

DEGREE-SCALE COSMIC MICROWAVE BACKGROUND POLARIZATION MEASUREMENTS FROM THREE YEARS OF BICEP1 DATA

BICEP1 COLLABORATION—D. BARKATS¹, R. AIKIN², C. BISCHOFF³, I. BUDER³, J. P. KAUFMAN⁴, B. G. KEATING⁴, J. M. KOVAC³, M. SU^{5,6}, P. A. R. ADE⁷, J. O. BATTLE⁸, E. M. BIERMAN⁴, J. J. BOCK^{2,8}, H. C. CHIANG⁹, C. D. DOWELL⁸, L. DUBAND¹⁰, J. FILIPPINI², E. F. HIVON¹¹, W. L. HOLZAPFEL¹², V. V. HRISTOV², W. C. JONES¹³, C. L. KUO^{14,15}, E. M. LEITCH¹⁶, P. V. MASON², T. MATSUMURA¹⁷, H. T. NGUYEN⁸, N. PONTIEU¹⁸, C. PRYKE¹⁹, S. RICHTER², G. ROCHA^{2,8}, C. SHEEHY¹⁶, S. S. KERNASOVSKIY^{14,15}, Y. D. TAKAHASHI¹², J. E. TOLAN^{14,15}, AND K. W. YOON^{14,15}

¹ Joint ALMA Observatory, ESO, Santiago, Chile; dbarkats@alma.cl

² Department of Physics, California Institute of Technology, Pasadena, CA 91125, USA

³ Harvard-Smithsonian Center for Astrophysics, 60 Garden Street MS 42, Cambridge, MA 02138, USA

⁴ Department of Physics, University of California at San Diego, La Jolla, CA 92093, USA

⁵ Department of Physics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

⁶ MIT-Kavli Center for Astrophysics and Space Research, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

⁷ Department of Physics and Astronomy, University of Wales, Cardiff, CF24 3YB Wales, UK

⁸ Jet Propulsion Laboratory, Pasadena, CA 91109, USA

⁹ Astrophysics and Cosmology Research Unit, University of KwaZulu-Natal, Durban, South Africa

¹⁰ SBT, Commissariat à l’Energie Atomique, Grenoble F-38041, France

¹¹ Institut d’Astrophysique de Paris, Paris, France

¹² Department of Physics, University of California at Berkeley, Berkeley, CA 94720, USA

¹³ Department of Physics, Princeton University, Princeton, NJ 08544, USA

¹⁴ Stanford University, Palo Alto, CA 94305, USA

¹⁵ Kavli Institute for Particle Astrophysics and Cosmology (KIPAC), Sand Hill Road 2575, Menlo Park, CA 94025, USA

¹⁶ University of Chicago, Chicago, IL 60637, USA

¹⁷ High Energy Accelerator Research Organization (KEK), Ibaraki 305-0801, Japan

¹⁸ Institut d’Astrophysique Spatiale, Université Paris-Sud, Orsay F-91405, France

¹⁹ Department of Physics, University of Minnesota, Minneapolis, MN 55455, USA

Received 2013 October 4; accepted 2014 January 7; published 2014 February 14

ABSTRACT

BICEP1 is a millimeter-wavelength telescope designed specifically to measure the inflationary *B*-mode polarization of the cosmic microwave background at degree angular scales. We present results from an analysis of the data acquired during three seasons of observations at the South Pole (2006–2008). This work extends the two-year result published in Chiang et al., with additional data from the third season and relaxed detector-selection criteria. This analysis also introduces a more comprehensive estimation of band power window functions, improved likelihood estimation methods, and a new technique for deprojecting monopole temperature-to-polarization leakage that reduces this class of systematic uncertainty to a negligible level. We present maps of temperature, *E*- and *B*-mode polarization, and their associated angular power spectra. The improvement in the map noise level and polarization spectra error bars are consistent with the 52% increase in integration time relative to Chiang et al. We confirm both self-consistency of the polarization data and consistency with the two-year results. We measure the angular power spectra at $21 \leq \ell \leq 335$ and find that the *EE* spectrum is consistent with Lambda cold dark matter cosmology, with the first acoustic peak of the *EE* spectrum now detected at 15σ . The *BB* spectrum remains consistent with zero. From *B*-modes only, we constrain the tensor-to-scalar ratio to $r = 0.03^{+0.27}_{-0.23}$, or $r < 0.70$ at 95% confidence level.

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

1. INTRODUCTION

In recent years, observational cosmology has produced enormous advances in our understanding of the universe. Observations of the cosmic microwave background (CMB) have played a central role in establishing what is now known as the standard cosmological model. Measurements of the CMB temperature anisotropies have reached sub-percent precision over the whole sky at a range of angular scales down to few-arcminute resolution. The angular power spectrum of temperature anisotropies has yielded tight constraints on the basic parameters of that cosmological model, referred to as Lambda cold dark matter (Λ CDM). Although temperature anisotropy experiments (e.g., Story et al. 2013; Hinshaw et al. 2013; Planck Collaboration et al. 2013a; Sievers et al. 2013) continue to test the validity of Λ CDM, the model by itself offers no solution to the

following mysteries: the high degree of flatness of the universe, the apparent large-scale correlations that suggest a larger particle horizon than allowed by the standard big bang scenario, the nearly scale-invariant spectrum of initial perturbations, and the lack of relic magnetic monopoles.

The inflationary scenario was proposed as an explanation to these observed properties of the universe (for review, see Liddle & Lyth 2000). One as-yet-unobserved prediction of inflation is a stochastic gravitational wave background that would imprint its signature on the anisotropies of the CMB. The most powerful method to search for this signature is to constrain the curl-mode (*B*-mode) polarization pattern of the CMB at degree angular scales (Seljak 1997; Seljak & Zaldarriaga 1997; Kamionkowski et al. 1997).

The CMB is polarized at the 10% level due to Thomson scattering at the surface of last scattering. The density perturbations

that give rise to the temperature anisotropies also cause the plasma to flow along gradients of this density field and so can only create gradient-mode, or E -mode, polarization with zero curl (Hu & White 1997). Since the first detection of E -mode polarization in 2002 (Kovac et al. 2002), several other experiments have refined the characterization of the EE and TE spectra (Montroy et al. 2006; Sievers et al. 2007; Wu et al. 2007; Bischoff et al. 2008; Brown et al. 2009; Chiang et al. 2010; QUIET Collaboration et al. 2011, 2012; Bennett et al. 2013). So far all have unambiguously confirmed the basic tenets of Λ CDM.

In addition to being a prediction of inflation, B -modes can also be generated through gravitational lensing of E -modes, producing a signature that is observationally distinct from inflationary B -modes and peaks at smaller angular scales. The lensing B -mode polarization has recently been detected, using cross-correlations formed with an external lensing template (Hanson et al. 2013). The inflationary B -mode pattern in the CMB polarization still remains elusive. A detection of primordial B -mode polarization would provide strong support to the inflationary scenario.

The amplitude of the B -mode signal is parameterized by the tensor-to-scalar ratio r . The best constraints on r are currently derived from CMB temperature anisotropies: $r < 0.11$ at 95% confidence for models that add only tensors to Λ CDM, or $r < 0.23$ for models allowing running of the spectral index (Story et al. 2013; Planck Collaboration et al. 2013b). Cosmic variance limits further improvements on r constraints using temperature anisotropies alone.

The best limit on r using only B -mode polarization is $r < 0.72$ at 95% CL, set using BICEP1 data (Chiang et al. 2010, hereafter C10). This previous result only used the first two years of observations and conservative detector-selection criteria. In this paper, we report measurements of the CMB polarization power spectra and improved constraints on r , using data from all three years, as well as relaxed detector-selection criteria. We also present an updated data-analysis method that includes an improved noise model, a more sophisticated calculation of the band power window functions, and new likelihood estimation techniques. Another unique addition to this analysis is a deprojection filter that suppresses instrumental temperature-to-polarization leakage from relative-gain miscalibration. Many of these techniques were developed jointly with successor experiments (BICEP2 and the Keck Array; Ogburn et al. 2012), and they will only grow in importance with improved instrumental sensitivity.

2. THE BICEP1 INSTRUMENT

In this section, we summarize the salient features of the BICEP1 instrument. More complete details are available in Takahashi et al. (2010, hereafter T10) and C10, as well as in several theses: Yoon (2008), Chiang (2009), Takahashi (2010), Bierman (2011), and Moyerman (2013).

BICEP1 is a bolometric polarimeter that is specifically designed to search for the signature of inflation in the B -mode of the CMB polarization. Its detectors, optical path, scan strategy, target region, and site are all designed to provide the highest possible sensitivity while minimizing polarization systematics.

The BICEP1 receiver consists of a focal plane of 49 polarization-sensitive-bolometer pairs (PSBs; Jones et al. 2003). The two detectors in each pair respond to orthogonal linear polarizations. We derive CMB temperature measurements from the summed pair response and polarization from the differenced pair response. A two-lens refracting telescope couples

the PSBs to the sky, providing FWHM angular resolution of 0°:93 and 0°:60 at 100 and 150 GHz, respectively, and an instantaneous field of view of 18°. The focal plane has 25 PSB pairs at 100 GHz, 22 at 150 GHz, and two at 220 GHz. These quantities reflect the 2007/2008 configuration. For 2006 only, six pairs (three at 100 GHz and three at 150 GHz) were used with Faraday rotation modulators (FRMs; Moyerman et al. 2013). The 220 GHz detectors were introduced in 2007 (Bierman et al. 2011).

BICEP1 was installed in the Dark Sector Laboratory (89°99 S, 44°65 W) at the Amundsen–Scott South Pole station to take advantage of the excellent millimeter transparency of the atmosphere above the cold polar plateau. The telescope mount provides three-axis motion: azimuth, elevation, and boresight rotation. The telescope is fully enclosed inside the warm lab with only the aperture exposed to the polar environment. The aperture is surrounded by a co-moving absorptive baffle and a large, fixed reflective ground screen to minimize any potential contamination from warm ground emission.

During its three seasons of operation, BICEP1 observed three fields, concentrating 85% of its observing time on one CMB region selected for low galactic dust emission. This region, called the “Southern Hole,” is located at a right ascension and declination range of $|\alpha| < 60^\circ$ and $-70^\circ < \delta < -45^\circ$. The telescope operated on a 48 hr observing cycle, containing four 9 hr “phases” targeting the Southern Hole. Each phase was further divided into 10 azimuth-fixed “scansets,” approximately 50 minutes long, during which the telescope scanned across the full 60° range of azimuth at a fixed elevation. Each scanset comprises 50 left-going and 50 right-going “half-scans.” Each scanset was bracketed by elevation nods, which are small (1°:2 peak-to-peak) excursions in telescope elevation used to calibrate relative detector gains from the atmospheric signal. The telescope stepped 0°:25 in elevation between each scanset and covered the full CMB field after two phases. The boresight rotation changed between observing cycles, stepping between four orientations (0°, 45°, 180°, and 225°) chosen to improve polarization angle coverage. The remaining 12 hr from each 48 hr cycle were spent on cryogenic service (6 hr) and Galactic field observations (6 hr; Bierman et al. 2011).

The focal plane, target field, scan strategy, observation cycle, and calibration methods remained unchanged from the 2007 to the 2008 season. As a result, for the analysis of the three-year data set, we use the same parameters for the detector transfer functions, relative gains, polarization orientation and efficiency, and beam shapes as those presented in C10. Similarly, we follow the same procedure for deriving the absolute gain calibration, boresight, and detector pointing as in C10.

3. DATA SELECTION

This analysis uses a data set that has been expanded since C10, most significantly by the inclusion of a third year of observations. The first two years of data include a total of 736 9 hr CMB phases, with 248 phases in 2006 and 488 phases in 2007; the 2008 season contributes another 270 phases, increasing the total by 37% to 1006 phases. The first and last season contribute less integration time due to time spent refining the observing schedule at the start of the 2006 season and time spent on final calibrations at the end of the 2008 season. As in C10, we exclude a small number of incomplete CMB observing phases from the analysis.

Table 1
BICEP1 Total Observing Efficiency

Activity	Days Spent ^a	Section
Installation, upgrades, calibration	271 days	3
Alternate observing modes	321 days	3
CMB observations	503 days	3
Down-selection	Fraction kept in CMB analysis ^b	...
Cryogenic service and galaxy observations ^c	75.0%	3.1
Scanset calibrations ^d	83.0%	3.1
Scan turn-arounds ^e	74.2%	3.1
Detector yield ^f	89.6%	3.2
Weather cut	92.5%	3.3
Half-scan cuts	96.8%	3.3
Total pass fraction	37.1%	...

Notes.

^a Number of days out of three calendar years, 2006–2008.

