The irreversible linear null: every surrogate family in the standard toolbox is time-reversible by construction, so no test built on one can fail on a time-asymmetry statistic

A Poisson shot-noise ensemble with an asymmetric two-exponential (FRED) kernel, matched per target on the marginal and on the spectral density and read by the instruments that produced the sealed departures, covers 22/22 Swift/BAT sources and 54/54 healthy-RR records — and is an order of magnitude *more* time-asymmetric than the heartbeat cohort it is matched to.

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

Fenopan standalone note — prose-complete draft (2026-09-15). Sealed record: Fenopan x13, record v20 (`10.5281/zenodo.22759725`, deposited 2026-09-15); the sibling records it reads beside are cited by DOI below. Epistemic tags per CONSTITUTION §0 travel with every number. Public-safe: quotes only sealed measurements and independently verified external sources; carries no strategy. Nothing under `out/v20` was modified in writing this note; the one figure is regenerated from the sealed JSON alone by `tools/make_figure_x13_note.py`, committed with it. This note's own deposit: Zenodo `10.5281/zenodo.22759959` (preprint; DOI reserved 2026-09-15; published 2026-09-15 on the owner's sign-off).

Abstract

The standard surrogate-data toolbox for testing nonlinearity has four workhorse null families — random shuffling, the phase-randomised Fourier surrogate, amplitude-adjusted and iteratively amplitude-adjusted Fourier surrogates (AAFT/IAAFT), and fitted linear models such as AR and ARFIMA. Every one of them is **time-reversible by construction**, so a test built on them cannot fail on a time-asymmetry statistic: a rejection says only *not reversible-linear-Gaussian*, strictly weaker than *nonlinear*, because a linear process can be irreversible. This note reports the missing null family built and raced: **Poisson shot noise through an asymmetric two-exponential (fast-rise/exponential-decay, FRED) kernel** at rise:decay ratios 1:2, 1:5, 1:10 and 1:30, matched per target on the marginal (exact rank remap) and on the spectral density (a zero-phase shaping filter), and read by the two vendored instruments that produced the program's sealed departures. The pre-remap member is exactly a linear filtered-Poisson process with the target's spectral density; irreversibility is the only property the ensemble adds over IAAFT. Under a rule frozen before the run ($K = 39$ members per target and ratio; *covered* iff $p = (1 + \#\{|t_m| \geq |t_{real}|\})/(K + 1) > 0.05$), **22/22 Swift/BAT sources are covered at every one of the four ratios** (union Clopper–Pearson 95 % **[0.846, 1.000]**) and **54/54 healthy-RR records at one or more ratios** (CP95 **[0.934, 1.000]**), with the **median**

coverage p equal to 1.000 at every ratio in both domains. The null does not merely reach the data: on the heartbeat cohort it overshoots by an order of magnitude, the matched members' per-record median asym z_agg running **556 to 1,172** against the cohort's sealed median of **73.1**. A second finding follows from re-deriving the isolated-singularity false-positive budget *on the matched controls themselves* rather than inheriting it: the omnibus flags **1.000** at every ratio in both domains and **every** scattering family flags at or above **0.935** — the cross-scale `C` arm a multiplicative-cascade claim specifically needs included — against sealed budgets of 2/51 and 0.010 measured on smooth linear controls. The consequence is methodological and general: **an excess of a time-asymmetry or wavelet-scattering statistic over IAAFT is not evidence of nonlinearity, multiplicativity or endogeneity until an irreversible *linear* null has been raced.** It is not evidence that any source or any heart *is* shot noise, and it touches no likelihood-based result. Alongside it we report an observation we do not interpret: healthy hearts are *less* time-asymmetric than the linear pulse trains matched to their own marginal and spectrum.

