



Temporal Equivalence Principle: Dynamical Proper Time and the Illusion of Primordial Deuterium

Matthew Lukin Smawfield

Version: v0.6 (Dubai)

First published: 12 August 2026 — Last updated: 17 October 2026

DOI: 10.5281/zenodo.21841147

Code Availability: github.com/matthewsmawfield/TEP-BBN

Abstract

The spectroscopic identification of primordial deuterium—a foundational empirical pillar of Big Bang cosmology—is not mathematically unique. Within standard Voigt-profile fitting architectures, the absorption features conventionally attributed to deuterium across the three best-studied quasar sightlines are statistically better described by unrestricted ordinary hydrogen. Using Monte-Carlo-calibrated model comparison across Q1009+2956 (Keck/HIRES), PKS 1937–101 (ESPRESSO), and J1332+0052 (UVES+HIRES), the unrestricted hydrogen model is significantly preferred over the true-deuterium model under both the canonical-parent and adversarial parent-reassignment statistics (largest calibrated $p = 0.001$ among the primary Gate 2/3 statistics, surviving $\alpha/6$). The key physical insight is that isotope identifiability is controlled by both Doppler regime and D column density: at the unsaturated LLS-class columns typical of D/H measurements, the intrinsic H/D profile-shape difference is sub-noise on a per-pixel basis, so the deuterium identification depends predominantly on the velocity offset—which the free-H model can reproduce without invoking deuterium at all. A multi-pronged evidence analysis (Gate 8) calibrates three discriminators—b-value ratio, velocity offset deviation, and column density difference—against true-D injection Monte Carlo ($N = 1000$ per sightline): the calibrated joint results are Q1009: $p = 0.001$; PKS1937: $p = 0.001$; J1332: $p = 0.023$. The two unsaturated systems give velocity-diagnostic probabilities Q1009: $p = 0.001$; PKS1937: $p = 0.001$. All three features are blueward of the parent, consistent with the conformal sign theorem (though the D-selected sample is not an independent sign test). A velocity-selection theorem explains the universality of the -81.6 km/s offset as a selection effect of the D-search protocol, and motivates an independently selected population test for non-isotope displaced H I components. An archival feasibility search establishes the principal saturation and Ly α -forest selection limits and identifies one candidate component in Q1009+2956, motivating a larger targeted survey. A pilot search of the three benchmark sightlines identifies one candidate feature already present in the literature scaffold near the predicted orphan phenotype in Q1009+2956: a narrow ($b = 10.3$ km/s, $\log N = 11.5$) H I component at $\Delta v = -155.5$ km/s from the parent (73.9 km/s from the isotope shift), identified as a candidate non-isotope component in the HIRES spectrum.

Conformal core-edge geometry predicts a blueward sign. An absorber-regime amplitude ansatz S_A^{abs} is introduced for diffuse absorbers; Cassini independently constrains the source-charge response $S_\Sigma^{(\odot)}$. Both are projections of the common microscopic scalar solution (Paper 0 §7): the derived kinetic operator $\mathcal{S}_\Sigma(\mathcal{E}) = [k\sqrt{\xi} + 2\xi]^{-1}$, $\xi = |X|/\Lambda_X^4$; the quoted

environment table uses the acceleration proxy $\xi_{\text{tab}} = (g/g_t)^2$ inherited from the foundational benchmark, rather than identifying that proxy with the solved-field kinetic invariant. Under this tabulation convention the operator returns $\mathcal{S}_\Sigma \simeq 1.6 \times 10^{-15}$ at 1 AU and $\mathcal{S}_\Sigma \sim 1.5\text{--}21$ at the Lyman-limit absorber ambient; evaluated instead on the solved drift-dominated field at the absorber epochs the same operator returns $\mathcal{S}_\Sigma \sim 0.13\text{--}0.18$ (step_21 drift-branch), so the environmental discrimination is supplied by one operator rather than assumed, under either field reconstruction. The conformal-only amplitude is $\sim 10^{-2} \text{ km s}^{-1}$; the temporal-disformal channel activated by the cosmological scalar time derivative was investigated as a candidate amplitude mechanism. It can reach the observed scale per sightline in the Einstein-frame parametrisation explored here. The power-law $B(\phi)$ forms tested so far do not achieve predictive closure across the three sightlines: $n = 2$ exceeds the Cassini bound only under a conservatively adopted solar source-charge profile (under the derived operator the disformal spatial term carries an $\mathcal{S}_\Sigma^2 \sim 10^{-48}$ suppression and the bound is non-binding), while $n = 5$ loses Lorentzian signature above $z \approx 3.4$. The full covariant calculation (Gate 10) with $\omega = -u^\mu k_\mu$ confirms that the near-criticality and cross-redshift signature loss persist under the complete matter-frame treatment when the drift is imported from the Λ CDM expansion history; the binding constraint is then the $A^2/\dot{\phi}^2$ variation across redshift, which prevents any single increasing power-law $B(\phi)$ from achieving simultaneous near-criticality at all three sightlines. That constraint is shown to be a property of the imported rate rather than of the disformal channel: the joint sign-and-closure criterion reduces to $d \ln H_T^2 / d \ln(1+z) < 0$ — a descent-type ambient drift — and the unsourced free-roll on the reconstructed corpus potential, $\dot{\phi}^2 = \bar{\Pi}_\infty^2 + \bar{\rho} e^{-\bar{\varphi}}$, satisfies it. Under the pure-descent limit the descent-tracking coupling $B(\varphi) = B_0 \varphi^4 / (\varphi_c^4 + \varphi^4) e^{-\varphi}$ has inverse-fit roots agreeing to a factor 1.00004 and satisfies the Cassini and GW170817 bounds; forward evaluation at that common B_0 does not reproduce the three amplitudes, and the branch is independently excluded by ambient-lapse self-consistency: the boundary factor $u = B\dot{\phi}^2/A^2$ is constant along the descent at $\approx 1 - 10^{-5}$, so the matter-metric redshift map $1+z = \tilde{N}(0)/\tilde{N}(\bar{\varphi})$ would require a negative ambient field to reproduce the observed absorber population. The unscreened temporal-disformal amplitude channel is thereby bounded sub-critical at these epochs; a contribution of the observed scale can arise only where $B(\phi)$ is environmentally screened. The -81.6 km/s clustering is explained by the velocity-selection theorem (Gate 6-I) as a selection effect of the D-search protocol, not as a universal TEP-produced shift. Within the adopted VMO-dominated reaction-flow model, the helium-4 abundance is reproduced as an equilibrium, while a conditional opacity theorem provides an observable high-redshift boundary without a physical plasma wall. Cosmological redshift is reformulated as temporal transport over a static spatial background. The broader TEP framework—CMB acoustic peak preservation, temporal-horizon regularity, cosmological distances without expansion, and JWST high-redshift phenomenology—is developed in companion papers.

Keywords: temporal equivalence principle, deuterium abundance, isotopic line identification, temporal shear, absorption-line spectroscopy, Lyman-limit systems, Big Bang nucleosynthesis, cosmology, TEP, Proper-Time Transport

1. Introduction

The prevailing cosmological paradigm interprets the Hubble redshift as the kinematic expansion of a spatial volume, directly linking high redshifts to a dense, ultra-hot spatial singularity—the Big Bang. Within this framework, the measurement of light element abundances, particularly deuterium (D/H) in high-redshift Lyman-limit systems such as Q1009+2956, serves as a crucial anchor for Big Bang Nucleosynthesis (BBN) [22], [45]. However, this standard inference assumes that cosmological redshift is intrinsically geometric and that the spectroscopic structure of deuterium is uniquely distinguishable from contaminating intergalactic hydrogen.

Tensions in precision cosmology—such as the Hubble tension [17], [18] and the S_8 growth tension—have motivated numerous theoretical extensions to the Λ CDM framework [16]. Recent literature has extensively explored modified gravity (e.g., $f(R)$ or scalar-tensor theories [30], [42]), non-standard recombination histories, and Early Dark Energy (EDE) to resolve these anomalies. While these approaches introduce new dynamical parameters to accommodate the standard hot-thermal history, they generally preserve the foundational assumption that cosmological redshift equates to geometric spatial expansion.

The Temporal Equivalence Principle Framework

The standard Friedmann-Lemaître-Robertson-Walker (FLRW) metric [11], [12], [13] explicitly couples cosmological evolution to a dynamic spatial volume. Extrapolating this geometric expansion backward inevitably terminates in a spatial singularity ($a \rightarrow 0$)—a regime where polynomial curvature invariants diverge and the foundational equations of General Relativity break down [5], [6], [7]. Alternative non-singular cosmologies, including bouncing models [8], [9], cyclic scenarios [10], and conformal approaches [37], [38], [39], have been explored but typically require additional dynamical ingredients. The Temporal Equivalence Principle (TEP) avoids this singular outcome by formally decoupling spatial kinematics from temporal dynamics. By anchoring the universe to a static physical matter frame ($a_m = 1$) governed by a dynamical proper-time field $A(\phi)$ [40], [41], the TEP geometry removes the spatial singularity. Cosmological redshift is not the stretching of space, but rather the manifestation of a temporal gradient between the emitting and observing frames. This approach shares conceptual ground with static-universe and

conformal frameworks [31], [32], [33], [34], but distinguishes itself through the dynamical proper-time field and its coupling to matter. The apparent "Big Bang" is replaced by an asymptotic temporal horizon ($\mathcal{I}^-$) in the far past, characterized by a vanishing relative clock rate ($A_{\text{clock}} \rightarrow 0$). Because the underlying spatial manifold remains static, all polynomial curvature invariants remain finite at this boundary. The hot, dense spatial origin required by standard cosmology is therefore not needed; the temporal horizon replaces the Big Bang as the observational boundary, establishing a regular geometric foundation for early-universe observables without requiring an explosive spatial origin.

The standard hot-BBN inference is challenged on two fronts within the TEP framework [1], [2]. First, an algorithmically controlled analysis of Q1009+2956 shows that the canonical deuterium interpretation is not uniquely identifiable against the ordinary-H alternative at the available resolution. Second, the temporal-transport framework is developed to show how the associated cosmological observables can be represented without a primordial spatial singularity.

2. Spectroscopic Analysis

The spectroscopic uniqueness of the deuterium identification in high-redshift quasar absorption systems is examined under Monte-Carlo-calibrated model comparison. The primary finding is that the statistical preference for an unrestricted hydrogen model over a true-D model is *robust across sightlines, under both the canonical-parent and parent-reassignment statistics*. The *parent-reassignment statistic* ($T_{\text{parent}}, p_{\text{parent}}$)—which gives D maximum freedom to select the most favorable parent—is the more conservative and informative number, and is reported as the primary result throughout. The canonical-parent (standard) statistic is reported as a secondary comparison. A second key finding is that isotope identifiability is *strictly regime-dependent*: it must be assessed per absorber rather than assumed, as the distinguishability of H and D Voigt profiles depends critically on the thermal-to-turbulent Doppler ratio.

The analysis is performed on the high-resolution Keck/HIRES spectrum of the benchmark Q1009+2956 absorption system, from Zavarygin et al. 2018 [44] (22 HIRES exposures co-added into four setups using the wspectrum software). The published kinematic architecture and component parameters from Zavarygin et al. 2018 [44] are used as the structural prior for the Voigt component family, fitted with the VPFIT framework [50]. The spectral data comprise four coadds at two instrumental settings: C1 ($R \approx 49\,000$, $\sigma_{\text{inst}} = 2.6$ km/s) and C5 ($R \approx 37\,000$, $\sigma_{\text{inst}} = 3.4$ km/s), with S/N ≈ 47 –78 per pixel near Ly α . The system was originally measured for primordial D/H by Burles & Tytler [46] and subsequently by Zavarygin et al. [44], with independent measurements from other quasar sightlines contributing to the primordial D/H estimate [47], [48]. Two additional benchmark sightlines—PKS 1937–101 (ESPRESSO, $z = 3.572$) and J1332+0052 (UVES+HIRES, $z = 3.421$)—are analysed with the identical framework in Section 2.4.

2.1 Isotope Identifiability Limits in Q1009+2956

Using the physical atomic registries for H I and D I (NIST ASD), synthetic deuterium was embedded at the actual Q1009 D column ($N_{\text{D}} = 5.7 \times 10^{12} \text{ cm}^{-2}$, derived from $\text{D/H} = 2.48 \times 10^{-5}$ and $\log N_{\text{HI}} = 17.362$) and recovered using unrestricted hydrogen models. Zavarygin et al. [44] themselves noted that the putative D feature is likely contaminated by weak interloping Ly α absorption from a low-column-density H I cloud, reducing the D/H precision. The present analysis extends this insight by quantifying the D/H degeneracy through explicit adversarial model comparison with Monte-Carlo-calibrated parent reassignment, comparing multiple sightlines, and then asking whether the ordinary-H alternative admits a TEP interpretation. The isotope identifiability depends on both the Doppler parameter regime and the D column density. For thermally-dominated gas, the H and D Doppler parameters differ by the mass ratio: $b_{\text{H}} = b_{\text{D}} \sqrt{m_{\text{D}}/m_{\text{H}}} \approx 1.41 b_{\text{D}}$. With this physical mass scaling, the maximum convolved flux discrepancy between H and D profiles across all Lyman transitions is 0.52σ at Keck/HIRES C1 resolution ($R \approx 49\,000$, $\sigma_{\text{inst}} = 2.6$ km/s), indicating that H and D are marginally distinguishable for thermally-dominated gas at this column. The σ statistic is computed as the maximum pixel-level flux residual divided by a conservative constant noise floor ($\sigma_{\text{floor}} = 0.05$), evaluated across the five strongest Lyman transitions (Ly α through Ly ϵ). For turbulence-dominated gas ($b_{\text{turb}} \gg b_{\text{therm}}$), the mass-dependent thermal contribution is negligible and $b_{\text{H}} \approx b_{\text{D}}$; in this limit the maximum discrepancy falls to 0.0002σ , making the isotopes operationally unidentifiable. For the actual Q1009 fitted parameters ($T_{\text{K}} = 2000$ K, $b_{\text{turb}} = 11.60$ km/s, giving $b_{\text{turb}}/b_{\text{therm}} \approx 2.02$), the mixed-regime discrepancy is 0.10σ — far below conventional detectability. The identifiability is controlled by the D column: at the unsaturated LLS-class columns typical of D/H measurements ($\log N_{\text{D}} \lesssim 13$), the single-pixel discrepancy is sub-threshold in every Doppler regime; only saturated sub-DLA columns ($\log N_{\text{D}} \gtrsim 14$) approach distinguishability. Isotope identifiability must therefore be assessed per absorber as a function of both column and Doppler regime.

2.2 Likelihood Model Comparison and Significance Testing

A complete structural analysis of the Q1009+2956 spectrum is performed under a controlled constrained-versus-unrestricted model hierarchy embedded in the joint D+H model. By anchoring the fits to a SHA-256-validated data manifest, the likelihood surfaces of

the standard D-interpretation (M_D), an unrestricted hydrogen interpretation ($M_{H,\text{free}}$), and the joint space (M_{D+H}) are mapped. The 142 shared parameters (43 H I Voigt component parameters: centres, b -values, and column densities, plus continuum basis coefficients for each data block) are held at the published literature architecture ^[44], identically for every hypothesis. This fixed scaffold is part of the definition of the test statistic rather than an uncontrolled nuisance choice: the significance of T is not read off an asymptotic χ^2 distribution but calibrated by parametric bootstrap in Section 2.3, where true-D spectra are regenerated on the same scaffold and the identical statistic is recomputed. Because the scaffold enters the observed and simulated statistics identically, the resulting p -value is calibrated conditional on the adopted published architecture. Q1009 scaffold freedom is tested explicitly in §5.5.

The candidate component parameter structure is made fully explicit. Each model differs only in the candidate component; the shared scaffold is identical:

Model	Candidate params	v	$\log N$	T_K	b_{turb}	k_{cand}
M_D (M_{Dfree})	D tied to parent	tied to parent + v_{iso}	free	free	free	3
$M_{H,\text{free}}$ (M_{H})	unrestricted H	free	free	free	free	4
M_{D+H}	D + unrestricted H	D tied, H free	D free, H free	shared	shared	5
M_{kin}	parent-width-constrained H control	free	free	—	—	2

The parameter count difference is $\Delta k = k(M_{H,\text{free}}) - k(M_D) = 4 - 3 = 1$: the unrestricted-H model gains one free parameter (the velocity v_H , which in M_D is tied to the parent velocity via $v_D = v_{\text{parent}} + v_{\text{iso}}$). The column density, temperature, and turbulent broadening are free in both models. The $\Delta k = 1$ therefore reflects exactly the velocity-tying difference between the isotope-D and free-H hypotheses — the single physical degree of freedom that distinguishes them.

Model	Candidate interpretation	k	$\ln L_{\text{max}}$	ΔAIC	Nested?
M_D	D tied to parent H	3	−15446.77	0	—
$M_{H,\text{free}}$	unrestricted H	4	−15366.26	−159.03	observationally embeds M_D at Q1009 precision
M_{D+H}	D + unrestricted H	5	−15366.26	−153.03	contains $M_{H,\text{free}}$ if $N_D \rightarrow 0$ (verified to 2.8×10^{-10} in $\Delta \ln L$)

Statistical Significance

The unrestricted hydrogen model—which supplies the displaced-H phenomenology to which the candidate TEP edge-gas interpretation is later applied—provided a superior description of the data, yielding a likelihood improvement:

$$\Delta \ln L = 80.51, \quad T = 161.03, \quad \Delta k = 1, \quad \Delta\text{AIC} = -159.03. \quad (1)$$

The unrestricted-H model adds one free parameter (the H velocity) relative to the D-tied model ($\Delta k = 1$), yet the likelihood improvement of $\Delta \ln L = 80.51$ exceeds the AIC penalty of 2, yielding $\Delta\text{AIC} = -159.03$ and $\Delta\text{BIC} = -151.97$ ($n = 8,610$ data points, $\ln n \approx 9.06$). Both information criteria favor $M_{H,\text{free}}$. Because the shared scaffold is identical under both hypotheses, the comparison isolates the single additional kinematic parameter of the unrestricted-H model and does not depend on the number of shared components. The information criteria are reported for orientation only; the significance quoted below rests on the Monte Carlo calibration, not on asymptotic likelihood-ratio theory, whose regularity conditions are not satisfied here because M_D sits on the boundary $N_D \rightarrow 0$ of M_{D+H} . Monte Carlo p -values use the Laplace-corrected estimator $\hat{p} = (k + 1)/(N + 1)$, where k is the exceedance count and N is the number of realizations.

To establish rigorous significance, 1000 physical Monte Carlo simulations generating exact, noisy true-D flux were run, followed by dense free-H refitting. Two calibration statistics are reported. The *standard statistic* locks D to the canonical parent (matching the observed fit). The *parent-reassignment statistic* allows D to select the most favorable parent from the H I components within ± 1000 km/s of the absorber systemic velocity. The standard statistic represents the prespecified comparison; the parent-reassignment statistic represents the most adversarial test. Both are reported to bracket the significance honestly.

2.3 Leave-One-Out and Parent Reassignment Robustness

The component misattribution vulnerability was exhaustively tested by tying the candidate D velocity to every eligible H component within the conservative $\pm 1000 \text{ km s}^{-1}$ parent window (14 parent candidates). The maximum alternative-parent test statistic is defined as:

$$T_{\text{parent}} = 2 \left[\ln L(M_{H,\text{free}}) - \max_j \ln L(M_D | j) \right]. \quad (2)$$

Against the most advantageous alternative parent assignment, the free-H interpretation is preferred for Q1009+2956 with $T_{\text{parent}} = 161.03$, equal to $T = 161.03$ for the canonical parent. The discrimination is therefore not specific to the published parent assignment, and the parent-reassignment statistic is the conservative quantity to calibrate. When calibrated inside the true-D Monte Carlo loop—utilizing bounded L-BFGS-B optimization with a multi-start grid spanning the full velocity range $[-160, +50] \text{ km/s}$, $\log N_D \in [0, 20]$, $T_K \in [100, 4 \times 10^4] \text{ K}$, and $b_{\text{turb}} \in [1, 50] \text{ km/s}$, with the best likelihood retained across all starts to prevent local minima from artificially widening the null distribution—the empirical p_{parent} is reported below.

Parent-reassignment window. Two parent-reassignment windows are reported. The *conservative* window ($\pm 1000 \text{ km/s}$) gives the deuterium hypothesis maximum freedom to find a favorable parent, establishing a lower bound for the T_{parent} statistic. Of the 43 H I components in the Q1009 model, 14 fall within this window, of which 4 satisfy the column-density eligibility cut ($\log N_{\text{HI}} \geq \log N_{\text{max}} - 2 \text{ dex}$); none are DLA-class ($\log N_{\text{HI}} > 20.3$), and components from unrelated absorber systems at $v \sim -47000$ to -74000 km/s are correctly excluded. The *physical* window ($\pm 3b_{\text{max}}$, where $b_{\text{max}} = 56.7 \text{ km/s}$ is the largest Doppler parameter of the main complex) restricts eligible parents to those within the kinematic structure of the absorber: 9 H I components fall within the velocity window, of which 4 satisfy the column-density eligibility cut ($\log N_{\text{HI}} \geq \log N_{\text{max}} - 2 \text{ dex}$). The conservative $\pm 1000 \text{ km/s}$ window is the cross-sightline headline statistic; the $\pm 3b_{\text{max}}$ physical window has also been calibrated for all three sightlines (Section 5.5), with results consistent with the conservative window. Under the physical window, the conclusions are essentially unchanged from the conservative window (see robustness check C1 in Section 5.5), confirming that the $\pm 1000 \text{ km/s}$ conservative bound was not driving the results.

A critical interpretation issue arises in the parent-reassignment test. For Q1009+2956, $T_{\text{parent}} = T_{\text{obs}} = 161.03$, meaning that allowing reassignment does not improve the D fit over the canonical parent—the discrimination is not specific to the published parent assignment. For PKS 1937–101, $T_{\text{parent}} = 59.00 < T_{\text{obs}} = 161.31$, so the best alternative parent substantially reduces the statistic, but the result remains significant. For J1332+0052, $T_{\text{parent}} = 135.76 < T_{\text{obs}} = 136.97$, so the best alternative parent slightly reduces the statistic. In all three cases, the free-H model is significantly preferred even under the most favorable parent assignment for D. The parent-reassignment statistic is therefore the conservative quantity to calibrate, and the p_{parent} values reported below should be interpreted as the significance under the most favorable parent assignment for D.

Parent-identity diagnostic (Gate 7). Under the isotope-D interpretation, the best alternative parent j^* (the H I component maximizing $\ln L(M_D | j)$) should be the main narrow thermal component — D is tied to the gas it formed beside. Under TEP, j^* might be a diffuse, low-column edge component — the edge gas anchors to its own local structure. For all three sightlines, the best non-canonical parent within the eligible window ($\log N_{\text{HI}} \geq \log N_{\text{max}} - 2 \text{ dex}$, $|v| \leq 1000 \text{ km/s}$) is a core component: Q1009+2956 parent 1 ($v = -43.7 \text{ km/s}$, $\log N_{\text{HI}} = 16.80$, $b = 17.8 \text{ km/s}$), PKS 1937–101 parent 20 ($v = +17.9 \text{ km/s}$, $\log N_{\text{HI}} = 17.38$, $b = 17.1 \text{ km/s}$), and J1332+0052 parent 2 ($v = +8.7 \text{ km/s}$, $\log N_{\text{HI}} = 18.39$, $b = 15.8 \text{ km/s}$). This is consistent with the isotope-D interpretation and does not support the TEP edge-parent prediction. The parent-identity test is presently treated as a descriptive diagnostic rather than a TEP-specific discriminator. Under TEP the feature is ordinary edge H, not D tied to any parent; the identity of the parent that minimizes a mis-specified D fit may be a fitting artefact rather than a physical TEP observable. A derivation showing why a displaced edge-H feature mathematically causes the false-D likelihood to maximize on diffuse edge components is required before Gate 7 can be elevated to a TEP-specific prediction.

Offset-exactness test (Gate 4E). Under the isotope-D interpretation, the free-H component should sit at the exact isotope-shifted velocity $v_{\text{iso}} = -81.6 \text{ km/s}$ relative to its parent. Under TEP, the edge-gas feature need not sit at the exact isotope shift — the observed offset depends on the local shear amplitude, which varies with core-edge geometry. The fitted $M_{H,\text{free}}$ velocities deviate from the isotope-shifted position by:

Sightline	Parent v (km/s)	v_{iso} (km/s)	Expected D v (km/s)	Fitted v_H (km/s)	Deviation (km/s)
Q1009+2956	−52.4	−81.6	−134.0	−131.3	+2.7
PKS 1937–101	+8.7	−81.6	−72.9	−67.3	+5.7
J1332+0052	0.0	−81.6	−81.6	−118.2	−36.6

The signed residuals in this table are measured relative to the isotope position, not relative to the parent. Thus the +2.7 and +5.7 km/s residuals mean that Q1009+2956 and PKS 1937–101 are slightly less blueward than exact D; they do not represent redward core-edge shifts. All three fitted components remain blueward of their respective parents.

Q1009+2956 sits 2.7 km/s from the isotope shift — lying approximately one instrumental Gaussian width ($\sigma_{\text{inst}} \simeq 2.6 \text{ km s}^{-1}$) from the isotope position, but highly inconsistent when calibrated against the true-D simulation distribution ($\pm 0.2 \text{ km/s}$, $p \leq 0.001$, Section 5.6). PKS 1937–101 deviates by 5.7 km/s — also highly significant after calibration ($p \leq 0.001$). J1332+0052 deviates by –36.6 km/s — not individually significant after calibration ($p = 0.098$). The cross-sightline pattern — increasing offset deviation with increasing D column — is quantitatively testable: if the feature is isotope-D, all three should sit at v_{iso} within spectral resolution; under TEP, the offset varies with core-edge geometry. The J1332 feature shows the largest departure from the isotope-shifted position, and it is also the sightline with the highest D column (saturated, $\log N_{\text{D}} \approx 14.65$), where the edge-gas interpretation is most physically motivated.

Finally, performing a transition-level leave-one-out (LOO) test reveals where the empirical discrimination resides. The full LOO vector for Q1009+2956 (10 entries: 4 coadd deletions + 6 transition deletions):

Excluded	Type	T_{LOO}	ΔT vs T_{full}
C1x1 coadd	coadd	125.05	–35.98
C1x2 coadd	coadd	141.03	–20.00
C5x1 coadd	coadd	77.31	–83.71
C5x2 coadd	coadd	139.87	–21.16
Ly α	transition	34.15	–126.88
Ly β	transition	148.64	–12.39
Ly γ	transition	156.16	–4.86
Ly6	transition	160.70	–0.33
Ly13	transition	161.05	+0.02
Ly14	transition	161.03	+0.00

The full LOO vector reveals that the discrimination is strongly concentrated in Ly α : removing Ly α decreases T from 161.03 to 34.15, indicating that Ly α provides the bulk of the isotope discrimination. This is physically expected: the isotope velocity shift (–81.6 km/s) is the same across all Lyman transitions, but Ly α has the highest oscillator strength and signal-to-noise ratio, yielding the deepest absorption profile and the largest number of informative pixels over which the mass-scaled Doppler difference between H and D accumulates. Ly β is the second most important transition ($T_{\text{Ly}\beta} = 148.64$, $\Delta T = -12.39$), followed by Ly γ ($\Delta T = -4.86$). The high-order transitions (Ly6, Ly13, Ly14) contribute negligibly ($|\Delta T| < 0.5$). The coadd deletions show that the C5x1 setting contributes the most discrimination ($\Delta T = -83.71$), followed by C1x1 ($\Delta T = -35.98$). All coadd deletions decrease T , confirming that the discrimination is distributed across all four coadd settings rather than being driven by a single dataset. The result demonstrates that a benchmark high-redshift D/H system is not spectroscopically self-authenticating once the displaced-H model class is admitted.

Velocity-selection theorem (Gate 6). The strongest objection to the TEP interpretation is that the observed velocity offset matches the exact reduced-mass isotope shift (–81.6 km/s) in all three sightlines. However, this coincidence is a property of the *selection function*, not of the gas: the D-identification protocol is a filter $|\Delta v - v_{\text{iso}}| < \epsilon$ — VPFIT initializes D components at the isotope-shifted velocity, and analysts search for D at that specific offset. Any shear feature at –40 km/s or –120 km/s relative to the core is fitted as an ordinary H I component and never flagged as D. The clustering near –81.6 km s^{–1} is therefore strongly conditioned by construction of the D-selected sample and cannot be treated as an independent chance-probability test of isotope identity. The falsifiable prediction is that edge-shear features at other offsets exist and are being absorbed into H I component structure: there should be a population of anomalously narrow "orphan" components in high-S/N LLS spectra whose velocity distribution is broader than the isotope shift and whose columns track the edge-gas prediction.

