

COMPARING *PLANCK* AND *WMAP*: MAPS, SPECTRA, AND PARAMETERS

D. LARSON¹, J. L. WEILAND¹, G. HINSHAW², AND C. L. BENNETT¹

Draft version March 30, 2015

ABSTRACT

We examine the consistency of the 9 yr *WMAP* data and the first-release *Planck* data. We specifically compare sky maps, power spectra, and the inferred Λ cold dark matter (Λ CDM) cosmological parameters. Residual dipoles are seen in the *WMAP* and *Planck* sky map differences, but their amplitudes are consistent within the quoted uncertainties, and they are not large enough to explain the widely noted differences in angular power spectra at higher l . We remove the residual dipoles and use templates to remove residual galactic foregrounds; after doing so, the residual difference maps exhibit a quadrupole and other large-scale systematic structure. We identify this structure as possibly originating from *Planck*'s beam sidelobe pick-up, but note that it appears to have insignificant cosmological impact. We develop an extension of the internal linear combination technique to find the minimum-variance difference between the *WMAP* and *Planck* sky maps; again we find features that plausibly originate in the *Planck* data. Lacking access to the *Planck* time-ordered data we cannot further assess these features. We examine Λ CDM model fits to the angular power spectra and conclude that the $\sim 2.5\%$ difference in the spectra at multipoles greater than $l \sim 100$ are significant at the $3\text{--}5\sigma$ level, depending on how beam uncertainties are handled in the data. We revisit the analysis of *WMAP*'s beam data to address the power spectrum differences and conclude that previously derived uncertainties are robust and cannot explain the power spectrum differences. In fact, any remaining *WMAP* errors are most likely to exacerbate the difference. Finally, we examine the consistency of the Λ CDM parameters inferred from each data set taking into account the fact that both experiments observe the same sky, but cover different multipole ranges, apply different sky masks, and have different noise. We find that, while individual parameter values agree within the uncertainties, the six parameters taken together are discrepant at the $\sim 6\sigma$ level, with $\chi^2 = 56$ for 6 degrees of freedom (probability to exceed, PTE = 3×10^{-10}). The nature of this discrepancy is explored: of the six parameters, χ^2 is best improved by marginalizing over $\Omega_c h^2$, giving $\chi^2 = 5.2$ for 5 degrees of freedom. As an exercise, we find that perturbing the *WMAP* window function by its dominant beam error profile has little effect on $\Omega_c h^2$, while perturbing the *Planck* window function by its corresponding error profile has a much greater effect on $\Omega_c h^2$.

Subject headings: cosmic background radiation – cosmological parameters – cosmology: observations
– space vehicles: instruments

1. INTRODUCTION

The March 2013 release of *Planck* data³ was another milestone in the fast-paced series of advances in observational cosmology. In December 2012 the *WMAP* team released an analysis of their full nine-year data set (Bennett et al. 2013; Hinshaw et al. 2013). The existence of two (almost completely) independent full-sky surveys of microwave emission using multiple overlapping frequency bands presents an enormous opportunity for cross-checking results.

There have been surprisingly few attempts to carefully evaluate the statistical consistency of the *WMAP* and *Planck* data. A thorough treatment requires accounting for both the similarities and differences between the two experiments: while both observe the same sky, they employ different frequency bands (with some partial overlap), they are sensitive to different multipole ranges, they apply different sky masks, and they have different in-

strument noise. While we take all of these effects into account in this paper, we are not the first to assess consistency. *Planck XV* (2014), *Planck XVI* (2014), and *Planck XXXI* (2014) report that the *WMAP* angular power spectrum is roughly 2.5% brighter than *Planck*'s for multipoles greater than $l \sim 100$ (1.25% in amplitude). Kovács et al. (2013) report that the *WMAP* angular power spectrum is about 2.6% higher than that of *Planck* “at a very high significance” on the same scales, while “at higher multipole moments there appears to be no significant bias.” Hazra & Shafieloo (2014b) tested consistency of the best-fit Λ CDM parameters: when they marginalize over the spectrum amplitude they find consistency, but when the amplitudes of each are fixed, they find tension at the 3σ level. They went on (Hazra & Shafieloo 2014a) to show that a concordance cosmological model fit to all data is consistent with the *Planck* data at only 2 to 3σ confidence level; the *Planck* spectrum had a dearth of power at both low- l and high- l relative to the concordance model. Spergel et al. (2013) found that the power spectrum from the *Planck* 217 \times 217 GHz detector set is responsible for some of the tension when compared with other cosmological measurements.

Figure 1 compares the best-fit Λ CDM models derived

Electronic address: dlarso11@jhu.edu

¹ Dept. of Physics & Astronomy, The Johns Hopkins University, 3400 N. Charles St., Baltimore, MD 21218-2686

² Dept. of Physics and Astronomy, University of British Columbia, Vancouver, BC Canada V6T 1Z1

³ <http://pla.esac.esa.int/pla/aio/planckProducts.html>

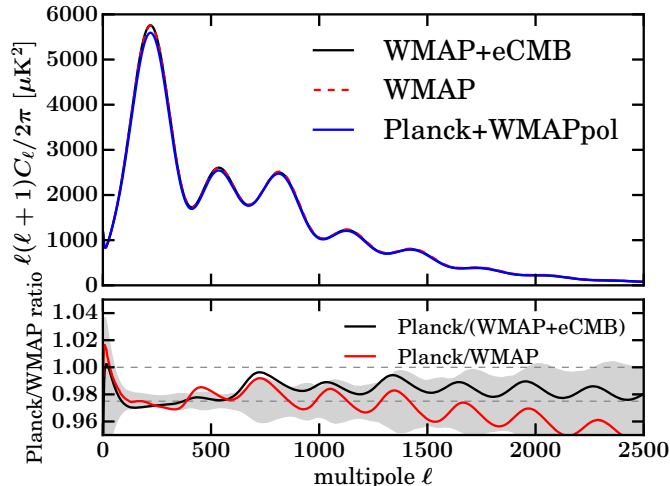


FIG. 1.— A comparison of the best-fit 6-parameter Λ CDM model spectra derived from recent CMB data sets. *Solid black* - fit to WMAP+eCMB data (SPT and ACT); *dashed red* - fit to WMAP (only) data; *solid blue* - fit to *Planck* plus WMAP low- l polarization data. The bottom panel shows the ratio of the *Planck*-based fits to the two WMAP-based fits. The dominant feature is a $\sim 2\%$ amplitude difference between the fits, though some structure is present in the ratio. A dashed horizontal line is shown at 0.975 to guide the eye. The WMAP-only forecast, in red, exhibits an additional $\sim 2\%$ difference at high l , well beyond the l range directly measured by WMAP ($l < 1200$). The gray band shows the 1σ uncertainty in the ratio accounting for the fact that WMAP and *Planck* observe the same sky, but apply different sky masks, cover different multipole ranges, and have independent noise. Given this uncertainty, we cannot distinguish a simple amplitude rescaling from a more complicated shift in parameters.

from *Planck* and WMAP data. The top panel shows the models and the bottom panel shows the model ratios. The error band plotted in Figure 1 (derived from simulations described in §4) indicates the spread in model fits that would be expected given differences in noise, multipole coverage, and sky masking between *Planck* and WMAP. The WMAP signal is clearly brighter than the *Planck* signal for multipoles $l \gtrsim 100$. Given the gray error band, for $l \lesssim 100$ the interpretation of the model fits is less clear. However, spectrum ratios derived purely from data suggest that the $\sim 2.5\%$ ratio persists to low l , see, for example, Figure A.1 of *Planck XVI* (2014). The range of best-fit models appears to encompass an overall rescaling of the data (and nothing more), but a richer set of parameter shifts is also allowed. A pure scaling difference would be surprising given that the initial *Planck* data release adopted the WMAP dipole signal as an absolute brightness standard.

In §4, we argue that the parameter sets making up each of these models are discrepant at the $\sim 6\sigma$ level. Curiously, the amplitude parameters A_s differ by less than 0.1σ , so the differences noted above are being absorbed by a combination of all the Λ CDM parameters. Even so, the most discrepant parameter, $\Omega_c h^2$, differs by only 1.7σ (when all other parameters are marginalized over). The announcement accompanying the first *Planck* data release emphasized that the universe is older and more massive than previously believed. This change is largely driven by the higher $\Omega_c h^2$ that *Planck* infers, though differences in the inferred Hubble constant (Bennett et al. 2014) may also contribute.

In this paper we examine differences between the *Planck* and WMAP data in detail. In §2.1 we study large angular scale differences in the delivered, single-frequency sky maps. In §2.2 we correct for residual foreground differences and again compare single-frequency map differences. In §2.3 we use a modified internal linear combination technique to examine multi-frequency cosmic microwave background (CMB) maps. The *Planck* multi-frequency SMICA map is compared with the WMAP internal linear combination (ILC) map as well.

In §3 we compare the angular power spectra from the two experiments and address calibration issues. We start in §3.1 with a comparison of the Λ CDM model fits (without regard to the inferred parameter values) as a means to enhance the sensitivity of the comparison. We argue that the shape of the power spectrum difference likely arises from errors in the transfer function of one or both experiments. Unfortunately we lack access to the requisite data needed to fully assess the *Planck* transfer function, but in §3.2 we re-examine the WMAP transfer function, specifically the WMAP main beam profiles and solid angle estimates in §3.2.1, and the WMAP far sidelobe response and its effect on dipole calibration in §3.2.2. These re-analyses are aimed at establishing the largest potential uncertainty in the WMAP measurements of the CMB power spectrum.

In §4 we compare the cosmological parameters inferred from the *Planck* and WMAP data. In addition to individual parameter comparisons, we study the consistency of the full 6-parameter Λ CDM model as a whole, taking into account the effects required to produce an accurate covariance matrix of parameter differences. The 6 Λ CDM parameters are the physical baryon density, $\Omega_b h^2$; the physical cold dark matter density, $\Omega_c h^2$; the amplitude of scalar fluctuations at $k_0 = 0.05 \text{ Mpc}^{-1}$, $A_{s,0.05}$; the scalar fluctuation power spectral index, n_s ; the Hubble constant, H_0 , in $\text{km s}^{-1} \text{ Mpc}^{-1}$; and the reionization optical depth, τ . Our conclusions are presented in §5.

2. SKY MAP COMPARISONS

We compare released versions of co-added single-frequency sky maps from *Planck* and WMAP, especially those derived from observations taken using similar pass bands. We also compare results taken from linear combinations of multi-frequency maps. All comparisons are performed with maps smoothed to a common resolution of 2° . Examination of maps on these scales is useful for visualizing large-scale effects such as far-sidelobe pickup.

2.1. Published single-frequency maps

The three *Planck* LFI bands and the 100 GHz HFI band have similar frequency coverage to the WMAP Ka, Q, V and W bands. We form *Planck*–WMAP difference maps from the pairs 30–Ka, 44–Q, 70–V, and 100–W in thermodynamic temperature units. If the absolute calibration of the *Planck* and WMAP maps agree to within a percent or so, the $l > 1$ CMB anisotropy signal should visually cancel in the difference maps. However, because the CMB dipole ($l = 1$) signal is relatively bright, even small calibration differences will be expected to produce residual dipoles in the difference maps. Further, since the *Planck* and WMAP band pairs have somewhat different




FIG. 2.— Differences between nearest-frequency *Planck* and *WMAP* sky maps, smoothed to 2° . The *Planck* bands are designated by their nominal frequencies (30, 44, 70 and 100 GHz), while *WMAP* bands are designated by Ka, Q, V and W, corresponding to nominal frequencies of 33, 41, 61 and 94 GHz, respectively. Differences are in thermodynamic units, on a $\pm 50 \mu\text{K}$ scale. Since the *Planck* and *WMAP* pass bands are not identical, Galactic foreground residuals are present; this is largest at the lower frequencies: the 30–Ka difference is saturated in the plane on this scale. The differences are always shown as *Planck* minus *WMAP* regardless of frequency ordering; this produces the expected negative Galactic residual in 44–Q and 70–V. Additionally, the *Planck* 100 GHz band encompasses emission from the 115 GHz CO $J = 1 - 0$ transition, which is seen in the Galactic plane in the 100–W difference. Although far-sidelobe (FSL) pickup of dipole and Galactic signals has been removed from the *WMAP* maps, it is still present in the currently available *Planck* maps at these frequencies. Of the three LFI bands, FSL pickup is strongest at 30 GHz. A residual dipole is evident in the 44–Q, 70–V and 100–W maps, indicative of calibration differences between the two experiments. A dipole residual similar to that seen in 44–Q and 70–V is also present in 30–Ka, but is obscured by the Galactic foreground residual, and by the presence of *Planck* 30 GHz far-sidelobe pickup of the CMB dipole. A deeper investigation of underlying differences between the two experiments at spatial scales $l > 1$ necessitates the removal of dipole and foreground residuals, and also ideally the *Planck* FSL contribution.

pass bands, we expect to see imperfect foreground signal cancellation in the difference maps. Figure 2 shows sky map differences formed from the nominal 15.5 month *Planck* DR1 release and the final nine-year *WMAP* release.

The [44,Q], [70,V] and [100,W] map differences in Figure 2 show clear dipole residuals. (The [30,Ka] difference map is dominated by Galactic signals and *Planck* far-sidelobe pickup, so it requires further processing to uncover the dipole residual.) These dipoles are well aligned with the CMB dipole direction, and show a systematic sign: all are positive in the south and negative in the north. This sign is consistent with the *WMAP* sky maps having a brighter temperature scale than *Planck*'s. We use the seven-year *WMAP* value of the CMB dipole (Jarosik et al. 2011) to estimate the fractional magnitude of the residual in each of these maps. We find the residual dipoles in the [44,Q], [70,V] and [100,W] differences are 0.3%, 0.3% and 0.5% of the CMB dipole, respectively. These values are within the quoted uncertainties for both the *Planck* and *WMAP* data (Planck III 2014; Planck VIII 2014; Bennett et al. 2013). Note that these differences are insufficient by at least a factor of 2 to explain the power spectrum differences at higher multipoles (see §3.1).

2.2. Foreground-reduced single frequency maps

Contributions from Galactic foregrounds must be removed (‘cleaned’) from the single-frequency difference maps before any residuals not associated with true sky signals may be seen. In this section, that foreground re-

moval is accomplished using spatial templates that trace known Galactic emission, fit as a linear combination to each difference map. The templates used for this purpose include some that are completely external to either the *Planck* or *WMAP* data, as well as pairwise differences between *WMAP* bands. The cleaned difference maps are largely independent of the exact templates used for cleaning; this is especially true of the cleaned 70–V and 100–W maps, where Galactic foregrounds are near their minimum.

