Detection of *B*-Mode Polarization at Degree Angular Scales by BICEP2

P. A. R. Ade,[1] R. W. Aikin,[2] D. Barkats,[3] S. J. Benton,[4] C. A. Bischoff,[5] J. J. Bock,[2,6] J. A. Brevik,[2] I. Buder,[5] E. Bullock,[7] C. D. Dowell,[6] L. Duband,[8] J. P. Filippini,[2] S. Fliescher,[9] S. R. Golwala,[2] M. Halpern,[10] M. Hasselfield,[10] S. R. Hildebrandt,[2,6] G. C. Hilton,[11] V. V. Hristov,[2] K. D. Irwin,[12,13,11] K. S. Karkare,[5] J. P. Kaufman,[14] B. G. Keating,[14] S. A. Kernasovskiy,[12] J. M. Kovac,[5,*] C. L. Kuo,[12,13] E. M. Leitch,[15] M. Lueker,[2] P. Mason,[2] C. B. Netterfield,[4,16] H. T. Nguyen,[6] R. O'Brient,[6] R. W. Ogburn IV,[12,13] A. Orlando,[14] C. Pryke,[9,7,†] C. D. Reintsema,[11] S. Richter,[5] R. Schwarz,[9] C. D. Sheehy,[9,15] Z. K. Staniszewski,[2,6] R. V. Sudiwala,[1] G. P. Teply,[2] J. E. Tolan,[12] A. D. Turner,[6] A. G. Vieregg,[5,15] C. L. Wong,[5] and K. W. Yoon[12,13]

(BICEP2 Collaboration)

[1]*School of Physics and Astronomy, Cardiff University, Cardiff, CF24 3AA, United Kingdom*
[2]*Department of Physics, California Institute of Technology, Pasadena, California 91125, USA*
[3]*Joint ALMA Observatory, Vitacura, Santiago, Chile*
[4]*Department of Physics, University of Toronto, Toronto, Ontario, M5S 1A7, Canada*
[5]*Harvard-Smithsonian Center for Astrophysics, 60 Garden Street MS 42, Cambridge, Massachusetts 02138, USA*
[6]*Jet Propulsion Laboratory, Pasadena, California 91109, USA*
[7]*Minnesota Institute for Astrophysics, University of Minnesota, Minneapolis, Minnesota 55455, USA*
[8]*Service des Basses Températures, Commissariat à l'Energie Atomique, 38054 Grenoble, France*
[9]*Department of Physics, University of Minnesota, Minneapolis, Minnesota 55455, USA*
[10]*Department of Physics and Astronomy, University of British Columbia, Vancouver, British Columbia, V6T 1Z1, Canada*
[11]*National Institute of Standards and Technology, Boulder, Colorado 80305, USA*
[12]*Department of Physics, Stanford University, Stanford, California 94305, USA*
[13]*Kavli Institute for Particle Astrophysics and Cosmology, SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, California 94025, USA*
[14]*Department of Physics, University of California at San Diego, La Jolla, California 92093, USA*
[15]*University of Chicago, Chicago, Illinois 60637, USA*
[16]*Canadian Institute for Advanced Research, Toronto, Ontario, M5G 1Z8, Canada*

We report results from the BICEP2 experiment, a cosmic microwave background (CMB) polarimeter specifically designed to search for the signal of inflationary gravitational waves in the *B*-mode power spectrum around $\ell \sim 80$. The telescope comprised a 26 cm aperture all-cold refracting optical system equipped with a focal plane of 512 antenna coupled transition edge sensor 150 GHz bolometers each with temperature sensitivity of $\approx 300 \ \mu K_{CMB} \sqrt{s}$. BICEP2 observed from the South Pole for three seasons from 2010 to 2012. A low-foreground region of sky with an effective area of 380 square deg was observed to a depth of 87 nK deg in Stokes Q and U. In this paper we describe the observations, data reduction, maps, simulations, and results. We find an excess of *B*-mode power over the base lensed-ΛCDM expectation in the range $30 < \ell < 150$, inconsistent with the null hypothesis at a significance of $> 5\sigma$. Through jackknife tests and simulations based on detailed calibration measurements we show that systematic contamination is much smaller than the observed excess. Cross correlating against WMAP 23 GHz maps we find that Galactic synchrotron makes a negligible contribution to the observed signal. We also examine a number of available models of polarized dust emission and find that at their default parameter values they predict power $\sim (5\text{--}10)\times$ smaller than the observed excess signal (with no significant cross-correlation with our maps). However, these models are not sufficiently constrained by external public data to exclude the possibility of dust emission bright enough to explain the entire excess signal. Cross correlating BICEP2 against 100 GHz maps from the BICEP1 experiment, the excess signal is confirmed with 3σ significance and its spectral index is found to be consistent with that of the CMB, disfavoring dust at 1.7σ. The observed *B*-mode power spectrum is well fit by a lensed-ΛCDM + tensor theoretical model with

tensor-to-scalar ratio $r = 0.20^{+0.07}_{-0.05}$, with $r = 0$ disfavored at 7.0σ. Accounting for the contribution of foreground, dust will shift this value downward by an amount which will be better constrained with upcoming data sets.

I. INTRODUCTION

The discovery of the cosmic microwave background (CMB) by Penzias and Wilson [1] confirmed the hot big bang paradigm and established the CMB as a central tool for the study of cosmology. In recent years, observations of its temperature anisotropies have helped establish and refine the "standard" cosmological model now known as ΛCDM, under which our universe is understood to be spatially flat, dominated by cold dark matter, and with a cosmological constant (Λ) driving accelerated expansion at late times. CMB temperature measurements have now reached remarkable precision over angular scales ranging from the whole sky to arcmin resolution, producing results in striking concordance with predictions of ΛCDM and constraining its key parameters to subpercent precision (e.g., [2–9]).

Inflationary cosmology extends the standard model by postulating an early period of nearly exponential expansion which sets the initial conditions for the subsequent hot big bang. It was proposed and developed in the early 1980s to resolve mysteries for which the standard model offered no solution, including the flatness, horizon, smoothness, entropy, and monopole problems ([10–17]; see [18] for a review). Inflation also explains the Universe's primordial perturbations as originating in quantum fluctuations stretched by this exponential expansion [19–24], and thus to be correlated on superhorizon scales. The simplest models further predict these perturbations to be highly adiabatic and Gaussian and nearly scale invariant (though typically slightly stronger on larger scales). These qualities, though not necessarily unique to the inflationary paradigm, have all been confirmed by subsequent observations (e.g., [25,26], and references above). Although highly successful, the inflationary paradigm represents a vast extrapolation from well-tested regimes in physics. It invokes quantum effects in highly curved spacetime at energies near 10^{16} GeV and time scales less than 10^{-32} s. A definitive test of this paradigm would be of fundamental importance.

Gravitational waves generated by inflation have the potential to provide such a definitive test. Inflation predicts that the quantization of the gravitational field coupled to exponential expansion produces a primordial background of stochastic gravitational waves with a characteristic spectral shape ([27–31]; also see [32,33]). Though unlikely to be directly detectable in modern instruments, these gravitational waves would have imprinted a unique signature upon the CMB. Gravitational waves induce local quadrupole anisotropies in the radiation field within the last-scattering surface, inducing polarization in the scattered light [34]. This polarization pattern will include a "curl" or B-mode component at degree angular scales that cannot be generated primordially by density perturbations. The amplitude of this signal depends upon the tensor-to-scalar ratio r, which itself is a function of the energy scale of inflation. The detection of B-mode polarization of the CMB at large angular scales would provide a unique confirmation of inflation and a probe of its energy scale [35–37].

The CMB is polarized with an amplitude of a few μK, dominated by the "gradient" or E-mode pattern that is generated by density perturbations at last scattering. These E modes peak at angular scales of $\sim 0.2°$, corresponding to angular multipole $\ell \approx 1000$. They were first detected by the DASI experiment [38]. Since then multiple experiments have refined measurements of the EE power spectrum, including CAPMAP [39,40], CBI [41,42], BOOMERANG03 [43], WMAP [2,44], MAXIPOL [45], QUAD [46,47], BICEP1 [48,49], and QUIET [50,51].

Gravitational lensing of the CMB's light by large scale structure at relatively late times produces small deflections of the primordial pattern, converting a small portion of E-mode power into B modes. The lensing B-mode spectrum is similar to a smoothed version of the E-mode spectrum but a factor ~ 100 lower in power, and hence also rises toward subdegree scales and peaks around $\ell \approx 1000$. The inflationary gravitational wave (IGW) B mode, however, is predicted to peak at multipole $\ell \approx 80$ and this creates an opportunity to search for it around this scale where it is quite distinct from the lensing effect. (This is the so-called "recombination bump." There is another opportunity to search for the IGW signal at $\ell < 10$ in the "reionization bump," but this requires observations over a substantial fraction of the full sky.)

A large number of current CMB experimental efforts now target B-mode polarization. Evidence for lensing B-mode polarization at subdegree scales has already been detected by two experiments in the past year, first by the SPT telescope [52] and more recently by POLARBEAR [53–55]. The search for inflationary B modes at larger scales will also be a goal of these experiments, as well as other ongoing experimental efforts in the U.S. that include the ABS [56], ACTPOL [57], and CLASS [58] ground-based telescopes and the EBEX [59], SPIDER [60], and PIPER [61] balloon experiments, each employing a variety of complementary strategies. It is also a major science goal of the ESA *Planck* satellite mission.

The BICEP/*Keck Array* series of experiments have been specifically designed to search for primordial *B*-mode polarization on degree angular scales by making very deep maps of relatively small patches of sky from the South Pole. The BICEP1 instrument initiated this series [62], observing from 2006 to 2008. Its initial results were described in Takahashi *et al.* [63] and Chiang *et al.* [48] (hereafter T10 and C10), and final results were recently reported in Barkats *et al.* [49] (hereafter B14) yielding a limit of $r < 0.70$ at 95% confidence.

In this paper we report results from BICEP2—a successor to BICEP1 which differed principally in the focal plane where a very large increase in the detector count resulted in more than an order of magnitude improvement in mapping speed. The observation field and strategy were largely unchanged, as were the telescope mount, observation site, etc. Using all three seasons of data taken with BICEP2 (2010–2012) we detect *B*-mode power in the multipole range $30 < \ell < 150$, finding this power to have a strong excess inconsistent with lensed ΛCDM at $> 5\sigma$ significance.

The structure of this paper is as follows. In Secs. II and III we briefly review the BICEP2 instrument, observations, and low-level data reduction deferring details to a related paper [64] (hereafter the Instrument Paper). In Sec. IV we describe our map-making procedure and present signal and signal-differenced "jackknife" *T*, *Q*, and *U* maps which have unprecedented sensitivity. This section introduces "deprojection" of modes potentially contaminated through beam systematics, which is an important new technique. In Sec. V we describe our detailed time stream-level simulations of signal and pseudosimulations of noise. In Sec. VI we describe calculation of the power spectra, including matrix-based *B*-mode purification. In Sec. VII we present the signal and jackknife power spectrum results for *TE*, *EE*, *BB*, *TB*, and *EB*. In Sec. VIII we discuss and summarize the many studies we have conducted probing for actual and potential sources of systematic contamination, and argue that residual contamination is much smaller than the detected *B*-mode signal. Full details are deferred to a related paper [65] (hereafter the Systematics Paper). In Sec. IX we investigate foreground projections and constraints based on external data and conclude that it is implausible that the *B*-mode signal which we see is dominated by synchrotron, and that the present data disfavor domination by dust or any other known foreground source. In Sec. X we take cross spectra of the BICEP2 maps with those from BICEP1 (as presented in B14) and find that the spectral signature of the signal is consistent with the CMB. Finally in Sec. XI we calculate some simple, largely phenomenological, parameter constraints, and conclude in Sec. XII.

II. THE BICEP2 INSTRUMENT

BICEP2 was similar to BICEP1 (see T10) reusing the same telescope mount and installation at the South Pole. Like BICEP1 the optical system was a simple 26 cm aperture all-cold refractor housed entirely in a liquid helium cooled cryostat. The main differences from BICEP1 were the use of a focal plane array of planar antenna-coupled devices [66] with voltage-biased transition-edge sensor (TES) detectors [67] and a multiplexed superconducting quantum interference device (SQUID) readout. BICEP2 observed at 150 GHz only. A very brief review of the instrument follows—for more details please refer to the Instrument Paper.

A. Optics

The optics were adapted from the original BICEP1 design [62]. Light entered the cryostat through a polypropylene foam window, passed through polytetrafluoroethylene filters cooled to 100 and 40 K, and then through polyethylene objective and eyepiece lenses cooled to 4 K. A 26.4 cm diameter aperture stop was placed at the objective lens and an additional nylon filter was placed on the sky side of the eyepiece lens. All the lenses and filters were antireflection coated and the interior of the optics tube was lined with microwave absorber. The optics were designed to be telecentric (flat focal plane) and the resulting beams had a full width at half maximum of $\approx 0.5°$. An absorptive fore baffle was mounted on the front of the telescope which was designed to prevent radiation from boresight angles greater than $\sim 20°$ entering the telescope. The telescope was located inside a large stationary reflective ground screen.

B. Focal plane

The BICEP2 focal plane employed monolithic arrays of antenna coupled TES detectors designed and fabricated at Caltech and JPL. Each pixel was composed of two interleaved 12×12 arrays of orthogonal slot antennas feeding beam-forming (phased-array) summing trees. The output of each summing tree was a microstrip which passed through a band-defining filter and deposited its power on a thermally isolated island. Changes in the power incident on this island were detected using a transition edge sensor (TES). There was an 8×8 array of pixels on each tile, and four such tiles were combined to form the complete focal plane unit. There were thus, in principle, 256 dual-polarization pixels in the focal plane for a total of 512 detectors, each with temperature sensitivity of $\approx 300~\mu K_{CMB}\sqrt{s}$. (Six pixels were deliberately disconnected between antenna and TES sensor to provide diagnostic "dark" channels.) The focal plane was cooled to 270 mK by a closed cycle three-stage sorption refrigerator.

C. Detector readout and data acquisition system

The TES detectors were read out through time-division SQUID multiplexing chips provided by NIST. A single readout channel was connected in rapid succession to 32 detectors, reducing wiring and heat load requirements.

These SQUID systems were biased and read out by a multi-channel electronics (MCE) crate external to the cryostat (provided by UBC). The sample rate stored to disk was 20 Hz. The housekeeping and readout electronics were connected to a set of Linux-based computers running a control system called GCP, which has been used by many recent ground-based CMB experiments [68].

D. Telescope mount

The receiver cryostat was mounted on a three-axis mount able to move in azimuth and elevation and to rotate the entire telescope about its boresight. Hereafter, we refer to the line-of-sight rotation angle as the "deck" angle. The window of the telescope looked out through an opening in a flexible environmental seal such that the cryostat, mount, and electronics were all located in a room temperature laboratory environment.

III. OBSERVATIONS AND LOW-LEVEL DATA REDUCTION

A. Observations

BICEP2 observed on a three day schedule locked to sidereal time. As in BICEP1, the basic unit of observation was a ≈50 minute "scan set" during which the telescope scanned back-and-forth 53 times at $2.8°$ s^{-1} in azimuth in a smooth turnaround triangle wave pattern, with a throw of ≈60°, at fixed elevation. We refer to each of the 106 motions across the field (either left- or right-going) as a "half scan." We do not use the turnaround portions of the scans in this analysis.

BICEP2 observed the same CMB field as BICEP1—a low foreground region centered at RA 0h, Dec. −57.5°. At the South Pole, the elevation angle is simply the negative of declination and azimuth maps to RA shifting by 15° per hour. The scan speed on the sky was thus ≈1.5° s^{-1} mapping multipole $\ell = 100$ into the time stream at ≈0.4 Hz. At the end of each scan set the elevation was stepped by 0.25° and the azimuth angle updated to recenter on RA 0h. The scans thus "slide" with respect to the sky during each scan set by ≈12.5° allowing us to subtract a scan fixed "template" from the time stream while leaving degree-scale sky structure only slightly attenuated (see Secs. IV A and VI C).

A total of 21 elevation offsets were used between Dec. of −55° and −60°. Note that since the field of view of the focal plane—~20°—is much larger in Dec., and somewhat larger in RA, than the region scanned by the boresight the final coadded map is naturally apodized. After a complete three-day schedule the instrument was rotated to a new deck angle, the refrigerator was recycled, and the process repeated. See the Instrument Paper for more details of the observation strategy.

The control system ran CMB observation schedules relentlessly between early 2010 and late 2012 collecting over 17 000 scan sets of data (≈590 days). (There were some breaks for beam mapping and other calibrations during the austral summers.) The raw data were transferred off site daily via satellite, allowing rigorous quality monitoring and ongoing analysis development. The analysis presented in this paper uses all of the CMB data taken by BICEP2.

B. Analysis pipeline

The analysis pipeline used in this paper is written in the MATLAB language and was originally developed for the QUAD experiment [46]. It was then adapted to BICEP1 data and became the secondary, and then primary, analysis pipeline for the C10 and B14 papers, respectively. For BICEP2 it has seen substantial further development including the addition of a sophisticated automatic data selection framework, full deprojection of beam systematics, and a map-based B-mode purification operation; these enhancements are detailed below.

