



BICEP2/Keck Array XI: Beam Characterization and Temperature-to-Polarization Leakage in the BK15 Data Set

P. A. R. Ade¹, Z. Ahmed², R. W. Aikin³, D. Barkats⁴, S. J. Benton⁵, C. A. Bischoff⁶, J. J. Bock^{3,7}, R. Bowens-Rubin⁴, J. A. Brevik³, I. Buder⁴, E. Bullock⁸, V. Buza^{4,9}, J. Connors⁴, J. Cornelison⁴, B. P. Crill⁷, M. Crumrine¹⁰, M. Dierickx⁴, L. Duband¹¹, J. P. Filippini^{12,13}, S. Fliescher¹⁰, J. Grayson¹⁴, G. Hall¹⁰, M. Halpern¹⁵, S. Harrison⁴, S. R. Hildebrandt^{3,7}, G. C. Hilton¹⁶, H. Hui³, K. D. Irwin^{2,14,16}, J. Kang¹⁴, K. S. Karkare^{17,4}, E. Karpel¹⁴, J. P. Kaufman¹⁸, B. G. Keating¹⁸, S. Kefeli³, S. A. Kernasovskiy¹⁴, J. M. Kovac^{4,9}, C. L. Kuo^{2,14}, N. A. Larsen¹⁷, K. Lau¹⁰, E. M. Leitch¹⁷, M. Lueker³, K. G. Megerian⁷, L. Moncelsi³, T. Namikawa¹⁹, C. B. Netterfield^{20,21}, H. T. Nguyen⁷, R. O'Brient^{3,7}, R. W. Ogburn IV^{2,14}, S. Palladino⁶, C. Pryke^{8,10}, B. Racine⁴, S. Richter⁴, A. Schillaci³, R. Schwarz¹⁰, C. D. Sheehy²², A. Soliman³, T. St. Germaine⁴, Z. K. Staniszewski^{3,7}, B. Steinbach³, R. V. Sudiwala¹, G. P. Teply^{3,18}, K. L. Thompson^{2,14}, J. E. Tolan¹⁴, C. Tucker¹, A. D. Turner⁷, C. Umiltà⁶, A. G. Vieregg^{17,23}, A. Wandui³, A. C. Weber⁷, D. V. Wiebe¹⁵, J. Willmert¹⁰, C. L. Wong^{4,9}, W. L. K. Wu¹⁷, H. Yang¹⁴, K. W. Yoon^{2,14}, and C. Zhang³

Keck Array and BICEP2 Collaborations

¹ School of Physics and Astronomy, Cardiff University, Cardiff, CF24 3AA, UK

² Kavli Institute for Particle Astrophysics and Cosmology, SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, CA 94025, USA

³ Department of Physics, California Institute of Technology, Pasadena, CA 91125, USA

⁴ Harvard-Smithsonian Center for Astrophysics, 60 Garden Street MS 42, Cambridge, MA 02138, USA; kkarkare@kicp.uchicago.edu

⁵ Department of Physics, Princeton University, Princeton, NJ 08544, USA

⁶ Department of Physics, University of Cincinnati, Cincinnati, OH 45221, USA

⁷ Jet Propulsion Laboratory, Pasadena, CA 91109, USA

⁸ Minnesota Institute for Astrophysics, University of Minnesota, Minneapolis, MN 55455, USA

⁹ Department of Physics, Harvard University, Cambridge, MA 02138, USA

¹⁰ School of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455, USA

¹¹ Service des Basses Températures, Commissariat à l'Energie Atomique, F-38054 Grenoble, France

¹² Department of Physics, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

¹³ Department of Astronomy, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

¹⁴ Department of Physics, Stanford University, Stanford, CA 94305, USA

¹⁵ Department of Physics and Astronomy, University of British Columbia, Vancouver, British Columbia, V6T 1Z1, Canada

¹⁶ National Institute of Standards and Technology, Boulder, CO 80305, USA

¹⁷ Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA

¹⁸ Department of Physics, University of California at San Diego, La Jolla, CA 92093, USA

¹⁹ Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, CB3 0WA, UK

²⁰ Department of Physics, University of Toronto, Toronto, Ontario, M5S 1A7, Canada

²¹ Canadian Institute for Advanced Research, Toronto, Ontario, M5G 1Z8, Canada

²² Physics Department, Brookhaven National Laboratory, Upton, NY 11973, USA

²³ Department of Physics, Enrico Fermi Institute, University of Chicago, Chicago, IL 60637, USA

Received 2019 April 9; revised 2019 July 23; accepted 2019 August 5; published 2019 October 16

Abstract

Precision measurements of cosmic microwave background (CMB) polarization require extreme control of instrumental systematics. In a companion paper we have presented cosmological constraints from observations with the BICEP2 and Keck Array experiments up to and including the 2015 observing season (BK15), resulting in the deepest CMB polarization maps to date and a statistical sensitivity to the tensor-to-scalar ratio of $\sigma(r) = 0.020$. In this work we characterize the beams and constrain potential systematic contamination from main beam shape mismatch at the three BK15 frequencies (95, 150, and 220 GHz). Far-field maps of 7360 distinct beam patterns taken from 2010–2015 are used to measure differential beam parameters and predict the contribution of temperature-to-polarization leakage to the BK15 *B*-mode maps. In the multifrequency, multicomponent likelihood analysis that uses BK15, *Planck*, and *Wilkinson Microwave Anisotropy Probe* maps to separate sky components, we find that adding this predicted leakage to simulations induces a bias of $\Delta r = 0.0027 \pm 0.0019$. Future results using higher-quality beam maps and improved techniques to detect such leakage in CMB data will substantially reduce this uncertainty, enabling the levels of systematics control needed for BICEP Array and other experiments that plan to definitively probe large-field inflation.

Key words: cosmic background radiation – cosmology: observations – gravitational waves – inflation – polarization

1. Introduction

Progress in understanding the physics of the early universe through measurement of the cosmic microwave background (CMB) has been driven by advances in instrumental sensitivity. Since its discovery over 50 years ago (Penzias & Wilson 1965),

successively finer spatial features in the CMB have been detected, including the 3 mK dipole (Conklin 1969), the $\sim 100 \mu\text{K}$ temperature fluctuations (Smoot et al. 1992), the $\sim 1 \mu\text{K}$ *E*-mode polarization anisotropies (Kovac et al. 2002), and most recently the $\sim 100 \text{ nK}$ *B*-mode polarization from

gravitational lensing of E modes (Hanson et al. 2013; Polarbear Collaboration 2014; Keisler et al. 2015; BICEP2/Keck Array Collaboration VIII 2016; Louis et al. 2017; Planck Collaboration et al. 2018). Confidence in these measurements requires constraining the effects of instrumental systematics to well below the statistical uncertainty, which is determined by factors such as detector count, scan strategy, and astrophysical component separation.

One of the next frontiers in CMB measurements is constraining degree-scale B -mode polarization that may have been imprinted on the surface of last scattering by gravitational waves—a distinctive feature of inflationary models (Kamionkowski & Kovetz 2016). Such a measurement is made challenging by the unknown (or potentially vanishing) amplitude of this signal, parametrized by the tensor-to-scalar ratio r , and the fact that gravitational lensing and Galactic foregrounds also generate B modes; see the CMB-S4 Science Book (2016) for a comprehensive review. The most stringent constraint to date, presented in the companion paper BICEP2/Keck Array Collaboration X (2018, hereafter **BK-X**), uses deep BICEP2 and Keck Array maps at 95, 150, and 220 GHz up to and including the 2015 season (hereafter denoted the BK15 data set) in conjunction with external maps from the Planck (Planck Collaboration et al. 2016) and the *Wilkinson Microwave Anisotropy Probe* (WMAP; Bennett et al. 2013) satellites to separate the CMB from foregrounds. The resulting BB power spectra are consistent with a combination of the expected lensing signal and Galactic dust, yielding a 95% upper limit of $r_{0.05} < 0.072$, which tightens to $r_{0.05} < 0.062$ when including CMB temperature and additional data. The total experimental statistical sensitivity is $\sigma(r) = 0.020$, which takes into account uncertainty in foreground separation.

The BICEP2/Keck Array CMB experiments, located at the Amundsen–Scott South Pole Station, are small-aperture refracting telescopes that measure polarization by differencing pairs of co-located orthogonally polarized detectors, resulting in extremely effective suppression of common-mode noise. The most prominent systematic in the measurement is leakage of the bright temperature sky into polarization ($T \rightarrow P$) due to mismatched beam shapes within a polarization pair. In this work, we report on optical characterization of the Keck Array receivers during the 2014 and 2015 observing seasons operating at 95, 150, and 220 GHz. We present high-fidelity far-field beam maps from which we measure Gaussian beam parameters and validate the “deprojection” procedure used to marginalize over the lowest-order main beam difference modes that produce the majority of $T \rightarrow P$ leakage. We then use these maps in specialized “beam map simulations” to derive upper limits on the higher-order undeprojected $T \rightarrow P$ leakage, and take cross spectra with the BK15 maps to estimate the systematic contribution to the real data. When this leakage is propagated through the BK15 multicomponent analysis, we recover a bias on r that is subdominant to the statistical sensitivity. The analysis techniques explored here offer an example of how we may treat systematics in future experiments with an order of magnitude greater sensitivity.

This paper is one in a series of publications by the BICEP2/Keck Array collaborations and accompanies the primary BK15 CMB results shown in **BK-X**. An overview of the BICEP2/Keck Array instruments is provided in BICEP2 Collaboration II (2014) and optical characterization of the array through 2013 is presented in BICEP2/Keck Array Collaboration IV (2015, hereafter **BK-IV**). Here we extend the beam

characterization in **BK-IV** to the 2014 and 2015 observing seasons, in which we reconfigured several receivers to operate at frequencies other than 150 GHz. Table 1 shows the configuration of BICEP2/Keck Array through 2015. The beam map simulations and constraints on $T \rightarrow P$ leakage presented here are based on those shown in BICEP2 Collaboration III (2015, hereafter **BK-III**) and include all beams contributing to the BK15 maps. In this work the technique is extended to explicitly test for leakage in the CMB maps and to account for the potential systematic contribution in multifrequency component separation.

