

Seven-Year *Wilkinson Microwave Anisotropy Probe* (*WMAP* ¹) Observations: Are There Cosmic Microwave Background Anomalies?

C. L. Bennett², R. S. Hill³, G. Hinshaw⁴, D. Larson², K. M. Smith⁵, J. Dunkley⁶, B. Gold², M. Halpern⁷, N. Jarosik⁸, A. Kogut⁴, E. Komatsu⁹, M. Limon¹⁰, S. S. Meyer¹¹, M. R. Nolte¹², N. Odegard³, L. Page⁸, D. N. Spergel^{5,13}, G. S. Tucker¹⁴, J. L. Weiland³, E. Wollack⁴, E. L. Wright¹⁵

ABSTRACT

A simple six parameter Λ CDM model provides a successful fit to *WMAP* data. This holds both when the *WMAP* data are analyzed alone or in combination with other cosmological data. Even so, it is appropriate to examine the data carefully to search for hints of deviations from the now standard model of cosmology, which includes inflation, dark energy, dark matter, baryons, and neutrinos. The cosmological community has subjected the *WMAP* data to extensive and varied analyses. While there is widespread agreement as to the overall success of the six parameter Λ CDM model, various “anomalies” have been reported relative to that model. In this paper we examine potential anomalies and present analyses and assessments of their significance. In most cases we find that claimed anomalies depend on posterior selection of some aspect or subset of the data. Compared with sky simulations based on the best fit model, one can select for

¹*WMAP* is the result of a partnership between Princeton University and NASA’s Goddard Space Flight Center. Scientific guidance is provided by the *WMAP* Science Team.

²Dept. of Physics & Astronomy, The Johns Hopkins University, 3400 N. Charles St., Baltimore, MD 21218-2686

³ADNET Systems, Inc., 7515 Mission Dr., Suite A100 Lanham, Maryland 20706

⁴Code 665, NASA/Goddard Space Flight Center, Greenbelt, MD 20771

⁵Dept. of Astrophysical Sciences, Peyton Hall, Princeton University, Princeton, NJ 08544-1001

⁶Astrophysics, University of Oxford, Keble Road, Oxford, OX1 3RH, UK

⁷Dept. of Physics and Astronomy, University of British Columbia, Vancouver, BC Canada V6T 1Z1

⁸Dept. of Physics, Jadwin Hall, Princeton University, Princeton, NJ 08544-0708

⁹Univ. of Texas, Austin, Dept. of Astronomy, 2511 Speedway, RLM 15.306, Austin, TX 78712

¹⁰Columbia Astrophysics Laboratory, 550 W. 120th St., Mail Code 5247, New York, NY 10027-6902

¹¹Depts. of Astrophysics and Physics, KICP and EFI, University of Chicago, Chicago, IL 60637

¹²Canadian Institute for Theoretical Astrophysics, 60 St. George St, University of Toronto, Toronto, ON Canada M5S 3H8

¹³Princeton Center for Theoretical Physics, Princeton University, Princeton, NJ 08544

¹⁴Dept. of Physics, Brown University, 182 Hope St., Providence, RI 02912-1843

¹⁵UCLA Physics & Astronomy, PO Box 951547, Los Angeles, CA 90095-1547

low probability features of the *WMAP* data. Low probability features are *expected*, but it is not usually straightforward to determine whether any particular low probability feature is the result of the *a posteriori* selection or non-standard cosmology. Hypothesis testing could, of course, always reveal an alternate model that is statistically favored, but there is currently no model that is more compelling. We find that two cold spots on the map are statistically consistent with random CMB fluctuations. We also find that the amplitude of the quadrupole is well within the expected 95% confidence range and therefore is not anomalously low. We find no significant anomaly with a lack of large angular scale CMB power for the best-fit Λ CDM model. We examine in detail the properties of the power spectrum data with respect to the Λ CDM model and find no significant anomalies. The quadrupole and octupole components of the CMB sky are remarkably aligned, but we find that this is not due to any single map feature; it results from the statistical combination of the full sky anisotropy pattern. It may be due, in part, to chance alignments between the primary and secondary anisotropy, but this only shifts the coincidence from within the last scattering surface to between it and the local matter density distribution. This alignment has been known for years and yet no theory has replaced Λ CDM as more compelling. We examine claims of a hemispherical or dipole power asymmetry across the sky and find that the evidence for these claim is not statistically significant. We confirm the claim of a strong quadrupolar power asymmetry effect, but there is considerable evidence that the effect is not cosmological. The likely explanation is an insufficient handling of beam asymmetries. We conclude that there is no compelling evidence for deviations from the Λ CDM model, which is generally an acceptable statistical fit to *WMAP* and other cosmological data.

Subject headings: cosmic microwave background, cosmology: observations, early universe, dark matter, space vehicles, space vehicles: instruments, instrumentation: detectors, telescopes

1. Introduction

The *WMAP* mission (Bennett et al. 2003a) was designed to make precision measurements of the CMB to place constraints on cosmology. *WMAP* was specifically designed to minimize systematic measurement errors so that the resulting measurements would be highly reliable within well-determined and well-specified uncertainty levels. The rapidly switched and highly symmetric differential radiometer system effectively makes use of the sky as a stable reference load and renders most potential systematic sources of error negligible. The spacecraft spin and precession paths on the sky create a highly interconnected set of differential data. Multiple radiometers and multiple frequency bands enable checks for systematic effects associated with particular radiometers and frequency dependencies. Multiple years of observations allow for checks of time-dependent systematic errors.

The *WMAP* team has provided the raw time ordered data to the community. It has also made full sky maps from these data, and these maps are the fundamental data product of the mission. If (and only if) the full sky CMB anisotropy represented in a map is a realization of an isotropic Gaussian random process, then the power spectrum of that map contains all of the cosmological information. The maps and cosmological parameter likelihood function based on the power spectrum are the products most used by the scientific community.

The *WMAP* team used realistic simulated time ordered data to test and verify the map-making process. The *WMAP* and *Cosmic Background Explorer (COBE)* maps, produced by independent hardware and with substantially different orbits and sky scanning patterns have been found to be statistically consistent. Freeman et al. (2006) directly verified the fidelity of the *WMAP* team’s map-making process (Hinshaw et al. 2003; Jarosik et al. 2007, 2010). It was indirectly verified by Wehus et al. (2009) as well. Finally, numerous CMB experiments have verified the *WMAP* sky maps over small patches of the sky (mainly with cross-correlation analyses), either to extract signal or to transfer the more precise *WMAP* calibration.

WMAP data alone are consistent with a six parameter inflationary Λ CDM model that specifies the baryon density $\Omega_b h^2$, the cold dark matter density $\Omega_c h^2$, a cosmological constant Ω_Λ , a spectral index of scalar fluctuations n_s , the optical depth to reionization τ , and the scalar fluctuation amplitude $\Delta_{\mathcal{R}}^2$ (Dunkley et al. 2009; Larson et al. 2010). This Λ CDM model is flat, with a nearly (but not exactly) scale-invariant fluctuation spectrum seeded by inflation, with Gaussian random phases, and with statistical isotropy over the super-horizon sky. When *WMAP* data are combined with additional cosmological data, the Λ CDM model remains a good fit, with a narrower range of allowed parameter values (Komatsu et al. 2010). It is remarkable that such diverse observations over a wide range of redshifts are consistent with the standard Λ CDM model.

There are three major areas of future investigation: (1) further constrain allowed parameter ranges; (2) test the standard Λ CDM model against data to seek reliable evidence for flaws; and (3) seek the precise physical nature of the components of the Λ CDM model: cold dark matter, inflation, and dark energy. It is the second item that we address here: are there potential deviations from Λ CDM within the context of the allowed parameter ranges of the existing *WMAP* observations?

A full sky sky map $T(\mathbf{n})$ may be decomposed into spherical harmonics Y_{lm} as

$$T(\mathbf{n}) = \sum_{l=0}^{\infty} \sum_{m=-l}^l a_{lm} Y_{lm}(\mathbf{n}) \quad (1)$$

with

$$a_{lm} = \int d\mathbf{n} T(\mathbf{n}) Y_{lm}^*(\mathbf{n}) \quad (2)$$

where \mathbf{n} is a unit direction vector. If the CMB anisotropy is Gaussian-distributed with random phases, then each a_{lm} is independent, with a zero-mean $\langle a_{lm} \rangle = 0$ Gaussian distribution with

$$\langle a_{lm} a_{l'm'}^* \rangle = \delta_{ll'} \delta_{mm'} C_l \quad (3)$$

where C_l is the angular power spectrum and δ is the Kronecker delta. C_l is the mean variance per multipole moment l that would be obtained if one could take and average measurements from every vantage point throughout the universe. We have only our one sample of the universe, however, and its spectrum is related to the measured a_{lm} coefficients by

$$C_l^{\text{sky}} = \frac{1}{2l+1} \sum_{m=-l}^l |a_{lm}|^2 \quad (4)$$

where $\langle C_l^{\text{sky}} \rangle = C_l$ if we were able to average over an ensemble of vantage points. There is an intrinsic cosmic variance of

$$\frac{\sigma_l}{C_l} = \sqrt{\frac{2}{2l+1}}. \quad (5)$$

In practice, instrument noise and sky masking complicate these relations. In considering potential deviations from the Λ CDM model in this paper, we examine the goodness-of-fit of the C_l model to the data, the Gaussianity of the a_{lm} derived from the map, and correlations between the a_{lm} values.

We recognize that some versions of Λ CDM (such as with multi-field inflation, for example) predict a weak deviation from Gaussianity. To date, the *WMAP* team has found no such deviations from Gaussianity. This topic is further examined by Komatsu et al. (2010, 2009). Statistical isotropy is a key prediction of the simplest inflation theories so any evidence of a violation of rotational invariance would be a significant challenge to the Λ CDM or any model based on standard inflation models.

Anomaly claims should be tested for contamination by systematic errors and foreground emission, and should be robust to statistical methodology. Statistical analyses of *WMAP* CMB data can be complicated, and simulations of skies with known properties are usually a necessary part of the analysis. Statistical analyses must account for *a posteriori* bias, which is easier said than done. With the large amount of *WMAP* data and an enormous number of possible ways to combine the data, some number of low probability outcomes are *expected*. For this reason, what constitutes a “significant” deviation from Λ CDM can be difficult to specify. While methods to reduce foreground contamination (such as sky cuts, internal linear combinations, and template-based subtractions) can be powerful, none is perfect. Since claimed anomalies often tend to be at marginal levels of significance (e.g., $2 - 3\sigma$), the residual foreground level may be a significant consideration.

The *WMAP* Science Team has searched for a number of different potential systematic effects and placed quantitative upper limits on them. The *WMAP* team has extensively examined systematic measurement errors with each of its data releases: Jarosik et al. (2003) and Hinshaw et al. (2003) for the first year data release, Hinshaw et al. (2007) and Page et al. (2007) for the three-year data release, Hinshaw et al. (2009) for the five-year data release, and Jarosik et al. (2010) for the current seven-year data release. Since those papers already convey the extensive systematic error analysis efforts of the *WMAP* team, this paper focuses on the consistency of the data with the Λ CDM model and relies on the systematic error limits placed in those papers. Some data analysis

techniques compute complicated combinations of the data where the systematic error limits must be fully propagated.

This is one of a suite of papers presenting the seven-year *WMAP* data. Jarosik et al. (2010) provide a discussion of the sky maps, systematic errors, and basic results. Larson et al. (2010) derive the power spectra and cosmological parameters from the *WMAP* data. Gold et al. (2010) evaluate the foreground emission and place limits on the foreground contamination remaining in the separated cosmic microwave background (CMB) data. Komatsu et al. (2010) present a cosmological interpretation of the *WMAP* data combined with other cosmological data. Weiland et al. (2010) analyze the *WMAP* observations of the outer planets and selected bright sources, which are useful both for an understanding of the planets and for enabling these objects to serve as more effective calibration sources for CMB and other millimeter-wave and microwave experiments.

This paper is organized as follows. In §2 we comment on the prominent large cold spot, nearby but offset from the Galactic Center region, that attracted attention when the first *WMAP* sky map was released in 2003. In §3 we comment on a cold spot in the southern sky that has attracted attention more recently. In §4 we assess the level of significance of the low value measured for the amplitude of the CMB quadrupole ($l = 2$) component. In §5 we discuss the lack of large-scale power across the sky. In §6 we assess the goodness-of-fit of the *WMAP* data to the Λ CDM model. In §7 we examine the alignment of the quadrupole and octupole. In §8 we assess claims of a hemispherical or dipole power asymmetry, and in §9 we assess claims of a quadrupolar power asymmetry. We summarize our conclusions in §10.

2. Cold Spot I, Galactic Foreground Emission, and the Four Fingers

When *WMAP* data were first released in 2003, an image of the full sky was presented in Galactic coordinates, centered on the Galactic Center (Bennett et al. 2003b). Galactic emission was minimized for this image by using an Internal Linear Combination (ILC) of *WMAP* data from independent frequency bands in such a way as to minimize signals with the frequency spectra of the Galactic foregrounds. The 7-year raw V-band map and ILC map are shown in Figure 1. A prominent cold (blue) spot is seen near the center of these maps, roughly half of which lies within the Kp2 Galaxy mask. For the portion within the mask, the ILC process removes $\sim 99\%$ of the V-band pixel-pixel variance, while making almost no difference to the variance in the portion outside the mask. In the ILC map, the variance in the two regions are nearly equal. Given its CMB-like spectrum and the fact that it is not centered on the Galactic center, this cold spot is very unlikely to be due to galactic foreground emission. Several years of modeling the separation of foreground emission from the CMB emission continues to support the conclusion that this cold spot is not dominated by Galactic foreground emission, but rather is a fluctuation of the CMB. The depth and extent of such a feature is not unusual for the best-fit Λ CDM model. Further, the probability that a randomly-located feature of this size would be near the Galactic Center is $\sim 5\%$ and the probability of such a feature overlapping the Galactic plane, at any longitude, is

much higher. This large central cold spot is a statistically reasonable CMB fluctuation within the context of the Λ CDM model.