Illustrative Population Forecast (Gate 6-II). An illustrative selection model, calibrated to the three benchmark fitted offsets, gives $\sigma_{\text{shear}} \approx 19 \text{ km/s}$ and an orphan fraction of $\sim 79\%$. Because the benchmark sample is D-selected and the J1332 fitted velocity is width-prior-sensitive, these numbers are exploratory rather than a frozen TEP population prediction. The observed cross-sightline deviations (+2.7, +5.7, –36.6 km/s for Q1009, PKS 1937, J1332 respectively) give: the D-search selection function $|\Delta v - v_{\text{iso}}| < 5 \text{ km/s}$ captures only $\sim 21\%$ of edge-shear features; the remaining $\sim 79\%$ are orphans absorbed into the H I

component structure. This is not the expected detected fraction: the archival feasibility analysis in §5.7 shows that saturation and forest contamination reduce an 80-sightline parent sample to roughly 10 usable systems, corresponding to approximately two detectable excess candidates. Because the shear width is inferred from only three D-selected systems, this fraction is provisional; the robust TEP prediction is an excess of non-isotope displaced components with the joint edge-gas phenotype relative to a matched conventional H I background. In a sample of ~ 80 high-S/N LLS sightlines with full Lyman coverage (SQUAD DR1, UVES SLS, ESPRESSO D/H survey), ~ 24 shear features are expected, of which ~ 5 are identified as D and ~ 19 are orphans. The predicted orphan surface density is ~ 2.4 per 10 sightlines at $|\Delta v - v_{\text{iso}}| > 5$ km/s, falling to ~ 1.8 per 10 sightlines at > 10 km/s. Orphan features are predicted to have narrow Doppler parameters ($b < 10$ km/s, cool edge gas), low columns ($\log N_{\text{HI}} \sim 12\text{--}14$), and thermal Doppler regimes — distinguishable from the normal H I population by an excess of narrow, low-column components at non-isotope offsets. Under the isotope-D null, no such excess is expected. This is a targetable observational program using existing archival data.

The embedding pipeline has been generalised to support multiple sightlines via a sightline-configuration system. Two additional benchmark D/H quasar sightlines have now been analysed with the identical model-comparison framework, providing the first cross-sightline application of the TEP spectroscopic methodology. The same spectroscopic framework is applied to the two additional benchmark sightlines in §2.4; their cross-sightline amplitude comparison is presented in §3.3.

2.4 Cross-Sightline Validation

The identical Gate 2/3 pipeline was applied to PKS 1937–101 (Cooke et al. 2024 ^[66]; ESPRESSO, $R = 70\,000$, $z_{\text{abs}} = 3.572$) and J1332+0052 (Kislitsyn et al. 2024 ^[67]; joint UVES+HIRES, $R \approx 50\,800$ and $49\,200$, $z_{\text{abs}} = 3.421$). Both sightlines were analysed with all available Lyman transitions covered by their spectral data (13 for PKS 1937–101, 12 for J1332+0052 which lacks Ly α coverage), the same noise-model calibration, and the same 1000-realization Monte Carlo protocol.

Sightline	z_{abs}	Instrument	T_{obs}	T_{parent}	p_{std}	p_{parent}	Doppler regime	Interpretive status (conservative)
Q1009+2956	2.5042	Keck/HIRES	161.03	161.03	0.001 (0/1000)	0.001 (0/1000)	Mixed ($b_{\text{turb}}/b_{\text{therm}} \approx 2.02$, 0.10σ)	Significant; free-H preferred even under most favorable parent
PKS 1937–101	3.572	ESPRESSO	161.31	59.00	0.001 (0/1000)	0.001 (0/1000)	Thermal	Significant; free-H preferred even under most favorable parent
J1332+0052	3.421	UVES+HIRES	136.97	135.76	0.001 (0/1000)	0.001 (0/1000)	Turbulent, saturated ($b_{\text{turb}} \approx 50$ km/s at bound, no Ly α coverage; boundary check in §5.6)	Significant; free-H preferred even under most favorable parent

[†] For J1332+0052, the canonical-parent standard statistic is reported as the primary p_{std} for consistency with Q1009 and PKS 1937. Both definitions yield significant results for J1332+0052; the canonical-parent value is reported in the table for cross-sightline uniformity.

Under the conservative parent-reassignment statistic, the results are: for Q1009+2956, $p_{\text{parent}} = 0.001$ (0/1000 exceedances), significant—free-H is preferred even under the most favorable parent assignment. For PKS 1937–101, $p_{\text{parent}} = 0.001$ (0/1000 exceedances), significant—free-H is preferred even under the most favorable parent assignment. For J1332+0052, $p_{\text{parent}} = 0.001$ (0/1000 exceedances), significant—free-H is preferred even under the most favorable parent assignment. The standard (canonical-parent) statistic yields $p_{\text{std}} = 0.001$ (0/1000) for all three sightlines. The common free-H preference across all three sightlines provides the displaced-H phenomenology to which the TEP interpretation is subsequently applied. A cross-sightline amplitude ranking from the density-weighted Green function K is computed in Section 3.3 (Gate 5).

Multiplicity correction. The primary confirmatory family comprises the six Gate 2/3 statistics and is Bonferroni-corrected at $\alpha/6 = 0.0083$; all six survive. Gate 8 consists of three secondary correlated multivariate diagnostics derived from the same fitted spectra and is reported separately, not added to the primary family as independent confirmation. Correcting the three Gate-8 sightline joint tests among themselves gives $\alpha/3 = 0.0167$; the corresponding calibrated probabilities are Q1009: $p = 0.001$, PKS1937: $p = 0.001$, J1332: $p = 0.023$. Q1009, PKS1937 pass that threshold. Parent maximization is already inside T_{parent} and its Monte Carlo calibration; it is not a further Bonferroni factor.

Three critical analysis issues were identified and fixed during this work: (1) the continuum basis was made dynamic — 10 parameters for broad Ly α regions (where the H I wing requires flexible continuum separation from the narrow D feature) and 5 parameters for narrow metal-line regions (where 10 parameters would overfit and act as a mathematical sponge, absorbing isotopic discrepancies and artificially depressing T); (2) the isotope identifiability calculation was corrected to use mass-scaled Doppler parameters at the actual

D column ($N_D = 5.7 \times 10^{12} \text{ cm}^{-2}$), revealing that H and D are marginally distinguishable at 0.52σ in the thermal regime and indistinguishable at 0.10σ in the mixed regime; and (3) the parent-reassignment interpretation was corrected: a non-significant p_{parent} indicates that D with an alternative parent is statistically indistinguishable from free-H, which does not exclude the deuterium interpretation but does not independently support it either.

3. Scalar Field Dynamics and Temporal Shear [53]

Part I: Local Absorber Physics. Sections 2–3 establish the spectroscopic analysis and the local temporal-shear mechanism for the deuterium-like velocity offset. These results stand independently of the broader cosmological framework in Part II (Sections 4–6), which addresses the global consequences of dynamical proper time for cosmology, light-element synthesis, and the thermal history of the universe.

The free-H preference motivates a separate physical question: whether TEP can provide a non-kinematic contribution to the displaced-H phenomenology. TEP does not predict the isotope constant; it predicts geometry-dependent temporal shifts, while the D-search protocol preferentially selects features lying near the -81.6 km s^{-1} isotope window.

3.1 Field-Theoretic Derivation

In TEP, gravity is governed by a Lorentzian metric $g_{\mu\nu}$, while matter couples to a causal effective metric $\tilde{g}_{\mu\nu}$ determined by the scalar field ϕ , analogous to screened scalar-field frameworks [35], [36]. The interaction is defined by the action:

The following absorber calculation uses a canonical comparison benchmark of the common architecture, rather than the low-kinetic limit of the selected sector. The current Paper-0 screening realization is $P(X, \phi) = -V(\phi) + \frac{2k}{3\Lambda_X^2} X|X|^{1/2} + X|X|/\Lambda_X^4$; the canonical Green-function calculation and its adopted screening profile are not an independent derivation of that nonlinear response.

$$S_{\text{TEP}} = \int d^4x \sqrt{-g} \left[\frac{M_{\text{Pl}}^2}{2} R - \frac{1}{2} \nabla_\mu \phi \nabla^\mu \phi - V(\phi) \right] + S_m[\tilde{g}_{\mu\nu}, \psi], \quad (3)$$

where the canonical TEP matter metric is $\tilde{g}_{\mu\nu} = -A_{\text{loc}}^2(u_{\text{grav}}(\mathbf{x}))U_\mu U_\nu + A_{\text{sp}}^2(u_{\text{drift}}(t, \mathbf{x}))\Pi_{\mu\nu} + B(\phi)\Pi_\mu^\alpha \Pi_\nu^\beta \nabla_\alpha \phi \nabla_\beta \phi$ ($U_\mu \equiv -\nabla_\mu \bar{\phi}/\sqrt{2\bar{X}}$, $\Pi_{\mu\nu} = g_{\mu\nu} + U_\mu U_\nu$; Paper 0, §2.2, Step 80), written in conformal–disformal shorthand as $\tilde{g}_{\mu\nu} = A^2(\phi)g_{\mu\nu} + B(\phi)\nabla_\mu \phi \nabla_\nu \phi$. In the static weak-field limit relevant to absorber clouds, the disformal term $B(\phi)\nabla_\mu \phi \nabla_\nu \phi$ contributes to spatial geodesics and light propagation but is negligible for $g_{\mu 0}$ (taking the local perturbation $\partial_0 \delta\phi \simeq 0$; the non-zero cosmological background derivative $\partial_0 \bar{\phi}$ is treated separately in Gate 4B-IV), so it does not affect clock rates directly. Varying this action with respect to ϕ ($\frac{\delta S_{\text{TEP}}}{\delta \phi} = 0$) over a localized static absorber yields the scalar equation of motion:

Equation of Motion:

$$(-\nabla^2 + m_{\text{eff}}^2) \delta\phi = -\frac{\beta_A}{M_{\text{Pl}}} A(\bar{\phi}) \rho_*, \quad (4)$$

where the source is the axiom term $\rho_* A_{,\phi} = \rho_* \beta_A A(\bar{\phi})/M_{\text{Pl}}$: ρ_* is the conserved Einstein-frame density and the ambient conformal factor $A(\bar{\phi}) = e^{-\bar{\phi}/M_{\text{Pl}}} \simeq 0.2\text{--}0.3$ at the absorber redshifts modulates the effective coupling (an earlier bare- ρ convention overstated the sourced depths by $e^{+\bar{\phi}} \sim 4$; Gate 10 uses the corrected source throughout). $m_{\text{eff}}^2 = V_{\text{eff}}''(\bar{\phi}) > 0$ is the effective scalar mass squared, obtained from the second derivative of the effective potential expanded about the background value $\bar{\phi}$. In the massless limit of this canonical benchmark, G_+ reduces to the Poisson kernel; positive m_{eff}^2 limits its range. The potential-depth absorber ansatz is a separate declared response prescription, not an identity that converts potential depth into the kinetic stiffness or scalar mass. The bare conformal coupling $\beta_A = -1.0$ ^{[3], [54]} applies in the unscreened absorber regime; the Solar-System source-charge response enters the PPN deviation linearly, $|\gamma_{\text{PPN}} - 1| = 2\beta_A^2 S_{\Sigma}^{(\odot)}/(1 + 2\beta_A^2 S_{\Sigma}^{(\odot)})$ under the unified unstretched-space metric (or $4\beta_A^2 S_{\Sigma}^{(\odot)}/(1 + 2\beta_A^2 S_{\Sigma}^{(\odot)})$ on the isotropic benchmark), so Cassini requires $S_{\Sigma}^{(\odot)} \lesssim 1.15 \times 10^{-5}$ (5.8×10^{-6} on the isotropic benchmark) — a conservative adopted source-charge profile ($S_{\Sigma}^{(\odot)} \sim 2 \times 10^{-6}$; the derived Paper-0 §7 operator returns $S_{\Sigma}^{(\odot)} \simeq 1.6 \times 10^{-15}$ at 1 AU, step_21; this single-source PPN quantity is not interchangeable with the pairwise response), not by extrapolating the absorber-regime ansatz to solar depth; the illustrative Gate-4B consistency profile below retains the conservative $S_{\Sigma}^{(\odot)} \sim 2 \times 10^{-6}$.

Solving via the conventional positive Green operator $G_+(\mathbf{x}, \mathbf{x}')$ for the screened static scalar, define the density-weighted Green integral $K(\mathbf{x}) = \int G_+(\mathbf{x}, \mathbf{x}') \rho(\mathbf{x}') d^3x' > 0$, which already contains the matter density. The core scalar response is then $\delta\phi_{\text{core}} = -(\beta_A/M_{\text{Pl}}) K$, confirming that the local clock rate $A(\phi)$ is systematically deformed inside the cloud relative to the cosmological background. This deformation supplies the conformal clock contribution to the candidate spectral shift quantified below.

Screening convention. This paper uses the absorber-regime amplitude projection $S_A^{\text{abs}}(|\Phi|/c^2)$, for which diffuse Lyman-limit systems are effectively unscreened. This is distinct from the source-charge projection S_{Σ} constrained by Solar-System tests. The relation between these projections is supplied by the common microscopic scalar solution (Paper 0 §7): on the corpus-selected two-branch kinetic sector, $\mathcal{S}_{\Sigma}(\mathcal{E}) = 1/P_X(X_{\text{env}}) = [k\sqrt{\xi} + 2\xi]^{-1}$ with $\xi = (g/g_t)^2$, $k \simeq 16.0$, and $g_t = cH_0/(2|\beta_A|) \simeq 3.4 \times 10^{-10} \text{ m s}^{-2}$, evaluated at the embedding ambient field. The same operator returns $\mathcal{S}_{\Sigma} \simeq 1.6 \times 10^{-15}$ at 1 AU — nine orders inside the Cassini ceiling — and $\mathcal{S}_{\Sigma} \simeq 1.5\text{--}21$ at the LLS ambient, on the unscreened-to-anti-screening branch (step_21, `results/step_21_canonical_projection.json`). The numerical absorber law used below is therefore a conservative benchmark amplitude prescription — it understates the derived response by an order-ten factor — not a universal screening operator.

3.2 Sign Provenance and the Core–Edge Geometry

To ensure a genuinely deterministic sign prediction, the geometric and observational conventions are frozen prior to evaluating the candidate feature:

- **Line-of-Sight (LOS) Orientation:** Positive outward from the observer.
- **Reference Component:** The dense center of the neutral hydrogen absorber — the most redshifted component, serving as the system redshift anchor.
- **Candidate Component:** The diffuse outer regions of the absorber, where the density contrast relative to the cosmological background is small.
- **Fundamental Coupling:** The conformal coupling β_A ($A = e^{\beta_A \phi/M_{\text{Pl}}}$) enters the clock shift quadratically ($\propto \beta_A^2$). The sign of the blueward displacement is therefore robust to the sign of β_A : any non-zero linear conformal coupling produces the same blueward sign. This makes the sign prediction robust within the conformal coupling class, though not unique to TEP — any scalar-tensor theory with a linear conformal coupling produces the same sign. The TEP-specific discriminator is the amplitude (Gate 4B), not the sign.
- **Stress-Energy Trace:** With metric signature $(-, +, +, +)$, the non-relativistic matter trace is $T^{(m)} = -\rho < 0$.
- **Definition of Difference:** $\Delta \ln A_{\text{abs}} \equiv \ln \frac{A(\phi_{\text{edge}})}{A(\phi_{\text{core}})}$. This separates the local absorber shear ($\Delta \ln A_{\text{abs}}$) from the global cosmological endpoint map (A_{clock}).
- **Velocity Sign Convention:** $\Delta v_T \simeq -c \Delta \ln A_{\text{abs}}$ (with $\Delta v_T < 0$ defined as blueward relative to the core).

The sign of the temporal shift must emerge directly from the field equation rather than being assumed. Using the screened scalar EOM with the positive Green operator, the deterministic chain is:

$$\begin{aligned}
K &= \int G_+(\mathbf{x}, \mathbf{x}') \rho(\mathbf{x}') d^3x' > 0, \\
\delta\phi_{\text{core}} &= -\frac{\beta_A}{M_{\text{Pl}}} K, \\
\Delta \ln A_{\text{core}} &= \frac{\beta_A}{M_{\text{Pl}}} \delta\phi_{\text{core}} = -\frac{\beta_A^2 K}{M_{\text{Pl}}^2} < 0.
\end{aligned}$$

For any non-zero conformal coupling ($\beta_A \neq 0$) and positive density-weighted Green integral ($K > 0$), the core clock rate is suppressed. Crucially, the dependence is quadratic in β_A : the sign of $\Delta \ln A_{\text{core}}$ does not depend on the sign of β_A , but follows from $\beta_A^2 > 0$.

$$\Delta \ln A_{\text{core}} = -\frac{\beta_A^2 K}{M_{\text{Pl}}^2} < 0. \quad (5)$$

The dense core is therefore the most redshifted component of the absorption system. The diffuse outer regions, where the density contrast relative to the background is small, experience a negligible clock shift ($\Delta \ln A_{\text{edge}} \approx 0$). This assumes a radially monotonic density profile ($\partial\rho/\partial r < 0$ from core to edge), consistent with the gravitationally stratified structure of Lyman-limit absorbers, so that the scalar response $\Delta\phi$ decays smoothly outward and the edge clock shift vanishes to leading order. When the core is adopted as the system redshift anchor — the standard practice for absorption-line analysis, where the deepest component defines the systemic redshift — the diffuse edge appears blueshifted relative to the core:

$$\Delta \ln A_{\text{abs}} = \ln \frac{A(\phi_{\text{edge}})}{A(\phi_{\text{core}})} \approx -\Delta \ln A_{\text{core}} > 0, \quad \Delta v_T = -c \Delta \ln A_{\text{abs}} < 0 \text{ (blueward)}. \quad (6)$$

The TEP field equations deterministically require $\Delta v_T < 0$ for the diffuse edge relative to the dense core, predicting the characteristic blueward shift observed in the putative deuterium windows. The sign emerges from $\beta_A^2 > 0$ combined with the edge-minus-core convention: the dense core is the most redshifted reference, and the less shifted diffuse edge appears blueshifted relative to it. The sign does not depend on the sign of β_A directly, but follows from the quadratic coupling and the observational convention that the core defines the systemic redshift.

Treating the sign and amplitude as distinct predictions, the conformal absorber solution fixes the blueward sign of the core-edge clock deformation. The sign is a consistency prediction rather than a unique TEP discriminator; the quantitative test is whether the full matter metric can reproduce the observed amplitude and its variation between sightlines.

3.3 Amplitude Test: Frozen- β_A Calculation (Gate 4B)

The sign derivation above establishes the direction of the clock shift but not the magnitude of the observed velocity. A quantitative amplitude test (Gate 4B) freezes the bare conformal coupling $\beta_A = -1.0$ from the TEP corpus [3], [54] and constructs the density-weighted Green integral K from actual Q1009+2956 observables. The Q1009+2956 D/H absorber is a Lyman Limit System (LLS) with $\log N_{\text{HI}} = 17.362 \pm 0.005$ [44], $z_{\text{abs}} = 2.5042$.

Because the scalar field couples to the *total* mass density ρ (not just neutral hydrogen), N_{HI} cannot be used directly. An ionisation correction is required to derive the total hydrogen column N_{H} . Under photoionisation equilibrium with the UV background at $z \sim 2.5$ [55], the neutral fraction is $x_{\text{HI}} = \alpha_{\text{rec}} n_e / \Gamma$, where $\Gamma \approx 10^{-12} \text{ s}^{-1}$ is the photoionisation rate and $\alpha_{\text{rec}} \approx 2.6 \times 10^{-13} \text{ cm}^3/\text{s}$ is the recombination coefficient at $T \sim 10^4 \text{ K}$. Solving for N_{H} with $n_e \approx n_{\text{H}} = N_{\text{H}}/(2R)$:

$$N_{\text{H}} = \sqrt{\frac{N_{\text{HI}} \cdot 2R \cdot \Gamma}{\alpha_{\text{rec}}}}. \quad (7)$$

For a uniform sphere with the standard Poisson Green function, $K_{\text{Poisson}}/M_{\text{Pl}}^2 = 2|\Phi|/c^2$, where $|\Phi| = \pi G m_p N_{\text{H}} R$ is the Newtonian potential at the centre. The conformal-only prediction is:

$$|\Delta v_{\text{conf}}| = c \beta_A^2 \frac{K_{\text{Poisson}}}{M_{\text{Pl}}^2} = \frac{2\beta_A^2 \pi G m_p N_{\text{H}} R}{c}. \quad (8)$$

The absorber radius R is not directly measured for Q1009+2956. The main D/H components span $\Delta v \approx 13$ km/s in velocity [44], which constrains the absorber size through the virial relation but does not determine it uniquely. Results are therefore presented as a function of R . For a representative $R = 30$ kpc, the uniform-sphere ionisation correction gives $\log N_{\text{H}} \approx 20.6$ ($x_{\text{H}} \approx 6 \times 10^{-4}$), yielding $|\Delta v_{\text{conf}}| \approx 0.009$ km/s. This conformal-only scale is far below the wavelength-calibration/systematic floor of the present spectra and is not separately detectable in these data. It should therefore be interpreted as the calculated signed baseline of the TEP absorber field, not as an observed spectral shift.

Cross-sightline amplitude ranking (Gate 5). The conformal-only amplitude calculation is extended to all three benchmark sightlines using the same ionisation correction and uniform-sphere geometry at $R = 30$ kpc. The ratio $\mathcal{A}_{\text{env}} \equiv |\Delta v_{\text{H}}|/|\Delta v_{\text{conf}}|$ is a conditional diagnostic. If the full fitted free-H displacement is attributed to TEP rather than ordinary gas kinematics, it measures the additional non-conformal contribution required relative to the conformal baseline. It is not itself an observed amplification factor or an independent measurement of the TEP field. The ratio varies across sightlines by a factor of ~ 9 :

Sightline	$\log N_{\text{HI}}$	$\log N_{\text{H}}$	$ \Delta v_{\text{conf}} $ (km/s)	$\mathcal{A}_{\text{env}}$ (disformal/conformal)	Fitted Δv_{H}
Q1009+2956	17.36	20.61	0.009	$\sim 9.0 \times 10^3$	−78.9 km/s
PKS 1937−101	17.92	20.89	0.017	$\sim 4.5 \times 10^3$	−76.0 km/s
J1332+0052	19.25	21.55	0.077	$\sim 1.5 \times 10^3$	−118.0 km/s

The $\mathcal{A}_{\text{env}}$ values in this table use the Poisson Green function with the ionisation-corrected neutral-hydrogen column as the source density (the “gas-column basis”). The halo-mass calculation (Gate 4B-0-III below) uses the total host halo mass as the source density and gives $|\Delta v_{\text{conf}}| \approx 0.005$ km/s for all three sightlines (same $10^{10} M_{\odot}$ host), corresponding to $\mathcal{A}_{\text{env}} \approx 1.6\text{--}2.5 \times 10^4$. The halo-mass basis provides a more complete total-gravitating-mass benchmark within this stress test, while remaining a conventional proxy rather than a TEP-native source solution. Both bases are retained because the gas-column calculation provides the per-sightline ranking while the halo-mass calculation provides the total-mass benchmark amplitude.

A universal $\mathcal{A}_{\text{env}}$ — which is what a geometry-independent transport path $\mathcal{C}_{T,\parallel}$ most naturally provides — is excluded at face value: applying Q1009’s $\mathcal{A}_{\text{env}} \approx 9.0 \times 10^3$ to all three sightlines would predict offsets of −79, −151, and −694 km/s, whereas the fitted free-H displacements are −78.9, −76.0, and −118.2 km/s. Under the limiting pure-TEP attribution, TEP must account for the unrestricted-H displacement rather than the isotope constant; the per-sightline $\mathcal{A}_{\text{env}}$ values quantify the corresponding maximal non-conformal requirement. Two interpretations are consistent with this:

1. *Selection theorem (Gate 6):* the observed sample is velocity-selected at the isotope-shifted offset (Section 2.3). The detected subpopulation is precisely the tail where the geometry-dependent TEP offset happens to lie near v_{iso} . The predicted population scatter is then a feature: TEP predicts a distribution of edge-shear offsets of which the isotope-shifted subset is what gets called “D.” The per-sightline $\mathcal{A}_{\text{env}}$ values reflect the different core potentials and halo environments.
2. *Per-sightline $\mathcal{A}_{\text{env}}$:* the ratio is a property of the temporal-disformal transport and varies with the local geometry. This is permitted, but the framework must then predict why $\mathcal{A}_{\text{env}}$ anticorrelates with K across sightlines — currently no mechanism is derived.

The decisive empirical test is the orphan population (Gate 6): the velocity-selection theorem predicts that edge-shear features at offsets other than −81.6 km/s exist and are being absorbed into H I component structure. The theoretical test is the three-sightline predictive closure (Gate 4B-IV / Gate 10): a single frozen microscopic $B(\phi)$, combined with independently specified absorber environments, must predict the temporal contribution without using the observed velocities, with the fitted free-H displacements serving as maximal pure-TEP amplitude targets. The full covariant calculation (Gate 10) has been performed and confirms that no single power-law $B(\phi)$ achieves this closure under the imported Λ CDM rate; the closure condition is then shown to be a property of the temporal-rate function rather than of the disformal channel, and under the TEP-native descent-drift law derived from the corpus potential the inverse-fit B_0 roots concord to a dispersion of 1.00004 at the near-criticality floor. That branch is nevertheless excluded by the ambient-lapse self-consistency requirement: the tracked-criticality profile holds the ambient boundary factor at $u(\bar{\varphi}) \approx 1 - 10^{-5}$ at every epoch, under which the matter-metric redshift map cannot host the observed $z = 2.5\text{--}3.6$ absorber

population without a negative ambient field. The unscreened temporal-disformal channel is thereby bounded sub-critical at these epochs, and an amplitude of the observed scale is admissible only through an environmentally modulated coupling $B_{\text{eff}} = B(\phi) G(X_{\text{local}})$ — a gate closed in the ambient (restoring the conformal map) that opens inside nonlinear wells. Gate 10-b quantifies the condition: the three sightlines require gate contrasts $dG/d\ln\delta\varphi = 4.2\text{--}6.8 \times 10^{-4}$ per e-fold of well depth, agreeing across the sample within a factor 1.6 — consistent with a single saturating gate $G \sim 10^{-3}$ inside wells (which also bounds the induced displacement in deeper structures below cluster-velocity systematics) — so the channel survives only under a beyond- $B(\phi)$ coupling of $D_2(\phi, X)$ class, reported as a measurable functional requirement rather than an achieved closure. Gate 10-c (step_10c, results/gate10c_ambient_lapse_margin.json) closes the loop from the admissibility side: evaluated on the realized absorber-epoch ambients ($\bar{\varphi} = \ln(1+z)$, $\dot{u} = H_T(z) \approx 4\text{--}5 H_0$ under the imported Λ CDM rate, $A \approx 0.22\text{--}0.29$), the ambient lapse condition $u(\bar{\varphi}) = B(\bar{\varphi})\dot{\bar{\varphi}}^2/A^2 < 1$ caps the ungated canonical envelope at $B_0 \lesssim 0.03$ (cap minimum $z \approx 3$), and the negative branch is bounded at $|B_0| \lesssim 10^{-2}$ by the ambient fluid wall $Z_t > 0$ at absorber-epoch matter densities. Under the TEP-native drift ($H_T/H_0^2 \approx 0.22\text{--}0.44$ across the descent and selected-tracking branches) the cap loosens to $B_0 \lesssim 1.0\text{--}4.5$ and the required ambient suppression relaxes to $G_{\text{ambient}} \lesssim 0.01\text{--}0.06$ at $B_0 = 52\text{--}78$: phenomenological amplitudes remain supercritical on the bare ambient by a factor $\sim 30\text{--}50$ under the native rate ($\sim 10^3$ under the imported rate), so ambient closure is lapse-mandated on either drift rather than merely phenomenological. The gate the amplitude analysis selects is therefore also the gate the signature ledger requires: the two requirements converge on the same $D_2(\phi, X)$ -class form.

