

Impact of Galactic dust non-Gaussianity on searches for B -modes from inflation

Irene Abril-Cabezas^{1,2}*, Carlos Hervías-Caimapo^{1,3}, Sebastian von Hausegger^{1,4,5}, Blake D. Sherwin^{1,2} and David Alonso^{1,5}

¹Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge CB3 0WA, UK

²Kavli Institute for Cosmology, University of Cambridge, Cambridge CB3 0HA, UK

³Instituto de Astrofísica and Centro de Astro-Ingeniería, Facultad de Física, Pontificia Universidad Católica de Chile, Av. Vicuña Mackenna 4860, 7820436 Macul, Santiago, Chile

⁴Department of Physics, University of Oxford, Rudolf-Peierls Centre for Theoretical Physics, Parks Road, OX1 3PU Oxford, UK

⁵Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK

Accepted 2023 November 13. Received 2023 November 8; in original form 2023 September 20

ABSTRACT

A key challenge in the search for primordial B -modes is the presence of polarized Galactic foregrounds, especially thermal dust emission. Power-spectrum-based analysis methods generally assume the foregrounds to be Gaussian random fields when constructing a likelihood and computing the covariance matrix. In this paper, we investigate how non-Gaussianity in the dust field instead affects CMB and foreground parameter inference in the context of inflationary B -mode searches, capturing this effect via modifications to the dust power-spectrum covariance matrix. For upcoming experiments such as the Simons Observatory, we find *no* dependence of the tensor-to-scalar ratio uncertainty $\sigma(r)$ on the degree of dust non-Gaussianity or the nature of the dust covariance matrix. We provide an explanation of this result, noting that when frequency decorrelation is negligible, dust in mid-frequency channels is cleaned using high-frequency data in a way that is independent of the spatial statistics of dust. We show that our results hold also for non-zero levels of frequency decorrelation that are compatible with existing data. We find, however, that neglecting the impact of dust non-Gaussianity in the covariance matrix can lead to inaccuracies in goodness-of-fit metrics. Care must thus be taken when using such metrics to test B -mode spectra and models, although we show that any such problems can be mitigated by using only cleaned spectrum combinations when computing goodness-of-fit statistics.

Key words: methods: data analysis – cosmic background radiation – cosmological parameters – diffuse radiation.

1 INTRODUCTION

The cosmic microwave background (CMB) was emitted when the Universe was approximately 380 000 yr old. The temperature and polarization fluctuations present in the CMB have informed our understanding of the origin of the Universe. In particular, in the standard cosmological model these primordial density fluctuations were produced by quantum fluctuations in the inflationary early Universe. The simplest inflation models make a key prediction that has not yet been confirmed: the production of a primordial gravitational wave background (Starobinsky 1979; Abbott & Wise 1984). This background would induce a unique B -mode polarization pattern in the CMB (Kamionkowski, Kosowsky & Stebbins 1997; Seljak 1997), which is an important signature of primordial tensor fluctuations. The amplitude of these fluctuations is parametrized by the tensor-to-scalar ratio r ; the tightest constraint to date sets $r < 0.032$ at 95 per cent confidence (Tristram et al. 2022). Although r could in principle be arbitrarily small, broad classes of well-

motivated theories (e.g. Starobinsky 1979) predict a value for r that would be within the observable range of future CMB experiments, given current constraints on the scalar spectral index (see Abazajian et al. 2016).

In this context, there are and will be multiple CMB experiments pursuing the search for the faint primordial B -mode signal, such as *Planck* (Planck Collaboration X 2020c), the BICEP2 and Keck array telescope (Buder et al. 2014), CLASS (Essinger-Hileman et al. 2014), SPT-3G (Benson et al. 2014), the POLARBEAR-2 and Simons Array experiments (Suzuki et al. 2016), the Simons Observatory (SO; Simons Observatory Collaboration et al. 2019), the CMB-S4 experiment (Abitbol et al. 2017; Barron et al. 2022), and *LiteBIRD* (Hazumi et al. 2019; LiteBIRD Collaboration 2023). In particular, SO will deploy a number of Small Aperture Telescopes (SATs) targeted at measuring large-scale B -modes to constrain r at a level of $\sigma(r = 0) \lesssim 0.003$ (Wolz et al. 2023), using an effective sky fraction $f_{\text{sky}} \sim 0.1$ with a white noise level of around $2 \mu\text{K-arcmin}$.

Unfortunately, the observed B -mode signal is contaminated by Galactic foreground emission. Polarized synchrotron emission sourced by the motion of charged particles in the Galactic magnetic field (GMF) overwhelms the B -mode signal at low ($\nu \lesssim 70 \text{ GHz}$)

* E-mail: ia404@cam.ac.uk

frequencies (see e.g. Page et al. 2007). In addition to this, thermal dust emission from our own Galaxy is dominant at frequencies $\nu \gtrsim 70$ GHz (see e.g. Planck Collaboration X 2016d). Dust grains, which are aligned with the GMF, are heated by interstellar radiation and subsequently re-emit partially polarized light that traces the GMF structure (Draine & Fraisse 2009). Dust is the most important foreground at most frequencies relevant for B -mode searches.

As a consequence, exquisite control over the foreground polarization signal is key to measuring primordial B -modes. Several component separation methods have been developed to disentangle the primordial signal from foreground contamination. They can be broadly classified into three groups: methods working in harmonic space (e.g. Cardoso et al. 2008; Dunkley et al. 2013; BICEP2/Keck Collaboration 2015; BICEP2 Collaboration 2018; Choi et al. 2020; Azzoni et al. 2021; Mangilli et al. 2021; The Keck Array and BICEP2 Collaborations 2021), map-based methods (e.g. Eriksen et al. 2006; Gratton 2008; Dunkley et al. 2009; Efstathiou, Gratton & Paci 2009; Remazeilles, Delabrouille & Cardoso 2011; Stompor, Errard & Poletti 2016; Chluba, Hill & Abitbol 2017; Hensley & Bull 2018; de Belsunce, Gratton & Efstathiou 2022; McBride, Bull & Hensley 2023; Vacher et al. 2023a) and hybrid methods (e.g. Fantaye et al. 2011; Azzoni et al. 2023). They all require B -mode polarization measurements at a set of different frequencies to perform the cleaning.

In this paper, we focus on the impact of Galactic dust non-Gaussianity on C_ℓ -based methods. These methods perform a likelihood analysis of the power-spectrum signal, which includes all BB auto- and cross-power spectra at a set of frequencies. Crucially, they require an estimation of the power-spectrum covariance matrix. At this stage, most works usually assume that all sky components can be described by statistically isotropic Gaussian random fields at the map level. In the full sky, and assuming statistical isotropy, this would give rise to a purely diagonal covariance matrix, directly related to the power spectrum of the different components. Off-diagonal elements coupling different power-spectrum multipoles would then only be caused by masking and filtering of the data (Knox 1997; Efstathiou 2004). In practice, however, Galactic foregrounds are both statistically anisotropic and exhibit rather non-Gaussian features, both in temperature (e.g. Ben-David, von Hausegger & Jackson 2015; Jung, Racine & van Tent 2018) and in polarization (e.g. Coulton & Spergel 2019; von Hausegger, Gammelgaard Ravnebjerg & Liu 2019), which inevitably leads to couplings between different scales. Given the magnetized and turbulent nature of the diffuse interstellar medium (ISM), which gives rise to highly non-linear motions in the ISM fluid, these couplings are indeed expected. At the map-level, they give rise to structures such as filaments. At the power-spectrum level, they introduce couplings between different ℓ -bins in the power-spectrum covariance, which itself is no longer completely determined by just the multicomponent power spectra. Although the problem of foreground non-Gaussianity has been briefly discussed in the literature (e.g. The Keck Array and BICEP2 Collaborations 2021, hereafter BK18), an in-depth study on how dust non-Gaussianities affect the power-spectrum covariance, and the full impact this has on CMB and foreground parameter inference is still outstanding.

Traditionally, component separation studies have made use of simulations based on observed sky templates. Most of these templates rely on polarization measurements that lack sensitivity on small scales. This issue is usually circumvented by filtering out these scales, and filling them in with Gaussian realizations of a given target power spectrum that follows the trend measured on large scales (see e.g. Hervías-Caimapo, Bonaldi & Brown 2016; Thorne et al. 2017). As a result, works using these models do not account correctly for non-

Gaussianity on small scales. In the last few years, several groups have produced foreground simulations including non-Gaussianities on small scales. These include magnetohydrodynamic simulations (e.g. Kritsuk, Flauger & Ustyugov 2018; Kim, Choi & Flauger 2019), phenomenological models based on temperature observations using a superposition of Galactic dust layers (e.g. Vansyngel et al. 2017; Martínez-Solaache, Karakci & Delabrouille 2018), models built from wavelet scattering transform statistics (Allys et al. 2019; Regalado-Saint Blancard et al. 2020), dust models based on neutral hydrogen (Clark & Hensley 2019), and filament-based ones (Huffenberger, Rotti & Collins 2020; Hervías-Caimapo & Huffenberger 2022). Similar studies have also been carried out in the context of Galactic synchrotron (e.g. Wang et al. 2020; Martire, Barreiro & Martínez-González 2022; Martire et al. 2023). BK18 explored some of these non-Gaussian foreground models. They used single dust realizations from Vansyngel et al. (2017), Kritsuk, Flauger & Ustyugov (2018), and Martínez-Solaache, Karakci & Delabrouille (2018), which they added to their modelled signal. However, with the analysis of only a limited number of realizations, it is unclear how reliably and precisely one can quantify the effect that foreground non-Gaussianity has on the C_ℓ covariance matrix and on the inference of r .

The paper is organized as follows. We present the general methodology of C_ℓ -based analyses in Section 2, with a particular focus on how to incorporate dust non-Gaussianity into the covariance matrix (2.4). In Section 3, we quantify the impact that dust non-Gaussianity has on parameter constraints, including in the presence of frequency decorrelation. We study the effect on parameter biases and variances, as well as goodness-of-fit metrics. We summarize our results and conclude in Section 4.

2 FORMALISM

2.1 Parameter inference

We focus on multifrequency power-spectrum-based analyses. Here, the data vector is composed of all BB auto- and cross-power spectra between all frequency maps. We use Monte Carlo Bayesian analysis, as implemented in EMCEE (Foreman-Mackey et al. 2013) to obtain posterior distributions of our sky model parameters (Section 2.2) given the mock data vector we generated (Section 2.3). This analysis pipeline has already been used in works such as Azzoni et al. (2021) and Wolz et al. (2023).

Often, C_ℓ -based analyses constraining primordial B -modes assume a Gaussian likelihood within the Bayesian analysis. This implicitly assumes that each sky component can be summarized by its power spectrum and its covariance (ignoring measurement noise and systematics). While this is certainly the case for the CMB, one also makes these assumptions for the foregrounds because each bandpower is an average over approximately $\Delta\ell \times (2\ell_{\text{eff}} + 1)$ modes, (neglecting correlations due to, e.g. masking), where ℓ_{eff} is the effective ℓ value at each bin of size $\Delta\ell$. Thanks to the high number of modes, the central limit theorem applies and one can assume that each C_ℓ follows a Gaussian distribution. However, this bandpower averaging does not remove the impact of non-Gaussian foregrounds, as the bandpower covariance matrix, with elements $\text{Cov}(C_\ell, C_{\ell'})$, includes contributions from the connected trispectrum of the non-Gaussian foregrounds.¹

¹As mentioned in Section 1, when only a fraction of the sky is used, $f_{\text{sky}} < 1$, off-diagonal elements in the covariance matrix will also be present due to the mask.

We assess the impact of the non-Gaussian dust foreground by comparing two separate runs of the parameter inference. For these we use two different models for the covariance matrix, keeping the data vector and the sky model fixed. In the first case, we assume a Gaussian dust foreground. In the second case, we describe the anisotropic and non-Gaussian nature of the dust sky by modelling the covariance matrix accordingly. We denote these covariances as Σ_G and Σ_{NG} , respectively. In Section 2.4, we describe how they are computed.

We run our analyses with the Hamimeche & Lewis (2008, HL) likelihood. It accounts for the non-Gaussian form of the likelihood on large scales, where there are fewer modes to average over. On the scales we are interested in ($30 \leq \ell \leq 300$), both the Gaussian and HL likelihoods yield the same results (Wolz et al. 2023). We explore the posterior distribution by Markov chain Monte Carlo sampling using 128 walkers and 10^4 steps, and compute the integrated auto-correlation time (Goodman & Weare 2010) to assess convergence and ensure we have a sufficient number of independent samples. We remove the initial 25 per cent of the chains (‘burn-in’) to perform inferences. We present our results in Section 3.