C. Transfer function correction and deglitching

Starting from the raw time streams, the first step of the pipeline is to deconvolve the temporal response of the instrument. The TES detectors themselves have a very fast and uniform response at all frequencies of interest. To correct for the effect of the digital low-pass filtering, which was applied to the data before it was down sampled for recording to disk, we apply an FIR deconvolution operation in the time domain (which also reapplies a zero-delay low pass filter). Glitches and flux jumps in the SQUID readout are also corrected and/or flagged at this point—they are relatively rare in these data. See the Instrument Paper for more details.

D. Relative gain calibration

At the beginning and end of each scan set an elevation nod or "elnod" was performed. The telescope was moved up-down-up or vice versa in a roughly sinusoidal excursion in time, injecting a signal proportional to the atmospheric opacity gradient into the detector time streams. In analysis, each elnod is regressed against the air-mass profile through which it was looking to derive a relative gain coefficient in SQUID feedback units per air mass. The time stream for each scan set is then divided by its own elnod coefficient and multiplied by the median over all good detectors. This roughly equalizes the gain of each channel and results in considerable cancellation of atmospheric fluctuations when taking the difference of detector pairs, thus making the data considerably easier to work with. The relative gain as determined using the atmospheric gradient is not necessarily the relative gain which minimizes leakage of CMB temperature anisotropy to polarization—see Sec. VIII. Absolute calibration is deferred until after the final coadded map is made—see Sec. IV H.

E. First round data cuts

At this point in the data reduction, individual channels are cut at per half-scan granularity. Reasons for removal include glitches and flux jumps in the channel in question, or its multiplex neighbors, and synchronization problems in the data acquisition system. BICEP2 data are very well behaved and over 90% of the data pass this stage.

IV. MAP MAKING

A. Time stream filtering

In the next step the sum and difference of each detector pair is taken, the pair sum being ultimately used to form maps of temperature anisotropy, and the pair difference to measure polarization. Each half scan is then subjected to a third-order polynomial filtering.

Each half scan is constrained to have the same number of time samples. In addition to the polynomial filtering we also perform a "template" subtraction of any scan-synchronous component by averaging together the corresponding points over a scan set and removing the result from each half scan. Forward and backward half scans are treated separately.

Within our simulation-based analysis framework we are free to perform any arbitrary filtering of the data which we choose. Although any given filtering implies some loss of sensitivity due to the removal and mixing of modes within the map these effects are corrected as described in Secs. VI B and VI C. We defer discussion of the particular filtering choices made in this analysis to Sec. VIII.

B. Pointing reconstruction

The pointing trajectory of the telescope boresight (i.e., the line-of-sight axis of rotation of the mount) is determined using a mount pointing model calibrated using a star camera as described in the Instrument Paper. To convert time stream into maps it is then necessary to know the pointing offset of each detector from this direction. To measure these we first make per channel maps assuming approximate offsets, and then regress these against the WMAP5 temperature map to determine corrections. Comparing maps made from left-going and right-going scans at each of the four deck angles, we estimate that this procedure is accurate to better than 0.05° absolute pointing uncertainty. The beam positions relative to the boresight are averaged over the scan directions and deck angles to produce a single reconstruction for each detector used in map making.

C. Construction of deprojection time stream

The two halves of each detector pair would ideally have identical angular response patterns (beams) on the sky. If this is not the case, then leakage of temperature anisotropy (pair sum) to polarization (pair difference) will occur [69].

One can resample an external map of the temperature sky and its derivatives to generate templates of the leakage resulting from specific differential beam effects. In this analysis we smooth the *Planck* 143 GHz map [70] using the average measured beam function and resample following the procedure described in Ref. [71] and the Systematics Paper. Our standard procedure is to calculate templates for the six modes which correspond to differences of elliptical Gaussian beams. In practice, we do not normally use all six—see Secs. IV F and VIII.

D. Binning into pair maps

At this point we bin the pair-sum and pair-difference signals into per-scan set, per-pair RA-Dec. pixel grids which we refer to as "pair maps." The pixels are 0.25° square at declination −57.5°. The data from each scan set are weighted by their inverse variance over the complete scan set (with separate weights for pair sum and pair difference). We note that while the pair-sum weights vary widely due to variation in atmospheric $1/f$ noise, the pair-difference weights are extremely stable over time—i.e., atmospheric fluctuations are empirically shown to be highly unpolarized. For pair difference a number of products of the time stream and the sine and cosine of the polarization angle are recorded to allow construction of Q and U maps as described in Sec. IV G below. The deprojection templates are also binned into pair maps in parallel with the pair-difference data.

We use per-pair detector polarization angles derived from a dielectric sheet calibrator (as described in the Instrument Paper). (These derived angles are within 0.2° rms of their design values, well within the required accuracy. However note that we later apply an overall rotation to minimize the high ℓ TB and EB spectra—see Sec. VIII B.) The measured polarization efficiency of our detectors is very high ($\approx$99%, see the Instrument Paper)— we perform a small correction to convert temperature-based gains to polarization gains.

E. Second round data cuts

The per-scan set, per-pair maps are recorded on disk to allow rapid recalculation of the coadded map while varying the so-called "second round" cut parameters. These include a variety of cuts on the bracketing elnods, including goodness of fit to the atmospheric cosecant model and stability in both absolute and pair-relative senses. We also make some cuts based on the behavior of the data themselves, including tests for skewness and noise stationarity. Many of these cuts identify periods of exceptionally bad weather and are redundant with one another. We also apply "channel cuts" to remove a small fraction of pairs— principally those with anomalous measured differential beam shapes. In general BICEP2 data are very well behaved and the final fraction of data retained is 63%. See the Instrument Paper for more details.

F. Accumulation of pair maps to phase, and template regression

Once the second round cuts have been made we accumulate the pair maps over each set of ten elevation steps (hereafter referred to as a "phase"). The deprojection templates are also accumulated. We then regress some of these binned templates against the data—i.e., we effectively find the best fit value for each nonideality, for each pair, within each phase. The templates scaled by the regression coefficients are then subtracted from the data, entirely removing that imperfection mode if present. This operation also filters real signal and noise due to chance correlation (and real TE-induced correlation in the case of signal). This filtering is effectively just additional time stream filtering like that already mentioned in Sec. IV A and we calibrate and correct for its effect in the same way (see Secs. VI B and VI C).

The choice of deprojection time scale is a compromise— reducing it guards against systematic modes which vary over short time scales (as relative gain errors might), while covering more sky before regressing reduces the filtering of real signal (the coefficient is fit to a greater number of pixels). In practice reduction of the filtering going from ten elevation steps to twenty is found to be modest and for this analysis we have deprojected modes on a per-phase basis.

We also have the option to fix the coefficients of any given mode at externally measured values, corresponding to a subtraction of the systematic with no additional filtering of signal. In this analysis we have deprojected differential gain and pointing, and have subtracted the effects of differential ellipticity—we defer discussion of these particular choices to Sec. VIII.

G. Accumulation over phases and pairs

We next proceed to coadded maps accumulating over phases and pairs. Full coadds are produced as well as many "jackknife" splits—pairs of maps made from two subsets of the data which might be hypothesized to contain different systematic contaminations. Some splits are strictly temporal (e.g., first half vs second half of the observations), some are strictly pair selections (e.g., inner vs outer part of each detector tile), and some are both temporal and pairwise
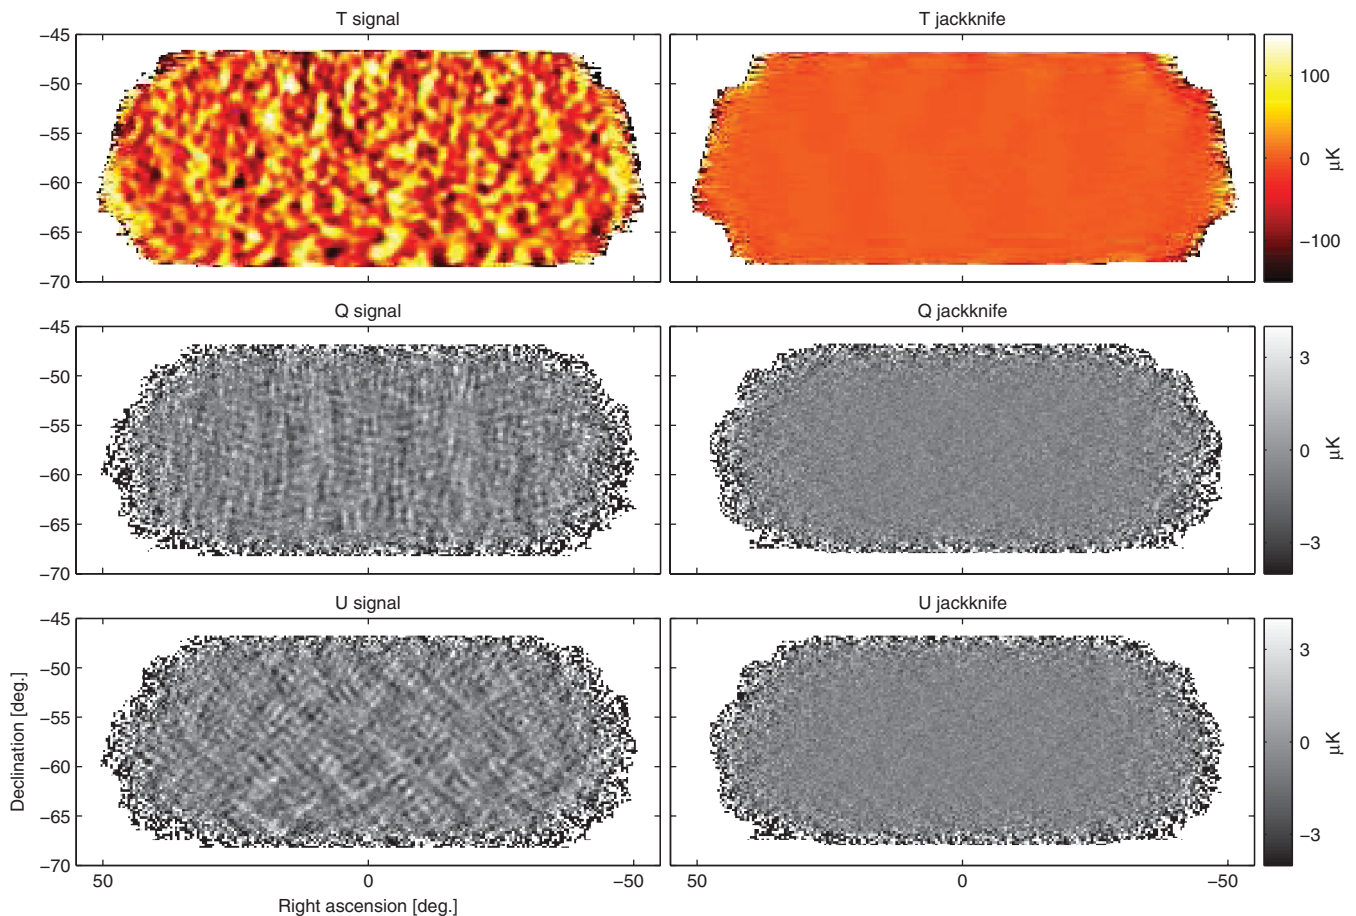

FIG. 1 (color).   BICEP2 T, Q, U maps. The left column shows the basic signal maps with 0.25° pixelization as output by the reduction pipeline. The right column shows difference (jackknife) maps made with the first and second halves of the data set. No additional filtering other than that imposed by the instrument beam (FWHM 0.5°) has been done. Note that the structure seen in the Q and U signal maps is as expected for an E-mode dominated sky.

(e.g., the so-called tile and deck jackknife)—see Sec. VII C for details.

Once the accumulation over all 590 days and $\approx$200 detector pairs is done the accumulated quantities must be converted to T, Q and U maps. For T this is as simple as dividing by the sum of the weights. For Q and U we must perform a simple 2×2 matrix inversion for each pixel. This matrix is singular if a given pixel has been observed at only a single value of the deck angle modulo 90°. In general for BICEP2 data we have angles 68°, 113°, 248°, and 293° as measured relative to the celestial meridian.

We perform absolute calibration by taking the cross spectrum of the T map with either the *Planck* 143 GHz map or the WMAP9 W-band T map as described in the Instrument Paper. We adopt an absolute gain value intermediate to these two measurements and assign calibration uncertainty of 1.3% in the map to account for the difference.

H. Maps

Figure 1 shows the BICEP2 T, Q, and U signal maps along with a sample set of difference (jackknife) maps. The "vertical-stripe-Q, diagonal-stripe-U" pattern indicative of an E-mode dominated sky is visible. Note that these maps are filtered by the relatively large beam of BICEP2 ($\approx$0.5° FWHM). Comparison of the signal and jackknife maps shows that the former are signal dominated—they are the deepest maps of CMB polarization ever made at degree angular scales with an rms noise level of 87 nK in (nominal) $1° \times 1°$ pixels.

V. SIMULATIONS

A. Signal simulations

As is common practice in this type of analysis, we account for the filtering which our instrument and data reduction impose on the underlying sky pattern through simulations [72]. Starting with input T, Q, and U sky maps we smooth using the average measured beam function and then resample along the pointing trajectory of each detector at each time stream sample. We have the option of perturbing to per-channel elliptical Gaussian beam shapes using the derivatives of the map (in a similar manner to the construction of the deprojection templates described in Sec. IV C above). However, for our standard simulations we include only differential pointing as this is our leading order beam imperfection (see Sec. VIII).

We perform three sets of signal-only simulations: (i) simulations generated from unlensed ΛCDM input spectra (hereafter "unlensed ΛCDM"), (ii) simulations generated from those same input skies, explicitly lensed in map space as described below (hereafter "lensed ΛCDM"), and (iii) simulations containing only tensor B modes with $r = 0.2$ (and $n_t = 0$).

1. Constrained input maps

The observing matrix and purification operator described in Sec. VI B are constructed for a specific assumed T sky map. Since its construction is computationally very expensive it is preferable to constrain the input T skies used for the simulations to be the same rather than to recalculate the operator for each simulation.

To construct constrained Q and U sky maps which respect the known ΛCDM TE correlation we start from a map of the well-measured temperature anisotropy, specifically the *Planck* needlet internal linear combination (NILC) T map [73]. We calculate the $a^T_{\ell m}$ using the SYNFAST software from the HEALPix [74] package [75], and then calculate sets of $a^E_{\ell m}$ using

$$a^E_{\ell m} = \frac{C^{TE}_\ell}{C^{TT}_\ell} a^T_{\ell m} + \sqrt{C^{EE}_\ell - (C^{TE}_\ell)^2/C^{TT}_\ell}\, n_{\ell m}, \quad (1)$$

where the C_ℓ's are ΛCDM spectra from CAMB [76] with cosmological parameters taken from *Planck* [9], and the $n_{\ell m}$ are normally distributed complex random numbers. For C^{TT}_ℓ we use a lensed-ΛCDM spectrum since the $a^T_{\ell m}$ from *Planck* NILC inherently contain lensing. We have found the noise level in the *Planck* NILC maps for our region of observation and multipole range to be low enough that it can be ignored.

Using the $a^E_{\ell m}$ we generate Nside = 2048 maps using SYNFAST. We substitute in the $a^T_{\ell m}$ from *Planck* 143 GHz so that the T map more closely resembles the T sky we expect to see with BICEP2. (This is also the map that is used in Sec. IV F to construct deprojection templates.) Additionally, we add in noise to the T map at the level predicted by the noise covariance in the *Planck* 143 GHz map, which allows us to simulate any deprojection residual due to noise in the *Planck* 143 GHz map.

2. Lensing of input maps

Lensing is added to the unlensed-ΛCDM maps using the LENSPIX [77] software [78]. We use this software to generate lensed versions of the constrained CMB input $a_{\ell m}$'s described in Sec. V A 1. Input to the lensing operation are deflection angle spectra that are generated with CAMB as part of the standard computation of ΛCDM spectra. The lensing operation is performed before the beam smoothing is applied to form the final map products. We do not apply lensing to the 143 GHz temperature $a^T_{\ell m}$ from *Planck* since these inherently contain lensing. Our simulations hence approximate lensed CMB maps ignoring the lensing correlation between T and E.

B. Noise pseudosimulations

The previous BICEP1 and QUAD pipelines used a Fourier based procedure to make simulated noise time streams, maintaining correlations between all channels

[46]. For the increased channel count in BICEP2 this is computationally very expensive, so we have switched to an alternate procedure adapted from SPT [79]. We perform additional coadds of the real pair maps randomly flipping the sign of each scan set. The sign-flip sequences are constructed such that the total weight of positively and negatively weighted maps is equal. We have checked this technique against the older technique, and against another technique which constructs map noise covariance matrices, and have found them all to be equivalent to the relevant level of accuracy. [In the lowest two band powers a difference can be detected within the available statistics between the sign-flip and traditional noise generators, with the sign flip predicting (10–15)% higher noise power. This is about one third of the fluctuation on the noise, and about 5% of the apparent signal. The sign-flip and matrix techniques agree to within the available statistics. Since the sign-flip sequences are 17 000 scan sets long the resulting maps are effectively uncorrelated. Separate sequences are used for each half of each temporal jack-knife.] By default we use the sign-flipping technique and refer to these realizations as "noise pseudosimulations."

