

This is the accepted manuscript made available via CHORUS. The article has been published as:

Polyspectra searches for sharp oscillatory features in cosmic microwave sky data

J. R. Fergusson, H. F. Gruetjen, E. P. S. Shellard, and B. Wallisch

Phys. Rev. D **91**, 123506 — Published 5 June 2015

DOI: [10.1103/PhysRevD.91.123506](https://doi.org/10.1103/PhysRevD.91.123506)

Polyspectra searches for sharp oscillatory features in cosmic microwave sky data

J.R. Fergusson, H.F. Gruetjen,^{*} E.P.S. Shellard, and B. Wallisch[†]

Centre for Theoretical Cosmology, DAMTP, University of Cambridge, CB3 0WA, UK

Despite numerous efforts, the search for oscillatory signatures in primordial spectra has not produced any convincing evidence for feature models to date. We undertake a thorough search for signatures of sharp features in the WMAP9 power spectrum and bispectrum as well as in the Planck power spectrum. For the first time, we carry out searches in both the power spectrum and bispectrum simultaneously, employing well-defined look-elsewhere statistics to assess significances in a rigorous manner. Developing efficient methods to scan power spectrum likelihoods for oscillatory features, we present results for the phenomenological bare sine and cosine modulations, allowing validation against existing Planck Likelihood surveys, as well as templates that include the correct sharp feature scaling. In particular, we study degeneracies between feature and cosmological parameters. For frequencies beyond the scale set by the acoustic peaks, the dependencies are realised through uninteresting adjustments of the comoving distance to last scattering. Hence, it is sufficient to keep cosmological parameters fixed and employ Gaussian approximations to the likelihood as a function of the feature model amplitude. In cases where results can be compared to the literature, our method shows excellent agreement. We supplement results from the Planck Likelihood with an analysis based on the Planck SMICA component separation map that, working on the assumption that the component separation algorithm is reliable, allows for the inclusion of a larger sky fraction. In principle, this allows us to place the most stringent constraints to date on the amplitudes of feature models in the temperature power spectrum. Invoking the WMAP bispectrum, we perform a combined power spectrum and bispectrum survey. We use and slightly generalise statistics developed in previous work to reliably judge the significance of large feature model amplitude estimates. We conclude that our results are entirely consistent with a featureless realisation of a Gaussian cosmic microwave background.

PACS numbers: 98.80.Cq, 98.80.-k, 98.80.Es

I. INTRODUCTION

In recent decades, significant advances have been made in our understanding of the early Universe. The inflationary paradigm has emerged as the best explanation of how our Universe began predicting it to be flat, isotropic and homogeneous with an approximately scale invariant power spectrum of primordial fluctuations. Apart from the tensor to scalar ratio r , which is constrained to be less than $r < 0.11$ (95% limit) [1], there is traditionally only one quantity, the spectral index n_s parametrisation deviations from scale invariance of the primordial power spectrum, that can be used to distinguish between differing inflationary models. Despite the observational evidence that $n_s \approx 0.96$ [1], constituting a deviation from scale invariance at the four-sigma level, this constraint still leaves a plethora of viable candidates.

An exciting possibility that could provide further insight into the physics driving inflation is the presence of an oscillatory scale dependence of the primordial spectra. A well-studied scenario causing such oscillatory features are violations of the slow-roll conditions during the era of horizon exit which do not spoil inflation (cf. e.g. the review article [2]). These can either arise due to sharp features in the slow-roll parameters (including the speed

of sound) [3–8], an oscillatory component in their evolution [9–12] or a combination of both (cf. e.g. Ref. [13, 14]). In particular, it has been shown that feature models can generate possibly observable non-Gaussianity with characteristic bispectrum shapes. This allows to look for signatures in higher-order correlation functions which can lower the threshold for a detection. It has been shown that the modulations to the power spectrum and the bispectrum are closely linked and typically oscillate with the same underlying frequency [2, 7–9, 15–19].

Inspired by these ideas many searches for feature models have been undertaken in WMAP and Planck cosmic microwave background (CMB) data. Most of these focused on the power spectrum either targeting specific models, e.g. Refs. [11, 20–32], or using model-independent approaches, e.g. Refs. [33–38]. A search for oscillatory signals in the bispectrum was first undertaken in WMAP data [39] and more recently using the Planck data [40]. So far, none of these searches have produced convincing evidence for the existence of features in the primordial spectra.

In this work, we will focus on the signatures of sharp features using templates that will be thoroughly discussed in Sec. II and App. A. We do not only work with phenomenological sine and cosine modulations to the spectra, but also investigate whether the correct k -dependent scalings predicted by the theory of sharp features impact the results.

The primary goals of this work are fourfold. First, we develop efficient pipelines to scan power spectrum likeli-

^{*} hfg22@cam.ac.uk

[†] b.wallisch@damtp.cam.ac.uk

hoods for the presence of sharp features. In the process, we study degeneracies between cosmological and feature model parameters and argue that, except for very low frequencies, it is sufficient to keep cosmological parameters fixed to their best-fit values in a power spectrum survey. Secondly, we construct a likelihood based on the Planck Spectral Matching Independent Component Analysis (SMICA) [41]. While this allows for a comparison and cross-validation of the two methods, the main advantage is that the SMICA map in principle provides us with a larger accessible sky fraction. Working on the assumption that the SMICA cleaning algorithm is reliable, we can extend the included sky fraction and, thus, significantly lower the error bars on feature models. This provides clues as to whether or not large signals observed in the Planck Likelihood could be candidates for actual features.

Thirdly, we combine the power spectrum results with WMAP bispectrum results in the spirit of a combined survey proposed in Ref. [42]. As will be discussed in detail below, finding large results at the same frequencies could provide us with further evidence for a feature. Finally, we employ and slightly generalise the statistics developed in Ref. [42] to rigorously assess the significance of the findings in the individual and combined surveys throughout this work.

The paper is organised as follows. We start in Sec. II by introducing and motivating the power spectrum and bispectrum templates we study. Further details are provided in App. A. We go on to describe and validate our pipelines in Sec. III. In particular, we study degeneracies between the cosmological and the additional feature model parameters in the power spectrum in Sec. III A. While Sec. III A 1 is focused on the Planck Likelihood, our power spectrum analysis of the Planck SMICA component separation map is detailed in Sec. III A 2. Section III B then provides an outline of the bispectrum pipeline employed on the WMAP data.

Our main results are presented in Sec. IV. Sections IV A and IV B discuss the results of individual Planck power spectrum and WMAP bispectrum surveys. In particular, we compare the power spectrum results obtained from the Planck Likelihood with those obtained from the SMICA map for different sky fractions to investigate whether the inclusion of more data reinforces the observed signals. Section IV C goes on to combine the power spectrum results from the WMAP Likelihood as well as the Planck Likelihood and SMICA map with the WMAP bispectrum. Throughout Sec. IV we employ the statistics developed in Ref. [42] to rigorously judge the significance of our findings and discuss to what extent the data is consistent with a featureless Gaussian CMB. To do this we slightly generalise the work in Ref. [42] as detailed in App. B. Finally, we summarise our results and conclude in Sec. V.

II. FEATURE MODELS: OSCILLATING POLYSPECTRA

As in a previous publication [42], we continue to study a linearly-spaced template with the power spectrum $P_{\mathcal{R}}(k)$ and the bispectrum shape $S(k_1, k_2, k_3)$ given by

$$\frac{\Delta P_{\mathcal{R}}}{P_{\mathcal{R},0}}(k) = A_P \sin(2\omega k + \phi_P), \quad (1)$$

$$S(k_1, k_2, k_3) = \frac{(k_1 k_2 k_3)^2}{\Delta_{\mathcal{R}}^4(k_*)} B(k_1, k_2, k_3) \\ = A_B \sin(\omega K + \phi_B), \quad (2)$$

where $K = k_1 + k_2 + k_3$, $P_{\mathcal{R},0}(k)$ is the power spectrum in the absence of any feature, $\Delta_{\mathcal{R}}^2(k) = k^3/(2\pi^2)P_{\mathcal{R},0}(k)$ is the dimensionless power spectrum and k_* is a fiducial momentum scale. We will refer to this template as PS1, Eq. (1), and BS1, Eq. (2), respectively, to distinguish them from the modified templates, PS2 and BS2, that will be introduced below. Note that for a given feature the same frequency ω appears in both the modulation to the power spectrum and the oscillatory running of the bispectrum. The frequency ω is a dimensionful quantity with units of Mpc, but, for brevity, we suppress units when quoting frequencies throughout this work. The phases ϕ_P and ϕ_B and especially the amplitudes A_P and A_B are typically model-dependent so that this template has five parameters.

Such oscillatory spectra are well-motivated theoretically. While other models are also known to produce oscillations¹, a well-explored example is the appearance of sharp features in the slow-roll parameters or the speed of sound during inflation. A rigorous treatment in the case of the power spectrum involves the Generalised Slow Roll (GSR) framework (cf. e.g. Ref. [4]) while the bispectrum is studied using the in-in formalism (cf. e.g. Ref. [2]) with possible GSR corrections to the mode functions [45].

In this work, we are mostly interested in the behaviour of the oscillatory spectra for $\omega k \gg 1$. The reason for this is that for extended oscillations, the contribution to the overall S/N (signal-to-noise) in the CMB from low multipoles is small, so that nearly all the S/N comes from $l > \mathcal{O}(10^2)$ corresponding to $k > \mathcal{O}(10^{-2})$. Low frequencies $\omega \lesssim 140$ are strongly degenerate with cosmological parameters as shown explicitly in Sec. III A 1 and we will be mainly interested in large frequencies with $\omega \gg 10^2$. Thus, for these models most of the S/N generically comes from regions with $\omega k \gg 1$ so that it is justified to scan for these models based on their behaviour for $\omega k \gg 1$.

Rather than rigorously calculating polyspectra employing the GSR method and the in-in formalism, we

¹ See for example Refs. [43, 44] for a discussion of non-standard vacuum choices that produce sinusoidal modulations to the power spectrum.

provide a simplified discussion in App. A that allows us to extract the leading-order behaviour for $\omega k \gg 1$ and motivates our templates. In the case of the power spectrum, simple solution matching across a sharp feature, as is done in App. A 1, shows that one generally expects a feature in single-field inflation to generate linearly-spaced oscillations where the leading-order behaviour for $\omega k \gg 1$ is given by a component $\sim \cos(2\omega k)$ and a suppressed component $\sim \sin(2\omega k)/(\omega k)$. This is consistent with rigorous results from the GSR approximation in the literature [5]. The relative magnitude of these terms is determined by the jumps in the slow-roll parameters at the location of the feature. This statement is general in the sense that it also applies to inflation with non-standard kinetic terms where sharp features in the speed of sound can occur. In fact, features in ϵ and c_s are (nearly) degenerate at the level of the power spectrum².

The case of the bispectrum is briefly discussed in App. A 2. Extracting the leading-order behaviour with the in-in formalism shows that the scaling of the bispectrum for a sharp feature in the slow-roll parameters is given by a component $\sim (\omega K)^2 \cos(\omega K)$ and a suppressed component $\sim (\omega K) \sin(\omega K)$ in single-field inflation. While steps in the speed of sound give rise to oscillatory signals in the bispectrum as well, they typically produce a different shape, hence breaking the near degeneracy found in the power spectrum. In this work, we focus on the former case.

We see that even though the templates PS1 and BS1, Eqs. (1) and (2), have the correct oscillatory behaviour that we encounter in the phenomenology of sharp features, they do not correctly capture the scalings of all the components. Hence, we also search for modulations that include the correct scalings using the templates

$$\frac{\Delta P_{\mathcal{R}}}{P_{\mathcal{R},0}} = A_P \left(\cos \phi_P \frac{f_P(\omega)}{\omega k} \sin(2\omega k) + \sin \phi_P \cos(2\omega k) \right), \quad (3)$$

$$S = A_B (\cos \phi_B f_B(\omega)(\omega K) \sin(\omega K) + \sin \phi_B (\omega K)^2 \cos(\omega K)), \quad (4)$$

where $f_P(\omega)$ and $f_B(\omega)$ are functions chosen to give equal S/N to the sine and the cosine components and the angles $\phi_P, \phi_B \in [0, \pi)$ parametrise the relative magnitudes. These templates will be referred to as PS2, Eq. (3), and BS2, Eq. (4), respectively.

We emphasise that while these templates should be reliable for $\omega k \gg 1$, the behaviour of the solutions for $\omega k \lesssim 1$ is much more complicated and requires a careful

treatment. We argued previously that this region contributes very little S/N. However, in the case of the template PS2, the $1/(\omega k)$ suppression can assign a larger fraction of the overall S/N to this region. To avoid this problem we will restrict our analysis for the template PS2 to multipoles $l \geq 50$, enforcing that for $\omega \gg 100$ only k with $\omega k \geq 1$ contribute to feature model amplitudes. This solution clearly comes with a loss in S/N for those feature models that have most of their support in the region $l \leq 50$. These models would benefit from a more rigorous treatment. We will discuss this point in more detail in Sec. IV A 2.

We conclude this section with a few remarks on the validity of the sharp feature limit (cf. App. A). As has been discussed in Refs. [46, 47], strictly speaking, the sharp feature limit is not under perturbative control. This can be seen naively by considering the ratio of the quadratic and cubic Lagrangian that should satisfy $\mathcal{L}_3/\mathcal{L}_2 \ll 1$. The couplings in the cubic Lagrangian diverge as the sharp feature limit is taken and violate this bound which indicates that the theory becomes strongly coupled. Hence, features that can be studied within the framework of perturbation theory can only have a small, but finite width. The finite width typically manifests itself as an exponentially decaying envelope multiplying the feature templates. Effects on wavenumbers k which where deep inside the horizon at the time of the feature are suppressed. For the frequency range we are studying it is sensible to assume that very sharp, but still perturbative, features produce a signature in the power spectrum that is largely unaffected by the envelope in the signal-dominated region of Planck. Thus, these types of features should be well described by the templates PS1 and PS2 without an envelope. We are using WMAP data with $l_{\max} = 600$ to study the bispectrum. In this case it is also sensible to assume that very sharp, but still perturbative, features are captured by the templates BS1 and BS2 without taking the envelope into account.

III. METHODS

A. Power spectrum: a dual pipeline

The analysis of the power spectrum is performed employing a dual pipeline based on the Planck 2013 data release. Our search for feature models in the power spectrum uses both the Planck power spectrum Likelihood as well as a pseudo- C_l (PCL) likelihood based on the Planck component separation maps [41]. While the former incorporates a more rigorous modelling of experimental effects such as noise anisotropies and beam uncertainties, the latter allows, for example, to easily change the included sky fraction and use larger parts of the CMB sky that, assuming successful component separation, should be clean of foregrounds. We provide an outline of the two approaches in Secs. III A 1 and III A 2, respectively.

As previously pointed out and confirmed in this work,

² In the context of the GSR framework this degeneracy between features in the inflaton potential and in the speed of sound has been demonstrated for example in Ref. [5]. The degeneracy is only broken by minor differences at small k that are observationally nearly irrelevant due to the poor S/N at low multipoles and do not contribute to our discussion of the $\omega k \gg 1$ behaviour.

the likelihood function has many local maxima associated with different frequencies ω . This is a serious obstacle for a standard Markov Chain Monte Carlo (MCMC) analysis as the likelihood is difficult to explore. Therefore, we introduce a grid in the frequency ω as has been done in previous works [23, 25, 26, 42]. The spacings were taken to be $\Delta\omega = 10$ unless stated otherwise. We also introduce a grid in the phase ϕ with $\Delta\phi = 0.1\pi$. However, as will be further discussed in Sec. III A 3, adopting the methods discussed below it usually suffices to perform the analysis only for the pure sine component, $\phi = 0$, and the pure cosine component, $\phi = \pi/2$. All other phases are straightforwardly related to these two cases.

For any given frequency ω and phase ϕ we only vary the amplitude A_P and find the best fit while keeping all cosmological, foreground and nuisance parameters fixed and set to their best-fit values without the presence of any feature. We will argue and verify in Sec. III A 1 that this is entirely sufficient for a search for linearly-spaced feature models with frequencies $\omega \gg 100$ which are beyond the oscillation patterns imprinted on the CMB due to the acoustic oscillations.

1. Planck Likelihood

a. Fast extraction of feature model amplitudes from the Planck Likelihood Unsurprisingly, keeping the six cosmological Λ CDM parameters $A_s, n_s, \Omega_b h^2, \Omega_c h^2, \theta_A$ and τ as well as the foreground and nuisance parameters fixed, the Planck Likelihood is very nearly Gaussian³ in the feature model amplitude A_P for any given ω and ϕ . We verified this empirically by plotting the likelihood, but it can also be understood in a straightforward fashion by noting that the high- l log-likelihood, which dominates the S/N, is based on the fiducial Gaussian approximation [48–50] given by

$$-2 \log \mathcal{L} = \chi^2 \equiv (\hat{C}_{l_1} - C_{l_1}) \Delta_{l_1 l_2} (\hat{C}_{l_2} - C_{l_2}), \quad (5)$$

where \hat{C}_l and $\Delta_{l_1 l_2} = \langle \hat{C}_{l_1} \hat{C}_{l_2} \rangle$ are the pseudo- C_l (PCL) estimates and covariance matrix, respectively. The covariance matrix is evaluated for a fiducial model and kept fixed, explaining the name. If we only vary the feature model amplitude, we have $C_l = C_{0,l} + A_P \delta C_l$ in linear theory and, therefore, it is evident that \mathcal{L} as a function of A_P is a Gaussian.

