Two instruments on one null: the reported multifractality of the human heartbeat, re-audited

On byte-identical IAAFT surrogates, MFDFA width finds healthy-cohort multifractality mostly absent (`19/54`) while a power-calibrated wavelet-scattering statistic finds decisive nonlinearity on `54/54` — the width's null result is a statement about the width estimator's power, not about the heartbeat.

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

Draft — methods/result paper (the two-instrument audit, `Q2`); the physiological instance of the width/IAAFT incoherence (canon §6.2). Target venue: arXiv first (physics.data-an / q-bio.NC), then a physiology/complexity-science venue — not a cardiology journal (`SPECS.md` §7); journal shortlist is an owner call (`OWNER-QUEUE.md` item 3). Every quantitative claim traces to a named, immutable results record under `out/`: the width audit `out/v0/` (`x1`) and the scattering audit `out/v3/` (`x7`), run on the byte-identical IAAFT null. The program-level mechanism, pincer, and known-truth power/specificity certificate are not re-derived here — they are owned by the sibling methods-core preprint (Fenocosm [`papers/08-width-power.md`], "The width and the null", doi:10.5281/zenodo.22180477 — published 2026-08-31), which this paper is the physiological instance of and cites; see §1.3. The apparatus gate (`./verify.sh` → `X0 PASS`) runs from a fresh clone; the scattering statistic's time-reversal anchor is asserted in `tests/test_scattering.py` (module `fenosoma/scattering.py`). Figure: `out/v0/fig_soma.png` (audit panel), regenerable by `python make_figures.py`.

Priority note: this finding has no DOI and no priority date. The live t_MF robustness line (Mangalam, Likens & Kelty-Stephen 2025) is active and is engaged head-on in §7. 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Correction, 2026-08-31: the priority note above is superseded on its first clause — this paper's Zenodo deposit was published on 2026-08-31 and its DOI, doi:10.5281/zenodo.22180531, now resolves; the publication date is the priority date (`PROVENANCE.md` §5). The sibling methods-core preprint's DOI (10.5281/zenodo.22180477, quoted above and in §1.3) was published the same day, and its two "reserved; publication pending" marks in this draft are updated in place accordingly. The rest of the note stands: the t_MF engagement is unchanged, and arXiv posting remains an owner action.

Abstract

A large clinical and complexity-science literature quotes multifractal detrended-fluctuation-analysis (MFDFA) *widths* of physiological signals as evidence of nonlinear, cascade-like dynamics, and reads a width that does **not** exceed its phase-randomised (IAAFT) surrogate band as evidence that a signal is monofractal, linear, or "not a cascade." **We show, on the human heartbeat, that this reading is an artifact of the width estimator's own low power, not a property of the heart.** We run two detection statistics on the *byte-identical* IAAFT null — regenerated from the same pinned surrogate seeds, so only the statistic differs — across four canonical PhysioNet RR cohorts. Instrument one is the field-standard MFDFA two-point width; instrument two is a wavelet-scattering-moment omnibus (sparsity, normalised second-order scattering, cross-scale modulus correlation, and one phase/modulus time-asymmetry coordinate proven exactly odd under time reversal), with a pre-registered specificity and power battery.

Under the width, healthy-cohort multifractality mostly *collapses* against the matched linear null (`nsr2db`: `19/54` records significant under NN editing, median IAAFT `z = 0.60`), while the congestive-heart-failure cohort retains it (`chf2db`: `18–25/29`). Under the scattering omnibus, on the *same* surrogates, the healthy collapse does not harden — **it inverts**: decisive excess on **`54/54`** healthy records (both editing policies), carried predominantly by *time-irreversibility* (median aggregate `z` `73` under NN, significant on `53/54`), with a real secondary intermittency component; the CHF cohort is `29/29`, carried instead by the multiplicative sparsity / second-order families. The specificity battery places the linear-control false-positive rate at `0.010` (nominal `0.030`) and the power curve reaches `16/16` at Holter length. The width's healthy null result was therefore a statement about the width estimator's power on this cohort; a sharper instrument, on the identical null, rejects the same records. This is the physiological instance of the program's finding that **the width/IAAFT pairing is incoherent** (canon §6.2): a null width-vs-IAAFT result is never, on its own, evidence of linearity. We read the live `t_MF` robustness counterclaim head-on against its own reported cascade split (§7). All results reproduce from a fresh clone behind a deterministic gate. **Nothing here is clinical**; time-irreversibility is a documented physiological feature, and every statement is a measurement about public, de-identified datasets.