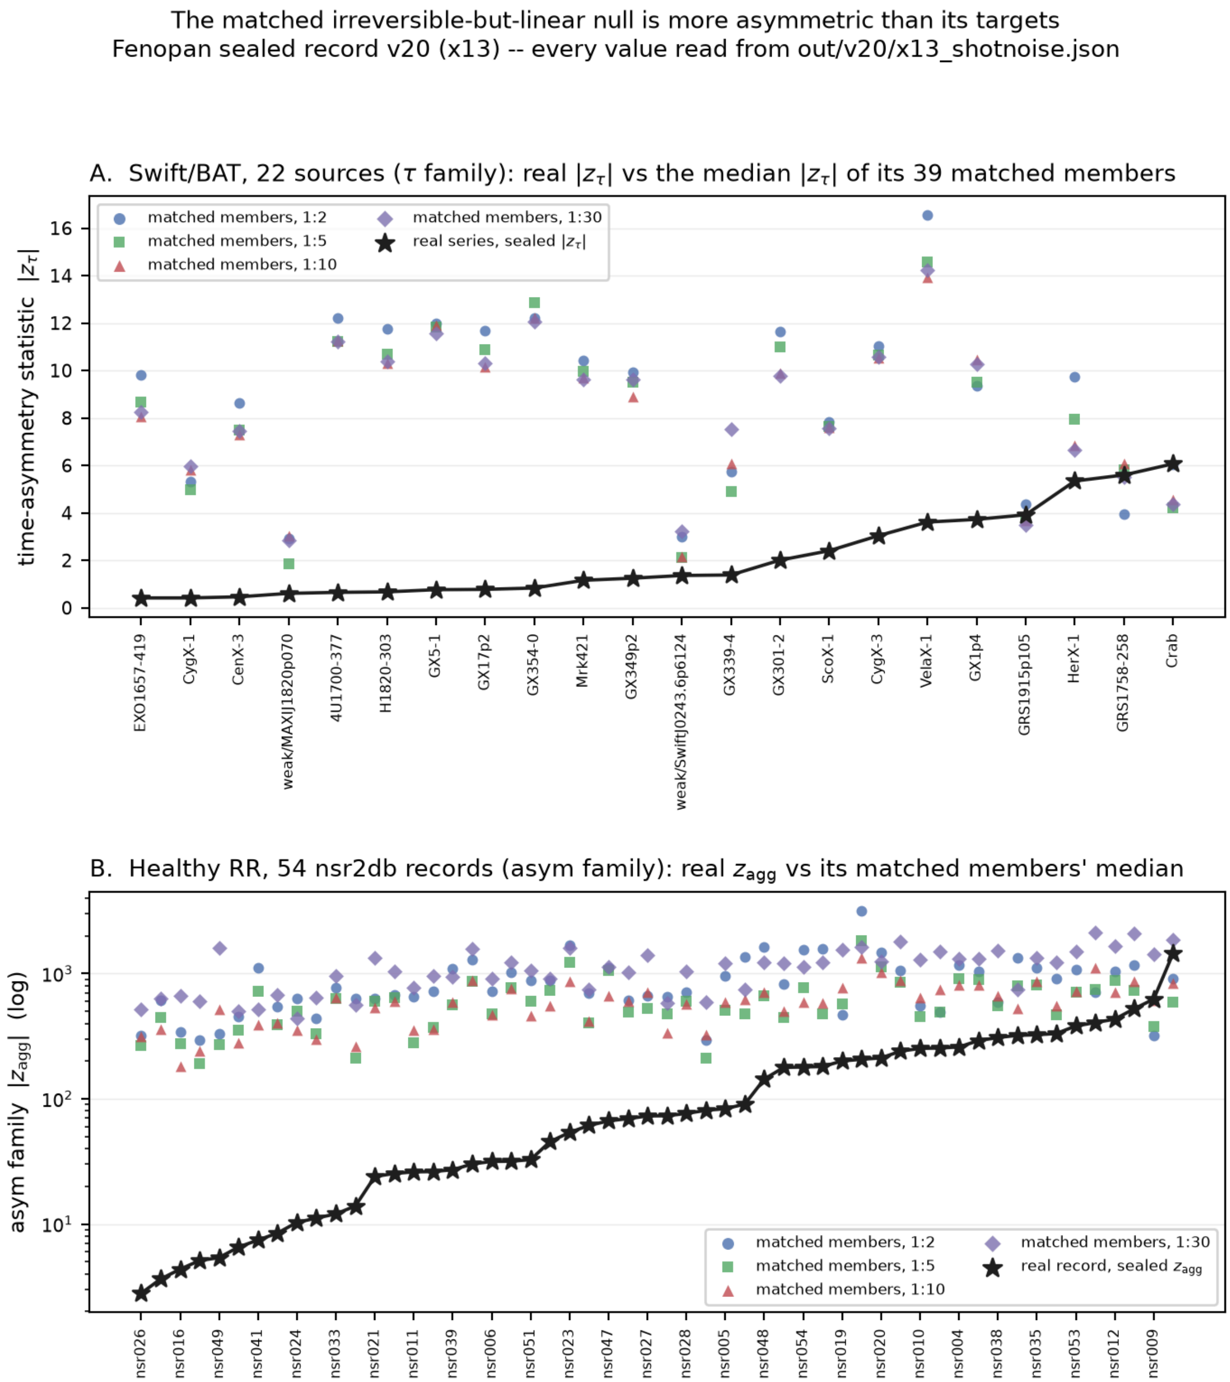

1. A null family that cannot fail

Surrogate data testing has had a stable shape since Theiler, Eubank, Longtin, Galdrikian and Farmer introduced it [VERIFIED — *Physica D* 58(1–4):77–94, 1992]: choose a discriminating statistic, generate an ensemble sharing with the data everything the null permits and nothing it forbids, and ask whether the data's statistic falls outside it. The method's honesty lives in the *construction* of that ensemble, because the null hypothesis is not a sentence in a paper — it is whatever the generator produces.

Four generators do nearly all the work. Random shuffling preserves the marginal and destroys everything else. The phase-randomised Fourier surrogate preserves the power spectrum and Gaussianises the marginal. AAFT and, iteratively refined, IAAFT preserve the marginal *and* the spectrum to a controlled tolerance — the construction of Schreiber and Schmitz [VERIFIED — *Phys. Rev. Lett.* 77(4):635–638, 1996] — and are the default null in most applied work; the multivariate extension is Prichard and Theiler's [VERIFIED — *Phys. Rev. Lett.* 73(7):951–954, 1994]. The fourth family is parametric: fit an AR, ARMA or ARFIMA model and simulate from it.

All four share a property that is rarely stated when a rejection is interpreted. **Each of them produces time-reversible processes.** A shuffle is exchangeable. A phase-randomised surrogate is Gaussian with a specified spectrum, and a stationary Gaussian process is reversible because its law is fixed by a covariance symmetric in the time index. IAAFT converges to a specified amplitude spectrum and a specified marginal from a randomised start; the ensemble carries no systematic arrow. And for the parametric family the result is a theorem: a stationary linear (ARMA) process is time-reversible essentially only in the Gaussian case — Weiss's characterisation [VERIFIED — *J. Appl. Prob.* 12(4):831–836, 1975] — so a *Gaussian* fitted AR or ARFIMA null is reversible, which is exactly the null that is fitted in practice.

The consequence is structural rather than statistical. If the null ensemble is reversible and the statistic is odd under time reversal, the null's expected statistic is zero and its ensemble is symmetric about zero; the test can *reject*, but it cannot **fail** in the sense that matters — it can never be that the null family was capable of producing the observed asymmetry and happened not to. A time-asymmetry excess over any of these four families licenses exactly one sentence: *the data is not reversible-linear-Gaussian*. It does not license *nonlinear*, because the other half of Weiss's characterisation is that a **non-Gaussian linear process is generically irreversible**. The gap between what is measured and what is usually claimed is the subject of this note.

That gap is not hypothetical in our own work. The Feno program's winning model class — a Markov-switching multifractal cascade with detailed-balance switching — is provably time-reversible, so its expected time-asymmetry statistic is an exact zero rather than a surrogate-estimated one; and yet the data it wins or fights on reads measurably irreversible. Fenocosm's scattering-moment null audit finds all **22** Swift/BAT sources departing the IAAFT null, with the time-asymmetry family dominant in **6** of them, GRS 1915+105 among them [REPO — Fenocosm v22, `10.5281/zenodo.22180491`]; Fenosoma's scattering-moment audit finds the healthy heart-rate cohort's excess "dominated by time-irreversibility", asym-dominant on **28/54** records and significant on 53/54 [REPO — Fenosoma v3, `10.5281/zenodo.22180546`]. Every null family the program owned was one of the four above, so no existing control could have failed on τ. The program's public trap registry already held that the two-point multifractal width has *zero power on time asymmetry by construction*, and that width and IAAFT are an incoherent pairing; what it did not hold was the mirror-image trap on the other jaw — a statistic with power to spare, raced against a family that cannot supply the alternative.

2. Linear is not reversible: the object the toolbox is missing

The missing null is not exotic. Take a homogeneous Poisson process of shot times, give each shot an independent random amplitude, and convolve with a fixed kernel h(t). The result

> $x(t) = \Sigma_i a_i \cdot h(t - t_i)$, $t_i \sim \text{Poisson}(\nu)$, $a_i \sim \text{iid}$

is a **linear** process in every sense that matters to the toolbox: a linear filter applied to an iid point-process input, with an exactly specifiable power spectrum, satisfying no nonlinear recursion, and with nothing in it multiplicative, self-exciting or state-dependent. It is also, whenever h is not symmetric, **time-irreversible** — reversing time reverses the kernel, and an ensemble built on h is statistically distinguishable from one built on h(−t). Asymmetry of the pulse is the whole of the arrow; the process is otherwise a superposition.

The asymmetric kernel of choice in astrophysics is the two-exponential

> $h(t) = \exp(-t/\tau_d) - \exp(-t/\tau_r)$, $t \geq 0$, $\tau_r < \tau_d$,

the fast-rise/exponential-decay or **FRED** pulse. It is the standard descriptive shape for gamma-ray burst pulses [VERIFIED — Norris et al., *ApJ* 459:393, 1996], and pulse superposition is the oldest quantitative model of X-ray binary variability, from Terrell's shot-noise reading of Cygnus X-1 [VERIFIED — *ApJ* 174:L35, 1972] to Lochner, Swank and Szymkowiak's phase-portrait fit of the shot parameters two decades later [VERIFIED — *ApJ* 376:295, 1991]. Solar flare hard X-ray time profiles are read the same way, as sequences of individually asymmetric bursts [VERIFIED — Collier et al., *A&A* 671:A79, 2023].

