



The Simons Observatory: Quantifying the impact of beam chromaticity on large-scale B -mode science

Nadia Dachlythra ^a, Kevin Wolz ^b, Susanna Azzoni ^c, David Alonso ^b,
Adriaan J. Duivenvoorden ^d, Alexandre E. Adler ^{e,f}, Jon E. Gudmundsson ^{g,h},
Carlo Baccigalupi ^{i,j,k}, Alessandro Carones ⁱ, Gabriele Coppi ^{a,l},
Samuel Day-Weiss ^c, Josquin Errard ^m, Nicholas Galitzki ^{n,o}, Martina Gerbino ^p,
Remington G. Gerras ^q, Carlos Hervias-Caimapo ^r, Selim C. Hotinli ^s,
Federico Nati ^{a,l}, Bruce Partridge ^t, Yoshinori Sueno ^c and Edward J. Wollack ^u

^aDepartment of Physics, University of Milano-Bicocca, Piazza della Scienza 3, 20126, Milano, Italy

^bDepartment of Physics, University of Oxford,
Denys Wilkinson Building, Keble Road, Oxford, OX1 3RH, U.K.

^cJoseph Henry Laboratories of Physics, Jadwin Hall, Princeton University,
Princeton, NJ 08544, U.S.A.

^dMax-Planck-Institut für Astrophysik, Karl-Schwarzschild Str. 1, 85741 Garching, Germany

^eDepartment of Physics, University of California, Berkeley,
366 LeConte Hall Berkeley, CA 94720, U.S.A.

^fPhysics Division, Lawrence Berkeley National Laboratory,
1 Cyclotron Road, Berkeley, CA 94720, U.S.A.

^gScience Institute, University of Iceland, 107 Reykjavik, Iceland

^hThe Oskar Klein Centre, Department of Physics, Stockholm University,
AlbaNova, SE-10691 Stockholm, Sweden

ⁱThe International School for Advanced Studies (SISSA), via Bonomea 265, I-34136 Trieste, Italy

^jThe National Institute for Nuclear Physics (INFN), via Valerio 2, I-34127, Trieste, Italy

^kThe Institute for Fundamental Physics of the Universe (IFPU),
Via Beirut 2, I-34151, Trieste, Italy

^lIstituto Nazionale di Fisica Nucleare, INFN, Sezione Milano-Bicocca,
Piazza della Scienza 3, 20126, Milano, Italy

^mUniversité Paris Cité, CNRS, Astroparticule et Cosmologie, F-75013 Paris, France

ⁿDepartment of Physics, University of Texas at Austin, Austin, TX, 78712, U.S.A.

^oWeinberg Institute for Theoretical Physics, Texas Center for Cosmology and Astroparticle Physics,
Austin, TX 78712, U.S.A.

^pINFN Sezione di Ferrara, Via Giuseppe Saragat, 1, 44121, Ferrara

^qDepartment of Physics, University of Southern California, Los Angeles, CA 90089, U.S.A.

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

^s*Perimeter Institute for Theoretical Physics, 31 Caroline St N, Waterloo, ON N2L 2Y5, Canada*

^t*Department of Physics and Astronomy, Haverford College,
370 Lancaster Ave, Haverford PA 19041 U.S.A.*

^u*NASA Goddard Space Flight Center, 8800 Greenbelt Road, Greenbelt, MD 20771, U.S.A.*

E-mail: konstantina.dachlythra@unimib.it, kevin.wolz@physics.ox.ac.uk,
sa5705@princeton.edu, David.Alonso@physics.ox.ac.uk,
adriaand@mpa-garching.mpg.de, aadler@lbl.gov, jegudmunds@hi.is,
bacci@sissa.it, acarones@sissa.it, gabriele.coppi@unimib.it,
dayweiss@princeton.edu, josquin@apc.in2p3.fr,
nicholas.galitzki@austin.utexas.edu, marti.gerbino@gmail.com,
gerras@usc.edu, carlos.hervias@uc.cl, selimcanhotinli@gmail.com,
federico.nati@unimib.it, bpartrid@haverford.edu, ys5857@princeton.edu,
edward.j.wollack@nasa.gov

ABSTRACT: The Simons Observatory (SO) Small Aperture Telescopes (SATs) will observe the Cosmic Microwave Background (CMB) temperature and polarization at six frequency bands. Within these bands, the angular response of the telescope (beam) is convolved with the instrument’s spectral response (commonly called bandpass) and the signal from the sky, which leads to the band-averaged telescope beam response, which is sampled and digitized. The spectral properties of the band-averaged beam depend on the natural variation of the beam within the band, referred to as beam chromaticity. In this paper, we quantify the impact of the interplay of beam chromaticity and intrinsic frequency scaling from the various components that dominate the polarized sky emission on the tensor-to-scalar ratio, r , and foreground parameters. We do so by employing a parametric power-spectrum-based foreground component separation algorithm, namely **BBPower**, to which we provide beam-convolved time domain simulations performed with the **beamconv** software while assuming an idealized version of the SO SAT optics. We find a small, 0.02σ , bias on r , due to beam chromaticity, which seems to mostly impact the dust spatial parameters, causing a maximum 0.77σ bias on the dust B -mode spectra amplitude, A_d , when employing Gaussian foreground simulations. However, we find all parameter biases to be smaller than 1σ at all times, independently of the foreground model. This includes the case where we introduce additional uncertainty on the bandpass shape, which accounts for approximately half of the total allowed gain uncertainty, as estimated in previous work for the SO SATs.

KEYWORDS: CMBR detectors, CMBR experiments, cosmological parameters from CMBR, gravitational waves and CMBR polarization

ARXIV EPRINT: [2503.01791](https://arxiv.org/abs/2503.01791)

Contents

1	Introduction	1
2	Spectral dependence of beam response	3
3	Methods	7
3.1	Sky model	7
3.2	Beam-convolved simulations	8
3.3	B -mode analysis pipeline	10
4	Results	12
4.1	Bias from beam systematics	12
4.2	Chromatic beams	15
4.3	Impact of bandpass uncertainty	18
5	Conclusion and discussion	20
A	Bias from the simulated beamconv approach	23
B	Cross-polarization impact on EB correlation	24
C	Dependence on the detector location	25
D	Chromaticity bias for different beam spectral response	27

1 Introduction

Measurements of the Cosmic Microwave Background (CMB) polarization are at the forefront of cosmological research, as they offer insights into the origins and evolution of the Universe. Of particular interest is the measurement of the CMB B -mode polarization anisotropies (the parity-odd component of the CMB polarization) as a probe of the very early Universe and a strong test of the inflationary paradigm. Cosmic inflation is the current leading theory explaining the initial density perturbations, and the Universe’s large-scale homogeneity and flatness. It postulates a brief period of exponential expansion of an initially small patch of the infant Universe into its current observable size. The primordial quantum fluctuations, stretched by this expansion to cosmological scales, leave traces in the spacetime metric in the form of both scalar and tensor perturbations. The latter, also referred to as gravitational waves, are the only primordial source of CMB B -mode fluctuations, with scalar modes only sourcing the parity-even E -modes. By measuring both components of the CMB polarization, we can quantify the amplitude of primordial gravitational waves generated by inflation, and thus probe the high-energy physical processes that caused it. This amplitude is commonly parametrized in terms of the ratio between the amplitudes of the power spectrum of primordial tensor and scalar perturbations, r [1–3].

Cosmological B -modes are significantly smaller than the E -mode perturbations. In addition, on the large, degree-sized scales where the sensitivity to the primordial signal is highest, the polarized sky is dominated by emission from Galactic foregrounds [4–6]. Thus, the accuracy with which we can measure r is tightly linked to our ability to properly characterize these foregrounds and, crucially, the impact of all instrumental effects on the measured signal. To cleanly separate the cosmological signal from any other contaminants, the CMB community has developed a large variety of component separation methods (e.g. [6–10]) that rely mostly on the frequency-dependent Spectral Energy Distributions (SEDs) of the foregrounds to distinguish them from the CMB’s nearly perfect blackbody spectrum [11]. However, the performance of any of these methodologies depends strongly on the impact of instrumental systematics on the data and on our ability to mitigate or model them. For B -mode experiments, great care must be taken when modeling the systematic effects related to the telescope’s optics. In particular, improper modeling of the telescope’s spatial response, namely the telescope’s Point-Spread-Function (PSF), commonly referred to as the telescope’s beam, can cause leakage of the stronger CMB temperature signal into the fainter polarization data, as well as mixing between the E - and B -modes [12].

An important aspect of beam modeling is linked to the beam’s intrinsic frequency dependence. As CMB telescopes typically observe radiation over wide frequency bands instead of monochromatic frequency channels, it is natural to expect some variation in the beam properties within each band. The impact of this beam ‘chromaticity’ is further complicated when the beam is convolved with frequency-dependent sky components with differing spectral signatures. It is worth noting that, due to these differing spectral signatures, the chromaticity effects will vary slightly for each component. Further complications arise when additional beam non-idealities are present. Representative examples include pronounced beam power at large angles away from the beam center (sidelobes), non-negligible polarization in the direction orthogonal to the detector’s polarization direction (cross-polarization), as well as beam asymmetry [12–14]. For sufficiently complex beams and foregrounds, if these effects are not taken into account in the analysis, they may give rise to significant biases in the inferred model parameters. Therefore, these effects must be quantified to determine the most reliable estimate of the instrumental beam to employ in the component separation analysis [15, 16].

Although these beam-related biases are typically sub-percent at the power spectrum level, they can still threaten the high-accuracy requirements of current and next-generation CMB missions [17] such as the Simons Observatory (SO) [18]. By observing the polarized microwave sky from the Chilean Atacama desert in six frequency bands centered between 27 and 280 GHz, the SO Small Aperture Telescopes (SATs) [19] aim to constrain r with a statistical uncertainty of $\sigma(r) \simeq 0.003$ or better. Given the broad observing bands of the SATs, with $\sim 25\%$ relative fractional bandwidths, it is crucial to model the impact of the anticipated beam frequency dependence on the large-scale B -mode spectra. We do so in this paper, starting from time-domain simulations convolved with chromatic beams that we then analyse using a power-spectrum-based component separation pipeline. An estimation of the corresponding impact on scientific objectives relevant to the SO Large Aperture Telescope (LAT) [20, 21] like the effective number of relativistic species, N_{eff} , and amplitude of unresolved radio sources in temperature, a_s , is presented in [22]. That work assumes a LAT-

like experiment and Gaussian beams with perturbed Full-Width-Half-Maximum (FWHM) values, which are applied to the simulated data at the power spectrum level.

The paper is structured as follows: section 2 introduces the individual components that form the telescope’s band-averaged beam and describes the specifics of the beam simulations. Section 3 elaborates on the sky model assumptions and methods employed for the beam-convolved simulations, the power spectra estimation, and the foreground component separation algorithm. The results are presented in section 4, including the impact of beam systematics of non-frequency-dependent beams, the inclusion of the beam chromaticity effect as a function of foreground complexity, and any further biases caused by the integration of bandpass uncertainty in the model. Finally, section 5 summarizes the paper and outlines potential directions for future studies.

2 Spectral dependence of beam response

The SO SATs observe the microwave sky at frequency bands of approximately 25% width around the center frequencies. Assuming a detector on the focal plane characterized by an instrumental bandpass, $\tau(\nu)$, and beam response, $B(\theta, \phi, \nu)$, at frequency ν , we can derive the band-averaged beam of the detector as follows:

$$B_S(\theta, \phi) = \frac{\int \tau(\nu) B(\theta, \phi, \nu) S(\nu) d\nu}{\int \tau(\nu) S(\nu) d\nu}. \quad (2.1)$$

This formula assumes integration only over frequency and incorporates the SED of the observed sky component denoted by $S(\nu)$ expressed in antenna temperature units. The beam is described in terms of both the polar angle, θ , and the azimuthal angle, ϕ , allowing for beam asymmetry by introducing a dependency on the position angle in the time-ordered data. Note that, while the source emission will not be spatially constant, we avoid adding its orientation dependence in this equation, which describes the synthesis of the band-averaged beam. The spatial dependence of $S(\nu)$ is taken into account properly in the sky simulations.

