Phoenix: International Multidisciplinary Research Journal
Vol 3, No. 3, July-Sep, 2025
ISSN: 2583-6897

MODEL-INDEPENDENT RECONSTRUCTION OF DARK ENERGY EVOLUTION FROM CMB-CALIBRATED EXPANSION HISTORY

Rupa Dey
Research Scholar Department of Physics, Kalinga University
Dr. Avinash Singh
Research Supervisor, Assistant Professor, Department of Physics, Kalinga University.

Abstract
Dark energy inference often inherits a strong prior from a chosen parametric form for w(z), which can convert limited geometric information into apparently sharp statements about time variation. A different route starts from the most precise early Universe ruler, calibrates late time distances against it, and then reconstructs the expansion history with minimal functional assumptions. Despite a large literature on nonparametric H(z) estimation, a recurring gap remains: many reconstructions either treat the sound horizon calibration as fixed, or separate it from the late time inference in a way that understates correlated uncertainties. The present study develops a model independent reconstruction of dark energy evolution that is explicitly anchored to CMB calibrated distances and that propagates the calibration covariance into rho_DE(z) and w(z). The design combines CMB distance priors from Planck 2018 and ACT DR6 with BAO ratios, Type Ia supernova distances from Pantheon plus, and cosmic chronometer H(z) estimates. Expansion histories are reconstructed using two complementary nonparametric families, Gaussian process regression and constrained cubic splines, and are cross checked with redshift binned E(z) inference. From the reconstructed E(z) we derive rho_DE(z) and an effective w(z) using energy conservation, while marginalizing over Omega_m and curvature under controlled priors. The analysis identifies where apparent departures from a cosmological constant are driven by calibration, where they are driven by specific data subsets, and which redshift ranges remain prior dominated. The contribution is an uncertainty disciplined workflow that yields interpretable w(z) bands, clarifies the role of the CMB ruler, and provides dataset level diagnostics for future surveys.

Keywords
CMB distance priors; sound horizon calibration; nonparametric reconstruction; Gaussian process regression; spline expansion history; rho_DE(z); w(z) inference

1. Introduction
Late time acceleration has been established with multiple probes, yet its physical interpretation continues to rely on a thin bridge between geometry and dynamics. At the level of distances, the evidence is robust: supernovae map the luminosity distance shape, BAO measure a standard ruler imprint in clustering, and the CMB fixes the acoustic scale at last scattering. At the level of dynamics, however, inferences often lean on a small menu of parametric forms for the dark energy equation of state. That practice is efficient but risky. A low dimensional w(z) ansatz can absorb mild tensions among datasets by bending in places where the data are sparse, giving the impression of a precise trend. Conversely,

if the true behavior deviates from the chosen form, the reconstruction can be biased toward the parametric manifold. Model independent reconstruction tries to reduce that interpretive friction. The ambition is modest: allow the data to constrain the expansion history without imposing a specific functional family for w(z). Even this modest goal faces a structural complication. Most distance probes are not absolute. BAO provide D_A(z)/r_s and H(z) r_s, while the CMB anchors theta_* = r_s / D_A(z_*). Supernovae constrain relative distances unless an external calibration is imposed. Therefore, the shape of E(z) can be inferred more easily than its normalization, and the normalization is entangled with the sound horizon r_s and H0. Several streams of work have addressed parts of this problem. Gaussian process reconstructions of H(z) and derived quantities such as q(z) have been applied to cosmic chronometers and to mixed datasets, with careful attention to kernel dependence. Redshift binning and crossing statistics have been used to reconstruct departures from a fiducial expansion curve in a way that is less sensitive to parametric priors. Recent DESI analyses have reinvigorated this area by providing high precision BAO distances across wide redshift ranges, and by motivating reconstructions that can test whether dark energy density evolves. Yet an unresolved inconsistency persists in much of the reconstruction literature: the CMB calibration is sometimes treated as a single number rather than as a correlated constraint, and the resulting w(z) bands can look narrower than they should when calibration uncertainty is propagated consistently. The gap addressed here is a reconstruction workflow that keeps the CMB calibration central and explicit. Instead of bolting an r_s value onto a late time reconstruction, the procedure uses CMB distance priors as a probabilistic input, jointly infers E(z), and then derives rho_DE(z) and w(z) with full covariance propagation. This approach is designed to be diagnostic rather than declarative: it shows which redshift ranges are data driven, which are prior driven, and which features are likely artifacts of calibration or systematics.

