

Running of the spectral index: Reconciling the CMB with the Lyman- α forest

Malcolm Fairbairn,^{*} Lucien Heurtier[†],[‡] and María Olalla Olea-Romacho[‡]

*Theoretical Particle Physics and Cosmology, King's College London,
Strand, London WC2R 2LS, United Kingdom*



(Received 30 November 2025; accepted 15 June 2026; published 6 August 2026)

We investigate the scale dependence of the primordial power spectrum by combining *Planck*, ACT DR6, SPT-3G, and eBOSS Lyman- α forest data, extending sensitivity to smaller comoving scales than those probed by the cosmic microwave background (CMB) alone. Within a parametrization based on a Taylor expansion around the pivot scale, we constrain the running of the spectral index α_s and its running β_s . By using eBOSS likelihoods exhibiting a suppression of small-scale power either in amplitude or spectral index, we show that the latter can be accommodated by correlated variations of (α_s, β_s) , leading to a preference for nonzero running. We show that inflationary potentials with localized features—such as Gaussian dips, bumps, or axion-monodromy modulations—can reproduce the inferred scale dependence while remaining compatible with current CMB constraints. We release the public PIPE code to enable systematic tests of inflationary potentials against current CMB datasets.

DOI: 10.1103/pwfl-mt4g

I. INTRODUCTION

Cosmic inflation provides the leading explanation for the initial conditions of our Universe, accounting for its homogeneity and flatness while generating the primordial fluctuations observed in the cosmic microwave background (CMB) and large-scale structure (LSS) [1,2]. Yet, the physical origin of this accelerated expansion—and the detailed shape of the inflaton potential that may drive it—remain unknown.

Since its original proposal [3,4], CMB observatories and LSS surveys have enabled increasingly precise tests of inflation. The beautiful successes of *WMAP* [5] and *Planck* [6] revealed the structure of the CMB fluctuations on small scales, providing evidence for a Universe fully consistent with the Λ CDM hypothesis—a spatially flat Universe containing only baryons, cold dark matter, and a cosmological constant, with a flat or nearly flat dimensionless power spectrum of scalar curvature perturbations that can be modeled with a single power law $\Delta_{\mathcal{R},\Lambda\text{CDM}}^2(k) \equiv A_s(k/k_\star)^{n_s-1}$. Cosmic inflation typically predicts a spectral index n_s close to unity; indeed, *Planck* measured $n_s = 0.9651 \pm 0.0044$ at the pivot scale $k_\star = 0.05 \text{ Mpc}^{-1}$ [6], confirming the success of the Λ CDM

model. In single-field slow-roll inflation, the dynamics of the inflaton—uniquely determined by the shape of the inflaton potential—predict small but generic deviations from this form of the primordial power spectrum (PPS), which can be captured around a pivot scale by Taylor-expanding its logarithm as

$$\ln \Delta_{\mathcal{R}}^2(k) \equiv \ln A_s + (n_s - 1) \ln \left(\frac{k}{k_\star} \right) + \frac{\alpha_s}{2} \ln \left(\frac{k}{k_\star} \right)^2 + \frac{\beta_s}{6} \ln \left(\frac{k}{k_\star} \right)^3 + \cdots \quad (1)$$

where α_s and β_s denote the running and running-of-the-running of the spectral index.

While CMB observations tightly constrain this spectrum around the pivot scale, their sensitivity decreases at smaller scales, limiting their ability to detect mild departures from a power-law spectrum. To probe these smaller scales, we incorporate data from eBOSS [7], which measured $\sim 2 \times 10^5$ quasar spectra with detectable Lyman- α forests.¹

¹Shortly after the completion of this work, a new Lyman- α forest analysis based on DESI DR1 data became available [8], reporting constraints consistent with a primordial spectrum with negligible running. The eBOSS likelihood of Chabanier *et al.* [7] adopted here does not incorporate the latest modeling improvements in Lyman- α analyses and should therefore be regarded as a benchmark scenario exhibiting a suppression of small-scale power. In this sense, our results are best interpreted as characterizing how such a suppression would map onto primordial parameters and inflationary dynamics, rather than as evidence for its presence in current data. Future measurements with improved precision and control of systematics will be crucial to establish whether such scale-dependent features are supported.

^{*}Contact author: malcolm.fairbairn@kcl.ac.uk

[†]Contact author: lucien.heurtier@kcl.ac.uk

[‡]Contact author: maria_olalla.olea_romacho@kcl.ac.uk

These observations probe wave numbers around $k \sim 1 h \text{ Mpc}^{-1}$, providing complementary information on the amplitude and slope of the matter power spectrum in a regime inaccessible to the CMB alone.

Recent CMB measurements from ACT DR6 [9,10] mildly favor a positive running, $\alpha_s > 0$. These experiments probe smaller angular scales than *Planck*, extending CMB sensitivity to higher wave numbers and thereby bridging the gap with the eBOSS Lyman- α forest regime. However, some analyses of the eBOSS data suggested a suppression of small-scale power relative to the nearly scale-invariant spectrum preferred by CMB measurements. In particular, Rogers and Poulin [11] reported a significant deviation from a single power-law primordial spectrum when combining *Planck* with eBOSS data, which can be interpreted as a preference for negative running of the spectral index, therefore pointing in the opposite direction of recent CMB measurements. In this paper, we explore how such results evolve when accounting for new CMB datasets, using different alternative likelihoods [12,13], but also allowing for the running-of-the-running β_s of the spectral index, which was not present in previous analyses.

