


Consistency of cosmic microwave background temperature measurements in three frequency bands in the 2500-square-degree SPT-SZ survey

To cite this article: L.M. Mocanu *et al* JCAP07(2019)038

View the [article online](#) for updates and enhancements.

Recent citations

- [Analytic Calculation of Covariance between Cosmological Parameters from Correlated Data Sets, with an Application to SPTpol](#)
Joshua A. Kable *et al.*



IOP | ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

Consistency of cosmic microwave background temperature measurements in three frequency bands in the 2500-square-degree SPT-SZ survey

L.M. Mocanu,^{a,b,c} T.M. Crawford,^{a,b,1} K. Aylor,^d B.A. Benson,^{e,a,b}
L.E. Bleem,^{f,a} J.E. Carlstrom,^{a,f,g,b,h} C.L. Chang,^{f,a,b} H.-M. Cho,ⁱ
R. Chown,ⁱ A.T. Crites,^{a,b,k} T. de Haan,^{l,l} M.A. Dobbs,^{i,m}
W.B. Everett,^q E.M. George,^{l,o} N.W. Halverson,^{n,p}
N.L. Harrington,^l J.W. Henning,^{f,a} G.P. Holder,^{i,m,q,r}
W.L. Holzapfel,^l Z. Hou,^{a,b} J.D. Hrubes,^s L. Knox,^d A.T. Lee,^{l,t}
D. Luong-Van,^s D.P. Marrone,^u J.J. McMahon,^y S.S. Meyer,^{a,b,h,g}
M. Millea,^d J.J. Mohr,^{w,x,y} T. Natoli,^b Y. Omori,^{i,z,aa} S. Padin,^{a,b}
C. Pryke,^{ab} C.L. Reichardt,^{ac} J.E. Ruhl,^{ad} J.T. Sayre,^{ad,n}
K.K. Schaffer,^{a,h,ae} E. Shirokoff,^{la,b} Z. Staniszewski,^{ad,af}
A.A. Stark,^{ag} K.T. Story,^{z,aa} K. Vanderlinde,^{ah,ai} J.D. Vieira,^{q,r}
R. Williamson and W.L.K. Wu^a

^aKavli Institute for Cosmological Physics, University of Chicago,
Chicago, IL, 60637 U.S.A.

^bDepartment of Astronomy and Astrophysics, University of Chicago,
Chicago, IL, 60637 U.S.A.

^cInstitute of Theoretical Astrophysics, University of Oslo,
P.O. Box 1029 Blindern, Oslo, NO-0315 Norway

^dDepartment of Physics, University of California,
Davis, CA, 95616 U.S.A.

^eFermi National Accelerator Laboratory, MS209,
P.O. Box 500, Batavia, IL, 60510 U.S.A.

^fHigh Energy Physics Division, Argonne National Laboratory,
Argonne, IL, 60439 U.S.A.

¹Corresponding author.

- ^gDepartment of Physics, University of Chicago,
Chicago, IL, 60637 U.S.A.
- ^hEnrico Fermi Institute, University of Chicago,
Chicago, IL, 60637 U.S.A.
- ⁱSLAC National Accelerator Laboratory,
2575 Sand Hill Road, Menlo Park, CA, 94025 U.S.A
- ^jDepartment of Physics and McGill Space Institute, McGill University,
Montreal, Quebec, H3A 2T8 Canada
- ^kCalifornia Institute of Technology,
Pasadena, CA, 91125 U.S.A.
- ^lDepartment of Physics, University of California,
Berkeley, CA, 94720 U.S.A.
- ^mCanadian Institute for Advanced Research, CIFAR Program in Cosmology and Gravity,
Toronto, ON, M5G 1Z8 Canada
- ⁿCenter for Astrophysics and Space Astronomy,
Department of Astrophysical and Planetary Sciences, University of Colorado,
Boulder, CO, 80309 U.S.A.
- ^oEuropean Southern Observatory,
Karl-Schwarzschild-Straße 2, Garching, 85748 Germany
- ^pDepartment of Physics, University of Colorado,
Boulder, CO, 80309 U.S.A.
- ^qAstronomy Department, University of Illinois at Urbana-Champaign,
1002 W. Green Street, Urbana, IL, 61801 U.S.A.
- ^rDepartment of Physics, University of Illinois Urbana-Champaign,
1110 W. Green Street, Urbana, IL, 61801 U.S.A.
- ^sUniversity of Chicago, Chicago, IL, 60637 U.S.A.
- ^tPhysics Division, Lawrence Berkeley National Laboratory,
Berkeley, CA, 94720 U.S.A.
- ^uSteward Observatory, University of Arizona,
933 North Cherry Avenue, Tucson, AZ, 85721 U.S.A.
- ^vDepartment of Physics, University of Michigan,
Ann Arbor, MI, 48109 U.S.A.
- ^wFaculty of Physics, Ludwig-Maximilians-Universität,
München, 81679 Germany
- ^xExcellence Cluster Universe, Garching, 85748 Germany
- ^yMax-Planck-Institut für extraterrestrische Physik,
Garching, 85748 Germany
- ^zKavli Institute for Particle Astrophysics and Cosmology, Stanford University,
452 Lomita Mall, Stanford, CA, 94305 U.S.A.
- ^{aa}Department of Physics, Stanford University,
382 Via Pueblo Mall, Stanford, CA, 94305 U.S.A.
- ^{ab}Department of Physics, University of Minnesota,
Minneapolis, MN, 55455 U.S.A.
- ^{ac}School of Physics, University of Melbourne,
Parkville, VIC 3010 Australia

^{ad}Physics Department, Center for Education and Research in Cosmology and Astrophysics,
Case Western Reserve University, Cleveland, OH, 44106 U.S.A.

^{ae}Liberal Arts Department, School of the Art Institute of Chicago,
Chicago, IL, 60603 U.S.A.

^{af}Jet Propulsion Laboratory, California Institute of Technology,
Pasadena, CA, 91109 U.S.A.

^{ag}Harvard-Smithsonian Center for Astrophysics,
Cambridge, MA, 02138 U.S.A.

^{ah}Dunlap Institute for Astronomy & Astrophysics, University of Toronto,
50 St George St, Toronto, ON, M5S 3H4 Canada

^{ai}Department of Astronomy & Astrophysics, University of Toronto,
50 St George St, Toronto, ON, M5S 3H4 Canada

E-mail: tcrawfor@kicp.uchicago.edu

Received May 1, 2019

Accepted July 7, 2019

Published July 24, 2019

Abstract We present an internal consistency test of South Pole Telescope (SPT) measurements of the cosmic microwave background (CMB) temperature anisotropy using three-band data from the SPT-SZ survey. These measurements are made from observations of $\sim 2500 \text{ deg}^2$ of sky in three frequency bands centered at 95, 150, and 220 GHz. We combine the information from these three bands into six semi-independent estimates of the CMB power spectrum (three single-frequency power spectra and three cross-frequency spectra) over the multipole range $650 < \ell < 3000$. We subtract an estimate of foreground power from each power spectrum and evaluate the consistency among the resulting CMB-only spectra. We determine that the six foreground-cleaned power spectra are consistent with the null hypothesis, in which the six cleaned spectra contain only CMB power and noise. A fit of the data to this model results in a χ^2 value of 236.3 for 235 degrees of freedom, and the probability to exceed this χ^2 value is 46%.

