

Characterizing cosmic birefringence in the presence of Galactic foregrounds and instrumental systematic effects

Baptiste Jost^{1,2}, Josquin Errard¹, and Radek Stompor^{2,1}

¹*Université de Paris Cité, CNRS, Astroparticule et Cosmologie, F-75013 Paris, France*

²*CNRS-UCB International Research Laboratory, Centre Pierre Binétruy, IRL2007, CPB-IN2P3, Berkeley, California 94720, USA*

 (Received 3 April 2023; accepted 12 September 2023; published 19 October 2023)

We study a possibility of constraining isotropic cosmic birefringence with help of cosmic microwave background polarization data in the presence of polarization angle miscalibration without relying on any assumptions about the Galactic foreground angular power spectra and in particular, on the correlation between their E- and B-modes. For this purpose, we propose a new analysis framework based on a generalized parametric component separation approach, which accounts simultaneously on the presence of Galactic foregrounds, relevant instrumental effects, and external priors. We find that upcoming multi-frequency cosmic microwave background (CMB) data with appropriate calibration priors will allow producing an instrumental-effect-corrected and foreground-cleaned CMB map, which can be used to estimate the isotropic birefringence angle and the tensor-to-scalar ratio, accounting on statistical and systematic uncertainties incurred during the entire procedure. In particular, in the case of a Simons Observatory-like, three small aperture telescopes, we derive an uncertainty on the birefringence angle of $\sigma(\beta_b) = 0.07^\circ$ (0.1°), assuming the standard cosmology and calibration priors for all (one) frequency channels with the precision of $\sigma(\alpha_i) = 0.1^\circ$ as aimed at by the near future ground-based multifrequency experiments. This implies that these experiments could confirm or disprove the recently detected value of $\beta_b = 0.35^\circ$ with a significance between 3 and 5σ . We furthermore explore the impact of precision of the calibration priors and of foreground complexity on our results and discuss requirements on the calibration precision. In addition, we also investigate constraints on the tensor-to-scalar ratio, r , which can be derived in the presence of isotropic birefringence and/or polarization angle miscalibration. We find that the proposed method allows setting constraints on r in such cases, even if no prior is available, and with only a minor increase of the final uncertainty as compared to cases without these effects.

DOI: [10.1103/PhysRevD.108.082005](https://doi.org/10.1103/PhysRevD.108.082005)

I. INTRODUCTION

Cosmic birefringence rotates the polarization angle of cosmic microwave background (CMB) photons as a consequence of some parity violating mechanism, which could be due to multiple reasons ranging from a violation of the Lorentz symmetry [1] to the impact of specific dark-energy models [2] or dark-matter axionlike particles through the Chern-Simons effect [3,4]. This parity violation would lead to nonzero correlation between the E- and B-modes (EB), even if the primordial EB correlation is zero as it is the case in the standard cosmology. Cosmic birefringence can take many forms, depending on the details of the underlying physical mechanism. It can be isotropic or anisotropic, frequency dependent (similar to Faraday rotation), time dependent or constant. In this paper, we focus on the case of time- and frequency- independent, isotropic birefringence. Hereafter, we will denote the direction-independent birefringence angle as β_b and assume no primordial EB correlation, i.e., $C_\ell^{\text{EB,CMB}} = 0$. The proposed

framework is however adaptable to any nonzero primordial EB correlations, as predicted by, e.g., anisotropic inflation models [5] or an asymmetry in primordial GW handedness [6,7] (chiral gravitational waves).

Detecting cosmic birefringence is of significant scientific importance as it could provide valuable hints about the physics beyond the standard model of particle physics. This, however, requires efficient means of breaking an inherent degeneracy between the birefringence angle and an effective orientation of the polarization-sensitive CMB detectors with respect to the sky coordinate. Cosmic birefringence and the polarization angle miscalibration can also affect the estimation of other cosmological parameters such as the tensor-to-scalar ratio, r [8].

Here, we propose a self-contained framework, which permits studying both these effects in a consistent and statistically robust manner. We focus on the approach where the degeneracy is broken with the help of calibration constraints assumed to be available for all or some of the single-frequency maps. The way these calibration

constraints determine the effective polarization angle of the recovered CMB map depends on the Galactic foregrounds and the details of the instrument design, which we study in the context of the generalized parametric component separation method.

A number of studies have been published recently, addressing some of the issues mentioned above. The most relevant is the study by [8] who in the case of the Simons Observatory (SO) estimate the precision with which the polarization angle needs to be known in order to meet the targeted precision goal on r of $\sigma(r) \approx 10^{-3}$ [9]. They obtain the value of a few tenths of a degree. Similarly, in LiteBIRD [10], aiming at constraining $r \leq 10^{-3}$, they find that the polarization angle precision should be even more stringent and ranges between an arcminute and a few tens of arcminutes depending on the frequency channel [11]. Our work generalizes these approaches by providing a general, self-contained framework for such analyses, allowing to correct for the angle miscalibration to the extent possible, and setting simultaneous constraints on the birefringence and r .

From observational perspective, there exists constraints on time dependent [12–14] and on anisotropic birefringence [15–21] derived from BICEP Keck, SPT, WMAP, POLARBEAR, SPT, and Planck, datasets. Such constraints are independent on the knowledge of the absolute polarization angle as they rely on the variability, temporal, or spatial, of the signal.

For isotropic birefringence, Minami *et al.* [22] proposed to lift the degeneracy between the polarization angle and birefringence by assuming a model for the EB power spectrum of the foregrounds. The method was further elaborated on in [23]. With this model, it is possible to estimate the polarization angle of the telescope by fitting the parameters of the foreground EB to observations. The first results using this method applied to the Planck datasets are very promising and potentially hint at nonzero isotropic birefringence with $\beta_b = 0.35^\circ \pm 0.14^\circ$ [24]. Diego-Palazuelos *et al.* [23] found similar results but had to assume more complex foreground models such as filament models from [25,26] as well as the COMMANDER sky model [27] to ensure that the found value of birefringence angle is independent on the area of observed sky. Even more recently, Eskilt and Komatsu [28] used this method on the Planck and WMAP data and found $\beta_b = 0.37^\circ \pm 0.14^\circ$ with $f_{\text{sky}} = 0.62$ and $\beta_b = 0.342^{+0.094}_{-0.091}^\circ$ with $f_{\text{sky}} = 0.92$. They also found the constraints consistent with a frequency independent birefringence angle in agreement with the Planck PR4 data only analysis [29]. While these results potentially hint at the nonzero birefringence, they all rely on assumptions about the EB angular power spectrum of the foregrounds for which reliable models or measurements are lacking at this time. The approach discussed here is thus complementary to these studies.

Other relevant methods include the so-called self-calibration method of [30] and the birefringence tomography proposed in [31]. The self-calibration method

constrains the polarization angle so that the EB correlations of the signal contained in the map vanishes. This allows us, in principle, to calibrate the angle with high precision, but, by assumption, it rules out any possible detection of isotropic cosmic birefringence. The birefringence tomography aims at measuring the difference between a birefringence angle induced at recombination and one at reionization (corresponding respectively to small and large angular scales of the CMB spectra) therefore constraining the change of the birefringence angle between the two corresponding redshifts. Both these methods should ideally be applied to the foreground-cleaned CMB maps as produced at the outcome of the component separation procedures. For multifrequency observations with potentially different miscalibration angles at different frequencies, this however requires understanding the impact of such effects on the component separation and the expected level of foreground residuals in the recovered maps. This work is therefore also relevant for these methods.

In this paper, we first generalize a parametric component separation method as applied to a multifrequency set of Stokes maps so it can correct jointly for instrumental effects and clean complex foregrounds. We use calibrations of the polarization angle of the telescopes to break the polarization angle-birefringence degeneracy, whatever are the spatial properties of the foregrounds or whether cosmic birefringence is present or not. We then assess the impact of this procedure on the joint estimation of the birefringence angle, β_b , and the tensor-to-scalar ratio, r . This approach can be seen as a generalization of the self-calibration method [30] extended to allow for a detection of isotropic birefringence, simultaneously with the amplitude of the primordial gravitational waves, and explicitly accounting on the foreground contaminations.

While the presented approach should be eventually implemented within an actual CMB data analysis pipeline, hereafter we recast it as a forecasting tool in order to derive realistic and robust, ensemble-averaged constraints on the cosmological parameters, β_b and r , and to provide meaningful precision requirements for the calibration priors.

II. METHOD

Our method is composed of two steps. The first step consists in a simultaneous estimation of foreground and instrumental parameters performed as part of the generalized parametric component separation described in [32]. It uses a generalized version of the so-called spectral likelihood and yields constraints on foreground and instrumental parameters. These are then used to derive estimates of the sky components including that of the CMB, as well as their generalized statistical uncertainties.

The second step then constraints cosmological parameters from the CMB maps and their covariances as derived on the first step. This is done with help of the likelihood obtained assuming that CMB signal is (nearly) Gaussian and isotropic, and its covariance is given by the CMB

power spectra, which in turn depend on cosmological parameters. Consequently, the covariance model does not account for the presence of foreground residuals in the CMB map estimated on the first step. This may then lead to systematic errors in the estimated parameters, which we study hereafter, together with their statistical uncertainties.

We describe the entire procedure in detail below.

A. Data model

The input for the component separation step are single frequency maps reconstructed from actual measurements of a CMB experiment. These maps are gathered in a single data vector \mathbf{d} ; for each sky pixel, we store the measured sky amplitudes for each frequency, so \mathbf{d} contains $n_{\text{stokes}} \times n_f$ maps. We model \mathbf{d} as

$$\mathbf{d} = \mathbf{X}\mathbf{A}\mathbf{B}\mathbf{c} + \mathbf{n}, \quad (1)$$

where \mathbf{n} stands for the noise in all maps concatenated together. \mathbf{c} is a vector composed of all maps of the sky signals. It contains therefore $n_{\text{stokes}} \times n_{\text{comp}}$ maps. \mathbf{B} acts on \mathbf{c} and models the impact of birefringence on the CMB. \mathbf{A} is the mixing matrix that acts on the component vector synthesizing them into the corresponding frequency maps for each observation channel of the considered telescope.

Finally, \mathbf{X} is the instrumental response matrix that models how the instrument's characteristics may affect the sky signal at each frequency.

We will specifically focus on the Q and U Stokes parameters, ignoring I. Indeed the observation of temperature anisotropies with SO SAT-like instruments is quite limited especially at large angular scales. This is primarily due to $1/f$ noise, which is greatly reduced in polarization thanks to the half wave plate. Consequently, we also exclude the TB correlation when determining the birefringence angle. The formalism can however be easily generalized and can include TT and TB signals.

We define the mixing matrix for each sky pixel p as $\mathbf{A}_p(\{\beta_{fg}\})$. It scales the sky component amplitudes in frequency and coadds them together. Each column of \mathbf{A} corresponds to a component, and each row to an observation frequency for each considered Stokes parameter, either Q or U in the following. The scaling of foreground components with respect to frequency is assumed to be parametrized by a set of spectral parameters $\{\beta_{fg}\}$. We adopt the CMB units so that the elements of \mathbf{A} acting on the CMB sky signal are all equal to 1. The mixing matrix has $(n_{\text{stokes}}n_f) \times (n_{\text{stokes}}n_{\text{comp}})$ elements. Here, the mixing matrix does not mix between different Stokes parameters and corresponding elements of the mixing matrix vanish, leaving only $n_{\text{stokes}}n_f n_{\text{comp}}$ nonzero elements, and the elements of \mathbf{A} corresponding to Q and U are equal.

Given that our focus is on isotropic cosmic birefringence \mathcal{B} , we consider polarization angle misalignment as the only

instrumental effects. Other effects should be considered in a more complete analysis as discussed in Sec. V. As we consider only Q and U Stokes parameter maps, both on the input and the output of the component separation procedure. Hereafter, $n_{\text{stokes}} = 2$. The general rotation matrix in Q and U space for the spin 2 polarization headless vector is denoted $\mathcal{R}(\alpha)$ and is defined as¹

$$\begin{aligned} \begin{pmatrix} Q_{\text{out}} \\ U_{\text{out}} \end{pmatrix} &= \begin{pmatrix} \cos(2\alpha) & \sin(2\alpha) \\ -\sin(2\alpha) & \cos(2\alpha) \end{pmatrix} \begin{pmatrix} Q_{\text{in}} \\ U_{\text{in}} \end{pmatrix} \\ &\equiv \mathcal{R}(\alpha) \begin{pmatrix} Q_{\text{in}} \\ U_{\text{in}} \end{pmatrix}. \end{aligned} \quad (2)$$

Effects of cosmic birefringence at the component map level can be represented as a block diagonal matrix, with each block corresponding to a different sky pixel. As birefringence acts only on the CMB Q and U maps, the CMB-CMB subblock of each block of \mathcal{B} is a rotation matrix $\mathcal{R}(\beta_b)$, where β_b the birefringence angle. The other components are unchanged by \mathcal{B} and the remaining diagonal blocks of the matrix are equal to the identity matrix, and all off diagonal subblocks vanish. For definiteness, we will assume that properties of the ‘‘primordial’’ CMB, i.e., as contained in the component vector, \mathbf{c} , are well-defined and known, and for simplicity throughout this paper, we will take the ‘‘primordial’’ CMB EB cross-correlation to be zero. This makes the definition of the birefringence angle (in the absence of other effects studied below) well defined. However, incorporating the cases with some specific, nonvanishing, ‘‘primordial’’ EB correlations is straightforward. We redefine the sky signal to an effective sky signal after birefringence as: $s_p \equiv \mathcal{B}(\beta_b)\mathbf{c}_p$. We also note that the framework could be generalized to the anisotropic case by simply allowing the birefringence angle to vary between sky pixels.