Rather than thoroughly exploring the likelihood using MCMC techniques, it is thus possible to simply calculate the best-fit amplitude \hat{A}_P and the variance $\sigma^2 = \langle \Delta \hat{A}_P^2 \rangle$

by fitting a Gaussian through three points. In practice, we sample the likelihood on a coarse grid in the amplitude A_P using a version of CAMB [51] modified to include the additional feature degrees of freedom. In order to correctly resolve the oscillations over the entire frequency range, we make sure that the accuracy settings are chosen appropriately⁴. In particular, we enforce calculation of the transfer functions at each l for large frequencies $\omega \geq 1000$. Then, we pick the amplitude with the lowest χ^2 from this small set of samples⁵ and two points approximately a distance σ to the left and right and calculate the corresponding Gaussian. The mean of this Gaussian is the best-fit amplitude \hat{A}_P and its variance gives $\langle \Delta \hat{A}_P^2 \rangle$. The whole process is very fast making scans over large frequency ranges feasible.

We also performed an MCMC exploration with CosmoMC [53] only varying A_P for frequencies up to $\omega = 2000$ enforcing a stringent convergence criterion of $R-1 < 0.01$ to exclude the possibility that non-Gaussian corrections to the likelihood have a large effect. We found excellent agreement of the mean and standard deviation of the posterior amplitude distribution with the Gaussian approximation described above in accordance with the expectation that the likelihood should be very nearly Gaussian.

To quantify our results, we assign a significance \bar{A}_P to the best-fit amplitudes according to

$$\bar{A}_P = \frac{\hat{A}_P}{\langle \Delta \hat{A}_P^2 \rangle^{\frac{1}{2}}}. \quad (6)$$

From a Bayesian point of view this significance measures how inconsistent the posterior amplitude distribution is with $A_P = 0$. However, this significance can also be interpreted from a frequentist point of view. Under the fiducial Gaussian approximation, the maximum-likelihood estimate \hat{A}_P is normally distributed with mean $\langle \hat{A}_P \rangle = 0$ under the null hypothesis ($A_P = 0$) with the same variance $\langle \Delta \hat{A}_P^2 \rangle$. Hence, Eq. (6) is also the frequentist significance corresponding to the p -value of measuring an amplitude at least as big as \hat{A}_P . Finally, we note for

³ While the high- l part of the likelihood should be Gaussian as pointed out below, there can be deviations due to the non-Gaussian nature of the low- l likelihood. Furthermore, non-linear corrections to the modulations of the C_l from lensing might cause very slight deviations from a perfect Gaussian shape of the likelihood as a function of A_P .

⁴ For increasing frequency, we adjust the CAMB accuracy parameters ‘accuracy_boost’ and ‘l_sample_boost’ to enforce a denser wavenumber sampling, a decrease in integration step sizes and a denser sampling in l when interpolating the C_l [52]. We explicitly checked that our settings are sufficient to resolve oscillations at a given ω by ensuring that our results are unaffected by a further increase in accuracy parameters. A detailed study of accuracy settings in the context of feature searches can be found in Ref. [23]. We generally chose to be more conservative using (accuracy_boost, l_sample_boost) = (3, 30) for $\omega < 600$, (4, 40) for $\omega < 1000$ and (8, 50) for $\omega \geq 1000$. Especially at lower frequencies this is likely to be excessively accurate and further optimisation is possible.

⁵ This is just a rough estimate for the actual best-fit amplitude due to the limited number of samples from the likelihood. We emphasise that the actual best-fit amplitude is calculated by fitting a Gaussian.

later reference that for the fiducial Gaussian approximation mentioned above, the likelihood improvement, which is often quoted, is simply given by

$$-2\Delta\log\mathcal{L} = -\bar{A}_P^2. \quad (7)$$

b. Comparison with a full MCMC analysis The analysis outlined above does not explore the possibility of further likelihood improvements coming from varying the cosmological parameters. Hence, one might be worried that the method does not identify big likelihood improvements or, equivalently, significant feature model amplitudes with large \bar{A}_P reliably. The purpose of this section is to show that for frequencies larger than the scale of acoustic oscillations, $\omega \gg 100$, varying the amplitude alone is sufficient to correctly determine the significant feature model signals and the maxima in likelihood improvement.

To do this, we compare results obtained from the Planck Likelihood by first fixing all cosmological, foreground and nuisance parameters to their best-fit values and then just varying the feature model amplitude as above to an analysis, where the six cosmological and 14 foreground and nuisance parameters⁶ are varied simultaneously with the amplitude A_P . For the latter, we assume the same priors on the cosmological, foreground and nuisance parameters as the Planck collaboration (see Tables 1 and 4 in Ref. [1]) and a flat prior with $-0.3 \leq A_P \leq 0.3$ on the feature amplitude. We run eight CosmoMC chains until they satisfied the convergence criterion $R - 1 < 0.02$ in the least-converged parameter with the actual convergence being closer to 0.01 and below for most cases. Additionally, we also checked that the results are unaffected by a more stringent convergence criterion.

Let us first formulate some expectations to interpret the results. As already discussed in Ref. [42], we expect strong degeneracies between feature models and the cosmological parameters below the scale set by the sound horizon at last scattering (LS), $\omega \approx 140$. The oscillations due to features effectively alter the acoustic peak structure which can mimic the effects of changing other parameters. These degeneracies are expected to largely disappear for $\omega > 140$. However, as has been mentioned elsewhere (see e.g. Ref. [25]), even for high-frequency oscillations there still is a further subtle effect that can lead to dependencies between the feature model amplitude and the cosmological parameters. For a feature model with primordial frequency ω in k -space, the frequency ω_l in l -space observed in the power spectrum is approximately given by

$$\omega_l \sim \frac{2\omega}{\eta_*}, \quad (8)$$

⁶ We note that the other foreground and nuisance parameters incorporated in the Planck Likelihood are analytically marginalised over as in Ref. [1].

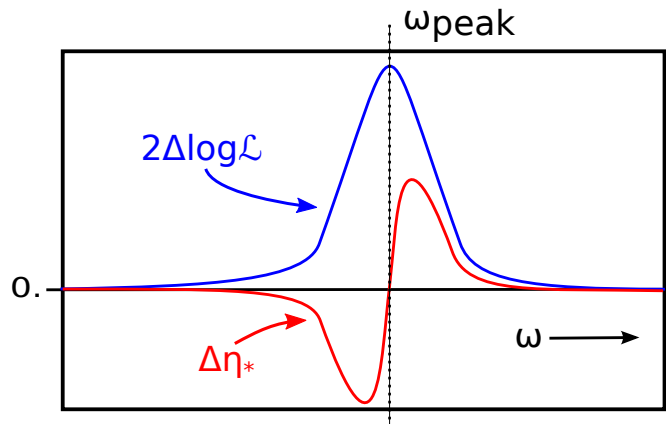


FIG. 1. An idealised sketch of the behaviour of the comoving distance to LS η_* in the vicinity of a peak in the feature model likelihood. The blue curve represents the likelihood improvement as a function of ω measured when only varying the amplitude and keeping all cosmological parameters fixed to their best-fit values. The red curve shows the behaviour of $\Delta\eta_* = \eta_* - \eta_{*,0}$ attempting to tune the resulting l -space frequency ω_l to match $\omega_{l,\text{peak}}$.

where η_* is the comoving distance to LS. The latter depends on the expansion history of the universe and thus on the Λ CDM parameters. For a given point on the grid of feature models, the frequency ω is fixed. However, if we decide to vary the cosmological parameters, we effectively allow some freedom in the frequency that is ultimately compared to the data. If there is a large signal at some $\omega_{l,\text{peak}}$ in the data, we see a peak at $\omega_{\text{peak}} \sim \eta_{*,0}\omega_{l,\text{peak}}/2$ in the computations where the cosmological parameters are kept fixed. Here, $\eta_{*,0}$ is the best-fit value of η_* without allowing for feature models that we used in these runs. For this particular value $\omega = \omega_{\text{peak}}$ we expect that allowing the cosmological parameters to vary will not have any impact on the best fit. It should still have the same feature model parameters and we should still find a best fit with $\eta_{*,BF} = \eta_{*,0}$. However, if the frequency is slightly lower than ω_{peak} , we expect that the best-fit cosmological parameters get shifted such that η_* decreases in an attempt to arrange $\omega_l \sim 2\omega/\eta_* \sim \omega_{l,\text{peak}}$. Similarly, if ω is slightly larger than ω_{peak} , we expect η_* to increase in order to produce a better fit to the feature in the data. These expectations are summarised in Fig. 1 where the behaviour of $\Delta\eta_* = \eta_* - \eta_{*,0}$ around a peak in the likelihood improvement is shown schematically.

With this heuristic picture in mind we proceed to study results extracted from the Planck Likelihood. In this section, we decrease the frequency spacing to $\Delta\omega = 5$ for $\omega \in [220, 450]$ (even smaller values are chosen in the vicinity of the likelihood peaks) in order to better resolve the peaks in the amplitude estimates and expected jumps in $\Delta\eta_*$. Figure 2 compares the amplitudes obtained with and without varying the cosmological, foreground and nuisance parameters. The amplitude estimates \bar{A}_P are plotted for $\phi = 0.5\pi$ where the feature template reduces

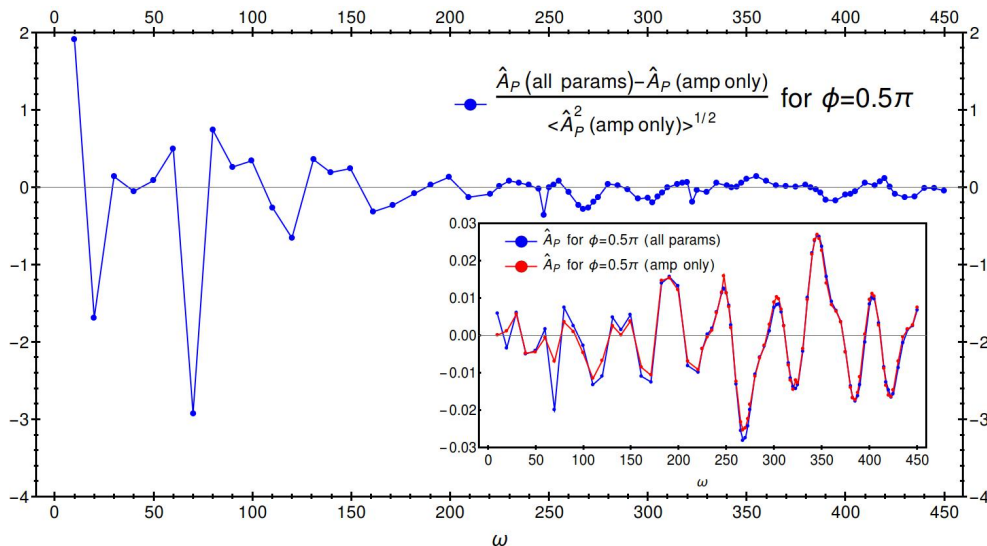


FIG. 2. Comparison of the oscillation amplitudes obtained from an MCMC analysis of the Planck Likelihood sampling over all cosmological, foreground and nuisance parameters (all params) and the fast method where all parameters are kept fixed and the likelihood is assumed to be Gaussian in the amplitude A_P (amp only). In the former case, we take \hat{A}_P to be the mean of the posterior amplitude distribution as we find this to be numerically more stable than the best fit. In the latter case, we use the best-fit amplitudes obtained from the Gaussian approximation. We checked that this gives almost exactly the same answer as the mean of the corresponding posterior distribution due to the near perfect Gaussianity of the likelihood in this case (see discussion in the main text). For each frequency the results are displayed for $\phi = 0.5\pi$. The main plot shows the difference between the amplitudes normalised by the standard deviation as calculated from the Gaussian approximation. A direct comparison of the amplitudes is shown as an inset. Good agreement is found for $\omega \gtrsim 200$.

to a bare cosine, $\Delta P_{\mathcal{R}} \sim \cos(2\omega k)$. For $\omega \gtrsim 200$ this shows good agreement between the values obtained from the two approaches. We emphasize that this is not only true for the slice shown, but also for all other values of the oscillation phase ϕ . Evidently, any remaining degeneracies have little effect on the measured amplitudes for large frequencies.

However, the degeneracy with η_* discussed above can indeed be observed at higher ω . Figure 3 shows the likelihood improvements for ω in the range $[230, 445]$ obtained from just varying the amplitude together with the values of η obtained from runs where all parameters are varied. The qualitative behaviour shown in Fig. 1 is clearly recognisable at the likelihood peaks for $\omega \gtrsim 300$. For smaller frequencies this effect can still be observed, but is less prominent due to more complicated degeneracies in this region not captured by this simple model. The presence of these residual low- ω degeneracies is also suggested by the small differences in the amplitude measurements near the peaks in the region $200 \lesssim \omega \lesssim 300$ in Fig. 2.

Summing up, for $\omega \gg 140$ a search for features where only the amplitude is varied is sufficient and any remaining degeneracies with cosmological parameters are small. Our results suggest that any further effects can be explained by an adjustment of the comoving distance to LS. Such an effect does not lead to bigger peaks in the likelihood improvement. It only allows frequencies close to a peak to benefit slightly, but does not change the likelihood improvement measured for the actual peak fre-

quency.

2. Fast quadratic estimate based on the SMICA map

To search for feature models in the Planck SMICA component separation map [41] we construct a simple PCL likelihood as in previous work [42]. Using the SMICA map has the advantage that we can extend the included sky fraction compared to the Planck Likelihood [1, 50] allowing in principle for more stringent tests of feature models⁷.

As has become standard, we use a PCL likelihood based on cross correlators to analyse the power spectrum [1, 50, 54]. The Planck component separation analysis published two maps for each foreground cleaning method as part of the 2013 data release. A map of the full CMB temperature sky (which is supposed to be clean of foregrounds) that we will refer to as $T(\hat{n})$ and a half-ring-half-difference (HRHD) map that we refer to as $N(\hat{n})$. The latter is obtained by running the foreground cleaning pipeline on the data from the first and second half of each stable pointing period and taking half the difference

⁷ Whether the results are reliable depends on the performance of the foreground cleaning scheme in regions of higher contamination. We will discuss this point in more detail below.

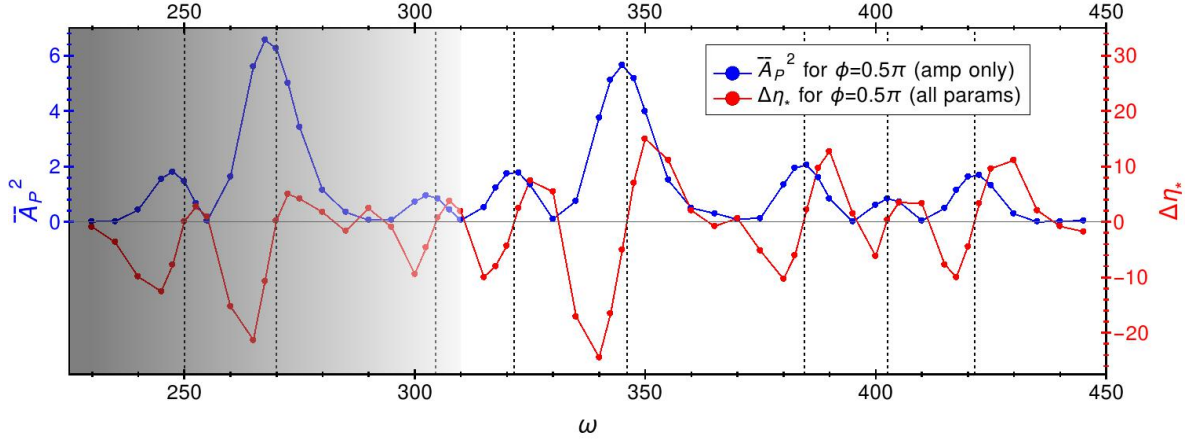


FIG. 3. The behaviour of the comoving distance to LS, η_* , as a function of frequency ω (red). η_* is obtained by varying all cosmological, foreground and nuisance parameters as well as the amplitude A_P . We plot $\Delta\eta_* = \eta_* - \eta_{*,0}$ where $\eta_{*,0}$ is the best-fit value setting $A_P = 0$. Also plotted is the likelihood improvement \bar{A}_P^2 obtained by keeping all parameters fixed and employing a Gaussian approximation to the likelihood as discussed in the main text (blue). For each frequency the result is plotted for $\phi = 0.5\pi$. The shaded area indicates where one expects further residual degeneracies with cosmological parameters and the idealised picture presented in Fig. 1 does not fully apply. For $\omega \gtrsim 300$ one can clearly identify the predicted behaviour of η_* in the vicinity of the likelihood peaks. Dashed vertical lines are drawn at the approximate locations of peak frequencies ω_{peak} as a guide.

of the results. It can be thought of as an estimate of the noise in the final temperature map [41].

