# An exact-likelihood test of the cascade hypothesis in heartbeat dynamics

**MSM-modulated durations lose to a matched linear roster on every one of 83 R–R records by exact held-out likelihood; the continuous-scale member is rejected on all of them; the heartbeat's log-volatility is rough, not cascading.**

Evan Tabak Atlas (ORCID 0009-0007-7374-2338)

*Draft — result paper (the refutation paper,* `SPECS.md` *§7). Target venue: arXiv first (physics.data-an / q-bio.QM); journal shortlist is an owner call (see* `OWNER-QUEUE.md` *item 3). Every quantitative claim traces to a named, immutable results record under* `out/`*: the founding tournament* `out/v0/` *(Q1, signatures/atlas/estimability/tournament) and the continuous-scale arm* `out/v2/` *(the family-level extension). The apparatus gate (*`./verify.sh` *→* `X0 PASS`*) runs from a fresh clone with only* `requirements.txt`*; the duration-channel closed forms are asserted against simulation in* `tests/test_durations.py` *(module* `fenosoma/durations.py`*). Figure:* `out/v0/fig_soma.png` *(tournament panel), regenerable by* `python make_figures.py` *from the sealed records.*

*Priority note: this finding has no DOI and no priority date. The alternative it confirms — a tempered-Lévy (linear) description of heartbeat increments — was proposed by Kiyono et al. (2006) and left untested for the lack of a model-comparison apparatus; the dedicated robustness sweep (editing-policy sensitivity as a first-class record, IG-vs-Gamma emission, cohort/record-length robustness, and the k-selection BIC path) is prompt P-2.1, whose record* `v5` *is not yet sealed — §8 reports what robustness evidence the sealed records already carry and marks the P-2.1 interval as forthcoming. arXiv posting is an owner action (endorsement).*

*Correction, 2026-08-31: the priority note above is superseded on its first clause — this paper's Zenodo deposit was published on 2026-08-31 and its DOI, doi:10.5281/zenodo.22180527, now resolves; the publication date is the priority date (*`PROVENANCE.md` *§5). The rest of the note stands: record* `v5` *is not yet sealed, and arXiv posting remains an owner action.*

---

## Abstract

For twenty-five years the multifractality of the human heartbeat has been read as the signature of a *multiplicative cascade* — a hierarchy of regime timescales whose products modulate the R–R interval, by analogy to the energy cascade of fluid turbulence (Lin &

Hughson 2001; Kiyono et al. 2004, 2005). The reading has never been closed, because its own authors found the increment distribution fit *equally well* by a tempered-Lévy (linear) alternative (Kiyono et al. 2006) and could not confirm the multiplicative mechanism (Kiyono & Bekki 2011). The descriptive tools of the field — multifractal detrended fluctuation analysis (MFDFA) widths, increment probability densities — are structurally unable to separate a cascade from a well-tuned linear process. **Exact-likelihood model comparison can, and this paper runs it.**

We treat the R–R stream as a renewal-type point process observed through its inter-event durations, with the durations conditionally inverse-Gaussian and their conditional mean modulated by a discrete Markov-switching multifractal (MSM) cascade of depth `k`; the `2^k`-state regime chain admits an *exact* forward-filter likelihood. On the two canonical PhysioNet cohorts — the Normal Sinus Rhythm (`nsr2db`, 54 subjects) and Congestive Heart Failure (`chf2db`, 29 subjects) 24-hour RR databases — we fit the cascade against a fixed, matched roster (the tempered-Lévy-type linear model in the spirit of Kiyono 2006; the Barbieri–Brown history-dependent inverse-Gaussian point process, reimplemented from the published equations; flat `K`-state hidden Markov models, `K ∈ {2,4,8}`; an Ornstein–Uhlenbeck-modulated duration model; and an i.i.d. baseline) under one protocol — identical held-out splits, complexity-penalised by BIC and scored held-out, parametric-bootstrapped on the headline verdict.

**The cascade loses on every record.** Held-out, the best cascade is worse than the best matched linear competitor by a median of `−27,246` nats/record on the healthy arm (`0/54` records won; 95% CI `[−32,552, −24,810]`; `−859` nats per 1,000 events) and `−27,492` nats on the CHF arm (`0/29`). The best held-out model is the tempered-Lévy-type linear process on `79/83` records (best BIC on `82/83`); an Ornstein–Uhlenbeck-modulated duration model — the same emission, one *fewer* parameter, no access to lagged durations — beats the depth-4 cascade held-out on `77/83` records, so the loss is not an artifact of history-aware rivals (on the six exceptions the cascade beats that one rival and still loses the tournament; §6). A parametric bootstrap places every observed cascade deficit at `p ≥ 0.54` for any cascade advantage while its power arm recovers cascade wins on cascade-simulated data, and the estimability battery certifies that the cascade was *identifiable* at these lengths — the recovery is clean from ~500 events, so this is not a starved-data artifact. We then remove the discrete scaffolding entirely: the continuous-scale log-correlated cascade (the log-normal MRW / log-S-fBM field) is GoF-rejected on all `83` records, with a fitted Hurst exponent significantly above its cascade limit (`Ĥ = 0.111`, 95% CI `[0.088, 0.131]` on the healthy cohort) — **the heartbeat's log-volatility reads as *rough*, not cascading.** The founding negative is a family-level statement: neither the discrete `2^k` ladder nor the continuous-scale field is the generative story of heartbeat durations, and the grid was not the reason. Every result is reproducible from a fresh clone behind a deterministic apparatus gate. **No statement here is clinical**; each is a model-comparison or measurement statement about public, de-identified datasets.

