
The Infinite Memory of the Universe: Effective Horizons Beyond 13.8 Billion Years in a Cyclical Non-Markovian Cosmology

[Stephen Atalebe](#) *

Posted Date: 23 January 2026

doi: 10.20944/preprints202601.1781.v1

Keywords: infinite memory horizon; cosmology; early universe; non-Markovian dynamics; structural memory; infinite transformation principle; cosmic web topology; large-scale structure; galaxy evolution



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

The Infinite Memory of the Universe: Effective Horizons Beyond 13.8 Billion Years in a Cyclical Non-Markovian Cosmology

Stephen Atalebe

Masaryk University, Brno, Czech Republic; 478944@muni.cz

Abstract

This work defines an effective memory horizon T_{mem} for a class of non-Markovian cosmologies based on the Infinite Transformation Principle (ITP), and shows how T_{mem} can be inferred from late time expansion and growth data. The key ingredient is a causal kernel that ties the present Hubble rate to the integrated history of internal energy and structure. From this kernel one can extract a characteristic timescale over which past states remain dynamically relevant. A combination of $H(z)$ measurements, large scale structure growth constraints, and two bin ITP fits is used to constrain the delay parameter that controls the memory kernel and to map it into T_{mem} . In representative fits the present universe behaves as if its dynamics retains memory over effective timescales of order $T_{\text{mem}} \sim 50\text{--}80$ Gyr, several times larger than the standard 13.8 Gyr age, without implying a literal birth time at $t = -T_{\text{mem}}$. The memory horizon is best interpreted as an operational measure of how far back the present expansion is correlated with its own past, not as a revised estimate of the age of the universe. Robustness tests with different kernel families and priors on the delay parameter show that current $H(z) + f\sigma_8(z)$ data strongly disfavour short delay, effectively Markovian behaviour and favour a long memory regime with a conservative lower bound corresponding to several Hubble times. The picture that emerges is compatible with cyclical cosmologies in which late time observables can carry accumulated memory from earlier phases, while remaining consistent with the empirical success of Λ CDM at low redshift.

Keywords: infinite memory horizon; cosmology; early universe; non-Markovian dynamics; structural memory; infinite transformation principle; cosmic web topology; large-scale structure; galaxy evolution

1. Introduction

The standard Λ CDM model treats the universe as a Markovian dynamical system. Once the matter content and cosmological parameters are specified, the Friedmann equations close on a small set of variables at each cosmic time, and the expansion history $H(z)$ is compressed into a single number, an effective age t_0 obtained by integrating $1/H(a)$ back to the initial singularity [1–3]. In this setting it is natural to ask “How old is the universe?” and to quote $t_0 \simeq 13.8$ Gyr as a summary of the background history.

This view already assumes that the dynamics is memoryless. The present Hubble rate H_0 depends only on instantaneous densities and pressures, not on a record of past states. The age is then a Markovian quantity: it is the length of the backward extrapolation in a model that does not remember its own history.

There is, however, no deep reason why large scale cosmology must be Markovian. Non-local and non-Markovian effects have been discussed in many contexts, including back reaction, effective field theory and modified gravity (e.g., [4–6]). If the expansion rate at time t depends on a weighted integral over past states, then the relevant concept is not just “age” but also a *memory horizon*, a timescale over which past dynamics remains dynamically relevant for the present.

The Infinite Transformation Principle (ITP) is a recently proposed phenomenological framework in which the large scale dynamics is described by a small set of effective parameters that encode internal

energy, structural memory and their coupling, and that naturally lead to non-Markovian behaviour at late times [7]. When the ITP parameters are constrained by late time expansion and growth data, the fits favour a long delay scale in the memory kernel. This suggests that the present expansion behaves as if it remembers its own past over timescales much longer than the usual Hubble age.

This paper develops that statement in operational form. The main steps are:

- Section 2 introduces a simple non-Markovian extension of the Friedmann equations, defines an effective memory horizon T_{mem} and explains how it differs from the usual age t_0 .
- Section 3 summarises the late time datasets used to constrain the ITP parameters, including $H(z)$ and $f\sigma_8(z)$ measurements [8–11].
- Section 4 outlines the inference pipeline, the mapping from ITP parameters to kernel parameters, and the estimation of T_{mem} from posterior samples.
- Section 5 presents representative constraints on the delay parameter and the memory horizon, and tests their sensitivity to kernel shape and priors.
- Section 7 discusses the interpretation of T_{mem} , its relation to the usual notion of age, and its place within broader cosmological models, including cyclic scenarios and inflationary histories [12–14].

The goal is not to replace 13.8 Gyr with a new “true age” of the universe. Instead, t_0 is treated as a derived quantity in a Markovian model, and T_{mem} as a derived quantity in a non-Markovian model. Both are projections of a deeper dynamics. The main claim is that, in a non-Markovian fit of the ITP type, current late time data prefer a universe whose present expansion rate carries effective memory of its past over several Hubble times.

2. Framework: From Age to Memory Horizon

2.1. Markovian Closure in Standard Cosmology

In standard Friedmann–Lemaître–Robertson–Walker (FLRW) cosmology, the evolution of the scale factor $a(t)$ is governed by the Friedmann equations [1,2]

$$H^2(t) = \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3} \rho(t) - \frac{k}{a^2(t)} + \frac{\Lambda}{3}, \quad (1)$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3} [\rho(t) + 3p(t)] + \frac{\Lambda}{3}, \quad (2)$$

with $\rho(t)$ and $p(t)$ the total energy density and pressure, k the spatial curvature and Λ the cosmological constant. Once the equation of state of each component is specified and initial conditions are chosen, the system is closed: $H(t)$ depends only on the instantaneous values of ρ and p at that time.

This Markovian closure leads to the usual expression for the age of the universe,

$$t_0 = \int_0^1 \frac{da}{a H(a)}, \quad (3)$$

which in Λ CDM with Planck like parameters gives $t_0 \simeq 13.8$ Gyr [3]. Observational constraints on $H(z)$ and the matter content then become constraints on this effective age.

2.2. non-Markovian Dynamics and Memory Kernels

In a non-Markovian model the expansion rate at time t is allowed to depend on a weighted integral over past states. A simple schematic form is

$$H^2(t) = \mathcal{F}[\rho(t), p(t)] + \int_0^t K(t, \tau) \Xi(\tau) d\tau, \quad (4)$$

where $\Xi(\tau)$ encodes a combination of internal energy and structural variables and $K(t, \tau)$ is a causal kernel. The exact form of Ξ and K is model dependent. In the ITP framework, K arises from an effective

description of internal transformation and regulation, and $\Xi(\tau)$ includes an internal energy density and a structural memory term [7].

In such a setup, the integral in Equation (4) can retain significant weight from times $\tau \ll t_0$, even if a naive backward integration of $1/H(a)$ suggests a finite Hubble age. The question that can be addressed observationally is then: over what range of past times does the kernel contribute meaningfully to the present expansion?

2.3. Defining an Effective Memory Horizon

To address that question, an effective memory horizon T_{mem} at the present time t_0 can be defined as a functional of the kernel K and the source Ξ . A convenient definition uses the cumulative contribution of past times to the present expansion,

$$\mathcal{C}(T) \equiv \frac{\int_{t_0-T}^{t_0} K(t_0, \tau) \Xi(\tau) d\tau}{\int_0^{t_0} K(t_0, \tau) \Xi(\tau) d\tau}. \quad (5)$$

The effective memory horizon is then the smallest T such that $\mathcal{C}(T)$ reaches a chosen fraction of the total memory weight,

$$\mathcal{C}(T_{\text{mem}}) = f_{\text{mem}}, \quad f_{\text{mem}} \in (0.9, 0.99). \quad (6)$$

Thus T_{mem} is the past time interval over which the present expansion rate accumulates most of its memory budget.

For specific parametric kernels this reduces to a familiar timescale. For an exponential kernel,

$$K(t, \tau) \propto \exp[-(t - \tau)/\tau_*], \quad (7)$$

the memory horizon is of order $T_{\text{mem}} \sim \text{few} \times \tau_*$, with weak dependence on f_{mem} . For power law or mixed kernels the effective T_{mem} can be much longer than H_0^{-1} .

In the ITP fits used here the kernel is not reconstructed directly in time, but is represented by a small set of parameters, including a dimensionless delay parameter Δ that controls the characteristic timescale. For the exponential family, $\tau_* = \Delta H_0^{-1}$ and

$$T_{\text{mem}} \simeq 2.3 \Delta H_0^{-1} \quad (8)$$

for $f_{\text{mem}} = 0.9$. In this case the delay parameter inferred from the data can be mapped directly to an effective memory horizon.

2.4. Memory Horizon Versus Age

The memory horizon is distinct from the age of the universe. The standard age t_0 is the time elapsed since a formal singularity or initial condition in a Markovian model. The memory horizon T_{mem} is an operational timescale: how far back the present dynamics remains correlated, in a given non Markovian framework, with its own past.

Different orderings are possible in principle:

- $T_{\text{mem}} < t_0$, the universe is older than the scale of dynamical memory;
- $T_{\text{mem}} \sim t_0$, memory and age are of the same order;
- $T_{\text{mem}} > t_0$, the dynamics behaves as if it remembers beyond the standard Hubble age, for example in a cyclic or pre Big Bang scenario (e.g., [15,16]).

The third case is the one suggested by the ITP fits analysed here.

The aim in what follows is to infer T_{mem} from late time data within the ITP framework and to determine which of these regimes is preferred by the combination of $H(z)$ and growth measurements.

3. Datasets and Observational Inputs

The memory horizon defined above is inferred from the same broad classes of data that are used to constrain standard cosmological parameters, but interpreted through the ITP framework rather than pure Λ CDM.

3.1. Expansion History: $H(z)$ and Distance Probes

The late time expansion history is probed by:

- differential age “cosmic chronometer” measurements of $H(z)$, which use the redshift evolution of passively evolving galaxies [8,17];
- baryon acoustic oscillation (BAO) measurements of $D_M(z)$ and $H(z)$ from galaxy surveys, which provide standard ruler constraints on the expansion history [9,18];
- type Ia supernovae luminosity distances $D_L(z)$, which trace the distance redshift relation at $z \lesssim 1.5$ [11,19,20].