We organize this paper as follows. In Section 2 we review the setup at the South Pole used to measure the far-field beam pattern of every detector contributing to the CMB maps. Differential Gaussian parameters extracted from these maps are presented in Section 3. Composite maps and array-averaged beam profiles are shown in Section 4. We then use the composite beam maps to predict the level of undeprojected $T \rightarrow P$ leakage in the BK15 data set in Section 5, and cross-correlate these predictions with the CMB maps. We analyze the impact of $T \rightarrow P$ leakage on the multicomponent likelihood analysis in Section 6, and conclude in Section 7.

2. Far-field Beam Measurements

Precision measurements of the BICEP/Keck Array beams in the far field are enabled by our small-aperture approach. The standard far-field distance criterion of $2D^2/\lambda$, where D is the aperture size (26.4 cm) and λ is the wavelength, yields 46, 70, and 103 m for Keck Array receivers at 95, 150, and 220 GHz, respectively. Since the Martin A. Pomerantz Observatory (where the Keck Array is located) and the Dark Sector Laboratory (where BICEP2 was located) are 210 m apart, a source mounted on either building is comfortably in the far field of a receiver on the opposite building as illustrated in Figure 1—see Section 4.1 for more details.

To view the source we install a 45° flat mirror above the mount to redirect the beams over the ground shield; when the telescope is at zenith, the beams point toward the horizon. The five Keck Array receivers are spread out over 1.5 m and redirecting all beams simultaneously would require a larger mirror than can be feasibly supported by the mount. Instead the 1.8×2.7 m aluminum honeycomb mirror is mounted so that at any instant several receivers are completely underneath the mirror. Rotation about the boresight axis then allows all receivers to be measured. Since it is valuable to measure each receiver at many boresight angles, we move the mirror to various positions above the mount so that different angles are accessible (see Figure 2). The co-moving absorptive forebaffles are removed so that the mirror can be mounted.

The source consists of a blade coated with Eccosorb HR-10 microwave absorber within an enclosure. As the blade spins, a beam pointed at the circular aperture on the enclosure alternately sees a hot load (the ~ 250 K ambient-temperature blade) or a cold load (a flat mirror behind the blade that redirects the beam up to the ~ 12 K sky). An optical encoder on the rotation axis of the blade is recorded and used to demodulate the detector timestreams so that only variation at the chop frequency and phase is interpreted as signal. In general a larger aperture offers higher signal and therefore faster mapping speed. Because the telescope scans in azimuth continuously while the chopper spins, a faster chop rate reduces systematic contribution from bright azimuth-fixed signals. In

Table 1
BICEP2/Keck Array Configuration 2010–2015

Rx	BICEP2 2010–2012	Keck 2012	Keck 2013	Keck 2014	Keck 2015
0	150 GHz, 512 (432)	150 GHz, 512 (326)	150 GHz, 512 (318)	95 GHz, 288 (224)	95 GHz, 288 (218)
1		150 GHz, 512 (408)	150 GHz, 512 (400)	150 GHz, 512 (400)	220 GHz, 512 (346)
2		150 GHz, 512 (314)	150 GHz, 512 (312)	95 GHz, 288 (248)	95 GHz, 288 (248)
3		150 GHz, 512 (340)	150 GHz, 512 (422)	150 GHz, 512 (398)	220 GHz, 512 (376)
4		150 GHz, 512 (392)	150 GHz, 512 (386)	150 GHz, 512 (388)	150 GHz, 512 (378)

Note. Center frequency, nominal detector count, and detector count used in cosmological analysis (in parentheses) for all receivers contributing to the BK15 maps. The nominal detector count includes a small number of dark detectors that are intentionally disconnected from the antenna (used to characterize sensitivity to temperature fluctuations and RF interference) and detectors lost to imperfect yield. The detector count contributing to the final maps includes only detectors that have passed all data quality cuts.

previous publications, beam maps were generated with 30 or 45 cm aperture choppers that spun at 18 and 10 Hz chop rates, respectively. For the 2015–2016 deployment season we used a chopper made with a carbon fiber composite to reduce the weight (to 55 lb), featuring a 60 cm aperture and spinning at a 14 Hz chop rate (Karkare et al. 2016). The aperture is closed with Zotefoam HD30 to prevent wind from affecting the motion. Figure 1 shows two choppers mounted on masts at the South Pole and in the lab.

Far-field beam maps are generated by scanning across the source in azimuth and stepping in elevation. A typical beam measurement spans 22.8° in azimuth and 20° in elevation with steps of 0.05° , takes ~ 8 hr, and—given the $\sim 16^\circ$ field of view—allows for measurement of all detectors in at least one Keck Array focal plane to $>2^\circ$ from each beam center. The specifics of the beam measurement are determined by factors such as chop rate, desired map pixelization, and sub-Kelvin refrigerator hold time. Each year several weeks are spent taking beam map data before CMB scanning begins. Since we mask out parts of the map with known ground-fixed contamination (see Section 4.2), maps are made at multiple boresight angles to allow measurement of all regions of the beam. Using several boresight angles also enables rigorous systematics checks. In a single year 40–50 of the scans described above are typical, spread out over 10 boresight angles.

The new measurements presented in this paper took place in February and March of 2014 and 2015; measurements from previous beam mapping campaigns are also included in the simulation results presented in Section 5. From 2010–2015 we measured 10,368 beam patterns to high precision for both BICEP2 and Keck Array, of which 3008 were repeats in which the instrument was unchanged from a previous season, allowing for consistency tests.²⁴

3. Gaussian Beam Parameters

BICEP/Keck Array beam shapes are well approximated by elliptical Gaussians. In BK-IV we presented Gaussian parameters for BICEP2 and Keck Array from 2010–2013, through which all receivers operated at 150 GHz. For the 2014 observing season two 150 GHz receivers were converted to 95 GHz, and for 2015 two more were converted to 220 GHz (Table 1). Here we present beam parameters for 2014 and 2015.

²⁴ These numbers do not reflect the total number of fielded detectors: in one year we installed additional baffling inside the cryostats, which affected the beam patterns of existing detectors (Buder et al. 2014).

3.1. Coordinate System

We begin by defining an instrument-fixed spherical coordinate system that is independent of the orientation of the instrument with respect to celestial coordinates, illustrated in Figure 2. A spatial pixel P containing two orthogonally polarized detectors is defined to be at a location (r, θ) from the boresight B , where r is the radial distance from B and θ is the counterclockwise angle looking out from the telescope toward the sky from the $\theta = 0^\circ$ ray. The $\theta = 0^\circ$ ray is defined with the choice of a fixed index angle on the instrument; we choose this to be, from the boresight, along Tiles 1 and 2 on the focal plane. Tiles are numbered counterclockwise looking directly down on the focal plane, and Tiles 1 and 2 are physically located on the side of the focal plane connecting the heat straps to the sub-Kelvin refrigerator. Each receiver’s $\theta = 0^\circ$ ray is different, depending on the clocking of the receiver in the mount.

For each pixel we then define a local (x', y') Cartesian coordinate system, where the positive x' axis is defined to be along the great circle passing through the pixel center that is an angle $-\theta$ from the \hat{r} direction of the pixel. The y' axis is defined to be the great circle that is $+90^\circ$ away from the x' axis. This (x', y') coordinate system is then projected onto a plane at the pixel center, and rotates with the instrument on the sky.

3.2. Beam Fitting and Statistics

Each measured beam is fit to a two-dimensional elliptical Gaussian $B(\mathbf{x})$ with six free parameters:

$$B(\mathbf{x}) = \frac{1}{\Omega} e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})}, \quad (1)$$

where $\mathbf{x} = (x', y')$ is the beam map coordinate, $\boldsymbol{\mu} = (x_0, y_0)$ is the location of the beam center, Ω is the normalization, and $\boldsymbol{\Sigma}$ is the covariance matrix parametrized as

$$\boldsymbol{\Sigma} = \begin{pmatrix} \sigma^2(1+p) & c\sigma^2 \\ c\sigma^2 & \sigma^2(1-p) \end{pmatrix}. \quad (2)$$

Here σ is the beamwidth, and p and c are the ellipticities in the “plus” and “cross” directions, respectively. An elliptical Gaussian with a major axis oriented along the x' or y' axes (see Figure 3) has $+p$ or $-p$ ellipticity, respectively, and one with a major axis oriented diagonally has $\pm c$ ellipticity. The total ellipticity is $e = \sqrt{p^2 + c^2}$. Differential parameters (i.e.,




Figure 1. Several weeks of each South Pole deployment season are spent generating far-field beam maps using the pictured setup. Left: Keck Array and BICEP3 taking far-field beam maps simultaneously at the South Pole. Both choppers and far-field flat mirrors are visible. Middle: carbon fiber chopper blade coated in Eccosorb HR-10, mounted in the lab; during operation it spins at 14 Hz. Right: carbon fiber enclosure holding the blade; the 60 cm aperture (white) is sealed with Zotefoam HD30. When the blade does not fill the aperture, the beam is redirected to the zenith with a 45° mirror.

within a polarization pair) are defined as differences of per-detector fits: differential beamwidth $d\sigma = \sigma_A - \sigma_B$, differential pointing $dx = x_{0,A} - x_{0,B}$ and $dy = y_{0,A} - y_{0,B}$, differential ellipticity $dp = p_A - p_B$ and $dc = c_A - c_B$. Here A and B refer to the orthogonally polarized detectors within a pair. Note that in this parameterization differential beamwidth and differential pointing are expressed in absolute units and not as fractions of the nominal beamwidth (Wong 2014).

