

Statistics of Financial Markets

Chapter 11: Fractal markets hypothesis and long memory



Daniel Traian PELE

Bucharest University of Economic Studies
IDA Institute Digital Assets
Blockchain Research Center
AI4EFin Artificial Intelligence for Energy Finance
Romanian Academy, Institute for Economic Forecasting
MSCA Digital Finance

Acronyms used in this material

- **ACF**: Autocorrelation Function
- **AI**: Artificial Intelligence
- **AMH**: Adaptive Markets Hypothesis
- **AR**: AutoRegressive (model); Abnormal Return
- **ARFIMA**: AutoRegressive Fractionally Integrated Moving Average
- **ARMA**: AutoRegressive Moving Average (model)
- **BET**: Bucharest Exchange Trading (index)
- **BET-TR**: BET Total Return (index)
- **BVB**: Bursa de Valori București (the Bucharest Stock Exchange)
- **DFA**: Detrended Fluctuation Analysis
- **EMH**: Efficient Market Hypothesis
- **EODHD**: EOD Historical Data (market-data provider)
- **FIGARCH**: Fractionally Integrated GARCH
- **FMH**: Fractal Market Hypothesis
- **GARCH**: Generalised AutoRegressive Conditional Heteroskedasticity
- **GPH**: Geweke–Porter-Hudak (log-periodogram estimator)
- **HAR**: Heterogeneous AutoRegressive (model)
- **IGARCH**: Integrated GARCH
- **MA**: Moving Average
- **OLS**: Ordinary Least Squares
- **SE**: Standard Error
- **SFE**: Statistics of Financial Markets, Exercises (Quantlet collection of the textbook)
- **US**: United States

Today's Question and Route

- **Question**: do today's returns remember what happened months ago?
  - the fractal market hypothesis adds a second theme: who trades, and at which horizon
  - Chapter 7: daily returns are almost uncorrelated at short lags; Chapters 8 and 9: volatility is persistent
  - this chapter asks about dependence over **long** lags and over **many** time scales
- **Route** of the chapter
  - the fractal market hypothesis of Peters; self-similarity and fractals
  - Hurst and the Nile; the R/S statistic; long memory, fractional Brownian motion, ARFIMA
  - estimators (R/S, Lo, DFA, GPH), Monte Carlo bands, spurious long memory
  - long memory in volatility; rolling Hurst exponents on six markets; consequences for risk

Learning Outcomes

- State the fractal market hypothesis and compare it with the efficient and the adaptive market hypotheses
- Define self-similarity, the Hurst exponent, fractional Brownian motion, fractional Gaussian noise and ARFIMA(0,d,0)
- Compute the R/S statistic step by step and estimate H with R/S, DFA and the GPH regression; apply Lo's modified R/S test
- Judge an estimate against Monte Carlo bands and recognise spurious long memory (breaks, volatility clustering)
- Separate long memory in returns from long memory in volatility on real data, with rolling estimates
- Explain how H changes the scaling of risk with the horizon

Reading and Tools

- Textbook: Franke, Härdle & Hafner (2019), Ch. 14 (14.1–14.5 definition, fractional integration, self-similarity, tests and estimators of long memory; 14.6 long memory models)
  - the fractal market hypothesis: Peters (1994); a survey of long memory: Baillie (1996)
- Python Quantlets of this chapter: Quantlets/Ch_11
  - ported from the SFE Quantlets SFEfbmplot, SFEfgnacf and SFEarfima
- Lecture notebook: open in Google Colab
- Video courses: Efficiency of cryptocurrency markets; Applied Time Series Analysis with Python
- Further reading on crashes and power laws in Romanian and crypto markets: Pele & Mazurencu-Marinescu (2012); Pele & Mazurencu-Marinescu-Pele (2019)

Limitations of the EMH

- EMH (efficient market hypothesis, Fama (1970); Chapter 7): prices reflect the available information, so returns cannot be forecast from past returns
  - the random walk of prices implies independent increments: in this chapter, $H = 0.5$
- Facts that the random walk with Normal increments does not explain
  - heavy tails (Chapters 2, 3, 5); volatility clustering (Chapters 8, 9)
  - crashes such as that of 19 October 1987 (a few slides ahead)
- The EMH says nothing about **who** trades and at which **horizon**: all investors are treated as one representative investor

Edgar Peters and the Fractal Market Hypothesis

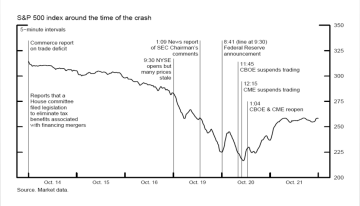
- **FMH** (fractal market hypothesis) Peters (1994): the market is made of investors with **many investment horizons**, from minutes to decades
- The five statements of Peters
  - the market is stable when it contains investors with many different horizons: this ensures liquidity
  - short horizons react to sentiment and technical information; long horizons react to fundamental information
  - if the fundamental information becomes doubtful, long-term investors leave or start to trade short term: horizons become uniform and the market becomes unstable
  - prices combine short-term technical trading and long-term fundamental valuation
  - an asset without a link to the economic cycle has no long-term trend: short-term trading and liquidity dominate

Investment Horizons and Liquidity: an Example

- A share falls by 3% in one morning
  - a day trader (horizon: hours) sees a large move and sells
  - a pension fund (horizon: ten years) sees a small move relative to its horizon and buys at a lower price
- The long-horizon investor supplies the liquidity that the short-horizon investor demands: the price stays orderly
- If the pension fund also starts to fear the next hours, nobody buys: there is no liquidity and the fall accelerates
- **Fractal**: the same pattern of behaviour repeats at every horizon, with statistically similar returns at different scales

The 1987 Crash in the Fractal Market Hypothesis

- 19 October 1987: the S&P 500 fell by about 20% in one day Carlson (2007)
- The FMH reading
  - portfolio insurance and program trading made long-term portfolios sell like short-term traders
  - all horizons became one horizon: buyers disappeared
- Testable consequence: in crises, short horizons dominate and the statistical structure across scales changes



The S&P 500 around the crash of October 1987

Chart: Mark Carlson, Federal Reserve Board (2006); public domain; Wikimedia Commons

Testing the FMH: Scaling, Horizons and Liquidity

- The FMH is a verbal theory; its statistical consequences can be tested
  - **self-similarity**: the distribution of h-day returns looks the same after rescaling by h^H (Section 2)
  - **memory**: the Hurst exponent H measures dependence over long horizons (Sections 3–5)
  - **change**: in crises, H and the share of short horizons should change (Section 8)
- Landmark evidence: Kristoufek (2012) and Kristoufek (2013) on the 2008 crisis
  - local Hurst exponents and wavelet power of US and European indices
  - the most turbulent periods coincide with a dominance of short investment horizons, as the FMH predicts

EMH, FMH and AMH

	EMH	**FMH**	**AMH**
Source	Fama (1970)	Peters (1994)	Lo (2004)
Investors	one representative investor	many horizons	species that compete and adapt
Stability	prices always fair	diversity of horizons gives liquidity	changes with the environment
Memory	$H = 0.5$	H may differ from 0.5	H varies over time
Crises	outside the theory	collapse of horizons	mass extinction of strategies

- The three are not rival tests: the FMH and the AMH explain when and why the EMH approximation fails

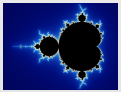
Benoît Mandelbrot and the Measurement of Roughness

- Mandelbrot (1963): cotton prices have heavy tails and look alike at daily, monthly and yearly scales
  - the α-stable distributions of Chapter 3
- Mandelbrot (1967): the measured length of a coast grows as the ruler shrinks; the growth rate defines a **fractal dimension**
- Mandelbrot & Van Ness (1968): fractional Brownian motion, the model of this chapter
- **Fractal**: an object whose parts resemble the whole, exactly or in distribution