To generate the SAT Physical Optics (PO) beam simulations, we assume the same optical configuration (three-lens refracting telescope) as for the work described in [23] and employ the TICRA TOOLS¹ software. Specifically, we produce monochromatic far-field beam maps for four frequency bands roughly centered at 93, 145, 225, and 280 GHz, respectively. Each of these Mid- and Ultra-High-Frequency (MF and UHF) bands has 25% relative width around its center frequency, and to each band correspond five monochromatic beam maps evenly spread across the full frequency range of the band. Note that we have two MF and two UHF bands as the SATs employ dichroic detectors. For the nominal case, we assume a detector located at the center of the focal plane. We then integrate across the full band in each case, employing the original or perturbed versions of the simulated bandpasses presented in [24]. The beams of the Low-Frequency (LF) bands, which are roughly centered at 27 and 39 GHz, are approximated using Gaussian curves of FWHM equal to 91’ and 63’, respectively. This approach is sufficient at this stage since the SO SATs will dedicate only a small part of their full scanning time to LF observations. Furthermore, we ensure in appendix A that the partial treatment of only the MF and UHF bands as chromatic PO beams does not impact

¹TICRA, Landemærket 29, Copenhagen, Denmark (<https://www.ticra.com>).

the analysis results. We do, however, plan to generate PO simulations for the LF beams and integrate them into the pipeline in future work.

The resulting beam models include both asymmetric and cross-polar components, as well as faint far sidelobes. Note, however, that these beam models are still the product of ideal simulations of the SO SAT optics.² In reality, we may expect more pronounced beam non-idealities arising from the optical setup which could impact the cosmological analysis. Appendices B and C discuss how the results could vary if we scale up the cross-polarization and ellipticity of the simulated beams. In appendix B, we present the EB spectra of a single sky realization convolved with only co-polar beams, the nominal beams including their simulated cross-polar components, and a perturbed version of the latter, where we dramatically increased the cross-polarization amplitude by one order of magnitude. We find the resulting EB correlation increasing non-negligibly only in the last case, and for multipoles $\ell \gtrsim 150$, while the degree scales are not largely impacted. To increase the beam ellipticity, we simulate one pixel at the edge of the focal plane, probing the asymmetrical illumination of the aperture. We compare the beam profiles for a center and edge pixel in appendix C, and find a minor bias from the increased ellipticity beam on the fitted parameters.

The logarithmic profiles of the monochromatic co-polar beams for the MF and UHF bands are shown in figure 1. From the figure, we observe the profiles ‘contracting’ with increasing frequency within each band, as expected for a diffraction limited beam. The asymptotic yellow curve in the bottom right panel, representing the highest frequency of the 280 GHz band, highlights the challenge posed by the numerical precision of the PO simulations. Since the computational costs of the simulations scale as frequency squared, the accuracy with which the simulation converges on a solution is also frequency-dependent. At the limits of the machine’s precision, very low beam power will, therefore, be approximated to be at the convergence threshold, creating a constant-looking sidelobe at large angles.

The smooth in-band variation of the beam reduces the number of sub-frequencies needed to adequately capture its chromaticity. The relative variation of the main beam between any pair of adjacent sub-frequencies in a band is roughly $\sim 6\%$. This holds true for all four frequency bands, independently of how far apart the adjacent simulations are in frequency space, enabling us to probe the chromaticity effect in all bands with the same number of sub-frequency simulations. In this way, we ensure we do not favor some frequency bands with respect to others. As a parametric component separation method like `BBPower` is expected to absorb any beam chromaticity bias in the parameters of the most challenging sky component to model, changing the number of sub-frequency samples consistently across all bands is unlikely to significantly affect the results. This has been confirmed by performing the chromaticity analysis only using the center and edge frequency beam simulations for each band (three sub-frequencies instead of five). Whether to employ a different number of sub-frequencies between bands to sufficiently characterize the CMB and foreground spectra separate from beam effects, is a valid question but one that extends beyond the scope of this paper. We therefore choose five sub-frequency simulations for all bands, as illustrated in figure 1. Note that, while these will be the beam models used throughout the paper, we

²More realistic optics simulations could, for example, include non-ideal baffling elements or reflective Half-Wave-Plates (HWPs), like the effects discussed in [25].

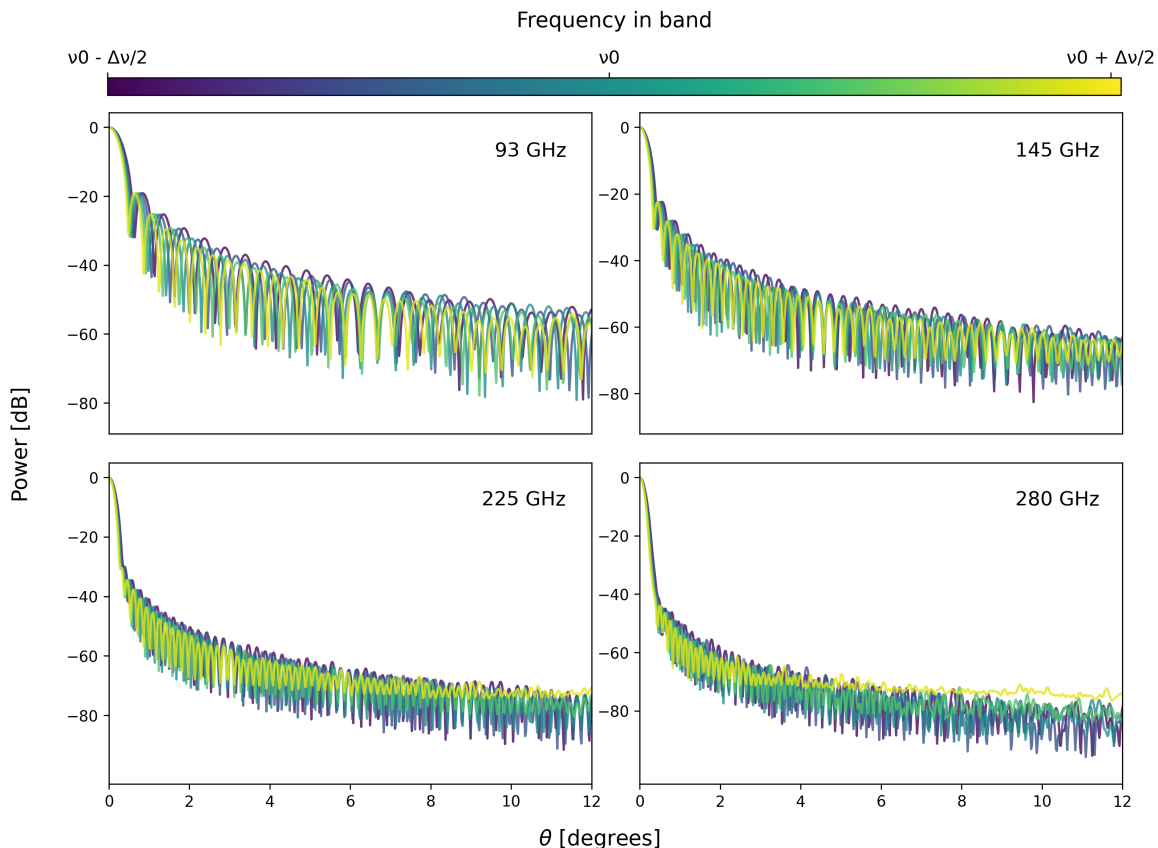


Figure 1. Logarithmic profiles of five monochromatic co-polar beam maps generated for each of the 93 (top left), 145 (top right), 225 (bottom left), and 280 GHz bands (bottom right), using TICRA TOOLS software. Each set of (five) maps is produced at frequencies of uniform spacing across the full frequency range of the corresponding band. All bands are assumed to have 25% fractional width, and all beam maps are generated for a detector on the center of the focal plane of a simulated three-lens refracting telescope.

also estimate the best-fit forecast parameters in the case where we have underestimated the chromaticity of the beam by a smaller and larger value of 10% and 25% change, respectively, within each band (see appendix D).

Figure 2 illustrates the frequency-dependent beam variation in the harmonic domain utilizing the monochromatic beams of the 93 GHz band (MF1), shown in the top left panel of the previous figure, as a representative example. The plots show the harmonic transforms for the first ten azimuthal beam modes, $m \leq 9$, spanning multipoles in the range $\ell = 30 - 300$. This is the ℓ range we will be employing throughout this paper for consistency with [26]. For the symmetric part of the beam ($m=0$), increasing the frequency widens the beam transfer function. The impact on the rest of the modes, nevertheless, is not as straightforward to capture. To decide on the number of beam modes we employ in the analysis, we refer to the expected calibration accuracy for the MF and UHF bands of the SO SATs. The beam of the 93 GHz band that is shown in this example is expected to be calibrated within a $\sim 1.2\%$ accuracy for multipoles spanning $\ell = 30 - 700$ and better than 0.6% for $\ell = 50 - 200$ (see

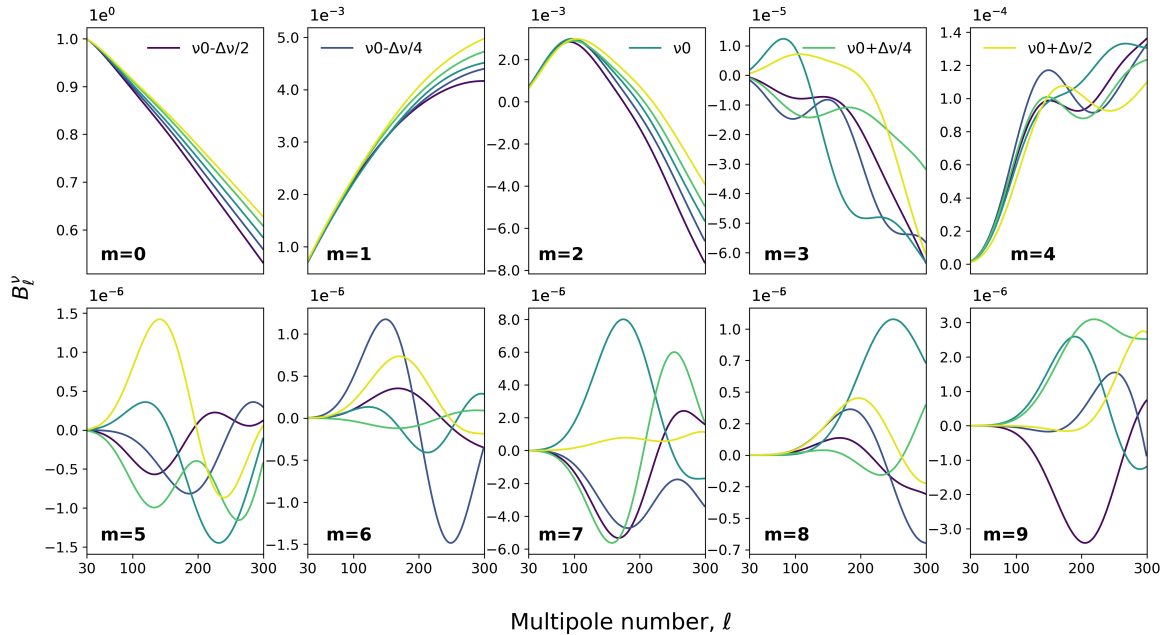


Figure 2. The harmonic transforms of the first ten azimuthal modes of five monochromatic beam simulations sampled uniformly across a 25%-relative-width frequency band centered on 93 GHz. The beam transfer functions have been truncated to a multipole range spanning $\ell=30\text{--}300$ and normalized with respect to the peak amplitude of the symmetric mode.

table 3 of [23]). In principle, this means we can ignore any modes that are smaller than three orders of magnitude with respect to the main beam. However, we choose to allow for a margin of error and include all azimuthal modes that contribute at least 10^{-4} relative to the $m=0$ beam. These correspond to the first five azimuthal modes for the MF1 beam. We ensure this approach is consistent across all four frequency bands.

The combined frequency scaling of the beam and the sky components results in a net frequency-dependent scaling of the band-averaged beams, as shown in equation 2.1. That depends on the observation and is generally different from simply averaging the beam over the bandpass. To illustrate the interplay between the intrinsic beam chromaticity and the frequency dependence of different sky components, we compute the band-averaged beam transfer function by assuming a source SED that corresponds either to pure Galactic dust, synchrotron radiation, CMB or planets (commonly used as beam calibrators). The bottom panel of figure 3 presents the percentage ratio of the frequency-scaled transfer function, B_ℓ^{fsc} , relative to the nominal band-averaged transfer function, B_ℓ , which assumes a constant SED. Both cases are produced employing the first five azimuthal beam modes, as displayed in figure 2. These modes are also shown in map space for the constant SED beam in the top panel of the same figure. In each case, the sky component’s SED is the one also used in the employed component separation pipeline (see section 3.1). Note that the TICRA TOOLS simulations employed in this work only include the beam intrinsic frequency dependence, corresponding to the case where the calibrated beam which will be used in the analysis has already been corrected for the SED of the calibration source (see section 4.5 of [27]).

