

# BICEP/Keck XX: Component-separated maps of polarized CMB and thermal dust emission using *Planck* and BICEP/Keck Observations through the 2018 Observing Season

BICEP/Keck Collaboration: P. A. R. Ade,<sup>1</sup> Z. Ahmed<sup>10</sup>,<sup>2,3</sup> M. Amiri<sup>10</sup>,<sup>4</sup> D. Barkats<sup>10</sup>,<sup>5</sup> R. Basu Thakur<sup>10</sup>,<sup>6</sup> C. A. Bischoff<sup>10</sup>,<sup>7</sup> D. Beck<sup>10</sup>,<sup>8,\*</sup> J. J. Bock,<sup>6,9</sup> H. Boenish,<sup>5</sup> V. Buza,<sup>10</sup> B. Cantrall<sup>10</sup>,<sup>8,2</sup> J. R. Cheshire IV<sup>10</sup>,<sup>6</sup> J. Connors,<sup>11</sup> J. Cornelison<sup>10</sup>,<sup>12</sup> M. Crumrine,<sup>13</sup> A. J. Cukierman,<sup>6</sup> E. Denison,<sup>11</sup> L. Duband,<sup>14</sup> M. Echter,<sup>5</sup> M. Eiben,<sup>5</sup> B. D. Elwood<sup>10</sup>,<sup>15,5</sup> S. Fatigoni<sup>10</sup>,<sup>6</sup> J. P. Filippini<sup>10</sup>,<sup>16,17</sup> A. Fortes,<sup>8</sup> M. Gao,<sup>6</sup> C. Giannakopoulos,<sup>7</sup> N. Goeckner-Wald,<sup>8</sup> D. C. Goldfinger<sup>10</sup>,<sup>8</sup> S. Gratton,<sup>18,19</sup> J. A. Grayson,<sup>8</sup> A. Greathouse,<sup>6</sup> P. K. Grimes<sup>10</sup>,<sup>5</sup> G. Hall,<sup>20</sup>,<sup>8</sup> G. Halal<sup>10</sup>,<sup>8</sup> M. Halpern,<sup>4</sup> E. Hand,<sup>7</sup> S. A. Harrison,<sup>5</sup> S. Henderson,<sup>2,3</sup> T. D. Hoang,<sup>13</sup> J. Hubmayr,<sup>11</sup> H. Hui<sup>10</sup>,<sup>6</sup> K. D. Irwin,<sup>8</sup> J. H. Kang<sup>10</sup>,<sup>6</sup> K. S. Karkare<sup>10</sup>,<sup>21,2,3</sup> S. Kefeli,<sup>6</sup> J. M. Kovac<sup>10</sup>,<sup>15,5</sup> C. Kuo,<sup>8</sup> K. Lasko<sup>10</sup>,<sup>13,20</sup> K. Lau<sup>10</sup>,<sup>6</sup> M. Lautzenhiser,<sup>7</sup> A. Lennox,<sup>17</sup> T. Liu<sup>10</sup>,<sup>8</sup> S. Mackey,<sup>10</sup> N. Maher,<sup>13</sup> K. G. Megerian,<sup>9</sup> L. Minutolo,<sup>6</sup> L. Moncelsi<sup>10</sup>,<sup>6</sup> Y. Nakato,<sup>8</sup> H. T. Nguyen,<sup>9,6</sup> R. O'Brient,<sup>9,6</sup> S. N. Paine,<sup>5</sup> A. Patel,<sup>6</sup> M. A. Petroff<sup>10</sup>,<sup>5</sup> A. R. Polish<sup>10</sup>,<sup>15,5</sup> T. Prouve,<sup>14</sup> C. Pryke<sup>10</sup>,<sup>13</sup> C. D. Reintsema,<sup>11</sup> S. Richter,<sup>5</sup> T. Romand,<sup>6</sup> M. Salatino,<sup>8</sup> A. Schillaci,<sup>6</sup> B. Schmitt,<sup>5</sup> R. Schwarz,<sup>13</sup> C. D. Sheehy,<sup>13</sup> B. Singari<sup>10</sup>,<sup>13,20</sup> A. Soliman,<sup>9,6</sup> T. St Germaine,<sup>5</sup> A. Steiger<sup>10</sup>,<sup>6</sup> B. Steinbach,<sup>6</sup> R. Sudiwala,<sup>1</sup> G. P. Teply,<sup>6</sup> K. L. Thompson,<sup>8</sup> C. Tucker<sup>10</sup>,<sup>1</sup> A. D. Turner,<sup>9</sup> C. Vergès<sup>10</sup>,<sup>22,5</sup> A. G. Vieregg,<sup>10</sup> A. Wandui<sup>10</sup>,<sup>6</sup> A. C. Weber,<sup>9</sup> J. Willmert<sup>10</sup>,<sup>13</sup> C. L. Wong,<sup>5,15</sup> W. L. K. Wu<sup>10</sup>,<sup>3,2,8</sup> H. Yang,<sup>8</sup> C. Yu<sup>10</sup>,<sup>10</sup> L. Zeng<sup>10</sup>,<sup>5</sup> C. Zhang<sup>10</sup>,<sup>8</sup> and S. Zhang<sup>6</sup>

<sup>1</sup>School of Physics and Astronomy, Cardiff University, Cardiff, CF24 3AA, United Kingdom

<sup>2</sup>Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, Stanford, CA 94305, USA

<sup>3</sup>SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

<sup>4</sup>Department of Physics and Astronomy, University of British Columbia, Vancouver, British Columbia, V6T 1Z1, Canada

<sup>5</sup>Center for Astrophysics, Harvard & Smithsonian, Optical and Infrared Astronomy Division, Cambridge, MA 02138, USA

<sup>6</sup>Department of Physics, California Institute of Technology, 1200 E. California Boulevard, Pasadena, CA 91125, USA

<sup>7</sup>Department of Physics, University of Cincinnati, Cincinnati, OH 45221, USA

<sup>8</sup>Department of Physics, Stanford University, Stanford, CA 94305, USA

<sup>9</sup>Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA

<sup>10</sup>Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA

<sup>11</sup>National Institute of Standards and Technology, Boulder, CO 80305, USA

<sup>12</sup>Argonne National Laboratory, High Energy Physics Division, Lemont, IL 60439, USA

<sup>13</sup>School of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455, USA

<sup>14</sup>Service des Basses Températures, Commissariat à l'Energie Atomique, 38054 Grenoble, France

<sup>15</sup>Department of Physics, Harvard University, Cambridge, MA 02138, USA

<sup>16</sup>Department of Physics, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

<sup>17</sup>Department of Astronomy, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

<sup>18</sup>Centre for Theoretical Cosmology, DAMTP, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, UK

<sup>19</sup>Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge CB3 0HA, UK

<sup>20</sup>Minnesota Institute for Astrophysics, University of Minnesota, Minneapolis, MN 55455, USA

<sup>21</sup>Department of Physics, Boston University, Boston, MA 02215, USA

<sup>22</sup>Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

(Dated: September 29, 2025)

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.

\*Electronic address: dobeck@stanford.edu

## I. INTRODUCTION

Measurements of the CMB have been transformative for our understanding of the Universe, leading to the establishment of the  $\Lambda$ CDM mode 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 [1] and the detection of CMB polarization, another two orders of magnitude smaller than the total intensity fluctuations [2].

Significant effort has been made towards 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 two orders of magnitude larger than the sought-after signal: Galactic foregrounds and weak gravitational lensing.

The most sensitive constraint on large-scale B-mode polarization comes from the BICEP/Keck collaboration [3], referred to as “BK18” in the following, setting a 95% upper-limit on the level of primordial B-mode power approximately one 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 Ref. [4], which compares data and theory at the level of power spectra, or bandpowers, is widely used in current and upcoming analyses of experiments [5–8].

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.

## II. 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 Refs. [9–11], this framework produces unbiased maps of the CMB and dust signals.

### A. Observation model

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

$$\mathbf{d} = \begin{pmatrix} \tilde{\mathbf{m}}_1 \\ \vdots \\ \tilde{\mathbf{m}}_N \end{pmatrix}. \quad (1)$$

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

$$\tilde{\mathbf{m}} = \begin{pmatrix} \tilde{\mathbf{m}}^Q \\ \tilde{\mathbf{m}}^U \end{pmatrix}, \quad (2)$$

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

The timestream filtering and deprojection of BICEP/Keck frequency maps,  $\tilde{\mathbf{m}}_\nu$ , can be modeled as a linear operation represented by the observing matrix  $\mathbf{R}$ . 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 [12]. This matrix can act on beam-convolved input maps to produce simulated maps as they would be observed by a BICEP/Keck frequency band  $\nu$

$$\tilde{\mathbf{m}}_\nu = \mathbf{R}\mathbf{B}\mathbf{m}_\nu + \mathbf{n}_\nu, \quad (3)$$

where  $\mathbf{B}$  is an operator convolving the  $Q$  and  $U$  maps with their respective symmetric beam functions,  $b_\ell$ . Both 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,  $\mathbf{n}$ .

