Isotonic-Thawing Quintessence with Canonical Perturbations: Transient Acceleration, Native CAMB Spectra, Growth, and a DESI DR2–Pantheon+–CMBComp Posterior

Chris Clifton*
Independent Researcher, Delaware, USA
(Dated: 1 September 2026)

We test whether a restrictive canonical, non-phantom dark-energy model can remain observationally viable while containing its own dynamical exit from cosmic acceleration. For the dimensionless scalar coordinate $\Phi = \phi_{\rm phys}/M_{\rm Pl}$, the potential

$$V(\Phi) = \Lambda_0 \left(\frac{\Phi}{2} - \frac{2}{\Phi}\right)^2, \qquad \Phi > 0,$$

is non-negative and has a unique regular zero-energy minimum at $\Phi = 2$. This coefficient locking removes any asymptotic cosmological-constant residue and makes the model's future fate part of the same dynamics tested at the present epoch.

The model is parameterized by the bounded displacement coordinate $\eta = 2/\Phi_i$, with the flat-ΛCDM limit nested at $\eta = 0$. For every finite $\eta > 0$ on the prepared branch, the same coordinate that fixes the observable thawing history also fixes a finite future acceleration-exit time; at every sampled point, the potential amplitude is solved by spatial flatness. A deterministic posterior using DESI DR2 BAO, an 11-knot Pantheon+ distance compression, and the CMBComp CMB-3w compression gives a continuous best fit at $\eta = 0.403405$ ($\Phi_i = 4.95779$), with $(w_0, w_a) = (-0.948924, -0.078448)$ and $\Delta\chi^2 = -1.9433$ relative to flat ΛCDM for one additional parameter. The corresponding ΔAIC $= +0.0567$ and ΔBIC $= +1.3525$ leave the models statistically tied. An independent DESI tracer-subset audit removes the $z_{\rm eff} = 0.510$ and 0.706 anisotropic BAO pairs separately and together; the posterior mean of η shifts by at most 0.27 of the baseline posterior standard deviation, while the scalar likelihood gain weakens rather than grows.

The posterior constrains a one-parameter curve rather than a free CPL plane, and the exact histories lie inside the Caldwell–Linder thawing band at the present epoch. Across the sampled range $0.01 \leq \eta \leq 0.60$, the numerically computed $\Delta t_{q=0}$ decreases monotonically with η and diverges as $\eta \to 0$, recovering the nested ΛCDM no-exit limit. Direct future integration gives the continuous-best-fit exit at $a = 13.8803$, 60.28 Gyr from the present, while the original $\Phi_i = 4$ benchmark exits at $a = 2.86230$ and 21.13 Gyr. This is a branch-conditioned prediction, not an observational detection of transient acceleration, because the ΛCDM boundary remains inside the posterior support. Near the minimum the canonical field enters damped oscillations with cycle-averaged $\langle w_\Phi \rangle \to 0$, so the asymptotic expansion is decelerating rather than de Sitter and has no permanent de Sitter event horizon.

The exact scalar histories are also executed in CAMB 1.6.6 using tabulated non-phantom $w(a)$, $c_s^2 = 1$, and zero anisotropic stress. For $30 \leq \ell \leq 2500$, the continuous best fit differs from matched ΛCDM by at most 0.217% in TT, 0.295% in EE, and 0.767% in lensing BB; the $\Phi_i = 4$ point reaches 0.484%, 0.660%, and 1.716%. With fixed $(A_s, n_s, \tau) = (2.1 \times 10^{-9}, 0.965, 0.0544)$, the present-day amplitudes are $\sigma_8 = 0.793894$ and 0.781132, versus 0.803341 for matched ΛCDM. A fixed-cosmology audit at $\eta = (0.05, 0.403405, 0.50)$ produces distinct exact histories, spectra, matter power, and growth products with no constant-w or PPF fallback. It also identifies and fixes a serialization-only redshift-ordering error in the archived growth writer; all growth values quoted here already use the corrected row pairing. Because A_s, n_s, τ, calibration, foregrounds, and $\sum m_\nu$ are not sampled, the quoted σ_8, S_8, and lensing amplitudes are conditional native predictions rather than a full CMB-likelihood posterior. In the ISW-sensitive range $2 \leq \ell \leq 30$, an independent audit of the archived CAMB 1.6.6 production spectra gives maximum lensed-TT residuals of 0.503% for the continuous best fit and 1.175% for $\Phi_i = 4$, both at $\ell = 4$ and far below the ideal full-sky per-multipole cosmic-variance scale. Under a constant-floor deformation $V/\Lambda_0 = \mathcal{F}(\Phi) + \delta$, percent-level positive floors shift the MAP exit time at the percent level, while $\delta_{\rm crit} \simeq 0.2433$ is required to eliminate the first finite $q = 0$ crossing; this quantifies sensitivity only and leaves radiative stability open. A robustness closeout finds all 18 tested CAMB floor/accuracy configurations successful for the two scalar histories, shows the upper η prior edge is posterior-inactive, and obtains a four-chain independent-likelihood MCMC posterior consistent with the deterministic quadrature result.

* Contact author: xtheperfectcellx@gmail.com

I. MOTIVATION AND SCOPE

Canonical thawing quintessence provides a restrictive physical alternative to a free Chevallier–Polarski–Linder (CPL) history: the field is initially Hubble frozen, remains non-phantom, and traces a one-dimensional locus once the potential and preparation surface are fixed [6, 7]. The central question here is a constraint-and-consequence question rather than a model-selection claim: can current observations constrain a single canonical branch strongly enough to calibrate both its observable thawing history and the future fate fixed by the same dynamical parameter? We therefore test whether the branch can remain close enough to ΛCDM to survive current background and compressed-CMB data, produce calculable perturbation and Boltzmann observables without phantom crossing, and, given the exact coefficient locking that sets $V_{\min} = 0$, map those present constraints onto its acceleration-exit time.

The potential is motivated first by its coefficient-locked mathematical structure. The standard isotonic oscillator combines a harmonic term with an inverse-square barrier [1]. Here those terms occur in the exact relation

$$\frac{V}{\Lambda_0} = \frac{\Phi^2}{4} + \frac{4}{\Phi^2} - 2 = \left(\frac{\Phi}{2} - \frac{2}{\Phi}\right)^2, \tag{1}$$

which enforces $V \geq 0$, a unique regular minimum at $\Phi = 2$, and $V_{,\Phi\Phi} > 0$ throughout the positive-field branch. The inverse-square term excludes the singular origin, the quadratic term controls the large-field flank, and the completed square removes any asymptotic cosmological-constant residue. In $Q = \Phi^2/4 = e^{\psi}$, the same structure is $Q + Q^{-1} - 2 = 2(\cosh\psi - 1)$, explaining the earlier "completed-cosh" description.

The historical route by which this potential was selected is separate from its evidentiary status. It began with Project Metatron, an unpublished exploratory vehicle-scale boundary-response design exercise. Chronologically, that exploration led first to Alpha-Lattice, which stripped its laboratory-facing question into a preregistered, mechanism-agnostic test of a driven electromagnetic state [3], and then to Procedural Vacuum Breakdown (PVB), which developed the gravitational question as a scalar-weighted three-form measure theory [4]. During the PVB-related scalar/measure exploration, the reciprocal structure $Q+Q^{-1}-2$ was isolated and subsequently detached from its parent interpretation to be studied as the standalone canonical quintessence model used here. This is a chronology of question refinement, not a claim that Alpha-Lattice tests PVB, that either predecessor supports this cosmology empirically, or that PVB derives the canonical action below. It also records that the functional form preceded the present DESI–Pantheon+–CMBComp likelihood analysis rather than being selected by scanning that likelihood.

The specifically relevant PVB result is narrow: after local constraint reduction, its fixed sector is Einstein gravity plus a scalar with an additive contribution $\Lambda_* e^{-\psi}$ [4]. That term motivates looking at reciprocal exponentials but does not force the matching e^{ψ} term, equal coefficients, the exact zero floor, or a canonical kinetic term in Φ. A separate branch-specific two-measure Palatini construction provides a conditional mathematical realization in which an algebraic measure constraint can produce

$$V_{\text{eff}} = \frac{(V_1 + M)^2}{4V_2}, \tag{2}$$

with $V_1 \propto e^{\psi}$, $V_2 \propto e^{\psi}$, and the integration-constant branch $M = -c$ yielding $V_{\text{eff}}/\Lambda_0 = e^{\psi} + e^{-\psi} - 2$ after canonical normalization [2]. This construction shows one way in which the reciprocal term and zero can arise together, but it is not inference machinery, a derivation from PVB, a unique parent theory, or a UV completion. The standalone analysis below therefore begins with the canonical action, coefficient lock, preparation surface, and one-dimensional branch as explicit model assumptions. No likelihood result is inherited from Project Metatron, Alpha-Lattice, PVB, or the optional two-measure construction; matter coupling, radiative stability, displacement selection, and parent-variable continuation through the minimum remain open.

The native spectra remain a conditional extension of the CMBComp posterior. The CMB compression does not sample A_s, n_s, τ, calibration parameters, foreground nuisance parameters, or $\sum m_\nu$. The native runs therefore fix those quantities and do not constitute a map-level or full-spectrum CMB likelihood analysis.

Recent DESI-era analyses of thawing quintessence reach differing conclusions about preference relative to ΛCDM depending on priors, parameterization, and supernova data combinations [16, 17]. The DESI DR2 extended analysis finds non-phantom alternatives disfavored relative to phantom-crossing histories but not excluded, while de Souza, Rodrigues, and Alcaniz obtain a statistically competitive parameterized thawing history that also predicts transient acceleration [18, 19]. A broader assessment of single-scalar-field dark energy likewise finds canonical quintessence only marginally distinguishable from Λ, with conclusions sensitive to dataset selection and prior assumptions [20]. These comparisons motivate treating the present $\Delta\chi^2$/AIC/BIC result as specific to the coefficient-locked isotonic branch and the adopted DESI DR2–Pantheon+–CMBComp likelihood rather than as a general verdict on thawing quintessence.

II. CANONICAL MODEL AND INFERENCE PARAMETERIZATION

The physical canonical field is $\phi_{\text{phys}} = M_{\text{Pl}}\Phi$; hence Φ, $u = d\Phi/dN$, and the fluctuation $\delta\Phi$ used below are dimensionless (in reduced-Planck units). Throughout, uppercase Φ denotes this scalar coordinate. Although the same glyph is often used for a Newtonian-gauge potential,

TABLE I. Core notation. Uppercase and lowercase symbols are distinct.

