A ladder of conclusions on the history of dark energy: an inception at $z \simeq 4$ preferred by Planck, DESI DR2, and Pantheon+ data

Yaroslav Ryabov

November 2026

Abstract

We confront a family of dark-energy histories parametrized by the cosmological event horizon, $\rho_{\rm de}(t) = f(t)/R_E^2(t)$, with Planck 2018 (TT, TE, EE), DESI DR2 baryon acoustic oscillations — including the 2026 Lyman-α full-shape update — and the Pantheon+ supernova Hubble diagram, using ΛCDM and the w_0–w_a tangent as external controls. The data build a ladder of conclusions: (1) a fit with dark energy is decisively better than one without; (2) an evolving density beats a constant by $\Delta\chi^2 = 4$–9, above the ± 3 cross-pipeline systematic; (3) dark energy that is *born* — absent at high redshift, activating at $z_c = 4.2^{+1.0}_{-0.3}$ ($z_c > 3.4$ at 95%) — beats every eternal form: an eternal screen loses by $+37$ at equal parameter count, and the born Avrami form leads the full three-probe battery, with $w_0 = -0.95$ to -0.96 pinned along the residual degeneracy of the fit (marginal posterior -0.96 ± 0.03) and an early tail $\Omega_{\rm de}(z{=}2.33) = 2.6$–$4.3\%$. (4) The data cannot select a model of dark-energy inception within the top class: the leading fits differ by at most 3.3 units of χ^2, comparable to the ± 3 cross-pipeline systematic, leaving open *how* the birth proceeded (nucleation versus global ordering) and *where* the component resides (a condensate on the horizon screen versus a bulk component reproducing the same density history). We record a mild preference for the horizon reading on grounds of interpretive economy, not of χ^2.

Contents

1 Introduction and motivation **2**

2 Data, priors, and fitting approach **3**

3 Models with physical motivation **5**

4 Fitting results **9**
4.1 Baseline battery: CMB + BAO 9
4.2 Second battery: the Lyman-α full-shape update 10
4.3 Third battery: adding the supernova Hubble diagram 11
4.4 Profile-likelihood uncertainties 13

5 The ladder of conclusions **17**

6 Discussion and open questions **19**

Appendix A: The de Sitter attractor and solver integrity **26**

Appendix B: Complete best-fit parameters **28**

Appendix C: The data behind the batteries, and the posterior **31**

References cited in the appendices **35**

1 Introduction and motivation

Two developments frame this work. The first is observational: the era of the Dark Energy Spectroscopic Instrument (DESI [1]) has turned "is dark energy dynamical?" from a philosophical preference into a data question, with combined analyses of the cosmic microwave background (CMB [2]), baryon acoustic oscillations (BAO [3]), and type-Ia supernovae (SNe Ia [4, 5]) reporting consistent preferences for an evolving equation of state over a cosmological constant [3]. The second is a question that the standard w_0–w_a language — the Chevallier–Polarski–Linder (CPL) tangent $w(a) = w_0 + w_a(1-a)$ for the dark-energy equation of state $w \equiv P_{\rm de}/\rho_{\rm de}$ (pressure over density) [6, 7] — cannot even pose cleanly: granted that the dark-energy density evolves, *what is the shape of its history*? In particular — has dark energy existed at all epochs, merely changing its value, or was there an epoch before which it did not exist? A tangent parameterization describes the neighborhood of today; it carries no notion of a beginning.

To pose the question operationally we need a family of histories in which "birth" is a natural, preferably one-parameter, operation — so that the value of the birth itself can be judged at near-equal parameter count — together with controls that isolate each logical step. We take the family

$$\rho_{\rm de}(t) \;=\; \frac{f(t)}{R_E^2(t)}, \qquad R_E(t) = a(t)\int_t^\infty \frac{c\,dt'}{a(t')}, \tag{1}$$

in which the density history is *parametrized by the future event horizon* $R_E(t)$: here $a(t)$ is the scale factor, t cosmic time, c the speed of light, and $R_E(t)$ is the proper radius of the cosmological event horizon [8] (the subscript stands for "event") — the farthest distance light emitted today will ever travel. Throughout the paper $H(t) \equiv \dot{a}(t)/a(t)$ is the Hubble rate, an overdot denoting the derivative with respect to t; $h \equiv H_0/(100\ \mathrm{km\,s^{-1}\,Mpc^{-1}})$ the reduced Hubble constant (H_0 its present value), z the redshift ($1+z = 1/a$, with $a_0 = 1$ today), and we use natural units $\hbar = k_B = 1$ (and $c = 1$ except where displayed). Two physical facts motivate this choice of scale. First, the event horizon is the one surface in cosmology on which semiclassical gravity places genuine thermodynamics — a temperature (defined below) and an entropy proportional to its area [9] — so it can act as a physical *screen* (a boundary carrying thermodynamic degrees of freedom) rather than a bookkeeping device. Second, $1/R_E^2(t)$ is the holographic ceiling for the vacuum energy: effective field theory in a region of size L whose states are required not to collapse into a black hole can carry an energy density of at most $\rho \sim M_{\rm pl}^2/L^2$ [10] ($M_{\rm pl}$ the Planck mass), and among the candidate cosmological lengths L only the future event horizon lets the saturated bound drive accelerated expansion [11, 12]. Equation (1) therefore reads: the dark-energy density is the fraction $f(t)$ of the maximal density that a horizon-sized region can hold. The parametrization fixes the scale of the density, not its residence; whether the component lives on the horizon or in the volume is taken up in Sec. 6. Thus, to address the central question of this work, we propose to explore the following progression of forms of the dimensionless filling function $f(t)$:

$f(t) \equiv 0$	no dark energy (control);
$f(t) \equiv c^2_{\rm HDE} = \text{const}$	eternal holographic screen (canonical holographic dark energy, HDE [11]; control) — $c_{\rm HDE}$ is the standard dimensionless HDE amplitude, written plain c in the HDE literature and subscripted here to keep it distinct from the speed of light;
$f(t) = s\,\alpha(I(t))$	causal-clock activation after a critical epoch;
$f(t) = s\,\alpha(T(t))$	thermodynamic activation as the horizon cools,

where s is a constant amplitude and $\alpha \in [0, 1]$ is the filling fraction of the screen. Throughout the paper α denotes this one physical quantity (not to be confused with the scale factor $a(t)$); the argument displayed — the elapsed causal time $I(t)$ of Sec. 3, or the horizon temperature $T(t)$ defined below — indicates which activation law drives it, and the laws themselves are specified in Sec. 3.

Three properties recommend this parameterization — and we stress that it is a parameterization, not a claim about the physical carrier. *First*, it makes birth natural: the Gibbons–Hawking temperature of the horizon [9], $T_{\rm dS}(t) = \hbar c/(2\pi k_B R_E(t))$, falls monotonically as the universe expands, so "activation when a critical temperature is crossed" is a one-parameter statement. (The subscript recalls the de Sitter origin of the expression, where it is exact; here it is the instantaneous generalization to the evolving event horizon. We write $T^0_{\rm dS}$ for its value today.) The horizon supplies the second possible clock as well: for a birth that proceeds by growth rather than by equilibrium, the elapsed time in horizon light crossings, $I(t)$, is the natural causal variable: unlike the cosmic time t, which counts seconds, I counts how many times a light signal could have traversed the horizon since birth — the only clock a process spreading causally across a horizon-sized region can follow. Thus both thermodynamic characters a transition can have (Sec. 3) are expressible inside one and the same parameterization. *Second*, it gives the measured numbers names (a critical temperature in units of today's horizon temperature; a front speed in units of light crossings; an amplitude fixed by an exact attractor identity, Sec. 3). *Third*, it contains its own null control: the same capacity $1/R^2_E(t)$ *without* a birth is precisely Li's holographic dark energy, so the value of the birth itself can be isolated at equal parameter count.

The program of the paper is then a data analysis, not a model advocacy: we fit the progression of dark-energy histories, together with ΛCDM and the CPL tangent as external controls, to three successively larger probe combinations, and we report which steps of the progression the data can separate and which they cannot. The main results are four conclusions stated in Sec. 5; everything else — including the specific best-fitting forms — is instrumentation.

2 Data, priors, and fitting approach

To execute the program of this work we use the following data.

CMB. Planck 2018 `plik_lite` nuisance-marginalized bandpowers [2] with the full covariance: the complete 613-bin vector over the temperature auto-spectrum (TT, 215 bins), the temperature–polarization cross-spectrum (TE, 199), and the E-mode polarization auto-spectrum (EE, 199), at multipoles $\ell \geq 30$. The choice of `plik_lite` over a full multi-nuisance high-ℓ likelihood is deliberate: it is Planck's own compression of the 2018 (PR3) release — foregrounds and instrumental nuisances pre-marginalized down to a single calibration parameter $A_{\rm Planck}$, the entire likelihood carried by two published files (bandpowers and covariance) — so it can be evaluated millions of times inside our own minimizer, shipped whole with the release scripts, and compared number for number with the bulk of the dark-energy literature built on the same baseline. The likelihood reader is

self-validating: before any fit it evaluates the official Planck-2018 best-fit point and refuses to run unless χ^2 reproduces the value independently verified in the `cobaya` framework [13] ($\chi^2 = 206.37$ on the 215-bin TT block); binning conventions and the polarization block ordering are auto-calibrated against this anchor rather than assumed. A cross-pipeline systematic was measured by refitting the Avrami–ΛCDM pair on bandpowers from Planck's fourth public release (PR4/NPIPE maps), through the CamSpec TTTEEE likelihood [14], with the $z = 2.33$ row in both conventions above: the fitted gap moves from -7.7 (`plik_lite`) to -6.8 (CamSpec) with the BAO-only row, and from -6.4 to -5.4 with the joint row — likelihood-choice shifts of $+0.9$ and $+1.0$ units of χ^2 (the four fits in Table 4B of Appendix B). A single model pair, however consistently it shifts across the two Lyman-α conventions, does not exhaust the distribution of likelihood-choice effects, so we conservatively adopt ± 3 units of χ^2 — three times the measured shift — as the systematic floor for every verdict in this paper. The floor is an overestimate by design: the typical shift is likely closer to the measured one unit than to three, so a difference below the floor is not thereby insignificant — it is a difference we decline to lean on in formulating the conclusions (Sec. 5).

BAO. DESI Data Release 2 (DR2; 13 numbers): the isotropic measurement from the bright galaxy sample (BGS) and six pairs (D_M/r_d, D_H/r_d) with published correlations [3], where $D_M(z)$ is the comoving angular-diameter distance, $D_H(z) = c/H(z)$ the Hubble distance, and r_d the sound horizon at the baryon drag epoch. The drag horizon uses the fitting formula

$$r_\mathrm{d} = 147.05\,\big(\omega_{cb}/0.1432\big)^{-0.23}\big(\omega_b/0.02236\big)^{-0.13}\ \mathrm{Mpc}, \tag{2}$$

a power-law approximation to Boltzmann-code outputs around the Planck point [15], as adopted in the DESI analyses [3]; here $\omega_b \equiv \Omega_b h^2$ and $\omega_{cb} \equiv \Omega_{cb} h^2$ are the physical densities of baryons and of baryons plus cold dark matter, with Ω_i the present-day density of species i in units of the critical density $\rho_\mathrm{crit} = 3H_0^2/(8\pi G)$ (G Newton's constant). The original formula carries a further factor $(N_\mathrm{eff}/3.04)^{-0.1}$, where N_eff is the effective number of relativistic species. We fix $N_\mathrm{eff} = 3.046$ — the Standard-Model prediction for three neutrino families with non-instantaneous decoupling [16], and the baseline of the Planck analysis we validate against — so this factor is a constant within 2×10^{-4} of unity and is omitted. For the second and third probe combinations the Lyman-α row is replaced by the July 2026 joint Alcock–Paczyński (AP [17]) + BAO measurement from the full shape of the DR2 Lyman-α correlations, $D_H/r_\mathrm{d} = 8.600 \pm 0.066$, $D_M/r_\mathrm{d} = 39.32 \pm 0.33$ at $z_\mathrm{eff} = 2.33$ (the effective redshift of the tracer sample) [18] — twice as tight as the BAO-only constraint from the same data (the correlation coefficient between the D_M/r_d and D_H/r_d measurements of this pair is kept at the BAO-only value -0.431 of the DR2 Lyman-α row [3], pending the published number for the joint analysis; all conclusions below are stable as the assumed correlation coefficient r_corr varies over $[0, -0.6]$: at fixed best-fit parameters — an upper bound on the sensitivity, since refitting absorbs part of any shift — the inter-model contrasts move by $\lesssim 1$ unit of χ^2, $\lesssim 2$ for the thermodynamic form that the full battery excludes, Sec. 4).

Supernovae. The Pantheon+ SNe Ia compilation [4, 19]: 1371 SNe Ia with $z_\mathrm{HD} > 0.023$, where z_HD is the catalog's Hubble-diagram redshift — the redshift in the CMB rest frame, corrected for the host galaxy's peculiar velocity [20] — and the Cepheid-calibrated objects of the SH0ES local-calibration program [21] are excluded; the absolute offset — degenerate between the supernovae's standardized absolute magnitude and H_0 — is marginalized analytically, so that only the *shape* of the diagram enters. The supernova block uses the full Pantheon+ statistical-plus-systematic (STAT+SYS) covariance (1701×1701, cut to the analysis subset), and the offset marginalization is performed with this full covariance; only the direction-setting frozen-cost probe of Sec. 4 retains diagonal errors.

Priors and counting. The baryon density ω_b is free within a flat range motivated by big-bang nucleosynthesis [22], $[0.020, 0.025]$; the optical depth to reionization carries a Gaussian prior $\tau = 0.0605 \pm 0.0046$, encoding the Planck low-ℓ polarization constraint [2]; the Planck absolute-calibration parameter carries its standard Gaussian prior $A_{\rm Planck} = 1 \pm 0.0025$ [2]; and the summed neutrino mass is floored at the normal-hierarchy minimum implied by the measured oscillation mass splittings [23], $\Sigma m_\nu \geq 58.19$ meV, imposed as a one-sided Gaussian wall of width 0.29 meV below the floor. These priors are common to every fit in this paper: all models, all three probe combinations, the profile likelihoods, and the MCMC run alike. Model comparison throughout uses the Akaike and Bayesian information criteria, $\mathrm{AIC} = \chi^2 + 2k$ and $\mathrm{BIC} = \chi^2 + k \ln N$ [24, 25], where k is the number of free fit parameters and N the number of fitted data points, quoted as differences ΔAIC and ΔBIC from the ΛCDM row of each table. The small-sample correction $\mathrm{AIC}_c - \mathrm{AIC} = 2k(k+1)/(N-k-1)$ is ≤ 0.23 per entry and ≤ 0.10 in any difference; we verified that it changes no conclusion-bearing ordering (its only effect anywhere is to swap two baseline rows separated by 0.03 units, on which nothing rests), and we ignore it. Quoted counts k exclude the two prior-dominated nuisances (τ, $A_{\rm Planck}$), whose values are set by the external priors rather than by the battery; since both are common to all models, this convention leaves every ΔAIC and ΔBIC unchanged.

