The aggregation law: why an index is less intermittent than its stocks, and why it cannot be why the index is rough

Evan Tabak Atlas — ORCID 0009-0007-7374-2338

Fenopan standalone note — prose-complete draft (2026-08-28), restated after an adversarial audit 2026-08-30 (see the audit note below). Sealed records: Fenopan x8 synthetic half (record v7) and real-data half (record v17). Epistemic tags per CONSTITUTION §0. Public-safe: quotes only sealed measurements and fetched external sources; carries no strategy. Registered in PUBLISHING §3 (batch J).

> **Dated note (2026-08-30) — what this revision changed.** An adversarial audit of record v17 found that several sentences in the 2026-08-28 draft were stronger than the measurements support. Every one is restated below, in the direction the measurements point: the real-data half is downgraded from "closes the necessity gap" to *consistent with, but not powered to exclude, an aggregate as rough as the index at this length*; "$\hat{H} = 0$ at every D (grid max 0.000)" becomes *the per-D median $\hat{H}$ is 0 at every D; 3 of 41 subset fits are non-zero, max 0.050*; the estimator-free band is stated both ways; the ρ interval becomes a boundary statement with its resampling movement; the comparator's cross-measurement nature and the two matched comparators are named; the "predicted by D_eff = 180" framing is withdrawn against the measured log-space D_eff; ρ is redefined as the quantity the estimator identifies; and the necessity argument is scoped as first-order with its measured second-order bound. **Nothing was softened by omission** — where the corrected reading is weaker, it is stated as weaker. The corrective computations, their seeds and their tables are in `aggregation-law-audit-addendum-2026-08-30.md`, which is the citable basis for every number this revision moved.

Abstract

Stock indices are **less intermittent** than their constituent stocks and, at the same time, **rougher**: the log-volatility of an index has a smaller intermittency coefficient λ^2 and a larger monofractal Hurst exponent H than the log-volatility of the individual names it is built from (Wu, Muzy & Bacry 2022). Two mechanisms have been proposed for the intermittency half — aggregation destroying log-correlation, or a realized-volatility measurement artifact — and a nested factor model with a rough common factor has been *fit* to reconcile both halves at once (Zarhali, Aubrun, Bacry, Bouchaud & Muzy 2025). This note reports a sharper result than a reconciling fit: a **closed-form aggregation law** with a **necessity structure**. We show, first analytically and on synthetic aggregates of known cascades, and then **estimator-free on the committed 180-stock S&P-500 panel**, that aggregating D imperfectly-correlated **H = 0** cascades into an equal-weight

aggregate variance drives the intermittency down by **λ^2_agg = $\rho\lambda^2_0$ + $(1-\rho)\lambda^2_0$/D_eff** — so the *lower intermittency of an index is aggregation* — but **cannot raise H**: a weighted sum of H = 0 fields is again H = 0, because linear mixing preserves the correlation exponent. The roughness half is therefore not aggregation, not iid microstructure noise (excluded separately), and not constituent-level liquidity (excluded here on the own panel): a **genuinely rough common volatility factor is necessary**, not merely one sufficient modelling choice. The structural argument is **first-order** in the (weak) intermittency, and the real panel prices its own second-order term: on the committed 180-stock S&P-500 panel the covariance-mixture law holds estimator-free to **[0.958, 1.016]** at the six dyadic lags the slopes are read on and **[0.894, 1.046]** across all integer lags 1…32. The calibration-validated slope-ratio is **$\rho \approx 0.58$** (calibrated 0.587) with a moving-block-bootstrap 95 % interval of **[0.43, 0.65]** at the sealed 400-replicate run; converged to 2,000 replicates at the same block length the interval is **[0.42, 0.65]** and the sealed synthetic comparator 0.65 falls just **outside** it, so the honest statement is that **0.65 sits on the boundary of the interval, and the boundary moves ±0.03 with the resampling choices** — and the comparator is itself cross-measurement (world indices over US stocks at unmatched Δ and N; the S&P 500 itself reads ratio 1.095, and the matched-N read is 1.03). The single-common-cascade approximation is measurably violated — $\rho(\tau)$ **falls with lag**, the fingerprint of exactly the rough, shorter-scale common factor the necessity argument forces. The delta over the nested-factor fit is thus a **necessity argument** where the prior work supplied a sufficiency fit — established structurally and confirmed on synthetic aggregates, and **consistent with, but not powered to exclude, an aggregate as rough as the index** on the real panel at its 1,259-day length.