Benoît Mandelbrot (1924–2010), Paris, 2006

Photo: David Monniaux (2006); CC BY-SA 3.0; Wikimedia Commons



The Mandelbrot set

Image: Wolfgang Beyer; CC BY-SA 3.0; Wikimedia Commons

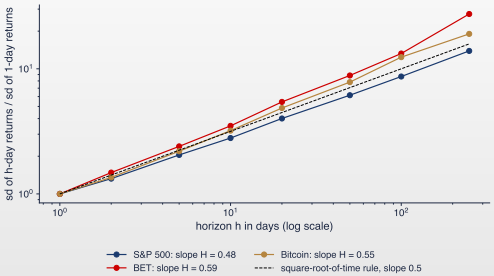
Fractal Dimension

- **Similarity dimension**: an object made of N copies of itself, each scaled by $1/s$, has $D = \ln N / \ln s$
  - a segment: 2 halves, $D = \ln 2 / \ln 2 = 1$; a square: 4 quarters of side $1/2$, $D = \ln 4 / \ln 2 = 2$
  - the Koch curve: 4 copies scaled by $1/3$, $D = \ln 4 / \ln 3 = 1.262$
  - the Sierpinski triangle: 3 copies scaled by $1/2$, $D = \ln 3 / \ln 2 = 1.585$
- The west coast of Britain: $D \approx 1.25$ Mandelbrot (1967)
- The graph of a price path: $D = 2 - H$
  - Brownian motion: $D = 1.5$; a smoother, trending path ($H > 0.5$) has $D < 1.5$

Self-Similar Processes

- **Self-similar process** Franke, Härdle & Hafner (2019), Def. 14.1: $X_{ct} \stackrel{d}{=} c^H X_t$ for every $c > 0$; H is the **Hurst exponent**
  - $\stackrel{d}{=}$: equal in distribution; stretching time by c equals stretching values by c^H
- Brownian motion (Chapter 4) is self-similar with $H = 1/2$: $W_{ct} \stackrel{d}{=} \sqrt{c}\, W_t$
- Consequence for log prices with self-similar increments
  - $\mathrm{sd}(r_t^{(h)}) = h^H \,\mathrm{sd}(r_t)$, where $r_t^{(h)}$ is the h-day return and sd the standard deviation
  - the square-root-of-time rule $\sqrt{h}$ is the special case $H = 1/2$
- A first estimator of H: the slope of $\ln \mathrm{sd}(r^{(h)})$ on $\ln h$

The Scaling of *h*-Day Returns



- Standard deviation of non-overlapping *h*-day returns divided by the daily one, $h = 1, \ldots, 250$, log scales; data to 18 September 2026
- Slopes: S&P 500 0.48, BET 0.59, Bitcoin 0.55; the square-root rule has slope 0.5

SFM_ch11_self_similarity

Interpretation: Scaling Exponents

- S&P 500: 10-day risk is 2.80 times the daily risk, below $\sqrt{10} = 3.16$
  - slight mean reversion: daily returns have a negative first-order autocorrelation, -0.099
- BET: 3.50 times, above $\sqrt{10}$: positive autocorrelation
  - thin trading in the early years makes the index adjust slowly (Chapter 7)
- Bitcoin: 3.20 times, close to $\sqrt{10}$
- Caution: there are only a few dozen non-overlapping 250-day blocks; the slope is imprecise, so it needs a proper estimator and a band

Harold Edwin Hurst and the Nile

- Hurst (1880–1978), a British hydrologist who worked in Egypt from 1906, designed reservoirs on the Nile
  - the capacity needed to deliver the average flow every year equals the range of the cumulative deviations of the inflows from their mean
- Hurst (1951): for independent inflows this range grows like $n^{0.5}$; in river flows, rainfall and tree rings it grows like n^K with K about 0.73 on average
  - years of high flow follow years of high flow: **long memory**



H.E. Hurst (1953)

Photo: Elliott & Fry (1953); public domain; Wikimedia Commons



The Nilometer on Roda Island, Cairo, built in 861

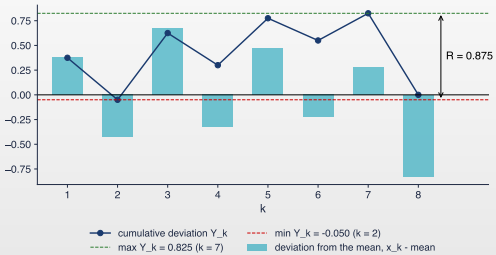
Photo: David Stanley (2013); CC BY 2.0; Wikimedia Commons

The R/S Statistic

- Block $x_1, \ldots, x_n$ with mean $\bar{x}$; **cumulative deviations** $Y_k = \sum_{j=1}^{k}(x_j - \bar{x})$, $k = 1, \ldots, n$
  - $Y_n = 0$ by construction
- **Range**: $R_n = \max_k Y_k - \min_k Y_k$; **standard deviation**: $S_n = \sqrt{\frac{1}{n}\sum_j (x_j - \bar{x})^2}$
- **R/S statistic** (rescaled range) Franke, Härdle & Hafner (2019), 14.4.1: $(R/S)_n = R_n / S_n$
  - dividing by S_n removes the unit: $(R/S)_n$ is comparable across assets
  - for i.i.d. data with finite variance, $(R/S)_n$ grows like $n^{1/2}$; under long memory like n^H, $H > 1/2$
- Robust to heavy tails: R and S are inflated by the same extreme values Mandelbrot & Wallis (1969)

Worked Example: R/S of Eight Returns

- Returns (%): 0.5; −0.3; 0.8; −0.2; 0.6; −0.1; 0.4; −0.7
- Mean $\bar{x} = 1.0/8 = 0.125$
- Deviations: +0.375, −0.425, +0.675, −0.325, +0.475, −0.225, +0.275, −0.825
- Y_k: 0.375, −0.050, 0.625, 0.300, 0.775, 0.550, 0.825, 0.000
- $R = 0.825 - (-0.050) = 0.875$; $S = \sqrt{1.915/8} = 0.489$
- $(R/S)_{\mathbf{8}} = 0.875/0.489 = 1.788$; one block gives one point, H needs many block sizes



SFM_ch11_rs_estimators

From R/S to the Hurst Exponent

- Steps Mandelbrot & Wallis (1969)
  - choose block sizes n between 10 and $N/4$, evenly spaced on a log scale
  - split the N returns into $\lfloor N/n \rfloor$ non-overlapping blocks; compute R/S in each block and average: $(R/S)_n$
  - regress $\ln(R/S)_n = \ln c + H \ln n + e_n$ by OLS (ordinary least squares): the slope $\hat{H}$ estimates H
  - N: the number of returns; n: the block size; c: a constant; e_n: the regression error
- The log-log plot is the main diagnostic: a straight line supports a single scaling exponent
- Applied to returns (stationary), not to prices: for prices, $(R/S)_n$ grows like n

Interpreting H

H	increments	memory	path ($D = 2 - H$)
$0 < H < 0.5$	negatively correlated	anti-persistent: reversals	rough, $D > 1.5$
$H = 0.5$	uncorrelated	none (random walk)	$D = 1.5$
$0.5 < H < 1$	positively correlated	persistent: trends	smoother, $D < 1.5$

- H is not a probability: $H = 0.7$ does not mean "70% chance that the sign repeats"
- Two effects named by Mandelbrot & Wallis (1968)
  - **Noah effect**: extreme values (heavy tails, the flood); **Joseph effect**: long runs of high or low values ("seven years of plenty, seven years of famine")
  - H measures the Joseph effect; the tail index of Chapter 5 measures the Noah effect