Since people are highly effective at detecting patterns, it is not surprising that a visual inspection of the *WMAP* sky map reveals interesting features. Four elongated cold (blue) fingers stretching from about the Galactic equator to the South Galactic pole are seen in Figure 2. There do not appear to be any similar fingers or features in the northern Galactic hemisphere aside from the northern-most extensions of the mostly southern fingers. Cold Spot I can be seen to be the northern part of one of the colder fingers.

A tendency may be to ask what causes the anomaly of the four fingers, for example, or of Cold Spot I. Simulated Λ CDM maps reveal that these sorts of large scale fluctuations are expected. Rather, it is the lack of these features in the Northern sky that may be the more unusual situation.

3. Cold Spot II

A detection of non-Gaussianity and/or phase correlations in the *WMAP* a_{lm} data would be of great interest. While a detection of non-Gaussianity could be indicative of an experimental systematic effect or of residual foregrounds, it could also point to new cosmological physics. There is no single preferred test for non-Gaussianity. Rather, different tests probe different types of non-Gaussianity; therefore, it is important to subject the data to a variety of tests.

Vielva et al. (2004) used a spherical Mexican hat wavelet (SMHW) analysis on the first year *WMAP* data to claim a detection of a non-Gaussian signal on a scale of a few degrees, independent of the *WMAP* observing frequency. Also applying the SMHW kernel, Mukherjee & Wang (2004) claimed to detect non-Gaussianity at $\sim 99\%$ significance. The signal is a positive kurtosis in the wavelet coefficients attributed to a larger than expected number of 3° to 5° cold spots in the southern Galactic hemisphere. Mukherjee & Wang (2004) found the same result for the ILC map. Following up on this, Cruz et al. (2006) reported that the kurtosis in the wavelet distribution could be exclusively attributed to a single cold spot, which we call Cold Spot II, in the sky map at Galactic coordinates ($l = 209^\circ$, $b = -57^\circ$), as indicated by the red curve in Figure 2. In an analysis of the *WMAP* 3-year data, Cruz et al. (2007a) reported only 1.85% of their simulations deviated from the *WMAP* data either in the skewness or in the kurtosis estimator at three different angular scales, a 2.35σ effect.

In replicating the SMHW approach described above, Zhang & Huterer (2010) found $\sim 2.8\sigma$ evidence of kurtosis and $\sim 2.6\sigma$ evidence for a cold spot. These values are in reasonable agreement with the earlier findings for individual statistics. Zhang and Huterer then allowed for a range of possibilities in disk radius, spatial filter shape, and the choice of non-Gaussian statistic in a “superstatistic”. They found that 23% of simulated Gaussian random skies are more unusual than the *WMAP* sky. Zhang & Huterer (2010) also analyzed the sky maps with circular disk and Gaussian filters of varying widths. They found no evidence for an anomalous cold spot at any

scale when compared with random Gaussian simulations. When requiring the SMHW spatial filter shape, 15% of simulated Gaussian random skies were more unusual than the *WMAP* sky using the constrained superstatistic.

With 1.85% to 15% of random Gaussian skies more deviant than *WMAP* data in a wavelet analysis (depending on the marginalization of posterior choices) potential physical interpretations have been proposed for this $1.45\sigma - 2.35\sigma$ effect. In theory, cold spots in the CMB can be produced by the integrated Sachs-Wolfe (ISW) effect as CMB photons traverse cosmic voids along the line of sight. If Cold Spot II is due to a cosmic void, it would have profound implications because Λ CDM does not produce voids of sufficient magnitude to explain it. Mota et al. (2008) examined void formation in models where dark energy was allowed to cluster and concluded that voids of sufficient size to explain Cold Spot II were not readily produced. Rudnick et al. (2007) examined number counts and smoothed surface brightness in the NRAO VLA Sky Survey (NVSS) radio source data. They claimed a 20-45% deficit in the NVSS smoothed surface brightness in the direction of Cold Spot II. However, this claim was refuted by Smith & Huterer (2010), who found no evidence for a deficit, after accounting for systematic effects and posterior choices made in assessing statistical significance. Further, Granett et al. (2010) imaged several fields within Cold Spot II on the Canada-France-Hawaii Telescope and attained galaxy counts that rule out a 100 Mpc radius spherical void at high significance, finding no evidence for a supervoid.

Cruz et al. (2007b) suggested that the cold spot could be the signature of a topological defect in the form of a cosmic texture rather than an adiabatic fluctuation. This suggestion was further discussed by Bridges et al. (2007) and Cruz et al. (2008). CMB power spectra combined with other cosmological data can be used to place limits on a statistical population of topological defects. Urrestilla et al. (2008) placed a 95% confidence upper limit of $G\mu < 1.8 \times 10^{-6}$ based on Hubble constant, nucleosynthesis, and CMB (including three-year *WMAP*) data. Textures at this level are compatible with the cold spot and are neither favored nor disfavored by parameter fits. Following the method described in Urrestilla et al. (2008), we now place a power spectrum based 95% confidence upper limit of $G\mu < 1.5 \times 10^{-6}$ using the 7-year *WMAP* data, finer scale CMB data, and the Hubble constant. Since the new 95% upper limit value from the power spectrum corresponds with the Cruz et al. estimate for cold spot feature (after correction for bias using simulations), the texture model is now disfavored at about the 95% confidence level.

In conclusion, there are two possible points of view. One is that the cold spot is a texture. This can not be ruled out with current data. The other view is that the $< 2.35\sigma$ (after posterior correction) statistical evidence for a cold spot feature is not compelling, and that the texture explanation is disfavored. Had the anomaly been significant at the part per million level instead of a part per thousand, the posterior marginalization issues would be moot: the feature would be considered real. If real, a void explanation would have been in strong conflict with the Λ CDM model, but a population of topological defects that contribute at a low level to the CMB power spectrum would not so much falsify the model as provide a small modification to it. Given the evidence to date, the WMAP Team is of the opinion that there is insufficient statistical

support to conclude that the cold spot is a CMB anomaly relative to Λ CDM.

4. The Low Quadrupole Amplitude

The CMB quadrupole is the largest observable structure in our universe. The magnitude of the quadrupole has been known to be lower than models predict since it was first measured by *COBE* (Bennett et al. 1992; Hinshaw et al. 1996). It is also the large-scale mode that is most prone to foreground contamination, owing to the disk-like structure of the Milky Way, thus estimates of the quadrupole require especially careful separation of foreground and CMB emission. The ILC method of foreground suppression is especially appropriate for large-scale foreground removal, since the ILC’s complex small-scale noise properties are unimportant in this context. With a sky cut applied, residual foreground contamination of the quadrupole in the ILC map is determined to be insignificant (Gold et al. 2010; Jarosik et al. 2010). In this section we reassess the statistical significance of the low quadrupole power in the ILC map.

A statistical analysis of the quadrupole must account for the highly non-Gaussian posterior distribution of the low- l ($l \lesssim 32$) multipoles. In this paper we use Gibbs sampling of the low- l power spectrum to evaluate the quadrupole (O’Dwyer et al. 2004; Dunkley et al. 2009). The *WMAP* ILC data are smoothed to 5° resolution (Gaussian, FWHM), degraded to HEALPix resolution level 5 ($N_{\text{side}} = 32$), and masked with the KQ85y7 mask. The Gibbs sampling produces a likelihood that has been marginalized over all other multipoles. A Blackwell-Rao estimate is used to produce the *WMAP* quadrupole likelihood shown in Figure 3. The peak of the likelihood is at $200 \mu\text{K}^2$, the median is at $430 \mu\text{K}^2$, and the mean is at $1000 \mu\text{K}^2$. The 68% confidence range extends from 80 to $700 \mu\text{K}^2$, and the 95% confidence range extends from 40 to $3200 \mu\text{K}^2$.

Figure 4 shows the cumulative quadrupole distribution derived from 300,000 Gibbs samples. The mean quadrupole predicted by the best-fit Λ CDM model lies at a cumulative probability of 0.824, well within the 95% confidence region allowed by the data. We conclude that the *WMAP* quadrupole measurement is not anomalously low. Further, while alternative models that predict a lower quadrupole will better match this specific part of the data, it is impossible to significantly favor such models on the basis of quadrupole power alone.

5. The Lack of Large-Scale Power

The angular correlation function complements the power spectrum by measuring structure in real space rather than Fourier space. It measures the covariance of pixel temperatures separated by a fixed angle,

$$C(\theta_{ij}) = \langle T_i T_j \rangle \tag{6}$$

where i and j are two pixels on the sky separated by an angle θ , and the brackets indicate an ensemble average over independent sky samples. Expanding the temperature in spherical harmonics, and using the addition theorem for spherical harmonics, it is straightforward to show that $C(\theta_{ij})$ is related to the angular power spectrum by

$$\langle C(\theta_{ij}) \rangle = \frac{1}{4\pi} \sum_l (2l+1) \langle C_l \rangle W_l P_l(\cos \theta) \quad (7)$$

where $\langle C_l \rangle$ is the ensemble-average angular power spectrum, W_l is the experimental window function, and P_l are the Legendre polynomials. If the CMB is statistically isotropic, $C(\theta_{ij}) \equiv C(\theta)$ depends only on the separation of pixels i and j , but not on their individual directions. Since we are unable to observe an ensemble of skies, we must devise estimates of $C(\theta)$ using the temperature measured in our sky. One approach is to assume that the CMB is ergodic (statistically isotropic) in which case $C(\theta)$ can be estimated by averaging all temperature pairs in the sky separated by an angle θ ,

$$C(\theta) = \langle T_i T_j \rangle |_{\angle ij = \theta} \quad (8)$$

where the brackets indicate an average over directions i and j such that $\angle ij = \theta$ (to within a bin). Another approach is to estimate the angular power spectrum C_l and to compute $C(\theta)$ using Eq. (7).

The angular correlation function over the full sky ILC map from Equation (7) is shown in Figure 5. As can be seen, $C(\theta)$ lies within the 95% confidence range of the best fit Λ CDM model for all θ , as determined by Monte Carlo simulations. This supports the conclusion that there is no statistically significant lack of large-scale power on the full sky.

Spergel et al. (2003) applied the pixel-pair estimator to the first-year *WMAP* data and found an almost complete lack of correlated structure at angles $> 60^\circ$ for the sky, but that calculation was with a Galactic foreground cut. A foreground cut was made because of the concern that additional power from within the Galactic cut may arise from foregrounds. For regions outside the cut, it was appreciated that systematic errors and residual Galactic foregrounds are far more likely to *add* correlated power to the sky maps than to remove it. They quantified the the lack of large-angular-scale power in terms of the statistic

$$S_{1/2} \equiv \int_{-1}^{1/2} C^2(\theta) d \cos \theta \quad (9)$$

and found that fewer than 0.15% of simulations had lower values of $S_{1/2}$. A low $S_{1/2}$ value persists in later *WMAP* sky maps.

Copi et al. (2007) and Copi et al. (2009) claimed that there is evidence that the *WMAP* temperature fluctuations violate statistical isotropy. They directly computed the angular correlation function from pixel pairs, as in Eq. (8), omitting from the sum pixel pairs where at least one pixel was within the foreground mask. The KQ85 foreground mask (at that time) removed 18% of the pixels from the full sky (now 22% for KQ85y7), while KQ75 removed 29%. Copi et al. found

p -values of $\approx 0.03\%$ for their computation of $S_{1/2}$, concluding that the data are quite improbable given the model. The exact p -value depended on the specific choice of CMB map and foreground mask. Cayón (2010) finds no frequency dependence to the effect.

Efstathiou et al. (2009) find that the value of $S_{1/2}$ is sensitive to the method of computation. For example, Efstathiou et al. (2009) computed the angular correlation function using the estimator

$$C(\theta) = \frac{1}{4\pi} \sum_l (2l+1) M_{ll'}^{-1} \tilde{C}_{l'} P_l(\cos\theta) \quad (10)$$

where

$$M_{ll'} = \frac{1}{2l+1} \sum_{mm'} |K_{lm'l'm'}|^2 \quad (11)$$

and $K_{lm'l'm'}$ is the coupling between modes (lm) and $(l'm')$ induced by the sky cut, and $\tilde{C}_{l'}$ is the pseudo-power spectrum obtained by transforming the sky map into spherical harmonics on the cut sky. This estimator produced a significantly larger value for $S_{1/2}$ than the estimator in Equation (6).

Efstathiou et al. (2009) also reconstructed the low- l multipoles across the foreground sky cut region in a manner that was numerically stable, without an assumption of statistical isotropy. Their method relied on the fact that the low multipole *WMAP* data are signal-dominated and that the cut size is modest. They showed that the small reconstruction errors introduce no bias and they did not depend on assumptions of statistical isotropy or Gaussianity. The reconstruction error only introduced a small “noise” to the angular correlation function without changing its shape.

The original use of a sky cut in calculating $S_{1/2}$ was motivated by concern for residual foregrounds in the ILC map. We now recognize that this precaution was unnecessary as the ILC foreground residuals are relatively small. Values of $S_{1/2}$ are smaller on the cut sky than on the full sky, but since the full sky contains the superior sample of the universe and the cut sky estimates suffer from a loss of information, cut sky estimates must be considered sub-optimal. It now appears that the Spergel et al. (2003) and Copi et al. (2007, 2009) low p -values result from both the *a posteriori* definition of $S_{1/2}$ and a chance alignment of the Galactic plane with the CMB signal. The alignment involves Cold Spot I and the northern tips of the other fingers, and can also be seen in the maps that will be discussed in §7.

Efstathiou et al. (2009) corrected the full-sky *WMAP* ILC map for the estimated Integrated Sachs-Wolfe (ISW) signal from redshift $z < 0.3$ as estimated by Francis & Peacock (2009). The result was a substantial increase in the $S_{1/2}$. Yet there is no large-scale cosmological significance to the orientation of the sky cut or the orientation of the local distribution of matter with respect to us, thus the result from Spergel et al. and Copi et al. must be a coincidence.

More generally, Hajian et al. (2005) applied their bipolar power spectrum technique and found no evidence for a violation of statistical isotropy at 95% CL for angular scales corresponding to multipole moments $l < 60$.

The low value of the $S_{1/2}$ integral over the large-angle correlation function on the cut-sky results from a posterior choice of the statistic. Further, it is a sub-optimal statistic in that it is not computed over the full sky. There is evidence for a chance alignment of the Galactic plane cut with the CMB signal, and there is also evidence of a chance alignment of the primary CMB fluctuations with secondary ISW features from the local density distribution. The full-sky angular correlation function lies within the 95% confidence range. For all of these reasons, we conclude that the large-angle CMB correlation function is consistent with Λ CDM.