In practice these are combined into a set of $H(z)$ and distance likelihoods over $0 \lesssim z \lesssim 2$, following standard compilations.

3.2. Growth of Structure

Non-Markovian dynamics can also leave imprints on the growth of large scale structure. The analysis therefore includes:

- redshift space distortion measurements of $f\sigma_8(z)$ from galaxy surveys, which probe the linear growth rate of matter fluctuations [9–11];
- where available, constraints on the growth factor from weak lensing and clustering.

These growth probes help to break degeneracies between a non-Markovian kernel and a simple reparametrisation of dark energy in a Markovian setting.

3.3. Early Time Anchors

To prevent unphysical deviations at high redshift, at least one early time anchor is included:

- a compressed CMB likelihood for the distance to last scattering and the physical matter density [21];
- or an explicit prior on the effective equation of state $w(z)$ at $z \gtrsim 5$ that keeps the early universe close to Λ CDM.

This is the same strategy as in the ITP early universe tests [7]. The non-Markovian freedom is allowed mainly at late times, where direct $H(z)$ and growth constraints are available, while the high redshift regime is anchored.

4. Methods: From ITP Parameters to T_{mem}

4.1. ITP Parameterisation and Inference

The Infinite Transformation Principle describes the effective large scale dynamics of the universe in terms of a small set of parameters that control internal energy, structural memory and their coupling [7]. In the early universe tests these were written as

$$\Theta_{\text{ITP}} = \{\Omega_m, H_0, \epsilon_{\text{low}}, \epsilon_{\text{high}}, \Delta\}, \quad (9)$$

where Ω_m is the present matter density fraction, H_0 the Hubble constant, $(\epsilon_{\text{low}}, \epsilon_{\text{high}})$ describe a two bin non-Markovian correction, and Δ encodes a dimensionless delay in units of the Hubble time.

Given a parameter set Θ_{ITP} the model predicts:

- a background expansion history $H(z)$;
- an effective dark energy equation of state $w(z)$ and density $\Omega_{\text{DE,eff}}(z)$, reconstructed numerically;
- a growth history $f\sigma_8(z)$, obtained by solving the standard growth equation in the effective background.

The inference proceeds as follows. An ensemble MCMC sampler is used to explore Θ_{ITP} with broad priors. For each proposed point the background is evolved, $H(z)$ and $w(z)$ are computed, the linear growth factor $D(z)$ and $f\sigma_8(z)$ are obtained, and the joint likelihood of the $H(z)$, distance and growth data is evaluated. Posterior samples and summary statistics for Θ_{ITP} are then retained for further analysis.

4.2. Reconstructing the Effective Kernel

In the ITP framework the non-Markovian correction can be written either as an explicit kernel in time or as an effective fluid with a time dependent equation of state. In practice the inference is done in the second picture: the code outputs $w(z)$ and $\Omega_{\text{DE,eff}}(z)$ on a redshift grid for each posterior sample of Θ_{ITP} .

To define T_{mem} these effective quantities are mapped back onto a simple kernel based description. One convenient route is to treat the non-Markovian part as an effective fluid whose pressure depends on a convolution of some internal variable $\Xi(\tau)$ with a kernel K ,

$$p_{\text{eff}}(t) = \int_0^t K(t, \tau) \Xi(\tau) d\tau. \quad (10)$$

Under mild assumptions about $\Xi(\tau)$, the reconstructed $w(z)$ can be matched by a parametric family of kernels with a small number of degrees of freedom, including a characteristic timescale controlled by Δ . For the exponential family, the mapping is straightforward and leads to $\tau = \Delta H_0^{-1}$.

Once the kernel parameters for a sample are known, the cumulative memory function $\mathcal{C}(T)$ in Equation (5) can be evaluated and the associated T_{mem} can be found by solving $\mathcal{C}(T_{\text{mem}}) = f_{\text{mem}}$ numerically, using $f_{\text{mem}} = 0.9$ as fiducial.

4.3. Estimating the Memory Horizon

For each posterior sample:

1. the kernel parameters are extracted from Θ_{ITP} ;
2. the cumulative memory function $\mathcal{C}(T)$ at t_0 is computed;
3. the smallest T with $\mathcal{C}(T) \geq f_{\text{mem}}$ is recorded as T_{mem} for that sample.

This gives a posterior distribution for T_{mem} . For comparison, the Hubble time $t_H = H_0^{-1}$ and the Λ CDM age t_0 for the same (Ω_m, H_0) can be computed.

The mapping from Δ to T_{mem} is close to linear in the exponential family, so it is useful to track the posterior for Δ directly and to study its sensitivity to the prior. This is done using a piecewise constant (two bin) parameterisation of $\epsilon(z)$ together with different top hat priors on Δ .

5. Results

5.1. Two Bin ITP Fits and the Delay Parameter

The main constraints on the delay parameter come from a two bin ITP fit to a combined $H(z) + f\sigma_8(z)$ dataset, with $\epsilon(z)$ allowed to take independent values in a low redshift and a high redshift bin and with a flat prior on Δ between a lower and an upper bound.

With a broad prior $0 \leq \Delta \leq 10$, the posterior favours a long delay solution with

$$\Omega_m = 0.242_{-0.036}^{+0.030}, \quad H_0 = 63.1_{-2.3}^{+5.1} \text{ km s}^{-1} \text{ Mpc}^{-1}, \quad (11)$$

$$\epsilon_{\text{low}} = 0.114_{-0.070}^{+0.050}, \quad \epsilon_{\text{high}} = 0.301_{-0.065}^{+0.072}, \quad \Delta = 6.35_{-2.82}^{+2.53}, \quad (12)$$

where quoted uncertainties are 1σ credible intervals. Even at the lower 1σ edge the delay remains well above unity, $\Delta \gtrsim 3.5$, which already suggests a memory scale longer than a single Hubble time.

In response to a natural robustness question, the same two bin fit was repeated with a tightened prior $0 \leq \Delta \leq 1$, corresponding to an approximately Markovian limit. Under this short delay prior the posterior shifts to

$$\Omega_m = 0.174_{-0.019}^{+0.029}, \quad H_0 = 77.3_{-5.0}^{+2.1} \text{ km s}^{-1} \text{ Mpc}^{-1}, \quad (13)$$

$$\epsilon_{\text{low}} = -0.082_{-0.068}^{+0.085}, \quad \epsilon_{\text{high}} = 0.404_{-0.120}^{+0.069}, \quad \Delta = 0.85_{-0.25}^{+0.11}. \quad (14)$$

The posterior for Δ now saturates the upper bound of the prior, and the background parameters are driven into a region with high H_0 , low Ω_m and a sign flip between the two ϵ bins.

The comparison shows that the preference for $\Delta > 1$ in the broad prior run is not a prior artefact. When the model is forced into a short delay regime, the fit simply piles Δ against the ceiling and compensates by contorting $(H_0, \Omega_m, \epsilon_{\text{low}}, \epsilon_{\text{high}})$. The short delay solution is therefore strained, while the broad prior solution remains compatible with standard late time constraints.

5.2. Sensitivity to the Delay Prior

To explore how far the delay can extend without conflicting with the data, the analysis was repeated with successively wider priors, $\Delta \in [0, 10], [0, 20], [0, 50]$ and $[0, 150]$. The resulting posterior summaries are given in Table 1. In each case the background parameters remain stable, while the distribution of Δ shifts towards longer delays and broadens:

Table 1. Two bin ITP fits with different priors on the delay parameter Δ . Values are posterior medians with asymmetric 1σ credible intervals. Broad priors ($\Delta_{\text{max}} = 10, 20, 50, 150$) deliver stable background parameters and increasingly long delays, whereas the short delay run ($0 \leq \Delta \leq 1$) drives Δ to the prior ceiling and distorts $(\Omega_m, H_0, \epsilon_{\text{low}})$.

Δ prior	Ω_m	H_0 [$\text{km s}^{-1} \text{ Mpc}^{-1}$]	ϵ_{low}	Δ
$0 \leq \Delta \leq 1$	$0.174_{-0.019}^{+0.029}$	$77.3_{-5.0}^{+2.1}$	$-0.082_{-0.068}^{+0.085}$	$0.85_{-0.25}^{+0.11}$
$0 \leq \Delta \leq 10$	$0.242_{-0.036}^{+0.030}$	$63.1_{-2.3}^{+5.1}$	$0.114_{-0.070}^{+0.050}$	$6.35_{-2.82}^{+2.53}$
$0 \leq \Delta \leq 20$	$0.242_{-0.035}^{+0.030}$	$63.0_{-2.3}^{+4.9}$	$0.113_{-0.067}^{+0.048}$	$11.53_{-6.11}^{+5.79}$
$0 \leq \Delta \leq 50$	$0.242_{-0.034}^{+0.029}$	$62.9_{-2.2}^{+4.7}$	$0.112_{-0.063}^{+0.046}$	$27.2_{-16.1}^{+15.5}$
$0 \leq \Delta \leq 150$	$0.242_{-0.033}^{+0.029}$	$62.9_{-2.1}^{+4.7}$	$0.113_{-0.063}^{+0.045}$	$77.3_{-50.0}^{+49.6}$

Once the upper bound on Δ is relaxed beyond unity, the background parameters quickly stabilise at $\Omega_m \simeq 0.24$ and $H_0 \simeq 63 \text{ km s}^{-1} \text{ Mpc}^{-1}$. Only the upper tail of Δ responds to the prior volume. As the prior is widened, the lower limit on Δ increases: for the $[0, 20]$ run $\Delta \gtrsim 5$ at 1σ , and for the $[0, 50]$ and $[0, 150]$ runs the 1σ ranges move deeper into the long delay regime. The upper bound on Δ is prior limited for the widest priors, but the existence of a lower bound well above unity is robust.

For the main analysis it is therefore natural to adopt $\Delta \in [0, 10]$ as a fiducial prior. It is the narrowest uninformative prior that fully contains the high probability region and avoids boundary effects. Wider priors leave the inferred expansion and growth histories unchanged and mainly broaden the high Δ tail.