1. The milieu inversion

The object throughout is the **log-volatility field** — the logarithm of the realized variance, `ω = ln RV`, the "full field" of Wu, Muzy & Bacry's log-S-fBM convention `M(dt) = exp(ω)dt` [VERIFIED — Wu, Muzy & Bacry 2022, *Physica A* 604:127919]. Two numbers summarise its second-order structure at matched (Δ, T, N): the **intermittency coefficient** λ^2 (the amplitude of the logarithmic decorrelation of the field, the multifractal / log-correlation strength) and the **Hurst exponent** H (0 for a purely log-correlated / MRW field, positive for a monofractally rough one; $\lambda^2 = H(1-2H)\nu^2$ is the recoverable combination, and H and ν^2 are separately identifiable only away from the $H \to 0$ ridge).

The empirical puzzle Wu, Muzy & Bacry named is an **inversion of milieu**: read through one estimator at matched scale, stock **indices** — the widest milieu — come out *less intermittent* (small λ^2) and *rougher* ($\hat{H} \approx$ 0.11–0.13) than the **individual stocks** that compose them (larger λ^2, $\hat{H} \approx 0$). The Feno program reproduced both halves with an independent variogram-GMM estimator (Fenopan x2, sealed record v1): 21 Oxford-Man world indices at $\hat{H} \approx 0.13$, 180 S&P-500 constituents at $\hat{H} \approx 0$, with the individual-stock intermittency about 1.5× the index's ($\lambda^2 \approx 0.058$ vs 0.038; ratio 0.65) [REPO — Fenopan v1].

What that 0.65 is, exactly [REPO — Fenopan v1, read out 2026-08-30]. It is a **cross-market, cross-measurement, cross-N** ratio: the median λ^2_sfbm over 21 Oxford-Man **world** indices (5-minute RV5, N ≈ 4,200, 2000–2016) over the median over 180 **US** stocks (daily Garman-Klass, N = 1,259, 2013–2018). It is produced by the non-US indices in that set; the US ones sit at or near the constituent median, and **the S&P 500 itself (SPX2) reads λ^2_sfbm = 0.0639 — a ratio of 1.095, above the stock median, not below it** (DJI2 1.02, IXIC2 0.99, RUT2 1.42). The same sealed record's **matched-N cell**, with both cohorts truncated to N = 1,259, reads **indices/stocks = 1.03** with overlapping CIs. So the milieu inversion in λ^2 is a **statement about the world index cohort at its own length**, not a within-market statement about the S&P 500 at the constituents' length — and every comparison against "0.65" later in this note carries that caveat.

Two mechanisms were on the table for the *intermittency* half:

- **Aggregation destroys log-correlation.** An index is a sum over many units; if those units are imperfectly correlated, the sum averages away the idiosyncratic part of each unit's log-fluctuation, leaving a weaker common signal — lower λ^2.
- **A realized-volatility measurement artifact.** The index's realized variance is measured on a smoother, better-sampled series; the estimate of λ^2 could be depressed by the measurement process rather than by anything about the volatility (Cont & Das 2022, "Rough volatility: fact or artefact?" — the standing objection that measured roughness/intermittency of realized volatility is largely microstructure noise and discretization) [VERIFIED — Cont & Das 2022].

And the *roughness* half — the index's positive H — had its own artifact objection (same family: measured roughness as a discretization / microstructure phenomenon) and its own genuine-mechanism candidate (a rough common factor). The program's §3 "Law of Genesis" clause (iv) — *the measured object is the unit where the multiplicative feedback closes, not an aggregate over many such units* — was [DERIVED / single-source], resting entirely on this one inversion. It could not be promoted until the aggregation mechanism was tested against the artifact null on objects where the truth was known.

2. The closed form

Write each constituent's log-variance field as a common component plus an idiosyncratic one on the same rate ladder,

> $\omega\_i = \omega\_\text{common} + \omega\_\text{idio},i$, with $\mathrm{var}(\omega\_\text{common}) = \rho\cdot\mathrm{var}(\omega_0)$, $\mathrm{var}(\omega\_\text{idio}) = (1-\rho)\cdot\mathrm{var}(\omega_0)$,

so that every constituent is exactly a cascade of the reference amplitude λ^2_0 (at H = 0) and, **in this noise-free model**, the contemporaneous cross-sectional correlation is exactly ρ. The **aggregate variance** is the (equal- or cap-) weighted sum of the constituent *variances*, `V_agg(t) = Σ_i w_i exp(ω_i(t))`, and the question is what `(Ĥ, λ̂²)` its log field carries.

> **What ρ is, as measured** [REPO — Fenopan v17 + the 2026-08-30 addendum]. Real log-RV carries idiosyncratic measurement noise (the fitted per-stock $\hat{\sigma}^2\_\varepsilon \approx 0.40$ in the sealed x2 record), and that noise sits **only on the within diagonal**. The estimator used throughout §5 — the ratio of the cross-covariance slope to the within-covariance slope — therefore identifies the **ratio of common-to-total log-correlated covariance amplitude, noise-excluded**, and *not* the contemporaneous cross-sectional correlation. The two differ by a factor of two on this panel: **$\rho(0) = \bar{C}\_cross(0)/\bar{C}\_within(0) = 0.285$**, and the mean pairwise lag-0 correlation of the 180 demeaned log-RV series is **0.287**, against $\rho(\tau \geq 1) \approx 0.51$–$0.56$. Excluding the noise is the right thing to do for a covariance-slope law; naming which ρ is meant is the honest thing to do, and this note means the noise-excluded amplitude ratio everywhere below.

