From Latent Cascade to Note Stream: How Scaling Exponents Survive Musical Quantization

Evan Tabak Atlas

Independent · ORCID 0009-0007-7374-2338

Evanatlas@gmail.com

> Preprint draft. Target venue: Chaos or Physica A (physics framing), with a shorter ISMIR/SMC variant possible (music-technology framing).

Datasets and analysis code. The grid datasets come from `foundry sweep` (`engine/crates/foundry`): `research/data/sweep-v1.csv` (5×5 Spread × Intensity grid × 8 seeds × 8192 ticks × 3 launch scaffolds; per-cell medians in the companion `*.medians.csv`), `research/data/sweep-v2.csv` (the same grid at a **32-seed** panel — the working dataset for the transport results below), `research/data/sweep-v3.csv` (the same grid re-cut with latent-lane MFDFA width columns — the width-compression dataset; §4.3), and the `*-gatepoint.csv` companions (the exact gate operating point at both panel sizes). The estimator-side null-model comparison lives at `research/data/nulls-v1.csv` (gate point) and `research/data/nulls-e2-spread500-int{000,250,500,750,1000}-v1.csv` (the per-cell E2 slice; §4.3), all from `foundry nulls`; the known-Hurst estimator anchors (dense recovery, sparse-binary attenuation, template-removal no-op) are mechanized in `engine/crates/core/tests/spectral.rs`; and the mechanism model and grid maps of §4.4–§5 are the analysis package `research/analysis/quantization-mechanism/` (offline, seeded, `out/RESULTS-mechanism-v1.md` and `out/RESULTS-maps-v1.md`). Provenance note: `sweep-v2` measured the launch packs at `scaffold_version 1.0.0`; `sweep-v3` and the per-cell nulls are cut at the current tips, and the beat-driven and cinematic emitted measurements are **exactly unchanged** across the intervening pack bumps (800/800 rows string-equal per pack), so those two packs' numbers read directly against `sweep-v2`; warm-ambient `1.3.0` is a genuinely retuned stream and is reported at its own version throughout.

ABSTRACT

Generative systems built on multifractal processes make a quiet assumption: that the scaling structure of the *latent* process survives the discrete, nonlinear observation path that turns it into events. In the Fenophone — an instrument whose latent core is a Markov-switching multifractal (MSM) — that path is explicit and deterministic: band products are read into `[0,1]` lanes, mixed through a routing matrix into per-voice drives, thresholded into onsets, quantized onto a metric grid, and gated by harmony. We measure, across a control grid and seed panel, the detrended-fluctuation (DFA-1) exponents of the four latent lanes alongside those of the emitted onset and velocity streams, and the MFDFA width of both. We characterize the transformation empirically — routing decides *which* latent scaling the note stream inherits (only the routed mid-band lane's α transports; fixed-effects coefficient ≈ +0.32 against ≈ 0 for the unrouted fast band), while the channel's contrast decides *how much* of it registers (emitted α climbs ≈ 0.51 → 0.83 across the Intensity axis with every latent lane α constant, including exact low-contrast dead zones where adjacent dial positions emit bit-identical streams) — and explain the contrast axis with a level-crossing transfer model: the channel's quantizer is a monotone integer staircase, hence Hermite rank 1, so the latent exponent is preserved asymptotically at every contrast, and what contrast actually sets is the amplitude of the exponent-carrying component against the quantization-noise floor — the crossover scale a fixed measurement ladder reads as an effective α. A no-free-parameter mixture predictor built on that decomposition reproduces the measured climb cell by cell (median $|\Delta\alpha|$ = 0.009 on the model's theory rung). Three further maps complete the channel's characterization: the per-cell surrogate comparison shows the mid-Intensity width dip is a collapse into the finite-size floor (the dip cell beats neither its shuffle nor its IAAFT null), the latent-vs-emitted width map **refutes** the natural hypothesis that thresholding compresses multifractal width (emitted ≈ lane width; median per-run ratios 0.88–1.18), and the degeneracy map locates the usable region (the entire primary control plane is non-degenerate at 32 seeds, with all exact dead zones on one pack). The result is a practical design rule for multifractal instruments: a fitted α(Intensity, Spread) surface (R^2 0.89–0.996) to invert for targets, which latent scaling survives musical quantization, and by what mechanism.

1. Introduction

The engine's spectral gate (A.8 #10) asserts that the *emitted* note stream keeps 1/f-family correlation, and the measured landscape sits at α(onset) ≈ 0.70–0.78 — but the latent lanes that drive the stream have their own scaling, and the pipeline between them is a stack of nonlinearities: a clamped logarithmic band map, a routing mix, an integer onset quantizer, a metric grid, and harmony gating. Each stage is individually innocuous; jointly they are the *quantization channel*, and nobody has quantified what it does to a scaling exponent. The gap matters beyond this instrument. Any sonification or generative system that thresholds a long-memory process into events inherits the same question, usually invisibly — and so does the inverse problem (the companion corpus-fitting paper), which must invert the channel to recover latent parameters from note streams.

The Fenophone makes the question answerable exactly. Because the event layer is integer-deterministic, the latent and emitted series are *jointly* observable, bit-reproducibly, at any operating point: one `foundry sweep` invocation renders both sides of the channel for a full control grid × seed panel, and the estimator that measures them is the CI gate's own lineage, anchored on known-Hurst ground truth. The result below is, to our knowledge, the first end-to-end empirical map of a musical quantization channel — which exponents transport (§4.1–4.2), what the contrast axis really does and why (§4.2, §5.1), what happens to multifractal width (§4.3), where the control space degenerates (§4.4) — together with the design rule the map implies (§5.2).

2. The channel, precisely

Stages (each with the file that implements it):

1. **Band products → lanes** (`cascade.rs`): `lane = clamp(1/2 + log2(Π M_i)/(4n), 0, 1)` — a monotone, *clamped* log map of the band product; timbre additionally one-pole smoothed.
2. **Lanes → per-voice drive** (`routing.rs`): weighted average over bands (monotone, normalization preserves range).
3. **Drive → onsets** (`pipeline.rs`, A.3): the onset model — the density lane, sampled at each bar downbeat, sets the bar's onset count through a monotone integer `onsets(density)` curve bounded by the coherence bounds, and the onsets fill the heaviest cells of the authored metric template (the deliberately periodic skeleton).
4. **Onsets → notes**: harmony gating (legal pools), contour snap, velocity from dynamics.

