

First Year *Wilkinson Microwave Anisotropy Probe* (WMAP¹) Observations: The Angular Power Spectrum

G. Hinshaw², D. N. Spergel³, L. Verde³, R. S. Hill⁴, S. S. Meyer⁵, C. Barnes⁶, C. L. Bennett², M. Halpern⁷, N. Jarosik⁶, A. Kogut², E. Komatsu³, M. Limon^{2,8}, L. Page⁶, G. S. Tucker^{2,8,9}, J. L. Weiland⁴, E. Wollack², E. L. Wright¹⁰

Gary.F.Hinshaw@nasa.gov

ABSTRACT

We present the angular power spectrum derived from the first-year *Wilkinson Microwave Anisotropy Probe* (WMAP) sky maps. We study a variety of power spectrum estimation methods and data combinations and demonstrate that the results are robust. The data are modestly contaminated by diffuse Galactic foreground emission, but we show that a simple Galactic template model is sufficient to remove the signal. Point sources produce a modest contamination in the low frequency data. After masking ~ 700 known bright sources from the maps, we estimate residual sources contribute $\sim 3500 \mu\text{K}^2$ at 41 GHz, and $\sim 130 \mu\text{K}^2$ at 94 GHz, to the power spectrum $[l(l+1)C_l/2\pi]$ at $l = 1000$. *Systematic errors are negligible compared to the (modest) level of foreground emission.* Our best estimate of the power spectrum is derived from 28 cross-power spectra of statistically independent channels. The final spectrum is essentially independent of the noise properties of an individual radiometer. The resulting spectrum provides a definitive measurement of the CMB power spectrum, with uncertainties limited by cosmic variance, up to $l \sim 350$. The spectrum clearly exhibits a first acoustic peak at $l = 220$ and a second acoustic peak at $l \sim 540$ (Page et al. 2003b), and it provides strong support for adiabatic initial conditions (Spergel et al. 2003). Kogut et al. (2003) analyze the C_l^{TE} power spectrum, and present evidence for a relatively high optical depth, and an early period of cosmic reionization. Among other things, this implies that the temperature power spectrum has been suppressed by $\sim 30\%$ on degree angular scales, due to secondary scattering.

Subject headings: cosmic microwave background, cosmology: observations, early universe, space vehicles: instruments

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.

² Code 685, Goddard Space Flight Center, Greenbelt, MD 20771

³ Dept of Astrophysical Sciences, Princeton University, Princeton, NJ 08544

⁴ Science Systems and Applications, Inc. (SSAI), 10210 Greenbelt Road, Suite 600 Lanham, Maryland 20706

⁵ Depts. of Astrophysics and Physics, EFI and CfCP, University of Chicago, Chicago, IL 60637

⁶ Dept. of Physics, Jadwin Hall, Princeton, NJ 08544

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

⁸ National Research Council (NRC) Fellow

⁹ Dept. of Physics, Brown University, Providence, RI 02912

¹⁰ UCLA Astronomy, PO Box 951562, Los Angeles, CA 90095-1562

1. INTRODUCTION

The *Wilkinson Microwave Anisotropy Probe (WMAP)* mission was designed to measure the CMB anisotropy with unprecedented precision and accuracy on angular scales from the full sky to several arc minutes by producing maps at five frequencies from 23 to 94 GHz. The *WMAP* satellite mission (Bennett et al. 2003a) employs a matched pair of 1.4m telescopes (Page et al. 2003c) observing two areas on the sky separated by $\sim 141^\circ$. A differential radiometer (Jarosik et al. 2003a) with a total of 10 feeds for each of the two sets of optics (Barnes et al. 2003; Page et al. 2003c) measures the difference in sky brightness between the two sky pixels. The satellite is deployed at the Earth-Sun Lagrange point, L_2 , and observes the sky with a compound spin and precession that covers the full sky every six months. The differential data are processed on the ground to produce full sky maps of the CMB anisotropy (Hinshaw et al. 2003b).

Full sky maps provide the smallest record of the CMB anisotropy without loss of information. They permit a wide variety of statistics to be computed from the data – one of the most fundamental is the angular power spectrum of the CMB. Indeed, if the temperature fluctuations are Gaussian, with random phase, then the angular power spectrum provides a *complete* description of the statistical properties of the CMB. Komatsu et al. (2003) have analyzed the first-year *WMAP* sky maps to search for evidence of non-Gaussianity and find none, aside from a modest level of point source contamination which we account for in this paper. Thus, the measured power spectrum may be compared to predictions of cosmological models to develop constraints on model parameters.

This paper presents the angular power spectrum obtained from the first-year *WMAP* sky maps. Companion papers present the maps and an overview of the basic results (Bennett et al. 2003b), and describe the foreground removal process that precedes the power spectrum analysis (Bennett et al. 2003c). Spergel et al. (2003), Verde et al. (2003), Peiris et al. (2003), Kogut et al. (2003), and Page et al. (2003b) discuss the implications of the *WMAP* power spectrum for cosmological parameters and carry out a joint analysis of the *WMAP* spectrum together with other CMB data and data from large-scale structure probes. Hinshaw et al. (2003b), Jarosik et al. (2003b), Page et al. (2003a), Barnes et al. (2003), and Limon et al. (2003) discuss the data processing, the radiometer performance, the instrument beam characteristics and the spacecraft in-orbit performance, respectively.

A sky map $\Delta T(\mathbf{n})$ defined over the full sky can be decomposed in spherical harmonics

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

with

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

where \mathbf{n} is a unit direction vector, and $Y_{lm}(\mathbf{n})$ are the spherical harmonic functions evaluated in the direction \mathbf{n} .

If the CMB temperature fluctuation ΔT is Gaussian distributed, then each a_{lm} is an independent Gaussian deviate with

$$\langle a_{lm} \rangle = 0, \quad (3)$$

and

$$\langle a_{lm} a_{l'm'}^* \rangle = \delta_{ll'} \delta_{mm'} C_l, \quad (4)$$

where C_l is the ensemble average power spectrum predicted by models, and δ is the Kronecker symbol. The actual power spectrum realized in our sky is

$$C_l^{\text{sky}} = \frac{1}{2l+1} \sum_{m=-l}^l |a_{lm}|^2. \quad (5)$$

In the absence of noise, and with full sky coverage, the right hand side of equation (5) provides an unbiased estimate of the underlying theoretical power spectrum, which is limited only by cosmic variance. However, realistic CMB anisotropy measurements contain noise and other sources of error that cause the quadratic estimator in equation (5) to be biased. In addition, while *WMAP* measures the anisotropy over the full sky, the data near the Galactic plane are sufficiently contaminated by foreground emission that only a portion of the sky ($\sim 85\%$) can be used for CMB power spectrum estimation. Thus, the integral in equation (2) cannot be evaluated as such, and other methods must be found to estimate C_l .

In Appendix A, we review two methods that have appeared in the literature for estimating the angular power spectrum in the presence of instrument noise and sky cuts. The first (Hivon et al. 2002) is a quadratic estimator that evaluates equation (2) on the *cut* sky yielding a “pseudo power spectrum” \tilde{C}_l . The ensemble average of this quantity is related to the true power spectrum, C_l by means of a mode coupling matrix $G_{ll'}$ (Hauser & Peebles 1973). The second method (Oh et al. 1999) uses a maximum likelihood approach optimized for fast evaluation with *WMAP*-like data. In Appendix A we demonstrate that the two methods produce consistent results.

A quadratic estimator offers the possibility of computing both an “auto-power” spectrum, proportional to $\sum_m |a_{lm}|^2$, and a “cross-power” spectrum, proportional to $\sum_m a_{lm}^i a_{lm}^{j*}$, where the a_{lm} coefficients are estimated from two independent CMB maps, i and j . This latter form has the advantage that, if the noise in the two maps is uncorrelated, the quadratic estimator is not biased by the noise. For all cosmological analyses, we use only the cross-power spectra between statistically independent channels. As a result, the angular spectra are, for all intents and purposes, independent of the noise properties of an individual radiometer. This is analogous to interferometric data which exhibits a high degree of immunity to systematic errors. The precise form of the estimator we use is given in Appendix A.

The plan of this paper is as follows. In §2 we review the properties of the *WMAP* instrument and how they affect the derived power spectrum. In §3 we present results for the angular power spectra obtained from individual pairs of radiometers, the cross-power spectra, and examine numerous consistency tests. In §4, in preparation for generating a final combined power spectrum, we present the full covariance matrix of the cross-power spectra. In §5 we present the methodology used to produce the final combined power spectrum and its covariance matrix. In §6 we compare the *WMAP* first-year power spectrum to a compilation of previous CMB measurements and to a prediction based on a combination of previous CMB data and the 2dFGRS data. We summarize our results in §7 and outline the power spectrum data products being made available through the Legacy Archive for Microwave Background Data Analysis (LAMBDA). Appendix A reviews two methods used to estimate the angular power spectrum from CMB maps. Appendix B describes how we account for point source contamination. Appendix C presents our approach to combining multi-channel data. Appendix D describes how the foreground mask correlates multipole moments in the Fisher matrix, and Appendix E collects some useful properties of the spherical harmonics.

2. INSTRUMENTAL PROPERTIES

The *WMAP* instrument is composed of 10 “differencing assemblies” (DAs) spanning 5 frequencies from 23 to 94 GHz (Bennett et al. 2003a). The 2 lowest frequency bands (K and Ka) are primarily Galactic foreground monitors, while the 3 highest (Q, V, and W) are primarily cosmological bands. There are 8 high frequency differencing assemblies: Q1, Q2, V1, V2, and W1 through W4. Each DA is formed from two differential radiometers which are sensitive to orthogonal linear polarization modes; the radiometers are designated 1 or 2 (e.g., V11 or W12) depending on which polarization mode is being sensed.

The temperature measured on the sky is modified by the properties of the instrument. The most important properties that affect the angular power spectrum are finite resolution and instrument noise. Let $C_l^{ii'}$ denote the auto or

cross-power spectrum evaluated from two sky maps, i and i' , where i is a DA index. Further, define the shorthand $\mathbf{i} \equiv (i, i')$ to denote a pair of indices, e.g., (Q1, V2). This spectrum will have the form

$$C_l^{\mathbf{i}} = w_l^{\mathbf{i}} C_l^{\text{sky}} + N_l^{\mathbf{i}}, \quad (6)$$

where $w_l^{\mathbf{i}} \equiv b_l^i b_l^{i'} p_l^2$ is the window function that describes the combined smoothing effects of the beam and the finite sky map pixel size. Here b_l^i is the beam transfer function for DA i , given by Page et al. (2003a) [note that they reserve the term “beam window function” for $(b_l^i)^2$], and p_l is the pixel transfer function supplied with the HEALPix package (Górski et al. 1998). $N_l^{\mathbf{i}}$ is the noise spectrum realized in this particular measurement. On average, the observed spectrum estimates the underlying power spectrum, C_l ,

$$\langle C_l^{\mathbf{i}} \rangle = w_l^{\mathbf{i}} C_l + \langle N_l^{\mathbf{i}} \rangle \delta_{ii'}, \quad (7)$$

where $\langle N_l^{\mathbf{i}} \rangle$ is the average noise power spectrum for differencing assembly i , and the Kronecker symbol indicates that the noise is uncorrelated between differencing assemblies. To estimate the underlying power spectrum on the sky, C_l , the effects of the noise bias and beam convolution must be removed. The determination of transfer functions and noise properties are thus critical components of any CMB experiment.

In §2.1 we summarize the results of Page et al. (2003a) on the *WMAP* window functions and their uncertainties. We propagate these uncertainties through to the final Fisher matrix for the angular power spectrum. In §2.2 we present a model of the *WMAP* noise properties appropriate to power spectrum evaluation. For cross power spectra ($i \neq i'$ above), the noise bias term drops out of equation (7) if the noise between the two DAs is uncorrelated. These cross-power spectra provide a nearly optimal estimate of the true power spectrum, essentially independent of errors in the noise model, thus we use them exclusively in our final power spectrum estimate. The noise model is primarily used to propagate noise errors through the analysis, and to test a variety of different power spectrum estimates for consistency with the combined cross-power spectrum.

2.1. Window Functions

As discussed in Page et al. (2003a), the instrument beam response was mapped in flight using observations of the planet Jupiter. The signal to noise ratio is such that the response, relative to the peak of the beam, is measured to approximately -35 dB in W band, the band with the highest angular resolution. The beam widths, measured in flight, range from $0^\circ 82$ at K band down to $0^\circ 20$ in some of the W band channels (FWHM). Maps of the full two-dimensional beam response are presented in Page et al. (2003a), and are available with the *WMAP* first-year data release. The radial beam profiles obtained from these maps have been fit to a model consisting of a sum of Hermite polynomials that accurately characterize the main Gaussian lobe and small deviations from it. The model profiles are then Legendre transformed to obtain the beam transfer functions b_l^i for each DA i . Full details of this procedure are presented in Page et al. (2003a), and the resulting transfer functions are also provided in the first-year data release. We have chosen to normalize the transfer function to 1 at $l = 1$ because *WMAP* calibrates its intensity response using the modulation of the CMB dipole ($l = 1$). This effectively partitions calibration uncertainty from window function uncertainties.

The beam processing described above provides a straightforward means of propagating the noise uncertainty directly from the time-ordered data through to the final transfer functions. The result is the covariance matrix $\Sigma_{b, ll'}^i$ for the normalized transfer function. Plots of the diagonal elements of $\Sigma_{b, ll'}^i$ are presented in Page et al. (2003a). The fractional error in the transfer functions b_l^i are typically 1-2% in amplitude. In the end, these window function uncertainties dominate the small off-diagonal elements of the final covariance matrix for the combined power spectrum (see §4). Additional observations of Jupiter will reduce these uncertainties.

An additional source of error in our treatment of the beam response arises from non-circularity of the main beam. The effects of this non-circularity are mitigated by *WMAP*'s scan strategy which results in most sky pixels being observed over a wide range of azimuth angles. The effective beam response on the sky is thus largely symmetrized. We estimate that the effects of imperfect symmetrization produce window function errors of $< 1\%$ relative to a perfectly symmetrized beam window function (Page et al. 2003a; Hinshaw et al. 2003b). This error is well within the formal uncertainty given in $\Sigma_{b,ll'}^i$. In §5 we infer the optimal power spectrum by combining the 28 cross-power spectra measured by the 8 high frequency DAs Q1 through W4. As part of this process we marginalize over the window function uncertainty, which automatically propagates these errors into the final covariance matrix for the combined power spectrum. Both the combined power spectrum and the corresponding Fisher matrix are part of the first-year data release.

2.2. Instrument Noise Properties

The noise bias term in equation (7) is the noise per a_{lm} mode on the sky. If auto-power spectra are used in the final power spectrum estimate, the noise bias term must be known very accurately because it exponentially dominates the convolved power spectrum at high l . If only cross correlations are used, the noise bias is only required for propagating errors. Our final best spectrum is based only on cross correlations, and is independent of this term. However, as an independent check of our results, we evaluate the maximum likelihood spectrum based on a combined Q+V+W sky map. The noise bias must be estimated accurately for this application.

In the limit that the time-ordered instrument noise is white, the noise bias is a constant, independent of l . If the time-ordered noise has a $1/f$ component, the bias term will rise at low l . In this subsection we estimate the noise bias properties for each of the high frequency *WMAP* radiometers based on the time-ordered noise properties presented in Jarosik et al. (2003b). While the *WMAP* radiometer noise is nearly white by design (Jarosik et al. 2003a) with $1/f$ knee frequencies of less than 10 mHz for 9 out of 10 differencing assemblies, one of the radiometers (W41) has a $1/f$ knee frequency of ~ 45 mHz. The latter is large enough that the deviations of $\langle N_l^i \rangle$ from a constant must be accounted for.

The most reliable way to estimate the effects of $1/f$ noise on the measured power spectra is by Monte Carlo simulation. Using the pipeline simulator discussed in Hinshaw et al. (2003b) we have generated a library of noise maps with flight-like properties. Specifically we have included flight-like $1/f$ noise in the simulated time-ordered data, and have run each full-year realization through the map-making pipeline, including the baseline pre-whitening discussed in Hinshaw et al. (2003b). We evaluate the power spectra of these maps using the quadratic estimator described in Appendix A with 3 different pixel weighting schemes. (See Appendix A.1.2 for definitions of the weights, and the l range in which each is used.) We define the effective noise as a function of l based on fits to these Monte Carlo noise spectra. For the analyses in this paper, we fit the spectra to a model of the form

$$\ln \langle N_l^i \rangle = \left[\sum_{n=0}^{n_{\max}} c_n^i (\ln l)^n \right]^{-1} \quad (8)$$

where the c_n^i are fit coefficients given in Table 1, with $n_{\max} = 3$ for $l < 200$, and $n_{\max} = 1$ for $l > 200$.

Figure 1 shows the noise spectrum derived from the simulations for each of the 8 high frequency DAs, using uniform weighting over the entire l range. For comparison, we also plot an estimate of the CMB power spectrum from §5 in grey. Note that the W4 spectrum is the only one of this set to exhibit deviations from white noise in an l range where the signal-to-noise is relatively low, and we believe this simulation slightly over-estimates the $1/f$ noise in the flight W4 differencing assembly (Hinshaw et al. 2003a).

2.3. Systematic Errors

Hinshaw et al. (2003a) present limits on systematic errors in the first-year sky maps. They consider the effects of absolute and relative calibration errors, artifacts induced by environmental disturbances (thermal and electrical), errors from the map-making process, pointing errors, and other miscellaneous effects. The combined errors due to relative calibration errors, environmental effects, and map-making errors are limited to $< 15 \mu\text{K}^2$ (2σ) in the quadrupole moment C_2 in any of the 8 high-frequency DAs. Tighter limits are placed on higher-order moments. We conservatively estimate the absolute calibration uncertainty in the first-year *WMAP* data to be 0.5%.

Random pointing errors are accounted for in the beam mapping procedure; the beam transfer functions presented by Page et al. (2003a) incorporate random pointing errors automatically. A systematic pointing error of $\sim 10''$ at the spin period is suspected in the quaternion solution that defines the spacecraft pointing. This is much smaller than the smallest beam width ($\sim 12'$ at W band), and we estimate that it would produce $< 1\%$ error in the angular power spectrum at $l = 1000$, thus we do not attempt to correct for this effect. Barnes et al. (2003) place limits on spurious contributions due to stray light pickup through the far sidelobes of the instrument. They place limits of $< 10 \mu\text{K}^2$ on spurious contributions to C_l , at Q through W band, due to far sidelobe pickup.

A detailed model of Galactic foreground emission based on the first-year *WMAP* data is presented by Bennett et al. (2003c) and is summarized in §3.1. We show that diffuse foreground emission is a modest source of contamination at large angular scales ($\gtrsim 2^\circ$). *Systematic errors on these angular scales are negligible compared to the (modest) level of foreground emission.* On smaller angular scale ($\lesssim 2^\circ$), the 1-3% uncertainty in the individual beam transfer functions is the largest source of uncertainty, while for multipole moments greater than ~ 600 , random white noise from the instrument is the largest source of uncertainty.

3. THE DATA

Figure 2 shows the cross-power spectra obtained from all 28 combinations of the 8 differencing assemblies Q1 through W4 using the quadratic estimator described in Appendix A.1. These spectra have been evaluated with the Kp2 sky cut described in Bennett et al. (2003c). The spectra are color coded by effective frequency, $\sqrt{\nu^i \nu^j}$, where ν^i is the frequency of differencing assembly i . The low frequency (41 GHz) data are shown in red, the high frequency (94 GHz) data in blue, with intermediate frequencies following the colors of the rainbow. The top panel shows $l(l+1)C_l/2\pi$ in μK^2 , while the bottom panel plots the ratio of each channel to our final combined spectrum presented in §5. The top panel shows a very robust measurement of the first acoustic peak with a maximum near $l \sim 220$ and a shape that is consistent with the predictions of adiabatic fluctuation models. There is also a clear indication of the rise to a second peak at $l \sim 540$. See Page et al. (2003b) for an analysis and discussion of the peaks and troughs in the first-year *WMAP* power spectrum.

The red data in the top panel show very clearly that the low frequency data are contaminated by diffuse Galactic emission at low l and by point sources at higher l . The higher frequency data show less apparent contamination, consistent with the foreground emission being dominated by radio emission, rather than thermal dust emission, as expected in this frequency range.

3.1. Galactic and Extragalactic Foregrounds

Bennett et al. (2003c) present a detailed model of the Galactic foreground emission based on a Maximum Entropy

analysis of all 5 *WMAP* frequency bands, in combination with external tracer templates. They demonstrate that the emission is well modeled by three distinct emission components. 1) Synchrotron emission from cosmic ray electrons, with a steeply falling spectrum in the *WMAP* frequency range: $T_A(\nu) \propto \nu^\beta$ with $\beta < -3$, steepening with increasing frequency. 2) Free-free emission from the ionized interstellar medium that is well traced by $H\alpha$ emission in regions where the dust extinction is low. 3) Thermal emission from interstellar dust grains with an emissivity index $\beta \sim 2.2$. The model has a Galactic signal minimum between V and W band.

