

Entropy as Record

A Locked-Source Shape Test of an SMBH-Dominated Cosmic Entropy Ledger Across Redshift

A T16 empirical demonstration: pre-specified shape criteria under a locked cosmological entropy ledger

Jeremy C. Jones

HoldingLight LLC

ORCID: 0009-0007-2515-3774

contact@universalcollapse.com

v1.0 · 2026-07-02

Abstract

Several interpretive frameworks, including Universal Collapse Theory's standards layer, treat entropy as an accumulated record of irreversible structure formation — durable constraints written into physical degrees of freedom (Jones 2026c). This motivates a shape-level empirical question: when standard cosmological entropy components are assembled into a time-indexed ledger under Λ CDM, does the resulting curve exhibit the smooth, monotone, parsimoniously structured growth a record-accumulation reading expects, or does it require complex multi-kink behavior inconsistent with that reading? We test this with a locked-source ledger: four primary sources fixed in advance with no post-hoc substitution (Egan and Lineweaver 2010 anchors; Shankar et al. 2009 SMBH demographics; Madau and Dickinson 2014 stellar history; Profumo et al. 2024 census), binned estimates of $y(t) = \log_{10} S_{\text{tot}}(t)$ over $z \in [0, 6]$ under a fixed Planck18 redshift-to-time transform, and two pre-specified models — a single linear trend (M1) and a one-break continuous piecewise linear alternative (M2). One construction fact governs interpretation throughout: under this ledger the total is BH-dominated, so the test's shape content is the SMBH growth history mapped through a fixed entropy scaling (§3.3).

The design was pre-specified and locked before execution: source lock file, fallback policies, evaluation criteria, and smoothness thresholds, with the executed state carried by a git-committed repository and a SHA-256-stamped artifact bundle, and deviations recorded in an Update Integrity Standard-style log (Jones 2026d). In the primary run ($\Delta z = 0.5$, full window), M2 is decisively preferred ($\Delta\text{BIC} = -26.67$ under the published parameter convention; -24.2 under the conservative convention that counts the grid-selected break as a free parameter; LOBO-CV RMSE improves by 1.18), with a breakpoint at $t_{\text{break}} \approx 1.96$ Gyr separating rapid early accumulation (4.56 dex/Gyr) from slow late accumulation (0.24 dex/Gyr). Smoothness diagnostics show zero large slope jumps at the pre-specified scale. A 2×2 window-by-resolution sensitivity suite shows breakpoint location is window-driven, not resolution-driven. One pre-specified criterion — C5, a BH rate-share direction check — fails and is reported as a failure.

The claim boundary is carried explicitly: under this construction the ledger is BH-dominated ($S \sim M^2$), and the non-BH sub-ledger is constant by construction, so the shape test's discriminating content is carried entirely by the SMBH growth history mapped through the entropy scaling. The result is a shape-consistency finding, not a discriminating test against conventional cosmology: it establishes that the assembled entropy record is monotone, smooth, and compressible to a single structural bend — the record-accumulation reading survives its pre-specified shape falsifiers — while identifying independent

$\rho_{\text{BH}}(z)$ reconstruction and a gas redshift series as the load-bearing next data. At T16 conversion, the archived bundle was re-verified byte-for-byte against its published SHA-256 and all four runs were independently re-fit from the archived data, reproducing every breakpoint exactly.

Keywords: Universal Collapse Theory; entropy; records; cosmology; shape test; locked sources; pre-specified analysis; model selection; SMBH; sensitivity analysis; reproducibility.

1. Introduction

Entropy is commonly treated as a global thermodynamic bookkeeping quantity. Several interpretive frameworks instead read entropy as an accumulated record of irreversible structure formation — durable constraints written into physical degrees of freedom. Within Universal Collapse Theory this reading is native: the Records standard (Jones 2026c) treats record formation as the load-bearing product of constraint-guided resolution, and the physics working paper (Jones 2026a) carries the cosmological case. That framing motivates a testable shape-level question that does not require adopting the framework: **does the time evolution of total cosmic entropy look like a smooth, monotone record-accumulation curve, or does it require complex, multi-kink behavior inconsistent with a parsimonious record interpretation?**

Rather than proposing new cosmological dynamics, this paper performs a pre-specified shape test on compiled estimates of the major cosmic entropy components. The goal is not cosmological parameter inference; it is to evaluate whether the record-accumulation interpretation is consistent with the empirically implied shape of $\log_{10} S_{\text{tot}}(t)$ over the data-supported epoch — with the failure modes named in advance.

