

This is the accepted manuscript made available via CHORUS. The article has been published as:

First CMB constraints on direction-dependent cosmological birefringence from WMAP-7

Vera Gluscevic, Duncan Hanson, Marc Kamionkowski, and Christopher M. Hirata

Phys. Rev. D **86**, 103529 — Published 26 November 2012

DOI: [10.1103/PhysRevD.86.103529](https://doi.org/10.1103/PhysRevD.86.103529)

First CMB Constraints on Direction-Dependent Cosmological Birefringence from WMAP-7

Vera Gluscevic¹, Duncan Hanson^{1,2}, Marc Kamionkowski^{1,3}, and Christopher M. Hirata¹

¹*California Institute of Technology, Mail Code 350-17, Pasadena, CA 91125, USA*

²*Jet Propulsion Laboratory, California Institute of Technology,*

4800 Oak Grove Drive, Pasadena CA 91109, USA

³*Department of Physics and Astronomy, Johns Hopkins University, Baltimore, MD 21218, USA*

A Chern-Simons coupling of a new scalar field to electromagnetism may give rise to cosmological birefringence, a rotation of the linear polarization of electromagnetic waves as they propagate over cosmological distances. Prior work has sought this rotation, assuming the rotation angle to be uniform across the sky, by looking for the parity-violating TB and EB correlations a uniform rotation produces in the CMB temperature/polarization. However, if the scalar field that gives rise to cosmological birefringence has spatial fluctuations, then the rotation angle may vary across the sky. Here we search for direction-dependent cosmological birefringence in the WMAP-7 data. We report the first CMB constraint on the rotation-angle power spectrum $C_L^{\alpha\alpha}$ for multipoles between $L = 0$ and $L = 512$. We also obtain a 68% confidence-level upper limit of $\sqrt{C_2^{\alpha\alpha}/(4\pi)} \lesssim 1^\circ$ on the quadrupole of a scale-invariant rotation-angle power spectrum.

I. INTRODUCTION

In this work, we use the cosmic microwave background (CMB) temperature and polarization maps of the Wilkinson Microwave Anisotropy Probe (WMAP) 7-year data release [1] to search for direction-dependent cosmological birefringence (CB). CB is a postulated rotation of the linear polarization of photons that propagate through cosmological distances [2]. It is present, for example, in models where a Nambu-Goldstone boson plays the role of quintessence [3], but also in models with new scalar degrees of freedom that have nothing to do with quintessence [4–7]. The rotation of the polarization is a consequence of the coupling of a scalar field to the electromagnetic Chern-Simons term, such that the rotation angle α is proportional to the total change $\Delta\phi$ of the field ϕ along the photon’s path.

Prior to this work, a rotation angle α that is uniform across the sky had been sought in the CMB [8], where it would induce parity-violating TB and EB temperature/polarization correlations [9]. CB has also been sought in quasar data [2, 10]. The tightest constraint currently comes from a combined analysis of the WMAP, Bicep [11], and QUAD experiments [12]; it is $-1.4^\circ < \alpha < 0.9^\circ$ at the 95% confidence level [13].

There are, however, a number of reasons to expand the search and look for a CB angle $\alpha(\hat{\mathbf{n}})$ that varies as a function of position $\hat{\mathbf{n}}$ on the sky. To begin with, a dynamical field ϕ that drives the rotation can have fluctuations, in which case the rotation angle varies across the sky [4–6]. Furthermore, if ϕ is some massless scalar, not necessarily quintessence, its background value does not necessarily evolve, and the uniform component of the rotation angle may vanish. The only way to look for CB in this scenario is through its direction dependence. Additionally, if $\alpha(\hat{\mathbf{n}})$ is measured with high significance, the exact shape of its power spectrum provides a window into the detailed physics of the new cosmic scalar ϕ . Currently,

the strongest limit on a direction-dependent CB angle comes from AGN [14], which constrain the root-variance of the rotation angle to be $\lesssim 3.7^\circ$.

In previous studies [15, 16], a formalism was developed to search for anisotropic CB rotation with the CMB. The sensitivity of WMAP data to this anisotropic rotation is expected to be competitive with that from AGN [15–17]. However, the CMB also allows individual multipoles $C_L^{\alpha\alpha}$ to be probed—the AGN data only constrain the variance—and is sensitive to higher L than AGN. The CMB also probes CB to a larger lookback time than AGN.

Here we apply the formalism developed earlier to the WMAP 7-year data. Within experimental precision, we report a non-detection of a direction-dependent cosmological birefringence. We obtain an upper limit on all the rotation-angle power-spectrum multipoles $C_L^{\alpha\alpha}$ up to $L = 512$. This result implies a 68% confidence-level upper limit on the quadrupole of a scale-invariant power spectrum of $\sqrt{C_2^{\alpha\alpha}/(4\pi)} \lesssim 1^\circ$,¹. As a check, we also find a constraint on the uniform rotation that agrees with the results of Ref. [13].

The rest of this paper is organized as follows. In §II, we review the physical mechanism for CB. In §III, we revisit the full-sky formalism to search for direction-dependent rotation, and discuss its implementation. WMAP data selection, our simulations, and the tests of the analysis method are described in §IV. Results are reported in §V, and we conclude in §VI. Appendix A contains a detailed explanation of the procedure we used to obtain an upper limit of the root-mean-squared rotation angle from the measurement of the TE correlation in the data; Ap-

¹ Here, the power spectrum is defined in the usual way, $C_L^{\alpha\alpha} \equiv \sum_M \alpha_{LM} \alpha_{LM}^*/(2L+1)$, where a spherical-harmonic decomposition of the rotation field provides the rotation-angle multipoles, $\alpha_{LM} \equiv \int Y_{LM}^*(\hat{\mathbf{n}}) \alpha(\hat{\mathbf{n}}) d\hat{\mathbf{n}}$.

pendix B contains a discussion of the geometrical properties of the rotation-angle estimator; Appendix C details the calculation of the L -dependence of the fractional correction for a scale-invariant power spectrum recovered from cut-sky maps; and Appendix D displays the analysis masks we used in this work.

II. PHYSICAL ORIGIN OF COSMOLOGICAL BIREFRINGENCE

Theories with a weakly broken global $U(1)$ symmetry provide a natural mechanism for producing a shallow potential for the pseudo-Nambu-Goldstone-boson field ϕ . From a cosmological perspective, a PNGB field with this property is a natural candidate for quintessence, since it can drive epochs of accelerated expansion [3]. In addition, many other extensions of the Standard Model of particle physics and Λ CDM cosmology abound in scalar fields descending from theories with shift symmetry and resembling the PNGB. Such fields generically couple to photons through the Chern-Simons term $F^{\mu\nu}\tilde{F}_{\mu\nu}$ of electromagnetism, while the underlying shift symmetry suppresses all other leading-order couplings to Standard-Model particles [3]. This way, the existence of a new degree of freedom ϕ could evade detection in colliders and other lab experiments, but could still be manifest in cosmology through CB. We now review in more detail the physical mechanism that gives rise to CB.

The Chern-Simons-modified electromagnetic Lagrangian reads

$$\mathcal{L} = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} - \frac{\beta}{2M}\phi F^{\mu\nu}\tilde{F}_{\mu\nu}, \quad (1)$$

where $F^{\mu\nu}$ is the electromagnetic field-strength tensor, $\tilde{F}^{\mu\nu}$ is its dual, β is a coupling constant, and the mass M is a vacuum expectation value of the spontaneously broken symmetry. The dispersion relation following from this modified Lagrangian has different solutions for the left- and right-handed photon polarizations, the net effect being the rotation of the linearly polarized electromagnetic wave that propagates through the vacuum with the evolving field ϕ . The direction of polarization is rotated by an amount

$$\alpha = \frac{\beta}{2M}\Delta\phi, \quad (2)$$

that depends on the total change $\Delta\phi$ along the photon's path. Since M can be arbitrarily large, perhaps on the order of the Planck mass, the accumulated change in ϕ must also be large in order for this angle to be measurable. This motivates the use of cosmological probes in search for CB. There are models in which $\Delta\phi$ is uniform across the sky (giving rise only to a uniform rotation angle), as well as models in which it has anisotropies [4–6]. In this work, we do not focus on any particular physical model for CB, but rather derive a model-independent constraint on the rotation-angle power spectrum $C_L^{\alpha\alpha}$.