For simplicity, we model polarization miscalibration with a single angle for each single frequency map. This can be modified as needed depending on specific experiment conditions, assigning one angle per focal plane wafer or on the contrary using one angle per entire multifrequency instrument, for instance. The miscalibration of polarization angles is described as a rotation matrix acting on each pixel of Q and U maps of a particular frequency channel. The corresponding instrumental response matrix, $\mathbf{X}(\{\alpha_1, \dots, \alpha_{n_f}\})$, is then a block diagonal matrix with a block assigned to each sky pixel and composed of frequency-specific subblocks given by a rotation matrix $\mathcal{R}(\alpha_i)$ acting on the Q and U Stokes parameters of the i^{th} frequency channel of the vector given by $\mathbf{A}_p(\{\beta_{fg}\})\mathcal{B}(\beta_b)\mathbf{c}_p$.

¹Note that some references use the other sign convention for the rotation matrix; however, this does not affect the results presented in this work.

The miscalibration angles are estimated at the same time as foreground spectral indices. We therefore introduce an effective mixing matrix defined as $\Lambda_p(\{\Gamma\}) \equiv \mathbf{X}(\{\alpha_1, \dots, \alpha_{n_f}\})\mathbf{A}_p(\{\beta_{fg}\})$, where $\{\Gamma\} \equiv \{\{\alpha\}, \{\beta_{fg}\}\}$ and $\{\alpha\}$ denotes $\{\alpha_1, \dots, \alpha_{n_f}\}$. A generic element of $\{\Gamma\}$ will be referred to as γ . We can now rewrite the data model in Eq. (3) as

$$\mathbf{d}_p = \underbrace{\mathbf{X}_p(\{\alpha_1, \dots, \alpha_{n_f}\})\mathbf{A}_p(\{\beta_{fg}\})}_{\Lambda_p(\{\Gamma\})} \underbrace{\mathbf{B}(\beta_b)\mathbf{c}_p}_{s_p} + \mathbf{n}_p. \quad (3)$$

In this perspective, the instrumental and foreground parameters are both merely unknowns of a global fitting problem; however, the manner in which they impact the entire procedure is rather different. The foreground parameters are sky-component specific and depend on the assumed model; hence, they do not depend on the number of available frequency channels. More channels permits in general better estimation of the foreground parameters. This is not always so for the instrumental parameters, as they are commonly specific to frequency channels and more frequencies typically mean more parameters. This is, for instance, the case of the polarization angle misalignment as discussed in detail later in this work. (See Ref. [33] for a counter example). In such cases, increasing the number of available channels may not improve the problem's stability and instead other means, such as priors, may need to be incorporated in the component separation formalism.

Let us consider an arbitrary instrumental angle α_0 , we can always write

$$\begin{aligned} \mathbf{X}_p(\{\alpha_1, \dots, \alpha_{n_f}\}) &= \mathbf{X}(\{\alpha_1 - \alpha_0, \dots, \alpha_{n_f} - \alpha_0\}) \\ &\quad \times \mathbf{X}_p(\{\alpha_0, \dots, \alpha_0\}) \\ &= \mathbf{X}_p(\{\alpha'_1, \dots, \alpha'_{n_f}\})\mathbf{X}_p(\{\alpha_0, \dots, \alpha_0\}), \end{aligned} \quad (4)$$

where all the angles with a prime include an extra common rotation by an angle $-\alpha_0$ and the rightmost factor in the last equation rotates all channels by the same angle α_0 to compensate for this. We note that rotating all frequency channels by the same angle is equivalent to rotating all sky components by the very same angle, i.e.,

$$\mathbf{X}_p(\{\alpha_0, \dots, \alpha_0\})\mathbf{A}_p\mathbf{s}_p = \mathbf{A}_p\tilde{\mathbf{X}}_p(\{\alpha_0, \dots, \alpha_0\})\mathbf{s}_p, \quad (5)$$

where $\tilde{\mathbf{X}}_p$ is a rotation operator analogous to \mathbf{X}_p but operating on the sky components instead of the frequency channels and for definiteness, we assume, as always, that CMB is the first component. We can therefore rewrite the signal term of our data model in Eq. (3) as

$$\begin{aligned} \mathbf{A}_p\tilde{\mathbf{X}}_p(\{\alpha_0, \dots, \alpha_0\})\mathbf{B}(\beta_b)\mathbf{c}_p &= \mathbf{A}_p\tilde{\mathbf{X}}_p(\{\alpha_0 + \beta_b, \dots, \alpha_0\})\mathbf{c}_p \\ &= \mathbf{A}_p\mathbf{B}(\alpha_0 + \beta_b)\mathbf{c}'_p. \end{aligned} \quad (6)$$

The new sky components, \mathbf{c}' , defined above, contain the same CMB signal as \mathbf{c} , and, in particular, its EB correlation continues to vanish as is required throughout here. The foreground signals are however modified due to the rotation by the angle α_0 . In the absence of any additional assumptions both \mathbf{c} and \mathbf{c}' and the corresponding birefringence angles, β_b and $\beta_b + \alpha_0$, are consistent with the data, \mathbf{d} . Consequently, the problem does not have a unique solution for foreground components and the birefringence angle. In practice, this implies that employing any of the standard methods of solving the inverse problem in Eq. (3) is going to have a degeneracy and that we will need some additional assumptions to break it. These extra assumptions could concern any of the foreground components, e.g., defining their EB cross-correlations, as done, for instance, in the approaches of [22,23], or provide some external constraints on the common rotation angle, α_0 . While both of these could be considered in the framework proposed here, this is the second option we focus on in the following.

B. Joint parametric component separation and systematic effects correction

1. Parameter estimation

We adapt the fiducial parametric component separation methods of [32] to take into account the generalized data model of Eq. (3). We replace the standard mixing matrix of [32] with the effective mixing matrix Λ_p and the sky signal vector with the effective one s_p , and we jointly fit for systematic effects and foreground parameters. Depending on the considered systematic effects some parameters might be degenerate, such as the absolute polarization angles of the detectors. As mentioned before, some of these degeneracies can be lifted using calibration priors. But the impact of those on the statistical error of the parameter estimation needs to be assessed and propagated correctly throughout the pipeline.

As in *verges2020*, adapting the spectral likelihood maximized over sky signals from [32] results in the following log-likelihood:

$$\begin{aligned} S &\equiv -2 \ln(\mathcal{L}(\{\Gamma\})) \\ &= \text{cst} - \sum_p \text{tr}(N_p^{-1}\Lambda_p(\Lambda_p^t N_p^{-1}\Lambda_p)^{-1}\Lambda_p^t N_p^{-1}\mathbf{d}_p\mathbf{d}_p^t), \end{aligned} \quad (7)$$

where N_p is the noise covariance matrix. Here it is assumed that the noise is uncorrelated between pixels. Relaxing this hypothesis in this formalism can be computationally very costly. And it could be addressed, for instance, by implementing the entire formalism in the time domain. We leave this for future work.

Optimizing this likelihood gives us an estimation of both foreground parameters β_{fg} and instrumental parameters $\{\alpha\}$. This is the likelihood we would use while analyzing a specific, actual or simulated, dataset. In the forecasting procedure, we average the likelihood over both CMB and noise realization similarly as in [32],

$$\langle S \rangle = -\sum_p \text{tr}((\mathbf{N}_p^{-1} - \mathbf{P}_p(\{\Gamma\}))\langle \mathbf{d}_p \mathbf{d}_p^t \rangle), \quad (8)$$

where \mathbf{P}_p is the projection operator defined as

$$\mathbf{P}_p(\{\Gamma\}) \equiv \mathbf{N}_p^{-1} - \mathbf{N}_p^{-1} \mathbf{\Lambda}_p (\mathbf{\Lambda}_p^t \mathbf{N}_p^{-1} \mathbf{\Lambda}_p)^{-1} \mathbf{\Lambda}_p^t \mathbf{N}_p^{-1}, \quad (9)$$

and $\langle \mathbf{d}_p \mathbf{d}_p^t \rangle$ is given by

$$\begin{aligned} \langle \mathbf{d}_p \mathbf{d}_p^t \rangle &= \langle \hat{\mathbf{\Lambda}}_p \hat{\mathbf{s}}_p \hat{\mathbf{s}}_p^t \hat{\mathbf{\Lambda}}_p^t \rangle + \langle \mathbf{n}_p \mathbf{n}_p^t \rangle \\ &= \mathbf{\Lambda}_p(\{\hat{\Gamma}\}) \langle \hat{\mathbf{s}}_p \hat{\mathbf{s}}_p^t \rangle \mathbf{\Lambda}_p^t(\{\hat{\Gamma}\}) + \mathbf{N}_p. \end{aligned} \quad (10)$$

In Eq. (10), $\hat{\mathbf{\Lambda}}_p$ is $\mathbf{\Lambda}_p$ evaluated at the true values of parameters $\{\hat{\Gamma}\}$ for both instrumental parameters $\{\hat{\alpha}\}$ and foreground parameters $\{\hat{\beta}_{fg}\}$, and $\hat{\mathbf{s}}_p$ is the true effective sky signal containing CMB and foregrounds.

We can express $\hat{\mathbf{\Lambda}}_p$ and $\hat{\mathbf{s}}_p$ in such a way as to distinguish between the CMB terms and the foregrounds term in Eq. (10). $\hat{\mathbf{s}}_p^{\text{cmb}}$ represents the two rows of $\hat{\mathbf{s}}_p$ that correspond to the Q and U CMB signals, and $\hat{\mathbf{s}}_p^{\text{fg}}$ to all the leftover

components. Similarly, $\hat{\mathbf{\Lambda}}_p^{\text{cmb}}$ is made of the two columns of $\hat{\mathbf{\Lambda}}_p$ that act on the CMB components in $\hat{\mathbf{s}}_p$, and $\hat{\mathbf{\Lambda}}_p^{\text{fg}}$ the other columns. Equation (10) becomes

$$\begin{aligned} \langle \mathbf{d}_p \mathbf{d}_p^t \rangle &= \hat{\mathbf{\Lambda}}_p^{\text{cmb}} \langle \hat{\mathbf{s}}_p^{\text{cmb}} \hat{\mathbf{s}}_p^{\text{cmb},t} \rangle \hat{\mathbf{\Lambda}}_p^{\text{cmb},t} \\ &\quad + \hat{\mathbf{\Lambda}}_p^{\text{fg}} \langle \hat{\mathbf{s}}_p^{\text{fg}} \hat{\mathbf{s}}_p^{\text{fg},t} \rangle \hat{\mathbf{\Lambda}}_p^{\text{fg},t} + \mathbf{N}_p. \end{aligned} \quad (11)$$

This expression can be straightforwardly generalized to the cases when the underlying foreground model is nonparametrizable. This can be done by replacing $\hat{\mathbf{\Lambda}}_p^{\text{fg}} \hat{\mathbf{s}}_p^{\text{fg}}$ by a vector of foreground signals at each frequency modified by the instrumental effects operator computed for the true values of the instrumental parameters, i.e., $\hat{\mathbf{X}}_p \hat{\mathbf{f}}_p$. The *parametric* model used in the effective mixing matrix $\mathbf{\Lambda}_p$ would then mismatch the data what could lead to foreground residuals and a bias in the cosmological parameters. The average over CMB and noise realizations does not affect $\hat{\mathbf{f}}_p$, and we use the output frequency maps from PySM to compute the $\hat{\mathbf{f}}_p \hat{\mathbf{f}}_p^t$ term.

The CMB term can be expressed as an average over pixel of the correlation between the Q and U CMB signals. Here, we focus on the case where the effective mixing matrix $\mathbf{\Lambda}_p$, the instrument matrix, $\hat{\mathbf{X}}_p$, and the noise covariance matrix \mathbf{N}_p are all pixel independent. We can then rewrite Eq. (8) as

$$\langle S \rangle = -\text{tr} \left\{ (\mathbf{N}^{-1} - \mathbf{P}) \left(n_{\text{pix}} \mathbf{N} + \hat{\mathbf{\Lambda}}^{\text{cmb}} \left(\sum_p \langle \hat{\mathbf{s}}_p^{\text{cmb}} \hat{\mathbf{s}}_p^{\text{cmb},t} \rangle \right) \hat{\mathbf{\Lambda}}^{\text{cmb},t} + \hat{\mathbf{X}} \sum_p \hat{\mathbf{f}}_p \hat{\mathbf{f}}_p^t \hat{\mathbf{X}}^t \right) \right\} \quad (12)$$

$$= -\text{tr} \{ n_{\text{pix}} (\mathbf{N}^{-1} - \mathbf{P}) (\mathbf{N} + \hat{\mathbf{\Lambda}}^{\text{cmb}} \mathbf{S}^{\text{cmb}} \hat{\mathbf{\Lambda}}^{\text{CMB},t} + \hat{\mathbf{X}} \mathbf{F} \hat{\mathbf{X}}^t) \}, \quad (13)$$

where n_{pix} is the total number of observed pixels over which the summation acts. And we dropped the subscript p for all the pixel-independent objects. We denote the pixel averages of the CMB sky component as \mathbf{S}^{cmb} and of foregrounds frequency maps as \mathbf{F} ,

$$\mathbf{S}^{\text{cmb}} \equiv \frac{1}{n_{\text{pix}}} \sum_p \langle \hat{\mathbf{s}}_p^{\text{cmb}} \hat{\mathbf{s}}_p^{\text{cmb},t} \rangle \quad (14)$$

$$\mathbf{F} \equiv \frac{1}{n_{\text{pix}}} \sum_p \hat{\mathbf{f}}_p \hat{\mathbf{f}}_p^t. \quad (15)$$

Only the projection matrix \mathbf{P} needs to be updated when exploring the likelihood, Eq. (12), which makes it more efficient to explore the parameter space $\{\Gamma\}$. We refer to the objects defined in Eqs. (14) and (15), as signal covariances

and together with the noise covariance; they provide a complete and necessary description of the input data, which is needed by our forecasting pipeline as shown in Fig. 1.