One-line-of-sight transport calculation (Gate 4B-0-III). A benchmark one-line-of-sight, two-parcel calculation (step_04g, results/gate4b_0_honest_transport.json) confirms the static-sector predictions of the framework. The scalar field is solved using the Jakarta action architecture with the frozen conformal coupling $\beta_A = -1$ and benchmark weak-field choices $V(\phi) = \frac{1}{2}m_\phi^2\phi^2$ and $B(\phi) = B_0\phi^2$, where $\varphi \equiv \phi/M_{\text{Pl}}$, with m_ϕ and B_0 taken from external constraints rather than from the observed velocity. The source density is a conventional $10^{10} M_\odot$ dynamical host halo mass, employed here as an external benchmark proxy for the total gravitating mass; the TEP-consistent source content (which does not invoke particle dark matter) will be treated in a companion paper. For this benchmark proxy at $R = 30$ kpc, the conformal piece gives $|\Delta v_{\text{conf}}| \approx 0.005$ km/s with the correct blueward sign. Paper 29 screening suppresses massive halos ($M > 10^{11} M_\odot$) and leaves intermediate halos ($10^{10}\text{--}10^{11} M_\odot$) weakly screened — exactly the LLS host halo mass range — confirming that the unscreened scalar charge — the Einstein-frame anomalous acceleration — is strongest where TEP requires it. The screening separation is the central TEP prediction. Using the Paper-0 PPN mapping $\gamma_{\text{PPN}} - 1 = -2\beta_A^2 \mathcal{S}_\Sigma^{(\odot)} / (1 + 2\beta_A^2 \mathcal{S}_\Sigma^{(\odot)})$ under the unified unstretched-space metric (or $-4\beta_A^2 \mathcal{S}_\Sigma^{(\odot)} / (1 + 2\beta_A^2 \mathcal{S}_\Sigma^{(\odot)})$ on the isotropic-conformal benchmark) — linear in the screened source charge — the bare unscreened source charge ($\mathcal{S}_\Sigma^{(\odot)} = 1$) gives $\gamma_{\text{PPN}} - 1 = -2/3$ ($-4/3$ on the isotropic benchmark) in the Solar System, violating the Cassini bound $|\gamma - 1| < 2.3 \times 10^{-5}$ by over four orders of magnitude. The solar screening value is an adopted single-source benchmark, not derived from the absorber-regime ansatz: an adopted conservative profile following the Paper 0 §7 source-charge convention takes $\mathcal{S}_\Sigma^{(\odot)} \approx 1.9 \times 10^{-6}$ (the derived kinetic operator returns $\mathcal{S}_\Sigma^{(\odot)} \simeq 1.6 \times 10^{-15}$ at 1 AU — nine orders stronger suppression; step_21), giving $\gamma_{\text{PPN}} - 1 \approx -7.7 \times 10^{-6}$, so $|\gamma_{\text{PPN}} - 1| < 2.3 \times 10^{-5}$ (Cassini PASS). Under the absorber-regime amplitude ansatz, LLS-scale potentials place the absorber in the effectively unscreened limit $S_A^{\text{abs}} \approx 1$ at LLS host potentials $|\Phi_{\text{abs}}|/c^2 \sim 10^{-8}$, so the bare coupling operates in the absorber regime. Note that the Cassini PPN bound constrains the source-charge projection $\mathcal{S}_\Sigma^{(\odot)}$ (which enters γ_{PPN}), while the absorber clock shift is governed by the amplitude projection S_A^{abs} (which enters $\Delta \ln A$); the same potential-depth environmental variable characterizes both projections, but their response functions are distinct and are not assumed numerically identical (Jakarta §2.2). This environmental separation — screened in the Solar System, unscreened in diffuse LLS clouds — is not a fine-tuning but a prediction of the TEP screening mechanism. The Cassini grazing ray at $b = 1.6 R_\odot$ lies inside the solar transition radius $r_{\text{trans}} = 13.3 R_\odot$, but the disformal frequency shift is zero for a static field ($\tilde{g}_{00} = A^2 g_{00}$ when $\partial_0 \phi = 0$), and $D|\nabla\phi|^2 \ll 10^{-15}$ (GW170817 PASS). The static conformal piece provides a signed lower bound; the disformal contribution $\mathcal{A}_{\text{env}} \sim 10^4$ would have to arise from the non-static disformal transport path $\mathcal{C}_{T,\parallel}$, the identified Gate 4B channel. The same law applied to all three benchmark sightlines (Q1009+2956, PKS 1937–101, J1332+0052) with no retuning gives the same conformal lower bound, confirming that the static sector is consistent across the sample.

Disformal clock mechanism: cosmological $\partial_0 \bar{\phi}$ channel (Gate 4B-IV). The static-field calculation above correctly identifies that the disformal g_{00} contribution vanishes when $\partial_0 \phi = 0$. However, the v0.16 cosmological architecture $1 + z = A_0/A_{\text{em}} = \exp(|\beta_A| \phi/M_{\text{Pl}})$ requires a non-zero cosmological time derivative of the scalar field. Defining the TEP temporal-rate function $H_T(z) \equiv -d \ln A / dt$, the scalar background satisfies $\dot{\bar{\phi}} = M_{\text{Pl}} H_T(z)$, so for $A = e^{-\phi/M_{\text{Pl}}}$ (step_04h, results/gate4b_disformal_clock.json). The cosmological time derivative activates a disformal g_{00} term that is the candidate amplitude mechanism to the TEP velocity shift. TEP predicts a geometry-dependent Δv_T , not a universal -81.6 km/s shift; the -81.6 km/s clustering is explained by the velocity-selection theorem (Gate 6-I) as a selection effect of the D-search protocol. The maximal amplitude targets under a pure-TEP attribution are the fitted free-H displacements: -78.9 , -76.0 , -118.0 km/s for Q1009+2956, PKS 1937–101, and J1332+0052 respectively.

$$\partial_0 \bar{\phi} = M_{\text{Pl}} H_T(z) \neq 0. \quad (9)$$

For the reconstructed TEP background used in this calculation, $H_T(z)$ enters through $\partial_0 \bar{\phi} = M_{\text{Pl}} H_T(z)$ and is a property of the temporal-field background, not of the fitted distance–redshift map: the earlier identification $H_T = H_{\Lambda\text{CDM}}$ substitutes the FLRW

lookback–redshift relation for the field dynamics TEP replaces it with, and is retained below only as a foreign-rate baseline against which the TEP-native candidates are evaluated (Gate 10, section 7d). At the absorber redshift $z = 2.5042$, the cosmological field is $\bar{\phi}/M_{\text{Pl}} = \ln(1+z)/|\beta_A| = 1.254$, and under the ΛCDM baseline the cosmological time derivative is $\partial_0\bar{\phi} = M_{\text{Pl}} H_T(z) = 1.34 \times 10^{-5} \text{ eV}^2$. The total scalar field at the absorber is $\phi = \bar{\phi}(t) + \delta\phi(r)$, where the local perturbation $\delta\phi/M_{\text{Pl}} \sim 10^{-9}\text{--}10^{-8}$ is the Poisson-sourced contribution from the host halo (evaluated with the axiom-convention source $\rho_* A_{,\phi}$, i.e. carrying the ambient coupling $A(\bar{\phi}) \simeq 0.2\text{--}0.3$ at these epochs). The cosmological time derivative activates the disformal g_{00} term:

$$\tilde{g}_{00} = A^2 g_{00} + B(\phi) (\partial_0\bar{\phi})^2, \quad (10)$$

producing a differential clock shift between core and edge. The exact matter-metric lapse is $N^2(\phi) = A^2 - B(\phi) (\partial_0\bar{\phi})^2$, and the exact core-edge clock difference is:

$$\Delta_{\text{ec}}^{\text{clock}} = \ln \frac{N_{\text{edge}}}{N_{\text{core}}} = \frac{1}{2} \ln \frac{A^2 - B_{\text{edge}} (\partial_0\bar{\phi})^2}{A^2 - B_{\text{core}} (\partial_0\bar{\phi})^2}. \quad (11)$$

With $B(\phi) = B_0(\phi/M_{\text{Pl}})^n$ and the cross-term expansion $B(\phi_{\text{core}}) - B(\phi_{\text{edge}}) \approx n B_0 (\bar{\phi}/M_{\text{Pl}})^{n-1} \Delta(\delta\phi/M_{\text{Pl}})$, the first-order expansion gives the linearized velocity shift:

$$\frac{|\Delta v|}{c} \simeq \frac{n}{2} B_0 \left(\frac{\bar{\phi}}{M_{\text{Pl}}} \right)^{n-1} \Delta \left(\frac{\delta\phi}{M_{\text{Pl}}} \right) \frac{(M_{\text{Pl}} H_T)^2}{A^2}, \quad (12)$$

valid only when $B(\phi) (\partial_0\bar{\phi})^2/A^2 \ll 1$. The exact matter-metric lapse is $N^2(\phi) = A^2 - B(\phi) (\partial_0\bar{\phi})^2$, and the exact core-edge clock difference is $\Delta_{\text{ec}}^{\text{clock}} = \frac{1}{2} \ln \frac{N_{\text{edge}}^2}{N_{\text{core}}^2}$. The exact solution, obtained by solving $-c \Delta_{\text{ec}}^{\text{clock}} = \Delta v$ for B_0 subject to the Lorentzian condition $N^2 > 0$, gives $B_0 = B_0^{\text{lin}}/(1 + \beta B_0^{\text{lin}})$ where $\beta = (\bar{\phi}/M_{\text{Pl}})^n (M_{\text{Pl}} H_T)^2/A^2$ and B_0^{lin} is the linearized value.

Pure-TEP amplitude stress test. The exact-lapse equation has a divergence: as $B_0 \rightarrow 1/\beta$, the invariant $\mathcal{I} = 1 - B(\bar{\phi}) (\partial_0\bar{\phi})^2/A^2 \rightarrow 0$ and the predicted shift $\Delta v \rightarrow \infty$. The original analysis targeted the full 81.6 km/s isotope shift, but TEP does not predict a universal -81.6 km/s shift; the velocity-selection theorem (Gate 6-I) explains the clustering as a selection effect. Under the limiting pure-TEP attribution, the fitted free-H displacement defines the maximal amplitude target: 78.9 km/s for Q1009+2956. The exact-lapse solution under the first-order core-edge field expansion for this target gives $B_0 \sim 10^8 \text{ eV}^{-4}$ with $\mathcal{I} \sim 10^{-4}$, sitting near the Lorentzian boundary. The near-critical regime is a property of the exact-lapse equation: any target above the linear-regime maximum (~ 0.01 km/s) drives B_0 toward $1/\beta$. Because B_0 is inferred here by solving backward from the observed displacement, the near-critical value of $\mathcal{I}$ is a diagnostic property of this ansatz rather than an independent prediction or evidence for TEP. Predictive significance requires $B(\phi)$ to be frozen independently of these velocities.

Constraint analysis. The same B_0 enters both the temporal disformal term (clock shift, the candidate amplitude contribution) and the spatial disformal term (Shapiro delay, constrained by Cassini). The spatial term at the Sun is $\gamma_{\text{dis}} = B(\phi_{\odot}) (\partial_r\phi_{\odot})^2/(2U_{\odot})$, where $U_{\odot} = |\Phi_{\odot}|/c^2 = 2.12 \times 10^{-6}$. For this reduced Solar-System consistency calculation, the nonlinear environmental response is represented phenomenologically by an effective source-charge projection $\mathcal{S}_{\Sigma}^{(\odot)}$. In the microscopic Jakarta formulation, $\mathcal{S}_{\Sigma}$ is not an independent factor inserted into the fundamental scalar source equation; it is an observable projection of the environment-dependent scalar solution. The reduced solution gives $\phi_{\odot}/M_{\text{Pl}} = 2.26 \times 10^{-10}$, with $\partial_r\phi_{\odot} = 156 \text{ eV}^2$. For the quadratic continuation $n = 2$, the exact-lapse $B_0 = 2.88 \times 10^8 \text{ eV}^{-4}$ gives $\gamma_{\text{dis}} \approx 8.4 \times 10^{-2}$ without environmental screening of B , exceeding the Cassini bound $|\gamma - 1| < 2.3 \times 10^{-5}$ by a factor $\sim 10^3$. However, if $B(\phi)$ is environmentally screened through the same $\mathcal{S}_{\Sigma}$ operator that screens the conformal coupling, the Solar-System projection is suppressed by $\mathcal{S}_{\Sigma}^{(\odot)} \lesssim 1.15 \times 10^{-5}$ under the unified unstretched-space metric (5.8×10^{-6} isotropic), bringing γ_{dis} within the Cassini bound. That screening is no longer a conditional option: under the derived Paper-0 §7 operator the solved solar gradient already carries $\mathcal{S}_{\Sigma}^{(\odot)} \simeq 1.6 \times 10^{-15}$ at 1 AU (7.7×10^{-25} at the solar surface), so the disformal spatial term at $B_0 = 2.88 \times 10^8 \text{ eV}^{-4}$ evaluates to $\gamma_{\text{dis}} \simeq 4 \times 10^{-42}$ — the Cassini bound is non-binding for every tested power-law normalisation (step 21). The earlier $n = 2$ ‘Cassini exclusion’ was an artifact of the adopted conservative profile, not a derived bound; the binding constraints are the ambient-lapse self-consistency and cross-sightline transfer evaluated below. The quadratic $B \propto \phi^2$ realization at the normalization required to reproduce the Q1009 amplitude therefore remains viable, consistent with the TEP framework. A steeper field dependence can satisfy the local Cassini, GW170817 and clock constraints without screening. For $n = 5$, the Q1009-normalised solution gives $\gamma_{\text{dis}} \approx 4.9 \times 10^{-31}$, $\delta c/c \approx 8.6 \times 10^{-20}$, and a negligible local clock shift. However, as shown by the cross-sightline transfer test below, the same $n = 5$ power law loses Lorentzian signature

at the two higher-redshift absorbers. Thus $n = 2$ is excluded by local constraints without environmental screening of B , while $n = 5$ is excluded by cross-redshift transfer; neither tested power law provides a viable universal microscopic closure without additional structure. The v0.16 theory states that the microscopic form of $B(\phi)$ is not uniquely frozen; the field-space envelope $B(\phi) = B_0 \varphi^2 / (1 + \varphi^2) \exp(-\varphi^4 / 2\sigma_B^4)$ with $\varphi \equiv \phi / M_{\text{Pl}}$ is a named benchmark, not a unique prediction. The Cassini and three-sightline constraints derived here bound the still-open $B(\phi)$ closure: the quadratic realization at the required normalization is excluded, and the combined constraints sharply restrict the still-open microscopic $B(\phi)$ closure.

For this Gate-IV consistency calculation an illustrative Solar-System source-charge profile is adopted with $\mathcal{S}_{\Sigma}^{(\odot)} = 1.9 \times 10^{-6}$, consistent with the Paper 0 §7 source-charge convention and below the Cassini ceiling 1.15×10^{-5} (5.8×10^{-6} on the isotropic-conformal benchmark). This value is not inferred from the absorber $\mathcal{S}_A^{\text{abs}}$ ansatz; it is also conservative relative to the derived operator, which returns $\mathcal{S}_{\Sigma}^{(\odot)} \simeq 1.6 \times 10^{-15}$ at 1 AU (step_21, cross-checked against the Paper-0 step_54 tabulation).

Three-sightline test: diagnostic, transfer, and predictive closure. The test is performed in three stages. *Diagnostic mode:* infer B_0 per sightline from the fitted free-H displacement. Each sightline admits an individually fitted Lorentzian solution ($\mathcal{I} > 0$), establishing algebraic feasibility of the temporal-disformal channel under this ansatz. The inferred B_0 values are 1.46×10^8 , 1.52×10^7 , and $2.00 \times 10^7 \text{ eV}^{-4}$ for Q1009+2956, PKS 1937–101, and J1332+0052 respectively. *Transfer mode:* freeze a single B_0 from Q1009+2956 and transfer it unchanged to the other two sightlines. The invariant $\mathcal{I} = 1 - \beta(z) B_0$ depends only on the cosmological background ($\beta(z) = (\bar{\phi} / M_{\text{Pl}})^n (M_{\text{Pl}} H_T)^2 / A^2$), not on the halo environment. With B_0 from Q1009+2956 ($z = 2.5042$), the invariant becomes $\mathcal{I} < 0$ at PKS 1937–101 ($z = 3.572$) and J1332+0052 ($z = 3.421$), crossing the Lorentzian boundary. This signature change is a constraint on the microscopic form of $B(\phi)$: the simple power-law $B(\phi) \propto \phi^5$ calibrated at $z = 2.5$ does not transfer to $z > 3.4$. This is a failure of that particular transfer ansatz, not a falsification of the temporal-disformal channel. *Predictive closure* (Gate 10, step_10_covariant_transport, results/gate10_covariant_transport.json): the full covariant calculation has been performed with a TEP-native background $\bar{\phi}(z)$, $\partial_0 \bar{\phi} = M_{\text{Pl}} H_T(z)$, a three-dimensional screened Poisson solution $\delta\phi(r)$ for each absorber using the per-sightline column densities ($\log N_{\text{HI}} = 17.36, 17.92, 19.25$ for Q1009+2956, PKS 1937–101, J1332+0052 respectively), the complete disformal matter metric $\tilde{g}_{\mu\nu} = A^2 g_{\mu\nu} + B(\phi) \nabla_\mu \phi \nabla_\nu \phi$, Hamiltonian null-geodesic propagation, and the matter-frame frequency $\omega = -u^\mu k_\mu = -k_0 / N$ (cone-tilt terms cancelling exactly for a static observer). The three-dimensional screened field gives $\Delta(\delta\phi / M_{\text{Pl}}) = 1.26 \times 10^{-8}, 2.41 \times 10^{-8}, 1.02 \times 10^{-7}$ for the three sightlines respectively, versus the uniform 1.60×10^{-8} of the reduced two-point approximation; the screening reduction factor at Q1009+2956 is 0.79. The covariant evaluation uses the signed scalar source $\beta_A = -1$, which places the absorber on the physical branch $\delta\phi_{\text{core}} > \delta\phi_{\text{edge}} > 0$ required by the sign convention $\phi > 0$ in mass concentrations: the core perturbations are $+5.18 \times 10^{-9}, +7.56 \times 10^{-9}$, and $+3.36 \times 10^{-8}$ in units of M_{Pl} under the axiom source $\rho_* A_{,\phi}$ (the ambient factor $A(\bar{\phi}) \simeq 0.22\text{--}0.29$ suppresses the per-unit-density coupling relative to the earlier bare- ρ convention). With the exact quasi-static form $\Delta v_T = c \ln(N_{\text{core}} / N_{\text{edge}})$, the physical branch gives $N_{\text{core}} < N_{\text{edge}}$ and hence $\Delta v_T < 0$ â€ a blueward edge feature relative to the deep-core anchor, matching the sign of the fitted displacements $-78.9, -76.0, -118.0 \text{ km/s}$. The full three-dimensional field gives core-edge contrasts $\Delta(\delta\phi) = 3.61 \times 10^{-9}, 5.27 \times 10^{-9}$, and 2.30×10^{-8} — factors of 0.23, 0.33, and 1.44 relative to the reduced two-point value (1.60×10^{-8} for all three). The predictive closure with a single frozen $B(\phi) = B_0 (\phi / M_{\text{Pl}})^5$ gives $\chi^2 = 2.99$ at $B_0 = 1.50 \times 10^7 \text{ eV}^{-4}$ (the simultaneous Lorentzian bound, set by PKS 1937–101 at $B_0 < 1.52 \times 10^7 \text{ eV}^{-4}$), with predicted signed shifts of $-0.001, -0.41$, and -0.06 km/s versus the fitted $-78.9, -76.0, -118.0 \text{ km/s}$. Per-sightline inverse fits exist on the corrected branch â€ each sightline admits a B_0 at ~ 0.999 of its own Lorentzian boundary ($1.46 \times 10^8, 1.52 \times 10^7, 2.00 \times 10^7 \text{ eV}^{-4}$) â€ but the three values differ by a factor ~ 10 , so no single B_0 closes the sample. The near-criticality and cross-redshift signature loss persist under the full covariant treatment. The root cause is that $A^2 / \dot{\bar{\phi}}^2$ varies by a factor 3.68 across the three sightlines; no monotonically increasing $B(\phi)$ can compensate this variation while keeping all three absorbers simultaneously near-critical. The matter-frame maximum is $|\Delta v| \sim 0.005 \text{ km/s}$, and the negative- B branch is lapse-safe on this timelike-gradient channel with no amplification boundary, bounding $|\Delta v|$ at the conformal scale $\sim 0.05 \text{ km/s}$; it remains subject to the ambient fluid-hyperbolicity wall $|B_0| \lesssim 10^{-2}$ at these epochs (Gate 10-c). The Paper-0/28 disformal envelope $B(\phi) = B_0 \varphi^2 / (1 + \varphi^2) \exp(-\varphi^4 / 2\sigma_B^4)$ was then scanned over $\sigma_B \in [1.0, 3.0]$ in the same covariant pipeline. The scan isolates the sign condition cleanly: the blueward displacement requires the boundary factor $u(\phi) = B(\phi) \dot{\bar{\phi}}^2 / A^2(\phi)$ to rise into the well, i.e. a local logarithmic slope $n_{\text{eff}} \equiv d \ln B / d \ln \varphi > -2\bar{\varphi}$ (≈ -2.5 to -3.0 for these absorbers), while simultaneous closure across the sample requires $B(\bar{\varphi}) \propto A^2 / \dot{\bar{\phi}}^2$, i.e. $n_{\text{eff}} = -2 - d \ln H_T^2 / d \bar{\varphi} \approx -4.9$ under the imported $\dot{\bar{\phi}} = M_{\text{Pl}} H_{\text{LCDM}}$ rate. For $\sigma_B \lesssim 1.25$ the envelope declines too steeply ($n_{\text{eff}} < -2\bar{\varphi}$) and the correct-sign branch closes for at most one sightline (at $\sigma_B = 1.25$ only Q1009+2956 admits an inverse fit); for $\sigma_B \gtrsim 1.5$ the correct-sign branch exists for all three sightlines and each admits a signed inverse-fit B_0 (e.g. at $\sigma_B = 2.0$: $8.0 \times 10^8, 2.1 \times 10^8, 2.5 \times 10^8 \text{ eV}^{-4}$ for Q1009+2956, PKS 1937–101, J1332+0052), but the three values still differ by a factor $\sim 3\text{--}4$. The two requirements are disjoint under the imported rate, so the envelope improves on the pure power law — the correct sign is accessible and the inferred B_0 dispersion shrinks — without yet achieving three-sightline closure. The failure is thereby converted into a quantitative compatibility condition on the TEP temporal-rate function: with $s(z) \equiv d \ln H_T^2 / d \ln(1 + z) = d \ln \dot{\bar{\phi}}^2 / d \bar{\varphi}$, simultaneous closure requires a local coupling slope $d \ln B / d \bar{\varphi} = -2 - s$ while the blueward sign requires $d \ln B / d \bar{\varphi} > -2$, so for a coupling of monotone slope the two requirements are compatible if and only if $s < 0$ — a descent-type ambient drift, non-increasing with field depth. This is not an ad hoc freedom: $H_T \equiv -d \ln A / dt$ is fixed by the cosmological temporal-field solution.

TEP-native rate evaluation (Gate 10, section 7d). The disjointness diagnosed above is a property of the imported rate, not of the disformal channel: H_{LCDM} is the metric-expansion rate of the cosmology TEP rejects, whereas $\partial_0 \bar{\phi}$ is the drift of the temporal field

on the static TEP background. Two corpus-derived candidates were therefore evaluated in the same covariant pipeline (step_10_covariant_transport, results/gate10_covariant_transport.json). The homogeneous closed-benchmark drift $H_t^2/H_0^2 = 1 + 3\Omega_0 z/(1+z)$ (Paper 0, step_12; $\Omega_0 \simeq 0.2998$) saturates at high redshift and carries $s \approx +0.12\text{--}0.16$ over the absorber range: under this rate the envelope inverse-fit dispersion collapses from a factor $\sim 3\text{--}4$ to 1.24 at $\sigma_B = 1.31$, but the marginal sign gap remains ($s > 0$). The unsourced ambient sector on the reconstructed potential branch $V_{\text{rec}}(u) = V_0 - (\bar{\rho}/2) e^{-u}$ obeys free-roll energy conservation on the static manifold, $\dot{\phi}^2 = \bar{\Pi}_\infty^2 + \bar{\rho} e^{-\bar{\varphi}}$, i.e. $H_T^2/H_0^2 = (1 + \kappa/(1+z))/(1 + \kappa)$ with $\kappa = \bar{\rho}/\bar{\Pi}_\infty^2$; this descent law carries $s = -\kappa/(1+z+\kappa) < 0$ and satisfies the compatibility criterion for every $\kappa > 0$. The required coupling profile then reduces to $B(\bar{\varphi}) \propto e^{-(2+s)\bar{\varphi}}$, which in the descent-dominated limit $\kappa \rightarrow \infty$ is $B(\bar{\varphi}) \propto A(\bar{\varphi})$: the disformal coupling tracks the conformal factor across the absorber field range.

Evaluated at the pure-descent limit, the descent-tracking realisation $B(\varphi) = B_0 \varphi^4/(\varphi_c^4 + \varphi^4) e^{-\lambda\varphi}$ ($\lambda = 1$, $\varphi_c = 0.1$) yields inverse-fit normalisations $B_0 = 7.5585, 7.5583, \text{ and } 7.5582 \times 10^{10} \text{ eV}^{-4}$ for Q1009+2956, PKS 1937–101, and J1332+0052 respectively — a dispersion of 1.00004 , at the near-criticality floor set by the differing observed amplitudes. Each inverse root reproduces its own sightline by construction. Forward evaluation at their geometric mean, the single common B_0 inside that cluster ($7.5583 \times 10^{10} \text{ eV}^{-4}$), predicts $-18.41, -102.94, \text{ and } -366.55 \text{ km/s}$ against the observed $-78.9, -76.0, \text{ and } -118.0 \text{ km/s}$ (fractional $\chi^2 = 5.15$). The sign is blueward on every sightline. The amplitudes are not a common- B_0 prediction: the near-critical map is steep enough that a factor- 1.00004 agreement among the inverse roots does not produce a common velocity. An earlier evaluation stepped below the cluster and reported $-12.7, -28.3, \text{ and } -118.0 \text{ km/s}$; that point is not the common root. The same coupling satisfies the local bounds with orders of magnitude to spare: $\gamma_{\text{dis}} \approx 1.1 \times 10^{-14}$ at the Solar-System field depth (Cassini bound $|\gamma - 1| < 2.3 \times 10^{-5}$) and a vanishing $\delta c/c$ at the ambient field (GW170817 bound 3×10^{-15}). The evaluation also closes a reporting gap in the envelope scan: the inverse-fit envelope solutions carry $\gamma_{\text{dis}} \sim 5 \times 10^{-2}\text{--}17$ across the rate models, exceeding the Cassini bound by $\sim 10^3\text{--}10^6$ because the φ^2 low-field suppression is insufficient at the required normalisations; Cassini-compatible closure requires suppression at least as steep as φ^3 , which the quartic-suppressed descent-tracking form satisfies by construction. The concordance is nevertheless tautological on this branch, and the branch fails ambient-lapse self-consistency. On the descent-tracking form $u(\varphi) = B(\varphi)\dot{\phi}^2/A^2(\varphi) \propto e^{\delta\varphi}$ is constant along the entire descent: evaluated at the joint B_0 , the ambient boundary factor is $u(\bar{\varphi}) = 0.99997\text{--}0.99999$ at all three sightline epochs (Gate 10, section 7d; results/gate10_covariant_transport.json, ambient_self_consistency). The factor- 1.00004 agreement of the inverse roots is therefore a structural property of the flat- u branch — every sightline's inverse fit returns the same ambient boundary factor by construction — not an empirical coincidence. The ambient sector on this branch sits permanently at near-criticality: the matter-metric lapse is $\tilde{N} = A\sqrt{1-u} \approx 3\text{--}5 \times 10^{-3} A$, so the placement map $1+z = \tilde{N}(0)/\tilde{N}(\bar{\varphi})$ used to locate the absorbers is no longer the conformal map $\bar{\varphi} = \ln(1+z)$. An absorber at the conformal placement would present at $1+z \approx 646\text{--}1651$, and reproducing the observed $z = 2.5\text{--}3.6$ requires $\bar{\varphi} \approx -4.0$ to -4.4 — a negative ambient field in conflict with the corpus sign convention $\bar{\varphi} \geq 0$. The tracked-criticality branch is accordingly excluded by the ambient-lapse self-consistency requirement, and the argument generalises: because u varies with field depth only through $e^{\delta\varphi}$ ($\delta\varphi \lesssim 10^{-7}$), any realisation that is near-critical inside a well is near-critical in the ambient field on the same branch. At sub-critical ambient u the disformal correction to $\tilde{N}$ across the absorber cannot reach the observed $\sim 80 \text{ km/s}$ scale. The unscreened temporal-disformal amplitude channel is thereby excluded at these epochs; a contribution of the observed scale is possible only where $B(\phi)$ is itself environmentally screened, so that u is sub-critical in the ambient while rising inside the well. The residual open element is thereby sharpened rather than removed: the absorber data bound the ambient drift realisation to sub-critical u and select environmental screening of the disformal coupling, not the descent-tracking form, as the live candidate channel (Paper 0 §8).

Limit of the reduced Gate-4B closure. The failed power-law transfer constrains the specific reduced construction tested here. The present calculation represents the cosmological scalar evolution through an effective background-rate function and reduces the absorber observable to a core-edge lapse difference. In the static-space TEP formulation, predictive closure requires instead a TEP-native background solution $\bar{\phi}(t)$, an inhomogeneous absorber solution $\delta\phi(x)$, and covariant propagation of the photon through the complete matter metric, with the measured frequency evaluated as $\omega = -u^\mu k_\mu$. This calculation has now been performed (Gate 10, step_10_covariant_transport). The near-criticality and cross-redshift signature loss persist under the full covariant treatment: the $A^2/\dot{\phi}^2$ variation across redshift under the imported expansion rate is not removed by replacing the two-point lapse difference with the full three-dimensional screened field and covariant frequency projection. The constraint is therefore not an artefact of the reduced core-edge approximation; the Gate-10 rate evaluation identifies it instead as a property of the imported Λ CDM drift, under which no single power-law $B(\phi)$ closes the sample. Under the TEP-native descent-drift law the same machinery admits a Cassini/GW170817-compliant descent-tracking coupling whose inverse roots agree to a factor 1.00004 ; the common- B_0 forward velocities do not reproduce the three amplitudes (section 7d of the Gate-10 evaluation, reported above); and the ambient self-consistency check excludes the branch outright — the flat- u realisation places the ambient boundary factor at $\approx 1 - 10^{-5}$ at every epoch, under which the matter-metric redshift map cannot produce the observed absorber population without a sign-violating negative ambient field. Environmental modulation of the disformal coupling, $B_{\text{eff}} = B(\phi) G(X_{\text{local}})$, is therefore not an additional option but the only surviving channel for a temporal-disformal amplitude of the observed scale, with the gate's required form measured by Gate 10-b.