6. The Goodness of Fit of the Standard Model

The power spectrum of the *WMAP* data alone places strong constraints on possible cosmological models (Dunkley et al. 2009; Larson et al. 2010). Plots of the χ^2 per degree of freedom of the temperature-temperature (TT) power spectrum data relative to the best-fit Λ CDM model are shown in Figure 6. In Figure 7 the cumulative probability of the *WMAP* data given the Λ CDM model is evaluated based on simulations. All of the simulated skies were calculated for the same input Λ CDM model, but each result was fit separately. The *WMAP* sky is statistically compatible with the model within 82.6% confidence, with an uncertainty of $\sim 5\%$.

The χ^2 can be elevated because of excess scatter within each multipole relative to the experimental noise variance. It could also be elevated because of an accumulation of systematic deviations of the model from the data across different multipoles, such as would happen if a parameter value were incorrect. Therefore, despite an acceptable overall χ^2 , we examine other aspects of the power spectrum data relative to the model that may have been masked by the inclusion of all of the data into a single χ^2 value. We examine both of these possibilities below.

To explore the l -to- l' correlation properties of C_l , we compute:

$$\begin{aligned} S(\Delta l) &\equiv \sum_l \frac{(C_l^{\text{data}} - C_l^{\text{bestfit}})(C_{l+\Delta l}^{\text{data}} - C_{l+\Delta l}^{\text{bestfit}})}{\sqrt{\frac{2(C_l^{\text{bestfit}} + N_l)^2}{(2l+1)f_{\text{sky},l}^2} \frac{2(C_{l+\Delta l}^{\text{bestfit}} + N_{l+\Delta l})^2}{(2l+2\Delta l+1)f_{\text{sky},l+\Delta l}^2}}} \\ &= \sum_l f_{\text{sky},l} f_{\text{sky},l+\Delta l} \sqrt{\left(l + \frac{1}{2}\right) \left(l + \Delta l + \frac{1}{2}\right)} \\ &\quad \times \frac{(C_l^{\text{data}} - C_l^{\text{bestfit}})(C_{l+\Delta l}^{\text{data}} - C_{l+\Delta l}^{\text{bestfit}})}{(C_l^{\text{bestfit}} + N_l)(C_{l+\Delta l}^{\text{bestfit}} + N_{l+\Delta l})}. \end{aligned}$$

When $\Delta l = 0$, this quantity is exactly χ^2 :

$$S(\Delta l = 0) = \chi^2.$$

For $l = 300 - 349$ and $f_{\text{sky}} \sim 0.8$,

$$S(\Delta l) \sim 200 \sum_{l=300}^{349} \frac{(C_l^{\text{data}} - C_l^{\text{bestfit}})(C_{l+\Delta l}^{\text{data}} - C_{l+\Delta l}^{\text{bestfit}})}{(C_l^{\text{bestfit}} + N_l)(C_{l+\Delta l}^{\text{bestfit}} + N_{l+\Delta l})}.$$

Since the data suggest $S(\Delta l = 1) \sim -20$ for $l = 300 - 349$, we find

$$\frac{(C_l^{\text{data}} - C_l^{\text{bestfit}})(C_{l+1}^{\text{data}} - C_{l+1}^{\text{bestfit}})}{(C_l^{\text{bestfit}} + N_l)(C_{l+1}^{\text{bestfit}} + N_{l+1})} \sim -0.002.$$

For this multipole range $(C_l^{\text{bestfit}} + N_l) \sim 3500 \mu\text{K}^2$; thus,

$$(C_l^{\text{data}} - C_l^{\text{bestfit}})(C_{l+1}^{\text{data}} - C_{l+1}^{\text{bestfit}}) \sim -(160 \mu\text{K}^2)^2.$$

Note that the power spectrum of the Finkbeiner et al. (1999) dust map in this multipole range is $< 10 \mu\text{K}^2$ in W band, i.e., more than an order of magnitude smaller.

Figure 8 shows the results of l -to- l' correlation calculations of C_l for different values of $l - l'$, calculated both for simulations and for the *WMAP* data. For the most part the data and simulations are in good agreement. The most discrepant correlations in C_l are for $\Delta l = 1$ near $l \sim 320$ and $\Delta l = 2$ near $l \sim 280$.

Motivated by the outlier $\Delta l = 1$ correlation at $l \sim 320$ seen in Figure 8, we further explore a possible even l versus odd l effect in this portion of the power spectrum. (Note that this is an *a posteriori* selection.) We define an even excess statistic, \mathcal{E}_ℓ , which compares the mean power at even values of ℓ with the mean power at odd values of ℓ , within a given ℓ -range. It is essentially a measure of anticorrelation between adjacent elements of the power spectrum, with a sign indicating the phase of the anticorrelation:

$$\mathcal{E}_\ell = \frac{\langle C_\ell^{\text{obs}} - C_\ell^{\text{th}} \rangle_{\text{even}} - \langle C_\ell^{\text{obs}} - C_\ell^{\text{th}} \rangle_{\text{odd}}}{\langle C_\ell^{\text{th}} \rangle},$$

where $C_\ell = \ell(\ell + 1)C_\ell/2\pi$, the superscript “obs” refers to the observed power spectrum, and the superscript “th” refers to a fiducial theoretical power spectrum used for normalization. From this definition, it follows that $\mathcal{E}_\ell > 0$ is an *even excess* and $\mathcal{E}_\ell < 0$ is an *odd excess*. The range of ℓ is small enough that the variation in C_ℓ^{th} is also small and convenient for normalization. We choose (*a posteriori*) to bin \mathcal{E}_ℓ by $\Delta\ell = 50$.

An apparent positive \mathcal{E}_ℓ in the *WMAP* power spectrum in the range $\ell \sim 200 - 400$ is investigated quantitatively using Monte Carlo simulations. Our analysis is limited to $33 \leq \ell < 600$, which is the part of the power spectrum that is flat-weighted on the sky and where the Monte Carlo Apodised Spherical Transform EstimatoR (MASTER) (Hivon et al. 2002) pseudo-spectrum is used. Because the binning is by $\Delta\ell = 50$, the actual ℓ range for this analysis is $50 - 599$. The Monte Carlo realizations are CMB sky map simulations incorporating appropriate *WMAP* instrumental noise and beam smoothing. Each power spectrum, whether from observed data or from the Monte Carlo generator, is coadded from 861 year-by-year cross spectra in the V and W bands, with weighting that accounts for the noise and the beam transfer functions.

Figure 9 shows that an even excess of significance $\sim 2.7\sigma$ is found for $\ell = 300 - 349$. If we combine the two adjacent bins between $\ell = 250$ and $\ell = 349$, the significance of \mathcal{E}_ℓ in the combined

bins is $\sim 2.9\sigma$, with a probability to exceed (PTE) of 0.26% integrated directly from the Monte Carlo set (Figure 10). However, it is important to account for the fact that this significance level is inflated by the posterior bias of having chosen the ℓ range to give a particularly high value.

We attempt to minimize the posterior bias by removing bin selection from the Monte Carlo test. Instead of focusing on one bin, the revised test is based on the distribution of the *maximum* value of the significance $\mathcal{E}_\ell/\sigma(\mathcal{E}_\ell)$ over *all* bins in each Monte Carlo realization. The 11436 Monte Carlo realizations are split into two groups: 4000 are used to compute the normalization $\sigma(\mathcal{E}_\ell)$ for each bin, and 7436 are used to compute the distribution of the maximum value of $\mathcal{E}_\ell/\sigma(\mathcal{E}_\ell)$, giving the histogram that is compared to the single observed value.

The results of the de-biased test are shown in Figure 11. In addition to the avoidance of bin selection, this test also incorporates negative excursions of \mathcal{E}_ℓ , which are excesses of power at odd ℓ . The test shows that the visually striking even excess in the $\ell = 300 - 350$ bin is actually of low significance, with a PTE of 5.1% (top of Figure 11). However, the test also shows that large excursions in odd- ℓ power are less frequent in the observed power spectrum than in the Monte Carlos, such that 99.3% of the Monte Carlos have a bin with greater odd- ℓ significance than the observed data (bottom of Figure 11). Thus there appears to be a modestly significant suppression of odd- ℓ power. This effect is only slightly relieved by accounting for the posterior selection of the enhanced even excess, as seen in the bottom panel of Figure 11.

We find no evidence for a radiometer-dependence of the effect. We are suspicious that the effect could arise from an interaction of the foreground mask with large-scale power in the map, but our simulation results were not nearly of sufficient amplitude to explain this effect.

Steps and other sharp features in the power spectrum $P(k)$ tend to be smeared out in translation to C_l space. For example, for the non-standard “meandering” cosmological inflation model of Tye & Xu (2010) the scalar mode responsible for inflation meanders in a multi-dimensional potential. This leads to a primordial CMB power spectrum with complicated small-amplitude variations with wavenumber k . Conversion to C_l has the effect of a significant amount of smoothing (see, for example, Figure 1.4 of Wright (2004)). It is not likely that any cosmological scenario can cause the observed odd excess deficit. Likewise, we are aware of no experimental effects that could cause an odd excess deficit. We therefore conclude that this $< 3\sigma$ effect is a statistical fluke.

7. Aligned Quadrupole and Octupole

The alignment of the quadrupole and octupole was first pointed out by Tegmark et al. (2003) and later elaborated on by Schwarz et al. (2004), Land & Magueijo (2005a), and Land & Magueijo (2005b). The fact of the alignment is not in doubt, but the significance and implications of the alignment are discussed here.

Do foregrounds align the quadrupole and dipole? Chiang, Naselsky & Coles (2007) conclude

that the lowest spherical harmonic modes in the ILC map are significantly contaminated by foregrounds. Park, Park, & Gott (2007) find that the residual foreground emission in a map resulting from their own independent foreground analysis is not statistically important to the large-scale modes of CMB anisotropy. The large scale modes of their map show anti-correlation with the Galactic foreground emission in the southern hemisphere, but they are agnostic on whether this is due to residual Galactic emission or by simply a matter of chance. Park, Park, & Gott (2007) also assess the WMAP Team’s ILC map and conclude that residual foreground emission in the ILC map does not affect the estimated large scale values significantly. Tegmark et al. (2003) also performed their own foreground analysis and conclude that their CMB map is clean enough that the lowest multipoles can be measured without any galaxy cut at all. They also point out that much of the CMB power falls within the Galaxy cut region, “seemingly coincidentally.” In other words, they conclude that the residual foregrounds are subdominant to the intrinsic CMB signal even without any Galaxy cut so long as a reduced foreground map is used. de Oliveira-Costa & Tegmark (2006) state that they believe it is more likely that the true alignment is degraded by foregrounds rather than created by foregrounds.

We determine the direction of the quadrupole by rotating the coordinate system until $|a_{\ell,\ell}|^2 + |a_{\ell,-\ell}|^2$ is maximized, where $\ell = 2$ and where $a_{\ell,m}$ are the spherical harmonic coefficients. This maximizes the power around the equator of the coordinate system. The optimization is done by numerically checking the value of this quantity where the z vector of the coordinate system is rotated to the center of each half-degree pixel in a res 7 ($N_{\text{side}} = 128$) HEALPix pixelization. The pixel where this value is maximized is taken to be the “direction” of the quadrupole. The pixel on the opposite side of the sphere has the same value; we arbitrarily pick one. Using a similar method, a “direction” is found for the octupole, with $\ell = 3$.

The probability that $l = 2$ and $l = 3$ multipoles would be aligned is shown in Figure 12. The $< 1^\circ$ alignment in our sky appears to be quite improbable based upon random simulations of the best-fit Λ CDM model.

A resolution level 7 map has 196608 pixels of about 0.5° diameter. The best fit alignment axis specifies two pixels directly opposite each other on the sphere. The probability of two axes randomly aligning in the same pair of pixels is then $2/196608 = 0.001\%$. The probability of getting an alignment within 0.25° of a given axis is 0.00095% , which is close to 0.001% above.

In an attempt to gain insight into the alignment we start with the ILC full sky temperature map. We then produce a map of ΔT^2 , which is a map of anisotropy power. This map is constructed by smoothing the 7-year ILC map by 10° , removing the mean value, and squaring. Since the ILC map is already smoothed to 1° , the total smoothing corresponds to a Gaussian with $\text{FWHM} = \sqrt{10^2 + 1^2} = 10.05^\circ$. We then create various masks to probe whether the dipole-quadrupole alignment can be attributed to one or perhaps two localized features in the map. The edges of the some of the masks are found from contours of the ΔT^2 map. The contours were selected by eye, from a grayscale Mollweide projection in an image editing program, and then converted to

HEALPix fits files. Others trial masks were chosen more randomly, again to probe the sensitivity of the alignment to different regions of the map.

For each mask, we take the 7 year ILC map that has been smoothed by 10° , zero the region inside the mask, and take a spherical harmonic transform. From the $a_{\ell m}$ coefficients, we determine the angle between the quadrupole and octupole.

Figure 13 shows the color smoothed, squared temperature map and the effect of various masks (shown in gray) on the quadrupole-octupole alignment. Masking Cold Spot I eliminates any significant alignment. However, keeping that region but masking other regions also significantly reduces the quadrupole-octupole alignment. The posterior selection of the particular masked regions is irrelevant as the point is only to demonstrate that no single region or pair of regions solely generates the $< 1^\circ$ alignment. Rather the high degree of quadrupole-octupole alignment results from the statistical distribution of anisotropy power over the whole sky. This rules out single-void models, a topological defect at some sky position, or any other such explanation. The alignment behaves as one would expect if it originates from chance random anisotropy amplitudes and phases. The alignment of the $l = 2$ and $l = 3$ multipoles are intimately connected with the large-scale cool fingers and intervening warm regions, discussed earlier, as can be seen in Figure 14. Although the alignment is indeed remarkable, current evidence is more compatible with a statistical combination of full sky data than with the dominance of one or two discrete regions.