2.2 Sky model

To create the mock data vector in this work, we use the same sky model described in Abitbol et al. (2021), which follows the fiducial BICEP2/Keck model (BK18; BICEP2/Keck Collaboration 2015; BICEP2 Collaboration 2018). We assume three components c for the sky signal: dust (d), synchrotron (s), and CMB. The foregrounds (d, s) autospectra are parametrized as a power law:

$$\frac{\ell(\ell+1)}{2\pi} C_\ell^{cc} = A_c \left(\frac{\ell}{\ell_0} \right)^{\alpha_c}, \quad c = \{d, s\}, \quad (1)$$

where A_c is the amplitude, α_c the tilt, and $\ell_0 = 80$ the pivot frequency. The model also allows for a cross-correlation between dust and synchrotron,

$$C_\ell^{ds} = \epsilon_{ds} \sqrt{C_\ell^{dd} C_\ell^{ss}}, \quad (2)$$

via a scale-independent parameter ϵ_{ds} . The CMB B -mode power spectrum is modelled as

$$C_\ell^{\text{CMB}} = A_{\text{lens}} C_\ell^{\text{lens}} + r C_\ell^{\text{tens}, r=1}, \quad (3)$$

where C_ℓ^{lens} and $C_\ell^{\text{tens}, r=1}$ are theoretical predictions obtained with CAMB (Lewis & Challinor 2011) for the B -mode power spectrum produced by gravitational lensing and primordial tensor perturbations with $r = 1$, respectively. Treating A_{lens} as a free parameter allows us to consider the need to marginalize over uncertainties in the residual lensing B -modes if delensing is present ($A_{\text{lens}} = 1$ recovers the case without delensing), which will be of importance to improve r constraints (see e.g. Seljak & Hirata 2004; Sherwin & Schmittfull 2015; Namikawa et al. 2022).

The multifrequency power spectrum is then given by

$$C_\ell^{vv'} = \sum_{c, c'} S_c^v S_{c'}^{v'} C_\ell^{cc'}, \quad (4)$$

where S_c^v is the spectral energy distribution (SED) of component c at frequency ν in thermodynamic temperature units. We model the SEDs in the same manner as Abitbol et al. (2021):

(i) **CMB.** For the CMB, the SED is a constant by definition:

$$S_\nu^{\text{CMB}} = 1. \quad (5)$$

Table 1. Priors used in the analysis, based on those used in Azzoni et al. (2021). The power-spectrum amplitude for tensor modes r is allowed to be negative to detect possible biases.

Parameter	Prior	Bounds
r	Top-hat	$[-1, 1]$
A_{lens}	Top-hat	$[0, 2]$
A_d	Top-hat	$[0, \infty)$
α_d	Top-hat	$[-1, 0]$
β_d	Gaussian	1.59 ± 0.50
A_s	Top-hat	$[0, \infty)$
α_s	Top-hat	$[-1, 0]$
β_s	Gaussian	-3.0 ± 0.6
ϵ_{ds}	Top-hat	$[-1, 1]$
B_d	Top-hat	$[-10, 10]$
γ_d	Top-hat	$[-6, -2]$
B_s	Top-hat	$[-10, 10]$
γ_s	Top-hat	$[-6, -2]$
Δ_d	Top-hat	$[0.9, 1.1]$

(ii) **Synchrotron.** We model the synchrotron spectrum as a simple power law (see e.g. Planck Collaboration IV 2020b):

$$S_\nu^s = g(x_{\nu_0^s}) \left(\frac{\nu}{\nu_0^s} \right)^{\beta_s}, \quad (6)$$

where β_s is the synchrotron spectral index and $\nu_0^s = 23$ GHz is the pivot frequency. $g(x_\nu)$ is the conversion factor between Rayleigh–Jeans brightness temperature units (commonly used for foregrounds) and thermodynamic temperature units (Planck Collaboration IX 2014):

$$g(x_\nu) = e^x \left(\frac{x}{e^x - 1} \right)^2, \quad x = \frac{h\nu}{k\Theta_{\text{CMB}}}, \quad (7)$$

where $\Theta_{\text{CMB}} = 2.7255\text{K}$ is the CMB temperature (Fixsen 2009).

(iii) **Dust.** Thermal dust emission is characterized by a modified blackbody spectrum (see e.g. Planck Collaboration IV 2020b):

$$S_\nu^d = g(x_{\nu_0^d}) \left(\frac{\nu}{\nu_0^d} \right)^{\beta_d} \frac{B_\nu(\Theta_d)}{B_{\nu_0^d}(\Theta_d)}, \quad (8)$$

where β_d is the dust spectral index, $\nu_0^d = 353$ GHz is the pivot frequency, and $B_\nu(\Theta)$ is the Planck blackbody spectrum, which expresses the spectral radiance of a blackbody for frequency ν at temperature Θ :

$$B_\nu(\Theta) = \frac{2h\nu^3}{c^2} \left[\exp\left(\frac{h\nu}{k\Theta}\right) - 1 \right]^{-1}, \quad (9)$$

and where k and h are the Boltzmann and Planck constants, respectively. The priors for all model parameters presented hereto are based on table 1 of Azzoni et al. (2021), which we reproduce in Table 1 for completeness.

The simple frequency dependence of equation (4) assumes that all components are perfectly correlated across frequencies. However, we expect Galactic foregrounds to present some level of frequency decorrelation, caused by variations in their spectral energy distribution across the sky and along the line of sight (see e.g. Pelgrims et al. 2021; Ritacco et al. 2023). To parametrize this decorrelation we will make use of a simple scale-independent dust decorrelation parameter, Δ_d , quantifying the map-level correlation coefficient between dust at 353 and 217 GHz (Planck Collaboration XXX 2016a). Specifically, it is defined as the ratio of the cross-spectrum between those two maps

and the geometric mean of the corresponding autospectra (see equation 1 in Planck Collaboration L 2017). Even though the physically meaningful range for Δ_d is $\Delta_d \leq 1$ (BICEP2 Collaboration 2018), we allow for Δ_d values greater than 1 to check for possible biases. We will also assess the impact of frequency decorrelation through the minimal moment expansion method of Azzoni et al. (2021), which is based on the ‘moments-based’ expansion (Chluba, Hill & Abitbol 2017) and effectively generalizes of the simple Δ_d parametrization, modelling the spatial fluctuations of the foreground spectral indices, $\delta\beta_c$, as Gaussian random fields with a power spectrum of the form

$$C_\ell^{\beta_c} = B_c \left(\frac{\ell}{\ell_0} \right)^{\gamma_c}, \quad c = \{d, s\}, \quad (10)$$

and associated frequency dependence (Azzoni et al. 2021)

$$\bar{S}_v^c(\delta\beta_c) = \log \left(\frac{v}{v_0^c} \right) S_v^c, \quad c = \{d, s\}, \quad (11)$$

Even though B_s, B_d correspond to autospectra amplitudes (see e.g. Mangilli et al. 2021; Vacher et al. 2022), we do not enforce positivity in their priors. This is to avoid a bias in the marginalized posterior distribution of r , which is found when data with no spectral index variations are analysed with the moments expansion and a positive prior $B_c \geq 0$ (Azzoni et al. 2021). In summary, our baseline sky model has 9 free parameters. The number of parameters is extended to 10 as we include decorrelation in the dust component, or to 13 for the moments expansion in both dust and synchrotron.

2.3 Mock data

We generate mock multifrequency power-spectrum data on which to run our analysis by evaluating the baseline model we described in the Section (2.2) with the following parameter values. For the CMB contribution, we assume no primordial B -modes, $r = 0$, and assume no delensing, $A_{\text{lens}} = 1$, in the first instance. For the foreground power spectra, we use $A_d = 28 \mu\text{K}^2$, $\alpha_d = -0.16$, $A_s = 2.0 \mu\text{K}^2$, and $\alpha_s = -0.60$, and assume no dust-synchrotron correlation ($\epsilon_{ds} = 0$). These values correspond to the best-fitting amplitudes found in Azzoni et al. (2021) from a suite of PYSM (Thorne et al. 2017) foreground simulations within the SO footprint.² For the SED parameters, we use the best-fitting parameters from *Planck*: $\Theta_d = 19.6\text{K}$, $\beta_d = 1.59$ (Planck Collaboration XXII 2015) and $\beta_s = -3.00$ (Dunkley et al. 2009; Planck Collaboration XXV 2016e).

We take the Simons Observatory experiment, as described in Ade et al. (2019), as our baseline instrument model. We generate multifrequency power spectra using the sky model above assuming six SO frequency bandpasses, centred around 27 and 39 GHz (low-frequency bands, LF), 93 and 145 GHz (mid-frequency bands, MF), and 225 and 280 GHz (ultra-high-frequency bands, UHF). This leads to a total of 21 unique multifrequency power spectra $C_\ell^{vv'}$.

To incorporate the contribution from instrumental noise to the covariance matrix, we include the noise autospectrum at each frequency, assuming no correlation of the noise between different frequencies. For this we follow Ade et al. (2019) and model the SO SATs’ noise spectrum as:

$$N_\ell^{vv'} = \delta_{vv'} N_{\text{white}}^v \left[1 + \left(\frac{\ell}{\ell_{\text{knee}}} \right)^{\alpha_{\text{knee}}^v} \right], \quad (12)$$

²See, however, nuances on A_d and α_d in Section 2.4.2.

where the first term accounts for white noise and the second for red $1/f$ noise, most prominent on large scales, and caused by atmospheric contributions and timestream filtering. This noise model is rather simplistic, and does not incorporate features associated with the specific scanning strategy, such as inhomogeneity or anisotropy. Although Wolz et al. (2023) find that noise inhomogeneity may lead to an increase in $\sigma(r)$ of up to $\sim \mathcal{O}(30)$ per cent, this level of precision is sufficient for the purposes of quantifying the impact of foreground non-Gaussianity. We use values of N_{white}^v and α_{knee}^v corresponding to the ‘goal-optimistic’ scenario, specified in tables 1 and 2 of Ade et al. (2019). This corresponds to the most optimistic noise model envisaged by SO.

We apply our results to BB -power-spectrum measurements on the SO patch, which we take to be our baseline analysis. We use the SAT footprint released in Ade et al. (2019) as our sky mask. Its edges are apodized on 5° scales using the ‘C1’ apodization method (Grain, Tristram & Stompor 2009) as implemented in NAMASTER.³ (Alonso et al. 2019), giving an effective sky fraction of $f_{\text{sky}} = 0.1$. We only use data in the range $30 \leq \ell \leq 300$, binned into $\Delta\ell = 30$ (9 bins in total). We choose a somewhat broad ℓ -binning to ensure that the covariance converges given the number of dust filament realizations we have (see Section 2.4.2), but we check that our results are robust against a finer $\Delta\ell = 15$ binning.

2.4 Dust non-Gaussian covariance

As we describe in Section 2.1, we capture the effect of the non-Gaussian dust foreground in the covariance matrix Σ_{NG} . This is then compared to the results obtained assuming a Gaussian dust foreground as per Σ_G . We first sketch how we compute the Gaussian covariance matrix Σ_G . Detailed calculations of all results obtained in this section are described in Appendix A.

We start with a Gaussian field $\delta(\hat{\mathbf{n}})^4$, which is modulated by the sky mask $w(\hat{\mathbf{n}})$:

$$d_G(\hat{\mathbf{n}}) = w(\hat{\mathbf{n}})\delta(\hat{\mathbf{n}}), \quad (13)$$

and define the corresponding pseudo- C_ℓ , \tilde{C}_ℓ^{dG} , for which one has

$$\langle \tilde{C}_\ell^{dG} \rangle = \sum_{\ell'} M_{\ell\ell'}^w C_{\ell'}^\delta, \quad (14)$$

where C_ℓ^δ is the power spectrum of $\delta(\hat{\mathbf{n}})$, and we have defined the C_ℓ -level mode-coupling matrix $M_{\ell\ell'}^w$, which can be computed from the sky mask alone (Hivon et al. 2002; Alonso et al. 2019). From equation (14), one can see that the effect of the sky mask is to couple different ℓ -modes.

We calculate the covariance of the power spectrum of the Gaussian field via its 4-point function. We assume that the sky mask is smooth, the mode-coupling coefficients are peaked around $\ell = \ell'$ and that C_ℓ^δ is slowly varying. The final expression is:

$$\text{Cov}(C_{\ell_1}^\delta, C_{\ell_2}^\delta) = \frac{2(\tilde{C}_{\ell_{12}}^\delta)^2}{\langle w^2 \rangle_\Omega (2\ell_2 + 1)} M_{\ell_1\ell_2}^{w^2}, \quad (15)$$

where $\tilde{C}_{\ell_{12}}^\delta$ denotes the geometric mean of $C_{\ell_1}^\delta$ and $C_{\ell_2}^\delta$, and $\langle \dots \rangle_\Omega$ denotes sky averaging. $M_{\ell_1\ell_2}^{w^2}$ is now the mode-coupling matrix for the sky mask squared, $w^2(\hat{\mathbf{n}})$, which we compute analytically with NAMASTER (Alonso et al. 2019). Under the given approximations, the $\langle w^2 \rangle_\Omega$ term serves as the conversion factor between \tilde{C}_ℓ^{dG} and C_ℓ^δ .

³<https://github.com/LSSTDESC/NaMaster/>

⁴In this work, we only work with the B -mode component, which we treat as a spin-0 scalar field.




