV_t - \sigma_t$. **Right:** Log-estimation error $\log RV_t - \log \sigma_t$

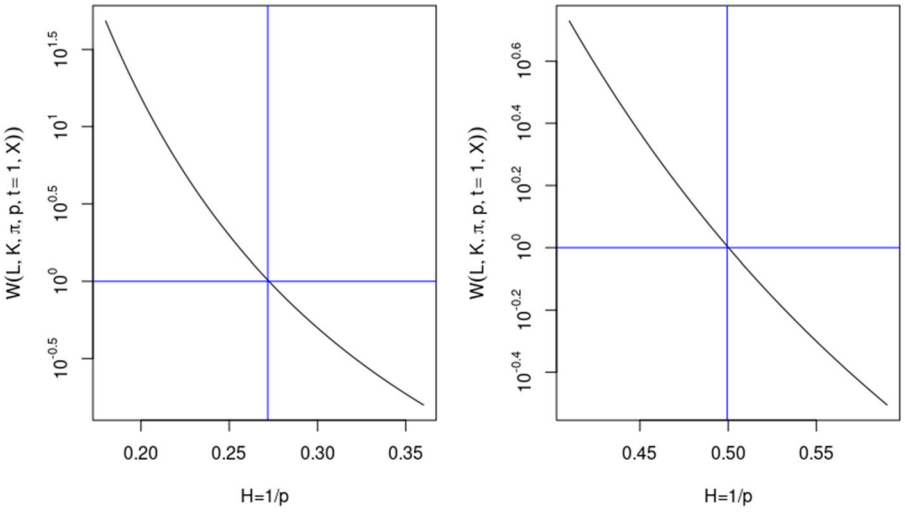

Figure 7: **Left:** Estimated roughness index $\widehat{H}_{L,K} \simeq 0.27$ (via normalized p-variation statistic with $L = 500 \times 500, K = 500$), for realized volatility derived from a Brownian diffusion model (15). **Right:** Estimated roughness index $\widehat{H}_{L,K} \simeq 0.49$ for spot volatility along the same price path

ing rougher behaviour of realized volatility. As in our simulation study we do not have any measurement errors, this roughness behaviour purely comes from estimation error. In some studies it is assumed that the estimation error or the log-estimation error is I.I.D. (see for example Fukasawa et al. (2022)). But as one can see from this diffusion example, both the estimation error and the log-estimation error is far from I.I.D. Hence the assumptions put forth for example in Fukasawa et al. (2022), is not very realistic for general diffusion models.

The solid black lines in Figs. 8 and 9 respectively represent the estimated roughness index $\widehat{H}_{L=300\times300,K}$ plotted against different values of K for the realized and instantaneous volatility (Model 15). The blue vertical line represents for $K = 500, L = 500 \times 500$. From the figures, we can observe that irrespective of the choice of K for the finite sample dataset of length $L = 500 \times 500$, the realized volatility is significantly rougher than the instantaneous volatility.

We now compare our roughness estimator with the log-regression method suggested in Gatheral et al. (2018) for the Model (15). It turns out that even with the log-regression model, similar 'rougher' realized volatility is observed

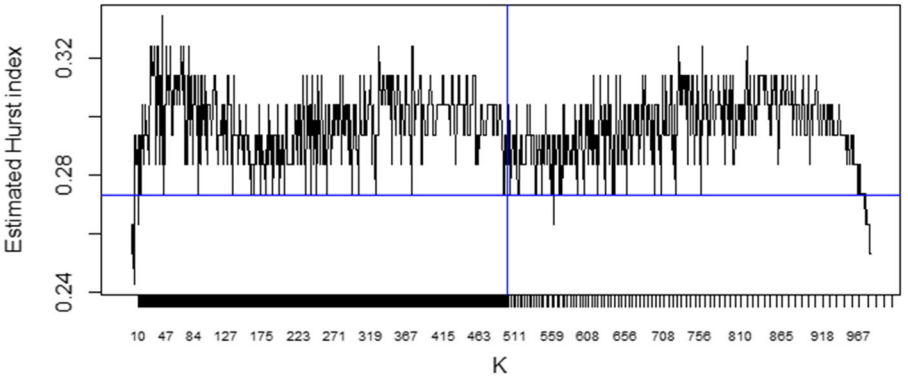

Figure 8: The solid black line represents the estimated roughness index $\widehat{H}_{L,K}$ via normalized p-th variation statistic $W(L = 500 \times 500, K, \pi, q, t = 1, X = RV)$ plotted against different values of K for the realized volatility shown in Fig. 6. The blue vertical line represents $L = 500 \times 500, K = 500$ whereas the blue horizontal line represents $\widehat{H} = 0.273$

even if the instantaneous volatility has Brownian diffusive behaviour. Figures 10 and 11 show that the realized volatility has a significant smaller roughness index than the instantaneous volatility even with respect to the log regression method. In this example it is clear that the roughness observed