Francis & Peacock (2009) estimated the local ($z < 0.3$) density field from the 2MASS and SuperCOSMOS galaxy catalogs and used that field to estimate the integrated Sachs-Wolfe (ISW) effect within this volume. Large-scale features were extrapolated across the Galactic plane. The effects of radial smearing from the photometric galaxy sample were reduced by taking only three thick redshift shells with $\Delta z = 0.1$. A linear bias was used to relate galaxies to density, independent of both scale and redshift within each of the three shells. They estimated that the $z < 0.3$ data limit contains $\sim 40\%$ of the total ISW signal. Francis & Peacock (2009) removed their estimate of the ISW effect from the *WMAP* map. One result was to raise the amplitude of the quadrupole while the octupole amplitude was relatively unchanged. More importantly, they reported that there remains no significant quadrupole-octupole alignment after the ISW removal. With the Francis & Peacock (2009) result, the quadrupole-octupole alignment shifts from an early universe property to a statistical fluke that the secondary anisotropy effect from the local density distribution happens to superpose on the primordial anisotropy in such a way as to align the quadrupole and octupole.

8. Hemispherical and Dipole Power Asymmetry

Claims of a dipole or hemispherical power asymmetry in *WMAP* maps have appeared in the literature since the release of the first-year *WMAP* data, with estimates of statistical significance ranging up to 3.8σ . We distinguish between a “hemispherical” power asymmetry, in which the power spectrum is assumed to change discontinuously across a great circle on the sky, and a

“dipole” power asymmetry in which the CMB is assumed modulated by a smooth cosine function across the sky, i.e. the CMB is assumed to be of the form

$$T(\mathbf{n})_{\text{modulated}} = (1 + \mathbf{w} \cdot \mathbf{n}) T(\mathbf{n})_{\text{unmodulated}}. \quad (12)$$

Previous analyses of *WMAP* data in the literature have fit for either hemispherical or dipolar power asymmetry, and the results are qualitatively similar: asymmetry is found with similar direction and amplitude in the two cases. However, analyses that use optimal estimators (as we do in our analysis below) have all studied the dipolar modulation (e.g. Hoftuft et al. (2009) and Hanson & Lewis (2009)). Furthermore, theoretical attempts to obtain cosmological power asymmetry by altering the statistics of the primordial fluctuations (Gordon 2007; Donogue, Dutta & Ross 2009; Erickcek, Kamionkowski & Carroll 2008; Erickcek et al. 2009) have all found a dipolar modulation rather than a hemispherical modulation. Therefore, we will concentrate on the dipolar modulation, defined by Eq. (12), for the sake of better comparison with both early universe models, and with similar analyses in the literature. Unambiguous evidence for power asymmetry would have profound implications for cosmology.

We revisit this analysis and conclude that this claimed anomaly is not statistically significant, after *a posteriori* choices are carefully removed from the analysis. When looking for a power asymmetry, the most significant issue is removing an arbitrary choice of scale, either specified explicitly by a maximum multipole l , or implicitly by a sequence of operations such as smoothing and adding extra noise that define a weighting in l .

The first claimed hemispherical power asymmetry appeared in Eriksen et al. (2004), based on the first-year *WMAP* data, where a statistic for power asymmetry was constructed, and its value on high-resolution *WMAP* data was compared to an ensemble of Monte Carlo simulations in a direct frequentist approach. They quoted a statistical significance of 95 – 99%, depending on the range of l selected. The details of this analysis contained many arbitrary choices. Hansen et al. (2004, 2009) also used a similar methodology. In Hansen et al. (2009), the significance of a $2 \leq l \leq 600$ hemispherical power asymmetry was quoted as 99.6%.

A second class of papers used a low-resolution pixel likelihood formalism to study power asymmetry. Eriksen et al. (2007) used this approach to search for a dipole power asymmetry in low-resolution 3-year *WMAP* data and a statistical significance of $\sim 99\%$ was claimed. Hoftuft et al. (2009) repeated the likelihood analysis at somewhat higher resolution and quoted a statistical significance of $3.5\sigma - 3.8\sigma$ for different choices of resolution. Although the likelihood estimator contained fewer arbitrary choices than the preceding class of papers, the low-resolution framework still contained an arbitrary choice of angular scale, which may be tuned (intentionally or unintentionally) to spuriously increase statistical significance. Hoftuft et al. (2009) introduced an explicit cutoff multipole l_{mod} and the CMB signal was assumed to be unmodulated for $l > l_{\text{mod}}$ and modulated for $l \leq l_{\text{mod}}$. Both Eriksen et al. (2007) and Hoftuft et al. (2009) downgraded the data in resolution, smoothed with a Gaussian window, and added extra white noise. These processing steps implicitly defined a weighting in l where the power asymmetry is estimated, and introduced arbitrary choices

into the analysis.

A third approach to the analysis, based on optimal quadratic estimators, recently appeared in Hanson & Lewis (2009). In this approach, the *WMAP* data were kept at full resolution and a minimum-variance quadratic estimator was constructed for each of the three vector components of the dipole modulation \mathbf{w}_i . Hanson and Lewis found $\approx 97\%$ evidence for a dipole power asymmetry at $2 \leq l \leq 40$, and $\approx 99.6\%$ evidence for dipole power asymmetry at $2 \leq l \leq 60$. However, the result was strongly dependent on changing the l range, and quickly went away for higher l . A significant shift was seen between the KQ75 and KQ85 masks, and between raw and clean maps, suggesting that foreground contamination was not negligible.

Comparing these methods, we find that the Hanson & Lewis (2009) optimal quadratic estimator has significant advantages over other analysis methods that have appeared in the literature:

1. There are no arbitrary choices (such as smoothing scale) in the optimal quadratic estimator. One can either look for power asymmetry in a range of multipoles $2 \leq l \leq l_{\text{mod}}$, or over the entire range of angular scales measured by *WMAP*, and the estimator is uniquely determined by the minimum-variance requirement in each case. None of the previously considered methods had this property.
2. There is no need to degrade the *WMAP* data, or include processing steps such as adding extra noise, since the optimal quadratic estimator can be efficiently computed at full *WMAP* resolution using the multigrid C^{-1} code from Smith et al. (2007).
3. Statistical significance can be assessed straightforwardly by comparing the estimator with an ensemble of Monte Carlo simulations. In particular, maximum likelihood analyses in the literature have assessed significance using Bayesian evidence, but schemes for converting the evidence integral into a frequentist probability are not a sufficient substitute for true Monte Carlo simulations, which directly give the probability for a simulation to be as anomalous as the data.

For these reasons, we have studied dipole power asymmetry using the optimal quadratic estimator. We introduce a cutoff multipole l_{mod} , and assume that the CMB is isotropic for $l > l_{\text{mod}}$ and dipole-modulated for $2 \leq l \leq l_{\text{mod}}$. There is an optimal quadratic estimator \hat{w}_i for each of the three components of the (vector) modulation \mathbf{w} (Eq. (12)), and an estimator $\hat{\kappa}_1$ for the (scalar) amplitude of the modulation. Implementation details of the estimators are presented in Appendix A, where we also comment on the relation with maximum likelihood.

Figure 15 shows the probability that the value of the dipole modulation statistic $\hat{\kappa}_1$ is larger than for the *WMAP* data, when evaluated by Monte Carlo simulation. There are choices of l_{mod} where the power asymmetry appears to have high significance. For example, when we chose the KQ85y7 mask and $l_{\text{mod}} = 67$, the probability for a simulation to have a value of $\hat{\kappa}_1$ that is larger than the data is 0.7%. This could be interpreted as 2.5σ evidence for a power asymmetry, but

such an interpretation would be inflating the statistical significance since the choice of l_{mod} is an *a posteriori* choice. Consider an analogous example for the 5-year C_l power spectrum. The power in $C_{l=512}$ is high by 3.7σ , but this is not really a 3.7σ anomaly. Rather, it reflects the fact that there are a large number of l values that could have been chosen.

Now we seek to assess the global statistical significance of the power asymmetry without making any *a posteriori* choices. Consider the probability for a Monte Carlo simulation to have a larger value of $\hat{\kappa}_1$ than the *WMAP* data, as a function of l_{mod} . This can be interpreted as the statistical significance for power asymmetry in the range $2 \leq l \leq l_{\text{mod}}$, for a fixed choice of l_{mod} . Let η be the minimum value of the probability, which we find to be $\eta = 0.012$ with the KQ75y7 mask (corresponding to $l_{\text{mod}} = 66$), or $\eta = 0.007$ with the KQ85y7 mask (corresponding to $l_{\text{mod}} = 67$). We now assess whether this value of η is anomalously low. To determine this, we compute η for an ensemble of Monte Carlo simulations where l_{mod} is chosen to maximize the value of η *independently for each simulation*.

We perform the maximization over the range $10 \leq l_{\text{mod}} \leq 132$. (The results depend only weakly on this choice of range; we have taken the upper limit of the range to be twice as large as the most anomalous l_{mod} in the *WMAP* data.) We find that the probability that a simulation has a value of η that is smaller than for *WMAP* is 13% for the KQ75y7 mask, and 10% for the KQ85y7 mask.

Motivated by the power asymmetry, Erickcek et al. (2009) presented a variation of the curvaton inflationary scenario in which the curvaton has a large-amplitude super-horizon spatial gradient that modulates the amplitude of CMB anisotropy, thereby generating a hemispherical power asymmetry that could match the CMB data. Hirata (2009) used high-redshift quasars to place a limit on the gradient in the amplitude of perturbations that would be caused in this scenario. Their limit ruled out the simplest version of this curvaton spatial gradient scenario. Our new CMB results, presented here, largely remove the initial motivation for this theory.

We conclude that there is no significant evidence for an anomalous dipole power asymmetry in the *WMAP* data.

9. Quadrupolar Dependence of the Two-Point Function

There is another effect, related to the dipole power asymmetry from §8, in which the two-point function of the CMB contains a component that varies as a quadrupole on the sky. Motivated by an anisotropic model of the early universe proposed by Ackerman et al. (2007) that predicts such a signal, Groeneboom & Eriksen (2009) used a Gibbs sampler to claim “tentative evidence for a preferred direction” of $(l, b) = (110^\circ, 10^\circ)$ in the five-year *WMAP* map. A theoretical model that predicts large-scale quadrupolar anisotropy was also proposed by Gordon et al. (2005). A similar effect could be caused by *WMAP*’s asymmetric beams, which were not accurately represented in this work, and an algebraic factor was missing in the analysis. In Hanson & Lewis (2009), the

missing algebraic factor was corrected and the effect was verified with high statistical significance, using an optimal cut-sky quadratic estimator. Optimal estimators had previously been constructed in the all-sky case by Pullen & Kamionkowski (2007) and Dvorkin et al. (2008).

Recently, Groeneboom et al. (2009) returned another fit, this time including polarization, the factor correction, examinations of beam asymmetries, noise mis-estimation, and zodiacal dust emission. The new claim was 9σ evidence of the preferred direction $(l, b) = (96^\circ, 30^\circ)$, which was quite far from the original alignment direction claimed. The new preferred direction was towards the ecliptic poles, strongly suggesting that this is not a cosmological effect. The claimed amplitude was frequency dependent, also inconsistent with a cosmological effect. Zodiacal dust emission was ruled out as the source of the alignment. Hanson & Lewis (2009) found that the beam asymmetry was a large enough effect to explain the signal, although Groeneboom et al. reported the opposite conclusion. The claimed statistical significance of the quadrupolar power asymmetry is so high that it seems impossible for it to be a statistical fluke or built up by posterior choices, even given the number of possible anomalies that could have been searched for.

We have implemented the optimal quadratic estimator following the approach of Hanson & Lewis (2009) and confirmed that the effect exists with high statistical significance. Rather than present an analysis that is tied to a particular model (either cosmological or instrumental), we have found it convenient to parameterize the most general quadrupolar power asymmetry using the language of the bipolar spectrum from Hajian & Souradeep (2003). This is reviewed in Appendix A. The summary is that the most general quadrupolar anomaly can be parameterized by two quantities A_{ll}^{2M} and $A_{l-2,l}^{2M}$ which are l -dependent and have 5 components corresponding to the degrees of freedom of a quadrupole. If statistical isotropy holds, then $A_{ll}^{2M} = A_{l-2,l}^{2M} = 0$. The special case where $A_{ll}^{2M} \approx A_{l-2,l}^{2M} \neq 0$ corresponds to an anisotropic model in which the local power spectrum varies across the sky [i.e. the quadrupolar analogue of the dipole modulation in Eq. (12)]. The special case where $A_{ll}^{2M} \approx -2A_{l-2,l}^{2M} \neq 0$ corresponds to an anisotropic model in which the local power spectrum is isotropic, but hot and cold spots have preferred ellipticity where the local magnitude and orientation varies across the sky. Thus there are two independent “flavors” of quadrupolar anomaly; the bipolar power spectrum distinguishes the two and also keeps track of the l dependence. A proposed model for the quadrupolar effect in the *WMAP* data can be tested by computing the bipolar power spectrum of the model, and comparing with estimates of the bipolar spectrum from data.

Figure 16 shows the components of the bipolar power spectrum of the *WMAP* data that point along the ecliptic axis (i.e. A_{ll}^{20} and $A_{l-2,l}^{20}$ in ecliptic coordinates). A nonzero bipolar power spectrum is seen with high statistical significance, even in a single bin with $\Delta l = 50$, confirming the existence of a quadrupolar effect.

We implemented a number of diagnostic tests to characterize the quadrupolar effect; our findings can be summarized as follows:

1. Only the components of the bipolar power spectrum that point in the ecliptic direction (i.e.

components $A_{l_1 l_2}^{2M}$ with $M = 0$ in ecliptic coordinates) contain a statistically significant signal. The components with $M = 1$ or $M = 2$ are consistent with zero within their statistical errors, even if we sum over all values of l to maximize signal-to-noise.

2. The effect is larger in W-band than V-band, which is inconsistent with a cosmological origin.
3. The angular dependence of the effect shows a bump at the scale of the first acoustic peak ($l \approx 220$), disfavoring an explanation from foregrounds or noise, which would not be expected to show acoustic peaks.
4. If we split the optimal quadratic estimator into contributions from cross correlations between differential assemblies (DA's) in *WMAP*, and auto correlations in which each DA is correlated with itself, then we find that the amplitude of the effect is consistent in the two cases, disfavoring instrumental explanations that are not highly correlated between channels (such as striping due to $1/f$ noise).
5. The bipolar power spectrum of *WMAP* satisfies $A_{\ell\ell}^{20} \approx -2A_{\ell-2,\ell}^{20}$, corresponding to a model in which the small-scale power spectrum is isotropic, but the shapes of hot and cold spots are not. (In fact, for this reason, we have used the term “quadrupolar effect” in this section rather than “quadrupolar power asymmetry”, which would suggest that the power spectrum is modulated. We favor the label “effect” over “anomaly” because it is only an anomaly in the absence of a plausible cause.)