### B. Multifrequency model

For previous BICEP/Keck results, the baseline multifrequency model assumes three components: CMB, Galactic dust, and Galactic synchrotron emission [5, 13]. Given that the current data set is limited to 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 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

$$\mathbf{m} = \begin{pmatrix} \mathbf{m}_1 \\ \vdots \\ \mathbf{m}_N \end{pmatrix} = \begin{pmatrix} \mathbf{1} & \mathbf{f}_1 \\ \vdots & \vdots \\ \mathbf{1} & \mathbf{f}_N \end{pmatrix} \begin{pmatrix} \mathbf{s}^{\text{CMB}} \\ \mathbf{s}^{\text{dust}} \end{pmatrix} \equiv \mathbf{F}\mathbf{s}, \quad (4)$$

where  $\mathbf{1}$  is the  $2n_p \times 2n_p$  identity matrix and  $\mathbf{f}_i \equiv f_i \mathbf{1}$  are scaling factors multiplied to that identity matrix. We assume a modified black-body spectrum for Galactic dust such that the scaling factor for the  $i$ th frequency band is given by

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

where  $b_i(\nu)$  is the respective bandpass.

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

$$\tilde{\mathbf{m}} = \mathbf{RBFs} + \mathbf{n} \equiv \mathbf{As} + \mathbf{n}. \quad (6)$$

### C. 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 timestream filtering and binning into maps as Gaussian random fluctuations, which can therefore be fully described in terms of their noise covariance matrix

$$\hat{\mathbf{N}} = \langle \mathbf{n} \mathbf{n}^T \rangle. \quad (7)$$

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

### D. Likelihood formalism

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

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

where both the foreground model, represented by the parameter  $\beta_d$ , and the sky signal  $\mathbf{s}$  are unknown. We will make use of the maximization of the profile likelihood  $-2 \log P(\mathbf{s}, \beta_d = \beta_d^* | \mathbf{d})$ , where we fix the parameter in the foreground model in  $\mathbf{A}$  to the value  $\beta_d^*$ . The maximization of this profile likelihood yields the generalized least squares estimator [14]

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

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

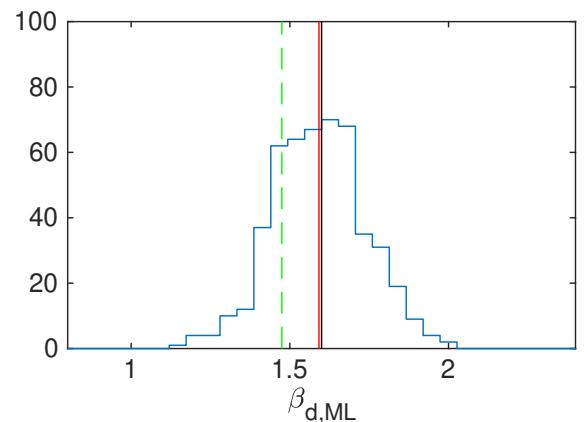


FIG. 1: Histogram of the maximum likelihood values of the dust spectral index  $\beta_d$  obtained in the baseline BK18 [3] auto-/cross-spectrum analysis. The black vertical line indicates the input value to the simulations, the (nearly coincident) red line indicates the mean of the recovered best-fit values, and the dashed green line marks the best-fit value of the real data.

### E. Foreground model

This study leverages the best-fit foreground parameters derived from the mainline BK18 analysis [3]. 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 CMB, dust, and instrumental noise. The distribution of the resulting best-fit values of the dust spectral index is shown in Fig. 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  $\beta_d$  for each simulation realization in the component map estimator in Eq. 9.

Possible extensions of this approach would be the inclusion of a consistent map-based fitting of  $\beta_d$  and  $\beta_s$  in a dedicated two-step approach following for example Ref. [11, 15]. This formalism also more directly allows for fitting spatially varying foreground parameters at map level, which would make this pipeline more robust against biases due to complex foregrounds. Moreover, Refs. [16, 17] have recently introduced and implemented a non-parametric maximum-likelihood framework for CMB foreground separation. We leave an exploration of these options for future work.

### F. 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 Eq. 8

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

Here,  $\mathbf{S}$  represents the signal covariance matrix, which incorporates the auto-correlations of CMB, dust, and synchrotron. These components are modeled in harmonic space using parametric representations of their respective power spectra. Marginalizing this likelihood over the sky signal  $\mathbf{s}$  and applying the arguments of Ref. [4] yields the baseline BK18 auto- and cross-spectra likelihood.

In practice, modeling the data at the level of multi-frequency spectra can lead to complications. In BK18, 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 BICEP2/*Keck* maps and a large-field mask for BICEP3 and external maps [3].

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 bandpowers 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.

### III. DATA

The BK18 dataset consists of observations from the BICEP2, *Keck Array*, and BICEP3 receivers located at the South Pole Station in Antarctica. The BICEP2 receiver observed at 150 GHz from 2010–2012 [18] with  $\approx 500$  bolometric detectors. The *Keck Array* 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 [19]. BICEP3 is a single, scaled up receiver which started science observations in 2016 [20] with  $\approx 2500$  detectors.

BICEP2 and *Keck Array* both mapped a region of sky centered at RA 0h, Dec.  $-57.5^\circ$  with an effective area of  $\approx 400$  square degrees [13]. BICEP3 has a larger instantaneous field of view and hence naturally maps a larger sky area with an effective area of  $\approx 600$  square degrees. This results in small-field maps at

95/150/220 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 it into maps of equirectangular pixelization, with  $0.25^\circ$  square pixels at declination  $-57.5^\circ$ . The map binning operation weights by the inverse variance of the timestream 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 [21]. As opposed to the mainline analysis of BK18, these maps are not reobserved prior to using them in the estimator.

## IV. IMPLEMENTATION

### A. Maximum-likelihood estimator

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

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

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  $\hat{\mathbf{s}}$ . This leads to a significant speed-up given that  $\mathbf{R}$ ,  $\mathbf{B}$  and  $\mathbf{F}$  are sparse operators either in pixel or harmonic spaces and thanks to fast routines available to transform between the two spaces. Fig. 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 sub-diagonals within each frequency-channel  $Q$  and  $U$  blocks. The beam convolution operator  $\mathbf{B}$  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  $\mathbf{F}$ -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 Eq. 9 with a preconditioned iterative method [22], employing a block Jacobi preconditioner. The use of

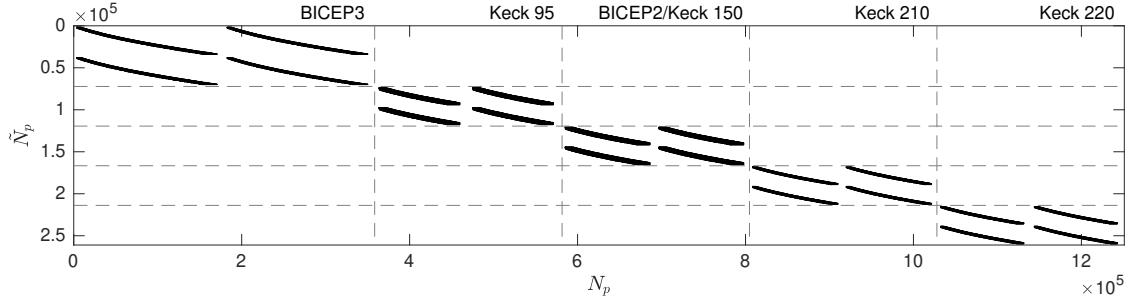


FIG. 2: Plot of the non-zero matrix elements of the BICEP/Keck internal maps' observing matrices, i.e., the BICEP/Keck-specific block of the matrix  $\mathbf{R}$  used in this work. Applying a vector including beam-convolved  $Q$  &  $U$  maps to the right of this matrix results in a vector of filtered  $Q$  &  $U$  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  $Q$  &  $U$  maps. This part of the matrix contains about 5 billion non-zero elements and is the biggest computational challenge in this analysis.

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

$$\mathbf{M} \equiv \mathbf{F}^T \hat{\mathbf{N}}^{-1} \mathbf{F}. \quad (12)$$

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.

