The Astrophysical Journal, 688:1–11, 2008 November 20

THE SCALAR PERTURBATION SPECTRAL INDEX n_s: *WMAP* SENSITIVITY TO UNRESOLVED POINT SOURCES

K. M. Huffenberger,[1,2] H. K. Eriksen,[3,4] F. K. Hansen,[3,4] A. J. Banday,[5] and K. M. Górski[1,2,6]

Received 2007 October 10; accepted 2008 June 25

ABSTRACT

Precision measurement of the scalar perturbation spectral index, n_s, from the *Wilkinson Microwave Anisotropy Probe* (*WMAP*) temperature angular power spectrum requires the subtraction of unresolved point-source power. Here we reconsider this issue, attempting to resolve inconsistencies found in the literature. First, we note a peculiarity in the *WMAP* temperature likelihood's response to the source correction: cosmological parameters do not respond to increased source errors. An alternative and more direct method for treating this error term acts more sensibly, and also shifts n_s by $\sim 0.3\,\sigma$ closer to unity. Second, we re-examine the source fit used to correct the power spectrum. This fit depends strongly on the Galactic cut and the weighting of the map, indicating that either the source population or masking procedure is not isotropic. Jackknife tests appear inconsistent, causing us to assign large uncertainties to account for possible systematics. Third, we note that the *WMAP* team's spectrum was computed with two different weighting schemes: uniform weights transition to inverse noise variance weights at $l = 500$. The fit depends on such weighting schemes, so different corrections apply to each multipole range. For the Kp2 mask used in cosmological analysis, we prefer source corrections $A = 0.012 \pm 0.005\ \mu\mathrm{K}^2$ for uniform weighting and $A = 0.015 \pm 0.005\ \mu\mathrm{K}^2$ for N_{obs} weighting. Correcting *WMAP*'s spectrum correspondingly, we compute cosmological parameters with our alternative likelihood, finding $n_s = 0.970 \pm 0.017$ and $\sigma_8 = 0.778 \pm 0.045$. This n_s is only $1.8\,\sigma$ from unity, compared to the $\sim 2.6\,\sigma$ *WMAP* 3 year result. Finally, an anomalous feature in the source spectrum at $l < 200$ remains in the 3 year data, most strongly associated with the W band. We note the implications of these results for the 5 year data.

Subject headings: cosmic microwave background — cosmological parameters — cosmology: observations — methods: data analysis

Online material: color figures

1. INTRODUCTION

Measuring n_s, the spectral index of initial scalar fluctuations, which is scale invariant ($n_s = 1$) in the Harrison-Zeldovich model and slightly shallower in inflation models, is difficult, primarily because experimental systematics require control over a broad range of spatial scales. In inflation, the deviation from unity closely relates to the inflationary potential and the number of e-folds of expansion, so a statistically robust measurement of $n_s \neq 1$ places compelling constraints on the physics of the inflationary epoch.

Because all-sky measurements of the cosmic microwave background (CMB) access the largest observable scales in the universe, the angular power spectrum of the CMB, with a long lever arm, is crucial to such studies. Indeed, the latest data release from the *Wilkinson Microwave Anisotropy Probe* (*WMAP*) claims $\sim 2.6\,\sigma$ deviation from the Harrison-Zeldovich spectrum (Spergel et al. 2007). Unfortunately, the CMB is not a totally clean measurement. For example, the well-known degeneracy with the optical depth since recombination (τ) makes precision measurement of n_s impossible using CMB temperature anisotropies alone, and polarization is required to break it. Complicated noise properties and hints of unknown systematics in the *WMAP* measurement of large-scale polarization indicate that the systematic uncertainty in both τ and n_s should still be considered significant (Eriksen et al. 2007).

Another important, but underappreciated, complication for the measurement of n_s is additional power in the angular spectrum from unresolved, and unmasked, point sources. At high l, this shot noise can significantly bias the power spectrum, and consequently n_s. The *WMAP* team devised a sensible prescription for dealing with this contaminant: (1) use the spectral energy distribution measured from detected sources (distinct from the CMB) to infer it for undetected ones; (2) measure the contamination using multifrequency data; (3) correct the spectrum; and (4) marginalize over the measurement error when computing the likelihood (Hinshaw et al. 2003, 2007).

Huffenberger et al. (2004) found a level of source contamination consistent with the level in the first *WMAP* data release (Hinshaw et al. 2003). However, based on the 3 year temperature data (Hinshaw et al. 2007), Huffenberger et al. (2006) measured a point-source spectrum with two irregularities. First, at $l > 200$ the spectrum is white, but with an amplitude below the value in the original preprint of Hinshaw et al. (2007). In the present work, we discovered a small error in the power spectra used for the Huffenberger et al. (2006) estimate, which should have reported $A = 0.013 \pm 0.001\ \mu\mathrm{K}^2$ instead of $A = 0.011 \pm 0.001\ \mu\mathrm{K}^2$, still below the original *WMAP* value of $A = 0.017 \pm 0.002\ \mu\mathrm{K}^2$. Prompted by our result, Hinshaw et al. (2007) re-examined the issue, revising their value down somewhat and increasing the error bars, to $A = 0.014 \pm 0.003\ \mu\mathrm{K}^2$. The Spergel et al. (2007) bispectrum analysis indicates a non-Gaussianity consistent with these values, but lacks the statistical power of the multifrequency power spectrum comparison. The second peculiarity is that the power at

[1] Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena CA 91109; huffenbe@jpl.nasa.gov.
[2] California Institute of Technology, Pasadena, CA 91125.
[3] Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, N-0315 Oslo, Norway.
[4] Centre of Mathematics for Applications, University of Oslo, P.O. Box 1053 Blindern, N-0316 Oslo, Norway.
[5] Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, Postfach 1317, D-85741 Garching bei München, Germany.
[6] Warsaw University Observatory, Aleje Ujazdowskie 4, 00-478 Warszawa, Poland.

$100 < l < 200$ in Huffenberger et al. (2006) was inconsistent, at strong statistical significance, with the rest of the white spectrum.

This paper again considers the power spectrum source correction procedure in detail. We begin in § 2 with a study of the impact of the source correction on the scalar spectral index through the likelihood. Following this, we probe the source amplitude in § 3, examining the dependence of the fit on the sky weighting, mask, year of observation, and frequency dependence, and present our best estimates of the cosmological parameters. These same tests test the robustness of the $l < 200$ feature. Finally, we conclude in § 4.

2. SOURCE CORRECTION IMPACT ON SPECTRAL INDEX

The final *WMAP* temperature power spectrum is a noise-weighted combination of cross-spectra computed from V-band (61 GHz) and W-band (94 GHz) maps. Prior to the computation of the angular spectrum, a foreground model is removed from the maps. The individual cross-spectra are corrected for the sky mask, instrument beams, and point-source contamination before combination (see Hinshaw et al. 2007). The combined spectrum is folded into the likelihood calculation,[7] which interfaces to a Markov chain Monte Carlo code such as CosmoMC (Lewis & Bridle 2002), yielding parameter estimates.

Assuming that the angular spectrum for sources is white, we wish to explore the dependence of n_s on the size of the source correction. In Figure 1, we show the correction made by Hinshaw et al. (2007). The correction is not white, because the combined spectrum gets a larger contribution from the W band at higher l, due to the noise weighting. To change the amplitude of the correction, we simply scale. Later, we discuss the estimation of this amplitude.

For real data, including nonuniform sky coverage, Galactic foregrounds, and nonideal noise, the approximation of the likelihood function is a vital, challenging, and still open research question. Here we explore two methods. The standard *WMAP* method calculates the likelihood for a theoretical power spectrum (C_l) at $l > 30$ using the power spectrum estimated from the data ($\hat{C}_l$) and the covariance matrix ($\Sigma_{ll'}$) (Verde et al. 2003; Hinshaw et al. 2003, 2007). Schematically, we write this as

$$-2 \log L = \mathcal{L}(C, \hat{C}, \Sigma). \tag{1}$$

A Gaussian likelihood, where $\mathcal{L}$ is quadratic, is perhaps the simplest likelihood of this form. Verde et al. (2003) constructed a Gaussian plus lognormal (GLN) approximation that better matches the true likelihood in simulations, which we denote by

$$\mathcal{L} = \mathcal{L}_{\rm GLN}(C, \hat{C}, \Sigma) \qquad \text{(alternative method)}. \tag{2}$$

We label this as an "alternative method," because the *WMAP* likelihood code's implementation actually differs slightly. In the *WMAP* likelihood code, the matrix is split explicitly between the different sources of covariance,

$$\Sigma = \Sigma_0 + \Sigma_1, \tag{3}$$

where Σ_0 contains cosmic variance, noise variance, and mask mode-coupling, and Σ_1 contains beam and point-source correction uncertainty, most important at high l. The *WMAP* likelihood code computes

$$\mathcal{L} = \mathcal{L}_{\rm GLN}(C, \hat{C}, \Sigma_0) + \mathcal{L}_1(C, \hat{C}, \Sigma_1) \qquad (\textit{WMAP} \text{ method}), \tag{4}$$

where $\mathcal{L}_1$, an update to the likelihood, assumes the Gaussian form for $\mathcal{L}$, and is calculated using the Woodbury formula.