We add the noise maps to the lensed-ΛCDM realizations to form signal plus noise simulations—hereafter referred to as lensed ΛCDM + noise.

VI. FROM MAPS TO POWER SPECTRA

A. Inversion to spectra

The most basic power spectrum estimation procedure which one can employ is to apply an apodization windowing, Fourier transform, construct E and B from Q and U, square, and take the means in annuli as estimates of the CMB band powers. A good choice for the window may be the inverse of the noise variance map (or a smoothed version thereof). Employing this simple procedure on the unlensed-ΛCDM simulations we find an unacceptable degree of E to B mixing. While such mixing can be corrected for in the mean using simulations, its fluctuation leads to a significant loss of sensitivity.

There are several things which can cause E to B mixing: (i) the "sky cut" implied by the apodization window (the transformation from Q and U to E and B is nonlocal so some of the modes around the edge of the map are ambiguous), (ii) the time stream (and therefore map) filtering which we have imposed in Secs. IV A and IV F, and (iii) the simple RA-Dec. map projection which we have chosen.

To correct for sky cut-induced mixing, improved estimators have been suggested. We first tried implementing the estimator suggested by Smith [80] which takes Fourier transforms of products of the map with various derivatives of the apodization window. However, testing on the unlensed-ΛCDM simulations revealed only a modest

improvement in performance since this estimator does not correct mixing caused by filtering of the map.

B. Matrix-based map purification

To overcome the E to B mixing described in the previous subsection we have introduced an additional purification step after the Q and U maps are formed. This step has to be performed in pixel space where the filtering takes place. In parallel with the construction of the pair maps and their accumulation we construct pixel-pixel matrices which track how every true sky pixel maps into the pixels of our final coadded map due to the various filtering operations. We take "true sky pixel maps" to be Nside = 512 HEALPix maps, whose pixel size (∼0.1° on a side) is smaller than our observed map pixels (0.25°). The act of simulating our various filtering operations becomes a simple matrix multiplication:

$$\tilde{\mathbf{m}} = \mathbf{R}\mathbf{m}, \tag{2}$$

where $\mathbf{m}$ is a vector consisting of $[Q, U]$ values for each HEALPix pixel and $\tilde{\mathbf{m}}$ is a $[Q, U]$ vector as observed by BICEP2 in the absence of noise.

Next, we "observe" an Nside = 512 HEALPix theoretical covariance matrix (constructed following Appendix A of Ref. [81]), $\mathbf{C}$, with $\mathbf{R}$:

$$\tilde{\mathbf{C}} = \mathbf{R}\mathbf{C}\mathbf{R}^{\mathrm{T}}. \tag{3}$$

We form $\tilde{\mathbf{C}}$ for both E-mode and B-mode covariances. These matrices provide the pixel-pixel covariance for E modes and B modes in the same observed space as the real data. However, the matrix $\mathbf{R}$ has made the two spaces nonorthogonal and introduced ambiguous modes, i.e., modes in the observed space which are superpositions of either E modes or B modes on the sky.

To isolate the pure B modes we adapt the method described in Bunn *et al.* [82]. We solve a generalized eigenvalue problem:

$$(\tilde{\mathbf{C}}_{\mathbf{B}} + \sigma^2 \mathbf{I})\mathbf{b} = \lambda_{\mathbf{b}}(\tilde{\mathbf{C}}_{\mathbf{E}} + \sigma^2 \mathbf{I})\mathbf{b}, \tag{4}$$

where $\mathbf{b}$ is a $[Q, U]$ eigenmode and σ^2 is a small number introduced to regularize the problem. By selecting modes corresponding to the largest eigenvalues $\lambda_{\mathbf{b}} \gg 1$, we can find the B-mode subspace of the observed sky which is orthogonal to E modes and ambiguous modes. The covariance matrices are calculated using steeply reddened input spectra ($\sim 1/\ell^2$) so that the eigenmodes are separated in angular scale, making it easy to select modes up to a cutoff ℓ set by the instrument resolution.

The matrix purification operator is a sum of outer products of the selected eigenmodes; it projects an input map onto this space of pure B modes:

$$\Pi_{\mathbf{b}} = \sum_i \mathbf{b}_i \mathbf{b}_i^{\mathrm{T}}. \tag{5}$$

It can be applied to any simulated map vector ($\tilde{\mathbf{m}}$) and returns a purified vector which contains only signal coming from B modes on the true sky:

$$\tilde{\mathbf{m}}_{\mathrm{pure}} = \Pi_{\mathbf{b}} \tilde{\mathbf{m}}. \tag{6}$$

This method is superior to the other methods discussed above because it removes the E-to-B leakage resulting from the filterings and the sky cut, because **R** contains all of these steps. After the purification, in the present analysis we use the simple power spectrum estimation described in the previous subsection, although in the future we may switch to a fully matrix-based approach.

Testing this operator on the standard unlensed-ΛCDM simulations (which are constructed entirely independently) we empirically determine that it is extremely effective, with residual false B modes corresponding to $r < 10^{-4}$. Testing the operator on the $r = 0.2$ simulations we find that it produces only a very modest increase in the sample variance—i.e., the fraction of mixed (ambiguous) modes is found to be small.

C. Noise subtraction and filter and beam correction

As is standard procedure in the MASTER technique [72], we noise-debias the spectra by subtracting the mean of the noise realizations (see Sec. V B). The noise in our maps is so low that this is a relevant correction only for BB, although we do it for all spectra. (The BB noise debias is 0.006 μK^2 in the $\ell \approx 75$ band power.)

To determine the response of each observed band power to each multipole on the sky we run special simulations with δ function spectra input to SYNFAST multiplied by the average measured beam function. Taking the mean over many realizations (to enable the 600 multipoles × 100 realizations per multipole required we do "short-cut" simulations using the observing matrix **R** mentioned in Sec. VI B above rather than the usual explicit time stream simulations; these are empirically found to be equivalent to high accuracy) we determine the "band power window functions" (BPWF) [83]. The integral of these functions is the factor by which each band power has been suppressed by the instrument beam and all filterings (including the matrix purification). We therefore divide by these factors and renormalize the BPWF to unit sum. This is a variant on the standard MASTER technique. (We choose to plot the band power values at the weighted mean of the corresponding BPWF instead of at the nominal band center.)

One point worth emphasizing is that when comparing the real data to our simulations (or jackknife differences thereof) the noise subtraction and filter or beam corrections have no effect since they are applied equally to the real data

and simulations. The BPWFs are required to compare the final band powers to an arbitrary external theoretical model and are provided with the data release.

The same average measured beam function is used in the signal simulations and in the BPWF calculation. In as much as this function does not reflect reality the real band powers will be under- or overcorrected at high ℓ. We estimate the beam function uncertainty to be equivalent to a 1.1% width error on a 31 arcmin FWHM Gaussian.

VII. RESULTS

A. Power spectra

Following the convention of C10 and B14 we report nine band powers, each $\approx$35 multipoles wide and spanning the range $20 < \ell < 340$. Figure 2 shows the BICEP2 power spectra [84]. With the exception of BB all spectra are consistent with their lensed-ΛCDM expectation values— the probability to exceed (PTE) the observed value of a simple χ^2 statistic is given on the plot (as evaluated against simulations—see Sec. VII C).

BB appears consistent with the lensing expectation at higher ℓ, but at lower multipoles there is an excess which is detected with high signal to noise. The χ^2 of the data is much too high to allow us to directly evaluate the PTE of the observed value under lensed ΛCDM using the simulations. We therefore "amplify" the Monte Carlo statistics by resampling band-power values from distributions fit to the simulated ones. For the full set of nine band powers shown in the figure we obtain a PTE of 1.3×10^{-7} equivalent to a significance of 5.3σ. Restricting to the first five band powers ($\ell \lesssim 200$) this changes to 5.2σ. We caution against over interpretation of the two high band powers at $\ell \approx 220$—their joint significance is $< 3\sigma$ (also see Fig. 9).

Figure 2 also shows the temporal-split jackknife—the spectrum produced when differencing maps made from the first and second halves of the data. The BB excess is not seen in the jackknife, which disfavors misestimation of the noise debias as the cause (the noise debias being equally large in jackknife spectra).

B. E and B maps

Once we have the sets of E and B Fourier modes, instead of collapsing within annuli to form power spectra, we can instead reinvert to make apodized E and B maps. In Fig. 3 we show such maps prepared using exactly the same Fourier modes as were used to construct the spectra shown in Fig. 2 filtering to the range $50 < \ell < 120$. In comparison to the simulated maps we see (i) BICEP2 has detected B modes with high signal-to-noise ratio *in the map*, and (ii) this signal appears to be evenly distributed over the field, as is the expectation for a cosmological signal, but generally will not be for a Galactic foreground.
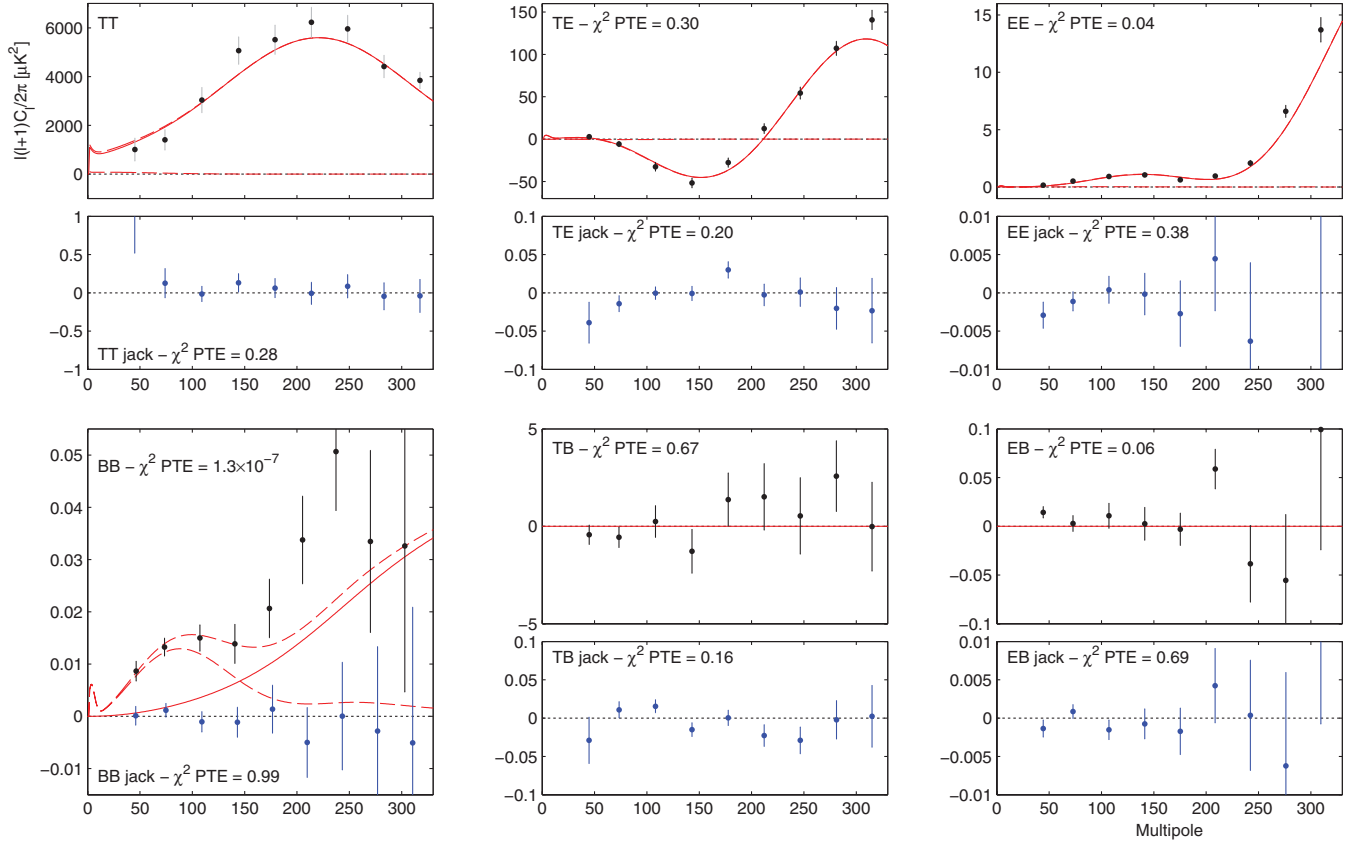
FIG. 2 (color).   BICEP2 power spectrum results for signal (black points) and temporal-split jackknife (blue points). The solid red curves show the lensed-ΛCDM theory expectations while the dashed red curves show $r = 0.2$ tensor spectra and the sum of both. The error bars are the standard deviations of the lensed-ΛCDM + noise simulations and hence contain no sample variance on tensors. The probability to exceed (PTE) the observed value of a simple χ^2 statistic is given (as evaluated against the simulations). Note the very different y-axis scales for the jackknife spectra (other than BB). See the text for additional discussion of the BB spectrum. (Note that the calibration procedure uses EB to set the overall polarization angle so TB and EB as plotted above cannot be used to measure astrophysical polarization rotation—see Sec. VIII B.)

C. Internal consistency tests

We evaluate the consistency of the jackknife spectra with their ΛCDM expectations by using a simple χ^2 statistic,

$$\chi^2 = (\mathbf{d} - \langle\mathbf{m}\rangle)^{\mathrm{T}}\mathbf{D}^{-1}(\mathbf{d} - \langle\mathbf{m}\rangle), \qquad (7)$$

where $\mathbf{d}$ is the vector of observed band-power values, $\langle\mathbf{m}\rangle$ is the mean of the lensed-ΛCDM + noise simulations (except where alternative signal models are considered), and $\mathbf{D}$ is the band-power covariance matrix as evaluated from those simulations. (Because of differences in sky coverage between the two halves of a jackknife split, in conjunction with filtering, the expectation value of some of the jackknifes is not quite zero—hence we always evaluate χ^2 versus the mean of the simulations. Because the BPWF overlap slightly adjacent band powers are $\lesssim 10\%$ correlated. We zero all but the main and first off-diagonal elements of $\mathbf{D}$ as the other elements are not measured above noise given the limited simulation statistics.) We also compute χ^2 for each of the simulations (recomputing $\mathbf{D}$ each time,

excluding that simulation) and take the probability to exceed (PTE) the observed value versus the simulated distribution. In addition to χ^2 we compute the sum of normalized deviations,

$$\chi = \sum_i \frac{d_i - \langle m_i\rangle}{\sigma_{m_i}}, \qquad (8)$$

where the d_i are the observed band-power values and $\langle m_i\rangle$ and σ_{m_i} are the mean and standard deviation of the lensed-ΛCDM + noise simulations. This statistic tests for sets of band powers which are consistently all above or below the expectation. Again we evaluate the PTE of the observed value against the distribution of the simulations.

We evaluate these statistics both for the full set of nine band powers (as in C10 and B14), and also for the lower five of these corresponding to the multipole range of greatest interest ($20 < \ell < 200$). Numerical values are given in Table I and the distributions are plotted in Fig. 4. Since we have 500 simulations the minimum
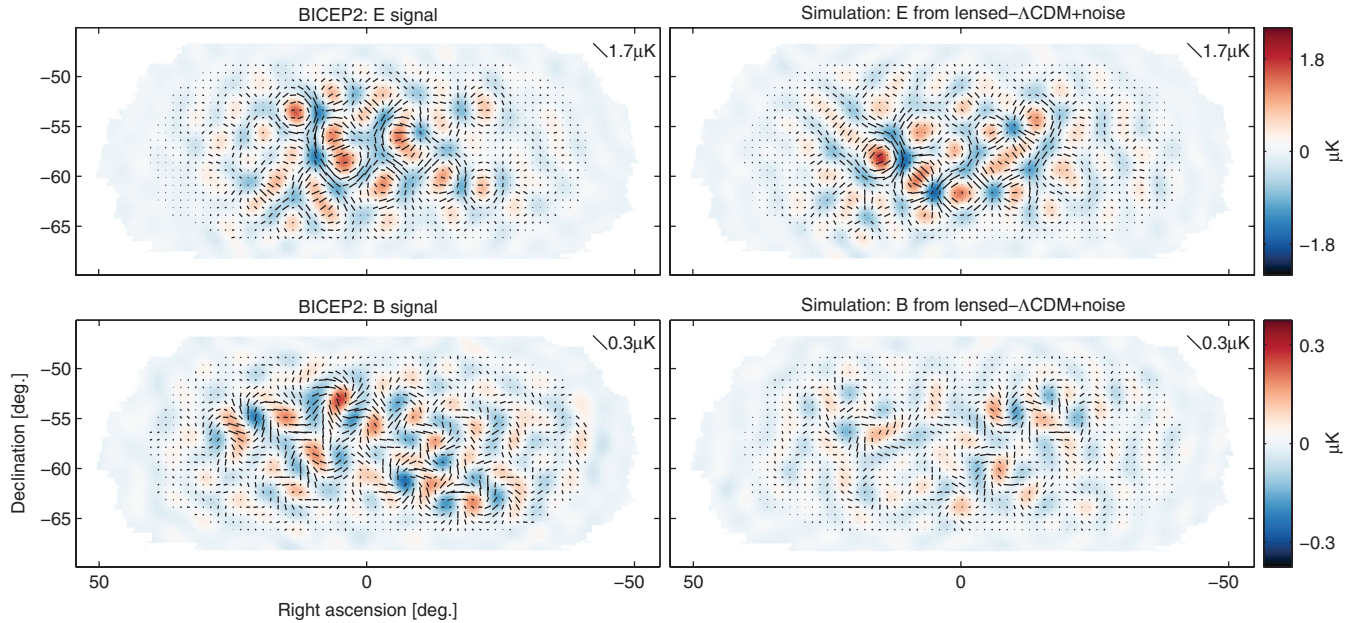
FIG. 3 (color).   Left: BICEP2 apodized *E*-mode and *B*-mode maps filtered to $50 < \ell < 120$. Right: The equivalent maps for the first of the lensed-ΛCDM + noise simulations. The color scale displays the *E*-mode scalar and *B*-mode pseudoscalar patterns while the lines display the equivalent magnitude and orientation of linear polarization. Note that excess *B* mode is detected over lensing+noise with high signal-to-noise ratio in the map ($s/n > 2$ per map mode at $\ell \approx 70$). (Also note that the *E*-mode and *B*-mode maps use different color and length scales.)

observable nonzero value is 0.002. Most of the TT, TE, and TB jackknifes pass, but following C10 and B14 we omit them from formal consideration (and they are not included in the table and figure). The signal-to-noise ratio in TT is $\sim 10^4$ so tiny differences in absolute calibration between the data subsets can cause jackknife failure, and the same is true to a lesser extent for TE and TB. Even in EE the signal to noise is approaching $\sim 10^3$ (500 in the $\ell \approx 110$ bin) and in fact most of the low values in the table are in EE. However, with a maximum signal-to-noise ratio of $\lesssim 10$ in BB such calibration differences are not a concern. All the BB (and EB) jackknifes are seen to pass, with the 112 numbers in Table I having one greater than 0.99, one less than 0.01 and a distribution consistent with uniform. Note that the four test statistics for each spectrum and jackknife are correlated this must be taken into account when assessing uniformity.

