

Bayesian inflationary reconstructions from *Planck* 2018 dataWill J. Handley^{1,2,3,*} Anthony N. Lasenby^{1,2,†} Hiranya V. Peiris^{4,5,‡} and Michael P. Hobson^{1,§}¹*Astrophysics Group, Cavendish Laboratory, J.J. Thomson Avenue,
Cambridge CB3 0HE, United Kingdom*²*Kavli Institute for Cosmology, Madingley Road, Cambridge CB3 0HA, United Kingdom*³*Gonville & Caius College, Trinity Street, Cambridge CB2 1TA, United Kingdom*⁴*The Oskar Klein Centre for Cosmoparticle Physics, Department of Physics, Stockholm University,
AlbaNova, Stockholm SE-106 91, Sweden*⁵*Department of Physics and Astronomy, University College London,
Gower Street, London WC1E 6BT, United Kingdom*

(Received 5 August 2019; published 11 November 2019)

We present three nonparametric Bayesian primordial reconstructions using *Planck* 2018 polarization data: linear spline primordial power spectrum reconstructions, cubic spline inflationary potential reconstructions, and sharp-featured primordial power spectrum reconstructions. All three methods conditionally show hints of an oscillatory feature in the primordial power spectrum in the multipole range $\ell \sim 20$ to $\ell \sim 50$, which is to some extent preserved upon marginalization. We find no evidence for deviations from a pure power law across a broad observable window ($50 \lesssim \ell \lesssim 2000$), but find that parametrizations are preferred which are able to account for lack of resolution at large angular scales due to cosmic variance, and at small angular scales due to *Planck* instrument noise. Furthermore, the late-time cosmological parameters are unperturbed by these extensions to the primordial power spectrum. This work is intended to provide a background and give more details of the Bayesian primordial reconstruction work found in the *Planck* 2018 papers.

DOI: [10.1103/PhysRevD.100.103511](https://doi.org/10.1103/PhysRevD.100.103511)**I. INTRODUCTION**

The final release of *Planck* satellite data [1–4] provides an unprecedented window onto the cosmic microwave background (CMB). These high-resolution CMB anisotropy data give constraints on the state of the Universe in its earliest observable stage. Assuming a theory of inflation, the primordial power spectrum of curvature perturbations provides an indirect probe of ultrahigh energy physics.

This paper focuses on nonparametric reconstructions of primordial physics. The aim of such analyses is to provide information on quantities and functions of interest that are arguably model-independent. While unambiguous scientific detections will only ever result from a consideration of specific, physically motivated models, results from reconstructions such as these can be used to inform and guide observational and theoretical cosmology, providing insight and evidence for interesting features not clearly visible in the data when using standard modeling assumptions.

Throughout we adopt a fully Bayesian framework, treating our non-parametric reconstruction functions using

priors, posteriors and evidences to marginalize out factors that are irrelevant to physical quantities of interest. We reconstruct both the inflationary potential and the primordial power spectrum directly using spline and feature-based reconstructions, in a manner related but not identical to the existing literature [5–20].

In Sec. II we review the relevant background theory in primordial cosmology, Bayesian inference, nonparametric reconstruction and CMB data. Section III reconstructs the primordial power spectrum directly using a linear interpolating spline. Section IV takes the analysis one step back and reconstructs the inflationary potential using a cubic spline, treating the primordial power spectrum as a derived quantity. Section V works with a parametrization that is more suited for reconstructing sharp features in the primordial power spectrum as a complementary approach to that of Sec. III. Section VII draws conclusions from all three analyses.

II. BACKGROUND**A. Primordial cosmology**

We begin by summarizing the background theory and establish notation. For a more detailed discussion of inflationary cosmology and perturbation theory, we recommend Mukhanov *et al.* [21] or Baumann [22,23].

*wh260@mrao.cam.ac.uk

†a.n.lasenby@mrao.cam.ac.uk

‡h.peiris@ucl.ac.uk

§mph@mrao.cam.ac.uk

The evolution equations for a spatially homogeneous, isotropic and flat universe filled with a scalar field ϕ with arbitrary potential $V(\phi)$ are

$$\ddot{\phi} + 3H\dot{\phi} + \frac{dV}{d\phi} = 0, \quad (1)$$

$$H^2 = \frac{1}{3M_{\text{p}}^2} \left(\frac{1}{2} \dot{\phi}^2 + V(\phi) \right), \quad (2)$$

where $H = \frac{\dot{a}}{a}$ is the Hubble parameter, a is the scale factor of the Universe, dots denote derivatives with respect to cosmic time $\dot{f} \equiv \frac{df}{dt}$ and M_{p} is the reduced Planck mass. For most potentials $V(\phi)$, solutions to Eqs. (1) and (2) rapidly converge on the attractor *slow roll* state, satisfying $\dot{\phi}^2 \ll V(\phi)$.

The evolution equations for the Fourier k -components of the gauge-invariant comoving curvature \mathcal{R} and tensor \mathcal{T} perturbations are

$$\mathcal{R}_k'' + 2\frac{z'}{z}\mathcal{R}_k' + k^2\mathcal{R}_k = 0, \quad (3)$$

$$\mathcal{T}_k'' + 2\frac{a'}{a}\mathcal{T}_k' + k^2\mathcal{T}_k = 0, \quad (4)$$

$$z = \frac{a\dot{\phi}}{H}, \quad \eta = \int \frac{dt}{a}, \quad (5)$$

where primes denote derivatives with respect to conformal time $f' \equiv \frac{df}{d\eta}$. These equations have the property that in-horizon solutions ($k \gg aH$) oscillate with time-varying amplitude and frequency, whilst out-of-horizon solutions ($k \ll aH$) freeze out. The dimensionless primordial power spectra of these perturbations are defined as

$$\mathcal{P}_{\mathcal{X}}(k) = \lim_{aH \gg k} \frac{k^3}{2\pi^2} |\mathcal{X}_k|^2, \quad \mathcal{X} \in \{\mathcal{R}, \mathcal{T}\}. \quad (6)$$

Initial conditions for the background Eqs. (1) and (2) may be set using the slow roll approximation:

$$H = \sqrt{\frac{V(\phi_{\text{init}})}{3M_{\text{p}}^2}}, \quad \dot{\phi} = -\frac{V'(\phi_{\text{init}})}{3H}, \quad (7)$$

where V' denotes the derivative of V with respect to ϕ . Whilst solutions set with these initial conditions do not lie precisely on the attractor state, they rapidly converge on it. Providing that ϕ_{init} is chosen self-consistently [24–26] with enough additional evolution so that any transient effects are lost, these initial conditions are equivalent to choosing the background solution to be the attractor.

For the perturbation Eqs. (3) and (4), Bunch-Davies initial conditions are chosen such that the Mukhanov variables match onto the de-Sitter vacuum solutions

$$\mathcal{R}_k = \frac{1}{z\sqrt{2k}} e^{-ik\eta}, \quad \mathcal{T}_k = \frac{1}{a\sqrt{2k}} e^{-ik\eta}. \quad (8)$$

Providing that the k -mode lies well within the horizon ($k \gg aH$), this is the canonical choice for initializing the perturbation spectrum, although other vacua are available [27–30].

Whilst Eqs. (1) to (4) take their simplest form using cosmic and conformal time as variables, for numerical stability it is more prudent to choose a timelike parameter which does not saturate during inflation, such as cosmic time t , or the number of e -folds $N = \log(a)$. We choose the logarithmic comoving horizon $\log(aH)$ as the independent variable for our analyses. In this form Eqs. (1)–(4) become complicated, so to avoid typographical errors we generate Fortran source code using the Maple [31] computer algebra package. The numerical integration of all differential equations was performed using the NAG library [32].

It should also be noted that Eqs. (3) and (4) are usually phrased in terms of the Mukhanov variables $v = z\mathcal{R}$ and $h = a\mathcal{T}$, but \mathcal{R} and \mathcal{T} prove to be more numerically stable as they have the attractive property that they explicitly freeze out.

B. Bayesian statistics

Once the primordial power spectra $\mathcal{P}_{\mathcal{R},\mathcal{T}}(k)$ have been determined, these form the initial conditions for Boltzmann codes [33,34]. For a universe described by a cosmological model M with corresponding late-time parameters Θ_{c} and primordial power spectra \mathcal{P} , a Boltzmann code computes CMB power spectra C_{ℓ} in both temperature and polarization. These CMB power spectra may then be fed into cosmological likelihood codes [2], which typically depend on additional nuisance parameters Θ_{n} associated with the experiment. The end result is a likelihood $P(D|\Theta, M)$ of the parameters $\Theta = (\mathcal{P}(k), \Theta_{\text{c}}, \Theta_{\text{n}})$ given CMB data D , and a cosmological model M .

We may formally invert the conditioning on θ in the likelihood using Bayes theorem

$$P(\Theta|D, M) = \frac{P(D|\Theta, M)P(\Theta, M)}{P(D|M)}, \quad (9)$$

$$P(D|M) = \int P(D|\Theta, M)P(\Theta|M)d\Theta, \quad (10)$$

where the first expression above should be read as “*posterior is likelihood times prior over evidence*”, and the second expression indicates that the evidence is the normalizing constant of Eq. (9), and is a multidimensional

marginalization of the likelihood over the prior. The evidence may also be used in a Bayesian model comparison, to assess the relative merits of a set of competing models $\{M_i\}$

$$P(M_i|D) = \frac{P(D|M_i)P(M_i)}{P(D)}, \quad (11)$$

$$P(D) = \sum_i P(D|M_i)P(M_i). \quad (12)$$

In the event of uniform priors over models, the evidence establishes the relative probability weighting to give to models describing the same data D .

Throughout this work, we use a modified version of CAMB [33] to compute C_ℓ power spectra and CosmoChord [35] (a modified version of CosmoMC [36,37]) to interface the likelihoods. To sample the posterior and compute evidences we make use of the nested sampling [38] algorithm PolyChord [39,40]. The default Metropolis–Hastings sampler in CosmoMC is insufficient both due to the complexity of the posteriors that must be navigated, and the requirement of evidence computation. Furthermore, PolyChord is required in place of the previous nested sampling algorithm MultiNest [41–43] due to the high dimensionality of the full *Planck* likelihood with nuisance parameters. As an added bonus, PolyChord has the ability to exploit the fast-slow cosmological hierarchy [37], which greatly speeds up the sampling. Most importantly all parameters associated with the primordial power spectrum are “semislow,” given that one does not need to recompute transfer functions upon changing the primordial power spectrum.

C. Functional inference

For our reconstructions, the quantities of interest are functions $f(k; \Theta_f)$ of wavenumber k , parametrized by a set of parameters Θ_f , which presents a challenge in both plotting and quantifying our results.

We utilize two related techniques to plot the posterior of a function $f(k; \Theta_f)$. First, we can generate equally weighted samples of Θ_f , and therefore of the function f , and plot each sample as a curve on the (k, f) plane. In general, we simultaneously plot prior samples in red, and posterior samples in black. An example of such a plot can be found in the upper-left panel of Fig. 4. For the second type of plot, we first compute the marginalized posterior distribution $P(f|k)$ of the dependent variable f conditioned on the independent variable k using Gaussian kernel density estimation. The isoprobability credibility intervals are then plotted in the (k, f) plane, with their mass converted to σ -values via an inverse error function transformation. An example of this kind of plot can be seen in the upper-right panel of Fig. 4. The code for producing such plots is published in Ref. [44].

To quantify the constraining power of a given reconstruction, we use the conditional Kullback-Leibler (KL) divergence [45] as exemplified by Hee *et al.* [10]. For two distributions $P(x)$ and $Q(x)$, the KL divergence is defined as

$$D_{KL}(P|Q) = \int \ln \left[\frac{P(x)}{Q(x)} \right] P(x) dx, \quad (13)$$

and may be interpreted as the information gain in moving from a prior Q to a posterior P [46–49]. For our reconstructions, we compute the KL divergence from prior to posterior for each distribution $P(f|k)$ conditioned on k . An example of such a plot can be found in the lower-right panel of Fig. 4.

Throughout this work, plots use an approximate correspondence between wavenumber k and multipole moment ℓ via the Limber approximation $\ell \approx k/D_A$, where $D_A = r_*/\theta_*$ is the *Planck* 2018 best-fit comoving angular distance to recombination at r_* .

D. Nonparametric reconstructions

Throughout this work, we explore various nonparametric functional forms for either the primordial power spectrum or the inflationary potential. “Nonparametric” is a slightly misleading terminology, as in general such reconstructions choose a function with a very large number of additional parameters. We prefer the terminology free-form [4,13], flexible [12] or adaptive [10,50]. The principle behind this is that the parametrization should have enough freedom to reconstruct any reasonable underlying function, independent of any underlying physical model.

For example, in this paper we work with variations on the *linear spline*, defined by parameters Θ_f , producing a mapping from the independent variable x to the dependent variable y thus

$$\text{Lin}(x; \Theta_f) = \sum_{i=1}^N \frac{y_i(x_{i+1} - x) + y_{i+1}(x - x_i)}{x_i - x_{i+1}} [x_i < x \leq x_{i+1}],$$

$$\Theta_f = (x_1, \dots, x_N, y_1, \dots, y_N). \quad (14)$$

Here we have used a compact notation for denoting piecewise functions espoused by Graham *et al.* [51] whereby $[R]$ is a logical truth function, yielding 1 if the relation R is true, and 0 if false. For consistency, we interpret the case $N = 1$ as having a constant value of y_1 for all x .