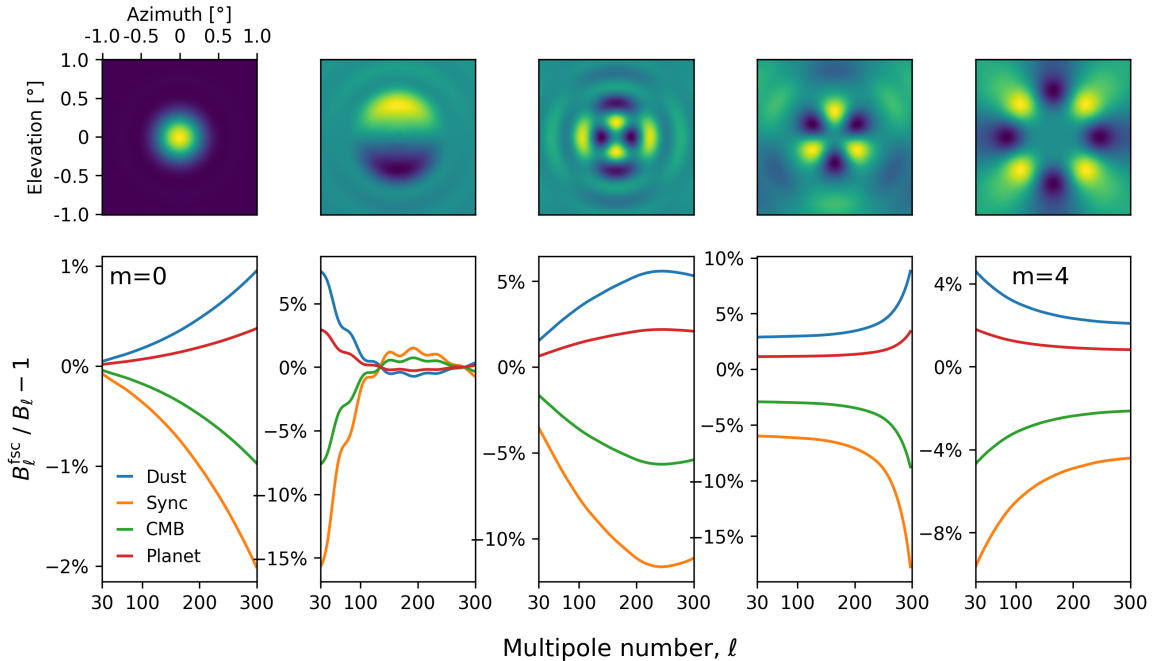


Figure 3. Top: the first five azimuthal modes of the band-averaged 93 GHz beam shown in map space. Bottom: the percent relative error between the frequency-scaled, B_ℓ^{fsc} , and constant-SED beam harmonic transforms, B_ℓ , per mode. The different curves represent the cases where the frequency scaling matches the dust (blue), synchrotron (orange), CMB (green) and planet (red) SEDs. The band-averaged beams with and without frequency scaling are estimated from the monochromatic beam modes of $m = 0 \dots 4$ shown in the first row of figure 2.

From the figure, we observe the symmetric beam component varying as much $\sim 2\%$ at the highest multipoles, whereas the frequency-scaled asymmetric modes show deviations from the corresponding azimuthal mode’s constant-SED beam of up to $\sim 10\% - 15\%$ (the non-frequency scaled $m=1-4$ modes remain, however, three to five orders of magnitude smaller than the corresponding $m=0$ beam).

3 Methods

To assess the impact that beam chromaticity has on the cosmological analysis, we produce beam-convolved simulations in the time domain, bin them into maps, and estimate their pseudo- C_ℓ , B -mode power spectra, which we then utilize to derive the best-fit values for the tensor-to-scalar ratio, r , and foreground parameters. We do so for both chromatic and achromatic beams and quote the bias between the two cases on the forecast parameters in terms of each parameter’s expected uncertainty in the achromatic beams scenario.

3.1 Sky model

The assumed sky model is comprised of CMB, Galactic dust, and synchrotron emission. To simulate CMB sky maps, we draw Gaussian random realizations of the CMB power spectra assuming the *Planck* best-fit Λ CDM parameters [28] and no tensor fluctuations ($r = 0$).

In our fiducial analysis, we also simulate dust and synchrotron as (uncorrelated) Gaussian realizations of their corresponding EE and BB spectra as measured by the *Planck* and WMAP experiments [4, 6, 29] (see section 3.3 of [26]). We do, however, repeat part of our analysis using realistic, non-Gaussian foreground realizations (see section 4.2.2), as this is particularly important for beams carrying pronounced far sidelobes that may pick up the signal from strongly emitting Galactic sources.

The sky SEDs are given by a blackbody law for the CMB, a modified blackbody law (at fixed temperature $T_d = 20$ K) for thermal dust, and a power law for synchrotron. The free parameters in the sky model, which will be later fitted for, are the tensor-to-scalar ratio, r , the amplitude of lensing B -modes, A_{lens} , the spectral indices for dust and synchrotron, β_d and β_s , the amplitude and spectral tilt of the dust and synchrotron B -mode power spectra, denoted as A_d , A_s , α_d and α_s , and the dust-synchrotron correlation parameter, ε_{ds} . The CMB power spectrum employed for the Gaussian realizations is computed as:

$$C_\ell^c = A_{\text{lens}} C_\ell^{\text{lens}} + r C_\ell^{\text{tens}}. \quad (3.1)$$

The templates employed for C_ℓ^{lens} and C_ℓ^{tens} correspond to lensing-only B -modes and tensor-only B -modes of amplitude $r = 1$, respectively, and are precomputed by CAMB [30]. The amplitudes and scaling factors for Galactic dust and synchrotron are linked through a power law that describes the input spectra for these two components:

$$C_\ell^{d/s} = \frac{A_{d/s}}{\ell(\ell+1)/(2\pi)} \left(\frac{\ell}{\ell_0} \right)^{\alpha_{d/s}}, \quad (3.2)$$

with $\ell_0 = 80$.

It is worth noting that the foreground models used here are relatively simple, particularly since they assume foreground sources that are perfectly correlated across frequencies. Frequency decorrelation, which could be sourced by spatially varying foreground spectra, is a major source of contamination for primordial B-mode searchers if not accounted for. We ignore this effect in this work by design, as it allows us to isolate the effective frequency decorrelation induced by beam chromaticity. As we shall see, the effects of beam chromaticity are relatively small and therefore would not modify the resulting constraints on r in more complex foreground scenarios. The impact of foreground frequency decorrelation and techniques to account for it are further discussed in [26, 31, 32].

3.2 Beam-convolved simulations

Random realizations of the different sky components are generated from a known model of their power spectra, as described above.³ These are produced by employing the `synfast` function of the `HEALPix`⁴ package [33, 34], which is used throughout the paper and convolution code. We generate maps of NSIDE=256 at five single frequencies uniformly sampled across a 25% bandwidth around the center frequency of each MF and UHF band. These single-frequency sky models are provided as input to `beamconv`⁵ [12] along with the corresponding simulated

³The code used for the sky simulations is available at <https://github.com/susannaaz/BBSims>.

⁴<http://healpix.sourceforge.net>.

⁵<https://github.com/AdriJD/beamconv>.

SAT beams for the same frequencies. The software is equipped with a ground-based scanning function, which we employ to approximate the scan strategy of the SATs by simulating a telescope located in the Atacama Desert in Chile, performing observations informed by the nominal SAT scan strategy (see section 2.3 of [35]). We use a wide Field-Of-View (FOV) of 35° matching the corresponding size of the actual telescope, populate the focal plane with two hundred detector pairs that are spread uniformly across a square grid, and simulate year-long observations sampling the sky at 50 Hz. While there is dedicated software for SO time-domain simulations, namely the Time-Ordered Astrophysics Scalable Tools (TOAST⁶) and `sotodlib`⁷ packages, we choose to employ `beamconv` for this study due to the latter’s reduced computational cost and the fact that no noise or atmospheric simulations are needed. At this stage, we choose not to simulate the SAT Half-Wave-Plate (HWP) since our main aim is to assess the impact of beam chromaticity independently of the systematics arising from the beam and HWP coupling. For HWPs comprised of three birefringent layers, as is the case for one of the SO SATs, the coupling between the beam chromaticity and the frequency-dependent angle offset associated with this type of HWP could introduce additional challenges [36].

The sky templates are produced only for the Stokes Q and U components while a zero temperature map is provided to `beamconv`. This is done to exclude any contributions from T -to- P leakage in the analysis and for consistency with the work of [26]. After the input polarization maps are beam-convolved in the time domain, they are transformed back into maps using a ‘naive’ mapmaker that simply bins the time-ordered data into pixels given the pointing. An example of a monochromatic map at 280 GHz produced this way is shown in Equatorial coordinates in the left panel of figure 4. In total, for each sky realization, we obtain six maps corresponding to the six SO frequency bands; four of them constructed using PO chromatic beams (MF, UHF) and two using achromatic Gaussian beams (LF). The band-averaged map of each of the MF and UHF bands is the result of applying the (nominal or perturbed) bandpass to the corresponding (five) monochromatic maps. The right panel of the same figure shows the apodized mask that is applied to the beam-convolved maps and is created from the detector hits of the simulated scan strategy. This mask is fairly similar to the SO SAT mask defined in [35] and used in [26]. However, the latter was an early estimate that did not employ the scan strategy, and, thus, we choose to use the mask that is informed from the hits map produced in `beamconv`. Masking the sky results in mode-mixing, which we need to estimate and correct for when reconstructing the angular spectra. Assuming no noise contributions to the data, the pseudo- C_ℓ power spectra, \tilde{C}_ℓ , can be connected with true spectra, C_ℓ , as follows [37, 38]:

$$\langle \tilde{C}_\ell \rangle = \sum_{\ell'} M_{\ell\ell'} B_{\ell'}^2 \langle C_{\ell'} \rangle. \quad (3.3)$$

The matrix $M_{\ell\ell'}$ is the mode-mixing kernel, capturing the statistical coupling between different harmonic-space modes caused by the sky mask while $B_{\ell'}$ represents the beam transfer function. We make use of the pseudo- C_ℓ power spectrum estimator as implemented in `NaMaster`⁸ [38]. In particular, we make use of B -mode purification [39], which projects out the contribution from

⁶<https://github.com/hpc4cmb/toast>.

⁷<https://github.com/simonsobs/sotodlib>.

⁸<https://github.com/LSSTDESC/NaMaster>.

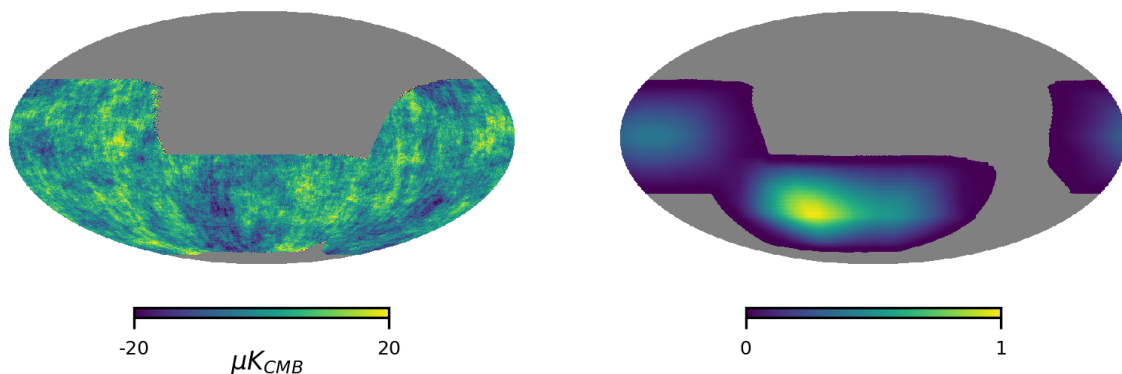


Figure 4. Left: beam-convolved Stokes Q map at 280 GHz produced by simulating the SAT scan strategy with `beamconv`, assuming an input sky model including CMB and Gaussian foregrounds. Right: the apodized custom mask that is generated from the detector hits for the simulated SAT-like scan strategy. Both maps are shown in Equatorial coordinates.

true (i.e. full-sky) E -modes that leak into the observed B -modes due to mode mixing caused by the sky mask. B -mode purification reduces the power spectrum variance on the large scales by order 50% or more, depending on the noise level, and is therefore crucial for achieving SO’s projected precision of $\sigma(r) = 0.002 - 0.003$ [35]. Beam convolution, on the other hand, reduces the observed signal power spectrum on the small scales, effectively up-weighting the instrumental noise contribution, and increasing the uncertainty on the parameters sensitive to $\ell > 150$, such as A_{lens} . The beam deconvolution in `NaMaster` is performed by including only the symmetric component of the beam in the mode-mixing matrix, $M_{\ell\ell'}$. Hence, we can expect a slight bias when estimating the power spectra of maps convolved with asymmetric beams using this method (see section 4.1), or when using the incorrect beam transfer function due to the chromaticity effects discussed above.