Method and its limits. All results are χ^2 minimizations (multi-start Nelder–Mead simplex [26] over the full parameter vector; CMB observables from the Boltzmann code CAMB (Code for Anisotropies in the Microwave Background) [27], with each model's background supplied as an equation-of-state table $w(a)$ through the parameterized post-Friedmann dark-energy module [28]), not posterior samples. For the horizon family the tabulated $w(a)$ is itself a derived quantity, never a fitted one: the fits sample the parameters of the activation law $f(t)$ in Eq. (1), the solver returns the density history $\rho_{\rm de}(t)$, and $w(a)$ is read off that solution. Uncertainties for the headline quantities come from the profile likelihoods of Sec. 4.4, and a Markov-chain Monte-Carlo (MCMC) ensemble run for the leading model validates them there (posterior medians on the profile ridge; marginal widths quoted alongside). We deliberately do not quote a Bayesian evidence: for phenomenological parameterizations its value is dominated by the arbitrary prior volumes assigned to s, λ, z_c, and their analogues, whereas profile likelihoods expose the nonlinear degeneracies (such as the z_c–λ direction and its invariants) explicitly and without that prior dependence; a reader who prefers an evidence comparison can attach priors of their own to the released likelihood machinery and compute it. The hand-verification gate on the likelihood anchor stays in force for every run, and the cross-pipeline systematic floor above brackets every sub-3-unit verdict. One methodological caution is recorded explicitly: the free-exponent fit reproduced the same false minimum ($p \simeq 2$) in four independent data configurations before the profile of Sec. 4.4 exposed it — multi-start minimization can repeat an optimizer attractor; only profiles arbitrate. A corrections ledger (seventeen entries) and a stake protocol (predicted intervals before every run, misses recorded alongside hits) accompany the campaign notes; all scripts, with the self-validation machinery and the verdict files underlying every number below, are released on GitHub `https://github.com/yryabov/horizon-ladder` and archived on Zenodo (DOI: 10.5281/zenodo.24069445).

3 Models with physical motivation

All horizon-parametrized histories solve the flat Friedmann equations with Eq. (1) self-consistently: because R_E is the *future event* horizon, the system is solved by global relaxation to a history in which past and future agree — a numerical implementation of the Novikov self-consistency principle [29],

the same construction as in Li's HDE [11]. This teleology is a boundary condition, not a signal. No information propagates backward — the future event horizon enters the equations only as a global constraint selecting *which* self-consistent history the universe is on, and the selection is made by present-day data (the solver keeps the self-consistent history whose present-day expansion rate equals the fitted reduced Hubble constant h), not by future events acting on the past — the same logic by which top-down cosmology conditions histories on late-time data [30]. Point (d) of Appendix A makes this concrete: the self-consistency loop — today's density depending on the future expansion — admits a one-parameter family of solutions (the choice of H_∞), today's boundary condition picks one member, and within that member the evolution equations are ordinary and local in time. This setup is therefore no more acausal than any boundary-value problem.

With this construction in place, we now specify each member of the progression of dark-energy histories in turn — the controls first, then the two birth classes — stating for each its physical motivation and its number of free fitting parameters k.

The controls. Two controls belong to the family of Eq. (1) itself. *No dark energy* (table row "No dark energy"; $k = 5$):

$$f(t) \equiv 0, \tag{3}$$

flat matter plus radiation, where flatness leaves no freedom in the reduced Hubble constant and instead fixes it from the fitted physical densities, $h^2 = \omega_{cb} + \omega_\nu + \omega_r$ (with ω_ν and ω_r the neutrino and radiation densities) — the zero point of the progression. *Eternal screen* (table row "Eternal screen (HDE)"; $k = 7$):

$$f(t) \equiv c_{\rm HDE}^2 = \text{const}, \tag{4}$$

with $c_{\rm HDE}$ free, integrated by the standard HDE evolution equations including radiation [11] — the control that isolates the birth itself: the same capacity, no activation. Each of these two controls is fitted once: the eternal screen on the baseline battery (Table 1), where its equal-parameter duel with the thermodynamic birth is staged, and no dark energy on the full three-probe battery (Table 3); neither is refit on the other batteries. Two more controls sit outside the family of Eq. (1) and anchor the comparison. ΛCDM (table row "ΛCDM"; $k = 6$): a cosmological constant, $\rho_{\rm de} = \text{const}$, i.e. $w \equiv -1$ — the standard cosmological model, and therefore the benchmark that any claim of an evolving dark energy must beat. *The CPL tangent* (table row "CPL"; $k = 8$): $w(a) = w_0 + w_a(1 - a)$ [6, 7] — the first-order Taylor expansion of any smooth equation-of-state history around the present epoch ($a = 1$), adopted by survey analyses (including DESI) as the model-agnostic measure of evolution; it asks *whether* w evolves while committing to no mechanism.

Causal-kinetic forms (first-order transition). For a model in which dark energy comes into existence at some epoch after the birth of the universe, the first possibility to consider is that this inception is a first-order transition. The framework of this work parametrizes such a transition by the cosmological event horizon, Eq. (1). In most physical settings a first-order transition has a hallmark kinetics: the old and the new phase are separated by an interface that costs energy, so the new phase cannot switch on everywhere at once — it starts at discrete random sites and takes over by spreading. On the screen this reads: after a critical epoch, domains of the new phase appear at random sites and spread until they tile the screen, and the filling α is the fraction of area they have covered. For randomly seeded, growing, mutually impinging domains the covered fraction follows the Kolmogorov–Johnson–Mehl–Avrami (KJMA) coverage law $\alpha = 1 - e^{-X}$, where X is the mean area the domains would sweep if free to overlap [31, 32]. The sweeping itself can be metered by two different clocks, and the two resulting forms are fitted independently, as separate members of

the progression. The first clock is the e-fold variable $x = \ln a$; it gives the *Avrami* form (table row "Avrami"; $k = 8$)

$$f(t) = s\Big[1 - \exp\big(-((x - x_c)/\lambda)^2\big)\Big], \tag{5}$$

with $x_c = -\ln(1 + z_c)$ the activation epoch and λ the dimensionless growth width. The second clock is the light-crossing count I; it gives the *front/clock* form (table row "Front/clock"; $k = 8$)

$$f(t) = s\Big[1 - \exp\big(-\pi u^2 I^2(x)\big)\Big], \qquad I(x) = \int_{x_c}^{x} \frac{dx'}{H(x')\,R_E(x')}, \qquad u = \sqrt{N_0}\,v, \tag{6}$$

where I is the elapsed time since birth in units of horizon light crossings — one unit of I is the time in which a light signal covers one horizon radius R_E — v is the front speed in units of c, and N_0 is the number of independent nucleation seeds per horizon (lengths on the screen are measured in units of R_E throughout, so a front moving at speed v sweeps a disc of radius vI). The clock I is "causal" because it counts the number of times a signal at speed c could have crossed the horizon — which is what paces any physical process spreading over the screen, whereas an interval of cosmic time t carries no such invariant meaning for it. The form of the exponent is then exact: each seed grows a causal front that sweeps the area $\pi(vI)^2$, so the exponent is minus the mean covered area of N_0 Poisson seeds. Thus both Eq. (5) and Eq. (6) are exponential forms of $f(t)$, but their exponents carry different physical interpretations: a Gaussian growth profile on the e-fold clock in the former, the exactly computed swept area of causal fronts in the latter. Two physical readings of Eq. (6) — causal fronts growing from sparse seeds, or a single nonlocal screen ordering on its own clock — are observationally degenerate at the background level, since only the combination $u = \sqrt{N_0}\,v$ enters the equation. The two forms of Eqs. (5) and (6) are also the survivors of a broader scan of growth laws (Appendix A). The front reading has a natural speed scale: second sound — the temperature/entropy wave of a two-fluid system — which in its relativistic holographic realization propagates at $c/\sqrt{2}$ [33]; fronts at that speed turn the fitted combination into a seed count, $N_0 = 2u^2$. Both return in Sec. 6, as part of the open question of *how* the birth proceeded.

Thermodynamic forms (continuous transition, second order). If instead the transition is continuous, there is no interface to pay for: the order parameter rises from zero, the surface tension between the phases vanishes with it, and the condensed and normal fractions coexist at every point of the screen instead of in separate domains. The screen — a physical two-dimensional system held at the horizon temperature $T_{\rm dS}$ — then condenses *uniformly* as the expansion cools it through a critical value T_c, exactly as a superconductor or superfluid does on cooling through its critical point: below T_c a condensed fraction grows at the expense of the normal fraction, and the dark energy is carried by this condensate. A two-fluid reading of the de Sitter state itself — vacuum and matter as the superfluid and normal components of a thermal bath at the horizon-scale temperature — has been developed by Volovik [34]; that form is eternal — the two components coexist at all epochs, and no critical temperature enters. Here, instead, the two fluids live on the horizon screen and their conversion has an onset; to instrument this continuous transition we adopt for the condensed fraction the classic two-fluid mean-field ansatz of Gorter and Casimir [35],

$$f(t) = \begin{cases} 0, & T_{\rm dS}(t) > T_c \quad \text{(the screen is still empty)}, \\ s\Big[1 - \big(T_{\rm dS}(t)/T_c\big)^p\Big], & T_{\rm dS}(t) \le T_c. \end{cases} \tag{7}$$

The critical exponent p probes the excitation spectrum of the screen ($p = 4$: linear-dispersion 2d modes; $p = 2$: quadratic dispersion). We fit this form twice: with the exponent fixed at its canonical $p = 4$ (one physical parameter; table row "Thermal GC"; $k = 7$), and with p free (table row "GC-p"; $k = 8$).

State versus history. The two classes of activation law are not arbitrary parameterizations but the two thermodynamic characters a transition can have. In a first-order transition the order parameter is a *functional of history*: once nucleated, domains need time to grow, so α depends on the time elapsed since birth, $\alpha(I(t))$, however the temperature could move meanwhile. In an equilibrium (second-order-like) condensation the order parameter is instead a *function of state*: the condensed and normal fractions stay in mutual thermodynamic equilibrium at the instantaneous horizon temperature — they equilibrate faster than the expansion cools the screen — so α tracks its equilibrium value at the current temperature, $\alpha(T(t))$: equilibrium knows no other clock. One structural consequence matters below. In the equilibrium form the single constant T_c sets both *when* activation starts and *how fast* α grows afterwards; since $T_{\rm dS}(t)$ falls monotonically, an early crossing of T_c would leave the screen today deep below its critical point, $T_{\rm dS} \ll T_c$, and Eq. (7) would then force $\alpha \simeq 1$ — a saturated, effectively eternal screen. In the kinetic forms the onset (z_c) and the growth rate (λ or u) are independent parameters, and the birth of dark energy can fall at any epoch.

Exact properties. Three identities are specified here: the first holds across the whole family of Eq. (1), the other two within its equilibrium class. On the de Sitter attractor the expansion rate and the horizon approach constants, $H \to H_\infty$ and $R_E \to R_\infty = c/H_\infty$, and the filling approaches its limit α_∞. Therefore the following holds. First, the Friedmann equation on the attractor, $H_\infty^2 = s\,\alpha_\infty/R_\infty^2$, forces

$$s\,\alpha_\infty = 1, \tag{8}$$

so for the models with a dark-energy birth the amplitude s of the filling function $f(t) = s\,\alpha$ is the inverse asymptotic filling, $s = 1/\alpha_\infty \geq 1$: $s = 1$ for a screen that completes ($\alpha_\infty = 1$, the kinetic forms), $s > 1$ for an equilibrium screen whose condensation halts at a temperature floor and never completes. In either case the product $s\,\alpha(t)$ is driven to exactly 1 in the future. In HDE language, the constant $c_{\rm HDE}^2$ (in eternal HDE a number set by hand) is promoted to the dynamical product $s\,\alpha(t)$, driven to unity by an attractor rather than by choice; this may be related to the fact that fits of eternal HDE typically return $c_{\rm HDE} \lesssim 1$, with the preferred value varying noticeably between data sets [36, 37]. It also removes the big-rip doomsday: in eternal HDE the asymptotic equation of state is $w_\infty = -1/3 - 2/(3c_{\rm HDE})$ [11], phantom — and rip-bound — whenever $c_{\rm HDE} < 1$, whereas here Eq. (8) pins $w_\infty = -1$ exactly (both statements are verified numerically in Appendix A, points (a) and (b)). Second, for the equilibrium class the filling is a function of the horizon alone ($\alpha = 1 - (T_{\rm dS}/T_c)^p$ with $T_{\rm dS} \propto 1/R_E$), so the effective equation of state $w \equiv P_{\rm de}/\rho_{\rm de}$ crosses the cosmological-constant value at a filling fixed by the critical exponent alone,

$$w = -1 \quad \Longleftrightarrow \quad \alpha = \frac{p}{p+2}, \tag{9}$$

independent of both T_c and s: at the canonical $p = 4$ the crossing sits at two thirds of the horizon budget, $\alpha = 2/3$. This is a direct consequence of the state-versus-history dichotomy above: a filling that is a function of state makes w an algebraic function of α, and the crossing parameter-free; the kinetic forms, whose filling is a functional of history, carry no such identity. Third, for the same class at the canonical $p = 4$, the uncondensed screen area $A \equiv (1-\alpha)A_{\rm hor}$ and the full horizon area $A_{\rm hor} = 4\pi R_E^2$ obey, in natural units (see Sec. 1), the parameter-free reciprocity

$$A \cdot A_{\rm hor} \;=\; \frac{1}{\pi^2 T_c^4} \;=\; \text{const.} \tag{10}$$

This is an invariant with no adjustable content: once T_c is known, the product is fixed at every epoch, giving the equilibrium class an internal consistency check beyond the fit itself. Physically

Table 1: Baseline battery: `plik_lite` TT+TE+EE + DESI DR2 BAO, $N = 626$ data points (613 CMB band powers + 13 BAO numbers), $\ln N = 6.44$. ΔAIC and ΔBIC relative to ΛCDM. Model names: Front/clock = the constant-front-speed form of Eq. (6); Avrami = the KJMA form; GC = the Gorter–Casimir form of Eq. (7) ($p = 4$ fixed; GC-p: p free); HDE = eternal holographic screen.

Model	k	$\chi^2_{\rm CMB}$	$\chi^2_{\rm BAO}$	$\chi^2_{\rm tot}$	ΔAIC	ΔBIC	key parameters
Front/clock	8	**581.00**	8.69	**589.82**	−4.90	+3.98	z_c=4.05, u=1.63
CPL	8	582.05	**8.06**	590.29	−4.43	+4.45	$w_0 = -0.708$, $w_a = -0.913$
Avrami	8	581.12	9.77	591.05	−3.68	+5.20	z_c=4.01, λ=0.90
GC-p	8	581.31	11.16	592.61	−2.11	+6.77	$T_c/T^0_{\rm dS}$=1.66, p=2.01
Thermal GC	7	581.54	13.06	594.64	−2.08	**+2.36**	$T_c/T^0_{\rm dS}$=1.50, z_c=2.02
ΛCDM	6	586.42	11.88	598.72	0	**0**	h=0.684
Eternal screen (HDE)	7	598.91	31.67	631.66	+34.9	+39.4	$c_{\rm HDE}$=0.54, h=0.733

it states a rigid trade: as the expansion grows the horizon, the uncondensed area shrinks in exact inverse proportion — the condensation outruns the growth of the very screen it is filling. The constant on the right-hand side is the square of the horizon area at birth, $A_c \equiv 1/(\pi T_c^2)$: the uncondensed and the full area are mutual reflections through the birth area, $A/A_c = A_c/A_{\rm hor}$, and the history enters this hyperbola at its self-dual point, $A = A_{\rm hor} = A_c$ — the activation epoch — and slides along it ever after. The quadratic structure also carries the Landau reading of the ansatz: the filling is the squared amplitude of an order parameter, $\alpha = |\psi|^2$ with $\psi = \pm[1 - (T_{\rm dS}/T_c)^p]^{1/2}$, real only below T_c, so the birth is the point where ψ turns real, and the sign pair is the usual degenerate choice of the broken phase. More generally the uncondensed area scales as $A \propto R_E^{2-p}$ — a race between the spectral exponent p, pacing how fast cooling kills the thermal population, and the two geometric powers of the growing area, with parity at $p = 2$, where the uncondensed area would freeze. The canonical $p = 4$ wins this race with exactly the margin that holds the product of Eq. (10) constant.

4 Fitting results

The progression of dark-energy histories detailed in Sec. 3 is fitted to three successively larger combinations of experimental data probes; each is an independent repetition of the comparison. Tables 1–3 of this section present the essential parameters and results of the data fits; the full fitted-parameter vectors behind every table row are listed in Appendix B (Tables 1B–3B).

