The Knotted Sky I: Planck constraints on the primordial power spectrum

Grigor Aslanyan,[a] **Layne C. Price,**[a] **Kevork N. Abazajian,**[b] **and Richard Easther**[a]

[a]Department of Physics, University of Auckland, Private Bag 92019, Auckland, New Zealand
[b]Department of Physics, University of California at Irvine, Irvine, CA 92697

E-mail: g.aslanyan@auckland.ac.nz, lpri691@aucklanduni.ac.nz, kevork@uci.edu, r.easther@auckland.ac.nz

Abstract. Using the temperature data from *Planck* we search for departures from a power-law primordial power spectrum, employing Bayesian model-selection and posterior probabilities. We parametrize the spectrum with n knots located at arbitrary values of $\log k$, with both linear and cubic splines. This formulation recovers both slow modulations and sharp transitions in the primordial spectrum. The power spectrum is well-fit by a featureless, power-law at wavenumbers $k > 10^{-3}\,\mathrm{Mpc}^{-1}$. A modulated primordial spectrum yields a better fit relative to ΛCDM at large scales, but there is no strong evidence for a departure from a power-law spectrum. Moreover, using simulated maps we show that a local feature at $k \sim 10^{-3}\,\mathrm{Mpc}^{-1}$ can mimic the suppression of large-scale power. With multi-knot spectra we see only small changes in the posterior distributions for the other free parameters in the standard ΛCDM universe. Lastly, we investigate whether the hemispherical power asymmetry is explained by independent features in the primordial power spectrum in each ecliptic hemisphere, but find no significant differences between them.

Contents

1 Introduction

The first cosmological data analysis by the *Planck* Science Team [1] confirmed the conventional ΛCDM model of cosmology with unprecedented precision. In particular, a scale-invariant primordial power spectrum (PPS) is excluded at $> 5\sigma$. Simple models of single-field inflation generically yield an almost scale-invariant PPS, but no inflationary models are yet favoured by Bayesian evidence relative to ΛCDM [2–5]. Conversely, models with relatively complex spectra, including oscillations or localized amplifications, are consistent with current cosmological data [6–11].

While it is clear that the Harrison-Zel'dovich spectrum does not provide an optimal fit to the data, it does not follow that the power-law PPS is preferred over all other possible forms. Furthermore, as cosmological constraints become sensitive to increasingly delicate signals in the PPS, it is important to check whether constraints on these parameters depend on the assumed form of the PPS. Model-independent approaches to reconstructing the PPS have been widely studied [12–39], and the approached used here closely parallels that of Ref. [36], which examines the seven year WMAP dataset.

We revisit this problem using *Planck* data, Bayesian model-selection based on evidence (or *marginalized likelihood*) ratios [33–38] and a flexible specification for the PPS. We use this formalism to test whether *Planck* constraints on cosmological parameters are weakened when permit a generic PPS, rather than usual, almost–scale-invariant power-law formulation. While parameter degeneracies with the PPS could, in principle affect the posteriors on other cosmological variables (e.g., [40]), we find that the constraints on these parameters do not change significantly when we allow generic forms of the PPS. Secondly, we use this formalism to determine whether the observed large-scale hemispherical asymmetry in the two ecliptic hemispheres can be attributed to differences in the form of the PPS. We find no difference in the structure of the power spectrum in the two hemispheres, either qualitatively or in the
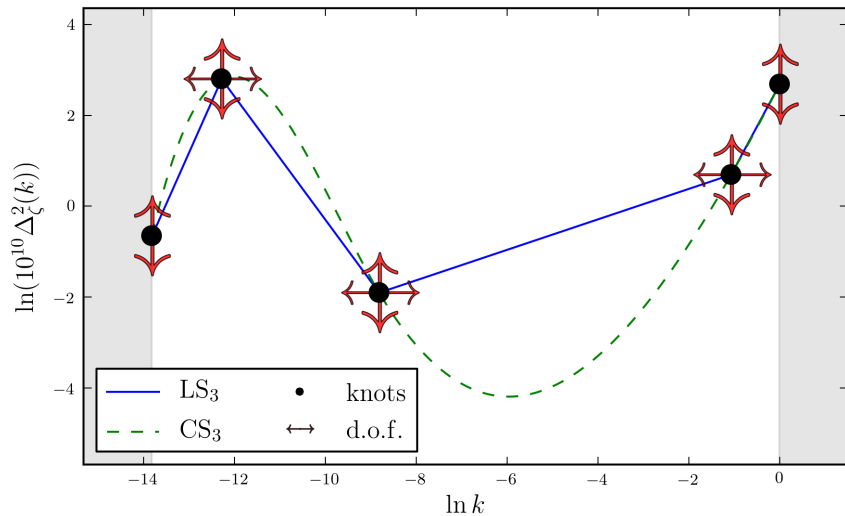
Figure 1. We illustrate the linear-spline LS_n and cubic-spline CS_n models for the primordial power spectrum, where n is the number of knots between the two endpoints. The white region denotes the range of k for which the spectrum is defined, with $k_{\min} = 10^{-6}\,\mathrm{Mpc}^{-1}$ and $k_{\max} = 1.0\,\mathrm{Mpc}^{-1}$. There are $2n + 2$ degrees of freedom for each choice of spline, since we vary the amplitude Δ_ζ^2 at the endpoints and allow the knots to move in both Δ_ζ^2 and k.

evidence ratios. Finally, the algorithm described here was implemented in COSMO++ [41], which is publicly available.

We compare Bayesian evidence for the non–power-law models to the evidence for the red-tilted PPS of ΛCDM. As we add more parameters to the PPS the Bayesian evidence does not change significantly, indicating the data cannot substantially distinguish between these models. However, most of the extra knots appear in the long wavelength section of the power spectrum, with $k \lesssim 10^{-3}\,\mathrm{Mpc}^{-1}$, suggesting that smaller scales are indeed well described by a power-law PPS.

Since no model-selection method can be completely non-parametric, we check our analysis by obtaining posterior probabilities for two different styles of non–power-law PPS. We compare both a linear- and a cubic-spline interpolation model, which capture sharp and smooth features in the PPS, respectively. The two models are illustrated in Fig. 1 and explained in detail in Section 2.3. We allow variation in the number of knots, their amplitudes, their positions in k-space, and the endpoint amplitudes. We see the maximum increase in evidence is $\Delta \ln Z = 0.7$ for the one knot linear spline model with varying foregrounds and $\Delta \ln Z = 2.2$ for the five knot linear spline, albeit with the foreground parameters in the Planck likelihood fixed to their best-fit values.

To test the effectiveness of our PPS parametrizations we attempt to recover nontrivial signals in simulated CMB temperature maps. The method clearly finds even small added features in the PPS, while the evidence ratio strongly favors models with added interior knots when the simulated feature is large enough. The analysis here considers only *Planck* data. In a separate paper we will consider the implications of the recent BICEP2 B-mode polarization detection [42, 43] for the scalar power spectrum using the knot-spline techniques developed here.

2 Modeling the primordial power spectrum

2.1 Bayesian evidence and model-selection

Cosmological data is intrinsically stochastic, so overfitting is a key danger for empirical reconstructions of the PPS. While many strategies have been suggested to prevent "fitting to noise", the Bayesian approach naturally disfavors models that yield a small region with high likelihood in a much larger multidimensional parameter space.

To compare two models M_i and M_j given data D, we evaluate the ratio of their posterior probabilities

$$K_{ij} = \frac{P(M_i \,|\, D)}{P(M_j \,|\, D)} = \frac{P(M_i)Z_i(D)}{P(M_j)Z_j(D)}\,, \tag{2.1}$$

where Z_i is the Bayesian evidence, or marginalized likelihood, given in terms of parameters θ and the data likelihood $\mathcal{L}_i(D\,|\,\theta) \equiv P(D\,|\,\theta, M_i)$ by

$$Z_i(D) \equiv P(D\,|\,M_i) = \int P(\theta\,|\,M_i)\mathcal{L}_i(D\,|\,\theta)d\theta. \tag{2.2}$$

The prior probabilities $P(M_i)$ and $P(\theta\,|\,M_i) \equiv P_i(\theta)$ incorporate our *a priori* information about the model M_i and the model's parameters θ, respectively. The evidence ratio or *Bayes factor* $B_{ij} = Z_i/Z_j$, is the ratio of the posterior model odds divided by the prior model odds and measures the extent to which the data "prefers" M_i over M_j. If the model priors are equal, $B_{ij} = K_{ij}$ and the Bayes' factor directly measures the posterior model odds. Generically, the integral in Eq. (2.2) is numerically challenging, but it is rendered tractable by multimodal nested sampling. We implement this algorithm using the MULTINEST package [44–46], which is widely employed for this purpose within cosmology.

Evidence ratios can be judged qualitatively using the Jeffreys' scale [47] or a more conservative "cosmology scale" [48], collated in Table 1. Negative evidences can be assessed by inverting the ratio before taking the logarithm in Eq. (2.1). Bayesian model-selection is typically more conservative than frequentist hypothesis testing [49]. For example, it is always possible to postulate a very complex PPS that will perfectly fit the data, thereby trivially maximizing $\mathcal{L}(D)$. In fact, it is possible to obtain a significant improvement in χ^2 or log-likelihood by letting the PPS have a lot of freedom [13, 17–19, 21–23, 25], but this usually results in a complex power spectrum with many curious wiggles and oscillations. However, in these circumstances it is likely that the best fit corresponds to a small subregion of the overall parameter space, and Bayesian evidence penalizes scenarios of this form.

While the Bayes factor quantifies the relative evidence between two models, it is not an "all-purpose tool." Bayesian model selection is often characterized as an implementation of Occam's razor, as it penalizes models with many parameters. In fact, a close inspection of Eq. (2.2) shows that it penalizes models that do not give sufficiently high likelihood over the integration volume. If a model adds parameters θ' that are only weakly-constrained by the likelihood $\mathcal{L}(\theta') \sim 1$, the integrals over θ' in Eq. (2.2) factorise and cancel when evidence ratios are computed. Furthermore, if we make the common choice that the prior probabilities for the models are equal, Bayesian model selection is not testing against a "null hypothesis", but comparing two models; a Bayes factor $|B_{ij}| \lesssim 1.0$ simply expresses the fact that the data used to construct the likelihood has no strong preference for either model. The simpler model, *i.e.*, the one with fewer parameters, is only preferred if we give it more *a priori* weight. In this paper we report both our quantitative calculations for B_{ij} and a qualitative analysis of the posterior probabilities on the model parameters.

$\ln K_{ij}$	K_{ij}	Jeffreys Scale	Cosmology Scale
0.0 to 1.0	1.0 to 2.7	Not worth more than a bare mention	
1.0 to 2.5	2.7 to 12.2	Substantial	Weak
2.5 to 5.0	12.2 to 148.4	Strong	Significant
> 5	> 148.4	Decisive	Strong

Table 1. Rough guideline for Bayesian evidence interpretation with the Jeffreys scale [47] and the re-scaled "cosmology scale" from Ref. [48]. Assuming two models $\mathcal{M}_i$ and $\mathcal{M}_j$ have the same *prior probability*, the Bayes factor B_{ij} is equivalent to the ratio of posterior probabilities K_{ij} and can be interpreted directly as the *posterior betting odds* for $\mathcal{M}_i$ over $\mathcal{M}_j$. The more conservative scale emphasizes that in cosmology there is often substantial uncertainty in the choice and form of model priors and the resulting evidences should therefore be interpreted more carefully.