The theory this construction should be read against is the theory of transformations of long-memory processes. Monotone static transforms preserve Hurst-type exponents of Gaussian long-memory processes (Taqqu-type results at Hermite rank 1); clamping and thresholding are not rank-1-smooth, but level-crossing and hard-threshold results predict exponent preservation for persistent processes, with variance compression (Taqqu 1975; Beran 1994; for level crossings of long-memory / fractional-Gaussian processes, Kratz 2006 and Vahabi, Jafari & Movahed 2011). Two honest caveats frame everything that follows: the lanes are not Gaussian (bounded, piecewise-constant-ish switching sums), and the theory is asymptotic while every measurement below lives on a finite ladder. The empirical channel map is therefore the contribution; the theory is the frame — though §5.1 shows the frame predicts more of the map than we initially expected.

3. Methods

Dataset generation (one command each, reproducible):

```
cargo run -p fenophone-foundry --release -- sweep \
  --all-launch --grid 5 --seeds 8 --ticks 8192 --out research/data/sweep-v1.csv
cargo run -p fenophone-foundry --release -- sweep \
  --all-launch --grid 5 --seeds 32 --ticks 8192 --out research/data/sweep-v2.csv
cargo run -p fenophone-foundry --release -- sweep \
  --all-launch --grid 1 --seeds 32 --ticks 8192 --out research/data/sweep-v2-gatepoint.csv
cargo run -p fenophone-foundry --release -- sweep \
  --all-launch --grid 5 --seeds 32 --ticks 8192 --out research/data/sweep-v3.csv
cargo run -p fenophone-foundry --release -- nulls \
  --all-launch --intensity <x> --spread 0.5 \
  --out research/data/nulls-e2-spread500-int<xxx>-v1.csv   # x ∈ {0, 0.25, 0.5, 0.75, 1}
```

(All outputs are committed under `research/data/`; the v1 dataset took ~35 s total on one core, v2 ≈ 4× that, and each per-cell nulls run ≈ 11 minutes — reruns and larger seed panels are cheap. v2's 32-seed panel is the working dataset for the transport results; v1 is retained for provenance; v3 adds the latent-lane width columns at the current pack tips, with the cross-version equality check reported in the header note above.)

- Grid: Spread × Intensity ∈ $\{0, .25, .5, .75, 1\}^2$, other positions pinned at the gate's default operating point (§5.2 positions); decorrelated seed panel (8 in v1, 32 in v2/v3); 8192 ticks.
- Estimators: the published protocol (DFA-1, bar-template removal on emitted series only; MFDFA q ∈ {−2, 2}; scale ladders 16–1024 onset / 16–256 velocity), anchored at white noise α ≈ 0.5 and random walk α ≈ 1.5 — same code lineage as the CI gate.
- Series per run: latent `dynamics`, `density`, `tension`, `timbre` lanes (no template removal — the latent process has no authored periodicity); emitted per-tick onset counts and mean velocities (template removed).
- Columns: `alpha_lane_{dynamics,density,tension,timbre}`, `alpha_onset`, `alpha_velocity`, `width_onset` (v3 adds `mfdfa_width_lane_*`), `n_onsets`, degeneracy flags; per-cell medians in the companion `*.medians.csv`.
- Mechanism model and grid maps: `research/analysis/quantization-mechanism/` — a stylized image of stages 1 and 3 (the engine's own lane map composed with a monotone integer staircase and template fill), driven by exact fractional Gaussian noise (Hosking) and by a k-layer toggle-cascade net state, measured with a verbatim copy of the shared Python protocol mirror whose anchors (fGn recovery, white/shuffle collapse, a closed-form indicator check) re-run in the package's tests.

4. Results

4.1 The channel map at the default operating point

Medians over the 8-seed panel at the gate operating point (Intensity 0.8, Spread 0.65), 8192 ticks; source `research/data/sweep-v1-gatepoint.medians.csv`:

series (DFA α, median)	beat-driven	warm-ambient	cinematic
tension lane (slow band)	1.015	0.926	0.959
timbre lane (smoothed whole)	1.004	0.961	0.991
density lane (mid band)	0.728	0.657	0.684
dynamics lane (fast band)	0.549	0.498	0.532
emitted onsets	**0.768**	**0.695**	**0.770**
emitted velocity	0.701	0.702	0.676
MFDFA width (onsets)	0.126	0.054	0.108

Three observations shape the rest of the paper:

1. **Cross-validation of the apparatus.** The foundry's α(onset) medians reproduce the CI gate's documented landscape (0.768 / 0.695 / 0.770) to three decimals from an independently compiled copy of the estimator — the measurement chain is sound. The 32-seed panel (sweep-v2) reproduces the map within ≈ ±0.012 in α and ≈ ±0.004 in width (0.779 / 0.686 / 0.758 α(onset); 0.126 / 0.058 / 0.106 width): the 8-seed medians were not a small-panel accident. (Two of these 32-seed medians — beat-driven 0.779 and ambient 0.686 — fall just outside the 8-seed descriptive band `0.695–0.770` reported in papers 01–03 and 05, by ≈ ±0.01; this is the ±0.012 seed-panel scatter above, not a regime change, and both remain well inside the CI gate interval `[0.62, 1.10]` that actually governs releases. Where the descriptive band is quoted as a fixed target across the program, read it as `0.70 ± 0.04` at this panel size.)
2. **The onset stream inherits the density lane's scaling, not the whole product's.** At the gate point, α(onset) sits ≈ 0.04–0.09 above the mid-band density lane (its routed driver in the default matrix) and far below the slow/smoothed lanes (≈ 0.93–1.02). The latent lanes span nearly white (fast band ≈ 0.5) to random-walk-adjacent (slow band ≈ 1.0); the channel hands the note stream the scaling of whichever band the routing feeds it, slightly raised. The grid regression (§4.2) qualifies this: routing decides *which* latent scaling the stream inherits, while the channel's contrast decides *how much* of it registers.
3. **Null verdicts at this point** (`research/data/nulls-v1.summary.csv`): every pack's measured α exceeds its shuffled null's p95 on 8/8 seeds (shuffles

collapse to ≈ 0.50), and the width row is a surrogate-controlled multifractality claim for beat-driven (0.126 vs. IAAFT p95 ≈ 0.080) and cinematic (0.108 vs. ≈ 0.069) but **not** for warm-ambient (0.054 vs. ≈ 0.092, 0/8 seeds) — at this length the ambient stream is consistent with a linear long-memory process with these marginals. Any channel story about *width* must treat the ambient column as a finite-size reading, not a nonlinearity measurement — the per-cell extension of this comparison is §4.3's.