## B. Numerical experiments

To focus on the problem of inverting the system matrix on the right-hand side of Eq. 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

$$\hat{\mathbf{s}}_{B3} = \left( \mathbf{R}^T \hat{\mathbf{N}}_{B3}^{-1} \mathbf{R} \right)^{-1} \mathbf{R}^T \hat{\mathbf{N}}_{B3}^{-1} \mathbf{d}_{B3}. \quad (13)$$

We know that this problem is mathematically ill-defined since the matrix  $\mathbf{R}^T \hat{\mathbf{N}}^{-1} \mathbf{R}$  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 has

negligible filtering suppression, by solving the estimator

$$\hat{\mathbf{s}}_{95 \text{ GHz}} = \left( \mathbf{R}^T \hat{\mathbf{N}}_{B3}^{-1} \mathbf{R} + \hat{\mathbf{N}}_P^{-1} \right)^{-1} \left( \mathbf{R}^T \hat{\mathbf{N}}_{B3}^{-1} \mathbf{d}_{B3} + \hat{\mathbf{N}}_P^{-1} \mathbf{d}_P \right). \quad (14)$$

In Fig. 3 we show the  $Q$  maps and the  $EE$  two-dimensional auto-power spectrum of a noise simulation for the BICEP3 and *Planck* 100 GHz input maps, as well as their combination  $\hat{\mathbf{s}}_{95 \text{ GHz}}$ . It illustrates how this estimator fills in the filtered-out modes in the BICEP3 map with noisier *Planck* modes in the “poly-trench” along the  $\ell_y$  direction. This causes characteristic stripes in the combined map.

In Fig. 4 we explore three different preconditioned iterative solvers for this problem applied to signal-and-noise maps: the Conjugate Gradient method (CG), the Generalized Minimal Residual method (GMRES), and the Biconjugate Gradient Stabilized method (Bi-CGSTAB) [22]. We find the latter to have superior convergence behavior and will use it in the following work.

In Fig. 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 simulations, the iterative solver converges much faster.

In Fig. 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 Eq. 13 will take a long time to converge, and a faithful estimate is not guaranteed due to the non-invertibility of

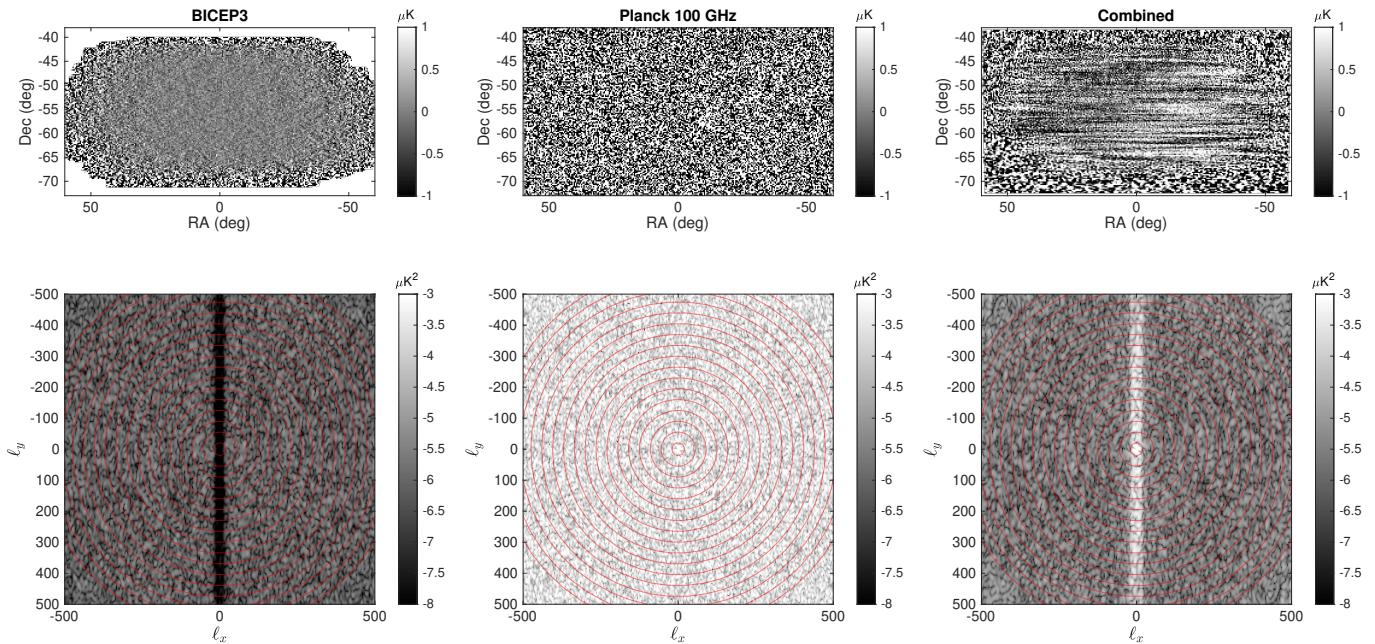


FIG. 3:  $Q$  maps (top row) and the  $EE$  two-dimensional auto-power spectrum (bottom row) of one noise simulation for BICEP3 (left column), *Planck* HFI 100 GHz (middle column), and their combination (right column). The combination estimator corrects for filtering suppression at the map-level and hence boosts the noise compared to the filtered BICEP3 map and fills in modes from *Planck* for small  $\ell_x$ . This is why the combined  $Q$  noise maps show strong horizontal stripes in the central BICEP3 map region.

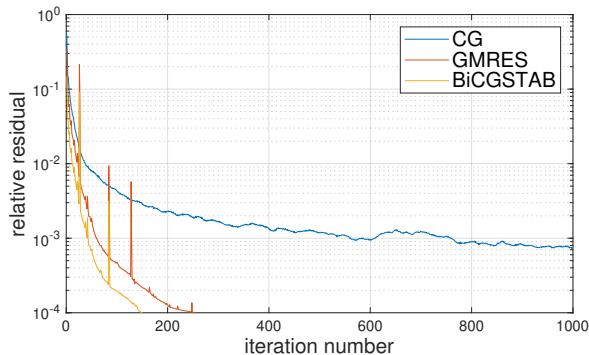


FIG. 4: The relative residual  $\|\mathbf{As} - \mathbf{b}\|/\|\mathbf{b}\|$ , where  $\mathbf{A} \equiv \mathbf{R}^T \hat{\mathbf{N}}_{B3}^{-1} \mathbf{R} + \hat{\mathbf{N}}_P^{-1}$  and  $\mathbf{b} = \mathbf{R}^T \hat{\mathbf{N}}_{B3}^{-1} \mathbf{d}_{B3} + \hat{\mathbf{N}}_P^{-1} \mathbf{d}_P$ , for each iteration of the preconditioned iterative method. We test three different iterative solvers: the classic CG, GMRES, and BiCGSTAB [22]. The spikes are due to numerical noise, which these iterative solvers are susceptible to [23].

the system matrix. The regularization with an external *Planck* map as described in Eq. 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 Sec. III to solve for a CMB component map reaches this criterion in  $\mathcal{O}(70)$  steps, noting that a larger data vector leads to larger runtime per iteration and higher memory requirement.

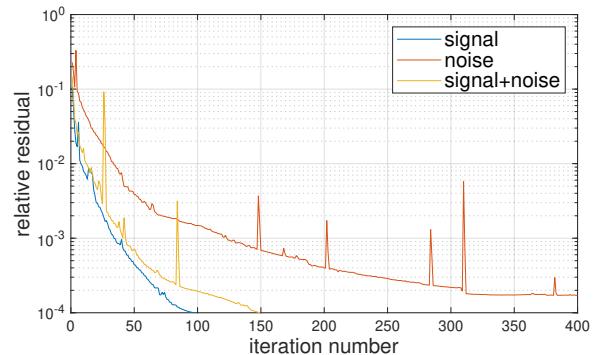


FIG. 5: The relative residual as defined in Fig. 4 for each iteration of the (Bi-)CGSTAB preconditioned conjugate gradient method. We show the convergence performance for a signal-only, noise-only, and a signal-and-noise simulation.

## V. NOISE PROPERTIES

The output of the estimator,  $\hat{\mathbf{s}}$ , 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

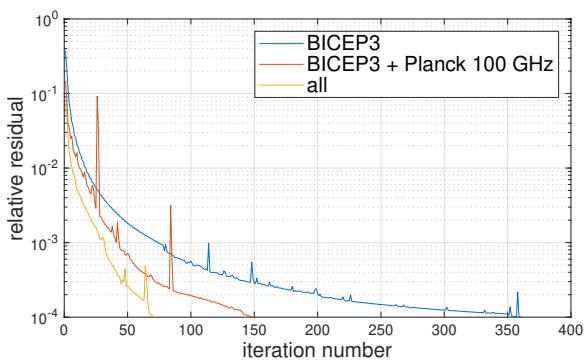


FIG. 6: The relative residual as defined in the caption of Fig. 4 for each PCG iteration step. We show a comparison to solve for the frequency-map solution for BICEP3 and the BICEP3+*Planck* combination, as well as the CMB component map solution given all input frequency maps as described in Sec. III.

maps of the  $Q$  and  $U$  variance in the CMB and dust component maps by taking the diagonal of

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

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 Ref. [3]. These weights are shown in Fig. 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  $(\ell, m)$  mode yields Figs. 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 declination: the higher the maximal declination, the higher the  $m$  modes with non-zero 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.