While the Bayes factor B_{ij} is one tool we can use to perform model-selection, our parameterization of the non–power-law model still plays a major role in the Bayesian evidence integral of Eq. (2.2). An ideal non–power-law PPS parameterization needs to be able to construct any type of *possible* deviation from a power-law PPS. This is not practical, since this requires an infinite-dimensional parameter space. There are too many possible forms for the PPS to explore and we inevitably must truncate the allowed degrees of freedom. Instead, we have to settle for a model that can reconstruct a wide range of *probable* PPS features. This clearly depends on what features we expect to be present in the data, which introduces the possibility of confirmation bias into the analysis.

2.2 The power-law power spectrum

As reviewed in Ref. [50], the primordial scalar perturbations can be expressed in terms of the gauge invariant curvature perturbations $\zeta(\mathbf{x})$ in real space. In Fourier space they become

$$\zeta(\mathbf{k}) = \int d^3x \, e^{-i\mathbf{k}\cdot\mathbf{x}} \zeta(\mathbf{x}) \,. \tag{2.3}$$

If the perturbations are statistically homogeneous, the two-point correlation function is

$$\langle \zeta(\mathbf{k})\zeta^*(\mathbf{k}') \rangle = (2\pi)^3 \delta^3(\mathbf{k} - \mathbf{k}') \left(\frac{2\pi^2}{k^3} \Delta_\zeta^2(k) \right) \tag{2.4}$$

where δ^3 denotes the delta-function in three dimensions and Δ_ζ^2 is the dimensionless *primordial power spectrum (PPS)*.

If the perturbations are Gaussian, they are fully described by Eq. (2.4). Given that no primordial non-Gaussianity has been detected [1, 51] we restrict our analysis to the Gaussian case. In the standard ΛCDM model the PPS is described by a simple power law with two free parameters

$$\Delta_\zeta^2(k) = A_s \left(\frac{k}{k_*} \right)^{n_s - 1} \,, \tag{2.5}$$

where A_s is the *amplitude* of the scalar perturbations; n_s is the *scalar spectral index*; and k_* is the *pivot scale*. The latest CMB temperature data from *Planck* [1], CMB polarization data from WMAP [52], small scale temperature data from ACT [53] and SPT [54], as well

as the large scale structure data are all in excellent agreement with this assumption. These data sets in combination [1] put the following bounds on the PPS parameters: $\ln(10^{10}A_s) = 3.091 \pm 0.025$, $n_s = 0.9608 \pm 0.0054$ at $k_* = 0.05\,\mathrm{Mpc}^{-1}$.

2.3 Relaxing the power-law

The goal of this paper is to explore the consequences of relaxing the assumption expressed by Eq. (2.5). We wish to let the PPS have a model independent form and reconstruct it from experimental data, and implement a version of the "knot-spline" reconstruction method developed in [36]. This approach is described qualitatively in in Fig. 1. Specifically, we assume that $P(\log k)$ is either a linear or cubic spline with n knots, defined between fixed endpoints. In the absence of knots, this model is functionally equivalent to a power-law, given that we are interpolating in $\log k$. We systematically add knots to allow for more complex features.[1]

The spectrum is generated via the following algorithm:

1. We fix $k_{\min} = 10^{-6}\,\mathrm{Mpc}^{-1}$ and $k_{\max} = 1.0\,\mathrm{Mpc}^{-1}$, allowing the PPS to vary only in amplitude $\Delta_\zeta^2(k)$ at the endpoints.

2. Add n knots, chosen with a uniform prior on $\ln k$, in the range $\ln k_{\min} < \ln k_i < \ln k_{\max}$ and a uniform prior on $-2 < A_i < 4$ for $A_i \equiv \ln(10^{10}\Delta_\zeta^2(k_i))$, with $i = 1, 2, ..., n$. We then order the set of knots so that $k_{i-1} \le k_i$.

3. Interpolate between the endpoints and the n ordered knots by a linear spline (LS_n) or cubic spline (CS_n). Perform the interpolation in logarithmic space for both k and Δ_ζ^2.

The endpoints adopted here define an overall range of k that maps to angular scales from $l = 2$ to $l \approx 14000$, with the rough relationship $l \sim kL_0$, where $L_0 \approx 14.4\,\mathrm{Gpc}$ is the distance to the last scattering surface. The high-end of this range is well past the l-values accessible by *Planck* or any other current or near-future experiment. However, this range coincides with the convention typically used in the Boltzmann-solvers CLASS [55, 56] and CAMB [57], even when C_l's less than $l \sim 2000$ are not evaluated. The logarithmic priors on the knots k_i and their amplitudes A_i indicate that these is no preference for the amplitude and location of any feature(s) and knots with arbitrary locations in k-space can capture both local features and gradual modulations in the PPS.

As Fig. 1 shows, different choices for the interpolation scheme between knots can result in wildly different power spectra. The linear-spline interpolation (LS_n) constructs a continuous but not differentiable PPS by connecting the endpoints and knots by simple line segments. On the other hand, the cubic spline interpolation (CS_n) constructs a PPS with continuous first and second derivatives, connecting knots and endpoints with segments of cubic polynomials. This lets us use the knot-interpolation approach to model and reconstruct many possible features, although a $P(k)$ with many turning points (such as a rapid modulation) would require a prohibitive number of knots. The linear-spline technique is best suited to sharp transitions in the PPS, while the cubic-spline models smooth deviations from the simple

---

[1]The simplest case with only one parameter, *i.e.*, a constant PPS is the familiar Harrison-Zel'dovich spectrum, and is disfavored by *Planck* at $> 5\sigma$ [1].

power law case.[2] Since the positions of the knots are not fixed, they can recover step-like features anywhere in k-space, especially with LS_n, as well as a cut-off PPS (e.g., Ref. [35]).

We have stipulated $k_{i-1} \leq k_i$, so the volume of parameter space corresponding to the knot locations is

$$V_n = \frac{1}{n!} \left(\Delta \ln k \right)^n .$$
(2.6)

If we had not rank-ordered the knot positions we would have $n!$ combinations of k_i and the corresponding parameter region would be an n-cube with volume $n!V_n$. However, in this case the marginalized likelihood integral splits into a sum of $n!$ identical terms after appropriate relabelings of the k_i and the two factors of $n!$ cancel. Consequently, the evidence values do not depend on whether the k_i are assumed to be ordered, but the posterior distributions of the individual k_i will depend strongly on this choice.

With no knots ($n = 0$) the PPS is described by only two parameters: the amplitudes of the PPS at the endpoints $k_{\min}$ and $k_{\max}$. In this limit, both choices LS_0 and CS_0 are equivalent to the power-law PPS of Eq. (2.5). This makes it very easy to compare the Bayesian evidence for the primordial power spectra with features to the standard power law case. Each additional knot adds two extra degrees of freedom to the model, its location k_i and its amplitude A_i, yielding $2n + 2$ total parameters in each PPS parameterization.

3 Data, likelihoods and basic checks

The remaining cosmological parameters $\Omega_b h^2$, $\Omega_c h^2$, h, and τ have uniform priors in the ranges $[0.020, 0.025]$, $[0.1, 0.2]$, $[0.55, 0.85]$, and $[0.02, 0.20]$, respectively. Here, Ω_b denotes the dimensionless baryon density, Ω_c is the dimensionless cold dark matter density, h defines the Hubble parameter H via $H = 100 \, \mathrm{km/s/Mpc}$, and τ is the reionization optical depth.

We use the publicly available code Cosmo++ [41] for our analysis. The CMB power spectra are calculated using the CLASS package [55, 56] and we use the multimodal nested sampling, implemented in the publicly available code MultiNest [44–46], to compute posteriors and Bayesian evidence. The publicly available $Planck$ likelihood code [58] is used for the PPS reconstruction in Section 5. We use the high-l CamSpec likelihood, the low-l Commander likelihood, and the lensing likelihood, thus incorporating all available $Planck$ data, but do not include information from any other sources. The hemispherical analysis in Section 5.2 uses the SMICA map [59] and the Cosmo++ likelihood code, which is also used in Section 4 for analyzing simulated maps. The analysis presented here was carried out using Cosmo++ package [41], and the knot-spline PPS is available in the current version of this package.

We test our tool-chain by estimating parameters for the power-law PPS, Eq. (2.5), using the $Planck$ likelihood code with lensing. We find excellent agreement with the corresponding $Planck$ constraints:[3] $\Omega_b h^2 = 0.02219 \pm 0.00032$, $\Omega_c h^2 = 0.1185 \pm 0.0030$, $h = 0.684 \pm 0.015$, $\tau = 0.089 \pm 0.030$, $n_s = 0.9625 \pm 0.0090$, $A_s = 3.085 \pm 0.054$. With no knots the LS_0 and the CS_0 are both equivalent to the standard case and, as expected, the parameter constraints for these models are almost indistinguishable from the constraints given above.

---

[2]The PPS is numerically represented as a cubic spline of a large number of sample points for either case. This means that the sharp corners of the linear spline interpolation will be smoothed out. This helps avoid numerical problems that could arise because of the discontinuity of the first derivative of the PPS for the linear spline case.

[3]In particular, compare these results to those of Table 9 in Ref. [1].

n	$A'_s = 1 \times 10^{-10}$	$A'_s = 3 \times 10^{-10}$
0	0.0	0.0
1	0.3	51.0
2	0.8	51.7

Table 2. Change in Bayesian evidence $\Delta \ln Z$ with respect to ΛCDM as a function of the number of knots for two simulated maps. The maps have an artificial sinusoidal primordial power spectrum of amplitude A'_s. The errors on $\ln Z$ are smaller than ± 0.3 in all cases.

The *Planck* likelihood code includes 14 nuisance parameters, which mostly relate to unmodelled foregrounds [58]. For the LS_n scenarios we have done the analysis with the foreground parameters both varying and fixed to the central values with a power law PPS. We have verified that the shapes of the posteriors do not change significantly when the foreground parameters are fixed. The computational cost of our procedure is substantially higher for the cubic spline than the linear spline and for convenience the cubic spline analysis was performed with the nuisance parameters fixed.

4   Should we trust our reconstruction?

Before we apply our approach to *Planck* data, we investigate its ability to recover a non-power-law PPS from a temperature map with an injected non–power-law PPS signal in the form of a sinusoidal modulation to the standard power-law PPS

$$\Delta^2_{\zeta,\mathrm{sim}}(k) = A_s \left(\frac{k}{k_*} \right)^{n_s-1} + A'_s \cos \left(\frac{2\pi \ln k}{\ln \lambda'} \right) . \tag{4.1}$$

For simplicity we use the standard ΛCDM parameters at *Planck* best-fit values: $\Omega_b h^2 = 0.022032$, $\Omega_c h^2 = 0.12038$, $h = 0.6704$, $\tau = 0.0925$, $A_s = 2.2154 \times 10^{-9}$, $n_s = 0.9619$, $k_* = 0.05\,\mathrm{Mpc}^{-1}$. For the modification, we use the wavelength $\lambda' = 0.01\,\mathrm{Mpc}$ and two different values for the amplitude: $A'_s = 1 \times 10^{-10}$ and $A'_s = 3 \times 10^{-10}$.