Finally, some degree of bias is to be expected when employing time-domain simulations instead of an analytic approach using the (true) component spectra. Simple binning mapmakers like the one `beamconv` uses are limited by the degree of cross-linking they can get in each pixel. They are also subject to sub-pixel errors [40]. The bias on the resulting BB spectra between the two approaches is shown in appendix A and is further quantified in terms of the corresponding impact on the tensor-to-scalar ratio and foreground parameters.

3.3 B -mode analysis pipeline

For the foreground component separation, we use a power-spectrum-based algorithm as implemented in `BBPower`,⁹ which is introduced in [24, 31, 32] and in [26], where it is referred to as ‘Pipeline A’. The code consists of two main stages, one to compute BB power spectra from maps, and another one to perform power-spectrum-based component separation and parameter inference. We use both stages when running on the binned time-domain simulations, and only the second stage when cross-checking with (true) analytical power spectra.

⁹<https://github.com/simonsobs/BBPower>.

Stage 1 uses **NaMaster** to compute B -mode-purified cross-frequency power spectra from a set of input maps of the partial sky. The purification algorithm requires a smoothed input sky mask, which we obtain by apodizing the hits map corresponding to our simulated SAT scanning strategy described above. As in [26], we choose an analysis range of $\ell \in [30, 300]$, and adopt multipole bins of width 10 to increase the numerical stability of the mode coupling matrix. One aspect to note is the fact that if the asymmetric modes of the beam are significant, then the standard purification would not be optimal since it would not clean up the ambiguous B -modes caused by the asymmetric beam. This could lead to additional variance in the estimator and potentially bias the B -mode power spectra.

Stage 2 carries out component separation and parameter inference at the same time. Specifically, given the full set of B -mode auto- and cross-power spectra between the six different SO SAT frequencies, **BBPower** builds a forward model for these measurements incorporating the contribution from both CMB and foregrounds. Using a likelihood for the measured spectra, it then derives the posterior distribution of all foreground and CMB parameters at the same time. The methodology is similar to that used by *BICEP-Keck* [41]. Here we use Gaussian likelihood to describe the distribution of the cross-frequency bandpower amplitudes.¹⁰ The forward model that constitutes the mean of this distribution is built by combining the contributions of all sky components (c : CMB, s : synchrotron, d : dust):

$$C_{\ell}^{\nu_i \nu_j} = \sum_{x \in \{c, s, d\}} S_x(\nu_i) S_x(\nu_j) C_{\ell}^x + \varepsilon_{\text{ds}} \sqrt{C_{\ell}^d C_{\ell}^s} [S_s(\nu_i) S_d(\nu_j) + S_d(\nu_i) S_s(\nu_j)]. \quad (3.4)$$

The first term above contains the autocorrelations of all sky components, while the second term describes the correlation between dust and synchrotron emission, quantified by the parameter ε_{ds} . The auto-spectra of the different components are described by the same models outlined in section 3.1. Note that we expect the combination of potential model mismatch with beam chromaticity to be worrisome only in cases of fairly complex foreground models like the ones including frequency de-correlation. The model mismatch in such cases can be compensated using the method of moments extension to the model [31] and dust marginalization techniques [38]. The small mismatch that can be seen between the Gaussian and non-Gaussian cases we study in this work can be absorbed by employing wide priors for the parameters of the (fairly generic) fitting model. The full set of nine parameters entering the likelihood is:

$$\{A_{\text{lens}}, r, \beta_{\text{d}}, \varepsilon_{\text{ds}}, \alpha_{\text{d}}, A_{\text{d}}, \beta_{\text{s}}, \alpha_{\text{s}}, A_{\text{s}}\}. \quad (3.5)$$

The theoretical prediction is evaluated at all integer multipoles and then convolved with the bandpower window functions provided by **NaMaster**. The bandpower covariance is the same as the one described in [26], estimated from 500 simulations of coadded Gaussian CMB, foregrounds, and SAT-like noise in the ‘baseline-optimistic’ scenario [35]. The latter assumes isotropic noise, incorporating contributions from both white and 1/f noise components, with

¹⁰Although the power spectrum of a Gaussian field generally follows a Wishart distribution, the binned bandpowers contain a sum over sufficiently many independent modes so that the central limit theorem guarantees approximate Gaussianity. As verified in [26], choosing a Gaussian likelihood in the present setup ($\ell > 30$, $\Delta\ell = 10$, SAT sky patch) is sufficiently accurate.

Parameter	Prior type	Bounds
A_{lens}	TH	[0.0, 2.0]
r	TH	[-0.1, 0.1]
β_d	G	1.54 ± 0.11
ϵ_{ds}	TH	[-1.0, 1.0]
α_d	TH	[-1., 0.]
A_d	TH	[0, ∞)
β_s	G	-3 ± 0.3
α_s	TH	[-2., 0.]
A_s	TH	[0, ∞)

Table 1. Parameter priors employed by `BBPower` for all foreground models assumed throughout the paper. Top-hat and Gaussian priors are labeled ‘TH’ and ‘G’, respectively.

observations covering 10% of the sky over a total duration of five years for the MF and UHF bands and one year for the LF bands, respectively. The posterior distribution is sampled using the `emcee` Markov Chain Monte-Carlo code [42], assuming prior distributions given in table 1.¹¹

4 Results

The results presented in this section are produced by employing simulations that are a mix of only CMB, Galactic dust, and synchrotron in all cases. To avoid mixing potential beam-related bias with the noise bias, we deliberately exclude any noise contributions from the simulations. The covariance estimate used by the component separation pipeline is the one described in [26] and includes the noise variance, cosmic variance, and foreground uncertainty estimates. After assessing the significance of non-ideal beam properties for this analysis, we demonstrate the impact of beam chromaticity on the forecast parameters for both Gaussian and non-Gaussian foreground models, considering cases both with the nominal and perturbed bandpasses.

4.1 Bias from beam systematics

In order to quantify the impact of different beam non-idealities in the inferred model parameters, we employ sky realizations that include CMB and Gaussian foregrounds and carry out four distinct types of `beamconv` simulations only for the center frequencies of the MF, UHF bands using (non-chromatic) beams of increasing complexity. The LF beams are also approximated as Gaussian beams in this case. The different types of simulations include (i) symmetric co-polar beams, (ii) symmetric co- and cross-polar beams, (iii) asymmetric co- and cross-polar beams, and (iv) asymmetric co- and cross-polar beams, including (low-amplitude) far sidelobes. For cases (i)-(iii), we truncate the beams to a maximum angle of 4° while

¹¹The (unphysical) negative values for r are employed to track volume effects from biases in other parameters (as explained in [26]). Volume effects in the posteriors can bias the cosmological parameters due to noisy (foreground) data and will be investigated in future work.

for case (iv), we use the full angular range of the beam simulations, extending to 12° . The beam transfer function employed for the deconvolution is the one of case (i) in all cases. We chose to use a small number of sky realizations for computational efficiency, and we find them to be enough to capture any significant parameter biases caused by beam non-idealities. This is also aided by the fact that the simulations used are noiseless, although they still include cosmic variance. The estimated best-fit parameters for the four beam cases, assuming only the center frequencies of the bands and averaged over ten different sky realizations,¹² are depicted in figure 5 through their one-dimensional posteriors. These are obtained by averaging the posterior distributions of each parameter across all sky realizations.

The inclusion of the simulated symmetric cross-polar beam (red) does not appear to change the estimated values of the forecast parameters as compared to the symmetric co-polar beam simulations (green). This is likely due to the small relative amplitude of the cross-polar component (see appendix B). However, the results for the asymmetric beam simulations do display small, but still visible, deviations, both for the narrow (magenta) and wide beams (blue). In particular, we can see two distinct sets of distributions forming for the spectral indices of dust and synchrotron (β_d and β_s), depending on whether or not the employed beams include asymmetric components. This is because an asymmetric beam induces a frequency-dependent asymmetry in the maps, and the most direct way for the likelihood to interpret this is through a shift in the spectral parameters of the foreground components. The bias induced by beam asymmetry is approximately 0.2σ for both β_d and β_s , where σ represents the average uncertainty across all sky realizations in the symmetric co-polar beam scenario. This type of bias significantly outweighs any additional bias from the beam’s far sidelobes, which are much fainter than the main beam (by six to eight orders of magnitude compared to the amplitude at the beam center). All other parameters are biased by less than 0.05σ , across the various cases studied. Note, however, that these statements solely refer to the beam systematics derived from our ideal optics simulations and may not accurately reflect the challenges present in real-world conditions. The in-band beam chromaticity, though, is a consequence of the diffraction-limited optics of the SATs and is not expected to change, even if, for example, we observe larger amplitudes in the far sidelobes when modeling the beams from planet observations. The reason for assessing the beam systematics present in the simulations was, thus, only to showcase they are not strong enough to mask the chromaticity effect we want to capture. Another relevant remark to make here is that beam asymmetry can introduce T -to- P leakage, which the HWP is expected to prevent. The contribution of the HWP in mitigating this type of leakage is thoroughly discussed in [43]. Moreover, achieving uniform coverage across multiple HWP angles is expected to reduce beam asymmetry, which in turn helps mitigate its impact on the spectral indices of the galactic foregrounds. Consequently, the previous results could be further verified by also simulating temperature scans and including the SATs HWP in the assumed optical setup.

Any residual bias between the symmetric co-polar beam simulations and the input parameter values is primarily due to the integration of `beamconv` into the pipeline and, to a

¹²We observed throughout the paper that while a larger number of simulations could improve the accuracy of the forecast parameters in the achromatic beam scenario, the bias on the parameters due to the inclusion of chromatic beams converges rapidly, typically within the range of five to ten simulations.




Figure 5. The posterior distributions of the nine model parameters as estimated from ten sets of six `beamconv` maps where the MF, UHF frequency maps were convolved with symmetric co-polar (green), symmetric co- and cross-polar (red), asymmetric co- and cross-polar (magenta) and wide asymmetric co- and cross-polar beams (blue). The true values for these parameters are also shown for reference (grey vertical lines). The component separation method is applied only to the center frequencies of each band for a net evaluation of the beam systematics impact that is decoupled from the beam chromaticity effect.

lesser extent, to deviations of these beams from perfect Gaussian distributions. The former is further examined in appendix A, where we observe small biases in all parameters, similar to what is seen in figure 4.1. Notably, when employing `beamconv` in the simulations, the posteriors of A_{lens} and α_d show a visible difference. While we expect a small difference in all parameters due to the use of a mask motivated by the scan strategy (as opposed to the simple estimate for which these parameters were recovered in [26]), A_{lens} and α_d are particularly affected by the loss of power at small scales which occurs due to the sparse distribution of detectors on the focal plane (we use two hundred detector pairs as opposed to the few thousand detectors present on the actual SAT focal plane). This choice was made with

respect to the computational efficiency we wanted to achieve.¹³ However, any bias between the analytic and `beamconv` cases is included both in the chromatic and non-chromatic beam simulations and does not, thus, impact the results.

4.2 Chromatic beams

4.2.1 Gaussian foregrounds

We now repeat our analysis, applying our component separation pipeline to ten sets of six `beamconv` maps (one per observing band) that were simulated employing the approximated SAT scanning strategy described in the previous section. The simulations are performed for the center frequencies of the LF bands assuming Gaussian beams and for five single frequencies within each of the MF and UHF bands, using monochromatic PO beams that include asymmetry, cross-polarization, and far sidelobes. The best-fit tensor-to-scalar ratio and foreground parameters are then derived from band-averaged maps that were created by applying the reference bandpasses from [24] to the MF, UHF monochromatic maps and compared against the case where only the center-frequency maps were used for all bands. Note that the beam used to deconvolve the power spectra in the current and following sections is the harmonic transform of the symmetric center-frequency or band-averaged beam (for the achromatic and chromatic cases, respectively) and includes both the co- and cross-polar response.

The average values of the best-fit parameters from these ten simulation sets are shown in table 2 and are denoted as ‘Chromatic’ or ‘C’. The corresponding parameters, where frequency bands are represented solely by their center frequencies, are also provided for reference. These are labeled as ‘Achromatic’ or ‘A’, and are presented alongside the input values assumed in the simulated sky configuration. The last column presents the bias between the chromatic and achromatic beams scenario in terms of the anticipated uncertainty for each parameter in the achromatic case.

Upon examining the table, we conclude that the most noticeable effects of beam chromaticity apply to the Galactic dust parameters, with the amplitude of the dust BB spectra inheriting a 0.77σ bias when chromatic beams are employed in the simulations. On the contrary, the tensor-to-scalar ratio, r , appears to be largely insensitive to the beam’s frequency dependence. This is due to the fact that the constraining power on r is at degree scales (around $\ell \sim 80$), thus the beam systematics that are more pronounced at the higher multipoles, like chromaticity (see figure 3), have very little effect. This is different from the impact on other parameters like the lensing amplitude, which does depend on what happens at $\ell \sim 300$. Furthermore, the small values associated with the synchrotron spatial parameters should be interpreted with caution throughout the analysis results in this and the following sections. This is because they are primarily informed by the LF bands, where the simulations did not account for the beam chromaticity effect. All parameters, however, show deviations smaller than 1σ .