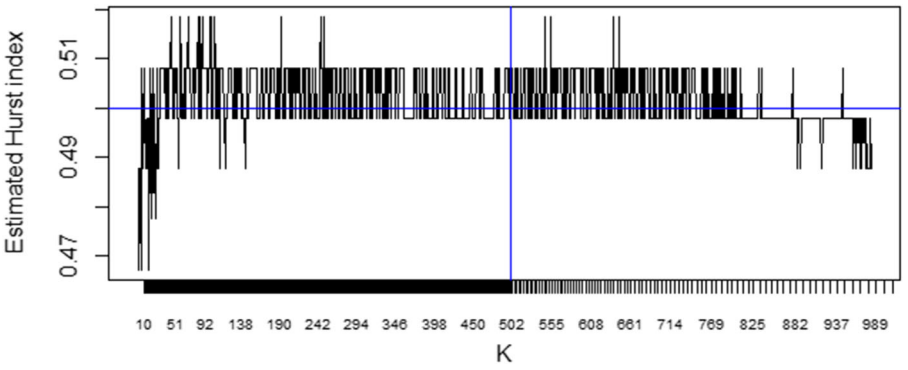

Figure 9: The solid black line represents the estimated roughness index $\widehat{H}_{L,K}$ via normalized p-th variation statistic $W(L = 500 \times 500, K, \pi, q, t = 1, X = IV)$ plotted against different values of K for the instantaneous volatility shown in Fig. 6. The blue vertical line represents $L = 500 \times 500, K = 500$ whereas the blue horizontal line represents true Hurst parameter $H = 0.5$

in realized volatility is attributable to the discretization error ('estimation error') and not the roughness of the spot volatility process, which is Brownian.

Next we consider a more realistic mean-reverting volatility model in which the volatility follows a Brownian Ornstein–Uhlenbeck process:

Example 6 (OU-SV model)

$$dS_t = S_t \sigma_t dB_t, \qquad \sigma_t = \sigma_0 e^{Y_t}, \quad dY_t = -\gamma Y_t dt + \theta dB'_t \tag{16}$$

where (B, W) *is a two-dimensional Brownian motions.*

In the simulations shown below, we use $\sigma_0 = 1, Y_0 = 0$ and $\gamma = \theta = 1$ and we take B, W independent. The left plot of Fig. 12 represents the realized volatility (respectively spot volatility) of the price process in black (respectively red). The center plot in Fig. 12 represents the corresponding estimation error, defined as the difference between the realized and the spot volatility. We observe that the estimation error has a complex dependence structure, and is certainly not IID. But, in contrast to Example 5, here the estimation error for the *log*-volatility, shown in Fig. 12 right, is closer to IID (this is supported by the theoretical developments in Fukasawa et al. (2022)).

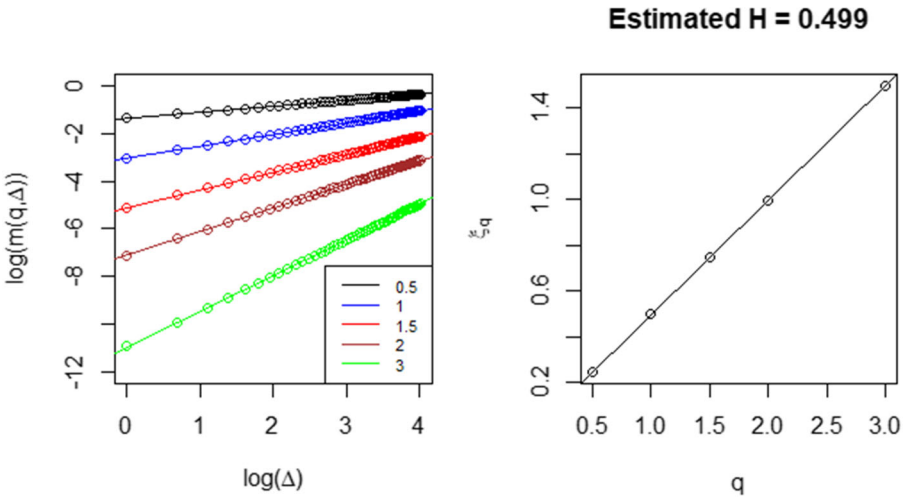

Figure 10: **Left:** Scaling analysis of instantaneous volatility of a simulated Brownian stochastic volatility model (Example 5) using the log-regression method used in Gatheral et al. (2018). **Right:** Regression coefficients ξ_q as a function of q. The estimated roughness index is via log-regression method $\widehat{H}^R = 0.499$

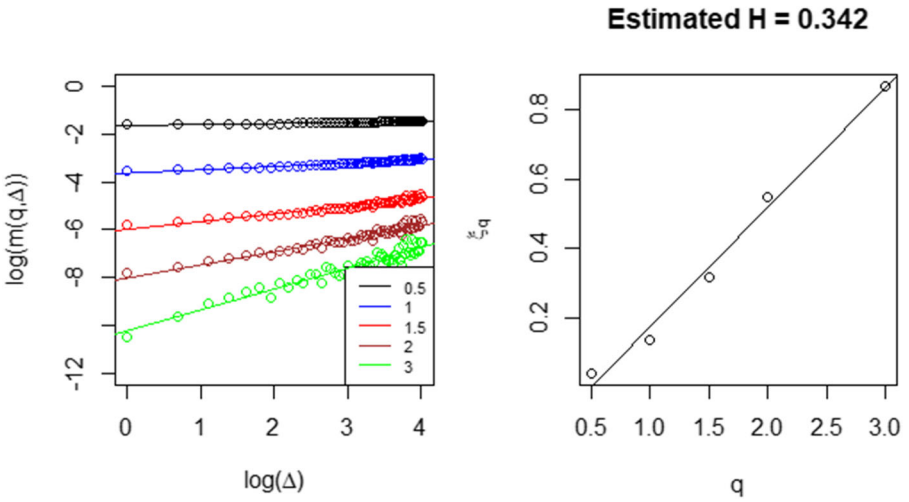

Figure 11: **Left:** Scaling analysis of realized volatility of a simulated Brownian stochastic volatility model (Example 5) using the log-regression method used in Gatheral et al. (2018). **Right:** Regression coefficients ξ_q as a function of q. The estimated roughness index via log-regression method is $\widehat{H}^R = 0.34$