To form the jackknife spectra we difference the maps made from the two halves of the data split, divide by two, and take the power spectrum. This holds the power spectrum amplitude of a contribution which is uncorrelated in the two halves (such as noise) constant, while a fully correlated component (such as sky signal) cancels perfectly. The amplitude of a component which appears only in one half will stay the same under this operation as it is in the fully coadded map and the apparent signal-to-noise will also stay the same. For a noise-dominated experiment this means that jackknife tests can only limit potential

contamination to a level comparable to the noise uncertainty. However, the BB band powers shown in Fig. 2 have signal-to-noise as high as 10. This means that jackknife tests are extremely powerful in our case—the reductions in power which occur in the jackknife spectra are empirical proof that the *B*-mode pattern on the sky is highly correlated between all data splits considered.

We have therefore conducted an unusually large number of jackknife tests trying to imagine data splits which might conceivably contain differing contamination. Here we briefly describe each of these:

BICEP2 observed at deck angles of 68°, 113°, 248° and 293°. We can split these in two ways without losing the ability to make *Q* and *U* maps (see Sec. IV G). The *deck* jackknife is defined as 68° and 113° vs 248° and 293° while the *alt. deck* jackknife is 68° and 293° vs 113° and 248°. Uniform differential pointing averages down in a coaddition of data including an equal mix of 180° complement angles, but it will be amplified in either of these jackknifes (as we see in our simulations). The fact that we are passing these jackknifes indicates that residual beam systematics of this type are subdominant after deprojection.

The *temporal-split* simply divides the data into two equal weight parts sequentially. Similarly, but at the opposite end of the time scale range, we have the *scan direction* jackknife, which differences maps made from the right and left going half scans, and is sensitive to errors in the detector transfer function.

TABLE I. Jackknife PTE values from χ^2 and χ (sum of deviation) tests.

Jackknife	Band powers 1–5 χ^2	Band powers 1–9 χ^2	Band powers 1–5 χ	Band powers 1–9 χ
Deck jackknife				
EE	0.046	0.030	0.164	0.299
BB	0.774	0.329	0.240	0.082
EB	0.337	0.643	0.204	0.267
Scan dir jackknife				
EE	0.483	0.762	0.978	0.938
BB	0.531	0.573	0.896	0.551
EB	0.898	0.806	0.725	0.890
Temporal split jackknife				
EE	0.541	0.377	0.916	0.938
BB	0.902	0.992	0.449	0.585
EB	0.477	0.689	0.856	0.615
Tile jackknife				
EE	0.004	0.010	0.000	0.002
BB	0.794	0.752	0.565	0.331
EB	0.172	0.419	0.962	0.790
Azimuth jackknife				
EE	0.673	0.409	0.126	0.339
BB	0.591	0.739	0.842	0.944
EB	0.529	0.577	0.840	0.659
Mux col jackknife				
EE	0.812	0.587	0.196	0.204
BB	0.826	0.972	0.293	0.283
EB	0.866	0.968	0.876	0.697
Alt deck jackknife				
EE	0.004	0.004	0.070	0.236
BB	0.397	0.176	0.381	0.086
EB	0.150	0.060	0.170	0.291
Mux row jackknife				
EE	0.052	0.178	0.653	0.739
BB	0.345	0.361	0.032	0.008
EB	0.529	0.226	0.024	0.048
Tile and deck jackknife				
EE	0.048	0.088	0.144	0.132
BB	0.908	0.840	0.629	0.269
EB	0.050	0.154	0.591	0.591
Focal plane inner or outer jackknife				
EE	0.230	0.597	0.022	0.090
BB	0.216	0.531	0.046	0.092
EB	0.036	0.042	0.850	0.838
Tile top or bottom jackknife				
EE	0.289	0.347	0.459	0.599
BB	0.293	0.236	0.154	0.028
EB	0.545	0.683	0.902	0.932
Tile inner or outer jackknife				
EE	0.727	0.533	0.128	0.485
BB	0.255	0.086	0.421	0.036
EB	0.465	0.737	0.208	0.168
Moon jackknife				
EE	0.499	0.689	0.481	0.679
BB	0.144	0.287	0.898	0.858
EB	0.289	0.359	0.531	0.307
A and B offset best and worst				
EE	0.317	0.311	0.868	0.709
BB	0.114	0.064	0.307	0.094
EB	0.589	0.872	0.599	0.790
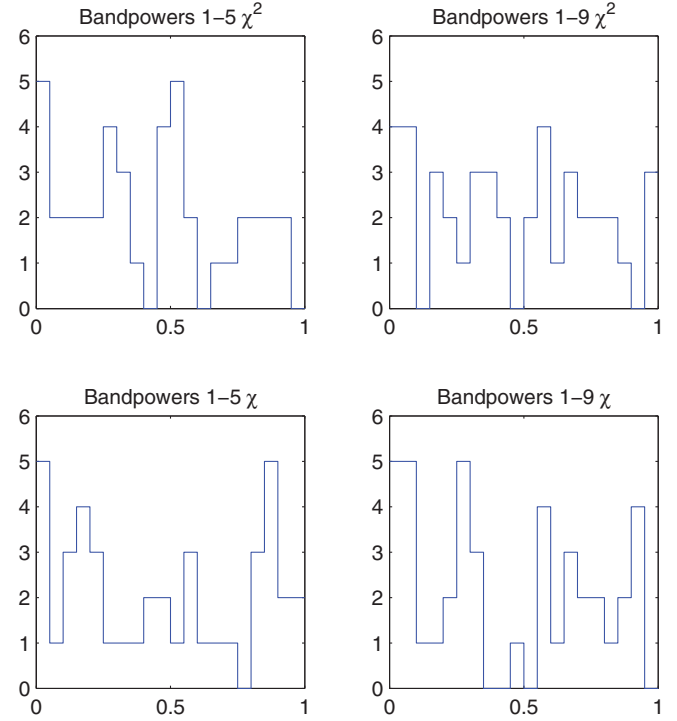
FIG. 4 (color).   Distributions of the jackknife χ^2 and χ PTE values over the 14 tests and three spectra given in Table I. These distributions are consistent with uniform.

The *azimuth* jackknife differences data taken over different ranges of telescope azimuth angle—i.e., with different potential contamination from fixed structures or emitters on the ground. A related category is the *moon* jackknife, which differences data taken when the moon is above and below the horizon.

A series of jackknifes tests if the signal originates in some subset of the detector pairs. The *tile* jackknife tests tiles 1 and 3 vs 2 and 4 (this combination being necessary to get reasonable coverage in the Q and U maps). Similarly the *tile inner or outer* and *tile top or bottom* jackknifes are straightforward. The *focal plane inner or outer* does as stated for the entire focal plane and is a potentially powerful test for imperfections which increase radially. The *mux row* and *mux column* jackknifes test for systematics originating in the readout system.

The *tile and deck* jackknife tests for a possible effect coming from always observing a given area of sky with detectors the "same way up," although due to the small range of the elevation steps it is limited to a small sky area.

Finally we have performed one test based on beam nonideality as observed in external beam map runs. The *differential pointing best or worst* jackknife differences the best and worst halves of the detector pairs as selected by that metric.

See the Systematics Paper for a full description of the jackknife studies.

VIII. SYSTEMATIC UNCERTAINTIES

Within the simulation-calibrated analysis framework described above we are free to perform any arbitrary filtering of the data which may be necessary to render the results insensitive to particular systematics. However, as such mode removal increases the uncertainty of the final band powers, we clearly wish to filter only systematics which might induce false B mode at relevant levels. Moreover, it may not be computationally feasible to construct simple time stream templates for some potential systematics. Therefore once we have made our selection as to which filterings to perform we must then estimate the residual contamination and either subtract it or show it to be negligible.

To guide our selection of mode removal we have two main considerations. First, we can examine jackknifes of the type described in Sec. VII C above—reduction in failures with increasing mode removal may imply that a real systematic effect is present. Second, and as we will see below more powerfully, we can examine external calibration data (principally beam maps) to directly calculate the false B mode expected from specific effects.

A. Simulations using observed per-channel beam shapes

As described in the Instrument Paper we have made extremely high signal-to-noise *in situ* measurements of the far-field beam shape of each channel. Fitting these beams to elliptical Gaussians we obtain differential parameters that correlate well with the mean value of the deprojection coefficients from Sec. IV F. One may then ask whether it would be better to subtract rather than deproject. In general it is more conservative to deproject as this (i) allows for the possibility that the coefficients are changing with time, and (ii) is guaranteed to completely eliminate the effect in the mean, rather than leaving a residual bias due to noise on the subtraction coefficients.

We use the per-channel beam maps as inputs to special T-only input simulations and measure the level of T to B mixing while varying the set of beam modes being deprojected. The beam maps do not provide a good estimate of differential gain so we substitute estimates which come from a per-channel variant of the absolute calibration procedure mentioned in Sec. IV G above. The left panel of Fig. 5 shows B-mode power spectra from these simulations under the following deprojections (i) none, (ii) differential pointing only, (iii) differential pointing and differential gain, (iv) differential pointing, differential gain and differential beam width, and (v) differential pointing, differential gain and differential ellipticity.

We see that differential pointing has the largest effect and so to be conservative we choose to deproject it. Differential gain is also seen to be a significant effect and we again deproject it—we lack independent subtraction coefficients, and it might plausibly be time variable. Differential beam width is a negligible effect and we do not deproject it. Differential ellipticity is also a small effect. We find in the simulations that deprojection of differential ellipticity interacts with real TE correlation in a complex manner slightly distorting the TE spectrum. We therefore choose to subtract this effect by fixing the coefficients to their beam map derived values in Sec. IV F. Whether differential ellipticity is deprojected or subtracted makes no significant difference to any of the spectra other than TE. Finally, we make a small correction for the undeprojected residual by subtracting the final curve in the left panel of Fig. 5 from the results presented in Sec. VII. (The correction is equivalent to $r = 0.001$.) We also increase the band power fluctuation to reflect the postcorrection upper limit on extended beam mismatch shown in the right panel of Fig. 5. See the Systematics Paper for details.

B. Overall polarization rotation

Once differential ellipticity has been corrected we notice that an excess of TB and EB power remains at $\ell > 200$ versus the ΛCDM expectation. The spectral form of this power is consistent with an overall rotation of the polarization angle of the experiment. While the detector-to-detector relative angles have been measured to differ from the design values by $< 0.2°$ we currently do not have an accurate external measurement of the overall polarization angle. We therefore apply a rotation of $\sim 1°$ to the final Q and U maps to minimize the TB and EB power [85,86]. We emphasize that this has a negligible effect on the BB band powers at $\ell < 200$. (The effect is 1.5×10^{-3} μK^2 at $\ell \sim 130$ and decreasing to lower ℓ.)

C. Other possible systematics

Many other systematics can be proposed as possibly leading to false B modes at a relevant level. Some possible effects will produce jackknife failure before contributing to the nonjackknife B-mode power at a relevant level. Limits on others must be set by external data or other considerations. Any azimuth fixed effect, such as magnetic pickup, is removed by the scan-synchronous template removal mentioned in Secs. III A and IV.

We have attempted an exhaustive consideration of all possible effects—a brief summary will be given here with the details deferred to the Systematics Paper. The right panel of Fig. 5 shows estimated levels of, or upper limits on, contamination from extended beam mismatch after the undeprojected residual correction, thermal drift in the focal plane, systematic polarization angle miscalibration, randomized polarization angle miscalibration, ghost beams, detector transfer function mismatch, and crosstalk. The upper limit for extended beam mismatch is the 1σ uncertainty on contamination predicted from beam map simulations identical to those described in Sec. VIII A but using
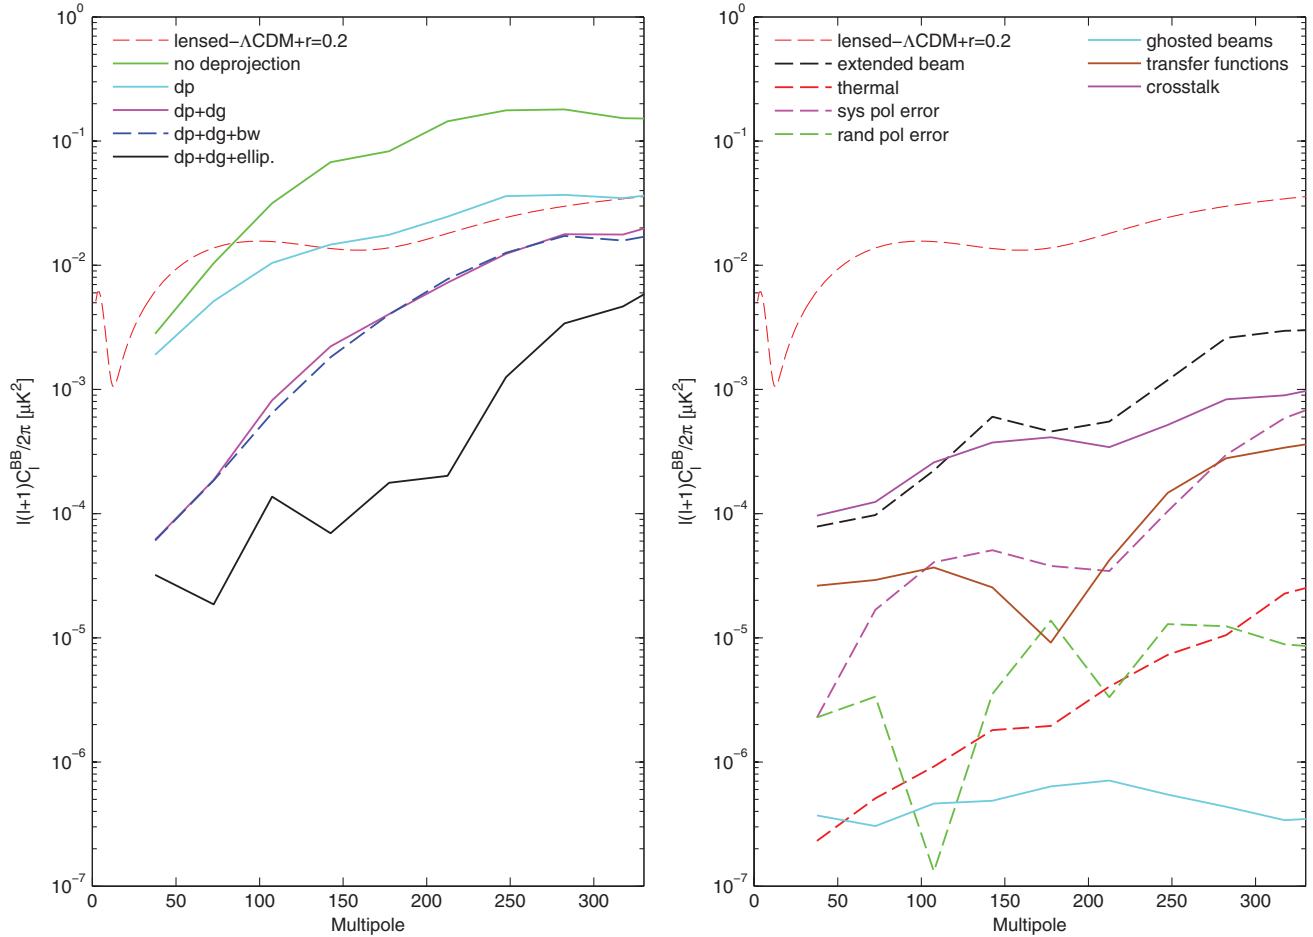
FIG. 5 (color).   Left: *BB* spectra from *T*-only input simulations using the measured per channel beam shapes compared to the lensed-ΛCDM $+ r = 0.2$ spectrum. From top to bottom the curves are (i) no deprojection, (ii) deprojection of differential pointing only (dp), (iii) deprojection of differential pointing and differential gain of the detector pairs (dp + dg), (iv) adding deprojection of differential beam width (dp + dg + bw), and (v) differential pointing, differential gain and differential ellipticity (dp + dg + ellip). Right: Estimated levels of other systematics as compared to the lensed-ΛCDM$+ r = 0.2$ spectrum. Solid lines indicate expected contamination. Dashed lines indicate upper limits. All systematics are comparable to or smaller than the extended beam mismatch upper limit.

a larger region of the beam. (Note that this will include beam or beamlike effects which are present in the beam mapping runs, including crosstalk and side lobes at $\lesssim 4°$.) For systematic polarization angle miscalibration it is the level at which such an error would produce a detectable TB signal with 95% confidence. For randomized polarization angle miscalibration, it is the leakage we would incur from assuming nominal polarization angles, i.e., no ability to measure per-pair relative polarization angles. For thermal drift, it is the noise floor set by the sensitivity of the thermistors that monitor focal plane temperature.