To first order in the (weak) intermittency the aggregate log field is a weighted mix of the constituent log fields, and the algebra of the sample covariance of a linear combination gives a **covariance-mixture identity**:

> $$C_{agg}(\tau) \approx (1/D)\cdot\bar{C}_{within}(\tau) + (1 - 1/D)\cdot\bar{C}_{cross}(\tau),$$

where `C̄_within` is the mean own-autocovariance of a constituent and `C̄_cross` the mean pairwise cross-covariance. Because the cross-covariance is the *common* component only (`C̄_cross = ρ·C₀`) and the within is the total (`C̄_within = C₀`), the aggregate covariance is a scalar multiple `[ρ + (1−ρ)/D_eff]` of the constituent covariance — and since λ^2 is (minus) the slope of that covariance in `ln(τ+1)`, the intermittency obeys

> $$\lambda^2_{agg} = \rho\cdot\lambda^2_0 + (1 - \rho)\cdot\lambda^2_0 / D_{eff}, \quad D_{eff} = 1 / \Sigma_i w_i^2$$ (the inverse Herfindahl / effective breadth).

Two readings fall out immediately. The **common part contributes a floor** $\rho\lambda^2_0$ that survives any amount of aggregation; the **idiosyncratic part is diversified away** as $1/D_{eff}$. As breadth grows the aggregate intermittency descends from the constituent value λ^2_0 to the floor $\rho\lambda^2_0$ — the lower intermittency of an index is, under this law, *exactly aggregation*, controlled not by the raw count D but by the effective breadth D_{eff} (which is why a cap-weighted index of hundreds behaves like a much smaller equal-weight basket).

The same first-order argument settles the roughness half in the opposite direction. A weighted sum of log-correlated **$H = 0$** fields is **again a log-correlated $H = 0$ field**: linear mixing preserves the correlation *exponent*, it only rescales the amplitude. Roughness ($H > 0$) is a different spectral shape, and no linear combination of $H = 0$ inputs can create it. So the law predicts, in one breath, that **λ^2 falls** (by the closed form) while **H stays exactly 0**.

The scope of that argument, stated plainly. The object actually measured is `ln Σ_i w_i exp(ω_i)`, a *nonlinear* functional; "linear mixing preserves the exponent" applies to its **linearisation**. The argument is therefore **first-order / perturbative in the intermittency**, not an exact theorem about the nonlinear object, and it is not stated as one here. What makes it usable is that the neglected second-order term is **measured, not assumed**: the ratio C_agg/C_mix *is* that term, and on the real panel it is within 4.2 % over the six dyadic slope lags, within **10.6 %** over all integer lags 1…32, and degrades to 28 % out to lag 64 (§5). A non-perturbative argument — that `ln Σ_i w_i e^{ω_i}` of H = 0 inputs is itself H = 0 — is **not supplied here** and is the cleanest available strengthening of this note. [gap — registered]

3. The synthetic grid

Fenopan x8's synthetic half (sealed record v7) is the two-world test of that prediction on aggregates of *known* cascades [REPO — Fenopan v7]. D constituent log-variance cascades of a known (λ^2_0, H = 0) are correlated at a swept $\rho \in [0, 1]$ by the additive common-factor construction above — exact at both pre-registered anchors (D = 1 and $\rho = 1$ reproduce the constituent by construction) — and the aggregate variance is re-estimated with the same vendored variogram-GMM estimator the whole amplitude ladder uses.

The result is the closed form, measured. The aggregate $\hat{\lambda}^2$ falls from the constituent value toward the floor $\rho\lambda^2_0$ across the whole 6 ρ × 5 D grid, captured by the first-order law to a **median 2.3 %** (and holding at a higher constituent amplitude too); the negative-control anchors are identities and the $\rho = 0$ positive control collapses $\hat{\lambda}^2$ toward zero as D grows. The observed index/stock intermittency ratio (≈ 0.65) is reproduced at a plausible equity log-volatility cross-correlation **ρ ≈ 0.65** [REPO — Fenopan v7]. Clause (iv) was promoted [DERIVED / single-source] → [REPO] for the intermittency claim: an index really is less intermittent than its units *because aggregation destroys log-correlation*.