Finally, it is worth noting that the beam chromaticity bias on r is not expected to increase even if a non-zero input value of r is chosen in the fiducial model. We would,

¹³A full set of chromatic beam simulations for a given foreground model can be performed within a week.

Parameter	Input	Achromatic (A)	Chromatic (C)	(C - A)/ σ (A)
A_{lens}	1	0.94 ± 0.04	0.93 ± 0.04	0.18
r (10^{-3})	0	1.320 ± 2.282	1.369 ± 2.305	0.02
β_d	1.54	1.540 ± 0.009	1.538 ± 0.008	0.17
ε_{ds} (10^{-3})	0	-3.375 ± 34.65	-0.35 ± 32.33	0.09
α_d	-0.16	-0.229 ± 0.035	-0.239 ± 0.031	0.27
A_d	28	27.8 ± 0.9	27.1 ± 0.7	0.77
β_s	-3.0	-2.992 ± 0.051	-2.979 ± 0.051	0.24
α_s	-0.93	-0.918 ± 0.197	-0.920 ± 0.194	0.01
A_s	1.6	1.57 ± 0.14	1.56 ± 0.14	0.06

Table 2. The average best-fit parameters as estimated from ten simulation sets of six `beamconv` maps presented for the ‘Chromatic’ and ‘Achromatic’ MF/UHF beams scenarios. The input values of the fitted parameters are also presented in the second column for reference. The last column presents the beam chromaticity bias on the best-fit values in terms of their expected uncertainty in the achromatic beams case.

however, find a larger value for $\sigma(r)$ due to the added cosmic variance. Moreover, the most significant biases are observed on the dust spatial parameters (A_d and α_d), which are known to be largely uncorrelated with r , unlike the spectral index β_d , which remains largely unaffected. Nevertheless, a follow-up study could ensure that this is the case. The interplay between the input r (0 versus 0.01) and foreground anisotropy was studied using the same pipeline in [26], see table 5.

4.2.2 Non-Gaussian foregrounds

We proceed by assuming non-Gaussian foregrounds and repeating the steps described in the previous section. In this case, we use the `PySM` module to produce template maps of dust emission and synchrotron radiation that describe, once again, a modified black-body distribution and power-law, respectively. These foreground models correspond to models ‘d0’ and ‘s0’ of the `PySM` package and are commonly used as representative non-Gaussian test models [26, 44]. The spectral indices of synchrotron and dust are constant across the sky, but the template maps now include the large-scale anisotropies caused by the Milky Way. Much of the emission from the Galactic plane is removed from the simulated maps by applying the mask described in section 3.2. Figure 6 shows the Stokes Q maps generated for the lowest (left) and the highest (right) band center frequencies, namely 27 and 280 GHz, using the ‘d0s0’ model. An important remark at this stage is that these non-Gaussian simulations display a larger dynamic range that could be picked up by the beam sidelobes. The best-fit parameters obtained by combining chromatic beams with non-Gaussian foreground models are presented in the same manner as for the Gaussian case. Table 3 summarizes the average values from ten simulation sets, utilizing both achromatic and chromatic beams. It also presents the biases between these two cases, estimated as a function of the expected uncertainty per parameter.

Parameter	Input	Achromatic (A)	Chromatic (C)	(C - A)/ σ (A)
A_{lens}	1	0.94 ± 0.04	0.93 ± 0.04	0.14
r (10^{-3})	0	0.916 ± 2.273	0.930 ± 2.283	0.01
β_d	1.54	1.539 ± 0.007	1.539 ± 0.007	0.01
ε_{ds} (10^{-3})	—	35.7 ± 49.5	46.7 ± 46.2	0.22
α_d	—	-0.318 ± 0.036	-0.335 ± 0.033	0.47
A_d	—	34.1 ± 1.1	33.4 ± 1.0	0.53
β_s	-3.0	-3.05 ± 0.09	-3.05 ± 0.09	0.15
α_s	—	-1.579 ± 0.303	-1.597 ± 0.293	0.06
A_s	—	0.63 ± 0.14	0.63 ± 0.14	0.02

Table 3. The average best-fit parameters as estimated from ten simulation sets of six `beamconv` maps produced using sky templates of CMB and non-Gaussian foreground models convolved with both ‘Achromatic’ (A) and ‘Chromatic’ beams (C) for the MF/UHF bands. Selected input parameters are presented in the second column while the last column describes the beam chromaticity bias on the best-fit values with respect to their expected uncertainty for achromatic beams.

Note that we do not provide input values for ε_{ds} , α_d , A_d , α_s , and A_s , in the second column as, in this case, the spatial properties of the dust and synchrotron maps are determined by the data on which `PySM` is based, rather than any input power spectrum. The best-fit values for the power law parameters of the dust and synchrotron power spectra instead depend on the particular sky region targeted. Furthermore, one should consider that, even though the Galactic plane is largely removed when masking the data, strong emission from that region may still leak into the unmasked footprint via the beam sidelobes (especially at the mask edges), potentially impacting the fitted parameters. However, it is important to note that the values of $\alpha_{d/s}$ and $A_{d/s}$ used to generate the Gaussian simulations discussed earlier are determined as the best-fit values derived from a large sample of ‘d0s0’ maps (without any additional beam convolution) and the analysis mask described in [26]. The choice to use a Gaussian likelihood also for non-Gaussian foregrounds is motivated by the same work, where the authors tested that the chi-squared statistic computed for non-Gaussian foregrounds and using the simulation-based covariance discussed earlier agreed with the expectations for Gaussian data (see appendix A of the same publication).

As in the case of Gaussian foregrounds, the largest bias corresponds to the dust spatial parameters, A_d and α_d , at approximately 0.5σ for both parameters. It is, however, interesting to notice the increased bias of the dust-synchrotron correlation parameter, ε_{ds} (0.22σ as compared to 0.09σ in the Gaussian foreground scenario). Independently of the foreground model assumed, the two components are convolved with the same beam and the fitting code can interpret the additional frequency scaling introduced by the beam as additional correlation between the foregrounds. This result is more obvious for the non-Gaussian simulations (that already include some level of dust-synchrotron correlation) described in table 3 but note that it also applies to the quoted values of table 2. The beam frequency scaling increases also the correlation between the Gaussian dust and synchrotron components. In fact, the slight

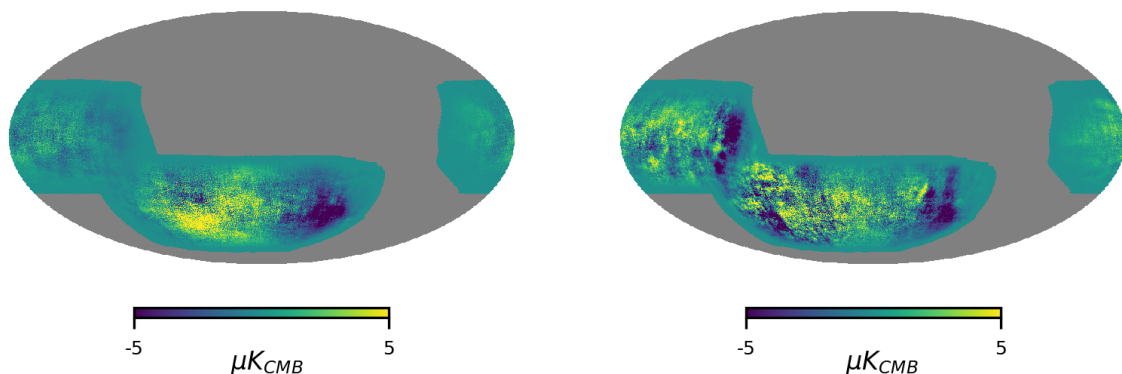


Figure 6. Stokes Q maps of sky simulations employing non-Gaussian foreground models at 27 GHz (left) and 280 GHz (right), after applying the custom mask.

anti-correlation observed in the achromatic beam scenario is nearly eliminated when beam chromaticity is included in the simulations. Overall, the chromaticity bias on all forecast parameters remains approximately below the $\sim 0.5\sigma$ level, when we assume non-Gaussian foreground models.

4.3 Impact of bandpass uncertainty

Bandpass uncertainty can arise due to variations in the fabrication process of the filters or from inherent systematics present in their measurements conducted using a Fourier transform spectrometer [45, 46]. It can also be the result of frequency-dependent reflections between elements in the optical system. Spectrally averaging over a finite bandwidth can help in mitigation. However, care needs to be employed to ensure the methodology used to characterize the receiver elements does not differ from the as-used configuration at a level that influences the calibration data acquired [47]. Time-varying atmospheric fluctuations further contribute to bandpass variations by inducing alterations in the frequency-dependent atmospheric transmission. These factors collectively influence fluctuations in both the effective gain and central frequency of the bandpass as explained in [48]. A detailed analysis of the bandpass calibration requirements for the SO SATs is provided in [24], where the allowed bandpass variation is quantified in terms of effective frequency and gain. According to this work, both of these parameters must be known to percent level or better. Note that these constraints are designed with respect to avoid biases of the order of $\Delta r \sim 10^{-3}$ on the tensor-to-scalar ratio, r , measuring which is the primary scientific goal for the SATs.

The gain can be expressed as the integrated product of the telescope’s effective area, A_{eff} , the instrumental bandpass, and the instrument’s beam response:

$$g = \int A_{\text{eff}} \tau(\nu) B(\theta, \phi, \nu) d\nu. \quad (4.1)$$

Since the authors of [24] employ achromatic beams in the analysis, we attribute the majority of gain variations to arise from bandpass variations. However, in this work, our interest is not in redefining the bandpass uncertainty requirements in terms of the r -constraint. Instead, we focus on quantifying potential biases across all fitted parameters arising from the interplay of




Figure 7. Ten perturbed versions of the reference SO SAT bandpasses from [24] for the MF and UHF frequency bands. The nominal bandpasses (shown with black dashed lines) have been perturbed by a slope ranging from -1 to 1.

bandpasses with systematic chromaticity and chromatic beams. We, thus, employ a simple model where we perturb the nominal bandpasses by adding a slope, parametrized by a spectral index, β_τ . In this way, we construct a set of ten perturbed bandpass curves, $\tau(\nu)^{\text{pert}}$, for each of the MF and UHF bands as follows:

$$\tau(\nu)^{\text{pert}} = \tau(\nu) \left(\frac{\nu}{\nu_0} \right)^{\beta_\tau}. \quad (4.2)$$

The parameter ν_0 refers to the same effective frequencies assumed throughout the analysis, while we assign β_τ ten values uniformly spanning the range $[-1,1]$.

Figure 7 shows the newly constructed perturbed bandpasses (curves of various colors) for all bands considered in this work, plotted against the reference bandpasses (black dashed lines). The changes in β_τ become clearly evident. The amplitude of the bandpass variations allowed by our parametrization is comparable with those originating from realistic atmospheric variations and the frequency-dependent scaling from the sky components presented in [48]. Specifically, by observing figure 4 of the same work, we deduce that the range of slopes covered in our case is significantly larger. The ratio between each perturbed bandpass integrated over the full frequency range and the corresponding value for the reference bandpass does not exceed 1.5% for the MF bands and 0.6% for the UHF bands. These upper boundaries represent approximately half of the allowed percent gain uncertainty determined in [24], ensuring that the bias on r does not exceed the desired $\Delta r = 10^{-3}$ from the bandpass perturbations alone. This will not necessarily be the case when employing both the above bandpasses and the chromatic beams used in the previous sections.

We construct ten new versions for each of the ten sets of band-averaged `beamconv` maps used in the analysis of sections 4.2.1 and 4.2.2, one for each of the new perturbed bandpasses.

We assume the same input sky realizations as before and quantify the combined impact of beam chromaticity and bandpass fluctuations on the forecast parameters for both the Gaussian and non-Gaussian foreground scenarios. We do so similarly to the previous sections by taking the bias of each parameter with respect to its expected mean and uncertainty. This time, however, we use the mean and uncertainty of the fitted parameters in the chromatic beams and nominal bandpass cases (namely case ‘C’ in tables 2, 3) for each foreground model. This is because we want to specifically quantify the added bias from the sloped bandpasses with respect to the already established bias due to chromatic beams.