Figure 1. Correlation of the dust power spectrum covariance at 280 GHz for different dust models in the SO patch. The ℓ range spans $\ell \in (30, 300)$, binned with $\Delta\ell = 30$, and is the same across all panels. The first panel corresponds to the covariance matrix of our initial Gaussian dust field, Σ_G , which already incorporates masking effects: the covariance is essentially diagonal with the exception of small off-diagonal terms close to the diagonal that appear due to the mask. The second panel corresponds to a Gaussian dust field modulated by a large-scale template (Section 2.4.1). One can see that the large-scale modulation moves power from the diagonal to the off-diagonal terms, thus further coupling neighbouring ℓ -bins. The third panel corresponds to the covariance of a filamentary small-scale dust population obtained from 200 realizations of the DUSTFILAMENTS (Hervías-Caimapo & Huffenberger 2022) model. The net effect of this population is to couple small scales to large scales. The last panel refers to the covariance matrix Σ_{NG} , which contains both large and small scale effects. We use it to capture the full anisotropic non-Gaussian nature of the dust field.

Equation (15) is essentially the Knox formula (Knox 1995, 1997) modified to properly include mask-coupling effects (Efstathiou 2004; García-García, Alonso & Bellini 2019). We show in the first panel of Fig. 1 the correlation matrix for the power-spectrum covariance of a Gaussian dust field at 280 GHz computed in this way. The covariance Σ_G is essentially diagonal with the exception of small coupling terms close to the diagonal that appear due to the mask.

However, as we already discussed, the nature of foregrounds is anisotropic and non-Gaussian, and quantifying the impact of this on parameter inference is the main goal of this paper. Previous analyses (e.g. Azzoni et al. 2023) have assumed Gaussian foregrounds or have only briefly discussed the question of how non-Gaussian foregrounds impact r inference (BK18). Focusing on dust, we compare how the covariance matrix of dust changes as we introduce non-Gaussianity to an initially Gaussian isotropic dust field in two regimes:

(i) **Large-scale modulation.** We introduce non-Gaussianity by modulating the Gaussian dust field by a large-scale template, which we take to be the *Planck* 353 GHz temperature map smoothed on small scales. This model was already proposed in Mak et al. (2017). The net effect of the modulation is to couple neighbouring bins, thus moving power from the diagonal to the nearby off-diagonal elements of the covariance.

(ii) **Small-scale dust filaments.** We use dust simulations based on filaments (DUSTFILAMENTS; Hervías-Caimapo & Huffenberger 2022) to estimate the covariance on small scales. These structures affect all scales and therefore introduce dominant off-diagonal elements away from the diagonal.

2.4.1 Large-scale modulating template

As found in Miville-Deschênes et al. (2007), the small-scale variance of the dust emission scales with the square of the large-scale intensity. This relation motivates a simple model to account for the statistically anisotropic nature of Galactic dust, as proposed in Mak et al. (2017). In this model, we modify equation (13) so that the Gaussian isotropic dust field $\delta(\hat{\mathbf{n}})$ is also modulated by a large-scale intensity field $t(\hat{\mathbf{n}})$:

$$d_{NG}(\hat{\mathbf{n}}) = w(\hat{\mathbf{n}})t(\hat{\mathbf{n}})\delta(\hat{\mathbf{n}}), \quad (16)$$

which enforces the results found empirically. There are multiple ways to estimate the template $t(\hat{\mathbf{n}})$ (Delabrouille et al. 2013; Remazeilles et al. 2015; Hervías-Caimapo, Bonaldi & Brown 2016; Mak et al. 2017; Thorne et al. 2017). In this work, we take as modulating template the 353 GHz temperature map from *Planck* (Planck Collaboration III 2020a), smoothed on 0.4° ($\ell \approx 400$) scales. This scale is chosen to be well above the map noise limit, which we estimate using the FFP10 simulations (Planck Collaboration III 2020a). We run our analysis for different smoothing scale values in the range $0.1^\circ - 40.0^\circ$ and check that our final results are robust against changes in the smoothing scale. The results we present in Section 3 correspond to a somewhat aggressive choice of smoothing scale (0.4°), where the non-Gaussian component is enhanced in comparison to the anisotropy that would be present in a map with stronger smoothing.

After smoothing, the template is normalized to ensure that both the Gaussian dust field and the modulated one have the same power spectrum (see Appendix A1). As a result, the non-Gaussianity only enters at the level of the power-spectrum covariance. We show the modulating template on the SO patch in Fig. 2.

In Appendix A, we derive how the dust power-spectrum covariance changes as we account for the modulating template. The idea is to capture the effect of $t(\hat{\mathbf{n}})$ as producing an effective sky mask $\tilde{w}(\hat{\mathbf{n}}) \equiv w(\hat{\mathbf{n}})t(\hat{\mathbf{n}})$ and repeat the procedure used for the masked Gaussian field. The template normalization ensures that $\langle w^2 \rangle_\Omega^2 = \langle \tilde{w}^2 \rangle_\Omega^2$. The new effective mask introduces new couplings between neighbouring ℓ -bins. This is in turn reflected in the covariance matrix of the dust field $D(\hat{\mathbf{n}}) = t(\hat{\mathbf{n}})\delta(\hat{\mathbf{n}})$:

$$\text{Cov}(C_{\ell_1}^D, C_{\ell_2}^D) = \frac{2(\bar{C}_{\ell_2}^\delta)^2}{\langle w^2 \rangle_\Omega^2 (2\ell_2 + 1)} M_{\ell_1 \ell_2}^{\tilde{w}^2}, \quad (17)$$

where the mode-coupling matrix now captures both the effect of the mask and the template via the effective mask squared, $\tilde{w}^2(\hat{\mathbf{n}})$. We show how it compares to the Gaussian case in the first two panels of Fig. 1. As can be seen in the second panel, the large-scale modulating template increases the coupling between neighbouring bins and introduces a larger variance with respect to the Gaussian case (first panel).

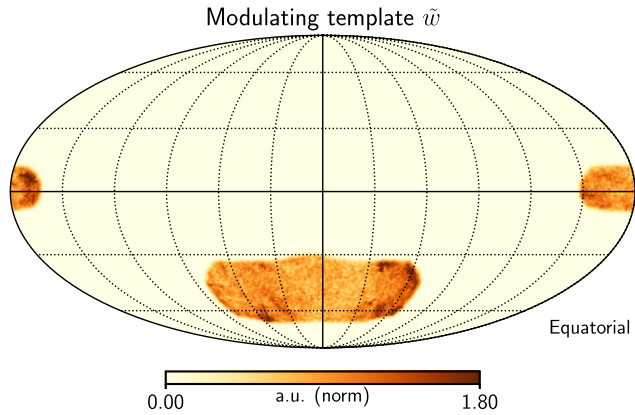


Figure 2. Large-scale modulating template in normalized arbitrary units (a.u.) for the SAT region of the SO experiment. It is the product of a uniform SO patch mask (apodized on 5° scales around the edges) with the large-scale field. The large-scale field corresponds to the *Planck* 353 GHz temperature map, smoothed on 0.4° ($\ell \approx 400$) scales. We take this map to modulate the Gaussian dust field, introducing anisotropic features which are reflected as couplings between neighbouring ℓ -bins in the power-spectrum covariance.

2.4.2 Small scales

The previous model assumes that the small-scale fluctuations are Gaussian distributed and modulated by a template. In practice, we know that this is a poor description for Galactic dust, which, for instance, exhibits a filamentary structure that cannot be reproduced from random Gaussian statistics. These filaments have been observed at millimetre wavelengths by *Planck* (e.g. Planck Collaboration XXXIII XXXVIII 2016b, c), among others, and also found to be traced by Galactic neutral hydrogen emission (Clark, Peek & Putnam 2014; Clark et al. 2015; Clark & Hensley 2019; Halal et al. 2023). To improve the modelling of small-scale fluctuations, we use the dust filament model DUSTFILAMENTS⁵ (Hervías-Caimapo & Huffenberger 2022). Filaments and their misalignment with respect to the Galactic magnetic field can quantitatively reproduce the properties of Galactic dust as observed by *Planck* at 353 GHz (Planck Collaboration XXX, XI 2016a, 2020d), such as the *BB/EE* power ratio or the dust *BB*-power-spectra power-law slope α_d (Huffenberger, Rotti & Collins 2020). Based on this idea, DUSTFILAMENTS constructs a full-sky simulation of the Galactic dust at millimetre frequencies by integrating the emission of millions of filaments in a pre-defined volume.

By construction, the filament population is calibrated to match the Large Region 71 (LR71) Galactic mask in Planck Collaboration XI (2020d), as described in Hervías-Caimapo & Huffenberger (2022). Since the polarization fraction changes depending on the portion of the sky considered, the DUSTFILAMENTS model would produce a different amplitude of polarization signal if calculated in a different region than LR71, e.g. the SO patch (see section 5 of Hervías-Caimapo & Huffenberger 2022, for an extended discussion of this point). Consequently, we re-scale the final DUSTFILAMENTS power spectra to match our chosen value for A_d in the SO patch. Moreover, the slope α_d in the simulations has the value found for the LR71 mask ($\alpha_d = -0.54$) by default. Our mock data therefore agree with this value, instead of the number previously quoted in Section 2.3.

The raw DUSTFILAMENTS model does not reproduce the large-scale features in polarization. Before computing the covariance

⁵<https://github.com/chervias/DustFilaments>

matrix from a set of 200 realizations, we employ a harmonic high-pass filter that removes the large-scale signal $\ell \lesssim 50$ from each map. The resulting dust power-spectrum covariance matrix on small scales is shown in the third panel of Fig. 1. This covariance is sourced only by the filamentary small-scale dust population, which induces important correlations between small and large scales. This effectively adds off-diagonal terms to the power-spectrum covariance and also leads to larger small-scale variance than predicted by the Gaussian covariance. We also check that the power spectra from which the covariance is derived follow a Gaussian distribution once they are binned in bandpowers of width $\Delta\ell = 30$. For this, we make histograms of the distribution of power spectra across the 200 simulations for the lowest and highest bandpower. We check that they closely follow a Gaussian distribution with a mean given by the bandpower mean and a standard deviation derived from the diagonal of the covariance matrix.

Since we have filtered out the large scales, we cannot estimate the covariance matrix at low ℓ directly from the simulations. Consequently, we merge the covariance obtained from small scales with the one obtained in the previous section for large scales. The procedure is described in Appendix A3. It involves computing power spectra at every ℓ using the usual *anafast* HEALPIX⁶ routine (Górski et al. 2005). We show the resulting correlation coefficient for the covariance matrix in the last panel of Fig. 1. This final covariance matrix Σ_{NG} is the one we use to represent the non-Gaussian anisotropic dust field.

3 RESULTS AND DISCUSSION

3.1 Baseline analysis

We show in Fig. 3 the 2D and 1D marginalized posterior distributions for the 9 theoretical model parameters. We compare the posterior distributions obtained when the likelihood incorporates the covariance matrix of a Gaussian dust field (Σ_G , filled solid blue) with the contours obtained when assuming a non-Gaussian anisotropic dust field (Σ_{NG} , unfilled dashed red). Note that the only difference between the two cases lies in the dust contribution to the covariance; the mock power spectra remain the same in both runs regardless of the employed covariance, as discussed in Section 2.1.

Judging from this plot we can describe how dust non-Gaussianities impact C_ℓ -based analyses on primordial *B*-modes. We mark in dotted grey the values used to generate the mock data, from which one can see that we obtain unbiased constraints on all parameters. This is expected since non-Gaussian dust only affects the covariance matrix in our analysis. More importantly, the uncertainties of both r and A_{lens} are insensitive to the anisotropic and non-Gaussian nature of the dust covariance matrix. We recover $\sigma(r) = 1.3 \times 10^{-3}$, in agreement with previous forecasts focused on SO that used the same noise model (Ade et al. 2019; Azzoni et al. 2021; Wolz et al. 2023)⁷ We note that this result is also consistent with the BK18 test, which also found no effect on their r constraints when using dust simulations whose amplitude is modulated by degree-scale *BB*power measured on the *Planck* 353 GHz map. However, for the dust power-spectrum parameters A_d and α_d , we do see a difference in their

⁶<http://healpix.sf.net>

⁷The $\sigma(r)$ value we report is slightly lower than other works forecasting SO performance, for which they assume a higher white noise level ($\sigma(r) = 1.7 \times 10^{-3}$, Azzoni et al. 2021) or inhomogeneous noise ($\sigma(r) = 2.1 \times 10^{-3}$; Wolz et al. 2023)




Figure 3. Comparison between posterior distributions obtained with a Gaussian and non-Gaussian dust covariance matrix (shown in filled solid blue and unfilled dashed red, respectively). The key result is that $\sigma(r)$ is *unaffected*, while the spatial dust parameter uncertainties are greater in the non-Gaussian case. When we use the covariance matrix capturing the non-Gaussian anisotropic nature of dust, we see an increase in $\sigma(\alpha_d)$ and $\sigma(A_d)$ by a factor 1.1 and 1.3, respectively. Therefore, using a Gaussian covariance to describe the statistics of dust leads to an underestimation of the dust spatial parameter uncertainties. True mock data values are shown as grey dotted lines. We recover unbiased parameter constraints in all cases. This is expected since non-Gaussian dust only affects the covariance matrix in our analysis.

posterior variance depending on the covariance matrix used. In the Gaussian case, we have $A_d = 28.08 \pm 0.93$, $\alpha_d = -0.542 \pm 0.035$, as opposed to $A_d = 28.1 \pm 1.2$, $\alpha_d = -0.543 \pm 0.040$ in the non-Gaussian case; the non-Gaussian uncertainties are larger by a factor 1.3 and 1.1, respectively. Therefore, using a Gaussian covariance to describe the statistics of dust leads to an underestimation of the dust spatial parameter uncertainties.

