The Astrophysical Journal, 999:89 (14pp), 2026 March 1
https://doi.org/10.3847/1538-4357/ae3e53

BICEP/Keck. XX. Component-separated Maps of the Polarized Cosmic Microwave Background and Thermal Dust Emission Using Planck and BICEP/Keck Observations through the 2018 Observing Season

P. A. R. Ade[1], Z. Ahmed[2,3], M. Amiri[4], D. Barkats[5], R. Basu Thakur[6], C. A. Bischoff[7], D. Beck[8], J. J. Bock[6,9], H. Boenish[5], V. Buza[10], B. Cantrall[2,8], J. R. Cheshire IV[6], J. Connors[11], J. Cornelison[12], M. Crumrine[13], A. J. Cukierman[6], E. Denison[11], L. Duband[14], M. Echter[5], M. Eiben[5], B. D. Elwood[5,15], S. Fatigoni[6], J. P. Filippini[16,17], A. Fortes[8], M. Gao[6], C. Giannakopoulos[7], N. Goeckner-Wald[8], D. C. Goldfinger[8], S. Gratton[18,19], J. A. Grayson[8], A. Greathouse[6], P. K. Grimes[5], G. Hall[8,20], G. Halal[8], M. Halpern[4], E. Hand[7], S. A. Harrison[5], S. Henderson[2,3], T. D. Hoang[13], J. Hubmayr[11], H. Hui[6], K. D. Irwin[8], J. H. Kang[6], K. S. Karkare[2,3,21], S. Kefeli[6], J. M. Kovac[5,15], C. Kuo[8], K. Lasko[13,20], K. Lau[6], M. Lautzenhiser[7], A. Lennox[17], T. Liu[8], S. Mackey[10], N. Maher[13], K. G. Megerian[9], L. Minutolo[6], L. Moncelsi[6], Y. Nakato[8], H. T. Nguyen[6,9], R. O'Brient[6,9], S. N. Paine[5], A. Patel[6], M. A. Petroff[5], A. R. Polish[5,15], T. Prouve[14], C. Pryke[13], C. D. Reintsema[11], S. Richter[5], T. Romand[6], M. Salatino[8], A. Schillaci[6], B. Schmitt[5], R. Schwarz[13], C. D. Sheehy[13], B. Singari[13,20], A. Soliman[6,9], T. St Germaine[5], A. Steiger[6], B. Steinbach[6], R. Sudiwala[1], G. P. Teply[6], K. L. Thompson[8], C. Tucker[1], A. D. Turner[9], C. Vergès[5,22], A. G. Vieregg[10], A. Wandui[6], A. C. Weber[9], J. Willmert[13], C. L. Wong[5,15], W. L. K. Wu[2,3,8], H. Yang[8], C. Yu[10], L. Zeng[5], C. Zhang[8], and S. Zhang[6]
(BICEP/Keck Collaboration)
[1] School of Physics and Astronomy, Cardiff University, Cardiff CF24 3AA, UK
[2] Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, Stanford, CA 94305, USA
[3] SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA
[4] Department of Physics and Astronomy, University of British Columbia, Vancouver, BC V6T 1Z1, Canada
[5] Center for Astrophysics, Harvard & Smithsonian, Cambridge, MA 02138, USA
[6] Department of Physics, California Institute of Technology, 1200 East California Boulevard, Pasadena, CA 91125, USA
[7] Department of Physics, University of Cincinnati, Cincinnati, OH 45221, USA
[8] Department of Physics, Stanford University, Stanford, CA 94305, USA; dobeck@stanford.edu
[9] Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA
[10] Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA
[11] National Institute of Standards and Technology, Boulder, CO 80305, USA
[12] Argonne National Laboratory, High Energy Physics Division, Lemont, IL 60439, USA
[13] School of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455, USA
[14] Service des Basses Temperatures, Commissariat a l'Energie Atomique, 38054 Grenoble, France
[15] Department of Physics, Harvard University, Cambridge, MA 02138, USA
[16] Department of Physics, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA
[17] Department of Astronomy, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA
[18] Centre for Theoretical Cosmology, DAMTP, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, UK
[19] Kavli Institute for Cosmology, Cambridge, Madingley Road, Cambridge CB3 0HA, UK
[20] Minnesota Institute for Astrophysics, University of Minnesota, Minneapolis, MN 55455, USA
[21] Department of Physics, Boston University, Boston, MA 02215, USA
[22] Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

Received 2025 October 14; revised 2026 January 22; accepted 2026 January 25; published 2026 February 25

Abstract

We present component-separated polarization maps of the cosmic microwave background (CMB) and Galactic thermal dust emission, derived using data from the BICEP/Keck experiments through the 2018 observing season and Planck. By employing a maximum-likelihood method that utilizes observing matrices, we produce unbiased maps of the CMB and dust signals. We outline the computational challenges and demonstrate an efficient implementation of the component map estimator. We show methods to compute and characterize power spectra of these maps, opening up an alternative way to infer the tensor-to-scalar ratio from our data. We compare the results of this map-based separation method with the baseline BICEP/Keck analysis. Our analysis demonstrates consistency between the two methods, finding an 84% correlation between the pipelines.

Unified Astronomy Thesaurus concepts: Cosmic inflation (319); Cosmology (343); Cosmic microwave background radiation (322); Maximum likelihood estimation (1901); Linear regression (1945)

1. Introduction

Measurements of the cosmic microwave background (CMB) have been transformative for our understanding of the Universe, leading to the establishment of the ΛCDM model as the standard cosmological paradigm. These successes were achieved through technological and analytical advances that overcame challenges posed by the faintness of the signal of interest: the detection of one part in 10^{-5} total intensity fluctuations on top of the isotropic thermal radiation at $T \approx 2.7$ K (G. F. Smoot et al. 1992) and the detection of CMB polarization, another 2 orders of magnitude smaller than the total intensity fluctuations (J. M. Kovac et al. 2002).

Significant effort has been made toward the next milestone in CMB measurements, the search for large-scale B-mode polarization at the nanokelvin level. This challenging measurement not only requires instruments with unprecedented sensitivity at millimeter wavelengths, but also experimental and analytical methods to remove astrophysical signals now approaching 2 orders of magnitude larger than the sought-after signal: Galactic foregrounds and weak gravitational lensing.

The most sensitive constraint on degree-scale B-mode polarization comes from the Bicep/Keck Collaboration XIII (2021, hereafter BK18), setting a 95% upper limit on the level of primordial B-mode power approximately 1 order of magnitude below the power of Galactic dust emission at 150 GHz within the observed patch. This is achieved by jointly fitting for an amplitude of primordial B modes and seven foreground parameters following a model of the frequency and spatial behavior of Galactic dust and synchrotron emission. This likelihood, based on S. Hamimeche & A. Lewis (2008), which compares data and theory at the level of power spectra, or band powers, is widely used in current and upcoming analyses of experiments (BICEP2/Keck & Planck Collaborations 2015; Planck Collaboration V 2020; CMB-S4 Collaboration 2022; K. Wolz et al. 2024).

In this paper, we explore an alternative approach to separating different sky components in real BICEP/Keck data with the goal of producing maps of the polarized CMB and Galactic dust emission. We investigate the differences and consistency with the baseline approach of modeling foregrounds and the CMB at the level of power spectra.

2. Map-based Component Separation

We first introduce the likelihood formalism underpinning our component-separation method. In a manner similar to the pixel-based spectral-fitting approaches of H. K. Eriksen et al. (2006), Planck Collaboration IV (2020), and R. Stompor et al. (2009), this framework produces unbiased maps of the CMB and dust signals.

2.1. Observation Model

The component maps we produce are maximum-likelihood estimates given the input of a data vector, $\boldsymbol{d}$. In our case, $\boldsymbol{d}$ is a stack of N frequency-domain polarization maps (Q and U for each frequency):

$$\boldsymbol{d} = \begin{pmatrix} \tilde{\boldsymbol{m}}_1 \\ \vdots \\ \tilde{m}_N \end{pmatrix}. \tag{1}$$

In this notation $\tilde{\boldsymbol{m}}$ contains column vectors of Q and U maps stacked together:

$$\tilde{\boldsymbol{m}} = \begin{pmatrix} \tilde{m}^Q \\ \tilde{\boldsymbol{m}}^U \end{pmatrix}, \tag{2}$$

so $\tilde{\boldsymbol{m}}$ is a vector of length $2n_p$, where n_p is the number of pixels, which is generally $\mathcal{O}(10^5)$.

The time-stream filtering and deprojection of BICEP/Keck frequency maps, $\tilde{\boldsymbol{m}}_\nu$, can be modeled as a linear operation represented by the observing matrix $\boldsymbol{R}$. This matrix is constructed alongside other BICEP/Keck data products and plays an essential role in achieving sufficient B-mode purity to infer r from B-mode power spectra (Keckarray and BICEP2 Collaborations VII 2016). This matrix can act on beam-convolved input maps to produce simulated maps as they would be observed by a BICEP/Keck frequency band ν:

$$\tilde{m}_\nu = \boldsymbol{RBm}_\nu + \boldsymbol{n}_\nu, \tag{3}$$

where $\boldsymbol{B}$ is an operator convolving the Q and U maps with their respective symmetric beam functions, b_ℓ. Both the observing matrix and beam function can be different from frequency channel to frequency channel. Further, each simulation of a BICEP/Keck Q and U frequency map includes an additive noise component, $\boldsymbol{n}$.

2.2. Multifrequency Model

For previous BICEP/Keck results, the baseline multifrequency model assumes three components: CMB, Galactic dust, and Galactic synchrotron emission (BICEP2/Keck & Planck Collaborations 2015; Keckarray and BICEP2 Collaborations X 2018). Given that the current dataset is limited to Wilkinson Microwave Anisotropy Probe (WMAP) and Planck observations at low frequencies and measures an amplitude of polarized synchrotron emission consistent with zero in the observation patch, we only attempt to explicitly separate the CMB and Galactic thermal dust in the present paper. We employ a parametric model for dust emission as a function of frequency, such that we can model each frequency map as a linear combination of CMB and dust:

$$\boldsymbol{m} = \begin{pmatrix} \boldsymbol{m}_1 \\ \vdots \\ \boldsymbol{m}_N \end{pmatrix} = \begin{pmatrix} \mathbf{1} & \boldsymbol{f}_1 \\ \vdots & \vdots \\ \mathbf{1} & \boldsymbol{f}_N \end{pmatrix} \begin{pmatrix} \boldsymbol{s}^{\mathrm{CMB}} \\ \boldsymbol{s}^{\mathrm{dust}} \end{pmatrix} \equiv \boldsymbol{Fs}, \tag{4}$$

where $\mathbf{1}$ is the $2n_p \times 2n_p$ identity matrix and $\boldsymbol{f}_i \equiv f_i \mathbf{1}$ are scaling factors multiplied to that identity matrix. We assume a modified blackbody (MBB) spectrum for Galactic dust such that the scaling factor for the ith frequency band is given by

$$f_i \sim \int b_i(\nu) \frac{\nu^{3+\beta_d}}{\exp\left(\frac{h\nu}{k_{\mathrm{B}} T_d}\right) - 1} d\nu, \tag{5}$$

where $b_i(\nu)$ is the respective bandpass. In this work, we assume a spatially constant β_d and do not treat it as a free parameter in our maximum-likelihood search.