IX. FOREGROUND PROJECTIONS

Having provided evidence that the detected *B*-mode signal is not an instrumental artifact, we now consider whether it might be due to a Galactic or extragalactic

foreground. At low or high frequencies Galactic synchrotron and polarized-dust emission, respectively, are the dominant foregrounds. The intensity of both falls rapidly with increasing Galactic latitude but dust emission falls faster. The equal amplitude crossover frequency therefore rises to $\gtrsim 100$ GHz in the cleanest regions [[87], Fig. 10]. The BICEP2 field is centered on Galactic coordinates $(l, b) = (316°, -59°)$ and was originally selected on the basis of exceptionally low contrast in the FDS dust maps [88]. In these unpolarized maps such ultraclean regions are very special—at least an order of magnitude cleaner than the average $b > 50°$ level.

Foreground modeling involves extrapolating maps taken at lower or higher frequencies to the CMB observation band, and there are inevitably uncertainties. Many previous studies have been conducted and projections made—see, for instance, Dunkley *et al.* [87], and references therein.

Such previous studies have generically predicted levels of foreground *B*-mode contamination in clean high latitude regions equivalent to $r \lesssim 0.01$—well below that which we observe—although they note considerable uncertainties.

A. Polarized dust projections

The main uncertainty in foreground modeling is currently the lack of a polarized dust map. (This will be alleviated soon by the next *Planck* data release.) In the meantime we have therefore investigated a number of existing models using typical or default assumptions for polarized dust, and have formulated a new one. A brief description of each model is as follows:

FDS: Model 8 [88], scaled with a uniform polarization fraction of 5%, is a commonly used all-sky baseline model (e.g.,[44,87]). We set $Q = U$.

BSS: Bisymmetric spiral (BSS) model of the Galactic magnetic field [89,90]. The polarization fraction varies across the sky in this model; by default it is scaled to match the 3.6% all-sky average reported by WMAP [91], giving a mean and standard deviation in the BICEP2 field of $(5.7 \pm 0.7)\%$.

LSA: Logarithmic spiral arm (LSA) model of the Galactic magnetic field [89,90]. The polarization fraction varies across the sky in this model; by default it is also scaled to match the 3.6% all-sky average reported by WMAP [91], giving a mean and standard deviation in the BICEP2 field of $(5.0 \pm 0.3)\%$.

PSM: Planck sky model (PSM) [92] version 1.7.8, run as a "Prediction" with default settings, including 15% dust intrinsic polarization fraction [93]. In this model, the intrinsic polarization fraction is reduced by averaging over variations along each line of sight. The resulting polarization fraction varies across the sky; its mean and standard deviation in the BICEP2 field are $(5.6 \pm 0.8)\%$.

DDM1: "Data driven model 1" (DDM1) constructed from publicly available *Planck* data products. The *Planck* dust model map at 353 GHz is scaled to 150 GHz assuming a constant emissivity value of 1.6 and a constant temperature of 19.6 K [94]. A nominal uniform 5% sky polarization fraction is assumed, and the polarization angles are taken from the PSM. This model will be biased down due to the lack of spatial fluctuation in the polarization fraction and angles, but biased up due to the presence of instrument noise and (unpolarized) cosmic infrared background anisotropy in the *Planck* dust model [95].

All of the models except FDS make explicit predictions of the actual polarized dust pattern in our field. We can therefore search for a correlation between the models and our signal by taking cross spectra against the BICEP2 maps. The upper panel of Fig. 6 shows the resulting *BB* auto and cross spectra—the autospectra are all below the level of our observed signal and no significant cross-correlation is found. [The cross spectra between each model and real data are consistent with the cross spectra

between that model and (uncorrelated) lensed-LCDM + noise simulations.] We note that the lack of cross-correlation can be interpreted as due to limitation of the models. To produce a power level from DDM1 auto comparable to the observed excess signal would require one to assume a uniform polarization fraction of ∼13%. While this is well above typically assumed values, models
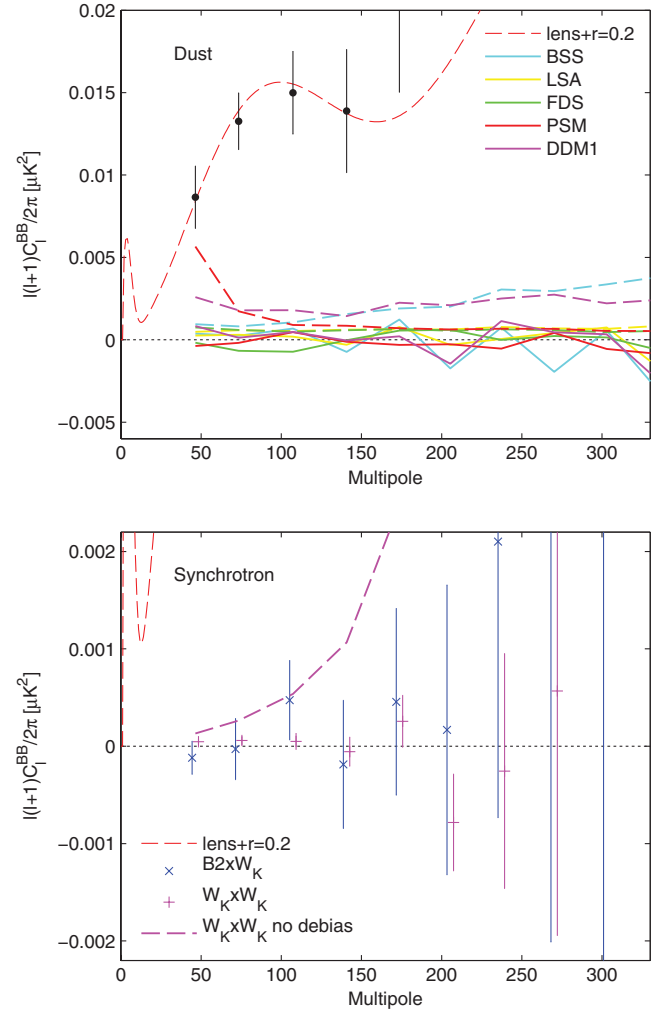

FIG. 6 (color). Upper: Polarized dust foreground projections for our field using various models available in the literature, and a new one formulated using the information officially available from *Planck*. Dashed lines show autospectra of the models, while solid lines show cross spectra between the models and the BICEP2 maps. The BICEP2 auto spectrum from Fig. 2 is also shown with the lensed-ΛCDM + $r = 0.2$ spectrum. Lower: Polarized synchrotron constraints for our field using the WMAP K band (23 GHz) maps projected to 150 GHz using the mean spectral index within our field ($\beta = -3.3$) from WMAP. The blue points with error bars show the cross spectrum between the BICEP2 and WMAP maps, with the uncertainty estimated from cross spectra against simulations of the WMAP noise. The magenta points with error bars and the dashed curve show the WMAP auto spectrum with and without noise debias. See the text for further details.

are not yet well-enough constrained by external public data
to exclude the possibility of emission at this level.

B. Synchrotron

To constrain the level of Galactic synchrotron in our field
we take the WMAP K-band (23 GHz) map, extrapolate it to
150 GHz, reobserve with our simulation pipeline, and take
the cross spectrum against the BICEP2 maps, with appro-
priate BICEP2 filtering and WMAP beam correction. In
our field and at angular scales of $\ell > 30$ the WMAP
K-band maps are noise dominated. We therefore also make
noise realizations and take cross spectra with these to assess
the uncertainty. The lower panel of Fig. 6 shows the
resulting cross spectrum and its uncertainty. Using the
MCMC Model f spectral index map provided by WMAP
[2] we obtain a mean value within our field of
$\beta = -3.3 \pm 0.16$. For this value, the resulting cross spec-
trum implies a contribution to our r constraint (calculated
as in Sec. XI) equivalent to $r_{\text{sync},150} = 0.0008 \pm 0.0041$,
while for a more conservative $\beta = -3.0$, $r_{\text{sync},150} =
0.0014 \pm 0.0071$. In contrast to analysis with the models
of polarized dust, cross spectra with the official WMAP
polarized maps can be confidently expected to provide an
unbiased estimate of signal correlated with synchrotron for
a given spectral index, with a quantified uncertainty. Note
that the assumed spectral index only enters as the first
power in these $\text{BICEP2} \times \text{WMAP}_K$ cross spectral con-
straints, and the uncertainty depends only weakly on the
model for WMAP noise. The WMAP_K auto spectrum, if
de-biased for noise, implies even tighter constraints on the
synchrotron contribution to our r parameter: for $\beta = -3.3$,
$r_{\text{sync},150} = 0.001 \pm 0.0006$, or for $\beta = -3.0$, $r_{\text{sync},150} =
0.003 \pm 0.002$, although these have a somewhat greater
dependence on assumptions about WMAP noise levels and
the spectral index.

C. Point sources

Extragalactic point sources might also potentially be a
concern. Using the 143 GHz fluxes for the sources in our
field from the *Planck* catalog [97], together with polariza-
tion information from ATCA [98] we find that the con-
tribution to the BB spectrum is equivalent to $r \approx 0.001$.
This is consistent with the projections of Battye *et al.* [99].

X. CROSS SPECTRA

A. Cross spectra with BICEP1

BICEP1 observed essentially the same field as BICEP2
from 2006 to 2008. While a very similar instrument in
many ways the focal plane technology of BICEP1 was
entirely different, employing horn-fed PSBs read out via
neutron transmutation-doped (NTD) germanium thermis-
tors (see T10 for details). The high-impedance NTD
devices and readouts have different susceptibility to micro-
phonic pickup and magnetic fields, and the shielding of

unwanted RFI and EMI was significantly different from
that of BICEP2. The beam systematics were also quite
different with a more conservative edge taper and smaller
observed pair centroid offsets (see T10 and the Instrument
Paper). BICEP1 had detectors at both 100 and 150 GHz.

Figure 7 compares the BICEP2 EE and BB auto spectra
with cross spectra taken against the 100 and 150 GHz maps
from BICEP1. For EE the correlation is extremely strong,
which simply confirms that the mechanics of the process
are working as expected. For BB the signal-to-noise is of
course much lower, but there appear to be positive
correlations. To test the compatibility of the BB auto
and cross spectra we take the differences and compare
to the differences of lensed-$\Lambda\text{CDM} + \text{noise} + r = 0.2$
simulations (which share common input skies). (For all
spectral difference tests we compare against lensed-
$\Lambda\text{CDM} + \text{noise} + r = 0.2$ simulations as the cross terms
between signal and noise increase the variance even for
perfectly common sky coverage.) Using band powers 1–5
the χ^2 and χ PTEs are midrange, indicating that the spectra
are compatible to within the noise. (This is also true
for EE.)

To test for evidence of excess power over the base
lensed-ΛCDM expectation we calculate the BB χ^2 and χ
statistics against this model. The $\text{BICEP2} \times \text{BICEP1}_{150}$
spectrum has PTEs of 0.37 and 0.05, respectively, while the
$\text{BICEP2} \times \text{BICEP1}_{100}$ spectrum has PTEs of 0.005 and
0.001. The latter corresponds to a $\approx 3\sigma$ detection of excess
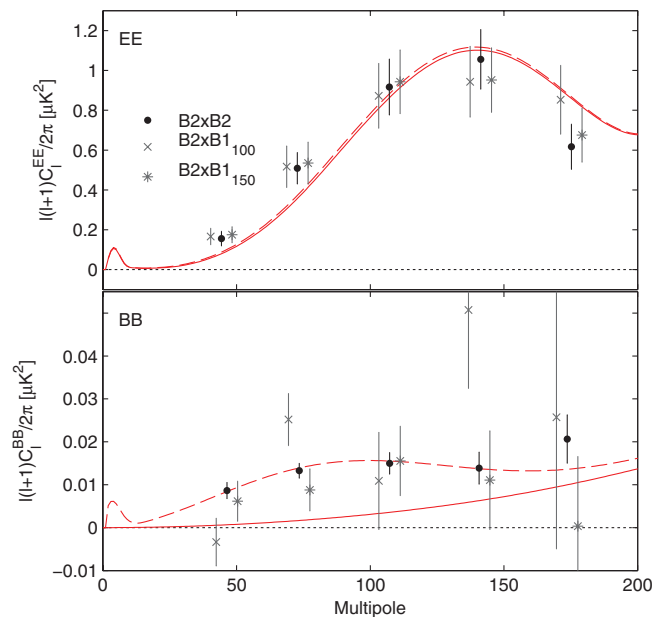

FIG. 7 (color).   The BICEP2 EE and BB auto spectra (as shown
in Fig. 2) compared to cross spectra between BICEP2 and the 100
and 150 GHz maps from BICEP1. The error bars are the standard
deviations of the lensed-$\Lambda\text{CDM} + \text{noise}$ simulations and hence
contain no sample variance on tensors. (For clarity the cross
spectrum points are offset horizontally.)

power. While it may seem surprising that one cross spectrum shows a stronger detection than the other, it turns out not to be unusual in lensed-ΛCDM + noise + $r = 0.2$ simulations. (Compared to such lensed-ΛCDM + noise + $r = 0.2$ simulations, χ^2 and χ PTEs are 0.92 and 0.74 for BICEP2 × BICEP1$_{150}$ and 0.18 and 0.23 for BICEP2 × BICEP1$_{100}$. These simulations also indicate that the BICEP2 × BICEP1$_{150}$ and BICEP2 × BICEP1$_{100}$ values are only weakly correlated. Therefore if $r = 0.2$ is the true underlying model then the observed BICEP2 × BICEP1$_{150}$ χ^2 and χ values appear to be modest downward fluctuations and the BICEP2 × BICEP1$_{100}$ values modest upward fluctuations—but they are compatible.)

B. Spectral index constraint

We can use the BICEP2 auto and BICEP2 × BICEP1$_{100}$ spectra shown in Fig. 7 to constrain the frequency dependence of the observed signal. If the signal at 150 GHz were due to synchrotron we would expect the frequency cross spectrum to be much larger in amplitude than the BICEP2 auto spectrum. Conversely, if the 150 GHz power were due to polarized dust emission we would not expect to see a significant correlation with the 100 GHz sky pattern.

Pursuing this formally, we use simulations of both experiments observing a common sky to construct a combined likelihood function for band powers 1–5 of the BICEP2 auto, BICEP1$_{100}$ auto, and their cross spectrum using the Hamimeche-Lewis [100] approximation (HL); see B14 for implementation details. As with all likelihood analyses we report, this procedure fully accounts for sample variance. We use this likelihood function to fit a six-parameter model parametrized by five 150 GHz band power amplitudes and a single common spectral index, β. We consider two cases, in which the model accounts for (1) the total BB signal or (2) only the excess over lensed ΛCDM, and we take the spectral index to be the power law exponent of this signal's antenna temperature as a function of frequency. We marginalize this six-parameter model over the band powers to obtain a one-parameter likelihood function over the spectral index.