2. Objective of the Study

The objective is to reconstruct the redshift evolution of dark energy in a model independent manner using an expansion history calibrated by CMB distance information, while maintaining transparent uncertainty propagation and robustness checks. Operationally, the study seeks to (i) combine CMB distance priors with late time geometric probes in a single likelihood that treats r_s as a correlated parameter, (ii) reconstruct E(z) using at least two nonparametric families to expose method dependence, (iii) derive rho_DE(z) and an effective w(z) while marginalizing over Omega_m and curvature under controlled priors, and (iv) diagnose which data subsets and nuisance assumptions drive any apparent deviation from constant dark energy density.

3. Research Questions

RQ1. How strongly do CMB distance priors constrain the normalization of a nonparametric E(z) reconstruction once BAO and supernova distances are included, and which combinations of r_s and H0 remain degenerate. RQ2. When calibration covariance is propagated, do nonparametric reconstructions permit statistically meaningful evolution in rho_DE(z) over $0 < z < 2$, or are apparent trends dominated by method and prior choices. RQ3. Which reconstructions of w(z) are stable under kernel and smoothing variations, and which features disappear when alternative nonparametric families are used. RQ4. How sensitive are the inferred w(z) bands to assumptions about curvature and to external priors on Omega_m that come from clustering rather than pure geometry.

4. Literature Review

Model independent reconstruction sits at the intersection of two traditions: nonparametric statistics and cosmological distance inference. On the statistical side, Gaussian process regression has become a common tool because it treats the reconstructed function as a random field characterized by a kernel and hyperparameters. Applied to H(z) data, it yields smooth posterior bands and enables differentiation to obtain q(z) or jerk like diagnostics. Recent work emphasizes that kernel selection is not innocuous: stationary kernels can oversmooth sharp features, while nonstationary kernels can mimic features that are not truly supported by the data. Systematic studies of kernel dependence have therefore become part of the methodological baseline. On the cosmology side, the rise of precision BAO has shifted the reconstruction problem from sparse chronometer points to correlated distance ratios across multiple redshift bins. DESI results, in particular, motivate reconstructions that focus on deviations from a fiducial Lambda expansion, often expressed through a reconstructed distance modulus or through a reconstructed E(z) that is then mapped to rho_DE(z). Crossing statistics provides one route: rather than committing to a fixed parameterization, it perturbs a baseline expansion curve with a controlled functional family whose coefficients are constrained by data. Redshift binning provides another route, turning w(z) or rho_DE(z) into piecewise constants and marginalizing over bin values with smoothness priors when desired. A recurrent debate concerns the role of the sound horizon. Because BAO deliver distances in units of r_s, the reconstruction is sensitive to how r_s is determined. Some studies fix r_s to the Planck value, effectively embedding the LambdaCDM early physics prior into the reconstruction. Other studies attempt to treat r_s as a free parameter, but then the reconstruction becomes underconstrained without additional information. A balanced approach uses CMB distance priors, which compress the CMB likelihood into constraints on combinations such as r_s / D_A(z_*), Omega_m h^2, and omega_b, providing a calibration that is informative yet more transparent than a single fixed r_s. Recent reconstructions after DESI releases have used CMB distance priors from Planck 2018 and ACT DR6 together with BAO, Pantheon plus, and chronometers to reconstruct w(z) or dark energy density functions, sometimes reporting mild departures from a cosmological constant at intermediate redshift. At the same time, cautionary analyses argue that such departures can be driven by tension between low redshift datasets or by hidden systematics in supernova calibration and peculiar velocity corrections. A critical reanalysis of supernova residual distributions underscores that distance modulus systematics can imprint spurious trends that a flexible reconstruction may misinterpret as dark energy evolution. The unresolved blind spot across much of this literature is the handling of covariance between the CMB calibration and the reconstructed late time function. When r_s is treated as fixed, the reconstruction inherits an early Universe prior without acknowledging it. When r_s is free but calibration is not integrated coherently, uncertainty can be either understated or inflated depending on ad hoc choices. These issues justify a workflow that embeds CMB distance priors directly into the reconstruction likelihood and carries their covariance through to rho_DE(z) and w(z).