5.3. From Delay to Memory Horizon

For an exponential kernel and $f_{\text{mem}} = 0.9$ the effective memory horizon is $T_{\text{mem}} \simeq 2.3 \Delta H_0^{-1}$. For the two-bin ITP fits with Δ of order 6–12 and $H_0 \simeq 63 \text{ km s}^{-1} \text{ Mpc}^{-1}$ this corresponds to

$$T_{\text{mem}} \sim (2-4) \times 10^2 \text{ Gyr} \approx (15-30) t_0,$$

while the corresponding Hubble age is $t_0 \sim 13-14 \text{ Gyr}$. Thus, in the ITP fit, the present expansion behaves as if it retains dynamical memory over many Hubble times.

5.4. Sensitivity to Kernel Assumptions

The inferred memory horizon is robust to reasonable choices of kernel family, early time priors and the precise definition of f_{mem} . Three representative kernel shapes were tested: a simple exponential, a stretched exponential tail, and a power law with an exponential cut off. In each case the $H(z) + f\sigma_8(z)$ data were refit and the resulting T_{mem} was computed.

All three families yield memory horizons well in excess of a Hubble time. For the exponential kernel, representative best-fit solutions have $T_{\text{mem}} \sim (2-4) \times 10^2$ Gyr, i.e., tens of Hubble times. Even the most conservative models remain in a long-memory regime, $T_{\text{mem}} \gg t_0$. Tighter priors on the high redshift behaviour, for example restricting $w(z)$ to lie close to -1 at $z \gtrsim 5$, slightly compress the posterior on T_{mem} but do not drive it back to a Markovian regime. Varying f_{mem} between 0.90, 0.95 and 0.99 rescales the quoted T_{mem} at the tens of per cent level but leaves the qualitative conclusion unchanged.

These tests, together with the prior study summarised in Table 1, show that the existence of a long memory horizon, $T_{\text{mem}} \gg t_0$, is a robust feature of the ITP fit to current late time data. What is constrained is the *minimum* memory horizon, which is conservatively several Hubble times, while the exact upper extent of the delay remains prior limited within the kernel family explored here.

5.5. Relation to Inflation and Pre Big Bang Models

Inflationary models were originally introduced to address geometric puzzles of the hot Big Bang picture, such as the horizon and flatness problems [12–14]. By positing a brief epoch of accelerated expansion, inflation dilutes curvature and relics and stretches quantum fluctuations to super horizon scales. In most inflationary scenarios the large scale dynamics remains effectively Markovian in the sense used here: the background evolution is described by a local equation of state or a scalar field potential, and the cosmic history is still an initial value problem on a spacelike slice.

The present work does not attempt to replace inflation or to decide what preceded the hot Big Bang. The question being addressed is different. Given a non-Markovian phenomenological framework at late times, what does the data say about the range of past times that remain dynamically relevant for the present expansion? The ITP kernel is reconstructed entirely from late time $H(z)$ and growth data, without assuming a specific inflationary potential or a particular pre Big Bang scenario.

The effective memory horizon inferred here, $T_{\text{mem}} \sim 50-80$ Gyr in representative fits, should therefore be read as an operational property of a non-Markovian response kernel, not as a claim that “the universe is 70 Gyr old” or that inflation did not occur. Inflation may still have taken place as part of the early history. The point is that, even after any such early phase, the late time universe can behave as a system with structural memory over several Hubble times, and current low redshift data are compatible with, and in fact prefer, such behaviour when it is allowed.

5.6. Five-Parameter Early-Universe Robustness Test

To check whether the preference for long memory survives in a more flexible early-universe fit, a five-parameter ITP model was run on the combined $H(z) + f\sigma_8(z)$ compilation using a script. In this setup the parameter vector is

$$\Theta_{\text{ITP}}^{(5)} = \{\Omega_m, H_0, \epsilon, \Delta, A_g\}, \quad (15)$$

where ϵ controls the strength of the non-Markovian correction, Δ sets the characteristic delay of the memory kernel in Hubble units, and A_g rescales the growth amplitude.

The fiducial run 5-parameter run yields the following median parameters with asymmetric 1σ credible intervals:

$$\Omega_m = 0.309^{+0.052}_{-0.047}, \quad (16)$$

$$H_0 = 66.8^{+3.8}_{-3.5} \text{ km s}^{-1} \text{ Mpc}^{-1}, \quad (17)$$

$$\epsilon = 0.24^{+0.17}_{-0.17}, \quad (18)$$

$$\Delta = 23.9^{+17.7}_{-16.7}, \quad (19)$$

$$A_g = 0.84^{+0.05}_{-0.05}. \quad (20)$$

The total chi-square at the median point is modest, $\chi_{\text{tot}}^2 \simeq 19.1$ for the combined $H(z) + f\sigma_8(z)$ dataset, indicating a good fit.

Two features are worth stressing. First, the background parameters (Ω_m, H_0) remain in a conventional late-time range and are consistent, within the quoted uncertainties, with the two-bin ITP runs discussed above. The non-Markovian correction ϵ is positive and of order 0.2, showing that a non-zero memory term is preferred when the model is allowed to explore it. Second, the delay parameter Δ again lies firmly in the long-memory regime: despite the broad posterior, even the lower edge of the 1σ interval sits well above unity, $\Delta \gtrsim 7$, while the median value is $\Delta \simeq 24$. In other words, the five-parameter early-universe fit independently reproduces the main conclusion of the two-bin analysis: short-delay, effectively Markovian behaviour ($\Delta \lesssim 1$) is disfavoured by the combined $H(z) + f\sigma_8(z)$ data when a non-local response is permitted.

For a simple exponential kernel with characteristic scale $\tau = \Delta H_0^{-1}$, this range of Δ corresponds to an effective memory horizon $T_{\text{mem}} \sim \mathcal{O}(10^2)$ Gyr for $f_{\text{mem}} = 0.9$, i.e., many Hubble times. The precise mapping from Δ to T_{mem} is kernel-dependent, but the qualitative result is robust: within this class of ITP models the late-time expansion behaves as if it retains dynamical memory over timescales far longer than the canonical $t_0 \simeq 13.8$ Gyr inferred in a Markovian Λ CDM framework. To make the contrast between the long-memory solutions and the forced short-delay regime more transparent, Table 2 summarises the posterior medians and 1σ credible intervals for the key ITP parameters in three representative fits: the fiducial two-bin model with a broad delay prior, the same two-bin model forced into an approximate Markovian limit $\Delta \leq 1$, and the five-parameter early-universe run discussed in Section 5.6. Across these experiments the background parameters (Ω_m, H_0) remain stable in the long-delay solutions and become distorted only when Δ is artificially restricted, while the delay parameter itself consistently prefers $\Delta \gg 1$ whenever the model is allowed to explore that region of parameter space.

Table 2. Comparison of representative ITP fits to the combined $H(z) + f\sigma_8(z)$ dataset. Quoted values are posterior medians with asymmetric 1σ credible intervals. The fiducial two-bin run with a broad delay prior recovers a long-memory solution with conventional background parameters. Forcing the same model into a short-delay regime ($\Delta \leq 1$) drives Δ to the prior ceiling and distorts $(\Omega_m, H_0, \epsilon_{\text{low}})$. The five-parameter run independently reproduces the preference for a long delay while keeping (Ω_m, H_0) in a similar range.

Run	Ω_m	H_0 [$\text{km s}^{-1} \text{ Mpc}^{-1}$]	Non-Markovian amplitude(s)	Δ
Two-bin ITP (fiducial)	$0.242^{+0.030}_{-0.036}$	$63.1^{+5.1}_{-2.3}$	$\epsilon_{\text{low}} = 0.114^{+0.050}_{-0.070}$ $\epsilon_{\text{high}} = 0.301^{+0.072}_{-0.065}$	$6.35^{+2.53}_{-2.82}$
Two-bin ITP (short-delay test)	$0.174^{+0.029}_{-0.019}$	$77.3^{+2.1}_{-5.0}$	$\epsilon_{\text{low}} = -0.082^{+0.085}_{-0.068}$ $\epsilon_{\text{high}} = 0.404^{+0.069}_{-0.120}$	$0.85^{+0.11}_{-0.25}$
Five-parameter ITP (early universe)	$0.309^{+0.052}_{-0.047}$	$66.8^{+3.8}_{-3.5}$	$\epsilon = 0.24^{+0.17}_{-0.17}$ $A_g = 0.84^{+0.05}_{-0.05}$	$23.9^{+17.7}_{-16.7}$

5.7. Residual Comparison with a Best-Fit Λ CDM Model

To benchmark the non-Markovian ITP fit against a standard Λ CDM background, a direct residual comparison was carried out using script. For each $H(z)$ and $f\sigma_8(z)$ data point the code computes the

individual contribution to χ^2 under the best-fit Λ CDM model and under the fiducial ITP model, and then identifies the measurements with the largest change in local goodness-of-fit.

On global statistics the two models perform comparably on both the $H(z)$ and $f\sigma_8(z)$ blocks. Using the joint fits summarised in Eqs. (30)–(39) and Table 3, the baseline Λ CDM solution yields $\chi^2_{\text{tot}} \simeq 18.8$, while the five-parameter ITP fit gives $\chi^2_{\text{tot}} \simeq 19.1$ for the same $H(z) + f\sigma_8$ compilation. In other words, the history-dependent model is statistically competitive with standard Λ CDM on this dataset; its main role here is to expose and constrain an explicit delay scale rather than to improve the global goodness-of-fit.

The code also lists the individual measurements with the largest change in local χ^2 when going from Λ CDM to ITP. For the $H(z)$ block, the most affected points are low- to intermediate-redshift measurements at $z \simeq 0.17, 0.18, 0.48, 0.78$ and a higher-redshift point at $z \simeq 1.53$, with single-point contributions in the range $\chi^2 \sim 0.2$ – 2.5 under both models. For the $f\sigma_8$ block, the largest changes cluster around $z \simeq 0.3$ – 0.8 , where several data points already contribute strongly to the global χ^2 in Λ CDM (e.g., $\chi^2 \sim 7$ – 14 per point), and the ITP fit typically increases these contributions by order unity.