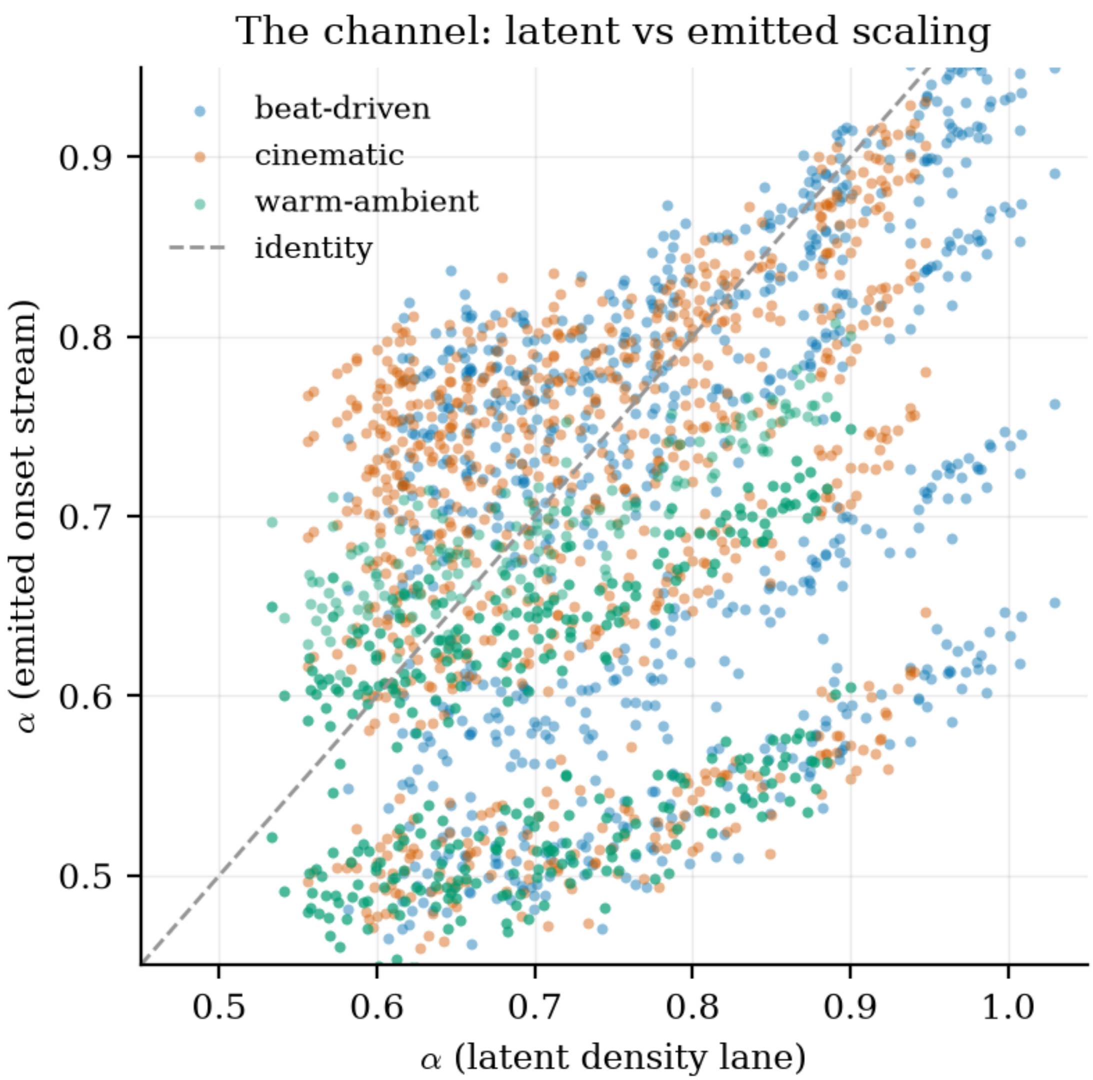

Figure 1. α(emitted onset) vs. α(latent density lane) across the full Spread × Intensity grid × seed panel (`sweep-v2.csv`), with the identity line for reference: the emitted stream inherits the routed mid-band lane's scaling, raised by the channel.

4.2 Exponent transport across the control grid

Results from the 32-seed grid (`sweep-v2.csv`; estimator caveats for this regime are pinned by the known-Hurst anchors — dense fGn recovers H to ≤ 0.01, and the sparse binary channel reads latent persistence *low* by ≈ 0.07, never high, so none of the effects below can be estimator inflation):

- **Regression** (fixed effects; 2400 runs): `alpha_onset ~ alpha_lane_density + alpha_lane_dynamics + scaffold (+ seed)`. Coefficients: **density lane +0.32 ± 0.06**, dynamics lane +0.06 ± 0.05 (indistinguishable from zero), ambient scaffold −0.082 ± 0.006, cinematic +0.005 ± 0.005; seed fixed effects add nothing (R^2 0.300 → 0.302), i.e. seed variance is idiosyncratic, not systematic. (The analysis package re-derives this independently: +0.33 ± 0.05, consistent.) The routed driver (density, mid band) is the only latent lane whose scaling transports — consistent with §4.1's "routed, not global" — **but the model's R^2 is only ≈ 0.30**: most of the emitted-α variance across the grid is *not* explained by latent scaling at all. It is explained by the channel:
- **The Intensity axis.** The naive first-order story (Intensity sets contrast/width, not correlation) is falsified at 32 seeds exactly as it was at 8: at Spread 0.5, emitted α rises monotonically 0.528 → 0.616 → 0.732 → 0.791 → 0.824 on `launch-beat-driven` (0.510 → 0.643 → 0.716 → 0.769 → 0.794 cinematic; 0.509 → 0.509 → 0.638 → 0.638 → 0.684 ambient) **while the latent lane α's are constant along the axis to three decimals** (the density lane sits at 0.775 in every Intensity cell of that slice). Monotone non-decreasing within 0.005 on **all 15 scaffold × Spread rows**. Higher multiplier contrast makes the *same* latent switching register as stronger long-range correlation after thresholding. This is a core result of the paper rather than a robustness check, and §5.1 supplies the mechanism: a level-crossing transfer model in which contrast sets the crossover scale between the exponent-carrying component and the quantization-noise floor, whose no-free-parameter predictor reproduces the climb.
- **Channel dead zones.** The ambient plateaus are exact: at Spread 0.5, Intensity 0 and 0.25 emit **bit-identical onset streams** (per-seed `n_onsets` and α agree to every decimal; likewise 0.5 and 0.75). At low contrast the lane moves too little for the pack's `onsets(density)` quantization to change any bar's onset count — the channel is a step function with dead zones, which any inverse map (corpus fitting) must respect. The dial is honest (non-decreasing) but not injective at the emitted-stream level. §4.4 maps the dead zones exhaustively.