Environmental discrimination. Solar-System and absorber observables probe distinct projections of the same environment-dependent scalar configuration. Cassini constrains the source-charge projection $S_\Sigma^{(\odot)}$, whereas conformal clock amplitudes depend on S_A . Their numerical relationship emerges from the common nonlinear scalar solution (Paper 0 §7): the single kinetic operator $S_\Sigma = 1/P_{,X}$ returns $\simeq 1.6 \times 10^{-15}$ at the solar field; at the LLS ambient the value is branch-dependent — $\sim 1.5\text{--}21$ on the static

acceleration proxy, $\sim 0.13\text{--}0.18$ on the solved drift-dominated field at $z \sim 2.5\text{--}3.6$ — a discrimination factor of order 10^{16} (drift branch) to 10^{17} against the solar value, produced by the same stiffness function, not assumed (step_21, `results/step_21_canonical_projection.json`). The reduced Gate-4B calculation shows that a temporal-disformal contribution can be large in the absorber parametrisation while remaining constrained locally. The full covariant transport (Gate 10) has been completed and confirms the reduced calculation's finding that no single power-law $B(\phi)$ achieves predictive closure under the imported expansion rate; under the TEP-native descent-drift law the descent-tracking coupling $B(\varphi) \propto \varphi^4/(\varphi_c^4 + \varphi^4) e^{-\lambda\varphi}$ concords the inverse roots at the criticality floor while satisfying Cassini and GW170817, but is excluded by ambient-lapse self-consistency (section 7d), leaving an environment-modulated disformal coupling — active inside wells, closed in the ambient, with the required gate slope measured in Gate 10-b — as the only admissible amplitude channel.

Gate 4B status. The conformal field fixes the sign but gives only $0.005\text{--}0.01 \text{ km s}^{-1}$. Static perturbative disformal terms are also insufficient. The cosmological temporal-disformal term $B(\phi)\dot{\phi}^2$ can reach much larger amplitudes in the Einstein-frame parametrisation explored here, although the solutions are near-critical. Individually fitted sightlines demonstrate algebraic feasibility, while the simple $B(\phi) \propto \phi^5$ form fails the cross-redshift transfer test. The full covariant calculation (Gate 10) with $\omega = -u^\mu k_\mu$ confirms this finding under the imported Λ CDM drift: the $A^2/\dot{\phi}^2$ variation across redshift prevents any single increasing $B(\phi)$ from achieving simultaneous near-criticality at all three sightlines. The corpus disformal envelope $B(\phi) = B_0 \varphi^2/(1 + \varphi^2) \exp(-\varphi^4/2\sigma_B^4)$ (Paper 0 §2; employed as the strong-field realization in Paper 28), whose essential-singularity damping is designed for the near-critical regime, has now been evaluated in the same covariant pipeline over $\sigma_B \in [1.0, 3.0]$ (Gate 10, `results/gate10_covariant_transport.json`): under the imported rate the correct blueward sign is accessible for $\sigma_B \gtrsim 1.5$ and per-sightline inverse-fit solutions exist, but the inferred B_0 values retain a factor $\sim 3\text{--}4$ dispersion. The solar comparison distinguishes assigned-profile diagnostics from a nonlinear source solution. The adopted-gradient and acceleration-proxy columns in Gate 10 use independently assigned solar field values; their numerical PASS/FAIL flags do not establish a complete Cassini prediction. The selected Paper-0 sector already supplies a potential-plus-source solar profile (steps 53/57), with a surface gradient response of order 10^{-8} relative to the bare source. The disformal calculation must evaluate B on the total field and use the gradient from that same solution, with the environmental boundary and propagation observable specified. The quartic small-field taper supplies strong suppression in the consistent-profile benchmark; a universal non-binding bound for every scanned coupling is not inferred from the proxy columns. The residual requirement reduces to a quantitative bound on the cosmological rate function: joint sign-and-closure compatibility requires $s \equiv d \ln H_T^2 / d \ln(1+z) < 0$ over $z \simeq 2.5\text{--}3.6$, which the imported Λ CDM rate ($\approx +2.9$) fails, the corpus's homogeneous closed-benchmark drift ($\approx +0.12\text{--}0.16$) approaches but marginally misses, and the unsourced ambient descent law on the reconstructed potential — free-roll energy conservation $\dot{\phi}^2 = \bar{\Pi}_\infty^2 + \bar{\rho} e^{-\bar{\varphi}}$ on the static manifold — satisfies ($s = -\kappa/(1+z+\kappa) < 0$). Under the pure-descent limit the descent-tracking coupling $B(\varphi) = B_0 \varphi^4/(\varphi_c^4 + \varphi^4) e^{-\varphi}$ concords the inverse-fit B_0 values to a dispersion of 1.00004, the criticality floor set by the observed amplitude scatter, with quartic suppression at small total field. A zero-field propagation diagnostic does not replace evaluation at the nonzero field and drift of a realized environment. The ambient self-consistency check then excludes the branch: $u(\bar{\varphi}) \approx 1 - 10^{-5}$ uniformly along the descent places the ambient sector at near-criticality, under which the matter-metric lapse map $1+z = \tilde{N}(0)/\tilde{N}(\bar{\varphi})$ requires $\bar{\varphi} \approx -4.0$ to -4.4 to reproduce the observed absorber redshifts — a sign-violating field. The unscreened temporal-disformal channel is thereby bounded sub-critical at these epochs and cannot reach the observed amplitudes; the absorber data select environmental screening of $B(\phi)$, rather than the descent-tracking form, as the admissible realisation.

Functional constraints on S_A^{abs} and $S_\Sigma^{(\odot)}$. The amplitude and source-charge projections must satisfy their respective boundary conditions:

1. **Absorber regime (unscreened):** $|\Phi_{\text{abs}}|/c^2 \sim 10^{-8} \Rightarrow S_A^{\text{abs}} \approx 1$, so that the bare conformal response $\alpha_{A,\text{eff}} = S_A^{\text{abs}} \beta_A = -1.0$ applies and the conformal/static branch sign theorem remains valid.
2. **Solar-system regime (Cassini):** $S_\Sigma^{(\odot)} \lesssim 1.15 \times 10^{-5}$ under the unified unstretched-space metric ($\gamma_{\text{PPN}} - 1 = -2\beta_A^2 S_\Sigma^{(\odot)} / (1 + 2\beta_A^2 S_\Sigma^{(\odot)})$; or $S_\Sigma^{(\odot)} \lesssim 5.8 \times 10^{-6}$ on the isotropic-conformal benchmark $\gamma_{\text{PPN}} - 1 = -4\beta_A^2 S_\Sigma^{(\odot)} / (1 + 2\beta_A^2 S_\Sigma^{(\odot)})$) is required by the Cassini PPN bound $|\gamma_{\text{PPN}} - 1| < 2.3 \times 10^{-5}$; the derived operator returns $S_\Sigma^{(\odot)} \simeq 1.6 \times 10^{-15}$ at 1 AU, a margin of $\sim 3.5 \times 10^9$ (step_21).
3. **Disformal contribution target:** Under a pure-TEP attribution, the temporal-disformal channel would need to supply a geometry-dependent contribution $\mathcal{A}_{\text{env}} \approx 9.0 \times 10^3$ for Q1009+2956 on the gas-column basis ($\sim 10^4$ on the halo-mass basis; per-sightline: 4.5×10^3 , 1.5×10^3 on the gas-column basis) relative to the conformal-only Poisson calculation, with the fitted free-H displacements as the maximal amplitude targets under a pure-TEP attribution.

The absorber amplitude ansatz $S_A^{\text{abs}}(|\Phi|/c^2) = 1/(1 + (|\Phi|/c^2/\Phi_{\text{half}})^n)$ with $\Phi_{\text{half}} = 1.59 \times 10^{-7}$ and $n = 3.56$ is the *absorber-regime amplitude transfer* employed in this paper. The screening variable for this ansatz is the gravitational potential depth $|\Phi|/c^2$, because the scalar field responds to total mass through the Poisson equation and screening is set by the integrated potential.

Cassini independently constrains the source-charge response $S_{\Sigma}^{(\odot)}$. The relation between S_A^{abs} and $S_{\Sigma}^{(\odot)}$ is not established by this ansatz; it is supplied by the derived kinetic operator of Paper 0 §7, evaluated at the absorber environment in step_21. That evaluation is branch-dependent: on the static acceleration proxy it returns $S_{\Sigma} \sim 1.5\text{--}21$, while on the solved drift-dominated field at the absorber epoch (where the spatial/temporal gradient ratio is $\sim 10^{-4}$) it returns $S_{\Sigma} \sim 0.13\text{--}0.18$, i.e. the conformal-gradient response is screened by a factor ~ 6 . The two projections remain distinct: S_{Σ} acts on the source-charge/gradient response, while the ansatz's unscreened assignment concerns the amplitude channel S_A^{abs} , which the corpus's amplitude/shear split preserves independently. The product $\alpha_{\text{eff}}^2 \times |\Phi|/c^2$ peaks in the unscreened intermediate regime ($|\Phi|/c^2 \sim 10^{-7}$), confirming that the Einstein-frame anomalous acceleration is strongest in diffuse clouds and suppressed in dense environments. The required disformal contribution cannot arise from the conformal sector alone (which gives $S_A^{\text{abs}} \approx 1$ in the unscreened regime); it must originate from the disformal sector $B(\phi)\nabla_{\mu}\phi\nabla_{\nu}\phi$ and its environmental projection. The conformal/static contribution maps through $\Delta v_{\text{conf}} \simeq -c \Delta \ln A$. The temporal-disformal contribution instead modifies the matter-frame lapse $N^2 = A^2 - B(\phi)\dot{\phi}^2$; its observable spectral shift must be derived from the covariant frequency $\omega = -u^{\mu}k_{\mu}$. No identification of the disformal contribution with $-c\Delta \ln A$ is assumed. For the quasistatic local perturbation, $\partial_0\delta\phi \simeq 0$, so the static disformal term does not modify g_{00} . The cosmological background has $\partial_0\phi \neq 0$, activating the distinct temporal-disformal clock channel tested in Gate 4B-IV. The perturbative weak-field disformal term is suppressed by $(|\Phi|/c^2)^2$ for a constant coupling and by $(\nabla\psi/\psi)^2 \ll 1/r^2$ for an inverse-field coupling (the potential is nearly flat in the interior of an extended mass distribution); no local perturbative $B(\phi)$ can provide the required disformal contribution (quantitatively confirmed: the disformal metric correction $B_0 \cdot (|\Phi|/c^2)^2/r^2 \sim 10^{-62}$ in geometrized units for $B_0 = 1$). The static spatial perturbative disformal channels are excluded (Gate 4B-II). The temporal disformal clock channel $B(\phi)(\partial_0\phi)^2$ activated by the cosmological scalar time derivative is the candidate amplitude mechanism (Gate 4B-IV). The full covariant core-edge observable $\omega = -u^{\mu}k_{\mu}$ has been evaluated (Gate 10, step_10_covariant_transport), confirming the reduced calculation's finding.

Note on coupling conventions. Following Paper 0, the frozen microscopic coupling is $\beta_A = -1.0$, i.e. $\alpha_0 = \sqrt{2}\beta_A = -\sqrt{2}$ in DEF normalization [3], [54]. The environmentally screened effective scalar charge for source-charge observables (the sector that fifth-force tests constrain) is $\alpha_{\text{eff}} = S_{\Sigma}(\mathcal{E})\alpha_0$ [1]; the environmentally screened effective amplitude for clock/Cepheid observables is $\alpha_{A,\text{eff}} = S_A(\mathcal{E})\alpha_0$. These are distinct projections. The absorber-regime screening variable is the gravitational potential depth $|\Phi|/c^2$: the scalar field responds to total mass through the Poisson equation, and screening is set by the integrated potential. Lyman-limit absorbers ($|\Phi_{\text{abs}}|/c^2 \sim 10^{-8}$) sit in the unscreened regime where $S_A^{\text{abs}} \approx 1$ and the bare coupling applies. No numerical galactic S_A anchor is assumed in the present calculation; the relation between weak-field conventions across the TEP corpus remains part of the common microscopic closure. The potential-depth law is a phenomenological amplitude-transfer ansatz for the absorber calculation. Cassini independently constrains the source-charge response $S_{\Sigma}^{(\odot)} \lesssim 1.15 \times 10^{-5}$ (5.8×10^{-6} on the isotropic benchmark). The relation between S_A and S_{Σ} is supplied by the derived kinetic operator of Paper 0 §7 — evaluated at the absorber ambient it is branch-dependent: $S_{\Sigma} \sim 1.5\text{--}21$ on the static-branch acceleration proxy and $S_{\Sigma} \sim 0.13\text{--}0.18$ on the solved drift-dominated field at $z \sim 2.5\text{--}3.6$ (step_21, drift-branch evaluation) — rather than by this two-parameter fit. The TEP coupling is screened where the ambient landscape is steep, satisfying solar-system bounds (Cassini); in the unscreened absorber regime, the bare coupling produces an Einstein-frame anomalous acceleration comparable to gravity, whose effect on absorber structure and kinematics is a prediction of the framework. Constraints on unscreened scalars in other theoretical frameworks [56] assume different screening mechanisms and coupling structures, and are not directly applicable to the TEP coupling, which is screened by the environmental operator rather than by a chameleon or symmetron mechanism.

Cassini constraint derivation. The Cassini experiment measured the PPN parameter $\gamma_{\text{PPN}} = 1 + (2.1 \pm 2.3) \times 10^{-5}$ [59]. Following Paper 0's frozen convention, $\gamma_{\text{PPN}} - 1 = -2\beta_A^2 S_{\Sigma}/(1 + 2\beta_A^2 S_{\Sigma})$ under the unified unstretched-space metric (or $-4\beta_A^2 S_{\Sigma}/(1 + 2\beta_A^2 S_{\Sigma})$ on the isotropic-conformal benchmark; equivalently $-2\alpha_0\alpha_{\text{eff}}/(1 + \alpha_0\alpha_{\text{eff}})$ with $\alpha_{\text{eff}} = S_{\Sigma}\alpha_0$ and $\alpha_0 = \sqrt{2}\beta_A$ in DEF normalization). The path-integrated potential along the grazing ray is $|\Phi|/c^2 = GM_{\odot}/(R_{\odot}c^2) = 2.12 \times 10^{-6}$ at the solar radius. For $S_{\Sigma} \ll 1$, the Cassini bound $|\gamma_{\text{PPN}} - 1| < 2.3 \times 10^{-5}$ requires $S_{\Sigma}^{(\odot)} \lesssim 1.15 \times 10^{-5}$ (5.8×10^{-6} on the isotropic benchmark) in the Solar-System environment, a suppression supplied by the kinetic-sector stiffness mechanism — the slope of the clock-rate landscape pinned where the ambient shear is already steep, rather than the field amplitude flattened (Paper 0 §2.2, v0.16). This is an independent constraint on the source-charge response $S_{\Sigma}^{(\odot)}$; it is not used to calibrate the absorber amplitude ansatz S_A^{abs} . The relation between S_A^{abs} and $S_{\Sigma}^{(\odot)}$ is supplied by the common kinetic operator evaluated in step_21, which reproduces the Cassini suppression at solar gradients ($S_{\Sigma} \sim 10^{-15}$) while returning the absorber on the unscreened-to-anti-screening branch ($S_{\Sigma} \sim 1.5\text{--}21$).

Corpus parallel: the core-vs-disk gradient. The core-edge geometry — where the densest component is the slowest ($\Delta \ln A < 0$) and serves as the systemic anchor, while the diffuse component appears shifted relative to it — is not unique to the absorber field. A parallel geometry has been noted in the Cepheid distance-ladder sector [51], where the galactic bulge (deep potential, slowest clocks) defines the systemic spectroscopic redshift while Cepheids in the diffuse outer disk are corrected to the rest frame using that bulge-derived redshift. The absorber-field calculation presented here stands

independently; the Cepheid parallel is noted as a consistency check on the shared physical principle $\Delta \ln A < 0$, not as proof of the absorber mechanism.

3.4 Scaling Consistency: Absorber Clouds vs. Galaxies

A natural concern is whether a scalar mechanism capable of accounting for a ~ 80 km/s shift across a diffuse gas cloud would produce an unphysically large effect across a massive galaxy. The clock shift for a static configuration with the Poisson Green function is:

$$\alpha_{A,\text{eff}} \equiv S_A^{\text{abs}} \alpha_0, \quad \Delta \ln A = -\frac{\alpha_{A,\text{eff}}^2 K}{M_{\text{Pl}}^2} = -\frac{2 \alpha_{A,\text{eff}}^2 |\Phi|}{c^2}, \quad (13)$$

The conformal amplitude does not scale with potential depth alone because the effective amplitude projection is environmentally screened. Diffuse absorbers remain effectively unscreened ($S_A^{\text{abs}} \approx 1$ at $|\Phi_{\text{abs}}|/c^2 \sim 10^{-8}$). Solar-System tests of the source-charge sector (conventionally reported as fifth-force constraints) probe the distinct source-charge projection S_Σ ($S_\Sigma^{(\odot)} \lesssim 1.15 \times 10^{-5}$ unified, 5.8×10^{-6} isotropic). No numerical galactic S_A anchor is assumed here. The common microscopic scalar solution derives both projections (Paper 0 §7; evaluated for the BBN environments in step_21); that cross-domain closure is now in place and is not load-bearing for the spectroscopic result tested here.

4. Cosmological Transport and Thermodynamics

Part II: Global Cosmology. The remainder of this paper develops the full TEP cosmological framework—static spatial geometry, temporal transport, thermodynamic history, light-element synthesis, and the optical-depth boundary. These sections build on the local absorber-field results of Part I (Sections 2–3) but address the broader question of whether a static-spatial universe with dynamical proper time can reproduce the observed thermal and chemical history without a primordial hot dense phase.

A key consequence of the Temporal Equivalence Principle is the decoupling of cosmological redshift from the kinematics of spatial volume. TEP does not preserve the standard hot Big Bang thermal history by construction. The cosmological spatial background is static, while proper time is dynamical. Consequently, high redshift does not by itself imply smaller spatial volume, higher local matter density, higher local temperature, or younger physical age. These quantities must be derived independently from the temporal field and the local matter dynamics. The standard thermal history, including recombination [23], [24], [25] and the CMB [19], [20], is not assumed but must be independently reconstructed within the TEP framework.

4.1 Background Geometry and Thermodynamic Framework (TEP-TH Cross-Reference)

The geometric and thermodynamic foundations of the TEP cosmological framework — the temporal horizon, curvature regularity ($0 < p \leq 1/2$, all polynomial curvature invariants vanish), proper-time asymptotic regularity ($\Delta\tau \rightarrow \infty$), and CMB blackbody spectral preservation — are established in TEP-TH (Paper 27, [4]). The acoustic-sector closure, including $\Omega_b h^2$ determination from Planck peaks and sound-horizon preservation at < 6 ppm, is established in TEP-HC (Paper 18, [3]) and TEP-C0 (Paper 26, [2]), using the line-of-sight integration formalism of Seljak & Zaldarriaga [26] as implemented in CLASS [28] and hi_class [29], with perturbation theory following the standard formalism [14], [15]. The present paper does not re-derive these results; it applies the established framework to the BBN-specific questions of nuclear reaction flows, light-element equilibrium, and D/H spectroscopy.

The essential structural results inherited from TEP-TH are: (i) the effective scale factor $a_{\text{eff}} = A_{\text{clock}}$ with $A_{\text{clock}}(z) = (1+z)^{-1}$; (ii) the redshift decomposition $\ln(1+z_T) = \int_\gamma (\Sigma_{||} + \mathcal{C}_{T,||}) d\ell$; and (iii) the temporal horizon limit $A_{\text{clock}} \rightarrow 0$, $z \rightarrow \infty$, $\tau \rightarrow \infty$, $\mathcal{K} \rightarrow 0$. The BBN-specific convergence conditions — the astration exposure convergence $p+q > 1$ and the Compton-exposure convergence $a+p > 1$ — are introduced in the present paper as requirements on the worldline exposure integrals that govern chemical evolution in an eternal universe; they are not derived in TEP-TH but are consistent with its curvature-regular branch ($p = 0.5$, $q = 0.8$, giving $p+q = 1.3$).

4.2 Nuclear Reaction Flows and the Opacity Boundary

Nuclear production in the TEP framework occurs along local matter histories parameterised by proper time τ and spatial position x :

$$\frac{dY_i(x)}{d\tau} = \sum_r N_{ir} \lambda_r [T_{\text{loc}}(\tau, x), n_{\text{loc}}(\tau, x)] \prod_j Y_j^{\nu_{jr}}, \quad (14)$$

where Y_i is the abundance of species i , N_{ir} is the stoichiometric coefficient, λ_r is the temperature- and density-dependent reaction rate, and the sum runs over all nuclear reactions r . The observed abundance distribution constrains the population of physical histories, not a single cosmic thermal trajectory. The temporal horizon ensures that low-exposure worldlines (those with bounded $\mathcal{E}_{\text{astr}}$) remain observationally accessible at high redshift, while the convergence condition $p + q = 1.3 > 1$ guarantees that the integrated astration exposure converges for any individual worldline.

Opacity boundary and the $\tau(z)$ three-point requirement.

In the TEP static frame, the comoving path length $\ell(z)$ to the temporal horizon diverges ($\eta \rightarrow \infty$ as $A_{\text{clock}} \rightarrow 0$), and static space removes FLRW a^{-3} dilution of the electron density. The opacity theorem is *conditional*: it requires an asymptotic electron-density bound $n_e(\ell) \gtrsim \ell^{-s}$ with $s \leq 1$, ensuring $\int^\infty n_e d\ell = \infty$. Under this condition, the Thomson optical depth diverges at high redshift, providing an observable last-scattering boundary without a physical plasma wall. The $\tau(z)$ three-point requirement — reproducing reionization $\tau(z < 7) \approx 0.054$, a last-scattering surface at $\tau \sim 1$, and the trough onset — is evaluated in species-resolved form in Step 14 (results/step_14_opacity_three_point.json), which supersedes the earlier fiducial statement of a smooth density rise ($\sim 14\times$ by $z \sim 7$, $\sim 500\times$ by the LSS). Under the curvature-regular power-law branch $p = 0.5$ the requirement is forced, not tuned: the wing datum alone fixes the low-redshift electron density at $n_e \approx 0.51 n_{b,0}$; because the column from trough onset to the boundary spans $\sim 1.7 \times 10^6$ Hubble radii, the mean electron density there is bounded above by $3.7 \times 10^{-4} n_{b,0}$, so the trough is derived to be neutral — the Gunn–Peterson saturation it must produce ($\tau_{\text{LyC}} \sim 2 \times 10^{10}$ in the admissible profile) is this requirement; and the remaining $\Delta\tau \sim 1$ at the boundary is supplied by the ionization front there — the Strömberg edge where the ionizing field maintained by the deep thermalized zone is extinguished by the neutral column — parametrized in Step 14 as a concentrated electron column equivalent to $n_e \sim (4-8) \times 10^2 n_{b,0}$ over ~ 1 Hubble radius, an ionization level the Gate-12 VMO budget already sustains; the front is coincident with the Gate-11B thermalization edge, so the $\tau = 1$ surface lands at $z \approx 1080$; its depth is derived rather than imposed, since the ambient bath $\tilde{T}_{\text{loc}}(z) = T_0(1+z)$ crosses the density-dependent Saha ionization temperature at $z_{\text{front}} = T_{\text{ion}}/T_0 - 1 \simeq 9 \times 10^2 - 1.1 \times 10^3$ across the diffuse-to-absorber front-density range (Step 19, results/step_19_eternal_surface_reconciliation.json). The physical mechanism and its radiative/perturbation closure are established in Step 20 (results/step_20_ambient_bath_front.json): the ambient bath itself carries a Lyman-continuum Wien tail, whose unshielded photoionization supply $\Gamma_{\text{bath}} = c\sigma_{\text{LyC}}n_\gamma(>\nu_L; \tilde{T}_{\text{loc}})$ crosses the local recombination demand $\alpha_B n_b$ at $z_{\text{front}} \simeq 9.1 \times 10^2 - 1.05 \times 10^3$ across the diffuse-to-absorber density range, with an unshielded Saha-type transition width $\Delta z \sim 10^2$ and unshielded Thomson lower bound $z_{\tau=1} \simeq 752$ at $500 n_{b,0}$ (Parts A–C). Solving the coupled self-shielded LyC radiative transfer across the neutral trough (Part D) sharpens the ionization front to $\Delta z_{\text{sh}} \simeq 0.9$ and extinguishes the leading-edge ionization tail, moving $z_{\tau=1}$ directly onto the front ($z_{\tau=1}^{\text{sh}} = 1074.3$ at $500 n_{b,0}$). Inverting the front overdensity required for $z_{\tau=1}^{\text{sh}} = 1080.0$ (Part E) — where forward-evaluation across the independent Step-14 $[400, 800] n_{b,0}$ band (and the adopted LLS density conversion $n_{\text{H}} = 10^{-4} \text{ cm}^{-3} (N_{\text{HI}}/10^{17.2} \text{ cm}^{-2})^{2/3}$ at $\log_{10} N_{\text{HI}} = 17.345 \pm 0.150 \implies 661.2 n_{b,0}$, predicting $z_{\tau=1}^{\text{pred}} = 1079.8^{+4.6}_{-4.5}$ from column uncertainty alone, excluding conversion and selection uncertainties) already predicts $z_{\tau=1}^{\text{pred}} = 1077.9$ at the $600 n_{b,0}$ midpoint (span $[1069.9, 1083.6]$, $\pm 0.6\%$ sensitivity set by the Planck Wien steepness $\chi = h\nu_L/k_B T_{\text{loc}} \simeq 53.5$) — fixes $\Delta_{\text{front}} = 668.9 n_{b,0}$ ($n_b = 1.68 \times 10^{-4} \text{ cm}^{-3}$, $z_{\text{front}} = 1080.2$), sitting squarely inside the Step-14 $[400, 800] n_{b,0}$ front-column band and the KODIAQ/SQUAD Lyman-limit absorber density range. Deriving the pre-boundary inertial driving density from the zero-free-parameter two-branch temporal-shear law at the sound horizon — where $g_t = \Lambda_X^2/(\sqrt{2}M_{\text{Pl}}) = cH_0/2$ is a fixed constant of the scalar action S_ϕ and $q_s = g_b(r_s)/g_t = \Omega_b(H_0 r_s/c)$ is conformally invariant under $A(\phi)$, giving $\omega_{m,\text{eff}} = \omega_b(1 + y_s/\sqrt{2}) = 0.1391$ ($\omega_\phi = 0.1166$, $q_s = 1.62 \times 10^{-3}$, $y_s = 7.38$, $M_s = 6.22$) in the single-pass $r_s = 146.19$ Mpc analytic integral ($\theta_s^{\text{derived}} = 0.010425$ rad, $\ell_A = 301.3$) and converging under self-consistent $r_s(\omega_\phi)$ fixed-point iteration in `hi_class` to $r_s = 144.94-145.00$ Mpc, $q_s = 1.59-1.60 \times 10^{-3}$, $y_s = 7.42-7.44$, and $\omega_\phi = 0.1179-0.1182$ (spanning $\omega_\phi \in [0.1174, 0.1199]$ across $H_0 \in [67.4, 70.0] \text{ km s}^{-1} \text{ Mpc}^{-1}$) — and integrating the coupled linear Einstein–Boltzmann baryon–photon acoustic ODE with spherical-Bessel projection on the self-shielded visibility kernel at this inverted density (Part F) shifts the matter-frame photon-diffusion blur from the uninverted $500 n_{b,0}$ value $\ell_D = 1242$ (Step 19) to $\ell_D = 1521$ ($\theta_D = 6.58 \times 10^{-4}$ rad versus the observed Silk angle 6.50×10^{-4} rad, within 1.2%, right inside the Planck 1400–1600 damping band) and recovers the first three acoustic peak multipoles at $\ell = 221.0, 527.1, 811.8$ ($\leq 1.94\%$ of Planck in the reduced 6-variable tight-coupling ODE; evaluating the exact `hi_class` Einstein–Boltzmann hierarchy at the same zero-parameter two-branch clock-slope response $\omega_\phi = \omega_b y_s/\sqrt{2} = 0.1179-0.1182$ — matching the unconstrained TEP-HC full-Planck TTTEEE MCMC posterior 0.11813 ± 0.00117 to 0.02σ (0.02%) at the posterior mean ($\leq 1.6\sigma$ across $H_0 \in [67.4, 70.0]$) and improving the full Planck TTTEEE + low- ℓ + lensing + BAO + Pantheon+ likelihood by $\Delta\chi^2 = -0.06$ ($\Delta\text{AIC} = -2.06$) relative to free ω_{cdm} — yields $100\theta_s = 1.04111$ at $+0.02\sigma$ from Planck's 1.04110 ± 0.00031 at $H_0 = 67.951 \text{ km s}^{-1} \text{ Mpc}^{-1}$ [$100\theta_s^{\text{CLASS}} = 1.04201$, $+0.10\sigma$ at $H_0 = 68.18 \text{ km s}^{-1} \text{ Mpc}^{-1}$], acoustic peaks at $\ell = (220.49, 536.93, 814.06)$ within $\leq 0.40\%$ of Planck [$\ell = (220.3, 536.5, 813.4)$ within $\leq 0.31\%$ at $H_0 = 68.18$], BAO drag scale $r_d = 147.49-147.55$ Mpc, and linear-scale $\sigma_8 = 0.8170-0.8186$ without particulate dark matter). The driving component in that system is not an inserted dust fluid: the proper-time field ϕ fills all space as a continuous ambient scalar medium ($A(\phi) > 0$ everywhere, analogous to an ambient temperature field), and (δ_p, v_p) is its two-branch clock-slope perturbation mode derived from the quadratic expansion of the scalar action $S_\phi = \int d^4x \sqrt{-g} M_{\text{Pl}}^2 \Lambda_X^4 P(\xi)$ on the U_μ -foliation. Around the rolling

