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

6 August 2026

Abstract

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

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

has the harmonic-plus-inverse-square structure of an isotonic oscillator, but with coefficient locking that completes a non-negative square and fixes an exact zero-energy minimum at $\Phi = 2$. In the positive weight variable $Q = \Phi^2/4 = e^{\psi}$, the same shape is $V/\Lambda_0 = Q + Q^{-1} - 2 = 2(\cosh\psi - 1)$. A branch-specific two-measure construction can generate this completed form algebraically, but the observational calculation requires only ordinary canonical quintessence and is not presented as a UV completion.

The model is parameterized by the bounded displacement coordinate $\eta = 2/\Phi_i$, with the flat-ΛCDM limit nested at $\eta = 0$. 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 $= +1.3525$ do not select the scalar model.

The exact histories lie inside the Caldwell–Linder thawing band at the present epoch: $(1 + w_\Phi, w') = (0.0510757, 0.0784478)$ for the continuous best fit and $(0.124912, 0.212846)$ for $\Phi_i = 4$, where $w' = dw_\Phi/d\ln a$. Thus the isotonic branch is a bona fide thawing model over the domain in which that bound applies, although its exact one-parameter locus is a curve rather than a free CPL plane.

Direct future integration of the same field equations, rather than extrapolation of a CPL proxy, gives a second transition $q = 0$ for the continuous best fit at $a = 13.8803$, 60.28 Gyr from the present. The original $\Phi_i = 4$ benchmark exits acceleration earlier, at $a = 2.86230$ and 21.13 Gyr. In both cases acceleration ends before the first crossing of the zero-energy minimum. Near that minimum the potential is quadratic, so the canonical field enters damped oscillations with cycle-averaged $\langle w_\Phi \rangle \to 0$; the asymptotic expansion is decelerating rather than de Sitter and has no permanent de Sitter event horizon. This is a branch-conditioned prediction, not an observational detection of transient acceleration, because the eternal-acceleration ΛCDM boundary remains inside the posterior support.

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 $30 \le \ell \le 2500$, 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.

Keywords: thawing quintessence; transient cosmic acceleration; isotonic potential; CMB anisotropies; structure growth; DESI DR2.

1 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 [4, 5]. The central question here is sharper than whether one more scalar potential can mimic a time-varying equation of state. We ask whether a single canonical branch can satisfy three requirements simultaneously: remain close enough to ΛCDM to survive current background and compressed-CMB data, produce calculable perturbation and Boltzmann observables without phantom crossing, and end the present accelerated era because its own potential has no residual vacuum-energy floor.

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.

There is also a conditional upstream motivation. In a branch-specific two-measure Palatini construction, an algebraic measure constraint can produce

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

with $V_1 \propto e^{\psi}$, $V_2 \propto e^{\psi}$, and the integration-constant branch $M = -c$ yielding exactly $V_{\rm eff}/\Lambda_0 = e^{\psi} + e^{-\psi} - 2$ after canonical normalization [2]. This candidate origin explains why the reciprocal term and additive zero can arise together rather than being independently fitted. It is not required by the present inference, is not claimed as a unique derivation or UV completion, and leaves matter coupling, radiative stability, displacement selection, and parent-variable continuation through the minimum as open problems.

The source background study reported a reproducible trajectory for $\Phi_i = 4$ but deferred perturbations, growth, likelihood-level inference, and the calibrated future fate. A first inference-ready revision supplied the exact background, canonical perturbation equations, a deterministic posterior, and a CAMB adapter, but did not execute native Boltzmann spectra. The present revision retains the native CAMB execution, exact future integration, Caldwell–Linder phase-plane test, quantitative compression audits, posterior exit-time push-forward, and numerical post-minimum check. It additionally seals a three-point dynamic-dependence audit, corrects the native growth-table writer, and preserves both pre-fix and corrected source hashes.

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.

2 Canonical model and inference parameterization

The physical canonical field is $\phi_{\rm phys} = M_{\rm Pl}\Phi$; 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, no metric potential Φ appears in this work. Write

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

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

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

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 $0.06\,\mathrm{eV}$ thermal neutrino. The exact closure and scalar equations are

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

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

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

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

The scalar equation of state is

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

The sampled variables are

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

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

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

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.

