

CMB multipole measurements in the presence of foregrounds

Angélica de Oliveira-Costa* and Max Tegmark

MIT Kavli Institute for Astrophysics and Space Research, Cambridge, Massachusetts 02139, USA

(Received 16 March 2006; published 13 July 2006)

Most analysis of cosmic microwave background spherical harmonic coefficients $a_{\ell m}$ has focused on estimating the power spectrum $C_\ell = \langle |a_{\ell m}|^2 \rangle$ rather than the coefficients themselves. We present a minimum-variance method for measuring $a_{\ell m}$ given anisotropic noise, incomplete sky coverage and foreground contamination, and apply it to the Wilkinson Microwave Anisotropy Probe (WMAP) data. Our method is shown to constitute lossless data compression in the sense that the widely used quadratic estimators of the power spectrum C_ℓ can be computed directly from our $a_{\ell m}$ -estimators. As the Galactic cut is increased, the error bars $\Delta a_{\ell m}$ on low multipoles go from being dominated by foregrounds to being dominated by leakage from other multipoles, with the intervening minimum defining the optimal cut. Applying our method to the WMAP quadrupole and octopole as an illustration, we investigate the robustness of the previously reported “axis of evil” alignment to Galactic cut and foreground contamination.

DOI: [10.1103/PhysRevD.74.023005](https://doi.org/10.1103/PhysRevD.74.023005)

PACS numbers: 98.70.Vc

I. INTRODUCTION

Cosmology has been revolutionized by the advent of precision maps of the cosmic microwave background (CMB), allowing accurate tests of the cosmological standard model (CSM) and measurements of its free parameters (e.g., [1–5]). Although the analysis of CMB maps has traditionally focused on the power spectrum, various surprises discovered in the data have triggered a growing interest in extracting additional information pertaining to possible non-Gaussianity. Specifically, the all-sky CMB temperature fluctuation map $\delta T(\hat{\mathbf{r}})$ is customarily expanded in spherical harmonics:

$$\delta T(\hat{\mathbf{r}}) = \sum_{\ell m} a_{\ell m} Y_{\ell m}(\hat{\mathbf{r}}). \quad (1)$$

These multipole coefficients $a_{\ell m}$ are treated as stochastic variables, varying randomly when extracted from CMB maps in widely separated Hubble volumes. According to the CSM, the CMB fluctuations are for all practical purposes Gaussian and statistically isotropic, which implies that the $a_{\ell m}$ -coefficients are independent Gaussian random variables with zero mean, i.e., that the CMB contains no cosmological information whatsoever except for its power spectrum $C_\ell = \langle |a_{\ell m}|^2 \rangle$.

Although some inflation models predict potentially observable departures from Gaussianity on small angular scales ($\ell \gtrsim 10^2$), no such non-Gaussianity has yet been found (see [6,7] and references therein). However, a number potential non-Gaussian anomalies have recently been discovered on large angular scales. The surprisingly low CMB quadrupole C_2 has intrigued the cosmology community ever since it was first observed by COBE/DMR [8], and was taken more seriously once the precision measurements of the Wilkinson Microwave Anisotropy Probe (WMAP) [9] showed that it could not be completely blamed on Galactic foreground contamination. Although the low

quadrupole does not involve non-Gaussianity and its statistical significance is debatable [1,10–12], it is a fly in the ointment of the CSM, and closer investigation of this has revealed numerous hints of non-Gaussianity on large angular scales.

The first reported hints of such non-Gaussianity involved the quadrupole and octopole. They were both found to be rather planar, i.e., with most of their hot and cold spots centered in a single plane in the sky, with their two preferred planes surprisingly closely aligned [13,14]. In the CSM context, the reported alignment required a 1-in-66 fluke and the octopole planarity an independent 1-in-20 coincidence [14]. These features were confirmed by other groups and examined with other techniques, finding puzzling alignments all the way up to $\ell = 5$ that may require as much as a 1-in-1000 fluke [14–26]. The preferred axis, dubbed the “axis of evil” by [20], points towards Virgo and is intriguingly close to the ecliptic poles [16,24,26]. Additional hints of low- ℓ non-Gaussianity have also turned up involving, e.g., asymmetries [15–25,27–44].

These purely observational results have triggered numerous papers on potential physical explanations. These include effects of Galactic foreground emission [12,45–48] and local structures [26,49–56] as well as theoretical explanations going beyond the CSM and involving compact cosmic topology [14,57–69], modified inflation [70–83], or other new physics [84–101].

A substantial fraction of this work has involved the $a_{\ell m}$ -coefficients from Eq. (1) with $\ell \leq 5$. Given this interest in the low- ℓ multipoles, it is timely and worthwhile to measure them as accurately as possible, with quantified error bars, further improving on approaches such as [13,28]. This is the purpose of the present paper. More accurate $a_{\ell m}$ -measurements can either build confidence in the CSM or teach us about interesting new physics.

Because of Galactic foreground contamination and anisotropic detector noise, the best way to measure the multipole coefficients $a_{\ell m}$ is not to simply apply the relation

*Electronic address: angelica@space.mit.edu




FIG. 1 (color online). Foreground contamination is clearly visible in the WMAP V-band map (top left). After taking a simulated WMAP CMB map (top right), adding simulated noise and foregrounds, applying the TOH cleaning method [13] and subtracting the CMB back out, foregrounds remain clearly visible in the residual map (bottom left). We apply our multipole measurement technique after masking out each of seven regions of decreasing cleanliness (bottom right) to optimize the trade-off between residual foreground contamination and limited sky coverage (which causes leakage from unwanted multipoles coupling to our estimator).

$a_{\ell m} = \int Y_{\ell m}(\hat{\mathbf{r}})^* \delta T(\hat{\mathbf{r}}) d\Omega$ to a CMB map (see Fig. 1). Rather, the customary approach is to discard the most contaminated part of the sky and use a more elaborate linear weighting on the remainder of the map. In Sec. II, we study how to optimize this procedure. We then apply our method to the latest WMAP data in Sec. III, investigate the implications for the quadrupole-octopole alignment in Sec. IV and summarize our conclusions in Sec. V.

II. BASIC METHOD

A. The problem

Given a CMB map with pixels $i = 1, \dots, N$, let x_i denote the observed temperature fluctuation δT in the pixel whose sky direction corresponds to the unit vector $\hat{\mathbf{r}}_i$. For the case of WMAP, there are $N = 12 \times 512^2 = 3\,145\,728$ pixels distributed according to the HEALPix scheme¹ [102,103], but we will only use those outside of some appropriate Galaxy cut below, and at a lower resolution. Suppose we wish to measure a particular set of multipole coefficients a_{ℓ, m_j} , $j = 1, \dots, M$, labeled simply a_j for brevity.² Grouping the pixels

¹The HEALPix package is available from <http://www.eso.org/science/healpix/>.

²For convenience, since CMB maps are real rather than complex-valued, we work with real-value spherical harmonics $Y_{\ell m}$ throughout this paper. These are obtained from the standard spherical harmonics by replacing $e^{im\phi}$ by $\sqrt{2} \sin m\phi$, 1, $\sqrt{2} \cos m\phi$ for $m < 0$, $m = 0$, $m > 0$, respectively. As detailed in Appendix A of [14], the corresponding real-valued coefficients $a_{\ell m}$ are related to the traditional complex-valued coefficients $\bar{a}_{\ell m}$ by a simple unitary transformation: for $m < 0$, $m = 0$ and $m > 0$, $\bar{a}_{\ell m}$ equals $\sqrt{2} \text{Im} a_{\ell m}$, $a_{\ell 0}$ and $\sqrt{2} \text{Re} a_{\ell m}$, respectively. This means that for $m < 0$, $m = 0$ and $m > 0$, the traditional complex coefficients $\bar{a}_{\ell m}$ equal $(-1)^m (a_{\ell m} - i a_{\ell, -m}) / \sqrt{2}$, $a_{\ell m}$ and $(a_{\ell m} + i a_{\ell, -m}) / \sqrt{2}$, respectively.

into an N -dimensional vector \mathbf{x} and the multipole coefficients into an M -dimensional vector \mathbf{a} , we can rewrite Eq. (1) as

$$\mathbf{x} = \mathbf{Y}\mathbf{a} + \mathbf{n}, \quad (2)$$

where the $N \times M$ spherical harmonic matrix \mathbf{Y} is defined by $Y_{ij} = Y_{\ell, m_j}(\hat{\mathbf{r}}_i)$ and the “noise” vector \mathbf{n} contains all contributions to the sky map except from the desired multipoles. In other words, \mathbf{n} contains not only detector noise, but also genuine sky signal contributed by other multipole coefficients that are not included in the \mathbf{a} -vector. We make the usual assumption that the noise has zero mean ($\langle \mathbf{n} \rangle = \mathbf{0}$) and define its covariance matrix $\mathbf{C} \equiv \langle \mathbf{n}\mathbf{n}^t \rangle$. If the detector noise covariance matrix is \mathbf{N} , then we can write this in the usual form

$$\mathbf{C} = \mathbf{N} + \mathbf{S}, \quad S_{ij} = \sum_{\ell} \frac{2\ell + 1}{4\pi} P_{\ell}(\hat{\mathbf{r}}_i \cdot \hat{\mathbf{r}}_j) C_{\ell}, \quad (3)$$

where \mathbf{S} is the CMB contribution, P_{ℓ} is a Legendre polynomial and the sum runs over those multipoles that are not included in our \mathbf{a} -vector. We ignore the issue of foreground contamination for now, but will cover this important issue in detail below in Sec. III.

