

Multiscaling in the Rough Bergomi Model: A Tale of Tails

Giuseppe Brandi¹ Tiziana Di Matteo^{2,3,4}

¹CoMENS, Northeastern University London

²King's College London

³Complexity Science Hub Vienna

⁴Centro Ricerche Enrico Fermi



Financial Engineering Workshop
Bayes Business School, London ■ 14 October 2026



1 Scaling and multiscaling in finance

what it is, how we measure it, and why its *source* is still debated

2 Rough volatility meets multiscaling

what we found for the rough Bergomi model, and the question it left open

3 A tale of tails

a two-stage surrogate test that separates distributional from temporal sources; validation on MRW and FLSM; results for rBergomi

4 What it means in practice

for users of rough volatility models, and as a diagnostic for any model or market

Financial markets are open systems in which many subunits interact nonlinearly in the presence of feedback. The governing rules are rather stable and the time evolution of the system is continuously monitored.

Two stylised facts are central to this talk:

- **heavy, power-law tails** of returns;
- **volatility clustering**: slowly decaying autocorrelation of $|r_t|$.

Scaling analysis asks how the statistical properties of price changes depend on the time horizon τ . Both facts leave a mark on that dependence, and the question of this talk is *which one* produces multiscaling.

For a (log-)price process $X(t)$ with stationary increments, the *generalised Hurst exponent* (GHE) approach studies

$$\Xi(\tau, q) \equiv \mathbb{E}[|X(t + \tau) - X(t)|^q] \sim K(q) \tau^{\zeta(q)}, \quad \zeta(q) = q H(q).$$

Uniscaling

- $H(q) \equiv H$: $\zeta(q)$ is linear in q
- one exponent describes fluctuations of every size
- e.g. Brownian motion ($H = \frac{1}{2}$), fBm

Multiscaling

- $H(q)$ varies with q : $\zeta(q)$ is concave
- small and large fluctuations scale differently (intermittency)
- e.g. multifractal cascades, MRW

$H(1)$ describes the scaling of absolute deviations, $H(2)$ that of the variance (autocorrelation).

T. Di Matteo, Quant. Finance 7 (2007) 21; B. Mandelbrot, A. Fisher & L. Calvet, Cowles Foundation DP 1164 (1997); L. Calvet & A. Fisher, Rev. Econ. Stat. 84 (2002) 381.

A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design


Method

Validation

Results


In practice

Financial data scale... but they do not uniscale



Moments of Nikkei 225 price changes against τ (log-log): clean power laws for every $q \Rightarrow$ **scaling**.

T. Di Matteo, Quant. Finance 7 (2007) 21.



$qH(q)$ vs q for equity indices, FX, Treasury and Eurodollar rates: not linear, $H(1) \neq H(2) \neq 0.5 \Rightarrow$ **multiscaling**.

A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice

Measuring multiscaling: the GHE regression



A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice

$$\ln \Xi(\tau, q) = \zeta(q) \ln \tau + \ln K(q), \quad \tau \in [\tau_{\min}, \tau_{\max}]$$

Multiscaling proxy, fitted across moment orders:

$$\zeta(q) \simeq Aq + Bq^2 \iff H(q) \simeq A + Bq, \quad \begin{cases} \hat{B} = 0 & \text{uniscaling} \\ \hat{B} < 0 & \text{multiscaling} \end{cases}$$

the more negative $\hat{B}$, the stronger the multiscaling

Two pitfalls:

- **Range of q .** Moments exist only for $q < \alpha$ (tail index), and their sample estimators need $q < \alpha/2$ for finite variance. Heavy tails bias $\hat{H}(q)$ at large q and can *mimic* multiscaling.
- **Range of τ .** Different horizons can give different answers.

Tails or memory? The debate on the source

Kantelhardt et al. (2002) pointed to two possible sources; the literature has not agreed since.

Distributional source

Multiscaling comes from the fat-tailed marginal of returns; temporal structure plays no essential role beyond producing the tails.