## VI. MAPS

In this section we will present simulated and real data CMB and dust component maps for BICEP, *Keck*, and *Planck* HFI data. For the reasons explained in the previous section, any plot of the map and computation of a power spectrum requires some weighting in the  $(\ell, m)$  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

$$\mathbf{m}^{\text{CMB/dust}} = \mathbf{R}_{\text{B3}} \cdot \hat{\mathbf{s}}^{\text{CMB/dust}}$$

This method also allows us to use the well-established pipeline from observed BICEP/*Keck* maps, to purified power spectra, to estimates of  $r$  [3, 12].

In Fig. 10 we show the CMB part of the component map estimator when running on a simulated realization of lensed  $\Lambda$ CDM, 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  $\beta_d$  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 signal 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 Figs. 11 and 12, respectively. Further, we show the dust component  $Q$  and  $U$  maps in comparison with a noise realization side-by-side in Fig. 13. The high-signal-to-noise filamentary dust structures are clearly visible in  $Q$  and  $U$ , consistent with features previously identified in external neutral-hydrogen data [24] as well as SPT-3G [25].

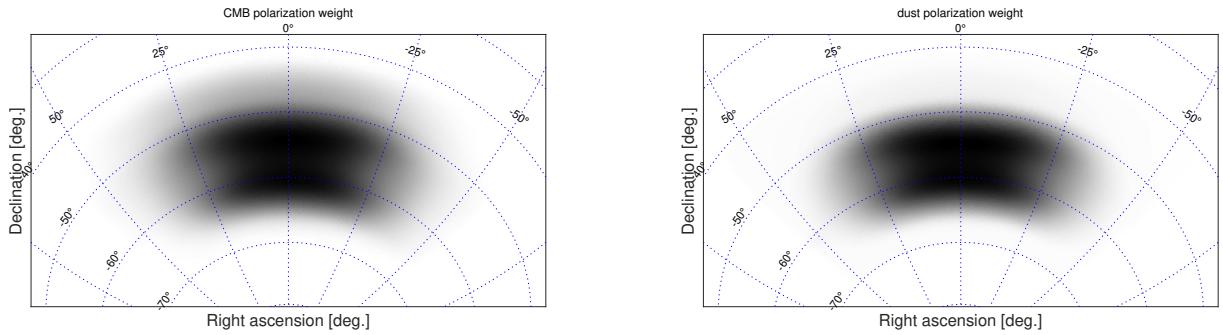


FIG. 7: The polarization weight defined as the inverse arithmetic mean of the Q and U noise variance assuming white noise per pixel for the CMB component (left) and the dust component (right) in arbitrary units. Due to the extended coverage of BICEP3, the CMB extends above  $50^\circ$  in declination,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

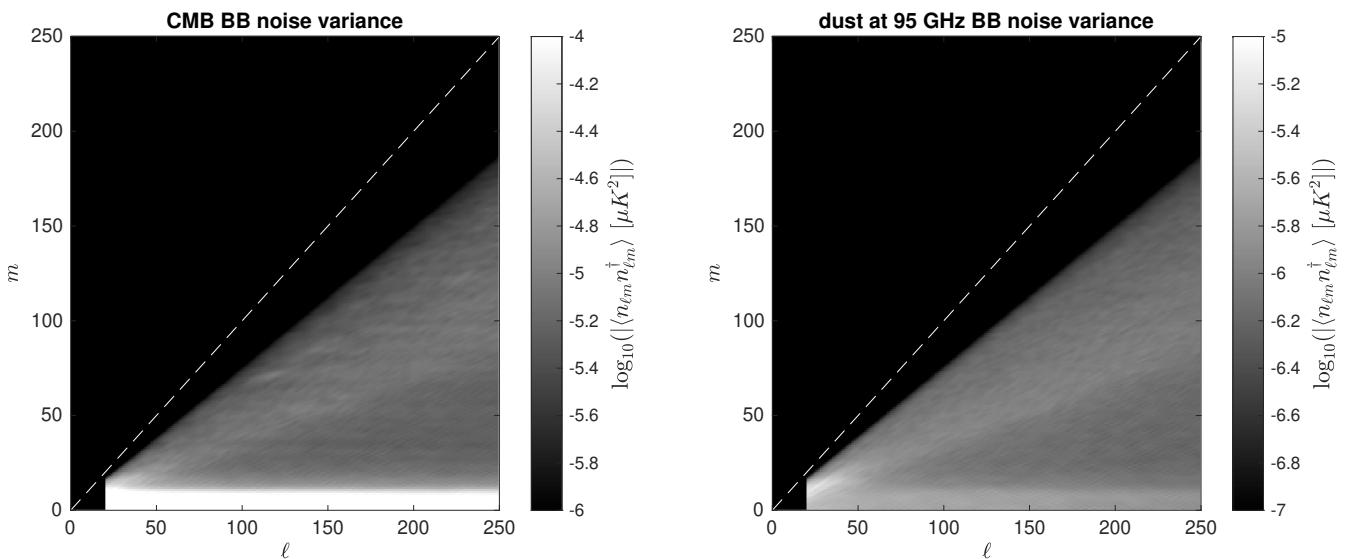


FIG. 8: The noise variance per  $(\ell, m)$  mode of the CMB component map obtained by averaging over 499 noise-only simulations, after applying the corresponding map-level weight shown in Fig. 7 before computing harmonic coefficients.

FIG. 9: Same as Fig. 8 but for the dust component map scaled to 95 GHz.

## A. Methods

### 1. Pseudo- $C_\ell$ method

In order to use the pseudo- $C_\ell$  method, which is the standard method for estimating power spectra of BICEP/Keck maps, 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- $C_\ell$ s of matrix-purified maps as outlined in Ref. [12].