In a Bayesian approach, one treats the additional degrees of freedom Θ_f of the nonparametric function as parameters in a posterior distribution, which one marginalizes out in order to obtain model-independent reconstructions. Typically there is a degree of choice as to how many parameters N to use, and a penalty is applied for larger N to avoid over-parametrization and noise fitting. In this work

we treat N in a Bayesian sense as well. Each reconstruction with a given number of parameters N is treated as an independent model. We can then marginalize over the number of models using the Bayesian evidence.

E. *Planck* data and cosmology

In this paper in almost all cases we focus our efforts on using the pure *Planck* 2018 polarization data baseline, referred to in [1–4] as TT, TE, EE + lowE + lensing. Throughout we use a flat cold-dark-matter with dark energy (Λ CDM) late-time cosmology; therefore there are four associated cosmological parameters which in the default CosmoMC basis take the form

$$\Theta_c = (\Omega_b h^2, \Omega_c h^2, 100\theta_{MC}, \tau). \quad (15)$$

TABLE I. The prior distributions on late-time cosmological parameters and *Planck* nuisance parameters for all analyses. The parameters of each Gaussian distribution are defined as $[\mu, \sigma]$, and the above distributions combine to make a truncated Gaussian distribution on the nuisance parameters A^{kSZ} and A_{143}^{tSZ} . The nuisance priors are the default ones in CosmoMC, whilst the cosmological priors are narrowed to speed up sampling, but remain sufficiently wide to effectively include the entire posterior mass. Also indicated is each parameter’s speed with respect to the CosmoMC fast-slow hierarchy.

Parameters	Prior type	Prior parameters	Speed
$\Omega_b h^2$	Uniform	[0.019, 0.025]	Slow
$\Omega_c h^2$	Uniform	[0.095, 0.145]	Slow
$100\theta_{MC}$	Uniform	[1.03, 1.05]	Slow
τ	Uniform	[0.01, 0.4]	Slow
y_{cal}	Gaussian	1 ± 0.0025	Semi-slow
A_{217}^{CIB}	Uniform	[0, 200]	Fast
$\xi^{tSZ-CIB}$	Uniform	[0, 1]	Fast
A_{143}^{tSZ}	Uniform	[0, 10]	Fast
A_{100}^{PS}	Uniform	[0, 400]	Fast
A_{143}^{PS}	Uniform	[0, 400]	fast
$A_{143 \times 217}^{PS}$	Uniform	[0, 400]	Fast
A_{217}^{PS}	Uniform	[0, 400]	Fast
A^{kSZ}	Uniform	[0, 10]	Fast
A_{100}^{dustTT}	Gaussian	8.6 ± 2	Fast
A_{143}^{dustTT}	Gaussian	10.6 ± 2	Fast
$A_{143 \times 217}^{dustTT}$	Gaussian	23.5 ± 8.5	Fast
A_{217}^{dustTT}	Gaussian	91.9 ± 20	Fast
A_{100}^{dustTE}	Gaussian	0.13 ± 0.042	Fast
$A_{100 \times 143}^{dustTE}$	Gaussian	0.13 ± 0.036	Fast
$A_{100 \times 217}^{dustTE}$	Gaussian	0.46 ± 0.09	Fast
A_{143}^{dustTE}	Gaussian	0.207 ± 0.072	Fast
$A_{143 \times 217}^{dustTE}$	Gaussian	0.69 ± 0.09	Fast
A_{217}^{dustTE}	Gaussian	1.938 ± 0.54	Fast
c_{100}	Gaussian	1.0002 ± 0.0007	Fast
c_{217}	Gaussian	0.99805 ± 0.00065	Fast
$A^{kSZ} + 1.6A_{143}^{tSZ}$	Gaussian	9.5 ± 3	Fast

Additionally there are 21 nuisance parameters associated with Galactic foregrounds and the *Planck* instrumentation:

$$\Theta_n = (y_{cal}, A_{217}^{CIB}, \xi^{tSZ-CIB}, A_{143}^{tSZ}, A_{100}^{PS}, A_{143}^{PS}, A_{143 \times 217}^{PS}, A_{217}^{PS}, A^{kSZ}, A_{100}^{dustTT}, A_{143}^{dustTT}, A_{143 \times 217}^{dustTT}, A_{217}^{dustTT}, A_{100}^{dustTE}, A_{100 \times 143}^{dustTE}, A_{100 \times 217}^{dustTE}, A_{143}^{dustTE}, A_{143 \times 217}^{dustTE}, A_{217}^{dustTE}, c_{100}, c_{217}). \quad (16)$$

We use the *Planck* PolyChord CosmoMC defaults as prior widths for all of these (indicated in Table I), but as they are common to all models considered and sufficiently wide to encompass the entire posterior bulk, any prior effects from these parameters have been shown theoretically [52] and in practice [53] to cancel out.

F. Sampling strategy

Throughout this paper, we sample over the full parameter space $\Theta = (\Theta_f, \Theta_c, \Theta_n)$ of reconstruction, cosmological and nuisance parameters. The resulting posteriors are in general multimodal with complicated degeneracies between many parameters, particularly when there are a large number of reconstruction parameters Θ_f . In all cases, plots in this paper have the unmentioned parameters implicitly marginalized out. Marginalizing the likelihood over the prior in order to compute evidences is in general even more challenging. Nested sampling is ideally suited to performing such tasks, with PolyChord providing the cutting-edge of such technology in a cosmological context [39,40], proving to be essential for sampling over these complicated parameter spaces with up to $\sim \mathcal{O}(50)$ dimensions with fast-slow parameter hierarchies.

III. PRIMORDIAL POWER SPECTRUM RECONSTRUCTION

Traditionally in a Λ CDM cosmology, the primordial power spectrum $\mathcal{P}_{\mathcal{R}}(k)$ is modeled by a two-parameter function with an amplitude A_s and spectral index $n_s - 1$

$$\ln \mathcal{P}_{\mathcal{R}}(k) = \ln A_s + (n_s - 1) \ln \left(\frac{k}{k_*} \right), \quad (17)$$

i.e., a straight line in the $(\ln k, \ln \mathcal{P})$ plane. The tensor spectrum $\mathcal{P}_{\mathcal{T}}(k)$ may be parametrized by its own independent amplitude A_t and index n_t , or via the tensor to scalar ratio $r = A_t/A_s$ and slow roll inflation consistency condition $n_t = -r/8$ [22]. Extensions to parametrization (17) can be made by adding quadratic (running) and cubic (running of running) terms, but no evidence is found that these are required to describe the primordial power spectrum in the k -window which *Planck* probes.

Extending Eq. (17) with runnings of the spectral index creates a stiff parametrization, with no ability to account for sharper features, or large deviations at low or high- k .




FIG. 1. We parametrize the primordial power spectrum reconstruction via a linear interpolating spline in the $(\ln k, \ln \mathcal{P})$ plane with N spline locations $(k_1, \mathcal{P}_1), \dots, (k_N, \mathcal{P}_N)$. The outermost k -locations are fixed, with the inner locations constrained by $k_1 < \dots < k_N$, and the entire spline constrained within the box indicated by the dashed line.

For our first primordial power spectrum reconstruction, we therefore parametrize as a logarithmic spline

$$\begin{aligned} \ln 10^{10} \mathcal{P}_{\mathcal{R}}(k) &= \text{Lin}(\log_{10} k; \Theta_{\mathcal{P}}) \\ \Theta_{\mathcal{P}} &= (\log_{10} k_1, \dots, \log_{10} k_N, \\ &\quad \ln 10^{10} \mathcal{P}_1, \dots, \ln 10^{10} \mathcal{P}_N). \end{aligned} \quad (18)$$

This represents a spline [Eq. (14)] that is linear in the $(\ln k, \ln \mathcal{P})$ plane as shown in Fig. 1. For the tensor power spectrum we have analyzed the cases where r is allowed to vary as a parameter, and also when $\mathcal{P}_{\mathcal{T}}(k)$ is given its own independent linear spline. Unsurprisingly, given that *Planck* measured an r consistent with 0, the addition of a tensor power spectrum makes no difference to the scalar reconstructions. For simplicity we assume $r = 0$ for the remainder of this section.

This technique has a history of being successfully applied to the primordial power spectrum [5–9], but has also been applied to dark energy equation of state by Hee *et al.* [10] and Vázquez *et al.* [11], to the cosmic reionization history by Millea and Bouchet [12] and to galaxy cluster profiles by Olamaie *et al.* [13]. Our work differs from previous primordial power spectrum reconstructions in both the data we use, the styling of the priors, and in the application of more modern inference tools such as functional posterior plotting [44], conditional Kullback-Leibler divergences [10] and PolyChord. The technique was applied to primordial power spectrum reconstruction from CORE simulated data in Sec. 6 of [7], where it was shown that this approach accurately reconstructs complicated injected features (or lack thereof).

A. Priors

For priors on the vertical spline location parameters, we choose them to be independently uniform in $2 < \log 10^{10} \mathcal{P} < 4$. This spans an almost maximally wide range, increasing their width further has little effect due to CosmoMC discarding unphysically normalized spectra.

For priors on the horizontal spline location parameters, we choose the outermost knots to be fixed at 10^{-4} and $10^{-0.3}$. This corresponds roughly to a C_ℓ multipole range $1 \lesssim \ell \lesssim 7000$, which fully encompasses the CMB window that *Planck* observes. For the remaining horizontal knots, we choose a prior which distributes the parameters logarithmically within this range, such that $k_2 < \dots < k_{N-1}$. This sorting procedure breaks the $(N-2)!$ implicit switching degeneracy, and is also termed a *forced identifiability prior* [40,54,55].

To implement this sorted prior in the context of nested sampling, we need to define the transformation from the unit hypercube to the physical space. Coordinates in the unit hypercube x_1, \dots, x_N , can be transformed to coordinates in the physical space $\theta_1, \dots, \theta_N$, such that they are distributed uniformly in $[\theta_{\min}, \theta_{\max}]$ and sorted so that $\theta_1 < \dots < \theta_N$ via the following reversed recurrence relation

$$\theta_n = \theta_{\min} + (\theta_{n+1} - \theta_{\min}) x_n^{1/n}, \quad \theta_{N+1} = \theta_{\max}, \quad (19)$$

which is equivalent to saying that θ_n is marginally distributed as the largest of n uniformly distributed variables within $[\theta_{\min}, \theta_{n+1}]$. Another method for breaking the switching degeneracy is to exclude the region of the parameter space which does not satisfy the sorting criterion. This becomes exponentially small as more knots are added, which makes the initial sampling from the prior more challenging. It is also more in keeping with the nested sampling methodology to explicitly transform the full hypercube onto the space of interest. In our case, given that k_1 and k_N are fixed, for $N \geq 4$, we sort the $N-2$ inner logarithmic coordinates $\log_{10} k_2 < \dots < \log_{10} k_{N-1}$.

We perform the reconstruction for $N = 1, \dots, 9$ and then marginalize using Bayesian evidences with an implicit equal weighting for each N . This is equivalent to sampling from a full joint posterior with a uniform prior on N , and could alternatively be accomplished using the method described in Hee *et al.* [50]. Our priors on the reconstruction parameters are summarized in Table II.

TABLE II. The prior distributions on early-time cosmological parameters for the primordial power spectrum reconstruction.

Parameters	Prior type	Prior range
N	Discrete uniform	[1, 9]
$\mathcal{P}_1, \dots, \mathcal{P}_N$	Log-uniform	$10^{-10}[e^2, e^4]$
$k_2 < \dots < k_{N-1}$	Sorted log-uniform	$[10^{-4}, 10^{-0.3}]$




FIG. 2. Equally-weighted sample plots of primordial power spectrum reconstructions, conditioned on the number of knots N . The outermost knots are fixed at the bounds of the figure, so $N = 1$ is equivalent to a scale-invariant primordial power spectrum, $N = 2$ is equivalent to Λ CDM, up to a small difference in prior and $N > 2$ has $N - 2$ knots capable of moving in both the k and \mathcal{P} directions. Prior samples are drawn in red, whilst posterior samples are indicated in black.

B. Results

We show results for our primordial power spectrum reconstruction using *Planck* 2018 TT, TE, EE + lowE + lensing data in Figs. 2–4.

Figures 2 and 3 show the prior and posterior conditioned on each value of N . The case $N = 1$ corresponds to a scale-invariant spectrum, whilst $N = 2$ is equivalent (up to a small difference in prior) to the standard Λ CDM

parametrization. As further knots are added, the reconstruction accounts for cosmic variance at low- k , and instrument noise at high- k . Furthermore, for large N , in a fraction of the samples there is a visibly clear oscillation characterized by a rise in power at $\ell \sim 50$, and a dip in power at $20 < \ell < 30$, as well as an overall suppression of power at low- k . For lower values of ℓ , cosmic variance sets in, and few conclusions can be drawn




FIG. 3. Same as Fig. 2, but plotted using isoproability credibility intervals as discussed in Sec. II C. Blue and red contours represent prior and posterior respectively.

from sampling differences between runs of different N at these values.

To determine the statistical significance of these features, one should consider the Bayesian evidence, as indicated in the lower-left panel of Fig. 4.

The first observation from Fig. 4 is that the $N = 1$ scale-invariant power spectrum is completely ruled out, with a logarithmic difference of $\ln \mathcal{B}_{N=1}^{N=2} \sim \mathcal{O}(33)$. This represents overwhelming evidence for a tilted power spectrum, one of the key predictions of the theory inflation. A gambler could

get odds of a quintillion to one against scale-invariance vs Λ CDM.