---

# 1. Introduction

## 1.1 A twenty-five-year dispute

Heartbeat variability is multifractal: the local scaling exponent of the R–R interval series varies with the moment order, and the singularity spectrum has finite width (Ivanov et al. 1999). The dominant generative reading of that width is a *multiplicative cascade* — the same mathematical object turbulence theory uses for the energy dissipation field — in which a small number of regime timescales, geometrically spaced, multiply together to produce a heavy-tailed, long-range-correlated, scale-invariant modulation of the interbeat interval (Lin & Hughson 2001; Kiyono et al. 2004, 2005). The picture is attractive: it explains the lognormal-like amplitude statistics and the broadband `1/f` structure at once.

It has never been closed. The canonical papers themselves report the obstacle: Kiyono et al. (2006) found the heartbeat increment PDF described *equally well* by a tempered-Lévy stable (linear) process, and Kiyono & Bekki (2011) stated plainly that the multiplicative mechanism could not be confirmed by the available statistics. The reason is structural. A multifractal width, an increment PDF, a partition-function slope — these are *descriptive summaries*. A well-tuned linear process with the right marginal and the right spectrum can reproduce all of them. The dispute is not resolvable by a sharper summary statistic; it is resolvable only by a *generative* comparison that penalises complexity and scores on data the models did not see.

## 1.2 What an exact likelihood buys

The Markov-switching multifractal (Calvet & Fisher 2001, 2004, 2008) is the member of the cascade family that carries an *exact* likelihood: a depth-`k` cascade is a hidden Markov chain over `2^k` states, so the forward algorithm evaluates its likelihood exactly, in closed form, per observation. That is the whole leverage of this program. It lets the cascade enter a tournament as a fully generative model — fit by maximum likelihood, penalised by BIC, scored on a held-out tail — against the linear alternatives the field could only compare *descriptively*. The verdict is then a likelihood statement, not a width statement.

This transposes directly the tournament the sibling astrophysics program (Fenocosm) ran on accreting compact objects, where the same MSM cascade was raced against the damped random walk. Here the object is the heartbeat and the `Q1` question is exactly the dispute above: *is heartbeat multifractality actually a multiplicative cascade, decided by exact likelihood rather than width?* Either verdict is a deliverable. The pre-registered kill/continue gate (`SPECS.md` §7) fixed the consequence in advance: **if the cascade cannot beat the matched linear competitor on the healthy arm's held-out score, this becomes the refutation paper** — the program does not spin the loss.

### 1.3 What this paper claims, and what it does not

This paper claims that a *multiplicative cascade in the conditional mean* of the inter-beat duration is not the generative description of the R–R streams tested, by exact likelihood and by a calibrated goodness-of-fit on the continuous-scale relaxation of the same hypothesis. It does **not** claim that heartbeat dynamics contain no multiplicative structure at any level: §5 and §6 name precisely what is measured (a rough, `Ĥ ≈ 0.11` log-volatility) and what is not, and the companion audit paper (Paper B, `papers/02-two-instrument-audit.md`) shows the healthy heartbeat is strongly nonlinear in a way *no* two-point width sees. This paper is about one hypothesis, tested one way, decisively.

### 1.4 Contributions

1. The first *exact-likelihood* model-comparison test of the cascade hypothesis in heartbeat dynamics, on the two canonical 24-hour cohorts, under a pre-registered fairness protocol and kill gate (§3, §6; record `v0`).
2. A pre-tournament *estimability* certificate (§4) showing the cascade is identifiable from ~500 events on this channel, so the negative is not a starved-data artifact — anchors before claims.
3. A physiology-specific *observation-channel atlas* (§5) measuring how beat editing and count-binning bias the recovered cascade amplitude, and the finding that the astrophysical "reads low, never high" law does **not** transfer.
4. The verdict, with its bootstrap and its honest scope (§6): `0/83`, tempered-Lévy best on `79/83`, the apples-to-apples OU control, and what "the cascade loses" does and does not mean.
5. A *family-level* extension (§7; record `v2`): removing the grid entirely, the continuous-scale log-correlated cascade is rejected on `83/83` records and the log-volatility reads rough (`Ĥ ≈ 0.11`) — so the discrete scaffolding was not the reason the cascade lost.

---

## 2. Data

*(Record: `out/v0/`, the founding run; editing policy and edited-beat fractions carried per number per `SPECS.md` §5 F4.)*

The founding tournament runs on two 24-hour RR-interval databases from PhysioNet: the **Normal Sinus Rhythm RR Interval Database** (`nsr2db`, `54` subjects, the healthy arm and the kill-gate cohort) and the **Congestive Heart Failure RR Interval Database** (`chf2db`, `29` subjects, the pathological contrast). Together they are `83` records (`5,739,633` and `3,103,627` R–R intervals respectively).