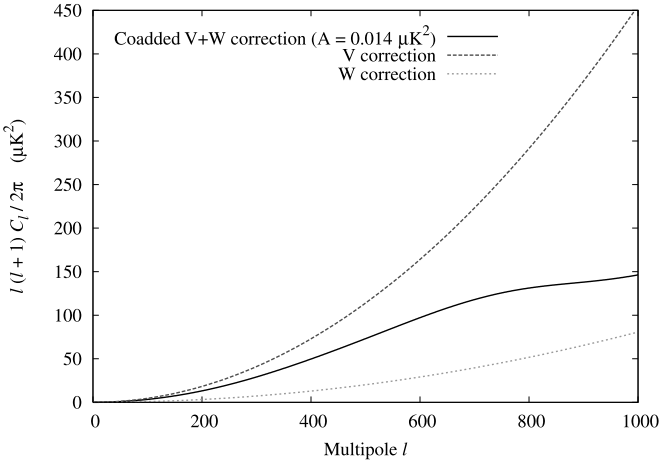

FIG. 1.— Size of the source correction in Hinshaw et al. (2007). [*See the electronic edition of the Journal for a color version of this figure.*]

Although the likelihood should be approximately Gaussian at high l, this assumption is perhaps flawed, for we find these two likelihood approaches (eqs. [2] and [4]) are not equivalent. Under the alternative method, the change in the GLN likelihood due to the inclusion of the beam/point-source term is -2.64, compared to $\mathcal{L}_1 = -1.22$ computed by the *WMAP* method, for the test theory spectrum included with the *WMAP* likelihood code (where $\mathcal{L} \sim 3541$). The discussions in Hinshaw et al. (2003, 2007) do not anticipate that this covariance split might alter the cosmological analysis, but we shall see that it does. However, recognizing that these approximate likelihoods are different does not demonstrate that either is acceptable, or which might be a closer approximation to the true likelihood.

The covariance split speeds the computation of that piece of the likelihood (since the beam uncertainty is well-approximated by a small number of modes), but here speed is not critical. It avoids the inversion of a $\sim 1000 \times 1000$ matrix per likelihood evaluation, but the low-l part of the code already inverts a more expensive $\sim 2000 \times 2000$ matrix. Performing the extra inversion and directly including the beam and source term with the other sources of error in the alternative likelihood (eq. [2]) is little additional burden.

We see the difference between the likelihood methods in the top panel of Figure 2, where we show the dependence of n_s on the source correction. We hold the errors fixed at the Hinshaw et al. (2007) value of $\sigma_A = 0.003\ \mu\mathrm{K}^2$, and marginalize over all the other parameters. Our alternative likelihood procedure shifts n_s higher by 0.005, or 0.3 σ. (The other parameters shift somewhat less, with σ_8 the next most sensitive.) For both likelihoods, as the source correction increases, the power spectrum at high l is lowered, and n_s decreases. Completely ignoring the source correction shifts n_s by ~ 0.01 higher, or about 0.6 σ.

Next we study the influence of error in the source correction. At very small and at very large error, we expect the parameter measurement to be independent of the error. For very small source error, the errors in other quantities dominate. For large errors, all modes that could be contaminated by point sources are effectively

[7] We are using version 2.2.2 of the *WMAP* likelihood, available at http://lambda.gsfc.nasa.gov/.