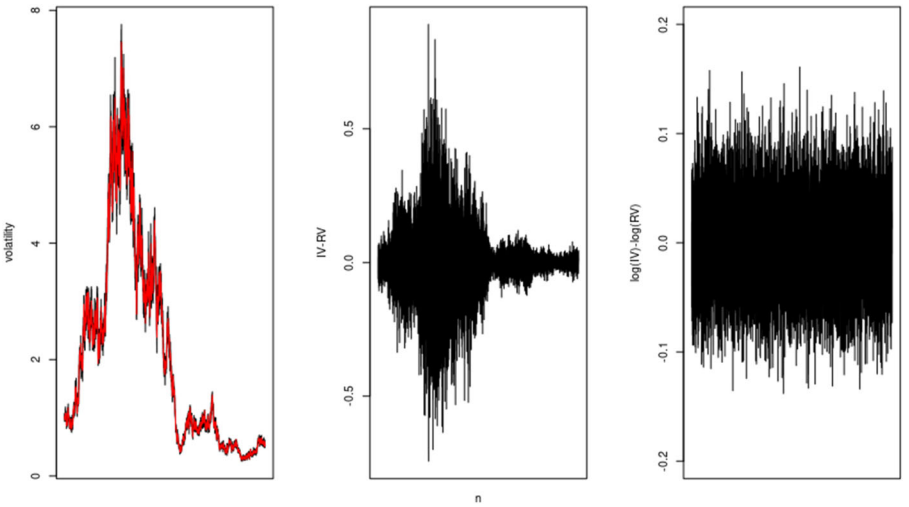

Figure 12: **Left:**Realized volatility (in black) vs instantaneous volatility (red) for the OU stochastic volatility model (16). **Center:** Volatility estimation error $RV_t - \sigma_t$. **Right:** Estimation error for log-volatility $\log RV_t - \log \sigma_t$

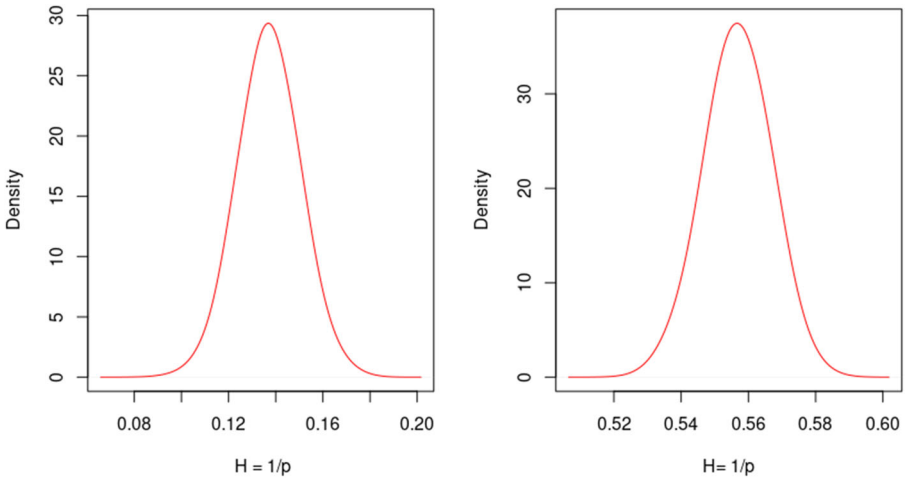

Figure 13: Distribution of the estimator $\widehat{H}_{L,K}$ for $(L = 300 \times 300, K = 300)$ across 2500 independent simulations for the OU-SV model (16) with $H = 0.5$ for spot volatility. **Left:** Estimated roughness for realized volatility: mean is ≈ 0.13. **Right:** Estimated roughness for spot volatility: mean is ≈ 0.55

Now we compare the distribution of the estimator $\widehat{H}_{L,K}$ with $(L = 300 \times 300, K = 300)$ for realized and instantaneous volatility across 2500 independent paths drawn from (16). The left plot in Fig. 13 is the distribution of estimated roughness index $\widehat{H}_{L,K}$ for the realized volatility while the right plot is for spot volatility. Table 4 provides summary statistics for the

Table 4: Summary statistics of estimated roughness index $\widehat{H}_{L,K}, L = 300 \times 300, K = 300$ for realized volatility and instantaneous volatility across 2500 independent sample paths from the diffusion model (16)

	Realized volatility	Instantaneous volatility
Min.	0.087	0.528
5% quantile	0.121	0.540
25% quantile	0.128	0.552
Median	0.136	0.556
Mean	0.137	0.557
75% quantile	0.148	0.563
95% quantile	0.158	0.569
Max.	0.181	0.581

roughness estimator $\widehat{H}_{L,K}$ with $L = 300 \times 300, K = 300$ across 2500 independent sample paths for realized volatility and instantaneous volatility, respectively. As one can observe there is an upward bias in the estimator for spot volatility and a strong negative bias for the roughness index based on realized volatility.