Keywords: CMBR experiments, cosmological parameters from CMBR

ArXiv ePrint: [1904.12995](https://arxiv.org/abs/1904.12995)

Contents

1	Introduction	1
2	Instrument, observations, and data reduction	3
3	Power spectrum estimation	4
4	Bandpower covariance matrix	6
4.1	Estimating and conditioning the noise and signal terms	7
4.2	Foreground-removal, beam, and calibration contributions	8
4.3	Relative size of covariance matrix contributions	8
5	Null tests	9
6	Foreground treatment	10
7	Inter-spectrum consistency test	12
7.1	Formalism	12
7.2	Treatment of sample variance	12
8	Results	13
9	Conclusion	16
A	Single-spectrum squared deviation	17
A.1	Expected and measured single-frequency weighted, squared deviation	17
A.2	Unequal sharing of degrees of freedom	18

1 Introduction

Measurements of the temperature anisotropy in the cosmic microwave background (CMB) have played a key role in establishing our current understanding of the Universe. Since the first detection of anisotropy in data from the Cosmic Background Explorer Differential Microwave Radiometer experiment [1], the power of CMB temperature anisotropy measurements to constrain cosmology has progressed steadily, culminating in measurements from the *Planck* satellite [2]. The *Planck* measurements achieve percent-level constraints on five of the six parameters of the so-called Lambda Cold Dark Matter model.

At this level of statistical precision, certain mild tensions have arisen among and within CMB datasets, and between CMB data and other cosmological measurements. One of these is a mild tension between the *Planck* best-fit cosmological parameters and those estimated from a combination of data from the South Pole Telescope (SPT) and the *Wilkinson Microwave Anisotropy Probe* (WMAP) in [3, hereafter S13]. These tensions have been noted in several publications including [4]. The cosmological constraints in S13 are tighter than any CMB constraints in the literature other than the *Planck* constraints, and it is important to understand whether disagreements in the two datasets arise from slightly unusual statistical fluctuations, from untreated systematic effects in either dataset, or from a breakdown

in our cosmological model. References [5] and [6] compared data from SPT and *Planck* in map, power spectrum, and cosmological parameter space in an attempt to distinguish among these possibilities and found no evidence that the parameter differences are due to unmodeled systematic errors in either data set.

In this work, we specifically target the possibility of unmodeled systematics in the SPT data by comparing data from three frequency bands in the 2500-square-degree SPT-SZ survey. We refer to the three bands as 95 GHz, 150 GHz, and 220 GHz, indicating their rough center frequencies. The 150 GHz SPT-SZ data provided the bulk of the statistical weight to the S13 cosmological constraints, so this comparison is directly relevant to the investigation of cosmological parameter differences between S13 and *Planck*. Additive and multiplicative systematic errors typical of CMB datasets are unlikely to be 100% correlated between observing bands, so a comparison of CMB temperature measurements in different bands can be a powerful tool for identifying such errors. Typical additive systematics include incomplete or inaccurate foreground modeling and subtraction, inaccurate noise modeling, and ground pickup (for ground-based telescopes). The most important multiplicative systematics for CMB temperature measurements are inaccurate modeling of the telescope beam and absolute response. All of these are likely to vary with observing frequency, depending on how beams, noise, and calibration are characterized.

For the particular case of SPT-SZ data, the most difficult residual systematics to rule out as potential causes for tension with *Planck* are estimates of foregrounds and beams. Systematic errors in both of these quantities are expected to have very different amplitudes in the three observing bands, because of the spectral dependence of millimeter-wave emission mechanisms and the different beam sizes in the three SPT-SZ bands. Thus, we would expect significant systematics in either of these quantities to lead to statistical disagreement among the foreground-cleaned CMB temperature estimates in the three bands.

We estimate the consistency among CMB measurements in the three SPT-SZ bands by creating six angular power spectra, including the three single-frequency power spectra and three cross-frequency spectra (95×150 , 95×220 , and 150×220). We subtract an estimate of foregrounds from each spectrum, and we fit the six foreground-cleaned spectra to a simple model in which the CMB is the only source of signal remaining in the spectra. The χ^2 of this fit and the probability to exceed (PTE) that value of χ^2 are our metrics for internal consistency of this dataset. We note that a version of this test has already been implicitly performed in the cosmological analysis of [7, hereafter G15]. In that work, the same three-band data were used to estimate the same six power spectra, and the spectra were fit to a combination CMB and foreground model in the multipole range $2000 < \ell < 11000$. The CMB part of the model nominally came from a six-parameter Λ CDM model, but the fit was performed jointly with *Planck* and baryon acoustic oscillation data, which effectively fixed the cosmological parameters and hence the predicted CMB power spectrum. The result of that fit was a χ^2 of 88 for 80 degrees of freedom and a PTE of 21.5%. The test in this paper differs from G15 in multipole range and model assumptions. In this work we examine the multipole range used for cosmological constraints in S13, namely $650 < \ell < 3000$, and we use only data above $\ell = 3000$ to construct the foreground model. Assuming nothing about the underlying cosmological model, we test whether the six foreground-cleaned power spectra yield consistent estimates of the power in each ℓ bin.

This paper is organized as follows. We describe the SPT telescope and SPT-SZ camera, the SPT-SZ survey observations, and the data analysis up to the map level in section 2. We describe our power spectrum pipeline in section 3, our construction of the ℓ -space covariance

matrix in section 4, and a series of null tests we perform on the data in section 5. We describe our treatment of foregrounds in section 6. The method we use to test for inter-band consistency is presented in section 7. We present the individual power spectra and the results of the consistency test in section 8, and we conclude in section 9.

2 Instrument, observations, and data reduction

The SPT is a 10-meter telescope located at the Amundsen-Scott South Pole station in Antarctica. The first camera on the SPT, the SPT-SZ camera, was a 960-pixel transition-edge-sensor bolometer receiver configured to observe in three frequency bands centered at roughly 95, 150, and 220 GHz. For more information on the telescope and camera, see [8] and [9]. From 2008–2011, the SPT-SZ camera was used to conduct a survey covering a ~ 2500 deg² region of sky between declinations of -40° and -65° and right ascensions of 20 h and 7 h. The SPT-SZ survey footprint is shown in, e.g., figure 1 of S13 and figure 1 of [5]. The SPT-SZ survey area was observed in 19 contiguous sub-patches, or fields, ranging in area from ~ 70 to ~ 225 deg².

The observations and data reduction methods used in this work are very similar to those described in [10, hereafter K11] and S13, and we refer the reader to those works for detailed descriptions of the observations and analysis. In the following, we give a brief overview of the most important features of the observations and data reduction.

The SPT-SZ fields are observed using a raster-scan pattern, in which the telescope is scanned back and forth across the field at constant elevation. After each right/left scan pair, a small step in elevation is taken, and the process is repeated until the full field is covered. Though the extent in elevation and the size of the elevation step varies field to field, a typical single-field observation takes roughly two hours. Each field is observed an average of ~ 200 times. Two SPT-SZ fields were also observed with elevation scans at constant azimuth, but only the azimuth-scan data are included in this analysis. The SPT-SZ detector array was upgraded between the 2008 and 2009 seasons; in 2008 the 95 GHz detectors did not produce science-quality data. The two fields that were observed in 2008 were reobserved later to gain 95 GHz sensitivity on these fields, but for ease of building on the S13 analysis we only use 2008 data on these fields in this work.

The raw time-ordered data (TOD) from the SPT-SZ survey are converted to maps of the sky using a simple bin-and-average process. Maps are made individually from each single observation of a field, using data from all detectors of a given observing frequency. For a given map pixel, the TOD samples from any detector of a given frequency pointed in that direction in the sky are combined using inverse-noise-weighted averaging. Each detector is assigned a single weight value in an observation, based on a combination of noise in a certain range of temporal frequencies and the detector response to celestial sources. This mapmaking procedure returns the minimum-variance result in the limit of detector noise that is “white” (uncorrelated between time samples) and uncorrelated between detectors. To reduce correlated noise among time samples for a given detector and among detectors, and to reduce aliasing of high-frequency noise when we bin into pixels, we filter the TOD before mapmaking. The TOD from each detector are high-pass filtered by projecting out across a single scan a set of low-order polynomials and sines and cosines. The high-pass filter cutoff corresponds to roughly $\ell = 300$ in the scan direction. The anti-aliasing filter is a Fourier-domain low-pass filter with a cutoff at a temporal frequency corresponding to roughly $\ell = 6600$ in the scan direction. Finally, the average over all detectors in each of

the six detector modules is subtracted from every detector at every time sample. In the pseudo- C power spectrum pipeline described in the next section, we account for the effects of this filtering through the filter transfer function.