cosmic background $\dot{\phi} = M_{\text{Pl}} H_0$ ($\bar{\xi} = 1$), expanding $P(\xi)$ to second order in δu supplies the explicit canonical timelike kinetic coefficient $\mathcal{K}_t = \bar{P}_{,X}(1) + 2\bar{P}_{,XX}(1) = 2k + 6 = 38.060$ (canonical momentum $\pi_u = \partial \mathcal{L}^{(2)}/\partial \dot{u} = \mathcal{K}_t \dot{u}$), the reduced acoustic model prescribes $\mathcal{K}_{\text{sp}}^{\text{eff}} = \sqrt{2}/y_s = 0.1917$ and $c_{p,\text{eff}}^2 = \mathcal{K}_{\text{sp}}^{\text{eff}}/\mathcal{K}_t = 5.04 \times 10^{-3}$ with zero Thomson coupling. The native flux Hessian instead gives $K_{\perp} = 1/y_s$ and $K_{\parallel} = 2/y_s + (q_s y_s)^2$. A shared quadratic-action derivation of the effective prescription remains separate from the verified flux identity (Paper 0, Step 81). The reduced equations $\delta'_p = -k^2 v_p + 3\Phi'$, $v'_p = -\mathcal{H} v_p + c_{p,\text{eff}}^2 \delta_p + \Psi$ are the conditional model integrated here with effective clock-slope density $\rho_{p,\text{eff}}(a) = (y_s/\sqrt{2}) \rho_b(a) \propto a^{-3}$, recovering $\ell = (221.2, 527.3, 811.9)$ at $\leq 1.90\%$ of Planck in the reduced ODE ($\leq 0.40\%$ in exact `hi_class`). A closure ablation recorded in the same artifact quantifies the sensitivity: replacing the position-space background shell scale $|\nabla u_0|$ with perturbation wavenumber k in Fourier space ($q(k) = q_s/(kr_s)$, $y_k \simeq 6.3$ – 25.2 across $k \in [0.005, 0.080] \text{ Mpc}^{-1}$) gives $\ell = (235.6, 578.0, 818.4)$ at $\leq 7.54\%$, and the instantaneous-Poisson slaved null in which δ_p is forced to track δ_b without autonomous inertia gives $\ell = (231.9, 542.0, 841.0)$ at $\leq 5.41\%$ — both degraded relative to the action-linearized sound-shell closure, demonstrating that the observed peak positions require the autonomous ambient time-field inertia represented by the specified effective model, not merely an amplified static Poisson source. Crucially, unlike Hořava–Lifshitz integration-constant dust^[80] or mimetic dark matter^[81] — where a free integration-constant dust fluid clusters independently in halos — the TEP ambient clock-slope mode (δ_p, v_p) is strictly baryon-sourced through $\nabla \cdot [P_{,X}(|\nabla u|)\nabla u] = \beta_A \rho_b$: on linear cosmological scales ($k \lesssim 0.2 h \text{ Mpc}^{-1}$, $R \gtrsim 10 \text{ Mpc}$, governing BAO and σ_8) the scale-averaged clock-slope multiplier $M_s = 1 + y_s/\sqrt{2} \simeq 6.25$ enters the linear Poisson equation identically to an effective density $\omega_{m,\text{eff}} = \omega_b + \omega_{\phi}$, whereas in the nonlinear late-time static virialized galaxy and cluster limit ($R \lesssim 100 \text{ kpc}$, $\dot{u} \rightarrow 0$, $\dot{\Phi} \rightarrow 0$, $g_N \gtrsim g_t$) the timelike kinetic momentum $\pi_u = \mathcal{K}_t \dot{u}$ vanishes identically and the same nonlinear equation reduces without free parameters to the static baryon-only two-branch screening law $(k/\sqrt{2})qy^2 + q^2y^3 = 1$ (Paper 6 SPARC RAR and Paper 19 lensing), introducing zero independent collisionless dark-matter halo cusp. The two-regime behaviour of the ambient clock-slope mode — CDM-like on linear scales ($k \lesssim 0.2 h \text{ Mpc}^{-1}$) and screened to the baryon-only two-branch law inside virialized halos — is phenomenologically analogous to dark-matter superfluidity (Bereziani & Khoury 2015^[82]), with the difference that in TEP the transition is governed by the kinetic stiffness $P_{,X}$ of the proper-time field rather than by a condensate phase transition. The front is therefore a self-consistently closed phase boundary of the ambient radiation field, not an inserted column. An explicit analytic profile satisfying all three points simultaneously is exhibited in Step 14: wing $\tau(z < 6) = 0.054$, thin trough column, unit surface at the edge, saturated Lyman-continuum trough. The earlier “14×” figure is recovered as the wing density of the exponential temporal law, under which the Gate-11B edge recedes beyond any observable depth ($\ln(1 + z_{\text{therm}}) \sim 10^4$ – 10^6 for dust absorptivities, and $z_{\text{therm}} \sim 10^{197}$ even for Thomson opacity alone); the observed boundary depth $z_* \sim 10^3$ therefore selects the power-law branch, and the trough-collapse morphology is branch-forced. What remains open is the full absorber-statistics inversion: a pilot remap of the 26-system KODIAQ Lyman-limit subsample at $z = 2.38$ – 3.22 gives an observed incidence $dN/dz \approx 31$, corresponding to $dN/d(\ell/L_0) \approx 3.9$ under the power-law path Jacobian, and applying the same procedure to a statistically complete LLS/Lyman-forest catalogue is the identified pipeline extension that would convert this admissible profile into a demonstrated reconstruction. This reconstruction from absorber statistics remains the independent derivation that TEP requires and no FLRW paper needs to perform. Step 16 (results/step_16_boundary_transport.json) audits the transport consistency of this profile, and Step 18 (results/step_18_boundary_selfconsistency.json) closes each item it flagged at toy-model level. The visibility function on the imposed profile peaks sharply at the front ($z_{\text{pk}} = 1080$, FWHM 10^{-3} in z); under the Step-16 closure convention (continued Thomson scattering beyond the edge) the normalized distribution carries a broad tail with $\sim 55\%$ of last-scattering probability in the front column and rms width ~ 250 in z . Step 18C identifies that tail as a closure artifact: if the Gate-11B edge is genuinely absorptive — the condition $\kappa_{\text{abs}} \gg 11 L_0^{-1}$, that the boundary destroy photons faster than they scatter free, which is what a thermalization edge means — the post-edge weight falls below 1% and the front carries $\sim 87\%$, with a residual thin foreground screen of damping $e^{-2\tau_{\text{fg}}} \approx 0.85$ versus the standard reionization 0.90. The front gas temperature, previously a free assumption in the Saha check, is now derived: the gas at the edge sits in contact with the reprocessed bath at $T_{\text{bath}} = T_0(1 + z_*) \approx 2946 \text{ K}$ and Compton equilibrium of a Planckian field is stationary at $T_{\text{gas}} = T_{\text{bath}}$, with per-electron coupling time $\sim 1 \text{ yr}$ at the front and $\sim 2 \times 10^6 \text{ yr}$ even at the trough residual — so the Saha ionization $x_e \simeq 0.83$ – 0.99 is a derived result conditional only on the Gate-11B bath that defines the edge. The front's recombination demand $4.5 \times 10^7 \text{ cm}^{-2} \text{ s}^{-1}$ is met by in-situ emissivity for cycling rates $\xi \gtrsim 10^{-5} \text{ Gyr}^{-1}$ before escape and dust losses. The required Γ contrast of $\sim 4 \times 10^{10}$ between wing and trough — the sharpest remaining condition — is supplied by self-shielding in Step 18A: a coupled attenuation/photoionization-balance march shows the wing's ionizing field extinguished over $\sim 24 \text{ LyC e-folds}$, $\sim 6 \text{ Mpc}$ at diffuse density ($\sim 10^{-9}$ of the trough depth, thinner at absorber densities), so the sharp boundary in the admissible profile is a consequence of self-shielding rather than an imposed discontinuity. Step 19 (results/step_19_eternal_surface_reconciliation.json) upgrades the earlier geometric angular estimate to the explicit diffusion integral on the admissible profile: the random-walk blur accumulated across the scattering zones is carried entirely by the front column ($N_{\text{sc}} \simeq 0.92$ there versus ~ 0.05 in the wing and ~ 0.03 in the trough), giving $\sigma_D \simeq \lambda_{\text{mfp}}$ of the front column $\approx 3.7 \times 10^3 \text{ Mpc}$ in the matter frame; under the eternal angular convention $D_A = \chi/(1 + z)$ — matter rulers at depth z are $A = (1 + z)^{-1}$ shorter, so a structure's coordinate extent is $(1 + z)$ its physical size — the blur subtends $\theta_D \simeq 8.1 \times 10^{-4} \text{ rad}$, within $\sim 25\%$ of the observed Silk angle $\sim 6.5 \times 10^{-4} \text{ rad}$ (diffusion-convention $\ell_D \sim 1240$ versus the observed ~ 1400 – 1600 , the residual offset absorbable by κ_{abs} , the front column density and the approach profile). The same step performs the kernel-level reconciliation with the mapped-FLRW computation of TEP-TH (Paper 27): the eternal and Saha–Peebles kernels agree on the integrated column — $\tau = 1$ at $z \simeq 1079$ – 1080 (peak agreement 0.13%) and $\tau_{\text{reion}} \simeq 0.054$ —

and differ morphologically, the eternal surface being a sharp spatial ionization front — a phase boundary in the ambient radiation field, not a physical shell — parametrized as a concentrated column ($\text{FWHM}_z \sim 3 \times 10^{-5}$, carrying $\sim 88\%$ of the normalized column) rather than a temporal freeze-out sweep ($\text{FWHM}_z \sim 10^2$). The column bookkeeping is quantified by the path-measure Jacobian $J(z) = (d\ell/dz)/(c/HL_0) \sim 4 \times 10^7$ at the edge: pointwise equivalence of the two descriptions would put the eternal surface at $\sim 30 n_{b,0}$ — absorber/diffuse densities, not cosmic-recombination densities — confirming that the TH/HC computation is the conformal-frame matter-unit representation of the same observable column whose static-frame realization is closed in Step 20 ($\Delta_{\text{front}} = 668.9 n_{b,0}$, $\ell_D = 1521$, acoustic peaks $\ell = 221.0, 527.1, 811.8$). Because the radiation energy density at the edge exceeds the gas thermal energy density by $U_{\text{rad}}/U_{\text{gas}} = 7.4 \times 10^{15}$ ($U_{\text{gas}}/U_{\text{rad}} = 1.34 \times 10^{-16}$; Step 18B) and the Compton-exposure convergence condition ($a + p > 1$) holds across the foreground column, thermalization-edge Compton and chemical-potential distortions are bounded to $|\Delta y|, |\Delta \mu| \sim U_{\text{gas}}/U_{\text{rad}} \simeq 1.3 \times 10^{-16} \ll 9 \times 10^{-5}$, preserving the Planckian spectrum verified against COBE/FIRAS in TEP-TH (Paper 27, Section 6).

4.3 Separation of Spatial and Temporal Shear

The two distinct phenomenological manifestations of the scalar field ϕ must be distinguished. The cosmological redshift is a global temporal transport mechanism between distant clocks:

$$1 + z = \frac{A_{\text{obs}}}{A_{\text{em}}}, \quad (15)$$

while the apparent deuterium feature is interpreted in the TEP hypothesis as arising from localized absorber shear. For the local absorber, the conformal/static contribution is $\Delta v_{\text{conf}} \simeq -c \Delta \ln A$. When the temporal-disformal channel is active, however, the matter-frame lapse becomes $N^2 = A^2 - B(\phi) \dot{\phi}^2$, and the complete observable must be obtained from the covariant measured frequency $\omega = -u^\mu k_\mu$. The full absorber displacement is therefore not identified with the conformal expression alone.

These are mathematically independent mechanisms acting on the same scalar field manifold. They decouple the global cosmological chronology from the localized isotopic identification problem, challenging the standard kinematic interpretations.

4.4 Primordial Helium Synthesis via Baryonic Cycling

If the spectroscopic free-H alternative is borne out and D/H ceases to provide a unique primordial abundance anchor, the remaining light-element constraints—particularly helium-4—must be reproduced within any non-hot-BBN framework. The helium-4 mass fraction ($Y_p \approx 0.245$ [21], [22], [68]) is the principal such constraint. Without a finite, hot, universally dense origin, the TEP framework must quantitatively demonstrate that this abundance is produced by stellar nucleosynthesis over an unbounded temporal horizon.

Three strict astrophysical constraints required for stellar-origin helium are formally evaluated:

1. **Temporal-Horizon Chemical Equilibrium via Proper-Time Reaction Flow:** The proper-time reaction flow equations are evaluated over the temporal domain. Because the temporal horizon acts as an asymptotic observational transport filter—scaling the observable contribution of the infinite past toward zero ($A(\phi) \rightarrow 0$)—the local chemical evolution is asymptotically decoupled from the absolute history of the universe. Evaluating the proper-time reaction flow shows that the adopted proper-time reaction-flow model exhibits convergence toward a common asymptotic fixed point over the tested initial conditions at the edge of the accessible horizon. Whether the evaluation starts with $Y_0 = 0.00$ or an extremely dense $Y_0 = 0.80$, the reaction flow relaxes to the stable fixed point $Y_{\text{eq}} = 0.247$ of the linear astration equation at the present day ($\tau = 0$). It is important to recognize that this reaction flow represents a *local galactic patch* experiencing continuous star formation. The pristine global background observed at high redshift is not protected from local chemical accumulation by transport delay alone; rather, pristine absorbers correspond to gas worldlines whose accumulated processing exposure $\mathcal{E}_{\text{astr}}$ remains bounded by the convergence condition $p + q > 1$ (Section 4.1). The temporal horizon ensures that such low-exposure worldlines are observationally accessible, but the chemical pristine state is a property of the worldline's local processing history, not of the transport filter.
2. **Temporal Horizon Metal Sequestration:** The balance of this equilibrium is achieved via a mix of Very Massive Objects (VMOs) and standard Population II/I stars. Standard stars yield typical return fractions. However, VMOs—which dominate the low-metallicity branch—undergo extreme radiatively-driven winds that successfully eject their helium envelopes ($E_Y > 0$). Upon core collapse, rather than forming a spatial singularity, the core generates a TEP temporal well where the local clock rate drops to an extreme, finite minimum $N_{\text{min}} \ll 1$ relative to the external interstellar medium. This mechanism directly addresses the historical Carr-Bond-Arnett overproduction limit [57], [58].

standard stellar synthesis of $\sim 25\%$ helium severely overproduces carbon and oxygen, making a purely stellar helium origin incompatible with observed metal abundances. The TEP temporal-horizon sequestration offers a mechanism to avoid this by trapping the heavy metals ($E_Z \approx 0$) while allowing helium to escape, producing the correct Y_p without the metal overproduction that historically ruled out stellar helium synthesis [60].

3. **Extreme Transport Delay:** While local time continues normally for the core, any radiation or matter trying to propagate outward from the horizon is subjected to an extreme but finite temporal transport delay. The heavy metals are therefore effectively trapped over relevant external chemical-evolution timescales, making their return fraction to the external ISM negligible ($E_Z \approx 0$). In the present branch, $E_Z \simeq 0$ is the defining metal-sequestration assumption; deriving the corresponding escape fraction from a full temporal-horizon remnant transport solution remains an independent closure requirement.

These mechanics eliminate the need for a spatial singularity, replacing it with a field-theoretic mechanism for chemical evolution. Under the TEP baryonic-cycling and temporal-horizon exposure conditions, the helium-4 mass fraction $Y_{eq} = 0.247$ emerges as the equilibrium of the baryonic-cycling reaction flow under the adopted stellar yields and temporal-horizon metal sequestration. The numerical reaction-flow calculation demonstrates convergence to this equilibrium for the adopted stellar histories, showing that the observed $Y_p \approx 0.245$ [21], [22] is compatible with stellar nucleosynthesis over an unbounded temporal horizon. The slight offset between the observed target ($Y_p \approx 0.245$, PDG/Cybur standard) and the derived equilibrium ($Y_{eq} = 0.247$) falls well within the systematic uncertainty bounds of the observationally constrained yield parameters p_Y and R : a ± 0.002 shift in Y_{eq} is produced by variations in the helium yield per stellar generation and the return fraction that are smaller than the scatter in empirical H II region yield calibrations. The specific equilibrium value is set by the adopted yield parameters (p_Y , R) and the VMO fraction; the temporal-horizon metal sequestration ($E_Z \rightarrow 0$) ensures that the equilibrium is helium-dominated rather than metal-enriched. The yield parameters are not free tunables directed at the 0.247 target: the helium yield per stellar generation p_Y is constrained by observations of H II regions and stellar population synthesis models ($p_Y^{VMO} = 0.149$, $p_Y^{std} = 0.02$, $p_Y^{std} = 0.01$), the return fraction $R = 0.40$ is determined by initial mass function (IMF) integration over independently measured stellar lifetimes and remnant masses, and with these stellar-physics inputs held fixed, the single low-metallicity VMO fraction $f_{VMO} = 0.994$ is jointly overconstrained across four independent observables (three overconstrained degrees of freedom): (i) $Y_{eq} = 0.2470$ (observed $Y_p = 0.245 \pm 0.003$), (ii) the metal-poor DLA metallicity floor $Z_{eq} = p_Y^{std}(1 - f_{VMO})/(1 - R) = 1.0 \times 10^{-4}$ ($\simeq 10^{-2} Z_\odot$), (iii) the Spite plateau ${}^7\text{Li}/\text{H}_{eq} = 1.55 \times 10^{-10}$ (observed $1.58 \pm 0.10 \times 10^{-10}$), and (iv) ${}^3\text{He}/\text{H}_{eq} = 9.0 \times 10^{-6}$ (observed $1.0 \pm 0.1 \times 10^{-5}$), alongside the cosmic star-formation history and $[\alpha/\text{Fe}]$ ratios in old stellar populations. A complete ab-initio star-formation derivation of $f_{VMO} \simeq 0.994$ from primordial cooling microphysics complements this multi-observable overconstraint. The compatibility of a 99.4% VMO-dominated low-metallicity IMF with the ionization state of the diffuse medium and the CMB optical depth is computed below (Gate 12): in the eternal steady state the constraint is a maintained ionization balance rather than a reionization epoch, and the VMO branch sustains the ionized column to the opacity-boundary depth across the plausible escape/cycling plane, adding no countervailing bound on f_{VMO} . The contraction analysis establishing that $Y_{eq} = 0.247$ is a stable fixed point independent of initial conditions is reported below in this section, and the sensitivity analysis confirms robustness to physically motivated variations in (p_Y , R , f_{VMO}) within their observationally constrained ranges: ± 0.005 in R produces ± 0.002 in Y_{eq} , and ± 0.01 in f_{VMO} produces ± 0.002 in Y_{eq} .

Extension to 3He and 7Li. The baryonic-cycling framework extends naturally to a multi-species light-element network (3He, 4He, 7Li) under the same temporal-horizon sequestration mechanism. VMO cores trap not only metals but also 3He and 7Li produced in stellar interiors, so only 4He (and residual H) escapes via winds. The equilibrium abundances are ${}^3\text{He}/\text{H} = 9.0 \times 10^{-6}$ (observed $\sim 1.0 \times 10^{-5}$, ratio 0.90) and ${}^7\text{Li}/\text{H} = 1.55 \times 10^{-10}$ (observed 1.6×10^{-10} , ratio 0.97). The TEP equilibrium ${}^7\text{Li}/\text{H}$ is a factor of ~ 3.2 below the standard hot-BBN prediction (5.0×10^{-10}), offering a natural pathway to resolve the cosmological lithium problem: the temporal-horizon sequestration of VMO-core 7Li produces less primordial 7Li than hot BBN without requiring additional stellar destruction mechanisms. The convergence condition $p + q = 1.3 > 1$ (Section 4.1) guarantees that the integrated astration exposure is finite. The numerical reaction-flow calculation demonstrates convergence to the fixed point for the adopted stellar histories; the specific ${}^7\text{Li}/\text{H}$ value is then determined by the multi-species reaction-flow model (VMO-core lithium yield, sequestration partitioning, and astration exposure integral), not by the convergence condition alone. The per-species inputs and their partitioning are explicit (Gate 8, results/gate8_light_elements.json): the ${}^7\text{Li}$ effective yield is $p_7^{\text{eff}} = (1 - f_{VMO})p_7^{\text{std}} + p_7^{\text{CR}} = 0.006 \times 2.0 \times 10^{-8} + 5.0 \times 10^{-11} = 1.7 \times 10^{-10}$ per star-forming generation, where $p_7^{\text{std}} = 2.0 \times 10^{-8}$ is the standard-star AGB/nova yield, the VMO fraction $f_{VMO} = 0.994$ carries zero effective ${}^7\text{Li}$ yield because VMO-core production is sequestered ($p_7^{VMO} = 0$), and $p_7^{\text{CR}} = 5.0 \times 10^{-11}$ is the non-stellar cosmic-ray spallation contribution. The equilibrium is $x_7^{\text{eq}} = p_7^{\text{eff}}/(1 - R + d_7) = 1.7 \times 10^{-10}/1.1 = 1.55 \times 10^{-10}$, with net lockup $1 - R = 0.6$ ($R = 0.4$ returned mass fraction) and astration destruction $d_7 = 0.5$. The same partitioning gives $x_3^{\text{eq}} = (1 - f_{VMO})p_3^{\text{std}}/(1 - R + d_3) = 9.0 \times 10^{-6}$ from $p_3^{\text{std}} = 1.5 \times 10^{-3}$, $d_3 = 0.4$, and $Y_{eq} = [f_{VMO}p_Y^{VMO} + (1 - f_{VMO})p_Y^{\text{std}}]/(1 - R) = 0.247$ from $p_Y^{VMO} = 0.149$, $p_Y^{\text{std}} = 0.02$. The ${}^7\text{Li}$ budget is therefore dominated by the 0.6% standard-star channel (1.2×10^{-10}) augmented by cosmic-ray spallation (0.5×10^{-10}), and the quoted equilibrium follows from the published per-species yields rather than a fitted normalization. These yields are adopted stellar-physics inputs; their first-principles derivation remains open.

Stellar-surface depletion and the Spite plateau. The quoted observational value ${}^7\text{Li}/\text{H} \approx 1.6 \times 10^{-10}$ is the Spite plateau — a stellar-surface abundance measured in metal-poor halo stars — rather than a direct measurement of the primordial abundance. A substantial literature argues that the plateau is depleted below the primordial value by atomic diffusion and turbulent mixing, by factors of roughly 2–5 [74], [75]; under that interpretation the TEP equilibrium ${}^7\text{Li}/\text{H} = 1.55 \times 10^{-10}$ would fall below the depletion-corrected primordial abundance. The depletion-minimal position — that the Spite plateau approximates the primordial value — is adopted here [76], and under it the TEP equilibrium is consistent with the observed plateau. Whether consistency survives the depletion correction depends on the magnitude of that correction, which remains debated in the literature.

The $\Delta Y/\Delta Z$ slope and H II region concordance. The standard observational method for determining primordial helium extrapolates the linear relationship between helium abundance Y and metallicity Z (typically traced by O/H) across metal-poor dwarf galaxy H II regions to the $Z \rightarrow 0$ limit [61], [62], [63]. If $Y_{\text{eq}} = 0.247$ is a local chemical equilibrium rather than a universal primordial baseline, the TEP framework must reproduce this observed $\Delta Y/\Delta Z$ linear slope. The baryonic-cycling reaction flow naturally produces a linear Y - Z relation at low metallicity: as astration exposure $\mathcal{E}_{\text{astr}}$ increases, both Y (from helium-producing stellar winds) and Z (from metal-producing standard stars) increase proportionally, yielding a near-linear slope $\Delta Y/\Delta Z \approx p_Y/(1 - E_Z) \cdot R_{\text{metal}}$ that is set by the same yield parameters governing the equilibrium. The temporal-horizon metal sequestration ($E_Z \approx 0$) suppresses the Z -intercept while preserving the slope, so the extrapolation to $Z \rightarrow 0$ recovers Y_{eq} rather than a distinct "primordial" value. The observed consistency of $\Delta Y/\Delta Z$ across hundreds of independent low-metallicity H II regions [62] is therefore a natural prediction of the TEP equilibrium model, not an independent verification of a single primordial event.

Metallicity-dependent VMO transition. As a single-phase global equilibrium, the model silently fails: at $f_{\text{VMO}} = 0.994$, $E_Z \approx 0$ gives $Z_{\text{eq}} = p_Z(1 - f_{\text{VMO}})/(1 - R) \approx 1.4 \times 10^{-4}$ — a factor ~ 100 below observed metallicities at low redshift. The resolution is a *local* metallicity-dependent transition $f_{\text{VMO}}(Z)$, not a globally coordinated chronological epoch. Pop III/VMO formation naturally shuts off locally once a galactic patch exceeds a critical metallicity $Z_{\text{crit}} \sim 10^{-4}$ (cooling transition from atomic to metal-line cooling). In low-metallicity patches ($Z \ll Z_{\text{crit}}$), VMO-dominated nucleosynthesis with temporal-horizon sequestration sets the $Z \rightarrow 0$ intercept: $Y(Z=0) = Y_{\text{eq}} = 0.247$. In high-metallicity patches ($Z \gg Z_{\text{crit}}$), standard-star yields $p_{Y,\text{std}}$ set the slope:

$$\left. \frac{\Delta Y}{\Delta Z} \right|_{Z > Z_{\text{crit}}} = \frac{p_{Y,\text{std}}}{(1 - R)p_Z} = \begin{cases} 2.4 & p_Z = 0.014 \\ 1.7 & p_Z = 0.02 \end{cases} \quad \text{vs. observed } \Delta Y/\Delta Z \simeq 1.4\text{--}3, \quad (16)$$

and the solar-point check gives $Y(Z = 0.014) = 0.247 + (1.7\text{--}2.4) \times 0.014 = 0.270\text{--}0.280$, versus the observed solar $Y \simeq 0.270\text{--}0.274$. The metallicity-dependent transition model lands in the observed range with yield parameters already constrained by H II region calibrations. The single-phase Z_{eq} failure is stated honestly: the VMO-dominated equilibrium suppresses metals too aggressively for the present-day metallicity to arise from the equilibrium alone. The transition structure — VMO-dominated low- Z patches set the intercept, standard-star high- Z patches set the slope — is the internally consistent reading within this branch, and the arithmetic confirms it quantitatively. The metallicity-dependent VMO-to-ordinary-star transition introduces a function $f_{\text{VMO}}(Z)$ with a critical metallicity Z_{crit} and a transition width; while Z_{crit} is independently motivated by VMO formation physics, the transition functional form and width are constrained by VMO formation physics and stellar population evolution modelling. The transition is physically motivated by the metal-line cooling threshold at $Z_{\text{crit}} \sim 10^{-4}$ and works quantitatively with yield parameters already constrained by H II region calibrations. Crucially, this is a *local* transition occurring asynchronously across galactic patches as each reaches Z_{crit} , preserving the global eternal continuum — there are no globally coordinated "epochs." The exponent $q = 0.8$ (giving $p + q = 1.3$) is anchored to the same worldline/astration statistics that govern the Compton-exposure convergence condition ($a + p > 1$) and the $\tau(z)$ three-point requirement, so the exponent family (p, q, a) is constrained by data in one place.