So the object is not a contrived counterexample invented to embarrass a test: it is the *standard model of the variability* in one of the two domains where the program measured irreversibility, and it had never been used as a null. A targeted search of the constellation before registration found no shot-noise or asymmetric-pulse null registered anywhere [REPO — Fenopan x13, record v20].

The time-reversibility testing literature has known for decades that irreversibility is a *different* property from nonlinearity, and has built statistics for it accordingly: Diks, van Houwelingen, Takens and DeGoede's phase-space criterion [VERIFIED — *Phys. Lett. A* 201(2–3):221–228, 1995], Daw, Finney and Kennel's symbolic measure [VERIFIED — *Phys. Rev. E* 62(2):1912–1921, 2000], and Lacasa, Nuñez, Roldán, Parrondo and Luque's visibility-graph divergence [VERIFIED — *Eur. Phys. J. B* 85(6):217, 2012]. Zanin and Papo's two reviews [VERIFIED — *Entropy* 23(11):1474, 2021; *Entropy* 27(2):126, 2025] survey a field rich in *statistics*. What is thin, on our reading, is the matching **null**: the ensembles those statistics are scored against are, once again, shuffles, Fourier surrogates and fitted linear models. The statistic side is well tooled — it now has a unified software package [VERIFIED — Zanin, *Entropy* 27(11):1146, 2025] — while the null side still offers nothing that can be irreversible without also being nonlinear.

3. The construction

Record v20 builds that null and races it. Four choices define it, and all four were frozen in code and printed before any cell ran.

The generator. Per-sample Poisson(ν) shot counts with iid Exp(1) marks, filtered by the FRED kernel at unit peak and truncated at twelve decay times, on a rise:decay grid of **1:2, 1:5, 1:10, 1:30** [`meta.budgets.RATIO_GRID`]. Per target, the decay time τ_d is the target's *own* 1/e correlation lag clipped to [2, 512] samples and the shot rate is ν = 1/τ_d — one shot per decay time, the overlap parameter N_OV = 1 that sets how far from Gaussian, hence how irreversible, the raw superposition is. On the Swift/BAT cohort τ_d runs from the clip floor of 2 days (ten sources; daily hard X-ray log-rates are noise-dominated at short lags) to 512 days on GRS 1915+105; on the heartbeat cohort every record's 1/e lag exceeds the clip, so τ_d is the ceiling of 512 beats on all 54. Each member has a **time-symmetrised twin** — the identical shots through (h(t) + h(−t))/2 — the negative control for the whole apparatus: same linearity, same spectrum family, no arrow.

The matching. A member is matched to its target on the two properties the standard nulls preserve, so that irreversibility is the *only* thing the ensemble adds over IAAFT. The spectral density is matched by a **zero-phase (that is, time-symmetric) linear shaping filter** to the target's log-frequency-smoothed periodogram, with the target's exact amplitude imposed at its line bins; the marginal is then matched **exactly** by rank remapping onto the target's sorted values, which fixes the mean — the mean rate of a rate series — and every other moment. Because the filter is zero-phase, the pre-remap member is *exactly a linear filtered-Poisson process carrying the target's spectral density*. On the real cohorts the members' residual against the smoothed density has median **0.023** (Swift/BAT) and **0.027** (heart rate), and line bins carry at most 0.25 % and 4×10^{-5} of the power [REPO — v20].

The instruments. Neither instrument is ours. Both are the sibling repositories' own code, vendored byte-identically and checksum-pinned: Fenocosm's scattering-moment vector with its symmetric leave-one-out omnibus, whose τ family is the normalised imaginary fine-envelope × coarse-coefficient cross-correlation, scored against the sealed convention's 50-member IAAFT ensemble; and Fenosoma's wavelet phase/modulus Im $\rho(j_1, j_2)$ asym family, scored against its 32-member IAAFT ensemble at the sibling's exact seeds. Their published definitions are Bruna, Mallat, Bacry and Muzy [VERIFIED — *Ann. Statist.* 43(1), 2015] and Morel, Rochette, Leonarduzzi, Bouchaud and Mallat [VERIFIED — *Appl. Comput. Harmon. Anal.* 75:101724, 2025]. Using the siblings' own instruments is what makes this a like-for-like race rather than a re-analysis.

The rule. K = **39** matched members per (target, ratio); every member scored by the same instrument against the *target's* sealed-convention IAAFT ensemble; coverage $p = (1 + \#\{|t_m| \geq |t_{real}|\})/(K + 1)$; **covered iff $p > 0.05$** — with K = 39, at least two members as extreme as the data. Magnitude is primary, since the sealed Swift/BAT τ signs are mixed across sources and the heartbeat statistic is a magnitude; the signed fast-rise/slow-decay orientation is secondary. A target counts as covered if it is covered at **one or more** grid ratios — the union, the reading most generous to the null and so the conservative one for the program's own irreversibility claims. Bearings were pre-registered per domain on that fraction: ≥ 2/3 → *largely covered*; ≤ 1/3 → *largely not covered*, obliging the program to register an irreversible cascade entrant; between → split.

The shared-ensemble step — scoring a member against the target's ensemble rather than its own — is an approximation, and it is measured rather than assumed: 264 Swift/BAT and 16 heartbeat members re-scored against their **own** ensembles agree at Spearman **0.888** and **0.921**, with τ-flag agreement **0.962** and **1.000** [REPO — v20]. The coverage margins below exceed those discrepancies by far.

4. Calibration before the real pass

No coverage number was read until the apparatus had been shown to work on worlds where the answer is known. Four checks ran first; one of them changed the apparatus.

The statistics are exactly odd. Both vendored τ statistics reproduce their reversal identity far inside the registered tolerance of 10^{-9}: Fenosoma's Im ρ at **1.3×10^{-14}**, and Fenocosm's τ(j) — whose upstream smoke test asserts only a loose bound — shown here exactly odd at **1.4×10^{-16}**, because its Morlet masks are real in frequency, so $\psi(-t) = \mathrm{conj}\ \psi(t)$ and the imaginary cross-term flips sign identically [REPO — v20].

Known truth: the generator fires, the twin is quiet. On pure shot noise with no matching, scored against its own IAAFT ensembles, the τ family fires on **8/8** seeds at every ratio on the Swift/BAT instrument (median z_τ rising from +8.54 at 1:2 to **+22.41** at 1:30) and on **6/6** at every ratio on the heartbeat instrument (median asym z_{agg} +2,502 to +3,245) — and stays **quiet on the time-symmetrised twin**, 0–1 flags of 8 and 0 of 6. Irreversibility is what the asymmetric kernel adds; nothing else in the pulse train carries it.

The two-world gate. The coverage machinery needs a world it must cover and a world it must not. **W+** is a target that genuinely *is* FRED shot noise: its matched ensemble at the same ratio covered **31 of 32** Swift/BAT cells and **6/6** heartbeat cells. **W−** is nonlinear irreversibility that no pulse superposition makes — a Poisson-reset relaxation oscillator — covered at any ratio on only **1/8** and **0/3** seeds, inside the registered bar of ≤ 0.25. Beside them sits the **Hénon map**, kept and reported as the *boundary* world: its magnitude sits inside the matched shot-noise band on every seed while its **sign** is covered on **0/8**. That boundary is a measured power limit every "covered" verdict in this note inherits, and §7 states it as such.