## VII. 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.

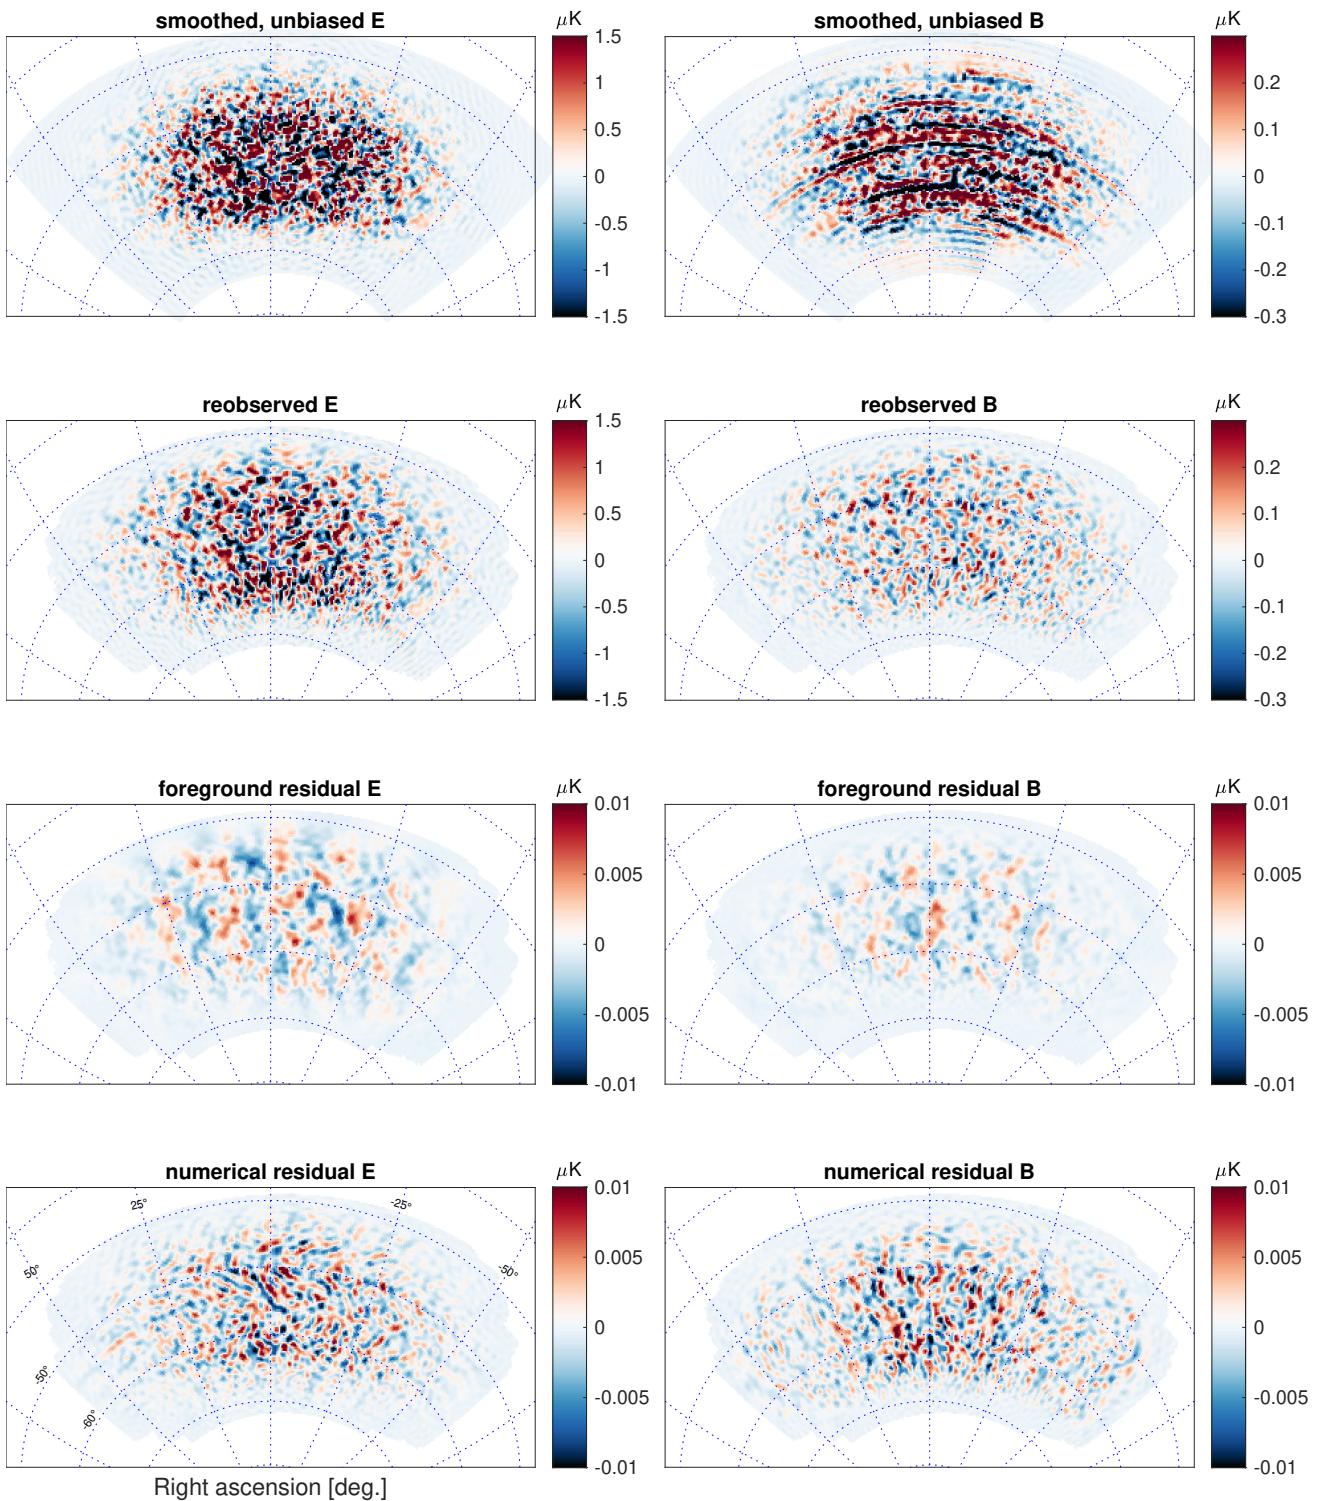


FIG. 10: The CMB part of the component map estimator when run on a simulation, after apodization and  $B$ -mode purification. *First row:* The input is a simulated realization of lensed  $\Lambda$ CDM, galactic dust and noise. The output has been convolved with a 20 arcmin Gaussian beam, multiplied with a pixel-space weighting and transformed into  $E$  and  $B$ -mode maps correcting for  $E$ -to- $B$ -leakage effects from the masking [26].

*Second row:* After applying the BICEP3 observing matrix and purifying using the corresponding purification matrix [12]. *Third row:* Subtracting the corresponding CMB-and-noise-only simulation from the maps in the previous row reveals the residual from foreground. Given that the foreground simulations used here are isotropic and Gaussian, the residual is entirely caused by statistical fluctuation of the  $\beta_d$  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 timestreams versus 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.

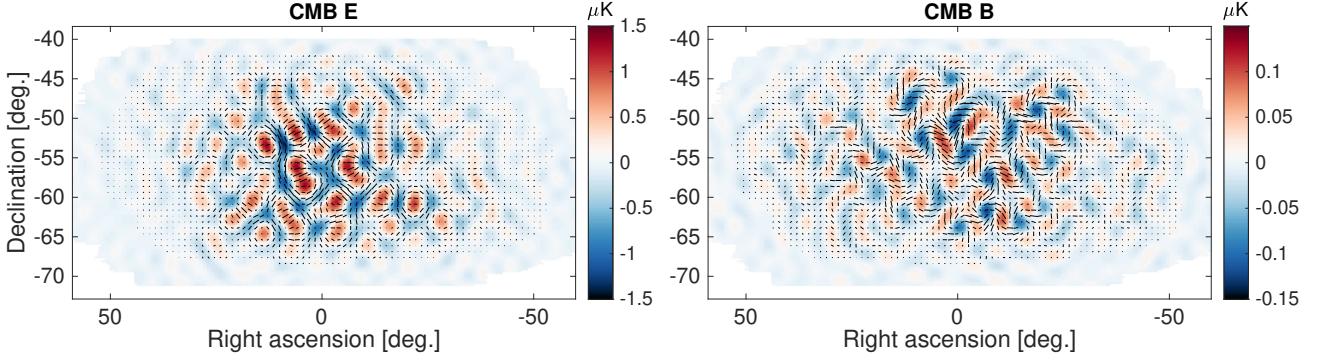


FIG. 11:  $E$ -mode (left) and  $B$ -mode (right) maximum-likelihood maps of the CMB in CMB units, beam-convolved and filtered like a BICEP3 map at 95 GHz with an additional bandpass filter to degree angular scales ( $50 < \ell < 120$ ). Note the differing color ranges; on the left, the  $E$  map is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

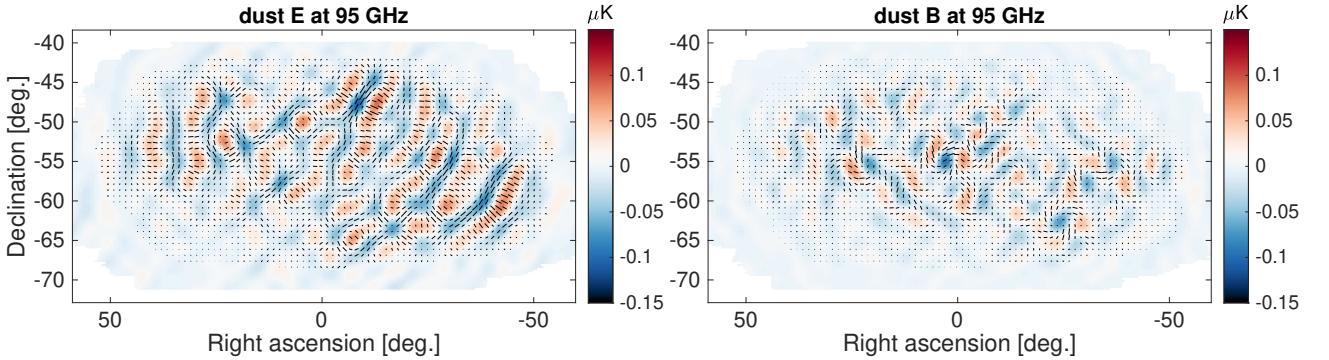


FIG. 12:  $E$ -mode (left) and  $B$ -mode (right) maximum-likelihood maps of thermal dust in CMB units, beam-convolved and filtered like a BICEP3 map at 95 GHz with an additional bandpass filter to degree angular scales ( $50 < \ell < 120$ ). The maps are apodized and  $B$ -modes are purified. The  $E$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