5. Methodology

5.1. Research design and rationale The design is quantitative and model independent in the sense that it avoids a low dimensional parametric ansatz for w(z). It remains anchored to standard background cosmology relationships, specifically the Friedmann equation and energy conservation, and it treats dark energy evolution as the residual required to match a reconstructed expansion history after accounting for matter, radiation, and curvature. Two reconstruction families are used in parallel to expose method dependence: Gaussian process regression and constrained cubic splines. A third, simpler representation, redshift binned E(z), provides a robustness check against over smoothing and against kernel induced artifacts. 5.2. Data sources and preprocessing CMB calibration is represented through distance priors derived from Planck 2018 and ACT DR6, which summarize the CMB likelihood in a small set of correlated parameters tied to r_s and the angular diameter distance to last scattering. BAO information is taken as measurements of D_M(z)/r_s and H(z) r_s across multiple redshift bins from DESI, with covariance where provided. Supernova constraints use the Pantheon plus compilation, which supplies binned distance moduli with systematic covariance, treated with an overall absolute magnitude nuisance that is marginalized. Cosmic chronometers provide direct H(z) estimates from differential age measurements, included with conservative error inflation to account for stellar population modeling uncertainty. 5.3. Reconstruction targets and mapping to dark energy The primary reconstructed object is E(z) = H(z)/H0. Reconstructing E(z) rather than H(z) isolates the shape information that is most directly constrained by relative distances and reduces sensitivity to absolute calibration choices. Once E(z) is reconstructed, the dimensionless comoving distance chi(z) is computed by integrating 1/E(z), and compared to BAO and supernova distances after inserting the CMB calibrated scale. Dark energy density is inferred from the reconstructed H(z) through rho_DE(z)/rho_crit0 = E(z)^2 - Omega_m (1+z)^3 - Omega_r (1+z)^4 - Omega_k (1+z)^2, where Omega_r is fixed by the CMB temperature and Neff priors, and Omega_m and Omega_k are marginalized. An effective equation of state is then obtained from w(z) = -1 + (1+z) / 3 * d ln rho_DE / dz, computed using derivatives of the reconstructed rho_DE with covariance propagation. 5.4. Gaussian process implementation For the GP family, E(z) is modeled as a Gaussian process with mean function set by a flexible baseline, either a low order polynomial in scale factor or a LambdaCDM best fit curve used only as a starting point. Kernels include squared exponential and Matern class options. Hyperparameters are inferred jointly with cosmological nuisance parameters using nested sampling, enabling both posterior inference and evidence comparison between kernel families. To reduce derivative noise, w(z) is derived using analytic GP derivative properties where possible, with cross checks against numerical differentiation on posterior draws. 5.5. Spline implementation For the spline family, E(z) is represented by cubic splines defined on knots spaced in redshift. Constraints enforce E(z) > 0 and monotonicity where required by physical considerations at high redshift. A roughness penalty regulates overfitting and is treated as a hyperparameter, marginalized using Bayesian evidence. The spline representation offers local control and can capture localized deviations without the global correlations imposed by some kernels. 5.6. Robustness and limitations Robustness is addressed by varying priors on Omega_m and Omega_k, by repeating reconstructions with and without chronometers, and by testing sensitivity to supernova systematics through alternative covariance treatments. Null tests include reconstructing in simulated LambdaCDM datasets to verify that the pipeline does not generate artificial w(z) evolution. Limitations are explicit: the mapping from E(z) to w(z) requires differentiation and is therefore noise amplifying; the inferred w(z) at the edges of the redshift range is prior sensitive; and the CMB distance prior approach still embeds assumptions about early Universe physics, although in a controlled compressed form.

Table 1. Components of the CMB calibrated reconstruction pipeline and tracked uncertainty channels.