---

1. Introduction

1.1 A field that quotes raw widths

The multifractal width of a physiological signal — the span of its singularity spectrum, measured by MFDFA — is one of the most-quoted nonlinear biomarkers in the complexity-science and clinical-physiology literatures. The canonical positive result is Ivanov et al. (1999): the healthy heartbeat is broadly multifractal, and the width narrows in disease. A large downstream literature reports *raw* widths, or widths compared only to shuffled surrogates, as evidence of cascade-like dynamics. But a raw width is not claim-bearing (`SPECS.md` §5 F2; canon F2): temporal structure of *any* kind — including purely linear correlation — inflates it. The disciplined test compares the width against a **matched IAAFT surrogate** ensemble that preserves the signal's marginal distribution and power spectrum while randomising its phases, so that only *nonlinear* cross-scale structure can produce excess.

1.2 The incoherent pairing

The sibling astrophysics program measured that most apparent multifractal width in an independent (astrophysical) sample *collapses* under IAAFT surrogates — i.e., is linear-correlation artifact. The natural inference — *width collapses under IAAFT ⇒ the process is linear* — is the one this paper, and the program, show to be void. The reason is structural and is stated at canon level:

> **The width/IAAFT pairing is incoherent** (canon §6.2): a null width-vs-IAAFT result is a statement about the width statistic's power, never evidence of linearity.

The MFDFA width is a *second-order, time-symmetric* functional — built from a variance of detrended fluctuations, invariant under time reversal. IAAFT surrogates return *exactly* the width-bearing part of a nonlinear signal (its marginal and spectrum) while destroying the fourth-order cross-scale modulus coupling and the time-asymmetry where multiplicative and nonlinear structure actually live. The width therefore cannot *see* the structure the IAAFT null *destroys*: their overlap — the test's detection power — can be near zero. Consequently (canon §7 rule 5) an amplitude/nonlinearity claim requires an **IAAFT-class null *plus* a power-calibrated statistic**; a bare width-vs-IAAFT result licenses nothing. This paper supplies the power-calibrated statistic on the heartbeat and reads the two instruments side by side on one null.

1.3 Scope: this paper and the methods-core preprint (recorded before drafting)

This finding has a program-level twin, and the two must not become one paper. The **methods-core** — the *mechanism* (the incoherent pairing), the *pincer* (the width reads too high on trends/finite-size/isolated singularities and too low on genuine cascades), the synthetic *known-truth power/specificity certificate* on the canonical binomial cascade, the program-level engagement of the `t_MF` counterclaim and its adversarial-collaboration offer, and the proposed reporting standard — is owned by the sibling preprint **Fenocosm [`papers/08-width-power.md`], "The width and the null"** (a `physics.data-an` methods preprint; doi:10.5281/zenodo.22180477 — published 2026-08-31). That paper cites *this* repository's sealed record `v3` as its "independent same-day replication" (its §5.6) and flags that those numbers must be verified at their source.

This paper (Paper B) is the physiological instance. Its scope is: the actual PhysioNet RR cohorts and their own sealed records (`v0`'s `x1` width audit and `v3`'s `x7` scattering audit); the instance-level numbers (`19/54 → 54/54` on `nsr2db`, `18–25/29 → 29/29` on `chf2db`, the τ-domination *in physiology*, the byte-identical IAAFT null); and the physiological interpretation. Paper B **cites** the methods-core preprint for the mechanism, the pincer, and the synthetic certificate rather than re-deriving them, reciprocally to how the preprint cites `v3`. The two are a methods-core / domain-instance pair, deliberately kept distinct. This scope split is recorded here and in `papers/README.md`.