^b Percentages in this column describe what fraction of the time spent on CMB observations is ultimately included in analysis. The fractions can be applied cumulatively to obtain the total pass fraction (bottom row).

^c “Cryogenic service and galaxy observations” combines the two 6 hr periods from each 48 hr cycle that did not target the CMB field.

^d “Scanset calibrations” refers to the fraction of time in each observing phase spent performing elevation nods, as well as brief scans over the Galactic field.

^e “Scan turn-arounds” refer to the periods of acceleration at either end of each azimuthal half-scan, which are cut from the analysis.

^f Fraction of 100 and 150 GHz detectors used for CMB analysis, weighted average across observing seasons.

3.1. Observing Efficiency

Table 1 describes BICEP1’s total observing efficiency, relative to a hypothetical experiment with perfect detector yield, no time spent on cryogenic service or calibration, and no weather cuts.

In the top section of the table, we divide up three calendar years (1095 days) into time spent on summer activities (deployment, upgrades, and the summer calibration described in T10), time spent on CMB observations of the Southern Hole that are used for this analysis, and alternate observing modes, including observations using FRM detectors, published in Moyerman et al. (2013), and Galactic field observations, published in Bierman et al. (2011). During its three years of operation, BICEP1 spent 46% (503 days out of 1095) of its time on the primary science target.

The bottom section of Table 1 describes the observing cycle efficiency during CMB observations. Although the final fraction of time spent observing the CMB seems low, this summary of BICEP1 operations describes an instrument that achieved a goal of extremely focused observation on its target field.

3.2. Detector Selection

Of the 49 optically active PSB pairs in the focal plane, the analysis in C10 excluded many due to various unexpected behaviors such as poorly behaved transfer functions. The transfer functions are used to deconvolve the raw detector timestream into a cleaned timestream usable for analysis and allow us to link the relative gain measured at 0.02 Hz via elevation nods to the science band (0.1–1 Hz). Although we deconvolve the measured transfer functions, a fit to a phenomenological model was used to identify poorly behaved detectors. C10 excluded all detectors with larger than 0.2% residuals to the model fit. For this analysis, we include those detectors. Although their transfer functions do not follow the common BICEP1 bolometer model,

they are measured sufficiently well over the frequency range of interest.

The analysis of C10 included 33 PSB pairs (19 at 100 GHz, 14 at 150 GHz) in 2006 and 37 pairs (22 at 100 GHz, 15 at 150 GHz) in 2007. For this analysis, the count increases to 36 PSB pairs (19 at 100 GHz, 17 at 150 GHz) in 2006 and 43 pairs (23 at 100 GHz, 20 at 150 GHz) in 2007 and 2008. We exclude the six detector pairs containing FRMs for the 2006 season and two 220 GHz pairs from the 2007 and 2008 seasons.

Averaged across the three observing seasons, the number of detector pairs included in the analysis increased by 12%. The addition of these detectors together with the third season of observations, brings the total increase in integration time over C10 to 52%.

3.3. Data Cuts

Starting from the expanded data set, we remove some data that suffer from bad weather or glitches in the pair-difference timestreams. The cut criteria are identical to those used in C10, but we briefly review them here and present updated cut fractions for the three-year data set.

The first cut is designed to remove entire phases affected by bad weather. For each PSB and phase, we compute the standard deviation of the 10 relative-gain measurements from elevation nods made during the phase. The median of these standard deviations is calculated separately for 100 GHz and 150 GHz PSBs. If either median value exceeds a threshold, selected to be 20% of the typical relative-gain value, then the entire phase is cut for all detectors at both frequencies. Applying these criteria reduces the number of phases from 1006 to 930, a cut fraction of 7.5%.

Next, we cut individual PSB pairs at a single half-scan level according to three criteria:

1. A detector pair is cut for all half-scans in a scanset if the A/B relative-gain ratio differs by more than 3% between the two elevation nods bracketing that scanset.
2. Pairs are cut for any individual half-scans where the pair-difference data show significant skewness or kurtosis.
3. Half-scans containing large glitches (in excess of 7σ from the mean) are cut.

If the combination of these cuts removes more than half of the data for a PSB pair in a particular scanset, we take that as evidence of unreliable behavior and cut that pair for the entire scanset. Altogether, the half-scan cuts remove 3.2% of the data, significantly less than the weather cut. The skew/kurtosis cut is the most important of the set; dropping it entirely would lower the cut fraction to 2%.

4. MAPMAKING

C10 presented results from two analysis pipelines. The primary result came from an analysis that utilized HEALPix map binning (Górski et al. 2005) and the Spice power-spectrum estimator (Chon et al. 2004). The “alternate analysis pipeline” of C10, which uses an equirectangular map pixelization and two-dimensional Fourier transform (2D FT) for power-spectrum estimation, is the only pipeline used for the current work. This pipeline has been derived from one originally developed for the QUAD experiment (Pryke et al. 2009).

4.1. Low-level Processing

The low-level timestream processing is unchanged from C10. Each 9 hr phase is cut to exclude elevation nods bracketing each

scanset and periods of acceleration at either end of each half-scan. The detector transfer function is deconvolved from each detector timestream, which is then low-pass filtered at 5 Hz and down-sampled to 10 Hz.

A relative-gain correction is calculated for each detector by measuring its response to atmospheric signal during the elevation nods and comparing to the mean response of all good detectors. After all of the detectors have been scaled to have the same response to the gradient of the airmass, the timestreams for each pair are summed and differenced to produce timestreams sensitive to temperature and polarization, respectively. At this point, we calculate the pointing and polarization angle of each pair and apply a polarization efficiency correction to the pair-difference data. For details of the pointing and polarization corrections, see T10.

4.2. Timestream Filtering

Pair-sum and pair-difference timestreams are filtered to remove atmospheric and azimuth-fixed noise prior to mapmaking. First, we fit and remove a third-order polynomial in azimuth from each half-scan. Next, for each pair-sum or pair-difference timestream, all half-scans in a scanset are binned together in azimuth and the resulting scan-synchronous structure is subtracted from the data. This azimuth-fixed filter targets contamination signals from the ground, which remain stationary over the scanset duration while sky signals rotate under sidereal motion. For discussion of how the filtering affects the signal power measured in the maps and the power-spectrum error bars, see Section 5.4.

4.3. Weighting and Map Binning

The filtered timestreams are binned into maps using the calculated pointing trajectories for each detector pair. Pair-sum timestreams contribute to a CMB temperature (T) map, while pair-difference timestreams contribute to maps of the Stokes Q and U parameters, where the particular linear combination of Q and U measured in each sample is calculated based on the measured polarization angle of the detector pair and the telescope boresight orientation. The T , Q , and U maps all use a common equirectangular pixelization, with 0.25° square pixels at declination -57.5° .

In the map binning operation, we weight each half-scan by its inverse variance, calculated as the variance of the time-ordered data after application of the polynomial and azimuth-fixed filters. We weight the pair-sum and pair-difference timestream separately. This choice of weighting is different from the one used in C10, which applied a uniform weighting for all half-scans for a particular detector pair in a scanset based on the power spectral density in the range 0.5–1 Hz. We build weight maps corresponding to the T , Q , and U maps, which are used as apodization masks for power-spectrum estimation (see Section 5.1).

4.4. Relative-gain Deprojection

As described in T10, a miscalibration in the relative gain of the two detectors in a pair has the effect of leaking the CMB temperature signal into the polarization maps constructed from that detector pair. Relative-gain mismatch can arise from a difference in spectral response between detectors within a pair. The atmospheric signal used to measure the relative gains has a different spectrum than the CMB; in the presence of small mismatches in the bandpass within a polarization pair,

the relative gains chosen to reject an unpolarized atmospheric signal do not perfectly remove the CMB temperature.

To mitigate systematic uncertainties arising from relative-gain mismatch, we implement a “relative-gain deprojection” analysis, described in detail in R. Aikin et al. (in preparation). Over each 9 hr phase, a template timestream for the CMB temperature signal is constructed from the *Wilkinson Microwave Anisotropy Probe* (WMAP) seven-year V-band map (Jarosik et al. 2011), which has been smoothed to the resolution of BICEP1. This template undergoes the same filtering operations as the real data. Then, we perform a linear regression of the 9 hr pair-difference timestream against the leakage template. The individual half-scans are weighted by their inverse variance for this regression, the same weighting that is used for mapmaking. The coefficient obtained from the regression is proportional to the amplitude of the CMB temperature signal that has leaked into the pair-difference data due to relative-gain mismatch. This temperature-to-polarization leakage signal is subtracted from the data before mapmaking.

While the regression coefficient for a particular detector pair in a single 9 hr phase is noisy, we can average over many phases to obtain a better measurement of the relative-gain mismatch for each pair. We find six detector pairs at 150 GHz and one pair at 100 GHz whose relative-gain regression coefficient exceeds 1%, which is consistent with the findings of T10. With the relative-gain leakage thus removed, we find that the residual from this systematic has been suppressed to negligible levels (see Section 7.2).

In addition to relative-gain deprojection, R. Aikin et al. (in preparation) also describes procedures for deprojection of other differential beam systematics within each polarization pair (differential pointing, beam width, and beam ellipticity) that cause temperature-to-polarization leakage. For this analysis, we deproject only the leakage from relative-gain mismatch, because it was demonstrated in T10 to be the most significant systematic for BICEP1.

4.5. Map Results

Applying the mapmaking algorithm described above to the three-year data set yields T , Q , and U maps at 100 and 150 GHz (available in the data release), which are qualitatively similar to those in Figure 2 of C10. We compare the noise level in these maps, estimated from noise simulations, to that of the two-year maps from C10. When using the same central region, which encompasses 27% (19.2%) or 305 (203) deg^2 of the 100 (150) GHz maps, we find that the rms noise is 0.68 (0.50) $\mu\text{K-deg}$. This is consistent with the 52% increase in integration time found in Section 3.2.

An alternate, and perhaps more natural, measure of map depth is obtained by calculating the weighted standard deviation of a jackknife map (see Section 6.2), using the weight map described in Section 4.3. The jackknife map, chosen to be the season-split temporal jackknife, has the same noise level as the three-year map, but no signal. The advantages of this technique are that it is tightly related to the procedure used for power-spectrum estimation, which uses the same weighting, and that there are no tunable parameters in the procedure (the first map depth calculation requires either a choice of a particular map area or else a threshold on the integration time per pixel). By this method, we find the noise level in the BICEP1 three-year polarization maps to be 0.90 (0.73) $\mu\text{K-deg}$ over an effective area of 446 (291) deg^2 for 100 (150) GHz. The noise levels quoted by this method are higher than those listed above because