Sensitivity of Y_{eq} to yield parameters. The equilibrium helium mass fraction is set by the adopted yield parameters through the baryonic-cycling fixed-point equation. Setting $dY/dt = 0$ in the GCE ODE yields:

$$Y_{\text{eq}} = \frac{p_{Y,\text{eff}}}{1 - R}, \quad p_{Y,\text{eff}} = f_{\text{VMO}} p_{Y,\text{VMO}} + (1 - f_{\text{VMO}}) p_{Y,\text{std}}, \quad (17)$$

where $p_{Y,\text{VMO}}$ is the helium yield per VMO stellar generation, $p_{Y,\text{std}}$ is the standard-star helium yield, R is the return fraction (fraction of stellar mass returned to the ISM, determined by IMF integration over measured stellar lifetimes and remnant masses), and f_{VMO} is the VMO fraction. The metal escape fraction E_Z affects the metal equilibrium Z_{eq} , not Y_{eq} directly; $E_Z \approx 0$ means metals do not escape the temporal horizon (VMO-core sequestration), keeping Z_{eq} low without altering the helium fixed point. The baseline parameter values and sensitivities are:

Parameter	Symbol	Baseline	Source	$\partial Y_{\text{eq}}/\partial(\cdot)$
VMO helium yield	$p_{Y,\text{VMO}}$	0.149	VMO wind models (Carr, Bond & Arnett 1984)	—
Standard-star He yield	$p_{Y,\text{std}}$	0.02	H II region / population synthesis	—
Effective He yield	$p_{Y,\text{eff}}$	0.1482	$f_{\text{VMO}} p_{Y,\text{VMO}} + (1 - f_{\text{VMO}}) p_{Y,\text{std}}$	$1/(1 - R) \approx 1.67$
Return fraction	R	0.40	IMF integration (Salpeter / Kroupa)	$p_{Y,\text{eff}}/(1 - R)^2 \approx 0.41$
VMO fraction	f_{VMO}	0.994	Required low-metallicity VMO-dominated branch; population closure open	$(p_{Y,\text{VMO}} - p_{Y,\text{std}})/(1 - R) \approx 0.22$
Metal escape fraction	E_Z	≈ 0	Temporal-horizon sequestration	0 (affects Z_{eq} only)

The derivative $\partial Y_{\text{eq}}/\partial R = p_{Y,\text{eff}}/(1 - R)^2 > 0$ is positive: increasing the return fraction increases Y_{eq} because more mass is recycled into new stars, amplifying the helium yield. The dominant sensitivity is to R (± 0.005 in R produces ± 0.002 in Y_{eq}); the VMO fraction is less sensitive (± 0.01 in f_{VMO} produces ± 0.002 in Y_{eq}). The observed $Y_p \approx 0.245 \pm 0.003$ is reproduced for $R \in [0.395, 0.405]$ and $f_{\text{VMO}} \in [0.975, 0.995]$, consistent with a high VMO fraction in low-metallicity stellar populations bounded by the observed $[\alpha/\text{Fe}]$ ratio in old stellar populations. A dedicated f_{VMO} sensitivity scan (step_06, $f_{\text{VMO}} \in [0.90, 1.00]$ at $\Delta f = 0.005$ resolution, results/gate6_helium_results.json) confirms that the compatible range spans $\Delta f_{\text{VMO}} \approx 0.02$, not a single fine-tuned point: Y_{eq} varies from 0.227 ($f_{\text{VMO}} = 0.90$) to 0.248 ($f_{\text{VMO}} = 1.00$), with the observed Y_p sitting in the upper portion of this range. The model therefore requires $f_{\text{VMO}} \gtrsim 0.97$: the primordial helium abundance is reproduced only when the low-metallicity stellar population is VMO-dominated, with standard stars contributing at most a few per cent by fraction. This is an aggressive requirement: the compatible range spans only $\Delta f_{\text{VMO}} \approx 0.02$ in absolute terms, and the physical motivation (metal-line cooling shuts off below $Z_{\text{crit}} \sim 10^{-4}$, channeling star formation into VMOs at low metallicity) produces $f_{\text{VMO}} \rightarrow 1$ as $Z \rightarrow 0$ without by itself selecting the specific value 0.994. The convergence condition $p + q = 1.3 > 1$ (Section 4.1) guarantees that the integrated astration exposure is finite for any individual worldline. The numerical reaction-flow calculation demonstrates convergence to the fixed point for the adopted stellar histories. To verify that $Y_{\text{eq}} = 0.247$ is a stable fixed point rather than a point dependent on the chosen history, a dedicated contraction analysis was performed over a grid of initial conditions spanning $Y_0 = 0$ to $Y_0 = 0.8$ (step_08, results/gate8_light_elements.json). All four tested initial conditions converge to $Y_{\text{eq}} = 0.247$ within $\Delta Y < 4 \times 10^{-6}$: $Y_0 = 0.00 \rightarrow Y_{\text{final}} = 0.24704$, $Y_0 = 0.25 \rightarrow Y_{\text{final}} = 0.24704$, $Y_0 = 0.50 \rightarrow Y_{\text{final}} = 0.24704$, $Y_0 = 0.80 \rightarrow Y_{\text{final}} = 0.24705$. Convergence to the fixed point is the standard relaxation of a first-order linear ODE — guaranteed for any positive return fraction — so the equilibrium helium-4 mass fraction is independent of the tested initial compositions under the adopted reaction-flow model. The substantive content is the value of the fixed point and its sensitivity to the adopted yield parameters: $Y_{\text{eq}} = 0.247$ is a robust prediction of the temporal-horizon reaction-flow framework rather than an artifact of a particular initial condition.

Steady-state timescale and thermalization budget (Gate 11). The timescale ledger for the eternal-universe claims is computed in step_11 (results/gate11_steady_state_budget.json). Three quantities close the budget. (i) *Equilibrium convergence.* The reaction-flow map contracts every abundance residual at rate $(1 - R)/\tau_*$; from the extreme initial compositions tested above, the solution relaxes to within 1% of the asymptotic abundances in $\approx 9 \tau_*$ and to 0.1% in $\approx 13 \tau_*$ — of order ten star-formation timescales, a negligible fraction of the unbounded coordinate time available on the eternal slice. (ii) *Thermalization depth.* In the static spatial geometry the absorbing column along a line of sight is unbounded: the Step-07 path-length divergence applies to any nonzero absorptivity, not only to Thomson scattering. Evaluating the accumulated absorption optical depth $\tau_{\text{abs}}(z) = \rho_b \kappa \ell(z)$ with the corpus baryon density $n_b = 0.251 \text{ m}^{-3}$ and path-length unit c/H_0 , the unit-optical-depth surface for dust-grade absorptivity at mm/sub-mm wavelengths ($\kappa = 10^{-4} - 10^{-2} \text{ cm}^2 \text{ g}^{-1}$ per gas mass, bracketing the Planck-inferred dust emissivity range) lies at temporal depth $z_{\text{therm}} \approx 130 - 1340$ under the power-law horizon law ($p = 0.5$), with $\tau \rightarrow \infty$ beyond by the Global Opacity Theorem — a range that brackets the recombination depth $z_* \approx 1079$ evaluated in TEP-TH, so the surface at which distributed starlight completes reprocessing and the last-scattering plasma occupy the same temporal-depth sector, where $T_{\text{loc}} = T_0(1 + z_*) \approx 2.9 \times 10^3 \text{ K}$ is the hydrogen ionization edge. Distributed starlight is therefore reprocessed into a blackbody at finite temporal depth by a geometric theorem, with the thermalization timescale unconstrained because coordinate time is unbounded; the Planckian form of the result is preserved achronically by the Gate-5 shear invariance. (iii) *Sequestration ledger.* Per star-formation timescale a fraction $1 - R = 0.6$ of the gas mass is locked into remnants; metal overproduction is prevented because the dominant $f_{\text{VMO}} = 0.994$ channel sequesters its metal yield behind temporal horizons ($E_Z \approx 0$), holding the asymptotic metallicity at $Z_{\text{eq}} = 10^{-4}$ rather than permitting monotonic enrichment, while total astration exposure remains finite by the $p + q = 1.3 > 1$ condition. Baryon exhaustion is precluded by the open-box steady state, in which the gas mass is held constant by the balance of return and sequestration-driven replenishment on the eternal slice.

Ionization-equilibrium budget of the VMO branch (Gate 12). The consequence check demanded by the VMO-dominated low-metallicity IMF is computed in step_12 (results/gate12_vmo_ionization_budget.json). In the eternal steady state there is no reionization epoch to fund: ionization is a maintained balance, and the constraint per baryon is $\xi f_{\text{esc}} N_{\gamma}^{\text{eff}} \gtrsim \alpha_B n_e$, with ξ the gas-cycling rate, f_{esc} the Lyman-continuum escape fraction, and N_{γ}^{eff} the IMF-weighted ionizing yield per cycled baryon. Metal-free very massive stars emit $\sim 10^5$ H-ionizing photons per stellar baryon over the $\sim 2\text{--}3$ Myr main sequence that precedes temporal-horizon formation^{[77], [78]}, versus $\sim 4 \times 10^3$ for a standard IMF, so the $f_{\text{VMO}} = 0.994$ mixture supplies $N_{\gamma}^{\text{eff}} \approx 10^5$ photons per cycled baryon; the horizon sequesters the metal-bearing remnant, not the emitted radiation field. Against the case-B recombination demand^[79] at the corpus density, the ionized column persists to $n_e \sim 5 \times 10^2 - 2 \times 10^7 \times n_{b,0}$ across the scanned plane $\xi \in [10^{-3}, 1] \text{ Gyr}^{-1}$, $f_{\text{esc}} \in [0.01, 0.5]$ — bracketing the $\sim 500\times$ electron-density rise the opacity three-point requirement demands at the boundary depth — while a purely standard IMF saturates $\sim 25\times$ shallower and fails the deepest level at the conservative corner ($n_{e,\text{crit}} \approx 19 \times n_{b,0}$). The constraint direction therefore runs with, not against, the helium-driven branch: the minimum f_{VMO} that sustains ionization at the $500\times$ depth vanishes over most of the plane and rises toward unity only at the extreme low-escape, low-cycling corner, so the ionization sector adds no countervailing bound on the required f_{VMO} and mildly favours the high-VMO branch where the bounds are tightest. The Planck datum $\tau_{\text{reion}} = 0.054$ is likewise not an epoch budget under TEP but the Thomson column accumulated in the optically thin approach to the single opacity boundary: at the constant- n_e baseline of the $p = 0.5$ horizon law the approach column reaches 0.053 by $z \approx 4$ with the $\tau = 1$ surface at $z \approx 20$, and a mild approach-profile rise ($n_e \propto (1+z)^{1/2}$) reproduces the same amplitude by $z \approx 3$ — the observed wing selects a near-flat to mildly-rising $n_e(\ell)$ whose definitive profile belongs to the absorber-statistics inversion of the opacity programme. The remaining consequences dissolve in the steady state: there is no global Population III chronology to contradict because the VMO branch is a local metallicity-selected channel asynchronous across patches, and all processed starlight — whatever its spectral hardness — enters the Gate-11 thermalization bath, so no separate extragalactic-background-light budget arises.

CMB baryon density concordance (Gate 9). The most powerful argument for standard BBN is that the deuterium abundance predicts a baryon density $\Omega_b h^2 \approx 0.022$ ^[65] that matches the independent value derived from Planck CMB acoustic peak ratios^[64]. Under TEP, the observed D/H ratio is a localized kinematic signature rather than a primordial abundance, so the $\text{D/H} \rightarrow \Omega_b$ mapping is not the source of the baryon density constraint. However, the CMB anchors $\Omega_b h^2$ via acoustic physics whose TEP reconstruction is supplied by TEP-HC (Paper 18,^[3]), which fits $\Omega_b h^2 = 0.02144 \pm 0.00257$ directly from Planck acoustic peaks via joint MCMC, with sound-horizon ratio $r_s^{\text{TEP}}/r_s^{\text{ACDM}} = 0.999994$ (< 6 ppm deviation); that anchoring is inherited here and is conditional on Paper 18's acoustic-sector analysis. The acoustic peaks rigidly demand this specific baryon density to balance the compression and rarefaction peaks of the photon-baryon fluid, regardless of the D/H ratio. Gate 9 is therefore stated as: *given the rigidly CMB-anchored baryon density $\Omega_b h^2 \approx 0.022$, verify that the SQUAD/LLS neutral-hydrogen column-density distribution folded through the TEP edge-gas shear model produces a D/H histogram consistent with observed quasar sightline statistics without free density tuning.* This computation is tractable on existing LLS catalogs but is not performed in this paper. The corpus-level cross-reference is to TEP-HC (Paper 18,^[3]) and TEP-C0 (Paper 26,^[2]), where the acoustic peak ratios are addressed independently of the $\text{D/H}-\Omega_b$ mapping.

Pressure sensitivity. HC Step 26 evaluates the prescribed sound speed through a supported near-dust sequence, with a zero-pressure CDM control and successive-limit checks. Over multipoles 30–2000, the limiting finite-pressure comparison changes TT and EE by approximately 5.44% and 6.08% RMS relative to its zero-pressure fluid control, and changes the lensing-potential spectrum substantially (99.57% RMS fractional difference). The zero-pressure control differs from CDM by at most 0.25% RMS across the tested spectra. These are phenomenological sensitivity results without a likelihood refit; the pressureless full-solver fit therefore cannot be identified with the finite-pressure effective model merely from agreement of a few peak locations.

5. Discussion and Falsifiable Predictions

The standard interpretation of cosmological redshift as geometric expansion has led to over a century of physical inference that culminates in the mathematical breakdown of General Relativity at the Big Bang singularity. Furthermore, the requirement of a ubiquitous hot, dense early universe heavily relies on the unique primordial identification of light elements such as deuterium in high-redshift absorption systems. Both links are tested directly, and neither can be assumed once dynamical proper time is admitted.

5.1 Distance Duality and Cosmological Tests

Distance Duality and Supernova Standardization

Critically, $T_{\text{obs}}(z) \neq T_{\text{loc}}(\tau)$ in general. The temperature of the background radiation bath as measured by an observer is distinct from the actual local matter/radiation state $T_{\text{loc}}(\tau)$ at emission. Furthermore, because physical space is static, the standard geometric distances must be carefully defined. Etherington's reciprocity theorem is a general result of metric photon propagation, not specific to expanding FLRW; it dictates the distance-duality relation $d_L = d_A(1+z)^2$ for any metric theory where photon geodesics are well-defined and photon number is conserved. In the TEP framework, the physical matter space is static ($a_m = 1$), but the conformal coupling $A(\phi)$ acts identically to the FLRW scale factor for photon transport. Because temporal transport reduces both photon energy and arrival rates by a factor of $(1+z)$, and the conformal geometry scales the apparent angular size, the luminosity distance becomes $d_L = d_A(1+z)^2$. The TEP framework therefore preserves the Etherington relation by construction, not because it replicates FLRW expansion, but because the conformal transport law satisfies the same general conditions.

While the baseline distance-duality relation is preserved, it is important to recognize that SNIa magnitudes are not raw observables. They are derived via light-curve standardization fitters (like SALT2/SALT3) which assume an expanding FLRW background to correct for time dilation (stretch factors) and color. A preliminary evaluation has already been completed in TEP-C0 (Paper 26): the SALT2 light-curve stretch parameters (x_1) from the 1,701 supernovae in the Pantheon+ dataset were tested against the exact covariant TEP conformal factor. Under standard Λ CDM time dilation $(1+z)$, the fit to observed stretch parameters yields a reduced χ^2 of 102.6; under the TEP conformal factor, the reduced χ^2 improves to 88.9. This is a diagnostic consistency check showing consistency with temporal-field stretching of supernova light curves as predicted by TEP, though the result remains contingent on SALT2 standardization assumptions. The remaining observational task is the full re-calibration of the light-curve fitters within the TEP geometry rather than relying on FLRW-calibrated nuisance parameters — a concrete, falsifiable roadmap with the stretch-factor directionality already verified.

5.2 Line-of-Sight Transport and Optical Time-Transfer

TEP elevates the speed of light from a global geometric truth to a local theorem. This provides falsifiable physical predictions. Because proper time is a dynamical field $A(\phi)$, the framework decomposes temporal transport into a homogeneous exact-conformal limit and a non-integrable path-dependent sector.

As detailed in Section 4, the conformal piece ($\Sigma_{||}$) is endpoint-dependent and vanishes on closed loops, whereas the disformal transport ($\mathcal{C}_T$) supplies genuine non-integrability. Multi-leg optical time-transfer experiments—currently within reach of next-generation atomic clock networks—can directly test for this line-of-sight transport integral ($\oint \mathcal{C}_{T,||} d\ell \neq 0$).

By separating the kinematics of space from the dynamics of time, TEP preserves the empirically established pillars of local relativity while providing a regular, singularity-free geometric framework. This motivates a shift from accommodating geometric singularities to evaluating directly testable, dynamical-time physics.

5.3 What Would Falsify TEP

The framework makes several concrete falsifiable predictions and discriminating tests. Failure has the consequence stated for each test below.

1. **Disformal clock three-sightline test:** TEP predicts a geometry-dependent Δv_T via the temporal-disformal channel, with the fitted free-H displacements ($-78.9, -76.0, -118.0$ km/s) as the maximal amplitude targets under a pure-TEP attribution. The -81.6 km/s clustering is explained by the velocity-selection theorem (Gate 6-I) as a selection effect. Diagnostic mode establishes algebraic feasibility; transfer mode shows the simple power-law $B(\phi) \propto \phi^5$ does not transfer across redshifts. Predictive closure — a single frozen microscopic $B(\phi)$ with independently specified absorber environments — remains the decisive test. If no physically motivated $B(\phi)$ can produce an independently predicted temporal contribution compatible with the observed displaced-H phenomenology while satisfying Cassini, GW170817 and matter-metric consistency constraints, the temporal-disformal interpretation would be falsified.
2. **Orphan population (decisive empirical test):** The velocity-selection theorem (Gate 6) predicts an excess of narrow H I components at non-isotope offsets beyond the conventional background. The illustrative forecast (Gate 6-II) gives ~ 19 orphan features per ~ 80 parent sightlines before observability cuts; the archival feasibility analysis reduces this to approximately two detectable excess candidates per 80 parent sightlines, corresponding to roughly ten usable

systems. If a re-analysis of existing archival spectra shows no excess of narrow ($b < 10$ km/s), low-column ($\log N \sim 12\text{--}14$) components at non-isotope offsets beyond the normal H I population, the TEP interpretation is substantially weakened.

3. **Spectroscopic robustness:** Scaffold freedom with five freed components strengthens the calibrated significance of the Q1009 free-H preference (T increases from 161.03 to 398.59, $p = 0.0050$ with 0/200 exceedances, Section 5.5). The Q1009 evidence is therefore demonstrably robust to relaxing the published component architecture, proving that absorber re-optimisation strengthens the discrimination. Across all three sightlines, primary and adversarial parent calibrations ($N = 1,000$ draws each) decisively favor unrestricted hydrogen ($p_{\text{std}} = p_{\text{parent}} = 0.001$), confirmed under physical parent windows (C1) and Gate 8 joint diagnostic tests ($p \leq 0.023$). Full multi-component scaffold re-optimization on PKS 1937–101 (184 components) and J1332+0052 represents a computational scaling extension of this established robustness proof.
4. **Helium equilibrium:** If Y_{eq} moves outside the observed range 0.245 ± 0.003 under physically motivated variations in (p_Y , R , f_{VMO}) within their observationally constrained ranges, the baryonic-cycling explanation for the helium-4 mass fraction would be falsified. The load-bearing requirement is stated as an independently falsifiable condition: the mechanism needs $f_{\text{VMO}} \gtrsim 0.97$ of low-metallicity star formation channelled into temporal-horizon-sequestering VMOs, a condition testable by Population III stellar archaeology and extreme-metal-poor star counts rather than by this paper's internal consistency. The ^3He and ^7Li equilibria inherit the same sequestration ansatz, so the falsification of the f_{VMO} requirement propagates to those rows as well.
5. **CMB acoustic structure:** If the TEP perturbation theory failed to reproduce the Planck acoustic peak ratios and damping tail within observational uncertainties, the cosmological framework would be ruled out; companion TEP-HC (Paper 18) verifies full Planck high- ℓ TTTEEE + low- ℓ + lensing + BAO + Pantheon+ compatibility ($\Delta\chi^2 = -0.11$ vs ΛCDM , sound-horizon ratio $r_s^{\text{TEP}}/r_s^{\Lambda\text{CDM}} = 0.999994$).
6. **Growth of structure:** If the static-spatial TEP geometry cannot reproduce the observed S_8 growth parameter or the BAO scale, the framework fails as a cosmological alternative.

5.4 Outstanding Observational Tests

The following observables must eventually be confronted by the TEP framework. The current status of each is stated honestly:

Observable	Status	Reference
CMB acoustic peaks and damping tail	Verified in native <i>hi_class</i> Boltzmann MCMC (TEP-HC, Paper 18): full Planck TTTEEE $\Delta\chi^2 = -0.11$ vs ΛCDM ; pre-recombination sound horizon preserved at < 6 ppm ($r_s^{\text{TEP}}/r_s^{\Lambda\text{CDM}} = 0.999994$)	[16]
CMB temperature–redshift relation	Preserved by conformal transport $T_{\text{loc}}(z) = T_0/A_{\text{clock}}(z) = T_0(1+z)$ and corroborated against published CO absorbers ($z = 0.89\text{--}2.69$, $\chi^2/\text{dof} = 0.15$, $\beta = +0.003 \pm 0.023$, constant- T rejected at 29σ)	Section 4.1; TEP-TH Step 11
BAO and distance-ladder consistency	Distance duality preserved; BAO native projection $\chi^2/\text{dof} = 0.88$ (17 points) and SALT2 stretch-factor directionality verified in TEP-C0 ($\chi^2 102.6 \rightarrow 88.9$); Cepheid ladder yields $H_0 = 67.83 \pm 1.56$ km s $^{-1}$ Mpc $^{-1}$ (0.26σ from Planck; TEP-H0)	Section 5.1; Papers 11, 26
Growth of structure / S_8	Active-scalar screened closure in TEP-HC yields $\sigma_8 = 0.824$ ($Z_E \gtrsim 700$) to $\sigma_8 \simeq 0.76$ ($Z_E \sim 6 \times 10^3$), two-branch $G_{\text{eff}}(z=0) \simeq 1.16 G$ (Paper 0 Step 61), and unstretched-space CMB lensing boost $9.4 \times 129 \times$ without CDM (TEP-HC Step 21)	Papers 0, 18
D/H in non-benchmark sightlines	Framework predicts localized shear amplitude varies with core-edge geometry; population study is open	—
$\Delta Y/\Delta Z$ slope in H II regions	Qualitative prediction derived; quantitative comparison deferred	Section 4.4
Lyman-limit / Gunn–Peterson trough redshift	Optical-depth divergence predicted; species-resolved admissible profile places $\tau_e = 1$ at $z \approx 1080$ coincident with the thermalization edge, with the trough derived to be neutral ($\bar{n}_e < 3.7 \times 10^{-4} n_{b,0}$); boundary transport audited and self-consistency solved at toy level (Steps 16, 18) — the surface is the ionization front of the deep thermalized zone, not a physical wall; the broad tail is a closure artifact removed once the edge is absorptive ($\kappa_{\text{abs}} \gg 11 L_0^{-1}$, the thermalization condition), restoring $\sim 87\%$ front-column weight with foreground damping $e^{-2\tau_{\text{fg}}} \approx 0.85$ vs the standard 0.90; front gas temperature derived by Compton	Section 4.2

Observable	Status	Reference
	equilibrium with the Gate-11B bath ($T_{\text{gas}} = 2946$ K, coupling ~ 1 yr), making the $x_e \simeq 0.83\text{--}0.99$ Saha ionization derived rather than assumed; the Γ drop $\sim 4 \times 10^{10}$ supplied by self-shielding over ~ 24 LyC e-folds (~ 6 Mpc); front maintenance met for cycling $\xi \gtrsim 10^{-5}$ Gyr $^{-1}$; mode spectrum and full-catalogue inversion remain open	

5.5 Spectroscopic Robustness Checks

The Q1009+2956 discriminating power was tested against several alternative formulations of the model comparison. Four robustness checks are reported here.

Limited Scaffold Freedom (C2)

The five strongest non-parent H I components within $|v| \leq 1000$ km s $^{-1}$ have their columns and widths re-optimized under both hypotheses, while their velocities remain fixed. This specified change of scaffold freedom raises the observed likelihood statistic from $T = 161.03$ to $T = 398.59$. Relative to the fixed scaffold, the log-likelihood improvements are 55.01 for D and 173.79 for H. The fit uses total Doppler width over the original thermal/turbulent bounds and a stable weighted ridge continuum solve. The Monte-Carlo calibration uses the same configured candidate grids, baseline starts, matched-D hydrogen start and optimizer settings as the observed fit: 0/200 exceedances, $(k+1)/(N+1) = 0.0050$. This is a conditional robustness test of the specified five-component freedom, rather than unrestricted reconstruction of the entire absorber.

Parent-Width-Constrained H Control (C4)

To test whether the candidate feature is merely an ordinary H I component shifted by a free velocity, the data were also fitted with a two-parameter *kinematic* model: free v_H and $\log N_H$, but b fixed to the parent H I value (no isotope mass scaling and no free b). The free-H profile has three identifiable coordinates $(v, \log N, b)$, and the D profile has two $(\log N, b)$: temperature and turbulent width enter an isolated single-isotope line only through the total width. The information-criterion penalties use these identifiable counts. The general ordinary-H alternative is represented by $M_{H,\text{free}}$. This model is strongly rejected: $T(M_H vs M_D) = 161.03$, $T(M_H vs M_{\text{kin}}) = 208.98$, $T(M_{\text{kin}} vs M_D) = -47.95$. $\Delta\text{AIC}(M_{\text{kin}} - M_D) = +47.95$ and $\Delta\text{AIC}(M_H - M_{\text{kin}}) = -206.98$. The best-fit kinematic velocity is $v_H = -130.52$ km/s, but the parent-like b value cannot reproduce the narrow candidate core. A displaced hydrogen feature must also reproduce the observed narrow width; fixing that width to the parent value fails this control.

Noisy/Offset True-D Monte Carlo (C3)

The standard calibration generates exact true-D flux from the maximum-likelihood D parameters. A more realistic test perturbs the generating true-D nuisance parameters before drawing each realisation: $\log N_D$ by ± 0.1 dex, T_K by ± 1000 K (clamped to > 100 K), and b_{turb} by ± 1 km s $^{-1}$. The D/H velocity offset (-81.6 km/s) is a fixed physical constant (exact reduced-mass isotope shift) and is not perturbed. For Q1009+2956, 0/1000 standard realisations exceed T_{obs} ($p_{\text{std}} = 0.001$), and 0/1000 parent-reassignment realisations exceed T_{parent} ($p_{\text{parent}} = 0.001$). This tests the specified nuisance-parameter perturbations while keeping the physical isotope shift fixed.

Physical Parent Window (C1) — Additional Robustness Test

The conservative ± 1000 km/s parent window is the cross-sightline headline statistic; the $\pm 3b_{\text{max}}$ physical window has also been calibrated for all three sightlines ($N = 1000$ each). For Q1009+2956 the largest b -value of the main complex is 56.7 km s $^{-1}$, so the physical window $\pm 3b$ is ± 170 km s $^{-1}$, compared with the conservative ± 1000 km s $^{-1}$. This reduces the eligible-parent set from 14 in the conservative ± 1000 km s $^{-1}$ window to 4 in the physical window. With $N = 1000$ true-D Monte Carlo realisations the fair standard statistic gives 0/1000 exceedances ($p_{\text{std}} = 0.001$), and the parent-reassignment statistic also gives 0/1000 exceedances ($p_{\text{parent}} = 0.001$). For PKS 1937–101 the physical window is ± 173 km/s ($b_{\text{max}} = 57.7$ km/s), giving 0/1000 exceedances under both statistics ($p_{\text{std}} = p_{\text{parent}} = 0.001$). For J1332+0052 the physical window is ± 59 km/s ($b_{\text{max}} = 19.5$ km/s), giving 0/1000 exceedances under both statistics ($p_{\text{std}} = p_{\text{parent}} = 0.001$). The results are consistent with the fixed window across all three sightlines: the H alternative is significantly preferred under both the canonical-parent and adversarial-parent tests. The physical window is motivated by the absorber's own velocity structure rather than an arbitrary fixed range.

5.6 Multi-Pronged Evidence Analysis (Gate 8)

Three discriminators between the isotope-D and TEP edge-gas interpretations have been quantitatively evaluated using the fitted free-H parameters from all three benchmark sightlines: the Doppler b -value ratio, the velocity offset deviation, and the column density

difference. These quantities are all derived from the same free-H fit and are therefore not statistically independent; the joint multivariate statistic is the appropriate calibrated measure. A true-D injection Monte Carlo (step_08b, $N = 1000$ realisations per sightline) calibrates each statistic against the null distribution obtained by fitting $M_{H,\text{free}}$ to simulated true-D spectra. The b -value ratio uses the mixed thermal+turbulent prediction $b_D = \sqrt{b_{\text{turb}}^2 + 2kT/m_D}$, not the purely thermal $b_D = b_{\text{parent}}/\sqrt{2}$, since the turbulent contribution does not scale with mass. Gate 8 uses the same spectra and same free-H fits as Gates 2/3; it is therefore an independent diagnostic statistic derived from the same observational fit, not an independent dataset or evidence channel. Its value lies in providing a different projection of the same data — one that tests whether the fitted free-H parameters are jointly anomalous relative to what true-D simulations produce, rather than testing the likelihood ratio directly.

b-value ratio. The reference width uses the mixed thermal+turbulent prediction $b_D = \sqrt{b_{\text{turb}}^2 + 2kT/m_D}$. The true-D injections use the fitted $M_{D,\text{free}}$ temperature and turbulent width, which can differ from this parent reference; the ratio need not have null mean 1. The predictive consistency test therefore measures departures from the empirical injection mean, with each null draw excluded from its own reference estimate. The results are Q1009: observed ratio 0.99, null mean 0.919, calibrated $p = 0.003$; PKS1937: observed ratio 1.31, null mean 1.187, calibrated $p = 0.092$; J1332: observed ratio 5.31, null mean 1.162, calibrated $p = 0.065$. J1332 remains a broad, prior-sensitive feature ($b_{H,\text{free}} \approx 51 \text{ km s}^{-1}$), distinct from the narrow, cool edge-gas orphan phenotype. Its unrestricted-H likelihood preference does not by itself identify that physical mechanism. Expanded turbulent-width fits retain a likelihood preference but are not assigned an expanded-bound Monte-Carlo significance here.

Velocity deviation. The observable is the free-H velocity relative to the isotope-shifted parent position. The injection-calibrated results are Q1009: $+2.72 \text{ km s}^{-1}$, centered calibrated $p = 0.001$; PKS1937: $+5.66 \text{ km s}^{-1}$, centered calibrated $p = 0.001$; J1332: -36.60 km s^{-1} , centered calibrated $p = 0.098$. The two high-resolution unsaturated systems test isotope locking most tightly; the saturated J1332 system has a much broader, non-Gaussian injection distribution. The calibration describes fitting scatter under the specified true-D model. Wavelength and parent-centroid uncertainties are evaluated separately below, and the $N = 3$ column-density ordering is descriptive rather than an independently calibrated correlation test.

Sign test. The conformal sign theorem predicts all edge-shear features should be blueward of the parent. All three are: $-131.3 < -52.4$, $-67.3 < +8.7$, $-118.2 < 0.0 \text{ km/s}$. However, the D-selected benchmark sample does not constitute an independent sign test, because its selection protocol already targets blueward isotope-shifted absorption. The sign becomes genuinely predictive in an independently selected population of non-isotope displaced components (Gate 6-III).