III. FULL-SKY FORMALISM AND ITS IMPLEMENTATION

In this Section we review the full-sky-estimator formalism of Ref. [16] for measuring direction-dependent CB. In order to apply this formalism to the cut sky (after masking out the Galaxy), all measured power spectra need to be corrected by a factor of $\sim 1/f_{\text{sky}}$.² Unless otherwise noted, f_{sky} is calculated as the fraction of the pixels that the mask admits; we include this factor when appropriate in the following derivation. We rewrite all the relevant formulas in a position-space form which is numerically efficient for the analysis of data.

In the presence of birefringence, the polarization field acquires a phase factor,

$$p(\hat{\mathbf{n}}) \equiv [Q + iU](\hat{\mathbf{n}}) = \tilde{p}(\hat{\mathbf{n}})e^{2i\alpha(\hat{\mathbf{n}})}, \quad (3)$$

where Q and U are the Stokes parameters for linear polarization, and tilde denotes the polarization in the absence of birefringence, which we refer to as the “primary polarization”. To obtain an estimate of the phase factor $e^{2i\alpha(\hat{\mathbf{n}})}$ from the polarization field in Eq. (3), we require a tracer of $\tilde{p}(\hat{\mathbf{n}})$. The primary polarization is generated by Thomson scattering of the local temperature quadrupole, so the observed temperature field $T(\hat{\mathbf{n}})$ may be used for this purpose. Due to projection effects, the local temperature quadrupoles at last scattering appear on the sky as a curvature of the temperature field. The estimator for the rotation angle then involves projecting the temperature field into a map as a spin-2 quantity (which evaluates the curvature), and looking for correlation with the polarization field which varies as a function of the position on the sky. We review a rigorous derivation of the estimator in the following subsections.

A. Rotation-induced B modes

On the full sky, the polarization field can be decomposed in terms of spin-2 spherical harmonics ${}_2Y_{lm}(\hat{\mathbf{n}})$ as

$$p(\hat{\mathbf{n}}) = -\sum_{lm} (E_{lm} + iB_{lm}) {}_2Y_{lm}(\hat{\mathbf{n}}), \quad (4)$$

where E and B modes represent polarization patterns of opposite parity [20, 21].

The primary E -mode polarization signal \tilde{E}_{lm} (sourced by the dominant scalar perturbations) is detected with high significance in WMAP-7 data [13], although primary

² When multipole coefficients are calculated from a map where a fraction $1 - f_{\text{sky}}$ of the pixels is masked (i.e. signal set to zero), the usual full-sky expression for their variance (i.e. the power spectrum; see Ref. [20] or Eq. (13) where $f_{\text{sky}} = 1$) is underestimated by a factor of f_{sky} , because the variance corresponding to the masked pixels is effectively zero.

B modes (sourced by the subdominant tensor perturbations) have only been constrained with upper limits. For this reason, most of the constraining power for CB in WMAP comes from the search for a CB-induced rotation of the primary E mode into an observed B mode. The induced B mode is given as [15, 16]

$$B_{lm} = \frac{i}{2} \int d\hat{\mathbf{n}} [\tilde{p}(\hat{\mathbf{n}}) e^{2i\alpha(\hat{\mathbf{n}})} {}_2Y_{lm}^*(\hat{\mathbf{n}}) - \tilde{p}(\hat{\mathbf{n}})^* e^{-2i\alpha(\hat{\mathbf{n}})} {}_2Y_{lm}^*(\hat{\mathbf{n}})]. \quad (5)$$

This B mode is correlated with the primary E mode (from which it originated), and through it also with the temperature anisotropies. The presence of rotation therefore gives rise to anomalous EB and TB correlations, and both these power spectra can be used to search for CB. It is, however, worth keeping in mind that individual multipoles of the E -mode polarization signal are still noise dominated, whereas the temperature is measured at $S/N > 1$, for a large number of multipoles, in every frequency band of the WMAP 7-year data. Therefore, at WMAP noise levels the temperature field makes a better tracer of the primary E -mode than the observed E mode itself. For this reason, on most angular scales, the search for a TB correlation, which we focus on in this work, provides the best constraint on CB [16]. Assuming the primary polarization is a pure E mode at the surface of last scatter, the CB-induced TB correlation reads [15, 16]

$$\begin{aligned} \langle B_{lm} T_{l'm'}^* \rangle = & \int d\hat{\mathbf{n}} \tilde{C}_{l'l}^{TE} \\ & \times \left[\frac{1}{2} \sin(2\alpha) [{}_2Y_{l'm'} {}_2Y_{lm}^* + {}_2Y_{l'm'} {}_2Y_{lm}^*] \right. \\ & \left. - \frac{i}{2} \cos(2\alpha) [{}_2Y_{l'm'} {}_2Y_{lm}^* - {}_2Y_{l'm'} {}_2Y_{lm}^*] \right], \quad (6) \end{aligned}$$

where we suppress the $\hat{\mathbf{n}}$ dependence for clarity. Power spectrum $\tilde{C}_{l'l}^{TE}$ is the correlation between the temperature and the primary E mode, which may be calculated with CAMB [22].

So far, we have not assumed anything about the magnitude of the rotation per pixel in CMB maps. Observations of quasars suggest an upper bound on the root-mean-squared (RMS) of $\alpha(\hat{\mathbf{n}})$ of just a few degrees [14], while the measurement of the TE correlation from WMAP-7 data implies a somewhat weaker constraint: $\langle \alpha(\hat{\mathbf{n}})^2 \rangle < 9.5^\circ$ (see Appendix A for details). Motivated by these results, in the rest of this paper we adopt a small-rotation-angle limit. The numerical results we present in §VA do not depend on the validity of this assumption, but their interpretation as an upper limit of the rotation-angle autocorrelation $C_L^{\alpha\alpha}$ does; this subtlety is discussed in more detail in §IV, §V, and Appendix A.

In the limit of small rotation angle, only the sine term

contributes to the observed TB which then reads

$$\begin{aligned} \langle B_{lm} T_{l'm'}^* \rangle \approx & \int d\hat{\mathbf{n}} \tilde{C}_{l'l}^{TE} \sum_{LM} \alpha_{LM} Y_{LM}(\hat{\mathbf{n}}) \\ & \times [{}_2Y_{l'm'}(\hat{\mathbf{n}}) {}_2Y_{lm}^*(\hat{\mathbf{n}}) + {}_2Y_{l'm'}(\hat{\mathbf{n}}) {}_2Y_{lm}^*(\hat{\mathbf{n}})], \quad (7) \end{aligned}$$

where parity condition $l+l'+L=\text{even}$ must be satisfied³ for any given L . A TB correlation generated by weak gravitational lensing of the CMB is of opposite parity, with $l+l'+L=\text{odd}$, and does not represent a source of bias for measuring a small-rotation signal. In addition, the effect of lensing is not internally observable at WMAP noise levels, even with an optimal estimator [23]. We therefore do not consider lensing further.

B. Minimum-variance quadratic estimator: $\hat{\alpha}_{LM}$

From Eq. (7), it is evident that scale-dependent birefringence induces correlations between temperature and polarization modes at different wavenumbers l, l' ; i.e., it produces a statistically anisotropic imprint on the covariance matrix of the observed CMB. Each ll' pair measured in the maps may therefore be used as an estimate of the rotation-angle multipole α_{LM} , provided that it satisfies the usual triangle inequalities, $L \leq l+l'$ and $L \geq |l-l'|$, as well as the parity condition $l+l'+L=\text{even}$. The prescription for combining all ll' estimates in order to produce a minimum-variance quadratic estimator is explained in detail in Ref. [16]. Here, we only present the final expressions for the TB estimator,

$$\begin{aligned} \hat{\alpha}_{LM} = & N_L \int d\hat{\mathbf{n}} Y_{LM}(\hat{\mathbf{n}}) \\ & \times \left[\sum_{lm} \bar{B}_{lm}^* Y_{lm}(\hat{\mathbf{n}}) \sum_{l'm'} \tilde{C}_{l'l}^{TE} \bar{T}_{l'm'} Y_{l'm'}^*(\hat{\mathbf{n}}) \right. \\ & \left. + (\text{complex conjugate}) \right], \quad (8) \end{aligned}$$

where N_L is an L -dependent normalization and the barred quantities represent inverse-variance filtered multipoles. For full-sky coverage and homogenous noise in pixel space, the expressions for these quantities read

$$\bar{B}_{lm} \equiv \frac{B_{lm}}{C_l^{BB}}, \quad \bar{T}_{l'm'} \equiv \frac{T_{l'm'}}{C_{l'}^{TT}}. \quad (9)$$