The data enters the likelihood through the power spectrum estimates

$$\hat{C}_l^X = \hat{C}_l^T - \hat{C}_l^N, \quad (9)$$

where \hat{C}_l^T and \hat{C}_l^N are the power spectra of the T and N maps and we use the standard PCL estimates obtained via

$$\hat{C}_{l_1} = (M^{-1})_{l_1 l_2} \tilde{C}_{l_2}, \quad \tilde{C}_l = \frac{1}{2l+1} \sum_m |a_{lm}|^2, \quad (10)$$

with $M_{l_1 l_2}$ the standard PCL coupling matrix for a given mask [55]. The subscript X indicates that we think of this power spectrum effectively as a cross spectrum. The rationale behind this is the following: If we define the maps

$$T_1(\hat{n}) = \frac{T(\hat{n}) + N(\hat{n})}{2}, \quad T_2(\hat{n}) = \frac{T(\hat{n}) - N(\hat{n})}{2} \quad (11)$$

and think of the T map as the half-ring-half-sum (HRHS) map, then the half-ring (HR) maps T_1 and T_2 correspond to the foreground cleaned maps obtained from the first and second half of each stable pointing period. The cross-correlator between these two maps is then given by

$$\begin{aligned} \tilde{C}_l^{T_1 \times T_2} &= \frac{\sum_m (a_{lm}^1)^* a_{lm}^2}{2l+1} \\ &= \frac{\sum_m ((a_{lm}^T)^* a_{lm}^T - (a_{lm}^N)^* a_{lm}^N)}{2l+1} \equiv \tilde{C}_l^X. \end{aligned} \quad (12)$$

This argument is obviously only entirely correct if the foreground cleaning procedures are exactly linear in all

input maps, which is not the case for all foreground cleaning schemes. However, for the purpose of this analysis we will assume this is a good approximation and treat \hat{C}_l^X as a cross-correlator that should then also carry no noise bias⁸.

As in Ref. [42] we apodise the masks by approximate convolution with a Gaussian beam of FWHM 0.5° to minimise leakage using the procedure outlined in Ref. [56]. We approximate the PCL log-likelihood with the fiducial Gaussian approximation [48–50] introduced in Eq. (5). As the fiducial model we simply use a smoothed version of the power spectrum \hat{C}_l^X together with a smoothed version of the power spectrum \hat{C}_l^N to model the noise contribution to the covariance⁹. We employ the analytic approximations from Ref. [48] to calculate the covariance matrices. These approximations assume an approximately constant power spectrum. To account for small

⁸ There is another caveat here that there can be correlations between the noise in the first and second half of the pointing periods. Such correlations spoil the independence of the noise realisation in the two HR maps and would thus lead to a noise bias in the cross-correlator. This minor issue leads to small underestimates of the noise in the HR maps and was discussed in Ref. [41]. We will ignore this here.

⁹ The fiducial model and the noise spectrum only enter the covariance approximation and should roughly correspond to the true power spectrum underlying the data and the actual noise power spectrum. To obtain good approximations, we simply use the PCL estimates of these spectra given by \hat{C}_l^X and \hat{C}_l^N . As we are only including multipoles with $l \geq 50$ the scatter in these estimates is not very significant. However, to further reduce this scatter, we smooth the PCL spectra $l(l+1)\hat{C}_l^X$ and \hat{C}_l^N over a width $\Delta l \approx 20$ by convolution with a Gaussian.

leakage effects we correct for slight underestimates of the variance using an improved analytic approximation [56].

In order to calculate the covariance matrices we have to assume a noise model for the HR maps. In the analysis below we just model the noise as isotropic and Gaussian with a power spectrum given by a smoothed version of $2\hat{C}_l^N$. While the noise is clearly not isotropic in reality¹⁰, we find that this assumption does not seem to affect the feature searches presented in this work at any significant level¹¹. Due to significant deviations from a Gaussian distribution at low l the fiducial Gaussian approximation is not reliable in this region. Hence, we only consider the multipole range $50 \leq l \leq 2000$. As we are mainly interested in extended oscillations, the loss in S/N from discarding the low multipoles is negligible. Naively, assuming $l_{\max} \sim 2000$, the region $2 \leq l \leq 50$ should contribute only a fraction of $\mathcal{O}(10^{-3})$ to the total S/N. However, this estimate does not take into account the specific shape of the oscillatory feature model templates that are increasingly damped in amplitude due to the convolution with the transfer functions and lensing. Furthermore, it is well-known that there is a dip in the power spectrum at $20 \lesssim l \lesssim 30$ [1, 50] that may have some weight. To quantify this, we explicitly checked the influence of the low- l likelihood on the results for the PS1 template in the Planck Likelihood analysis and found shifts in the amplitudes of at most ~ 0.3 sigma and usually far below. This is a rather small effect although larger than the naive S/N consideration suggests, which is likely due to the anomaly mentioned previously.

In linear theory, especially ignoring lensing, the observed CMB power spectrum given the six Λ CDM parameters p_i and a feature model is given by

$$C_l(p_i, A) = C_l^{\Lambda\text{CDM}}(p_i) + A_P \delta C_l(p_i, \omega, \phi). \quad (13)$$

If $A_P \delta C_l$ is a small contribution to the power spectrum, which it always is in this work, the lensed power spectrum can be written in the same way by linearising the effect of lensing and defining $\delta C_l := \partial C_l^{\text{lensed}} / \partial A_P$. In principle, determining the maximum-likelihood estimate for the amplitude $\hat{A}_P \equiv \hat{A}_P^{\text{ML}}$ requires all parameters to be varied simultaneously. As we showed in Sec. III A 1, it is sufficient for the purpose of this study to set the Λ CDM parameters to their best-fit values obtained by assuming a featureless model and only vary the amplitude A_P for any given ω and ϕ . We thus compute the

lensed feature model templates for each point on the grid assuming the best-fit Λ CDM cosmology employing CAMB [51] with sufficiently high precision settings to ensure that the oscillations are accurately resolved.

As we only vary the amplitude A_P , the best fit can be found as a simple quadratic estimate¹² given by

$$\hat{A}_P = \frac{2}{N_P} \delta C_{l_1}(\Delta^{-1})_{l_1 l_2} \left(\hat{C}_{l_2} - C_{l_2}^{\Lambda\text{CDM}} \right), \quad (14)$$

$$N_P = 2 \delta C_{l_1}(\Delta^{-1})_{l_1 l_2} \delta C_{l_2} \quad (15)$$

for any ω and ϕ . The variance of the estimates is simply given by $\langle \hat{A}_P^2 \rangle = 2/N_P$ so that we can extract the normalised amplitude estimates, Eq. (6), via

$$\bar{A}_P = \left(\frac{2}{N_P} \right)^{\frac{1}{2}} \delta C_{l_1}(\Delta^{-1})_{l_1 l_2} \left(\hat{C}_{l_2} - C_{l_2}^{\Lambda\text{CDM}} \right). \quad (16)$$

This leads to very quick scans over the full frequency range. The estimator is essentially the optimal quadratic estimator for the amplitude A_P except for well-known and small suboptimalities due to the slightly lossy data compression in PCL power spectrum estimates (cf. Ref. [57] and references within).

3. A comment on the phase ϕ

Even though the methods discussed above are sufficiently fast that one can simply introduce a grid in the phase ϕ as well and obtain significances $\bar{A}_P(\omega, \phi)$ for each point on a sufficiently dense (ω, ϕ) -grid, we emphasise that this is typically not necessary. This is particularly evident from Eq. (16). For all feature models considered in this work it can be written as

$$\bar{A}_P(\phi) = \frac{N_P(0)^{\frac{1}{2}}}{N_P(\phi)^{\frac{1}{2}}} \cos \phi \bar{A}_P(0) + \frac{N_P(\pi/2)^{\frac{1}{2}}}{N_P(\phi)^{\frac{1}{2}}} \sin \phi \bar{A}_P(\pi/2) \quad (17)$$

suppressing the frequency ω . The normalisation factors can be easily calculated and the dependence on the data is only through the two estimates $\bar{A}_P(0)$ and $\bar{A}_P(\pi/2)$. Further simplifications occur due to the fact that the sine and cosine components are in all cases very nearly uncorrelated already for moderately high frequencies, $\langle \hat{A}_P(\omega, 0) \hat{A}_P(\omega, \pi/2) \rangle \approx 0$, and, furthermore, we have $\langle \hat{A}_P^2(\omega, 0) \rangle \approx \langle \hat{A}_P^2(\omega, \pi/2) \rangle$. The latter holds automatically for the bare sine and cosine modulations of the template PS1 and is arranged through the $f_P(\omega)$ factor in the case of the template PS2. This implies

¹⁰ Anisotropy in the noise arises at the very least due to the highly anisotropic scanning strategy of Planck. To a lesser extent it is also conceivable that the foreground cleaning methods introduce anisotropies in the noise patterns of the final maps.

¹¹ To test this, we also constructed a likelihood assuming an anisotropic noise pattern based on the average hit counts of the 143 and 217 GHz maps. While the χ^2 values shift slightly, the significances with which various feature models are detected seem to be largely unaffected. Therefore, we stick to a simple isotropic model.

¹² We introduce a redundant factor of 2 in the definition of the quadratic estimator here for consistency with the standard optimal power spectrum estimator and the optimal bispectrum estimator later on. This definition implies $\langle \hat{A}_P^2 \rangle = 2!/N_P$ in line with $\langle \hat{A}_B^2 \rangle = 3!/N_B$ for the bispectrum.

$N_P(\phi) \approx N_P(0) \approx N_P(\pi/2)$ so that we arrive at the simple relation

$$\begin{aligned} \bar{A}_P(\omega, \phi) &= \frac{\hat{A}_P(\omega, \phi)}{\langle \hat{A}_P^2(\omega, \phi) \rangle^{\frac{1}{2}}} \\ &= \cos \phi \bar{A}_P(\omega, 0) + \sin \phi \bar{A}_P(\omega, \pi/2). \end{aligned} \quad (18)$$

Even though we made use of the form of the quadratic estimator, we expect that the same reasoning can be applied to the Planck Likelihood owing to the fact that it is largely based on the fiducial Gaussian approximation that gave rise to Eq. (14). We checked this explicitly and for all the results presented in Sec. IV this is an extremely good approximation.

Note that with minimal modifications, this discussion

also applies to the bispectrum estimator that will be discussed below. Refs. [58, 59] make use of this property in their search for oscillatory models in the bispectrum where it is sufficient to evaluate the estimator for the sine and the cosine component only.

B. Bispectrum: optimal f_{NL} estimator

To constrain feature models via the bispectrum we use a modified version of the modal polynomial pipeline that was used in the 2013 Planck analysis [40, 60–62]. This is an implementation of the standard optimal bispectrum estimator. The optimal estimator for the bispectrum amplitude of a feature model \hat{A}_B , in the diagonal covariance approximation, reads

$$\hat{A}_B = \frac{1}{N_B} \sum_{l_i m_i} \frac{\mathcal{G}_{m_1 m_2 m_3}^{l_1 l_2 l_3} b_{l_1 l_2 l_3} (a_{l_1 m_1} a_{l_2 m_2} a_{l_3 m_3} - 3 \langle a_{l_1 m_1} a_{l_2 m_2} \rangle a_{l_3 m_3})}{C_{l_1} C_{l_2} C_{l_3}}, \quad (19)$$

where b is the theoretical bispectrum of the feature model defined by

$$\langle a_{l_1 m_1} a_{l_2 m_2} a_{l_3 m_3} \rangle = B_{m_1 m_2 m_3}^{l_1 l_2 l_3} = \mathcal{G}_{m_1 m_2 m_3}^{l_1 l_2 l_3} b_{l_1 l_2 l_3} \quad (20)$$

and \mathcal{G} is the Gaunt integral, which is the projection of the angular part of the primordial delta function. The normalisation of the estimator is

$$N_B \equiv \sum_{l_i l_i} \frac{(\mathcal{G}_{m_1 m_2 m_3}^{l_1 l_2 l_3} b_{l_1 l_2 l_3})^2}{C_{l_1} C_{l_2} C_{l_3}}. \quad (21)$$

The pipeline breaks the bispectrum being constrained into a set of orthonormal separable basis bispectra which dramatically reduces computation time and allows us to constrain all frequencies, within resolution, simultaneously. The approach was first described in Ref. [63] and a fully realised version was first implemented in Ref. [64]. It was recently extended to polarisation in preparation for the next round of Planck papers in Ref. [65], which also included many other small advances. Here, we use the temperature only version of this pipeline on WMAP9 year data [66] restricting ourselves to $l_{\text{max}} = 600$. This reduction in range coupled with the increase in the number of basis functions from 600 to 2000 allows us to cover a frequency range six times larger than in the first Planck analysis extending up to $\omega = 1000$. For the WMAP data we use the weighted average of the V and W channels with weights 1.0 and 0.9, respectively. The linear term and variance is computed from 500 simulations generated with the fiducial power spectrum combined with white anisotropic noise created from the coadded hit-count maps. The simulations were then masked and diffusively inpainted to reduce mode coupling in the multipoles. We use a frequency grid with a stepwidth of

$\Delta\omega = 20$, which is sufficient for a WMAP-type survey as discussed in Sec. IV B and App. B.

IV. RESULTS

A. Power spectrum surveys

1. Bare sine and cosine: results for the template PS1

We present the results for the template PS1 for both the Planck Likelihood and the SMICA map masked with the U73 mask in Fig. 4. Each plot consists of a top panel with a density plot of the normalised amplitude \bar{A}_P as defined in Eq. (6) and a bottom panel showing the maximum \bar{A}_P^2 found for any phase ϕ at a given frequency ω . We remind the reader that the latter corresponds to the maximum likelihood improvement observed at that frequency.

The results show clear similarities. However, there are evidently differences. We emphasise that due to the different sky fractions included in the analysis one should not expect the results to match exactly¹³. Even if we assume that a given method exactly accounts for all systematics and noise properties of the data, inclusion of more data necessarily shifts the observed peaks. The Planck Likelihood is based on the CL49 and CL31 mask

¹³ In addition, the low- l likelihood is included in the Planck Likelihood search, but not in the SMICA analysis. As already discussed in Sec. III A 2, this can give rise to further minor deviations in the measured amplitudes.

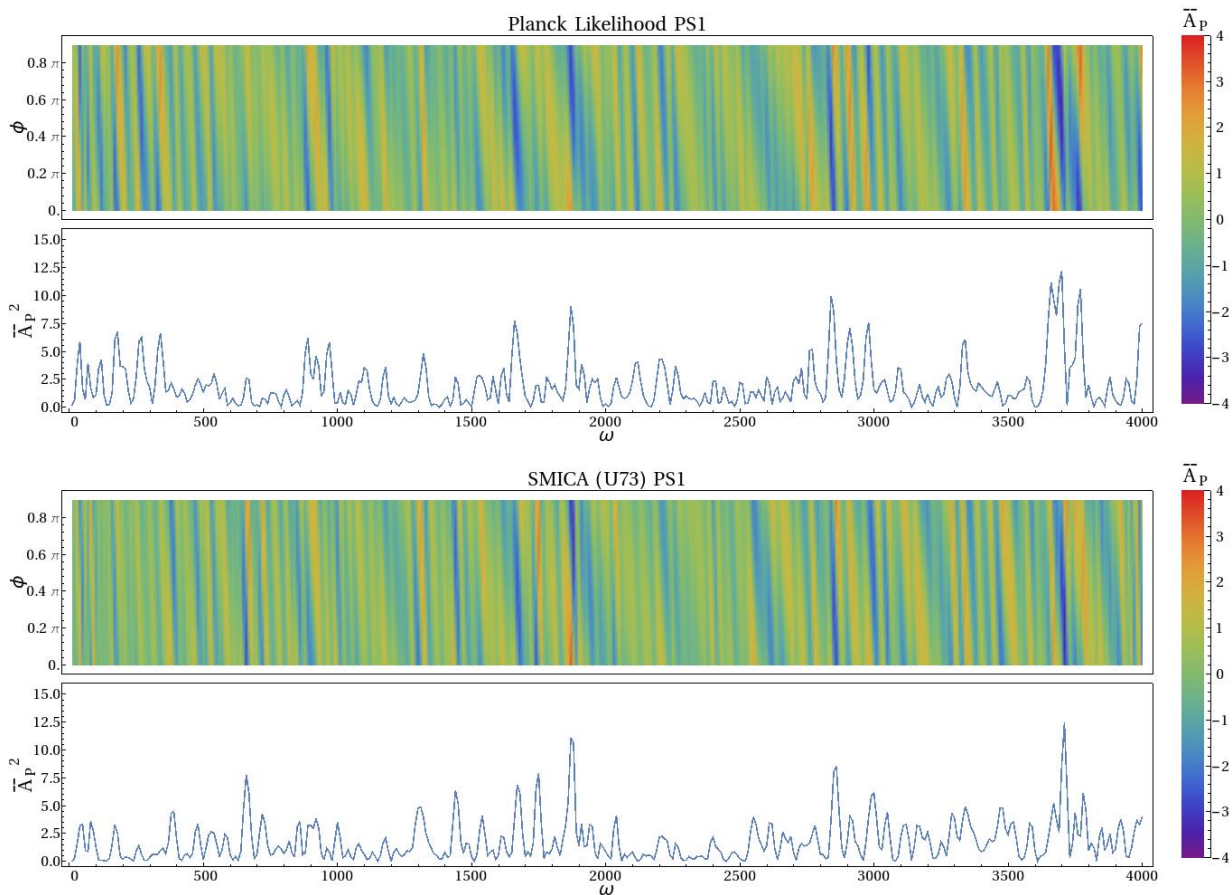


FIG. 4. Amplitudes \bar{A}_P for the template PS1 up to $\omega = 4000$ as obtained from the Planck Likelihood (top) and based on the SMICA map masked with an apodised version of the U73 mask with $f_{\text{sky}} = 0.61$ (bottom). In each plot the bottom panel shows the maximum \bar{A}_P^2 at a given ω , corresponding to the maximum likelihood improvement at that frequency.

with a sky fraction of $f_{\text{sky}} = 0.49$ and $f_{\text{sky}} = 0.31$, respectively, so that significantly more data is included¹⁴ in the SMICA analysis with $f_{\text{sky}} = 0.61$. Note that this implies that, assuming the foreground cleaning procedure is reliable, tighter constraints on feature models can be obtained from the SMICA map (cf. Fig. 7). We quantified how correlated the results are and arrived at the conclusion that the differences between the Planck Likelihood and SMICA results can be explained by the differences between the masks used in the two analyses.