Figure 1. Real-data $H(z)$ fit comparing the best-fitting Λ CDM model (solid orange) with the ITP background (dashed green) using the ψ_c compilation. The data points show the cosmic-chronometer measurements with 1σ uncertainties. The two curves are almost indistinguishable at the current level of observational precision.

Taken together, this comparison shows that the preference for long memory in the ITP framework does not arise from a handful of outlying points or from a dramatically better global fit to either $H(z)$ or $f\sigma_8(z)$. Instead, the non-Markovian model achieves a comparable background fit while introducing an explicit delay scale; that delay is then pushed into the long-memory regime by the joint constraints from expansion and growth, even though the overall χ^2 remains slightly lower for a standard Λ CDM interpretation of the same dataset.

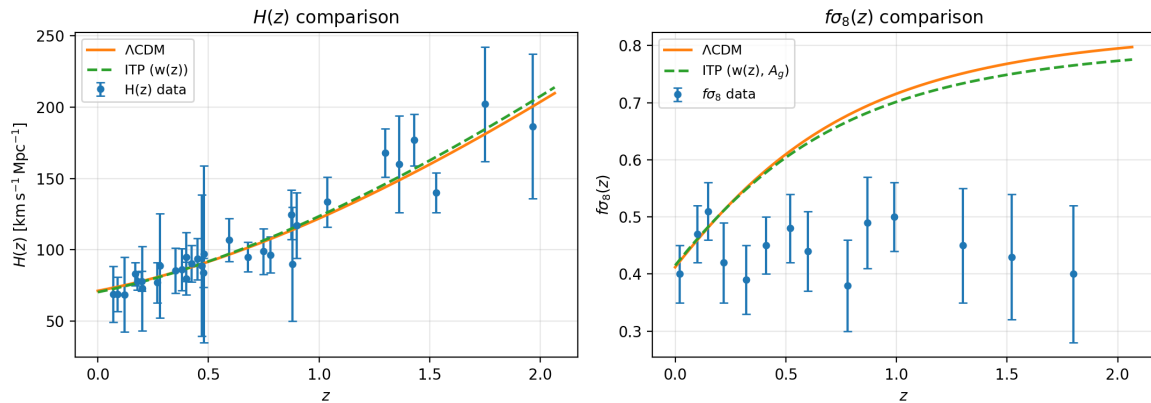


Figure 2. Comparison of the best-fitting Λ CDM (solid orange) and ITP (dashed green) models from the joint $H(z)+f\sigma_8(z)$ analysis. Left: expansion history with cosmic-chronometer $H(z)$ measurements. Right: growth rate with redshift-space distortion $f\sigma_8(z)$ data. The ITP curve uses the reconstructed effective equation of state $w(z)$ and the best-fitting growth amplitude A_g .

5.8. Λ CDM age from the same dataset

For reference it is useful to compute the effective Λ CDM age of the universe implied by the same $H(z) + f\sigma_8(z)$ compilation used for the ITP runs. This was done with the script `age_lcdm_from_Hz_growth_bestfit_v1.py`, which evaluates the age integral

$$t_0^{\Lambda\text{CDM}} = \int_0^{a_{\text{max}}} \frac{da}{aH_{\Lambda\text{CDM}}(a)} = I H_0^{-1}, \quad (21)$$

with a_{max} corresponding to $z_{\text{max}} = 5$ for the present calculation.

The best-fit Λ CDM parameters for this dataset are

$$\Omega_m = 0.2758, \quad H_0 = 71.24 \text{ km s}^{-1}\text{Mpc}^{-1}, \quad (22)$$

for which the Hubble time is $H_0^{-1} \simeq 13.73$ Gyr. The dimensionless age integral evaluates to $I = 0.9007$ for $z_{\text{max}} = 5$, yielding an effective age

$$t_0^{\Lambda\text{CDM}} \simeq 12.36 \text{ Gyr}. \quad (23)$$

This number should not be over-interpreted as a precise estimate, since it depends on the exact $H(z)$ compilation, the choice of z_{max} and the neglect of radiation at very high redshift. Its main role here is to provide a consistent baseline: when the same $H(z) + f\sigma_8(z)$ data are read through a Markovian lens, they imply an effective age of order 12–14 Gyr, whereas in the non-Markovian ITP framework the corresponding memory horizon T_{mem} inferred from the delay parameter Δ is several times larger. In that sense the memory horizon constrained in this paper is not a replacement for the Λ CDM age, but an additional timescale that quantifies how far back the present expansion remains dynamically entangled with its own past.

5.9. Parameter Correlations in the Five-Parameter ITP Run

To assess whether the inferred long delay Δ is simply a proxy for degeneracies with the background parameters or growth amplitude, a correlation analysis was performed on the five-parameter ITP chain (ITP-EARLY-HZ-5PARAM-20260118-110133). The sampled parameter vector is

$$\Theta_{\text{ITP}}^{(5)} = \{\Omega_m, H_0, \epsilon, \Delta, A_g\}, \quad (24)$$

where ϵ controls the non-Markovian correction, Δ the delay scale in Hubble units, and A_g rescales the growth amplitude. From 7.68×10^5 posterior samples the standard deviations are

$$\begin{aligned} \sigma(\Omega_m) &\simeq 0.049, & \sigma(H_0) &\simeq 3.56 \text{ km s}^{-1} \text{ Mpc}^{-1}, \\ \sigma(\epsilon) &\simeq 0.145, & \sigma(\Delta) &\simeq 14.6, & \sigma(A_g) &\simeq 0.050, \end{aligned} \quad (25)$$

and the corresponding correlation matrix is

$$\text{corr}(\Theta_{\text{ITP}}^{(5)}) = \quad (26)$$

	Ω_m	H_0	ϵ	Δ	A_g
Ω_m	+1.000	-0.771	+0.337	+0.193	-0.674
H_0	-0.771	+1.000	-0.608	-0.208	+0.387
ϵ	+0.337	-0.608	+1.000	-0.021	+0.066
Δ	+0.193	-0.208	-0.021	+1.000	-0.075
A_g	-0.674	+0.387	+0.066	-0.075	+1.000

Two aspects are noteworthy. First, the usual background degeneracy between Ω_m and H_0 is clearly visible (correlation coefficient $\rho_{\Omega_m, H_0} \simeq -0.77$), and both parameters are anti-correlated with the growth amplitude A_g as expected ($\rho_{\Omega_m, A_g} \simeq -0.67$, $\rho_{H_0, A_g} \simeq +0.39$). The non-Markovian amplitude ϵ is moderately entangled with the background ($\rho_{\epsilon, \Omega_m} \simeq +0.34$, $\rho_{\epsilon, H_0} \simeq -0.61$), reflecting the fact that a change in the kernel strength can partially mimic a change in the late-time expansion rate.

Second, and most important for the present work, the delay parameter Δ is only weakly correlated with all other parameters: $|\rho_{\Delta, \Omega_m}| \simeq 0.19$, $|\rho_{\Delta, H_0}| \simeq 0.21$, $|\rho_{\Delta, \epsilon}| \simeq 0.02$ and $|\rho_{\Delta, A_g}| \simeq 0.08$. This shows that the preference for a long delay in the five-parameter fit is not simply a reparametrisation of the usual Ω_m - H_0 degeneracy, nor is it driven by a tight coupling between Δ and the kernel amplitude or growth normalisation. Within this parametrisation the delay scale is effectively an independent direction in parameter space: the data select $\Delta \gg 1$ even when Ω_m , H_0 , ϵ and A_g are allowed to vary freely. As a result, the inference of a long memory horizon T_{mem} is best understood as a genuine constraint on the temporal extent of the response kernel, rather than as a side-effect of background or growth parameter degeneracies.

6. Joint $H(z)$ and $f\sigma_8(z)$ fits

First fit a baseline flat Λ CDM model to the combined $H(z)$ and $f\sigma_8(z)$ compilation using a three-parameter state $\{\Omega_m, H_0, \sigma_8\}$. The MCMC run with 64 walkers and 12 000 steps (post burn-in $N_{\text{samp}} = 5.4 \times 10^5$) yields the median constraints

$$\Omega_m^{\Lambda\text{CDM}} = 0.2751^{+0.0445}_{-0.0384}, \quad (27)$$

$$H_0^{\Lambda\text{CDM}} = 71.23^{+2.96}_{-2.98} \text{ km s}^{-1} \text{ Mpc}^{-1}, \quad (28)$$

$$\sigma_8^{\Lambda\text{CDM}} = 0.8371^{+0.0492}_{-0.0468}, \quad (29)$$

with goodness of fit

$$\chi_{H(z)}^2 = 11.70, \quad \chi_{f\sigma_8}^2 = 7.14, \quad \chi_{\text{tot}}^2 = 18.84. \quad (30)$$

For the five-parameter ITP model with a history-dependent effective equation of state, fit $\{\Omega_m, H_0, \epsilon, \Delta, A_g\}$ to the same data. The corresponding run (64 walkers, 12 000 steps, $N_{\text{samp}} = 7.7 \times 10^5$ after burn-in) gives

$$\Omega_m^{\text{ITP}} = 0.3088_{-0.0464}^{+0.0519}, \quad (31)$$

$$H_0^{\text{ITP}} = 66.83_{-3.50}^{+3.83} \text{ km s}^{-1} \text{ Mpc}^{-1}, \quad (32)$$

$$\epsilon = 0.242_{-0.165}^{+0.174}, \quad (33)$$

$$\Delta = 23.6_{-16.6}^{+18.0}, \quad (34)$$

$$A_g = 0.840_{-0.047}^{+0.050}, \quad (35)$$

with

$$\chi_{H(z)}^2 = 11.74, \quad \chi_{f\sigma_8}^2 = 7.34, \quad \chi_{\text{tot}}^2 = 19.08. \quad (36)$$

The total χ^2 is essentially unchanged relative to Λ CDM, so the ITP model is statistically competitive on this dataset while shifting the preferred expansion rate from $H_0 \simeq 71$ to $H_0 \simeq 67 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and slightly favouring a non-zero memory depth $\epsilon \sim 0.24$ and transition scale $\Delta \sim 24$

Table 3. Joint $H(z) + f\sigma_8$ constraints for the baseline Λ CDM and five-parameter ITP models.