Beam parameter statistics are generated by finding the best-fit x , σ , p , and c for each beam in each mapping run. We then remove measurements in which the beam was not redirected toward the source due to the mirror geometry (i.e., the receiver in question was not physically underneath the mirror), the detector was not operating normally, the fit did not converge for both A and B , the fit did not fall in a physically acceptable range, or the residual between the fit and the measured beam showed obvious artifacts. In a given year a typical Keck Array beam is measured 10 times, though this varies significantly across the focal plane.

For each detector and pair, we take the median across all measurements as the best estimate of each parameter and take half the width of the central 68% of the distribution of those measurements as the measurement uncertainty.²⁵ The characteristic uncertainty for an individual measurement—taken to be the median of the measurement uncertainties for all detectors/pairs across the array—is denoted “individual measurement uncertainty.” In general the uncertainty on per-detector ellipticity is somewhat large because random artifacts occasionally escape the automated cuts and cause the fitting routine to choose an ellipticity that is a poor fit to the beam. Measurement uncertainties of differential parameters are typically smaller than those of per-detector parameters. Common-mode effects, such as artifacts in the timestreams or systematics associated with boresight rotation, tend to affect both detectors in a pair equally and thus bias parameter estimates in the same way. We defer a detailed discussion of noise and systematics in the beam measurement to Section 4.1.

To characterize the distribution of parameters we also find the median across the focal plane (“FPU median”) and quantify the variation across the focal plane as half the width of the central 68% of the distribution of best estimates for each detector/pair (“FPU scatter”). Note that FPU scatter measures the spread of best-estimate parameters across the array and is not a

measurement uncertainty. We correct for the non-negligible size of the thermal source aperture when reporting σ (Section 4.3).

3.3. Measured Beam Parameters

Tables 2 and 3 show per-detector and differential beam parameters for Keck Array receivers in 2014 and 2015, respectively, in the following format: FPU median \pm FPU scatter \pm individual measurement uncertainty.

Beamwidths for the three frequencies in the BK15 data set are roughly $0^\circ.305$ (43' FWHM, 95 GHz), $0^\circ.210$ (30' FWHM, 150 GHz), and $0^\circ.141$ (20' FWHM, 220 GHz). Variability across the focal plane is not detected above the individual measurement uncertainty. The central values of per-detector p and c ellipticities are generally close to zero and are spread equally above and below; this is because the optical design places the optimal focus on an annulus of detectors located a median distance from the center of the focal plane, causing enhanced ellipticity toward the edge. Such a pattern is roughly azimuthally symmetric, so that p and c average to small values.

Differences in beam shapes between the orthogonally polarized A and B detectors contribute to $T \rightarrow P$ leakage (Hu et al. 2003). The majority of the power in the BICEP/Keck difference beams is encapsulated in a second-order expansion of the beam profile, which couples to the CMB temperature sky and its first and second derivatives. The modes corresponding to these couplings are either “deprojected” from the CMB maps by scaling and removing the best-fit templates of the Planck sky map and its derivatives from pair difference data, or “subtracted” by removing these templates with amplitudes determined directly from beam maps. We typically subtract differential ellipticity instead of deprojecting it, because deprojecting would preferentially filter the TE and EE spectra. BICEP/Keck polarization maps are therefore largely insensitive to $T \rightarrow P$ leakage entering through these modes; see BK-III for more details.²⁶

Nevertheless, since differential Gaussian parameters—which correspond roughly to the modes that are deprojected²⁷—often probe optical and detector fabrication effects, it is worthwhile

²⁵ This statistic is relatively insensitive to outliers, and would equal 1σ for a Gaussian distribution of measurements.

²⁶ This assumes the Planck temperature maps are free of systematics. We have also simulated the effect of Planck noise in the templates and found it to be negligible (BK-III). Noise in beam measurements could affect the accuracy of subtraction; based on the repeatability of dp and dc , the uncertainty in differential ellipticity could contribute $T \rightarrow P$ leakage at the $\rho < 1 \times 10^{-6}$ level (see Section 5.2).

²⁷ Deprojection coefficients are obtained by regressing the pair difference data against templates of the Planck temperature sky and its derivatives that have been smoothed by our array-averaged beam profile. If the beam profile were perfectly Gaussian the coefficients would exactly match the parameters obtained by differencing Gaussian fits.

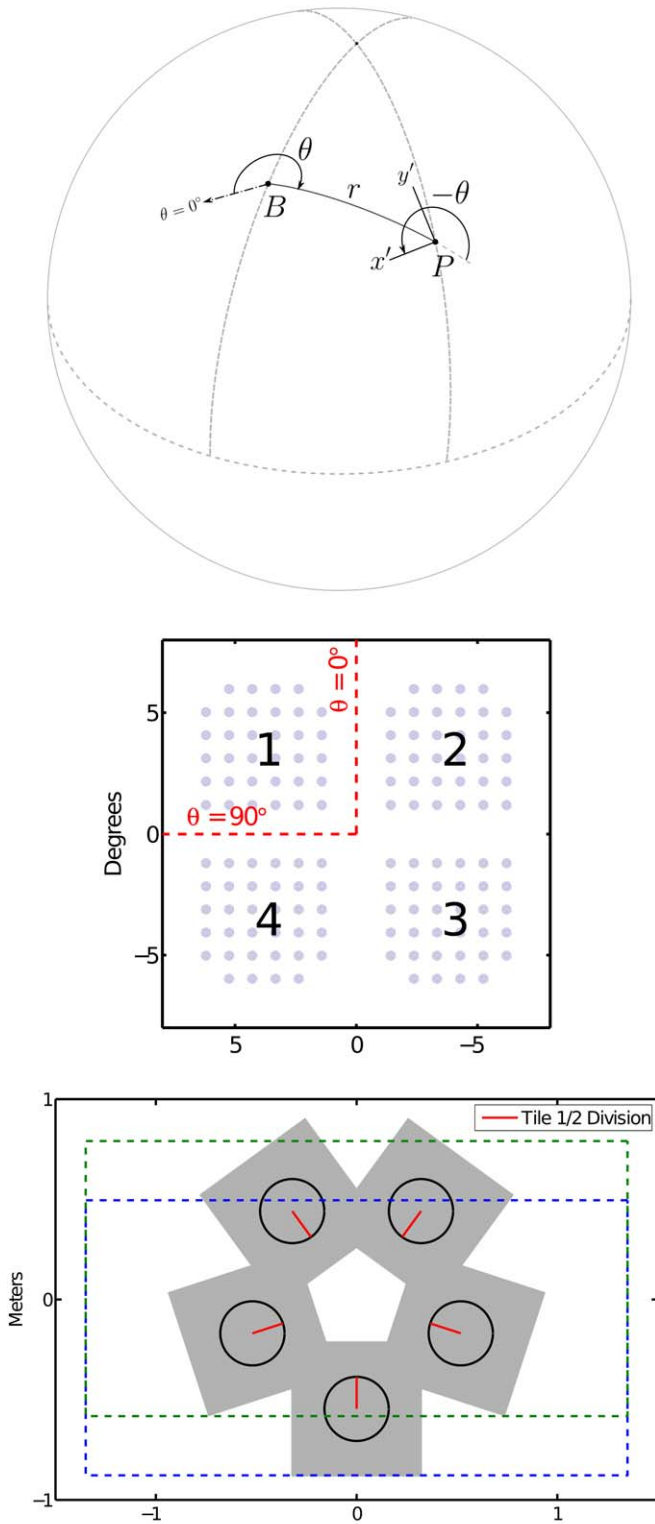


Figure 2. Top: the (x', y') coordinate system is centered locally for each pixel P at a location (r, θ) from the boresight B , and is referenced to the $\theta = 0^\circ$ ray. Middle: the $\theta = 0^\circ$ ray is referenced to the orientation of detector tiles, here depicted looking directly down onto the focal plane. Bottom: schematic of the five Keck receiver apertures in the mount, clocked at 72° with respect to each other. Each receiver’s Tile 1/2 dividing line (corresponding to the $\theta = 0^\circ$ ray) is indicated. Two footprints of the far-field flat mirror (hoisted 3.7 m above the apertures) are depicted as dashed boxes. The rough extent to which beams have diverged by the time they intercept the mirror is depicted in gray. Multiple mirror positions are necessary to reflect beams from all receivers at many boresight angles.

to quantify and plot them in focal plane coordinates to understand their causes and how they propagate through the final analysis. Here we discuss differential pointing and ellipticity, which if left uncorrected could contribute significantly to $T \rightarrow P$ leakage. Differential beamwidth is generally small enough to be negligible.

Figure 3 shows differential pointing dx, dy and ellipticity dp, dc for all Keck Array 2014 and 2015 receivers in focal plane layout. Differential pointing is usually the beam shape mismatch mode that causes the most $T \rightarrow P$ leakage.²⁸ Several trends stand out: first, measurements are consistent between years in which the receiver did not change (e.g., Rx0 in 2014 and 2015). Second, the absolute magnitude is somewhat frequency-dependent—95 GHz receivers show larger offsets than both 150 and 220 GHz, which are similar to each other—but there is large variability. Finally, it is typical to have a coherent differential pointing direction across a detector tile. Significant effort at the design and fabrication stages has been spent in reducing this mode, resulting in tiles with factors of ~ 2 – 4 lower differential pointing than earlier, BICEP2-era tiles (O’Brien et al. 2012). Differential ellipticity is similarly consistent from year to year and coherent within tiles; $-dp$ appears to dominate.

We have checked that the deprojection coefficients—which scale templates of the Planck temperature sky and its derivatives to best fit the CMB data—match measured beam parameters and exhibit scatter consistent with their noise. On-sky data confirm that the CMB maps and far-field beam maps measure the same dominant beam mismatch modes (BK-III).

It is the higher-order residuals, not captured by the above second-order expansion of the beam profile, which could contribute $T \rightarrow P$ leakage to the final polarization maps. To move beyond differential Gaussians we require sensitive beam maps capable of capturing the potentially complex higher-order difference beam morphology.