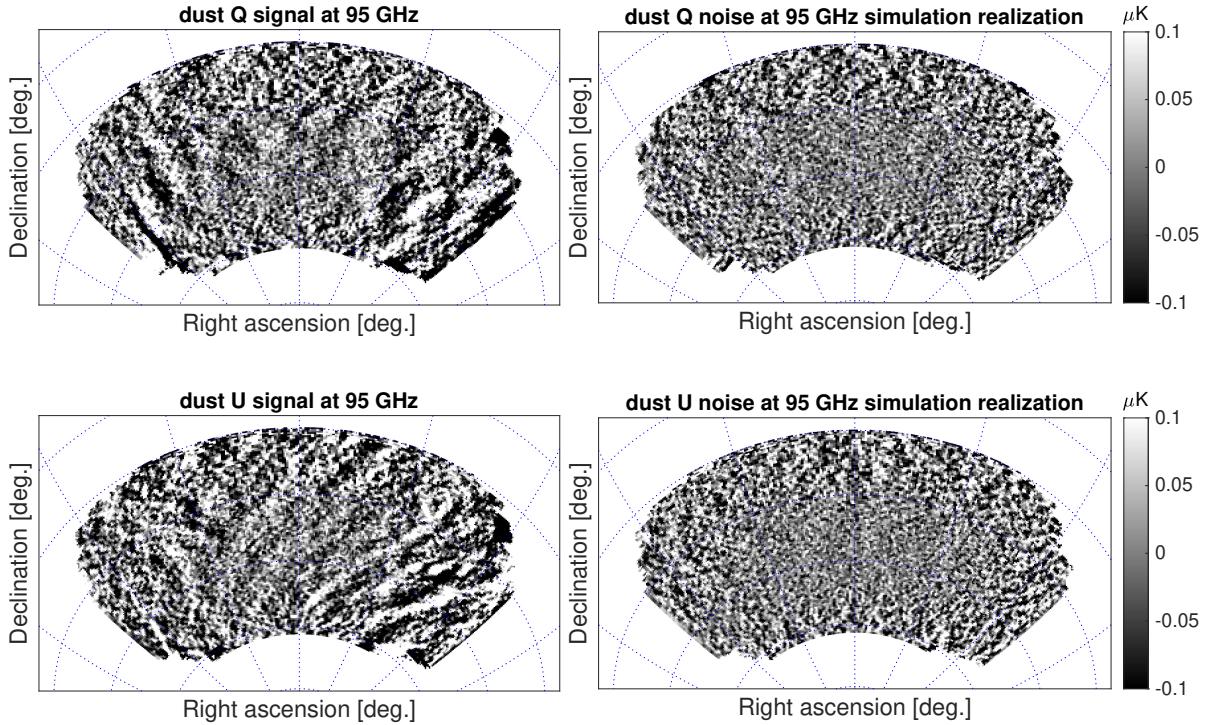


FIG. 13: Comparing the thermal dust Q (top) and U (bottom) signal on the left to a realization of noise on the right. The noise is lower in the central BICEP2/*Keck* region due to the high-frequency data from *Keck*. Structure in the dust polarization becomes apparent, especially in the outer parts of the observing field.

## 2. Optimal quadratic maximum likelihood (QML) method

Following Refs. [27, 28], we can construct an optimal power spectrum estimator as the solution of a likelihood maximization by estimating a given bandpower at bin  $b$  as

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

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

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

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

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

We model the signal covariance to be diagonal in harmonic space, given our baseline fiducial cosmological model derived from *Planck* 2013 cosmological parameters [29]

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

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

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

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

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

such that we are never required to actually build  $\bar{\mathbf{N}}$ . The inverse is computed using singular value decomposition. In Ref. [28], a pure-B QML estimator was proposed by introducing a free parameter  $\alpha$  in the signal covariance matrix

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

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

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

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

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

## B. Simulation validation

We run the power spectrum estimators introduced in the previous section on the standard set of 499 BK18

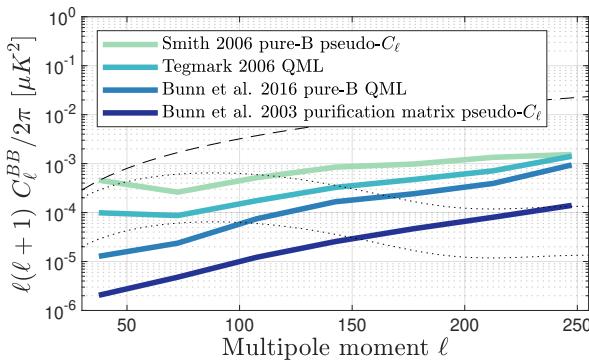


FIG. 14: Comparing purification performance for different power spectrum estimators considered in this work: a simple pseudo- $C_\ell$  pure- $B$  estimator of Ref. [26], optimal quadratic maximum likelihood (QML) methods of Refs. [27] and [28], and a purification-matrix-based method described in Ref. [31] and used in previous BICEP/Keck analyses [12]. The dashed line corresponds to the lensed- $B$  power of our baseline  $\Lambda$ CDM model, while the dotted lines correspond to a primordial gravitational wave signal of  $r = 10^{-2}$  and  $r = 10^{-3}$ .

simulations. We test for the level of  $E$ -to- $B$  leakage by estimating a  $B$ -mode auto-power spectrum on simulation maps which have no  $B$ -mode power in the input maps. In particular, we use unlensed  $\Lambda$ CDM simulations. The residual  $B$ -mode power spectra are shown in Fig. 14.

We find that simply accounting for the mask-induced mode-coupling with the pure- $B$  estimator of Ref. [26, 30] 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 Fig. 15 in terms of the noise power  $N_\ell$  and effective degrees of freedom. For the simple pseudo- $C_\ell$  estimator, the noise weighting is highly sub-optimal. 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 timestream 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 Fig. 16 we show the

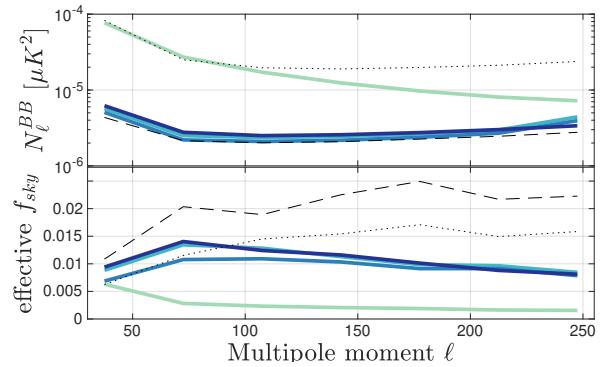


FIG. 15: *First row:* The noise spectra for the different power-spectrum estimators run on the CMB component map. Colors are the same as in Fig. 14. The BICEP3 95 GHz noise and the Keck 220 GHz noise are shown in the black dashed and dotted lines, respectively. The spectra are shown after correction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

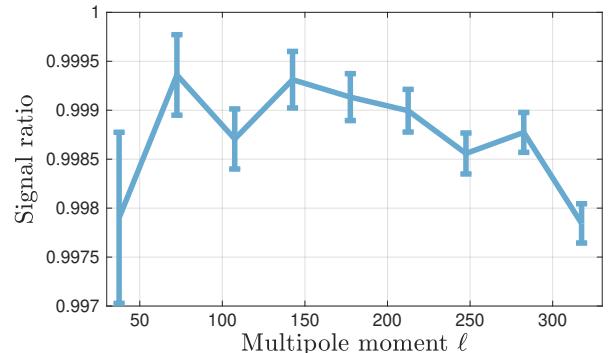


FIG. 16: Suppression factor of the CMB component map computed as the ratio between the mean of power spectra of 499 signal-only simulations and the expectation for each band-power. Error bars show the standard error of this mean.

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, sub-percent 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 residual [32]. The former is caused by the statistical uncertainty of the dust model

used to build the estimator. Specifically, the estimated  $\beta_d$  scatters around the true simulation-input value of  $\beta_d^{\text{true}} = 1.6$ . Propagating this scatter to the CMB  $B$ -mode power spectrum using dust-only simulations as input results in the statistical residuals shown in Fig. 17. This statistical residual will be included in the error budget of the CMB  $B$ -mode bandpowers 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 [3]. In Fig. 18 we show the systematic residuals for all eight alternate dust models described in Refs. [3, 13]. 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 [33, 34]. The MHDv2 model simulates non-Gaussian dust and synchrotron emission from a 3D Galactic magnetic field [35]. The MKD model adds 3D dust structure with varying density, temperature, and spectral index [36], while the Vansyngel model constructs Q/U maps by integrating over multiple magnetic field layers with fixed intensity and varying polarization [37].