In principal we could subtract the above model from each *WMAP* channel and recompute the power spectrum. However, since the model is based on *WMAP* data that have been smoothed to an angular resolution of 1° , the resulting maps would have complicated noise properties. For the purposes of power spectrum analysis, we adopt a more straightforward approach based on fitting foreground tracer templates to the Q, V, and W band data. The details of this procedure, the resulting fit coefficients, and a comparison of the fits to the Maximum Entropy model are given in Bennett et al. (2003c). They estimate that the template model removes $\sim 85\%$ of the foreground emission in Q, V, and W bands and that the remaining emission constitutes less than $\sim 2\%$ of the CMB variance (up to $l = 200$) in Q band, and $\lesssim 1\%$ of the CMB variance in V and W bands.

The contribution from extragalactic radio sources has been analyzed in three separate ways. Bennett et al. (2003c) directly fit for sources in the sky maps. The result of this analysis is that we have identified 208 sources in the *WMAP* data with sufficient signal to noise ratio to pass the detection criterion (we estimate that ~ 5 of these are likely to be spurious). The derived source count law is consistent with the following model for the power spectrum of the unresolved sources

$$C_l^{i,src} = A \left(\frac{\nu_i}{\nu_0} \right)^\beta \left(\frac{\nu_{i'}}{\nu_0} \right)^\beta w_l^i, \quad (9)$$

with $A = 0.015 \mu\text{K}^2 \text{ sr}$ (measured in thermodynamic temperature), $\beta = -2.0$, and $\nu_0 \equiv 45 \text{ GHz}$. Komatsu et al. (2003) evaluate the bispectrum of the *WMAP* data and are able to fit a non-Gaussian source component based on a particular configuration of the bispectrum. They find the same source model, equation (9), fits the bispectrum data. For the remainder of this section, we adopt this model for correcting the cross-power spectra. At Q band (41 GHz) the correction to $l(l+1)C_l/2\pi$ is 868 and $3468 \mu\text{K}^2$ at $l=500$ and 1000 , respectively. At W band (94 GHz), the correction is only 31 and $126 \mu\text{K}^2$ at the same l values. For comparison, the CMB power in this l range is $\sim 2000 \mu\text{K}^2$. Later, when we derive a final combined spectrum from the multi-frequency data, we adopt equation (9) as a model with A as a free parameter. We simultaneously fit for a combined CMB spectrum and source amplitude and marginalize over the residual uncertainty in A . The best-fit source amplitude from this process is consistent with the other two methods.

Figure 3 shows the cross-power spectra obtained from the same 28 combinations as in Figure 2, this time with the Galactic template model and source model subtracted. The bottom panel of the Figure shows the ratio of the 28 channels to the combined spectrum obtained in §5. The 28 cross-power spectra are consistent with each other at the 5 to 20% level over the l range 2 – 500. Similar scatter is seen in Monte Carlo simulations of an ensemble of 28 cross-power spectra with *WMAP*'s beam and noise properties. The only significant deviation lies in the Q band data at low l which is $\sim 10\%$ below the higher frequency bands at $l < 20$. This is consistent with the accuracy estimated above for the Galactic template model, see also Figure 11 of Bennett et al. (2003c) for images of the maps after Galactic template subtraction. Since the *WMAP* data are not noise limited at low l , we use only V and W band data in the final combined spectrum for $l < 100$.

The subtraction of the source model, equation (9), brings the Q band spectrum into good agreement with the other cross-power spectra up to $l \sim 500$. At higher l , the Q band data contributes very little to the final combined spectrum because the (normalized) Q band window function has dropped to less than 5% (Page et al. 2003a). As discussed in §5, we marginalize over the source amplitude uncertainty, δA , when obtaining the final power spectrum estimate and associated covariance matrix. Thus the uncertainty is also accounted for in subsequent cosmological parameter fits

(Verde et al. 2003; Spergel et al. 2003; Peiris et al. 2003).

Figure 4 shows a close-up of the 28 cross-power spectra in Figure 3 up to $l = 100$. The top panel shows the raw (unbinned) data which has correlations of $< 2\%$ between neighboring points in this l range (see §4). These data are strikingly consistent with each other and support the conclusion that systematic errors at low l are insignificant. To assess the level of scatter that does exist between the spectra, we have generated a Monte Carlo simulation in which we compute the rms scatter among the 28 spectra at each l , relative to the measured power. The bottom panel of Figure 4 shows the results of this simulation, averaged over 1000 realizations, compared to the relative rms scatter in the data. The agreement is excellent, indicating that the uncertainty in the measured power spectrum in this l range is a few percent and is consistent with a combination of instrument noise and mode coupling due to the 15% sky cut.

Another striking feature is the low amplitude of the observed quadrupole, and the sharp rise in power, almost linear in l , to $l = 5$. Bennett et al. (2003b) quote a value for the rms quadrupole amplitude, $Q_{rms} = \sqrt{(5/4\pi)C_2} = 8 \pm 2 \mu\text{K}$, where the uncertainty is largely due to Galactic model uncertainty. This is consistent with the amplitude measured by the *COBE-DMR* experiment, $Q_{rms} = 10_{-4}^{+7} \mu\text{K}$ (Bennett et al. 1996). The fast, nearly linear rise to $l = 5$ produces an angular correlation function with essentially no power on angular scales $\gtrsim 60^\circ$, again in excellent agreement with the *COBE-DMR* correlation function (Bennett et al. 2003b; Hinshaw et al. 1996). In the context of a standard ΛCDM model, the probability of observing this little power on scales greater than 60° is $\sim 2 \times 10^{-3}$ (Bennett et al. 2003b).

4. THE FULL COVARIANCE MATRIX

In §5 we derive our best estimate of the angular power spectrum by optimally combining the 28 cross-power spectra discussed above. The procedure for combining spectra requires the full covariance matrix of the individual cross-power spectra – in this section we outline the salient features of this matrix. There are six principal sources of variance for the measured spectra, C_l^i : cosmic variance, instrument noise, mode coupling due to the foreground mask, point source subtraction errors, uncertainty in the beam window functions, and an overall calibration uncertainty. We ignore uncertainties in the diffuse foreground correction since they are everywhere sub-dominant to the cosmic variance uncertainty (see §3.1).

We may write the covariance matrix as

$$(\Sigma_{\text{full}})_{ll'}^{ij} = \left\langle [C_l^i - (C_l + AS^i)w_l^i][C_{l'}^j - (C_{l'} + AS^j)w_{l'}^j] \right\rangle, \quad (10)$$

where the angle brackets represent an ensemble average, C_l is the true underlying power spectrum, w_l^i is the window function of spectrum \mathbf{i} , and AS^i is the point source contribution to spectrum \mathbf{i} . Here we have defined a point source spectral function, S^i , as

$$S^i \equiv \left(\frac{\nu_i}{\nu_0} \right)^\beta \left(\frac{\nu_{l'}}{\nu_0} \right)^\beta, \quad (11)$$

where ν_i , ν_0 , and β are as defined after equation (9). Note that $(\Sigma_{\text{full}})_{ll'}^{ij}$ is symmetric in both (ll') and (ij) .

In the process of forming the combined spectrum we will estimate a best-fit point source amplitude, \bar{A} , and subtract the corresponding source contribution from each spectrum \mathbf{i} . We thus rewrite Σ_{full} as

$$(\Sigma_{\text{full}})_{ll'}^{ij} = \left\langle [\bar{C}_l^i - (C_l + \delta AS^i)w_l^i][\bar{C}_{l'}^j - (C_{l'} + \delta AS^j)w_{l'}^j] \right\rangle, \quad (12)$$

where $\bar{C}_l^i \equiv C_l^i - \bar{A}S^i w_l^i$ is the source-subtracted spectrum, and $\delta A \equiv A - \bar{A}$ is the residual source amplitude, which we will marginalize over as a nuisance parameter.

We may expand the covariance matrix as

$$\Sigma_{\text{full}} = \Sigma_{\text{cv}} + \Sigma_{\text{mask}} + \Sigma_{\text{src}} + \Sigma_{\text{win}}. \quad (13)$$

We discuss each of these contributions in more detail below.

Cosmic Variance, Noise, and Mode Coupling.— The first two terms, $\Sigma_{\text{cv}} + \Sigma_{\text{mask}}$, incorporate the combined uncertainty due to cosmic variance, instrument noise, and mode coupling due to the foreground mask,

$$(\Sigma_{\text{cv}} + \Sigma_{\text{mask}})_{ll'}^{\text{ij}} = \left\langle [C_l^{\text{i}} - C_l w_l^{\text{i}}] [C_{l'}^{\text{j}} - C_{l'} w_{l'}^{\text{j}}] \right\rangle, \quad (14)$$

where w_l^{i} is fixed at its measured value. We have split this contribution into two pieces to mimic the procedure we actually use to compute the covariance matrix. As outlined in §5.1, we start with Σ_{cv} , then incorporate the effects of point source error and window function error. We do not add the effects of mode coupling, Σ_{mask} , until the very end of the computation. We consider this term in more detail in Appendix D.

Point Source Subtraction Errors.— The third term, Σ_{src} , is due to uncertainty in the point source amplitude determination

$$(\Sigma_{\text{src}})_{ll'}^{\text{ij}} = S^{\text{i}} w_l^{\text{i}} \sigma_{\text{src}}^2 w_{l'}^{\text{j}} S^{\text{j}}, \quad (15)$$

where $\sigma_{\text{src}}^2 \equiv \langle (A - \bar{A})^2 \rangle$ is the variance in the best-fit amplitude \bar{A} , and we assume that the frequency dependence, S^{i} , is perfectly known. In practice, we do not explicitly evaluate Σ_{src} as given above, rather we employ a method based on marginalizing a Gaussian likelihood function, $\mathcal{L}(C_l^{\text{i}}, A)$, over a nuisance parameter A . This process, which is discussed in Appendix B, effectively yields $\Sigma_{\text{cv}} + \Sigma_{\text{src}}$.

Window Function and Calibration Errors.— The fourth term, Σ_{win} , is due to uncertainty in the beam window function, w_l^{i} . This term arises from fluctuations in the window function which cause the measured spectrum, C_l^{i} to differ from our estimate of the convolved spectrum, $C_l w_l^{\text{i}}$, where w_l^{i} is the estimated window function. This contribution has the form

$$(\Sigma_{\text{win}})_{ll'}^{\text{ij}} = C_l \left\langle \Delta w_l^{\text{i}} \cdot \Delta w_{l'}^{\text{j}} \right\rangle C_{l'}. \quad (16)$$

Recall from §2 that $w_l^{\text{i}} = b_l^{\text{i}} b_l^{\text{i}'} p_l^2$ where b_l^{i} is the beam transfer function for DA i and p_l^2 is the pixel window function. Define $u_l^{\text{i}} \equiv \Delta b_l^{\text{i}} / b_l^{\text{i}}$ to be the fractional error in b_l^{i} , then to first order in u_l^{i} we have

$$\left\langle \Delta w_l^{\text{i}} \cdot \Delta w_{l'}^{\text{j}} \right\rangle = w_l^{\text{i}} \left\langle u_l^{\text{i}} u_{l'}^{\text{j}} + u_l^{\text{i}} u_{l'}^{\text{j}'} + u_l^{\text{i}'} u_{l'}^{\text{j}} + u_l^{\text{i}'} u_{l'}^{\text{j}'} \right\rangle w_{l'}^{\text{j}}. \quad (17)$$

For *WMAP* the uncertainty in b_l^{i} is uncorrelated between DAs, thus the above expression reduces to

$$\left\langle \Delta w_l^{\text{i}} \cdot \Delta w_{l'}^{\text{j}} \right\rangle = w_l^{\text{i}} \left[\Sigma_{\text{u},ll'}^{\text{i}} (\delta_{ij} + \delta_{i'j'}) + \Sigma_{\text{u},ll'}^{\text{i}'} (\delta_{i'j} + \delta_{ij'}) \right] w_{l'}^{\text{j}}, \quad (18)$$

where we have defined

$$\Sigma_{\text{u},ll'}^{\text{i}} \equiv \left\langle u_l^{\text{i}} u_{l'}^{\text{i}} \right\rangle = (b_l^{\text{i}} b_{l'}^{\text{i}})^{-1} \Sigma_{\text{b},ll'}^{\text{i}}, \quad (19)$$

and where $\Sigma_{\text{b},ll'}^{\text{i}}$ is the covariance matrix of the beam transfer function for DA i given by Page et al. (2003a). When generating the combined spectrum in the next section, we add Σ_{win} to the above contributions, giving $\Sigma_{\text{cv}} + \Sigma_{\text{src}} + \Sigma_{\text{win}}$.

The *WMAP* absolute calibration uncertainty is 0.5%. We do not explicitly incorporate this contribution in the covariance matrix. Instead, we propagate a 0.5% uncertainty into the normalization of the final power spectrum amplitude (Spergel et al. 2003).

5. THE COMBINED POWER SPECTRUM

In §3 we demonstrated that the three high frequency bands of *WMAP* data produced consistent estimates of the angular power spectrum, after a modest correction for diffuse Galactic emission and extragalactic point sources. It is therefore justifiable to combine these data into a single “optimal” estimate of the angular power spectrum of the CMB. In this section, we provide an overview of two methods we use to generate a single combined spectrum. The first is a multi-step process that simultaneously fits the 28 cross-power spectra presented above to a single CMB power spectrum and a point source model, equation (9), while correctly propagating beam and residual point source uncertainties through to a final Fisher matrix. This spectrum constitutes our best estimate of the CMB power spectrum from the first-year *WMAP* data. The second spectrum, which serves as a cross check of the first, is based on forming a single co-added sky map from the Q1 through W4 maps, and using the quadratic estimator with noise bias subtraction. We compare the two spectra in §5.3.

5.1. Method I – Optimal Combination of Cross-Power Spectra

Since this method is relatively complicated, we outline the basic procedure here and relegate the details to Appendices, as indicated. We present the result in §5.3. The steps are as follows.

1. Subtract best-fit Galactic foreground templates from each of the maps Q1 through W4, using the coefficients given in Table 3 of Bennett et al. (2003c).
2. Evaluate the 28 cross-power spectra from the maps Q1 through W4, where each spectrum has been evaluated using the quadratic estimator of Appendix A.1 with the weighting scheme defined in Appendix A.1.2.
3. Collect the noise bias estimate, $\langle N_l^i \rangle$, for each DA from §2.2. These estimates are used in the calculation of the covariance matrix for the combined spectrum, and in setting the relative weight of each cross-power spectrum in the final combined spectrum.
4. Apply the procedure presented in Appendix B.1 to obtain an estimate for the point source amplitude. The value obtained is $\bar{A} = 0.0155 \pm 0.0017$, roughly independent of l_{\max} in equation (B6). This value is in good agreement with an estimate based on the bispectrum (Komatsu et al. 2003), and on an extrapolation of point source counts (Bennett et al. 2003b). Subtract the point source contribution from the cross-power spectra: $\bar{C}_l^i = C_l^i - \bar{A} S_l^i w_l^i$.
5. Compute an approximate form of the full covariance matrix discussed in §4, $\tilde{\Sigma}_{\text{full}}$. The procedure we use produces a covariance matrix which includes cosmic variance, instrument noise, source subtraction uncertainties, and window function uncertainties. At this stage of the process, it does not yet include the effects of mode coupling. More details are given in §4 and Appendix B.
6. Invert the approximate covariance matrix for use in computing the optimal spectrum. This is the most computationally intensive step in the process.
7. Compute the final combined spectrum from the 28 C_l^i as per the procedure given in Appendix C. In particular, assume a fiducial cosmological model (as specified in the Appendix), and use equation (C6), with $\Sigma_{\text{full } l l'}^{\text{ij}} = \tilde{\Sigma}_{\text{full } l l'}^{\text{ij}}$. This produces a final spectrum which is very nearly optimal.

Note: for $l < 100$ we use a surrogate procedure for computing the combined spectrum. In order to minimize Galactic foreground contamination, we use only V and W band data. Moreover, because the statistics of the C_l are mildly non-Gaussian, and because point source contamination and window function uncertainties are small,

the “optimal” machinery developed above is unnecessary. Rather we simply form a weighted average spectrum from the V and W band C_l^i .

8. Compute the approximate inverse-covariance matrix, \tilde{Q} , for the combined spectrum using equation (C7). This matrix is approximate in two ways. 1) It does not yet incorporate the effects of mode coupling – this is added below. 2) It has been evaluated for a fiducial cosmological model, $\tilde{Q}(C_l^{\text{fid}})$, while in a likelihood application, we need to evaluate $\tilde{Q}(C_l^{\text{th}})$ for an arbitrary model, C_l^{th} – we add this next.
9. Introduce the dependence on cosmological model into \tilde{Q} as follows. Invert \tilde{Q} to obtain the approximate covariance matrix of the combined spectrum, $\tilde{\Sigma}$. The off-diagonal terms of $\tilde{\Sigma}$ are small and weakly dependent on cosmological model, so we expand $\tilde{\Sigma}$ as

$$\tilde{\Sigma}_{ll'} = D_l \delta_{ll'} + \epsilon_{ll'}, \quad (20)$$

where $\epsilon_{ll'}$ encodes the mode coupling due to window function and source subtraction uncertainties. This relation defines $\epsilon_{ll'}$, which we take to be zero on the diagonal. D_l is dominated by cosmic variance and noise; in order to separate the two contributions, we introduce an *ansatz* of the form

$$D_l \equiv \frac{2}{2l+1} \frac{1}{f_{\text{sky}}} (C_l^{\text{fid}} + N_l^{\text{eff}})^2, \quad (21)$$

where C_l^{fid} is the fiducial model used to generate the combined spectrum (see Appendix C), and N_l^{eff} is the effective noise in the combined spectrum, which is defined by this equation. The dependence on cosmological model is introduced in the covariance matrix by re-computing D_l with $C_l^{\text{fid}} \rightarrow C_l^{\text{th}}$, leaving N_l^{eff} fixed.

10. Estimate the coupling induced by the foreground mask as described in Appendix D. The effect of the mask on the off-diagonal elements of the Fisher (inverse-covariance) matrix can be written as

$$(\Sigma_{\text{mask}}^{-1})_{ll'} \equiv r_{ll'} (D_l' D_{l'}')^{-1/2}, \quad (22)$$

where $D_l' \equiv D_l f_{\text{sky}}^{-1}$. This expression parameterizes, and defines, the off-diagonal mode coupling in the form of a correlation matrix, $r_{ll'}$. Note that $\tilde{F}_{ll'}$, defined in Appendix D.2, is related to $r_{ll'}$ by $\tilde{F}_{ll'} \equiv \delta_{ll'} + r_{ll'}$.

11. The final Fisher (curvature) matrix, $Q_{ll'}$, is obtained using

$$Q_{ll'} = (D_l')^{-1} \delta_{ll'} - \epsilon_{ll'} (D_l' D_{l'}')^{-1/2} + r_{ll'} (D_l' D_{l'}')^{-1/2}, \quad (23)$$

We further calibrate D_l' with Monte Carlo simulations. This calibration process, and a description of how the curvature matrix is used in a maximum likelihood determination of cosmological parameters, is given in §2 of Verde et al. (2003). As part of the first-year data release, we provide a Fortran 90 subroutine to evaluate the likelihood of a given cosmological model, C_l^{th} , given the *WMAP* data (supplied in the routine). The code also optionally returns the Fisher (inverse-covariance) matrix for the combined spectrum.

5.2. Method II – Combined Sky Map

Our second approach is to form a single co-added map from the Q1 through W4 maps,

$$T = \frac{\sum_{i=3}^{10} T_i / \sigma_{0,i}^2}{\sum_{i=3}^{10} 1 / \sigma_{0,i}^2}, \quad (24)$$

where T_i' is the sky map for DA i with the best-fit Galactic template model subtracted, and $\sigma_{0,i}^2$ is the noise per observation for DA i , given by Bennett et al. (2003b). We evaluate the power spectrum of this map on the Kp2 cut sky using the quadratic estimator in Appendix A.1. An effective noise model is obtained by using the same approach as described in §2.2: we generate co-added noise maps from the library of end-to-end simulations, evaluate their average spectra, then fit a noise model. The noise bias model is then subtracted from the power spectrum of the combined temperature map. We have performed this analysis with 3 distinct pixel weighting schemes (see Appendix A.1.2) and 3 corresponding noise models. The results are shown in Figure 5 where it is seen that the three cases are virtually indistinguishable.

