SECOND AND THIRD SEASON QUaD COSMIC MICROWAVE BACKGROUND TEMPERATURE AND POLARIZATION POWER SPECTRA

C. Pryke[1], P. Ade[2], J. Bock[3,4], M. Bowden[2,5], M. L. Brown[6,9], G. Cahill[7], P. G. Castro[6,10], S. Church[5], T. Culverhouse[1],
R. Friedman[1], K. Ganga[8], W. K. Gear[2], S. Gupta[2], J. Hinderks[5,11], J. Kovac[4], A. E. Lange[4], E. Leitch[3,4], S. J. Melhuish[2,12],
Y. Memari[6], J. A. Murphy[7], A. Orlando[2,4], R. Schwarz[1], C. O' Sullivan[7], L. Piccirillo[2,12], N. Rajguru[2,13],
B. Rusholme[5,14], A. N. Taylor[6], K. L. Thompson[5], A. H. Turner[2], E. Y. S. Wu[5], AND M. Zemcov[2,3,4] (QUAD COLLABORATION)
[1] Kavli Institute for Cosmological Physics, Department of Astronomy & Astrophysics, Enrico Fermi Institute, University of Chicago, 5640 South Ellis Avenue,
Chicago, IL 60637, USA
[2] School of Physics and Astronomy, Cardiff University, Queen's Buildings, The Parade, Cardiff CF24 3AA, UK
[3] Jet Propulsion Laboratory, 4800 Oak Grove Dr., Pasadena, CA 91109, USA
[4] California Institute of Technology, Pasadena, CA 91125, USA
[5] Kavli Institute for Particle Astrophysics and Cosmology and Department of Physics, Stanford University, 382 Via Pueblo Mall, Stanford, CA 94305, USA
[6] Institute for Astronomy, University of Edinburgh, Royal Observatory, Blackford Hill, Edinburgh EH9 3HJ, UK
[7] Department of Experimental Physics, National University of Ireland Maynooth, Maynooth, Co. Kildare, Ireland
[8] Laboratoire APC/CNRS, Bâtiment Condorcet, 10, rue Alice Domon et Léonie Duquet, 75205 Paris Cedex 13, France

ABSTRACT

We report results from the second and third seasons of observation with the QUaD experiment. Angular power spectra of the cosmic microwave background are derived for both temperature and polarization at both 100 GHz and 150 GHz, and as cross-frequency spectra. All spectra are subjected to an extensive set of jackknife tests to probe for possible systematic contamination. For the implemented data cuts and processing technique such contamination is undetectable. We analyze the difference map formed between the 100 and 150 GHz bands and find no evidence of foreground contamination in polarization. The spectra are then combined to form a single set of results which are shown to be consistent with the prevailing LCDM model. The sensitivity of the polarization results is considerably better than that of any previous experiment—for the first time multiple acoustic peaks are detected in the *E*-mode power spectrum at high significance.

Key words: cosmic microwave background – cosmology: observations – polarization

1. INTRODUCTION

The anisotropy of the cosmic microwave background (CMB) gives us direct insight into the structure of the universe when it was a tiny fraction of its current age, and is one of the central pillars of the enormously successful standard cosmological model. The temperature anisotropy power spectrum has now been measured to good precision from the largest angular scales down to a small fraction of a degree (e.g., Reichardt et al. 2009) the expected series of acoustic peaks is present and fitting the spectrum yields tight constraints on the basic parameters of the cosmological model (e.g., Dunkley et al. 2009).

The CMB is expected to be polarized at the ~10% level principally because of motions in the material at the time of last scattering. Since the plasma flows along gradients in the density field the resulting observable polarization pattern has gradients (*E*-modes), but zero curl (*B*-modes) (e.g., Hu & White 1997). Given a standard cosmological model fit to the temperature spectrum (*TT*), the *E*-mode spectrum (*EE*), and temperature–*E*-mode cross-spectrum (*TE*), are nearly deterministically predicted—only at the largest angular scales is there additional information. It is important to remember that, although very successful, the standard cosmological model (which we will refer to throughout as LCDM) contains several components for which we have only circumstantial evidence (dark matter and dark energy). Measuring the *EE* and *TE* spectra is thus a crucial test of the overall theoretical paradigm.

As the CMB travels to us through the developing large-scale structure subtle deflections due to gravitational lensing occur (e.g., Hu 2003). This converts some fraction of the *E*-mode pattern into the so-called lensing *B*-modes—this effect is most important at smaller angular scales. In addition, if the cosmogenic theory known as inflation is correct, there must also be large angular scale *B*-modes caused by gravity waves propagating through the primordial plasma (e.g., Seljak & Zaldarriaga 1997).

The polarization of the CMB was first detected by the Degree Angular Scale Interferometer (DASI) experiment (Kovac et al. 2002), and since then several experiments have reported measurements of the *EE* and *TE* spectra (Barkats et al. 2005; Readhead et al. 2004; Montroy et al. 2006; Sievers et al. 2007; Page et al. 2007; Bischoff et al. 2008; Nolta et al. 2009). Thus far all reported measurements are consistent with LCDM, although precision remains limited. *B*-mode polarization has not yet been detected—all results so far are upper limits. We previously reported preliminary results from QUaD in Ade et al. (2008). In this paper, we report considerably improved results using data from the second and third seasons of observations.

---

[9] Current address: Cavendish Laboratory, University of Cambridge, J.J. Thomson Avenue, Cambridge CB3 OHE, UK.
[10] Current address: CENTRA, Departamento de Física, Edifício Ciência, Piso 4, Instituto Superior Técnico (IST), Universidade Técnica de Lisboa, Av. Rovisco Pais 1, 1049-001 Lisboa, Portugal.
[11] Current address: NASA Goddard Space Flight Center, 8800 Greenbelt Road, Greenbelt, Maryland 20771, USA.
[12] Current address: School of Physics and Astronomy, University of Manchester, Manchester M13 9PL, UK.
[13] Current address: Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK.
[14] Current address: Infrared Processing and Analysis Center, California Institute of Technology, Pasadena, CA 91125, USA.
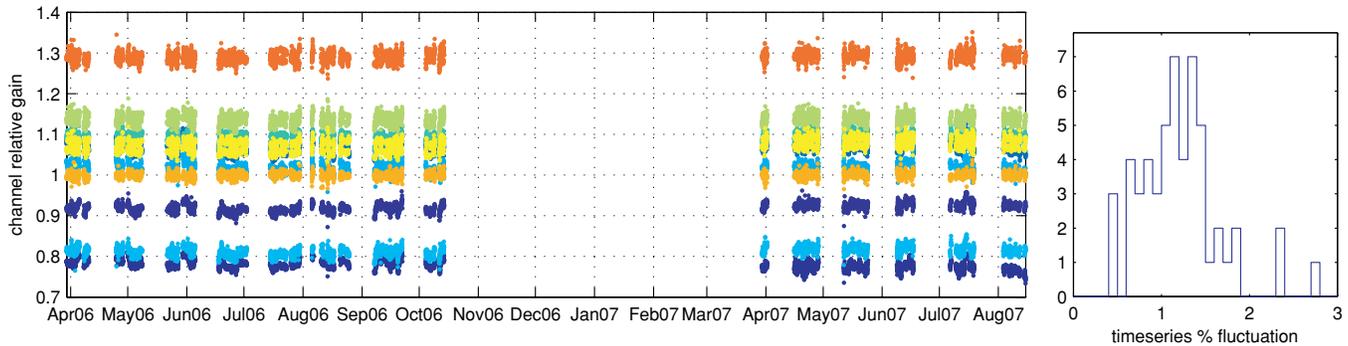
Figure 1. Left: the relative gains of the first ten 150 GHz detector channels as measured by elevation nods every half hour over two seasons. Right: a histogram of the percentage fluctuations of the timeseries at left including all channels.

This paper is structured as follows: in Section 2 we briefly review the instrument and detail the observations, Section 3 describes the low level data processing and calibration, Section 4 outlines the steps used to make timestream data into maps, Section 5 describes the simulation methodology, Section 6 converts the maps into power spectra, Section 7 gives the results of jackknife tests, Section 8 describes foreground studies, Section 9 gives the final combined power spectrum results, Section 10 contains some further investigations of systematic effects, and Section 11 states our conclusions.

2. INSTRUMENT AND OBSERVATIONS

The design, implementation, and performance of the QUaD experiment are described in detail in a companion paper (Hinderks et al. 2009) hereafter referred to as the "Instrument Paper"—only a very brief summary will be given here. QUaD was a 2.6 m Cassegrain radio telescope on the mount originally constructed for the DASI experiment (Leitch et al. 2002). This is an az/el mount with a third axis allowing the entire optics and receiver to be rotated around the line of sight (referred to as "deck" rotation). The QUaD receiver consisted of 31 pairs of polarization-sensitive bolometers (PSBs; Jones et al. 2003), 12 at 100 GHz and 19 at 150 GHz. The detector pairs were arranged in two orientation angle groups separated by 45°. The mount is enclosed in a large bowl-shaped reflective ground shield on top of a tower approximately 1 km from the geographic South Pole. The bolometers were read out using AC bias electronics, and digitized by a 100 Hz, 16 bit ADC. The raw data were staged on disk at Pole and transferred out daily via satellite. QUaD was decommissioned in late 2007.

The observations reported on in this paper were made during the Austral winter seasons of 2006 and 2007—the QUaD telescope was not able to observe during the summer due to contamination from the Sun. Complete CMB observation runs occurred on 171 days of 2006 and 118 days of 2007 (defined as an uninterrupted run with the Sun below the horizon). Of these available 289 days 44 were rejected after initial low level processing—mostly due to very bad weather, with a few due to instrumental problems. A further 43 days show obvious signs of contamination by the Moon, and 59 more fail a very conservative Moon proximity cut—Moon contamination is discussed in the Instrument Paper and Section 10.7. This leaves us with a total of 143 days of data which are used in the current analysis. For simplicity the cut granularity is very coarse—we only consider complete days, and if there is anything wrong with any part of a day we cut the entire day. It would certainly be possible to include somewhat more data with additional work. In Figure 1,

we can see the resulting set of days used—the monthly gaps are due to the rising and setting of the Moon.

The QUaD telescope is mounted on a tower at one end of the MAPO observatory building. At Pole the celestial sphere rotates about the zenith every 24 hr. Therefore to minimize the potential for contamination from the building, or the heat plume from the furnace it contains, each day of observation starts at a fixed local sidereal time (LST) such that our chosen CMB field (centered on R.A. 5.5h, decl. $-50°$) has just cleared the laboratory building. The observations are split into two blocks, each of 8 hr, with special calibration observations before and after each block. Between the two blocks the entire telescope is rotated by 60° around the line of sight, and then approximately 30 minutes are allowed for thermal stabilization. The total observation schedule takes about 19 hr, with the remainder of the 24 hr period being taken up by fridge cycling.

Each 8 hr block of CMB observations is divided into 16 half hour periods. Each starts with an observation of the internal calibration source followed by an "elevation nod"— the telescope is moved up and then down again by one degree in elevation injecting a large signal into the data stream due to the atmospheric gradient (see the Instrument Paper for details). The telescope is then scanned back and forth five times over a $7°.5$ throw in azimuth, with the scan being applied as a modulation on top of sidereal tracking. The scan rate is $0°.25$ per second in azimuth translating to around $0°.16$ per second on the sky at our observing elevation. The pointing position is then stepped by $0°.02$ in declination and the process repeated four times. Including time for moves and settling, these calibration observations, plus the 40 "half-scans" of 30 s each, take half an hour.

During the first half hour of each pair, the pointing center is R.A. 5.25h. This is then set to R.A. 5.75h and the declination offsets repeated during the second half hour. The observations are thus taken in a lead–trail manner where the scanning pattern is identical in ground fixed azimuth-elevation coordinates between the two half hours of each pair. By subtracting the lead and trail data one therefore cancels any signal coming from the ground which is constant over the half hour, while producing a difference map of the sky. We do see significant ground pickup and lead–trail field differencing is used throughout the analysis presented in this paper (see Section 3.5).

From hour to hour the declination offsets are cumulative, but they are reset before the second 8 hr block. The telescope therefore scans a "letter box" region $0°.64$ high in decl. twice per day, at two different line-of-sight rotation angles. Each day the observation region steps by $0°.64$ in decl. to cover the entire field as rapidly as possibly (in ≈ 9 days), and then cycles around
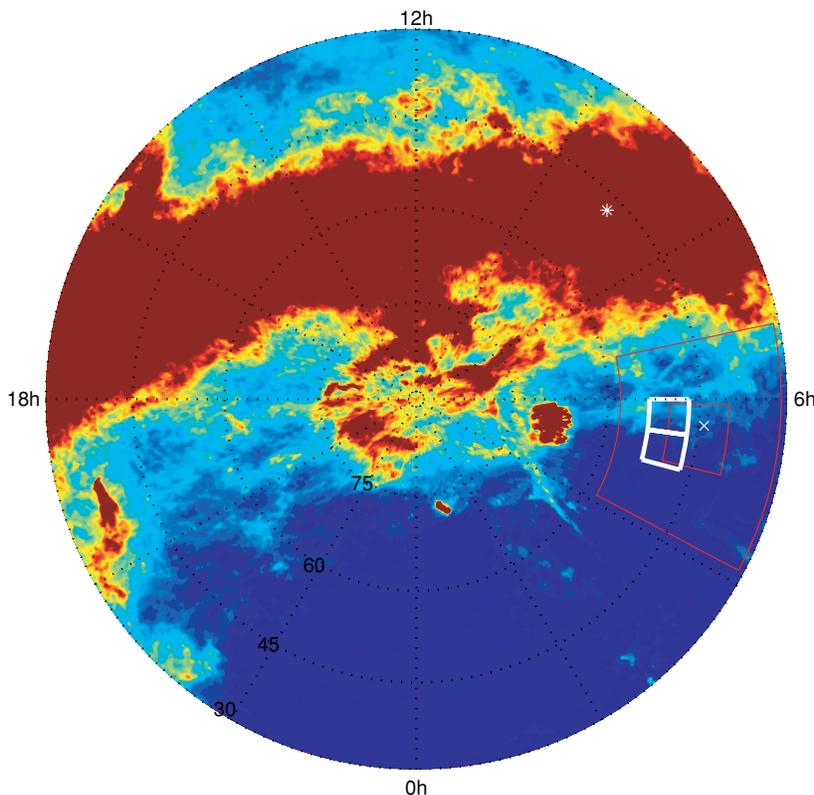
Figure 2. Location of the QUaD lead–trail fields delineated in white on an equal area azimuthal projection about the SCP. The color map is the prediction of dust emission intensity at 150 GHz from FDS model 8 (Finkbeiner et al. 1999). The color scale is linear from 0 to 100 μK, and is heavily saturated. (The B03 deep and shallow regions are delineated in red, while the white asterisk and cross show the locations of RCW38 and PKS0537–441 respectively.)

with a $0°16$ offset to generate an eventual four-fold interleaved coverage pattern (after 37 days). This observation pattern was repeated throughout the 2006 and 2007 seasons a total of ≈ 8.5 times. Figure 2 shows the location of the QUaD field.

Note that only azimuth scanning is used. Since the telescope is only ~ 1 km from the Earth's rotational axis essentially zero "cross-linking" of the map occurs. For the multipole range presented in this paper this has only a small negative impact on the final CMB power spectrum results.

The bright galactic H ɪɪ region RCW38 was observed on 11 days distributed through 2006 and 2007, to monitor the beam offset angles and shapes, and the bright quasar PKS0537–441 was observed during the 2007 season to further study the beam shapes. These observations are used in the analysis below (see Sections 4.2 and 5.1.1).

We also observed several other discrete sources (Cen A, Galactic center, Moon), as well as conducted a survey of part of the galactic plane. These observations will be described in future papers.

3. LOW LEVEL DATA PROCESSING

This section describes the low level data processing which occurs before the timestream is binned into maps (Section 4) and analyzed to generate simulations (Section 5). The data are deconvolved to remove the effects of detector time constants, relative gain calibrated, and field differenced.

3.1. Deconvolution of Detector Temporal Response

The initial stages of the data reduction are performed on each day of data individually. The first step is deconvolution of the detector time constants—the bolometers have noninstantaneous response to changes in incident optical power, and hence the output data timestream is a low-pass filtered version of the desired input waveform. With knowledge of the detector time constants this filtering can be undone at the price of increased high-frequency noise.

We measured the temporal response of our detectors in situ using an external Gunn oscillator source as described in the Instrument Paper. We find that many of the detectors are well fitted by a simple single time-constant model. The median primary time constant is 16 ms. However, a substantial fraction ($\sim 50\%$) require a second additive time constant to obtain a good fit—we are hypothesizing that some fraction of the incident heat goes into a second reservoir which has a weaker coupling to the thermal bath. In a couple of cases, the second time constant is several seconds long but the two time-constant model is still a good fit—we retain these detectors. In two other cases, the dual time-constant model is not a good fit and we reject these detectors leading to the loss of two channel pairs (both at 100 GHz).

To check that the deconvolution process is working, and that the time constants are stable over time, we examine forward–backward jackknife maps of the bright compact source RCW38 taken on 11 days distributed through 2006 and 2007. These show no detectable residuals and hence cancellation to $\ll 1\%$.

As part of the deconvolution process, the timestream is also low-pass filtered to < 5 Hz. After deconvolution, we de-glitch to remove cosmic ray hits etc. (losing $\ll 1\%$ of the data).