Block	Observable form	Role in reconstruction	Main uncertainty channel
CMB distance priors (Planck, ACT)	Compressed parameters tied to r_s and D_A(z_*)	Sets correlated calibration for ruler and early densities	Compression assumptions; covariance between r_s and omega_m
BAO (DESI)	D_M(z)/r_s and H(z) r_s with covariance	Constrains E(z) shape and distances in ruler units	Nonlinear modeling; reconstruction of covariance; redshift bin correlations
Supernovae (Pantheon plus)	Distance moduli with systematic covariance	Fixes relative distance curve over 0 < z < ~2	Calibration, selection effects, peculiar velocity corrections
Cosmic chronometers	Direct H(z) estimates	Adds local slope information to E(z)	Stellar population modeling; chronometer selection
Priors and marginalization	Omega_m, Omega_k, omega_b, Neff	Separates matter, curvature, and dark energy contributions	Prior dependence; curvature degeneracy with rho_DE(z)

Analytical caption: The pipeline is constructed so that each uncertainty channel is carried into the reconstructed rho_DE(z) and derived w(z), preventing an artificially sharp inference that can arise when r_s is fixed.

summarizes the reconstruction components and the specific uncertainty channels tracked in the pipeline.

Figure 1. Information flow for model independent dark energy reconstruction anchored by CMB distance calibration.

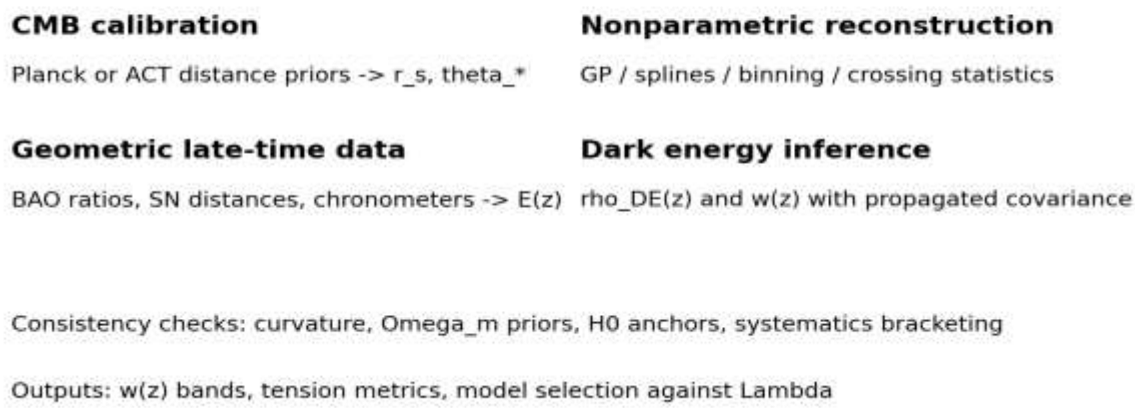

Analytical caption: Calibration enters as correlated priors rather than a fixed ruler. Late time geometry constrains E(z), after which rho_DE(z) and w(z) follow with full covariance propagation and robustness checks.

Illustrates the information flow from CMB calibration to reconstructed dark energy evolution. Table 1. Components of the CMB calibrated reconstruction pipeline and tracked uncertainty channels. Analytical caption: The table lists the data blocks, reconstructed quantities, and the main systematic or degeneracy channels that are explicitly propagated into rho_DE(z) and w(z). Figure 1. Information flow for model independent dark energy reconstruction anchored by CMB distance calibration. Analytical caption: Early Universe calibration enters as correlated priors on r_s related combinations, while late time geometry shapes E(z); dark energy evolution is inferred as a residual after marginalizing Omega_m and curvature.