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

Table 1. 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
$\mathscr{F}(\Phi)$	Dimensionless potential shape, V/Λ_0	Canonical model
Q	Positive weight coordinate $Q = \Phi^2/4 = e^{\psi}$	Abstract
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 $\Lambda_0/(3M_{\rm Pl}^2H_0^2)$	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
$\mathscr{S}(T)$	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

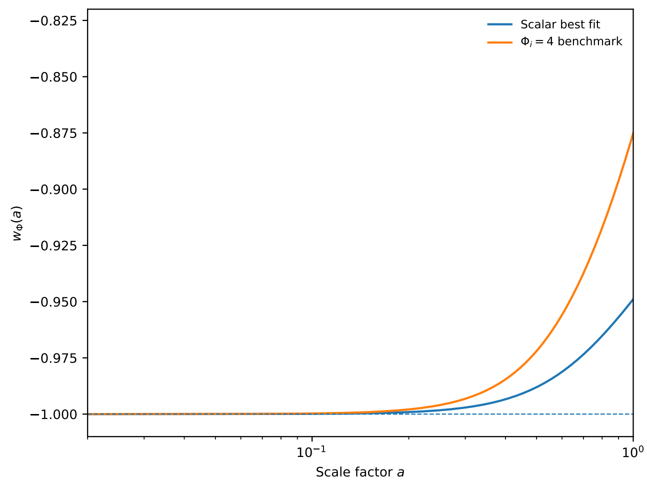

Figure 1. Exact scalar equation-of-state histories reconstructed from the standalone field equations. Both trajectories remain non-phantom and approach $w = -1$ on the frozen early-time surface.

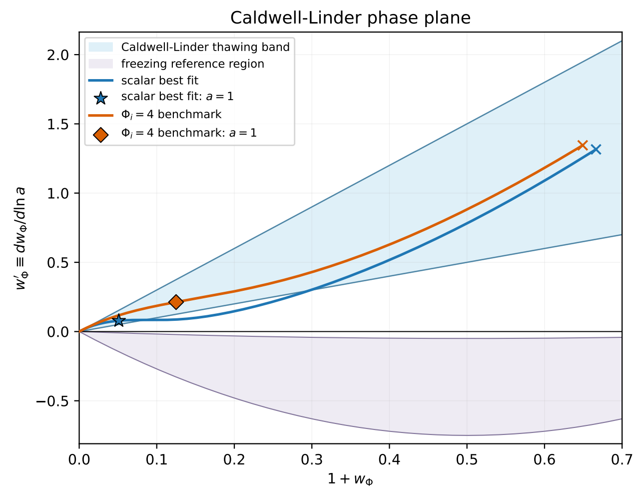

Figure 2. Exact isotonic trajectories in the Caldwell–Linder plane, from the frozen epoch to the first future acceleration exit. The present epoch is marked separately on each curve and crosses mark the $q = 0$ endpoints. The positive-w' thawing wedge and the negative-w' freezing reference region are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

3 Exact Caldwell–Linder phase-plane test

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

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

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 $3w_\Phi(1 + w_\Phi) \le w' \le 0.2\,w_\Phi(1 + w_\Phi)$ [4].

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

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

and the $\Phi_i = 4$ coordinates are

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

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 $0 \le z \le 3$. 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.

4 Canonical scalar perturbations

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

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

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 [6, 7]

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

The scalar density and pressure perturbations are

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

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

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

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

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.

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

4.1 Native CAMB implementation

CAMB provides a tabulated dark-energy equation-of-state interface and a rest-frame sound-speed parameter [8]. 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

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

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. A second floor-convergence run remains a recommended code-level check before a final journal archive.

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.

5 Data and likelihood

The posterior combines three statistically independent blocks.

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 [9]. The drag horizon is evaluated consistently from a CAMB calibration grid in (ω_m, ω_b) for the fixed neutrino sector.

Pantheon+. The supernova likelihood is the 11-knot Gaussian compression of $\log r_p(z)$ [10, 11]. The compression removes the absolute distance normalization, so no supernova absolute-magnitude nuisance parameter is sampled.