Depending on the number and specific nature of instrumental and foreground parameters considered in the problem, some degeneracies between the parameters may arise, see, e.g., Eq. (6) and a discussion there, preventing a robust determination of at least some of them. To deal with those, prior knowledge of some parameters may be required. It can come from instrumental calibration, for instance. We allow for priors on the miscalibration angles, $\{\alpha_i\}$ and assume them to be Gaussians with a mean, $\tilde{\alpha}_i$, corresponding to the value of the parameter as measured during the calibration campaign and the dispersion, σ_{α_i} , reflecting the error of the measurement. This can be straightforwardly generalized to other parameters. The complete log-likelihood is then given by

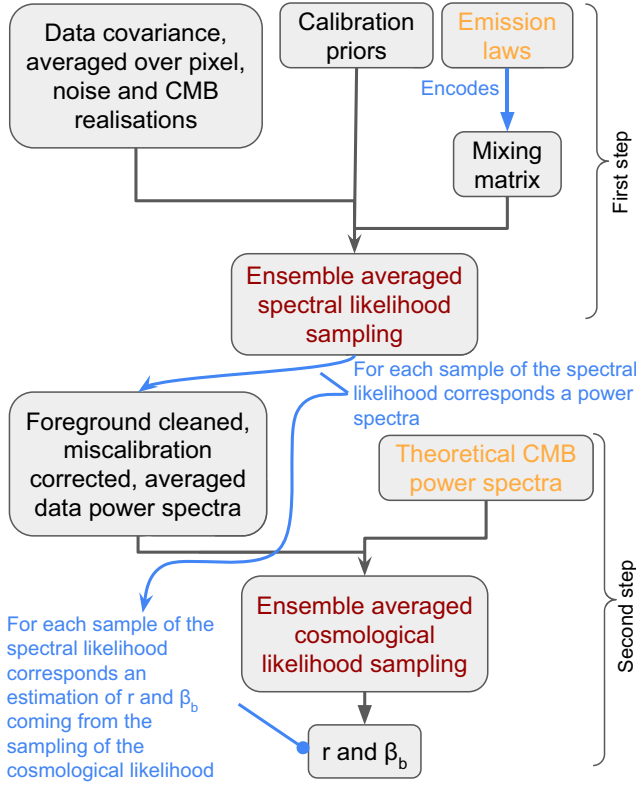


FIG. 1. Diagram depicting the main features of the method in its forecasting rendition. In red, there are the two main steps of the algorithm, in orange—the basic assumptions going in constructing the likelihoods, and in black—the various inputs and outputs of each of the main steps.

$$S' \equiv \langle S \rangle + \sum_{\alpha_i} \frac{(\alpha_i - \tilde{\alpha}_i)^2}{\sigma_{\alpha_i}^2}. \quad (16)$$

In the following, we will assume that our priors are unbiased in a sense that if multiple calibration campaigns were to be performed, the best-fit values obtained from each measurement will be drawn from a Gaussian with the mean corresponding to the true value of the parameter and the dispersion set by the measurement error. We will then average our likelihood over the ensemble of the calibration procedures. The effective log-likelihood we will use hereafter therefore reads

$$S' \equiv \langle S \rangle + \sum_{\alpha_i} \frac{(\alpha_i - \tilde{\alpha}_i)^2}{2\sigma_{\alpha_i}^2}. \quad (17)$$

2. Residuals

For each set of sampled spectral and instrumental parameters Γ , we compute the effective mixing matrix Λ and use it to get the estimate of the separated sky components,

$$\begin{aligned} m_p &= (\Lambda^t N^{-1} \Lambda)^{-1} \Lambda^t N^{-1} d_p \equiv W_p d_p \\ &= W_p \hat{\Lambda}^{\text{cmb}} \hat{s}_p^{\text{cmb}} + W_p \hat{X} \hat{f}_p + n_p \\ &= W_p \hat{\Lambda} \hat{s}_p + n_p, \end{aligned} \quad (18)$$

where the last equality is only true if the actual foreground signal at the required frequencies can be represented as a linear combination of the foreground component templates. The sought-after CMB map corresponds then to the first element of the sky component estimates, m . The latter provides an unbiased (over the statistical ensemble of noise realizations) estimate of the true sky components if $W_p \hat{\Lambda} = 1$ for all pixels p . This will be in general only true if the foregrounds can indeed be modelled as a linear combination of some templates [as in the last line of Eq. (18)], and the true mixing matrix known, so $\Lambda = \hat{\Lambda}$. However, if Λ needs to be recovered from data then, in the best case, the equality above will hold only on average and, case-by-case, the estimates of the sky components will include contributions from the others. Hereafter, we refer to these additional contributions as residuals.

We split the component vector, s , the mixing matrix, Λ , and the map-making operator, W_p into a CMB and foreground parts. For W_p the split is performed row-wise. We can then express noise-free CMB map estimate as

$$s_p^{\text{cmb}} = W_p^{\text{cmb}} (\hat{\Lambda}^{\text{cmb}} \hat{s}_p^{\text{cmb}} + \hat{X} \hat{f}_p). \quad (19)$$

Note again that here s_p^{cmb} has two elements corresponding to two Stokes parameters.

As highlighted by Eq. (19), there are two types of effects which affect the estimation of the CMB map:

- (i) A *multiplicative effect* coming from the action of $W_p^{\text{cmb}} \hat{\Lambda}^{\text{cmb}}$ on \hat{s}_p^{cmb} . Without miscalibration $W_p^{\text{cmb}} \hat{\Lambda}^{\text{cmb}}$ would be the identity. However, this is not so here due to the *a priori* unknown instrumental response matrix X . Consequently, and unlike in [32], in our case not all of the actual CMB signal is bound to end up in the CMB map estimate. Instead, part of it may leak to the estimates of the other components, and the CMB signal found in the CMB estimate may be corrupted.
- (ii) An *additive effect* coming from the contamination of foregrounds in $W_p^{\text{cmb}} \hat{X} \hat{f}_p$ which is closely related to the usual definition of residuals in the context of parametric component separation.

Both these effects will in general give rise to a residual in the recovered CMB map either by adding some spurious foreground signal—the additive effect—or by directly misestimating the CMB signal—the multiplicative effect. This residual can subsequently potentially bias the estimation of cosmological parameters. We note that due to the presence of the multiplicative effect, related to the inclusion of the instrumental effects, the expression for the residual becomes more complex in our case than it was in the

original formalism of [32], see, e.g., [33]. Consequently, a direct generalization of that formalism is rather more cumbersome. In the formalism proposed hereafter, we therefore do not perform Taylor expansion of the residuals with respect to the parameters. Instead, while computing the data matrix, we compute analytically only the term due to the multiplicative effect while the additive effect is computed numerically as in Eq. (19) for any value of instrumental and foreground parameters and averaging over these is performed with the help of sampling of the spectral likelihood. This simplifies the formalism significantly and makes no assumption that the errors on spectral parameters derived from the spectral likelihood are small. A possible downside of this approach is that we lose some insight into the morphology and sources of the residuals.

C. Cosmological parameter estimation

We use the cleaned and corrected CMB map obtained in the previous step, Eq. (18), in order to estimate the cosmological parameters that we are interested in, using the standard cosmological likelihood. The estimated CMB maps consist of multiple contributions as detailed in Eq. (19). We use it then to characterize the statistical properties of the estimated maps averaged over the ensemble of CMB and noise realizations, thus including the effects due to the component separation. These are then incorporated in a construction of the ensemble averaged cosmological likelihood, which is subsequently used to derive forecasts concerning cosmological parameters.

In the remainder of this section, we detail the procedure and mathematical framework in its most general rendition, specializing it to the case of the joint estimation of the tensor-to-scalar ratio, r , and the cosmic birefringence angle, β_b , only at the end.

1. CMB covariances in harmonic domain

We build the cosmological likelihood in the harmonic domain as this is convenient and sufficient for the purpose of the forecasting pipeline. However, the analogous constructions can be performed in the pixel domain as could be more appropriate for actual, case-by-case applications of the procedure. In the harmonic domain, we represent sky maps via their harmonic expansion coefficients denoted as \mathbf{a}_j , where j is related to the multipole numbers (ℓ, m) as $j \equiv \ell^2 + \ell + m$. As before, we use a hat to distinguish the true values from their estimates. We collect the harmonic coefficients for the combined foreground signals for all frequency bands in a single vector denoted, $\hat{\mathbf{a}}^{\text{fg}}$, and those of the CMB signals into a single vector $\hat{\mathbf{a}}^{\text{cmb}}$. The CMB signal estimate in the harmonic domain after the generalized component separation is then computed *case by case* for each set of values of Γ , e.g., a sample from the generalized spectral likelihood, it is given by

$$\mathbf{a}_j^{\text{cmb}} = \mathbf{W}^{\text{cmb}} \hat{\mathbf{\Lambda}}^{\text{cmb}} \hat{\mathbf{a}}_j^{\text{cmb}} + \mathbf{W}^{\text{cmb}} \hat{\mathbf{X}} \hat{\mathbf{a}}_j^{\text{fg}} + \mathbf{W}^{\text{cmb}} \hat{\mathbf{a}}_j^{\text{noise}}. \quad (20)$$

Our cosmological likelihood is averaged over the CMB and noise realizations, and therefore the data are represented by the covariance of the recovered CMB maps after generalized component separation defined as $\mathbf{E}_{jj'} \equiv \langle \mathbf{a}_j \mathbf{a}_{j'}^t \rangle$. Using Eq. (20) and assuming that there is no cross-correlations between the CMB, foregrounds and the noise, we get

$$\begin{aligned} \mathbf{E}_{jj'} &= \mathbf{W}^{\text{cmb}} \hat{\mathbf{\Lambda}}^{\text{cmb}} \mathbf{C}_\ell^{\text{cmb}} \delta_{jj'} \hat{\mathbf{\Lambda}}^{\text{cmb},t} \mathbf{W}^{\text{cmb},t} \\ &+ \mathbf{W}^{\text{cmb}} \hat{\mathbf{X}} \hat{\mathbf{a}}_j^{\text{fg}} \hat{\mathbf{a}}_{j'}^{\text{fg},t} \hat{\mathbf{X}}^t \mathbf{W}^{\text{cmb},t} + \mathbf{C}_\ell^{\text{noise}} \delta_{jj'}, \end{aligned} \quad (21)$$

where $\mathbf{C}_\ell^{\text{cmb}} \equiv \frac{1}{2\ell+1} \sum_m \langle \hat{\mathbf{a}}_j^{\text{cmb}} \hat{\mathbf{a}}_j^{\text{cmb},t} \rangle$, and $\mathbf{C}_\ell^{\text{noise}}$ stand for the CMB and noise spectra, respectively. As we consider Q and U polarization information only, the harmonic coefficients are either of E or B type and the spectra contain the auto-, EE, BB, and cross-, EB, spectra. As we treat the foreground as templates, $\hat{\mathbf{a}}_j^{\text{fg}}$ is not affected by the averaging over noise and CMB realizations and the matrix $\hat{\mathbf{a}}_j^{\text{fg}} \hat{\mathbf{a}}_{j'}^{\text{fg},t}$ contains products of all multipole coefficients of polarized foreground component. Computing those may pose a significant challenge as they may depend on fine details of the foreground models. However, as discussed in [32] and shown below, in the computation of the likelihood we in fact only need the auto- and cross- spectra of all foreground signals. This not only speeds up the calculations but as the spectra are generally much better known, our predictions are more robust and reliable.

The first term of $\mathbf{E}_{jj'}$ in Eq. (21) depends explicitly on \mathbf{W}^{cmb} reflecting the effect of the estimation of the instrumental parameters in the generalized likelihood on the CMB content in the estimated CMB signal. As mentioned earlier, see also [32], in the case without instrumental effects, $\mathbf{W}^{\text{cmb}} \hat{\mathbf{\Lambda}}^{\text{cmb}} = \mathbf{1}$, this term would reduce to $\mathbf{C}_\ell^{\text{cmb}} \delta_{jj'}$. The instrumental parameters also affect the second term of Eq. (21). This term produces a nonvanishing contribution even in the absence of instrumental effects; it is however modified if they are present.