UCT positioning (scope note). UCT functions here only as a hypothesis generator for which curve-level discriminators are worth testing: trend, smoothness, and breakpoint stability. The paper does not depend on UCT's broader ontological commitments; the results can be interpreted entirely within conventional thermodynamics and cosmology. One boundary is stated at the outset: this is a shape-consistency test of an assembled literature ledger, not a discriminating test against standard cosmology and not an inference about entropy microphysics. What it can do is falsify a parsimonious record reading at the ledger level — via non-monotonicity, multi-kink spikiness, or irreducible model complexity — and none of those failure modes occurred.

This paper joins the T16 empirical corpus (Rice Hysteresis, COGITATE iEEG, and the AI-substrate demonstration; Jones 2026b) as the physics-wing entry, and is the v1.0 merge of the original February 2026 paper and its Registered Analysis Report into a single document under the family standard.

2. Pre-Specified Design and Registration Mechanism

2.1 Questions, criteria, and models

The analysis plan pre-specified four primary questions and five success/failure criteria (numbered C1–C5 in the plan), targeting: monotone positive growth of the total; a strict complexity cap (at most one breakpoint); smoothness — no "merger-spiky" multi-kink structure at the pre-specified bin scale; model selection between the two pre-specified models; and a BH rate-share direction check (C5: the BH share of the entropy accumulation rate increases toward late times). Two models were fixed in advance: **M1**, a single linear trend $y(t) = a + bt$; and **M2**, a one-break continuous piecewise linear model with the breakpoint selected by grid search over candidate bin midpoints. Evaluation was fixed as information

criteria (AICc/BIC), leave-one-bin-out cross-validation RMSE (LOBO-CV), and smoothness metrics computed from finite-difference slopes against a pre-specified jump threshold recorded per run. Fits are unweighted least squares on the point ledger — verified at conversion, where an independent unweighted re-fit reproduces the archived fits exactly (§6). AICc and BIC serve as descriptive model-selection scores under an independent-Gaussian residual convention on the point ledger; they are not a formal observational likelihood, and bins share reconstruction systematics through the locked sources.

2.2 Registration mechanism

Registration was internal and artifact-anchored, in the same class as the frozen-configuration mechanism used elsewhere in the T16 family: a locked source file (sources.md / sources.yaml) fixing the four primaries, allowable fallbacks, and the no-post-hoc-substitution rule; a git-committed repository carrying the executed state (commit 5a0d430dc7f50cf0157f555cdd40d639afe61675); a SHA-256-stamped artifact bundle (§10) generated 2026-02-17 19:59:06 containing inputs, scripts, outputs, and QC narratives; and an Update Integrity Standard-style deviations log (Jones 2026d) recording every departure from the plan. No third-party timestamped registry was used; "pre-specified" throughout this paper refers to commitments fixed in that locked, committed state before the reported runs.

3. Data Sources and Ledger Construction

3.1 Locked source set

Source	Role in pipeline	Extracted anchors / series
Egan & Lineweaver (2010)	Present-day entropy budget anchors ($z \approx 0$ normalizations)	$\log_{10} S_{\text{BH}}(0) = 104.49$; $\log_{10} S_{\text{star}}(0) = 80.98$; $\log_{10} S_{\text{rad}} = 90.02$ (photons + relic neutrinos); Table 1 provenance archived
Shankar et al. (2009)	SMBH demographic evolution (BH component shape)	$\rho_{\text{BH}}(z)$ reconstructed from Eqs. 1–5 + Table 1 parameters (deviation D1, §7)
Madau & Dickinson (2014)	Stellar history proxy (stellar component shape)	Stellar mass density from SFRD integration (Eq. 15, return fraction $R = 0.27$)
Profumo et al. (2024)	Modern census framing; gas/plasma treatment	$\log_{10} S_{\text{gas}}(0) = 83.15$ (diffuse baryons); no gas z -series available (fallback, §7)

Table 1. Locked source set. No post-hoc substitution permitted; sensitivity variants may be added only as explicitly labeled variants. Per-source extraction provenance (tables, methods, and QC images) is archived in the reproducibility bundle.