Physiology's observation channel is beat detection and editing, and it is declared, not assumed (`SPECS.md` §5 F4). We take the committed PhysioNet beat annotations as the primary event source — this program writes no beat detectors — and report every duration number under the **normal-to-normal (NN)** editing view, with the raw-RR view available for the audit (Paper B) and the fraction of edited beats carried per record. The edited-beat fraction is small and cohort-dependent: median `0.0005` on `nsr2db`, `0.0125` on `chf2db` (the implausible-interval fraction, from the declared `0.2–4.0` s plausibility stage, is `≤ 0.0007` in both). No interpolation is applied by default.

Every artifact using these data cites both the source database and PhysioNet itself (Goldberger et al. 2000). These are public, de-identified research datasets; nothing in this paper is a diagnostic, predictive, or individual-level statement.

## 3. The model and the roster

*(Record: `out/v0/x4_tournament.json`; roster and fairness rules per `SPECS.md` §4; protocol constants in `fenosoma/protocol.py`, block `p1`.)*

**The latent object.** A `k`-level ladder of independent two-state multipliers with geometrically spaced switching rates, state space `2^k`, evaluated by an exact recursive forward filter — the vendored MSM theory module, unmodified. Nothing new is invented in the cascade; that is the point.

**Channel B — durations (the new modelling piece).** The R–R stream is a renewal-type point process observed through its inter-event durations. Durations are conditionally **inverse-Gaussian** (the Barbieri–Brown emission, chosen deliberately so the field's smooth-intensity competitor is nested in spirit; a Gamma emission is the sensitivity alternative, §8), with conditional mean modulated by the cascade product. Because the regime chain is discrete, the exact `2^k` filter applies unchanged; the derivation of how regime transitions are handled between unevenly spaced events is documented in the module docstring (`fenosoma/durations.py`) and self-checked by closed-form stationary duration moments in the apparatus gate.

**The roster (fixed before the tournament, `SPECS.md` §4).** Nine models were fit per record:

- **Cascade** — MSM-modulated durations at `k ∈ {2, 4}` (`cascade_k2`, `cascade_k4`).
- **Linear group** — (1) the **tempered-Lévy-type linear model** in the spirit of Kiyono (2006): an AR(4) on log-durations with normal-inverse-Gaussian innovations, scored on durations through the exact Jacobian (`templevy_ar4`); (2) the **Barbieri–Brown** history-dependent inverse-Gaussian point process, reimplemented from the 2005 equations, not vendored (`bb`); (3) an **Ornstein–Uhlenbeck / DRW-modulated** inverse-Gaussian duration model on the same event clock (`ou_ig`); (4) an **i.i.d. inverse-Gaussian** baseline (`iid_ig`).
- **Flat group** — flat `K`-state IG-HMMs, `K ∈ {2, 4, 8}` (`hmm_k2/k4/k8`), the "many states, one layer" alternative to a hierarchy.

**Fairness rules (verbatim, `SPECS.md` §4).** Identical emission error models where applicable; identical held-out splits (a contiguous `30%` tail per record); complexity-penalised — **BIC and held-out score, both reported**; parametric bootstrap on the verdict. The self-excitation (Hawkes) rival named in the roster is a *count-stream* rival and is raced on Channel-A event streams in the methodological bridge that seals record `v6` (P-2.3, `FINDINGS.md` F9); it does not enter the duration tournament.

## 4. Estimability before the tournament

*(Record: `out/v0/x3_estimability.json`, finding F3; anchors before claims, `SPECS.md` §5 F1.)*

A negative tournament is only a statement about the *hypothesis* if the apparatus could have found the cascade had it been there. We certify that before touching the verdict (the x3-before-x4 discipline). Simulating cascades at the founding operating point (`k = 4`, `m_hi = 1.4`, geometric switching, six seeds per cell) and re-estimating by exact-filter maximum likelihood across observation counts from `500` to `100,000` events, on both channels:

| `N` (durations) | 500 | 1,000 | 2,000 | 5,000 | 10,000 | 30,000 | 100,000 |
|---|---|---|---|---|---|---|---|
| `m̂_hi` median | 1.397 | 1.405 | 1.399 | 1.401 | 1.401 | 1.399 | 1.400 |
| median `\|err\|` | 0.0027 | 0.0067 | 0.0021 | 0.0040 | 0.0027 | 0.0013 | 0.0005 |
| collapsed | 0/6 | 0/6 | 0/6 | 0/6 | 0/6 | 0/6 | 0/6 |
| ladder rungs alive (of 4) | 2 | 3 | 3 | 4 | 4 | 4 | 4 |

**Table 1 — Recovery vs. observation count** (record `v0`, F3; duration channel, 6-seed medians; true dwell lengths ≈ `[500, 108, 24, 5.5]` events per layer).

There is **no collapse anywhere**. The cascade multiplier is recoverable on the duration channel from `~500` events (median `|err| = 0.003`); the full four-rung ladder is identifiable from `~5,000` events on durations (`~10,000` on counts). The rungs that die first as `N` shrinks are the fastest/deepest (layer 4, dwell ~5.5 events, becomes identifiable last) — the honest `k`-identifiability statement, not a hidden collapse. Both cohorts here carry millions of intervals; at Holter length the amplitude error is `0.0005`. **The cascade was seeable at these lengths.** Whatever the tournament decides, it does not decide it because the data were too short.