4. The refuted Ĥ-half — and why that is a necessity proof

The grid's second finding is the one that matters most here, and it arrived as a refutation of the experiment's own pre-registered prediction (the measurement outranks the prompt). The registered discriminator stated that the aggregation mechanism would make Ĥ **rise** as well as $\hat{\lambda}^2$ fall. It does not. Across the entire grid the aggregate Ĥ(sfbm) has **median 0.000** (q75 0.000; worst single cell 0.0019), and this survives a heterogeneous-constituent probe (per-constituent depth and amplitude drawn independently) and the matched stock-scale length. **Aggregation destroys the intermittency amplitude but does not manufacture roughness** [REPO — Fenopan v7].

That negative is the load-bearing result, because it converts a *sufficiency* question into a *necessity* one. The literature's move is to *posit* a rough common factor and show it *can* reproduce the inversion — a sufficiency argument (§6). The refuted Ĥ-half says something stronger and in the other direction: whatever ρ and D one chooses, aggregation of H = 0 constituents **cannot** produce a rough aggregate, so the index roughness has

no aggregation explanation at all. Every synthetic cell that reproduces the observed index λ^2 sits at $\hat{H} \approx 0$, two orders of magnitude below the observed index $\hat{H} \approx 0.13$ — and, crucially, *on synthetic aggregates the truth is known and the instrument's power is not in question*: the grid is long enough and the inputs are $H = 0$ by construction, so a null there is a measurement. Aggregation is thereby *eliminated* as a candidate for the roughness — not disfavoured, eliminated, at first order, by a structural argument (linear mixing preserves the exponent) whose second-order term the same record prices (§2). Combined with the independent exclusion of the iid-microstructure-noise artifact for the index roughness (x2: the index reads rough through the noise-robust realized kernel, RK ≈ RV5), the field of candidates for the index's positive H narrows to a **genuinely rough common volatility factor** — the milieu inversion splits into an **aggregation half (λ^2)** and a **separate, non-aggregation roughness half (H)**.

The necessity argument had one gap: it was established on *synthetic* aggregates. A sceptic could still hope that real constituents, aggregated for real, would somehow read rough — that the synthetic $H = 0$ inputs were the artifice. The real-data half addresses that gap; it does **not close it**, and this note does not claim it does.

Why not, said before the result [REPO — Fenopan v1, `ladder.matched_N_financial`]. The real-data half works at the constituents' own length, **N = 1,259 days**. At exactly that length the program's own sealed record measures the index cohort at **$\hat{H}$ = 0.0245** — about 19 % of its full-N read of ≈ 0.13 — and states verbatim that "at matched N = 1259 neither the index nor the stock length resolves the index roughness (H_index collapses) … the index roughness is unresolvable at the stock length". So a null $\hat{H}$ on a 1,259-day aggregate is **what the program's own predecessor predicts would be observed even if that aggregate were as rough as the index**. The real-data half is therefore read below against the **matched-N comparator $\hat{H}$ = 0.0245**, never against the full-N 0.13, and its verdict is **consistent with, but not powered to exclude, an aggregate as rough as the index at this length**. The measurement that would change that — an H-detection positive control: inject a known-H ≈ 0.11–0.13 log-S-fBM common factor into the panel at the measured $\hat{\sigma}^2\_\varepsilon \approx 0.40$ and N = 1,259, and report the recovered $\hat{H}$ distribution — has **not been run**, and is the single most valuable next cell for this lane. [gap — registered]

5. The real-data confirmation

Fenopan x8's real-data half (sealed record v17) runs the law on the **committed 180-stock panel** — daily log-Garman-Klass realized variance for 180 S&P-500 constituents over 1,259 trading days on one matched date grid [REPO — Fenopan v17]. The matched index is not fetched from outside; it is the **equal-weight aggregate variance of the constituents themselves**, `V_agg = (1/D)Σ_i RV_i`, an index-analog at *identical* (Δ, T, N). (The cap-weight arm is dropped: no market caps are committed, so constituent ranks are unidentified.) Five things are then measured — and, per §4, all of them at **N = 1,259**, the length at which the program's own sealed ladder puts the index cohort at $\hat{H}$ = 0.0245 rather than 0.13.

The per-stock refits reproduce the sealed rung — a determinism check, not an independent one. The 180 constituents refit to a cross-sectional λ^2 with median **0.0583** and cross-sectional **mean λ^2_0 = 0.0633**; their $\hat{H}$ is ≈ 0 (86 % below 0.02). The stocks are the pure cascade, exactly as the inversion says. The median is *bit-identical* to the sealed x2 stock rung (0.0583376130574831 in both records, with the $\hat{H}$ medians agreeing to all 17 digits) — same data, same SHA-pinned estimator, same fit seed, so this is a **regression check that the pipeline is unchanged**, and it carries no independent corroborative weight. It is reported here as that, not as a second measurement.

