A hierarchical regime cascade versus the damped random walk on 21 years of GRS 1915+105 and Cyg X-1

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

Draft — result paper. Companion to the methods paper `papers/01-correspondence-metrology.md`*. Target venue: MNRAS or ApJ; arXiv first (astro-ph.HE). Every quantitative claim traces to a named, immutable results record under* `out/` *(records v0–v5); figures regenerate with* `python -m fenocosm.paper_figures`*. Results provided by the Swift/BAT Hard X-ray Transient Monitor team (Krimm et al. 2013).*

Abstract

The damped random walk (DRW; a single-timescale, linear-Gaussian Ornstein–Uhlenbeck process) is the field-default fitted model for accreting- compact-object light curves. It captures red noise with one break frequency but is additive: it produces neither the lognormal flux distribution nor the rms–flux relation that characterise accretion variability. We ask whether a **hierarchy** of regime timescales — a multiplicative Markov-switching multifractal (MSM) cascade with an exact 2^k-state likelihood — is a better description than a single regime switch or a single-timescale linear process, on two touchstone sources: **GRS 1915+105**, famous for its zoo of discrete variability states, and **Cyg X-1**, the clean two-state source.

Using ~21 years of Swift/BAT hard X-ray (15–50 keV) daily light curves, we run a model tournament with a common, *fair* observation-noise treatment: every model carries the per-point Swift/BAT measurement errors. The entries span iid, AR(1), DRW, DRW+per-point-errors, CARMA(2,1) / OU-mixtures, and the MSM cascade. We find: (i) within the switching family a **hierarchy of timescales decisively beats a single switch** on both sources; (ii) on **GRS 1915+105 the cascade beats every linear-Gaussian competitor**, including the best two-timescale CARMA(2,1), by ΔBIC ≈ +840 with fewer parameters, and the DRW-null is rejected by a parametric bootstrap at $p < 0.01$ at full power; (iii) on **Cyg X-1 the contest is a genuine near-tie** — the fair DRW wins at cascade depth k = 6, but the cascade overtakes it once the ladder is deep enough (k ≥ 8) and by the 4-day rebin, and its fitted ladder collapses to a single timescale, the honest signature of a two-state source. The verdict is **robust across a factor-16 span of cadence** (its limit is sub-daily, where per-bin noise dominates), and the cascade recovers a consistent slowest timescale of a few hundred days from daily to orbital sampling. The cascade's advantage scales with how genuinely multi-timescale the source is — largest on GRS 1915+105, marginal on Cyg X-1 — which is exactly the physically honest boundary.

---

1. Introduction

Accreting black-hole and neutron-star binaries vary aperiodically across a wide range of timescales. Two sources anchor the extremes of variability complexity:

- **GRS 1915+105** is the archetype of *multi-state* variability. Belloni et al. (2000) classified its light curve into a dozen distinct variability classes, transitions between a few underlying spectral states on timescales from seconds to days. If any source is generated by a *hierarchy* of regimes, it is this one.

- **Cyg X-1** is the archetype of *two-state* variability: a hard state and a soft state, with transitions between them. Its aperiodic variability is the textbook target for a single-break red-noise model.

The default *estimable* model applied to such light curves is the damped random walk (Kelly, Bechtold & Siemiginowska 2009): a continuous-time AR(1) process with a single characteristic timescale and an exact Kalman likelihood. The DRW is linear and Gaussian; by construction it has no lognormal marginal and no rms–flux relation (the multiplicative fingerprints; see the companion methods paper). The question of this paper is whether replacing the single break of the DRW with a **hierarchy of regime timescales** — a multiplicative cascade — is a statistically preferred description, and if so, on which sources and by how much.

We use the MSM-Cox cascade of the companion paper: *k* independent two-state layers on a geometric switching-rate ladder, whose product is the (log-) intensity, with an exact 2^k-state forward-algorithm likelihood. Because it is a proper state-space model, we can put it into a like-for-like BIC + held-out tournament against the linear-Gaussian field.

This paper reports a five-stage hardening of a founding result (the founding run found the cascade beating the DRW by ΔBIC ≈ 1,157 on GRS 1915+105). Each stage is an immutable record: fair per-point errors (v1), deeper ladders and bootstrap significance (v2), stronger linear competitors (v3), the population context (v4), and cadence robustness (v5). The founding headline is superseded — halved by the fair noise model, then shown significant and conservative — and the surviving verdict is the subject here.

---

2. Data

We use the **Swift/BAT Hard X-ray Transient Monitor** (Krimm et al. 2013), 15–50 keV, daily-binned light curves (columns: MJD, rate, error, data flag). The two bright persistent sources of the tournament, over ~7,700 days (≈ 21 years):

Table 1 — The bright pair (record `out/v0/`, F4).

source	days	coverage	mean rate (ct cm^{-2} s^{-1})	S/N	span
Cyg X-1	7,697	0.92	0.129	23.8	≈21 yr
GRS 1915+105	7,697	0.91	0.043	13.4	≈21 yr

Cyg X-1 additionally has 95,211 orbital-cadence points used for the sub-daily cadence test (§7). Two fainter sources (GX 339-4, Mrk 421) are used only for the signature battery, not the tournament.