The amendment, made before any real number was read. The registration had specified the spectral match as the Schreiber–Schmitz IAAFT loop *seeded with the shot draw* rather than with a shuffle — appealing, since the phases would then descend from the shot process. The W+ world measured that it does not work: the loop **erodes** the members' phase structure. On a FRED target at ratio 1:10 the target reads z_τ **+9.59 / +9.18** on two seeds while the iteratively matched members' median is **+2.80 / +2.14** (and +4.43 / +3.75 even with only three iterations after a smooth shaping), so W+ was **not covered on 4/4 pilot cells**. A null that weak by construction is biased toward the "not covered" bearing — that is, toward flattering the program's own irreversibility claims — so the primary matcher became the zero-phase shaping filter of §3, under which the same cells read members' median +7.41 / +9.00 and are covered while the relaxation W− world stays beyond. The registered loop was kept as a sensitivity cell, and the erosion is anchored as a standing inequality: on the smoke fixture the shaping matcher keeps z_τ **8.55** where the iterative loop keeps **0.58** [REPO — v20]. The run that had started under the registered matcher was stopped inside its battery, before any real cell; nothing from it is sealed.

Two things should be said plainly about that amendment. It is a change of apparatus made after a measurement and before the data, logged with its pilot numbers in the sealed record; the honest reading is that the pre-registration was *wrong about a construction detail and the two-world gate caught it*. And it changed no verdict — as a sensitivity cell the eroded loop still covers **21/22** Swift/BAT sources — only whether the apparatus's negative would have been calibrated at all.

The sealed numbers reproduced before any new number. Read through the verbatim sibling pipelines, the 22 Swift/BAT light curves reproduce the sealed per-source τ statistic on **22/22** to a maximum absolute difference of **0.005** — the sealed table's two-decimal resolution — with the sealed omnibus p and dominant family on 22/22, so the re-derived count of τ-dominant sources is the sealed **6**; the 54 heartbeat records reproduce **byte-for-byte on 54/54** (omnibus statistic to **2.1×10^{-15}** relative), the re-derived count of asym-dominant records being the sealed **28** [REPO — v20].

5. The two real cohorts

Swift/BAT: 22 of 22 sources are covered at every one of the four ratios (union Clopper–Pearson 95 % **[0.846, 1.000]**), with the **median coverage p equal to 1.000 at every ratio** [REPO — v20]. The margins are not marginal: the real cohort's median |z_τ| is **1.31**, while the matched members' per-source median |z_τ| at ratio 1:10 runs from **2.18** (SwiftJ0243.6+6124) to **13.92** (Vela X-1). The smallest margins belong where they should — to the three strongest sealed departures and the two short weak-catalog sources. The Crab, the cohort's largest sealed departure at z_τ = **+6.07**, reads coverage p of 0.475 / 0.150 / 0.175 / **0.075** along the grid, covered at 1:30 by exactly the two members the rule requires; GRS 1758−258 (−5.60) reads 0.275–0.575 and GRS 1915+105 (+3.92) 0.450–0.575. The other seventeen sources sit at p ≥ 0.90 everywhere.

Healthy RR: 54 of 54 records are covered at one or more ratios (union CP95 **[0.934, 1.000]**) — 54/54 at 1:2 and 1:30, **53/54** at 1:5 and 1:10 — again with the **median p equal to 1.000 at every ratio**. The single record that escapes at two ratios is **nsr042**, which is the record with the *largest* sealed asym statistic in the cohort (z_agg **1,454.9**); it is uncovered at 1:5 (one member above, p 0.050) and at 1:10 (none, p 0.025), and covered at 1:2 (p 0.300) and 1:30 (p 0.925).

The heartbeat domain does not merely reach coverage; it **overshoots by an order of magnitude**. The matched members' per-record median asym z_agg, taken as a median over records, is **749 / 556 / 594 / 1,172** at the four ratios against the cohort's sealed median of **73.1** (minimum 2.8, maximum 1,454.9); even the members' per-record *minimum*, again median over records, runs **173 / 266 / 360 / 649**. Figure 1, panel B, is that statement drawn: the black trace is the cohort, the coloured markers its own matched linear nulls, and the nulls sit above the data across essentially the whole cohort.

The signed secondary reading is reported and is not the claim: fast-rise/slow-decay pulses cover **15/22** Swift/BAT sources in sign, which is exactly the 15 sources whose sealed τ is positive. The remaining 7 read, under this instrument, as slow-rise/fast-decay pulse trains — the time reversal of the same ensemble, hence covered in magnitude by construction. Whether pulse orientation can be recovered from a daily-binned hard X-ray light curve at all is a question this record raises and does not answer [gap — registered].

Sensitivity cells, reported and never claim-bearing, move the picture in one place only. At ratio 1:10, N_OV ∈ {0.25, 4} gives 22/22 and 21/22 Swift/BAT and 8/8 of the eight probed heartbeat records; τ_d × 4 gives 22/22 and 8/8; the eroded iterative matcher 21/22. The cell that moves is **τ_d × 0.25**, covering **18/22**: quartering the clip-floor decay time of two days makes the pulses sub-sample (median τ_d 0.75 samples), so the null is nearly white, its members carry a median |z_τ| of only 1.86, and four sources escape it. That is informative rather than adverse — coverage is carried by pulses that are actually resolved, not by an artefact of the matching.

6. The budget re-derived: not only the τ family

Sharp pulses are isolated singularities, and the scattering omnibus has a known high side against those. So this arm's false-positive budget was **re-derived on the matched controls themselves and never inherited** from the sealed budgets, which were measured on smooth linear controls — AR(1) and white noise for the Swift/BAT instrument, fractional Gaussian noise, AR(4) and Student-t for the heartbeat one.

The re-derived budget is the note's second finding, and it is larger than the first. On the matched members the **omnibus flags 1.000 at every ratio in both domains**, and **every family flags at or above 0.935**:

rise:decay	omnibus (both)	BAT: S1 · ρ · S2n · C · τ	RR: sparsity · norm_s2 · env_corr · asym
1:2	1.000	1.000 · 0.974 · 1.000 · 0.993 · 0.951	1.000 · 1.000 · 0.997 · 1.000
1:5	1.000	1.000 · 0.963 · 1.000 · 0.988 · 0.935	1.000 · 1.000 · 0.994 · 1.000

rise:decay	omnibus (both)	BAT: S1 · ρ · S2n · C · τ	RR: sparsity · norm_s2 · env_corr · asym
1:10	1.000	1.000 · 0.967 · 1.000 · 0.993 · 0.951	1.000 · 1.000 · 0.989 · 1.000
1:30	1.000	1.000 · 0.985 · 1.000 · 0.993 · 0.965	1.000 · 1.000 · 0.980 · 1.000

against sealed budgets of **2/51 = 0.039** [REPO — Fenocosm v22 `10.5281/zenodo.22180491`; v33 `10.5281/zenodo.22180499`, whose mask-transported reads are 0.045 and 0.036] and **0.010** [REPO — Fenosoma v3 `10.5281/zenodo.22180546`] — a matched linear pulse train flags the omnibus at roughly 26 times the Swift/BAT budget and 100 times the heartbeat one.