Zhou (2009)

Temporal source

Multiscaling comes from memory, i.e. temporal dependence, and appears even in long-range correlated Gaussian processes.

Green, Hanan & Heffernan (2014)

... or both, acting jointly (Stanley 2003; see also Barunik, Aste, Di Matteo & Liu 2012)

- **How to tell them apart:** shuffling the returns keeps the marginal and destroys memory. Multiscaling that survives shuffling is distributional.
- **Why rough volatility:** rough variance generates both heavy tails and volatility clustering, so it is a natural test case.

J. W. Kantelhardt et al., Physica A 316 (2002) 87; W.-X. Zhou, EPL 88 (2009) 28004; E. Green, W. Hanan & D. Heffernan, Eur. Phys. J. B 87 (2014) 129; H. E. Stanley, Physica A 318 (2003) 279; J. Barunik et al., Physica A 391 (2012) 4234.

A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice

- Increments of log-realised volatility scale with $H \approx 0.1$ across assets and periods:
volatility is rough (Gatheral, Jaisson & Rosenbaum 2018)
- Rough models fit the implied-volatility surface, notably the steep short-maturity ATM skew, with few parameters (Bayer, Friz & Gatheral 2016)
- Microstructural foundations from nearly unstable Hawkes order flow and no-arbitrage (El Euch, Fukasawa & Rosenbaum 2018; Jusselin & Rosenbaum 2020)
- Model-free support: LSTM forecasters recover the same scaling (Rosenbaum & Zhang 2024)
- **Nuance**: the ATM skew is not a single power law across maturities, and non-rough two-factor models also fit SPX well (Rømer 2022; Guyon & El Amrani 2023; Delemotte, De Marco & Segonne 2023)

The rough Bergomi model



A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice

$$\frac{dS(t)}{S(t)} = \sqrt{v(t)} \left(\rho dW(t) + \sqrt{1 - \rho^2} dW^\perp(t) \right),$$

$$v(t) = \xi_0(t) \exp\left(\eta \widehat{W}^H(t) - \frac{1}{2}\eta^2 t^{2H}\right), \quad t \in [0, T],$$

$$\widehat{W}^H(t) = \sqrt{2H} \int_0^t (t-s)^{H-1/2} dW(s), \quad H \in (0, \tfrac{1}{2}).$$

- ρ : leverage correlation; η : vol-of-vol; $\xi_0(t)$: forward variance curve; H : roughness, paths become rougher as $H \rightarrow 0$.
- The canonical rough volatility model, used for SPX and VIX derivatives; calibrations to SPX options give $H \approx 0.05$ – 0.15 with $\rho \approx -0.9$.

Our earlier work: multiscaling and rough volatility



A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice

Motivation

- Volatility behaves like a fractional process with very small H .
- Prices are multiscaling, with different degrees across assets.

Questions

- Can rough volatility models generate multiscaling prices?
- How are volatility roughness $\hat{H}^{(v)}$ and price multiscaling $\hat{B}^{(P)}$ related?
- Are the models in line with the empirical evidence?

Approach: simulate the rBergomi model and study the interplay between volatility roughness and price multiscaling; then compare with stock indices.



G. Brandi & T. Di Matteo, Int. Rev. Financial

Analysis 84 (2022) 102324

rBergomi simulations: roughness generates multiscaling



A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea


Design

Method


Validation

Results

In practice




$\hat{B}^{(P)}$ as a function of H and the price-vol correlation λ



Pearson correlation between $\hat{B}^{(P)}$ and $\hat{H}^{(v)}$

- Price multiscaling appears only for small H , i.e. strong volatility roughness.
- Positive and significant correlation between $\hat{H}^{(v)}$ and $\hat{B}^{(P)}$: rougher volatility, stronger price multiscaling.

Averages over 100 datasets, surfaces smoothed by interpolation. λ in these plots is the correlation denoted ρ on the previous slide.