- **The Spread axis: both sides move, and the decomposition closes.** Pooled medians (the median of the pooled seed × Intensity samples at each Spread, $n = 160$) fall monotonically on every scaffold (beat-driven 0.841 → 0.700, cinematic 0.809 → 0.691, ambient 0.693 → 0.601 across Spread $\in \{0, \ldots, 1\}$) — the gate ordering writ large. Unlike Intensity, Spread moves the latent side too, and *faster*: the density lane's pooled α falls at −0.328 / −0.271 / −0.296 per unit Spread (beat-driven / ambient / cinematic) against the emitted stream's −0.148 / −0.088 / −0.113 (`maps_spread_slopes_v1.csv`). Adding the dial positions to the transport regression separates the two routes (density lane +0.62 ± 0.03, Intensity +0.261 ± 0.002, Spread +0.077 ± 0.008, R^2 0.901): the additive decomposition `dα_onset/dS ≈ β_lane·(dα_lane/dS) + β_S` gives lane-mediated terms of −0.205 / −0.169 / −0.185 plus a small **positive** direct term (+0.077), summing to −0.128 / −0.092 / −0.107 — reproducing the measured slopes per scaffold. Reading: Spread acts on the emitted exponent almost entirely **through** the routed lane (widening the span pushes the mid-band rates — the corners of the density lane — higher, decorrelating it), and the channel *damps* the latent shift, with the residual direct effect slightly offsetting it (conditioned on the lane's α, faster mid-band rates mean more crossings per unit time at fixed latent persistence). Intensity's coefficient, by contrast, is pure channel — the contrast effect above, not transport.

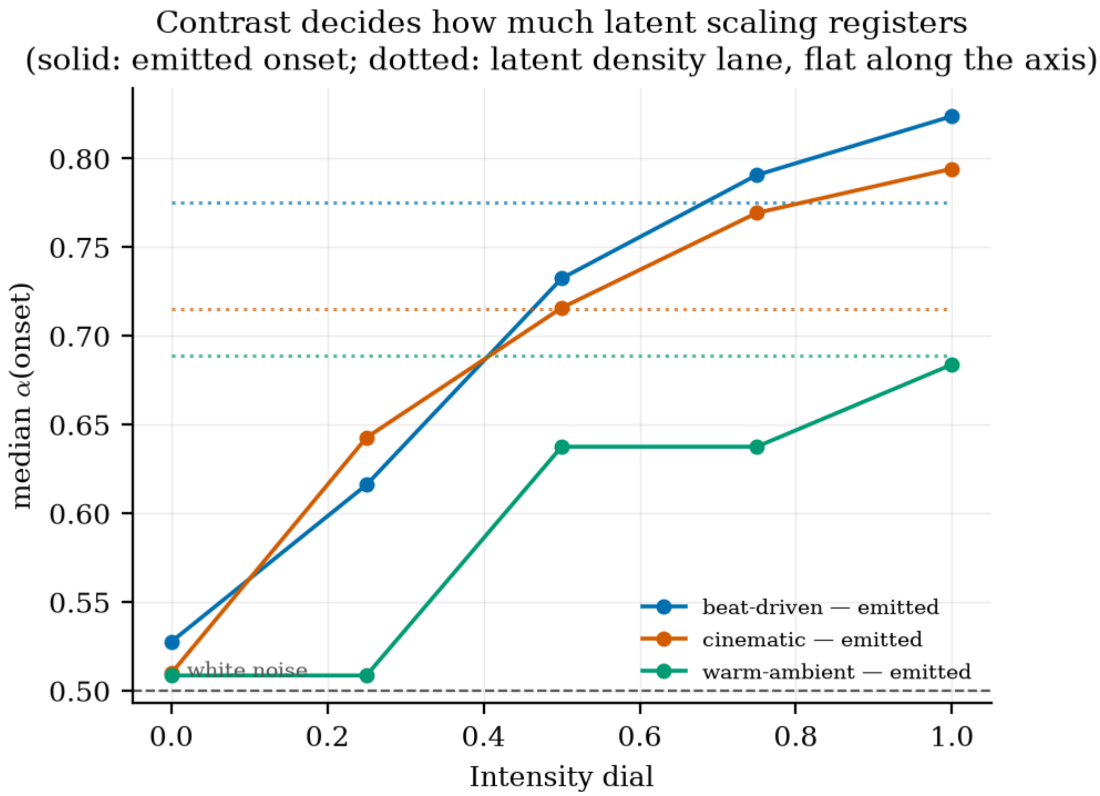

Figure 2. The Intensity axis at Spread 0.5 (`sweep-v2.medians.csv`): median emitted α(onset) climbs monotonically on every scaffold (solid) while the routed density lane's α is flat along the axis (dotted) — contrast decides how much of the latent scaling registers.

4.3 Width through the channel

- `width_onset` vs Intensity (the Intensity→width map — doubles as the E2 stimulus map). The 8-seed slice looked non-monotone (0.097, 0.124, 0.059, 0.131, 0.146 across Spread 0.5) and was provisionally blamed on panel size; **the 32-seed rerun keeps the same shape** (0.093, 0.119, 0.053, 0.124, 0.142), so the mid-Intensity dip is reproducible structure, not seed noise. The same dip appears on the other scaffolds (ambient 0.089 → 0.017 → 0.044; cinematic dips at low-mid Intensity).
- **The dip lives in the negative-q branch (diagnosed 2026-07-19, `foundry spectrum`).** Running the full generalized-Hurst ladder at the Spread 0.5 Intensity slice (8-seed panel; derivations in `research/analysis/spectrum-ensemble-v1/`) reproduces the dip at curve level — beat-driven two-point width `h(−2) − h(2)` = 0.090, 0.123, **0.056**, 0.125, 0.147 across Intensity {0, .25, .5, .75, 1} — and resolves *which moments* collapse. Decomposing the two-point width into a negative-q half `h(−2) − h(0)` and a positive-q half `h(0) − h(2)` (and the full ±4 ladder into `h(−4) − h(0)` and `h(0) − h(4)`), at the dip cell the **negative-q (small-fluctuation)** span falls to 0.045 from ≈ 0.11–0.12 at the flanking cells (a drop of ≈ 0.07), while the

positive-q span falls to 0.046 from ≈ 0.075–0.08 (a drop of ≈ 0.03): ≈ ⅔–70 % of the width loss is on the negative-q branch — the small-fluctuation Hurst exponents flatten toward the bulk while the large-fluctuation side barely moves. This is the branch that interacts with the negative-q sparse-window exclusion, so a definitive attribution (genuine narrowing vs. an exclusion artifact) needs the per-window fluctuation distributions — a remaining follow-up.