Small Samples: the Bias of R/S

- For i.i.d. data, $(R/S)_n$ approaches $c\, n^{1/2}$ only for large n; for small n it grows faster Anis & Lloyd (1976)
  - the fitted slope over $n = 10, \ldots, N/4$ is above 0.5 even without any memory
- Monte Carlo, 500 i.i.d. Normal series
  - mean $\hat{H}_{R/S}$: 0.588 for $N = 250$, 0.562 for $N = 1000$, 0.547 for $N = 5000$
  - 95% band for $N = 1000$: $[0.50, 0.63]$
- Rule: compare $\hat{H}$ with the distribution of $\hat{H}$ under the null hypothesis for the same N, not with 0.5 (Section 5)

Short and Long Memory

- Stationary X_t with autocorrelations $\rho(k)$ (ACF, autocorrelation function); **long memory** Franke, Härdle & Hafner (2019), 14.1:
  - $\sum_k |\rho(k)| = \infty$: the autocorrelations are not summable
  - the usual form: hyperbolic decay $\rho(k) \sim C\, k^{2d-1}$ as $k \to \infty$, with the **memory parameter** $0 < d < 0.5$; it implies non-summability, the converse does not hold in general
  - in frequency: the spectral density has a pole at zero, $f(\lambda) \sim C\, \lambda^{-2d}$ as $\lambda \to 0$
  - $C > 0$: a constant; $\sim$: the ratio of the two sides tends to 1; $f(\lambda)$: the share of variance at frequency λ (long cycles: λ near 0)
- **Short memory**: ARMA (autoregressive moving average) models, e.g. AR(1) with $\rho(k) = \phi^k$
  - exponential decay: $\sum_k |\phi|^k = 1/(1 - |\phi|) < \infty$
- Link with the Hurst exponent: $d = H - 1/2$

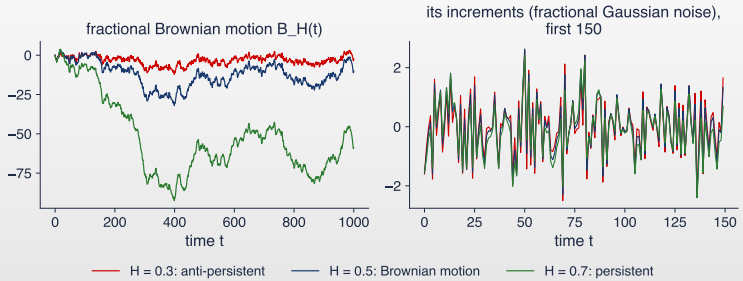
Worked Example: Autocorrelation Decay under Long and Short Memory

- Long memory: fractional Gaussian noise (next slides) with $H = 0.8$
  - $\rho(1) = 2^{2H-1} - 1 = 2^{0.6} - 1 = 0.516$
  - $\rho(10) = 0.191$, $\rho(50) = 0.100$, $\rho(100) = 0.076$
- Short memory: AR(1) with the same first autocorrelation, $\phi = 0.516$
  - $\rho(10) = \phi^{10} = 0.0013$; $\rho(100) = \phi^{100}$, practically zero
- Same short-run behaviour, very different long run: after 100 days, a shock still matters only under long memory

Fractional Brownian Motion

- **Fractional Brownian motion** (fBm) $B_H(t)$, $0 < H < 1$ Mandelbrot & Van Ness (1968); Franke, Härdle & Hafner (2019), Def. 14.2: a Gaussian process with continuous paths, self-similar with exponent H and with stationary increments
- Properties (Franke, Härdle & Hafner (2019), Th. 14.1)
  - $B_H(0) = 0$, $E[B_H(t)] = 0$, $E[B_H(t)^2] = \sigma^2 |t|^{2H}$ (σ^2: the variance at $t = 1$)
  - $\mathrm{Cov}(B_H(t), B_H(s)) = \frac{\sigma^2}{2} \left(|t|^{2H} + |s|^{2H} - |t - s|^{2H} \right)$
- $H = 1/2$: Brownian motion, independent increments; $H \neq 1/2$: the increments are correlated at all lags
- Simulation: exact, from the covariance matrix of the increments (circulant embedding, as in the notebook)

fBm Paths for Three Hurst Exponents



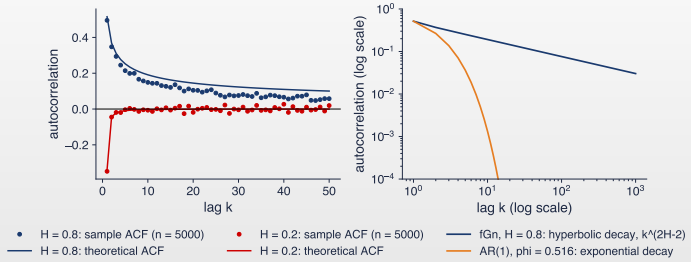
- The same random numbers, $H = 0.3, 0.5, 0.7$ (as in SFEfbmplot); right: the first 150 increments
- $H = 0.7$: long excursions and a wide range; $H = 0.3$: frequent reversals, a narrow range; lag-1 autocorrelation of the increments 0.32 and -0.23 (theory 0.32 and -0.24)

SFM_ch11_fbm_fgn

Fractional Gaussian Noise

- **Fractional Gaussian noise** (fGn): the increments $Y_k = B_H(k+1) - B_H(k)$, a stationary series Franke, Härdle & Hafner (2019), 14.3
  - $\rho(k) = \frac{1}{2}\left(|k+1|^{2H} - 2|k|^{2H} + |k-1|^{2H}\right)$
  - for large k: $\rho(k) \approx H(2H-1)\,k^{2H-2}$
- Three regimes
  - $1/2 < H < 1$: $\rho(k) > 0$ and not summable: long memory with $d = H - 1/2 \in (0, 1/2)$
  - $H = 1/2$: white noise
  - $0 < H < 1/2$: $\rho(k) < 0$, summable: anti-persistence
- fGn is the model behind R/S and DFA: for fGn data, both estimate H

The ACF of Fractional Gaussian Noise



- Left: sample and theoretical ACF for $H = 0.8$ and $H = 0.2$ (as in SFEfgnacf); $\rho(1) = 0.516$ and -0.340
- Right, log scales: hyperbolic decay (a straight line) against the exponential decay of an AR(1) with the same $\rho(1)$
- The sample ACF of a long-memory series lies below the theoretical one: estimating the mean removes part of the slow component

SFM_ch11_fbm_fgn

ARFIMA(0,d,0): Fractional Differencing

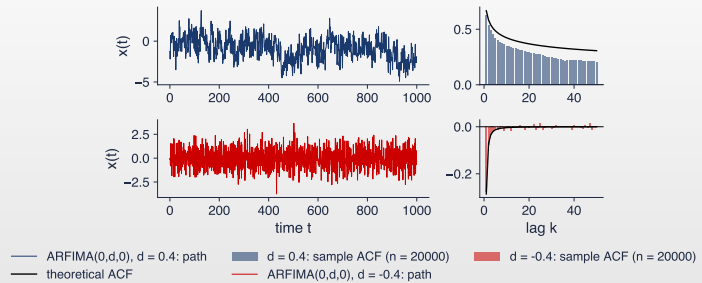
- **ARFIMA**(0,d,0) Granger & Joyeux (1980); Hosking (1981): $(1-L)^d X_t = \varepsilon_t$, ε_t i.i.d. $(0, \sigma^{\mathbf{2}})$, L the lag operator ($LX_t = X_{t-\mathbf{1}}$)
  - ARFIMA: autoregressive fractionally integrated moving average; $d = 0$: white noise, $d = 1$: random walk (Chapter 7)
- Binomial expansion Franke, Härdle & Hafner (2019), 14.2: $(1-L)^d = \sum_{j \geq \mathbf{0}} \pi_j L^j$, $\pi_{\mathbf{0}} = 1$, $\pi_j = \pi_{j-\mathbf{1}} (j - 1 - d)/j$
  - $X_t = \varepsilon_t - \sum_{j \geq \mathbf{1}} \pi_j X_{t-j}$: the whole past matters, with slowly decaying weights
  - MA(∞) form: $X_t = \sum_{j \geq \mathbf{0}} \psi_j \varepsilon_{t-j}$, $\psi_{\mathbf{0}} = 1$, $\psi_j = \psi_{j-\mathbf{1}} (j - 1 + d)/j$
- Autocorrelations: $\rho(k) = \dfrac{\Gamma(1-d)\,\Gamma(k+d)}{\Gamma(d)\,\Gamma(k+1-d)}$, i.e. $\rho(k) = \rho(k-1)\,\dfrac{k-1+d}{k-d}$
  - $\Gamma(\cdot)$: the gamma function, $\Gamma(n) = (n-1)!$ for integers; $\rho(1) = d/(1-d)$
  - for large k, $\rho(k) \sim C k^{\mathbf{2}d-\mathbf{1}}$: long memory for $0 < d < 1/2$