The Gaussian-decorrelation and MHDv3 models lead to significantly large residuals. The former incorporates a decorrelation parameter of  $\Delta_d = 0.85$  following the parametrization presented in Ref. [13], which is much lower than what is allowed by data in the BICEP/*Keck* patch [3]. Hence such a residual is expected and could be accounted for by marginalizing over a decorrelation parameter in the baseline BICEP/*Keck* analysis [3, 38]. The MHDv3 model contains a significant amount of polarized synchrotron emission, about a factor four larger than the 95% C.L. upper limit reported in Ref. [3]. 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 SED, are well below the  $1\sigma$  error bars in each bandpower and comparable to the statistical dust residual.

Finally, we show in Fig. 19 the bandpowers of the real data CMB and dust component maps together with the theory power spectrum of  $\Lambda\text{CDM}$  and the best-fit dust model of BK18 [3]. Using 499 simulations, we compute the probability-to-exceed (PTE) of the  $\chi^2$  values between the data bandpowers 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 bandpowers are consistent with the best-fit  $\Lambda\text{CDM}$  and dust models, respectively. Notably, this implies that the dust bandpowers are consistent with a power-law shaped spectrum.

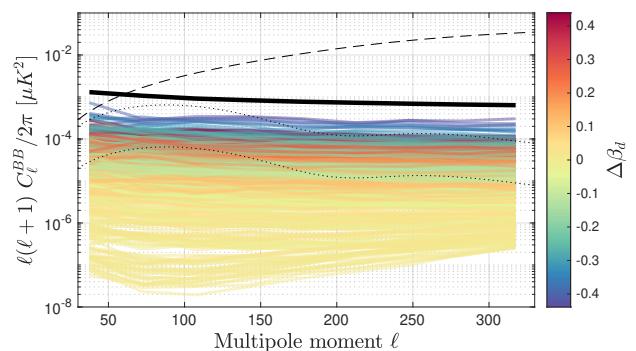


FIG. 17: Foreground residual in the CMB component  $B$ -mode auto-power spectrum for 499 Gaussian-dust-only simulations. The color of the lines show the deviation of the respective best-fit value of  $\beta_d$  for the specific realization from the simulation input of  $\beta_d^{\text{true}} = 1.6$ . The thick solid black line indicates the dust level at 95 GHz, without any foreground cleaning, as measured in Ref. [3]. The dashed line corresponds to the lensed- $B$  power of our baseline  $\Lambda\text{CDM}$  model while the dotted lines correspond to a primordial gravitational wave signal of  $r = 10^{-2}$  and  $r = 10^{-3}$ .

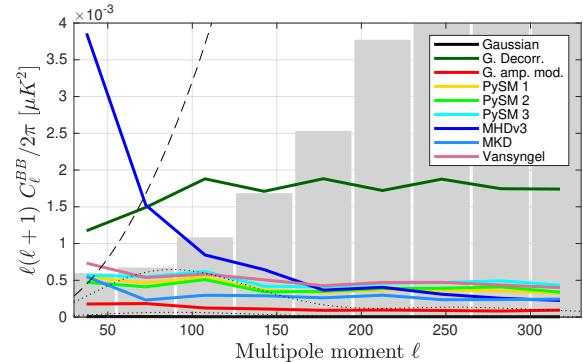


FIG. 18: Systematic bias from one realization of the Gaussian and the alternate dust models as described in Refs. [3, 13] and Sec. VII B. The gray bars represent the standard deviation of  $\Lambda\text{CDM} + \text{dust} + \text{noise}$  simulations in each bandpower. Most models lead to residuals smaller than this scatter representing cosmic variance. The dashed line corresponds to the lensed- $B$  power of our baseline  $\Lambda\text{CDM}$  model.

## VIII. CONSISTENCY WITH MULTI-FREQUENCY-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 estimated parameters. In the baseline analysis, we estimate  $r$  together with seven foreground parameters using multi-frequency power spectra computed from

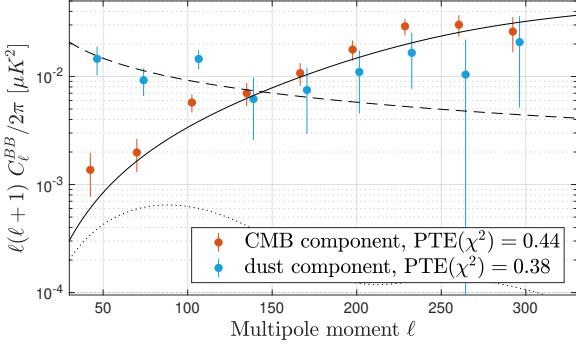


FIG. 19: Real data bandpowers of the CMB (orange points) and dust (blue points) component maps. The solid black line is the fiducial  $\Lambda$ CDM B-mode power spectrum of CMB lensing, the dotted black line is the primordial B-mode power spectrum for  $r = 10^{-2}$ , and the dashed black line corresponds to the best-fit dust model of Ref. [3] at 150 GHz. This figure is similar to Fig. 16 of BK18 [3]. The PTEs of the  $\chi^2$  values for the CMB and dust components, comparing real data bandpowers to the  $\Lambda$ CDM model and the best-fit dust model of BK18 [3], are 0.44 and 0.38, respectively.

BICEP/Keck, WMAP, and *Planck* frequency maps [3, 13]. 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 auto-power spectra, as well as the cross-spectrum between the two.

We use the parametric likelihood function described in Ref. [5] for the two auto-power spectra and one cross-power spectrum computed from the CMB and the dust map. The bandpower covariance matrix is derived from 499 simulations of signal and noise, explicitly setting covariances between the CMB and dust signal-only bandpowers 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  $\Lambda$ CDM cosmology [39], to the CMB auto-power spectrum

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

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 auto-power spectrum, referenced at a frequency  $\nu = 353 \text{ GHz}$ ,

$$D_\ell^{BB} = A_d \cdot \left( \frac{\ell}{80} \right)^{\alpha_d}. \quad (17)$$

In Fig. 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:

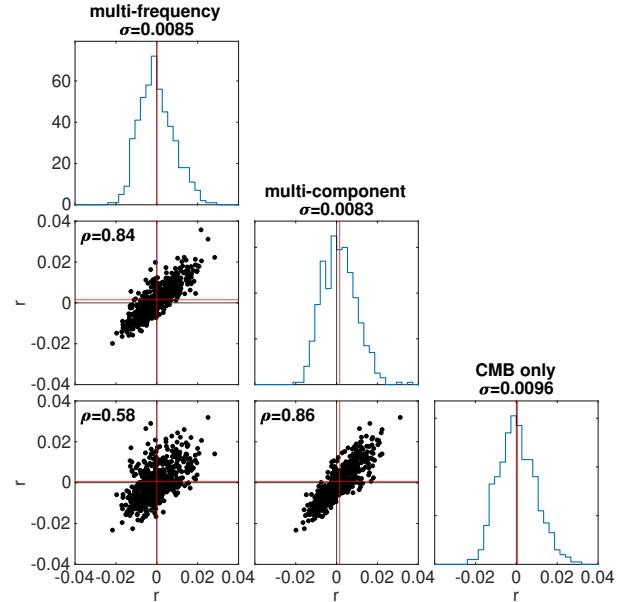


FIG. 20: Histograms and 2D scatter plots between best-fit  $r$  values obtained with the multi-frequency-spectrum-based method of Ref. [3] (with the modification of fixing the synchrotron component in the model and aligning the selected frequency bands with those used in this paper) and the multi-component-spectrum-based 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 histogram's title. The Pearson correlation coefficient,  $\rho$ , ranges from 58% to 86%.

- “multi-frequency”: The baseline likelihood of the BK18 analysis [3], 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,
- “multi-component”: The likelihood introduced in the paragraph above, based on computing auto- and cross-power spectra between the CMB and dust component maps,
- “CMB only”: The same likelihood as above, but using only the CMB auto-power spectra.

The histograms show an unbiased recovery of the tensor-to-scalar ratio, with comparable sensitivity between the multi-frequency and multi-component 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%.

The Pearson correlation coefficient between these three distributions is 84% for the multi-frequency and multi-component approaches. In [40], 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 multi-frequency approach, each frequency map is filtered differently, whereas in the multi-component 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 (multi-component) approach to deliver constraints on the tensor-to-scalar ratio that are consistent with those obtained using the baseline multi-frequency method. For the latest and most stringent upper-limit constraints on  $r$ , we therefore refer the reader to BK18 [3].

## IX. 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 non-trivial 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 multi-frequency power-spectrum method of Ref. [3]. We find an 84% correlation between 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 [3] 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 [41].

## 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, & 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 BICEP1 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.

---

[1] G. F. Smoot, C. L. Bennett, A. Kogut, E. L. Wright, J. Aymon, N. W. Boggess, E. S. Cheng, G. de Amici, S. Gulkis, M. G. Hauser, et al., *Astrophys. J. Lett.* **396**, L1 (1992).