Our simulated maps have similar characteristics to the SMICA map from Planck [59]. We set the HEALPIX parameter $N_{\mathrm{side}} = 2048$ and add white noise to the simulated maps which is similar to the instrumental noise in SMICA. We combine the low-l pixel space likelihood code in the COSMO++ package [41] with the high-l likelihood code in C_l space. The pixel space code is used for $l = 2 \rightarrow 30$ and the C_l-space code is used for $l = 31 \rightarrow 1750$.[4] For the pixel space analysis we reduce the resolution of the simulated maps to HEALPIX $N_{\mathrm{side}} = 16$ by first transforming the maps to harmonic space with the `anafast` routine from HEALPIX, then smooth with a Gaussian kernel with a FWHM of $10°$ using the `alteralm` routine, and finally transform back to pixel space with a lower resolution using `synfast`. For the high-l likelihood calculation we first derive the C_l values from the data using the implementation of the MASTER algorithm [60] in COSMO++ [41]. We combine the SMICA, NILC, and SEVEM masks by Planck to get an approximation of the U73 mask, which is not publicly available. For the pixel space analysis we reduce the resolution of the mask by first

---

[4]The simulated maps become noise dominated after $l \approx 1750$, like the SMICA map.
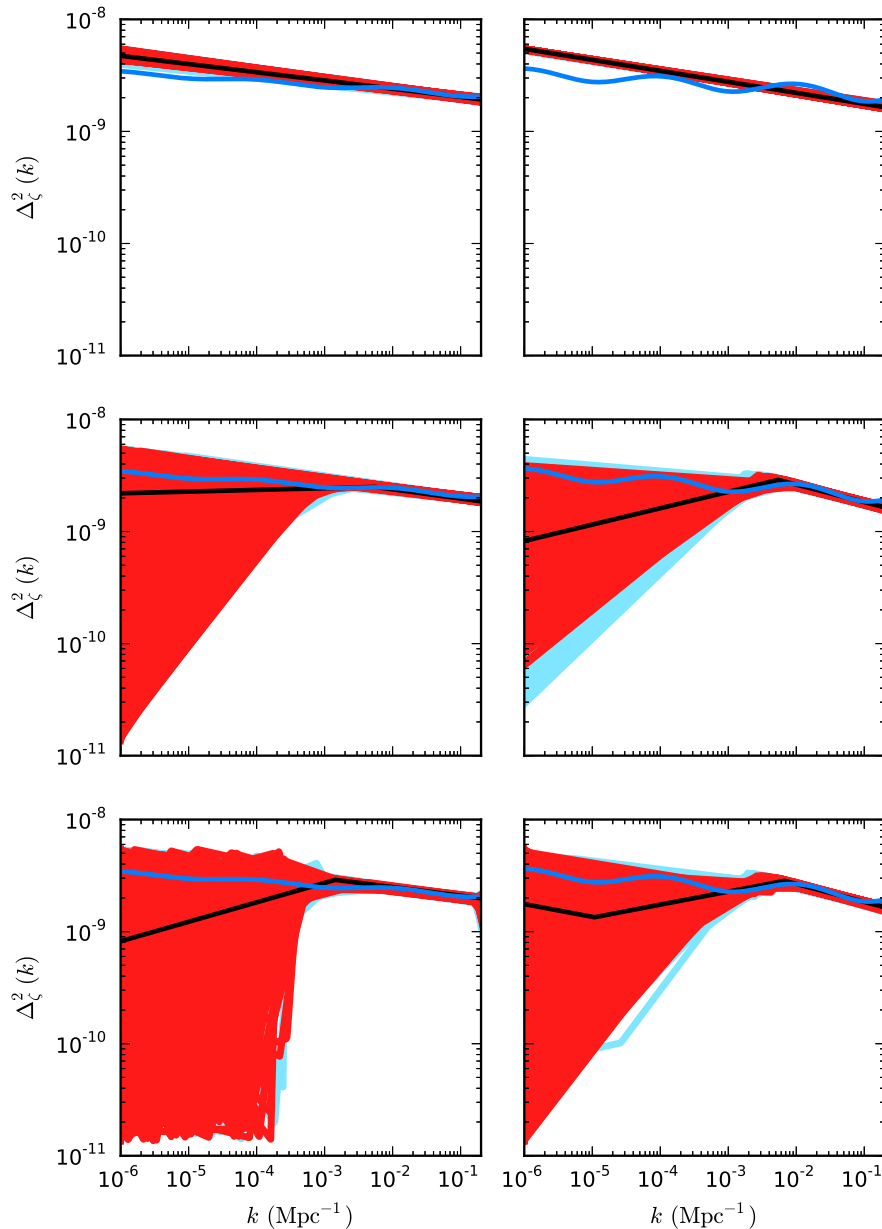
Figure 2. The reconstructed primordial power spectrum (PPS) with simulated data. The columns correspond to (*left*) $A'_s = 1 \times 10^{-10}$ and (*right*) $A'_s = 3 \times 10^{-10}$. The number of knots increases from 0 to 2 from top to bottom. The black solid lines show the best-fit PPS, the red lines are the PPS in the 68% CI, and the light blue lines are the PPS in the 95% CI. The blue solid line shows the modified PPS used for the simulation. Linear spline interpolation is used for all of the plots.

smoothing it with a Gaussian kernel with a FWHM of 10°, then downgrading it with the `ud_grade` routine from HEALPIX, then masking out all the pixels with a final value lower than 0.8.

We apodize the combined mask with a 30′ cosine function for the high-l C_l calculation. We use this mask for the simulated maps to make the analysis similar to the real SMICA

map analysis. The priors on all of the parameters are the same as in Section 3, except for the optical depth of reionization τ. Since no polarization or lensing data is included in this analysis, we use an informative prior of $\tau = 0.0851 \pm 0.014$. This follows the analysis of foreground cleaned temperature maps for *Planck* [59].

The posterior probability distributions for the knot positions, amplitudes, and the best-fit PPS, are shown in Fig. 2. With no knots (equivalent to the power law PPS) the best fit spectrum is a good match to the simulated spectrum at small scales, *i.e.*, $k \gtrsim 10^{-3} \, \mathrm{Mpc}^{-1}$. With $A'_s = 1 \times 10^{-10}$ we recover a power-law with $n_s = 0.924$, while for $A'_s = 3 \times 10^{-10}$ we get $n_s = 0.900$. The added modulation increases the tilt on small scales, making both spectra redder than *Planck*'s. With a single knot we detect the local maximum in Eq. (4.1) at $k \approx 10^{-2} \, \mathrm{Mpc}^{-1}$ or $l \sim 150$. Adding more knots makes less qualitative difference to the posterior distributions as there is only one local maximum on small scales. However, models with knots give posterior probabilities that indicate a suppression of the PPS for $k \lesssim 10^{-3} \, \mathrm{Mpc}^{-1}$ for both simulations, although the assumed PPS in Eq. (4.1) has no such feature. The model reconstructs the slight depression at $10^{-3} \, \mathrm{Mpc}^{-1} \lesssim k \lesssim 10^{-2} \, \mathrm{Mpc}^{-1}$, but is not able to capture the small-k behavior, due to cosmic variance. Consequently, a preference for PPS suppression at $k \lesssim 10^{-3} \, \mathrm{Mpc}^{-1}$ (or $l \lesssim 15$) can be degenerate with a strong, local feature at scales that are better constrained by the likelihood.

Table 2 shows Bayesian evidence ratios for the simulated maps. For a very mild modulation ($A'_s = 1 \times 10^{-10}$) Bayesian evidence is not able to distinguish between models with 0, 1, or 2 knots. However, for a stronger modulation ($A'_s = 3 \times 10^{-10}$) adding a single knot yields a large increase in evidence, $\Delta \ln Z = 51$, conclusively favoring the one knot reconstruction. However models with two or more knots do not give significant extra improvement, due to cosmic variance on large scales.

These simulated maps show that our method can detect features in the power spectrum, even if the features are relatively small. Bayesian evidence clearly distinguishes between models with and without features for the case of the large modulation. However, we have deliberately chosen an uninformative prior, which allows a weakly constrained search but dilutes the ability of the evidence calculation to confirm the presence of smaller modulations.[5] Finally, we must be careful when interpreting posterior probabilities for the PPS that show a decrease in power for $k < 10^{-3} \, \mathrm{Mpc}^{-1}$ as this suppression can be mimicked by a local feature.

## 5 Reconstructing with *Planck*

5.1 Primordial power spectrum

Figure 3 gives the posterior probability distributions for the "knot-spline" primordial power spectrum (PPS) reconstruction. We use the linear-spline (LS_n) model with up to $n = 5$ knots and the cubic-spline (CS_n) model with up to $n = 4$ knots[6] and show the best-fit PPS in black. The 68% and 95% CI are shown in red and light blue, respectively. The constraints on the positions of the knots are given in Fig. 4 and the numerical values are given explicitly in Table 3 of Appendix A.

None of the posteriors show any features beyond $k \gtrsim 2 \times 10^{-3}$ (roughly $l \gtrsim 30$), which is the region best-constrained by the *Planck* likelihoods. The position of maximum likelihood

---

[5]See Ref. [3] for a discussion of this problem in the context of inflationary model selection.