Figure 8 shows the resulting estimates of the spectral index, with approximate 1σ uncertainty ranges. We evaluate the consistency with specific values of β using a likelihood ratio test. Both the total and the excess observed BB signal are consistent with the spectrum of the CMB ($\beta = -0.7$ for these bands and conventions). The spectrum of the excess BB signal has a CMB-to-peak likelihood ratio of $L = 0.75$. Following Wilks [101] we take $\chi^2 \approx -2\log L$ and evaluate the probability to exceed this value of χ^2 (for a single degree of freedom). A synchrotron spectrum with $\beta = -3.0$ is disfavored for the excess BB ($L = 0.26$, PTE 0.10, 1.6σ); although the BICEP2 × WMAP$_K$ spectrum offers a much stronger constraint. The preferred whole-sky dust spectrum from *Planck* [94], which corresponds under these conventions to $\beta \approx +1.5$, is also disfavored as an
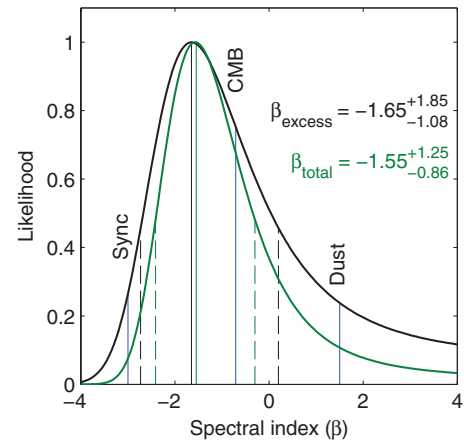

FIG. 8 (color). The constraint on the spectral index of the BB total signal (green) and excess signal over lensed ΛCDM (black), based on joint consideration of the BICEP2 auto, BICEP1$_{100}$ auto, and BICEP2 × BICEP1$_{100}$ cross spectra. The curve shows the marginalized likelihood as a function of assumed spectral index. The vertical solid and dashed lines indicate the maximum likelihoods and the $\pm 1\sigma$ intervals. The blue vertical lines indicate the equivalent spectral indices under these conventions for the CMB, synchrotron, and dust. The observed signal is consistent with a CMB spectrum, while synchrotron and dust are both disfavored.

explanation for the excess BB ($L = 0.24$, PTE 0.09, 1.7σ). We have also conducted a series of simulations applying this procedure to simulated data sets with CMB and dust spectral indices. These simulations indicate that the observed likelihood ratios are typical of a CMB spectral index but atypical of dust. [For the dust simulations we simulate power spectra for our sky patch using the HL likelihood function, assuming the observed BICEP2 power spectrum at 150 GHz and extrapolating to 100 GHz using a spectral index of $+1.5$ for the excess above lensing. For each simulation we compute this likelihood function and calculate the likelihood ratio of $L(1.5)/L(\text{CMB})$. In 45 of 500 such simulations we find a likelihood ratio smaller than that in our actual data.]

In the analysis above, the 100 GHz auto spectrum contributes little statistical weight, so what is being constrained is effectively the spectral index of the component of the 100 GHz sky pattern which correlates with the 150 GHz pattern. A mixture of synchrotron and dust, summing to the level of the observed BB excess, could in principle be constructed to achieve any intermediate effective spectral index. Spatial correlation between the two patterns is an additional potential degree of freedom. Considering a scenario with no such correlation and nominal dust and synchrotron spectral indices ($\beta_{\text{sync}} = -3.0$), reproducing the maximum likelihood effective $\beta = -1.65$ (see Fig. 8) would require a nearly equal mix of dust and synchrotron BB power at 150 GHz. In this scenario, the synchrotron contribution in the BICEP2

auto spectrum would be $r_{\mathrm{sync},150} \approx 0.10$. However, the corresponding constraint from the BICEP2 $\times$ WMAP$_K$ cross spectrum (Sec. IX B, $r_{\mathrm{sync},150} = 0.0014 \pm 0.0071$) rules this scenario out at 13.5σ. Calculating the BICEP1$_{100}$ $\times$ WMAP$_K$ cross spectrum yields a similar but slightly weaker constraint: for $\beta_{\mathrm{sync}} = -3.0$, $r_{\mathrm{sync},150} = -0.0005 \pm 0.0076$, disfavoring this scenario at 12.6σ.

In a scenario with 100% correlation between synchrotron and dust, an effective index $\beta = -1.65$ can be produced with a lower synchrotron contribution, but the assumption of dust correlation adds to this scenario's predicted level for BICEP2 $\times$ WMAP$_K$, so that the actual measured cross spectrum also disfavors this scenario at $> 13\sigma$. More generally, scenarios which mix dust with synchrotron ($\beta_{\mathrm{sync}} = -3.0$) with any assumed degree of correlation from (0–100)%, in ratios needed to produce an effective $\beta < 0.2$, are disfavored by the BICEP2 $\times$ WMAP$_K$ cross spectrum constraint at $> 3\sigma$. Scenarios which would approximate a CMB-like index ($\beta = -0.7$) with a mixture of dust and synchrotron are therefore unlikely.

C. Additional cross spectra

Having seen that the BICEP2 auto spectrum is compatible with both the BICEP2 $\times$ BICEP1$_{100}$ and the BICEP2 $\times$ BICEP1$_{150}$ cross spectra we proceed to combine the latter. (We combine using weights which minimize the variance of the lensed-ΛCDM + noise simulations as described in B14.) Figure 9 compares the result to the BICEP2 auto spectrum from Fig. 2. Taking the difference of these spectra and comparing to the differences of the lensed-ΛCDM + noise + $r = 0.2$ simulations the band power 1–5 χ^2 and χ PTEs are midrange indicating compatibility.

Comparing the BICEP2 $\times$ BICEP1$_{\mathrm{comb}}$ spectrum to the lensed-ΛCDM expectation the χ^2 and χ values have PTE of 0.005 and 0.002, respectively, corresponding to $\approx 3\sigma$ evidence of excess power. The compatibility of the BICEP2 auto and BICEP2 $\times$ BICEP1$_{\mathrm{comb}}$ cross spectra combined with the detection of excess power in the cross spectra provides yet more evidence against a systematic origin of the nominal signal given the significant differences in focal plane technology and beam imperfections.

The successor experiment to BICEP2 is the *Keck Array* which consists of five BICEP2-like receivers [102]. The *Keck Array* data analysis is not yet complete and will be the subject of future publications. However, as an additional systematics check we show in Fig. 9 a cross spectrum between BICEP2 and preliminary *Keck Array* 150 GHz maps from the 2012 and 2013 seasons. This cross spectrum also shows obvious excess *BB* power at low ℓ.

XI. COSMOLOGICAL PARAMETER CONSTRAINTS

We have shown that our observed *B*-mode spectrum (i) is not explained by known systematics (jackknifes, beam-map
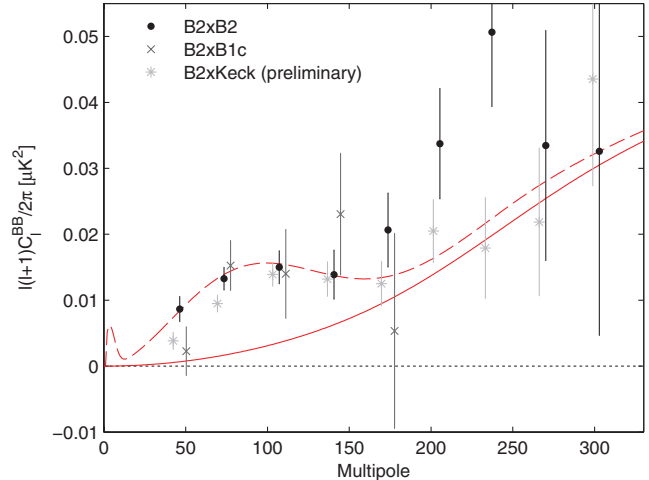
FIG. 9 (color). Comparison of the BICEP2 *BB* auto spectrum and cross spectra taken between BICEP2 and BICEP1 combined, and BICEP2 and *Keck Array* preliminary. The error bars are the standard deviations of the lensed-ΛCDM + noise simulations and hence contain no sample variance on tensors. (For clarity the cross spectrum points are offset horizontally and the BICEP2 $\times$ BICEP1 points are omitted at $\ell > 200$.)

simulations, other systematics studies, and cross spectra with BICEP1$_{150}$), and (ii) domination by foregrounds is disfavored (dust model projections, dust model cross correlations, synchrotron constraints, and spectral index constraints from cross spectra with BICEP1$_{100}$). In this section we do some basic fitting of cosmological parameters while noting again that all the band powers and ancillary data are available for download so that others may conduct fuller studies.

A. Lensed-ΛCDM + tensors

In Fig. 2 we see a substantial excess of *BB* power in the region where an inflationary gravitational wave (IGW) signal would be expected to peak. We therefore proceed to find the most likely value of the tensor-to-scalar ratio r using the "direct likelihood" method introduced in B14. We first form additional sets of simulations for many values of r by combining the lensed-ΛCDM and scaled $r = 0.2$ simulations. (Hence we assume always $n_t = 0$ making the value of r independent of the tensor pivot scale.) We then combine the band powers of these and the real band powers with s/n weighting where s is the IGW spectrum for a small value of r and n is the variance of the lensed-ΛCDM + noise simulations. Arranging the simulation pdf values as rows we can then read off the likelihood curve for r as the columns at the observed combined bandpower value.

The result of this process is shown in Fig. 10. Defining the confidence interval as the equal likelihood contour which contains 68% of the total likelihood we find $r = 0.20^{+0.07}_{-0.05}$. This uncertainty is driven by the sample
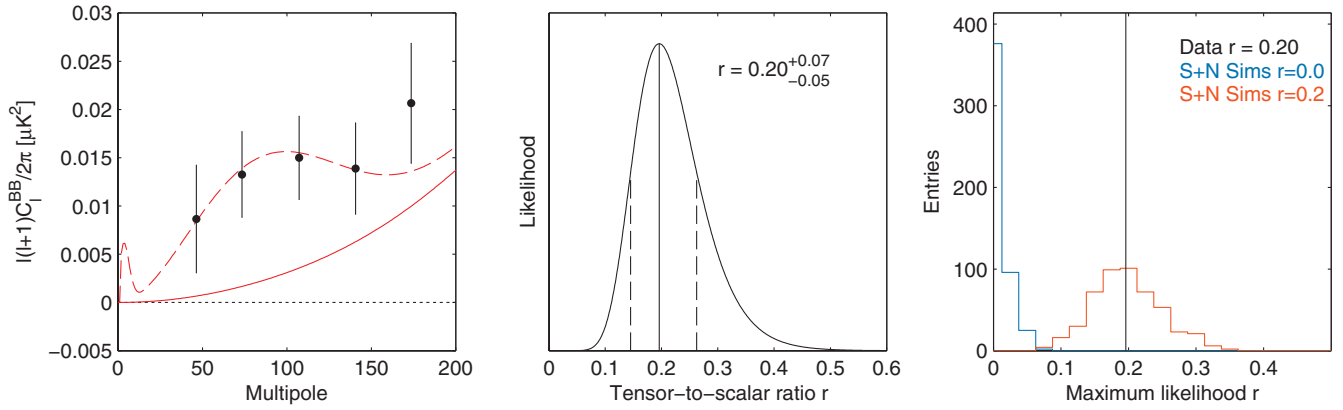
FIG. 10 (color). Left: The BICEP2 band powers plotted with the maximum likelihood lensed-ΛCDM $+ r = 0.20$ model. The uncertainties are taken from that model and hence include sample variance on the r contribution. Middle: The constraint on the tensor-to-scalar ratio r. The maximum likelihood and $\pm 1\sigma$ interval is $r = 0.20^{+0.07}_{-0.05}$, as indicated by the vertical lines. Right: Histograms of the maximum likelihood values of r derived from lensed-ΛCDM $+$ noise simulations with $r = 0$ (blue) and adding $r = 0.2$ (red). The maximum likelihood value of r for the real data is shown by the vertical line.

variance in our patch of sky, and the likelihood falls off very steeply towards $r = 0$. The likelihood ratio between $r = 0$ and the maximum is 2.9×10^{-11} equivalent to a PTE of 3.3×10^{-12} or 7.0σ. The numbers quoted above are for bins 1–5 although due to the weighting step they are highly insensitive to the inclusion of the higher band powers. (Absolute calibration and beam uncertainty are included in these calculations but have a negligible effect.)

Evaluating our simple χ^2 statistic between band powers 1–5 and the lensed-ΛCDM $+$ noise $+ r = 0.2$ simulations yields a value of 1.1, which for 4 degrees of freedom has a PTE of 0.90. Using all nine band powers χ^2 is 8.4, which for 8 degrees of freedom has a PTE of 0.40. The model is therefore a perfectly acceptable fit to the data.

In Fig. 11 we recompute the r constraint subtracting each of the dust models shown in Fig. 6. For the auto spectra the range of maximum likelihood r values is 0.15–0.19, while for the cross it is 0.19–0.21 (random fluctuations in the cross can cause shifts up as well as down). The probability that each of these models reflects reality is hard to assess. To explain the entire excess BB signal with dust requires increasing the power predicted by the auto spectra of the various models by factors ranging from $\sim (5$–$10)\times$. For example, in the context of the DDM1 model the preferred value of r varies as $r \sim 0.20$–$13p^2$, so that increasing this model's assumed uniform polarization fraction from $p = 5\%$ to $p \sim 13\%$ would explain the full excess under this model.

The dust foreground is expected to have a power law spectrum which slopes modestly down $\propto \ell^{\sim -0.6}$ in the usual $\ell(\ell + 1)C_\ell/2\pi$ units [87]—although how this might fluctuate from small field to small field at high Galactic latitude has not been investigated. We note that the s/n band-power weighting scheme described above weights the

first bin highly, and it is here that the foreground models equal the largest fraction of the observed signal. Therefore if we were to exclude the first band power the difference between the unsubtracted and foreground subtracted model lines in Fig. 11 would be smaller; i.e.. while dust may contribute significantly to our first band power it seems less able to explain band powers two through five. Reevaluating the base r constraint using band powers 2–5 yields $r = 0.19^{+0.07}_{-0.05}$ with $r = 0$ ruled out at 6.4σ.

Computing an r constraint using the BICEP2 $\times$ BICEP1$_{\text{comb}}$ cross spectrum shown in Fig. 9 yields $r = 0.19^{+0.11}_{-0.08}$. The likelihood ratio between $r = 0$ and
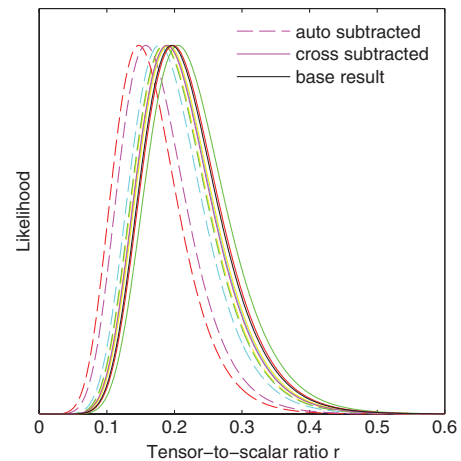

FIG. 11 (color). Modified constraints on the tensor-to-scalar ratio r when subtracting each of the foreground models shown in Fig. 6 from the BICEP2 BB band powers. The line styles and colors match Fig. 6 with dashed for auto spectra and solid for cross spectra. The probability that each of these models reflects reality is hard to assess—see the text for discussion.
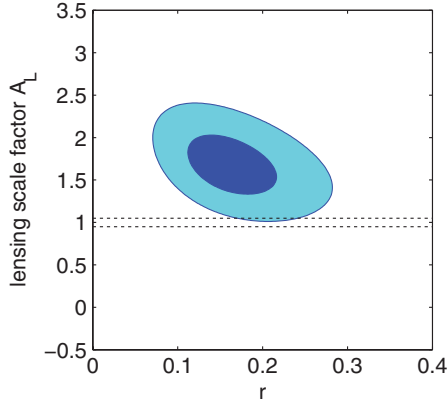
FIG. 12 (color). Joint constraints on the tensor-to-scalar ratio r and the lensing scale factor A_L using the BICEP2 BB band powers 1–9. One and two σ contours are shown. The horizontal dotted lines show the 1σ constraint from *Planck* [9]. The B-mode lensing signal is detected at 5.5σ, with an amplitude $\sim 2\sigma$ higher than the expected value.

the maximum is 2.0×10^{-3} equivalent to a PTE of 4.2×10^{-4} or 3.5σ.

B. Scaled-lensing + tensors

Lensing deflections of the CMB photons as they travel from last scattering remap the patterns slightly. In temperature this leads to a slight smoothing of the acoustic peaks, while in polarization a small B mode is introduced with a spectrum similar to a smoothed version of the EE spectrum a factor ~ 100 lower in power. Using their own and other data *Planck* [9] quote a limit on the amplitude of the lensing effect versus the ΛCDM expectation of $A_L = 0.99 \pm 0.05$.

Figure 12 shows a joint constraint on the tensor-to-scalar ratio r and the lensing scale factor A_L using our BB band powers 1–9. As expected there is a weak anticorrelation—one can partially explain the low ℓ excess by scaling up the lensing signal. However, the constraint is mostly driven by band powers six through nine where the IGW signal is small. The maximum likelihood scaling is ≈ 1.75, $\sim 2\sigma$ from unity. Marginalizing over r the likelihood ratio between peak and zero is 3×10^{-7}, equivalent to a PTE of 4.7×10^{-8} or a 5.5σ detection of lensing in the BICEP2 BB auto spectrum. We note again that the high values of band powers six and seven are not present in the preliminary cross spectra against *Keck Array* shown in Fig. 9.

C. Compatibility with temperature data

If present at last scattering, tensor modes will add power to all spectra including TT. For an r value of 0.2 the contribution to TT at the largest angular scales ($\ell < 10$) would be $\approx 10\%$ of the level measured by WMAP and *Planck*. The theoretical ΛCDM power level expected at these scales is dependent on several cosmological parameters including the spectral index of the initial scalar perturbations

n_s and the optical depth to the last scattering surface τ. However, by combining temperature data taken over a wide range of angular scales indirect limits on r have been set. A combination of WMAP + SPT data [4] yields $r < 0.18$ (95% confidence) tightening to $r < 0.11$ when also including measurements of the Hubble constant and baryon acoustic oscillations (BAO). More recently *Planck* [9] quote $r < 0.11$ using a combination of *Planck*, SPT and ACT temperature data, plus WMAP polarization (to constrain τ).