We estimate the filter transfer function using simulated observations. We create simulated skies with lensed CMB fluctuations (with an underlying power spectrum from the best-fit *Planck* 2015 TT,TE,EE+lowP+lensing cosmological model [2]) and foregrounds, and we mock-observe these skies using the real detector pointing, weights, and filtering. We create maps from these simulated data using the same procedure as for real data. We calculate the power spectrum of these mock maps and compare it to the known input power spectrum. We define the filter transfer function F_ℓ following eq. 18 in [11]. We create a separate filter transfer function for each observing band.

The SPT-SZ beams are measured with a combination of planet observations and sources in the SPT-SZ fields. The main lobes of the beams are measured using bright sources in the fields, the sidelobes are measured using observations of Jupiter, and these two measurements are connected using observations of Venus. Individual beam profiles are estimated for each frequency band and each survey year (because the receiver was modified between every observing season). We approximate the beams as azimuthally symmetric and create a one-dimensional ℓ -space beam function B_ℓ for each frequency band and year.

The relative calibration between detectors and between individual observations is accomplished using a combination of observations of the Galactic HII region RCW38 and the response of detectors to a blackbody source mounted behind the secondary mirror. An approximate calibration for each frequency band is obtained using the known brightness of RCW38 in these bands. The final absolute calibration for each band is obtained either by comparing the fully coadded maps to *Planck* or by comparing power spectra of those maps to *Planck* power spectra. At 150 GHz, we use the absolute calibration from [5], which was obtained using a cross-spectrum analysis between the SPT and *Planck* maps. At 95 and 220 GHz, we use an absolute calibration obtained by comparing SPT power spectra from G15 with full-sky *Planck* 2015 power spectra. The fractional uncertainty on the absolute calibration is 2.1%, 0.3%, and 4.5% in power in the 95, 150, and 220 GHz bands, respectively. We use these uncertainties to derive the calibration contribution to the covariance described in section 4.

We note that we could in principle achieve tighter priors on the 95 and 220 GHz calibrations by doing a map-based comparison to *Planck* similar to that done in [5] for 150 GHz. We can achieve higher precision, however, by transferring the 150 GHz calibration to 95 and 220 GHz using our own data. If we were to perform the fit in this paper with the 95 and 220 GHz calibrations as free parameters (instead of folding the calibration uncertainty into the bandpower covariance matrix), the posteriors on those parameters would be our best possible calibration in those bands. We have in fact performed such a fit, the posteriors on the best-fit 95 and 220 GHz calibrations relative to the priors are 1.016 ± 0.004 and 1.046 ± 0.009 (in power), respectively, and we will use these calibrations in future analyses of 95 and 220 GHz SPT-SZ data.

3 Power spectrum estimation

In this section, we describe the pipeline used to compute the angular power spectrum from the maps. This analysis follows the methods used in [12, 13], K11, S13, and G15, and we refer the reader to those works for more detail. We adopt the flat-sky approximation, in which angular

wavenumber k is equivalent to multipole number ℓ , and spherical harmonic transforms are replaced by Fourier transforms. The six power spectra are calculated independently in each field and then combined. We report power spectra in terms of D_ℓ , defined as

$$D_\ell = \frac{(\ell + 1)}{2\pi} C_\ell. \quad (3.1)$$

Throughout this work, we will often refer to the various single-frequency and cross-frequency spectra using the shorthand “95 × 95” for the 95 GHz spectrum, “95 × 150” for the 95 GHz–150 GHz spectrum, etc.

We use a cross-spectrum estimator similar to that described in [14]; specifically, we cross-correlate maps of different observations of the same field. These maps can be in the same observing band or different observing bands. We assume that the noise is uncorrelated between different observations, so the cross-spectra are free of noise bias. We multiply the map of observation A at frequency ν by a mask, zero-pad it to the same size for all fields, and calculate its Fourier transform $m^A(\nu_1)$, then we do the same for observation B at frequency ν_2 . The mask $\mathbf{W}$ rolls off the noisy edges of the field and has apodized holes at the locations of bright point sources (sources with flux density above 50 mJy at 150 GHz). We calculate the average cross spectrum between these two maps within an ℓ -bin b :

$$\mathbb{D}_b^{AB}(\nu_1, \nu_2) \equiv \frac{(\ell + 1)}{2\pi} H_\ell \text{Re}[m^A(\nu_1) m^{B*}(\nu_2)]_{\ell \in b}. \quad (3.2)$$

Here, ℓ is the two-dimensional ℓ -space vector, and H is a two-dimensional weight array that is described below. We average all cross-spectra $\mathbb{D}_b^{AB}$ for $A \neq B$ to calculate a binned power spectrum $\mathbb{D}_b$ for each field and frequency combination. We will often refer to binned power spectrum values D_b as *bandpowers*.

Noise in the SPT-SZ survey maps is statistically anisotropic. Modes of a given angular frequency ℓ that oscillate perpendicular to the scan direction ($\ell_y \neq 0$; $\ell_x = 0$) are noisier than modes that oscillate along the scan direction ($\ell_x \neq 0$; $\ell_y = 0$). To combine the power from different modes in each ℓ bin more optimally, we define a two-dimensional weight array $H_\ell(\nu_1, \nu_2)$ for each frequency combination:

$$H_\ell(\nu_1, \nu_2) \propto \frac{C_\ell^{\text{th}}}{C_\ell^{\text{th}} + N_\ell(\nu_1, \nu_2)^{i_{-2}}}, \quad (3.3)$$

where C_ℓ^{th} is the theoretical power spectrum used in simulations, and N_ℓ is the two-dimensional calibrated, beam-deconvolved noise power. The noise power is calculated for each field and frequency as the average two-dimensional cross-power spectrum of 100 noise realizations for each frequency in the cross-spectrum pair. Noise realizations are created under the assumption of stationarity (statistically identical noise in each individual observation) by splitting the individual observations into two halves, multiplying one half by -1 , and averaging all observations. Multiple semi-independent realizations are created by randomizing which individual observations go into each half. We smooth the weight array and normalize it to the maximum value in each annulus. In addition to this smooth weighting, which is identical to the treatment in S13, we also mask a set of modes in the spectra involving 95 GHz with ℓ values that correspond to frequencies of 1 Hz in the time-ordered data. If these modes are not masked, we find anomalously low PTE values in one of the null tests described in section 5.

The raw bandpowers $\mathbb{D}_b$ are a biased estimate of the true sky power, D_b . The biased and unbiased estimates are related by

$$\mathbb{D}_b \equiv K_{b\ell} D_b, \quad (3.4)$$

where the K matrix accounts for the effects of the binning, windowing, TOD filtering, pixelization, and beams. Following [11], K can be written as

$$K_{bb'} = P_{b'} M_{\ell\ell'} [\mathbf{W}] F_{\ell} B_{\ell}^2 Q_{\ell\ell'}, \quad (3.5)$$

where $Q_{\ell\ell'}$ is the binning operator and $P_{b'}$ its reciprocal, $M_{\ell\ell'}$ is the matrix describing the mode-coupling induced by the mask, and B and F are the beam and filter transfer functions described in section 2.

After performing the analysis described above, we are left with 19 sets of $6N_{\text{bins}}$ bandpowers, one per SPT-SZ survey field. The total ℓ range covered is $650 < \ell < 3000$ with $\Delta\ell = 50$, resulting in $N_{\text{bins}} = 47$ bins. The value of $\Delta\ell$ follows the choice in K11 and S13, chosen to minimize bin-to-bin correlation in a typical-size field while preserving resolution of peaks in the CMB power spectrum. We combine the bandpowers from each field using area-based weights. The combined bandpowers are given by

$$D_b(\nu_1, \nu_2) = \sum_i D_b^i(\nu_1, \nu_2) w^i \quad (3.6)$$

where

$$w^i = \frac{A^i}{\sum_j A^j}, \quad (3.7)$$

and A^i is the area of field i . For notation convenience, we will subsequently refer to the full $6N_{\text{bins}}$ unbiased bandpowers as D_α , where α runs over the N_{bins} bins as the fast index and the six frequency combinations as the slow index.