- Pearson $\rho = -0.43$, Spearman $\rho_S = -0.51$
- outlier-robust: $\tilde{\rho} = -0.61$, $\tilde{\rho}_S = -0.65$
- significant at 5%

The sign is opposite to the simulations.

The open question



A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice

- We showed numerically that rBergomi prices become multiscaling as $H \rightarrow 0$; Forde, Fukasawa, Gerhold & Smith (2022) characterised the small- H limit theoretically, where the integrated variance becomes extremely heavy-tailed.
- But *why* is there multiscaling? Is it
 - a genuine **temporal** property of the price dynamics, or
 - inherited from the **fat-tailed marginal** of returns that rough variance produces?
- The answer matters for how we read the “complexity” of rough models, for scaling-based risk extrapolations, and for the mismatch with data.

This talk

A statistical procedure that answers this question, and what it says about rBergomi.

Multiscaling in the rough Bergomi model: a tale of tails

G. Brandi & T. Di Matteo



A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation


Results

In practice

Contributions

- 1 A two-stage surrogate-based test:** (1) is multiscaling present, against a matched uniscaling fBm? (2) can the return distribution alone reproduce it (shuffled surrogates)?
- 2 Sharper estimation:** weighted least squares for B ; data-driven ranges for q and τ .
- 3 Validation** on processes whose source is known: MRW (temporal) and FLSM (distributional).
- 4 Finding:** two regimes in rBergomi. Tail-driven multiscaling as $H \rightarrow 0$, clustering-driven multiscaling at moderate H , with a crossover around $H \approx 0.05\text{--}0.10$.


Key idea: surrogates that switch off one ingredient



Shuffling preserves the exact marginal (all moments, tails) and destroys temporal dependence at every lag. If $\hat{B}$ survives shuffling, the multiscaling is *distributionally attributable*.

- **Parameters** in line with SPX calibrations: $\rho = -0.9$, $\eta = 1.9$, flat forward variance $\xi_0 = 0.1$.
- **Roughness grid** $H \in \{0.001, 0.005, 0.01, 0.05, 0.10, 0.20\}$: from the rough limit to beyond the calibrated range $H \approx 0.05$ – 0.15 .
- **Size**: $N = 10,000$ steps per path and $R = 1000$ Monte Carlo paths for each H ; hybrid scheme of Bennedsen, Lunde & Pakkanen (2017).
- Analysis on the log-price $X(t) = \log S(t)$.

Rougher volatility, wilder prices



Price (top), log-returns (middle) and spot volatility (bottom) for $H = 0.3, 0.1, 0.05$, same random seed. As H decreases: larger and more clustered price excursions (paths are continuous: these are not jumps), heavier-tailed returns, rougher and more clustered volatility.

A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design


Method

Validation

Results

In practice

A tale of tails: the return distribution



H	$\hat{\alpha}$ (Hill)	Exc. kurt.
0.30	2.822 ± 0.089	4.08
0.10	2.422 ± 0.077	25.42
0.05	2.050 ± 0.065	35.30

Top: QQ-plots against the Normal.
Bottom: CCDF of $|r_t|$ with power-law fit (Hill estimator).

As $H \rightarrow 0$ the tail index drifts towards 2 and kurtosis explodes.

A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice

Estimating the multiscaling proxy

Moments. From non-overlapping increments, normalise by $\tau = 1$ and take the q -th root, so all orders share one scale:

$$\Xi(\tau, q) = \left(\frac{\hat{\Xi}(\tau, q)}{\hat{\Xi}(1, q)} \right)^{1/q} \sim \tau^{H(q)}, \quad \log \Xi(\tau, q) = H(q) \log \tau + \varepsilon_{\tau, q} \Rightarrow \hat{H}_n(q), \hat{\sigma}_{H(q)}$$

Proxy. Fit $H(q) = A + Bq$ across moments by weighted least squares:

$$(\hat{A}, \hat{B}) = \arg \min_{A, B} \sum_k w_k [\hat{H}_n(q_k) - (A + Bq_k)]^2, \quad w_k = 1/\hat{\sigma}_{H(q_k)}^2$$

- Heavy tails make the high- q estimates the noisiest; the weights stop them from driving $\hat{B}$.
- This matters even more on market data, where the noise differs more across moments than in simulations.