The equal-weight aggregate is less intermittent, and its per-D median $\hat{H}$ is 0 at every breadth. The full aggregate reads λ^2 = **0.0350** at $\hat{H}$ = **0.000**, well below the constituent mean; across a breadth sweep D ∈ {2 … 180} the intermittency falls monotonically toward the floor while the $\hat{H}$ **median stays 0 at every D**. The per-cell picture, which the record's summary statistic hides: of the **41 individual subset fits, 3 are non-zero — $\hat{H}$ = 0.050 (D = 5), 0.019 and 0.003 (D = 10)** — and the largest of them is twice the matched-N index read of 0.0245. So the honest sentence is **"the per-D median $\hat{H}$ is 0 at every D; 3 of 41 subset fits are non-zero, max 0.050"**, not "grid max 0.000". This is the necessity argument's synthetic H = 0 inputs *replaced by real constituents*: at this length the real equal-weight aggregate reads not rough — which, per §4, is consistent with an aggregate as rough as the index and does not exclude it.

The covariance-mixture law holds with no estimator in the loop — over the lags the slopes are read on. The measured autocovariance of the *nonlinear* `ln V_agg` matches the covariance-mixture `(1/D)C̄_within + (1−1/D)C̄_cross` closely: the ratio C_agg/C_mix sits in **[0.958, 1.016]** at the **six dyadic lags {1, 2, 4, 8, 16, 32} the λ^2 slopes are read on**, and in **[0.894, 1.046]** — a **10.6 %** band — across **all integer lags 1…32**, degrading to [0.838, 1.282] out to lag 64. Both must be stated: the first is the claim the slope reads rest on, the second is what an interval claim about "out to lag 32" asserts. Two further scope notes. The mixture **identity** is *algebraic by construction* — C̄_cross is computed as (D·C_linear − C̄_within)/(D−1), so the mixture reproduces C_linear to round-off (2.8×10^{-17}) by definition, and that residual validates nothing; what is measured is the **log-of-sum vs mean-of-logs (Jensen) gap**, i.e. exactly the second-order term §2 needs bounded. And the synthetic positive control for this test is run on a *homogeneous* panel and returns [0.995, 0.999], an order of magnitude tighter than the real band — so it shows the test manufactures no gap where there is none, but it does not anchor the test in the regime it is applied to. With those scopes, the aggregation law is a *directly measured* property of the real panel, prior to and independent of any GMM fit.

The slope-ratio reads ρ ≈ 0.58, and 0.65 sits on the boundary of its interval. Reducing the panel to the closed form's single free parameter, the estimator the two-world calibration validates is the *pure, self-consistent slope-ratio* — the common-factor covariance slope over the within-stock covariance slope, both read the same way — and it gives **ρ ≈ 0.581** (calibrated **0.587**), with a moving-block-bootstrap 95 % interval of **[0.43, 0.65]** at the sealed run's 400 replicates and block length 63. The like-for-like GMM/GMM analog of how the sealed 0.65 was formed (a fitted aggregate over the fitted cross-sectional constituent amplitude, exactly x2's index/stock construction) reads **0.55–0.60**. So every principled estimate sits at ~0.55–0.60, *below* the sealed 0.65. **Four qualifications belong beside that, and none of them was in the 2026-08-28 draft.**

(i) The interval is a boundary statement, not a containment one. The sealed interval `[0.4286, 0.6507]` contains 0.65 by **0.0007**. Raised to **n_boot = 2000 at the same declared block length**, the raw interval is **[0.4217, 0.6492]** and the calibration-inverted one **[0.4335, 0.6480]** — **both exclude 0.65**; the sealed calibrated interval already excluded it (hi95 = 0.6494). Across block lengths {21, 63, 126, 252} at 2000 replicates the upper endpoint runs 0.623 → 0.685 and the bootstrap median 0.523 → 0.599, tracking the fraction of long-lag pairs destroyed by the block joins. **0.65 sits on the boundary of the interval, and the boundary moves ±0.03 with the resampling choices**; at the declared block length, converged, the interval excludes it.

(ii) ρ is band-conditional. With the identical estimator and only the slope band changed: **0.581** (declared dyadic 1…32), **0.655** (dropping lag 1), 0.555 (dyadic 1…16), 0.531 (integers 1…8), **0.610** (all integers 1…32). Dropping one lag reaches the comparator. The reason is visible in the fit quality: over the declared band the *within* curve fits ln(τ+1) at $R^2 = 0.926$ with a systematic +0.029 residual at lag 1, while the *cross* curve fits at $R^2 = 0.970$ — the log-linear form the slope read assumes is misspecified for the denominator at short lags. Neither the bootstrap nor the calibration carries any lag-band uncertainty.