Given the strong ecliptic alignment and that the ecliptic plane was the symmetry axis of the *WMAP* observations, and the non-blackbody frequency dependence, we conclude that this is not a Λ CDM anomaly. It seems very likely that the observed quadrupolar effect is the result of incomplete handling of beam asymmetries. Beam asymmetry generates an instrumental bipolar power spectrum that is consistent with all 5 items above, and it is difficult to think of any other instrumental contribution that satisfies these properties. However, we have not yet simulated the effects of asymmetric beams to confirm this explanation. A full investigation of the effect of beam asymmetry will be difficult and computer intensive.

While a detailed explanation of the quadrupolar effect is pending, it is important to have as much confidence as possible that a large anomaly in *WMAP* does not bias the estimated power spectrum. It is reassuring (item 5 above) that the angle-averaged power spectrum appears to be statistically isotropic, suggesting that the power spectrum is “blind” to the effect (or, less sensitive to beam asymmetries, assuming that is the cause). Furthermore, if beam asymmetry does turn out to explain the quadrupolar effect, then the analysis in Appendix B of Hinshaw et al. (2007) shows independently that the power spectrum bias due to beam asymmetry is small. Nevertheless, it is important to follow up on the studies to date, and we plan to do so in the future.

10. Conclusions

In the context of this paper we take an “anomaly” to refer to a statistically unacceptable fit of the Λ CDM model to the C_l data, a statistically significant deviation of the a_{lm} s from Gaussian random phases, or correlations between the a_{lm} . We are not concerned here with the current uncertainty range of parameter values allowed by the Λ CDM model or with whether an alternative model is also an acceptable fit to the data.

Numerous claims of *WMAP* CMB anomalies have been published. We find that there are a few valuable principles to apply to assess the significance of suspected anomalies: (1) Human eyes and brains are excellent at detecting visual patterns, but poor at assessing probabilities. Features seen in the *WMAP* maps, such as the large Cold Spot I near the Galactic center region, can stand out as unusual. However, the likelihood of such features can not be discerned by visual inspection of our particular realization of the universe. (2) Monte Carlo simulations are an invaluable way to determine the expected deviations within the Λ CDM model. Claims of anomalies without Monte Carlo simulations are necessarily weak claims. (3) Some parameters are weak discriminants of cosmology because they take on a broad range of values for multiple realizations of the same model. (4) *A posteriori* choices can have a substantial effect on the estimated significance of features. For example, it is not unexpected to find a 2σ feature when analyzing a rich data set in a number of different ways. However, to assess whether a particular 2σ feature is interesting, one is often tempted to narrow in on it to isolate its behavior. That process involves *a posteriori* choices that amplify the apparent significance of the feature.

Shortly after the *WMAP* sky maps became available, one of the authors (L.P.) noted that the initials of Stephen Hawking appear in the temperature map, as seen in Figure 17. Both the “S” and “H” are beautifully vertical in Galactic coordinates, spaced consistently just above the $b = 0$ line. We pose the question, what is the probability of this occurrence? It is certainly infinitesimal; in fact, much less likely than several claimed cosmological anomalies. Yet, we do not take this anomaly seriously because it is silly. The Stephen Hawking initials highlight the problem with *a posteriori* statistics. By looking at a rich data set in multiple different ways, *unlikely events are expected*. The search for statistical oddities must be viewed differently from tests of pre-determined hypotheses. The data have the power to support hypothesis testing rooted in ideas that are independent of the *WMAP* data. We can ask which of two well-posed theoretical ideas is best supported by the data. Much of the *WMAP* analysis happens in a different context asking, “What oddities can I find in the data.”

For example, no one had predicted that low- l multipoles might be aligned. Rather, this followed from looking into the statistical properties of the maps. Simulations, both by the *WMAP* and others agree that this is a highly unusual occurrence for the standard Λ CDM cosmology. Yet, a large fraction of simulated skies will likely have some kind of oddity. The key is whether the oddity is specified in advance.

The search for oddities in the data is essential for testing the model. The success of the model

makes these searches even more important. A detection of any highly significant *a posteriori* feature could become a serious challenge for the model. The less significant features discussed in this paper provided the motivation for considering alternative models and developing new analysis of *WMAP* (and soon *Planck*) data. The oddities have triggered proposed new observations that can further test the models.

It is often difficult to assess the statistical claims. It may well be that an oddity could be found that motivates a new theory, which then could be tested as a hypothesis against Λ CDM. The data support these comparisons. Of course, other cosmological measurements must also play a role in testing new hypotheses. No CMB anomaly reported to date has caused the scientific community to adopt a new standard model of cosmology, but claimed anomalies have been used to provoke thought and to search for improved theories.

We find that Cold Spot I does not result from Galactic foregrounds, but rather forms the northernmost part of one of four cool “fingers” in the southern sky. Its amplitude and extent are not unusual for Λ CDM. In fact, structures with this nature are expected.

We find that Cold Spot II is at the southernmost end of a different one of the southern fingers, and it has been shown not to be an anomalous fluctuation.

We find that the amplitude of the $l = 2$ quadrupole component is not anomalously low, but well within the 95% confidence range.

We conclude that there is no lack of large-scale CMB power over the full *WMAP* sky. The low value of the S -statistic integral over the large-angle correlation function on the cut sky results from a posterior choice of a sub-optimal (i.e, not full sky) statistic ($S_{1/2}$, a chance alignment of the Galactic plane cut with CMB signal, and a chance alignment of primary CMB fluctuation features with secondary ISW features from the local density distribution).

We find that the quadrupole and octupole are aligned to a remarkable degree, but that this alignment is not due to a single feature in the map or even a pair of features. The alignment does not appear to be due to a void, for example. We find that the alignment is intimately associated with the fingers of the large scale anisotropy visible in the southern sky, and it results from the statistical combination of fluctuations over the full sky. There is also evidence that the alignment is due, in part, to a coincidental alignment of the primary anisotropy with the secondary anisotropy from the local density distribution through the ISW effect. At the present time the remarkable degree of alignment appears to be no more than a chance occurrence, discovered *a posteriori* with no motivating theory. A new compelling theory could change this conclusion.

There is a portion of the power spectrum where there is a marginally significant lack of odd multipole power relative to even multipole power, but overall the *WMAP* data are well-fit by the Λ CDM model. There is no systematic error that we are aware of that could cause the even power excess, nor are there any cosmological effects that would do so. We conclude that the even excess is likely a statistical fluctuation that was found *a posteriori*. No motivating theory for this

phenomenon is known.

We find that claims of hemispherical and/or dipole asymmetries have suffered from *a posteriori* choices. After carrying out an analysis in a manner that avoids *a posteriori* bias, we find that the evidence for a hemispherical power asymmetry is weak.

Evidence has been reported for a significant quadrupolar power asymmetry that does not appear to be cosmological in origin, and most likely results from an incomplete propagation of beam asymmetries. A careful analysis will be a subject of future work.

We have examined selected claims of CMB anomalies, but this paper is not a comprehensive review article and we have not attempted to address every anomaly paper in the literature. However, we can extend our results by recognizing that various claims of anomalies are not necessarily independent from those that have been examined. For example, there are smaller-scale consequences of the fact that the value of the quadrupole on our sky is low (but not anomalously low) compared with the maximum likelihood expected value for the best-fit Λ CDM model. The amplitude of moderate to small scale hot and cold spots are expected to be reduced statistically relative to their predicted amplitude in the best-fit Λ CDM model due to the reduced contribution from the quadrupole component. That is, the temperature anisotropy in a specific spot has contributions from a range of multipoles, including the quadrupole contribution. Since the power in the multipoles scales as l^{-2} , the low quadrupole value in our sky statistically reduces the amplitude stretch of hot and cold spots in the map of our sky. This should help, at least in part, to explain the results of Hou et al. (2009), Monteserin et al. (2008), Ayaita, Weber & Wetterich (2010), and Larson & Wandelt (2004). Another example where the results from this paper may relate to a claimed anomaly is our report of a quadrupolar power asymmetry. This may be the effect that was detected by Wiaux et al. (2006) through an analysis with a second gaussian derivative.

The *WMAP* mission is made possible by the support of the NASA Science Mission Directorate. This research has made use of NASA’s Astrophysics Data System Bibliographic Services. We thank Jon Urrestilla for sharing the texture power spectrum with us. We also acknowledge use of the HEALPix (Gorski et al. 2005), CAMB (Lewis et al. 2000), and CMBFAST (Seljak & Zaldarriaga 1996) packages. We are very grateful to Duncan Hanson for useful discussions.

APPENDIX A

DIPOLAR AND QUADRUPOLAR POWER ASYMMETRY ESTIMATORS

A. Estimators for dipolar and quadrupolar anomalies

In this appendix, we present the detailed construction of the estimators used to study dipole power asymmetry in §8 and quadrupolar dependence of the two-point function in §9.

A.1. Bipolar power spectrum

If statistical isotropy is assumed, then the two-point function of the CMB is parameterized by the power spectrum C_ℓ :

$$\langle a_{\ell_1 m_1} a_{\ell_2 m_2} \rangle = (-1)^{m_1} C_{\ell_1} \delta_{\ell_1 \ell_2} \delta_{m_1, -m_2} \quad (\text{A1})$$

The bipolar power spectrum, introduced in Hajian & Souradeep (2003), is a formalism for analogously parameterizing the two-point function if the assumption of statistical isotropy is relaxed. If we decompose the two-point function $\langle a_{\ell_1 m_1} a_{\ell_2 m_2} \rangle$ into a sum of terms which transform under rotations with total angular momentum L , then we arrive at the following expansion:

$$\langle a_{\ell_1 m_1} a_{\ell_2 m_2} \rangle = \sqrt{(2\ell_1 + 1)(2\ell_2 + 1)} \sum_{LM} A_{\ell_1 \ell_2}^{LM*} \sqrt{2L + 1} \begin{pmatrix} \ell_1 & \ell_2 & L \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & M \end{pmatrix} \quad (\text{A2})$$

This equation defines the bipolar power spectrum $A_{\ell_1 \ell_2}^{LM}$. (Note that our normalization and sign convention differs from Hajian & Souradeep (2003); this will simplify some of the equations that follow.)

The bipolar power spectrum has the following properties:

1. $A_{\ell_1 \ell_2}^{LM}$ vanishes unless $-L \leq M \leq L$, the triple (ℓ_1, ℓ_2, L) satisfies the triangle inequality, and $(\ell_1 + \ell_2 + L)$ is even.
2. Under rotations, $A_{\ell_1 \ell_2}^{LM}$ transforms as a spin- L object; under the parity operation $\hat{\mathbf{n}} \rightarrow (-\hat{\mathbf{n}})$, it transforms as $A_{\ell_1 \ell_2}^{LM} \rightarrow (-1)^L A_{\ell_1 \ell_2}^{LM}$.
3. Symmetry: $A_{\ell_2 \ell_1}^{LM} = A_{\ell_1 \ell_2}^{LM}$.
4. Reality: $A_{\ell_1 \ell_2}^{LM*} = (-1)^M A_{\ell_1 \ell_2}^{L, -M}$.

To get some intuition for the bipolar power spectrum, we now consider a series of increasingly complicated models and compute the bipolar power spectrum in each case.

Our first trivial example will be an isotropic model with power spectrum C_ℓ . In this case comparison with Eq. (A2) shows that the bipolar power spectrum is given by

$$A_{\ell_1 \ell_2}^{LM} = C_{\ell_1} \delta_{\ell_1 \ell_2} \delta_{L0} \delta_{M0} \quad (\text{A3})$$

In general, the $L = 0$ component $A_{\ell \ell}^{00}$ of the bipolar power spectrum is equal to the angle-averaged power spectrum C_ℓ . (Note that the properties above imply that the only component of the bipolar power spectrum with $L = 0$ is $A_{\ell \ell}^{00}$.) Components with $L > 0$ will parameterize deviations from statistical isotropy.

Our next example (from §8) is an anisotropic model which is obtained by applying a dipolar “sky” modulation to an isotropic CMB with power spectrum C_ℓ :

$$T(\hat{\mathbf{n}}) = \left(1 + \sum_{M=-1}^1 w_{1M} Y_{1M}(\hat{\mathbf{n}}) \right) T(\hat{\mathbf{n}})_{\text{iso}} \quad (\text{A4})$$

where $T_{\text{iso}}(\hat{\mathbf{n}})$ is an isotropic CMB realization. To first order in the modulation w_{1M} , a short calculation shows that the bipolar power spectrum is given by

$$A_{\ell\ell}^{00} = C_\ell \quad (\text{A5})$$

$$A_{\ell-1,\ell}^{1M} = A_{\ell,\ell-1}^{1M} = \frac{w_{1M}(C_{\ell-1} + C_\ell)}{(4\pi)^{1/2}} \quad (\text{A6})$$

with all other components zero. The modulation does not change the sky-averaged power spectrum $A_{\ell\ell}^{00}$, but the power spectrum in local patches near the two poles will be different.

A very similar example is the quadrupolar “sky” modulation:

$$T(\hat{\mathbf{n}}) = \left(1 + \sum_{M=-2}^2 w_{2M} Y_{2M}(\hat{\mathbf{n}}) \right) T(\hat{\mathbf{n}})_{\text{iso}} \quad (\text{A7})$$

with bipolar power spectrum given by:

$$A_{\ell\ell}^{00} = C_\ell \quad (\text{A8})$$

$$A_{\ell\ell}^{2M} = \frac{w_{2M} C_\ell}{\pi^{1/2}} \quad (\text{A9})$$

$$A_{\ell-2,\ell}^{2M} = A_{\ell,\ell-2}^{2M} = \frac{w_{2M}(C_{\ell-2} + C_\ell)}{(4\pi)^{1/2}} \quad (\text{A10})$$

In the quadrupolar case, the most general anisotropic two-point function is parameterized by two ℓ -dependent quantities ($A_{\ell-2,\ell}^{2M}$ and $A_{\ell\ell}^{2M}$), in contrast to the dipole case. A short calculation shows that the quadrupolar anisotropy in the power spectrum is proportional to $2A_{\ell-2,\ell} + A_{\ell\ell}$, so that a model which satisfies $A_{\ell\ell} \approx -2A_{\ell-2,\ell}$ has a roughly isotropic power spectrum, even though the two-point function contains a component which transforms under rotations as a quadrupole.