In effect, this analysis uses both the auto- and cross-power spectra. We view this as a useful check of the more sophisticated procedure described in §5.1, but we do not rely on it for a final result. Uncertainties in the noise model only effect the fourth moment of the cross-power spectra, but they effect the second moment of the auto-power spectra and potentially bias the final result. The $\sim 6\%$ sensitivity advantage gained by including auto-power spectra was not deemed worth the effort required to guarantee that the final result was not biased.

5.3. Comparison of Results

Figure 6 compares the power spectra obtained from methods I and II above. The combined cross-power spectrum from §5.1 is shown in black, the auto-power spectrum obtained in §5.2 from the co-added map is shown in grey. The two methods agree extremely well with the only notable deviation being at the highest l range probed by the first-year data. This is the regime where the auto-power spectrum will be most sensitive to the noise bias subtraction. As can be seen in the error estimates shown in Figure 8, the deviation is less than 1σ .

A separate test of robustness is to compute the angular power spectrum in separate regions of the sky to see if the spectrum changes. We have computed the power spectrum in two subsets of the sky – the ecliptic poles, and the ecliptic plane, using the quadratic estimator with the combined sky map. The results are shown in Figure 7 where the pole data is shown in grey and the plane data in black. The two spectra are very consistent overall, but some of the features that appear in the combined spectrum, such as the “peak” at $l \sim 40$ and the “bite” at $l \sim 210$, are not robust to this test, thus we consider these features to be of marginal significance. There is also no evidence that beam ellipticity, which would be more manifest in the plane than in the poles, systematically biases the spectrum. This is consistent with estimates of the effect given by Page et al. (2003a).

6. DISCUSSION

Our best estimate of the angular power spectrum of the CMB is shown in Figure 8. Also shown is the best-fit Λ CDM model from Spergel et al. (2003) which is based on a fit to the this spectrum plus a compilation of additional CMB and large-scale structure data. The *WMAP* data points are plotted with measurement errors based on the diagonal elements of the Fisher matrix presented in Appendix D. The cosmic variance errors, which include the effects of the sky cut, are plotted as a 1σ band around the best-fit model. As discussed in Spergel et al. (2003), the model is an excellent fit to the data. The combined spectrum provides a definitive measurement of the CMB power spectrum with uncertainties limited by cosmic variance up to $l \sim 350$. The spectrum clearly exhibits a first acoustic peak at $l = 220.1 \pm 0.8$ and a second acoustic peak at $l = 546 \pm 10$. Page et al. (2003b) present an analysis and interpretation of the peaks and troughs in the first-year *WMAP* power spectrum.

Figure 9 compares the first-year *WMAP* spectrum to a compilation of recent balloon and ground-based measure-

ments. In order to make this Figure meaningful, we plot the best-fit model spectrum to represent the *WMAP* results. The data points are plotted with errors that include both measurement uncertainty and cosmic variance, so no error band is included with the model curve. (Since individual groups report band power measurement with different bandwidths, it is not possible to represent a single cosmic variance band that applies to all data sets.) The model spectrum fit to *WMAP* agrees very well with the ensemble of previous observations.

Wang et al. (2002a) have recently distilled a CMB power spectrum from an optimal combination of the extent pre-*WMAP* data. In Figure 10 we plot their derived band power points alongside the *WMAP* data. To make this comparison meaningful, we plot the *WMAP* data with cosmic variance plus measurement errors and omit the error band from the model spectrum. The distilled spectrum is notably lower than the *WMAP* data in the vicinity of the first acoustic peak. In a previous version of this work (Wang et al. 2002b) the authors noted that the first peak of their combined spectrum was lower than a significant fraction of their input data. They attribute this to their formalism allowing for a renormalization of individual experiments within their respective calibration uncertainties. Figure 1 in Bennett et al. (2003a) presents a similarly distilled spectrum from the data extent in late 2001 and found a first peak amplitude that was more intuitively consistent with the bulk of the input data, and which is now seen to be consistent with the *WMAP* power spectrum.

Figure 11 shows the *WMAP* combined power spectrum compared to the locus of predicted spectra, in red, based on a joint analysis of pre-*WMAP* CMB data and 2dFGRS large-scale structure data (Percival et al. 2002). As in Figure 8, the *WMAP* data are plotted with measurement uncertainties, and the best-fit Λ CDM model (Spergel et al. 2003) is plotted with a 1σ cosmic variance error band. Percival et al. (2002) predict the location of the first peak should occur at $l = 221.8 \pm 2.4$, which is quite consistent with the value reported by Page et al. (2003b) of $l = 220.1 \pm 0.8$. The height of the first peak was predicted to be in the range $4920 \pm 170 \mu\text{K}^2$, while Page et al. (2003b) report a measured height of $5580 \pm 75 \mu\text{K}^2$, about 13% higher. Unlike the position, the amplitude of the first peak has a complicated dependence on cosmological parameters. Percival et al. (2002) report best-fit parameters for a Λ CDM model that are mostly consistent with those reported by Spergel et al. (2003) for the same class of models. The only mildly disparate comparison lies in the combination of normalization, σ_8 , and optical depth, τ . Percival et al. (2002) report the product $\sigma_8 e^{-\tau} = 0.72 \pm 0.03 \pm 0.02$, where the first error is a “theory” error and the second is measurement error. While Spergel et al. (2003) does not report a maximum likelihood range for this explicit parameter combination, the product of their maximum likelihood values for σ_8 and τ yields $\sigma_8 e^{-\tau} = 0.74$, which is consistent with Percival et al. (2002), but would make the first peak a few percent higher. Small differences in $\Omega_b h^2$, $\Omega_m h^2$, and n_s , may also contribute to the difference.

7. CONCLUSIONS

We present measurements of the angular power spectrum of the cosmic microwave background from the first-year *WMAP* data. The eight high-frequency sky maps from DAs Q1 through W4 were used to estimate 28 cross-power spectra, which are largely independent of the noise properties of the experiment. These data were tested for consistency in §3, then used in §5 as input to a final combined spectrum, discussed in §6. The procedure for estimating the uncertainties in the final combined spectrum were discussed in §4 and in numerous Appendices.

The combined spectrum provides a definitive measurement of the CMB power spectrum, with uncertainties limited by cosmic variance up to $l \sim 350$, and a signal to noise per mode > 1 up to $l \sim 650$. The spectrum clearly exhibits a first acoustic peak at $l = 220.1 \pm 0.8$ and a second acoustic peak at $l = 546 \pm 10$. Page et al. (2003b) present an analysis and interpretation of the peaks and troughs in the first-year *WMAP* power spectrum. Spergel et al. (2003), Verde et al. (2003), and Peiris et al. (2003) analyze the combined spectrum in the context of cosmological models.

They conclude that the data provide strong support adiabatic initial conditions, and they give precise measurements of a number of cosmological parameters. Kogut et al. (2003) analyze the correlation between *WMAP*'s temperature and polarization signals, the C_l^{TE} spectrum, and present evidence for a relatively high optical depth, and an early period of cosmic reionization. Among other things, this result implies that the temperature power spectrum is suppressed by $\sim 30\%$ on degree angular scales, due to secondary scattering.

A variety of first-year *WMAP* data products are being made available by NASA's new Legacy Archive for Microwave Background Data Analysis (LAMBDA). In addition to the sky maps and calibrated time-ordered data, we are providing the 28 cross power spectra used in this paper (with diffuse foregrounds subtracted), the combined spectrum from §5.1, and a Fortran 90 subroutine to compute the likelihood of a given cosmological model, (the code will also optionally return the Fisher (inverse-covariance) matrix for the combined spectrum.) The LAMBDA URL is <http://lambda.gsfc.nasa.gov/>.

The *WMAP* mission is made possible by the support of the Office of Space Sciences at NASA Headquarters and by the hard and capable work of scores of scientists, engineers, technicians, machinists, data analysts, budget analysts, managers, administrative staff, and reviewers. We thank Mike Nolta for helpful comments on an earlier draft of this manuscript. LV is supported by NASA through Chandra Fellowship PF2-30022 issued by the Chandra X-ray Observatory center, which is operated by the Smithsonian Astrophysical Observatory for an on behalf of NASA under contract NAS8-39073. We acknowledge use of the HEALPix package.

REFERENCES

- Barnes, C. et al. 2003, ApJ, submitted
- Bennett, C. L., Banday, A. J., Górski, K. M., Hinshaw, G., Jackson, P., Keegstra, P., Kogut, A., Smoot, G. F., Wilkinson, D. T., & Wright, E. L. 1996, ApJ, 464, L1
- Bennett, C. L., Bay, M., Halpern, M., Hinshaw, G., Jackson, C., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S., Page, L., Spergel, D. N., Tucker, G. S., Wilkinson, D. T., Wollack, E., & Wright, E. L. 2003a, ApJ, 583, 1
- Bennett, C. L., Halpern, M., Hinshaw, G., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S., Page, L., Spergel, D. N., Tucker, G. S., Wollack, E., Wright, E. L., Barnes, C., Greason, M., Hill, R., Komatsu, E., Nolta, M., Odegard, N., Peiris, H., Verde, L., & Weiland, J. 2003b, ApJ, submitted
- Bennett, C. L. et al. 2003c, ApJ, submitted
- Górski, K. M., Hivon, E., & Wandelt, B. D. 1998, in *Evolution of Large-Scale Structure: From Recombination to Garching*
- Gupta, S. & Heavens, A. F. 2002, MNRAS, 334, 167
- Hauser, M. G. & Peebles, P. J. E. 1973, ApJ, 185, 757
- Hinshaw, G., Banday, A. J., Bennett, C. L., Górski, K. M., Kogut, A., Lineweaver, C. H., Smoot, G. F., & Wright, E. L. 1996, ApJ, 464, L25
- Hinshaw, G. F. et al. 2003a, ApJ, submitted
- . 2003b, ApJ, submitted

- Hivon, E., Górski, K. M., Netterfield, C. B., Crill, B. P., Prunet, S., & Hansen, F. 2002, *ApJ*, 567, 2
- Jarosik, N. et al. 2003a, *ApJS*, 145
- . 2003b, *ApJ*, submitted
- Kogut, A. et al. 2003, *ApJ*, submitted
- Komatsu, E. et al. 2003, *ApJ*, submitted
- Limon, M., Wollack, E., Bennett, C. L., Halpern, M., Hinshaw, G., Jarosik, N., Kogut, A., Meyer, S. S., Page, L., Spergel, D. N., Tucker, G. S., Wright, E. L., Barnes, C., Greason, M., Hill, R., Komatsu, E., Nolte, M., Odegard, N., Peiris, H., Verde, L., & Weiland, J. 2003, Wilkinson Microwave Anisotropy Probe (*WMAP*): Explanatory Supplement, http://lambda.gsfc.nasa.gov/data/map/doc/WMAP_supplement.pdf
- Muciaccia, P. F., Natoli, P., & Vittorio, N. 1997, *ApJ*, 488, L63+
- Oh, S. P., Spergel, D. N., & Hinshaw, G. 1999, *ApJ*, 510, 551
- Page, L. et al. 2003a, *ApJ*, submitted
- . 2003b, *ApJ*, submitted
- . 2003c, *ApJ*, 585, in press
- Peiris, H. et al. 2003, *ApJ*, submitted
- Percival, W. J., Sutherland, W., Peacock, J. A., Baugh, C. M., Bland-Hawthorn, J., Bridges, T., Cannon, R., Cole, S., Colless, M., Collins, C., Couch, W., Dalton, G., De Propris, R., Driver, S. P., Efstathiou, G., Ellis, R. S., Frenk, C. S., Glazebrook, K., Jackson, C., Lahav, O., Lewis, I., Lumsden, S., Maddox, S., Moody, S., Norberg, P., Peterson, B. A., & Taylor, K. 2002, *MNRAS*, 337, 1068
- Spergel, D. N. et al. 2003, *ApJ*, submitted
- Tegmark, M., Taylor, A. N., & Heavens, A. F. 1997, *ApJ*, 480, 22
- Verde, L. et al. 2003, *ApJ*, submitted
- Wang, X., Tegmark, M., Jain, B., & Zaldarriaga, M. 2002a, *Phys. Rev. D*, submitted (astro-ph/0212417)
- Wang, X., Tegmark, M., & Zaldarriaga, M. 2002b, *Phys. Rev. D*, 65, 123001

A. POWER SPECTRUM ESTIMATION METHODS

For the analysis of *WMAP*'s first-year data, we have chosen two distinct methods for inferring the power spectrum. The first is a fast and accurate quadratic method for estimating the power spectrum of a partial sky map (Hivon et al. 2002). We summarize the basic approach here, highlighting the aspects of the method that are especially pertinent to *WMAP*, and refer the reader to Hivon et al. (2002) for details. The second is a maximum likelihood method that provides an independent estimate of the spectrum measured by *WMAP* (Oh et al. 1999).

We have applied both of these methods to the *WMAP* data. The results are shown in Figure 12, which shows spectra estimated from the V band map, up to $l = 200$, for the two methods. The maximum likelihood estimate has slightly lower uncertainties at low l , because the method optimally weights the data with a pixel-pixel covariance $\mathbf{C} = \mathbf{S} + \mathbf{N} \approx \mathbf{S}$ where \mathbf{S} is the covariance of the CMB signal and \mathbf{N} is the covariance of the noise (see Appendix A.2). Our quadratic estimator uses uniform pixel weights at low l (see Appendix A.1.2) though it is clear from the Figure that the difference is not significant. At high l , where the *WMAP* data are noise dominated, the two estimators give essentially identical results because they effectively weight the data in the same way.

To obtain our “best” estimate of the *WMAP* power spectrum, we adopt the quadratic estimator because it can be readily applied to pairs of *WMAP* radiometers in a way that is nearly independent of the properties of the instrument noise. In §5 we discuss our methodology for combining spectra from pairs of radiometers and present the final combined spectrum.

A.1. Quadratic Estimation

Hivon et al. (2002) start with the full-sky estimator in equation (2), add a position dependent weight, $W(\mathbf{n})$, and set W to zero in the regions where the sky is contaminated. In other regions, W is chosen to optimize the sensitivity of the estimator. A temperature map $\Delta T(\mathbf{n})$ on which a weight $W(\mathbf{n})$ is applied can be decomposed in spherical harmonic coefficients as

$$\tilde{a}_{lm} = \int d\Omega_{\mathbf{n}} \Delta T(\mathbf{n}) W(\mathbf{n}) Y_{lm}^*(\mathbf{n}) \quad (\text{A1})$$

$$\approx \Omega_p \sum_p \Delta T(p) W(p) Y_{lm}^*(p), \quad (\text{A2})$$

where the integral over the sky is approximated by a discrete sum over map pixels, each of which subtends solid angle Ω_p . Hivon et al. (2002) then define the “pseudo power spectrum” \tilde{C}_l as

$$\tilde{C}_l = \frac{1}{2l+1} \sum_{m=-l}^l |\tilde{a}_{lm}|^2. \quad (\text{A3})$$

The pseudo power spectrum \tilde{C}_l , given by the weighted spherical harmonic transform of a map, is clearly different from the full sky angular spectrum, C_l^{sky} , but the *ensemble averages* of the two spectra can be related by

$$\langle \tilde{C}_l \rangle = \sum_{l'} G_{ll'} \langle C_{l'}^{\text{sky}} \rangle \quad (\text{A4})$$

where $G_{ll'}$ describes the mode coupling resulting from the weight function $W(\mathbf{n})$ (Hauser & Peebles 1973). Hivon et al. (2002) give the following expression for the coupling matrix, which depends only on the geometry of the weight function $W(\mathbf{n})$

$$G_{l_1 l_2} = \frac{2l_2 + 1}{4\pi} \sum_{l_3} (2l_3 + 1) \mathcal{W}_{l_3} \begin{pmatrix} l_1 & l_2 & l_3 \\ 0 & 0 & 0 \end{pmatrix}^2, \quad (\text{A5})$$

where the final term in parentheses is the Wigner 3- j symbol, and \mathcal{W}_l is the angular power spectrum of the weight function

$$\mathcal{W}_l = \frac{1}{2l+1} \sum_m |w_{lm}|^2, \quad (\text{A6})$$

where

$$w_{lm} = \int d\Omega_{\mathbf{n}} W(\mathbf{n}) Y_{lm}^*(\mathbf{n}). \quad (\text{A7})$$

Upon inverting the coupling matrix $G_{ll'}$ and making the identification $\langle C_l^{\text{sky}} \rangle = C_l$, we obtain the following estimator of the power spectrum

$$C_l^{\text{obs}} = \sum_{l'} G_{ll'}^{-1} \tilde{C}_{l'}. \quad (\text{A8})$$

The computation of equation (A2) for each (l, m) up to $l = l_{\text{max}}$ would scale as $N_{\text{pix}} l_{\text{max}}^2$ if performed on an arbitrary pixelization of the sphere, where N_{pix} is the number of sky map pixels. However, for a pixelization scheme with iso-latitude pixel centers, fast FFT methods may be employed to speed up the evaluation, so it scales like $N_{\text{pix}}^{1/2} l_{\text{max}}^2$ (Muciaccia et al. 1997). The *WMAP* sky maps have been produced using the HEALPix layout (Górski et al. 1998) which supports such fast spherical harmonic transforms. In particular, the HEALPix routine `map2alm` evaluates equation (A2).

A.1.1. Auto- and Cross-Power Spectra from the WMAP Data

For a multi-channel experiment like *WMAP* it is quite powerful to evaluate the power spectra from different maps and compare results. In particular, the quadratic estimator described above may be used on 1 or 2 maps at a time by generalizing the expression for the pseudo power spectrum equation (A3) as

$$\tilde{C}_l = \frac{1}{2l+1} \sum_{m=-l}^l \tilde{a}_{lm}^i \tilde{a}_{lm}^{j*} \quad (\text{A9})$$

where \tilde{a}_{lm}^i refers to the transform of map i and \tilde{a}_{lm}^{j*} refers to map j , which needn't be the same as map i . As discussed in §2, if $i \neq j$ and the noise in the two maps is uncorrelated, the estimator equation (A9) provides an unbiased estimate of the underlying power spectrum.

We have tested the auto- and cross-power estimator extensively with Monte Carlo simulations of the first-year *WMAP* data. Selected results from this testing are shown in Figure 13. The auto- and cross-power spectra that obtain from the flight data are presented in detail in §3. The cross-power spectra form the basis for our final combined spectrum, presented in §5.

A.1.2. Choice of Weighting

We seek a weighting scheme that mimics the maximum likelihood estimation (Appendix A.2), which effectively weights the data by the full inverse covariance matrix $\mathbf{C}^{-1} = (\mathbf{S} + \mathbf{N})^{-1}$. For the combined spectrum presented in §5, we use three distinct weighting functions in three separate l ranges.

1. For $l < 200$ we give equal weight to all un-cut pixels,

$$W(p) = M(p) \quad (\text{A10})$$

where $M(p)$ is the Kp2 sky mask, defined by Bennett et al. (2003c). It takes values of 0 within the mask and 1 otherwise.

2. For $l > 450$ we use inverse-noise weighting,

$$W(p) = M(p)N_{\text{obs}}(p) \quad (\text{A11})$$

where $N_{\text{obs}}(p)$ is the number of observations of pixel p .

3. For $200 < l < 450$ we use a transitional weighting,

$$W(p) = \frac{M(p)}{1/\langle N_{\text{obs}} \rangle + 1/N_{\text{obs}}(p)} \quad (\text{A12})$$

where $\langle N_{\text{obs}} \rangle$ is the mean of N_{obs} evaluated over the cut sky.

Verde et al. (2003) discuss the choice of weighting in more detail.

A.2. Maximum Likelihood Estimation

This Appendix provides a summary of the maximum likelihood approach to power spectrum estimation originally presented by Oh et al. (1999). If the temperature fluctuations are Gaussian, and the *a priori* probability of a given set of cosmological parameters is uniform, then the power spectrum may be estimated by maximizing the multi-variate Gaussian likelihood function

$$\mathcal{L}(C_l | \mathbf{m}) = \frac{\exp(-\frac{1}{2} \mathbf{m}^T \mathbf{C}^{-1} \mathbf{m})}{\sqrt{\det \mathbf{C}}} \quad (\text{A13})$$

where \mathbf{m} is a data vector (see below) and \mathbf{C} is the covariance matrix of the data, which has contributions from both the signal and the instrument noise, $\mathbf{C} = \mathbf{S} + \mathbf{N}$. We can work in whatever basis is most convenient. In the pixel basis the data are the sky map pixel temperatures, and in the spherical harmonic basis the data are the a_{lm} coefficients of the map. In the former basis the noise covariance is nearly diagonal, while in the latter, the signal covariance is.