Masked days (data-flag ≠ 0, or gaps) are treated as uninformative throughout; every model is scored on the same informative-tick set. The late-mission BAT noise increase (post-2023) inflates the per-point errors for all models equally.

---

3. Models and tournament design

3.1 The entries

All models are fit on the daily **log-flux** with a masked likelihood.

- **iid Gaussian** (2 params) — no temporal structure; the floor.
- **AR(1)** (3) — a single short decorrelation timescale, no observation noise.
- **DRW + noise** (4) — Ornstein–Uhlenbeck + a scalar observation-noise term; the field default (Kelly et al. 2009), a single break timescale.
- **DRW + per-point errors** (4) — the DRW with the *actual* Swift/BAT per-point measurement errors as the observation-noise term (the fair linear baseline; §3.2).
- **CARMA(2,1) / OU-mixtures** (6 / 8) — the field's *stronger* linear-Gaussian models: a damped stochastic oscillator + first-order MA (equivalently a two-component Ornstein–Uhlenbeck mixture), and a three-timescale OU mixture, each with per-point errors. Two independent implementations of the two-timescale model (CARMA(2,1) and OU-mixture-2) agree to the digit, a cross-check.
- **MSM cascade, depth k** (5) — the multiplicative regime cascade with a **heteroscedastic lognormal emission**: per-tick log-flux variance `σ_intrinsic² + σ_meas,t²`, where `σ_meas,t = ERROR_t/RATE_t` is the Swift/BAT error propagated into log space (delta method). Free parameters `(m_hi, f_slow, S, μ, σ_intrinsic)`; κ fixed = 2 for identifiability. The exact 2^k-state forward pass is Kronecker-factored so the scan can reach k = 9.

Model selection is by **BIC** (effective sample size = informative-tick count for every model) plus an out-of-sample **held-out log-likelihood** on a contiguous final 30% of each series. Parameter counts are carried in every table.

3.2 Fairness: the measurement errors go to *both* families

The founding tournament lumped intrinsic scatter and measurement error into a single constant σ for the cascade, while the DRW carried an explicit observation- noise term. Before any model comparison we measured the effect of the per-point errors *alone*: giving the **DRW** the actual Swift/BAT per-point errors drops its BIC by ~1,300 (Cyg X-1) and ~6,400 (GRS 1915+105) (record `out/v1/`). The noise model, not the dynamics, is the dominant term in the BIC landscape. A heteroscedastic-cascade-vs-scalar-DRW fight would therefore have manufactured a spurious cascade blowout that is really just the noise model. **The fair tournament gives every model the same per-point errors** — the measurement outranks the convenience of an unchanged baseline. This roughly halves the founding margin (see §4) and is the single most important methodological choice in the paper.

---

4. Results I: the tournament

(Records `out/v1/`, `out/v3/`; Figure `figures/fig_tournament.png`

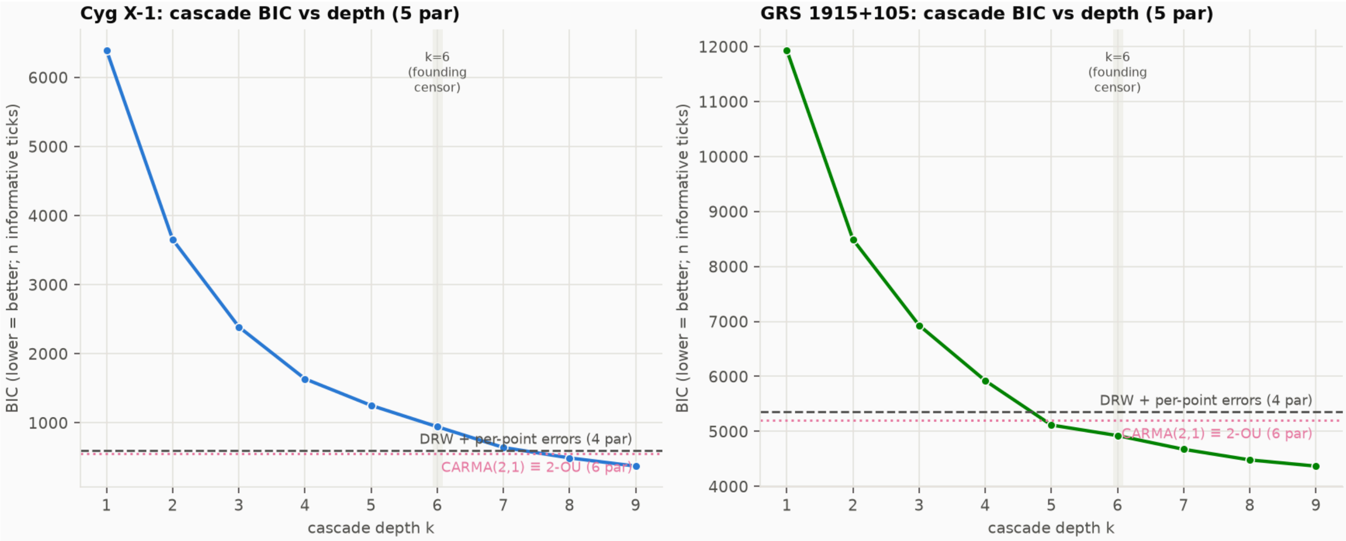

figures/fig_tournament.png

.)