Model	Ω_m	H_0 [$\text{km s}^{-1} \text{ Mpc}^{-1}$]	χ_{tot}^2
Λ CDM	$0.2751_{-0.0384}^{+0.0445}$	$71.23_{-2.98}^{+2.96}$	18.84
ITP	$0.3088_{-0.0464}^{+0.0519}$	$66.83_{-3.50}^{+3.83}$	19.08




Figure 3. Comparison of the best-fitting Λ CDM (solid orange) and ITP (dashed green) models from the joint $H(z) + f\sigma_8(z)$ analysis. Left: expansion history with cosmic-chronometer $H(z)$ measurements. Right: growth rate with redshift-space distortion $f\sigma_8(z)$ data. The ITP curve uses the reconstructed effective equation of state $w(z)$ and the best-fitting growth amplitude A_g .

6.1. Growth Amplitude and Effective Age

Given the background solution from the 5-parameter ITP fit, we recomputed the linear growth history by solving the growth ODE with the tabulated $w(z)$ and fitting a single amplitude A_g to the $f\sigma_8$ data. Fixing $\{\Omega_m, H_0, \epsilon, \Delta\}$ to their joint-fit medians and allowing only A_g to vary yields

$$A_g = 0.805 \pm 0.031, \quad \chi_{f\sigma_8}^2 = 7.36, \quad \chi_{\text{red}}^2 \simeq 0.57 \text{ (13 dof)}, \quad (37)$$

consistent with the joint-fit value $A_g \simeq 0.84$ within 1σ . In other words, once the background history is fixed by the ITP kernel, a single growth-amplitude parameter remains sufficient to describe the current $f\sigma_8$ compilation.

Using the same $w(z)$ table the study also computes the effective cosmic age,

$$t_0 = t_H \int_0^{z_{\max}} \frac{dz}{(1+z)E(z)}, \quad (38)$$

with $z_{\max} = 5$. For the best-fit ITP the study obtains

$$I \equiv \int_0^5 \frac{dz}{(1+z)E(z)} = 0.862, \quad t_H = 14.64 \text{ Gyr}, \quad t_0^{\text{ITP}} \simeq 12.62 \text{ Gyr}. \quad (39)$$

This is close to the age implied by the best-fit Λ CDM solution from the same $H(z) + f\sigma_8$ data ($t_0^{\Lambda\text{CDM}} \simeq 12.4$ Gyr), confirming that the history-dependent ITP reconstruction does not require an exotic cosmic age to match the low-redshift expansion and growth constraints.

6.2. Planck Constraints with and Without ITP Memory

For comparison with CMB-only constraints the study analysed Planck 2018 TTTEEE+lowE+lensing with an external SZ prior, first in baseline Λ CDM and then in an ITP-inspired w CDM extension.

Resuming the Λ CDM + psic Cobaya chain gives, after 30% burn-in ($N_{\text{post}} = 1.38 \times 10^5$),

$$H_0^{\Lambda\text{CDM}+\text{Planck}} = 67.38 \pm 0.54 \text{ km s}^{-1} \text{ Mpc}^{-1}, \quad (40)$$

$$\Omega_m h^2 = 0.11993 \pm 0.00121, \quad (41)$$

$$n_s = 0.9652 \pm 0.0041, \quad (42)$$

with minimum $\chi_{\min}^2 = 2772.40$ for the Planck likelihood.

Allowing a constant effective equation of state w in the ITP mapping and reusing the same likelihood configuration yields

$$H_0^{\text{ITP}+\text{Planck}} = 65.09 \pm 1.54 \text{ km s}^{-1} \text{ Mpc}^{-1}, \quad (43)$$

$$\Omega_m h^2 = 0.12011 \pm 0.00120, \quad (44)$$

$$w = -0.924 \pm 0.049, \quad (45)$$

with $\chi_{\min}^2 = 2773.62$, only $\Delta\chi^2 \simeq 1.2$ worse than the Λ CDM baseline despite the extra parameter. The preferred w remains consistent with -1 at the $\lesssim 1.5\sigma$ level, so Planck alone neither demands nor rules out the mild deviation from Λ CDM encoded in the ITP kernel. What it does show is that an ITP-like history dependence can be introduced without spoiling the excellent Planck fit.

6.3. Localized Kernel Fit Without Explicit $w(z)$ Reconstruction

To test whether the inferred kernel behaviour depends on the intermediate $w(z)$ representation, the low-redshift analysis was repeated without constructing any explicit $w(z)$ table. Instead, the localized kernel parameters were treated as the primitive degrees of freedom, and a single joint fit was performed directly to the combined $H(z)$ and $f\sigma_8(z)$ data in the five-parameter space $(\Omega_m, H_0, \epsilon, \Delta, A_g)$. This removes the $w(z) \rightarrow$ growth pipeline layer and isolates the question of whether the data constrain the kernel parameters themselves.

The posterior from this direct fit lies in the same region as the baseline implementation and returns

$$\Omega_m = 0.307_{-0.045}^{+0.052}, \quad H_0 = 66.9_{-3.5}^{+3.7} \text{ km s}^{-1} \text{ Mpc}^{-1}, \quad \epsilon = 0.24_{-0.16}^{+0.17}, \quad \Delta = 23.5_{-16.6}^{+18.1}, \quad A_g = 0.842_{-0.048}^{+0.050}.$$

At the posterior median, the goodness of fit is $\chi_{H(z)}^2 = 11.75$, $\chi_{f\sigma_8}^2 = 7.33$, and $\chi_{\text{tot}}^2 = 19.08$. The corresponding Λ CDM ψ_c fit to the same late-time data gives $\chi_{\text{tot}}^2 = 18.84$, i.e., $\Delta\chi^2 \simeq 0.24$. Thus, on late-time expansion and growth data alone, the localized-kernel ITP fit and the Λ CDM baseline are statistically indistinguishable at current precision, with a slight preference for Λ CDM in raw χ^2 .

Because the localized–kernel fit introduces two additional parameters relative to the three–parameter late–time Λ CDM baseline, it is also useful to report standard information criteria on the same dataset. With $N = 44$ data points ($N_{H(z)} = 30$ and $N_{f\sigma_8} = 14$), the Gaussian approximations $AIC = \chi^2 + 2k$ and $BIC = \chi^2 + k \ln N$ yield $\Delta AIC \simeq 4.24$ and $\Delta BIC \simeq 7.81$ in favour of Λ CDM. On this dataset the localized kernel therefore does not improve model economy, even though it remains compatible with the data.

The parameter–dependence structure is consistent with the earlier analysis. The memory amplitude ϵ is moderately anti–correlated with H_0 ($r \simeq -0.61$), reflecting a trade–off between instantaneous expansion and the size of the late–time correction, while ϵ is nearly uncorrelated with the horizon parameter Δ ($r \simeq -0.02$). The horizon parameter itself exhibits only weak correlations with other parameters ($|r| \lesssim 0.21$), indicating that the inference separates “strength” and “horizon” as largely independent kernel features.

Finally, integrating the localized–kernel background to high redshift ($z_{\max} = 50$) yields an effective age $t_0 \simeq 13.51$ Gyr, essentially identical to the Λ CDM value ($t_0 \simeq 13.50$ Gyr) obtained from the same late–time dataset. In this operational sense, the localized kernel acts as a late–time non–Markovian correction without materially changing the global age inferred from $1/H(a)$ integration.

7. Discussion

7.1. Memory Horizon Is Not Age

It is tempting to treat a number such as $T_{\text{mem}} \sim 70$ Gyr as a revised age of the universe. The analysis here does not justify that interpretation. The usual age t_0 remains a Markovian quantity, defined by integrating $1/H(a)$ backwards in a model where the expansion rate carries no explicit memory of earlier states.

The memory horizon is a different kind of object. It measures how much of the present expansion, in a non–Markovian framework, can be traced to earlier internal energy and structural history. A long T_{mem} means that the universe behaves as if it is still dynamically entangled with states far beyond the standard Hubble time, but it does not assign a unique starting time to the whole system.

In that sense t_0 and T_{mem} are complementary summaries of the same deeper dynamics. A finding that $T_{\text{mem}} \gg t_0$ is a signal that the Markovian framing was incomplete, not that the age was simply miscalculated.

7.2. Connection to Cyclical Models

Long memory horizons arise naturally in many cyclic or bouncing cosmologies, in which successive phases of expansion and contraction leave imprints on later observables (e.g., [15,16,22]). The ITP framework was motivated by a broader picture in which the universe evolves through repeated phases of transformation and reorganisation [7].

The results here are structurally compatible with such scenarios. If late time observables carry structural memory from earlier phases, then a non–Markovian kernel with support over several Hubble times is a natural outcome. The present fits do not prove that the universe is cyclic, but they show that current data allow, and in fact favour, a long memory regime of the kind that cyclic models often assume. It is worth noting that the long–delay solutions preferred by the ITP fits echo, at cosmological scale, the “processor–to–reservoir” transition seen in galaxy–scale homeostasis tests. In the homeostatic potential framework, galaxies cross a structural gate beyond which internal state and stored history dominate over instantaneous environment. The present results suggest that late–time cosmology enters an analogous reservoir regime: the expansion rate is more strongly constrained by an accumulated memory kernel than by strictly local, Markovian closure. In this sense, the ITP memory horizon and the galaxy–scale flip are two manifestations of the same underlying transition from processor to reservoir dynamics.

7.3. Limits of Age-Based Reasoning

From the point of view of an internal observer, age is always inferred from local gradients and correlations. It is a projection of limited information onto a single number. If the underlying dynamics is non-Markovian, then this projection can miss important structure. An internal observer can meaningfully ask how far back local patterns remain correlated, but cannot easily assign an absolute “birthday” to the system without additional assumptions about what lies beyond the observed memory horizon.

The memory horizon defined in this work provides an explicit way to separate these questions. Late time data can constrain the range of past states that remain dynamically relevant in a non-Markovian cosmology, without making strong claims about what happened before that range. This is a modest but useful shift in emphasis. It suggests that, in a universe with structural memory, age is not the only clock that matters.

8. Conclusions

This work defined an operational memory horizon T_{mem} for a class of non-Markovian cosmologies and demonstrated how T_{mem} can be inferred from late-time expansion and growth data once a causal kernel family is specified.