We also find the same qualitative behaviour in a more realistic case that includes delensing at the level expected for SO, $A_{\text{lens}} = 0.3$ (Namikawa et al. 2022). Although the off-diagonal entries of the B -mode power-spectrum covariance will become somewhat more relevant after delensing, the lensing B -modes are typically small compared to the dust foreground power, so this result is not surprising. In this paper, we do not investigate the impact of dust non-Gaussianity in the delensing tracer itself, as the results will be dependent on which tracer is used and the bias arising from non-Gaussian dust in delensing has already been investigated in other works (Beck, Errard & Stompor 2020; Baleato Lizancos et al. 2022).

We also check that this result holds for larger sky coverages ($f_{\text{sky}} \sim 0.6$) in a noise configuration without atmospheric noise

(similar to a *LiteBIRD*-like experiment). In this case, the increase in sky fraction and decrease in noise levels improves the error on r , $\sigma(r) = 5.0 \times 10^{-4}$, for both Σ_G and Σ_{NG} . Remarkably, we observe the same behaviour in this configuration even when we make a large change to the dust covariance matrix, artificially boosting every off-diagonal dust covariance matrix element to a level corresponding to a (very large) 10 per cent correlation matrix: the error on r remains at $\sigma(r) = 5.0 \times 10^{-4}$, whereas now the errors on the spatial dust parameters A_d and α_d increase by a factor 2.8 and 1.3, respectively. We show the corresponding posteriors in Fig. 4.

We also explore whether dust non-Gaussianities (incorporated via the power-spectrum covariance matrix) introduce biases in the recovered parameter values. We generate 10^5 mock power-spectrum observations \hat{C}_ℓ by performing a multivariate Gaussian sampling of the non-Gaussian power-spectrum covariance matrix Σ_{NG} . For each mock observation, we find the best-fitting sky model parameters. We find that the histogram of best-fitting values found for r and the spatial dust parameters are the same regardless of whether we use Σ_{NG} or Σ_G when evaluating the likelihood, and these are centred around the true model parameter values (Section 2.2). We thus expect dust non-Gaussianities to introduce no bias in our measurements.

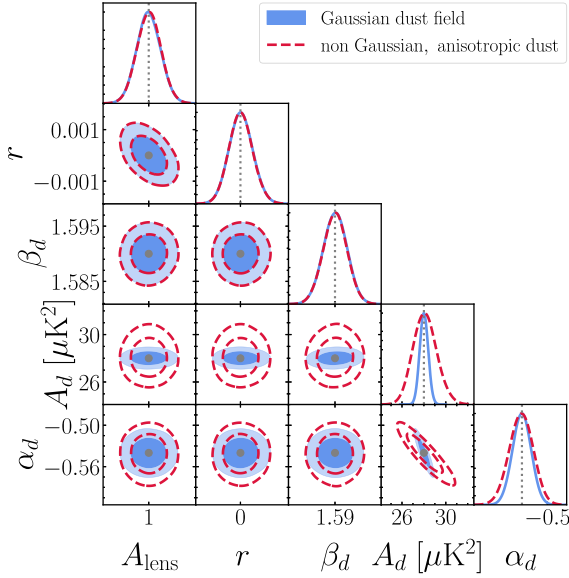


Figure 4. Comparison between posterior distributions obtained with a Gaussian and non-Gaussian dust covariance matrix (shown in filled solid blue and unfilled dashed red, respectively). We only show the most relevant parameters included in the sky model. We employ a configuration similar to a *LiteBIRD*-like experiment and artificially boost Σ_{NG} to a level corresponding to a (very large) 10 per cent correlation matrix. We find the same results as our baseline analysis: we recover the same unbiased parameter constraints in both (Σ_{G} , Σ_{NG}) cases, with the exception of the dust spatial parameters (A_d , α_d), whose variance increases. The larger sky fraction and decrease in noise levels improves the error on r , $\sigma(r) = 5.0 \times 10^{-4}$.

We thus see that the presence of significant dust non-Gaussianity does not seem to affect the final constraints on r , and its impact is reduced to those parameters describing the spatial distribution of dust. Understanding this result is the goal of the next section.

3.2 Theoretical model for observed behaviour

To gain insight into the insensitivity of the constraints on r to dust non-Gaussianity, we employ a simple Fisher forecast to estimate parameter constraints analytically. We start by writing down a simple model (Efstathiou & Gratton 2019) for an experiment which measures only two frequencies: M (mid-frequency band, where the primordial CMB signal is present) and H (high-frequency band, dominated by dust). Ignoring any other foregrounds and lensing for simplicity, we write the data vector at the power-spectrum level as:

$$\vec{\mu} = \begin{pmatrix} C_{\ell}^{MM} \\ C_{\ell}^{MH} \\ C_{\ell}^{HH} \end{pmatrix} = \begin{pmatrix} f^2 C_{\ell}^{dd} + N_{\ell}^M + r C_{\ell}^{BB} \\ f C_{\ell}^{dd} \\ C_{\ell}^{dd} + N_{\ell}^H \end{pmatrix}, \quad (18)$$

where we have included a noise signal N_{ℓ}^v on each channel, which captures other terms we are neglecting, such as the lensing B -mode power in the mid-frequency channel or the B -mode tensor contribution in the high-frequency channel. This is a two-parameter model. f controls the ratio of the dust SED between both frequency channels, while r is the amplitude of the primordial BB -power spectrum. As before, we invoke the central limit theorem and write down a Gaussian likelihood for which the Fisher matrix elements take the form (see e.g. Heavens 2016):

$$F_{\alpha\beta} = \vec{\mu}_{,\alpha}^T \Sigma^{-1} \vec{\mu}_{,\beta}, \quad (19)$$

where Σ is the (diagonal) covariance matrix, μ is the mean of the data and where we have denoted partial derivatives with respect to model parameters as ‘ $_{,\alpha}$ ’. We have also ignored parameter dependencies in the covariance matrix, as is appropriate (Carron 2013).

Assuming Gaussian fields, we calculate the (diagonal) covariance using the Wick’s theorem (see e.g. Allison et al. 2015):

$$\text{Cov}(C_{\ell}^{v_1 v_2}, C_{\ell}^{v_3 v_4}) \propto C_{\ell}^{v_1 v_3} C_{\ell}^{v_2 v_4} + C_{\ell}^{v_1 v_4} C_{\ell}^{v_2 v_3}, \quad (20)$$

where the proportionality factor includes terms dependent on ℓ and f_{sky} . When inverting the covariance matrix, we neglect the contribution from tensor modes in the mid-frequency channel. We finally invert the Fisher matrix to obtain the marginalized errors on each parameter (see e.g. Verde 2007) – for the full expressions of these variances see Appendix B. We now take the limit $N_{\ell}^M \ll C_{\ell}^{dd}$, as is justified for an experiment such as SO, where the noise in the mid-frequency channel is at most 1 per cent of the dust power spectrum in the high-frequency channel at $30 \lesssim \ell \lesssim 150$. The expressions for the variance per ℓ simplify to:

$$\sigma_{\ell}^2(r) \propto \frac{2C_{\ell}^{dd} (f^2 N_{\ell}^H)^2 (C_{\ell}^{dd} + 2N_{\ell}^H)}{[C_{\ell}^{BB} (C_{\ell}^{dd} + N_{\ell}^H)]^2}, \quad N_{\ell}^M / C_{\ell}^{dd} \ll 1, \quad (21)$$

$$\sigma_{\ell}^2(f) \propto \frac{f^2 N_{\ell}^H}{C_{\ell}^{dd}}, \quad N_{\ell}^M / C_{\ell}^{dd} \ll 1, \quad (22)$$

where the proportionality factor is $1/[f_{\text{sky}}(2\ell + 1)]$, as there are $f_{\text{sky}}(2\ell + 1)$ modes per multiple ℓ . In the further limit $N_{\ell}^H / C_{\ell}^{dd} \ll 1$ ⁸, the expression for $\sigma(r)$ becomes:

$$\sigma_{\ell}^2(r) = \frac{2f^4}{f_{\text{sky}}(2\ell + 1)} \left[\frac{N_{\ell}^H}{C_{\ell}^{BB}} \right]^2 + \mathcal{O}((N_{\ell}^H / C_{\ell}^{dd})^2). \quad (23)$$

A direct evaluation of this expression for a SO-like configuration ($f_{\text{sky}} = 0.1$, H referring to the 225 GHz ultra-high-frequency channel) gives $\sigma(r) = 1.3 \times 10^{-3}$ at $\ell = 80$, in agreement with the value found in our baseline MCMC analysis. Moreover, we explicitly see from equation (23) that $\sigma(r)$ is independent of the dust power spectrum up to second order in $N_{\ell}^H / C_{\ell}^{dd}$. Moreover, as N_{ℓ}^H goes to zero, the error on r gets arbitrarily small thanks to sample variance cancellation (Schmittfull & Seljak 2018). This limit also results in a small error in the dust parameter f . Thus, we see that the error on r is unaffected by the statistical properties of the dust (making only reasonable assumptions about the levels of dust and noise).

The intuition for these findings is that when measuring the power spectra in the high-frequency channels, we are measuring the actual dust fluctuations. This means that one can simply subtract them without knowing the specific statistical realization of the dust, simply because the fluctuations are the same on both high- and mid-frequency channels (provided negligible levels of frequency decorrelation). For this to work, we need an estimate of f that is not limited by our knowledge of the dust statistics. This was proven in our forecast, for which the error on f tends to zero as the noise in the high-frequency channel becomes insignificant. Another way of understanding this result is that, with a good enough estimate of the dust frequency-scaling f , the MF channel (where the CMB signal is most prominent) can be cleaned from dust using the UHF channel (where dust is dominant) by a simple subtraction at the map level. The power-spectrum residuals in the MF map will scale as $\sigma_{\ell}^2(f)$ (equation 22) times C_{ℓ}^{dd} , which is indeed independent of C_{ℓ}^{dd} itself. This map-level subtraction holds for whatever the dust realization

⁸The noise in the high-frequency channel of a SO-like experiment will be at most 10 per cent of the dust power-spectrum signal at $30 \lesssim \ell \lesssim 150$.

may be (Gaussian or not). As a result, $\sigma(r)$ is not affected. In fact, this technique has been implicitly used in most map-based component separation methods (Efstathiou, Gratton & Paci 2009; Efstathiou & Gratton 2019; de Belsunce, Gratton & Efstathiou 2022, 2023). We have now explicitly shown why these methods are immune to the statistics of the foreground fields; given this result, it appears that this operation at the map level is being ‘implicitly’ performed in the C_ℓ -likelihood. For the toy model at hand, where all parameters are linear, one can show that the power-spectrum likelihood indeed yields the same results as a map-based likelihood. Although r is unaffected, from equation (22) we can see that our ability to constrain the dust field properties is directly determined by our measurement of the dust field itself. Hence, the absence of a reliable method to quantify the level of foreground non-Gaussianity directly impacts the error estimation on the spatial dust parameters A_d and α_d .

3.3 Impact on goodness of fit

We have seen that C_ℓ -based foreground cleaning constraints on r are not impacted by a non-Gaussian, anisotropic dust field. However, we see that the uncertainties on the dust spatial parameters (A_d , α_d) can become significantly underestimated in the Gaussian approximation. We now investigate how this impacts goodness-of-fit χ^2 values, which are common way to assess the suitability of the model being used in the likelihood (e.g. BK18; Wolz et al. 2023) or to search for systematic errors in the measurements.

We take the 10^5 mock power-spectrum observations \hat{C}_ℓ generated by performing a multivariate Gaussian sampling of Σ_{NG} . We study two cases. In the first case, we directly compute the χ^2 value, defined as:

$$\chi^2 = (\hat{C}_\ell - C_\ell^{\text{model}})^T \Sigma^{-1} (\hat{C}_\ell - C_\ell^{\text{model}}), \quad (24)$$

where C_ℓ^{model} refers to our initial data vector (Section 2.3) that we generated with the sky model presented in Section 2.2. The χ^2 value is computed twice, for both $\Sigma = \Sigma_G$ and $\Sigma = \Sigma_{\text{NG}}$. In the second case, we first refit the mock observations and obtain best-fitting values C_ℓ^{bf} . The χ^2 value is then computed for the mock data vector with respect to (w.r.t.) the best-fitting model:

$$\chi_{\text{bf}}^2 = (\hat{C}_\ell - C_\ell^{\text{bf}})^T \Sigma^{-1} (\hat{C}_\ell - C_\ell^{\text{bf}}), \quad (25)$$

for both $\Sigma = \Sigma_G$, Σ_{NG} . Finally, we obtain the probability to exceed (PTE) of the observed χ^2 value for each simulation. In order to convert from a χ^2 value to a PTE, we assume the numbers of degrees of freedom (d.o.f.) to be the same as the number of data points in the first case, and the number of data points minus the number of model parameters in the second case.