The unbarred B and T are the observed temperature and polarization multipoles corrected for the combined instrumental beam and pixelization transfer function W_l , and the TT and BB power spectra are analytic estimates

³ The two contributions to the correlation, sin and cos, have opposite parities, where only terms that satisfy $l+l'+L=\text{even}$, and $l+l'+L=\text{odd}$, respectively, contribute to any given multipole L .

of the total signal + noise power spectrum in a given frequency band,

$$\begin{aligned} C_l^{TT} &\equiv \tilde{C}_l^{TT} + C_l^{TT, \text{ noise}}/W_l^2, \\ C_l^{BB} &\equiv \tilde{C}_l^{BB} + C_l^{BB, \text{ noise}}/W_l^2. \end{aligned} \quad (10)$$

In the idealized case of full-sky coverage and homogeneous instrumental noise, the estimator normalization N_L is calculable analytically and is equal to the inverse of the estimator variance,

$$N_L = \left(\sum_{l'} \frac{(2l+1)(2l'+1)}{4\pi} \frac{(\tilde{C}_{l'}^{TE})^2}{C_l^{BB} C_{l'}^{TT}} (H_{l'}^L)^2 \right)^{-1}, \quad (11)$$

where

$$H_{l'}^L = \begin{pmatrix} l & L & l' \\ -2 & 0 & 2 \end{pmatrix} + \begin{pmatrix} l & L & l' \\ 2 & 0 & -2 \end{pmatrix}. \quad (12)$$

The objects in parentheses are Wigner-3j symbols.

In the non-idealized case of real data, the simple inverse-variance filters (IVFs) presented above are sub-optimal, in the sense that the associated estimator variance is not truly minimized. To obtain a true minimum-variance estimate, computationally more involved filters are required [23]. In practice, however, we find that the full-sky expressions for the estimator in Eq. (8) provide a very good approximation to its behavior on the cut sky. Namely, the analytic expression for its variance, given by Eq. (11), is consistent with the full variance recovered from a suite of Monte Carlo simulations (described in detail in §IV B) when the simple IVFs of Eq. (10) are used in the presence of sky cuts; the appropriate correction for the fraction of the sky admitted by the analysis masks must be included in this case. This result motivates us to continue using the simple IVFs and the corresponding analytic expressions for the estimator normalization.

The insensitivity to the presence of the galaxy masks that we observe here can be interpreted as a consequence of the following properties. First, the estimator of Eq. (8) is a product of inverse-variance filtered T and B maps in real space, which are local functions of the data. The IVFs are local in pixel space (they resemble Gaussians with a width of a few arcmins, corresponding to the resolution in a given frequency band), and so the mask boundaries remain localized after filtering. Additionally, the estimator is an even function of the temperature map (see Eq. (8)—it contains a second derivative of the temperature field performed by ${}_2Y_{l'm'}$), and so it is relatively insensitive to the discontinuities introduced by the analysis mask. These properties put the rotation estimator $\hat{\alpha}_{LM}$ in sharp contrast with the estimators for the gravitational-lensing potential, where the dependence on the gradient of the temperature field renders the lensing reconstruction very sensitive to sky cuts [24].

C. Power-spectrum estimator: $\hat{C}_L^{\alpha\alpha}$

Once the rotation-angle multipoles are measured, their autocorrelation can be estimated as

$$C_L^{\hat{\alpha}\hat{\alpha}} \equiv \frac{1}{f_{\text{sky}}(2L+1)} \sum_M \hat{\alpha}_{LM} \hat{\alpha}_{LM}^*. \quad (13)$$

This represents a sum over the $\langle TBTB \rangle$ trispectrum, where “T” and “B” denote the temperature and B-mode multipole moments. This estimator for $C_L^{\alpha\alpha}$ is non-zero even in the absence of CB-induced rotation, due to the presence of Wick contractions from the primary CMB and the instrumental noise (discussed in §III B). They produce the “noise bias” $C_L^{\alpha\alpha, \text{ noise}}$ and must be subtracted from the measurement of $C_L^{\hat{\alpha}\hat{\alpha}}$, in order to recover an estimate of the CB-induced signal $C_L^{\alpha\alpha}$,

$$\hat{C}_L^{\alpha\alpha} = C_L^{\hat{\alpha}\hat{\alpha}} - C_L^{\alpha\alpha, \text{ noise}}. \quad (14)$$

For Gaussian CMB+noise fluctuations, the noise bias can be identified with the three disconnected Wick contractions of the trispectrum which $C_L^{\hat{\alpha}\hat{\alpha}}$ probes:

$$\begin{aligned} (a) & \quad \overline{\overline{TB} \overline{TB}} \\ (b) & \quad \overline{\overline{TB} \overline{TB}} \\ (c) & \quad \overline{\overline{TB} \overline{TB}}. \end{aligned} \quad (15)$$

where we neglect contraction (a), which only couples to the $L = 0$ mode of $\hat{\alpha}_{LM}$, and also contraction (c), as it is negligible⁴. In the absence of statistical anisotropy (i.e. for full-sky coverage and homogeneous instrumental noise), the contraction (b) between two real fields with multipoles $lm, l'm'$ carries a set of delta functions $\delta_{ll'} \delta_{mm'}$, and the realization-dependent noise bias may be written explicitly in terms of the observed power spectra. If $C_L^{\hat{\alpha}\hat{\alpha}}$ is evaluated by cross-correlating the f_1, f_2, f_3 , and f_4 frequency-band maps, the analytic expression for this “isotropic bias” follows from Eq. (8),

$$\begin{aligned} C_L^{\alpha^{f_1 f_2} \alpha^{f_3 f_4}, \text{ noise, iso}} &\equiv \langle \hat{\alpha}_{LM} \hat{\alpha}_{LM}^* \rangle_{\text{Gauss, iso}} = \\ &\sum_{l'} \frac{(2l+1)(2l'+1)}{4\pi} (H_{l'}^L \tilde{C}_{l'}^{TE})^2 \\ &\times \frac{C_{l'}^{TT, f_1 f_3 \text{ maps}} C_l^{BB, f_2 f_4 \text{ maps}}}{(C_{l'}^{TT, f_1 f_1} C_l^{BB, f_2 f_2} C_{l'}^{TT, f_3 f_3} C_l^{BB, f_4 f_4})}, \end{aligned} \quad (16)$$

where the power spectra in the denominator of Eq. (16) are the simple analytic IVFs. The $C_{l'}^{TT, f_1 f_3 \text{ maps}}$ and $C_l^{BB, f_2 f_4 \text{ maps}}$ are measured by cross-correlating data maps in the frequency bands f_1 and f_3 , or f_2 and f_4

⁴ In our simulations, we verify that this term has indeed a negligible numerical contribution.

respectively, and corrected by the factor of $1/f_{\text{sky}}^T$ and $1/f_{\text{sky}}^P$, corresponding to the temperature and polarization analysis mask, respectively. Most of the power in temperature comes from CMB fluctuations, and the B -mode power is mostly noise if $f_2 = f_4$, and negligible otherwise. Therefore, since the instrumental noise is independent for different frequency bands, the largest contribution to the noise bias can be eliminated by cross-correlating estimates of $\hat{\alpha}_{LM}$ obtained from two different bands.

In reality, we work with a masked sky which has been observed with inhomogeneous noise levels, and Eq. (16) does not provide a perfect description of the noise bias, although it is an excellent first approximation. This leads us to adopt a two-stage debiasing procedure in which we subtract both the isotropic bias of Eq. (16), and an additional Monte-Carlo-based correction, in order to correct for the effects of sky cuts and inhomogeneity of the instrumental noise. The total noise bias $C_L^{\alpha\alpha, \text{ noise}}$ is the sum of the two contributions,

$$C_L^{\alpha\alpha, \text{ noise}} \equiv C_L^{\alpha\alpha, \text{ noise, iso}} + C_L^{\alpha\alpha, \text{ noise, MC}}. \quad (17)$$

We estimate $C_L^{\alpha\alpha, \text{ noise, MC}}$ from a set of WMAP realizations generated with no birefringence signal (described in §IV B), analyzed in the same way as the data itself. For each realization, we calculate the appropriate $C_L^{\alpha\alpha, \text{ noise, iso}}$ and averaging over many realizations obtain $C_L^{\alpha\alpha, \text{ noise, MC}}$ as

$$C_L^{\alpha\alpha, \text{ noise, MC}} \equiv \langle \hat{C}_L^{\alpha\alpha} - C_L^{\alpha\alpha, \text{ noise, iso}} \rangle_{\text{sims}}. \quad (18)$$

This two-stage procedure reduces the sensitivity of our estimator to uncertainties in the CMB and instrumental-noise model, as compared to the case where the entire bias is recovered from Monte Carlo analysis. With the two-stage procedure, the largest (isotropic) contribution to the bias is evaluated directly from the power spectra of the observed maps, and is specific to the CMB realization at hand; subtracting it from the bispectrum naturally takes care of any noise (bias) contribution that might arise from the uncertainty in the background cosmology or in the noise description used to generate Monte Carlo simulations.