The reading is not that the sealed audits were run incorrectly; they measured specificity honestly against the controls they had. The reading is that **specificity does not transfer between null worlds**. A budget measured on AR(1), white noise and fractional Gaussian noise says what the instrument does on *smooth* linear worlds, and nothing about what it does on a *pulse* world — where, measured here, it flags almost everything. Two consequences follow.

First, the finding widens from the τ family to the whole scattering departure. The sealed 22/22 and 54/54 omnibus departures — including the cross-scale `C` family a multiplicative-cascade claim specifically needs, which flags 0.988–0.993 here — sit *inside* what a matched irreversible pulse train produces. The known-truth battery sharpens this: the **time-symmetrised twin**, a reversible linear filtered-Poisson process, already flags the omnibus on 8/8 and 6/6 seeds, carried by the S2n and C families at flag rate 1.00 with τ quiet. Pulse structure alone, with no arrow at all, fails the scattering omnibus at the sealed lengths.

Second, family *ordering* is a different object from family *magnitude*, and the two come apart here. On the Swift/BAT instrument the matched members are τ-dominant on **0.268–0.290** of draws — the sealed data's 6/22 = 0.273 to the second decimal, a coincidence we report and do not build on. On the heartbeat instrument the members are asym-dominant on only **0.023–0.162** of draws against the cohort's 28/54, because their norm_s2 excess (median 3,956 at 1:10) outruns their asym excess (594) by more than the real records' does (58.9 against 73.1). So "dominated by time-irreversibility" is a statement about which family is *relatively* largest under IAAFT; in absolute terms the cohort's asym excess is an order of magnitude below a matched pulse train's. That ordering difference is handed forward, not read [gap — registered].

7. What the result does not say

Four limits belong in the same breath as the result.

It is not a claim that any source or any heart is shot noise. The test is a coverage test: it asks whether a matched linear irreversible ensemble *can* produce what was observed, not whether it *did*. The Hénon boundary world settles the question of what "covered" is worth: a low-dimensional deterministic nonlinear map, matched the same way, is also covered in magnitude on every seed [REPO — v20]. "Covered" therefore means **not beyond a linear irreversible null in magnitude**, and never "linear". Any mechanistic reading of these coverages would be exactly the inference this note argues against, run backwards.

It says nothing about likelihood-based results. x13 raced no tournament, fit no model and scored no held-out likelihood. The program's likelihood verdicts — which model class wins on held-out data, at what sample size, against which rivals — are untouched by it, as are the sealed features' use as discriminative biomarkers, since a feature needs no mechanism to classify.

The magnitude test has a measured power limit, and it is the Hénon world. The sign reading is more specific — Hénon's sign is covered on 0/8 seeds, and the relaxation oscillator is beyond the null in both readings — but the sign reading is also the one that fails on 7 of 22 Swift/BAT sources, and the sealed signs are mixed, which is why magnitude was pre-registered as primary. A more powerful orientation-aware test is the obvious next instrument.

What survives is not nothing. "Not reversible-linear-Gaussian" remains true and untouched: the sealed departures are real departures from the null that was actually raced. What dies is the step from that sentence to *nonlinear*, *multiplicative* or *endogenous* on the scattering evidence alone. In the program's own canon the pincer of §6.2 — the two-point width reads too high when nothing is there and too low when something is — gains a third jaw of the same kind: **the scattering vector reads nonlinear when a linear irreversible process would do.** The general form of the rule, which is what we propose to the wider literature, is one sentence: *an IAAFT excess is evidence of nonlinearity only after an irreversible linear null has been raced.*

8. The fourth corner, and a deficit left open

The standing classification argument in seismology, market microstructure, epidemiology, neuroscience and crisis studies is a trichotomy: is the clustering **exogenous** (regimes, weather), **endogenous** (self-excitation, dominoes), or **rough drift**? The program's flagship treats it as a metrology problem rather than a classification problem, because the corners provably merge near the boundary, and its sharpest pre-registered conjecture was that **time-irreversibility is the trichotomy's actual discriminator** — the cascade ladder being reversible where self-excitation is not. Fenopan's fine-resolution test of that conjecture returned *uninterpretable at the arena tested*, both failing predicates measured rather than asserted [REPO — Fenopan v19, `10.5281/zenodo.22759721`].

This record adds something the trichotomy did not contain. **Linear superposition of asymmetric pulses is a fourth corner**: irreversible, and so indistinguishable from the endogenous corner by any time-asymmetry statistic raced against a reversible null; linear, and so excluded from the endogenous corner by construction; and not smooth, so not rough drift either. It is not a corner case in the pejorative sense — it is the standard model of X-ray binary variability and of gamma-ray burst and solar flare morphology, which is to say the default description of the very data where the program measured its irreversibility. *[Interpretation: that a fourth corner exists does not make the other three wrong, and nothing here locates any real system in it. The claim is about the discriminating power of a class of tests, not about mechanism.]*

Alongside it we report a direction we deliberately do not interpret. The matched linear pulse trains are **more** time-asymmetric than the healthy hearts they are matched to — persistently, by roughly an order of magnitude in the asym family's aggregate statistic, across all four rise:decay ratios and across essentially the whole cohort. Read literally, healthy human heart-rate variability is *less* time-asymmetric than a linear pulse train carrying its own marginal and its own spectral density. That is a **deficit**, not an excess — the direction opposite to the one every claim in this area is built on. We state it as an open question rather than a finding, for three reasons. It is measured through one instrument's aggregate statistic, and would need a second, independent irreversibility statistic to be believed as a property of the data rather than of the statistic. Its size depends on a generator constant — the overlap N_OV, pinned at 1 — that was matched to nothing, and the N_OV sensitivity arm moves the members' magnitudes. And a deficit against a null unmatched in mechanism has no accepted interpretation; one would first have to ask what class of linear irreversible processes reaches the observed level, which is a different experiment. *[Interpretation: if it survives those checks, the interesting sentence is not "hearts are irreversible" but "hearts are markedly less irreversible than the simplest linear process carrying their marginal and spectrum". We have not earned that sentence here.]* [gap — registered]

9. Positioning

Theiler et al. 1992 [VERIFIED] is the method this note works inside, not against: their paper is explicit that a surrogate test tests a *specific* null, and the discipline they ask for — state the null the ensemble actually embodies — is exactly what is missing when an IAAFT rejection is reported as nonlinearity. Our contribution is to supply a null family their framework has room for and the practice does not contain.

Schreiber & Schmitz 1996 [VERIFIED] is the null actually raced in both sealed sibling audits, and it is used here twice over — as the ensemble the real targets were scored against, and as the ensemble the matched members are scored against. The irony is that IAAFT's great virtue, preserving the marginal *and* the spectrum so that a rejection can be blamed on neither, is precisely what leaves the arrow of time unconstrained and unsupplied. **Prichard & Theiler 1994** [VERIFIED] extends the construction to several simultaneously measured variables; the reversibility property extends with it, and so does the gap.