Because we have used nearly identical analysis methods and data for the 150×150 spectrum as used in S13, we can compare the S13 bandpowers to the 150×150 bandpowers in this work as a check. Once we account for the slightly different calibration in the two analyses (the two calibrations differ by roughly 0.9% in temperature), the fractional difference in bandpowers is less than 1% in most ℓ bins and less than 2% everywhere. This residual difference is attributable to small differences in cuts, apodization, and ℓ -space weighting between the two analyses¹.

4 Bandpower covariance matrix

The bandpower covariance matrix quantifies the uncertainty on the $6N_{\text{bins}}$ bandpowers, as well as the correlations among uncertainties in different ℓ bins and frequency combinations. It includes terms describing the contributions from instrument and atmospheric noise, sample variance, and uncertainty in our knowledge of the instrument beam and absolute calibration. The bandpower covariance matrix $C_{\alpha\beta}$ is a set of $36 N_{\text{bins}}$ -by- N_{bins} blocks, with the six on-diagonal blocks corresponding to the bin-to-bin covariances within the six power spectra (95×95 , 95×150 , etc.), and the off-diagonal blocks corresponding to the covariances between one frequency combination and another. The indices of $C_{\alpha\beta}$ run over ℓ bins and frequency combinations, as defined in section 3. We describe how we estimate each of the contributions to the bandpower covariance matrix, and our method for numerically conditioning the matrix, below.

¹The 150×150 bandpowers used here are derived from nearly identical data and cut and weight settings as the bandpowers used in [5], but that work used the cross-spectrum of a single pair of half-depth maps rather than the mean cross-spectrum of all pairs of single-observation maps. The bandpowers used in [5] were compared to the S13 bandpowers, and no evidence of inconsistency was found.

4.1 Estimating and conditioning the noise and signal terms

In S13 and G15 and earlier SPT analyses, we split the estimation of the signal contribution to the bandpower covariance matrix from the estimation of the noise and noise $\times$ signal parts. The signal term, or sample variance, is calculated from the same (noise-free) simulations used to estimate the filter transfer function, while the noise and noise $\times$ signal terms are estimated from the data, namely from the distribution of cross-spectrum bandpowers over all the map pairs used in eq. (3.2). As discussed in detail in section 7.2, we do not include sample variance in estimating the bandpower covariance matrix for the fitting procedure described in that section, though we do include the noise $\times$ signal contribution. To accomplish this, we simply omit from the covariance the term estimated from noise-free simulations.

In the limit of full sky coverage, no gravitational lensing, no masking, and stationary Gaussian noise, the measurement of the power spectrum at every value of multipole ℓ is uncorrelated from the measurement at every other value of ℓ , and the covariance matrix of our binned, unbiased estimate $\hat{D}$ would be diagonal for a single frequency combination. Even in this limit, however, the noise between different combinations that share a frequency would be correlated, so we must include off-diagonal blocks in the bandpower covariance matrix. Additionally, the finite sky coverage, point-source mask, and non-Gaussian foregrounds induce correlations between bins in a single frequency combination, so we construct the full $6N_{\text{bins}}\text{-by-}6N_{\text{bins}}$ covariance matrix.

The initial estimate of the bandpower covariance matrix is noisy for the off-diagonal elements, both within a frequency-combination block and in the off-diagonal blocks. To condition this noisy estimate, we compute the correlation matrices of the six on-diagonal blocks and note that their shape is determined primarily by the mode-coupling matrix and depends only on the distance from the diagonal. We condition the on-diagonal blocks by averaging the off-diagonal elements at a fixed separation from the diagonal and set elements that are a distance $\ell > 250$ from the diagonal to zero. This method of conditioning does miss some real correlation structure in the sample covariance. At very high ℓ in some frequency combinations, the covariance becomes dominated by the brightest point sources and is hence strongly correlated across all ℓ bins. We account for this when we condition the sample covariance by fitting for and removing a fully correlated part of the covariance at high ℓ and replacing it after the conditioning described above. There are also potential off-diagonal terms in the sample covariance from the effects of lensing; we expect these to be small for the temperature power spectrum and do not attempt to recover them in the conditioning. We note that the sample-variance portion of the covariance matrix is omitted in the fitting procedure described in section 7.1, so the exact method of conditioning that part of the matrix does not affect the main result of this work.

For the off-diagonal blocks, the uncertainty on the computed covariance can be large compared to the true covariance. Therefore, we calculate the elements of the off-diagonal blocks by applying the average correlation matrix computed from the corresponding on-diagonal blocks to them.

We calculate and condition the sample and noise covariances on a by-field basis and then combine the results using the same weighting used to combine bandpowers. The combined covariance matrix is given by

$$\mathbf{C}_{\alpha\beta} = \sum_i \mathbf{C}_{\alpha\beta}^i (w^i)^2, \quad (4.1)$$

where w^i are the field weights defined in section 3.

4.2 Foreground-removal, beam, and calibration contributions

We add a contribution to the covariance matrix to account for the uncertainty in the foreground components which were subtracted to generate our CMB-only band powers. See section 6 for details on the foreground removal process and the covariance estimate.

As in S13 and G15 and earlier SPT papers, we also add components to the covariance matrix to account for beam and calibration uncertainty. We follow the treatment in [6] and create the beam and calibration covariance matrices using a theoretical prediction for the measured bandpowers rather than the measured bandpowers themselves. For the theoretical prediction, we use CMB only, because the additions to the covariance matrix from uncertainty in subtracted foreground power include beam and calibration contributions. We follow the beam uncertainty treatment described in detail in K11. We note that the calibration uncertainty at 150 GHz is uncorrelated from the calibration uncertainty in the other two bands because of the different method of estimation (comparison of SPT maps to *Planck* maps in the SPT-SZ observing region for 150 GHz, comparison of SPT power spectra to full-sky *Planck* power spectra for 95 and 220 GHz, for details see section 2).

4.3 Relative size of covariance matrix contributions

The (square root of the) diagonal parts of various contributions to the field-combined covariance matrices are shown in figure 1 for the six single-frequency and cross-frequency spectra. The “foreground” curves are the sum of the contribution from uncertainties in the G15 foreground estimates and uncertainties on the estimates of power from point sources masked in G15 and not in this work (see section 6 for details). While the foreground component only applies to the foreground-cleaned, CMB-only bandpowers, and we do not include sample variance in the consistency test described in section 7.1, we include both the sample variance and the foreground components in the figure for informational purposes.

It is also worth mentioning some of the correlation structure of the various parts of the covariance matrix. In the regions of ℓ space in which the CMB is the largest source of signal ($\ell > 2500$ or 3000 , depending on frequency combination — see figure 3), the sample variance is highly correlated between frequency combinations but mostly uncorrelated between ℓ bins. At very high ℓ in some frequency combinations (particularly 95×95), the sample variance becomes dominated by the brightest point sources and is thus highly correlated between ℓ bins and, to a degree, among frequency combinations. The noise variance is mostly uncorrelated between bins and frequency combinations. The calibration covariance contributions are 100% correlated between bins and partially correlated between frequency combinations (particularly among 95×95 , 95×220 , and 220×220 because of the method used to calibrate the 95 GHz and 220 GHz data, described in section 2). The beam contributions are strongly correlated between bins, strongly correlated between frequency combinations that share a band, and partially correlated between frequency combinations that do not share a band.