3.2 Binning, cosmology transform, and totals

The primary run uses uniform redshift bins of $\Delta z = 0.5$ spanning $z \in [0, 6]$ (12 bins; midpoints 0.25 ... 5.75), with cosmic time computed at bin midpoints under Planck18 (Planck Collaboration 2020). Components: BH entropy from SMBH mass density with Bekenstein–Hawking scaling $S \sim M^2$, normalized to the $z \approx 0$ anchor; stars scaled by the stellar-mass-density proxy and normalized to the $z \approx 0$ anchor; radiation treated as constant comoving entropy; gas held constant in comoving accounting under the pre-specified fallback (no z -series in the locked census source), with $\sigma_{\log_{10}} = 0.3$ dex. Totals are computed by stable log-sum, and uncertainty is propagated via Monte Carlo draws as specified in the plan. The default $\sigma_{\log_{10}} = 0.3$ dex applies where per-bin uncertainty series were not extractable.

3.3 Construct scope: a BH-dominated ledger

One construction fact governs the interpretation of everything that follows, and is stated here rather than discovered in discussion. Under this ledger, the total is BH-dominated across the full window, and the non-BH sub-ledger is **constant by construction**: radiation is constant bookkeeping (90.02), gas is held constant under the fallback (83.15), and stellar entropy — the only independently evolving non-BH component — sits 9–11 dex below radiation and is therefore invisible in the log-sum (verified from the archived processed data at conversion: the stars + gas + radiation sub-total is 90.024 at every bin, a span of 0.000 dex). The shape test's discriminating content is therefore carried entirely by the SMBH growth history mapped through $S \sim M^2$. This is a scope condition of the construct, not a defect discovered late: it defines what the test can and cannot say (§8), and it makes independent $\rho_{\text{BH}}(z)$ reconstruction and a gas redshift series the load-bearing next data (§9).

The BH mapping, precisely. The operational rule, read from the archived build script and verified against the archived series at conversion (implied exponent 2.000 on every bin step), is $\log_{10} S_{\text{BH}}(z) = \log_{10} S_{\text{BH}}(0) + 2 \cdot [\log_{10} \rho_{\text{BH}}(z) - \log_{10} \rho_{\text{BH}}(0)]$, with the stellar component mapped analogously at exponent 1. Because population black-hole entropy tracks the second moment of the mass function (the sum of squared masses), which is not determined by total mass density alone, the exponent-2 rule encodes a definite assumption: a fixed comoving population growing self-similarly, with all density evolution attributed to the growth of existing holes and none to number evolution. The BH curve is therefore a locked demographic proxy for entropy-shape evolution under a fixed mapping, not a reconstruction of the evolving second moment. The proxy's influence on the conclusions is bounded and was measured at conversion: any fixed positive exponent is an affine transform of $\log_{10} \rho_{\text{BH}}$, so breakpoint location, M2-versus-M1 preference, monotonicity, and the smoothness pass are invariant to the exponent choice — an exponent-1 counterfactual reproduces $t_{\text{break}} = 1.959$ Gyr and ΔBIC to three decimals, with only the slopes and span rescaling (early slope $4.56 \rightarrow 2.28$ dex/Gyr; span $6.7 \rightarrow 3.4$ dex). What the fixed-exponent family cannot represent is an evolving mass-function shape — a time-varying effective exponent — which is exactly what the second-moment falsifier of §9.2 targets.