CMBComp CMB-3w. We use the correlated vector (R, ℓ_A, ω_b) from CMBComp [12]. 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.

Table 2. Exact dark-energy fractions. For the scalar rows $\Omega_{\rm DE} = \Omega_\Phi$; for the flat-$\Lambda$CDM 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

5.1 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 2. 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 $N_{\rm eff} = 3.044$, $\sum m_\nu = 0.06\,\text{eV}$ 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 3. 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

Table 4. Posterior summaries under the stated uniform priors.

Parameter	Mean $\pm$ 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

6 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 [13].

Table 3 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,

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

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 5. 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

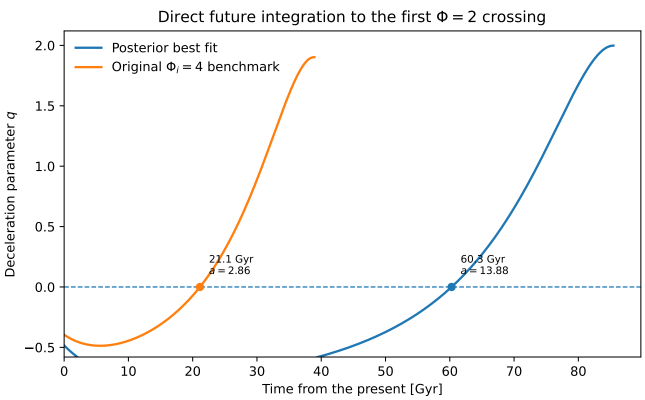

Figure 3. Deceleration parameter from direct future integration. Each curve terminates at its first crossing of the zero-energy minimum. The $q = 0$ roots occur before $\Phi = 2$, so the finite-time end of acceleration does not depend on how a speculative parent construction is continued through its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

7 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

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

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 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.

7.1 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

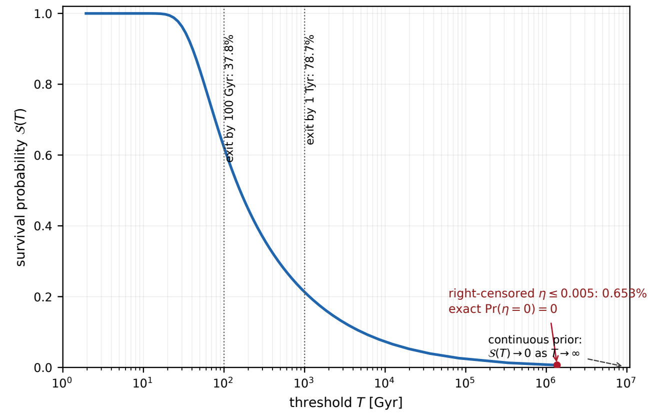

Figure 4. Posterior survival function for the exact acceleration-exit time. The arrow at the long-time edge denotes the unresolved $\eta \leq 0.005$ half-cell, not a literal ΛCDM point mass: under the continuous prior the exact $\Pr(\eta = 0)$ is zero and the true survival tends to zero as $T \to \infty$.

(ω_m, ω_b, h) at that node, and aggregate the archived order-7 Gauss–Hermite weights at fixed η. Within each cell we use the piecewise-linear posterior density that reproduces the trapezoidal integral and linearly interpolate the profiled nuisance cosmology before each exact threshold solve; literal node sums are retained as grid controls. The exit time is strictly monotone decreasing over $0.01 \leq \eta \leq 0.60$ and diverges as $\eta \to 0$. This monotonicity permits the weighted η posterior to be mapped to the survival function $\mathscr{S}(T) \equiv \Pr(\Delta t_{q=0} > T)$ 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 $\eta_{\min} = 0.005$, half the posterior grid spacing, and treat $\Pr(\eta \leq \eta_{\min}) = 0.006530$ as an unresolved right-censored near-boundary contribution. The resolved push-forward gives

$$\Pr(\Delta t_{q=0} < 100\,\text{Gyr}) = 0.3782, \quad \Pr(\Delta t_{q=0} < 1\,\text{Tyr}) = 0.7870. \tag{28}$$

