Rough volatility: fact or artefact?

Rama CONT and Purba DAS

Depertment of Mathematics, University of Michigan

February 8, 2023

UNIVERSITY OF MICHIGAN

LSA MATHEMATICS UNIVERSITY OF MICHIGAN

Outline

Fractional processes in finance: long-range to 'rough vol.'

Since Mandelbrot & VanNess (1968) Fractional Brownian motion have been used as a building block in stochastic models.

- Long-range dependence, as measured by the slow decay $\sim T^{2H-2}$ of autocorrelation functions of increments, where $0 < H < 1$ is the Hurst exponent.
- Roughness: Generate trajectories which have varying levels of Hölder regularity.

For fBM, the two properties are linked through self-similarity.

- In early applications to financial data, fractional processes were used to model *long range dependence* in financial time series (see **Baillie 1996**). Long-range dependence is modeled by choosing $1 > H > 1/2$ (Comte and Renault 1998)
- A recent strand of literature, starting with Gatheral et al.(2018) has suggested the use of fractional Brownian models with $H < 1/2$ for modeling volatility. Gatheral et al (2018) present empirical data on volatility estimators suggesting that volatility is 'rough' i.e. has an (Hölder) roughness $< 1/2$.

Follow-up studies based on **parametric models** (Fukasawa et al 2022, Pakkanen et al 2020).
We re-examine empirical evidence for these claims with a new model-free methodology.
Idea: assess the roughness of a signal/process from a (discrete) sample of its path.

Roughness and Hölder regularity

A function $X : [0, T] \to \mathbb{R}$ has Hölder exponent α if

$$|X(t+\Delta) - X(t)| \leq C\ \Delta^{\alpha}, \qquad i.e. \qquad \sup_{t,\Delta} \frac{|X(t+\Delta) - X(t)|}{\Delta^{\alpha}} < \infty$$

The ‘roughness’ of X is measured by the smallest $\alpha > 0$ for which this holds: then

$$|X(t+\Delta) - X(t)| \underset{\Delta \to 0}{\sim} c\Delta^{H}$$

In that case if we choose $t_j^n = jT/n$ (grid with step $1/n$) then the sum

$$\sum_{j=1}^{n} \left| X(t_{j+1}^n) - X(t_j^n) \right|^p \sim c\, \frac{n}{n^p H}$$

will go to zero for $p > 1/H$,
will converge to a non-zero limit for $p = 1/H$
while it will blow up to ∞ for $p < 1/H$.

p-th variation along a sequence of partitions

Consider a sequence $\pi = (\pi^n)_{n\geq 1}$ of partitions of $[0, T]$

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

representing observation times,with resolution $|\pi^n| = \sup(t_{i+1}^n - t_i^n) \to 0$.

Definition (p-th variation)

Let $x \in C^0([0, T], \mathbb{R})$. x has finite p-th variation along π if there exists a continuous function $[x]_\pi^{(p)}$ such that

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

If this property holds, then the convergence in (1) is uniform. Define,

$$V_\pi^p([0, T], \mathbb{R}) = \left\{ x \in C^0([0, T], \mathbb{R}) \,\middle|\, x \text{ has finite p-th variation along } \pi \right\}.$$

Variation and roughness index of a path

Definition (Variation index)

For a continuous function $x \in C^0([0,T],\mathbb{R})$ we define the variation index as the smallest $p \geq 1$ for which p-th variation along π is finite:

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

Definition (Roughness index)

For a continuous function $x \in C^0([0,T],\mathbb{R})$ the roughness index is defined as:

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

- Fractional Brownian motion B^H has variation index $p^\pi(B^H) = \frac{1}{H}$ and roughness index H.

Existence of Variation index

Define, $q^+(x,\pi)$ and $q^-(x,\pi)$ as follows:

$$q^+(x,\pi) := \sup_{p>0} \left\{ \limsup_{n\uparrow\infty} \sum_{\pi^n} \left| x(t_{i+1}^n) - x(t_i^n) \right|^p = \infty \right\} \text{ and,}$$

$$q^-(x,\pi) := \sup_{p>0} \left\{ \liminf_{n\uparrow\infty} \sum_{\pi^n} \left| x(t_{i+1}^n) - x(t_i^n) \right|^p = \infty \right\}.$$

Lemma (Existence of variation index)

Let $x \in C^0([0,T],\mathbb{R})$. For any partition sequence π with vanishing mesh, the variation index $p^\pi(x)$ exists if and only if $q^+(x,\pi) = q^-(x,\pi)$. In this case, $p^\pi(x) = q^+(x,\pi) = q^-(x,\pi)$.

p-th variation and roughness

For a continuous path, the pathwise p-th variation plays an important role in determining the 'roughness' of a function.

Lemma

Let π be a sequence of partitions with vanishing mesh $|\pi^n| \downarrow 0$. If $x \in C^0([0,T],\mathbb{R})$ has variation index $p^\pi(x)$ and $x \in V_\pi^{p^\pi(x)}([0,T],\mathbb{R})$ then:

$$[x]_\pi^{(q)}(t) = \begin{cases} 0 & \text{if } q > p^\pi(x) \\ 0 \leq [x]_\pi^{(q)} < \infty & \text{if } q = p^\pi(x) \\ \infty & \text{if } q < p^\pi(x) \end{cases} . \tag{2}$$

- For two continuous functions x, y if $p^\pi(x) < p^\pi(y)$, then x is smoother than y with respect to the sequence of partition π.