The end result is a model relating sky component maps to observed BICEP/Keck and external frequency maps, which is given by

$$\tilde{\boldsymbol{m}} = \boldsymbol{RBFs} + \boldsymbol{n} \equiv \boldsymbol{As} + \boldsymbol{n}. \tag{6}$$

2.3. Noise Model

Apart from the signal, each frequency map also includes an additive noise term. One simple and straightforward choice is to characterize the noise fluctuations after time-stream filtering and binning into maps as Gaussian random fluctuations, which can therefore be fully described in terms of their noise covariance matrix:

$$\hat{\boldsymbol{N}} = \langle \boldsymbol{nn}^T \rangle. \tag{7}$$

In this work, we approximate this matrix to be diagonal in pixel space, with the variance maps of each frequency channel populating its diagonal. The off-diagonal noise correlations in

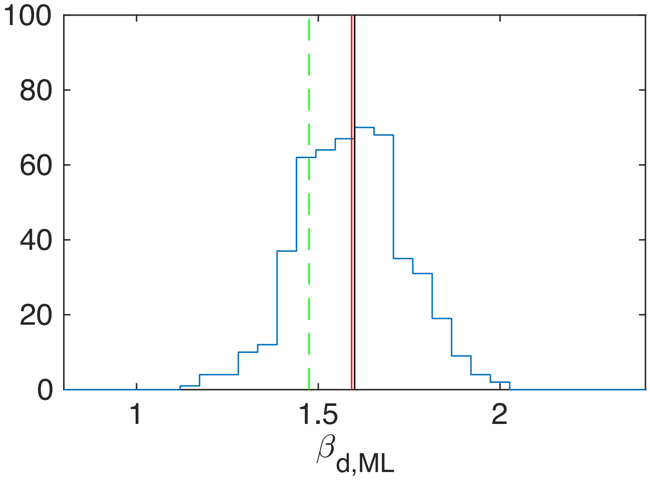

Figure 1. Histogram of the maximum-likelihood values of the dust spectral index β_d obtained in the baseline BK18 auto- and cross-spectrum analysis. The black vertical line indicates the input value to the simulations, the (nearly coincident) red line indicates the mean of the recovered best-fit values, and the dashed green line marks the best-fit value of the real data.

map space are on the 5% level and are neglected here. Note that a mismatch between the true noise covariance and this estimate will not cause a biased estimate as long as the noise is Gaussian. However, it can make the estimator suboptimal.

2.4. Likelihood Formalism

We start out with the problem of maximizing the full likelihood:

$$-2\log P(\boldsymbol{s},\, \beta_d|\boldsymbol{d}) = (\boldsymbol{d} - \boldsymbol{A}(\beta_d)\boldsymbol{s})^T\hat{\boldsymbol{N}}^{-1}(\boldsymbol{d} - \boldsymbol{A}(\beta_d)\boldsymbol{s}), \quad (8)$$

where both the foreground model, represented by the parameter β_d, and the sky signal $\boldsymbol{s}$ are unknown. We will make use of the maximization of the profile likelihood $-2\log P(\mathbf{s},\, \beta_d = \beta_d^*|\boldsymbol{d})$, where we fix the parameter in the foreground model in $\boldsymbol{A}$ to the value β_d^*. The maximization of this profile likelihood yields the generalized least squares estimator (A. C. Aitken 1936):

$$\hat{\boldsymbol{s}} = (\boldsymbol{A}^T\hat{\boldsymbol{N}}^{-1}\boldsymbol{A})^{-1}\boldsymbol{A}^T\hat{\boldsymbol{N}}^{-1}\boldsymbol{d}. \quad (9)$$

This makes the problem of separating out components a computationally expensive problem, since the inversion of the large system matrix $\boldsymbol{A}^T\hat{\boldsymbol{N}}^{-1}\boldsymbol{A}$ is expensive.

2.5. Foreground Model

This study leverages the best-fit foreground parameters derived from the mainline BK18 analysis. There, we fit a model of CMB, dust, and synchrotron to auto- and cross-frequency spectra of the real data and 499 simulated skies, including the CMB, dust, and instrumental noise. The distribution of the resulting best-fit values of the dust spectral index is shown in Figure 1 and can be used to propagate the statistical uncertainty of this parameter to the component-separated maps. We use the respective best-fit value for β_d for each simulation realization in the component map estimator in Equation (9).

Possible extensions of this approach would be the inclusion of a consistent map-based fitting of β_d and β_s in a dedicated two-step approach, following, for example, R. Stompor et al. (2009) and R. de Belsunce et al. (2022). This formalism also more directly allows for fitting spatially varying foreground parameters at the map level, which would make this pipeline more robust against biases due to complex foregrounds. Moreover, C. Leloup et al. (2023) and M. Morshed et al. (2024) have recently introduced and implemented a nonparametric maximum-likelihood framework for CMB foreground separation. We leave an exploration of these options for future work.

2.6. Comparison to the BICEP/Keck Baseline Foreground Cleaning

The baseline r analysis of BICEP/Keck accounts for foreground contamination at the power-spectrum level by modeling the foreground bias from dust and synchrotron using a spatial and frequency-space model. In the likelihood framework above, this corresponds to adding a signal prior to the full likelihood in Equation (8):

$$-2\log P(\boldsymbol{s},\, \beta_d|\boldsymbol{d}) = \ldots + \boldsymbol{s}^T\boldsymbol{S}^{-1}\boldsymbol{s} + \log\det\boldsymbol{S}. \quad (10)$$

Here, $\boldsymbol{S}$ represents the signal covariance matrix, which incorporates the autocorrelations of the CMB, dust, and synchrotron emission. These components are modeled in harmonic space using parametric representations of their respective power spectra. Marginalizing this likelihood over the sky signal $\boldsymbol{s}$ and applying the arguments of S. Hamimeche & A. Lewis (2008) yields the baseline BK18 auto- and cross-spectra likelihood.

In practice, modeling the data at the level of multifrequency spectra can lead to complications. In the BK18 analysis, for example, the analysis has to account for two significantly different observation footprints caused by the difference in field of view between BICEP2/Keck and BICEP3. This was solved by computing cross-spectra between the maps using a small-field mask for the BICEP2/Keck maps and a large-field mask for the BICEP3 and external maps (BK18).

Furthermore, introducing a prior on the foreground behavior may result in biased foreground parameter estimates. In practice, the traditional prior on the foreground behavior in harmonic space is a power-law power spectrum with an amplitude and spectral slope. While measurements of thermal dust emission and synchrotron emission suggest that this assumption is approximately true, failure to capture the actual shape of the band powers dust and synchrotron happen to produce in our patch can lead to misleading constraints on foreground spectral parameters. Hence, an alternative pipeline can improve robustness and build confidence in our foreground-cleaning capability.

3. Data

The BK18 dataset consists of observations from the BICEP2, Keckarray, and BICEP3 receivers located at the South Pole Station in Antarctica. The BICEP2 receiver observed at 150 GHz from 2010 to 2012 (BICEP2 Collaboration II 2014) with ≈500 bolometric detectors. The Keckarray, consisting of five copies of BICEP2-size receivers running from 2012–2019, initially observed at 150 GHz but switched over time to 95 and 220 GHz (Keckarray and BICEP2 Collaborations V 2015). BICEP3 is a single, scaled-up receiver at 95 GHz, which started science observations in 2016 (BICEP/Keck Collaboration XV 2022) with ≈2500 detectors.

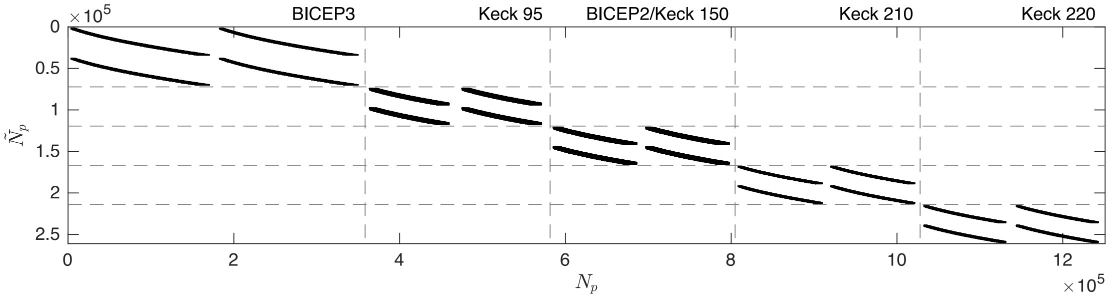

Figure 2. Plot of the nonzero matrix elements of the BICEP/Keck internal maps' observing matrices, i.e., the BICEP/Keck-specific block of the matrix $\boldsymbol{R}$ used in this work. Applying a vector including beam-convolved Q and U maps to the right of this matrix results in a vector of filtered Q and U maps in the flat pixelization used in BK18. The top-left block corresponds to the BICEP3 map, while the following blocks along the diagonal contain the observing matrices for the Keck 95 GHz, BICEP2/Keck 150 GHz, Keck 210 GHz, and Keck 220 GHz channels. Within each block on the diagonal, the top row produces a BICEP/Keck Q map and the bottom row a BICEP/Keck U map from a vector of stacked Q and U maps. This part of the matrix contains about five billion nonzero elements and is the biggest computational challenge in this analysis.

BICEP2 and Keckarray both mapped a region of sky centered at R.A. $= 0^{\mathrm{h}}$, decl. $= -57\overset{\circ}{.}5$ with an effective area of $\approx 400\,\mathrm{deg}^2$ (Keckarray and BICEP2 Collaborations X 2018). BICEP3 has a larger instantaneous field of view and hence naturally maps a larger sky area with an effective area of $\approx 600\,\mathrm{deg}^2$. This results in small-field maps at 95, 150, and 220 GHz and a large-field map at 95 GHz. In this paper, we make use of the standard BK18 maps and simulations.

These maps were produced with a filter-and-bin map maker, removing a third-order polynomial, scan synchronous signal, and T-to-P leakage templates from the data before accumulating them into maps of equirectangular pixelization, with $0\overset{\circ}{.}25$ square pixels at decl. $= -57\overset{\circ}{.}5$. The map binning operation weights by the inverse variance of the time-stream data. These weights, binned into maps, are used to get an estimate of the noise variance in T, Q, and U.

Additionally, we include external data to leverage the high-frequency observations of the Planck satellite mission covering the entire sky. In this paper, we make use of the 100, 143, 217, and 353 GHz Q and U maps from the NPIPE processing of the Planck data (Planck Collaboration LVII 2020). As opposed to the mainline analysis of BK18, these maps are not reobserved prior to using them in the estimator.

4. Implementation

4.1. Maximum-likelihood Estimator

The primary computational challenge is solving the large linear system in the estimator in Equation (9) via matrix inversion. We employ an iterative method to numerically solve for $\hat{s}$ in the linear equation

$$(\boldsymbol{F}^T\boldsymbol{B}^T\boldsymbol{R}^T\hat{\boldsymbol{N}}^{-1}\boldsymbol{R}\boldsymbol{B}\boldsymbol{F})\hat{s} = \boldsymbol{F}^T\boldsymbol{B}^T\boldsymbol{R}^T\hat{\boldsymbol{N}}^{-1}\boldsymbol{d}. \tag{11}$$

This method allows us to never have to explicitly construct the large $4n_p \times 4n_p$ matrix on the left-hand side. It is merely sufficient to build a routine that applies the matrix to the solution vector $\hat{s}$. This leads to a significant speedup given that $\boldsymbol{R}$, $\boldsymbol{B}$, and $\boldsymbol{F}$ are sparse operators either in pixel or harmonic spaces, and thanks to fast routines available to transform between the two spaces. Figure 2 highlights the sparsity pattern of the internal BICEP/Keck observing matrices, showing the blocks corresponding to the BICEP/Keck frequency bands. This matrix is heavily concentrated at the subdiagonals within each frequency-channel Q and U blocks. The beam convolution operator $\boldsymbol{B}$ convolves a map with a beam function, which is a simple multiplication in harmonic space and should hence be performed after spherically transforming the map. In this work, we assume no spatial variation of the foregrounds and hence the $\boldsymbol{F}$ matrix is simple in either space. Such variations could be easily incorporated, however, as long as they are diagonal in pixel or harmonic space.