As we show in §V, we find consistency of results obtained with either (i) the calculation of the trispectrum as a four-point autocorrelation of the maps in the same band, (ii) the calculation of the trispectrum from cross-band correlations, which have an almost negligible noise bias.

IV. DATA AND SIMULATIONS

A. Data

Our main results are based on the full-resolution (corresponding to HEALPix resolution of $N_{\text{side}} = 512$) co-

added seven-year sky maps that contain foreground-reduced measurements of the Stokes I , Q , and U parameters in three frequency bands: Q (41 GHz), V (61 GHz), and W (94 GHz), available at the LAMBDA website [25]. A summary of the instrumental parameters most relevant to this analysis is provided in Table I. We apply

Band	FWHM	Δ_T [$\mu\text{Karcmin}$]	Δ_P [$\mu\text{Karcmin}$]
Q (41 GHz)	34'	316	544
V (61 GHz)	24'	387	589
W (94 GHz)	22'	467	693

Table I: Relevant instrumental parameters: beam full-width half maximum (FWHM) and approximate map noise for temperature and polarization for the three frequency bands we use in the analysis [25].

the seven-year temperature $KQ85y7$ mask with 78.27% of the sky admitted, and a polarization $P06$ mask with 73.28% of the sky admitted. These masks are constructed to remove diffuse emission based on the data in K and Q bands, and on a model of thermal dust emission. Point sources are masked based on a combination of external catalog data and WMAP-detected sources. (For more information about the exclusion masks, see Appendix D and Ref. [27].)

B. Simulations

We produce a suite of simulated WMAP observations, both to test the normalization of our α_{LM} estimates as well as to estimate their variance for the subtraction of $C_L^{\alpha\alpha, \text{ noise}}$ in Eq. (14). We produce simple simulations of the WMAP data with the following procedure:

1. Generate CMB-sky temperature and polarization realizations for the best-fit “ $\Lambda\text{CDM}+\text{SZ}+\text{ALL}$ ” WMAP-7 cosmology of Ref. [25].
2. Convolve the simulated CMB skies with a symmetric experimental beam. For the WMAP band maps we use an effective beam calculated as the average beam transfer function for all differencing assemblies at the given frequency.
3. Add simulated noise realizations based on the published I, Q, U covariance matrices within each pixel. We do not make any attempt to generate noise with pixel-to-pixel noise correlations, although we do exclude multipoles with $l < 100$ from our analysis, as this is where most of this correlated noise resides. In §V A we demonstrate the consistency of $\hat{C}_L^{\alpha\alpha}$ estimates constructed from auto- and cross-correlations of maps with independent noise realizations, and so are justified in neglecting correlated noise in our analysis.

We do not include Galactic foreground residuals or unresolved point sources in our simulations, but we address their possible impact on our results in §VC.

C. Test runs

In order to demonstrate the recovery of the CB signal using the minimum-variance estimator formalism and the de-biasing method discussed in the previous sections, we generate a suite of simulations that include a CB signal, i.e. where the polarization maps are rotated by realizations of a scale-invariant power spectrum of α ,

$$C_L^{\alpha\alpha} = AC_L^{\alpha\alpha, \text{fiducial}} \equiv A \frac{131 \text{deg}^2}{L(L+1)}, \quad (19)$$

where we choose the amplitude of this fiducial model so that it gives a S/N ratio of order 1 at low L for WMAP V band, and an RMS rotation-angle on the sky of 10° , satisfying the small-angle approximation⁵. We apply analysis masks to each simulated map, and then analyze the map cross-correlations (as described in previous sections), recovering $\hat{\alpha}_{LM}$ multipoles; we then compute the power spectrum using Eq. 13. Due to the interaction of the power distribution at different scales in the map with the geometry of the analysis mask, the f_{sky} factor is in principle a function of the multipole moment L , which typically starts smaller than the average⁶ value at low L 's, and converges towards the average value at high L 's. Since most of the signal for this model (which we come back to in §VB) comes from low L 's, we evaluate the exact L dependence, and substitute the $f_{\text{sky}}(L)$ function in Eq. 13 (for more details on $f_{\text{sky}}(L)$, see Appendix C). In Fig. 1, show the results of this test, comparing the input $C_L^{\alpha\alpha}$ power spectrum to the mean of the reconstructed power from a large number of simulations, and demonstrate a successful recovery of the signal. In the following Section, we apply the same signal-reconstruction method to WMAP-7 data.

To conclude this section, we note one subtlety necessary for the correct interpretation of the results of our analysis. The expression for the estimator of Eq. (8) only recovers the rotation-angle multipole in the small-angle regime. In the general case of arbitrarily large rotation, Eq. (8) provides an exact estimate of the multipoles of another “observable” quantity: $\frac{1}{2}\sin[2\alpha(\hat{\mathbf{n}})]$. Strictly speaking, our de-biasing procedure also relies on the small-angle approximation, since $C_L^{\alpha\alpha, \text{noise, MC}}$ is calculated

from a suite of null simulations. It is therefore necessary to inquire which regime corresponds to a particular model of rotation before interpreting our results as a constraint on such a model. However, the fiducial model we use as an example here (and which we come back to in §VB) satisfies this assumption (producing an RMS rotation of $\sim 10^\circ$). In this particular case, the difference between the two power spectra of α and of $\frac{1}{2}\sin[2\alpha]$ is mainly contained in the 15% difference in their amplitudes. It thus is possible to recover the rotation-angle power spectrum by simply rescaling the measured power—the fact we use in §VB to constrain this model from WMAP data.




Figure 1: The recovery of the CB signal with the analysis method presented in §III is demonstrated using a suite of simulations that include realization of a rotated sky. Blue dashed line is the input signal power-spectrum of $C_L^{\alpha\alpha}$, red thick line represents the power spectrum of $\frac{1}{2}\sin[2\alpha(\hat{\mathbf{n}})]$, and the thin black line is the mean recovered power from the simulations; the gray region is a 1σ -confidence-level interval calculated from the same suite of simulations.

V. RESULTS

A. Model-independent constraints

Before continuing, let us first clarify our notation. The rotation-angle power spectra are marked with four frequency bands as $[f_1 f_2][f_3 f_4]$. This means that the two estimates of α_{LM} needed to evaluate the power spectrum are obtained by cross-correlating band f_1 with f_2 , and f_3 with f_4 , respectively. Here, the temperature multipoles are measured from f_1 and f_3 , and the B modes are obtained from the maps in f_2 and f_4 bands. We measure five different band cross-correlations: $[\text{VV}][\text{VV}]$, $[\text{QV}][\text{QV}]$, $[\text{QQ}][\text{VV}]$, $[\text{WV}][\text{WV}]$, and $[\text{WW}][\text{VV}]$, but since the results for all of them are qualitatively the same, here we only show plots for a characteristic subset.

⁵ It is important that the model satisfies the small-angle approximation, as our calculation of the bias from Monte Carlo analysis is based on the null-assumption. In regime where this approximation is not satisfied, higher-order corrections will be necessary to recover the rotation-angle power spectrum from the measured $\langle T B T B \rangle$ trispectrum.

⁶ The “average” here is the usual f_{sky} fraction associated with a mask, equal to the fraction of the pixels that the mask admits.

Figs. 2, 3, and 4 show the measurement of the rotation-angle autocorrelation, before and after debiasing, and different components of the noise bias described in §III C. The blue and gray areas in the middle panels represent 1σ and 3σ confidence-level intervals, respectively, derived from the null-hypothesis (no rotation) Monte Carlo analysis described in §IV B. We see no significant deviations from zero in any of the five band cross-correlations—our results are consistent with $\alpha_{LM} = 0$ to within 3σ , at each multipole in the range from $L = 0$ to $L = 512$. We bin the power and list the measurements for all multipoles in Table II. As an additional consistency check, the upper limit we obtain on the uniform rotation angle, given as $\alpha \equiv \alpha_{00}/\sqrt{4\pi}$, is in good agreement with previous WMAP results [13]; see Table III.