4.2 Fractional Stochastic Volatility Model In both previous examples, instantaneous volatility follows a diffusive behaviour driven by Brownian motion with $H = \frac{1}{2}$, yet the realized volatility exhibits "rough" behaviour with an estimated roughness index significantly smaller than 0.5. We now consider the more general case of a fractional stochastic volatility model (1):

Example 7 *(Fractional OU stochastic volatility model* (Comte and Renault 1998)*)*

$$dS_t = \sigma_t S_t dB_t, \quad \sigma_t = \sigma_0 e^{Y_t}; \qquad dY_t = -\gamma Y_t dt + \theta dB_t^H, \tag{17}$$

where $\gamma = \theta = \sigma_0 = 1$, B *is a Brownian motion and* B^H *a fractional Brownian motion with* $H \in (0,1)$.

In this model the spot volatility σ_t may have any roughness index $H \in (0,1)$. We explore how the roughness index of realized volatility relates to the roughness index H of spot volatility.

We compute the realized volatility $RV_t(\pi^L)$ as in (12) by taking 300 consecutive price data points, which corresponds to a 5-minute moving window. Further, we compare the estimated roughness index $\widehat{H}_{L,K}$ (with, $L = 300 \times 300, K = 300$) of instantaneous and realized volatility in the following table.

H	Instantaneous volatility	Realized volatility
0.10	0.130	0.190
0.20	0.215	0.250
0.30	0.310	0.258
0.40	0.413	0.207
0.50	0.507	0.130
0.60	0.601	0.087
0.70	0.678	0.061
0.80	0.756	0.052

Figure 14 shows sample paths for the price, realized volatility and the spot volatility for different values of $H = \{0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8\}$. We observe that for smaller H, the instantaneous volatility is rougher

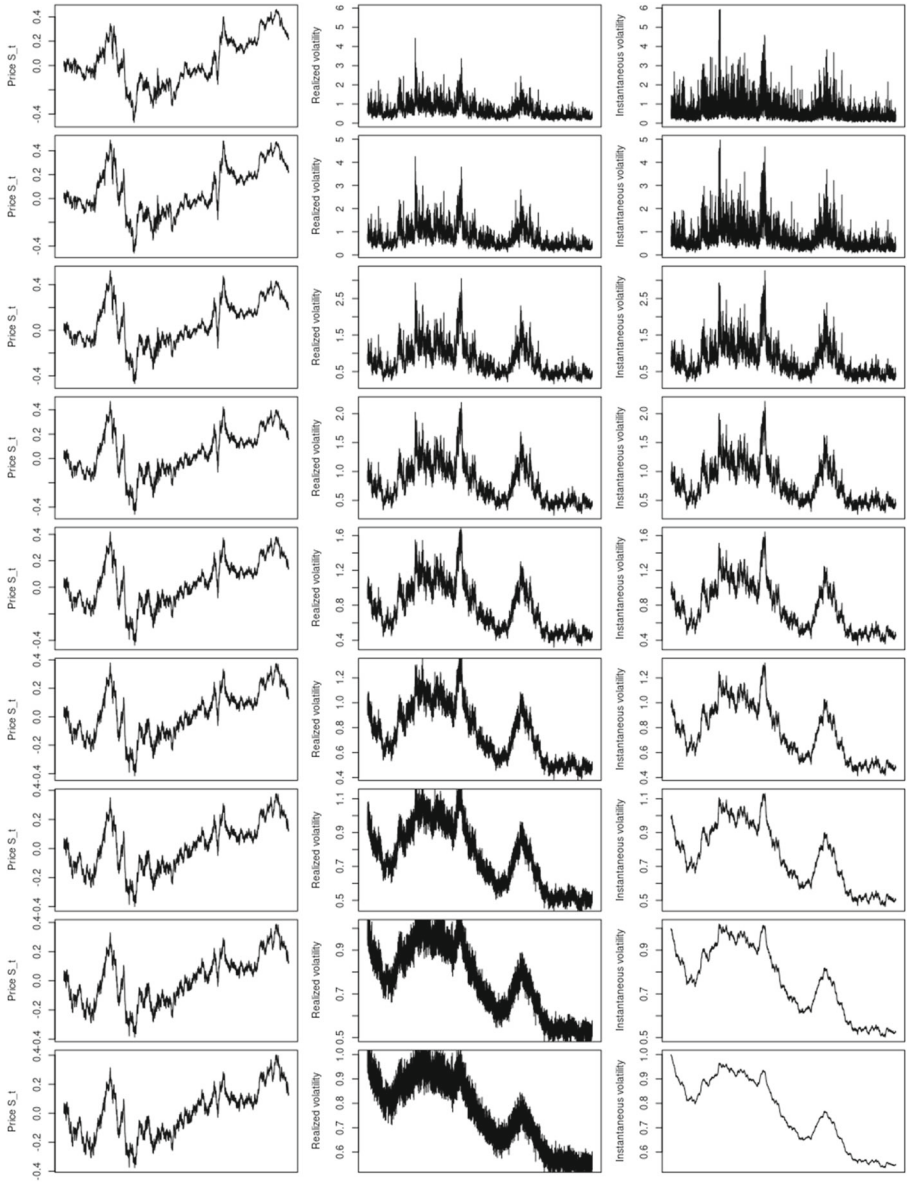

Figure 14: **Left:** Simulated price paths for fractional OU stochastic volatility model (17) for H={0.05,0.1,0.2,0.3,0.4,0.5,0.6,0.7,0.8} **Center:** Realized volatility. **Right:** Spot volatility