[6]The cubic-spline reconstruction is computationally more expensive. For this reason we restrict our analysis with this model up to 4 knots.
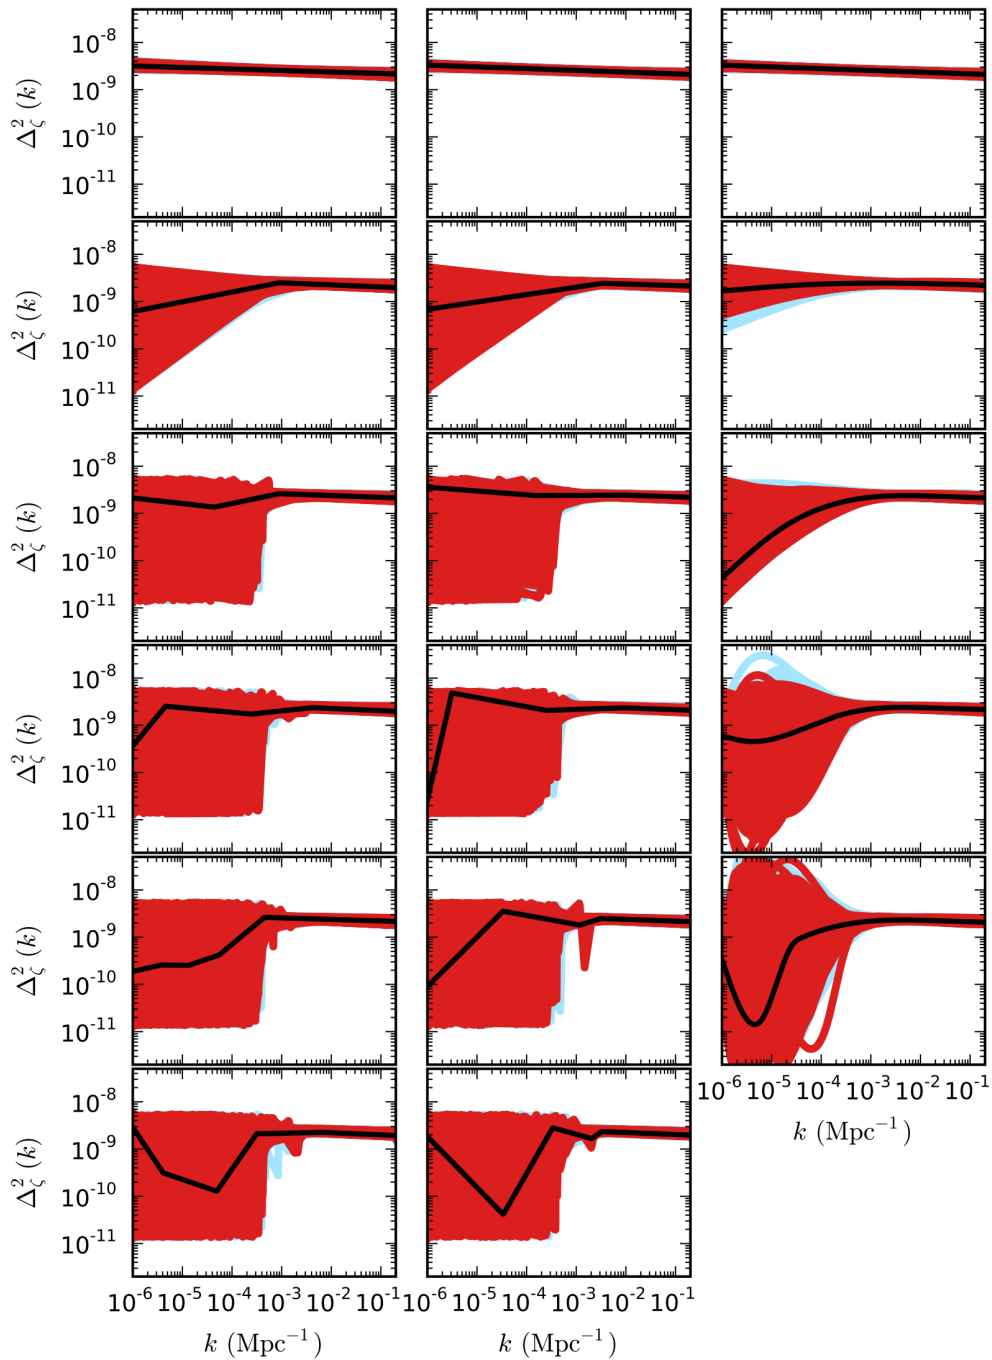
Figure 3. The reconstructed primordial power spectrum (PPS) from the *Planck* data. The columns correspond to (*left*) linear-spline interpolation with n knots (LS_n), while varying the foreground cosmology parameters; (*middle*) LS_n with fixed foreground parameters; and (*right*) cubic-spline interpolation with fixed foreground parameters. The number of knots increases from 0 to 5 from top to bottom. The black solid lines show the best-fit PPS, the red lines are the PPS in the 68% CI, and the light blue lines are the PPS in the 95% CI.
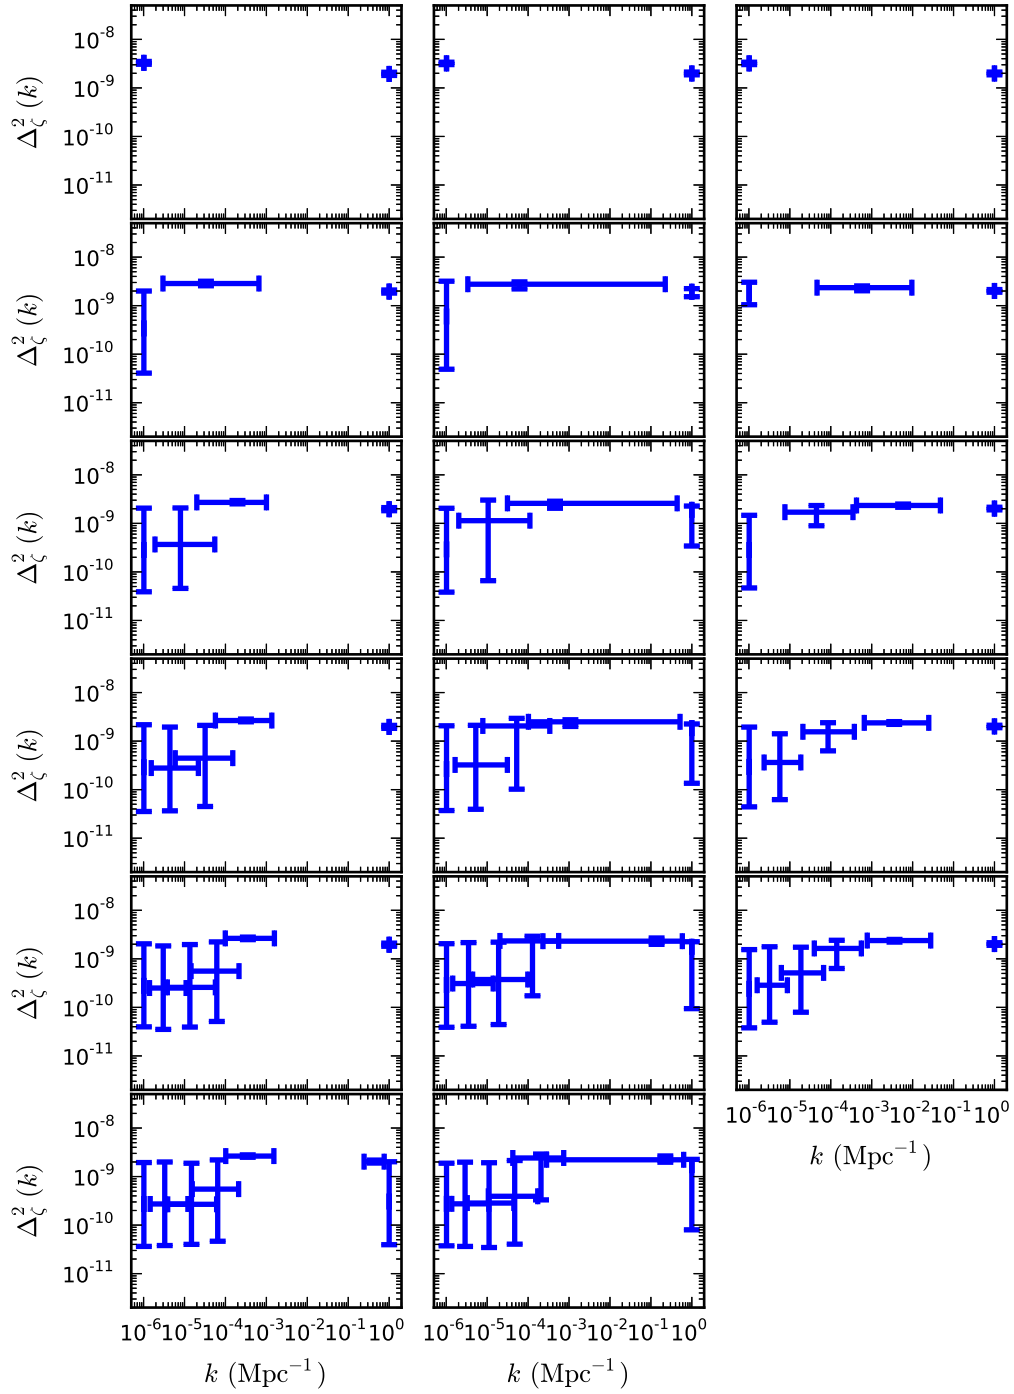
Figure 4. Constraints on the location of knots for the non–power-law primordial power spectrum at the 68% CI. There are n knots, whose location varies in k_i and Δ_ζ^2, and two endpoints that vary only in Δ_ζ^2. The columns are: (*left*) a linear-spline (LS_n) and varying foreground parameters; (*middle*) LS_n with fixed foreground parameters; and (*right*) cubic-spline (CS_n) with fixed foreground parameters. The number of knots increases from 0 to 5 from top to bottom.
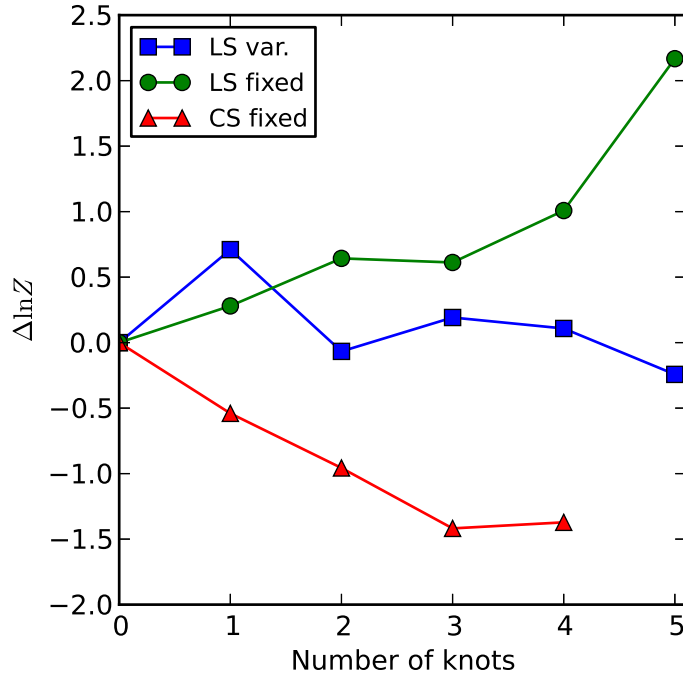
Figure 5. Change in Bayesian evidence $\Delta \ln Z$ with respect to ΛCDM as a function of the number of knots in the primordial power spectrum model. The blue squares correspond to a linear-spline (LS_n) without fixing the foreground parameters; the green circles correspond to LS_n with fixed foreground parameters; and the red triangles correspond to a cubic-spline (CS_n) with fixed foreground parameters.

for the last knot is relatively stable at $k \sim 5 \times 10^{-3}\,\mathrm{Mpc}^{-1}$, implying there is little evidence for global features. Furthermore, we do not see any pair of knots (or knot/endpoint) that consistently retains its position as more knots are added, indicating that there are no localized features in the PPS for scales that are well-constrained by data.

However, most of the reconstructed power spectra give preference to suppressed power below $k \lesssim 10^{-3}\,\mathrm{Mpc}^{-1}$, in agreement with previous results [12, 36]. The cosmic variance, however, is largest on these large scales. Furthermore, as described in Section 4, even if there were significant evidence in favor of a non–power-law model, this low-k suppression should be treated with care, since it could also indicate a local feature near $k \sim 10^{-3}\,\mathrm{Mpc}^{-1}$.