The second observation is that the evidence for $N = 3$ is greater than $N = 2$, namely a model that is able to account for cosmic variance at low- k and instrument noise at high- k is, in a Bayesian sense, preferred to the simpler Λ CDM parametrization. Up until *Planck 2018*, the data had not been quite powerful enough for us to define the window that we observe in the primordial power spectrum in this Bayesian sense.




FIG. 4. Bottom-left: Bayesian evidence as a function of number of knots N for the primordial power spectrum reconstruction. Top: Marginalized primordial power spectrum plot. These are produced by taking Figs. 2 and 3 and weighting each panel by their respective evidence. Bottom-right: Marginalized conditional Kullback-Leibler divergence.

The third observation is that whilst $N = 3$ is maximal in evidence, in fact $N = 4, 5, 6$ are competitive, and $N = 7, 8, 9$ are far from ruled out. With this lack of knowledge, the correct Bayesian approach is to marginalize over all models, using the Bayesian evidence as the relative weighting. Doing so, we can compute the marginalized spectrum and KL divergence as shown in Fig. 4. We find that the observable window is now clearly defined, and hints of the low- k features survive this marginalization.

C. Historical context

Figure 5 shows reconstructions using the same methodology¹ but now on data from a historical sequence of CMB experiments

¹For the historical data prior to WMAP, we needed to significantly widen the priors on cosmological late-time parameters

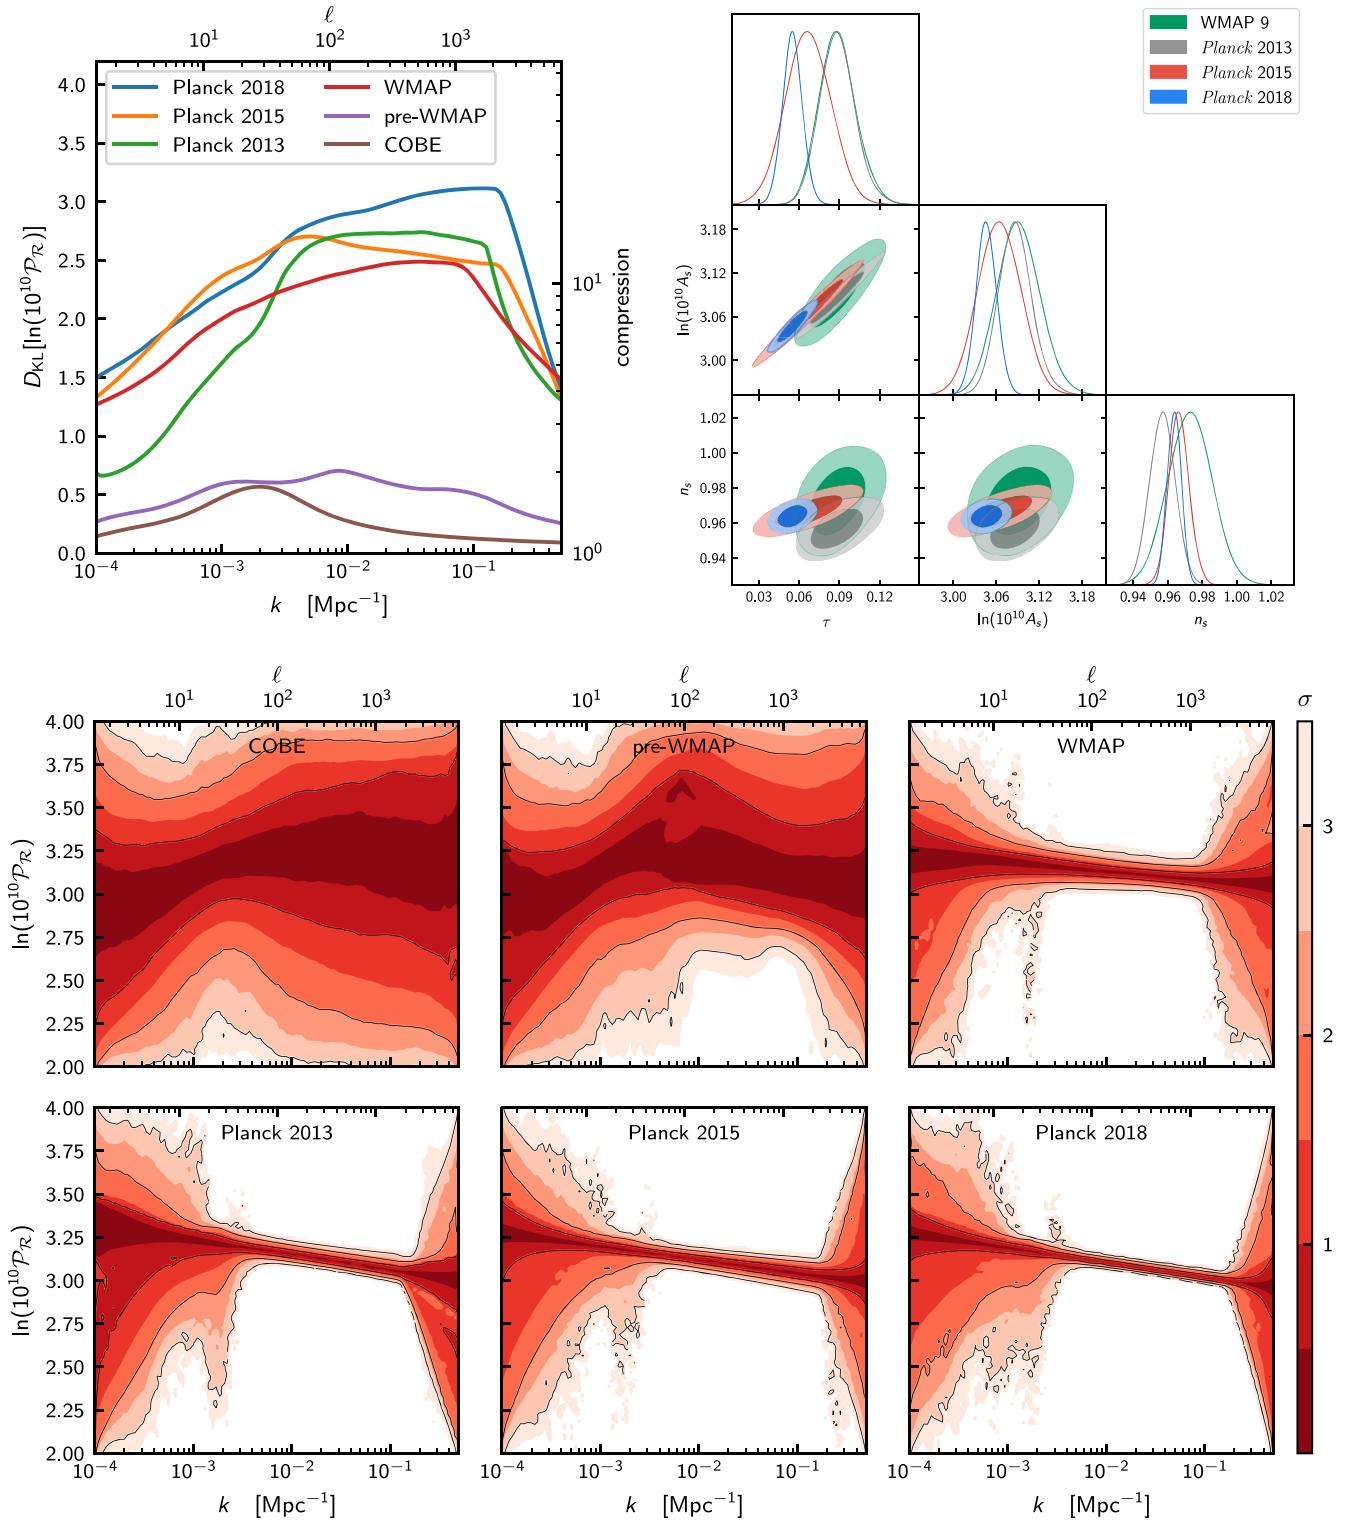


FIG. 5. Historical primordial power spectrum reconstructions. Top-left: Conditional Kullback-Leibler divergences. Top-right: inflationary power spectrum summary parameters, and the influence of τ on *Planck* constraints. Bottom: Marginalized power spectrum plots for each dataset.

- (1) COBE [56],
- (2) “pre-WMAP” (COBE [56], BOOMERANG [57], MAXIMA [58], DASI [59], VSA [60] and CBI [61]),
- (3) WMAP [62,63],
- (4) *Planck* 2013 (TT + lowlike + lensing) [64],
- (5) *Planck* 2015 (TT + lowTEB + lensing) [65],
- (6) *Planck* 2018 (TT, TE, EE + lowE + lensing) [1].

The hint of an $20 < \ell < 30$ feature becomes visible after WMAP, is strengthened in switching to *Planck*, and remains stable as the *Planck* data are updated.

Examining the KL divergences in Fig. 5, the vertical axis shows a greater overall constraint on the primordial power spectrum as improved cosmological constraints are obtained, and the horizontal axis shows that the k -window increases as the angular resolution of the experiments increases. The only alteration to this trend is the change from *Planck* 2013 to *Planck* 2015. In this case, the constraint on the primordial power spectrum is actually lowered, whilst the k -window increases. This is due to the fact that the τ constraint widened from 2013 to 2015, as can be seen in the top right panels of the Fig. 5.

Planck 2018 provides the best constraints on the primordial power spectrum, both via its high-accuracy measurement of τ , and in its small-scale angular resolution.

IV. INFLATIONARY POTENTIAL RECONSTRUCTION

In contrast to the analysis from the previous section, instead of parametrizing the primordial power spectrum directly, here we take the pipeline one stage backward and perform a nonparametric reconstruction of the inflaton potential $V(\phi)$. The scalar and tensor primordial power spectra $\mathcal{P}_{\mathcal{R},\mathcal{T}}$ are then derived from $V(\phi)$ via the procedure indicated in Sec. II A.

To reconstruct the inflationary potential $V(\phi)$, it is more appropriate to work first with $\ln V$, as in general $V(\phi)$ can *a-priori* span a many scales, and it is typically $\frac{d}{d\phi} \ln V$ that drives much of the evolution of the inflaton during inflation.

More importantly, one cannot parametrize $V(\phi)$ via a linear interpolating spline as was the case in Sec. III. The equations of motion [Eqs. (1)–(4)] in general depend on first (and sometimes second) derivatives of $V(\phi)$. Parametrizing the potential using a linear spline will typically yield primordial power spectra with (arguably) unphysical ringing effects.

One should therefore use a spline with continuous first derivatives, and it is natural to choose a cubic spline as the conceptually simplest smooth interpolator. It is tempting to try to do this directly by taking the locations of the knots of the spline as free parameters. Cubic splines, however, are stiff, yielding complicated posteriors that are very difficult to navigate, interpret, and set priors on [5].




FIG. 6. For the inflationary potential reconstruction, we parametrize the second derivative of the logarithmic potential via a linear interpolating spline, and then integrate twice to recover the logarithmic potential (seen schematically from bottom to top of the figure). This introduces two additional parameters: a gradient and global offset.

Cubic splines have the property that their derivative is a smooth piecewise quadratic, and their second derivative is a piecewise linear spline. This suggests that the cleanest way to reconstruct the potential is to parametrize the second derivative as a linear spline, and then integrate this function twice to get the log-potential. Two additional parameters are created by this double integration, a gradient term $\frac{d \ln V_*}{d\phi}$ and an overall offset $\ln V_*$. These two free parameters function as an alternative constraint choice in comparison with natural or clamped splines. Our reconstruction function is therefore

TABLE III. The prior distributions on early-time cosmological parameters for the inflationary potential reconstruction. $\tilde{\phi}_{\min}$ and $\tilde{\phi}_{\max}$ are defined by the observable window of the unperturbed potential. There is a further prior constraint in that we require that the inflaton should evolve in an inflating phase throughout the observable window and that the inflaton should be rolling downhill from negative to positive ϕ throughout.

Parameters	Prior type	Prior range
N	Discrete uniform	[0, 8]
$\ln V_*$	Uniform	[-25, -15]
$\frac{d \ln V_*}{d\phi}$	Log-uniform	$[10^{-3}, 10^{-0.3}]$
$\frac{d^2 \ln V_1}{d\phi^2}, \dots, \frac{d^2 \ln V_N}{d\phi^2}$	Uniform	[-0.5, 0.5]
ϕ_1, \dots, ϕ_N	Sorted uniform	$[\tilde{\phi}_{\min}, \tilde{\phi}_{\max}]$
$\ln 10^{10} \mathcal{P}_{\mathcal{R}}(k)$	Indirect constraint	[2, 4]




FIG. 7. Bottom-left: Bayesian evidence as a function of number of knots N for the inflationary potential reconstruction. Top: Marginalized functional posteriors for the primordial power spectrum. These are produced by taking Fig. 8 and weighting each panel by their respective evidence. Bottom-right: Marginalized conditional Kullback-Leibler divergence.

$$\begin{aligned}
 \ln V &= \ln V_* + (\phi - \phi_*) \frac{d \ln V_*}{d\phi} \\
 &\quad + \int_{\phi_*}^{\phi} d\phi' \int_{\phi_*}^{\phi'} d\phi'' \text{Lin}(\phi''; \Theta_V) \\
 \Theta_V &= \left(\phi_1, \dots, \phi_N, \frac{d^2 \ln V_1}{d\phi^2}, \dots, \frac{d^2 \ln V_N}{d\phi^2}, \frac{d \ln V_*}{d\phi}, \ln V_* \right),
 \end{aligned} \tag{20}$$

which can be viewed as working through Fig. 6 in reverse.