The bias of each parameter averaged over all sky realizations scales linearly and smoothly as a function of β_τ around its expected value in the ‘C’ scenario for both Gaussian and non-Gaussian foregrounds. We will thus quote these biases only at the extreme values of the bandpass spectral index. Note that the reason for considering as many β_τ values was to ensure that the pipeline was not finding spurious degeneracies between the frequency scaling from the sky, the beam and the bandpasses for any of the assumed bandpass slopes. At $\beta_\tau = \mp 1$ and assuming Gaussian foregrounds, we compute biases of $\pm 0.05\sigma$, $\mp 0.03\sigma$, $\pm 0.07\sigma$, $\pm 0.24\sigma$, $\mp 0.42\sigma$, $\mp 0.83\sigma$, $\pm 0.27\sigma$, $\mp 0.01\sigma$, and $\mp 0.01\sigma$ for A_{lens} , r , β_{d} , ε_{ds} , α_{d} , A_{d} , β_{s} , α_{s} , and A_{s} , respectively. Note once again that we now assume $\sigma = \sigma(\text{C})$. These are symmetric biases around the $\beta_\tau = 0$ nominal case. We observe comparable results for the non-Gaussian case with average biases of $\pm 0.02\sigma$, $\pm 0.1\sigma$, $\mp 0.04\sigma$, ∓ 0.02 , ∓ 0.48 , ∓ 0.67 , ∓ 0.02 , ∓ 0.01 , and ± 0.04 for the same parameters.

The tensor-to-scalar ratio, r , appears to be insensitive to the simulated bandpass perturbations; even for the non-Gaussian case, the bias is at the 0.1σ level. Once again, most of the impact is absorbed in the dust spatial parameters independently of the assumed foreground model. In the Gaussian foreground scenario, we also observe smaller biases in the spectral indices of dust and synchrotron ($\beta_{\text{d}/\text{s}}$). These are consistent with the results of sections 4.1, 4.2.1, where we demonstrated that the pipeline aims to account for the convolution of non-Gaussian beam modes with Gaussian sky realizations through shifts in the spectral parameters of the foreground components. In conclusion, the results indicate that even in the fairly exaggerated cases of $\beta_\tau = \mp 1$, the added bias from bandpass variations does not exceed the expected uncertainty of any of the fitted parameters estimated assuming chromatic beams for the SATs (and r in particular is remarkably robust). Note again that the synchrotron parameters should be interpreted carefully, as they are predominantly informed by the LF bands, the maps of which are only convolved with achromatic beams.

5 Conclusion and discussion

In this paper, we present the method and associated pipeline for quantifying the beam chromaticity impact on the large-scale B -mode power spectra of the SO SATs, through the constraints on the tensor-to-scalar ratio, r , and foreground parameters derived from them. The method entails:

- Simulating five monochromatic PO beams for each of the MF and UHF frequency bands, at frequencies evenly spanning the band’s full frequency range.

- Performing the time-domain convolution of simulated sky templates with these beams while assuming a SAT-like scan strategy, using `beamconv`.
- Estimating the power spectra of these beam-convolved maps after masking and minimizing, correcting for E -to- B leakage, and employing B -mode purification techniques with `NaMaster`.
- Using these spectra to determine the best-fit values for the tensor-to-scalar ratio, lensing amplitude, and foreground parameters, employing `BBPower`.

We perform the analysis assuming both Gaussian (section 4.2.1) and non-Gaussian foreground models (section 4.2.2) and also study the interplay between beam chromaticity and bandpass variations (section 4.3).

The beam simulations used in this analysis are generated assuming a three-lens refracting telescope, representing the fundamental optical configuration of the SO SATs (see section 2). We find that most non-idealities present in the achromatic beam simulations, like asymmetry or cross-polarization, are not expected to impact significantly the best-fit values and accuracy of the tensor-to-scalar ratio and foreground parameters (as shown in section 4.1). Note, however, that these beam non-idealities will be further modeled from real planet observations, and we will then assess again their impact, though we do not expect any major effect for this analysis. The ten sky realizations provided to `beamconv` contain CMB, Galactic dust, and synchrotron at the same frequencies as the beam simulations and the convolution happens assuming a SAT-like scan strategy.

When considering Gaussian foregrounds, we find that beam chromaticity leads to a negligible bias on the tensor-to-scalar ratio (0.02σ). In turn, most of the effect is absorbed by the parameters characterizing the spatial structure of Galactic dust, recovering a bias on the B -mode power spectrum amplitude, A_d , of 0.77σ (see table 2). We find similar results when repeating the analysis on more realistic, non-Gaussian foregrounds (PySM model ‘d0s0’). The largest impact is again on the spatial dust parameters, with both A_d and α_d biased at the $\sim 0.5\sigma$ level (see table 3). Nevertheless, the constraints on r remain comparably unaffected. This is not unexpected considering that beam chromaticity impacts mostly the higher multipoles (see figure 3) while most of the constraining power for r is at large scales ($\ell \sim 80$). Furthermore, the dust-synchrotron correlation increases when beam chromaticity is accounted for in the simulations, and this effect is independent of the employed foreground model.

In addition to quantifying the net effect of beam chromaticity, we also consider variations in the bandpass shape and evaluate how these couple to the beam frequency dependence. We do so by applying a slope perturbation to the fiducial SO SAT bandpasses (the impact of uncertainties in the effective frequencies was studied in [24]), characterized by a spectral index β_τ . We compute the bias of the best-fit parameters for a given sky realization when $\beta_\tau = \mp 1$ and then average over all realizations. The biases are expressed with respect to the expected uncertainty in the chromatic beams, nominal bandpass scenarios. For the Gaussian case, the largest biases are of the order of $\mp 0.83\sigma$ and $\mp 0.42\sigma$ for A_d and α_d , respectively, while the spectral indices of dust and synchrotron carry smaller but non-negligible biases from the convolution with asymmetric beams. For the non-Gaussian case, the bandpass variations are fully absorbed in A_d and α_d with corresponding biases of $\mp 0.67\sigma$ and $\mp 0.48\sigma$.

While this analysis indicates that in-band beam chromaticity is unlikely to cause significant bias on the tensor-to-scalar ratio and foreground parameters, there are still future directions to explore within the same scope. A key example is the development of more advanced beam simulations to establish a clear link between the complexity of the beam and the resulting bias in the foreground parameters. Another important component in this context is the SAT HWP, which is not yet included in the simulation setup. Any frequency-dependent systematics associated with the HWP (like the frequency-dependent angle offset that multi-layer HWPs often carry [36, 49, 50]) can contribute to the total, in-band instrument chromaticity if not properly modeled in advance or as part of the component separation pipeline. The latter will be studied in [51] for a SAT-like instrument. The performance of the HWP is also vital for mitigating T -to- P leakage and consequently any resulting bias from the interplay of this leakage with beam chromaticity. We do, however, expect only a small second-order effect from this interplay. Moreover, one could repeat the analysis employing other component separation approaches, particularly those acting at the map level (e.g. ‘Pipeline B’ and ‘Pipeline C’ in [26]) which have been shown to recover constraints on r that are consistent with those found using **BBPower**. Therefore, although we would expect similar results to hold in those cases (e.g. the impact of isotropic beam chromaticity would be absorbed by the foreground components), this would require a more detailed investigation. Finally, the analysis will be verified using SO SATs beam chromaticity data we are obtaining, since December 2024, by drone-borne calibrators with tunable emission frequencies [52–55].

Acknowledgments

This work was supported in part by a grant from the Simons Foundation (Award #457687, B.K.). ND and JEG acknowledge support from the Swedish Research Council (Reg. no. 2019-03959). ND, GC and FN acknowledge funding from the European Union (ERC, POLOCALC, 101096035). Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the EU or the ERC. Neither the EU nor the granting authority can be held responsible for them. KW is supported by the STFC, grant ST/X006344/1, and by a Gianturco Junior Research Fellowship of Linacre College, Oxford. SA is partially funded by a grant from the Simons Foundation (Award #457687, B.K.). DA acknowledges support from UKSA under grant ST/Y005902/1, and from the Becroft Trust. AEA is supported by the National Science Foundation (Award No. 2153201, UEI GM1XX56LEP58). JEG acknowledges support from the Swedish National Space Agency (SNSA/Rymdstyrelsen) and the Icelandic Research Fund (Grant number: 2410656-051). Funded in part by the European Union (ERC, CMBBeam, 101040169). CB acknowledges partial support by the Italian Space Agency LiteBIRD Project (ASI Grants No. 2020-9-HH.0 and 2016-24-H.1-2018), and the Italian Space Agency Euclid Project, as well as the InDark and LiteBIRD Initiative of the National Institute for Nuclear Physics, and the RadioForegroundsPlus Project HORIZON-CL4-2023-SPACE-01, GA 101135036. JE acknowledges funding from the SCIPOL project¹⁴ funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (PI: Josquin Errard, Grant agreement No. 101044073). MG

¹⁴<https://scipol.in2p3.fr/>.

is funded by the European Union (ERC, RELiCS, project number 101116027) and by the PRIN (Progetti di ricerca di Rilevante Interesse Nazionale) number 2022WJ9J33. RG would like to acknowledge support from the University of Southern California. CHC acknowledges ANID FONDECYT Postdoc Fellowship 3220255 and BASAL CATA FB210003. SCH is supported by P. J. E. Peebles Fellowship at the Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Research, Innovation and Science.

A Bias from the simulated `beamconv` approach

To assess the bias induced by integrating `beamconv` into the pipeline, we run two different kinds of simulations. Both sets refer only to the center frequencies of the SO bands (there is no chromaticity effect included) and employ Gaussian beams. In the first case, the beam convolution is performed using Gaussian smoothing, while in the second, the Gaussian beams are convolved with the maps in `beamconv` assuming the scan strategy described in section 3.2. In an additional set of simulations, we provide to `beamconv` Gaussian beams only for the LF bands while the MF and UHF beams now correspond to the symmetric part of the simulated PO beams which are used throughout the paper. This last set aims at capturing the significance of the deviation of the symmetric PO beams from their Gaussian approximation. The maps generated in all ways described above are masked using the same custom mask discussed in section 3.2, and the resulting beam-deconvolved spectra are provided to `BBPower` to obtain the best-fit estimates for the nine model parameters. For consistency with the tests performed in section 4.1, we employ Gaussian foregrounds and present the maximum likelihood posteriors of the tensor-to-scalar ratio and other parameters, averaged over ten different sky realizations in figure 8.

From the figure, we observe biases when including `beamconv` in the pipeline (red and blue curves) with respect to posteriors of the Gaussian smoothing case (black curves) which are closely centered on the input parameters (vertical gray dashed lines). The largest discrepancies refer to the lensing amplitude, A_{lens} , and scaling factor of the dust B -mode spectra, α_d , and to a lesser extent to the tensor-to-scalar ratio, r . The main reason for these biases lies in the sparsity of the focal plane of the simulated optical configuration. In order to reduce the required simulation time, we choose to only use two hundred detector pairs across the 35° simulated focal plane as opposed to the six thousand pairs employed for the actual SAT focal plane. This inevitably leads to suppression of the smaller scales that carry most of the constraining power for the lensing amplitude and the dust spatial parameters. Part of this bias is also leaking to the tensor-to-scalar ratio as the scale suppression starts around $\ell \sim 130$ and hence at scales still somewhat relevant to primordial B -modes (expected to peak at $\ell \sim 80$). The replacement of the Gaussian beams applied to the MF, UHF bands with the symmetric co-polar ones (blue curves) does not seem to shift the parameter posteriors. Note that the biases induced with `beamconv` are not a threat to the success of this analysis. Since we will be exclusively working within the context of simulations performed with this software, this shift in the estimated parameters will be consistently included as a systematic