Symbol	Definition	First defined in
Φ, $\phi_{\rm phys}$	Dimensionless scalar coordinate and physical canonical field, with $\phi_{\rm phys} = M_{\rm Pl}\Phi$	Canonical model
$\mathcal{F}(\Phi)$	Dimensionless potential shape, V/Λ_0	Canonical model
Q	Positive weight coordinate $Q = \Phi^2/4 = e^{\psi}$	Motivation and scope
N	E-fold time $N = \ln a$	Canonical model
u	Scalar velocity $u = d\Phi/dN$	Canonical model
E	Reduced expansion rate $E = H/H_0$	Canonical model
α	Potential amplitude $\Lambda_0/(3M_{\rm Pl}^2 H_0^2)$	Canonical model
η	Sampled displacement coordinate $2/\Phi_i$	Abstract
h	Reduced Hubble constant, $H_0/(100\,{\rm km\,s^{-1}\,Mpc^{-1}})$	Canonical model
h_L	Synchronous-gauge scalar metric-trace perturbation	Perturbations
q	Deceleration parameter $-1 - H_{,N}/H$	Future fate
$\mathcal{S}(T)$	Exit-time survival probability $\Pr(\Delta t_{q=0} > T)$	Future fate
n_ρ	Same-phase dilution index, $\rho_\Phi \propto a^{-n_\rho}$	Future fate

no metric potential Φ appears in this work. Write

$$\mathcal{F}(\Phi) = \frac{\Phi^2}{4} + \frac{4}{\Phi^2} - 2, \tag{3}$$

$$V(\Phi) = \Lambda_0 \mathcal{F}(\Phi), \tag{4}$$

$$\mathcal{F}_{,\Phi} = \frac{\Phi}{2} - \frac{8}{\Phi^3}, \qquad \mathcal{F}_{,\Phi\Phi} = \frac{1}{2} + \frac{24}{\Phi^4} > 0. \tag{5}$$

This is a coefficient-locked member of the isotonic family rather than an arbitrary sum of two terms: changing the constant independently would lift the minimum and remove the transient-acceleration prediction. The point $\Phi = 2$ is a regular zero-energy minimum of the canonical model; it is not an endpoint of field space. We use $N = \ln a$, $u = d\Phi/dN$, and

$$\alpha \equiv \frac{\Lambda_0}{3M_{\rm Pl}^2 H_0^2}. \tag{6}$$

All energy densities are normalized to $3M_{\rm Pl}^2 H_0^2$. Let $B(N)$ and $P(N)$ denote the combined background density and pressure of baryons, cold dark matter, photons, massless neutrinos, and one 0.06 eV thermal neutrino. The exact closure and scalar equations are

$$E^2 \equiv \frac{H^2}{H_0^2} = \frac{B + \alpha\mathcal{F}(\Phi)}{1 - u^2/6}, \tag{7}$$

$$\Phi_{,N} = u, \tag{8}$$

$$u_{,N} = -\left(3 + \frac{H_{,N}}{H}\right) u - \frac{3\alpha\mathcal{F}_{,\Phi}}{E^2}, \tag{9}$$

$$\frac{H_{,N}}{H} = -\frac{3(B+P)}{2E^2} - \frac{u^2}{2}. \tag{10}$$

The scalar equation of state is

$$w_\Phi = \frac{E^2u^2/6 - \alpha\mathcal{F}}{E^2u^2/6 + \alpha\mathcal{F}}. \tag{11}$$

The sampled variables are

$$\boldsymbol{\theta} = (\eta, \omega_m, \omega_b, h), \qquad \eta \equiv \frac{2}{\Phi_i}, \tag{12}$$

where $\omega_m \equiv \Omega_m h^2$ is the total physical matter density and $\omega_b \equiv \Omega_b h^2$ is the physical baryon density; neither symbol denotes a fractional density parameter. The CAMB interface therefore forms the cold-dark-matter density only after subtracting the baryon and massive-neutrino physical densities. The uniform priors are

$$\begin{aligned} 0 \le \eta \le 0.6, \qquad & 0.133 \le \omega_m \le 0.147, \\ 0.0208727 \le \omega_b \le 0.0248, \qquad & 0.55 \le h \le 0.82. \end{aligned} \tag{13}$$

For every point, α is solved from $E(0) = 1$. Thus neither $\Omega_{\Phi 0}$ nor Λ_0 is fixed to the diagnostic value used in the source manuscript. The bounded coordinate η makes the nested ΛCDM boundary finite and explicit. The upper value $\eta = 0.6$ is a branch-domain choice rather than a posterior-derived field-space boundary: it corresponds to $\Phi_i = 10/3$ and already reaches a rapidly thawing, still-accelerating profiled solution with $w_0 = -0.6062$ and $q_0 = -0.0879$. The robustness closeout verifies that this ceiling is inactive in the inference. Under the archived continuous marginal density, only 3.53×10^{-6} of the posterior lies at $\eta \ge 0.55$; renormalizing after moving the ceiling inward to 0.55 changes the mean from 0.28252525 to 0.28252430. Even the much stronger truncation $\eta \le 0.50$ removes only 0.5804% of the posterior and shifts the mean to 0.281205. We therefore retain 0.6 as the declared range of the prepared thawing branch rather than silently extending the model into a more rapidly rolling regime.

Figure 1 shows the reconstructed exact equation-of-state histories used in the native CAMB runs.

III. EXACT CALDWELL–LINDER PHASE-PLANE TEST

For a canonical scalar, the exact phase-plane velocity can be evaluated without differencing a CPL surrogate:

$$w' \equiv \frac{dw_\Phi}{d\ln a} = -3(1 - w_\Phi^2) - \frac{2u\,\alpha\mathcal{F}_{,\Phi}}{E^2u^2/6 + \alpha\mathcal{F}}. \tag{14}$$

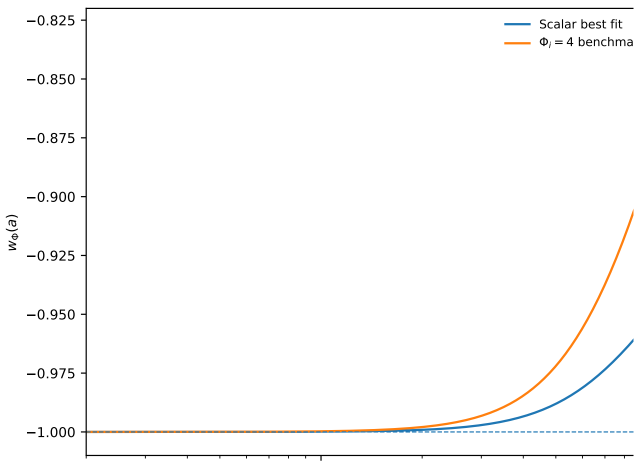

FIG. 1. Exact scalar equation-of-state histories reconstructed from the standalone field equations. Both trajectories remain non-phantom and approach $w = -1$ on the frozen early-time surface.

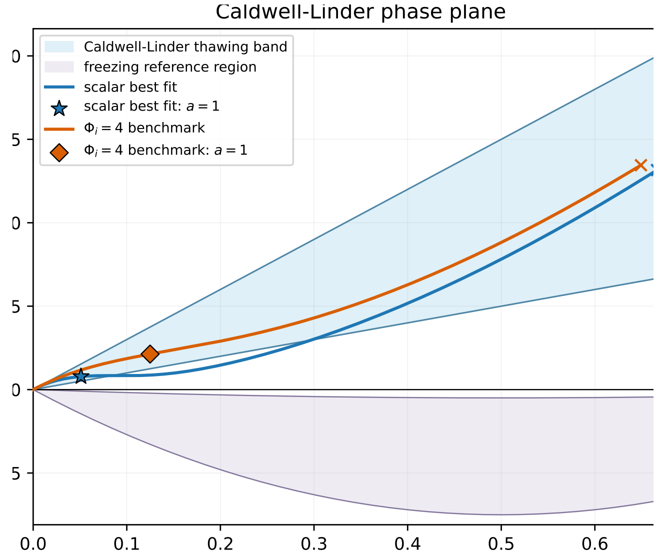

FIG. 2. Exact isotonic trajectories in the Caldwell–Linder plane, from the frozen epoch to the first future acceleration exit. The present epoch is marked separately on each curve and crosses mark the $q = 0$ endpoints. The positive-w' thawing wedge and the negative-w' freezing reference region are shaded; neither shaded construction is an asymptotic theorem outside its stated slow-roll-like domain.

We apply equation (14) to the posterior-calibrated field solutions from the frozen preparation surface through the first future $q = 0$ crossing. Figure 2 compares the resulting tracks with the practical thawing bounds

$$1 + w_\Phi < w' < 3(1 + w_\Phi) \tag{15}$$

and, for orientation, the freezing region $3w_\Phi(1 + w_\Phi) \leq w' \leq 0.2w_\Phi(1 + w_\Phi)$ [6].

At $a = 1$, the exact best-fit coordinates are

$$\begin{aligned}(1 + w_\Phi, w') &= (0.051075730,\ 0.078447842),\\ 3(1 + w_\Phi) &= 0.153227190.\end{aligned} \tag{16}$$

and the $\Phi_i = 4$ coordinates are

$$\begin{aligned}(1 + w_\Phi, w') &= (0.124911727,\ 0.212846477),\\ 3(1 + w_\Phi) &= 0.374735182.\end{aligned} \tag{17}$$

Both present-day points therefore satisfy the thawing inequality. Relative to the printed tangent-CPL summaries $-w_a = 0.078448$ and 0.212847, the exact w'_0 values differ by -1.58×10^{-7} and -5.23×10^{-7}, respectively (about 0.0002%). This near identity is expected: in the released inference code w_a is defined from the local present derivative, not from a global fit over $0 \leq z \leq 3$. The CPL history used only as a CMB-compression diagnostic is consequently a present-tangent proxy; no CPL continuation enters the phase-plane or future calculations.

The frozen initial condition, convex downhill roll, non-phantom stress tensor, and present scalar fractions $\Omega_\Phi = 0.69145$ and 0.68251 satisfy the assumptions behind the thawing classification today; the faster $\Phi_i = 4$ trajectory is closer to the edge of the usual $w \lesssim -0.8$ regime. Those assumptions fail as Ω_Φ approaches unity and w_Φ approaches $-1/3$ near the future exit. Correspondingly, the best-fit track drops below the lower practical line for $1.4875 \lesssim a \lesssim 7.7526$, after it has left the usual $\Omega_\Phi \lesssim 0.8$ domain, and re-enters before $q = 0$; the $\Phi_i = 4$ track remains within the plotted wedge through its exit. The isotonic family is therefore a bona fide one-parameter thawing curve over the observationally relevant era, not a two-parameter CPL plane. Its coefficient-locked $V_{\min} = 0$ then prevents de Sitter settling and carries the field beyond slow-roll thawing into the post-exit $w_\Phi \to 0$ oscillatory regime.