4. Composite Beam Maps and Beam Profiles

To fully assess the impact of beams in CMB analysis we use deep far-field measurements extending to several degrees away from the beam center. In this section we describe how the beam maps generated in Section 2—hereafter denoted “component” maps—are combined to form high-fidelity “composite” maps. Using composite beam maps reduces noise and systematics in the beam measurement and allows us to measure the main beam at all azimuthal angles.

4.1. Noise and Systematics in Beam Measurements

Typical noise levels in a single pair’s component beam maps (i.e., from one measurement made with the 45 cm chopper) would induce a bias on the prediction of $T \rightarrow P$ leakage from mismatched beams equivalent to $\rho \sim 0.02 \pm 0.01$ at 150 GHz even if the beams were perfectly matched²⁹ (see Section 5.2).

²⁸ While relative gain is also important, it is not measured in beam mapping because beam map normalization effectively deprojects it. For CMB analysis, relative gain is corrected at the timestamp level by normalizing by the response to a small change in atmospheric loading (“elevation nods”), and at the map level by deprojection.

²⁹ Beam map noise is measured by taking the standard deviation of pixel values in a region of the map far away from the main beam and ground contamination (i.e., more than 10° above the beam). For 150 GHz detectors measured with the 45 cm chopper, the typical rms noise level in 0.1 pixels is $1/800$ of the main beam amplitude. The bias on r is estimated by generating Gaussian noise simulations with these amplitudes and running them through the beam map simulation pipeline.




Figure 3. Differential pointing and ellipticity for the Keck Array receivers in the 2014 and 2015 seasons. Parameters are plotted in focal plane coordinates as projected on the sky. A/B detectors are polarized along the y'/x' axes, respectively, indicated in the bottom-right panel. Detectors at 95 GHz are in red, 150 GHz in green, and 220 GHz in blue; those in gray are not used in the analysis. Parameters are consistent across years when the receiver did not change—Rx0, Rx2, and Rx4 stayed the same in 2014 and 2015. Left columns: differential pointing rendered as an arrow pointing from the A detector location to the B detector location. The arrow length indicates the degree of mismatch, multiplied by 20. Right columns: differential ellipticity rendered as ellipses; major axes are proportional to $\sqrt{dp^2 + dc^2}$, a measure of the magnitude of the differential ellipticity (minor axes are fixed). The magnitude has been multiplied by 75 for visibility. **BK-IV** presents analogous plots for BICEP2 and Keck Array in previous seasons.

Table 2
Keck Array 2014 Beam Parameter Summary Statistics

	Rx0 (95 GHz)	Rx1 (150 GHz)	Rx2 (95 GHz)	Rx3 (150 GHz)	Rx4 (150 GHz)
σ ($^\circ$)	$0.303 \pm 0.003 \pm 0.003$	$0.213 \pm 0.004 \pm 0.003$	$0.306 \pm 0.002 \pm 0.003$	$0.216 \pm 0.004 \pm 0.003$	$0.208 \pm 0.002 \pm 0.002$
p (+)	$-0.010 \pm 0.013 \pm 0.019$	$0.005 \pm 0.021 \pm 0.023$	$-0.004 \pm 0.011 \pm 0.016$	$0.008 \pm 0.030 \pm 0.021$	$0.002 \pm 0.018 \pm 0.019$
c (\times)	$-0.001 \pm 0.013 \pm 0.020$	$0.005 \pm 0.022 \pm 0.025$	$-0.008 \pm 0.011 \pm 0.019$	$-0.003 \pm 0.012 \pm 0.021$	$0.005 \pm 0.014 \pm 0.023$
dx ($''$)	$0.46 \pm 0.76 \pm 0.05$	$0.10 \pm 0.97 \pm 0.06$	$0.34 \pm 0.64 \pm 0.06$	$-0.11 \pm 0.35 \pm 0.08$	$0.18 \pm 0.42 \pm 0.05$
dy ($''$)	$-0.57 \pm 0.79 \pm 0.05$	$-0.57 \pm 0.52 \pm 0.05$	$-0.52 \pm 0.65 \pm 0.06$	$-0.08 \pm 0.46 \pm 0.08$	$-0.12 \pm 0.32 \pm 0.05$
$d\sigma$ ($^\circ$)	$0.001 \pm 0.002 \pm 0.001$	$0.001 \pm 0.001 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$
dp (+)	$-0.013 \pm 0.014 \pm 0.002$	$-0.009 \pm 0.010 \pm 0.002$	$-0.006 \pm 0.004 \pm 0.002$	$-0.007 \pm 0.012 \pm 0.003$	$-0.020 \pm 0.004 \pm 0.002$
dc (\times)	$0.002 \pm 0.004 \pm 0.002$	$-0.010 \pm 0.005 \pm 0.003$	$0.002 \pm 0.003 \pm 0.002$	$-0.002 \pm 0.008 \pm 0.002$	$-0.002 \pm 0.002 \pm 0.002$

Note. Beam parameters are listed as FPU median \pm FPU scatter \pm individual measurement uncertainty. See Section 3 for a detailed discussion of the uncertainties. BK-IV presents analogous statistics for BICEP2 and Keck Array in previous seasons.

This bias and uncertainty can be reduced by combining maps from multiple measurements. The average pair has ~ 10 good maps per year (Section 3), so assuming Gaussian noise and a systematic error-free measurement, we expect to measure leakage with an uncertainty of $\sigma(r) \sim 0.001$ for a *single* pair in a given year after averaging. Similarly—since our estimate of $T \rightarrow P$ leakage for the CMB map is made by coadding the predicted Q/U leakage maps from all pairs—if noise is uncorrelated across detectors, the uncertainty on the final leaked BB estimate scales down with the number of pairs. Therefore, if there is real structure in the higher-order difference beams at a level that could be important for CMB analysis, we fully expect the beam maps to measure it. Since the 45 cm chopper noise level already outperforms our $\sigma(r)$ requirements, the current 60 cm thermal chopper will be well below the requirement for the next generation of BICEP experiments.

The most prominent systematic in the beam map measurement is ground-fixed contamination, which leaks into the in-phase demodulated signal at a low level.³⁰ We conservatively choose to mask out all regions of the map with known ground-fixed structure. To measure the regions of the beam that were masked, the receiver is rotated about the boresight axis and another map is made.

Other potential systematics include curvature in the redirecting mirror, the finite distance of the source, and nonlinearity in the detector response at the relatively high loading conditions used in calibration measurements. Using the measured mirror curvature³¹ we find that the absolute beam shape experiences a very small change— $B(\ell)$ changes by $<0.01\%$ for the 220 GHz beams at $\ell = 200$ —and the differential beam shapes will see even less distortion because both detectors in a pair reflect off the mirror nearly identically. Through simulation of the beam shape error expected for a measurement at 210 m instead of at infinity, we find that BICEP3 with a nominal far-field distance of 171 m experiences a deviation of $<0.2\%$ at $\ell = 200$; the effect will be smaller for all Keck Array bands. Finally, we constrain nonlinearity on the high-loading aluminum transition sensor to be $\lesssim 1\%$ using load

curves in which the detector bias voltage is ramped, providing an I–V characteristic.

Nonrepeatable contamination, i.e., non-Gaussian noise, is more worrisome and most often consists of single-sample glitches, which manifest as “hot spots” in the demodulated timestreams. A single glitch in the timestream can have an amplitude much larger than that of the main beam, resulting in a leakage prediction orders of magnitude larger than the true value even if the rest of the beam map were accurate.

For this reason we regard the differential beam maps as upper limits on the $T \rightarrow P$ leakage from mismatched beams: they should capture all true leakage within the beam map radius, but could also contain non-Gaussian noise or measurement systematics that are difficult to quantify. Looking toward future results, we have demonstrated a reduction in beam map contamination through improved low-level deglitching and data quality cuts. These improvements will find application both in reanalysis of existing beam map data and in future data sets.

4.2. Composite Beam Map Generation

Composite beam maps are generated by combining component maps that have passed data quality cuts and in which known, spatially fixed contamination has been removed. We begin with the set of component maps from which beam statistics were calculated (Section 3). We then apply a spatial mask: the ground ($>2^\circ$ below the source), the mast, and the South Pole Telescope are removed. Maps are centered on the common pair centroid, rotated to account for boresight angle, and peak-normalized. Finally the composite is made by assembling a stack of all good maps and taking a median in each spatial pixel. We take a median instead of a mean because occasionally high-amplitude artifacts escape the automated cuts. In current Keck Array maps, median-filtering results in demonstrably lower noise levels and fewer artifacts than mean-filtering. In future data sets we plan to improve low-level deglitching and spatial coverage to the point at which a mean may be taken.

Figure 4 shows sample composite beam maps and the component maps used to form them at the three Keck Array frequencies. The component maps have been masked for ground-fixed signal. Noise reduction and rejection of spurious contaminating signals are evident in the composites.

In principle it is best to use measurements that extend far away from the beam center to capture as much of the optical response as possible. We note, however, that in standard CMB

³⁰ This contamination arises when scanning across a change in temperature, such as from cold sky to the (relatively) warm South Pole Telescope. Part of that temperature change is interpreted by the deconvolution kernel as signal from the thermal chopper. Ground-fixed signal usually enters into the measurement at the -20 dB level or lower.

³¹ Using in situ photogrammetry we have constrained the mirror’s flatness to better than 1.5 mm across its surface.