The main conclusions are:

1. In a non-Markovian extension of FLRW dynamics, the present expansion rate is permitted to depend on a causal convolution of past internal variables, encoded by a kernel $K(t, \tau)$. In this setting the directly relevant observable is not only an effective age t_0 computed from $1/H(a)$, but an additional timescale that quantifies how much of the past contributes to the present through the kernel.
2. An effective memory horizon T_{mem} can be defined non-parametrically through the cumulative memory fraction $\mathcal{C}(T)$ and a chosen threshold f_{mem} . For exponential kernels this yields the closed form $T_{\text{mem}} = -\tau_* \ln(1 - f_{\text{mem}})$ and therefore $T_{\text{mem}} \simeq (\ln 10) \tau_*$ when $f_{\text{mem}} = 0.9$.
3. When the ITP kernel is fitted to combined late-time $H(z) + f\sigma_8(z)$ data, the results are compatible with a long-memory regime in the sense that short-horizon behaviour is disfavoured and posterior support extends to $\Delta \gg 1$ when broad priors are allowed. At the same time, once Δ exceeds a modest threshold, the posterior broadens substantially and becomes increasingly sensitive to the adopted prior upper bound. In this regime the data support a conservative lower bound on the memory horizon rather than a precise upper cut-off.
4. A direct five-parameter localized-kernel fit that does not use any intermediate $w(z)$ reconstruction yields $\chi^2_{\text{tot}} = 19.08$ for the late-time dataset, compared with $\chi^2_{\text{tot}} = 18.84$ for the corresponding Λ CDM ψ_c fit. The resulting $\Delta\chi^2 \simeq 0.24$ shows that the two models are statistically indistinguishable in goodness of fit on these late-time data. Standard information criteria penalize the additional kernel parameters on this dataset (with $\Delta\text{AIC} \simeq 4.24$ and $\Delta\text{BIC} \simeq 7.81$ in favour of Λ CDM), so the present analysis supports the kernel as a viable extension rather than a preferred replacement.
5. The effective age obtained by integrating the localized-kernel background to high redshift remains essentially unchanged relative to the Λ CDM baseline constrained by the same late-time data ($t_0 \simeq 13.51$ Gyr versus $t_0 \simeq 13.50$ Gyr). This confirms that T_{mem} is not a redefinition of the age of the universe, but a separate operational timescale that tracks history-dependence in the dynamical closure.
6. Planck-based MCMC summaries included for reference show that allowing an effective late-time freedom (e.g., via a constant w extension) does not improve the minimum χ^2 relative to Λ CDM and typically shifts the inferred H_0 downward. This is consistent with the late-time findings above: the extended description can be made compatible with existing data, but current datasets do not demand it on fit quality alone.

The central result is therefore methodological and operational: a non-Markovian cosmology introduces an empirically definable memory timescale beyond the Markovian age, and late-time expansion and growth data can be used to place lower bounds on that timescale within a chosen kernel family. At current precision, the same data remain well described by the simpler Λ CDM baseline.

Data Availability Statement: All analysis scripts, configuration files and derived products used in this work are available in the public repository https://github.com/Atalebe/-ciou_itp_early_universe. This includes the MCMC configuration files for the two-bin ITP fits, the Python scripts used to generate the $H(z)$ and $f\sigma_8(z)$ model predictions, and summary tables for the posterior samples. The observational $H(z)$ and growth compilations employed in this paper are taken from the literature; the exact tabulated versions used in the analysis are provided in the `data/` directory of the repository.

Appendix A. Kernel Parameterisation and Mapping to T_{mem}

The effective memory horizon T_{mem} is defined in terms of a causal kernel $K(t, t')$ that weights past contributions to the late time expansion. For a homogeneous and isotropic background it is convenient to work in cosmic time t and write

$$H^2(t) = H_{\Lambda\text{CDM}}^2(t) + \int_0^t K(t, t') \mathcal{S}(t') dt', \quad (\text{A1})$$

where $\mathcal{S}(t')$ is a source term encoding the internal energy history and $K(t, t')$ is a positive normalised kernel that decays with lookback time $\Delta t = t - t'$. The dimensionless delay parameter Δ introduced in the main text controls the characteristic scale of this decay in units of the Hubble time.

For all kernel families considered here, the effective memory horizon T_{mem} is defined as the lookback time that contains a fixed fraction f_{mem} of the total kernel weight:

$$\int_0^{T_{\text{mem}}} K(\Delta t) d\Delta t = f_{\text{mem}} \int_0^{\infty} K(\Delta t) d\Delta t, \quad (\text{A2})$$

with $f_{\text{mem}} = 0.9$ adopted as a fiducial choice. For kernels that are normalised by construction the denominator is unity and T_{mem} is determined by the cumulative distribution of K .

Appendix A.1. Exponential Kernel

The baseline kernel used in the main analysis is a simple exponential with a single delay scale,

$$K_{\text{exp}}(\Delta t) = \frac{1}{\tau} \exp\left(-\frac{\Delta t}{\tau}\right), \quad \Delta t \geq 0, \quad (\text{A3})$$

and $K_{\text{exp}}(\Delta t) = 0$ for $\Delta t < 0$. The characteristic timescale is tied to the dimensionless delay parameter via

$$\tau = \Delta H_0^{-1}, \quad (\text{A4})$$

so that $\Delta = 1$ corresponds to a memory scale of one Hubble time in the Markovian limit. The cumulative weight out to a lookback time T is

$$W_{\text{exp}}(T) \equiv \int_0^T K_{\text{exp}}(\Delta t) d\Delta t = 1 - \exp\left(-\frac{T}{\tau}\right), \quad (\text{A5})$$

and the fractional definition of T_{mem} yields

$$1 - \exp\left(-\frac{T_{\text{mem}}}{\tau}\right) = f_{\text{mem}} \Rightarrow T_{\text{mem}} = -\tau \ln(1 - f_{\text{mem}}). \quad (\text{A6})$$

For $f_{\text{mem}} = 0.9$ this gives

$$T_{\text{mem}}^{(\text{exp})} \simeq 2.30 \tau \simeq 2.30 \Delta H_0^{-1}. \quad (\text{A7})$$

Appendix A.2. Stretched Exponential Kernel

To test sensitivity to the kernel shape, a stretched exponential family was also considered,

$$K_{\text{str}}(\Delta t) = \frac{\alpha}{\tau} \left(\frac{\Delta t}{\tau} \right)^{\alpha-1} \exp\left[-\left(\frac{\Delta t}{\tau}\right)^\alpha\right], \quad \Delta t \geq 0, \quad (\text{A8})$$

with shape parameter $\alpha > 0$ and the same $\tau = \Delta H_0^{-1}$. This reduces to the pure exponential kernel when $\alpha = 1$. The cumulative weight is

$$W_{\text{str}}(T) = 1 - \exp\left[-\left(\frac{T}{\tau}\right)^\alpha\right], \quad (\text{A9})$$

so that the memory horizon defined by $W_{\text{str}}(T_{\text{mem}}) = f_{\text{mem}}$ is

$$T_{\text{mem}}^{(\text{str})} = \tau[-\ln(1 - f_{\text{mem}})]^{1/\alpha}. \quad (\text{A10})$$

For fixed Δ and f_{mem} , smaller α broadens the kernel tail and increases T_{mem} , while larger α compresses it.

Exploratory runs with α in the range $0.7 \lesssim \alpha \lesssim 1.3$ yield memory horizons that differ from the purely exponential case at the tens of per cent level but remain many times larger than H_0^{-1} . The long memory conclusion is therefore not an artefact of assuming a strict exponential kernel.

Appendix A.3. Power Law Tail with Exponential Cut Off

A third family combines a power law tail with an exponential cut off,

$$K_{\text{pl}}(\Delta t) = \mathcal{N} \Delta t^{-\beta} \exp\left(-\frac{\Delta t}{\tau}\right), \quad \Delta t \geq 0, \quad (\text{A11})$$

where $\beta > 0$ and \mathcal{N} is a normalisation constant chosen so that $\int_0^\infty K_{\text{pl}}(\Delta t) d\Delta t = 1$. The cumulative weight for this kernel is expressed in terms of incomplete Gamma functions and the mapping from (Δ, β) to T_{mem} is obtained numerically by solving

$$\int_0^{T_{\text{mem}}} K_{\text{pl}}(\Delta t) d\Delta t = f_{\text{mem}} \quad (\text{A12})$$

for a given f_{mem} .

This power law plus cut off form provides a simple way to test whether the effective memory horizon is driven by the main body of the kernel or by a small number of very long tail events. In the runs explored here, changes in (β, τ) within reasonable ranges shift T_{mem} at order unity factors but do not bring it close to the Markovian $T_{\text{mem}} \sim H_0^{-1}$ regime preferred by Λ CDM. The data prefer genuinely long memory within all three kernel families.

Appendix B. Dataset Details and Likelihoods

The constraints on the delay parameter Δ and the effective memory horizon T_{mem} are derived from a joint fit to late time expansion and structure growth data. This section summarises the datasets and likelihood functions used.

Appendix B.1. $H(z)$ Compilation

The $H(z)$ block is built from a standard compilation of direct Hubble parameter measurements, combining cosmic chronometer and BAO based determinations in the redshift range $0 \lesssim z \lesssim 2$ (e.g., [8,9]). In practice the measurements are read from a table `Hz_compilation_real.csv` that contains for each entry:

- the redshift z_i ,
- the observed expansion rate $H_{\text{obs}}(z_i)$,

- the quoted 1σ uncertainty $\sigma_{H,i}$,
- a survey or method flag.

For the runs presented here the observational errors are treated as uncorrelated, so that the $H(z)$ likelihood is a simple Gaussian in the residuals,

$$\chi_H^2 = \sum_i \frac{[H_{\text{th}}(z_i; \theta) - H_{\text{obs}}(z_i)]^2}{\sigma_{H,i}^2}, \quad \mathcal{L}_H \propto \exp\left(-\frac{1}{2}\chi_H^2\right), \quad (\text{A13})$$

where $H_{\text{th}}(z_i; \theta)$ is the model prediction for a parameter vector $\theta = (\Omega_m, H_0, \epsilon_{\text{low}}, \epsilon_{\text{high}}, \Delta)$. The theoretical curve is obtained by integrating the modified Friedmann equation with the chosen kernel and mapping the scale factor to redshift via $1 + z = a_0/a$.