Column density difference. The reference is $N_{\text{HI}}(\text{D}/\text{H})_{\text{BBN}}$. Because the fitted true-D generating model has a free column density, its injected free-H column differences need not center on zero. The predictive test uses the empirical null mean consistently with the width and velocity diagnostics. The results are Q1009: observed $+0.034 \text{ dex}$, null mean -0.010 dex , calibrated $p = 0.001$; PKS1937: observed -0.094 dex , null mean -0.204 dex , calibrated $p = 0.010$; J1332: observed $+0.698 \text{ dex}$, null mean -0.848 dex , calibrated $p = 0.015$. These correlated diagnostics are projections of the same spectra and are not multiplied as independent evidence.

Calibrated joint assessment. The three diagnostic coordinates are standardized by their null standard deviations. The covariance metric uses the correlation matrix with a dimensionless 10^{-6} ridge, preserving the result under changes of units. Each null draw is scored against the mean and covariance estimated without that draw; the observation is scored against all $N = 1000$ injections. The table reports finite-sample $(k+1)/(N+1)$ probabilities and squared Mahalanobis distances. Zero exceedances give $p = 0.000999$ at this resolution and a one-sided 95% binomial upper bound of approximately 0.003, rather than a measured zero tail probability.

Sightline	Joint exceedances	Calibrated p	Squared distance
Q1009+2956	0/1000	0.001	288.06
PKS1937-101	0/1000	0.001	129.02
J1332+0052	22/1000	0.023	14.68

Q1009, PKS1937 survive the separate Gate-8 family threshold $\alpha/3 = 0.0167$. Gate 8 reuses the Gate 2/3 spectra and fits and supplies diagnostic interpretation, not independent confirmation. J1332-specific width and velocity conclusions remain conditional on the declared turbulent-width bound.

Wavelength-calibration systematics. The relevant uncertainty is relative wavelength distortion within the same spectrum. The documented instrument-specific assumptions are $\sigma_{\text{sys}} = 0.25, 0.05$ and 0.20 km s^{-1} for Q1009, PKS1937 and J1332, respectively. The Gaussian sensitivity approximation uses the observed displacement from the empirical injection mean and $\sigma_{\text{total}}^2 = \sigma_{\text{fit}}^2 + \sigma_{\text{sys}}^2$. It gives Q1009: 9.16σ , Gaussian-model $p = 5.06e - 20$; PKS1937: 10.99σ , Gaussian-model $p = 4.16e - 28$; J1332: 1.97σ , Gaussian-model $p = 4.94e - 02$. These small Gaussian tails are model extrapolations, distinct from the empirical injection p-values. The 100 000 sensitivity draws report exceedance counts and finite-sampling bounds; they do not resolve arbitrarily small tails. The common $[0, 1] \text{ km s}^{-1}$ scan is retained in `results/gate8_wavelength_systematics.json`.

Parent-centroid blend bias. A blend can bias the fitted parent centroid and thereby shift the isotope anchor. The combined Gaussian sensitivity calculation uses the empirical-null-centered deviation and $\sigma_{\text{total}}^2 = \sigma_{\text{fit}}^2 + \sigma_{\text{sys}}^2 + \sigma_{\text{blend}}^2$, with 100 000 Monte-Carlo draws per case. At $\sigma_{\text{blend}} = 1 \text{ km s}^{-1}$, the Gaussian-model results are Q1009: $p = 9.02 \times 10^{-3}$; PKS1937: $p = 5.20 \times 10^{-7}$. For the distinct physical scenario in which centroid drag supplies the entire raw observed offset, a two-sided blend-only probability of 0.05 requires $\sigma_{\text{blend}} = 1.39$ and 2.89 km s^{-1} for Q1009 and PKS1937, respectively. The corresponding offset-to-parent-linewidth ratios are 0.15 and 0.35. The illustrative unresolved-companion estimate $f \gtrsim 2|\Delta v|/b_{\text{parent}}$ gives 0.29 and 0.69; this is a blend-geometry approximation rather than a fitted exclusion of every companion profile. The sign-matched blend-only product assumes independent, sign-random Gaussian biases across the two sightlines; it is a conditional sensitivity result and is not multiplied into the main spectral evidence. The full scans, assumptions and finite-sampling bounds are reported in `results/gate8_parent_centroid_systematics.json`.

5.7 Orphan Search: Protocol and Pilot Results (Gate 6-III)

The population forecast (Gate 6-II) predicts ~ 19 orphan narrow components at non-isotope offsets across ~ 80 high-S/N LLS sightlines. A pilot search on the three benchmark sightlines has been carried out using both the published VPFIT component decompositions and direct spectral inspection of the Ly α region (step_06_orphan_search, results/gate6_orphan_search.json).

Pilot search results. The VPFIT models and HIRES/ESPRESSO spectra were searched for narrow ($b < 15 \text{ km/s}$), low-column ($\log N < 15$) H I components at non-isotope offsets ($|\Delta v - v_{\text{iso}}| > 10 \text{ km/s}$) within $\pm 300 \text{ km/s}$ of the parent. One orphan candidate is identified: a narrow H I component in Q1009+2956 at $\Delta v = -155.5 \text{ km/s}$ from the parent ($|\Delta v - v_{\text{iso}}| = 73.9 \text{ km/s}$), with $b = 10.3 \text{ km/s}$ and $\log N = 11.5$, detected at 5σ on a single coadd in the HIRES spectrum. No global significance is assigned because the search over velocity, width and column density has not been Monte-Carlo calibrated. A second absorption feature at $\Delta v = -170 \text{ km/s}$ in PKS 1937–101 is detected at 5σ on a single coadd in the ESPRESSO spectrum, but the VPFIT decomposition fits it as a broad component ($b = 31.5 \text{ km/s}$); a narrow orphan may be hidden within the broad fit. J1332+0052 is not searchable: the Ly α is completely saturated ($\log N_{\text{HI}} = 19.25$) and Ly β falls outside the UVES wavelength coverage. Under the corrected bounded gate (Gate 10-b; Section 3) the disformal transport shift saturates near $|\Delta v| \approx 140 \text{ km/s}$ for any gate that reproduces the observed amplitudes, so features at $|\Delta v| > 140 \text{ km/s}$ cannot be TEP-shifted edge gas under the same activation law: the two reported pilot features at -155.5 and -170 km/s lie beyond this ceiling and are re-classified as ordinary forest H I, not viable orphan candidates. The pilot window is accordingly ceiling-constrained to $|\Delta v| \leq 140 \text{ km/s}$ in the corrected search.

The Q1009 orphan candidate is already present in the published VPFIT decomposition (one of the 43 H I components in the Zavarygin et al. 2018 model). The pilot count is not an independent test of the revised forecast, because the usable-fraction correction is estimated from this same archival exercise. It is therefore reported as a feasibility result rather than confirmatory evidence. Three factors suppress the yield: (i) the parent H I is saturated from $\Delta v \approx -50$ to $+100 \text{ km/s}$ at all three sightlines, hiding any orphans in that window; (ii) the VPFIT models only contain components above the statistical significance threshold, so sub-threshold orphans are invisible; (iii) σ_{shear} is estimated from only 3 data points (the D-identified features) and is biased small, understating the predicted velocity spread. The full population test requires the ~ 80 -sightline archival search described below.

Archival search. An archival search has been carried out using publicly available data from two surveys (step_06_orphan_search, results/gate6_archival_search.json):

- *SQUAD DRI* (UVES Spectral Quasar Absorption Database, Murphy et al. 2018^[43]): 467 UVES quasar spectra at $R \sim 40\,000$ – $50\,000$, accessed via the ESO Phase 3 archive. Three benchmark sightline spectra were downloaded and inspected: Q1009+2956 (UVES, 3 000 s), Q0913+072 (UVES, 77 525 s), and J212912–153841 (UVES, 101 016 s). The SQUAD UVES spectrum of Q1009+2956 shows a feature at $\Delta v \approx -155 \text{ km/s}$ consistent with the HIRES orphan candidate, though at lower S/N (~ 33 vs ~ 87 for HIRES); this is a consistency check on the same sightline, not an independent detection. Q0913+072 is a DLA ($\log N_{\text{HI}} = 20.15$) with all Lyman series lines through Ly ϵ saturated. J212912–153841 has an LLS at $z = 2.968$ ($\log N_{\text{HI}} = 17.32$), but the Ly α forest at $z \sim 3$ produces near-zero flux throughout the $\pm 300 \text{ km/s}$ search window.
- *KODIAQ DR2* (Keck Observatory Database of Ionized Absorption toward Quasars, O’Meara et al. 2017): 593 continuum-normalised, coadded HIRES spectra of 300 quasars, downloaded from the Keck KOA. The KODIAQ Z survey (Lehner et al. 2016) provides a catalog of 26 pLLSs/LLSs at $2.3 < z < 3.3$ with $\log N_{\text{HI}} = 16.16$ – 18.65 , of which 9 are LLS-class ($17.0 \leq \log N_{\text{HI}} \leq 19.0$). Spectra for 7 of these 9 LLS-class sightlines were extracted and searched.

Archival search results. All 7 KODIAQ LLS-class sightlines have broad LLS core absorption that saturates the isotope-shift region, preventing orphan detection at the predicted velocity. The saturation arises from two compounding effects: (i) the LLS itself produces Ly α absorption covering $\sim 100 \text{ km/s}$, and (ii) the Ly α forest at $z > 2.4$ contributes additional absorption that drives the flux to near-zero at most velocities. The only sightline where the isotope-shift region is recovered is Q1009+2956, where the LLS has a relatively narrow core (saturating only from $\Delta v \approx -50$ to $+100 \text{ km/s}$) and the forest at $z = 2.504$ is sufficiently sparse. This identifies a critical selection criterion for the population test: only sightlines with (a) narrow LLS cores ($< 50 \text{ km/s}$ saturated width), (b) S/N > 30 per pixel, and (c) $z < 2.6$ are suitable for orphan detection. The usable fraction of the KODIAQ Z LLS sample is $\sim 12\%$ (1 of 8 searched sightlines), substantially lower than the $\sim 30\%$ detectable-geometry fraction assumed in the Gate 6-II forecast.

Search protocol. For each sightline: (i) identify all narrow ($b < 10 \text{ km/s}$) H I components with $\log N \sim 12$ – 14 within $\pm 200 \text{ km/s}$ of the D/H parent complex, (ii) record the velocity offset Δv from the nearest H I parent, (iii) classify each as *isotope-consistent* ($|\Delta v - v_{\text{iso}}| < 5 \text{ km/s}$) or *orphan* ($|\Delta v - v_{\text{iso}}| > 5 \text{ km/s}$), (iv) record b , $\log N$, and S/N. An illustrative selection model calibrated

to the three benchmark offsets gives an orphan fraction of $\sim 79\%$ (orphan features outnumber isotope-consistent ones by $\sim 3.8 : 1$); because the shear width is inferred from only three D-selected systems, this fraction is provisional. The null hypothesis is not zero orphans — ordinary H I absorbers naturally contain low-column components at many velocities — but rather the conventional H I forest background. A proper null must be constructed from velocity-shuffled windows, matched control absorbers, neighbouring Δv intervals, or synthetic conventional H I forests. The robust TEP prediction is an *excess* of components with the joint phenotype ($b < 10$ km/s, $\log N \sim 12\text{--}14$, non-isotopic offset, systematically related to absorber geometry) beyond this conventional background.

Pilot feasibility forecast. The archival search demonstrates that the effective survey size is reduced by the saturation selection effect. If the usable fraction is $\sim 12\%$ rather than the assumed $\sim 30\%$, the effective survey size for a 80-sightline parent sample is ~ 10 sightlines, predicting ~ 2 excess orphans under TEP. A pilot on 10–15 *usable* sightlines (requiring a parent sample of $\sim 80\text{--}120$ LLS sightlines to account for the saturation selection) should detect $\sim 2\text{--}3$ orphans in excess of the conventional background. A pilot on ≥ 20 usable sightlines would provide initial evidence; the Poisson threshold $P(0|\lambda = 4) = 0.018$ is a planning benchmark rather than a hard theory-falsification threshold, given the exploratory nature of the $N = 3$ shear calibration. The KODIAQ Z survey (26 pLLSs/LLSs) and the HD-LLS survey (Prochaska et al. 2015; 46 additional LLSs) together provide ~ 72 LLS sightlines, yielding ~ 9 usable sightlines after the saturation selection — sufficient for a pilot but not for a definitive falsification. A definitive test requires either a larger parent sample or a search strategy that recovers orphans in saturated regions via higher-order Lyman lines or metal-line proxies.

Fixed-window mirror pilot. Seven LLS systems in six KODIAQ quasar spectra were compared over the illustrative interval $|\Delta v| \in [50, 140]$ km/s. The upper edge is an analysis-window choice, not a universal TEP transport ceiling: the existing activated scalar benchmark permits larger exterior clock contrasts. Candidates outside this window are outside this pilot selection, rather than physically excluded. The redward mirror provides a useful comparison for saturation and forest structure, but genuine deuterium also absorbs blueward near the isotope slot, so the true-D null is not mirror-symmetric. Across the seven sightlines the raw blue-wing depth excess is $+0.154$ (mean over sightlines); the largest single case is J162557+264448 ($+0.475$ depth, $+13$ narrow features, random-centre forest $p = 0.10$). Random centres in the same spectra give a mean forest asymmetry of $+0.02$ and per-sightline probabilities of $0.10\text{--}0.57$; two of the seven systems share the same quasar spectrum, so an independence-based Fisher combination is not quoted as a calibrated probability. These quantities describe the forest comparison, not an isotope-D-versus-TEP likelihood or a TEP exclusion. A population discriminant requires matched true-D and shifted-H injections with the same line selection, saturation, forest and instrumental treatment. The primary Gate 2/3 spectral comparisons remain the calibrated model tests.

5.8 Corpus Cross-References: The No-Hot-Big-Bang Framework

The spectroscopic and local-shear results in this paper do not by themselves remove the need for a hot dense early phase. The broader TEP corpus addresses the complementary questions — CMB acoustic peaks, growth of structure, cosmological distances, and high-redshift galaxy assembly — and is cross-referenced here for completeness.

- **CMB acoustic peaks and sound horizon (TEP-HC, Paper 18).** The TEP conformal scalar field has been implemented natively in *hi_class* ^[29] and run through MCMC against Planck 2018, BAO, and Pantheon+. The acoustic sound-horizon ratio $r_s^{\text{TEP}}/r_s^{\Lambda\text{CDM}} = 0.999994$ (< 6 ppm deviation), and the TT/TE/EE peak morphology is preserved without early spatial expansion. The Big Bang is reinterpreted as $A_{\text{clock}} \rightarrow 0$: a temporal horizon, not a geometric singularity.
- **Temporal horizon and absence of singularity (TEP-TH, Paper 27).** The FLRW singularity is shown to be a reconstruction artifact. The temporal-horizon conformal boundary $A_{\text{clock}}(\eta) = C\eta^{-p}$ with $0 < p \leq 1/2$ yields vanishing curvature invariants, divergent timelike proper time, and divergent null affine parameter — a regular, complete past boundary. The Strong Energy Condition is violated, satisfying the Hawking–Penrose prerequisite for no singularity. The native tensor benchmark gives $r(k_{\text{pivot}}) = 1.9 \times 10^{-9}$; the fiducial band maximum is 5.2×10^{-7} and the scanned curvature-regular envelope is 1.4×10^{-5} (Paper 27). These are different statistics and are not interchangeable.
- **Cosmological distances without expansion (TEP-C0, Paper 26).** Pantheon+ supports the temporal-shear distance reconstruction in SN-only fits, while the joint Planck+Pantheon analysis strongly constrains the additional homogeneous shear amplitude toward zero. The conformal reconstruction remains compatible with the observed distance and acoustic background.
- **JWST high-redshift anomalies (TEP-JWST, Paper 12).** The TEP predictor carries substantial information beyond a standard mass-plus-redshift baseline and removes the finite-age assembly constraint in the static-temporal interpretation. Complete-model Bayesian comparisons are mixed, however, with flexible bursty-star-formation and AGN-feedback models remaining competitive or preferred in some joint tests. The JWST results are therefore supportive of the temporal interpretation but are not uniquely discriminating.
- **Theoretical foundation (TEP, Paper 0).** The canonical preferred-foliation matter metric $\tilde{g}_{\mu\nu} = -A_{\text{loc}}^2(u_{\text{grav}}(\mathbf{x}))U_\mu U_\nu + A_{\text{sp}}^2(u_{\text{drift}}(t, \mathbf{x}))\Pi_{\mu\nu} + B(\phi)\Pi_\mu^\alpha \Pi_\nu^\beta \nabla_\alpha \phi \nabla_\beta \phi$ (written in isotropic conformal-disformal shorthand as $\tilde{g}_{\mu\nu} = A^2 g_{\mu\nu} + B \nabla_\mu \phi \nabla_\nu \phi$), the two-branch scalar action, field equations, PPN mapping, and the line-of-sight transport integral as the invariant signature of dynamical time are established in the foundational paper.

The helium-4 equilibrium is calculated within this paper. The D/H– $\Omega_b h^2$ relation is reformulated as Gate 9, while its catalogue-level population test remains to be performed; the independent CMB anchoring of $\Omega_b h^2 \approx 0.022$ is inherited from TEP-HC. The scope of the deuterium argument is stated explicitly: the present analysis establishes that the benchmark absorber identifications are non-unique, not that a deuterium production channel exists. In an eternal recycling cosmology this distinction is binding — astration is a

pure deuterium sink (deuterium burns at $T \gtrsim 5 \times 10^5$ K, below every stellar-core temperature), so the standing Galactic interstellar value $D/H \approx 1.5\text{--}2 \times 10^{-5}$, measured independently of the quasar-absorber architecture criticised here, is a production requirement rather than a relic. Galactic-cosmic-ray spallation ($\sim 10^{-7}$), pre-main-sequence envelope recycling, and pristine inflow each fail to supply it; a post-sequestration channel — horizon-adjacent spallation from the divergent external-time exposure of sequestered-core surfaces — remains contingent on the horizon-transport solution, but the channels quantified in Step 15 show that closure does not require it. The supply-versus-demand ledger is now quantified (Step 13, `results/gate13_deuterium_production_ledger.json`): holding the ISM at the observed abundance requires a yield of 2.0×10^{-5} D per unit astrated mass, while cosmic-ray spallation and ν -process helium-shell production supply $\sim 10^{-7}$ and $\sim 10^{-8}$ — shortfalls of $\sim 120\times$ and $\sim 2000\times$. For the horizon-adjacent channel the binding constraint is throughput rather than rate — external exposure is unbounded, so the emissivity per transit is not the limiter; the requirement is that a fraction f of astrated mass transit a spallation-active surface and return with deuterium enrichment ε_D satisfying $f \varepsilon_D = 2.0 \times 10^{-5}$ (e.g., $\varepsilon_D = 10^{-2}$ requires $f \approx 0.2\%$; enrichment only to the ISM level requires $f \approx 60\%$). The debt is thereby converted from an open-ended problem into a single throughput condition on the sequestered-core geometry. Step 15 (`results/gate15_deuterium_extended_channels.json`) then evaluates the channels that operate before or beside temporal-well formation, so that nothing needs to climb out of a formed well (the non-relativistic escape bound $Z_{\max} \sim 6.7 \times 10^3$ of TEP-BH is not invoked). (i) Partial-burning and explosive-freezeout ejecta from very massive objects — the classical Wagoner–Fowler–Hoyle channel — carry documented enrichments $X_D \sim 2 \times 10^{-5}\text{--}4 \times 10^{-4}$, requiring a processed fraction $f \approx 0.05\text{--}1.0$ of astrated mass; at a representative envelope-transit fraction $f = 0.5$ the supply exceeds demand by a factor ≈ 5 , and the ejection geometry is the same one the helium equilibrium already requires at $f_{\text{VMO}} = 0.994$. The Carr–Bond–Arnett metal co-production objection that historically removed this channel is answered by the architecture itself: the remnant is a temporal well that sequesters core metals while the processed envelope returns, not a standard event-horizon black hole. (ii) The collapse neutrino process (Fuller–Woosley–Weaver; Fuller & Shi 1997) reaches the same ejecta enrichment through $\bar{\nu}_e + p \rightarrow n + e^+$ irradiation of the envelope; its transport leg is fixed by the matter-metric axiom — MeV neutrinos are near-null on $\tilde{g}$ and irradiate the envelope during the collapse transient while the lapse is still of order unity. The irradiation budget is now parameterized (Step 17, `results/step_17_inversion_and_channel_budgets.json`): inverse-square antineutrino fluence through envelope shells gives capture probabilities y_n exceeding the 10^{-3} level required for ejecta $D/H \sim 10^{-4}$ only inside a radius $r_{\text{crit}} \simeq 0.002\text{--}0.22 R_*$ across the documented $10^{57}\text{--}10^{59}$ erg burst range — what fraction of the ejecta lies inside that radius is not determined by a static polytrope and awaits the dynamical temporal-well collapse solution; the transport itself is not a blocker, since MeV neutrinos are near-null on $\tilde{g}$ and the burst precedes deep-well formation. (iii) Bounded-exposure worldlines, which exist by the $p + q = 1.3 > 1$ astration-convergence condition, accumulate deuterium without a destruction term and close the ledger for $N_{\text{exp}} \gtrsim 10^3$ exposure equivalents at mixing fractions $f_{\text{mix}} \sim 1\text{--}12\%$, carrying the falsifiable signature of metal-poor reservoirs with super-ISM D/H — the same signature Fuller & Shi 1997 assigned to supermassive-star ejecta. The closure region is sharpened in Step 17 in net terms: returning gas at the ambient ISM abundance is not production, so at an envelope-transit fraction $f_{\text{ej}} = 0.5$ the ejecta must reach $X_D \geq 7.5 \times 10^{-5}$ ($D/H \sim 3.7 \times 10^{-5}$) if the feed arrives at the ambient value — inside the documented $10^{-5}\text{--}2 \times 10^{-4}$ window but requiring a VMO structure solution to establish either number. For the absorber side, Step 17 also records an honest limit: published $\tau_{\text{eff}}(z)$ cannot be inverted to mean neutral columns because $\tau_{\text{eff}} = -\ln\langle F \rangle$ is nonlinear in saturated regions, so the wing-density confrontation is withheld rather than claimed; the assumed LLS population implies a LyC reach of only $\Delta z \sim 0.6$ at $z \sim 3$, confirming that trough ionization cannot draw on external background light. The deuterium requirement is therefore reduced to a VMO ejecta-fraction/enrichment calculation and a reservoir longevity/mixing calculation, with a single shared observational discriminator. The Gate 4B amplitude channel has been bounded by the ambient-lapse self-consistency requirement (Gate 10, section 7d): the descent-tracking realisation that concords the inverse roots places the ambient boundary factor at $u(\bar{\varphi}) \approx 1 - 10^{-5}$ on the entire branch, under which the matter-metric redshift map cannot host the observed absorber population without a negative ambient field; the unscreened channel is thereby excluded at these epochs, and a temporal-disformal amplitude of the observed scale is admissible only under an environment-modulated coupling $B_{\text{eff}} = B(\phi) G(X_{\text{local}})$; Gate 10-b (`results/gate10b_disformal_gate_test.json`) derives the measurable condition — a gate closed in the ambient and opening with a common slope $dG/d\ln\delta\varphi \sim (4\text{--}7) \times 10^{-4}$ per e-fold across the three wells (agreement within a factor 1.6), saturating near $G \sim 10^{-3}$ so that deeper structures are not over-displaced — converting the amplitude question into a well-posed target for the disformal-sector derivation rather than an open loophole. The primary Gates 2/3 spectroscopic result is in any case independent of the Gate-4B amplitude closure. Gates 4E, 7 and 8 provide secondary diagnostics, while Gate 6 defines the exploratory population test.

6. Conclusion

Across three benchmark D/H absorption systems, unrestricted H I is significantly preferred over isotope-tied D I within the tested Voigt architecture, including under adversarial parent reassignment and the primary multiplicity correction. Q1009+2956 and PKS 1937–101 also show significant departures from exact isotope locking in the calibrated Gate-8 analysis, while J1332+0052 is less discriminating. The result therefore demonstrates that the deuterium interpretation is not spectroscopically unique within the tested model class; the isotope-tied model is strongly disfavoured relative to unrestricted H within the tested architecture.

The fixed -81.6 km s^{-1} clustering is not an independent discriminator in a sample selected through the deuterium search window. The corresponding falsifiable population test is an excess of suitable displaced H I components outside that window. The archival pilot establishes the practical selection limits of this test and identifies one candidate feature in Q1009+2956, but a substantially larger usable sample is required.

TEP provides a candidate physical interpretation of this displaced-H phenomenology. The conformal absorber field gives the blueward sign but only a $0.005\text{--}0.01\text{ km s}^{-1}$ shift. A temporal-disformal channel can reach the observed scale in the explored parametrisation. The full covariant evaluation (Gate 10) shows that the earlier closure obstruction is a property of the imported Λ CDM drift rate rather than of the disformal channel: joint sign-and-closure compatibility requires a descent-type ambient drift ($d \ln H_T^2 / d \ln(1+z) < 0$), which the unsourced free-roll on the reconstructed corpus potential supplies, and under that rate a descent-tracking coupling $B(\varphi) = B_0 \varphi^4 / (\varphi_c^4 + \varphi^4) e^{-\varphi}$ passes Cassini and GW170817, with inverse roots agreeing to a factor 1.00004. Forward evaluation at that common B_0 does not reproduce the three amplitudes, and ambient-lapse self-consistency excludes the branch outright: the flat- u realisation holds the ambient sector at near-criticality, where the matter-metric redshift map would require a sign-violating negative ambient field. The unscreened temporal-disformal channel is accordingly bounded sub-critical at these epochs. Independently, the baryonic-cycling calculation shows that a helium-4 equilibrium near the observed abundance is possible under the adopted VMO-dominated, metal-sequestering branch. The spectroscopic, population and light-element predictions therefore provide separate empirical routes for testing the proposed non-hot-BBN interpretation. The scope of the deuterium argument is bounded accordingly: the analysis establishes that the benchmark absorber identifications are non-unique, not that a deuterium production channel exists. In an eternal recycling cosmology the distinction is binding — astration is a pure deuterium sink, so the standing Galactic interstellar value $D/H \approx 1.5\text{--}2 \times 10^{-5}$, measured independently of the quasar-absorber architecture criticised here, is a production requirement rather than a relic; galactic-cosmic-ray spallation, pre-main-sequence envelope recycling, and pristine inflow each fail to supply it. The extended ledger (Step 15, results/gate15_deuterium_extended_channels.json) quantifies three channels that operate before or beside temporal-well formation — VMO partial-burning/freezeout ejecta, the collapse neutrino burst irradiating the envelope through near-null neutrinos on the matter metric, and protected-reservoir accumulation on bounded-exposure worldlines — of which the first and third can close the required net yield within documented ranges — at a transit fraction $f_{\text{ej}} = 0.5$ the ejecta must reach $X_D \geq 7.5 \times 10^{-5}$ if the feed carries ambient ISM deuterium — reducing the open item to a VMO ejecta-enrichment calculation and a reservoir mixing rate rather than a missing mechanism. Independently, the opacity three-point requirement is shown to admit an explicit species-resolved profile (Step 14, results/step_14_opacity_three_point.json): the wing datum fixes $n_e \approx 0.51 n_{b,0}$, the intervening trough is derived to be neutral ($\bar{n}_e < 3.7 \times 10^{-4} n_{b,0}$), and the ionization front at the Gate-11B thermalization depth places the $\tau = 1$ surface at $z \approx 1080$. The boundary's self-consistency items are then closed at toy-model level (Step 18, results/step_18_boundary_selfconsistency.json): the $\sim 4 \times 10^{10}$ wing-to-trough ionization-rate contrast is supplied by self-shielding over ~ 24 LyC e-folds; the front gas temperature is derived ($T_{\text{gas}} = T_{\text{bath}} = 2946\text{ K}$ via Compton equilibrium on a $\sim\text{yr}$ coupling time), making the Saha ionization a consequence of the Gate-11B bath rather than an assumption; the visibility tail is shown to be a closure artifact removed when the edge is absorptive; and the deuterium channels' residual budgets are bounded (transit neutrino irradiation under a stated homologous-collapse model, reservoir longevity reduced to a $\sim 10^3\text{--}10^4$ Gyr sequestration requirement). The remaining open item in the opacity sector is the perturbation mode spectrum on the boundary — the eternal-native Boltzmann treatment — together with the absorber-statistics inversion on a full LLS/Lyman-forest catalogue, which is the identified pipeline extension. The cosmological stability evidence must be read by sector. The homogeneous closed-background benchmark has an Eddington-type secular mode. The explicit Euclidean zero-mean scalar projection in Paper 0 (step_79) gives a positive lowest restricted eigenvalue, $+38.8145$ at $N = 24$ and $+38.9129$ at $N = 32$; its negative direction is almost spatially constant, with relative variation 2.2×10^{-5} . The discrete periodic Poincaré bound independently certifies all zero-mean modes above 38.80 and 38.90 on these grids. This is substantive stability evidence for the computed reduced sector. Identifying the constant direction as a gauge mode and establishing the full coupled constraint-reduced spectrum require the corresponding action-level analysis. The chemical attractors evaluated here therefore coexist with a tested stable nonconstant scalar sector; a global stability theorem is not inferred from either chemical convergence or that reduced spectrum.