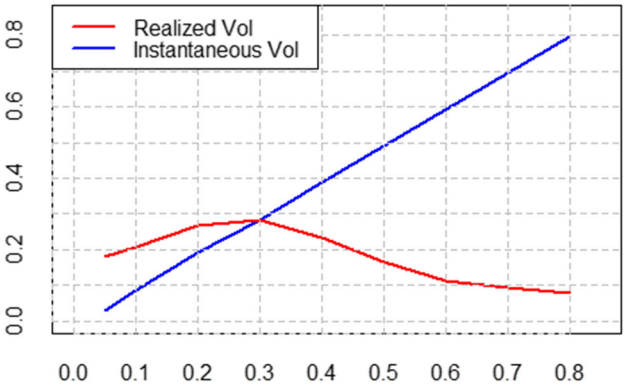

Figure 15: Estimated roughness index $\widehat{H}_{L,K}$ for realized volatility (red) and spot volatility (blue) for a high-frequency sample from the fractional-OU stochastic volatility model (17). X-axis: H. Y-axis: estimator $\widehat{H}_{L,K}$, $L = 300 \times 300, K = 300$

than realized volatility but as we increase H the realized volatility shows significantly rougher behaviour than the instantaneous volatility.

Figure 15 shows the estimated roughness index $\widehat{H}_{L,K}(RV)$ and $\widehat{H}_{L,K}(\sigma)$ respectively in red and blue. Though the roughness index of instantaneous volatility (represented in blue line) gives an accurate estimate of H, the roughness index for realized volatility is always below 0.3, regardless of whether $H < 0.5$ or $H > 0.5$! In particular, when the instantaneous volatility is smoother than Brownian motion (corresponding to $H \geq 0.5$) the estimated roughness index of realized volatility turns out to be an extremely poor estimate for H, giving *qualitatively incorrect* information about the roughness of volatility. In this regime, using roughness esimators based on realized volatility would lead us to incorrectly infer that (spot) volatility is rough ($H < 0.5$) whereas in fact the opposite is true ($H > 0.5$)!

Figure 17 shows the distribution of the estimator $\widehat{H}_{L,K}$ for spot volatility and realized volatility in the case $H = 0.55$, across 500 sample paths. Although $H > 0.5$, realized volatility exhibits an apparent rough behavior with an empirical value $\widehat{H}_{L,K}$ in the range $[0.02, 0.06]$.

This behavior is generically observed across the whole range $0 < H < 1$. Figure 16 shows the corresponding estimators $\widehat{H}_{L,K}(RV)$ and $\widehat{H}_{L,K}(\sigma)$ for 100 independent simulated price paths from (17). The bold black lines represent the mean across 100 independent simulations whereas the dotted lines represent the corresponding 25% and 75% confidence intervals. For

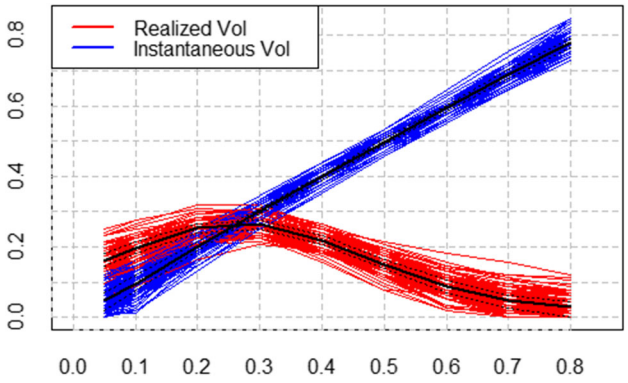

Figure 16: Estimated roughness index $\widehat{H}_{L,K}$ for realized volatility (red) and spot volatility (blue) for a high-frequency sample from the fractional-OU stochastic volatility model (17), for 100 independent price paths

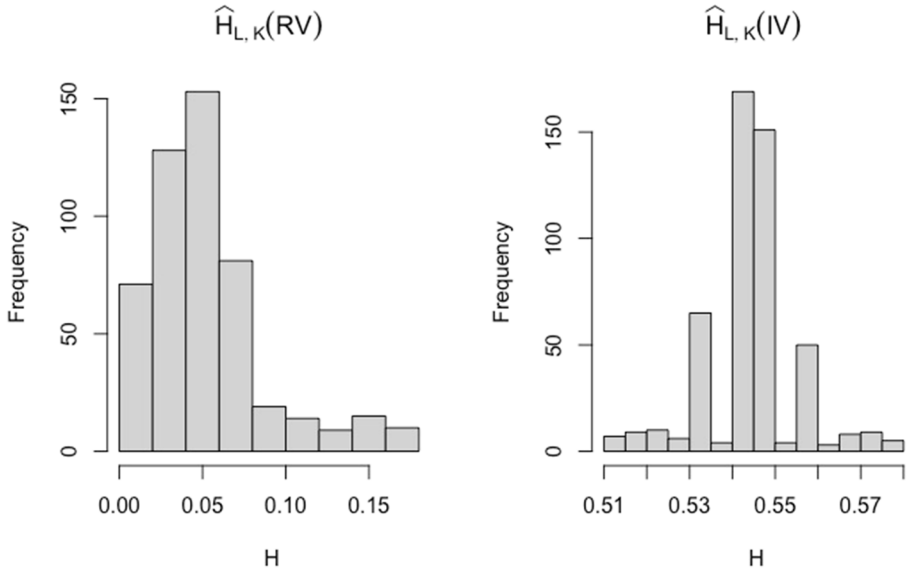

Figure 17: Histogram of the estimator $\widehat{H}_{L,K}$ for $(L = 300 \times 300, K = 300)$ across 500 sample paths from the Comte-Renault fractional OU-SV model (1) with $H = 0.55$. **Left:** Estimated roughness for realized volatility: mean is ≈ 0.05. **Right:** Estimated roughness for instantaneous volatility: mean is ≈ 0.55