The results for the Planck Likelihood are in very good agreement with the corresponding Fig. 3 in Ref. [26] where a similar search has been performed using a different method. Their approach relies on a Taylor expansion

of the power spectrum in the cosmological parameters which allows for faster sampling when varying all parameters as the derivatives of the transfer functions can be precomputed [25]. Then, a Metropolis-Hastings algorithm is used to find the best fit. By adopting sampling schemes tailored to the problem, further improvements in computational efficiency are possible [27]. Note that we argued previously in Sec. III A 1 b that this is not necessary for linearly-spaced oscillations and accurate amplitude estimates as well as likelihood improvements can be extracted using the numerically very efficient and reliable methods proposed in Sec. III that keep cosmological parameters fixed to their best-fit values. The striking agreement with Ref. [26] is a further validation of this claim.

¹⁴ One should keep in mind that sky fractions do not fully reflect the difference in the underlying datasets. The CL31 and CL49 masks are not simply larger versions of the U73 mask, but there are parts of the sky that are masked by the U73 mask, but not by the CL31 or CL49 mask. This leads to larger differences between the results than might be expected just based on the ratio of sky fractions.

We find the overall best fit at $\omega \sim 3710$ using both our approaches with a significance of about 3.5 sigma corresponding to $2\Delta \log \mathcal{L} \approx 12$. Reference [26] reports the maximum likelihood improvement to be $2\Delta \log \mathcal{L} \approx 13$ at $\omega = 3670$ (note that their frequency definition differs from ours by a factor of two), which is part of the same

peak structure¹⁵.

The natural question that arises when studying the results is whether or not the large signals in various places present significant evidence for feature models. A generic property of feature model surveys is that one scans over a large number of effectively independent models. Each of these can fit fluctuations in the noise¹⁶ by chance so that we expect large results in a large survey simply because we compared a vast number of models to the data. The look-elsewhere effect that arises for the oscillatory feature models under investigation was studied in Ref. [42]. There it was argued that the distribution of the maximum significances observed in an individual survey can be well-described analytically and gives rise to a look-elsewhere-adjusted significance according to

$$S = 2^{\frac{1}{2}} \text{Erf}^{-1} \left[(F_{\chi, 2}(\bar{A}_P))^{N_{\text{eff}}} \right]. \quad (22)$$

Here, N_{eff} quantifies the effective number of independent models and $F_{\chi, 2}(x)$ is the cumulative distribution function (CDF) of the χ -distribution with two degrees of freedom. Equivalently, one can define an effective step width in frequency $\Delta\omega_{\text{eff}}$ according to

$$\Delta\omega_{\text{eff}} = \frac{\omega_{\text{max}} - \omega_{\text{min}}}{N_{\text{eff}} - 1}. \quad (23)$$

This effective step width is related to the degree of correlation between nearby frequencies and can be understood as a rough estimate of the separation in frequency ω at which models become effectively independent. For the type of models under investigation $\Delta\omega_{\text{eff}}$ is independent of frequency and depends mainly on l_{max} of the experiment.

In App. B we show that we have $\Delta\omega_{\text{eff}} \approx 13$ for a Planck-like set-up and, therefore, we arrive at $N_{\text{eff}} \approx 300$ for our survey range. Note that this also justifies our choice of frequency step width $\Delta\omega = 10$. Due to the strong correlations between frequencies separated by less than $\Delta\omega_{\text{eff}}$ this step width should be sufficiently small to resolve all peaks in the likelihood improvement. Figure 5 shows the analytic relation between the maximum raw significance observed and the corresponding look-elsewhere-adjusted significance for $N_{\text{eff}} = 300$ with the respective maximum values from the Planck Likelihood and SMICA analysis highlighted. The look-elsewhere-adjusted significances in both cases are clearly below the one sigma level. Loosely speaking this implies that one

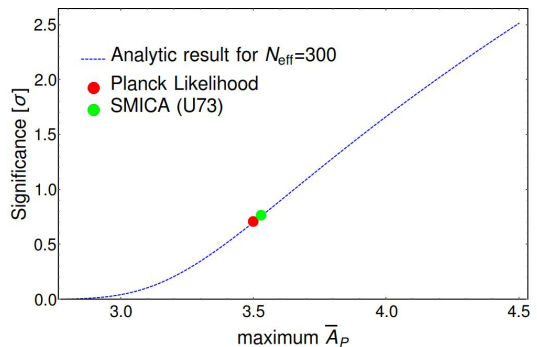


FIG. 5. Analytic result for the distribution of maximum significance feature model amplitudes in a survey with $N_{\text{eff}} = 300$. The values obtained from the Planck Likelihood and the SMICA analysis are highlighted.

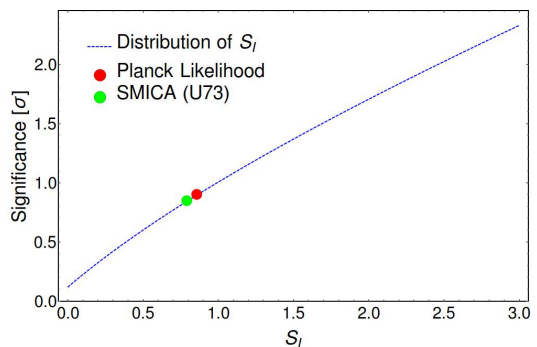


FIG. 6. Prediction of the distribution of the integrated statistic S_I . The values obtained from the Planck Likelihood and the SMICA analysis assuming $N_{\text{eff}} = 300$ are highlighted.

expects roughly every other random realisation of a featureless CMB to give rise to a maximum significance at least as big as what we observe in our CMB. Hence, we cannot conclude that the maximum likelihood improvements in the power spectrum on their own present convincing evidence for feature models at the respective frequencies.

In Ref. [42] a further test was suggested that addresses the possibility that the data could present evidence for feature models that give rise to modulations with multiple well-separated frequencies. This is different from simply looking at the maximum significance found in a survey as the height of all other peaks are taken into account as well. It was found that the integrated statistic S_I given by

$$S_I^2 = \frac{\Delta\omega}{\Delta\omega_{\text{eff}}} \sum_{\omega} 2 \text{Erf}^{-1} \left[(F_{\chi, 2}(\bar{A}_{P, \omega}))^{N_{\text{eff}}} \right]^2 \quad (24)$$

produces significances that agree well with a rigorous look-elsewhere analysis for multi-frequency models. Here, $\Delta\omega$ is the step width in frequency of the survey and the sum reduces to an integral over frequency in the $\Delta\omega \rightarrow 0$ limit. Figure 6 shows the distribution of S_I

¹⁵ The peak observed at this frequency in the Planck Likelihood analysis is only marginally smaller than our best fit and recorded at 3.3 sigma ($2\Delta \log \mathcal{L} \approx 11$) making it a competing maximum. In fact, the authors of Ref. [26] also obtained this structure with three nearby large peaks, whose likelihood improvements differ slightly leading to a different overall best-fit value.

¹⁶ Noise in this context refers to both the cosmic variance and experimental noise. Both cause scatter in the PCL estimates that can give rise to good feature model fits by chance.

taken from¹⁷ Ref. [42] with the values obtained from the Planck Likelihood and the SMICA analysis highlighted. Again, both results are below the one sigma level indicating that the survey results are consistent with a random realisation of a featureless Gaussian CMB.

To sum up, the results above indicate that neither the maximum likelihood improvements nor the abundance of further large peaks present convincing evidence for a detection of features in the power spectrum based on the Planck Likelihood or the SMICA map. Of course, this does not exclude the possibility that the primordial power spectrum exhibits oscillatory features and with the inclusion of more data a different outcome with a positive detection might be reached.

There are several ways to make progress at this stage. In Sec. IV C we will combine our survey with bispectrum results. As feature models should also produce signatures in the bispectrum this can provide further evidence. Without invoking the bispectrum one can only include more data, for example the Planck polarisation data, to lower the error bars on feature model amplitudes¹⁸. Without the polarisation data at hand, the only way to increase the amount of data is the inclusion of a larger sky fraction. This is not possible in the framework of the Planck Likelihood as, by construction, it only operates on the cleanest parts of the sky where foregrounds can be modelled as effective contributions to the power spectrum.

However, the SMICA component separation algorithm produces valid¹⁹ results on more than 80% of the sky, defined by the validation mask²⁰ of the algorithm. The validation mask was published as part of the 2013 Planck data release and we constructed an apodised version of it with $f_{\text{sky}} = 0.81$. The larger sky fraction evidently leads to a significant reduction in error bars as can be seen in Fig. 7. If the large results we observe at various frequencies are in fact signatures of real features of the primordial power spectrum, we expect their significances to rise. This expectation obviously necessitates that we trust the SMICA algorithm over the region previously excluded by the U73 mask, but not by the validation mask.

¹⁷ We note that the fit stated in Ref. [42] contains a minor error. The CDF of the statistic S_I , $F_{S_I}(x)$, is well described by

$$F_{S_I}(x) = 1 - \exp(-(ax^2 + bx + c)) \quad (25)$$

where $c = 0.102$ as in Ref. [42] but $a = 0.108$ (rather than $a = 0.092$) and $b = 0.950$ (rather than $b = 0.876$). We use the correct values in this work even though the small changes make little difference in the cases considered here.

¹⁸ See for example the discussions in Refs. [42, 67, 68] and references therein.

¹⁹ As measured by a somewhat subjective criterion as described in Ref. [41] and references therein.

²⁰ The U73 mask is the union of the validation masks of the four component separation algorithms used in the Planck analysis. Some of these have significantly smaller sky fractions resulting in the smaller sky fraction of the union mask.

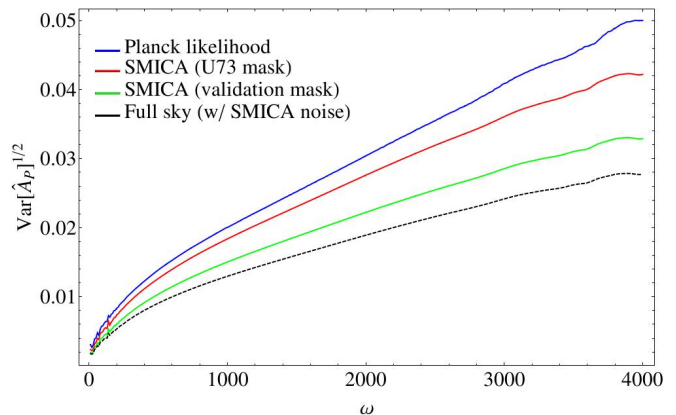


FIG. 7. The standard deviations of the amplitude measurements for the various analysis methods presented in this work. The Planck Likelihood uses the CL31 ($f_{\text{sky}} = 0.31$) and CL49 ($f_{\text{sky}} = 0.49$) masks while the SMICA analysis was carried out with an apodised version of the U73 mask ($f_{\text{sky}} = 0.61$) and the validation mask ($f_{\text{sky}} = 0.81$). The standard deviation evidently decreases with growing sky fraction or equivalently more included data.

The results obtained using the validation mask in an otherwise identical analysis that led to the SMICA results in Fig. 4 are presented in Fig. 8. The large peaks observed in Fig. 4 shrink rather than grow. This leads us to conclude that either the regions that were previously excluded are not faithful representations of the actual CMB due to a failure of the component separation algorithm to isolate the CMB component or the large signals are simply the result of fitting the scatter of the \hat{C}_l by chance.

2. Including the sharp feature scaling: results for the template PS2

As pointed out in Sec. II, the template PS1 given by Eq. (1) that we used in the previous section is not entirely appropriate when looking for the signatures of sharp features in the power spectrum. For these kinds of features the sine component of the signal is generically suppressed by a factor of $1/(\omega k)$ giving rise to the template PS2, Eq. (3). Hence, it is interesting to see whether the large peaks observed in the previous section carry over to this case.

Before presenting the results we would like to draw attention to an important qualifier. The way we set up the search for signatures of feature models places some restrictions on the type of sharp features we can look for. First of all, the templates are only supposed to capture the behaviour accurately for $\omega k \gg 1$. As explained in Sec. II for the template PS1 and frequencies $\omega \gg 100$ this is never really a problem as all the S/N comes from wavenumbers k satisfying this condition. While this conclusion carries over to the cosine part of the PS2 template, the decaying sine part is more problematic. In particular, those models of the PS2 template

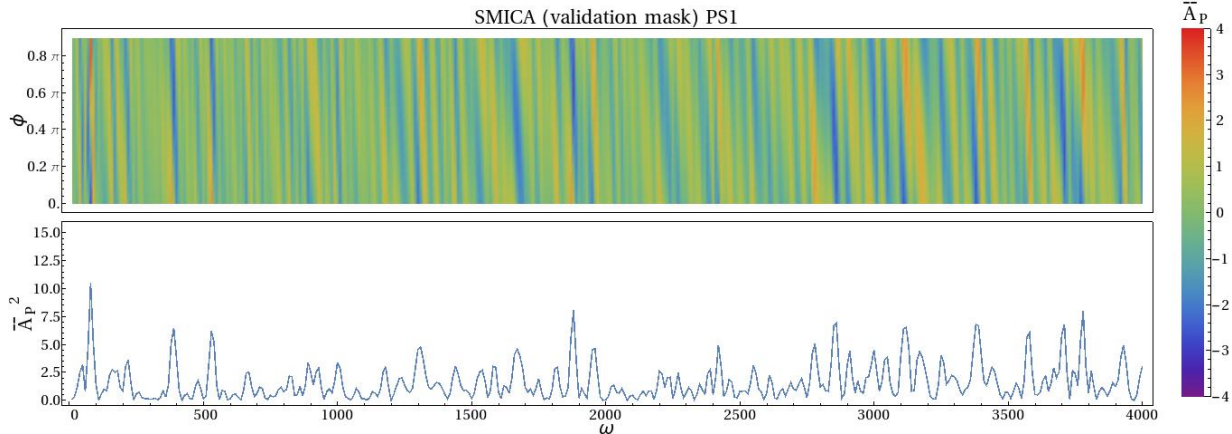


FIG. 8. Amplitudes \bar{A}_P for the template PS1 up to $\omega = 4000$ based on the SMICA map masked with an apodised version of the SMICA validation mask with $f_{\text{sky}} = 0.81$. The bottom panel shows the maximum \bar{A}_P^2 at a given ω , corresponding to the maximum likelihood improvement at that frequency.

that have mostly a decaying sine component ($\phi \sim 0$ or $\phi \sim \pi$) might in fact receive a non-negligible contribution to their S/N that comes from wavenumbers k that violate $\omega k \gg 1$. The SMICA analysis above discards multipoles with $l < 50$ so that for $\omega \gg 100$ this region is not included in the estimates. To achieve the same in the Planck analysis we discard the low- l likelihood and only use the Planck high- l likelihood for the PS2 survey. Then, just like in the SMICA analysis, only multipoles with $l \geq 50$ are taken into account. For those models of the PS2 template that receive significant S/N contributions from $l < 50$ (i.e. those with $\phi \sim 0$ or $\phi \sim \pi$) this is evidently a suboptimal solution and an analysis that uses templates that capture the exact low- l behaviour would produce better constraints on these models.

A further closely related limitation comes from the fact that we assume that the primordial power spectrum is linear in the feature model amplitude. While this should be a good approximation for amplitudes up to $\mathcal{O}(0.1)$, it clearly breaks down for amplitudes that would imply order unity modulations. In extreme cases these would lead to negative values of the primordial power spectrum. The correct templates in these cases have to be obtained from an appropriate non-linear GSR approximation [4]. Again, this problem never occurs when looking for the template PS1 as can be easily seen from the standard deviations on the amplitude \hat{A}_P shown in Fig. 7. For these models the modulations of the primordial power spectrum never become of order unity²¹. However, for the decaying sine component of the template PS2 the sharp rise of the modulation amplitude towards small k can

be problematic. The only models that are significantly affected are those with $\phi \sim 0$ or $\phi \sim \pi$ that contain a large decaying sine contribution. For large frequencies $\omega \gg 2000$ a highly significant result exceeding four sigma would imply order unity modulations at small k for these models. Such results should be interpreted with care as non-linear corrections to the template might have non-negligible effects on the results.

Summing up, when studying the figures in this section care should be taken when interpreting the results near $\phi = 0$ or $\phi = \pi$ as the results might not accurately reflect the signature of a corresponding sharp feature or at least can be significantly improved by an analysis that takes into account the correct small- k behaviour.

Figure 9 shows the results for the template PS2 for both the Planck Likelihood and the SMICA map masked with an apodised version of the U73 mask. We see that the large signal in the SMICA analysis at $\omega \sim 1880$ slightly benefits from the inclusion of the decaying sine and is now observed at 3.7 sigma, whereas the significance remains basically unchanged in the Planck Likelihood. The best fit is still at $\phi \sim 0.7\pi$ as in Fig. 4, so is dominated by the unsuppressed cosine and should be well in the parameter region where our analysis is valid.

On the other hand, the peak at $\omega \sim 3710$ shifts slightly down in significance and is seen at 3.3 sigma in both searches. This could imply that an interpretation in terms of a sharp feature signal is potentially problematic. The largest peak ($\omega \sim 3770$) in the Planck Likelihood analysis is now detected at 3.8 sigma corresponding to a likelihood improvement of $2\Delta \log \mathcal{L} \approx 14$. However, we emphasise that it receives a considerable contribution from the decaying sine as it is seen at $\phi \sim 0.8\pi$, so an interpretation in terms of a sharp feature might require further investigation. Finally, the significance of the peak at $\omega \sim 2840$ in the Planck Likelihood remains unchanged at 3.2 sigma.