The left column of Figure 3 shows the *Planck*–*WMAP* single-frequency difference maps after removal of foreground templates and a dipole component. Each LFI–*WMAP* difference is cleaned using a linear combination of 3 foreground templates: two in common and one specific to a given frequency pair. The first of the two common templates is the FDS dust emission model 8 (Finkbeiner et al. 1999) evaluated at 94 GHz which, for our purposes, serves as either a thermal or spinning dust template. The second common template is a surrogate for free-free emission: we employ a modified version of the Finkbeiner (2003) composite H α template which has been corrected for Galactic extinction (Bennett et al. 2013). A third template is formed from the difference between bracketing *WMAP* frequencies: for example, K–Q was used as a template for the 30–Ka pair. This choice permits a realistic near-frequency foreground template at the expense of including *WMAP* instrument noise. For 100–W, only the two common foreground templates were used, but high CO emission near the Galactic plane in the *Planck* 100 GHz map necessitated masking this




FIG. 3.— Template-cleaned *Planck*–*WMAP* difference maps; the color scale is $\pm 10 \mu\text{K}$. The maps have been smoothed to 2° resolution. Each row corresponds to a single-frequency difference, from top to bottom: 30–Ka, 44–Q, 70–V, and 100–W. *Left*: residuals after removal of foreground templates and a dipole component. The template fits were performed over the entire sky for the LFI–*WMAP* differences: thus as expected, foreground removal is imperfect near the Galactic plane. Regions with strong CO emission (grey) were excluded from the 100–W template fit. *Right*: residuals after removal of foreground templates and both a dipole and quadrupole component. Grey areas indicate regions masked from the fit, which was kept the same for all maps in this column. Inclusion of quadrupole terms in the template fits removes much of the residual signal from the LFI–*WMAP* differences. No far sidelobe correction has been applied to *Planck* data; far sidelobe pickup of the CMB dipole is removed to first order by inclusion of a dipole fitting component. Far sidelobe Galactic pickup remains; this is most evident in 30–Ka, where Galactic far sidelobe pickup produces a $\sim 85^\circ$ “ring” about the Galactic center.

region from the template fit.

The LFI–*WMAP* differences in the left column of Figure 3 exhibit a clear quadrupolar signature. This originates in part from the treatment of the kinematic quadrupole in the delivered *Planck* and *WMAP* map products. Both HFI and *WMAP* data retain the kinematic quadrupole in the map products (Planck VIII 2014; Jarosik et al. 2007) whereas LFI intended to fully remove it (Planck II 2014). However, the *Planck* Collaboration later noted that a code bug had unintentionally caused the retention of half of the kinematic quadrupole (Planck V 2014, Appendix A.1) in the 2013 DR1 LFI maps, leaving a residual quadrupole with a peak-to-trough amplitude of $\sim 2 \mu\text{K}$. We have additionally fit for and removed a quadrupole from the difference maps and the result is shown in the right column of Figure 3. This removal leaves a recognizable *Planck* FSL signature in the 30–Ka difference map, but the majority of structure in the 44–Q and 70–V difference maps appears to have been removed. The quadrupole fit to the 30–Ka map

has a peak-to-trough amplitude near $8 \mu\text{K}$ and is most likely influenced by residual FSL and foregrounds. The quadrupoles fit to the 44–Q and 70–V maps have similar peak-to-trough amplitudes of roughly $4 \mu\text{K}$ and are fairly well aligned with the residual kinematic quadrupole, but have amplitudes twice that expected based on the documented treatment of the kinematic quadrupole in LFI and *WMAP* data processing. It is possible, however, that the residual quadrupole we observe is associated with a systematic which just happens to align with the kinematic quadrupole. A remaining known systematic is the LFI FSL pickup which was not removed in the delivered sky maps; we next attempt to remove this contribution based on published LFI estimates.

Figure 4 illustrates results similar to Figure 3 but uses only the *Planck* Survey 1 and Survey 2 LFI maps, rather than the nominal 15.5 month mission maps. The reason for using only Survey 1 and 2 maps is two-fold: 1) we can remove digitally scanned versions of LFI far-sidelobe estimates for Survey 1 and 2 in order to estimate the

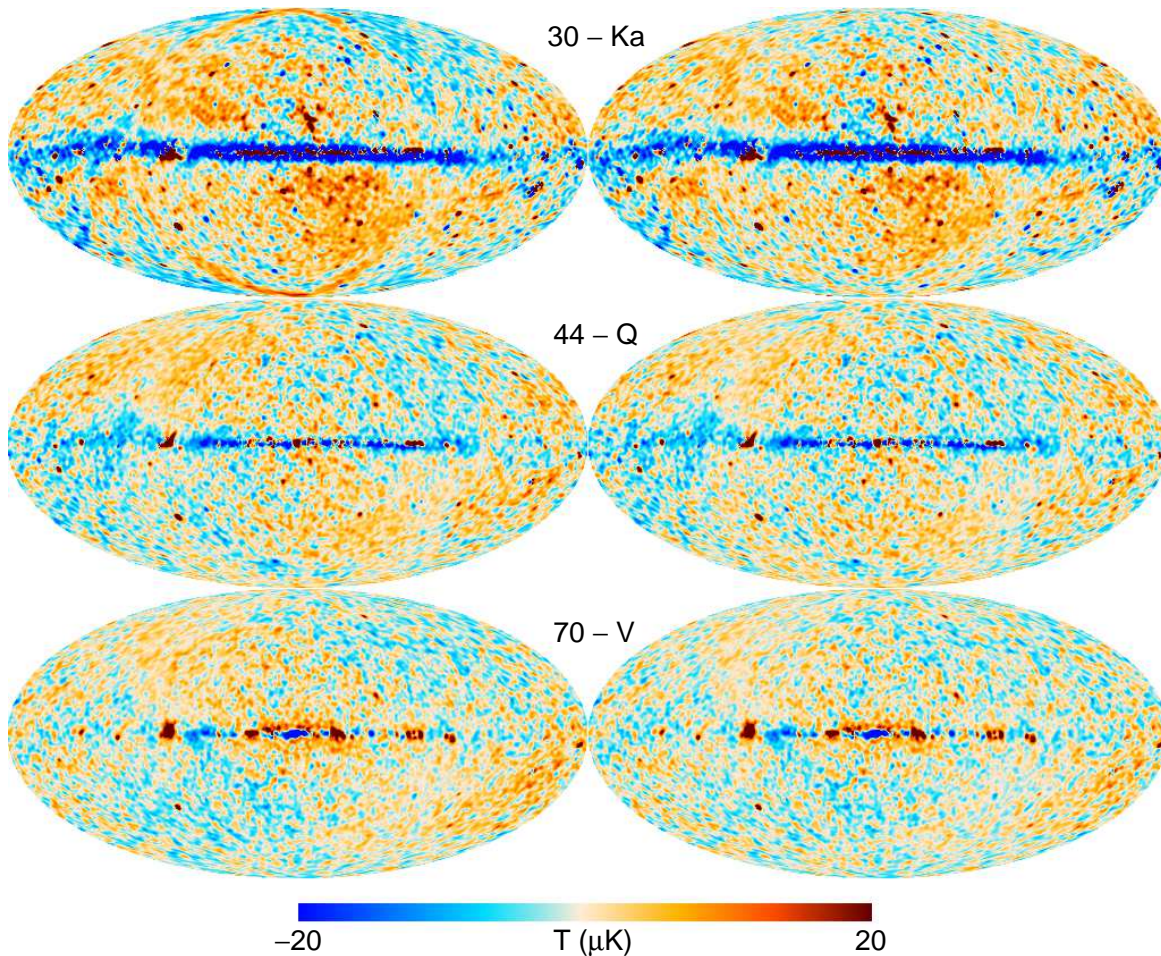


FIG. 4.— Combined *Planck* (Survey 1+Survey 2)–*WMAP* sky maps; the color scale is $\pm 20 \mu\text{K}$. The maps have been smoothed to 2° resolution. Each row corresponds to a single-frequency difference: from top to bottom these are 30–Ka, 44–Q, and 70–V. The left column shows differences after removal of Galactic emission, dipole and monopole components. The right column is identical to the left, with the exception that an estimate of FSL emission and a half-strength kinematic quadrupole component have been removed. The *Planck* FSL contribution was obtained using digitized versions of the estimates published in [Planck III \(2014\)](#). The quadrupolar structure seen in [Figure 3](#) is present in these difference pairs as well.

effect of FSL removal, and 2) estimates of LFI systematics are not currently provided for Survey 3 ([Planck III 2014](#)). The Galactic foreground templates and coefficients used to clean the (Survey 1+Survey 2)–*WMAP* data are identical to those used to clean the nominal 15.5 month *Planck*–*WMAP* difference maps. As seen in the right-hand column of [Figure 4](#), removal of the digitized FSL pickup does not completely remove high-latitude, large-scale structure in the LFI–*WMAP* differences. We conclude that unaccounted-for systematics remain, but at this stage cannot attribute them unambiguously to a particular instrument.

The 100–W difference map has a different residual signature than that seen in the LFI–*WMAP* maps. The residuals are shown both in [Figure 3](#) and in the top panel of [Figure 5](#). [Figure 3](#) demonstrates that visible structure in 100–W is little affected by removal of a quadrupole component. While searching for possible explanations of the structure in this map, we coincidentally examined the *Planck*-estimated maps of CO emission, in particular the Type 2 and 3 CO maps, which are generated solely from *Planck* data ([Planck XIII 2014](#)). The bottom panel of [Figure 5](#) shows high latitude, low-level residuals not

associated with CO emission present in the Type 2 CO map, which is produced from *Planck* 70, 100, 143 and 353 GHz data. (The Type 3 CO map has similar structure, but is not shown.) The similarity of the top and bottom images gives us reasonable confidence that most of the structure in the 100–W map arises from the 100 GHz map. We discuss this further in the next section.

2.3. Multi-frequency difference maps

2.3.1. Extended ILCs

In this section we consider another approach to foreground subtraction. As originally implemented, the Internal Linear Combination (ILC) method ([Bennett et al. 2003](#)) was designed to produce a foreground-cleaned CMB map from a linear combination of multi-frequency maps, each in thermodynamic temperature units. The single-frequency map weights were chosen to minimize the variance of the combined map (within a large multipixel region) subject to the constraint that the weights sum to 1.0, thus preserving the CMB signal.

In this section we extend the ILC algorithm to form what may be regarded as the difference between two




FIG. 5.— *Top*: *Planck*–*WMAP*, for the 100–*W* pair, after template subtraction of foreground emission and smoothing to 2° resolution; the color scale is $\pm 10 \mu\text{K}$. The grey region masks strong CO emission present in the *Planck* 100 GHz map. There is clear residual structure in excess of the instrument noise. *Bottom*: The *Planck* Type 2 CO $J = 1-0$ map, degraded to $N_{\text{side}} = 32$ (1.8° pixels). As in the top image, the grey region masks most of the known CO emission. The remaining large-scale features are unrelated to CO emission and may be considered residual systematics. A logarithmic scale is used to emphasize these features. The similarity of the structure in these two maps suggests that the structure in the difference map originates in the *Planck* 100 GHz data, rather than the *W*-band map, since the CO map is made entirely with *Planck* data.

ILC’s. The algorithm, described in greater detail in Appendix B, still uses a minimum variance criteria to minimize foreground signals, but it breaks the input maps into two groups: the weights for the first group are constrained to sum to $+1$, the second group to -1 . Ideally, the map formed in this way would completely cancel all CMB and foreground signals, yielding a null map. Systematic effects that are common between the two groups of maps would also tend to cancel, leaving behind signals that are distinct between the two sets.

We form extended ILC (eILC) maps from three groupings of data: LFI–*WMAP*, HFI–*WMAP*, and HFI–LFI. We use all the available LFI and *WMAP* frequencies, but only the 100, 143 and 217 GHz maps from HFI. We further limit the weight computation to use only pixels outside of the *WMAP*9 KQ75 mask (Bennett et al. 2013), which avoids strong foreground emission in the Galactic plane, in particular strong CO contamination in the HFI 100 and 217 GHz maps. In addition, we add back the half-strength kinematic quadrupole removed in the LFI processing (Planck V 2014) to ensure quadrupole consistency among the groupings.

The eILC maps from the groupings LFI–*WMAP*, HFI–*WMAP*, and HFI–LFI are shown in the first row of Figure 6. The weights used to form these maps are given in the first column of Table 1, under the columns labeled *a*. Evidently a dipole component must be removed from these maps in order to view finer details. To facilitate this, we add dipole components to the eILC computa-

tion, as described in Appendix B, and recompute the weights (the columns labeled *b* in Table 1). The results are shown in the middle row of Figure 6. The last row in Figure 6 similarly shows the result of including both dipole and quadrupole components in the eILC computation (the weights are given in the columns labeled *c* in Table 1).

The sky maps in the middle row of Figure 6 are similar in appearance to their counterparts constructed from template-cleaned, single-frequency difference maps, as discussed in §2.2. The LFI–*WMAP* eILC exhibits a quadrupole structure similar to that noted in the 30–Ka, 44–Q and 70–V differences. The HFI–*WMAP* eILC shows a similar ring-like structure as seen in 100–*W*. The *Planck*-only HFI–LFI eILC combination appears to contain both the broad ring-like structure and the inverse of the quadrupole signature noted in the *Planck*–*WMAP* combinations. Taken together, these features argue for the presence of residual systematic effects in the *Planck* maps, at the level of $\sim 5 \mu\text{K}$.

The presence of low-level residuals in the *Planck* DR1 maps is not unexpected. Both the LFI and HFI papers note incomplete removal of far-sidelobe pickup from the delivered maps (Planck III 2014; Planck XIV 2014). Other systematics in the LFI data are discussed in Planck III (2014), which states that the main effects are due to “stray-light pick-up from far sidelobes and imperfect photometric calibration”. We cannot definitively attribute residuals to a single cause. It is tempting to associate the broad ring-like structure in the 100–*W* difference map with *Planck* FSL pickup, since there is a known ring from primary spillover located roughly 85° from the Galactic center. However, it seems unlikely that the eILC residual is due to direct FSL pickup because the feature is not sharp. It appears more likely to us to be an indirect effect of FSL pickup on calibration, such as that illustrated in Figure 11 of Planck V (2014). Alternatively, these features may arise from time-dependent calibration residuals due to ADC nonlinearity or thermal effects (see, e.g., Figures 15 and 28 of Planck III (2014)).

2.3.2. *Planck* SMICA and *WMAP* ILC

We next compare CMB maps produced by the *Planck* and *WMAP* teams, respectively. The *WMAP* product is referred to simply as the ILC map (Bennett et al. 2013). Planck XII (2014) discuss four different methods for producing foreground-cleaned CMB maps. Of these, we adopt the SMICA map as representative of this class of *Planck* DR1 products, based on the *Planck* Collaboration’s recommendation (Planck XII 2014). We obtain the SMICA map and the *WMAP* nine-year ILC map at the same HEALPix resolution, smooth both to a common 2° FWHM resolution, and difference them. We further remove a dipole outside of a mask, and degrade to HEALPix $N_{\text{side}} = 16$ (3.6° pixels) to reduce the noise. This difference map is shown in the top image of Figure 7.