We wish to find an estimator $\hat{\mathbf{a}}$ of the true multipole vector \mathbf{a} that is unbiased ($\langle \hat{\mathbf{a}} \rangle = \mathbf{a}$) and whose elements have as low variance as possible.

B. The solution

For our low- ℓ applications, the linear system given by Eq. (3) is normally greatly overdetermined with $N \gg M$, i.e., with many more pixels than desired multipole coefficients. Mathematically equivalent linear problems frequently occur in the CMB literature in the contexts of mapmaking and foreground removal, and the optimal solution is well known to be [104]

$$\hat{\mathbf{a}} = \mathbf{W}\mathbf{x}, \quad \mathbf{W} \equiv [\mathbf{Y}^t \mathbf{C}^{-1} \mathbf{Y}]^{-1} \mathbf{Y}^t \mathbf{C}^{-1}, \quad (4)$$

with covariance matrix

$$\Sigma \equiv \langle \hat{\mathbf{a}} \hat{\mathbf{a}}^t \rangle - \langle \hat{\mathbf{a}} \rangle \langle \hat{\mathbf{a}} \rangle^t = [\mathbf{Y}^t \mathbf{C}^{-1} \mathbf{Y}]^{-1}. \quad (5)$$

To make this method useful in practice, we need to augment it with a prescription for how to optimize the Galaxy cut when faced with foreground contamination—we will do this in Sec. III. Before that, let us now provide some intuition for how the method works.

First note that since the spherical harmonics are orthogonal over the full sphere, complete sky coverage with uniform detector noise per pixel will make the covariance matrix Σ diagonal, reducing our method to simply expanding the map in spherical harmonics the usual way. The measurement of a given harmonic $a_{\ell m}$ is then the exact sky signal plus a detector noise contribution.

If part of the sky is missing due to foreground masking or lack of observation, or if the pixel weighting varies

because of nonuniform detector noise, then the spherical harmonics are no longer orthogonal, which is reflected in Σ from Eq. (5) having off-diagonal elements. However, the matrix algebra corrects for this coupling between the multipole coefficients. Equation (4) implies that $\mathbf{W}\mathbf{Y} = \mathbf{I}$ so that the measurement error $\hat{\mathbf{a}} - \mathbf{a} = \mathbf{W}(\mathbf{Y}\mathbf{a} + \mathbf{n}) - \mathbf{a} = \mathbf{W}\mathbf{n}$, independent of \mathbf{a} . For example, if \mathbf{a} contains the 16 multipole coefficients with $\ell \leq 3$, then the measured quadrupole coefficient \hat{a}_{20} will equal the true full-sky value a_{20} plus a noise contribution. This noise will include detector noise and leakage from sky multipoles with $\ell > 3$, but there will be no contribution whatsoever from the monopole, dipole or octopole, or from quadrupole coefficients a_{2m} with $m \neq 0$. The method will generically be able to solve for these 16 unknowns as long as there are at least 16 pixels, but the error bars will clearly grow as the Galaxy cut is increased, since this makes it more difficult to disentangle the different multipoles. Numerically, we indeed find that the matrix $[\mathbf{Y}'\mathbf{C}^{-1}\mathbf{Y}]$ remains nonsingular for all the Galaxy cuts we consider, so no regularization techniques are needed. However, we will see that the error bars grow sharply as the cut increases and the matrix becomes less well-conditioned, particularly because multipoles with large $|m|$ live predominantly in the masked-out regions near the Galactic plane.

C. Relation to other methods

The multiple estimation paper most closely related to this one is that of [28]. They consider two alternative linear techniques that, cast in the notation of the present paper, simply use different \mathbf{W} -matrices than the one given by Eq. (4), perform a careful and detailed study of their statistical properties, and apply them to the WMAP data. The first technique they explore is Wiener filtering, defined by $\mathbf{W} \equiv \mathbf{Y}'\mathbf{S}^{-1}[\mathbf{Y}'\mathbf{C}^{-1}\mathbf{Y}]^{-1}$. Using numerical simulations, they confirm that this causes a systematic suppression of power with increasing ℓ (decreasing signal-to-noise) as expected. They then focus on the power equalization (PE) filter defined by a \mathbf{W} -matrix based on Cholesky decomposing the spherical harmonic coupling matrix $\mathbf{Y}'\mathbf{Y}$. The PE filter has the attractive property of eliminating all contributions to a given multipole estimator from lower multipoles, but it is neither unbiased (in the sense that $\mathbf{W}\mathbf{Y} \neq \mathbf{I}$) nor minimum-variance. For the special case where we include all multipoles up to some ℓ in the \mathbf{a} -vector, our estimator for the last $a_{\ell m}$ -coefficient will also be independent of the lower multipoles, merely with a lower variance than the PE estimator [since the estimator defined by Eq. (4) by construction gives the smallest variance of any estimator with this property].

D. Relation to power spectrum estimation

The most common technique in the CMB community for measuring CMB power spectra uses so-called quadratic estimators [105,106], largely because they have been

shown to be information-theoretically optimal, giving the smallest possible error bars. Eqs. (4) and (5) imply that

$$\hat{\mathbf{a}}' \Sigma^{-2} \hat{\mathbf{a}} = \mathbf{x}' \mathbf{C}^{-1} \mathbf{Y} \mathbf{Y}' \mathbf{C}^{-1} \mathbf{x}. \quad (6)$$

When \mathbf{Y} contains all spherical harmonics for a given ℓ , the right hand side is precisely the quantity that a quadratic estimator of C_ℓ extracts from the map \mathbf{x} . Equation (6) therefore shows that we can calculate optimal estimators of the power spectrum from our measured multipoles $\hat{\mathbf{a}}$ without recourse to the CMB map. In other words, our method can be viewed as a form of lossless data compression, with all information from the map \mathbf{x} about the power spectrum coefficient C_ℓ retained in the measured coefficients $\hat{\mathbf{a}}$.

III. APPLICATION TO WMAP

Let us now apply our method to the WMAP data [9]. After briefly describing the data and foreground masks used, we present results for a variety of Galactic cuts. To quantify the foreground contribution and select the best cut, we then create and analyze simulated CMB and foreground maps. This also allows us to optimize other practical aspects of our method such as the pixel size.

A. Data and foreground masks used

Our analysis is based on the 1 yr WMAP data described in [9] with the maps from the five observing frequencies combined into the single foreground-cleaned map TOH map described in [13] and downloadable at <http://www.lambda.gsfc.nasa.gov>. The TOH map is an all-sky CMB map with the same resolution as the highest-frequency WMAP channel (about 12.6' FWHM), HEALPix-pixelized at resolution level $n_{\text{side}} = 512$. The TOH foreground cleaning algorithm assumes that the CMB has a blackbody spectrum, but is otherwise completely blind, making no assumptions about the CMB power spectrum, the foregrounds, WMAP detector noise or external templates.

We explore the same series of increasingly conservative foreground masks used in [13] (Fig. 1, bottom right). Their construction is described in detail in [13], and involves the following key steps. One first converts the maps at the five WMAP frequency bands K, Ka, Q, V and W to a common angular resolution and forms four difference maps W-V, V-Q, Q-Ka and Ka-K, thereby obtaining maps guaranteed to be free of CMB signal that pick up any signals with a non-CMB spectrum. One then creates a combined ‘‘junk map’’ (foreground map) by taking the largest absolute value of these four maps at each pixel. Finally, with appropriate smoothing, one creates sky regions based on contour plots of this map. We use cuts that are roughly equispaced on a logarithmic scale, corresponding to thresholds of 30000, 10000, 3000, 1000, 300 and 100 μK (Fig. 1, right bottom). We label our cuts as masks 0, 1, . . . , 8, where mask 8 uses

pixels with less than $100 \mu\text{K}$ in the junk map and mask 0 uses the entire sky. Masks 1 & 2 merely cut out small blobs in the Galactic plane as described in [13]. For masks 0, . . . , 8, the total sky percentages removed are 0%, 0.03%, 0.08%, 0.25%, 0.64%, 2.0%, 6.6%, 24% and 64%, respectively.

B. Basic results and their interpretation

For our basic results, we work at HealPix pixel level $n_{\text{side}} = 16$ and use the \mathbf{C} -matrix corresponding to noise and CMB measurements by the WMAP team [9], with C_ℓ in Eq. (3) given by the best fit ΛCDM model from [1]. As detailed below, our low- ℓ results are quite insensitive to these choices.