4.1 Baseline battery: CMB + BAO

We start the data analysis with the simplest combination of probes, the CMB and BAO data. The essential results of this analysis are presented in Table 1. Every born form improves on ΛCDM by $\Delta\chi^2 = 4.1$–8.9 at one or two extra parameters; the causal-clock form is the first physical parameterization of the analysis to overtake the free tangent, CPL, at equal k; and the one-parameter duel at equal $k = 7$ — thermodynamic birth versus eternal screen — ends 594.6 against 631.7: removing the birth while keeping the holographic capacity costs +37, with the eternal screen falling +33 *below* ΛCDM. (The eternal screen is not refit on the later batteries: an added independent block can only raise χ^2, and the supernova Hubble diagram tightens precisely the late-time expansion its best fit already distorts, $h = 0.733$.) The birth-date reproduces across the kinetic forms, $z_c = 4.01$–4.05, and survives the likelihood swap of Sec. 2 ($z_c = 4.11$ on PR4/CamSpec TTTEEE; full vectors in Appendix B).

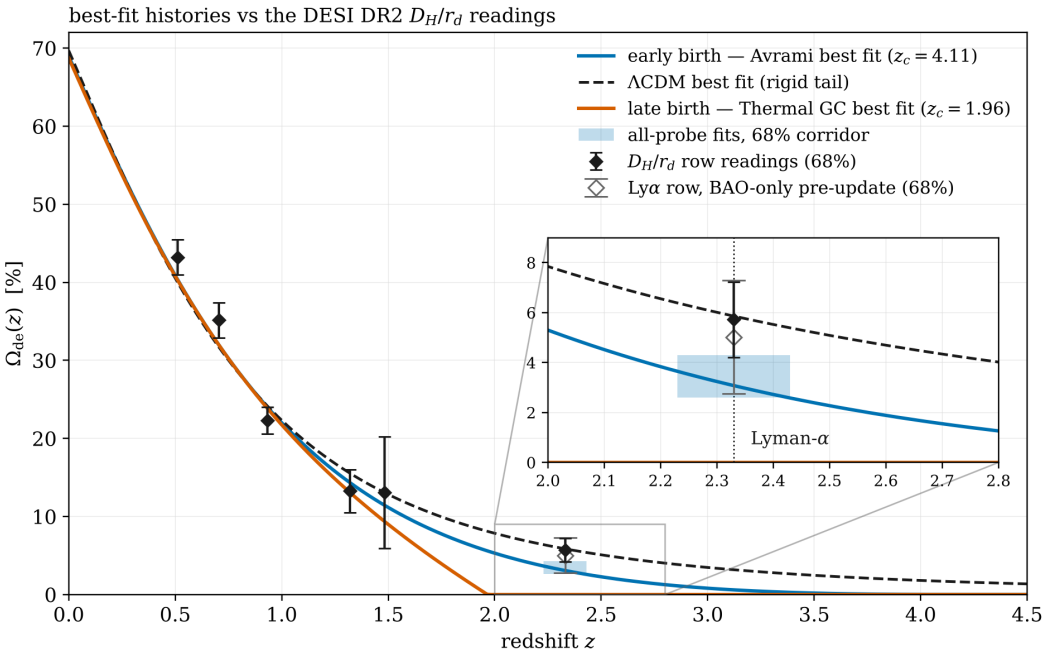

Figure 1: The Lyman-α lever at the three-probe best fits. The early dark-energy tail $\Omega_{\rm de}(z)$ at the Table 3 best fits of the three classes, against the single-row inversions (Sec. 4.2) of the six DESI DR2 radial BAO measurements D_H/r_d with ω_m and r_d CMB-pinned (diamonds: the galaxy and quasar bins at $z = 0.510$–1.484 and the joint Lyman-α row at $z = 2.33$; the transverse D_M/r_d rows and the D_V bin, fitted in the battery, do not invert into $\Omega_{\rm de}$ row by row and are not shown). Below $z \approx 1$ the three histories coincide and the readings test only their common fit; the classes separate where the data thin out. Inset (the lever region): the joint row reading $\Omega_{\rm de}(2.33) = 5.7 \pm 1.5\%$ (filled diamond) excludes a vanishing tail at 3.8σ on its own; the pre-update BAO-only row $5.0 \pm 2.3\%$ (open diamond) and the all-probe 68% corridor of Sec. 4.4 (shaded box) bracket the arbitration between the rigid ΛCDM tail, the born Avrami tail, and the vanished Thermal GC tail — a contest of about two of the row's standard deviations. At these best fits the battery's BAO block (13 numbers) gives $\chi^2_{\rm BAO} = 14.15$ (ΛCDM), 12.10 (Avrami), 19.89 (Thermal GC): for the early-birth form the galaxy rows and the Lyman-α row nearly cancel — the balance is the lever shown here — while the late-birth form, its tail vanished at $z = 2.33$, is charged by the Lyman-α row alone (≈ 12 of its ≈ 20).

4.2 Second battery: the Lyman-α full-shape update

The recent full-shape analysis of the DESI DR2 Lyman-α forest [18] provides a joint AP+BAO measurement of the $z = 2.33$ point, with a shifted center and roughly twice tighter uncertainties than the BAO-only row of the baseline. To evaluate the effect of this improvement we performed the fits detailed in Table 2 (full parameter vectors in Table 2B of Appendix B). The new Lyman-α center taxes *every* form, ΛCDM included (+1.3 units of χ^2 for replacing a single bin) — this should be said before any comparison.

The physics of the lever is worth spelling out. Since $D_H/r_{\rm d} = c/(H\, r_{\rm d})$, the shifted and twice-tightened value is in effect a direct measurement of the expansion rate $H(z{=}2.33)$. By the Friedmann equation, $H^2(2.33) \propto \rho_m(2.33) + \rho_{\rm de}(2.33)$ (the radiation term is omitted: at this epoch it contributes only $\sim 10^{-3}$ of H^2). The matter density and $r_{\rm d}$ are pinned by the CMB acoustic peaks — raising ω_m would displace *every* BAO point, not this one. The only free lever at $z = 2.33$ is therefore the early dark-energy tail: the fraction of the critical density, $\Omega_{\rm de}(z) \equiv \rho_{\rm de}(z)/\rho_{\rm crit}(z)$, still present at that epoch. The row is thus the first direct measurement of this tail. Its single-row inversion, with ω_m and r_d CMB-pinned, gives $\Omega_{\rm de}(2.33) = 5.7 \pm 1.5\%$ of the critical density at that

Table 2: Second battery: the effect of the Lyman-α row update on the fits, compared with the baseline battery of Table 1 — the $z = 2.33$ row is replaced by the joint AP+BAO measurement [18]. "Price" is the model's refitted $\chi^2_{\rm tot}$ minus its own Table 1 total.

Model	k	$\chi^2_{\rm CMB}$	$\chi^2_{\rm BAO}$	$\chi^2_{\rm tot}$	price	key drifts
CPL	8	581.81	9.96	**592.13**	+1.84	$w_0 = -0.712$, $w_a = -0.893$
Avrami	8	581.57	11.97	593.60	+2.55	z_c: 4.01 → 4.17
Front/clock	8	581.11	12.44	593.63	+3.81	z_c=4.02
GC-p	8	581.66	14.48	596.34	+3.73	p=2.03, z_c=2.75
ΛCDM	6	586.10	13.53	600.02	+1.30	—
Thermal GC	7	582.35	19.89	602.26	+7.62	z_c=1.96

epoch.

Figure 1 draws the lever itself: the tail $\Omega_{\rm de}(z)$ along the competing histories (drawn at the full-battery best fits of Table 3 ahead), with the measured row at $z = 2.33$. The inset magnifies the decisive interval, $z = 2$–2.8. There the joint reading stands against the same inversion applied to the pre-update, BAO-only row: $\Omega_{\rm de}(2.33) = 5.0 \pm 2.3\%$. The update shifts the center by a third of the old error bar and cuts the uncertainty by a third.

A history activating as late as the canonical thermodynamic form ($z_c \simeq 2$, Tables 1–2) carries identically *zero* dark energy at $z = 2.33$ — no response is possible — and the measured tail excludes that vanishing value at 3.8σ. The cosmological constant remains compatible: its tail is fixed and non-adjustable, $\Omega_\Lambda/[\Omega_\Lambda + \Omega_m(1+z)^3] \simeq 6\%$. The pull toward a smaller tail comes from the battery as a whole: the full fits settle at $\Omega_{\rm de}(2.33) \simeq 3\%$ (Sec. 4.4), a value only the kinetic class, activating early ($z_c \simeq 4$), can dial freely.

Hereafter we call the two classes *late-birth* and *early-birth*, after where their fits land: the canonical thermodynamic form at $z_c \simeq 2$, the kinetic forms at $z_c \simeq 4$, the free exponent between. Either birth leaves the recombination epoch of the CMB ($z \simeq 1100$) untouched — dark energy is still absent at $z \gtrsim 5$ — and the D_M half of the joint measurement anchors the integrated low-redshift side. The consequences follow. The canonical Gorter–Casimir form ($p = 4$ fixed) absorbs the full blow (+7.6 units of χ^2, falling below ΛCDM). The free exponent buys part of it off by shifting its birth earlier (z_c: 2.60 → 2.75; price +3.7, comparable to the kinetic forms). The early-birth dates drift upward to $z_c = 4.0$–4.2, and the tangent — least punished — briefly regains the lead. The free thermodynamic exponent lands on the integer again ($p = 2.03$; baseline 2.01) — a reproducibility whose discussion we return to in Sec. 4.4.

4.3 Third battery: adding the supernova Hubble diagram

Finally, we add the Pantheon+ supernova Hubble diagram, completing the set of batteries. The fits are detailed in Table 3 (full parameter vectors in Table 3B of Appendix B). A frozen-background probe decides the direction before any refit: the supernova χ^2 alone, scored with the diagonal errors of the Pantheon+ covariance at the Table 2 best fits. On pure Hubble-diagram shape, every born form is *cheaper* than ΛCDM by 4.0–4.5 units of χ^2. The Lyman-α–optimized tangent *pays* +3.05. That is a ~ 7.5-unit gap on the supernova leg alone. The full refit with the complete STAT+SYS covariance (Table 3) confirms and sharpens it. The born Avrami form takes the top row: $\Delta\chi^2 = -8.4$ against ΛCDM, the strongest AIC evidence of the analysis ($\Delta{\rm AIC} = -4.4$). The CPL tangent survives only by *migration*: squeezed by all three probes, it abandons its Lyman-α optimum $(-0.71, -0.91)$ for $(-0.873, -0.488)$, deep in the born family's corner ($w_0 \simeq -0.96$), and trails the born leader by 1.7 units of χ^2.

Table 3: Three-probe battery: fits with the full data set, TTTEEE + [BAO with the joint Lyman-α row] + Pantheon+, $N = 1997$, $\ln N = 7.60$. The supernova block uses the full Pantheon+ STAT+SYS covariance with the offset marginalized analytically.

Model	k	$\chi^2_{\rm CMB}$	$\chi^2_{\rm BAO+SN}$	$\chi^2_{\rm tot}$	ΔAIC	key parameters
Avrami	8	581.27	**1222.11**	**1803.48**	**−4.41**	z_c=4.11, λ=0.903
GC-p	8	581.46	1222.85	1804.65†	−3.24	p=0.8 (profile†), $z_c \simeq$3.0
CPL	8	582.11	1222.70	1805.15	−2.74	$w_0 = -0.873$, $w_a = -0.488$
Front/clock	8	**581.13**	1225.46	1806.73	−1.16	z_c=4.06, u=1.620
ΛCDM	6	585.57	1225.93	1811.89	0	—
Thermal GC	7	582.38	1229.86	1812.25	+2.36	z_c=1.96
No dark energy	5	2302.4	6662.2	8994.7	+7181	flat matter, h=0.450 derived

†The GC-p row quotes the class minimum located by the p-profile of Sec. 4.4 ($p \simeq 0.8$, $z_c \simeq 3.0$). The free multi-start fit had stalled on the shallow exponent direction at an optimizer attractor ($p \simeq 2.0$, $\chi^2 = 1806.40$).

Figure 1 puts this verdict against the expansion-rate data themselves. At the Table 3 best fits the three classes are indistinguishable below $z \approx 1$ — precisely where the DESI radial rows are most precise. The five galaxy and quasar readings ride all three histories equally well and arbitrate nothing: at those redshifts the three curves are images of the same CMB-pinned background, so the precise galaxy rows test the shared matter sector, not the dark-energy law. Sharper low-redshift BAO cannot change this. Only the thinly measured tail can — which is why the five-year Lyman-α row, not more galaxy precision, carries the next arbitration (see Sec. 6). The histories part only in the tail, and there the joint Lyman-α row is the single reading inside the fork (inset of Fig. 1). It excludes the vanished Thermal GC tail at 3.8σ. It sits on the rigid ΛCDM tail. It pulls 1.8σ above the born best fit, whose all-probe corridor it overlaps at one standard deviation. The whole contest between the surviving dark-energy histories is thus a percent-level question about the tail the Lyman-α row measures — which is why that single row carries the arbitration. (The front-runners against the other two data blocks — the raw CMB bandpowers and the supernova Hubble diagram — are collected in Appendix B, together with the marginalized posterior behind the intervals of Sec. 4.4.)

The Lyman-α update of Sec. 4.2 can now be priced at these very best fits. Frozen at the Table 3 vectors, the BAO block scores $\chi^2_{\rm BAO} = 13.00$ / 10.43 / 9.58 (ΛCDM / CPL / Avrami) on the pre-update row, and 14.15 / 11.83 / 12.10 on the corrected one — a tax of +1.15, +1.40 and +2.52 units of χ^2 respectively. The correction charges most exactly the form that empties the tail fastest. On the old row the born form led the BAO block outright; on the corrected one the tangent edges it within the block by a quarter-unit. The battery's verdict absorbs the blow: the Avrami lead of Table 3 stands on a correction priced *against* it — worth knowing when the five-year row arrives. The leading gaps of Table 3 — 1.7 to the migrated tangent, 3.3 across the top four — are comparable to or below the ± 3 cross-pipeline systematic floor of Sec. 2.

The covariance upgrade *strengthens* rather than erases the born-form lead: from 0.5 units of χ^2 with diagonal errors to 1.7 with the full STAT+SYS matrix. For the first time it also mildly separates the two early-birth kinetic forms, Avrami and Front/clock: 3.3 units on the supernova leg, favoring the Avrami coverage over constant-speed fronts (directional only, at the level of the systematic floor). Their dates and scales hold: $z_c = 4.06$–4.11, $u = 1.620$. The free exponent again settles near $p \simeq 2.0$; whether this is a data preference or an optimizer artifact is a question for the profile analysis of Sec. 4.4, whose located class minimum the GC-p row of Table 3 quotes (table note). The canonical Gorter–Casimir form ($p = 4$ fixed) falls below ΛCDM (ΔAIC = +2.4). The

no-dark-energy control fails catastrophically: $\chi^2_{\rm tot} = 8994.7$ at the Einstein–de Sitter point (the flat matter-only solution, $h_{\rm derived} = 0.450$), with every free parameter pinned at its prior bound — so $\chi^2 = 8995$ is an *optimistic* bound for a universe without dark energy. (This control is fitted once, on the fullest battery — the most generous configuration, with the most data over which the optimizer could spread the contradiction. Its results appear as the "No dark energy" row of Table 3 only; a failure of this size justifies forgoing the per-battery replication.)

4.4 Profile-likelihood uncertainties

We now attach uncertainties to the fitted parameters. The intervals come from profile likelihoods on the full three-probe battery: one parameter is held on a grid while all remaining parameters are re-minimized, warm-started outward from the best fit.