²¹ Unless we detect a signal at a significance of $\gtrsim 20$ sigma. If such a strong signature was detected, it should be studied separately using exact predictions for the modulations induced by corresponding feature models. Unfortunately, we never observe a signal that is even remotely close to this level of significance.

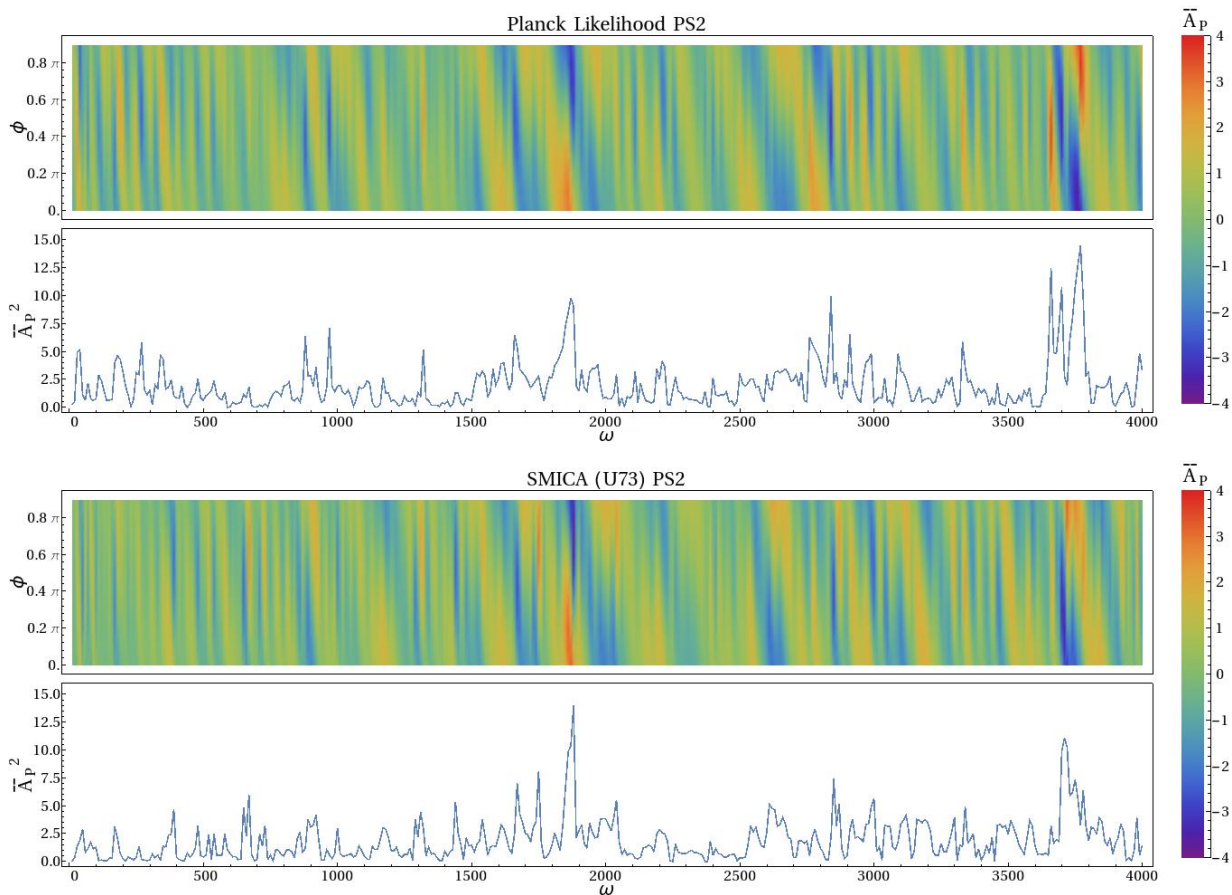


FIG. 9. Amplitudes \bar{A}_P for the template PS2 (decaying sine) up to $\omega = 4000$ based on the Planck Likelihood (top) and the SMICA map masked with an apodised version of the U73 mask with $f_{\text{sky}} = 0.61$ (bottom). In each plot the bottom panel shows the maximum \bar{A}_P^2 at a given ω , corresponding to the maximum likelihood improvement at that frequency.

Figure 10 displays the best-fit models for the SMICA analysis at $\omega \sim 1880$ and $\omega \sim 3710$, respectively. Also plotted are the corresponding PS1 templates. One can clearly see how the decaying sine alters the low- l behaviour. At high l the best-fit PS1 and PS2 templates are identical up to a shift in phase and amplitude of the modulations.

The statistics to discuss the look-elsewhere-adjusted significance were developed focusing on the template PS1 [42]. There are differences in the correlation structure of the templates PS1 and PS2 as can be seen directly by comparison of Figs. 4 and 9. While the cross section through the density plot at $\phi = 0.5\pi$ must be the same (the templates are identical in this case), the decaying sine is more correlated in frequency which results in broader peak patterns around $\phi = 0$ and $\phi = \pi$ ²². With this caveat in mind it should be clear from Fig. 5

that none of the results will give rise to significant improvements once the look-elsewhere effect is taken into account. In particular, the highest result with $\bar{A}_P \approx 3.8$ at $\omega \sim 3770$ in the Planck Likelihood analysis is roughly at the one sigma level.

Finally, repeating the same analysis for the validation mask as in the previous section does not lead to an expected increase in significance as can be seen in Fig. 11. Both large results observed in the SMICA analysis above (cf. Fig. 9) decrease in significance. In particular, the strongest result seen in Fig. 9 at $\omega \sim 1880$ drops below the three sigma level. Again, we emphasise that the meaningfulness of these results depends on the reliability of the component separation method over the entire region of the sky not excluded by the validation mask.

B. WMAP bispectrum survey

1. Bare sine and cosine: results for the template BS1

The results for the template BS1 extracted from WMAP 9-year data up to $\omega = 1000$ are shown on the

²² Heuristically, the suppression of oscillations with increasing l can be thought of as effectively introducing a lower l_{max} beyond which oscillations are negligible. This implies a larger effective step width $\Delta\omega_{\text{eff}}$ and, hence, broader peak patterns (cf. App. B).

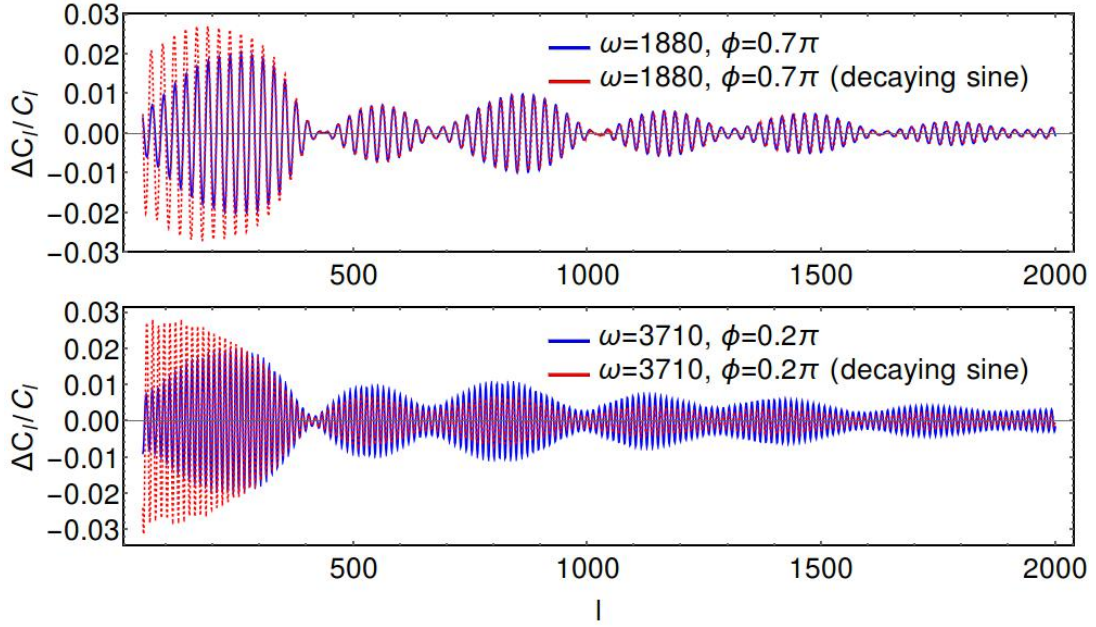


FIG. 10. Best-fit models based on the SMICA analysis for the template PS1 at $\omega = 1880$ with $\phi = 0.7\pi$ and $\omega = 3710$ with $\phi = 0.2\pi$, respectively, with the corresponding fits of the template PS2.

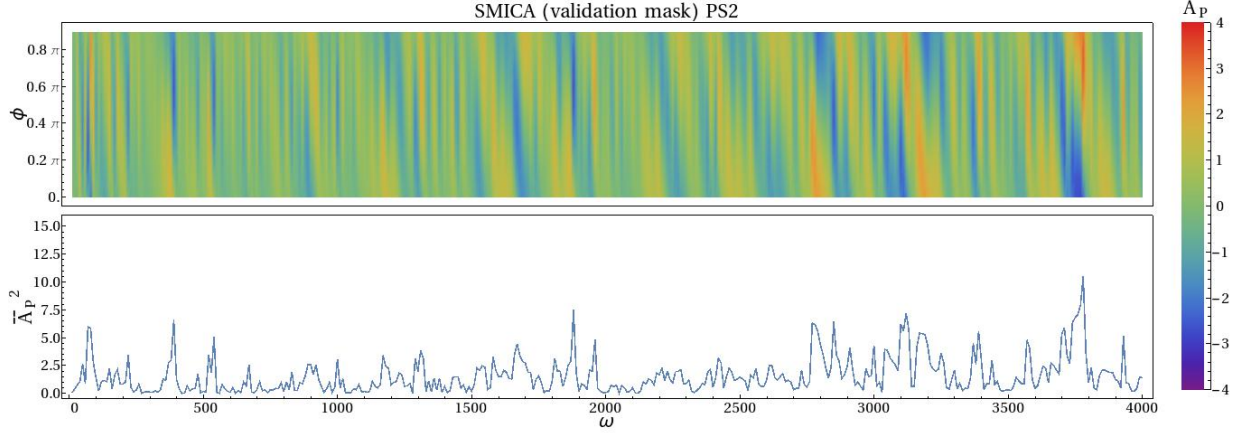


FIG. 11. Amplitudes \bar{A}_P for the template PS2 (decaying sine) up to $\omega = 4000$ based on the SMICA map masked with an apodised version of the validation mask with $f_{\text{sky}} = 0.81$. The bottom panel shows the maximum \bar{A}_P^2 at a given ω , corresponding to the maximum likelihood improvement at that frequency.

left of Fig. 12. As in the corresponding plot for the power spectrum, the upper panel shows a density plot of the normalised amplitudes \bar{A}_B , while the lower panel shows the maximum measured \bar{A}_B^2 at a given frequency. The highest peak is found at $\omega \sim 800$ with a significance of about 2.8 sigma.

To judge whether the observed peaks are at a significant level after the look-elsewhere effect has been taken into account, we can again make use of the statistic Eq. (22) with an appropriate choice of N_{eff} . For an analysis with $l_{\text{max}} = 600$ and WMAP noise level it is shown in

App. B that we have²³ $\Delta\omega_{\text{eff}} \approx 50$. This implies that for a survey covering a frequency range up to $\omega = 1000$ we have $N_{\text{eff}} \approx 20$. The significance of a given \bar{A}_B based on Eq. (22) is presented in Fig. 13 and the value determined from the WMAP survey presented in this work is highlighted. The measured amplitudes are evidently below

²³ Note that the ratio of the values of $\Delta\omega_{\text{eff}}$ for WMAP and Planck, 50 and 13 respectively, is in reasonable agreement with a rough estimate of an effective l_{max} of the two experiments, ~ 600 and ~ 2000 , as one would expect based on how oscillatory templates should be correlated on a given domain (cf. App. B).

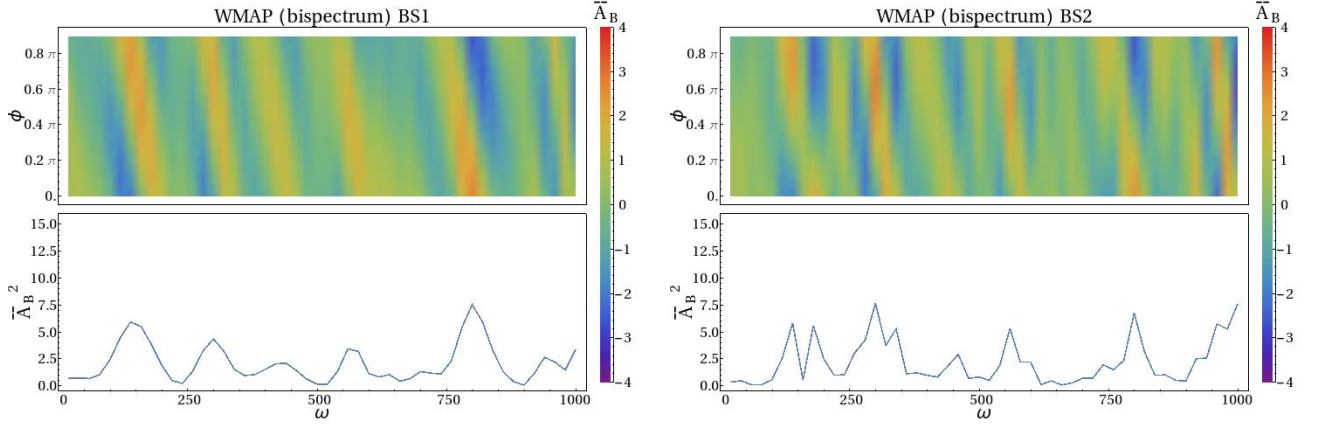


FIG. 12. Amplitudes \bar{A}_B for the template BS1 (left) and BS2 (right) up to $\omega = 1000$ as obtained from the WMAP data. In each plot the bottom panel shows the maximum \bar{A}_B^2 at a given ω .

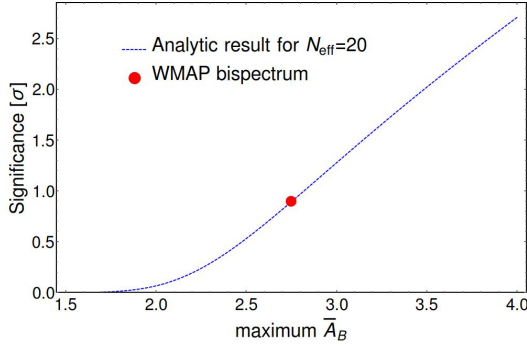


FIG. 13. Analytic model for the distribution of the maximum significance feature model amplitude \bar{A}_B in a survey with $N_{\text{eff}} = 20$. The value obtained from the WMAP data is highlighted.

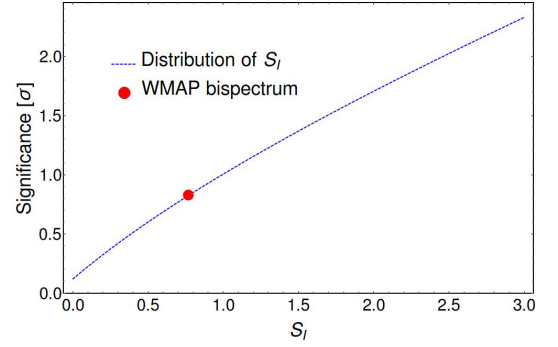


FIG. 14. Distribution of the integrated statistic S_I . The value obtained from the WMAP bispectrum survey assuming $N_{\text{eff}} = 20$ is highlighted.

the one sigma level implying that none of the observed peaks on their own present evidence for an oscillation in the bispectrum.

As in the power spectrum analysis above, we also consider the integrated statistic S_I from Eq. (24) evaluated on the WMAP bispectrum data which is shown in Fig. 14. Just as in the case of the single peak statistic, we arrive at a value below the one sigma level meaning that the abundance of large peaks is entirely consistent with a Gaussian featureless CMB.

2. Including the sharp feature scaling: results for the template BS2

As in the case of the power spectrum, we pointed out in Sec. II, that the bispectrum template BS1 given by Eq. (2) that we used in the previous section does not correctly capture the sharp feature scaling of the shape. The sine component of the signal generically comes with a factor of ωK while the cosine component is multiplied

by a factor of $(\omega K)^2$. We included these scalings in our template BS2, Eq. (4). We go on to investigate whether the inclusion of these scalings changes the results from the previous section.

The results for the template BS2 extracted from WMAP 9-year data up to $\omega = 1000$ are shown on the right of Fig. 12. There are clear similarities to the results for the template BS1 without the correct scaling on the left. This is particularly true for the sine component ($\phi \sim 0$ and $\phi \sim \pi$) showing very correlated patterns in the two plots. However, the cosine component seems to be strongly affected by the inclusion of the $(\omega K)^2$ scaling causing qualitative differences. Generally, the correlation width in frequency seems to have decreased.

None of the peaks have substantially gained in significance due to the inclusion of the sharp feature scaling with the maximum amplitude still being $\bar{A}_B \approx 2.8$. If we use the same statistic as for the BS1 template, it should be clear from Fig. 13 that the results do not exceed the one sigma level. Again with the caveat in mind that the distribution of the maximum amplitudes might be slightly affected due to the different correlation struc-

ture of the template BS2 (cf. Sec. IV A 2), we are led to conclude that no convincing evidence for the presence of sharp features in the WMAP bispectrum alone can be extracted in this frequency range.