- **The dip cell fails both of its own nulls (the per-cell run).** The `foundry nulls` protocol, previously run only at the gate point, is now extended to every cell of the Spread-0.5 Intensity slice (`research/data/nulls-e2-spread500-int{000,250,500,750,1000}-v1.csv`; 8-seed panels, 99 IAAFT + 99 shuffle surrogates per seed; digest in `maps_e2_null_excess_v1.csv`). The earlier spectrum diagnosis showed the dip cell's width (0.056) sitting **at its shuffle floor** (0.058); the per-cell IAAFT comparison sharpens the verdict — at the beat-driven dip cell the measured width (median 0.059) beats its matched IAAFT p95 (0.089) on **0/8 seeds** (excess −0.030), while the flanking high-Intensity cells beat theirs on 8/8 with excesses of +0.049 and +0.066. Cinematic shows the same structure at its own dip cells (I = 0.25: width 0.020 vs. p95 0.102, 0/8; I = 0.5: 0.054 vs. 0.086, 0/8; I ≥ 0.75: 7/8 with excesses ≈ +0.05). So the mid-Intensity "dip" is, in full: the two-point width collapsing into the finite-size floor, not a nonlinearity change — at the dip cells the stream's multifractal spread is not separable from a linear (or even order-destroyed) surrogate of the same values, exactly the excess-over-null caveat the E2 stimulus design required. Every E2 candidate cell now carries its width as an excess over its own nulls rather than a raw number. (Warm-ambient, measured here at its retuned `1.3.0`: the mid-grid cells beat their IAAFT p95 on 8/8 seeds — the retune's documented lift, off the gate point as well as on it — while Intensity 0 is degenerate at this version; §4.4.)
- **Width compression: the hypothesis is refuted.** The natural expectation — thresholding compresses multifractal width, so the emitted stream should read narrower than its driving lane — is what `sweep-v3`'s paired lane/onset width columns were cut to test, and it does not survive measurement. Per run, the emitted onset width is *comparable to* the routed density lane's width, not systematically below it: median per-run ratios width(onset)/width(lane) of **1.07** (beat-driven), **1.18** (warm-ambient), **0.88** (cinematic), with pack-level medians of 0.108 vs. 0.096, 0.112 vs. 0.095, and 0.083 vs. 0.085 respectively, and per-cell median compression in only 12 / 9 / 15 of 25 grid cells (`maps_width_compression_v1.csv`). The honest reading is a negative: at this length and protocol, the quantization channel neither reliably compresses

nor reliably amplifies the two-point width, and no width-compression correction should be applied when inverting the channel. One clause now travels with the ratio itself: both sides are measured at the same length, the ratio is uninformative when either side sits near its finite-size floor, and the width statistic's measured low power against genuine cascade structure (Fenocosm v22 / Fenosoma x7) makes near-floor the expected case rather than an edge case — the ratio bounds the channel's *width* effect only, not its effect on the nonlinear structure width cannot see. The caveat from §4.1 carries over — the dense-lane and sparse-count width estimators have different finite-size floors (the per-cell nulls above are the context for the latter), so this comparison bounds the *systematic* effect, not the per-cell one.

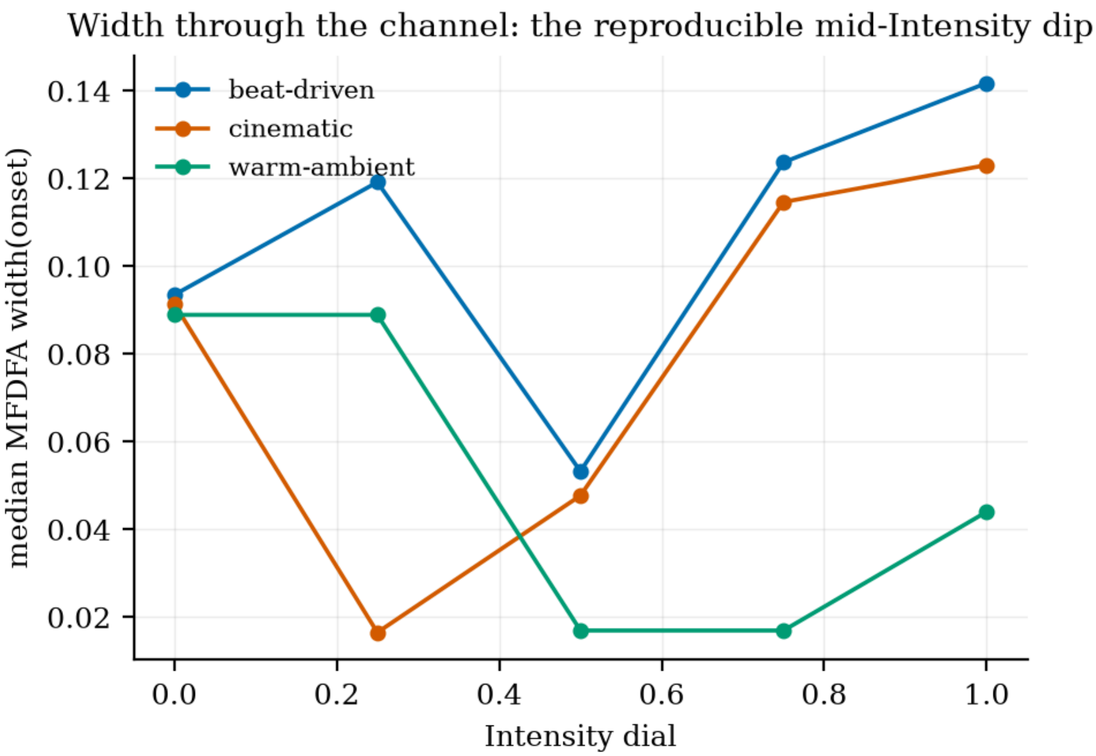

Figure 3. Median MFDFA width(onset) vs. Intensity at Spread 0.5 (`sweep-v2.medians.csv`): the reproducible mid-Intensity dip. The per-cell null comparison (§4.3) shows the dip cells are width-at-the-floor readings, not narrower nonlinear structure.