The Memory Parameter d

d	memory	mean reversion	variance	ACF
$(-0.5, 0)$	anti-persistent	yes	finite	hyperbolic, negative
0	short	yes	finite	exponential (ARMA)
$(0, 0.5)$	long, stationary	yes	finite	hyperbolic
$[0.5, 1)$	long, non-stationary	yes	infinite	hyperbolic
1	unit root	no	infinite	does not decay

- Source: Franke, Härdle & Hafner (2019), Table 14.1; $H = d + 1/2$ for the stationary range
- ARFIMA(p,d,q): ARMA(p,q) dynamics for the short run, d for the long run Franke, Härdle & Hafner (2019), 14.6.1
- Fractional integration fills the gap between the stationary $d = 0$ and the unit root $d = 1$ of Chapter 7

ARFIMA(0,d,0) with $d = 0.4$ and $d = -0.4$



- Left: 1000 simulated values (as in SFEarfima); right: sample ACF (20000 values) and theoretical ACF
- $d = 0.4$: $\rho(1) = 0.667$, $\rho(10) = 0.424$, $\rho(50) = 0.307$: slow waves around the mean; $d = -0.4$: $\rho(1) = -0.286$, then almost zero

SFM_ch11_arfima

Worked Example: ARFIMA(0, 0.3, 0)

- Weights of $(1-L)^{0.3}$: $\pi_1 = -d = -0.300$, $\pi_2 = \pi_1(1-d)/2 = -0.105$
  - MA weights: $\psi_1 = d = 0.300$, $\psi_2 = \psi_1(1+d)/2 = 0.195$
- Autocorrelations
  - $\rho(1) = 0.3/0.7 = 0.429$; $\rho(2) = \rho(1) \times 1.3/1.7 = 0.328$
  - $\rho(10) = 0.173$, $\rho(100) = 0.069$
- AR(1) with $\phi = 0.429$: $\rho(10) = 0.0002$
- The autocorrelation falls below 0.05 after 4 lags for the AR(1) and after 222 lags for ARFIMA; $H = d + 0.5 = 0.8$

Summary of the Long-Memory Methods

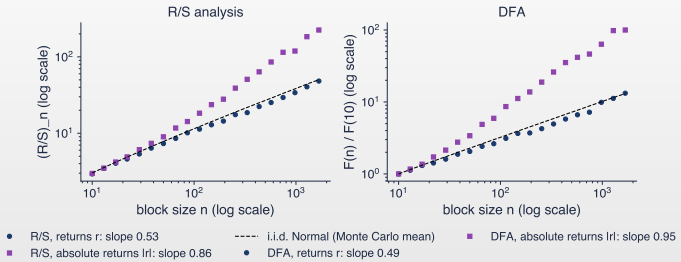
Tool	Domain	What it gives	Weak point
R/S	time	$\hat{H}$, slope of a log-log plot	biased upwards in small samples; sensitive to short memory
Lo's modified R/S	time	a test of short memory	low power against long memory
DFA	time	$\hat{H}$ (exponent α)	needs block sizes; affected by breaks
GPH	frequency	$\hat{d}$ with a standard error	large variance; choice of bandwidth m

- Other estimators: the local Whittle (Gaussian semiparametric) estimator Robinson (1995), wavelets, exact maximum likelihood for ARFIMA (Franke, Härdle & Hafner (2019), 14.5)
- Good practice: report several estimators, each with its uncertainty

Detrended Fluctuation Analysis (DFA)

- **DFA** Peng et al. (1994): designed for signals with slowly changing trends
  - 1. profile $Y_k = \sum_{j \le k}(x_j - \bar{x})$
  - 2. split Y into blocks of size n; in each block, fit a straight line by OLS (linear detrending)
  - 3. $F(n)$ = root mean square of the deviations from the lines, over all blocks
  - 4. $F(n) \propto n^{\alpha}$: the slope of $\ln F(n)$ on $\ln n$ is $\hat{\alpha}$
- Interpretation
  - for fGn, $\alpha = H$: $\alpha = 0.5$ white noise, $0.5 < \alpha < 1$ long memory; for fBm, $\alpha = H + 1$
  - the local line removes linear trends inside each block; $F(n)$ is unreliable for n below about 10 Kantelhardt et al. (2001)

R/S and DFA: S&P 500 since 2000



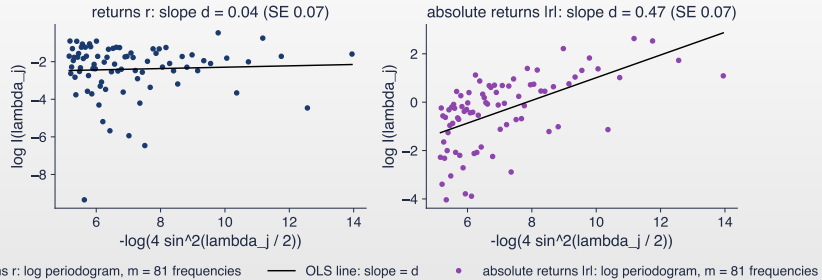
- 20 block sizes from 10 to 1678 days; dashed: mean curve of i.i.d. Normal series of the same length
- Returns: R/S slope 0.53 (i.i.d. series: 0.55), DFA 0.49 (i.i.d.: 0.50): no long memory
- Absolute returns: R/S 0.86, DFA 0.95: strong long memory in volatility (Section 7)

SFM_ch11_rs_estimators

The GPH Log-Periodogram Regression

- **Periodogram** at the Fourier frequencies $\lambda_j = 2\pi j/T$: $I(\lambda_j) = \frac{1}{2\pi T}\left|\sum_{t=1}^{T} x_t e^{-i\lambda_j t}\right|^2$
  - T: the number of observations; i: the imaginary unit; $|\cdot|^2$: the squared modulus
  - it estimates the spectral density, which behaves like $C\lambda^{-2d}$ near zero
- **GPH** Geweke & Porter-Hudak (1983); Franke, Härdle & Hafner (2019), 14.5.2:
  $\ln I(\lambda_j) = c - d\,\ln\left(4\sin^2(\lambda_j/2)\right) + e_j$, $j = 1, \ldots, m$
  - m: the number of low frequencies used; c: a constant; e_j: the error
  - $\hat{d}$ = OLS slope on $-\ln(4\sin^2(\lambda_j/2))$; $\hat{H} = \hat{d} + 0.5$
  - asymptotic standard error $\pi/\sqrt{24m}$; the bandwidth $m = \lfloor T^{0.5} \rfloor$
- Trade-off: a small m keeps only the long-run frequencies (small bias, large variance); a large m lets short-run dynamics in (bias)