1.4 Contributions

1. A two-instrument audit of heartbeat multifractality on the *byte-identical* IAAFT null — the width statistic and a power-calibrated scattering omnibus on the same surrogate realisations, so the comparison isolates the statistic (§4, §5).
2. The measured *inversion* (§5): where the width finds healthy multifractality mostly absent (`19/54`), the scattering statistic finds decisive nonlinearity on `54/54` of the same records, dominated by time-irreversibility — the physiological instance of canon §6.2.
3. A pre-registered specificity/power certificate *for the physiological channel* (§4.3): linear-control false-positive rate `0.010` (nominal `0.030`); power `16/16` at Holter length.
4. A *converse* null (§5.1): the tournament's own fitted linear winner (Paper A) leaves the scattering structure unexplained on `51/51` healthy and `26/26` CHF records — a stronger null than IAAFT, and along the way a diagnosis that `6/83` fitted tempered-Lévy MLEs are not stationary generators at all.
5. A head-on engagement with the live `t_MF` robustness counterclaim, read against its own reported cascade split, as the physiological instance (§7).

---

2. Data and the shared null

(Record: `out/v0/x1_audit.json` and `out/v3/x7_scattering.json`; editing policy per number, `SPECS.md` §5 F4; the IAAFT surrogates regenerated from `x1`'s exact pinned seeds.)

The audit runs on four PhysioNet RR databases: **Normal Sinus Rhythm** (`nsr2db`, `54` subjects, healthy), **Congestive Heart Failure** (`chf2db`, `29`, the Ivanov contrast), **BIDMC CHF** (`chfdb`, `15`), and **Fantasia** (`fantasia`, `40`, the aging contrast) — `138` records in all, each run under **both** editing policies (raw-RR and NN) side by side (`SPECS.md` §5 F4), with the edited-beat fraction carried per record. Committed PhysioNet beat annotations are the only event source; every artifact cites the source database and PhysioNet (Goldberger et al. 2000).

The decisive design choice is the shared null. Both instruments are tested against a matched ensemble of `32` IAAFT surrogates per record; for the scattering audit the surrogates are **regenerated from the width audit's exact pinned seeds**, so the null realisations are byte-identical between the two instruments. **Only the detection statistic differs.** Any difference in verdict is therefore attributable to the statistic, not to the null.

3. The two instruments

(Statistic pinned in `fenosoma/protocol.py` block `p3` before the first record was read; module `fenosoma/scattering.py`, reimplemented from the published equations — nothing vendored.)

Instrument 1 — the width. MFDFA singularity-spectrum width, `h(−4) − h(+4)`, the field-standard statistic. It is second-order and time-symmetric by construction. Significance is IAAFT-excess at `z ≥ 2`, with the shuffle ensemble reported beside it for decomposition (`SPECS.md` §5 F2).

Instrument 2 — the scattering omnibus. An analytic dyadic log-normal wavelet (`ξ0 = 0.75π`, `σ_ln = 0.55`, `J = 12` scales, centre periods ≈ `2.7 … 5,461` beats) gives, per record: first-order **sparsity** `S1(j)/σ_j` (the Rayleigh value `√(π/4) ≈ 0.886` for a linear-Gaussian coefficient, below it under intermittency); the normalised second-order scattering `S2(j1,j2)/S1(j1)` (the intermittency carrier); the wavelet-modulus **cross-scale correlation**; and one **time-asymmetry** coordinate `Im ρ(j1,j2)`, *proven exactly odd under circular time reversal* (the module anchor, asserted in the gate), hence a pure time-irreversibility probe that is zero for any exactly reversible sample and for a stationary Gaussian. The aggregate excess is a per-coefficient `z` against the surrogate ensemble combined into `T = Σz²`; the **calibrated significance is the rank of the real `T` among the surrogates' own leave-one-out aggregates**, so with `32` IAAFT

surrogates the per-test false-positive rate is exactly `1/33 ≈ 0.030` by exchangeability. (The aggregate magnitude `z_agg` reflects the *number* of jointly significant coefficients and is reported as a magnitude only; the significance is the rank-`p` and the specificity budget, not `z_agg`.)