Activation epoch. The defining parameter of the kinetic class is the activation epoch z_c. The profile over z_c for the leading Avrami form (21 grid points, $2.8 \le z_c \le 6.3$) gives

$$z_c = 4.2^{+1.0}_{-0.3}\ (68\%), \qquad z_c > 3.4\ (95\%), \tag{11}$$

with the upper 95% bound beyond the probed window ($\Delta\chi^2 = 3.2$ at $z_c = 6.3$). The interval is strongly asymmetric: the data forbid activation later than $z_c \simeq 3.9$ and are nearly indifferent to how much earlier it occurred, the residual direction being degenerate with the growth rate (λ runs $0.68 \to 1.22$ along the profile; Fig. 4C(a)). Two quantities are nearly invariant along the ridge of this degeneracy and are therefore the actual ridge measurements: today's equation of state, $w_0 = -0.95$ to -0.96 (0.5% spread inside the 68% window), and the early dark-energy tail, $\Omega_{\rm de}(z{=}2.33) = 2.6$–$4.3\%$ (68%) — the direct target of the five-year Lyman-α measurement. In the wall region ($z_c \le 3.4$) the re-optimizations converge to within ~ 1 unit, and the 95% lower bound carries that convergence uncertainty.

Thermodynamic exponent. Models built on a continuous, second-order transition are specified above all by their critical exponent p. The profile over p for the free-exponent thermodynamic form (16 points, $0.6 \le p \le 2.8$) overturns the free-fit result: $\chi^2(p)$ decreases monotonically below the free fits' $p \simeq 2.0$ and settles on a flat plateau at $p \lesssim 1$ (minimum $\chi^2 = 1804.65$ at $p = 0.8$; the entire range 0.6–1.0 lies within $\Delta\chi^2 < 0.03$). The result is an upper limit,

$$p < 1.7\ (68\%), \qquad p < 2.7\ (95\%), \tag{12}$$

with $p = 2$ disfavored at $\Delta\chi^2 = 1.7$ and the Gorter–Casimir $p = 4$ excluded. The repeated landings of the free fits at $p \simeq 2.00$ across the three batteries of this analysis are thereby exposed as an optimizer attractor on a shallow direction ($\Delta\chi^2 \sim 1$ per 0.5 in p), not a data preference; we record this as a methodological caution. At its minimum the thermodynamic form migrates into the same corner as every other survivor — earlier, flatter activation ($z_c \simeq 3.0$, onset $\propto (1 - T/T_c)^{0.8}$), with the present-day product $s\,\alpha_0 \simeq 0.93$ ($\alpha_0 \equiv \alpha$ today) holding the family invariant. In the context of the state-dependent models of a continuous transition, the plateau sits below the parity exponent $p = 2$ of Sec. 3 ($A \propto R_E^{2-p}$): at the preferred soft exponents the uncondensed area grows in absolute terms — the data select a condensation that no longer outruns the expansion of its own screen, with an exponent matching no standard dispersion law (see Sec. 3).

Posterior validation (MCMC). The full posterior analysis is performed for the Avrami form only — the winner of the three-probe battery (Table 3), on which every headline number of this paper rests. The profile intervals follow the ridge of the likelihood; an independent posterior sampling tests what honest marginalization over that ridge does to them. An `emcee` ensemble run (the affine-invariant ensemble sampler of Goodman and Weare [38], in the implementation of Foreman-Mackey et al. [39]) on the full battery (32 walkers × 1500 steps, burn-in 500; acceptance 0.33; integrated autocorrelation times 80–133 steps, i.e. an effective sample size of $\sim$250–400 — adequate for 68% intervals, thin for extreme tails) confirms the profile structure and prices the marginalization (Fig. 4C). The activation-epoch posterior is $z_c = 4.45$ (median) with 68% interval $[3.61, 5.61]$, one-sided exactly as the profile found (the 97.5% point, $z_c = 7.5$, abuts the prior edge at 8 — the posterior image of the profile's unbounded upper arm); the (z_c, λ) degeneracy appears as a sample correlation of $+0.95$; and the ridge invariants hold at the center while acquiring their honest marginal widths — re-solving the background on a random subsample of 320 posterior samples (the background solve is the expensive step; the same 320 re-solved histories supply the posterior bands of Figs. 2 and 3) gives $w_0 = -0.96 \pm 0.03$ and $\Omega_{\rm de}(2.33) = 3.5\%$ with 68% interval $[2.2, 4.5]\%$, against the ridge-bottom corridors -0.951 to -0.962 and 2.6–4.3%; panel (b) of Fig. 4C overlays the two, with the profile corridors sitting inside the Bayesian 68% contour — the frequentist and Bayesian pipelines read the same ridge. The posterior intervals are wider than the profile's, as marginalization over the ridge must make them; we keep the profile intervals as the headline (they describe the ridge the data actually constrain) and quote the marginal widths alongside as the forecast-grade numbers. One genuine difference surfaced: the summed neutrino mass, wall-pinned at the normal-hierarchy floor in every minimization, lifts off in the posterior to $\Sigma m_\nu = 87^{+39}_{-21}$ meV ($P(\Sigma m_\nu < \text{floor}) = 0.9\%$) — a volume effect typical of a shallow one-sided direction, consistent with the known relaxation of neutrino-mass bounds in evolving-dark-energy backgrounds; we flag it as such, not as a detection.

From the physical perspective, these observations can be summarized as follows: what the data pin are not the internal dials of any single form — the date–rate ridge and the exponent plateau both stay soft — but two shared projections of the history itself (Fig. 4C(b)): a present-day equation of state frozen just above the cosmological constant, $w_0 = -0.96 \pm 0.03$ (Fig. 2), and an early dark-energy tail $\Omega_{\rm de}(2.33) \simeq 3.5\%$. The first is the forecast-grade number of this analysis; the second is the quantity that the five-year DESI Lyman-α measurement, expected to be released plausibly by mid-2027, will read directly.

The apparent phantom. The evolution of the dark-energy equation of state, $w(z)$, along the history most consistent with the data — the Avrami best fit of the fullest battery (Table 3) against the CPL tangent refitted on the same battery — is presented in Fig. 2. Its phantom side, $w < -1$, is kinematic, not a field pathology. For any component of the cosmic energy budget on a Friedmann background, the continuity equation ties the effective equation of state to the sign of the density trend — a standard identity of dark-energy reconstruction (e.g. [40]),

$$1 + w \;=\; -\frac{1}{3}\,\frac{d\ln\rho}{d\ln a}. \tag{13}$$

A density still growing toward its attractor value reads $w < -1$, however healthy the dynamics behind it. The born histories therefore enter the diagram on the phantom side due to a race of rates. With $\rho_{\rm de} \propto \alpha/R_E^2$, the young component reads $w < -1$ while its activated fraction grows faster than the area of its screen. It exits the phantom at the moment the expanding screen overtakes the filling — the density maximum of Fig. 3: the fitted Avrami form crosses at $z = 0.58$ (corridor

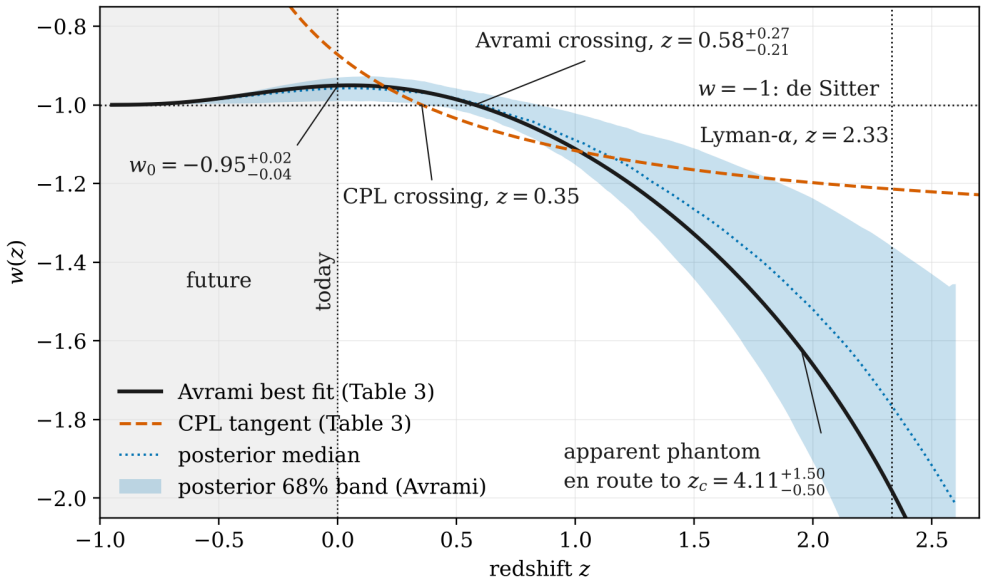

Figure 2: The fitted equation-of-state history: $w(z)$ of the leading Avrami form — the Table 3 best fit with the posterior median and 68% band — against the CPL tangent refitted on the same battery. The axis continues into the future ($z < 0$, shaded): there every posterior history has long been born, the band closes onto the attractor, and the two histories part for good — the Avrami form settles exactly on $w = -1$ (the de Sitter attractor), while the CPL tangent, extrapolated far outside its Taylor validity, rises without bound and leaves the frame just past today. On the past side the band ends at $z = 2.6$, the latest redshift at which every posterior history has already activated: the ridge admits births as late as $z_c \simeq 2.7$, and above its own z_c a history carries $\rho_{\rm de} \equiv 0$ — there is no equation of state to quote for a component not yet born. The cut is thus the lower edge of the data-allowed birth dates, not a boundary of the data; the best-fit curve, with its own $z_c = 4.11$, continues through it. The features named in the text are marked: today's equation of state, anchored at the best fit, $w_0 = -0.95^{+0.02}_{-0.04}$ (68%; the asymmetry is the one-sided ridge — the posterior prefers slightly more negative values; the median-centered marginal reads -0.96 ± 0.03, the number quoted in the text); the two exits from the apparent phantom — the Avrami best fit crosses $w = -1$ at $z = 0.58^{+0.27}_{-0.21}$ (the uncertainty read off the band-edge crossings), the CPL tangent at $z = 0.35$, permanently; and the apparent-phantom transient en route to activation. The tangent leaves the Avrami posterior band both at high redshift (a rigid asymptote against the deepening transient) and at $z = 0$ (its $w_0 = -0.873$ sits above the band).

0.37–0.85 across the posterior band), the equilibrium class at the parameter-free filling of Eq. (9). The filling itself never stops ($\alpha \to 1$ only asymptotically); what turns over is the race.

Figure 3 makes the statement literal by plotting the density itself: in units of its de Sitter plateau, $\rho_{\rm de}(z)$ is one gentle hill — identically zero above the birth at $z_c = 4.11$, a monotone climb to a maximum of $1.15^{+0.07}_{-0.11}\,\rho_{\rm dS}$ at $z = 0.58$, and a 13% relaxation onto the plateau, passing through today's $1.091\,\rho_{\rm dS}$ on the way down. By Eq. (13), the top of the hill *is* the $w = -1$ crossing, and the whole phantom side of Fig. 2 is nothing but the hill's rising flank read through a logarithmic derivative; nothing in the density history is singular, negative, or runaway, and the overshoot that the word "phantom" dramatizes amounts to fifteen per cent of the eternal level. The CPL tangent instead carries its phantom side as a permanent asymptote ($w \to w_0 + w_a$ at early times) and exits it only at $z = 0.35$.

The two exits open onto opposite futures ($z < 0$ in Figs. 2 and 3): the born history settles onto its attractor — w exactly at -1, the density on its de Sitter plateau — while the CPL tangent, read forward far outside its Taylor validity, has no attractor to land on: its $w(a)$ rises without

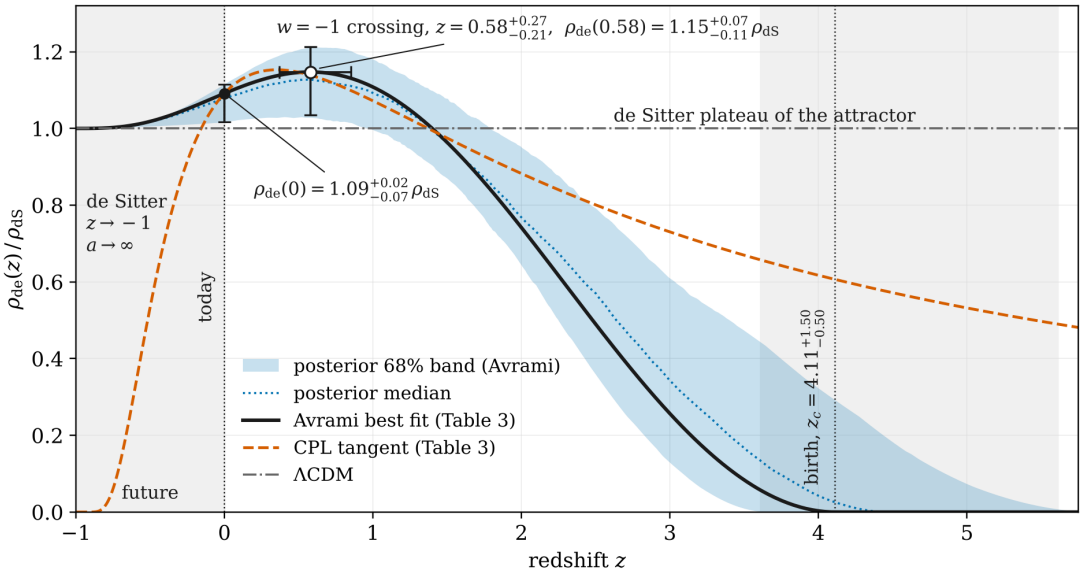

Figure 3: The fitted density history: $\rho_{de}(z)$ of the leading Avrami form in units of its de Sitter plateau ρ_{dS} — the Table 3 best fit, with the posterior median and 68% band re-solved on the same 320-sample subset as Fig. 2, each history normalized to its own plateau. Unlike $w(z)$, the density is defined everywhere: $\rho_{de} \equiv 0$ before birth is a value, not a gap, so the band needs no activation window, and the axis runs to $z = -1$, the infinite future ($a \to \infty$), where every history settles on its plateau and the band closes. The history is a single gentle hill: identically zero above the birth at $z_c = 4.11$, a monotone climb to a maximum of $1.15^{+0.07}_{-0.11}\,\rho_{dS}$ at $z = 0.58$ (the vertical whiskers span the 68% interval of the posterior histories' own maxima) — by the identity $1 + w = -\frac{1}{3}\, d\ln\rho_{de}/d\ln a$ the density maximum *is* the $w = -1$ crossing of Fig. 2, and the whiskers on the maximum span that crossing's 68% corridor ($z = 0.37$–0.85) — and a 13% relaxation onto the plateau, passing through today's value on the way down. Today's value carries its posterior interval, anchored at the best fit: $\rho_{de}(0) = 1.09^{+0.02}_{-0.07}\,\rho_{dS}$ (68%, vertical whiskers; the asymmetry is the one-sided ridge again — the posterior prefers slightly lower values than the best fit). The grey band on the right spans the posterior 68% interval of the birth, $z_c \in [3.61, 5.61]$ (median 4.45, one-sided toward the prior edge exactly as the profile found). The entire apparent-phantom episode of Fig. 2 is the rising flank of this hill. ΛCDM, a constant, coincides identically with the plateau line; the CPL tangent, matched to the best fit at $z = 0$, falls monotonically into the past, with no birth; extrapolated into the future — far outside its Taylor validity, drawn only to make the contrast visible — it has no plateau either: its $w(a) = w_0 + w_a(1 - a)$ rises without bound and the density collapses to zero.

bound just past today and its density collapses to zero. An extrapolation of a history into the future cannot be confronted with data, so this contrast arbitrates nothing; it shows a difference in kind — the tangent records the present slope of a history, the born form is a history whole, with an endpoint of its own. Either way, the benign reading is available only because the family of Eq. (1) parameterizes the density on the horizon: $\rho_{de}(t)$ is specified directly, and no propagating field stands behind it. That a horizon parameterization walks through the phantom divide free of field pathologies has been part of holographic dark energy since its canonical form [11]: for $c_{HDE} < 1$ the effective equation of state crosses $w = -1$ with no ghost behind it. It is nevertheless worth noting that in the original HDE treatment the phantom side is the history's destination — the asymptote $w_\infty = -1/3 - 2/(3c_{HDE}) < -1$, singular in the far future — whereas the born family turns it into a strictly bounded transient steered by the activation dynamics: entered at birth, exited at the density maximum, closed by the attractor exactly at $w = -1$.