Our final example is the anisotropic early universe model from Ackerman et al. (2007), in which the initial adiabatic curvature fluctuation $\zeta(\mathbf{k})$ is modulated in Fourier space as follows:

$$\zeta(\mathbf{k}) = \left[1 + \sum_{M=-2}^2 w_{2M} Y_{2M}(\hat{\mathbf{k}}) \right] \zeta(\mathbf{k})_{\text{iso}} \quad (\text{A11})$$

The bipolar power spectrum of this model is

$$A_{\ell_1\ell_2}^{2M} = \frac{i^{\ell_1-\ell_2}}{(4\pi)^{1/2}} w_{2M} \int \frac{2k^2 dk}{\pi} \Delta_{\ell_1}(k) \Delta_{\ell_2}(k) P(k) \quad (\text{A12})$$

where $\Delta_\ell(k)$ is the angular CMB transfer function.

A.2. Estimators: General Construction

In this subsection we will construct estimators for the bipolar power spectrum and related quantities. We will assume that the bipolar power spectrum of the CMB is a linear combination of

N “template” shapes A_1, A_2, \dots, A_N .

$$A_{\ell_1 \ell_2}^{LM} = \sum_{i=1}^N t_i (A_i)_{\ell_1 \ell_2}^{LM} \quad (\text{A13})$$

where the coefficients t_i are to be estimated from data. (The choice of template shapes will be discussed shortly.) We assume that our template shapes satisfy $A_{\ell \ell}^{00} = 0$, i.e. the templates parameterize deviations from isotropy, not changes in the power spectrum.

A lengthy but straightforward calculation shows that the minimum-variance estimator \hat{t}_i for the template coefficients which is unbiased (i.e. $\langle \hat{t}_i \rangle = t_i$) is given by

$$\hat{t}_i[a] = F_{ij}^{-1} (\mathcal{E}_j[a] - \mathcal{N}_j) \quad (\text{A14})$$

where the quantity $\mathcal{E}_j[a]$ is defined by

$$\begin{aligned} \mathcal{E}_j[a] = & \frac{1}{2} \sum_{\ell_1 m_1 \ell_2 m_2 LM} (A_j)_{\ell_1 \ell_2}^{LM*} \sqrt{(2\ell_1 + 1)(2\ell_2 + 1)(2L + 1)} \begin{pmatrix} \ell_1 & \ell_2 & L \\ 0 & 0 & 0 \end{pmatrix} \\ & \times \begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & M \end{pmatrix} (C^{-1}a)_{\ell_1 m_1}^* (C^{-1}a)_{\ell_2 m_2}^* \end{aligned} \quad (\text{A15})$$

and the matrix F_{ij} and vector \mathcal{N}_j are defined by

$$F_{ij} = \text{Cov}(\mathcal{E}_i[a], \mathcal{E}_j[a]) \quad (\text{A16})$$

$$\mathcal{N}_j[a] = \langle \mathcal{E}_j[a] \rangle \quad (\text{A17})$$

where the covariance and expectation value are taken over isotropic realizations of the noisy CMB $a_{\ell m}$.

The estimator in Eq. (A14) can be specialized to measure different types of statistical isotropy by making different choices of template shapes (A_i). For example, consider the dipole modulation, parameterized as in Eq. (12) by a three-vector w_i . We assume that multipoles $2 \leq \ell \leq \ell_{\text{mod}}$ are modulated, and multipoles $\ell > \ell_{\text{mod}}$ are unmodulated. The bipolar power spectrum is given by

$$A_{\ell_1 \ell_2}^{LM} = \sum_{i=1}^3 w_i (A_i)_{\ell_1 \ell_2}^{LM} \quad (\text{A18})$$

where the template shapes A_i are defined by

$$(A_i)_{\ell_1 \ell_2}^{LM} = \begin{cases} (\sqrt{4\pi}/3)(C_{\ell_1} + C_{\ell_2})Y_{1M}^*(\hat{e}_i) & \text{if } |\ell_1 - \ell_2| = 1 \text{ and } \ell_i \leq \ell_{\text{mod}} \\ 0 & \text{otherwise} \end{cases} \quad (\text{A19})$$

When specialized to these three template shapes, Eq. (A14) gives the optimal quadratic estimator \hat{w}_i for the three components of the modulation. (A closely related estimator has also been constructed

in Dvorkin et al. (2008).) We also construct an estimator κ_1 for the total amplitude of the dipole modulation, irrespective of its direction, by

$$\kappa_1 = \sum_i \hat{w}_i^2 \quad (\text{A20})$$

We have used this estimator in §8 to assess statistical significance of the dipole modulation, by comparing the WMAP value of κ_1 to an ensemble of simulations.

The quadrupole modulation can be treated analogously. For example, in Fig. 16 we have shown estimates of $A_{\ell-2,\ell}^{20}$ and $A_{\ell\ell}^{20}$ in each bin, by taking $(10N_{\text{bins}})$ template shapes, corresponding to the five components of the two bipolar power spectra in each ℓ bin. As another example, if an optimal estimator (unbinned in ℓ) for the primordial modulation in Eq. (A11) is desired, one would take five template shapes corresponding to the five components of w_{2M} , given by Eq. (A12).

A.3. Relation to likelihood formalism

We conclude this appendix by showing how the optimal estimator is related to the maximum likelihood formalism, for the dipole and quadrupole cases.

First consider the dipole modulation. The likelihood function $\mathcal{L}[a|w]$ for the modulation w_i , given noisy CMB data $a_{\ell m}$, is given by:

$$\mathcal{L}[a|w] = (\text{const.}) \times \text{Det}^{-1/2}[M(w)SM(w) + N] \exp\left(-\frac{1}{2}a^\dagger[M(w)SM(w) + N]^{-1}a\right) \quad (\text{A21})$$

where $M(w)$ is the operator which applies the modulation to a harmonic-space map, defined by:

$$(M(w)a)_{\ell m} = \int d^2\hat{\mathbf{n}} Y_{\ell m}^*(\hat{\mathbf{n}}) (1 + w_i \hat{\mathbf{n}}_i) \left(\sum_{\ell' m'} a_{\ell' m'} Y_{\ell' m'}(\hat{\mathbf{n}}) \right) \quad (\text{A22})$$

We will show that the modulation w which maximizes the likelihood $\mathcal{L}[w|a]$ is equal to value of the optimal quadratic estimator \hat{w}_i defined in the previous subsection, under two approximations that will be discussed further.

First suppose that the maximum likelihood modulation is small, so that the Taylor expansion of $(\log \mathcal{L})$ to second order in w_i is an accurate approximation near maximum likelihood. The Taylor expansion is given by:

$$\log \mathcal{L}[a|w] \approx -\frac{1}{2} \mathcal{H}_{ij}[a] w_i w_j + \mathcal{G}_i[a] w_i + (\text{const.}) \quad (\text{A23})$$

where we have defined:

$$\mathcal{H}_{ij}[a] = -\left(\frac{\partial^2}{\partial w_i \partial w_j} \right)_{w=0} (\log \mathcal{L}[a|w]) \quad (\text{A24})$$

$$\begin{aligned}
 &= -a^\dagger C^{-1} M_i S M_j C^{-1} a + a^\dagger C^{-1} \{M_i, S\} C^{-1} \{M_j, S\} C^{-1} a \\
 &\quad + \text{Tr}(C^{-1} M_i S M_j) - \frac{1}{2} \text{Tr}(C^{-1} \{M_i, S\} C^{-1} \{M_j, S\}) \quad (\text{A25})
 \end{aligned}$$

$$\mathcal{G}_i[a] = a^\dagger C^{-1} M_i S C^{-1} a - \text{Tr}[C^{-1} M_i S] \quad (\text{A26})$$

and $M_i = \partial M(w)/\partial w_i$. Second, we assume that we can approximate the $a_{\ell m}$ -dependent quantity $\mathcal{H}_{ij}[a]$ by its expectation value:

$$\mathcal{H}_{ij}[a] \approx \langle \mathcal{H}_{ij} \rangle = \frac{1}{2} \text{Tr}(C^{-1} \{M_i, S\} C^{-1} \{M_j, S\}) \quad (\text{A27})$$

Under these approximations the maximum likelihood modulation is given by

$$(w_i)_{\text{ML}} = \langle \mathcal{H}_{ij} \rangle^{-1} \mathcal{G}_j[a] \quad (\text{A28})$$

We would like to compare this expression to the optimal quadratic estimator \hat{w}_i defined in Eq. (A14), in the special case where the template shapes are given by Eq. (A19). In this special case, a short calculation shows that $\langle \mathcal{H}_{ij} \rangle = F_{ij}$, and $\mathcal{G}_j[a] = \mathcal{E}_j[a] - \mathcal{N}_j[a]$, which implies $(w_i)_{\text{ML}} = \hat{w}_i$.

We comment briefly on the two approximations made in this subsection, namely the second-order Taylor approximation in Eq. (A23) and the approximation $\mathcal{H}_{ij}[a] \approx \langle \mathcal{H}_{ij} \rangle$. By the central limit theorem, both approximations are expected to become accurate in the limit where the number of CMB modes which contribute to the estimators is large. The analyses in §8, 9 have taken wide ℓ bins (either the ℓ range $2 \leq \ell \leq \ell_{\text{mod}}$ or a series of bins with $\Delta\ell = 50$) so these approximations should be very accurate and there should be little difference in practice between a likelihood estimator and the optimal quadratic estimator, although we have not tested this directly. In any case, the optimal quadratic estimator has several important advantages over a likelihood estimator, as described in §8.

REFERENCES

- Ackerman, L., Carroll, S. M., & Wise, M. B. 2007, *Phys. Rev. D*, 75, 083502
- Ayaita, Y., Weber, M. & Wetterich, C. 2010, *Phys. Rev. D*, 81, 023507
- Bennett, C. L., et al. 1992, *ApJ*, 396, L7
- Bennett, C. L. et al. 2003a, *ApJ*, 583, 1
- Bennett, C. L., et al. 2003b, *ApJS*, 148, 1
- Bridges, M., McEwen, J. D., Cruz, M., Hobson, M. P., Lasenby, A. N., Vielva, P., & Martinez-Gonzalez, E. 2007, arXiv:0712.1789
- Cayón, L. 2010, arXiv:1001.4680

- Chiang, L.-Y. and Naselsky, P. D. & Coles, P. 2007, *ApJ*, 664, 8
- Copi, C. J., Huterer, D., Schwarz, D. J., & Starkman, G. D. 2007, *Phys. Rev. D*, 75, 023507
- . 2009, *MNRAS*, 399, 295
- Cruz, M., Cayón, L., Martínez-González, E., Vielva, P., & Jin, J. 2007a, *ApJ*, 655, 11
- Cruz, M., Turok, N., Vielva, P., Martínez-González, E., & Hobson, M. 2007b, *Science*, 318, 1612
- Cruz, M., Martínez-González, E., Vielva, P., Diego, J. M., Hobson, M., & Turok, N. 2008, *MNRAS*, 390, 913
- Cruz, M., Tucci, M., Martínez-González, E., & Vielva, P. 2006, *MNRAS*, 369, 57
- Donogue, J. F., Dutta, K. & Ross, A., *Phys. Rev. D*, 80, 023526
- Dunkley, J., et al. 2009, *ApJS*, 180, 306
- Dvorkin, C., Peiris, H. V., & Hu, W. 2008, *Phys. Rev. D*, 77, 063008
- Efstathiou, G., Ma, Y., & Hanson, D. 2009, arXiv:0911.5399
- Erickcek, A. L., Hirata, C. M., & Kamionkowski, M. 2009, *Phys. Rev. D*, 80, 083507
- Erickcek, A. L., Kamionkowski, M. & Carroll, S. M. 2008, *Phys. Rev. D*, 78, 123520
- Eriksen, H. K., Banday, A. J., Górski, K. M., Hansen, F. K., & Lilje, P. B. 2007, *ApJ*, 660, L81
- Eriksen, H. K., Hansen, F. K., Banday, A. J., Górski, K. M., & Lilje, P. B. 2004, *ApJ*, 605, 14
- Finkbeiner, D. P., Davis, M., & Schlegel, D. J. 1999, *ApJ*, 524, 867
- Francis, C. L. & Peacock, J. A. 2009, arXiv:0909.2495
- Freeman, P. E., Genovese, C. R., Miller, C. J., Nichol, R. C., & Wasserman, L. 2006, *ApJ*, 638, 1
- Gold, B. et al. 2010, in preparation
- Gordon, C. 2007, *Astrophys. J.*, 656, 636
- Gordon, C., Hu, W., Huterer, D. & Crawford, T. 2005, *Phys. Rev. D*, 72, 103002
- Gorski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, *ApJ*, 622, 759
- Granett, B. R., Szapudi, I., & Neyrinck, M. C. 2010, *ApJ*, 714, 825
- Groeneboom, N. E., Ackerman, L., Kathrine Wehus, I., & Eriksen, H. K. 2009, arXiv:0911.0150