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

$$V(\Phi) = \Lambda_0 \left[\varphi^2 + \mathscr{O}(\varphi^3)\right] = \tfrac{1}{2} M_{\rm Pl}^2 m_\Phi^2 \varphi^2 + \mathscr{O}(\Lambda_0 \varphi^3),$$
$$m_\Phi^2 = \frac{2\Lambda_0}{M_{\rm Pl}^2}, \qquad \frac{m_\Phi^2}{H_0^2} = 6\alpha. \tag{29}$$

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

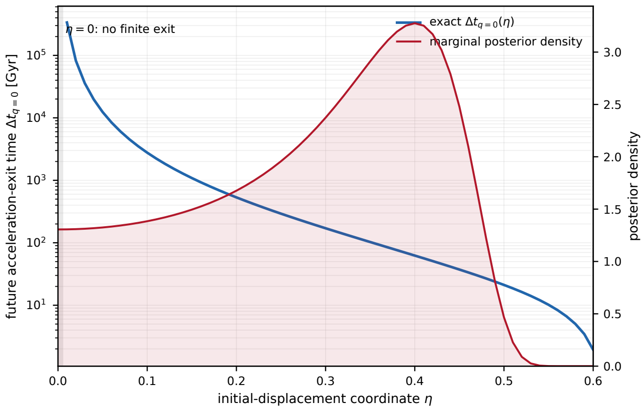

Figure 5. Exact acceleration-exit time across the positive η grid, with the marginal posterior density on the second axis. The monotone divergence toward the nested $\eta = 0$ boundary motivates the survival-function and censoring convention used in Figure 4.

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 $\rho_\Phi \propto a^{-3}$ [3]. 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,

$$\int_t^\infty \frac{dt'}{a(t')} \sim \int_t^\infty dt'\, t'^{-2/3} = \infty, \tag{30}$$

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.

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.

8 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 $\ell_{\max} = 3000$ at accuracy setting 1.0. The scalar runs used the exact tabulated histories after the early frozen-epoch regularization described above. Table 6 summarizes the executed residuals relative to the separately best-fitting flat-ΛCDM cosmology.

Table 6. 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 $\max|D^{TE}_{\ell,\Lambda}|$ to avoid zero-crossing singularities. The $P_m(k)$ entry is the maximum absolute linear-power residual for $k \le 0.5\,h\,\mathrm{Mpc}^{-1}$ at $z = 0$.

Model	TT [%]	EE [%]	lensing BB [%]	TE norm. [%]	$P_m(k, z=0)$ [%]
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

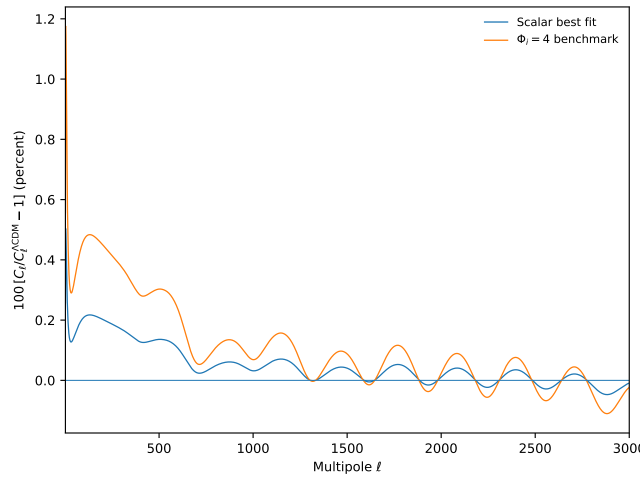

Figure 6. Native CAMB lensed TT residuals relative to the matched flat-ΛCDM run. The continuous scalar best fit remains within approximately 0.22% for $30 \le \ell \le 2500$; the original $\Phi_i = 4$ trajectory reaches approximately 0.48%.

The primary TT residual is sub-percent across the measured acoustic range (Figure 6). EE follows the same pattern at slightly larger amplitude, while lensing BB is more responsive to the accumulated growth suppression (Figure 7). 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.

The linear matter-power residual is shown in Figure 8. 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.

