

Local Time Acceleration: A Phenomenological Solution to the Hubble Tension Driven by Kinematic Anomalies in the Local Void

January 11, 2026

Insan Eman
LocalTimeAcceleration@gmail.com

For the benefit of humankind, we stand on the shoulders of giants.

Abstract

The tension between local ($H_0 \approx 73$ km/s/Mpc) and early-universe ($H_0 \approx 67$ km/s/Mpc) measurements of the Hubble constant remains the defining crisis of modern cosmology. In this work, we investigate whether this tension can be ameliorated by a phenomenological “Local Time Acceleration” (LTA) parametrization, which modifies the observed redshift via a bounded, observer-anchored decay function. We fit this model to the Pantheon+SH0ES dataset using the full released covariance matrix and a ladder-anchored likelihood. We report a statistically significant preference for the LTA model over standard Λ CDM ($\Delta\chi^2 \approx 13.65$, corresponding to $\sim 3.7\sigma$), with a Monte Carlo calibrated tail probability $p_{\text{MC}} \approx 8 \times 10^{-4}$ under $N = 10000$ baseline null injections. Tomographic analysis reveals that this signal is driven entirely by supernovae in the local volume ($z < 0.03$), effectively correcting a known “dip” in the Hubble residuals often attributed to peculiar velocities in the Local Void. Crucially, we perform a rigorous Jackknife resampling of all major constituent surveys (CfA, CSP, PS1, etc.) and find the signal is robust to the exclusion of any single telescope group, ruling out instrument-specific calibration systematics. While the effect vanishes at higher redshifts, its persistence and coherence in the local volume suggest that the local Hubble flow exhibits structured kinematic deviations that are better described by a smooth temporal decay than by standard peculiar velocity corrections alone.

Keywords: cosmology, Type Ia supernovae, BAO, local time acceleration, redshift mapping, negentropy

1 Introduction

The persistent mismatch between late-time (local) determinations of the Hubble constant ($H_0 \simeq 73$ km s $^{-1}$ Mpc $^{-1}$) and early-universe inferences under minimal flat Λ CDM ($H_0 \simeq 67$ km s $^{-1}$ Mpc $^{-1}$) has become the central anomaly in precision cosmology [1, 2]. What makes the “Hubble tension” unusually stubborn is that it lives at the intersection of multiple tightly constrained datasets: attempts to modify early-time physics must preserve the CMB acoustic peak structure, while late-time modifications to the expansion history are strongly constrained by the inverse distance ladder (BAO + SNe Ia) once the sound horizon scale is fixed.

A less explored possibility is that part of the discrepancy is not best described as a global change in the background expansion $H(z)$, but rather as a small, observer-linked distortion in the mapping between observed redshift and the cosmological redshift that enters background distances. This is not a semantic distinction: the SH0ES-style distance-ladder inference is dominated by nearby Type Ia supernovae and is therefore maximally sensitive to the detailed kinematics of the local Hubble flow. The observer resides in a structured environment (including a local underdensity often discussed under the umbrella of “Hubble bubble” / “Local Void” phenomenology), and the lowest-redshift portion of the Hubble diagram exhibits coherent residual structure that is not naturally captured by purely stochastic peculiar-velocity scatter alone. If the local Hubble flow contains an additional smooth, causal, line-of-sight effect that is not modeled in the standard pipeline, it should manifest precisely in the redshift range where the distance ladder is anchored and where local-flow systematics are most acute.

In this paper we introduce and test a phenomenological extension of Λ CDM which we call *Local Time Acceleration* (LTA). Operationally, LTA is implemented as an accumulated fractional frequency drift along

the photon trajectory, producing an additive contribution to $\ln(1+z)$ and therefore a multiplicative correction to $(1+z)$:

$$1 + z_{\text{obs}} = (1 + z_{\text{cos}}) \exp[I(\chi)]. \quad (1)$$

The defining modeling choice is that the drift strength is evaluated using an *Earth-retarded lookback time* argument, so that the effect is explicitly causal and tied to the observer’s lightcone rather than to a global function of the scale factor. We further impose boundedness (the drift saturates and becomes negligible beyond a finite activation epoch) and monotonicity of the mapping (so that $z_{\text{cos}}(z_{\text{obs}})$ remains single-valued), making the model both numerically well-posed and falsifiable.

Our goal in this work is deliberately specific: *does a causal, bounded, observer-anchored redshift mapping of the LTA type improve the joint fit to late-time distance–redshift data when those data are analyzed using their released full covariances?* We emphasize what this statement does and does not entail. A detection of a preference for LTA would be an empirical statement about the adequacy of a particular mapping for describing the Hubble diagram and BAO observables; it would not, by itself, uniquely identify a microphysical origin. At the same time, this is not an unconstrained curve-fitting exercise: the mapping is nested ($s_{\text{anchor}} \rightarrow 0$ recovers ΛCDM exactly), it is restricted by causality and boundedness, and it is tested jointly against BAO, whose radial observable depends directly on the redshift Jacobian implied by the mapping.

Our analysis is intentionally conservative in data handling. For SNe Ia we use the Pantheon+SH0ES public release [1, 2], restricting to the `shoes_global` selection and adopting the released *full* covariance matrix, including the nontrivial cross-correlations between calibrators and Hubble-flow supernovae. Because calibrator/Hubble-flow correlations are substantial, we construct a ladder-anchored SN likelihood in which calibrators provide an absolute-distance anchor and the cosmology-sensitive Hubble-flow likelihood is obtained by conditioning on calibrators via a Schur complement (Sec. 6.2). We combine this with a BAO consensus likelihood (Sec. 6.3)[3], treating the sound horizon through a single scaling parameter α_{rd} . Crucially, BAO provides an independent lever arm on the mapping: the LTA redshift Jacobian modifies the radial BAO prediction in a way that a purely SN-driven distortion generally cannot evade.

For inference we compare baseline flat ΛCDM against $\Lambda\text{CDM}+\text{LTA}$ under the same data treatment and numerical safeguards. Because the LTA extension is nested and we impose physical bounds on the LTA amplitude, the null is a boundary point and standard asymptotic χ^2 -difference heuristics are not guaranteed to be reliable. We therefore calibrate the detection significance using a Monte Carlo parametric bootstrap (null injections) under the baseline model, generating mock realizations with the *same* released covariances and refitting both models end-to-end (Sec. 6.6 and Sec. 8). In addition, we perform falsification-oriented robustness tests, including leave-one-group-out (jackknife) refits by telescope/survey grouping to check whether the preference is driven by any single subset, and supplementary cross-validation diagnostics designed for strongly correlated data (App. A).

Roadmap. Section 3 explains the LTA ansatz as a causal frequency-drift mechanism and makes explicit the model’s falsifiability requirements. Section 4 defines the $\Lambda\text{CDM}+\text{LTA}$ mapping, the Earth-retarded time construction, and the Jacobian needed for BAO. Sections 5–6 describe the datasets and the ladder-anchored likelihood used to respect the released covariance structure. Section 7 presents the primary fits and residual diagnostics, including the localization of the improvement in redshift. Section 8 summarizes jackknife robustness and Monte Carlo null calibration. Finally, Section 9 discusses interpretation and independent tests that can confirm or decisively falsify an LTA-like observer-local redshift-mapping effect.

2 Context: Relation to existing approaches to the Hubble tension

The discrepancy between late-time determinations of the Hubble constant and early-universe inferences under minimal flat ΛCDM has generated an exceptionally broad literature. Comprehensive reviews emphasize that proposed resolutions fall into three overlapping categories: (i) unrecognized systematics in one or more datasets and inference pipelines; (ii) new physics that modifies early-universe calibration, typically by reducing the sound horizon r_d and thereby raising the CMB-inferred H_0 ; and (iii) new physics (or unmodeled astrophysics) that modifies the late-time distance–redshift relation or the interpretation of low- z anchors (e.g. dynamical dark energy, modified gravity, or local-structure effects) [5, 6]. Because the observational inputs are heterogeneous (SNe Ia, Cepheids/TRGB, BAO, CMB anisotropies, lensing, time delays, masers,

and standard sirens), it is useful to state explicitly where the LTA framework tested here sits within this landscape, and how it relates to the failure modes that have been most discussed.

At a high level, the LTA hypothesis explored in this paper is *not* an attempt to modify the global expansion history $H(z)$ as a free function, nor is it designed to alter recombination-era physics to shift r_d . Instead, it is a phenomenological proposal that the mapping from observed redshift to the background cosmological redshift entering homogeneous FLRW distances may acquire an additional *observer-anchored, causal, bounded* contribution in the local Universe. This places LTA closest in spirit to the class of explanations in which part (or all) of the SH0ES–Planck discrepancy is driven by effects that are specific to the *local Hubble-flow anchor regime*, including local kinematics, local structure, or other observer-linked biases that are not captured by standard peculiar-velocity modeling.

2.1 Early-universe resolutions: shifting the sound horizon

A large fraction of proposed “new physics” solutions aim to reduce the sound horizon at the drag epoch, r_d , while preserving the excellent fit of Λ CDM to the CMB angular power spectra and other early-time observables. If r_d is smaller than in minimal Λ CDM, then the inverse distance ladder (BAO + SNe Ia), when anchored by CMB physics, can accommodate a larger inferred H_0 . Representative ideas include early dark energy (EDE) episodes that briefly contribute non-negligible energy density near matter–radiation equality [7], additional relativistic degrees of freedom (larger N_{eff}), modifications to recombination, and scenarios in which neutrino physics is altered (e.g. self-interactions) in ways that reshape the CMB damping tail and inferred parameters [8]. The review literature emphasizes that these models are tightly constrained: they must preserve the CMB peak structure, lensing, and BBN consistency, and they are increasingly challenged by the combined CMB+BAO+large-scale-structure dataset hierarchy [5].

The LTA analysis in this paper is largely orthogonal to this program. Operationally, we introduce an explicit sound-horizon scaling parameter α_{rd} in the BAO likelihood, and we test whether an *observer-anchored redshift mapping* can improve the fit in the regime that dominates local distance-ladder anchoring, without requiring a shift of early-time physics. Consequently, if future work finds a residual tension after accounting for observer-local effects, LTA-like mappings could in principle be combined with early-time modifications rather than being in direct competition with them. Conversely, if an early-time solution fully resolves the discrepancy across all late-time probes, it would predict that no additional observer-local mapping distortion is required; this provides a clear empirical discriminant.

2.2 Late-universe modifications: dynamical dark energy and modified gravity

A second broad class of proposals modifies late-time cosmology by changing the expansion history or the growth of structure relative to minimal Λ CDM. These include dynamical dark energy parameterizations ($w \neq -1$, evolving $w(z)$), interacting dark sector models, and modified-gravity theories in which the relation between geometry and stress–energy differs from GR on cosmological scales [5, 6]. Such models are constrained not only by the SNe Ia Hubble diagram but also by BAO (both transverse and radial), weak lensing, redshift-space distortions, cluster abundances, and CMB lensing.

The empirical signature isolated in this paper differs from what these models typically target. Our best-fit LTA profiles are deliberately *bounded and saturating*, and the improvement is tomographically localized to $z \lesssim 0.03$, leaving the best-fit background (H_0, Ω_m) and the higher- z Hubble diagram close to baseline. This behavior is more naturally interpreted as a correction to the *local* mapping between z_{obs} and the effective redshift argument used in distances, rather than as a genuine modification of $H(z)$ over a broad redshift lever arm. Accordingly, LTA does not aim to replace late-time dark energy phenomenology globally; rather, it tests whether the low- z ladder anchor regime exhibits a structured deviation that is better captured by a smooth, causal, observer-local mapping than by standard local-flow scatter models.

2.3 Systematics and independent late-time anchors

A persistent theme in the literature is that the SH0ES–Planck discrepancy could reflect residual systematics in one or more components of the distance ladder (Cepheid calibration, SN Ia standardization, selection effects, dust, host correlations) or in early-universe inference. For this reason, independent determinations

of H_0 have been heavily emphasized. These include TRGB-calibrated ladders [11], strong-lensing time-delay cosmography [12, 13], geometric megamaser distances [14], and gravitational-wave standard sirens [15]. A key point from review work is that the most informative cross-checks are those that change both the astrophysical systematics and the inference architecture, rather than merely reprocessing the same ladder ingredients [5, 6].

The contribution of this paper is not to adjudicate between all of these external H_0 anchors. Instead, we take the Pantheon+SH0ES public release at face value and analyze it using the published *full* covariance structure (including calibrator–Hubble-flow cross-correlations), and we ask a narrower question: within that released statistical model, is there evidence for a local, smoothly varying mapping distortion that improves the Hubble-flow residual structure, and does that distortion survive falsification-oriented robustness tests (jackknife, null injections) and remain consistent with BAO once the mapping Jacobian is propagated? In that sense, LTA should be read as an *internal-consistency hypothesis test* on a specific late-time dataset combination, rather than as an attempt to supersede the full cross-probe literature on H_0 .