Large-scale structure in the SMICA–ILC difference map has been noted previously, e.g. Kovács et al. (2013), but we are unaware of any explanation of it in the literature. Based on our findings in previous sections, we conjecture that the dominant large-scale features arise from residual systematics in the *Planck* maps. We can

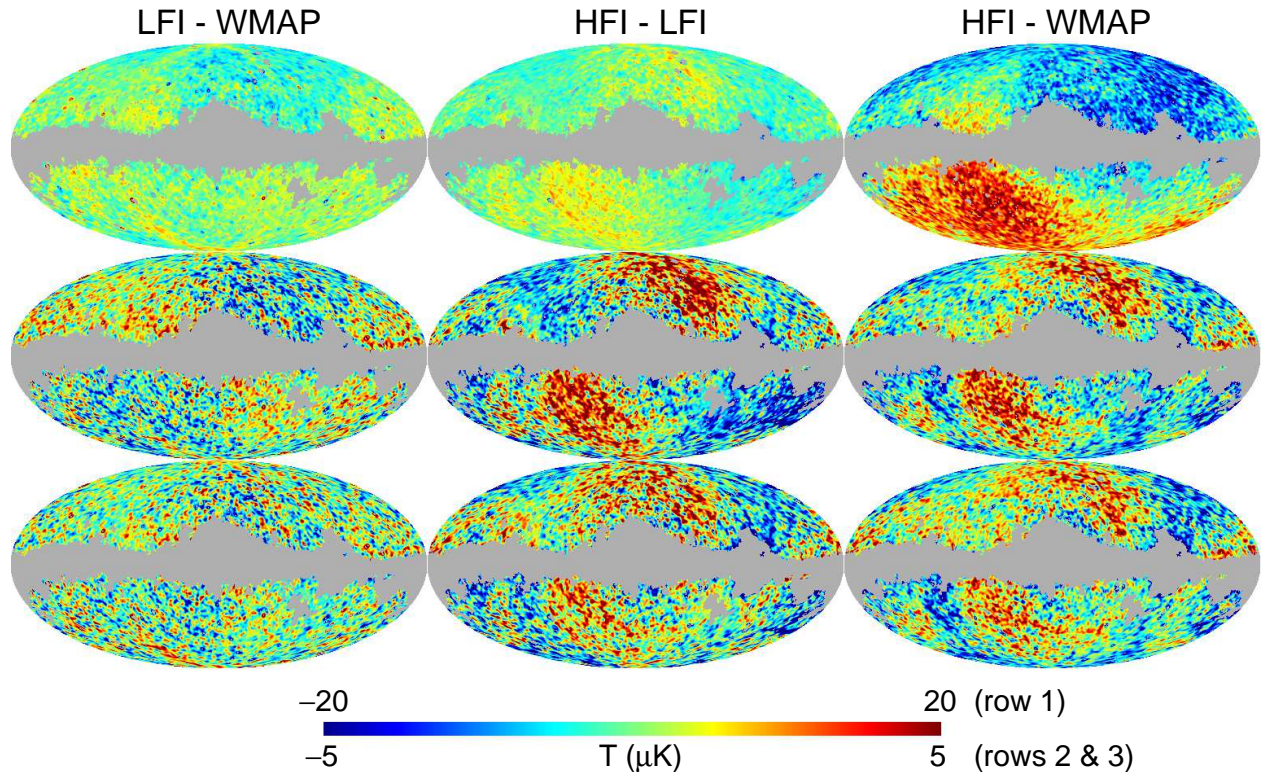


FIG. 6.— Extended ILC maps computed using maps smoothed to 2° resolution and applying a Q75 mask (grey regions). In the absence of systematics, these maps should resemble noise, but this is clearly not the case. The left column is LFI–WMAP, the middle is HFI–LFI, and the right is HFI–WMAP. Only 100, 143, and 217 GHz maps are used for the HFI calculations. The top row has no dipole component removed and uses a color scale of $\pm 20 \mu\text{K}$. The middle row has a dipole component removed and uses a color scale of $\pm 5 \mu\text{K}$. The bottom row has a dipole and quadrupole removed and uses a color scale of $\pm 5 \mu\text{K}$. The weights are given in Table 1. The *Planck*–only HFI–LFI combination indicates that systematic structure seen in the *Planck*–WMAP combinations arises primarily from the *Planck* maps.

test this hypothesis by forming maps of estimated systematics at each *Planck* frequency and combining them using the published SMICA weights to form an “artifact SMICA”. If the hypothesis is plausible, then the “artifact SMICA” will be similar to the SMICA–ILC map itself.

To estimate the residual systematics at each *Planck* frequency, we start with a 100–70 GHz difference map, and remove residual foregrounds using a linear combination of the FDS dust model and the Finkbeiner H α template, as in §2.2. The resultant foreground-cleaned map looks very much like the HFI–LFI eILC in Figure 6. A quadrupole component is additionally removed to form the 100 GHz template. This template resembles the cleaned 100–W difference map, but is free of any WMAP data. Next, candidate templates for 143, 217 and 353 GHz are constructed by differencing each of these band maps against the 100 GHz map, cleaning foregrounds with the above dust and free-free templates, and then adding back the 100 GHz artifact template. For 545 and 857 GHz, we adopt the far-sidelobe maps provided in the *Planck* DR1 release, but with the zodiacal contribution removed. For 30, 44 and 70 GHz, we adopt the small excess quadrupole in conjunction with the digitized LFI FSL pickup maps from §2.2.

Once we have candidate templates for each *Planck* frequency, we combine them using SMICA weights applicable to $l < 50$. We digitally scan these values from Figure D.1 of *Planck* XII (2014). Weights are specified in units such that a SMICA map in thermodynamic tem-

perature units is produced from a linear combination of maps in Rayleigh-Jeans units. The weights we use are listed in Table 2. The frequencies with the most weight are 100, 143, 217 and 353 GHz. Our estimated SMICA artifact map is shown in the bottom panel of Figure 7, and may be directly compared to the SMICA–ILC residual map above it. We conclude that this model for the SMICA–ILC residual is plausible; but other plausible models may exist.

3. POWER SPECTRUM COMPARISON

The WMAP and *Planck* temperature measurements can be compared in a variety of ways over a range of angular scales. The results appear to present a complex and sometimes conflicting story: at large to medium angular scales, the *Planck* data calibration produces a fainter sky than WMAP, while at the smallest angular scales within the 30–70 GHz bands, *Planck* returns a slightly brighter surface temperature for Jupiter than does WMAP, (Planck V 2014), though they are within the 0.5% to 1% uncertainties.

On the largest scale, the dipole residual in the WMAP – *Planck* difference map is $< 0.5\%$ of the CMB dipole (§2.1). Other systematic low- l structure, on the scale of a few μK , is present in both the LFI and HFI *Planck* maps (§2), but these effects are confined to roughly $l \leq 10$, and they contribute only a few μK^2 to the angular power spectrum. This is not a major source of discrepancy with WMAP. At intermediate angular scales, the discrepancy

TABLE 1
EXTENDED ILC COEFFICIENTS

Band/ Component	HFI–LFI			LFI–WMAP			HFI–WMAP		
	a	b	c	a	b	c	a	b	c
K	–	–	–	–0.237	–0.002	–0.010	–0.239	0.039	0.034
Ka	–	–	–	–0.133	–0.176	–0.187	1.485	0.060	0.067
Q	–	–	–	–0.383	–0.315	–0.299	–1.313	–0.172	–0.184
V	–	–	–	–0.145	–0.332	–0.315	–0.808	–0.482	–0.488
W	–	–	–	–0.102	–0.174	–0.189	–0.124	–0.445	–0.429
30	0.202	0.122	0.107	0.669	0.149	0.175	–	–	–
44	–0.552	–0.276	–0.285	–0.101	0.306	0.304	–	–	–
70	–0.650	–0.846	–0.823	0.432	0.545	0.521	–	–	–
100	0.448	0.452	0.389	–	–	–	0.232	0.402	0.360
143	0.812	0.784	0.854	–	–	–	1.055	0.821	0.868
217	–0.261	–0.237	–0.243	–	–	–	–0.287	–0.223	–0.228
d_x	–	–0.225	–0.229	–	–0.675	–0.657	–	–1.130	–1.124
d_y	–	–2.138	–2.176	–	–6.401	–6.242	–	–10.735	–10.678
d_z	–	2.409	2.452	–	7.223	7.035	–	12.098	12.034
$\tilde{a}_{2,-2}$	–	–	0.077	–	–	0.200	–	–	1.279
$\tilde{a}_{2,-1}$	–	–	4.034	–	–	–1.764	–	–	2.102
$\tilde{a}_{2,0}$	–	–	–1.838	–	–	1.890	–	–	–0.691
$\tilde{a}_{2,1}$	–	–	–1.195	–	–	–0.139	–	–	–0.448
$\tilde{a}_{2,2}$	–	–	1.136	–	–	–0.851	–	–	0.682

NOTE. — The weights in the first group (HFI or LFI) sum to +1, and in the second group (LFI or WMAP) to –1. The weights for the band maps are dimensionless, the weights for the multipole components have units of μK . Harmonic conventions are defined in Appendix A.

^a No multipole components were included in the weight computation.

^b Three dipole components were additionally included in the weight computation.

^c Five quadrupole components were additionally included in the weight computation.




FIG. 7.— *Top*: the SMICA–ILC map at HEALPix $N_{\text{side}} = 16$ (3.6° pixels); the color scale is $\pm 20 \mu\text{K}$. *Bottom*: the “SMICA artifact map” constructed from *Planck* data as described in the text. This artifact map produces similar large-scale structure to that seen in the SMICA–ILC difference map.

is more significant, as noted in [Planck XXXI \(2014\)](#); see also [Figure 1](#) above.

3.1. Power spectrum model fits

The largest difference between *WMAP* and *Planck* lies in the temperature power spectrum at $l \gtrsim 100$. To quantify the significance of this difference in a way that captures the full error estimates provided by each team, we evaluate the quantity $D_{220} \equiv l(l+1)C_l^{\text{T T}}/(2\pi)|_{l=220}$, the

TABLE 2
LOW MULTIPOLE SMICA
WEIGHTS^a

Band [GHz]	Weight [$\mu\text{K}^{\text{CMB}}/\mu\text{K}^{\text{RJ}}$]
30	+0.009708
44	–0.1029
70	–0.1145
100	–0.3281
143	+1.95
217	+1.303
353	–2.242
545	+0.8252
857	–0.13

^a These weights have been digitally scanned from [Figure D.1 of Planck XII \(2014\)](#). They are applicable only to $l < 50$ and subject to digitization error.

model spectrum amplitude near the first acoustic peak, using the Markov chains released by each team. We choose to compare the model amplitude (assuming flat ΛCDM) because it effectively smooths the data over a range of l , substantially reducing the uncertainties in the comparison. For the *WMAP* chains, the value of D_{220} was reported in the chains, so we adopt those values. For the *Planck* chains we evaluate D_{220} from the cosmological parameters using CAMB, including the Gaussian lensing approximation but excluding all foreground contributions (SZ effect, galactic foregrounds, etc.) since we seek to compare the CMB component of the power spectrum returned by each experiment.

From the first 10,000 points in each *Planck* chain we collect every 10th point to obtain 1000 relatively independent samples. Given the values of $\Omega_b h^2$, $\Omega_c h^2$, H_0 , $A_{s,0.05}$, n_s , τ , $Y_{\text{He,used}}$, and $\Omega_\nu h^2$, we recalculate

TABLE 3
COMPARISON OF D_{220} ^a VALUES

Data	D_{220} [μK^2]	Source
WMAP	5746 ± 35	WMAP9
WMAP	5745 ± 36	WMAP9, re-run by <i>Planck</i>
<i>Planck</i>	5588 ± 35	<i>Planck</i> ^b + WMAP9 low- l pol.
<i>Planck</i>	5587 ± 36	<i>Planck</i> alone ^b
<i>Planck</i>	5598 ± 35	<i>Planck</i> + τ prior ^c

^a $D_{220} \equiv l(l+1)C_l/(2\pi)|_{l=220}$.

^b $2 \leq l \leq 2500$

^c $50 \leq l \leq 2500$

TABLE 4
WMAP BEAM CHARACTERIZATION
REGIMES^a

Band	Cutoff (dBi / dB)	Transition Radius (degrees)
K	2.0 / -45.0	7.0
Ka	3.0 / -46.4	5.5
Q	5.0 / -46.3	5.0
V	6.0 / -48.8	4.0
W	9.0 / -48.8	3.5

^a From [Bennett et al. \(2013\)](#).

the power spectrum assuming one massive neutrino and 2.046 massless neutrinos. We use the March 2013 version of CAMB⁴, matching the version used to generate these chains. The mean and standard deviation are computed using the weights reported in the *Planck* chains.

The results for D_{220} are given in Table 3: we find the *Planck* CMB spectrum to be significantly lower than WMAP’s. While there is some variation between the three *Planck* chains explored, the range is small. Based on only 1000 samples, we expect the sampling error to be about 3% of the quoted uncertainty, so $1 \mu\text{K}^2$ differences in the mean are not significant. Furthermore, CAMB does not claim an accuracy of more than $\sim 0.1\%$, so even $6 \mu\text{K}^2$ variations should not be considered highly significant. As a cross check of our calculations, we extract D_{220} from the best-fit spectra released with the chains and find values consistent with those we derived: $5747.6 \mu\text{K}^2$ for the WMAP chain run by *Planck*, $5586.9 \mu\text{K}^2$ for the *Planck*+(WMAP low- l pol.) chain, $5588.4 \mu\text{K}^2$ for *Planck* alone, and $5597.5 \mu\text{K}^2$ for the *Planck*+ τ prior chain.

If the WMAP and *Planck* errors were uncorrelated, the WMAP D_{220} would be 3.2σ higher than *Planck*’s. Since these chain-based uncertainties include calibration error (which will be correlated because *Planck* used the WMAP solar dipole calibration) and cosmic variance error (which will be correlated between WMAP and *Planck* due to overlapping sky coverage), we expect some correlation in the D_{220} errors. Specifically, the uncertainty in the *difference* between the WMAP and *Planck* will be lower than if the errors were uncorrelated; this will increase the significance of the discrepancy. Using the simulations presented in §4 we estimate that the D_{220} uncertainties have a Pearson correlation coefficient of $r = 0.69$, not including l -to- l correlations in the window function

⁴ <http://camb.info/>

uncertainties. If we include this value of r , we estimate the difference, ΔD_{220} , is 5.7σ rather than 3.2σ , but this probably over-estimates the significance slightly due to the window function treatment. In the end, we estimate that the first peak amplitude is discrepant at somewhere between 3 and 5σ significance.

Plausible explanations for the *Planck* – WMAP spectrum difference include an overall brightness scale difference and/or beam characterization errors. However the former would be puzzling because the *Planck* brightness calibration was based on the WMAP dipole and, as discussed in §2, the measured dipoles are relatively consistent. [Hazra & Shafieloo \(2014b\)](#) find that a rescaled *Planck* concordance model is consistent with the WMAP data for $l < 1200$, but it does not follow that the measured discrepancy is actually due to a mis-calibration of either experiment’s brightness scale. We note that the recent *Planck* consistency paper ([Planck XXXI 2014](#)) claims some l dependence in the *Planck*/WMAP ratio and suggests that the cause is not simply a calibration error. Further, that paper reports a partial reduction in the *Planck*/WMAP spectrum difference following a revised analysis of the *Planck* beam response. Since the Collaboration has not yet released a new likelihood code or Markov chains, we cannot determine the effect it has on D_{220} . Further, given the restricted scope of the first *Planck* data release, we cannot independently assess the *Planck* Collaboration’s estimates of systematic errors. Instead, we revisit potential sources of systematic error in the WMAP beam analysis, especially the possibility that the WMAP solid angle has been systematically over-estimated, as would be required to explain the sign of the spectrum discrepancy.