Spot volatility and realized volatility

- Unlike prices, volatility is not directly observable and must be estimated from prices.
- The *realized volatility* of a price process S over time interval $[t, t+\delta]$ sampled along the time partition π^n is defined as:

$$RV_{t,t+\Delta}(\pi^n) = \sqrt{\sum_{\pi^n \cap [t,t+\Delta]} \left(X(t_{i+1}^n) - X(t_i^n)\right)^2} = \sqrt{[X]_{\pi^n}(t+\Delta) - [X]_{\pi^n}(t)} \qquad (3)$$

  where $X = \log S$.
- The realized variance is sum of square increments of log price $= \left(RV_{t,t+\Delta}(\pi^n)\right)^2$.
- If the price S_t follows a stochastic volatility model with **instantaneous volatility** σ_t:

$$dS_t = \sigma_t S_t dB_t + \mu_t S_t dt$$

  then realized variance converges to the quadratic variation of $\log S$ (also called '*integrated variance*') as sampling frequency increases:

$$RV_t(\pi^n)^2 \underset{n\to\infty}{\overset{\mathbb{P}}{\to}} IV_t := \int_0^t \sigma_u^2 du, \qquad RV_{t,t+\Delta}(\pi^n) \underset{n\to\infty}{\overset{\mathbb{P}}{\to}} \sqrt{IV_{t,t+\Delta}} = \sqrt{\int_t^{t+\Delta} \sigma_u^2 du}. \qquad (4)$$

No Hölder norm convergence

- Sample paths of BM are almost surely continuous.
- $\exists$ **piecewise linear** functions B^n which converge to BM

$$B^n \to BM.$$

- (Hölder) roughness of B^n is of order 1 Vs. (Hölder) roughness of BM is $\frac{1}{2}-$.
- (Hölder) roughness of B^n $\not\to$ (Hölder) roughness of BM.

‘Volatility is rough’: Gatheral et al. (2018)

- Linear regression to compute the roughness of S&P500 realized volatility:

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

The exponent ξ_q are shown to behave linearly in q: $\xi_q \approx Hq$ with $\hat{H} = 0.13$.

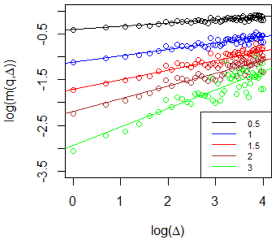

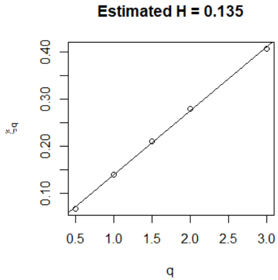

Based on this they propose a fractional SDE for volatility: $d\log\sigma^2_t \approx \eta dB^H_t$.

'Things we think we know': Rogers (2019)

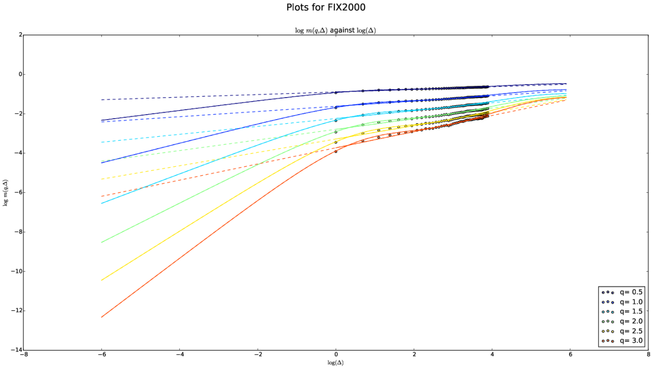

Rogers (2019) showed that roughness estimators based on linear-regression of $p-$th variation have poor accuracy: they exhibit similar behavior over a range of time scales even in a simple Brownian OU model for volatility (so: not rough!), so this cannot be taken at face value as evidence of 'rough volatility'.
Similar evidence of the lack of accuracy of such linear-regressions in shown in Takabatake, Fukusawa and Westphal (2021).

p-th variation as roughness: Theory Vs reality

Lemma

For any sequence of partition π *and for any* $x \in C^0([0,T],\mathbb{R})$,

$$[x]_\pi^{(q)}(t) = \begin{cases} 0 & \text{if } q > p^\pi(x) \\ 0 \le [x]_\pi^{(q)} < \infty & \text{if } q = p^\pi(x) \\ \infty & \text{if } q < p^\pi(x) \end{cases}. \tag{5}$$

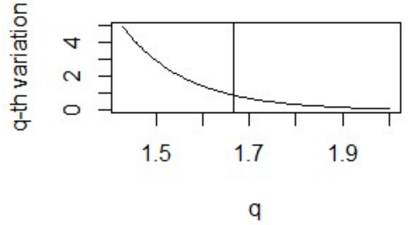

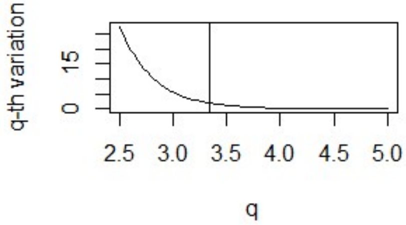

Practical example: N=250000 **Left:** fBM with $H = 0.6$ **Right:** fBM with $H = 0.3$.

Estimating the roughness of a path:
Normalized p-th variation

Normalized p-th variation

We propose an alternative estimator for measuring the roughness of a path from discrete observations.