2.4 Local structure, cosmic variance, and redshift interpretation

A third major theme, directly relevant to the present work, is whether the local environment of the observer can bias low- z determinations of H_0 . This includes (i) the possibility that the observer resides in a local underdensity or otherwise atypical region, affecting inferred distances and velocities (“Hubble bubble” / Local Void phenomenology), and (ii) irreducible cosmic variance in the locally measured Hubble flow even in a perfect experiment [9]. Detailed analyses in the SN context have generally concluded that a simple spherically symmetric local void of realistic amplitude is insufficient, by itself, to reconcile SH0ES with Planck [10], though local-flow modeling and sample selection remain active topics.

LTA is deliberately aligned with the *spirit* of this class of explanations—namely, that the anchor regime of the distance ladder may encode structured local effects—but it differs in its implementation. Rather than modeling a particular density profile or velocity field, we introduce a nested, bounded, monotonic mapping in which photons acquire an accumulated contribution to $\ln(1+z)$ parameterized as a line-of-sight integral and evaluated with an Earth-retarded time argument. This choice yields a concrete, testable prediction beyond a generic “local effect” statement: the mapping implies a redshift Jacobian $dz_{\text{obs}}/dz_{\text{cos}}$ that directly affects radial BAO. In other words, LTA is not merely an empirical correction to the SN Hubble diagram; it is a mapping whose consistency can be independently checked in observables that depend on redshift derivatives.

2.5 What is novel in the present analysis

Within the broader literature, the main methodological distinctions of this work are:

1. **A specific observer-anchored redshift-mapping ansatz.** We implement the hypothesis as a multiplicative correction to $(1+z)$ (additive in $\ln(1+z)$), accumulated along the photon trajectory, with explicit bounded activation and an Earth-retarded time construction to enforce causal bookkeeping.
2. **Full released-covariance treatment with calibrator cross-correlations respected.** The Pantheon+SH0ES release contains nontrivial calibrator–Hubble-flow covariance coupling; we construct a ladder-anchored likelihood via a Schur complement so that calibrators act as an absolute-distance anchor without discarding the published correlation structure.
3. **Significance calibration under a boundary-nested null.** Because the LTA amplitude is bounded ($s_{\text{anchor}} \geq 0$), the null is a boundary point; we therefore calibrate $\Delta\chi^2$ using parametric bootstrap null injections that preserve the released covariance and rerun the full refitting pipeline.
4. **A joint SN+BAO test that propagates the mapping Jacobian.** BAO is not used only as a late-time distance lever arm; it provides a falsification channel because radial BAO depends explicitly on the redshift Jacobian implied by the mapping.
5. **Tomographic localization of the effect.** The preference is driven by a narrow low- z shell ($z \lesssim 0.03$), consistent with a local kinematic/mapping effect and inconsistent with a broad modification of the late-time expansion law.

Taken together, these points place LTA in a well-defined position relative to the existing solution space reviewed in [5]: it is a *phenomenological, observer-local mapping extension* designed to address structured low- z Hubble-flow residuals while remaining explicitly testable against BAO through the Jacobian constraint. If the preference survives further external scrutiny (alternative SN compilations, alternative flow models, sky-split tests, independent anchors), then it motivates a broader program of checking whether part of the Hubble tension is best understood as a local redshift-interpretation bias rather than as a single global modification of early- or late-time cosmology.

3 Conceptual Framework: Local Time Acceleration

The motivating hypothesis behind this work is unconventional, so it behoves us to be explicit about what is being proposed at the observational level, what requirements we impose to keep it well-posed, and what claims the present analysis does and does not make. Our aim here is not to argue for a specific microphysical origin, but to define a concrete and falsifiable phenomenological mapping between cosmological redshift and observed redshift, and then to test whether that mapping is preferred by late-time distance–redshift data analyzed using their released covariances.

3.1 LTA as an accumulated, observer-local contribution to $\ln(1+z)$

In standard FLRW cosmology, the observed redshift is determined by the ratio of scale factors, $1+z = a(t_{\text{obs}})/a(t_{\text{emit}})$, up to conventional local corrections (peculiar velocities, gravitational potentials, etc.) that are treated as separate effects. LTA modifies this assumption by introducing an additional *line-of-sight accumulated* fractional frequency drift. The observational statement is that photons acquire a small multiplicative factor in $(1+z)$ beyond the cosmological contribution:

$$1 + z_{\text{obs}} = (1 + z_{\text{cos}}) \exp[I(\chi)]. \quad (2)$$

Writing the effect in $\ln(1+z)$ is not a stylistic choice: it is the natural representation for a process that composes multiplicatively and accumulates additively along a trajectory. The quantity $I(\chi)$ is constructed as the line-of-sight integral of an effective drift-rate field $s(\chi)$ (with units of a Hubble rate), so that s/c acts as a fractional frequency-drift per unit comoving distance (Sec. 4.2).

This is not a coordinate relabeling. A pure coordinate transformation cannot change dimensionless local observables without additional physical content. In contrast, LTA is defined precisely as an additional physical contribution to the observed frequency ratio relative to the baseline Λ CDM prediction, parameterized in a way that is nested, bounded, and testable.

3.2 Why an Earth-retarded time argument appears

The distinctive structural element of the LTA ansatz is that the drift-rate profile is evaluated as a function of an *Earth-retarded lookback time* t_{ret} rather than as a generic function of cosmic time or scale factor. The purpose of this construction is causal bookkeeping: if the source of the effect is local to the observer’s environment (broadly construed), then any influence on the observed lightcone must be consistent with propagation at or below the speed of light. We therefore tie the effective drift strength at a point along the photon path to a retarded-time assignment based on the observer’s worldline and the background geometry (Sec. 4.3).

Importantly, “Earth-sourced” here should be read as *observer-anchored* rather than as a commitment to a particular terrestrial mechanism. The present analysis does not require (and does not assume) that the relevant state variable is literally biological. Instead, the Earth-retarded construction is a compact proxy for a wide class of observer-local causal influences whose effective strength depends on the observer’s local history or environment. The empirical question tested below is therefore: does an *observer-anchored, causal, bounded* mapping of this type improve the description of the data?

3.3 Bounded activation and controlled degrees of freedom

To remain physically viable and empirically falsifiable, the LTA contribution must not grow without limit. We therefore impose a bounded activation through a function $g(t) \in [0, 1]$ such that the effect decays to zero at a finite cutoff epoch t_{life} (Sec. 4.4). This ensures that the additional drift saturates and becomes negligible for sufficiently large lookback times, preventing the model from contaminating high-redshift cosmology where early-universe constraints are exceptionally tight.

A second requirement is that the redshift mapping remain monotonic:

$$\frac{dz_{\text{obs}}}{dz_{\text{cos}}} > 0, \quad (3)$$

so that the inverse relation $z_{\text{cos}}(z_{\text{obs}})$ is single-valued and the Hubble diagram remains well-defined (Sec. 4.5). This is not merely a numerical convenience: non-monotonic mappings would imply multi-valued distance predictions at fixed observed redshift and would be directly incompatible with the interpretation of both SNe Ia and BAO as functions of redshift.

Finally, the model is explicitly *nested*. In the zero-amplitude limit ($s_{\text{anchor}} \rightarrow 0$), the integral $I(\chi) \rightarrow 0$ and the mapping reduces exactly to baseline Λ CDM. This nesting allows us to compare Λ CDM and Λ CDM+LTA at fixed data treatment using standard likelihood-ratio-style test statistics, while still calibrating their null distribution via Monte Carlo because the null lies on a boundary (Sec. 6.6).

3.4 Why BAO matters for falsification

Although the primary low-redshift leverage comes from SNe Ia, BAO plays a crucial conceptual role in this test. The LTA mapping implies not only a shift in the distance–redshift relation but also a redshift Jacobian $dz_{\text{obs}}/dz_{\text{cos}}$. Radial BAO observables depend explicitly on $d\chi/dz$ and therefore directly probe the Jacobian induced by the mapping (Sec. 6.3). This provides an independent check against mappings that might otherwise fit SNe Ia by distorting redshift in a way that would be inconsistent with a standard ruler measured in the radial direction. For this reason, joint SN+BAO fitting is an essential part of rendering the LTA hypothesis falsifiable with late-time data.

3.5 What this paper claims and what it does not

The restricted empirical claim supported by this paper (if the preference persists under the full set of robustness checks) is: *within the adopted data treatment and model family, an observer-anchored, causal, bounded LTA mapping provides an improved fit to late-time distance–redshift data relative to baseline Λ CDM.*

This is not a claim of unique microphysical identification. The same phenomenology could in principle arise from unmodeled local kinematics, residual calibration structure, or other observer-local effects not captured in the baseline pipeline. Accordingly, our strategy is to (i) implement a precise mapping with clear causal and boundedness requirements, (ii) fit it using the released full covariance structure (including calibrator/HF cross-covariance), and (iii) report falsification-oriented diagnostics that test whether the preference is distributed across the data, robust to leave-out refits, and unlikely under baseline null injections.

In this sense, LTA is treated here as a *testable phenomenological hypothesis*: the cosmology data are used to evaluate whether the mapping is statistically preferred, while the interpretation of its ultimate origin is left to independent constraints and future targeted tests (Sec. 9).

4 Model definition: Λ CDM background plus LTA

This section specifies the forward model used throughout: a standard flat Λ CDM background for distances and lookback times, coupled to an LTA redshift mapping in which the observed redshift z_{obs} differs from the background cosmological redshift z_{cos} by a bounded, causal, observer-anchored accumulated drift.

4.1 Baseline flat Λ CDM background

We assume a spatially flat Λ CDM background with parameters (H_0, Ω_m) and define

$$E(z) \equiv \sqrt{\Omega_m(1+z)^3 + (1 - \Omega_m)}, \quad H(z) \equiv H_0 E(z). \quad (4)$$

The line-of-sight comoving distance is

$$\chi(z) = \frac{c}{H_0} \int_0^z \frac{dz'}{E(z')}, \quad (5)$$

and the standard proper lookback time is

$$t_{\text{lb}}(z) = \frac{1}{H_0} \int_0^z \frac{dz'}{(1+z')E(z')}. \quad (6)$$

In the implementation, χ is evaluated in Mpc using $c = 299\,792.458 \text{ km s}^{-1}$ and H_0 in $\text{km s}^{-1} \text{ Mpc}^{-1}$; t_{lb} is converted to Gyr for use in the LTA activation.

4.2 LTA as an accumulated fractional frequency drift

We model *Local Time Acceleration* (LTA) as an additional, line-of-sight accumulated, *multiplicative* contribution to $(1+z)$, which is naturally represented as an additive term in $\ln(1+z)$.

We define the dimensionless accumulated integral

$$I(\chi) \equiv \frac{1}{c} \int_0^\chi s(\chi)' d\chi', \quad (7)$$

where $s(\chi)$ has units of $\text{km s}^{-1} \text{ Mpc}^{-1}$ (the same dimensions as a Hubble rate). Operationally, $s(\chi)/c$ acts as a fractional frequency-drift per unit comoving distance.

The mapping between cosmological redshift and observed redshift is

$$1 + z_{\text{obs}} = (1 + z_{\text{cos}}) \exp[I(\chi)(z_{\text{cos}})], \quad (8)$$

where $I(\chi)(z_{\text{cos}}) \equiv I(\chi(z_{\text{cos}}))$. For $I(\chi) \ll 1$ this reduces to

$$z_{\text{obs}} \approx z_{\text{cos}} + (1 + z_{\text{cos}})I(\chi), \quad (9)$$

but we retain the exponential form to preserve positivity, improve numerical conditioning, and make the nested limit manifest: $s_{\text{anchor}} \rightarrow 0$ implies $I(\chi) \rightarrow 0$ and recovers baseline Λ CDM exactly.

Differentiating Eq. (8) gives an identity used repeatedly for Jacobians and diagnostics:

$$\frac{dI(\chi)}{dz_{\text{cos}}} = \frac{1}{c} s(\chi)(z_{\text{cos}}) \frac{d\chi}{dz_{\text{cos}}} = \frac{s(\chi)(z_{\text{cos}})}{H(z_{\text{cos}})}. \quad (10)$$

4.3 Earth-retarded lookback time construction

To encode the motivating “observer-sourced” (Earth-anchored) character in a causally consistent way, we evaluate the activation along the photon path using an *Earth-retarded lookback time* argument $t_{\text{ret}}(\chi)$.

For a given line-of-sight distance χ (corresponding to z_{cos}), we define an “Earth emission” redshift z_e implicitly by

$$\chi(z_e) = 2\chi, \quad (11)$$

and then set

$$t_{\text{ret}}(\chi) \equiv t_{\text{lb}}(z_e). \quad (12)$$

Intuitively, $2\chi/c$ represents a round-trip light-travel construction: it assigns to each point on the line of sight a retarded time at which an influence originating locally could causally reach that point in the background geometry.¹

¹At low redshift, where $\chi \approx (c/H_0)z$ and $t_{\text{lb}} \approx H_0^{-1}z$, Eq. (11) implies $z_e \approx 2z$ and $t_{\text{ret}} \approx 2t_{\text{lb}}$, making the causal interpretation explicit.