For this analysis we also exclude two additional channel pairs: one (at 150 GHz) due to a time evolving scan synchronous signal and a second (at 100 GHz) due to strongly atypical pair
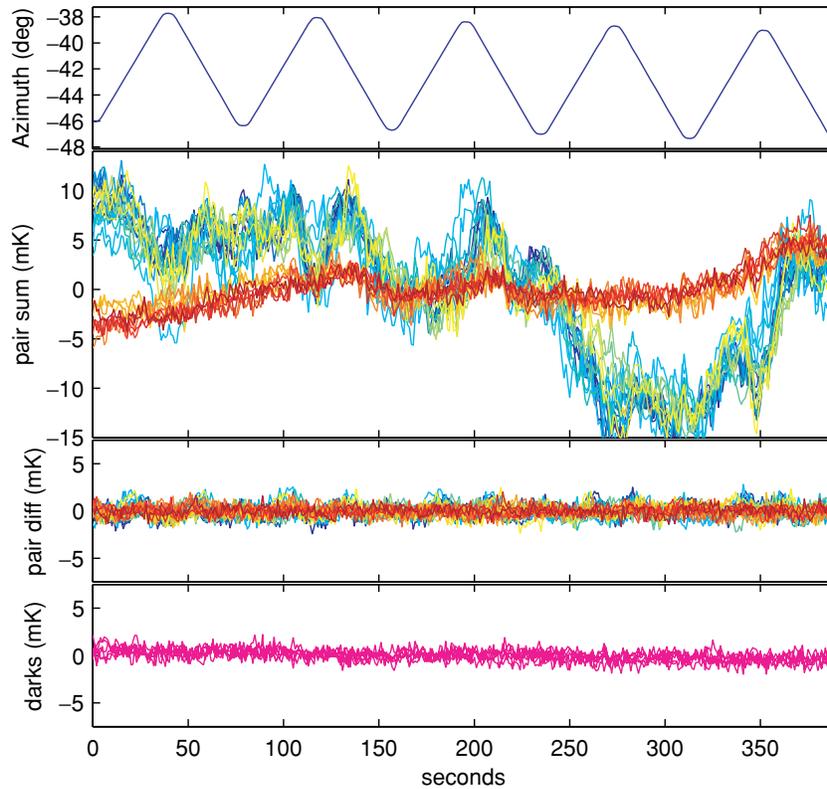
Figure 3. Timestream data for a sample scan set. The top panel shows the azimuth angle, and the middle panels the pair sum and difference detector timestreams with relative gain calibration applied; red/orange colors are the 100 GHz pairs, while blue/green colors are the 150 GHz. The bottom panel shows the dark channels. In all cases an approximate scaling to temperature units has been applied. For the purposes of this illustration the timestreams have been heavily low-pass filtered (to $\leqslant$ 1 Hz).

differenced noise. We are thus left with nine of the possible 12 pairs at 100 GHz, and 18 of the 19 pairs at 150 GHz.

3.2. Relative Gain Calibration via Elevation Nods

We measure the relative gains of the detector channels using the elevation nod method mentioned in Section 2. The air mass through which each channel pair was looking is calculated from the elevation encoder reading and then regressed against the observed signal to yield a calibration factor in volts per airmass. We then simply normalize each channel's gain to the mean of all channels:

$$V_i'(t) = V_i(t)\frac{\overline{g}}{g_i}, \qquad (1)$$

where the subscript i is a loop over channels, $V(t)$ are the detector timeseries data, and g are the elevation nod gain factors. This equalizes the gains both within and between channel pairs.

The nominal accuracy of each elevation nod gain measurement is $\ll$ 1%. Weak trends in the apparent relative gains are observed over short and long timescales, the cause of which is unknown. However, as we see in Figure 1 this leads to fluctuation over the entire two season time span of only $\approx$ 1% root mean square (rms). In this analysis, we choose to regard these variations as real and the relative gains derived from each elevation nod are applied to the subsequent 40 half-scans.

However even if these apparent variations are false, random errors of this magnitude will cause negligible leakage of total intensity to polarization (hereafter T to pol. leakage) as, for example, T will leak sometimes into $+Q$ and sometimes into $-Q$, averaging down in the final maps over both time and detector pairs.

Possible systematic errors in the relative gains are of much greater concern and are discussed and simulated further in Section 10.3. Imperfect deconvolution would lead to the relative gain of a detector pair being a function of temporal frequency. Since the elevation nods measure the gain at an effective frequency well below the CMB measurement band this would result in systematic T to pol. leakage. Using the Gunn oscillator derived time constants and deconvolution procedure described in Section 3.1, the lack of "monopole" residuals in pair difference jackknife maps of the bright source RCW38 indicates that the low-frequency elevation nod gains are accurate to better than 1% at high (beam scale) frequencies.

3.3. Examination of Timestream and Noise Spectra

The sum of the signals from a PSB pair measures the total intensity of the incident radiation, while the difference measures polarization. Having performed the relative gain calibration and sum/difference operations, Figure 3 shows a typical scan set at this stage of the processing. Atmospheric emission at 100 and 150 GHz is dominated by oxygen and water vapor, respectively. The intensity of the radiation received is an integral over the temperature and density along the atmospheric column through which the telescope is looking. Water vapor is poorly mixed in the atmosphere leading to much greater fluctuations at 150 GHz. From the lack of an obvious scan synchronous component, we can infer that the wind blows inhomogeneities through the beam faster than the telescope scans. Since the QUaD beams do not diverge until high in the atmosphere the pair sum timestreams are highly correlated across the array. As expected, the atmospheric
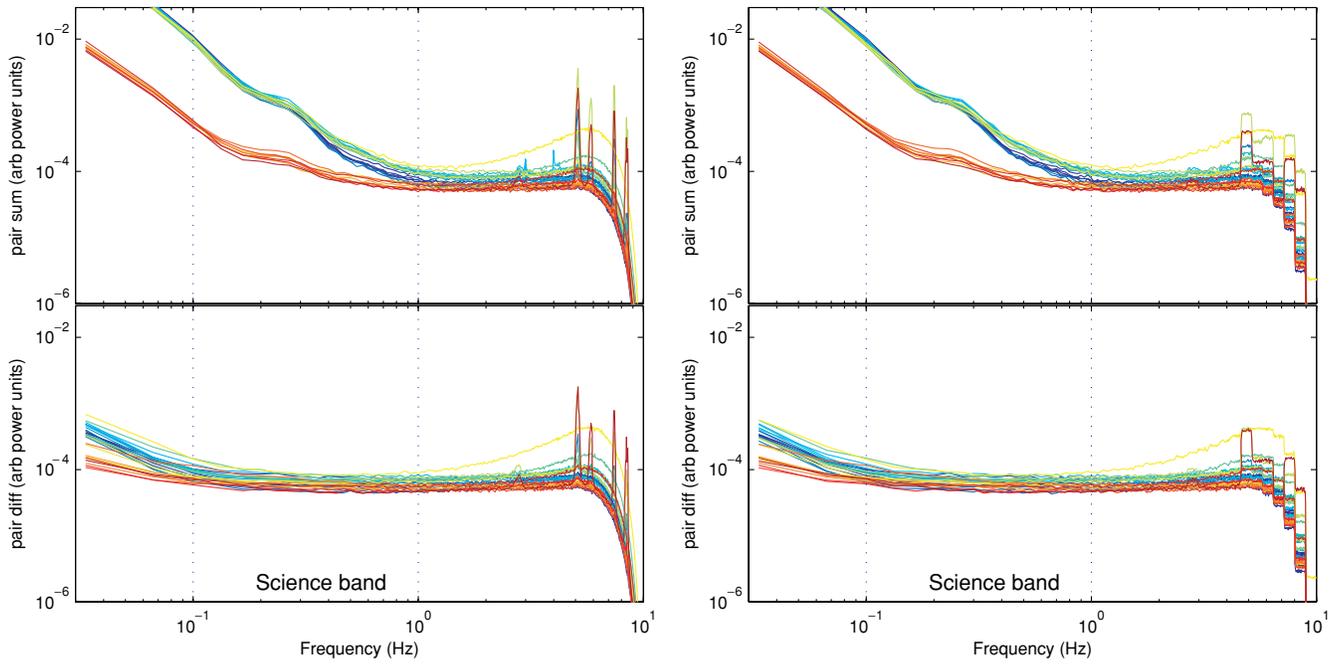
Figure 4. Left: timestream PSDs averaged over a single day of data. Pair sum and pair difference spectra are shown; red/orange colors are the 100 GHz pairs, while blue/green colors are the 150 GHz. The range of frequencies corresponding to $200 < \ell < 2000$ is enclosed within the dotted lines. See the text for further details. Right: the same thing for simulated noise timestream (see Section 5.2).
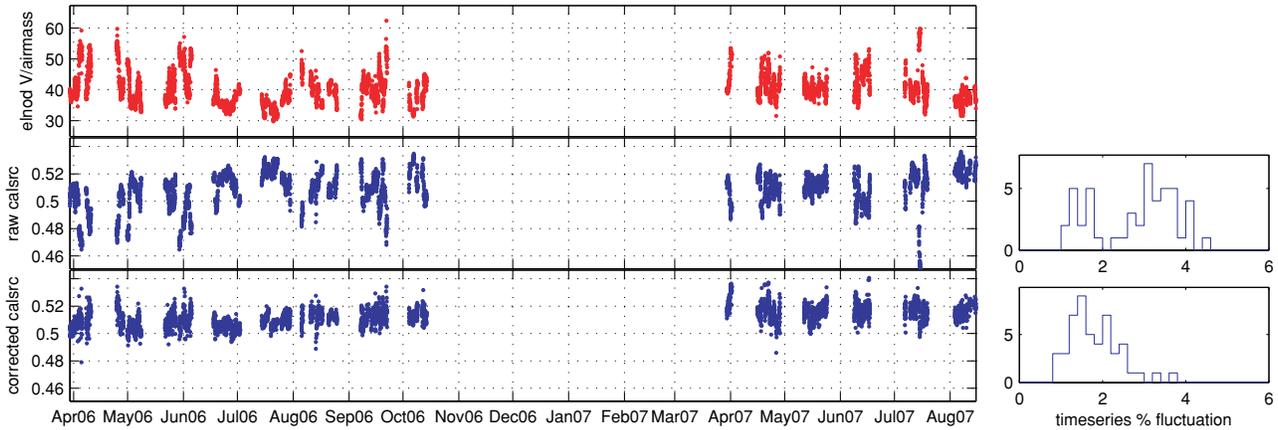
Figure 5. Top: the atmospheric emission measured by elevation nods in volts per airmass for a sample detector. Middle: the sinusoidal modulation amplitude of the calibration source as observed by the same channel. Bottom: the calibration source amplitude after application of the loading gain suppression correction (see the text). Histograms over all channels of the percentage fluctuation in the pre- and post-corrected timeseries are shown to the right.

emission is clearly very weakly polarized leading to strong cancellation in the pair difference timestreams.

Figure 4 shows the mean power spectral densities (PSDs) of the detector timestreams. To make this plot spectra were computed for each channel and half-scan and then averaged over a day of data. We see substantial $1/f$ noise in the pair sum data. The $1/f$ knee shifts up in bad weather and down in good—the plot is for a day of intermediate quality. For our scanning speed ($0°.25\,\mathrm{s}^{-1}$ in azimuth) and observing elevation ($\sim 50°$) the conversion from timestream frequency f to multipole on the sky is $\ell \sim 2000\,f$, giving a "science band" of $0.1 < f < 1$ Hz. Even on the worst days used in this analysis, the pair difference spectra remain close to white within this range. See the Instrument Paper for further details of the sensitivity. Note the lack of narrow line noise within the science band—some channels show microphonic lines at much higher frequencies. The roll-off above 5 Hz is imposed in the initial processing as part of the

deconvolution procedure. Also note the uniform noise properties of the detectors within each frequency group. (The "roll up" toward higher frequencies observed in some pairs is due to the deconvolution of exceptionally long time-constant detectors.)

3.4. Long Term Gain Equalization via Calibration Source

The QUaD telescope was equipped with a battery powered, remote-controlled calibration source mounted behind the secondary mirror, inside the foam cone. Before each half hour of observations the $45°$ flip mirror was commanded down and the polarizer grid rotated several times injecting a sinusoidally modulated signal into the detector timestreams—see the Instrument Paper for details. The low level analysis measures the modulation amplitude for each channel. Figure 5 shows the time series for one channel over the whole season and also the volts per airmass as measured by the elevation nods. There is a clear

anti-correlation—as the atmosphere becomes more opaque the atmospheric loading goes up, suppressing the detector gains. (The source temperature is monitored and shows no correlation with the external temperature.) For the good weather data used in this analysis the suppression is $\leqslant 10\%$. We make a linear regression of the mean calibration source timeseries against the mean elevation nod timeseries within each frequency band and use this to apply a correction. After application of this correction we believe the absolute gain of the QUaD system to be stable at the few percent level over the entire 2006 and 2007 seasons.

3.5. Field Differencing

If we do not field difference ground pickup produces obvious artifacts in the final co-added maps. In the timestream this pickup is normally not visible above the detector and atmospheric noise as we see in Figure 3. However on the very worst days, the very worst pairs can show pickup equivalent to as much as 5 mK CMB across a half-scan in the pair differenced data. We are confident that this is ground signal because it is not present in the dark channels, is fixed in azimuth angle, is worst in the bottom row pixels, and correlates with the amount of snow present on the ground shield—see the Instrument Paper for further details.

We note that the ground signal does not show structure smaller than a few degrees which is expected since it is near field pickup. It may be possible to remove the ground signal sufficiently well using a template-based approach, and we are investigating this option. However, for the analysis presented here we have applied simple lead–trail differencing throughout. In doing so we make the assumption that the ground signal, which clearly does change from day to day, is stable to the relevant level of accuracy over the half-hour timescale. In the field differenced maps no artifacts are visible. Ultimately, the jackknife tests presented in Section 7 are the most sensitive test for residual ground contamination.

The start/end points of each half-scan are carefully tuned to give a best fit to the ideal linear scan motion. We then apply field differencing point-by-point in the timestream by subtracting from each half-scan its partner occurring half an hour later. The delta R.A. between the differenced points has a standard deviation of $0°007$.

Note that for a Gaussian random field, such as the CMB, the power spectrum of a difference field (in the limit of no correlations) is twice that of an undifferenced field, with an associated increase in the sample variance due to the reduced effective sky area. The QUaD field differencing is explicitly modeled in our analysis pipeline.

4. FROM TIMESTREAM TO MAPS

To make polarization maps three more ingredients are required—knowledge of the overall pointing of the telescope, of the relative pointing of each of the detectors, and of the polarization angles and efficiencies of the detectors.

4.1. Overall Telescope Pointing

The QUaD mount used a nine-parameter online pointing model derived from optical and radio observations as described in the Instrument Paper. During special radio pointing runs this was shown to have an absolute accuracy over the hemisphere of $\sim 0'.5$ rms.

In addition, a pointing check was performed on RCW38 before and after each 8 hr block of CMB observations. These also indicate a $\sim 0'.5$ rms wander in the absolute pointing. Attempts were made to use the measured offsets to make an

offline pointing correction but it was not possible to demonstrate any clear improvement, although a few days were rejected due to abnormally large pointing errors. Note that the effect of a pointing wander of the observed magnitude is included in our simulations below.

4.2. Detector Offset Angles

On 11 days throughout the 2006 and 2007 seasons full day observations were conducted of RCW38 at three line-of-sight orientations ("deck" angles). For each day, the data from each channel were fitted to the six-parameter model of an elliptical Gaussian beam with free centroid positions, orientation angle, and widths. We did the fits both in the timestream and in maps, yielding equivalent results. The scatter in the centroid positions for a given detector over the set of days is $\sim 0'.5$ consistent with the overall pointing wander of $\sim 0'.5$ rms discussed above. There is hence no evidence for systematic changes in the detector offset angles over time. We therefore take the detector offset angles as the mean over the set of observed values in the RCW38 runs and estimate their uncertainty as a negligible $\sim 0'.15$.

4.3. Detector Polarization Angle and Efficiency

Our best measurements of the polarization angles of the detectors come from in situ measurements of an external source as described in the Instrument Paper. A chopped thermal source was placed behind a polarizing grid and observed with the telescope at many rotation angles, the signal from each PSB tracing out a sinusoid. The phase of these sinusoids gives the detector polarization angle, the fitted values agreeing with the design values with a scatter of around one degree rms. Since this is compatible with the estimated measurement uncertainty we have used the design values in this analysis when constructing maps.

The degree to which the sinusoid mentioned above fails to reach zero represents the response of a measurement channel to anti-aligned radiation. This ratio of minimum-to-maximum response is conventionally referred to as the cross polar leakage ϵ. Our measured values of ϵ have a mean of 0.08 with an rms scatter of 0.015. In this analysis we have assumed the mean value to apply to all channels when constructing maps, and included scatter in the simulations (see Section 5.1). Note that for an experiment of this type cross polar leakage does not imply T to pol. leakage—it is simply a small loss of efficiency, which must be corrected by an additional calibration factor applied only to the pair difference data. The effect of systematic error on ϵ is discussed in Section 10.6.

4.4. Map Making

To make maps we perform the following operations: sum and difference the detector timestreams for each pair, remove a third-order polynomial across each 30 s half-scan, and bin into a grid of pixels weighting by the inverse variance of each half-scan. In this analysis the pixelization is in R.A. and decl., and the pixels are $0°02$ square. The polynomial subtraction removes the bulk of the atmospheric $1/f$ noise allowing the simplicity of "naive" map making without incurring a large noise penalty.