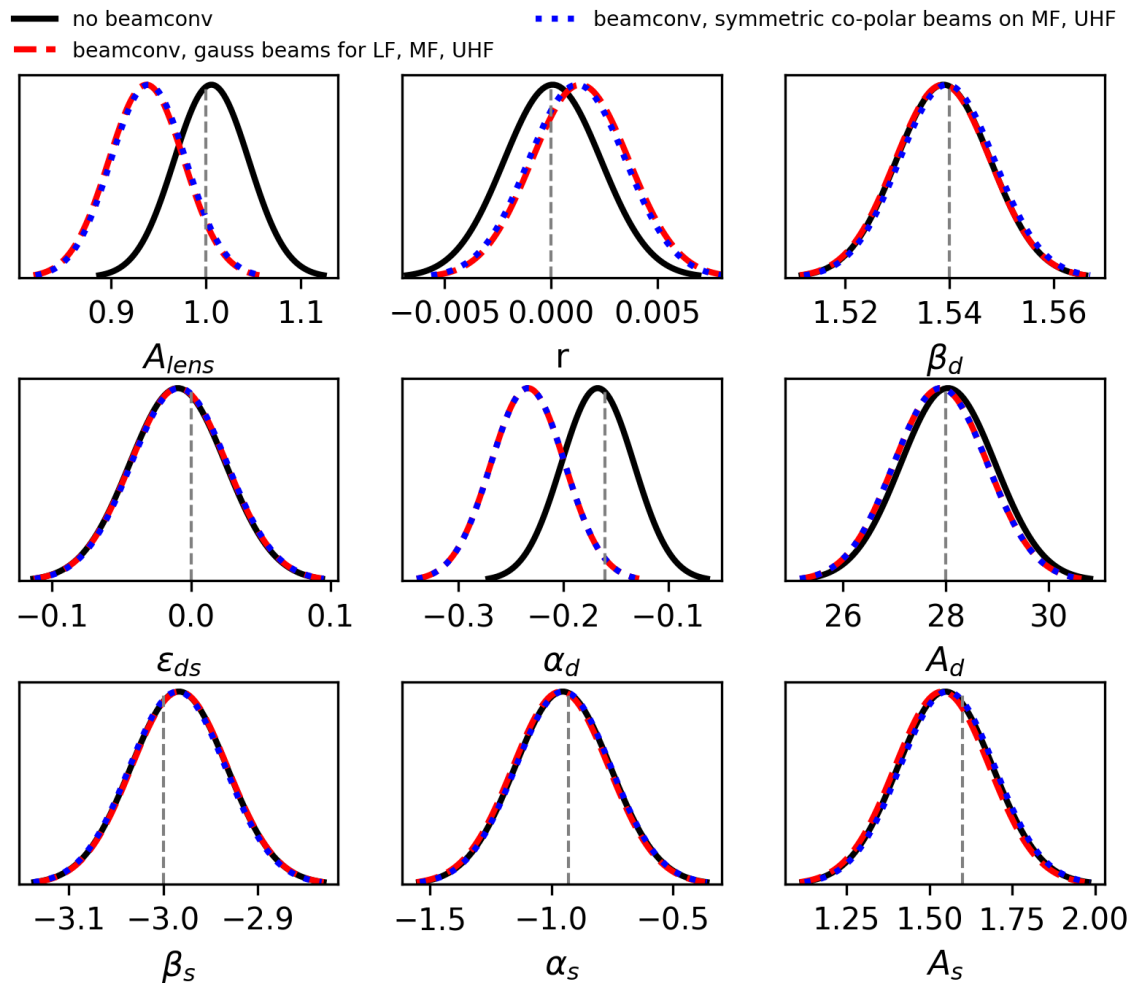


Figure 8. The parameter posterior distributions estimated from simulations where the beamconvolution was performed (i) as Gaussian smoothing to all frequency bands (black), (ii) using `beamconv` simulations with Gaussian beams for all bands (red), and (iii) using `beamconv` simulations with symmetric co-polar PO beams for the MF, UHF bands and Gaussian smoothing for the LF bands (blue).

in both chromatic and achromatic beam scenarios. Consequently, it will also be consistently eliminated when computing the difference in the best-fit parameters between the two cases.

B Cross-polarization impact on EB correlation

As discussed in section 4.1, the component separation analysis results do not seem to vary significantly when introducing the cross-polar component of the simulated beam. Nevertheless, it is important to assess the cross-polarization impact on the resulting EB spectra. A rotation of the EB spectra due to cross-polarization could interfere with cosmic birefringence measurements or affect the polarization angle calibration [12]. In contrast to beam asymmetry, the effect of beam cross-polarization can not be mitigated in the case where a HWP is included in the instrumental setup [36]. Figure 9 shows the auto- and cross-channel EB

spectra estimated from six center-frequency `beamconv` maps of a single sky realization with `NaMaster`. In the case of MF and UHF bands, the maps have been convolved either with strictly co-polar beams (blue) or with beams carrying both the co- and cross-polar components. The LF beams are still approximated with Gaussian distributions. We run simulations using both the cross-polar beams generated with `TICRA TOOLS` (orange) and a modified version, where the cross-polarization amplitude is intentionally exaggerated by an order of magnitude (green) for comparison.

The sky maps, whose spectra are shown in this figure, consist of CMB and Gaussian foregrounds. They have been masked and corrected for mode-mixing. While it would be more appropriate to use simulations of a CMB-only sky model for this assessment, the spectra convolved with beams that include the unscaled cross-polarization component appear identical to those convolved with only co-polar beams. Non-negligible EB values for these two cases are observed only due to foregrounds and cut-sky effects. The contribution of foregrounds to the EB spectra of maps convolved with either co-polar beams or beams that include the nominal cross-polar component becomes apparent when examining their auto-spectra. Notably, the spectra from the lowest (27 GHz) and highest frequency channels (280 GHz) exhibit the largest amplitudes. The MF channels (93 and 145 GHz), which refer to the least foreground-contaminated frequency regime, exhibit EB spectra amplitudes on the order of $10^{-2} \mu K_{\text{CMB}}$ and do not appear to deviate significantly from the spectra shown in [56]. To quantify any systematic E -to- B leakage from `beamconv` simulations, the best-fit birefringence angle will be estimated in the near future from maps that do not contain any foreground emission.

When we scale up the amplitude of the cross-polar beam by an order of magnitude we observe the EB correlation notably increasing, in particular for the highest frequencies. After the scaling, the cross-polar beam becomes roughly just an order of magnitude fainter than the co-polar beam, resulting in some E -to- B leakage in the power spectra, primarily impacting the higher multipoles. The degree-scale spectra, where B -modes are expected to peak, do not appear to be significantly threatened by this (large) increase in the cross-polarization amplitude.

C Dependence on the detector location

The beam model depends on the position of the detector on the telescope’s focal plane. A detector positioned at the edge of the SAT focal plane produces a significantly more elliptical beam pattern compared to a detector placed on the boresight, due to the asymmetrical illumination of the telescope aperture. Specifically, SAT beam simulations indicate a 1 – 2% variation in beam size and nearly a 50% increase in beam ellipticity between the center and edge pixel cases [23]. Figure 10 shows the logarithmic beam profiles of two detectors located at the boresight and edge of the simulated SAT focal plane for the center frequencies of the MF and UHF frequency bands, which have been uniformly averaged over the beam azimuthal angle, ϕ . Upon examining the various sub-panels of this figure, we deduce that the edge-pixel beam profiles appear more diffuse compared to the sharper main peaks we see for the center pixels. Note that the figure illustrates the power closer to the center of the beam where the main contribution comes from. For the time-domain simulations, however, we use the wide edge-pixel beams (extending to 12°) to achieve a direct comparison with the analysis described in section 4.2.2 where the input beam models were simulated for boresight pixels only.




Figure 9. The EB auto- and cross-channel spectra of a single set of six beamconv maps. The MF and UHF maps have been convolved with beams that either have no cross-polar components (blue), incorporate the original simulated cross-polar components (orange), or include cross-polar components with an amplitude ten times greater than the original simulation (green). The LF beams are assumed to be Gaussian, and the illustrated spectra refer only to the center frequencies of the SO bands. The blue and orange spectra appear identical due to the small cross-polarization in the simulated beams (2–3 orders of magnitude smaller than the co-polar component).

Parameter	Input	Chromatic	10% Chromatic	25% Chromatic
A_{lens}	1	0.93 ± 0.04	0.93 ± 0.04	0.77 ± 0.07
r (10^{-3})	0	1.369 ± 2.305	2.081 ± 2.352	1.120 ± 2.0
β_d	1.54	1.538 ± 0.008	1.537 ± 0.008	1.25 ± 0.06
ε_{ds} (10^{-3})	0	-0.35 ± 32.33	-9.677 ± 42.31	-5.314 ± 29.49
α_d	-0.16	-0.239 ± 0.031	-0.244 ± 0.031	-0.198 ± 0.034
A_d	28	27.1 ± 0.7	27.3 ± 0.8	22.1 ± 1.68
β_s	-3.0	-2.979 ± 0.051	-2.866 ± 0.061	-2.388 ± 0.154
α_s	-0.93	-0.920 ± 0.194	-1.540 ± 0.246	-0.892 ± 0.220
A_s	1.6	1.56 ± 0.14	0.79 ± 1.30	1.08 ± 0.53

Table 4. The average best-fit parameters over ten sky realizations for the cases where we increased the spectral slope of the MF/UHF beams by 10% and 25%. The fitted parameters for the nominal case are also shown together with their input values.

Applying the same component separation method on the B -mode spectra of non-Gaussian foreground maps, which have been convolved with edge-pixel beam models, yields parameter values that closely match those presented in table 3. The most significantly affected parameter is β_d , for which the beam chromaticity bias is now estimated as 0.13σ , compared to just 0.01σ in the center-pixel case. This minor difference, however, may not fully capture the realistic scenario, as the pixels at the edges of the SAT focal plane could also experience additional sidelobe power due to reflections from filters and other surface elements of the optical setup, which are not accounted for in the simulations.

D Chromaticity bias for different beam spectral response

Let us now assume that we have underestimated the intrinsic beam frequency scaling. In the simplest parametrization, we only account for the main beam which we describe solely by its best-fit FWHM. Specifically, we fit a first-order polynomial to the beam sizes of the five monochromatic frequencies employed for the MF and UHF frequency bands. We then increase the fitted slope by 10% and by an exaggerated 25% with respect to its original value while keeping the FWHM value of the center frequency of each band fixed. In this way, we increase the relative beam size difference inside the bands with respect to the corresponding band centers. We run `beamconv` simulations for the two newly constructed cases employing Gaussian beams and assuming Gaussian foregrounds. Table 4 presents the best-fit parameters in the cases where the beam spectral slope increases by 10% and 25%, averaged over ten sky realizations. The corresponding parameters for the nominal assumed spectral dependence of the beam are also shown as a reference, along with their input values.

For the more reasonable case of the 10% increase in chromaticity, we observe the largest biases in the synchrotron parameters. This is due to the component separation algorithm assigning the additional frequency scaling to errors in the assumed spectral dependence of synchrotron which more closely resembles the scaled spectral dependence of the beam than

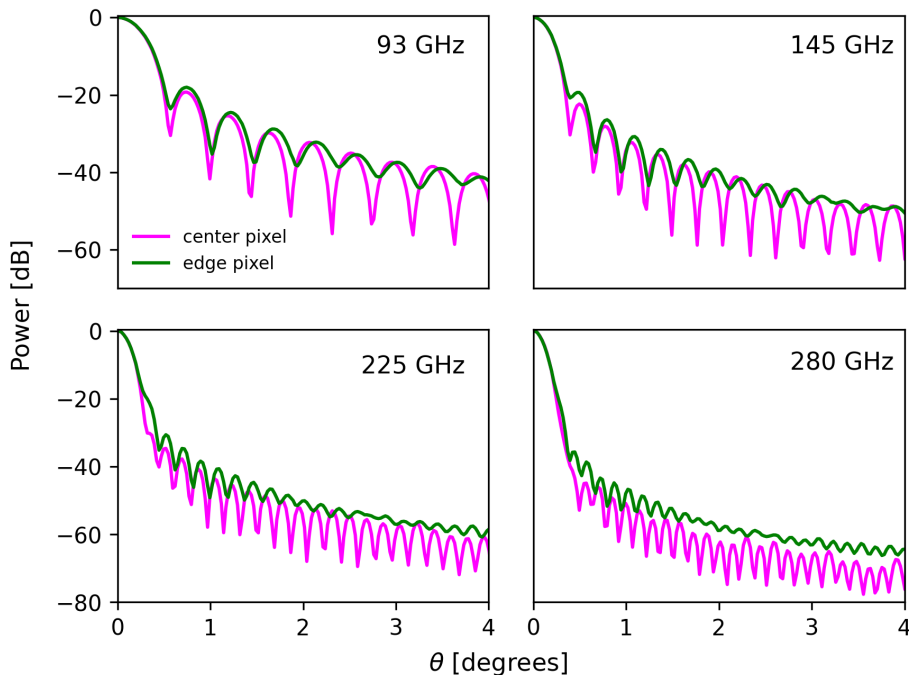


Figure 10. Azimuthally averaged radial profiles of far-field beams simulated assuming a detector at the center (purple curve) and a detector at the edge of the SAT focal plane (green curve) for the center frequencies of the four bands shown in logarithmic scale.

does that of dust. On the other hand, for the exaggerated case of 25%, it becomes evident that we can no longer recover the lensing amplitude and foreground parameters. The scaled beam spectral dependence is now very different from the ones of both dust and synchrotron and, thus, the code can not properly integrate it in the fitting model. The spectral indices of dust and synchrotron absorb most of the chromaticity bias. The tensor-to-scalar ratio increases in the first case (corresponding to a bias of $\sim 0.2\sigma$ compared to achromatic beams), yet remains largely unaffected in the second. Note that, even though we assumed Gaussian beams, we did the beam-convolution in the time-domain for consistency in the used method and associated biases throughout the paper.