4.1 Signatures on the real data

Both bright sources show the accretion fingerprints, and — decisively — their daily log-flux distributions are **not single-lognormal but left-skewed** (skew(log-flux) = **−1.16** for Cyg X-1, **−1.18** for GRS 1915+105): the textbook signature of a *discrete multi-state (hard/soft) mixture*, i.e. direct evidence for regime structure on these timescales, which is what a cascade models. Persistence is strongly red (DFA α ≈ 1.19 Cyg, 1.22 GRS). The rms–flux relation is present but weak on the daily bright pair (slope 0.09 Cyg, 0.14 GRS; binned r 0.64 / 0.72); it is far stronger on the transient GX 339-4 (slope 0.59, rank-r 0.73) and on Cyg X-1's 95k-point *orbital* light curve (slope 0.183, binned r 0.761). The signatures motivate the cascade; the tournament tests it.

4.2 The tournament table

Table 2 — The fair tournament (bright pair, daily log-flux, sorted by BIC; records v1 & v3). BIC lower is better; held-out log-likelihood higher is better.

model	params	Cyg X-1 BIC	Cyg held-out	GRS 1915+105 BIC	GRS held-out
iid Gaussian	2	16,299	−2,457	23,139	−7,462
AR(1) (τ ≈ 11 d)	3	3,687	−865	13,352	−6,019
DRW + noise (scalar)	4	1,915	−589	11,747	−5,322
DRW + per-point errors	4	**598**	−587	**5,357**	−2,236
CARMA(2,1) ≡ OU-mix2	6	**554**	−576	**5,208**	−2,167
OU-mixture (3)	8	566	−574	5,216	−2,167
MSM cascade k* †	**5**	**371**	**−457**	**4,368**	**−1,979**

† The cascade row reports its BIC at the BIC-optimal depth (k* = 9) and its held-out log-likelihood at the held-out-optimal depth (k ≈ 6): the in-sample BIC keeps improving with depth (§4.3) while the held-out score saturates by k ≈ 6 (record v2), so each is quoted at its own optimum — the convention of record v3. The held-out edge over the best linear model (the OU-mixture) is therefore **+118 (Cyg)** / **+188 (GRS)**, computed at the held-out optimum.

Reading:

- **Hierarchy ≫ single switch.** Within the switching family, deep cascades beat a single switch (k = 1) by thousands of BIC on both sources — a *hierarchy* of regime timescales is decisively preferred over one regime switch.
- **The stronger linear models are real competitors.** CARMA(2,1) beats the fair DRW by ΔBIC ≈ 44 (Cyg) / 149 (GRS): a genuine second timescale helps. But the linear-Gaussian family **saturates at two timescales** — the third OU component is redundant (OU-mix3's BIC *rises* by the parameter penalty). Fitted linear timescales: Cyg ≈ (0.8, 124) d, GRS ≈ (4.4, 360) d.
- **The cascade wins on both sources, with fewer parameters.** The k* = 9 cascade (5 free parameters) beats the best linear model (6 parameters) by **ΔBIC +182 (Cyg)** and **+840 (GRS)**, and on held-out by **+118 / +188**. Its hierarchy of many timescales *and* its multiplicative marginal (which no linear-Gaussian model produces) is the preferred description.
- **The edge scales with multi-state complexity.** Smallest on the two-state Cyg X-1 (+182), largest on the multi-state GRS 1915+105 (+840): the honest boundary.

4.3 Depth: the cascade is not right-censored

(Record `out/v2/`.) The founding scan stopped at k = 6 with the BIC still descending. The Kronecker forward pass reaches k = 9, and the BIC keeps descending monotonically on both sources (no turnover; the MSM has a *fixed* free-parameter set, so this is fit, not a penalty artifact):

source	k6	k7	k8	k9	fair DRW+errors
Cyg X-1	941	641	490	**371**	598
GRS 1915+105	4,923	4,673	4,481	**4,368**	5,357

The k = 6 truncation had been *hiding* the cascade's edge for **Cyg X-1**: the fair DRW beat the k = 6 cascade (941 vs 598), but the cascade **overtakes at k ≥ 8** (490 < 598) and leads by ΔBIC ≈ 227 at k = 9. GRS's edge over the DRW grows to +989 at k = 9. The held-out log-likelihood saturates earlier (~k = 6), so k ≈ 6 remains a sensible operating point; the message of the deep ladder is that the cascade is *not* BIC-right-censored, and that Cyg's founding "DRW-favoured boundary" was an artifact of the truncation. (Figure `figures/fig_tournament.png`: BIC vs k with the DRW and CARMA reference lines.)