The last term quantifies the contribution due to the noise in the CMB map obtained after generalized component separation and expressed via its power spectrum given by

$$\begin{aligned} \mathbf{C}_\ell^{\text{noise}} &= \mathbf{W}^{\text{cmb}} \frac{1}{2\ell+1} \sum_m \langle \hat{\mathbf{a}}_j^{\text{noise}} \hat{\mathbf{a}}_j^{\text{noise},t} \rangle \mathbf{W}^{\text{cmb},t} \\ &= [(\hat{\mathbf{\Lambda}}^t \mathbf{N}_\ell^{-1} \hat{\mathbf{\Lambda}})^{-1}]_{\text{cmb} \times \text{cmb}}, \end{aligned} \quad (22)$$

where the power spectrum \mathbf{N}_ℓ , for each frequency band is defined using the characteristics of the considered instrument; for instance, for the Simons Observatory, it is given by the following formula with the SAT's $1/f$ power index [9],

$$N_\ell^i \equiv (w_i)^{-1} e^{\left(\ell(\ell+1) \frac{\text{FWHM}_i^2}{8 \log^2}\right)} \left(\left(\frac{\ell}{\ell_{\text{knee}}^i} \right)^{-2.4} + 1 \right), \quad (23)$$

with w_i^{-1} the sensitivity of the frequency channel i in $(\mu\text{K} \cdot \text{rad})^2$. Here, we also take into account the effect of the beam and the correlated noise. We note that, as mentioned in [32], our procedure is not fully consistent concerning the beam and correlated noise treatment; as for simplicity in the generalized spectral likelihood, the noise is assumed to be whitelike all the way down to the pixel scale, and no beams are accounted for. Here, FWHM_i stands for the full-width half maximum for the i^{th} channel, in radians. $1/f$ noise is also included via the last term and ℓ_{knee}^i is the position of the knee in harmonic space for the i^{th} frequency band. We therefore include these two effects only once the generalized mixing matrix parameters are estimated. We have assumed no frequency to frequency correlations here.

2. Cosmological likelihood

All the terms that compose our recovered CMB data in harmonic domain, Eq. (19), notwithstanding, our cosmological likelihood accounts only on the CMB signal and the noise. The corresponding covariance matrix, \mathbf{C} , is then given by

$$\mathbf{C}_{jj'} = \mathbf{C}_\ell^{\text{cmb,model}}(\theta) \delta_{jj'} + \mathbf{C}_\ell^{\text{noise}} \delta_{jj'}, \quad (24)$$

where the CMB covariance includes birefringence effect, described by the birefringence angle β_b , and the B-modes signal, described by the tensor-to-scalar ratio, r , and reads

$$\begin{aligned} & \mathbf{C}_\ell^{\text{cmb,model}}(r, \beta_b) \\ & \equiv \mathcal{R}(\beta_b) \begin{pmatrix} \mathbf{C}_\ell^{EE,p} & 0 \\ 0 & r \cdot \mathbf{C}_\ell^{BB,p} + A_L \cdot \mathbf{C}_\ell^{BB,\text{lens}} \end{pmatrix} \mathcal{R}^{-1}(\beta_b). \end{aligned} \quad (25)$$

Here, $\mathbf{C}_\ell^{BB,p}$ is a primordial B-mode spectrum computed for $r = 1$ and $\mathbf{C}_\ell^{BB,\text{lens}}$ stands for the lensed B mode power

spectrum. A_L encodes the delensing, where $A_L = 1$ means no delensing and $A_L = 0$ means total delensing. In the following, we only consider $A_L = 1$. $\mathbf{C}_\ell^{EE,p}$ is the E-mode spectrum including lensing, and, as usual, the *primordial* CMB EB cross spectra is set to zero. Models predicting nonzero primordial EB could be accounted on here in the covariance matrix, but we choose to ignore such models for simplicity and consider them in future work. We ignore the effect of r and delensing on the EE power spectrum.

The equation above assumes that isotropic birefringence acts on the lensed CMB spectra. This does not imply any loss of generality as isotropic birefringence commutes with lensing since it is scale independent. This is related to the fact that the lensing effect does not depend on the coordinate frame while the effect of isotropic birefringence can be seen as merely a coordinate change. The effect of a nonzero birefringence angle on the CMB power spectra is depicted in Fig. 2.

We can then finally input our model and data in the cosmological likelihood, which is computed case by case for each set of Γ parameters drawn from the generalized spectral likelihood distribution, e.g., Tegmark *et al.* [34],

$$\begin{aligned} \langle S^{\text{cos}} \rangle &= \text{tr} \mathbf{C}^{-1} \mathbf{E} + \ln \det \mathbf{C} \\ &= f_{\text{sky}} \sum_{\ell=\ell_{\text{min}}}^{\ell_{\text{max}}} \frac{(2\ell+1)}{2} (\text{Tr}(\mathbf{C}_\ell^{-1} \mathbf{E}_\ell) + \ln(\det(\mathbf{C}_\ell))), \end{aligned} \quad (26)$$

where

$$\mathbf{E}_\ell \equiv \frac{1}{2\ell+1} \sum_m \mathbf{E}_{jj}, \quad \text{and} \quad j = \ell^2 + \ell + m. \quad (27)$$

This shows that we only need to know m -averaged, diagonal elements of the data matrix, \mathbf{E} , which are fully defined by the cross spectra of all the foreground signals as well as the CMB power spectra. We note that this




FIG. 2. Impact of an isotropic birefringence angle on the CMB lensed spectra. The Galactic foregrounds power spectra, unaffected by birefringence and estimated here at 93 GHz on a $f_{\text{sky}} \approx 10\%$ SO SAT-like patch [9], are represented in dark red.

conclusion as well as Eq. (26) hold only if the noise is isotropic.

D. Implementation

The overall implementation of the method follows its main stages as described in the Introduction of this section and as shown in Fig. 1 in the case of the forecasting framework. Here, we provide a few more details concerning the implementation of some of the key stages for each of the two steps of the method.

On the first step, in order to evaluate spectral likelihood, we need to estimate \mathbf{S}^{cmb} , \mathbf{F} [defined in Eqs. (14) and (15), respectively] and \mathbf{N} for the CMB, foreground and noise signals. Those values are averaged over observed sky pixels as well, resulting in a 2×2 matrix encoding QQ , UU , and QU correlations only. Here we are focusing on polarized signal only, but it could be easily generalized to intensity signal as well. In practice, to get \mathbf{S}^{cmb} , we average 1000 CMB map realizations generated using the synfast function in HEALPY² [35,36] and using as input power spectra those generated by CAMB³ assuming the cosmological parameters estimated in the Planck 2018 release [37], with $\ell_{\text{min}} = 2$ and $\ell_{\text{max}} = 4000$. In this paper, we will use two sets of (r, β_b) parameters that will be described in Sec. III B 2. The assumed resolution, sky coverage, and the noise covariance, \mathbf{N} , depend on the instrument for which the forecast is performed.

The foregrounds covariance matrix, \mathbf{F} , is computed for noiseless frequency maps produced by PySM [38] and subsequently averaged over all observed pixels. The experimental and cosmological parameters, as well as foreground models are detailed in Sec. III.

The component separation code used here is based on the ForeGroundBuster (FGBuster) library⁴ and has been adapted to account for systematic effects such as polarization angles and the addition of priors as mentioned earlier. The sampling of the generalized spectral likelihood is performed using the EMCEE package [39]. We used two walkers per dimension, with 13,000 steps, and we burned the 5,000 first ones. For the fiducial SO SAT-like case explored in this paper with six miscalibration angles and two foreground parameters, it totals 128,000 samples.

On the second step, in order to evaluate the cosmological likelihood, we need to construct observed and model CMB power spectra $\mathcal{C}_\ell^{\text{cmb}}$ we use the same Planck CMB power spectra templates as the one used for the generation of CMB maps in the first step and simulations.

We use the same frequency maps generated by PySM that were utilized in the first step to compute all the auto- and cross spectra needed to compute $\hat{\mathbf{a}}_j^{\text{fg}} \hat{\mathbf{a}}_j^{\text{fg},t}$. Once computed for a given sky model and instrumental characteristics, the

relevant $\mathbf{W}^{\text{cmb}} \hat{\mathbf{X}}$ factors are applied to evaluate the contribution of foregrounds to the data after generalized component separation. It is only this last step that needs to be done for each sample of the spectral likelihood. As in the first step, we are only using polarized power spectra EE, BB, and EB, but the total intensity could easily be added if desired.

In order to propagate the statistical uncertainties incurred on the first step all the way to the estimation of cosmological parameters, we perform double sampling in which for each sample (after burning) of the spectral likelihood we compute the corresponding CMB correlation matrix \mathbf{E}_ℓ , Eq. (21) and subsequently, draw a sample from the corresponding cosmological likelihood Eq. (26). To avoid any bias coming from initial values on each of these steps when drawing from the cosmological likelihood, we use once again EMCEE with 300 steps and a burn of 299 so that we only keep the last point. Both samples put together constitute a single sample drawn from an effective joint distribution of spectral, instrumental, and cosmological parameters. We resort to this rather intricate way of sampling in order to alleviate biases due the method itself, first on the spectral/instrumental parameters and then, as a consequence, on the cosmological ones, [40]. The downside, in addition to the computational cost, is that the effective joint distribution we sample from is not merely a product of the spectral and cosmological likelihoods.

III. APPLICATION

In this section, we discuss the application of the forecasting method described in the previous section. We focus here on the case of a typical CMB ground-based telescope of third generation demonstrating the proposed framework and its performance on a specific experimental setup as described in the next section. The framework is however general and can be applied to any other CMB experiment. Below we first describe the instrument configuration, followed by the sky simulation used in the application and finally the specific analysis assumptions that we consider in this work, such as the modeling of the instrumental response matrix \mathbf{X} .

A. Instrument specifications

For concreteness, we use the configuration and noise specifications of the upcoming Simons Observatory's (SO) small aperture telescopes (SAT) as described in [9]. The three SO SATs are planned to observe the sky in six frequency channels: 27, 39, 93, 145, 225, and 280 GHz. This will help to separate the CMB signal from the astrophysical foregrounds.

The SATs will observe $f_{\text{sky}} \approx 10\%$ of the sky and generate sky maps with a typical resolution of $n_{\text{side}} = 512$, which corresponds to ~ 6.8 arcmin using the HEALPix convention [35,36]. This results in around 3×10^5 observed sky pixels.

²<http://healpix.sf.net>.

³<https://camb.info>.

⁴<https://github.com/fgbuster/fgbuster>.

TABLE I. SO SAT baseline white noise levels, $1/f$ noise properties, and FWHMs.

Frequency channel [GHz]	27	39	93	145	225	280
Polarization sensitivity [μK -arcmin]	49	30	3.8	4.7	9.0	23
ℓ_{knee}	15	15	25	25	35	40
FWHM [arcmin]	91	63	30	17	11	9

Given the resolution, sky coverage, and noise property, we fix the multipole scales at $\ell_{\text{min}} = 30$ and $\ell_{\text{max}} = 300$.

We use the publicly available code `v3calc`⁵ and the sensitivities from the SO science goals and forecast paper [9] to compute the sensitivity per frequency w_i^{-1} after five years of observation for the high frequencies focal plane (225 and 280 GHz), 5 and 4 years for two middle frequencies focal planes (93, 145 GHz), and 1 year of observation for the low frequencies focal plane (27, 39 GHz). The resulting sensitivities of the baseline white noise case used are detailed in Table I. We then define the noise per sky pixel for each frequency band i , N_p^i , used in Eq. (10), as

$$N_p^i = \frac{w_i^{-1}}{\Omega_{\text{pix}}^2}, \quad (28)$$

where we converted the sensitivities from Table I to $(\mu K \cdot \text{rad})^2$ to get w_i^{-1} , as in Eq. (23). And Ω_{pix}^2 is the area of a pixel in the sky in rad^2 . Such noise covariance corresponds to the white noise, which is uncorrelated between pixels and frequencies. This turns out to be a reasonable assumption given the latest available SO simulations [41] and sufficient for the spectral likelihood evaluation.

For the cosmological likelihood, we take into account beam effects and $1/f$ noise in Eq. (23). For the computation of N_p^i , we use the baseline sensitivity and so-called optimistic $1/f$ modes of SO SAT, as well as the SAT's beams as detailed in Table I.

For the calibration priors, we typically assume a fiducial precision of $\sigma(\alpha) = 0.1^\circ$. This is on the conservative side for a drone-borne calibrator currently validated on several telescopes in the Atacama such as ACT [42,43] and CLASS [44] and also planned to be applied to SO SATs. Current forecast for this method is $0.01^\circ \leq \sigma(\alpha) \leq 0.1^\circ$, which should be achievable in several frequency channels [45]. Other approaches, such as a mobile rotating wire grid or astrophysical sources, give typically comparable but somewhat worse precision, $< 1^\circ$ for the grid [46] and $\gtrsim 0.3^\circ$ for Tau A [47,48]. Therefore, in the following, we discuss the impact of the calibration precision on our conclusions.

⁵https://github.com/simonsobs/so_noise_models.