To evaluate the integrated likelihood for the non–power-law models, we report the Bayesian evidence in Fig. 5. We use the Bayesian evidence for the PPS models with no knots as a reference, since this case corresponds to the standard power-law PPS of Eq. (2.5) and the ΛCDM model. There is a slightly increased evidence for the linear-spline model with one knot (LS_1), both with and without varying the foreground parameters, with a maximum of $\Delta \ln Z = 0.7$. For more knots, the Bayesian evidence is smaller than the power-law PPS for the cubic-spline model and the linear-spline model with varying foreground parameters. The CS_n reconstruction gives smaller Bayes factors compared to the standard case for any n. When a linear-spline interpolation is used with fixed foreground parameters, we see a slightly increased evidence for a higher numbers of knots. However, in this case we are ignoring the impact of marginalising over the foreground parameters, which introduces a nontrivial

uncertainty into Bayesian evidence calculation. When the foreground parameters are varied, this evidence does not increase with n. Therefore, we conclude that we see no significant preference for a higher number of knots.

A frequentist search for broad features in the power spectrum using Planck temperature data was performed in Ref. [14]. The authors used binned power spectra with a fixed value of n_s with up to four bins, and a varying n_s with up to two bins. Their two-bin Model-C reconstruction is equivalent to our linear spline reconstruction with 1 knot. Our best-fit LS_1 power spectrum agrees very well with the best-fit case of [14] and the overall conclusions of [14] are consistent with with our Bayesian results.

Figure 6 presents the posterior probabilities for the standard parameters $\Omega_b h^2$, $\Omega_c h^2$, h, and τ for both the linear-spline (LS_n) and cubic-spline (CS_n) models. The constraints on the parameters do not change significantly when more freedom is given to the PPS. The constraints with fixed foreground parameters are in good agreement with those found when they are free to vary; the error bars shrink slightly with fixed foreground parameters, as expected. We do not see any significant changes in the constraints of the cosmological parameters with up to 5 knots in the PPS, implying that the parameter constraints by Planck [1] are robust and do not depend sensitively on the strong power-law assumption on the form of the PPS.

5.2   Hemispherical power asymmetry

The latest results from Planck [61] confirm the presence [62, 63] of a power asymmetry in different ecliptic hemispheres of the sky. Only about 4 simulations out of 500 have a higher level of power asymmetry than that found in the data [61]. While some authors (see e.g., Ref. [64]) have questioned the significance of the power asymmetry on small scales, it is a persistent anomaly on large scales. There are many suggestions [65–70] as to how such an asymmetry might arise, and we study the possibility of the power asymmetry resulting from having more structure in the primordial power spectrum (PPS) in one hemisphere compared to the other.

Using the SMICA map from Planck [59] we reconstruct the PPS in each hemisphere, including the mask as described in Section 4. The likelihood is computed using the same process employed with the simulated maps in Section 4. The apodized masks for the two hemispheres used in our analysis are shown in Fig. 7 and we test our mask and statistical methods by reproducing the Planck power spectra on different hemispheres of the SMICA map (cf. Fig. 28 in Ref. [61]). We also made a full analysis of the SMICA map with our combined mask and a standard power-law PPS as in Eq. (2.5). The resulting constraints on the cosmological parameters are very similar to those obtained by Planck [59] (see Ref. [41] for a detailed comparison).

We reconstruct the PPS on the full sky, as well as the northern and the southern ecliptic hemispheres, using linear-spline interpolation (LS_n) with $n \leq 5$ knots. The resulting posterior distributions for the PPS are shown in Fig. 8, and the Bayesian evidence ratios as a function of the number of knots are shown in Fig. 9. Although small differences are seen in the posterior distributions of the PPS and the best-fit PPS between the two hemispheres, Bayesian evidence does not give significant preference to extra features in either hemisphere. The evidence ratios for the two hemispheres are very close to each other, and to the full sky case. The largest evidence value seen is $\Delta \log Z = 2.3$ for the northern hemisphere but this was found with the foreground cleaned SMICA map. Given that the fixing of the foreground parameters for Planck led to a spurious increase in evidence on the full sky, we interpret
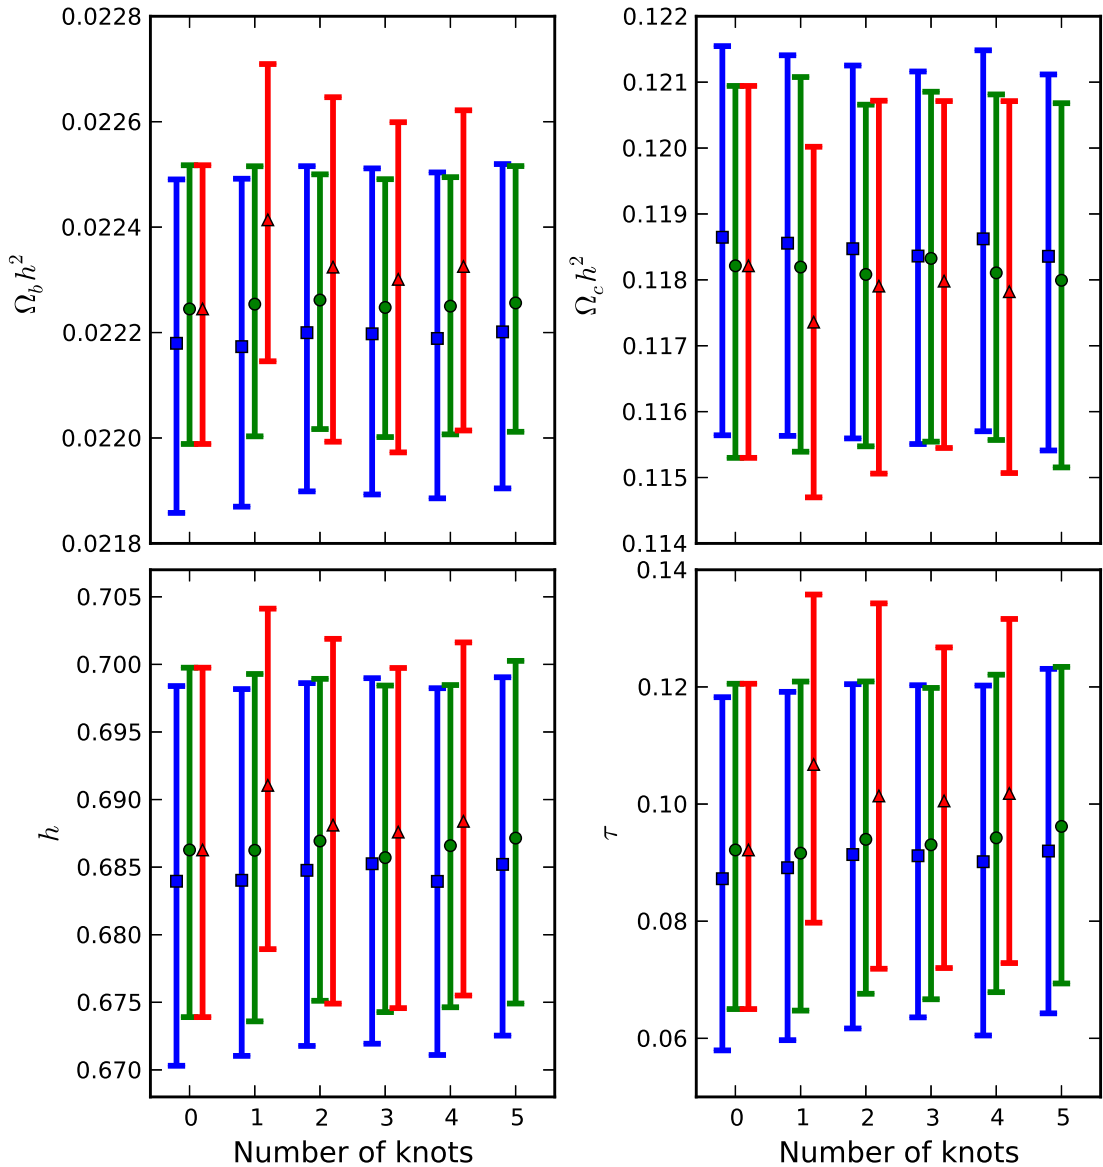
Figure 6. Cosmological parameter constraints as a function of the number of knots n at the 68% CI. The blue squares correspond to a power spectrum linear-spline interpolation and varying foreground parameters, the green circles correspond to a linear-spline interpolation with fixed foreground parameters, and the red triangles represent the results with a cubic-spline interpolation and fixed foreground parameters.

these results as indicating that there is no difference in the structure of the PPS in the two hemispheres.

6 Summary

We have applied the "knot-spline" reconstruction method [36] to the Planck temperature data. This paper breaks new ground by checking the algorithm's ability to recover the PPS
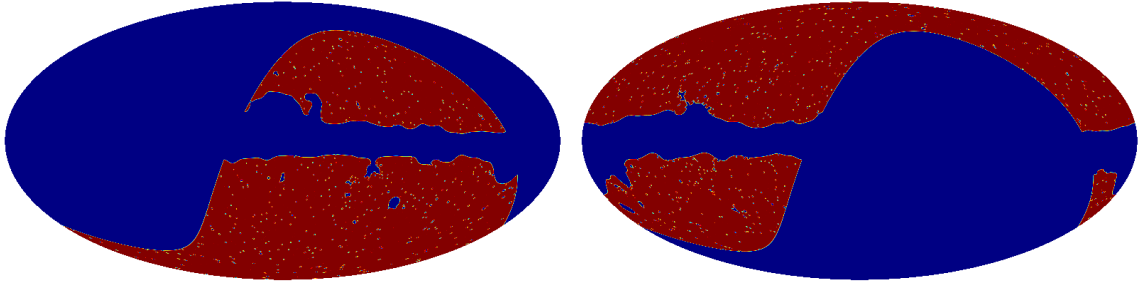
**Figure 7**. Apodized masks for the southern (*left*) and northern (*right*) ecliptic hemispheres.

from simulated maps with artificially introduced features and by confirming that cosmological parameter constraints obtained from the Planck data are not diluted when the usual assumptions about the form of the primordial power spectrum are relaxed. Furthermore, we investigate whether the hemispherical power asymmetry visible in the WMAP and Planck temperature maps is correlated with differences in the primordial power spectrum recovered from each hemisphere, finding that the two power spectra are in good agreement. Finally, the numerical tools needed to reproduce or extend this analysis are now included with the COSMO++ library [41].

The PPS reconstruction method used here allows the location of the knots to vary in both k-space and amplitude, and allows us to capture both gentle variations in the spectrum as well as a broad class of localized features. Increasing the possible complexity of the PPS necessarily improves the fit to the data, and we must guard against the "look-elsewhere" effect or "fitting the noise". Determining the optimal number of knots can be posed as a model selection problem, and we use Bayesian evidence ratios to safeguard against overfitting.

We applied our methods to simulated maps with Planck characteristics to check the reliability of the method, and to estimate the amplitude of possible features in the PPS that the method can detect. We were able to recover modulations which modified an underlying power law spectrum by less than 5%. Typically, specific modulations have well-defined thresholds above which they are very easy to detect; for example, a long wavelength modulation with an amplitude a factor of three beyond the threshold of detectability yields an improvement in evidence of $\Delta \ln Z = 51$. Because the *Planck* likelihoods are more sensitive to features at $k > 10^{-3}\,\mathrm{Mpc}^{-1}$, the posteriors on the knots' positions (Fig. 2) show a decrement of power at $k < 10^{-3}\,\mathrm{Mpc}^{-1}$, although this is not a feature of the simulated data, and cautions us against over-interpreting an apparent decrement in large scale power in the actual sky maps. More generally, the weak evidence computed for the smaller modulation is partly driven by the use of "uninformative" priors for the modulated spectrum [3]. Consequently, Bayesian evidence does not permit a strictly algorithmic solution to cosmological model-selection problems and with maximum entropy priors similar to those used here, and nuanced physical analyses of the improvement in the maximum likelihood along with cross-checks against other datasets will remain important.