(iii) The comparator is cross-measurement. Per §1, the sealed 0.65 is a world-index median over a US-stock median at unmatched Δ and N; the panel's own market reads **SPX2 ratio 1.095**, and the matched-N read is **1.03**. v17's 0.58 (equal-weight aggregate of S&P constituents against those same constituents) is therefore not being scored against a like quantity.

(iv) The "predicted offset" is withdrawn. The 2026-08-28 draft explained the 0.07 shortfall as the internal aggregate's higher effective breadth (D_eff = 180). It does not work quantitatively, and 180 is the wrong D_eff for this object: the aggregate is equal-weight **in variance**, so its log-space weights are ∝ RV level, and the measured inverse Herfindahl is **135** (mean-RV weights) or **72.5** (instantaneous median). Under the record's own closed form, moving from D_eff = 180 to the measured values shifts ρ by ≤ **0.004**, and even an implausible D_eff = 20 shifts it by 0.019 — against a gap of **0.069**. **The closed form attributes at most ~0.02 of the 0.07 gap to breadth**; the gap is not explained here.

(A fifth, smaller point: substituting the GMM constituent amplitude under the covariance-slope numerator — a mixed estimator the calibration does not back — inflates the ratio to 0.62–0.67 and reaches 0.65; those reads are the estimator-sensitivity picture, not the headline. And the two-world calibration validates the estimator on a **homogeneous** synthetic world; the real panel's factor-loading dispersion is CV(β) = 0.144, which attenuates a slope ratio by $1/(1+CV^2)$ — about **0.012 in ρ**, uncalibrated, and in the direction that manufactures this gap. A heterogeneous-loading calibration arm is now available in the apparatus; the sealed v17 calibration predates it.)

A measured refinement: the single-ρ closed form is violated. The lag-resolved correlation $\rho(\tau) = \bar{C}\_cross(\tau)/\bar{C}\_within(\tau)$ is **not flat** — a single common cascade of one integral scale would make it constant. It **falls with lag** (0.509 at $\tau = 1$, peaking 0.557 at $\tau = 4$, down to 0.422 at $\tau = 32$ and 0.406 at $\tau = 64$): the common factor **decorrelates faster than the idiosyncratic part**. (The contemporaneous value $\rho(0) = 0.285$ sits far below all of them; that step is the measurement noise on the within diagonal, per §2, not part of the decay.) The aggregation *amplitude* law is measurable estimator-free at the slope lags; the single-common-cascade *shape* is only a first approximation, and the direction of its failure is diagnostic — it is precisely the fingerprint of a **rougher, shorter-scale common factor**, the same object the necessity argument (§4) says the index roughness requires. The two independent lines — "aggregation cannot make H" and "the common factor is rougher and shorter-scale" — point at the same missing piece.

6. Positioning

Wu, Muzy & Bacry 2022 [VERIFIED] is the source of the inversion and of the log-S-fBM field convention this whole analysis is written in; the note's estimator is an independent reimplementation validated against their closed forms. Our contribution is not the phenomenon but its *decomposition*.

Cont & Das 2022 [VERIFIED] is the microstructure-noise objection: measured roughness and intermittency of realized volatility may be artifacts of discretization and noise rather than genuine dynamics. The program addresses it for the intermittency half in x8's artifact null (iid RV noise leaves λ^2 and H flat — the estimator's own noise variance absorbs it, so the *falling* λ^2 of an aggregate cannot be an iid-noise effect) and for the roughness half in x2 (the index reads rough through the noise-robust realized kernel). The two artifact nulls and the aggregation mechanism are separable in the (λ^2, H) plane.

Bennedsen, Lunde & Pakkanen 2022 ("Decoupling the Short- and Long-Term Behavior of Stochastic Volatility", *J. Financial Econometrics* 20(5):961–1006) [VERIFIED] is the careful rough-volatility modelling that a measurement/liquidity objection to roughness estimates has to be argued against: are low-H estimates genuine short-scale roughness, or an imprint of liquidity and measurement error? We address the constituent-level version of that objection **on the own panel**: per-stock $\hat{H}$ is ≈ 0 across the cross-section, and its association with the microstructure-noise proxy (the fitted measurement-noise variance $\hat{\sigma}^2\_\varepsilon$) runs the **wrong way for the artifact**. The Pearson correlation is −0.065, but 80 % of the per-stock $\hat{H}$ sit at machine zero, so a Pearson coefficient is the wrong summary of this cross-section; the **rank** statistic is **Spearman −0.21 (p = 0.004)** — a statistically significant association, and a **protective** one: *noisier constituents read LESS rough*, the opposite of what the artifact hypothesis predicts. (The inverse-liquidity proxy, the mean log-RV level, is +0.24 Pearson / +0.15 Spearman, p = 0.044 — also small, also not the artifact's direction of interest.) If the measured roughness were a liquidity/measurement artifact, the noisier or less-liquid constituents would read systematically rougher. They read *less* rough — so the roughness that appears in the common factor is not manufactured at the constituent level by liquidity. The 2026-08-28 draft's "essentially uncorrelated (−0.065)" was the wrong summary; the directional statement is both the correct one and the stronger one.