4.4 Degeneracy boundary

Where in the grid does the stream go silent, wall, or template-only — the gate's non-degeneracy failures, i.e. the usable region of the control space? The 32-seed grid answers exhaustively (`maps_degeneracy_v1.csv`):

- **The primary control plane is entirely usable.** Across all 2400 sweep-v2 runs — every launch scaffold × 25 cells × 32 seeds — there are **zero** silent and **zero** wall rows. The gate's hard failure modes live outside the Spread × Intensity plane (they require the pinned Change-Rate/Layers axes or hostile pack data), which is itself a design fact about the shipped packs: the two primary performance dials cannot drive a launch scaffold off the usable region.

- **The boundary structure on the plane is the exact dead zones.** 10 of the 60 adjacent-Intensity cell pairs are bit-identical on all 32 seeds (per-seed `n_onsets` and α equal to every decimal) — and all ten are warm-ambient's: `I = 0 ≡ 0.25` and `I = 0.5 ≡ 0.75` at every Spread value. The least intermittent pack has the shallowest `onsets(density)` staircase, so quarter-dial Intensity steps fall inside single quantizer steps across the whole Spread range; beat-driven and cinematic have none.
- **The boundary is pack-version-dependent.** At the current tips (`sweep-v3`), the warm-ambient `1.3.0` retune pushes the Intensity-0 row into `template_only` degeneracy — an exactly bar-periodic stream (constant per-bar count; α undefined after bar adjustment) on all 32 seeds at every Spread, a row that was measurable at `1.0.0`. The retune's convexity remap bought its documented mid-grid width lift at the price of the low-Intensity corner's variation, and the `template_only` flag catches exactly this. The degeneracy map is therefore a per-version artifact to re-cut with the packs, not a fixed property of the engine.
- **The mechanism is the staircase.** The stylized channel model (§5.1) reproduces both behaviors with no engine machinery: when the lane's dynamic range at a given contrast is smaller than one staircase step, the emitted count is constant (template-only), and quarter-dial steps that stay within one staircase step are exact dead zones. Degeneracy and dead zones are the low-contrast limit of the same quantizer that §5.1's transfer model describes everywhere else.

5. Discussion

5.1 The mechanism: a level-crossing transfer model under varying contrast

The Intensity-axis result — emitted α climbing 0.51 → 0.83 while every latent lane's α is constant to three decimals — demands a mechanism, and the channel's own structure supplies one. Strip the channel to its load-bearing pieces (the engine's lane map `clamp(1/2 + c·s, 0, 1)` with contrast `c = log2(m_hi)/(4n)` — the Intensity knob's image — composed with the monotone integer `onsets(density)` staircase and the template fill of §2), drive it with a persistent latent process, and measure the output with the published estimator: this is the analysis package `research/analysis/quantization-`

`mechanism/` (`out/RESULTS-mechanism-v1.md`), run on three rungs — exact fractional Gaussian noise at the density lane's measured exponent (H = 0.775, the theory rung), a k = 5 toggle-cascade net state (the model rung), and a per-tick engine-shaped rung that adds template fill and bar adjustment under the full 16–1024 protocol ladder.

The theory rung reproduces the phenomenon quantitatively: emitted bar-count α climbs 0.586 → 0.828 as contrast sweeps 0.02 → 0.18 while the latent α is constant at 0.827, and the engine-shaped rung climbs 0.506 → 0.771 — the measured Intensity axis's shape, in a model with no routing, no harmony, and no engine. The mechanism it isolates is a level-crossing transfer: because the composed lane-map-plus-staircase is a **monotone** step function, its Hermite rank is 1, so level-crossing theory (Taqqu 1975; Kratz 2006; Vahabi et al. 2011) predicts the latent exponent is preserved **asymptotically at every contrast** — contrast neither creates nor destroys long memory. What contrast does set is the *amplitude* of the exponent-carrying (rank-1) component relative to the quantization-noise floor: writing `g` for the composed quantizer and `c1 = E[g(X)·X]` for its first Hermite coefficient, the emitted fluctuation function is approximately the mixture `F²(s) ≈ c1²·F²_fGn(s) + (Var g − c1²)·F²_white(s)`, whose crossover scale falls as contrast raises `c1²/Var g` (measured 0.02 → 0.97 across the sweep) — and a fixed measurement ladder reads the mixture as an effective α between 0.5 and the latent value. The predictor built from this decomposition has **no free parameters** — `c1` and `Var g` are closed-form sums of normal density/tail values over the staircase's jump points, and the pure-fGn/pure-white curves are measured on the same seeds — and it reproduces the measured climb cell by cell: median |Δα| = 0.009 over the non-degenerate theory-rung cells (the maximum, 0.209, sits at the near-dead-zone contrast 0.02, where the count varies through a handful of rare excursions and single-seed DFA is at its noisiest). Robustness holds at H = 0.65 and H = 0.85, and the MSM rung shows the same climb measured-only (0.680 → 0.768 across the engine's own `m_hi ∈ [1.1, 4]` contrast range, saturating early because the net state's ±k span is wide).

Two further measured details close loops opened elsewhere in the paper. The model's low-contrast limit *is* §4.4's degeneracy: below one staircase step the count is constant (the model rung's `m_hi = 1.1` cell is degenerate, as is the theory rung's c = 0.02 on some seeds). And at the top of the contrast sweep the theory rung's measured α eases back (0.828 → 0.802) as the **clamp** begins to saturate — the rank-1 share falls from 0.97 to 0.93 because clipping flattens the largest excursions — a first quantitative answer to "what does the clamp cost": near the dial's top, saturation trades a few hundredths of effective α back. The frame's limits are stated as measured:

the predictor is exact only for a Gaussian latent (the MSM rung is bounded and discrete, reported measured-only), and the engine-shaped rung's climb is steeper than the bar-count rung's because template fill and bar adjustment add their own sub-bar structure — the stylized model brackets the engine's channel, it does not replace it.