Pixel basis:	Spherical harmonic basis:
$\mathbf{m} \rightarrow T_i$	$\mathbf{m} \rightarrow a_{lm}$
$\mathbf{S} \rightarrow \sum_l \frac{(2l+1)}{4\pi} C_l P_l(\cos \theta_{ij})$	$\mathbf{S} \rightarrow \text{diag}(C_2, C_2, \dots, C_3, C_3, \dots)$
$\mathbf{N} \rightarrow \sigma_i^2 \delta_{ij}$	$\mathbf{N} \rightarrow N_{(lm)(lm)'} \text{ (see below).}$

For *WMAP*, the length of the data vector, N_{data} , is 2,672,361, the number of $7'$ sky map pixels (HEALPix $N_{\text{side}} = 512$) that survive the Galaxy cut, so it is necessary to find methods for evaluating \mathcal{L} that do not require a full inversion of the covariance matrix \mathbf{C} , which requires $O(N_{\text{data}}^3)$ operations. We use an iterative method for evaluating the likelihood that exploits the ability to find an approximate inverse $\tilde{\mathbf{C}}^{-1}$. The most important features in the data that make this possible are that *WMAP* observes the full sky and the Galaxy cut is predominantly azimuthally symmetric in Galactic coordinates (Bennett et al. 2003b). Of secondary importance for this pre-conditioner is that *WMAP*'s noise per pixel does not vary strongly across the sky (Bennett et al. 2003a). We discuss the pre-conditioner in more detail below.

Defining $f \equiv -2 \ln \mathcal{L}$ and $\mathbf{P}_l \equiv \frac{\partial \mathbf{C}}{\partial C_l}$, we maximize the likelihood by solving

$$\frac{\partial f}{\partial C_l} = 0 = \mathbf{m}^T \mathbf{C}^{-1} \mathbf{P}_l \mathbf{C}^{-1} \mathbf{m} + \text{tr}(\mathbf{C}^{-1} \mathbf{P}_l) \quad (\text{A14})$$

using a Newton-Raphson root finding method that generates an iterative estimate of the angular power spectrum

$$C_l^{(n+1)} = C_l^{(n)} - \frac{1}{2} \sum_{l'} F_{ll'} \frac{\partial f}{\partial C_{l'}} \quad (\text{A15})$$

at each step. Here $F_{ll'}$ is the Fisher matrix

$$F_{ll'} = -\langle \partial^2 \ln \mathcal{L} / \partial C_l \partial C_{l'} \rangle = \frac{1}{2} \text{tr}(\mathbf{C}^{-1} \mathbf{P}_l \mathbf{C}^{-1} \mathbf{P}_{l'}). \quad (\text{A16})$$

To implement the solution in equation (A15) we need a fast way to evaluate the following three components of $\sum_{l'} F_{ll'} \frac{\partial f}{\partial C_{l'}}$:

1. $\mathbf{m}^T \mathbf{C}^{-1} \mathbf{P}_l \mathbf{C}^{-1} \mathbf{m}$
2. $\text{tr}(\mathbf{C}^{-1} \mathbf{P}_l)$
3. $\text{tr}(\mathbf{C}^{-1} \mathbf{P}_l \mathbf{C}^{-1} \mathbf{P}_{l'})$.

We use the spherical harmonic basis in which the data vector consists of the a_{lm} coefficients of the map obtained by least squares fitting on the cut sky. The signal covariance is diagonal in this basis, while the noise matrix is obtained from the normal equations for the a_{lm} fit

$$\sum_{(lm)'} N_{(lm)(lm)'}^{-1} a_{(lm)'} = y_{(lm)} \quad (\text{A17})$$

where

$$N_{(lm)(lm)'}^{-1} \equiv \sum_i \frac{Y_{(lm)}(\hat{n}_i) Y_{(lm)'}(\hat{n}_i)}{\sigma_i^2} \quad (\text{A18})$$

$$y_{(lm)} \equiv \sum_i \frac{T_i Y_{(lm)}(\hat{n}_i)}{\sigma_i^2}. \quad (\text{A19})$$

The sums are over all sky map pixels that survive the Galaxy cut, and we have assumed that the noise is uncorrelated from pixel to pixel.

A.2.1. Evaluation of $\mathbf{C}^{-1} \mathbf{m}$

The term $\mathbf{C}^{-1} \mathbf{m}$ appears repeatedly in the evaluation of equation (A15). We compute this by solving $\mathbf{C} \mathbf{z} = (\mathbf{S} + \mathbf{N}) \mathbf{z} = \mathbf{m}$ for \mathbf{z} . A more numerically tractable system is obtained by multiplying both sides by $\mathbf{S}^{\frac{1}{2}} \mathbf{N}^{-1}$, so

$$(\mathbf{I} + \mathbf{S}^{\frac{1}{2}} \mathbf{N}^{-1} \mathbf{S}^{\frac{1}{2}}) \mathbf{S}^{\frac{1}{2}} \mathbf{z} = \mathbf{S}^{\frac{1}{2}} \mathbf{N}^{-1} \mathbf{m} = \mathbf{S}^{\frac{1}{2}} \mathbf{y} \quad (\text{A20})$$

where \mathbf{y} is the spherical harmonic transform of the map, defined in equation (A19). Note that \mathbf{y} can be quickly computed in any pixelization scheme that has iso-latitude pixel centers with fixed longitude spacing. We then solve equation (A20) using an iterative conjugate gradient method with a pre-conditioner for the matrix $\mathbf{A} \equiv (\mathbf{I} + \mathbf{S}^{\frac{1}{2}} \mathbf{N}^{-1} \mathbf{S}^{\frac{1}{2}})$. We find the following block-diagonal form of \mathbf{A} to be a good starting point

$$\tilde{\mathbf{A}} = \begin{pmatrix} \mathbf{I} + \mathbf{S}^{\frac{1}{2}} \tilde{\mathbf{N}}^{-1} \mathbf{S}^{\frac{1}{2}} & 0 \\ 0 & \text{diag}(\mathbf{I} + \mathbf{S}^{\frac{1}{2}} \tilde{\mathbf{N}}^{-1} \mathbf{S}^{\frac{1}{2}}) \end{pmatrix} \quad (\text{A21})$$

where $\tilde{\mathbf{N}}^{-1}$ is a block-diagonal approximation of the noise matrix discussed below. The lower-right block of $\tilde{\mathbf{A}}$ occupies the high l portion of the matrix where the signal to noise ratio $\mathbf{S}^{\frac{1}{2}} \mathbf{N}^{-1} \mathbf{S}^{\frac{1}{2}}$ is low, so a diagonal approximation is adequate. The upper-left block occupies the low l portion of the matrix where the signal dominates the noise, so we need a better estimate of \mathbf{N}^{-1} . In practice we find this split works well at $l = 512$ for the WMAP noise levels. As for

the approximate form of \mathbf{N}^{-1} , defined in equation (A19), note that the dominant off-diagonal contributions arise from the azimuthally symmetric Galaxy cut, which couples different l modes, but not m modes. Thus \mathbf{N}^{-1} is approximately block diagonal, with perturbations induced by the non-uniform sky coverage of *WMAP*. We therefore use a block diagonal form of \mathbf{N}^{-1} as the pre-conditioner,

$$\tilde{N}_{(lm)(lm')}^{-1} = N_{(lm)(lm')}^{-1} \delta_{mm'}. \quad (\text{A22})$$

Using the pre-conditioner equation (A21) we find that the conjugate gradient solution of equation (A20) converges in approximately six iterations and requires only cpu-minutes of processing on an SGI Origin 2000.

B. POINT SOURCE SUBTRACTION

In this Appendix we describe the procedure we use to estimate and subtract the point source contribution directly from the multi-frequency cross-power spectra. We then show how we incorporate the source model uncertainty into the covariance matrix of the source-subtracted spectra by marginalizing a Gaussian likelihood function over the source model amplitude parameter.

This marks the starting point of the multi-frequency analysis which will lead to the combined power spectrum, discussed in §5.1. In order to generate the combined spectrum, we need the full covariance matrix of the cross-power spectra (see §4). Our approach to generating the full covariance is to start with the ideal, full-sky form, which only includes cosmic variance and instrument noise, then we incorporate additional effects step by step, as outlined in §5.1 and in these Appendices. For an ideal experiment with full sky coverage, no point source contamination, and no beam uncertainty the covariance matrix is

$$\Sigma_{ll'}^{ii'jj'} = \frac{1}{(2l+1)} \left[\left(C_l^{\text{th}} w_l^{ij} + n^i n^j \delta_{ij} \right) \left(C_l^{\text{th}} w_l^{i'j'} + n^{i'} n^{j'} \delta_{i'j'} \right) + \left(C_l^{\text{th}} w_l^{ij'} + n^i n^{j'} \delta_{ij'} \right) \left(C_l^{\text{th}} w_l^{i'j} + n^{i'} n^j \delta_{i'j} \right) \right] \quad (\text{B1})$$

where δ_{ij} denotes the Kronecker symbol, $w_l^{ij} \equiv b_l^i b_l^j p_l^2$ is the window function, and $n^i n^{i'} \delta_{ii'} = N^i$.

B.1. Estimating the Point Source Amplitude

We start by assuming a Gaussian likelihood for the sky model, given the *WMAP* data

$$-2 \ln \mathcal{L}(A, C_l | C_l^{\text{i}}) = \sum_{\mathbf{ij} l} [C_l^{\text{i}} - (C_l + AS^{\text{i}}) w_l^{\text{i}}] (\Sigma^{-1})_l^{\mathbf{ij}} [C_l^{\text{j}} - (C_l + AS^{\text{j}}) w_l^{\text{j}}], \quad (\text{B2})$$

where C_l^{i} is cross-power spectrum \mathbf{i} , w_l^{i} is the window function for spectrum \mathbf{i} , C_l is the true CMB power spectrum, and AS^{i} is the source model defined in equations (9) and (11). Here we assume the diagonal form of the covariance matrix in equation (B1).

To determine the best-fit source amplitude, A , we marginalize this likelihood over the CMB spectrum, C_l . First we expand equation (B2) as

$$\begin{aligned} -2 \ln \mathcal{L} &= \sum_l \sum_{\mathbf{ij}} (C_l^{\text{i}} - AS^{\text{i}} w_l^{\text{i}}) (\Sigma^{-1})_l^{\mathbf{ij}} (C_l^{\text{j}} - AS^{\text{j}} w_l^{\text{j}}) \\ &- \sum_l 2C_l \sum_{\mathbf{ij}} (C_l^{\text{i}} - AS^{\text{i}} w_l^{\text{i}}) (\Sigma^{-1})_l^{\mathbf{ij}} w_l^{\text{j}} \\ &+ \sum_l C_l^2 \sum_{\mathbf{ij}} w_l^{\text{i}} (\Sigma^{-1})_l^{\mathbf{ij}} w_l^{\text{j}}, \end{aligned} \quad (\text{B3})$$

which is a quadratic form in C_l , $\mathcal{L} \propto \prod_l \exp[-\frac{1}{2}(aC_l^2 - 2bC_l + c)]$, with

$$\begin{aligned} a &\equiv \sum_{\mathbf{ij}} w_l^{\mathbf{i}} (\Sigma^{-1})_{\mathbf{l}}^{\mathbf{ij}} w_l^{\mathbf{j}}, \\ b &\equiv \sum_{\mathbf{ij}} (C_l^{\mathbf{i}} - AS^{\mathbf{i}} w_l^{\mathbf{i}}) (\Sigma^{-1})_{\mathbf{l}}^{\mathbf{ij}} w_l^{\mathbf{j}}, \\ c &\equiv \sum_{\mathbf{ij}} (C_l^{\mathbf{i}} - AS^{\mathbf{i}} w_l^{\mathbf{i}}) (\Sigma^{-1})_{\mathbf{l}}^{\mathbf{ij}} (C_l^{\mathbf{j}} - AS^{\mathbf{j}} w_l^{\mathbf{j}}). \end{aligned} \quad (\text{B4})$$

We wish to evaluate the marginalized likelihood, $\mathcal{L}_{C_l}(A) = \int dC_l \mathcal{L}(A, C_l)$. This is readily evaluated using the substitution $C_l \rightarrow C_l - b/a$, giving $\mathcal{L}_{C_l} \propto \prod_l \exp[-\frac{1}{2}(c - b^2/a)]$, where we drop a multiplicative term proportional to a , which is independent of A . The marginalized likelihood function is thus

$$\begin{aligned} -2 \ln \mathcal{L}_{C_l} &= \sum_l \sum_{\mathbf{ij}} (C_l^{\mathbf{i}} - AS^{\mathbf{i}} w_l^{\mathbf{i}}) (\Sigma^{-1})_{\mathbf{l}}^{\mathbf{ij}} (C_l^{\mathbf{j}} - AS^{\mathbf{j}} w_l^{\mathbf{j}}) \\ &- \sum_l \frac{1}{\sum_{\mathbf{i}'\mathbf{j}'} w_l^{\mathbf{i}'} (\Sigma^{-1})_{\mathbf{l}}^{\mathbf{i}'\mathbf{j}'} w_l^{\mathbf{j}'}} \left[\sum_{\mathbf{ij}} (C_l^{\mathbf{i}} - AS^{\mathbf{i}} w_l^{\mathbf{i}}) (\Sigma^{-1})_{\mathbf{l}}^{\mathbf{ij}} w_l^{\mathbf{j}} \right]^2. \end{aligned} \quad (\text{B5})$$

Setting $\partial \mathcal{L}_{C_l}(A) / \partial A = 0$ gives the most likely point sources amplitude

$$\bar{A} = \frac{\sum_{\mathbf{ij}l} C_l^{\mathbf{j}} (\Sigma^{-1})_{\mathbf{l}}^{\mathbf{ij}} H_l^{\mathbf{i}} w_l^{\mathbf{i}}}{\sum_{\mathbf{ij}l} S^{\mathbf{j}} w_l^{\mathbf{j}} (\Sigma^{-1})_{\mathbf{l}}^{\mathbf{ij}} H_l^{\mathbf{i}} w_l^{\mathbf{i}}}, \quad (\text{B6})$$

where

$$H_l^{\mathbf{i}} = S^{\mathbf{i}} - \frac{\sum_{\mathbf{i}'\mathbf{j}'} S^{\mathbf{i}'} (\Sigma^{-1})_{\mathbf{l}}^{\mathbf{i}'\mathbf{j}'} w_l^{\mathbf{j}'}}{\sum_{\mathbf{i}'\mathbf{j}'} w_l^{\mathbf{i}'} (\Sigma^{-1})_{\mathbf{l}}^{\mathbf{i}'\mathbf{j}'} w_l^{\mathbf{j}'}}. \quad (\text{B7})$$

The standard error on the best fit value of A is

$$\sigma_{\text{src}}^2 = \left[\sum_{\mathbf{ij}l} H_l^{\mathbf{i}} w_l^{\mathbf{i}} (\Sigma^{-1})_{\mathbf{l}}^{\mathbf{ij}} S^{\mathbf{j}} w_l^{\mathbf{j}} \right]^{-1}. \quad (\text{B8})$$

For *WMAP*, the off-diagonal terms in the covariance matrix are small. Here, neglecting them changes the inferred point source amplitude by less than 1%.

B.2. Marginalizing over Point Source Amplitude

The source subtraction procedure discussed above is uncertain. In this section we incorporate this uncertainty into the full covariance matrix for the cross-power spectra. We again assume a Gaussian likelihood function of the form

$$-2 \ln \mathcal{L} = \sum_{\mathbf{ij}l} [\bar{C}_l^{\mathbf{i}} - (C_l^{\text{th}} + \delta A S_l^{\mathbf{i}}) w_l^{\mathbf{i}}] (\Sigma^{-1})_{\mathbf{l}}^{\mathbf{ij}} [\bar{C}_l^{\mathbf{j}} - (C_l^{\text{th}} + \delta A S_l^{\mathbf{j}}) w_l^{\mathbf{j}}] \quad (\text{B9})$$

where $\bar{C}_l^{\mathbf{i}}$ is the source-subtracted cross-power spectrum for DA pair \mathbf{i} , obtained above, $w_l^{\mathbf{i}}$ is the window function for spectrum \mathbf{i} , C_l^{th} is now a fixed CMB model spectrum, and δA is the residual source amplitude.

We can marginalize the likelihood function over the residual point source amplitude as follows. Expand equation (B9) as

$$\begin{aligned}
-2 \ln \mathcal{L} &= \sum_{\mathbf{ij} \mathbf{l} \mathbf{l}'} (\bar{C}_l^{\mathbf{i}} - C_l^{\text{th}} w_l^{\mathbf{i}}) (\Sigma^{-1})_{\mathbf{l} \mathbf{l}'}^{\mathbf{ij}} (\bar{C}_{l'}^{\mathbf{j}} - C_{l'}^{\text{th}} w_{l'}^{\mathbf{j}}) \\
&- 2(\delta A) \sum_{\mathbf{ij} \mathbf{l} \mathbf{l}'} (\bar{C}_l^{\mathbf{i}} - C_l^{\text{th}} w_l^{\mathbf{i}}) (\Sigma^{-1})_{\mathbf{l} \mathbf{l}'}^{\mathbf{ij}} S^{\mathbf{i}} w_{l'}^{\mathbf{j}} \\
&+ (\delta A)^2 \sum_{\mathbf{ij} \mathbf{l} \mathbf{l}'} S^{\mathbf{i}} w_l^{\mathbf{i}} (\Sigma^{-1})_{\mathbf{l} \mathbf{l}'}^{\mathbf{ij}} S^{\mathbf{j}} w_{l'}^{\mathbf{j}}, \tag{B10}
\end{aligned}$$

which is a quadratic form in δA , $\mathcal{L} \propto \exp[-\frac{1}{2}(a(\delta A)^2 - 2b(\delta A) + c)]$, with

$$\begin{aligned}
a &\equiv \sum_{\mathbf{ij} \mathbf{l} \mathbf{l}'} S^{\mathbf{i}} w_l^{\mathbf{i}} (\Sigma^{-1})_{\mathbf{l} \mathbf{l}'}^{\mathbf{ij}} S^{\mathbf{j}} w_{l'}^{\mathbf{j}}, \\
b &\equiv \sum_{\mathbf{ij} \mathbf{l} \mathbf{l}'} (\bar{C}_l^{\mathbf{i}} - C_l^{\text{th}} w_l^{\mathbf{i}}) (\Sigma^{-1})_{\mathbf{l} \mathbf{l}'}^{\mathbf{ij}} S^{\mathbf{j}} w_{l'}^{\mathbf{j}}, \\
c &\equiv \sum_{\mathbf{ij} \mathbf{l} \mathbf{l}'} (\bar{C}_l^{\mathbf{i}} - C_l^{\text{th}} w_l^{\mathbf{i}}) (\Sigma^{-1})_{\mathbf{l} \mathbf{l}'}^{\mathbf{ij}} (\bar{C}_{l'}^{\mathbf{j}} - C_{l'}^{\text{th}} w_{l'}^{\mathbf{j}}). \tag{B11}
\end{aligned}$$

We wish to evaluate the marginalized likelihood, $\mathcal{L}_A = \int \mathcal{L} d(\delta A)$. This is readily evaluated using the substitution $\delta A \rightarrow \delta A - b/a$, giving $\mathcal{L}_A \propto \exp[-\frac{1}{2}(c - b^2/a)]$, where we drop a multiplicative term proportional to a , which is independent of $\bar{C}_l^{\mathbf{i}}$. The marginalized likelihood function is thus

$$-2 \ln \mathcal{L}_A = \sum_{\mathbf{ij} \mathbf{l} \mathbf{l}'} (\bar{C}_l^{\mathbf{i}} - C_l^{\text{th}} w_l^{\mathbf{i}}) (\Sigma^{-1})_{\mathbf{l} \mathbf{l}'}^{\mathbf{ij}} (\bar{C}_{l'}^{\mathbf{j}} - C_{l'}^{\text{th}} w_{l'}^{\mathbf{j}}) \tag{B12}$$

$$- \frac{1}{a} \sum_{\mathbf{ij} \mathbf{l} \mathbf{l}'} (\bar{C}_l^{\mathbf{i}} - C_l^{\text{th}} w_l^{\mathbf{i}}) (\Sigma^{-1})_{\mathbf{l} \mathbf{l}'}^{\mathbf{ij}} S^{\mathbf{i}} w_{l'}^{\mathbf{j}} \sum_{\mathbf{i}' \mathbf{j}' \mathbf{l}'' \mathbf{l}'''} (\bar{C}_{l''}^{\mathbf{i}'} - C_{l''}^{\text{th}} w_{l''}^{\mathbf{i}'}) (\Sigma^{-1})_{\mathbf{l}'' \mathbf{l}'''}^{\mathbf{i}' \mathbf{j}'} S^{\mathbf{j}'} w_{l'''}^{\mathbf{j}'}. \tag{B13}$$

This expression may be recast in the form

$$-2 \ln \mathcal{L}_A = \sum_{\mathbf{ij} \mathbf{l} \mathbf{l}'} (\bar{C}_l^{\mathbf{i}} - C_l^{\text{th}} w_l^{\mathbf{i}}) (F^{\text{src}})_{\mathbf{l} \mathbf{l}'}^{\mathbf{ij}} (\bar{C}_{l'}^{\mathbf{j}} - C_{l'}^{\text{th}} w_{l'}^{\mathbf{j}}), \tag{B14}$$

where $(F^{\text{src}})_{\mathbf{l} \mathbf{l}'}^{\mathbf{ij}}$ is

$$(F^{\text{src}})_{\mathbf{l} \mathbf{l}'}^{\mathbf{ij}} \equiv (\Sigma^{-1})_{\mathbf{l} \mathbf{l}'}^{\mathbf{ij}} - \frac{1}{a} B_{\mathbf{l} \mathbf{l}'}^{\mathbf{ij}}, \tag{B15}$$

with

$$B_{\mathbf{l} \mathbf{l}'}^{\mathbf{ij}} \equiv \sum_{\mathbf{i}' \mathbf{j}' \mathbf{l}'' \mathbf{l}'''} (\Sigma^{-1})_{\mathbf{l} \mathbf{l}''}^{\mathbf{i}' \mathbf{j}'} S^{\mathbf{i}'} w_{l''}^{\mathbf{i}'} (\Sigma^{-1})_{\mathbf{l}'' \mathbf{l}'''}^{\mathbf{j}' \mathbf{i}''} S^{\mathbf{j}'} w_{l'''}^{\mathbf{j}'}. \tag{B16}$$

The superscript ‘‘src’’ indicates that the Fisher matrix so obtained includes point source subtraction uncertainty, in addition to whatever effects have been included in $(\Sigma^{-1})_{\mathbf{l} \mathbf{l}'}^{\mathbf{ij}}$ already, in this case only cosmic variance and noise. Note that equation (B15) neglects a term proportional to $\ln \det a$ which has a weak dependence on cosmological parameters. In the actual calculation, as in the previous section, we assume the diagonal form of the covariance matrix in equation (B1).