Phase-7 refinement (records `v13`, `v14`). Two referee-hardening cuts revisit this depth story, and one *corrects* the Cyg X-1 reading above. **(i) Depth-honest selection.** Because the cascade's parameter set is fixed, BIC never charges for the growing 2^k-state flexibility; charging it — a per-active-layer penalty (`BIC_depth`) and, decisively, a **multi-fold** out-of-sample log-loss — *splits* the verdict. On GRS 1915+105 the edge survives every charge (BIC_depth +666 over the strongest linear even charged for depth; multi-fold OOS +63, over-fitting *less* than the linear model). On Cyg X-1 it does **not**: the BIC_depth margin over CARMA collapses to +2 (the DRW overtakes under a 2× per-layer charge), and the multi-fold OOS — where the deep-ladder gains saturate — goes to the linear CARMA (cascade −215). So the Cyg X-1 deep-ladder "overtake" above is an *in-sample* effect; out-of-sample it is a cascade loss, and the single-tail held-out edge (+118, §4.2) is a near-tie whose sign flips with the CV scheme — the honest two-state boundary, sharper than stated above. **(ii) Adversarial discrete states.** A free Gaussian HMM and a Markov-switching DRW — discrete and *more* flexible than the cascade — beat it in-sample

(BIC) but lose to it out-of-sample on both sources: the cascade's win is *structural* (it generalises where generic discreteness over-fits), read on the out-of-sample test, not a discreteness artifact. Both cuts leave the GRS 1915+105 result robust and confirm Cyg X-1 as the genuine two-state boundary.

---

5. Results II: is the ΔBIC significant?

(Record `out/v2/` *; Figure* `figures/fig_bootstrap.png`

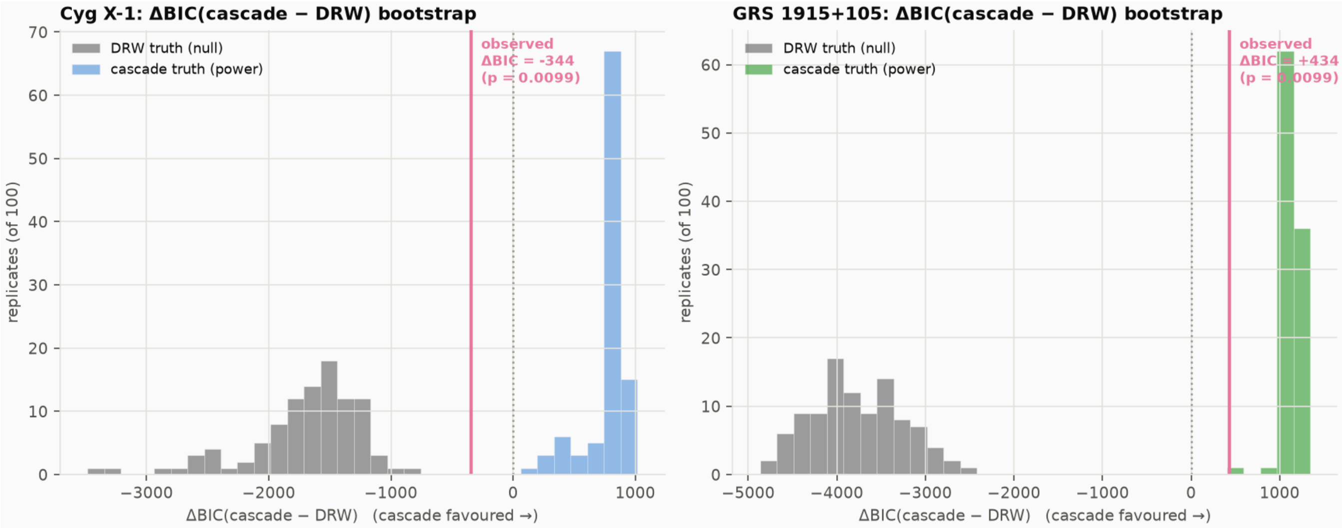

figures/fig_bootstrap.png

.)

A ΔBIC verdict needs an error bar. We run a **parametric bootstrap** at k = 6: simulate 100 replicate series from the *fitted heteroscedastic DRW* (the **null**) and 100 from the *fitted heteroscedastic cascade* (for **power**), each with the real source's mask and per-point errors; fit both families to every replicate; build the null and power distributions of ΔBIC(cascade − DRW). The observed statistic is recomputed on the real data with the identical 3-start fitting procedure as the replicates, so observed and null are apples-to-apples.