8.1 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

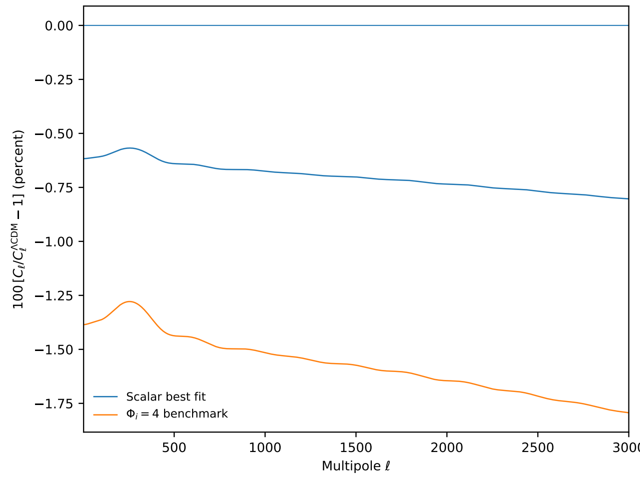

Figure 7. Native CAMB lensing-BB residuals relative to matched flat ΛCDM. The larger response compared with primary TT and EE reflects the integrated late-time growth and lensing-potential difference.

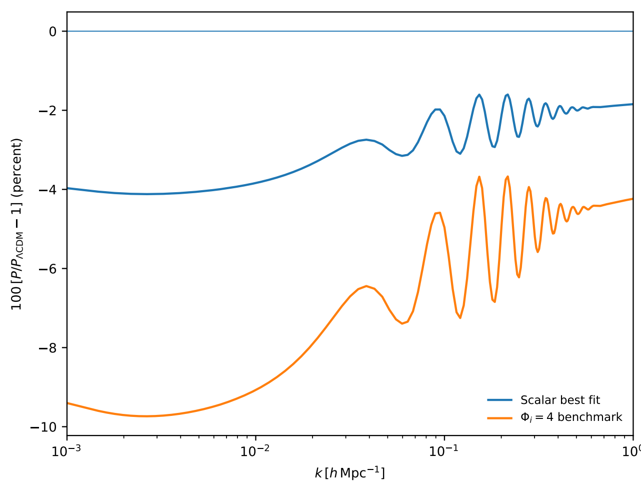

Figure 8. Native CAMB linear matter-power residual at $z = 0$, using the fixed primordial normalization of the conditional spectra.

pairing before producing Table 7 and Figure 9; 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

$$\Lambda\mathrm{CDM} : (0.414534, 0.467200, 0.428453, 0.324384), \tag{31}$$

$$\text{scalar best} : (0.413669, 0.459841, 0.422108, 0.320612), \tag{32}$$

$$\Phi_i = 4 : (0.412871, 0.450034, 0.413528, 0.315444). \tag{33}$$

These predictions provide the direct inputs for a future

Table 7. 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

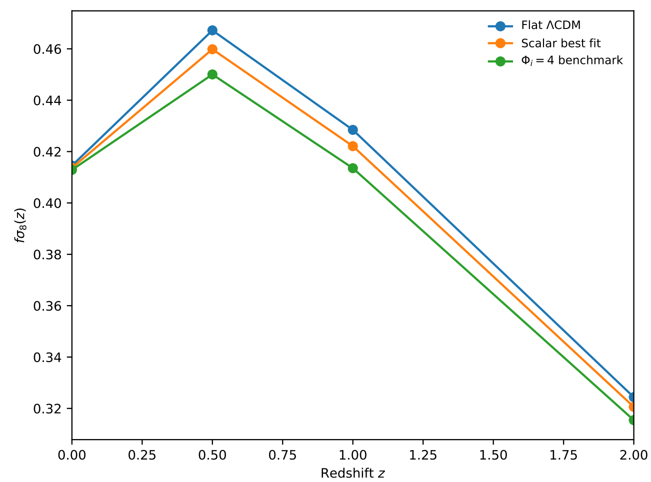

Figure 9. Native CAMB $f\sigma_8(z)$ predictions for the three matched cosmologies. The difference is small at $z = 0$ and grows toward intermediate redshift for the faster-thawing benchmark.

covariance-aware RSD likelihood, but no RSD data are included in the present posterior.