We show the PTE histograms in the form of a probability density function (PDF) in Fig. 5. The top panel corresponds to the first case (equation 24), and the bottom panel corresponds to the second case (equation 25), where the χ^2 is calculated with respect to each best-fitting model). The dashed line represents the analysis using Σ_{NG} , for which we obtain a uniform distribution since the χ^2 is computed with the same covariance that generated the data. On the other hand, the solid black line represents the PTEs obtained with Σ_G . This PTE distribution is skewed towards low values, meaning that there is a substantial number of simulations where the χ^2 are very high and significantly differ from the expected value for the data: 8 per cent of simulations have $\text{PTE} < 0.05$ (6 per cent if compared to the best fit) as opposed to the expected 5 per cent value. This would be interpreted as the model being unable to provide a good description of the data,




Figure 5. Histogram of PTE values obtained for 10^5 power-spectrum mock simulations drawn from Σ_{NG} , then analysed with Σ_{NG} (dashed line) and with Σ_G (solid line). In the top panel, the PTE is obtained directly from the raw simulations with respect to the true model. In the bottom panel, we first perform a best fit against which we compare the mock data. Different colours correspond to the inclusion of different channels. With the exception of the LF channel, which barely contains any dust, the PTE distributions are skewed towards low values. This would be interpreted as the model being unable to provide a good description of the data, whereas the actual problem is a mis-estimation of the data vector uncertainties.

whereas the actual problem is a mis-estimation of the data vector uncertainties.

Similar results were found in Azzoni et al. (2023), where the authors assumed a Gaussian covariance when analysing simulated data in which the dust component, generated using the model of Vansyngel et al. (2017), was significantly non-Gaussian. They found that, while the theoretical model used to describe the power spectra was able to obtain unbiased constraints on r , the best-fitting χ^2 values were consistently high.

We further calculate the PTE distributions corresponding to specific frequency channels (LF, MF, and UHF). When comparing to the best-fitting model, the d.o.f. to use is unclear a priori because the best fit was found including all frequency channels. We therefore resort to the distributions obtained with Σ_{NG} . For each band, we find the effective d.o.f. that produces a uniform distribution. This is then used for the equivalent calculation using Σ_G . The results are included as differently colour coded solid lines in Fig. 5. We are only able to recover the true distribution for the LF channels. This is because they barely contain any dust. As a result, a non-Gaussian dust covariance has a negligible effect on them. The strongest effect is seen for the UHF channels, which are dust dominated. Even though comparing to the best fit alleviates the differences, neglecting the impact of dust non-Gaussianity in the covariance matrix still leads to poor goodness-of-fit metrics.

3.3.1 Recovering good PTEs with cleaned maps

The reason for the degradation in the χ^2 PTE described in the previous section is our inability to fully describe the statistical uncertainties of the dust power spectrum. However, as we have shown in Section 3.1, and justified in Section 3.2, this mostly affects the final constraints on parameters related to the dust power spectrum, and does not impact the cosmological and spectral parameters. It should therefore be possible to select a sector of the data vector where the effects of dust non-Gaussianity are minimized, and for which standard goodness-of-fit metrics can be used as a reliable diagnostic of model suitability. This can be achieved by extracting a ‘cleaned’ CMB power spectrum. A similar approach at the map-based level was developed by de Belsunce, Gratton & Efstathiou (2022) and applied to *Planck* polarization maps.

Consider the harmonic coefficients of maps observed at \mathcal{N}_v different frequencies, and let us model them as arising from a linear combination of components with spectra S_c^v and instrumental noise:

$$m_{v,\ell m} = \sum_{c=1}^{\mathcal{N}_v} S_c^v A_{\ell m}^c + n_{v,\ell m}, \quad (26)$$

where $A_{\ell m}^c$ are the harmonic coefficients of component c , and $n_{v,\ell m}$ is the noise contribution. Assuming we know the spectra of all components, the maximum-likelihood estimator for the component amplitudes is

$$\vec{A}_{\ell m} = (S^\top N_\ell^{-1} S)^{-1} S^\top N_\ell^{-1} \vec{m}_{\ell m}, \quad (27)$$

where the matrix S , of size $\mathcal{N}_c \times \mathcal{N}_v$, contains the SEDs of all components, N_ℓ is the $\mathcal{N}_v \times \mathcal{N}_v$ noise power spectrum matrix, and we have promoted m_v and A^c to vectors of sizes \mathcal{N}_v and \mathcal{N}_c , respectively. If we are only interested in the CMB amplitude, we can express it as a linear combination of the input maps:

$$A_{\ell m}^{\text{CMB}} = \sum_v Q_{v,\ell} m_{v,\ell m}, \quad (28)$$

where

$$Q_{v,\ell} \equiv \sum_c (S^\top N_\ell^{-1} S)^{-1}_{\text{CMB},c} (S^\top N_\ell^{-1})_v^c. \quad (29)$$

Then, the power spectrum of the cleaned CMB map is

$$C_\ell^{\text{CMB}} = \sum_{v,v'} Q_{v,\ell} Q_{v',\ell} C_\ell^{vv'}, \quad (30)$$

and its covariance is related to the full multifrequency covariance via

$$\text{Cov}(C_\ell^{\text{CMB}}, C_{\ell'}^{\text{CMB}}) = \sum_{\substack{v_1, v_2, \\ v_3, v_4}} Q_{v_1,\ell} Q_{v_2,\ell} Q_{v_3,\ell'} Q_{v_4,\ell'} \\ \times \text{Cov}(C_\ell^{v_1 v_2}, C_{\ell'}^{v_3 v_4}), \quad (31)$$

which is the same expression derived in Azzoni et al. (2023) within the context of residual covariances in component separation methods.

We take our 10^5 mock power-spectrum observations. With their best-fitting values for the SED parameters, we compute the elements Q_v , specific to the SO frequency channels and noise spectra. We then clean the mock observations and the best-fitting signal estimates C_ℓ^{bf} with equation (30), and obtain the clean covariance matrix (equation 31). This is performed twice, once for the best fit and covariance




Figure 6. Similar to the bottom panel of Fig. 5, but we now show the PTE distributions obtained after cleaning the 10^5 power-spectrum mock simulations with the procedure outlined in Section 3.3.1. The cleaning procedure nulls the foregrounds and therefore minimizes the effect of their non-Gaussian statistics. We recover a uniform distribution for Σ_G , which reflects how it correctly summarizes the statistics of the data after cleaning. This contrasts with the originally skewed PTE distribution found for Σ_G (solid black line, bottom panel Fig. 5), added here as a solid grey line for comparison.

derived from Σ_G and once for the corresponding quantities using Σ_{NG} . We compute the corresponding χ_{bf}^2 values for each cleaned mock observation. In order to convert these to PTE values, we proceed as in the previous section and compute the effective number of d.o.f. that return a uniform PTE distribution in the Σ_{NG} case, which is the non-Gaussian covariance from which the mock observations were generated.

We show the corresponding PTE distributions in Fig. 6. The distribution we recover for Σ_G (solid red line) exactly matches that of Σ_{NG} (blue dashed line), in contrast with the previous disagreement found when no cleaning was performed (Section 3.3) and which we overplot as a solid grey line for comparison. This highlights the suitability of the procedure to null the non-Gaussian foregrounds. It can therefore be used to quantify the ability of the model to describe the data when considering only the sources whose statistical distribution we have more under control (the CMB, in this case). It is another goodness-of-fit metric that should be explored in addition to the standard χ^2 , to ward off issues with non-Gaussianities.

3.4 Foreground extensions

Perhaps the most important challenge for component separation is foreground frequency decorrelation, which may arise from inaccurate modelling of the foreground SEDs, as well as through the spatial variation of these SEDs. This has been noted in the context of map-based techniques (e.g. Armitage-Caplan et al. 2012; Kogut & Fixsen 2016; McBride, Bull & Hensley 2023), as well as in C_ℓ -based methods (BICEP2/Keck Collaboration 2015; Planck Collaboration L 2017). It can be addressed by extending the foreground model used in component separation, to directly account for these spatial variations at the map level (Eriksen et al. 2004, 2008; Alonso et al. 2017; Errard & Stompor 2019) or by forward-modelling the impact of this contamination on to the multifrequency power spectra, either through a moment expansion (Chluba, Hill & Abitbol 2017; Azzoni et al. 2021; Mangilli et al. 2021; Remazeilles, Rotti &

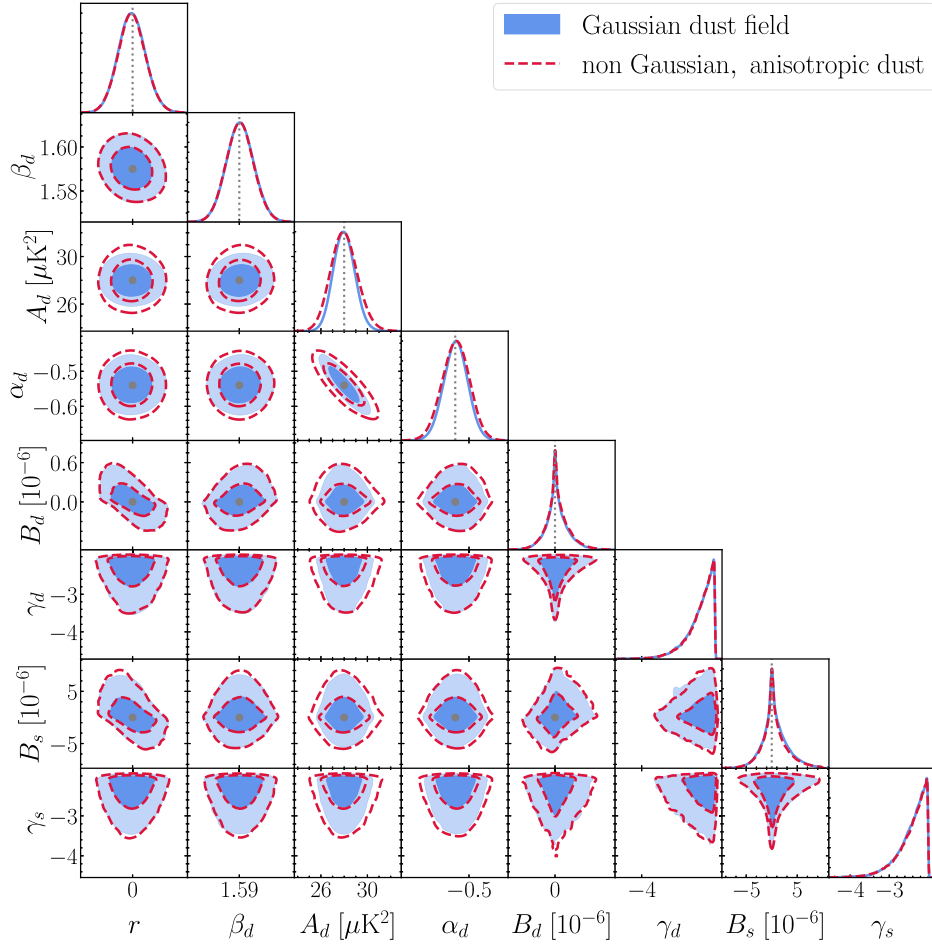


Figure 7. Similar to Fig. 3, we show the comparison between posterior distributions obtained with a Gaussian and non-Gaussian dust covariance matrix (shown in filled solid blue and unfilled dashed red, respectively). In this case, the data have been analysed with the minimal moments expansion method of Azzoni et al. (2021) to assess the impact of frequency decorrelation. This adds 4 more parameters to our model (see equation 10), although we only show a reduced data set for clarity. The amplitude of the spectral index power-spectrum fluctuations (B_d , B_s) is consistent with zero. As a result, we ignore the prior volume effects that appear on the corresponding γ_d , γ_s . We also find the same results as our baseline analysis: we recover the same unbiased parameter constraints in both (Σ_G , Σ_{NG}) cases, with the exception of the dust spatial parameters (A_d , α_d), whose variance increases. Due to the complexity of this model, the error on r increases by more than 50 per cent, $\sigma(r) = 2.0 \times 10^{-3}$, with respect to the baseline analysis.

Chluba 2021; Rotti & Chluba 2021; Sponseller & Kogut 2022; Vacher et al. 2022, 2023b), or through an effective description of frequency decorrelation (BICEP2/Keck Collaboration 2015). We here investigate whether the presence of decorrelation would affect our results, since the qualitative explanation in Section 3.2 relies on the assumption of 100 per cent correlation between frequency channels.