4.4 Bounded activation and parameterization of $s(\chi)$

We parameterize the drift-rate profile as a bounded activation scaled by an amplitude defined at an explicit *anchor time* t_{anchor} :

$$s(\chi) = s_{\text{anchor}} \frac{g(t_{\text{ret}}(\chi))}{g(t_{\text{anchor}})}. \quad (13)$$

Here $g(t) \in [0, 1]$ is an activation function of Earth-retarded lookback time, normalized so that $g(0) = 1$ and $g(t) = 0$ for $t \geq t_{\text{life}}$, where t_{life} is a finite cutoff epoch. The normalization by $g(t_{\text{anchor}})$ ensures that the reported parameter s_{anchor} corresponds to the drift rate at the chosen anchor time: $s(t_{\text{anchor}}) \equiv s_{\text{anchor}}$.

Anchor-time choice (leakage control). To avoid any dependence of the LTA normalization on the calibrator subset, we compute t_{anchor} using only the Hubble-flow subset (excluding calibrators). Concretely, we define t_{anchor} as the 5th percentile of the t_{ret} distribution over Hubble-flow objects. This rule is stable under resampling and keeps the normalization tied to the Hubble-flow regime that drives the cosmological fit.

Power-law activation family (fiducial choice). In this work we adopt a smooth power-law decay. For $0 \leq t \leq t_{\text{life}}$ define

$$r(t) = \left(1 + \frac{t}{B}\right)^{-p}, \quad (14)$$

with timescale $B > 0$ (in Gyr) and slope $p \geq 0$, and set

$$g(t) = \frac{r(t) - r(t_{\text{life}})}{1 - r(t_{\text{life}})}. \quad (15)$$

This guarantees $g(0) = 1$ and $g(t_{\text{life}}) = 0$, yielding a continuous, bounded activation with a tunable early-time scale (B) and tail steepness (p).

Fiducial degrees of freedom (what is actually fit). Because LTA enters through a line-of-sight integral, there are strong degeneracies between the overall amplitude and detailed curve-shape parameters. Our primary analysis therefore fixes the activation *shape* to representative smooth values and fits only the amplitude s_{anchor} as the additional LTA degree of freedom beyond $(H_0, \Omega_m, \alpha_{\text{rd}})$. Unless otherwise stated, we use fixed (t_{life}, B, p) values in the power-law family and treat the Λ CDM+LTA model as a one-parameter nested extension in model-selection metrics. (Where we report exploratory runs with additional free shape parameters, we state that explicitly.)

Optional modulation exponents (fixed in this work). The analysis code supports additional multiplicative modulation exponents (e.g. “complexity” and “life” weights) that rescale the activation family. In the fiducial results emphasized in this paper, these exponents are fixed to unity (via runtime flags) so that they do not contribute additional degrees of freedom and do not affect AIC/BIC accounting.

4.5 Inverse mapping and redshift Jacobian

Data are indexed by z_{obs} (or by a measured redshift proxy such as z_{HD} for SNe), whereas background distances are naturally functions of z_{cos} . Therefore, for each observed redshift z_{obs} we obtain $z_{\text{cos}}(z_{\text{obs}})$ by numerically inverting Eq. (8) under the monotonicity condition

$$\frac{dz_{\text{obs}}}{dz_{\text{cos}}} > 0. \quad (16)$$

Differentiating Eq. (8) gives

$$\frac{dz_{\text{obs}}}{dz_{\text{cos}}} = \exp[I(\chi)] \left[1 + (1 + z_{\text{cos}}) \frac{dI(\chi)}{dz_{\text{cos}}} \right] = \exp[I(\chi)] \left[1 + (1 + z_{\text{cos}}) \frac{s(\chi)(z_{\text{cos}})}{H(z_{\text{cos}})} \right], \quad (17)$$

where the final equality uses Eq. (10). This Jacobian enters directly in radial BAO predictions and serves as an important diagnostic of mapping regularity.

4.6 Distances and effective expansion observables as functions of z_{obs}

The comoving distance as a function of observed redshift is evaluated as

$$\chi(z_{\text{obs}}) \equiv \chi(z_{\text{cos}}(z_{\text{obs}})). \quad (18)$$

For radial derivatives,

$$\frac{d\chi}{dz_{\text{obs}}} = \frac{d\chi}{dz_{\text{cos}}} \frac{dz_{\text{cos}}}{dz_{\text{obs}}} = \frac{c}{H(z_{\text{cos}})} \left(\frac{dz_{\text{obs}}}{dz_{\text{cos}}} \right)^{-1}. \quad (19)$$

Accordingly, the *effective* Hubble rate inferred from the $z_{\text{obs}}-\chi$ relation is

$$H_{\text{eff}}(z_{\text{obs}}) \equiv \frac{c}{d\chi/dz_{\text{obs}}} = H(z_{\text{cos}}) \left(\frac{dz_{\text{obs}}}{dz_{\text{cos}}} \right). \quad (20)$$

Here “effective” means “as inferred from the z_{obs} -distance mapping in the presence of LTA”; it does *not* redefine the background expansion law $H(z)$.

In the joint SN+BAO analysis below, transverse distances depend primarily on $\chi(z_{\text{obs}})$, while radial observables depend explicitly on the Jacobian factor $dz_{\text{obs}}/dz_{\text{cos}}$ through $H_{\text{eff}}(z_{\text{obs}})$ (Sec. 6.3).

5 Data

5.1 Pantheon+SH0ES: calibrators + Hubble-flow subset with full covariance

We use the Pantheon+SH0ES public release with the corrected SN Ia magnitude observable

$$y \equiv m_b^{\text{corr}}, \quad (21)$$

and the accompanying *full* covariance matrix (including cross-correlations between calibrators and Hubble-flow objects). Our selected sample is the SH0ES global subset consisting of

$$N_{\text{SN}} = 354 \text{ total SNe, with } N_{\text{cal}} = 77 \text{ calibrators, } N_{\text{HF}} = 277 \text{ Hubble-flow objects.} \quad (22)$$

The selected Hubble-diagram redshifts satisfy

$$z_{\text{HD}} \in [0.00122, 0.14898], \quad (23)$$

and we retain the standard distinction between heliocentric redshift z_{HEL} (for the $(1 + z_{\text{HEL}})$ luminosity-distance prefactor) and the Hubble-diagram redshift z_{HD} (for the cosmology-dependent distance).

For calibrators, we use the externally determined distance moduli provided by the release (e.g. Cepheid-based; column CEPH_DIST in our implementation), spanning approximately

$$\mu_{\text{cal}} \in [29.177, 34.526], \quad (24)$$

and we keep calibrators present in every fit (and in every refit / cross-validation split when relevant) to preserve the absolute-distance anchor.

Covariance diagnostics. The released covariance is used in full (file-provided mode). As a compact indicator of correlation strength, we report an “effective dimension”

$$N_{\text{eff}} \equiv \frac{(\sum_i \lambda_i)^2}{\sum_i \lambda_i^2}, \quad (25)$$

where $\{\lambda_i\}$ are eigenvalues of the SN covariance; for the selected sample we obtain $N_{\text{eff}} \approx 149.5$ (out of 354), illustrating that correlations materially reduce the number of effectively independent modes.

5.2 BAO consensus compilation

We use the DESI 2024 Gaussian BAO compilation (GCcomb)[3], which provides D_V/r_d , D_M/r_d , and D_H/r_d constraints in multiple tracer/redshift bins, with a provided covariance.

We jointly fit a consensus BAO compilation providing transverse and radial distance information at discrete redshifts. The BAO data vector is treated as Gaussian with a provided covariance matrix. In our parameterization we introduce a single scaling parameter α_{rd} to absorb uncertainty in the sound horizon:

$$\alpha_{\text{rd}} \equiv \frac{r_d}{r_{d,\text{fid}}}. \quad (26)$$

To match the convention used in our code, the BAO observables are taken in the rescaled form

$$D_M^{(\text{data})}(z) \equiv \left(\frac{D_M(z)}{r_d} \right) r_{d,\text{fid}}, \quad H^{(\text{data})}(z) \equiv (H(z) r_d) \frac{1}{r_{d,\text{fid}}}, \quad (27)$$

so that the theory predictions become simple multiplicative scalings by α_{rd} (see Sec. 6.3). We assume the BAO and SN datasets are independent so that the joint likelihood covariance is block diagonal between SN and BAO.

6 Likelihood and inference

6.1 Supernova distance model and residual vector

For each SN, we compute the luminosity distance using the comoving distance evaluated at the cosmological redshift corresponding to the *observed* Hubble-diagram redshift:

$$d_L = (1 + z_{\text{HEL}}) \chi(z_{\text{cos}}(z_{\text{HD}})), \quad \mu_{\text{pred}} = 5 \log_{10} \left(\frac{d_L}{\text{Mpc}} \right) + 25. \quad (28)$$

The corrected magnitude observable is modeled as

$$m_{b,i}^{\text{corr}} = \mu_{\text{pred}}(z_i; \boldsymbol{\theta}) + M + \epsilon_i, \quad (29)$$

where M is an intercept (absolute magnitude) and ϵ is a correlated Gaussian error with the released covariance.

Let $\mathbf{y}$ be the full SN data vector (stacked m_b^{corr}) and $\boldsymbol{\mu}(\boldsymbol{\theta})$ the corresponding model vector (stacked μ_{pred}). The raw residual vector at fixed M is

$$\mathbf{r}(\boldsymbol{\theta}, M) \equiv \mathbf{y} - \boldsymbol{\mu}(\boldsymbol{\theta}) - M\mathbf{1}. \quad (30)$$

6.2 Ladder-anchored SN likelihood with calibrators and full cross-covariance

A key feature of the Pantheon+SH0ES release is that calibrators and Hubble-flow objects are not independent: the covariance includes nontrivial cross-blocks. We therefore construct an *anchored* likelihood that (i) retains the full covariance structure, (ii) uses calibrators as an absolute-distance anchor, and (iii) yields a Hubble-flow likelihood that can be evaluated repeatedly during cosmological optimization.

Calibrator / Hubble-flow partition. We partition the SN vector into calibrators (C) and Hubble-flow objects (H):

$$\mathbf{y} = \begin{bmatrix} \mathbf{y}_C \\ \mathbf{y}_H \end{bmatrix}, \quad \boldsymbol{\mu}(\boldsymbol{\theta}) = \begin{bmatrix} \boldsymbol{\mu}_C \\ \boldsymbol{\mu}_H(\boldsymbol{\theta}) \end{bmatrix}, \quad \mathbf{C} = \begin{bmatrix} \mathbf{C}_{CC} & \mathbf{C}_{CH} \\ \mathbf{C}_{HC} & \mathbf{C}_{HH} \end{bmatrix}. \quad (31)$$

For calibrators we take $\boldsymbol{\mu}_C$ from the external distance ladder (e.g. Cepheid-based distance moduli provided by the release), so $\boldsymbol{\mu}_C$ is fixed across cosmological models.

Step 1: calibrator-only posterior for M . Using calibrators alone, we compute the Gaussian posterior for the intercept:

$$\sigma_M^2 = (\mathbf{1}_C^\top \mathbf{C}_{CC}^{-1} \mathbf{1}_C)^{-1}, \quad M_{\text{cal}} = \sigma_M^2 \mathbf{1}_C^\top \mathbf{C}_{CC}^{-1} (\mathbf{y}_C - \boldsymbol{\mu}_C). \quad (32)$$

In the representative run reported in Sec. 7, this yields

$$M_{\text{cal}} = -19.248885 \pm 0.030106 \quad (1\sigma, \text{ calibrators only}). \quad (33)$$

This posterior is then used as an informative Gaussian prior on M when evaluating the Hubble-flow likelihood.

Step 2: condition Hubble-flow on calibrators (Schur complement). Define the cross-block regression operator

$$\mathbf{K} \equiv \mathbf{C}_{HC} \mathbf{C}_{CC}^{-1}, \quad (34)$$

and the conditional covariance (Schur complement)

$$\mathbf{S} \equiv \mathbf{C}_{HH} - \mathbf{C}_{HC} \mathbf{C}_{CC}^{-1} \mathbf{C}_{CH}. \quad (35)$$

The Hubble-flow data conditioned on calibrators is represented by the shifted residual vector

$$\tilde{\mathbf{d}}(\boldsymbol{\theta}) \equiv (\mathbf{y}_H - \boldsymbol{\mu}_H(\boldsymbol{\theta})) - \mathbf{K}(\mathbf{y}_C - \boldsymbol{\mu}_C). \quad (36)$$

At fixed intercept M , the conditional model is linear in M :

$$\tilde{\mathbf{d}}(\boldsymbol{\theta}) = \mathbf{a} M + \boldsymbol{\eta}, \quad \mathbf{a} \equiv \mathbf{1}_H - \mathbf{K} \mathbf{1}_C, \quad \boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \mathbf{S}). \quad (37)$$

Step 3: marginalize M using the calibrator posterior. We combine the conditional Gaussian likelihood in Eq. (37) with the calibrator posterior $M \sim \mathcal{N}(M_{\text{cal}}, \sigma_M^2)$ and analytically integrate M . Up to additive constants independent of $\boldsymbol{\theta}$, this produces an effective Hubble-flow negative log-likelihood (written as a χ^2 contribution):

$$\chi_{\text{HF}}^2(\boldsymbol{\theta}) \equiv \tilde{\mathbf{d}}^\top \mathbf{S}^{-1} \tilde{\mathbf{d}} - \frac{\left(\mathbf{a}^\top \mathbf{S}^{-1} \tilde{\mathbf{d}} + M_{\text{cal}} / \sigma_M^2 \right)^2}{\mathbf{a}^\top \mathbf{S}^{-1} \mathbf{a} + 1 / \sigma_M^2}. \quad (38)$$

The calibrator term itself contributes a constant (model-independent) offset once M_{cal} and σ_M are fixed, so the cosmology-sensitive SN contribution in the optimizer is $\chi_{\text{HF}}^2(\boldsymbol{\theta})$.