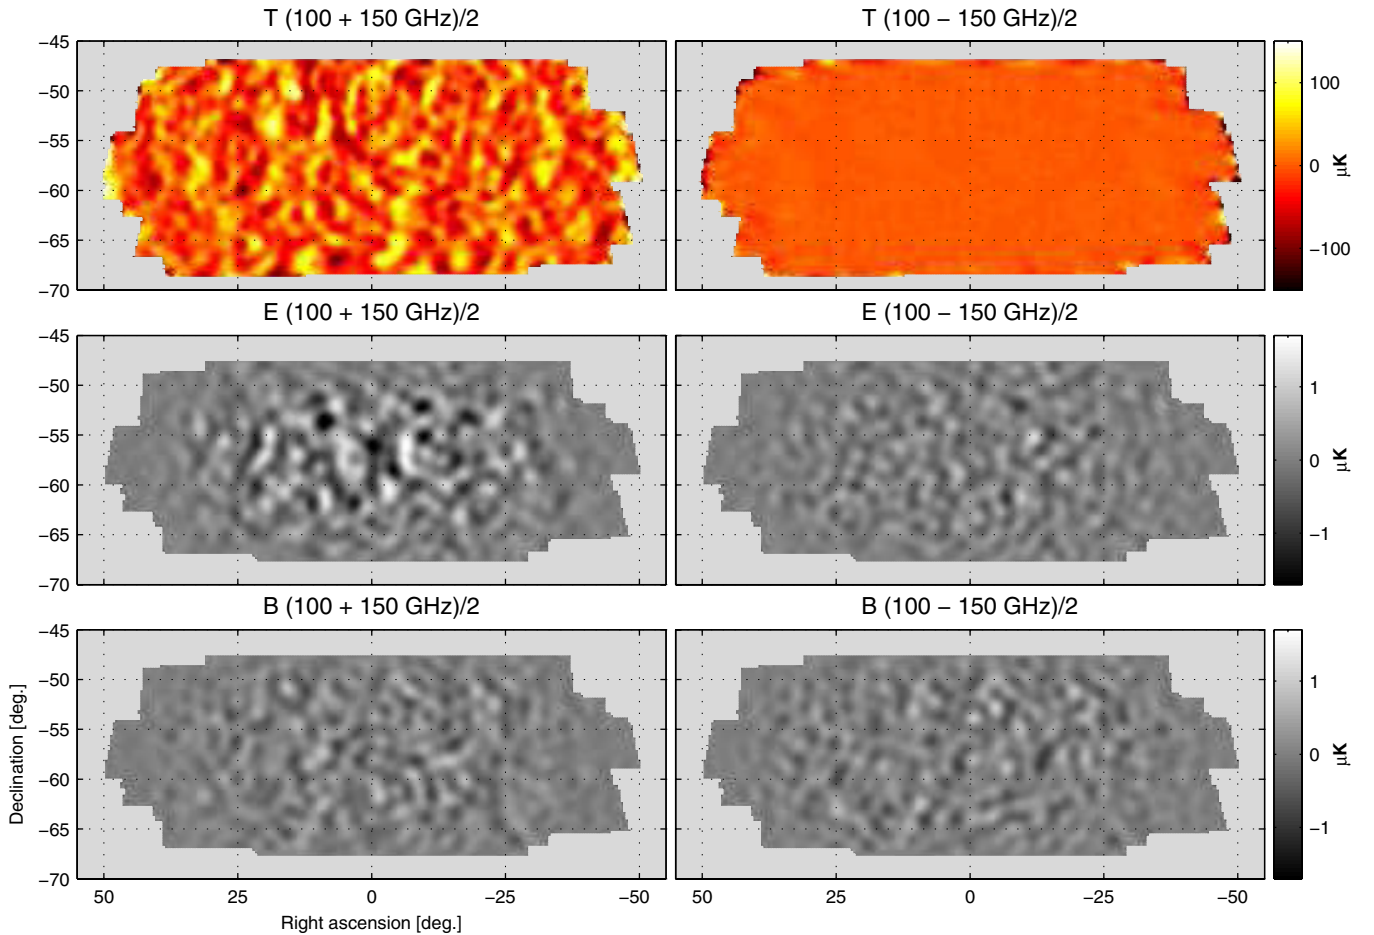


Figure 1. Maps of CMB temperature and E/B -mode polarization generated from three years of BICEP1 observations. The 100 and 150 GHz maps are each smoothed to a common 1° (FWHM) beam size before taking the frequency sum (left column) and difference (right column). The E and B polarization maps have been additionally filtered to remove power outside the range $30 < \ell < 200$, in order to emphasize the angular scales of interest for BICEP1. As in C10, BICEP1 detects E -mode polarization with high significance while the B -mode signal is consistent with noise. The E - and B -mode frequency-difference maps are consistent with noise, indicating that they do not suffer from significant contamination by polarized foregrounds.

they are determined from a larger fraction of the maps, including pixels near the edge of the field with less integration time.

The 100 and 150 GHz maps are calibrated to common units (μK_{CMB}) and combined to produce the T , E , and B maps shown in Figure 1. The left column shows frequency-sum maps and the right column frequency-difference jackknife maps. The stark contrast between the frequency-sum and frequency-difference temperature maps demonstrates that the CMB temperature anisotropy is observed at high signal-to-noise ratio. For visual clarity, the E and B maps have been filtered in Fourier space to include only power in the range $30 < \ell < 200$. The E -mode frequency-sum map shows the expected signal, while the frequency-difference map is consistent with noise, a confirmation that polarized foregrounds are not detected. Based on noise and signal simulations (Sections 5.2 and 5.3), we find that the signal-to-noise ratio in the E -mode map exceeds unity for the first four band powers (up to $\ell \sim 160$) and peaks at a value of 3.3 for the $90 < \ell < 125$ band. The B -mode frequency-sum and frequency-difference maps are both consistent with noise.

5. POWER-SPECTRUM ANALYSIS

The angular power spectra of the CMB are estimated from 2D FT of the temperature and polarization maps following the technique described in Pryke et al. (2009). Specifically, we make estimates of \mathcal{D}_b^{XY} , a binned version of $\mathcal{D}_\ell^{XY} = \ell(\ell+1)\mathcal{C}_\ell^{XY}/2\pi$.

The indices X and Y denote the two maps used to calculate a particular power spectrum, either the 100 or 150 GHz T , E , or B maps. This yields a total of 21 separate power spectra—three spectra for each of TT , EE , and BB ; four spectra for each of TE , TB , and EB .

Power measured directly from the maps, $\tilde{\mathcal{D}}_b^{XY}$, can be related to the true CMB power spectra as

$$\tilde{\mathcal{D}}_b^{XY} = F_b^{XY} \mathcal{D}_b^{XY} + \tilde{N}_b^{XY}. \quad (1)$$

The suppression factor, F_b^{XY} , accounts for power removed from the maps by filtering, including relative-gain deprojection, as well as smoothing of small-scale power due to the instrumental beam. Instrumental noise in the maps introduces an additive noise bias, \tilde{N}_b^{XY} , to the observed power spectra.

We use a simulation-based technique to derive F_b^{XY} and \tilde{N}_b^{XY} , and to ultimately solve for the CMB power spectra. Full timestream simulations of instrumental noise and/or cosmological signal are processed identically to the real data, including filtering, mapmaking, and power-spectrum estimation. We make these simulations as realistic as possible, to fully and transparently account for the effect of our analysis pipeline on the data.

5.1. From Maps to Power Spectra

The 2D FT is applied to the temperature and polarization (Stokes Q and U) maps after they have been apodized by the weight maps. Because the sky coverage is slightly different for the Q and U maps, the inverse of the mean of the Q and U variance in each pixel is used as a common apodization for those FTs.

The coordinates of points in the Fourier space maps are ℓ_x and ℓ_y , which represent modes with wave vector in the direction of right ascension and declination, respectively. The transformed Q and U maps can be rotated into maps representing the even-parity E -modes and odd-parity B -modes,

$$E(\ell_x, \ell_y) = +Q(\ell_x, \ell_y) \cos 2\phi + U(\ell_x, \ell_y) \sin 2\phi, \quad (2)$$

$$B(\ell_x, \ell_y) = -Q(\ell_x, \ell_y) \sin 2\phi + U(\ell_x, \ell_y) \cos 2\phi, \quad (3)$$

where $\phi = \arctan(\ell_y/\ell_x)$.

After transforming from Q and U to E and B , the power spectra, \tilde{D}_b^{XY} , are calculated by multiplying Fourier map X with the complex conjugate of Fourier map Y . This product is scaled by $\ell(\ell+1)/2\pi$, where $\ell = \sqrt{\ell_x^2 + \ell_y^2}$, and then averaged in annular bins. We report the BICEP1 results in nine bins of uniform width $\Delta\ell = 35$, with the first bin spanning $20 \leq \ell < 55$ and the ninth bin spanning $300 \leq \ell < 335$.

5.2. Noise Simulations

To recover the underlying true power spectra, the first step consists of subtracting the noise bias, \tilde{N}_b^{XY} . We form a noise model based on the correlations between detectors in the real data, generate noise-only simulated timestreams, and process them through the same timestream filtering, mapmaking, and power-spectrum estimation as the real data. The resulting simulated noise spectra are then averaged over many realizations to estimate the noise bias.

In C10, the noise covariance matrix was based on the correlations between filtered pair-sum and pair-difference timestreams, accumulated over all half-scans in a scanset. For this analysis, we have chosen instead to use the noise model described in Section 5.3 of Pryke et al. (2009), which calculates a noise covariance matrix for unfiltered individual detector timestreams over an uninterrupted scanset, including the scan turnarounds that are ultimately cut from both the real and simulated timestreams. This noise model more closely follows the analysis philosophy of faithfully simulating the real data and then treating the simulated and real timestreams symmetrically, including the filtering step.

Additionally, by deriving the noise model from a full scanset length timestream, instead of a large number of half-scan length segments, we capture the low-frequency atmospheric noise, which persists over many half-scans and is heavily correlated between detectors. A direct comparison of the noise model used in this analysis with that of C10 is presented in Section 6.1.2.

We generate 499 independent realizations of BICEP1 noise. Each realization consists of timestream data for all detectors across all three years of observation. These simulated timestreams pass through the low-level processing and filtering operations, are binned into maps, and are reduced to power spectra. The noise bias is calculated simply as the ensemble average of the spectra from these 499 noise-only simulations. While the noise bias is generally close to zero for cross-spectra, those terms are still calculated and subtracted.

5.3. Signal Simulations

We generate two classes of signal-only simulations, used to characterize the response of the analysis to CMB signals. The first set of signal-only simulations, hereafter referred to as “ E -no- B ,” use input CMB spectra generated by CAMB (Lewis et al. 2000) based on *WMAP* five-year best-fit cosmological parameters²⁰ (Komatsu et al. 2009). As the name implies, these theoretical power spectra have E -mode power but no B -modes. For the second set of signal-only simulations, hereafter referred to as “ B -no- E ,” we include primordial tensor perturbations corresponding to $r = 0.1$, but explicitly null the EE (and TE) power. The E -no- B simulations primarily characterize the leakage between the E and B polarization maps induced by our pipeline, while the B -no- E show our pipeline’s effect on an input BB signal.

Realizations of these signal simulations are generated from the theory spectra using the `synfast` utility included in the `HEALPix` package.²¹ At this stage, the CMB signal is smoothed by convolution with a Gaussian beam; for each simulated sky, two sets of T , Q , and U maps are produced, corresponding to BICEP1 100 GHz and 150 GHz beam widths. The resolution of the `HEALPix` maps is $0^\circ 11$ ($N_{\text{side}} = 512$). Using the pointing data from the telescope, we simulate timestreams for every detector by performing a second-order Taylor expansion interpolation from the nearest pixel center to the actual detector pointing. Signal-only timestreams are processed through the analysis pipeline, passing through the same filtering and mapmaking steps as the real data. For each of the two classes of signal simulation, we apply this process to 499 independent sky realizations drawn from the same underlying CMB power spectra to produce sets of simulated signal-only maps.

When we calculate the power spectra of the E -no- B maps, we measure non-zero BB spectra due to $E \rightarrow B$ leakage due to limited sky coverage and filtering. The ensemble average of BB spectra recovered from the E -no- B simulations is subtracted to debias the measured BB spectrum, exactly analogous to the noise bias, but for BB spectra only. The amplitude of the leakage bias corresponds to $\mathcal{D}_b^{BB} \sim 0.02 \mu\text{K}^2$ for $\ell \sim 100$. This value appears significantly larger than the $E \rightarrow B$ leakage reported in C10 because the `Spice` estimator includes an analytic debiasing operation for the sky cut effect; the level of the leakage quoted in C10 only accounts for the residual after this analytic debiasing step. For the 2D FT estimator, the $E \rightarrow B$ leakage is entirely measured from signal simulations. After debiasing, only the sample variance of the $E \rightarrow B$ leakage signal is important, as this can contribute additional uncertainty to the BB band powers.²² For this analysis, the leakage contribution to BB band power uncertainty is $4 \times 10^{-3} \mu\text{K}^2$ (see Figure 7), which is similar to the value obtained using `Spice` in C10 and subdominant to instrumental noise.

The power spectra of the E -no- B and B -no- E signal-only maps are used to correct the band power window functions and compute the suppression factors, F_b^{XY} , as described in Section 5.4. Additionally, we can combine the signal-only maps with noise-only maps to create complete simulations of the real BICEP1 data. The combination of E -no- B signal simulations plus noise is used to derive band power error bars, as described in

²⁰ The difference between *WMAP* five-year cosmology and updated parameters from *WMAP* nine-year or *Planck* is negligible for the purposes of calculating the suppression factors and signal variance.

²¹ <http://healpix.sourceforge.net>

²² The excess variance from $E \rightarrow B$ leakage can be reduced through the use of improved estimators (e.g., Smith 2006).

Section 5.5. By further addition of scaled versions of the B -no- E maps to the E -no- B signal and noise maps, we can construct map simulations of a cosmology containing inflationary B -modes at arbitrary values of r ; these simulations are used in Section 9.3 to derive a constraint on r from our data. Finally, the signal simulation infrastructure is capable of introducing a wide variety of instrumental systematics to the simulated data, such as temperature-to-polarization leakage due to mismatched relative gain or beam imperfections. These systematic uncertainties are not included in the fiducial set of 499 simulations, but they are included in alternate simulations to characterize power-spectrum systematic uncertainties, in both Section 7 of this paper and T10.