Figure 2 shows the result of applying our method to this data to measure the components of the quadrupole and octopole. The results are plotted as a function of Galactic cut, with the thin black lines showing the measurements $\hat{\mathbf{a}}_{\ell m}$ and the shaded bands illustrating the 1σ error range $\hat{a}_{\ell m} \pm \Delta\hat{a}_{\ell m}$. These plotted error bars are $\Delta\hat{a}_{\ell m_j} = \Delta\hat{a}_j = \Sigma_{jj}^{1/2}$. The results for more multipoles are listed in Table I. The measurements using the full sky (Mask 0) match those reported in [13,14]. Table II shows the corresponding results for the 3-year WMAP data [107–110], which was released after the original submission of this paper. The results are quite similar—we discuss the differences below in Secs IV and V.³ The quadrupole moments were computed by taking \mathbf{a} to include all multipole coefficients with $\ell \leq 2$. It is important to always include the monopole ($\ell = 0$) and dipole ($\ell = 1$) so that our method makes the measurements that we are interested in (in this case $\ell = 2$) completely independent of these totally unknown quantities (the maps released by the WMAP team have already been approximately purged of a 3K monopole corresponding to the mean CMB temperature and a 6 mK dipole corresponding to Earth’s motion). Similarly, the octopole measurements were computed by taking \mathbf{a} to include all multipole coefficients with $\ell = 3$ and $\ell \leq 1$ (We will use the alternative $\ell \leq 3$ option in Sec. IV where we study quadrupole-octopole alignment and want uncorrelated measurements of the two.)

Note that the error bars $\Delta\hat{a}_{\ell m}$ on, say, the quadrupole coefficients shown in Fig. 2 do *not* include quadrupole sample variance the way an error bar on a measurement of C_2 customarily would. This is why they are seen to be so tiny for Mask 0, reflecting only the detector noise contribution. The error bar $\Delta\hat{a}_{\ell m}$ simply reflects the uncertainty in our measurement of the coefficient $\hat{a}_{\ell m}$ on the sky seen from our particular vantage point in space. These coefficients vary randomly with a variance C_2 between widely




FIG. 2 (color online). The left and right panels show how the $a_{\ell m}$ -coefficients (phases) of the quadrupole (left) and the octopole (right) extracted from the TOH map change as more sky is masked out. Black curves give me the measurements for each mask number, shaded bands reflect 1-sigma errors from noise and multipole leakage. For example, as explained in detail in Sec. III B, the quadrupole error bars (left) do *not* include quadrupole sample variance the way an error bar on a measurement of C_2 customarily would, since we are interested in the actual $a_{\ell m}$ -values rather than the underlying power spectrum, but they *do* include sample variance aliased from higher multipoles. Based on our simulations, the coefficient most susceptible to foreground contamination from the Galactic plane is a_{20} (top left band), so its mask 6 value is probably closer to the truth than its mask 0 value.

separated Hubble volumes, and this quadrupole sample variance therefore enters only in the next data analysis step where our five locally measured quadrupole coefficients are used to make inferences about the value of C_2 .

As mentioned in Sec. II, however, the error bars $\Delta\hat{a}_{\ell m}$ do include sample variance from the other multipoles that were excluded from the \mathbf{a} -vector, i.e., $\ell \geq 3$ for our quadrupole example. This is why the error bars in Fig. 2 are seen to flare up dramatically towards the right: as large portions of the sky get masked out, the severely broken orthogonality between the spherical harmonics makes higher multipoles contribute to quadrupole estimators. Since the values of these higher multipole coefficients are unknown to us and are not measured by the method in this case, their contribution is counted as noise—which can be quantified since their variance C_ℓ is known. Since our method tries to minimize the error bars, it makes the best trade-offs it can to minimize the net effect of such leakage from nonincluded multipoles. For multipoles included in \mathbf{a} , in contrast, our method forces the leakage to be exactly zero, at

³We foreground-cleaned the 3 yr WMAP data with the exact same TOH algorithm [13] as was used for the 1-year data, and the interested reader can download the cleaned map from <http://space.mit.edu/home/tegmark/wmap.html>.

TABLE I. Measured multipole coefficients in μK . More coefficients are available online. The relation between these real-valued spherical harmonic coefficients and the usual complex ones is given in footnote ².

ℓ	m	Mask 0		Mask 5		Mask 6		
		Fit all	Fit 0, 1, ℓ	Fit 0–3	Fit 0–5	Fit 0, 1, ℓ	Fit 0–3	Fit 0–5
2	-2	-21.29 ± 0.20	-20.48 ± 1.21	-20.56 ± 1.22	-20.63 ± 1.22	-24.07 ± 2.49	-24.28 ± 2.49	-24.36 ± 2.53
2	-1	5.94 ± 0.20	5.76 ± 0.23	5.76 ± 0.23	5.76 ± 0.23	5.89 ± 0.40	5.90 ± 0.40	5.93 ± 0.40
2	0	10.61 ± 0.22	13.64 ± 1.11	13.69 ± 1.11	13.86 ± 1.12	14.04 ± 2.29	14.10 ± 2.30	14.32 ± 2.34
2	1	-8.30 ± 0.20	-8.21 ± 0.26	-8.21 ± 0.26	-8.20 ± 0.26	-9.14 ± 0.59	-9.15 ± 0.59	-9.08 ± 0.60
2	2	-19.39 ± 0.20	-20.99 ± 1.48	-20.97 ± 1.48	-21.16 ± 1.49	-17.81 ± 3.10	-17.70 ± 3.11	-18.53 ± 3.18
3	-3	40.56 ± 0.20	43.17 ± 1.49	43.19 ± 1.50	43.32 ± 1.51	41.84 ± 3.00	41.94 ± 3.00	42.15 ± 3.06
3	-2	2.54 ± 0.20	2.29 ± 0.27	2.29 ± 0.27	2.28 ± 0.27	2.87 ± 0.50	2.86 ± 0.50	2.84 ± 0.50
3	-1	1.00 ± 0.20	2.91 ± 1.03	2.89 ± 1.03	2.78 ± 1.04	2.55 ± 2.10	2.50 ± 2.10	1.86 ± 2.14
3	0	-6.47 ± 0.22	-6.34 ± 0.29	-6.34 ± 0.29	-6.36 ± 0.29	-7.24 ± 0.65	-7.23 ± 0.65	-7.28 ± 0.66
3	1	-12.90 ± 0.20	-16.45 ± 1.21	-16.49 ± 1.21	-16.72 ± 1.22	-12.05 ± 2.47	-12.19 ± 2.48	-12.95 ± 2.55
3	2	30.37 ± 0.20	30.76 ± 0.26	30.76 ± 0.26	30.76 ± 0.26	31.53 ± 0.72	31.51 ± 0.72	31.45 ± 0.73
3	3	-19.32 ± 0.20	-19.02 ± 1.41	-18.98 ± 1.41	-18.97 ± 1.42	-20.76 ± 2.95	-20.63 ± 2.96	-20.94 ± 3.04
4	-4	-9.69 ± 0.20	-12.12 ± 1.51		-12.49 ± 1.52	-9.27 ± 3.07		-10.47 ± 3.10
4	-3	-29.85 ± 0.20	-29.92 ± 0.33		-29.93 ± 0.33	-29.98 ± 0.69		-30.00 ± 0.69
4	-2	5.24 ± 0.20	4.68 ± 1.03		4.82 ± 1.04	7.65 ± 2.07		7.97 ± 2.09
4	-1	11.38 ± 0.20	11.66 ± 0.28		11.66 ± 0.28	11.32 ± 0.58		11.26 ± 0.58
4	0	18.81 ± 0.23	15.95 ± 1.10		15.67 ± 1.11	15.47 ± 2.20		14.87 ± 2.23
4	1	-8.95 ± 0.20	-9.04 ± 0.32		-9.05 ± 0.32	-7.57 ± 0.88		-7.66 ± 0.89
4	2	11.75 ± 0.20	13.09 ± 1.26		13.20 ± 1.27	10.98 ± 2.52		11.45 ± 2.57
4	3	10.23 ± 0.20	9.89 ± 0.25		9.89 ± 0.25	8.42 ± 0.74		8.45 ± 0.75
4	4	-2.46 ± 0.20	-0.44 ± 1.53		-0.44 ± 1.54	3.64 ± 3.10		3.67 ± 3.14
5	-5	25.01 ± 0.20	26.42 ± 1.61		26.53 ± 1.61	26.81 ± 3.20		27.09 ± 3.22
5	-4	11.98 ± 0.20	11.49 ± 0.34		11.48 ± 0.34	11.95 ± 0.77		11.84 ± 0.77
5	-3	7.70 ± 0.20	5.98 ± 1.21		5.72 ± 1.22	7.05 ± 2.34		6.19 ± 2.38
5	-2	2.46 ± 0.20	2.78 ± 0.31		2.78 ± 0.31	2.00 ± 0.66		2.04 ± 0.66
5	-1	5.31 ± 0.21	3.46 ± 1.01		3.49 ± 1.01	4.31 ± 2.01		4.46 ± 2.02
5	0	15.61 ± 0.23	15.34 ± 0.34		15.35 ± 0.34	16.46 ± 0.83		16.44 ± 0.83
5	1	30.49 ± 0.21	34.32 ± 1.18		34.40 ± 1.18	31.00 ± 2.30		31.23 ± 2.35
5	2	-11.21 ± 0.20	-11.76 ± 0.32		-11.78 ± 0.32	-12.68 ± 0.95		-12.73 ± 0.96
5	3	26.11 ± 0.20	25.84 ± 1.15		25.82 ± 1.16	26.67 ± 2.28		26.81 ± 2.33
5	4	-6.84 ± 0.20	-7.14 ± 0.28		-7.15 ± 0.28	-5.49 ± 0.83		-5.55 ± 0.83
5	5	14.00 ± 0.20	12.74 ± 1.53		12.76 ± 1.54	14.22 ± 3.10		14.29 ± 3.12

the cost of larger error bars and more leakage from non-included multipoles.