C. Combined power spectrum and bispectrum survey

In Ref. [42] a natural statistic was introduced to identify evidence for feature models in a combined survey. We expect feature models to exhibit the same frequency ω in both the power spectrum and the bispectrum (cf. App. A). However, other parameters such as the relative amplitude of the signal in the bispectrum compared to the power spectrum, A_B/A_P , are very model-dependent. Thus, it is sensible to construct the maximum significance joint amplitude estimate \bar{A} at a given ω by maximising over the ratio A_B/A_P (and the phases of the oscillation, ϕ_P and ϕ_B). This results in [42]

$$\bar{A} = (\bar{A}_P^2 + \bar{A}_B^2)^{\frac{1}{2}}, \quad (26)$$

where \bar{A}_P and \bar{A}_B are the individual (normalised) power spectrum and bispectrum amplitudes at that frequency.

An analytic model for the distribution of the maximum joint amplitude estimate found in a combined survey, analogous to Eq. (22) that holds in the case of an individual survey, was presented in Ref. [42]. The essential difference is that in this case the CDF of the χ -distribution with two degrees of freedom is replaced by the corresponding distribution with four degrees of freedom giving

$$S = 2^{\frac{1}{2}} \text{Erf}^{-1} \left[(F_{\chi,4}(\bar{A}))^{N_{\text{eff}}} \right]. \quad (27)$$

Similarly, for a combined survey the integrated statistic S_I , analogous to Eq. (24), is given by

$$S_I^2 = \frac{\Delta\omega}{\Delta\omega_{\text{eff}}} \sum_{\omega} 2 \text{Erf}^{-1} \left[(F_{\chi,4}(\bar{A}_{\omega}))^{N_{\text{eff}}} \right]^2. \quad (28)$$

In both of these definitions an appropriate choice of N_{eff} is required. In the case of combining surveys with identical $\Delta\omega_{\text{eff}}$ as in Sec. IV C 1 the choice is obvious. The choice of N_{eff} for a combined survey with different $\Delta\omega_{\text{eff}}$ is discussed in Sec. IV C 2.

1. WMAP bispectrum and WMAP power spectrum

The results for the template PS1 up to $\omega = 1000$ as extracted from the WMAP likelihood using the efficient methods described in Sec. III A 1 are shown in Fig. 15. They agree very well with the corresponding figure in Ref. [25] as expected.

Before proceeding to perform a combined analysis we would like to point out that the WMAP power spectrum

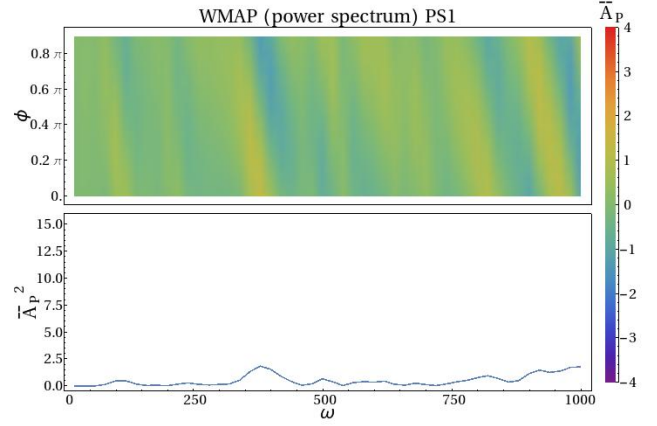


FIG. 15. Amplitudes \bar{A}_P for the template PS1 up to $\omega = 1000$ based on the WMAP likelihood. As in previous figures, the bottom panel shows the maximum \bar{A}_P^2 at a given ω , corresponding to the maximum likelihood improvement at that frequency.

results below $\omega = 1000$ are curiously low with a maximum of $\bar{A}_P \approx 1.4$. We remind the reader that we expect $N_{\text{eff}} = 20$ for WMAP and the given frequency range as discussed above. Simply employing the analytic model for the distribution of the maximum amplitude from Ref. [42] we find that such a low maximum should only occur roughly once in 10 000 realisations which would make the absence of peaks in this region a four sigma anomaly. The analytic distribution is not entirely accurate for judging very low-significance results as it was not designed for this purpose. Closer inspection using MC sampling shows that this anomaly is likely around the three sigma level. At the present stage it is unclear whether this is simply due to an unlikely realisation of the low-frequency scatter around Λ CDM or due to a step in the data processing that systematically eliminates low-frequency oscillations.

The values of the squared combined amplitude estimates \bar{A}^2 obtained from combining the WMAP bispectrum survey with the WMAP power spectrum up to $\omega = 1000$ are shown in Fig. 16. Due to the absence of large peaks in the WMAP power spectrum in this region most of the contributions to \bar{A}^2 come from the bispectrum. Furthermore, the small peaks that can be seen in the power spectrum do not match those in the bispectrum. Hence, we do not expect this combined survey to present us with more significant evidence.

The statistics discussed above confirm this. The result from the statistic for the maximum joint amplitude estimate, Eq. (27), is shown in Fig. 17. It is well below the one sigma level and noticeably lower than the corresponding result for the bispectrum only analysis, Fig. 13. The result from evaluating the integrated statistic, Eq. (28), is shown in Fig. 18. Again, the result is low and dropped compared to its value for the bispectrum only analysis in Fig. 14. The decrease in significance as

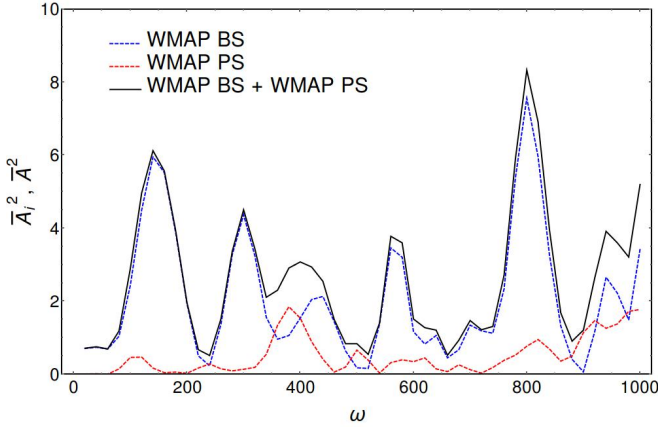


FIG. 16. Individual amplitudes \bar{A}_P^2 and \bar{A}_B^2 and combined amplitude estimates \bar{A}^2 at a given ω obtained from combining the WMAP bispectrum survey with the WMAP power spectrum survey up to $\omega = 1000$.

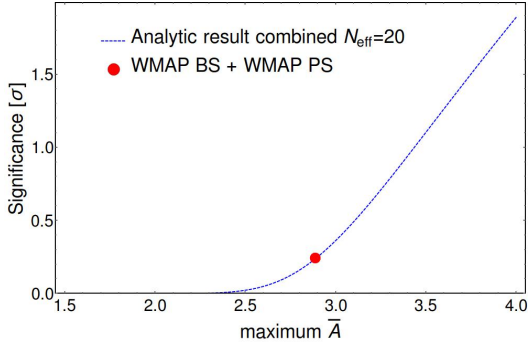


FIG. 17. Analytic model for the distribution of the maximum amplitude estimate \bar{A} in a combined survey with $N_{\text{eff}} = 20$. The value for the combined WMAP bispectrum and WMAP power spectrum analysis is highlighted.

judged by these statistics is due to the absence of large results in the power spectrum in this region and the mismatch of peaks in the power spectrum and bispectrum. The larger look-elsewhere effect that arises in a combined search with more parameters is in this case not matched by corresponding larger observed significances \bar{A} .

2. WMAP bispectrum and Planck/SMICA power spectrum

The values of \bar{A}^2 obtained from combining the WMAP bispectrum survey with the Planck Likelihood up to $\omega = 1000$ are shown on the left of Fig. 19. The corresponding results for the SMICA power spectrum survey are shown on the right. As \bar{A}^2 is simply the sum of the contributions from the individual surveys, it inherits the highly irregular shape with many local maxima.

Again, we would like to decide whether or not any of these joint amplitude estimates present significant evidence. In the present case, there is a further complica-

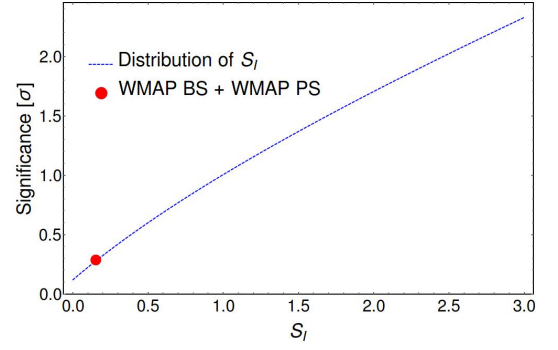


FIG. 18. Distribution of the integrated statistic S_I . The values obtained from the combined WMAP bispectrum and WMAP power spectrum survey assuming $N_{\text{eff}} = 20$ is highlighted.

tion related to the fact that we are combining a WMAP bispectrum survey with Planck power spectrum surveys. These two types of surveys have different $\Delta\omega_{\text{eff}}$ due to their different noise levels and l_{max} . Hence, it is not immediately clear which value to plug into Eq. (23) to extract the correct value of N_{eff} . In App. B it is shown that Eq. (27) is an excellent model for the distribution of the maximum \bar{A} if N_{eff} is taken to be the arithmetic mean of the values of N_{eff} for the individual surveys. Equivalently, the $\Delta\omega_{\text{eff}}$ of the combined survey is the harmonic mean of the values of $\Delta\omega_{\text{eff}}$ of the individual surveys. This gives $\Delta\omega_{\text{eff}} \approx 21$ for the combined survey. For a frequency range up to $\omega = 1000$ this corresponds to $N_{\text{eff}} \approx 48$.

The corresponding distribution with the values for the two combined surveys highlighted is shown in Fig. 20. Both of these combined surveys give a result well below the one sigma level. Note that this result is a combination of the fact that the individual surveys show no significant peaks below $\omega = 1000$ as is evident from Figs. 4 and 12 and the fact that the largest peaks in this region occur at different frequencies and, hence, do not enhance each other. More precisely, there are peaks with $\bar{A}_P^2, \bar{A}_B^2 \approx 8$ in both the power spectrum and the bispectrum survey. If these were located at the same ω , they would produce a joint estimate approaching the two sigma level.

Using the appropriate values for the combined survey, $\Delta\omega_{\text{eff}} \approx 21$ and $N_{\text{eff}} \approx 48$, we arrive at the results presented in Fig. 21. The integrated statistic S_I produces low values in both cases.

V. SUMMARY AND CONCLUSIONS

In this paper, we undertook a thorough search for signatures of sharp features in Planck and WMAP9 data. For the first time, we carried out searches in both the power spectrum and bispectrum simultaneously, employing well-defined look-elsewhere statistics to assess significances in a rigorous manner.

We developed highly efficient methods to scan the

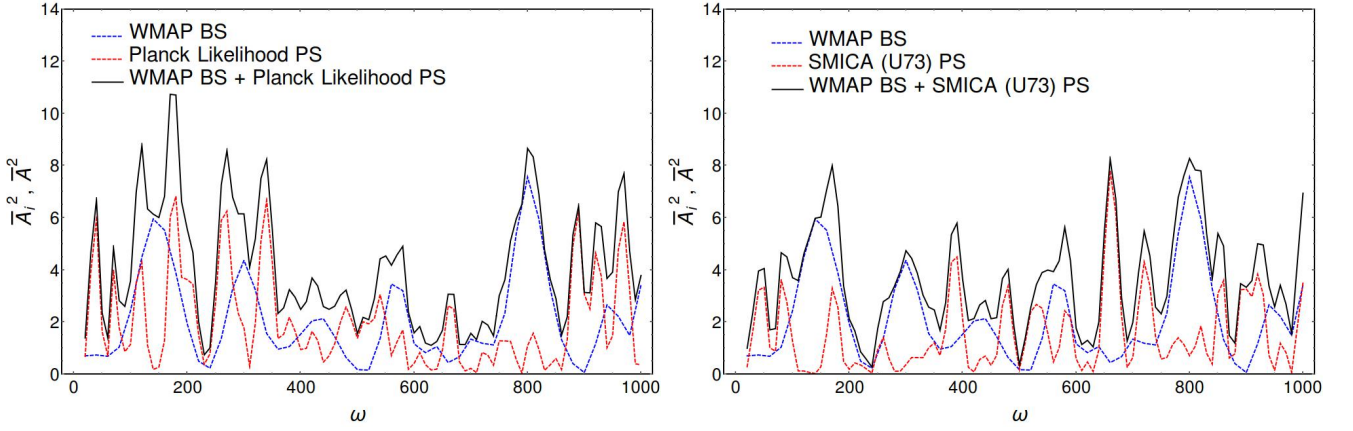


FIG. 19. Individual amplitudes \bar{A}_P^2 and \bar{A}_B^2 and combined amplitude estimates \bar{A}^2 at a given ω obtained from combining the WMAP bispectrum survey with either the Planck Likelihood (left) or the SMICA power spectrum survey (right) up to $\omega = 1000$.

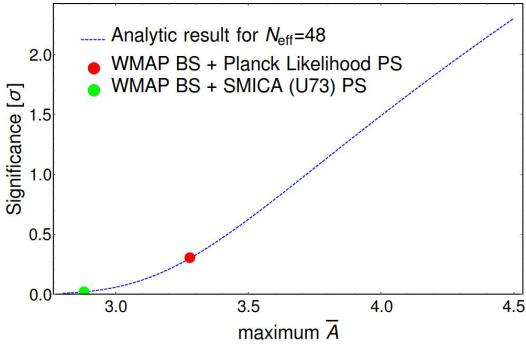


FIG. 20. Analytic model for the distribution of the maximum amplitude estimate \bar{A} in a combined survey with $N_{\text{eff}} = 48$. The values for the two combined surveys considered in this work are highlighted.

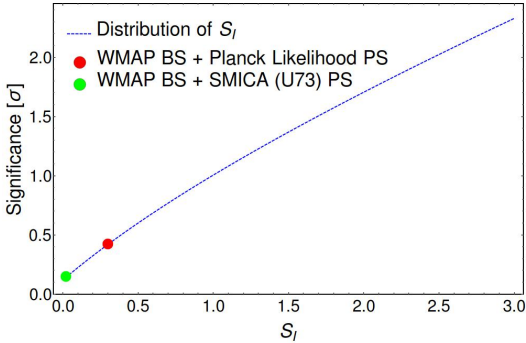


FIG. 21. Distribution of the integrated statistic S_I . The values obtained from the two combined surveys considered in this work assuming $N_{\text{eff}} = 48$ are highlighted.

power spectrum for sharp oscillatory features with frequencies larger than the comoving sound horizon at LS, $\omega \gg 140$. Demonstrating that in this case the only degeneracies of feature degrees of freedom with cosmological parameters are realised through uninteresting adjust-

ments of the comoving distance to LS η_* , we argue that it is entirely sufficient to introduce a grid in frequency, keep the cosmological parameters fixed and only vary the feature model amplitude. In particular, varying η_* does not produce bigger maximum likelihood improvements so that there is no risk of missing possibly interesting results adopting this simplified procedure. The only effect of a change is that frequencies close to a given peak in the likelihood improvement can benefit by an appropriate adjustment that changes the resulting effective oscillation in C_l to match the one observed in the data.

Given that power spectrum likelihoods are very nearly Gaussian in the amplitude of feature models, the use of time consuming MCMC runs is not required and the best-fit amplitude and a corresponding significance can be extracted by fitting a Gaussian. We validated against a full MCMC analysis and found excellent agreement.

We employed these methods to scan the Planck Likelihood and also a likelihood based on the SMICA component separation maps for the signatures of sharp features. The latter has the advantage that, working on the assumption that the foreground cleaning is reliable, a larger sky fraction can be included providing in principle the strongest constraints on oscillatory features in the temperature power spectrum to date.

For our search we used both the phenomenological bare sine and cosine modulations (template PS1) and also included the correct sharp feature scaling (template PS2) up to $\omega = 4000$. The Planck Likelihood scans in the case of the template PS1 agree very well with the corresponding results in Ref. [26] further validating our method. The SMICA map produces consistent results with comparably large likelihood improvements in the same places. To investigate further, we also used the SMICA validation mask for the analysis including a substantially larger sky fraction and, therefore, lowering the error bars by nearly a factor of two. Rather than gaining in significance the large peaks disappear. If we trust the cleaning procedure in regions of higher contamination, this should

be interpreted as evidence that the large peaks are the result of fitting the scatter in the power spectrum estimates by chance and not a true signal.

The latter point of view is supported by the statistics developed in Ref. [42] to judge the look-elsewhere-adjusted significances of findings in feature model surveys. According to these statistics, neither the maximum significance nor the abundance of large peaks in the survey are in any way surprising with significances below the one sigma level. This implies that our realisation of the CMB is entirely consistent with a featureless primordial power spectrum. The inclusion of the sharp feature scaling has an effect on the results, but does not lead to very significant gains.

We went on to study the WMAP bispectrum up to frequencies $\omega = 1000$. While various peaks can be identified, neither the phenomenological bare sine and cosine modulations nor the inclusion of the correct sharp feature scaling in the bispectrum give rise to significant results after look-elsewhere adjustment. To perform a combined search in the slightly simpler case of two surveys with the same effective frequency step width $\Delta\omega_{\text{eff}}$, we first used WMAP power spectrum data. In the process, we noticed that there is a curious absence of peaks in the power spectrum likelihood improvement over this frequency range constituting a roughly three sigma anomaly. With this in mind, it is not surprising that this joint analysis does not lead to any evidence for features.