5.4. Band power Window Functions and Suppression Factors

Band power window functions (Knox 1999) are used to relate theoretical input spectra, \mathcal{D}_ℓ^{XY} , to expectation values for the band powers measured by BICEP1,

$$\langle \mathcal{D}_b^{XY} \rangle = \sum_\ell w_{b,\ell}^{XY} \mathcal{D}_\ell^{XY}. \quad (4)$$

The window functions, $w_{b,\ell}^{XY}$, are defined to have unit integral over ℓ . The shape of the window function is primarily determined by the apodization mask, which mixes power between angular scales. In Fourier space, this effect can be understood as a convolution of the sharp-edged annulus used by the power-spectrum estimator with a smoothing kernel given by the FT of the mask. The window functions are also modified by timestream filtering and smoothing of the sky signal by the beam and the map pixelization.²³ This is most significant for the first bin, as the timestream filtering preferentially removes power from the largest angular scales in that range.

To calculate the band power window functions, we start by calculating the window function corresponding only to the apodization mask, $m_{b,\ell}^X$, following the procedure described in Challinor & Chon (2005). To better capture variations in the suppression factor, we calculate this “mask window function” for ℓ -bins, b' , that are smaller than the usual bins by a factor of four (annuli in the Fourier plane with width $\Delta\ell = 8.75$). Additionally, we define these bins across a much wider range of angular scales, from the origin of the Fourier plane out to $\ell \sim 500$, to capture the leakage of power between angular scales. The mask window function and signal-only simulations are used to make a preliminary estimate of the suppression factor for the fine angular bins,

$$F_{b'}^{XY} = \frac{\langle \tilde{\mathcal{D}}_{b'}^{XY} \rangle}{\sum_\ell m_{b',\ell}^{XY} \mathcal{D}_\ell^{XY}}. \quad (5)$$

Here, the expectation value in the numerator is an ensemble average over power spectra calculated in fine angular bins from the simulated signal-only maps, while the term in the denominator comes from applying the mask window function to the input spectrum used to generate those simulations. The suppression factor describes how the telescope and analysis pipeline remove power at each angular scale through beam smoothing and timestream filtering. The E -no- B simulations are used to make this calculation for the TT and EE suppression factors. The BB suppression factors are calculated using




Figure 2. Shown in the top panel is the suppression factor, $F_b^{150E \times 150E}$, including contributions from beam smoothing and our filtering choices. At large angular scales, power suppression is dominated by our conservative choice of polynomial and azimuth-fixed filters, shown both individually (plus sign and downward triangles) and combined (squares). At small angular scales, beam smoothing (dotted line; calculated analytically) dominates. The total suppression factor (circles) also includes small contributions from relative-gain deprojection and the pixel window function. The suppression factor should not be mistaken for a measure of the low ℓ performance of the experiment (see text). The bottom two panels show the actual impact of our filtering choices on the signal-to-noise ratio of EE and BB band powers. The total sensitivity loss from polynomial plus azimuth-fixed filtering is shown by squares, with the gray shaded region indicating the 1σ uncertainty on this calculation. The plus signs and downward triangles indicate the loss of sensitivity from polynomial or azimuth-fixed filtering individually. While the two filters suppress similar amounts of power in the maps (top panel), only the azimuth-fixed filter has a significant effect on signal-to-noise ratio (bottom panels).

B -no- E simulations. For the TE , TB , and EB cross-spectra, we use the geometric mean of the suppression factors for the two corresponding auto-spectra (e.g., the TE suppression factor is the geometric mean of the TT and EE suppression factors).

Next, the band power window functions are corrected by multiplying the mask window functions with a smoothly interpolated version of the suppression factor. After renormalizing the band power window functions, this procedure can be iterated, with the more accurate band power window functions substituted in place of the initial mask window functions in Equation (5). In practice, this procedure converges very quickly; we perform three iterations, but essentially all of the modification to the window functions occurs in the first iteration.

At the end of this procedure we have final suppression factors and corresponding corrected band power window functions. The finely binned band power window functions and suppression factors are then merged to recover equivalent functions appropriate for the nine $\Delta\ell = 35$ bins used to report BICEP1 results. The top panel of Figure 2 shows the suppression factor for the EE spectrum, as well as individual contributions from polynomial and azimuth-fixed filtering, relative-gain deprojection, and beam smoothing. Figure 3 shows the band power window

²³ The contribution of pixelation to the window functions and suppression factor is very small for our maps, so we will ignore it for the following discussion, although it is accounted for in our analysis.




Figure 3. BICEP1 band power window functions for the 150 GHz BB auto-spectrum. The gray dashed line shows the mask window function for the first ℓ -bin. Band power window functions for other spectra are visually similar.

functions for BB only.²⁴ For the first ℓ -bin, we also plot the mask window function, to highlight the change in that bin due to filtering.

We emphasize that the power suppression due to filtering is not a measure of the loss of information; signal and noise are both suppressed by filtering. The suppression factor is also not an indicator of the noise spectrum; for our pipeline the filtering choices are set by hand and have been chosen conservatively (see below). The suppression factor is not a meaningful measure of low ℓ performance.

To measure the true impact of filtering, we run simulations with the filtering relaxed by either reducing the order of the polynomial filter from three to one, turning off the azimuth-fixed filter, or both. The fractional change in signal-to-noise ratio (bottom panels of Figure 2) is calculated relative to a “minimal filtering” analysis that uses a first-order polynomial filter and no azimuth-fixed filtering. This statistic depends on both the signal-to-noise regime and the assumed shape of the signal spectra. Uncertainty in the statistic, due to the finite number of realizations, is estimated based on the variance between subsets of the realizations. For EE , the combination of polynomial and azimuth-fixed filtering reduces the first bin sensitivity by 15%. For BB , the first bin sensitivity is reduced by 35%. This factor is larger for BB than EE because the BB spectrum is noise dominated and because the EE spectrum rises steeply across the first bin ℓ range, so loss of information from the lowest multipoles is less important.

Comparing the loss of sensitivity for the combined polynomial and azimuth-fixed filtering (squares) with the similar factors for polynomial (plus signs) or azimuth-fixed (triangles) filtering only, we find that nearly all of the information loss in the first bin is due to the azimuth-fixed filtering. The azimuth-fixed filtering targets ground-fixed or scan-synchronous contamination but also attenuates long-wavelength sky signals. The scale at which the azimuth-fixed filter affects band power sensitivity is a direct consequence of the choice of timescale (one scanset, ~ 50 minutes) used to construct the azimuth-fixed template. Polynomial filtering affects signal and noise more equally. Third-order filtering was chosen based on our temperature data. As seen in Figure 2 (top panel, triangles), relaxing the polynomial filtering to first order for our polarization data would change the suppression factor dramatically at low ℓ , but brings no significant benefit in signal-to-noise ratio (bottom panels), so for simplicity we retain the same filtering choices for both temperature and polarization data.

5.5. Power-spectrum Results

After calculating \tilde{N}_b^{XY} and F_b^{XY} , we can solve for the underlying CMB power spectra, \mathcal{D}_b^{XY} , using Equation (1). As discussed in Section 5.3, the average BB signal from E -no- B signal simulations is subtracted from the BB band powers along with the noise bias.

By combining the noise-only maps with the E -no- B signal-only maps, we create a set of 499 signal-plus-noise maps; each of these maps represents a full simulation of the three years BICEP1 observations with independent noise and signal realizations. We calculate the power spectra of the signal-plus-noise maps and process them identically to the real data, subtracting the noise bias and dividing by the suppression factor. This set of simulated power spectra is used to directly determine the covariance matrix of the band powers. All 21 BICEP1 power spectra are shown in Figure 4, with error bars given by the square root of the diagonal elements of the band power covariance matrix.

The full statistical power of BICEP1 is realized by combining the frequency auto- and cross-spectra. For each bin of each spectrum, we take a weighted average between the corresponding band powers from three (or four) different frequency combinations. Weights are calculated as

$$w_i = \sum_j \mathcal{M}_{ij}^{-1}, \quad (6)$$

where \mathcal{M}_{ij} is the appropriate 3×3 (or 4×4) block of the band power covariance matrix. After performing the same frequency combination on the simulated spectra, we can calculate a new band power covariance matrix for the frequency-combined spectra. The frequency combined power spectra are shown in Figures 5 and 9. For the chosen bin width, $\Delta\ell = 35$, band powers in adjacent ℓ -bins have 8%–10% positive correlations. Elements of the covariance matrix that are separated by more than two ℓ -bins are too small to be well measured from our finite sample of simulations, so we set them to zero.

6. CONSISTENCY TESTS

During the analysis that led to C10, two different analysis pipelines worked in parallel and compared results systematically up to the final result on the tensor-to-scalar ratio, r . Only one of these pipelines was used for the three years BICEP1 analysis, so this side-by-side confirmation is not available. We have therefore tested that the results are insensitive to both the algorithmic and the data-selection changes between this result and those presented in C10.

In addition, we subject all the spectra to the same set of jackknife null tests as in C10 to probe for systematic contamination. The TE and TB jackknives merit special discussion to demonstrate that the non-jackknife spectra are valid despite some formal null-test failures.

6.1. Consistency with Two-year Results

Here, we review all the algorithmic and data-selection differences between this analysis and that of C10 and present the effect of these changes on the polarization band powers.

6.1.1. Consistency between Analysis Pipelines

As described in Section 4, the two separate pipelines used in C10 differ mainly in map format (HEALPix versus equirectangular map pixelization) and power-spectrum estimator (Spice

²⁴ The other window functions are calculated and included in the BICEP1 data release, but are not included in Figure 3.