Having tested our methods, we reconstructed the PPS from *Planck* CMB temperature and lensing data. We found no evidence for deviations from the standard power law PPS on scales with $k \gtrsim 10^{-3}\,\mathrm{Mpc}^{-1}$. Although on larger scales the data is not able to distinguish between models with or without features due to cosmic variance, the extensions to ΛCDM do not have sufficient Bayesian evidence to favor them over a standard power-law PPS. Furthermore, the posteriors for the "standard" cosmological parameters did not differ substantially
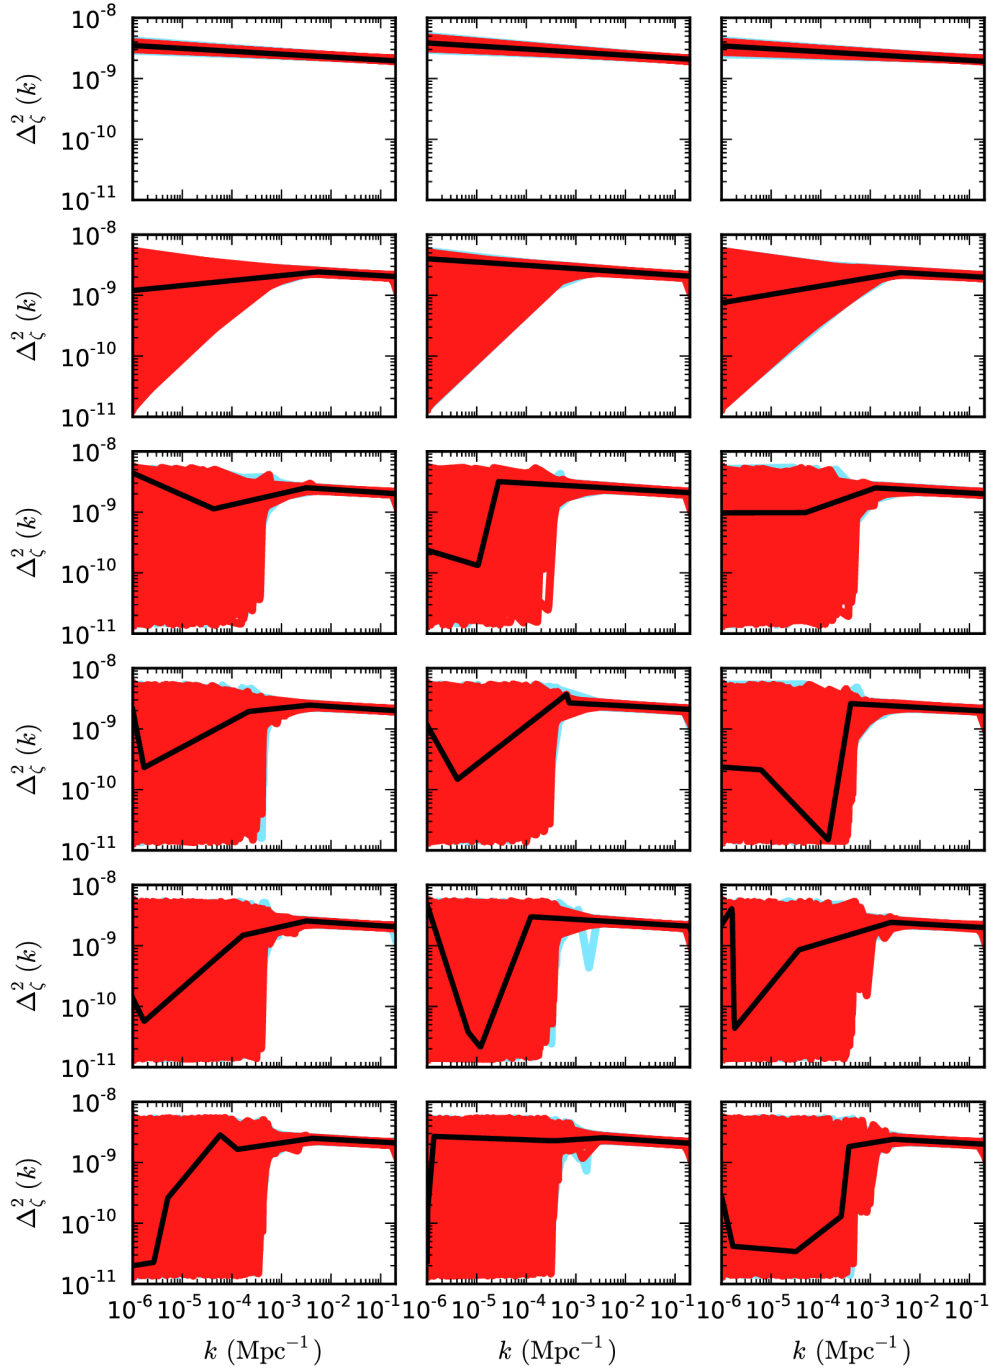
Figure 8. The reconstructed primordial power spectrum (PPS) on the full sky compared to different ecliptic hemispheres. The columns correspond to (*left*) the full sky; (*middle*) the southern hemisphere; and (*right*) the northern hemisphere. The number of knots increases from 0 to 5 from top to bottom. The black solid lines show the best-fit PPS, the red lines are the PPS in the 68% CI, and the light blue lines are the PPS in the 95% CI. All of the plots have been obtained from the SMICA map from *Planck* using the linear-spline interpolation model with n knots.
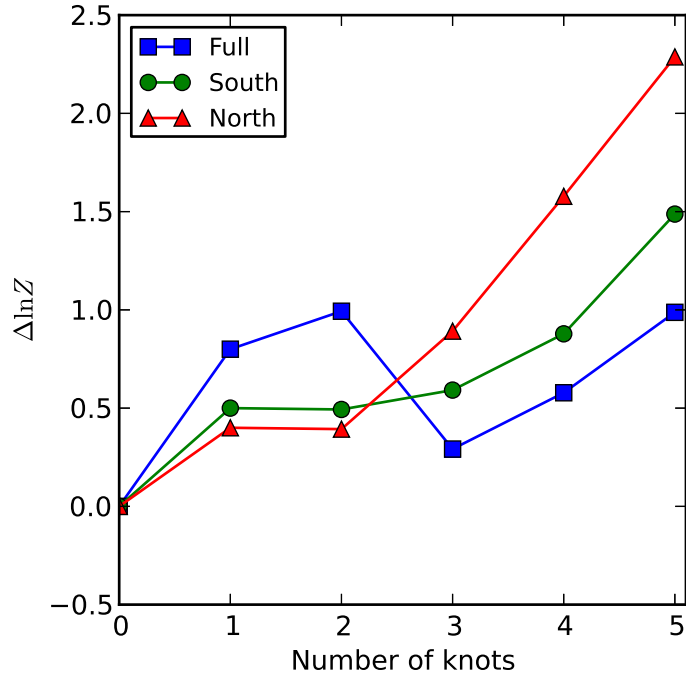
Figure 9. Change in Bayesian evidence $\Delta \ln Z$ with respect to ΛCDM as a function of the number of knots used in the primordial power spectrum reconstruction. The blue squares correspond to the full sky, the green circles correspond to the southern hemisphere (in ecliptic coordinates), and the red triangles correspond to the northern hemisphere. All of the plots have been obtained from the SMICA map from *Planck* using the linear-spline interpolation model with n knots.

from the power-law case did not change significantly when a more general PPS was allowed and we can conclude that the *Planck* constraints on these parameters are robust. Finally, we performed a PPS reconstruction on each individual hemisphere, but found no systematic difference between the results showing that any "hemispherical anomaly" is not associated with differences in the underlying power spectrum.

This paper is the first in a sequence of analyses of non-standard power spectra. In particular we will investigate the implications of the recent detection of B-mode polarization by the BICEP2 telescope [42, 43] for the scalar power spectrum [71], and in a third paper we will study whether permitting a non–power-law PPS changes the estimated values of derived parameters such as σ_8 or modifies estimated constraints on the neutrino sector.

Acknowledgments

A   Numerical data

In this appendix we present tables of numerical values to further quantify our analysis. Table 3 reports the numerical values for the posterior probabilities on the location of knots in the reconstructed primordial power spectrum in Figs 3 and 4. Table 4 shows the constraints on the cosmological parameters in Fig. 6.

References

[1] **Planck** Collaboration, P. Ade *et. al.*, *Planck 2013 results. I. Overview of products and scientific results*, arXiv:1303.5062.

[2] J. Martin, C. Ringeval, and R. Trotta, *Hunting Down the Best Model of Inflation with Bayesian Evidence*, Phys.Rev. **D83** (2011) 063524, [arXiv:1009.4157].

[3] R. Easther and H. V. Peiris, *Bayesian Analysis of Inflation II: Model Selection and Constraints on Reheating*, Phys.Rev. **D85** (2012) 103533, [arXiv:1112.0326].

[4] **Planck** Collaboration, P. Ade *et. al.*, *Planck 2013 results. XXII. Constraints on inflation*, arXiv:1303.5082.

[5] J. Martin, C. Ringeval, R. Trotta, and V. Vennin, *The Best Inflationary Models After Planck*, arXiv:1312.3529.

[6] A. Ashoorioon and A. Krause, *Power Spectrum and Signatures for Cascade Inflation*, hep-th/0607001.

[7] A. Ashoorioon, A. Krause, and K. Turzynski, *Energy Transfer in Multi Field Inflation and Cosmological Perturbations*, JCAP **0902** (2009) 014, [arXiv:0810.4660].

[8] R. Flauger, L. McAllister, E. Pajer, A. Westphal, and G. Xu, *Oscillations in the CMB from Axion Monodromy Inflation*, JCAP **1006** (2010) 009, [arXiv:0907.2916].

[9] A. Achucarro, J.-O. Gong, S. Hardeman, G. A. Palma, and S. P. Patil, *Features of heavy physics in the CMB power spectrum*, JCAP **1101** (2011) 030, [arXiv:1010.3693].

[10] H. Peiris, R. Easther, and R. Flauger, *Constraining Monodromy Inflation*, JCAP **1309** (2013) 018, [arXiv:1303.2616].

[11] R. Easther and R. Flauger, *Planck Constraints on Monodromy Inflation*, JCAP **1402** (2014) 037, [arXiv:1308.3736].

[12] X. Wang and G.-B. Zhao, *Reconstructing primordial power spectrum using Planck and SDSS-III measurements*, arXiv:1309.6624.

[13] D. K. Hazra, A. Shafieloo, and T. Souradeep, *Primordial power spectrum: a complete analysis with the WMAP nine-year data*, JCAP **1307** (2013) 031, [arXiv:1303.4143].