Definition (Normalized p-th variation along a partition sequence (Cont & Das))

Let π be a sequence of partitions of $[0, T]$ with mesh $|\pi^n| \to 0$. We say a path $x \in C^0([0, T], \mathbb{R})$ has finite normalized p-th variation along a sequence of partition π if and only if there exists a continuous function $w(x, p, \pi) : [0, T] \to \mathbb{R}$ such that:

$$\forall t \in [0, T], \qquad w^n(x, p, \pi)(t) := \sum_{\pi^n \cap [0,t]} \frac{|x(t^n_{i+1}) - x(t^n_i)|^p}{[x]^{(p)}_\pi(t^n_{i+1}) - [x]^{(p)}_\pi(t^n_i)} \times (t^n_{i+1} - t^n_i) \xrightarrow{n \to \infty} w(x, p, \pi)(t). \tag{6}$$

- Notation: The class of all functions with finite Normalized p-th variation along π is denoted as $N^p_\pi([0, T], \mathbb{R})$.

Properties of normalized p-th variations

Theorem (Normalized p-th variations and roughness (Cont & Das))

Let $x \in C^0(([0,T],\mathbb{R}))$.
$\bar{p} = p^{\pi}(x)$ is the variation index of x. Then:

$$\forall t \in [0,T], \quad w(x,q,\pi)(t) = \begin{cases} \infty & \text{if } q > \bar{p} \\ 0 & \text{if } q < \bar{p} \end{cases}. \tag{7}$$

Furthermore, If $\bar{p}$-th variation $[x]_{\pi}^{(\bar{p})}$ is strictly increasing and the derivative $\frac{d}{du}[x]_{\pi}^{(\bar{p})}(u)$ exists and continuous then:

$$\forall p \in [1,\bar{p}] : \quad x \in N_{\pi}^{p}([0,T],\mathbb{R}) \quad \text{and,} \quad \forall t \in [0,T], \quad w(x,q,\pi)(t) = \begin{cases} \infty & \text{if } q > \bar{p} \\ t & \text{if } q = \bar{p} \\ 0 & \text{if } q < \bar{p} \end{cases}. \tag{8}$$

BM, fBM and stochastic integrals

- Let B be a Wiener process on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, $T > 0$ and $(\pi^n)_{n\geq 1}$ a sequence of partitions of $[0, T]$ with $|\pi^n| \log(n) \to 0$. Then:

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

- Let B^H be a fBM with Hurst parameter H, on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, $T > 0$. Then for any sequence of partitions π with mesh $|\pi| \to 0$, we have

$$w\left(B^H, \frac{1}{H}, \pi^n\right)(t) \xrightarrow{n\to\infty} t \text{ in probability.}$$

  Furthermore for dyadic partition $\mathbb{T} = (\mathbb{T}^n)_{n\geq 1}$ the convergence is in almost sure sense. ie.

$$\mathbb{P}\left(\lim_{n\to\infty} w\left(B^H, \frac{1}{H}, \pi^n\right)(t) = t\right) = 1.$$

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

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

Normalized p-th variation: examples

Schied (2016) and Schied and Mishura (2016) provide several examples of functions with prescribed p-th variation. Schied defines a class of functions $\mathcal{X}^p$

$$\mathcal{X}^H = \left\{ x \in C^0([0,1],\mathbb{R}) \;\middle|\; x(t) = \sum_{m=0}^{\infty} 2^{m(\frac{1}{2}-H)} \sum_{k=0}^{2^m-1} \theta_{m,k} e_{m,k}(t) \text{ for coeff. } \theta_{m,k} \in \{-1,+1\} \right\}.$$

The graph of $e_{m,k}$ is a wedge with height $2^{-\frac{m+2}{2}}$, width 2^{-m}, centred at $c = 2^{\frac{k-\frac{1}{2}}{2^m}}$.

Lemma (Generalized Takagi functions have unit normalized p-th variation)

Let $\mathbb{T}$ be the dyadic partition sequence. For any generalized Takagi function $x \in \mathcal{X}^H$ the normalized quadratic variation is given by t:

$$\forall x \in \mathcal{X}^H,\ \forall t \in [0,1], \qquad w(x, \frac{1}{H}, \mathbb{T}) = t.$$

Graph of $e^{\pi}_{m,k}$

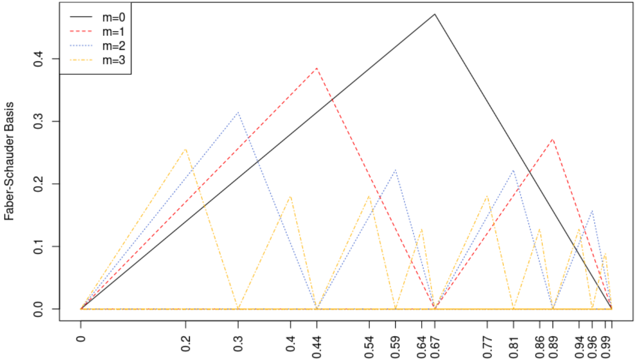

Figure: Plot of Schauder basis $e^{\pi}_{m,k}$ for $m = 0, 1, 2, 3$ along non-uniform non balanced doubly refining partition sequence π.

Normalized p-th variation statistics

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:

$$W(L, K, \pi, p, t, X) := \sum_{\pi^K \cap [0,t]} \frac{\left| X(t_{i+1}^K) - X(t_i^K) \right|^p}{\sum_{\pi^L \cap [t_i^K, t_{i+1}^K]} \left| X(t_{j+1}^L) - X(t_j^L) \right|^p} \times \left(t_{i+1}^K - t_i^K \right). \quad (9)$$

The definition of the statistic (9) involves two frequencies: a 'block' frequency K and a sampling frequency $L \gg K$.

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.