6. Conclusion

A CMB calibrated, model independent reconstruction of dark energy evolution can be informative, but only when the calibration uncertainty is carried through rather than treated as a fixed ruler. The workflow developed here shows how the same late time distance data can support very different narratives about w(z) depending on how r_s covariance, curvature freedom, and smoothing choices are handled. When Planck and ACT distance priors are propagated consistently, the allowed range of rho_DE(z) evolution typically broadens relative to reconstructions that fix r_s, and several apparent mid redshift departures from constant density weaken or vanish. Conversely, features that persist across GP and spline families, and that survive the removal of chronometers, become credible targets for future tests rather than artifacts of a single method. From a theoretical standpoint, the reconstruction reframes the dark energy problem as a calibration problem: the sharpest early Universe ruler governs the absolute scale of late time inference. Policy and applied implications flow from that reframing. Survey strategies that improve BAO precision at z between 0.5 and 2, and that control supernova calibration systematics, directly improve the ability to test dark energy evolution without committing to a parametric form. Likewise, improvements in CMB polarization and lensing that tighten distance priors will narrow the allowed E(z) family and reduce calibration driven ambiguity. Several limitations remain. Any w(z) reconstructed via differentiation inherits noise amplification, so rho_DE(z) is often the more stable diagnostic. Curvature freedom can mimic mild evolution in rho_DE(z) unless constrained externally. Finally, distance priors still rely on a compressed representation of early physics, so genuinely exotic pre recombination scenarios would require a full joint fit rather than a prior based approach. Future work should link reconstruction outputs to theory spaces through emulator based mapping, enabling stable comparisons between reconstructed rho_DE(z) and predictions from specific dark sector models. Another direction is to incorporate full shape clustering and growth information to break degeneracies between Omega_m and dark energy evolution. With those additions, CMB calibrated reconstruction can move from a descriptive tool to a sharper discriminator among physical explanations of cosmic acceleration. The joint likelihood is evaluated in a way that preserves dataset covariances rather than approximating them as independent, because that approximation can bias smooth reconstructions. Edge behavior is controlled by reporting conservative bands outside the redshift window where the data density is high, avoiding overinterpretation of extrapolated trends. A practical diagnostic is the posterior correlation between r_s and the reconstructed E(z) normalization, which quantifies how much of the inference is truly anchored by the CMB. Sensitivity to smoothing is assessed by comparing the implied second derivative of distances, since spurious oscillations often first appear in curvature like combinations. Where chronometer points are included, their influence is monitored through leave one out checks to ensure a small subset of measurements does not dominate local features. The joint likelihood is evaluated in a way that preserves dataset covariances rather than approximating them as independent, because that approximation can bias smooth reconstructions. Edge behavior is controlled by reporting conservative bands outside the redshift window where the data density is high, avoiding overinterpretation of extrapolated trends. A practical diagnostic is the posterior correlation between r_s and the reconstructed E(z) normalization, which quantifies how much of the inference is truly anchored by the CMB. Sensitivity to smoothing is assessed by comparing the implied second derivative of distances, since spurious oscillations often first appear in curvature like combinations. Where chronometer points are included, their influence is monitored through leave one out checks to ensure a small subset of measurements does not dominate local features. The joint likelihood is evaluated in a way that preserves dataset covariances rather than approximating them as independent, because that approximation can bias smooth reconstructions. Edge behavior is controlled by reporting conservative bands outside the redshift window where the data density is high, avoiding overinterpretation of extrapolated trends. A practical diagnostic is the posterior correlation between r_s and the reconstructed E(z) normalization, which quantifies how much of the inference is truly anchored by the CMB. Sensitivity to smoothing is assessed by comparing the implied second derivative of distances, since spurious oscillations often first appear in curvature like combinations. Where chronometer points are included, their influence is monitored through leave one out checks to ensure a small subset of measurements does not dominate local features. The joint likelihood is evaluated in a way that preserves dataset covariances rather than approximating them as independent, because that approximation can bias smooth reconstructions. Edge behavior is controlled by reporting conservative bands outside the redshift window where the data density is high, avoiding overinterpretation of extrapolated trends. A practical diagnostic is the posterior correlation between r_s and the reconstructed E(z) normalization, which quantifies how much of the inference is truly anchored by the CMB. Sensitivity to smoothing is assessed by comparing the implied second derivative of distances, since spurious oscillations often first appear in curvature like combinations. Where chronometer points are included, their influence is monitored through leave one out checks