- Groeneboom, N. E. & Eriksen, H. K. 2009, *ApJ*, 690, 1807
- Hajian, A. & Souradeep, T. 2003, *Astrophys. J.*, 597, L5
- Hajian, A., Souradeep, T., & Cornish, N. 2005, *ApJ*, 618, L63
- Hansen, F. K., Banday, A. J., Górski, K. M., Eriksen, H. K., & Lilje, P. B. 2009, *ApJ*, 704, 1448
- Hansen, F. K., Cabella, P., Marinucci, D., & Vittorio, N. 2004, *ApJ*, 607, L67
- Hanson, D. & Lewis, A. 2009, *Phys. Rev. D*, 80, 063004
- Hinshaw, G., Banday, A. J., Bennett, C. L., Górski, K. M., Kogut, A., Smoot, G. F., & Wright, E. L. 1996, *ApJ*, 464, L17
- Hinshaw, G., et al. 2003, *ApJS*, 148, 63
- . 2007, *ApJS*, 170, 288
- . 2009, *ApJS*, 180, 225
- Hirata, C. M. 2009, *Journal of Cosmology and Astro-Particle Physics*, 9, 11
- Hivon, E., Górski, K. M., Netterfield, C. B., Crill, B. P., Prunet, S., & Hansen, F. 2002, *ApJ*, 567, 2
- Hoftuft, J., Eriksen, H. K., Banday, A. J., Górski, K. M., Hansen, F. K., & Lilje, P. B. 2009, *ApJ*, 699, 985
- Hou, Z., Banday, A. J., & Gorski, K. M. 2009, *MNRAS*, 396, 1273
- Jarosik, N., et al. 2003, *ApJS*, 148, 29
- Jarosik, N., et al. 2007, *ApJS*, 170, 263
- Jarosik, N. et al. 2010, in preparation
- Komatsu, E., et al. 2009, *ApJS*, 180, 330
- Komatsu, E. et al. 2010, in preparation
- Land, K. & Magueijo, J. 2005a, *MNRAS*, 357, 994
- . 2005b, *Phys. Rev. D*, 72, 101302
- Larson, D. et al. 2010, in preparation
- Larson, D. L. & Wandelt, B. D. 2004, *ApJ*, 613, L85
- Lewis, A., Challinor, A., & Lasenby, A. 2000, *ApJ*, 538, 473

- Monteserin, C., Barreiro, R. B., Vielva, P., Martínez-González, E., Hobson, M.P., & Lasenby, A. N. 2008, *MNRAS*, 387, 209
- Mota, D. F., Shaw, D. J., & Silk, J. 2008, *ApJ*, 675, 29
- Mukherjee, P. & Wang, Y. 2004, *ApJ*, 613, 51
- O’Dwyer, I. J. et al. 2004, *Astrophys. J.*, 617, L99
- de Oliveira-Costa, A. & Tegmark, M. 2006, *Phys. Rev. D*, 74, 023005
- Page, L. et al. 2007, *ApJ*, 170, 335
- Park, C.-G., Park, C. & Gott, III, J. R. 2007, *ApJ*, 660, 959
- Pullen, A. P. & Kamionkowski, M. 2007, *Phys. Rev. D*, 76, 103529
- Rudnick, L., Brown, S., & Williams, L. R. 2007, *ApJ*, 671, 40
- Schwarz, D. J., Starkman, G. D., Huterer, D., & Copi, C. J. 2004, *Physical Review Letters*, 93, 221301
- Seljak, U. & Zaldarriaga, M. 1996, *ApJ*, 469, 437
- Smith, K. M. & Huterer, D. 2010, *MNRAS*, 403, 2
- Smith, K. M., Zahn, O., & Doré, O. 2007, *Phys. Rev. D*, 76, 043510
- Spergel, D. N., et al. 2003, *ApJS*, 148, 175
- Tegmark, M., de Oliveira-Costa, A., & Hamilton, A. J. 2003, *Phys. Rev. D*, 68, 123523
- Tye, S. & Xu, J. 2010, *Phys. Lett. B*, 683, 326
- Urrestilla, J. et al. 2008, *JCAP*, 0807, 010
- Vielva, P., Martínez-González, E., Barreiro, R. B., Sanz, J. L., & Cayón, L. 2004, *ApJ*, 609, 22
- Wiaux, Y., Vielva, P., Martínez-González, E. & Vandergheynst, P. 2006 *Phys. Rev. Lett.*, 96, 151303
- Wehus, I. K., Ackerman, L., Eriksen, H. K., & Groeneboom, N. E. 2009, *ApJ*, 707, 343
- Weiland, J. L. et al. 2010, in preparation
- Wright, E. L. 2004, in *Measuring and Modeling the Universe*, ed. W. L. Freedman, 291
- Zhang, R. & Huterer, D. 2010, *Astropart. Phys.*, 33, 69




Fig. 1.— (*top*) A large colder than average region, highlighted by the white curve, appears prominently on the raw V-band temperature map. The full sky map is shown with the Galactic plane horizontal across the center of the map with the Galactic Center at the center of the displayed projection. (*bottom*) With the foreground signals strongly suppressed by the ILC technique, the highlighted cold spot is seen to be at least as prominent. It is offset from the Galactic Center in both latitude and longitude. This fact, combined with the fact that the clearly effective ILC foreground reduction does not diminish this feature, establishes that this is a CMB fluctuation and not a foreground effect. This feature is not anomalous in that simulated realizations of Λ CDM model skies routinely produce features like this.




Fig. 2.— A visual inspection of the ILC map reveals four elongated valleys of cooler temperature that stretch from about the Galactic equator to nearly the south Galactic pole. Ridges of warmer than average temperature lie between the cooler fingers. These features are a consequence of large scale power in the southern sky. It is more difficult to discern as much large scale power in the northern sky. Cold Spot I is located near the northernmost part of one of the fingers, while Cold Spot II (within the red curve) is near the southernmost part of another finger.




Fig. 3.— The curve is a Blackwell-Rao estimate of the relative likelihood of the quadrupole power $l(l+1)C_2/2\pi$ in μK^2 from *WMAP*. The *WMAP* ILC data were smoothed to 5° and the KQ85y7 mask was used, both degraded to res 5. The Gibbs sampling produced a likelihood that has been marginalized over all other multipoles. The highly non-Gaussian likelihood distribution is characteristic of the lowest- l multipoles. For $l > 32$ the curves become nearly Gaussian. The vertical line with the label “ ΛCDM ” is the expected quadrupole from the full power spectrum ΛCDM model best-fit to the *WMAP* data. The maximum likelihood of the *WMAP* measured $l = 2$ quadrupole is at the vertical dotted line. These two values are consistent to well within the 95% confidence region. The *WMAP* quadrupole is not anomalously low.




Fig. 4.— The cumulative distribution function of the quadrupoles from the Gibbs sampling based on 300,000 points. The vertical line is the predicted Λ CDM model quadrupole value. The cumulative probability is 0.824 where the vertical line crosses the cumulative distribution function . Since the expected quadrupole from the Λ CDM model is well within the 95% confidence range of the measured quadrupole, accounting for noise and cosmic variance, we conclude that the measured quadrupole is not anomalously low.




Fig. 5.— The angular correlation function of the full sky *WMAP* ILC map is shown (heavy black curve). For comparison, the angular correlation function for the best fit ΛCDM model is also shown (thin black curve), along with the associated 68% and 95% confidence ranges, as determined by Monte Carlo simulations. The angular correlation function of the full sky map is seen to be within the 95% confidence range of the best fit ΛCDM model. This angular correlation function was computed from the C_l power spectrum, but is nearly indistinguishable from a pixel pair computation. Either way, there is no evidence of a lack of large-scale power.

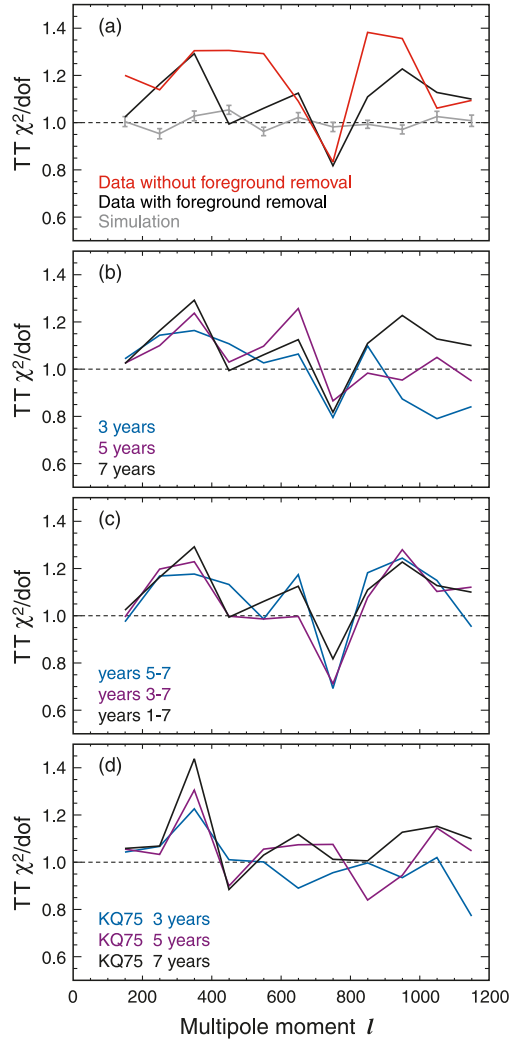


Fig. 6.— (a) The χ^2 per degree of freedom of the 7 year temperature-temperature power spectrum data relative to the best fit Λ CDM model. Light gray points are from 50 simulations that used the same Λ CDM model with the appropriate noise and cosmic variance included, where the error bars are driven by the number of simulations. The data and simulations were run through the same data analysis pipeline. These simulations were used to help fit the effective fraction of the sky, f_{sky} to use for the data analysis. (b) The χ^2 per degree of freedom is compared for the 3-year, 5-year, and 7-year maps. Small differences in the fit model have a negligible effect on these plots. The χ^2 per degree of freedom for $l \approx 300$ has been slightly growing with additional data, while other multipole moment ranges are more random with additional data. (c) The χ^2 per degree of freedom had the same *WMAP* data been taken in a reverse time order. The $l < 400$ region appears robust, while the χ^2 variations for $l > 400$ appear more random. (d) Variation of χ^2 with the choice of Galactic foreground mask appears random for $l > 400$.




Fig. 7.— The black curve is the cumulative probability of the *WMAP* temperature data based on 499 simulations. All of the simulations were drawn from the same Λ CDM model, but χ^2 was evaluated with respect to the best-fit model for each realization. Of these, 412 (82.6%) had a lower χ^2 than the vertical line at 1224.6. Thus, the *WMAP* power spectrum is statistically compatible with the model. The red curve is a χ^2 distribution with 1170 degrees of freedom, shown for comparison.

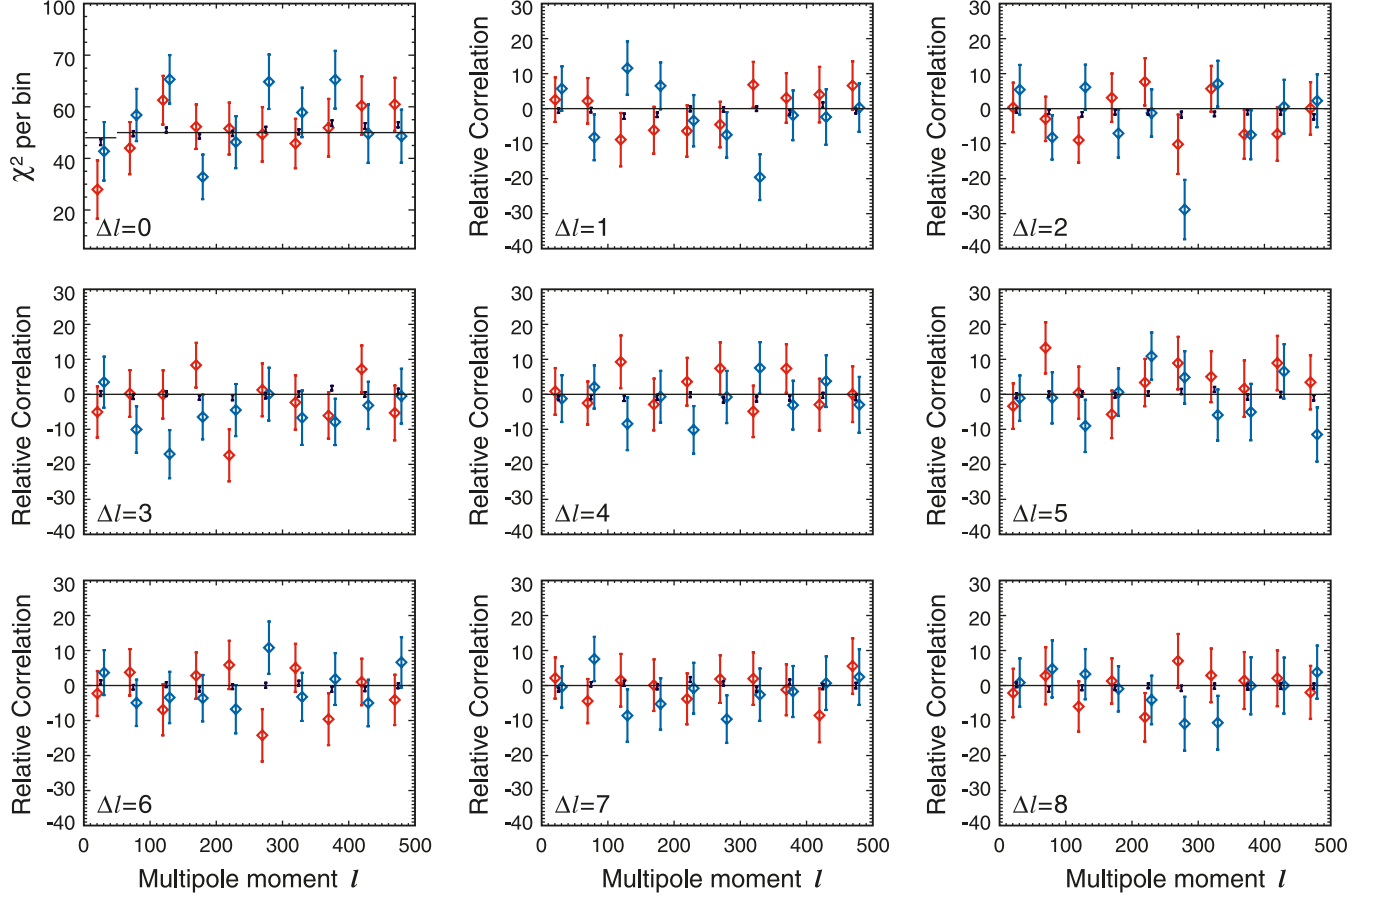


Fig. 8.— We compute $\Delta l = l - l'$ correlation properties of C_l for nearby multipoles for the *WMAP* data (blue) and comparable simulations (red). For the most part the data and simulations are in good agreement. The most discrepant correlations in C_l are for $\Delta l = 1$ near $l \sim 320$ and $\Delta l \sim 2$ near $l \sim 280$.