To do this, we first repeat our analysis using an extended foreground model, through the so-called minimal moment expansion method of Azzoni et al. (2021). The model allows for spatially varying spectral indices for both dust and synchrotron, and parametrizes those variations as uncorrelated Gaussian random fields with power-law-like power spectra. Since we are only interested in the changes that foreground non-Gaussianities introduce at the level of the power-spectrum covariance, we analyse data without any spectral index variations. This is equivalent to evaluating the parameters of the minimal moment expansion model at B_c , B_d , γ_d , $\gamma_c = 0$. The results obtained using this model for component separation are presented in Fig. 7. We find the same results as our baseline analysis: we recover

the same unbiased parameter constraints in both (Σ_G , Σ_{NG}) cases, with the exception of the dust spatial parameters α_d and A_d whose variance increases a factor 1.2 and 1.3, respectively. We obtain null amplitudes for the spectral index power-spectrum fluctuations (B_s , B_d). Due to the complexity of this model⁹, the error on r increases in both cases by more than 50 per cent, $\sigma(r) = 2.0 \times 10^{-3}$, with respect to the baseline analysis where no moments method was used. This agrees with the increase found in Azzoni et al. (2021).

The results obtained so far have been recovered from a data vector which itself contained no frequency decorrelation (even if the moment expansion model allows for it). To test if the presence of significant decorrelation in the data would change our results, we generate mock data that include decorrelation via the scale-independent decorrelation parameter Δ_d presented in Section 2.2. It is worth noting that the moment expansion model is a generalization

⁹We are now marginalizing over a model that allows for frequency decorrelation, and thus expect the additional uncertainty to leak into r .

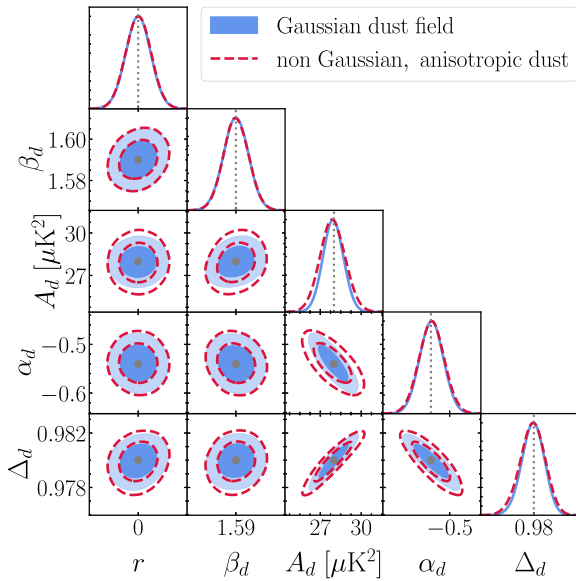


Figure 8. Similar to Fig. 3, we show the comparison between posterior distributions obtained with a Gaussian and non-Gaussian dust covariance matrix (shown in filled solid blue and unfilled dashed red, respectively) for the most relevant parameters included in the sky model with simple Δ_d decorrelation. The scale-independent parameter Δ_d characterizes the level of decorrelation between the two UHF channels. In this case, the dust power-spectrum signal itself contains 2 per cent decorrelation, although this level is small enough not to be captured in the covariance. We also find the same results as our baseline analysis: we recover the same unbiased parameter constraints in both (Σ_G , Σ_{NG}) cases, with the exception of the dust spatial parameters (A_d , α_d), whose variance increases. The complexity of decorrelation increases the error on r by 30 per cent, $\sigma(r) = 1.7 \times 10^{-3}$.

of this simple constant-decorrelation model, the latter corresponding to the limit in which spectral index variations are uncorrelated between pixels (Vansyngel et al. 2017; Azzoni et al. 2023). The amount of decorrelation present in data is still highly disputed. Sheehy & Slosar (2018) found negligible levels of decorrelation between the 217 and 353 GHz *Planck* channels, a result also found independently in Planck Collaboration XI (2020d). No significant evidence of decorrelation was found in the in the BICEP/Keck region either (BK18), a result that was further supported in an analysis including correlations with neutral hydrogen (Ade et al. 2023). Nevertheless, decorrelation between the foreground emission spectral laws at different frequencies can be detected on individual lines of sight, as noted by Pelgrims et al. (2021), and hence it is worth exploring the potential impact of its presence on our conclusions.

We choose to generate data with 2 per cent decorrelation between the 217 and 353 GHz channels for the SO SAT region. This level of decorrelation is compatible with existing data (Planck Collaboration XI 2020d; Pelgrims et al. 2021). While the power-spectrum values change, this small level of decorrelation allows us to use the same covariance matrix as in all previous tests. This incurs a 2 per cent error on the cross-frequency correlations and is expected to be irrelevant when computing the covariance matrix.

We run the MCMC inference with both data and model now including scale-independent decorrelation in the dust component. We show our posteriors in Fig. 8. While the uncertainty on r has now increased by 30 per cent, $\sigma(r) = 1.7 \times 10^{-3}$, we recover the same behaviour observed in the baseline analysis (Fig. 3). First, we obtain unbiased constraints for all parameters. Secondly, we find that

the constraint on $\sigma(r)$ is unaffected by the non-Gaussian statistics of dust at this level of decorrelation. Finally, we see a 1.1 and 1.2 factor increase in the errors of the spatial dust parameters α_d and A_d , respectively, when the data are analysed with Σ_{NG} . Thus, as long as it is included in the modelling, this level of decorrelation does not hamper our ability to employ dust maps at one frequency to constrain the dust emission at another one.

If decorrelation is not taken into account within the model, we see an increase in $\sigma(r)$. In particular, when we analyse data containing $\Delta_d = 0.98$ with our baseline model (which does not include decorrelation), we obtain $\sigma(r) = 3.5 \times 10^{-3}$ (which corresponds to a factor of 2 increase).¹⁰ This illustrates how the interpretation obtained with our simple toy model (equation 18) only applies when all frequency channels are perfectly correlated. If a significant level of decorrelation is present, a potential solution would be to break up the sky into different patches, within which the spectral index is constant and the implicit map-level subtraction can be performed (provided the patches are signal dominated and the spectral fluctuations are slowly varying).

4 CONCLUSIONS

A detection of the primordial gravitational wave background is of paramount importance: it would provide a conclusive test of inflation and open a whole new window to probe physics at the highest energies and earliest times. However, exquisite control over the polarized foreground signal is key for a reliable detection of primordial B -modes. In this paper, we addressed the standard assumption made in C_ℓ -based cleaning methods, namely that foregrounds can be approximated as Gaussian (e.g. BK18). Focusing on dust, its covariance matrix is modified to include the effects of the anisotropic and non-Gaussian nature of this field. Both effects lead to additional couplings between different ℓ bins, as well as modifying the power-spectrum uncertainties.

We studied two regimes of non-Gaussianity, introduced by (1) a large-scale modulating field, and (2) small-scale filamentary dust structures. In (1), the large-scale template, which we took to be the *Planck* 353 GHz temperature map smoothed on small scales, introduces couplings between neighbouring ℓ bins close to the dust covariance matrix diagonal. In (2), we obtained the covariance on small-scales using 200 realizations of the DUSTFILAMENTS model (Hervías-Caimapo & Huppenberger 2022). These filamentary structures couple large and small scales, and also modify the diagonal elements of the covariance matrix.

We compared the CMB and dust parameter inference obtained with a dust covariance matrix that captures these non-Gaussianities against one assuming Gaussian foregrounds. We selected noise levels and sky coverage such that we replicate the SO experiment, although we checked that our results are robust against other configurations more similar to a *LiteBIRD*-like experiment. We found that the final constraints on r are *not* affected by non-Gaussian dust foregrounds, although the error on the spatial dust parameters increases by a factor ~ 1.2 on average. This reflects how their uncertainties can be significantly underestimated in the Gaussian approximation. We believe these results hold in general, as we recovered the same behaviour for r when we artificially set all off-diagonal covariance matrix elements to have different, very large values.

¹⁰In practice, however, this value for $\sigma(r)$ is not reliable. This model would never be used to analyse the data set as the PTE goodness-of-fit metric is less than 10^{-9} .

We explained these results by arguing that C_ℓ -based constraints on r are nearly immune to the statistics of the dust field, as one can always clean the CMB signal from dust by a simple subtraction of the particular realization of dust on the sky using a dust-dominated high-frequency channel. We checked this conclusion using a simple analytical Fisher forecast that assumes 100 per cent correlation between frequency channels, although we found the same behaviour numerically for data that include decorrelation using an MCMC analysis. We also expect map-based component separation methods to be immune to the statistics of the foreground fields, given that similar arguments about removing the exact realization of the foreground field, without requiring knowledge of its statistics, apply to map-based cleaning methods.

We note, however, that even though r constraints are not affected, dust non-Gaussianity noticeably impacts goodness-of-fit analyses. This behaviour was also found by de la Hoz et al. (2022) in the context of determining polarization angles in CMB experiments. The χ^2 values computed with a covariance representing Gaussian foregrounds are degraded not because our model is unable to provide a good description of the data, but because we are unable to fully describe the statistical fluctuations of the dust power-spectrum through the Gaussian covariance matrix. We explored how this can be mitigated by using only foreground-cleaned spectrum combinations when computing goodness-of-fit statistics.

In this paper, we have considered non-Gaussianity in the foreground amplitudes, but not in the spectral index variations. The latter could add additional contributions to the non-Gaussian covariance that produce a non-trivial frequency structure. Without better data to model these SED variations, it is difficult to construct a well-educated model, so we leave this task for future work.

The arguments presented in this work should also apply to non-Gaussianities present in the synchrotron background, although we leave such a study to future work. In addition, our work highlights the importance of taking into account non-Gaussianities in any kind of Galactic science study. Although the work presented here focused on the Simons Observatory, its results will be relevant for any experiment that interprets B -mode constraints on the tensor-to-scalar ratio in the absence of a reliable method to quantify the level of foreground non-Gaussianity.

ACKNOWLEDGEMENTS

The authors acknowledge the efforts of Susanna Azzoni and Kevin Wolz, who contributed to the development of the component separation likelihood used in this work. We thank Steven Gratton and Anthony Challinor for helpful input and discussions, and Erik Rosenberg for sharing his expertise on using *Planck* data. We also thank Susanna Azzoni, Antón Baleato Lizancos, and Kevin Wolz for providing insightful comments on an early paper draft. IAC acknowledges support from Fundación Mauricio y Carlota Botton and the Cambridge International Trust. CHC acknowledges funding by Agencia Nacional de Investigación y Desarrollo, Fondo Nacional de Desarrollo Científico y Tecnológico Postdoc fellowship 3220255 and by Agencia Nacional de Investigación y Desarrollo, BASAL, FB210003. SvH was supported by the Carlsberg Foundation, acknowledges support by the Beecroft Trust, and wishes to thank Linacre College for the award of a Junior Research Fellowship. BDS acknowledges support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 851274). BDS further acknowledges support from an Science and Technology Facilities Council Ernest Rutherford Fellowship. DA acknowledges support

from the Science and Technology Facilities Council through an Ernest Rutherford Fellowship, grant reference ST/P004474.

Some of the results in this paper have been derived using the HEALPY (Zonca et al. 2019) and HEALPIX (Górski et al. 2005) packages. This research has also made extensive use of the NAMASTER (Alonso et al. 2019), SACC (Joezuntz et al. 2020), ASTROPY (Astropy Collaboration 2013), NUMPY (van der Walt, Colbert & Varoquaux 2011), SCIPY (Virtanen et al. 2020), and EMCEE (Foreman-Mackey et al. 2013) packages. We also acknowledge use of the MATPLOTLIB (Hunter 2007) and GETDIST (Lewis 2019) packages to produce the plots in this paper. This research also used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231 using NERSC award HEP-ERCAP-mp107. The Geryon cluster at the Centro de Astro-Ingeniería UC was used for some of the simulations performed in this paper. ANID BASAL project FB21000, BASAL CATA PFB-06, the Anillo ACT-86, FONDEQUIP AIC-57, and QUIMAL 130008 provided funding for several improvements to the Geryon cluster.

DATA AVAILABILITY

The *Planck* 353 GHz map was accessed from the *Planck*.¹¹ Legacy Archive at <https://pla.esac.esa.int/>. The small-scale dust filament maps were generated using the DUSTFILAMENTS (Hervías-Caimapo & Huffenberger 2022) simulation. The code to generate all other data is available at <https://github.com/iabrilcabezas/DustNGWBBP>, and will be shared on reasonable request to the corresponding author.