The disklike geometry of our Galaxy provides some intuition for the behavior of the error bars. Because the Galaxy cuts are approximately symmetric under reflection (parity-even), different ℓ -values couple mainly if they are separated by an even number. To the extent that the Galaxy cut is azimuthally symmetric (a crude approximation at best), different m -values do not couple. For example, the quadrupole estimator \hat{a}_{21} therefore picks up noise mainly from a_{41} , a_{61} , etc. Moreover, since spherical harmonics with high $|m|$ live mainly near the Galactic equator and rapidly approach zero towards the Galactic poles, they are the ones that suffer most as the cut is increased and therefore have the largest error bars in Table I.

The above-mentioned sample variance considerations also shed light on how the method is affected by changes in the \mathbf{C} -matrix, i.e., on how robust our method is to assumptions about unwanted noise and CMB signals. Since the method is required to faithfully (without bias) recover the multipoles \mathbf{a} whatever they are, one would intuitively expect that \mathbf{W} and therefore $\hat{\mathbf{a}}$ are independent of contributions from signal and noise in these multipoles to the \mathbf{C} -matrix. With some matrix algebra, one can prove that this is indeed the case, specifically that adding to \mathbf{C} a contribution of the form $\mathbf{Y}\Sigma_S\mathbf{Y}^t$ for some matrix Σ_S leaves \mathbf{W} and $\hat{\mathbf{a}}$ unchanged and simply increases Σ by Σ_S . For example, if we use our method to measure the five quadrupole coefficients, the result $\hat{\mathbf{a}}$ will be independent of our assumptions about the true value of C_2 .

TABLE II. Same as previous table, but for the 3-year WMAP data.

		Mask 0		Mask 5		Mask 6		
ℓ	m	Fit all	Fit 0, 1, ℓ	Fit 0–3	Fit 0–5	Fit 0, 1, ℓ	Fit 0–3	Fit 0–5
2	-2	-24.29 ± 0.20	-23.06 ± 1.21	-23.14 ± 1.22	-23.20 ± 1.22	-25.88 ± 2.49	-26.09 ± 2.49	-26.20 ± 2.53
2	-1	6.83 ± 0.20	6.70 ± 0.23	6.69 ± 0.23	6.69 ± 0.23	6.65 ± 0.40	6.66 ± 0.40	6.69 ± 0.40
2	0	3.22 ± 0.22	3.88 ± 1.11	3.94 ± 1.11	4.11 ± 1.12	3.79 ± 2.29	3.86 ± 2.30	4.14 ± 2.34
2	1	0.37 ± 0.20	0.51 ± 0.26	0.51 ± 0.26	0.52 ± 0.26	-0.62 ± 0.59	-0.62 ± 0.59	-0.54 ± 0.60
2	2	-21.06 ± 0.20	-19.54 ± 1.48	-19.52 ± 1.48	-19.72 ± 1.49	-15.31 ± 3.10	-15.18 ± 3.11	-16.11 ± 3.18
3	-3	45.13 ± 0.20	47.03 ± 1.49	47.06 ± 1.50	47.24 ± 1.51	44.50 ± 3.00	44.59 ± 3.00	44.97 ± 3.06
3	-2	2.36 ± 0.20	2.11 ± 0.27	2.11 ± 0.27	2.10 ± 0.27	2.62 ± 0.50	2.62 ± 0.50	2.59 ± 0.50
3	-1	8.44 ± 0.20	10.58 ± 1.03	10.55 ± 1.03	10.46 ± 1.04	8.72 ± 2.10	8.66 ± 2.10	8.07 ± 2.14
3	0	-8.45 ± 0.22	-8.27 ± 0.29	-8.27 ± 0.29	-8.28 ± 0.29	-9.52 ± 0.65	-9.50 ± 0.65	-9.54 ± 0.66
3	1	-15.88 ± 0.20	-15.97 ± 1.21	-16.00 ± 1.21	-16.24 ± 1.22	-11.94 ± 2.47	-12.03 ± 2.48	-12.86 ± 2.55
3	2	32.65 ± 0.20	32.96 ± 0.26	32.96 ± 0.26	32.96 ± 0.26	33.71 ± 0.72	33.70 ± 0.72	33.61 ± 0.73
3	3	-16.76 ± 0.20	-18.59 ± 1.41	-18.55 ± 1.41	-18.52 ± 1.42	-20.71 ± 2.95	-20.60 ± 2.96	-20.79 ± 3.04
4	-4	-12.24 ± 0.20	-14.39 ± 1.51		-14.81 ± 1.52	-11.83 ± 3.07		-13.10 ± 3.10
4	-3	-27.82 ± 0.20	-27.77 ± 0.33		-27.79 ± 0.33	-27.60 ± 0.69		-27.62 ± 0.69
4	-2	5.74 ± 0.20	4.79 ± 1.03		4.92 ± 1.04	7.10 ± 2.07		7.43 ± 2.09
4	-1	9.79 ± 0.20	9.99 ± 0.28		10.00 ± 0.28	9.89 ± 0.58		9.84 ± 0.58
4	0	17.25 ± 0.23	16.71 ± 1.10		16.43 ± 1.11	16.62 ± 2.20		16.02 ± 2.23
4	1	-14.43 ± 0.20	-14.63 ± 0.32		-14.64 ± 0.32	-12.94 ± 0.88		-13.02 ± 0.89
4	2	16.33 ± 0.20	15.01 ± 1.26		15.12 ± 1.27	12.45 ± 2.52		12.89 ± 2.57
4	3	12.17 ± 0.20	11.91 ± 0.25		11.90 ± 0.25	9.89 ± 0.74		9.92 ± 0.75
4	4	-5.72 ± 0.20	-3.22 ± 1.53		-3.23 ± 1.54	1.14 ± 3.10		1.19 ± 3.14
5	-5	27.10 ± 0.20	28.39 ± 1.61		28.52 ± 1.61	28.69 ± 3.20		29.01 ± 3.22
5	-4	11.39 ± 0.20	10.89 ± 0.34		10.88 ± 0.34	11.41 ± 0.77		11.28 ± 0.77
5	-3	3.89 ± 0.20	2.69 ± 1.21		2.40 ± 1.22	4.57 ± 2.34		3.63 ± 2.38
5	-2	2.56 ± 0.20	2.89 ± 0.31		2.90 ± 0.31	2.21 ± 0.66		2.26 ± 0.66
5	-1	0.90 ± 0.21	-1.15 ± 1.01		-1.14 ± 1.01	0.98 ± 2.01		1.09 ± 2.02
5	0	15.78 ± 0.23	15.45 ± 0.34		15.45 ± 0.34	16.98 ± 0.83		16.94 ± 0.83
5	1	34.10 ± 0.21	34.50 ± 1.18		34.58 ± 1.18	31.71 ± 2.30		31.94 ± 2.35
5	2	-9.74 ± 0.20	-10.19 ± 0.32		-10.21 ± 0.32	-11.06 ± 0.95		-11.10 ± 0.96
5	3	24.12 ± 0.20	25.59 ± 1.15		25.55 ± 1.16	26.41 ± 2.28		26.53 ± 2.33
5	4	-5.72 ± 0.20	-5.95 ± 0.28		-5.97 ± 0.28	-3.74 ± 0.83		-3.83 ± 0.83
5	5	13.47 ± 0.20	12.37 ± 1.53		12.41 ± 1.54	13.25 ± 3.10		13.37 ± 3.12

C. Simulations with foregrounds

The key remaining issue is how to quantify the contribution of residual foreground contamination and how to optimize the Galaxy cut to minimize this contribution. We will now address this issue with simulations.

The residual foreground contamination present after foreground cleaning can be quite important. All five linearly cleaned WMAP-maps used in the literature have their problems. The WMAP team ILC map [9] comes with the disclaimer that it should not be used for scientific analysis—needless to say, the science questions are so interesting that this has not stopped large numbers of groups from using it anyway, particularly when all-sky coverage was helpful. The TOH cleaned WMAP map [13] has the advantage of having lower total noise and residual foregrounds, but like the ILC map, it lacks error bars accurately quantifying these residuals, as do all other

foreground-cleaned maps published to date. The SMICA map [111] is generated with a method very similar that used for the TOH map [13,112], and the result is indeed encouragingly similar [111]. The WI-FIT map of [113] is also closely related (since this method cleans with linear combinations of linear combinations (pair-wise differences) of maps, it too is a linear and WMAP-internal technique just like the ILC, TOH and SMICA methods—the only fundamental difference between the four methods is that the approximate variance minimization is performed in pixel space for ILC, in (masked) spherical harmonic space for TOH and SMICA and in wavelet space for WI-FIT (and optionally SMICA). Finally, using external template maps alone [9] suffers from both dubious extrapolation and from the problem that they (e.g., the Haslam-based synchrotron map) have striping or other systematic errors at a substantially larger level than the

WMAP maps, and that these errors will propagate into the cleaned map, perhaps making it less clean than it was to start with in the most systematics-affected modes—template-cleaned residual WMAP foregrounds have never been accurately quantified in the literature.