As representative examples, slices through the covariance matrix along the rows corresponding to the ℓ -space bin centered at $\ell = 1525$ in the 95×150 and 150×150 spectra (C_{64} and $C_{\alpha 158}$, respectively, in the numbering scheme introduced above) are plotted in figure 2. We choose these two spectra because they have the lowest total covariance and thus the highest weight in the fit described in section 7. We show slices through the total covariance and through the sample variance, noise, calibration, beam, and foreground components individually. All slices are normalized by the square root of the product of the diagonal of the total covariance and the $C_{64\ 64}$ or $C_{158\ 158}$ element.




Figure 1. Square root of the diagonal elements of the bandpower covariance matrix for each frequency combination, broken down into contributions from five different sources. We note that the sample variance contributions are not included in the fitting procedure described in section 7. The covariance matrices are binned at the same $\Delta_\ell = 50$ resolution as the data, so the plotted diagonal values also have that resolution.

5 Null tests

While the purpose of this analysis is to test for unmodeled systematics in the final reported SPT-SZ bandpowers by comparing the different single-frequency and cross-frequency spectra, we can also check the individual spectra for contamination by splitting the data along some axis that is expected to maximize the sensitivity to a known source of contamination,




Figure 2. Slices through the covariance matrix along the rows corresponding to the $\ell = 1525$ bin in 95×150 (left panel) and 150×150 (right panel). The x-axis corresponds to bin index α (cf. section 3), which runs over ℓ -space bins from $\ell = 650$ to $\ell = 3000$ six times, once each for 95×95 , 95×150 , etc. (so bin 64 is $\ell = 1525$ in 95×150 , and bin 158 is $\ell = 1525$ in 150×150). All slices are normalized by the square root of the product of the diagonal of the total covariance and the $C_{64\ 64}$ or $C_{158\ 158}$ element. We note that the sample variance contributions are not included in the fitting procedure described in section 7.1.

calculating the power spectrum of the differenced data, and comparing to the expected value of such a difference from simulations. We follow S13, G15, and previous SPT analyses in performing these null tests. Briefly, in each null test, we divide a set of maps into two halves and subtract one set from the other, removing most of the signal (some residual signal is expected due to slight differences in weights, filtering, and calibration in individual observations). We then repeat the cross-spectrum analysis described in section 3 but with the subtracted “null maps”, and we compare the result to the expectation value of residual signal, using the variance over many cross-spectra to test whether any difference from the expectation value is consistent with noise.

The data splits are chosen to maximize the sensitivity to certain families of systematic uncertainties typical of ground-based CMB data. In this work, we use the “Time”, “Scan Direction”, “Azimuthal Range”, and “Sun” tests defined in S13. We perform them on the 95, 150, and 220 GHz data individually, for a total of 12 tests. The PTEs from these tests are shown in table 1. They range from 2% to 78%, three of the PTE values are below 10%, and ten of the PTE values are below 50%. In the absence of significant contamination, we would expect a minimum PTE value among the 12 tests of 2% or less roughly 20% of the time, three values below 10% roughly 10% of the time, and a distribution this asymmetric roughly 2% of the time. This takes into account the correlation among the null test values expected from the fact that they are not perfectly orthogonal splits of the data. (The largest expected correlation is 15–20% between the Time and Sun splits.) We make the subjective decision that 1-2 σ fluctuations in statistics defined a posteriori are not strong evidence of contamination, and we proceed to subtract an estimate of foregrounds from the data and perform the consistency test.

6 Foreground treatment

Our bandpowers include contributions from primary CMB temperature anisotropies and several foreground components. To obtain primary CMB-only bandpowers, we subtract

Band	Left-right	1st-half-second-half	sun	azimuth
95 GHz	0.08	0.02	0.19	0.07
150 GHz	0.44	0.51	0.78	0.1
220 GHz	0.49	0.13	0.22	0.14

Table 1. Probability to exceed (PTE) the χ^2 value obtained in each null test in each band. The individual tests are not independent; tests of the overall distribution of PTEs are discussed in the text.

estimates for foreground bandpowers from our data. We also add a contribution to the bandpower covariance matrix (section 4) to account for the uncertainty in the subtracted foregrounds.

We adopt the baseline model from G15 to estimate foreground contributions to these bandpowers. The baseline model includes contributions from the thermal and kinematic Sunyaev-Zel'dovich effects (tSZ and kSZ), spatially clustered and unclustered contributions from the cosmic infrared background (CIB), and a contribution from spatially unclustered radio galaxies (i.e., galaxies emitting synchrotron radiation from an active galactic nucleus). In the baseline model, eight parameters are derived for those foregrounds: the amplitudes of tSZ, kSZ power, unclustered CIB power, and clustered CIB power; two parameters describing the frequency dependence of the CIB terms; the tSZ-CIB correlation; and the spectral index of radio galaxies. The amplitude of radio galaxy power is fixed to a prior value based on the model in [15]. The analysis in G15 used the same raw data as this analysis; we refit the foreground model cutting G15 data at $\nu < 3000$, so that the foreground constraints are independent of the data in this work.

The analysis in G15 used a point source mask with a threshold of 5σ at 150 GHz, corresponding to a flux density of roughly 6.4 mJy at our survey depth. This analysis uses a masking threshold of 50 mJy. Consequently, the best-fit foreground model from our reanalysis of G15 data does not include the contribution from point sources between these two thresholds (roughly 3000 sources total, or roughly one per square degree and 150 per individual field), and we need to estimate and remove the power from these sources separately.

To estimate the power from sources between the two masking thresholds, we make use of the fact that we have full posterior deboosted flux density distributions for every source masked in G15 but not masked in this work. These distributions are calculated as part of the source count analysis in [16, 17] and W. Everett et al. (in preparation). We take 50,000 mock source catalogs drawn from these posterior flux distributions and calculate the contribution to the cross-spectrum between observing frequencies ν_1 and ν_2 from each of those mocks as:

$$C_{i,m}^{\text{PS } \nu_1 \times \nu_2} = \frac{1}{A_i} \frac{dB_{\nu_1}}{dT} \frac{dB_{\nu_2}}{dT}^{-1} \sum_j S_j^{\nu_1} S_j^{\nu_2}, \quad (6.1)$$

where i is the field index (19 fields total), A is the area of the field in steradians, m is the mock catalog index, and S_j^ν is the flux of source j (roughly 150 sources per individual field) from the list of selected sources in mock catalog m at observing frequency ν . We then estimate the power from these sources $C_{i,m}^{\text{PS } \nu_1 \times \nu_2}$ in each field as the average of the power from all mocks and derive error bars on this quantity from the standard deviation of power among the mocks. We subtract the resulting point source power in every field.

This approach constitutes our best possible estimate of the power from the sources we intend to mask, but it does not account for power from sources unintentionally masked. We

do not expect a significant amount of power from unintentionally masked sources, because sources at millimeter wavelengths between flux densities of 6 and 50 mJy are nearly all radio galaxies, and the clustering of these sources at these flux density levels has been measured to be very small (e.g., G15). Nevertheless, as a cross-check we take a second approach to estimating the difference in source power between G15 and this work. In this approach, we extend the multipole range in this analysis to $\ell_{\text{max}} = 4000$, we directly measure the difference in power between G15 and this work in the range $3000 < \ell < 4000$, we fit the result to a constant C_ℓ spectrum, and we subtract this value from the combined bandpowers in each single-frequency and cross-frequency combination. As in the first approach, we also add a contribution to the covariance matrix from the uncertainty in this fit. The final results using this approach are statistically indistinguishable from those obtained using our primary method.

7 Inter-spectrum consistency test

7.1 Formalism

Our primary goal in this paper is to test the consistency among the six sets of foreground-cleaned bandpowers. We wish to perform this test independently of any assumptions about a cosmological model, so we choose to find the one number for each of the 47 $\Delta\ell = 50$ bins which best fits the data, and we then evaluate the goodness of fit for the six bandpower sets to this computed best-fit set. We apply a generalized least-squares fit to the six sets of bandpowers, obtaining a single set of best-fit bandpowers.