These limits appear to be in moderately strong tension with interpretation of our B-mode measurements as primarily due to tensors. One possibility is a larger than anticipated contribution from polarized dust, but as our present data disfavor this one can ask what additional extensions to the standard model might resolve the situation.

One obvious modification is to allow the initial scalar perturbation spectrum to depart from the simple power law form which is assumed in the base ΛCDM model. A standard way in which this is done is by introducing a "running" parameter $dn_s/d\ln k$. In *Planck* XVI [9] the constraint relaxes to $r < 0.26$ (95% confidence) when running is allowed with $dn_s/d\ln k = -0.022 \pm 0.010$ (68%) (for the *Planck* + WP + highL data combination). In Fig. 13 we show the constraint contours when allowing running as taken from Fig. 23 of [9], and how these change when the BICEP2 data are added. The red contours on the plot are simply the Monte Carlo Markov chains (MCMC)
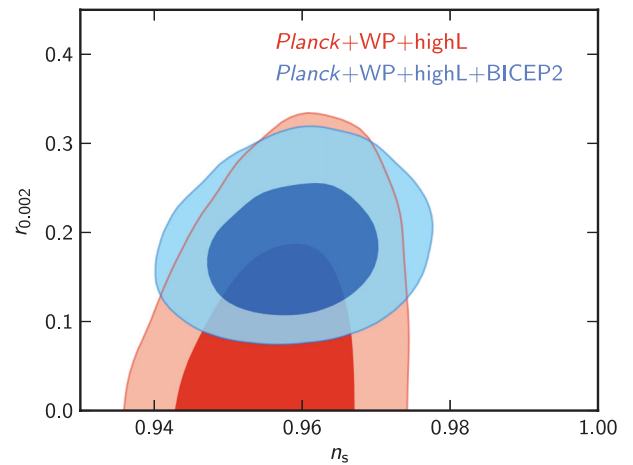

FIG. 13 (color). Indirect constraints on r from CMB temperature spectrum measurements relax in the context of various model extensions. Shown here is one example, following *Planck* XVI [9] Fig. 23, where tensors and running of the scalar spectral index are added to the base ΛCDM model. The contours show the resulting 68% and 95% confidence regions for r and the scalar spectral index n_s when also allowing running. The red contours are for the "Planck + WP + highL" data combination, which for this model extension gives a 95% bound $r < 0.26$ [9]. The blue contours add the BICEP2 constraint on r shown in the center panel of Fig. 10. See the text for further details.

[103,104] provided with the *Planck* data release [105] (and are thus identical to those shown in that *Planck* paper). We then apply importance sampling [106] to these chains using our r likelihood as shown in Fig. 10 to derive the blue contours, for which the running parameter constraint shifts to $dn_s/d\ln k = -0.028 \pm 0.009$ (68%).

The point of Fig. 13 is not to endorse running as the correct explanation of the observed deficit of low ℓ TT power. It is simply to illustrate one example of a simple model extension beyond standard ΛCDM + tensors which can resolve the apparent tension between previous TT measurements and the direct evidence for tensors provided by our B-mode measurements—probably there are others. Of course, one might also speculate that the tension could be reduced within the standard ΛCDM + tensors model, for example if τ or other parameters were allowed to shift. We anticipate a broad range of possibilities will be explored.

XII. CONCLUSIONS

We have described the observations, data reduction, simulation, and power spectrum analysis of all three seasons of data taken by the BICEP2 experiment. The polarization maps presented here are the deepest ever made at degree angular scales having noise level of 87 nK deg in Q and U over an effective area of 380 square deg.

To fully exploit this unprecedented sensitivity we have expanded our analysis pipeline in several ways. We have added an additional filtering of the time stream using a template temperature map (from *Planck*) to render the results insensitive to temperature to polarization leakage caused by leading order beam systematics. In addition we have implemented a map purification step that eliminates ambiguous modes prior to B-mode estimation. These deprojection and purification steps are both straightforward extensions of the kinds of linear filtering operations that are now common in CMB data analysis.

The power spectrum results are consistent with lensed ΛCDM with one striking exception: the detection of a large excess in the BB spectrum in the ℓ range where an inflationary gravitational wave signal is expected to peak. This excess represents a 5.2σ excursion from the base lensed-ΛCDM model. We have conducted a wide selection of jackknife tests which indicate that the B-mode signal is common on the sky in all data subsets. These tests offer strong empirical evidence against a systematic origin for the signal.

In addition, we have conducted extensive simulations using high fidelity per channel beam maps. These confirm our understanding of the beam effects, and that after deprojection of the two leading order modes, the residual is far below the level of the signal which we observe.

Having demonstrated that the signal is real and "on the sky" we proceeded to investigate if it may be due to foreground contamination. Polarized synchrotron emission

from our galaxy is estimated to be negligible using low frequency polarized maps from WMAP. For polarized dust emission public maps are not yet available. We therefore investigate a number of commonly used models and one which uses information which is currently officially available from *Planck*. At default parameter values these models predict auto spectrum power well below our observed level. However, these models are not yet well constrained by external public data, which cannot empirically exclude dust emission bright enough to explain the entire excess signal. In the context of the DDM1 model, explaining the entire excess signal would require increasing the predicted dust power spectrum by 6×, for example by increasing the assumed uniform polarization fraction in our field from 5% (a typical value) to ∼13%. None of these models show significant cross-correlation with our maps (although this may be interpreted simply as due to limitations of the models).

Taking cross spectra against 100 GHz maps from BICEP1 we find significant correlation and set a constraint on the spectral index of the B-mode excess consistent with CMB and disfavoring dust by 1.7σ. The fact that the BICEP1 and *Keck Array* maps cross correlate with BICEP2 is powerful further evidence against systematics.

An economical interpretation of the B-mode signal which we have detected is that it is largely due to tensor modes—the IGW template is an excellent fit to the observed excess. We therefore proceed to set a constraint on the tensor-to-scalar ratio and find $r = 0.20^{+0.07}_{-0.05}$ with $r = 0$ ruled out at a significance of 7.0σ, with no foreground subtraction. Multiple lines of evidence suggest that the contribution of foregrounds (which will lower the favored value of r) is subdominant: (i) direct projection of the available foreground models using typical model assumptions, (ii) lack of strong cross-correlation of those models against the observed sky pattern (Fig. 6), (iii) the frequency spectral index of the signal as constrained using BICEP1 data at 100 GHz (Fig. 8), and (iv) the power spectral form of the signal and its apparent spatial isotropy (Figs. 3 and 10).

Subtracting the various dust models at their default parameter values and re-deriving the r constraint still results in high significance of detection. As discussed above, one possibility that cannot be ruled out is a larger than anticipated contribution from polarized dust. Given the present evidence disfavoring this, these high values of r are in apparent tension with previous indirect limits based on temperature measurements and we have discussed some possible resolutions including modifications of the initial scalar perturbation spectrum such as running. However, we emphasize that we do not claim to know what the resolution is, if one is in fact necessary.

Figure 14 shows the BICEP2 results compared to previous upper limits. We have pushed into a new regime of sensitivity, and the high-confidence detection of B-mode
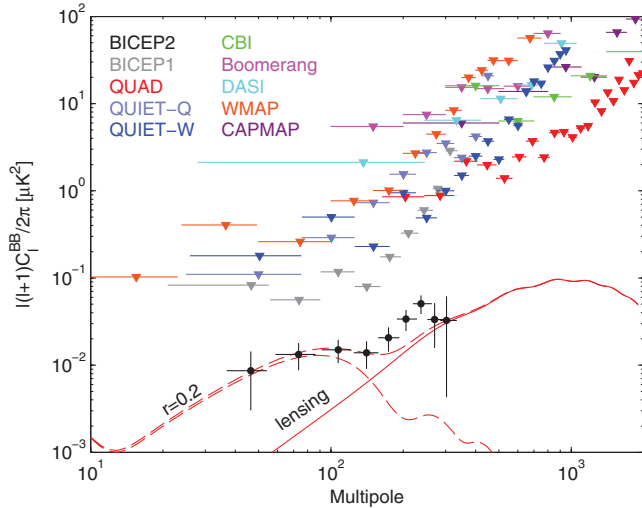
FIG. 14 (color). BICEP2 *BB* auto spectra and 95% upper limits from several previous experiments [2,40,42,43,47,49–51,107]. The curves show the theory expectations for $r = 0.2$ and lensed ΛCDM. The BICEP2 uncertainties include sample variance on an $r = 0.2$ contribution.

polarization at degree angular scales brings us to an exciting juncture. If the origin is in tensors, as favored by the evidence presented above, it heralds a new era of *B*-mode cosmology. However, if these *B* modes represent evidence of a high-dust foreground, it reveals the scale of the challenges that lie ahead.

Note added

Since we submitted this paper new information on polarized dust emission has become available from the *Planck* experiment in a series of papers [96,108–110]. While these confirm that the modal polarization fraction of dust is ∼4%, there is a long tail to fractions as high as 20% (see Fig. 7 of [96]). There is also a trend to higher polarization fractions in regions of lower total dust emission [see Fig. 18 of [96] noting that the BICEP2 field has a column density of $\sim(1–2) \times 10^{20}$ H cm^{-2}]. We note that these papers restrict their analysis to regions of the sky where "systematic uncertainties are small, and where the dust signal dominates total emission," and that this excludes 21% of the sky that includes the BICEP2 region. Thus while these papers do not offer definitive information on the level of dust contamination in our field, they do suggest that it may well be higher than any of the models considered in Sec. IX.

In addition there has been extensive discussion of our preprint in the cosmology community. Two preprints [111,112] look at polarized synchrotron emission in our field and conclude that at 150 GHz it is very small, in broad agreement with our analysis in Sec. IX. Several preprints also examine the new information from *Planck* and raise the same concerns discussed above—that the polarized dust emission may be stronger than any of the models discussed in Sec. IX [112,113]. Given these concerns these studies also reexamine our spectral index constraint described in Sec. X B, since this offers (weak) evidence that the signal is not dust dominated. Both point out that our initial analysis gave the effective spectral index of the whole signal—including the lensing component. Figure 8 now shows an additional curve for the excess over lensing only—the maximum likelihood value is nearly unchanged while the evidence against dust is somewhat weakened. However, Flauger *et al.* [112] also question whether sample variance is properly included in our spectral index analysis, and whether noise in foreground templates could systematically suppress our estimates of template cross correlation. In fact, sample variance is naturally included in the HL-based formalism on which our spectral analysis is based, and the template cross spectra we report are not subject to bias from noise.

More data are clearly required to resolve the situation. We note that cross-correlation of our maps with the *Planck* 353 GHz maps will be more powerful than use of those maps alone in our field. Additional data are also expected from many other experiments, including *Keck Array* observations at 100 GHz in the 2014 season.

PHYSICAL REVIEW LETTERS

*jmkovac@cfa.harvard.edu
†pryke@physics.umn.edu

[1] A. A. Penzias and R. W. Wilson, Astrophys. J. **142**, 419 (1965).

[2] C. L. Bennett, D. Larson, J. L. Weiland, N. Jarosik, G. Hinshaw, N. Odegard, K. M. Smith, R. S. Hill, B. Gold, M. Halpern, E. Komatsu, M. R. Nolta, L. Page, D. N. Spergel, E. Wollack, J. Dunkley, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, and E. L. Wright, Astrophys. J. Suppl. Ser. **208**, 20 (2013).

[3] G. Hinshaw, D. Larson, E. Komatsu, D. N. Spergel, C. L. Bennett, J. Dunkley, M. R. Nolta, M. Halpern, R. S. Hill, N. Odegard, L. Page, K. M. Smith, J. L. Weiland, B. Gold, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, E. Wollack, and E. L. Wright, Astrophys. J. Suppl. Ser. **208**, 19 (2013).

[4] K. T. Story, C. L. Reichardt, Z. Hou, R. Keisler, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H.-M. Cho, T. M. Crawford, A. T. Crites, T. de Haan, M. A. Dobbs, J. Dudley, B. Follin, E. M. George, N. W. Halverson, G. P. Holder, W. L. Holzapfel, *et al.*, Astrophys. J. **779**, 86 (2013).

[5] Z. Hou, C. L. Reichardt, K. T. Story, B. Follin, R. Keisler, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H.-M. Cho, T. M. Crawford, A. T. Crites, T. de Haan, R. de Putter, M. A. Dobbs, S. Dodelson, J. Dudley, E. M. George, N. W. Halverson, *et al.*, Astrophys. J. **782**, 74 (2014).

[6] J. L. Sievers, R. A. Hlozek, M. R. Nolta, V. Acquaviva, G. E. Addison, P. A. R. Ade, P. Aguirre, M. Amiri, J. W. Appel, L. F. Barrientos, E. S. Battistelli, N. Battaglia, J. R. Bond, B. Brown, B. Burger, E. Calabrese, J. Chervenak, D. Crichton, S. Das, M. J. Devlin, *et al.*, J. Cosmol. Astropart. Phys. 10 (2013) 060.

[7] S. Das, T. Louis, M. R. Nolta, G. E. Addison, E. S. Battistelli, J. R. Bond, E. Calabrese, D. Crichton, M. J. Devlin, S. Dicker, J. Dunkley, R. Dünner, J. W. Fowler, M. Gralla, A. Hajian, M. Halpern, M. Hasselfield, M. Hilton, A. D. Hincks, R. Hlozek, *et al.*, J. Cosmol. Astropart. Phys. 4 (2014) 014.

[8] Planck Collaboration XV, arXiv:1303.5075.

[9] Planck Collaboration XVI, arXiv:1303.5076.

[10] R. Brout, F. Englert, and E. Gunzig, Ann. Phys. (N.Y.) **115**, 78 (1978).

[11] A. A. Starobinsky, Phys. Lett. B **91**, 99 (1980).

[12] D. Kazanas, Astrophys. J. Lett. **241**, L59 (1980).

[13] K. Sato, Mon. Not. R. Astron. Soc. **195**, 467 (1981).

[14] A. H. Guth, Phys. Rev. D **23**, 347 (1981).

[15] A. D. Linde, Phys. Lett. B **108**, 389 (1982).

[16] A. D. Linde, Phys. Lett. B **129**, 177 (1983).

[17] A. Albrecht and P. J. Steinhardt, Phys. Rev. Lett. **48**, 1220 (1982).

[18] Planck Collaboration XXII, arXiv:1303.5082.

[19] V. F. Mukhanov and G. V. Chibisov, Sov. J. Exp. Theor. Phys. Lett. **33**, 532 (1981).

[20] S. W. Hawking, Phys. Lett. B **115**, 295 (1982).

[21] A. H. Guth and S.-Y. Pi, Phys. Rev. Lett. **49**, 1110 (1982).

[22] A. A. Starobinsky, Phys. Lett. B **117**, 175 (1982).

[23] J. M. Bardeen, P. J. Steinhardt, and M. S. Turner, Phys. Rev. D **28**, 679 (1983).

[24] V. F. Mukhanov, Zh. Eksp. Teor. Fiz. Pisma Red. **41**, 402 (1985).

[25] D. N. Spergel and M. Zaldarriaga, Phys. Rev. Lett. **79**, 2180 (1997).

[26] H. V. Peiris, E. Komatsu, L. Verde, D. N. Spergel, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, G. S. Tucker, E. Wollack, and E. L. Wright, Astrophys. J. Suppl. Ser. **148**, 213 (2003).

[27] L. P. Grishchuk, Sov. J. Exp. Theor. Phys. **40**, 409 (1975).

[28] A. A. Starobinsky, Zh. Eksp. Teor. Fiz. Pisma Red. **30**, 719 (1979).

[29] V. A. Rubakov, M. V. Sazhin, and A. V. Veryaskin, Phys. Lett. B **115**, 189 (1982).

[30] R. Fabbri and M. D. Pollock, Phys. Lett. B **125**, 445 (1983).

[31] L. F. Abbott and M. B. Wise, Nucl. Phys. **B244**, 541 (1984).

[32] A. Ashoorioon, P. S. Bhupal Dev, and A. Mazumdar, arXiv:1211.4678.

[33] L. M. Krauss and F. Wilczek, Phys. Rev. D **89**, 047501 (2014).

[34] A. G. Polnarev, Sov. Astron. **29**, 607 (1985).

[35] U. Seljak, Astrophys. J. **482**, 6 (1997).

[36] M. Kamionkowski, A. Kosowsky, and A. Stebbins, Phys. Rev. Lett. **78**, 2058 (1997).

[37] U. Seljak and M. Zaldarriaga, Phys. Rev. Lett. **78**, 2054 (1997).

[38] J. M. Kovac, E. M. Leitch, C. Pryke, J. E. Carlstrom, N. W. Halverson, and W. L. Holzapfel, Nature (London) **420**, 772 (2002).