Summary. The SN objective used in fitting can be viewed as

$$\chi_{\text{SN}}^2(\boldsymbol{\theta}) = \chi_{\text{cal, const}}^2 + \chi_{\text{HF}}^2(\boldsymbol{\theta}), \quad (39)$$

where $\chi_{\text{cal, const}}^2$ is constant across models, and $\chi_{\text{HF}}^2(\boldsymbol{\theta})$ incorporates both the calibrator anchor and the full cross-covariance via the Schur complement.

6.3 BAO likelihood with LTA Jacobian

Because BAO measurements are indexed by observed redshift z_{obs} while comoving distances are functions of z_{cos} , the LTA mapping affects both transverse and radial predictions.

Transverse BAO. Using the rescaled data convention described in Sec. 5.2, the model prediction is

$$D_M^{(\text{model})}(z_{\text{obs}}) = \frac{\chi(z_{\text{cos}}(z_{\text{obs}}))}{\alpha_{\text{rd}}}. \quad (40)$$

Radial BAO. Radial BAO depends on $d\chi/dz$ and therefore requires the redshift Jacobian. Using Eq. (20), the effective Hubble rate at observed redshift is

$$H_{\text{eff}}(z_{\text{obs}}) = H(z_{\text{cos}}(z_{\text{obs}})) \left(\frac{dz_{\text{obs}}}{dz_{\text{cos}}} \right), \quad (41)$$

and therefore

$$H^{(\text{model})}(z_{\text{obs}}) = H_{\text{eff}}(z_{\text{obs}}) \alpha_{\text{rd}} = H(z_{\text{cos}}(z_{\text{obs}})) \left(\frac{dz_{\text{obs}}}{dz_{\text{cos}}} \right) \alpha_{\text{rd}}, \quad (42)$$

where $dz_{\text{obs}}/dz_{\text{cos}}$ is given by Eq. (17).

BAO χ^2 . Let $\mathbf{y}_{\text{BAO}}$ be the stacked BAO data vector and $\mathbf{m}_{\text{BAO}}(\boldsymbol{\theta})$ the stacked model vector obtained from Eqs. (40)–(42). With BAO covariance $\mathbf{C}_{\text{BAO}}$, the contribution is

$$\chi_{\text{BAO}}^2(\boldsymbol{\theta}) = (\mathbf{y}_{\text{BAO}} - \mathbf{m}_{\text{BAO}}(\boldsymbol{\theta}))^T \mathbf{C}_{\text{BAO}}^{-1} (\mathbf{y}_{\text{BAO}} - \mathbf{m}_{\text{BAO}}(\boldsymbol{\theta})). \quad (43)$$

6.4 Correlated Planck-chain Gaussian prior

Optionally (and in the representative run reported here), we include a correlated Gaussian prior derived from a Planck-chain subset over the parameters

$$\boldsymbol{\phi} \equiv (H_0, \Omega_m, \alpha_{\text{rd}}). \quad (44)$$

The prior is

$$\chi_{\text{prior}}^2(\boldsymbol{\phi}) = (\boldsymbol{\phi} - \boldsymbol{\mu}_{\text{P}})^T \boldsymbol{\Sigma}_{\text{P}}^{-1} (\boldsymbol{\phi} - \boldsymbol{\mu}_{\text{P}}), \quad (45)$$

with mean

$$\boldsymbol{\mu}_{\text{P}} = \begin{bmatrix} 67.35987729 \\ 0.315284 \\ 0.99533613 \end{bmatrix}, \quad (46)$$

and covariance

$$\boldsymbol{\Sigma}_{\text{P}} = \begin{bmatrix} 2.8666 \times 10^{-1} & -3.8905 \times 10^{-3} & 6.6134 \times 10^{-4} \\ -3.8905 \times 10^{-3} & 5.3718 \times 10^{-5} & -1.0157 \times 10^{-5} \\ 6.6134 \times 10^{-4} & -1.0157 \times 10^{-5} & 3.2006 \times 10^{-6} \end{bmatrix}, \quad (47)$$

corresponding to $(\sigma_{H_0}, \sigma_{\Omega_m}, \sigma_{\alpha_{\text{rd}}}) = (0.5354, 0.007329, 0.001789)$ and the reported correlation coefficients [4].

6.5 Total objective and parameter sets

The total objective minimized is

$$\chi_{\text{tot}}^2(\boldsymbol{\theta}) = \chi_{\text{SN}}^2(\boldsymbol{\theta}) + \chi_{\text{BAO}}^2(\boldsymbol{\theta}) + \chi_{\text{prior}}^2(\boldsymbol{\phi}). \quad (48)$$

Baseline Λ CDM parameters:

$$\boldsymbol{\theta}_{\Lambda\text{CDM}} = (H_0, \Omega_m, \alpha_{\text{rd}}). \quad (49)$$

LTA extension parameters (power-law family):

$$\boldsymbol{\theta}_{\text{LTA}} = (H_0, \Omega_m, \alpha_{\text{rd}}, t_{\text{life}}, s_{\text{anchor}}, B, p), \quad (50)$$

with t_{life} fixed in the baseline analysis to the chosen life-epoch cutoff (e.g. $t_{\text{life}} = 3.8$ Gyr in the representative run). The model is nested: setting $s_{\text{anchor}} = 0$ yields $I(\chi) \equiv 0$ and recovers baseline Λ CDM exactly.

6.6 Optimization and test statistic

We obtain best fits by direct numerical minimization of Eq. (48) using a derivative-free method (Powell-type directional search) subject to physical bounds (e.g. $s_{\text{anchor}} \geq 0$, $B > 0$, $p \geq 0$) and numerical safeguards for monotonic inversion of $z_{\text{obs}}(z_{\text{cos}})$.

Our primary global comparison statistic is the improvement

$$\Delta\chi^2 \equiv \chi_{\Lambda\text{CDM}}^2 - \chi_{\Lambda\text{CDM}+\text{LTA}}^2, \quad (51)$$

with $\Delta\chi^2 > 0$ favoring LTA. In the representative run, we find $\Delta\chi^2 \simeq 13.7$.

To calibrate the null distribution of $\Delta\chi^2$ beyond asymptotic χ^2 approximations, we also perform Monte Carlo null injections (parametric bootstrap) under the baseline model using the same covariance and refitting both models. In an $N = 10000$ null injection run, we obtain a Monte Carlo tail probability $p_{\text{MC}} \equiv P(\Delta\chi_{\text{null}}^2 \geq \Delta\chi_{\text{obs}}^2) \approx 8 \times 10^{-4}$.

Fiducial analysis choices. Unless otherwise stated, all quoted results use: (i) the Pantheon+SH0ES `shoes_global` selection ($N_{\text{SN}} = 354$) with the released full covariance, analyzed via the ladder-anchored Schur-complement likelihood of Sec. 6.2 (calibrators retained as the absolute anchor and the intercept marginalized using the calibrator posterior); (ii) the BAO consensus compilation of Sec. 5.2 with a single sound-horizon scaling parameter α_{rd} , evaluated at observed redshift using the LTA mapping and Jacobian of Sec. 6.3; (iii) the correlated Planck-chain Gaussian prior on $(H_0, \Omega_m, \alpha_{\text{rd}})$ described in Sec. 6.4 when explicitly stated; and (iv) the LTA(powerlaw) activation family with fixed $t_{\text{life}} = 3.8$ Gyr, bounds $s_{\text{anchor}} \geq 0$, $B > 0$, $p \geq 0$, and t_{anchor} recomputed from the Hubble-flow subset only to avoid leakage. Numerical inversion enforces $dz_{\text{obs}}/dz_{\text{cos}} > 0$ over the redshift domain of interest.

7 Primary Results

This section reports the primary empirical outcomes of the fiducial ΛCDM and $\Lambda\text{CDM}+\text{LTA}(\text{powerlaw})$ fits, emphasizing (i) the global goodness-of-fit improvement under the *full released covariance* and ladder-anchored likelihood, and (ii) where in redshift space that improvement is sourced. We then present reconstructed mapping diagnostics (the LTA profile, $\Delta \ln(1+z)$ saturation, and the redshift Jacobian) and the BAO consistency checks implied by the mapping.

7.1 Fiducial analysis configuration and strict nesting sanity check

Unless otherwise stated, results refer to run-tag 20260112-013808, using: (i) the Pantheon+SH0ES `shoes_global` selection ($N_{\text{SN}} = 354$, with $N_{\text{cal}} = 77$ calibrators and $N_{\text{HF}} = 277$ Hubble-flow SNe; $z_{\text{HD}} \in [0.00122, 0.14898]$); (ii) the ladder-anchored, full cross-covariance SN likelihood of Sec. 6.2 (calibrator posterior for M and Schur-complement conditioning); (iii) the BAO consensus likelihood evaluated at observed redshift using the LTA mapping and Jacobian (Sec. 6.3); and (iv) the correlated Planck-chain Gaussian prior on $(H_0, \Omega_m, \alpha_{\text{rd}})$ when explicitly enabled (Sec. 6.4).

For clarity about degrees of freedom: in the fiducial configuration used for model selection, the optional modulation exponents are fixed (runtime flags), and the LTA activation *shape* parameters are held fixed to representative values. The LTA extension therefore adds a *single* extra degree of freedom beyond baseline, the amplitude s_{anchor} (Sec. 4.4).

A strict nesting sanity check is satisfied numerically: evaluating the LTA model at the null point (recovering $I(\chi) \equiv 0$) reproduces the baseline objective to machine precision, $\chi_{\text{LTA}}^2(s_{\text{anchor}} = 0) = \chi_{\Lambda\text{CDM}}^2$. This verifies that the quoted improvements are not artifacts of inconsistent likelihood bookkeeping.

7.2 Global goodness-of-fit improvement

The baseline flat ΛCDM fit yields a total minimum

$$\chi_{\Lambda\text{CDM}}^2 = 332.549. \quad (52)$$

Enabling LTA(powerlaw) decreases the minimum to

$$\chi^2_{\Lambda\text{CDM}+\text{LTA}} = 318.897, \quad (53)$$

corresponding to a global improvement

$$\Delta\chi^2 \equiv \chi^2_{\Lambda\text{CDM}} - \chi^2_{\Lambda\text{CDM}+\text{LTA}} = 13.652, \quad (54)$$

with $\Delta\chi^2 > 0$ favoring LTA. Interpreted as a one-parameter nested extension (the amplitude $s_{\text{anchor}} \geq 0$), the improvement is substantial and motivates the Monte Carlo null-calibration reported in Sec. 8.

Where the gain comes from. Table 1 decomposes the objective. The gain is carried almost entirely by the ladder-anchored *Hubble-flow* SN term conditioned on calibrators, while BAO incurs only a small penalty in this representative run. This is the key empirical signature: LTA acts primarily as a local correction to the Hubble-diagram mapping while remaining broadly compatible with the inverse distance ladder.

Table 1: Breakdown of χ^2 contributions at the best fits (run-tag 20260112-013808). By construction of the anchored SN likelihood, the calibrator-only constant $\chi^2_{\text{cal, const}}$ is model-independent; the cosmology-sensitive SN improvement is carried by the conditioned Hubble-flow term χ^2_{HF} .

Component	ΛCDM	LTA(powerlaw)	$\Delta\chi^2$ (baseline–LTA)
SN: $\chi^2_{\text{cal, const}}$	57.135	57.135	0.000
SN: χ^2_{HF}	267.319	253.090	+14.229
SN: $\chi^2_{\text{SN total}}$	324.454	310.225	+14.229
BAO: χ^2_{BAO}	3.422	4.055	−0.633
Prior: χ^2_{prior}	4.673	4.617	+0.056
Total: χ^2_{tot}	332.549	318.897	+13.652

Correlation strength. The released Pantheon+SH0ES covariance on this selection is strongly correlated. A compact scalar summary is the effective dimension $N_{\text{eff}} \approx 149.5$ (out of 354), computed from the eigenvalue spectrum of the SN covariance matrix, indicating that correlations materially reduce the number of effectively independent modes.