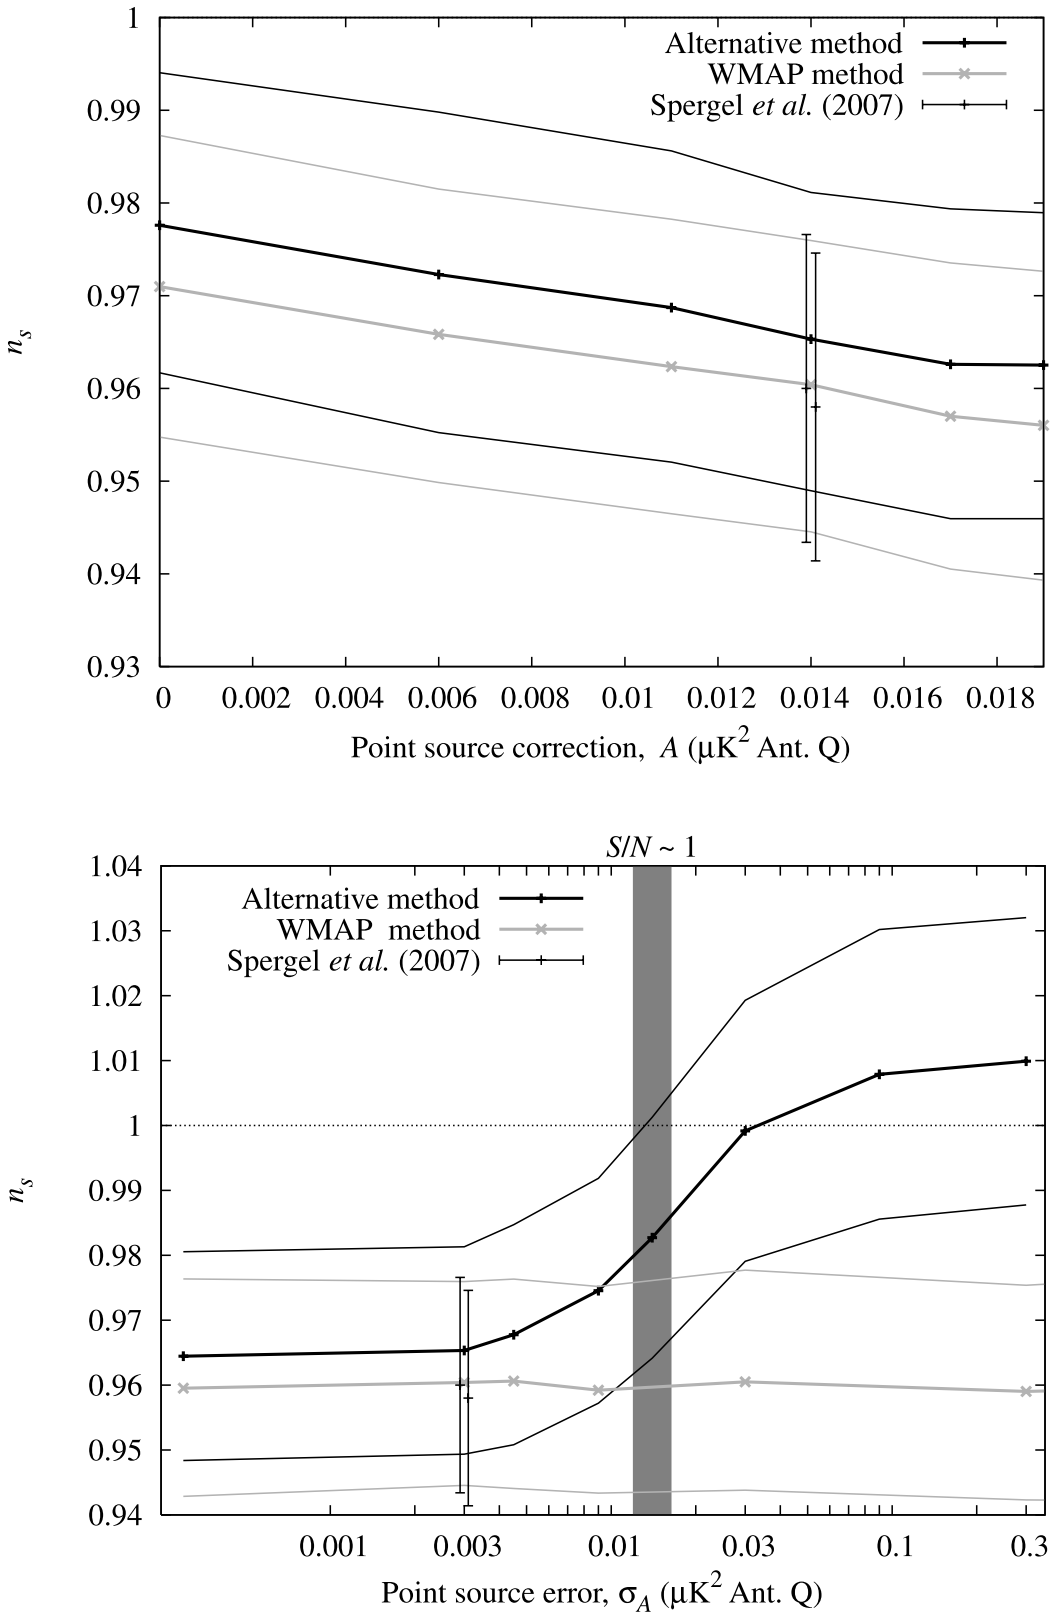

FIG. 2.—*Top*: Change in n_s for given source corrections, A. We compare the *WMAP* team's likelihood code to our modified version, which computes the high-l likelihood in a slightly different way. The error in the source correction is fixed at the Hinshaw et al. (2007) value used in Spergel et al. (2007). The thin lines bound the 68% probability interval. The Spergel et al. (2007) values are slightly offset horizontally for visibility, and either ignore (higher n_s) or include (lower) a correction for Sunyaev-Zeldovich effect contamination. The other n_s values ignore SZ. *Bottom*: Using the *WMAP* team's source correction, this plot shows the dependence of n_s on the point-source error, σ_A. [*See the electronic edition of the Journal for a color version of this figure*.]

projected out, and the parameter measurement is again independent. Near S/N $\sim$ 1, the measurement should undergo a transition, where the measurement error increases and (possibly) the mean value changes. In the bottom panel of Figure 2, we show the dependence of n_s on the error of the source-correction amplitude. The *WMAP* likelihood method shows a surprising result: the size of the point-source error has no effect on the measurement of n_s. (This holds true even when the point-source error rivals the size of the acoustic peaks.) The values of the likelihood change, but not the distribution of points in the Markov chain. This seems like a clear indication of a problem. The dependence of n_s on A in the top panel implies that at least the error bars should increase as σ_A increases.

On the other hand, our alternative likelihood shows more expected behavior. As the source-correction error is made very large, the errors increase by about 38%, and the mean value moves above 1. For this likelihood, the modes not subject to contamination by sources actually prefer $n_s > 1$, and to conclude that $n_s < 1$, the source contamination must be reliably measured. We also note that the errors on the source measurement do not make much difference, as long as $\sigma_A < 0.003$. In the next section, we examine this measurement in more detail.

3. UNRESOLVED POINT SOURCE SPECTRA

3.1. *Method*

The point-source spectrum can be estimated via a linear combination of the individual cross-spectra at several frequencies, a combination which projects out the CMB component. Huffenberger et al. (2004, 2006) examined the unresolved source component in *WMAP* data, using a generalized version of the method for the same task from Hinshaw et al. (2003).

In addition to the V and W bands used for cosmological measurement, the source analysis uses the Q band (41 GHz), because the contaminating sources are much brighter at lower frequencies. There are 276 cross-spectra in these three bands, accounting for all combinations of differencing assemblies per channel, and treating the 3 years of observation separately. These are combined as

$$A_L = \sum_{L'} \sum_{i \neq j} W^{ij}_{LL'} D^{ij}_{L'} \tag{5}$$

to give the point-source amplitude A_L in a multipole band denoted by L, where the weight is W^{ij}_L for the binned cross spectrum estimate D^{ij}_L, made from maps i and j. The map cross-spectra D^{ij}_L include the CMB power spectrum and the contribution from sources. As cross-spectra ($i \neq j$), they are noisy but lack a noise bias. The weights are based on the frequency independence of the CMB signal (in thermodynamic temperature units), the spectral energy distribution of the sources (measured for bright sources as $\beta \sim -2.0$, $S \propto \nu^{\beta+2}$; Hinshaw et al. 2003, 2007; Trushkin 2003), and the estimated noise covariance in the cross-spectrum measurements (see Huffenberger et al. 2004, 2006 for details).

The weights obey the constraint

$$\sum_{L'} \sum_{ij} W^{ij}_{LL'} = 0, \tag{6}$$

which means that no CMB will leak into the point-source estimate if the maps are properly calibrated and the instrumental beams are perfectly deconvolved from the spectra. If the source spectral energy distribution is correct, the weights also provide an unbiased estimate of the point-source power spectrum even if the noise covariance is wrong (although an incorrect noise covariance leads to suboptimal estimates and incorrect error bars).

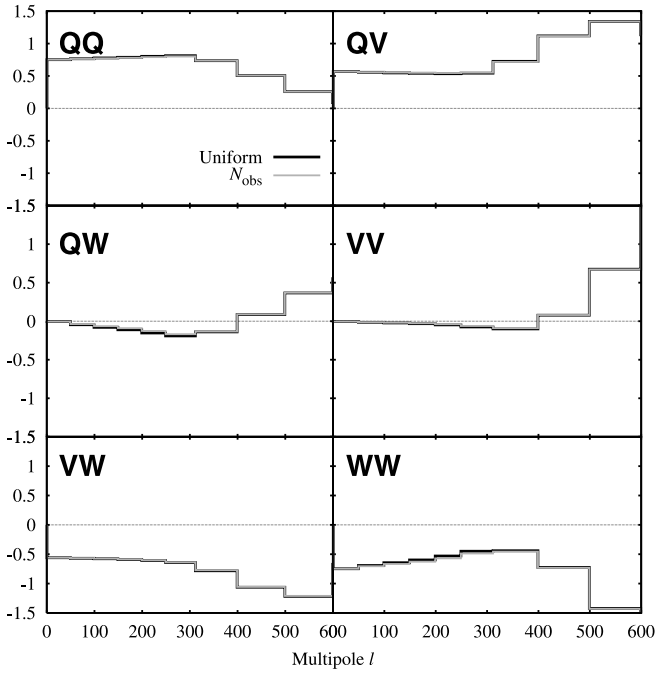

FIG. 3.— Weights for various cross-spectra in the source measurement, using the Kp2 mask, and two different map weightings. The individual weights for the 276 cross-spectra were summed together based on frequency coverage, yielding the 6 combinations shown. [*See the electronic edition of the Journal for a color version of this figure*.]

The covariance of the source power spectrum estimate is

$$\langle A_L A_{L'} \rangle = (\boldsymbol{W} \boldsymbol{\Sigma}_D \boldsymbol{W})_{LL'}, \tag{7}$$

where $\boldsymbol{\Sigma}_D$ is the estimate of the cross-spectrum covariance matrix. If the covariance matrix is diagonal in cross-spectra and multipole bin, then the weights are diagonal in multipole. Under this assumption, we plot some example weights in Figure 3.

The expected shot noise angular spectrum of sources in the *WMAP* data is flat in C_l, so we plot our measured spectra as l versus C_l. Throughout, we normalize the spectrum at the Q band in antenna temperature, which gets a larger signal than V or W.

In the following subsections, we describe several tests of the point-source measurement to explore the robustness of these features and their origin. Tests included changing the weights on the map for computing the spectrum, modifying the mask, fitting for the Galactic hemispheres separately, fitting year by year, and changing the assumed source spectral energy distribution.

3.2. *Spectrum at $l > 200$*

Whatever the cause of the excess power at $l < 200$ reported by Huffenberger et al. (2006), it is very likely not relevant for correcting the spectrum at high l. We therefore first concentrate on the correction required for the cosmological analysis, considering only $l \geq 200$, and then return to the anomalous low-l feature in § 3.3.

In a maximum likelihood estimate of the power spectrum, the map is weighted by the pixel-pixel inverse signal plus noise covariance. At small scales this estimate is computationally impractical, and Hinshaw et al. (2007) instead approximate it using two weighting schemes in the computation of the power spectra. In the signal-dominated regime at $l < 500$, they use spectra where every pixel is weighted evenly; in the noise-dominated regime at $l > 500$, the maps are inverse noise weighted (i.e., weighted by the number of observations, N_{obs}).