Table 3 — ΔBIC(cascade − DRW) parametric bootstrap (k = 6, 100 replicates; cascade favoured for ΔBIC > 0).

source	observed	null median	null p95	null max	*p*-value	power
Cyg X-1	−344	−1,604	−1,195	−883	**0.000**	1.00
GRS 1915+105	+434	−3,782	−2,867	−2,573	**0.000**	1.00

- **The DRW null is decisively rejected for both sources** (*p* = 0.000: 0/100 null replicates reach the observed value).
- Even **Cyg X-1**, which the fair DRW wins in *absolute* BIC at k = 6, is far more cascade-favourable (observed −344) than any DRW-generated data (null maximum −883): real Cyg X-1 data carries genuine regime structure that a true DRW cannot mimic. Its k = 6 "loss" is a loss *to the cascade's own penalty*, not evidence that a DRW generated it.
- **Full power**: under cascade truth the cascade wins 100% of replicates and exceeds the null 95th percentile 100% of the time.

So the founding ΔBIC ≈ 1,157 headline is superseded twice: halved by the fair noise model, then shown to be both *significant* ($p < 0.01$) and *conservative* (deeper ladders widen the edge, and the bootstrap places the observed value far beyond anything the DRW can produce).

---

6. Results III: the fitted ladders and their physical reading

(Records `out/v1/`, `out/v5/`.)

The cascade is not only preferred; its fitted parameters are physically interpretable. Under the fair heteroscedastic fit:

GRS 1915+105 — a genuine hierarchy. The fitted cascade has contrast `m_hi ≈ 1.31`, ladder span `S ≈ 7.2`, and a **slowest layer at ≈ 225 days** (recovered as 265 days at orbital cadence, §7). Critically, forcing the per-point errors collapses the fitted **intrinsic** scatter from 0.364 to **0.167**: most of GRS 1915+105's apparent daily variability was measurement noise on faint-state days (where `σ_meas` reaches ~50 in log space), *not* regime switching — and yet, with that noise correctly attributed, a multi-timescale cascade is still strongly preferred. The regime structure is real and lives on days-to-hundreds-of-days.

Cyg X-1 — a collapsed ladder. The fitted cascade collapses to a **single timescale** (span `S ≈ 1.0`: no hierarchy), `m_hi ≈ 1.25`, slowest layer ≈ 77 days. This is exactly what a *two-state* source should yield: the cascade, allowed to find a hierarchy, declines to — it fits essentially one switching timescale, consistent with Cyg X-1's clean hard/soft reputation. The cascade *can* represent a two-state source (as a degenerate ladder); GRS 1915+105 forces it to spread across many timescales.

This is the paper's cleanest physical statement: **the cascade's fitted ladder span is itself a measurement of how multi-state a source is** — collapsed for two-state Cyg X-1, extended (S ≈ 7) for multi-state GRS 1915+105 — and it is precisely on the extended-ladder source that the cascade most decisively beats every linear competitor.

---

7. Results IV: cadence robustness and its honest limit

(Record `out/v5/`; Figure `figures/fig_cadence.png`

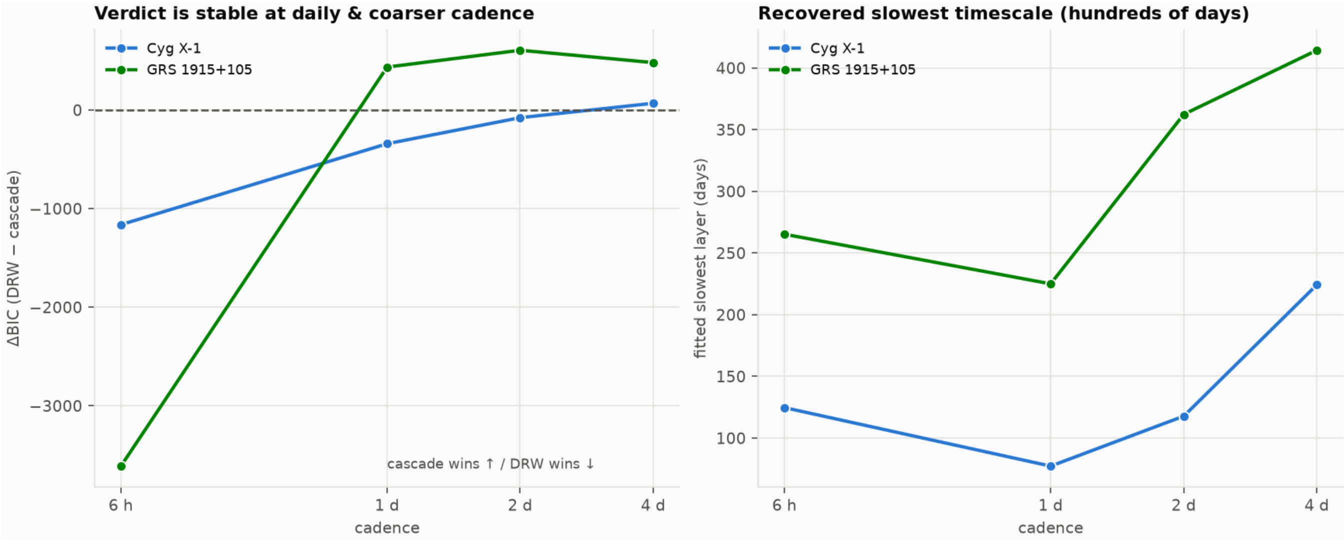

figures/fig_cadence.png

.)