Zarhali, Aubrun, Bacry, Bouchaud & Muzy 2025 — *A Nested Factor Model for Equity Markets: Reconciling Multifractal Stock Returns and Rough Index Volatilities* (arXiv:2505.02678) [VERIFIED — fetched] — is the closest prior work and the one against which the novelty claim must be stated precisely. It builds a **nested factor model** on the *same* S&P-500 data, with a **rough common (factor) volatility** ($H \approx 0.11$) and **multifractal, super-rough idiosyncratic residuals** ($H \approx 0$), and shows this structure **reconciles** the multifractal stock returns with the rough index volatilities. It is, in the terms of this note, a **sufficiency result**: a rough-common-factor model *can* reproduce the inversion, and its fitted exponents (common $H \approx 0.11$, residual $H \approx 0$) agree strikingly with the program's own independent reads (index/common $\hat{H} \approx 0.13$, stock $\hat{H} \approx 0$). The nested factor structure is the natural generative account, and the agreement is corroboration, not conflict.

The **delta** is the direction of the argument. Zarhali et al. *posit* the rough common factor and demonstrate it is **sufficient**. This note argues it is **necessary**: the closed-form aggregation law — established structurally at first order, measured on synthetic aggregates where the truth is known, and confirmed estimator-free on the real panel at the slope lags — says the *other* candidate for the index roughness, plain aggregation, **cannot** produce it, while the two artifact nulls (iid microstructure, constituent liquidity) are excluded separately. A rough common factor is then not one modelling choice among several that happens to fit; it is what remains after the alternatives are eliminated. The nested factor model supplies the mechanism; the aggregation law supplies the argument that a mechanism of that kind is required — and the measured single-ρ violation ($\rho(\tau)$ falling with lag) supplies an independent, model-free signature of exactly its rough, shorter-scale character.

Two honest limits on that delta. (a) The elimination is carried by the *structural* argument and the *synthetic* grid, both first-order; the real-panel half is **consistent with** it but, at $N = 1,259$, not powered to exclude an aggregate as rough as the index (§4). (b) The delta is a claim about what the body of Zarhali et al. does and does not argue. The abstract was independently re-fetched 2026-08-30 and is framed as sufficiency, with no aggregation-cannot-make-roughness argument — but **this repository carries no committed cite pack** (DOI, fetch date, and a quoted extract for each load-bearing attribution), so the delta cannot be re-checked offline the way Fenocrisis's cite pack or Fenocosm's paper anchors can. Building one is a **pre-deposit requirement** for this note. [gap — registered]

7. The reframed novelty claim, and what is open

The naive novelty claim — "aggregation explains the milieu inversion" — is both Zarhali-adjacent and only half true, and this note does not make it. The claim it makes is the **necessity decomposition**:

> Aggregation of imperfectly-correlated $H = 0$ units accounts for the **intermittency half** of the milieu inversion, by a closed-form law that holds estimator-free on real data at the lags the slopes are read on (validated slope-ratio $\rho \approx 0.58$, block-bootstrap CI [0.43, 0.65] at the sealed 400-replicate run — an interval on whose **boundary** the synthetic comparator 0.65 sits, and which excludes it once converged at the declared block length); and aggregation is, **to first order in the intermittency and with the**

> **second-order term measured**, **incapable** of accounting for the **roughness half**, so — with the microstructure-noise and constituent-liquidity artifacts excluded independently — a **genuinely rough common volatility factor is necessary**. The contribution over the prior nested-factor *fit* is this *necessity*, and a model-free signature of the factor's character (ρ(τ) falls with lag).
>
> **Scope, in the same breath:** the necessity argument is carried by the structure and by the synthetic grid, where the truth is known. The real-panel half is **consistent with, but not powered to exclude, an aggregate as rough as the index at N = 1,259** — the program's own matched-N read of the index cohort is $\hat{H}$ = 0.0245, and no H-detection positive control has been run on this panel.

Three questions are handed forward with this note; the first two are formally in this lane (the ROADMAP index-roughness orphan is closed here), the third is the pre-deposit repair of the real-data half:

1. **Estimate the rough common factor directly.** The single-ρ violation says the common factor is rougher and shorter-scale than a single log-correlated cascade; the natural next step is to fit its (H, λ^2, integral scale) from the cross-covariance structure and compare to the nested-factor exponent (H ≈ 0.11).
2. **The cap-weight extension.** The committed panel carries no market caps, so only the equal-weight aggregate could be built. A license-clean market-cap or cap-weighted-index derivative would let the D_eff axis be read as a weighting contrast (the closed form predicts a cap-weighted index sits nearer the constituent than an equal-weight basket of the same breadth), and would place the real ρ against the sealed 0.65 through the *same* aggregate the synthetic grid was calibrated to.
3. **The H-detection positive control (pre-deposit).** Inject a known-H ≈ 0.11–0.13 log-S-fBM common factor into the committed panel at the measured per-stock σ^2_ε ≈ 0.40 and N = 1,259, and report the recovered $\hat{H}$ distribution. Only if the instrument recovers ≳ 0.10 at this length does the real-data null carry necessity weight; until it runs, §4's downgraded reading stands. A committed cite pack for the four external references (§6) travels with it.