Slightly generalising the statistics to allow for different $\Delta\omega_{\text{eff}}$, we combined the WMAP bispectrum data with Planck Likelihood and SMICA power spectrum surveys. Again, none of the results are at a statistically significant level. This is a combination of the fact that the individual surveys show no highly significant results, but also due to the fact that peaks do not occur at matching frequencies.

Summing up, neither the inclusion of more sky fraction in the framework of the SMICA analysis nor the inclusion of the correct sharp feature scaling produce power spectrum likelihood improvements that are significant after accounting for the look-elsewhere effect and can thus be interpreted as evidence for features. Invoking the WMAP bispectrum up to $\omega = 1000$ in a combined survey does not change this conclusion. The results are all perfectly consistent with fitting the scatter of the power spectrum estimates assuming a featureless Gaussian CMB.

It will be interesting to see how the upcoming Planck polarisation data will change these results. Oscillatory features should be more prominent in polarisation due to less severe suppression by the transfer functions so that we can expect substantially lower error bars. A scan of the Planck bispectrum should also be available in due course, allowing a combined analysis to be carried out over the entire frequency range and providing further discovery potential.

We expect that similar methods to those presented in this work can be employed to search for other feature model templates. In particular, the sharp feature limit modulations could be generalised to allow for envelopes

in the power spectrum and bispectrum. Furthermore, while we focused on the bispectrum shape that arises from features in the slow-roll parameter ϵ , it is also of some interest to look for the characteristic shape generated by features in the speed of sound.

ACKNOWLEDGMENTS

We would especially like to thank Michele Liguori who was instrumental in developing and validating the modal pipeline which was used to produce the bispectrum results presented in this paper. We would also like to thank Daniel Baumann, Xingang Chen and Yi Wang for valuable discussions and comments. We are grateful to Juha Jäykkä and James Briggs for outstanding computational support. This work is partly based on observations obtained with Planck (<http://www.esa.int/Planck>), an ESA science mission with instruments and contributions directly funded by ESA Member States, NASA, and Canada. HFG and BW gratefully acknowledge the support of the Studienstiftung des deutschen Volkes and STFC studentships. BW also acknowledges support from a Starting Grant of the European Research Council (ERC STG grant 279617) and from a Cambridge European Scholarship of the Cambridge Trust. This work was supported by an STFC consolidated grant ST/L000636/1. It was undertaken on the COSMOS Shared Memory system at DAMTP, University of Cambridge operated on behalf of the STFC DiRAC HPC Facility. This equipment is funded by BIS National E-infrastructure capital grant ST/J005673/1 and STFC grants ST/H008586/1, ST/K00333X/1. We acknowledge use of the HEALPix package [69].

Appendix A: Sharp features in single field inflation

The aim of this section is to briefly motivate the templates we use to search for sharp feature signatures in the power spectrum and bispectrum. More detailed studies of the effects of sharp features can be found elsewhere (cf. Refs. [2, 4, 7, 8] and references therein). We focus on single field inflation with action

$$S = \int d^4x \sqrt{-g} \left(\frac{R}{2} + P(X, \phi) \right), \quad (\text{A1})$$

where P is an arbitrary function of $X = -1/2(\partial_\mu \phi)^2$ and ϕ , R is the Ricci scalar and we set the reduced Planck mass $M_{\text{Pl}} = 1$ for convenience.

In all of this appendix we assume that the sharp features can be dealt with perturbatively. As already discussed at the end of Sec. II this places an upper bound on the sharpness of features that can be studied [46, 47]. The finite width has the effect of introducing an envelope that exponentially suppresses the modulations in the power spectrum for wavenumbers that were deep inside

the horizon at the time of the feature. The treatment in this appendix does not take the effects of a finite width into account and, thus, the envelopes are absent. We argued in Sec. II that this is a reasonable idealisation when looking for the signatures of very sharp, but still perturbative features over the multipole ranges considered in this work.

1. Power spectrum

From Eq. (A1) one can deduce the quadratic part of the action for the scalar curvature perturbation ζ

$$S_2 = \int d^4x \left(\frac{a^3 \epsilon}{c_s^2} \dot{\zeta}^2 - a \epsilon (\partial_i \zeta)^2 \right) \quad (\text{A2})$$

$$= \frac{1}{2} \int d^3x ds \left((v')^2 - (\partial_i v)^2 + \frac{z''}{z} v^2 \right), \quad (\text{A3})$$

where we introduced the variable s following Ref. [70] with $ds = -c_s d\tau = -c_s/a dt$ and defined $v = z\zeta$ with

$$z^2 = \frac{2\epsilon a^2}{c_s}. \quad (\text{A4})$$

Intuitively, the variable s measures the comoving distance sound can travel until the end of inflation at $\tau = 0$. Here and in what follows $'$ denotes derivatives with respect to s . Varying the action we obtain the Mukhanov-Sasaki equation of motion for the Fourier modes v_k :

$$v_k'' + \left(k^2 - \frac{z''}{z} \right) v_k = 0. \quad (\text{A5})$$

To lowest order in slow-roll we simply have $z''/z \sim 2/s^2$. Nonetheless, this term can become very large if there is a sharp feature present at some s_0 in either of these parameters as it also contains the first and second derivatives of ϵ and c_s .

To study the behaviour of the solution for arbitrary deviations from slow-roll one can make use of the generalised slow roll (GSR) technique. However, since we are mainly interested in sharp features that cause high-frequency oscillations, there is an easy way to get insight into the generic behaviour of the resulting power spectra. Let us assume that $ks_0 \gg 1$ so that the z''/z -term is unimportant except in a vicinity of the sharp feature, where the derivatives can become large. We can then think of the effect of this term on such a mode as²⁴

$$\frac{z''}{z} \sim \frac{2}{s^2} + \frac{A}{s} \delta(s - s_0) + \frac{B}{2} \delta'(s - s_0) \quad (\text{A6})$$

for some real coefficients A and B , where A receives contributions from jump discontinuities in ϵ and c_s or their

derivatives. B incorporates contributions proportional to δ' and is only affected by jump discontinuities in ϵ or c_s , but not their derivatives. With this picture it is easy to deduce the effect of the feature. Choosing Bunch-Davies initial conditions for $s \rightarrow \infty$ we have

$$v(s) = \begin{cases} \frac{1}{\sqrt{2k}} \exp(iks), & s > s_0 \\ \frac{C_1}{\sqrt{2k}} \exp(iks) + \frac{C_2}{\sqrt{2k}} \exp(-iks), & s < s_0 \end{cases} \quad (\text{A7})$$

around s_0 . Here, we made use of the assumption $ks_0 \gg 1$ so that the z''/z term can be ignored except at the location of the feature and the solutions to the Mukhanov-Sasaki equation are plane waves.

To match the solutions at s_0 we need two boundary conditions. In order for v_k'' to be proportional to $\delta'v_k$ we need a jump in v_k itself,

$$v_k|_{s_0^+} - v_k|_{s_0^-} = \frac{B}{2} v_k|_{s_0^+}. \quad (\text{A8})$$

This assumes that the jump is not too big so that it is justified to take $v_k|_{s_0^+}$ on the right-hand side of the equation. The second boundary condition can be obtained from integrating the equation across the step and taking the limit of vanishing integration range. Again assuming that the jump is small enough so that we can safely take $v_k \sim v_k|_{s_0^+}$ we arrive at

$$v_k'|_{s_0^+} - v_k'|_{s_0^-} = \frac{A}{s_0} v_k|_{s_0^+} - \frac{B}{2} v_k'|_{s_0^+}. \quad (\text{A9})$$

These two boundary conditions result in

$$C_1 = 1 + \frac{iA}{2ks_0}, \quad (\text{A10})$$

$$C_2 = -\frac{1}{2} \left(B + \frac{iA}{ks_0} \right) \exp(2iks_0). \quad (\text{A11})$$

Including the correct $s \rightarrow 0$ behaviour of $v(s)$ the power spectrum is given by

$$P_{\mathcal{R}}(k) = \lim_{ks \rightarrow 0} \left| \frac{v(s)}{z} \right|^2 \sim \frac{H^2}{4k^3 \epsilon c_s} \Big|_{ks \ll 1} |C_1 - C_2|^2 \quad (\text{A12})$$

so that we arrive at

$$\begin{aligned} P_{\mathcal{R}}(k) &= P_{\mathcal{R},0}(k) (1 + \Delta P_{\mathcal{R}}) \\ &= P_{\mathcal{R},0}(k) \left(1 - \frac{A}{ks_0} \sin(2ks_0) + B \cos(2ks_0) + \dots \right). \end{aligned} \quad (\text{A13})$$

This simple calculation shows that in general single field inflation, the dominant modulations to the power spectrum due to a sharp feature in the $ks_0 \equiv \omega k \gg 1$ limit are a constant cosine and a sine that is suppressed by a factor $1/(\omega k)$. Here, we identified the feature location s_0 with the frequency ω , that we used to parametrise the oscillatory feature templates in our analysis. Eq. (A13) can be rewritten in the form of Eq. (3) by introducing the overall amplitude A_P and the phase ϕ_P .

²⁴ Cf. Ref. [71] where a similar approach was taken to derive an analytic approximation to the power spectrum modulations in the context of sharp steps in brane inflation.

The result is general in the sense that it applies independent of whether the feature arises due to sharp changes in the slow-roll parameter ϵ or in the speed of sound c_s . This is consistent with the rigorous GSR results in Ref. [5]. For $\omega k \lesssim 1$ the behaviour of the solution is much more complicated and requires a rigorous GSR treatment. The S/N in the CMB at multipoles with low l is poor so that nearly all the S/N for extended oscillations comes from $l > \mathcal{O}(10^2)$ corresponding to $k > \mathcal{O}(10^{-2})$. As we are mainly interested in large frequencies, $\omega \gg 10^2$, this means that most of the S/N generically comes from regions with $\omega k \gg 1$ so that it is justified to scan for these models in the power spectrum assuming this limit.

2. Higher order correlators: the bispectrum

Higher order correlators are extracted using the in-in formalism treating the interaction terms in the higher order actions perturbatively. In the spirit of the discussion above we only provide a brief discussion here that motivates the bispectrum shapes which are studied in this work. Thorough treatments can be found elsewhere [2, 72].

The tree-level bispectrum in the in-in formalism is obtained via

$$\begin{aligned} & \langle \zeta_{\vec{k}_1}(t_e) \zeta_{\vec{k}_2}(t_e) \zeta_{\vec{k}_3}(t_e) \rangle \\ & := (2\pi)^3 \delta(\vec{k}_1 + \vec{k}_2 + \vec{k}_3) B(k_1, k_2, k_3) \\ & = 2\Re \left[-i \int_{-\infty}^{t_e} dt \langle \zeta_{\vec{k}_1}(t_e) \zeta_{\vec{k}_2}(t_e) \zeta_{\vec{k}_3}(t_e) H_I(t) \rangle \right], \quad (\text{A14}) \end{aligned}$$

where t_e denotes cosmic time at the end of inflation and H_I is the interaction Hamiltonian. The leading-order bispectrum arises from the cubic action S_3 . The degeneracy between sharp features in the speed of sound and features in ϵ is broken at the level of the bispectrum. In this work we specialise on features in ϵ . In this case the important term is [2, 45]

$$S_3 \supset \int d\tau d^3x \frac{1}{2} a^2 \epsilon \eta' \zeta^2 \zeta', \quad (\text{A15})$$

where we simply set $c_s = 1$ for simplicity so that $s = -\tau$. The leading-order bispectrum arising from the corresponding interaction Hamiltonian then is

$$\begin{aligned} & B(k_1, k_2, k_3) \\ & = \Re \left[i \left(\prod_i u_{k_i}(\tau_e) \right) \int_{-\infty}^{\tau_e} d\tau \frac{\epsilon}{\tau^2 H^2} \eta' \left(\prod_i u_{k_i}^* \right)' \right], \quad (\text{A16}) \end{aligned}$$

where $\tau_e = \tau(t_e)$ is the conformal time at the end of inflation, the $u_k(\tau)$ are the standard slow-roll mode functions

$$u_k(\tau) = \frac{\tau H}{\sqrt{4\epsilon k}} \left(1 - \frac{i}{k\tau} \right) \exp(-ik\tau) \quad (\text{A17})$$

and we made use of the fact that $a \sim -\tau/H$ during inflation.

As in the previous section we investigate the case where the slow-roll parameters acquire singular behaviour so that η' has a δ and a δ' component,

$$\eta' \sim C\delta(\tau - \tau_0) + D\tau\delta'(\tau - \tau_0). \quad (\text{A18})$$

Here, C and D are again largely arbitrary coefficients related to the jump in the first and second derivative of ϵ . The first term gives a contribution

$$B \sim \frac{1}{(k_1 k_2 k_3)^{\frac{3}{2}}} \Re \left[\frac{\epsilon}{\tau^2 H^2} \left(\prod_i u_{k_i}^* \right)' \Big|_{\tau_0} \right], \quad (\text{A19})$$

where we discarded factors that only affect the amplitude and do not contribute to the scale dependence.

We are interested in the behaviour for $K\tau_0 \equiv (k_1 + k_2 + k_3)\tau_0 \gg 1$ as in the case of the power spectrum. In this case the mode functions can be well approximated as

$$u_k(\tau) \sim \frac{\tau H}{\sqrt{4\epsilon k}} \exp(-ik\tau) \quad (\text{A20})$$

and the leading-order behaviour is obtained by letting the derivative act on the exponential. This results in

$$\begin{aligned} B & \sim \frac{\tau_0 K}{(k_1 k_2 k_3)^2} \Re [i \exp(-ik\tau_0)] \\ & \sim \frac{(\tau_0 K) \sin(\tau_0 K)}{(k_1 k_2 k_3)^2} \quad (\text{A21}) \end{aligned}$$

up to an overall factor independent of k .

The second term in Eq. (A18) containing δ' gives

$$B \sim \frac{1}{(k_1 k_2 k_3)^{\frac{3}{2}}} \Re \left[\left(\frac{\epsilon}{\tau H^2} \left(\prod_i u_{k_i}^* \right)' \right)' \Big|_{\tau_0} \right]. \quad (\text{A22})$$

Again isolating the leading-order behaviour in $\tau_0 K$ by letting the derivatives only act on the exponential in the mode functions, we obtain

$$\begin{aligned} B & \sim \frac{(\tau_0 K)^2}{(k_1 k_2 k_3)^2} \Re [\exp(-ik\tau_0)] \\ & \sim \frac{(\tau_0 K)^2 \cos(\tau_0 K)}{(k_1 k_2 k_3)^2}. \quad (\text{A23}) \end{aligned}$$

Note that the second term also produces terms $\sim (\tau_0 K) \sin(\tau_0 K)$ similar to the first term. We can simply add these terms to the contribution arising from the first term and write the net leading-order bispectrum as

$$\begin{aligned} B(k_1, k_2, k_3) & = \frac{A_B}{(k_1 k_2 k_3)^2} (\cos \phi_B f_B(\omega)(\omega K) \sin(\omega K) \\ & \quad + \sin \phi_B (\omega K)^2 \cos(\omega K)), \quad (\text{A24}) \end{aligned}$$

where we identified the feature location with the frequency $\omega = -\tau_0$ as in the previous section. This shows that a given feature produces modulations with the same ω in both the power spectrum and the bispectrum. Here, $f_B(\omega)$ is a factor introduced to give equal S/N to the two terms and ϕ_B parametrises their relative contribution as in Sec. II.

Appendix B: Statistics

For the statistics that we apply to judge the significance of findings, the knowledge of the effective number of frequencies N_{eff} in a survey or, equivalently, the effective frequency step width $\Delta\omega_{\text{eff}}$, related to N_{eff} via Eq. (23), is required. The purpose of this appendix is to determine $\Delta\omega_{\text{eff}}$ for a WMAP-like and Planck-like survey. For the combined power spectrum and bispectrum survey in Sec. IVC we also need to determine an overall $\Delta\omega_{\text{eff}}$ when combining two surveys that do not have identical $\Delta\omega_{\text{eff}}$. The latter is a straightforward extension of the methods already presented in Ref. [42].

Based on the results in Ref. [42], we generally assume that the effective step width is the same for a power spectrum and bispectrum survey based on the same CMB experiment (i.e. identical noise level and l -range). Hence, to determine $\Delta\omega_{\text{eff}}$ it is sufficient to focus on the computationally simpler power spectrum. Sky coverage does not have any significant effects on the correlations between models of different frequency. Hence, we create 10 000 full sky Gaussian CMB realisations for both the WMAP and Planck scenario. In each case we use a standard concordance Λ CDM power spectrum, multiply by the beam function of the given experiment and add the appropriate noise level. In the WMAP case we use multipoles up to $l_{\text{max}} = 600$ while in the Planck case we set $l_{\text{max}} = 2000$, which is the l_{max} used in our SMICA analysis²⁵. We then extract feature model amplitudes employing the fast quadratic estimator discussed in Sec. III A 2 covering frequencies up to $\omega = 4000$.

Throughout this section we focus on the template PS1, i.e. the bare sine and cosine modulations. The corresponding distribution of the maximum significance \bar{A} found in the mock surveys is shown in Fig. 22.