B. Input sky simulations

1. Foregrounds

As the fiducial foreground case, we take the “d0s0” model of `PySM` [38]. For dust, this assumes a modified blackbody parametrized by its spatially constant temperature T_d and spectral index β_d . The synchrotron emission is modeled as a power law characterized by the constant spectral index β_s . The spectral emission densities for those two components are expressed as

$$Y_{\text{sync,p}}(\nu, \beta_s) = Y_{\text{sync,p}}(\nu_{0,s}) \left(\frac{\nu}{\nu_{0,s}} \right)^{\beta_s}, \quad (29)$$

$$Y_{\text{dust,p}}(\nu, \beta_d, T_d) = Y_{\text{dust,p}}(\nu_{0,d}) B(\nu, T_d) \left(\frac{\nu}{\nu_{0,d}} \right)^{\beta_d}, \quad (30)$$

where Y can be a Q or U map of a component, expressed in MJy/sr. $Y_{\text{sync,p}}(\nu_{0,s})$ and $Y_{\text{dust,p}}(\nu_{0,d})$ are the template maps for synchrotron and dust at their respective reference frequencies $\nu_{0,s} = 23$ GHz and $\nu_{0,d} = 353$ GHz that are then scaled at the frequencies of interest for observations. $B(\nu, T_d)$ is a blackbody at temperature T_d and frequency ν . In the `PySM` “d0s0” simulation, the spatially constant spectral parameters are based on Planck results [49] and are given by

$$\beta_d = 1.54, \quad T_d = 20 \text{ K}, \quad \beta_s = -3. \quad (31)$$

As examples of a more complex foreground model, we use the “d1s1” model which is similar to “d0s0” but allows for spatial variability of the spectral parameters, $\beta_d(p)$, $T_d(p)$ and $\beta_s(p)$, as well as the model referred to as “d7s3”. Here “s3” denotes the synchrotron model which adds an extra curvature term, \mathcal{C} , to the standard power law frequency scaling, which is constant over the sky,

$$Y_{\text{sync,p}}(\nu, \beta_s) = Y_{\text{sync,p}}(\nu_{0,s}) \left(\frac{\nu}{\nu_{0,s}} \right)^{\beta_s(p) + \mathcal{C} \ln(\nu/\nu_0)}. \quad (32)$$

We also use `PySM` dust model “d7”, which uses as a template the same Planck 353 GHz map as the other two models but the frequency scaling used is based on dust grain models with different physical properties, shape, size, and temperature described in detail in [50,51]. This dust model does not have an analytic function to describe the frequency scaling of the dust template and is therefore a good benchmark to test the parametric component separation since it does not trivially conform with the assumptions of the method.

2. CMB

For the CMB, we set all the parameters to the best-fit values provided by Planck [37] and only vary two

parameters, the tensor-to-scalar ratio r and the birefringence angle β_b . We do not consider delensing in this work, i.e., $A_L = 1$ in Eq. (25). However this can be straightforwardly included in our framework. The input CMB power spectra used to get the frequency maps (as described in the previous section) are computed using the same equation as the model CMB in Eq. (25). We choose two sets of cosmological parameters in addition to the fiducial Λ CDM cosmological parameters from Planck 2018 [37],

$$r = 0.0, \quad \beta_b = 0.0^\circ \quad (33)$$

$$r = 0.01, \quad \beta_b = 0.35^\circ. \quad (34)$$

With $\beta_b = 0.35^\circ$ corresponding to the central value found in the work of Minami and Komatsu [24].

3. Instrumental effects

In the simulation, we consider the effect of a potential miscalibration of the polarization angle of each of the frequency bands of the telescope. We model this effect assuming that each recovered single frequency map has its own independent polarization angle. This angle is then to be understood as an effective angle resulting from detector-level miscalibration averaged during the map-making procedure. This assumption can be adjusted as needed: our approach is generalizable to a miscalibration angle per focal plane, one per wafer, one per pixel, etc. We also assume that the miscalibration angle is the same for all pixels of the considered maps. Again this can be relaxed if needed. In this work, for concreteness, we assume some specific, true values of the miscalibration angle for each map. They are summarized in Table II. But the presented results do not depend on the specific values assumed.

C. Analysis model and priors

For the analysis, we model the data assuming the foreground scaling models as in Eq. (30) with both dust and synchrotron parameters assumed constant over the sky. Moreover, we fix dust temperature to $T_d = 20$ K as suggested by Planck results [49] given that the SO SATs do not have enough high frequency observation bands to discriminate between T_d and β_d . Importantly, this is in agreement with one of the “d0s0” assumption, Eq. (31) and therefore, the foreground model assumed in the analysis allows for an accurate description of the data simulated in the case of the “d0s0” model of PySM. This is however not the case for the “d1s1” and “d7s3” models. For the instrumental matrix X , we assume an isotropic rotation

of the polarization angle for each of the frequencies, exactly as for the sky simulations with rotations angles at each frequency treated as free parameters. We use Gaussian priors with a precision of $\sigma_{\alpha_i} = 0.1^\circ$, Eq. (17), in the spectral likelihood to break the degeneracies between these parameters. Unless specified otherwise, the priors are centred on the true input polarization angles.

IV. RESULTS

First, we present the results of the first analysis step that retrieves the spectral indices and the miscalibration angles from noisy, foreground-contaminated, and miscalibrated frequency maps. Then, we present the results of the second step that constrains cosmological parameters. We discuss various examples to demonstrate the effects of the different contaminants on the cosmological parameters estimation, such as biases or precision loss. We then investigate the dependence of the constraint on cosmological parameters as a function of prior precision. And finally, we explore the case where calibration priors are biased and their impact on instrumental and cosmological parameters estimations.

A. Fiducial case: “d0s0”, $r = 0$, $\beta_b = 0^\circ$

We start with the simplest case of the “d0s0” PySM model that assumes constant foreground spectral parameters across the sky. Since consistent assumptions were used to build the mixing matrix A , Eq. (3)—which is used throughout this work—should accurately describe the input foreground maps.

We explore two cases, one with a prior on each frequency channel and the other with a prior on the 93 GHz channel only. In both cases, the prior precision is $\sigma_{\alpha_i} = 0.1^\circ$.

1. Prior on a single channel

First, we consider a prior on the 93 GHz channel only with a precision of $\sigma_{\alpha_{93}} = 0.1^\circ$. We choose this channel as at this frequency the foregrounds amplitude is close to minimal as compared to the CMB signal and this is where most of the calibration effort is currently being allocated.

Generalized spectral likelihood results.—As described in Sec. IID, we explore the generalized spectral likelihood, Eq. (17) with help of MCMC sampling. The results are shown in Fig. 3, where the orange contours are obtained from the MCMC samples, the purple dashed line represents the Gaussian prior on the 93 GHz channel, and the gray dashed lines the input parameters. The 1σ statistical errors of the parameter estimations are detailed in Table III.

We notice that with only one prior on one polarization angle we are able to have an unbiased estimate for all six polarization angles and two spectral indices. Indeed the fact that we use all six frequency maps simultaneously in the generalized spectral likelihood allows for deriving tight constraints on the relative angles of all the considered frequency channels with respect to a global “instrument”

TABLE II. Input polarization angle per frequency bands.

Frequency channel [GHz]	27	39	93	145	225	280
Input polarization angle [°]	1	1.66	2.33	3	3.66	4.33

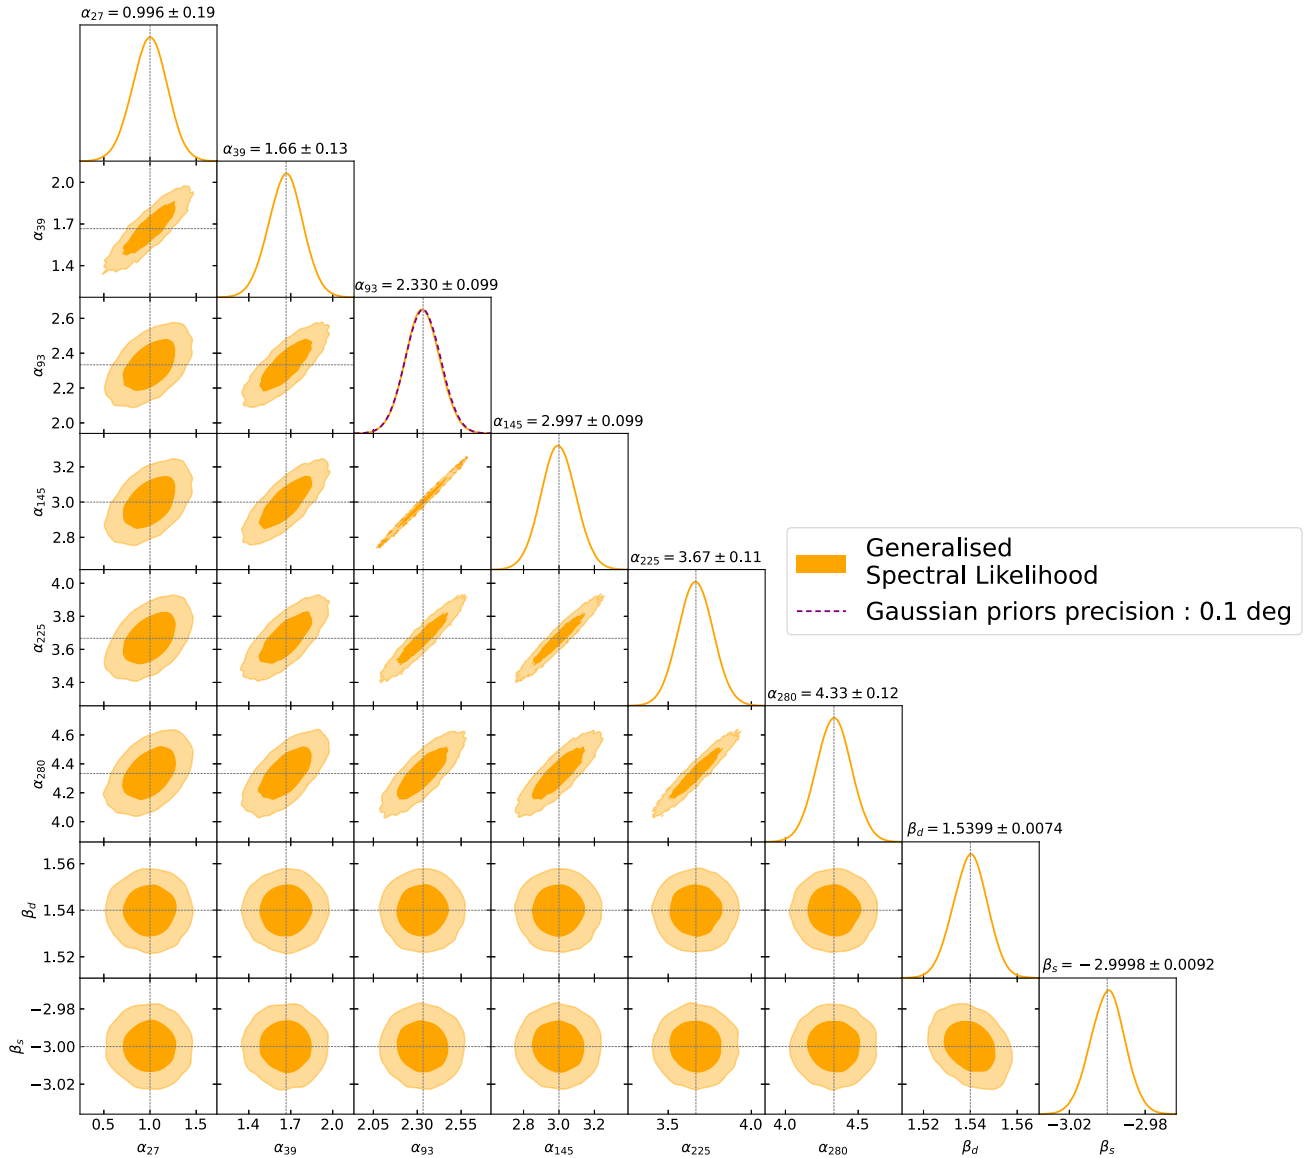


FIG. 3. Results of the generalized spectral likelihood with “d0s0” input foregrounds model. There is only one prior on the 93 GHz channel, with a precision of $\sigma_{\alpha_{93}} = 0.1^\circ$. The dashed purple lines correspond to the Gaussian priors. The orange contours correspond to the sampling of the generalized spectral likelihood. The gray dotted lines correspond to the input values.

orientation angle. The role of the prior is then to constrain the global angle which is necessary and sufficient to break the likelihood degeneracy.

The accuracy with which we can estimate the absolute polarization angles for any of the channels is therefore limited by the prior precision as summarized in Table III. The 93 GHz channel achieves the best precision on polarization angle, $\sigma(\alpha_{93}) = 0.099^\circ$, which corresponds to the prior precision (within the accuracy provided by the sampling). All other channels show larger errors as they include the error on the relative angle as set by the likelihood problem. The overall increase of the error is subdominant as compared to the prior-driven constraint on the global angle, showing that indeed the relative angles are constrained with high precision. The presence of the strong

correlations between polarization angles recovered for different frequency channels reflects the fact that the data can only constrain internally the relative angles between the frequency channels and not a common, global miscalibration error. This latter is only constrained with help of the priors. For the channels for which the relative angle uncertainty is much smaller than the prior, the constraint on the polarization angles are strongly elongated in the direction corresponding to the common change in both the angle. The effect is less pronounced, and the error contours more round, in the cases when the constraints on the relative angles are more comparable to the prior precision. The relative angle precision depends on the noise levels in both single frequency maps but also how different the underlying sky is in both the cases.

TABLE III. Summary of results for different input foreground models and instrumental parameters. All the priors used here have the precision $\sigma_{\alpha_i} = 0.1^\circ$. ★ In the last column (“d0s0”), the input cosmological parameters are $r = 0.01$ and $\beta_b = 0.35^\circ$.