As we pointed out in §IV C, in the general case of arbitrarily large rotation, our method provides an exact estimate of the autocorrelation of the following quantity: $\frac{1}{2}\sin[2\alpha(\hat{\mathbf{n}})]$, rather than the rotation angle itself; when the small-angle approximation is satisfied, this quantity and its power spectrum asymptote to $\alpha(\hat{\mathbf{n}})$ and $C_L^{\alpha\alpha}$, respectively. In order to establish the regime corresponding to a particular model, we note that the RMS fluctuation typical of realization of a power spectrum $C_L^{\alpha\alpha}$ is given by

$$\langle\alpha(\hat{n})^2\rangle^{1/2} = \sqrt{\sum_L \frac{2L+1}{4\pi} C_L^{\alpha\alpha}}. \quad (20)$$

In the event of a breakdown in the small-angle approximation, the values in Table II should be interpreted as constraints on the autocorrelation of $\frac{1}{2}\sin[2\alpha(\hat{\mathbf{n}})]$, rather than α itself. Evaluating Eq. (20) for the uncertainty levels quoted in Table II would erroneously lead to a conclusion that a large RMS rotation is allowed by the WMAP data; we show that the upper limit on the RMS rotation is roughly 9.5° (see Appendix A), and we again note that previous studies of quasar data imply an even stronger constraint of $\sim 4^\circ$ [14].

B. Constraints on a scale-invariant power spectrum

The null result shown in §V A can be translated into an upper limit on the amplitude of any model of rotation. As a generic example, we focus on a scale-invariant rotation-angle power spectrum of Eq. (19). The best-fit amplitude is evaluated from all multipoles in the range $0 \leq L \leq 512$, using a minimum-variance estimator [26]

$$\hat{A} = \sigma_A^2 \sum_L \frac{C_L^{\alpha\alpha, \text{fiducial}} \hat{C}_L^{\alpha\alpha}}{\text{var}(\hat{C}_L^{\alpha\alpha})}, \quad (21)$$

where

$$\sigma_A = \left(\sum_L \frac{(C_L^{\alpha\alpha, \text{fiducial}})^2}{\text{var}(\hat{C}_L^{\alpha\alpha})} \right)^{-1/2}, \quad (22)$$

L bin	[VV][VV]	[QV][QV]	[QQ][VV]	[WV][WV]	[WW][VV]
26	2.65±1.87	1.61±2.44	1.05±1.62	0.72±2.03	-0.43±1.34
77	1.86±2.58	0.70±2.84	1.57±2.36	1.03±2.70	0.17±2.04
128	1.07±1.33	1.00±1.36	0.27±1.17	3.04±1.35	0.96±1.02
179	1.40±1.49	-1.29±1.65	-0.31±1.15	-0.40±1.48	0.66±1.13
230	-1.90±1.76	-4.47±1.96	1.87±1.36	-3.36±1.97	-0.69±1.33
282	4.31±2.23	3.17±2.42	2.04±2.21	2.14±2.42	-0.20±1.90
333	1.98±2.39	-0.25±2.60	4.59±1.80	2.62±2.45	-1.11±1.96
384	0.81±1.78	-1.71±1.93	1.97±1.51	1.22±1.71	1.93±1.52
435	-0.40±1.64	-0.19±1.80	1.53±1.26	-1.03±1.74	-1.65±1.30
486	3.22±1.75	0.78±1.93	1.02±1.39	2.69±1.84	-0.28±1.27

Table II: Results for the measurement of $\hat{C}_L^{\alpha\alpha}$ [degrees²] are listed, as recovered from five different band cross-correlations. The 1σ confidence intervals are calculated with a suite of Gaussian sky simulations, described in §IV B. The results are binned, with the central L value of each bin listed in the table; the width of each bin is $\Delta L \sim 51$.

$[f_1 f_2]$	$\alpha \pm 1\sigma$ [°]
[VV]	-0.9 ± 2.3
[QV]	-0.5 ± 2.4
[QQ]	0.9 ± 2.8
[WV]	-2.2 ± 2.4
[WW]	-1.8 ± 2.7

Table III: Uniform-rotation angle α with a 1σ confidence interval, from five cross-band correlations of WMAP-7; the correction factor of $1/f_{\text{sky}}$ is applied to each measurement here. The uncertainties are consistent with the $\pm 1.4^\circ$ uncertainty on the uniform-rotation angle reported by the WMAP team [13] for a joint analysis of the Q, V and W-band data, after accounting for the fact that we analyze the bands individually (resulting in slightly larger error bars).

is the analytic expression for the variance of \hat{A} , and $\text{var}(\hat{C}_L^{\alpha\alpha})$ is the variance of the null-hypothesis rotation-angle power spectrum, estimated from a suite of null-hypothesis simulations. We note that the measured $\hat{C}_L^{\alpha\alpha}$ have been corrected by $f_{\text{sky}}(L)$ (see §IV C and Appendix C; the correction is calculated specifically for this model) only in this Subsection; for the presentation of the model-independent results, we use the average value, $f_{\text{sky}} \sim 0.68$. Most of the constraint here comes from low L 's; 50% of the sum in Eq. (22) comes from $L = 1$, and 90% from $L < 10$.

Even though the analytic expression above provides a good estimate of the statistical variance, because the constraint comes primarily from low- L modes the probability distribution of \hat{A} is significantly non-Gaussian. To capture this non-Gaussianity in our analysis, we again generate a suite of null-hypothesis Monte Carlo simulations and recover the 68% and 99% confidence-level intervals from these simulations. The corresponding probability




Figure 2: Top panel: Measurement of the rotation-angle power spectrum from V band, shown before debiasing, along with the components of the noise bias: Monte-Carlo measurement of the null-hypothesis mean $\langle \widehat{C}_L^{\alpha\alpha} \rangle$ (solid green), isotropic noise bias (blue dashed), and the mean isotropic bias (magenta dashed). Middle panel: The same power spectrum after debiasing, with 1σ and 3σ confidence interval. Bottom panel: binned version of the middle-panel power spectrum. The results are consistent with zero within 3σ .

distributions for \widehat{A} are shown in Figure 5.

The best-fit values for the quadrupole amplitude $\widehat{C}_2^{\alpha\alpha}$ and associated confidence intervals are listed in Table IV; consistency with zero is apparent within 3σ for all band-cross correlations we analyzed. The tightest constraint on the quadrupole amplitude of a scale-invariant rotation-angle power spectrum comes from [WW][VV]; it is $\sqrt{C_2^{\alpha\alpha}/(4\pi)} \lesssim 1^\circ$ with 68% confidence⁷.




Figure 3: Same as Figure 2, for $[f_1 f_2][f_3 f_4] = [\text{WW}][\text{VV}]$.

$[f_1 f_2][f_3 f_4]$	$\widehat{C}_2^{\alpha\alpha} \pm 1\sigma (\pm 3\sigma) [\text{deg}^2]$
[VV][VV]	$11.4^{+15.8}_{-16.9} (+79.0, -27.7)$
[QV][QV]	$29.6^{+18.8}_{-18.3} (+70.3, -33.4)$
[QQ][VV]	$19.8^{+14.3}_{-13.9} (+51.6, -46.6)$
[WV][WV]	$16.8^{+15.9}_{-16.9} (+79.0, -27.7)$
[WW][VV]	$3.0^{+14.0}_{-13.9} (+43.3, -42.9)$

Table IV: Measurement of the quadrupole amplitude of a scale-invariant rotation-angle power spectrum for different cross-band correlations, with 68% and 99% confidence-level intervals, recovered from a suite of null-hypothesis simulations. Consistency with zero within 3σ is apparent for all band cross-correlations, and the tightest constraint comes from [WW][VV].

⁷ Note that the conversion between the amplitude A and the

quadrupole is $C_2 = A \times 131 \text{ deg}^2/6$.




Figure 4: Same as Figure 2, for $[f_1 f_2][f_3 f_4] = [WV][WV]$.