Table 3
Keck Array 2015 Beam Parameter Summary Statistics

	Rx0 (95 GHz)	Rx1 (220 GHz)	Rx2 (95 GHz)	Rx3 (220 GHz)	Rx4 (150 GHz)
σ (°)	$0.304 \pm 0.003 \pm 0.003$	$0.141 \pm 0.002 \pm 0.002$	$0.307 \pm 0.002 \pm 0.004$	$0.142 \pm 0.002 \pm 0.003$	$0.207 \pm 0.003 \pm 0.002$
p (+)	$-0.007 \pm 0.015 \pm 0.019$	$0.003 \pm 0.020 \pm 0.028$	$-0.003 \pm 0.011 \pm 0.017$	$0.001 \pm 0.022 \pm 0.031$	$0.002 \pm 0.022 \pm 0.017$
c (×)	$-0.002 \pm 0.013 \pm 0.019$	$-0.003 \pm 0.021 \pm 0.031$	$-0.005 \pm 0.012 \pm 0.019$	$0.002 \pm 0.028 \pm 0.035$	$0.005 \pm 0.017 \pm 0.016$
dx (′)	$0.42 \pm 0.76 \pm 0.04$	$0.32 \pm 0.26 \pm 0.04$	$0.33 \pm 0.68 \pm 0.05$	$-0.49 \pm 0.19 \pm 0.04$	$0.19 \pm 0.40 \pm 0.04$
dy (′)	$-0.56 \pm 0.78 \pm 0.04$	$-0.51 \pm 0.23 \pm 0.04$	$-0.54 \pm 0.65 \pm 0.05$	$0.43 \pm 0.12 \pm 0.03$	$-0.11 \pm 0.33 \pm 0.03$
$d\sigma$ (°)	$0.001 \pm 0.002 \pm 0.001$	$0.000 \pm 0.000 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$	$0.000 \pm 0.000 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$
dp (+)	$-0.013 \pm 0.013 \pm 0.002$	$-0.015 \pm 0.009 \pm 0.006$	$-0.006 \pm 0.004 \pm 0.002$	$-0.017 \pm 0.005 \pm 0.005$	$-0.019 \pm 0.005 \pm 0.002$
dc (×)	$0.002 \pm 0.004 \pm 0.002$	$0.001 \pm 0.006 \pm 0.006$	$0.002 \pm 0.002 \pm 0.002$	$0.004 \pm 0.005 \pm 0.005$	$-0.002 \pm 0.002 \pm 0.002$

Note. Beam parameter are listed as FPU median \pm FPU scatter \pm individual measurement uncertainty. See Section 3 for a detailed discussion of the uncertainties. Only Rx0/Rx2/Rx4 are in common with Table 2. BK-IV presents analogous statistics for BICEP2 and Keck Array in previous seasons.

observations the co-moving forebaffle absorbs off-axis response at large angles (generally $>10^\circ$ from the main beam). Since the forebaffles are removed to make room for hoisting the far-field flat mirror, response measured in the far-field beam maps at very large angles may not correspond to real beam pickup during CMB scans.³² The composite maps also become noisier further away from the beam center because there are fewer hits per pixel as a result of the spatial masking. For the beam map simulation we use composites made out to 2° from the beam center (Section 5).

4.3. Beam Profiles

To generate array-averaged beam profiles we coadd the composite maps over all detectors contributing to CMB data, inversely weighted by CMB-derived per-detector noise. Figure 5 shows the averaged radial profiles and the equivalent $B(\ell)$ window functions. To remove the effect of the finite size of the source aperture, we divide $B(\ell)$ by $\frac{2J_1(\ell D/2)}{\ell D/2}$, where J_1 is the Bessel function of the first kind and D is the angular diameter of the source as seen from the telescope. Typical detector-to-detector variations at $\ell = 100$ are 2.6% (95 GHz), 2.2% (150 GHz), and 3.7% (220 GHz); the statistical uncertainty on individual beam profiles is small compared to this variation. We use these beam profiles in CMB analysis to smooth input maps for signal simulations and to smooth the Planck temperature sky and its derivatives, which serve as deprojection templates.

5. Simulation of Temperature-to-Polarization Leakage

To estimate the $T \rightarrow P$ leakage present in the BK15 maps we run “beam map simulations” with composite beam maps such as those presented in Section 4.³³ In this section we review how the simulations are generated and show the predicted leaked polarization maps and power spectra. The auto power spectra represent upper limits to the $T \rightarrow P$ leakage in the BK15 maps due to mismatched beam shapes within polarization pairs. We then take the cross spectra of the beam

map simulations with the real BK15 maps, which represent our best estimate of the leakage present in CMB data.

5.1. Simulation Methodology and Leakage Estimates

The standard BK15 pipeline is used to run beam map simulations. We inject $T \rightarrow P$ leakage at the timestream level by convolving the Planck temperature sky with the per-detector beam maps and sampling from the resulting smoothed maps using real detector pointing data.³⁴ We intentionally set the Planck Q/U maps to zero so that any measured polarization is a result of $T \rightarrow P$ leakage from mismatched beams. The timestreams are then processed into Q/U maps just as are the real data, including identical cuts and detector weighting. Deprojection of differential gain/pointing and subtraction of differential ellipticity are also applied³⁵ (see Section 3). As in our standard B -mode analysis, bandpowers are measured after applying a matrix-based purification that reduces $E \rightarrow B$ leakage from partial sky coverage and filtering (BICEP2/Keck Array Collaboration VII 2016) so that only the leakage modes relevant for our actual CMB maps are used. All beams contributing to the BK15 maps, including BICEP2 and Keck 2012–2013 (BK-IV), are accounted for in these results.

We assess the impact of a systematic contribution to a single-frequency BB spectrum with a quadratic estimator ρ that is the equivalent r level of the contamination. The estimator is

$$\rho = \frac{\langle C \rangle^T N^{-1} \hat{C}}{\langle C \rangle^T N^{-1} \langle C \rangle}, \quad (3)$$

where \hat{C} are the systematic bandpowers predicted from the beam map simulation, N is the bandpower covariance matrix from signal + noise simulations (“signal” refers to our fiducial lensed- Λ CDM + dust skies and “noise” matches that of the BK15 maps), and $\langle C \rangle$ are the BB bandpower expectation values for an $r = 1$ signal. This effectively weights the systematic bandpowers by the ratio of the expected signal to the noise variance in each ℓ bin. For BICEP/Keck Array, most of the

³² In separate measurements, we use an amplified microwave source to measure far sidelobes. Comparing such maps made with and without the forebaffle, we see no perceptible difference in the main beam region. We expect any such difference to be extremely small, given the small amount of power intercepted ($\sim 0.7\%$, see BK-IV) and the very efficient absorption of this power by the forebaffle.

³³ Although the term “simulation” is used, all inputs are based on real data: Planck maps, detector pointing, and on-sky beams measured in situ at the South Pole.

³⁴ The detector-centered, locally flat beam maps are scaled and rotated using the full spherical geometry appropriate for each detector’s pointing trajectory prior to convolution with a flat-sky projection of the Planck map. The Planck beam has been deconvolved from the temperature map.

³⁵ There is a small mismatch between the simulation timestreams (formed from flat-sky convolutions) and the deprojection templates (formed from curved HEALPix maps), but we have demonstrated that it corresponds to $\rho \sim 10^{-6}$ and is negligible; see BK-III for details.