The likelihood [7] used in Ref. [11] does not reflect the most recent developments in the analysis of eBOSS Lyman- α data, for which updated methodologies have since become standard in the literature [12,13]. In particular, as illustrated in Fig. 1, different eBOSS likelihoods, such as the ones derived in [7,12], compare fairly differently to CMB likelihoods, populating different regions in the $(\Delta_{\text{lin}}^2, n_{\text{lin}})$ plane, where $\Delta_{\text{lin}}^2 \equiv \frac{k^3 P_{\text{lin}}(k_p, z_p)}{2\pi^2}$ and $n_{\text{lin}} \equiv \frac{d \ln P_{\text{lin}}(k, z)}{d \ln k} |_{k_p, z_p}$ denote the amplitude and local slope of the linear matter power spectrum $P_{\text{lin}}(k, z)$, evaluated at $z_p = 3$




FIG. 1. Posteriors for $(\Delta_{\text{lin}}^2, n_{\text{lin}})$ when comparing ACT + P (left panel) and ACT + P + SPT (right panel) to the eBOSS likelihoods from Chabanier *et al.* [7] and Fernandez *et al.* [12], in the Λ CDM model (blue contours), and when including the running (purple contours) and running-of-the-running (teal contours).

and $k_p = 0.009 \text{ km}^{-1}$ (see Sec. II for details). Indeed, whereas the likelihood [12] favors a strong suppression of the spectrum amplitude only, the one from Ref. [7] points toward a milder suppression of the amplitude but a clear suppression of the tilt.

In this work, we assume that these deviations originate from primordial physics [14] and perform a joint analysis of *Planck*, ACT DR6, SPT-3G [15], and eBOSS data using both the likelihoods from Ref. [7] and Ref. [12]. We use these likelihoods as representative benchmarks of what could be potential suppressions in the matter power spectrum shape to study how this type of suppression can be translated into correlated variations of α_s and β_s and to identify which inflationary features can be described within this framework. As we will see in the next sections, the Taylor-expanded parametrization used in Eq. (1) only captures a restricted class of smooth distortions of the primordial power spectrum and will appear to be better suited to describe deviations of the type associated with the likelihood of Ref. [7]. We then test representative scenarios—including monomial, axion-monodromy, and potentials with localized features—and determine which can reproduce the inferred small-scale behavior while remaining consistent with CMB constraints. We also release the public PIPE code to evaluate arbitrary inflationary potentials against current CMB data.

II. POWER SPECTRUM RUNNING

A. Likelihood

To constrain the PPS, we use the most recent likelihoods from *Planck* [6] and the ACT DR6 Lite likelihood released by the ACT collaboration [9,10].² Our setup follows the *Planck*-ACT (P-ACT) cut configuration described in Ref. [9], which combines large-scale *Planck* data with high-resolution ACT measurements. Specifically, we include the low- ℓ TT spectrum from *Planck* [16] and the low- ℓ EE spectrum from the S_{ROLL2} reanalysis [17], while at high multipoles we use ACT DR6 CMB-only data in the *Planck*-ACT cut, retaining *Planck* data only up to $\ell = 1000$ in TT and $\ell = 600$ in TE/EE. ACT data dominate at smaller scales, covering the range $600 < \ell < 6500$ in TT, TE, and EE, with foregrounds marginalized over. We also include the foreground-marginalized likelihood SPT-3G D1 Lite published recently by SPT [15].³ Similarly to Ref. [11], we incorporate Lyman- α forest constraints on the linear matter power spectrum using a compressed two-dimensional Gaussian likelihood constructed from the eBOSS (Sloan Digital Sky Survey Data Release 14; SDSS DR14) Lyman- α forest flux power spectrum [7]. In this approach, the cosmological information contained in the flux power spectrum is projected onto the two effective parameters: the amplitude Δ_{lin}^2 and local spectral tilt n_{lin} of

²<https://github.com/ACTCollaboration/DR6-ACT-lite>.

³https://github.com/SouthPoleTelescope/spt_candl_data.

the linear matter power spectrum $P_{\text{lin}}(k, z)^4$, evaluated at the pivot redshift $z_p = 3$ and wave number $k_p = 0.009 \text{ h Mpc}^{-1}$. We first consider the compressed likelihood derived in Ref. [18], which is based on the eBOSS DR14 Lyman- α analysis of Ref. [7] and based on the BOSS/eBOSS Lyman- α forest flux power spectrum measurements reported in Ref. [19]. In this case, the mean values of Δ_{lin}^2 and n_{lin} are 0.310 and -2.340 , with standard deviations 0.020 and 0.006, respectively, and a correlation coefficient $\rho = 0.512$ [18].

In addition, we also consider the alternative Lyman- α likelihood proposed in Ref. [12], which adopts a more conservative treatment of astrophysical and modeling uncertainties affecting the inference of the linear matter power spectrum. For consistency with the analysis described above, we also approximate this likelihood by a two-dimensional Gaussian in the $(\Delta_{\text{lin}}^2, n_{\text{lin}})$ plane. In this Gaussian approximation, the mean values are taken to be $\Delta_{\text{lin}}^2 = 0.267^{+0.018}_{-0.023}$ and $n_{\text{lin}} = -2.288^{+0.020}_{-0.020}$, with a correlation coefficient of 0.4. These values are chosen to reproduce the central values, uncertainties, and degeneracy directions reported in Ref. [12].

B. Monte Carlo Markov chain

We perform Monte Carlo Markov chain (MCMC) sampling using the publicly available code COBAYA [20] interfaced with the Boltzmann solver CAMB [21] to extract constraints on the amplitude and shape of the PPS given in Eq. (1). We perform two separate analyses: one where we sample the scalar amplitude $\ln(10^{10} A_s)$ and spectral index n_s at the pivot scale $k_* = 0.05 \text{ h Mpc}^{-1}$, together with the running α_s , and one in which we also include the running-of-the-running β_s . These parameters are varied jointly with the remaining four baseline ΛCDM parameters: the optical depth to reionization τ , the physical baryon density ω_b , the physical cold dark matter density ω_c , and the angular size of the sound horizon θ_s . In addition, we vary the nuisance parameters T_{cal} and E_{cal} , which account for the absolute calibration of the temperature and polarization spectra, respectively. For all parameters, we adopt the same flat priors as used by the ACT collaboration in the ACT-DR6.02 “nrun” analysis.⁵ To ensure convergence, we run multiple chains (using 12 CPUs when including only α_s in the analysis and 15 CPUs when also including β_s) in parallel and monitor the Gelman-Rubin R statistic, requiring $R - 1 < 0.05$ for all sampled parameters.

C. Results

Figure 1 shows the regions favored by ACT + P + SPT for ΛCDM (blue), running (α_s , purple), and running-of-the-running (α_s, β_s , teal), compared to those preferred by

⁴We show in Appendix A that this approach remains valid across scales for the cases considered throughout this paper.