Foreground input	“d0s0”		“d1s1”		“d7s3”	“d0s0” ★
	1	6	1	6	6	6
Number of priors	1	6	1	6	6	6
$\alpha_{27} [^\circ]$	1.0 ± 0.2	1.00 ± 0.08	1.0 ± 0.2	1.01 ± 0.08	1.01 ± 0.08	1.00 ± 0.08
$\alpha_{39} [^\circ]$	1.7 ± 0.1	1.67 ± 0.05	1.7 ± 0.1	1.67 ± 0.05	1.67 ± 0.05	1.67 ± 0.05
$\alpha_{93} [^\circ]$	2.3 ± 0.1	2.33 ± 0.05	2.3 ± 0.1	2.33 ± 0.05	2.33 ± 0.05	2.33 ± 0.05
$\alpha_{145} [^\circ]$	3.0 ± 0.1	3.00 ± 0.05	3.0 ± 0.1	3.00 ± 0.05	3.00 ± 0.05	3.00 ± 0.05
$\alpha_{225} [^\circ]$	3.7 ± 0.1	3.67 ± 0.05	3.7 ± 0.1	3.66 ± 0.05	3.66 ± 0.05	3.66 ± 0.05
$\alpha_{280} [^\circ]$	4.3 ± 0.1	4.33 ± 0.06	4.3 ± 0.1	4.33 ± 0.06	4.33 ± 0.06	4.33 ± 0.06
β_d	1.540 ± 0.007	1.540 ± 0.007	1.575 ± 0.008	1.575 ± 0.007	1.377 ± 0.007	1.540 ± 0.007
β_s	-3.000 ± 0.009	-3.000 ± 0.009	-3.006 ± 0.009	-3.006 ± 0.009	-3.046 ± 0.009	-3.000 ± 0.009
r	0.000 ± 0.002	0.000 ± 0.002	0.002 ± 0.002	0.002 ± 0.002	0.002 ± 0.002	0.010 ± 0.002
$\beta_b [^\circ]$	0.0 ± 0.1	0.00 ± 0.07	0.0 ± 0.1	0.00 ± 0.07	0.00 ± 0.07	0.35 ± 0.07

Furthermore, we are able to retrieve the foreground spectral parameters with a precision comparable with the standard version of parametric component separation using FGBuster applied to a SO-like case but without instrumental parameter estimation [9].

As expected, the generalized spectral likelihood yields unbiased estimates of instrumental and spectral parameters, which we then use on the next step: the estimation of cosmological parameters.

Cosmological likelihood results.—As detailed in Sec. II D, for each sample of the spectral likelihood displayed in Fig. 3, we draw one sample of the corresponding cosmological likelihood. This approach allows us to efficiently sample the full distribution, without reintroducing any method-related biases. The result is shown in Fig. 4, where the orange contours are the MCMC samples obtained assuming a single 93 GHz prior, and the gray dashed lines the input parameters. The estimations of both r and β_b are unbiased and with $r = 0.0002^{+0.0015}_{-0.0017}$, which is compatible with the SO SATs published forecasts [9], and $\sigma(\beta_b) = 0.11^\circ$, which is 10% bigger than the error bars expected from prior precision alone, which can be explained by the presence of noise and cosmic variance as we will see in Sec. IV D.

We conclude that in the case with simple foreground SEDs that match our model, and with nonzero polarization angles, the method leads to unbiased estimates of spectral and hardware parameters, and provides competitive results on r and an estimation of the birefringence angle β_b limited only by the prior precision.

2. Priors on all six channels

We now investigate the case with priors on all six polarization angles. We proceed in a similar fashion as in the previous case.

Generalized spectral likelihood results.—Figure 5 shows the results of the MCMC sampling of the generalized spectral likelihood. Comparing the priors (dashed purple) and the samples (orange) we see that, contrary to the previous case, the precision of the polarization angle estimation is better than the prior precision, $\sigma_{\alpha_i} = 0.1^\circ$,

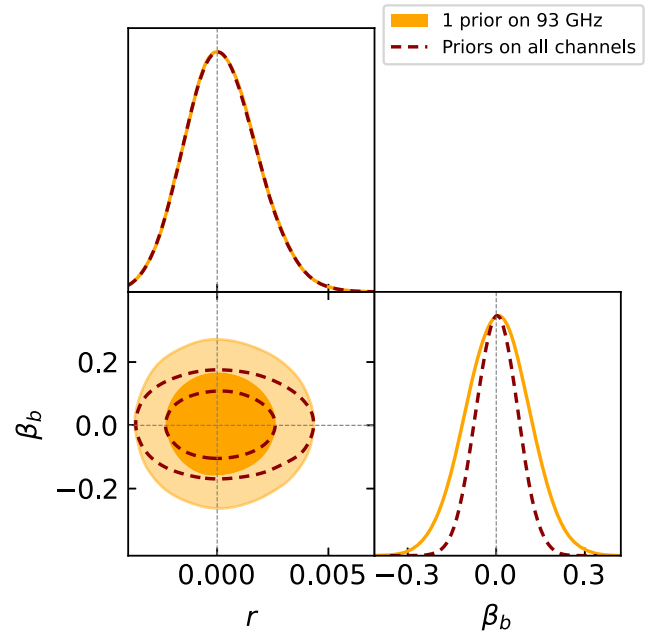


FIG. 4. Cosmological likelihood sampling, Eq. (26), after foreground cleaning and systematic effect correction with PySM “d0s0” as input and only *one* prior on the 93 GHz polarization angle (orange). Dashed dark-red contours correspond to the case with priors on all polarization angles. The gray dashed lines correspond to the input values. The central values and error bars are in Table III.




FIG. 5. Results of the generalized spectral likelihood with “d0s0” as input foregrounds model. The priors are on all six polarization angles, with a precision of $\sigma_{\alpha_i} = 0.1^\circ$. The dashed purple lines correspond to the Gaussian priors. The orange contours correspond to the sampling of the generalized spectral likelihood. The gray dotted lines correspond to the input values.

assumed for all frequency channel. This is consistent with the fact that the likelihood itself sets tight constraints on the relative angles for each frequency channel. So while the priors concern different objects, polarization angles for their respective frequency channel, each of them effectively constrains the very same global polarization angle. We thus expect that the actual constraint on this angle goes down roughly as one over square root of the number of frequency channels (however as the precision of priors gets better other contributions to the error bar, such as noise, become dominant as in this case, we explore this in more details in paragraph IV D). As the global angle uncertainty constitutes the biggest contribution to the uncertainty of polarization angle for each channel, we expect that the errors on these angles also decrease with the number of frequency

channels in roughly the same way. Here, and in other figures using six priors, the correlation between polarization angles is modified with respect to the one prior case. Indeed, in channels where the relative angles were less determined and hence the correlation was less apparent, the addition of priors reduced the correlation further.

Cosmological likelihood results.—The distribution of cosmological parameters after the generalized component separation with six priors is presented as dashed dark-red curves in Fig. 4. Again, the estimations of both r and β_b are unbiased and with a precision of $\sigma(r) \approx 2 \times 10^{-3}$ and $\sigma(\beta_b) \approx 0.07^\circ$ as mentioned in Table III. The estimation of r is therefore unchanged with respect to the previous case but the estimation of β_b has improved as a consequence of the

improvement of the polarization angle estimation in the first step. Having polarization angle calibration on multiple frequency bands would therefore improve $\sigma(\beta_b)$ without necessarily requiring a large improvement of the calibration precision itself which can be very challenging.

B. Complex foregrounds, $r=0$, $\beta_b=0^\circ$

1. Spatially varying foreground SEDs: “*dls1*”

This foreground model used in the simulated *data* implements spatially varying spectral indices. However, the model we use to describe the data still assumes constant spectral indices. This may potentially lead to bias on cosmological parameters induced by the mismatch between

foreground model and data. For conciseness, we focus on the case with *priors on all polarization angles*. The results obtained with one prior on the 93 GHz channel are detailed in Table III.

Generalized spectral likelihood results.—Figure 6 shows the results of the generalized spectral likelihood sampling. The estimation of polarization angles are not significantly affected by the more complex foregrounds and results are similar to the previous, 6-priors case. The spectral likelihood still manages to estimate effective values of spectral indices, even if in the simulated data they vary from pixel to pixel.

Cosmological likelihood results.—Results are shown in Fig. 7. The estimation of the birefringence angle β_b is not




FIG. 6. Results of the generalized spectral likelihood with *dls1* as input foregrounds model with the calibration priors on all six polarization angles, with a precision of $\sigma_{\alpha_i} = 0.1^\circ$. The dashed purple lines show these Gaussian priors, while the orange contours correspond to the sampling of the generalized spectral likelihood. The gray dotted lines mark the input values.




FIG. 7. Cosmological likelihood sampling after foreground cleaning and systematic effect correction with priors on *all* polarization angles. Results obtained for *d1s1* (respectively, *d7s3*) are shown as orange contours (respectively dashed dark red). The gray dashed lines correspond to the input values.

significantly affected by the complex foregrounds and by the mismatch between data and model. This seems consistent with the unbiased estimate of polarization angles in the previous step, and the foreground leakage to the recovered CMB EB correlation seems under-control. For r the estimation is slightly biased with $r \sim 0.0016$, but it is still 1σ compatible with the input, $r = 0$ value. The relatively small effects of foreground SEDs mismatch is mostly thanks to the small sky fraction observed by SO SATs ($f_{\text{sky}} = 0.1$), their limited frequency coverage, and their large angular scale, $\ell_{\text{min}} = 30$. Spatially varying spectral indices on, e.g., a larger sky patch or with a larger frequency coverage would certainly bias more significantly both parameters.

2. Nonparametric dust model and curved synchrotron: “d7s3”

As described in Sec. III B 1, the “d7” dust model is nonparametric and should therefore pose particular problems to our generalized parametric component separation based on the simplest, pixel-independent scaling relations. The “s3” synchrotron model adds complexity as well with a global curvature term not taken into account in our data model. Nevertheless, we find that polarization angles do not seem to be affected by nonparametric input foregrounds as can be seen in Table III. For the cosmological likelihood, once again the limited sky fraction used by SO SATs helps to alleviate the impact of spatial variations and the

generalized spectral likelihood sampling gives similar results as in the “d1s1” case. The impact on the cosmological likelihood is therefore limited as well, as shown in Fig. 7 as dashed dark-red contours. The error bars are slightly different as described in Table III with a $\sim 5\%$ increase in the upper error bar on r with respect to the “d1s1” case, and a $\sim 1.4\%$ decrease on $\sigma(\beta_b)$, and with no detectable biases.

To mitigate the effects of these complex foregrounds many approaches are possible. For instance, Errard and Stompór [52] proposed to mitigate foreground residuals, including the foreground power spectra retrieved after component separation in the cosmological likelihood covariance matrix. This method has been used in Wolz *et al.* [53] on realistic SO simulations and showed robustness against complex foregrounds. For bigger surveys, spatial variation of spectral indices can be handled using a multipatch approach that divides the sky in patches in which the spectral parameters are independently estimated from one patch to the other [52,54,55]. In power-spectra space, a moment method can be used to address the same issues [56]. Implementing those methods are not treated in this work and are left for future work.

C. Nonzero cosmological parameters:

$$r = 0.01, \beta_b = 0.35^\circ$$

For completeness, we now focus on a case where the two cosmological parameters considered here are nonzero. We also use the input foreground model “d0s0” and priors on all channels. We do not show the results of the generalized spectral likelihood as they are essentially identical to the ones presented in Fig. 5.

Figure 8 displays the cosmological constraints and indicates the possibility of detecting $r = 0.01$ with a $\sim 5\sigma$ precision, consistent with previous forecast [9]. A $\sim 5\sigma$ detection of the value of $\beta_b = 0.35^\circ$ as derived recently from the Planck data [23,24] seems achievable as well.

The results obtained for the other foreground models, “d1s1”, “d7s3”, or with only a single prior are analogous to the corresponding cases with $r = 0$ and $\beta_b = 0^\circ$, and we do not show the likelihood plots here again. In particular, we find that the results on r are becoming progressively more biased, but the biases never exceed 1σ error bars. The results on β_b are consistently unbiased with error bars going from 0.1° to 0.07° for the one and six priors cases, respectively.

D. Cosmological parameters estimation as a function of prior precision

To better understand the dependence of the results on the prior precision and on the number of priors, we perform the previous analyses with several values for the calibration precision. Conversely, this also provides specification for the calibration campaigns given predefined science goals.

We limit ourselves to the case where $r = 0$, $\beta_b = 0^\circ$ and use the simple “d0s0” foregrounds in the input data.




FIG. 8. Cosmological likelihood sampling after foreground cleaning and systematic effect correction with “d0s0” as foreground input and priors on *all* polarization angles. Input cosmological parameters are $r = 0.01$ and $\beta_b = 0.35^\circ$. The gray dashed lines correspond to the input values. This figure is analogous to Fig. 4 but assumes a nonzero birefringence angle.

We allow the prior precision to change in the range $0.001^\circ \leq \sigma_{\alpha_i} \leq 5^\circ$ and consider cases with a single prior on the 93 GHz channel and with priors on all polarization angles for all channels. The dependence of $\sigma(\beta_b)$ on σ_{α_i} is displayed in Fig. 9. The blue points correspond to the single prior case: it is clear that for large values of $\sigma(\alpha_i) \gtrsim 0.05^\circ$ the obtained precision on β_b is determined by the prior precision. For smaller values of $\sigma_{\alpha_i} \lesssim 0.05^\circ$, the precision on β_b saturates and reaches a plateau at $\sigma(\beta_b) \approx 0.045^\circ$.