A useful recent approach has been to quantify ILC errors by reapplying the ILC cleaning algorithm to simulated sky maps where the correct answer is known and the residuals can thus be computed [114]. This has confirmed that the residual foreground contribution is substantial. We will adopt a similar procedure for the TOH map, but with an important difference that makes our approach more conservative.

1. Mock foregrounds

It is absolutely crucial to use a foreground model that does not artificially mimic the cleaning algorithm. For example, as acknowledged in Eriksen *et al.* (2005) [114], their simulations underestimate the residual foregrounds by assuming that the frequency dependence of each physical component is the same in all directions in the sky—this assumption of perfect frequency coherence [115] is empirically known to be incorrect for synchrotron radiation (see, e.g., [116] for a recent treatment), and allows the simulated foregrounds to be perfectly removed by the cleaning method. To avoid reaching overly optimistic conclusions, we therefore make simulated foreground templates at the five WMAP frequencies in a different way: straight from data by subtracting the TOH CMB maps from the five WMAP frequency channels. The resulting five foreground maps that we use are shown in Fig. 1 of [13]. This procedure is probably overly conservative because some foregrounds and noise get double-counted: for instance, our *K*-band “foreground” map includes detector noise from all five WMAP frequencies because it was constructed by subtracting the TOH map (with its noise) from the *K*-band observations.

2. Mock CMB

We generate a simulated CMB-only sky map (Fig. 1, top right) by using the HEALPix software and the best fit WMAP power spectrum from [1], smoothed to have the same angular resolution as the TOH map (corresponding to the beam of the W-band, the WMAP channel with the highest resolution). As input to the cleaning algorithm, we also smoothed it to the four lower resolutions corresponding to the *K*-, *Ka*-, *Q*- and *W*-band beams [117].

3. Mock detector noise

We construct simulated WMAP noise maps at the five channels by assuming that the noise is Gaussian and uncorrelated both between different pixels and between different channels, which is a good approximation according to [9]. We set the noise amplitude in each map so that the

rms fluctuation levels match those estimated from the WMAP team (Table 5 in [118]). We confirmed that this noise normalization matched both the rms in various same-frequency difference maps such as $W = (W1 - W2 + W3 - W4)/4$ and the total power spectrum normalization at very high ℓ where CMB and foregrounds are negligible. Specifically, we used rms noise values $N_{ii}^{1/2}$ of 150 μK , 132 μK , 122 μK , 149 μK and 179 μK for *K*, *Ka*, *Q*, *V* and *W*, respectively. This corresponds to an ℓ -independent noise power spectrum of order $C_\ell \sim 0.1 \mu\text{K}^2$ in each channel, i.e., orders of magnitude below the large-scale CMB signal ([13] estimates $C_2 \sim 200 \mu\text{K}^2$ and $C_2 \sim 500 \mu\text{K}^2$), so the details of our noise modeling have no impact whatsoever on our results—rather, our uncertainties are dominated by sample variance and residual foregrounds. We tried both crude noise maps with equal variance in all pixels and more accurate ones with the variance modulated in inverse proportion to the number of WMAP observations of each pixel. As expected from the fact that noise is negligible at the low ℓ -values of interest here, these two approaches give for all practical purposes indistinguishable results—we adopt the former for simplicity.

4. Residuals

For each of the five WMAP channels, we sum the mock CMB, noise and foreground maps described above. We then apply the TOH foreground cleaning method to these five input maps exactly as described in [13] to obtain a single foreground-cleaned map. Finally, we subtract the “true” simulated CMB sky map from the cleaned map, obtaining the residual map shown in Fig. 1 (bottom left). This figure visually illustrates that the TOH method provides an unbiased recovery of the CMB, since the residual map shows no patterns whatsoever in common with the input CMB map (upper right). Comparing the residual map with the *V*-band map (top left) also illustrates that the TOH method reduces the levels of both foregrounds and detector noise. Nonetheless, residual foregrounds are clearly visible, particularly near the Galactic plane in the parts that the “junk map” (bottom right) suggests are highly contaminated.

If we knew that this was the *actual* residual in our TOH map, we could of course eliminate the foreground problem completely by subtracting it out. This is unfortunately not the case. To be conservative, will make no attempt to further clean our input map, and will merely use the residual map to quantify the errors in our multipole measurements.

What does the residual map depend on? By repeating the entire procedure with independent CMB simulations, we found that the residual map was rather independent of what CMB map was used. This is expected, since the TOH method is completely CMB-independent except for the small effect that CMB sample variance can have on the computation of optimal weights for the cleaning process.

Repeating the entire procedure with independent noise simulations alters the residual map mainly on the very smallest angular scales. The large-scale patterns visible in Fig. 1 (bottom left) are therefore determined mainly by properties of the foregrounds themselves. For example, foregrounds whose spectral index varies notably across the sky cannot be perfectly subtracted out by the TOH method. Since these large-scale residual patterns are all that matter for our present focus on low ℓ , there is no point in making more than one simulation of our type, i.e., it is pointless to make an ensemble of simulations with different CMB and noise realizations.

D. Optimizing the galaxy cut

To quantify the effect of residual foregrounds on our multipole measurements, we repeat the analysis from Sec. III using our simulated maps. Specifically, we measure the $a_{\ell m}$ -coefficients from the simulated CMB-only map, from a map with CMB and noise (made without introducing any foregrounds before the cleaning process) and from the cleaned map including CMB, noise and foregrounds. The results are shown in Fig. 3 with the correct (simulated) $a_{\ell m}$ -values subtracted off.

The fact that the black and red curves are almost identical confirms the above-mentioned claim that WMAP detector noise makes essentially no difference on these large angular scales. The fact that the red curves agree well with the shaded regions confirms that our method and our software implementing it are working as they should, i.e., that the multipoles are faithfully recovered with accuracy consistent with the predicted noise + leakage error bars $\Delta \hat{a}_{\ell m}$.

The key new information in Fig. 3 is contained in the blue curves, which include errors caused by residual foregrounds. These curves reveal two interesting facts:

- (1) Residual foreground contamination near the Galactic plane afflicts the a_{20} -component at a level around ten μK .
- (2) Residual foreground contamination at higher latitudes appears to afflict most coefficients at the level of a few μK .

The fact that a_{20} is special follows from the well-known fact that this is the component most similar in shape to the Galactic emission, as pointed out already in the COBE/DMR analysis [8]. Specifically, in the very crude approximation that the Galactic plane has both parity and azimuthal symmetry, it would contaminate only the multiple coefficients a_{20} , a_{40} , a_{60} , etc. Figure 3 shows that most of this residual contamination comes from the innermost parts of the plane, so that cutting out merely the dirtiest 7% of the sky (Mask 6) cuts the foreground contribution to a_{20} from about 9 μK to 3 μK .

In contrast, and in good agreement with these symmetry considerations, none of the other quadrupole or octopole coefficients show evidence of contamination from the inner




FIG. 3 (color online). The left and right panels show how the measurement errors $\Delta a_{\ell m} \equiv \hat{a}_{\ell m} - a_{\ell m}^{\text{true}}$ in μK for the quadrupole (left) and octopole (right) extracted from a simulated map change with increasing Galactic cut. The three curves in each panel correspond to errors from CMB multipole leakage alone (black) CMB plus detector noise (red) and CMB plus noise plus residual foregrounds (blue). The shaded/green areas show the 1σ error bars corresponding to leakage and noise. Note that these foreground-related errors are typically an order of magnitude smaller than the expected CMB signal (the ΛCDM concordance model gives $C_2^{1/2} = 30 \mu\text{K}$ and $C_3^{1/2} = 20 \mu\text{K}$), and that they are even more subdominant for power spectrum estimation where they add only in quadrature.

Galactic plane: the foreground contributions shown by the blue curves do not tend to shift closer to the zero as the cut is increased up to Mask 6. Beyond this, the leakage errors are seen to exceed the residual foreground errors. This small residual foreground contamination is seen to typically be 1–3 μK , of similar magnitude to the a_{20} -contamination that remains for Mask 6.

So based on these results, what is the optimal galaxy cut when measuring multipole coefficients? We have seen that too small a Galaxy cut leads to unnecessarily high foreground contamination, whereas too large a cut leads to unnecessarily high leakage from other multipoles. The natural optimum is seen to be Mask 5 or 6, where the typical errors $\Delta a_{\ell m}$ are of order a μK from leakage and a few μK from foregrounds, and we have therefore listed results for both in Table I. For the more aggressive Mask 7, sample variance strongly dominates, and for the less aggressive Mask 4, the a_{20} foreground contamination is noticeably larger. If the reader wishes to perform an analysis that in no way involves leakage from the black sheep a_{20} , another natural choice is Mask 0.