For the pair sum data only the signal times the weights, and the weights themselves, must be accumulated to form the temperature map. For the pair difference, the product of the data and the sine and cosine of the detector angle as projected on the sky are accumulated. Then a 2×2 matrix inversion is performed for each pixel to produce Q and U maps. We emphasize that for
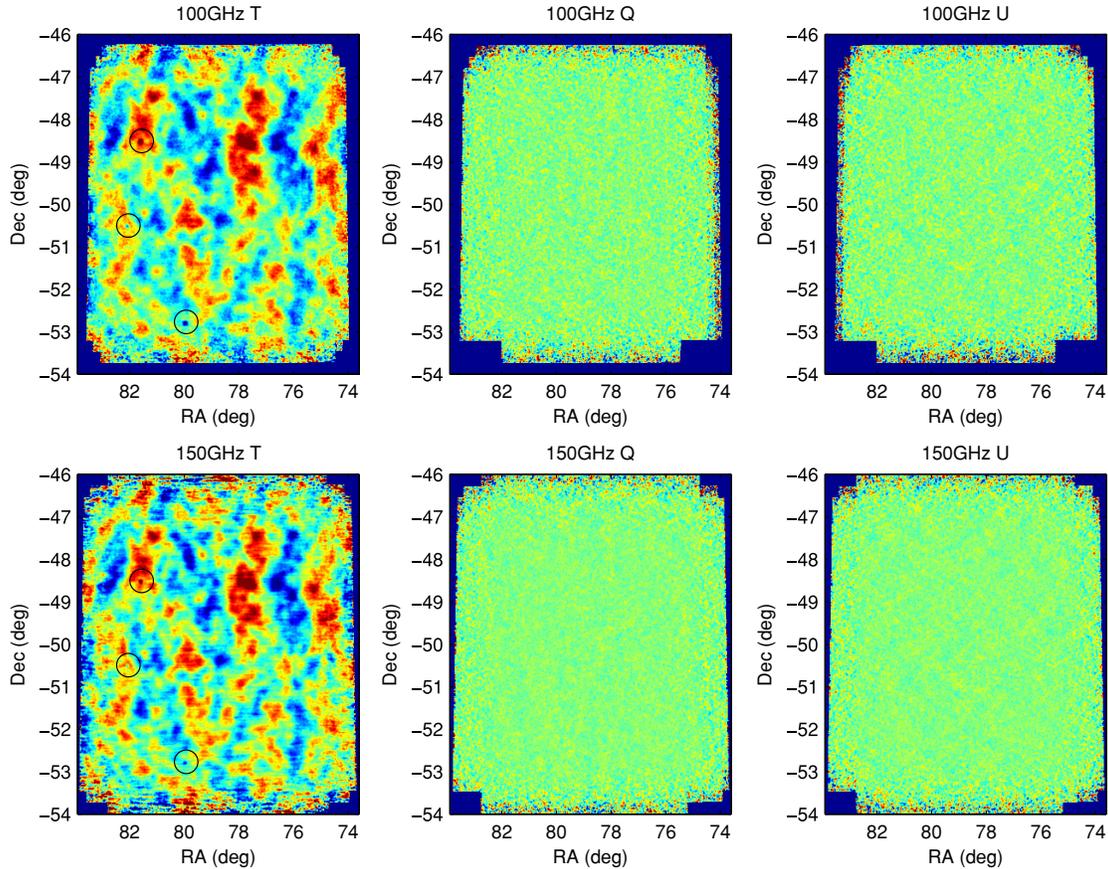
Figure 6. QUaD 100 and 150 GHz T, Q, and U field difference maps. The color scale is $\pm 200~\mu$K in all cases, and the maps have been smoothed with a 5′ FWHM Gaussian kernel. (The black circles indicate discrete sources which are removed in the power spectrum analysis.)

this matrix to be invertible any given pixel needs to have been measured at only two distinct grid angles—in the presence of noise, angles separated by 45° are optimal. The only gain from having a more uniform distribution of angles is in averaging down a limited number of systematic effects (such as beam size or pointing mismatch). In practice since we have detector pairs of two orientation angles, and observe at two deck angles, pixels in the central region of our maps have been measured at four angles.

4.5. Absolute Calibration

At this point, we have T, Q, and U maps at each frequency in raw detector units (volts). To scale these into μK, we perform a correlation analysis versus two noise independent temperature maps from the 2003 flight of the Boomerang experiment (B03)—which in turn have been calibrated against the *Wilkinson Microwave Anisotropy Probe* (*WMAP*) with a 2% stated uncertainty (Masi et al. 2006). The B03 maps are first passed through the QUaD simulation pipeline (described in Section 5.1). The raw QUaD and "B03 as seen by QUaD" maps are then both apodized by the QUaD sensitivity mask (see Section 6.1), Fourier transformed, and cross-spectra taken between the QUaD map and one B03 map, and between the two B03 maps. For each bandpower b, we then calculate our absolute calibration factor as

$$a_b = \frac{w_{Q,b}}{w_{C,b}} \frac{\langle m_R m_C \rangle}{\langle m_R m_Q \rangle}, \qquad (2)$$

where w is the Fourier transform of the beam, m are the modes of the Fourier transform of the apodized map, and R, C, and Q refer to the "reference," "calibration," and QUaD maps, respectively, and the mean is taken over the modes in a given annulus of the Fourier plane.

This calibration factor should ideally be a constant value for each bandpower (multipole range). In practice, we find that it is for $200 < \ell < 800$ where the B03 beam correction (and therefore uncertainty on that correction) is modest, and we therefore take the average value across that range. We have also performed the same operation using various combinations of the *WMAP* Q, V, and W band maps. Due to the much larger beams of *WMAP*, the beam corrections are very large, and the signal-to-noise low. However based on these results, and the point-to-point scatter in the B03 analysis, we estimate a 5% uncertainty in our primary B03-derived absolute calibration factors.

4.6. Map Results

Figure 6 shows the 100 and 150 GHz T, Q, and U maps as generated using the process described in Section 4.4. The signal to noise in the T maps is extremely high as is evident from the excellent spatial correlation of the pattern between the two frequencies. We deliberately plot the Q and U maps on the same color scale to visually emphasize how small the polarization of the CMB is compared to the degree scale structure in T. Three discrete sources are weakly detected in the temperature maps, but are not detected in polarization.

4.7. Jackknife Maps

The CMB polarization signal is extremely small necessitating extreme attention to detail. To probe for the presence of contaminating signal that does not originate on the sky, we perform a set of jackknife tests. For each test we split the timestream data into two approximately equal subsets which should contain (nearly) the same sky signal, but which might contain different false signal. We then generate maps for each data subset and difference them to produce jackknife maps:

$$M^J = (M^1 - M^2)/2 \qquad (3)$$
$$V^J = (V^1 + V^2)/4, \qquad (4)$$

where M^J is the jackknife map, M^1 and M^2 are the split maps, and V are the corresponding variance maps. Note that we form jackknife maps with this normalization by analogy to the nonjackknife map which is effectively $(M^1 + M^2)/2$.

Plots of the jackknife maps are not presented here as they are not particularly informative—they look like noise. However Section 6.3 includes plots of the Fourier transform of some sample jackknife maps, and Section 7 probes for any hint of structure in these maps which is inconsistent with noise.

5. GENERATION OF SIMULATED TIMESTREAM

Power spectra of maps such as those shown in Figure 6 need to be corrected for two effects: the "noise bias" (principally affecting auto spectra), and the suppression of power due to the effects of timestream filtering and finite beam size. In addition, we need to estimate the size of the final bandpower uncertainties due to sample and noise variance. We do this broadly following the MASTER technique (Hivon et al. 2002), which requires accurate timestream level simulations of signal and noise.

5.1. Signal Simulations

To construct simulations of signal timestream we start with LCDM power spectra generated using the CMBFAST program (Zaldarriaga & Seljak 2000) using the *WMAP3* cosmological parameters given under the heading "Three Year Mean" in Table 2 of Spergel et al. (2007), and feed these into the "synfast" generator (part of the HEALPix package[15]) to yield curved sky maps of T, Q, and U at a resolution of $0.'4$ (NSIDE of 8192).

We then read each day of data in turn, and loop over detectors. For each we calculate the sky map M_d which would be seen by a detector of the given angle and polarization efficiency by combining the T, Q, and U input maps according to

$$M_d = \frac{1}{2}\left(M_T + \frac{1-\epsilon}{1+\epsilon}\left(M_Q \cos 2\theta + M_U \sin 2\theta\right)\right), \qquad (5)$$

where ϵ is the cross polar leakage and θ is the detector polarization angle. This ideal sky map is then convolved with an elliptical Gaussian smoothing kernel with the appropriate parameters to simulate the effect of the beam. Finally we interpolate off the smoothed map along the pointing trajectory for the given detector, computed as the observed telescope pointing direction plus the detector offset angle.

In Section 4.3, we mentioned that fixed nominal values are used for the detector polarization angles and efficiencies when constructing maps. For each simulation realization, we generate values for each detector normally distributed about

these nominal values with the measured rms scatter ($1°$ and 0.015, respectively). The simulated data are then re-mapped assuming the nominal values as usual.

To simulate the small pointing wander mentioned in Section 4.1, we generate Gaussian random numbers with zero mean and σ of $0.'5$ as the R.A. and decl. offsets at the start and end of each 8 hr block. We then linearly interpolate these offsets to each time step and add them in to the observed pointing trajectory. Although in reality the wander has a more complex behavior given the observation strategy, and averaged over hundreds of days of data, this will lead to a broadening of the effective beam width of the correct (very small) amount.

The detector offset angles are taken as fixed at our best estimate values—we have no evidence for time variation as mentioned in Section 4.2. Note that the measured beam centroid positions show repeatable offsets between the two halves of each detector pair with an rms magnitude of $\sim 0.'1$. When we construct maps using pair sum and difference data, we use the mean for each pair. However, when we sample from the sky to generate simulated timestream we use the measured individual detector values. Hence, any T to pol. mixing which occurs due to beam centroid mismatch is fully included in the simulations—and found to be negligible—see Section 10.4.

Note that the generation of simulated signal timestreams takes place before the field differencing operation—effects such as actual signal correlations and mismatch of pointing coordinates between lead and trail fields will therefore be included.

5.1.1. Measurement of the Beam Widths

To convolve the ideal sky map with the beam we need to know the beam shape for each detector. We have obtained this from nine days of special observations of the bright quasar PKS0537–441, which is a point source at the angular resolution of our experiment (Fey & Charlot 2000). The mean of the measured major and minor FWHMs is $5.'5$ at 100 GHz and $3.'5$ at 150 GHz with uncertainties of $\approx 2\%$. There is evidence for a small degree of variation in width between the detectors of a given frequency band, and for $\leqslant 10\%$ elongation (mismatch between the major and minor widths), but for any given detector these are subdominant to the measurement uncertainty. As when determining the detector offset angles from measurements of RCW38 (see Section 4.2), we have performed the elliptical Gaussian fits in both the timestream and in maps yielding nearly identical beam widths and angles. For further details see the Instrument Paper and our Optics Paper (O'Sullivan et al. 2008).

The effective beam width in our CMB field co-added maps is measured with low signal to noise by the quasar PKS0524–485 (the brightest of the circled sources in Figure 6). This source appears as a Gaussian peak with width consistent with the single day observations of PKS0537–441.

5.2. Noise Simulations

As shown in Figures 3 and 4, the detector timestream is dominated by heavily correlated low-frequency atmospheric noise. The goal of the timestream noise generator is to reproduce simulated half-scans which are indistinguishable from the real under a battery of tests in both the time and frequency domains.

To achieve this we find it necessary to measure and re-generate longer pieces of timestream, and then cut them down to the half-scans which are actually used when constructing the maps. First, the complete time period spanning each five scan set is Fourier transformed (including turn arounds). Then for each of a set of logarithmically spaced frequency bins we

---

[15] See http://healpix.jpl.nasa.gov/index.shtml and Górski et al. (2005).

take the covariance matrix of the Fourier modes between all channels. This matrix is Cholesky decomposed and used to mix uncorrelated random numbers to re-generate the observed degree of covariance. This process is repeated for each frequency bin and then the resulting sets of Fourier modes are inverse transformed to yield simulated timestream. Possible correlations between the real and imaginary parts of each Fourier mode are preserved by using complex covariance and Cholesky matrices, but this process assumes that the Fourier modes are correlated only between channels, and not between modes. Breaking the simulated timestream down into half-scans reintroduces such correlations and we find it has equivalent temporal and spectral characteristics to the real data under a variety of statistical tests.

As an example the right part of Figure 4 shows the power spectra of the resulting simulated timestream—comparing to the real PSDs in the left part of the figure we see an excellent match. Note that the simulation produces non-pair differenced timestream which has then been differenced to generate the spectra shown in Figure 4—it is clear that the channel–channel correlations are being reproduced to high accuracy. The narrow lines in the real spectra are re-generated as broader boxcar features due to the frequency bin width used in the simulation process—this is irrelevant since these frequencies are far above the science band.

The signal-to-noise ratio in the timestream data is sufficiently low that we do not need to subtract the signal contribution before taking the noise spectra. (The PSD of signal only LCDM timestream is more than 2 order of magnitude below the noise at all frequencies for the pair sum, and 4 orders of magnitude for the pair difference.) To confirm that the noise modeling operates correctly we have performed a "sim the sim" study where the output of a single signal plus noise timestream simulation is used as the input to a complete simulation set, passing all the jackknife tests described in Section 7 as expected.

5.3. Simulated Maps

We generate many realizations of signal and noise timestream and produce the full set of unsplit and jackknife split maps from these using the same code used to make maps of the real timestream in Section 4.4. For each realization we then add the signal and noise maps to form signal plus noise maps— since the map making process is linear this is equivalent to adding the timestreams, and computationally more efficient. The generation of simulation realizations is computationally costly and hence their number is relatively small (500 in this analysis).

6. FROM MAPS TO POWER SPECTRA

This section will describe the various steps which take us from the real and simulated maps to angular power spectra.

6.1. The Apodization Mask

In addition to the signal maps $M_{T,Q,U}$ the process described in Section 4.4 also produces maps of estimated variance $V_{T,Q,U}$ based on the variance of each co-added half-scan, assuming that the noise is white. Empirically these are found to be close to correct for the polarization maps and a 10% (20%) underestimate for the 100 GHz (150 GHz) T maps, as expected since the pair sum data contain significant $1/f$ noise. We take the inverse of these variance maps as the apodization masks:

$$A_X = \frac{1}{V_X}. \tag{6}$$

However due to the partially overlapping "tiles" of coverage resulting from our observation strategy and the differing detector offset angles there are sharp steps in the masks (of $\sim 10\%$ magnitude). If these steps are not smoothed out then the product of the map and mask (the quantity which is about to get Fourier transformed) will also have sharp steps. Such steps correspond to a mixing of power from large to small scales which is particularly undesirable in the TT spectra where even a small fractional contribution from the very large low multipole C_ℓ values would dominate over the intrinsic power in the damping tail region of the spectrum.[16] To mitigate this effect we convolve the inverse variance mask with a Gaussian shaped smoothing kernel of FWHM $0\overset{\circ}{.}5$:

$$A'_X = A_X * G. \tag{7}$$

We emphasize that an arbitrary apodization mask can be used without biasing the results and, in fact, the inverse variance mask is only optimal in the limit of low signal to noise which is not the case for the T maps. In principle an optimal mask could be generated for each spectrum, and in fact each bandpower within each spectrum, based on its signal to noise, but we have not pursued this complication in this analysis.

After smoothing we inject Gaussian-shaped "divots" with an FWHM of $0\overset{\circ}{.}5$ into the apodization masks at the locations of three discrete sources circled in Figure 6 to null out any effect they might otherwise have on the results:

$$A''_x = A'_X \left(1 - \sum_i \exp\left(-\left(\frac{(x - x_{o,i})^2 + (y - y_{o,i})^2}{\sigma^2}\right)\right)\right), \tag{8}$$

where x and y are the pixel coordinates, $x_{o,i}$ and $y_{o,i}$ are the source locations, and the loop i is over the set of sources to be nulled. Note that although the sources are visible only in the T maps we apply the same masking in polarization also.

The final apodization masks imply an effective sky area for this analysis of ≈ 25 deg^2.

6.2. Fourier Transform and Power Spectra

We next make the flat-sky approximations and take the two-dimensional discrete Fourier transform of the product of the map and apodization mask:

$$m_X = c \, \text{FT}(M_X A''_X), \tag{9}$$

where c is a normalization constant which acts to make the apodization and Fourier transform operations total power preserving.

In the Fourier domain, the transform from Q, U to E, B is simply

$$m_E = + m_Q \cos 2\phi + m_U \sin 2\phi \tag{10}$$
$$m_B = - m_Q \sin 2\phi + m_U \cos 2\phi, \tag{11}$$

where ϕ is the polar angle of each mode m with respect to the Fourier plane origin.

We then take products of the (complex) modes within and between each set

---

[16] Note that the relevant ratio is in C_ℓ, not $\ell(\ell + 1)C_\ell$, which is the quantity conventionally plotted—for the TT spectrum under LCDM $C_{200}/C_{2000} \approx 2500$.

$$p_{XY} = m_X m_Y^*, \qquad (12)$$

where X and Y can be T, Q or U at either 100 or 150 GHz.

Each set of mode products p_{XY} is then multiplied by $d = \ell(\ell + 1)/2\pi$ where the multipole $\ell = 2\pi u$, and u is the Fourier conjugate variable to angular distance from the map center. Finally we take the binned angular power spectra

$$b_{XY,i} = \langle d p_{XY,i} \rangle, \qquad (13)$$

where the mean is taken over the modes in each annulus i of the Fourier plane. The multiplication of the map by the mask in real space corresponds to a convolution in Fourier space—hence the Fourier modes m_X are correlated and so are the bandpower estimates b_{XY}. For this analysis, we set the bin spacing to $\Delta\ell = 81$ which results in correlation of $\sim 20\%$. No importance should be attached to the exact value 81—it was an arbitrary choice made early on and remained unchanged throughout the development of the analysis.

The above description is a slight simplification of the actual procedure used. In reality, for each cross spectrum XY we enforce a common apodization mask $A''_{XY} = \sqrt{A''_X A''_Y}$ before taking the Fourier transforms and proceeding to power spectra.