8.2 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 3. Table 8 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 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.

9 Validation and numerical controls

The following checks support the reconstruction and native execution.

Table 8. 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

1. The reconstructed 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}.

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 reconstructed standalone and CAMB expansion rates agree at approximately the part-per-million level (Figure 10). 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 clean-room standalone reconstruction 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 $P_m(k)$ and σ_8 differences reported here are therefore background and transfer effects, not direct scalar clustering.

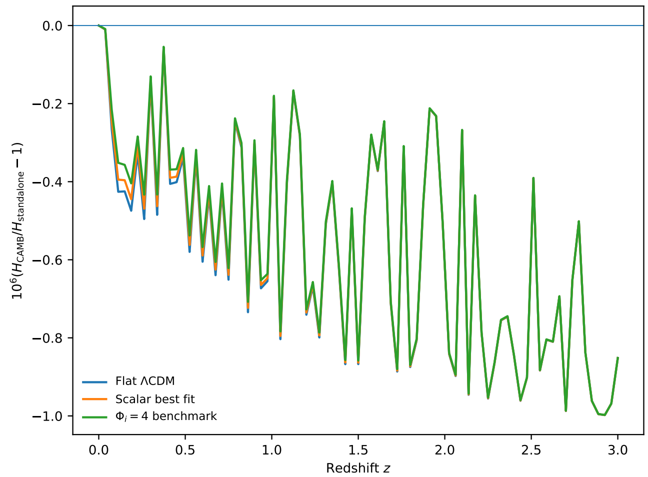

Figure 10. CAMB-versus-standalone expansion-rate residual for $0 \leq z \leq 3$. All three cosmologies agree at approximately the part-per-million level in the late-time domain relevant to the scalar dynamics.

9. The future $q = 0$ and first-minimum events were repeated with four tolerance/step combinations spanning $r_{\mathrm{tol}} = 10^{-8}$ to 2×10^{-11} and maximum N-steps from 0.02 to 0.003. All digits quoted in Table 5 are stable; the event locations change by less than 10^{-8} fractionally across the sweep.

10 Limitations and publication status

This revision resolves the previous absence of native Boltzmann spectra, transfer observables, a posterior-calibrated future integration, and an explicit multi-η dynamic-dependence test. It does not convert the CMBComp posterior into 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. A competitive precision-inference analysis would sample these sectors with genuine Planck/ACT/SPT likelihoods and should repeat the native runs across a documented $1 + w$ floor and CAMB `AccuracyBoost` convergence table. The three-point audit validates the chosen 10^{-6} floor and accuracy-1.0 pipeline dynamically; it is not a substitute for that convergence matrix.

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. Prior sensitivity to the upper limit on η, alternative physical measures, an MCMC or nested-sampling cross-check, supernova-compilation swaps, and a matched $m^2\phi^2$ thawing baseline remain valuable robustness tests.

The finite future times in Table 5 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 $\eta \leq 0.005$. Its 0.6530% near-boundary contribution is therefore right-censored. 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 is also open. 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.

The dynamical source accompanying this revision remains a clean-room reconstruction from the printed equations and numerical benchmarks after the original working directory was lost. The earlier deterministic-posterior archive recovered for v4 supplies the normalized quadrature weights and profiled η grid used by the push-forward, but it does not resolve the provenance of the pre-loss source package. Revision v5 preserves the original native outputs and the pre-fix writer rather than rewriting historical evidence; the corrected writer and versioned dynamic milestone are shipped alongside them. The preparation surface, inputs, hashes, and machine-readable audits are disclosed explicitly. This is adequate for an auditable phenomenology revision; the clean-room provenance caveat remains a journal-submission limitation.

11 Conclusion

The isotonic-thawing branch survives a current background-plus-compressed-CMB confrontation but is not preferred over flat ΛCDM. 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 one-parameter likelihood gain is too small to overcome standard information penalties, and the ΛCDM boundary remains inside the posterior support.