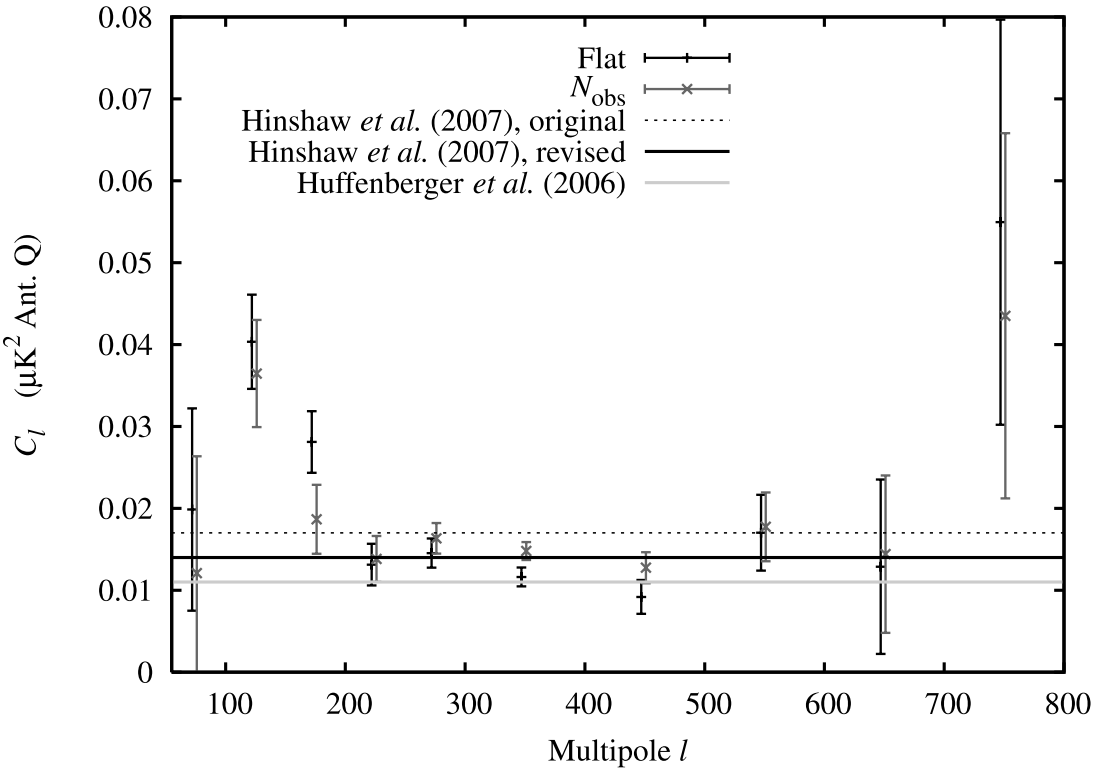

FIG. 4.—Fits for the unmasked source component, comparing spectra computed from maps using uniform weight outside of the mask with ones computed using N_{obs} (inverse variance) weighting, plotted as the Q-band amplitude. The lowest l-bin has no detection and is not shown. [*See the electronic edition of the Journal for a color version of this figure*.]

We compute the cross-spectrum using these two schemes and estimate the source contribution, plotted in Figure 4. The flat portion of the spectrum is notably higher in the N_{obs} weighted case. Specifically, for a white noise fit including $l > 200$, $A = 0.015 \pm 0.001$ μK^2 versus $A = 0.012 \pm 0.001$ μK^2 for the flat weighting. (The single-amplitude fits for all combinations we tried are compiled in Table 3.) This leads immediately to a key point: since the combined spectrum is built from two weighting schemes at different l's, the source correction must be different as well. Thus, we should have a smaller source correction for $l < 500$ and a larger correction for $l > 500$. The weighting scheme used by Hinshaw et al. (2007) is not specified.

This source level difference is a strong indication that the source population contaminating the *WMAP* spectrum is not isotropic. One possible scenario is as follows. Because source positions are stochastic, and source power strongly favors brighter sources

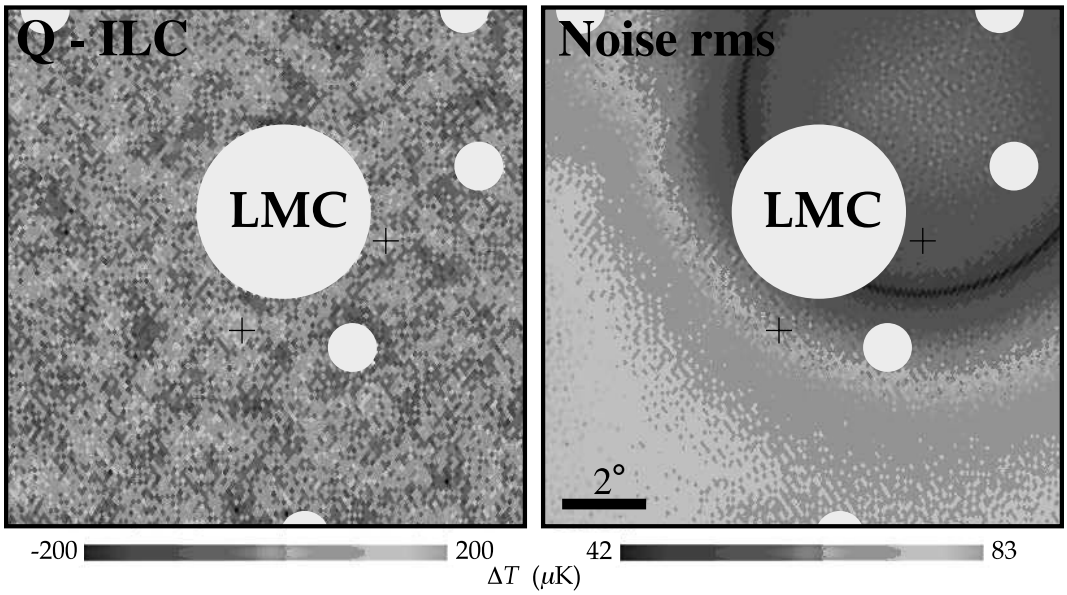

FIG. 5.—Difference map, Q−ILC, showing two pointlike objects near the LMC (*left*), compared to the Q-band noise rms for the same region (*right*), based on the number of observations reported by the *WMAP* team. [*See the electronic edition of the Journal for a color version of this figure*.]

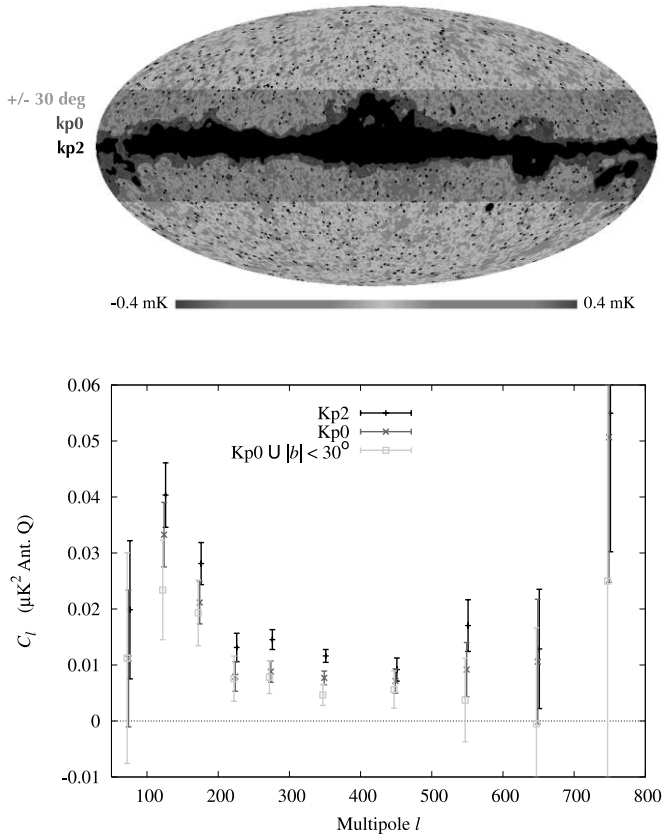

FIG. 6.—*Top*: Three masks used in this analysis. The smallest and most aggressive mask is the *WMAP* Kp2 mask, intermediate is Kp0, and most conservative is the union of Kp0 and a $|b| < 30^\circ$ cut. *Bottom*: The source power spectrum, computed with three different masks. As the mask becomes larger, the power declines, suggesting a concentration of unmasked source power near the Galactic plane. [*See the electronic edition of the Journal for a color version of this figure*.]