3.1 The battery is anchored before the data

(Record: `out/v3/x7_scattering.json`, `anchor_battery`; anchors before claims, `SPECS.md` §5 F1.)

Specificity. Fractional Gaussian noise at four Hurst exponents, an AR(4)-Gaussian, and i.i.d. Student-`t` — linear-Gaussian and heavy-tailed controls with **no** multiplicative or nonlinear structure, at manifest-matched `N = 100,000` — flag `1/96` significant against their *own* IAAFT ensembles: a false-positive rate of **`0.010`** against the expected `0.030` (fGn `1/64`, AR(4) `0/16`, i.i.d.-`t` `0/16`) — consistent with the nominal 0.030 (1 of 96 significant; expected 2.9; P = 0.21). Each instance is an exact rank test of per-test size `1/(n_sur + 1) = 0.0303`, so the expected count under the null is `96 × 0.0303 = 2.91` and `P(X ≤ 1 | Binom(96, 1/33)) = 0.21`: the battery certifies that the instrument is **not anti-conservative**, and has no power to establish that it is conservative. *Power.* The vendored cascade at the founding operating point shows detection fraction rising with `N`:

`N`	2,000	5,000	10,000	30,000	100,000
detection fraction	`0.625`	`1.00`	`1.00`	`1.00`	**`16/16`**

Table 1 — Scattering power and specificity (record `v3`; power on cascade truths, specificity `0.010` on linear/heavy-tailed controls). The excess is carried by the sparsity and normalised-second-order families (the multiplicative carriers), with the asymmetry family flat on a symmetric-switching cascade — exactly as expected. This is the power calibration canon §7 rule 5 demands, and it travels with every number below.

4. Instrument 1: the width audit

(Record: `out/v0/x1_audit.json`, finding F1.)

cohort (policy)	n	IAAFT-significant	median IAAFT z	median shuffle z	median raw width
nsr2db (raw)	54	21/54	0.90	12.2	0.109
nsr2db (NN)	54	**19/54**	0.60	11.8	0.116
chf2db (raw)	29	25/29	5.37	11.2	0.265
chf2db (NN)	29	18/29	3.02	12.9	0.249
chfdb (raw / NN)	15	10/15 / 4/15	2.69 / 0.06	5.5 / 3.6	0.096 / 0.100
fantasia (raw / NN)	40	12/40 / 10/40	1.03 / 0.31	4.3 / 4.8	0.221 / 0.177

Table 2 — The width audit (record v0, F1; z ≥ 2 = IAAFT-significant; 32 IAAFT + 32 shuffle surrogates).

Two facts. First, temporal structure is never in doubt: shuffle z ≈ 4–13 everywhere. But the IAAFT-corrected picture *inverts the naive one* — **healthy-cohort width mostly collapses under the matched linear null** (median IAAFT z ≈ 0.6–0.9; 19/54 NN significant), while chf2db retains a genuinely nonlinear component under both policies (median IAAFT z 3.0–5.4). Second, the canonical Ivanov contrast — that multifractality *narrows* in CHF — **fails to reproduce in raw width on these 24-hour streams**: chf2db's median raw width (0.25–0.27) is more than twice nsr2db's (0.11–0.12); CHF is *wider*, not narrower. Read on its own, the healthy result would license the inference "healthy heartbeat multifractality is mostly linear artifact." That inference is what the second instrument voids.