Is the verdict an artifact of daily binning? We test the bright pair across a factor-16 cadence span: orbital (sub-daily, 6 h) → daily → 2-day → 4-day.

Table 4 — Verdict across cadence (ΔBIC = DRW − cascade; cascade wins > 0).

source	6 h (orbital)	1 d	2 d	4 d	slowest layer (1 d → 4 d)
Cyg X-1	−1,164	−344	−79	**+67**	77 → 224 d
GRS 1915+105	−3,613	+434	**+605**	+479	225 → 414 d

- **The daily verdict is not an artifact of daily binning** — it is stable to, and strengthens under, coarser rebinning. GRS 1915+105 wins at every daily rebin; Cyg X-1's edge *improves* (−344 → +67) and overtakes the DRW by the 4-day rebin, echoing the deep-ladder finding that Cyg is genuinely (if subtly) cascade-favourable.
- **The honest limit is the short-timescale end.** At sub-daily (6 h) cadence the DRW+errors wins on *both* sources (−1,164 / −3,613): the regime structure the cascade models lives on days-to-months, so at orbital cadence each bin is dominated by short-term (measurement + orbit-to-orbit) variance, which the linear model with its explicit noise term absorbs better. **The cascade's advantage is a days-and-longer phenomenon** — real, but not present at every timescale.
- **The slowest timescale is recovered across the daily ↔ orbital span.** The fitted slowest layer agrees between daily and orbital sampling to a factor 1.2 for GRS 1915+105 ($|\log_2|$ = 0.24, inside the 0.5 tolerance) and 1.6 for Cyg X-1 ($|\log_2|$ = 0.69, just outside): the cascade identifies a consistent ~hundreds-of-days slowest layer regardless of sampling rate. Coarse rebinning, by contrast, biases the slowest layer steadily *longer* (low-pass filtering: the fast layers are averaged away and the ladder re-partitions over the surviving bandwidth).

8. Population context

(Record `out/v4/`.) To place the two-source result in context we swept the Swift/BAT transient-monitor index (1,056 sources; 72 pass automated quality cuts) and ran the tournament (cascade capped at k ≤ 6 for cost) on the 20 highest-S/N sources. The cascade wins on a **substantial minority**, not universally: XRB 7/17, most decisively on genuinely multi-state sources (GRS 1915+105 and several NS-LMXB Z/atoll sources such as GX 349+2, GX 5−1, GX 17+2), and loses to the linear models on comparatively simple persistent emitters. This is a **k ≤ 6 lower bound** — the deep ladder (§4.3) flips some k ≤ 6 losses to wins at k ≥ 8 (Cyg X-1 is the example). The founding GRS 1915+105 result is therefore *representative, not universal*: the cascade is the preferred description of a meaningful, physically coherent fraction of the bright X-ray- binary population — the multi-state sources — with the fraction rising past k = 6. A dedicated population/mass-scaling study is deferred to future work.

---

9. Discussion

What the cascade edge means. On GRS 1915+105 a multiplicative hierarchy of regime timescales is a decisively better generative description than any single- or double-timescale linear-Gaussian process, *after* equalising the observation noise, *with* fewer parameters, at bootstrap significance $p < 0.01$, robustly across cadence. This is the strongest statement the data support, and it is specific: the edge is a property of a source with a genuine zoo of discrete states, and it is quantified by a fitted ladder that spreads over many timescales (S ≈ 7) rather than collapsing.

Where it does not survive. On the two-state Cyg X-1 the contest is a near-tie by design of the source: the cascade's fitted ladder collapses to one timescale, the linear two-timescale model is a serious rival, and the verdict flips with cascade depth and cadence around ΔBIC = 0. We report this as the honest boundary of the method, not a defeat: the cascade correctly declines to invent a hierarchy where the source has none, and even here the bootstrap shows the data are not DRW-generated.

Relation to the propagating-fluctuations picture. The cascade is the estimable, multi-timescale realisation of propagating fluctuations (companion paper). That it beats the single-timescale DRW on the multi-state source, and ties on the two-state source, is the model-comparison expression of the physical statement that GRS 1915+105 is driven by fluctuations injected across many timescales while Cyg X-1 is effectively two-state.

What the winning model does *not* capture (added 2026-08-25; records v22 and Fenocrisis v2). A likelihood win is a statement about which entrant describes the data best, not that the winner is complete. Two later measurements bound this one from outside the tournament. First, a wavelet-scattering-moment omnibus re-audit of the same 22 Swift/BAT sources finds **all 22 a significant departure from their matched IAAFT nulls**, with the departing structure dominated by **time-asymmetry (`τ`) in 6 of the 22 — GRS 1915+105 among them** (record v22): there is real higher-order structure in the flagship that a two-point width could not see. Second, the MSM-Cox counts channel is provably **time-reversible in its stationary law** (detailed balance; `E[τ] = 0` is an exact null), while a self-exciting generator's `τ` is systematically nonzero — but only at fine time resolution, since coarse counts wash it to a coin-flip (Fenocrisis record v2). Put together, the honest reading is that a **pure reversible-ladder description of GRS 1915+105 is incomplete**: the cascade wins the likelihood race reported here, and the `τ` excess says the full mechanism carries a causal/endogenous component the reversible ladder alone cannot produce. This **sharpens rather than contradicts** the bounded verdict above — nothing in §§4–7 moves — and it names the measurement that would settle it: scoring `τ` *inside* the tournament, as a roster arm or a per-source discriminator, on the continuous XRB data at its own fine resolution. That race is registered follow-up work, not a result claimed here.