6.3. Fourier Plane Masking

The strong theoretical expectation is that the CMB is isotropic on the sky and therefore also in the two-dimensional Fourier plane. The noise however is certainly not—atmospheric variation injects strong low-frequency fluctuations ($1/f$ noise) along the scan direction and hence in a band around the perpendicular direction in the Fourier domain. We are perfectly at liberty to down weight or excise portions of the Fourier plane when we take power spectra—no bias will result so long as the weight/cut criteria are independent of the mode values of the real maps.

Figure 7 shows the 150 GHz T and Q Fourier modes for the deck jackknife map. The third-order polynomial filter which has been applied to each half-scan removes power along the scan direction producing the dark vertical bands down the middle of the plots. Atmospheric $1/f$ which survives the filtering appears as a brighter band either side of the dark band in the T plot, but due to the unpolarized nature of the atmosphere does not appear in the Q plot (see Section 5.2). Because the atmospheric noise is highly correlated amongst adjacent detector pairs in the array, we in fact see a complex-structured pattern of nonuniform noise in the T plot—our noise simulations fully reproduce this pattern.

Section 6.1 mentioned the need to smooth the inverse variance derived apodization mask to reduce up-mixing of power from small to large multipoles. There is an additional similar, but smaller, effect. Due to the detector offset angles each has observed a slightly different rectangular patch of the CMB sky. We polynomial filter each detector's timestream for each half-scan to reduce atmospheric noise, and so the CMB modes removed from each "tile" are slightly different. Hence, when we co-add the data from all detectors small step artifacts are present in the overall map at the tile overlap boundaries. Examining the nonjackknife Fourier T modes using plots similar to those shown in Figure 7 we see up-mixed power becoming dominant over the intrinsic for multipoles at $\ell > 1800$ in a narrow band around the scan direction axis—as expected, since the artifacts are vertical step edges in the maps. This spurious power is also seen in the signal only simulations confirming its origin is up-mixing. Note that this effect is only significant in T due to the much larger ratio between low and high multipoles than in the polarization spectra. However, we excise these modes when generating all
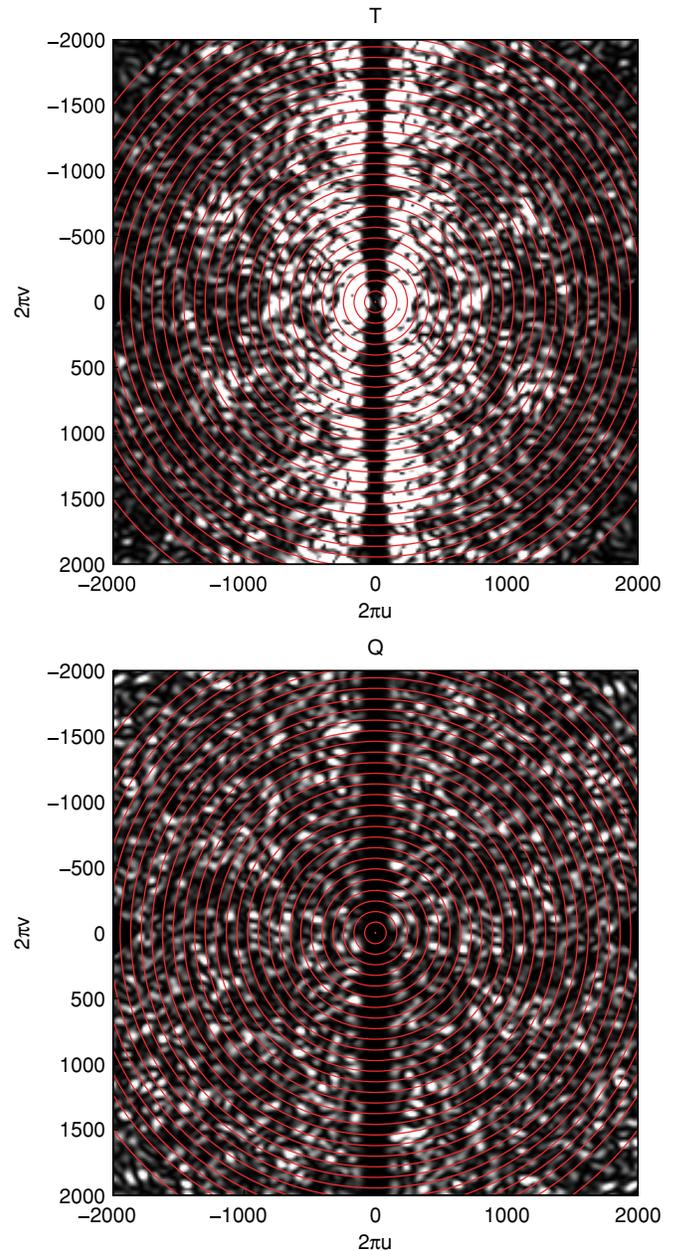
Figure 7. Fourier transforms of the apodized 150 GHz T and Q deck jackknife maps. The real part of the square of the Fourier modes is shown with a linear color stretch. The red circles indicate the annuli within which means are taken to generate bandpower values. Note that the horizontal direction in this plot corresponds to the scan direction—see the text for further details.

power spectra resulting in a trivial increase in the uncertainty of our highest bandpowers ($< 3\%$). This effect will be more important in a future analysis extending to $\ell = 3000$.

One could weight the Fourier modes when combining to form power spectra based on their signal-to-noise ratio as measured in the signal and noise simulations. This would be (at least a partial) substitute for the polynomial half-scan filtering—one either removes the atmospheric $1/f$ in the timestream or in the two-dimensional Fourier plane. Looking at Figure 7, such Fourier plane weighting would clearly be beneficial for the TT spectra. Although we have experimented with such weights and cuts we have not implemented them in the analysis presented here—the benefit for the essentially white noise polarization spectra is close to zero.
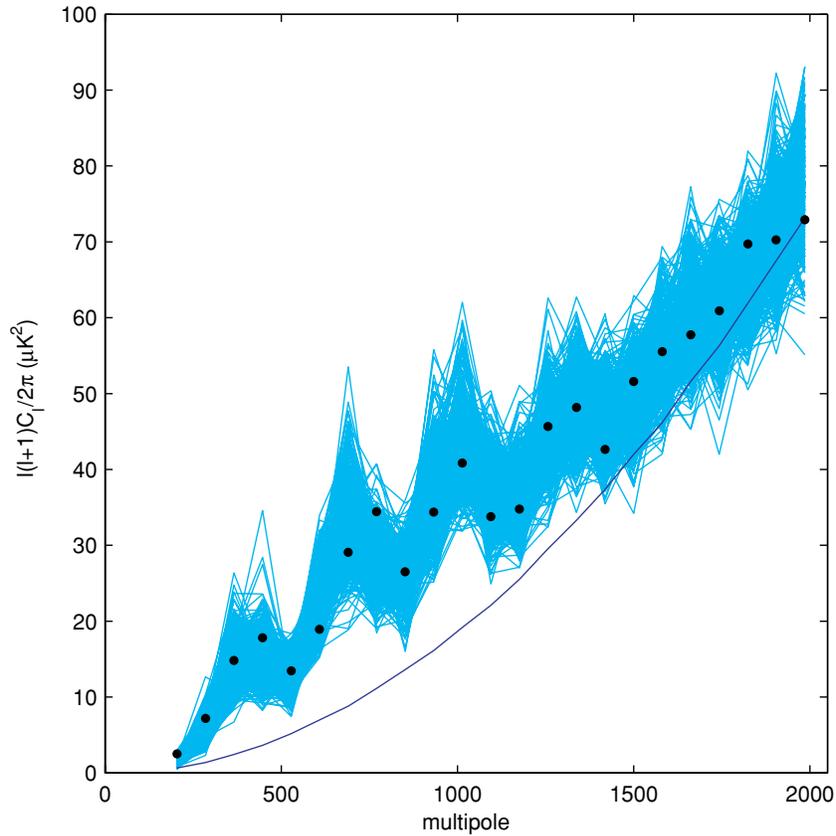
Figure 8. Comparison of real and simulated raw 150 GHz *EE* spectra. The cyan curves are the ensemble of signal plus noise simulations, the black points are the observed bandpower values, and the blue curve is the mean of the ensemble of noise only simulations.

6.4. Raw Spectra Compared to Simulations

By taking the mean of the masked Fourier mode products within each annulus, we generate raw power spectra of the real and simulated maps. Figure 8 compares the 150 GHz *EE* spectra and is in a sense the fundamental result of our analysis—is the observed spectrum consistent with being a realization of LCDM plus noise? To test that hypothesis, we could simply perform χ^2 tests at this point. However, for presentation purposes, and to allow our spectra to be used for cosmological parameter analyses, we proceed to noise and filter correct them as follows.

6.5. Removing the Noise Bias

At this point in the analysis, we have sets of bandpower values calculated using Equation (13) for each of the auto and cross spectra. We have these for the real spectra, and each realization of the signal only, noise only, and signal plus noise simulations. We will denote these by $\mathbf{r}_{XY}$, $\mathbf{s}_{XY}$, $\mathbf{n}_{XY}$, and $\mathbf{sn}_{XY}$, respectively.

To remove the noise bias we take the mean of the ensemble of noise only simulations and subtract from the real spectra,

$$\mathbf{r}'_{XY} = \mathbf{r}_{XY} - \overline{\mathbf{n}_{XY}}. \tag{14}$$

In Figure 8, the blue curve is subtracted from the black points. Correlated noise, presumably from the atmosphere, is fully simulated by the process described in Section 5.2. Such noise can result in the mean noise cross spectra—where cross spectrum means either within or across frequency bands—being nonzero. In practice, the levels are very small compared to the auto spectra but for completeness we subtract them anyway.

6.6. Bandpower Window Functions

As mentioned in Section 6.2, multiplication by the apodization mask in image space corresponds to convolution in the Fourier plane by the Fourier transform of that mask. In Figure 7 we see that adjacent Fourier modes (pixels) are highly correlated, and that the resulting annular mean bandpower estimates will therefore also be correlated. We thus need to know how much each multipole on the sky contributes to each experimental bandpower—the so-called bandpower window functions (BPWFs; Knox 1999). This could be determined by running sets of signal only simulations, each with nonzero input power only in a narrow band of multipoles. However the computational cost of doing this in practice is too high, and we hence use an alternate, much faster method.

We take each narrow annulus in the Fourier plane in turn and convolve this with the Fourier transform of the mask. For *Q* and *U*, we generate annuli corresponding to pure *E* or *B* modes. Taking the power spectrum of the convolved annuli measures the response of each experimental bandpower to sky power in the given annulus. Placing these spectra into the rows of a matrix, the columns are then the desired BPWFs (interpolating to every multipole). For *BB* we have two functions per bandpower representing the response to true *BB* sky power, and also the response to *EE*. In practice, the mixing is $\sim 10\%$ in our lowest bandpower and falls rapidly with increasing ℓ (see Section 10.1). Figure 9 shows some example BPWFs.

6.7. Determining the Filter/Beam Suppression Factor

Using the BPWFs described above, we can calculate the expectation values $\mathbf{e}_{XY}$ of the observed bandpowers given
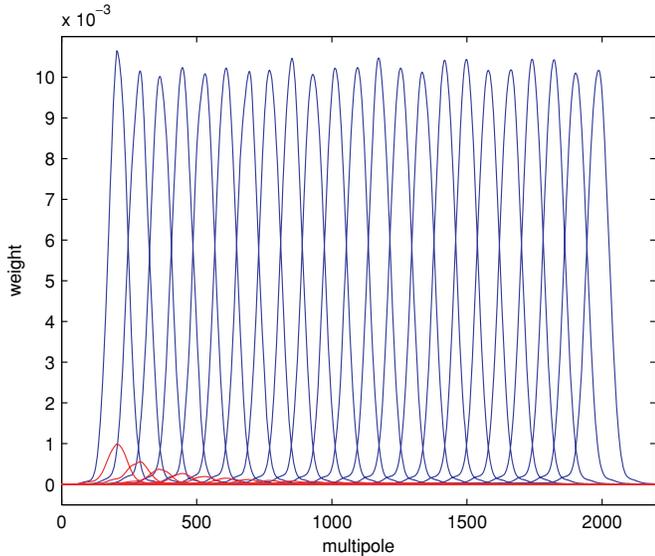
Figure 9. BPWFs for the 150 GHz *BB* spectrum. Each blue curve shows the relative response to *BB* sky power, while the corresponding red curves show the response to *EE*.

the input LCDM power spectra used to generate the signal simulations. We then take the ratio of the mean of the signal only simulated spectra to these expectation values to empirically determine the suppression of power which has occurred due to the convolution of the sky by the telescope beam pattern, and the polynomial filtering of the timestream

$$\mathbf{f}_{XY} = \frac{\overline{\mathbf{s}_{XY}}}{\mathbf{e}_{XY}}, \tag{15}$$

this is the factor by which the real spectra must be divided to yield an unbiased estimate of the true sky power. Note that this process will automatically include any simulated suppression effect (including the so-called pixel window functions). Figure 10 illustrates this step—the curve in the lower panel approaches the "beam window function" W_ℓ at higher ℓ.

We now divide the spectra by their respective suppression factors

$$\mathbf{r}''_{XY} = \frac{\mathbf{r}'_{XY}}{\mathbf{f}_{XY}}. \tag{16}$$

In practice, we apply the $\mathbf{f}_{TT}$ correction to the *TT* spectra and the $\mathbf{f}_{EE}$ correction to the *EE*, *BB*, and *EB* spectra. For the *TE*, *TB* spectra, we use the geometric mean of $\mathbf{f}_{TT}$ and $\mathbf{f}_{EE}$.

6.8. Power Spectrum Results

We also apply the noise de-bias and filter/beam correction operations of Equations (14) and (16) to each signal plus noise realization. Their fluctuation then provides an estimate of the uncertainty of the real bandpower values. This uncertainty will only be correct in as much as the theory spectrum used as input to the signal simulations matches reality—if there is any significant disagreement one should iterate the entire process until it converges. As we will see in Section 9.2, our results are in fact perfectly compatible with the *WMAP3*-based model used as input to the signal simulations (see Section 5.1) so there is no need to iterate.

Figure 11 shows the full set of 21 signal spectra: within each frequency band we have *TT*, *TE*, *EE*, *BB*, *TB*, and *EB*, while across frequency bands we can form $T_{100}T_{150}$, $T_{100}E_{150}$,
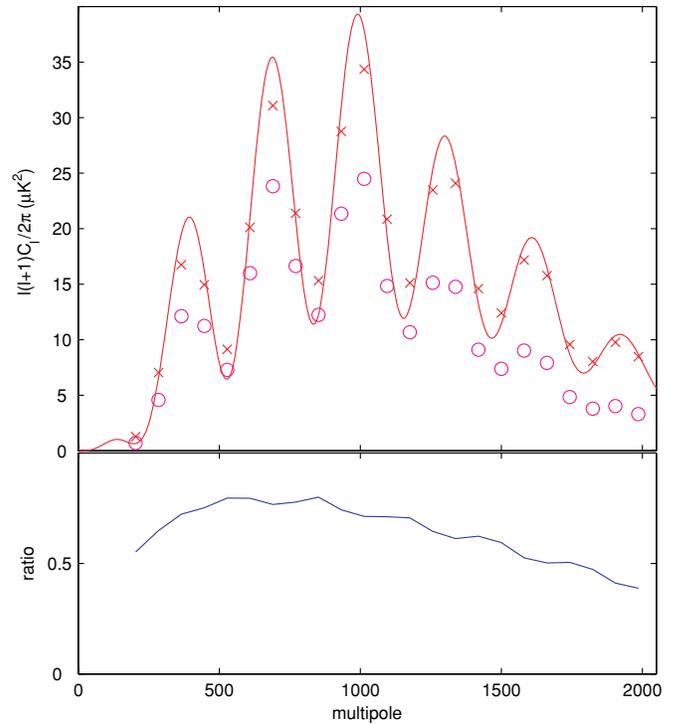
Figure 10. Upper: input *EE* spectrum (red line), calculated expectation values (red crosses), and mean simulated signal only spectrum (magenta circles). Lower: the ratio of the two sets of points in the upper panel—the filter/beam suppression factor. (This plot is for the 150 GHz spectrum.)

$E_{100}T_{150}$, $E_{100}E_{150}$, $B_{100}B_{150}$, $T_{100}B_{150}$, $B_{100}T_{150}$, $E_{100}B_{150}$, and $B_{100}E_{150}$. We see that the 150 GHz spectra have better sensitivity than the 100 GHz—this is partly due to the larger number of detectors at the higher frequency, but mostly due to the smaller beam size (and hence smaller beam correction).

6.9. Alternate Analysis

The main analysis in this paper comes from an evolution of "pipeline 2" in our previous paper Ade et al. (2008). A second pipeline exists, the low level parts of which are derived from our previous "pipeline 1," but which, for the results presented here, is now also performing flat-sky power spectrum estimation. This pipeline currently has noise modeling that is somewhat less sophisticated than that described in Section 5.2. The two pipelines were independently written and share no code. Although the algorithms implemented are intended to be basically the same they are sure to differ in some (hopefully unimportant) details. The fact that the final results agree very closely therefore adds considerable additional confidence. Figure 12 compares the 150 GHz spectra from the two pipeline.