($C_l \propto \int dS\, S^2\, dN/dS$), brighter sources which, by chance, fell in the best-observed regions could boost the spectrum in the $N_{\rm obs}$ weighting over uniform weighting. However, the size of this effect is much too small to be a viable explanation. We computed power spectra of Monte Carlo realizations of a noiseless map containing isotropically distributed faint sources, based on differential source counts $dN/dS \propto S^{-2.3}$ (White & Majumdar 2004; Cleary et al. 2005), normalized to *WMAP* source counts at 1 Jy, setting an upper flux limit to reproduce a reasonable amount of power. Over 1000 simulations, for a bin $100 < l < 150$, the rms fluctuation of the difference in power between the uniform and $N_{\rm obs}$ spectra is 1.8%, and should be smaller at higher l. Thus, the odds are very slight that the change in the power is due to chance alignments of sources with the well-observed part of the sky; sources at the appropriate flux level are too numerous. Alternatively, the observed anisotropy in the source population could either be representative of the real sources, or indicate a problem with the masking procedure, with bright sources slipping through.

The $N_{\rm obs}$ weighting emphasizes the ecliptic poles, so these are the best places to look for suspect sources. In particular, visual inspection of a Q−ILC difference map (Fig. 5) shows two bright sources near the Large Magellanic Cloud, in the highly observed portion of the sky near the south ecliptic pole. These sources are as bright as some of the sources found by the *WMAP* source-detection algorithm, and indeed the brighter one is included in the subsequent catalogs of López-Caniego et al. (2007) and Nie & Zhang (2007) based on *WMAP* data. The presence of the nearby LMC may have interfered with the source-finding procedure.

TABLE 1
SOURCES FOUND IN Q−ILC

R.A. (J2000.0)	Decl. (J2000.0)	Galactic Longitude (deg)	Galactic Latitude (deg)	ID
00 43 14	−73 16 30	303.7	−44.0	PMN J0047−7308*,†
01 04 11	−72 08 02	301.5	−45.0	PMN J0059−7210*,†
04 51 56	−69 31 10	276.2	−33.7	...
05 20 56	−66 10 24	281.0	−35.6	PMN J0506−6109*,†
06 49 57	−16 56 12	227.9	−8.0	PMN J0650−1637*
20 50 59	+28 51 35	72.5	−9.7	...†
20 52 02	+31 55 18	75.1	−8.0	...
20 57 54	+31 29 08	75.6	−9.2	...†

NOTES.—Pointlike objects, found visually in a Q−ILC difference map. Units of right ascension are hours, minutes, and seconds, and units of declination are degrees, arcminutes, and arcseconds. Sources marked with asterisk and/or dagger (*,†) were found respectively in the catalog of López-Caniego et al. (2007) and Nie & Zhang (2007).

Changing the sky cut gives further indications of an anisotropic source population. In addition to the Kp2 cut used in the cosmological analysis, we recomputed the point-source fit using cross-spectra generated with two additional masks and uniform weighting: the more conservative Kp0 mask and a very conservative mask consisting of the union of Kp0 with a $|b| < 30^\circ$ Galactic cut (see Fig. 6).

As seen in Table 3, the source power drops significantly as we expand the masks, indicating that *unmasked* sources are brighter or more common near the plane. This raises two possibilities: either some of the sources are Galactic in origin, or sources are less efficiently found and masked near the plane.

From a visual inspection of the Q−ILC difference map, we find ∼6 more bright pointlike objects, many near the Galactic plane, and all in the southern hemisphere (see Table 1). Many of these objects have also already been noted in the López-Caniego et al. (2007) and Nie & Zhang (2007) catalogs. Computing the power spectrum in hemispheres (with the Kp2 mask), the source amplitudes straddle the value for the whole sky. At $l < 200$, the northern hemisphere has slightly less source power and the southern slightly more (See Fig. 7).

The *WMAP* source mask, built from a variety of catalogs (Hinshaw et al. 2003, 2007), may not cover the sky evenly, because of differences in sensitivity in different regions. It may be more difficult to identify sources amid large Galactic foregrounds.

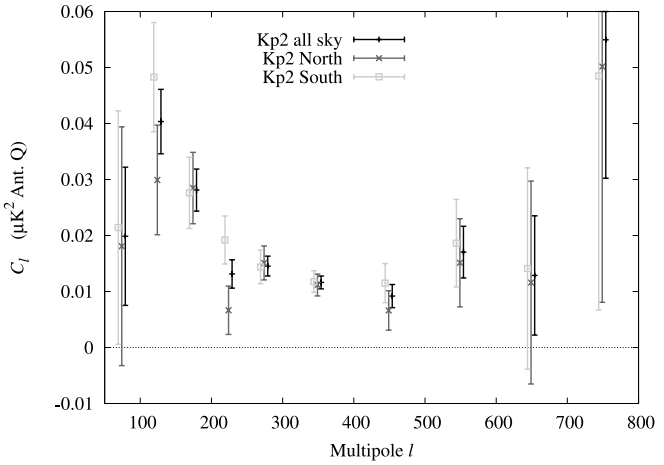

FIG. 7.—Source power spectrum, computed in hemispheres with the Kp2 mask. [*See the electronic edition of the Journal for a color version of this figure*.]

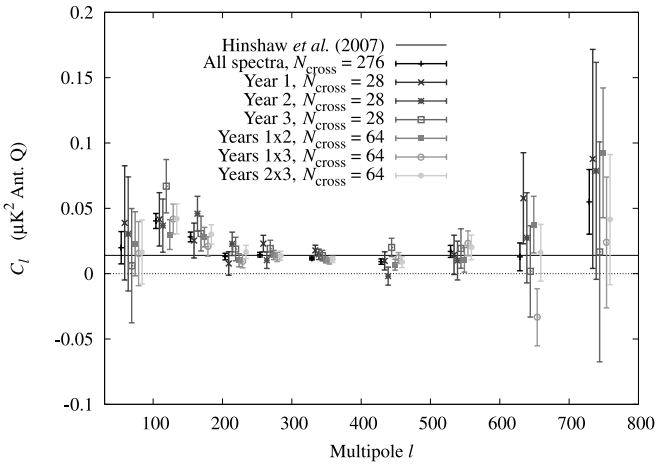

FIG. 8.—Year by year comparisons of the point-source power spectrum. [*See the electronic edition of the Journal for a color version of this figure.*]

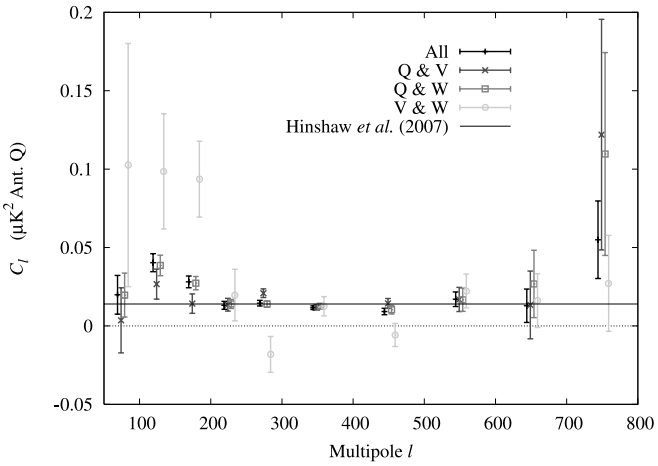

FIG. 9.—Source power spectrum estimate for combinations of bands. [*See the electronic edition of the Journal for a color version of this figure.*]

This would be particularly true for flat-spectrum sources in lower frequency radio surveys. The *WMAP* source mask has noticeable gaps near the Galactic plane, but outside the Kp2 and Kp0 masks.

Two jackknife tests, breaking the data into subgroups, also show some peculiarities. First, we divided up the cross spectra based on the years of observation (Fig. 8). Each individual year accounts for 28 cross-spectra, while a pair of years accounts for 64. At $l > 200$, the fits for individual years are always as greater than for pairs of different years. The largest estimate is the (year 3 $\times$ year 3) fit, at $A = 0.017 \pm 0.003$ μK^2, while the smallest is (year 1 $\times$ year 2), at 0.011 ± 0.002 μK^2. This seems unlikely to be due to chance, and could have a number of causes. For example, a slight cross-correlation between Q1 and Q2 or between Q and V, introducing a small noise bias in the cross-spectra, could have this effect. Removing either Q1 or Q2 from the source estimate, the amplitude drops to $A = 0.010 \pm 0.001$ μK^2.