A. Priors

The priors in this analysis proved to be critically important for recovering sensible results. To harmonize with the analysis of the primordial power spectrum, our first requirement is that any primordial power spectrum generated from a potential $V(\phi)$ resides in the range $2 < \ln 10^{10} \mathcal{P}_{\mathcal{R}}(k) < 4$.

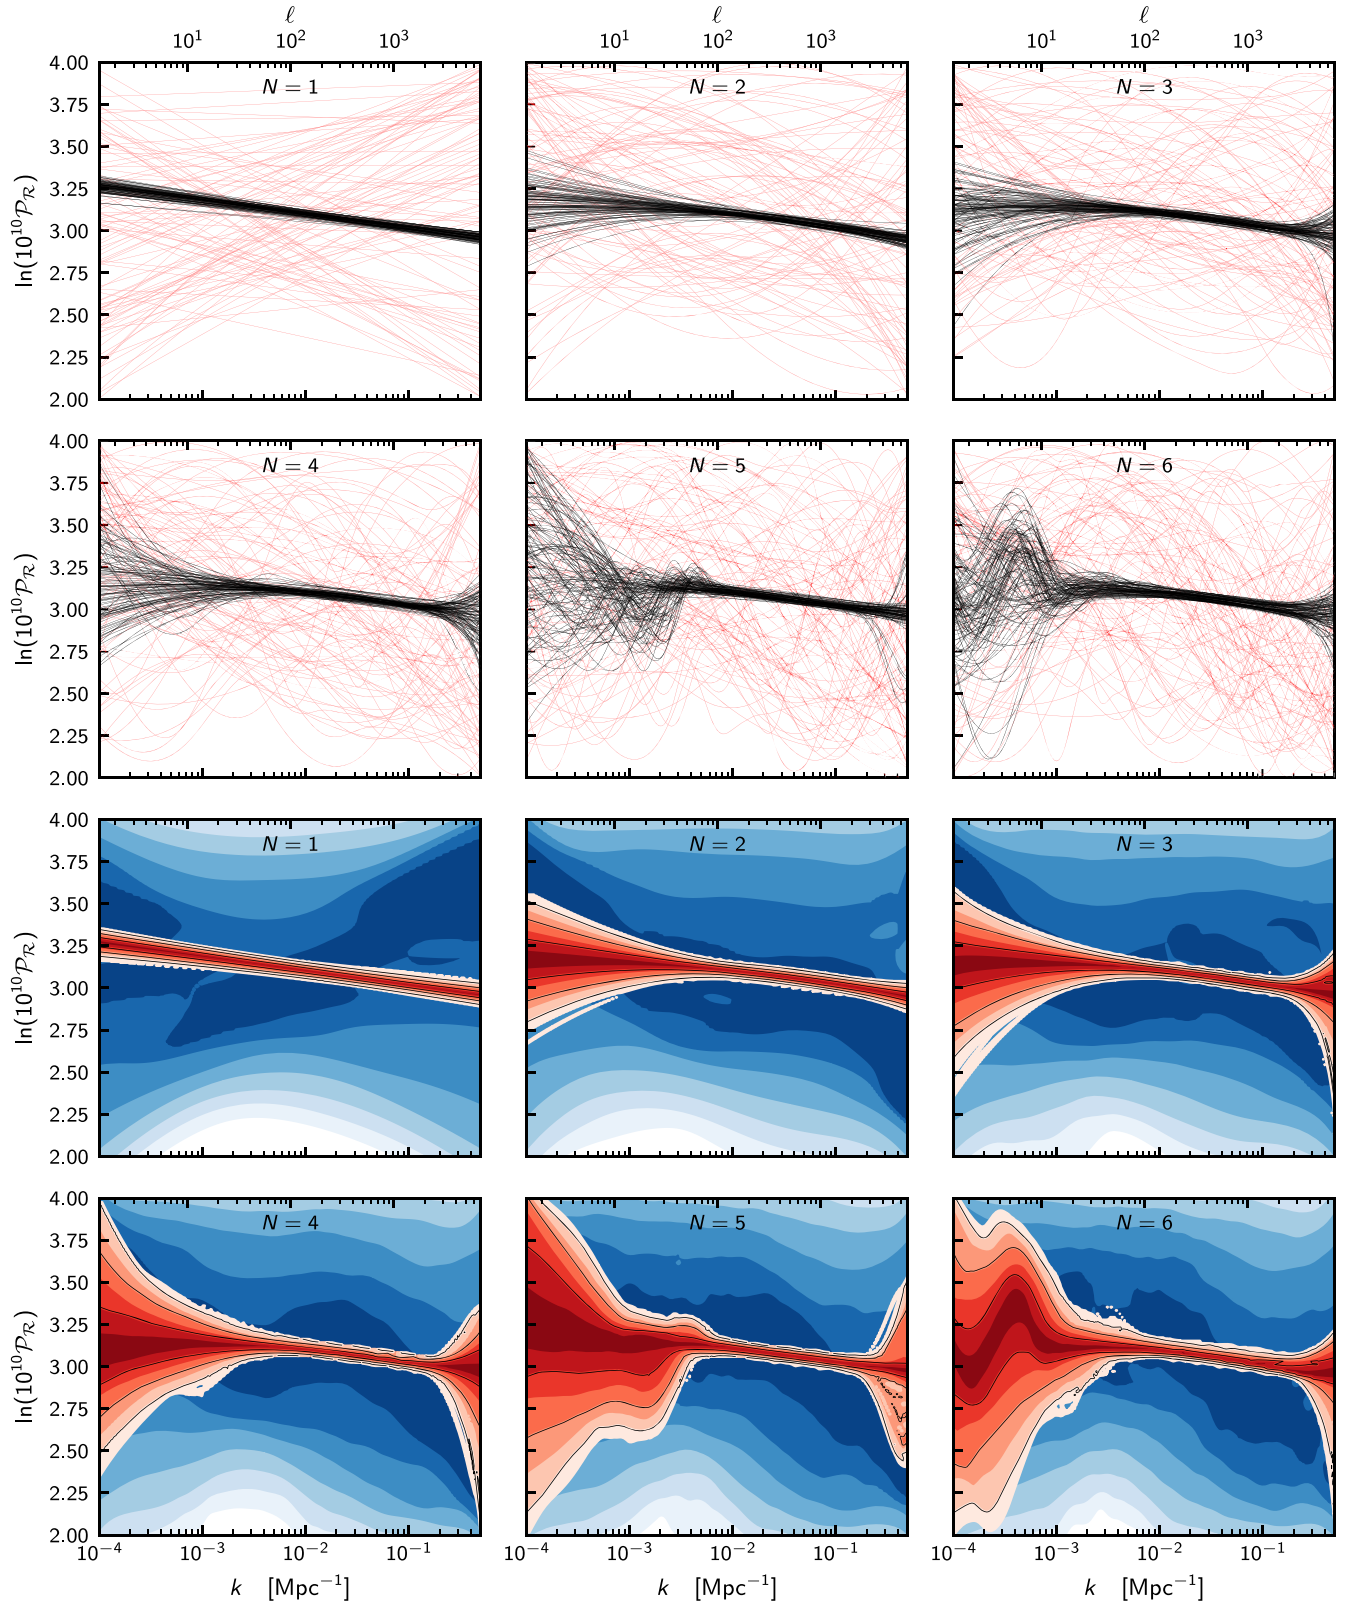


FIG. 8. Equally-weighted sample plots of the functional posterior of the primordial power spectrum from the inflationary potential reconstruction, conditioned on the number of knots N . $N = 0, 1$ have a potential equivalent to a first and second-order Taylor expansion respectively, whilst $N \geq 2$ provide the ability to reconstruct broad features in the underlying potential. Prior samples are drawn in red, whilst posterior samples are indicated in black.

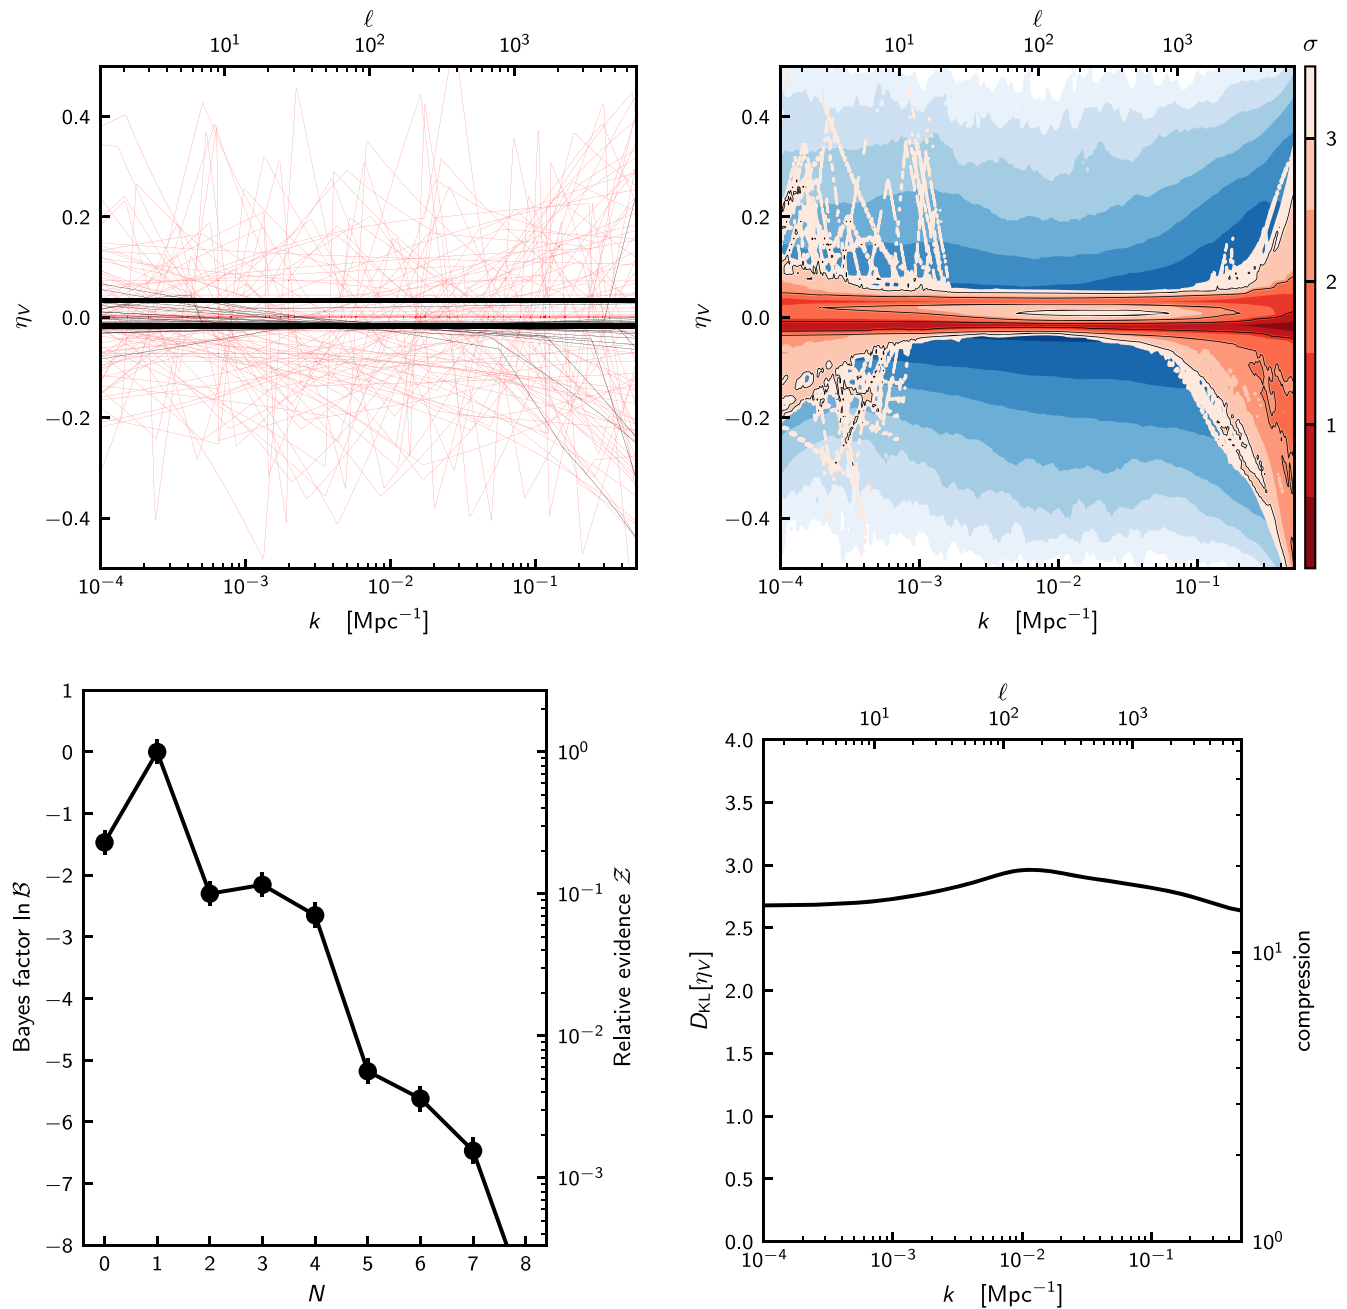


FIG. 9. Bottom-left: Bayesian evidence as a function of number of knots N for the inflationary potential reconstruction. Top: Marginalized functional posterior of the inflationary parameter η_V . These are produced by taking Fig. 10 and weighting each panel by their respective evidence. Bottom-right: Marginalized conditional Kullback-Leibler divergence.