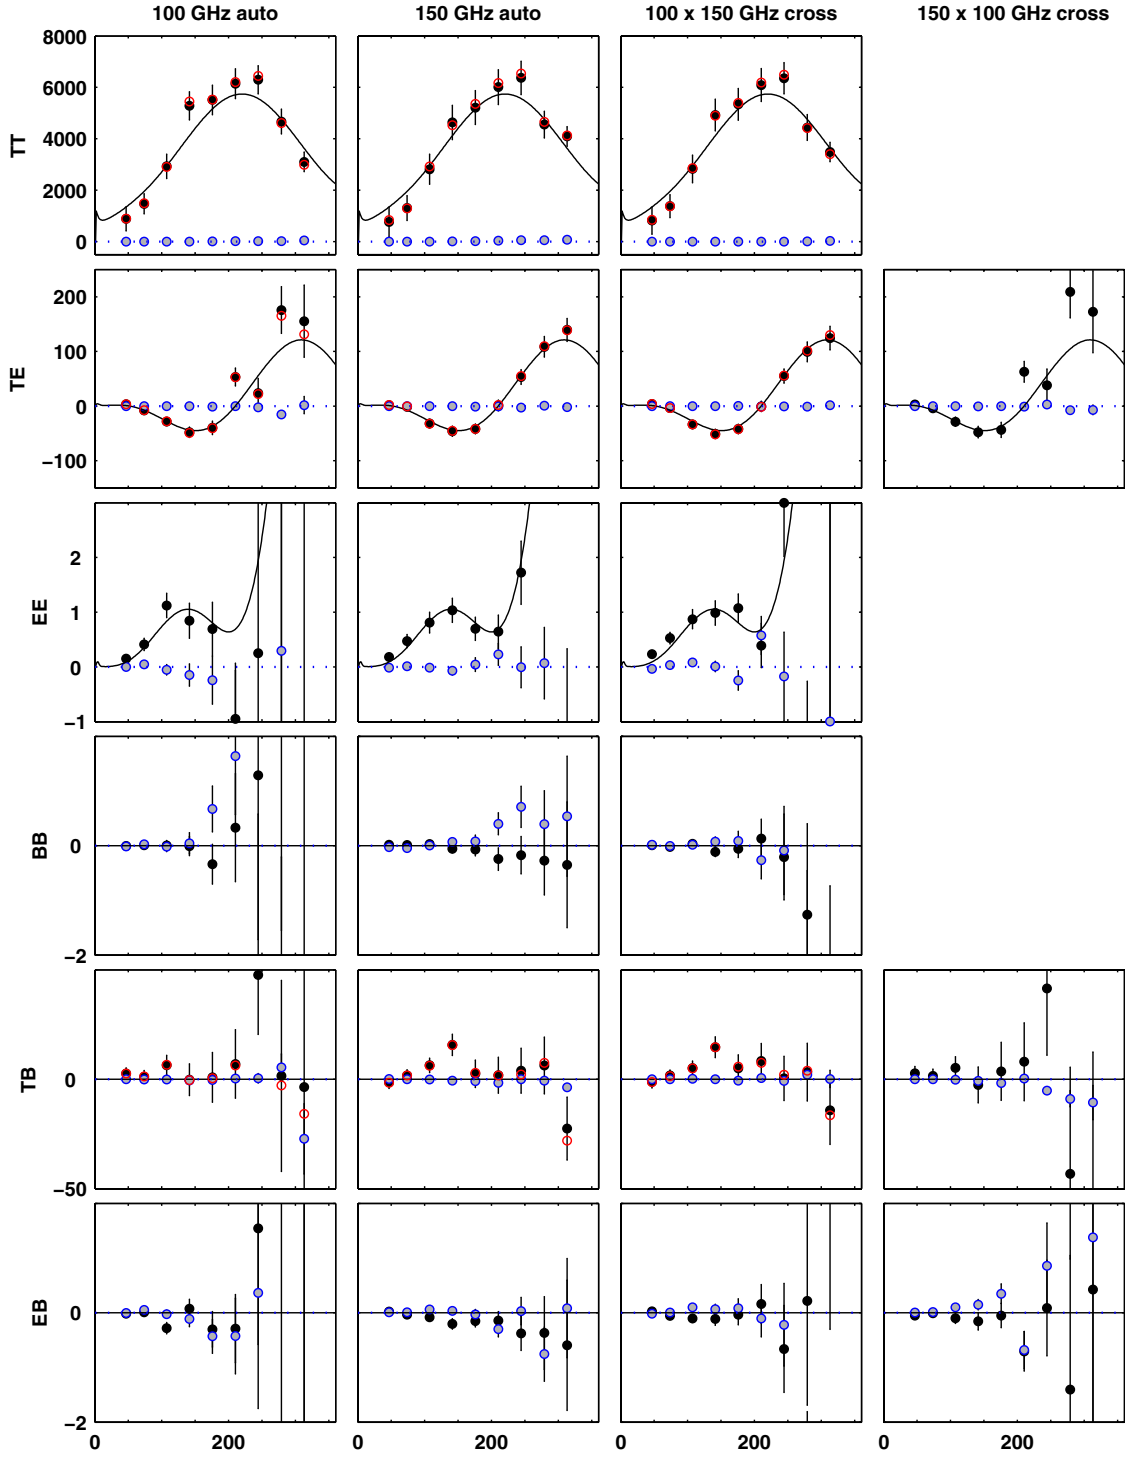


Figure 4. BICEP1 individual-frequency CMB power spectra. The horizontal axis is multipole ℓ , and the vertical axis is $\ell(\ell + 1)C_\ell/2\pi$ in units of μK_{CMB} . Black points show the full set of BICEP1 power spectra up to $\ell = 350$ with statistical error bars (including sample variance) only. The spectra agree well with a Λ CDM model (black lines) derived from *WMAP* five-year data and $r = 0$. The blue points correspond to the boresight-angle jackknife. The red open circles show the TT , TE , and TB spectra calculated using one half of the BICEP1 boresight-angle jackknife maps as the temperature map, as described in Section 6.2.

versus 2D FT), so we expect minor differences in their respective maps and derived power spectra. In C10 we have carefully cross-checked those differences and found that they are negligible compared to the statistical uncertainty of the spectra. Specifically, focusing on the frequency-combined BB spectrum (see Figure 11 in C10), we found that the band powers from the two pipelines used in C10 agree to better than 1σ with better than 0.2σ agreement in the first four bins.

6.1.2. Noise Model Consistency

As discussed in Section 5.2, this analysis uses a more sophisticated noise model to better capture the details of instrumental and atmospheric noise. The noise bias calculated from simulations is subtracted from the band powers, so noise misestimation leads directly to bias in the power spectra. To assess the impact of this change, we compare sets of 100 noise




Figure 5. BICEP1 frequency-combined power spectra (black points) are in excellent agreement with a Λ CDM model (black lines) derived from *WMAP* five-year data. The χ^2 (for 9 dof) and PTE values from a comparison of the data with the model are listed in the plots. Gray crosses denote band power expectation values for the model. Power-spectrum results from C10 are shown by open circles and are offset in ℓ for clarity. In both cases, the error bars are the square root of diagonal elements in the frequency-combined band power covariance matrix described in Section 5.5 and do not include systematic uncertainties.

realizations that differ only in the choice of noise model: either the model described in Section 5.2 or the one discussed in Section 6.1 of C10. We process these two sets as described in Section 5 and calculate the noise bias for each model, focusing on the *BB* spectrum for which a change in noise bias would have the most significant impact. We find that the two calculations differ by at most $0.03 \mu\text{K}^2$, with a largest fractional shift of 6%. Moreover, the noise power does not increase or decrease uniformly across ℓ -bins, so the small difference is averaged down further when constraining r . Using the C10 noise model would change the observed value of r by just 0.03. As a test of the accuracy of band power error bars derived from noise simulations, we recalculate the suite of jackknife null tests described in Section 6.2 and find that they are indistinguishable when we use either set of noise simulations.

6.1.3. Band power Window Function Consistency

The band power window functions reported in C10 are derived from the *Spice* kernel, making it difficult to provide a direct comparison to the window functions of the current analysis. However, the procedure described in Section 5.4, which modifies the window functions to account for the effects of filtering and beam smoothing, was not applied in C10. This change is significant for the first ℓ -bin only. For that bin, the

timestream filtering shifts the window function to higher ℓ and leads to a smaller value for the suppression factor (i.e., increased suppression of power). The instrumental noise contribution to the band powers is scaled by the inverse of the suppression factor, as can be seen from Equation (1), so the net result is a larger error bar for the first bin band powers. If we compare the error bars obtained for the *BB* spectrum using the mask window function in place of the band power window function, we find that they are underestimated by 40% for the first ℓ -bin only. This test is only applicable for the comparison between window function treatments for the 2D FT power-spectrum estimator; a comparison to the error bars reported in C10 includes a different power-spectrum estimator, which leads to an entirely separate change in the error bar estimate.

Tests of consistency between observed and simulated band powers, for either jackknife tests or comparison with Λ CDM cosmology, are independent of the suppression factor and are not affected by this change.

6.1.4. Relative-gain Deprojection

This analysis incorporates a deprojection technique to remove any temperature leakage due to relative-gain mismatch (see Sections 4.4 and 7.2 for details). We perform two complete analyses with and without this deprojection and find that

Table 2
Jackknife PTE Values from Polarization-only χ^2 Tests

Jackknife	100 GHz	150 GHz	100×150	150×100
Scan direction				
<i>EE</i>	0.756	0.124	0.575	
<i>BB</i>	0.244	0.246	0.327	
<i>EB</i>	0.679	0.804	0.148	0.391
Elevation coverage				
<i>EE</i>	0.341	0.471	0.581	
<i>BB</i>	0.106	0.581	0.319	
<i>EB</i>	0.335	0.639	0.273	0.764
Boresight angle				
<i>EE</i>	0.733	0.952	0.192	
<i>BB</i>	0.493	0.257	0.836	
<i>EB</i>	0.489	0.251	0.104	0.026
Season split				
<i>EE</i>	0.495	0.156	0.804	
<i>BB</i>	0.230	0.042	0.525	
<i>EB</i>	0.471	0.421	0.918	0.898
Focal plane <i>QU</i>				
<i>EE</i>	0.986	0.411	0.383	
<i>BB</i>	0.287	0.834	0.451	
<i>EB</i>	0.279	0.244	0.784	0.541

the final *BB* band powers shift by less than 0.1σ except for the bin at $\ell = 107.5$, which shifts down by 0.6σ . Although this change is small, the application of relative-gain deprojection significantly reduces this source of systematic uncertainty.

6.1.5. Consistency with Inclusion of Additional Data

In this analysis, the inclusion of the third season of data and previously discarded detectors has increased the total integration time by 52% over C10. We include the new data incrementally and use sets of 100 signal-plus-noise realizations to determine whether the resulting shifts in band powers are statistically significant. We find that re-including the set of detectors with abnormal transfer functions produces negligible shifts in all band powers, much smaller than the difference between the two pipelines, for example. Including the third season of data produces shifts that are consistent with the expectation from simulations. As additional confirmation that the third season is consistent with the first two, we calculate a modified version of the season-split temporal jackknife (see Section 6.2) where the two jackknife halves are the 2006+2007 seasons and 2008 season. This jackknife test passes, though we do not include it in our standard suite.

Examining the ratio of the error bars from two simulation sets that differ only by the additional data, we confirm that the sensitivity improvement for the *BB* spectra is purely proportional to the increase in integration time, matching the expectation for the noise-dominated case.

6.2. Jackknife Null Tests

We perform “jackknife” tests to verify that the power spectra are free of systematic contamination. These jackknives are statistical tests in which the data are split into two halves, processed to form maps, and then the maps are differenced. The power spectra of the differenced, or jackknife, map are tested for consistency with zero power, to within the uncertainty derived from simulations. For this analysis, we adopt the same five jackknife splits that were tested in C10: scan direction, elevation coverage, boresight rotation, season-split temporal, and focal plane *QU* jackknife. They are designed specifically to




Figure 6. Probabilities-to-exceed from χ^2 tests of 50 polarization-only (*EE*, *BB*, and *EB*) jackknives are consistent with a uniform distribution between zero and one (dashed line).

probe for instrumental systematic effects. We did not perform the eight-day temporal jackknife, and the season-split temporal jackknife has been updated to split the three seasons into two even halves (except for the test noted in Section 6.1.5). For brevity, we refer the reader to the description of these jackknife tests in Section 8.1 of C10.

To test the jackknife spectra, we evaluate the χ^2 goodness of fit, with 9 degrees of freedom (dof), to the null hypothesis. Due to differences in filtering and sky coverage between the two halves of a jackknife map, we expect small levels of residual signal. We account for this by evaluating the probability-to-exceed (PTE) for the real data against the distribution of χ^2 values from simulations, which should contain the same residuals, rather than using a theoretical χ^2 distribution. We evaluate the results of the jackknife tests by the following measures:

1. The fraction of jackknife spectra with PTE smaller than 5% should not be significantly larger than 5%.
2. None of jackknife spectra should have a PTE that is excessively small ($\ll 1\%$).
3. The PTE from all jackknives should be uniformly distributed between zero and one.

The PTEs from jackknife χ^2 tests of all frequency combinations for the polarization-only (*EE*, *BB*, *EB*) spectra are presented in Table 2. Only two of the tests (out of 50) have PTE values less than 5%, and neither one is exceedingly low (2.6% for the 150×100 *EB* spectrum from the boresight angle jackknife and 4.2% for the 150 GHz *BB* season-split jackknife). Figure 6 shows that the histogram of the PTE values follows the expected uniform distribution. The boresight rotation jackknife spectra are plotted in Figure 4 (blue points) for comparison with the non-jackknife spectra.

In addition to the χ^2 tests, we compare the jackknife band power deviations, defined as the ratio of band power values to their error bars, against the simulations. This provides a strong test for a coherent bias in the band powers, which could be caused by misestimation of the noise bias. For all the polarization jackknife spectra, the sum of band power deviations is found to be consistent with the simulated distributions.

The jackknife PTE values for those power spectra including the *T* map (*TT*, *TE*, and *TB*) are not shown in Table 2; they display significant failures. The PTE distributions for *TE* and *TB* have an excess of values between 0.05 and 1×10^{-5} , and most of the *TT* jackknives have extremely small PTE values. We hypothesize that the *TE* and *TB* failures are caused by imperfect signal cancelation in the temperature jackknife maps, as opposed to the polarization maps.

Working from this hypothesis, we have built estimates of TE and TB jackknife contamination derived from the observed TT jackknife failures. These estimates are indeed consistent with the observed TE and TB jackknife band powers.

Additionally, we explore a special type of TE or TB jackknife, referred to as “half-jackknives,” where we calculate the cross-spectrum between the jackknife polarization (E or B) map and a full, i.e., non-jackknife, temperature map. In these half-jackknives, an inconsistency in the E or B map would still show up as a jackknife failure in excess of the simulations, while an inconsistency driven by the temperature map would disappear. All the PTE for the 20 TE and TB half-jackknife tests exceed 5%, confirming again that the temperature maps are the likely source of the TE and TB jackknife failures.

As a final check, we confirm that the formal failures in the TE and TB jackknife tests reflect contamination that is at a negligible level compared to the noise level of those non-jackknife spectra. In C10, this was done by comparing the BICEP1-only TE and TB spectra to spectra constructed using the WMAP temperature map. Here, we perform an equivalent test by calculating alternate TT , TE , and TB spectra using just one half of any of the temperature jackknife maps. If the jackknife failures are indicating contamination at a level that significantly affects the BICEP1 power spectra, we would see a large difference between power spectra calculated from the discrepant halves of the jackknife map. These alternate spectra are shown for the case of the boresight angle jackknife as red open circles in Figure 4. The difference between the two is consistently less than 5% of the error bars for the TT , TE , or TB non-jackknifed spectra.

7. SYSTEMATIC UNCERTAINTIES

For a comprehensive study of systematic uncertainties, we refer the reader to T10. Here we address only the dominant sources of systematic uncertainties as identified in T10 and C10, specifically updating the estimates of relative gain and polarization orientation uncertainties. Figure 7 summarizes the statistical and systematic uncertainty estimates for BICEP1 power spectra.