4. Primary Results (A-full: $\Delta z = 0.5$, $z \leq 6$)

4.1 Model selection

The total entropy ledger rises 6.75 dex across the analyzed window ($t \approx 0.98$ – 10.75 Gyr). M2 is decisively preferred over M1 on every pre-specified comparison:

Model	R ²	RMSE	AICc	BIC	CV_RMSE	Selected
M1 (linear)	0.579	1.354	41.54	41.17	1.983	—
M2 (one break)	0.963	0.402	16.05	14.50	0.800	yes

Table 2. Primary-run fit summary (from the archived fit tables). M2 parameters: $t_{\text{break}} = 1.959$ Gyr (bin midpoint $z_{\text{mid}} = 3.25$); early slope 4.560 dex/Gyr; slope change -4.324 , giving a late slope of 0.236 dex/Gyr. ΔBIC (M2 – M1) = -26.669 ; $\Delta\text{CV_RMSE} = -1.183$.

Parameter-counting convention. The published $\Delta\text{BIC} = -26.669$ uses the convention identified from the archived fit tables ($k = 2$ for M1, $k = 3$ for M2): the grid-selected breakpoint is not counted as a free parameter. Under the conservative alternative that counts it ($k = 4$ for M2), the independently recomputed ΔBIC is -24.2 . M2 is decisively preferred under either convention; the published value is retained as the artifact value with the convention stated.

The interpretation of the selected shape is a two-regime accumulation profile: rapid early record accumulation (≈ 4.56 dex/Gyr before $t_{\text{break}} \approx 1.96$ Gyr) transitioning to slow late accumulation (≈ 0.24 dex/Gyr), with the transition landing in the early universe under the full window.

This is a point-ledger model-selection result within the pre-specified M1/M2 class: it licenses "not globally linear, and compressible to one structural bend within that class," not a claim that one-break piecewise linear is the unique best-fitting growth form. Smooth saturating alternatives (logistic- or Gompertz-type curves, monotone splines) were not in the pre-specified class and are untested here.




Figure 1. Primary run (A-full), regenerated from the archived processed data at conversion. Left: ledger components and total — the BH component carries the total; radiation and gas are constant by construction; stars evolve but sit far below radiation. Right: binned total with $\pm \sigma$, the rejected M1 trend, and the selected M2 fit with its breakpoint.

4.2 Smoothness

At the pre-specified scale the curve is not jagged: the number of large slope jumps is **0**, with a maximum absolute slope jump of 5.135 against the pre-specified threshold of 10.103 (per-run thresholds are recorded in each run's smoothness metrics file). The deviation from global linearity is a single structural bend, not multi-kink spikiness — the pre-specified smoothness criterion passes.

4.3 The pre-specified C5 criterion fails

C5 pre-specified that the BH share of the entropy accumulation rate increases toward late times. Under the strict early/late windows of the plan it fails: median share ≈ 1.003 early ($z \approx 1.5$ – 2.5) against ≈ 0.982 late ($z \approx 0$ – 0.5); `passes_C5` = False. This is recorded as a pre-specified failure, full stop. Two diagnostic observations accompany the record without softening it. First, under a BH-dominated ledger with gas held constant, rate shares of the total are numerically fragile where dS_{tot}/dt is small — C5 as specified was a weak instrument for its target under this construction. Second, a dominance-stable log-derivative share was computed afterward as a labeled post-hoc descriptive (§5.4); it does not rescue, replace, or re-adjudicate the criterion. A redesigned rate-share criterion belongs to the new-data variant (§9), not to reinterpretation of this run.

4.4 Pre-specified outcome ledger

Pre-specified target	Outcome	Status
----------------------	---------	--------

Monotone positive growth of the total	Total increases across the window in all four runs (verified at conversion)	Pass
Complexity cap: at most one breakpoint	One break suffices; no evidence of further structure at the tested scales	Pass
Smoothness: no multi-kink spikiness at the pre-specified scale	0 large slope jumps in all four runs	Pass
Model selection between the pre-specified models	M2 selected over M1 in all four runs	Pass
BH rate-share increases toward late times (C5)	Early median 1.003 vs. late 0.982 under strict windows	Fail

Table 3. Outcome ledger for the pre-specified targets. The analysis plan numbers its criteria C1–C5; only C5's number is unambiguous in the surviving documents, so targets are listed by name with C5 flagged explicitly. The plan document's numbered wordings are part of the project lineage.