Estimating the roughness index

Lemma (Cont & Das)

The statistic (9) converges to the normalized p-th variation (6) as the sampling frequency L and block frequency K increase:

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

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 $p^{\pi}_{L,K}(X)$,

$$W(L,K,\pi,\widehat{p}^{\pi}_{L,K}(X),T,X) = T. \tag{10}$$

One can either fix a window length T or solve (10) in a least squares sense across several values of T.

Consistency of estimators

Theorem (Pathwise consistency of estimator)

For any balanced partition sequence π, under some regularity assumptions on X there exists sequence (L_n, K_n) such that $L_n > K_n$ and

$$\lim_{n\to\infty} \widehat{H}^{\pi}_{L_n,K_n}(X) = H^{\pi}(X) \ \ \textit{and,}$$

$$\lim_{n\to\infty} \widehat{p}^{\pi}_{L_n,K_n}(X) = p^{\pi}(X)$$

Balanced $\sim$ step sizes are asymptotically comparable.

Finite sample behaviour of the roughness estimator

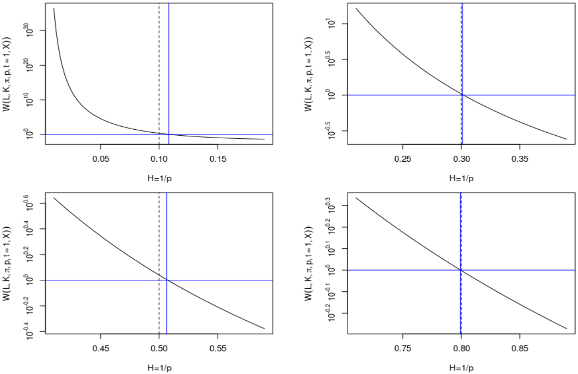

Figure: For $H \in \{0.1, 0.3, 0.5, 0.8\}$ we simulate a fBM with Hurst parameter H. The black line is log of normalized p-th variation statistics plotted against $H = 1/p$. The blue vertical line represents the estimated H value using the normalized p-th variation statistics ($K = 300$, $L = 300 \times 300$), whereas the green line represents the true value of H, with which we simulated the fBM.

FBM: Distribution of estimator for Hurst index

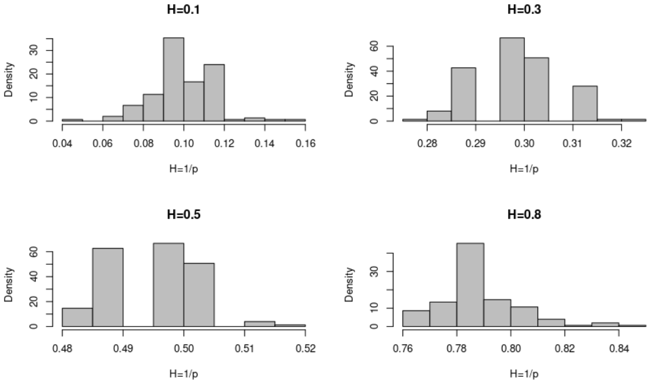

Figure: For $H \in \{0.1, 0.3, 0.5, 0.8\}$ we simulate a fBM with Hurst parameter H. Then use our normalized p-th variation statistics to estimate H with ($K = 300$, $L = 300 \times 300$). The histogram is generated from 150 independent runs.

fBM with Hurst index $H = 0.1$

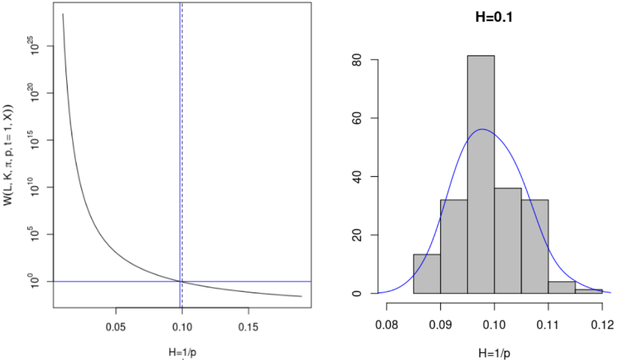

Figure: fBM simulated with $H = 0.1$. **Left:** The log of normalized p-th variation statistic is plotted against $H = 1/p$ in black. The blue vertical line represents the estimated roughness index $\hat{H}_{L,K}$ (with $M = 2000, L = 2000 \times 2000$), whereas the green line represents for true Hurst index $H = 0.1$. **Right:** Histogram of estimated roughness index $\hat{H}_{L,K}$ generated by simulating $n = 150$ independent fractional Brownian motion with Hurst parameter 0.1. The blue line represents the corresponding kernel plot generated by Gaussian density.

Summery statistics for roughness index

For $K = 300$ and $L = 300 \times 300$, total simulation size 150:

H	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
0.1	0.0450	0.0920	0.1030	0.1009	0.1100	0.1440
0.3	0.2730	0.2940	0.2980	0.2976	0.3020	0.3180
0.5	0.4820	0.4940	0.4980	0.4978	0.5020	0.5140
0.8	0.7570	0.7820	0.7900	0.7891	0.7940	0.8220

For $K = 2000$ and $L = 2000 \times 2000$, total simulation size 150:

H	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
0.1	0.086	0.096	0.099	0.099	0.103	0.117