A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice

Data-driven ranges for q and τ

A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice

Moment orders q

- Moments exist only for $q < \alpha$; beyond that, heavy tails create spurious curvature in $\hat{H}(q)$.
- Estimate the tail index from the series itself, via an α -stable fit, and set

$$q_{\max} = 0.8 \hat{\alpha}_{\text{stable}}.$$

- The range adapts to each series: narrower for heavy tails, wider for light ones. In rBergomi, $q_{\max} \in [1, 2)$.

Time scales τ

- For each candidate $\tau_{\max}$, check how well every moment scales (R^2 of the GHE regression).
- Keep the largest range on which *all* moments scale cleanly:

$$\tau_{\max}^* = \max \left\{ \tau_{\max} : \min_q R^2(q, \tau_{\max}) \geq 0.95 \right\}$$

- No hand-picked horizons: the same rule applies to any model or market.

Stage 1: is multiscaling present?



A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice

$\mathcal{H}_0^{(1)}$: the process is uniscaling fBm with $H(q) \equiv H(1)$

$\mathcal{H}_A^{(1)}$: its scaling is inconsistent with uniscaling fBm

- A t -test of $B = 0$ ignores that uniscaling processes give $\hat{B} \neq 0$ in finite samples $\Rightarrow$ over-rejection.
- Simulate $M = 1000$ fBm paths with $H = \hat{H}_n(1)$ and run the *identical* pipeline on them.
- One-sided empirical p -value, reject at 5%:

$$p_{\text{presence}} = \frac{1}{M} \sum_{m=1}^M \mathbb{I}(\hat{B}_{\text{fBm},m} \leq \hat{B}_{\text{rBergomi}})$$

Intuition

The benchmark behaves like the data at $q = 1$ and is uniscaling by construction, so it shows how negative $\hat{B}$ can get by chance alone.

Stage 2: where does it come from?



A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice

$\mathcal{H}_0^{(2)}$: $\hat{B}_{\text{rBergomi}}$ is consistent with what shuffled surrogates produce

- $J = 1000$ random permutations of the returns: same marginal, no memory.
- Distance from the median of the shuffled $\hat{B}$, robust to skewed null distributions:

$$\tilde{B} = \text{median}_j \hat{B}_{\text{shuf},j}, \quad d_{\text{orig}} = |\hat{B}_{\text{rBergomi}} - \tilde{B}|, \quad d_j = |\hat{B}_{\text{shuf},j} - \tilde{B}|,$$

$$p_{\text{source}} = \frac{1}{J} \sum_{j=1}^J \mathbb{I}(d_j \geq d_{\text{orig}}).$$

Not rejected

No detectable temporal contribution:
distributionally attributable.

Rejected

Temporal dependence contributes: it **enhances** ($\hat{B} < \tilde{B}$) or **attenuates** ($\hat{B} > \tilde{B}$) multiscaling.

The procedure at a glance



A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea


Design

Method

Validation

Results

In practice



Repeated for $R = 1000$ replicates per parameter value.

Sig (%): Stage 1 rejection rate.

Distributional / Temporal (%): Stage 2 outcomes among detected cases.

Validation I: Multifractal Random Walk (positive control)

Multiscaling is genuine and *temporal* by construction: a log-volatility cascade with intermittency λ^2 (fat tails also arise, but the cascade dominates for larger λ).

λ	Multiscaling	Source of multiscaling		Statistics	
	Sig (%)	Distributional (%)	Temporal (%)	Mean $\hat{B}$	SD($\hat{B}$)
0.05	18.8	95.2	4.8	-0.0018	0.0039
0.15	78.5	75.3	24.7	-0.0094	0.0053
0.25	97.7	21.5	78.5	-0.0248	0.0081

- Weak cascade ($\lambda = 0.05$): rarely detected, near-Gaussian. These few detections are the right tail of the noise, not a distributional source.
- Strong cascade ($\lambda = 0.25$): detected in 97.7% of replicates, and 78.5% of them are correctly attributed to **temporal dependence**.