Under a bulk reading — a single, minimally coupled field in the volume reproducing the same

history — the phantom side is *real* and carries its known pathologies: $w < -1$ requires violating the null energy condition, generically through a ghost kinetic term that destabilizes the vacuum against runaway particle production [41, 42], and, if sustained, drives the expansion to a Big Rip singularity in finite time [43, 44]. The generalized second law draws the same line thermodynamically — indeed, under the GSL the approach to de Sitter is a holographic no-hair theorem [45]. A genuine, sustained phantom fluid drives the event horizon to contract and its entropy $S \propto A_{\rm hor}$ to fall — the GSL-violating, rip-bound reading. Here, instead, the horizon grows monotonically through the entire phantom-side transient ($\min \dot{R}_E = +2.4 \times 10^{-4}$ over the observable epoch; Appendix A, point (f)) and approaches marginal saturation — $\dot{R}_E \to 0$, the frozen de Sitter area — as the attractor pins $w_\infty = -1$ exactly (Eq. (8)).

The apparent phantom thus costs no horizon entropy. Nor does the transient disturb the acceleration chronology: by the epoch at which $\ddot{a}$ changes sign the equation of state has already relaxed to $w = -1.02$, so the onset of acceleration falls at $z_{\rm acc} = 0.69$ with a dark-energy share of 32.6% — all but indistinguishable from the ΛCDM best fit ($z_{\rm acc} = 0.66$, the classic $\Omega_\Lambda = 1/3$ threshold [46]). The deep apparent phantom, under the horizon reading of the dark-energy history, therefore lives at redshifts where dark energy is still dynamically subdominant, and leaves no fingerprint on when the universe began to accelerate.

5 The ladder of conclusions

The results presented above can be summarized as the following ladder of conclusions, beginning with the most general. Each conclusion is established by a comparison of its own — a dedicated control or a matched model pair — so no rung of the ladder borrows its evidence from another.

Conclusion 1: a fit with dark energy is better than a fit without it. The no-dark-energy control (flat matter+radiation, h derived from the densities) fails catastrophically on every probe: $\chi^2_{\rm tot} = 8994.7$ versus 1811.9 for ΛCDM on the three-probe battery (Table 3; $\Delta\chi^2 = +7183$), at the Einstein–de Sitter point $h_{\rm derived} = 0.450$ and with every free parameter pinned at its prior bound. The physical content of the failure is geometric and gross, not statistical: the CMB acoustic angle forces a matter-only universe to $h \simeq 0.45$, while the supernova and BAO distances at $z \lesssim 1$ require the expansion history of $h \simeq 0.68$ with late acceleration — no reshuffling of the remaining parameters can satisfy both, and the $\sim$7000 units of χ^2 are the price of that contradiction. This rung is of course long established [47, 48]; we include our own measurement of it so that the entire ladder stands on one validated battery. The nontrivial addendum from our family is that *holographic anchoring alone is not the winning ingredient*: the eternal screen — dark energy present at all times with the capacity $1/R_E^2(t)$ — is *worse* than a plain cosmological constant by +33 (Table 1).

Conclusion 2: a fit with an evolving dark energy is better than a fit with a constant. Every dynamical row improves on the constant (ΛCDM): by 4.1–8.9 units on the baseline (Table 1), 3.7–7.9 after the Lyman-α update (Table 2), and 5.2–8.4 on the full three-probe battery (Table 3; ΔAIC down to −4.4) — with the single exception of the canonical $p = 4$ thermodynamic form, which the Lyman-α update drops below the constant (2.2 behind, Table 2) and the full battery leaves at its level (0.4 behind, Table 3). That exception is not a counterexample to this statement: what fails there is the form's late birth date, not its evolution — the same capacity with the exponent freed stays ahead of ΛCDM. The physical pull behind the preference is modest and localized. Every surviving fit settles at $w_0 \simeq -0.95 > -1$, which by Eq. (13) is the statement that the dark-energy density was a few per cent *higher* in the recent past than it is today. The tilt has a

definite observable signature: with the matter sector and the acoustic scale pinned by the CMB, a few per cent more dark energy at $0.3 \lesssim z \lesssim 1$ raises $H(z)$ there and shortens the distances to those same redshifts — supernovae appear slightly brighter, the mid-range $D_M/r_{\rm d}$ rows sit slightly lower — and this is exactly the region where the BAO mid-range and the supernova Hubble-diagram shape overlap at their most precise. A cosmological constant cannot produce the tilt by definition: its density at $z = 0.5$ *is* its density at $z = 0$, and with the CMB holding the matter side fixed, raising Ω_Λ moves the recent past and today together, never one against the other. This reproduces, on an independent single-pipeline battery, the central claim of the DESI-era combined analyses. That claim is the subject of an active and unsettled literature — asserted by the DESI analyses themselves [3], reproduced by independent groups [49], and challenged on supernova-systematics grounds [50].

Conclusion 3: a fit in which dark energy is born is better than any eternal fit. A closer look at the data of Tables 1–3 resolves this statement into three pairwise contrasts, each at matched or near-matched parameter count.

(i) *Born versus eternal screen* (Table 1): the thermodynamic birth beats canonical HDE by 37 units at equal $k = 7$ — the birth, not the capacity, is the load-bearing element. In eternal HDE the single constant $c^2_{\rm HDE}$ must set *both* how much dark energy there is today *and* how fast its fraction grew over history, while the data want a Λ-like present with almost no dark energy deep in the matter era — one dial cannot tune two values. The born forms decouple exactly this: the amplitude s sets the present level, the activation $\alpha(t)$ sets the history.

(ii) *Born versus eternal-smooth tangent* (Table 3): on the full battery with the complete supernova covariance the born forms hold the two top rows (Avrami 1803.5; profile-corrected GC-p 1804.7), the tangent third at 1.7 units behind the leader — and it approaches even that only by migrating deep into the born shape ($w_0 = -0.87$, $w_a = -0.49$). Physically the migration is the tangent tracing the born history inside the observed window: a straight line in $(1 - a)$ can match a curved history only locally, and what it cannot supply — a birth on the early side, an attractor on the late — is what it pays for at the window's edges.

(iii) *Early birth versus late birth* (Tables 2 and 3): the Lyman-α full shape measures the early tail directly (the $D_H \to H(2.33) \to \Omega_{\rm de}(2.33)$ lever of Sec. 4) and excludes the late-birth equilibrium form (+7.6, falling below ΛCDM; Table 2), while the early-birth dates hold a tight corridor $z_c = 4.0$–4.2 that reproduces across two CMB pipelines, three kinetic forms, and three probe combinations (Tables 1–3 and Table 4B). The physics is the tail itself: the expansion rate at $z = 2.33$ carries a small but nonvanishing dark-energy share ($\simeq 3\%$ in the full fits), and a component already present at $z = 2.33$ must have been born comfortably earlier — the data place that birth at $z_c \simeq 4$, where the tail it leaves behind at 2.33 is of exactly the measured size.

Taken together, the three contrasts read as one finding: the absence of dark energy at high redshift is not just an assumption of the born family but a preference of the existing cosmological data.

Conclusion 4: the data build the ladder but do not select a model within its top class. The leading fits of Table 3 span 1.7 units from the born leader to the migrated tangent, and 3.3 across the top four, with born forms first and second. These spreads are comparable to or below the ± 3 cross-pipeline systematic of Sec. 2, so the analysis ends in a statistical tie among born histories and a migrated tangent, with born forms on top. The full STAT+SYS supernova covariance, far from erasing this ordering, strengthened it; the profile likelihoods of Sec. 4.4 supply the headline intervals. The ladder and the tie are the same arithmetic read twice: the gaps that

carry Conclusions 2 and 3 exceed the measured likelihood-choice shift (Table 4B) severalfold, while the intra-class spread sits within the conservative floor — which is the tie. The tie is not an accident of noise but a structural fact: a smooth background measures only the one function $\rho_{\rm de}(z)$, every surviving form has converged on the *same* history (gentle activation from $z_c \simeq 4$, $s\,\alpha_0 \simeq 0.93$, $w_0 \simeq -0.95$; Tables 1–3), and mechanisms that differ only in their microscopic reading of one and the same curve are indistinguishable at the background level by construction. Nor do the data decide the two structural questions underneath: *how* the birth proceeded (domain nucleation and front growth, or global ordering) and *where* the born component resides (a condensate of the horizon screen, or a bulk component). Both questions, and the null probes that bound them, are taken up in Sec. 6.

6 Discussion and open questions

The ladder of Sec. 5 sharpens the DESI-era statement "dark energy evolves" into a statement about the *shape* of its history. Among the histories tested, the data prefer one in which dark energy did not exist at high redshift and activated at $z_c \simeq 4$, growing on the causal clock of the horizon scale. That history wins over constants, eternal holographic screens, and late thermodynamic activations — and, by a statistically insignificant but directionally consistent margin, over the free tangents, which the three probes jointly push into the born family's corner. That statement is the central result of this paper. In this closing section we nevertheless allow ourselves deliberations on a somewhat wider circle of questions.

Independent corroboration. A substantial part of the conclusions of this work rests on comparing the fits obtained for the original DESI DR2 data against those for the same data with the updated Lyman-α reduction (Sec. 4). That replacement costs $+1.15$ to $+2.52$ units of χ^2 to precisely the corner of parameter space our fits occupy. It is therefore natural to ask whether the pull survives outside this pipeline. A contemporaneous, independent analysis of the same DR2 Lyman-α measurement combined with CMB and galaxy BAO [49] finds deviations from ΛCDM of up to 3.1σ, entirely in that quadrant ($w_0 > -1$, $w_a < 0$, $w_0 + w_a < -1$) — the apparent-phantom corner of Sec. 4.4. Two differences are instructive. First, in that analysis adding supernovae *dilutes* the signal below 2σ, which its authors read as inconclusive; in our battery the supernovae sharpen it — because here they test a prediction rather than adjust coefficients. The divergence prices exactly the difference between fitting parameterizations and testing mechanisms: a parameterization has no stake in the shape of the Hubble diagram, and marginalizing its freedom spreads the signal, whereas a mechanism — here, a condensation or an order–disorder transition of the horizon screen — predicts that shape: the born form collects $\Delta\chi^2_{\rm SN} = -1.8$ where its own tangent collects -0.9 (Fig. 3C). Second, that analysis uses the pre-update Lyman-α row, so the quadrant agrees across the row choice while the significance should be read as row-sensitive. The corroboration we take from it is deliberately minimal: wherever the DR2 Lyman-α point is combined with CMB and BAO, by whatever pipeline, the data walk away from the cosmological constant in the same direction.

Chronological context. The activation epoch itself invites an astrophysical connection: $z_c = 4.2^{+1.0}_{-0.3}$ (Sec. 4.4) falls in the busiest stretch of cosmic history, with reionization just completed ($z \simeq 5.3$–6 [51]) and the global star-formation rate still climbing toward its $z \simeq 2$ maximum [52]. We note the coincidence of dates and immediately fence it: at activation the dark-energy fraction is a few per cent ($\Omega_{\rm de}(z{=}2.33) = 2.6$–$4.3\%$, Sec. 4.4) — a $\sim 1.5\%$ perturbation to the expansion

rate, far too small to inhibit collapse-scale astrophysics, whose bound systems are decoupled from the expansion in any case [53]. And no known astrophysical process can source a density locked to $1/R_E^2(t)$. At face value the overlap is therefore a chronological coincidence — distinct processes that happen to unfold within the same few e-folds — and we record it as chronology, not as a hint of mechanism. The only firm causal statement is standard and late: once dark energy dominates ($z \simeq 0.3$–0.7), accelerated expansion freezes the growth of new structure, and the decline of star formation from its $z \simeq 2$ peak unfolds in the same era — but the born histories inherit that chronology from ΛCDM essentially unchanged, so it does not discriminate among the ladder's survivors.

The horizon reading: a mild preference, honestly priced. We chose the horizon parameterization for economy, and the economy is real. More than that, it names every measured number: z_c as the redshift of the dark-energy birth; T_c as a fraction of the Gibbons–Hawking temperature; u as a speed in light crossings (see Sec. 3); the bounded exponent p as a spectral index of the condensing modes; and so on. It also offers a possible resolution of two classic puzzles of cosmological numerology: the value of the cosmological constant — $\rho_\Lambda^{1/4} = \sqrt{E_{\rm pl} T_{\rm dS}}/\sqrt{2\pi}$, with ρ_Λ the observed dark-energy density and $E_{\rm pl}$ the Planck energy, so the meV scale is a geometric mean, not a neutrino coincidence; and the HDE $c_{\rm HDE} < 1$ preference, which is the mid-transition reading of a quantity driven to unity. A bulk reading could carry the same numbers, but only as coincidences. Moreover, as a matter of accuracy, a bulk reading cannot be built on minimal quintessence: the fitted histories cross $w = -1$ (the equilibrium class at the parameter-free filling of Eq. (9)), and a single minimally coupled scalar field cannot cross the phantom divide [54]. So the volume carrier would have to be something more elaborate: a two-field (quintom), non-minimally coupled, or interacting construction. The horizon family exhibits the apparent phantom crossing by construction — and, as Sec. 4.4 shows, at no cost in horizon entropy. Nor does the reading stand alone: de Sitter as a two-fluid thermal state is an independent, active line of work [34]; the present family differs in placing the two fluids on the horizon screen itself, with a critical temperature the data can measure. This is a preference of interpretation, not of χ^2, and we leave it at that strength.

The Hubble tension. One of the standing problems in the interpretation of cosmological data is the so-called Hubble tension: the difference between the H_0 values inferred from the CMB-anchored fits ($h \simeq 0.68$) and measured locally, through Cepheid-calibrated supernovae (0.730 ± 0.010 [21]). This analysis neither addresses that tension nor is contaminated by it, both by construction. It is not contaminated by virtue of the data selection: the SH0ES calibrators are excluded and the supernova absolute offset is marginalized analytically (Sec. 2), so the battery measures the *shape* of the expansion history, never the local H_0. Nor is the tension addressed, since dark energy is absent at $z \gtrsim 5$, recombination and the drag horizon $r_{\rm d}$ are untouched, so the CMB-anchored Hubble constant is inherited essentially unchanged ($h = 0.679$ for the leading fit against 0.684 for ΛCDM; Table 3B of Appendix B); and with the acoustic angular scale θ_* fixed by the CMB and the DESI DR2 distances pinning the mid-redshift expansion, no late-time history retains the freedom to raise H_0 without breaking the very BAO agreement it was fitted on. Our battery thus confirms, on born histories rather than tangents, the established verdict that no late-time ($z \lesssim 5$) dark-energy dynamics can resolve the tension [55]; whatever does must act earlier, and would, to first order, move every row of our tables coherently, leaving the ladder's internal contrasts intact.

Open questions. The task of this paper is sharply delimited: to test a family of physically motivated dark-energy histories against the available cosmological data. As such, it leaves open

several substantial questions.

(i) *Where*: screen versus bulk. By the bulk we mean, precisely, any field-theoretic dynamics in the volume whose $\rho_{\rm de}(z)$ reproduces the same history, with no need to reference the horizon at all: the background measures only $\rho(t)$, and $R_E^2(t)$ in Eq. (1) may be a physical capacity or merely a scale such a field has locked onto. We have touched upon this question above — the bulk reading of the measured numbers earlier in this section, and the apparent phantom of Sec. 4.4. Here we only repeat: at present the data cannot answer it. For that reason it remains a question of theory, out of the scope of the present phenomenological study.