Dependence on K

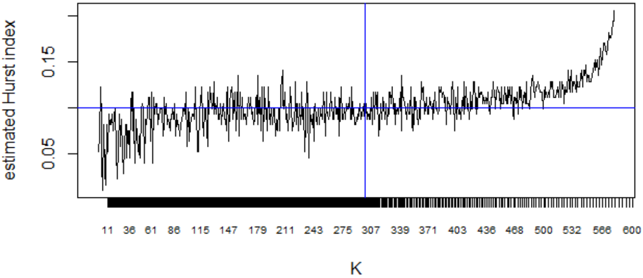

Figure: The solid line represents the estimated p-th variation statistic $W(L = 300 \times 300, K, \pi, q, t = 1)$ plotted against different values of K for a simulated fBM with $H = 0.1$. The blue vertical line represents for $K = 300, L = 300 \times 300$.

$$W(L, K, \pi, q, t) := \sum_{\pi^K} \frac{\left|x(t_{i+1}^K) - x(t_i^K)\right|^q}{\sum_{\pi^L \cap [t_i^K, t_{i+1}^K]} \left|x(t_{j+1}^L) - x(t_j^L)\right|^q} \times \left(t_{i+1}^K \wedge t - t_i^K \wedge t\right)$$

Highlights of roughness estimators

- Pathwise (Model free) estimator with no prior assumption on underlying distributions.
- Can estimate roughness of data observed on a irregular time scale π^n.
- Scale invariant.
- Invariant under smooth transformations.

Roughness index of realized volatility: Numerical experiments

Simulated Brownian diffusion model

Example

Consider the following price process where volatility follows a simple Brownian diffusion:

$$dS_t = \sigma_t S_t dB_t, \quad \text{with} \quad \sigma_t = |B'_t|, \tag{11}$$

where S_t is the price of the underline asset at time t and B_t, B'_t are two Brownian paths independent of each other.

Realized volatility and estimation

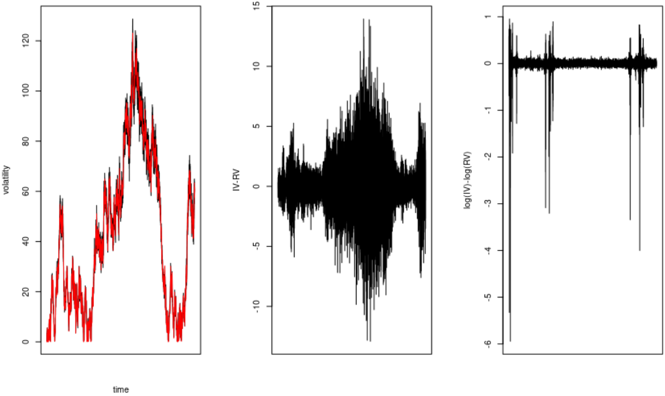

Figure: Simulation model: $\sigma_t = \mid B_t \mid, dS_t = S_t \sigma_t dB'_t$, where B_t and B'_t are Brownian motions independent of each other. **Left:** The red line is the plot of instantaneous volatility σ_t whereas the black line represents realized volatility RV_t from Model 11. **Right:** Corresponding estimation error for the simulated sample-path.

Roughness of estimated-realized vs instantaneous volatility

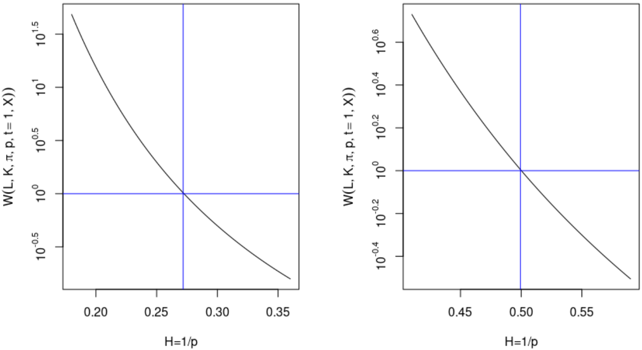

Figure: **Left:** Estimated roughness index $\hat{H}_{L,K}$ (via normalized p-variation statistic with $K = 500, L = 500 \times 500$), for realized volatility derived from a Brownian diffusion model. **Right:** Estimated roughness index $\hat{H}_{L,K}$ for instantaneous volatility of the same price path.

Varying the window size K: realized volatility

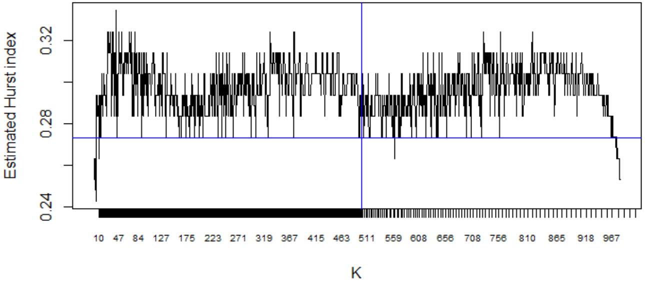

Figure: The solid black line represents the estimated p-th variation statistic $W(L = 500 \times 500, K, \pi, q, t = 1)$ plotted against different values of K for the realized volatility shown in Figure 6. The blue vertical line represents for $K = 500$, $L = 500 \times 500$ whereas the blue horizontal line represents $H = 0.273$.

$$W(L, K, \pi, q, t) := \sum_{\pi^K} \frac{\left|x(t_{i+1}^K) - x(t_i^K)\right|^q}{\sum_{\pi^L \cap [t_i^K, t_{i+1}^K]} \left|x(t_{j+1}^L) - x(t_j^L)\right|^q} \times \left(t_{i+1}^K \wedge t - t_i^K \wedge t\right)$$