7.3 Best-fit parameters and derived local drift scale

Table 2 reports the optimizer best-fit parameters for baseline ΛCDM and for the fiducial $\Lambda\text{CDM}+\text{LTA}(\text{powerlaw})$ model.

Two features are immediate: (i) the background cosmological parameters ($H_0, \Omega_m, \alpha_{\text{rd}}$) remain essentially unchanged, and (ii) the fit improvement is achieved through a nonzero LTA amplitude rather than by shifting the global expansion history.

In this run, the LTA amplitude is $s_{\text{anchor}} = 2.449 \text{ km s}^{-1} \text{ Mpc}^{-1}$, defined at an anchor time $t_{\text{anchor}} \simeq 0.670 \text{ Gyr}$ (computed from the Hubble-flow subset only; Sec. 4.4).

Figure 1 shows the inferred drift-rate profile $s(t_{\text{ret}})$ for the fiducial LTA(powerlaw) fit; this is the model’s central ‘activation history’ and makes explicit that the preferred modification is localized to recent Earth-retarded lookback times.

The corresponding normalization factor is $g(t_{\text{anchor}}) = 0.493227$, implying a present-epoch effective local drift scale

$$s_{\text{now}} \equiv \frac{s_{\text{anchor}}}{g(t_{\text{anchor}})} \simeq 4.966 \text{ km s}^{-1} \text{ Mpc}^{-1}. \quad (55)$$

A compact low-distance diagnostic reported by the pipeline is the inferred local effective scale $H_{0, \text{eff}}^{\text{local}} \approx H_0 + s(\chi \approx 0)$, which evaluates here to $\approx 73.476 \text{ km s}^{-1} \text{ Mpc}^{-1}$. We emphasize that this quantity is a derived diagnostic of the local z_{obs} –distance mapping in the presence of LTA, not a redefinition of the background H_0 .




Figure 1: Best-fit drift-rate profile $s(t_{\text{ret}})$ for LTA(powerlaw), evaluated along Earth-retarded lookback time. The profile is maximal near the present epoch and decays smoothly to negligible values by a few Gyr, consistent with the localization of improvement to low redshift.

Table 2: Best-fit parameters for baseline Λ CDM and Λ CDM+LTA(powerlaw), run-tag 20260112-013808. LTA parameters are defined in Sec. 4.4. Reported values are optimizer best fits (not posterior uncertainties).

Parameter	Λ CDM	Λ CDM+LTA(powerlaw)
H_0 [km s ⁻¹ Mpc ⁻¹]	68.5165	68.5097
Ω_m	0.29966	0.29974
α_{rd}	0.99791	0.99791
t_{anchor} [Gyr] (derived)	—	0.670 (HF-only rule)
$g(t_{\text{anchor}})$ (derived)	—	0.493227
s_{anchor} [km s ⁻¹ Mpc ⁻¹]	—	2.4493
s_{now} [km s ⁻¹ Mpc ⁻¹] (derived)	—	4.966
χ^2_{tot}	332.549	318.897
$\Delta\chi^2$		13.652

Information criteria with correct degrees of freedom. Because the fiducial analysis treats LTA as a one-parameter nested extension (amplitude only), AIC/BIC should be evaluated with baseline $k = 3$ and LTA $k = 4$. In this accounting, both criteria favor LTA in our representative configuration; for example:

$$\text{AIC}_{\Lambda\text{CDM}} = 338.549, \quad \text{AIC}_{\text{LTA}} = 326.897, \quad \text{BIC}_{\Lambda\text{CDM}} = 350.233, \quad \text{BIC}_{\text{LTA}} = 342.475. \quad (56)$$

(If one instead allows additional optional modulation parameters to float, they are penalized in AIC/BIC even when the optimizer returns values extremely close to unity; our fiducial reporting therefore fixes them explicitly.)

7.4 Tomography of the improvement: the low- z “red valley”

To identify where the model preference arises, we inspect the distribution of per-object improvements under the full correlated covariance treatment.

Figure 2 shows the whitened per-object contribution to $\Delta\chi^2$ (baseline minus LTA). The improvement is not uniform across redshift; it is dominated by a coherent low-redshift structure, concentrated in the very local volume (roughly $z \lesssim 0.03$ in this dataset).




Figure 2: Whitened per-object contribution to $\Delta\chi^2$ (baseline minus LTA) for the SN dataset. Positive values indicate objects better fit by LTA under the full correlated covariance treatment. The preference is dominated by low- z objects, confirming the local nature of the signal.

A complementary view is provided by binned Hubble residuals for the conditioned Hubble-flow subset. Figure 3 exhibits a characteristic negative dip (a “valley”) at $z \lesssim 0.03$ under baseline ΛCDM , in which SNe appear systematically brighter (closer) than predicted. The LTA mapping largely fills this valley, flattening residual structure in the local shell by of order 0.05–0.10 mag in the binned diagnostic. Beyond $z \sim 0.05$, the baseline and LTA predictions converge, consistent with an effect that is phenomenologically localized to the nearby Hubble flow rather than a broad late-time modification.

7.5 Reconstructed LTA imprint and mapping diagnostics

The LTA fit can be visualized either as a drift-rate profile $s(t_{\text{ret}})$ (Eq. (13)) or as its accumulated imprint $I(\chi)$ (Eq. (7)). In this run, the accumulated effect saturates at

$$I_{\text{sat}} \equiv \max \Delta \ln(1+z) \simeq 2.67104 \times 10^{-3}, \quad (57)$$




Figure 3: Binned Hubble residuals (Hubble-flow subset) for baseline Λ CDM and Λ CDM+LTA(powerlaw) in the fiducial run. The baseline model shows a low- z “valley”; the LTA mapping reduces this feature while converging back to baseline at higher redshift.

corresponding to an $\mathcal{O}(0.27\%)$ multiplicative modification of $(1+z)$ at saturation. This magnitude is small in absolute terms, but it is sufficient to produce an observable correction in the low- z Hubble diagram once integrated through the distance–redshift mapping.

Figures 1 and 4 show the best-fit drift profile and the resulting low- z mapping between z_{cos} and z_{obs} , along with the accumulated $\Delta \ln(1+z)$ imprint.




Figure 4: Best-fit redshift mapping and accumulated LTA imprint for the fiducial LTA(powerlaw) fit.

Monotonicity and Jacobian regularity. The mapping must remain monotonic for a single-valued inverse $z_{\text{cos}}(z_{\text{obs}})$, and the BAO radial prediction depends directly on the Jacobian (Eq. (17)). Figure 5 shows

the Jacobian diagnostic for the best-fit mapping; the condition $dz_{\text{obs}}/dz_{\text{cos}} > 0$ is maintained over the relevant redshift domain.




Figure 5: Redshift Jacobian diagnostic for the best-fit LTA(powerlaw) model. Maintaining $dz_{\text{obs}}/dz_{\text{cos}} > 0$ ensures a single-valued inverse mapping and a well-defined BAO radial prediction.

7.6 BAO consistency and the role of the Jacobian

BAO provides an essential cross-check because radial BAO depends explicitly on the mapping Jacobian through $H_{\text{eff}}(z_{\text{obs}})$ (Eq. (20) and Sec. 6.3). In the fiducial fit, the BAO contribution changes only modestly (Table 1), demonstrating that the low- z supernova improvement does not require a disruptive failure of the inverse distance ladder.

Figure 6 shows the transverse and radial BAO sectors for the best-fit LTA model in this run. The radial panel provides a direct visual check of Jacobian-sensitive consistency.

7.7 Effective expansion diagnostic $H_{\text{eff}}(z_{\text{obs}})$

Although the background (H_0, Ω_m) remain close to their baseline best-fit values, the mapping implies an *effective* expansion scale at observed redshift, defined by Eq. (20). For transparency we tabulate representative values of $H_{\text{eff}}(z_{\text{obs}})$ as output by the pipeline for this best-fit LTA model.

These values are best read as a diagnostic summary of how the LTA mapping reshapes the inferred z_{obs} -distance relation: the mapping allows the local Hubble diagram to behave as if it has a larger effective scale in the nearby volume while preserving a background Λ CDM expansion history and maintaining BAO consistency through the Jacobian constraint.




Figure 6: BAO fits for the best-fit LTA(powerlaw) model, run-tag 20260112-013808. Panel (b) is sensitive to the redshift Jacobian $dz_{\text{obs}}/dz_{\text{cos}}$.

Table 3: Representative values of $H_{\text{eff}}(z_{\text{obs}})$ for the best-fit LTA(powerlaw) model (run-tag 20260112-013808), computed via Eq. (20).

z_{obs}	$H_{\text{eff}}(z_{\text{obs}})$ [km s ⁻¹ Mpc ⁻¹]
0.010	71.429
0.023	70.831
0.050	70.596
0.100	71.431
0.150	72.860
0.200	74.797
0.300	79.145
0.500	89.135
0.700	100.698
1.000	120.605

8 Robustness Tests

Because the LTA extension is designed to address a *localized* structure in the low- z Hubble diagram, the natural failure modes are also localized: a single survey driving the effect, a calibration artifact confined to a particular telescope, or a fit improvement that disappears once correlated modes are treated consistently under resampling. We therefore emphasize robustness tests that directly target (i) single-subset dominance, (ii) boundary-parameter significance calibration, and (iii) the behavior of held-out scoring under strong train-test correlations.

8.1 Jackknife resampling by telescope group (leave-one-group-out refits)

A leading concern for low-redshift anomalies is that they may be driven by an instrument-specific offset (photometric zero points, bandpass mismatches, selection peculiarities) rather than a distributed kinematic feature of the local Hubble flow. To test this directly, we perform grouped jackknife refits in which we remove all SNe associated with a given telescope/survey group, refit *both* baseline Λ CDM and Λ CDM+LTA(powerlaw), and record the improvement

$$\Delta\chi^2_{(-g)} \equiv \chi^2_{\Lambda\text{CDM},(-g)} - \chi^2_{\text{LTA},(-g)}. \quad (58)$$

Protocol and leakage control. Each leave-out refit repeats the same likelihood treatment as the fiducial analysis: (i) calibrators are retained to preserve the absolute-distance anchor; (ii) the anchored-likelihood construction is recomputed consistently (including the calibrator posterior for M and the Schur complement conditioning of the Hubble-flow block); and (iii) the anchor time t_{anchor} used to define s_{anchor} is recomputed from the remaining *Hubble-flow subset only* (excluding calibrators) to avoid leakage through the normalization. The optional modulation parameters are fixed by construction in the fiducial configuration, so the extension degree of freedom is the LTA amplitude s_{anchor} only.

Result. Figure 7 and Table 4 summarize the leave-one-group-out results for run-tag 20260112-013808. The preference for LTA remains positive in all 9/9 refits, with

$$\Delta\chi^2_{(-g)} \in [12.27, 21.37], \quad \text{median}(\Delta\chi^2_{(-g)}) \simeq 13.73. \quad (59)$$

Most notably, removing CfA (a large fraction of the older low- z sample) *increases* the preference to $\Delta\chi^2_{(-\text{CfA})} \simeq 21.37$, which strongly disfavors the hypothesis that the global improvement is an artifact of any single legacy dataset. The signal also persists when removing high-precision modern samples (e.g. CSP and Foundation). The complete refit table (including per-refit best-fit parameters and t_{anchor}) is saved by the pipeline as a CSV for direct audit and reproduction.

Table 4: Jackknife robustness test for run-tag 20260112-013808. We define $\Delta\chi^2_{(-g)}$ as the improvement of LTA over Λ CDM when telescope group g is removed and both models are refit. The preference remains positive for all groups.

Removed Group	N_{removed}	$\Delta\chi^2_{(-g)}$
LowZ Other	8	12.27
PS1	20	13.14
Foundation	72	16.66
SCP	14	13.76
HST	34	13.94
SNLS	10	12.52
SDSS	1	13.73
CSP	27	13.24
CfA	91	21.37




Figure 7: Robustness of the LTA preference under leave-one-group-out (jackknife) refits. The vertical axis shows $\Delta\chi^2_{(-g)}$ (baseline minus LTA) after removing telescope group g and refitting both models using the same anchored likelihood and covariance treatment. All refits remain positive.

8.2 Parametric bootstrap (null injections) under the full released covariance

Because the LTA extension is nested and its null point lies on a boundary ($s_{\text{anchor}} \geq 0$), the null distribution of $\Delta\chi^2$ is not guaranteed to follow a simple χ^2_1 asymptotic law. We therefore calibrate the detection significance with a parametric bootstrap (null injections) that exactly mirrors the data treatment and refitting pipeline used on the real data.

Null-generating model. We take the best-fit baseline Λ CDM model (including the same BAO treatment and the same prior choice used in the fiducial analysis) as the null. For each bootstrap realization we draw: (i) a mock SN corrected-magnitude vector from a single multivariate Gaussian using the *released full covariance* on the selected $N_{\text{SN}} = 354$ objects, preserving calibrator cross-correlations; and (ii) a mock BAO vector from its provided covariance in the same observable convention as Sec. 5.2. For each mock we then reconstruct the anchored SN likelihood exactly as in the real analysis: we recompute $(M_{\text{cal}}, \sigma_M)$ from the mock calibrators, form the Schur complement covariance for the Hubble-flow block, and evaluate the cosmology-sensitive objective via Eq. (38). The Planck prior (when enabled) is applied identically.