⁵https://lambda.gsfc.nasa.gov/product/act/act_dr6.02/act_dr6.02_chains_nrun_get.html.

eBOSS Lyman- α analyses. The red contours correspond to the likelihood of Chabanier *et al.* [7] used in this work, while the yellow contours show an alternative analysis based on the PRIYA simulations [12], which rely on a more up-to-date modeling of the Lyman- α forest and yield $\Delta_{\text{lin}}^2 = 0.267^{+0.018}_{-0.023}$ and $n_{\text{lin}} = -2.288^{+0.020}_{-0.020}$, with correlation coefficient 0.4. In Appendixes D and E, we also present the full cosmological parameter inference and the corresponding results for σ_8 in comparison with recent observations.

While both analyses exhibit a tension with ΛCDM , they correspond to different types of deviations: the Chabanier *et al.* likelihood is primarily driven by a change in the local slope of the spectrum at the eBOSS scale, whereas the Fernandez *et al.* analysis is dominated by a shift in amplitude, with a tilt closer to the CMB-preferred value.

Allowing for running and running-of-the-running induces correlated shifts in $(\Delta_{\text{lin}}^2, n_{\text{lin}})$ that can partially alleviate the tension with the Chabanier *et al.* likelihood but cannot reach the region preferred by the Fernandez *et al.* analysis. This reflects the limited flexibility of the Taylor-expanded parametrization, which only captures smooth, correlated distortions of the primordial spectrum.

Figure 2 shows the constraints obtained when allowing only the running of the spectral index. Figure 3 extends to the running-of-the-running, β_s , showing 68% and 95% credible regions from Planck+ACT DR6 (yellow) and +SPT-3G (blue). Upon including eBOSS Ly- α data, the credible regions obtained using the Chabanier *et al.* (purple) [7] and Fernandez *et al.* (teal) [12] likelihoods are shown for comparison. In both cases, adding eBOSS data tightens the bounds, but it is remarkable that the corresponding favored regions for the spectral index, its running and running-of-the-running are fairly different.




FIG. 2. Joint and marginalized posterior probability distributions for n_s and its running $\alpha_s = dn_s/d \ln k$, derived from the three likelihood combinations analyzed in this work. Shaded contours represent the 68% and 95% credible regions. The dashed gray curve corresponds to the ACT + P constraint from Ref. [10], obtained when fitting only α_s .




FIG. 3. Same as FIG. 2 but including $\beta_s = d^2 n_s / d \ln k^2$. Median values and corresponding 68% credible intervals for these parameters are reported in the Appendix F.

The Chabanier *et al.* likelihood shift both n_s and α_s ; the data disfavor $\alpha_s = 0$ at $> 2\sigma$, already hinting at higher-order scale dependence. The inclusion of eBOSS sharply contracts the contours and drives them along the α_s - β_s toward more negative values, with $\beta_s = 0$ disfavored at about 2σ . The corresponding shift of n_s toward larger values suggests a slightly bluer tilt once small-scale Ly- α data are included. On the other hand, the Fernandez *et al.* likelihood appears consistent with $\alpha_s = \beta_s = 0$ in the corner plots, despite favoring a region in the $(\Delta_{\text{lin}}^2, n_{\text{lin}})$ plane that is clearly offset from the CMB contours (see Fig. 1). This does not indicate that zero running is preferred. Rather, it reflects the limited flexibility of the Taylor-expanded parametrization: the type of deviation favored by likelihood in Ref. [12] cannot be reproduced by smooth variations controlled by (α_s, β_s) . As a result, the fit remains close to the Λ CDM region in the (n_s, α_s, β_s) space. In practice, reproducing this behavior would require either additional parameters in the running expansion or a sharper feature in the primordial power spectrum.

Since our approach is able to reconcile CMB measurements primarily with the likelihood from Ref. [7], we will use the latter as an illustrative reference case in what follows. Although more recent analyses rely on improved modeling [12,13], this likelihood provides a concrete realization of a slope-driven small-scale suppression that can be consistently mapped onto correlated variations of (α_s, β_s) within the Taylor-expanded framework.

Using posteriors in the $(\Delta_{\text{lin}}^2, n_{\text{lin}})$ plane (Fig. 1), we quantify the tension between CMB data and the eBOSS Lyman- α constraint of Chabanier *et al.* with the

TENSIOMETER package [22] (see Appendix F for details). The combined ACT + P + SPT data increase the CMB-eBOSS discrepancy beyond the previously reported 4.9σ for *Planck* alone [11]. Allowing for running reduces this to ~ 3 – 3.3σ (vs $\sim 1\sigma$ with *Planck* only), and including β_s further lowers it to ~ 2.8 – 3σ . In Fig. 3, this decreasing tension manifests as competing preferences: *Planck* yields a mildly negative central value of α_s , and ACT DR6 favors $\alpha_s > 0$, whereas the eBOSS likelihood requires a strong suppression of small-scale power, corresponding to $\alpha_s < 0$. Consequently, CMB data collectively prefer $\alpha_s > 0$ and $\beta_s > 0$, while the eBOSS constraint drives both toward negative values. Allowing for $\beta_s < 0$ partially alleviates the discrepancy by steepening the small-scale suppression but does not recover the positive running preferred by the CMB.

III. IMPLICATIONS FOR INFLATION

Recent constraints favoring higher values of the scalar spectral index have revived interest in inflationary scenarios beyond canonical plateau potentials. Models such as exponential α -attractors [23–25], Starobinsky [26], and Higgs inflation [27] typically struggle to accommodate $n_s \gtrsim 0.97$ [28–32]. This trend motivates the exploration of alternatives that favor potentials approaching the plateau polynomially rather than exponentially [28]. We therefore consider several single-field potentials, beginning with a monomial baseline $V_\alpha = V_0 |\phi/M_p|^\alpha$, where α controls the slope and $\alpha < 1$ is preferred by recent data [10]. We then introduce localized or oscillatory deviations to capture transient dynamics.