[39] D. Barkats, C. Bischoff, P. Farese, L. Fitzpatrick, T. Gaier, J. O. Gundersen, M. M. Hedman, L. Hyatt, J. J. McMahon, D. Samtleben, S. T. Staggs, K. Vanderlinde, and B. Winstein, Astrophys. J. Lett. **619**, L127 (2005).

[40] C. Bischoff, L. Hyatt, J. J. McMahon, G. W. Nixon, D. Samtleben, K. M. Smith, K. Vanderlinde, D. Barkats, P. Farese, T. Gaier, J. O. Gundersen, M. M. Hedman, S. T. Staggs, B. Winstein, and CAPMAP Collaboration, Astrophys. J. **684**, 771 (2008).

[41] A. C. S. Readhead, S. T. Myers, T. J. Pearson, J. L. Sievers, B. S. Mason, C. R. Contaldi, J. R. Bond, R. Bustos, P. Altamirano, C. Achermann, L. Bronfman, J. E. Carlstrom, J. K. Cartwright, S. Casassus, C. Dickinson, W. L. Holzapfel, J. M. Kovac, E. M. Leitch, J. May, S. Padin, *et al.*, Science **306**, 836 (2004)

[42] J. L. Sievers, C. Achermann, J. R. Bond, L. Bronfman, R. Bustos, C. R. Contaldi, C. Dickinson, P. G. Ferreira, M. E. Jones, A. M. Lewis, B. S. Mason, J. May, S. T. Myers, N. Oyarce, S. Padin, T. J. Pearson, M. Pospieszalski, A. C. S. Readhead, R. Reeves, A. C. Taylor, and S. Torres, Astrophys. J. **660**, 976 (2007).

[43] T. E. Montroy, P. A. R. Ade, J. J. Bock, J. R. Bond, J. Borrill, A. Boscaleri, P. Cabella, C. R. Contaldi, B. P. Crill, P. de Bernardis, G. De Gasperis, A. de Oliveira-Costa, G. De Troia, G. di Stefano, E. Hivon, A. H. Jaffe, T. S. Kisner, W. C. Jones, A. E. Lange, S. Masi, *et al.*, Astrophys. J. **647**, 813 (2006).

[44] L. Page, G. Hinshaw, E. Komatsu, M. R. Nolta, D. N. Spergel, C. L. Bennett, C. Barnes, R. Bean, O. Doré,

J. Dunkley, M. Halpern, R. S. Hill, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, N. Odegard, H. V. Peiris, G. S. Tucker, L. Verde, *et al.*, Astrophys. J. Suppl. Ser. **170**, 335 (2007).

[45] J. H. P. Wu, J. Zuntz, M. E. Abroe, P. A. R. Ade, J. Bock, J. Borrill, J. Collins, S. Hanany, A. H. Jaffe, B. R. Johnson, T. Jones, A. T. Lee, T. Matsumura, B. Rabii, T. Renbarger, P. L. Richards, G. F. Smoot, R. Stompor, H. T. Tran, and C. D. Winant, Astrophys. J. **665**, 55 (2007).

[46] C. Pryke, P. Ade, J. Bock, M. Bowden, M. L. Brown, G. Cahill, P. G. Castro, S. Church, T. Culverhouse, R. Friedman, K. Ganga, W. K. Gear, S. Gupta, J. Hinderks, J. Kovac, A. E. Lange, E. Leitch, S. J. Melhuish, Y. Memari, J. A. Murphy, *et al.*, Astrophys. J. **692**, 1247 (2009).

[47] M. L. Brown, P. Ade, J. Bock, M. Bowden, G. Cahill, P. G. Castro, S. Church, T. Culverhouse, R. B. Friedman, K. Ganga, W. K. Gear, S. Gupta, J. Hinderks, J. Kovac, A. E. Lange, E. Leitch, S. J. Melhuish, Y. Memari, J. A. Murphy, A. Orlando, *et al.*, and QUa D Collaboration, Astrophys. J. **705**, 978 (2009).

[48] H. C. Chiang, P. A. R. Ade, D. Barkats, J. O. Battle, E. M. Bierman, J. J. Bock, C. D. Dowell, L. Duband, E. F. Hivon, W. L. Holzapfel, V. V. Hristov, W. C. Jones, B. G. Keating, J. M. Kovac, C. L. Kuo, A. E. Lange, E. M. Leitch, P. V. Mason, T. Matsumura, H. T. Nguyen, *et al.*, Astrophys. J. **711**, 1123 (2010).

[49] D. Barkats, R. Aikin, C. Bischoff, I. Buder, J. P. Kaufman, B. G. Keating, J. M. Kovac, M. Su, P. A. R. Ade, J. O. Battle, E. M. Bierman, J. J. Bock, H. C. Chiang, C. D. Dowell, L. Duband, J. Filippini, E. F. Hivon, W. L. Holzapfel, V. V. Hristov, W. C. Jones, *et al.*, Astrophys. J. **783**, 67 (2014).

[50] QUIET Collaboration, Astrophys. J. **741**, 111 (2011).

[51] QUIET Collaboration, Astrophys. J. **760**, 145 (2012).

[52] D. Hanson, S. Hoover, A. Crites, P. A. R. Ade, K. A. Aird, J. E. Austermann, J. A. Beall, A. N. Bender, B. A. Benson, L. E. Bleem, J. J. Bock, J. E. Carlstrom, C. L. Chang, H. C. Chiang, H.-M. Cho, A. Conley, T. M. Crawford, T. de Haan, M. A. Dobbs, W. Everett, *et al.*, Phys. Rev. Lett. **111**, 141301 (2013).

[53] POLARBEAR Collaboration, Phys. Rev. Lett. **112**, 131302 (2014).

[54] POLARBEAR Collaboration, arXiv:1312.6646 [Phys. Rev. Lett. (to be published)].

[55] POLARBEAR Collaboration, arXiv:1403.2369.

[56] T. Essinger-Hileman, J. W. Appel, J. A. Beal, H. M. Cho, J. Fowler, M. Halpern, M. Hasselfield, K. D. Irwin, T. A. Marriage, M. D. Niemack, L. Page, L. P. Parker, S. Pufu, S. T. Staggs, O. Stryzak, C. Visnjic, K. W. Yoon, and Y. Zhao, AIP Conf. Proc. **1185**, 494 (2009).

[57] M. D. Niemack, P. A. R. Ade, J. Aguirre, F. Barrientos, J. A. Beall, J. R. Bond, J. Britton, H. M. Cho, S. Das, M. J. Devlin, S. Dicker, J. Dunkley, R. Dünner, J. W. Fowler, A. Hajian, M. Halpern, M. Hasselfield, G. C. Hilton, M. Hilton, J. Hubmayr, *et al.*, Proc. SPIE Int. Soc. Opt. Eng. **7741**, 77411S (2010).

[58] J. R. Eimer, C. L. Bennett, D. T. Chuss, T. Marriage, E. J. Wollack, and L. Zeng, Proc. SPIE Int. Soc. Opt. Eng. **8452** , 845220 (2012).

[59] B. Reichborn-Kjennerud, A. M. Aboobaker, P. Ade, F. Aubin, C. Baccigalupi, C. Bao, J. Borrill, C. Cantalupo,

D. Chapman, J. Didier, M. Dobbs, J. Grain, W. Grainger, S. Hanany, S. Hillbrand, J. Hubmayr, A. Jaffe, B. Johnson, T. Jones, T. Kisner, *et al.*, Proc. SPIE Int. Soc. Opt. Eng. **7741**, 77411C (2010).

[60] A. A. Fraisse, P. A. R. Ade, M. Amiri, S. J. Benton, J. J. Bock, J. R. Bond, J. A. Bonetti, S. Bryan, B. Burger, H. C. Chiang, C. N. Clark, C. R. Contaldi, B. P. Crill, G. Davis, O. Doré, M. Farhang, J. P. Filippini, L. M. Fissel, N. N. Gandilo, S. Golwala, *et al.*, J. Cosmol. Astropart. Phys. 4 (2013) 047.

[61] A. Kogut, P. A. R. Ade, D. Benford, C. L. Bennett, D. T. Chuss, J. L. Dotson, J. R. Eimer, D. J. Fixsen, M. Halpern, G. Hilton, J. Hinderks, G. F. Hinshaw, K. Irwin, C. Jhabvala, B. Johnson, J. Lazear, L. Lowe, T. Miller, P. Mirel, S. H. Moseley, S. Rodriguez, *et al.*, Proc. SPIE Int. Soc. Opt. Eng. **8452**, 84521J (2012).

[62] B. G. Keating, P. A. R. Ade, J. J. Bock, E. Hivon, W. L. Holzapfel, A. E. Lange, H. Nguyen, and K. W. Yoon, Proc. SPIE Int. Soc. Opt. Eng. **4843**, 284 (2003).

[63] Y. D. Takahashi, P. A. R. Ade, D. Barkats, J. O. Battle, E. M. Bierman, J. J. Bock, H. C. Chiang, C. D. Dowell, L. Duband, E. F. Hivon, W. L. Holzapfel, V. V. Hristov, W. C. Jones, B. G. Keating, J. M. Kovac, C. L. Kuo, A. E. Lange, E. M. Leitch, P. V. Mason, T. Matsumura, *et al.*, Astrophys. J. **711**, 1141 (2010).

[64] BICEP2 Collaboration II, arXiv:1403.4302.

[65] BICEP2 Collaboration III (to be published).

[66] J. J. Bock, Proc. SPIE Int. Soc. Opt. Eng. **4843**, 314 (2003).

[67] K. D. Irwin, Appl. Phys. Lett. **66**, 1998 (1995).

[68] K. Story, E. Leitch, P. Ade, K. A. Aird, J. E. Austermann, J. A. Beall, D. Becker, A. N. Bender, B. A. Benson, L. E. Bleem, J. Britton, J. E. Carlstrom, C. L. Chang, H. C. Chiang, H.-M. Cho, T. M. Crawford, A. T. Crites, A. Datesman, T. de Haan, M. A. Dobbs, *et al.*, Proc. SPIE Int. Soc. Opt. Eng. **8451**, 84510T (2012).

[69] M. Shimon, B. Keating, N. Ponthieu, and E. Hivon, Phys. Rev. D **77**, 083003 (2008).

[70] http://irsa.ipac.caltech.edu/data/Planck/release_1/all-sky-maps/previews/HFI_SkyMap_143_2048_R1.10_nominal/index.html.

[71] C. D. Sheehy, Ph.D. dissertation, University of Chicago, 2013.

[72] E. Hivon, K. M. Górski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen, Astrophys. J. **567**, 2 (2002).

[73] http://irsa.ipac.caltech.edu/data/Planck/release_1/all-sky-maps/previews/COM_CompMap_CMB-nilc_2048_R1.20/index.html.

[74] http://healpix.sourceforge.net/.

[75] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, Astrophys. J. **622**, 759 (2005).

[76] http://camb.info/.

[77] http://cosmologist.info/lenspix/.

[78] A. Lewis, Phys. Rev. D **71**, 083008 (2005).

[79] A. van Engelen, R. Keisler, O. Zahn, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. M. Cho, T. M. Crawford, A. T. Crites, T. de Haan, M. A. Dobbs, J. Dudley, E. M. George, N. W. Halverson, G. P.

Holder, W. L. Holzapfel, S. Hoover, Z. Hou, *et al.*, Astrophys. J. **756**, 142 (2012).

[80] K. M. Smith, Phys. Rev. D **74**, 083002 (2006).

[81] M. Tegmark and A. de Oliveira-Costa, Phys. Rev. D **64**, 063001 (2001).

[82] E. F. Bunn, M. Zaldarriaga, M. Tegmark, and A. de Oliveira-Costa, Phys. Rev. D **67**, 023501 (2003).

[83] L. Knox, Phys. Rev. D **60**, 103516 (1999).

[84] These band powers along with all the ancillary data, noise information, likelihood code, etc. required to use them are available for download at http://bicepkeck.org/.

[85] B. G. Keating, M. Shimon, and A. P. S. Yadav, Astrophys. J. Lett. **762**, L23 (2013).

[86] J. P. Kaufman, N. J. Miller, M. Shimon, D. Barkats, C. Bischoff, I. Buder, B. G. Keating, J. M. Kovac, P. A. R. Ade, R. Aikin, J. O. Battle, E. M. Bierman, J. J. Bock, H. C. Chiang, C. D. Dowell, L. Duband, J. Filippini, E. F. Hivon, W. L. Holzapfel, V. V. Hristov, *et al.*, Phys. Rev. D **89**, 062006 (2014).

[87] J. Dunkley, A. Amblard, C. Baccigalupi, M. Betoule, D. Chuss, A. Cooray, J. Delabrouille, C. Dickinson, G. Dobler, J. Dotson, H. K. Eriksen, D. Finkbeiner, D. Fixsen, P. Fosalba, A. Fraisse, C. Hirata, A. Kogut, J. Kristiansen, C. Lawrence, A. M. Magalhães, M. A. Miville-Deschenes, *et al.*, AIP Conf. Proc. **1141**, 222 (2009).

[88] D. P. Finkbeiner, M. Davis, and D. J. Schlegel, Astrophys. J. **524**, 867 (1999).

[89] D. T. O'Dea, C. N. Clark, C. R. Contaldi, and C. J. MacTavish, Mon. Not. R. Astron. Soc. **419**, 1795 (2012).

[90] http://www3.imperial.ac.uk/people/c.contaldi/fgpol.

[91] A. Kogut, J. Dunkley, C. L. Bennett, O. Doré, B. Gold, M. Halpern, G. Hinshaw, N. Jarosik, E. Komatsu, M. R. Nolta, N. Odegard, L. Page, D. N. Spergel, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright, Astrophys. J. **665**, 355 (2007).

[92] J. Delabrouille, M. Betoule, J.-B. Melin, M.-A. Miville-Deschênes, J. Gonzalez-Nuevo, M. Le Jeune, G. Castex, G. de Zotti, S. Basak, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, J.-P. Bernard, F. R. Bouchet, D. L. Clements, A. da Silva, C. Dickinson, F. Dodu, K. Dolag, *et al.*, Astron. Astrophys. **553**, A96 (2013) .

[93] http://www.apc.univ-paris7.fr/~delabrou/PSM/psm.html.

[94] Planck Collaboration XI, arXiv:1312.1300.

[95] In the preprint version of this paper an additional DDM2 model was included based on information taken from *Planck* conference talks. We noted the large uncertainties

on this and the other dust models presented. In the *Planck* dust polarization paper [96] which has since appeared the maps have been masked to include only regions "where the systematic uncertainties are small, and where the dust signal dominates total emission." This mask excludes our field. We have concluded the information used for the DDM2 model has unquantifiable uncertainty. We look forward to performing a cross-correlation analysis against the *Planck* 353 GHz polarized maps in a future publication.

[96] Planck Collaboration, arXiv:1405.0871v1.

[97] Planck Collaboration XXVIII, arXiv:1303.5088.

[98] M. Massardi, R. D. Ekers, T. Murphy, E. Mahony, P. J. Hancock, R. Chhetri, G. de Zotti, E. M. Sadler, S. Burke-Spolaor, M. Calabretta, P. G. Edwards, J. A. Ekers, C. A. Jackson, M. J. Kesteven, K. Newton-McGee, C. Phillips, R. Ricci, P. Roberts, R. J. Sault, L. Staveley-Smith, *et al.*, Mon. Not. R. Astron. Soc. **412**, 318 (2011).

[99] R. A. Battye, I. W. A. Browne, M. W. Peel, N. J. Jackson, and C. Dickinson, Mon. Not. R. Astron. Soc. **413**, 132 (2011).

[100] S. Hamimeche and A. Lewis, Phys. Rev. D **77**, 103013 (2008).

[101] S. S. Wilks, Ann. Math. Stat. **9**, 60 (1938).

[102] C. D. Sheehy, P. A. R. Ade, R. W. Aikin, M. Amiri, S. Benton, C. Bischoff, J. J. Bock, J. A. Bonetti, J. A. Brevik, B. Burger, C. D. Dowell, L. Duband, J. P. Filippini, S. R. Golwala, M. Halpern, M. Hasselfield, G. Hilton, V. V. Hristov, K. Irwin, J. P. Kaufman, *et al.*, Proc. SPIE Int. Soc. Opt. Eng. **7741**, 77411R (2010).

[103] D. Gamerman and H. F. Lopes, *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference* (Chapman and Hall/CRC, London, U.K., 2006).

[104] A. Lewis and S. Bridle, Phys. Rev. D **66**, 103511 (2002).

[105] As downloaded from http://www.sciops.esa.int/wikiSI/planckpla, section "Cosmological Parameters."

[106] W. Hastings, Biometrika **57**, 97 (1970).

[107] E. M. Leitch, J. M. Kovac, N. W. Halverson, J. E. Carlstrom, C. Pryke, and M. W. E. Smith, Astrophys. J. **624**, 10 (2005).

[108] Planck Collaboration, arXiv:1405.0872v1.

[109] Planck Collaboration, arXiv:1405.0873v1.

[110] Planck Collaboration, arXiv:1405.0874v2.

[111] U. Fuskeland, I. K. Wehus, H. K. Eriksen, and S. K. Næss, arXiv:1404.5323v1.

[112] R. Flauger, J. C. Hill, and D. N. Spergel, arXiv:1405.7351v1.

[113] M. J. Mortonson and U. Seljak, arXiv:1405.5857v1.