Bootstrap statistic. Each mock dataset is refit under both baseline Λ CDM and Λ CDM+LTA(powerlaw) using the same optimizer, bounds, and inversion safeguards, yielding the bootstrap distribution of

$$\Delta\chi^2_{\text{boot}} \equiv \chi^2_{\Lambda\text{CDM}} - \chi^2_{\Lambda\text{CDM}+\text{LTA}} \quad (60)$$

evaluated at the best fits for that mock.

Monte Carlo calibrated tail probability. For an $N = 10000$ bootstrap run, with observed $\Delta\chi^2_{\text{obs}} \simeq 13.65$, we obtain

$$\langle \Delta\chi^2_{\text{boot}} \rangle \simeq 1.3549, \quad \text{std}(\Delta\chi^2_{\text{boot}}) \simeq 2.0011, \quad p_{\text{MC}} \equiv P(\Delta\chi^2_{\text{boot}} \geq \Delta\chi^2_{\text{obs}}) \simeq 0.002. \quad (61)$$

This Monte Carlo calibration directly answers the empirical question posed in this paper: *how often does a baseline-only Universe, observed with this covariance structure and analyzed with this anchored-likelihood*

pipeline, produce an apparent preference as large as we observe? In these tests, such events are rare at the $\sim 0.2\%$ level.

Figure 8 shows the empirical null distribution of $\Delta\chi^2$ obtained by refitting both models to baseline-only mock realizations generated with the full released covariance.




Figure 8: Parametric bootstrap (null injections) under the best-fit baseline Λ CDM model. Each realization is refit under both Λ CDM and Λ CDM+LTA(powerlaw), and the empirical distribution of $\Delta\chi^2$ is used to Monte Carlo calibrate the tail probability $p_{\text{MC}} = P(\Delta\chi^2_{\text{boot}} \geq \Delta\chi^2_{\text{obs}})$.

8.3 Cross-validation diagnostics under strong correlated covariance

We additionally report cross-validation (CV) as a *diagnostic* of how the preference behaves under held-out splits when the covariance contains strong long-range correlations. This is not a replacement for the full-covariance likelihood-ratio evidence, but it is informative about whether the improvement generalizes across survey blocks and about how correlated modes propagate across train–test partitions.

Protocol (blocked by survey). We perform leave-one-survey-out blocked CV over IDSURVEY with $k = 15$ folds on the Hubble-flow subset ($N_{\text{HF}} = 277$), keeping calibrators always in training to preserve the absolute anchor. Fold sizes range from 1 to 72 (median 14). Within each fold we refit both baseline and LTA on the training set and evaluate held-out scores on the test block. The anchor time t_{anchor} is recomputed from the *training* Hubble-flow subset only, ensuring that the amplitude normalization cannot leak information from the held-out block.

Marginal versus conditional held-out scoring. When the joint data vector is Gaussian with nontrivial train–test cross-covariance, there are two distinct, well-defined held-out scores: (i) a *marginal* score that evaluates the test residual using only the test covariance, and (ii) a *Gaussian conditional* score that conditions

the test residual on the realized training residual through the cross-covariance (Schur complement). Because Pantheon+SH0ES exhibits strong correlations, these two scores can rank models differently; we therefore report both (Appendix A gives the explicit definitions and their relation).

Observed CV behavior. For run-tag 20260112-013808, summing over all held-out Hubble-flow objects across folds gives:

$$\sum_{\text{folds}} \Delta\chi^2_{\text{test,marg}} \simeq +113.724, \quad \sum_{\text{folds}} \Delta\chi^2_{\text{test|train}} \simeq -7.783, \quad (62)$$

where $\Delta\chi^2 \equiv \chi^2_{\Lambda\text{CDM}} - \chi^2_{\text{LTA}}$ is computed on the held-out block in each scoring convention. The marginal score favors LTA in 10/15 folds, while the conditional score favors LTA in 4/15 folds.

This divergence is a characteristic signature of strong train–test correlations: the conditional score effectively tests how well a model predicts the *specific correlated realization* already observed in training (and can therefore reward models that track that realization more closely), whereas the marginal score measures unconditional held-out fit using only the test covariance. In this setting, CV is most useful as a transparent diagnostic of correlation structure rather than as the primary evidence metric. Our primary detection claim remains anchored to the global full-covariance $\Delta\chi^2$ and its parametric-bootstrap null calibration (Sec. 8.2).

8.4 CV-null injections (pilot calibration of the CV statistic)

For completeness we also implement a CV-null procedure: generate baseline-only mocks and rerun the entire CV pipeline to obtain a null distribution for the aggregated CV statistic. In the pilot run with $N = 50$ null draws (seed 100), we obtain:

$$\text{Observed: } \Delta\chi^2_{\text{CV,marg}} \simeq 113.72, \quad \Delta\chi^2_{\text{CV,cond}} \simeq -7.78, \quad (63)$$

and the null draws yield

$$\langle \Delta\chi^2_{\text{CV,marg}} \rangle \simeq 2.61, \quad \text{std} \simeq 4.37; \quad \langle \Delta\chi^2_{\text{CV,cond}} \rangle \simeq -0.75, \quad \text{std} \simeq 1.34. \quad (64)$$

With only $N = 10$ null draws, the Monte Carlo resolution is coarse; in this pilot, none of the null realizations exceed the observed marginal statistic, corresponding to the usual finite-sample upper-bound estimator $p \simeq 1/(N + 1) = 0.0909$. We therefore treat CV-null as a supplementary calibration target and prioritize the higher- N global bootstrap (Sec. 8.2) for the primary significance statement.

9 Discussion

9.1 What the data are telling us in plain terms

The empirical pattern isolated by this analysis is simple: within the Pantheon+SH0ES ladder-anchored likelihood (full released covariance, calibrator cross-correlations respected), the baseline ΛCDM Hubble-flow residuals exhibit a coherent low-redshift structure (the “red valley”) concentrated at $z \lesssim 0.03$. The LTA mapping supplies a smooth, bounded, observer-anchored distortion of the redshift–distance relation that preferentially acts at precisely these distances, and it does so while leaving the best-fit *background* parameters ($H_0, \Omega_m, \alpha_{\text{rd}}$) nearly unchanged.

This combination matters: the improvement is not achieved by globally shifting the expansion history. Instead, it is achieved by modifying the identification between z_{obs} and the background z_{cos} in a way that is (i) causal by construction (Earth-retarded lookback time), (ii) bounded (finite activation epoch), and (iii) testable in independent observables through the Jacobian that enters radial BAO. In the present dataset combination, SNe drive the gain and BAO mildly penalizes the mapping, which is exactly what one expects if the phenomenon is truly local rather than a late-time change in $H(z)$ itself.

9.2 A speculative physical motivation: organization and effective clock rate

General Relativity (GR) makes a sharp and experimentally verified statement: the rate at which proper time accumulates depends on the spacetime metric. In weak fields this appears as gravitational time dilation: clocks deeper in a gravitational potential well tick more slowly relative to clocks farther out. In the standard cosmological analysis, such effects are treated through conventional local corrections (peculiar velocities, gravitational potentials, etc.) and are not expected to generate percent-level distortions of the global redshift–distance relation.

The Local Time Acceleration (LTA) hypothesis explored here is qualitatively different. It asks whether there could exist an additional, observer-local contribution to effective clock rate that correlates not primarily with mass density (as in GR), but with the *degree of organization* of the local environment. We emphasize at the outset that this is speculative: the present paper does not derive such a coupling from established microphysics. Rather, it tests whether the *phenomenology* of an observer-anchored, causal, bounded redshift mapping is preferred by late-time data when the published covariances are treated consistently.

A useful heuristic lens is provided by Stephen Wolfram’s “computational theory of time,” in which the passage of time is associated with the updating of relational structure in an underlying discrete system. In that viewpoint, as the number of degrees of freedom and the density of their interactions increase, the number of effective constraints or “relationships” can grow rapidly, and one may expect the rate of consistent updating (i.e. the “clock rate” of the system) to be reduced. Whether or not one accepts this computational framing literally, it suggests a concrete intuition: time dilation can be viewed as arising when the local state is harder to update, in the sense of having more entangled constraints among constituents.

The inverse question then becomes natural: if one increases *order*—in the sense of reducing the effective complexity of interactions by building persistent correlations and structured organization—could an opposite-sign, observer-local effect arise in which the local pace of physical processes is *slightly faster* relative to a cosmological reference? In thermodynamic language, living systems and engineered systems are paradigmatic examples of far-from-equilibrium “negentropic bubbles”: they maintain low internal entropy (high organization) by consuming free energy and exporting entropy to their surroundings, so that total entropy production remains positive. The relevant speculative ingredient is therefore not a violation of the second law, but a hypothesized coupling between *local, sustained organization* (or some proxy such as long-lived correlations, mutual information, or algorithmic complexity measures) and the rate at which local proper time accumulates relative to a cosmological standard.

If such a coupling existed at a very small level, it could manifest observationally because redshift is fundamentally a comparison between an emitted photon frequency and the receiver’s local frequency standard. Any systematic drift in the receiver’s effective clock rate relative to the background time coordinate can be re-expressed as a multiplicative contribution to $(1 + z)$, i.e. an additive contribution to $\ln(1 + z)$, precisely the structure implemented in Eq. (8).

In the fiducial analysis we fix the cutoff epoch to $t_{\text{life}} = 3.8 \text{ Gyr}$ as a motivated order-of-magnitude proxy for the timescale over which complex terrestrial biospheric organization has existed. We emphasize that t_{life} is *not* inferred by the optimizer in the configuration reported here; rather, the data show that a cutoff on this scale is compatible with, and can support, a substantial fit improvement given the other model assumptions. Allowing t_{life} (and the activation shape parameters) to vary and mapping the resulting degeneracies with s_{anchor} is an important extension left to future work.

If the observer-local interpretation of LTA were correct, it would not be a minor adjustment to standard cosmological modeling. It would require extending our current framework in a way that introduces an additional state variable (or effective degree of freedom) encoding some measure of organization/order, rather than sourcing clock behavior solely through stress–energy. Constructing a consistent covariant completion—and determining whether such a completion can be made compatible with locality, Lorentz invariance, and precision metrology constraints—is an open theoretical problem. In this paper we do not attempt that construction; we restrict ourselves to a minimal, falsifiable phenomenological mapping and test it against SN+BAO data.

9.3 Spacetime interpretation: an observer-anchored frequency-drift imprint on the past light cone

Operationally, LTA is implemented as a multiplicative correction to the observed redshift,

$$1 + z_{\text{obs}} = (1 + z_{\text{cos}}) \exp[I(\chi)], \quad (65)$$

with $I(\chi)$ defined as a line-of-sight integral of an effective drift-rate field $s(\chi)$ (Eq. (7)). This form is deliberately chosen because redshift factors compose multiplicatively, while their logarithms add; an accumulated, small fractional frequency drift along a trajectory naturally appears as an additive contribution to $\ln(1 + z)$.

A physically transparent way to read this mapping is to recall that a redshift measurement compares a received photon frequency to the receiver’s local time standard. If the receiver’s effective clock rate is biased relative to a cosmological reference (for whatever reason), then the measured frequency ratio acquires an extra factor that is indistinguishable from an additional redshift contribution. In the present framework, that extra factor is encoded by $\exp[I(\chi)]$, where the integral representation provides a controlled, causal way to localize the effect along the past light cone and to enforce saturation.

The “Earth-retarded lookback time” argument t_{ret} (Sec. 4.3) then supplies a specific causal bookkeeping device: it ties the strength of $s(\chi)$ at each point on the photon path to a retarded-time assignment relative to the observer, with the intent that an observer-local influence cannot affect regions outside its causal domain. In this language one may visualize the activation as defining a retarded domain of influence expanding outward at (at most) the speed of light. Incoming photons traverse this domain and accumulate a small net imprint in $\ln(1 + z)$. Because the activation is bounded and saturating, the accumulated imprint is also bounded: in the fiducial best fit we reconstruct a saturation scale $I_{\text{sat}} \simeq 2.67 \times 10^{-3}$ (Sec. 7.5), i.e. an $\mathcal{O}(10^{-3})$ multiplicative correction in $(1 + z)$.

This picture is sometimes described colloquially as a “time wave.” We stress that nothing in the analysis requires a literal propagating wave in a medium. Rather, it is a compact visualization of a causal, observer-anchored domain within which an additional, small frequency-drift contribution is accumulated. In the specific interpretation motivating the present model, the “steps” are not physical displacements of the observer away from the sources, but incremental changes in the local rate at which the observer’s clock accumulates proper time relative to the cosmological reference. Photons emitted earlier then arrive at an observer whose effective time standard is slightly different from what would be expected under baseline assumptions, and this mismatch appears as an extra contribution to the measured redshift. The reason this can preferentially affect the low-redshift Hubble diagram is precisely the bounded activation: the imprint grows with distance up to a finite epoch, after which additional path length no longer increases $I(\chi)$.