A. Gaussian bumps and dips

Transient features—e.g. from localized particle production or modulations [33–36]—can modify the dynamics of $\dot{\phi}$, impact CMB observables, or even trigger phases of ultraslow roll leading to primordial black hole formation [37]. We model such effects by modulating the baseline potential with

$$V = V_\alpha \left[1 \pm A \exp \left(-\frac{(\phi - \phi_0)^2}{2\sigma^2} \right) \right], \quad (2)$$

where the + sign produces a bump and the – a dip. The parameters A , ϕ_0 , and σ set the amplitude, position, and width, controlling the induced running and feature scale.

B. Axion monodromy

Inspired by axionlike inflation in string theory [38,39], monodromy potentials also include periodic modulations on a monomial background:

$$V = V_0 \left[\left| \frac{\phi}{M_p} \right|^\alpha + A \cos \left(\gamma + \frac{M_p}{f} \left| \frac{\phi}{M_p} \right|^{1+p_f} \right) \right], \quad (3)$$

with amplitude A , phase γ , axion decay constant f , and index p_f . These oscillations naturally generate small-scale running and resonant signatures in the PPS.

C. Inflation observables

For each benchmark model and parameter set, we computed the PPS amplitude A_s , spectral index n_s , tensor-to-scalar ratio r , and scale dependence parameters α_s and β_s by integrating the background equations numerically. The slow-roll parameters were evaluated analytically from the potential and its derivatives at horizon exit to leading order in the slow-roll expansion [40]. The PPS was then reconstructed on all relevant scales using the slow-roll approximation. We restrict to parameter choices for which the Taylor expansion of Eq. (1) provides a reliable description of the primordial power spectrum, requiring percent-level agreement with the full numerical spectrum at the eBOSS scale (see Appendix B for details).

D. Likelihood

To compare model predictions with CMB and LSS data, we build a custom likelihood over the effective PPS parameters $\{A_s, n_s, \alpha_s, \beta_s\}$. Posterior samples from the MCMC chains are processed with GetDist [41] and approximated by a multidimensional Gaussian kernel density estimator (KDE), defining a smooth empirical likelihood $\ln \mathcal{L}_{\text{KDE}}(A_s, n_s, \alpha_s, \beta_s)$. These procedures are implemented in our public code PIPE (potential inflation posterior emulator), which provides a fast surrogate for the full CMB and eBOSS likelihoods (see Appendix C).

Although the analysis focuses on the scalar spectrum, PIPE also enforces consistency with current bounds on the tensor-to-scalar ratio by adding a BK18-based prior $\ln \pi_r$ [42] to the log likelihood.

E. Best-fit search

For each dataset, we maximized the PIPE log-posterior $\ln \mathcal{P} \equiv \ln \mathcal{L}_{\text{KDE}} + \ln \pi_r$ under broad, uninformative priors

TABLE I. Best-fit parameters, KDE log-likelihood ($\ln \mathcal{L}_{\text{KDE}}$), and tensor prior contribution ($\ln \pi_r$) for each dataset/model/feature combination from the updated PIPE runs. The total posterior is $\ln \mathcal{L} = \ln \mathcal{L}_{\text{KDE}} + \ln \pi_r$; the Taylor penalty is zero for all listed best fits.

Dataset	Model	$\log_{10} V_0$	α	$\log_{10} A$	$\phi_0[M_{\text{Pl}}]$	$\sigma[M_{\text{Pl}}]$	γ	$\log_{10} f$	p_f	$\ln \mathcal{L}_{\text{KDE}}$	$\ln \pi_r$	$\ln \mathcal{L}$
ACT + P	Power law	-9.3565	0.3935	...	...	...	...	...	...	32.032	2.835	34.867
	Power law + bump	-9.5142	0.1698	-3.8106	-4.287	0.161	...	...	...	34.024	3.603	37.626
	Power law + dip	-9.5350	0.1951	-2.3266	-5.594	0.599	...	...	...	34.024	3.603	37.626
	Axion monodromy	-9.5085	0.1662	-4.2186	...	...	7.416	-1.053	0.0053	34.051	3.603	37.654
+SPT	Power law	-9.3829	0.3067	...	...	...	...	...	...	31.730	3.334	35.065
	Power law + bump	-9.4408	0.0826	-1.5910	-4.096	0.938	...	...	...	35.099	3.601	38.699
	Power law + dip	-9.3917	0.0316	-1.3712	-1.071	0.988	...	...	...	35.099	3.601	38.699
	Axion monodromy	-9.5012	0.1566	-3.2910	...	...	0.492	0.077	0.829	35.099	3.603	38.702
+eBOSS	Power law	-10.2659	1.6448	...	...	...	...	...	...	6.859	-19.494	-12.636
	Power law + bump	-10.2095	0.8137	-1.0900	-8.282	0.964	...	...	...	37.305	3.603	40.908
	Power law + dip	-9.8199	0.5068	-1.3842	-7.776	0.774	...	...	...	37.305	3.603	40.908
	Axion monodromy	-9.6943	0.3672	-1.7747	...	...	24.70	-0.372	-0.0060	37.305	3.603	40.908

TABLE II. Verified spectral parameters from the best-fit potentials. All quantities are evaluated at the pivot scale $k_* = 0.05 \text{ Mpc}^{-1}$.

Dataset	Model	A_s	r	n_s	α_s	β_s
ACT + P	Power law	2.1367×10^{-9}	0.0291	0.9779	-4.09×10^{-4}	-8.81×10^{-6}
	Power law + Gaussian bump	2.1109×10^{-9}	0.0125	0.9680	-4.15×10^{-4}	-6.46×10^{-6}
	Power law + Gaussian dip	2.1110×10^{-9}	0.0125	0.9680	-4.12×10^{-4}	-6.39×10^{-6}
	Axion monodromy	2.1235×10^{-9}	0.0125	0.9713	$+3.70 \times 10^{-3}$	-5.13×10^{-4}
ACT + P + SPT	Power law	2.1092×10^{-9}	0.0227	0.9787	-3.94×10^{-4}	-8.27×10^{-6}
	Power law + Gaussian bump	2.1209×10^{-9}	0.0130	0.9681	$+2.94 \times 10^{-3}$	-2.92×10^{-4}
	Power law + Gaussian dip	2.1209×10^{-9}	0.0130	0.9681	$+2.96 \times 10^{-3}$	-2.96×10^{-4}
	Axion monodromy	2.1209×10^{-9}	0.0125	0.9681	$+2.96 \times 10^{-3}$	-2.95×10^{-4}
ACT + P + SPT + eBOSS	Power law	2.1547×10^{-9}	0.1210	0.9665	-6.16×10^{-4}	-1.65×10^{-5}
	Power law + Gaussian bump	2.1293×10^{-9}	0.0125	0.9693	-8.66×10^{-3}	-2.56×10^{-3}
	Power law + Gaussian dip	2.1293×10^{-9}	0.0125	0.9693	-8.66×10^{-3}	-2.56×10^{-3}
	Axion monodromy	2.1293×10^{-9}	0.0125	0.9693	-8.66×10^{-3}	-2.56×10^{-3}