FIG. 4 (color online). Effects of pixelization. The panels show how the measurement errors $\Delta a_{\ell m} \equiv \hat{a}_{\ell m} - a_{\ell m}^{\text{true}}$ in μK extracted from a simulated map change with increasing Galactic cut. The black, blue and red curves correspond to Healpix pixelization at angular resolution level $n_{\text{side}} = 4, 8$ and 16 , respectively. The shaded/green bands show the 1σ errors from detector noise and leakage.

E. Optimizing the pixel size

This subsection discusses practical details useful for readers interested in applying this method. Other readers may wish to skip straight to Sec. IV.

Applying Eqs. (4) and (5) using the native WMAP pixel resolution ($n_{\text{side}} = 512$) is unfeasible, since it would involve the inversion of an $N \times N$ -dimensional \mathbf{C} -matrix

with N of order 3×10^6 . Since we are only interested in the lowest multipoles, there are fortunately two simple ways to overcome this problem. The first approach involves using matrix identities to transform the problem into one involving smaller matrices (as for the so-called Woodbury formula). The second approach, which we used above, is to simply smooth the map onto a coarser pixelization and correct measurements for this.

Figure 4 shows the errors on the recovered quadrupole coefficients from our simulated map (with neither detector noise nor foregrounds) for pixelizations with $n_{\text{side}} = 4, 8$ and 16 , showing that as long as we work at resolution 16 , pixelization errors are negligibly small compared to noise and sample variance for the masks of interest.

This pixelization corresponds to averaging the underlying sky map across the area of each pixel. If the map were subjected to Gaussian smoothing, then the multipoles $a_{\ell m}$ would be suppressed by the well-known factor $e^{\theta^2 \ell(\ell+1)/2}$, where $\theta = \text{FWHM}/\sqrt{8 \ln 2}$ and FWHM is the full-width half-maximum of the Gaussian smoothing kernel. For our adopted resolution level 16 , the pixels are approximately squares of side 3.7° . A very crude estimate therefore suggests that the corresponding multipole suppression is of order $e^{-(\ell/40)^2/2}$, i.e., a 0.3% quadrupole suppression and a 2% suppression for $\ell = 5$. Since this smoothing effect is negligible compared to the measurement errors caused by leakage and foregrounds, we have not corrected for it in Table I.

In order to quantify the errors due to pixelization, we calculate the five components of the quadrupole for different n_{side} s of $4, 8$ and 16 , and also studied how these values change as we increase the mask number (or increase the Galactic cut).

Although resolution 8 already gives us a good estimate of the lower multipoles, we decided (to be conservative) to use $n_{\text{side}} = 16$ in all calculations done below.




FIG. 5 (color online). Change in the direction of the quadrupole and octopole “axis of evil” directions as a function of the mask number. The two crosses correspond to the true quadrupole and octopole axes in our CMB simulation, chosen to match the WMAP measured values. The departures of the shaded dots from the corresponding crosses reflect the effect of noise, multipole leakage and foreground contamination.




FIG. 6 (color online). The top panel shows the separation angle (in degrees) between the quadrupole and octopole vectors and how they change as the mask number increases. The lower panel shows the probability for this alignment to happen. In both figures the red, blue and green lines represent the coefficient values for a map with CMB only, CMB plus noise, and CMB plus noise plus residual foreground, respectively.

IV. IMPLICATIONS FOR THE QUADRUPOLE-OCTOPOLE ALIGNMENT

As a simple example of an application of our results, we use our new multipole measurements to revisit the significance of the WMAP quadrupole-octopole alignment. The two crosses in Fig. 5 show the “axis of evil” from [13,14],

extracted from a simulated CMB map where we inserted the exact same multipole and quadrupole coefficients as measured from WMAP. We then added noise and foregrounds as above and extracted the “axis of evil” for our galactic cut. Confirming the conclusion of [22], the octopole is seen to be quite robust, whereas the quadrupole moves around somewhat more (it is clearly more fragile due to its intrinsically lower amplitude). As seen in Fig. 6, this causes the apparent measured alignment to be somewhat less significant than the true one, but makes no dramatic difference. Similarly, we find that replacing Mask 0 by Mask 6 (the joint quadrupole/octopole fit in Table I) degrades the alignment significance only slightly, from a 1-in-60 fluke to a 1-in-40 fluke.

V. CONCLUSIONS

We have presented a minimum-variance method for measuring the CMB multipole coefficients $a_{\ell m}$ given anisotropic noise, incomplete sky coverage and foreground contamination. Our method constitutes lossless data compression in the sense that the widely used quadratic estimators of the power spectrum C_ℓ can be computed directly from our $a_{\ell m}$ -estimators. We illustrated the method by applying it to the WMAP data. As the Galactic cut is increased, the error bars $\Delta a_{\ell m}$ on low multipoles go from being dominated by foregrounds to being dominated by sample variance from other multipoles, with the intervening minimum defining the optimal cut.

Because WMAP detector noise is negligible on these large angular scales, one would not expect the improved sensitivity of second WMAP data release to significantly modify our results *per se*. Comparing Table I and II shows that the multipoles that change significantly are those most sensitive to foregrounds, notable a_{20} , suggesting the interpretation that the lower noise levels improved the fore-




FIG. 7 (color online). Changes in the direction of the preferred quadrupole and octopole axes as a function of mask number for the WMAP 1-year (left) and 3-year (right) data. In both panels, the quadrupole and octopole results are represented by balls and stars, respectively, and the colors black, cyan, red, blue, magenta, green and yellow correspond to the results for masks from 0 to 6, respectively. The quadrupole axis is seen to be more robust in the 3-year data, suggesting that the lower noise level may have improved the performance of the foreground cleaning process, reducing the level of residual low-latitude foreground contamination.

ground cleaning. The corresponding mask 6 quadrupole and octopole “axes of evil” change from $(l, b) = (251.5^\circ, 52.1^\circ)$ and $(238.4^\circ, 63.0^\circ)$ for WMAP1 to $(l, b) = (222.4^\circ, 75.5^\circ)$ and $(236.4^\circ, 64.9^\circ)$ for WMAP3, respectively. In other words, the octopole axis is essentially unchanged whereas the quadrupole axis moves to the other side of the octopole axis, getting still closer (from 13.0° off to 11.5°), requiring a 1-in-50 fluke—see Fig. 7. Moving the quadrupole axis in a random direction would typically *degrade* the alignment, so this is consistent with the hypothesis that the true alignment is still better and that this is partially masked by foregrounds.

Improved sensitivity from WMAP and Planck and also new cleaning techniques will hopefully allow better removal and quantification of residual foreground contamination, which would reduce the main source of uncertainty

in our measurements and shed further light on the low- ℓ CMB puzzles.

ACKNOWLEDGMENTS

We thank the WMAP team for producing such a superb data set and for making it public available from <http://lambda.gsfc.nasa.gov>. We thank Krzysztof Górski and collaborators for creating the HEALPIX package which we used both for harmonic transforms and map plotting. This work was supported by NSF Grants AST-0607597 and AST-0134999 and NASA Grant No. NAG5-11099. MT was supported by the David and Lucile Packard Foundation and a Cottrell Scholarship from Research Corporation.