GPH: S&P 500 Returns and Absolute Returns



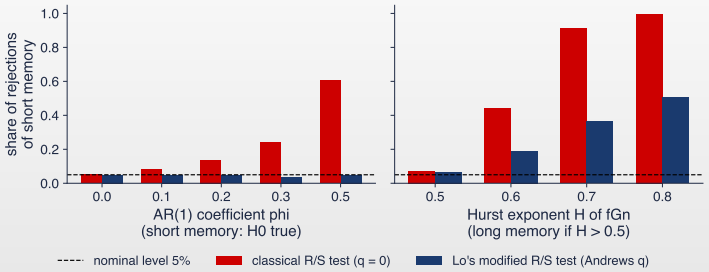
- $m = 81$ lowest frequencies; each point is one frequency
- Returns: $\hat{d} = 0.036$ (SE 0.071): not different from 0; absolute returns: $\hat{d} = 0.470$ (SE 0.071), $\hat{H} = 0.97$

SFM_ch11_rs_estimators

Lo's Modified R/S Test

- Problem: short memory (e.g. an AR(1)) also raises R/S, so classical R/S mistakes short memory for long memory
- Lo (1991): $V_N(q) = \dfrac{R_N}{\sqrt{N}\,\hat{\sigma}_N(q)}$, $\hat{\sigma}_N^{\mathbf{2}}(q) = \hat{\gamma}_{\mathbf{0}} + 2\sum_{j=\mathbf{1}}^{q}\left(1 - \frac{j}{q+\mathbf{1}}\right)\hat{\gamma}_j$
  - N: the number of returns; R_N: the range of the cumulative deviations over the whole sample
  - $\hat{\gamma}_j$: sample autocovariances; $\hat{\sigma}_N^{\mathbf{2}}(q)$: the Newey–West long-run variance Newey & West (1987); $q = 0$: classical R/S
  - $q = \lfloor (3N/2)^{\mathbf{1}/\mathbf{3}} \left(2\hat{\rho}_{\mathbf{1}}/(1-\hat{\rho}_{\mathbf{1}}^{\mathbf{2}})\right)^{\mathbf{2}/\mathbf{3}} \rfloor$ ($\hat{\rho}_{\mathbf{1}}$: first-order autocorrelation), data-dependent (Andrews' rule, used by Lo)
- $H_{\mathbf{0}}$: short memory; at 5%, reject if $V_N(q) \notin [0.809, 1.862]$
  - Lo's finding: no long memory in US stock index returns once short memory is allowed for

Classical against Modified R/S: Size and Power



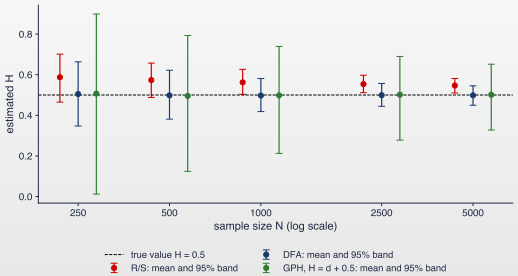
- 500 series of 1000 values. Left, AR(1) (short memory, H_0 true): with $\phi = 0.5$ the classical test rejects in 60% of cases, the modified test in 5%
- Right, fGn (long memory if $H > 0.5$): with $H = 0.7$, rejection rates 91% and 36%: the correction costs power Teverovsky, Taqqu & Willinger (1999)

SFM_ch11_monte_carlo

Small Samples and Monte Carlo Bands

- Asymptotic results are unreliable for R/S and DFA at the sample sizes of finance; use simulation Weron (2002)
  - 1. simulate many i.i.d. Normal series with the same length N as the data
  - 2. apply exactly the same estimator (same block sizes, same bandwidth)
  - 3. the 2.5% and 97.5% quantiles of $\hat{H}$ form the 95% band under $H = 0.5$
- Reject $H = 0.5$ only if $\hat{H}$ lies outside the band
- i.i.d. Normal is the simplest null; a fitted short-memory model (AR, GARCH) gives a stricter null (Section 6)

Monte Carlo: Bias and Spread of the Estimators



- 500 i.i.d. Normal series for each N; true $H = 0.5$
- R/S: mean 0.562 at $N = 1000$ (biased upwards); DFA: 0.497; GPH: 0.498, unbiased but wide
- Width of the 95% band at $N = 1000$: R/S 0.12, DFA 0.16, GPH 0.53; at $N = 250$: DFA 0.32

SFM_ch11_monte_carlo

S&P 500: All Estimators

	R/S	DFA	GPH $\hat{d}$ (SE)	Lo $V_N(q)$	q
returns r_t	0.53	0.49	0.04 (0.07)	1.56	7
absolute returns $\lvert r_t \rvert$	0.86	0.95	0.47 (0.07)	2.66	15
95% band, i.i.d.	[0.51, 0.57]	[0.46, 0.54]	[0.36, 0.65]	[0.809, 1.862]	

- $N = 6\,714$ daily returns, 2000 to 18 September 2026; GPH band given for $\hat{H} = \hat{d} + 0.5$
- Returns: every estimate inside its band, $V_N(q)$ inside $[0.809, 1.862]$
- Absolute returns: every estimate far above its band, Lo rejects short memory

Interpretation: S&P 500

- The direction of returns has no long memory since 2000
  - consistent with the weak form of the EMH (Chapter 7) and with Lo (1991)
  - R/S alone (0.53) would look like "$H > 0.5$"; its band starts at 0.51
- The size of returns has long memory: $\hat{d} = 0.47$ for $|r_t|$
  - the risk of a calm or turbulent period persists for months: the long-run version of volatility clustering
- Before concluding "long memory", rule out spurious long memory: next section

Structural Breaks and Spurious Long Memory

- A series with a few shifts in the mean has a sample ACF that decays slowly, like long memory
  - two regimes with means $\mu_1 \neq \mu_2$: every pair of observations from the same regime shares the deviation $\mu_i - \bar{x}$
- Landmark results
  - Diebold & Inoue (2001): rare regime switches produce estimates of d indistinguishable from true long memory
  - Granger & Hyung (2004): part of the long memory of S&P 500 absolute returns comes from occasional breaks
  - Lobato & Savin (1998): long memory in squared S&P 500 returns survives the checks; in returns, it is not found
- Tests of true against spurious long memory: Franke, Härdle & Hafner (2019), 14.4.3

The Nile Series: the Structural Break of 1898

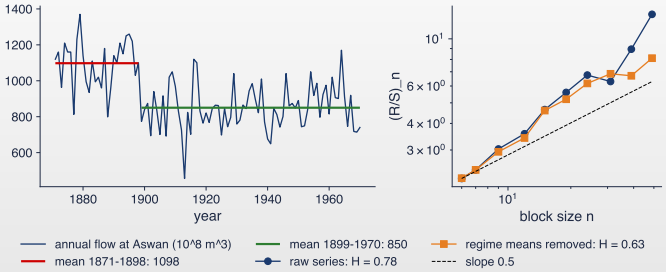
- Cobb (1978): the annual flow of the Nile at Aswan, 1871–1970, changes its mean after 1898
- Mean flow: 1098 before, 850 after (in 10^8 m^3), a drop of 23%
- The change coincides with the building of the Aswan Low Dam (1898–1902)
- A river record with a break is the same problem as a market with a change of regime



The Aswan Low Dam on the Nile

Photo: Karelj (2010); public domain; Wikimedia Commons

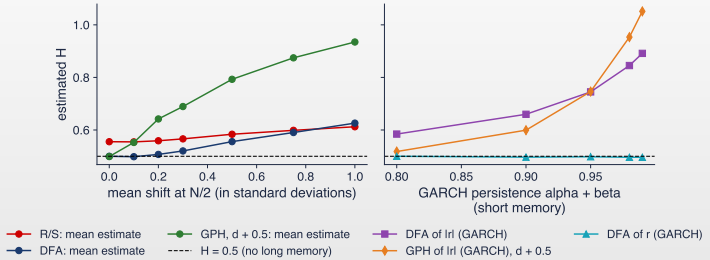
The Nile: R/S before and after Removing the Break