Varying window size K: spot volatility

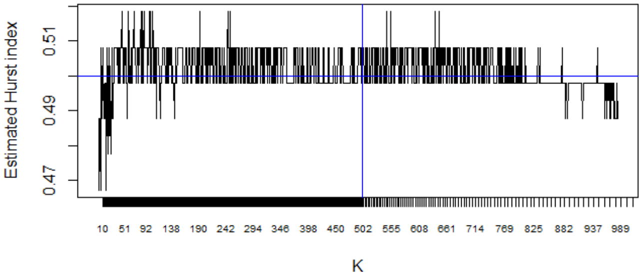

Figure: The solid black line represents the estimated p-th variation statistic $W(L = 500 \times 500, K, \pi, q, t = 1)$ plotted against different values of K for the instantaneous volatility shown in Figure 6. The blue vertical line represents for $K = 500$, $L = 500 \times 500$ whereas the blue horizontal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

Gatheral et al.'s Linear regression method: spot volatility

Applying the log-regression method to **instantaneous volatility** gives $H \approx 0.5$

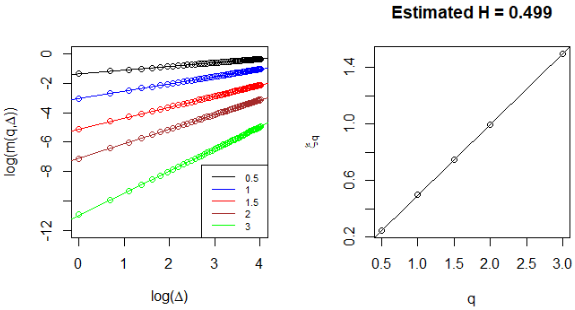

Figure: **Left:** Scaling analysis of instantaneous volatility simulated using a Brownian stochastic volatility model using the method of Gatheral et al. (2014). **Right:** linear regression coefficients ξ_q as a function of q. The estimated roughness index is $\hat{H} = 0.499$.

Gatheral et al.'s linear-regression method: realized vol

Applying the linear-regression method to **realized vol** for same data gives $H = 0.34$!

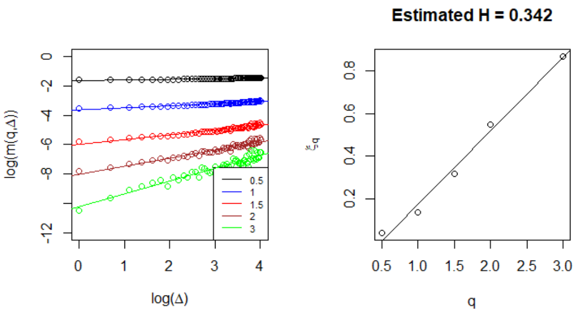

Figure: **Left:** Scaling analysis of realized volatility estimated from simulated paths from the Brownian stochastic volatility model using the method of Gatheral et al. (2014). **Right:** linear regression coefficients ξ_q as a function of q. The estimated roughness index is $\hat{H} = 0.342$.

Simulated price model: Brownian OU process

Example

Take a stochastic volatility model where the volatility follows a Ornstein–Uhlenbeck process.

$$dS_t = S_t \sigma_t dB_t, \quad \text{where,}$$
$$\sigma_t = \sigma_0 e^{Y_t}, \quad dY_t = -\gamma Y_t dt + \theta dB'_t \tag{12}$$

B_t and B'_t are two Brownian motions independent of each other. For our simulation $\sigma_0 = 1$, $Y_0 = 0$ and $\gamma = \theta = 1$.

Price process and estimation error

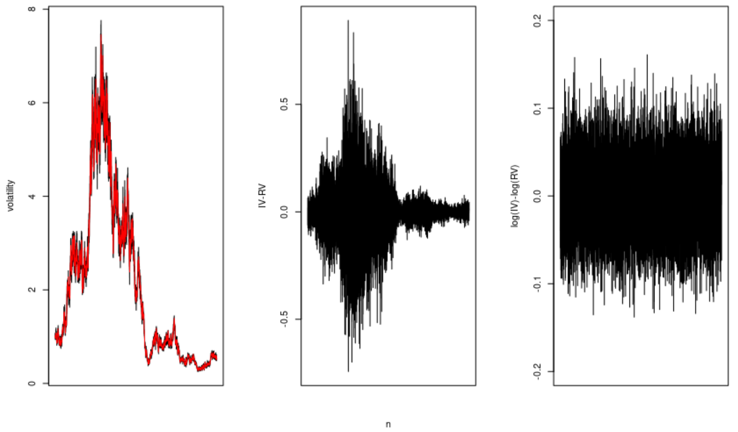

Figure: Left: Black = realized volatility. Red= instantanenous (spot) volatility. Right: estimation error for OU stochastic volatility model.

Kernel plot

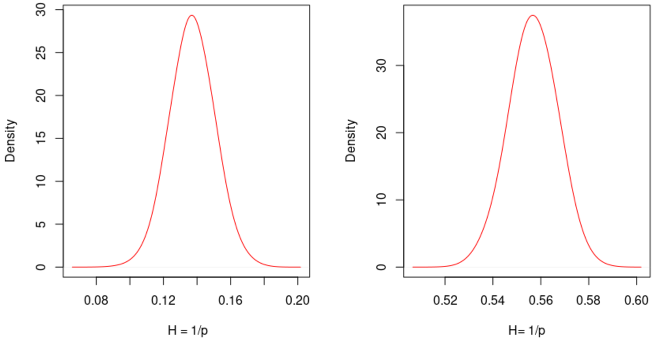

Figure: Distribution of the estimated roughness index $\hat{H}_{L,K}$ for ($K = 300, L = 300 \times 300$) across 2500 independent simulations for the OU-SV model (12). True value is $H = 0.5$.
Left: realized volatility. Right: instantaneous volatility respectively.