We create a design matrix A_{ab} consisting of $N_{\text{bins}} = 47$ vectors, each $6N_{\text{bins}}$ elements long. Each vector in the design matrix is equal to 1 in the six elements corresponding to the measurement of the power in a single ℓ -space bin in the six spectra and 0 otherwise. Our best estimate of the true CMB variance on our patch of the sky in one bin is then

$$\bar{D}_b = A_{b\alpha}^T W_{\alpha\beta} A_{\beta b}^{-1} A_{b\gamma}^T W_{\gamma\delta} D_\delta, \quad (7.1)$$

where (as defined in section 3) Greek indices run over the full $6N_{\text{bins}}$ bandpowers while Roman indices run over the N best-fit bandpowers, $W_{\alpha\beta} = C_{\alpha\beta}^{-1}$ is the inverse bandpower covariance matrix (see section 4 for details), and summation over repeated indices is assumed. Assuming Gaussian-distributed likelihoods and a covariance matrix that does not depend on the model, the probability of obtaining our data given the model (which is simply that all six power spectra consist of a 100% common signal and noise described by our covariance matrix) is $L = e^{-\chi^2/2}$, where

$$\chi^2 = (D_\alpha - M_\alpha)^T W_{\alpha\beta} (D_\beta - M_\beta), \quad (7.2)$$

and $M_\alpha = A_{ab} \bar{D}_b$. The value of χ^2 and the associated PTE are the primary results of this paper.

7.2 Treatment of sample variance

When comparing measured bandpowers to a cosmological model, we are interested in the true, underlying variance of the CMB anisotropy. This true variance is impossible to measure perfectly with a finite number of samples, so the bandpower covariance matrix used in fitting to cosmology is the sum of contributions from measurement uncertainty and the variance

in the signal itself (and the cross-terms). When comparing two different measurements of variance, however, the contribution of signal or sample variance to the covariance matrix of the comparison depends on the fraction of sky modes the two measurements have in common. In the limit that the same modes on the sky are used with the same weighting to estimate bandpowers, and the CMB is the only signal on the sky (or the two measurements are at the same observing frequency), the sample variance disappears entirely. Put another way, the fitting procedure described in the previous section is formally insensitive to any component of the covariance matrix that is perfectly correlated between the six single-frequency and cross-frequency spectra.

We expect some residual uncorrelated sample variance among the six power spectra measured in this work for two reasons: 1) at high ℓ , the sky signal begins to be dominated by foregrounds, not CMB, and the intensities of the various foreground sources will be different in different bands; 2) the weighting of modes in an ℓ -space bin (eq. (3.3)) is slightly different for each spectrum. The relative contribution of this residual uncorrelated sample variance to the final covariance matrix should be small, though, for several reasons. One is that the CMB dominates all contributions to foreground power in G15 even at the highest multipoles measured in this work ($\ell = 3000$) for all spectra except 220×220 , and that spectrum is noise- and beam-error-dominated at all multipoles (see figure 1). The extra unclustered point-source power in the data used in this work compared to G15 (see section 6 for details) does overtake the CMB power at $\ell < 3000$ in 95×95 (see figure 3), but the procedure used to remove this power (which derives from knowledge of the actual flux densities of individual sources contributing to this power) also removes the sample variance. Furthermore, the sample variance on this signal is highly correlated between bins, because it is dominated by the brightness of the brightest few sources in the map, so the process of removing this power by measuring it at $\ell > 3000$ will also remove the bulk of the sample variance. Finally, the difference in mode weighting is empirically found to be very small across the six power spectra.

All of these effects can be properly taken into account in constructing the sample covariance matrix, but we have empirically found that the large correlated component causes numerical instability in the inversion of the full matrix, even after the conditioning described in section 4.1. For this reason, and based on the arguments that the residual uncorrelated sample variance should be small, we choose to ignore sample variance entirely in the fit in section 7.1. This will result in a small underestimate of the true variance and make the consistency test marginally more difficult to pass.

8 Results

Figure 3 shows the six sets of power spectra before foreground cleaning. Error bars are derived from the diagonal of the noise part of the bandpower covariance matrix (no sample-variance, beam, calibration, or foreground components are included). To make the error bars more visible, the 3σ range is shown (rather than the typical 1σ). Overplotted are the best-fit CMB-only bandpowers D_b from equation (7.1), the best-fit foreground model based on G15, the best-fit extra unclustered source power calculated using equation (6.1) (see section 6 for details), and the sum of these three components. The bottom panel of each plot in figure 3 shows the residual of the data after subtracting these three components, divided by the square root of the diagonal elements of the noise covariance matrix.




Figure 3. The six sets of single-frequency and cross-frequency bandpowers before foreground cleaning. In each plot the top panel shows a single power spectrum; error bars derived from the diagonal of the noise part of the bandpower covariance matrix (multiplied by a factor of three to make them more visible); the best-fit CMB-only bandpowers, foreground model from G15, and extra unclustered source power calculated using equation (6.1); and the sum of these three components. The bottom panel of each plot shows the residual of the data after subtracting these three components, divided by the square root of the diagonal elements of the noise covariance matrix.

In the limit that noise is the only contribution to the bandpower covariance, we expect the plotted noise-variance-scaled residuals to be centered at zero with unit scatter and little correlation between points. Many of the non-noise sources of variance in the data are strongly correlated between bins, however, and contributions from these sources of variance will be visible as long-wavelength structure in the residuals, especially in the lowest-noise frequency combinations. A calibration factor has been applied to the data (except 150×150) to remove the overall calibration-related residual; these factors are well within the calibration uncertainty used to create the calibration contribution to the covariance matrix. By eye, there are no obvious discrepancies between the data and the sum of the three components plotted. There is some discernible long-wavelength structure in some of the residuals (for example a negative drift in 150×150). If these long-wavelength residuals are consistent with, for example, our measured beam uncertainty or the uncertainty in the foreground model that has been subtracted, then they will be properly down-weighted in the χ^2 calculation (equation 7.2) and will not result in an elevated χ^2 value; if they arise from some other, unmodeled source, they will potentially result in an elevated χ^2 value and low PTE.

The six foreground-cleaned power spectra, and the best-fit bandpowers $\bar{D}_b$, are shown in figure 4. The predicted temperature power spectrum from the best-fit *Planck* 2015 TT+lowTEB+lensing cosmological model is overplotted as a guide, but we emphasize that neither this nor any other cosmological model was used in the fitting procedure. As in figure 3, a calibration factor has been applied to the data for all spectra except 150×150 . Again, the individual spectra look consistent by eye with the best-fit combined spectrum.

To make this statement more quantitatively, we compute the χ^2 and associated PTE for the model that the six foreground-cleaned power spectra consist of a single common sky signal (the best-fit combined spectrum) and noise described by the bandpower covariance matrix. The total number of points in the six spectra is 282 (47 ℓ -space bins per spectrum), and there are 47 free parameters in the fit (one value per bin in the combined spectrum), so we have 235 remaining degrees of freedom. The χ^2 , reduced χ^2 , and PTE for the fit are:

$$\begin{aligned}\chi^2 &= 236.3 \\ \chi^2/\text{dof} &= 1.006 \\ \text{PTE} &= 0.464.\end{aligned}\tag{8.1}$$

Thus we find that the six power spectra are consistent with our simple model, and we find no evidence of unmodeled systematics in the three-band SPT-SZ dataset.

We emphasize that, given the signal-to-noise on the power spectra in all six frequency combinations, and given the removal of the “protection” of the large sample variance contribution to the total covariance, this is a very stringent test of the consistency among the three SPT frequencies. The raw signal-to-noise in the six cross-spectra range from 10 to over 150 per bin, with a quadrature-summed signal-to-noise of over 1000. To give a sense of the level of unmodeled systematics the test in this work is sensitive to, we have determined empirically that, for example, an unmodeled foreground contribution to the 150 GHz data at the level of a few μK^2 at $\ell \sim 3000$ would cause an unacceptable PTE, as would an unmodeled instrumental systematic in one of the bands at the tens of μK^2 level at $\ell \sim 1000$.