## 5. The observation channel

*(Record: `out/v0/x2_channel_atlas.json`, finding F2; correction curves wired into `fenosoma/x2.py`, Standing Rule 6.)*

Beat editing and binning are physiology's photon channel: they distort the recovered cascade amplitude, and the distortion must be measured before any exponent is compared across cohorts or cadences. Passing simulated cascades (baseline recovery `m̂_hi = 1.399` against truth `1.4`) through parameterised corruption channels and re-estimating gives the attenuation atlas:

| corruption axis (strongest level) | recovered `m̂_hi` shift vs. baseline | monotone | direction |
|---|---|---|---|
| beat deletion (10%) | `+0.021` | yes | **reads high** |
| spurious insertion (5%) | `+0.018` | yes | **reads high** |
| ectopy left raw (10%) | `−0.006` | yes | flat |
| ectopy + NN editing (10%) | `+0.0002` | yes | flat |
| count-channel binning (30 s) | `−0.140` | yes | reads low |

**Table 2 — The physiological attenuation atlas** (record `v0`, F2; 6-seed medians).

The sibling astrophysics program measured a "reads low, never high" law for the Poisson counting channel. **It does not transfer wholesale to physiology's channel.** Duration-domain corruptions that merge or split intervals inflate the recovered multiplier (`+0.021`, `+0.018`); only binning to counts attenuates it (`−0.140`). The one axis that matters most for the tournament — ectopy — is *fully restored by NN editing* (`+0.0002`), which is exactly the editing view the tournament uses, so the recovered amplitudes it fits are un-attenuated on the axis that could have biased them. No cascade-amplitude comparison in this program ships without the named correction curves.

## 6. The tournament (Q1 decided)

*(Record: `out/v0/x4_tournament.json`, finding F4; the kill gate of `SPECS.md` §7. Figure: `out/v0/fig_soma.png`, tournament panel.)*

Per-record maximum-likelihood fits on `nsr2db` and `chf2db`, NN view, identical `30%` held-out tails, BIC and held-out score both reported:

| cohort | median Δ held-out (cascade − best linear), nats | 95% CI | per 1,000 events | cascade held-out wins |
|---|---|---|---|---|
| `nsr2db` (the kill gate) | **`−27,246`** | `[−32,552, −24,810]` | `−859` | **`0/54`** |
| `chf2db` | `−27,492` | `[−32,835, −21,836]` | `−817` | `0/29` |

**Table 3 — The tournament verdict** (record `v0`, F4; per-record medians, best cascade vs. best linear-group competitor held-out).

The cascade loses on every record, by a wide margin, on both cohorts. Four readings make the verdict robust to the obvious objections:

- **The winner is the field's own untested alternative.** The best held-out model is the tempered-Lévy-type linear process on `79/83` records (`50/54` healthy, `29/29` CHF); BIC agrees on `82/83`. The Barbieri–Brown smooth-intensity point process takes the small remainder. Kiyono et al.'s (2006) tempered-Lévy alternative, made generative and penalised, is the better description of **79 of 83 records** in the two cohorts.
- **The loss is not about history models reading lags.** The Ornstein–Uhlenbeck / DRW-modulated duration model — the *same* inverse-Gaussian emission, the same event clock, **one fewer parameter**, and no access to lagged durations — beats the depth-4 cascade held-out on `77/83` records. The six exceptions are five `nsr2db` records (`nsr006`, `nsr028`, `nsr034`, `nsr037`, `nsr043`) and `chf2db`'s `chf206`: there `ou_ig` loses to `cascade_k4` head-to-head, but the cascade still loses the tournament on each of them — to the tempered-Lévy process on five, and to Barbieri–Brown on `nsr006`, where the tempered-Lévy held-out score degenerates. A single continuous latent factor outscores the hierarchical discrete one on all but those six; the cascade's structure is not buying it anything the data reward.
- **The verdict is bootstrapped.** A parametric bootstrap (24 replicates/record on a pinned 8-record subset, the observed statistic recomputed at the replicate budget) finds the cascade winning `0` of the surrogate fits under the linear null and places every observed deficit at `p ≥ 0.54` for any cascade advantage (most at `p = 1.0`). The power arm — simulating from each record's *own* fitted cascade — recovers cascade wins, so the bootstrap is not simply blind to a cascade when one is present.

- **It is not a starved-data artifact** (§4): the cascade is recoverable from ~500 events; these records carry millions. The cascade was seeable; the data preferred the linear story.

The pre-registered kill gate — *does the exact filter beat the matched linear competitor held-out on the healthy arm?* — fires: **no, `0/54`.** As `SPECS.md` §7 provides, this is the refutation paper.

**What the verdict means, precisely.** Channel B routes *all* serial dependence through the regime ladder, and R–R streams carry strong short-memory structure (respiratory sinus arrhythmia at lags 1–5 beats) that a latent ladder of per-event switch probabilities cannot imitate. The honest claim is therefore that *a multiplicative cascade in the conditional mean does not describe heartbeat durations* — not that heartbeat dynamics contain no cascade at any level. Two sequels bound it: the amplitude structure the cascade misses is named in §7 (a rough log-volatility), and the strong nonlinearity the healthy heartbeat *does* carry is the subject of the companion audit (Paper B).