7. JACKKNIFE TESTS

To probe the results shown in Figure 11 for systematic contamination we conduct a battery of jackknife tests. As mentioned above split, data set maps of the real data and each simulation realization are generated. Each split divides the timestream data into two approximately equal subsets which should contain (nearly) identical sky signal, but which might contain different contaminating signal. In this analysis we consider four such data splits which we call the deck angle jackknife, scan direction jackknife, split season jackknife, and
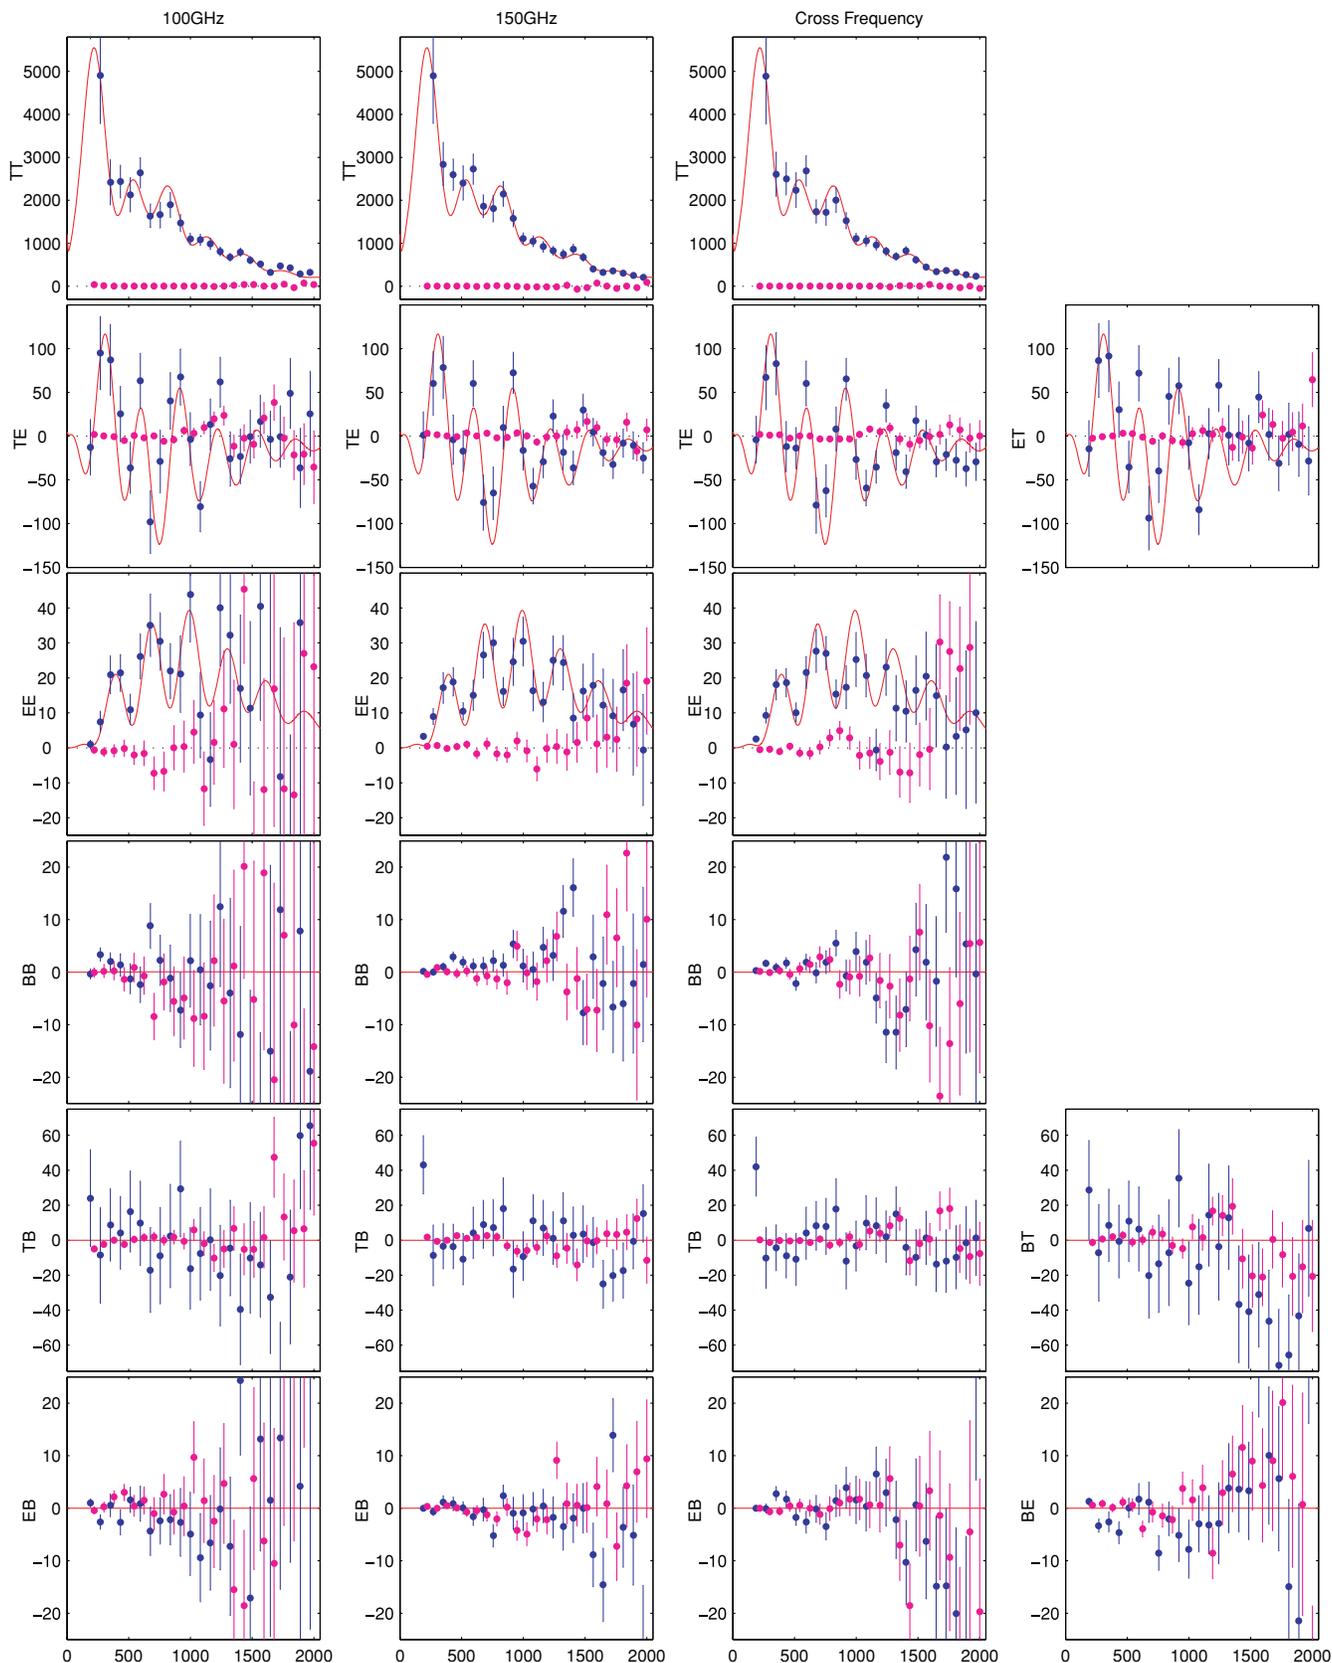
Figure 11. Full set of QUaD power spectrum results. The blue points are the signal spectra, the magenta the deck jackknife (see the text for definition), and the red curves a conventional LCDM model. The horizontal and vertical axes are multipole ℓ and $\ell(\ell + 1)C_l/2\pi$ (μK^2), respectively. Each bandpower is $\sim 20\%$ correlated with its neighbors and, where signal dominated, there are strong correlations between bandpowers across frequency bands. The error bars are the standard deviation of the simulated signal plus noise bandpower values, and the two sets of points have been offset by ± 15 in ℓ from their nominal values for clarity. Note that absolute calibration and beam size uncertainty are not included in the error bars. The jackknife spectra are consistent with null—see Table 1 and the text.
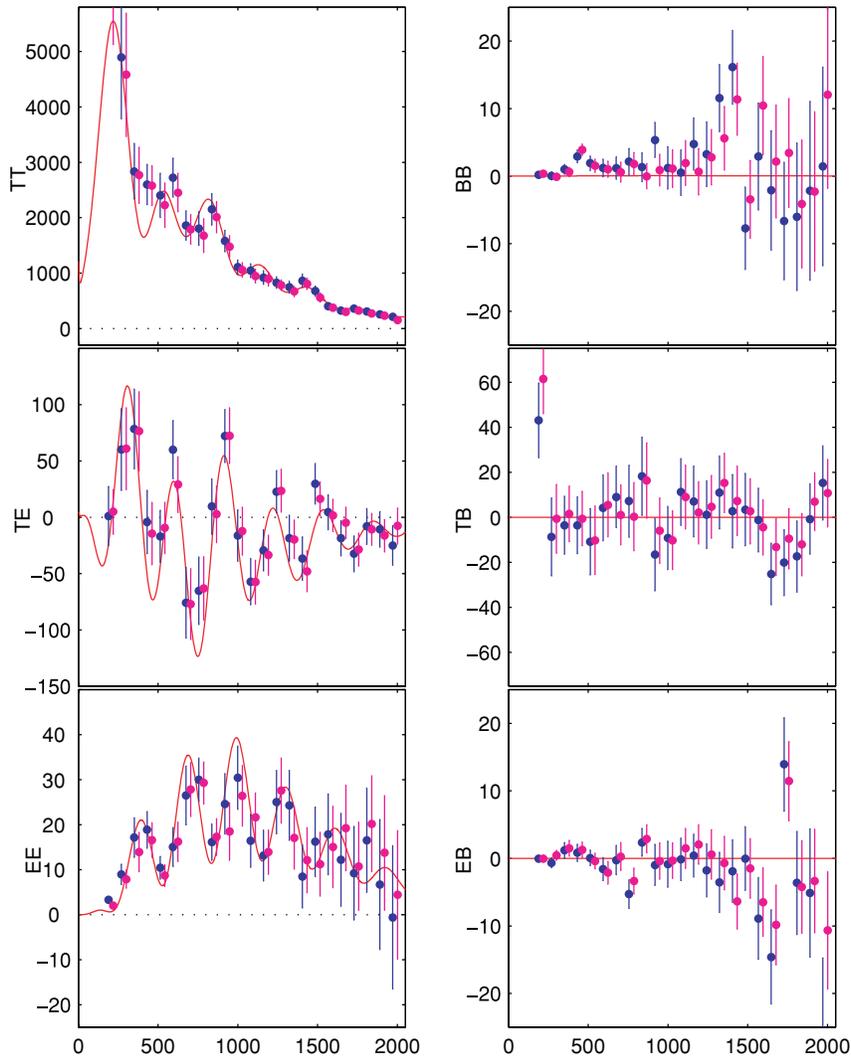
Figure 12. Comparison of 150 GHz spectra for the main (blue points) and alternate (magenta points) pipelines. The horizontal and vertical axes are multipole ℓ and $\ell(\ell+1)C_l/2\pi$ (μK^2), respectively, and the two sets of points have been offset by ± 15 in ℓ from their nominal values for clarity.

focal plane jackknife. Each of these is described in more detail below.

For each jackknife the split maps are subtracted from one another, divided by two, and we then proceed to power spectrum estimation as usual. For the real spectra we then calculate the χ^2 versus the null model,

$$\chi^2 = \mathbf{r}'' \mathbf{C}^{-1} \mathbf{r}''^t, \tag{17}$$

where $\mathbf{C}$ is the bandpower covariance matrix as estimated from the ensemble of signal plus noise simulations. We also calculate χ^2 for each signal plus noise simulation. The generation of simulation realizations is computationally costly and hence their number is relatively small (500 in this analysis). This has two implications. Firstly, since bandpower correlations beyond nearest neighbor are sufficiently weak as to be lost in the measurement noise, we set all but the main and first two off-diagonals of the bandpower covariance matrix to zero. Second, when calculating the χ^2 values for the signal plus noise realizations extreme fluctuations have the opportunity to partially self-compensate by injecting extra covariance into the matrix which they will be measured against, biasing the resulting χ^2 values low. To get around this we re-calculate the covariance

matrix excluding each realization in turn before calculating the χ^2 for that realization.

We find that the simulated jackknife χ^2 distributions for *TT* do not follow the analytical expectation. Examining Figure 13, which shows the bandpower deviations of the deck jack spectra, we see that there are two reasons why. Firstly, the simulated bandpower distributions are significantly non-Gaussian—the colored lines do not sit at -2 through $+2$. Second, and more importantly, there is significant predicted imperfect cancellation—the simulated distribution deviates strongly from a median of zero at large angular scales. The reason for this latter effect was mentioned in Section 6.3—in any split where the sky coverage "tiling" for the two subsets is nonidentical the interaction of the true sky brightness distribution and the half-scan polynomial filtering generates slightly different output maps. Consistent with our overall simulation-based approach we therefore take the probability to exceed the real spectra χ^2 values versus the simulated distributions, rather than the analytical.

Table 1 shows the full set of probability to exceed (PTE) values for all the jackknifes and spectra—there are no strong indications of problems. These values are expected to be uniformly distributed between zero and one, and in Figure 14
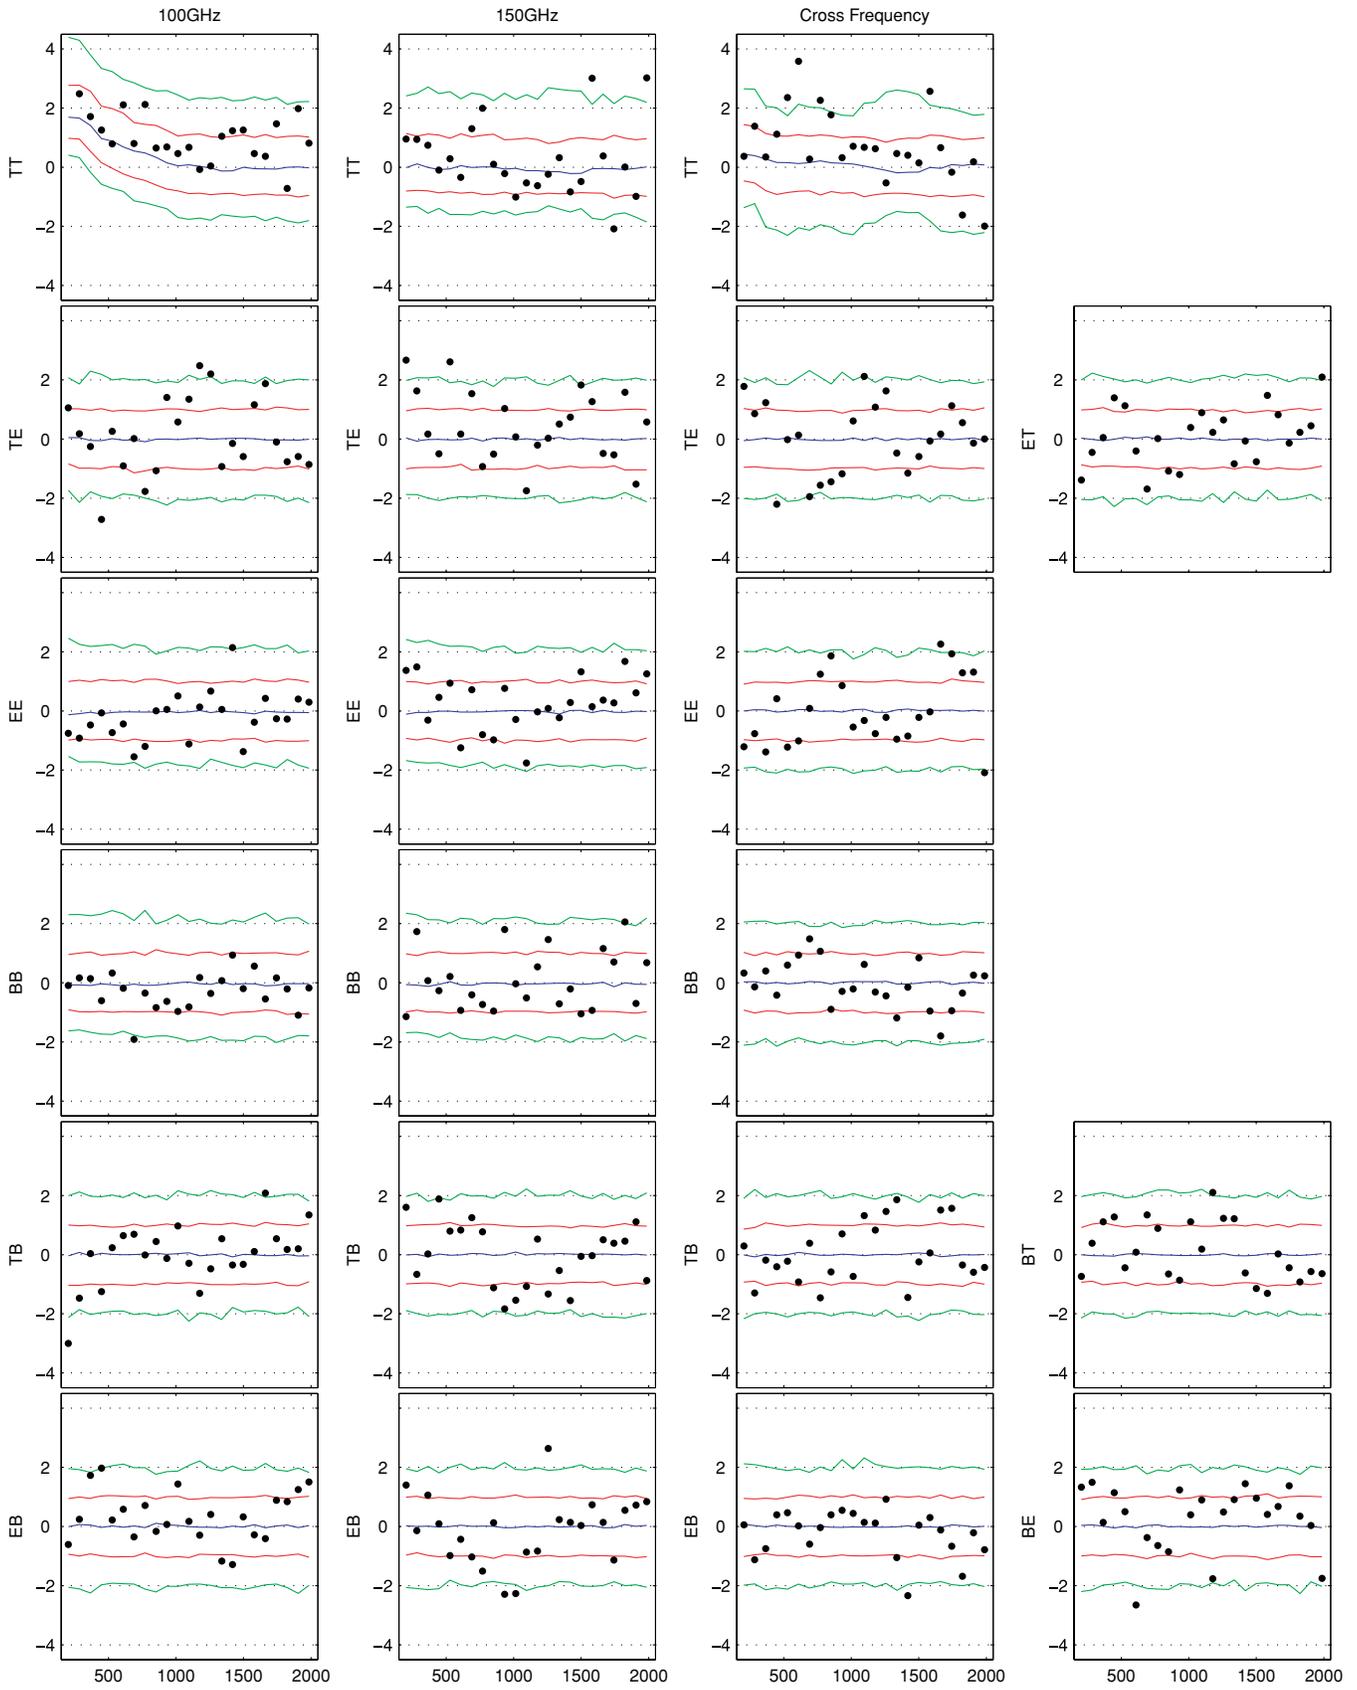
Figure 13. Bandpower deviations of the deck jackknife spectra. The horizontal axis is multipole number. The spectra shown as magenta points in Figure 11 have been divided by their error bars to investigate the contribution of each bandpower to χ^2. The green, red, and blue lines show the 2.3%, 15.9%, 50%, 84.1%, and 97.7% points of the integral distribution of the simulated signal plus noise bandpower values.
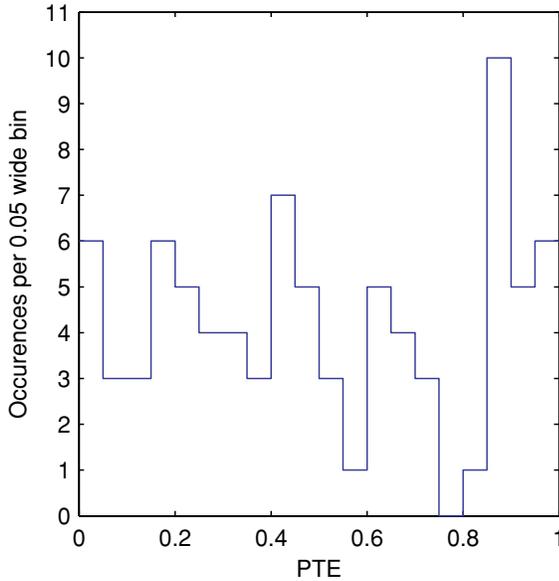
Figure 14. Distribution of the χ^2 PTE values from Table 1.