We solve Equation (9) with a preconditioned iterative method (R. Barrett et al. 1996), employing a block Jacobi preconditioner. The use of a preconditioner significantly speeds up the convergence of the iterative method by approximating the linear operator with an easily invertible matrix. We construct this preconditioner by ignoring the observing matrix and beam in the system matrix:

$$\boldsymbol{M} \equiv \boldsymbol{F}^T\hat{\boldsymbol{N}}^{-1}\boldsymbol{F}. \tag{12}$$

Adding external data from Planck not only extends the frequency coverage but also stabilizes the numerical problem by filling in modes that are in the null space of the BICEP/Keck observing matrices.

4.2. Numerical Experiments

To focus on the problem of inverting the system matrix on the right-hand side of Equation (9), we perform numerical experiments on a more simplified problem: we correct for the effects of filtering and deprojection at the map level and produce "unbiased" maps for a single frequency only. This effectively amounts to inverting a single-frequency observing matrix by solving the estimator, e.g., for BICEP3 ("$B3$") only:

$$\hat{\boldsymbol{s}}_{B3} = (\boldsymbol{R}^T\hat{\boldsymbol{N}}_{B3}^{-1}\boldsymbol{R})^{-1}\boldsymbol{R}^T\hat{\boldsymbol{N}}_{B3}^{-1}\boldsymbol{d}_{B3}. \tag{13}$$

We know that this problem is mathematically ill-defined since the matrix $\boldsymbol{R}^T\hat{\boldsymbol{N}}^{-1}\boldsymbol{R}$ is strictly not invertible. We can regularize this problem by adding external data with more complete mode coverage than BICEP/Keck, such as Planck. For example, we can produce a combined map at around 100 GHz with BICEP3 and Planck 100 GHz (P) data, assuming the latter have negligible filtering suppression, by solving the estimator

$$\hat{s}_{95\ \mathrm{GHz}} = (\boldsymbol{R}^T\hat{\boldsymbol{N}}_{B3}^{-1}\boldsymbol{R} + \hat{\boldsymbol{N}}_P^{-1})^{-1}(\boldsymbol{R}^T\hat{\boldsymbol{N}}_{B3}^{-1}\boldsymbol{d}_{B3} + \hat{\boldsymbol{N}}_P^{-1}\boldsymbol{d}_P). \tag{14}$$

In Figure 3, we show the Q maps and the EE 2D autopower spectrum of a noise simulation for the BICEP3 and Planck 100 GHz input maps, as well as their combination $\hat{\boldsymbol{s}}_{95\ \mathrm{GHz}}$. It

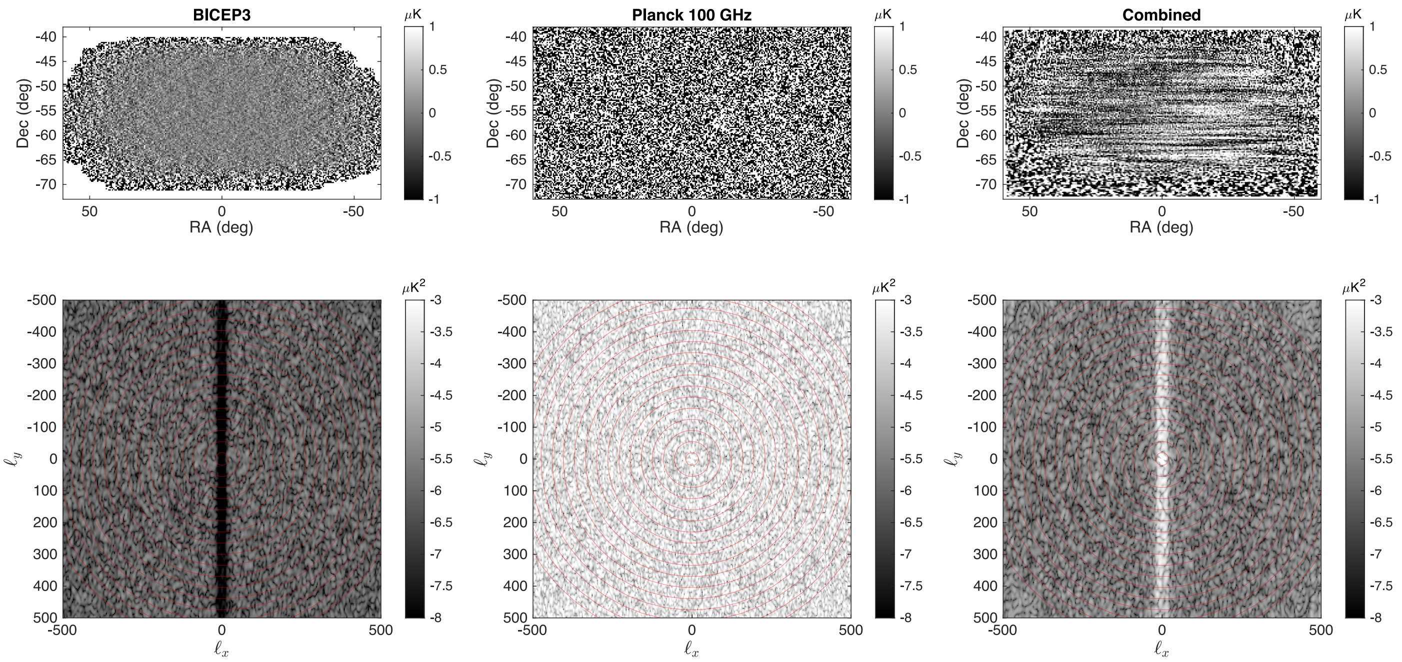

Figure 3. Q maps (top row) and the logarithmic (base 10) EE 2D autopower spectrum (bottom row) of one noise simulation for BICEP3 (left column), Planck High Frequency Instrument (HFI) 100 GHz (middle column), and their combination (right column). The combination estimator corrects for filtering suppression at the map level and hence boosts the noise compared to the filtered BICEP3 map and fills in modes from Planck for small ℓ_x. This is why the combined Q noise maps show strong horizontal stripes in the central BICEP3 map region.

illustrates how this estimator fills in the filtered-out modes in the BICEP3 map with noisier Planck modes in the "poly-trench" along the ℓ_y direction. This causes characteristic stripes in the combined map.

In Figure 4, we explore three different preconditioned iterative solvers for this problem applied to signal-and-noise maps: the conjugate gradient (CG) method, the generalized minimal residual (GMRES) method, and the biconjugate gradient stabilized (Bi-CGSTAB) method (R. Barrett et al. 1996). We find the latter to have superior convergence behavior and will use it in the following work. It should be noted that Bi-CGSTAB requires two matrix-vector multiplications per iteration, whereas GMRES requires only one. However, we find that the effective wall-clock time difference between the two is small, as the faster convergence rate of Bi-CGSTAB often compensates for the additional computational cost per step.

In Figure 5, we show the convergence behavior of the iterative solver for a signal-only, a noise-only, and a signal-and-noise simulation. Due to noise inhomogeneity, the iterative method struggles to converge. For a homogeneous input map such as a signal-only simulation, the iterative solver converges much faster.

In Figure 6, we show numerical experiments for a signal-and-noise simulation for different sizes of the input map vector. A single-frequency BICEP3 solution as in Equation (13) will take a long time to converge, and a faithful estimate is not guaranteed due to the noninvertibility of the system matrix. The regularization with an external Planck map as described in Equation (14) decreases the number of iterations needed to reach the convergence criterion of a relative residual of 10^{-4}. Lastly, the full problem of combining all frequency maps described in Section 3 to solve for a CMB component map reaches this criterion in $\mathcal{O}(70)$ steps, noting that a larger data vector leads to a larger runtime per iteration and higher memory requirements.

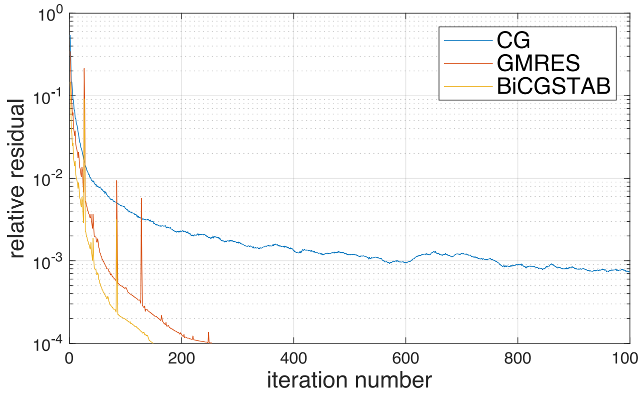

Figure 4. The relative residual $||\boldsymbol{A}\boldsymbol{s} - \boldsymbol{b}||/||\boldsymbol{b}||$, where $\boldsymbol{A} \equiv \boldsymbol{R}^T \hat{\boldsymbol{N}}_{B3}^{-1} \boldsymbol{R} + \hat{\boldsymbol{N}}_P^{-1}$ and $\boldsymbol{b} = \boldsymbol{R}^T \hat{\boldsymbol{N}}_{B3}^{-1} \boldsymbol{d}_{B3} + \hat{\boldsymbol{N}}_P^{-1} \boldsymbol{d}_P$, for each iteration of the preconditioned iterative method. We test three different iterative solvers: the classic CG, GMRES, and Bi-CGSTAB (R. Barrett et al. 1996). The spikes are due to numerical noise, which these iterative solvers are susceptible to (R. B. Sidje & N. Winkles 2011).

5. Noise Properties

The output of the estimator, $\hat{\boldsymbol{s}}$, is an unbiased map of CMB and dust, meaning the signal suppression due to filtering and deprojection is corrected for at the map level. This comes at the expense of elevated, inhomogeneous noise in the final map.

We build a map-level weight from the Q and U variance maps produced in the BICEP/Keck map making. Using our fiducial foreground model, we can build approximate maps of the Q and U variance in the CMB and dust-component maps

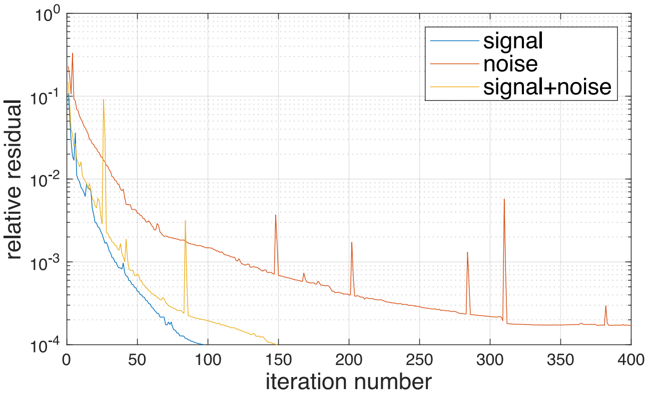

Figure 5. The relative residual as defined in Figure 4 for each iteration of the (Bi-CGSTAB) preconditioned CG method. We show the convergence performance for a signal-only, noise-only, and a signal-and-noise simulation.