[14] D. K. Hazra, A. Shafieloo, and G. F. Smoot, *Reconstruction of broad features in the primordial spectrum and inflaton potential from Planck*, JCAP **1312** (2013) 035, [arXiv:1310.3038].

[15] C. Gauthier and M. Bucher, *Reconstructing the primordial power spectrum from the CMB*, JCAP **1210** (2012) 050, [arXiv:1209.2147].

[16] P. Paykari and A. H. Jaffe, *Optimal Binning of the Primordial Power Spectrum*, Astrophys.J. **711** (2010) 1–12, [arXiv:0902.4399].

[17] G. Nicholson, C. R. Contaldi, and P. Paykari, *Reconstruction of the Primordial Power Spectrum by Direct Inversion*, JCAP **1001** (2010) 016, [arXiv:0909.5092].

[18] G. Nicholson and C. R. Contaldi, *Reconstruction of the Primordial Power Spectrum using Temperature and Polarisation Data from Multiple Experiments*, JCAP **0907** (2009) 011, [arXiv:0903.1106].

[19] R. Nagata and J. Yokoyama, *Band-power reconstruction of the primordial fluctuation spectrum by the maximum likelihood reconstruction method*, Phys.Rev. **D79** (2009) 043010, [arXiv:0812.4585].

[20] R. Nagata and J. Yokoyama, *Reconstruction of the primordial fluctuation spectrum from the five-year WMAP data by the cosmic inversion method with band-power decorrelation analysis*, Phys.Rev. **D78** (2008) 123002, [arXiv:0809.4537].

[21] A. Shafieloo and T. Souradeep, *Estimation of Primordial Spectrum with post-WMAP 3 year data*, Phys.Rev. **D78** (2008) 023511, [arXiv:0709.1944].

[22] A. Shafieloo, T. Souradeep, P. Manimaran, P. K. Panigrahi, and R. Rangarajan, *Features in the Primordial Spectrum from WMAP: A Wavelet Analysis*, Phys.Rev. **D75** (2007) 123502, [astro-ph/0611352].

[23] D. Tocchini-Valentini, Y. Hoffman, and J. Silk, *Non-parametric reconstruction of the primordial power spectrum at horizon scales from wmap data*, Mon.Not.Roy.Astron.Soc. **367** (2006) 1095–1102, [astro-ph/0509478].

[24] S. M. Leach, *Measuring the primordial power spectrum: Principal component analysis of the cosmic microwave background*, Mon.Not.Roy.Astron.Soc. **372** (2006) 646–654, [astro-ph/0506390].

[25] A. Shafieloo and T. Souradeep, *Primordial power spectrum from WMAP*, Phys.Rev. **D70** (2004) 043523, [astro-ph/0312174].

[26] S. Hannestad, *Reconstructing the primordial power spectrum - A New algorithm*, JCAP **0404** (2004) 002, [astro-ph/0311491].

[27] G. Efstathiou, *Myths and truths concerning estimation of power spectra*, Mon.Not.Roy.Astron.Soc. **349** (2004) 603, [astro-ph/0307515].

[28] P. Mukherjee and Y. Wang, *Model-independent reconstruction of the primordial power spectrum from WMAP data*, Astrophys.J. **599** (2003) 1–6, [astro-ph/0303211].

[29] S. Hannestad, *Reconstructing the inflationary power spectrum from CMBR data*, Phys.Rev. **D63** (2001) 043009, [astro-ph/0009296].

[30] L. Verde and H. V. Peiris, *On Minimally-Parametric Primordial Power Spectrum Reconstruction and the Evidence for a Red Tilt*, JCAP **0807** (2008) 009, [arXiv:0802.1219].

[31] H. V. Peiris and L. Verde, *The Shape of the Primordial Power Spectrum: A Last Stand Before Planck*, Phys.Rev. **D81** (2010) 021302, [arXiv:0912.0268].

[32] S. Bird, H. V. Peiris, M. Viel, and L. Verde, *Minimally Parametric Power Spectrum Reconstruction from the Lyman-alpha Forest*, Mon.Not.Roy.Astron.Soc. **413** (2011) 1717–1728, [arXiv:1010.1519].

[33] M. Bridges, A. Lasenby, and M. Hobson, *A bayesian analysis of the primordial power spectrum*, Mon.Not.Roy.Astron.Soc. **369** (2006) 1123–1130, [astro-ph/0511573].

[34] M. Bridges, A. Lasenby, and M. Hobson, *WMAP 3-year primordial power spectrum*, Mon.Not.Roy.Astron.Soc. (2006) [astro-ph/0607404].

[35] M. Bridges, F. Feroz, M. Hobson, and A. Lasenby, *Bayesian optimal reconstruction of the primordial power spectrum*, arXiv:0812.3541.

[36] J. A. Vazquez, M. Bridges, M. Hobson, and A. Lasenby, *Model selection applied to reconstruction of the Primordial Power Spectrum*, JCAP **1206** (2012) 006, [arXiv:1203.1252].

[37] J. A. Vazquez, A. Lasenby, M. Bridges, and M. Hobson, *A Bayesian study of the primordial power spectrum from a novel closed universe model*, *Mon.Not.Roy.Astron.Soc.* **422** (2012) 1948–1956, [`arXiv:1103.4619`].

[38] J. A. Vazquez, M. Bridges, Y.-Z. Ma, and M. Hobson, *Constraints on the tensor-to-scalar ratio for non-power-law models*, *JCAP* **1308** (2013) 001, [`arXiv:1303.4014`].

[39] R. de Putter, E. V. Linder, and A. Mishra, *Inflationary Freedom and Cosmological Neutrino Constraints*, *arXiv:1401.7022* **astro-ph.CO** (2014).

[40] W. H. Kinney, *How to fool cosmic microwave background parameter estimation*, *Phys.Rev.* **D63** (2001) 043001, [`astro-ph/0005410`].

[41] G. Aslanyan, *Cosmo++: An Object-Oriented C++ Library for Cosmology*, `arXiv:1312.4961`.

[42] **BICEP2** Collaboration, P. Ade *et. al.*, *BICEP2 I: Detection Of B-mode Polarization at Degree Angular Scales*, `arXiv:1403.3985`.

[43] **BICEP2** Collaboration, P. A. R. Ade *et. al.*, *BICEP2 II: Experiment and Three-Year Data Set*, `arXiv:1403.4302`.

[44] F. Feroz and M. Hobson, *Multimodal nested sampling: an efficient and robust alternative to MCMC methods for astronomical data analysis*, *Mon.Not.Roy.Astron.Soc.* **384** (2008) 449, [`arXiv:0704.3704`].

[45] F. Feroz, M. Hobson, and M. Bridges, *MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics*, *Mon.Not.Roy.Astron.Soc.* **398** (2009) 1601–1614, [`arXiv:0809.3437`].

[46] F. Feroz, M. Hobson, E. Cameron, and A. Pettitt, *Importance Nested Sampling and the MultiNest Algorithm*, `arXiv:1306.2144`.

[47] H. Jeffreys, *Theory of Probability.* Oxford University Press, 1961.

[48] M. P. Hobson, A. H. Jaffe, A. R. Liddle, P. Mukherjee, and D. Parkinson, *Bayesian methods in cosmology.* Cambridge University Press, 2010.

[49] J. O. Berger and J. Mortera, *Interpreting the stars in precise hypothesis testing*, *International Statistical Review* **59** (1991), no. 3 337–353.

[50] D. Baumann, *TASI Lectures on Inflation*, `arXiv:0907.5424`.

[51] **Planck** Collaboration, P. Ade *et. al.*, *Planck 2013 Results. XXIV. Constraints on primordial non-Gaussianity*, `arXiv:1303.5084`.

[52] C. Bennett, D. Larson, J. Weiland, N. Jarosik, G. Hinshaw, *et. al.*, *Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results*, `arXiv:1212.5225`.

[53] J. L. Sievers, R. A. Hlozek, M. R. Nolta, V. Acquaviva, G. E. Addison, *et. al.*, *The Atacama Cosmology Telescope: Cosmological parameters from three seasons of data*, `arXiv:1301.0824`.

[54] K. Story, C. Reichardt, Z. Hou, R. Keisler, K. Aird, *et. al.*, *A Measurement of the Cosmic Microwave Background Damping Tail from the 2500-square-degree SPT-SZ survey*, *Astrophys.J.* **779** (2013) 86, [`arXiv:1210.7231`].

[55] J. Lesgourgues, *The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview*, `arXiv:1104.2932`.

[56] D. Blas, J. Lesgourgues, and T. Tram, *The Cosmic Linear Anisotropy Solving System (CLASS) II: Approximation schemes*, *JCAP* **1107** (2011) 034, [`arXiv:1104.2933`].

[57] A. Lewis, A. Challinor, and A. Lasenby, *Efficient computation of CMB anisotropies in closed FRW models*, *Astrophys.J.* **538** (2000) 473–476, [`astro-ph/9911177`].

[58] **Planck** Collaboration, P. Ade *et. al.*, *Planck 2013 results. XV. CMB power spectra and likelihood*, `arXiv:1303.5075`.

[59] **Planck** Collaboration, P. Ade *et. al.*, *Planck 2013 results. XII. Component separation*, `arXiv:1303.5072`.

[60] E. Hivon, K. Gorski, C. Netterfield, B. Crill, S. Prunet, *et. al.*, *Master of the cosmic microwave background anisotropy power spectrum: a fast method for statistical analysis of large and complex cosmic microwave background data sets*, *Astrophys.J.* **567** (2002) 2, [`astro-ph/0105302`].

[61] **Planck** Collaboration, P. Ade *et. al.*, *Planck 2013 results. XXIII. Isotropy and statistics of the CMB*, `arXiv:1303.5083`.

[62] H. Eriksen, F. Hansen, A. Banday, K. Gorski, and P. Lilje, *Asymmetries in the Cosmic Microwave Background anisotropy field*, *Astrophys.J.* **605** (2004) 14–20, [`astro-ph/0307507`].

[63] F. Hansen, A. Banday, K. Gorski, H. Eriksen, and P. Lilje, *Power Asymmetry in Cosmic Microwave Background Fluctuations from Full Sky to Sub-degree Scales: Is the Universe Isotropic?*, *Astrophys.J.* **704** (2009) 1448–1458, [`arXiv:0812.3795`].

[64] S. Flender and S. Hotchkiss, *The small scale power asymmetry in the cosmic microwave background*, *JCAP* **1309** (2013) 033, [`arXiv:1307.6069`].

[65] A. L. Erickcek, M. Kamionkowski, and S. M. Carroll, *A Hemispherical Power Asymmetry from Inflation*, *Phys.Rev.* **D78** (2008) 123520, [`arXiv:0806.0377`].

[66] L. Dai, D. Jeong, M. Kamionkowski, and J. Chluba, *The Pesky Power Asymmetry*, *Phys.Rev.* **D87** (2013), no. 12 123005, [`arXiv:1303.6949`].

[67] A. R. Liddle and M. Cortês, *Cosmic microwave background anomalies in an open universe*, *Phys.Rev.Lett.* **111** (2013) 111302, [`arXiv:1306.5698`].

[68] J. McDonald, *Hemispherical Power Asymmetry from a Space-Dependent Component of the Adiabatic Power Spectrum*, `arXiv:1403.2076`.