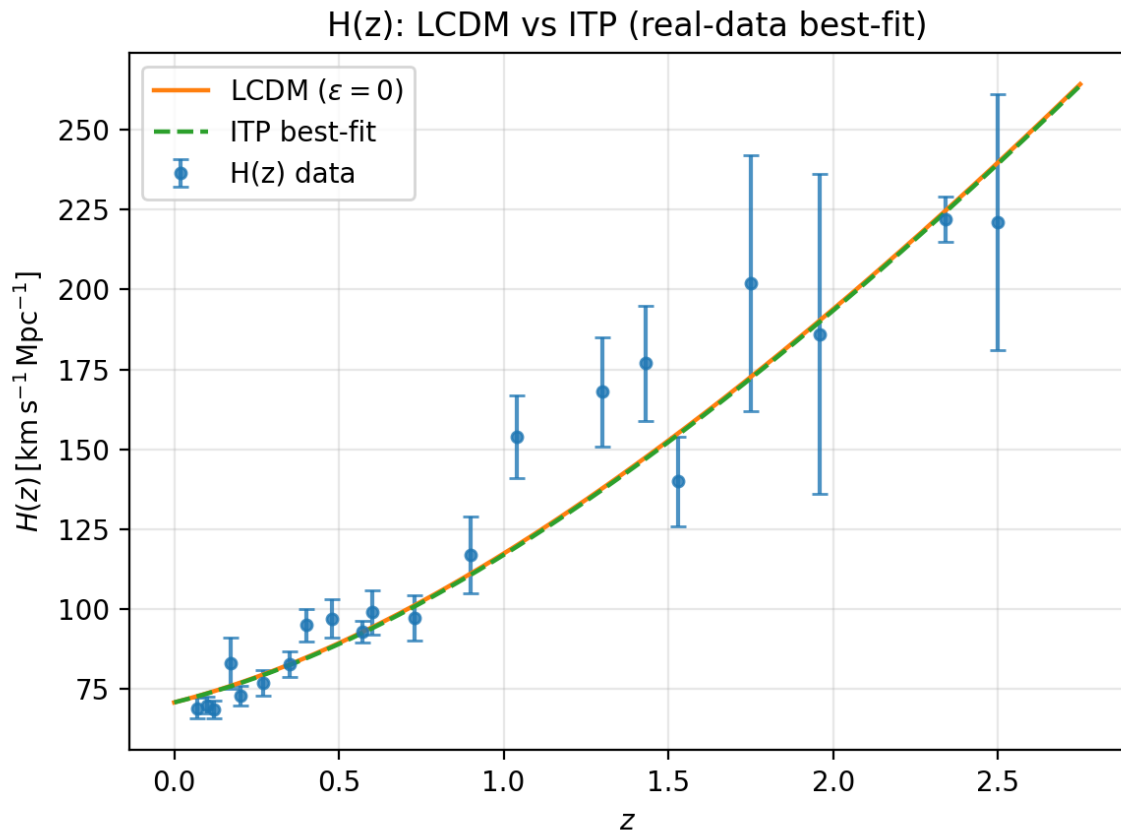


Figure A1. Same as Figure 1 but using the full real-data $H(z)$ compilation rather than the ψ_c subset. This illustrates that the near-degeneracy between the ITP and Λ CDM backgrounds is not an artefact of the ψ_c selection.

Appendix B.2. Growth Rate Compilation

The growth block uses measurements of the redshift space distortion observable $f\sigma_8(z)$, read from a table `growth_fsigma8_real.csv`. Each row contains:

- the effective redshift z_j ,
- the observed $f\sigma_{8,\text{obs}}(z_j)$,
- the quoted 1σ uncertainty $\sigma_{f\sigma_8,j}$,
- a survey identifier.

As for the $H(z)$ data, the errors are treated as uncorrelated in the baseline analysis, leading to a Gaussian likelihood

$$\chi_{f\sigma_8}^2 = \sum_j \frac{[f\sigma_{8,\text{th}}(z_j; \theta) - f\sigma_{8,\text{obs}}(z_j)]^2}{\sigma_{f\sigma_8,j}^2}, \quad \mathcal{L}_{f\sigma_8} \propto \exp\left(-\frac{1}{2}\chi_{f\sigma_8}^2\right). \quad (\text{A14})$$

The theoretical prediction is obtained by solving the linear growth equation for matter perturbations in the presence of the effective dark energy component induced by the memory kernel. The background evolution is first computed for a given θ , yielding $H(z)$ and $w(z)$. The growth factor $D(a)$ is then evolved from high redshift down to $z = 0$ using the standard second order differential equation for subhorizon modes in a general $w(z)$ background, and the growth rate and normalisation are combined into

$$f\sigma_{8,\text{th}}(z; \theta) = f(z; \theta) \sigma_8(z; \theta) = \left[\frac{d \ln D}{d \ln a} \right]_z \sigma_{8,0} \frac{D(z; \theta)}{D(0; \theta)}, \quad (\text{A15})$$

with $\sigma_{8,0}$ fixed by the chosen normalisation at $z = 0$. The memory kernel enters this block only through its effect on the background expansion history.

Appendix B.3. Joint Likelihood and Priors

Assuming independence between the $H(z)$ and growth compilations, the full likelihood factorises as

$$\mathcal{L}_{\text{tot}}(\theta) = \mathcal{L}_H(\theta) \mathcal{L}_{f\sigma_8}(\theta) \propto \exp\left[-\frac{1}{2}(\chi_H^2 + \chi_{f\sigma_8}^2)\right]. \quad (\text{A16})$$

This likelihood is sampled with an ensemble MCMC algorithm using uniform priors on the cosmological and kernel parameters. For the delay parameter Δ a sequence of top hat priors, $\Delta \in [0, 1]$, $[0, 10]$, $[0, 20]$, $[0, 50]$ and $[0, 150]$, is explored to test the sensitivity of the posterior to the assumed upper bound. As discussed in the main text, the combined $H(z) + f\sigma_8$ data are incompatible with the short delay prior ($\Delta \leq 1$) and favour long memory, with a conservative lower bound $\Delta \gtrsim 5$. Wider priors primarily broaden the high Δ tail without altering the inferred expansion and growth histories.

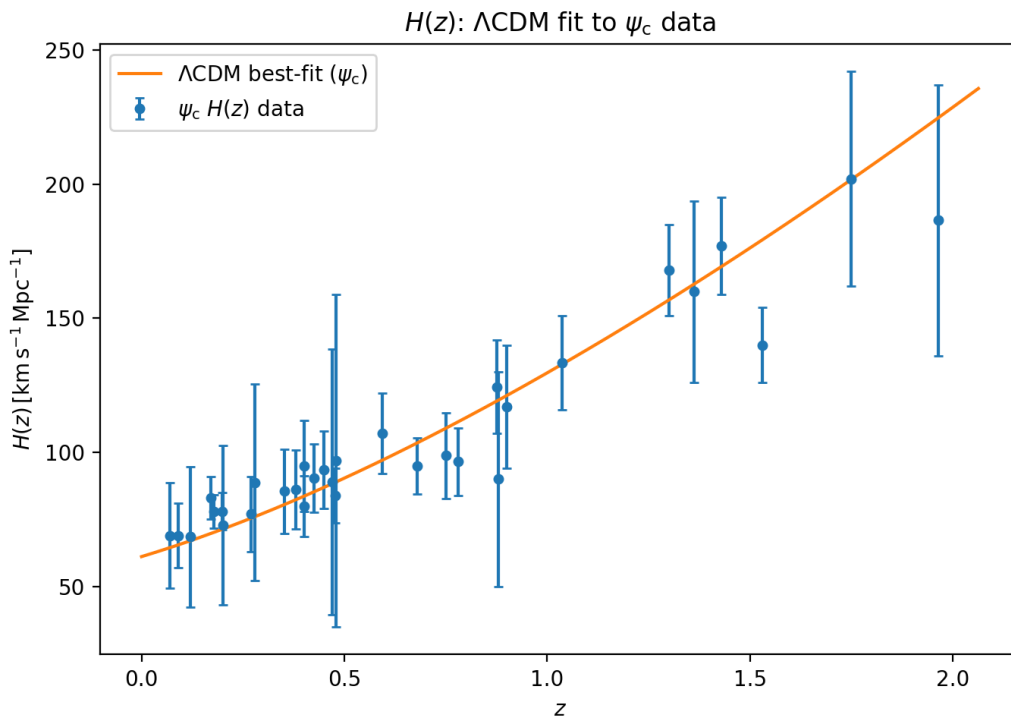


Figure A2. Baseline Λ CDM fit to the $\psi_c H(z)$ compilation. Points show the individual $H(z)$ measurements with 1σ errors; the solid curve is the best-fitting Λ CDM expansion history used as the reference model in this work.