3.2. WMAP beam uncertainties revisited

For the purposes of WMAP data processing, the treatment of beam effects is divided into three regimes ([Hill et al. 2009](#)), each of which are characterized and processed differently:

Main beam: In this regime, the beam response is determined entirely by in-flight measurements of Jupiter. The extent of the main beam is defined by an amplitude cutoff, measured in dBi. See Table 4 for the band-specific cutoff values.

Near sidelobe: In this regime, the response is determined by a physical optics model fit to in-flight measurements of Jupiter. The fit is performed simultaneously to all data channels. The outer limits of this regime are defined by a fixed “transition radius” for each band (see Table 4), while the inner limit is defined by the amplitude cutoff noted above. In this regime the beam is characterized by a mixture of Jupiter data and the physical model, governed by the amplitude of the model.

Far sidelobe: In the first two regimes, the A-side and B-side beams were determined separately and treated as independent antennas. In the far sidelobe regime (the full sphere outside the transition radius of each beam) the response is treated as a proper differential antenna which carries a sign. The data informing the response comes from four sources, listed in order of importance:




FIG. 8.— WMAP far sidelobe beam maps for V1-W4 in spacecraft coordinates, where the top of the map is sunward and vice-versa. The main beam and near sidelobe response, within the “transition radius”, has been masked in grey and lies $\sim 20^\circ$ below the mid-plane of the map. Positive response (from the A side) is shaded brown, and negative response (from the B side) is shaded blue. The color scale is in linear gain units, $G = 10^{(\text{dBi}/10)}$ and it saturates at -10 dBi, far below the forward gain of $55 - 58$ dBi, in V and W band. Creme indicates regions where $|G| < -30$ dBi. The prominent rings on either side arise from radiation that spills past the edge of the primary reflector; the sign change in that response depends on whether or not such radiation also reflects off the large instrument radiator.

1) ground-based measurements of the beam pedestal (within 10°) using the reflector evaluation unit (REU) in the Goddard Electromagnetic Anechoic Chamber (GEMAC) at the Goddard Space Flight Center, 2) in-flight measurements of the response using the moon as a source, 3) measurements taken on the roof of Jadwin Hall at Princeton University using the REU, and 4) a physical optics model computed with the code DADRA. The compiled far sidelobe maps are shown in Figure 8, based on Fig. 2 of Barnes et al. (2003). We estimate the uncertainty of the higher gain regions to be $\sim 30\%$, based on in-flight measurements of the moon using in-flight calibration data, and matching this to the GEMAC data.

In the data processing, the regime inside the transition radius (main beam and near sidelobe) is treated as a point spread function (PSF) with a radial profile which treats the beam as circularly symmetric. This PSF convolves the angular power spectrum with a beam window function given by the Legendre transform of the radial profile. Appendix B of Hinshaw et al. (2007) analyzed the effects of beam asymmetry on the convolution and concluded that for WMAP’s scan strategy and sensitiv-

ity levels, the effects of beam asymmetry on power spectrum convolution were $< 1\%$, in power units, for V and W bands. This topic was revisited in Bennett et al. (2013) with the development of an explicit map-making method that deconvolved the effects of beam asymmetry, see especially Figure 9 and related text. This work upheld the previous conclusions quoted above.

The WMAP processing pipeline treats the far sidelobe regime differently in that it estimates the differential far sidelobe contribution for each time-ordered data point using a model sky map based on the data. This correction is incorporated into the intensity calibration algorithm so that it self-consistently accounts for the $l = 1$ (dipole) response of the far sidelobe when estimating the intensity scale. We describe this processing step more fully later in this section.

3.2.1. Beam profiles and solid angle estimates

As noted above, the far sidelobe response is treated as a differential signal correction in the WMAP data processing. However, for study purposes it is useful to estimate the fraction of the beam solid angle that is attributable to the far sidelobes. This attribution is not unambiguous




FIG. 9.— The azimuthally-averaged radial profile of the WMAP beam response for the A-side V and W band differencing assemblies, normalized to a peak value of 1. The blue portion indicates the main beam and near sidelobe response within the transition radius; the black and red portions are two representations of the (differential) far sidelobe response binned by distance from the A-side boresight. The red trace includes all structure shown in Figure 8, irrespective of sign, while the black trace excludes spill past the B-side reflector, as described in the text and as shown in Figure 11. In the nominal WMAP data processing, the sidelobe response is corrected for in the differential data, prior to map-making.

because certain far sidelobe features cross the mid-plane dividing the A and B sides of WMAP and thus complicate a naive accounting of the solid angle. Nonetheless, we can obtain useful insights and so we proceed.

Figure 9 shows the azimuthally averaged radial profile of the 6 V and W-band A-side DA's, compiled from the sources listed above. The blue portion indicates the main beam and near sidelobe: the regime that is accounted for using a window function. The red portion shows the result of averaging the far sidelobe response (Figure 8) in radial bins centered on the A-side boresight. The black curve is a modified form of the red curve where unambiguous negative B-side contributions have been masked from the far sidelobe response to assess the magnitude of their contribution. Figure 10 shows the fractional cumulative solid angle derived from these profiles as a function of distance from the A-side boresight. These integrals are evaluated all the way out to 180° for study purposes, but they are only used inside the transition radius (3.5° and 4.0° for W and V bands, respectively). We note the following items in Figure 10:

1. The fraction of solid angle contributed by the far sidelobe response (outside the transition radius) is between 0.5% and 1%, but this result does not fully account for the differential nature of the far sidelobe response. Beyond about $\sim 20^\circ$ from each boresight, we begin to mix A-side (positive) and B-side (negative) response due to spill past the primary. We explore alternative accountings below. In the nominal WMAP data processing, the sidelobe response is corrected for in the differential data, prior




FIG. 10.— Cumulative solid angle of the A-side V and W band beam profiles, normalized to 1 at 180° . The dashed line indicates the transition radius, inside of which the WMAP beams are modeled using a combination of in-flight Jupiter data and a physical optics model. With the exception of W2 and W3, the response within the transition radius accounts for $\sim 99.5\%$ of the beam solid angle. The grey band shows the effects of altering the amplitude of the model contribution by $\pm 100\%$: the cumulative solid angle changes by $\sim 0.1\%$. Results for the B side are similar. The normalization at 180° is established using the masked sidelobe response (black trace in Figure 9); we explore the unmasked treatment in Figure 11. In the nominal WMAP data processing, the far sidelobe response (outside the transition radius) is corrected for in the differential data, prior to map-making.

to map-making. We revisit that procedure below as well.

2. The fraction of solid angle in the response between the transition radius and 20° is at most 0.1% for any DA. This limits the plausible range over which errors in this response could be affecting the far sidelobe correction.
3. The solid angle in the region between $\sim 1^\circ$ and the transition radius arises from a combination of Jupiter data and a physical optics model. The uncertainty in the model contribution is indicated by the grey band which shows the affect of varying the amplitude of the model by $\pm 100\%$: the solid angle changes by $\sim 0.1\%$. In particular, setting the model contribution to zero only reduces the beam solid angle within the transition radius by 0.1%.
4. The remainder of the beam solid angle within the transition radius is characterized by in-flight measurements of Jupiter. These data and their uncertainties are very well understood and are reflected in the published WMAP beam covariance matrices.

Recall that WMAP's measured power spectrum is higher than that reported by Planck. One way to reduce

TABLE 5
WMAP FAR SIDELOBE RECALIBRATION
FACTORS

Band/ sub-band	WMAP5+ ^a (%)	WMAP3 ^b (%)
V	+0.05	+0.10
W14	+0.23	+0.44
W23	-0.05	-0.12

^a WMAP five-, seven- and nine-year data releases (Hinshaw et al. 2009).

^b WMAP three-year release (Jarosik et al. 2007).

the measured spectrum is to reduce the estimated beam solid angle, which reduces the factor by which one deconvolves the raw spectrum. Since the physical optics model is contributing at most 0.1% to the overall solid angle, this limits the degree to which the final power spectrum amplitude can be reduced ($\sim 0.2\%$). We must look to the far sidelobes, especially item 1 above, for any remaining beam-related sources of error, beyond what is already published in the WMAP beam covariance matrices.

3.2.2. WMAP far sidelobe response and calibration

As discussed earlier in §3.2, the far sidelobe response maps are compiled from a variety of sources (Barnes et al. 2003; Hill et al. 2009), see Figure 8. They are dominated by power at $l < 40$ with the main features being a pedestal from the transition radius out to $\sim 20^\circ$, and direct spill past the primary mirror. We estimate an overall uncertainty in the highest-gain features in these maps to be about 30%.

An estimate of the far sidelobe response is removed in the time-ordered data as follows (Hinshaw et al. 2009): we evaluate a hyper-cube of responses to a dipole and a model sky map (based on earlier data) for all possible spacecraft Euler angles, with $\sim 1^\circ$ spacing. For each time-ordered data point, we evaluate and apply the correction, in raw counts, prior to evaluating a dipole-based gain solution, so that our response to the dipole is, in principle, only convolved by the beam response within the transition radius. We have evaluated the overall gain solutions with and without applying this correction and have tabulated the differences between them in Table 5. For five-year WMAP data release and beyond, the gain corrections range between -0.05% and $+0.23\%$ for V and W band. For reference, Table 5 also gives the factors for the three-year release, which used a slightly different transition radius and treatment of the sidelobe correction. The results are very similar, but generally a little larger.

In the vicinity of first acoustic peak, where the power spectrum discrepancy is most significant, the WMAP V and W band data contribute roughly equally to the combined CMB spectrum. In that regime, we estimate the net recalibration applied to the combined data is $\sim 0.08\%$. The far sidelobe response uncertainty of 30% implies a comparable fractional uncertainty in the net recalibration factor. Even if we were to inflate the uncertainty to 100%, corrections to the recalibration would be insufficient to close the observed gap between the WMAP and Planck spectra.




FIG. 11.— A summary of far sidelobe tests applied to the WMAP V1 differencing assembly. (a) The full-sky differential beam map, as in Figure 8 but with the main beam and near sidelobe regions unmasked. The white cross indicates the centroid of the A-side beam, and lines of constant angular distance from the A-side centroid are shown as white circles, in 20° increments. Regions outlined in black indicate relatively bright sidelobes of the B-side beam that are optionally masked when analyzing the A-side solid angle. (b) The azimuthally-averaged beam response, as in Figure 9. (c) Zoom of (b) in linear units. This illustrates the effect of including (red) or excluding (black) the B-side primary spill (outlined region) in the A-side solid angle computation. The response within 12° of the B-side beam centroid is excluded in both cases. (d) The cumulative solid angle at large angles, normalized to 1 at $\theta = 180^\circ$. The nominal treatment excludes the B-side primary spill (black trace); the red trace indicates that including the negative B-side spill would reduce the A-side solid angle by 0.5%. 99.5% of the solid angle is contributed by the main beam and near sidelobe, as shown in blue. Very large fractional and conspiring errors would be required to produce a significant change in our estimate of the solid angle from the far sidelobe. See the text for further discussion.

4. COSMOLOGICAL PARAMETER COMPARISON

In this section we consider the standard 6-parameter Λ CDM model and compare the best-fit parameters obtained from the WMAP nine-year data (alone) and the initial Planck data plus the WMAP polarization data (“WMAPpol”). The best-fit parameters for each data set are given in Table 6, while the best-fit power spectra derived from each model are shown and compared in Figure 1.

TABLE 6
MEAN AND MAXIMUM LIKELIHOOD COSMOLOGICAL PARAMETERS

Parameter	WMAP mean	Planck mean ^a	WMAP m.l. ^b	Planck m.l. ^c	WMAP–Planck m.l. ^d
$\Omega_b h^2$	0.02264 ± 0.00050	0.02205 ± 0.00028	0.02262	0.02194	0.00068 ± 0.00050 (1.4 σ)
$\Omega_c h^2$	0.1138 ± 0.0045	0.1199 ± 0.0027	0.1139	0.1220	-0.0081 ± 0.0049 (1.7 σ)
H_0 [km s ⁻¹ Mpc ⁻¹]	70.0 ± 2.2	67.3 ± 1.2	70.0	66.8	3.2 ± 2.3 (1.4 σ)
n_s	0.972 ± 0.013	0.9603 ± 0.0073	0.974	0.953	0.020 ± 0.014 (1.4 σ)
$10^9 A_{s,0.05}$ ^e	2.203 ± 0.067	2.196 ± 0.055	2.203	2.200	0.004 ± 0.042 (0.1 σ)
τ	0.089 ± 0.014	$0.089^{+0.012}_{-0.014}$	0.0875	0.0873	0.0002 ± 0.0031 (0.1 σ)

NOTE. — The results in the column “Planck mean” assume a non-zero neutrino mass, as per Planck XVI (2014); the remaining columns set the neutrino mass sum to 0 and the helium fraction to $Y_{\text{He}} = 0.24$.

^a “Planck+WP” results, from Table 2 of Planck XVI (2014). This uses Planck data for $2 \leq l \leq 2500$, as well as the low- l WMAP polarization likelihood (WP).

^b The WMAP maximum likelihood parameters were determined by fitting a quadratic function to the peak of the log likelihood, based on the released WMAP9 Λ CDM chain. The SZ amplitude was constrained to be positive.

^c The Planck maximum likelihood parameters were computed by maximizing the Planck TT likelihood code ($50 \leq l \leq 2500$) over 20 parameters: the 6 Λ CDM parameters and 14 foreground and calibration parameters. We assumed an optical depth prior of $\tau = 0.085 \pm 0.01$.

^d Single-parameter differences with marginalized uncertainties. The maximum likelihood parameters are used and the quoted errors account for correlations between the experiments using the rescaled simulation covariance matrix discussed in the text.

^e The WMAP team reports the spectrum normalization at $k_0 = 0.002$ Mpc⁻¹ while the Planck Collaboration adopts $k_0 = 0.05$ Mpc⁻¹. We report $A_{s,0.05}$, calculated from released chains, and we note that $A_{s,0.002} = A_{s,0.05} \times 25^{1-n_s}$ for the simple Λ CDM model without a running spectral index.

At the level of single parameters, the inferred values are generally consistent within their (marginalized) errors. For example, WMAP estimates the Hubble constant to be $H_0 = 70.0 \pm 2.2$ km s⁻¹ Mpc⁻¹ while Planck estimates $H_0 = 67.3 \pm 1.2$ km s⁻¹ Mpc⁻¹ (Bennett et al. 2014). If these errors were uncorrelated and Gaussian, we would expect the measurements to differ by 2.5 km s⁻¹ Mpc⁻¹ (1 σ), making the observed difference 1.1 σ . However, this comparison ignores the fact that a substantial fraction of the uncertainty in each parameter is due to cosmic variance, which is common between the two experiments; thus it over-estimates the consistency. (Looking ahead, the final column of Table 6 quotes single parameter differences with uncertainties that account for these correlations, using the methodology described below.) More importantly, a complete test of parameter consistency requires that we simultaneously assess all 6 Λ CDM parameter differences accounting for their parameter covariance.