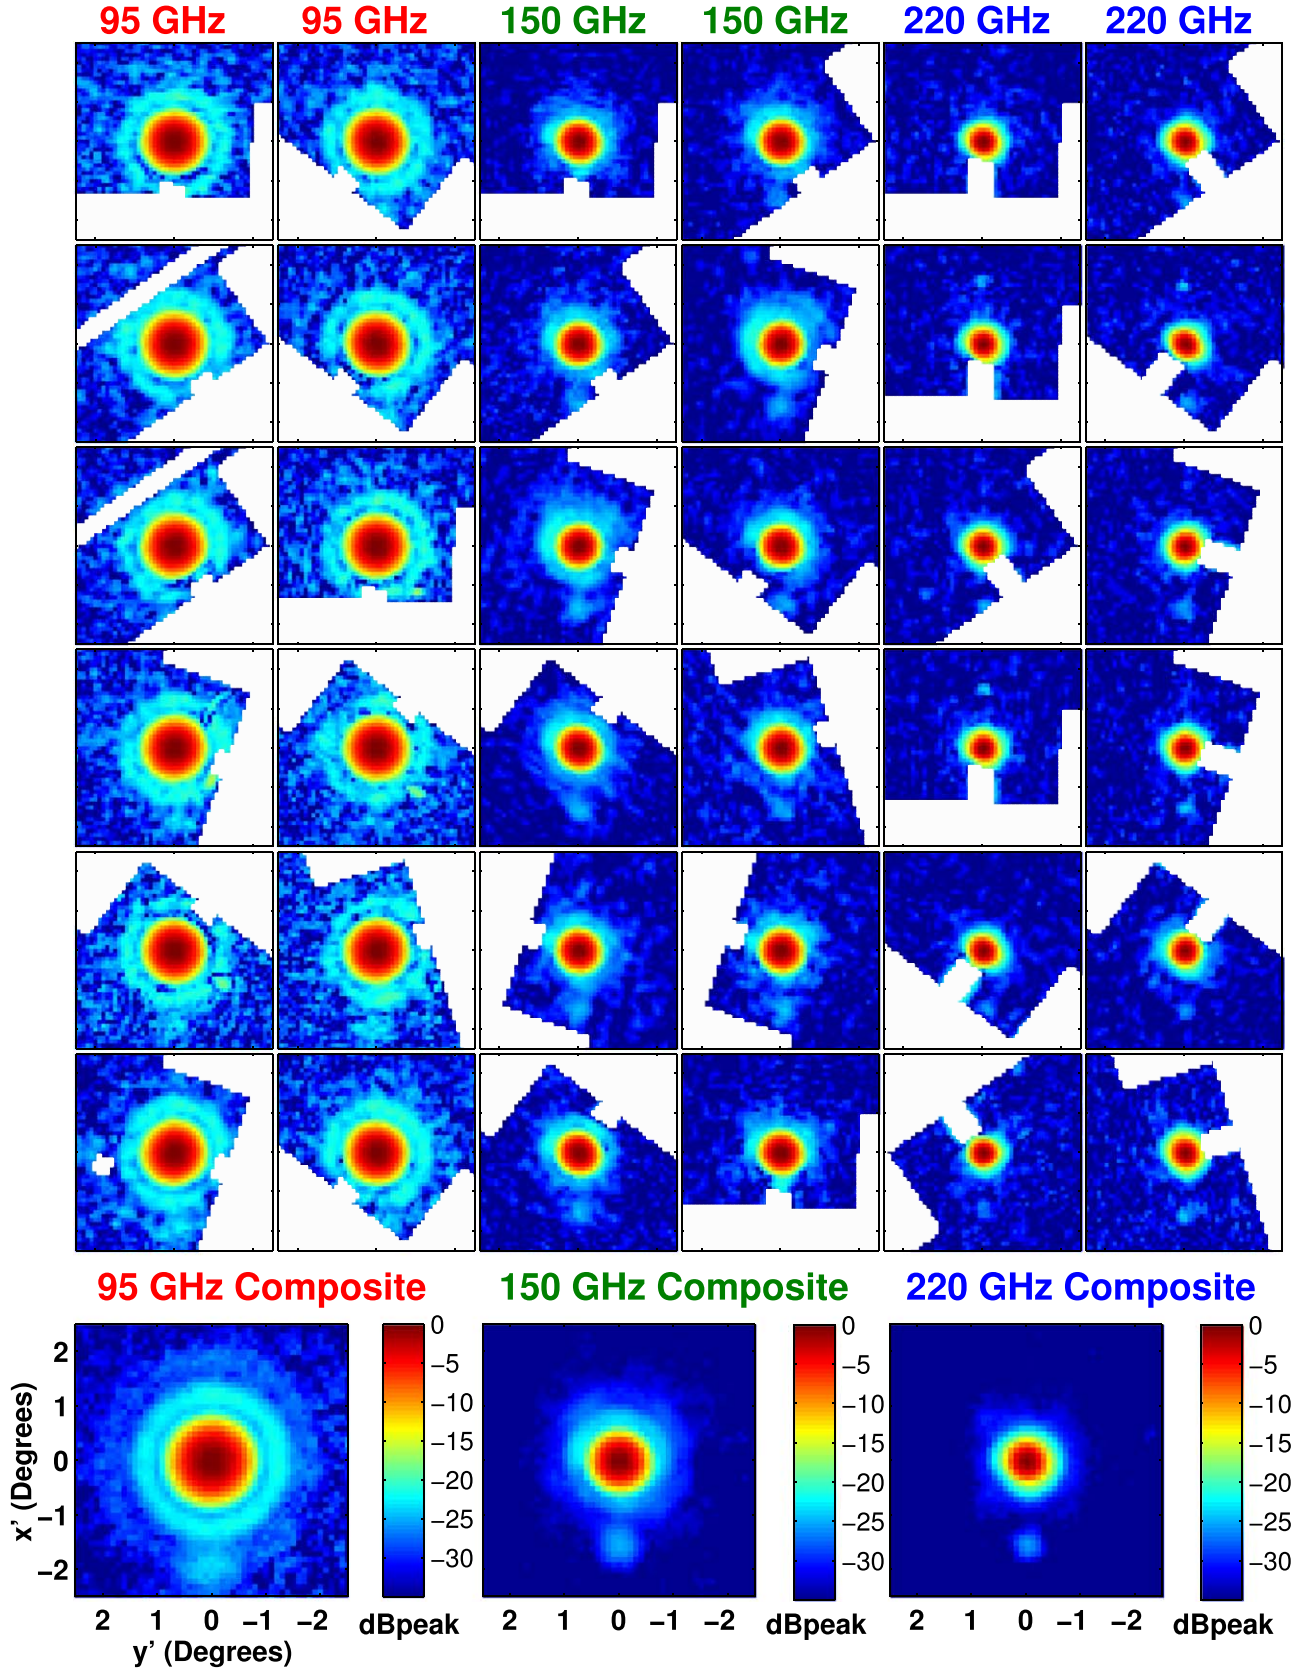


Figure 4. Large panels: example composite beam maps for individual 95 GHz (left), 150 GHz (middle), and 220 GHz (right) detectors. Small panels: the component maps that contribute to the composite, taken at multiple boresight angles—the ground, mast, and SPT have been masked out. The circular feature at the ~ -23 dB level, which in all cases here appears below the main beam, is due to crosstalk in the readout system and has been extensively characterized (BICEP2 Collaboration II 2014; BK-III). The maps are in dB relative to peak amplitude and share the same color scale; because beams of different widths have different peak amplitudes, the noise in the lower-frequency beams appears inflated.




Figure 5. Left: azimuthally averaged beam profiles for Keck Array at 95, 150, and 220 GHz, coadded over all operational channels in 2015. The maps are normalized relative to an isotropic radiator under the assumption that all beam power is contained within the $r < 4^\circ$ maps. Right: the $B(\ell)$ window functions, equivalent to Fourier transforms of the beam profiles. The effect of the finite source size has been removed. BK-IV presents analogous plots for BICEP2 and Keck Array in previous seasons.

statistical power in the ρ estimate comes from the first three bins (BICEP1 Collaboration 2014).

5.2. Simulation Results

Only the component of the pair difference beam remaining after deprojection (the “undeprojected residual”) is relevant for assessing the main beam $T \rightarrow P$ leakage in the BK15 map. To illustrate typical difference beams, in Figure 6 we plot example composite maps for both A and B detectors in a pair, the pre-deprojection difference map, the difference map with differential pointing removed, and the final undeprojected residual. Amplitudes of features in typical undeprojected residual maps are $\sim 0.2\%$ of the main beam peak.

Far-field beam maps routinely identify features in individual detector pairs that contribute excess $T \rightarrow P$ leakage. To illustrate, we show a 95 GHz beam that contains a large feature in the undeprojected residual, caused by a well-characterized anomalous interaction between the detector tile corrugations and the tile edge detectors. The impact on the undeprojected residuals and the subsequent remedy is discussed in more detail in Section 5.3.

The Q maps of residual $T \rightarrow P$ leakage predicted using the composite beam maps are shown in Figure 7. An estimate of the beam map noise is also shown, and discussed in more detail below. We expect all of the true leakage to be captured in these maps, but again emphasize that there is likely a non-negligible contribution from systematics in the beam measurement.

While in general there is no expectation that undeprojected residuals should prefer E or B , it is not uncommon for the leakage to contaminate one over the other. For example, the 95 GHz tile corrugation feature shown in Figure 6 leaks primarily to E modes because it is aligned with the detector polarization axes (the Q beam)—leakage aligned at 45° would leak to B (see BK-III for more details).

The central region of the 150 GHz map contains a band of lower leakage, which is due to the complex interaction of the rotational symmetry of undeprojected residuals (which determines whether they cancel under boresight rotation), their distribution across the focal plane (which determines whether they cancel when coadded with other detectors), and the

combination of the many distinct receivers that contribute to the map.

The BB spectra corresponding to these maps are shown in Figure 8. We present both noise-debiased auto spectra and cross spectra of the beam map simulations with the BK15 maps. The interfrequency cross spectra, corresponding to leakage that correlates across frequencies, are consistent with zero within the uncertainties and are not plotted here.

The beam map auto spectra (black lines) represent upper limits to the single-frequency leakage since the beam measurement may contain low-level systematics. We estimate the noise contribution by forming a “beam map jackknife” for each detector pair. For each beam, we randomly divide the component maps into two halves and form two separate composite maps. The difference between these maps is effectively a jackknife in which the signal is removed and the noise remains (Figure 7 right column). This measure of noise bias has been subtracted from the beam map auto spectrum. The ρ estimates for the upper limits and the jackknives are shown in Table 4 and indicate that the noise contribution to the beam maps is subdominant compared to the combination of true beam mismatch and potential systematics in the measurement.

The cross spectra with the BK15 maps (teal crosses; see BK-X Figures 7–9) offer an unbiased estimate of $T \rightarrow P$ leakage in the real data and should be insensitive to systematics in the beam measurement; ρ estimates are shown in Table 4. Uncertainties in the cross spectra arise from several sources: instrumental and atmospheric noise, the true-sky lensing and dust B modes, and noise in the beam map measurement. To estimate the uncertainty from the CMB half of the cross spectrum, we use our ensemble of 499 B -mode lensing, dust, and noise (i.e., instrumental and atmospheric) simulations—see BK15 for more details. Using the single beam map simulation, we take the 499 cross spectra:

$$(\text{fixed } T \rightarrow P) \times (\text{lensing} + \text{dust} + \text{noise}).$$

The variance of these cross spectra constitute the error bars in Figure 8, which are propagated through to uncertainties on the ρ estimates (“cross spectrum σ ” in Table 4). Since the beam




Figure 6. Example differential beam maps at 95 GHz (left), 150 GHz (middle), and 220 GHz (right). In each column, the top two panels show composite maps of *A* and *B* detectors in a pair; the color scale is in decibels relative to the peak amplitude. The third panel shows the difference between *A* and *B*, dominated by differential pointing. The color scale is linear relative to the pair sum peak amplitude ($\pm 2\%$)—this has effectively had differential gain deprojected. The fourth panel shows the residual after differential pointing is deprojected; note the color scale has changed ($\pm 0.5\%$). The last panel has the same color scale as the fourth and shows the undeprojected residual after differential pointing, beamwidth, and ellipticity have been removed; this contributes to the $T \rightarrow P$ leakage discussed in Section 5.2. The large feature in the 95 GHz map is discussed in more detail in the text.

map jackknife is negligible compared to this, uncertainty in the beam measurement is not included here.

5.3. Discussion

Tracing $T \rightarrow P$ leakage from input beam maps to the final *BB* spectra provides insight into detector and optics fabrication effects that can be improved. For example, the large tile corrugation feature in the 95 GHz pixels (Figure 6) impacted about half of the detectors contributing to the map, and has been corrected in focal planes produced subsequently. The contribution of these detectors is extremely anomalous; our present 95 GHz ρ results are over an order of magnitude larger than what they would be if the tile edge detectors were excluded (Karkare 2017).