- Raw series: $\hat{H}_{R/S} = 0.78$, DFA 0.84; after removing the two regime means: 0.63 and 0.50
- With $N = 100$ the i.i.d. band for R/S is wide, $[0.45, 0.75]$, centred at 0.60: most of the apparent memory was the break

SFM_ch11_spurious

Simulated Spurious Long Memory



- 300 series of 2000 values. Left: i.i.d. $N(0,1)$ with a mean shift at the middle; a shift of 0.5 standard deviations gives GPH $\hat{H} = 0.79$, DFA 0.56
- Right: GARCH(1,1), a short-memory model (Chapter 9); for $|r_t|$, DFA gives 0.74 at $\alpha + \beta = 0.95$ and 0.89 at 0.99; the returns stay at 0.50

SFM_ch11_spurious

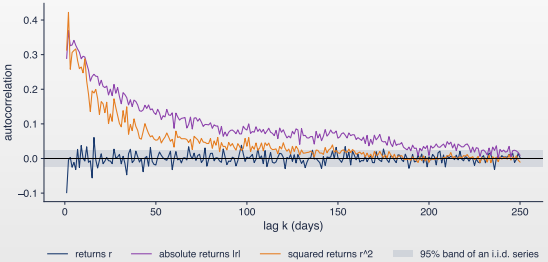
Volatility Clustering and Apparent Long Memory

- In GARCH(1,1), the ACF of r_t^2 decays like $(\alpha + \beta)^k$: short memory
  - with $\alpha + \beta$ close to 1, the decay is so slow that a sample of a few thousand days cannot separate it from k^{2d-1}
- Consequence: a high $\hat{H}$ for $|r_t|$ is compatible with a persistent GARCH; it is not by itself proof of fractional dynamics
- For returns, GARCH does not create memory: the estimates stay at 0.5 (right panel)
- Better null hypothesis: simulate the fitted GARCH and compare $\hat{H}$ of the data with the GARCH band

Checks against Spurious Long Memory

- **Shuffle test**: permute the series at random; the distribution stays, the order disappears; $\hat{H}$ should fall to 0.5
- **Subsamples**: estimate before and after known breaks (2008, 2020); true long memory appears in each subsample
- **Several estimators and bandwidths**: R/S, DFA, GPH with $m = T^{0.5}$ and $T^{0.6}$ should agree
- **Lo's modified R/S**: removes the effect of short memory
- **Model-based null**: Monte Carlo bands from a fitted AR or GARCH instead of i.i.d. noise

S&P 500: the ACF of r_t, $|r_t|$ and r_t^2



- $|r_t|$: $\rho(1) = 0.289$, $\rho(50) = 0.127$, $\rho(100) = 0.082$; 227 of the first 250 lags above the band 0.024
- r_t^2: $\rho(1) = 0.313$, $\rho(50) = 0.053$; r_t: $\rho(1) = -0.099$, then small ($|\rho(k)| \leq 0.061$), but 31 of the lags 2–250 lie outside the band (about 12 expected by chance; largest at lags 15, 16 and 34)
- Ding, Granger & Engle (1993): $|r_t|$ is more persistent than r_t^2, which is more affected by single extreme days

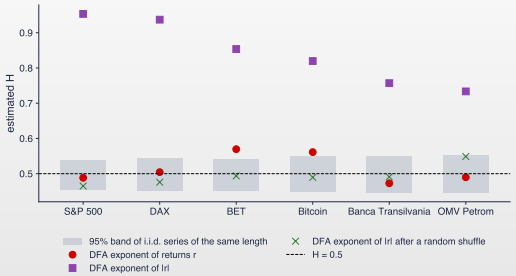
SFM_ch11_volatility_memory

Six Markets: Returns and Absolute Returns

	N	returns r_t			absolute returns $\lvert r_t \rvert$			DFA band
		R/S	DFA	GPH $\hat{d}$	DFA	GPH $\hat{d}$	Lo V	
S&P 500	6 714	0.53	0.49	0.04	0.95	0.47	2.66	[0.46, 0.54]
DAX	6 786	0.55	0.50	−0.06	0.94	0.48	3.43	[0.45, 0.54]
BET	6 670	0.58	0.57	0.17	0.85	0.47	4.47	[0.45, 0.54]
Bitcoin	4 384	0.59	0.56	0.05	0.82	0.33	3.94	[0.45, 0.55]
TLV	3 788	0.54	0.47	−0.06	0.76	0.25	2.47	[0.45, 0.55]
SNP	3 815	0.55	0.49	−0.02	0.73	0.21	2.39	[0.45, 0.55]

- Daily data to 18 September 2026: S&P 500, DAX, BET since 2000, Bitcoin since 2014, TLV and SNP since 2010; standard error of GPH about 0.07–0.08
- BET returns: DFA 0.57 above its band, GPH $\hat{d} = 0.17$, Lo $V = 1.96$ rejects short memory
- Absolute returns: $\hat{d}$ from 0.21 to 0.48, Lo rejects for all six

Six Markets: DFA Exponents and the Shuffle Test



- Shaded: 95% band of i.i.d. series of the same length; circles: returns; squares: $|r_t|$; crosses: $|r_t|$ after a random shuffle
- $|r_t|$: DFA from 0.73 to 0.95; after shuffling: from 0.46 to 0.55, inside the bands: the memory lies in the order of the days

SFM_ch11_volatility_memory

Interpretation: Predictability of Returns and Volatility

- **Returns**: no long memory in the S&P 500, DAX, TLV and SNP; Bitcoin at the edge of its band
  - BET: persistence, as in the variance-ratio tests of Chapter 7; partly the thin trading of the early years (Section 8)
- **Volatility**: long memory everywhere, strongest in the large equity indices
  - the risk regime of today is informative for months: useful for VaR (value at risk) and option pricing
- Weak-form efficiency (unpredictable direction) and predictable risk coexist: the FMH adds a reason, investors with different horizons react to volatility on different time scales

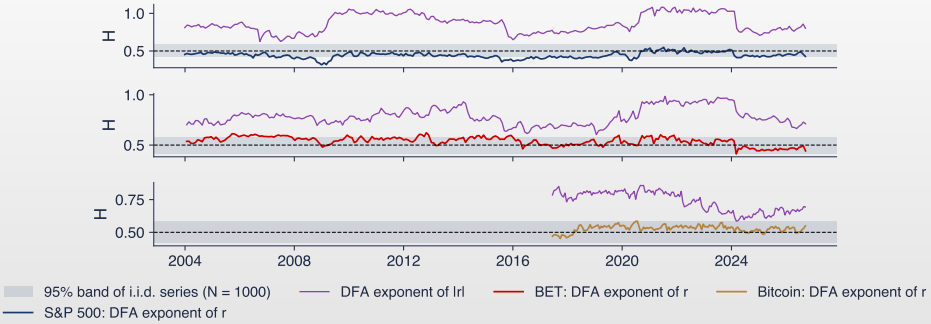
Models for Long Memory in Volatility

- **FIGARCH** (fractionally integrated GARCH) Baillie, Bollerslev & Mikkelsen (1996); Franke, Härdle & Hafner (2019), 14.6.2: $(1-L)^d$ applied to ε_t^2
  - shocks to volatility decay hyperbolically; $d = 0$ gives GARCH, $d = 1$ gives IGARCH (integrated GARCH, Chapter 9)
- **Heterogeneous horizons**, the FMH idea in volatility models
  - Müller et al. (1997): market components with different time resolutions (intraday traders to long-term investors)
  - Andersen & Bollerslev (1997): many volatility components with different persistence add up to long memory
  - Corsi (2009): HAR (heterogeneous autoregressive) model, tomorrow's volatility on the daily, weekly and monthly volatility: three horizons, approximate long memory