(ii) *How*: a kinetic arbitration performed on the self-consistent backgrounds — five growth laws tested against the fitted filling curve — excluded diffusive growth, curvature coarsening, and undercooling-driven fronts. A relaxational time-dependent Ginzburg–Landau (TDGL) analysis excluded uniform lagged dynamics both ways (both analyses are documented in Appendix A). The surviving kinetic readings of Eq. (6) are degenerate between front-mediated filling from N_0 seeds and global ordering (Sec. 3). The sky is silent on seams — the boundaries where fronts from neighboring seeds would meet, leaving a directional deficit of filling and hence of $\rho_{\rm de}$: the limits are $\sim 15\,\mu$K at $\ell \leq 8$ [2], while realistic seams imprint only $\sim 1\,\mu$K — a seam couples to the CMB solely through the late integrated Sachs–Wolfe effect ($\sim 10\,\mu$K at these ℓ), scaled down by its core deficit $\delta \sim 0.1$ — under cosmic variance forever. The expansion-anisotropy channel is bounded at $|\delta| < 0.06$ along the one direction a hint briefly suggested (refuted at 3.5σ by an independent catalog [5] within a day of being posed). So the *how*, unlike the *where*, remains a question of data: the background cannot separate the surviving readings, the CMB channel is closed by cosmic variance, and the one channel still live — directional expansion anisotropy at the sub-percent level — awaits the next generation of surveys.

(iii) *The missing theoretical ingredient of the horizon reading.* As noted in the previous item, for the histories motivated by a first-order transition the key test is observational: the search for directional signatures of the hypothetical boundaries between the horizon phases — no longer in the CMB, whose channel is closed by cosmic variance, but in the expansion anisotropy $H(\hat{n})$. For the histories motivated by a continuous, second-order transition the missing ingredient is instead theoretical: a single number, $T_c/T^0_{\rm dS}$, whose computation would convert the best rows of Table 3 from parameterization into prediction. At least three candidate routes to T_c can be proposed: pairing on the screen; a Kosterlitz–Thouless (KT) transition [56] of the screen's superfluid stiffness ρ_s, with the fitted bound $\rho_s(0) > 4.9\,T_c$; and a holographic-superconductor calculation, where T_c follows from the instability threshold of a bulk scalar in the anti–de Sitter (AdS) case [57, 58]. The dimensionality itself argues for the KT route: on a two-dimensional screen the Mermin–Wagner theorem [59] forbids true long-range order at any finite temperature, so an ordering transition there can only be topological — a gas of logarithmically bound vortex pairs gradually unbinding — and such a transition carries exactly the soft excitation spectrum the p-profile selects ($p \lesssim 1$), disfavoring the rigid canonical $p = 4$ on first principles, independently of the data that also exclude it. The holographic route supplies the mechanism in its best-understood form: in AdS/CFT a bulk scalar turns unstable once the horizon cools below a critical temperature, and a condensate forms at the boundary [57, 58]; the thermodynamic class of Sec. 3 is that mechanism read on the de Sitter horizon, where the corresponding instability threshold is precisely the number still missing.

Despite the arguments above, in this work we deliberately refrain from a final verdict on the mechanism, repeating once more: such a choice can only be made with the extended set of experimental data discussed below.

Forthcoming experimental data. Many of the questions posed in this work can be resolved by the new cosmological data now expected. Most directly, the DESI five-year Lyman-α measurement [1] (plausibly mid-2027) will test the early-birth tail, now a two-sided band ($\Omega_{\rm de}(2.33) = 2.6$–$4.3\%$ at 68%), and the activation bound $z_c > 3.4$ (both intervals from Sec. 4.4). The supernova surveys of the Vera C. Rubin Observatory [60] will test today's equation of state, where the born family now predicts $w_0 = -0.95$ to -0.96 along the profile ridge (marginal posterior -0.96 ± 0.03; Sec. 4.4) against the migrated tangent's -0.87 (Table 3B of Appendix B). Finally, the anticipated precision gains in the expansion-anisotropy channel (sub-percent $H(\hat{n})$ from the Rubin and DESI supernova and peculiar-velocity programs) may help resolve the class of the mechanism behind the dark-energy inception: the global-ordering and equilibrium readings predict strictly smooth, isotropic dark energy — zero expansion anisotropy at any precision — while the front reading admits seam-type directional modulation, bounded so far at $\delta \lesssim 0.06$–0.1 (item (ii) above). A confirmed $H(\hat{n})$ anisotropy would therefore not refute the family but decide *how* dark energy was born — in favor of a first-order transition.

Acknowledgments. The computational campaign — likelihood construction and validation, solver design, battery runs, sky and supernova tests, and the session notes from which this draft was assembled — was carried out in an extended human–AI collaboration (Anthropic Claude), with all physical claims, strategic decisions, and final text responsibility resting with the author.

References

[1] DESI Collaboration, "The DESI experiment Part I: science, targeting, and survey design," arXiv:1611.00036 (2016).

[2] Planck Collaboration, "Planck 2018 results. V. CMB power spectra and likelihoods," Astron. Astrophys. **641**, A5 (2020).

[3] DESI Collaboration, "DESI DR2 results II: Measurements of baryon acoustic oscillations and cosmological constraints," Phys. Rev. D **112**, 083515 (2025).

[4] D. Brout et al., "The Pantheon+ analysis: cosmological constraints," Astrophys. J. **938**, 110 (2022).

[5] DES Collaboration, "The Dark Energy Survey: cosmology results with ∼1500 new high-redshift type Ia supernovae using the full 5-year data set," Astrophys. J. Lett. **973**, L14 (2024).

[6] M. Chevallier, D. Polarski, "Accelerating universes with scaling dark matter," Int. J. Mod. Phys. D **10**, 213 (2001).

[7] E. V. Linder, "Exploring the expansion history of the universe," Phys. Rev. Lett. **90**, 091301 (2003).

[8] W. Rindler, "Visual horizons in world models," Mon. Not. R. Astron. Soc. **116**, 662 (1956).

[9] G. W. Gibbons, S. W. Hawking, "Cosmological event horizons, thermodynamics, and particle creation," Phys. Rev. D **15**, 2738 (1977).

[10] A. G. Cohen, D. B. Kaplan, A. E. Nelson, "Effective field theory, black holes, and the cosmological constant," Phys. Rev. Lett. **82**, 4971 (1999).

[11] M. Li, "A model of holographic dark energy," Phys. Lett. B **603**, 1 (2004).

[12] S. D. H. Hsu, "Entropy bounds and dark energy," Phys. Lett. B **594**, 13 (2004).

[13] J. Torrado, A. Lewis, "Cobaya: code for Bayesian analysis of hierarchical physical models," JCAP **05**, 057 (2021).

[14] E. Rosenberg, S. Gratton, G. Efstathiou, "CMB power spectra and cosmological parameters from Planck PR4 with CamSpec," Mon. Not. R. Astron. Soc. **517**, 4620 (2022).

[15] S. Brieden, H. Gil-Marín, L. Verde, "A tale of two (or more) h's," JCAP **04**, 023 (2023).

[16] G. Mangano, G. Miele, S. Pastor, T. Pinto, O. Pisanti, P. D. Serpico, "Relic neutrino decoupling including flavour oscillations," Nucl. Phys. B **729**, 221 (2005).

[17] C. Alcock, B. Paczyński, "An evolution free test for non-zero cosmological constant," Nature **281**, 358 (1979).

[18] DESI Collaboration, "DESI DR2 Results IV: Alcock–Paczyński measurements from the Lyman alpha forest and cosmological constraints," (2026), arXiv:2607.27410.

[19] D. Scolnic et al., "The Pantheon+ analysis: the full data set and light-curve release," Astrophys. J. **938**, 113 (2022).

[20] A. Carr et al., "The Pantheon+ analysis: improving the redshifts and peculiar velocities of type Ia supernovae used in cosmological analyses," Publ. Astron. Soc. Aust. **39**, e046 (2022).

[21] A. G. Riess et al., "A comprehensive measurement of the local value of the Hubble constant," Astrophys. J. Lett. **934**, L7 (2022).

[22] R. J. Cooke, M. Pettini, C. C. Steidel, "One percent determination of the primordial deuterium abundance," Astrophys. J. **855**, 102 (2018).

[23] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, T. Schwetz, A. Zhou, "The fate of hints: updated global analysis of three-flavor neutrino oscillations," JHEP **09**, 178 (2020).

[24] H. Akaike, "A new look at the statistical model identification," IEEE Trans. Autom. Control **19**, 716 (1974).

[25] G. Schwarz, "Estimating the dimension of a model," Ann. Statist. **6**, 461 (1978).

[26] J. A. Nelder, R. Mead, "A simplex method for function minimization," Comput. J. **7**, 308 (1965).

[27] A. Lewis, A. Challinor, A. Lasenby, "Efficient computation of CMB anisotropies in closed FRW models," Astrophys. J. **538**, 473 (2000).

[28] W. Fang, W. Hu, A. Lewis, "Crossing the phantom divide with parameterized post-Friedmann dark energy," Phys. Rev. D **78**, 087303 (2008).

[29] J. Friedman, M. S. Morris, I. D. Novikov, F. Echeverria, G. Klinkhammer, K. S. Thorne, U. Yurtsever, "Cauchy problem in spacetimes with closed timelike curves," Phys. Rev. D **42**, 1915 (1990).

[30] S. W. Hawking, T. Hertog, "Populating the landscape: a top-down approach," Phys. Rev. D **73**, 123527 (2006).

[31] A. N. Kolmogorov, "On the statistical theory of crystallization of metals," Izv. Akad. Nauk SSSR, Ser. Mat. **1**, 355 (1937).

[32] M. Avrami, "Kinetics of phase change. I. General theory," J. Chem. Phys. **7**, 1103 (1939).

[33] C. P. Herzog, P. K. Kovtun, D. T. Son, "Sound modes in holographic superfluids," Phys. Rev. D **79**, 066002 (2009).

[34] G. E. Volovik, "From Landau two-fluid model to de Sitter Universe," Phys. Usp. **68**, 1074 (2025).

[35] C. J. Gorter, H. Casimir, "On supraconductivity I," Physica **1**, 306 (1934).

[36] Q.-G. Huang, Y. Gong, "Supernova constraints on a holographic dark energy model," JCAP **08**, 006 (2004).

[37] S. Wang, Y. Wang, M. Li, "Holographic dark energy," Phys. Rept. **696**, 1 (2017).

[38] J. Goodman, J. Weare, "Ensemble samplers with affine invariance," Comm. Appl. Math. Comput. Sci. **5**, 65 (2010).

[39] D. Foreman-Mackey, D. W. Hogg, D. Lang, J. Goodman, "`emcee`: The MCMC Hammer," Publ. Astron. Soc. Pac. **125**, 306 (2013).

[40] V. Sahni, A. Starobinsky, "Reconstructing dark energy," Int. J. Mod. Phys. D **15**, 2105 (2006).

[41] S. M. Carroll, M. Hoffman, M. Trodden, "Can the dark energy equation-of-state parameter w be less than -1?," Phys. Rev. D **68**, 023509 (2003).

[42] J. M. Cline, S. Jeon, G. D. Moore, "The phantom menace: constraints on low-energy effective ghosts," Phys. Rev. D **70**, 043543 (2004).

[43] R. R. Caldwell, "A phantom menace? Cosmological consequences of a dark energy component with super-negative equation of state," Phys. Lett. B **545**, 23 (2002).

[44] R. R. Caldwell, M. Kamionkowski, N. N. Weinberg, "Phantom energy: dark energy with $w < -1$ causes a cosmic doomsday," Phys. Rev. Lett. **91**, 071301 (2003).

[45] S. M. Carroll, A. Chatwin-Davies, "Cosmic equilibration: a holographic no-hair theorem from the generalized second law," Phys. Rev. D **97**, 046012 (2018).

[46] S. Weinberg, *Cosmology* (Oxford University Press, Oxford, 2008).

[47] A. G. Riess et al., "Observational evidence from supernovae for an accelerating universe and a cosmological constant," Astron. J. **116**, 1009 (1998).

[48] S. Perlmutter et al., "Measurements of Ω and Λ from 42 high-redshift supernovae," Astrophys. J. **517**, 565 (1999).

[49] S. Capozziello, H. Chaudhary, G. Mustafa, S. K. J. Pacif, "Evidence of dynamical dark energy found via the DESI DR2 Lyman-α forest," Astron. Astrophys. **709**, A258 (2026).

[50] G. Efstathiou, "Evolving dark energy or supernovae systematics?," Mon. Not. R. Astron. Soc. **538**, 875 (2025).

[51] S. E. I. Bosman et al., "Hydrogen reionization ends by $z = 5.3$: Lyman-α optical depth measured by the XQR-30 sample," Mon. Not. R. Astron. Soc. **514**, 55 (2022).

[52] P. Madau, M. Dickinson, "Cosmic star-formation history," Annu. Rev. Astron. Astrophys. **52**, 415 (2014).

[53] M. Carrera, D. Giulini, "Influence of global cosmological expansion on local dynamics and kinematics," Rev. Mod. Phys. **82**, 169 (2010).

[54] A. Vikman, "Can dark energy evolve to the phantom?," Phys. Rev. D **71**, 023515 (2005).

[55] L. Knox, M. Millea, "Hubble constant hunter's guide," Phys. Rev. D **101**, 043533 (2020).

[56] J. M. Kosterlitz, D. J. Thouless, "Ordering, metastability and phase transitions in two-dimensional systems," J. Phys. C **6**, 1181 (1973).

[57] S. S. Gubser, "Breaking an Abelian gauge symmetry near a black hole horizon," Phys. Rev. D **78**, 065034 (2008).

[58] S. A. Hartnoll, C. P. Herzog, G. T. Horowitz, "Building a holographic superconductor," Phys. Rev. Lett. **101**, 031601 (2008).

[59] N. D. Mermin, H. Wagner, "Absence of ferromagnetism or antiferromagnetism in one- or two-dimensional isotropic Heisenberg models," Phys. Rev. Lett. **17**, 1133 (1966).

[60] Ž. Ivezić et al. (LSST Collaboration), "LSST: from science drivers to reference design and anticipated data products," Astrophys. J. **873**, 111 (2019).

[61] L. P. Gor'kov, G. M. Eliashberg, "Generalization of the Ginzburg–Landau equations for non-stationary problems in the case of alloys with paramagnetic impurities," Zh. Eksp. Teor. Fiz. **54**, 612 (1968) [Sov. Phys. JETP **27**, 328 (1968)].

Appendix A: The de Sitter attractor and solver integrity

In the family of Eq. (1) the screen radius R_E is the *future event* horizon, so today's density depends on the expansion still to come (Sec. 3). The solver resolves this closed loop not by integrating backward from the future but by relaxation on a fixed grid in $x = \ln a$ with the analytic de Sitter tail $\int_{t_f}^{\infty} dt/a = 1/(a_f H_f)$ at the endpoint. Six checks, run at the Avrami best fit of the full battery (`attractor_check_v1.py`; pure background, no Boltzmann code), make the resulting de Sitter attractor and the solver's integrity quantitative.

(a) The node is reached. On a grid extended to $x_{\max} = 25$ ($a \sim 7 \times 10^{10}$), the diagnostics $y \equiv R_E H \to 1$ and $\alpha \to 1$ hold to machine precision from $x \simeq 5$ onward ($|y - 1| = 2 \times 10^{-16}$ at $x = 20$), with $w \to -1$ limited only by grid differentiation ($|w + 1| = 2 \times 10^{-6}$) and H constant at the 10^{-4} level over twenty e-folds.