[69] P. K. Rath and P. Jain, *Relating the anisotropic power spectrum to the CMB hemispherical anisotropy*, `arXiv:1403.2567`.

[70] G. Aslanyan and R. Easther, *Large Scale Anomalies in the Microwave Background: Causation and Correlation*, *Phys.Rev.Lett.* **111** (2013) 261301, [`arXiv:1308.6593`].

[71] K. N. Abazajian, G. Aslanyan, R. Easther, and L. C. Price, *The Knotted Sky II: Does BICEP2 require a nontrivial primordial power spectrum?*, `arXiv:1403.5922`.

n	LS$_n$ var.		LS$_n$ fixed		CS$_n$ fixed	
	$\ln k_i$	A_i	$\ln k_i$	A_i	$\ln k_i$	A_i
0	-13.8	$3.493^{+0.064}_{-0.067}$	-13.8	$3.466^{+0.054}_{-0.056}$	-13.8	$3.466^{+0.054}_{-0.056}$
	0.0	$2.967^{+0.078}_{-0.073}$	0.0	$2.987^{+0.067}_{-0.064}$	0.0	$2.987^{+0.067}_{-0.064}$
1	-13.8	$1.2^{+1.8}_{-2.1}$	-13.8	$1.8^{+1.7}_{-2.5}$	-13.8	$2.91^{+0.49}_{-0.56}$
	$-10.2^{+2.9}_{-2.4}$	$3.35^{+0.11}_{-0.11}$	$-9.7^{+8.1}_{-2.8}$	$3.32^{+0.11}_{-0.23}$	$-7.5^{+2.8}_{-2.5}$	$3.157^{+0.077}_{-0.146}$
	0.0	$2.970^{+0.074}_{-0.071}$	0.0	$2.981^{+0.097}_{-0.166}$	0.0	$3.000^{+0.065}_{-0.059}$
2	-13.8	$1.0^{+2.0}_{-1.9}$	-13.8	$1.1^{+1.9}_{-2.0}$	-13.8	$1.1^{+1.6}_{-1.8}$
	$-11.7^{+1.9}_{-1.4}$	$1.3^{+1.7}_{-2.0}$	$-11.4^{+2.3}_{-1.6}$	$2.42^{+0.97}_{-2.81}$	$-10.0^{+2.0}_{-1.8}$	$2.84^{+0.31}_{-0.65}$
	$-8.5^{+1.6}_{-2.3}$	$3.298^{+0.089}_{-0.069}$	$-7.8^{+6.8}_{-2.6}$	$3.253^{+0.099}_{-0.233}$	$-5.1^{+2.1}_{-2.6}$	$3.153^{+0.064}_{-0.085}$
	0.0	$2.974^{+0.076}_{-0.071}$	0.0	$2.97^{+0.11}_{-1.62}$	0.0	$3.004^{+0.063}_{-0.061}$
3	-13.8	$1.1^{+2.0}_{-2.1}$	-13.8	$1.0^{+2.0}_{-1.9}$	-13.8	$1.1^{+1.8}_{-1.9}$
	$-12.2^{+1.6}_{-1.0}$	$1.0^{+1.9}_{-2.0}$	$-12.1^{+1.8}_{-1.1}$	$1.2^{+1.8}_{-2.1}$	$-12.07^{+1.16}_{-0.84}$	$1.3^{+1.4}_{-1.8}$
	$-10.3^{+1.5}_{-1.7}$	$1.5^{+1.5}_{-2.3}$	$-9.8^{+1.8}_{-1.9}$	$3.07^{+0.29}_{-3.00}$	$-9.4^{+1.5}_{-1.4}$	$2.75^{+0.41}_{-0.90}$
	$-8.1^{+1.4}_{-1.7}$	$3.278^{+0.069}_{-0.060}$	$-7.0^{+6.3}_{-2.2}$	$3.219^{+0.094}_{-0.219}$	$-5.6^{+1.9}_{-1.7}$	$3.168^{+0.056}_{-0.068}$
	0.0	$2.976^{+0.073}_{-0.068}$	0.0	$2.95^{+0.12}_{-2.60}$	0.0	$2.999^{+0.055}_{-0.061}$
4	-13.8	$1.0^{+1.9}_{-1.9}$	-13.8	$1.1^{+1.9}_{-2.0}$	-13.8	$0.9^{+1.8}_{-1.8}$
	$-12.62^{+1.26}_{-0.73}$	$0.9^{+2.0}_{-1.9}$	$-12.47^{+1.34}_{-0.87}$	$1.1^{+1.9}_{-2.0}$	$-12.62^{+0.96}_{-0.64}$	$1.0^{+1.8}_{-1.7}$
	$-11.2^{+1.4}_{-1.3}$	$1.0^{+2.0}_{-1.9}$	$-10.9^{+1.6}_{-1.4}$	$1.3^{+1.7}_{-2.1}$	$-10.9^{+1.2}_{-1.1}$	$1.6^{+1.2}_{-1.9}$
	$-9.7^{+1.2}_{-1.4}$	$1.7^{+1.4}_{-2.3}$	$-9.0^{+1.4}_{-1.8}$	$3.19^{+0.16}_{-2.56}$	$-8.9^{+1.4}_{-1.2}$	$2.80^{+0.37}_{-0.95}$
	$-8.0^{+1.5}_{-1.3}$	$3.274^{+0.059}_{-0.060}$	$-5.0^{+4.4}_{-3.3}$	$3.14^{+0.15}_{-0.15}$	$-5.6^{+2.0}_{-1.5}$	$3.170^{+0.052}_{-0.073}$
	0.0	$2.972^{+0.075}_{-0.073}$	0.0	$2.91^{+0.16}_{-2.94}$	0.0	$3.003^{+0.063}_{-0.058}$
5	-13.8	$1.0^{+1.9}_{-1.9}$	-13.8	$1.0^{+1.9}_{-2.0}$		
	$-12.53^{+1.27}_{-0.78}$	$1.0^{+2.0}_{-1.9}$	$-12.65^{+1.29}_{-0.74}$	$1.0^{+1.9}_{-2.0}$		
	$-11.1^{+1.4}_{-1.3}$	$1.0^{+1.9}_{-1.9}$	$-11.4^{+1.4}_{-1.2}$	$1.0^{+1.9}_{-2.1}$		
	$-9.6^{+1.2}_{-1.4}$	$1.7^{+1.4}_{-2.4}$	$-10.0^{+1.3}_{-1.5}$	$1.4^{+1.7}_{-2.2}$		
	$-8.0^{+1.5}_{-1.2}$	$3.274^{+0.057}_{-0.056}$	$-8.5^{+1.2}_{-1.6}$	$3.22^{+0.14}_{-1.96}$		
	$-0.84^{+0.52}_{-0.58}$	$3.010^{+0.075}_{-0.066}$	$-1.54^{+1.00}_{-6.65}$	$3.10^{+0.18}_{-0.11}$		
	0.0	$1.1^{+1.9}_{-2.0}$	0.0	$2.90^{+0.18}_{-3.08}$		

Table 3. Numerical values for the constraints on the location of knots. Companion to Figs 3 and 4.

n	$\Omega_b h^2$	$\Omega_c h^2$	h	τ
	$0.02218^{+0.00031}_{-0.00032}$	$0.1186^{+0.0029}_{-0.0030}$	$0.684^{+0.015}_{-0.013}$	$0.087^{+0.031}_{-0.029}$
0	$0.02225^{+0.00027}_{-0.00025}$	$0.1182^{+0.0027}_{-0.0029}$	$0.686^{+0.013}_{-0.012}$	$0.092^{+0.028}_{-0.027}$
	$0.02225^{+0.00027}_{-0.00025}$	$0.1182^{+0.0027}_{-0.0029}$	$0.686^{+0.013}_{-0.012}$	$0.092^{+0.028}_{-0.027}$
	$0.02217^{+0.00031}_{-0.00030}$	$0.1186^{+0.0028}_{-0.0029}$	$0.684^{+0.014}_{-0.013}$	$0.089^{+0.030}_{-0.029}$
1	$0.02226^{+0.00025}_{-0.00025}$	$0.1182^{+0.0029}_{-0.0028}$	$0.686^{+0.013}_{-0.013}$	$0.091^{+0.029}_{-0.026}$
	$0.02241^{+0.00029}_{-0.00026}$	$0.1173^{+0.0026}_{-0.0026}$	$0.691^{+0.013}_{-0.012}$	$0.107^{+0.029}_{-0.027}$
	$0.02220^{+0.00031}_{-0.00030}$	$0.1184^{+0.0028}_{-0.0028}$	$0.685^{+0.014}_{-0.013}$	$0.091^{+0.029}_{-0.030}$
2	$0.02226^{+0.00023}_{-0.00024}$	$0.1181^{+0.0025}_{-0.0026}$	$0.687^{+0.012}_{-0.012}$	$0.094^{+0.027}_{-0.026}$
	$0.02233^{+0.00032}_{-0.00033}$	$0.1179^{+0.0028}_{-0.0028}$	$0.688^{+0.014}_{-0.013}$	$0.101^{+0.033}_{-0.029}$
	$0.02220^{+0.00031}_{-0.00030}$	$0.1184^{+0.0028}_{-0.0028}$	$0.685^{+0.014}_{-0.013}$	$0.091^{+0.029}_{-0.028}$
3	$0.02225^{+0.00023}_{-0.00024}$	$0.1183^{+0.0025}_{-0.0028}$	$0.686^{+0.012}_{-0.011}$	$0.093^{+0.027}_{-0.026}$
	$0.02230^{+0.00029}_{-0.00032}$	$0.1180^{+0.0027}_{-0.0025}$	$0.688^{+0.012}_{-0.013}$	$0.101^{+0.026}_{-0.028}$
	$0.02219^{+0.00031}_{-0.00030}$	$0.1186^{+0.0029}_{-0.0029}$	$0.684^{+0.014}_{-0.013}$	$0.090^{+0.030}_{-0.029}$
4	$0.02225^{+0.00024}_{-0.00024}$	$0.1181^{+0.0026}_{-0.0025}$	$0.687^{+0.011}_{-0.012}$	$0.094^{+0.028}_{-0.027}$
	$0.02232^{+0.00029}_{-0.00030}$	$0.1178^{+0.0029}_{-0.0027}$	$0.688^{+0.013}_{-0.013}$	$0.102^{+0.030}_{-0.029}$
	$0.02220^{+0.00032}_{-0.00030}$	$0.1184^{+0.0028}_{-0.0029}$	$0.685^{+0.014}_{-0.013}$	$0.092^{+0.031}_{-0.027}$
5	$0.02225^{+0.00026}_{-0.00024}$	$0.1180^{+0.0027}_{-0.0028}$	$0.687^{+0.013}_{-0.012}$	$0.096^{+0.027}_{-0.026}$

Table 4. Cosmological parameter constraints as a function of the number of knots n at the 68% CI. For each value of n the first line corresponds to linear-spline interpolation (LS_n) with varying foreground paramters, the second line is LS_n with fixed foreground parameters, and the third line is cubic-spline interpolation (CS_n) with fixed foreground parameters. Companion to Fig. 6.