Table 1
Jackknife PTE Values from χ^2 Tests

Jackknife	100 GHz	150 GHz	Cross	Alt. Cross
Deck angle				
TT	0.102	0.072	0.030	
TE	0.066	0.040	0.174	0.586
EE	0.874	0.624	0.210	
BB	0.990	0.316	0.946	
TB	0.506	0.410	0.520	0.632
EB	0.736	0.210	0.952	0.180
Scan direction				
TT	0.650	0.238	0.982	
TE	0.316	0.444	0.870	0.718
EE	0.656	0.490	0.372	
BB	0.882	0.156	0.982	
TB	0.274	0.430	0.154	0.254
EB	0.540	0.332	0.826	0.418
Split season				
TT	0.202	0.020	0.408	
TE	0.962	0.296	0.300	0.472
EE	0.116	0.038	0.066	
BB	0.896	0.036	0.100	
TB	0.946	0.494	0.896	0.166
EB	0.860	0.612	0.972	0.890
Focal plane				
TT	0.352	0.472	0.232	
TE	0.858	0.872	0.448	0.936
EE	0.654	0.174	0.944	
BB	0.696	0.642	0.708	
TB	0.888	0.422	0.008	0.456
EB	0.932	0.640	0.354	0.344

they are seen to be so. See the following subsections for detailed discussion of these results.

7.1. Deck Jackknife

The deck angle jackknife is perhaps the most powerful. Due to the locking of our daily observations to LST, the two halves of this split contain data taken over completely different ranges of telescope azimuth. In addition, and equally importantly, the entire telescope is rotated by $60°$ around the line of sight between the two observation sessions. The two data subsets are therefore separated from one another in time, azimuth angle, and detector polarization angle as projected on the sky. It is very hard to conceive of a source of contamination which would be common in the polarized maps for the two halves of this split. We see no significant indication of problems with the deck angle jackknife in Figures 11 and 13, or in Table 1. One might perhaps worry about the rather low values for 150 GHz and cross-frequency *TT* but we caution against overinterpreting such numbers—the table contains 84 numbers so on average four numbers below 0.05 are to be expected. In addition, looking at Figure 13 we see that the bandpower making the strongest contribution to the cross-frequency χ^2 has a very low deviation in the 150 GHz spectra. Note that Figure 14 shows no obvious excess of low PTE values.

7.2. Scan Direction Jackknife

The scan direction jackknife splits the data into the outgoing and returning half of the scans. In terms of external contamination this is perhaps the easiest test to pass—only something very rapidly varying would cause it to fail. However, it is a stringent test for internal instrumental effects. Any scan synchronous false signal, caused perhaps by motion of the liquid cryogens in the tanks provoked by the telescope motion, would likely fail to cancel in this jackknife. Also any failure to adequately deconvolve the temporal response of the detector channels would cause this test to fail. Looking at Table 1 we see no problems.

7.3. Split Season Jackknife

For the split season jackknife, we have divided the time-ordered list of days in half (not in fact into 2006 and 2007—the

split occurs in late August 2006). We then make maps with each set of days and difference them. This test would fail if there were a significant shift in the absolute calibration of the telescope system beyond the atmospheric loading effect whose correction is described in Section 3.4. Looking at Table 1, we see a few numbers below 0.05 but nothing highly significant. Examining the bandpower deviations (a plot not shown in this paper analogous to Figure 13), we find that for each spectrum with a low PTE it is caused by a different bandpower(s), giving no further hint of any problem.

7.4. Focal Plane Jackknife

This jackknife splits the detectors into the two orientation groups which are separated by $45°$—referred to as instrument-*Q* and instrument-*U*. Due to co-adding across the two deck angles each map pixel has still been observed at two polarization angles allowing the construction of *Q* and *U* maps as usual. This test is perhaps the weakest, but might reveal problems with instrumental false signal in a subset of the pairs (although it is hard to see how that would not also show up in the deck angle jackknife). Looking at Table 1 we see no problems. Although this jackknife includes the lowest number in the table (0.008) one such number is not improbable and we note again that the PTE distribution shown in Figure 14 is consistent with uniform.

8. FOREGROUND STUDIES

8.1. Frequency Difference Maps and Spectra

We can also take the difference between our 100 GHz and 150 GHz maps. It is important to be clear that this is not a jackknife in the sense of Section 7—the true sky brightness distribution may in fact differ at these two frequencies. Therefore
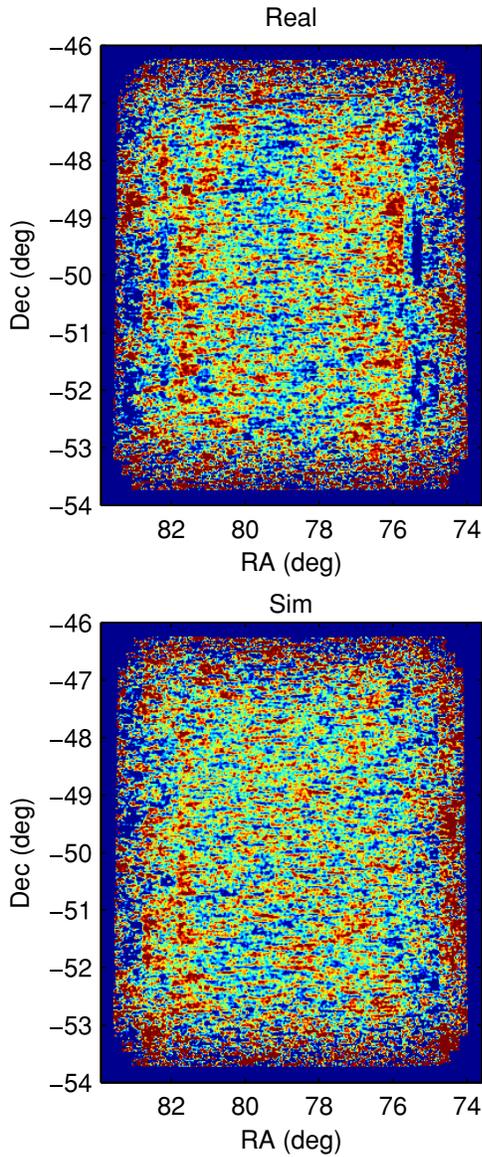
Figure 15. Upper: the difference between the 100 and 150 GHz T maps shown in Figure 6 on a $\pm 20\,\mu$K color stretch. Lower: the same thing for a signal plus noise simulation realization.

Figure 16. Left: frequency difference spectra; right: the associated bandpower deviations. χ^2 is taken using the bandpower covariance matrix and converted to the PTE value shown on each panel. (For further explanation see the captions to the analogous Figures 11 and 13.)

any failure to cancel might be due to the presence of astrophysical foregrounds, as well as instrumental systematics or contamination.

As described in Section 4.5, we find the absolute calibration scalings for our 100 and 150 GHz maps by cross-correlating them against the same B03 maps (which are at 150 GHz). However, if the sky pattern differs at these two frequencies then this will still show up in the frequency difference maps. Figure 15 shows the difference between the real 100 and 150 GHz T maps, and the same thing for a signal plus noise simulation realization.

As mentioned in Section 7, when subtracting maps with different sky coverage "tiling" a small degree of mismatch is expected due to the polynomial filtering. This is the cause of the vertical "step edges" observed in the upper panel of Figure 15. The simulation realizations show similar effects, although in the real map they do appear to be unusually strong. Note that we also expect to see cancellation failure at smaller angular scales due to the differing beam sizes for the two frequency bands.

We next take power spectra of the frequency difference maps and compare them to simulation as was done for the jackknifes in Section 7. Figure 16 shows the resulting bandpower deviations—PTE values analogous to those in Table 1 are shown in the figure. The TT spectrum shows larger deviations at low multipoles than expected from the simulations, and the probability that these are caused by the differing "tiling" effect alone is low. However, the absolute value of these bandpowers is $\approx 15\,\mu$K^2 to be compared to the thousands of μK^2 in the undifferenced map, i.e. the fractional cancellation failure is very small and completely irrelevant compared to the sample variance in

the *TT* spectra. It is not clear whether the excess cancellation failure is due to instrumental effects or real foreground signal.

8.2. Predicted Diffuse Foreground Levels

Our field was chosen to partially overlap the B03 deep field to allow absolute calibration against that map. They in turn chose the location based in part on the position of the Sun during their balloon flight. Although low, the foreground emission in this region is not the lowest available on the sky.

To obtain estimates for the expected level of foreground dust emission we use the Finkbeiner, Davis and Schlegel (FDS) model 8 extrapolation of IRAS maps (Finkbeiner et al. 1999)— this is shown in Figure 2. For synchrotron emission, we use an extrapolation of 408 MHz maps (Haslam et al. 1981; D. P. Finkbeiner 2001, private communication). We pass these maps through the QUaD simulation pipeline, including the field differencing and filtering operations. For dust, the maximum of the resulting *TT* spectra is $\sim 4 \, \mu K^2$ (150 GHz, $\ell = 200$) while the synchrotron maximum is a negligible 0.03 μK^2 (100 GHz, $\ell = 200$). Although these models are possibly not the current best available data it is clear that such extrapolations will not give detectable levels in either temperature or polarization, in the presence of CMB, and given the sensitivity of QUaD.

8.3. Point Sources

Three discrete sources are visible in the 100 and 150 GHz *T* maps shown in Figure 6. As described in Section 6.1, these are masked before taking the power spectra. Turning off the masking of these sources we find that the 100 GHz *TT* bandpowers increase by $\approx 10\%$ at the highest multipole considered in this paper (2000) with the increase falling off to lower multipoles. The 150 GHz *TT* spectrum shows an increase of $\approx 3\%$ at ℓ of 2000 and a similar falloff. The *EE* and *BB* bandpowers show changes of $\ll 1\%$.

The source flux distribution dN/ds is typically a power law with the majority of the anisotropic power being contributed by the brightest few sources. We therefore estimate the residual point source contribution to be < 3% in the highest bin of 100 GHz *TT*, < 1% in the highest bin of 150 GHz *TT* and negligible in all other spectra.

8.4. Template Cross Correlation

To test for the possibility of emission correlated with thermal dust, but stronger than expected on the basis of extrapolation, we have carried out a template cross-correlation study. After passing the FDS dust maps through our pipeline, the resulting maps were cross-correlated with the corresponding QUaD CMB maps. Though we might expect some nonzero correlation simply by a chance alignment of large-scale structure (Chiang et al. 2008), we find that compared to simulations there is no evidence of contamination. Indeed analyzing the cross power spectra at both 100 and 150 GHz between dust and CMB on a per bandpower basis, against simulations, reveals no problems, nor any suggestion that dust foregrounds are responsible for the cancellation failure seen in the lower bandpowers of the frequency difference *TT* spectrum.

9. COMBINED SPECTRA

We now wish to form a single combined set of spectra from the 100 GHz, 150 GHz, and frequency-cross results presented above. For each bandpower of each spectrum, we take the 3×3
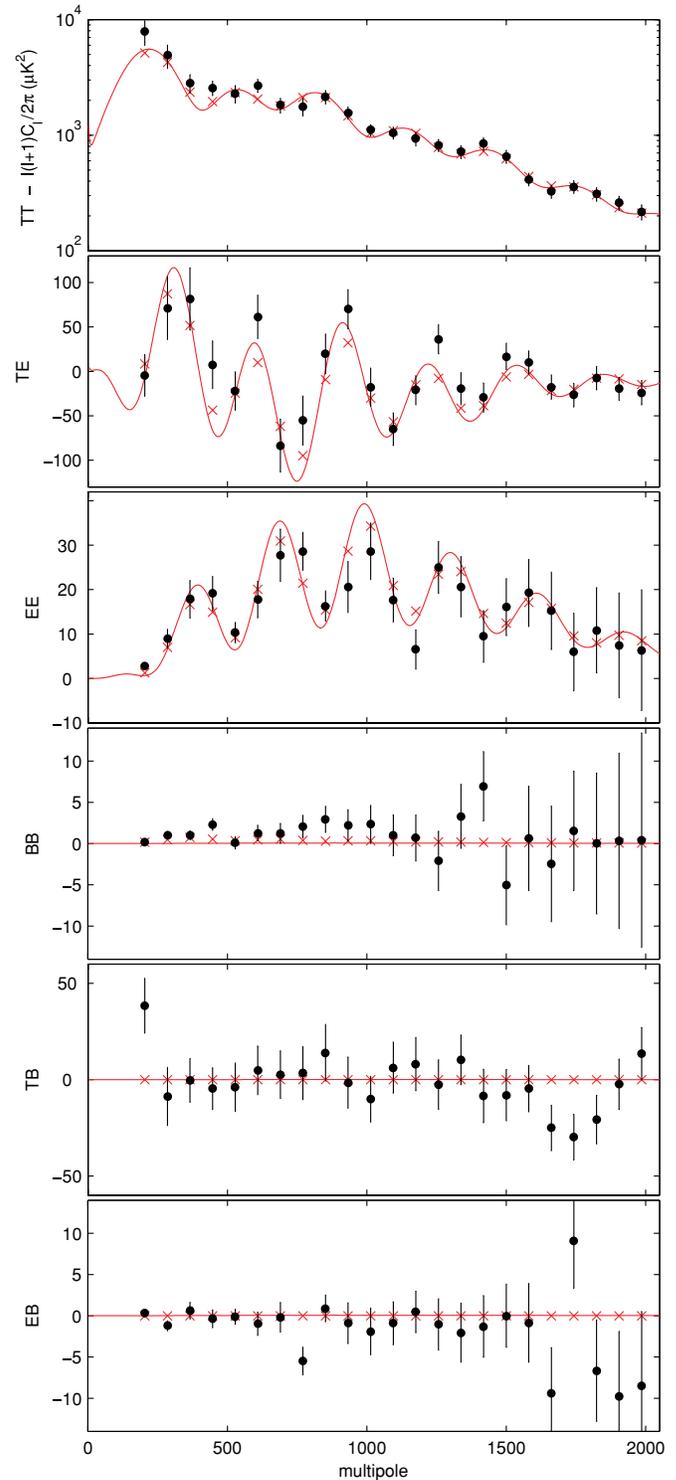
Figure 17. Combined QUaD power spectra shown as black points with error bars. The red crosses are the expectation values for each bandpower given the LCDM model plotted as a red line. Note the differing *y*-axis scales.

covariance matrix over the ensemble of signal plus noise simulations (4×4 for *TE*, *TB*, and *EB*). The combination weights are the column (or row) sums of the inverse of this matrix. The improvement over the 150 GHz bandpower uncertainties is between zero and 30% depending on whether the bandpower is signal or noise dominated. The BPWFs are also combined. Figure 17 shows the combined spectra as compared to their

expectation values under LCDM. The plotted bandpowers, together with their covariance matrices and BPWFs, are available in numerical form at http://quad.uchicago.edu/quad. In contrast to the jackknife spectra, for the signal spectra we find the simulated bandpower distributions to be Gaussian (see Section 7).