REFERENCES

- Abazajian K. N. et al., 2016, preprint (arXiv:1610.02743)
 Abbott L. F., Wise M. B., 1984, *Nucl. Phys. B*, 244, 541
 Abitbol M. H. et al., 2017, preprint (arXiv:1706.02464)
 Abitbol M. H. et al., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 032
 Ade P. et al., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 056
 Ade P. A. R. et al., 2023, *ApJ*, 945, 72
 Allison R., Caucal P., Calabrese E., Dunkley J., Louis T., 2015, *Phys. Rev. D*, 92, 123535
 Allys E., Levrier F., Zhang S., Colling C., Regalado-Saint Blancard B., Boulanger F., Hennebelle P., Mallat S., 2019, *A&A*, 629, A115
 Alonso D., Dunkley J., Thorne B., Naess S., 2017, *Phys. Rev. D*, 95, 043504
 Alonso D., Sanchez J., Slosar A., *LSST Dark Energy Science Collaboration*, 2019, *MNRAS*, 484, 4127
 Armitage-Caplan C., Dunkley J., Eriksen H. K., Dickinson C., 2012, *MNRAS*, 424, 1914
 Astropy Collaboration, 2013, *A&A*, 558, A33
 Azzoni S., Abitbol M. H., Alonso D., Gough A., Katayama N., Matsumura T., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 047
 Azzoni S., Alonso D., Abitbol M. H., Errard J., Krachmalnicoff N., 2023, *J. Cosmol. Astropart. Phys.*, 2023, 035
 BICEP2 Collaboration, 2018, *Phys. Rev. Lett.*, 121, 221301
 BICEP2/Keck Collaboration, 2015, *Phys. Rev. Lett.*, 114, 101301
 Baleato Lizancos A., Challinor A., Sherwin B. D., Namikawa T., 2022, *MNRAS*, 514, 5786
 Barron D. R. et al., 2022, in Zmuidzinas J., Gao J.-R. eds, *Proc. SPIE Conf. Ser. Vol. 12190, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI*. SPIE, Bellingham, p. 121900B

¹¹ *Planck* (<https://www.esa.int/Planck>) is an ESA science mission with instruments and contributions directly funded by ESA Member States, NASA, and Canada.

- Beck D., Errard J., Stompor R., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 030
- de Belsunce R., Gratton S., Efstathiou G., 2022, *MNRAS*, 517, 2855
- de Belsunce R., Gratton S., Efstathiou G., 2023, *MNRAS*, 518, 3675
- Ben-David A., von Hausegger S., Jackson A. D., 2015, *J. Cosmol. Astropart. Phys.*, 2015, 019
- Benson B. A. et al., 2014, in Holland W. S., Zmuidzinas J. eds, Proc. SPIE Conf. Ser. Vol. 9153, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII. SPIE, Bellingham, p. 91531P
- Buder I. et al., 2014, in Holland W. S., Zmuidzinas J. eds, Proc. SPIE Conf. Ser. Vol. 9153, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII. SPIE, Bellingham, p. 915312
- Cardoso J.-F., Martin M., Delabrouille J., Betoule M., Patanchon G., 2008, preprint (arXiv:0803.1814)
- Carron J., 2013, *A&A*, 551, A88
- Chluba J., Hill J. C., Abitbol M. H., 2017, *MNRAS*, 472, 1195
- Choi S. K. et al., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 045
- Clark S. E., Hensley B. S., 2019, *ApJ*, 887, 136
- Clark S. E., Peek J. E. G., Putman M. E., 2014, *ApJ*, 789, 82
- Clark S. E., Hill J. C., Peek J. E. G., Putman M. E., Babler B. L., 2015, *Phys. Rev. Lett.*, 115, 241302
- Coulton W. R., Spergel D. N., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 056
- Delabrouille J. et al., 2013, *A&A*, 553, A96
- Draine B. T., Fraisse A. A., 2009, *ApJ*, 696, 1
- Dunkley J. et al., 2009, *ApJ*, 701, 1804
- Dunkley J. et al., 2013, *J. Cosmol. Astropart. Phys.*, 2013, 025
- Efstathiou G., 2004, *MNRAS*, 349, 603
- Efstathiou G., Gratton S., 2019, preprint (arXiv:1910.00483)
- Efstathiou G., Gratton S., Paci F., 2009, *MNRAS*, 397, 1355
- Eriksen H. K. et al., 2004, *ApJS*, 155, 227
- Eriksen H. K. et al., 2006, *ApJ*, 641, 665
- Eriksen H. K., Jewell J. B., Dickinson C., Banday A. J., Górski K. M., Lawrence C. R., 2008, *ApJ*, 676, 10
- Errard J., Stompor R., 2019, *Phys. Rev. D*, 99, 043529
- Essinger-Hileman T. et al., 2014, in Holland W. S., Zmuidzinas J. eds, Proc. SPIE Conf. Ser. Vol. 9153, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII. SPIE, Bellingham, p. 915311
- Fantaye Y., Stivoli F., Grain J., Leach S. M., Tristram M., Baccigalupi C., Stompor R., 2011, *J. Cosmol. Astropart. Phys.*, 2011, 001
- Fixsen D. J., 2009, *ApJ*, 707, 916
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
- García-García C., Alonso D., Bellini E., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 043
- Goodman J., Wearé J., 2010, *Comm. Appl. Math. Comput. Sci.*, 5, 65
- Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, 622, 759
- Grain J., Tristram M., Stompor R., 2009, *Phys. Rev. D*, 79, 123515
- Gratton S., 2008, preprint (arXiv:0805.0093)
- Halal G., Clark S. E., Cukierman A., Beck D., Kuo C.-L., 2023, preprint (arXiv:2306.10107)
- Hamimeche S., Lewis A., 2008, *Phys. Rev. D*, 77, 103013
- von Hausegger S., Gammelgaard Ravnebjerg A., Liu H., 2019, *MNRAS*, 487, 5814
- Hazumi M. et al., 2019, *J. Low Temp. Phys.*, 194, 443
- Heavens A., 2016, *Entropy*, 18, 236
- Hensley B. S., Bull P., 2018, *ApJ*, 853, 127
- Hervías-Caimapo C., Huffenberger K. M., 2022, *ApJ*, 928, 65
- Hervías-Caimapo C., Bonaldi A., Brown M. L., 2016, *MNRAS*, 462, 2063
- Hivon E., Górski K. M., Netterfield C. B., Crill B. P., Prunet S., Hansen F., 2002, *ApJ*, 567, 2
- de la Hoz E., Diego-Palazuelos P., Martínez-González E., Vielva P., Barreiro R. B., Bilbao-Ahedo J. D., 2022, *J. Cosmol. Astropart. Phys.*, 2022, 032
- Huffenberger K. M., Rotti A., Collins D. C., 2020, *ApJ*, 899, 31
- Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
- Joezuntz et al., 2020, *LSSTDESC/sacc: 0.4.3*, Zenodo
- Jung G., Racine B., van Tent B., 2018, *J. Cosmol. Astropart. Phys.*, 2018, 047
- Kamionkowski M., Kosowsky A., Stebbins A., 1997, *Phys. Rev. Lett.*, 78, 2058
- Kim C.-G., Choi S. K., Flauger R., 2019, *ApJ*, 880, 106
- Knox L., 1995, *Phys. Rev. D*, 52, 4307
- Knox L., 1997, *ApJ*, 480, 72
- Kogut A., Fixsen D. J., 2016, *ApJ*, 826, 101
- Kritsuk A. G., Flauger R., Ustyugov S. D., 2018, *Phys. Rev. Lett.*, 121, 021104
- Lewis A., 2019, preprint (arXiv:1910.13970)
- Lewis A., Challinor A., 2011, Astrophysics Source Code Library, record ascl:1102.026
- LiteBIRD Collaboration, 2023, *Prog. Theor. Exp. Phys.*, 2023, 042F01
- Mak D. S. Y., Challinor A., Efstathiou G., Lagache G., 2017, *MNRAS*, 466, 286
- Mangilli A., Aumont J., Rotti A., Boulanger F., Chluba J., Ghosh T., Montier L., 2021, *A&A*, 647, A52
- Martínez-Solaache G., Karacki A., Delabrouille J., 2018, *MNRAS*, 476, 1310
- Martire F. A., Barreiro R. B., Martínez-González E., 2022, preprint (arXiv:2203.06724)
- Martire F. A., Banday A. J., Martínez-González E., Barreiro R. B., 2023, *J. Cosmol. Astropart. Phys.*, 2023, 049
- McBride L., Bull P., Hensley B. S., 2023, *MNRAS*, 519, 4370
- Miville-Deschênes M. A., Lagache G., Boulanger F., Puget J. L., 2007, *A&A*, 469, 595
- Namikawa T. et al., 2022, *Phys. Rev. D*, 105, 023511
- Nicola A., García-García C., Alonso D., Dunkley J., Ferreira P. G., Slosar A., Spergel D. N., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 067
- Page L. et al., 2007, *ApJS*, 170, 335
- Pelgrims V., Clark S. E., Hensley B. S., Panopoulou G. V., Pavlidou V., Tassis K., Eriksen H. K., Wehus I. K., 2021, *A&A*, 647, A16
- Planck Collaboration IX, 2014, *A&A*, 571, A9
- Planck Collaboration XXII, 2015, *A&A*, 576, A107
- Planck Collaboration XXX, 2016a, *A&A*, 586, A133
- Planck Collaboration XXXIII, 2016b, *A&A*, 586, A136
- Planck Collaboration XXXVIII, 2016c, *A&A*, 586, A141
- Planck Collaboration X, 2016d, *A&A*, 594, A10
- Planck Collaboration XXV, 2016e, *A&A*, 594, A25
- Planck Collaboration L, 2017, *A&A*, 599, A51
- Planck Collaboration III, 2020a, *A&A*, 641, A3
- Planck Collaboration IV, 2020b, *A&A*, 641, A4
- Planck Collaboration X, 2020c, *A&A*, 641, A10
- Planck Collaboration XI, 2020d, *A&A*, 641, A11
- Regaldo-Saint Blancard B., Levrier F., Allys E., Bellomi E., Boulanger F., 2020, *A&A*, 642, A217
- Remazeilles M., Delabrouille J., Cardoso J.-F., 2011, *MNRAS*, 418, 467
- Remazeilles M., Dickinson C., Banday A. J., Bigot-Sazy M. A., Ghosh T., 2015, *MNRAS*, 451, 4311
- Remazeilles M., Rotti A., Chluba J., 2021, *MNRAS*, 503, 2478
- Ritacco A., Boulanger F., Guillet V., Delouis J.-M., Puget J.-L., Aumont J., Vacher L., 2023, *A&A*, 670, A163
- Rotti A., Chluba J., 2021, *MNRAS*, 500, 976
- Schmittfull M., Seljak U., 2018, *Phys. Rev. D*, 97, 123540
- Seljak U., 1997, *ApJ*, 482, 6
- Seljak U., Hirata C. M., 2004, *Phys. Rev. D*, 69, 043005
- Sheehy C., Slosar A., 2018, *Phys. Rev. D*, 97, 043522
- Sherwin B. D., Schmittfull M., 2015, *Phys. Rev. D*, 92, 043005
- Simons Observatory Collaboration, 2019, *Bull. Am. Astron. Soc.*, 51, 147
- Sponseller D., Kogut A., 2022, *ApJ*, 936, 8
- Starobinsky A. A., 1979, *J. Exp. Theor. Phys. Lett.*, 30, 682
- Stompor R., Errard J., Poletti D., 2016, *Phys. Rev. D*, 94, 083526
- Suzuki A. et al., 2016, *J. Low Temp. Phys.*, 184, 805
- The Keck Array and BICEP2 Collaborations, 2021, *Phys. Rev. Lett.*, 127, 151301
- Thorne B., Dunkley J., Alonso D., Naess S., 2017, *MNRAS*, 469, 2821
- Tristram M. et al., 2022, *Phys. Rev. D*, 105, 083524
- Vacher L., Aumont J., Montier L., Azzoni S., Boulanger F., Remazeilles M., 2022, *A&A*, 660, A111
- Vacher L., Chluba J., Aumont J., Rotti A., Montier L., 2023a, *A&A*, 669, A5

- Vacher L., Aumont J., Boulanger F., Montier L., Guillet V., Ritacco A., Chluba J., 2023b, *A&A*, 672, A146
- Vansyngel F. et al., 2017, *A&A*, 603, A62
- Verde L., 2007, preprint (arXiv:0712.3028)
- Virtanen P. et al., 2020, *Nat. Methods*, 17, 261
- van der Walt S., Colbert S. C., Varoquaux G., 2011, *Comput. Sci. Eng.*, 13, 22
- Wang J., Jaffe T. R., Enßlin T. A., Ullio P., Ghosh S., Santos L., 2020, *ApJS*, 247, 18
- Wolz K. et al., 2023, preprint (arXiv:2302.04276)
- Zonca A., Singer L., Lenz D., Reinecke M., Rosset C., Hivon E., Gorski K., 2019, *J. Open Source Softw.*, 4, 1298

APPENDIX A: NON-GAUSSIAN ANISOTROPIC DUST COVARIANCES

A1 Large-scale modulating field

As discussed in the main text (Section 2.4), we write down a dust model where the Gaussian isotropic dust field $\delta(\hat{\mathbf{n}})$ is modulated by a large-scale intensity field $t(\hat{\mathbf{n}})$ (Mak et al. 2017):

$$d_{\text{NG}}(\hat{\mathbf{n}}) = w(\hat{\mathbf{n}})t(\hat{\mathbf{n}})\delta(\hat{\mathbf{n}}), \quad (\text{A1})$$

where $w(\hat{\mathbf{n}})$ is the sky mask. This model ensures that small-scale variance of the dust emission scales with the square of the large-scale intensity, as found empirically (Miville-Deschênes et al. 2007).