7.1. Absolute Calibration and Beam Width Uncertainty

We follow the same procedure as in C10 to derive the absolute gains for the three-year maps. Given the similarity of the results, the value and uncertainties of the absolute calibration and beam width remain unchanged from C10. These systematic uncertainties are multiplicative in the observed power spectra; they can lead to incorrect scaling of an observed signal but will not produce false B -modes. The sum of the absolute calibration and beam width systematics is plotted in Figure 7. For all polarization spectra, these uncertainties are a small fraction of the total statistical uncertainty.

7.2. Relative-gain Mismatch

For the C10 analysis, relative-gain mismatch caused by imperfectly differenced detector pairs was judged the leading source of possible BB contamination, estimated to produce a bias on $r \leq 0.17$ (T10). Although this exceeded the $r = 0.1$ benchmark set, it remained small compared to the statistical uncertainty of the two-year result. With the addition of a third year of data and additional detectors, we are motivated to use the relative-gain deprojection technique described in Section 4.4 (with further details in R. Aikin et al. in preparation) to keep the

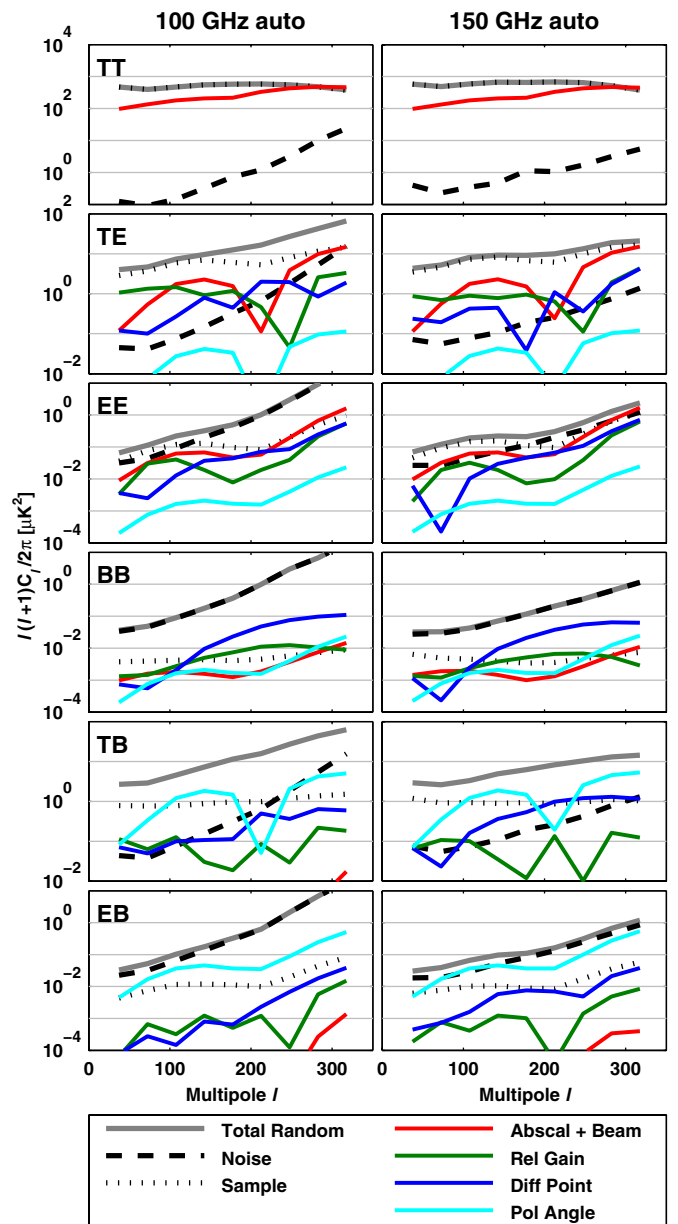


Figure 7. Summary of statistical and systematic uncertainties for the BICEP1 three-year result. Total random (solid gray line) indicates the overall statistical uncertainty due to both instrumental noise (black dashed) and CMB sample variance (black dotted). Systematic uncertainty contributions from absolute calibration and beam width (red line, Section 7.1), relative-gain mismatch (green line, Section 7.2), differential pointing (blue line, Section 7.3), and polarization angle calibration (cyan line, Section 7.4) are also included. Sample variance dominates the TT and TE spectra. The EE (and TB) spectra are dominated by sample variance at low ℓ and noise at high ℓ . Noise dominates the BB (and EB) uncertainty at all angular scales. Polarization angle calibration is a significant source of uncertainty for both TB and EB .

potential contamination from relative-gain leakage well below the statistical uncertainty.

To quantify the remaining leakage after deprojection, we compare three sets of simulated data: the first two include relative-gain mismatch at the level observed for BICEP1 and either do or do not employ deprojection to remove the resulting temperature-to-polarization leakage. The third set of simulations features perfectly matched relative gains and does not use relative-gain deprojection.

Comparing the simulations with leakage to the idealized case, we find that the excess EE and BB band power is greatly reduced by the deprojection technique. Figure 7 includes a residual leakage component that is at least an order of magnitude below the statistical uncertainty for BB and corresponds to a potential bias on r of less than 0.03.

7.3. Beam Mismatch

As in C10, differential pointing, an offset in the beam centers for paired detectors, is the leading systematic from beam mismatch. Other beam mismatch terms, differential beam ellipticity and differential beam width, are far below the $r = 0.1$ benchmark established in T10. Using simulations that contain the observed BICEP1 differential pointing offsets, we measure the resulting excess band power shown in Figure 7 and confirm that the possible bias from the differential pointing is less than $r = 0.02$. R. Aikin et al. (in preparation) include a description of deprojection techniques that can be applied to correct differential pointing, ellipticity, and beam width, but we find these corrections to be unnecessary for BICEP1.

7.4. Polarization Orientation

An error in the orientation of detectors can lead to rotation of E - into B -modes. The BB spectrum is affected only at second order, but the TB and EB spectra are more sensitive to such an effect. In Kaufman et al. (2013) we revisit the systematic uncertainty on our standard dielectric-sheet-based polarization orientation calibration by comparing it to three alternative calibrations, increasing this uncertainty from the previously reported 0.7 (T10) to 1.3. Figure 7 shows the systematic uncertainty on band powers due to the updated calibration error.

The TB and EB spectra can also be used to “self-calibrate” the polarization angle from the CMB itself (Keating et al. 2013), also as described in Kaufman et al. (2013). Applying each of these four alternative calibrations produces small shifts in the BB spectra. The maximum shift among these cases in the estimate of r is less than 0.04. The “self-calibrated” case produces a 95% upper limit $r < 0.65$ (versus $r < 0.70$). We consider these shifts small for all results based on the BB spectra and therefore keep the original estimate of polarization orientation angle for the main result of this paper.

8. FOREGROUNDS

In C10, we estimated the level of foreground contamination in BICEP1 maps and found it to be negligible. Here we update those estimates for the three-year analysis and with recently available foreground models. Given the modest improvement in sensitivity, we expect polarized foregrounds to remain undetected. Nevertheless, we provide new upper limits on possible contamination from Galactic diffuse emission and compact sources. As further evidence that our spectra are free of significant contamination, we present a 100–150 GHz frequency jackknife.

To estimate the effects of Galactic diffuse emission, we use simulated Planck Sky Model (Delabrouille et al. 2013, PSM v. 1.7.7) polarization maps including thermal dust and synchrotron emission. We process these maps through the BICEP1 pipeline to estimate the contamination in our field. The result is more than an order of magnitude smaller than our upper limits on the BB spectrum, even in the worst case of dust contamination at 150 GHz (Figure 8).

We also test for foreground contamination by cross-correlating the BICEP1 polarization maps with various




Figure 8. Estimated Galactic dust and synchrotron emission in the BICEP1 field is well below current BB upper limits. These foreground emission estimates come from processing the Planck Sky model foreground maps (Delabrouille et al. 2013) through the BICEP1 pipeline. The BB upper limits are derived from the band power likelihoods calculated in Section 9.1.

foreground templates. We compute the cross power spectra between BICEP1 and the O’Dea et al. (2012) dust models, FDS Model 8²⁵ (Finkbeiner et al. 1999), and the PSM dust and synchrotron predictions. We find no statistically significant correlations between the BICEP1 maps and any of the foreground models.

We search for compact “point” sources using external catalogs. We use the same method as C10, but with an updated list of catalogs including *Planck* (Planck Collaboration et al. 2013c), *WMAP* (Bennett et al. 2013), and ACBAR (Reichardt et al. 2009). We find at most three (two) catalog source locations with excess power between 2σ and 3σ in the BICEP1 100 (150) GHz map. However, we find a similar number of 2σ to 3σ spots in signal-plus-noise simulation maps with no point sources. Therefore, we conclude that compact sources are not significant compared to the BICEP1 noise level.

As a generic test for foreground contamination we perform a 100–150 GHz frequency jackknife. The χ^2 PTEs for the EE , BB , and EB spectra are 25.3%, 99.9%, and 84.2%, respectively. The PTEs for TE and TB are 37.1% and 65.9%, respectively. We conclude that the 100 and 150 GHz maps are consistent and that there is no evidence for foregrounds.

9. RESULTS AND DATA PRODUCTS

This section describes the main results on CMB band powers, likelihoods, and the tensor-to-scalar ratio, r . All the data products described here are available online at http://bicep.rc.fas.harvard.edu/bicep1_3yr/.

9.1. Band power Likelihood Functions

For BICEP1 the effective number of dof per ℓ -bin is small, so the non-Gaussianity of the likelihood is significant, especially in the lowest ℓ -bins. Therefore, we do not use a Gaussian likelihood approximation, but calculate band power likelihoods from our power spectra using the likelihood approximation of Hamimeche & Lewis (2008). The use of this approximation is a change from C10, which used the then-standard offset-lognormal approximation (Bond et al. 2000). In the course of this work we compared likelihoods on r derived using the offset-lognormal distribution with those derived from a new

²⁵ For the FDS model, the polarization maps are constructed by assuming 5% fractional polarization split evenly between Q and U .

direct likelihood (Section 9.3.1), and we found that the offset-lognormal likelihoods resulted in biases at the $r \sim 0.1$ level. The r estimates derived using offset-lognormal likelihoods also have significant variance among realizations from the same underlying measured value of r , compared to the other two methods. The Hamimeche & Lewis (2008) approximation greatly reduced such bias and scatter in constraints on r , and we therefore adopt it as a better approximation to the true band power likelihood functions.

The likelihood approximation is

$$-2 \log \mathcal{L}(\mathcal{D}_b | \hat{\mathcal{D}}_b) = X_c \mathcal{M}_{cc'}^{-1} X_{c'}, \quad (7)$$

where \mathcal{D}_b are the model band powers and $\hat{\mathcal{D}}_b$ are the data. In Equations (7) and (11), indices c and c' run over all 54 combinations of ℓ -bins and the six spectra (TT , EE , BB , TE , EB , TB). Index b runs only over ℓ -bins. The expression for the log-likelihood is similar to a χ^2 statistic, but calculated using X_c , a vector of band powers that have undergone a transformation to correct the shape of the likelihood.

$$\begin{pmatrix} X_b^{TT} \\ X_b^{EE} \\ X_b^{BB} \\ X_b^{TE} \\ X_b^{EB} \\ X_b^{TB} \end{pmatrix} = \text{vecp}((\mathcal{D}_b^f)^{1/2} \mathbf{U}_b g(\mathbf{D}_b) \mathbf{U}_b^\dagger (\mathcal{D}_b^f)^{1/2}), \quad (8)$$

where \mathcal{D}_b^f are fiducial band powers from the mean of Λ CDM signal-plus-noise simulations. The role of the fiducial model is to incorporate the band power covariance; \mathcal{D}_b^f and $\mathcal{M}_{cc'}$ are calculated from the same simulations. The function vecp gives the vector of unique elements in a symmetric matrix, and \mathcal{D}_b , $\hat{\mathcal{D}}_b$, and \mathcal{D}_b^f are symmetric matrices constructed from the band powers at each ℓ -bin, b . For example,

$$\mathcal{D}_b = \begin{pmatrix} \mathcal{D}_b^{TT} & \mathcal{D}_b^{TE} & \mathcal{D}_b^{TB} \\ \mathcal{D}_b^{TE} & \mathcal{D}_b^{EE} & \mathcal{D}_b^{EB} \\ \mathcal{D}_b^{TB} & \mathcal{D}_b^{EB} & \mathcal{D}_b^{BB} \end{pmatrix}. \quad (9)$$

The band powers used for this approximation are not debiased for noise or $E \rightarrow B$ leakage. Matrices \mathbf{U}_b and \mathbf{D}_b are the eigenvectors and eigenvalues of the matrix product $\mathcal{D}_b^{-1/2} \hat{\mathcal{D}}_b \mathcal{D}_b^{-1/2}$. The function

$$g(x) = \text{sign}(x - 1) \sqrt{2(x - \ln x - 1)} \quad (10)$$

is applied to the diagonal matrix \mathbf{D}_b to form $g(\mathbf{D}_b)$.