Summary Statistics

The following table provides summary statistics across 2500 independent samples for the roughness exponent estimator $\hat{H}_{L=300\times300,K=300}$ for realized volatility and spot volatility.

	Realized volatility	Instantaneous volatility
Min.	0.087	0.528
1st Quantile	0.128	0.552
Median	0.136	0.556
Mean	0.137	0.557
3rd Quantile	0.148	0.563
Max.	0.181	0.581

Fractional stochastic volatility process

Example (Fractional OU volatility process)

Consider the following price process where the volatility is coming from a fractional Ornstein–Uhlenbeck process.

$$dS_t = \sigma_t S_t dB_t, \quad \text{where,}$$
$$\sigma_t = \sigma_0 e^{Y_t}; \;\; dY_t = -\gamma Y_t dt + \theta dB_t^H, \tag{13}$$

where B and B^H are respectively Brownian motion and fractional Brownian motion correspond to Hurst index $H \in (0, 1)$.

Then spot volatility has roughness index $0 < H < 1$
We use $\gamma = \theta = \sigma_0 = 1$ and $Y_0 = 0$.

Fractional OU volatility process

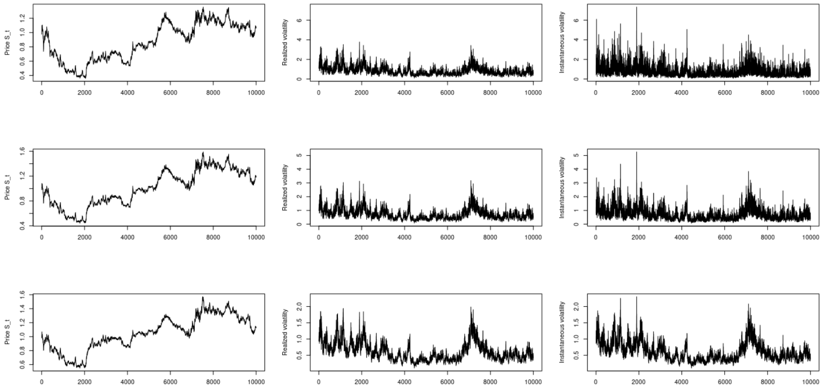

Figure: **Left:** OU process with H={0.05,0.1,0.2} respectively, **Middle:** Realized Vol., **Right:** Instantaneous Vol.

Fractional OU volatility process

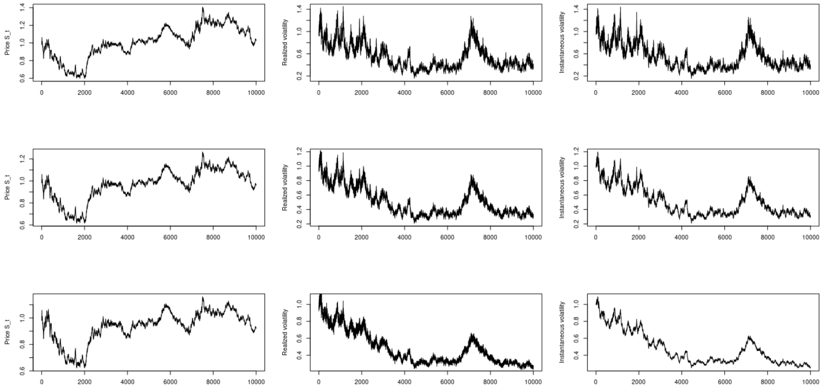

Figure: **Left:** OU process with H={0.3,0.4,0.5} respectively, **Middle:** Realized Vol., **Right:** Instantaneous Vol.

Fractional OU volatility process

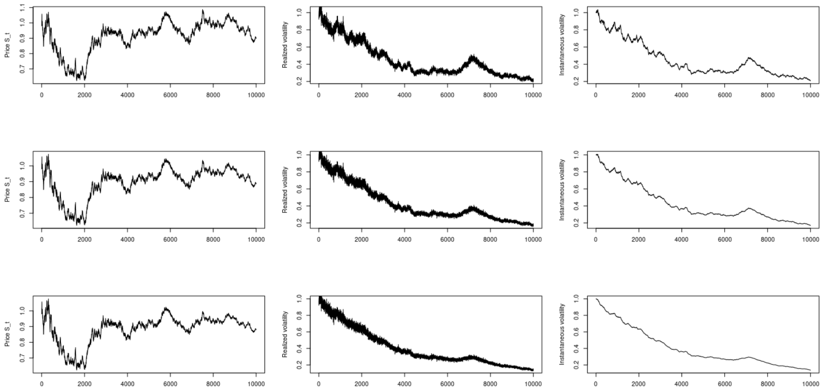

Figure: **Left:** OU process with H={0.6,0.7,0.8} respectively, **Middle:** Realized Vol., **Right:** Instantaneous Vol.

Summary statistics

The roughness index $\hat{H}_{L,K}$ (with, $L = 300 \times 300, K = 300$) of instantaneous and realized volatility are compared in the following table.

H	$\hat{H}_{L,K}$ of Instantaneous volatility	$\hat{H}_{L,K}$ of 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

Comparing $\hat{H}$ for realized Vol. and instantaneous Vol.