A simple test of 6-dimensional parameter consistency is to examine how much “overlap” exists between points in a WMAP parameter chain and a corresponding Planck chain. It turns out that a single hyper-plane in parameter space segregates the chain points from each experiment at the 0.1% level. Specifically, the parameter combination

$$\begin{aligned}
 P \equiv & \Omega_b h^2 - 0.874789 \Omega_c h^2 - 0.000736964 H_0 \\
 & - 0.051314 n_s + 0.0710558 \times 10^9 A_{s,0.05} \\
 & - 0.316342 \tau + 0.052037 \quad (1)
 \end{aligned}$$

where H_0 is in km s⁻¹ Mpc⁻¹ and the other parameters are dimensionless, splits the experiments as follows: $< 0.1\%$ of the Planck-released WMAP chain weight has $P < 0$, while $< 0.1\%$ of the Planck+WMAPpol chain weight has $P > 0$.

4.1. Simulating experiment covariance

We next proceed to quantify the covariance expected between the parameters inferred from each experiment.

To do so, we generate simulations that approximate the full parameter recovery process for each experiment and we use them to generate the full covariance matrix of parameter differences expected given the effects of independent instrument noise, different sky masks, and different multipole ranges probed by the two experiments.

To make the problem computationally tractable, we do not perform a full Markov chain analysis for each simulation realization but rather we estimate the most likely parameter value for each realization (and data set), using the `scipy` implementation of the Powell method. We then compare the measured parameter differences to the statistical distribution of simulated parameter differences to assess the consistency of the two measured parameter sets. To further speed up the calculation we ignore 14 calibration and foreground nuisance parameters used in the full Planck likelihood. Since our simulated maps do not include any modeled foreground residuals or calibration artifacts, ignoring these parameters will not introduce any bias.

The simulation program proceeds as follows: for each realization in our simulation, we generate a CMB map from a fixed Λ CDM power spectrum whose parameters are given in Table 6. We compute the a_{lm} coefficients of this (masked) map using PolSpice⁵ (Szapudi et al. 2001; Chon et al. 2004); for the WMAP simulation we apply the KQ85y9 temperature analysis mask (Bennett et al. 2013) while for Planck we use the CL49 temperature mask for $50 \leq l \leq 1200$ and the CL31 mask for $l > 1200$ (Planck XV 2014). We add white noise to each spectrum as described below, then we form a simple Gaussian, diagonal likelihood for each experiment using,

$$-2 \ln \mathcal{L} = \chi^2 = \sum_{l=l_{\min}}^{l_{\max}} \frac{(C_l^{\text{sim}} - C_l^{\text{th}})^2}{\sigma_l^2}, \quad (2)$$

where C_l^{sim} is the simulated C_l realization, C_l^{th} is the model spectrum to be fit, and σ_l^2 is the variance per l .

⁵ <http://www2.iap.fr/users/hivon/software/PolSpice/>

For *WMAP* we take $32 \leq l \leq 1200$ and for *Planck* we take $50 \leq l \leq 2500$. In addition to equation 2, we account for low- l polarization data by imposing a Gaussian prior on τ , as discussed in more detail below. In the end, we generate a database of 2×2000 maximum-likelihood parameters for each experiment and simulation realization.

Before discussing the results, we describe the construction of our (diagonal) power spectrum covariance matrix in more detail, and establish that it provides an adequate approximation to the full likelihood function for each experiment. The form of our matrix neglects mask-induced anti-correlations between nearby l values, but since the theory power spectra, C_l^{th} , are smooth functions of l , the effect of the off-diagonal elements can be approximated by summing the variance contributions in the neighborhood of each diagonal element. Given the full covariance matrix, $N_{ll'}$ for a given experiment, we can form the approximate diagonal matrix as

$$\sigma_l^2 = \sum_i N_{l-i, l+i} + \frac{1}{2} \sum_i N_{l-i+1, l+i} + \frac{1}{2} \sum_i N_{l-i, l+i+1}. \quad (3)$$

For *WMAP*, the full covariance is taken from the nine-year *WMAP* likelihood code and the sum over i runs over the full extent of the matrix (excluding $l = 0, 1$). For *Planck*, we form a frequency-averaged, foreground-removed covariance matrix from the full matrix provided in their likelihood code (described in Appendix C), and we restrict the sum to $-10 \leq i \leq +10$ for the first term, $-9 \leq i \leq +10$ for the second term, and $-10 \leq i \leq +9$ for the last term.

This approximate diagonal matrix correctly accounts for off-diagonal contributions to the variance in the limit that structure in the spectra are smooth. In order to determine the self-consistent level of instrument noise to add to each simulated spectrum, we estimate the cosmic variance contribution, $\sigma_{l, cv}^2$, then add Gaussian random noise with variance $\sigma_{l, n}^2 = \sigma_l^2 - \sigma_{l, cv}^2 \geq 0$. The effective cosmic variance is determined by filtering each simulated spectrum with a $\Delta l = 21$ boxcar filter (the smoothing scale applied in equation 3), multiplying by $\sqrt{21}$ to correct for the *rms* suppression induced by smoothing, then taking a standard deviation at each l over all 2000 realizations for each experiment.

Because the *Planck* Collaboration has not yet completed their low- l polarization analysis, they utilize *WMAP* polarization maps to constrain the optical depth, τ . As a result, the Gaussian prior we impose has the same peak for both experiments, so our τ statistics are highly correlated between the two. Further, since we use only temperature data in our approximate likelihood, the simulations only constrain the combination $A_s \exp(-2\tau)$, and not A_s or τ independently. As a result, we find that the recovered value of τ in any given realization is essentially set by the peak of the input τ prior. In order to mimic the true uncertainty in τ , while avoiding the (costly) use of low- l polarization data, we vary the τ prior applied to both likelihoods in each realization by randomly selecting the peak value from a Gaussian distribution with a mean value of 0.0851 and a standard deviation of 0.013. Note that the width of our prior is unimportant because we always recover the peak value,

but the scatter from one realization to the next correctly approximates the true uncertainty for each experiment.

The diagonal likelihood function is rapidly evaluated once C_l^{sim} and C_l^{th} are known. However, we require a concomitant speedup in the evaluation of $C_l^{\text{th}}(\mathbf{p})$ given the cosmological parameters \mathbf{p} in order to make the Monte Carlo-based comparison feasible. To achieve this, we use a custom python implementation of the PICO algorithm which casts $C_l^{\text{th}}(\mathbf{p})$ as a 4th-order polynomial within a specified volume of \mathbf{p} around the best fit region. We describe the details of the implementation in Appendix D and note here that our polynomial form of $C_l^{\text{th}}(\mathbf{p})$ was found to be within $0.09 \mu\text{K}^2$ (median absolute deviation) of the spectrum evaluated with CAMB.

4.2. Simulation testing

Figures 12 and 13 provide a visualization of our simulation fidelity by comparing our simulated parameter covariance matrices to those derived from the flight data, for *WMAP* and *Planck*, respectively. The black traces are derived from the Markov chains supplied by each experiment, the grey traces are derived from our simulations, as described below. The diagonal panels show Gaussian PDFs with a variance given by the corresponding elements of the covariance matrix. The off-diagonal panels show the 68% confidence region of the Gaussian PDF, marginalized over the four remaining parameters; these panels highlight the correlations among the parameters.

The simulation covariances shown in Figures 12 and 13 were derived from a Fisher matrix analysis of the diagonal likelihood in Equation 2. Enumerating the cosmological parameters as p_α for $\alpha = 1 \dots 6$, we form the matrix of spectrum derivatives,

$$M_{\alpha l} \equiv \left. \frac{\partial C_l}{\partial p_\alpha} \right|_{p_\beta \text{ fixed}}, \quad (4)$$

over the range of l used in each likelihood, and then form the 6×6 inverse parameter covariance matrix as

$$N_{\alpha\beta}^{-1} = M_{\alpha l} N_{ll'}^{-1} (M^T)_{l'\beta}. \quad (5)$$

The confidence regions derived from $N_{\alpha\beta}$ are shown in Figures 12 and 13 for *WMAP* and *Planck*.

To test the validity of our simulations further, we rescale the simulation covariance matrices as described below and evaluate the sensitivity of the inferred χ^2 to our choice of rescaling. We start by evaluating the full 12×12 covariance matrix, \mathbf{C} , of the combined *WMAP* and *Planck* parameters, including the off-diagonal blocks that measure the correlations between the two. We generate scaling factors by comparing the simulation variances to the data, $s_i^2 = C'_{ii}/C_{ii}$, as described below, and rescale the full covariance matrix according to $C'_{ij} = C_{ij} s_i s_j$. To project this down to the parameter differences we form the 6×6 matrix, $\Delta\mathbf{C}'$ as

$$\Delta\mathbf{C}' = (\mathbf{I} \quad -\mathbf{I}) \begin{pmatrix} \mathbf{C}'_{WW} & \mathbf{C}'_{WP} \\ \mathbf{C}'_{PW} & \mathbf{C}'_{PP} \end{pmatrix} \begin{pmatrix} \mathbf{I} \\ -\mathbf{I} \end{pmatrix}, \quad (6)$$

where each term is a 6×6 matrix, \mathbf{I} is the identity matrix, and the subscripts W and P refer to *WMAP* and *Planck*, respectively. Results with and without this rescaling of

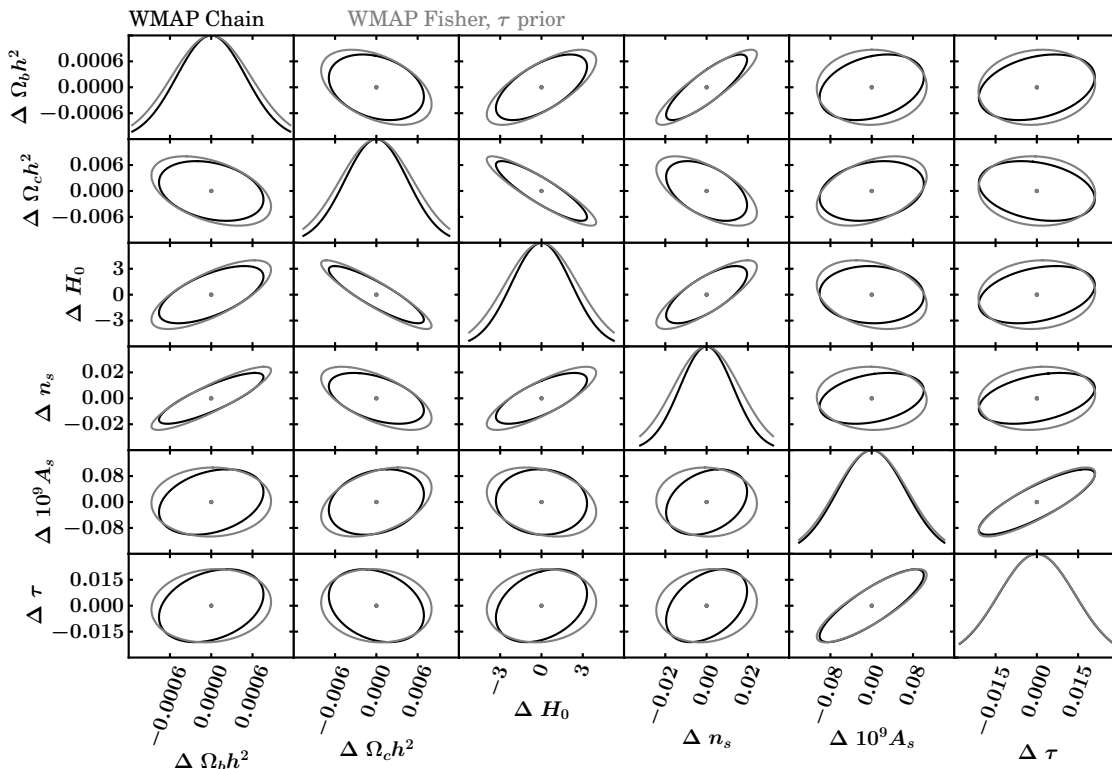


FIG. 12.— WMAP parameter uncertainties derived from the full WMAP Λ CDM chain (black) and from a Fisher analysis of the approximate diagonal WMAP likelihood used in this paper (grey). The approximate likelihood includes a τ prior with scatter σ_τ chosen to mimic the WMAP low l polarization likelihood. The scalar amplitude A_s is measured at $k_0 = 0.05 \text{ Mpc}^{-1}$ and the Hubble constant H_0 has units of $\text{km s}^{-1} \text{ Mpc}^{-1}$. Note that the projected simulation uncertainties are slightly larger than those derived from the full likelihood, hence these uncertainties are likely conservative.

the parameter difference covariance matrix are summarized in §4.3 below.

4.3. Results

We use the 6×6 covariance matrix of parameter differences, with and without rescaling, to assess the consistency of the measured parameter difference. The results are shown in Figure 14. When we do not rescale the simulations at all we find $\chi^2 = 85.3$. Upon rescaling each of the 12 standard deviations to match the published errors, we find $\chi^2 = 56$. To assess possible systematic uncertainty in the scaling, we also sample the 12 parameter standard deviations uniformly in the interval between the published data value and the simulation result. The resulting distribution of χ^2 values is well approximated by a Gaussian with a mean of 69.2 and a standard deviation of 6.6. The lowest χ^2 observed in this test was 38.8. All of these values have extremely low probabilities to exceed (PTE): $\chi^2 = 38.8$ has $\text{PTE} < 10^{-6}$ for six degrees of freedom. Our best estimate is $\chi^2 = 56$, which lies well within the extrema of 39 to 85.

We conclude that the differences between the two parameter measurements cannot be attributed to the effects of different sky cuts, multipole ranges, or instrument noise. The most striking pair-wise discrepancy lies in the sub-space $(\Omega_c h^2, A_{s,0.05})$, where *Planck* observes a higher value of $\Omega_c h^2$ than WMAP, but a nearly identical value of $A_{s,0.05}$, causing the pair-wise difference to lie well outside the 68% CL ellipse. To probe this difference further, we marginalize χ^2 over these two parameters, both individually and together, to assess the consistency

of the remaining parameters; the results are summarized in Table 7. For reference, the first row shows the high χ^2 value for the full 6 parameter set. Next, we show the results of marginalizing over $A_{s,0.05}$ and $\Omega_c h^2$: marginalizing over either single parameter drops χ^2 to a reasonable value. Marginalizing over any of the other 4 single parameters does not substantially reduce χ^2 . Next we show that marginalizing over both parameters gives a low χ^2 (but with a still-reasonable PTE). Finally, we show that marginalizing over the 4 parameters $\{\Omega_b h^2, H_0, n_s, \tau\}$ leaves a high χ^2/ν with a PTE of just 0.26% for the pair $\{\Delta A_{s,0.05}, \Delta \Omega_c h^2\}$.

The last column of Table 6 shows single parameter differences and their uncertainties, each after marginalizing over the other 5 cosmological parameters. We note that they all individually agree, with the largest difference being 1.7σ for $\Omega_c h^2$; curiously, $A_{s,0.05}$ differs by only 0.1σ . Collectively, the WMAP and *Planck* best-fit parameters are discrepant by $\sim 6\sigma$, with $\Omega_c h^2$ being the most discrepant single parameter.

Table A.1 of *Planck XVI* (2014) compares cosmological parameters derived from WMAP data to parameters derived from the *Planck* data restricted to $l < 1000$. In that comparison, they find that $\Omega_c h^2$ is more consistent ($\Delta \Omega_c h^2 = 0.0005$) and that the amplitude parameters differ by 1.2%, but still not the 2.5% one would naively expect from the spectrum ratios. The remainder of the spectrum difference is accounted for by small changes in the remaining Λ CDM parameters. Thus, at least some of the parameter tension appears to arise from multipoles $l \gtrsim 1000$.