The band power covariance matrix used in Equation (7) is related to the band power covariance matrix calculated from signal-plus-noise simulations (see Section 5.5), $\mathcal{M}_{cc'}$, by

$$\mathcal{M}_{cc'} = \mathcal{M}_{cc'} + G_c G_{c'} \hat{\mathcal{D}}_c \hat{\mathcal{D}}_{c'} + S_c S_{c'} \hat{\mathcal{D}}_c \hat{\mathcal{D}}_{c'}. \quad (11)$$

The additional terms account for systematic uncertainty from absolute gain (G_c) and beam width (S_c) calibration. Incorporating the systematic uncertainty in this way is an approximation to the likelihood obtained by introducing a systematic uncertainty nuisance parameter and marginalizing over it. Detailed checks of this approximation will be in an upcoming paper on likelihood methods. As in C10, we use only the terms of $\mathcal{M}_{cc'}$ that are two or fewer ℓ -bins apart.




Figure 9. Close-up of the EE and BB spectra from Figure 5. BICEP1 measures EE polarization (open circles) with high signal-to-noise ratio at degree angular scales (Section 9.2). The BB spectrum (black points) is consistent with zero. Theoretical BB and EE spectra with $r = 0.1$ are shown in solid and dashed gray lines, respectively. The gray crosses are the band power expectation value for the EE spectrum. They diverge from the Λ CDM curve because of the detailed shape of the band power window functions. The inset shows the low- ℓ region in more detail.

9.2. Consistency with Λ CDM

To assess the consistency of our results with the Λ CDM model, we use the band power likelihood described above to create a likelihood-based consistency test. For each frequency-combined auto-spectrum (TT , EE , and BB) we calculate $\chi^2 \equiv -2 \ln \mathcal{L}$ for the theory spectrum used for the E -no- B signal simulations.²⁶ For cross-spectra, TE , TB , and EB , the Hamimeche and Lewis (HL) likelihood model only allows us to calculate a total χ^2 including the related auto-spectra in the likelihood; we then subtract the auto-only χ^2 to get the final statistic. For example, $\chi_{TE}^2 \equiv \chi_{TE+TT+EE}^2 - \chi_{TT}^2 - \chi_{EE}^2$. For each spectrum χ^2 , we compute the PTE as the fraction of signal-plus-noise simulations having larger χ^2 than the real data. We list the χ^2 and PTE for each spectrum in Figure 5; these values show no inconsistency with Λ CDM.

We also use this likelihood approximation to calculate the significance of our detection of E -mode power (Figure 9). Using only the ℓ -bins around the first peak of the EE spectrum ($56 \leq \ell \leq 195$), we calculate the χ^2 for a model with zero power to be 241. This corresponds to a 15σ detection of power in the region of the first peak. Our EE detection significance using all nine ℓ -bins is 18σ . Our TE detection significance is 14σ for all ℓ and 8σ for the $56 \leq \ell \leq 160$ region, indicating the detection of superhorizon adiabatic fluctuations (Peiris et al. 2003) first detected by *WMAP* (Kogut et al. 2003).

The 95% confidence upper limits on the BB spectrum (Figure 13) come from applying this likelihood approximation to each BB band power individually and excluding all other band powers from the calculation. We then apply a uniform

²⁶ We have checked that using *WMAP*-9 cosmological parameters instead of *WMAP*-5 makes a negligible difference.

positive prior on the band power and integrate the resulting posterior probability distribution function (PDF) to find the limit containing 95% of the probability.

9.3. Constraints on Tensor-to-Scalar Ratio, r

The primary motivation for the BICEP1 measurement of the BB spectrum is to constrain the tensor-to-scalar ratio, r . Following standard practice, we define r as the ratio of power in primordial gravitational waves to curvature perturbations at a pivot scale $k_0 = 0.002 \text{ Mpc}^{-1}$. To model the BB spectrum at a specific value, r_* , we simply scale the $r = 0.1$ model spectrum, described in Section 5.3, by $(r_*/0.1)$. This method, which uses a fixed template shape for the BB spectrum and scales the amplitude, does not technically satisfy the slow roll consistency relation, $n_T = -r/8$ (Kinney 1998), but it provides a convenient and model-independent measure of sensitivity to a tensor-type BB signal and is consistent with the treatments in C10 and publications from the QUIET Collaboration.

9.3.1. Direct Likelihood Calculation

End-to-end signal simulations containing a tensor-type BB spectrum allow us to directly compute the one-dimensional likelihood for r , without band power likelihood approximations as an intermediate step. This method involves the definition of a quadratic estimator,

$$\rho = \alpha + \beta \sum_b \mathcal{W}_b \mathcal{D}_b^{BB}, \quad (12)$$

where \mathcal{D}_b^{BB} are frequency-combined BB band powers, \mathcal{W}_b are weights selected to target the B -mode signature of inflation, and α and β are calibrated from simulations so that ρ is an unbiased estimator of r .

The weights used to combine BB band powers are calculated as

$$\mathcal{W}_b = \sum_{b'} \mathcal{M}_{bb'}^{-1} \mathcal{A}_{b'}, \quad (13)$$

where $\mathcal{M}_{bb'}$ is the 9×9 BB block of the band power covariance matrix and $\mathcal{A}_{b'}$ are signal expectation values calculated by applying the BB band power window functions to the template BB spectrum. Note that the covariance matrix used here is calculated from signal-plus-noise simulations with a standard Λ CDM theory spectrum and $r = 0$, meaning that the estimator ρ is optimized for the case where r is not detected. BICEP1 is designed to target the peak in the BB spectrum occurring at $\ell \sim 80$. As expected, the weights chosen by this method strongly emphasize the first three ℓ -bins, which contribute 37%, 44%, and 15%, respectively, to the total of \mathcal{W}_b . This choice does not bias our estimate of r , but merely imposes a negligible penalty to the noise of our estimator in the case of non-zero true r .

Next, we generate simulated maps for a range of r values by combining maps from the standard E -no- B signal-plus-noise simulations with B -no- E signal-only Q and U maps that have been scaled by $\sqrt{r/0.1}$. These simulations are inherently restricted to the physically meaningful range, $r \geq 0$. It is necessary to add B -modes to the maps, rather than simply adding a scaled BB power spectrum, because while the B -modes from the B -no- E signal simulations have no correlation with the signal-plus-noise maps, the cross terms do contribute additional variance to the band powers.

BB band powers are calculated from the maps, combined across frequencies, and then further summed according to the

weights derived above to obtain a “raw” version of the ρ -statistic, unscaled by β . Because the signal and noise are uncorrelated, the ensemble average of ρ is linearly proportional to r , allowing us to fit for α and β . The BB band powers have already been debiased for contributions from instrument noise and $E \rightarrow B$ leakage, so the fit value of α is small. However, our simulations do not include the B -mode signal generated by gravitational lensing of E -modes; we correct for this by debiasing the ρ -statistic calculated from real data by an amount corresponding to $r = 0.03$, which is the value obtained by applying the ρ estimator to the expected lensing BB spectrum. Applying the calibrated ρ estimator to the real BICEP1 data, we obtain $\hat{\rho} = 0.038 \pm 0.233$. The 1σ error bar on this estimator is given by the square root of the variance of ρ values simulated for the fiducial ($r = 0$) model. It is important to note that $\hat{\rho}$ is not a maximum likelihood estimate of r (maximum likelihood estimates are presented in Section 9.3.3). Rather, it is similar to the band powers and error bars shown in Figure 5, which are direct measurements of power in the map, but scaled and with error estimates from simulations.

By running the Λ CDM+ r simulations described above, we can determine the probability density of our estimator ρ as a function of the input r . We model this probability density function as a scaled and shifted χ^2 distribution, which fits the simulated histograms well. The shift in the distribution can be calculated from the known noise and $E \rightarrow B$ leakage biases, which had previously been subtracted from the band powers; the scaling and dof parameters are estimated from the mean and variance of the simulated ρ values. Including absolute gain and beam width calibration uncertainties modifies the distribution, slightly increasing its variance.

The BICEP1 likelihood function for r is obtained directly by calculating the probability of obtaining the observed value, $\hat{\rho}$, as a function of model parameter r . This likelihood function is shown as the red curve in the left panel of Figure 10. The tabulated likelihood computed by this method is available as part of the BICEP1 three-year data release. We consider it to be the most reliable description of our constraint on r as it avoids band power likelihood approximations.

9.3.2. Alternate Likelihood Calculation

In addition to the direct likelihood computed above, it is useful to derive an estimator that is distributed symmetrically about the true value of r . To that end we construct an alternate likelihood for r based upon the band power likelihood approximation of Section 9.1. This alternative makes more assumptions than the direct method but has the advantage of being defined in the unphysical region of negative r . Therefore, we use it to calculate the maximum likelihood r and associated 68% confidence interval, which we allow to extend into the negative r region. We calculate this likelihood using a theory spectrum template calculated from an $r = 0.1$ model and using information from the BB spectrum only. We include the effect of gravitational lensing B -modes by adding a constant lensing spectrum consistent with Λ CDM to the theory model at every r . We include the systematic uncertainty as described in Section 9.1. The resulting maximum likelihood and minimum width 68% interval (uniform prior) are $r = 0.03^{+0.27}_{-0.23}$ (Figure 10).

9.3.3. Upper Limit and Confidence Intervals

The left panel of Figure 10 shows the BICEP1 likelihood function for r calculated using both the direct method (simulation




Figure 10. Likelihood for r calculated from the BICEP1 BB spectrum is shown in the left panel. The red curve comes from a direct likelihood calculation described in Section 9.3.1. The blue curve comes from an alternate calculation based on the band power likelihood approximation (Section 9.3.2). The maximum likelihood value and 1σ interval, $r = 0.03^{+0.27}_{-0.23}$, are shown as the blue solid and dashed lines. A histogram of maximum likelihood r values derived from 499 signal-plus-noise simulations (with $r = 0$ input) is shown in the central panel. In the right panel, we derive 95% confidence upper limits on r from simulated likelihoods (gray histogram) and real data likelihood. BICEP1 obtains an upper limit of $r < 0.70$ (red line), which lies within the simulated distribution. The gray dashed line shows the median ($r < 0.65$) of the upper limits derived from simulations.

based) and alternate method (via HL band power likelihood approximation). The most notable difference is that the alternate likelihood calculation extends to negative non-physical values of r . Because of this feature, we choose to derive the maximum likelihood estimate of r from the alternate likelihood calculation, as the likelihood peak will always exist, regardless of whether the data contain a high or low noise fluctuation. For the specific case of the BICEP1 three-year results, the likelihood peaks at a slightly positive value, $r = 0.03$, and the two likelihood calculations agree on the peak position; this agreement is generally quite good for all simulated results with maximum likelihood r above zero. The 1σ error bar quoted on the maximum likelihood estimate is a minimum width 68% interval, calculated assuming a uniform prior on r (positive and negative). The center panel of Figure 10 shows the distribution of maximum likelihood estimates obtained from a set of 499 simulations with input $r = 0$.

To set a 95% confidence upper limit, we adopt a uniform prior for $r \geq 0$ only and calculate the one-sided 95% credible interval. This construction, which was previously used in C10, as well as many other experiments in the literature, has the welcome property that it will not yield arbitrarily low (or even negative) upper limits even in the case of unlikely downward fluctuations in the data. Since the calculation involves only the parts of the likelihood with $r \geq 0$, we can use the direct likelihood calculation, which diverges from the alternate likelihood specifically in the tails of the distribution. The upper limit from the direct likelihood is more conservative than the same limit calculated from the alternate likelihood, both for the specific case of the BICEP1 three-year data and also for simulations of that data. We believe that the direct likelihood is more accurate, though the agreement in the region of the likelihood peak shows that the HL band power likelihoods are an excellent choice for most purposes. The 95% confidence upper limit from three years of BICEP1 observations is $r < 0.70$. The right panel of Figure 10 shows this limit along with the distribution of upper limits obtained from simulations. The median upper limit, a useful benchmark of experimental sensitivity, is $r < 0.65$ at 95% confidence. For both likelihood methods (direct and HL band power likelihood), we estimate the




Figure 11. Bayesian 95% upper limit (dashed line) and Feldman-Cousins 95% confidence interval (dotted lines) on r for the BICEP1 three-year result, as a function of the value of ρ . The shaded image shows the probability density of ρ as a function of r , derived from simulations; each horizontal slice of the image yields a normalized PDF for ρ given a particular theory. Vertical slices correspond to likelihood functions for r . The solid vertical line indicates the value of ρ measured by BICEP1.