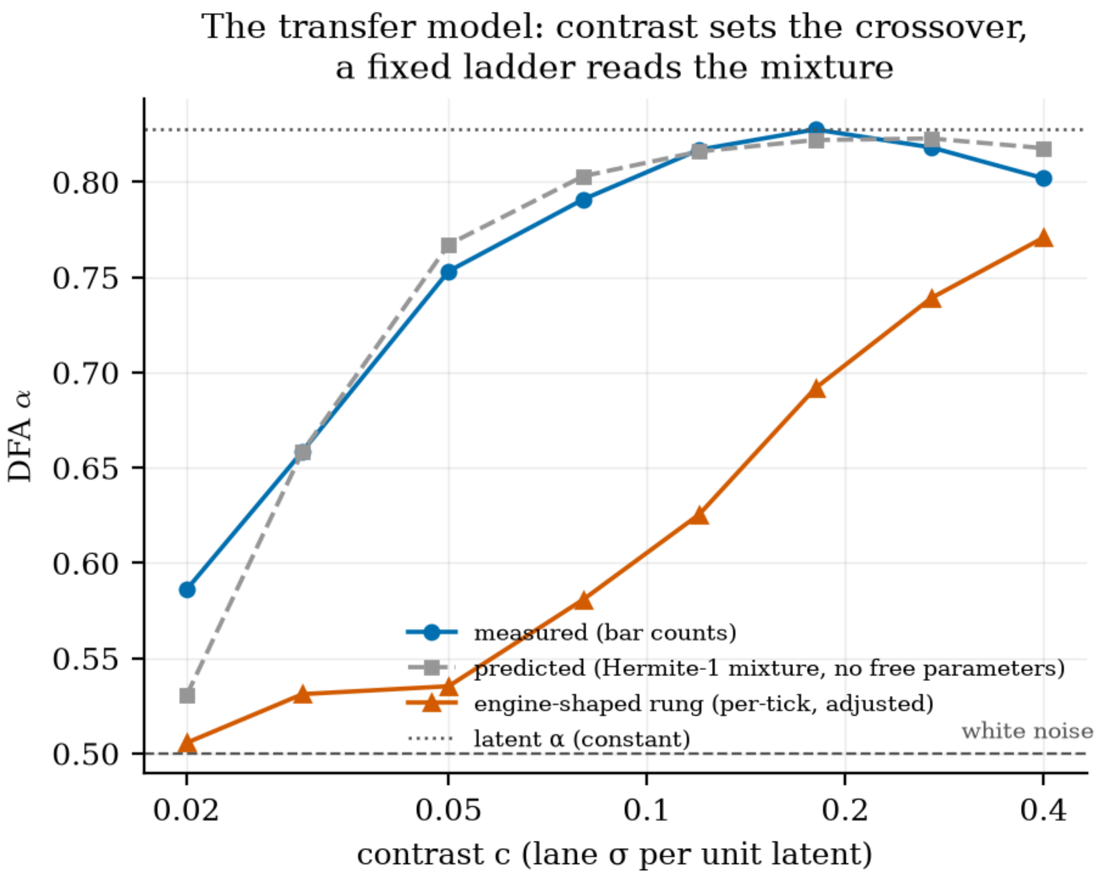

Figure 4. The transfer model (`quantization-mechanism/out/mechanism_v1.csv`, theory rung, $H = 0.775$): measured emitted α (circles) vs. the no-free-parameter Hermite-1 mixture prediction (squares) across contrast, with the constant latent α (dotted) and the engine-shaped per-tick rung (triangles). Contrast sets the crossover scale; a fixed ladder reads the mixture.

5.2 The design rule

- **To hit a target emitted α, use the fitted surface.** Quadratic fits of the sweep-v2 cell medians, $\alpha(I, S) = b0 + b\_I \cdot I + b\_S \cdot S + b\_IS \cdot I \cdot S + b\_I^2 \cdot I^2 + b\_S^2 \cdot S^2$ (`maps_inverse_design_v1.csv`), are accurate to well under the seed-panel scatter: beat-driven $0.618 + 0.536 \cdot I - 0.268 \cdot S - 0.043 \cdot I \cdot S - 0.201 \cdot I^2 + 0.142 \cdot S^2$ (R^2 0.991, RMSE 0.012), cinematic $0.587 + 0.566 \cdot I - 0.218 \cdot S - 0.048 \cdot I \cdot S - 0.262 \cdot I^2 + 0.138 \cdot S^2$ (R^2 0.996, RMSE 0.007), warm-ambient $0.543 + 0.247 \cdot I - 0.145 \cdot S - 0.031 \cdot I \cdot S - 0.037 \cdot I^2 + 0.081 \cdot S^2$ (R^2 0.887, RMSE 0.027 — the staircase's dead zones are the residual). Given a target α, its iso-α contour in the (Intensity, Spread) plane is the design rule — quoted to grid resolution, because §4.4's dead zones make the map a staircase at the emitted-stream level (state targets to ±half a grid step, and treat warm-ambient's low-Intensity half as quantized in quarter-dial steps).

- **Authored latent shifts transport at a fraction.** A pack author moving the routed lane's scaling should expect the emitted α to move by the transport coefficient — quoted as an observational range from grid variation (+0.32 pooled; +0.62 conditioned on the dial position), not as an authored-pack intervention, which remains the direct experiment to run. Combined with §5.1, the design intuition is: **author the latent scaling for the target's neighborhood, then let Intensity choose how much of it registers** — remembering that the top of the dial buys width (Figure 3) at a small cost in α (the clamp, §5.1), and the bottom of the dial is a staircase.