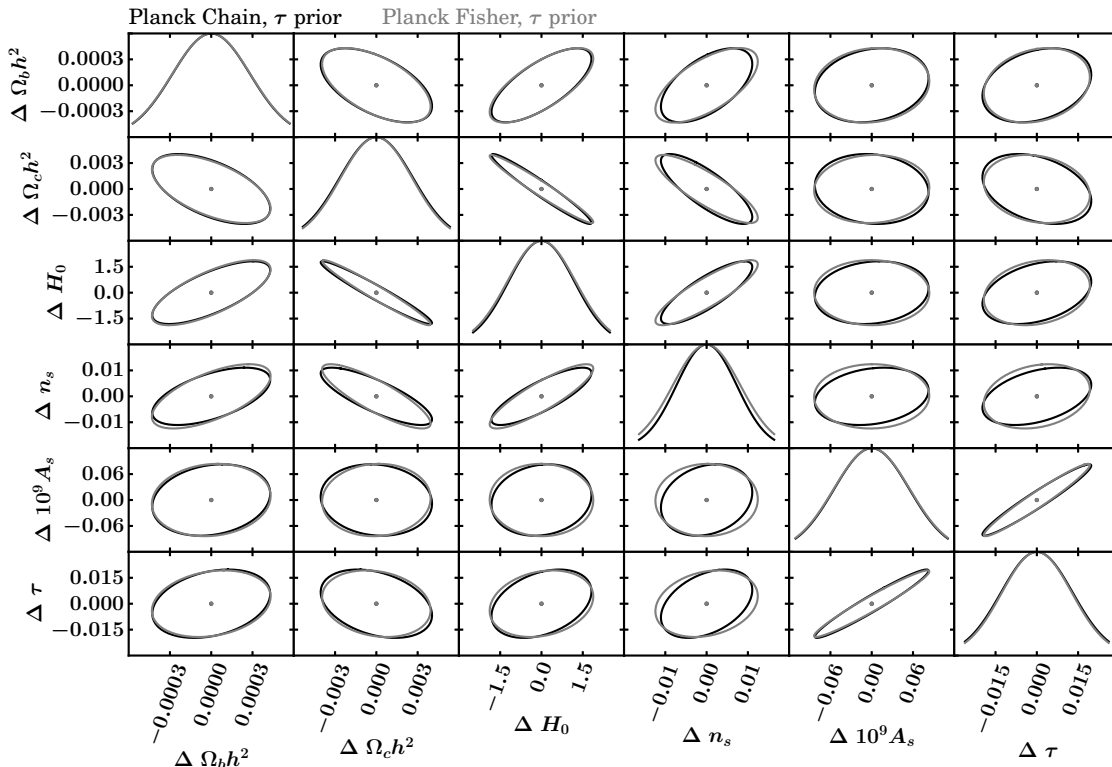


FIG. 13.— *Planck* parameter uncertainties derived from the full *Planck*+WMAPpol Λ CDM chain “PLA/base/planck_lowl_lowLike” (black) and from a Fisher analysis of the approximate diagonal *Planck* likelihood used in this paper (grey). The Fisher analysis includes a τ prior with scatter $\sigma_\tau = 0.013$ to mimic the *WMAP* low l polarization likelihood used in the chain, and it also includes marginalization over the thermal SZ template provided in the *Planck* likelihood. The scalar amplitude A_s is measured at $k_0 = 0.05 \text{ Mpc}^{-1}$ and the Hubble constant H_0 has units of $\text{km s}^{-1} \text{ Mpc}^{-1}$. The approximate uncertainties we derive from the diagonal likelihood agree closely with those obtained from the full *Planck* chain, showing that the diagonal likelihood is a useful approximation to the full likelihood.

TABLE 7
PARAMETER CONSISTENCY TESTS

Parameter(s) marginalized	χ^2	ν	χ^2/ν	PTE ^a
None	55.9	6	9.3	$< 10^{-6}$
$A_{s,0.05}$	6.3	5	1.3	0.28
$\Omega_c h^2$	5.2	5	1.0	0.39
$A_{s,0.05}, \Omega_c h^2$	2.2	4	0.6	0.70
$\Omega_b h^2, H_0, n_s, \tau$	11.9	2	6.0	0.0026

^a Probability to exceed, accounting for degrees of freedom, ν .

4.4. Discussion

For full-sky experiments like *WMAP* and *Planck*, measurements of the angular power spectrum rely on calibrating the overall brightness scale to the CMB dipole. This calibration is transferred to smaller angular scales by means of a window function, W_l , which mainly depends on the beam response, but it can also depend on the filtering applied to the time-ordered data (e.g. the *Planck* cosmic ray correction procedure). Given the similarity of the *WMAP* and *Planck* dipoles, the differences between the two inferred power spectra are most likely explained by errors in the window functions of one or both experiments.

For *WMAP*, the window function is dominated by the beam response of the experiment. In §3.2 we review the steps that lead to the *WMAP* brightness calibration and window function and conclude that the power spectrum cannot be altered by more than a fraction of its current

uncertainty of $\sim 0.6\%$ in amplitude ($\sim 1.2\%$ in power), and that the most likely sources of error would *increase* the power spectrum amplitude slightly.

Given its more sensitive instruments, the *Planck* window function must account for more effects than does *WMAP*’s. We do not have access to the data required to assess these effects, but in §2.2 we argue that certain features in the sky map differences resemble recognized far sidelobe features in the *Planck* data. We speculate that this may indicate that residual errors remain in the overall *Planck* window function. We note that *WMAP* was able to measure its beams in flight using 17 seasons of crisscrossed Jupiter observations, while the *Planck* observing pattern was more restricted, leading to a somewhat higher reliance on beam models. Further, the HFI detectors saturate on Jupiter, so the HFI beam measurements must employ a combination of sky sources to cover the full dynamic range of the beam response.

To explore the sensitivity of cosmological parameters to *WMAP* and *Planck* beam errors, we perform the following exercise: we take the largest beam error mode from each experiment’s likelihood function, and form a modified window function for each experiment. For *Planck* we use the 143×143 GHz beam error provided with the *Planck* likelihood, since it covers the largest l range; the *WMAP* likelihood provides a single set of beam error modes, so no choice is required. For each case, we refit Λ CDM parameters to the modified spectra. The *WMAP* beam error couples primarily to $A_{s,0.05}$ and τ , due to the steepness of the beam error rise at $l \sim 100$. The *Planck* beam error couples to almost all of

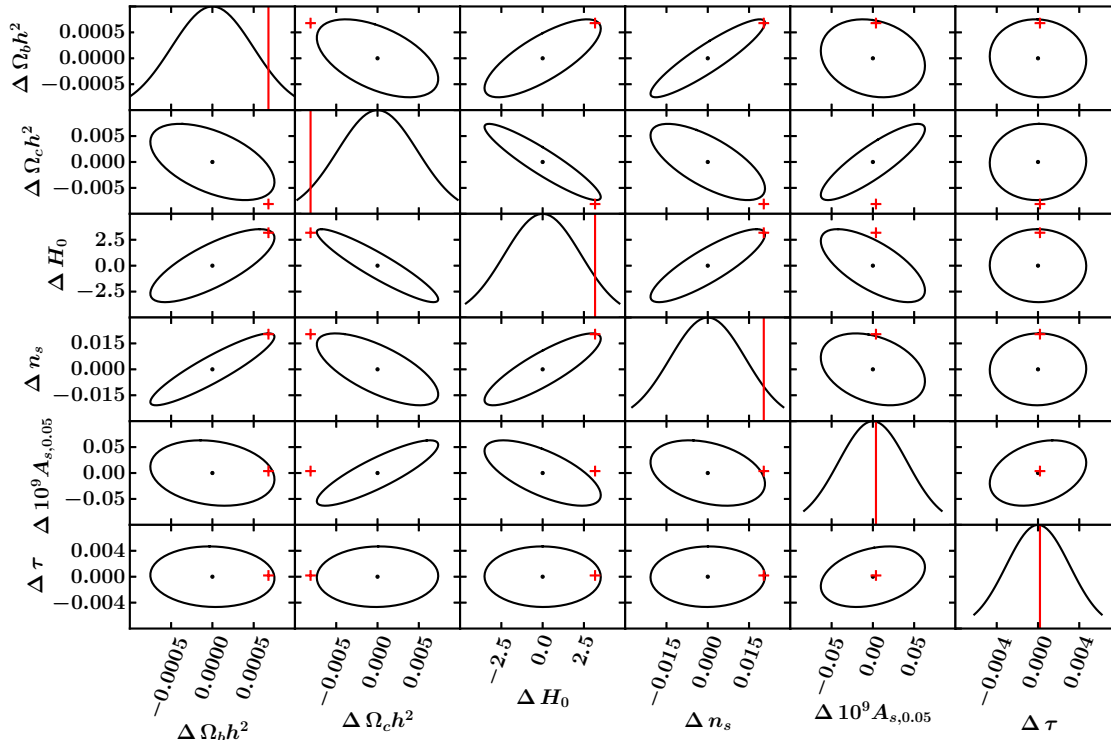


FIG. 14.— A comparison of the 6 Λ CDM model parameters derived from the *Planck* and *WMAP* data, respectively. The vertical lines in the diagonal panels show the measured difference for each parameter separately; the plus symbols in the off-diagonal panels show the measured parameter differences pairwise. The curves in the diagonal panels show the likelihood of the single parameter difference, taking into account the fact that both satellites observe a common sky, but employ different sky cuts, have different l ranges, and have different instrument noise. The ovals in the off-diagonal panels show the pair-wise covariance of the expected parameter differences. Note that each single measured parameter is consistent within the individual uncertainties, but several pairs of parameter differences lie outside their 68% CL ellipses, particularly the pair $A_{s,0.05}$ and $\Omega_c h^2$. Taken as a whole, the χ^2 for the *Planck*–*WMAP* parameter difference is 56 for 6 degrees of freedom, which has a probability to exceed of $\sim 3 \times 10^{-10}$. This indicates that the two experiments infer significantly different sets of Λ CDM parameters.

the parameters at a similar level, including $\Omega_c h^2$. Based on this exercise, it seems unlikely that a *WMAP* dominant mode beam error would affect $\Omega_c h^2$, while a *Planck* beam error might.

5. CONCLUSIONS

In this paper we report on consistency checks of the *Planck* and *WMAP* data. At first glance, the *WMAP* and *Planck* data agree remarkably well, however, there exist small but significant differences between them. We summarize our conclusions below.

(1) Examining the raw *Planck* – *WMAP* difference maps 30–Ka, 44–Q, 70–V and 100–W (in thermodynamic temperature units after smoothing to a common resolution of 2°), we find that much of the CMB anisotropy signal cancels, but calibration differences between the two experiments leave residual dipole differences because the dipole is a relatively strong signal. The [44,Q], [70,V] and [100,W] dipole differences are well aligned with the CMB dipole direction, with amplitudes of 0.3%, 0.3% and 0.5% of the CMB dipole, respectively. All are positive in the south and negative in the north, consistent with the *WMAP* maps being brighter than the *Planck* maps. These residuals are within the quoted uncertainties for both the *Planck* and *WMAP* data (Planck III 2014; Planck VIII 2014; Bennett et al. 2013) and are insufficient to explain the power spectrum differences at

higher multipoles, by at least a factor of 2. Since the effective frequencies for each pair of maps are not identical there is also imperfect cancellation of foregrounds in the [44,Q], [70,V] and [100,W] difference maps.

(2) After foreground template cleaning, the LFI–*WMAP* differences exhibit a quadrupole component higher than expected based on the different treatments of the kinematic quadrupole. The HFI 100 GHz – *WMAP* W band difference map exhibits large-scale systematic structure that is not solely limited to a quadrupolar signature. We identify this systematic structure as associated with the *Planck* map, possibly originating from beam side-lobe pick-up effects, either directly or indirectly. Note that the *Planck* Collaboration did not remove far side-lobe pickup from the delivered 100, 143, 217, or 353 GHz maps so some residual is expected.

(3) We form a template-independent, foreground-reduced difference map using an extension of the internal linear combination technique. Again, we find features that are plausibly associated with residual *Planck* systematic errors. Lacking the full *Planck* time-ordered data we cannot fully address these features.

(4) Over a wide range of angular scales, the fluctuations measured by *WMAP* are $\sim 1.2\%$ brighter than measured by *Planck* ($\sim 2.5\%$ in power). On the largest scale

probed—the dipole moment—the experiments agree to within the uncertainties. This agreement may persist up to multipoles of a few 10’s, as seen in Figure 1, but residual foreground differences arising from bandpass differences make it difficult to draw definitive conclusions on these scales. In model fits, the amplitudes of the first acoustic peak at $l = 220$ differ by $\sim 150 \mu\text{K}^2$ ($\sim 2.7\%$) between the two experiments. Using simulations, we examine the best-fit value of $D_{220} \equiv l(l+1)C_l/(2\pi)|_{l=220}$ for an ensemble of ΛCDM spectra fit to the *WMAP* and *Planck* data, and estimate the significance of this difference to be $3 - 5\sigma$, depending on the treatment of the window function uncertainties.

(5) We have re-examined the *WMAP* beam measurements and analysis seeking an explanation for the observed difference between the *Planck* and *WMAP* power spectra at intermediate angular scales. We conclude that residual errors in characterizing the *WMAP* beam response cannot plausibly explain the $\sim 2.5\%$ power spectrum difference.

(6) We have evaluated the consistency of the cosmological parameters inferred from each experiment, taking into account the similarities and differences between them. They observe the same sky but they cover different multipole ranges, employ different sky masks, and have different instrument noise. We develop simulations to account for these effects and, for the first time to our knowledge, derive a 6×6 covariance matrix of parameter differences. The 6 ΛCDM parameters derived from the two experiments are individually consistent with each other, after marginalizing over the remaining 5 parameters. The largest difference we find is $\Delta\Omega_c h^2 = -0.0081 \pm 0.0049$ (1.7σ). However the full set of 6 ΛCDM parameters are discrepant at the $\sim 6\sigma$ level, with a $\chi^2 = 56$ for 6 degrees of freedom (PTE = 3×10^{-10}). We have estimated the uncertainty of this χ^2 in §4.3.

(7) The $\sim 6\sigma$ *Planck*-*WMAP* parameter discrepancy is driven by two parameters: the cold dark matter density, $\Omega_c h^2$, and the power spectrum amplitude, $A_{s,0.05}$. The χ^2 of the pair-wise difference $\{\Delta A_{s,0.05}, \Delta\Omega_c h^2\}$ is 11.9 for two degrees of freedom, with a PTE of just 0.26%. Figure 14 shows that $\Omega_c h^2$ and $A_{s,0.05}$ are somewhat degenerate over the range of multipoles probed by *WMAP*, and that the fits prefer to boost the value of $\Omega_c h^2$ rather than decrease $A_{s,0.05}$ to fit the lower *Planck*

spectrum. Marginalizing over either $A_{s,0.05}$ or $\Omega_c h^2$ separately renders the remaining five parameters consistent. Our analysis results are consistent with [Hazra & Shafieloo \(2014b\)](#), but demonstrate that other parameter combinations provide comparable or statistically preferable fits.