We have also calculated the analogue to χ^2 and PTE for each of the six individual frequency combinations against the best-fit model. The interpretation of these quantities is not entirely straightforward, and we discuss these values and their interpretation in detail in appendix A. The general conclusion from the appendix is that the distribution of χ^2 values among the six individual spectra is also consistent with expectations.




Figure 4. The six individual foreground-subtracted single-frequency and cross-frequency spectra with the best-fit combined spectrum and the *Planck* 2015 TT+lowTEB+lensing best-fit model overplotted.

9 Conclusion

We have conducted a consistency test of three-band data from the 2500-square-degree SPT-SZ survey. Using 95, 150, and 220 GHz maps — effectively the same data used in previous SPT-SZ power spectrum analyses (including S13 and G15)—we have computed the six single-frequency and cross-frequency power-spectra among these three bands over the multipole range $650 < l < 3000$. We have subtracted a model of foreground power from the spectra. The foreground model is based on the best-fit foreground model in G15 with power added to account for sources that were masked in G15 but not in this analysis. Using a bandpower covariance matrix with contributions from noise, uncertainties in the subtracted foreground model, and beam and calibration uncertainties, we conducted a linear least-squares fit of the six spectra to a model in which each spectrum consists solely of a common sky signal (assumed to be CMB anisotropy) and noise described by the covariance matrix. The reduced χ^2 of that fit is 1.006 for 235 degrees of freedom, for a PTE of 0.464. We conclude from this result that there is no evidence of unmodeled frequency-dependent systematic error in the three-band SPT-SZ data over this multipole range. Together with the null results from the comparison of the 150 GHz SPT-SZ data set to *Planck* data in two recent publications [5, 6], this result bolsters the conclusion that unmodeled systematics in SPT-SZ data are unlikely to be the cause of any differences in cosmological parameters derived from SPT and *Planck*.

Acknowledgments

The South Pole Telescope program is supported by the National Science Foundation through grant PLR-1248097. Partial support is also provided by the NSF Physics Frontier Cen-

ter grant PHY-0114422 to the Kavli Institute of Cosmological Physics at the University of Chicago, the Kavli Foundation, and the Gordon and Betty Moore Foundation through grant GBMF#947 to the University of Chicago. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, and the resources of the University of Chicago Computing Cooperative (UC3), supported in part by the Open Science Grid, NSF grant NSF PHY 1148698. CR acknowledges support from an Australian Research Council's Future Fellowship (FT150100074). Some of the results in this paper have been derived using the HEALPix [18] package. We acknowledge the use of the Legacy Archive for Microwave Background Data Analysis (LAMBDA). Support for LAMBDA is provided by the NASA Office of Space Science.

A Single-spectrum squared deviation

In this appendix, we discuss the analogue to χ^2 and PTE for the six single-frequency or cross-frequency power spectra compared to the best-fit model. We discuss how we create expectation values for these χ^2 analogues from simulations, we report these expectation values and the values from the data, and we calculate and discuss the comparison PTEs.

A.1 Expected and measured single-frequency weighted, squared deviation

As discussed in section 8, we fit a model with 47 free parameters to a data set with 282 semi-independent data points, and we expect the sum of the weighted, squared residuals of the fit to follow a χ^2 distribution with $282 - 47 = 235$ degrees of freedom. We can also calculate the weighted, squared residuals for any subset of the data. It is of particular interest to calculate the residuals for each of the six sets of 47 bandpowers from the individual single-frequency or cross-frequency power spectra (95×95 , 95×150 , etc.). The χ^2 value for the entire data set is (cf. equation 7.2)

$$\chi^2 = (D_\alpha - M_\alpha)^T W_{\alpha\beta} (D_\beta - M_\beta). \quad (\text{A.1})$$

If we want to preserve all of the information about correlations among the individual spectra, we can create the residual vector $r = D_\alpha - M_\alpha$ and its weighted counterpart $\mathcal{Y} = W_{\alpha\beta} (D_\beta - M_\beta)$. The full χ^2 is the sum of the product of r and $\mathcal{Y}$, but we can also extract sums of subsets of that product as the analogue to χ^2 for the individual spectra. One might naively expect each set of these weighted, squared residuals to follow a χ^2 distribution with $235/6 \approx 39$ degrees of freedom. This is, however, only the case in the scenario in which the data points in the six sets of spectra are statistically independent, and each individual power spectrum has equal weights in the fit. The second point is elaborated upon in the next section.

We calculate expectation values for the weighted, squared deviation for each individual spectrum from simulations, and we compare the data to those expectations. The “simulations” we use for this purpose are simply the sum of mock CMB bandpowers and Gaussian realizations of the 282×282 -element covariance matrix $C_{\alpha\beta}$ used in the full fit. For each realization, we create a 47-element mock CMB bandpower vector (using a Gaussian realization of a theory power spectrum, binned using the same binning we use in the data) and add six copies of it to the 282-element realization of the covariance matrix. We perform the fit described in section 7 on each of these sets of mock bandpowers, and we calculate the total

spectrum	expected WSD	actual WSD	PTE
95 × 95	39.58	27.50	0.83
95 × 150	31.50	25.14	0.70
95 × 220	44.97	54.47	0.19
150 × 150	26.82	41.87	0.03
150 × 220	44.45	58.61	0.12
220 × 220	46.76	28.73	0.94

Table 2. Expected and actual total weighted, squared deviation values for the six individual spectra, and the fraction of simulations with higher values than the data for that spectrum.

squared deviation for each of the six (47-element) individual spectra:

$$\text{WSD}_i = \frac{\sum_{\alpha=47i}^{47i+46} r_{\alpha}^w}{47} \quad (\text{A.2})$$

We then perform the same calculation on the real data and compare the individual squared deviation values to the expected values for that spectrum. Table 2 shows for each individual spectrum the mean value of squared deviation from simulations, the value in the real data, and the fraction of simulations with higher squared deviation than the real data.

Because we already know from the main result of the paper that the ensemble data is consistent with expectations, we do not need to examine the overall distribution of these squared deviation and PTE values but need only look for individual outliers. The 150 × 150 spectrum has the PTE farthest from 50%, with only 3% of simulations having a higher squared deviation than the data. Using a simple trials-factor correction (from looking at six individual PTEs, also known as the Bonferroni correction, e.g., [19]), we would only flag this PTE value as significant if our threshold for single-PTE values was 18%. Thus we conclude the individual-spectrum deviations are also consistent with expectations.

A.2 Unequal sharing of degrees of freedom

The somewhat counter-intuitive result that the individual-spectrum weighted, squared deviations would not all be expected to be equal is easiest to understand in the simplified case of “fitting” N uncorrelated variables to the model in which they have the same underlying value — i.e., taking the weighted average of N numbers. For uncorrelated variables y_i with uncertainty σ_i , the inverse-noise-weighted average is

$$\bar{y} = \frac{\sum_i y_i / \sigma_i^2}{\sum_i 1 / \sigma_i^2} \quad (\text{A.3})$$

$$\equiv \frac{1}{W_{\text{tot}}} \sum_i \frac{y_i}{\sigma_i^2} \quad (\text{A.4})$$