Consider the slow roll parameters [66], and their relation to the second derivative of the log potential

$$\varepsilon_V = \frac{1}{2} \left(\frac{1}{V} \frac{dV_*}{d\phi} \right)^2, \quad \eta_V = \frac{1}{V} \frac{d^2 V_i}{d\phi^2} \Rightarrow \frac{d^2 \ln V}{d\phi^2} = \eta_V - 2\varepsilon_V.$$

We therefore take the priors on the second derivatives of the log potential $\frac{d^2 \ln V_i}{d\phi^2}$ to be uniformly distributed, and the gradient $\frac{d \ln V_*}{d\phi}$ is taken to be negatively log-uniform.

Negativity forces the inflaton to roll downhill from negative to positive ϕ , breaking a symmetric degeneracy. We take the potential offset to vary across a wide range $\ln V_*$. Widening any of these priors detailed in Table III further has no effect, as any primordial power spectrum generated outside these bounds lies outside the range [2, 4].

Particular care must be taken with the horizontal locations of the knots. Any reconstruction of the potential will be sensitive only to the observable window of inflation $[\phi_{\min}, \phi_{\max}]$, defined as when the largest and smallest




FIG. 10. Same as Fig. 8, but now for the inflationary slow roll parameter $\eta_V(k)$, with independent variable defined by an effective wave number k , which sits in one-to-one correspondence with ϕ via the size of the comoving Hubble radius at that moment in the field's evolution.




FIG. 11. For the sharp features reconstruction, we parametrize the primordial power spectrum via traditional amplitude-tilt (A_s, n_s) parametrization, with N top-hat features. We constrain the spectrum to be within the dashed box.

observable scales k_{\min} and k_{\max} exit the horizon. As in Sec. III, we take $(k_{\min}, k_{\max}) = (10^{-4}, 10^{-0.3}) \text{ Mpc}^{-1}$.

Unfortunately, the bounds of the window $[\phi_{\min}, \phi_{\max}]$ are strongly dependent on the other primordial parameters. One cannot therefore take an arbitrarily wide range in ϕ for the horizontal locations, as the reconstruction is then dominated by the prior effect of unconstrained knot parameters. The locations ϕ_1, \dots, ϕ_N of the reconstruction knots should instead be distributed throughout the observable window. Whilst the locations ϕ_1, \dots, ϕ_N and heights $\frac{d^2 \ln V_1}{d\phi^2}, \dots, \frac{d^2 \ln V_N}{d\phi^2}$ themselves influence the size of the observable window, a reasonable approach is to first estimate it using the unperturbed potential (i.e., setting $N = 0$), giving an alternative window $[\tilde{\phi}_{\min}, \tilde{\phi}_{\max}]$. In a similar manner as the horizontal knots in Sec. III, we take the N horizontal ϕ -knot locations to be sorted and uniform throughout this window.

Finally, we require that the inflaton should evolve in an inflating phase throughout the observable window, and that it should be rolling (not necessarily in slow roll) downhill from negative to positive ϕ throughout. These priors are summarized in Table III.

TABLE IV. The prior distributions on early-time cosmological parameters for the sharp feature reconstructions.

Parameters	Prior type	Prior range
N	Discrete uniform	$[0, 8]$
A_s	Uniform	$10^{-10}[e^2, e^4]$
n_s	Uniform	$[0.8, 1.2]$
h_1, \dots, h_N	Uniform	$[-1, 1]$
$k_2 < \dots < k_{N-1}$	Sorted log-uniform	$[10^{-4}, 10^{-0.3}]$
$\Delta_1, \dots, \Delta_N$	Uniform	$[0, 1]$
$\ln 10^{10} \mathcal{P}_{\mathcal{R}}(k)$	Indirect constraint	$[2, 4]$

Alternative methodologies for direct reconstruction of the potential exist in the literature. One approach is to expand the potential $V(\phi)$ as a Taylor series [24,67]. Another is to expand $H(\phi)$ as a Taylor series [26,68], and then derive the potential analytically via $V(\phi) = 3M_{\text{p}}^2 H^2 - 2M_{\text{p}}^4 H'^2$. Both of these approaches have been successfully applied in the *Planck* inflation papers [4,8,9].

B. Results

Due to the strong dependency of the ϕ -window on the potential itself, it is not particularly illuminating to plot $V(\phi)$ directly. Instead, in the spirit of the other two sections we start by plotting the functional posterior of the primordial power spectrum $\mathcal{P}_{\mathcal{R}}(k)$, shown in Figs. 7 and 8. Viewed in this manner, one can think of these primordial power spectrum reconstructions as having an alternative prior complementary to Sec. III, motivated by the assumption that the primordial power spectrum is derived from a smooth underlying potential.

In the same manner as Sec. III, Fig. 8 and is consequently a form of exponential potential. Regardless, it recovers a primordial power spectrum with an appropriate amplitude and tilt and minimal running, almost identical to the traditional A_s, n_s parametrization. $N = 1$ adds a constant second derivative term to the Taylor expansion, and produces a similar primordial power spectrum. As more knots are added, the potential has greater freedom, and the corresponding primordial power spectrum begins to gain similar features to the results in Sec. III; a loss of constraint at low and high- k . Intriguingly, there is also the same preference for an oscillation with a peak at $\ell \sim 50$ and trough at $20 < \ell < 30$. In the line plots of Fig. 8 the oscillation is now smooth, on account of the physical potential-based prior created by this reconstruction. It should be noted that while such oscillations are characteristic of this integrated inflationary potential parametrization, these were also partially recovered *a priori* in the free-form approach from Sec. III.

Examining the evidences in Fig. 7, we can see that despite the similarities in primordial power spectra $N = 1$ is preferred over $N = 0$. The reason for this is that the restrictive form of potential for $N = 0$ forces $r \approx 0.2$, which is now ruled out by *Planck*. Allowing a second derivative for the $N = 1$ relaxes the r constraint, resulting in a Bayesian preference for the $N = 1$ case, consistent with the results of the Planck Collaboration [4]. Adding further knots causes the evidence to drop, indicating that from a Bayesian standpoint, no further complexity is required by the data. The marginalized plots in Fig. 7 show similar attributes to those of the primordial power spectrum in Fig. 4, but in this case the stiffness of the primordial power spectrum reconstruction results in a slightly poorer recovery of the relative lack of power spectrum constraint at low and high- k .




FIG. 12. Bottom-left: Bayesian evidence as a function of number of knots N for the sharp features reconstruction. Top: Marginalized primordial power spectrum plot. These are produced by taking Fig. 13 and weighting each panel by their respective evidence. Bottom-right: Marginalized conditional Kullback-Leibler divergence.

Our second set of plots detail results for the inflationary slow roll parameter $\eta_V(k)$, shown in Figs. 9 and 10. Instead of using ϕ as the independent variable (which suffers from the dependency of the window widths on the underlying potential), we use the effective wavenumber $k(\phi)$, which is monotonically related to ϕ in our reconstruction, defined to be the size of the comoving Hubble radius at that moment in the field's evolution.

Figure 10 reveals that the oscillations in the primordial power spectrum at low k are created by a partial breakdown in

the slow roll conditions. For $N \geq 5$, $\eta_V \sim 0.5$, which are the same values of N at which oscillations become apparent in Fig. 8. This will be of particular interest for just enough inflation models [69–72], models with singularities and discontinuities [73–75], multifield phase-transitions [76–80], M-theory [81,82], or supergravity [83] models, to name a few examples. There is a long history of confronting such models with data [84–91]. Figure 9 details the marginalized results.

Planck provides only a weak upper bound on the other slow roll parameter $\epsilon_V \approx \frac{r}{16}$, meaning ϵ_V is nearly

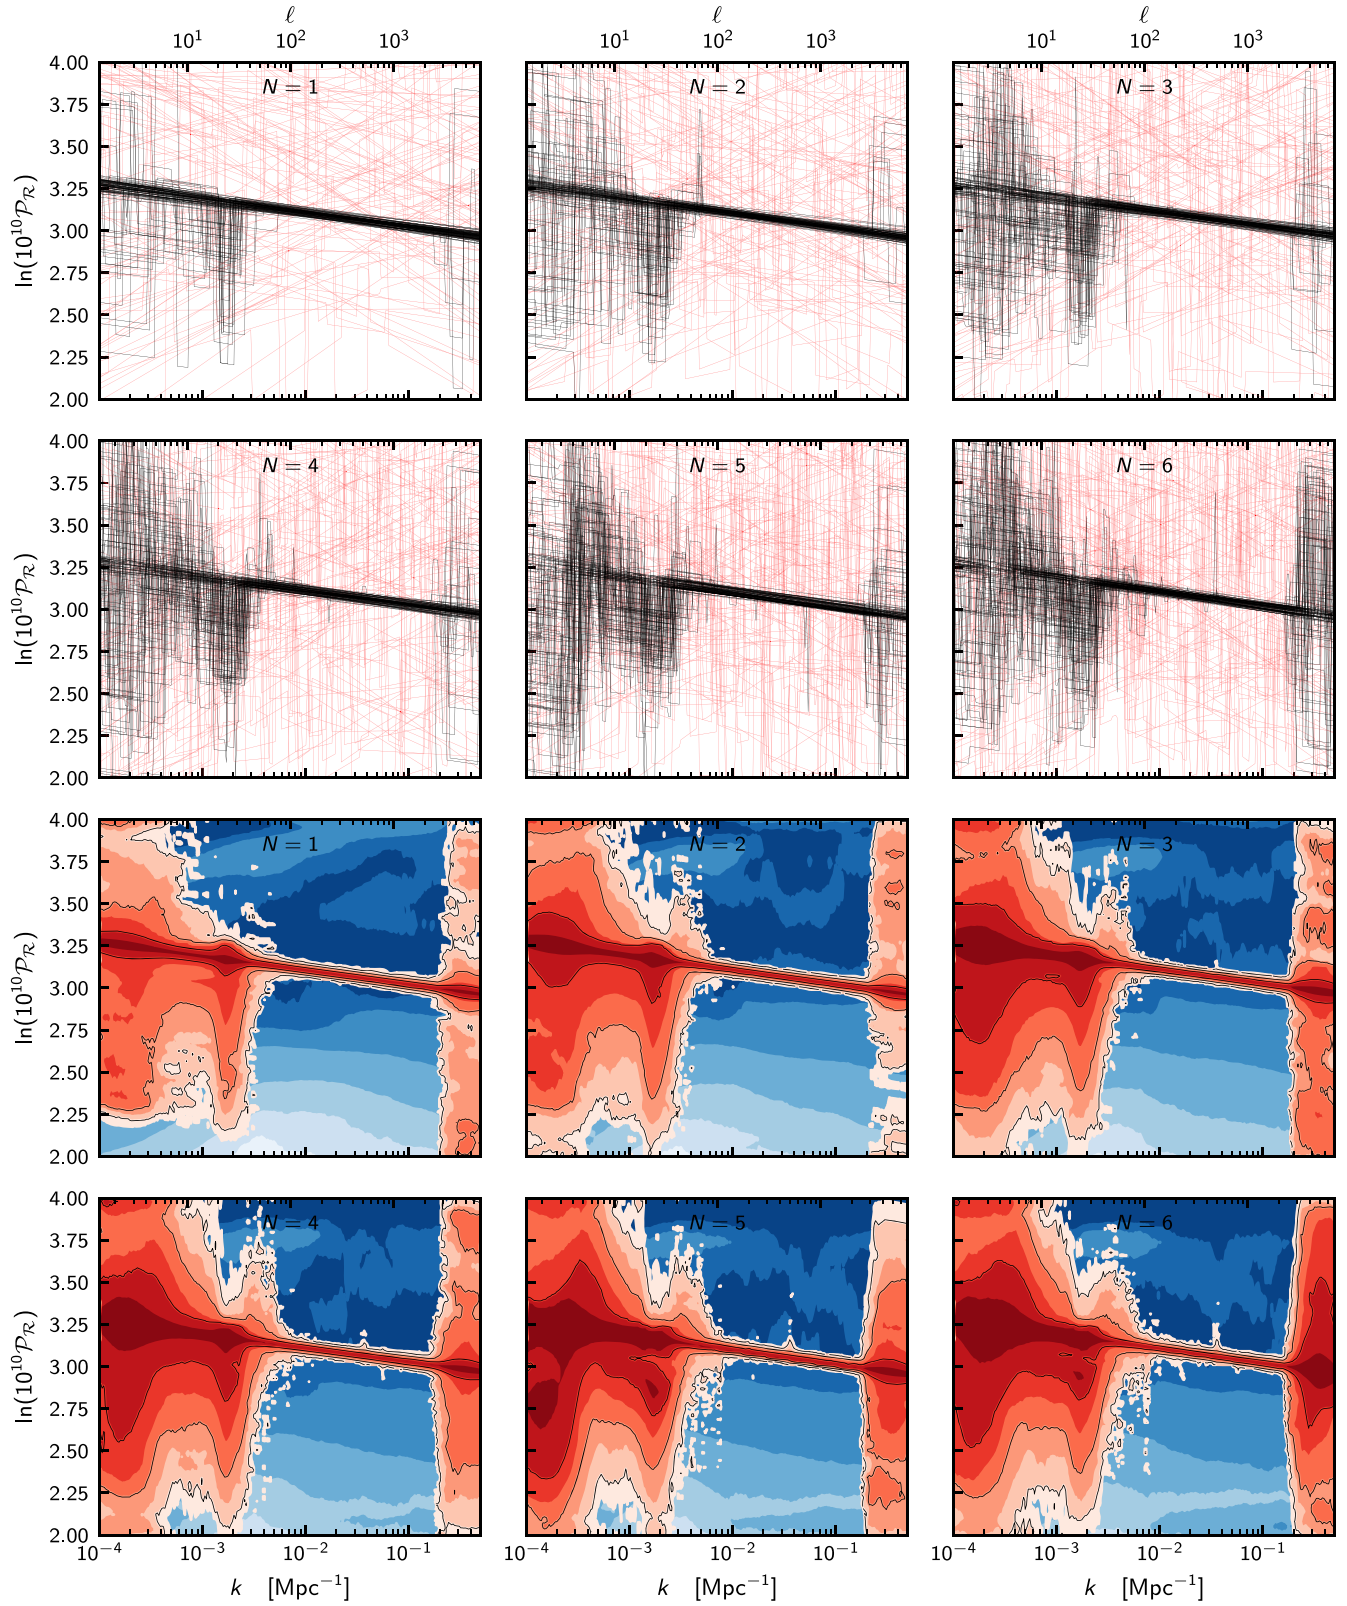


FIG. 13. Equally-weighted sample plots of sharp features reconstruction, conditioned on the number of knots N . $N = 0$ is exactly equivalent to a standard Λ CDM parametrization. Prior samples are drawn in red, whilst posterior samples are indicated in black.