5. Sensitivity Analyses

5.1 The 2×2 window-by-resolution suite

Because breakpoint location can depend on binning and epoch coverage, four locked-source runs isolate the two axes: window (full $z \leq 6$ vs. restricted $z \leq 2.5$) crossed with resolution ($\Delta z = 0.5$ vs. 0.25).

Run	Window	Δz	Bins	t_{break} (Gyr)	ΔBIC (M2–M1)	$\Delta \text{CV_RMSE}$	Large jumps
A-full	$z \leq 6$	0.5	12	1.959	–26.669	–1.183	0
B-full	$z \leq 6$	0.25	24	2.048	–58.662	–1.102	0
A-restricted	$z \leq 2.5$	0.5	5	4.961	–10.512	–0.187	0
B-restricted	$z \leq 2.5$	0.25	10	4.594	–24.014	–0.235	0

Table 4. 2×2 sensitivity suite (archived values; parameter convention as in §4.1). M2 is selected and smoothness passes in every run.

Key inference: the breakpoint shift is window-driven, not resolution-driven. Holding the window fixed, halving the bin width moves the breakpoint only slightly (full: $1.96 \rightarrow 2.05$ Gyr; restricted: $4.96 \rightarrow 4.59$ Gyr). Changing the window moves it materially, consistent with the restricted window removing the early steep regime entirely — the fit then localizes the remaining curvature later. The structural conclusions (M2 selection, monotonicity, smoothness) are invariant across all four runs.




Figure 2. The 2×2 sensitivity suite, each panel regenerated from its archived processed data at conversion: binned totals $\pm \sigma$ with the selected M2 fit and breakpoint. Within a fixed window, resolution barely moves the break; across windows, it moves materially.

5.2 Degrees-of-freedom honesty on the restricted runs

A-restricted fits a four-parameter M2 (counting the break) to five bins — one residual degree of freedom — so its ΔBIC of -10.5 should be read as directional only. The restricted-window conclusion is carried by B-restricted (10 bins, $\Delta\text{BIC} = -24.0$). The full-window runs are comfortably determined (12 and 24 bins).

5.3 Stars-only shape check (post-hoc, descriptive)

Because the total's shape content is BH-carried (§3.3), the one independently evolving non-BH component provides a limited but genuine secondary check. Applied to the stellar component alone (Madau–Dickinson-derived, 2.0 dex of monotone growth), the identical M1/M2 procedure prefers the one-break model ($\Delta\text{BIC} \approx -29.5$ under the conservative convention) with the same qualitative structure: rapid early accumulation (0.73 dex/Gyr) transitioning to slow late accumulation (0.07 dex/Gyr) at $t_{\text{break}} \approx 2.9$ Gyr. Two independent literature reconstructions — SMBH demographics and stellar history — thus both exhibit the rapid-early / slow-late one-break shape class under this ledger's mapping. This check was not pre-specified; it is computed from the archived processed data at conversion and reported as descriptive — the stellar sub-ledger shows the same qualitative shape class post hoc but does not materially affect the total under this construction (§3.3).

5.4 Dominance-stable rate share (post-hoc, descriptive)

As a labeled non-pre-specified stability check, a log-derivative share — $[d/dt \log_{10} S_{\text{BH}}] / [d/dt \log_{10} S_{\text{tot}}]$ — avoids the numerical fragility of raw rate shares under dominance. Across runs it remains near unity with a slight late decrease. It is reported alongside, not within, the pre-specified outcomes, and does not alter the C5 record of §4.3.

5.5 Monte Carlo model-selection robustness (post-hoc, computed at conversion)

The plan's Monte Carlo machinery propagated the ledger's stated uncertainties into bands but did not tally model selection across draws. That tally was computed at conversion from the archived primary-run data, under a stated convention: independent Gaussian per-bin draws in log space at the recorded σ , 10,000 draws, both parameter-counting conventions. Results: M2 is preferred in 100.00% of draws under either convention; $\Delta\text{BIC} < -10$ (strong preference) in 100.00%; the selected breakpoint falls at or before 3.0 Gyr in 99.92% of draws, with median 1.959 Gyr and interquartile range [1.959, 2.359]. Within the ledger's own uncertainty model, neither the model selection nor the early-universe localization of the break is fragile. The check is post-hoc, inherits the per-bin independence convention (which cannot represent shared reconstruction systematics), and does not upgrade the pre-specified record.