5.3 Implications for the inverse problem

Corpus fitting must model the channel, not assume identity, and the first external-corpus run (2026-07-12, `research/analysis/msm-corpus-fit/out/RESULTS-corpus-v1.md`) demonstrates the stakes in both directions. (i) The engine's own emission channel presents **near-Poisson marginals** that mask the latent long memory: exact-ML MSM-Cox fits of engine streams collapse to a flat cascade (`m_hi^κ ≈ 1.0`) even though the streams measure α ≈ 0.80. (ii) Real solo-piano performances (1,260 MAESTRO pieces) present an **overdispersed** count channel (Fano ≈ 1.55), and the same estimator recovers a non-flat cascade (`m_hi^κ` median 1.619) whose resynthesis retains the persistence (α 0.691 vs. 0.717). (The null calibration — `analysis/msm-corpus-fit/out/RESULTS-nullfloor-v2.md`, 2026-08-30, superseding the 2026-08-26 v1 — shows `m_hi^κ` carries a large **zero-correlation floor** set by that overdispersion. The **real-data** decorrelation arm is **GiantMIDI**, not MAESTRO: on 56 audio-transcribed GiantMIDI performances, at a different marginal from MAESTRO's (Fano median 1.797, mean 0.379/tick against 1.552 and 0.405), the data's own shuffle and an iid negative-binomial null matched to each piece's own (mean, Fano) both fit ≈ 1.76–1.81 against real 1.80. **MAESTRO's own shuffle / Fano-matched-negbin null has not been run** — it is queued in `research/TODO.md`, and no earlier statement of this paragraph should be read as reporting it. At MAESTRO's own operating point what is measured is the **synthetic** iid-negbin floor curve: mean- and Fano-matched to the corpus marginal (mean 0.4046/tick, Fano 1.552, k = 5) it fits `m_hi^κ` ≈ **1.605**, so the sealed 1.619 sits **at** its own zero-correlation floor. So `m_hi^κ = 1.619` is a marginal statistic, and the temporal-structure claim rests on the DFA α 0.717 vs. shuffled 0.500 dissociation, not on the bare multiplier.) (iii) The drum idiom (198 Groove takes) shows the channel can carry the cascade on a **different observable entirely**: beat-grid onset counts are near-white ($\alpha \approx 0.475$, flat fitted cascade) while the velocity channel fits non-flat (`vel_m_hi ≈ 1.36`; piano ≈ 1.52). The channel is not a nuisance correction — it decides which observable inherits the latent scaling, exactly this paper's thesis, now visible in real corpora and not only in the sweep grid. This paper adds the quantitative half the inverse problem needs: the transfer model says *how much* latent persistence a count channel of given contrast lets through (§5.1), the dead-zone map says where the inverse is set-valued rather than point-valued (§4.4), and the width result says no width-compression correction is warranted (§4.3).

5.4 Limitations

- **Where theory and measurement part ways.** The Hermite-rank-1 frame predicts the contrast axis and the asymptotic exponent preservation, but it is a Gaussian-latent, single-stage idealization: the real lanes are bounded non-Gaussian switching sums, the real channel stacks routing and harmony gating on top of the staircase, and the engine-shaped rung shows those extra stages steepen the climb. The stylized model brackets the channel; the sweep datasets remain the measurement of record.
- DFA on count series; finite-size MFDFA bias at 8192 ticks (the per-cell nulls of §4.3 are the guard); the negative-q attribution of the dip awaits the per-window fluctuation distributions.
- Two control axes swept, two pinned (Change Rate and Layers at the gate point); the degeneracy result is bounded to the swept plane.
- The width-compression comparison pairs estimators with different finite-size floors (dense lane vs. sparse counts); it bounds the systematic effect only.
- Pack-version provenance: sweep-v2 measures `1.0.0` packs; the per-cell nulls and sweep-v3 measure the current tips. For beat-driven and cinematic the emitted measurements are string-identical across versions (header note); warm-ambient's two versions are reported as distinct streams, and its `1.3.0` degeneracy row (§4.4) is a fact about the retune, not about `1.0.0`.

6. Reproducibility

Everything is a `cargo` invocation on pinned seeds plus two seeded offline Python drivers; the CSV schema is documented in `engine/crates/foundry/README.md`, the analysis package's protocol in `research/analysis/quantization-mechanism/README.md`, and every command appears in §3. The estimators are the CI gate's own lineage, anchored in unit tests on both sides of the language boundary (known-Hurst recovery, sparse-channel attenuation, template-removal no-op in `spectral.rs`; fGn recovery, white/shuffle collapse, and a closed-form staircase check in the package's pytest suite). Because the event layer is bit-identical across platforms, every number in §4 is a deterministic function of the repository at the recorded scaffold versions.

References

Verified during the pre-submission citation pass; the machine-readable bibliography with DOIs is `references.bib` (shared across papers 01–04).

1. Peng et al. (1994), *Phys. Rev. E* 49:1685–1689 — detrended fluctuation analysis (DFA).
2. Kantelhardt et al. (2002), *Physica A* 316:87–114 — multifractal DFA (MFDFA).
3. Calvet & Fisher (2001, *J. Econometrics* 105:27; 2004, *J. Financial Econometrics* 2:49; 2008, *Multifractal Volatility*, Academic Press) — the Markov-switching multifractal.
4. Taqqu (1975), *Z. Wahrsch. Verw. Gebiete* 31:287–302; Beran (1994), *Statistics for Long-Memory Processes*, Chapman & Hall — transformations of long-memory processes.
5. Kratz (2006), *Probability Surveys* 3:230–288; Vahabi, Jafari & Movahed (2011), *J. Stat. Mech.* P11021 — level crossings of stationary Gaussian / fractional-Gaussian processes (the channel's exponent-preservation anchor).
6. Voss & Clarke (1975, *Nature* 258:317; 1978, *JASA* 63:258) — 1/f fluctuations in music.
7. Bouchaud, Potters & Meyer (2000), *Eur. Phys. J. B* 13:595–599 — apparent multifractality (the null-comparison rationale).
8. Schreiber & Schmitz (1996), *Phys. Rev. Lett.* 77:635–638 — IAAFT surrogates.
9. Theiler et al. (1992), *Physica D* 58:77–94 — the surrogate-data method.
10. Hosking (1984), *Water Resour. Res.* 20:1898–1908 — exact fractional-Gaussian-noise simulation (the known-Hurst estimator anchors and the mechanism model's theory rung).
11. Hu et al. (2001), *Phys. Rev. E* 64:011114; Bryce & Sprague (2012), *Sci. Rep.* 2:315 — DFA caveats the anchors quantify.

Companion records of the Feno program (DOIs published 2026-08-31)

*Every record below is deposited on Zenodo — staged as a private draft with a **reserved** DOI 2026-08-30, **published 2026-08-31**: every DOI below is live and resolves. The deposit map is `research/zenodo/records.json`, and the source repositories are private, so each record carries its own reproduction materials. Cross-citations follow that map's `related_identifiers`.*

- Atlas, E. T. 2026. "The Fenophone: A Musical Instrument Built on the Markov-Switching Multifractal, with CI-Verified Spectral Invariants" (Fenophone research series, paper 01). Zenodo DOI 10.5281/zenodo.22180423 (published 2026-08-31).
- Atlas, E. T. 2026. "The Markov-Switching Multifractal as a Generative Model of Musical Intensity: An Instrument, Measurements, and a Research Programme" (Fenophone research series, paper 03). Zenodo DOI 10.5281/zenodo.22180429 (published 2026-08-31).

- Atlas, E. T. 2026. "Does the Markov-Switching Multifractal Describe Real Music? A Fitting Pipeline and a Statistical Turing Test" (companion analysis paper, `analysis/msm-corpus-fit/PAPER.md`). Zenodo DOI 10.5281/zenodo.22180435 (published 2026-08-31).