The time-reversibility test literature — Diks et al. 1995, Daw, Finney & Kennel 2000, Lacasa et al. 2012, and Zanin & Papo's two reviews (2021, 2025) [all VERIFIED] — has, if anything, been clearer than surrogate-testing practice that irreversibility is its own property with its own statistics. Zanin & Papo 2025's multiscale review concludes that the tests "yield conflicting results on the same data" and that no single solution dominates [VERIFIED — abstract fetched 2026-09-15]; that is a statement about the *statistic* axis. The result here is on the orthogonal axis: however good the statistic, scoring it against a reversible ensemble makes its rejection uninformative about nonlinearity, because an irreversible *linear* alternative was never on the table. A matched irreversible-but-linear ensemble would be a natural addition to exactly the kind of package that now makes the statistic side easy to use, and the construction in §3 is small enough to be one.

The FRED and shot-noise literature — Norris et al. 1996 for gamma-ray burst pulses, Terrell 1972 and Lochner, Swank & Szymkowiak 1991 for X-ray binary shot-noise modelling, Collier et al. 2023 for solar flare hard X-ray profiles [all VERIFIED] — supplies the null rather than the target, and the relationship is unusual: the model this note uses as a *control* is, in that literature, a *hypothesis about the mechanism*. Our use of it is agnostic about that hypothesis — we need only that the ensemble be linear, irreversible and matched — and that the astrophysics independently considers such processes plausible is what makes the control fair rather than a strawman.

Bruna, Mallat, Bacry & Muzy 2015 and **Morel et al. 2025** [both VERIFIED] are the published definitions of the instruments, vendored here verbatim. Nothing here is a criticism of the scattering representation, which does exactly what it is designed to do — detect structure a marginal-and-spectrum null does not carry. The pulse world *is* such structure. The criticism is of the inferential step downstream of the flag.

10. Reproducibility

Every number quoted above is read back out of one sealed JSON file, `out/v20/x13_shotnoise.json`, and the Appendix binds each headline to its key in that file. The record is deterministic: all seeds derive from a single base (`meta.seed_base` = 52,000,000) partitioned into disjoint sub-blocks, and the consumed sibling seed conventions are what make the sealed ensembles reproduce bit-for-bit. The sealed run took **114.8 minutes at four workers** [`meta.runtime_sec`] under Python 3.11.15, numpy 2.4.6 and scipy 1.17.1 [`meta.environment`]. The record carries its own regeneration block, ending in the record-level prose gate that binds its narrative to its own JSON.

The one figure is regenerated by `tools/make_figure_x13_note.py`, committed alongside this note: it reads the sealed JSON and nothing else, writes only `notes/figures/fig-irreversible-linear-null.png`, draws no random number and touches no network. The PDF is built by `tools/build_pdf.py`, to which the note is registered as a manifest entry.

One reproducibility caveat is a property of the *data* rather than of the code: the Swift/BAT transient monitor updates daily in principle, but its products were frozen at MJD 61112 at both the sibling seal and this run, which is why the leading-n truncation and the current snapshot coincide on 22/22 sources. A later reprocessing of the archive would move that anchor; the sealed per-source length is the pin.

11. Data availability

Two external cohorts are read, and **neither is redistributed**. The 22 Swift/BAT daily light curves come from the Swift/BAT Hard X-ray Transient Monitor [VERIFIED — Krimm et al., *ApJS* 209(1):14, 2013], fetched through the sibling repository's own fetch machinery; the acknowledgment travels with the record. The 54 healthy heart-rate records are the PhysioNet normal-sinus-rhythm beat-annotation files [VERIFIED — Goldberger et al., *Circulation* 101(23):e215, 2000; nsr2db v1.0.0, ODC-By], fetched through the sibling's own machinery and cross-checked against this repository's own 54 SHA-256 pins, 54/54. Per-file checksums, URLs and fetch times are sealed in `meta.data_manifests`; the fetched data itself is never committed.

The sealed record this note quotes is **Fenopan v20** — `README.md`, `RESULTS.md`, `x13_shotnoise.json` and its figure — deposited as a standalone Zenodo dataset at **`10.5281/zenodo.22759725`** (2026-09-15), carrying its own regeneration block so that every number here can be reconstructed from the deposited object without repository access. The sibling records it reads *beside*, never over, are **Fenocosm v22** (`10.5281/zenodo.22180491`) and **Fenocosm v33** (`10.5281/zenodo.22180499`) for the Swift/BAT instrument, its sealed per-source statistics and its specificity budget; **Fenocosm v7** (`10.5281/zenodo.22180487`) for the log-rate concatenate-and-censor convention the Swift/BAT loader reimplements; **Fenosoma v3** (`10.5281/zenodo.22180546`) for the heartbeat instrument and its sealed rows; and **Fenopan v19** (`10.5281/zenodo.22759721`) for the fine-resolution reading of §8. Nothing under `out/v0`–`out/v19` was touched by the record, and no sibling repository was modified by it.

This note is itself deposited as a standalone Zenodo preprint at **`10.5281/zenodo.22759959`** (Markdown, PDF and figure; DOI reserved 2026-09-15; published 2026-09-15 on the owner's sign-off), wired to v20 and the sibling records above by `isSupplementTo`/`cites` relations in its metadata.

The repository itself is private and remains so; requests for the apparatus should be addressed to the author.

Appendix A — provenance: every headline number and its sealed JSON key

Every quantitative claim in this note traces to `out/v20/x13_shotnoise.json` (Fenopan record v20, `10.5281/zenodo.22759725`) or to a named sibling record. "Reduce" entries are computed from the per-item values rather than from a summary field, which is the house rule for a headline: bind to the items, never only to the summary. Paths use dotted notation, `*` for a collected list, `[cell=grid]` for the grid cells of a domain block, and `[Crab]` / `[nsr042]` for the list entry with that `slug` or `record`.

#	headline number	value	sealed JSON key
1	BAT sources covered at ≥	22 of 22	`verdict.domains.bat.n_covered_any_ratio` / `.n_admitted`

#	headline number	value	sealed JSON key
	1 ratio		
2	BAT union Clopper–Pearson 95 %	[0.846, 1.000]	`verdict.domains.bat.ci95`
3	BAT covered per ratio	22 / 22 / 22 / 22	`bat.summary.per_ratio.{1:2,1:5,1:10,1:30}.n_covered`
4	BAT median coverage p per ratio	1.000 at every ratio	`bat.summary.per_ratio.*.median_p`
5	RR records covered at ≥ 1 ratio	54 of 54	`verdict.domains.rr.n_covered_any_ratio` / `.n_admitted`
6	RR union Clopper–Pearson 95 %	[0.934, 1.000]	`verdict.domains.rr.ci95`
7	RR covered per ratio	54 / 53 / 53 / 54	`rr.summary.per_ratio.{1:2,1:5,1:10,1:30}.n_covered`
8	RR median coverage p per ratio	1.000 at every ratio	`rr.summary.per_ratio.*.median_p`
9	Bearing, per domain	(a) largely covered, both	`verdict.domains.{bat,rr}.bearing`
10	Members per (target, ratio)	K = 39	`meta.budgets.N_MEMBERS`
11	Coverage bar	$p > 0.05$	`meta.budgets.COVER_P` ; `{bat,rr}.summary.thresholds.covered_iff_p_gt`
12	Rise:decay grid	1:2, 1:5, 1:10, 1:30	`meta.budgets.RATIO_GRID`
13	Decay-time clip; overlap	[2, 512] samples; N_OV = 1	`meta.budgets.TAU_D_CLIP` ; `meta.budgets.N_OV`
14	RR ensemble decay time	$\tau_d = 512$ beats (all 54)	`rr.tau_d_median`
15	Cohort sealed	73.1 (min 2.8, max	reduce: median of `rr.targets[*].asym_z_agg`