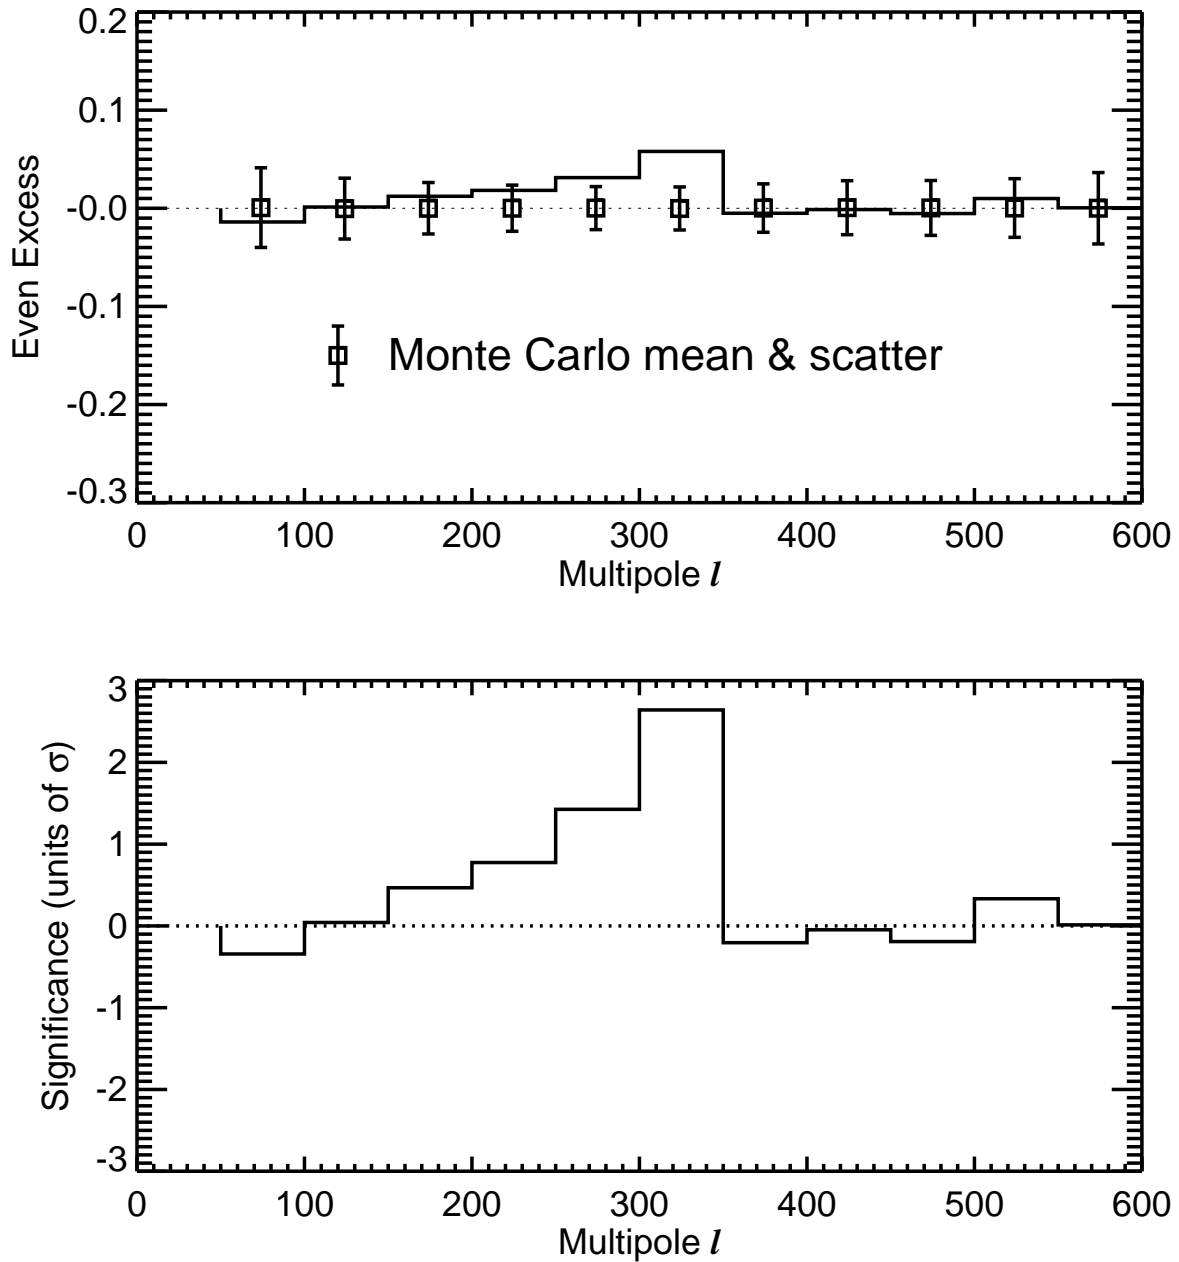


Fig. 9.— Top: Even excess \mathcal{E}_ℓ in the observed power spectrum, in bins of $\Delta\ell = 50$, compared to the mean and scatter from 11436 Monte Carlo realizations. Bottom: \mathcal{E}_ℓ as in the top plot, converted to significance units by normalizing to the Monte Carlo scatter in each bin. Only the $\ell = 250 - 299$ and $\ell = 300 - 349$ bins show a significance greater than 1σ .




Fig. 10.— Even excess \mathcal{E}_ℓ in the observed power spectrum, in the ℓ range 250 – 349, compared to a histogram of the same value as computed from 11436 Monte Carlo realizations. The probability for a random realization to exceed the observed value (PTE) is 0.26%, interpolated directly from the Monte Carlo cumulative distribution. A Gaussian with the same area and standard deviation as the histogram is over-plotted, with $\pm 1\sigma$ indications.

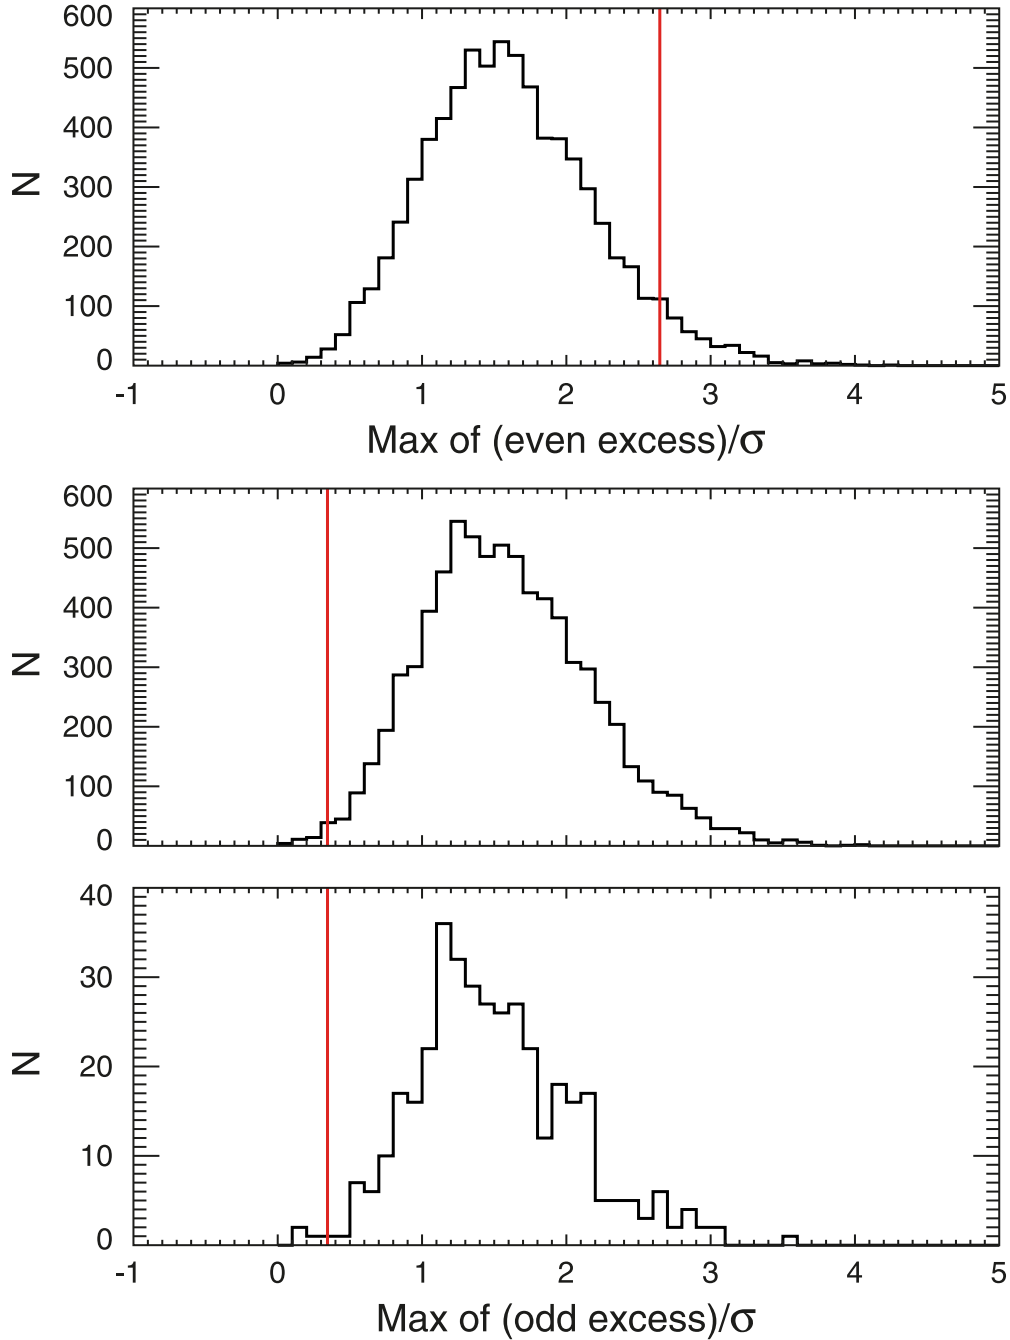


Fig. 11.— Top: Histogram of $\max(\mathcal{E}_\ell/\sigma(\mathcal{E}_\ell))$ for each Monte Carlo realization. The maximum is taken over 11 bins of $\Delta\ell = 50$, for $\ell = 50 - 599$. The red vertical line is the observed value of 2.65, with a probability to exceed of 5.11%. Middle: Histogram of $\max(-\mathcal{E}_\ell/\sigma(\mathcal{E}_\ell))$, which is the greatest significance found in any ℓ bin having an odd excess. Although the histograms are not normal distributions, the observed maximum even excess is equivalent to a $\sim 1.6\sigma$ result, and the observed maximum odd excess of 0.34 at the vertical red line is equivalent to a $\sim 2.5\sigma$ result. Bottom: The histogram is the same as the one above, but restricted to a pre-selection of the simulations where the maximum of the significance of the even excess is greater than or equal to the observed even excess.




Fig. 12.— The probability of a greater quadrupole-octupole alignment is given as a function of the alignment angle, in degrees. The probability given here does not account for the *a posteriori* selection of a multipole alignment search, nor does it take into account the choice of the quadrupole and the octupole moments in particular.




Fig. 13.— We smooth and square the ILC map, as described in the text, to make a map of $(\Delta T)^2$ to roughly visualize the anisotropy power in a sky map (A). Cold Spot I is the highest power region at this level of resolution, but several other regions also contribute substantial power at this scale. For map (A), the quadrupole and octupole are aligned to within $\sim 1^\circ$. To evaluate what regions contribute to this alignment, we mask selected gray areas as shown in the sky maps (B) through (J). Each map is labeled by the degree of quadrupole-octupole alignment remaining after the gray mask is applied. Masking Cold Spot I in (B) or (H) eliminates any significant alignment. However, keeping those regions and masking other regions, also breaks the alignment to a significant degree. We conclude that no single region or pair of regions generates the alignment. Rather the combination of power from a substantial fraction of the full sky map combine to cause the high degree of quadrupole-octupole alignment. Note that the chance alignment of CMB power with the Galactic cut region discussed in §5 is apparent in map (A).




Fig. 14.— The $l = 2$ quadrupole map and $l = 3$ octupole map are added. The combined map is then shown superposed on the ILC map from Figure 2. Note that the quadrupole and octupole components arrange themselves to match the cool fingers and the warm regions in between. The fingers and the alignment of the $l = 2$ and $l = 3$ multipoles are intimately connected.




Fig. 15.— The probability for a Monte Carlo simulation to have a larger dipole modulation than the co-added V+W *WMAP* data, as measured by the statistic κ_1 , is shown as a function of the maximum multipole moment l_{mod} that is assumed to be modulated. This can be interpreted as the statistical significance of power asymmetry for a fixed value of l_{mod} , if one does not account for possible *a posteriori* bias when choosing l_{mod} .




Fig. 16.— The quadrupolar bipolar power spectra, binned with $\Delta\ell = 50$, is shown separately for V-band and W-band *WMAP* data, using the KQ75y7 mask. Only the components of the bipolar power spectra that point along the ecliptic axis are shown (i.e. components with $M = 0$ in ecliptic coordinates). A statistically significant quadrupolar effect is seen, even for a single frequency band in a single angular bin.




Fig. 17.— The “SH” initials of Stephen Hawking are shown in the ILC sky map. The “S” and “H” are in roughly the same font size and style, and both letters are aligned neatly along a line of fixed Galactic latitude. A calculation would show that the probability of this particular occurrence is vanishingly small. Yet, there is no case to be made for a non-standard cosmology despite this extraordinarily low probability event. It is clear that the combined selection of looking for initials, these particular initials, and their alignment and location are all *a posteriori* choices. For a rich data set, as is the case with *WMAP*, there are a lot of data and a lot of ways of analyzing the data. Low probability events are guaranteed to occur. The *a posteriori* assignment of a likelihood for a particular event detected, especially when the detection of that event is “optimized” for maximum effect by analysis choices, does not result in a fair unbiased assessment. This is a recurrent issue with CMB data analysis, and is often a tricky issue and one that is difficult to overcome.