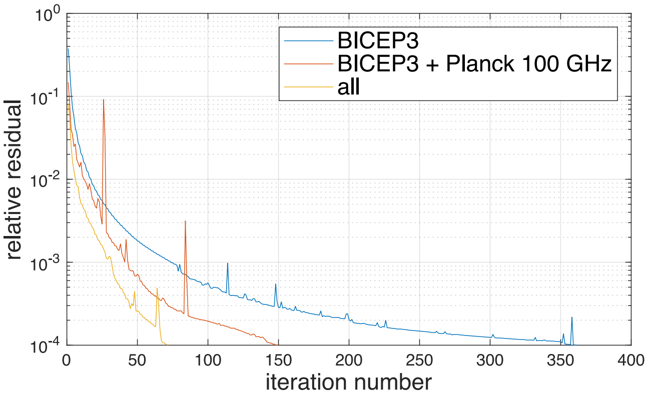

Figure 6. The relative residual, as defined in the caption of Figure 4, for each preconditioned conjugate gradient (PCG) iteration step. We show a comparison to solve for the frequency-map solution for BICEP3 and the BICEP3 + Planck combination, as well as the CMB component map solution given all input frequency maps as described in Section 3.

by taking the diagonal of

$$\tilde{N} = (\boldsymbol{F}^T\boldsymbol{N}^{-1}\boldsymbol{F})^{-1}. \tag{15}$$

We then weight the CMB and dust-component maps by the inverse of the arithmetic mean of their Q and U variance maps, just like we do for the BICEP/Keck frequency maps in BK18. These weights are shown in Figure 7. We observe that the CMB component weight extends to higher declinations due to the larger extent of the BICEP3 map, which significantly contributes to the CMB reconstruction. The dust map, however, is mostly limited to the Keck 220 GHz patch.

We multiply the Q and U maps with this weight in order to downweight noisy pixels and compute harmonic coefficients. Computing the noise variance per (ℓ, m) mode yields Figures 8 and 9. The noise variance for low-m modes is significantly boosted in the CMB component, as they correspond to modes along the BICEP/Keck scan direction, which are taken out by filtering. The wedge patterns can be explained by the correspondence of constant-m modes and the most negative observed decl.: the higher the maximal decl., the higher the m modes with nonzero power. As expected, the BICEP/Keck noise is significantly lower than the Planck noise when comparing low- and high-m modes in the CMB component. For the dust component, however, Keck 220 GHz and Planck 353 GHz contribute about equally to the sensitivity, and we obtain a more homogeneous noise variance.

6. Maps

In this section, we will present simulated and real-data CMB and dust-component maps for BICEP, Keck, and Planck High Frequency Instrument (HFI) data. For the reasons explained in the previous section, any plot of the map and computation of the power spectrum requires some weighting in the (ℓ, m) plane to avoid being dominated by large noise modes at low m.

One option is to reobserve the CMB or dust-component map again with one of the BICEP/Keck observing matrices. This will produce a CMB or dust map as observed by one of the BICEP/Keck receivers. For this, we choose the BICEP3 matrix since it has the largest footprint:

$$\boldsymbol{m}^{\mathrm{CMB/dust}} = \boldsymbol{R}_{\mathrm{B3}} \cdot \hat{\boldsymbol{s}}^{\mathrm{CMB/dust}}.$$

This method also allows us to use the well-established pipeline from observed BICEP/Keck maps to purified power spectra and estimates of r (Keckarray and BICEP2 Collaborations VII 2016; BK18).

In Figure 10, we show the CMB part of the component map estimator when running on a simulated realization of lensed ΛCDM, Galactic dust, and noise. The first row is the raw result, and the second row applies the BICEP3 observing and purification matrices. The observing matrix filters the noise-dominated low-m modes in the B-mode map, making the lensing B modes visible. In the third row, by subtracting the corresponding CMB-and-noise-only simulation, we obtain the foreground residual map caused by the statistical fluctuation of the β_d estimate. In the fourth row, we subtract a CMB-only single-frequency simulation from the corresponding CMB-only simulation run through the component-separation pipeline to obtain the numerical residual introduced by the implementation of the component-separation estimator. Both foreground and numerical residuals are well below the E- and B-mode signals and noise in the CMB component map. We will quantify their amplitude and impact on science results at the level of power spectra below.

The real-data E- and B-mode maps of the CMB and the dust components are shown in Figures 11 and 12, respectively. Figure 13 presents the derived dust-component Q and U maps alongside a noise realization and a comparison with Planck 353 GHz data. The high signal-to-noise ratio filamentary dust structures are clearly visible in Q and U, consistent with features previously identified in external neutral-hydrogen data (BICEP/Keck Collaboration XVI 2023) as well as SPT-3G (J. A. Zebrowski et al. 2025).

7. Power-spectrum Estimation

As described in the previous section, the inhomogeneous noise in the CMB and dust-component maps requires some kind of weighting in harmonic space in order to downweight noisy modes. Considering this, we will investigate in this section how to optimally compute power spectra from these maps.

7.1. Methods

7.1.1. Pseudo-C_ℓ Method

In order to use the pseudo-C_ℓ method, which is the standard method for estimating power spectra of BICEP/Keck maps,

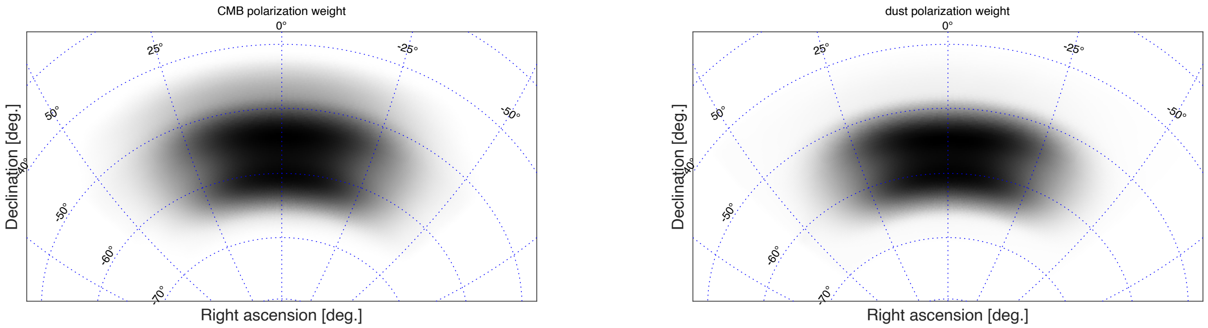

Figure 7. The polarization weight, defined as the inverse arithmetic mean of the Q and U noise variance, assuming white noise per pixel for the CMB component (left) and the dust component (right) in arbitrary units. Due to the extended coverage of BICEP3, the CMB extends above 50° in decl., while the dust-component map is most sensitive in the BICEP2/Keck region.

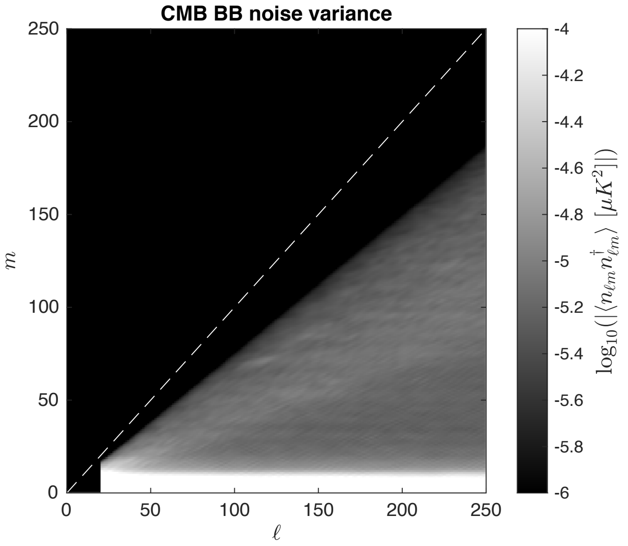

Figure 8. The noise variance per (ℓ, m) mode of the CMB component map obtained by averaging over 499 noise-only simulations, after applying the corresponding map-level weight shown in Figure 7 before computing the harmonic coefficients.

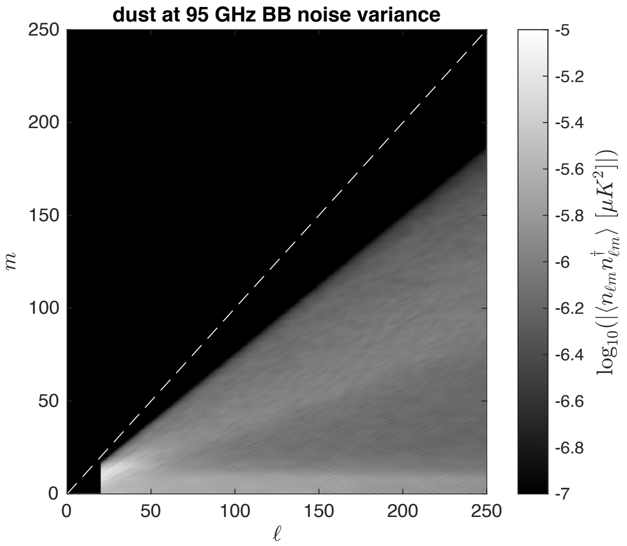

Figure 9. Same as Figure 8, but for the dust-component map scaled to 95 GHz.

we need to introduce some additional harmonic-space weighting to the standard pipeline. One natural way, as discussed in the previous section, is to apply a BICEP/Keck observing matrix to the CMB and dust-component maps. This not only applies a low-m filter but also renders these maps into the right format to use the existing power-spectrum estimation pipeline based on computing pseudo-C_ℓs of matrix-purified maps as outlined in Keckarray and BICEP2 Collaborations VII (2016).

7.1.2. Optimal Quadratic Maximum-likelihood Method

Following M. Tegmark (1997) and E. F. Bunn & B. Wandelt (2017), we can construct an optimal power-spectrum estimator as the solution of a likelihood maximization by estimating a given band power at bin b as

$$\boldsymbol{D}_b^{BB} \sim \boldsymbol{m}^T \boldsymbol{C}^{-1} \boldsymbol{P}_b \boldsymbol{C}^{-1} \boldsymbol{m},$$

where $\boldsymbol{m}$ are the input maps and the total covariance

$$\boldsymbol{C} \equiv \boldsymbol{S} + \bar{\boldsymbol{N}}$$

is the sum of the signal and noise covariance. The operator $\boldsymbol{P}_b$ is given by

$$\boldsymbol{P}_b = \frac{\partial \boldsymbol{C}}{\partial D_b^{BB}}.$$

We model the signal covariance to be diagonal in harmonic space, given our baseline fiducial cosmological model derived from Planck 2013 cosmological parameters (Planck Collaboration XVI 2014):

$$\boldsymbol{S} = \boldsymbol{Y}^{\dagger} \mathrm{diag}(\boldsymbol{C}_\ell^{EE}, \boldsymbol{C}_\ell^{BB}) \boldsymbol{Y},$$

where $\boldsymbol{Y}$ and $\boldsymbol{Y}^{\dagger}$ are forward and backward spherical harmonic transformations, respectively. The noise covariance is modeled using noise-only simulations as

$$\bar{\boldsymbol{N}}^{-1} = \boldsymbol{w} \boldsymbol{Y}^{\dagger} \mathrm{diag}(N_{\ell m})^{-1} \boldsymbol{Y} \boldsymbol{w},$$