[2] J. M. Kovac, E. M. Leitch, C. Pryke, J. E. Carlstrom, N. W. Halverson, and W. L. Holzapfel, *Nature* **420**, 772–787 (2002), ISSN 1476-4687, URL <http://dx.doi.org/10.1038/nature01269>.

[3] BICEP/Keck Collaboration XIII, Phys. Rev. Lett. **127** (2021), 2110.00483, URL <https://doi.org/10.1103/2Fphysrevlett.127.151301>.

[4] S. Hamimeche and A. Lewis, Phys. Rev. D **77** (2008), URL <https://doi.org/10.1103%2Fphysrevd.77.103013>.

[5] BICEP2/Keck, Planck Collaborations, Phys. Rev. Lett. **114** (2015), 1502.00612, URL <https://doi.org/10.1103%2Fphysrevlett.114.101301>.

[6] Planck Collaboration V, Astron. Astrophys. **641**, A5 (2020), ISSN 1432-0746, 1907.12875, URL <http://dx.doi.org/10.1051/0004-6361/201936386>.

[7] CMB-S4 Collaboration, Astrophys. J. **926**, 54 (2022), ISSN 1538-4357, 2008.12619, URL <http://dx.doi.org/10.3847/1538-4357/ac1596>.

[8] K. Wolz, S. Azzoni, C. Hervías-Caimapo, J. Errard, N. Krachmalnicoff, D. Alonso, C. Baccigalupi, A. Baleato Lizancos, M. L. Brown, E. Calabrese, et al., Astron. Astrophys. **686**, A16 (2024), ISSN 1432-0746, 2302.04276, URL <http://dx.doi.org/10.1051/0004-6361/202346105>.

[9] H. K. Eriksen, C. Dickinson, C. R. Lawrence, C. Baccigalupi, A. J. Banday, K. M. Gorski, F. K. Hansen, P. B. Lilje, E. Pierpaoli, M. D. Seiffert, et al., The Astrophysical Journal **641**, 665–682 (2006), ISSN 1538-4357, URL <http://dx.doi.org/10.1086/500499>.

[10] Planck Collaboration IV, Astron. Astrophys. **641**, A4 (2020), ISSN 1432-0746, URL <http://dx.doi.org/10.1051/0004-6361/201833881>.

[11] R. Stompor, S. Leach, F. Stivoli, and C. Baccigalupi, Mon. Not. R. Astron. Soc. **392**, 216 (2009), URL <https://doi.org/10.1111%2Fj.1365-2966.2008.14023.x>.

[12] Keck Array and BICEP2 Collaborations VII, Astrophys. J. **825**, 66 (2016), 1603.05976.

[13] Keck Array and BICEP2 Collaborations X, Phys. Rev. Lett. **121**, 221301 (2018), 1810.05216.

[14] A. C. Aitken, Proceedings of the Royal Society of Edinburgh **55**, 42–48 (1936).

[15] R. de Belsunce, S. Gratton, and G. Efstathiou, Mon. Not. R. Astron. Soc. **517**, 2855–2866 (2022), ISSN 1365-2966, URL <http://dx.doi.org/10.1093/mnras/stac2825>.

[16] C. Leloup, J. Errard, and R. Stompor, Phys. Rev. D **108** (2023), ISSN 2470-0029, 2308.03412, URL <http://dx.doi.org/10.1103/PhysRevD.108.123547>.

[17] M. Morshed, A. Rizzieri, C. Leloup, J. Errard, and R. Stompor, Phys. Rev. D **110** (2024), ISSN 2470-0029, 2405.18365, URL <http://dx.doi.org/10.1103/PhysRevD.110.103521>.

[18] BICEP2 Collaboration II, Astrophys. J. **792**, 62 (2014), 1403.4302.

[19] Keck Array and BICEP2 Collaborations V, Astrophys. J. **811**, 126 (2015), 1502.00643.

[20] BICEP/Keck Collaboration XV, Astrophys. J. **927**, 77 (2022), 2110.00482.

[21] Planck Collaboration LVII, Astron. Astrophys. **643**, A42 (2020), ISSN 1432-0746, URL <http://dx.doi.org/10.1051/0004-6361/202038073>.

[22] R. Barrett, M. Berry, T. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. Van der Vorst, Mathematics of Computation **64** (1996).

[23] R. B. Sidje and N. Winkles, J. Comput. Appl. Math. **235**, 1956 (2011), ISSN 0377-0427, URL <https://www.sciencedirect.com/science/article/pii/S0377042710005376>.

[24] BICEP/Keck Collaboration XVI, Astrophys. J. **945**, 72 (2023), ISSN 1538-4357, 2210.05684, URL <http://dx.doi.org/10.3847/1538-4357/acb64c>.

[25] J. A. Zebrowski, C. L. Reichardt, A. J. Anderson, B. Ansarinejad, M. Archipley, L. Balkenhol, P. Barry, K. Benabed, A. N. Bender, B. A. Benson, et al., *Constraints on inflationary gravitational waves with two years of spt-3g data* (2025), 2505.02827, URL <https://arxiv.org/abs/2505.02827>.

[26] K. M. Smith, Phys. Rev. D **74**, 083002 (2006), [astro-ph/0511629](https://arxiv.org/abs/astro-ph/0511629).

[27] M. Tegmark, Phys. Rev. D **55**, 5895–5907 (1997), ISSN 1089-4918, URL <http://dx.doi.org/10.1103/PhysRevD.55.5895>.

[28] E. F. Bunn and B. Wandelt, Phys. Rev. D **96** (2017), ISSN 2470-0029, URL <http://dx.doi.org/10.1103/PhysRevD.96.043523>.

[29] Planck Collaboration XVI, Astron. Astrophys. **571**, A16 (2014), 1303.5076.

[30] J. Grain, M. Tristram, and R. Stompor, Phys. Rev. D **79** (2009), ISSN 1550-2368, URL <http://dx.doi.org/10.1103/PhysRevD.79.123515>.

[31] E. F. Bunn, M. Zaldarriaga, M. Tegmark, and A. de Oliveira-Costa, Phys. Rev. D **67** (2003), ISSN 1089-4918, URL <http://dx.doi.org/10.1103/PhysRevD.67.023501>.

[32] R. Stompor, J. Errard, and D. Poletti, Phys. Rev. D **94** (2016), ISSN 2470-0029, URL <http://dx.doi.org/10.1103/PhysRevD.94.083526>.

[33] B. Thorne, J. Dunkley, D. Alonso, and S. Næss, Mon. Not. R. Astron. Soc. **469**, 2821–2833 (2017), ISSN 1365-2966, URL <http://dx.doi.org/10.1093/mnras/stx949>.

[34] The Pan-Experiment Galactic Science Group, *Full-sky models of galactic microwave emission and polarization at sub-arcminute scales for the python sky model* (2025), 2502.20452, URL <https://arxiv.org/abs/2502.20452>.

[35] A. G. Kritsuk, R. Flauger, and S. D. Ustyugov, Phys. Rev. Lett. **121** (2018), ISSN 1079-7114, URL <http://dx.doi.org/10.1103/PhysRevLett.121.021104>.

[36] G. Martínez-Solaeche, A. Karakci, and J. Delabrouille, Mon. Not. R. Astron. Soc. **476**, 1310–1330 (2018), ISSN 1365-2966, URL <http://dx.doi.org/10.1093/mnras/sty204>.

[37] F. Vansyngel, F. Boulanger, T. Ghosh, B. Wandelt, J. Aumont, A. Bracco, F. Levrier, P. G. Martin, and L. Montier, Astron. Astrophys. **618**, C4 (2018), ISSN 1432-0746, URL <http://dx.doi.org/10.1051/0004-6361/201629992e>.

[38] S. Azzoni, D. Alonso, M. Abitbol, J. Errard, and N. Krachmalnicoff, J. Cosmol. Astropart. Phys. **2023**, 035 (2023), ISSN 1475-7516, URL <http://dx.doi.org/10.1088/1475-7516/2023/03/035>.

[39] Planck Collaboration VI, Astron. Astrophys. **641**, A6 (2020), ISSN 1432-0746, URL <http://dx.doi.org/10.1051/0004-6361/201833910>.

[40] F. Bianchini, D. Beck, W. L. K. Wu, Z. Ahmed, S. Belkner, J. Carron, B. S. Hensley, C. L. Pryke, and C. Umlita, *Cmb-s4: Foreground-cleaning pipeline comparison for measuring primordial gravitational waves* (2025), 2502.04300, URL <https://arxiv.org/abs/2502.04300>.

[41] BICEP/Keck Collaboration XVIII (2025), 2410.12089, URL <https://arxiv.org/abs/2410.12089>.