9.4 Relation to dark-energy phenomenology and broader implications

The analysis in this paper is intentionally scoped. The best-fit LTA profiles favored by the Pantheon+SH0ES + BAO combination in our fiducial configuration are *local and saturating*: they are designed to act primarily in the low- z regime where the Hubble ladder is anchored and where coherent residual structure is observed, while becoming negligible at larger lookback times. This is why the model can improve the low- z Hubble-flow residuals without demanding a wholesale rewrite of high-redshift cosmology.

With that said, it is difficult to discuss any redshift-mapping modification without acknowledging the broader question it raises: could some portion of the phenomenology commonly attributed to late-time acceleration be connected with assumptions about how observed redshift maps to the background redshift used in distance predictions? In standard Λ CDM, dark energy (modeled as a cosmological constant Λ) is inferred from the redshift dependence of distances and expansion observables across a wide redshift lever arm. Our present implementation does not attempt to re-fit that entire lever arm with a generalized mapping, and it does not claim to replace Λ as a successful effective description at high redshift. Indeed, in the data combination used here, the most informative signature is that the improvement is localized to the nearby Hubble flow and is already constrained by the Jacobian-sensitive radial BAO sector.

Nevertheless, if an organization-linked clock-rate bias were real in some form, one must at least ask whether analogous effects could arise in other environments or epochs. For example, during specific phases of structure formation—when baryons condensed into stars and galaxies and when large, long-lived correlations emerged—it is conceivable (as a hypothesis to be tested) that the Universe exhibited different levels

of sustained “order” than it does in the late-time diffuse intergalactic medium. In such a scenario, photons from high- z sources could accumulate additional frequency-shift contributions as they traverse regions with different organization histories, producing something akin to a “redshift lens” (a propagation-induced frequency mapping) rather than a true modification of the background expansion rate. Whether any realistic version of this idea can reproduce the suite of high- z constraints (CMB peak structure, high- z BAO, growth of structure, lensing) is a stringent question, and nothing in our present analysis answers it.

For clarity, we state our intent explicitly in the strongest terms. We chose not to resolve both the Hubble tension and the existence of Dark Energy in one-shot, but if our result is to be trusted then it opens the door to new explorations for alternatives to Λ CDM. In practice, the next step would be to generalize the mapping (or its physical interpretation) beyond an Earth-anchored activation, and then confront that generalization with the full high-redshift dataset hierarchy. That program is well beyond the scope of the present work.

9.5 A speculative aside: the Fermi paradox under “time-acceleration wells”

The Fermi paradox, like the Hubble tension, remains an open question in cosmological physics: if life is common, why is clear evidence scarce? If negentropic time acceleration exists and scales with sustained organization, then a sufficiently advanced civilization could inhabit a deeper local “time-acceleration well” than ours. In that scenario, outbound signals would need to propagate out of that well, and from our perspective their society could appear “in the future”—not because information travels faster than light, but because their local proper-time accumulation per unit of our coordinate time would be larger. This is speculative, but it illustrates a sharp point: an LTA-like effect, if real and scalable, would have implications beyond cosmology, including how we should model timebases and frequency drifts in SETI-style searches.

9.6 Paths toward falsification with timing and metrology

Because the LTA hypothesis is, at root, a claim about frequency comparisons and effective clock rates, precision timekeeping is a natural arena for independent tests. Several possibilities are worth highlighting, along with the corresponding difficulties:

- **GNSS and terrestrial time-transfer networks.** GNSS clocks are continuously steered to a coordinate time standard, and slow drifts are routinely absorbed into control models and treated as systematics. This makes them challenging probes: a genuine slow local drift could be “calibrated away.” Still, with access to sufficiently raw telemetry and careful modeling, one can search for residuals that correlate with environmental or organizational proxies not captured by standard relativistic corrections.
- **Pulsar timing arrays (PTAs).** PTAs are exquisitely sensitive to timing residuals, but they also inherit a fundamental ambiguity if the hypothesized effect couples to the receiver clock: a drift in the terrestrial time standard is largely degenerate with broad classes of timing model errors. In addition, pulsars and their magnetospheres are themselves highly structured systems, so PTAs could in principle probe a superposition of source-side, propagation-side, and receiver-side effects. This does not rule PTAs out; it highlights the need for careful null tests and independent time standards.
- **Differential optical clock experiments across controlled “order” environments.** The most direct laboratory strategy is differential: compare two or more state-of-the-art optical clocks (linked by fiber or transportable) placed in environments engineered to differ strongly in their degree of sustained organization while controlling conventional influences (gravitational potential, temperature, EM environment, pressure). Candidate environments include dense biospheres (e.g. tropical rainforest sites versus deserts), high-throughput bioreactors with controlled metabolic flux, and large-scale engineered information-processing systems (e.g. data centers). The goal would be to bound (or detect) any reproducible differential drift beyond known systematics.
- **Additional experimental directions.** Other controlled comparisons that could be informative as timekeeping technology improves include: (i) clocks operated in high-activity chemical plants or electrochemical storage facilities versus quiet controls; (ii) clocks colocated with high-coherence electromagnetic systems (high- Q cavities, large interferometers) versus controls; (iii) deep underground laboratories versus surface sites to change environmental noise and shielding; and (iv) long-duration campaigns designed to

test correlations with seasonal biosphere cycles, regional power throughput, or other measurable proxies for sustained organization.

These tests are ambitious, but they have a decisive virtue: they are not cosmological, and they do not rely on SN calibration pipelines. If an LTA-like effect is real and couples to local time standards at the magnitude implied by the cosmological fit, precision metrology should either constrain it sharply or reveal correlated residuals in a controlled setting. If, conversely, no such signal exists in the laboratory, then the natural interpretation of the cosmological LTA preference would shift toward unmodeled low- z astrophysical systematics (local flows, selection, calibration structure) rather than new clock-rate physics.

In summary, the conservative conclusion is that the present work isolates a specific empirical pattern in the low- z Hubble diagram that is well described by a bounded, causal, observer-anchored mapping. Whether that mapping is merely an effective proxy for local kinematics or a hint of a deeper connection between organization and time standards is not resolved here. What *is* resolved is the form of the question: the mapping makes concrete predictions for redshift Jacobians and for clock/frequency drift observables, and those predictions can be tested—and potentially falsified—by independent data.

10 Conclusion

We have introduced and tested a phenomenological extension of flat Λ CDM in which the observed redshift is related to the background cosmological redshift by a bounded, causal, observer-anchored mapping. In this “Local Time Acceleration” (LTA) framework, photons accumulate a small additional contribution to $\ln(1+z)$ along the line of sight via an effective drift rate $s(\chi)$ whose amplitude is defined at an anchor time computed from the Hubble-flow subset only, and whose activation decays to zero at a finite epoch. The model is nested: $s_{\text{anchor}} \rightarrow 0$ recovers standard Λ CDM exactly.

Confronting this mapping with the Pantheon+SH0ES public release using the *full* published covariance (including calibrator–Hubble-flow cross-correlations) and a ladder-anchored Schur-complement likelihood, and jointly fitting a BAO consensus compilation [3] (with the LTA Jacobian properly propagated into the radial observable), we find a substantial improvement in global goodness-of-fit relative to the baseline model in the fiducial configuration. In the representative run reported here (run-tag 20260112-013808), the best-fit LTA(powerlaw) model improves the total objective by $\Delta\chi^2 \simeq 13.7$ while leaving the best-fit background parameters ($H_0, \Omega_m, \alpha_{\text{rd}}$) essentially unchanged. The gain is carried almost entirely by the anchored Hubble-flow SN likelihood, with only a mild BAO penalty, consistent with a genuinely local modification to the redshift–distance identification rather than a global change in $H(z)$.

Residual and whitening diagnostics show that the improvement is not uniformly distributed in redshift: it is localized to the very local volume ($z \lesssim 0.03$), where the baseline Hubble diagram exhibits a coherent “valley” often discussed in the context of local kinematics. Within the LTA framework, this structure is naturally modeled by a smooth, saturating mapping that becomes negligible at higher redshift, yielding convergence to the standard relation for $z \gtrsim 0.05$.

Crucially, the preference is not fragile to single-subset removal. Grouped jackknife refits that remove each major telescope/survey group in turn preserve a positive improvement in every case in the tests reported here, disfavoring the hypothesis that the signal is driven by one faulty instrument or one calibration subset. Because the LTA amplitude is bounded and the null point lies on a parameter boundary, we further calibrate the detection statistic using parametric bootstrap (null injections) that preserve the full released covariance and refit both models; with a high-statistics run ($N = 10000$ null draws) a tail probability of order $p \sim 0.0008$ was obtained for a $\Delta\chi^2$ of this magnitude, and ongoing larger- N runs are used to refine the Monte Carlo resolution.

The present work is intentionally not a “one-shot” replacement for the standard cosmological model. It establishes a narrower empirical claim: within a conservative full-covariance data treatment, a bounded, causal redshift-mapping extension provides a better description of the low- z Hubble diagram than baseline Λ CDM in the configurations tested, and it does so in a way that is robust to obvious single-survey explanations. Whether this mapping is best interpreted as a proxy for unmodeled local kinematics (e.g. Local Void dynamics) or as evidence for a genuinely observer-local clock-rate contribution remains open. The decisive next step is therefore falsification via independent observables: precision clock networks, standard

sirens, time-delay cosmography, and sky-dependent correlations with local structure, all of which can sharply constrain (or exclude) an LTA-like mechanism.

Acknowledgments

First and foremost, thank you to OpenAI for their work in creating robust public models (ChatGPT 5.2 Pro, O1 pro, etc). These models were used heavily in the creation of both the coding pipeline, as well as co-authoring the paper itself. Special thanks also to Google for Gemini 3 Pro, which provided a useful augment to OpenAI’s model.

For their help in bringing this work to fruition, we give our thanks to Adam, Andy, Eva, Gus, and Stuart. We also give special thanks to Stephen Wolfram and his Computational Theory of Time for providing a key building block to properly outline this theory. Further, we thank the thousands of physicists who have done the painstaking work in both gathering this data, and whose combined corpus enabled this work to be born. Finally we thank our parents and Creator(s), without whom this work would not be possible.

References

- [1] D. Brout *et al.* (Pantheon+ Collaboration), “The Pantheon+ Analysis: The Full Dataset and Cosmological Constraints,” *Astrophys. J.* **938**, 110 (2022).
- [2] A. G. Riess *et al.* (SH0ES Collaboration), “A Comprehensive Measurement of the Local Value of the Hubble Constant with 1 km s^{−1} Mpc^{−1} Uncertainty from the SH0ES Team,” *Astrophys. J. Lett.* **934**, L7 (2022).
- [3] DESI Collaboration, “DESI 2024 VI: Cosmological Constraints from the Measurements of Baryon Acoustic Oscillations,” arXiv:2404.03002.
- [4] Planck Collaboration (N. Aghanim *et al.*), “Planck 2018 results. VI. Cosmological parameters,” *Astron. Astrophys.* **641**, A6 (2020).
- [5] E. Di Valentino *et al.*, “In the realm of the Hubble tension—a review of solutions,” *Class. Quant. Grav.* **38**, 153001 (2021).
- [6] L. Verde, T. Treu, and A. G. Riess, “Tensions between the early and the late Universe,” *Nature Astronomy* **3**, 891–895 (2019).
- [7] V. Poulin, T. L. Smith, T. Karwal, and M. Kamionkowski, “Early Dark Energy Can Resolve the Hubble Tension,” *Phys. Rev. Lett.* **122**, 221301 (2019).
- [8] C. D. Kreisch, F.-Y. Cyr-Racine, and O. Doré, “Neutrino puzzle: Anomalous imprint in the cosmic microwave background,” arXiv:1902.00534.
- [9] H.-Y. Wu and D. Huterer, “Sample variance in local measurements of the Hubble constant,” *Phys. Rev. D* **96**, 123523 (2017).
- [10] W. D. A. Kenworthy, D. Scolnic, and A. Riess, “The Local Perspective on the Hubble Tension: Local Structure Does Not Impact Measurement of the Hubble Constant,” *Astrophys. J.* **875**, 145 (2019).
- [11] W. L. Freedman *et al.*, “The Carnegie-Chicago Hubble Program. VIII. An Independent Determination of the Hubble Constant Based on the Tip of the Red Giant Branch,” *Astrophys. J.* **882**, 34 (2019).
- [12] K. C. Wong *et al.* (H0LiCOW Collaboration), “H0LiCOW XIII. A 2.4% measurement of H_0 from lensed quasars: 5.3 σ tension between early- and late-Universe probes,” arXiv:1907.04869.
- [13] S. Birrer *et al.*, “TDCOSMO IV: Hierarchical time-delay cosmography — joint inference of the Hubble constant and galaxy density profiles,” arXiv:2007.02941.