9.1. Absolute Calibration and Beam Systematics

In addition to the sample and noise variance error bars shown in Figure 17 there are two major sources of systematic uncertainty. The first of these is the uncertainty on our absolute calibration against the B03 maps. As mentioned in Section 4.5, we estimate this uncertainty as 5% in temperature units (10% in power). To estimate this uncertainty one simply multiplies the bandpower expectation values by 0.1, and takes the outer product as an addition to the bandpower covariance matrix:

$$\mathbf{C}_a = \sigma_a^2 \mathbf{e}^t \mathbf{e}, \tag{18}$$

where $\sigma_a = 0.1$.

The second major systematic effect is uncertainty on the width of the telescope beam. It is highly unlikely that the widths used in the simulations are significantly broader than the true values—they are very close to the results of physical optics calculations (O'Sullivan et al. 2008). However, it is conceivable that we have somehow underestimated the pointing wander and associated effects discussed in Section 5.1, causing the effective overall beam width in the simulations to be narrower than that in the real maps, and the suppression factor curve plotted in Figure 10 to be higher than it should be. This would cause the corrected bandpower values to be biased low with increasing ℓ.

In addition, we have not yet carried out exhaustive investigations of the beam shape and measurement uncertainties. Therefore, we very conservatively assign a beam width uncertainty of 10% for this analysis with the expectation that this will be improved upon in a future analysis dedicated to high ℓ TT. The effective beam FWHM for the combined spectra is $4'.1$ and the fractional shift in the bandpower values which would result from increasing this is well approximated by

$$S_\ell = \frac{W_\ell}{W_\ell'} - 1 = e^{\sigma_b^2 (\delta^2 + 2\delta)\ell(\ell+1)} - 1, \tag{19}$$

where $\sigma_b = \theta_{\mathrm{FWHM}}/\sqrt{8 \ln 2}$, and $\delta = 0.1$ is the fractional beam error. To estimate this uncertainty one multiplies the bandpower expectation values by the S_ℓ factors calculated at the band center ℓ values, and takes the outer product as an addition to the bandpower covariance matrix:

$$\mathbf{C}_w = (\mathbf{e}S_\ell)^t (\mathbf{e}S_\ell). \tag{20}$$

Figure 18 shows the magnitude of these uncertainties, as well as the sample and noise variance, for each of our combined spectra. All spectra except BB are sample variance dominated at lower ℓ, and all but TT are noise dominated at high ℓ. Absolute calibration uncertainty is subdominant to random uncertainty for all spectra at all ℓ. Beam uncertainty becomes the dominant effect for TT at high ℓ.

9.2. Comparison to LCDM

As mentioned in Section 9, for the signal spectra we find the simulated bandpower distributions to be Gaussian. In addition, we find the simulated χ^2 distributions to be consistent with the analytical expectation. Therefore, in contrast to the jackknife
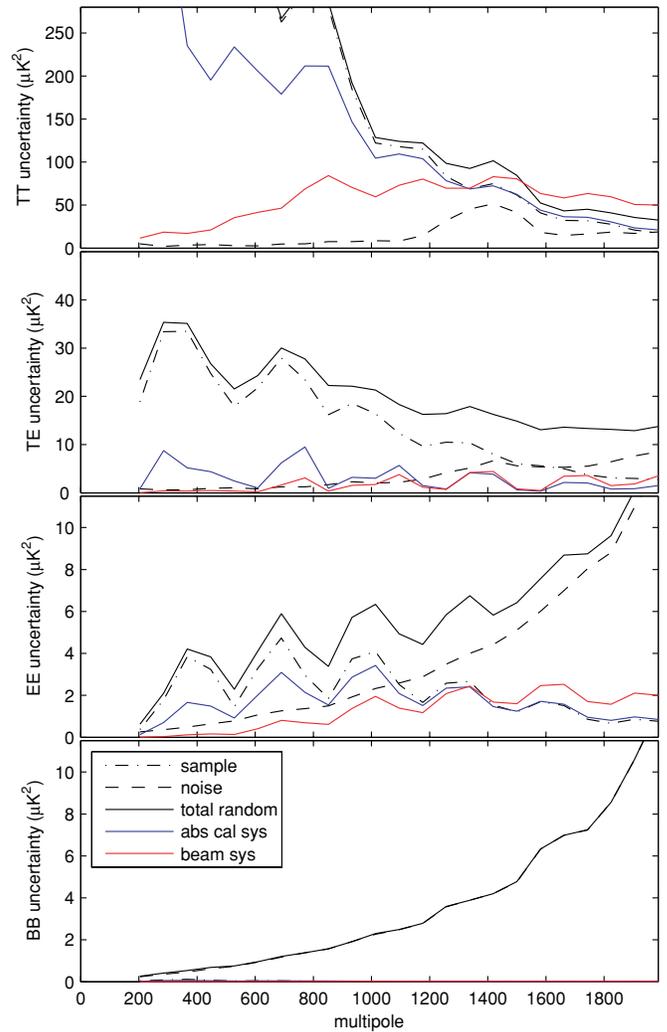
Figure 18. Contributions of the various sources of uncertainty for the combined bandpowers.

spectra, for the signal spectra we quote the analytical probability to exceed. Figure 19 shows the bandpower deviations comparing the combined spectra to LCDM and the null model. Our results are clearly perfectly compatible with LCDM and crushingly incompatible with the no polarization hypothesis. Note that ~ 10 EE bandpowers have $> 4\sigma$ significance.

9.3. Investigation of EE Peaks

Looking at Figure 17, we appear to see three or four of the expected acoustic peaks in the EE spectrum. It is interesting to ask at what significance these have been detected. One way to do this is to determine the χ^2 of the observed bandpowers against a version of the LCDM model which has been smoothed sufficiently to remove the peaks. Figure 20 shows the result—the LCDM model has been convolved with a Gaussian with $\sigma_\ell = 150$. We see that the probability that such a model is correct is very low, the χ^2 PTE being 0.001.

In LCDM, the TT and EE peaks are approximately half a cycle out of phase with one another. If ℓ_s is the peak spacing, and n is the peak number (starting from one), then the approximate locations of the TT peaks are $(n - 1/4)\ell_s$, while the EE peaks are at $(n - 3/4)\ell_s$. To investigate how well our EE bandpowers constrain the peak spacing, phase, and amplitude, we carry out an analysis similar to that of Readhead et al. (2004). Subtracting
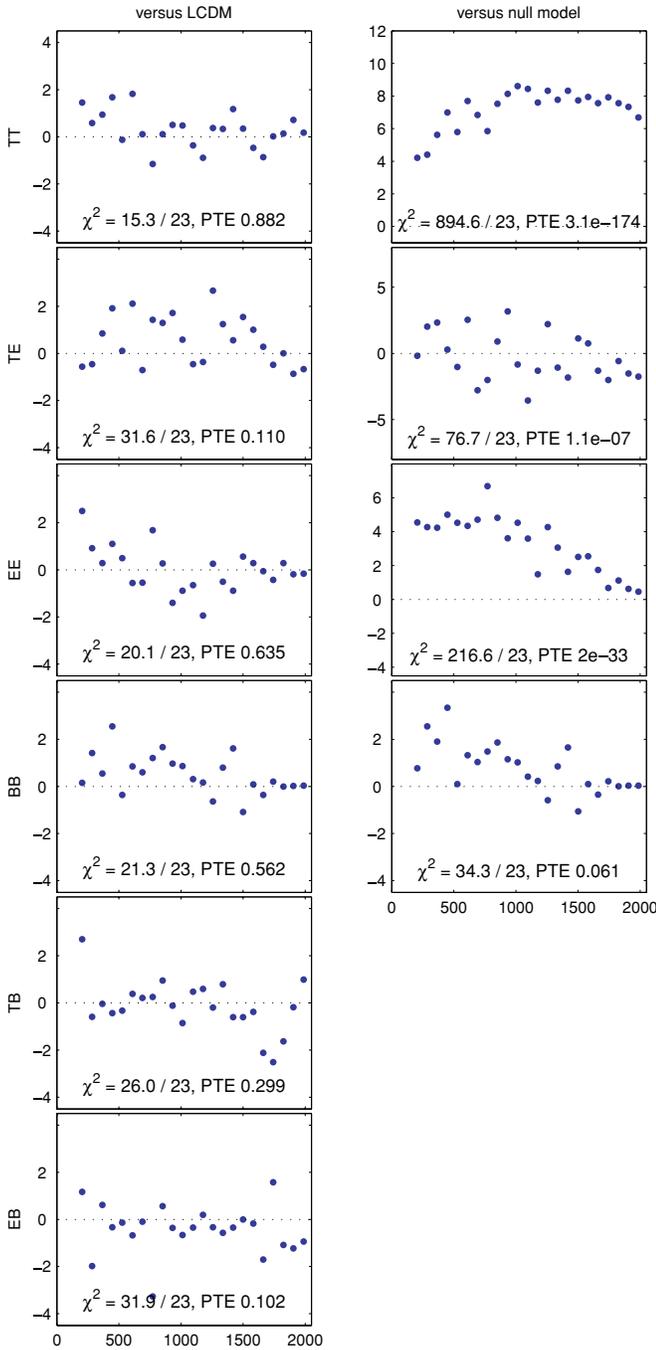
Figure 19. Bandpower deviations comparing the combined spectra to left: LCDM; and right: the null model. The horizontal axis is multipole number. The respective model expectation values have been subtracted from the spectra shown in Figure 17 and the result divided by the error bars. Noted on each panel is the χ^2/ndf and the probability to exceed this value by chance (χ^2 being calculated using the bandpower covariance matrix). Note the differing y-axis ranges in the right column. Note also that the *BB* LCDM expectation values contain a significant leakage contribution and this is why the left and right *BB* panels differ so much—see Section 10.1 for details.

the smoothed version of the LCDM *EE* spectrum from the unsmoothed (i.e., subtracting the green from the red curves in Figure 20) results in a series of approximately sinusoidal modulations whose fractional amplitude dies away close to linearly from the 2nd to the 9th peaks. This fact allows us to generate "toy-model" *EE* spectra which follow the envelope of
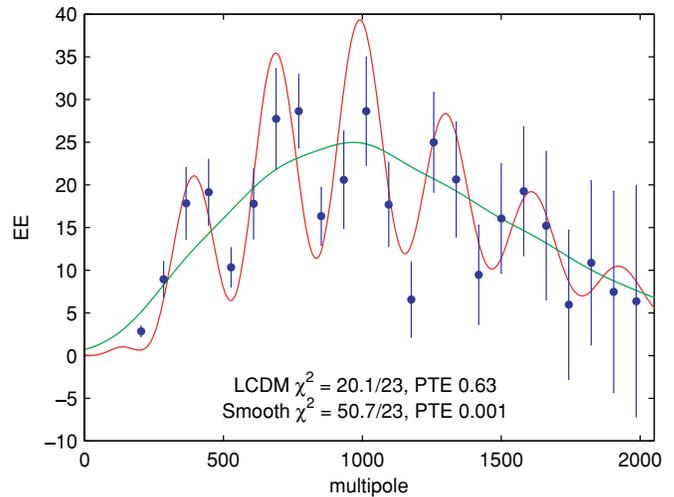
Figure 20. Comparison of the combined *EE* spectrum (blue points) to LCDM (red curve) and a smooth curve without peaks (green curve). The χ^2/ndf, and the probability to exceed this value by chance, is noted for each model (χ^2 being calculated using the bandpower covariance matrix).

LCDM using

$$t(\ell) = s(\ell) + a\,v(\ell)\,\sin\left(2\pi\frac{\ell}{\ell_s} + p\right), \qquad (21)$$

where $s(\ell)$ is the smoothed version of the LCDM spectrum, $v(\ell)$ is the linear falloff, a is a re-scaling of the amplitude, and p is the phase. We then fit this model to the data—Figure 21 shows the results. We find $\ell_s = 306 \pm 10$, $p = 13° \pm 33°$, and $a = 0.86 \pm 0.17$, consistent with LCDM, as is shown in the right part of the figure. Using *WMAP TT* data, Page et al. (2003) found the spacing between the first and second peaks to be 315 ± 2. The consistency of peak phases and spacings between temperature data and these new QUaD *EE* results constitutes yet another confirmation of the acoustic oscillation paradigm of CMB anisotropies. Readhead et al. (2004) allowed only the phase and amplitude to be free parameters finding $p = 24 \pm 33°$—making this restriction we find $p = 9 \pm 13°$.

9.4. BB Limits and Comparison to other Experiments

The *BB* results shown in Figure 17 are consistent with zero sky power, and we therefore interpret these results as upper limits. To convert the observed values into confidence limits, we find the 95% integral point of the positive part of the bandpower probability density function (which is assumed to be Gaussian with mean and spread as in Figure 17). Figure 22 shows a comparison of these limits, and our other spectra, to published results from other experiments. In this figure *EE* is shown on a log scale, and hence to avoid clutter only bandpowers whose center value is more than twice the distance between the center value and the lower end of the 68% confidence limit are shown. For *TE*, *EE*, and *BB* QUaD breaks new ground—for *TT* the high ℓ precision is comparable to ACBAR, although the beam uncertainty is larger in the present analysis.