IV. CANONICAL SCALAR PERTURBATIONS

For the physical canonical field $\phi_{\rm phys} = M_{\rm Pl}\Phi$ and Lagrangian $\mathcal{L} = X - V(\phi_{\rm phys}/M_{\rm Pl})$,

$$c_s^2 = \frac{\mathcal{L}_{,X}}{\mathcal{L}_{,X} + 2X\mathcal{L}_{,XX}} = 1, \tag{18}$$

and the linear anisotropic stress vanishes. In synchronous gauge, comoving with the pressureless baryon-plus-CDM component, let h_L denote the scalar metric-trace perturbation.[1] The canonical field fluctuation obeys [8, 9]

$$\delta\Phi_{,NN} + \left(3 + \frac{H_{,N}}{H}\right)\delta\Phi_{,N} + \left[\left(\frac{kc}{aH}\right)^2 + \frac{V_{,\Phi\Phi}}{M_{\rm Pl}^2 H^2}\right]\delta\Phi + \frac{1}{2}h_{L,N}u = 0. \tag{19}$$

The scalar density and pressure perturbations are

$$\frac{\delta\rho_\Phi}{3M_{\rm Pl}^2 H_0^2} = \frac{E^2 u\,\delta\Phi_{,N}}{3} + \alpha\mathcal{F}_{,\Phi}\delta\Phi, \tag{20}$$

$$\frac{\delta p_\Phi}{3M_{\rm Pl}^2 H_0^2} = \frac{E^2 u\,\delta\Phi_{,N}}{3} - \alpha\mathcal{F}_{,\Phi}\delta\Phi. \tag{21}$$

For the trace equation, define the normalized perturbations once and for all by

$$\delta B_{cb} \equiv \frac{\delta\rho_{cb}}{3M_{\rm Pl}^2 H_0^2}, \qquad \delta B_\Phi \equiv \frac{\delta\rho_\Phi}{3M_{\rm Pl}^2 H_0^2}, \qquad \delta P_\Phi \equiv \frac{\delta p_\Phi}{3M_{\rm Pl}^2 H_0^2}. \tag{22}$$

Together with $\delta_{cb,N} = -h_{L,N}/2$, the trace metric equation is

$$h_{L,NN} + \left(2 + \frac{H_{,N}}{H}\right)h_{L,N} = -\frac{3}{E^2}\left(\delta B_{cb} + \delta B_\Phi + 3\delta P_\Phi\right). \tag{23}$$

The independent validation solver begins at $z = 49$ with the matter-era growing mode and the regular frozen-field perturbation $\delta\Phi = \delta\Phi_{,N} = 0$. It remains restricted to $k \geq 0.01\,h\,\mathrm{Mpc}^{-1}$ and omits photon and neutrino perturbation hierarchies. Its role is to isolate the small clustered-versus-smooth quintessence correction; CAMB is authoritative for CMB anisotropies and total-matter transfer functions.

A. Native CAMB implementation

CAMB provides a tabulated dark-energy equation-of-state interface and a rest-frame sound-speed parameter [10]. The standalone solver supplies the exact $w(a)$ history to CAMB 1.6.6 through the fluid module with $c_s^2 = 1$. The executed scalar runs report `uses_tabulated_w=true`, `constant_w_fallback=false`, and `ppf_fallback=false`.

The field is numerically indistinguishable from $w = -1$ during the frozen early epoch. CAMB's fluid transfer system became stiff when entries were regularized only to $1 + w = 10^{-10}$. The successful execution uses the documented floor

$$1 + w \geq 10^{-6}, \tag{24}$$

[1] The subscript distinguishes the metric trace h_L from the reduced Hubble constant h, including its use in units such as $h\,\mathrm{Mpc}^{-1}$.

TABLE II. Executed CAMB floor/accuracy stress test. "Pass" means all three `AccuracyBoost` values $1.0, 1.5, 2.0$ completed with `DarkEnergyFluid`, tabulated $w(a)$ and $c_s^2 = 1$.

Model	$1+w$ floor	altered samples	$A = 1, 1.5, 2$
Best fit	10^{-6}	2881	Pass
Best fit	10^{-7}	2406	Pass
Best fit	10^{-8}	1934	Pass
$\Phi_i = 4$	10^{-6}	2705	Pass
$\Phi_i = 4$	10^{-7}	2231	Pass
$\Phi_i = 4$	10^{-8}	1762	Pass

while retaining the exact reconstructed history whenever its physical value exceeds the floor. The maximum absolute modification is 10^{-6} and occurs where Ω_Φ is negligible. This is a numerical non-phantom regularization, not a constant-w, CPL, or PPF replacement.

The archived Colab execution also contains an explicit floor/accuracy stress matrix for both displayed scalar histories. Floors $1 + w \geq (10^{-6}, 10^{-7}, 10^{-8})$ were crossed with `AccuracyBoost` $= (1.0, 1.5, 2.0)$, giving 18 independent `DarkEnergyFluid` executions. All 18 completed with tabulated $w(a)$, $c_s^2 = 1$, and no PPF or constant-w fallback. Table II gives the floor-dependent number of modified table entries; the result is independent of accuracy setting. A separate archived validation report records the maximum floor-induced change in C_ℓ power across 10^{-6}–10^{-8} as below 0.001%. The matrix did not serialize a floor-by-floor $f\sigma_8$ comparison, so no separate growth-floor bound is claimed here.

The native spectra use fixed primordial and reionization values

$$A_s = 2.1 \times 10^{-9}, \qquad n_s = 0.965, \qquad \tau = 0.0544, \tag{25}$$

with $\sum m_\nu = 0.06\,\mathrm{eV}$, $N_{\rm eff} = 3.044$, and one massive thermal species. The total physical matter density from the posterior is partitioned consistently as $\omega_m = \omega_b + \omega_c + \omega_\nu$. At the continuous best fit the executed values are $\omega_m = 0.140584$, $\omega_b = 0.0225253$, $\omega_\nu = 0.000644867$, and $\omega_c = 0.117413833$, which close to the stated total at floating-point precision.

V. DATA AND LIKELIHOOD

The posterior combines three statistically independent blocks.

a. DESI DR2 BAO. We use the 13-observable DR2 BAO vector and covariance, including D_M/r_d, D_H/r_d, and D_V/r_d measurements [11]. The drag horizon is evaluated consistently from a CAMB calibration grid in (ω_m, ω_b) for the fixed neutrino sector.

b. Pantheon+. The supernova likelihood is the 11-knot Gaussian compression of $\log r_p(z)$ [12, 13]. The compression removes the absolute distance normalization, so no supernova absolute-magnitude nuisance parameter is sampled.

TABLE III. Exact dark-energy fractions. For the scalar rows $\Omega_{\rm DE} = \Omega_\Phi$; for the flat-ΛCDM row it is Ω_Λ.

z	Scalar best	$\Phi_i = 4$	Flat ΛCDM
1100	1.3650×10^{-9}	1.4455×10^{-9}	1.3068×10^{-9}
300	8.1377×10^{-8}	8.6223×10^{-8}	7.7875×10^{-8}
50	1.7922×10^{-5}	1.8994×10^{-5}	1.7148×10^{-5}
0	0.691449	0.682509	0.697439

c. CMBComp CMB-3w. We use the correlated vector (R, ℓ_A, ω_b) from CMBComp [14]. CMBComp was constructed from SPT-3G D1, ACT DR6, Planck PR3 primary anisotropy, and Planck PR4/NPIPE lensing chains and was validated in smooth late-time model spaces. It is a compressed geometric likelihood, not a direct map-level likelihood.

A. Quantitative domain of the compressions

The exact isotonic histories are smooth, canonical and non-phantom, with w_Φ increasing monotonically from its frozen value to the present. Their early dark-energy fractions are shown in Table III. The 10^{-9}-level fractions at recombination and the absence of sharp transitions in the observed era satisfy the smoothness and early-energy prerequisites of the late-time domain for which CMBComp was constructed. These are necessary checks, not by themselves a guarantee that a particular non-CPL history is represented faithfully.

That domain check is not uniform over the two displayed scalar points. At the continuous best fit, the exact isotonic history and its present-tangent CPL proxy differ by at most 1.66×10^{-4} in $E(z)$ for $0 \le z \le 3$; their CMBComp coordinate offsets are 0.006 marginal standard deviations in R and 0.044 in ℓ_A. For $\Phi_i = 4$, however, the offsets are 0.225σ in R and 1.77σ in ℓ_A. The benchmark therefore lies near or outside the regime in which a CPL-calibrated compression can be trusted. This is precisely why its native CAMB execution is necessary; the compressed result for that point is retained only as a diagnostic, not promoted to a full-CMB validation.

The 11-knot Pantheon+ likelihood is lossless for this model class only to the extent that the exact $\log r_p(z)$ is represented by the prescribed natural cubic spline. Reconstructing each exact distance, sampling it at the fixed knot at $z = 0$ plus the 11 published knots, and comparing the spline with the exact function over $0 \le z \le 2.27$ gives maximum absolute natural-log residuals 4.71×10^{-5} for the scalar best fit and 4.00×10^{-5} for $\Phi_i = 4$. Thus the required interpolation condition is numerically satisfied at far below the marginal knot uncertainties; this test concerns representation fidelity, not the observational adequacy of the supernova data themselves.

The BAO, supernova, and CMB-compression vectors are treated as three independent data blocks. The parameter ω_b appears in the sampled vector and as one component of the CMBComp datum: it enters once as a parameter and once as a datum, not twice as a datum. The DESI drag horizon $r_d(\omega_m, \omega_b)$ is evaluated from the CAMB calibration grid with the same fixed $N_{\rm eff} = 3.044$, $\sum m_\nu = 0.06\,\mathrm{eV}$ sector used by the exact background and native runs. Conversely, the compression cannot constrain or marginalize over A_s, n_s, τ, calibration, foreground, or neutrino-mass freedom; all reported amplitude observables are conditional on their fixed values.

TABLE IV. High-accuracy background-plus-compressed-CMB best fits. The original $\Phi_i = 4$ row is profiled over (ω_m, ω_b, h).