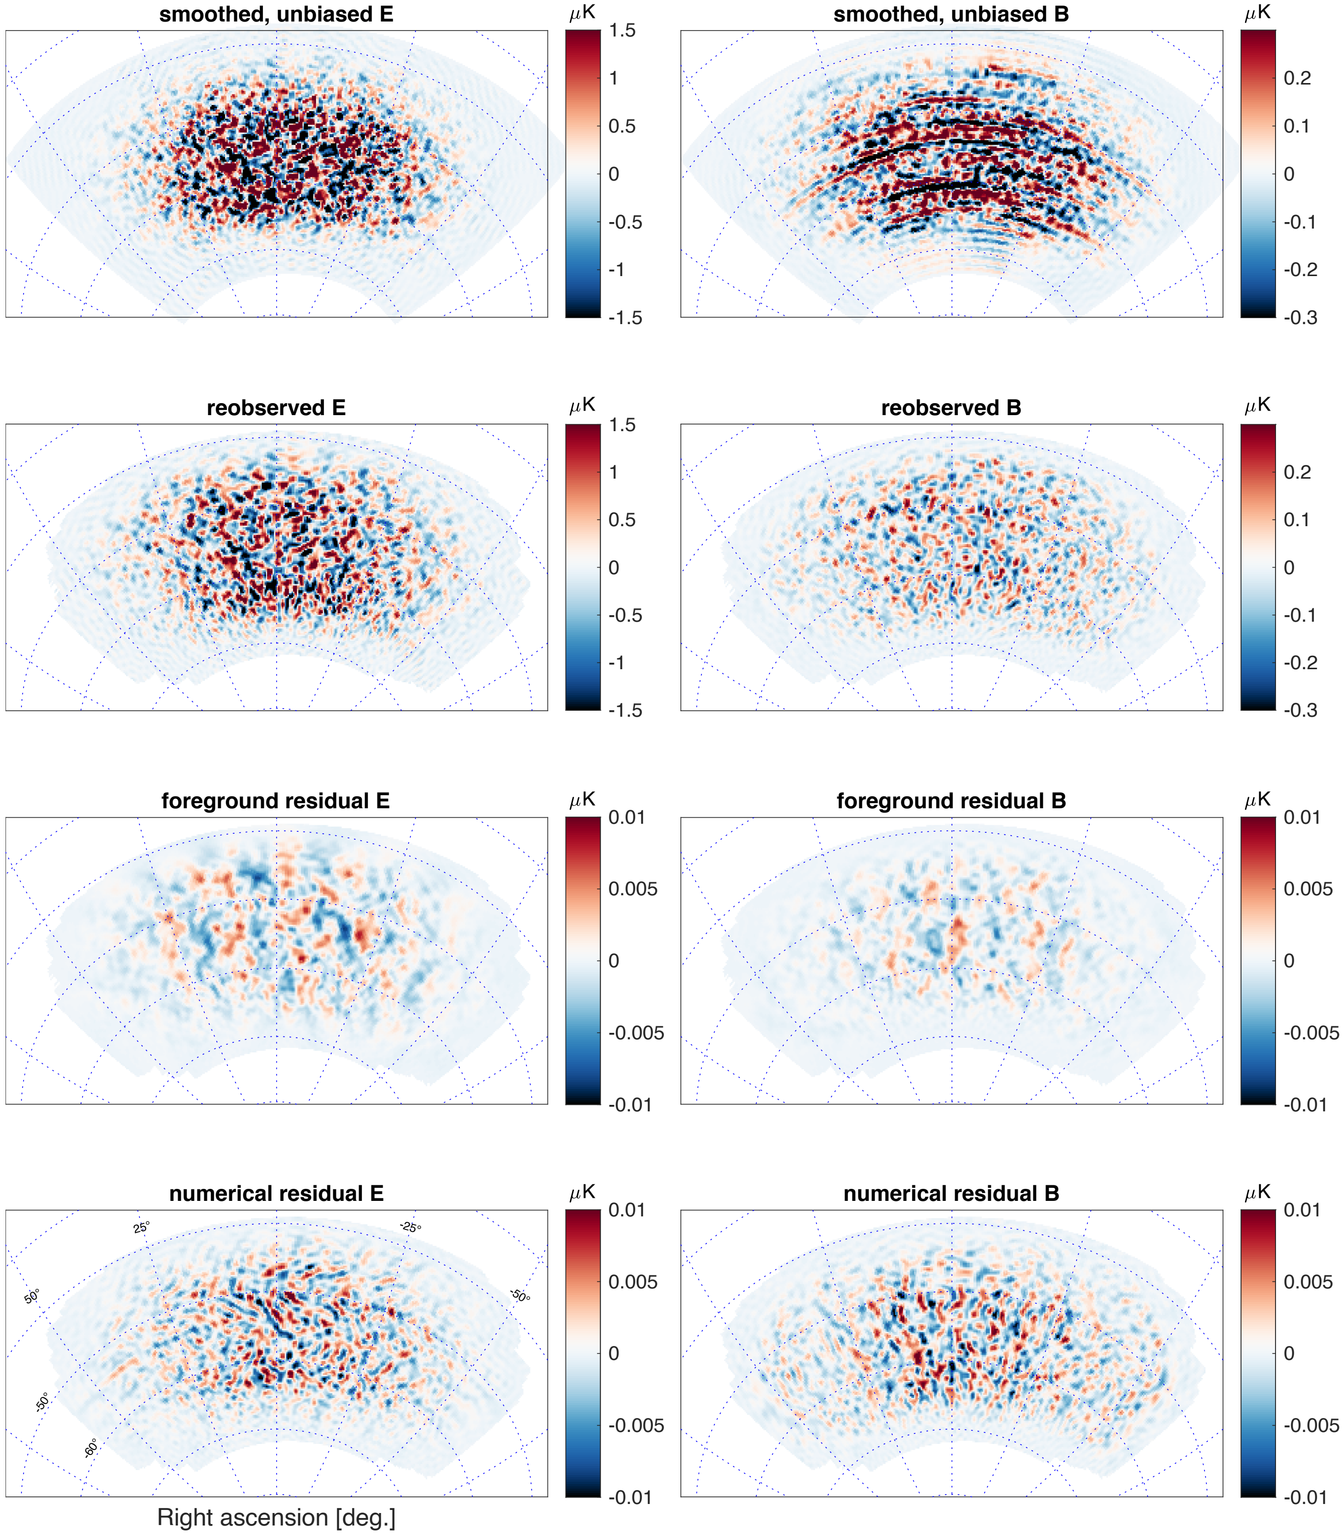

Figure 10. The CMB part of the component map estimator, when run on a simulation, after apodization and *B*-mode purification. First row: the input is a simulated realization of lensed ΛCDM, Galactic dust, and noise. The output has been convolved with a 20′ Gaussian beam, multiplied with a pixel-space weighting, and transformed into *E*- and *B*-mode maps correcting for *E*-to-*B*-leakage effects from the masking (K. M. Smith 2006). Second row: after applying the BICEP3 observing matrix and purifying using the corresponding purification matrix (Keckarray and BICEP2 Collaborations VII 2016). Third row: subtracting the corresponding CMB-and-noise-only simulation from the maps in the previous row reveals the residual from the foreground. Given that the foreground simulations used here are isotropic and Gaussian, the residual is entirely caused by the statistical fluctuation of the β_d fit. Fourth row: to get an estimate of the numerical residual, we subtract the corresponding BICEP3 CMB-only simulation from the CMB component map simulation. This contains effects from the small differences in the filtering done on actual time streams vs. what is incorporated into the observing matrix, and numerical errors coming from the spherical harmonic transforms and the iterative solution method applied to obtain the CMB component map estimate. Note the much-reduced color range in rows three and four.

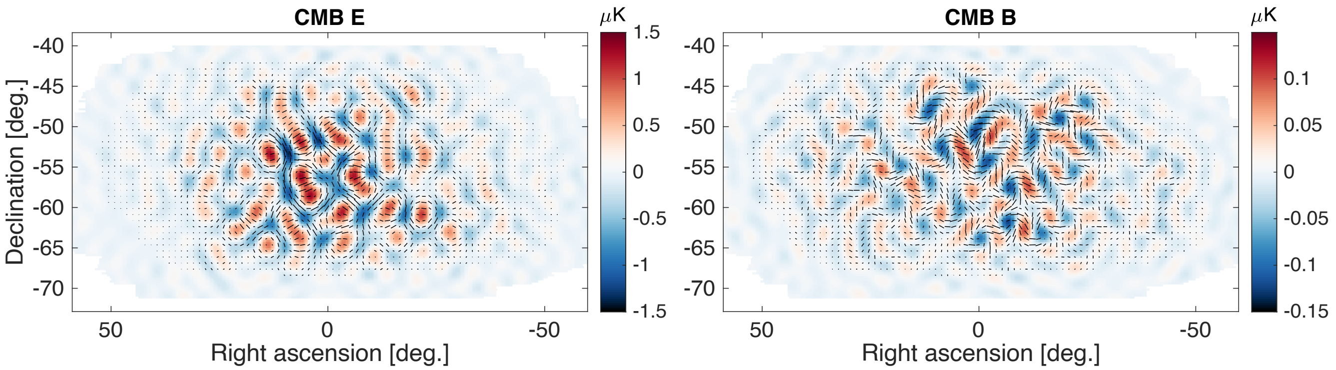

Figure 11. *E*-mode (left) and *B*-mode (right) maximum-likelihood maps of the CMB in CMB units, beam convolved and filtered like a BICEP3 map at 95 GHz with an additional bandpass filter to degree angular scales ($50 < \ell < 120$). Note the differing color ranges. On the left, the *E* map is dominated by the ΛCDM signal, whereas on the right, the *B* map is approximately equal parts lensed-ΛCDM signal and noise.

where $\boldsymbol{w}$ is the pixel weight in Figure 7, and $N_{\ell m}$ is the noise variance per (ℓ, m) mode in Figure 8. In the actual implementation of the quadratic maximum likelihood (QML), we construct the inverse of the total covariance as

$$\boldsymbol{C}^{-1} = (\bar{\boldsymbol{N}}^{-1}\boldsymbol{S} + \boldsymbol{1})^{-1}\bar{\boldsymbol{N}}^{-1},$$

such that we are never required to actually build $\bar{\boldsymbol{N}}$. The inverse is computed using singular value decomposition. In E. F. Bunn & B. Wandelt (2017), a pure-*B* QML estimator was proposed by introducing a free parameter α in the signal covariance matrix:

$$\boldsymbol{S} = \boldsymbol{Y}^{\dagger}\mathrm{diag}(\alpha \boldsymbol{C}_{\ell}^{EE}, \boldsymbol{C}_{\ell}^{BB})\boldsymbol{Y}.$$

In the limit where α tends to infinity, we can write the inverse of this matrix as

$$\boldsymbol{S}^{-1} = \boldsymbol{Y}^{\dagger}\mathrm{diag}(0, 1/\boldsymbol{C}_{\ell}^{BB})\boldsymbol{Y}.$$

With this inverse and the inverse of the noise covariance matrix, we can write the required inverse of the total covariance matrix as

$$\boldsymbol{C}^{-1} = \boldsymbol{S}^{-1}(\bar{\boldsymbol{N}}^{-1} + \boldsymbol{S}^{-1})^{-1}\bar{\boldsymbol{N}}^{-1}.$$

7.2. Simulation Validation

We run the power-spectrum estimators introduced in the previous section on the standard set of 499 BK18 simulations. We test for the level of *E* to *B* leakage by estimating a *B*-mode autopower spectrum on simulation maps that have no *B*-mode power in the input maps. In particular, we use unlensed ΛCDM simulations. The residual *B*-mode power spectra are shown in Figure 14.

We find that simply accounting for the mask-induced mode coupling with the pure-*B* estimator of K. M. Smith (2006) and J. Grain et al. (2009) is insufficient for our sensitivity. *B*-mode purity can be improved with a QML estimator, in addition to the obvious advantage of optimal sensitivity. The purification-matrix-based method performs best in this purification test.