The measurement, as always, outranks the note: the closed form, the estimator-free confirmation at the slope lags, and the eliminations are the deliverable; the rough common factor is named, bounded, and left open for the fit that will follow — and where the evidence is weaker than the 2026-08-28 draft said, this revision says so rather than saying less.

Data availability

The only data read is `data/finance/stocks_logrv.npz` — a committed, license-clean **derived** log-Garman-Klass realized-variance series for 180 S&P-500 constituents over 1,259 trading days, with its per-source license basis, extract tool and pinned source SHA-256s recorded in `data/finance/PROVENANCE.md`. No raw price, quote or OHLC table is redistributed. The sealed records this note quotes are Fenopan **v1**, **v7** and **v17**, each an append-only directory carrying its own JSON, `RESULTS.md` and regeneration commands.

The repository itself is private and remains so (owner decision, 2026-08-30). The records and this note are deposited as **standalone Zenodo objects** — the note carries no repository links, and the deposit descriptions carry each record's own regeneration block instead, so a reader can reconstruct every number from the deposited object without repository access. Requests for the derived panel or the apparatus should be addressed to the author.

> **Dated note (2026-08-30) — DOI wiring.** Zenodo drafts with reserved DOIs now exist for the program roster; the owner presses Publish later, and staging is not publication. Of the three sealed records this note quotes, only **v17** was in the batch-0 reserve-and-wire roster and carries a reserved DOI: **10.5281/zenodo.22180458** (reserved; publication pending). **v1** and **v7** were outside that roster and had no reserved DOI as of this writing; both have been added to the deposit manifest this session (`FENOPAN_V1`, `FENOPAN_V7`) so they join the reserve set at the next sitting — see `PUBLISHING.md` §5.6. Neither carries a DOI yet.
>
> **Dated note (2026-08-31) — published.** The owner pressed Publish: the program's full reserve set (**48 records**) went live on 2026-08-31, resolution verified via the Zenodo API the same day (LEDGER 0043). Of this note's own objects, **v17 is published at 10.5281/zenodo.22180458** and **this note itself at 10.5281/zenodo.22180460** — so "reserved; publication pending" in the note above and in the reference list below is stale as of that date, kept as the record of what was true at the 2026-08-30 wiring. The **deposited PDF** of this note predates this correction and carries the stale phrasing; the fix, if taken, is an optional version-2 re-deposit under the same concept DOI (`PUBLISHING.md` §5.7) — cosmetic, reference notes only. v1's and v7's own reservation status stays tracked in `PUBLISHING.md` §5.6, not here.

References

- Bennedsen, M., Lunde, A. & Pakkanen, M. S. (2022). Decoupling the Short- and Long-Term Behavior of Stochastic Volatility. *Journal of Financial Econometrics* 20(5), 961–1006. (arXiv:1610.00332) [VERIFIED]
- Cont, R. & Das, P. (2022). Rough volatility: fact or artefact? [VERIFIED]
- Wu, P., Muzy, J.-F. & Bacry, E. (2022). From rough to multifractal volatility: The log S-fBM model. *Physica A* 604, 127919. [VERIFIED]

- Zarhali, O., Aubrun, C., Bacry, E., Bouchaud, J.-P. & Muzy, J.-F. (2025). A Nested Factor Model for Equity Markets: Reconciling Multifractal Stock Returns and Rough Index Volatilities. arXiv:2505.02678. [VERIFIED — fetched]
- Atlas, E. T. (2026). *Aggregation-law audit addendum — the corrective computations (2026-08-30)*. Fenopan, `notes/aggregation-law-audit-addendum-2026-08-30.md`. [REPO] — the citable basis for every number this revision restated.
- Fenopan sealed records **v1** (x2, the matched-scale amplitude ladder; DOI reserve pending, `FENOPAN_V1`), **v7** (x8 synthetic half; DOI reserve pending, `FENOPAN_V7`), **v17** (x8 real-data half; **10.5281/zenodo.22180458**, reserved — publication pending). [REPO] *(Dated correction, 2026-08-31: v17's record is published and resolves as of 2026-08-31 — the "reserved — publication pending" annotation is stale; see the dated note in §Data availability above.)*