Do the beam map simulations correlate with the real BK15 maps and detect $T \rightarrow P$ leakage? At 95 and 150 GHz the cross spectra are generally positive and follow the auto spectra, but are a factor of ~ 2 lower. Given the error estimates— $\rho = (4 \pm 4) \times 10^{-3}$ at 95 GHz and $\rho = (5 \pm 4) \times 10^{-3}$ at 150 GHz—there is

only tentative evidence for leakage in the real BK15 maps. The suppression of the cross with respect to the auto could indicate non-negligible systematics in the beam map measurement, which would bias the upper limits high and fail to correlate with the real maps.

At 220 GHz there is more power in the real map (due to dust and higher noise levels; see BK-X), so the large fluctuations in the cross with real are not unexpected. It seems likely that the uncertainties are underestimated. While Table 4 indicates that the beam map statistical errors are a factor of 20 lower than the auto spectrum, it is possible that systematics in the beam map measurement contribute extra variance to the cross spectrum that is not captured in the jackknife estimate. For example, visual inspection of the beam maps shows that contamination from the ground and mast is much more prominent at 220 GHz than at 95 and 150 GHz, and may leak into pixels that are not spatially masked at a level below the noise of an individual map. Future work with more rigorous masking of fixed structure is likely to improve the quality of these maps. Given these caveats, the recovered $\rho = (2.4 \pm 1.4) \times 10^{-2}$ should not be taken as evidence for leakage.

Taken together, the beam map simulations do not definitively detect $T \rightarrow P$ leakage in cross-correlation with CMB. The 95 and 150 GHz maps show a $\sim 1\sigma$ excess, which cannot exclude zero leakage. At 220 GHz the cross spectra fluctuate enough to suggest underestimated error bars, and no real indication of leakage within the uncertainty. We also note that because this frequency mostly measures dust emission, the sensitivity of r recovery on its leakage is weak.

6. The Effect of Undeprojected Residuals on Parameter Recovery

The BICEP/Keck likelihood analysis uses maps at several frequencies to separate the CMB from foregrounds (see BK-X). The single-frequency ρ estimators of $T \rightarrow P$ leakage that we used in Section 5 are therefore not representative of the effective bias on r that we may incur from this potential additive systematic once all spectra are accounted for. In this Section, we discuss several methods of dealing with systematics in analysis. Using the auto and cross spectra presented in Section 5.2, we then simulate the potential effect of this leakage on r recovery in the multicomponent likelihood analysis.

6.1. Treating Systematics in Analysis

How we deal with systematics in analysis depends on the uncertainty in the systematic’s form. In the case of a well-characterized effect—i.e., in which the specific morphology of the $T \rightarrow P$ leakage is known—we would simply subtract the leaked signal in the time or map domain.

If we had intermediate knowledge of the systematic—such as a well-characterized amplitude, but not specific form—we would calculate the leakage bandpowers and debias them from the real data. The uncertainty on the debias would need to be included, e.g., by inflating the bandpower covariance matrix.

Finally, if there were substantial uncertainty on both the form and amplitude—e.g., if the uncertainty on the systematic is comparable to the estimated level of the systematic itself—there is little argument for debiasing. In this case, we would run simulations to determine whether the likely amplitude is small compared to the experiment’s statistical uncertainty.




Figure 7. Apodized Q maps of $T \rightarrow P$ leakage predicted by the beam map simulations (left column) and a beam map noise realization (right column). The visible signal is due to higher-order undeprojected residuals in the measured differential beam maps and potentially systematics in the beam measurement. The leakage at 95 GHz is mostly E -mode in character and is due to the anomalous feature in the 95 GHz tile edge pixels, shown in Figure 6. The band of lower leakage at the center of the 150 GHz map is due to the complex cancellation of the 17 receiver-years that compose the map.

In Section 5.3 we analyzed the cross spectra of the beam map simulations with the BK15 maps and found that the uncertainty on the predicted $T \rightarrow P$ leakage is comparable to its amplitude (Table 4). Since we cannot verify the amplitude of this leakage in the real CMB maps at high confidence, we will show that when propagated through the multicomponent likelihood, this potential systematic is subdominant compared to the statistical errors.

6.2. Simulating $T \rightarrow P$ Leakage in the Multicomponent Analysis

Here we use simulations to find the bias on r that the BK15 likelihood analysis would incur if $T \rightarrow P$ leakage levels consistent with either the beam map auto spectra or the cross spectra between beam maps and CMB data existed in the real maps. We also test the possibility of fitting for and marginalizing over a leakage template.

The BK15 likelihood analysis uses all available BB auto and cross spectra between the BK15 95, 150, and 220 GHz maps and external Planck and WMAP maps to generate a joint likelihood of the data for a particular parametric model. For details on the likelihood implementation see BICEP2/Keck Array/Planck Collaboration (2015), Keck Array/BICEP2 Collaboration VI (2016), and BK-X.

To gauge the impact of $T \rightarrow P$ leakage on r recovery, we analyze the shift in the maximum-likelihood r value (r_{ML}) for sets of 499 simulated bandpowers that have had bias added corresponding to one of the $T \rightarrow P$ leakage estimates. Note that we calculate Δr for the peak of the multi-dimensional likelihood, not the peak of the marginalized posterior probability density function; we find that this provides a more direct view of biases in the analysis and is easily extended to nonphysical negative values of r to avoid truncation of the distribution. Histograms of r_{ML} from simulations are used as validation in the baseline BK15 analysis (Figure 20 of BK-X),

and we use the standard deviation of that distribution as a measure of experimental sensitivity, $\sigma(r) = 0.020$ for BK15.

For each $T \rightarrow P$ leakage scenario, we look at the distribution of realization-by-realization shifts in r_{ML} relative to a baseline that does not include leakage in the simulation bandpowers and does not consider leakage parameters in likelihood analysis (this baseline analysis exactly corresponds to Figure 20 of BK-X). For the upper limit scenarios we report the median of the 499 shifts, while for the others we report the 16th, 50th, and 84th percentiles to reflect potentially asymmetric distributions. Details of each scenario are listed below; the results are illustrated in Figure 9.

1. Two “upper limit-driven” scenarios: We inject leakage consistent with the beam map simulation auto spectra (black lines in Figure 8), and ignore it in the likelihood analysis. Recall that the auto spectra are upper limits since they may be biased by systematics in the beam measurement. If we simply add the full beam map auto spectra (Figure 8 black curves) to the simulations—a situation we consider extremely unlikely—we find a median bias of $\Delta r = 0.0084$ (Scenario 1). To better reflect our belief in the auto spectra as 95% upper limits, in a second scenario we inject leakage with the same shape as the auto spectra, but with variable amplitude drawn from a Gaussian distribution centered at zero with $\sigma = 0.5$ of the nominal amplitude, and truncated at zero so only positive leakage can be added. The median bias is $\Delta r = 0.0042$ (Scenario 2).
2. A “CMB data-driven” scenario: We inject leakage consistent with the cross spectra between the beam map simulations and the BK15 maps, and ignore it in likelihood analysis. Specifically, each realization is biased by a leakage contribution that is randomly drawn from the Figure 8 teal points and error bars. This is an attempt to model the leakage that appears to actually exist




Figure 8. *BB* power spectra from beam map simulations, corresponding to the maps shown in Figure 7. The black lines are the per-frequency auto spectra, which have been noise-debiased using the beam map jackknife maps. They should be considered upper limits on $T \rightarrow P$ leakage. The teal crosses show the cross spectra of the beam map simulations with the BK15 maps. The error bars are derived from the cross spectra of the fixed beam map simulation with 499 CMB lensing + dust + instrumental/atmospheric noise simulations. Noise and systematics in the beam map measurement are not included in this error estimate.




Figure 9. Shifts in maximum-likelihood r with respect to the baseline analysis for a set of 499 simulations, which have had varying levels of $T \rightarrow P$ leakage injected and with various recovery scenarios. The numbers refer to scenarios listed in Section 6.2. For scenarios 3–6, the 16th, 50th, and 84th percentiles are plotted. The inset panel shows histograms of r_{ML} for the baseline analysis (blue) and Scenario 3 (black) to emphasize that the potential bias due to $T \rightarrow P$ leakage is much smaller than the BK15 statistical error, $\sigma(r) = 0.020$.

Table 4
Beam Map Simulation ρ Estimates

	95 GHz	150 GHz	220 GHz
Upper limit	1.3×10^{-2}	1.0×10^{-2}	8×10^{-3}
Beam map jackknife	1×10^{-6}	1×10^{-5}	4×10^{-4}
BK15 cross spectrum	4×10^{-3}	5×10^{-3}	2.4×10^{-2}
Cross spectrum σ	4×10^{-3}	4×10^{-3}	1.4×10^{-2}

Note. Single-frequency ρ estimates (i.e., equivalent r level) for the beam map simulation auto spectra (upper limit), a probe of the beam map uncertainty (beam map jackknife), the cross spectra with real BK15 maps, and uncertainty in the cross spectrum arising from the CMB map (cross spectrum σ).

in the maps, albeit at marginal significance. The recovered bias is $\Delta r = 0.0027 \pm 0.0019$ (Scenario 3).

3. Three “recovery” scenarios: Using knowledge of the leakage from beam map simulations, we attempt to marginalize over it in the likelihood. For each frequency we add a new parameter to the likelihood analysis, which scales the leakage bias contribution to the bandpowers. For example, the parameter $\eta_{95} = 0$ represents zero $T \rightarrow P$ leakage in the 95 GHz map, while $\eta_{95} = 1$ indicates leakage equal to that shown in the left panel of Figure 8. For this analysis, we must select one of the two leakage estimates to use as a template—either the beam map simulation auto spectrum or the beam map simulation cross spectrum with BK15 maps. We use a flat prior on η in the range of $[0, 2]$. If the template used in the analysis matches the leakage added to the simulations, the resulting bias on r is small: $\Delta r = -0.0004^{+0.0024}_{-0.0022}$ (Scenario 4). If the *wrong* template is used (e.g., inject the beam map/CMB cross spectrum, but fit for the beam map auto spectrum), we incur a negative bias: $\Delta r = -0.0013^{+0.0035}_{-0.0042}$ (Scenario 5). Finally, if *no leakage* is injected but we attempt to fit for it, we still find negative bias with $\Delta r = -0.0014^{+0.0014}_{-0.0027}$ (Scenario 6).