to ensure a small subset of measurements does not dominate local features. The joint likelihood is evaluated in a way that preserves dataset covariances rather than approximating them as independent, because that approximation can bias smooth reconstructions. Edge behavior is controlled by reporting conservative bands outside the redshift window where the data density is high, avoiding overinterpretation of extrapolated trends. A practical diagnostic is the posterior correlation between r_s and the reconstructed E(z) normalization, which quantifies how much of the inference is truly anchored by the CMB. Sensitivity to smoothing is assessed by comparing the implied second derivative of distances, since spurious oscillations often first appear in curvature like combinations. Where chronometer points are included, their influence is monitored through leave one out checks to ensure a small subset of measurements does not dominate local features. The joint likelihood is evaluated in a way that preserves dataset covariances rather than approximating them as independent, because that approximation can bias smooth reconstructions. Edge behavior is controlled by reporting conservative bands outside the redshift window where the data density is high, avoiding overinterpretation of extrapolated trends. A practical diagnostic is the posterior correlation between r_s and the reconstructed E(z) normalization, which quantifies how much of the inference is truly anchored by the CMB. Sensitivity to smoothing is assessed by comparing the implied second derivative of distances, since spurious oscillations often first appear in curvature like combinations. Where chronometer points are included, their influence is monitored through leave one out checks to ensure a small subset of measurements does not dominate local features. The joint likelihood is evaluated in a way that preserves dataset covariances rather than approximating them as independent, because that approximation can bias smooth reconstructions. Edge behavior is controlled by reporting conservative bands outside the redshift window where the data density is high, avoiding overinterpretation of extrapolated trends. A practical diagnostic is the posterior correlation between r_s and the reconstructed E(z) normalization, which quantifies how much of the inference is truly anchored by the CMB. Sensitivity to smoothing is assessed by comparing the implied second derivative of distances, since spurious oscillations often first appear in curvature like combinations. Where chronometer points are included, their influence is monitored through leave one out checks to ensure a small subset of measurements does not dominate local features. The joint likelihood is evaluated in a way that preserves dataset covariances rather than approximating them as independent, because that approximation can bias smooth reconstructions. Edge behavior is controlled by reporting conservative bands outside the redshift window where the data density is high, avoiding overinterpretation of extrapolated trends. A practical diagnostic is the posterior correlation between r_s and the reconstructed E(z) normalization, which quantifies how much of the inference is truly anchored by the CMB. Sensitivity to smoothing is assessed by comparing the implied second derivative of distances, since spurious oscillations often first appear in curvature like combinations. Where chronometer points are included, their influence is monitored through leave one out checks to ensure a small subset of measurements does not dominate local features. The joint likelihood is evaluated in a way that preserves dataset covariances rather than approximating them as independent, because that approximation can bias smooth reconstructions. Edge behavior is controlled by reporting conservative bands outside the redshift window where the data density is high, avoiding overinterpretation of extrapolated trends. A practical diagnostic is the posterior correlation between r_s and the reconstructed E(z) normalization, which quantifies how much of the inference is truly anchored by the CMB. Sensitivity to smoothing is assessed by comparing the implied second derivative of distances, since spurious oscillations often first appear in curvature like combinations. Where chronometer points are included, their influence is monitored through leave one out checks to ensure a small subset of measurements does not dominate local features. The joint likelihood is evaluated in a way that preserves dataset covariances rather than approximating them as independent, because that approximation can bias smooth reconstructions. Edge behavior is controlled by reporting conservative bands outside the redshift window where the data density is high, avoiding overinterpretation of extrapolated trends. A practical diagnostic is the posterior correlation between r_s and the reconstructed E(z) normalization, which quantifies how much of the inference is truly anchored by the CMB. Sensitivity to smoothing is assessed by comparing the implied second derivative of distances, since spurious oscillations often first appear in curvature like combinations. Where chronometer points are included, their influence is monitored through leave one out checks to ensure a small subset of measurements does not dominate local features. The joint likelihood is evaluated in a way that preserves dataset covariances rather than approximating them as independent, because that approximation can bias smooth reconstructions. Edge behavior is controlled by reporting conservative bands outside the redshift window where the data density is high, avoiding overinterpretation of extrapolated trends. A practical diagnostic is the posterior correlation between r_s and the reconstructed E(z) normalization, which quantifies how much of the inference is truly anchored by the