-
- [1] D.N. Spergel, L. Verde, H.V. Peiris, E. Komatsu, M.R. Nolta, C.L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S.S. Meyer, L. Page, G.S. Tucker, J.L. Weiland, E. Wollack, and E.L. Wright, *Astrophys. J. Suppl. Ser.* **148**, 175 (2003).
- [2] M. Tegmark, M. Strauss, M.R. Blanton, K. Abazajian, S. Dodelson, H. Sandvik, X. Wang, D.H. Weinberg, I. Zehavi, N. Bahcall, F. Hoyle, D. Schlegel, R. Scoccimarro, M. Vogeley, A. Berlind, T. Budavari, A. Connolly, D. Eisenstein, D. Finkbeiner, J.A. Frieman, J. Gunn, L. Hui, B. Jain, D. Johnston, S. Kent, H. Lin, R. Nakajima, R. Nichol, J. Ostriker, A. Pope, R. Scranton, U. Seljak, R. Sheth, A. Stebbins, A. Szalay, I. Szapudi, Y. Xu, J. Annis, J. Brinkmann, S. Burles, F. Castander, I. Csabai, J. Loveday, M. Doi, M. Fukugita, G. Hennessy, D. Hogg, Z. Ivezic, G. Knapp, D. Lamb, B. Lee, R. Lupton, T. McKay, P. Kunszt, J. Munn, L. O’Connell, J. Peoples, J. Pier, M. Richmond, C. Rockosi, D. Schneider, C. Stoughton, D. Tucker, D. Vanden Berk, B. Yanny, and D. York, *Phys. Rev. D* **69**, 103501 (2004).
- [3] U. Seljak, A. Makarov, P. McDonald, S. Anderson, N. Bahcall, J. Brinkmann, S. Burles, R. Cen, M. Doi, J. Gunn, Z. Ivezic, S. Kent, R. Lupton, J. Munn, R. Nichol, J. Ostriker, D. Schlegel, M. Tegmark, D. Van den Berk, D. Weinberg, and D. York, *Phys. Rev. D* **71**, 103515 (2005).
- [4] S. Cole, W.J. Percival, J.A. Peacock, P. Norberg, C.M. Baugh, C.S. Frenk, I. Baldry, J. Bland-Hawthorn, T. Bridges, R. Cannon, M. Colless, C. Collins, W. Couch, N.J.G. Cross, G. Dalton, V.R. Eke, R. De Propris, S.P. Driver, G. Efstathiou, R.S. Ellis, K. Glazebrook, C. Jackson, A. Jenkins, O. Lahav, I. Lewis, S. Lumsden, S. Maddox, D. Madgwick, B.A. Peterson, W. Sutherland, and K. Taylor, *Mon. Not. R. Astron. Soc.* **362**, 505 (2005).
- [5] C. MacTavish, P. Ade, J. Bock, J. Bond, J. Borrill, A. Boscaleri, P. Cabella, C. Contaldi, B. Crill, P. de Bernardis, G. De Gasperis, A. de Oliveira-Costa, G. De Troia, G. Di Stefano, E. Hivon, A. Jaffe, W.C. Jones, T. Kisner, A. Lange, S. Masi, P. Matuszewska, A. Melchiorri, T. Montroy, P. Natoli, B. Netterfield, E. Pascale, F. Piacentini, D. Pogosyan, G. Polenta, S. Prunet, S. Ricciardi, G. Romeo, J. Ruhl, P. Santini, M. Tegmark, M. Veneziani, and N. Vittorio (unpublished).
- [6] E. Komatsu, A. Kogut, M. Nolta, C.L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, M. Limon, S.S. Meyer, L. Page, D.N. Spergel, G.S. Tucker, L. Verde, E. Wollack, and E.L. Wright, *Astrophys. J. Suppl. Ser.* **148**, 119 (2003).
- [7] P. Creminelli, A. Nicolis, L. Senatore, M. Tegmark, and M. Zaldarriaga, astro-ph/0509029.
- [8] G.F. Smoot, C.L. Bennett, A. Kogut, E.L. Wright, J. Aymon, N.W. Boggess, E.S. Cheng, G. de Amici, S. Gulkis, M.G. Hauser, G. Hinshaw, P.D. Jackson, M. Janssen, E. Kaita, T. Kelsall, P. Keegstra, C. Lineweaver, K. Loewenstein, P. Lubin, J. Mather, S.S. Meyer, S.H. Moseley, T. Murdock, L. Rokke, R.F. Silverberg, L. Tenorio, R. Weiss, and T. Wilkinson, *Astrophys. J.* **396**, L1 (1992).
- [9] C.L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S.S. Meyer, L. Page, D.N. Spergel, G.S. Tucker, E. Wollack, E.L. Wright, C. Barnes, M.R. Greason, R.S. Hill, E. Komatsu, M.R. Nolta, N. Odegard, H.V. Peirs, L. Verde, and J.L. Weiland, *Astrophys. J. Suppl. Ser.* **148**, 1 (2003).
- [10] E. Gaztañaga, J. Wagg, T. Multamaki, A. Montaña, and D.H. Hughes, *Mon. Not. R. Astron. Soc.* **346**, 47 (2003).
- [11] G. Efstathiou, *Mon. Not. R. Astron. Soc.* **346**, L26 (2003).
- [12] A. Slosar and U. Seljak, *Phys. Rev. D* **70**, 083002 (2004).
- [13] M. Tegmark, Max A. de Oliveira-Costa, and A.J. Hamilton, *Phys. Rev. D* **68**, 123523 (2003).
- [14] A. de Oliveira-Costa, M. Tegmark, M. Zaldarriaga, and A.J. Hamilton, *Phys. Rev. D* **69**, 063516 (2004).
- [15] H.K. Eriksen, F.K. Hansen, A.J. Banday, K.M. Gorski, and P.B. Lilje, *Astrophys. J.* **605**, 14 (2004).
- [16] C.J. Copi, D. Huterer, and G.D. Starkman, *Phys. Rev. D*, **70**, 043515 (2004).

- [17] F.K. Hansen, A.J. Banday, and K.M. Gorski, *astro-ph/0404206*.
- [18] K. Land and J. Magueijo, *Mon. Not. R. Astron. Soc.*, **357**, 994 (2005).
- [19] J.R. Weeks, *astro-ph/0412231*.
- [20] K. Land and J. Magueijo, *Phys. Rev. Lett.* **95**, 071301 (2005).
- [21] K. Land and J. Magueijo, *Mon. Not. R. Astron. Soc.* **362**, 838 (2005).
- [22] P. Bielewicz, H.K. Eriksen, A.J. Banday, K.M. Gorski, and P.B. Lilje, *Astrophys. J.* (2005).
- [23] K. Land, and J. Magueijo, *Phys. Rev. D* **72**, 101302 (2005).
- [24] C.J. Copi, D. Huterer, and G.D. Starkman, *Mon. Not. R. Astron. Soc.* **367**, 79 (2006).
- [25] K. Land and J. Magueijo, *Mon. Not. R. Astron. Soc.* **367**, 1714 (2006).
- [26] D.J. Schwarz, G.D. Starkman, D. Huterer, and C.J. Copi, *Phys. Rev. Lett.* **93**, 221301 (2004).
- [27] C.-G. Park, *Mon. Not. R. Astron. Soc.* **349**, 313 (2004).
- [28] P. Bielewicz, K.M. Gorski, and A.J. Banday, *Mon. Not. R. Astron. Soc.* **355**, 1283 (2004).
- [29] P.D. Naselsky, L.-Y. Chiang, P. Olesen, and O.V. Verkhodanov, *Astrophys. J.* **615**, 45 (2004).
- [30] A. Hajian, T. Souradeep, and N. Cornish, *Astrophys. J.* **618**, L63 (2005).
- [31] S. Prunet, J.-P. Uzan, F. Bernardeau, and T. Brunier, *Phys. Rev. D* **71**, 083508 (2005).
- [32] H.K. Eriksen, A.J. Banday, K.M. Gorski, and P.B. Lilje, *Astrophys. J.* **622**, 58 (2005).
- [33] M. Gluck and C. Pisano, *astro-ph/0503442*.
- [34] G. Chen and I. Szapudi, *Astrophys. J.* **635**, 743 (2005).
- [35] P.E. Freeman, C.R. Genovese, C.J. Miller, R.C. Nichol, and L. Wasserman, *Astrophys. J.* **638**, 1 (2006).
- [36] A. Bernui, T. Villela, C.A. Wuensche, R. Leonardi, and I. Ferreira, *astro-ph/0601593*.
- [37] A. Hajian and T. Souradeep, *Astrophys. J.* **597**, L5 (2003).
- [38] O. Dore, G.P. Holder, and A. Loeb, *Astrophys. J.* **612**, 81 (2004).
- [39] J. Portsmouth, *Phys. Rev. D* **70**, 063504 (2004).
- [40] C. Skordis and J. Silk, *astro-ph/0402474*.
- [41] H.K. Eriksen, I.J. O'Dwyer, J.B. Jewell, B.D. Wandelt, D.L. Larson, K.M. Gorski, S. Levin, A.J. Banday, and P.B. Lilje, *Astrophys. J. Suppl. Ser.* **155**, 227 (2004).
- [42] A. Stannard and P. Coles, *Mon. Not. R. Astron. Soc.* **364** (2005) 929.
- [43] N. Seto and E. Pierpaoli, *Phys. Rev. Lett.* **95**, 101302 (2005).
- [44] A. Bernui, B. Mota, M.J. Reboucas, and R. Tavakol, *astro-ph/0511666*.
- [45] H.K. Eriksen, A.J. Banday, K.M. Gorski, and P.B. Lilje, *Astrophys. J.* **612**, 633 (2004).
- [46] P.D. Naselsky, O.V. Verkhodanov, L.-Y. Chiang, and I.D. Novikov, *Int. J. Mod. Phys. D* **14**, 1273 (2005).
- [47] G. Bernardi, E. Carretti, R. Fabbri, C. Sbarra, and S. Cortiglioni, *Mon. Not. R. Astron. Soc.* **364**, L5 (2005).
- [48] J. Medeiros and C.R. Contaldi, *Mon. Not. R. Astron. Soc.* **367**, 39 (2006).
- [49] L.R. Abramo and L. Sodre, Jr., *astro-ph/0312124*.
- [50] M.O. Calvao, G.I. Gomero, B. Mota, and M.J. Reboucas, *Classical Quantum Gravity* **22**, 1991 (2005).
- [51] K. Tomita, *Phys. Rev. D* **72**, 043526 (2005).
- [52] C. Vale, *astro-ph/0509039*.
- [53] A. Cooray and N. Seto, *J. Cosmol. Astropart. Phys.* **12** (2005) 004.
- [54] A. Racic, S. Rasanen, and D.J. Schwarz, *astro-ph/0601445*.
- [55] S. Burles and S. Rappaport, *astro-ph/0601559*.
- [56] K. T. Inoue and J. Silk, *astro-ph/0602478*.
- [57] B. McInnes, *J. High Energy Phys.* **09** (2003) 009.
- [58] B. Mota, G.I. Gomero, M.J. Reboucas, and R. Tavakol, *Classical Quantum Gravity* **21**, 3361 (2004).
- [59] N.J. Cornish, D.N. Spergel, G.D. Starkman, and E. Komatsu, *Phys. Rev. Lett.* **92**, 201302 (2004).
- [60] G.I. Gomero, *astro-ph/0310749*.
- [61] B.F. Roukema, B. Lew, M. Cechowska, A. Marecki, and S. Bajtlik, *Astron. Astrophys.* **423**, 821 (2004).
- [62] J. Levin, *Phys. Rev. D* **70**, 083001 (2004).
- [63] A. Hajian and T. Souradeep, *astro-ph/0501001*.
- [64] B. Mota, M.J. Reboucas, and R. Tavakol, *Int. J. Mod. Phys. A* **20**, 2415 (2005).
- [65] M. Makler, B. Mota, and M.J. Reboucas, *astro-ph/0507116*.
- [66] B. Mota, M. Makler, and M.J. Reboucas, *Astron. Astrophys.* **446**, 805 (2006).
- [67] W.S. Hipolito-Ricaldi and G.I. Gomero, *Phys. Rev. D* **72**, 103008 (2005).
- [68] J.G. Cresswell, A.R. Liddle, P. Mukherjee, and A. Riazuelo, *Phys. Rev. D* **73**, 041302 (2006).
- [69] M.J. Reboucas and J.S. Alcaniz, *astro-ph/0603206*.
- [70] B. Feng, M. Li, R.-J. Zhang, and X. Zhang, *Phys. Rev. D* **68** 103511 (2003).
- [71] Q.-G. Huang and M. Li, *J. Cosmol. Astropart. Phys.* **11** (2003) 001.
- [72] X.-J. Bi, B. Feng, and X. Zhang, *hep-ph/0309195*.
- [73] P.G. Castro, M. Douspis, and P.G. Ferreira, *Phys. Rev. D* **68**, 127301 (2003).
- [74] Y.-S. Piao, B. Feng, and X. Zhang, *Phys. Rev. D* **69**, 103520 (2004).
- [75] C.J. Hogan, *Phys. Rev. D* **70**, 083521 (2004).
- [76] N. Kaloper, *Phys. Lett. B* **583**, 1 (2004).
- [77] M. Liguori, S. Matarrese, M. Musso, and A. Riotto, *J. Cosmol. Astropart. Phys.* **08** (2004) 011.
- [78] C.J. Hogan, *astro-ph/0406447*.
- [79] R.V. Buniy, *Int. J. Mod. Phys. A* **20**, 1095 (2005).
- [80] P. Hunt and S. Sarkar, *Phys. Rev. D* **70**, 103518 (2004).
- [81] A. Linde, *J. Cosmol. Astropart. Phys.* **10** (2004) 004.
- [82] Y.-S. Piao, *Phys. Rev. D* **71**, 087301 (2005).
- [83] R.V. Buniy, A. Berera, and T.W. Kephart, *Phys. Rev. D* **73**, 063529 (2006).
- [84] A. Kostelecky, *Phys. Rev. D* **69**, 105009 (2004).
- [85] T. Multamaki and O. Elgaroy, *Astron. Astrophys.* **423**, 811 (2004).
- [86] S. Hannestad and L. Mersini-Houghton, *Phys. Rev. D* **71**, 123504 (2005).
- [87] C. Armendariz-Picon, *J. Cosmol. Astropart. Phys.* **07** (2004) 007.
- [88] C. Gordon and W. Hu, *Phys. Rev. D* **70**, 083003 (2004).
- [89] J.W. Moffat, *J. Cosmol. Astropart. Phys.* **10** (2005) 012.
- [90] R.A. Battye and A. Moss, *J. Cosmol. Astropart. Phys.* **06** (2005) 001.