A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice

Validation II: Fractional Lévy Stable Motion (specificity check)

Fractional integration of α -stable noise ($\alpha = 1.90$): heavy tails and long-range dependence, but *not* multiscaling. At $H = 0.5$ it is a pure α -stable Lévy motion, so any detection must be distributional.

H	Multiscaling	Source of multiscaling		Statistics	
	Sig (%)	Distributional (%)	Temporal (%)	Mean $\hat{B}$	SD($\hat{B}$)
0.10	23.0	93.9	6.1	-0.0003	0.0058
0.50	23.3	94.0	6.0	-0.0020	0.0078
0.90	27.3	65.2	34.8	-0.0107	0.0103

- Modest detection rates (23–27%): finite-sample artefacts of heavy tails.
- For $H \leq 0.5$ about 94% are attributed to the **marginal**: the test does not manufacture temporal structure out of tails.
- $H = 0.9$: the temporal share rises to 34.8% (fractional integration), the most ambiguous case.

What the two benchmarks establish



A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice


Benchmark	True by construction	What the procedure says
MRW, $\lambda = 0.25$	genuine multiscaling, temporal (cascade)	detected in 97.7%; 78.5% temporal ✓
FLSM, $\alpha = 1.9$, $H = 0.5$	heavy tails, <i>no</i> multiscaling	23.3% spurious detections; 94% distributional ✓
missing case	genuine multiscaling, distributional	future work: Dai Pra & Pigato (2015)

- **Sensitivity**: the procedure finds temporal multiscaling when it is there.
- **Specificity**: it does not manufacture temporal structure out of heavy tails.

The two benchmarks bracket the two ways the attribution could fail, so the rBergomi results can be read in both directions.

P. Dai Pra & P. Pigato, Stoch. Proc. Appl. 125 (2015) 3725.

rBergomi: multiscaling is present and grows as $H \rightarrow 0$



$\hat{B}$ over 1000 Monte Carlo replicates for each H . Stage 1 detection: **100%** for $H \leq 0.01$, 99.9% at $H = 0.05$, 98.6% at $H = 0.10$, 71.5% at $H = 0.20$.

- In the calibrated range ($H \approx 0.05$ – 0.10) multiscaling is milder, mean $\hat{B}$ from -0.048 to -0.029 , but still detected in almost every path.

A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation


Results

In practice

	H	Multiscaling	Source of multiscaling		Statistics	
		Sig (%)	Distributional (%)	Temporal (%)	Mean $\hat{B}$	SD($\hat{B}$)
very rough	0.001	100.0	95.3	4.7	-0.0802	0.0058
	0.005	100.0	90.3	9.7	-0.0767	0.0064
	0.010	100.0	78.4	21.6	-0.0747	0.0068
moderate	0.050	99.9	29.1	70.9	-0.0482	0.0102
	0.100	98.6	10.0	90.0	-0.0287	0.0101
	0.200	71.5	3.5	96.5	-0.0095	0.0087

- **Very rough ($H \leq 0.01$):** strong multiscaling, $|\hat{B}| \approx 0.08$, always detected, **78–95% distributional**.
- **Moderate roughness:** weaker multiscaling; the temporal share rises from 71% ($H = 0.05$) to 97% ($H = 0.20$).

Two regimes, two mechanisms



grey labels: Stage 1 detection rate

Strong, tail-driven ($H \rightarrow 0$)

$|\hat{B}| \approx 0.08$, detected in every path,
inherited from the heavy-tailed marginal.

Mild, clustering-driven (moderate H)

$|\hat{B}| \approx 0.03\text{--}0.05$ in the calibrated range,
driven by volatility clustering.

Crossover around $H \approx 0.05\text{--}0.10$.

Calibrated models sit on the temporal side.