C. Potential systematics

In addition to the statistical error reported here, there is also a systematic error for the measurement of the uniform rotation angle, owing to uncertainty in the detector polarization angles [13]. This systematic uncertainty should only apply to the monopole of α . The direction-dependent part is only sensitive to the extent that it affects the statistical noise bias, and this is mitigated by our data-dependent debiasing procedure. There are, however, other sources of systematic error that can potentially bias our estimates and add uncertainties to the rotation-angle measurements. In this Section we investigate the impact of asymmetry of the instrumental beams, unresolved polarized point sources, and foreground residuals from unremoved/unmasked Galactic emission.

1. Beam Asymmetries

The fast spin and precession rates of the WMAP scan strategy, coupled with the yearly motion of the satellite around the ecliptic, enforce that any bias to $\hat{\alpha}_{LM}$ originating from scan-strategy related systematics, like beam asymmetry, must be confined to $M = 0$ modes in ecliptic coordinates [29]. Furthermore, the smoothness of the scan strategy on large scales (dictated by the 85-degree opening angle of the detectors, and the large 22.5-degree amplitude of the hourly satellite precession), ensure that any such bias falls off quickly as a function of L . The estimates of C_2 , as we have discussed in the previous section, are most sensitive to the low- L modes of $\hat{C}_L^{\alpha\alpha}$, so to test for the presence of beam-asymmetry contamination, we rotate our coordinate system to ecliptic coordinates, and re-derive a constraint on C_2 from $L < 10$, $M = 0$ modes. We see no departure from the null hypothesis in this case where it should be maximal, and so conclude that beam asymmetries are not a significant source of bias for our measurements.

2. Unresolved point sources

To test the impact of unresolved point sources on our results, we repeat the analysis after unmasking the portion of the maps which are associated with detected point sources (but not Galactic contamination). Compared to our fiducial analysis, the measurement points for $\hat{C}_L^{\alpha\alpha}$ shift by $\lesssim 1\sigma$, where σ represents the statistical error from our foreground-free Monte Carlo analysis; see Figure 6. This shift provides a conservative upper limit on the systematic uncertainty that point-source residuals can produce, assuming that the bright detected population has comparable polarization properties to those of the fainter sources. The unresolved point-source power at WMAP frequencies is dominated by unclustered radio sources, with fluxes close to the detection threshold, and so this is a reasonable assumption. We note that there is no overall bias, as the direction of scatter does not appear to be correlated for different multipoles. Of course, the contribution of radio point sources to the map is a steep function of the flux cut, and by unmasking all detected point sources our estimate of potential bias and uncertainty is overly conservative. Given a model for radio-source number counts dN/dS , we can scale these results to the levels of contamination expected at the actual WMAP source detection threshold of (conservatively) ~ 1 Jy. Any bias $\Delta\hat{C}_L^{\alpha\alpha}$ (which we do not see evidence for, even in the unmasked map) will scale with the point-source trispectrum as

$$\Delta\hat{C}_L^{\alpha\alpha} \propto \int_{S=0}^{S_{\text{cut}}} S^4 \frac{dN}{dS} dS, \quad (23)$$




Figure 5: Probability distributions of the best-fit amplitude \hat{A} of the scale-invariant rotation power $C_L^{\alpha\alpha}$ recovered from a suite of null-hypothesis realizations are shown for some of the cross-band correlations. The gray-shaded areas denote a 68% confidence interval around a mean value; the red vertical line represents the measurement of the best-fit \hat{A} for a given band-cross correlation. We find consistency with zero within 3σ for all five measurements.

while the uncertainty on $\hat{C}_L^{\alpha\alpha}$ will scale with the point-source power as

$$\sigma(\hat{C}_L^{\alpha\alpha}) \propto \left(\int_{S=0}^{S_{\text{cut}}} S^2 \frac{dN}{dS} dS \right)^2. \quad (24)$$

Evaluating these terms for the dN/dS model of Ref. [28], we find that Δ and σ should be suppressed by factors of 0.005 and 0.06 respectively when moving from a flux cut of 10 Jy (no masking) to 1 Jy. We find even smaller (though comparable in magnitude) results using the simpler $dN/dS \propto S^{-2.15}$ model of Ref. [30]. This implies that any bias from unresolved sources should be completely negligible, and any increase in uncertainty due to their contribution to the observed power should be $\lesssim 0.1\sigma$, where σ represents the statistical error from our point-source-free Monte Carlo analysis. In conclusion, we expect the unresolved point sources to produce a negligible systematic uncertainty in the measurement of $\hat{C}_L^{\alpha\alpha}$.

3. Foreground residuals

An additional conceivable source of systematic uncertainty might result from Galactic foregrounds. To explore the extent to which such uncertainty might affect our results, we perform two tests. In the first, we repeat our analysis on non-foreground-reduced maps, to test the effect of the presence of unsubtracted foregrounds. In the second, we repeat the foreground-reduced analysis using a mask which excludes a larger fraction of the low-Galactic-latitude sky. We construct this conservative mask by combining the fiducial *KQ85y7* analysis mask with the extended mask of Ref. [25], and additionally masking out pixels with Galactic latitudes in the range of $\pm 40^\circ$; the mask admits only about 33% of the sky, approximately half the sky admitted in our fiducial analysis (where $\sim 68\%$ of the sky is admitted; see Appendix D). The function of this test is to explore the effect of residuals left by the foreground subtraction procedure, which




Figure 6: Measurement of $\hat{C}_L^{\alpha\alpha}$ from $[WW][VV]$. Results shown in black are obtained by using the analysis mask that covers all the point sources brighter than ~ 1 Jy, while the results in red (empty circles) are obtained after unmasking all the point sources. There is no apparent bias and the difference in every bin is less than the statistical uncertainty, despite the extreme variation in the source contamination. Scaling arguments in §V C 2 imply that the unresolved point sources have a negligible contribution to the estimated measurement uncertainty for the most constraining band-cross correlation.

should be stronger close to the Galactic plane. The measurement of $\hat{C}_L^{\alpha\alpha}$ is scaled appropriately to account for the fractional sky coverage and the results from the two modified analyses are compared with the results of the fiducial analysis in Figure 7. In the first case, the change in the measurements is negligible compared to the statistical uncertainty. In the second case, the scatter between the two results is consistent with the difference in sky coverage (producing up to 40% larger scatter for the extended-mask data points). The measurements show no apparent bias in either case. These results imply that foregrounds and foreground residuals are not likely to

make a large systematic contribution to our estimated statistical uncertainty, at least for the case of the most constraining band-cross correlation [WW][VV].




Figure 7: Measurement of $\hat{C}_L^{\alpha\alpha}$ from [WW][VV] band cross-correlation, with the corresponding statistical uncertainty obtained from a suite of null-hypothesis foreground-free simulations. Black filled circles are measurements obtained from the foreground-reduced maps after applying the fiducial analysis mask, and they represent our fiducial results of Figure 3. The colored data points (and the associated error bars) correspond to the two test analyses: the green x 's are obtained from the maps prior to foreground subtraction, but using the fiducial mask, while the red empty circles are measurements obtained from foreground-reduced maps after applying an extended mask (with an additional $\sim 35\%$ of the sky covered around the Galactic plane). The results from the two tests show no apparent bias. For the case of foreground-non-reduced analysis, the difference from the fiducial measurements is negligible compared to the statistical uncertainty; for the extended-mask case, the scatter between the two results is consistent with the difference in sky coverage. This implies that foregrounds and foreground residuals should not have a drastic impact on the estimated measurement uncertainty.

VI. SUMMARY AND CONCLUSIONS

In this work, we implement the minimum-variance quadratic-estimator formalism of Ref. [16] to search for direction-dependent cosmological birefringence with WMAP 7-year data. We derive the first CMB measurement of the rotation-angle power spectrum in the range $L = 0 - 512$, finding consistency with zero at each multipole, within 3σ . We estimate an upper limit on each power-spectrum multipole by simulating a suite of Gaussian sky realizations with no rotation, including symmetric beams and noise realizations appropriate for each WMAP frequency band, and also Q - U correlations and sky cuts. We investigate the impact of foregrounds

and polarized diffuse point sources on the reported constraints and come to the conclusion that they are not significant sources of systematic error for the rotation-angle estimates. Finally, we use the null result to get a 68% confidence-level upper limit of $\sqrt{C_2^{\alpha\alpha}/(4\pi)} \lesssim 1^\circ$ on the quadrupole of a scale-independent rotation-angle power spectrum. Even though the CMB constraint turns out to be comparable to that derived from quasar measurements, the CMB analysis has a significant advantage: it provides a measurement of the rotation-angle power at each individual multipole L and has better sensitivity to models with significant power at high multipoles.