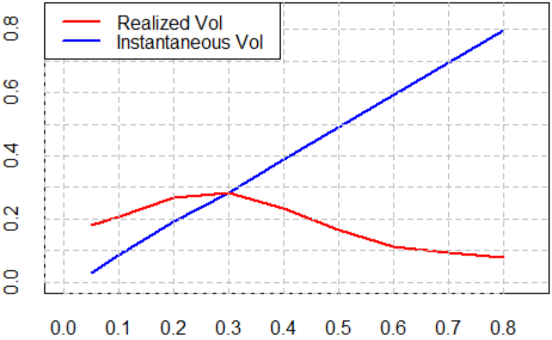

Figure: Estimated values of roughness exponent $\hat{H}$ from high-frequency realized volatility for a fractional-OU stochastic volatility model with different values of H. X axis: True Hurst index H. Y axis: Estimated roughness index $\hat{H}_{L,K}$

Comparing $\hat{H}$ of RV and IV for 200 simulations

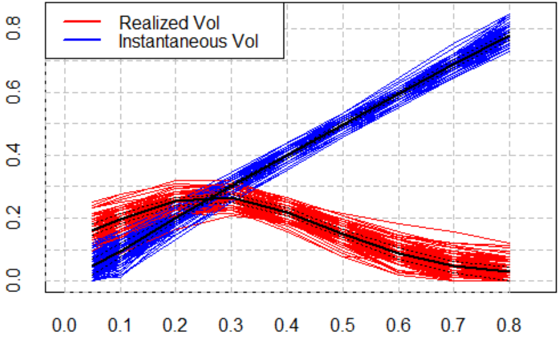

Figure: Estimated values of roughness exponent $\hat{H}$ from high-frequency realized volatility for 200 simulated fractional-OU stochastic volatility model with different values of H.

Application to high-frequency financial data

AAPL (high frequency) stock price data: year 2016

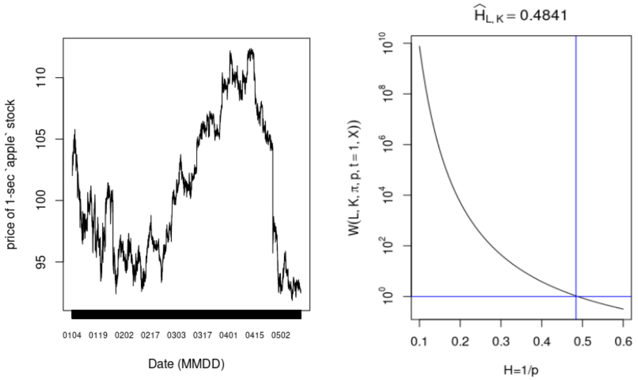

Figure: Left: plot of 1-sec price of AAPL 04/Jan/2016 - 11/May/2016 (90 days). Right: Estimation of $\hat{H}_{L,K}$(via normalized p-variation statistic with $L = 1400 \times 1400$, $K = 1400$) for the apple stock price plotted in left figure.

Roughness: AAPL (2016) realized volatility

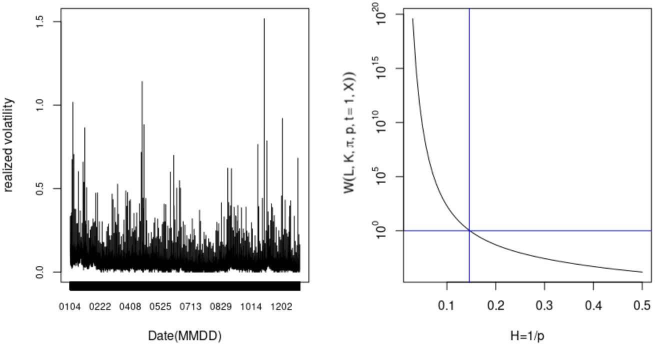

Figure: Left: plot of 1-min realised volatility of 'apple'(year 2016). Right: Estimation of $\hat{H}_{L,K}$(via normalized p-variation statistic with $L = 310 \times 310$, $K = 310$) for the 1-min realised volatility (estimated $H \in [.08 - .22]$).

Roughness: S&P 500 (Source: Oxford-Man Institute)

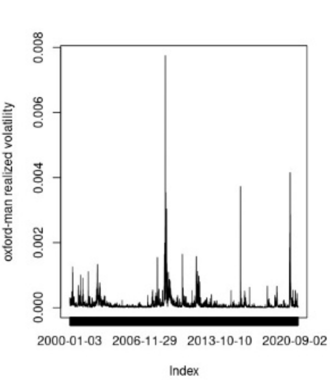

Figure: Left: plot of 5-min realised volatility . Right: normalized p-variation statistic with $L = 70 \times 70, K = 70$ for SPX 5-min realised volatility, as a function of $H = 1/p \in [.05 - .25]$.

Rough volatility ... or estimation noise?

- When one has a powerful hammer, everything looks like a nail...
- but it is important to check the robustness of 'stylized facts' to estimation error before jumping into complex models to 'explain' them.
- We introduce a nonparametric method for estimating the roughness of a path based on the notion of **normalized p-th variation**.
- For stochastic-volatility diffusion models driven by Brownian motion (so $H = 1/2$), the realized volatility exhibits an estimated roughness index $\hat{H}_{L,K} \approx 0.3$ so seems to exhibit significantly 'rougher' behaviour than spot volatility, both in terms of normalized p-th variation and in terms of the linear-regression method used by Gatheral et al. (2018). In this case roughness in realized vol is a pure "statistical artefact" i.e. entirely attributable to estimation error.
- These results suggest that the regression method is not robust to estimation noise: one cannot take the roughness observed in realized volatility as evidence of similar behaviour in spot volatility, as implicitly assumed in the 'rough volatility' literature.
- As shown in fOU example, 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.
- 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 diffusion models explain the empirical observations just as well.