6. Independent Verification at Conversion

This T16 deposit adds a verification layer executed at conversion (2026-07), independent of the original pipeline. First, the archived artifact bundle was re-hashed and matches its published SHA-256 byte-for-byte (§10). Second, the M1/M2 analysis was re-implemented from scratch (independent least-squares code path; no reuse of the archived analysis script) and re-fit to the archived processed data of all four runs. Breakpoints reproduce exactly (1.959, 2.048, 4.961, 4.594 Gyr); early and late slopes and $\Delta\text{CV_RMSE}$ reproduce to reported precision; the smoothness maxima reproduce exactly (e.g., 5.135 on A-full with zero threshold exceedances); and ΔBIC reproduces to within 0.11 of the published values under the published parameter convention, with the conservative-convention values reported in §4.1. Third, the constancy of the non-BH sub-ledger (§3.3), the exponent counterfactual (§3.3), the stars-only descriptive (§5.3), and the Monte Carlo model-selection tally (§5.5) were computed from the same archived data. The February results are therefore not merely archived but independently re-derived.

7. Deviations

All deviations were pre-specified fallbacks or documented access constraints, recorded in the Update Integrity Standard-style log within the bundle; no source substitutions occurred. **D1 — BH evolution reconstruction:** the Shankar et al. (2009) $\rho_{\text{BH}}(z)$ series was reconstructed from the paper's equations (Eqs. 1–5) and parameter table because full electronic evolution tables were not embedded in the locally available document; the reconstruction stays within the locked source. **D2 — uncertainty defaults:** per-bin uncertainty series were not extractable for the BH and stellar inputs; the lock file's fallback $\sigma_{\log_{10}} = 0.3$ dex was applied. **D3 — gas handling:** the locked census source provides a $z \approx 0$ diffuse-baryon anchor but no gas/plasma redshift series; gas was held constant in comoving accounting at $\log_{10} S_{\text{gas}} = 83.15$ with $\sigma_{\log_{10}} = 0.3$, per the pre-specified fallback. D3 is the deviation with interpretive weight: it contributes directly to the construct scope of §3.3 and to the fragility of C5.

8. Interpretation (Bounded)

What the result licenses. At the level of this locked-source ledger, the assembled cosmic entropy record is monotone increasing, smooth at the pre-specified scale (zero large slope jumps in all four runs), and compressible to a single structural bend within the pre-specified model class (§4.1) — rapid early accumulation transitioning to slow late accumulation — with the model selection stable across a window-by-resolution sensitivity suite and the breakpoint's location attributable to epoch coverage rather than binning. The record-accumulation reading survives every pre-specified shape falsifier it faced: no non-monotonicity, no multi-kink spikiness, no irreducible complexity beyond one break. The simplest globally-linear-in-time expectation is not supported; the constrained expectation — directional, parsimoniously structured accumulation — is.

What failed. The pre-specified C5 rate-share criterion failed and stands as a failure (§4.3). The honest diagnosis — that C5 was a weak instrument under a BH-dominated, gas-constant construction — motivates its redesign for the new-data variant; it does not convert the failure into a pass.

What the result does not establish. It does not discriminate the record interpretation from conventional thermodynamics and cosmology: any monotone, smoothly integrated astrophysical growth history produces a curve of this general character, and the paper's scope note (§1) stands. It does not establish anything independent of the Shankar et al. (2009) reconstruction: under this construct the shape content is the SMBH growth history through $S \sim M^2$ (§3.3), so the primary result is best read as "the literature-implied SMBH-dominated entropy ledger has the pre-specified record shape," not as a multi-component consilience. It does not constrain gas thermodynamic evolution (held constant by fallback), does not perform cosmological parameter inference, does not establish that one-break piecewise linear is the best of all smooth monotone growth forms (untested classes, §4.1), and does not bear on entropy microphysics or on UCT's ontological claims. The stars-only descriptive (§5.3) is the one genuinely independent shape corroboration in the ledger, and it is post-hoc.