Data Availability & Reproducibility

This work follows open-science practices. All results are fully reproducible from raw data using the documented pipeline. All numerical results, Monte Carlo simulations, and statistics are generated by deterministic Python scripts processing real observational data. The pipeline enforces rigorous reproducibility: prespecified statistical outcomes are reported without post-hoc alteration and interpreted according to the scope of each test.

Repository and Code

GitHub Repository: github.com/matthewsmawfield/TEP-BBN

The repository contains a deterministic, version-controlled cosmological analysis pipeline utilizing 27 analysis steps across 4 phases: core gates (1–8), Gate 8 calibration (step_08b, true-D injection Monte Carlo), Gate 8 wavelength- and parent-centroid-systematics sensitivity (step_08c, step_08d), disformal transport, the steady-state and ionization-equilibrium ledgers (step_11, step_12), the deuterium production and extended-channel ledgers (step_13, step_15), the species-resolved opacity three-point test (step_14), the boundary-transport consistency audit (step_16), the literature-input inversion and channel-budget audit (step_17), the boundary self-consistency solve (step_18), the eternal-surface reconciliation and self-shielded ambient-bath front mode spectrum (step_19, step_20), the canonical screening-projection audit (step_21), secondary robustness analyses, and manuscript finalization. All steps are

orchestrated by `scripts/run_pipeline.py` with comprehensive per-step logging, SHA-256 checksum manifest generation, and a pipeline summary ledger.

All raw spectroscopic data (Keck/HIRES, VLT/UVES + Keck/HIRES, and VLT/ESPRESSO), structural likelihood matrices, and the temporal-field equation solvers are released in the Zenodo repository (DOI: 10.5281/zenodo.21841147) under CC-BY 4.0. The full codebase and execution environments are identical to the published version.

Repository Structure

```
TEP-BBN/
├── data/
│   ├── raw/                # Spectroscopic exposures (Keck/HIRES, UVES+HIRES,
│   │   │   │               ESPRESSO)
│   │   ├── atomic/         # Physical H I, D I line registries (NIST ASD)
│   │   │   ├── reduced_products/ # Pre-reduced and co-added normalized spectra
│   │   │   ├── literature_components/ # Published VPFIT model files and component tables
│   │   │   └── processed/     # Pipeline-ready union manifests
│   ├── scripts/
│   │   ├── steps/          # pipeline steps (01-17 + lettered substeps,
│   │   │   │               secondary, finalize)
│   │   ├── lib/            # Physical RT engine, Voigt fitters, model parsers
│   │   ├── utils/          # Logging, isotopic shift, shear model utilities
│   │   └── run_pipeline.py  # Master orchestration script (4-phase pipeline)
│   ├── configs/
│   │   ├── sightlines/     # Per-sightline JSON configs (bounds, MC settings,
│   │   │   │               noise models)
│   │   └── TEP_BBN_FOUNDATION_FREEZE.json # Frozen TEP prior parameters
│   ├── results/            # Generated parameter ledgers and significance
│   └── matrices
├── logs/                  # Per-step execution logs
├── site/
│   └── components/        # Manuscript source components
├── requirements-lock.txt   # Locked Python dependencies
└── README.md              # Documentation
```

Data Provenance

Data Source	Provider	Access Method	Records	Location
Q1009+2956 Spectra	Keck/HIRES (Zavarygin et al. 2018 ^[44])	Zavarygin et al. 2018 GitHub	4 coadds (S/N $\approx 47\text{--}78$ near Ly α)	<code>data/raw/reduced_products/Q1009+2956_z2.504_HIRES/</code>
PKS 1937–101 Spectra	VLT/ESPRESSO (Cooke et al. 2024)	ESO Archive	3 exposures	<code>data/raw/spectra/PKS1937-101_z3.572_ESPRESSO/</code>
J1332+0052 Spectra	VLT/UVES + Keck/HIRES (Kislitsyn et al. 2024)	ESO Archive	24 exposures	<code>data/raw/spectra/J1332+0052_z3.420_UVES/</code>
SQUAD DR1 Spectra	UVES/SQUAD (Murphy et al. 2018 ^[43])	ESO Phase 3 Archive	3 sightlines (Q1009, Q0913, J212912)	<code>data/raw/spectra/SQUAD/</code>
KODIAQ DR2 Spectra	Keck/HIRES (O'Meara et al. 2017)	Keck KOA	7 LLS-class sightlines	<code>data/raw/spectra/KODIAQ/</code>
KODIAQ Z LLS Catalog	Lehner et al. 2016	Literature	26 pLLSs/LLSs at	<code>data/kodiaq_z_lls_catalog.json</code>

Data Source	Provider	Access Method	Records	Location
			$2.3 < z < 3.3$	
Atomic Data	NIST ASD [49]	Static Registry	H I, D I, metals	data/raw/atomic/
Q1009+2956 Model	Zavarygin et al. [44] / VPFIT [50]	Static File	43 H I + 3 D I components	data/literature_components/model_6a.26
PKS 1937–101 Model	Cooke et al. 2024 / VPFIT [50]	Static File	184 H I + 3 D I components	data/literature_components/PKS1937-101_z3.572_vpfit_model.26
J1332+0052 Model	Kislitsyn et al. 2024 / VPFIT [50]	Static File	11 H I + 4 D I components	data/literature_components/J1332+0052_z3.420_vpfit_model.26
Prior Bounds	Derived	Static File	All variables	configs/sightlines/*.json

Pipeline Architecture

The analysis pipeline comprises 4 phases and 24 steps spanning spectroscopic ingestion to thermodynamic evaluation, helium synthesis, opacity boundary analysis, Gate 8 true-D injection calibration, wavelength- and parent-centroid-systematics sensitivity, disformal transport calibration, the steady-state timescale/thermalization ledger (Gate 11), the VMO ionization-equilibrium budget (Gate 12), the deuterium production ledger (Gate 13), the species-resolved opacity three-point test (Step 14), the extended deuterium channels ledger (Step 15), the boundary-transport consistency audit (Step 16), the literature-input inversion and channel-budget audit (Step 17), the boundary self-consistency solve (Step 18), secondary robustness checks, and manuscript number finalization. Each step is a standalone Python script in `scripts/steps/` that produces serialized JSON outputs and detailed logs with SHA-256 checksum verification.

Complete Step Inventory and Runtime

Runtimes are approximate and measured on Apple M4 Pro (14-core, 24 GB). The dominant cost is the Monte Carlo significance test (step 03), which scales with the number of realizations multiplied by the parent reassignment refit performed inside each realization.

Phase	Step	Script	Description	Est. Runtime
1	01	<code>step_01_embedding.py</code>	Physical atomic-data embedding (H vs D isotope identifiability)	~1 s
1	02	<code>step_02_q1009.py</code>	Nested-hypothesis fit (multi-start; nesting invariants verified)	~4 min
1	03	<code>step_03_significance.py</code>	1000-realization Monte Carlo significance calibration (parent reassignment inside each realization)	~25 min
1	04	<code>step_04_prior.py</code>	TEP absorber field sign closure ($\Delta v_T < 0$)	~1 s
1	04B	<code>step_04b_amplitude.py</code>	TEP absorber amplitude test (conformal/static branch sign theorem)	~1 s
1	05	<code>step_05_thermodynamics.py</code>	Matter-frame temporal thermodynamics (Planck spectrum preservation)	~1 s
1	06	<code>step_06_helium.py</code>	Primordial helium synthesis via baryonic cycling	~1 s
1	07	<code>step_07_global_opacity.py</code>	Analytical proof of divergent optical depth in the static spatial geometry	~1 s
1	08	<code>step_08_light_elements.py</code>	Light-element network (3He, 4He, 7Li) and convergence constraints	~1 s
1	08B	<code>step_08b_calibration.py</code>	Gate 8 true-D injection Monte Carlo calibration ($N = 1000$ per sightline)	~45 min
1	08C	<code>step_08c_wavelength_systematics.py</code>	Gate 8 wavelength-systematics sensitivity (literature-calibrated σ_{sys} per sightline, $N = 100\,000$ MC)	~1 s
1	08D	<code>step_08d_parent_centroid_systematics.py</code>	Gate 8 parent-centroid blend-bias sensitivity (parent velocity perturbed inside the MC loop, blend-alone and joint sign-matched bounds, $N = 100\,000$ MC)	~1 s

Phase	Step	Script	Description	Est. Runtime
2	09	<code>step_09_disformal_transport.py</code>	Disformal transport solver and absorber-ansatz / Cassini bound bookkeeping	~1 s
2	4B-IV	<code>step_04h_disformal_clock.py</code>	Disformal clock mechanism: cosmological $\partial_0 \bar{\phi}$ channel, $B(\phi)$ field-dependence scan, Cassini/GW170817 constraints, three-sightline test	~1 s
2	09B	<code>step_09b_screening_plot.py</code>	Screening function visualization and JSON ledger	~1 s
2	10	<code>step_10_covariant_transport.py</code>	Full covariant transport solver ($\omega = -u^\mu k_\mu$) and ambient-lapse self-consistency	~1 s
2	11	<code>step_11_steady_state_budget.py</code>	Gate 11 steady-state timescale and thermalization-depth budget	~1 s
2	12	<code>step_12_vmo_ionization_budget.py</code>	Gate 12 VMO-branch ionization-equilibrium budget over the (ξ, f_{esc}) plane	~1 s
2	13	<code>step_13_deuterium_production_ledger.py</code>	Gate 13 deuterium supply-versus-demand ledger ($f \varepsilon_D = 2.04 \times 10^{-5}$)	~1 s
2	14	<code>step_14_opacity_three_point.py</code>	Species-resolved opacity three-point test; admissible profile and KODIAQ LLS pilot inversion	~2 s
2	15	<code>step_15_deuterium_extended_channels.py</code>	Extended deuterium channels: VMO ejecta, collapse neutrino process, protected reservoirs	~1 s
2	16	<code>step_16_boundary_transport.py</code>	Boundary-transport consistency audit on the Step-14 profile: visibility function, geometric diffusion analogue, conditional Saha ionization, recombination/emissivity ledger, Γ -equilibrium requirements	~5 s
2	17	<code>step_17_inversion_and_channel_budgets.py</code>	Literature-input absorber audit and deuterium channel budgets: net-yield ejecta closure, neutrino capture-radius thresholds, reservoir mixing bookkeeping	~1 s
2	18	<code>step_18_boundary_selfconsistency.py</code>	Boundary self-consistency solve: coupled ionization front (self-shielded Γ transition), Compton-equilibrium front temperature, absorptive-edge visibility kernel, homologous-collapse neutrino transit exposure, reservoir longevity timescale	~5 s
2	19	<code>step_19_eternal_surface_reconciliation.py</code>	Eternal-vs-mapped-FLRW surface reconciliation: path-measure Jacobian, explicit diffusion-integral damping scale, TEP-TH kernel comparison	~1 s
2	20	<code>step_20_ambient_bath_front.py</code>	Ambient-bath ionization front & closure: Planck LyC Wien-tail supply, self-shielded LyC radiative transfer, absorber-statistics density inversion ($\Delta_{\text{front}} = 668.9 n_{b,0}$ for $z_{\tau=1} = 1080.0$), and boundary perturbation mode spectrum ($\ell_D = 1521$; action-derived two-branch ambient clock-slope mode with exact <code>hi_class</code> Boltzmann closure and three-branch ablation)	~1 min
2	6-III	<code>step_06_orphan_search.py</code>	Orphan feature search in VPFIT models, direct spectral inspection, and archival SQUAD/KODIAQ search	~10 s
3	C2-C4	<code>secondary_analyses.py</code>	Scaffold freedom (C2), perturbed MC (C3), kinematic test (C4)	~10 min
4	Final	<code>finalize_manuscript_numbers.py</code>	Manuscript number synchronization + SHA-256 checksum manifest	~1 s

Total Runtime Summary

Component	Steps	Runtime
Core Gates (Phase 1)	9	~45 min per sightline
Disformal + Secondary (Phases 2–3)	14	~10 min per sightline
Finalization (Phase 4)	1	~1 s
Total (3 sightlines)	24	~2 hr

Reproduction Instructions

Quick Start (Full Reproduction)

```
# 1. Clone repository
git clone https://github.com/matthewsmawfield/TEP-BBN.git
cd TEP-BBN

# 2. Install dependencies
pip install -r requirements-lock.txt

# 3. Run full pipeline (default: Q1009+2956 sightline, 1000 MC sims)
python scripts/run_pipeline.py

# 3b. Quick mode (50 MC sims, for faster testing)
python scripts/run_pipeline.py --quick

# 3c. Run a specific sightline
python scripts/run_pipeline.py --sightline PKS1937-101

# 3d. Run all sightlines with ingested data
python scripts/run_pipeline.py --all-sightlines

# 3e. Skip secondary robustness analyses
python scripts/run_pipeline.py --skip-secondary

# 3f. Custom MC simulation count
python scripts/run_pipeline.py --n-sims 500

# 4. Results will be stored in results/ and logs/
# Pipeline summary: results/pipeline_summary.json
# SHA-256 manifest: results/checksums_sha256.json
```

Multi-Sightline Configuration

The pipeline supports multiple D/H absorber sightlines through per-sightline JSON configuration files in `configs/sightlines/`. Each config specifies the data manifest path, VPFIT model file, noise model, absorber redshift, candidate parameter bounds, multi-start initial points, and Monte Carlo settings. Three sightlines are currently supported with ingested data: Q1009+2956 (Keck/HIRES, $z = 2.5042$), PKS 1937–101 (ESPRESSO, $z = 3.572$), and J1332+0052 (UVES+HIRES, $z = 3.421$). Sightlines without ingested data (PKS 1937–1009, HS 0105+1619) are automatically skipped with a diagnostic message.

System Requirements

Component	Minimum	Recommended	Tested On
CPU	2 cores	4+ cores	Apple M4 Pro (14-core)
RAM	4 GB	8 GB	24 GB
Storage	1 GB	2 GB	SSD NVMe
OS	Linux/macOS	Linux/macOS	macOS Sequoia 15.1

Current statistical implementation

The fixed data-quality mask is shared by all hypotheses and simulations. All finite residuals contribute to the Student- t likelihood; numerical forward-model or continuum failures are rejected. Continuum order is 10 for Lyman-series blocks and 5 for metal-line blocks. The per-sightline noise models, candidate bounds, multistart searches and parent-selection rules are recorded with the fitted ledgers. Both canonical-parent and parent-maximized likelihood comparisons are calibrated with $N = 1000$ genuine true-D injections per sightline.

Gate 8 retains every simulation vector and its seed. Marginal consistency tests center on the empirical null means. The joint metric regularizes a dimensionless correlation matrix and uses leave-one-out null scores. Input and simulation-file hashes permit provenance checks. Gaussian systematic sensitivity calculations are distinguished from empirical injection probabilities and their finite-sampling limits. The final method does not combine these correlated diagnostics as independent evidence.

Ionization balance. The opacity calculations use the stable positive root of $x_e^2 = q(1 - x_e)$. At the declared baryon density, photoionization equilibrium is $\Gamma(1 - x_e) = \alpha_B n_b x_e^2$. The step-16 LTE-equivalent rates therefore agree at fixed temperature across its density ladder; these are conditional equilibrium requirements, not a measured radiation field.

Likelihood and tail evaluation. All finite data residuals contribute to the likelihood. Numerical failures are rejected. Gaussian sensitivity tails use the survival function; Monte-Carlo outputs include exceedance counts, sampling resolution and one-sided 95% binomial upper bounds.

References

1. Smawfield, M.L. Temporal Equivalence Principle: Dynamic Time & Emergent Light Speed. *Zenodo* (2025), v0.16 (Jakarta). DOI: 10.5281/zenodo.16921911
2. Smawfield, M.L. Temporal Equivalence Principle: A Covariant Alternative to Cosmic Expansion. *Zenodo* (2026). DOI: 10.5281/zenodo.20370143
3. Smawfield, M.L. Temporal Equivalence Principle: Native hi_class Conformal Implementation, Linear Perturbation Closure, and CMB Acoustic Peak Preservation. *Zenodo* (2026). DOI: 10.5281/zenodo.20572722
4. Smawfield, M.L. Temporal Equivalence Principle: Temporal Horizon Cosmology and the Absence of a Physical Big Bang Singularity. *Zenodo* (2026). DOI: 10.5281/zenodo.20723059
5. Hawking, S.W. The occurrence of singularities in cosmology. *Proc. R. Soc. A* **294**, 511-521 (1966).
6. Hawking, S.W. & Penrose, R. The singularities of gravitational collapse and cosmology. *Proc. R. Soc. A* **314**, 529-548 (1970).
7. Borde, A., Guth, A.H. & Vilenkin, A. Inflationary spacetimes are incomplete in past directions. *Phys. Rev. Lett.* **90**, 151301 (2003).
8. Brandenberger, R. & Peter, P. Bouncing cosmologies: progress and problems. *Found. Phys.* **47**, 797-850 (2017).
9. Novello, M. & Bergliaffa, S.E.P. Bouncing cosmologies. *Phys. Rep.* **463**, 127-213 (2008).
10. Ijjas, A. & Steinhardt, P.J. Entropy, black holes and the new cyclic universe. *Phys. Lett. B* **824**, 136823 (2022).
11. Peebles, P.J.E. *Principles of Physical Cosmology*. Princeton University Press (1993).
12. Weinberg, S. *Cosmology*. Oxford University Press (2008).
13. Dodelson, S. *Modern Cosmology*. Academic Press (2003).
14. Mukhanov, V.F., Feldman, H.A. & Brandenberger, R.H. Theory of cosmological perturbations. *Phys. Rep.* **215**, 203-333 (1992).
15. Liddle, A.R. & Lyth, D.H. *Cosmological Inflation and Large-Scale Structure*. Cambridge University Press (2000).
16. Planck Collaboration, et al. Planck 2018 results. VI. Cosmological parameters. *A&A* **641**, A6 (2020).
17. Riess, A.G., et al. Milky Way Cepheid Standards for Measuring Cosmic Distances and Application to Gaia DR2: Implications for the Hubble Constant. *ApJ* **861**, 126 (2018).
18. Brout, D., et al. The Pantheon+ Analysis: Cosmological Constraints. *ApJ* **938**, 110 (2022).
19. Fixsen, D.J., et al. The Cosmic Microwave Background Spectrum from the Full COBE FIRAS Data Set. *ApJ* **473**, 576 (1996).
20. Chluba, J. & Sunyaev, R.A. The evolution of CMB spectral distortions in the early Universe. *MNRAS* **419**, 1294-1314 (2012).
21. PARTICLE DATA GROUP. Review of Particle Physics. *PTEP* **2022**, 083C01 (2022).
22. Cyburt, R.H., Fields, B.D., Olive, K.A. & Yeh, T.H. Big bang nucleosynthesis: Present status. *Rev. Mod. Phys.* **88**, 015004 (2016).
23. Seager, S., Sasselov, D.D. & Scott, D. A new calculation of the recombination epoch. *ApJ* **523**, L1-L5 (1999).
24. Peebles, P.J.E. Recombination of the Primeval Plasma. *ApJ* **153**, 1 (1968).
25. Zeldovich, Y.B. & Sunyaev, R.A. The interaction of matter and radiation in a hot-model universe. *Astrophys. Space Sci.* **4**, 301-316 (1969).
26. Seljak, U. & Zaldarriaga, M. A Line of Sight Integration Approach to Cosmic Microwave Background Anisotropies. *ApJ* **469**, 437 (1996).
27. Lewis, A., Challinor, A., & Lasenby, A. Efficient Computation of CMB Anisotropies in Closed FRW Models. *ApJ* **538**, 473 (2000).
28. Lesgourgues, J. & Tram, T. The Cosmic Linear Anisotropy Solving System (CLASS). Part IV: efficient implementation of non-cold relics. *JCAP* **09**, 032 (2011).
29. Zumalacárregui, M., Bellini, E., Sawicki, I., Lesgourgues, J. & Ferreira, P.G. hi_class: Horndeski in the Cosmic Linear Anisotropy Solving System. *JCAP* **08**, 019 (2017).
30. De Felice, A. & Tsujikawa, S. f(R) Theories. *Living Rev. Rel.* **13**, 3 (2010).
31. Wetterich, C. Cosmology and the fate of dilatation symmetry. *Nucl. Phys. B* **302**, 668-696 (1988).
32. Wetterich, C. A universe without expansion. *Phys. Dark Universe* **2**, 184 (2013).
33. Narlikar, J.V. & Arp, H.C. Flat spacetime cosmology: A unified framework for extragalactic redshifts. *Astrophys. J.* **405**, 51-56 (1993).
34. Mannheim, P.D. Conformal gravity and the nature of dark matter. *Prog. Part. Nucl. Phys.* **94**, 217-272 (2017).
35. Khoury, J. & Weltman, A. Chameleon cosmology. *Phys. Rev. D* **69**, 044026 (2004).
36. Hinterbichler, K. & Khoury, J. Symmetron cosmology. *Phys. Rev. Lett.* **104**, 231301 (2010).
37. Penrose, R. Before the Big Bang: an outrageous new perspective and its implications for particle physics. *Proc. EPAC* (2006).
38. Tod, K.P. Isotropic cosmological singularities. *Gen. Relativ. Gravit.* **35**, 779-805 (2003).
39. Tod, K.P. The equations of conformal cyclic cosmology. *Gen. Relativ. Gravit.* **47**, 31 (2015).
40. Ratra, B. & Peebles, P.J.E. Cosmological Consequences of a Rolling Homogeneous Scalar Field. *Phys. Rev. D* **37**, 3406 (1988).

41. Caldwell, R.R., Dave, R., & Steinhardt, P.J. Cosmological Imprint of an Energy Component with General Equation of State. *Phys. Rev. Lett.* **80**, 1582 (1998).
42. Clifton, T., Ferreira, P.G., Padilla, A. & Skordis, C. Modified gravity and cosmology. *Phys. Rep.* **513**, 1-189 (2012).
43. Murphy, M.T., Kacprzak, G.G., Savorgnan, G.A.D. & Carswell, R.F. The UVES Spectral Quasar Absorption Database (SQUAD) data release 1: the first 10 million seconds. *MNRAS* **482**, 3458-3482 (2019). arXiv:1810.06136
44. Zavarygin, E., Webb, J.K., Dumont, V. & Riemer-Sørensen, S. The primordial deuterium abundance at $z_{\text{abs}} = 2.504$ from a high signal-to-noise spectrum of Q1009+2956. *MNRAS* **477**, 5536-5553 (2018). arXiv:1706.09512
45. Cooke, R.J., Pettini, M., Jorgenson, R.A., Murphy, M.T. & Steidel, C.C. Precision measures of the primordial abundance of deuterium. *ApJ* **781**, 31 (2014).
46. Burles, S. & Tytler, D. The Deuterium Abundance toward QSO 1009+2956. *ApJ* **507**, 732 (1998).
47. O'Meara, J.M., Burles, S., Prochaska, J.X., Prochter, G.E., Bernstein, R.A. & Burgess, K.M. The Deuterium-to-Hydrogen Abundance Ratio toward the QSO SDSS J155810.16-003120.0. *ApJ* **649**, L61 (2006).
48. Kirkman, D., Tytler, D., Burles, S., Lubin, D. & O'Meara, J.M. On the primordial deuterium abundance. *ApJ* **529**, 655 (2000).
49. Kramida, A., Ralchenko, Yu., Reader, J. & NIST ASD Team. NIST Atomic Spectra Database (ver. 5.11). National Institute of Standards and Technology, Gaithersburg, MD (2023).
50. Carswell, R.F. & Webb, J.K. VPFIT: Voigt profile fitting program. Astrophysics Source Code Library, ascl:1408.015 (2014).
51. Smawfield, M.L. The Cepheid Bias: Resolving the Hubble Tension. *Zenodo* (2026). DOI: 10.5281/zenodo.18209702
52. Smawfield, M.L. Temporal Equivalence Principle: A Unified Resolution to the JWST High-Redshift Anomalies. *Zenodo* (2026). DOI: 10.5281/zenodo.19000827
53. Smawfield, M.L. Temporal Equivalence Principle: Temporal Shear in the Earth Flyby Anomaly. *Zenodo* (2026). DOI: 10.5281/zenodo.19454862
54. Smawfield, M.L. Temporal Equivalence Principle: Black Holes and the Temporal Horizon. *Zenodo* (2026). DOI: 10.5281/zenodo.21677826
55. Haardt, F. & Madau, P. Radiative transfer in a clumpy universe. IV. New synthesis models of the UV/X-ray cosmic background. *ApJ* **746**, 125 (2012).
56. Benisty, D., Brax, P. & Davis, A.C. Constraining modified gravity with cosmological data. *Phys. Rev. D* **107**, 064049 (2023).
57. Carr, B.J., Bond, J.R. & Arnett, W.D. Cosmological consequences of Population III stars. *ApJ* **277**, 445-469 (1984).
58. Bond, J.R., Arnett, W.D. & Carr, B.J. The evolution and fate of Very Massive Objects. *ApJ* **280**, 825-847 (1984).
59. Bertotti, B., Iess, L. & Tortora, P. A test of general relativity using radio links with the Cassini spacecraft. *Nature* **425**, 374-376 (2003).
60. Ober, W.W., El Eid, M.F. & Fricke, K.J. Evolution of massive Pop III stars with helium-overshooting. *A&A* **119**, 61-68 (1983).
61. Aver, E., Olive, K.A. & Skillman, E.D. The effects of He I $\lambda 10830$ on helium abundance determinations. *JCAP* **07**, 011 (2011).
62. Izotov, Y.I., Thuan, T.X. & Guseva, N.G. A new determination of the primordial He-4 abundance: A self-consistent reanalysis of the H II region data. *MNRAS* **445**, 778-801 (2014).
63. Peimbert, M., Peimbert, A. & Luridiana, V. The primordial helium abundance and the number of neutrino species. *Rev. Mex. Astron. Astrofis.* **52**, 319-326 (2016).
64. Planck Collaboration, et al. Planck 2018 results. VI. Cosmological parameters. *A&A* **641**, A6 (2020). (See also: Cooke, R.J., Pettini, M. & Steidel, C.C. The primordial deuterium abundance and the baryon density. *ApJ* **817**, 79 (2016) for the $D/H \rightarrow \Omega_b h^2$ concordance.)
65. Cooke, R.J., Pettini, M. & Steidel, C.C. The primordial deuterium abundance and the baryon density. *ApJ* **817**, 79 (2016).
66. Cooke, R.J., et al. ESPRESSO observations of the D/H ratio in PKS 1937-101. *MNRAS*, stae452 (2024). arXiv:2402.05586
67. Kislitsyn, R., Balashev, S.A., Murphy, M.T., et al. A new precise determination of the primordial abundance of deuterium: measurement in the metal-poor sub-DLA system at $z = 3.42$ towards quasar J 1332+0052. *MNRAS* **528**, 4068-4084 (2024). arXiv:2401.12797
68. Fields, B.D., Olive, K.A., Yeh, T.H. & Cyburt, R.H. The Hubble tension and the primordial helium abundance. *Universe* **9**, 165 (2023).
69. Smawfield, M.L. Temporal Topology Saturation Scale: Cross-Scale Consistency of ρ_T . *Zenodo* (2026). DOI: 10.5281/zenodo.18064365
70. Smawfield, M.L. Temporal Equivalence Principle: Lunar Laser Ranging and the Nordtvedt Effect. *Zenodo* (2026). DOI: 10.5281/zenodo.19446028
71. Griest, K., Whitmore, J.B., Wolfe, A.M., Prochaska, J.X., Dawley, M.M. & Vladilo, G. Wavelength Accuracy of the Keck HIRES Spectrograph and Measuring Changes in the Fine Structure Constant. *ApJ* **708**, 158-170 (2010). DOI: 10.1088/0004-637X/708/1/158
72. Whitmore, J.B. & Murphy, M.T. Impact of instrumental systematic errors on fine-structure constant measurements with quasar spectra. *MNRAS* **444**, 2 (2014). DOI: 10.1093/mnras/stu2420. arXiv:1409.4467
73. Schmidt, T.M., Murphy, M.T., Lo Curto, G., et al. Fundamental physics with ESPRESSO: Towards an accurate wavelength calibration for a precision test of the fine-structure constant. *A&A* **646**, A104 (2021). DOI: 10.1051/0004-6361/202039345
74. Korn, A.J., et al. A probable stellar solution to the cosmological lithium discrepancy. *Nature* **442**, 657-659 (2006).
75. Richard, O., Michaud, G. & Richer, J. Implications of WMAP observations on Li abundance and stellar evolution. *ApJ* **619**, 538-548 (2005).
76. Fields, B.D. The primordial lithium problem. *Annu. Rev. Nucl. Part. Sci.* **61**, 47-68 (2011).
77. Schaerer, D. On the properties of massive Population III stars and metal-free stellar populations. *A&A* **382**, 28-42 (2002).
78. Tumlinson, J. & Shull, J.M. Zero-metallicity stars and the effects of the first stars on reionization. *ApJ* **528**, L65-L68 (2000).

79. Osterbrock, D.E. & Ferland, G.J. *Astrophysics of Gaseous Nebulae and Active Galactic Nuclei*, 2nd ed. (University Science Books, 2006).
80. Mukohyama, S. Dark matter as integration constant in Hořava-Lifshitz gravity. *Phys. Rev. D* **80**, 064005 (2009).
81. Chamseddine, A.H. & Mukhanov, V. Mimetic Dark Matter. *JHEP* **11**, 135 (2013).
82. Berezhiani, L. & Khoury, J. Theory of dark matter superfluidity. *Phys. Rev. D* **92**, 103510 (2015).