References

- [1] R. Crittenden, R.L. Davis and P.J. Steinhardt, *Polarization of the microwave background due to primordial gravitational waves*, *Astrophys. J. Lett.* **417** (1993) L13 [[astro-ph/9306027](#)] [[INSPIRE](#)].
- [2] W. Hu and M.J. White, *A CMB polarization primer*, *New Astron.* **2** (1997) 323 [[astro-ph/9706147](#)] [[INSPIRE](#)].
- [3] M. Zaldarriaga, D.N. Spergel and U. Seljak, *Microwave background constraints on cosmological parameters*, *Astrophys. J.* **488** (1997) 1 [[astro-ph/9702157](#)] [[INSPIRE](#)].
- [4] A. Kogut et al., *Three-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: foreground polarization*, *Astrophys. J.* **665** (2007) 355 [[arXiv:0704.3991](#)] [[INSPIRE](#)].

- [5] PLANCK collaboration, *Planck intermediate results. XXII. Frequency dependence of thermal emission from galactic dust in intensity and polarization*, *Astron. Astrophys.* **576** (2015) A107 [[arXiv:1405.0874](#)] [[INSPIRE](#)].
- [6] PLANCK collaboration, *Planck 2018 results. IV. Diffuse component separation*, *Astron. Astrophys.* **641** (2020) A4 [[arXiv:1807.06208](#)] [[INSPIRE](#)].
- [7] J. Aumont and J.F. Macías-Pérez, *Blind component separation for polarized observations of the CMB*, *Mon. Not. Roy. Astron. Soc.* **376** (2007) 739 [[astro-ph/0603044](#)] [[INSPIRE](#)].
- [8] M. Galloway et al., *BEYONDPLANCK — III. Commander3*, *Astron. Astrophys.* **675** (2023) A3 [[arXiv:2201.03509](#)] [[INSPIRE](#)].
- [9] J. Delabrouille et al., *A full sky, low foreground, high resolution CMB map from WMAP*, *Astron. Astrophys.* **493** (2009) 835 [[arXiv:0807.0773](#)] [[INSPIRE](#)].
- [10] R. Stompor, S.M. Leach, F. Stivoli and C. Baccigalupi, *Maximum likelihood algorithm for parametric component separation in CMB experiments*, *Mon. Not. Roy. Astron. Soc.* **392** (2009) 216 [[arXiv:0804.2645](#)] [[INSPIRE](#)].
- [11] D.J. Fixsen, *The temperature of the Cosmic Microwave Background*, *Astrophys. J.* **707** (2009) 916 [[arXiv:0911.1955](#)] [[INSPIRE](#)].
- [12] A.J. Duivenvoorden, J.E. Gudmundsson and A.S. Rahlin, *Full-sky beam convolution for Cosmic Microwave Background applications*, *Mon. Not. Roy. Astron. Soc.* **486** (2019) 5448 [[arXiv:1809.05034](#)] [[INSPIRE](#)].
- [13] BICEP2 collaboration, *BICEP2 III: instrumental systematics*, *Astrophys. J.* **814** (2015) 110 [[arXiv:1502.00608](#)] [[INSPIRE](#)].
- [14] M. Lungu et al., *The Atacama Cosmology Telescope: measurement and analysis of 1D beams for DR4*, *JCAP* **05** (2022) 044 [[arXiv:2112.12226](#)] [[INSPIRE](#)].
- [15] N.J. Miller, M. Shimon and B.G. Keating, *CMB polarization systematics due to beam asymmetry: impact on cosmological birefringence*, *Phys. Rev. D* **79** (2009) 103002 [[arXiv:0903.1116](#)] [[INSPIRE](#)].
- [16] M. Shimon, B. Keating, N. Ponthieu and E. Hivon, *CMB polarization systematics due to beam asymmetry: impact on inflationary science*, *Phys. Rev. D* **77** (2008) 083003 [[arXiv:0709.1513](#)] [[INSPIRE](#)].
- [17] LITEBIRD collaboration, *Impact of beam far side-lobe knowledge in the presence of foregrounds for LiteBIRD*, *JCAP* **06** (2024) 011 [[arXiv:2312.09001](#)] [[INSPIRE](#)].
- [18] SIMONS OBSERVATORY collaboration, *The Simons Observatory: science goals and forecasts*, *JCAP* **02** (2019) 056 [[arXiv:1808.07445](#)] [[INSPIRE](#)].
- [19] SO collaboration, *The Simons Observatory: design, integration, and testing of the small aperture telescopes*, *Astrophys. J. Suppl.* **274** (2024) 33 [[arXiv:2405.05550](#)] [[INSPIRE](#)].
- [20] N. Zhu et al., *The Simons Observatory large aperture telescope receiver*, *Astrophys. J. Suppl.* **256** (2021) 23 [[arXiv:2103.02747](#)] [[INSPIRE](#)].
- [21] N. Zhu et al., *The Simons Observatory large aperture telescope receiver*, *Astrophys. J. Suppl.* **256** (2021) 23 [[arXiv:2103.02747](#)] [[INSPIRE](#)].
- [22] S. Giardiello et al., *Modeling beam chromaticity for high-resolution CMB analyses*, *Phys. Rev. D* **111** (2025) 043502 [[arXiv:2411.10124](#)] [[INSPIRE](#)].
- [23] N. Dachlythra et al., *The Simons Observatory: beam characterization for the small aperture telescopes*, *Astrophys. J.* **961** (2024) 138 [[arXiv:2304.08995](#)] [[INSPIRE](#)].

- [24] M.H. Abitbol et al., *The Simons Observatory: gain, bandpass and polarization-angle calibration requirements for B-mode searches*, *JCAP* **05** (2021) 032 [[arXiv:2011.02449](#)] [[INSPIRE](#)].
- [25] J. Didier et al., *Intensity-coupled polarization in instruments with a continuously rotating half-wave plate*, *Astrophys. J.* **876** (2019) 54 [[arXiv:1711.01314](#)] [[INSPIRE](#)].
- [26] K. Wolz et al., *The Simons Observatory: pipeline comparison and validation for large-scale B-modes*, *Astron. Astrophys.* **686** (2024) A16 [[arXiv:2302.04276](#)] [[INSPIRE](#)].
- [27] PLANCK collaboration, *Planck 2015 results VII. High frequency instrument data processing: time-ordered information and beams*, *Astron. Astrophys.* **594** (2016) A7 [[arXiv:1502.01586](#)] [[INSPIRE](#)].
- [28] PLANCK collaboration, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6 [*Erratum ibid.* **652** (2021) C4] [[arXiv:1807.06209](#)] [[INSPIRE](#)].
- [29] PLANCK collaboration, *Planck 2018 results. XI. Polarized dust foregrounds*, *Astron. Astrophys.* **641** (2020) A11 [[arXiv:1801.04945](#)] [[INSPIRE](#)].
- [30] A. Lewis, A. Challinor and A. Lasenby, *Efficient computation of CMB anisotropies in closed FRW models*, *Astrophys. J.* **538** (2000) 473 [[astro-ph/9911177](#)] [[INSPIRE](#)].
- [31] S. Azzoni et al., *A minimal power-spectrum-based moment expansion for CMB B-mode searches*, *JCAP* **05** (2021) 047 [[arXiv:2011.11575](#)] [[INSPIRE](#)].
- [32] S. Azzoni et al., *A hybrid map- C_ℓ component separation method for primordial CMB B-mode searches*, *JCAP* **03** (2023) 035 [[arXiv:2210.14838](#)] [[INSPIRE](#)].
- [33] K.M. Górski et al., *HEALPix — a framework for high resolution discretization, and fast analysis of data distributed on the sphere*, *Astrophys. J.* **622** (2005) 759 [[astro-ph/0409513](#)] [[INSPIRE](#)].
- [34] A. Zonca et al., *healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in python*, *J. Open Source Softw.* **4** (2019) 1298 [[INSPIRE](#)].
- [35] SIMONS OBSERVATORY collaboration, *The Simons Observatory: science goals and forecasts*, *JCAP* **02** (2019) 056 [[arXiv:1808.07445](#)] [[INSPIRE](#)].
- [36] A.J. Duivenvoorden et al., *Probing frequency-dependent half-wave plate systematics for CMB experiments with full-sky beam convolution simulations*, *Mon. Not. Roy. Astron. Soc.* **502** (2021) 4526 [[arXiv:2012.10437](#)] [[INSPIRE](#)].
- [37] A.E. Gambrel et al., *The XFaster power spectrum and likelihood estimator for the analysis of cosmic microwave background maps*, *Astrophys. J.* **922** (2021) 132 [[arXiv:2104.01172](#)] [[INSPIRE](#)].
- [38] LSST DARK ENERGY SCIENCE collaboration, *A unified pseudo- C_ℓ framework*, *Mon. Not. Roy. Astron. Soc.* **484** (2019) 4127 [[arXiv:1809.09603](#)] [[INSPIRE](#)].
- [39] K.M. Smith, *Pseudo- C_ℓ estimators which do not mix E and B modes*, *Phys. Rev. D* **74** (2006) 083002 [[astro-ph/0511629](#)] [[INSPIRE](#)].
- [40] S. Naess and T. Louis, *Large-scale power loss in ground-based CMB mapmaking*, *Open J. Astrophys.* **6** (2023) 2210.02243 [[arXiv:2210.02243](#)] [[INSPIRE](#)].
- [41] BICEP2 and PLANCK collaborations, *Joint analysis of BICEP2/Keck Array and Planck data*, *Phys. Rev. Lett.* **114** (2015) 101301 [[arXiv:1502.00612](#)] [[INSPIRE](#)].
- [42] D. Foreman-Mackey, D.W. Hogg, D. Lang and J. Goodman, *emcee: the MCMC hammer*, *Publ. Astron. Soc. Pac.* **125** (2013) 306 [[arXiv:1202.3665](#)] [[INSPIRE](#)].

- [43] S. Takakura et al., *Performance of a continuously rotating half-wave plate on the POLARBEAR telescope*, *JCAP* **05** (2017) 008 [[arXiv:1702.07111](#)] [[INSPIRE](#)].
- [44] B.S. Hensley et al., *The Simons Observatory: galactic science goals and forecasts*, *Astrophys. J.* **929** (2022) 166 [[arXiv:2111.02425](#)] [[INSPIRE](#)].
- [45] A. Shitvov et al., *Broadband coated lens solutions for FIR-mm-wave instruments*, in the proceedings of the *Millimeter, submillimeter, and far-infrared detectors and instrumentation for astronomy XI*, (2022) [[DOI:10.1117/12.2629968](#)].
- [46] PLANCK collaboration, *Planck 2013 results. IX. HFI spectral response*, *Astron. Astrophys.* **571** (2014) A9 [[arXiv:1303.5070](#)] [[INSPIRE](#)].
- [47] M.R. Burleigh et al., *Spectrometer baseline control via spatial filtering*, *Appl. Opt.* **55** (2016) 8201.
- [48] J.T. Ward et al., *The effects of bandpass variations on foreground removal forecasts for future CMB experiments*, *Astrophys. J.* **861** (2018) 82 [[arXiv:1803.07630](#)] [[INSPIRE](#)].
- [49] T. Matsumura et al., *Analysis of performance of three- and five-stack achromatic half-wave plates at millimeter wavelengths*, *Appl. Opt.* **48** (2009) 3614 [[arXiv:0806.1518](#)] [[INSPIRE](#)].
- [50] C.A. Hill et al., *Design and development of an ambient-temperature continuously-rotating achromatic half-wave plate for CMB polarization modulation on the POLARBEAR-2 experiment*, *Proc. SPIE Int. Soc. Opt. Eng.* **9914** (2016) 99142U [[arXiv:1607.07399](#)] [[INSPIRE](#)].
- [51] E. Tsang King Sang et al., *Half-wave-plate non idealities propagated to component separated CMB B-modes*, in preparation (2025).
- [52] F. Nati et al., *POLOCALC: a novel method to measure the absolute polarization orientation of the cosmic microwave background*, *J. Astron. Inst.* **06** (2017) 1740008 [[arXiv:1704.02704](#)] [[INSPIRE](#)].
- [53] R. Dunner et al., *Drone-based polarization calibration source for mm-wave telescopes*, in the proceedings of the *2021 15th European Conference on Antennas and Propagation (EuCAP)*, (2021) [[DOI:10.23919/eucap51087.2021.9411058](#)].
- [54] G. Coppi et al., *PROTOCOLC: an artificial calibrator source for CMB telescopes*, *Proc. SPIE Int. Soc. Opt. Eng.* **12190** (2022) 540 [[arXiv:2207.07595](#)] [[INSPIRE](#)].
- [55] G. Coppi et al., *PROTOCOLC, a W-band polarized calibrator for CMB telescopes: application to Simons Observatory and CLASS*, *Astrophys. J. Suppl.* **279** (2025) 30 [[arXiv:2502.14473](#)] [[INSPIRE](#)].
- [56] Y. Minami and E. Komatsu, *New extraction of the cosmic birefringence from the Planck 2018 polarization data*, *Phys. Rev. Lett.* **125** (2020) 221301 [[arXiv:2011.11254](#)] [[INSPIRE](#)].