	Scalar best	$\Phi_i = 4$	Flat ΛCDM
η	0.403405	0.500000	0
Φ_i	4.95779	4	∞
ω_m	0.140584	0.140126	0.140960
ω_b	0.0225253	0.0225464	0.0225079
h	0.675088	0.664431	0.682649
w_0	−0.948924	−0.875088	−1
w_a	−0.078448	−0.212847	0
$\chi^2_{\rm BAO}$	12.2109	12.2482	13.0476
$\chi^2_{\rm SN}$	17.9926	19.4442	20.5701
$\chi^2_{\rm CMB}$	4.1699	6.4208	2.6991
$\chi^2_{\rm total}$	34.3734	38.1132	36.3167

VI. POSTERIOR ALGORITHM AND BACKGROUND-LIKELIHOOD RESULTS

At each η, the conditional nuisance cosmology (ω_m, ω_b, h) is optimized. Its local Gauss–Newton covariance defines a change of variables for product Gauss–Hermite integration. The final posterior uses order 7 and $\Delta\eta = 0.01$, followed by trapezoidal integration in η. Two controls use order 5 on the same grid and order 5 with $\Delta\eta = 0.02$. This is a deterministic posterior rather than an MCMC chain; quadrature-order and grid-refinement tests replace Gelman–Rubin and effective-sample-size diagnostics [15].

Table IV gives the continuous high-accuracy fits. The scalar improves the total chi-square by 1.9433 for one additional parameter, which is insufficient to overcome standard complexity penalties.

For the scalar relative to ΛCDM,

$$\Delta\chi^2 = -1.9433, \qquad \Delta\mathrm{AIC} = +0.0567, \\ \Delta\mathrm{BIC} = +1.3525. \tag{26}$$

The scalar and ΛCDM are essentially tied by AIC, while BIC mildly favors ΛCDM. There is no evidence for selecting the scalar model.

The original $\Phi_i = 4$ point is $\Delta\chi^2 = 3.7398$ above the scalar best fit and 1.7965 above ΛCDM. It lies barely inside the conventional 95% one-parameter profile contour but outside the 95% equal-tailed posterior. The posterior probability $\Pr(\eta \ge 0.5) = 0.00580$.

TABLE V. Posterior summaries under the stated uniform priors.

Parameter	Mean ± s.d.	2.5%	Median	97.5%
η	0.2825 ± 0.1337	0.020	0.310	0.480
ω_m	0.140748 ± 0.000637	0.139486	0.140792	0.142028
ω_b	0.0225181 ± 0.0000920	0.0223093	0.0225157	0.0227261
h	0.67822 ± 0.00487	0.66712	0.67894	0.68592
w_0	-0.97009 ± 0.02888	-0.99994	-0.97805	-0.89766
w_a	-0.04641 ± 0.04714	-0.16978	-0.03243	-0.00009
Ω_m	0.30603 ± 0.00478	0.29779	0.30568	0.31638

TABLE VI. DESI LRG-subset robustness. The full row is the independent reconstruction used as the control; the published order-7 quadrature posterior remains authoritative. $\Delta\chi^2$ is scalar minus matched flat-ΛCDM within the same subset analysis.

BAO set	MAP η	Mean η	s.d.	$\Delta\chi^2$
Full control	0.4037	0.2824	0.1334	-1.945
−LRG1	0.3853	0.2637	0.1329	-1.391
−LRG2	0.3842	0.2621	0.1332	-1.314
−LRG1,2	0.3627	0.2462	0.1314	-0.885

A. DESI tracer-subset robustness

The DESI contribution is sufficiently modular to test whether the inferred branch is being driven by the two low-redshift LRG anisotropic BAO pairs. As an independent robustness reconstruction, we delete the correlated $(D_M/r_d, D_H/r_d)$ pair at $z_{\rm eff} = 0.510$ (LRG1), the pair at $z_{\rm eff} = 0.706$ (LRG2), or both pairs, retain the corresponding principal covariance submatrix for the remaining BAO observables, and re-optimize the scalar and flat-ΛCDM nuisance cosmologies. The subset marginal in η is reconstructed with a fixed-η Laplace marginalization. Before deleting any tracer, this implementation recovers $\eta = 0.28238 \pm 0.13343$, median 0.3072, and 95% interval $[0.0192, 0.4767]$, compared with the authoritative order-7 quadrature result 0.2825 ± 0.1337, median 0.310, and $[0.020, 0.480]$. The high-accuracy background solver was then used to re-evaluate each subset MAP.

Removing either pair shifts the posterior mean by only 0.14–0.15 of the baseline posterior standard deviation; removing both shifts it by 0.27 standard deviations. The posterior width changes by less than 1.5%, and the 97.5th percentile moves only from 0.477 in the control to 0.461 when both pairs are removed. At the same time the already modest scalar likelihood gain decreases in magnitude. Thus the branch constraint and the conclusion of no model selection are not artifacts of either LRG1 or LRG2; if anything, those bins contribute part of the small preference for the scalar best fit.

B. Sampling-method robustness closeout

As a sampling-method check, we independently sampled the compressed likelihood with four random-walk Metropolis chains. Each chain contains 28,000 steps; after discarding the first 3,000 per chain, 100,000 draws remain. The acceptance fractions are 0.190–0.198, the largest split-$\hat{R}$ among $(\eta, \omega_m, \omega_b, h)$ is 1.008, and the smallest autocorrelation-based effective sample size is 967. The independent-likelihood MAP is $\eta = 0.40340587$ with $\chi^2 = 34.3683$, compared with the archived deterministic values $\eta = 0.403405$ and $\chi^2 = 34.3734$. The small χ^2 offset between the independent MCMC implementation and the archived deterministic pipeline is negligible for this cross-check; the MCMC result is not substituted for the archived quadrature posterior.

TABLE VII. Deterministic posterior and independent-likelihood MCMC cross-check. The MCMC w_0 summary uses a reproducible 5000-draw derived-parameter subset.

Method	η mean ± s.d.	95% η	w_0 mean ± s.d.
Quadrature	0.2825 ± 0.1337	$[0.020, 0.480]$	-0.97009 ± 0.02888
MCMC	0.27597 ± 0.13625	$[0.0183, 0.4751]$	-0.97093 ± 0.02892

TABLE VIII. Direct future integration of the posterior best fit and the original benchmark. The acceleration-exit column is the first future root of $q = 0$. The minimum column is the first crossing of $\Phi = 2$. Times are measured from the present.

Model	q_0	$a_{q=0}$	$\Delta t_{q=0}$ [Gyr]	$\Phi_{q=0}$	$a_{\Phi=2}$	$\Delta t_{\Phi=2}$ [Gyr]
Scalar best	-0.48416	13.8803	60.28	2.83162	21.5800	85.46
$\Phi_i = 4$	-0.39584	2.86230	21.13	2.81334	4.46342	38.94

The agreement is sufficient for its intended purpose: the posterior location and width, the nested-boundary support, and the lack of model selection do not depend on the deterministic quadrature algorithm. The deterministic posterior remains authoritative because it is the likelihood object used throughout the published push-forward and tables.

VII. TRANSIENT ACCELERATION AND FUTURE FATE

The exact zero-energy minimum changes the physical content of the model. A finite-displacement scalar branch does not approach a positive vacuum-energy plateau: after the present thawing phase the field continues down the convex flank and the total expansion eventually crosses from acceleration back to deceleration. We locate this transition by extending the same thermal-neutrino background and scalar equations used in the posterior, with no CPL extrapolation. The deceleration parameter and elapsed cosmic time are

$$q(N) = -1 - \frac{H_{,N}}{H}, \qquad \Delta t(N) = H_0^{-1} \int_0^N \frac{dN'}{E(N')}. \quad (27)$$

The future solve starts from the reconstructed $N = 0$ state and stops at the first downward crossing of $\Phi = 2$.

At the transition, the total equation of state is $w_{\rm eff} = -1/3$ by definition. The posterior best fit is already

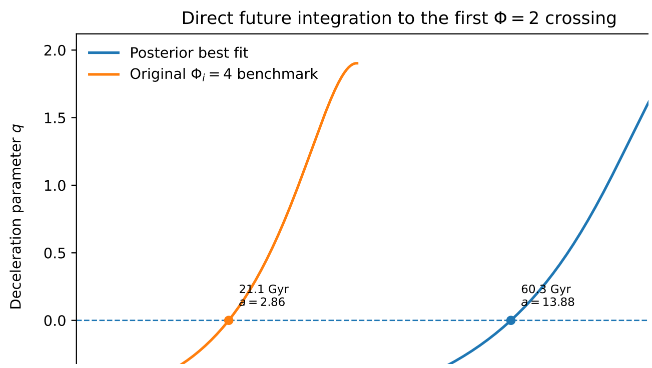

FIG. 3. Deceleration parameter from direct future integration. Each curve terminates at its first crossing of the zero-energy minimum. The $q = 0$ roots occur before $\Phi = 2$, so the finite-time end of acceleration does not depend on how a speculative parent construction is continued through its endpoint variables.

scalar dominated, with $\Omega_\Phi = 0.99912$ and $w_\Phi = -0.33363$ at $q = 0$. For the $\Phi_i = 4$ benchmark, $\Omega_\Phi = 0.95019$ and $w_\Phi = -0.35081$. Thus the old statement that acceleration ends on an approximately 20 Gyr timescale applies to the original $\Phi_i = 4$ trajectory, not to the more frozen continuous best fit. The updated best-fit prediction is approximately 60 Gyr.

A. Posterior push-forward of the exit time

We evaluate $\Delta t_{q=0}(\eta)$ by direct field integration at every positive node of the existing $\Delta\eta = 0.01$ grid, using the profiled (ω_m, ω_b, h) at that node and the archived order-7 Gauss–Hermite weights. Within each cell, piecewise-linear posterior density and linearly interpolated nuisance cosmology reproduce the trapezoidal integration while literal node sums provide grid controls. Over the sampled range $0.01 \le \eta \le 0.60$, the numerically computed exit time decreases monotonically with η and diverges as $\eta \to 0$. This is the present-to-future predictive map of the branch: no independent future-fate parameter is introduced beyond η. Constraints that drive η toward zero push allowed exits later and recover the nested ΛCDM no-exit boundary; evidence for $\eta > 0$ would instead imply a finite exit-time distribution, given the exact coefficient locking. The monotonicity statement for the fully re-profiled exit time is numerical over the sampled branch, not an analytic theorem, but the potential supplies a simple analytic reason for its direction. On the frozen preparation surface $\Phi_i = 2/\eta$,

$$s_i \equiv \frac{V_{,\Phi}}{V}\bigg|_i = \frac{\eta(1+\eta^2)}{1-\eta^2}, \qquad \frac{ds_i}{d\eta} = \frac{1+4\eta^2-\eta^4}{(1-\eta^2)^2} > 0 \tag{28}$$

for $0 < \eta < 1$. Thus increasing η monotonically increases the initial logarithmic downhill slope and favors earlier, faster thawing toward the zero-energy minimum. Resolving α and the nuisance cosmology prevents this local argument from being a proof of global $\Delta t_{q=0}$ monotonicity, but it explains the sign of the numerically verified map.