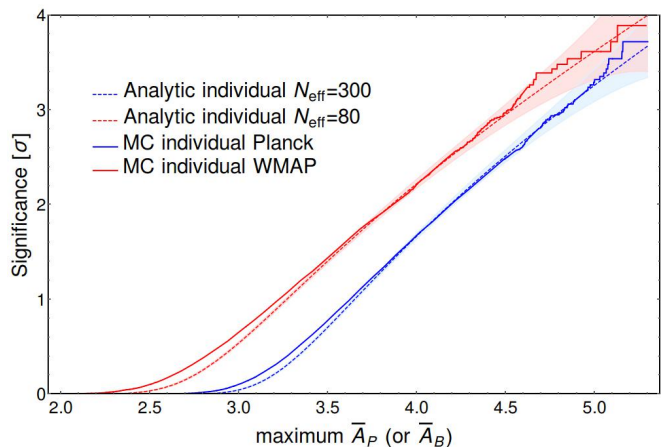


FIG. 22. Distribution of the maximum amplitude estimate \bar{A}_P (or \bar{A}_B) in a power spectrum (or bispectrum) only survey covering frequencies up to $\omega = 4000$ and using the templates PS1 (or BS1). We plot MC results for a WMAP-like and Planck-like survey together with the analytic models of the distributions with $N_{\text{eff}} = 300$ and $N_{\text{eff}} = 80$ respectively.

²⁵ The Planck Likelihood includes multipoles up to $l = 2500$ in some frequency bands, but we do not expect such high multipoles to contribute to the feature S/N due to noise and foregrounds that swamp the already heavily suppressed (due to transfer functions and lensing) signal. Hence, choosing $l_{\text{max}} = 2000$ to extract N_{eff} should be reliable.

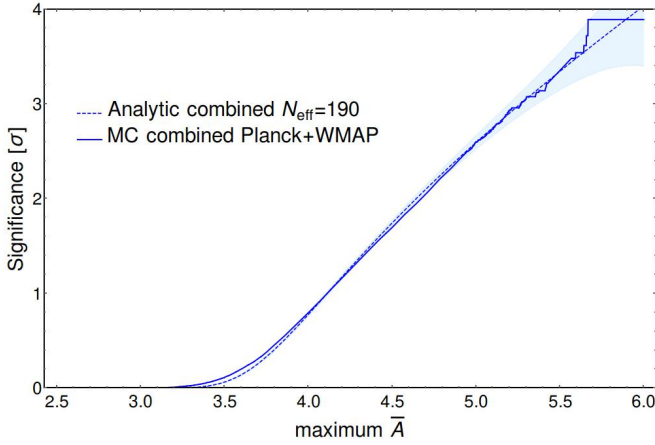


FIG. 23. Distribution of the maximum amplitude estimate \bar{A} in a combined survey up to $\omega = 4000$ where one survey is based on WMAP and the other on Planck data. We plot MC results together with an analytic model of the distribution of \bar{A} in a combined survey with $N_{\text{eff}} = 190$. Note that this choice of N_{eff} is the mean of the corresponding values of the individual surveys.

We also plot the analytic models of the distribution according to Eq. (22) for appropriate choices of the parameter N_{eff} . We observe very good agreement between the analytic models and the MC results as already reported in Ref. [42]. For Planck, the extracted value of $N_{\text{eff}} = 300$ for this survey corresponds to $\Delta\omega_{\text{eff}} \approx 4000/300 \approx 13$ while for WMAP $N_{\text{eff}} = 80$ gives $\Delta\omega_{\text{eff}} \approx 50$. Note that these values are in good agreement with the gen-

eral expectation that the correlation width in frequency of oscillatory modes should decay as $1/l_{\text{max}}$.

To address the question of the appropriate choice of N_{eff} for a combined survey we remind the reader that the statistic we are using to detect evidence for feature models is given by the maximum significance joint amplitude estimate $\bar{A} = (\bar{A}_P^2 + \bar{A}_B^2)^{1/2}$ obtained for any relative amplitude at a given ω . Just as in the case of the individual surveys we can extract an MC estimate of the distribution of the maximum \bar{A} in a combined survey by using the 10 000 WMAP-like and Planck-like realisations described above. Again, we scan for feature models up to $\omega = 4000$. The corresponding results are shown in Fig. 23. We also plot the analytic model for the distribution of \bar{A} in a combined survey, Eq. (27), for the appropriate choice of N_{eff} . First of all, there is again very good agreement between MC simulations and the analytic model, showing that the model is also valid for combining surveys with different $\Delta\omega_{\text{eff}}$. This is a slight generalisation of the results presented in Ref. [42]. Furthermore, it is evident that the overall N_{eff} for the combined survey can be taken to be the arithmetic mean of the values of N_{eff} of the two individual surveys. In particular, in the present case we obtain $N_{\text{eff}} = (300 + 80)/2 = 190$ for the combined survey. This means that the overall $\Delta\omega_{\text{eff}}$ for a combined survey is obtained as the harmonic mean of the effective step widths of the individual surveys $\Delta\omega_{\text{eff},1}$ and $\Delta\omega_{\text{eff},2}$, i.e.

$$\Delta\omega_{\text{eff}} = \frac{2}{\Delta\omega_{\text{eff},1}^{-1} + \Delta\omega_{\text{eff},2}^{-1}}. \quad (\text{B1})$$

-
- [1] P. A. R. Ade *et al.* (Planck Collaboration), *Astron. Astrophys.* **571**, A16 (2014), arXiv:1303.5076 [astro-ph.CO].
 - [2] X. Chen, *Adv. Astron.* **2010**, 638979 (2010), arXiv:1002.1416 [astro-ph.CO].
 - [3] X. Chen, R. Easther, and E. A. Lim, *J. Cosmol. Astropart. Phys.* **0706**, 023 (2007), arXiv:astro-ph/0611645 [astro-ph].
 - [4] C. Dvorkin and W. Hu, *Phys. Rev. D* **81**, 023518 (2010), arXiv:0910.2237 [astro-ph.CO].
 - [5] V. Miranda, W. Hu, and P. Adshead, *Phys. Rev. D* **86**, 063529 (2012), arXiv:1207.2186 [astro-ph.CO].
 - [6] A. Achúcarro, V. Atal, B. Hu, P. Ortiz, and J. Torrado, *Phys. Rev. D* **90**, 023511 (2014), arXiv:1404.7522 [astro-ph.CO].
 - [7] N. Bartolo, D. Cannone, and S. Matarrese, *J. Cosmol. Astropart. Phys.* **10**, 038 (2013), arXiv:1307.3483 [astro-ph.CO].
 - [8] P. Adshead, C. Dvorkin, W. Hu, and E. A. Lim, *Phys. Rev. D* **85**, 023531 (2012), arXiv:1110.3050 [astro-ph.CO].
 - [9] X. Chen, *J. Cosmol. Astropart. Phys.* **1012**, 003 (2010), arXiv:1008.2485 [hep-th].
 - [10] X. Chen, R. Easther, and E. A. Lim, *J. Cosmol. Astropart. Phys.* **0804**, 010 (2008), arXiv:0801.3295 [astro-ph].
 - [11] R. Flauger, L. McAllister, E. Pajer, A. Westphal, and G. Xu, *J. Cosmol. Astropart. Phys.* **1006**, 009 (2010), arXiv:0907.2916 [hep-th].
 - [12] R. Flauger and E. Pajer, *J. Cosmol. Astropart. Phys.* **1101**, 017 (2011), arXiv:1002.0833 [hep-th].
 - [13] X. Chen and M. H. Namjoo, *Phys. Lett. B* **739**, 285 (2014), arXiv:1404.1536 [astro-ph.CO].
 - [14] X. Chen, M. H. Namjoo, and Y. Wang, *J. Cosmol. Astropart. Phys.* **2**, 027 (2015), arXiv:1411.2349 [astro-ph.CO].
 - [15] F. Arroja, A. E. Romano, and M. Sasaki, *Phys. Rev. D* **84**, 123503 (2011), arXiv:1106.5384 [astro-ph.CO].
 - [16] J. Martin and L. Sriramkumar, *J. Cosmol. Astropart. Phys.* **1201**, 008 (2012), arXiv:1109.5838 [astro-ph.CO].
 - [17] A. Achúcarro, J.-O. Gong, G. A. Palma, and S. P. Patil, *Phys. Rev. D* **87**, 121301 (2013), arXiv:1211.5619 [astro-ph.CO].
 - [18] G. A. Palma, *J. Cosmol. Astropart. Phys.* **4**, 035 (2015), arXiv:1412.5615 [hep-th].
 - [19] J.-O. Gong, K. Schalm, and G. Shiu, *Phys. Rev. D* **89**, 063540 (2014), arXiv:1401.4402 [astro-ph.CO].
 - [20] J. Martin and C. Ringeval, *Phys. Rev. D* **69**, 083515 (2004), arXiv:astro-ph/0310382 [astro-ph].
 - [21] C. Pahud, M. Kamionkowski, and A. R. Liddle, *Phys. Rev. D* **79**, 083503 (2009), arXiv:0807.0322 [astro-ph].

- [22] L. Covi, J. Hamann, A. Melchiorri, A. Slosar, and I. Sorbera, *Phys. Rev. D* **74**, 083509 (2006), arXiv:astro-ph/0606452 [astro-ph].
- [23] P. D. Meerburg, R. A. M. J. Wijers, and J. P. van der Schaar, *Mon. Not. R. Astron. Soc.* **421**, 369 (2011), arXiv:1109.5264 [astro-ph.CO].
- [24] H. Peiris, R. Easther, and R. Flauger, *J. Cosmol. Astropart. Phys.* **1309**, 018 (2013), arXiv:1303.2616 [astro-ph.CO].
- [25] P. D. Meerburg, D. N. Spergel, and B. D. Wandelt, *Phys. Rev. D* **89**, 063536 (2014), arXiv:1308.3704 [astro-ph.CO].
- [26] P. D. Meerburg, D. N. Spergel, and B. D. Wandelt, *Phys. Rev. D* **89**, 063537 (2014), arXiv:1308.3705 [astro-ph.CO].
- [27] P. D. Meerburg, D. N. Spergel, and B. D. Wandelt, (2014), arXiv:1406.0548 [astro-ph.CO].
- [28] A. Achúcarro, V. Atal, P. Ortiz, and J. Torrado, *Phys. Rev. D* **89**, 103006 (2014), arXiv:1311.2552 [astro-ph.CO].
- [29] P. A. R. Ade *et al.* (Planck Collaboration), *Astron. Astrophys.* **571**, A22 (2014), arXiv:1303.5082 [astro-ph.CO].
- [30] M. Benetti, *Phys. Rev. D* **88**, 087302 (2013), arXiv:1308.6406 [astro-ph.CO].
- [31] V. Miranda and W. Hu, *Phys. Rev. D* **89**, 083529 (2014), arXiv:1312.0946 [astro-ph.CO].
- [32] B. Hu and J. Torrado, *Phys. Rev. D* **91**, 064039 (2015), arXiv:1410.4804 [astro-ph.CO].
- [33] J. Hamann, A. Shafieloo, and T. Souradeep, *J. Cosmol. Astropart. Phys.* **1004**, 010 (2010), arXiv:0912.2728 [astro-ph.CO].
- [34] G. Nicholson and C. R. Contaldi, *J. Cosmol. Astropart. Phys.* **0907**, 011 (2009), arXiv:0903.1106 [astro-ph.CO].
- [35] G. Nicholson, C. R. Contaldi, and P. Paykari, *J. Cosmol. Astropart. Phys.* **1001**, 016 (2010), arXiv:0909.5092 [astro-ph.CO].
- [36] D. Tocchini-Valentini, M. Douspis, and J. Silk, *Mon. Not. Roy. Astron. Soc.* **359**, 31 (2005), arXiv:astro-ph/0402583 [astro-ph].
- [37] L. Verde and H. V. Peiris, *J. Cosmol. Astropart. Phys.* **0807**, 009 (2008), arXiv:0802.1219 [astro-ph].
- [38] C. Dvorkin and W. Hu, *Phys. Rev. D* **84**, 063515 (2011), arXiv:1106.4016 [astro-ph.CO].
- [39] J. R. Fergusson, M. Liguori, and E. P. S. Shellard, *J. Cosmol. Astropart. Phys.* **1212**, 032 (2012), arXiv:1006.1642 [astro-ph.CO].
- [40] P. A. R. Ade *et al.* (Planck Collaboration), *Astron. Astrophys.* **571**, A24 (2014), arXiv:1303.5084 [astro-ph.CO].
- [41] P. A. R. Ade *et al.* (Planck Collaboration), *Astron. Astrophys.* **571**, A12 (2014), arXiv:1303.5072 [astro-ph.CO].
- [42] J. R. Fergusson, H. F. Gruetjen, E. P. S. Shellard, and M. Liguori, *Phys. Rev. D* **91**, 023502 (2015), arXiv:1410.5114 [astro-ph.CO].
- [43] U. H. Danielsson, *Phys. Rev. D* **66**, 023511 (2002), arXiv:hep-th/0203198 [hep-th].
- [44] U. H. Danielsson, *J. High Energy Phys.* **0207**, 040 (2002), arXiv:hep-th/0205227 [hep-th].
- [45] P. Adshead, W. Hu, C. Dvorkin, and H. V. Peiris, *Phys. Rev. D* **84**, 043519 (2011), arXiv:1102.3435 [astro-ph.CO].
- [46] P. Adshead and W. Hu, *Phys. Rev. D* **89**, 083531 (2014), arXiv:1402.1677 [astro-ph.CO].
- [47] D. Cannone, N. Bartolo, and S. Matarrese, *Phys. Rev. D* **89**, 127301 (2014), arXiv:1402.2258 [astro-ph.CO].
- [48] G. Efstathiou, *Mon. Not. R. Astron. Soc.* **349**, 603 (2004), arXiv:astro-ph/0307515 [astro-ph].
- [49] S. Hamimeche and A. Lewis, *Phys. Rev. D* **79**, 083012 (2009), arXiv:0902.0674 [astro-ph.CO].
- [50] P. A. R. Ade *et al.* (Planck Collaboration), *Astron. Astrophys.* **571**, A15 (2014), arXiv:1303.5075 [astro-ph.CO].
- [51] A. Lewis, A. Challinor, and A. Lasenby, *Astrophys. J.* **538**, 473 (2000), arXiv:astro-ph/9911177 [astro-ph].
- [52] C. Howlett, A. Lewis, A. Hall, and A. Challinor, *J. Cosmol. Astropart. Phys.* **1204**, 027 (2012), arXiv:1201.3654 [astro-ph.CO].
- [53] A. Lewis and S. Bridle, *Phys. Rev. D* **D66**, 103511 (2002), arXiv:astro-ph/0205436 [astro-ph].
- [54] D. Larson *et al.* (WMAP Collaboration), *Astrophys. J. Suppl. Ser.* **192**, 16 (2011), arXiv:1001.4635 [astro-ph.CO].
- [55] E. Hivon, K. M. Górski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen, *Astrophys. J.* **567**, 2 (2002), arXiv:astro-ph/0105302 [astro-ph].
- [56] H. F. Gruetjen, J. F. Fergusson, M. Liguori, and E. P. S. Shellard, To be submitted.
- [57] H. F. Gruetjen and E. P. S. Shellard, *Phys. Rev. D* **89**, 063008 (2014), arXiv:1212.6945 [astro-ph.CO].
- [58] M. Münchmeyer, F. Bouchet, M. Jackson, and B. Wandelt, *Astron. Astrophys.* **570**, A94 (2014), arXiv:1405.2550 [astro-ph.CO].
- [59] M. Münchmeyer, P. D. Meerburg, and B. D. Wandelt, *Phys. Rev. D* **91**, 043534 (2015), arXiv:1412.3461 [astro-ph.CO].
- [60] P. A. R. Ade *et al.* (Planck Collaboration), *Astron. Astrophys.* **571**, A25 (2014), arXiv:1303.5085 [astro-ph.CO].
- [61] P. A. R. Ade *et al.* (Planck Collaboration), *Astron. Astrophys.* **571**, A23 (2014), arXiv:1303.5083 [astro-ph.CO].
- [62] P. A. R. Ade *et al.* (Planck Collaboration), *Astron. Astrophys.* **571**, A19 (2014), arXiv:1303.5079 [astro-ph.CO].
- [63] J. R. Fergusson and E. P. S. Shellard, *Phys. Rev. D* **76**, 083523 (2007), arXiv:astro-ph/0612713 [astro-ph].
- [64] J. R. Fergusson, M. Liguori, and E. P. S. Shellard, *Phys. Rev. D* **82**, 023502 (2010), arXiv:0912.5516 [astro-ph.CO].
- [65] J. R. Fergusson, *Phys. Rev. D* **90**, 043533 (2014), arXiv:1403.7949 [astro-ph.CO].
- [66] C. L. Bennett *et al.* (WMAP Collaboration), *Astrophys. J. Suppl. Ser.* **208**, 20 (2013), arXiv:1212.5225 [astro-ph.CO].
- [67] M. J. Mortonson, C. Dvorkin, H. V. Peiris, and W. Hu, *Phys. Rev. D* **79**, 103519 (2009), arXiv:0903.4920 [astro-ph.CO].
- [68] V. Miranda, W. Hu, and C. Dvorkin, *Phys. Rev. D* **91**, 063514 (2015), arXiv:1411.5956 [astro-ph.CO].
- [69] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, *Astrophys. J.* **622**, 759 (2005), arXiv:astro-ph/0409513 [astro-ph].
- [70] W. Hu, *Phys. Rev. D* **84**, 027303 (2011), arXiv:1104.4500 [astro-ph.CO].
- [71] R. Bean, X. Chen, G. Hailu, S.-H. H. Tye, and J. Xu, *J. Cosmol. Astropart. Phys.* **0803**, 026 (2008), arXiv:0802.0491 [hep-th].
- [72] J. M. Maldacena, *J. High Energy Phys.* **0305**, 013 (2003), arXiv:astro-ph/0210603 [astro-ph.CO].