#	headline number	value	sealed JSON key
	median asym z_agg	1,454.9)	
16	Members' median asym z_agg, per ratio	749 / 556 / 594 / 1,172	reduce: median over records of `rr.cells[cell=grid, ratio=r].member_fam_stat_median.asym`
17	Members' per-record minimum asym z_agg	173 / 266 / 360 / 649	reduce: median over records of min| `rr.cells[cell=grid, ratio=r].member_tau_stat` |
18	RR members' norm_s2 vs asym at 1:10	3,956 vs 594	reduce: median of `rr.cells[cell=grid, 1:10].member_fam_stat_median.{norm_s2,asym}`
19	Cohort sealed norm_s2 vs asym median	58.9 vs 73.1	reduce: median of `rr.targets[*].sealed.families.{norm_s2,asym}.z_agg`
20	BAT real median |z_τ|	1.31	reduce: median of | `bat.targets[*].kinds.trunc.tau_z` |
21	BAT members' median |z_τ| at 1:10, range over sources	2.18 to 13.92	reduce: median of | `bat.cells[cell=grid, 1:10].member_tau_stat` |
22	Crab: sealed z_τ; coverage p at 1:30	+6.07; 0.075	`bat.targets[Crab].sealed.z_tau`; `bat.cells[Crab, 1:30].coverage.p`
23	nsr042: sealed asym z_agg; p at 1:5, 1:10	1,454.9; 0.050, 0.025	`rr.targets[nsr042].asym_z_agg`; `rr.cells[nsr042, {1:5,1:10}].coverage.p`
24	Signed (fast-rise/slow-decay) coverage	15 of 22, every ratio and union	reduce: count of `bat.cells[cell=grid].coverage_signed.covered`

#	headline number	value	sealed JSON key
25	Omnibus flag rate on matched controls	1.000, every ratio, both domains	`{bat,rr}.summary.per_ratio.*.omnibus_flag_rate_mean`
26	BAT family flag rates on matched controls	S1 1.000 · ρ 0.963–0.985 · S2n 1.000 · C 0.988–0.993 · τ 0.935–0.965	`bat.summary.per_ratio.*.fam_flag_rate_mean.{S1,rho,S2n,C,tau}`
27	RR family flag rates on matched controls	sparsity 1.000 · norm_s2 1.000 · env_corr 0.980–0.997 · asym 1.000	`rr.summary.per_ratio.*.fam_flag_rate_mean.{sparsity,norm_s2,env_corr,asym}`
28	Sealed specificity budgets quoted beside them	2/51 = 0.039; 0.010	[CITED] Fenocosm v22 / v33; Fenosoma v3 — no value imported into any computation
29	BAT members τ-dominant fraction	0.268–0.290 (sealed data 6/22 = 0.273)	`bat.summary.per_ratio.*.tau_dominant_member_frac`
30	RR members asym-dominant fraction	0.023–0.162 (sealed data 28/54)	`rr.summary.per_ratio.*.tau_dominant_member_frac`
31	Re-derived dominant counts	BAT τ-dominant 6; RR asym-dominant 28	`bat.dominance.rederived_tau_dominant`; `rr.dominance.rederived_asym_dominant`
32	BAT sealed z_τ reproduced	22 of 22; max |Δ| 0.005	`bat.n_trunc_reproduces_sealed`; `bat.concordance_trunc.max_abs_dz`
33	RR sealed rows reproduced	54 of 54; max rel. diff 2.1×10^{-15}	`rr.n_reproduces_sealed`; `rr.max_rel_dT_obs`

#	headline number	value	sealed JSON key
34	Reversal-oddness anchors	1.3×10^{-14} (Fenosoma); 1.4×10^{-16} (Fenocosm)	`anchors.fenosoma_reversal_odd`; `anchors.fenocosm_tau_odd`
35	Erosion inequality (shaping vs iterative)	z_τ 8.55 vs 0.58	`anchors.erosion_z_shape`; `anchors.erosion_z_iter`
36	Known truth, BAT: FRED fires; median z_τ	8/8 at every ratio; +8.54 → +22.41	`battery.known_truth.bat.per_ratio.*.fred.{tau_flag_count,tau_stat_median}`
37	Known truth, BAT: twin quiet	0–1 flags of 8	`battery.known_truth.bat.per_ratio.*.twin.tau_flag_count`
38	Known truth, RR: FRED fires; median z_agg	6/6 at every ratio; +2,502 → +3,245	`battery.known_truth.rr.per_ratio.*.fred.{tau_flag_count,tau_stat_median}`
39	Known truth, RR: twin quiet	0 flags of 6	`battery.known_truth.rr.per_ratio.*.twin.tau_flag_count`
40	Twin (reversible) omnibus flag rate	1.00 on 8/8 and 6/6; S2n and C at 1.00	`battery.known_truth.{bat,rr}.per_ratio.*.twin.{omnibus_flag_rate,fam_flag_rate}`
41	W+ (target is FRED) covered	31 of 32 BAT cells; 6 of 6 RR	`battery.two_worlds.bat.W_plus.*.n_covered`; `battery.two_worlds.rr.W_plus.*.n_covered`
42	W− relaxation oscillator covered	1 of 8 BAT seeds; 0 of 3 RR	`battery.two_worlds.{bat,rr}.W_minus.relax.n_covered_any_ratio`
43	Hénon boundary world	covered 8/8 in magnitude; 0/8 in sign	`battery.two_worlds.bat.W_minus.henon.n_covered_any_ratio`; `.frac_covered_signed_any_ratio`
44	Amendment pilot: target	z_τ +9.59/+9.18 vs median	`meta.registration.amendments[0].reason`

#	headline number	value	sealed JSON key
	vs eroded members	+2.80/+2.14; W+ 0/4	
45	Amendment: shaping matcher on the same cells	members' median +7.41/+9.00, covered	`meta.registration.amendments[0].reason`
46	Spectral-match residual vs smoothed density	median 0.023 (BAT), 0.027 (RR)	reduce: median of `{bat,rr}.cells[cell=grid].budget.resid_median`
47	Line-bin power fraction	≤ 0.25 % (BAT); ≤ 4×10^{-5} (RR)	reduce: max of `{bat,rr}.cells[cell=grid].budget.line_power_frac`
48	Shared-vs-own ensemble validity	Spearman 0.888 / 0.921; τ-flag agreement 0.962 / 1.000	`{bat,rr}.validity.{spearman_shared_vs_own,tau_flag_agreement}`
49	Sensitivity, BAT at 1:10	N_OV 0.25 → 22/22; N_OV 4 → 21/22; τ_d×4 → 22/22; τ_d×0.25 → 18/22 (members' median |z_τ| 1.86); iterative matcher → 21/22	`bat.sensitivity.{n_ov=0.25,n_ov=4,tau_d_x4,tau_d_x0.25,iaaft_iter}.{n_covered,member_tau_stat_abs_median}`
50	Seed base; sealed runtime	52,000,000; 6,890 s = 114.8 min at 4 workers	`meta.seed_base`; `meta.runtime_sec`; `meta.workers`
51	Instrument ensembles	50 IAAFT surrogates	`meta.budgets.BAT_N_SURR`; `meta.budgets.RR_N_SUR`

#	headline number	value	sealed JSON key
		(BAT); 32 (RR)	

Numbers cited from sibling records and *not* recomputed here carry their own record: the sealed per-source τ statistics, omnibus p and dominant families of Fenocosm v22 (reproduced in v20 at item 32), the sealed per-record asym statistics and accounting of Fenosoma v3 (item 33), and the sealed specificity budgets of Fenocosm v22/v33 and Fenosoma v3 (item 28, quoted beside the re-derived budget and never reused in it).