## 7. The family-level extension: rough, not cascading

*(Record: `out/v2/x6_continuous.json`, finding F6. Amplitudes in the identifiable sfbm reading `λ² = H(1−2H)ν²` per canon §4; the H↔ν² ridge is this repo's own F6 measurement.)*

A skeptic can still object that the tournament tested one *parameterisation* — a `b^k` geometric grid, binary per-layer multipliers, `k+1` distinct conditional means — rather than the cascade hypothesis itself. Record `v2` removes the scaffolding. On the event-index clock, model the log-durations as `ln d_i = μ + ω_i + ε_i` with `Cov(ω_i, ω_j) = λ² · ln⁺(T/(|i−j|+1))`: the log-normal multifractal random walk of Bacry, Delour & Muzy (2001), generalised to the log-S-fBM of Wu, Muzy & Bacry (2022), whose Hurst parameter `H` interpolates from a log-correlated cascade proper (`H → 0`) to a rough monofractal log-volatility (`0 < H < ½`). **No grid, no branching ratio, no depth choice.** Each member is fit per record by GMM on the log-duration variogram; because it carries no exact likelihood it never enters the tournament, and its verdict is delivered by a parametric-bootstrap goodness-of-fit and an energy score over a pre-declared statistic vector.

The estimator is anchored before the data: on `240` seeded synthetic truths at Holter length the `H = 0` null's fitted `Ĥ` never exceeds `0.021` (q95 `0.016`), and the identifiability degeneracy is measured directly — `corr(Ĥ, ln ν²) = −0.98`, so `λ²` (not `Ĥ` or `ν²` separately) is the stable coordinate (now a canon constant, §4). The verdict, per cohort:

| cohort | `n` | median `Ĥ` (95% CI) | `H = 0` (cascade) GoF rejected | `Ĥ > 0` | rough member also rejected |
|---|---|---|---|---|---|
| `nsr2db` | 54 | **`0.111` `[0.088, 0.131]`** | **`54/54`** | `50/54` | `46/54` |
| `chf2db` | 29 | `0.122` `[0.068, 0.159]` | **`29/29`** | `28/29` | `26/29` |

**Table 4 — The continuous-scale verdict** (record `v2`, F6; GMM fits per record, 200-replicate bootstrap GoF).

**The log-correlated member — the continuous-scale cascade proper — is rejected on every 24-hour record in both cohorts**, with the fitted `H` sitting far outside its `H = 0` null: healthy RR log-volatility reads as *rough*, not log-correlated-cascading. The finding is honestly bounded — at `N ≈ 10^5` the GoF has power against any imperfect variogram shape, so the pure log-S-fBM member is itself rejected as a *complete* second-order description on most records; the informative reads are the direction (`Ĥ ≫ 0`) and the comparative score, not a claim that the rough member is the true model. The amplitude is small throughout: in the identifiable sfbm reading, `λ²_sfbm = 0.00044` (healthy) and `0.00016` (CHF); the `H = 0` (MRW) reading of the same field, `λ²_MRW = 0.00109` (healthy), is quoted only beside it and is **roughness-inflated** by the measured program-level curve (×2.32 at `Ĥ ≈ 0.126`), never used alone or compared across domains in `Ĥ` (canon §4).

**Independent reproduction by a different repo's harness.** The healthy-RR amplitude rung above has since been reproduced *outside this repository*, by an instrument we did not write and did not tune. **Fenopan x12, record v12** (the Feno program's meta-repo, canon 1.12.0) runs a turnkey cross-domain amplitude harness against its own sealed landmarks, fitting **from raw admissible streams with no per-domain tuning**; on `nsr2db` it returns a cohort median `λ²_sfbm = 4.46×10⁻⁴` and `Ĥ = 0.115`, landing on this paper's sealed `v2` values (`λ²_sfbm = 0.00044`, median `Ĥ` `0.111` `[0.088, 0.131]`). We cite this as **external corroboration — one reproduction of our measurement, not a second measurement of ours**: no number in this paper is sourced from it, and it enters no interval, count or test here. Its evidentiary weight is that the amplitude rung is a property of the data and the declared estimator rather than of this repository's pipeline, which is precisely the objection a single-repo GMM fit invites.

**What F6 changes.** The founding negative is now a *family-level* statement: neither the discrete `2^k` ladder (record `v0`, exact likelihood) nor the continuous-scale log-correlated field (record `v2`, calibrated GoF on `54/54 + 29/29` records) describes heartbeat durations — the grid was not the reason the cascade lost. And the *direction* of the failure is itself a positive result handed forward: `Ĥ ≈ 0.11`, rough not log-correlated, is an alternative hypothesis with a coordinate, not merely the absence of the cascade.