FIG. 9. Dependence of $\sigma(\beta_b)$ on prior precision for different calibration and noise scenarios as discussed in Sec. IV D.

This plateau is due to the cosmic and noise variances, which dominate the error budget over the prior precision. A similar behavior is observed in the six priors case. However, for large σ_{α_i} , the obtained values if $\sigma(\beta_b)$ are now roughly a factor of $\sqrt{\#\text{prior}}$ lower than in the single prior case. As discussed already earlier in Sec. IV A 2, this is because for large calibration uncertainties the relative errors are fixed internally and the dominant uncertainty is due to the global polarization angle. Each prior corresponds effectively to an independent measurement of the global polarization angle, and the effective error on it therefore decreases with the number of channels. As the error is smaller for large prior uncertainties compared to the one prior case, the dependence starts reaching the plateau somewhat earlier (i.e., for larger values of the prior precision) as its level remains the same whatever is the number of priors.

This result can help us with future calibration requirements and suggests that as long as we are in the prior dominated regime, to achieve a given $\sigma(\beta_b)$, one must either improve the precision of the calibration method, or up to a certain limit depending on the number of channels, calibrate several frequency channels to get the same results. As improving the absolute precision of calibration is quite challenging; multiplying calibration campaigns to other frequency channels seems to be a reasonable option.

Furthermore, to see how both noise and cosmic variance account for the level of the plateau, we performed a *noiseless* analysis with six priors represented by the orange dots in the figure. The cosmic-variance limit reaches $\sigma(\beta_b) \approx 0.026^\circ$. This seems to indicate that for the sky coverage and noise levels of SO SATs the noise accounts for $\sim 40\%$ of the plateau’s amplitude. To improve on the level of the plateau one then needs either to lower the noise, to increase f_{sky} or to perform delensing. However, with a larger sky survey the spatial variability of foreground SEDs will potentially become a bigger issue for the component separation and might bias the estimation of cosmological parameters.

E. Biased priors

Up to now we have only considered cases where the priors, when averaged over an ensemble of calibration procedures, are centered on the true values of miscalibration angles. In this section, we explore how the framework performs with biased priors. These could for instance arise due to a systematic effect that would not vanish even after the averaging over an ensemble of the calibration procedures as assumed previously. In particular, we would like to get some insight on how different biases at different frequency channels would typically impact our conclusions. For this exploratory work, we use a simple foreground “d0s0” with cosmological parameters set at $r = 0$ and $\beta_b = 0^\circ$. All the priors used in this section have a precision of $\sigma_{\alpha_i} = 1^\circ$. We then draw a random set of biases

TABLE IV. The three different cases of input biased prior centers used for each frequency channel. The numbers in parentheses correspond to the value of the bias in degrees.

Input angle [°]	Input one prior [°]	Input six priors, equal biases [°]	Input six priors, different biases [°]
$\alpha_{27} = 1.00$		1.13 (0.13)	1.28 (0.28)
$\alpha_{39} = 1.67$		1.80 (0.13)	0.88 (-0.78)
$\alpha_{93} = 2.33$	2.46 (0.13)	2.46 (0.13)	2.46 (0.13)
$\alpha_{145} = 3.00$		3.13 (0.13)	3.30 (0.30)
$\alpha_{225} = 3.67$		3.80 (0.13)	5.12 (1.46)
$\alpha_{280} = 4.33$		4.46 (0.13)	4.43 (0.09)

for each frequency channel. They are drawn from a Gaussian distribution centered at 0° with a standard deviation of 1° following the prior precision. Hereafter, we explore three cases: first, with only one prior on the 93 GHz channel, which is biased by 0.13° ; a second case with priors on all channels, where they are all biased by 0.13° ; and the third case with priors on all channels, and all biased by different random amounts. Table IV summarizes all these cases.

The results of the generalized ensemble averaged spectral likelihood sampling for the three cases are displayed in Table V, which shows the recovered biases significance (expressed as a ratio of the bias over the corresponding sigma) for the estimated parameters.

In the case of the single prior, the *recovered* miscalibration angles for all the channels have the same overall bias of $\Delta\alpha = 0.13^\circ$. As we are in a case with 1° prior precision the final precision on polarization angles is lower than in previous case and here $\sigma(\alpha_i) \approx 1^\circ$. The bias is in the end of $\sim 0.1\sigma$ (see Table V). In the end, this bias was directly imposed by the calibration on the 93 GHz channel. This is because all these angles are calibrated *relative* to the channel with a prior. We also find that the spectral indices are estimated correctly as their estimates do not depend on

TABLE V. Resulting bias over standard deviation (bias significance) of each studied parameters in the three biased cases explored here.

Δ/σ	One biased prior	Six equally biased priors	Six differently biased priors
α_{27}	0.10	0.29	0.55
α_{39}	0.10	0.30	0.58
α_{93}	0.10	0.30	0.60
α_{145}	0.10	0.30	0.60
α_{225}	0.10	0.30	0.60
α_{280}	0.10	0.30	0.59
β_d	-0.01	0.01	0.01
β_s	-0.01	0.04	-0.02
r	0.14	0.14	0.14
β_b	-0.10	-0.30	-0.60

the coordinate choice given that the scaling laws assumed for the Q and U Stokes parameters are the same. In the cosmological likelihood, this bias affects the estimation of β_b leading to a bias of -0.13° (corresponding to a $\sim -0.1\sigma$ bias). Estimation of r is however not significantly impacted by the bias compared to previous cases. Indeed, its bias was of the order $\sim 10^{-4}$ while its precision was of $\sim 10^{-3}$. This is because we estimate r and β_b jointly and assume no EB correlations in the CMB covariance. This allows us to separate the B-mode signal due to the birefringence or miscalibration from the primordial signal at minimal loss of precision. We note that the method is similar to self-calibration, which however is performed on a foreground-cleaned and hopefully, miscalibration-corrected signal and includes the extra uncertainty due to the marginalization over the birefringence angle. As we discussed above the method is robust for r even in the case where miscalibration is not well corrected for—as it is the case here.

In the second case with six priors and a 0.13° bias on each of them, the estimation of the miscalibration angles are biased by the same value ($\Delta\alpha = 0.13^\circ$). The bias on β_b is -0.13° as well. But with the addition of priors, the overall precision on α_i and β_b improves to $\sigma(\alpha) \approx 0.41^\circ$ as well as for β_b . This results in an increase in biases significance for both polarization and birefringence angle (Table IV). r is not affected here. Finally, in the case where all biases are different, we see that all miscalibration angles estimations are biased with the *same* value of $\sim 0.25^\circ$, which corresponds to the average of all bias values from the right most column of Table IV. Here, the error bars remain constant with respect to the previous case, but the biases increases, resulting in a $\sim 0.6\sigma$ bias on α_i . Consequently, the bias on miscalibration angles is carried to β_b leading to a $\sim -0.25^\circ$ bias or $\sim -0.6\sigma$ (see Table V) and r is again not affected.

We note that as a single prior is sufficient for us to solve the problem, having multiple priors allows for a number of consistency tests, which in actual data analysis practice could shed some light on underlying (albeit unknown) biases. In particular, having priors from different calibration methods (wire grids, drone, observations of Tau-A) would allow for systematic checks and correction between priors.

V. CONCLUSIONS

We propose and demonstrate on simulations a extension of the component separation method allowing to remove the contributions due to Galactic foregrounds and simultaneously accounting on polarization-angle misalignment and allowing for an inclusion of calibration priors. The method generalizes the pixel-based parametric component separation method introduced in [32,40]. The instrumental effects are represented via an instrumental response matrix X incorporated directly in the data model. The calibration priors are included as multiplicative terms to the generalized spectral likelihood and the method propagates statistical and systematic errors due to the data, assumed

models, and the priors all the way to cosmological parameters. We focus specifically on the tensor-to-scalar ratio, r , and birefringence angle, β_b , and we fit simultaneously for these two parameters to the CMB signal as recovered from the data on the initial, component separation step.

We specialize this method to forecast performance of a typical future ground based multifrequency CMB experiments. For this, we employ likelihoods semianalytically averaged over CMB and noise realizations, which permits obtaining statistically meaningful predictions and their uncertainties averaged over the same ensemble.

We use the Simons Observatory small aperture telescopes characteristics as our fiducial experimental setup and assume a single, pixel-independent, miscalibration error for each single frequency map. These angles are parameters of the instrumental response matrix, \mathbf{X} , and are fitted for on the component separation step together with the parameters describing the foregrounds. We consider different foreground models, including those where the assumed foreground model matches the actual foreground signal as used in the simulations, and models where we allow for the mismatch between the two. We then investigate the performance of the method from the perspective of the biases and statistical errors on the cosmological parameters.

We show that the data on its own set strong constraints on relative polarization angles between different single frequency maps. Consequently, a single prior on a polarization angle of one of the single frequency maps is sufficient to allow setting meaningful constraints first on polarization angles for all frequency channels and foreground spectral indices, and later the cosmological parameters. Using multiple calibration priors on different single frequency maps is beneficial in terms of the resulting statistical uncertainties of the recovered polarization angles but also allows for robustness tests of the derived results. In our fiducial study cases, we find that for a single prior with precision $\sigma_{\alpha_3} = 0.1^\circ$ on the polarization angle of the sky map at 93 GHz, the polarization angles for all maps can be derived without any biases and with the precision equal to, for the 93 GHz channel, and only slightly worse than, for all other channels, than the assumed prior precision. We find that there is little impact of the more involved data model employed in this work on the estimation of the foreground indices, and our results for r , $r = 0.0002_{-0.0018}^{+0.0015}$, are in agreement with the SO SAT forecasts with FGBuster [9], which neglect the polarization angle misalignment. Allowing for the foreground model mismatch does not affect significantly statistical errors but may lead to a bias in the estimated values of r . However, in the cases studied here, the biases were never larger than 1σ statistical uncertainties with the most significant bias on r found in the case of the spatially varying foreground model “d1s1”, $r = 0.0016_{-0.0018}^{+0.0016}$. We find that the estimates of r are largely independent of the assumed priors and that we

can set meaningful constraints on r even in their absence. The proposed method can therefore be considered as a self-calibration approach.

Priors are necessary however in order to constrain the birefringence parameter. The foreground model mismatch does not bias the estimates of β_b , and for the case of a single prior with the precision of $\sigma_\alpha = 0.1^\circ$, we get the uncertainty on β_b to be $\simeq 0.1^\circ$. For six priors with the same precision, this improves to $\simeq 0.07^\circ$. In general, the better the priors, the better the final uncertainty on β_b ; however, the latter saturates once the calibration precision gets sufficiently low and the uncertainty on β_b starts being dominated by the signal and noise variance. For the studied instrumental setup, this happens for calibration precision of $\sim 0.07^\circ$ for a single and $\sim 0.1^\circ$ for six calibration priors. Overall, we conclude that the next generation of the CMB polarization experiments, aiming at the precision of their angle calibration of $\sim 0.1^\circ$, should be capable of rejecting or confirming the value of 0.35° suggested by some recent analyses of the Planck data [23,24] with $\sim 3-5\sigma$ significance depending on the number of the priors.

We also find the biases on the birefringence angle arise in the cases when the calibration priors themselves are biased. We find however that this does not affect the estimates of r .

The instrumental model assumed here is clearly idealized. Most importantly, it neglects bandpasses, beams, and gains. These would affect both the actual sky signal but also polarization angle calibration in the way which will depend on their effective spectral dependence. To first order, this will lead to biases on the priors, and therefore, including such effects is a key to any claim about the detection of birefringence. This could also affect the r constraints but mostly via their impact on the foreground residuals. Other effects which could affect the polarization angle, such as smoothly rotating half wave plate or sinuous antennas, are also relevant and should be taken into account. We leave such extensions to future work.

Similarly, unrealistic is the assumption of a single pixel-independent polarization angle per a single frequency map. Indeed, the miscalibration angle should be more of a property of a detector or of a focal plane wafer; this will generally lead to a pixel-domain effective polarization angle on the map level due to the fact that different wafers/detectors typically observe the sky differently. This formalism is easily adaptable to using as an input maps produced for every wafer or detector, each with a specific polarization angle. This may however lead to proliferation of the instrumental degrees of freedom in the spectral likelihood problem with potential effects on the precision of derived constraints. More studies are needed to assess whether this can be successfully controlled.

Complex foregrounds do not seem to significantly affect the estimation of polarization angles; however, the small sky and frequency coverage of the SATs limits the impact of complex foregrounds with respect to simpler ones.

This forecasting method could be adapted to an analysis pipeline applied to existing and future datasets including SO, Planck, LiteBIRD etc. Planck offers more modes, potentially increasing precision on β_b , but it lacks sensitivity. This will be addressed by LiteBIRD, but its resolution will be limited. Robust calibration of polarization angles remains challenging for space missions. However, leveraging Planck's ground-based calibration [57–59] and LiteBIRD's requirements [11] as priors can improve β_b constraints. Future work will explore these applications.