TABLE III. Statistical deviation of the best-fit feature and axion-monodromy models relative to the baseline power-law potential for each dataset combination. Negative ΔAIC values indicate statistical preference over the baseline model. The likelihood difference is defined as $\Delta\chi^2 = -2(\ln \mathcal{L}_{\text{model}} - \ln \mathcal{L}_{\text{Power law}})$.

Dataset	Model	$\Delta\chi^2$	ΔAIC
ACT + P	Power law + bump	-7.27	-1.27
	Power law + dip	-7.27	-1.27
	Axion monodromy	-7.27	0.73
+SPT	Power law + bump	-5.52	0.48
	Power law + dip	-5.52	0.48
	Axion monodromy	-5.57	2.43
+eBOSS	Power law + bump	-107.09	-101.09
	Power law + dip	-107.09	-101.09
	Axion monodromy	-107.09	-99.09

on the model parameters. The optimization combined a genetic algorithm (EVORTAN [43]) with a final Nelder-Mead refinement. For every dataset-model-feature combination we repeated the GA with different random seeds and retained the best posterior value. This procedure yields robust potential constraints and enables consistency checks across the P – ACT, P – ACT + SPT, and P – ACT + SPT + eBOSS combinations.

F. Results

The best-fit potentials and spectral parameters are listed in Tables I and II. Because multiple parameter combinations yield nearly identical effective spectra ($A_s, n_s, \alpha_s, \beta_s$), the KDE posterior exhibits flat regions with equivalent




FIG. 4. Relative modulation of the potential $(V - V_\alpha)/V_\alpha$, as compared to the baseline model $V_\alpha \equiv V_0|\phi/M_p|^\alpha$, for each set of best-fit parameters ($V_0, \alpha, \dots$) listed in Table I. Vertical, colored, dashed lines denote the field value corresponding to where perturbations exit the horizon at the eBOSS scale.

maxima; we therefore quote a representative best fit per model and dataset. Likelihood-ratio tests relative to the baseline power law give the $\Delta\chi^2$ and ΔAIC values in Table III. For CMB data alone (ACT + P and ACT + P + SPT), the extensions provide only mild improvements, with $|\Delta\chi^2| \simeq 5\text{--}7$ and small $|\Delta\text{AIC}| \lesssim 2$. When eBOSS data are included, the power-law potential becomes strongly disfavored, while all featureful and axion-monodromy models achieve comparable maximum posteriors around $\ln \mathcal{L} \simeq 40.9$. This degeneracy reflects the fact that different potentials can produce nearly identical primordial spectra over the range of scales probed. As illustrated in Fig. 4, the corresponding best-fit features induce different localized modulations of the potential around the field value associated with the eBOSS scale, yet lead to equivalent effective parameters ($A_s, n_s, \alpha_s, \beta_s$). Within the Taylor-expanded parametrization, these scenarios are therefore observationally indistinguishable.

IV. SUMMARY AND DISCUSSION

In this work, we have investigated the scale dependence of the PPS by combining CMB measurements from *Planck*, ACT DR6, and SPT-3G with Lyman- α forest constraints from eBOSS. Extending the analysis beyond the scales probed by the CMB alone allows us to test departures from a simple power-law PPS and to assess their interpretation within inflationary models.

Using a Taylor expansion of $\ln \Delta_R^2(k)$ around a pivot scale, we parametrized deviations from scale invariance in terms of the running α_s and its running β_s . We showed that, within this framework, correlated variations of (α_s, β_s) shift the CMB-preferred region in the $(\Delta_{\text{lin}}^2, n_{\text{lin}})$ plane along a well-defined direction. This direction aligns with the type of deviation associated with the eBOSS likelihood of Chabanier *et al.*, which is characterized by a suppression of small-scale power driven primarily by a change in the local spectral slope.

When adopting this likelihood as a benchmark, we find that allowing for running partially alleviates the tension between CMB and Lyman- α data, reducing it from $\sim 5\sigma$ to $\sim 3\sigma$, and further to $\sim 2.8\text{--}3\sigma$ when including β_s . This improvement reflects the ability of the Taylor-expanded parametrization to capture smooth, slope-driven distortions of the PPS. However, the tension is not fully resolved, as the CMB data favor mildly positive running, while the eBOSS likelihood requires a strong suppression at small scales, corresponding to negative α_s and β_s .

A key result of this work is that not all Lyman- α inferences can be mapped onto this parametrization. In particular, more recent analyses, such as those based on the PRIYA simulations, favor deviations that are predominantly amplitude driven rather than slope driven. These cannot be reproduced within the limited flexibility of the Taylor expansion, and therefore do not translate into a preference for nonzero (α_s, β_s) in parameter space. This highlights that




FIG. 5. Schematic inflaton potential (left panel) featuring both a smooth feature, which can be captured by the parametrization used in this paper, and a sharp feature, which lies beyond its range of validity. The corresponding primordial power spectrum is shown in the right panel. Smooth enough features lead to a smooth running that can be well approximated by the Taylor expansion of Eq. (1), whereas sharper features can induce more pronounced variations that cannot be fully captured within this framework and are left for future work.

the running parametrization probes only a restricted class of smooth deformations of the PPS. Indeed, as it is represented in Fig. 5, only inflaton potentials with smooth enough features lead to primordial power spectra that can accurately be approximated by the Taylor expansion of Eq. (1). In the contrary case of sharp features, leading to more pronounced variations of the primordial power spectrum, one would need either to push the Taylor expansion to higher and higher orders or to perform statistical analyses at the level of the potential parameters directly while solving the Mukhanov-Sasaki equations for each set of parameters consistently. Doing so would go beyond the scope of this paper, which we leave for future work.