Time Variation of Market Efficiency

- The AMH and the FMH both predict that market memory changes over time
  - Cajueiro & Tabak (2004): rolling Hurst exponents of emerging stock markets, to test whether they become more efficient
  - Bariviera (2017): rolling DFA on Bitcoin returns and volatility
- Design used here
  - windows of 1000 observations (about four years of trading days), moved by 21 observations (about one month)
  - DFA on r_t and on $|r_t|$; each estimate dated at the last day of its window, so it uses only past data
  - 95% Monte Carlo band for $N = 1000$: $[0.42, 0.58]$

Rolling DFA Exponents: S&P 500, BET, Bitcoin



- Top to bottom: S&P 500, BET, Bitcoin; coloured line: r_t; purple line: $|r_t|$; shaded: i.i.d. band
- $|r_t|$ always far above the band; its exponent rises in and after crises (2008–2009, 2020)

SFM_ch11_rolling_hurst

Interpretation: Rolling Exponents

- S&P 500: 273 windows; $\hat{H}$ of returns between 0.32 and 0.55; 24% of windows below the band, none above
  - the minimum, 31 December 2008, follows the 2008 crash: daily reversals (anti-persistence) in a panic
- BET: 20% of windows above the band; mean $\hat{H}$ 0.56 in the first third of the windows, 0.52 in the last third
  - the maximum, 31 October 2012; a market that becomes more efficient as liquidity grows
- Bitcoin: windows from 13 June 2017; $\hat{H}$ between 0.45 and 0.59, 1% of windows above the band

Case Study: Efficiency of the Bitcoin Market

- Urquhart (2016): several tests, among them R/S, on Bitcoin 2010–2016
  - inefficient over the full sample, moving towards efficiency in the later years
- Nadarajah & Chu (2017): the same data after a power transformation of the returns: consistent with weak-form efficiency
- Bariviera (2017): rolling DFA: persistent returns until 2014, close to $H = 0.5$ afterwards; long memory in volatility throughout
- Our windows start in 2017 and agree: returns near the band, $|r_t|$ far above it



A Bitcoin cash machine in Prague (2016)

Photo: Perituss (2016); CC0; Wikimedia Commons

Limitations of Rolling Estimates

- **Multiple testing**: 273 windows at 5% each give about 14 windows outside the band by chance
  - look for long runs outside the band, not for isolated windows
- **Overlap**: consecutive windows share 979 of their 1000 observations; the path is smooth by construction
- **Window length**: shorter windows react faster but have wider bands (DFA, $N = 250$: width 0.32)
- **Volatility regimes**: a crisis inside a window shifts the estimate for four years
- **Dating**: date each estimate at the end of its window, never at the middle (no look-ahead)

Scaling Risk with h^H

- If returns are self-similar: h-day VaR 1% $= -z_{0.01}\,\sigma\,h^H$, with $z_{0.01} = -2.326$ (Normal case, mean 0)
  - $z_{0.01}$: the 1% quantile of $N(0,1)$; σ: the daily volatility; h^H replaces the factor $\sqrt{h}$
  - daily volatility $\sigma = 1.2\%$: one-day VaR 1% $= 2.79\%$
- Ten days
  - $H = 0.45$: factor 2.82, VaR 7.87%; $H = 0.5$: 3.16, 8.83%; $H = 0.55$: 3.55, 9.91%
  - $H = 0.55$ instead of 0.5: 12% more risk at 10 days, 32% more at 250 days
- Data: the 10-day ratio is 2.80 for the S&P 500 and 3.50 for the BET, against $\sqrt{10} = 3.16$: the $\sqrt{h}$ rule overstates S&P 500 risk and understates BET risk (Chapter 10)
- Pricing: with fBm prices and $H \neq 1/2$, a frictionless market admits arbitrage Rogers (1997); long memory enters option models through volatility instead

An Open Research Question

- **Does a change in the Hurst exponent come before a crisis, or only after it?**
  - the FMH predicts a collapse of horizons; in our data, $\hat{H}$ of $|r_t|$ rises in and after 2008 and 2020
  - a rolling window contains the crisis only after it starts: is there any signal before?
- Earlier evidence: Kristoufek (2013) (horizons in 2008); Cajueiro & Tabak (2004) (emerging markets)
- Why it is open: few crises, overlapping windows, volatility regimes, many possible estimators
- AI tools can speed up such a study, but its results must still be checked Wang et al. (2023)

Possible Uses of AI Tools

- **Literature**: a list of studies on Hurst exponents as early-warning indicators, with their data and estimators
- **Code**: a first draft of rolling R/S, DFA and GPH with Monte Carlo bands
- **Robustness**: other window lengths, wavelet estimators, a GARCH-based null, other crises (2011, 2022)
- Example prompt
  - `Write Python code that computes the DFA exponent of daily S&P 500 absolute log returns on rolling windows of 500 and 1000 days, dated at the window end, adds a 95% band from 500 simulated GARCH(1,1) paths, and marks the start of the 2008 and 2020 crises.`

What to Check

- Dating: an estimate dated at the middle of its window uses future data
- Conversions: $d = H - 0.5$, not $H + 0.5$; DFA on prices gives $H + 1$, not H
- Bands: every $\hat{H}$ needs a band for its own N; R/S is biased upwards
- Returns or volatility: memory in $|r_t|$ does not make the direction of returns predictable
- GPH bandwidth: $m = T/2$ (all frequencies) is not "more accurate"; report $T^{0.5}$ and $T^{0.6}$
- References: every cited paper must exist; check the DOI

Project Idea

- **Question**: does the rolling Hurst exponent of volatility rise before crises on the BVB and in Bitcoin?
  - data: BET, BET-TR, TLV, SNP and Bitcoin (EODHD), daily
- Steps
  - rolling DFA and GPH of r_t and $|r_t|$, windows of 500 and 1000 days, dated at the end
  - bands from i.i.d. and from fitted GARCH(1,1) simulations
  - event windows around 2008, 2020 and 2022: the exponent in the 60 days before and after
- Deliverable: one table, one chart, and a paragraph on what the data can and cannot show
- Declare any AI use, and list the errors of the AI that you corrected

Key Takeaways

- FMH: stability needs investors with many horizons; crises are collapses of horizons
- H measures memory and scaling: $H = 0.5$ none, $H > 0.5$ persistence, $H < 0.5$ anti-persistence; $d = H - 0.5$
- R/S, DFA, GPH and Lo's test, always with a Monte Carlo band for the same N
- Breaks and persistent GARCH volatility create spurious long memory: shuffle, split, compare
- Returns: no long memory in large markets (BET: persistence); volatility: long memory everywhere
- Multi-day risk scales like h^H

Key Formulas

Quantity	Formula
Self-similarity	$X_{ct} \overset{d}{=} c^H X_t; \quad \mathrm{sd}(r^{(h)}) = h^H \mathrm{sd}(r)$
R/S	$(R/S)_n = [\max_k Y_k - \min_k Y_k]/S_n, \quad Y_k = \sum_{j\le k}(x_j - \bar{x}); \quad (R/S)_n \propto n^H$
fGn	$\rho(k) = \frac{1}{2}(\lvert k+1\rvert^{2H} - 2\lvert k\rvert^{2H} + \lvert k-1\rvert^{2H}) \approx H(2H-1)k^{2H-2}$
ARFIMA(0,d,0)	$(1-L)^d X_t = \varepsilon_t, \quad \rho(k) = \rho(k-1)\frac{k-1+d}{k-d}, \quad d = H - 0.5$
DFA	$F(n) \propto n^{\alpha}, \quad \alpha = H$ for fGn
GPH	$\ln I(\lambda_j) = c - d\ln(4\sin^2(\lambda_j/2)) + e_j, \quad j \le m = T^{0.5}, \quad \mathrm{SE} = \pi/\sqrt{24m}$
Lo	$V_N(q) = R_N/(\sqrt{N}\hat{\sigma}_N(q))$; reject at 5% outside $[0.809, 1.862]$
VaR 1%	$-z_{0.01}\,\sigma\,h^H$