We take the CMB data-driven Scenario 3, $\Delta r = 0.0027 \pm 0.0019$, as our best estimate of the bias on r from $T \rightarrow P$ leakage. Figure 9 illustrates that in all cases the bias is subdominant compared to the BK15 statistical error, $\sigma(r) = 0.020$. Quadrature addition of the worst-case upper limit bias inflates $\sigma(r)$ by 8%.

In the standard COSMOMC analysis of e.g., BK15, physical priors are imposed. If we introduce a component that (i) is partially degenerate with the signal of interest (r), (ii) is not clearly detected at the available noise level, and (iii) has a positive-only prior applied, r tends to be biased downwards. Since we do not clearly detect leakage in the current data, we decide not to marginalize over the η parameters in the r constraint analysis at this time. This is analogous to the treatment of the dust decorrelation parameter as discussed in Appendix F of BK-X. We will continuously review this situation going forward.

7. Conclusions

In this paper we presented far-field optical characterization of Keck Array detectors at 95, 150, and 220 GHz contributing to the BK15 data set. From an extensive far-field beam mapping campaign we measured differential Gaussian parameters with high precision, which are repeatable from year to year. For each detector we formed deep composite maps covering all azimuthal angles, and from them generated array-averaged beam profiles.

The composite beam maps were then used to predict the residual main beam $T \rightarrow P$ leakage expected in the BK15 polarization maps after deprojection of the lowest-order beam mismatch modes. From an auto spectrum analysis of beam map simulations, noting that there may exist contributions from low-level systematics in the beam measurement, we presented upper limits to the leakage. Cross spectra between the beam map simulations and the BK15 CMB maps offer tentative evidence for leakage, but the uncertainties are large enough that zero leakage cannot be excluded.

We have run simulations using the BK15 multicomponent likelihood analysis to test the effect of undeprojected $T \rightarrow P$ leakage on r recovery. When leakage consistent with the cross spectra between beam map simulations and real CMB maps is added, the bias is $\Delta r = 0.0027 \pm 0.0019$. It is possible to marginalize over this contribution in the multicomponent analysis, but this admits the possibility of a small negative bias if the wrong template is used with physical priors. Because the leakage is not clearly detected in the CMB data, we do not marginalize over it in the current constraint on r . All of the biases presented are small compared to the BK15 statistical uncertainty $\sigma(r) = 0.020$.

BICEP's sensitivity to r will continue to steadily improve: with data taken through the 2017 season we expect $\sigma(r) \sim 0.010$, and with the BICEP Array experiment under construction we anticipate $\sigma(r) < 0.005$ within 5 years (Hui et al. 2018). Constraints on beam systematics will need to similarly tighten. Future effort will focus on three aspects of the problem: the intrinsic $T \rightarrow P$ leakage level, the measurement thereof, and treatment in analysis.

High-fidelity beam maps point to features in individual beams that contribute significant $T \rightarrow P$ leakage, which can then be remedied in hardware. Such feedback has been critical to ensuring that leakage is reduced in later generations of receivers. For example, the 95 GHz difference beam shown in Figure 6 highlighted an anomalous interaction between the tile corrugations and edge detectors. Compared to detectors that were not affected, the tile edge detectors drove up the estimated $T \rightarrow P$ leakage for the BK15 results by over an order of magnitude. We corrected this effect in focal planes produced subsequently, and expect that this feedback cycle will continue.

Looking beyond the BK15 result, we have already improved the beam map reduction compared to the maps presented in this paper: non-Gaussian noise has been significantly reduced through more optimal low-level deglitching and demodulation. Re-reduction of existing data will produce cleaner maps and allow us to mean-filter the component maps when generating the composites. In future beam mapping campaigns we also plan to modify the raster scan strategy to produce robust noise estimates, e.g., by using out-and-back scans at the same elevation to form a “scan direction” jackknife. Maps generated from these data will be used to generate beam map noise

realizations without dividing the composite maps into two halves. Their enhanced statistical properties will facilitate comparison of $T \rightarrow P$ leakage to CMB data.

While in this paper we have shown a basic attempt to detect $T \rightarrow P$ leakage by cross-correlating the final coadded leakage and real CMB maps (in which much of the leakage has canceled), this comparison can be improved. In our next results we plan to form cross spectra that detect the predicted leakage at high significance. We will isolate high-leakage detector subsets or form combinations of on-sky data that are expected to enhance the leakage (e.g., boresight angles that do not cancel) in particular sub-maps compared to the full data set. Given these higher-confidence estimators of $T \rightarrow P$ leakage in the CMB maps, we will explore several options to remove the effect. If the number of high-leakage detectors is small, excluding or de-weighting them will dramatically improve systematic control. With improved statistical properties of the beam maps, debiasing as discussed in Section 6.1 would also be reasonable. Finally, we can further reduce $T \rightarrow P$ residuals by deprojecting additional modes of this leakage for each detector pair. This can be done using modes drawn from bases independent of the beam maps, by using the beam maps to directly predict the template of each pair's undeprojected residual, or by using these maps to guide definition of a small subset of modes to be deprojected.

The Keck Array project has been made possible through support from the National Science Foundation under grants ANT-1145172 (Harvard), ANT-1145143 (Minnesota), and ANT-1145248 (Stanford), and from the Keck Foundation (Caltech). The development of antenna-coupled detector technology was supported by the JPL Research and Technology Development Fund and grant Nos. 06-ARPA206-0040 and 10-SAT10-0017 from the NASA APRA and SAT programs. The development and testing of focal planes were supported by the Gordon and Betty Moore Foundation at Caltech. Readout electronics were supported by a Canada Foundation for Innovation grant to UBC. The computations in this paper were run on the Odyssey 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 Robert Schwarz and Steffen Richter. We thank all those who have contributed past efforts to the BICEP-Keck Array series of experiments, including the BICEP1 team. We also thank the Planck and WMAP teams for the use of their data.

ORCID iDs

P. A. R. Ade  <https://orcid.org/0000-0002-5127-0401>
D. Barkats  <https://orcid.org/0000-0002-8971-1954>
C. A. Bischoff  <https://orcid.org/0000-0001-9185-6514>
B. P. Crill  <https://orcid.org/0000-0002-4650-8518>
M. Dierickx  <https://orcid.org/0000-0002-3519-8593>
J. P. Filippini  <https://orcid.org/0000-0001-8217-6832>
M. Halpern  <https://orcid.org/0000-0002-1760-0868>
K. S. Karkare  <https://orcid.org/0000-0002-5215-6993>
L. Monceli  <https://orcid.org/0000-0002-4242-3015>
W. L. K. Wu  <https://orcid.org/0000-0001-5411-6920>

References

- Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, *ApJS*, **208**, 20
- BICEP1 Collaboration 2014, *ApJ*, **783**, 67
- BICEP2 Collaboration II 2014, *ApJ*, **792**, 62
- BICEP2 Collaboration III 2015, *ApJ*, **814**, 110
- BICEP2/Keck Array Collaboration IV 2015, *ApJ*, **806**, 206
- BICEP2/Keck Array Collaboration VI 2016, *PhRvL*, **116**, 031302
- BICEP2/Keck Array Collaboration VII 2016, *ApJ*, **825**, 66
- BICEP2/Keck Array Collaboration VIII 2016, *ApJ*, **833**, 228
- BICEP2/Keck Array Collaboration X 2018, *PhRvL*, **121**, 221301
- BICEP2/Keck Array/Planck Collaboration 2015, *PhRvL*, **114**, 101301
- Buder, I., Ade, P. A. R., Ahmed, Z., et al. 2014, *Proc. SPIE*, **9153**, 915312
- CMB-S4 Science Book 2016, arXiv:1610.02743
- Conklin, E. K. 1969, *Natur*, **222**, 971
- Hanson, D., Hoover, S., Crites, A., et al. 2013, *PhRvL*, **111**, 141301
- Hu, W., Hedman, M. M., & Zaldarriaga, M. 2003, *PhRvD*, **67**, 043004
- Hui, H., Ade, P. A. R., Ahmed, Z., et al. 2018, *Proc. SPIE*, **10708**, 1070807
- Kamionkowski, M., & Kovetz, E. D. 2016, *ARA&A*, **54**, 227
- Karkare, K. S. 2017, PhD thesis, Harvard Univ., <https://dash.harvard.edu/handle/1/40046430>
- Karkare, K. S., Ade, P. A. R., Ahmed, Z., et al. 2016, *Proc. SPIE*, **9914**, 991430
- Keisler, R., Hoover, S., Harrington, N., et al. 2015, *ApJ*, **807**, 151
- Kovac, J. M., Leitch, E. M., Pryke, C., et al. 2002, *Natur*, **420**, 772
- Louis, T., Grace, E., Hasselfield, M., et al. 2017, *JCAP*, **2017**, 031
- O'Brient, R., Ade, P. A. R., Ahmed, Z., et al. 2012, *Proc. SPIE*, **8452**, 84521G
- Penzias, A. A., & Wilson, R. W. 1965, *ApJ*, **142**, 419
- Planck Collaboration, Adam, R., Ade, P. A. R., et al. 2016, *A&A*, **594**, A1
- Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2018, *A&A*, in press, arXiv: 1807.06210
- Polarbear Collaboration 2014, *ApJ*, **794**, 171
- Smoot, G. F., Bennett, C. L., Kogut, A., et al. 1992, *ApJL*, **396**, L1
- Wong, C. L. 2014, PhD thesis, Harvard Univ., <https://pqdtopen.proquest.com/doc/1625428252.html?FMT=ABS>