ACKNOWLEDGMENTS

We thank the reviewer for their insightful comments that helped improve the quality of this article. We also thank Clara Vergès, Clément Leloup, Hamza El Bouhargani,

Magdy Morshed, Arianna Rizzieri, and Simon Biquard for useful discussions. This research 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. The authors acknowledge support of the French National Research Agency (Agence National de Recherche) Grants No. ANR BxB (ANR-17-CE31-0022) and No. B3DCMB (ANR-17-CE23-0002). This work is also part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (PI: Josquin Errard, Grant Agreement No. 101044073). Some of the results in this paper have been derived using the HEALPY, NumPy, and PySM packages. Some of the figures in this article have been created using GetDist.

-
- [1] S. M. Carroll, G. B. Field, and R. Jackiw, Limits on a Lorentz- and parity-violating modification of electrodynamics, *Phys. Rev. D* **41**, 1231 (1990).
- [2] S. M. Carroll, Quintessence and the Rest of the World: Suppressing Long-Range Interactions, *Phys. Rev. Lett.* **81**, 3067 (1998).
- [3] F. Finelli and M. Galaverni, Rotation of linear polarization plane and circular polarization from cosmological pseudoscalar fields, *Phys. Rev. D* **79**, 063002 (2009).
- [4] M. A. Fedderke, P. W. Graham, and S. Rajendran, Axion dark matter detection with CMB polarization, *Phys. Rev. D* **100**, 015040 (2019).
- [5] M. Watanabe, S. Kanno, and J. Soda, Imprints of the anisotropic inflation on the cosmic microwave background, *Mon. Not. R. Astron. Soc.* **412**, L83 (2011).
- [6] A. Lue, L. Wang, and M. Kamionkowski, Cosmological Signature of New Parity-Violating Interactions, *Phys. Rev. Lett.* **83**, 1506 (1999).
- [7] S. Saito, K. Ichiki, and A. Taruya, Probing polarization states of primordial gravitational waves with cosmic microwave background anisotropies, *J. Cosmol. Astropart. Phys.* **09** (2007) 002.
- [8] M. H. Abitbol, D. Alonso, S. M. Simon *et al.*, The Simons Observatory: Gain, bandpass and polarization-angle calibration requirements for B-mode searches, *J. Cosmol. Astropart. Phys.* **05** (2021) 032.
- [9] P. Ade, J. Aguirre, Z. Ahmed *et al.*, The Simons Observatory: Science goals and forecasts, *J. Cosmol. Astropart. Phys.* **02** (2019) 056.
- [10] E. Allys, K. Arnold *et al.* (LiteBIRD Collaboration), Probing cosmic inflation with the litebird cosmic microwave background polarization survey, *Prog. Theor. Exp. Phys.* **2023**, 042F01 (2023).
- [11] P. Vielva, E. Martínez-González *et al.* (The LiteBIRD Collaboration), Polarization angle requirements for CMB B-mode experiments. Application to the LiteBIRD satellite, *J. Cosmol. Astropart. Phys.* **04** (2022) 029.
- [12] P. A. R. Ade, Z. Ahmed, M. Amiri *et al.* (BICEP/Keck Collaboration), BICEP/Keck XII: Constraints on axionlike polarization oscillations in the cosmic microwave background, *Phys. Rev. D* **103**, 042002 (2021).
- [13] P. A. R. Ade, Z. Ahmed, M. Amiri *et al.* (BICEP/Keck Collaboration), BICEP/Keck xiv: Improved constraints on axionlike polarization oscillations in the cosmic microwave background, *Phys. Rev. D* **105**, 022006 (2022).
- [14] K. R. Ferguson *et al.* (SPT-3G Collaboration), Searching for axion-like time-dependent cosmic birefringence with SPT-3G, *Phys. Rev. D* **106**, 042011 (2022).
- [15] V. Gluscevic, D. Hanson, M. Kamionkowski, and C. M. Hirata, First CMB constraints on direction-dependent cosmological birefringence from WMAP-7, *Phys. Rev. D* **86**, 103529 (2012).
- [16] P. Ade, Z. Ahmed, R. Aikin *et al.*, BICEP2 / Keck array IX: New bounds on anisotropies of CMB polarization rotation and implications for axionlike particles and primordial magnetic fields, *Phys. Rev. D* **96**, 102003 (2017).
- [17] P. A. Ade, K. Arnold, M. Atlas *et al.*, POLARBEAR constraints on cosmic birefringence and primordial magnetic fields, *Phys. Rev. D* **92**, 123509 (2015).
- [18] T. Namikawa, Y. Guan, O. Darwish *et al.*, Atacama cosmology telescope: Constraints on cosmic birefringence, *Phys. Rev. D* **101**, 083527 (2020).
- [19] F. Bianchini, W. Wu, P. Ade *et al.*, Searching for anisotropic cosmic birefringence with polarization data from SPTpol, *Phys. Rev. D* **102**, 083504 (2020).
- [20] A. Gruppuso, D. Molinari, P. Natoli, and L. Pagano, Planck 2018 constraints on anisotropic birefringence and its cross-correlation with CMB anisotropy, *J. Cosmol. Astropart. Phys.* **11** (2020) 066.

- [21] P. A. R. Ade, Z. Ahmed, M. Amiri *et al.*, BICEP/Keck. XVII. line-of-sight distortion analysis: Estimates of gravitational lensing, anisotropic cosmic birefringence, patchy reionization, and systematic errors, *Astrophys. J.* **949**, 43 (2023).
- [22] Y. Minami, H. Ochi, K. Ichiki, N. Katayama, E. Komatsu, and T. Matsumura, Simultaneous determination of the cosmic birefringence and miscalibrated polarisation angles from CMB experiments, *Prog. Theor. Exp. Phys.* **2019**, 083E02 (2019).
- [23] P. Diego-Palazuelos, J. Eskilt, Y. Minami *et al.*, Cosmic Birefringence from the Planck Data Release 4, *Phys. Rev. Lett.* **128**, 091302 (2022).
- [24] Y. Minami and E. Komatsu, New Extraction of the Cosmic Birefringence from the Planck 2018 Polarization Data, *Phys. Rev. Lett.* **125**, 221301 (2020).
- [25] S. E. Clark, C.-G. Kim, J. C. Hill, and B. S. Hensley, The origin of parity violation in polarized dust emission and implications for cosmic birefringence, *Astrophys. J.* **919**, 53 (2021).
- [26] K. M. Huffenberger, A. Rotti, and D. C. Collins, The power spectra of polarized, dusty filaments, *Astrophys. J.* **899**, 31 (2020).
- [27] Y. Akrami, M. Ashdown, J. Aumont *et al.*, Planck 2018 results: IV. Diffuse component separation, *Astron. Astrophys.* **641**, A4 (2020).
- [28] J. R. Eskilt and E. Komatsu, Improved constraints on cosmic birefringence from the WMAP and Planck cosmic microwave background polarization data, *Phys. Rev. D* **106**, 063503 (2022).
- [29] J. R. Eskilt, Frequency-dependent constraints on cosmic birefringence from the LFI and HFI Planck data release 4, *Astron. Astrophys.* **662**, A10 (2022).
- [30] B. G. Keating, M. Shimon, and A. P. S. Yadav, Self-calibration of cosmic microwave background polarization experiments, *Astrophys. J.* **762**, L23 (2012).
- [31] B. D. Sherwin and T. Namikawa, Cosmic birefringence tomography and calibration-independence with reionization signals in the CMB, *Mon. Not. R. Astron. Soc.* **520**, 3298 (2023).
- [32] R. Stompor, J. Errard, and D. Poletti, Forecasting performance of CMB experiments in the presence of complex foreground contaminations, *Phys. Rev. D* **94**, 083526 (2016).
- [33] C. Vergès, J. Errard, and R. Stompor, Framework for analysis of next generation, polarized CMB data sets in the presence of galactic foregrounds and systematic effects, *Phys. Rev. D* **103**, 063507 (2021).
- [34] M. Tegmark, A. N. Taylor, and A. F. Heavens, Karhunen-Loeve Eigenvalue problems in cosmology: How should we tackle large data sets?, *Astrophys. J.* **480**, 22 (1997).
- [35] A. Zonca, L. Singer, D. Lenz, M. Reinecke, C. Rosset, E. Hivon, and K. Gorski, HEALPY: Equal area pixelization and spherical harmonics transforms for data on the sphere in PYTHON, *J. Open Source Software* **4**, 1298 (2019).
- [36] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere, *Astrophys. J.* **622**, 759 (2005).
- [37] N. Aghanim, Y. Akrami *et al.* (Planck Collaboration), Planck 2018 results: VI. Cosmological parameters, *Astron. Astrophys.* **641**, A6 (2020).
- [38] B. Thorne, J. Dunkley, D. Alonso, and S. Næss, The PYTHON sky model: Software for simulating the galactic microwave sky, *Mon. Not. R. Astron. Soc.* **469**, 2821 (2017).
- [39] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, EMCEE: The MCMC hammer, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
- [40] R. Stompor, S. Leach, F. Stivoli *et al.*, Maximum likelihood algorithm for parametric component separation in cosmic microwave background experiments, *Mon. Not. R. Astron. Soc.* **392**, 216 (2008).
- [41] K. Wolz, S. Azzoni, C. Hervias-Caimapo *et al.*, The simons observatory: Pipeline comparison and validation for large-scale B-modes, [arXiv:2302.04276](https://arxiv.org/abs/2302.04276).
- [42] J. W. Fowler, M. D. Niemack, S. R. Dicker *et al.*, Optical design of the Atacama cosmology telescope and the millimeter bolometric array camera, *Appl. Opt.* **46**, 3444 (2007).
- [43] R. J. Thornton, P. A. R. Ade, S. Aiola *et al.*, The Atacama Cosmology Telescope: The polarization-sensitive ACTPol Instrument, *Astrophys. J. Suppl. Ser.* **227**, 21 (2016).
- [44] T. Essinger-Hileman, A. Ali, M. Amiri *et al.*, CLASS: The cosmology large angular scale surveyor, in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, edited by W. S. Holland and J. Zmuidzinas (SPIE, 2014).
- [45] F. Nati, M. J. Devlin, M. Gerbino, B. R. Johnson, B. Keating, L. Pagano, and G. Teply, POLOCALC: A novel method to measure the absolute polarization orientation of the cosmic microwave background, *J. Astron. Instrum.* **06**, 1740008 (2017).
- [46] S. A. Bryan, G. P. Teply, S. M. Simon *et al.*, Development of calibration strategies for the simons observatory, in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, edited by J. Zmuidzinas and J.-R. Gao (SPIE, 2018).
- [47] J. Aumont, J. F. Macías-Pérez, A. Ritacco, N. Ponthieu, and A. Mangilli, Absolute calibration of the polarisation angle for future CMB B-mode experiments from current and future measurements of the Crab nebula, *Astron. Astrophys.* **634**, A100 (2020).
- [48] A. Ritacco, J. F. Macías-Pérez, N. Ponthieu *et al.*, NIKA 150 GHz polarization observations of the Crab nebula and its spectral energy distribution, *Astron. Astrophys.* **616**, A35 (2018).
- [49] R. Adam, P. A. Ade, N. Aghanim *et al.*, Planck 2015 results: X. Diffuse component separation: Foreground maps, *Astron. Astrophys.* **594**, A10 (2016).
- [50] B. T. Draine and B. Hensley, Magnetic nanoparticles in the interstellar medium: Emission spectrum and polarization, *Astrophys. J.* **765**, 159 (2013).
- [51] B. Hensley, On the nature of interstellar grains, Ph.D. thesis, Princeton University, 2015.
- [52] J. Errard and R. Stompor, Characterizing bias on large scale CMB B-modes after galactic foregrounds cleaning, *Phys. Rev. D* **99**, 043529 (2019).

- [53] K. Wolz, S. Azzoni, C. Hervias-Caimapo *et al.*, The Simons Observatory: Pipeline comparison and validation for large-scale B-modes, [arXiv:2302.04276](https://arxiv.org/abs/2302.04276).
- [54] G. Puglisi, G. Mihaylov, G. V. Panopoulou, D. Poletti, J. Errard, P. A. Puglisi, and G. Vianello, Improved galactic foreground removal for B-mode detection with clustering methods, *Mon. Not. R. Astron. Soc.* **511**, 2052 (2022).
- [55] E. Allys, Arnold *et al.* (LiteBIRD Collaboration), Probing cosmic inflation with the LiteBIRD cosmic microwave background polarization survey, *Prog. Theor. Exp. Phys.* **2023**, 042F01 (2022).
- [56] S. Azzoni, M. Abitbol, D. Alonso, A. Gough, N. Katayama, and T. Matsumura, A minimal power-spectrum-based moment expansion for CMB B-mode searches, *J. Cosmol. Astropart. Phys.* **05** (2021) 047.
- [57] C. Rosset, M. Tristram, N. Ponthieu *et al.*, Planck pre-launch status: High frequency instrument polarization calibration, *Astron. Astrophys.* **520**, A13 (2010).
- [58] A. Gruppuso, M. Gerbino, P. Natoli, L. Pagano, N. Mandolesi, A. Melchiorri, and D. Molinari, Constraints on cosmological birefringence from Planck and BICEP2/Keck data, *J. Cosmol. Astropart. Phys.* **06** (2016) 001.
- [59] N. Aghanim, M. Ashdown *et al.* (Planck Collaboration), Planck intermediate results—XLIX. parity-violation constraints from polarization data, *Astron. Astrophys.* **596**, A110 (2016).