A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice

The mechanism: tails fade, clustering emerges

A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design


Method

Validation


Results

In practice

Kurtosis of returns (log scale)



Volatility clustering ($\sum_{k=1}^{10} \text{ACF}_{|r|}(k)$)



- Kurtosis falls by an order of magnitude: median ~ 50 at $H = 0.001$, ~ 5 at $H = 0.2$.
- Clustering rises from ≈ 0 for $H \leq 0.01$ to ~ 3 at $H = 0.2$.
- In the calibrated range ($H \approx 0.05$ – 0.10) clustering is already strong, while kurtosis is two to four times lower than in the rough limit.

Boxplots over 1000 Monte Carlo replicates per H .

What it means for rBergomi users

Use	Relevant H	What drives multiscaling	Take-away
Option pricing and calibration	$\approx 0.05\text{--}0.15$	volatility clustering	mild multiscaling, consistent with stylised facts; pricing value untouched
Very rough fits, $H \rightarrow 0$	≤ 0.01	fat-tailed marginal	strong multiscaling that is a tail effect (kurtosis ≈ 50)
Risk across horizons	either	changes with H	know which channel a scaling law extrapolates before using it

Bottom line

The same value of $\hat{B}$ can come from tails or from memory. Matching the *level* of multiscaling does not validate a model; matching its *source* does.

A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice

A diagnostic for any model or market



A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice

The two-stage test needs only a time series, so it can be run on simulated paths and on market data alike.

- **Model validation:** run it on paths of a calibrated model and on the market series. Do presence and source agree?
- **Model choice:** rough and non-rough two-factor models can fit SPX options comparably (Rømer 2022). The source of their multiscaling is a further way to tell them apart.
- **Risk across horizons:** before extrapolating a scaling law to long horizons, check whether it reflects persistence (temporal) or tails (distributional).
- **Next step:** the empirical application, which also revisits the opposite-sign result we found on real data.

- **Market data:** decompose the multiscaling of real returns into distributional and temporal parts.
- **Other rough models:** rough Heston, quadratic rough Heston, multivariate rough volatility with asset-specific H . Is the regime transition generic?
- **A distributional but genuine benchmark:** Dai Pra & Pigato (2015). Lévy-driven volatility with power-law tails gives analytically characterised multiscaling: moments scale as $h^{q/2}$ below q^* and as $h^{A(q)}$, $A(q) < q/2$, above it.
- **Joint variation of (H, η)** along the calibration ridge, rather than fixing η .

- Multiscaling in the rough Bergomi model has **two sources, not one**:
 - very rough ($H \leq 0.01$): strong and **tail-driven**, 78–95% distributional;
 - moderate ($H \geq 0.10$): mild and **clustering-driven**, 90–97% temporal;
 - crossover around $H \approx 0.05$ – 0.10 , explained by falling kurtosis and rising volatility clustering.
- Our earlier finding is confirmed and reinterpreted: the dramatic multiscaling at small H is largely *distributional*.
- Methodology: a two-stage surrogate test (matched fBm, then shuffles with a distance-based p -value), with WLS and data-driven q and τ ranges, validated on MRW and FLSM.

Take-home message

- Test against the right null: uniscaling processes give $\hat{B} \neq 0$ in finite samples.
- Heavy tails can masquerade as multiscaling; shuffled surrogates tell them apart.
- In rough volatility, *how rough* decides *which* multiscaling you get.

"The map is not the territory. The only usefulness of a map depends on its similarity of structure to the territory."

— A. Korzybski

"It doesn't matter how beautiful your theory is, it doesn't matter how smart you are. If it doesn't agree with experiment, it's wrong."

— R. Feynman

Models, however complex and realistic, must be tested for **falsification**, not just confirmation. Points of failure are often more informative than points of success: a model that fits everywhere but fails at the critical moment is the most dangerous map of all.

Thank you



A Tale of Tails

G. Brandi

Scaling

Rough vol.

Tale of tails

Idea

Design

Method

Validation

Results

In practice