This model is compared against one where no modulation is present:

$$d_{\text{G}}(\hat{\mathbf{n}}) = w(\hat{\mathbf{n}})\delta(\hat{\mathbf{n}}). \quad (\text{A2})$$

One can understand the effect of $t(\hat{\mathbf{n}})$ as producing an effective sky mask $\tilde{w}(\hat{\mathbf{n}}) \equiv w(\hat{\mathbf{n}})t(\hat{\mathbf{n}})$. For both cases, we can define the corresponding pseudo- C_ℓ , whose ensemble average can be expressed as

$$\langle \tilde{C}_\ell^{d_{\text{G}}} \rangle = \sum_{\ell'} M_{\ell\ell'}^w C_{\ell'}^\delta, \quad (\text{A3})$$

$$\langle \tilde{C}_\ell^{d_{\text{NG}}} \rangle = \sum_{\ell'} M_{\ell\ell'}^{\tilde{w}} C_{\ell'}^\delta, \quad (\text{A4})$$

where C_ℓ^δ is the power spectrum of $\delta(\hat{\mathbf{n}})$, and we have defined the C_ℓ -level mode-coupling matrices $M_{\ell\ell'}^x$, which depend on the mask $x(\hat{\mathbf{n}})$ field alone and can be computed analytically with NAMASTER (Alonso et al. 2019). One property of $M_{\ell\ell'}^{\tilde{w}}$ is that

$$\sum_{\ell'} M_{\ell\ell'}^{\tilde{w}} = \int \frac{d\hat{\mathbf{n}}}{4\pi} \tilde{w}^2(\hat{\mathbf{n}}) \equiv \langle \tilde{w}^2 \rangle_\Omega, \quad (\text{A5})$$

where $\langle \dots \rangle_\Omega$ denotes sky averaging. If C_ℓ^δ is smooth, it can be pulled out of the sums in equations (A3) and (A4) to yield:

$$\langle \tilde{C}_\ell^{d_{\text{NG}}} \rangle = C_\ell^\delta \langle \tilde{w}^2 \rangle_\Omega. \quad (\text{A6})$$

With this approximation, $\langle \tilde{w}^2 \rangle_\Omega$ serves as the conversion factor between the two power spectra, without the need to invert the mode-coupling matrix (see e.g. Nicola et al. 2021).

We want to reflect the presence of the anisotropic correlating template at the covariance level, but not at the power-spectrum level. As a result, we need to normalize the initial large-scale template $t_0(\hat{\mathbf{n}})$ (which in our case we take to be the *Planck* 353 GHz map smoothed over 0.4° scales) with the following operation:

$$t(\hat{\mathbf{n}}) = t_0(\hat{\mathbf{n}}) \sqrt{\frac{\langle w^2 \rangle_\Omega}{\langle w^2 t_0^2 \rangle_\Omega}}. \quad (\text{A7})$$

This then guarantees $\langle \tilde{C}_\ell^{d_{\text{NG}}} \rangle = \langle \tilde{C}_\ell^{d_{\text{G}}} \rangle = \langle w^2 \rangle_\Omega C_\ell^\delta$.

A2 Covariance of large-scale modulated dust field

We start by considering the harmonic coefficients of both d_{G} and d_{NG} :

$$d_{\text{G},\ell m} = \sum_{\ell'm'} W_{\ell m, \ell' m'}^w \delta_{\ell' m'}, \quad (\text{A8})$$

$$d_{\text{NG},\ell m} = \sum_{\ell'm'} W_{\ell m, \ell' m'}^{\tilde{w}} \delta_{\ell' m'}, \quad (\text{A9})$$

where $W_{\ell m, \ell' m'}^x$ are the map-level mode-coupling coefficients for $x = w, \tilde{w}$. They are such that if x is slowly varying, then $W_{\ell m, \ell' m'}^x$ will have a narrow peak around $\ell = \ell', m = m'$. They are also directly related to the C_ℓ -level mode-coupling matrix:

$$M_{\ell\ell'}^x \equiv \frac{1}{2\ell + 1} \sum_{mm'} |W_{\ell m, \ell' m'}^x|^2. \quad (\text{A10})$$

We calculate the covariance of $\tilde{C}_\ell^{d_{\text{NG}}}$ via computing the 4-point function of $d_{\text{NG}}(\hat{\mathbf{n}})$, $\langle |d_{\text{NG},\ell_1 m_1}|^2 |d_{\text{NG},\ell_2 m_2}|^2 \rangle$ using the decomposition (A9). We use Wick's theorem, assuming $\delta(\hat{\mathbf{n}})$ to be Gaussian, to transform the $\langle \delta_{\ell_1 m_1} \delta_{\ell_1' m_1'}^* \delta_{\ell_2 m_2} \delta_{\ell_2' m_2'}^* \rangle$ term into different $C_\ell^\delta C_{\ell'}^\delta$ contributions. From the 4-point function we can derive the covariance of $\tilde{C}^{d_{\text{NG}}}$, $\text{Cov}(\tilde{C}_{\ell_1}^{d_{\text{NG}}}, \tilde{C}_{\ell_2}^{d_{\text{NG}}})$, which for a given pair of ℓ values involves sums over all other ℓ and corresponding m values. We use the property that for a smooth \tilde{w} , the mode-coupling coefficients are peaked around the diagonal and also assume that the C_ℓ^δ coefficients are slowly varying. As a result, we approximately pull them out of the sums in the form $\tilde{C}_{\ell_{12}}^\delta = \sqrt{C_{\ell_1}^\delta C_{\ell_2}^\delta}$ ¹². This, in turn, allows us to group the remaining terms into the mode-coupling matrix definition (A10). The final covariance reads:

$$\text{Cov}(\tilde{C}_{\ell_1}^{d_{\text{NG}}}, \tilde{C}_{\ell_2}^{d_{\text{NG}}}) = \frac{2(\tilde{C}_{\ell_{12}}^\delta)^2}{2\ell_2 + 1} M_{\ell_1 \ell_2}^{\tilde{w}^2}, \quad (\text{A11})$$

where $M^{\tilde{w}^2}$ is the mode-coupling matrix for the mask $\tilde{w}^2(\hat{\mathbf{n}}) = \tilde{w}(\hat{\mathbf{n}})\tilde{w}(\hat{\mathbf{n}})$. In the case of d_{G} , the derivation is analogous, with the only difference arising from the updated sky mask, which no longer contains the modulating template $t(\hat{\mathbf{n}})$:

$$\text{Cov}(\tilde{C}_{\ell_1}^{d_{\text{G}}}, \tilde{C}_{\ell_2}^{d_{\text{G}}}) = \frac{2(\tilde{C}_{\ell_{12}}^\delta)^2}{2\ell_2 + 1} M_{\ell_1 \ell_2}^{w^2}. \quad (\text{A12})$$

We generalize these expressions to the case of multi-frequency power spectra with the transformation:

$$2(\tilde{C}_{\ell_{12}}^\delta)^2 \rightarrow \tilde{C}_{\ell_{12}}^{\nu_1 \nu_3} \tilde{C}_{\ell_{12}}^{\nu_2 \nu_4} + \tilde{C}_{\ell_{12}}^{\nu_1 \nu_4} \tilde{C}_{\ell_{12}}^{\nu_2 \nu_3}. \quad (\text{A13})$$

Finally, we correct equation (A11) to obtain the covariance of the dust field itself, $D(\hat{\mathbf{n}}) = t(\hat{\mathbf{n}})\delta(\hat{\mathbf{n}})$. Since the normalization we perform on $t(\hat{\mathbf{n}})$ (A7) ensures that $C_\ell^\delta = C_\ell^D$, we can also write

$$\langle \tilde{C}_\ell^{d_{\text{NG}}} \rangle = C_\ell^D \langle w^2 \rangle_\Omega. \quad (\text{A14})$$

With this relation, the final covariance for the modulated dust field reads:

$$\text{Cov}(C_{\ell_1}^D, C_{\ell_2}^D) = \frac{\text{Cov}(\tilde{C}_{\ell_1}^{d_{\text{NG}}}, \tilde{C}_{\ell_2}^{d_{\text{NG}}})}{\langle w^2 \rangle_\Omega^2}. \quad (\text{A15})$$

A3 Large-scale covariance for DUSTFILAMENTS simulations

As discussed in the main text, the DUSTFILAMENTS (Hervías-Caimapo & Huffenberger 2022) simulations do not reproduce the

¹²Note this is an approximation, and the arithmetic mean could have also been chosen, $\tilde{C}_{\ell_{12}}^\delta = \frac{C_{\ell_1}^\delta + C_{\ell_2}^\delta}{2}$.

large-scale polarization features and cannot be used to estimate the covariance matrix at low ℓ . We now describe how to add a large-scale covariance (derived in the previous section) to the covariance computed from these simulations (which provides accurate answers at small scales).

Our final dust map can be split into 2 components:

$$m_T = m_L + m_S, \quad (\text{A16})$$

where m_L is a fixed large-scale template and m_S describes the small-scale component, obtained by filtering out the large scales in the DUSTFILAMENTS realization. We measure the covariance arising from the small scales, $\text{Cov}(C_\ell^{SS}, C_{\ell'}^{SS})$, from a set of DUSTFILAMENTS simulations.

We can approximate the covariance of the total map as:

$$\text{Cov}(C_\ell^{TT}, C_{\ell'}^{TT}) = \widehat{\text{Cov}}(C_\ell^{TT}, C_{\ell'}^{TT}) - \widehat{\text{Cov}}(C_\ell^{SS}, C_{\ell'}^{SS}) + \text{Cov}(C_\ell^{SS}, C_{\ell'}^{SS}), \quad (\text{A17})$$

where $\widehat{\text{Cov}}(C_{\ell_1}^{TT}, C_{\ell_2}^{TT})$ is the covariance of the large-scale modulated dust field,

$$\widehat{\text{Cov}}(C_{\ell_1}^{TT}, C_{\ell_2}^{TT}) = \frac{2\langle \bar{C}_{\ell_{12}}^{TT} \rangle^2}{2\ell_2 + 1} M_{\ell_1 \ell_2}^{w^2}, \quad (\text{A18})$$

for which $\langle \bar{C}_{\ell_{12}}^{TT} \rangle$ is simply the analytical power-spectra used in this work (Section 2.3). This ensures that non-Gaussianities are introduced only at the level of the power-spectrum covariance. The term $\widehat{\text{Cov}}(C_{\ell_1}^{SS}, C_{\ell_2}^{SS})$ is the Gaussian contribution from the small scales,

$$\widehat{\text{Cov}}(C_{\ell_1}^{SS}, C_{\ell_2}^{SS}) = \frac{2\langle \bar{C}_{\ell_{12}}^{SS} \rangle^2}{2\ell_2 + 1} M_{\ell_1 \ell_2}^{w^2}, \quad (\text{A19})$$

for which $\langle \bar{C}_{\ell_{12}}^{SS} \rangle$ is the average power spectra, computed over simulation ensembles of m_S maps. We only bin the resulting covariance

matrices at the last step. Therefore, in order to obtain $\langle \bar{C}_{\ell_{12}}^{SS} \rangle$ at each ℓ , we follow the method described in Nicola et al. (2021): we compute the power spectra at every ℓ with the usual anafast HEALPIX routine (Górski et al. 2005) on the masked maps, and then divide this quantity by $\langle w^2 \rangle_\Omega$. We verified that the covariance matrices estimated with this method were both symmetric and positive-definite in all cases.

APPENDIX B: TOY MODEL AND FISHER FORECAST

In Section 3.2, we devised a simple model (18) to gain insight into the observed behaviour for the posterior distributions (Fig. 3). The model has two parameters: r , the amplitude of the primordial BB signal and f , the dust emission ratio between the mid- and high-frequency channels. We assume Gaussian fields and compute the covariance matrix using the Knox formula (Knox 1995). The full marginalized errors on each parameter, obtained from the Fisher forecast are:

$$\sigma^2(f) = \frac{f^2 N_\ell^H C_\ell^{dd} + A_\ell}{(C_\ell^{dd})^2}, \quad (\text{B1})$$

$$\sigma^2(r) = \frac{2(f^2 N_\ell^H C_\ell^{dd} + A_\ell) [f^2 N_\ell^H (C_\ell^{dd} + 2N_\ell^H) + A_\ell]}{[C_\ell^{BB} (C_\ell^{dd} + N_\ell^H)]^2}, \quad (\text{B2})$$

where $A_\ell = (C_\ell^{dd} + N_\ell^H) N_\ell^M$. Taking into account that there are $f_{\text{sky}}(2\ell + 1)$ modes per multiple ℓ , one can check that these expressions simplify to equations (21) and (22) in the limit $N_\ell^M \ll C_\ell^{dd}$. This approximation holds for an experiment such as SO, where the noise in the mid-frequency channel is at most 10 per cent of the dust power spectrum in the high-frequency channel.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.