We show the sensitivity to the CMB component for each estimator in Figure 15 in terms of the noise power N_ℓ and effective degrees of freedom. For the simple pseudo-C_ℓ estimator, the noise weighting is highly suboptimal. The other estimators perform very similarly. The noise is only a little elevated compared to the noise in the BICEP3 95 GHz frequency map. For low multipoles, the effective f_{sky} is around 1%–1.5% and hence lies between the effective sky fraction of the larger BICEP3 and the smaller Keck footprints, as expected. This motivates us to select the matrix-based purification method as the baseline for the remainder of this paper, as this method allows us to cut ambiguous modes quite aggressively without sacrificing a noticeable amount of sensitivity.

The CMB component map is supposed to be unbiased, meaning the suppression from time-stream filtering and deprojection should be corrected for at the map level. This is different compared to the baseline approach in the BICEP/Keck analysis, where filtering is corrected for at the power-spectrum level. In Figure 16, we show the residual suppression factor by plotting the ratio between the mean of the output power spectra for signal-only simulations and the expectation in each band power (which is likewise computed from signal simulations that are not run through the component-separation estimator). We find a negligible, subpercent residual signal suppression caused by the component-separation estimator.

We further can estimate the residual dust power in the CMB component map. We distinguish between statistical and systematic residuals (R. Stompor et al. 2016). The former is caused by the statistical uncertainty of the dust model used to build the estimator. Specifically, the estimated β_d scatters around the true simulation-input value of $\beta_d^{\mathrm{true}} = 1.6$. Propagating this scatter to the CMB *B*-mode power spectrum using dust-only simulations as input results in the statistical residuals shown in Figure 17. This statistical residual will be included in the error budget of the CMB *B*-mode band powers and covariance.

We estimate the potential systematic residual, i.e., the residual caused by any mismatch between the true dust model and the model assumed in the estimator, by running the component separation on our suite of more complex alternate dust models (BK18). In Figure 18, we show the systematic residuals for all eight alternate dust models described in Keckarray and BICEP2 Collaborations X (2018) and BK18. The PySM models 1–3 use Planck- and WMAP-derived templates for dust and synchrotron with varying spectral models and added small-scale Gaussian structure, though they overpredict dust levels in the BICEP/Keck field (B. Thorne et al. 2017; The Pan-Experiment Galactic Science Group 2025). The MHDv2 model simulates non-Gaussian dust and synchrotron emission from a 3D Galactic magnetic

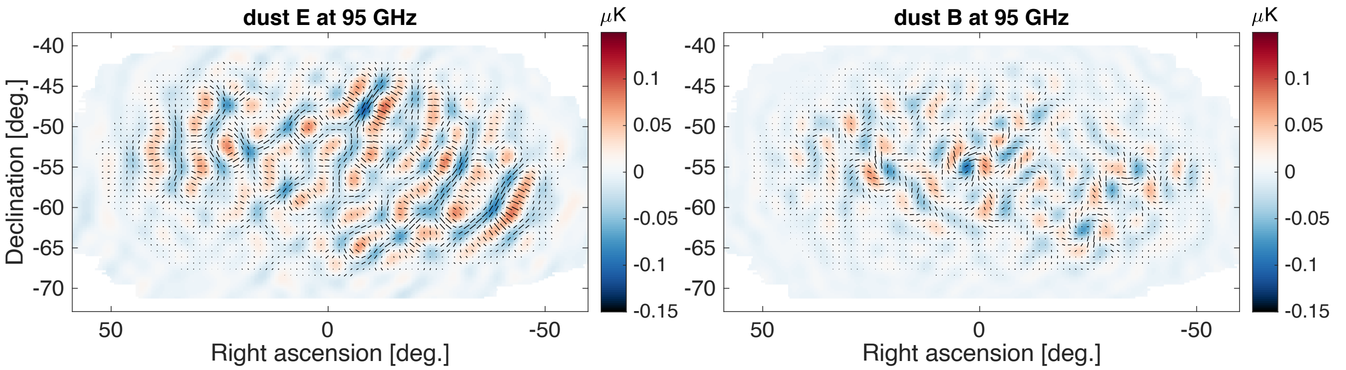

Figure 12. *E*-mode (left) and *B*-mode (right) maximum-likelihood maps of thermal dust in CMB units, beam convolved and filtered like a BICEP3 map at 95 GHz with an additional bandpass filter to degree angular scales ($50 < \ell < 120$). The maps are apodized, and *B* modes are purified. The *E* modes are visibly brighter than the *B* modes.

field (A. G. Kritsuk et al. 2018). The MKD model adds 3D dust structure with varying density, temperature, and spectral index (G. Martínez-Solaeche et al. 2018), while the Vansyngel model constructs Q and U maps by integrating over multiple magnetic field layers with fixed intensity and varying polarization (F. Vansyngel et al. 2018).

The Gaussian-decorrelation and MHDv3 models lead to significantly large residuals. The former incorporates a decorrelation parameter of $\Delta_d = 0.85$ following the parameterization presented in Keckarray and BICEP2 Collaborations X (2018), which is much lower than what is allowed by data in the BICEP/Keck patch (BK18). Hence, such a residual is expected and could be accounted for by marginalizing over a decorrelation parameter in the baseline (BICEP/Keck analysis; S. Azzoni et al. 2023; BK18). The MHDv3 model contains a significant amount of polarized synchrotron emission, about a factor 4 larger than the 95% confidence limit upper limit reported in BK18. Given that we do not presently model a synchrotron component in the map-based component separation, such a residual is expected. All other alternate foreground models, most notably the PySM and Vansyngel models incorporating spatial variation of the dust spectral energy distribution, are well below the 1σ error bars in each band power and comparable to the statistical dust residual.

Finally, we show in Figure 19 the band powers of the real-data CMB and dust-component maps together with the theoretical power spectrum of ΛCDM and the best-fit dust model of BK18. Using 499 simulations, we compute the probability to exceed (PTE) of the χ^2 values between the data band powers and the best-fit model, obtaining PTEs of 0.44 for the CMB component and 0.38 for the dust component. We therefore conclude that these band powers are consistent with the best-fit ΛCDM and dust models, respectively. Notably, this implies that the dust band powers are consistent with a power-law-shaped spectrum.

8. Consistency with Multifrequency Spectra Likelihood

The component-separation method presented in this work allows us to build an alternative pipeline for estimating the tensor-to-scalar ratio r. We check for consistency between the two pipelines at the level of the estimated parameters. In the baseline analysis, we estimate r together with seven foreground parameters using multifrequency power spectra computed from BICEP/Keck, WMAP, and Planck frequency maps (Keckarray and BICEP2 Collaborations X 2018; BK18). In this work, we use a subset of these maps, in particular, the BICEP/Keck and Planck HFI maps, to compute CMB and dust-component maps, from which we can compute CMB and dust B-mode autopower spectra, as well as the cross-spectrum between the two.

We use the parametric likelihood function described in BICEP2/Keck & Planck Collaborations (2015) for the two autopower spectra and one cross-power spectrum computed from the CMB and the dust maps. The band power covariance matrix is derived from 499 simulations of signal and noise, explicitly setting covariances between the CMB and dust signal-only band powers to zero, but allowing for noise covariances between the two components. We fit a one-parameter model consisting of a primordial tensor component with varying amplitude r and a fixed component from CMB lensing, derived from the Planck best-fit ΛCDM cosmology (Planck Collaboration VI 2020), to the CMB autopower spectrum:

$$D_\ell^{BB} = r \cdot D_\ell^{BB\ \text{tensor}} + D_\ell^{BB\ \text{lensing}}. \tag{16}$$

The tensor-to-scalar ratio r is evaluated at a pivot scale of $0.05\ \text{Mpc}^{-1}$. At the same time, we fit a power-law model to the dust autopower spectrum, referenced at a frequency $\nu = 353$ GHz:

$$D_\ell^{BB} = A_\text{d} \cdot \left(\frac{\ell}{80}\right)^{\alpha_\text{d}}. \tag{17}$$

In Figure 20, we show the distribution of best-fit values of r, computed for our set of 499 signal and noise simulations, for three likelihood variations.

1. *Multifrequency.* The baseline likelihood of the BK18 analysis, modified by fixing all synchrotron parameters in the model to their fiducial values and aligning the frequency bands with those used in this paper, to enable a direct comparison with the map-based approach presented here.
2. *Multicomponent.* The likelihood introduced in the paragraph above, based on computing auto- and cross-power spectra between the CMB and dust-component maps.
3. *CMB only.* The same likelihood as above, but using only the CMB autopower spectra.