Motivated by these findings, we explored the inflationary origin of such deviations by reconstructing the PPS from explicit single-field potentials. We considered a monomial baseline supplemented by localized features (Gaussian bumps and dips) and oscillatory modulations (axion monodromy). We found that while CMB data alone mildly favor featureful extensions over a pure power-law potential, the statistical significance remains limited. In contrast, when including the eBOSS (Chabanier *et al.*) likelihood, the power-law potential becomes strongly disfavored, and models with features or oscillations provide a substantially better fit.

Despite this improvement, we observe a pronounced degeneracy at the level of inflationary potentials: different microphysical realizations—bumps, dips, or oscillatory modulations—lead to nearly identical effective parameters ($A_s, n_s, \alpha_s, \beta_s$) and therefore indistinguishable observational signatures within the Taylor-expanded framework.

These results should be interpreted with care. The eBOSS likelihood adopted here does not incorporate the

most recent modeling developments, and newer analyses, including DESI, are consistent with negligible running within current uncertainties. Our findings therefore do not constitute evidence for a departure from Λ CDM, but rather illustrate how a class of small-scale suppressions—if confirmed by future measurements—could be consistently mapped onto primordial parameters and inflationary dynamics.

More broadly, this work clarifies the connection between late-time observables and inflationary physics in the presence of scale-dependent features, and upcoming Lyman- α measurements will determine whether the hints of small-scale structure suppression seen with eBOSS persist. If confirmed, they would provide a rare window into inflationary dynamics beyond the slow-roll paradigm; if not, they will further reinforce the robustness of the minimal Λ CDM description across an extended range of scales.

ACKNOWLEDGMENTS

The authors would like to thank Hengameh Bagherian, Jo Dunkley, and Mudit Jain for helpful discussions throughout this project, as well as Dobby Marsh and Keir Rogers for valuable comments on our manuscript. They are also grateful to Thomas Biekötter for his assistance with the EVORTRAN code. This work was supported by the STFC under UKRI Grant No. ST/X000753/1. L.H. also thanks the Institute for Particle Physics Phenomenology at Durham University for access to its high-performance computing facilities.

DATA AVAILABILITY

The software that supports the findings of this article is openly available at the following address: [44].




FIG. 6. Linear matter power spectrum, rescaled both in amplitude and scale compared to the Λ CDM best fit from *Planck* [6] at the eBOSS scale, and compared to the eBOSS range of frequencies and credible region (gray-shaded area) as defined in Ref. [45].

APPENDIX A: VALIDITY OF THE eBOSS LIKELIHOOD

To assess the validity of the compressed eBOSS likelihood we used in the analysis, we checked that the linear matter power spectrum obtained remains sufficiently close to a power law within the 1σ validity band of the eBOSS measurement. To do so, we used a similar approach as in Ref. [11], by taking the best-fit linear matter power spectra for each dataset combination and rescaling them such that they have the same values of $\Delta_{\text{lin}}^2, n_{\text{lin}}$ as the best-fit Λ CDM model from *Planck* [6] at the eBOSS scale. Results are depicted in Fig. 6, where one can see that each best-fit point leads to a linear matter power spectrum that is contained within the eBOSS range of frequencies and credible region, as calculated in Ref. [45] (gray-shaded area), meaning that the linear matter power spectrum obtained is sufficiently power-law-like for the eBOSS compression to remain valid over the appropriate range of scales.

APPENDIX B: TAYLOR EXPANSION

When confronting power spectra obtained from slow-roll inflation potentials, we used the Taylor expansion of Eq. (1) at the pivot scale and assumed that this expansion is valid up to the eBOSS scale $k \simeq 1 \text{ h Mpc}^{-1}$, so we can use the constraints on $\{A_s, n_s, \alpha_s, \beta_s\}$ obtained from CMB and Ly- α likelihoods to constrain potential parameters. To assess the validity of this assumption, we calculated the PPS using the slow-roll approximation numerically and compared its value at the eBOSS scale. We imposed the relative difference to be within $\delta \equiv |\Delta_{\mathcal{R}, \text{Taylor}}^2 - \Delta_{\mathcal{R}}^2| / \Delta_{\mathcal{R}}^2 < \delta_{\text{thresh}}$, where we chose $\delta_{\text{thresh}} = 1\%$ so that the corresponding error made on the linear matter power spectrum at the eBOSS scale lies

within the 1σ sensitivity bound of the eBOSS measurement that is of order $\mathcal{O}(10)\%$ [45]. To enforce this condition while guiding the genetic algorithm to the right region of the parameter space, we added to the likelihood an extra contribution

$$\Delta\mathcal{L}_{\text{Taylor}} = \begin{cases} \left(\frac{\delta - \delta_{\text{thresh}}}{\delta_{\text{thresh}}}\right)^p & \text{if } \delta > \delta_{\text{thresh}} \\ 0 & \text{if } \delta < \delta_{\text{thresh}} \end{cases}. \quad (\text{B1})$$

Furthermore, because the eBOSS data span a larger range of frequencies, we also demanded that this error does not exceed 3% at $k_{\text{eBOSS, max}} \approx 3.5 \text{ h Mpc}^{-1}$. We checked that every best fit found by the algorithm lies within the region where $\Delta\mathcal{L}_{\text{Taylor}} = 0$, so that this guiding penalty does not contribute to the value of the likelihood. In our analysis, we used $p = 2$, although our code allows for choosing arbitrary values of $p, \delta_{\text{thresh}}$, and the possibility of adding an additional prefactor to the penalty.