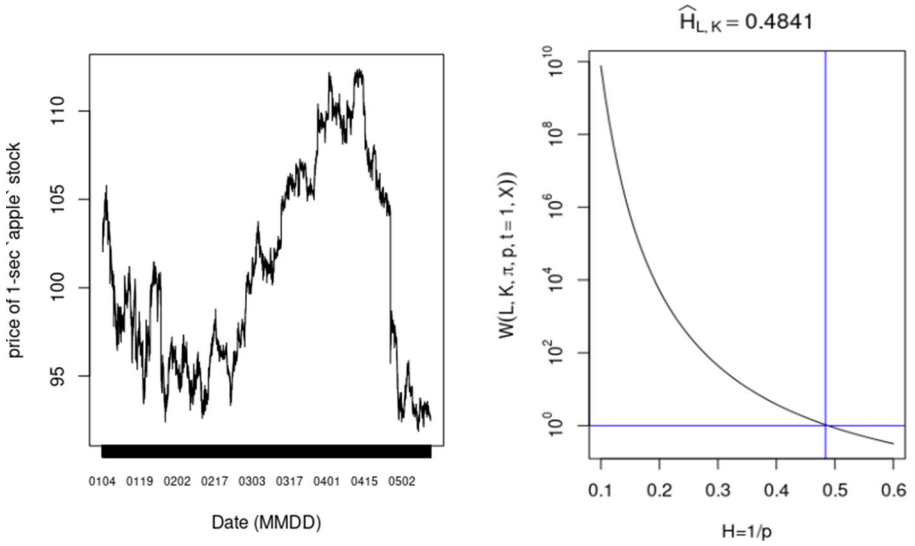

Figure 18: **Left:** Price of AAPL 04/Jan/2016 - 11/May/2016 (90 days). **Right:** From normalized p-th variation estimator the estimated roughness index $\widehat{H}_{L,K}$ (with $L = 1400 \times 1400, K = 1400$) for AAPL stock price data is approx 0.5

the price process (17), no matter what the value of the Hurst exponent for instantaneous volatility, the roughness index of realized volatility $\widehat{H}_{L,K}(RV)$ is always estimated in the range $[0.05, 0.3]$ (Fig. 16).

These examples illustrate our point: one *cannot conclude* from the (apparent) rough behavior of realized volatility that (spot) 'volatility is rough' i.e. reject the null hypothesis $H = 1/2$. Indeed, realized volatility exhibits 'rough' behaviour with $\widehat{H}_{L,K}(RV) < 0.5$ or $\widehat{H}^R < 0.5$, *even* when $H > 0.5$.

5 Application to High-Frequency Financial Data

Having extensively explored the performance of our roughness estimator $\widehat{H}_{L,K}$ based on the normalized p-th variation statistic for spot and realized volatility on simulated price paths, we now apply it to high-frequency financial time series.

5.1 AAPL Figure 18 (left) shows the second-by-second record of AAPL stock price. The right graph of Fig. 18 plots $\log(W(L = 1900 \times 1900, K = 1900, \pi, p, t = 1, X))$ against Hurst parameter $H = 1/p$ for

'AAPL' price. The plot suggests estimated roughness index for Apple second-by-second price path is 0.48.

Figure 19(left) represents 1-minute realized volatility of 'AAPL' in 2016. The right graph of Fig. 19 plots $\log(W(L = 310 \times 310, K = 310, \pi, p, t = 1, X = RV))$ against Hurst exponent $H = 1/p$ for the 1-min AAPL realized volatility. Fixing the value of $L = 310 \times 310$, if we deviate the value of block frequency K a little, then the estimated roughness index for 1-min realized volatility varies between 0.08 to 0.22. This is consistent with the results of Gatheral et al. (2018) regarding realized volatility. But as our simulation study suggests, the roughness index of realized volatility may be very different from that of spot volatility which is the quantity modelled in continuous-time stochastic volatility models.

5.2 SP500 Several studies on rough volatility, including the original study Gatheral et al. (2018), are based on the Oxford-Man Institute Realized Volatility dataset [3]. Figure 20 represents the plot of 5-minute realized volatility of $SP500$. The right graph of Fig. 20 plots $\log(W(L = 70 \times 70, K = 70, \pi, p, t = 1, X = RV))$ against Hurst parameter $H = 1/p$ for the 5-min Oxford-Man institute realized volatility data. Fixing the value of $L = 70 \times 70$, if we vary the value of K, the estimated roughness index $\widehat{H}_{L,K}$ of realized volatility varies between 0.05 to 0.25. This finding is consistent with the values observed in Gatheral et al. (2018).

Overall, the picture that emerges from SP500 and AAPL data is quite similar to the one observed in simulations of diffusion-type stochastic volatility models discussed in Section 4.1. As observed in Section 2.4, these observations are fully compatible with a diffusion-type stochastic volatility model such as (16). This suggests one cannot reject the null hypothesis $H(\sigma) = 1/2$ based on the observation that realized volatility exhibits 'rough' behaviour with $\widehat{H}_{L,K}(RV) < \frac{1}{2}$ or $\widehat{H}^R < 1/2$, *even* though these estimators (including the log-regression estimator used by Gatheral et al. (2018)) exhibit values as low as 0.1.