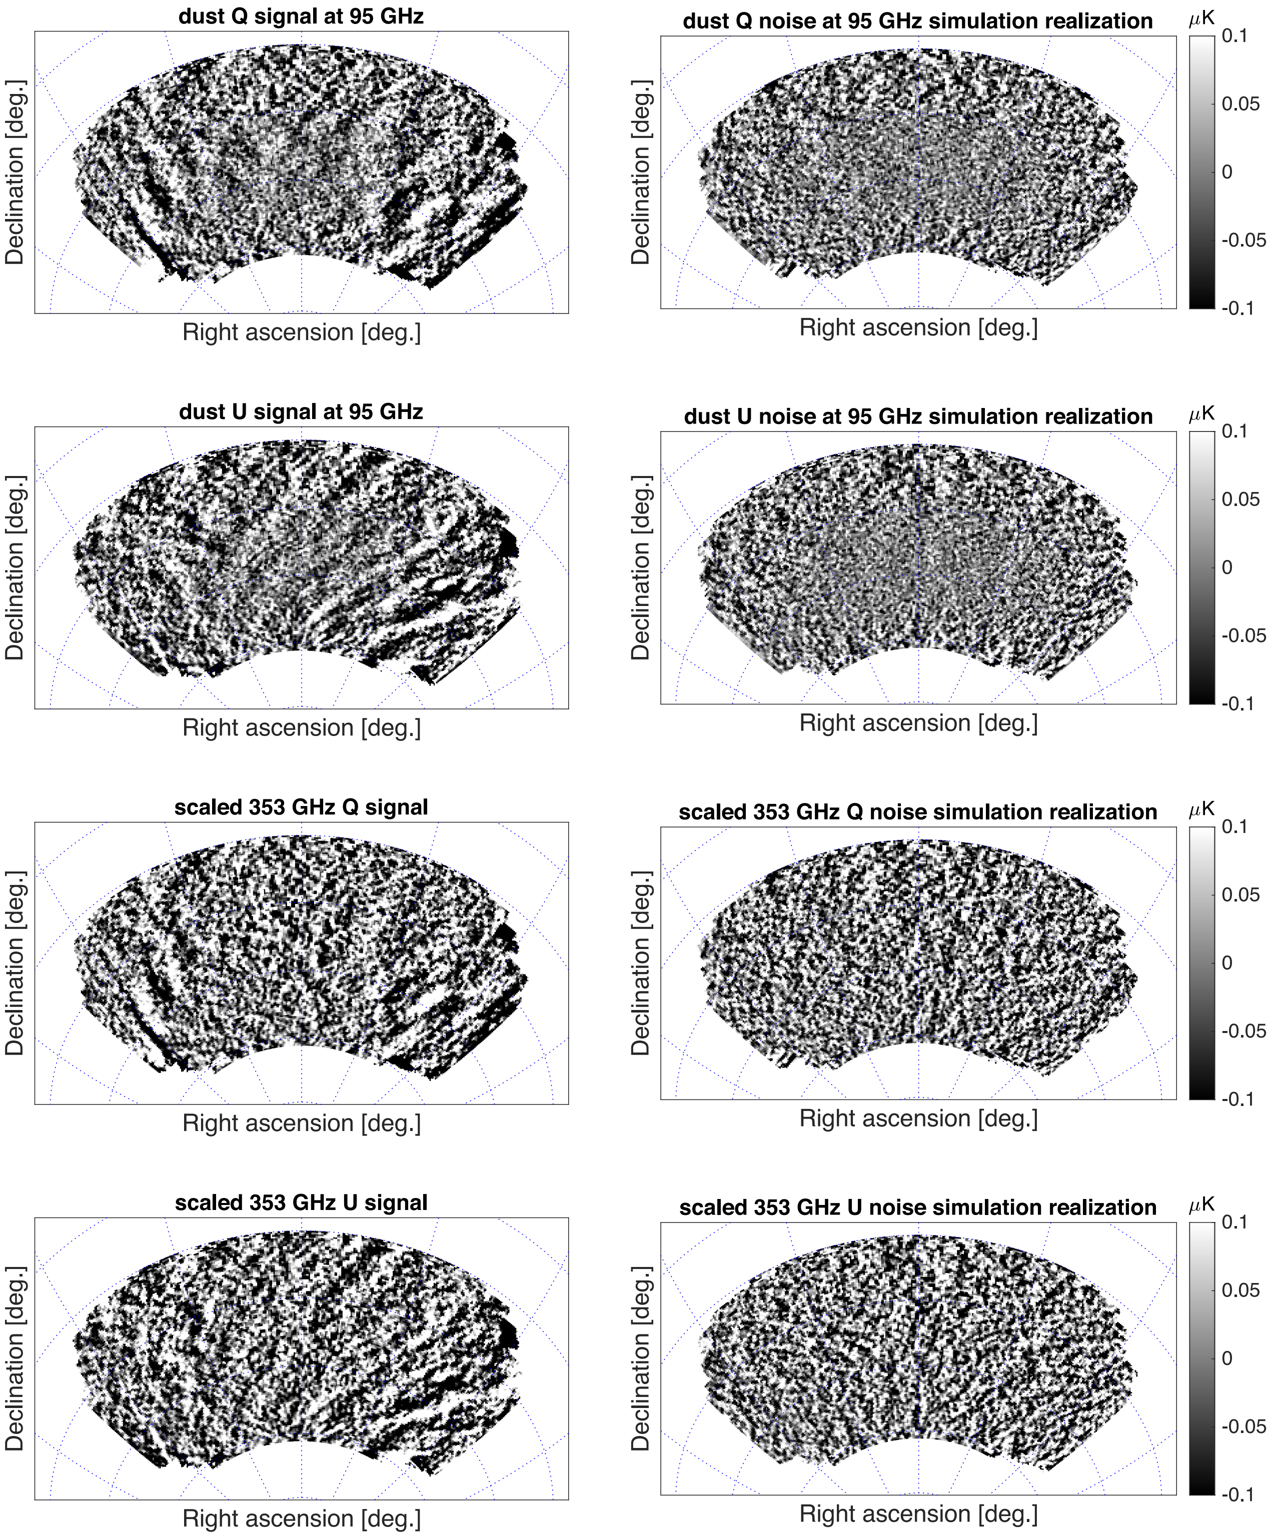

Figure 13. The first and second rows show the derived thermal dust-component maps of this paper for the Q and U signals (left), respectively, compared to a realization of noise (right). The noise is lower in the central BICEP2/Keck region due to the high-frequency data from Keck. Structure in the dust polarization is apparent, particularly in the outer regions of the observing field. The third and fourth rows show equivalent Q and U maps for Planck 353 GHz data (left) and one noise simulation (right), scaled to 95 GHz using an MBB model with $\beta_d = 1.5$. These Planck maps have been convolved to the same resolution as the derived component maps. All maps are reobserved with the BICEP3 observing matrix.

The histograms show an unbiased recovery of the tensor-to-scalar ratio, with comparable sensitivity between the multi-frequency and multicomponent likelihood approaches. Excluding any information about the dust power spectrum, and thus rendering it agnostic to the particular spectral shape of dust, degrades the sensitivity by about 16%.

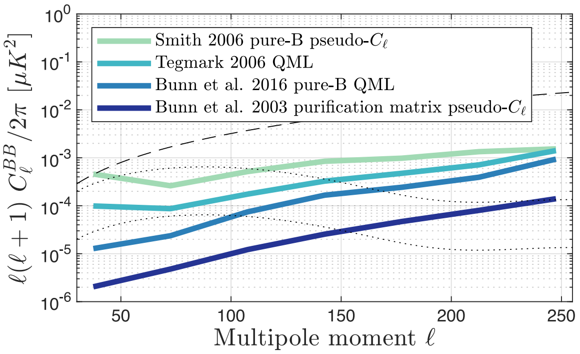

Figure 14. Comparing purification performance for different power-spectrum estimators considered in this work: the simple pseudo-C_ℓ pure-B estimator of K. M. Smith (2006), the optimal QML methods of M. Tegmark (1997) and E. F. Bunn & B. Wandelt (2017), and the purification-matrix-based method described in E. F. Bunn et al. (2003) and used in previous BICEP/Keck analyses (Keckarray and BICEP2 Collaborations VII 2016). The dashed line corresponds to the lensed-B power of our baseline ΛCDM model, while the dotted lines correspond to primordial gravitational-wave signals of $r = 10^{-2}$ and $r = 10^{-3}$.

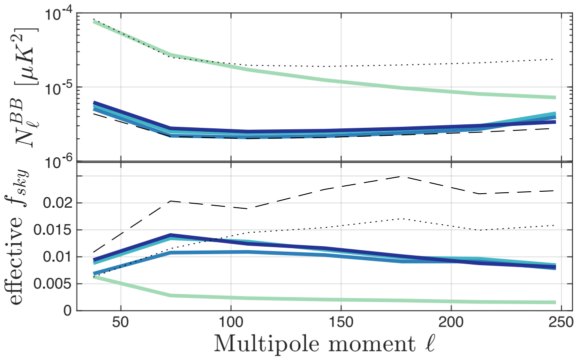

Figure 15. First row: the noise spectra for the different power-spectrum estimators run on the CMB component map. Colors are the same as in Figure 14. The BICEP3 95 GHz noise and the Keck 220 GHz noise are shown in the black dashed and dotted lines, respectively. The spectra are shown after correction for the filtering of the signal, which occurs due to the beam roll-off, time-stream filtering, and B-mode purification. (Note that no ℓ^2 scaling is applied.) Second row: the effective sky fraction as calculated from the ratio of the mean noise realization band powers to their fluctuation $f_{\rm sky}(\ell) = \frac{1}{2\ell\Delta\ell}\left(\frac{\sqrt{2}\bar{N}_b}{\sigma(N_b)}\right)^2$, i.e., the observed number of B-mode degrees of freedom divided by the nominal full-sky number.

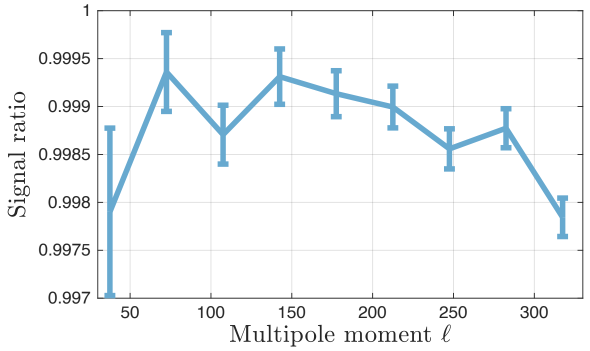

Figure 16. Suppression factor of the CMB component map computed as the ratio between the mean of the power spectra of 499 signal-only simulations and the expectation for each band power. Error bars show the standard errors of the mean.

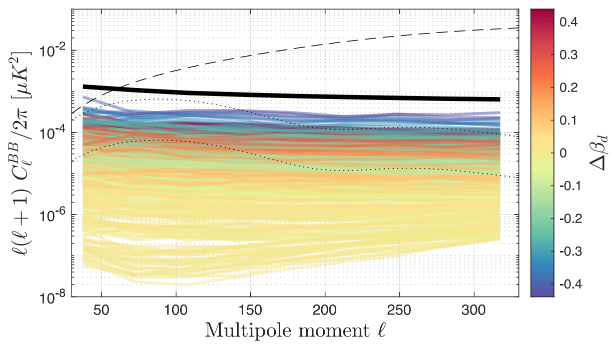

Figure 17. Foreground residual in the CMB component B-mode autopower spectrum for 499 Gaussian-dust-only simulations. The color of the lines shows the deviation of the respective best-fit value of β_d for the specific realization from the simulation input of $\beta_d^{\rm true} = 1.6$. The thick solid black line indicates the dust level at 95 GHz, without any foreground cleaning, as measured in BK18. The dashed line corresponds to the lensed-B power of our baseline ΛCDM model, while the dotted lines correspond to primordial gravitational-wave signals of $r = 10^{-2}$ and $r = 10^{-3}$.

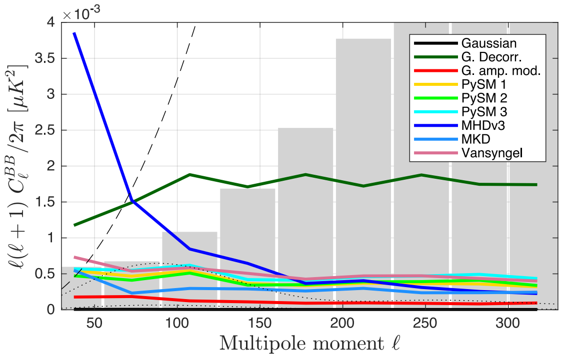

Figure 18. Systematic bias from one realization of the Gaussian and the alternate dust models as described in Keckarray and BICEP2 Collaborations X (2018), BK18, and Section 7.2. The gray bars represent the standard deviation of ΛCDM + dust + noise simulations in each band power. Most models lead to residuals smaller than this scatter, representing cosmic variance. The dashed line corresponds to the lensed-B power of our baseline ΛCDM model.

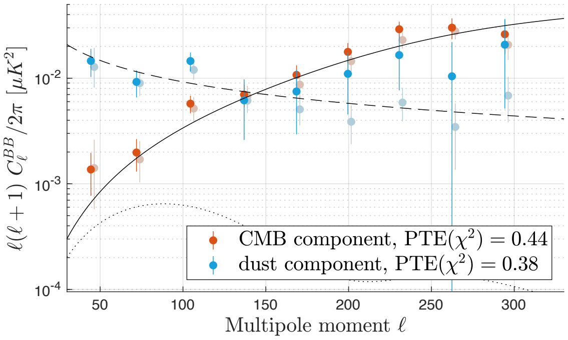

Figure 19. Real-data band powers of the CMB (orange points) and dust (blue points) component maps. The solid black line is the fiducial ΛCDM B-mode power spectrum from CMB lensing, the dotted black line is the primordial B-mode power spectrum for $r = 10^{-2}$, and the dashed black line corresponds to the best-fit dust model of BK18 at 150 GHz. The faint points are derived using the cross-frequency likelihood method, where results represent a full Markov Chain Monte Carlo sampling of the posterior within each individual band power bin; these are equivalent to Figure 16 of BK18 but with marginalization over synchrotron parameters removed. The χ^2 PTE values for the CMB and dust components, comparing the primary data band powers to the ΛCDM and best-fit dust models, are 0.44 and 0.38, respectively.