9. Limitations and Falsifier Conditions

9.1 Scope conditions

BH dominance (construct scope). Stated in §3.3 and carried through §8: the test's discriminating content is the BH component, mapped through the fixed-exponent demographic proxy of §3.3. Every conclusion is conditioned on this construction.

Gas held constant. The absence of a locked gas z-series means one physically evolving component enters as a constant. A future locked source providing gas entropy evolution could, in principle, introduce structure the present ledger cannot show — including the multi-kink behavior that would count against the record reading.

Literature-reconstruction basis. The "data" are published reconstructions and anchors mapped through fixed transforms, not a new observational likelihood. The test is theory-neutral in its criteria but inherits every systematic of its sources.

Model-preference robustness under uncertainty. The executed February analysis propagated uncertainties into bands but did not tally model selection across Monte Carlo draws. That tally now exists as a conversion-added post-hoc check (§5.5) and is strongly favorable. The residual limitation is the check's convention: independent Gaussian per-bin draws cannot represent systematics shared across bins through the locked reconstructions. A correlated-systematics variant belongs to the new-data pass.

Restricted-run degrees of freedom. Per §5.2, A-restricted is near-saturated (one residual degree of freedom) and is carried as directional only.

9.2 Falsifiers

The interpretation in §8 should be revised or withdrawn if: (i) an independent reconstruction of the SMBH entropy shape — ideally the evolving mass function's second moment rather than mass density alone (e.g., Soltan-type AGN luminosity-density integration with mass-function modeling) — run through the identical locked criteria breaks monotonicity, smoothness, or the one-break cap, or reveals the time-varying effective exponent that the fixed-exponent proxy of §3.3 cannot represent; (ii) a locked gas-entropy redshift series introduces multi-kink structure into the total at the pre-specified scale; (iii) a Monte Carlo ensemble model comparison shows M2 preference is not robust to the ledger's stated uncertainties; or (iv) the identical criteria applied to an updated census ledger select irreducibly higher complexity. Of these, (i) — the independent $\rho_{\text{BH}}(z)$ variant — is the load-bearing next test and is ordered first; it addresses the construct scope of §3.3 directly rather than incrementally.

10. Provenance and Reproducibility

The executed state is reconstructible from SHA-pinned artifacts, and the deposit-time verification of §6 confirms the chain end-to-end.

Artifact	Identity / location
Artifact bundle	EAR_artifacts_v0_1_2b_clean.zip — SHA-256 745d5249f878969abc95d317de27d35e48530de4e427455f3607cc7d54e24fb 4 (re-verified byte-for-byte at T16 conversion, 2026-07)
Repository state	git commit 5a0d430dc7f50cf0157f555cdd40d639afe61675; bundle generated 2026-02-17 19:59:06 (MANIFEST.txt, RUN_ID.txt)
Lock and plan	sources.md / sources.yaml (locked source set, fallback policies, no-post- hoc-substitution rule); TRANSFORMS.md; REPRODUCIBILITY_NOTE.md; environment.txt; requirements.txt
Deviations	deviations.md — Update Integrity Standard-style log (entries D1–D3, §7)
Inputs	final/entropy_components.csv (+ per-run variants for the 2×2 suite); bin definitions bins_A/B/C; raw_extracted/ per-source extractions with QC provenance images
Outputs (per run)	outputs/, outputs_B_full/, outputs_A_restricted/, outputs_binsB/ — figures, fit tables (table_fit_summary.csv, table_model_comparison.csv), smoothness_metrics.csv, rate-share tables, processed CSVs
Reproduction	python analysis.py --data final/entropy_components.csv --out outputs --mc 10000 --seed 1 (per-run variants recorded in the Registered Analysis Report, retained in the project's append-only archive; the analysis driver is carried at the pinned repository commit)

Table 5. Provenance chain. The bundle carries the inputs, extraction and QC scripts, outputs, QC narratives, lock and plan files, and the deviations log; the analysis driver is pinned by the repository commit.