However, this dependence on the Q band could also be due to the shape of the source spectrum. For the second jackknife test, we broke the cross-spectra up by band. We considered combinations with Q and V bands only, Q and W only, and V and W only. These spectra are shown in Figure 9.

At $l > 200$, the V+W combinations are very noisy, and have the lowest source fit, $A = 0.006 \pm 0.004$ μK^2, although the χ^2 is poor. The Q+V combination has the largest source amplitude, $A = 0.014 \pm 0.001$ μK^2. This discrepancy could mean that the source spectrum is incorrect: if the true spectrum is steeper than the $\beta = -2$ we have used, or the spectrum steepens between V and W, one could observe this effect. In this case, the Q+V combination gives the most accurate measurement of the Q-band source contamination, but this is not used in the cosmological analysis. To get the correction in the V and W bands, we are better off to scale from the lower amplitude, although even this is not completely satisfactory. Using the wrong spectrum for sources means that the shape of the source correction will be somewhat incorrect. We gradually steepened the source spectrum, and refit. The Q+V, Q+W, and V+W fits can be brought within 1 σ by setting $\beta = -2.5$. These can be made equal (at $A = 0.012$ μK^2) by setting $\beta \sim -2.9$, which compared to $\beta = -2.0$ would represent a 50% smaller amplitude in V and an 80% smaller amplitude in W. If true, the source correction should be set correspondingly smaller, driving n_s higher. (We continue to quote values of A in the Q band, assuming $\beta = -2.0$ scaling to V and W.)

Also implicit in this source-fitting method is that the same sources are the cause of the excess power in each of the bands. Statistical cleaning of multiple source populations from CMB maps can be difficult, and may require detailed information about the statistics and frequency dependence of each population (Huffenberger & Seljak 2005). If only one population is present, the cross-band power spectra due to sources will be (on average) the geometric means of the auto-band spectra. If multiple populations with distinct frequency scaling are present, sources in one band will not correlate with sources in the other, and the cross-band power estimates will be lower than the auto-band cross-correlations, leading to a poor estimate of the size of the correction. In *WMAP*, the estimate of the source power based only on cross-band data combinations (QV+QW+VW) is lower than the auto-band combinations (QQ+VV+WW; Table 3), but only at the 1.4 σ level, if we trust the statistical errors. This analysis, while suggestive, does not provide compelling evidence for multiple source populations which could disrupt the source correction, and the difference is within the error bars we adopt.

In the Appendix, we discuss the impact of beam errors. These tend to have little impact on the point-source spectrum at $l > 200$, since the resulting CMB leakage is large where the CMB is large, at lower l.

TABLE 2
COSMOLOGICAL PARAMETERS

Parameter	*WMAP*	Source/Likelihood Corrected
	WMAP Data Only	
$\Omega_b h^2$	0.0223 ± 0.0007	0.0220 ± 0.0008
Ω_m	0.237 ± 0.034	0.242 ± 0.037
h	0.735 ± 0.032	0.731 ± 0.034
τ	0.088 ± 0.030	0.090 ± 0.030
n_s	0.951 ± 0.016	0.968 ± 0.017
σ_8	0.742 ± 0.051	0.780 ± 0.052
	WMAP + ACBAR + BOOMERANG	
$\Omega_b h^2$	0.0232 ± 0.0007	0.0224 ± 0.0007
Ω_m	0.233 ± 0.034	0.234 ± 0.032
h	0.739 ± 0.033	0.742 ± 0.032
τ	0.088 ± 0.032	0.092 ± 0.030
n_s	0.951 ± 0.016	0.970 ± 0.017
σ_8	0.739 ± 0.051	0.778 ± 0.045

NOTES.—Comparison of marginalized parameter results obtained from Table 5 of Spergel et al. (2007) (col. [2]) and after our modifications to the source correction and likelihood (col. [3]).

TABLE 3
Fits for Source Amplitude

Mask	Weight	Subset	A (μK^2)	σ_A	χ^2/dof
		Diagonal Covariance			
30kp0.....	flat	all	0.006	0.001	0.25
kp2.....	flat	VW	0.006	0.004	1.99
kp0.....	flat	all	0.008	0.001	0.54
kp2.....	flat	QV+QW+VW	0.010	0.001	1.09
kp2.....	flat	noQ1	0.010	0.001	2.02
kp2.....	flat	noQ2	0.010	0.001	0.71
kp2.....	flat	yr12	0.011	0.002	1.02
kp2.....	flat	yr13	0.011	0.002	1.10
kp2.....	flat	all (N. hemi.)	0.011	0.001	0.92
kp2.....	flat	all	0.012	0.001	1.39
kp2.....	flat	yr22	0.012	0.003	1.16
kp2.....	flat	yr23	0.012	0.002	0.50
kp2.....	flat	QQ+VV+WW	0.013	0.001	0.79
kp2.....	flat	QW	0.013	0.001	0.70
kp2.....	flat	all (S. hemi.)	0.013	0.001	0.68
kp2.....	flat	noW1W2	0.013	0.001	2.48
kp2.....	flat	noW3W4	0.013	0.001	1.23
kp2.....	flat	QV	0.014	0.001	1.65
kp2.....	nobs	all	0.015	0.001	0.69
kp2.....	flat	yr11	0.017	0.003	0.87
kp2.....	flat	yr33	0.017	0.003	0.20
		Diagonal + Beams Errors			
kp2.....	flat	noQ1	0.009	0.001	1.87
kp2.....	flat	VW	0.010	0.004	2.07
kp2.....	flat	noQ2	0.010	0.001	0.73
kp2.....	flat	yr12	0.011	0.002	1.05
kp2.....	flat	yr13	0.011	0.002	1.10
kp2.....	flat	yr22	0.011	0.003	1.11
kp2.....	flat	yr23	0.011	0.002	0.34
kp2.....	flat	all	0.012	0.001	1.53
kp2.....	flat	QW	0.013	0.001	0.72
kp2.....	nobs	all	0.015	0.001	0.79
kp2.....	flat	QV	0.016	0.001	2.30
kp2.....	flat	yr33	0.017	0.003	0.25
kp2.....	flat	yr11	0.019	0.003	0.85

Notes.—Fits for source power spectrum amplitude, including data at $l > 200$. A values normalized to the Q band antenna temperature.

Finally, we must decide by how much to correct the combined *WMAP* spectrum. There is a strong detection of unresolved power in the *WMAP* spectra, but as we have seen through the above tests, our knowledge of the character of these sources is poor. Our best-fit amplitudes including all the data are $A = 0.012$ μK^2 for flat weighting and $A = 0.015$ μK^2 for $N_{\rm obs}$ weighting. But the spectral effects (from considering V and W without Q) indicate it might be much lower, while individual year fits are much higher. We then are forced to the unsatisfying course of artificially inflating the error bars beyond their nominal statistical values, in order to account for these possibilities. Therefore, we set the error on the source estimates at $\sigma_A = 0.005$ μK^2.

To obtain our final estimate for cosmological parameters, we correct the spectrum for point sources using the two amplitudes quoted above (uniform weights for $l < 500$ and $N_{\rm obs}$ weights for $l > 500$), and our modified likelihood code. The marginalized parameters from the resulting Markov chains are given in Table 2, both considering *WMAP* data alone and including data from ACBAR (Kuo et al. 2004) and BOOMERanG (Jones et al. 2006; Montroy et al. 2006; Piacentini et al. 2006). Driven by the alternative likelihood and dual correction, this new value of $n_s = 0.970 \pm 0.017$ is only 1.8 σ away from 1. As another consequence, σ_8 increases to 0.778 ± 0.045. (Had we retained the smaller *WMAP* errors, $\sigma_A = 0.003$ μK^2, we would get $n_s = 0.968 \pm 0.017$, 1.9 σ away from 1.)