A stronger semi-analytic check follows in the scalar-dominated slow-roll limit. For the pre-minimum branch $\Phi > 2$,

$$\frac{\sqrt{\mathcal{F}(\Phi)}}{\mathcal{F}_{,\Phi}} = \frac{\Phi^2}{\Phi^2+4}, \qquad G(\Phi) \equiv \Phi - 2\tan^{-1}\left(\frac{\Phi}{2}\right), \tag{29}$$

so the slow-roll time has the shape

$$X(\eta) \equiv \frac{G(2/\eta) - G(\Phi_{\rm exit})}{h(\eta)\sqrt{\alpha(\eta)}}. \tag{30}$$

Normalizing this approximation once to the exact MAP integration gives

$$\Delta t_{\rm SA}(\eta) = 60.28\,\text{Gyr}\,\frac{X(\eta)}{X(0.403405)}, \qquad \Phi_{\rm exit} = 2.83. \tag{31}$$

Against fresh direct integrations at $\eta = (0.20, 0.30, 0.40, 0.45, 0.50)$, the fractional errors are $(+1.33, +0.16, -0.03, +0.97, +4.48)\%$, respectively. The degradation at $\eta = 0.50$ is expected for the faster-thawing benchmark as the slow-roll assumptions become less accurate. The approximation is therefore an intuition and cross-check, not a replacement for the exact event integration.

It also makes the nested-boundary divergence transparent. On the frozen preparation surface,

$$\mathcal{F}(2/\eta) = \frac{(1-\eta^2)^2}{\eta^2}, \tag{32}$$

and flatness gives $\alpha \propto \eta^2$ as $\eta \to 0$, while $G(2/\eta) \sim 2/\eta$. Up to the slowly varying profiled h, the slow-thawing scaling is consequently

$$\Delta t_{q=0} \propto \eta^{-2} \qquad (\eta \to 0), \tag{33}$$

which supplies a semi-analytic explanation for the numerical divergence. The weighted η posterior can therefore be mapped directly to $\mathcal{S}(T) \equiv \Pr(\Delta t_{q=0} > T)$ without fitting or extrapolating an equation-of-state proxy.

The boundary requires care. With the stated continuous uniform prior, $\Pr(\eta = 0) = 0$ exactly; a trapezoidal endpoint is not a physical ΛCDM atom. We set $\eta_{\rm min} = 0.005$, half the posterior grid spacing, and treat $\Pr(\eta \le \eta_{\rm min}) = 0.006530$ as an unresolved right-censored near-boundary contribution. The resolved push-forward gives

$$\begin{aligned} \Pr(\Delta t_{q=0} < 100\,\text{Gyr}) &= 0.3782, \\ \Pr(\Delta t_{q=0} < 1\,\text{Tyr}) &= 0.7870. \end{aligned} \tag{34}$$

Conditional on $\eta > \eta_{\rm min}$, the 16th, 50th, and 84th percentile exit times are 50.0 Gyr, 155.8 Gyr, and 1.73 Tyr;

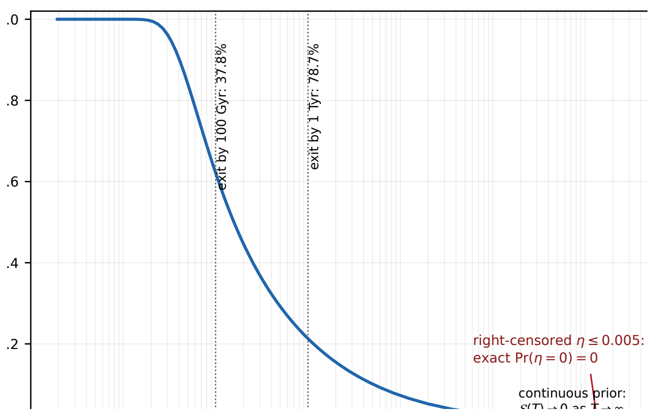

FIG. 4. Posterior survival function for the exact acceleration-exit time. The arrow at the long-time edge denotes the unresolved $\eta \leq 0.005$ half-cell, not a literal ΛCDM point mass: under the continuous prior the exact $\Pr(\eta = 0)$ is zero and the true survival tends to zero as $T \to \infty$.

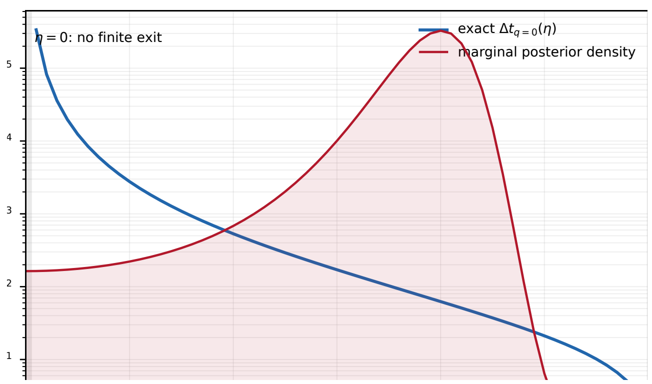

FIG. 5. Exact acceleration-exit time across the positive η grid, with the marginal posterior density on the second axis. The monotone divergence toward the nested $\eta = 0$ boundary motivates the survival-function and censoring convention used in Figure 4.

the outer 2.5th–97.5th percentile interval is 27.4 Gyr–56.8 Tyr. These results are a censored push-forward, never a mean-plus-standard-deviation summary and never posterior odds for an exact ΛCDM model.

The first minimum crossing is kinetic dominated instantaneously, but the scalar does not remain in a stiff $w_\Phi = 1$ state. Writing $\varphi = \Phi - 2$, the potential near the minimum is

$$\begin{aligned} V(\Phi) &= \Lambda_0 \left[\varphi^2 + O(\varphi^3)\right] \\ &= \frac{1}{2} M_{\rm Pl}^2 m_\Phi^2 \varphi^2 + O(\Lambda_0 \varphi^3), \\ m_\Phi^2 &= \frac{2\Lambda_0}{M_{\rm Pl}^2}, \qquad \frac{m_\Phi^2}{H_0^2} = 6\alpha. \end{aligned} \tag{35}$$

At the continuous MAP point, $\alpha = 0.172573$ and therefore $m_\Phi^2/H_0^2 = 1.03544$. We also continue the exact equations through three complete same-phase cycles after the first downward $\Phi = 2$ crossing. For the scalar best fit the cosmic-time averages are $\langle w_\Phi \rangle = (3.19 \times 10^{-3}, 8.32 \times 10^{-4}, 3.81 \times 10^{-4})$; for $\Phi_i = 4$ they are $(3.02 \times 10^{-3}, 7.85 \times 10^{-4}, 3.58 \times 10^{-4})$. Defining the same-phase dilution index by $\rho_\Phi \propto a^{-n_\rho}$, the corresponding sequences are $n_\rho = (3.0265, 3.0071, 3.0032)$ and $(3.0260, 3.0069, 3.0032)$. By the third cycle, the change in $\rho_\Phi a^3$ is below 7×10^{-4} for either history. This numerical continuation confirms the standard quadratic-minimum limit $\langle w_\Phi \rangle \to 0$ and $\rho_\Phi \propto a^{-3}$ [5]. Matter and the oscillating scalar therefore approach a dustlike decelerating state with $a(t) \propto t^{2/3}$ and $q \to 1/2$. The future conformal-time integral then diverges,

$$\int_t^\infty \frac{dt'}{a(t')} \sim \int_t^\infty dt'\, t'^{-2/3} = \infty, \tag{36}$$

so the canonical model has no asymptotic de Sitter event horizon. It avoids the permanent causal isolation characteristic of eternal de Sitter expansion, although this statement does not by itself eliminate every thermodynamic route to a cold late universe.

B. Constant-floor sensitivity of the fate map

The exact zero of the coefficient-locked potential is structurally important, so we also test a controlled deformation that lifts only the additive floor,

$$\frac{V(\Phi)}{\Lambda_0} = \mathcal{F}(\Phi) + \delta, \qquad \delta \equiv \frac{\delta V_0}{\Lambda_0} > 0. \tag{37}$$

This is deliberately narrower than a radiative-stability calculation: it asks how strongly the calibrated MAP future changes if the zero-energy minimum is replaced by a small positive constant floor, without claiming a protective symmetry or a generic loop correction. Repeating the future integration gives the absolute fractional changes in the first MAP acceleration-exit time shown in Table IX. Small lifts perturb the fate map continuously; a finite first $q = 0$ crossing persists until $\delta_{\rm crit} \simeq 0.2433$. At that threshold the residual constant floor is approximately 3.9% of today's critical density. Thus percent-level lifts shift the benchmark exit time at the percent level, while eliminating the finite MAP exit requires a substantially larger positive floor.

The prediction is conditional in two distinct senses. First, $\eta = 0$ is exactly ΛCDM and never exits acceleration, so the exit time diverges toward the nested boundary; the present posterior does not constitute evidence that the finite branch is realized. Second, the post-minimum oscillatory statement is a result of the canonical model. If the optional two-measure parent is imposed, its variables require an additional continuation rule at the measure-degeneracy surface. The observationally sharper result—the first future $q = 0$ crossing—occurs before that surface in both calibrated trajectories.

TABLE IX. Constant-floor deformation of the continuous MAP future. The second column is the absolute fractional change of the first future $q = 0$ time relative to the coefficient-locked $\delta = 0$ result. This is a constant-lift sensitivity test, not a radiative-stability proof.

$\delta = \delta V_0/\Lambda_0$	$\lvert \Delta t_{q=0}(\delta)/\Delta t_{q=0}(0) - 1 \rvert$
10^{-4}	$\simeq 0.01\%$
10^{-3}	$\simeq 0.10\%$
10^{-2}	$\simeq 0.96\%$
3×10^{-2}	$\simeq 2.94\%$
0.1	$\simeq 10.6\%$

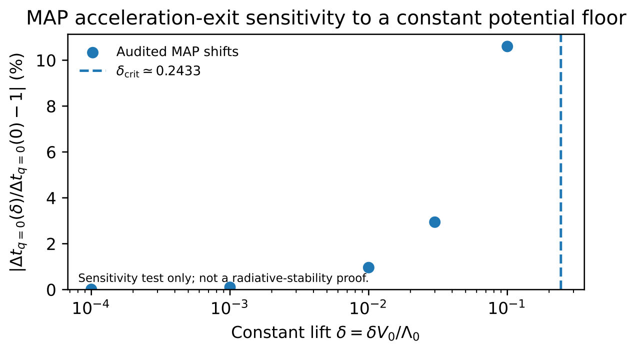

FIG. 6. Audited MAP acceleration-exit sensitivity to a positive constant potential floor. Points show only the executed deformation values listed in Table IX; the dashed line marks the numerically determined $\delta_{\rm crit} \simeq 0.2433$ at which the finite first MAP $q = 0$ crossing disappears.