## 8. Robustness

*(Records: `out/v0/` (bootstrap, two editing policies, estimability), `out/v2/` (bootstrap CI, GoF). The dedicated robustness record `v5` — prompt P-2.1 — is not yet sealed; this section reports the robustness already carried by the sealed records and marks the P-2.1 interval as forthcoming.)*

The verdict is large, and several robustness checks already travel inside the sealed records:

- **Sampling robustness of the verdict.** The parametric bootstrap (§6) puts the cascade advantage at `p ≥ 0.54` throughout, with `0` cascade wins under the linear null and the power arm confirming the bootstrap detects a cascade when one is simulated. The margin (`−27,246` nats/record, CI `[−32,552, −24,810]`) is far from any threshold a sweep could plausibly flip.
- **Editing-policy robustness.** Both editing policies (raw-RR and NN) are run and reported per `SPECS.md` §5 F4; the recovered cascade amplitude is un-attenuated under NN editing on the ectopy axis (§5), so the tournament's fits are not biased by the editing choice on the axis that could have mattered.
- **Continuous-arm agreement.** The family-level rejection (§7) is bootstrapped independently (`Ĥ = 0.111 [0.088, 0.131]`, `54/54` GoF-rejected), a second, grid-free route to the same negative.
- **Cross-repository reproduction of the amplitude rung.** The strongest robustness check on §7's fitted amplitude is not one we ran: **Fenopan x12, record v12** reproduced the `nsr2db` rung from raw admissible streams with a different repository's harness and no per-domain tuning (cohort median `λ²_sfbm` `4.46×10⁻⁴`, `Ĥ` `0.115`; see §7). It is quoted as an independent reproduction of a sealed record, never as an input to a number here — which is why it appears in neither Table 4 nor Appendix A. Note the scope: it corroborates the *amplitude and roughness* reads, not the tournament verdict of §6, which no external harness has re-run.

What the sealed records do **not** yet contain is the *dedicated* robustness sweep that P-2.1 defines and that a submission of this paper requires: (i) the editing-policy sensitivity promoted to a first-class record rather than a per-number annotation; (ii) an **emission-family** sensitivity (inverse-Gaussian vs. Gamma conditional density); (iii) cohort and record-length robustness of the headline deficit; and (iv) the **k-selection honesty** check — a BIC path across cascade depth `k`, the one degree of freedom the founding record pinned by protocol (at `k = 4` the binomial cascade offers only `k+1 = 5` distinct conditional means against the OU competitor's `15` grid levels, and whether a deeper ladder narrows the gap is exactly what the sweep must measure). The apparatus for all six axes is *built* (`fenosoma/x9.py`, the P-2.1 sweep, which quotes `v0`–`v4` and recomputes robustness intervals with the identical `v0` aggregation machinery), but its record `v5` is **not yet sealed** — and by design it promotes to the findings any verdict that flips under a sweep (Standing Rule 2; P-2.1's definition of done). **Every headline number in this paper cites a sealed record; the P-2.1 robustness interval is the one forthcoming addition, and this section reserves its slot rather than quoting a number that does not yet exist.** Per the program's own rule (`PROMPTS.md` §0.2, *the measurement wins*), we mark the gap in the open rather than paper over it.

## 9. Discussion and limitations

The result is a clean, pre-registered negative on a twenty-five-year question: by exact likelihood, a multiplicative regime cascade in the conditional mean is not the generative description of the human heartbeat's inter-beat durations on the canonical cohorts, and the field's own untested tempered-Lévy alternative is the better model of every record. Removing the discrete scaffolding does not rescue the cascade — the continuous-scale log-correlated field is rejected on all `83` records — so the failure is family-level, and the log-volatility that remains is *rough* (`Ĥ ≈ 0.11`), a positive alternative with an amplitude coordinate rather than a void.

Honest limitations:

- **Scope of the claim.** The negative is about a cascade *in the conditional mean of the duration*. It is not a statement that no multiplicative or nonlinear structure exists in the heartbeat; the companion audit (Paper B) shows the healthy heartbeat is strongly nonlinear in a way width statistics cannot see, and §7 names the roughness the cascade misses.
- **`k` is pinned.** The founding record fixes cascade depth by protocol; the `k`-selection BIC path is P-2.1's job (§8), and until `v5` is sealed the paper's depth is a fixed, declared choice rather than a swept one.
- **Emission family.** Durations are modelled inverse-Gaussian (Barbieri–Brown by design); the Gamma-emission sensitivity is part of the pending P-2.1 record.
- **Amplitude coordinates.** Cross-domain amplitude statements, where made, are in the identifiable `λ²_sfbm` reading at matched `(Δ, T, N)`; `Ĥ` is never compared across domains (canon §4), and the deep-cascade comparison is bounded by the estimator's representable-depth ceiling (`k ≈ 20`, `X6_T_MAX`).
- **Downstream program uses are not pre-empted here.** The direction of the failure (`Ĥ = 0.111`) is an input to program-level work elsewhere (a kernel-exponent bridge and a time-irreversibility discriminator); this paper reports the measurement and does not run or claim those flagship experiments.

Every output of this program is a model-comparison or measurement statement about public, de-identified datasets. Nothing here is a medical device, a diagnostic, or an individual-level prediction, and no scheduling- or actuation-flavoured mechanism is described or claimed (the live-instrument arm is deliberately unscheduled pending an owner IP decision).