Comparing with the distribution of $\widehat{H}_{L,K}$ under the OU stochastic volatility model (16), shown in Fig. 13, and the quantiles reported in Table 2, we observe that the empirical values of $\widehat{H}^R$ and $\widehat{H}_{L,K}(RV)$ are in fact close to the mean of $\widehat{H}_{L,K}(RV)$ under the null hypothesis (16). In other words, one *cannot* reject the OU stochastic volatility model based on these values.

[3] https://realized.oxford-man.ox.ac.uk/data

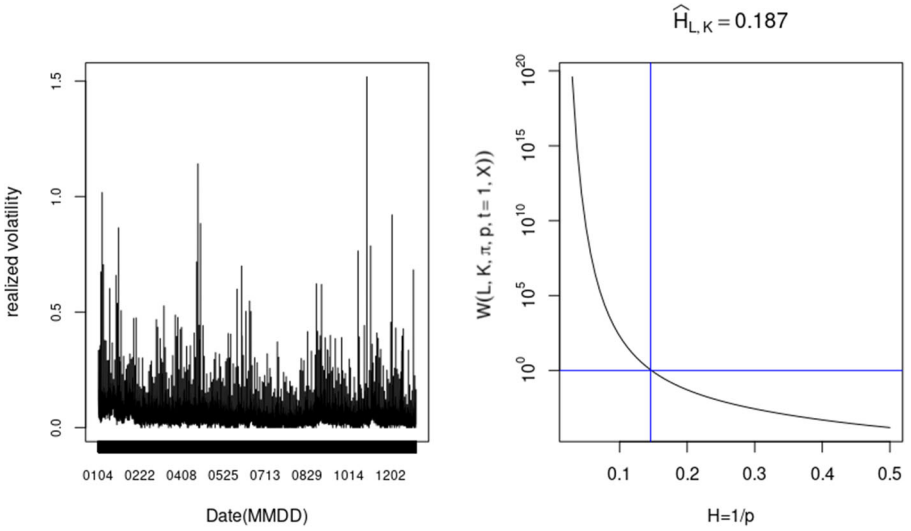

Figure 19: **Left:** Plot of 1-min realized volatility of 'AAPL' (year 2016). **Right:** Estimated roughness index $\widehat{H}_{L,K}$ (with $L = 310 \times 310, K = 310$) for the 1-min realized volatility for 'AAPL' is in the interval $[.08, 0.22]$)

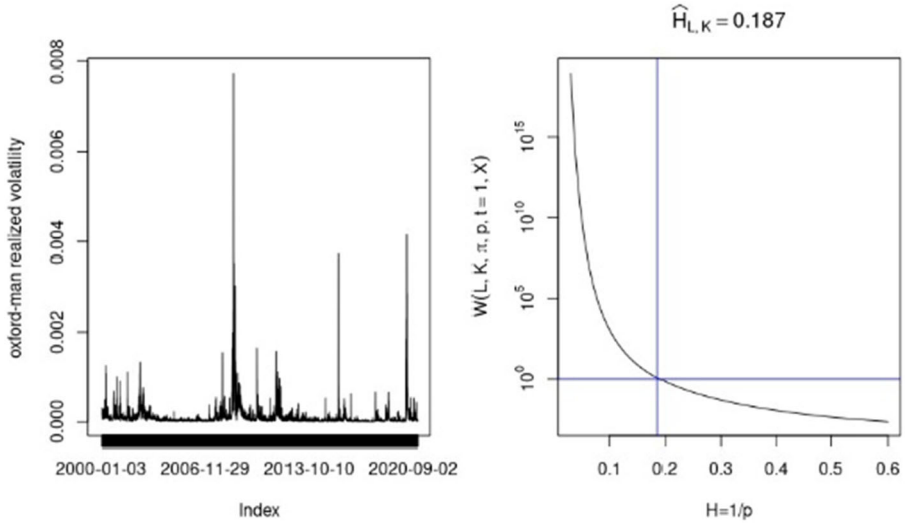

Figure 20: **Left:** S&P500 5-min realized volatility. **Right:** Estimated roughness index (via p-th variation statistics) of S&P500 realized volatility $\widehat{H}_{L,K} \in [.05 - .25]$ with $L = 70 \times 70, K = 70$

6 Rough Volatility ... or Estimation Error?

Given the large literature on 'rough volatility' in quantitative finance, it is somewhat surprising that the initial claim Gatheral et al. (2018) that one needs to model the spot volatility process using a 'rough' fractional noise with Hurst exponent $H < 1/2$ has not been examined more closely. Especially given that several follow-up studies (Fukasawa et al. 2022, Rogers 2019) point to the fact that the observations in Gatheral et al. (2018) may well be compatible with a Brownian diffusion model for volatility.

Our detailed examples illustrate that, for stochastic-volatility diffusion models driven by Brownian motion as described in Examples 5 and 6, the realized volatility has a roughness index $\widehat{H_{L,K}} \approx 0.3$ so exhibits an 'apparent roughness' which instantaneous volatility does not have, both in terms of normalized p-th variation statistics and also in terms of the log-regression method used by Gatheral et al. (2018). Clearly in these simulation examples this is entirely due to the discretization error or 'estimation error'.

These results suggest that one cannot hastily conclude that the roughness observed in realized volatility is an indicator of similar behaviour in spot volatility, as implicitly assumed in the 'rough volatility' literature. The observations in high-frequency financial data are in fact compatible with a stochastic volatility model driven by Brownian motion. So the origin of this apparent roughness in realized volatility may very well lie in estimation error rather than the noise process driving spot volatility.