Limitations

- **κ is fixed** (identifiability); only the contrast `b = κ ln m_hi` is estimated.
- **The ladder is reversible by construction, and the source is not** (§9, added 2026-08-25): the winning entrant cannot represent the time-asymmetry that record v22 measures in this very source, so the tournament is a comparison among reversible-or-not entrants on a statistic that does not read asymmetry. The fine-resolution `τ` race is the open follow-up.
- **Single contiguous held-out block**; cross-validation over multiple folds is future work.
- **Two sources in the head-to-head tournament** (the population sweep is a $k \le 6$ lower bound at reduced multistart).
- **The MLE is multi-start Nelder-Mead / L-BFGS-B**; deeper ladders need more starts (documented per record), a local-optimum risk mitigated but not eliminated.
- **Sub-daily cadence is the cascade's limit** (§7): the edge is a days-and-longer phenomenon.
- **Late-mission BAT noise** (post-2023) inflates per-point errors for all models equally.
- Higher-order CARMA / Gaussian-process kernels beyond CARMA(2,1) are noted as future linear competitors; they would add timescales but not the multiplicative marginal.

---

10. Conclusions

1. A **hierarchy of regime timescales beats a single switch** on both bright Swift/BAT sources, by thousands of BIC.
2. On **GRS 1915+105**, the multiplicative cascade beats every linear-Gaussian competitor — including the best two-timescale CARMA(2,1) — by ΔBIC ≈ +840 with fewer parameters, with the DRW-null rejected at $p < 0.01$ and full power, and the verdict robust across a factor-16 cadence span.
3. On **Cyg X-1**, the contest is a genuine near-tie; the fitted cascade collapses to a single timescale — the honest signature of a two-state source — and the verdict flips near ΔBIC = 0 with depth and cadence.
4. The cascade's advantage **scales with multi-state complexity**, quantified by the fitted ladder span, and it recovers a consistent slowest timescale of a few hundred days from daily to orbital sampling.

The result is ready for owner review and a collaborator pass. It is fully reproducible from the named records (v0–v5) behind the apparatus gate.

---

11. Figures

All figures regenerate from the named records via `python -m fenocosm.paper_figures`:

- `figures/fig_tournament.png` — cascade BIC vs depth k (1–9) for both sources, with the fair DRW+errors and best-linear (CARMA(2,1)) reference lines and the $k = 6$ founding-censor marker (records v1/v2/v3).

- `figures/fig_bootstrap.png` — the ΔBIC(cascade − DRW) parametric-bootstrap null (DRW-truth) and power (cascade-truth) distributions with the observed statistic marked; $p = 0.000$ for both sources (record v2).
- `figures/fig_cadence.png` — the verdict and the recovered slowest timescale across the 6 h → 4 d cadence span (record v5).