VIII. NATIVE CMB SPECTRA AND TRANSFER OBSERVABLES

The three-run comparison suite (flat ΛCDM, the scalar best fit, and the re-profiled $\Phi_i = 4$ benchmark) completed successfully in CAMB 1.6.6 to $\ell_{\rm max} = 3000$ at accuracy setting 1.0. The scalar runs used the exact tabulated histories after the early frozen-epoch regularization described above. Table X summarizes the executed residuals relative to the separately best-fitting flat-ΛCDM cosmology.

The primary TT residual is sub-percent across the measured acoustic range (Figure 7). EE follows the same pattern at slightly larger amplitude, while lensing BB is more responsive to the accumulated growth suppression (Figure 9). These spectra confirm that the allowed best-fit branch remains close to ΛCDM in the primary CMB while producing a larger, though still percent-level, late-time lensing response.

A. Low-ℓ lensed-TT audit in the ISW-sensitive range

The acoustic-range summary above begins at $\ell = 30$. To close the separate low-ℓ question, we recomputed the same-cosmology lensed-TT fractional residuals over $2 \le \ell \le 30$ against the archived CAMB 1.6.6 production spectra. Table XI gives the independently checked production-version extrema. The maximum occurs at $\ell = 4$: 0.503% for the continuous best fit and 1.175% for the $\Phi_i = 4$ benchmark. By $\ell = 30$ the residuals have fallen to approximately 0.128% and 0.292%, respectively. The maximum residuals are only about 0.0115 and 0.0271 of the ideal full-sky per-multipole cosmic-variance scale. The late-time ISW contribution makes this range physically sensitive to the thawing history, but the plotted quantity is total lensed TT; these numbers therefore do not support a standalone low-ℓ detectability claim. Fractional ratios are unchanged by the conventional same-ℓ conversion between CAMB D_ℓ and C_ℓ.

The linear matter-power residual is shown in Figure 10. At fixed primordial amplitude, the continuous best fit gives a roughly 4% suppression over conventional linear scales, while the faster-thawing $\Phi_i = 4$ point approaches 10%. This model-to-model power difference includes the separately optimized background parameters and transfer-function shape; it should not be confused with the much smaller full-versus-smooth scalar-clustering correction on an identical background.

TABLE X. Native CAMB residual summary relative to matched flat ΛCDM. TT, EE, and BB entries are maximum absolute fractional differences for $30 \le \ell \le 2500$. TE is normalized by max $\lvert D^{TE}_{\ell,\Lambda} \rvert$ to avoid zero-crossing singularities. The $P_{\rm m}(k)$ entry is the maximum absolute linear-power residual for $k \le 0.5\,h\,{\rm Mpc}^{-1}$ at $z = 0$.

Model	TT [%]	EE [%]	lensing BB [%]	TE norm. [%]	$P_{\rm m}(k, z = 0)$ [%]
Scalar best	0.217	0.295	0.767	0.205	4.138
$\Phi_i = 4$	0.484	0.660	1.716	0.456	9.625

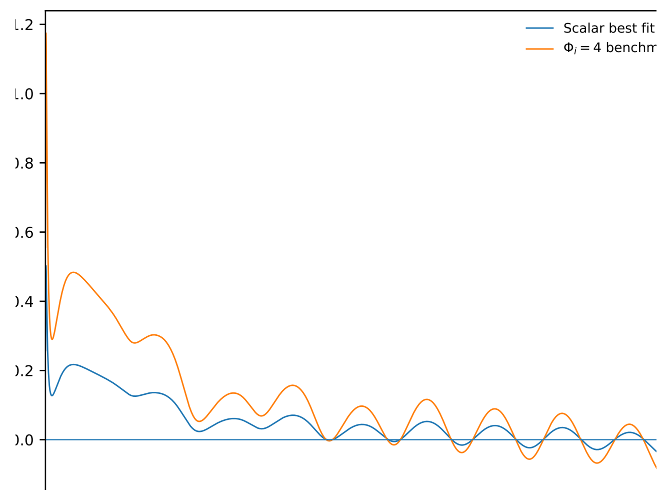

FIG. 7. Native CAMB lensed TT residuals relative to the matched flat-ΛCDM run. The continuous scalar best fit remains within approximately 0.22% for $30 \le \ell \le 2500$; the original $\Phi_i = 4$ trajectory reaches approximately 0.48%.

TABLE XI. Low-ℓ lensed TT residuals (ISW-sensitive range) from the archived CAMB 1.6.6 production-spectrum audit. "CV ratio" is the maximum residual divided by the ideal full-sky per-multipole cosmic-variance scale at the same multipole.

Model	max $\lvert\Delta TT/TT\rvert$ [%]	$\ell_{\max}$	$\lvert\Delta TT/TT\rvert_{\ell=30}$ [%]	CV ratio
Scalar best	0.503	4	$\simeq 0.128$	$\simeq 0.0115$
$\Phi_i = 4$	1.175	4	$\simeq 0.292$	$\simeq 0.0271$

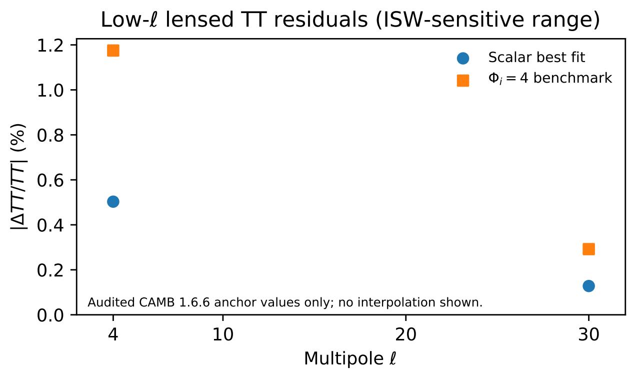

FIG. 8. Low-ℓ lensed TT residuals in the ISW-sensitive range, showing only the independently checked CAMB 1.6.6 anchor values available from the closeout audit: the maximum at $\ell = 4$ and the value at $\ell = 30$. No interpolated low-ℓ curve is shown.

B. Growth normalization and $f\sigma_8$

CAMB returns its growth arrays in the order of `results.Params.Transfer.PK_redshifts`, which is high to low redshift for these runs. The original native writer instead attached an independently specified low-to-high label vector without reordering the values. Thus the numerical σ_8 and $f\sigma_8$ arrays were correct but their raw CSV redshift labels were reversed. The independent assessment corrected the pairing before producing Table XII and Figure 11; the release writer now zips values to CAMB's own redshift array and sorts complete rows. The continuous best fit lowers $\sigma_8(z = 0)$ by 1.176% relative to matched ΛCDM, while the re-profiled $\Phi_i = 4$ point lowers it by 2.765%. Because the scalar fits have slightly larger Ω_m, the corresponding S_8 values remain close: 0.8050 and 0.8035 versus 0.8067 for ΛCDM.

The predicted $f\sigma_8$ values at $z = (0, 0.5, 1, 2)$ are respectively

$$\Lambda\text{CDM} : (0.414534,\ 0.467200,\ 0.428453,\ 0.324384), \tag{38}$$

$$\text{scalar best} : (0.413669,\ 0.459841,\ 0.422108,\ 0.320612), \tag{39}$$

$$\Phi_i = 4 : (0.412871,\ 0.450034,\ 0.413528,\ 0.315444). \tag{40}$$

These predictions provide the direct inputs for a future covariance-aware RSD likelihood, but no RSD data are included in the present posterior.

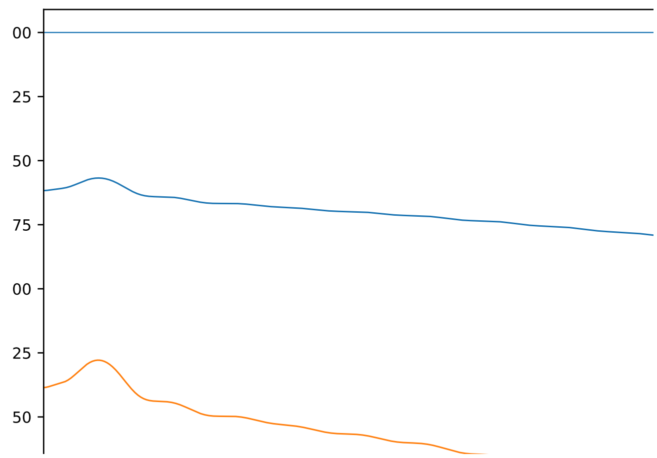

FIG. 9. Native CAMB lensing-BB residuals relative to matched flat ΛCDM. The larger response compared with primary TT and EE reflects the integrated late-time growth and lensing-potential difference.

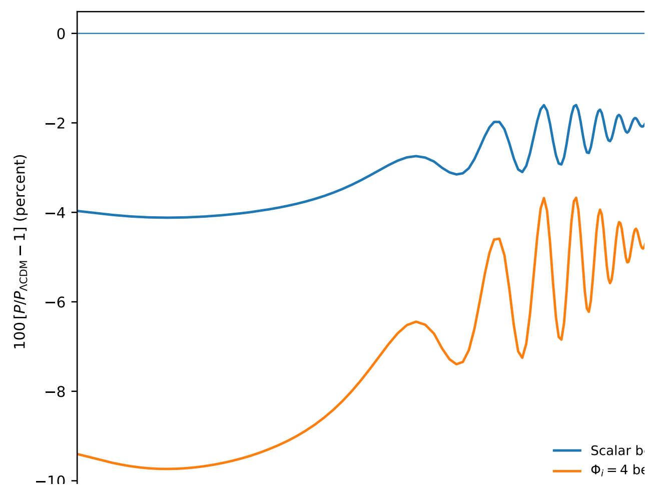

FIG. 10. Native CAMB linear matter-power residual at $z = 0$, using the fixed primordial normalization of the conditional spectra.