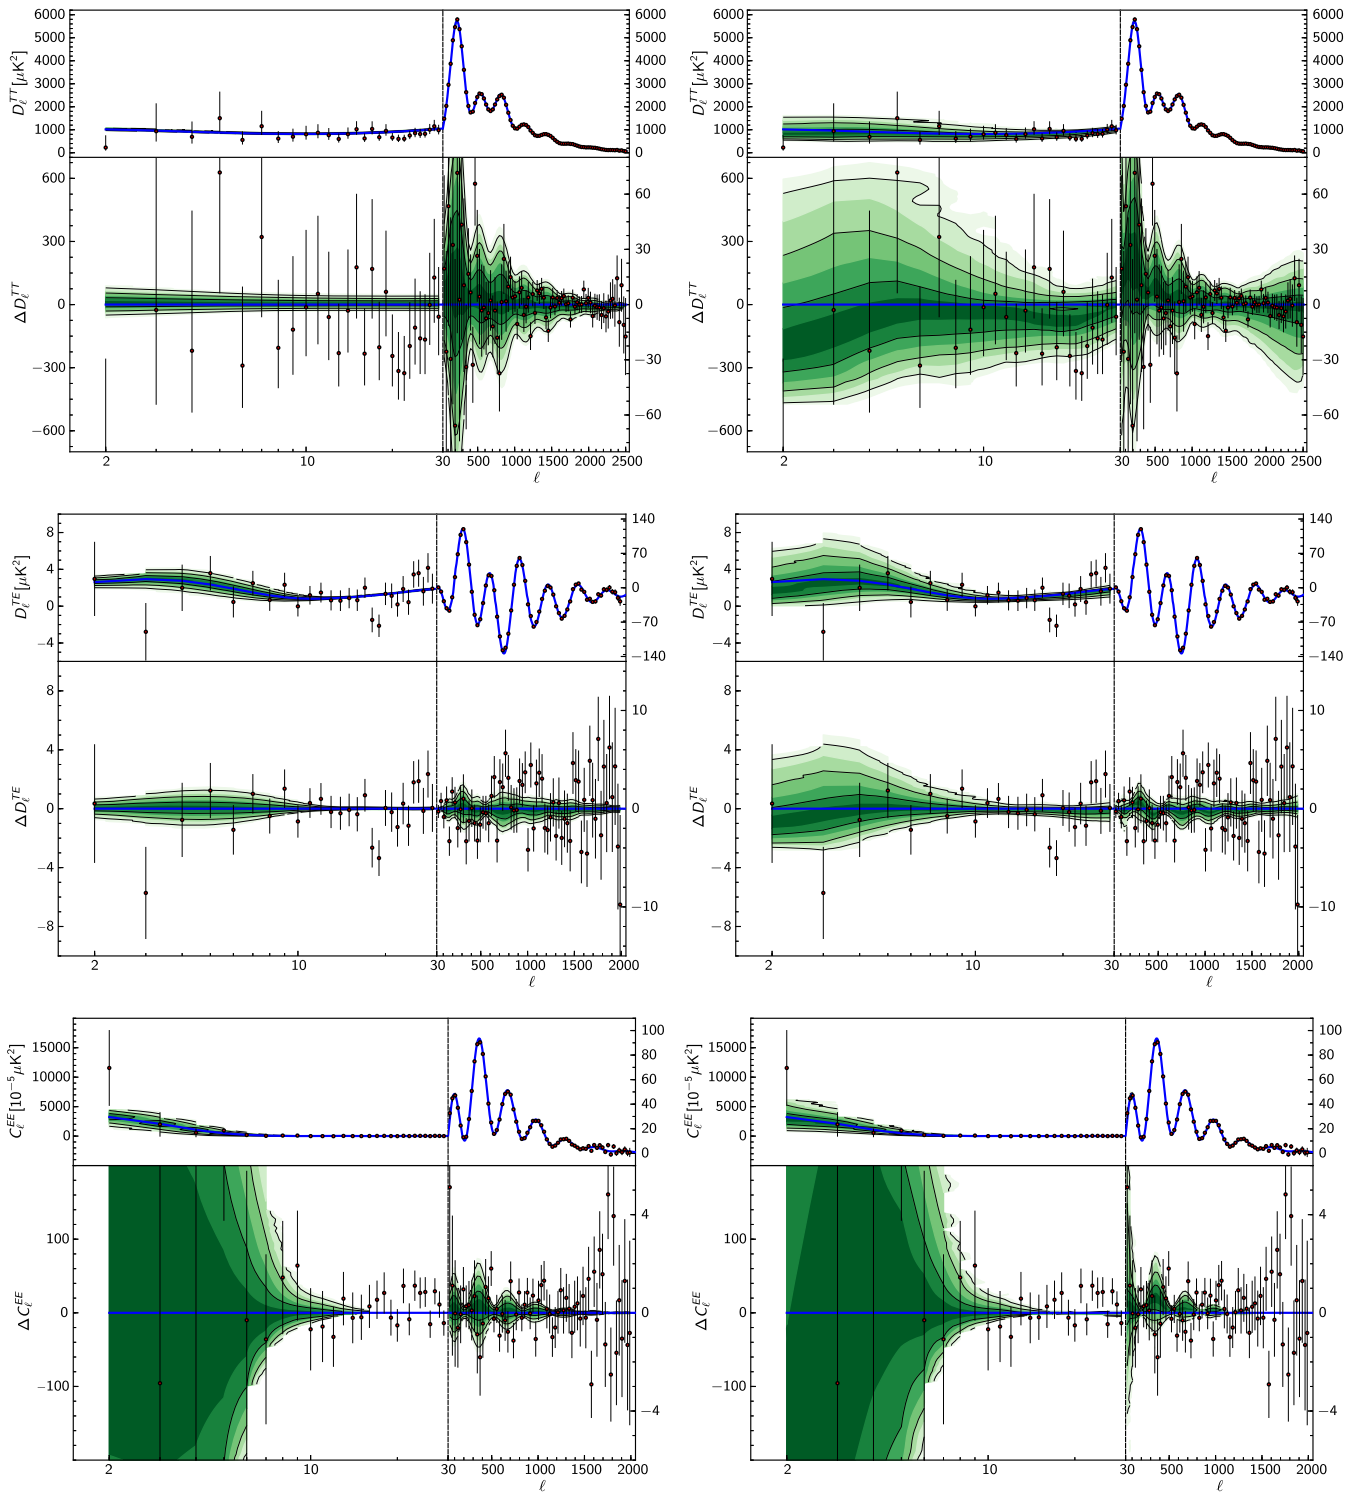


FIG. 14. Functional posterior distribution for the C_ℓ spectra from the sharp features reconstruction. Left hand column of panels is for $N = 0$, i.e., a Λ CDM parametrization. Right hand column of panels is for $N = 8$ features. Residual plots are with respect to the *Planck* 2018 best-fit Λ CDM cosmology.

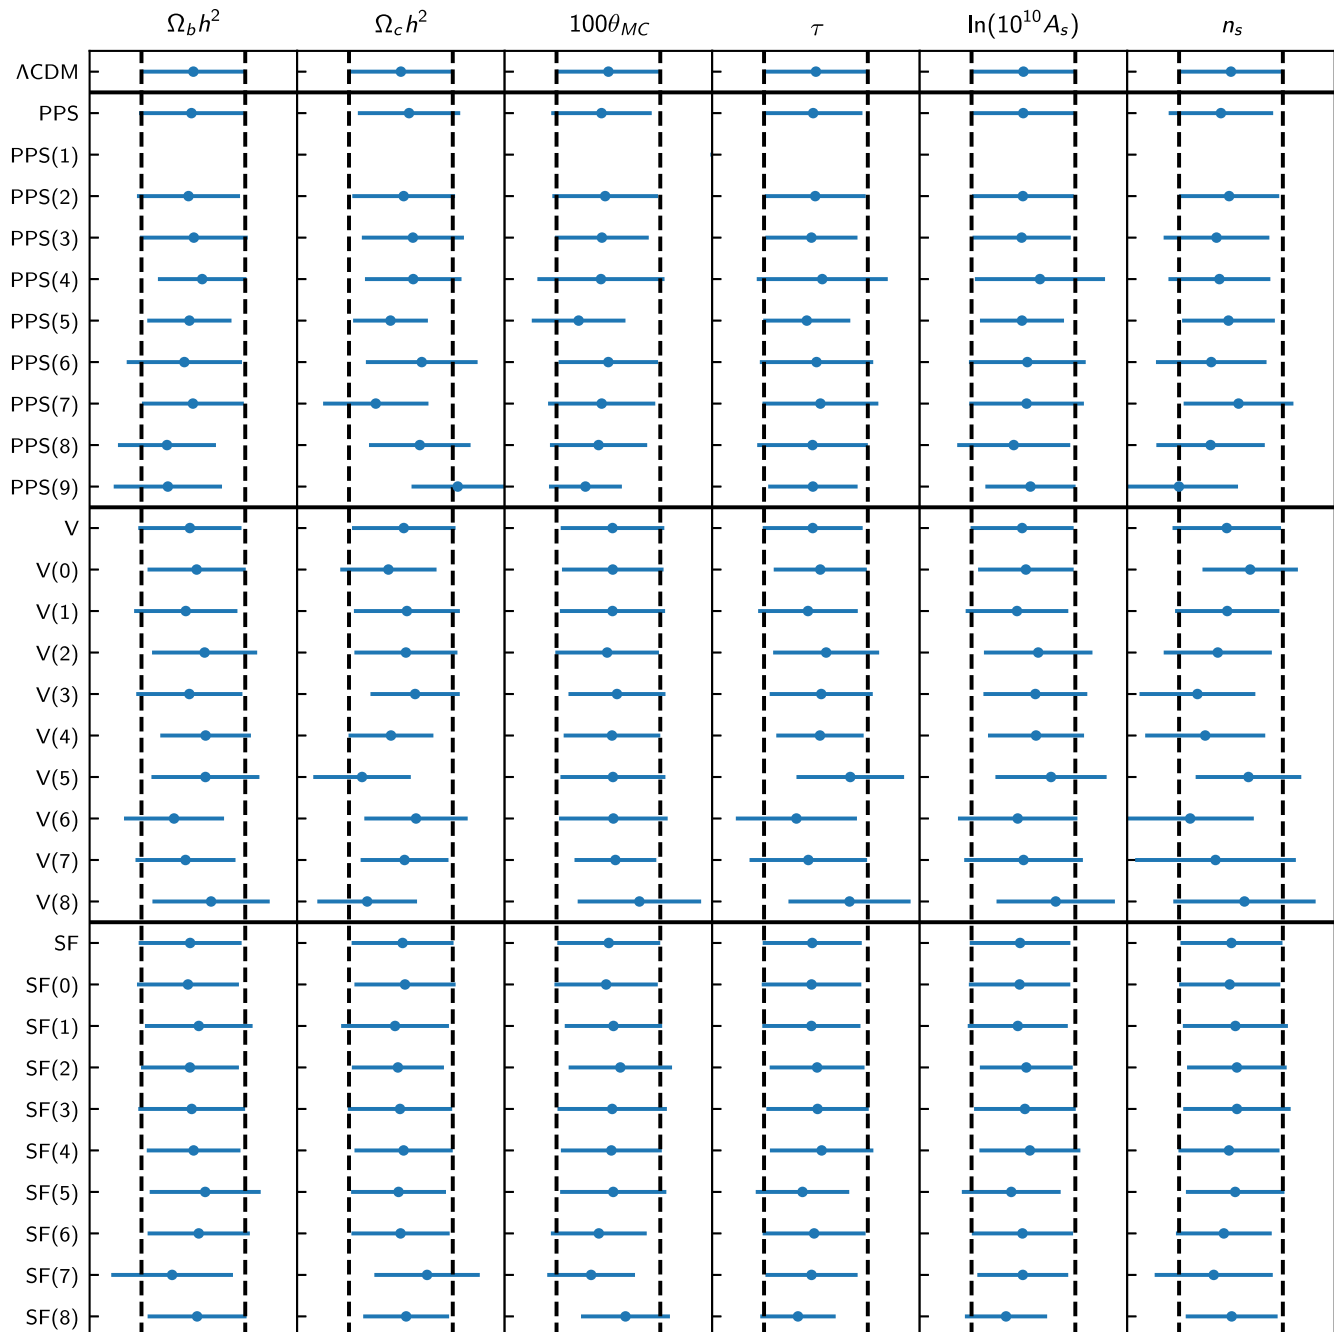


FIG. 15. Stability of the cosmological parameters for the primordial power spectrum reconstruction (PPS), the potential reconstruction (V) and the sharp features reconstruction (SF). Parameters are shown for the reconstructions conditioned on N , and for the marginalized case. For all cases except the highly-disfavored $N = 1$ PPS reconstruction (equivalent to an $n_s = 1$ scale-invariant power spectrum), the parameters agree with the baseline Λ CDM parameters. Note that for PPS and V the parameters n_s and A_s are derived parameters.

indistinguishable from its logarithmic prior. Given the slow roll relations, we find that for our reconstructions $\eta_V \approx \frac{d^2 \log V}{d\phi^2}$, so plots of η_V are nearly identical to plots of the equivalent underlying linear second-derivative reconstruction parameters.