5. Instrument 2: the scattering audit

(Record: out/v3/x7_scattering.json, finding F7; the byte-identical x1 IAAFT null.)

cohort (policy)	n	scattering significant	width significant (Table 2)	dominant excess family (median z_agg)
nsr2db (raw)	54	**54/54**	21/54	time-asymmetry (156)
nsr2db (NN)	54	**54/54**	19/54	time-asymmetry (73), 53/54 records
chf2db (raw)	29	**29/29**	25/29	sparsity / 2nd-order (≈712 agg)
chf2db (NN)	29	**29/29**	18/29	sparsity / 2nd-order (456 / 467)

Table 3 — The scattering audit, beside the width (record v3, F7; same null as Table 2).

The healthy-cohort collapse does not harden — it inverts. Where the width found most healthy records consistent with their linear-Gaussian IAAFT null, the scattering statistic finds decisive excess on **every one of the `54` `nsr2db` records**, against the same surrogates. The healthy excess is carried predominantly by the **time-asymmetry** family (NN median `z_agg = 73`, significant on `53/54` records) with a real secondary intermittency/sparsity component (sparsity significant on `50/54`, second-order on `54/54`); the record-level check is direct — on `nsr001` the observed median `|Im ρ|` is ≈`5×` its IAAFT surrogate's (`0.13` vs `0.03`), and the observed sparsity sits *below* the Rayleigh line where the surrogate's sits on it. Healthy R–R is not linear-Gaussian; it is strongly nonlinear in a way a symmetric two-point width functional cannot see — heart-rate time-irreversibility is a documented physiological feature, orthogonal to the spectrum and marginal IAAFT preserves. The CHF cohort retains excess under scattering too (`29/29`, both policies), but carried by the **sparsity / normalised-second-order** families (the multiplicative carriers, matching the cascade's power signature and the width finding) rather than by asymmetry — the two cohorts fail the linear null through *different* structure, and the scattering decomposition is what separates them.

5.1 The converse null

(Record: `out/v3/x7_scattering.json`, converse arm.)

Beside IAAFT we run a sharper null: per record, the scattering vector against `48` seeded replicas from that record's *own* fitted linear held-out winner (the tempered-Lévy AR(4)-NIG model of Paper A). This exposes a property of the tournament's winner first: **`6` of the `83` fitted tempered-Lévy MLEs are not stationary generators at all** (the AR coefficients sum above one, or the NIG shape collapses so the draws are constant) — `nsr009/022/042` and `chf207/211/227`; they win *predictive* likelihood but are ill-defined as simulation nulls and are reported separately. Over the records with a generatively valid winner, **the fitted winner leaves scattering structure systematically unexplained on `51/51` healthy and `26/26` CHF records**: none sit inside their winner's ensemble. This names exactly what the tournament's winner misses — the cross-scale envelope and time asymmetry — against a *stronger* null than IAAFT. Paper A's held-out likelihood verdict is a different functional and is untouched: the tempered-Lévy model remains the best one-step predictive density; it is simply not a complete generative account of the record's multiscale envelope and time asymmetry.

6. The incoherent pairing, priced

(Mechanism, pincer, and synthetic power/specificity certificate: Fenocosm [`papers/08-width-power.md`], cited not re-derived — §1.3.)

The two instruments on one null are the physiological instance of the general result. The width's healthy null was not evidence of linearity; it was the width estimator's low power on this cohort, and a power-calibrated statistic on the identical surrogates rejects the same records. This is canon §6.2 in the heart, and it is the reason canon §7 rule 5 requires an IAAFT-class null *plus* a power-calibrated statistic before any amplitude/nonlinearity claim: the width alone, however carefully surrogate-corrected, supplies only half of that and can return a confident null on a strongly nonlinear signal. The general mechanism (the width is second-order and time-symmetric; IAAFT returns the marginal that carries it; the discriminating structure is fourth-order and time-asymmetric), the two-sided *pincer* (too high on trends and isolated singularities, too low on genuine cascades), and the synthetic known-truth certificate on the canonical binomial cascade are established at program level in the methods-core preprint; this paper does not restate them, and confines its claims to the sealed physiological records `v0` and `v3`.