C. OPTIMAL WEIGHTING OF MULTI-CHANNEL SPECTRA

We use the 8 high-frequency differencing assemblies Q1 through W4 to generate the final combined spectrum. In this Appendix we show how we combine these spectra into a single estimate of the angular power spectrum.

The ultimate goal of the *WMAP* analysis is to produce a likelihood function for a set of cosmological parameters, $\{\vec{\alpha}\}$, given the data, C_l^i . Specifically

$$\mathcal{P}(\vec{\alpha}|C_l^i) = \mathcal{L}(C_l^i|C_l^{\text{th}}(\vec{\alpha}))\mathcal{P}(\vec{\alpha}), \quad (\text{C1})$$

where $\mathcal{P}(\vec{\alpha}|C_l^i)$ is the probability of $\{\vec{\alpha}\}$ given the data, $\mathcal{L}(C_l^i|C_l^{\text{th}}(\vec{\alpha}))$ is the likelihood of the data given the model, $C_l^{\text{th}}(\vec{\alpha})$, and $\mathcal{P}(\vec{\alpha})$ is the prior probability of the parameter set (Verde et al. 2003). To this end, we seek a combined spectrum, \hat{C}_l , that estimates the power spectrum in our sky, C_l^{sky} , with the property that $\mathcal{P}(\vec{\alpha}|\hat{C}_l) = \mathcal{P}(\vec{\alpha}|C_l^i)$, and hence $\mathcal{L}(\hat{C}_l|C_l^{\text{th}}(\vec{\alpha})) = \mathcal{L}(C_l^i|C_l^{\text{th}}(\vec{\alpha}))$.

To estimate the combined spectrum, we approximate the likelihood function for the cross-power spectra as Gaussian

$$-2 \ln \mathcal{L}(\bar{C}_l^i|C_l^{\text{th}}) = \sum_{\mathbf{ij} l l'} (\bar{C}_l^i - C_l^{\text{th}} w_l^i) (\Sigma_{\text{full}}^{-1})_{ll'}^{\mathbf{ij}} (\bar{C}_l^j - C_l^{\text{th}} w_l^j), \quad (\text{C2})$$

where \bar{C}_l^i is the spectrum with the best-fit source model subtracted, defined after equation (12), w_l^i is the window function of spectrum \mathbf{i} , defined after equation (6), and Σ_{full} is the covariance matrix of the 28 cross-power spectra. Note that the treatment in the remainder of this section does not depend on any specific property of the covariance, so we use generic notation for readability. However, when we generate the *WMAP* first-year combined spectrum, the actual form of the covariance used at this step is $(\tilde{\Sigma}_{\text{full}})_{ll'}^{\mathbf{ij}}$, where $\tilde{\Sigma}$ indicates the approximate covariance, obtained in §4, that has not yet had the effects of the foreground mask accounted for.

We seek a spectrum \hat{C}_l such that

$$-2 \ln \mathcal{L}(\hat{C}_l|C_l^{\text{th}}) \equiv \sum_{l l'} (\hat{C}_l - C_l^{\text{th}}) Q_{ll'} (\hat{C}_l - C_l^{\text{th}}) = -2 \ln \mathcal{L}(\bar{C}_l^i|C_l^{\text{th}}), \quad (\text{C3})$$

where $Q_{ll'}$ is the inverse-covariance matrix of the combined spectrum which comes with the estimate of \hat{C}_l . Suppose, for simplicity, that Σ_{full} is diagonal, $(\Sigma_{\text{full}})_{ll'}^{\mathbf{ij}} = (\Sigma_{\text{full}})_{l l'}^{\mathbf{ij}} \delta_{ll'}$, then it is straightforward to show that the deconvolved, weighted-average spectrum

$$\hat{C}_l = \frac{\sum_{\mathbf{ij}} \bar{C}_l^i (\Sigma_{\text{full}}^{-1})_{ll'}^{\mathbf{ij}} w_l^j}{Q_l}, \quad (\text{C4})$$

with

$$Q_l = \sum_{\mathbf{ij}} w_l^i (\Sigma_{\text{full}}^{-1})_{ll'}^{\mathbf{ij}} w_l^j, \quad (\text{C5})$$

is the desired spectrum. Substituting equations (C4) and (C5) into equation (C3) produces equation (C2) up to a term which has a weak dependence on cosmological model, which we ignore. This combined spectrum is equivalent to the result we would obtain using the ‘‘optimal data compression’’ of Tegmark et al. (1997).

If the inverse covariance matrix is not diagonal in l , it can be shown that the optimal combined spectrum is given by

$$\hat{C}_l = \frac{\sum_{\mathbf{ij} l'} \bar{C}_l^i (\Sigma_{\text{full}}^{-1})_{ll'}^{\mathbf{ij}} C_{l'}^{\text{fid}} w_{l'}^j}{\sum_{l'} Q_{ll'} C_{l'}^{\text{fid}}}, \quad (\text{C6})$$

with

$$Q_{ll'} = \sum_{\mathbf{ij}} w_l^i (\Sigma_{\text{full}}^{-1})_{ll'}^{\mathbf{ij}} w_{l'}^j, \quad (\text{C7})$$

and where C_l^{fid} is a fiducial cosmological model which we take to be a flat Λ CDM model with $\Omega_b h^2 = 0.021$, $\Omega_c h^2 = 0.129$, $\Omega_{\text{tot}} = 1$, $h = 0.68$, $n_s = 1.2$, and $\tau = 0.2$. While this model has parameters values that are substantially different than the best-fit *WMAP* model obtained (afterwards) by Spergel et al. (2003), the parameter degeneracies are such that C_l^{fid} is close to the best-fit model for $l > 100$ where this estimator is actually used (see §5.1). The combined spectrum

is optimal if the fiducial model chosen is the correct one; otherwise it is still unbiased but slightly sub-optimal (Gupta & Heavens 2002).

D. CUT-SKY FISHER MATRIX

The *WMAP* sky maps have nearly diagonal pixel-pixel noise covariance (Hinshaw et al. 2003a) which greatly simplifies the properties of the power spectrum Fisher matrix. In this Appendix we present an analytic derivation of the effect of a sky cut and non-uniform pixel weighting. In D.1, we assume that we have an optimal estimator of the power spectrum. In the noise dominated limit, we can obtain an exact expression, while in the signal dominated limit, we need to approximate the signal correlation matrix to obtain an analytic expression. In D.2, we interpolate the Fisher matrix between the signal and noise-dominated regimes and show that it agrees with numerical estimates. In D.3, we estimate the power spectrum covariance matrix from Monte Carlo simulations of the sky and calibrate the interpolation formula.

D.1. Cut-Sky Fisher Matrix: Analytic Evaluation

We equate the Fisher matrix of the power spectrum to the curvature of the likelihood function, equation (A16), then develop an approximate form for the covariance matrix, $\mathbf{C} = \mathbf{S} + \mathbf{N}$, where \mathbf{S} is the signal matrix and \mathbf{N} is the noise matrix. We split the noise matrix into two pieces: a weight term and a mask term. In the limit that the pixel noise is diagonal, the weight term has the form

$$(\mathbf{N}_w^{-1})_{ij} = \frac{n_i}{\sigma_0^2} \delta_{ij} \equiv w_i \delta_{ij} \quad (\text{D1})$$

where i and j are pixel indices, n_i is the number of observations of pixel i , σ_0 is the rms noise of a single observation, and, by definition, w_i is the weight of pixel i . The mask, M_i , is defined so that M_i equals 0 in pixels that are not used due to foreground contamination, and equals 1 otherwise. The noise matrix can then be written as the product of the two terms

$$\mathbf{N}^{-1} = \mathbf{N}_w^{-1} \mathbf{M} = \mathbf{M} \mathbf{N}_w^{-1} = \frac{\hat{n}_i}{\sigma_0^2} \delta_{ij} = \hat{w}_i \delta_{ij} \quad (\text{D2})$$

where $\hat{n}_i = n_i$ and $\hat{w}_i = w_i$ in the unmasked pixels, and are set to 0 otherwise. We thus define the covariance matrix over the full sky, which allows us to exploit the orthogonality properties of the spherical harmonics. Note that $\mathbf{M}^2 = \mathbf{M}$.

The inverse of the full covariance matrix can now be written as

$$\mathbf{C}^{-1} = (\mathbf{S} + \mathbf{N})^{-1} = \mathbf{N}^{-1} (\mathbf{S} \mathbf{N}^{-1} + \mathbf{I})^{-1} \quad (\text{D3})$$

$$= \mathbf{M}^{-1} (\mathbf{S} \mathbf{M}^{-1} + \mathbf{N}_w)^{-1}. \quad (\text{D4})$$

The covariance matrix has two limits. In the noise dominated limit, $\mathbf{S} \mathbf{N}^{-1} \ll \mathbf{I}$,

$$\mathbf{C}^{-1} \rightarrow \mathbf{N}^{-1}. \quad (\text{D5})$$

In the signal dominated limit, $\mathbf{S} \mathbf{M}^{-1} \ll \mathbf{N}_w$, only the mask alters the covariance matrix, so

$$\mathbf{C}^{-1} \rightarrow \mathbf{M} \mathbf{S}^{-1} \mathbf{M}, \quad (\text{D6})$$

where we have set the inverse of the mask to zero where there is no data, i.e. $\mathbf{M}^{-1} \equiv \mathbf{M}$.

D.1.1. Cut-Sky Fisher Matrix in the Noise Dominated Limit

In the noise dominated limit, the Fisher matrix, equation (A16), is approximately

$$F_{l'l'}^N = \frac{1}{2} \text{tr} (\mathbf{C}^{-1} \mathbf{P}_l \mathbf{C}^{-1} \mathbf{P}_{l'}) \rightarrow \frac{1}{2} \text{tr} (\mathbf{N}^{-1} \mathbf{P}_l \mathbf{N}^{-1} \mathbf{P}_{l'}). \quad (\text{D7})$$

Using

$$(\mathbf{P}_l)_{ij} \equiv \frac{\partial \mathbf{C}}{\partial C_l} = \frac{(2l+1)}{4\pi} w_l P_l(\cos \theta_{ij}) = w_l \sum_m Y_{lm}^*(p_i) Y_{lm}(p_j) \quad (\text{D8})$$

where θ_{ij} is the angle between pixels i and j , $Y_{lm}(p_i)$ is a spherical harmonic evaluated in the direction of pixel i , and we have employed the addition theorem for spherical harmonics in the last equality. Then the Fisher matrix takes the form

$$F_{l'l'}^N = \frac{1}{2} w_l w_{l'} \sum_{mm'} \sum_{ij} \hat{w}_i Y_{lm}^*(p_i) Y_{lm}(p_j) \hat{w}_j Y_{l'm'}^*(p_j) Y_{l'm'}(p_i). \quad (\text{D9})$$

Now expand the weight array as

$$\hat{w}_i = \sum_{lm} \hat{w}_{lm} Y_{lm}(p_i) \quad (\text{D10})$$

and substitute this into equation (D9) to yield

$$F_{l'l'}^N = \frac{1}{2} w_l w_{l'} \sum_{mm'} \sum_{l''m''} \sum_{l''''m''''} \hat{w}_{l''m''} \hat{w}_{l''''m''''} \sum_{ij} Y_{l''m''}(p_i) Y_{l''m''}^*(p_i) Y_{lm}(p_j) Y_{l''m''}(p_j) Y_{l''m''}^*(p_j) Y_{l'm'}(p_i). \quad (\text{D11})$$

Since \hat{w}_i was defined over the full sky, the sum over pixels may be expressed in terms of products of Wigner 3- j symbols. As shown in Appendix E, we may use an orthogonality property of the 3- j symbols to reduce the expression for the Fisher matrix to

$$F_{l'l'}^N = \frac{1}{2} w_l w_{l'} \frac{(2l+1)(2l'+1)}{4\pi \Omega_p^2} \sum_{l''} \hat{W}_{l''} \begin{pmatrix} l & l' & l'' \\ 0 & 0 & 0 \end{pmatrix}^2 \quad (\text{D12})$$

where $\hat{W}_l \equiv \sum_m |\hat{w}_{lm}|^2$, and Ω_p is the solid angle per pixel.

D.1.2. Cut-Sky Fisher Matrix in the Signal Dominated Limit

In the signal dominated limit, the Fisher matrix, equation (A16), is approximately

$$F_{l'l'}^S = \frac{1}{2} \text{tr} (\mathbf{C}^{-1} \mathbf{P}_l \mathbf{C}^{-1} \mathbf{P}_{l'}) \rightarrow \frac{1}{2} \text{tr} (\hat{\mathbf{S}}^{-1} \mathbf{P}_l \hat{\mathbf{S}}^{-1} \mathbf{P}_{l'}), \quad (\text{D13})$$

where $\hat{\mathbf{S}}^{-1} \equiv M_i S_{ij}^{-1} M_j$. Using equation (D8) we can write $F_{l'l'}^S$ in the form

$$F_{l'l'}^S = \frac{1}{2} w_l w_{l'} \sum_{mm'} \sum_{ijk} M_i M_j S_{ij}^{-1} Y_{lm}^*(p_j) Y_{lm}(p_k) \sum_{kp} M_k M_p S_{kp}^{-1} Y_{l'm'}^*(p_p) Y_{l'm'}(p_i). \quad (\text{D14})$$

Substituting a normalized expression for \mathbf{S}^{-1} in the pixel basis,

$$S_{ij}^{-1} = \Omega_p^2 \sum_{lm} \frac{1}{C_l w_l} Y_{lm}^*(p_i) Y_{lm}(p_j), \quad (\text{D15})$$

into equation (D14), one obtains a term

$$\Omega_p^2 \sum_{l''} \frac{1}{C_{l''} w_{l''}} \sum_{m''} \sum_i M_i Y_{l'' m''}(p_i) Y_{l'' m''}^*(p_i) \sum_j M_j Y_{l'' m''}(p_j) Y_{l'' m''}^*(p_j). \quad (\text{D16})$$

Since our mask cuts only $\sim 15\%$ of the sky, the coupling sum, $\sum_i M_i Y_{l'' m''}(p_i) Y_{l'' m''}^*(p_i)$, peaks very sharply at $|l'' - l''| \ll l''$. Therefore, one may approximate equation (D16) with

$$\begin{aligned} & \frac{\Omega_p^2}{(C_l w_l C_{l'} w_{l'})^{1/2}} \sum_{l'' m''} \sum_i M_i Y_{l'' m''}(p_i) Y_{l'' m''}^*(p_i) \sum_j M_j Y_{l'' m''}(p_j) Y_{l'' m''}^*(p_j) \\ &= \frac{\Omega_p}{(C_l w_l C_{l'} w_{l'})^{1/2}} \sum_i M_i Y_{l' m'}(p_i) Y_{l' m'}^*(p_i), \end{aligned} \quad (\text{D17})$$

where, in the last equality, we have used the completeness relation for the spherical harmonics

$$\sum_{lm} Y_{lm}^*(p_i) Y_{lm}(p_j) = \Omega_p^{-1} \delta_{ij}. \quad (\text{D18})$$

With this approximation equation (D14) reduces to

$$F_{l'l'}^S \simeq \frac{1}{2} \frac{\Omega_p^2}{C_l C_{l'}} \sum_{mm'} \sum_{ij} M_i Y_{lm}^*(p_i) Y_{lm}(p_j) M_j Y_{l' m'}^*(p_j) Y_{l' m'}(p_i). \quad (\text{D19})$$

Following the same steps outlined above for the derivation of equation (D12) yields

$$F_{l'l'}^S = \frac{1}{2} \frac{1}{C_l C_{l'}} \frac{(2l+1)(2l'+1)}{4\pi} \sum_{l''} M_{l''} \begin{pmatrix} l & l' & l'' \\ 0 & 0 & 0 \end{pmatrix}^2 \quad (\text{D20})$$

where $M_l \equiv \sum_m |m_{lm}|^2$, and m_{lm} is the spherical harmonic transform of the mask. Note that $M_0 = 4\pi f_{\text{sky}}^2$, where f_{sky} is the fraction of the sky that survives the mask.

D.2. Interpolating the Cut-Sky Fisher Matrix

We can combine the two limiting cases of the Fisher matrix, obtained in the previous section, into a single expression

$$F_{l'l'}^{\text{mask}} = \frac{\sqrt{(2l+1)(2l'+1)} f_{\text{sky}}^2}{2(C_l + N_l)(C_{l'} + N_{l'})} \tilde{F}_{l'l'}, \quad (\text{D21})$$

where $N_l \equiv \Omega_p \sigma_0^2 / (\bar{n}_{\text{obs}} w_l)$ is the deconvolved noise power spectrum. Using equations (D12) and (D20) derived in the previous section, $\tilde{F}_{l'l'}$ can be expressed in the low l (signal-dominated) and high l (noise-dominated) limits as

$$\tilde{F}_{l'l'}^S = \frac{\sqrt{(2l+1)(2l'+1)}}{4\pi f_{\text{sky}}^2} \sum_{l''} M_{l''} \begin{pmatrix} l & l' & l'' \\ 0 & 0 & 0 \end{pmatrix}^2, \quad (\text{D22})$$

and

$$\tilde{F}_{l'l'}^N = \frac{\sqrt{(2l+1)(2l'+1)}}{4\pi f_{\text{sky}}^2 \bar{n}_{\text{obs}}^2 / \sigma_0^4} \sum_{l''} \hat{W}_{l''} \begin{pmatrix} l & l' & l'' \\ 0 & 0 & 0 \end{pmatrix}^2. \quad (\text{D23})$$

By construction, these matrices are normalized to 1 on the diagonal, $\tilde{F}_{ll}^S = \tilde{F}_{ll}^N = 1$, since $M_0 = 4\pi f_{\text{sky}}^2$ and $\hat{W}_0 = 4\pi f_{\text{sky}}^2 \bar{n}_{\text{obs}}^2 / \sigma_0^4$.

We have computed $F_{l'l'}^{\text{mask}}$ directly, for selected l values, by evaluating $Y_{lm}\mathbf{C}^{-1}Y_{l'm'}$ using the pre-conditioner code described by Oh et al. (1999). We find that deviations between the numerical and analytical results are consistent with numerical noise in the Fisher matrix estimate. Figure 14 shows the estimated Fisher matrix for the Kp2 cut in the noise and the signal dominated limits.

We interpolate between the two regimes with the following expression:

$$F_{l'l'}^{\text{mask}} = \frac{\left(C_l\sqrt{F_{l'l'}^S} + N_l\sqrt{F_{l'l'}^N}\right)\left(C_{l'}\sqrt{F_{l'l'}^S} + N_{l'}\sqrt{F_{l'l'}^N}\right)}{(C_l + N_l)(C_{l'} + N_{l'})}. \quad (\text{D24})$$

Note that the form of the Fisher matrix primarily depends on Δl , but is weakly dependent on l – there is more coupling in the noise dominated limit.