- [91] T. R. Jaffe, A. J. Banday, H. K. Eriksen, K. M. Gorski, and F. K. Hansen, *Astrophys. J.* **629**, L1 (2005).
- [92] L. Knox, *Phys. Rev. D* **73**, 023503 (2006).
- [93] C. Gordon and D. Wands, *Phys. Rev. D* **71**, 123505 (2005).
- [94] C. Gordon, W. Hu, D. Huterer, and T. Crawford, *Phys. Rev. D* **72**, 103002 (2005).
- [95] I. Sawicki and S. M. Carroll, astro-ph/0510364.
- [96] P. Jain, M. S. Modgil, and J. P. Ralston, astro-ph/0510803.
- [97] T. Koivisto and D. F. Mota, *Phys. Rev. D* **73**, 083502 (2006).
- [98] K. Subramanian, astro-ph/0601570.
- [99] L. Cayon, A. J. Banday, T. Jaffe, H. K. Eriksen, F. K. Hansen, K. M. Gorski, and J. Jin, astro-ph/0602023.
- [100] R. A. Battye and A. Moss, astro-ph/0602377.
- [101] P. S. Apostolopoulos, N. Brouzakis, N. Tetradis, and E. Tzavara, astro-ph/0603234.
- [102] K. M. Górski, E. Hivon, and B. D. Wandelt, astro-ph/9812350.
- [103] K. M. Górski, B. D. Wandelt, F. K. Hansen, and A. J. Banday, astro-ph/9905275.
- [104] M. Tegmark, *Phys. Rev. D* **56**, 4514 (1997).
- [105] M. Tegmark, *Phys. Rev. D* **55**, 5895 (1997).
- [106] J. R. Bond, A. H. Jaffe, and L. E. Knox, *Astrophys. J.* **533**, 19 (2000).
- [107] G. Hinshaw, M. R. Nolta, C. L. Bennett, R. Bean, O. Dore', M. R. Greason, M. Halpern, R. S. Hill, N. Jarosik, A. Kogut, E. Komatsu, M. Limon, N. Odegard, S. S. Meyer, L. Page, H. V. Peiris, D. N. Spergel, G. S. Tucker, L. Verde, J. L. Weiland, E. Wollack, and E. L. Wright, astro-ph/0603451.
- [108] L. Page, G. Hinshaw, E. Komatsu, M. R. Nolta, D. N. Spergel, C. L. Bennett, C. Barnes, R. Bean, O. Dore', M. Halpern, R. S. Hill, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, N. Odegard, H. V. Peiris, G. S. Tucker, L. Verde, J. L. Weiland, E. Wollack, and E. L. Wright, astro-ph/0603450.
- [109] D. N. Spergel, R. Bean, O. Dore', M. R. Nolta, C. L. Bennett, G. Hinshaw, N. Jarosik, E. Komatsu, L. Page, H. V. Peiris, L. Verde, C. Barnes, M. Halpern, R. S. Hill, A. Kogut, M. Limon, S. S. Meyer, N. Odegard, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright, astro-ph/0603449.
- [110] N. Jarosik, C. Barnes, M. R. Greason, R. S. Hill, M. R. Nolta, N. Odegard, J. L. Weiland, R. Bean, C. L. Bennett, O. Dore', M. Halpern, G. Hinshaw, A. Kogut, E. Komatsu, M. Limon, S. S. Meyer, L. Page, D. N. Spergel, G. S. Tucker, E. Wollack, and E. L. Wright, astro-ph/0603452.
- [111] G. Patanchon, J.-F. Cardoso, J. Delabrouille, and P. Vielva, *Mon. Not. R. Astron. Soc.* **364**, 1185 (2005).
- [112] M. Tegmark and G. Efstathiou, *Mon. Not. R. Astron. Soc.* **281**, 1297 (1996).
- [113] F. K. Hansen, A. J. Banday, H. K. Erikssen, K. M. Górski, and P. B. Lilje, astro-ph/0603308.
- [114] H. K. Eriksen, A. J. Banday, K. M. Gorski, and P. B. Lilje, astro-ph/0508196.
- [115] M. Tegmark, D. J. Eisenstein, W. Hu, and A. de Oliveira-Costa, *Astrophys. J.* **530**, 133 (2000).
- [116] R. D. Davies, C. Dickinson, A. J. Banday, T. R. Jaffe, K. M. Gorski, and R. J. Davis, astro-ph/0511384.
- [117] L. Page, C. Barnes, G. Hinshaw, D. N. Spergel, J. L. Weiland, E. Wollack, C. L. Bennett, M. Halpern, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, and E. L. Wright, *Astrophys. J. Suppl. Ser.* **148**, 39 (2003).
- [118] G. Hinshaw, D. N. Spergel, L. Verde, R. S. Hill, S. S. Meyer, C. Barnes, C. L. Bennett, M. Halpern, N. Jarosik, A. Kogut, E. Komatsu, M. Limon, L. Page, G. S. Tucker, J. Weiland, E. Wollack, and E. L. Wright, *Astrophys. J. Suppl. Ser.* **148**, 135 (2003).