What this does and does not say about cascades. Scattering excess is a *broader* object than "multiplicative cascade": time-irreversibility and envelope intermittency are nonlinearities IAAFT destroys, but neither is, on its own, evidence of a `2^k` or continuous-scale *cascade*. Paper A's verdicts — the cascade loses `0/83` by exact likelihood, and the log-volatility is rough not log-correlated (`54/54` GoF-rejected) — are untouched here. What this audit overturns is narrower and specific: the inference from "healthy multifractal width collapses under IAAFT" to "healthy R–R multifractality is mostly linear artifact" does not survive a more powerful statistic on the same null. The nonlinearity is there; it is not width-shaped.

7. Direct engagement with the t_MF robustness counterclaim

(The physiological instance of the program-level engagement in Fenocosm [`papers/08-width-power.md`] §6; the counterclaim the canon registers as a live gap.)

A live line of work argues the opposite of this paper: that multifractal nonlinearity, distilled into the one-sample statistic `t_MF`, is a **robust** estimator of multiplicative cascade dynamics (Mangalam, Likens & Kelty-Stephen 2025, *Phys. Rev. E* **111**, 034126; arXiv:2312.05653). Because that claim and this one cannot both be read naively, we engage it directly and read its formula first.

Their statistic, quoted, not paraphrased. The multifractal-nonlinearity `t`-statistic compares the original series' singularity-spectrum width to the widths of its IAAFT surrogates and standardises the difference. In the lineage's own words (Bell et al. 2019, *Front. Physiol.* **10**:998):

> `t_MF = (W_MF of original series − Average W_MF of surrogate-data series) / (Standard error of W_MF of surrogate-data series)`

with the surrogate ensemble generated by IAAFT and a significant *positive* `t_MF` read as evidence of nonlinear cross-scale interactions.

Their own numbers are the calibration. The robustness paper simulates cascade generations and reports, verbatim, the distribution of `t_MF` across `7000` simulated cascade generations:

> "there were 3592 significant positive values of `t_MF` … for **51.31% of the simulated generations**."
>
> "we found 1895 such significant negative `t_MF`, **27.07% of the total simulated generations**."

Read as a power/specificity certificate on genuine cascades — which is exactly what it is — this confirms the pincer rather than refuting it:

- **Detection power ≈ `0.51` on the positive control.** On series *known* to be multiplicative cascades, the statistic returns the correct (significant-positive) verdict barely more than half the time.
- **A `27%` wrong-sign rate.** On more than a quarter of genuine cascades the statistic is significant in the *negative* direction — declaring the cascade's spectrum *narrower* than its linear surrogates, reading a true cascade as sub-linear / monofractal. This is precisely the low jaw of the pincer, and the sign flip Lee & Kelty-Stephen (2017) warned of. (The remaining ≈ `21.6%` are null.)

The disagreement is not about the numbers — we accept theirs — but about their interpretation. A statistic correct `51%` of the time and confidently-wrong `27%` of the time on its own positive control is not thereby shown robust; it has the two-sided bias the program names. The `t_MF` framework is a fine *feature extractor*; the inference "significant `t_MF` ⇒ cascade / null `t_MF` ⇒ not a cascade" is what fails, in both directions, by their own measurement — **and the heartbeat is exactly where it fails.** The healthy width null (`19/54`, §4) is what a `~51%`-power statistic produces on a genuinely nonlinear cohort; the scattering audit (`54/54`, §5) is the power-calibrated instrument canon §7 rule 5 requires, and it finds the nonlinearity the width missed.

Adversarial collaboration, offered. The productive path is not duelling preprints. We offer an adversarial collaboration: a jointly pre-registered battery — agreed cascade constructions, agreed physiological lengths and surrogate budgets, agreed decision rule

— on which both `t_MF` and the scattering omnibus report power and specificity, with the sign convention and the `Δα`-vs-`h(q)` width definition fixed in advance. (Whether to open that collaboration is an owner decision — this paper only offers it.)