D.3. Power Spectrum Covariance: Monte Carlo Evaluation

As discussed in Appendix A.1.2 we compute the C_l using three different pixel weightings. The uniform weighting and the inverse-noise inverse weighting are optimal in the signal-dominated regime and the noise-dominated regime, respectively. In between these limits the transitional weighting performs better. In order to determine which ranges in l correspond to which regimes, and to calibrate our *ansatz* for the covariance matrix, equation (D21), we proceed as follows. Using 100,000 Monte Carlo simulations of signal plus noise (with *WMAP* noise levels and symmetrized beams), we compute the diagonal elements of the covariance matrix for the three different weighting schemes, evaluated with the Kp2 sky cut. Denote these estimates D_l^{sim}

We find that the uniform weighting produces the smallest D_l^{sim} below $l = 200$. Inverse-noise weighting is the best above $l = 450$, and the transitional weighting produces the lowest variance in between. In each of these regimes we use the resulting D_l^{sim} to calibrate our *ansatz* for the diagonal elements of the covariance matrix

$$D_l = \frac{2}{2l+1} \frac{1}{f_{\text{sky}}^{\text{eff}}} (C_l + N_l^{\text{eff}})^2, \quad (\text{D25})$$

as illustrated in Verde et al. (2003). These calibrations produce a smooth correction to equation (D24) of at most 6%. No correction at all is required in the signal dominated regime.

E. SOME USEFUL PROPERTIES OF SPHERICAL HARMONICS

The derivation of the form of the Fisher matrix in the signal and noise dominated limits led to expressions which included a term

$$\sum_{mm'} \sum_{ij} Y_{l'm''}(p_i) Y_{lm}^*(p_i) Y_{lm}(p_j) Y_{l''m''}(p_j) Y_{l'm'}^*(p_j) Y_{l'm'}(p_i), \quad (\text{E1})$$

where the sum is effectively a double integral over the full sky. This can be evaluated in terms of the Wigner 3- j symbols, defined as

$$\int d\Omega_{\mathbf{n}} Y_{lm}(\mathbf{n}) Y_{l'm'}^*(\mathbf{n}) Y_{l''m''}(\mathbf{n}) \equiv (-1)^{m'} \sqrt{\frac{(2l+1)(2l'+1)(2l''+1)}{4\pi}} \begin{pmatrix} l & l' & l'' \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} l & l' & l'' \\ m & -m' & m'' \end{pmatrix}. \quad (\text{E2})$$

Substituting this into equation (E1) gives

$$\sum_{mm'} (-1)^{m+m'} \frac{(2l+1)(2l'+1)}{4\pi} \sqrt{(2l''+1)(2l''' + 1)} \begin{pmatrix} l & l' & l'' \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} l & l' & l''' \\ 0 & 0 & 0 \end{pmatrix}$$

$$\times \begin{pmatrix} l & l' & l'' \\ m & -m' & m'' \end{pmatrix} \begin{pmatrix} l & l' & l''' \\ m & -m' & m''' \end{pmatrix}. \quad (\text{E3})$$

We simplify this using the Wigner 3- j orthogonality condition

$$\sum_{mm'} (-1)^{m+m'} \begin{pmatrix} l & l' & l'' \\ -m & m' & m'' \end{pmatrix} \begin{pmatrix} l & l' & l''' \\ m & -m' & m''' \end{pmatrix} = \frac{1}{2l''+1} \delta_{l'' l'''} \delta_{m'' m'''} \delta(l, l', l''), \quad (\text{E4})$$

where $\delta(l, l', l'') = 1$ for $|l-l'| \leq l'' \leq l+l'$ and is 0 otherwise. This reduces equation (E1) to

$$\frac{(2l+1)(2l'+1)}{4\pi\Omega_p^2} \begin{pmatrix} l & l' & l'' \\ 0 & 0 & 0 \end{pmatrix}^2, \quad (\text{E5})$$

where the factor of Ω_p^2 accounts for the fact that equation (E1) is a double sum over pixels, instead of a double integral.

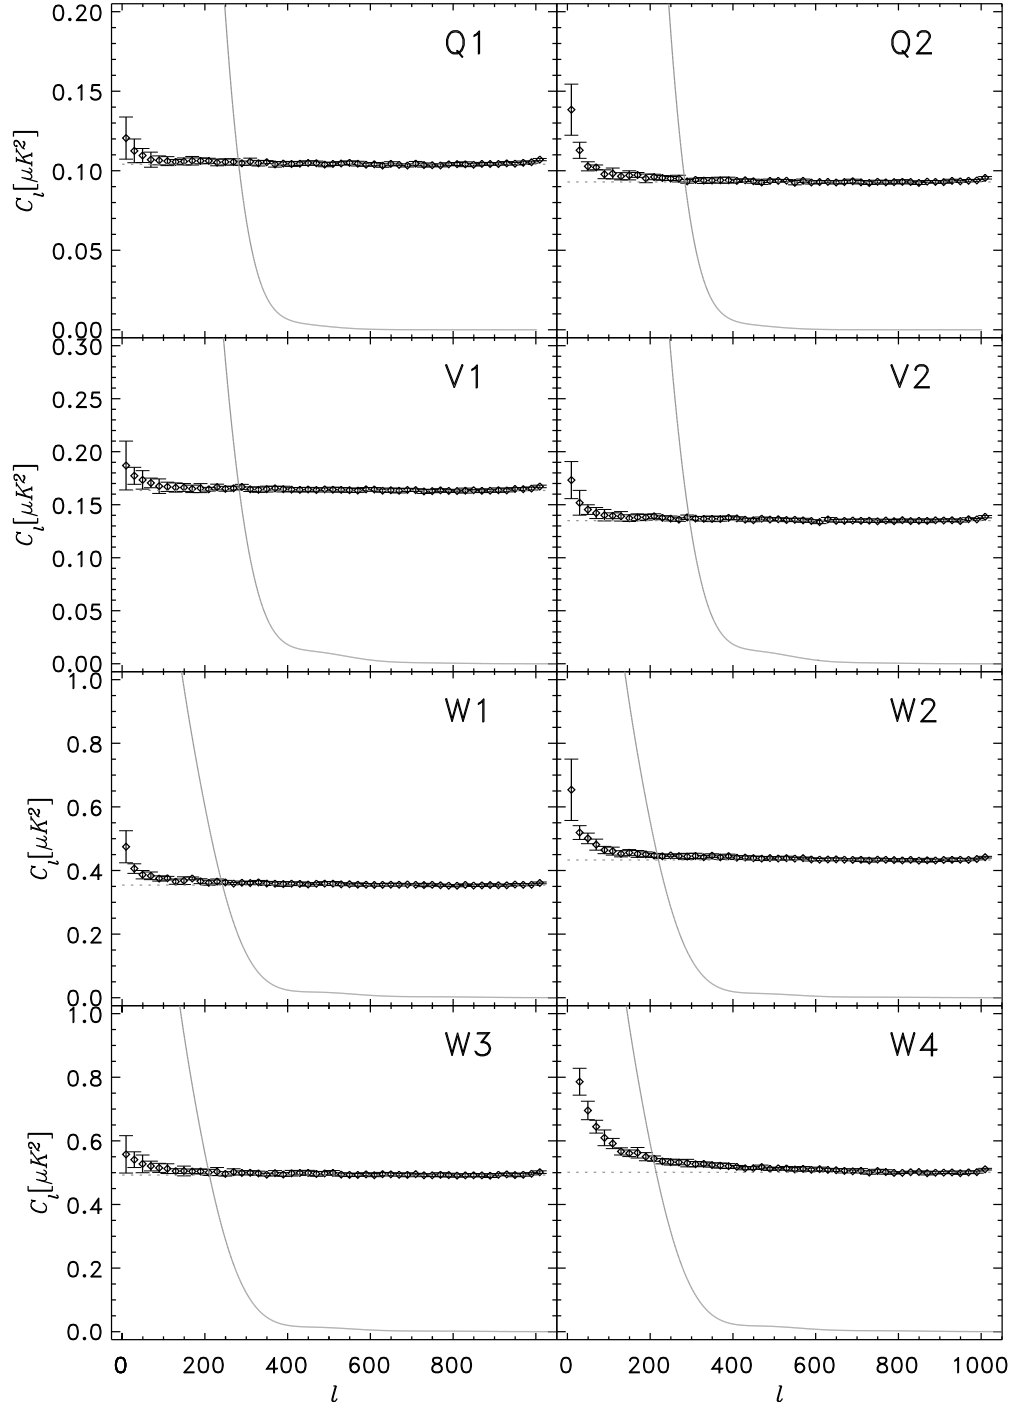


Fig. 1.— The effective noise as a function of l for the 8 differencing assemblies used in the combined power spectrum analysis. These spectra were computed from end-end simulations of noise maps, as discussed in §2.2. For illustration, the spectra shown here were computed using the quadratic estimator with uniform pixel weights. The actual noise model uses three separate weighting schemes in three separate l ranges, and thus has discontinuities where the effective noise level changes. The weights are defined in Appendix A.1.2.

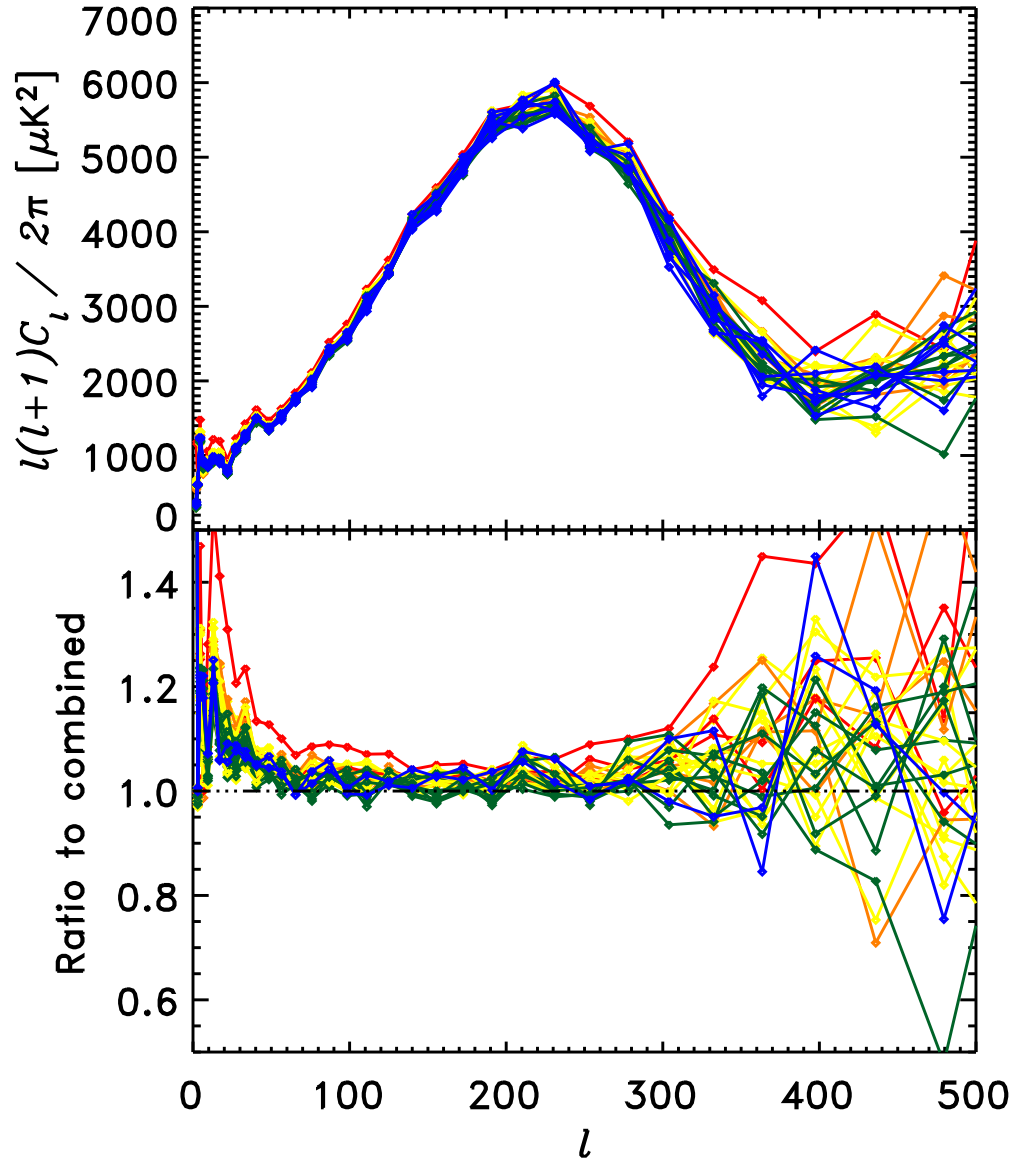


Fig. 2.— The full set of individual cross-power spectra for $l < 500$, computed from the 8 high frequency differencing assemblies Q1 through W4, 28 spectra in all. The spectra were evaluated from the uncorrected sky maps (no Galaxy model subtracted) using the Kp2 sky cut with uniform weighting. The data are plotted in color by effective frequency $\sqrt{\nu_i \nu_j}$ with red corresponding to 41 GHz and blue to 94 GHz. The top panel shows a very robust measurement of the first acoustic peak with a maximum near $l \sim 220$. There is also a clear indication of the rise to a second peak at $l \sim 540$ as discussed in §6. The bottom panel shows the ratio of each channel to the combined spectrum presented in §5. This clearly shows the residual foreground emission due to diffuse Galactic radio emission at low l and to point sources at higher l . The level of contamination, which is strongest at Q band, is consistent with the expected level of foreground emission. See Figure 3 for the spectra after foreground subtraction.

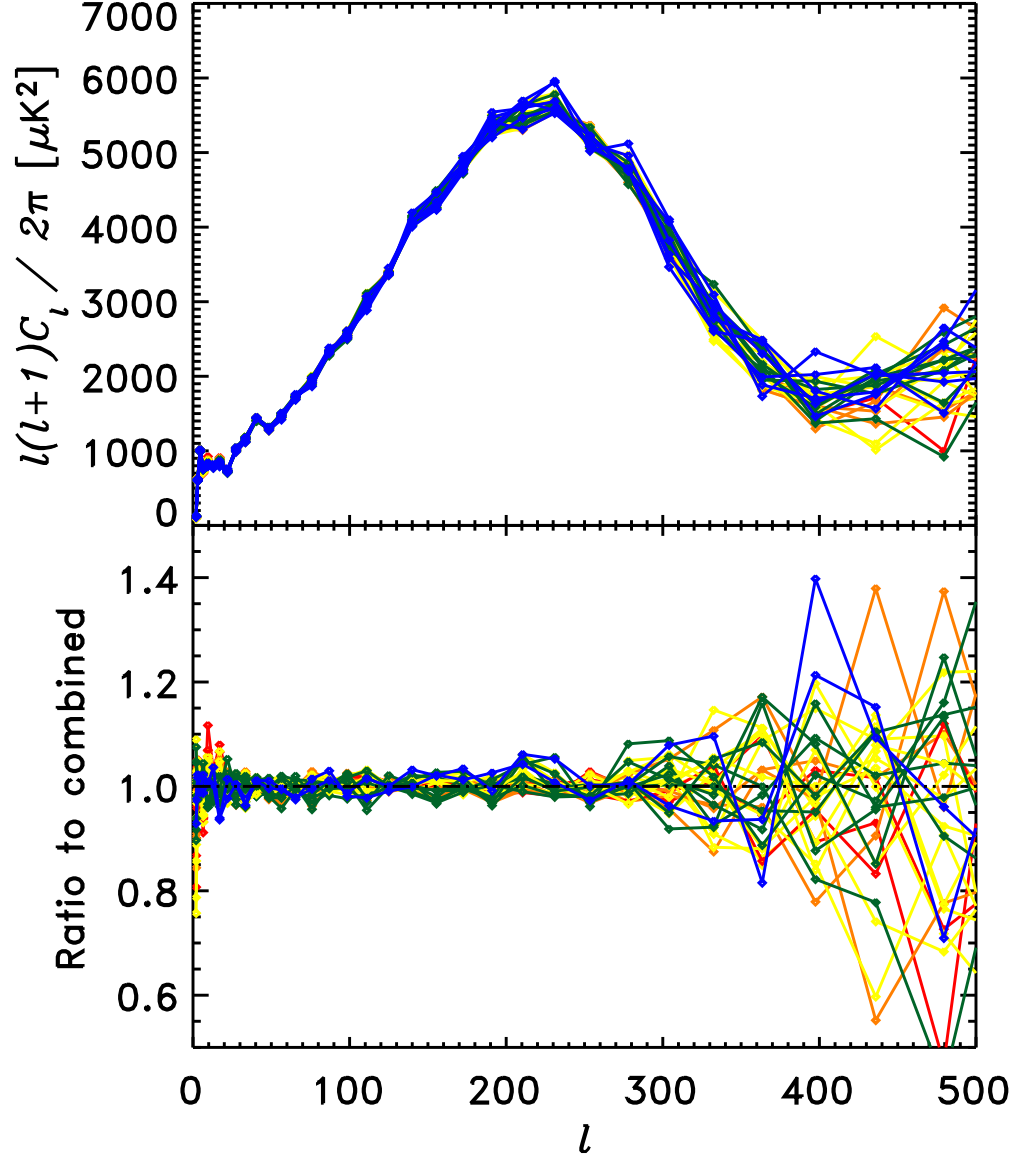


Fig. 3.— The same set of cross-power spectra as shown in Figure 2. Here, the foreground model discussed in §3.1 has been subtracted from each channel. The bottom panel shows the ratio of each of the 28 cross-power spectra to the combined spectrum presented in §5. Aside from a $\sim 10\%$ discrepancy in the Q band data at $l < 20$, the data are consistent with each other to the sensitivity limits of the individual spectra. Because the *WMAP* data are not sensitivity limited at low l , we use only V and W band data in the final combined spectrum for $l < 100$ (see §5) to minimize residual Galactic contamination.

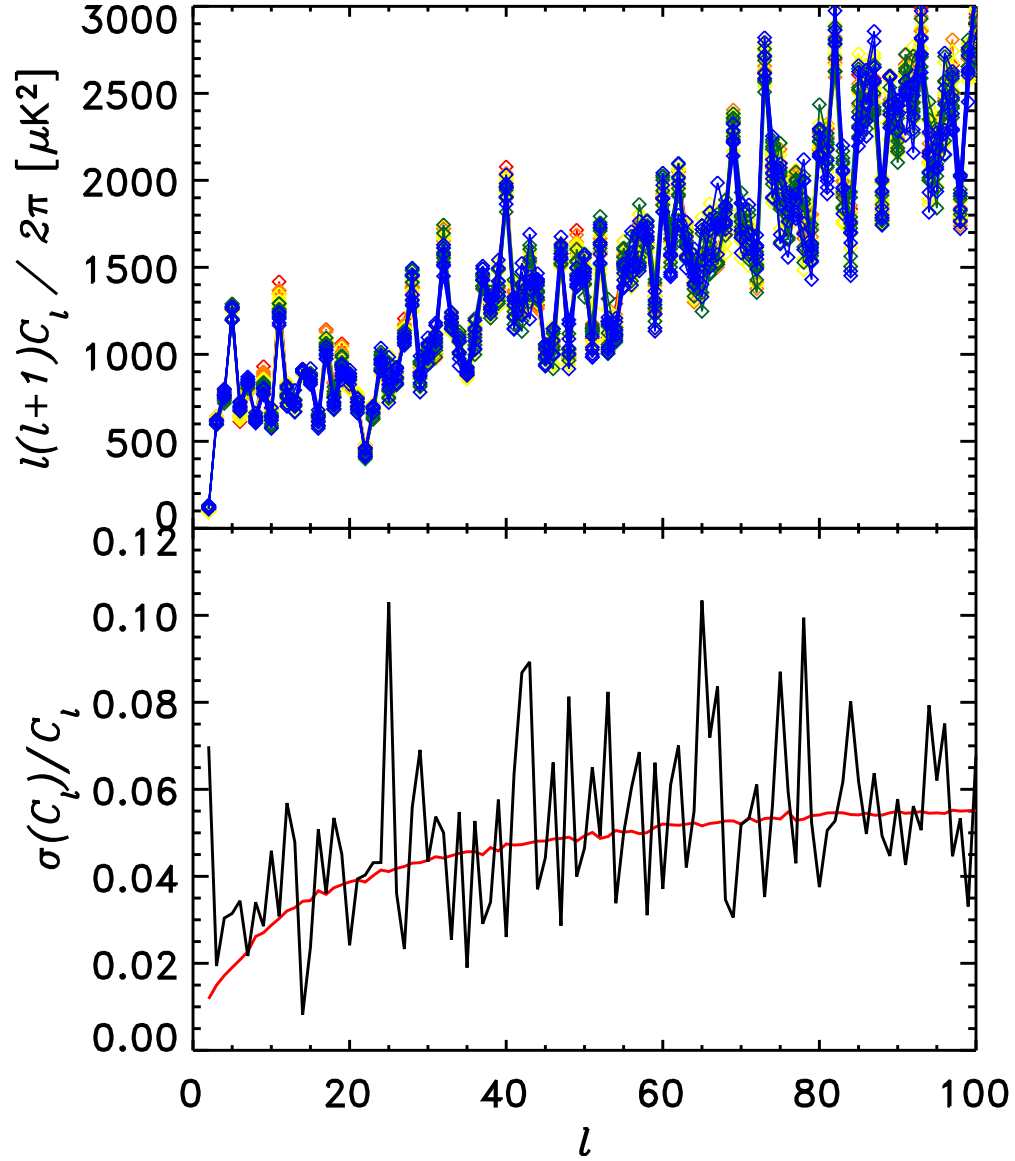


Fig. 4.— The same set of cross-power spectra as shown in Figure 3, but showing the low l spectrum unbinned. The agreement between the individual spectra is striking. The low value of the quadrupole moment, C_2 , that was first seen by *COBE-DMR* is also seen in the *WMAP* data. The steep, nearly linear rise in the spectrum from $l = 2$ to 5 translates to a near absence of power in the angular correlation function at separations larger than $\sim 60^\circ$ (Spergel et al. 2003; Bennett et al. 2003b). This was also seen in the *COBE-DMR* data, but it is now clear that this is not due to Galaxy modeling errors. The bottom panel shows the fractional rms among the 28 *WMAP* cross-power spectra in black, while the red curve shows the same statistic averaged over an ensemble of 1000 Monte Carlo realizations. Based on this we estimate the measurement error on the combined spectrum to be $< 2\text{-}5\%$ for $l < 100$.