9.5. E and B maps

In Figure 17, it is clear that we detect dramatically more *E*-mode power on the sky than *B*-mode. Another way to visualize this is in the image plane. Having converted the apodized *Q* and *U* maps to *E* and *B* Fourier modes as described in Section 6,
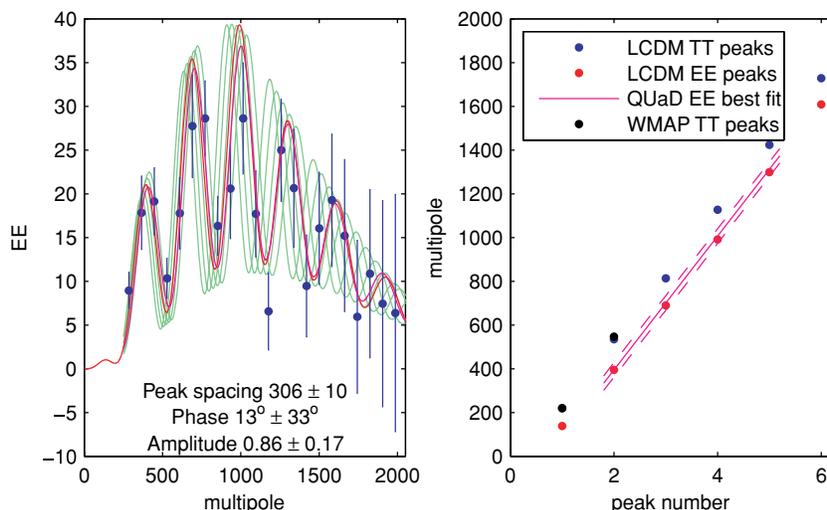
Figure 21. Left: fitting a "toy model" to the combined *EE* spectrum to determine the peak spacing, phase and amplitude, and the uncertainties thereon. The red line is the initial LCDM model, which is then smoothed and sinusoidally modulated with a range of peak spacings from 280 to 320 to generate the example family of curves shown in green. The best-fitting model curve is shown in magenta, and the associated parameter constraints indicated. Right: the location of the *TT* and *EE* peaks under LCDM as compared to the best-fit model and its uncertainty range. The *WMAP* points are from Page et al. (2003).
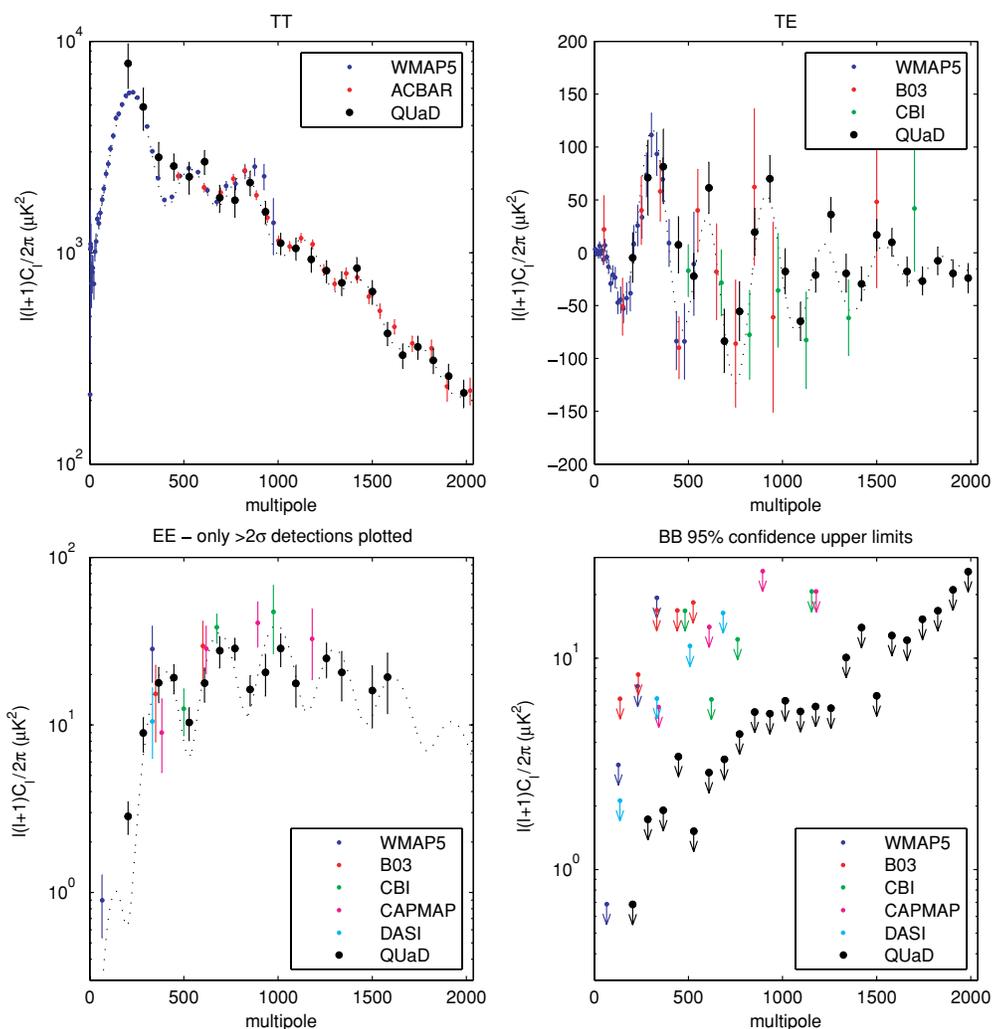
Figure 22. QUaD power spectra compared to results from *WMAP* (Nolta et al. 2009), ACBAR (Reichardt et al. 2009), B03 (Piacentini et al. 2006; Montroy et al. 2006), CBI (Sievers et al. 2007), CAPMAP (Bischoff et al. 2008), and DASI (Leitch et al. 2005). The *BB* upper limits are stated values where provided, and otherwise the 95% point of the positive part of the bandpower pdf.
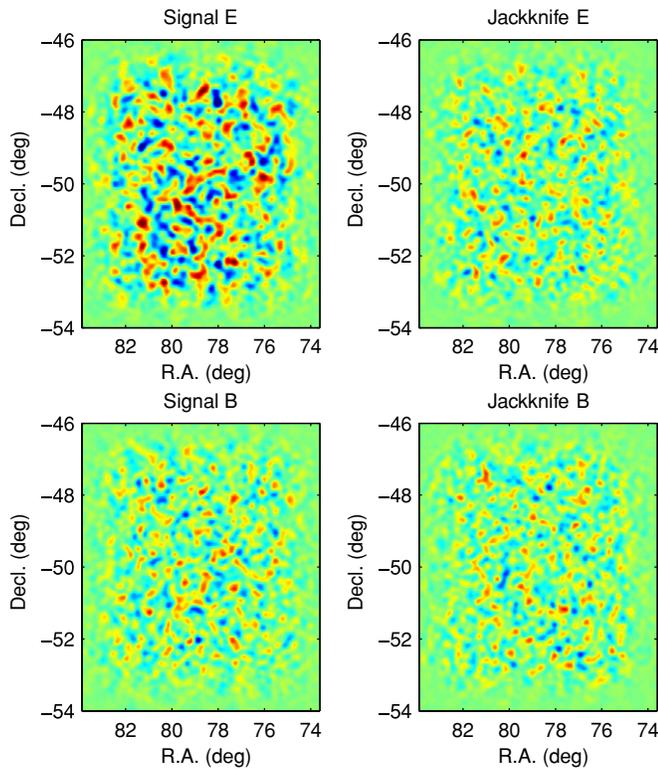
Figure 23. *E* and *B* signal and (deck) jackknife maps. The color scale is ±20 μK. These maps have been apodized and filtered to enhance signal to noise—see the text for details.

we can take the inverse Fourier transform to produce *E* and *B* maps. To enhance the signal to noise we apply a Fourier space filter.[17] Figure 23 shows the result—the *E* map contains far more structure than the *B* map. To confirm that the residual *B* signal is consistent with noise the figure also shows equivalent deck jackknife maps.

10. SYSTEMATIC ISSUES

There are a wide range of systematic effects which potentially mix *T* into pol. and/or *EE* into *BB*. In practice, as seen in Section 9.2, our results are consistent with the LCDM prediction of lensing *BB*, which is an order of magnitude smaller than our noise induced bandpower uncertainty, and hence effectively zero for the purposes of our experiment.

While there are many ways to produce false *BB* signal, it is virtually impossible that true *BB* power could somehow be canceled out through systematic effects. The fact that our *BB* results are consistent with zero is therefore powerful evidence that systematic mixing effects have been controlled to the required level of precision.

Note that from an instrumental point of view this is a fortunate accident of cosmology—the theoretical prediction is that the CMB sky presents us with a high purity "test pattern" against which we can validate our experiment. If we did see *BB* power at a level greater than the LCDM prediction extensive investigation of possible systematics would be required to attempt to verify that it was real. But since we do not see any *BB* further investigation of mixing systematics is arguably

---

[17] Here we use a Weiner filter assuming the LCDM *EE* spectrum, but a band pass filter $200 < \ell < 2000$ produces a qualitatively similar result.

unnecessary. However, for completeness, we present some additional discussion below.

10.1. EE to BB Mixing Due to the Sky Cut

Since the transform from *Q*, *U* to *E*, *B* is nonlocal, for less than full sky coverage leakage occurs from the *EE* spectrum to *BB* (and vice versa, although this is irrelevant under LCDM). In this analysis, we deal with this effect through the cross-spectra BPWFs discussed in Section 6.6. For LCDM, our *BB* expectation values peak at 0.7 μK² for $\ell = 360$ and fall rapidly to higher multipoles. At the current level of sensitivity including the cross-spectral BPWFs is just starting to become necessary— as we see from the increased χ^2 value going from the LCDM to null model *BB* panels in Figure 19.

Our signal only simulations include several other effects which potentially produce false *BB* signal (including beam centroid mismatch and grid angle uncertainty—see Section 5.1). However, for the standard simulation parameters we find that the sky cut effect is dominant since the mean *BB*-simulated spectrum follows the expectation values calculated using the cross-spectral BPWFs.

The original MASTER approach (Hivon et al. 2002) includes a "de-mixing" correction, and this was subsequently extended to polarization (e.g., Brown et al. 2005). Such techniques reduce spectral mixing in the mean, and hence also reduce the cross-spectra BPWFs. More recently, a technique was proposed by Smith & Zaldarriaga (2007) which results in much lower mixing from *E* to *B* within any given realization—we have implemented this in the flat-sky case, and confirmed with simulations that it works. However since this complication is unnecessary we do not include it in this paper.

10.2. Curved Sky Versus Flat Sky

The analysis presented here uses flat-sky power spectrum estimation while most recent CMB analyses have been conducted in the spherical harmonic basis. By using full curved sky maps as the input to our signal timestream simulation, our power estimation is normalized to recover the curved-sky input power. In addition, these simulations empirically test for problems associated with these flat-sky estimators, which are found to be negligible for our small ($\sim 6 \times 6°$) patch of sky. As mentioned above, *EE* to *BB* leakage is dominated by the cut sky effect and is well reproduced in the simulations by the cross-spectral BPWFs—for QUAD the use of flat-sky power estimators does not contribute significantly to this leakage.

10.3. Relative Gain Calibration

In Section 3.2, the elevation nod based method we use to equalize the detector gains was described. This method appears to be very accurate, and we do not include any scatter in the standard simulations. However, since error in the pair relative gains leads to *T* to pol. leakage there are a couple of issues which one might worry about.

The elevation nods integrate over the full beam including any far sidelobes, whereas when mapping the CMB, it is the ratio of the main lobe gains which we wish to know. Sky dips extending over a much larger zenith angle range (5°–45°) follow the expected sec(θ) dependence very closely (< 1% rms residual), indicating that the elevation nods are not distorted by sidelobe response.

The two detectors of each pair share a common feed horn and filters, but one might worry that the bandpasses of the fore-and-aft detectors might still differ. Since the atmospheric emission has a different frequency spectrum from the CMB this might lead to a systematic error in the relative gain within each pair.

To test how such errors would play out in practice we ran some special simulations where the input sky maps contain T only, and where the detector pair gains are deliberately systematically distorted such that $g_{fore}/g_{aft} = 1.03$—far worse than the $< 1\%$ constraint on possible mismatch which we derive from RCW38 observations. We find that the resulting EE and BB spectra follow the same form as the input TT spectrum with a peak of 0.4 μK^2 at $\ell = 200$. It is important to note that even for a systematic error like this there is still considerable averaging down as pairs of different angles, at different telescope orientations, leak T into, for instance, both $+Q$ and $-Q$ within a given map pixel. Making the gain ratio errors a random 3% across the focal plane (but fixed over time) the averaging down is much more effective and the peak leakage becomes 0.1 μK^2.

10.4. Pair Beam Mismatch

As mentioned in Section 5.1, we observe repeatable beam centroid offsets between the two halves of each detector pair, fixed in the instrument frame, with rms magnitude of $\sim 0'.1$. T to pol. leakage introduced by this effect is included in the standard simulations but makes a negligible contribution. However, for interest we run some special simulations with only T input introducing random pair centroid offsets with rms magnitude of $2'$—twenty times the observed value. We find that the resulting leakage has a broad peak at ~ 2 μK^2 around $\ell \sim 1000$ for 100 GHz, and ~ 1 μK^2 for 150 GHz. In addition, we run simulations under the totally unrealistic scenario that all pairs are systematically offset in the same direction by $1'$ and find a leakage of 6 μK^2 at $\ell \sim 1000$.

For the standard simulations we use the individual channel major/minor fit widths and orientation angles as measured using PKS0537–441. Some of the apparent variation between channels in these observations is measurement noise, but there is some real variation, and $\leqslant 10\%$ elongation. To investigate the impact of differential beam size, we run special T only simulations with both the major and minor axis FWHMs for one half of each pair systematically inflated by $1'.4$—the resulting leakage peaks at ~ 1.5 μK^2 above $\ell \sim 1500$. For differential elongation, where we inflate only one axis of one half of each pair (with a common orientation angle for all pairs), we find ~ 1 μK^2 above $\ell \sim 1500$.

10.5. Polarization Angle

As seen in Section 9.2, our results are consistent with LCDM—we detect considerable EE power and no BB. This is in itself confirmation that the assumed polarization angles of the detectors are known to sufficient accuracy. If we did see significant BB we might suspect that it was false signal due to incorrect angles. But it would be nearly impossible for a sky which truly had B-mode power to appear not to due to incorrect detector angles.

To confirm this we re-generated the real maps using detector angles systematically biased from the best estimate values by far more than the $1°$ estimated uncertainty (see Section 4.3 and the Instrument Paper). For a $5°$ bias there is almost no effect—the χ^2 PTE versus LCDM for the BB spectrum (as shown in Figure 19) falls from 0.56 to 0.21. Only at $10°$ do we start to

see significant extra signal in the BB bandpowers, and χ^2 failure versus LCDM for BB (and EB).

10.6. Polarization Efficiency

As mentioned in Section 4.3, and our Instrument Paper, we measure the polarization efficiency of our detectors to be $\epsilon = 0.08$ with an rms scatter of 0.015. An additional calibration factor $\gamma = (1 - \epsilon)/(1 + \epsilon)$ is then applied to the pair difference data. Random errors in ϵ will average down, while a systematic error will translate into a shift in the absolute calibration of the polarization power spectra ~ 4 times as large (including the additional doubling when going from units of temperature to power). Since we estimate the uncertainty on $\bar{\epsilon}$ to be < 0.02 the implied uncertainty on the polarization spectra is subdominant to the overall absolute calibration uncertainty.

10.7. Moon Pickup

QUaD has a variety of far sidelobes as described in the accompanying Instrument Paper. It is not clear which of these produces the bulk of the highly polarized ground pickup mentioned in Section 3.5—however since the field differencing is so effective at removing this contamination this is probably of academic interest only. However, any source of contamination which moves with respect to the ground will not be removed by field differencing.

Part of our basic low level data reduction infrastructure involves making single pair sum and difference maps for each 8 hr block of observations. On certain days when the Moon is high above the horizon, but at a very large angle from the telescope pointing direction, we see obvious stripes in these maps. Even after cutting these visibly contaminated days from the analysis, we saw a strong peak toward zero in the PTE distribution as shown in Figure 14, and some χ^2 values were much too large.

Extensive effort has been required to elucidate the coupling mechanism by which the Moon enters the CMB field data. We have determined that radiation reflects off the inside surface of the foam cone creating a polarized ring sidelobe at $\approx 100°$ from the main beam. The shape of the pickup across any given scan has been successfully modeled using the position of the Moon relative to the telescope pointing direction, the feed offset angle, and the detector polarization angle (see the Instrument Paper for details).

Using this model we cut data periods which are potentially contaminated by the Moon. For this analysis the cut is simple and quite aggressive—for any day where the Moon passes within a generous band about the ring sidelobe we simply discard the entire day. A future analysis could retain slightly more data by selectively cutting channels and using a time granularity shorter than a whole day. Performing this cut we reject an additional 59 days of observation but as seen in Section 7 the jackknife tests pass, meaning that we can be confident of the final power spectrum results.

11. CONCLUSIONS

We have described the observations, data reduction, simulation, and power spectrum analysis of the QUaD experiment. The results reported here are from 143 days of data taken in the second and third (final) seasons of observation, employing a conservative Moon cut, and simple lead–trail differencing. A future analysis may be able to include more data, and/or reduce the information lost in the ground removal.

The three sets of power spectra, 100 GHz, 150 GHz, and frequency cross, have been subjected to an extensive set of jackknife tests, and residual systematic contamination has been shown to be undetectable above the instrumental noise.

The combined spectra improve very considerably in sensitivity over previous results, and are consistent with LCDM—the standard cosmological model has passed yet another stringent test. Furthermore, we find that a smooth curve is a very poor fit to the observed *EE* spectrum—acoustic peaks in the *EE* spectrum have been detected with high significance for the first time. The impact of possible instrumental systematics has been considered in detail, but in fact the tight upper limits on *BB* power obtained are in themselves extremely powerful evidence that such effects are adequately controlled.

REFERENCES

Ade, P., et al. 2008, ApJ, 674, 22
Barkats, D., et al. 2005, ApJ, 619, L127
Bischoff, C., et al. 2008, ApJ, 684, 771
Brown, M. L., Castro, P. G., & Taylor, A. N. 2005, MNRAS, 360, 1262
Chiang, L.-Y., Naselsky, P. D., & Coles, P. 2008, Mod. Phys. Lett. A, 23, 1489
Dunkley, J., et al. 2009, ApJS, 180, 306
Fey, A. L., & Charlot, P. 2000, ApJS, 128, 17
Finkbeiner, D. P., Davis, M., & Schlegel, D. J. 1999, ApJ, 524, 867
Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, ApJ, 622, 759
Haslam, C. G. T., Klein, U., Salter, C. J., Stoffel, H., Wilson, W. E., Cleary, M. N., Cooke, D. J., & Thomasson, P. 1981, A&A, 100, 209
Hinderks, J., et al. 2009, ApJ, 692, 1221
Hivon, E., Górski, K. M., Netterfield, C. B., Crill, B. P., Prunet, S., & Hansen, F. 2002, ApJ, 567, 2
Hu, W. 2003, Ann. Phys., 303, 203
Hu, W., & White, M. 1997, New Astron., 2, 323
Jones, W. C., Bhatia, R., Bock, J. J., & Lange, A. E. 2003, in Proc. SPIE 4855, Millimeter and Submillimeter Detectors for Astronomy, ed. T. G. Phillips & J. Zmuidzinas (Bellingham, WA: SPIE), 227
Knox, L. 1999, Phys. Rev. D, 60, 103516
Kovac, J. M., Leitch, E. M., Pryke, C., Carlstrom, J. E., Halverson, N. W., & Holzapfel, W. L. 2002, Nature, 420, 772
Leitch, E. M., Kovac, J. M., Halverson, N. W., Carlstrom, J. E., Pryke, C., & Smith, M. W. E. 2005, ApJ, 624, 10
Leitch, E. M., et al. 2002, ApJ, 568, 28
Masi, S., et al. 2006, A&A, 458, 687
Montroy, T. E., et al. 2006, ApJ, 647, 813
Nolta, M. R., et al. 2009, ApJS, 180, 296
O'Sullivan, C., et al. 2008, Infrared Phys. Technol., 51, 277
Page, L., et al. 2003, ApJS, 148, 233
Page, L., et al. 2007, ApJS, 170, 335
Piacentini, F., et al. 2006, ApJ, 647, 833
Readhead, A. C. S., et al. 2004, Science, 306, 836
Reichardt, C. L., et al. 2009, ApJ, submitted (arXiv:0801.1491)
Seljak, U., & Zaldarriaga, M. 1997, Phys. Rev. Lett., 78, 2054
Sievers, J. L., et al. 2007, ApJ, 660, 976
Smith, K. M., & Zaldarriaga, M. 2007, Phys. Rev. D, 76, 043001
Spergel, D. N., et al. 2007, ApJS, 170, 377
Zaldarriaga, M., & Seljak, U. 2000, ApJS, 129, 431