3.3. *Excess at $l < 200$*

We now turn to the $l < 200$ feature in the source power spectrum (Fig. 4), which is inconsistent with a white spectrum. This is present in both the uniform and $N_{\rm obs}$ weighted spectra, and in each year and pair of years. Because of its shape and prominence at low l, we initially considered two specific explanations. The first possibility concerned misestimation of the overall multiplicative map calibration of each DA map: if two maps are calibrated differently, the weights W_L^{ij} in equation (5) would not cancel the CMB component precisely. Thus, one would observe a leakage from the CMB signal into the point-source spectrum, with a signature resembling the CMB power spectrum. A similar effect would be caused by beam uncertainties. In the Appendix, we present the formalism to take these uncertainties into account in

the method described in § 3.1, and the results from the corresponding analysis of the *WMAP* data are presented in the bottom section of Table 3. The conclusion from these computations is that the uncertainties quoted by the *WMAP* team for the calibration and beam errors are too small to explain this effect.

A third hypothesis is residual Galactic foregrounds, which should show through the mask and frequency dependency. As the Galactic cut widens from Kp2 to the wide cut (Fig. 6), the bin from $l = 100$ to 150 drops about 0.017 μK^2. At the same time, the fit for the white-noise level drops by 0.006 μK^2. Subtracting off the white-noise levels for each, the component in the excess has dropped by about 40%. The power in the excess is still significant, even for this broad cut.

Next, we turn to the estimates using only two bands. In this case, we see that the feature is strongly enhanced in the V+W combinations, about the same in the Q+W combination, but clearly diminished in the Q+V combination. The latter appears consistent with the flat-source spectrum. This may indicate that the excess is associated with the W band. On the other hand, because the W-band spectra tend to carry negative weight in these estimates (Fig. 3), this excess would represent a deficit of power in W, which is peculiar. It may indicate an oversubtraction of the Galactic foreground template in the W band, which could have consequences for the cosmological analysis. At this point, it is difficult to be definitive.

4. CONCLUSIONS

We have reanalyzed the source-correction procedure for the 3 year data release of *WMAP*. First, we considered the impact of this procedure in the *WMAP* likelihood code. Surprisingly, we found that the *WMAP* likelihood does not react to changes in the point-source correction error. We explore an alternative likelihood which does respond as expected, although we note that more work is needed to validate this approach, as it couples to the important problem of how to approximate the likelihood with fitting formulas over a wide multipole range. To conclude $n_s < 1$, a precision measurement of the source contamination is required. We note that the modes not contaminated favor n_s consistent with unity.

Second, we found several indications that the unmasked source population in the *WMAP* data is anisotropic. This implies that the combined spectrum should be corrected differently in two multipole regions, based on the weighting of the map. Anisotropy in the unmasked sources is unexpected, but can be turned to an advantage: by very carefully masking near the ecliptic poles and galaxy, or employing a wide Galactic cut, the point-source contamination can be cut substantially. This gain must be weighed against the reduction of the sky area.

We note irregularities in jackknife tests of the source fit, grouped by time of observation or frequency band. This prompted us to adopt large errors on the source fit, to account for systematics beyond the estimate of statistical error which accompanies our measurement. This step is necessary to treat the source correction conservatively. The modified likelihood, reduced source amplitude from ignoring the excess at $l < 200$, dual correction due to map weighting, and the enlarged error bars are responsible for raising our values of n_s and σ_8.

Finally, the previously noted anomalous $l < 200$ feature is still present, shows signs of being spatially associated with the Galaxy, and is most strongly associated with the W band. It may represent an oversubtraction of the foreground template in W, although further investigation is warranted. However, the immediate conclusion is that this part of the spectrum should not be used to infer the point source amplitude at higher l's.

5. NOTE ON THE *WMAP* 5 YEAR RELEASE

While this paper was in review, the *WMAP* 5 year data was released, and we look forward to examining the new data set. The correction methodology for unresolved point sources is the same as for the 3 year data (Nolta et al. 2008), so the broad conclusions of our analysis, regarding the likelihood, source masking, and use of two corrections, should be similar. With more data, the catalog of detected and masked sources improved (Wright et al. 2008), which makes the source correction slightly smaller. More of the Galactic plane has been masked (Gold et al. 2008), which should also help, as noted above.

In response to a preprint of this work, on the issue of the likelihood form, Nolta et al. (2008) show by explicit integration that the *WMAP* approximation yields the same likelihood value as a Gaussian treatment of the beam and point-source errors for a test spectrum, although this does not directly address the difference between the two likelihood methods (eqs. [2] and [4]), or necessarily show which is closer to the correct likelihood. They employ another Gaussian-based modified likelihood, which yields parameter values similar to their approximation, but responds more sensibly to increases in the source and beam error.

In addition, a substantial revision in the beams (Hill et al. 2008) seems to have eliminated the source anomaly at $l < 200$ (Nolta et al. 2008), which almost certainly means that it represented leakage from the CMB due to beam errors.

We thank the *WMAP* team for useful discussions and for providing additional data. In particular we thank Gary Hinshaw, Michael Nolta, and Lyman Page, who suggested examining the effect of beam uncertainties. We acknowledge the helpful suggestion of an anonymous referee to consider the effect of multiple source populations. HEALPix software (Górski et al. 2005) was used to deriving some results in this paper. We also acknowledge use of the Legacy Archive for Microwave Background Data Analysis (LAMBDA). H. K. E. acknowledges financial support from the Research Council of Norway. This work was partially performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with NASA.

APPENDIX

LEAKAGE FROM CALIBRATION AND BEAM ERRORS

In this appendix, we estimate the leakage from CMB into the point source estimate. We consider two sources of leakage, the overall calibration of the map and the errors in the beam deconvolution. Each of these can cause the CMB spectrum to have a slightly different amplitude or shape in each of the cross-spectra. This means that the CMB will not cancel itself in the weighted linear combination to give the source estimate. This leakage can be accounted for in the cross-spectrum covariance. The calibration is a constant function of l, and the beam error is nearly so over a wide range of scales, so the shape of the leakage term strongly follows the CMB power spectrum, which drops rapidly with l.

The calibration uncertainty is simpler, so we begin there. In the signal-dominated (s) regime, without including point sources, we can model the calibration errors with

$$d_i = (1 + g_i)s. \tag{A1}$$

where the calibration in map d_i is quantified with dimensionless factors g_i, with $\langle g_i \rangle = 0$ and $\langle g_i g_j \rangle = \sigma_g^2 \delta_{ij}$. Then the estimated cross-spectra are

$$D_l^{ij} = (1 + g_i + g_j + g_i g_j)\hat{C}_l, \tag{A2}$$

where $\hat{C}_l$ is the power spectrum of the particular sky realization (although in the actual calculation, we substitute the *WMAP* best-fit spectrum). In the cross-spectrum $i \neq j$, so this is an unbiased estimator, $\langle D_l^{ij} \rangle = \hat{C}_l$. The covariance is

$$\mathrm{Cov}(D_l^{ij}, D_{l'}^{pq}) = \hat{C}_l \hat{C}_{l'} \left[\sigma_g^2 (\delta_{ip} + \delta_{iq} + \delta_{jp} + \delta_{jq}) + \sigma_g^4 (\delta_{ip}\delta_{jq} + \delta_{iq}\delta_{jp}) \right]. \tag{A3}$$

Essentially, the variance only gets a contribution in terms where the two cross-spectra share a map.

The point-source estimate is

$$A_l = \sum_{l'} \sum_{(ij)} W_{ll'}^{ij} D_{l'}^{ij} = \sum_{l'} \hat{C}_{l'} \sum_{(ij)} W_{ll'}^{ij} (g_i + g_j + g_i g_j), \tag{A4}$$

where one term has dropped out because the weights sum to zero. Because we have not included any source contribution, we have $\langle A_l \rangle = 0$, where the ensemble average is over the calibration factors. For plotting, its useful to compute the variance due to the calibration factors

$$\langle A_l^2 \rangle = \sum_{l'l''} \hat{C}_{l'} \hat{C}_{l''} \sum_{(ij)(pq)} W_{ll'}^{ij} W_{ll''}^{pq} \left[\sigma_g^2 (\delta_{ip} + \delta_{iq} + \delta_{jp} + \delta_{jq}) + \sigma_g^4 (\delta_{ip}\delta_{jq} + \delta_{iq}\delta_{jp}) \right]. \tag{A5}$$

A word about binning the spectrum is appropriate here. In practice, we bin the spectrum because the large number of cross-spectra (276) and multipoles ($\sim$800) slows the computation, and the signal-to-noise ratio per multipole is low. We can define a binning matrix G_{Ll} which averages quantities in non-overlapping bands indexed by L, and compute our estimate from the binned cross spectra $\sum_l G_{Ll} D_l^{ij}$. The quantity we are now computing is the variance in a bin. Since the weights are also computed in bands, this reduces to replacing all l with L, and replacing each power spectrum term $\hat{C}_l$ with $\sum_l G_{Ll} \hat{C}_l$.