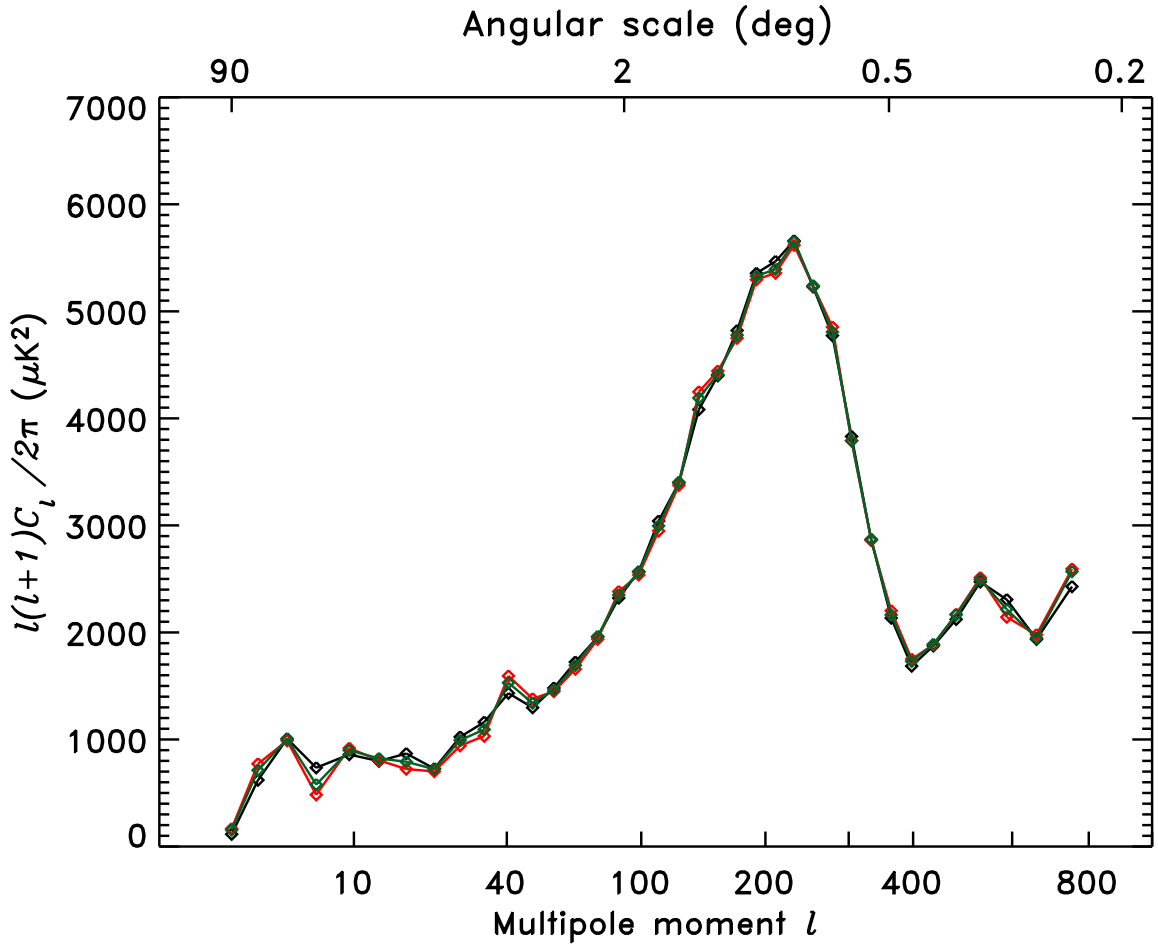


Fig. 5.— The auto-power spectrum of the combined Q+V+W map evaluated with the three weighting schemes defined in Appendix A.1.2. In each case, the spectrum was computed over the entire l range, and black shows uniform weights, red shows inverse-noise weights, and green shows transitional weights. The agreement is excellent.

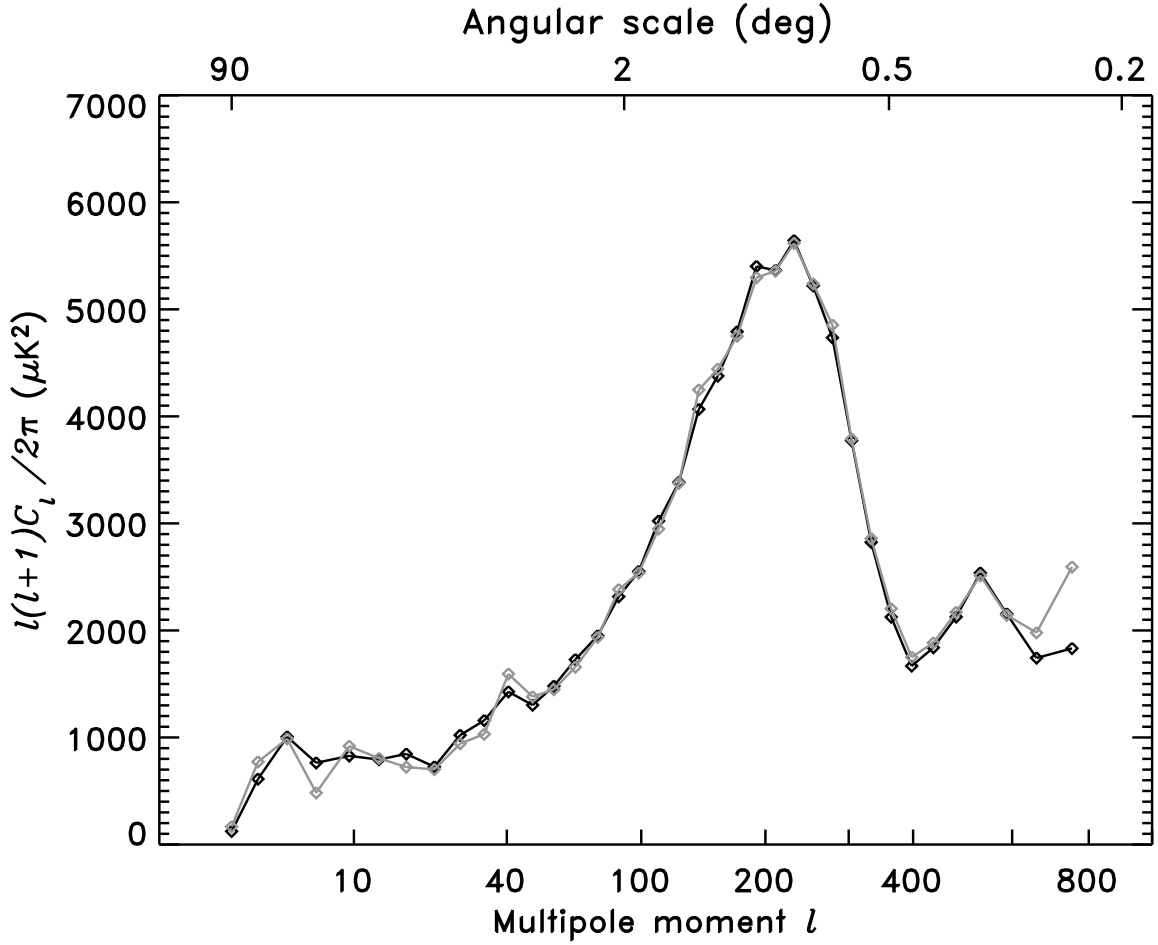


Fig. 6.— This figure compares the auto-power spectrum computed from the combined Q+V+W map (grey) to the optimally combined cross-power spectrum (black). The close agreement indicates that the noise properties of the first-year WMAP data are well understood

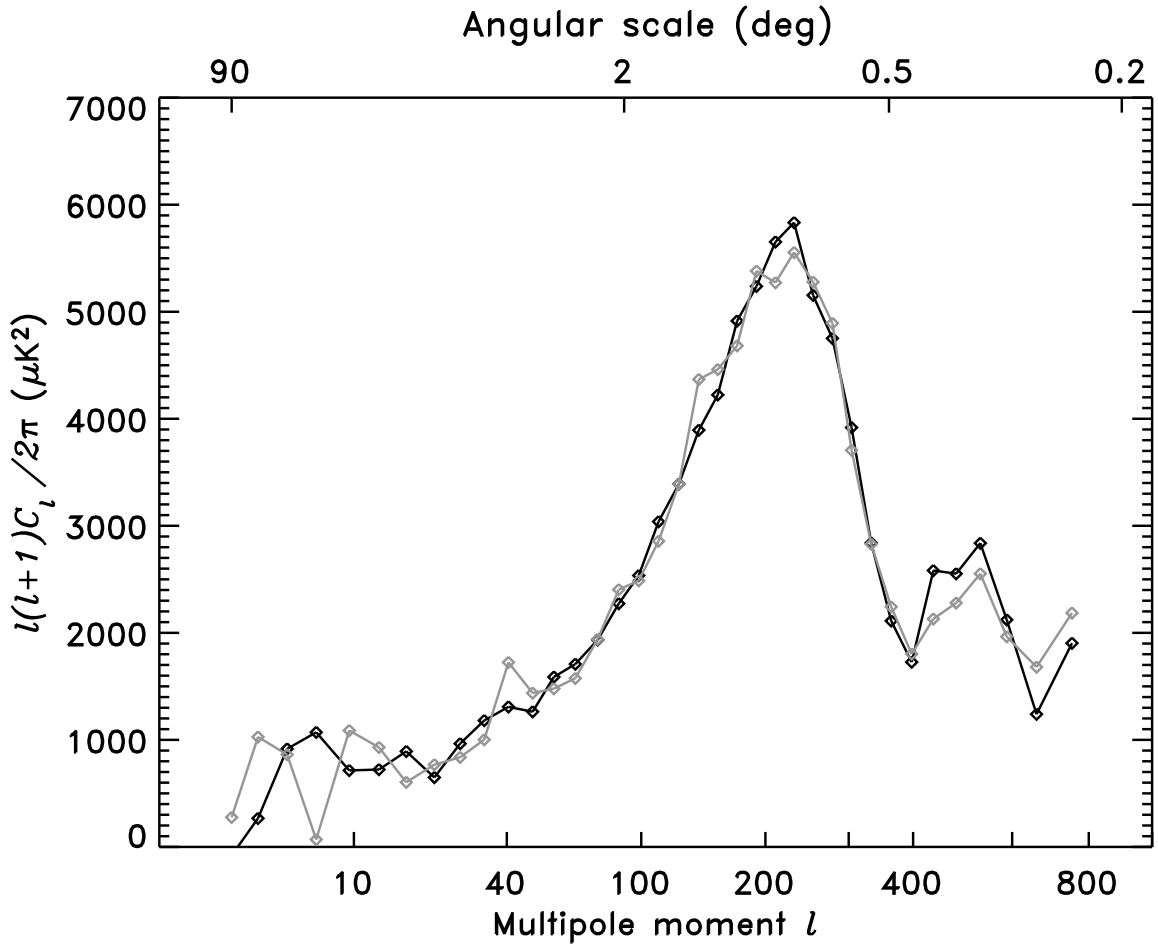


Fig. 7.— A comparison of the power spectrum computed with data from the ecliptic plane (black) vs. data from the ecliptic poles (grey). Note that some of the “bite” features that appear in the combined spectrum are not robust to data excision. There is also no evidence that beam ellipticity, which would be more manifest in the plane than in the poles, systematically biases the spectrum. This is consistent with estimates of the effect given by Page et al. (2003a).

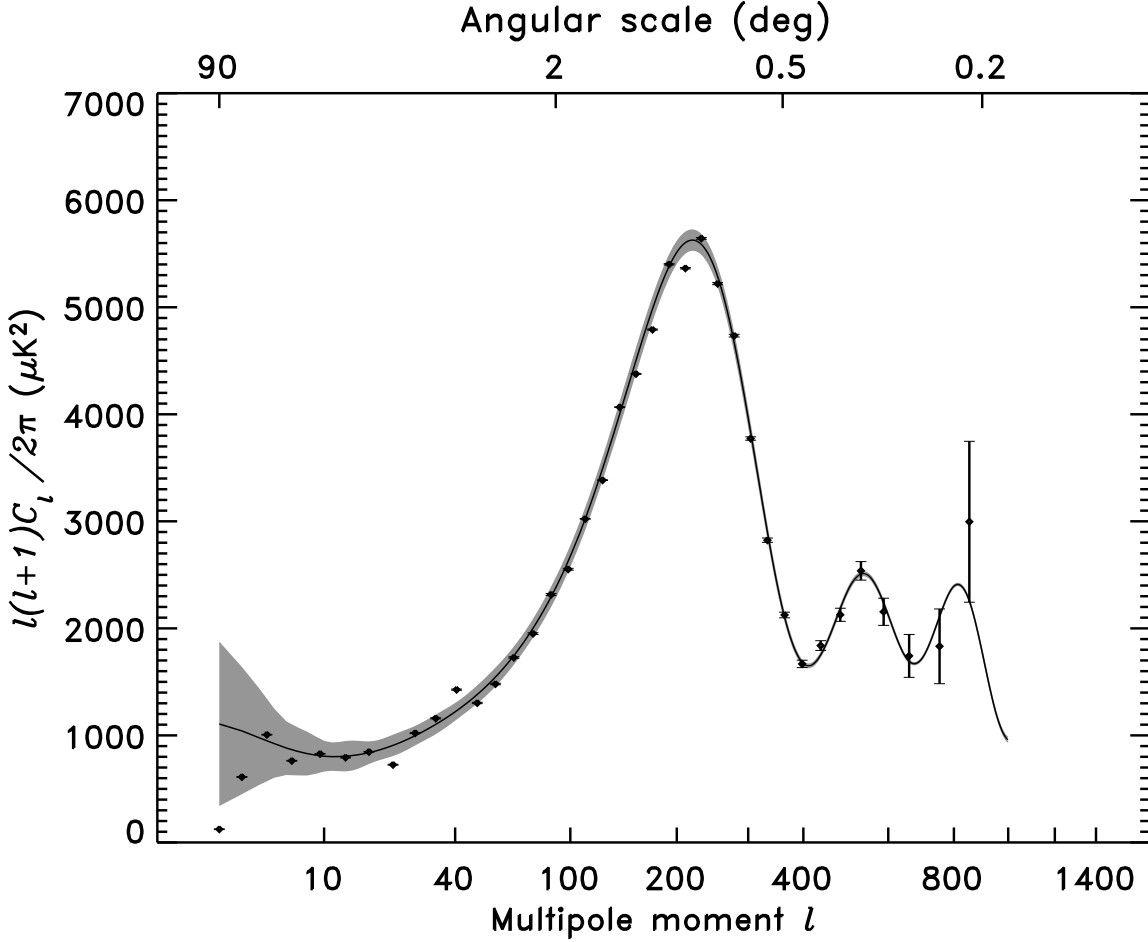


Fig. 8.— The final angular power spectrum, $l(l+1)C_l/2\pi$, obtained from the 28 cross-power spectra, as described in §5. The data are plotted with 1σ measurement errors only which reflect the combined uncertainty due to noise, beam, calibration, and source subtraction uncertainties. The solid line shows the best-fit Λ CDM model from Spergel et al. (2003). The grey band around the model is the 1σ uncertainty due to cosmic variance on the cut sky. For this plot, both the model and the error band have been binned with the same boundaries as the data, but they have been plotted as a splined curve to guide the eye. On the scale of this plot the unbinned model curve would be virtually indistinguishable from the binned curve except in the vicinity of the third peak.

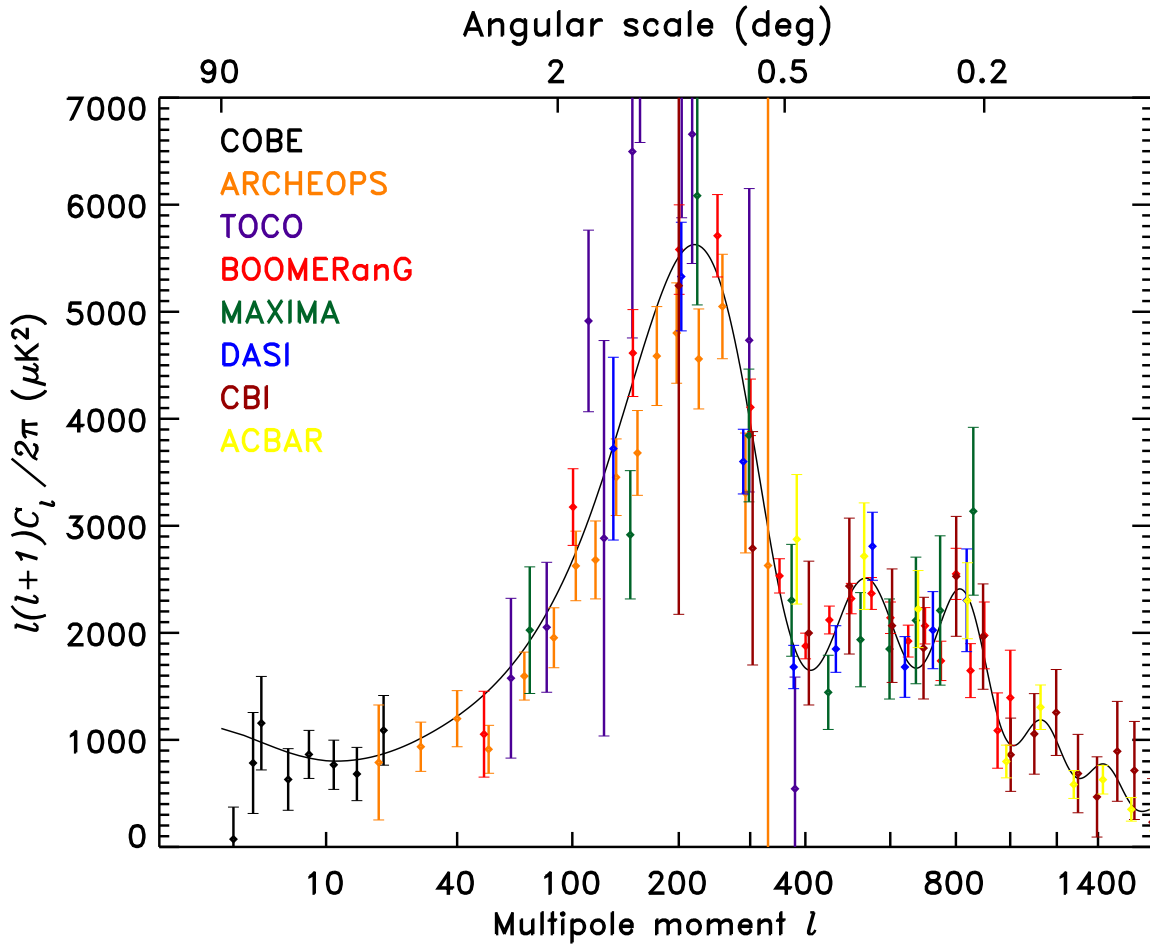


Fig. 9.— A compilation of recent CMB power spectrum measurements compared to the best-fit Λ CDM model from the first-year *WMAP* data. The data points include noise and cosmic variance uncertainty (but not calibration uncertainty) thus we omit the cosmic variance band from the model curve in the Figure. On average, the pre-*WMAP* data agree well with the *WMAP* power spectrum.