(8) We explored the sensitivity of cosmological parameters to *WMAP* and *Planck* beam errors with a simple exercise: we constructed a modified window function for each experiment (§4.4) and refit the ΛCDM parameters to correspondingly modified power spectra. We found that the dominant *WMAP* beam error coupled mostly to $A_{s,0.05}$ and τ , while the dominant *Planck* beam error coupled to most of the parameters, including $\Omega_c h^2$.

The disagreement between the *Planck* and *WMAP* data could signal a failure of the 6-parameter ΛCDM model, and hence point to evidence for new physics. However, it is of primary importance to first fully understand and account for systematic measurement errors such as those apparent in the sky map and power spectrum differences reported here and by the *Planck* Collaboration. We note that [Planck XXXI \(2014\)](#) reports some changes in the *Planck* window function relative to the 2013 versions used in this paper. This update brings the *Planck* spectrum closer to *WMAP*’s, but a $\sim 2\%$ discrepancy remains near the first acoustic peak. Once this difference is resolved, it will be interesting to see if small differences in the ΛCDM parameters persist.

This research was supported in part by NASA grant NNX14AF64G and by the Canadian Institute for Advanced Research (CIFAR). We thank the *Planck* team for providing us with a copy of [Planck XXXI \(2014\)](#) in advance of its publication. We acknowledge use of the HEALPix ([Gorski et al. 2005](#)) and CAMB ([Lewis et al. 2000](#)) packages. This research has made use of NASA’s Astrophysics Data System Bibliographic Services. We acknowledge the use of the Legacy Archive for Microwave Background Data Analysis (LAMBDA), part of the High Energy Astrophysics Science Archive Center (HEASARC). HEASARC/LAMBDA is a service of the Astrophysics Science Division at the NASA Goddard Space Flight Center. We also acknowledge use of the *Planck* Legacy Archive. *Planck* is an ESA science mission with instruments and contributions directly funded by ESA Member States, NASA, and Canada.

APPENDIX

A. HARMONIC CONVENTIONS

In this paper we quote coefficients for monopoles, dipoles, and quadrupoles using real-valued spherical harmonics. To make the numbers more readily understandable, we use basis functions that peak at ± 1 for the monopole and dipole. Our monopole basis function is 1 everywhere, and not $Y_{0,0} = 1/\sqrt{4\pi}$, so that when we quote a monopole coefficient it is the monopole offset of the map, and not that offset divided by $\sqrt{4\pi}$. Likewise, our x , y , and z dipole basis functions are $\cos(\phi)\sin(\theta)$, $\sin(\phi)\sin(\theta)$, and $\cos(\theta)$, respectively, where θ and ϕ are the usual co-latitude and longitude angular coordinates. This means that a dipole-only map represented by $T(\theta, \phi) = a x(\theta, \phi) + b y(\theta, \phi) + c z(\theta, \phi)$ has a peak value of $\sqrt{a^2 + b^2 + c^2}$ on the sky. This matches the conventions in the *WMAP* papers where 3.355 mK is the peak value of the measured dipole on the sky ([Hinshaw et al. 2009](#)).

For the quadrupole and higher modes, we switch to the usual orthonormalization convention for the spherical

harmonics. This is in an attempt to match the sign conventions in Table 7 of [Hinshaw et al. \(2007\)](#)⁶ and Table III of [de Oliveira-Costa et al. \(2004\)](#).

A software package such as Healpy is able to calculate spherical harmonic coefficients a_{lm} so that the temperature field can be described as

$$T(\theta, \phi) = \sum_{l=0}^{l_{\max}} \sum_{m=-l}^l a_{lm} Y_{lm}(\theta, \phi) \quad (\text{A1})$$

See for example Section 5.2 of [Varshalovich et al. \(1988\)](#) for a definition of the spherical harmonics $Y_{l,m}(\theta, \phi)$, and Section 5.13 of that book for their formulas written out explicitly. For real-valued fields T , we have the relation

$$a_{l,-m} = (-1)^m a_{l,m}^* \quad (\text{A2})$$

which means the a_{lm} coefficients for negative m are completely determined by those for positive m . Healpy only stores a_{lm} for $m \geq 0$, to save space. For ease of use, we define real-valued coefficients as follows.

$$\tilde{a}_{lm} = \begin{cases} \sqrt{2} \operatorname{Im}[a_{l|m|}] & \text{if } m < 0 \\ a_{lm} & \text{if } m = 0 \\ \sqrt{2} \operatorname{Re}[a_{lm}] & \text{if } m > 0 \end{cases} \quad (\text{A3})$$

For an alternative definition of real-valued spherical harmonics, see [Chisholm \(1976\)](#) and [Blanco et al. \(1997\)](#), although Chisholm adds an extra “c” or “s” index to the harmonics instead of allowing negative values of m , so he does not specify a natural ordering of the coefficients.

To determine the basis functions corresponding to these \tilde{a}_{lm} coefficients, consider a real-valued map that contains only two modes for a specific l , and for m and $-m$, with $m > 0$.

$$T = a_{lm} Y_{lm} + a_{l,-m} Y_{l,-m} \quad (\text{A4})$$

Since the temperature field is real-valued, we have

$$T = a_{lm} Y_{lm} + a_{l,m}^* Y_{l,m}^* = 2 \operatorname{Re}[a_{lm} Y_{lm}] = 2 \operatorname{Re}[a_{lm}] \operatorname{Re}[Y_{lm}] - 2 \operatorname{Im}[a_{lm}] \operatorname{Im}[Y_{lm}] \quad (\text{A5})$$

This reduces to

$$T = \tilde{a}_{lm} \cdot \sqrt{2} \operatorname{Re}[Y_{lm}] + \tilde{a}_{l,-m} \cdot (-\sqrt{2} \operatorname{Im}[Y_{lm}]) = \tilde{a}_{lm} \tilde{Y}_{lm} + \tilde{a}_{l,-m} \tilde{Y}_{l,-m} \quad (\text{A6})$$

where we have defined real-valued basis functions \tilde{Y}_{lm} as:

$$\tilde{Y}_{lm} = \begin{cases} -\sqrt{2} \operatorname{Im}[Y_{l|m|}] = (-1)^m \sqrt{2} \operatorname{Im}[Y_{l,-|m|}] & \text{if } m < 0 \\ Y_{lm} & \text{if } m = 0 \\ \sqrt{2} \operatorname{Re}[Y_{lm}] & \text{if } m > 0 \end{cases} \quad (\text{A7})$$

This allows us to expand a general real-valued temperature field in real-valued harmonics up to some maximum multipole l_{\max} as

$$T(\theta, \phi) = \sum_{l=0}^{l_{\max}} \sum_{m=-l}^l \tilde{a}_{lm} \tilde{Y}_{lm}(\theta, \phi). \quad (\text{A8})$$

Our dipole x, y, z and quadrupole $\tilde{Y}_{2,m}$ modes are shown in [Figure 15](#).

⁶ Errata for Table 7: Numbers in table for $l = 2, 3$ and $m \neq 0$ do not have $\sqrt{2}$ multiplied in. The $l < 0$, $l = 0$, and $l > 0$ in the notes should be $m < 0$, $m = 0$, and $m > 0$, respectively. We should have $\tilde{a}_{lm} = \sqrt{2} \operatorname{Im} a_{l|m|}$ for $m < 0$, using standard a_{lm} values from Healpy. Dipole normalization is for x, y, z harmonics that peak at ± 1 , not $\tilde{a}_{1,m}$ values, for numbers in table. Error bars on l, b for (d, l, b) triple should be switched in footnote a .




FIG. 15.— The real-valued dipole and quadrupole modes used in this paper. In the top row are the y , z , and x dipole modes, with the order chosen to align with the quadrupole modes in the bottom row. The dipole modes reach their maximum value of $+1$ (colored orange-brown) in the directions of the positive y , z , and x axes, respectively. From this it is clear that our Mollweide projection uses the astronomical convention that the observer is inside the sphere: θ increases downwards, and ϕ increases when moving right to left in this image. On the bottom row, real-valued quadrupole modes $\tilde{Y}_{2,m}$, defined in Appendix A, are ordered from $m = -2$ on the left to $m = 2$ on the right. Color scale is ± 1 for each image. Note that the dipole basis functions have a larger amplitude than the quadrupole basis functions due to our different choices of normalization. Also note that the $m = -1$ quadrupole mode is prominent in various *Planck*/*WMAP* differences, in galactic coordinates.

To relate the normalization convention for the \tilde{Y}_{1m} modes to our preferred dipole modes, we have:

$$x(\theta, \phi) = -\sqrt{\frac{4\pi}{3}}\tilde{Y}_{1,1}(\theta, \phi) \quad (\text{A9})$$

$$y(\theta, \phi) = \sqrt{\frac{4\pi}{3}}\tilde{Y}_{1,-1}(\theta, \phi) \quad (\text{A10})$$

$$z(\theta, \phi) = \sqrt{\frac{4\pi}{3}}\tilde{Y}_{1,0}(\theta, \phi) \quad (\text{A11})$$

Our new basis of \tilde{Y}_{lm} real-valued harmonics still has a similar normalization to the usual Y_{lm} harmonics,

$$\iint \tilde{Y}_{lm}\tilde{Y}_{l'm'} d\cos\theta d\phi = \delta_{ll'}\delta_{mm'}, \quad (\text{A12})$$

so we can quote the power in the quadrupole as the usual C_l value,

$$C_2 = \frac{1}{5} \sum_{m=-2}^2 (\tilde{a}_{2,m})^2 \quad \text{in } [\mu\text{K}^2]. \quad (\text{A13})$$

We also briefly discuss the kinematic quadrupole, which is of the form $Y_{2,0}$ and is already real-valued, in an appropriately rotated coordinate system. Since $Y_{2,0}(\theta, \phi) = \frac{1}{4}\sqrt{\frac{5}{\pi}}(3\cos^2(\theta) - 1)$, the peak-to-trough amplitude for this harmonic is $\frac{3}{4}\sqrt{\frac{5}{\pi}} = 0.946175$. Ignoring the frequency-dependent effects discussed in [Planck XXVII \(2014\)](#), the peak-to-trough amplitude of the kinematic quadrupole is $T_{\text{CMB}}(v/c)^2 \approx 4.1 \mu\text{K}$, where v is the velocity of the observer with respect to the CMB rest frame, and c is the speed of light. If the peak-to-trough distance is $4.1 \mu\text{K}$, the $a_{2,0}$ value is $4.1/0.946175 \approx 4.3 \mu\text{K}$. Since the kinematic quadrupole has no $m \neq 0$ power in our rotated reference frame, we have only one term in our sum for the kinematic quadrupole power: $C_2 = \frac{1}{5}a_{2,0}^2 \approx 3.8 \mu\text{K}^2$.

B. EXTENDED ILC METHOD

To extend our ILC method to look at systematic differences between experiments, we follow much the same algebra as in the appendix of [Eriksen et al. \(2004\)](#). As they do, we minimize the variance of a linear combination of maps, but we use the different constraints that the sum of weights for the first set of maps must be $+1$, and the sum of the weights for the second set of maps must be -1 .

The variance that we minimize is still their Equation A1:

$$\text{Var}(T) = \mathbf{w}^T \mathbf{C} \mathbf{w} \quad (\text{B1})$$

where \mathbf{w} is a vector of weights and \mathbf{C} is the covariance matrix of the maps T_i (which must all be in the same units) given by their Equation A2.

$$C_{ij} = \langle \Delta T_i \Delta T_j \rangle = \frac{1}{N_{\text{pix}}} \sum_{p=1}^{N_{\text{pix}}} (T^i(p) - \bar{T}^i)(T^j(p) - \bar{T}^j) \quad (\text{B2})$$

We wish to minimize their Equation A3

$$f(\mathbf{w}) = \sum_{ij} w_i C_{ij} w_j \quad (\text{B3})$$

subject to two constraints:

$$g(\mathbf{w}) = \sum_{i=1}^{n_1} w_i = 1, \quad (\text{B4})$$

and

$$h(\mathbf{w}) = \sum_{i=n_1+1}^{n_2} w_i = -1. \quad (\text{B5})$$

The vector \mathbf{w} has indices 1 through n_1 corresponding to the bands of the first experiment, and indices $n_1 + 1$ through n_2 corresponding to the bands of the second experiment. Any additional templates which should be removed from the ILC (and which presumably contain no CMB) can be added in with indices above n_2 ; these have unconstrained weights. If the vector \mathbf{w} has n indices, then indices $n_2 + 1$ to n correspond to extra templates. Eriksen's Equation A5 is now modified for our constraints:

$$\nabla f(\mathbf{w}_0) = \lambda_1 \nabla g(\mathbf{w}_0) + \lambda_2 \nabla h(\mathbf{w}_0) \quad (\text{B6})$$

We convert this, in addition to the two direct constraints on the weights above, into the $n + 2$ by $n + 2$ matrix equation written out in blocks:

$$\begin{pmatrix} 2\mathbf{C} & -\mathbb{1} \\ \Leftrightarrow & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{w} \\ \lambda_1 \\ -\lambda_2 \end{pmatrix} = \begin{pmatrix} \mathbf{0} \\ 1 \\ -1 \end{pmatrix}, \quad (\text{B7})$$

where \Leftrightarrow is the $2 \times n$ matrix of 1s and 0s that encodes the constraints on the weights, and $\mathbb{1} = \Leftrightarrow^T$. For example, if we have 5 *WMAP* bands so that $n_1 = 5$, and we have 7 more *Planck* bands so that $n_2 = 12$, and we have three templates so that $n = 15$, then we use

$$\Leftrightarrow = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 \end{pmatrix}. \quad (\text{B8})$$

The matrix in Equation B7 can be constructed directly and inverted to solve for \mathbf{w} , λ_1 , and λ_2 . Since the matrix is so small (17×17 in the above example), the matrix inversion is a rapid operation.

If the matrix in Equation B7 is singular, then there is probably a linear dependence among the maps or templates ΔT_i , making the matrix \mathbf{C} singular as well. One way to remedy this is to remove one or more of the maps or templates from the calculation until the matrix \mathbf{C} becomes invertible again, allowing a unique solution to Equation B7. This is only likely to be a problem if a map has been duplicated in the calculation of \mathbf{C} ; the independent noise from maps at different frequencies tends to make the maps linearly independent.

C. SIMPLIFYING THE PLANCK LIKELIHOOD

Planck has released a set of cross power spectra, for 100×100 , 143×143 , 143×217 , and 217×217 GHz, extractable from their released likelihood code. These use l ranges of 50–1200, 50–2000, 500–2500, and 500–2500, respectively. *Planck* provides the measured values of these power spectra and their 7104×7104 element covariance matrix.