As shown in Example 7, the rough behaviour of realized volatility does not lead us to reject the hypothesis that the underlying spot volatility may be modeled with a Brownian diffusion model or even a smoother model with long-range dependence and $H > 1/2$. We are thus drawn to concur with Rogers (2019) that "the notion that volatility is rough, that is, governed by a fractional Brownian motion (with $H < 1/2$), is not an incontrovertible established fact; simpler models explain the observations just as well."

More precisely, we have shown that the low empirical values observed for the roughness index of realized volatility are entirely compatible with diffusion-type stochastic volatility models such as the (well-studied) Ornstein-Uhlenbeck model (16). This observation, together with "Occam's razor", pleads for the use of diffusion-based stochastic volatility models which seem compatible with the empirical evidence but are far more tractable.

Acknowledgements. We thank Peter Friz, Masaaki Fukasawa, James-Michael Leahy, Mathieu Rosenbaum and Alexander Schied and participants at QuantMinds 2023 for helpful comments and discussions

Funding Information Funding information is not applicable as no funding was received.

Statements and Declarations

Conflicts of interest There is no conflict of interest.

References

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica*, **71**(2), 579–625.

Baillie, R. T. (1996). Long memory processes and fractional integration in econometrics. *Journal of Econometrics*, **73**(1), 5–59.

Barndorff-Nielsen, O. E. and Shephard, N. (2002). Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **64**(2), 253–280.

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.

Beran, J. (1994). *Statistics for long-memory processes*, vol 61 of *Monographs on Statistics and Applied Probability*. Chapman and Hall, New York.

Bolko, A. E., Christensen, K., Pakkanen, M. S., and Veliyev, B. (2022). A GMM approach to estimate the roughness of stochastic volatility. *Journal of Econometrics*, in press.

Bollerslev, T. and Ole Mikkelsen, H. (1996). Modeling and pricing long memory in stock market volatility. *Journal of Econometrics*, **73**(1),151–184.

Breidt, F., Crato, N., and de Lima, P. (1998). The detection and estimation of long memory in stochastic volatility. *Journal of Econometrics*, **83**(1),325–348.

Comte, F. and Renault, É. (1998). Long memory in continuous-time stochastic volatility models. *Mathematical Finance*, **8**, 291–323.
Cont, R. (2005). Long range dependence in financial markets. In *Fractals in Engineering*, (J. Lévy-Véhel and E. Lutton, eds.) pages 159–179, London. Springer London.
Cont, R. (2007). Volatility clustering in financial markets: Empirical facts and agent-based models. In *Long Memory in Economics*, (G. Teyssière and A.P. Kirman, eds.) pages 289–309. Springer, Berlin, Heidelberg.
Cont, R. and Das, P. (2024). Measuring the roughness of a signal. *Working Paper*.
Cont, R. and Mancini, C. (2011). Nonparametric tests for pathwise properties of semimartingales. *Bernoulli*, **17**(2), 781 – 813.
Cont, R. and Perkowski, N. (2019). Pathwise integration and change of variable formulas for continuous paths with arbitrary regularity. *Transactions of the American Mathematical Society*, **6**, 134–138.
Dudley, R. M. (1973). Sample functions of the gaussian process. *Ann. Probab.*, **1**(1), 66–103.
Fleming, J. and Kirby, C. (2011). Long memory in volatility and trading volume. *Journal of Banking and Finance*, **35**(7), 1714–1726.
Fukasawa, M., Takabatake, T., and Westphal, R. (2022). Consistent estimation for fractional stochastic volatility model under high-frequency asymptotics. *Mathematical Finance*, **32**(4), 1086–1132.
Gatheral, J., Jaisson, T., and Rosenbaum, M. (2018). Volatility is rough. *Quantitative Finance*, **18**, 933–949.
Han, X. and Schied, A. (2022). The hurst roughness exponent and its model-free estimation.
Jacod, J. and Protter, P. (2012). *Discretization of processes*, volume 67 of *Stochastic Modelling and Applied Probability*. Springer, Heidelberg.
Lahiri, A. and Sen, R. (2020). Fractional brownian markets with time-varying volatility and high-frequency data. *Econometrics and Statistics*, **16**, 91–107.
Lévy-Véhel, J., Lutton, E., and Tricot, C. (2005). *Fractals in Engineering*. Springer, Berlin.
Lewis, A. L. (2000). Option valuation under stochastic volatility. *Finance Press*.
Mandelbrot, B. B. and Van Ness, J. W. (1968). Fractional Brownian motions, fractional noises and applications. *SIAM Review*, **10**(4), 422–437.
Podolskij, M. and Vetter, M. (2009). Estimation of volatility functionals in the simultaneous presence of microstructure noise and jumps. *Bernoulli*, **15**(3), 634 – 658.
Rogers, L. C. G. (1997). Arbitrage with fractional Brownian motion. *Mathematical Finance*, **7**, 95–105.
Rogers, L. C. G. (2019). Things we think we know. . https://www.skokholm.co.uk/wp-content/uploads/2019/11/TWTWKpaper.pdf.
Viitasaari, L. (2019). Necessary and sufficient conditions for limit theorems for quadratic variations of gaussian sequences. *Probab. Surveys*, **16**, 62–98.
Willinger, W., Taqqu, M. S., and Teverovsky, V. (1999). Stock market prices and long-range dependence. *Finance and stochastics*, **3**(1), 1–13.

Rama Cont
Mathematical Institute, University of Oxford, Oxford, UK

Purba Das
Department of Mathematics, King's College London, London, UK
E-mail: purba.das@kcl.ac.uk

Paper received: 13 September 2023; accepted 12 January 2024.