Monte Carlo uncertainty by repeating the above calculations for the first 250 and the last 249 realizations separately and find that the 95% confidence upper limits and the maximum likelihood estimate differ by $\Delta r < 0.04$ between each half.

The direct likelihood procedure, involving ρ values calculated across a range of input r models, lends itself naturally to the construction of frequentist confidence intervals. As an alternative to the one-sided 95% credible interval that we use for the headline upper limit on r , we also offer a frequentist 95% confidence interval following the construction described in Feldman & Cousins (1998). This interval construction is chosen because it handles the physical constraint, $r \geq 0$, in a natural way. Figure 11 shows the probability distribution for ρ as a function of the theory, with both the Bayesian upper limit and the Feldman-Cousins confidence interval shown. It is a feature of the Feldman-Cousins construction that the 95% confidence interval for a bounded theory can be either one-sided or two-sided depending on the data; for the BICEP1 three-year result, we obtain a one-sided 95% confidence interval, $r < 0.62$.




Figure 12. Histogram of the shift in the 95% confidence upper limit on r from simulations upon including the additional 52% data of the full BICEP1 observations. A negative value indicates that the three-year upper limit is tighter than the two-year limit, but 7% of realizations show a positive value. The dashed black line indicates the median of the distribution (-0.27). The solid black line indicates the value of this shift for the real data (-0.10).

10. CONCLUSIONS

In summary, we present improved measurements of the degree-scale CMB polarization from BICEP1. Compared to the previous data release (C10), we include 52% more data with a corresponding decrease in statistical uncertainty. We dramatically reduce systematic uncertainty by developing and implementing the relative-gain deprojection technique. We also implement two new likelihood calculations: a band power

likelihood based on the existing Hamimeche & Lewis (2008) approximation and a new direct simulation-based likelihood for r . Both likelihoods are available as part of our data release. We propose these methods as standards for future inflationary B -mode search experiments.

We support the new results with an extensive suite of consistency tests. First, we show that the new results are consistent with C10. The differences due to the change of analysis pipeline, noise model, band power window function calculation, and data selection are within expectation. Second, jackknife null tests confirm the internal consistency of the data and analysis.

The most important results of BICEP1 are the CMB power spectra band powers. Overall, the spectra are consistent with the Λ CDM cosmological model. We detect E -mode power in the first acoustic peak at 15σ , the most significant such detection to date (Figure 9). We confirm the TE superhorizon fluctuations, first detected by WMAP (Kogut et al. 2003), at 8σ . The total detection significance for non-zero EE power is 18σ , and 14σ for TE power.

The primary goal of BICEP1 is to search for the inflationary B -mode signal. The B -mode spectra are consistent with zero, and we place the strongest upper limits to date in the $50 < \ell < 200$ region where the signal from inflation is expected to peak (Figure 13). We also report this result as a constraint on the tensor-to-scalar ratio, r . This constraint is $r = 0.03^{+0.27}_{-0.23}$ (68% CI) or $r < 0.70$ (95% CL). The corresponding upper limit from C10 is $r < 0.72$.




Figure 13. BICEP1's EE and BB power spectra complement existing data from other CMB polarization experiments (Leitch et al. 2005; Montroy et al. 2006; Sievers et al. 2007; Bischoff et al. 2008; Brown et al. 2009; QUIET Collaboration et al. 2011, 2012; Bennett et al. 2013). For visual clarity, we only display the experiments where at least one of the EE band powers has a center value that is greater than twice the distance between the center value and the lower end of the 68% confidence interval. Theoretical spectra from a Λ CDM model with $r = 0.1$ are shown for comparison; the BB curve is the sum of the inflationary and gravitational lensing components. At degree angular scales, BICEP1's constraints on BB are the most powerful to date.

One might naively expect the upper limit to improve by a larger factor when adding 52% more data. We have confirmed that the relatively small decrease is not a result of the set of analysis refinements discussed in Section 6.1. The changes in mapmaking pipeline, noise model, band power window function calculation, and deprojection each produce small shifts in band powers and error bars as described above. But when we apply all of these changes together to the original C10 data set, the resulting r constraint derived using the same offset-lognormal likelihood approximation as C10 is $r < 0.71$. In other words, the net effect of these changes on the r upper limit is close to zero for this data set.

The relatively small decrease in the new upper limit is explained by two factors. First, the offset-lognormal likelihood approximation used in C10 resulted in a negative bias on the upper limit for those specific BB band powers (we find that the r constraint derived from offset-lognormal band power likelihoods is biased low for some cases and high for others). Applying the more accurate direct likelihood calculation to the reanalyzed C10 data set shifts the upper limit in this case from $r < 0.71$ to $r < 0.80$. Second, upon including the new data in this analysis, the upper limit fluctuates somewhat high compared to the average of simulations. With an $r = 0$ input model, simulated data sets run through our final analysis yield upper limits on r that decrease by a median of 0.27 when including the additional data of the full three years (Figure 12). The corresponding decrease seen in the real data is only 0.10 (from 0.80 to 0.70). Although this decrease is smaller than average, it is not an unlikely result; 17% of the simulations saw even less of a decrease, and in 7% of the simulations the upper limit actually increases when adding the additional data.

Interesting constraints can be placed on cosmic birefringence from the BICEP1 TB and EB spectra, which are predicted to be zero by the Λ CDM model. This topic will be explored in detail in Kaufman et al. (2013).

Measurement of CMB B -mode polarization remains the most promising approach for testing the inflationary paradigm. BICEP1 has provided the lowest upper limits on inflationary B -modes to date (Figure 13). In this data release, we demonstrate a deprojection technique that will enable future experiments to cope with the increasingly important temperature-to-polarization leakage and develop a direct likelihood calculation for converting band power results into constraints on r without approximations. BICEP2 has completed three years of observation with an order of magnitude better mapping speed at 150 GHz than BICEP1. Keck Array operations are ongoing, with two full years of observation completed by 2013 November (Ogburn et al. 2012; Aikin et al. 2010). BICEP3 will begin observing in 2014 and 2015. Measurements of B -mode polarization from these and other experiments, using the new analytical tools we have demonstrated here, have the potential to test inflationary cosmology with unprecedented precision.

BICEP1 was supported by NSF grant No. OPP-0230438, Caltech President's Discovery Fund, Caltech President's Fund PF-471, JPL Research and Technology Development Fund, and the late J. Robinson. This analysis was supported in part by NSF CAREER award No. AST-1255358 and the Harvard College Observatory, and J.M.K. acknowledges support from an Alfred P. Sloan Research Fellowship. B.G.K. acknowledges support from NSF PECASE Award No. AST-0548262. We thank the

South Pole Station staff for helping make our observing seasons a success. We also thank our colleagues in ACBAR, BOOMERANG, QUAD, BOLOCAM, SPT, WMAP, and Planck for advice and helpful discussions, and Kathy Deniston and Irene Coyle for logistical and administrative support. We thank Patrick Shopbell for computational support at Caltech and the FAS Science Division Research Computing Group at Harvard University for providing support to run all the computations for this paper on the Odyssey cluster.

REFERENCES

- Aikin, R. W., Ade, P. A., Benton, S., et al. 2010, *Proc. SPIE*, 7741, 77410V
- Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, *ApJS*, 208, 20
- Bierman, E. M. 2011, PhD thesis, Univ. California, San Diego
- Bierman, E. M., Matsumura, T., Dowell, C. D., et al. 2011, *ApJ*, 741, 81
- Bischoff, C., Hyatt, L., McMahon, J. J., et al. 2008, *ApJ*, 684, 771
- Bond, J. R., Jaffe, A. H., & Knox, L. 2000, *ApJ*, 533, 19
- Brown, M. L., Ade, P., Bock, J., et al. 2009, *ApJ*, 705, 978
- Challinor, A., & Chon, G. 2005, *MNRAS*, 360, 509
- Chiang, H. C. 2009, PhD thesis, California Institute of Technology
- Chiang, H. C., Ade, P. A. R., Barkats, D., et al. 2010, *ApJ*, 711, 1123
- Chon, G., Challinor, A., Prunet, S., Hivon, E., & Szapudi, I. 2004, *MNRAS*, 350, 914
- Delabrouille, J., Betoule, M., Melin, J.-B., et al. 2013, *A&A*, 553, A96
- Feldman, G. J., & Cousins, R. D. 1998, *PhRvD*, 57, 3873
- Finkbeiner, D. P., Davis, M., & Schlegel, D. J. 1999, *ApJ*, 524, 867
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759
- Hamimeche, S., & Lewis, A. 2008, *PhRvD*, 77, 103013
- Hanson, D., Hoover, S., Crites, A., et al. 2013, *PhRvL*, 111, 141301
- Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, *ApJS*, 208, 19
- Hu, W., & White, M. 1997, *NewA*, 2, 323
- Jarosik, N., Bennett, C. L., Dunkley, J., et al. 2011, *ApJS*, 192, 14
- Jones, W. C., Bhatia, R., Bock, J. J., & Lange, A. E. 2003, *Proc. SPIE*, 4855, 227
- Kamionkowski, M., Kosowsky, A., & Stebbins, A. 1997, *PhRvL*, 78, 2058
- Kaufman, J. P., Miller, N. J., Shimon, M., et al. 2013, arXiv:1312.7877
- Keating, B. G., Shimon, M., & Yadav, A. P. S. 2013, *ApJL*, 762, L23
- Kinney, W. H. 1998, *PhRvD*, 58, 123506
- Knox, L. 1999, *PhRvD*, 60, 103516
- Kogut, A., Spergel, D. N., Barnes, C., et al. 2003, *ApJS*, 148, 161
- Komatsu, E., Dunkley, J., Nolte, M. R., et al. 2009, *ApJS*, 180, 330
- Kovac, J. M., Leitch, E. M., Pryke, C., et al. 2002, *Natur*, 420, 772
- Leitch, E. M., Kovac, J. M., Halverson, N. W., et al. 2005, *ApJ*, 624, 10
- Lewis, A., Challinor, A., & Lasenby, A. 2000, *ApJ*, 538, 473
- Liddle, A. R., & Lyth, D. H. 2000, *Cosmological Inflation and Large-Scale Structure* (Cambridge: Cambridge Univ. Press)
- Montroy, T. E., Ade, P. A. R., Bock, J. J., et al. 2006, *ApJ*, 647, 813
- Moyerman, S. 2013, PhD thesis, Univ. California, San Diego
- Moyerman, S., Bierman, E., Ade, P. A. R., et al. 2013, *ApJ*, 765, 64
- O'Dea, T. D., Clark, C. N., Contaldi, C. R., & MacTavish, C. J. 2012, *MNRAS*, 419, 1795
- Ogburn, R. W., Ade, P. A. R., Aikin, R. W., et al. 2012, *Proc. SPIE*, 8452, 84521A
- Peiris, H. V., Komatsu, E., Verde, L., et al. 2003, *ApJS*, 148, 213
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013a, arXiv:1303.5076
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013b, arXiv:1303.5082
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013c, arXiv:1303.5088
- Pryke, C., Ade, P., Bock, J., et al. 2009, *ApJ*, 692, 1247
- QUIET Collaboration, Araujo, D., Bischoff, C., et al. 2012, *ApJ*, 760, 145
- QUIET Collaboration, Bischoff, C., Brizius, A., et al. 2011, *ApJ*, 741, 111
- Reichardt, C. L., Ade, P. A. R., Bock, J. J., et al. 2009, *ApJ*, 694, 1200
- Seljak, U. 1997, *ApJ*, 482, 6
- Seljak, U., & Zaldarriaga, M. 1997, *PhRvL*, 78, 2054
- Sievers, J. L., Achermann, C., Bond, J. R., et al. 2007, *ApJ*, 660, 976
- Sievers, J. L., Hlozek, R. A., Nolte, M. R., et al. 2013, *JCAP*, 10, 060
- Smith, K. M. 2006, *PhRvD*, 74, 083002
- Story, K. T., Reichardt, C. L., Hou, Z., et al. 2013, *ApJ*, 779, 86
- Takahashi, Y. D. 2010, PhD thesis, Univ. California, Berkeley
- Takahashi, Y. D., Ade, P. A. R., Barkats, D., et al. 2010, *ApJ*, 711, 1141
- Wu, J. H. P., Zuntz, J., Abroe, M. E., et al. 2007, *ApJ*, 665, 55
- Yoon, K. W. 2008, PhD thesis, California Institute of Technology