## 10. Reproducibility

Every number above traces to an immutable, versioned results record under `out/`; the founding record is `out/v0/`, and the family-level extension is `out/v2/`, each pinning the exact environment, seeds, and per-artifact command lines. The apparatus gate (`./verify.sh` → `X0 PASS`) runs from a fresh clone with only `requirements.txt` installed — vendored-import checks, the duration-channel closed-form self-checks, DFA anchors, per-experiment smokes, and a ledger citation check that enforces that every headline number in the claim surface traces to a record. A `--run-slow` byte-for-byte reproduction test replays one pinned tournament fit.

```
python3 -m venv venv && venv/bin/pip install -r requirements.txt
./verify.sh                                  # apparatus gate -> X0 PASS
venv/bin/python -m fenosoma.durations        # Channel B closed forms, standalone
./fetch_data.sh                              # PhysioNet WFDB annotations -> data/
venv/bin/python x3_estimability.py           # §4 recovery vs N
venv/bin/python x2_channel.py                # §5 attenuation atlas
venv/bin/python x4_tournament.py             # §6 the tournament (hours; checkpoint-resumed)
venv/bin/python x6_continuous.py             # §7 the continuous-scale arm
venv/bin/python make_figures.py              # out/v0/fig_soma.png
```

Environment of record: Python 3.11.15, numpy 2.4.6, scipy 1.17.1, wfdb 4.3.1, matplotlib 3.11.1, numba 0.67.0.

## 11. References

*(Reference verification pass run 2026-08-30 against the Crossref REST API (`api.crossref.org`): every DOI below was resolved and its title, container, year, volume and page range checked against the entry. Corrections applied in that pass: the 2004 Kiyono entry gained its DOI and full author list; the 2006 Kiyono entry, previously carrying "exact title/venue to be pinned at submission", is now pinned to the verified item; and the Wu–Muzy–Bacry entry's title was wrong ("Log-S-fBM: log S-fractional Brownian motion and its estimation") and is corrected to the published title with its DOI. One entry could not be discharged and is flagged inline: no work by* Kiyono & Bekki (2011) *exists in Crossref under any author-pair or bibliographic query — see the entry for the verified near neighbours and the two in-text attributions that depend on it. The Calvet–Fisher 2008 monograph is a book and carries no Crossref DOI; its ISBN is quoted unverified by this pass. Everything else in this list is verified.)*

- **Bacry, E., Delour, J., & Muzy, J.-F. 2001**, "Multifractal random walk," *Physical Review E*, 64, 026103. doi:10.1103/PhysRevE.64.026103
- **Barbieri, R., Matten, E. C., Alabi, A. A., & Brown, E. N. 2005**, "A point-process model of human heartbeat intervals: new definitions of heart rate and heart rate variability," *American Journal of Physiology — Heart and Circulatory Physiology*, 288, H424–H435. doi:10.1152/ajpheart.00482.2003
- **Calvet, L. E., & Fisher, A. J. 2001**, "Forecasting multifractal volatility," *Journal of Econometrics*, 105, 27–58. doi:10.1016/S0304-4076(01)00069-0
- **Calvet, L. E., & Fisher, A. J. 2004**, "How to Forecast Long-Run Volatility: Regime Switching and the Estimation of Multifractal Processes," *Journal of Financial Econometrics*, 2, 49–83. doi:10.1093/jjfinec/nbh003
- **Calvet, L. E., & Fisher, A. J. 2008**, *Multifractal Volatility: Theory, Forecasting, and Pricing* (Academic Press / Elsevier). ISBN 978-0-12-150013-9.
- **Goldberger, A. L., et al. 2000**, "PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals," *Circulation*, 101, e215–e220. doi:10.1161/01.CIR.101.23.e215
- **Ivanov, P. Ch., Amaral, L. A. N., Goldberger, A. L., et al. 1999**, "Multifractality in human heartbeat dynamics," *Nature*, 399, 461–465. doi:10.1038/20924
- **Kiyono, K., Struzik, Z. R., Aoyagi, N., Togo, F., & Yamamoto, Y. 2005**, "Phase transition in a healthy human heart rate," *Physical Review Letters*, 95,
- doi:10.1103/PhysRevLett.95.058101
- **Kiyono, K., Struzik, Z. R., Aoyagi, N., Sakata, S., Hayano, J., & Yamamoto, Y. 2004**, "Critical Scale Invariance in a Healthy Human Heart Rate," *Physical Review Letters*, 93, 178103. (the multiplicative log-amplitude cascade picture of the healthy heartbeat) doi:10.1103/PhysRevLett.93.178103
- **Kiyono, K., Struzik, Z. R., Aoyagi, N., & Yamamoto, Y. 2006**, "Multiscale Probability Density Function Analysis: Non-Gaussian and Scale-Invariant Fluctuations of Healthy Human Heart Rate," *IEEE Transactions on Biomedical Engineering*, 53(1), 95–102 — the heartbeat-increment lineage entry: the non-Gaussian / tempered-Lévy-type (linear) alternative found to describe the increment distribution equally well, *the* Q1 rival, made generative here. doi:10.1109/TBME.2005.859804 *(Bibliographic details verified 2026-08-30. The "tempered-Lévy" characterisation is this paper's reading of that result, not a phrase quoted from its title.)*
- **Kiyono, K., & Bekki, N. 2011**, on the difficulty of confirming the multiplicative mechanism in heartbeat dynamics from increment statistics alone. ⚠️ **Unverified — not discharged by the 2026-08-30 Crossref pass.** No work by this author pair in this year is indexed in Crossref (searched by author pair, by author + topic, and by bibliographic query). The nearest **verified** items making the same methodological point are *Kiyono, Struzik & Yamamoto 2007*, "Estimator of a non-Gaussian parameter in multiplicative log-normal models," *Physical Review E* 76,

041113 (doi:10.1103/PhysRevE.76.041113) and *Kiyono, Struzik & Hayano 2011*, "Asymmetric intermittency observed in human heart rate dynamics," *Proc. IEEE EMBC 2011*, 7743–7746 (doi:10.1109/IEMBS.2011.6091908). **This entry, and the two in-text attributions that rest on it (§1 and §1's historical paragraph), must be resolved to a real source or dropped before submission** — neither carries a number used anywhere in this paper.