(b) No big rip. In eternal HDE the asymptotic equation of state is $w_\infty = -1/3 - 2/(3c_{\rm HDE})$, phantom for the observationally preferred $c_{\rm HDE} < 1$ [11]. Here the effective amplitude $c_{\rm HDE}^{\rm eff} \equiv \sqrt{s\,\alpha}$ runs from 0.98 today to exactly 1 as $z \to -1$, so $w_\infty = -1$: the $c_{\rm HDE} < 1$ reading is a mid-transition snapshot, not an asymptotic constant. The activation edge does produce a formally divergent w (the log-derivative of a near-zero density, at $\Omega_{\rm de} \sim 10^{-59}$). The divergence is a ratio artifact, not a physical one: by the continuity identity the dark-energy pressure is $-\rho - \frac{1}{3}\, d\rho/dx$, so at the Avrami edge, where $\rho \propto (x - x_c)^2$, density and pressure both vanish (the pressure linearly) and only their ratio w diverges. For the thermodynamic class the edge zero of the density is linear at any value of the critical exponent $p > 0$ (Eq. (7) opens as $s\,p\,\varepsilon$ near T_c), so the pressure instead tends to a finite negative constant. In either class nothing that gravitates diverges. A rip would require the *density* to diverge in the *future*; here the only divergence is the ratio w at the past edge, where the density vanishes. The fitting pipeline nevertheless clips the $w(a)$ table to $[-3, -1/3]$ before it reaches the Boltzmann code — a clip acting only in this numerically empty region, where dark energy has strictly no effect on the expansion or the perturbations, so no $\Omega_{\rm de}(z)$ integral is biased — and the physically meaningful phantom transient is bounded: $w_{\min} = -1.36$ where $\Omega_{\rm de} > 10\%$ (at $z \simeq 1.6$), relaxing to $w_0 = -0.95$ today.

(c) Stability of the shape. A localized perturbation injected into the future (+10% and +50% bumps at $x = 10$) decays geometrically under the relaxation map, returning the attractor diagnostics to $|y - 1| < 10^{-7}$ within 160 iterations with the matched h unchanged. This once again demonstrates numerically that the de Sitter node is a stable attractor of the self-consistency map.

(d) The teleological family and its resolution. A uniform rescaling of the far-future density is *exactly* self-consistent: in the pure de Sitter regime the map $\rho \mapsto \alpha/R_E^2$ is scale-covariant ($\rho \to \mu\rho$ for an arbitrary constant $\mu > 0$ gives $H \to \sqrt{\mu}\, H$ and $R_E \to R_E/\sqrt{\mu}$), so the teleological loop admits a one-parameter family of solutions — the choice of H_∞. This is the expected freedom, and the data fix it: the solver's bisection matches h today, and the converged h is strictly monotone in the seed — the uniform initial density from which the relaxation starts — across its full range (eight seeds spanning $[0.02, 2.4]$ give $h \in [0.43, 1.57]$ monotonically), so the matched branch exists and is unique. Stability in (c) is transverse to this family; along it there is no dynamics to stabilize, only a boundary condition to impose.

(e) Insensitivity to the future cutoff. Moving the grid end from $x_{\max} = 13$ (the value used in every fit of this paper) to 20 and 25 shifts $R_E(z{=}0)$ by 2×10^{-5} relative — today's observables do not depend on how much numerical future is attached beyond the analytic tail.

(f) Generalized second law. The generalized second law requires the total entropy — horizon ($S \propto$ horizon area [9]) plus matter — not to decrease. Given the negligible entropy contribution of cosmic matter at late times, the requirement is asymptotically driven by the geometric growth of the horizon area, i.e. $\dot{R}_E \geq 0$. The same law is known to enforce the approach to de Sitter as a holographic no-hair theorem [45], so checks (a) and (c) are, with this one, the numerical face of one statement. It holds through the phantom transient with margin ($\min \dot{R}_E = +2.4 \times 10^{-4}$ over the observable epoch $z < 25$); asymptotically $\dot{R}_E \to 0$ — marginal saturation, horizon area frozen in de Sitter — with the residual -3×10^{-6} plateau shrinking fourfold under grid doubling ($O(\Delta x^2)$ discretization, a numerical zero).

Kinetic arbitration of the growth laws. Beyond the six integrity checks, the same background machinery supports the kinetic arbitration quoted in Sec. 6. The arbitration proceeds in two stages.

The first, described here, is a parameter-free shape test — it arbitrates the kinetics of the filling, not the cosmology: five growth laws for the covered fraction α were integrated on the same self-consistent background (the Avrami best fit of the baseline battery) and normalized to the common half-filling epoch, after which a law has no freedom left — its filling curve either tracks the fitted Avrami curve or fails outright. The comparison metrics over the data window $z < 2.6$ are the effective Avrami exponent $n_{\rm eff}$ (the log–log slope of $-\ln(1-\alpha)$ against $x - x_c$) and the root-mean-square (RMS) deviation in α over that window.

The results are as follows (`kinetics_arbitration_v1.py`, in the script release at https://github.com/yryabov/horizon-ladder). A front advancing a fixed fraction of the horizon per e-fold ($v \propto HR_E$) reproduces the curve exactly ($n_{\rm eff} = 2.00$, RMS < 0.001). A constant physical front speed survives at a level the data cannot distinguish ($n_{\rm eff} = 2.66$, RMS 0.062 — its battery fit is the Front/clock row of the tables). Undercooling-driven speed $v \propto \varepsilon$ ($n_{\rm eff} = 4.26$, RMS 0.131), diffusive growth $L = \sqrt{Dt}$ read as an angular size on the growing screen ($n_{\rm eff} = 0.94$, RMS 0.135), and curvature coarsening at fixed comoving scale ($n_{\rm eff} = 1.12$, RMS 0.112) are excluded. Their failure is one of shape, not size: where the fitted filling demands an exponent $n \approx 2$, they supply $n \approx 1$ or $n \approx 4$ — a curvature of the history that the one remaining normalization cannot repair — and their residual distortions of α, at the 0.1 level across the window, are twice the largest deviation the batteries certify as viable (the Front/clock row).

Uniform (frontless) relaxational dynamics was tested separately, as time-dependent Ginzburg–Landau relaxation toward the equilibrium curve with critical slowing, $\dot{\alpha} = -[\alpha - \alpha_{\rm eq}(T)]/\tau$, $\tau = \kappa/(T_c\,\varepsilon)$, where the dimensionless coefficient κ measures the relaxation time in units of the thermal time $\hbar/(k_B T_c)$. Matching the fitted shape demands $\kappa \simeq 0.04$, several times below both the causal (Planckian) floor and the value $\kappa = \pi/8$ derived from Bardeen–Cooper–Schrieffer (BCS) theory [61]. At the shape-fitted critical temperature a causally honest κ destroys the fit (RMS 0.3–0.5). Uniform lagged dynamics is thus excluded in both directions.

The second stage is the batteries of Sec. 4 themselves: the two laws that pass the shape test are precisely the two surviving forms of Sec. 3 — growth on the e-fold clock (Eq. (5)) and the causal front/clock form (Eq. (6), which itself carries the two degenerate readings discussed in Secs. 3 and 6) — and both were fitted in full — the Avrami and Front/clock rows of the tables — where they tie within the systematic floor. The three excluded laws never reach the second stage: they fail at the shape level, with no parameter left to repair them.

Appendix B: Complete best-fit parameters

For each battery of Sec. 4 the tables below list the complete parameter vector at every model's best fit: the sampled cosmological parameters (ω_b, ω_c, Σm_ν, h, A_s, n_s), the sampled nuisances (τ, $A_{\rm Planck}$; excluded from k, Sec. 2), the sampled model parameters, and derived quantities (marked ‡: h of the no-dark-energy control from flatness; s from the attractor identity, Eq. (8); z_c of the equilibrium forms from the T_c crossing; $T_c/T^0_{\rm dS}$ from the fitted T_c and the model's own present-day horizon). The sampled count reproduces each row's k; the per-fit verdict files (JSON) ship with the released repository (`https://github.com/yryabov/horizon-ladder`).

The cross-pipeline pairs behind the ± 3 floor of Sec. 2 are collected in Table 4B: the Avrami form and ΛCDM, each refitted on the Planck PR4/NPIPE CamSpec TTTEEE likelihood plus DESI DR2 BAO, with the $z = 2.33$ row in both conventions and the nine CamSpec calibration and foreground parameters free. The fitted gaps — -6.8 with the BAO-only row, -5.4 with the joint row — stand against the `plik_lite` gaps of -7.7 (Table 1) and -6.4 (Table 2): likelihood-choice shifts of $+0.9$ and $+1.0$ units of χ^2. The birth date is pipeline-stable at the headline value — $z_c = 4.11$, $\lambda = 0.92$ on both rows (quoted in Sec. 4) — with $h = 0.681$ and $\Sigma m_\nu = 0.063$ eV; the matching ΛCDM refits sit at $h = 0.682$. The pipeline is validated in-process: the direct spectrum injection reproduces the `cobaya` anchor at the reference point exactly (zero pipeline delta), and the BAO blocks of all four fits reproduce under the released solver to 10^{-4}. The full 18- and 16-parameter vectors ship as verdict files with the repository.

Finally, a Fisher cross-check of the uncertainties. Fisher-matrix estimates (Gaussian uncertainties from finite-difference second derivatives of χ^2) at the best fits agree with the profile likelihoods of Sec. 4.4 on the constrained direction ($\sigma(z_c) = 0.20$ against the profile's lower arm 0.26) but are unreliable for the solver-based models along the CMB-degenerate directions (the self-consistent solver's bisection tolerance corrupts cross-derivatives); they are retained only for the analytic-background controls, in the released verdict files, and no number in the paper rests on them.

Table 1B: Full parameter vectors at the Table 1 best fits (baseline battery).

	Front/clock	CPL	Avrami	GC-p	Thermal GC	ΛCDM	HDE
ω_b	0.02237	0.02243	0.02240	0.02241	0.02229	0.02254	0.02265
ω_c	0.12042	0.11934	0.12000	0.11972	0.12133	0.11772	0.11585
Σm_ν [eV]	0.0603	0.0624	0.0589	0.0611	0.0582	0.0582	0.0886
h	0.6707	0.6607	0.6827	0.6777	0.6795	0.6840	0.7330
$10^9 A_s$	2.135	2.130	2.133	2.131	2.136	2.127	2.125
n_s	0.9636	0.9661	0.9646	0.9653	0.9615	0.9700	0.9745
τ	0.0620	0.0621	0.0621	0.0620	0.0614	0.0631	0.0649
$A_{\rm Planck}$	1.0004	1.0006	1.0005	1.0005	1.0003	1.0008	1.0011
z_c	4.05	—	4.01	$2.60^\ddagger$	$2.02^\ddagger$	—	—
λ	—	—	0.901	—	—	—	—
u	1.628	—	—	—	—	—	—
α_0	—	—	—	0.639	0.800	—	—
$T_c/T^0_{\rm dS}$	—	—	—	$1.66^\ddagger$	$1.50^\ddagger$	—	—
p	—	—	—	2.008	4 (fixed)	—	—
s	$1^\ddagger$	—	$1^\ddagger$	$1.474^\ddagger$	$1.188^\ddagger$	—	—
w_0	—	−0.708	—	—	—	—	—
w_a	—	−0.913	—	—	—	—	—
$c_{\rm HDE}$	—	—	—	—	—	—	0.542

Table 2B: Full parameter vectors at the Table 2 best fits (second battery).

	CPL	Avrami	Front/clock	GC-p	ΛCDM	Thermal GC
ω_b	0.02243	0.02242	0.02235	0.02244	0.02254	0.02226
ω_c	0.11944	0.11967	0.12065	0.11940	0.11783	0.12181
Σm_ν [eV]	0.0612	0.0582	0.0611	0.0605	0.0582	0.0589
h	0.6604	0.6792	0.6714	0.6742	0.6835	0.6805
$10^9 A_s$	2.135	2.126	2.134	2.131	2.127	2.132
n_s	0.9659	0.9655	0.9631	0.9661	0.9698	0.9603
τ	0.0629	0.0608	0.0617	0.0624	0.0630	0.0600
$A_{\rm Planck}$	1.0007	1.0006	1.0004	1.0005	1.0007	1.0002
z_c	—	4.17	4.02	$2.75^\ddagger$	—	$1.96^\ddagger$
λ	—	0.904	—	—	—	—
u	—	—	1.605	—	—	—
α_0	—	—	—	0.664	—	0.789
$T_c/T^0_{\rm dS}$	—	—	—	$1.71^\ddagger$	—	$1.48^\ddagger$
p	—	—	—	2.034	—	4 (fixed)
s	—	$1^\ddagger$	$1^\ddagger$	$1.424^\ddagger$	—	$1.200^\ddagger$
w_0	−0.712	—	—	—	—	—
w_a	−0.893	—	—	—	—	—

Table 3B: Full parameter vectors at the Table 3 best fits (three-probe battery). The GC-p column quotes the profile-located class minimum of Sec. 4.4. *Parameter pinned at its prior bound (the Einstein–de Sitter point of Sec. 4).

	Avrami	GC-p	CPL	Front/clock	ΛCDM	Thermal GC	No dark energy
ω_b	0.02240	0.02242	0.02244	0.02235	0.02252	0.02226	0.018*
ω_c	0.12001	0.11958	0.11927	0.12071	0.11802	0.12183	0.180*
Σm_ν [eV]	0.0591	0.0638	0.0584	0.0608	0.0582	0.0586	0.40*
h	0.6791	0.6796	0.6773	0.6703	0.6827	0.6807	0.450‡
$10^9 A_s$	2.131	2.135	2.134	2.136	2.128	2.132	2.264
n_s	0.9646	0.9656	0.9662	0.9629	0.9694	0.9604	0.90*
τ	0.0616	0.0629	0.0630	0.0617	0.0630	0.0601	0.040*
$A_{\rm Planck}$	1.0005	1.0006	1.0006	1.0007	1.0008	1.0000	0.992
z_c	4.113	2.97‡	—	4.055	—	1.96‡	—
λ	0.903	—	—	—	—	—	—
u	—	—	—	1.620	—	—	—
α_0	—	0.364	—	—	—	0.788	—
$T_c/T^0_{\rm dS}$	—	1.76‡	—	—	—	1.48‡	—
p	—	0.8	—	—	—	4 (fixed)	—
s	1‡	2.558‡	—	1‡	—	1.201‡	—
w_0	—	—	−0.873	—	—	—	—
w_a	—	—	−0.488	—	—	—	—

Table 4B: The cross-pipeline pairs behind the ±3 systematic floor (Sec. 2): both models refitted on the PR4/NPIPE CamSpec TTTEEE likelihood + DESI DR2 BAO, with the $z = 2.33$ row in its BAO-only (baseline) and joint AP+BAO conventions. Totals include the Gaussian priors of Sec. 2; gap = Avrami − ΛCDM.

Lyman-α row	Model	$\chi^2_{\rm CamSpec}$	$\chi^2_{\rm BAO}$	$\chi^2_{\rm tot}$	gap
BAO-only	Avrami	10541.61	9.62	10551.24	−6.82
BAO-only	ΛCDM	10545.50	12.46	10558.06	
joint AP+BAO	Avrami	10541.61	12.26	10553.88	−5.42
joint AP+BAO	ΛCDM	10545.39	13.83	10559.30	

Appendix C: The data behind the batteries, and the posterior

This appendix collects the figures that document the batteries' data blocks and the posterior validation — material that supports, but does not alter, the conclusions of the main text. Figure 1C shows the three best fits against the Planck `plik_lite` bandpowers, residual by residual and bin by bin; Figure 2C the bandpowers themselves, the 613 CMB numbers of every battery; Figure 3C the Pantheon+ Hubble diagram in residuals, under the battery's supernova conventions; and Figure 4C the marginalized posterior of the leading form behind the forecast-grade intervals of Sec. 4.4. At the scale of the data the front-runners are all but indistinguishable point by point — and this is expected: the batteries' verdicts are sums over hundreds of small, correlated differences, each far below the scatter of the data, not any single visible feature.