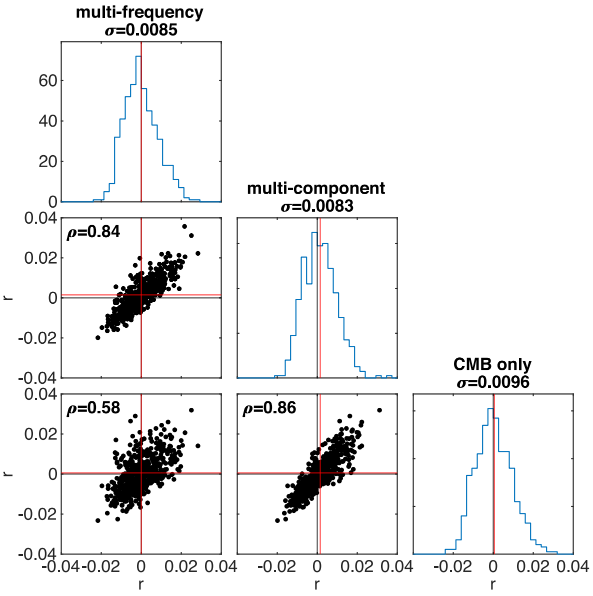

Figure 20. Histograms and 2D scatterplots between best-fit r values obtained with the multifrequency-spectrum-based method of BK18 (with the modification of fixing the synchrotron component in the model and aligning the selected frequency bands with those used in this paper) and the multicomponent-spectrum-based method presented in this paper, either including or excluding the dust-component map in the power-spectrum computation. Black vertical and horizontal lines show the simulation input of $r = 0$, while the corresponding red lines show the distribution mean. The distribution standard deviations are shown in the respective histograms' titles. The Pearson correlation coefficient, ρ, ranges from 58% to 86%.

The Pearson correlation coefficient between these three distributions is 84% for the multifrequency and multicomponent approaches. In F. Bianchini et al. (2025), the correlation between the two methods was found to be higher, indicating that the departure from unit correlation in our analysis arises from the filtering applied to the real data. In the multifrequency approach, each frequency map is filtered differently, whereas in the multicomponent approach, we attempt to undo this filtering, supplementing the missing modes with information from Planck. This results in effective weighting differences, both in pixel and harmonic space, between the two methods. They are ultimately sensitive to slightly different modes. When the dust-component channel is removed from the likelihood, the correlation declines further to 56%.

Given the results presented in this section, we expect the map-based (multicomponent) approach to deliver constraints on the tensor-to-scalar ratio that are consistent with those obtained using the baseline multifrequency method. For the latest and most stringent upper-limit constraints on r, we therefore refer the reader to BK18.

9. Conclusions

We present, for the first time, component-separated maps derived from BICEP and Keck data. Constructing such maps from frequency data with differing beams, filtering, and sky coverage requires accounting for mode coupling at the map level. We develop and validate a method to recover unbiased CMB and Galactic dust-component maps, demonstrating optimal procedures to extract power spectra and cosmological parameters from these maps while carefully treating their nontrivial noise properties. Applying this approach to real data, we present the resulting CMB and dust power spectra and compare constraints on the tensor-to-scalar ratio, r, to those obtained with the baseline multifrequency power-spectrum method of BK18. We find an 84% correlation between the recovered r values, with a comparable $\sigma(r)$ from 499 simulations.

We find that including the dust component in the likelihood is essential for the map-level cleaning approach to achieve r sensitivity comparable to the baseline BK18 method. Omitting the dust channel instead provides a constraint that is agnostic to the assumed dust power-spectrum shape. A further strength of the map-based method lies in the production of high-fidelity CMB and dust-component maps, which reveal the BICEP2/Keck field with striking visual clarity, both in the clean CMB fluctuations and in the detailed Galactic dust structure, and enable analyses beyond simple power-spectrum estimation, such as higher-order statistics and cross-correlation studies. The dust-component map produced here also serves as a high-quality foreground template, which has already proven valuable for assessing the impact of Galactic dust on cosmic birefringence measurements with BICEP3 (BICEP/Keck Collaboration XVIII 2025).

Acknowledgments

The BICEP/Keck projects have been made possible through a series of grants from the National Science Foundation, most recently including 2220444-2220448, 2216223, 1836010, and 1726917. The development of antenna-coupled detector technology was supported by the JPL Research and Technology Development Fund and by NASA grants 06-ARPA206-0040, 10-SAT10-0017, 12-SAT12-0031, 14-SAT14-0009, and 16-SAT-16-0002. The development and testing of focal planes was supported by the Gordon and Betty Moore Foundation at Caltech. Readout electronics were supported by a Canada Foundation for Innovation grant to UBC. Support for quasi-optical filtering was provided by UK STFC grant ST/N000706/1. The computations in this paper were run on the Odyssey/Cannon cluster supported by the FAS Science Division Research Computing Group at Harvard University. The analysis effort at Stanford and SLAC is partially supported by the U.S. DOE Office of Science.

We thank the staff of the U.S. Antarctic Program, and in particular the South Pole Station, without whose help this research would not have been possible. Most special thanks go to our heroic winter-overs during the observing seasons up until 2018: Robert Schwarz, Steffen Richter, Sam Harrison, Grantland Hall, and Hans Boenish. We thank all those who have contributed past efforts to the BICEP/Keck series of experiments, including the BICEPone team. We also thank the Planck and WMAP teams for the use of their data.

We are grateful to Federico Bianchini and Radek Stompor for insightful discussions and to Josquin Errard for valuable comments on the manuscript.

ORCID iDs

Z. Ahmed https://orcid.org/0000-0002-9957-448X
M. Amiri https://orcid.org/0000-0001-6523-9029
D. Barkats https://orcid.org/0000-0002-8971-1954
R. Basu Thakur https://orcid.org/0000-0002-3351-3078
C. A. Bischoff https://orcid.org/0000-0001-9185-6514
D. Beck https://orcid.org/0000-0003-0848-2756
B. Cantrall https://orcid.org/0000-0003-4541-7080

J. R. Cheshire IV https://orcid.org/0000-0002-1630-7854
J. Cornelison https://orcid.org/0000-0002-2088-7345
B. D. Elwood https://orcid.org/0000-0003-4117-6822
S. Fatigoni https://orcid.org/0000-0002-3790-7314
J. P. Filippini https://orcid.org/0000-0001-8217-6832
D. C. Goldfinger https://orcid.org/0000-0001-5268-8423
P. K. Grimes https://orcid.org/0000-0001-9292-6297
G. Halal https://orcid.org/0000-0003-2221-3018
H. Hui https://orcid.org/0000-0001-5812-1903
J. H. Kang https://orcid.org/0000-0002-3470-2954
K. S. Karkare https://orcid.org/0000-0002-5215-6993
J. M. Kovac https://orcid.org/0009-0003-5432-7180
K. Lasko https://orcid.org/0000-0002-4540-1495
K. Lau https://orcid.org/0000-0002-6445-2407
T. Liu https://orcid.org/0000-0001-5677-5188
L. Moncelsi https://orcid.org/0000-0002-4242-3015
S. N. Paine https://orcid.org/0000-0003-4622-5857
M. A. Petroff https://orcid.org/0000-0002-4436-4215
A. R. Polish https://orcid.org/0000-0002-7822-6179
C. Pryke https://orcid.org/0000-0003-3983-6668
B. Singari https://orcid.org/0000-0001-7387-0881
A. Steiger https://orcid.org/0000-0003-0260-605X
C. Tucker https://orcid.org/0000-0002-1851-3918
C. Vergès https://orcid.org/0000-0002-3942-1609
A. Wandui https://orcid.org/0000-0002-8232-7343
J. Willmert https://orcid.org/0000-0002-6452-4693
W. L. K. Wu https://orcid.org/0000-0001-5411-6920
C. Yu https://orcid.org/0000-0002-8542-232X
L. Zeng https://orcid.org/0000-0001-6924-9072
C. Zhang https://orcid.org/0000-0001-8288-5823

References

Aitken, A. C. 1936, RSEPS, 55, 42
Azzoni, S., Alonso, D., Abitbol, M., Errard, J., & Krachmalnicoff, N. 2023, JCAP, 2023, 035
Barrett, R., Berry, M., Chan, T., et al. 1996, MaCom, 64, 1349
Bianchini, F., Beck, D., Wu, W. L. K., et al. 2025, ApJ, 993, 105
BICEP2 Collaboration II 2014, ApJ, 792, 62
BICEP/Keck Collaboration XIII 2021, PhRvL, 127, 151301
BICEP/Keck Collaboration XV 2022, ApJ, 927, 77
BICEP/Keck Collaboration XVI 2023, ApJ, 945, 72
BICEP/Keck Collaboration XVIII 2025, PhRvD, 111, 063505
BICEP2/KeckPlanck Collaborations 2015, PhRvL, 114, 101301
Bunn, E. F., & Wandelt, B. 2017, PhRvD, 96, 043523
Bunn, E. F., Zaldarriaga, M., Tegmark, M., & de Oliveira-Costa, A. 2003, PhRvD, 67, 023501
CMB-S4 Collaboration 2022, ApJ, 926, 54
de Belsunce, R., Gratton, S., & Efstathiou, G. 2022, MNRAS, 517, 2855
Eriksen, H. K., Dickinson, C., Lawrence, C. R., et al. 2006, ApJ, 641, 665
Grain, J., Tristram, M., & Stompor, R. 2009, PhRvD, 79, 123515
Hamimeche, S., & Lewis, A. 2008, PhRvD, 77, 103013
Keckarray and BICEP2 Collaborations V 2015, ApJ, 811, 126
Keckarray and BICEP2 Collaborations VII 2016, ApJ, 825, 66
Keckarray and BICEP2 Collaborations X 2018, PhRvL, 121, 221301
Kovac, J. M., Leitch, E. M., Pryke, C., et al. 2002, Natur, 420, 772
Kritsuk, A. G., Flauger, R., & Ustyugov, S. D. 2018, PhRvL, 121, 021104
Leloup, C., Errard, J., & Stompor, R. 2023, PhRvD, 108, 123547
Martínez-Solaeche, G., Karakci, A., & Delabrouille, J. 2018, MNRAS, 476, 1310
Morshed, M., Rizzieri, A., Leloup, C., Errard, J., & Stompor, R. 2024, PhRvD, 110, 103521
Planck Collaboration IV 2020, A&A, 641, A4
Planck Collaboration LVII 2020, A&A, 643, A42
Planck Collaboration V 2020, A&A, 641, A5
Planck Collaboration VI 2020, A&A, 641, A6
Planck Collaboration XVI 2014, A&A, 571, A16
Sidje, R. B., & Winkles, N. 2011, JCAM, 235, 1956
Smith, K. M. 2006, PhRvD, 74, 083002
Smoot, G. F., Bennett, C. L., Kogut, A., et al. 1992, ApJL, 396, L1
Stompor, R., Errard, J., & Poletti, D. 2016, PhRvD, 94, 083526
Stompor, R., Leach, S., Stivoli, F., & Baccigalupi, C. 2009, MNRAS, 392, 216
Tegmark, M. 1997, PhRvD, 55, 5895
The Pan-Experiment Galactic Science Group 2025, ApJ, 991, 23
Thorne, B., Dunkley, J., Alonso, D., & Naess, S. 2017, MNRAS, 469, 2821
Vansyngel, F., Boulanger, F., Ghosh, T., et al. 2018, A&A, 618, C4
Wolz, K., Azzoni, S., Hervías-Caimapo, C., et al. 2024, A&A, 686, A16
Zebrowski, J. A., Reichardt, C. L., Anderson, A. J., et al. 2025, PhRvD, 112, 123520