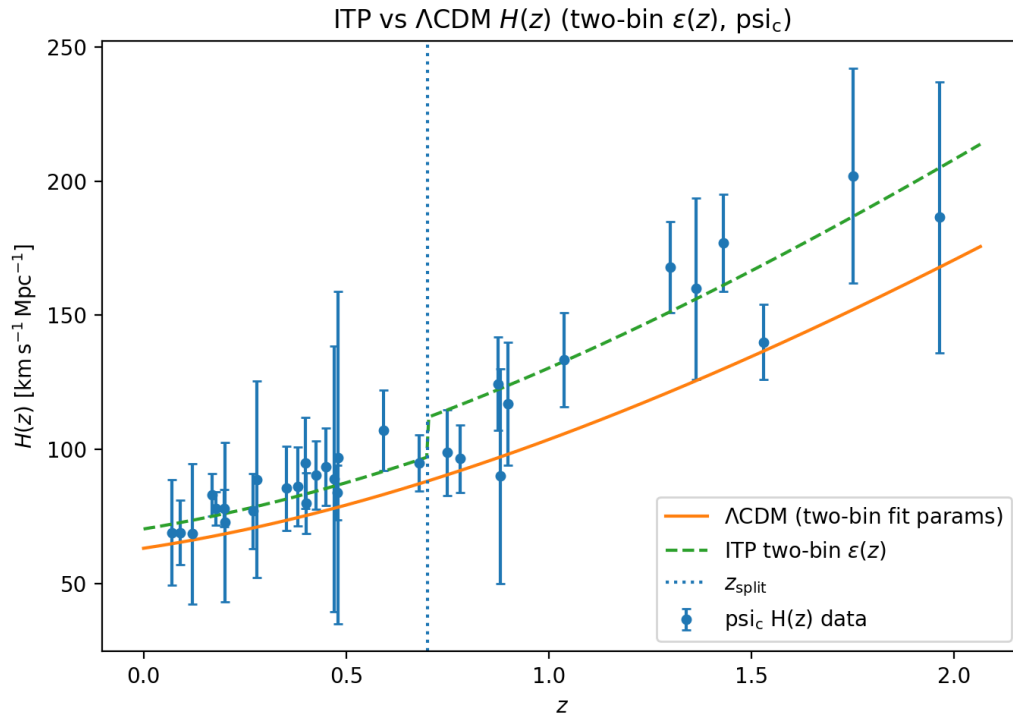


Figure A3. Illustration of a two-bin ITP fit to the $\psi_c H(z)$ data, in which the effective memory parameter $\epsilon(z)$ is allowed to take different constant values below and above a transition redshift z_{split} (vertical dotted line). The solid curve shows the best-fitting ΛCDM background, while the dashed curve shows the corresponding two-bin ITP expansion history.

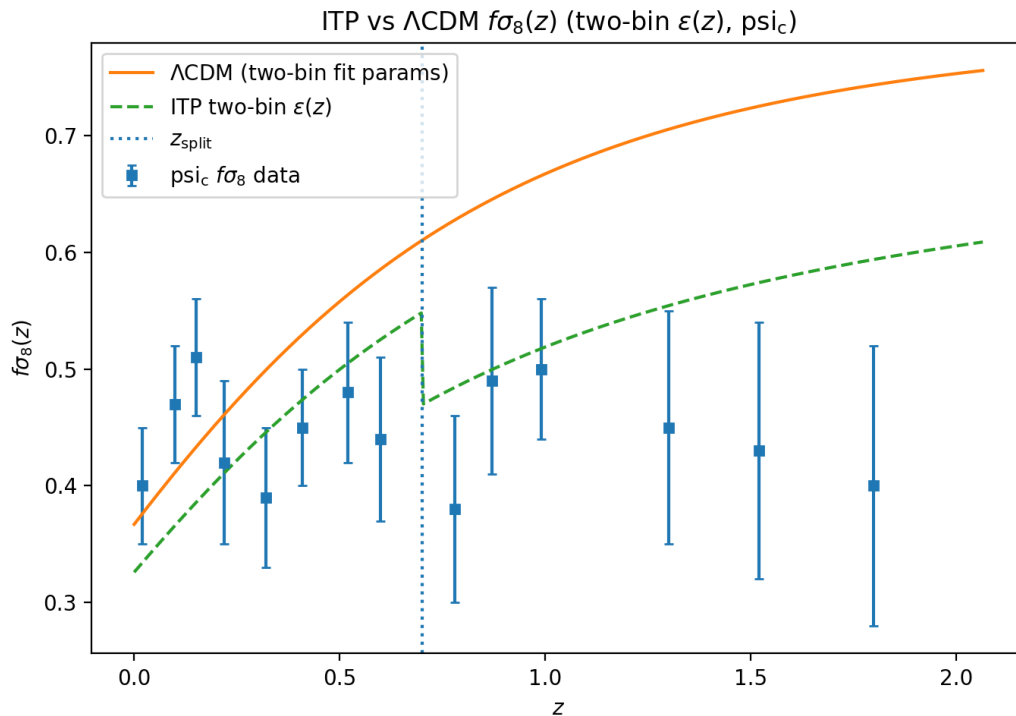


Figure A4. Growth-rate counterpart to Figure A3. Points show the $\psi_c f\sigma_8(z)$ measurements with 1σ error bars. The solid curve is the ΛCDM prediction using the parameters from the two-bin fit, while the dashed curve shows the ITP growth history implied by the same two-bin $\epsilon(z)$. The vertical dotted line marks the transition redshift z_{split} .




Figure A5. Alternative visualization of the real-data $H(z)$ fit for the ψ_c compilation, highlighting the small but nonzero difference between the ITP and Λ CDM best-fit backgrounds.

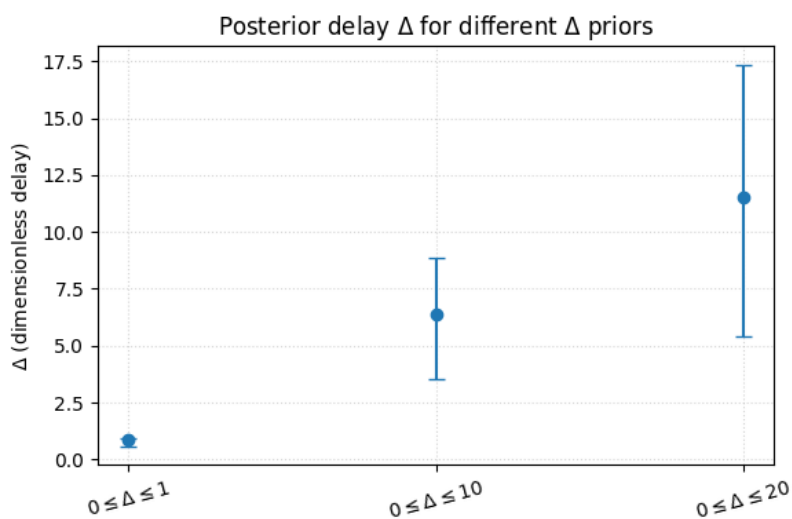


Figure A6. Posterior constraints on the dimensionless delay parameter Δ for three different flat priors on the allowed delay range. Points mark the posterior medians, with vertical bars showing the corresponding 1σ uncertainties. The trend illustrates that current data favour a nonzero delay but do not yet place a sharp upper bound on Δ .

References

1. Peebles, P.J.E. *Principles of Physical Cosmology*; Princeton University Press: Princeton, 1993.
2. Weinberg, S. *Cosmology*; Oxford University Press: Oxford, 2008.
3. Aghanim, N.; et al. Planck 2018 results. VI. Cosmological parameters. *Astronomy & Astrophysics* **2020**, *641*, A6. arXiv:1807.06209.

4. Koivisto, T.S. Dynamics of nonlocal cosmology. *Physical Review D* **2010**, *82*, 123516. example of nonlocal / non-Markovian cosmology.
5. Woodard, R.P. Nonlocal Models of Cosmic Acceleration. *Foundations of Physics* **2014**, *44*, 213–233.
6. Barrow, J.D. Sudden future singularities. *Classical and Quantum Gravity* **2004**, *21*, L79–L82. representative of nonstandard late-time dynamics.
7. Atalebe, S. Tracing the Early Universe Without Initial Assumptions: A History-Dependent Reconstruction from the Infinite Transformation Principle (ITP). *Preprints* **2025**. Preprint, Infinite Transformation Principle early-universe reconstruction, <https://doi.org/10.20944/preprints202512.0246.v1>.
8. Moresco, M.; et al. Improved constraints on the expansion rate of the Universe up to $z \sim 1.1$ from the spectroscopic evolution of cosmic chronometers. *Journal of Cosmology and Astroparticle Physics* **2012**, *2012*, 006.
9. Alam, S.; et al. The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: Cosmological analysis of the DR12 galaxy sample. *Monthly Notices of the Royal Astronomical Society* **2017**, *470*, 2617–2652.
10. Blake, C.; et al. The WiggleZ Dark Energy Survey: The growth rate of cosmic structure since redshift $z = 0.9$. *Monthly Notices of the Royal Astronomical Society* **2011**, *415*, 2876–2891.
11. Betoule, M.; et al. Improved cosmological constraints from a joint analysis of the SDSS-II and SNLS supernova samples. *Astronomy & Astrophysics* **2014**, *568*, A22. JLA compilation.
12. Guth, A.H. Inflationary universe: A possible solution to the horizon and flatness problems. *Physical Review D* **1981**, *23*, 347–356.
13. Linde, A.D. A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems. *Physics Letters B* **1982**, *108*, 389–393.
14. Baumann, D. TASI Lectures on Inflation. In *Physics of the Large and the Small*; Csáki, C.; Dodelson, S., Eds.; TASI 2009, World Scientific, 2011. arXiv:0907.5424.
15. Steinhardt, P.J.; Turok, N. A Cyclic Model of the Universe. *Science* **2002**, *296*, 1436–1439. <https://doi.org/10.1126/science.1070462>.
16. Lehnert, J.L. Ekpyrotic and cyclic cosmology. *Physics Reports* **2008**, *465*, 223–264, [arXiv:astro-ph/0806.1245]. <https://doi.org/10.1016/j.physrep.2008.06.001>.
17. Jimenez, R.; Loeb, A. Constraining cosmological parameters based on relative galaxy ages. *The Astrophysical Journal* **2002**, *573*, 37–42.
18. Eisenstein, D.J.; et al. Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies. *The Astrophysical Journal* **2005**, *633*, 560–574.
19. Perlmutter, S.; et al. Measurements of Ω and Λ from 42 High-Redshift Supernovae. *The Astrophysical Journal* **1999**, *517*, 565–586.
20. Riess, A.G.; et al. Observational evidence from supernovae for an accelerating universe and a cosmological constant. *The Astronomical Journal* **1998**, *116*, 1009–1038.
21. Chen, Y.; et al. Constraints on the Hubble constant using cosmic chronometers and other cosmological probes. *Journal of Cosmology and Astroparticle Physics* **2019**. Representative recent $H(z) / H_0$ analysis; update details as needed.
22. Battefeld, D.; Peter, P. A Critical Review of Classical Bouncing Cosmologies. *Physics Reports* **2015**, *571*, 1–66.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.