The same formalism we use in this work can be applied to forecast the sensitivity of upcoming and future-generation CMB satellites to detecting direction-dependent cosmological birefringence. With 7 years worth of integration time with WMAP, we are able to constrain the uniform component of the rotation to less than about a degree; it will be interesting to see the results of this analysis method applied to the upcoming data from Planck satellite [31], where the sensitivity to rotation angles on the order of a few arcminutes [16] is expected.

Acknowledgments

This work was supported by DoE DE-FG03-92-ER40701, NASA NNX12AE86G, DoE (DOE-DE-SC0006624), and the David and Lucile Packard Foundation. Part of the research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Some of the results in this paper have been derived using the HEALPix package [32].

Appendix A: Constraints on rms rotation from WMAP

If the primordial B mode is small compared to the primordial E mode, and the rotation field is independent of the CMB, the measured TE correlation reads (see also Figure 8)

$$C_l^{TE} = \langle \cos[2\alpha(\hat{\mathbf{n}})] \rangle \tilde{C}_l^{TE}, \quad (\text{A1})$$

where the mean is taken over all realizations of the rotation field, and it does not depend on the direction $\hat{\mathbf{n}}$. In the case the probability distribution for α is a Gaussian centered at zero and with a width $\langle \alpha(\hat{\mathbf{n}})^2 \rangle^{1/2}$, the expectation value in Eq. (A1) is simply related to the pixel-variance of α ,

$$\langle \cos[2\alpha(\hat{\mathbf{n}})] \rangle = e^{-2\langle \alpha(\hat{\mathbf{n}})^2 \rangle}. \quad (\text{A2})$$

Therefore, an estimate of this expectation value and its uncertainty, obtained from the TE measurement as compared to the primordial power spectrum \hat{C}_l^{TE} provides an

upper limit of the rotation-angle pixel-variance. Adopting the expressions for a minimum-variance estimator and its variance (analogous to those of Eqs. (21) and (22)), we obtain: $\langle \cos[2\alpha(\hat{\mathbf{n}})] \rangle = 0.997 \pm 0.050$ (note that this constraint follows from the 21σ confidence-level detection of the TE correlation reported by Ref. [13]), implying an upper limit on the rotation rms $\langle \alpha(\hat{\mathbf{n}})^2 \rangle^{1/2} \lesssim 9.5^\circ$.




Figure 8: TE correlation measured from WMAP-7 data (black) is compared to the primordial power spectrum, generated using CAMB for the best-fit cosmology with no rotation (red, solid); both power spectra are obtained from Ref. [25]. The uncertainty on this measurement (gray) leaves room for a maximal rotation-angle rms of roughly 9.5° .

Appendix B: Visualizing the CB kernels

To illustrate the shape in harmonic space of the statistical anisotropy introduced by CB, we plot here the power-spectrum kernel as well as the geometric Wigner-3j contributions to the CB kernels in the ll' space from Eq. (8); see Figs. 9 and 10. The structure of the power-spectrum kernel originates from the polarization and temperature power spectra; the terms that correspond to the acoustic peaks in the primordial TT and TE power spectra have the largest contribution to the sum over ll' . The geometric weight dictates the shape of the l, l', L triangles which are generated by CB at a scale L . The terms where either l or l' is close in value to L have the largest contribution. The combination of the geometric weight and the power-spectra weight dictates the size of the noise bias at any particular L . The interplay of the two, for example, produces a peak at $L \sim 270$, apparent in all the plots of the noise bias presented in this work. The local maximum in the variance of $\hat{C}_L^{\alpha\alpha}$ also appears at this scale.

Appendix C: $f_{\text{sky}}(L)$

In order to evaluate the exact L -dependence of f_{sky} used to reconstruct the scale-invariant rotation-angle power from the CMB maps (see §IV C), we generate a large number of $\alpha(\hat{\mathbf{n}})$ realizations of the power-spectrum model of Eq. (19), mask the sky with the fiducial analysis mask, and then recover the input power spectrum in the usual way, i.e. take the pseudo- C_L of the masked map. The $f_{\text{sky}}(L)$ shown in Fig. 11 is the average ratio of the output to the input power, as a function of the multipole moment L .

Appendix D: Analysis masks

Here we visualize all the analysis masks we used in this paper: the fiducial temperature and polarization masks, as well as the two test masks used in §V C.




Figure 9: The Wigner-3j geometric factors in the summands of Eqs. (8) and (16) for three different values of L are shown. The geometric factor is non-zero only in the region of the ll' space where the triangle inequalities are satisfied. The dominant contribution comes from the triangles in which $L \sim l$, or $L \sim l'$, i.e. where either the temperature or the polarization mode has a length scale comparable to the length scale of the rotation-angle mode.




Figure 10: The power-spectrum kernel of the summands in Eqs. (8) and (16) is shown. The types of triangles that contribute the most to the isotropic bias of $C_L^{\hat{\alpha}\hat{\alpha}}$ are set by the geometric properties of spin-2 Wigner-3j symbols illustrated by the kernel shown in Figure 9, which is modulated by this kernel to produce summands in the expression for the bias.




Figure 11: The L dependence of the f_{sky} factor used for the reconstruction of the scale-invariant rotation-angle power spectrum in §IV C and §V B. The horizontal (green) line at $f_{\text{sky}} = 0.68$ represents the fraction of pixels admitted by the mask.

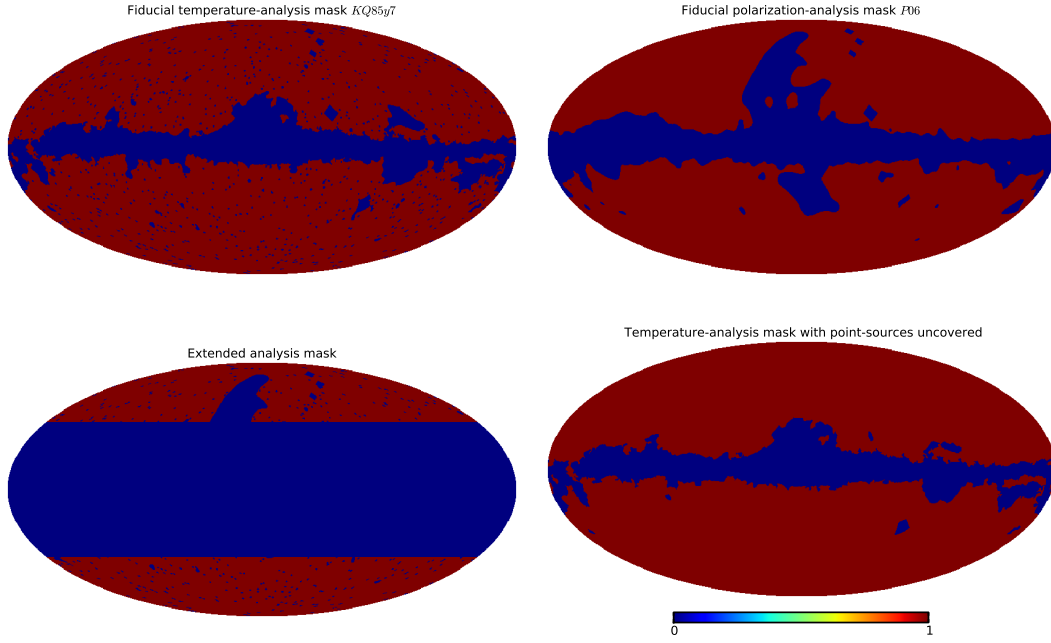


Figure 12: Fiducial temperature-analysis mask (with $\sim 78\%$ of the sky admitted), fiducial polarization-analysis mask ($\sim 73\%$), extended mask ($\sim 33\%$), and the temperature-analysis mask with point sources uncovered ($\sim 82\%$). Fraction of the sky admitted of the combined fiducial masks for polarization and temperature is $\sim 68\%$, which is twice the sky admitted as compared to the extended mask.