Self-Assessment

- **Question**: an R/S estimate on 1000 i.i.d. returns is $\hat{H} = 0.56$. Is this evidence of long memory?
  - **Answer**: no: the i.i.d. band for $N = 1000$ is $[0.50, 0.63]$, centred near 0.562
- **Question**: GPH gives $\hat{d} = 0.35$ for $|r_t|$. What is $\hat{H}$?
  - **Answer**: $\hat{H} = \hat{d} + 0.5 = 0.85$: long memory in volatility, not in returns
- **Question**: after a random shuffle of $|r_t|$, DFA falls from 0.9 to 0.5. What does this show?
  - **Answer**: the memory comes from the order of the days, not from the distribution (heavy tails)
- Next: Chapter 12, scoring models

References (1/3)

- Andersen, T.G., Bollerslev, T. (1997). Heterogeneous information arrivals and return volatility dynamics: uncovering the long-run in high frequency returns. *Journal of Finance*, 52(3), 975–1005.
- Anis, A.A., Lloyd, E.H. (1976). The expected value of the adjusted rescaled Hurst range of independent normal summands. *Biometrika*, 63(1), 111–116.
- Baillie, R.T. (1996). Long memory processes and fractional integration in econometrics. *Journal of Econometrics*, 73(1), 5–59.
- Baillie, R.T., Bollerslev, T., Mikkelsen, H.O. (1996). Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 74(1), 3–30.
- Bariviera, A.F. (2017). The inefficiency of Bitcoin revisited: a dynamic approach. *Economics Letters*, 161, 1–4.
- Cajueiro, D.O., Tabak, B.M. (2004). The Hurst exponent over time: testing the assertion that emerging markets are becoming more efficient. *Physica A*, 336(3–4), 521–537.
- Carlson, M. (2007). A brief history of the 1987 stock market crash with a discussion of the Federal Reserve response. Finance and Economics Discussion Series 2007-13, Federal Reserve Board.
- Cobb, G.W. (1978). The problem of the Nile: conditional solution to a changepoint problem. *Biometrika*, 65(2), 243–251.
- Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2), 174–196.
- Diebold, F.X., Inoue, A. (2001). Long memory and regime switching. *Journal of Econometrics*, 105(1), 131–159.
- Ding, Z., Granger, C.W.J., Engle, R.F. (1993). A long memory property of stock market returns and a new model. *Journal of Empirical Finance*, 1(1), 83–106.
- Fama, E.F. (1970). Efficient capital markets: a review of theory and empirical work. *Journal of Finance*, 25(2), 383–417.
- Franke, J., Härdle, W.K., Hafner, C.M. (2019). *Statistics of Financial Markets: An Introduction*, 5th ed. Springer.
- Geweke, J., Porter-Hudak, S. (1983). The estimation and application of long memory time series models. *Journal of Time Series Analysis*, 4(4), 221–238.

References (2/3)

- Granger, C.W.J., Hyung, N. (2004). Occasional structural breaks and long memory with an application to the S&P 500 absolute stock returns. *Journal of Empirical Finance*, 11(3), 399–421.
- Granger, C.W.J., Joyeux, R. (1980). An introduction to long-memory time series models and fractional differencing. *Journal of Time Series Analysis*, 1(1), 15–29.
- Hosking, J.R.M. (1981). Fractional differencing. *Biometrika*, 68(1), 165–176.
- Hurst, H.E. (1951). Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers*, 116(1), 770–799.
- Kantelhardt, J.W., Koscielny-Bunde, E., Rego, H.H.A., Havlin, S., Bunde, A. (2001). Detecting long-range correlations with detrended fluctuation analysis. *Physica A*, 295(3–4), 441–454.
- Kristoufek, L. (2012). Fractal markets hypothesis and the global financial crisis: scaling, investment horizons and liquidity. *Advances in Complex Systems*, 15(6), 1250065.
- Kristoufek, L. (2013). Fractal markets hypothesis and the global financial crisis: wavelet power evidence. *Scientific Reports*, 3, 2857.
- Lo, A.W. (1991). Long-term memory in stock market prices. *Econometrica*, 59(5), 1279–1313.
- Lo, A.W. (2004). The adaptive markets hypothesis. *Journal of Portfolio Management*, 30(5), 15–29.
- Lobato, I.N., Savin, N.E. (1998). Real and spurious long-memory properties of stock-market data. *Journal of Business & Economic Statistics*, 16(3), 261–268.
- Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business*, 36(4), 394–419.
- Mandelbrot, B. (1967). How long is the coast of Britain? Statistical self-similarity and fractional dimension. *Science*, 156(3775), 636–638.
- Mandelbrot, B.B., Van Ness, J.W. (1968). Fractional Brownian motions, fractional noises and applications. *SIAM Review*, 10(4), 422–437.
- Mandelbrot, B.B., Wallis, J.R. (1968). Noah, Joseph, and operational hydrology. *Water Resources Research*, 4(5), 909–918.

References (3/3)

- Mandelbrot, B.B., Wallis, J.R. (1969). Robustness of the rescaled range R/S in the measurement of noncyclic long run statistical dependence. *Water Resources Research*, 5(5), 967–988.
- Müller, U.A., Dacorogna, M.M., Davé, R.D., Olsen, R.B., Pictet, O.V., von Weizsäcker, J.E. (1997). Volatilities of different time resolutions: analyzing the dynamics of market components. *Journal of Empirical Finance*, 4(2–3), 213–239.
- Nadarajah, S., Chu, J. (2017). On the inefficiency of Bitcoin. *Economics Letters*, 150, 6–9.
- Newey, W.K., West, K.D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708.
- Pele, D.T., Mazurencu-Marinescu, M. (2012). Modelling stock market crashes: the case of Bucharest Stock Exchange. *Procedia – Social and Behavioral Sciences*, 58, 533–542.
- Pele, D.T., Mazurencu-Marinescu-Pele, M. (2019). Metcalfe's law and log-period power laws in the cryptocurrencies market. *Economics*, 13(1), 20190029.
- Peng, C.-K., Buldyrev, S.V., Havlin, S., Simons, M., Stanley, H.E., Goldberger, A.L. (1994). Mosaic organization of DNA nucleotides. *Physical Review E*, 49(2), 1685–1689.
- Peters, E.E. (1994). *Fractal Market Analysis: Applying Chaos Theory to Investment and Economics*. Wiley.
- Robinson, P.M. (1995). Gaussian semiparametric estimation of long range dependence. *Annals of Statistics*, 23(5), 1630–1661.
- Rogers, L.C.G. (1997). Arbitrage with fractional Brownian motion. *Mathematical Finance*, 7(1), 95–105.
- Teverovsky, V., Taqqu, M.S., Willinger, W. (1999). A critical look at Lo's modified R/S statistic. *Journal of Statistical Planning and Inference*, 80(1–2), 211–227.
- Urquhart, A. (2016). The inefficiency of Bitcoin. *Economics Letters*, 148, 80–82.
- Wang, H., Fu, T., Du, Y., Gao, W., et al. (2023). Scientific discovery in the age of artificial intelligence. *Nature*, 620, 47–60.
- Weron, R. (2002). Estimating long-range dependence: finite sample properties and confidence intervals. *Physica A*, 312(1–2), 285–299.