Four spectra are useful when studying foregrounds in our galaxy, but for cosmology we only want one: the spectrum of the primordial CMB, characterized by the frequency dependence of a blackbody of constant temperature. Therefore, we find the best fit power spectrum given the above information. This is the spectrum \hat{C} which is the best fit to the below matrix equation, explicitly written out in blocks:

$$\begin{pmatrix} \mathbb{I}_{450} & 0_{450 \times 701} & 0_{450 \times 800} & 0_{450 \times 500} \\ 0_{701 \times 450} & \mathbb{I}_{701} & 0_{701 \times 800} & 0_{701 \times 500} \\ \mathbb{I}_{450} & 0_{450 \times 701} & 0_{450 \times 800} & 0_{450 \times 500} \\ 0_{701 \times 450} & \mathbb{I}_{701} & 0_{701 \times 800} & 0_{701 \times 500} \\ 0_{800 \times 450} & 0_{800 \times 701} & \mathbb{I}_{800} & 0_{800 \times 500} \\ 0_{701 \times 450} & \mathbb{I}_{701} & 0_{701 \times 800} & 0_{701 \times 500} \\ 0_{800 \times 450} & 0_{800 \times 701} & \mathbb{I}_{800} & 0_{800 \times 500} \\ 0_{500 \times 450} & 0_{500 \times 701} & 0_{500 \times 800} & \mathbb{I}_{500} \\ 0_{701 \times 450} & \mathbb{I}_{701} & 0_{701 \times 800} & 0_{701 \times 500} \\ 0_{800 \times 450} & 0_{800 \times 701} & \mathbb{I}_{800} & 0_{800 \times 500} \\ 0_{500 \times 450} & 0_{500 \times 701} & 0_{500 \times 800} & \mathbb{I}_{500} \end{pmatrix} \begin{pmatrix} \hat{C}_{50-499} \\ \hat{C}_{500-1200} \\ \hat{C}_{1201-2000} \\ \hat{C}_{2001-2500} \end{pmatrix} = \begin{pmatrix} C_{100 \times 100}^{50-499} \\ C_{100 \times 100}^{500-1200} \\ C_{143 \times 143}^{50-499} \\ C_{143 \times 143}^{500-1200} \\ C_{143 \times 143}^{1201-2000} \\ C_{143 \times 217}^{500-1200} \\ C_{143 \times 217}^{500-1200} \\ C_{143 \times 217}^{1201-2000} \\ C_{143 \times 217}^{1201-2000} \\ C_{2001-2500}^{217 \times 217} \\ C_{217 \times 217}^{500-1200} \\ C_{217 \times 217}^{500-1200} \\ C_{1201-2000}^{217 \times 217} \\ C_{217 \times 217}^{2001-2500} \end{pmatrix}, \quad (\text{C1})$$

where $0_{n \times m}$ is a matrix of zeros with n rows and m columns, and \mathbb{I}_n is an $n \times n$ identity matrix. Also, \hat{C}_{n-m} is a column vector with $m - n + 1$ elements, corresponding to a best fit power spectrum with multipole l values ranging

from n to m , inclusive. The frequency dependent cross spectra $C_{n-m}^{x \times y}$ are cross spectra of x GHz and y GHz maps also with multipole l values ranging from n to m , inclusive. More compactly, this is

$$A\hat{C} = C, \quad (\text{C2})$$

where A is a 7104×2451 matrix of 1s and 0s, \hat{C} is the 2451 element column vector for which we are solving, and C is a 7104 element column vector of measured C_l values at various frequencies. If N is the covariance matrix of C , then the linear least squares solution is:

$$\hat{C} = (A^T N^{-1} A)^{-1} A^T N^{-1} C. \quad (\text{C3})$$

Since the measured power spectrum at different frequencies has different levels of foregrounds, we choose to consider foreground-related spectra to be completely unconstrained modes. This means making sure that N^{-1} acting on a pure-foreground-spectrum mode zeros it. Let F be a matrix with foreground spectra in columns. This assumes that the *Planck* foregrounds can be constructed by a linear combination of templates. Since there are some nonlinear features to the *Planck* likelihood foreground model (for example, the beam mode template and calibration factors multiply the other foreground templates), we linearize around reasonable values of these calibration and foreground parameters. To include these foreground templates into Equation C3, we replace N with $N + \infty F F^T$, which means replacing N^{-1} with $(N + \infty F F^T)^{-1}$. This should be done with the Sherman-Morrison-Woodbury formula, since we already have N^{-1} .

$$N' = N + \sigma_f^2 F F^T \quad (\text{C4})$$

$$(N')^{-1} = N^{-1} - N^{-1} F (\sigma_f^{-2} \mathbb{I} + F^T N^{-1} F)^{-1} F^T N^{-1} \quad (\text{C5})$$

We take the limit $\sigma_f \rightarrow \infty$ and find

$$(N')^{-1} = N^{-1} - N^{-1} F (F^T N^{-1} F)^{-1} F^T N^{-1}. \quad (\text{C6})$$

Note that this makes the matrix $A^T (N')^{-1} A$ singular, so we have to take the pseudo-inverse when finding the linear least squares solution. This is equivalent to projecting the foreground modes out of our final spectrum.

We explicitly construct $\hat{N} = (A^T (N')^{-1} A)^{-1}$, since this is the (singular) covariance matrix of our power spectrum \hat{C} . With the pseudo-inverse, foreground modes are zeroed by this covariance matrix⁷, but it is otherwise accurate and useful.

For reference, we expand the solution for the foreground-removed single power spectrum \hat{C} in terms of the original matrices A , N , and F :

$$\hat{C} = (A^T [N^{-1} - N^{-1} F (F^T N^{-1} F)^{-1} F^T N^{-1}] A)^{-1} A^T [N^{-1} - N^{-1} F (F^T N^{-1} F)^{-1} F^T N^{-1}] C. \quad (\text{C7})$$

D. POLYNOMIAL ESTIMATION OF C_L^{TH}

We randomly sample 10,000 sets of six parameters, using a uniform distribution over the following intervals: $\Omega_b h^2 \in [0.020, 0.025]$, $\Omega_c h^2 \in [0.10, 0.15]$, $H_0 \in [60, 80]$ km s⁻¹ Mpc⁻¹, $n_s \in [0.8, 1.1]$, $A_s \in [2.0 \times 10^{-9}, 2.5 \times 10^{-9}]$ for $k_0 = 0.05$ Mpc⁻¹, and $\tau \in [0.03, 0.20]$. These intervals are chosen to encompass current parameter measurements, while being small enough to allow a fourth-order polynomial to fit the C_l^{th} values over the given range. For each set of parameters, we calculate the lensed spectrum using CAMB, as described below. For each multipole in the range $2 \leq l \leq 2500$, we fit $l(l+1)C_l^{\text{th}}/2\pi$ to a fourth-order polynomial, using the first 9000 members of the sampled parameters; we save the remaining 1000 members for testing, as described below. With 6 variables, a fourth-order polynomial has 210 terms⁸, so it is well constrained when fit to 9000 samples. We check the accuracy of the polynomial prediction for CAMB with the remaining 1000 samples. We find the maximum absolute value of the error (for any l) to be $|\Delta l(l+1)C_l^{\text{th}}/2\pi| < 4.5 \mu\text{K}^2$, with the median absolute error being much smaller, at $0.09 \mu\text{K}^2$. Evaluating this 4th order polynomial is much faster than running CAMB itself, since the polynomial can be evaluated in only about a million floating point operations (210 times 2500 times 2, plus a few), which incidentally is fast enough to be almost real-time in a web browser window if implemented in JavaScript. Another useful benefit is that the quartic polynomial is differentiable to machine precision, unlike CAMB, which is only differentiable to the accuracy needed for MCMC parameter estimation, usually about 0.1%. This increased differentiability is very useful for several different optimization methods. We note that this fit was inspired by and follows the same philosophy as that of the PICO code⁹ (Fendt & Wandelt 2007b,a), and PICO has been shown to work for a much larger range of parameters by using multiple polynomial fits.

When computing with CAMB, we use standard values for parameters not in our six-parameter model. For example, we set the curvature, Ω_k , and tensor-to-scalar ratio, r , both to be zero, and we use `massless_neutrinos = 3.046` and `massive_neutrinos = 0`. We use the July 2013 version of CAMB, with the following also set in the parameter file: `l_max_scalar = 3000`, `k_eta_max_scalar = 7500`, `temp_cmb = 2.7255`, `helium_fraction = 0.24`, `high_accuracy_default = T`, `accuracy_boost = 1`, `l_accuracy_boost = 1`, and `l_sample_boost = 1`.

We run CAMB individually for each set of parameters, for simplicity, although further speed-ups could be obtained in the initial 10,000 runs of CAMB by computing C_l for several different values of the parameters affecting the primordial

$P(k)$, namely A_s and n_s , within a single run of CAMB. This can be done by only modifying the CAMB input parameter `.ini` file.

⁷ A CMB power spectrum fluctuation that precisely matches the shape of a foreground template is therefore not measurable.

⁸ With six variables, the 0th order part has 1 term; the 1st order part has 6 terms; the 2nd order part has 21 terms; 3rd has 56 terms; and 4th has 126 terms.

⁹ <https://sites.google.com/a/ucdavis.edu/pico/home>

REFERENCES

- Barnes, C., Hill, R. S., Hinshaw, G., Page, L., Bennett, C. L., et al. 2003, *ApJS*, 148, 51 [3.2, 3.2.2]
- Bennett, C. L., Bay, M., Halpern, M., Hinshaw, G., Jackson, C., et al. 2003, *ApJ*, 583, 1 [2.3.1]
- Bennett, C. L., Larson, D., Weiland, J. L., & Hinshaw, G. 2014, *ArXiv e-prints*, 1406.1718 [1, 4]
- Bennett, C. L., Larson, D., Weiland, J. L., Jarosik, N., Hinshaw, G., et al. 2013, *ApJS*, 208, 20, 1212.5225 [1, 2.1, 2.2, 2.3.1, 2.3.2, 4, 3.2, 4.1, 5]
- Blanco, M. A., Flrez, M., & Bermejo, M. 1997, *Journal of Molecular Structure: {THEOCHEM}*, 419, 19 [A]
- Chisholm, C. D. H. 1976, *Group Theoretical Techniques in Quantum Chemistry* (London: Academic Press) [A]
- Chon, G., Challinor, A., Prunet, S., Hivon, E., & Szapudi, I. 2004, *Mon. Not. Roy. Astron. Soc.*, 350, 914, astro-ph/0303414 [4.1]
- de Oliveira-Costa, A., Tegmark, M., Zaldarriaga, M., & Hamilton, A. 2004, *Phys. Rev.*, D69, 063516, astro-ph/0307282 [A]
- Eriksen, H. K., Banday, A. J., Górski, K. M., & Lilje, P. B. 2004, *ApJ*, 612, 633 [B]
- Fendt, W. A. & Wandelt, B. D. 2007a, *ArXiv e-prints*, 0712.0194 [3]
- . 2007b, *ApJ*, 654, 2, astro-ph/0606709 [3]
- Finkbeiner, D. P. 2003, *ApJS*, 146, 407, accepted (astro-ph/0301558) [2.2]
- Finkbeiner, D. P., Davis, M., & Schlegel, D. J. 1999, *ApJ*, 524, 867, astro-ph/9905128 [2.2]
- Gorski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., et al. 2005, *ApJ*, 622, 759, astro-ph/0409513 [5]
- Hazra, D. K. & Shafieloo, A. 2014a, *J. Cosmology Astropart. Phys.*, 1, 043, 1401.0595 [1]
- . 2014b, *Phys. Rev. D*, 89, 043004, 1308.2911 [1, 3.1, 5]
- Hill, R. S., Weiland, J. L., Odegard, N., Wollack, E., Hinshaw, G., et al. 2009, *ApJS*, 180, 246, 0803.0570 [3.2, 3.2.2]
- Hinshaw, G., Larson, D., Komatsu, E., Spergel, D. N., Bennett, C. L., et al. 2013, *ApJS*, 208, 19, 1212.5226 [1]
- Hinshaw, G., Nolta, M. R., Bennett, C. L., Bean, R., Doré, O., et al. 2007, *ApJS*, 170, 288, arXiv:astro-ph/0603451 [3.2, A]
- Hinshaw, G., Weiland, J. L., Hill, R. S., Odegard, N., Larson, D., et al. 2009, *ApJS*, 180, 225, 0803.0732 [5, 3.2.2, A]
- Jarosik, N., Barnes, C., Greason, M. R., Hill, R. S., Nolta, M. R., et al. 2007, *ApJS*, 170, 263, arXiv:astro-ph/0603452 [2.2, 5]
- Jarosik, N., Bennett, C. L., Dunkley, J., Gold, B., Greason, M. R., et al. 2011, *ApJS*, 192, 14, 1001.4744 [2.1]
- Kovács, A., Carron, J., & Szapudi, I. 2013, *MNRAS*, 436, 1422, 1307.1111 [1, 2.3.2]
- Lewis, A., Challinor, A., & Lasenby, A. 2000, *ApJ*, 538, 473 [5]
- Planck II, Aghanim, N., Armitage-Caplan, C., Arnaud, M., Ashdown, M., et al. 2014, *A&A*, 571, A2, 1303.5063 [2.2]
- Planck III, Aghanim, N., Armitage-Caplan, C., Arnaud, M., Ashdown, M., et al. 2014, *A&A*, 571, A3, 1303.5064 [2.1, 4, 2.2, 2.3.1, 5]
- Planck V, Aghanim, N., Armitage-Caplan, C., Arnaud, M., Ashdown, M., et al. 2014, *A&A*, 571, A5, 1303.5066 [2.2, 2.3.1, 3]
- Planck VIII, Ade, P. A. R., Aghanim, N., Armitage-Caplan, C., Arnaud, M., et al. 2014, *A&A*, 571, A8, 1303.5069 [2.1, 2.2, 5]
- Planck XII, Ade, P. A. R., Aghanim, N., Armitage-Caplan, C., Arnaud, M., et al. 2014, *A&A*, 571, A12, 1303.5072 [2.3.2, 2]
- Planck XIII, Ade, P. A. R., Aghanim, N., Alves, M. I. R., Armitage-Caplan, C., et al. 2014, *A&A*, 571, A13, 1303.5073 [2.2]
- Planck XIV, Ade, P. A. R., Aghanim, N., Armitage-Caplan, C., Arnaud, M., et al. 2014, *A&A*, 571, A14, 1303.5074 [2.3.1]
- Planck XV, Ade, P. A. R., Aghanim, N., Armitage-Caplan, C., Arnaud, M., et al. 2014, *A&A*, 571, A15, 1303.5075 [1, 4.1]
- Planck XVI, Ade, P. A. R., Aghanim, N., Armitage-Caplan, C., Arnaud, M., et al. 2014, *A&A*, 571, A16, 1303.5076 [1, 6, 4.3]
- Planck XXVII, Aghanim, N., Armitage-Caplan, C., Arnaud, M., Ashdown, M., et al. 2014, *A&A*, 571, A27, 1303.5087 [A]
- Planck XXXI, Ade, P. A. R., Arnaud, M., Ashdown, M., Aumont, J., et al. 2014, *A&A*, 571, A31 [1, 3, 3.1, 5]
- Spergel, D., Flauger, R., & Hlozek, R. 2013, *ArXiv e-prints*, 1312.3313 [1]
- Szapudi, I., Prunet, S., Pogosyan, D., Szalay, A. S., & Bond, J. R. 2001, *ApJ*, 548, L115 [4.1]
- Varshalovich, D. A., Moskalev, A. N., & Khersonskii, V. K. 1988, *Quantum Theory of Angular Momentum* (Singapore: World Scientific) [A]