References

- Bruna, J., Mallat, S., Bacry, E. & Muzy, J.-F. (2015). Intermittent process analysis with scattering moments. *The Annals of Statistics* 43(1). doi:10.1214/14-AOS1276 [VERIFIED — Crossref, fetched 2026-09-15]
- Collier, H., Hayes, L. A., Battaglia, A. F., Harra, L. K. & Krucker, S. (2023). Characterising fast-time variations in the hard X-ray time profiles of solar flares using Solar Orbiter's STIX. *Astronomy & Astrophysics* 671, A79. doi:10.1051/0004-6361/202245293 [VERIFIED — Crossref, fetched 2026-09-15]
- Daw, C. S., Finney, C. E. A. & Kennel, M. B. (2000). Symbolic approach for measuring temporal "irreversibility". *Physical Review E* 62(2), 1912–1921. doi:10.1103/PhysRevE.62.1912 [VERIFIED — Crossref, fetched 2026-09-15]
- Diks, C., van Houwelingen, J. C., Takens, F. & DeGoede, J. (1995). Reversibility as a criterion for discriminating time series. *Physics Letters A* 201(2–3), 221–228. doi:10.1016/0375-9601(95)00239-Y [VERIFIED — Crossref, fetched 2026-09-15]
- Goldberger, A. L., Amaral, L. A. N., Glass, L., Hausdorff, J. M., Ivanov, P. Ch., Mark, R. G., Mietus, J. E., Moody, G. B., Peng, C.-K. & Stanley, H. E. (2000). PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals. *Circulation* 101(23), e215. doi:10.1161/01.CIR.101.23.e215 [VERIFIED — Crossref, fetched 2026-09-15]
- Krimm, H. A., Holland, S. T., Corbet, R. H. D., Pearlman, A. B., Romano, P., Kennea, J. A., Bloom, J. S. et al. (2013). The Swift/BAT Hard X-ray Transient Monitor. *The Astrophysical Journal Supplement Series* 209(1), 14. doi:10.1088/0067-0049/209/1/14 [VERIFIED — Crossref, fetched 2026-09-15]
- Lacasa, L., Nuñez, A., Roldán, É., Parrondo, J. M. R. & Luque, B. (2012). Time series irreversibility: a visibility graph approach. *The European Physical Journal B* 85(6), 217. doi:10.1140/epjb/e2012-20809-8 [VERIFIED — Crossref, fetched 2026-09-15]
- Lochner, J. C., Swank, J. H. & Szymkowiak, A. E. (1991). Shot model parameters for Cygnus X-1 through phase portrait fitting. *The Astrophysical Journal* 376,
- doi:10.1086/170280 [VERIFIED — Crossref, fetched 2026-09-15]

- 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 [VERIFIED — Crossref, fetched 2026-09-15]
- Norris, J. P., Nemiroff, R. J., Bonnell, J. T., Scargle, J. D., Kouveliotou, C., Paciesas, W. S., Meegan, C. A. & Fishman, G. J. (1996). Attributes of Pulses in Long Bright Gamma-Ray Bursts. *The Astrophysical Journal* 459, 393. doi:10.1086/176902 [VERIFIED — Crossref, fetched 2026-09-15]
- Prichard, D. & Theiler, J. (1994). Generating surrogate data for time series with several simultaneously measured variables. *Physical Review Letters* 73(7), 951–954. doi:10.1103/PhysRevLett.73.951 [VERIFIED — Crossref, fetched 2026-09-15]
- Schreiber, T. & Schmitz, A. (1996). Improved Surrogate Data for Nonlinearity Tests. *Physical Review Letters* 77(4), 635–638. doi:10.1103/PhysRevLett.77.635 [VERIFIED — Crossref, fetched 2026-09-15]
- Terrell, N. J., Jr. (1972). Shot-Noise Character of Cygnus X-1 Pulsations. *The Astrophysical Journal* 174, L35. doi:10.1086/180944 [VERIFIED — Crossref, fetched 2026-09-15]
- Theiler, J., Eubank, S., Longtin, A., Galdrikian, B. & Farmer, J. D. (1992). Testing for nonlinearity in time series: the method of surrogate data. *Physica D: Nonlinear Phenomena* 58(1–4), 77–94. doi:10.1016/0167-2789(92)90102-S [VERIFIED — Crossref, fetched 2026-09-15]
- Weiss, G. (1975). Time-reversibility of linear stochastic processes. *Journal of Applied Probability* 12(4), 831–836. doi:10.1017/S0021900200048804 [VERIFIED — Crossref, fetched 2026-09-15]
- Zanin, M. (2025). irreversibility: A Python Package for Assessing and Manipulating the Time Irreversibility of Real-World Time Series. *Entropy* 27(11), 1146. doi:10.3390/e27111146 [VERIFIED — Crossref, fetched 2026-09-15]
- Zanin, M. & Papo, D. (2021). Algorithmic Approaches for Assessing Irreversibility in Time Series: Review and Comparison. *Entropy* 23(11), 1474. doi:10.3390/e23111474 [VERIFIED — Crossref, fetched 2026-09-15]
- Zanin, M. & Papo, D. (2025). Algorithmic Approaches for Assessing Multiscale Irreversibility in Time Series: Review and Comparison. *Entropy* 27(2), 126. doi:10.3390/e27020126 [VERIFIED — Crossref, fetched 2026-09-15]

Sealed records (Feno program; deposited standalone Zenodo datasets, author Evan Tabak Atlas, ORCID 0009-0007-7374-2338).

- Fenopan sealed record **v20** — x13, the irreversible-but-linear null for the program's time-asymmetry findings (2026-09-11). doi:10.5281/zenodo.22759725 (deposited 2026-09-15). [REPO] — the basis for every measured number in this note.
- Fenopan sealed record **v19** — x10, T4 on the ant colonies at fine resolution. doi:10.5281/zenodo.22759721 [REPO]
- Fenocosm sealed record **v22** — the scattering-moment null audit (P-3.4). doi:10.5281/zenodo.22180491 [REPO]
- Fenocosm sealed record **v33** — the mask-transported monofractal specificity control (P-3.5). doi:10.5281/zenodo.22180499 [REPO]

- Fenocosm sealed record **v7** — the null audit (P-3.2). doi:10.5281/zenodo.22180487 [REPO]
- Fenosoma sealed record **v3** — the scattering-moment audit (2026-08-19). doi:10.5281/zenodo.22180546 [REPO]

(All six deposit records were resolved and their metadata checked against the DataCite API on 2026-09-15.)