In Jarosik et al. (2007), the calibration uncertainty is quoted as 0.5%, so $\sigma_g = 0.005$. As written, we are assuming the all maps are uncorrelated, as is the case if the calibration fluctuations are dominated by noise, as expected (L. Page 2007, private communication). In this case, fluctuations in the calibration are fairly small compared to the source estimate. If the calibration is dominated by something else, say foregrounds, then whole bands could be correlated, and the calibration leakage can be substantial compared to the source fit.

In a similar way, we can estimate the beam uncertainties. We begin with

$$d_{lm}^i = \hat{B}_l^i (1 + E_l^i) s_{lm}, \tag{A6}$$

where $\hat{B}_l^i$ is an unbiased measurement of the beam, and E_l^i is the (small) fractional difference between the true beam and the measured beam. We assume all beams have independent errors, and define a δ-function-like object δ_{ij}^B, which is unity when maps i and j share a beam, and zero otherwise. With this definition, we can describe the beam error with a Gaussian distribution with $\langle E_l^i \rangle = 0$ and

$$\langle E_l^i E_l^j \rangle = \Sigma_{ll'}^{B,i} \delta_{ij}^B. \tag{A7}$$

That is, the beam errors are correlated in l but not between beams. The beam-deconvolved cross spectra are

$$D_l^{ij} = (1 + E_l^i)(1 + E_l^j)\hat{C}_l. \tag{A8}$$

Following the same procedure as with gain calibrations, we find a similar expression for the covariance:

$$\mathrm{Cov}(D_l^{ij}, D_{l'}^{pq}) = \hat{C}_l \left[\Sigma_{ll'}^{B,i} (\delta_{ip}^B + \delta_{iq}^B) + \Sigma_{ll'}^{B,j} (\delta_{jp}^B + \delta_{jq}^B) + \Sigma_{ll'}^{B,i} \Sigma_{ll'}^{B,j} (\delta_{ip}^B \delta_{jq}^B + \delta_{iq}^B \delta_{jp}^B) \right] \hat{C}_{l'}. \tag{A9}$$

Hinshaw et al. (2003, 2007) found that the beam uncertainty can be well represented as a small number of orthogonal modes ($\sim$10):

$$\Sigma_{ll'}^{B,i} = \sum_r U_{rl}^i U_{rl'}^i, \tag{A10}$$

where r denotes the mode. Here the binning is a little more complicated than for the calibration uncertainty. Instead of simply $\hat{C}_l$, the quantities that must be binned are $\hat{C}_l U_{rl}^i$ (for the first-order terms) and $\hat{C}_l U_{rl}^i U_{tl}^j$ (for the second-order term).

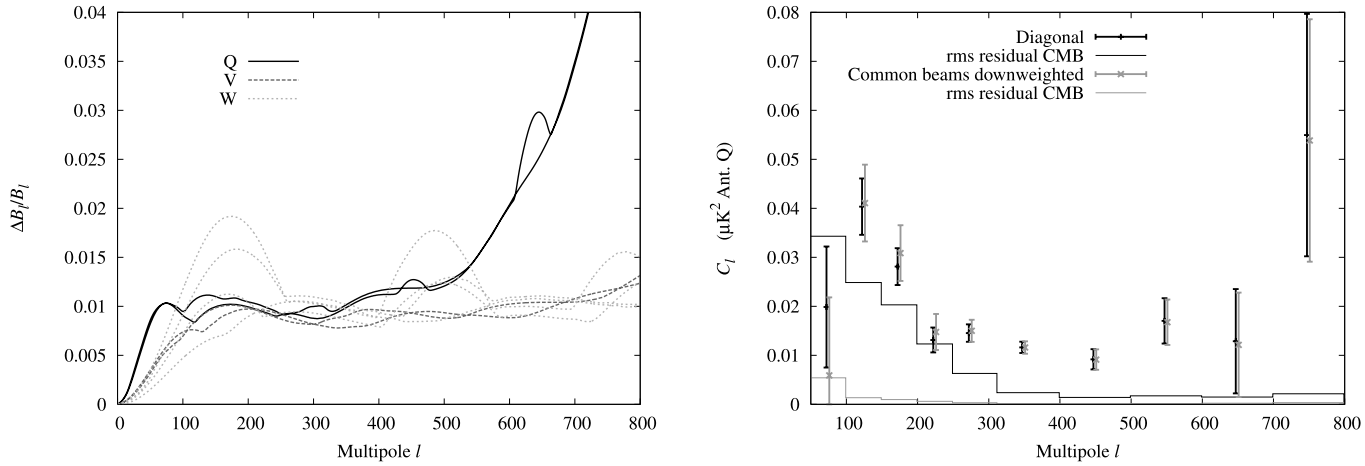

FIG. 10.—*Left*: Beam errors taken from Jarosik et al. (2007). The Q-band errors rise smoothly to ~0.08 at $l = 800$. *Right*: rms leakage of the CMB into the point-source measurement, with and without taking the contribution to the beam covariance into account. [*See the electronic edition of the Journal for a color version of this figure*.]

At this time, the beam modes for each *WMAP* differencing assembly are not public. (Only the modes for the combined spectrum are included in the likelihood code.) To approximate the beam modes, we take the beam errors from Jarosik et al. (2007; shown in Fig. 10). We treat each of these as a single mode for the associated beam, correlating all the multipoles in that beam. This should provide a conservative estimate of the beam covariance, although it does not capture all of its properties.

With these approximate modes, we can compute the beam variance and the rms CMB leakage for any set of weights. This is plotted in the right panel of Figure 10. For our standard diagonal covariance, the rms beam leakage at first looks like a promising explanation for the $l < 200$ excess. It has a similar shape, and strong correlations bin-to-bin. However, when we recompute the weights taking the beam covariance into account, we produce a similar estimate for point sources in a combination that allows very little CMB leakage. Given the estimated size of the rms residual CMB, it is surprising that the source estimates are so similar. Perhaps this is an indication that the beam covariance is badly estimated. In any case, either the $l < 200$ is not due to beams, *WMAP*'s quoted beam errors are incorrect, or our approximation for the beam errors is very poor.

REFERENCES

Cleary, K. A., et al. 2005, MNRAS, 360, 340
Eriksen, H. K., Huey, G., Banday, A. J., Górski, K. M., Jewell, J. B., O'Dwyer, I. J., & Wandelt, B. D. 2007, ApJ, 665, L1
Gold, B., et al. 2008, ApJS, submitted (arXiv:0803.0715)
Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, ApJ, 622, 759
Hill, R. S., et al. 2008, ApJS, submitted (arXiv:0803.0570)
Hinshaw, G., et al. 2003, ApJS, 148, 135
———. 2007, ApJS, 170, 288
Huffenberger, K. M., Eriksen, H. K., & Hansen, F. K. 2006, ApJ, 651, L81
Huffenberger, K. M., & Seljak, U. 2005, NewA, 10, 491
Huffenberger, K. M., Seljak, U., & Makarov, A. 2004, Phys. Rev. D, 70, 063002
Jarosik, N., et al. 2007, ApJS, 170, 263
Jones, W. C., et al. 2006, ApJ, 647, 823
Kuo, C. L., et al. 2004, ApJ, 600, 32
Lewis, A., & Bridle, S. 2002, Phys. Rev. D, 66, 103511
López-Caniego, M., González-Nuevo, J., Herranz, D., Massardi, M., Sanz, J. L., De Zotti, G., Toffolatti, L., & Argüeso, F. 2007, ApJS, 170, 108
Montroy, T. E., et al. 2006, ApJ, 647, 813
Nie, J.-Y., & Zhang, S.-N. 2007, Chinese J. Astron. Astrophys., 7, 199
Nolta, M. R., et al. 2008, ApJS, submitted (arXiv:0803.0593)
Piacentini, F., et al. 2006, ApJ, 647, 833
Spergel, D. N., et al. 2007, ApJS, 170, 377
Trushkin, S. A. 2003, Bull. Special Astrophys. Obs., 55, 90
Verde, L., et al. 2003, ApJS, 148, 195
White, M., & Majumdar, S. 2004, ApJ, 602, 565
Wright, E. L., et al. 2008, ApJS, submitted (arXiv:0803.0577)