The model nevertheless makes a sharper physical prediction than a generic exercise in numerical rigor. Because the isotonic coefficients complete an exact square with $V_{\min} = 0$, the finite scalar branch cannot settle into eternal de Sitter acceleration. Direct integration gives the best-fit 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. 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 $\eta_{\min} = 0.005$ and a 0.6530% censored near-boundary contribution. These are conditional branch probabilities, not evidence against the nested ΛCDM limit. The subsequent canonical evolution approaches matterlike scalar oscillations, decelerating expansion, and no asymptotic de Sitter event horizon.

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. 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 mathematically motivated non-phantom scalar with a data-calibrated but non-detected transient-acceleration prediction, native CMB spectra, and growth outputs. Full CMB likelihood sampling, floor and accuracy convergence, RSD and weak-lensing inclusion, sampling-method and supernova-compilation cross-checks, low-ℓ ISW residuals, and a completed, radiatively stable parent mechanism define the next tier.

Code and data availability

The revision bundle contains the standalone background solver, isotonic potential, thermal-neutrino table, corrected native CAMB adapter, preserved pre-fix adapter, automated checks, exact $w(a)$ tables, executed spectra and transfer outputs, and machine-readable audits. Standalone scripts generate the Caldwell–Linder track, compression-fidelity diagnostics, the weighted exit-time survival function, the $\Delta t_{q=0}(\eta)$ curve, and the three-cycle oscillatory check. Their JSON/CSV results, posterior weights, profiled η grid, and all ten vector figures are included and hashed. The sealed dynamic milestone is included both expanded and as `Dynamic_Theory_Milestone_20260806_v2_1_metadatafix.zip`, whose SHA-256 is `8a0e6fe05634c09049a3786da84ef213561044d273f7a2ba0726579d898a2db7`. The native execution used CAMB 1.6.6 in Python 3.12 with $\ell_{\max} = 3000$ and accuracy setting 1.0.

References

[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," Journal of Physics A: Mathematical and Theoretical **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] M. S. Turner, "Coherent scalar-field oscillations in an expanding universe," Physical Review D **28**, 1243–1247 (1983), doi:10.1103/PhysRevD.28.1243.

[4] R. R. Caldwell and E. V. Linder, "The limits of quintessence," Physical Review Letters **95**, 141301 (2005), arXiv:astro-ph/0505494, doi:10.1103/PhysRevLett.95.141301.

[5] E. J. Copeland, M. Sami, and S. Tsujikawa, "Dynamics of dark energy," International Journal of Modern Physics D **15**, 1753–1936 (2006), arXiv:hep-th/0603057.

[6] C.-P. Ma and E. Bertschinger, "Cosmological perturbation theory in the synchronous and conformal Newtonian gauges," Astrophysical Journal **455**, 7–25 (1995), arXiv:astro-ph/9506072.

[7] J. Weller and A. M. Lewis, "Large-scale cosmic microwave background anisotropies and dark energy," Monthly Notices of the Royal Astronomical Society **346**, 987–993 (2003), arXiv:astro-ph/0307104.

[8] A. Lewis, A. Challinor, and A. Lasenby, "Efficient computation of cosmic microwave background anisotropies in closed Friedmann–Robertson–Walker models," Astrophysical Journal **538**, 473–476 (2000), arXiv:astro-ph/9911177.

[9] DESI Collaboration, "DESI DR2 Results II: Measurements of baryon acoustic oscillations and cosmological constraints," Physical Review D **112**, 083515 (2025), arXiv:2503.14738.

[10] D. Brout et al., "The Pantheon+ analysis: Cosmological constraints," Astrophysical Journal **938**, 110 (2022), arXiv:2202.04077, doi:10.3847/1538-4357/ac8e04.

[11] Z. Wang and Y. Wang, "Lossless compression of cosmological information from type Ia supernova distance measurements," preprint (2026), arXiv:2605.19188.

[12] A. Srivastav, P. Bansal, and D. Huterer, "CMBComp: A simple and accurate compressed CMB likelihood for dark energy, curvature, and massive neutrinos," preprint, version 2 (2026), arXiv:2606.18455.

[13] A. Lewis and S. Bridle, "Cosmological parameters from CMB and other data: a Monte Carlo approach," Physical Review D **66**, 103511 (2002), arXiv:astro-ph/0205436.