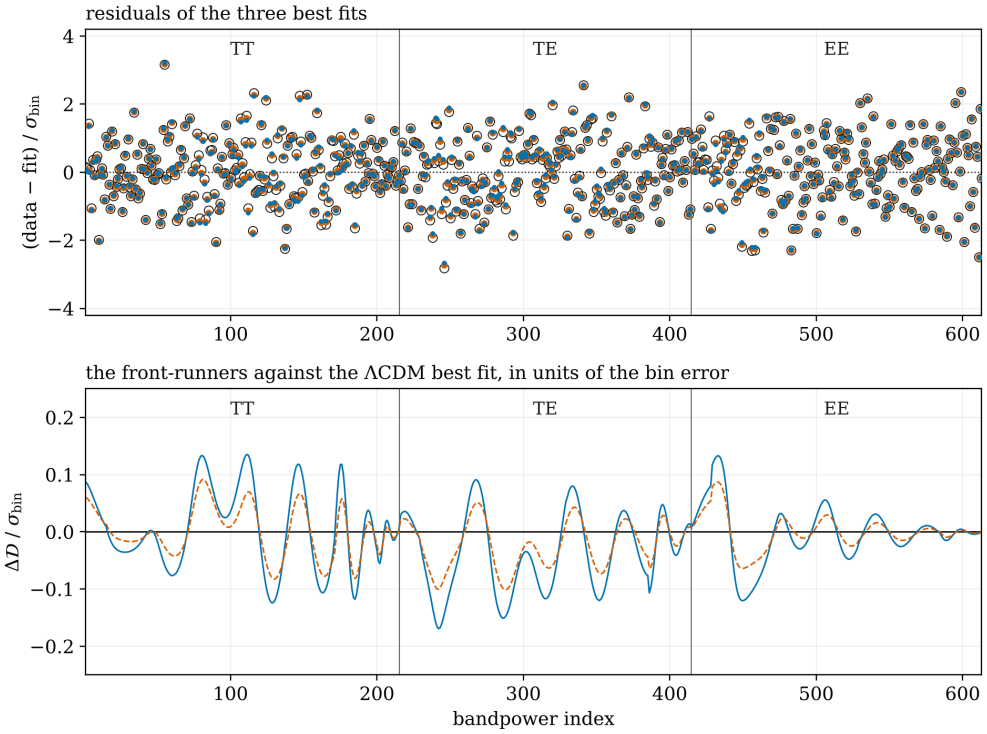

Figure 1C: The three-probe best fits against the CMB data of Fig. 2C. Top: the residuals of the three best fits, drawn as layered markers — ΛCDM a large open circle, CPL a filled orange disc inside it, Avrami a filled blue disc on top — so that coincident fits read as one concentric symbol and any divergence shows as an off-center ring. The three sets coincide to within the marker size, and their diagonal χ^2 are indistinguishable (552.9 / 551.7 / 551.7): point by point, the fits are the same fit. Bottom: the front-runners against the ΛCDM best fit, bin by bin (solid blue: Avrami − ΛCDM; dashed orange: CPL − ΛCDM) — each model's binned spectrum, with its own calibration $A^2_{\rm Planck}$ divided out, minus ΛCDM's, in units of the bin error. No single bin moves by more than 0.17σ (Avrami) or 0.10σ (CPL): the CMB leg of the three-probe verdict, $\Delta\chi^2_{\rm CMB}$(Avrami − ΛCDM) = −4.3, is a coherent sum over the acoustic-phase pattern visible here, not the work of any single multipole. The pattern reads directly: CPL tracks the ΛCDM baseline at half the Avrami amplitude — of the two front-runners it is the nearer to ΛCDM — yet it is Avrami that is the nearer to the data: $\chi^2_{\rm CMB}$ = 585.6 (ΛCDM), 582.1 (CPL), 581.3 (Avrami). Per spectrum, with each block taken against its own covariance (the cross-spectrum correlations are then dropped, so the blocks do not sum to the full χ^2): TT 207.2 / 204.7 / 204.0, TE 185.0 / 185.7 / 186.7, EE 198.0 / 197.1 / 196.5 (ΛCDM / CPL / Avrami) — the born form's gain lives in TT and EE, while TE mildly prefers ΛCDM.

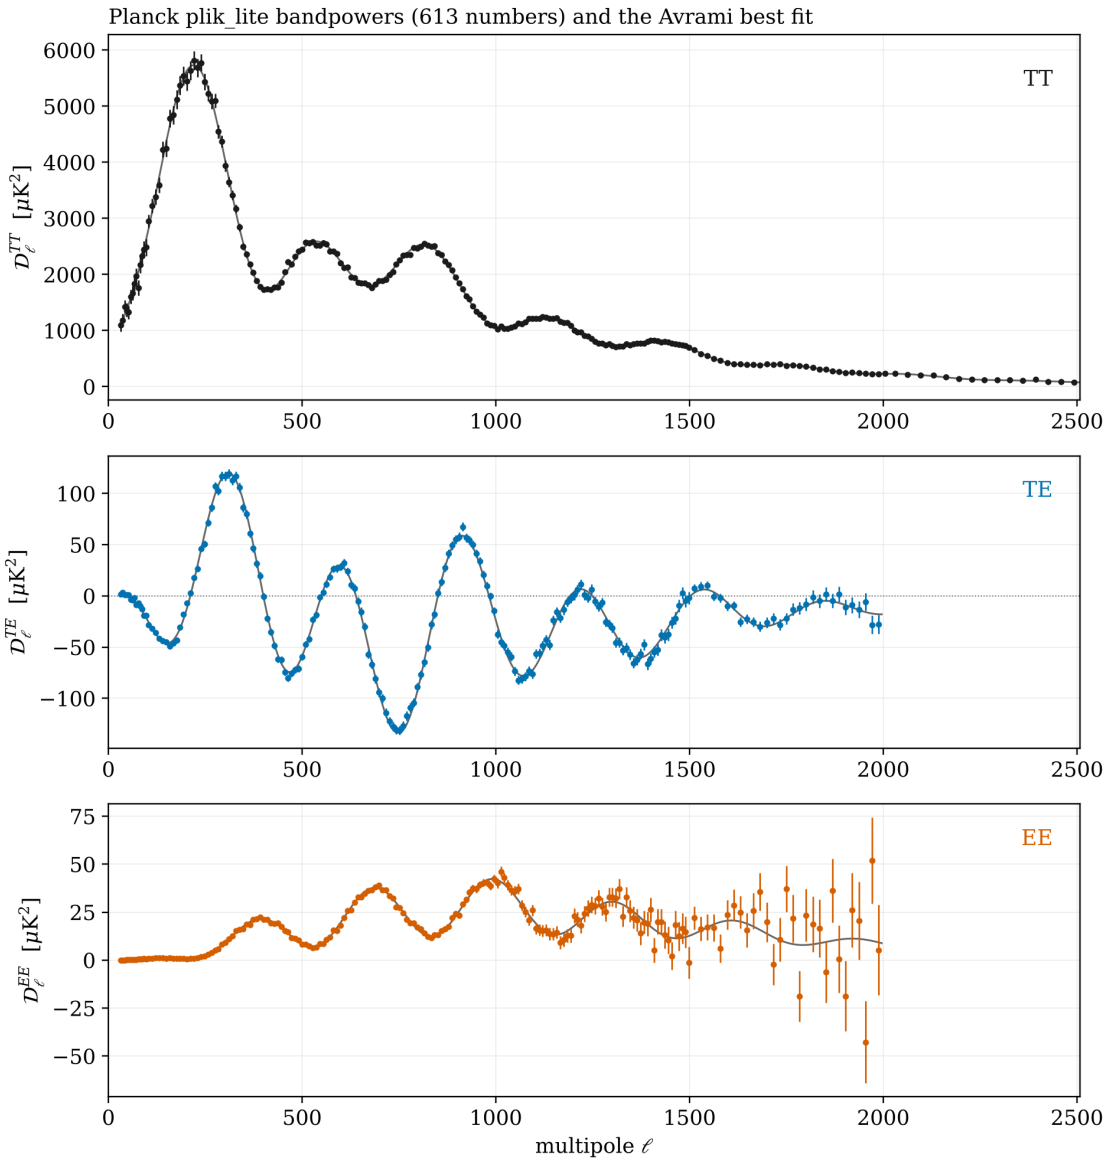

Figure 2C: The Planck `plik_lite` TTTEEE bandpowers (2018/PR3, v22): TT (215 bins, $\ell = 30$–2508), TE and EE (199 bins each, $\ell = 30$–1996) — the 613 CMB numbers entering every battery (Sec. 2). Error bars are the diagonal of the published bandpower covariance (the full covariance is used in the fits); the grey curve is the Avrami best fit of Table 3 — at this scale the three best fits are indistinguishable. Spectra from the released dump script; rebinned through the released `plik_lite` reader, they reproduce the verdict χ^2 to 585.57 / 582.11 / 581.34 (exact / exact / 0.07 from the verdict — the residual of re-tabulating the solved $w(a)$ for the dump, not a fit difference).

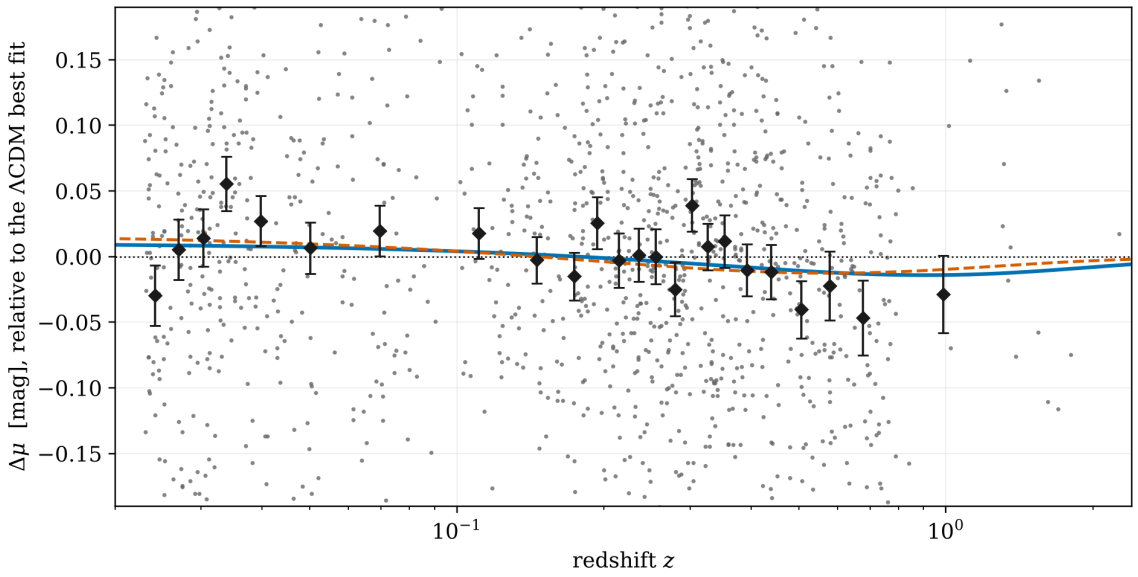

Figure 3C: The Pantheon+ Hubble diagram in residuals relative to the ΛCDM best fit of Table 3, under the battery conventions: the cosmology cut $z_{\rm HD} > 0.023$ (1371 supernovae), the full STAT+SYS covariance, the absolute magnitude profiled analytically per model. Grey points: individual supernovae (diagonal errors). Black diamonds with error bars: 24 equal-count bins and their uncertainties, weighted by the full covariance restricted to each bin (cross-bin correlations not shown); the dotted horizontal line is the ΛCDM baseline (zero residual). The bins enter no fit — every χ^2 in this paper is computed from the 1371 individual supernovae with the full covariance; the diamonds are a guide for the eye in the strict sense, a lossless-weighted compression of the data drawn only to make the ensemble mean visible against the scatter. Solid blue curve: the Avrami best fit of Table 3 minus the ΛCDM one — the difference of the two models' distance-modulus predictions, each with its own profiled M, i.e. how much brighter or fainter a supernova at that redshift should appear if the born form, not ΛCDM, is the truth; dashed orange: the same for the CPL tangent. On this battery $\chi^2_{\rm SN} = 1211.8$ (ΛCDM), 1210.9 (CPL), 1210.0 (Avrami): the entire supernova leg of the three-probe verdict, $\Delta\chi^2_{\rm SN}$(Avrami − ΛCDM) = −1.8 (−0.9 for the CPL tangent), lives in the $\sim$ 0.01–0.02 mag corridor between the curves.

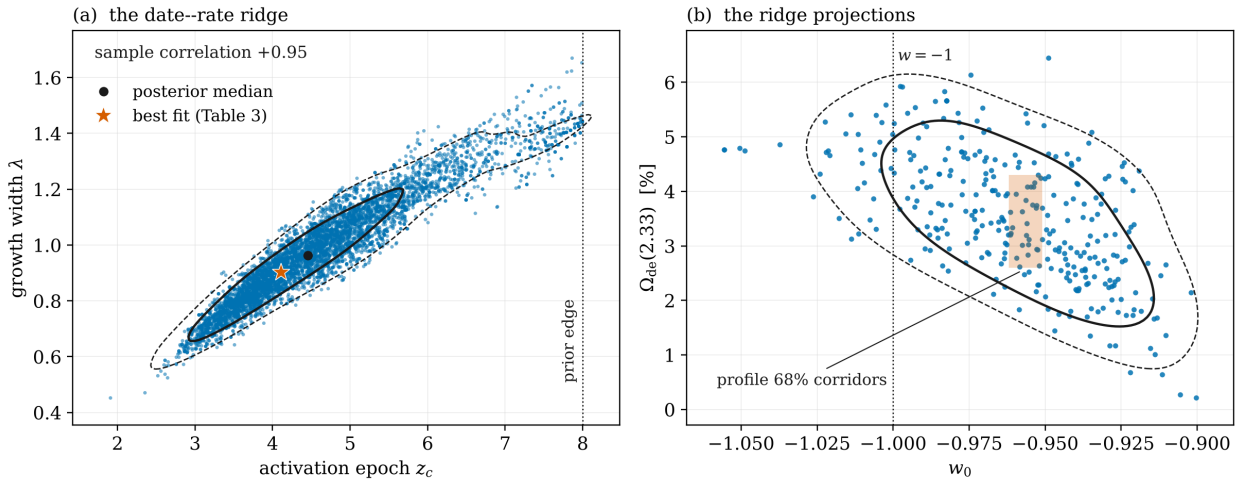

Figure 4C: The posterior of the leading Avrami form (blue points: thinned `emcee` samples; solid/dashed black: 68/95% KDE contours). (a) The date–rate ridge: the (z_c, λ) samples (annotated sample correlation +0.95) with the posterior median (black dot), the Table 3 best fit (orange star), and the prior edge at $z_c = 8$ (dotted) that the upper arm abuts; the activation epoch marginalizes to $z_c = 4.45$ with 68% interval $[3.61, 5.61]$, one-sided exactly as the profile found. (b) The ridge projections $(w_0, \Omega_{de}(2.33))$, re-solved on a posterior subsample of 320; the orange box marks the profile ridge corridors ($w_0 = -0.951$ to -0.962, $\Omega_{de}(2.33) = 2.6$–4.3%), and the dotted line the phantom divide $w = -1$. Marginalization widens them to the forecast-grade $\Omega_{de}(2.33) = 3.5\%$ with 68% interval $[2.4, 4.7]\%$, and to the $w_0 = -0.96 \pm 0.03$ quoted in the text (Fig. 2 marks the best-fit-anchored $-0.95^{+0.02}_{-0.04}$).

References cited in the appendices

[9] G. W. Gibbons, S. W. Hawking, "Cosmological event horizons, thermodynamics, and particle creation," Phys. Rev. D **15**, 2738 (1977).
[11] M. Li, "A model of holographic dark energy," Phys. Lett. B **603**, 1 (2004).
[45] S. M. Carroll, A. Chatwin-Davies, "Cosmic equilibration: a holographic no-hair theorem from the generalized second law," Phys. Rev. D **97**, 046012 (2018).
[61] L. P. Gor'kov, G. M. Eliashberg, "Generalization of the Ginzburg–Landau equations for non-stationary problems in the case of alloys with paramagnetic impurities," Zh. Eksp. Teor. Fiz. **54**, 612 (1968) [Sov. Phys. JETP **27**, 328 (1968)].