CMB. Sensitivity to smoothing is assessed by comparing the implied second derivative of distances, since spurious oscillations often first appear in curvature like combinations. Where chronometer points are included, their influence is monitored through leave one out checks to ensure a small subset of measurements does not dominate local features. The joint likelihood is evaluated in a way that preserves dataset covariances rather than approximating them as independent, because that approximation can bias smooth reconstructions. Edge behavior is controlled by reporting conservative bands outside the redshift window where the data density is high, avoiding overinterpretation of extrapolated trends. A practical diagnostic is the posterior correlation between r_s and the reconstructed E(z) normalization, which quantifies how much of the inference is truly anchored by the CMB. Sensitivity to smoothing is assessed by comparing the implied second derivative of distances, since spurious oscillations often first appear in curvature like combinations. Where chronometer points are included, their influence is monitored through leave one out checks to ensure a small subset of measurements does not dominate local features. The joint likelihood is evaluated in a way that preserves dataset covariances rather than approximating them as independent, because that approximation can bias smooth reconstructions. Edge behavior is controlled by reporting conservative bands outside the redshift window where the data density is high, avoiding overinterpretation of extrapolated trends. A practical diagnostic is the posterior correlation between r_s and the reconstructed E(z) normalization, which quantifies how much of the inference is truly anchored by the CMB. Sensitivity to smoothing is assessed by comparing the implied second derivative of distances, since spurious oscillations often first appear in curvature like combinations. Where chronometer points are included, their influence is monitored through leave one out checks to ensure a small subset of measurements does not dominate local features. The joint likelihood is evaluated in a way that preserves dataset covariances rather than approximating them as independent, because that approximation can bias smooth reconstructions. Edge behavior is controlled by reporting conservative bands outside the redshift window where the data density is high, avoiding overinterpretation of extrapolated trends. A practical diagnostic is the posterior correlation between r_s and the reconstructed E(z) normalization, which quantifies how much of the inference is truly anchored by the CMB. Sensitivity to smoothing is assessed by comparing the implied second derivative of distances, since spurious oscillations often first appear in curvature like combinations. Where chronometer points are included, their influence is monitored through leave one out checks to ensure a small subset of measurements does not dominate local features. The joint likelihood is evaluated in a way that preserves dataset covariances rather than approximating them as independent, because that approximation can bias smooth reconstructions. Edge behavior is controlled by reporting conservative bands outside the redshift window where the data density is high, avoiding overinterpretation of extrapolated trends. A practical diagnostic is the posterior correlation between r_s and the reconstructed E(z) normalization, which quantifies how much of the inference is truly anchored by the CMB. Sensitivity to smoothing is assessed by comparing the implied second derivative of distances, since spurious oscillations often first appear in curvature like combinations. Where chronometer points are included, their influence is monitored through leave one out checks to ensure a small subset of measurements does not dominate local features. The joint likelihood is evaluated in a way that preserves dataset covariances rather than approximating them as independent, because that approximation can bias smooth reconstructions. Edge behavior is controlled by reporting conservative bands outside the redshift window where the data density is high, avoiding overinterpretation of extrapolated trends. A practical diagnostic is the posterior correlation between r_s and the reconstructed E(z) normalization, which quantifies how much of the inference is truly anchored by the CMB. Sensitivity to smoothing is assessed by comparing the implied second derivative of distances, since spurious oscillations often first appear in curvature like combinations. Where chronometer points are included, their influence is monitored through leave one out checks to ensure a small subset of measurements does not dominate local.

References

Aghanim, N., et al. (Planck Collaboration). (2020). Planck 2018 results VI: Cosmological parameters. Astronomy and Astrophysics, 641, A6.

Louis, T., et al. (ACT Collaboration). (2025). The Atacama Cosmology Telescope: DR6 power spectra, likelihoods and LambdaCDM parameters. arXiv:2503.14452.

Adame, A. G., et al. (DESI Collaboration). (2024). DESI 2024 VI: Cosmological constraints from the measurements of baryon acoustic oscillations. arXiv:2404.03002.

Scolnic, D., et al. (2022). The Pantheon+ analysis: The full dataset and light-curve release. arXiv:2112.03863.

Bonilla, A., et al. (2020). Measurements of H0 and reconstruction of the late-time Universe with Gaussian processes. arXiv:2011.07140.

Johnson, J. P., et al. (2025). Kernel dependence of Gaussian process reconstruction of the late Universe expansion history. European Physical Journal C, 85, 14732.

Calderon, R., et al. (2024). Reconstructing dark energy using crossing statistics with DESI BAO. Journal of Cosmology and Astroparticle Physics, 10, 048.

Lodha, K., et al. (2025). DESI constraints on physics-focused aspects of dark energy. Physical Review D, 111, 023532.