C. Three-point dynamic-dependence audit

As a leakage control, a second native suite holds the continuous-best-fit nuisance cosmology fixed at $(\omega_m, \omega_b, h) = (0.140584, 0.0225253, 0.675088)$ and changes only η. This is deliberately not the re-profiled $\Phi_i = 4$ comparison in Table IV. Table XIII gives the reconstructed present equation of state and corrected growth normalization for the near-ΛCDM, MAP, and stronger-thawing points. The monotone change in the outputs, together with distinct file hashes for every pair

TABLE XII. Executed native CAMB growth outputs after correcting the array/redshift ordering.

Model	Ω_m	$\sigma_8(0)$	$f\sigma_8(0)$	S_8
Flat ΛCDM	0.302483	0.803341	0.414534	0.806659
Scalar best	0.308472	0.793894	0.413669	0.805025
$\Phi_i = 4$	0.317409	0.781132	0.412871	0.803476

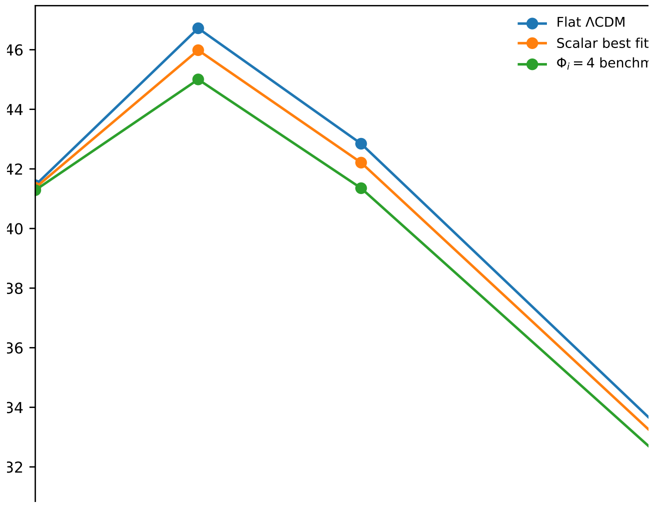

FIG. 11. Native CAMB $f\sigma_8(z)$ predictions for the three matched cosmologies. The difference is small at $z = 0$ and grows toward intermediate redshift for the faster-thawing benchmark.

of $w(a)$ tables, background checks, lensed spectra, matter spectra, and growth tables, rules out accidental reuse of a static MAP history.

The MAP rerun reproduces the frozen non-growth products: the background, lens-potential, matter-power, unlensed-spectrum, and tabulated-w files are byte-identical; the lensed and total spectra agree to a maximum fractional difference of 3.60×10^{-14}. The corrected MAP growth rows are $(z, \sigma_8, f\sigma_8) = (0, 0.793894, 0.413669)$, $(0.5, 0.613570, 0.459841)$, $(1, 0.486070, 0.422108)$, and $(2, 0.335434, 0.320612)$. No reference file was modified to obtain this agreement.

IX. VALIDATION AND NUMERICAL CONTROLS

The following checks support the numerical implementation and native execution.

1. The standalone solver reproduces the reported best-fit w_0 to 2.7×10^{-7} and w_a to 3.3×10^{-5}. For $\Phi_i = 4$, the corresponding differences are 2.7×10^{-7} and 5.3×10^{-6}.

2. The scalar fractions are negligible before recombination. For the continuous best fit, $\Omega_\Phi(z = 50) = 1.79 \times 10^{-5}$, $\Omega_\Phi(z = 300) = 8.14 \times 10^{-8}$, and $\Omega_\Phi(z = 1100) = 1.37 \times 10^{-9}$. The $\Phi_i = 4$ values are 1.90×10^{-5}, 8.62×10^{-8}, and 1.45×10^{-9}.

TABLE XIII. Fixed-cosmology dynamic audit. Only η changes between rows; $w_a = -w_0'$ is the local present tangent and $\sigma_8(0)$ uses the corrected redshift pairing.

η	Φ_i	w_0	w_a	$\sigma_8(0)$
0.050000	40.0000	−0.999622	−0.000547	0.799459
0.403405	4.95780	−0.948924	−0.078481	0.793894
0.500000	4.00000	−0.871684	−0.216327	0.785394

3. Native CAMB succeeded for the ΛCDM control and both scalar points. The scalar audit files record tabulated $w(a)$, $c_s^2 = 1$, and no constant-w, CPL, PPF, or ΛCDM fallback.

4. The dynamic audit separately succeeds at $\eta = (0.05, 0.403405, 0.50)$ with the same physical-density partition and fixed primordial sector. Pairwise-distinct outputs establish actual dependence on η, including a near-ΛCDM point that executes without the earlier density-bookkeeping failure.

5. File identity is fixed by the release manifest. The corrected and preserved pre-fix adapter SHA-256 hashes begin `bf27ac11aa5f` and `e8726e6266cd`, respectively; the unchanged standalone-background hash begins `23a48bdc3204`. Full digests appear in `DYNAMIC_THEORY_MILESTONE.md` and the machine-readable manifests.

6. For $z \leq 3$, the standalone and CAMB expansion rates agree at approximately the part-per-million level (Figure 12). The maximum distance residual through $z = 1100$ is below 3.5×10^{-6}.

7. The largest $H(z)$ residual over the full validation grid is approximately 1.78×10^{-4} at $z = 1100$. The same offset occurs in the ΛCDM control, identifying it as a small radiation/neutrino-background convention difference between the standalone implementation and CAMB rather than an isotonic-specific failure.

8. The independent late-time perturbation solver recovers the ΛCDM full/smooth growth ratio to 4.4×10^{-16} and finds the canonical clustering correction below one part per million at $k = 0.1\,h\,\mathrm{Mpc}^{-1}$ for the posterior-grid maximum. The larger native $P_\mathrm{m}(k)$ and σ_8 differences reported here are therefore background and transfer effects, not direct scalar clustering.

9. An independent check of the archived CAMB 1.6.6 low-ℓ production spectra finds maxima of 0.503% and 1.175% at $\ell = 4$ for the best fit and $\Phi_i =$

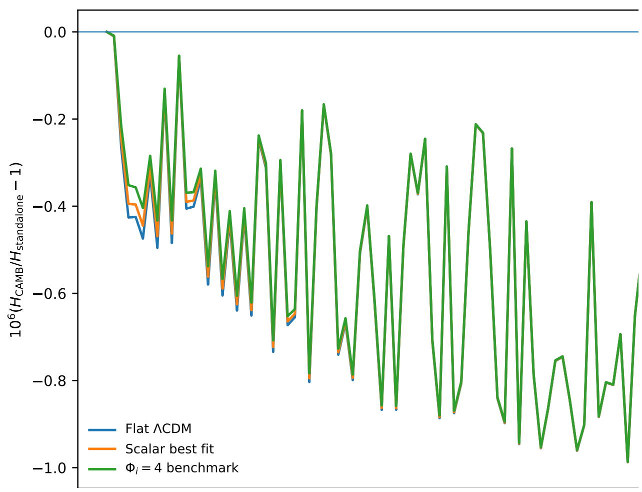

FIG. 12. CAMB-versus-standalone expansion-rate residual for $0 \le z \le 3$. All three cosmologies agree at approximately the part-per-million level in the late-time domain relevant to the scalar dynamics.

4 benchmark, respectively; both are far below the ideal per-multipole cosmic-variance scale.

10. The constant-floor deformation test changes the MAP exit time continuously and retains a finite first $q = 0$ crossing until $\delta_{\rm crit} \simeq 0.2433$; this test quantifies sensitivity to an additive lift but does not establish radiative stability.

11. The future $q = 0$ and first-minimum events were repeated with four tolerance/step combinations spanning rtol $= 10^{-8}$ to 2×10^{-11} and maximum N-steps from 0.02 to 0.003. All digits quoted in Table VIII are stable; the event locations change by less than 10^{-8} fractionally across the sweep.

A. Generative-AI-assisted research workflow

Substantive generative-AI assistance was used throughout the research workflow. The recorded systems and interface model labels were OpenAI ChatGPT 5.6; Fable 5; Google Gemini 3.1 and 3.5; xAI Grok 4.5; Anthropic Claude Opus 5; and the PaperReview.ai Stanford Agentic Review service as accessed in August 2026. They were used iteratively to translate author-specified physical ideas into candidate mathematical formulations; propose and critique derivations; generate and debug Python, likelihood, CAMB-adapter, plotting, and audit code; organize numerical runs; inspect tabulated outputs; synthesize literature; simulate adversarial referee review; and draft and edit the manuscript.

The author chose the research questions, model assumptions, datasets, acceptance criteria, and final claims, and directed every iteration. AI-generated equations, code, citations, and interpretations were treated as provisional until checked against source literature or executable tests. All quoted numerical results come from deterministic code execution and were accepted only after the solver comparisons, native CAMB runs, independent perturbation checks, fixed-cosmology multi-η tests, tolerance sweeps, growth audit, and file/hash controls described in this paper. Cross-model agreement was not treated as independent validation. Generative AI assisted with plotting code, but every figure is a deterministic rendering of archived numerical data; no generative-image model was used.

X. LIMITATIONS AND OPEN SCOPE

The baseline inference does not constitute a full map-level CMB analysis. The quantities A_s, n_s, τ, calibration and foreground nuisances, and $\sum m_\nu$ are not sampled. Consequently the reported σ_8, S_8, lensing, and power-spectrum amplitudes are conditional, not marginalized constraints. The 18-configuration CAMB matrix tests transfer-integrator floor and accuracy sensitivity over 10^{-6}–10^{-8} and `AccuracyBoost` 1.0–2.0, with an archived C_ℓ floor sensitivity below 0.001%. It does not provide a separately serialized floor-by-floor $f\sigma_8$ table. A competitive precision-inference analysis would still sample the primordial and nuisance sectors with genuine Planck/ACT/SPT likelihoods.

The data combination also remains background dominated: it includes no covariance-aware RSD likelihood and no weak-lensing likelihood. The executed $f\sigma_8$ predictions now make such tests straightforward, but incorporating them is a separate inference extension. The upper edge of the declared η prior is demonstrably inactive and the independent-likelihood MCMC agrees with the deterministic posterior at the level relevant here. The DESI subset audit additionally shows that deleting LRG1 or LRG2, or both together, leaves the η posterior width nearly unchanged; the largest mean shift is 0.27 baseline standard deviations and the scalar $\Delta\chi^2$ weakens to about -0.89 when both are removed. Extending the branch definition beyond $\eta = 0.6$, alternative physical measures, supernova-compilation swaps, a direct full-CMB likelihood, and a matched $m^2\phi^2$ thawing baseline remain valuable distinct extensions rather than prerequisites for the present claims.