APPENDIX C: PIPE: POTENTIAL INFLATION POSTERIOR EMULATOR

To streamline the comparison between theoretical predictions and observational constraints, we developed a public PYTHON package called PIPE, which can be found in

$$\text{gitlab.com/cosmoPipe/Pipe-inflation}. \quad (\text{C1})$$

PIPE provides a fast, data-driven likelihood for inflationary models by emulating the joint posterior of effective spectral parameters $(A_s, n_s, \alpha_s, \beta_s)$ from the P-ACT, P-ACT + SPT, and P-ACT + SPT + eBOSS datasets. The code reads in GetDist MCMC chains of effective parameters and constructs a smoothed empirical likelihood via a Gaussian KDE with a bandwidth tuned for each dataset according to Scott's rule. It then allows the user to specify either built-in or fully custom inflationary potentials; for each parameter point θ , PIPE evolves the background equations and returns the predicted spectral parameters $[A_s, r, n_s, \alpha_s, \beta_s]$ at leading order in slow roll, following Ref. [40]. These predictions are directly evaluated against the KDE likelihood to obtain $\ln \mathcal{L}_{\text{KDE}}$, optionally combined with an r -prior $\ln \pi_r$ from BK18 via `LogL_r.py`, yielding the total log-posterior $\ln \mathcal{L}_{\text{KDE}} + \ln \pi_r$. A Taylor-expansion penalty is also implemented (as described in Appendix B).

Internally, PIPE consists of

- (1) `pipe/Inflation_evo.py`, which computes $[A_s, r, n_s, \alpha_s, \beta_s]$ for supported and user-provided potentials.
- (2) `pipe/LogL_KDE.py`, which constructs the KDE likelihood directly from GetDist chains.
- (3) `pipe/LogL_r.py`, which configures the BK18 r -prior [42].




FIG. 7. Comparison for the P-ACT + SPT + eBOSS dataset of the 1D and 2D posterior distributions for the effective inflationary parameters $\{A_s, n_s, \alpha_s, \beta_s\}$ obtained directly from the GetDist chains (black outlines) and from a Gaussian KDE resampling (purple fill) using our public PIPE code. The band width has been set according to the Scott's rule.

- (4) `example/example.py`, an end-to-end demonstration showing how to use the built-in Powerlaw model with the same custom potential and compute $\ln \mathcal{L}_{\text{KDE}}$ all the contributions to the likelihood.

Built-in templates include Powerlaw (with optional Gaussian dip or bump features) and Monodromy; new models can be added by providing the potential and its first four derivatives.

Figure 7 shows, for the P-ACT+EPT+eBOSS dataset, a direct comparison between the raw GetDist chains (black outlines) and their KDE approximation (purple fill) produced with PIPE, illustrating how the KDE accurately reproduces the full MCMC posterior used in our analysis.

APPENDIX D: EFFECT ON σ_8

We survey here to which extent the running of the PPS's spectral index favored by the CMB datasets considered affect the value of σ_8 and compare it to the existing constraints from KiDS Legacy [46]. As one can see from Fig. 8, inclusion of the running and running-of-the-running affect very mildly the σ_8 and remains consistent with current data, with SPT asking for a slightly larger σ_8 than ACT. The discrepancy remains between the CMB values and the other weak lensing surveys, but the tension is not significantly altered when we include the running considered in this work.

APPENDIX E: FULL COSMOLOGICAL RESULTS

We report in Figs. 9 and 10 the full 1D and 2D cosmological posteriors obtained when including either α_s or β_s in our analysis.




FIG. 8. Posteriors for Ω_m, S_8 when including the running (left panel) and running-of-the-running (right panel) of the spectral index. For comparison, green contours depict the credible region, as preferred by KiDS Legacy [46].




FIG. 9. Full cosmological 1D and 2D posteriors obtained when including the running of the spectral index α_s .




FIG. 10. Full cosmological 1D and 2D posteriors obtained when including the running of the spectral index α_s , and the running-of-the-running β_s .

APPENDIX F: PRIMORDIAL POWER SPECTRUM MEDIAN VALUES AND STATISTICAL TENSIONS WITH Λ CDM

The median values and 68% credible intervals (16th and 84th percentiles) for n_s , α_s , and β_s obtained from different dataset combinations are listed in Table IV.

For the case of Λ CDM, and when including either the running alone (α_s) or both the running and running-of-the-running (β_s) of the spectral index, we used the package

TABLE IV. Median values and 68% credible intervals (16th and 84th percentiles) for n_s , α_s , and β_s in three data combinations: ACT+ *Planck* (ACT + P), with added SPT data (+SPT), and with both SPT and eBOSS data (+eBOSS).

Dataset	n_s	α_s	β_s
ACT + P	$0.96501^{+0.00490}_{-0.00493}$	$0.01195^{+0.00623}_{-0.00628}$	$0.01869^{+0.00954}_{-0.00967}$
+SPT	$0.96437^{+0.00453}_{-0.00449}$	$0.00804^{+0.00567}_{-0.00571}$	$0.01477^{+0.00884}_{-0.00876}$
+eBOSS	$0.97101^{+0.00391}_{-0.00388}$	$-0.00323^{+0.00390}_{-0.00388}$	$-0.00755^{+0.00346}_{-0.00347}$

TENSIOMETER [22] to estimate the statistical tension between CMB observatories and eBOSS measurements. This package allows us to calculate the statistical tension between distributions that are not necessarily Gaussian, which is exactly what is needed in our case. In each case, we tested that the result obtained is stable when varying internal settings used by TENSIOMETER, such as the smoothing scale used in the estimate of the KDE shift—integrated square error (MISE), adaptive bandwidth, asymptotic MISE estimator, or maximum bandwidth (MAX)—or the *boost* sampling parameter, which we varied to larger and larger values until obtaining a stable result. Our results are presented in Table V, in which, for each entry, a subset of CMB likelihoods (row entries) is tensioned against the eBOSS likelihood from Ref. [7] given the model chosen

TABLE V. Statistical tension between different CMB dataset combinations: *Planck*+ACT (P+ACT), with added SPT data (+SPT), and eBOSS data (+eBOSS).