8. What this audit does and does not license

The audit is a metrology contribution to a field with a known reproducibility problem: it shows that a null width-vs-IAAFT result on the heartbeat carries no information about linearity, and it supplies the power-calibrated instrument that does. Honest bounds:

- **Not a cascade claim.** `54/54` decisive scattering excess is evidence of *nonlinear* structure (dominated by time-irreversibility in the healthy cohort), not of a multiplicative cascade; Paper A's exact-likelihood and continuous-scale verdicts against the cascade stand.
- **The width is not useless.** It remains a valid *likelihood/biomarker* feature and a correct detector of the multiplicative families in the CHF cohort; what dies is the *converse inference* ("null ⇒ linear"). Positive width rejections become under-claims, not errors.
- **Time-irreversibility is physiological, not pathological.** Its dominance in the healthy cohort is a documented feature of cardiovascular control; nothing here is a diagnostic, predictive, or individual-level statement (`SPECS.md` §5 F6). Every output is a model-comparison or measurement statement about public, de-identified datasets.

9. Reproducibility

Every number above traces to an immutable, versioned results record under `out/`: the width audit `out/v0/` and the scattering audit `out/v3/`, each pinning environment, seeds, and per-artifact command lines. The scattering audit regenerates the width audit's exact IAAFT surrogate seeds, so the two instruments share a byte-identical null. The apparatus gate (`./verify.sh` → `X0 PASS`) runs from a fresh clone with only `requirements.txt`; the scattering statistic's exact-odd-under-time-reversal anchor and the specificity/power budget are asserted in the test suite and the ledger citation check enforces that every headline number traces to a record.

```
python3 -m venv venv && venv/bin/pip install -r requirements.txt
./verify.sh                                  # apparatus gate -> X0 PASS
./fetch_data.sh                              # PhysioNet WFDB annotations -> data/
venv/bin/python x1_signatures.py             # §4 the width audit (Q2)
venv/bin/python x7_scattering.py             # §5 the scattering audit (same null)
venv/bin/python make_figures.py              # out/v0/fig_soma.png (audit panel)
```

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.

10. 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. Three corrections applied in that pass — the Bell et al. 2019 title was wrong (the entry carried an invented subtitle, "The case of surrogate-based `t_MF` in physiological time series"); the Lee & Kelty-Stephen 2017 entry carried no title at all; and the Mangalam et al. 2025 title was mis-transcribed as "Robustness of multifractal nonlinearity as a `t_MF` estimator …". All three are corrected below against the registered metadata, and the Mangalam preprint's title was additionally checked against arXiv:2312.05653, which agrees. This list is fully discharged: every entry is verified, and no unresolved item remains.)

- **Bell, C. A., Carver, N. S., Zbaracki, J. A., & Kelty-Stephen, D. G. 2019**, "Non-linear Amplification of Variability Through Interaction Across Scales Supports Greater Accuracy in Manual Aiming: Evidence From a Multifractal Analysis With Comparisons to Linear Surrogates in the Fitts Task," *Frontiers in Physiology*, 10:998. doi:10.3389/fphys.2019.00998
- **Bruna, J., Mallat, S., Bacry, E., & Muzy, J.-F. 2015**, "Intermittent process analysis with scattering moments," *Annals of Statistics*, 43(1), 323–351. doi:10.1214/14-AOS1276
- **Goldberger, A. L., et al. 2000**, "PhysioBank, PhysioToolkit, and PhysioNet," *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
- **Lee, H., & Kelty-Stephen, D. G. 2017**, "Cascade-Driven Series with Narrower Multifractal Spectra Than Their Surrogates: Standard Deviation of Multipliers Changes Interactions across Scales" (allied prior art on the bidirectional `w_orig < w_surr` reading), *Complexity*, 2017:7015243, 1–8. doi:10.1155/2017/7015243
- **Mangalam, M., Likens, A. D., & Kelty-Stephen, D. G. 2025**, "Multifractal nonlinearity as a robust estimator of multiplicative cascade dynamics," *Physical Review E*, 111, 034126; arXiv:2312.05653. doi:10.1103/PhysRevE.111.034126
- **Morel, R., Rochette, G., Leonarduzzi, R., Bouchaud, J.-P., & Mallat, S. 2025**, "Scale dependencies and self-similar models with wavelet scattering spectra," *Applied and Computational Harmonic Analysis*, 75, 101724. doi:10.1016/j.acha.2024.101724
- **Schreiber, T., & Schmitz, A. 1996**, "Improved surrogate data for nonlinearity tests," *Physical Review Letters*, 77, 635–638. doi:10.1103/PhysRevLett.77.635