V. SHARP FEATURE RECONSTRUCTION

In this section, we return to a direct analysis of the primordial power spectrum. Inspired by the oscillatory features present in both the linear spline primordial power spectrum reconstruction (Sec. III) and in the functional

posterior of the primordial power spectrum from the inflationary potential reconstruction (Sec. IV), we now consider a parametrization of the primordial power spectrum which favors sharp features. Attempts at explaining these features has a long history in the literature, initially investigated in Refs. [84,92–94].

We introduce sharp features into the parametrization of the spectrum by placing a variable number N of top-hat functions with varying widths Δ , heights h , and locations $\log_{10} k$ on top of the traditional A_s, n_s parametrization

$$\ln \mathcal{P}_{\mathcal{R}}(k) = \ln A_s + (n_s - 1) \ln \left(\frac{k}{k_*} \right) + \sum_{i=1}^N h_i \left[\left| \log_{10} k - \log_{10} k_i \right| < \frac{\Delta_i}{2} \right],$$

$$\Theta_{\mathcal{P}} = (\Delta_1, \dots, \Delta_N, h_1, \dots, h_N, k_1, \dots, k_N). \quad (21)$$

where the square brackets (as in Sec. II D) denote a logical truth function [51]. This parametrization is indicated schematically in Fig. 11, and Table IV provides a summary of the priors that we use. Readers are referred to the priors sections of Secs. III and IV for further details.

A. Results

The results for the sharp feature reconstructions can be found in Figs. 12–14. Figure 13 shows that in the marginalized plots the oscillations and lack of power spectrum constraints at low- and high- k are once again recovered. Visually, the features are even more striking in these reconstructions, on account of the ability for this parametrization to localize in k more precisely. There are also hints of features at high- k in this reconstruction, which were smoothed out by the parametrizations of the two previous sections.

Marginalization in Fig. 12 shows that there is little Bayesian evidence to support the introduction of more than two features, but the low- k oscillation still comes through clearly in the fully marginalized plot.

Finally, we examine the effects of these reconstructions on the C_{ℓ} spectra by considering the functional posterior in Fig. 14. By comparing the Λ CDM case $N = 0$ with the case $N = 8$ we see that the features at low- k in the PPS correspond to both a suppression of power at low- ℓ in the TT spectrum, as well as a more specific localized reduction of power in the $20 < \ell < 30$ region. There seems to be no obvious correspondence with possible features seen in the polarization TE or EE spectra.

VI. COSMOLOGICAL PARAMETER STABILITY

Finally in Fig. 15 we show that the underlying cosmological parameter constraints remain effectively unchanged for all three of the analyses in Secs. III to V, in spite of the

additional degrees of freedom we have given to the primordial power spectrum.

The only exception is the PPS $N = 1$ reconstruction. This model is highly-disfavored (Fig. 4), since *Planck* rules out a Harrison-Zeldovich scale-invariant ($n_s = 1$) spectrum [4,8,65]. Requiring $n_s = 1$ gives a poorly fit model which forces the cosmological parameters into locations discordant with Λ CDM.

This overall parameter stability for models consistent with the data demonstrates that one can explain features in CMB power spectra via modifications to the primordial cosmology, without the need to alter late-time cosmological parameters.

VII. CONCLUSIONS

In this work, we have reconstructed the primordial Universe three ways. In Sec. III, we reconstructed the primordial power spectrum using a linear spline. In Sec. IV, we reconstructed the inflationary potential using a cubic spline. In Sec. V, we probed sharp features in the primordial power spectrum by superimposing top-hat functions on top of the traditional Λ CDM power spectrum.

We showed that the Bayesian odds of a scale-invariant power spectrum are around a quintillion to one against in comparison to the Λ CDM cosmology. This agrees with the Planck Collaboration’s conclusions [4] that there is decisive evidence for $n_s \neq 1$ —one of the key predictions of the theory of inflation.

All methods reconstruct a featureless tilted power law consistent with a simple (A_s, n_s) parametrization across a broad observable window ($50 \lesssim \ell \lesssim 2000$). In addition, all reconstructions demonstrate that in a Bayesian sense it is preferable to have models which are able to recover the lack of power spectrum constraints at low- k due to cosmic variance, and at high- k due to *Planck* instrument noise, reflected in the evidences and marginalized plots (Figs. 4, 7 and 12).

All large N conditional reconstructions partially recover oscillatory features in the primordial power spectrum, with a peak at $\ell \sim 50$ and a trough at $20 < \ell < 30$, which manifest themselves in the functional posteriors of the C_{ℓ} spectra (Fig. 14). The inflationary potential reconstruction (Sec. IV) shows that this oscillation could be due to a breakdown in slow roll near the start of the inflationary window (Fig. 9), which is relevant for a wide variety of inflationary models [69–91]. However, the oscillations do not survive marginalization over N , indicating that the Bayesian evidence is not strong enough from *Planck* data to indicate a significant detection of such a feature.

The renewed upper bound on r from *Planck* 2018 now has enough discriminatory power to begin reconstructing potentials, as shown by the preference for the $N = 1$ case in the inflationary potential reconstructions.

As shown in Fig. 15, in all cases, the distributions on the late-time cosmological parameters remain unperturbed by

the additional degrees of freedom on the primordial cosmology provided by these reconstructions, indicating that any conclusions using late-time parameters are unlikely to be affected by modifying the primordial cosmology.

There is scope for inflationary models which *a-priori* predict these low- k features to be preferred over the Λ CDM cosmology, particularly if such models are capable of producing sharper features in the C_ℓ spectra at $20 < \ell < 30$. Additionally, in light of further CMB data [7], or failing that, strong τ characterization, it is likely that these hints of features will sharpen and provide further discriminatory power in constructing better models of the primordial Universe.

ACKNOWLEDGMENTS

W. J. H. was supported via STFC PhD studentship RG68795, the European Research Council (ERC) under the European Community's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement number 306478-CosmicDawn, and a Gonville & Caius College Research Fellowship. W. J. H. and A. N. L. thank Fabio Finelli, Martin Bucher and Julien Lesgourgues for numerous useful comments and suggestions for improvements

over the course of this work. H. V. P. acknowledges useful conversations with George Efstathiou, Matt Johnson, and Keir Rogers. H. V. P. was supported by the European Research Council (ERC) under the European Community's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement number 306478-CosmicDawn. This work was performed in part at the Aspen Center for Physics, which is supported by National Science Foundation Grant No. PHY-1607611. This work was also partially supported by a grant from the Simons Foundation. This work was performed using the Darwin Supercomputer of the University of Cambridge High Performance Computing Service, provided by Dell Inc. using Strategic Research Infrastructure Funding from the Higher Education Funding Council for England and funding from the Science and Technology Facilities Council, as well as resources provided by the Cambridge Service for Data Driven Discovery (CSD3) operated by the University of Cambridge Research Computing Service, provided by Dell EMC and Intel using Tier-2 funding from the Engineering and Physical Sciences Research Council (capital grant EP/P020259/1), and DiRAC funding from the Science and Technology Facilities Council.

-
- [1] Planck Collaboration, Planck 2018 results. I. Overview and the cosmological legacy of Planck, [arXiv:1807.06205](#).
 - [2] Planck Collaboration, Planck 2018 results. V. CMB power spectra and likelihoods, [arXiv:1907.12875](#).
 - [3] Planck Collaboration, Planck 2018 results. VI. Cosmological parameters, [arXiv:1807.06209](#).
 - [4] Planck Collaboration, Planck 2018 results. X. Constraints on inflation, *Astrophys. Space Sci.* **364**, 69 (2019).
 - [5] J. A. Vázquez, M. Bridges, M. P. Hobson, and A. N. Lasenby, Model selection applied to reconstruction of the primordial power spectrum, *J. Cosmol. Astropart. Phys.* **06** (2012) 006.
 - [6] G. Aslanyan, L. C. Price, K. N. Abazajian, and R. Easther, The Knotted Sky I: Planck constraints on the primordial power spectrum, *J. Cosmol. Astropart. Phys.* **08** (2014) 052.
 - [7] CORE Collaboration, Exploring cosmic origins with CORE: Inflation, *J. Cosmol. Astropart. Phys.* **04** (2018) 016.
 - [8] Planck Collaboration, Planck 2013 results. XXII. Constraints on inflation, *Astron. Astrophys.* **571**, A22 (2014).
 - [9] Planck Collaboration, Planck 2015 results. XX. Constraints on inflation, *Astron. Astrophys.* **594**, A20 (2016).
 - [10] S. Hee, J. A. Vázquez, W. J. Handley, M. P. Hobson, and A. N. Lasenby, Constraining the dark energy equation of state using Bayes theorem and the Kullback-Leibler divergence, *Mon. Not. R. Astron. Soc.* **466**, 369 (2017).
 - [11] J. A. Vázquez, M. Bridges, M. P. Hobson, and A. N. Lasenby, Reconstruction of the dark energy equation of state, *J. Cosmol. Astropart. Phys.* **09** (2012) 020.
 - [12] M. Millea and F. Bouchet, Cosmic microwave background constraints in light of priors over reionization histories, *Astron. Astrophys.* **617**, A96 (2018).
 - [13] M. Olamaie, M. P. Hobson, F. Feroz, K. J. B. Grainge, A. Lasenby, Y. C. Perrott, C. Rumsey, and R. D. E. Saunders, Free-form modelling of galaxy clusters: A Bayesian and data-driven approach, *Mon. Not. R. Astron. Soc.* **481**, 3853 (2018).
 - [14] C. Gauthier and M. Bucher, Reconstructing the primordial power spectrum from the CMB, *J. Cosmol. Astropart. Phys.* **10** (2012) 050.
 - [15] C. Dvorkin and W. Hu, Generalized slow roll approximation for large power spectrum features, *Phys. Rev. D* **81**, 023518 (2010).
 - [16] C. Dvorkin and W. Hu, CMB constraints on principal components of the inflaton potential, *Phys. Rev. D* **82**, 043513 (2010).
 - [17] C. Dvorkin and W. Hu, Complete WMAP constraints on band-limited inflationary features, *Phys. Rev. D* **84**, 063515 (2011).
 - [18] V. Miranda, W. Hu, and C. Dvorkin, Polarization predictions for inflationary CMB power spectrum features, *Phys. Rev. D* **91**, 063514 (2015).