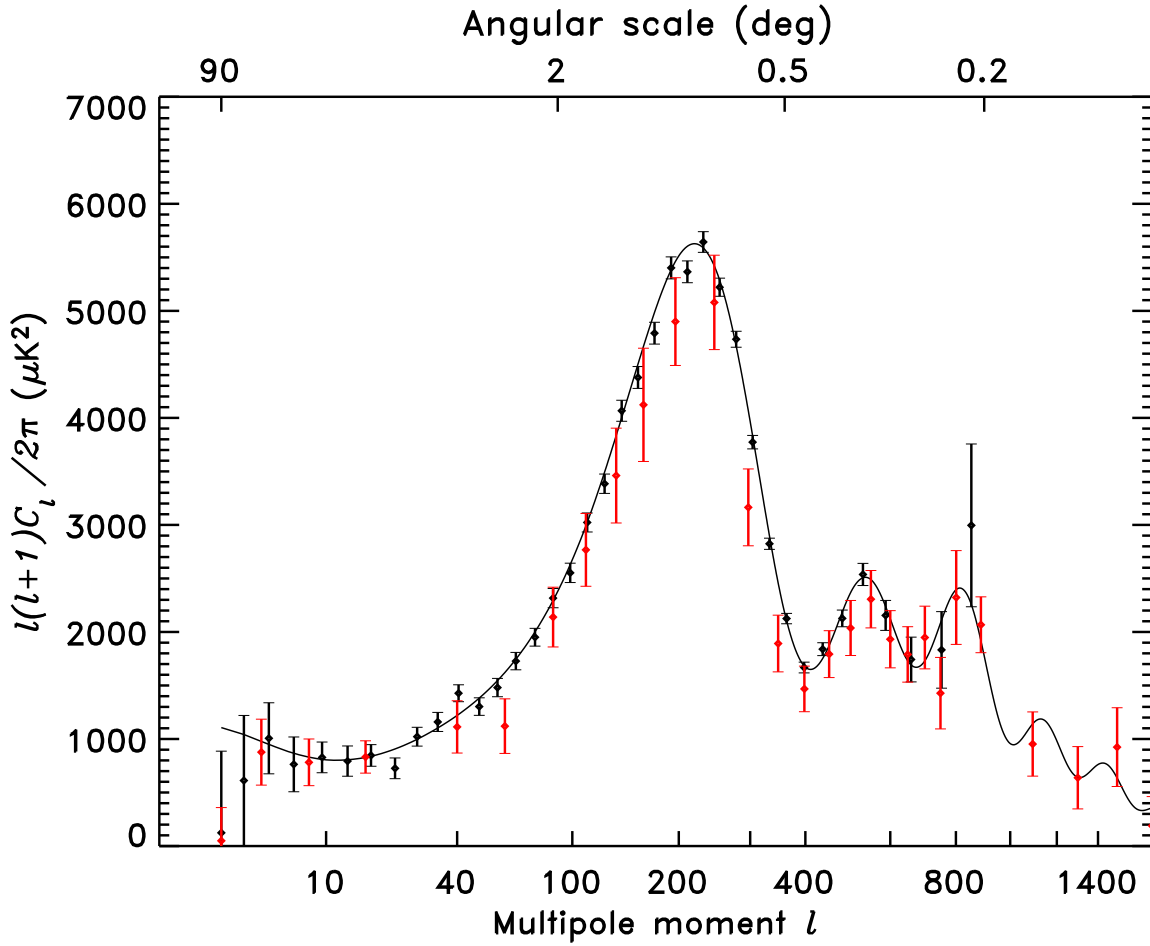


Fig. 10.— The *WMAP* combined power spectrum, in black, compared to a compilation of all CMB data published prior to *WMAP* from Wang et al. (2002a), in red. The *WMAP* data are plotted with cosmic variance plus measurement uncertainties here in order to facilitate a comparison with the compiled data which is reported in this way. The data agree well on *COBE* scales, $l < 20$, (but note that the *WMAP* cosmic variance errors are computed from the best-fit model rather than the data, thus they appear larger than the *COBE* errors at the quadrupole). However, the overall normalization of the *WMAP* spectrum is $\sim 10\%$ higher on smaller scales.

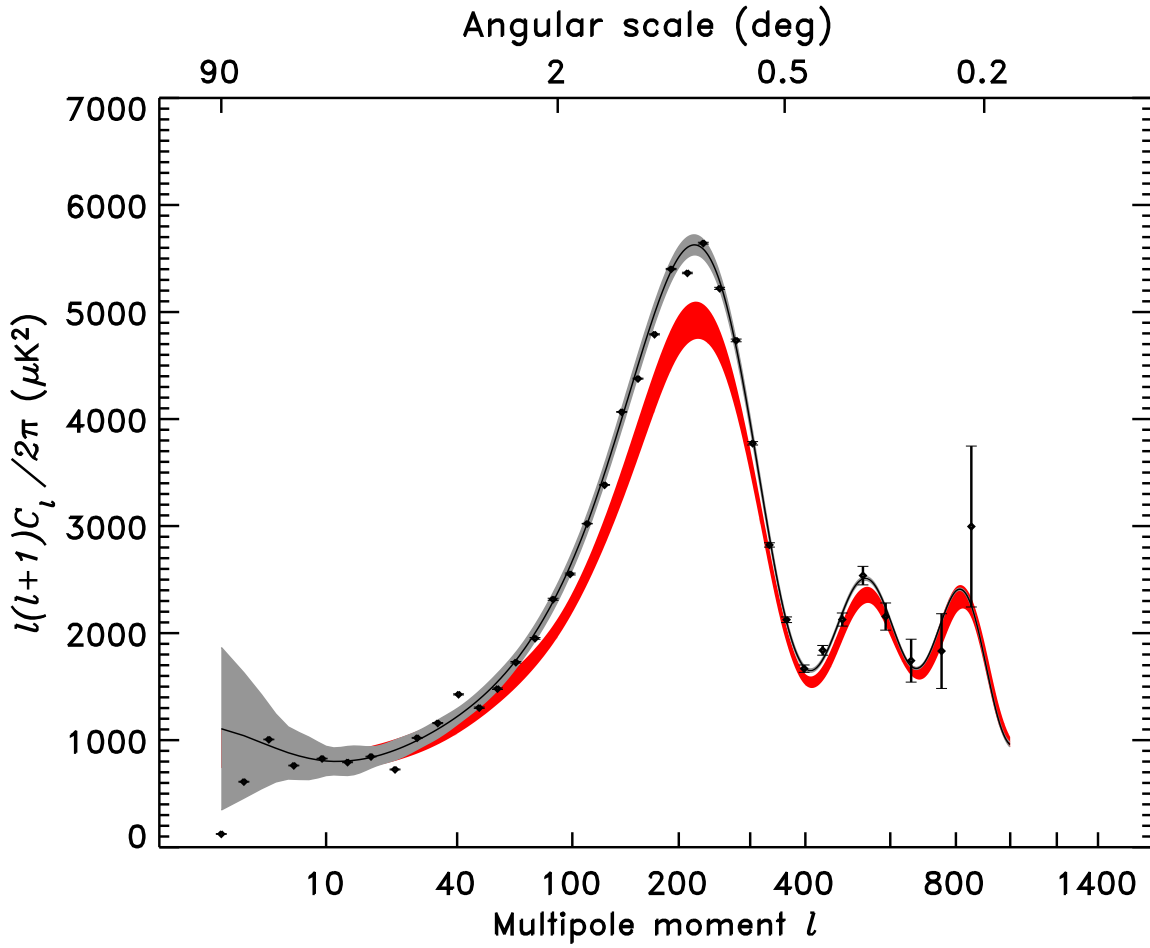


Fig. 11.— The *WMAP* combined power spectrum compared to the locus of predicted spectra, in red, based on a joint analysis of pre-*WMAP* CMB data and 2dFGRS large-scale structure data (Percival et al. 2002). As in Figure 8, the *WMAP* data are plotted with measurement uncertainties, and the best-fit Λ CDM model (Spergel et al. 2003) is plotted with a 1σ cosmic variance error band. The locus of predicted spectra lie systematically below the *WMAP* data at intermediate l .

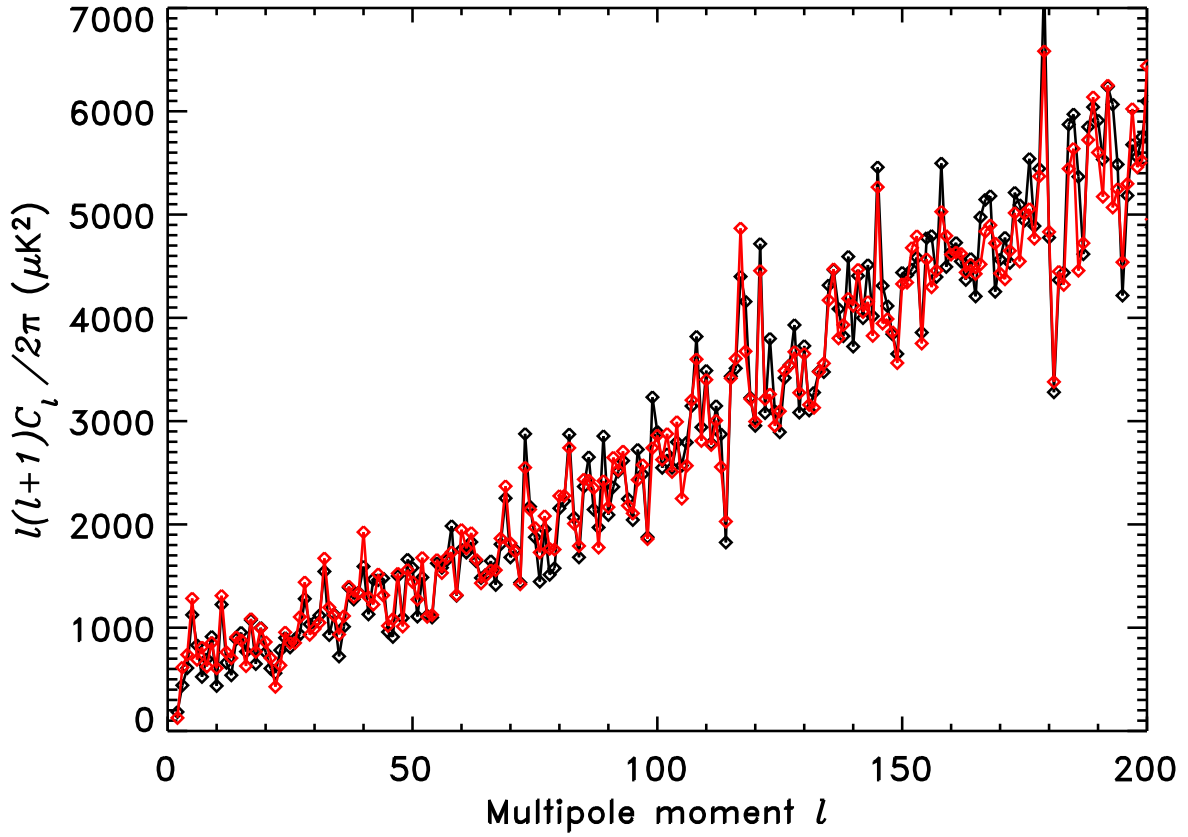


Fig. 12.— A comparison of two power spectra computed from a single V band map. The black points result from the maximum likelihood method, the red points from the quadratic estimator computed with uniform pixel weighting. In both cases a noise model has been assumed to treat the noise bias. At higher l , the two spectra would be nearly identical since both impose inverse noise weighting on the data.

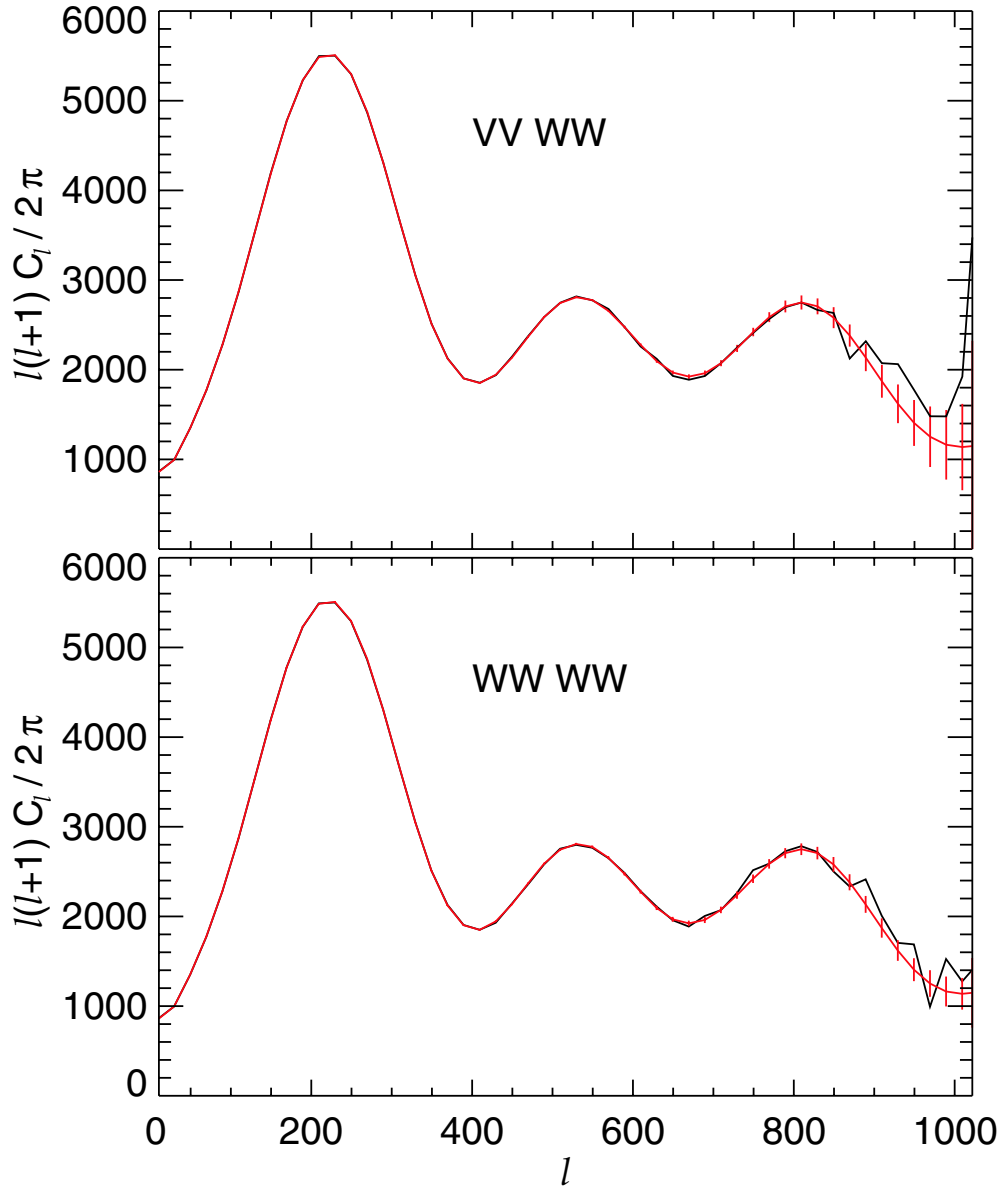


Fig. 13.— The quadratic estimator discussed in §A.1 has been extensively tested with Monte Carlo simulations of first-year *WMAP* data. The top panel shows a model spectrum in red, and the mean of 600 realizations of the cross-power spectrum, computed from the V and W band-averaged maps, in black. The red error bars show the uncertainty in the mean in bins with $\Delta l = 20$. The cross-power spectrum estimator is unbiased. The bottom panel is similar, but with auto-power spectra computed from the W band-averaged map. Here the noise bias term was estimated directly from the high l tail of the computed spectrum. This simple auto-power spectrum estimator is unbiased if the noise is white, and is suitable for V band data, but see Figure 1 for examples of W band noise spectra.

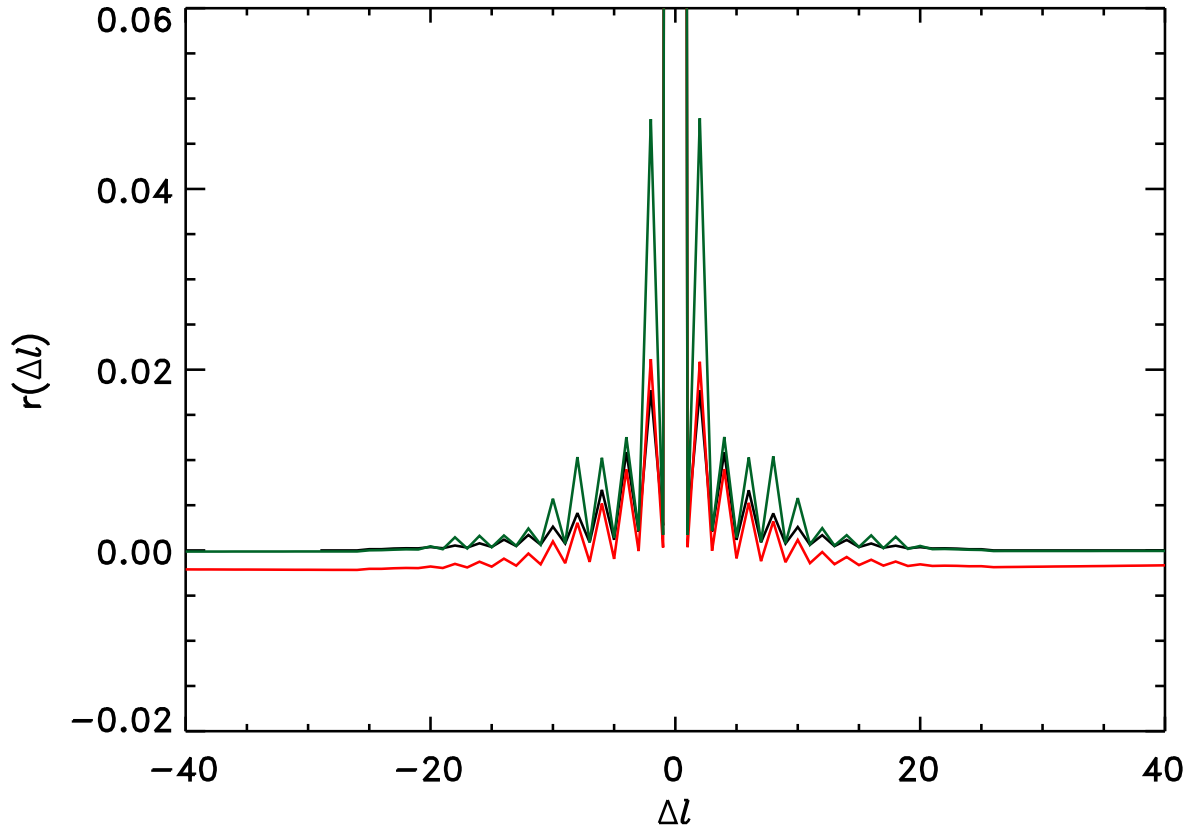


Fig. 14.— Slices of the Fisher (or curvature) matrix normalized as $F_{ll'}/\sqrt{F_{ll}F_{l'l'}}$, plotted vs. $\Delta l = l - l'$, from the Fisher matrix for the combined spectrum. Black is $l = 30$, red is $l = 300$, green is $l = 600$.

Table 1. WMAP Power Spectrum Noise Model^a

Parameter	Q1	Q2	V1	V2	W1	W2	W3	W4
$l < 200$								
$c_0(\times 10^2)$	-6.604	-6.971	-6.823	-6.674	-7.174	-7.439	-7.044	-7.989
$c_1(\times 10^3)$	1.934	4.244	2.497	1.395	3.060	3.721	1.459	4.820
$c_2(\times 10^4)$	-3.798	-8.521	-5.203	-1.483	-6.021	-6.901	-2.677	-5.870
$c_3(\times 10^5)$	2.549	5.921	3.715	0.242	4.147	4.540	1.750	2.167
$200 < l < 450$								
$c_0(\times 10^2)$	-6.143	-6.268	-6.320	-6.192	-6.618	-6.833	-6.768	-6.849
$c_1(\times 10^4)$	-1.334	1.341	-0.800	-1.351	0.769	2.582	0.885	2.443
$l > 450$								
$c_0(\times 10^2)$	-6.202	-6.167	-6.338	-6.264	-6.564	-6.704	-6.688	-6.772
$c_1(\times 10^5)$	3.288	4.085	3.327	5.979	7.342	12.970	8.851	24.430

^aBest-fit coefficients for the noise model in equation (8). The units of the output noise are mK², thermodynamic temperature.