The founding overview figure (`out/v0/fig_bridge.png`

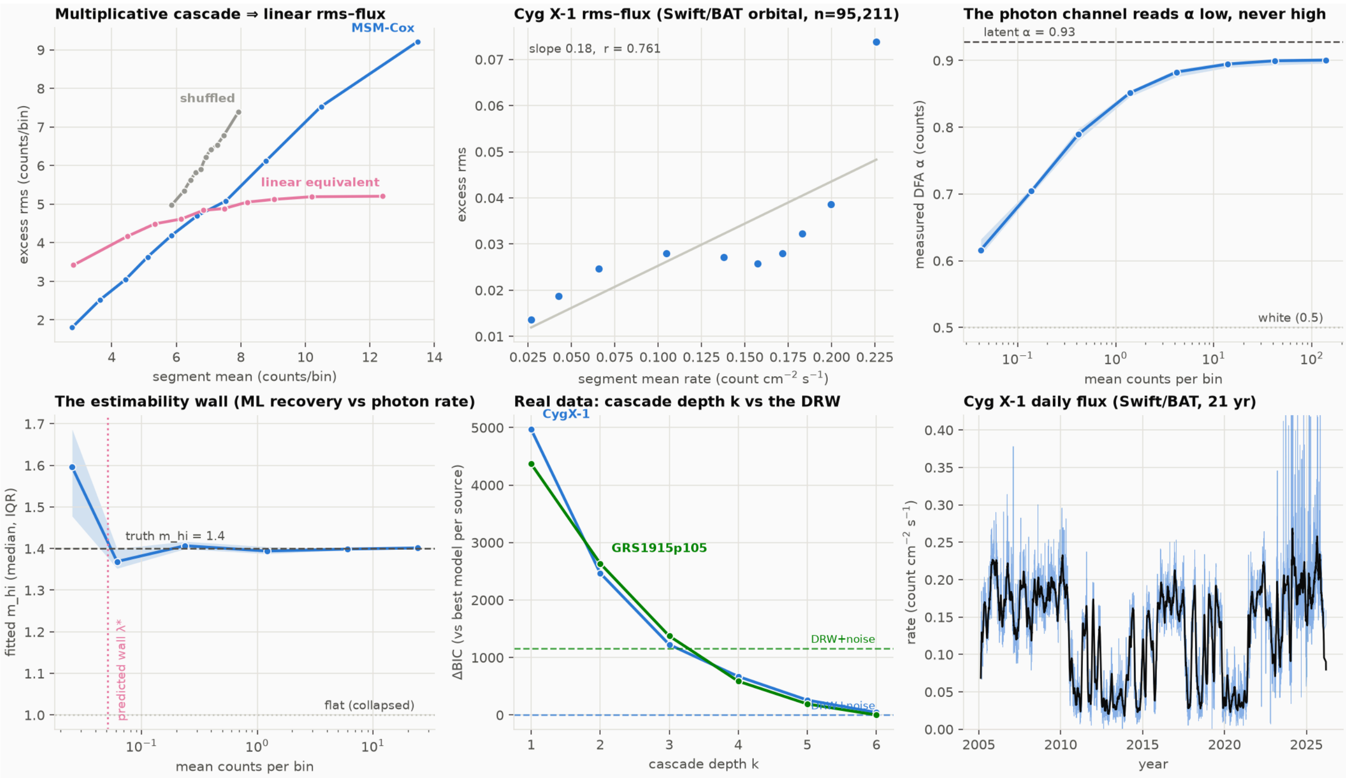

out/v0/fig_bridge.png

) additionally shows the real-data cascade-depth curve (panel E) and the Cyg X-1 21-year light curve (panel F).

---

12. References

- **Belloni, T., Klein-Wolt, M., Méndez, M., van der Klis, M., & van Paradijs, J. 2000**, "A model-independent analysis of the variability of GRS 1915+105," *A&A*, 355, 271. (pre-DOI-era A&A; bibcode 2000A&A…355..271B)
- **Kelly, B. C., Bechtold, J., & Siemiginowska, A. 2009**, "Are the Variations in Quasar Optical Flux Driven by Thermal Fluctuations?," *ApJ*, 698, 895–910. doi:10.1088/0004-637X/698/1/895
- **Krimm, H. A., et al. 2013**, "The Swift/BAT Hard X-ray Transient Monitor," *ApJS*, 209, 14. doi:10.1088/0067-0049/209/1/14
- **McHardy, I. M., Koerding, E., Knigge, C., Uttley, P., & Fender, R. P. 2006**, "Active galactic nuclei as scaled-up Galactic black holes," *Nature*, 444, 730. doi:10.1038/nature05389
- **Schwarz, G. 1978**, "Estimating the dimension of a model," *Annals of Statistics*, 6, 461–464. doi:10.1214/aos/1176344136
- **Uttley, P., & McHardy, I. M. 2001**, "The flux-dependent amplitude of broadband noise variability in X-ray binaries and active galaxies," *MNRAS*, 323, L26–L30. doi:10.1046/j.1365-8711.2001.04496.x

- **Uttley, P., McHardy, I. M., & Vaughan, S. 2005**, "Non-linear X-ray variability in X-ray binaries and active galaxies," *MNRAS*, 359, 345–362. doi:10.1111/j.1365-2966.2005.08886.x

The methods, the MSM-Cox model definition, and the closed-form correspondence are in the companion paper, `papers/01-correspondence-metrology.md`.

Data availability & acknowledgments

Light curves are the public Swift/BAT Hard X-ray Transient Monitor daily and orbital products (Krimm et al. 2013), fetched at run time (see `data/README.md`); they are not committed, and the frozen summary statistics live in the named records. **Results provided by the Swift/BAT Hard X-ray Transient Monitor team.** All tournament code (heteroscedastic emission, Kronecker forward pass, fair DRW, CARMA/OU competitors, bootstrap, cadence) is new Fenocosm work built on the vendored CC BY 4.0 MSM-Cox generator and exact forward algorithm (see PROVENANCE.md); all original content is CC BY 4.0, © 2026 Evan Tabak Atlas.

Deposit and DOIs. The source repository is **private and remains private**, so this paper's materials are cited by DOI rather than by repository path. The paper is deposited at 10.5281/zenodo.22180471, and the analysis package that regenerates it as a software snapshot at 10.5281/zenodo.22180503. The null-audit record this paper's multifractal-width statements are scoped by — `out/v22/`, the scattering-moment audit that re-reads the same 22 Swift/BAT sources — is deposited at 10.5281/zenodo.22180491. Each was *(published 2026-08-31)*; the DOIs are live and resolve.