The finite future times in Table VIII remain calibrated benchmark predictions, not a detection of an acceleration-exit epoch. The posterior push-forward now treats the divergence directly, but it inherits the $\Delta\eta = 0.01$ grid and profiled nuisance cosmology and cannot resolve $\eta \le 0.005$. Its 0.6530% near-boundary contribution is therefore right-censored. The semi-analytic slow-roll expression reproduces the direct exit times at the percent level across the main thawing region and yields the asymptotic $\Delta t_{q=0} \propto \eta^{-2}$ scaling, but it remains an approximation and is not used to generate the quoted posterior push-forward. Under the continuous prior the exact ΛCDM point has zero probability mass; assigning it finite model odds would require an explicit spike-and-slab model prior that is not used here.

The first $q = 0$ crossing is reached before the field touches the zero-energy minimum and is therefore a clean prediction of the canonical pre-minimum dynamics. Three-cycle integration verifies the onset of matterlike oscillations in the canonical theory; the arbitrarily late asymptote additionally uses the standard quadratic-oscillator result. The optional two-measure parent description is less complete: its measure variables become degenerate at the same surface, so a parent-level junction or continuation rule remains necessary. The radiative stability of the coefficient locking remains open: no protective symmetry is claimed here. The explicit constant-floor deformation test now quantifies one narrow version of that sensitivity: $\delta = 10^{-2}$ shifts the MAP exit time by about 0.96%, while the finite MAP exit disappears only near $\delta_{\rm crit} \simeq 0.2433$, corresponding to a residual floor of roughly 3.9% of today's critical density. This does not establish non-renormalization or generic loop stability. Likewise, the absence of a de Sitter event horizon should not be conflated with a proof that all forms of thermodynamic heat death are avoided.

XI. CONCLUSION

Current observations do not select the isotonic-thawing branch over flat ΛCDM, but they do constrain its single dynamical parameter. The continuous best fit moves from the original $\Phi_i = 4$ benchmark to $\Phi_i \simeq 4.96$ and a more frozen present equation of state; its likelihood gain remains too small to overcome standard information penalties, and the ΛCDM boundary remains inside the posterior support. Deleting the DESI LRG1 and LRG2 anisotropic BAO pairs separately or together shifts the reconstructed posterior mean by at most 0.27 baseline standard deviations and weakens rather than strengthens the scalar likelihood gain, so this conclusion is not tied to either tracer bin. Given the exact coefficient locking, the same displacement coordinate η fixes both the present thawing history and the future acceleration-exit time. Across the sampled branch the numerical exit-time map is monotone, and $\Delta t_{q=0}$ diverges as $\eta \to 0$, recovering the nested ΛCDM no-exit limit. Present-day constraints on the branch thus propagate directly into its allowed future histories.

For the current best fit, direct integration gives the end of acceleration at $a = 13.8803$, approximately 60.28 Gyr from now; the original $\Phi_i = 4$ benchmark exits at $a = 2.86230$, approximately 21.13 Gyr from now. Both transitions occur before the first minimum crossing. A semi-analytic scalar-dominated slow-roll reduction, normalized only at the MAP, reproduces fresh direct integrations to about 1.4% or better over $0.20 \le \eta \le 0.45$ and exposes the asymptotic $\Delta t_{q=0} \propto \eta^{-2}$ divergence; the exact integration remains authoritative. The posterior push-forward assigns 37.8% probability to an exit within 100 Gyr and 78.7% within 1 Tyr, with a conditional median of 155.8 Gyr above $\eta_{\rm min} = 0.005$ and a 0.6530% censored near-boundary contribution. These are branch-conditioned probabilities under the stated continuous prior, not evidence that the finite branch is realized or evidence against the nested ΛCDM limit. The subsequent canonical evolution approaches matterlike scalar oscillations, decelerating expansion, and no asymptotic de Sitter event horizon. Under the controlled additive-floor deformation, percent-level lifts perturb the MAP exit time at the percent level and a finite MAP exit persists until $\delta_{\rm crit} \simeq 0.2433$; this quantifies constant-floor sensitivity while leaving the radiative stability of the exact coefficient locking open.

The exact phase-plane tracks establish the branch as Caldwell–Linder thawing in the observational domain, while making clear that its one-parameter curve is not a free CPL plane. The native CAMB execution then fixes the present observable scale: the allowed best-fit branch changes primary TT and EE spectra by only a few parts in 10^3, while the more growth-sensitive lensing BB spectrum reaches a sub-percent difference. The dedicated $2 \le \ell \le 30$ audit finds a larger fractional TT response at the very lowest multipoles—0.503% for the best fit and 1.175% for $\Phi_i = 4$ at $\ell = 4$—but still far below the ideal per-multipole cosmic-variance scale. At fixed primordial normalization, the best fit lowers σ_8 by about 1.2% and the original re-profiled $\Phi_i = 4$ point by about 2.8%. A separate fixed-cosmology three-η audit confirms that the native observables change with the exact scalar history and that the corrected growth ordering reproduces the assessed values. These amplitudes remain conditional on the fixed primordial, nuisance, and neutrino sectors.

The result is therefore a focused, falsifiable phenomenology note: a restrictive non-phantom scalar whose present-day constraints and future acceleration-exit time are linked through one parameter, alongside native CMB spectra and growth outputs. Improved measurements can progressively push the allowed branch toward the nested no-exit ΛCDM limit or, if $\eta > 0$ is preferred, narrow a finite exit-time distribution. The robustness closeout now covers the CAMB floor/accuracy matrix, upper-prior inactivity, independent sampling, DESI LRG-subset deletion, low-ℓ production-spectrum residuals, semi-analytic exit-time scaling, and the controlled constant-floor fate sensitivity. The remaining high-value requests—full map-level CMB likelihood sampling, covariance-aware RSD and weak-lensing inclusion, supernova-compilation swaps, a dedicated floor-by-floor growth table, comparison with matched alternative thawing potentials, Bayesian evidence, and a completed radiatively stable parent mechanism—are larger-scope inference or theory extensions rather than repairs to the present calculation.

DATA AVAILABILITY STATEMENT

The observational inputs used in this work are publicly available through the DESI DR2, Pantheon+, and CMBComp sources cited in the text. The numerical code and derived data products supporting the figures, tables, and validation audits are available from the author upon reasonable request.

ACKNOWLEDGMENTS

The substantive generative-AI use, including named systems and versions, research tasks, human direction, and verification, is documented in Sec. IX A. The model concept, project lineage, and research direction originated with the author, who selected the assumptions and claims and assumes sole responsibility for all calculations, interpretations, errors, and text. AI systems are not authors, and their agreement was not treated as scientific validation.

[1] J. F. Cariñena, A. M. Perelomov, M. F. Rañada, and M. Santander, A quantum exactly solvable non-linear oscillator related with the isotonic oscillator, J. Phys. A: Math. Theor. **41**, 085301 (2008), arXiv:0711.4899.
[2] E. I. Guendelman and A. B. Kaganovich, Some cosmological applications of two measures theory, arXiv:gr-qc/0403017 (2004).
[3] C. Clifton, Twist-angle control of driven electromagnetic response in van der Waals heterostructures: Alpha-Lattice preregistered three-rung verification ladder, aiXiv:aixiv.260203.000003 (2026).
[4] C. Clifton, Procedural Vacuum Breakdown: A three-form completion of the gravitational core, aiXiv:aixiv.260612.000001 (2026).
[5] M. S. Turner, Coherent scalar-field oscillations in an expanding universe, Phys. Rev. D **28**, 1243 (1983), doi:10.1103/PhysRevD.28.1243.
[6] R. R. Caldwell and E. V. Linder, The limits of quintessence, Phys. Rev. Lett. **95**, 141301 (2005), arXiv:astro-ph/0505494, doi:10.1103/PhysRevLett.95.141301.
[7] E. J. Copeland, M. Sami, and S. Tsujikawa, Dynamics of dark energy, Int. J. Mod. Phys. D **15**, 1753 (2006), arXiv:hep-th/0603057.
[8] C.-P. Ma and E. Bertschinger, Cosmological perturbation theory in the synchronous and conformal Newtonian gauges, Astrophys. J. **455**, 7 (1995), arXiv:astro-ph/9506072.
[9] J. Weller and A. M. Lewis, Large-scale cosmic microwave background anisotropies and dark energy, Mon. Not. R. Astron. Soc. **346**, 987 (2003), arXiv:astro-ph/0307104.
[10] A. Lewis, A. Challinor, and A. Lasenby, Efficient computation of cosmic microwave background anisotropies in closed Friedmann–Robertson–Walker models, Astrophys. J. **538**, 473 (2000), arXiv:astro-ph/9911177.
[11] DESI Collaboration, DESI DR2 Results II: Measurements of baryon acoustic oscillations and cosmological constraints, Phys. Rev. D **112**, 083515 (2025), arXiv:2503.14738, doi:10.1103/tr6y-kpc6.
[12] D. Brout *et al.*, The Pantheon+ analysis: Cosmological constraints, Astrophys. J. **938**, 110 (2022), arXiv:2202.04077, doi:10.3847/1538-4357/ac8e04.
[13] Z. Wang and Y. Wang, Lossless compression of cosmological information from type Ia supernova distance measurements, arXiv:2605.19188 (2026).
[14] A. Srivastav, P. Bansal, and D. Huterer, CMBComp: A simple and accurate compressed CMB likelihood for dark energy, curvature, and massive neutrinos, arXiv:2606.18455, version 2 (2026).
[15] A. Lewis and S. Bridle, Cosmological parameters from CMB and other data: A Monte Carlo approach, Phys. Rev. D **66**, 103511 (2002), arXiv:astro-ph/0205436.
[16] D. Shlivko, Thawing quintessence: Priors, evidence, and likely trajectories, Phys. Rev. D **113**, 083518 (2026), arXiv:2512.20832, doi:10.1103/xt9s-jx2s.
[17] G. Payeur, E. McDonough, and R. Brandenberger, Do observations prefer thawing quintessence?, Phys. Rev. D **111**, 123541 (2025), arXiv:2411.13637, doi:10.1103/bggr-61nr.
[18] K. Lodha *et al.* (DESI Collaboration), Extended dark energy analysis using DESI DR2 BAO measurements, Phys. Rev. D **112**, 083511 (2025), arXiv:2503.14743, doi:10.1103/w4c6-1r5j.
[19] R. de Souza, G. Rodrigues, and J. Alcaniz, Thawing quintessence and transient cosmic acceleration in light of DESI, Phys. Rev. D **112**, 083533 (2025), arXiv:2504.16337, doi:10.1103/2tjq-dtbc.
[20] C. García-García, P. G. Ferreira, and W. J. Wolf, The status of single scalar field dark energy, arXiv:2607.07777 (2026).