Data availability & acknowledgments

All results reproduce from committed records and pinned seeds. The heartbeat data are the public PhysioNet **Normal Sinus Rhythm** (`nsr2db`), **Congestive Heart Failure** (`chf2db`), **BIDMC CHF** (`chfdb`), and **Fantasia** (`fantasia`) RR-interval databases, fetched at run time and never committed; results use committed beat annotations only, under both declared editing policies. Every artifact using these data cites the source database and PhysioNet (Goldberger et al. 2000). The vendored MSM and surrogate modules are CC BY 4.0 (see PROVENANCE.md); the wavelet scattering statistic and the IAAFT surrogates are reimplemented from the published equations (Bruna, Mallat, Bacry & Muzy 2015; Morel et al. 2025; Schreiber & Schmitz 1996) and say so.

Appendix A — Provenance table (every headline number → its sealed record)

number (as used)	record	file / section
width `nsr2db` `21/54` (raw) / `19/54` (NN), median IAAFT `z` `0.90` / `0.60`	v0	`out/v0/x1_audit.json`, `nsr2db`
width `chf2db` `25/29` / `18/29`, median IAAFT `z` `5.37` / `3.02`	v0	`out/v0/x1_audit.json`, `chf2db`
width `chfdb` `10/15` / `4/15`; `fantasia` `12/40` / `10/40`	v0	`out/v0/x1_audit.json`
shuffle `z ≈ 4–13` all cohorts; median raw width `nsr2db 0.11`, `chf2db 0.25–0.27`	v0	`out/v0/x1_audit.json`, shuffle / width_obs
scattering `nsr2db` `54/54` both policies; `chf2db` `29/29`	v3	`out/v3/x7_scattering.json`, `cohorts`
healthy dominant family time-asymmetry, median `z_agg` `156` (raw) / `73` (NN), `53/54`	v3	`out/v3/x7_scattering.json`, families
CHF dominant family sparsity / 2nd-order, median `z_agg` `456` / `467` (NN)	v3	`out/v3/x7_scattering.json`, families
specificity FPR `0.010` (`1/96`; fGn `1/64`, AR(4) `0/16`, iid-`t` `0/16`)	v3	`out/v3/x7_scattering.json`, `anchor_battery`
power curve `0.625 → 16/16` (`N` `2,000 → 100,000`)	v3	`out/v3/x7_scattering.json`, `anchor_battery`
converse: `6/83` degenerate MLEs; winner unexplained `51/51` + `26/26`	v3	`out/v3/x7_scattering.json`, converse arm
`Im ρ` exactly odd under time reversal (module anchor)	—	`fenosoma/scattering.py`; `tests/test_scattering.py`
mechanism / pincer / synthetic power-specificity certificate	*external (methods-core)*	Fenocosm `papers/08-width-power.md` (cited, §1.3, §6)
`t_MF` split `51.31%` positive / `27.07%` significant-negative (7000 gens)	*external*	Mangalam et al. 2025, PRE 111, 034126 (arXiv:2312.05653)
`t_MF` formula	*external*	Bell et al. 2019, Front. Physiol. 10:998
prior art `w_orig < w_surr`, bidirectional reading	*external*	Lee & Kelty-Stephen 2017, Complexity 2017:7015243