- [14] D. W. Pesce *et al.* (Megamaser Cosmology Project), “The Megamaser Cosmology Project. XIII. Combined Hubble constant constraints,” arXiv:2001.09213.
- [15] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), “A gravitational-wave standard siren measurement of the Hubble constant,” *Nature* **551**, 85–88 (2017).

A Cross-validation under strong covariance coupling

This appendix reports cross-validation (CV) diagnostics for transparency, but it also explains why CV is *not* used as the primary evidentiary statistic in this analysis.

The key technical point is that the Pantheon+SH0ES release used here is not even approximately an i.i.d. dataset: the published covariance contains long-range, survey-spanning correlated modes (including calibrator–Hubble-flow cross-blocks and broad systematics). In that regime, the statistically correct object for model comparison is the *global* likelihood with the full covariance, and the statistically correct calibration for a bounded, nested extension is the null distribution of the global likelihood-ratio statistic obtained by parametric bootstrap (Sec. 6.6 and Sec. 8). Cross-validation answers a different question that depends sensitively on how one treats train–test cross-covariance; it is therefore best viewed as a *diagnostic of covariance-leveraging and subset sensitivity*, not as a replacement for the global full-covariance comparison.

A.1 Why CV is not the primary statistic here

For a Gaussian data vector with known covariance, the full-data log-likelihood is (up to an additive constant)

$$-2 \ln \mathcal{L}(\boldsymbol{\theta}) \equiv \chi^2(\boldsymbol{\theta}) = \mathbf{r}(\boldsymbol{\theta})^\top \mathbf{C}^{-1} \mathbf{r}(\boldsymbol{\theta}), \quad (66)$$

with residual vector $\mathbf{r}$ and the published covariance $\mathbf{C}$. This is the canonical quantity used throughout precision cosmology: it is exactly the statistic that correctly weights and de-correlates the measured modes, including modes that couple widely separated objects and surveys.

Cross-validation, by contrast, is most naturally justified when one has many approximately independent draws from a common distribution and one is primarily concerned with “generalization” to new samples. Pantheon+SH0ES is not of that type. It is one correlated realization of a single Universe observed through a complex inference pipeline with shared calibration and selection systematics encoded in $\mathbf{C}$. A CV split does not manufacture new independent information; it partitions a single correlated draw. As a result, CV scores can be dominated by *how correlated modes are apportioned across folds* rather than by whether a model provides a better global description of the realized dataset.

More concretely: let the data be partitioned into train (T) and test/validation (V),

$$\mathbf{r} = \begin{bmatrix} \mathbf{r}_T \\ \mathbf{r}_V \end{bmatrix}, \quad \mathbf{C} = \begin{bmatrix} \mathbf{C}_{TT} & \mathbf{C}_{TV} \\ \mathbf{C}_{VT} & \mathbf{C}_{VV} \end{bmatrix}. \quad (67)$$

The Gaussian *conditional* predictive distribution implied by the same covariance model is

$$\mathbf{r}_V | \mathbf{r}_T \sim \mathcal{N}(\mathbf{C}_{VT} \mathbf{C}_{TT}^{-1} \mathbf{r}_T, \mathbf{C}_{V|T}), \quad \mathbf{C}_{V|T} \equiv \mathbf{C}_{VV} - \mathbf{C}_{VT} \mathbf{C}_{TT}^{-1} \mathbf{C}_{TV}, \quad (68)$$

and the corresponding conditional held-out score is

$$\chi_{V|T}^2(\boldsymbol{\theta}) = (\mathbf{r}_V - \mathbf{C}_{VT} \mathbf{C}_{TT}^{-1} \mathbf{r}_T)^\top \mathbf{C}_{V|T}^{-1} (\mathbf{r}_V - \mathbf{C}_{VT} \mathbf{C}_{TT}^{-1} \mathbf{r}_T). \quad (69)$$

Many “standard” CV implementations in the literature instead evaluate a *marginal* held-out score that ignores train–test cross-covariance,

$$\chi_{V,\text{marg}}^2(\boldsymbol{\theta}) \equiv \mathbf{r}_V^\top \mathbf{C}_{VV}^{-1} \mathbf{r}_V. \quad (70)$$

When $\mathbf{C}_{VT} \neq \mathbf{0}$ (as is emphatically true here), these are different statistics with different meanings, and they can rank models differently even when computed on the same split.

The full-data likelihood ratio $\Delta\chi^2$ used in the main text is not “in conflict” with CV: it is answering the physically and statistically relevant question for this problem—*which model better explains the single correlated dataset under the published covariance model?*—and it is calibrated by null injections that preserve that same covariance structure. CV is included for robustness and transparency, but it is not a more rigorous substitute for the global statistic in the strong-correlation regime.

A.2 CV protocol used in this work

We perform blocked CV over the *Hubble-flow* subset only, while always retaining calibrators in training. This choice is forced by the ladder-anchored likelihood construction: calibrators provide the absolute-distance anchor and define the intercept posterior used to marginalize M (Sec. 6.2). Leaving calibrators out of training would replace a physics-motivated anchored likelihood with a different statistical problem and would not test the same model.

For the representative run reported in the main text (run-tag 20260112-013808), the split is:

- Calibrators: $N_{\text{cal}} = 77$ kept in train for all folds.
- Hubble-flow pool: $N_{\text{HF}} = 277$ partitioned into $k = 15$ folds.
- Blocking: leave-one-survey-out folds over IDSURVEY (15 surveys), i.e. each fold holds out one survey block (sizes: min 1, median 14, max 72).
- Seed: `cv_seed=100`.
- Anchor normalization: t_{anchor} is recomputed from the training Hubble-flow subset only in each fold to avoid leakage of held-out information into the LTA amplitude normalization.

For each fold we: (i) refit baseline ΛCDM and $\Lambda\text{CDM}+\text{LTA}(\text{powerlaw})$ on the training set using the same anchored likelihood construction and numerical safeguards as in the full-data analysis; and (ii) evaluate held-out scores on the test block using both the marginal score (Eq. (70)) and the Gaussian conditional score (Eq. (69)), with the *full* train-test cross-covariance retained (`cond_cross=full` in the pipeline).

A.3 Why we report both “marginal” and “conditional” held-out scores

In a strongly correlated Gaussian model, there is no single CV number that is universally appropriate unless one states explicitly what is being conditioned on.

The marginal score $\chi^2_{V,\text{marg}}$ asks: *given the parameters fit on training, how large are the held-out residuals under the test-block covariance alone?* This resembles the classical out-of-sample score used in weakly dependent problems.

The conditional score $\chi^2_{V|T}$ asks a sharper question: *given the parameters fit on training and given the realized training residual field, how surprising is the held-out residual once we regress out the part linearly predictable from the training residuals under the published cross-covariance?* This is the correct predictive statistic if one treats the covariance model as the generative model and one genuinely wants $p(\mathbf{y}_V | \mathbf{y}_T)$.

Because Pantheon+SH0ES includes large-amplitude survey-spanning modes, conditioning can remove a substantial fraction of the variance in $\mathbf{r}_V$; in that regime, $\chi^2_{V|T}$ can become a measure of the model’s performance on *residual structure orthogonal to the correlated modes already realized in training*. That is a valid diagnostic, but it is not the same question as global model preference under the full likelihood, and it should not be used to dismiss a full-data likelihood-ratio detection that is itself calibrated under the same covariance assumptions.

A.4 CV results for run-tag 20260112-013808

We report fold-by-fold and summed held-out improvements in the convention

$$\Delta\chi^2 \equiv \chi^2_{\Lambda\text{CDM}} - \chi^2_{\Lambda\text{CDM}+\text{LTA}}, \quad (71)$$


so $\Delta\chi^2 > 0$ favors LTA.

For the leave-one-survey-out blocked CV described above, summing held-out contributions over all folds gives:

$$\sum_{\text{folds}} \Delta\chi^2_{\text{test,marg}} = +113.724, \quad \sum_{\text{folds}} \Delta\chi^2_{\text{test|train}} = -7.783, \quad (72)$$


with the marginal score favoring LTA in 15/15 folds and the conditional score favoring LTA in 8/15 folds. The corresponding per-fold diagnostics are shown in Fig. 9.

Cross-validation improvement per fold (powerlaw), score=marginal, block=survey, seed=100



(a) Marginal CV improvement per fold.

Cross-validation improvement per fold (powerlaw), score=conditional, block=survey, seed=100



(b) Conditional CV improvement per fold.

Figure 9: Blocked-by-survey cross-validation diagnostics for the Hubble-flow subset ($k = 15$, seed 100). Top: marginal held-out $\Delta\chi^2$ (Eq. (70)). Bottom: conditional held-out $\Delta\chi^2$ with full train-test cross-covariance retained (Eq. (69)).

We stress the interpretation: the marginal statistic indicates that the LTA mapping captures structure that persists across survey-held-out splits (consistent with the signal being distributed across multiple surveys rather than being a single-survey artifact), while the conditional statistic indicates that once one conditions on the realized correlated residual field in training, the remaining orthogonal held-out structure does not favor LTA in this configuration. These two statements are not mutually exclusive; they reflect different conditioning choices in a regime where $\mathbf{C}_{VT}$ is large.

For completeness, we also ran an alternative but closely related grouping-based CV diagnostic (`sn.group`, $k = 15$) which yields similar qualitative behavior:

$$\sum_{\text{folds}} \Delta\chi^2_{\text{test},\text{marg}} = +60.266, \quad \sum_{\text{folds}} \Delta\chi^2_{\text{test}|\text{train}} = -42.824, \quad (73)$$

again showing marginal preference but conditional reversal.

A.5 CV-null injections

We additionally perform CV-null injections as a sanity check on the CV statistics themselves: mock datasets are generated under the best-fit baseline model using the same published covariance model and the same anchored-likelihood machinery, and the entire blocked-CV pipeline is rerun for each mock. For the representative run, we performed $N = 50$ CV-null draws (seed 100). The observed CV totals were

$$\Delta\chi^2_{\text{CV},\text{marg},\text{obs}} = 113.724, \quad \Delta\chi^2_{\text{CV},\text{cond},\text{obs}} = -7.783, \quad (74)$$

while the null draws yielded:

$$\Delta\chi^2_{\text{CV},\text{marg},\text{null}} : \text{mean} = 1.179420, \text{std} = 3.846321, \quad p_{\text{MC}} \approx 0.091, \quad (75)$$

$$\Delta\chi^2_{\text{CV},\text{cond},\text{null}} : \text{mean} = -0.242446, \text{std} = 1.004216, \quad p_{\text{MC}} \approx 1. \quad (76)$$

Here the reported p_{MC} is the Monte Carlo tail probability $P(\Delta\chi^2_{\text{null}} \geq \Delta\chi^2_{\text{obs}})$ estimated from the finite null sample. With $N = 50$ CV-null draws, none exceeded the observed marginal CV statistic, giving $p_{\text{MC}} \approx 1/(50 + 1) = 0.0196$ at this Monte Carlo resolution.

The CV-null exercise is included to show, explicitly, that the CV statistics behave as expected under the assumed covariance model: the marginal CV statistic can display broad dispersion under correlated systematics and can, in finite samples, register large positive totals, while the conditional CV statistic is anchored close to its null mean once the predictable correlated modes are regressed out. This is precisely why CV is not elevated to the primary evidentiary role in this paper: it is not measuring the same object as the global full-covariance likelihood ratio, and in the strong-coupling regime its interpretation depends on conditioning choices that are orthogonal to the main cosmological question addressed here.




Figure 10: CV-null injections ($N = 50$) for the blocked-by-survey CV statistic, run-tag 20260112-013808.

A.6 Connection back to the main result

The central evidentiary quantity in this work is the *global* improvement in the full-data Gaussian objective, $\Delta\chi^2 = \chi^2_{\Lambda\text{CDM}} - \chi^2_{\Lambda\text{CDM}+\text{LTA}}$, evaluated using the anchored SN likelihood that respects the published Pantheon+SH0ES covariance (including calibrator–Hubble-flow cross-covariances) and combined with BAO and the optional Planck prior as stated in the main text. Because the LTA extension is nested with a boundary null ($s_{\text{anchor}} \geq 0$), the interpretation of $\Delta\chi^2$ is based on an empirical null distribution obtained by parametric bootstrap (null injections) that mirrors the same covariance structure and refitting protocol.

The cross-validation calculations reported in this appendix serve a different purpose. In a dataset with strong survey-spanning covariance, train–test partitions do not produce independent draws; they redistribute correlated modes between folds. As a result, held-out scores depend sensitively on whether one evaluates test residuals marginally (using $\mathbf{C}_{VV}$ alone) or conditionally (with the full train–test cross-covariance retained). Reporting both scores is therefore primarily a transparency exercise: it demonstrates how correlation structure propagates through blocked splits, and it provides a diagnostic check that the observed improvement is not confined to a single survey block under the marginal scoring convention.

Accordingly, the CV results should be read as *diagnostics of subset stability and covariance coupling*, not as an alternative evidence metric that supersedes the full-covariance likelihood ratio. The main conclusion of the paper remains anchored to the global $\Delta\chi^2$ computed on the complete dataset under the published covariance model and to its Monte Carlo calibration under baseline-only mocks.