Dataset	Λ CDM	$+\alpha_s$	$+(\alpha_s, \beta_s)$
P+ACT	$\gg 5\sigma$	$3.375^{+0.002}_{-0.002}\sigma$	$3.074^{+0.001}_{-0.001}\sigma$
+SPT	$\gg 5\sigma$	$3.038^{+0.001}_{-0.001}\sigma$	$2.813^{+0.001}_{-0.001}\sigma$

(column entries). Note that, in the case of Λ CDM, TENSIOMETER failed at returning either a stable or finite result, due to the too small overlap between the CMB and eBOSS distributions. Some finite results returned were above 5.7σ , which allowed us to safely claim a tension larger than 5σ both for ACT + P and ACT + P + SPT.

-
- [1] A. R. Liddle and D. H. Lyth, *Phys. Rep.* **231**, 1 (1993).
 - [2] D. Baumann, *Cosmology* (Cambridge University Press, Cambridge, England, 2022).
 - [3] A. A. Starobinsky, *JETP Lett.* **30**, 682 (1979).
 - [4] A. H. Guth, *Phys. Rev. D* **23**, 347 (1981).
 - [5] E. Komatsu *et al.* (WMAP Collaboration), *Astrophys. J. Suppl. Ser.* **192**, 18 (2011).
 - [6] N. Aghanim *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A6 (2020); **652**, C4(E) (2021).
 - [7] S. Chabanier *et al.* (eBOSS Collaboration), *J. Cosmol. Astropart. Phys.* 07 (2019) 017.
 - [8] J. Chaves-Montero *et al.* (2026), 10.5281/zenodo.16943722.
 - [9] T. Louis *et al.* (ACT Collaboration), *J. Cosmol. Astropart. Phys.* 11 (2025) 062.
 - [10] E. Calabrese *et al.* (ACT Collaboration), *J. Cosmol. Astropart. Phys.* 11 (2025) 063.
 - [11] K. K. Rogers and V. Poulin, *Phys. Rev. Res.* **7**, L012018 (2025).
 - [12] M. A. Fernandez, S. Bird, and M.-F. Ho, *J. Cosmol. Astropart. Phys.* 07 (2024) 029.
 - [13] M. Walther *et al.*, *J. Cosmol. Astropart. Phys.* 05 (2025) 099.
 - [14] S. Hannestad, S. H. Hansen, F. L. Villante, and A. J. S. Hamilton, *Astropart. Phys.* **17**, 375 (2002).
 - [15] E. Camphuis *et al.* (SPT-3G Collaboration), *Phys. Rev. D* **113**, 083504 (2026).
 - [16] N. Aghanim *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A5 (2020).
 - [17] L. Pagano, J. M. Delouis, S. Mottet, J. L. Puget, and L. Vibert, *Astron. Astrophys.* **635**, A99 (2020).
 - [18] S. Goldstein, J. C. Hill, V. Iršič, and B. D. Sherwin, *Phys. Rev. Lett.* **131**, 201001 (2023).
 - [19] H. A. R. El-Ella, A. Huck, and U. L. Andersen, *Phys. Rev. B* **100**, 214407 (2019).
 - [20] J. Torrado and A. Lewis, *J. Cosmol. Astropart. Phys.* 05 (2021) 057.
 - [21] A. Lewis, A. Challinor, and A. Lasenby, *Astrophys. J.* **538**, 473 (2000).
 - [22] M. Raveri and C. Doux, *Phys. Rev. D* **104**, 043504 (2021).
 - [23] R. Kallosh, A. Linde, and D. Roest, *J. High Energy Phys.* 11 (2013) 198.
 - [24] R. Kallosh, A. Linde, and D. Roest, *J. High Energy Phys.* 08 (2014) 052.
 - [25] R. Kallosh and A. Linde, *J. Cosmol. Astropart. Phys.* 10 (2013) 033.
 - [26] A. A. Starobinsky, *Phys. Lett.* **91B**, 99 (1980).
 - [27] F. L. Bezrukov and M. Shaposhnikov, *Phys. Lett. B* **659**, 703 (2008).
 - [28] E. G. M. Ferreira, E. McDonough, L. Balkenhol, R. Kallosh, L. Knox, and A. Linde, *Phys. Rev. D* **113**, 043524 (2026).
 - [29] M. Drees and Y. Xu, *Phys. Lett. B* **867**, 139612 (2025).
 - [30] M. R. Haque, S. Pal, and D. Paul, *Phys. Lett. B* **869**, 139852 (2025).
 - [31] D. S. Zharov, O. O. Sobol, and S. I. Vilchinskii, *Phys. Rev. D* **112**, 023544 (2025).
 - [32] Y.-Y. Ye and B.-M. Gu, *Phys. Lett. B* **873**, 140204 (2026).
 - [33] N. Barnaby and Z. Huang, *Phys. Rev. D* **80**, 126018 (2009).
 - [34] N. Barnaby, *Phys. Rev. D* **82**, 106009 (2010).
 - [35] D. J. H. Chung, E. W. Kolb, A. Riotto, and I. I. Tkachev, *Phys. Rev. D* **62**, 043508 (2000).
 - [36] J. Fumagalli, S. Renaux-Petel, and L. T. Witkowski, *J. Cosmol. Astropart. Phys.* 08 (2021) 030.
 - [37] S. S. Mishra and V. Sahni, *J. Cosmol. Astropart. Phys.* 04 (2020) 007.
 - [38] R. Flauger, L. McAllister, E. Pajer, A. Westphal, and G. Xu, *J. Cosmol. Astropart. Phys.* 06 (2010) 009.
 - [39] R. Flauger, L. McAllister, E. Silverstein, and A. Westphal, *J. Cosmol. Astropart. Phys.* 10 (2017) 055.
 - [40] P. Auclair and C. Ringeval, *Phys. Rev. D* **106**, 063512 (2022).

-
- [41] A. Lewis, *J. Cosmol. Astropart. Phys.* **08** (2025) 025.
[42] M. Tristram *et al.*, *Phys. Rev. D* **105**, 083524 (2022).
[43] T. Biekötter, *SciPost Phys. Codebases* **64**, 1 (2026).
[44] <https://gitlab.com/cosmoPipe/pipe-inflation>
[45] S. Chabanier, M. Millea, and N. Palanque-Delabrouille, *Mon. Not. R. Astron. Soc.* **489**, 2247 (2019).
[46] A. H. Wright *et al.*, *Astron. Astrophys.* **703**, A158 (2025).