-
- [1] N. Jarosik, C. L. Bennett, J. Dunkley, B. Gold, M. R. Greason, M. Halpern, R. S. Hill and G. Hinshaw *et al.*, *Astrophys. J. Suppl.* **192**, 14 (2011) [arXiv:1001.4744 [astro-ph.CO]].
- [2] S. M. Carroll, G. B. Field and R. Jackiw, *Phys. Rev. D* **41**, 1231 (1990).
- [3] S. M. Carroll, *Phys. Rev. Lett.* **81**, 3067 (1998) [arXiv:astro-ph/9806099].
- [4] M. Li and X. Zhang, *Phys. Rev. D* **78**, 103516 (2008) [arXiv:0810.0403 [astro-ph]].
- [5] M. Pospelov, A. Ritz, C. Skordis, A. Ritz and C. Skordis, *Phys. Rev. Lett.* **103**, 051302 (2009) [arXiv:0808.0673 [astro-ph]].
- [6] R. R. Caldwell, V. Gluscevic and M. Kamionkowski, *Phys. Rev. D* **84**, 043504 (2011) [arXiv:1104.1634 [astro-ph.CO]].
- [7] F. Finelli and M. Galaverni, *Phys. Rev. D* **79**, 063002 (2009) [arXiv:0802.4210 [astro-ph]].
- [8] B. Feng, M. Li, J. Q. Xia, X. Chen and X. Zhang, *Phys. Rev. Lett.* **96**, 221302 (2006) [arXiv:astro-ph/0601095]; P. Cabella, P. Natoli and J. Silk, *Phys. Rev. D* **76**, 123014 (2007) [arXiv:0705.0810 [astro-ph]]; T. Kahniashvili, R. Durrer and Y. Maravin, *Phys. Rev. D* **78**, 123009 (2008) [arXiv:0807.2593 [astro-ph]]; J. Q. Xia, H. Li, G. B. Zhao and X. Zhang, *Astrophys. J.* **679**, L61 (2008) [arXiv:0803.2350 [astro-ph]]; J. Q. Xia, H. Li, X. l. Wang and X. m. Zhang, *Astron. Astrophys.* **483**, 715 (2008) [arXiv:0710.3325 [hep-ph]]; L. Pagano *et al.*, *Phys. Rev. D* **80**, 043522 (2009) [arXiv:0905.1651 [astro-ph.CO]]. A. Gruppuso, P. Natoli, N. Mandolesi, A. De Rosa, F. Finelli and F. Paci, *JCAP* **1202**, 023 (2012) [arXiv:1107.5548 [astro-ph.CO]]; R. Antonucci, *hep-ph/0110206*.
- [9] A. Lue, L. M. Wang and M. Kamionkowski, *Phys. Rev. Lett.* **83**, 1506 (1999) [arXiv:astro-ph/9812088]; N. F. Lepora, arXiv:gr-qc/9812077; G. -C. Liu, S. Lee and K. -W. Ng, *Phys. Rev. Lett.* **97**, 161303 (2006) [astro-ph/0606248].
- [10] B. Nodland and J. P. Ralston, *Phys. Rev. Lett.* **78**, 3043 (1997) [arXiv:astro-ph/9704196]; T. J. Loredo, E. E. Flanagan and I. M. Wasserman, *Phys. Rev. D* **56**, 7507 (1997) [arXiv:astro-ph/9706258]; D. J. Eisenstein and E. F. Bunn, *Phys. Rev. Lett.* **79**, 1957 (1997) [arXiv:astro-ph/9704247]; J. F. C. Wardle, R. A. Perley and M. H. Cohen, *Phys. Rev. Lett.* **79**, 1801 (1997) [arXiv:astro-ph/9705142]; J. P. Leahy, arXiv:astro-ph/9704285. S. M. Carroll and G. B. Field, *Phys. Rev. Lett.* **79**, 2394 (1997) [arXiv:astro-ph/9704263]; A. Cimatti, S. di Serego Alighieri, G. B. Field and R. A. E. Fosbury, *Astrophys. J.* **422**, 562 (1994); S. d. S. Alighieri, F. Finelli and M. Galaverni, *Astrophys. J.* **715**, 33 (2010) [arXiv:1003.4823 [astro-ph.CO]]; P. P. Kronberg, C. C. Dyer, and H.-J. Röser, *Astrophys. J.* **472**, 115 (1996); D. Horns, L. Maccione, A. Mirizzi and M. Roncadelli, arXiv:1203.2184 [astro-ph.HE]; L. Shao and B. -Q. Ma, *Phys. Rev. D* **83**, 127702 (2011) [arXiv:1104.4438 [astro-ph.HE]].
- [11] H. C. Chiang, P. A. R. Ade, D. Barkats, J. O. Battle, E. M. Bierman, J. J. Bock, C. D. Dowell and L. Duband *et al.*, *Astrophys. J.* **711**, 1123 (2010) [arXiv:0906.1181 [astro-ph.CO]].
- [12] E. Y. S. Wu *et al.* [QUaD Collaboration], *Phys. Rev. Lett.* **102**, 161302 (2009) [arXiv:0811.0618 [astro-ph]].
- [13] E. Komatsu *et al.* [WMAP Collaboration], *Astrophys. J. Suppl.* **192**, 18 (2011) [arXiv:1001.4538 [astro-ph.CO]].
- [14] M. Kamionkowski, *Phys. Rev. D* **82**, 047302 (2010) [arXiv:1004.3544 [astro-ph.CO]].
- [15] M. Kamionkowski, *Phys. Rev. Lett.* **102**, 111302 (2009) [arXiv:0810.1286 [astro-ph]].
- [16] V. Gluscevic, M. Kamionkowski and A. Cooray, *Phys. Rev. D* **80**, 023510 (2009) [arXiv:0905.1687 [astro-ph.CO]].
- [17] A. P. S. Yadav, R. Biswas, I. M. Su and M. Zaldarriaga, *Phys. Rev. D* **79**, 123009 (2009) [arXiv:0902.4466 [astro-ph.CO]].
- [18] <http://healpix.jpl.nasa.gov>.
- [19] D. Larson, J. Dunkley, G. Hinshaw, E. Komatsu, M. R. Nolta, C. L. Bennett, B. Gold and M. Halpern *et al.*, *Astrophys. J. Suppl.* **192**, 16 (2011) [arXiv:1001.4635 [astro-ph.CO]].
- [20] M. Kamionkowski, A. Kosowsky and A. Stebbins, *Phys. Rev. D* **55**, 7368 (1997) [arXiv:astro-ph/9611125]; M. Kamionkowski, A. Kosowsky and A. Stebbins, *Phys. Rev. Lett.* **78**, 2058 (1997) [arXiv:astro-ph/9609132].
- [21] M. Zaldarriaga and U. Seljak, *Phys. Rev. D* **55**, 1830 (1997) [arXiv:astro-ph/9609170]; U. Seljak and M. Zaldarriaga, *Phys. Rev. Lett.* **78**, 2054 (1997) [arXiv:astro-ph/9609169].
- [22] A. Lewis, A. Challinor and A. Lasenby, *Astrophys. J.* **538**, 473 (2000) [astro-ph/9911177].
- [23] K. M. Smith, O. Zahn and O. Dore, *Phys. Rev. D* **76**, 043510 (2007) [arXiv:0705.3980 [astro-ph]].
- [24] C. M. Hirata, N. Padmanabhan, U. Seljak, D. Schlegel and J. Brinkmann, *Phys. Rev. D* **70**, 103501 (2004) [astro-ph/0406004].
- [25] <http://lambda.gsfc.nasa.gov/>
- [26] V. Gluscevic and M. Kamionkowski, *Phys. Rev. D* **81**, 123529 (2010) [arXiv:1002.1308 [astro-ph.CO]].
- [27] B. Gold, N. Odegard, J. L. Weiland, R. S. Hill, A. Kogut, C. L. Bennett, G. Hinshaw and J. Dunkley *et al.*, *Astrophys. J. Suppl.* **192**, 15 (2011) [arXiv:1001.4555 [astro-ph.GA]].
- [28] G. De Zotti, M. Massardi, M. Negrello and J. Wall, *Astron. Astrophys. Rev.* **18**, 1 (2010) [arXiv:0908.1896 [astro-ph.CO]].
- [29] D. Hanson, A. Lewis and A. Challinor, *Phys. Rev. D* **81**, 103003 (2010) [arXiv:1003.0198 [astro-ph.CO]].
- [30] R. C. Bolton, G. Cotter, G. G. Pooley, J. M. Riley, E. M. Waldram, C. J. Chandler, B. S. Mason and T. J. Pearson *et al.*, *Mon. Not. Roy. Astron. Soc.* **354**, 485 (2004) [astro-ph/0407228].
- [31] <http://www.rssd.esa.int/PLANCK>.
- [32] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke and M. Bartelman, *Astrophys. J.* **622**, 759 (2005) [astro-ph/0409513].