- **Lin, D. C., & Hughson, R. L. 2001**, "Modeling heart rate variability in healthy humans: A turbulence analogy," *Physical Review Letters*, 86, 1650–1653. doi:10.1103/PhysRevLett.86.1650
- **Schreiber, T., & Schmitz, A. 1996**, "Improved surrogate data for nonlinearity tests," *Physical Review Letters*, 77, 635–638. doi:10.1103/PhysRevLett.77.635
- **Wu, P., Muzy, J.-F., & Bacry, E. 2022**, "From rough to multifractal volatility: The log S-fBM model," *Physica A: Statistical Mechanics and its Applications*, 604, 127919 (the multifractal-to-rough interpolation of the log-field). doi:10.1016/j.physa.2022.127919

## Data availability & acknowledgments

All results reproduce from committed records and pinned seeds; the estimability and atlas experiments (§4–5) require no external data. The heartbeat data are the public PhysioNet **Normal Sinus Rhythm** (`nsr2db`) and **Congestive Heart Failure** (`chf2db`) RR-interval databases, fetched at run time and never committed; results use the committed beat annotations only, under the declared NN editing policy. Every artifact using these data cites the source database and PhysioNet (Goldberger et al. 2000). The vendored MSM theory/estimator/filter modules and the surrogate module are CC BY 4.0 (see PROVENANCE.md); the Barbieri–Brown point process and the IAAFT surrogates are reimplemented from the published equations (Barbieri & Brown 2005; Schreiber & Schmitz 1996) and say so. 

## Appendix A — Provenance table (every headline number → its

## sealed record)

| number (as used) | record | file / section |
|---|---|---|
| `nsr2db` `0/54` held-out wins, median deficit `−27,246` nats, CI `[−32,552, −24,810]`, `−859`/1000 ev | v0 | `out/v0/x4_tournament.json`, `kill_gate.nsr2db` |
| `chf2db` `0/29`, median deficit `−27,492` nats, CI `[−32,835, −21,836]` | v0 | `out/v0/x4_tournament.json`, `kill_gate.chf2db` |
| tempered-Lévy best held-out `79/83`, best BIC `82/83` | v0 | `out/v0/x4_tournament.json`, `best_heldout_counts` / `best_bic_counts` |
| OU model beats `cascade_k4` held-out on `77/83` records (exceptions `nsr006`, `nsr028`, `nsr034`, `nsr037`, `nsr043`, `chf206`) | v0 | `out/v0/x4_tournament.json`, `ou_ig` vs `cascade_k4` |
| bootstrap `0` cascade wins under null, `p ≥ 0.54`; power arm recovers wins | v0 | `out/v0/x4_tournament.json`, `bootstrap` |
| recovery from ~500 events, full ladder ~5,000 (durations); `0/6` collapsed | v0 | `out/v0/x3_estimability.json`, `summary` |
| atlas: deletion `+0.021`, insertion `+0.018`, ectopy+NN `+0.0002`, binning `−0.140` | v0 | `out/v0/x2_channel_atlas.json`, per-axis verdicts |
| edited-beat fraction `0.0005` (nsr2db) / `0.0125` (chf2db) | v0 | `out/v0/x4_tournament.json`, `cohorts.*.median_edited_beat_fraction` |
| `H = 0` GoF rejected `54/54` + `29/29`; `Ĥ > 0` on `50/54` + `28/29` | v2 | `out/v2/x6_continuous.json`, route (i) |
| median `Ĥ` `0.111 [0.088, 0.131]` (nsr2db), `0.122 [0.068, 0.159]` (chf2db) | v2 | `out/v2/x6_continuous.json`, `H_hat_*` |
| `λ²_sfbm` `0.00044` / `0.00016`; `λ²_MRW` `0.00109` (inflated ×2.32) | v2 | `out/v2/x6_continuous.json`, `lam2_hat_*` |
| `corr(Ĥ, ln ν²) = −0.98`; `H = 0` null `Ĥ` q95 `0.016` at `N=10^5` | v2 | `out/v2/x6_continuous.json`, calibration/degeneracy |
| robustness interval (editing-policy record, IG-vs-Gamma, k- | v5 *(pending — P-2.1 apparatus built,* | `fenosoma/x9.py` |

| number (as used) | record | file / section |
|---|---|---|
| selection BIC path) | *record not yet sealed; §8)* | |
| `t_MF` / scattering nonlinearity of the same cohorts | Paper B / v3 *(companion)* | `papers/02-two-instrument-audit.md` |