- [19] G. Obied, C. Dvorkin, C. Heinrich, W. Hu, and V. Miranda, Inflationary features and shifts in cosmological parameters from Planck 2015 data, *Phys. Rev. D* **96**, 083526 (2017).
- [20] G. Obied, C. Dvorkin, C. Heinrich, W. Hu, and V. Miranda, Inflationary versus reionization features from Planck 2015 data, *Phys. Rev. D* **98**, 043518 (2018).
- [21] V. F. Mukhanov, H. A. Feldman, and R. H. Brandenberger, Theory of cosmological perturbations, *Phys. Rep.* **215**, 203 (1992).
- [22] D. Baumann, TASI lectures on inflation, [arXiv:0907.5424](https://arxiv.org/abs/0907.5424).
- [23] D. Baumann, TASI lectures on primordial cosmology, [arXiv:1807.03098](https://arxiv.org/abs/1807.03098).
- [24] M. J. Mortonson, H. V. Peiris, and R. Easther, Bayesian analysis of inflation: Parameter estimation for single field models, *Phys. Rev. D* **83**, 043505 (2011).
- [25] R. Easther and H. V. Peiris, Bayesian analysis of inflation. II. Model selection and constraints on reheating, *Phys. Rev. D* **85**, 103533 (2012).
- [26] J. Noreña, C. Wagner, L. Verde, H. V. Peiris, and R. Easther, Bayesian analysis of inflation. III. Slow roll reconstruction using model selection, *Phys. Rev. D* **86**, 023505 (2012).
- [27] C. Armendáriz-Picón and E. A. Lim, Vacuum choices and the predictions of inflation, *J. Cosmol. Astropart. Phys.* **12** (2003) 006.
- [28] W. J. Handley, A. N. Lasenby, and M. P. Hobson, Novel quantum initial conditions for inflation, *Phys. Rev. D* **94**, 024041 (2016).
- [29] W. Handley, Primordial power spectra for curved inflating universes, [arXiv:1907.08524](https://arxiv.org/abs/1907.08524).
- [30] F. A. Agocs, L. T. Hergt, W. H. Handley, A. N. Lasenby, and M. P. Hobson, Investigating the gauge invariance of quantum initial conditions for inflation, *Phys. Rev. D* (to be published).
- [31] Toronto: Maplesoft, a division of Waterloo Maple Inc. Maple User Manual, Maplesoft, 2018.
- [32] The Numerical Algorithms Group (NAG), Oxford, United Kingdom. The NAG Library, <http://www.nag.com>.
- [33] A. Lewis, A. Challinor, and A. Lasenby, Efficient computation of CMB anisotropies in closed FRW models, *Astrophys. J.* **538**, 473 (2000).
- [34] D. Blas, J. Lesgourgues, and T. Tram, The Cosmic Linear Anisotropy Solving System (CLASS) II: Approximation schemes, *J. Cosmol. Astropart. Phys.* **07** (2011) 034.
- [35] W. J. Handley, Cosmochord 1.15, <http://dx.doi.org/10.5281/zenodo.2552056> (2019).
- [36] A. Lewis and S. Bridle, Cosmological parameters from CMB and other data: A Monte Carlo approach, *Phys. Rev. D* **66**, 103511 (2002).
- [37] A. Lewis, Efficient sampling of fast and slow cosmological parameters, *Phys. Rev. D* **87**, 103529 (2013).
- [38] J. Skilling, Nested sampling for general bayesian computation, *Bayesian Anal.* **1**, 833 (2006).
- [39] W. J. Handley, M. P. Hobson, and A. N. Lasenby, POLYCHORD: Nested sampling for cosmology, *Mon. Not. R. Astron. Soc.* **450**, L61 (2015).
- [40] W. J. Handley, M. P. Hobson, and A. N. Lasenby, POLYCHORD: Next-generation nested sampling, *Mon. Not. R. Astron. Soc.* **453**, 4384 (2015).
- [41] F. Feroz and M. P. Hobson, Multimodal nested sampling: An efficient and robust alternative to Markov Chain Monte Carlo methods for astronomical data analyses, *Mon. Not. R. Astron. Soc.* **384**, 449 (2008).
- [42] F. Feroz, M. P. Hobson, and M. Bridges, MULTINEST: An efficient and robust Bayesian inference tool for cosmology and particle physics, *Mon. Not. R. Astron. Soc.* **398**, 1601 (2009).
- [43] F. Feroz, M. P. Hobson, E. Cameron, and A. N. Pettitt, Importance nested sampling and the MultiNest algorithm, [arXiv:1306.2144](https://arxiv.org/abs/1306.2144).
- [44] W. Handley, Fgiveness: Functional posterior plotter, *J. Open Source Software* **3**, 849 (2018).
- [45] S. Kullback and R. A. Leibler, On information and sufficiency, *Ann. Math. Stat.* **22**, 79 (1951).
- [46] A. Hosoya, T. Buchert, and M. Morita, Information Entropy in Cosmology, *Phys. Rev. Lett.* **92**, 141302 (2004).
- [47] L. Verde, P. Protopapas, and R. Jimenez, Planck and the local Universe: Quantifying the tension, *Phys. Dark Universe* **2**, 166 (2013).
- [48] S. Grandis, S. Seehars, A. Refregier, A. Amara, and A. Nicola, Information gains from cosmological probes, *J. Cosmol. Astropart. Phys.* **05** (2016) 034.
- [49] M. Raveri, M. Martinelli, G. Zhao, and Y. Wang, Information gain in cosmology: From the discovery of expansion to future surveys, [arXiv:1606.06273](https://arxiv.org/abs/1606.06273).
- [50] S. Hee, W. J. Handley, M. P. Hobson, and A. N. Lasenby, Bayesian model selection without evidences: Application to the dark energy equation-of-state, *Mon. Not. R. Astron. Soc.* **455**, 2461 (2016).
- [51] R. L. Graham, D. E. Knuth, and O. Patashnik, *Concrete Mathematics: A Foundation for Computer Science*, 2nd ed. (Addison-Wesley Longman Boston, MA, 1994).
- [52] R. Trotta, Applications of Bayesian model selection to cosmological parameters, *Mon. Not. R. Astron. Soc.* **378**, 72 (2007).
- [53] W. Handley and P. Lemos, Quantifying tensions in cosmological parameters: Interpreting the DES evidence ratio, *Phys. Rev. D* **100**, 043504 (2019).
- [54] R. D. Hall, S. J. Thompson, W. Handley, and D. Queloz, On the feasibility of intense radial velocity surveys for Earth-Twin discoveries, *Mon. Not. R. Astron. Soc.* **479**, 2968 (2018).
- [55] R. Buscicchio, E. Roebber, J. M. Goldstein, and C. J. Moore, The label switching problem in Bayesian analysis for gravitational wave astronomy, *Phys. Rev. D* **100**, 084041 (2019).
- [56] C. L. Bennett, A. J. Banday, K. M. Gorski, G. Hinshaw, P. Jackson, P. Keegstra, A. Kogut, G. F. Smoot, D. T. Wilkinson, and E. L. Wright, Four-year COBE DMR cosmic microwave background observations: Maps and basic results, *Astrophys. J. Lett.* **464**, L1 (1996).
- [57] C. B. Netterfield *et al.*, A measurement by BOOMERANG of multiple peaks in the angular power spectrum of the cosmic microwave background, *Astrophys. J.* **571**, 604 (2002).
- [58] S. Hanany *et al.*, MAXIMA-1: A measurement of the cosmic microwave background anisotropy on angular scales of 10'-5deg, *Astrophys. J. Lett.* **545**, L5 (2000).
- [59] N. W. Halverson, E. M. Leitch, C. Pryke, J. Kovac, J. E. Carlstrom, W. L. Holzapfel, M. Dragovan, J. K. Cartwright, B. S. Mason, S. Padin, T. J. Pearson, A. C. S. Readhead, and

- M. C. Shepherd, Degree angular scale interferometer first results: A measurement of the cosmic microwave background angular power spectrum, *Astrophys. J.* **568**, 38 (2002).
- [60] P. F. Scott *et al.*, First results from the very small array—III. The cosmic microwave background power spectrum, *Mon. Not. R. Astron. Soc.* **341**, 1076 (2003).
- [61] T. J. Pearson *et al.*, The anisotropy of the microwave background to $l = 3500$: Mosaic observations with the cosmic background imager, *Astrophys. J.* **591**, 556 (2003).
- [62] C. L. Bennett *et al.*, Nine-year Wilkinson microwave anisotropy probe (WMAP) observations: Final maps and results, *Astrophys. J. Suppl. Ser.* **208**, 20 (2013).
- [63] G. Hinshaw *et al.*, Nine-year Wilkinson microwave anisotropy probe (WMAP) observations: Cosmological parameter results, *Astrophys. J. Suppl. Ser.* **208**, 19 (2013).
- [64] P. A. R. Ade, N. Aghanim, M. I. R. Alves, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, H. Aussel *et al.* (Planck Collaboration), Planck 2013 results. I. Overview of products and scientific results, *Astron. Astrophys.* **571**, A1 (2014).
- [65] R. Adam, P. A. R. Ade, N. Aghanim, Y. Akrami, M. I. R. Alves, F. Argüeso, M. Arnaud, F. Arroja, M. Ashdown *et al.* (Planck Collaboration), Planck 2015 results. I. Overview of products and scientific results, *Astron. Astrophys.* **594**, A1 (2016).
- [66] A. R. Liddle, P. Parsons, and J. D. Barrow, Formalizing the slow-roll approximation in inflation, *Phys. Rev. D* **50**, 7222 (1994).
- [67] J. Lesgourgues and W. Valkenburg, New constraints on the observable inflaton potential from WMAP and SDSS, *Phys. Rev. D* **75**, 123519 (2007).
- [68] J. Lesgourgues, A. A. Starobinsky, and W. Valkenburg, What do WMAP and SDSS really tell us about inflation?, *J. Cosmol. Astropart. Phys.* **01** (2008) 010.
- [69] E. Ramirez and D. J. Schwarz, Predictions of just-enough inflation, *Phys. Rev. D* **85**, 103516 (2012).
- [70] W. J. Handley, S. D. Brechet, A. N. Lasenby, and M. P. Hobson, Kinetic initial conditions for inflation, *Phys. Rev. D* **89**, 063505 (2014).
- [71] L. T. Hergt, W. J. Handley, M. P. Hobson, and A. N. Lasenby, Case for kinetically dominated initial conditions for inflation, *Phys. Rev. D* **100**, 023502 (2019).
- [72] L. T. Hergt, W. J. Handley, M. P. Hobson, and A. N. Lasenby, Constraining the kinetically dominated universe, *Phys. Rev. D* **100**, 023501 (2019).
- [73] A. A. Starobinskij, Spectrum of adiabatic perturbations in the universe when there are singularities in the inflationary potential, *Sov. J. Exp. Theor. Phys. Lett.* **55**, 489 (1992).
- [74] J. Adams, B. Cresswell, and R. Easther, Inflationary perturbations from a potential with a step, *Phys. Rev. D* **64**, 123514 (2001).
- [75] M. Joy, V. Sahni, and A. A. Starobinsky, New universal local feature in the inflationary perturbation spectrum, *Phys. Rev. D* **77**, 023514 (2008).
- [76] J. Silk and M. S. Turner, Double inflation, *Phys. Rev. D* **35**, 419 (1987).
- [77] R. Holman, E. W. Kolb, S. L. Vadas, and Y. Wang, Plausible double inflation, *Phys. Lett. B* **269**, 252 (1991).
- [78] D. Polarski and A. A. Starobinsky, Spectra of perturbations produced by double inflation with an intermediate matter-dominated stage, *Nucl. Phys.* **B385**, 623 (1992).
- [79] J. A. Adams, G. G. Ross, and S. Sarkar, Multiple inflation, *Nucl. Phys.* **B503**, 405 (1997).
- [80] P. Hunt and S. Sarkar, Multiple inflation and the WMAP “glitches”, *Phys. Rev. D* **70**, 103518 (2004).
- [81] J. Lesgourgues, Features in the primordial power spectrum of double D-term inflation, *Nucl. Phys.* **B582**, 593 (2000).
- [82] C. P. Burgess, R. Easther, A. Mazumdar, D. F. Mota, and T. Multamäki, Multiple inflation, cosmic string networks and the string landscape, *J. High Energy Phys.* **05** (2005) 067.
- [83] A. Ashoorioon and A. Krause, Power Spectrum and Signatures for Cascade Inflation, [arXiv:hep-th/0607001](https://arxiv.org/abs/hep-th/0607001).
- [84] H. V. Peiris, E. Komatsu, L. Verde, D. N. Spergel, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, G. S. Tucker, E. Wollack, and E. L. Wright, First-year Wilkinson microwave anisotropy probe (WMAP) observations: Implications for inflation, *Astrophys. J. Suppl. Ser.* **148**, 213 (2003).
- [85] L. Covi, J. Hamann, A. Melchiorri, A. Slosar, and I. Sorbera, Inflation and WMAP three year data: Features are still present, *Phys. Rev. D* **74**, 083509 (2006).
- [86] J. Hamann, L. Covi, A. Melchiorri, and A. Slosar, New constraints on oscillations in the primordial spectrum of inflationary perturbations, *Phys. Rev. D* **76**, 023503 (2007).
- [87] P. Hunt and S. Sarkar, Multiple inflation and the WMAP “glitches”. II. Data analysis and cosmological parameter extraction, *Phys. Rev. D* **76**, 123504 (2007).
- [88] M. Joy, A. Shafieloo, V. Sahni, and A. A. Starobinsky, Is a step in the primordial spectral index favoured by CMB data?, *J. Cosmol. Astropart. Phys.* **06** (2009) 028.
- [89] J. Martin and C. Ringeval, Superimposed oscillations in the WMAP data?, *Phys. Rev. D* **69**, 083515 (2004).
- [90] M. Kawasaki, F. Takahashi, and T. Takahashi, Making waves on CMB power spectrum and inflaton dynamics, *Phys. Lett. B* **605**, 223 (2005).
- [91] R. K. Jain, P. Chingangbam, J.-O. Gong, L. Sriramkumar, and T. Souradeep, Punctuated inflation and the low CMB multipoles, *J. Cosmol. Astropart. Phys.* **01** (2009) 009.
- [92] G. Hinshaw, D. N. Spergel, L. Verde, R. S. Hill, S. S. Meyer, C. Barnes, C. L. Bennett, M. Halpern, N. Jarosik, A. Kogut, E. Komatsu, M. Limon, L. Page, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright *et al.*, First-year Wilkinson microwave anisotropy probe (WMAP) observations: The angular power spectrum, *Astrophys. J. Suppl. Ser.* **148**, 135 (2003).
- [93] D. N. Spergel, L. Verde, H. V. Peiris, E. Komatsu, M. R.olta, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright, First-year Wilkinson microwave anisotropy probe (WMAP) observations: Determination of cosmological parameters, *Astrophys. J. Suppl. Ser.* **148**, 175 (2003).
- [94] M. J. Mortonson, C. Dvorkin, H. V. Peiris, and W. Hu, CMB polarization features from inflation versus reionization, *Phys. Rev. D* **79**, 103519 (2009).