Conversion note. This v1.0 merges the original paper and its Registered Analysis Report; both are retained in the project's append-only archive as lineage. The independent re-fit code used for §6 verification accompanies the deposit alongside the original bundle.

References

- Egan, C. A., and Lineweaver, C. H. (2010). A Larger Estimate of the Entropy of the Universe. *The Astrophysical Journal* 710(2): 1825–1834. <https://doi.org/10.1088/0004-637X/710/2/1825>.
- Jones, J. C. (2026a). Collapse in Physics: Coherence as Law from Cosmology to Matter (WP02, v1.0). HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/3NYMP>.
- Jones, J. C. (2026b). Empirical Demonstration of S_1 , S_2 , and S_3 in Deployed AI Systems (v1.0). HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/JPXCU>.
- Jones, J. C. (2026c). Records Across Nature, Life, and Mind (v2.0). HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/7H6DY>.
- Jones, J. C. (2026d). Update Integrity Standard (UIS) (v1.0). HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/DWM29>.
- Madau, P., and Dickinson, M. (2014). Cosmic Star-Formation History. *Annual Review of Astronomy and Astrophysics* 52: 415–486. <https://doi.org/10.1146/annurev-astro-081811-125615>.
- Planck Collaboration (Aghanim, N., et al.) (2020). Planck 2018 results. VI. Cosmological parameters. *Astronomy & Astrophysics* 641: A6. <https://doi.org/10.1051/0004-6361/201833910>.
- Profumo, S., Colombo-Murphy, L., Huckabee, G., Diaz Svensson, M., Garg, S., Kollipara, I., and Weber, A. (2024). A New Census of the Universe's Entropy. *arXiv:2412.11282v2*.
- Shankar, F., Weinberg, D. H., and Miralda-Escudé, J. (2009). Self-Consistent Models of the AGN and Black Hole Populations: Duty Cycles, Accretion Rates, and the Mean Radiative Efficiency. *The Astrophysical Journal* 690(1): 20–41. <https://doi.org/10.1088/0004-637X/690/1/20>.

Document status. This is the v1.0 T16 deposit version of Entropy as Record, merging the original paper and internal Registered Analysis Report ("registered" in the project's locked-source sense of §2.2) of 2026-02-17 into a single document under the T16 empirical-family standard. The analysis, criteria, outcomes, and deviations are those of the February 2026 execution; conversion added the claim-architecture hardening, the construct-scope statement of §3.3, the post-hoc descriptives of §5.3–5.4 (labeled as such), and the independent verification of §6. The independent $\rho_{\text{BH}}(z)$ new-data variant is the designed extension. This paper deposits under the T16 empirical umbrella alongside Rice Hysteresis, COGITATE iEEG, and the AI-substrate demonstration (Jones 2026b).

Library note. This paper is part of the Universal Collapse Theory library, published by HoldingLight LLC. It is the physics-wing entry of the T16 empirical corpus. For a reading guide and full architecture, visit universalcollapse.com/roadmap.

AI Disclosure. AI tools were used to assist with manuscript preparation, drafting, organization, and editorial refinement. In this T16 empirical paper, AI assistance additionally extended to data evaluation: analysis-pipeline development and source-extraction tooling in the original February 2026 execution, and, at the July 2026 conversion, the independent re-fit verification of §6 (a from-scratch implementation run against the archived data), the conversion-added checks of §3.3 and §5.5, and regeneration of the figures from the archived artifacts. No AI system is a subject of study in this paper. The underlying theory, structural decisions, analysis, criteria, and conclusions are the author's own, and the author takes full responsibility for the manuscript and its contents.

Citation. Jones, J. C. (2026). Entropy as Record: A Locked-Source Shape Test of an SMBH-Dominated Cosmic Entropy Ledger Across Redshift (v1.0). HoldingLight LLC.

Contact. Inquiries about methodology, factual corrections, or replication results should be directed to contact@universalcollapse.com.