Assuming the uncertainties are Gaussian-distributed, i.e., $P(y) = N(y_{\text{true}}, \sigma_i^2)$, where $N(\mu, \sigma)$ is the normal distribution with mean μ and variance σ^2 , the probability distribution of $\bar{y}$, as

expected, is

$$\begin{aligned}
P(\bar{y}) &= P \left(\frac{1}{W_{\text{tot}}} \sum_i y_i \right) \\
&= P \left(\frac{1}{W_{\text{tot}}} \frac{y_0}{\sigma_0^2} \right) \sim P \left(\frac{1}{W_{\text{tot}}} \frac{y_1}{\sigma_1^2} \right) \sim \dots \sim P \left(\frac{1}{W_{\text{tot}}} \frac{y_N}{\sigma_N^2} \right) \\
&= N \left(\frac{y_{\text{true}}}{W_{\text{tot}} \sigma_0^2}, \frac{1}{W_{\text{tot}} \sigma_0^2} \right) \sim N \left(\frac{y_{\text{true}}}{W_{\text{tot}} \sigma_1^2}, \frac{1}{W_{\text{tot}} \sigma_1^2} \right) \sim \dots \sim N \left(\frac{y_{\text{true}}}{W_{\text{tot}} \sigma_N^2}, \frac{1}{W_{\text{tot}} \sigma_N^2} \right) \\
&= N \left(\frac{y_{\text{true}}}{W_{\text{tot}} \sigma_f^2}, \frac{1}{W_{\text{tot}}^2 \sigma_f^2} \right) \\
&= N \left(y_{\text{true}}, \frac{1}{W_{\text{tot}}} \right),
\end{aligned} \tag{A.5}$$

where $\sim$ indicates convolution. The probability distribution of the residual between any one y_i and the best-fit $\bar{y}$ is:

$$\begin{aligned}
P(y_i - \bar{y}) &= P \left(1 - \frac{1}{W_{\text{tot}} \sigma_f^2} y_i - \frac{1}{W_{\text{tot}}} \sum_{j \neq i} \frac{y_j}{\sigma_j^2} \right) \\
&= N \left(1 - \frac{1}{W_{\text{tot}} \sigma_f^2} y_{\text{true}}, 1 - \frac{1}{W_{\text{tot}} \sigma_f^2} \sigma_i^2 \right) \sim N \left(-\frac{y_{\text{true}}}{W_{\text{tot}} \sigma_0^2}, \frac{1}{W_{\text{tot}} \sigma_0^2} \right) \sim \dots \\
&= N \left(0, \left[1 - \frac{1}{W_{\text{tot}} \sigma_f^2} \sigma_f^2 + \frac{1}{W_{\text{tot}}^2} \sum_{j \neq i} \frac{1}{\sigma_j^2} \right]^{1/2} \right) \\
&\equiv N \left(0, (1 - f_1)^2 \sigma_f^2 + \frac{f_2}{W_{\text{tot}}} \right)^{1/2} \\
&= N \left(0, f_2^2 \sigma_f^2 + f_2 \sigma_{\bar{y}}^2 \right)^{1/2},
\end{aligned} \tag{A.6}$$

where $f_1 = 1/(W_{\text{tot}} \sigma_f^2)$ is the fraction of weight in sample i , $f_2 = 1/W_{\text{tot}} \sum_{j \neq i} 1/\sigma_j^2 = 1 - f_1$ is the fraction of weight in the other samples, and $\sigma_{\bar{y}} = 1/W_{\text{tot}}$.

If all the weights are equal, i.e., $\sigma = \sigma_0$ for all i , then $f_2 = (N-1)/N$, $\sigma_{\bar{y}} = \sigma_0 / \sqrt{N}$, and

$$\begin{aligned}
P(y_i - \bar{y}) &= N \left(0, \frac{N-1}{N} \sigma_0^2 + \frac{N-1}{N^2} \sigma_0^2 \right)^{1/2} \\
&= N \left(0, \frac{1}{N} \sigma_0^2 \right)^{1/2},
\end{aligned} \tag{A.7}$$

and every variable “shares” the loss of the one degree of freedom equally. But if the weights for the different y_i are not identical, the sharing is not equal. In particular, in the limiting case of one variable dominating the weights, for that y_i , $f_2 \rightarrow 1$, $\sigma_{\bar{y}} \rightarrow \sigma_i$, and

$$P(y_i - \bar{y}) \rightarrow N \left(0, \sigma_i^2 \right), \tag{A.8}$$

i.e., the variance of the residuals of this particular y_i relative to the best-fit y are heavily suppressed compared to variance of y_i itself (because y is dominating the fit).

Note that the total χ^2 of the fit over all variables is not affected by the distribution of weights. I.e., it can be shown from the variance of $y_i - \bar{y}$ in equation (A.6) that

$$\chi^2 = \sum_i \frac{(y_i - \bar{y})^2}{\sigma_i^2} = N - 1 \quad (\text{A.9})$$

regardless of the distribution of σ_i .

References

- [1] COBE collaboration, *Structure in the COBE differential microwave radiometer first year maps*, *Astrophys. J.* **396** (1992) L1 [[INSPIRE](#)].
- [2] Planck collaboration, *Planck 2015 results. XIII. Cosmological parameters*, *Astron. Astrophys.* **594** (2016) A13 [[arXiv:1502.01589](#)] [[INSPIRE](#)].
- [3] K.T. Story et al., *A Measurement of the Cosmic Microwave Background Damping Tail from the 2500-square-degree SPT-SZ survey*, *Astrophys. J.* **779** (2013) 86 [[arXiv:1210.7231](#)] [[INSPIRE](#)].
- [4] Planck collaboration, *Planck 2013 results. XVI. Cosmological parameters*, *Astron. Astrophys.* **571** (2014) A16 [[arXiv:1303.5076](#)] [[INSPIRE](#)].
- [5] Z. Hou et al., *A Comparison of Maps and Power Spectra Determined from South Pole Telescope and Planck Data*, *Astrophys. J.* **853** (2018) 3 [[arXiv:1704.00884](#)] [[INSPIRE](#)].
- [6] SPT collaboration, *A Comparison of Cosmological Parameters Determined from CMB Temperature Power Spectra from the South Pole Telescope and the Planck Satellite*, *Astrophys. J.* **850** (2017) 101 [[arXiv:1706.10286](#)] [[INSPIRE](#)].
- [7] E.M. George et al., *A measurement of secondary cosmic microwave background anisotropies from the 2500-square-degree SPT-SZ survey*, *Astrophys. J.* **799** (2015) 177 [[arXiv:1408.3161](#)] [[INSPIRE](#)].
- [8] S. Padin et al., *South Pole Telescope optics*, *Appl. Opt.* **47** (2008) 4418.
- [9] J.E. Carlstrom et al., *The 10 Meter South Pole Telescope*, *Publ. Astron. Soc. Pac.* **123** (2011) 568 [[arXiv:0907.4445](#)] [[INSPIRE](#)].
- [10] R. Keisler et al., *A Measurement of the Damping Tail of the Cosmic Microwave Background Power Spectrum with the South Pole Telescope*, *Astrophys. J.* **743** (2011) 28 [[arXiv:1105.3182](#)] [[INSPIRE](#)].
- [11] E. Hivon, K.M. Górski, C.B. Netterfield, B.P. Crill, S. Prunet and F. Hansen, *Master of the cosmic microwave background anisotropy power spectrum: a fast method for statistical analysis of large and complex cosmic microwave background data sets*, *Astrophys. J.* **567** (2002) 2 [[astro-ph/0105302](#)] [[INSPIRE](#)].
- [12] M. Lueker et al., *Measurements of Secondary Cosmic Microwave Background Anisotropies with the South Pole Telescope*, *Astrophys. J.* **719** (2010) 1045 [[arXiv:0912.4317](#)] [[INSPIRE](#)].
- [13] E. Shirokoff et al., *Improved constraints on cosmic microwave background secondary anisotropies from the complete 2008 South Pole Telescope data*, *Astrophys. J.* **736** (2011) 61 [[arXiv:1012.4788](#)] [[INSPIRE](#)].
- [14] G. Polenta et al., *Unbiased estimation of angular power spectrum*, *JCAP* **11** (2005) 001 [[astro-ph/0402428](#)] [[INSPIRE](#)].
- [15] G. De Zotti et al., *Predictions for high-frequency radio surveys of extragalactic sources*, *Astron. Astrophys.* **431** (2005) 893 [[astro-ph/0410709](#)] [[INSPIRE](#)].

- [16] J.D. Vieira et al., *Extragalactic millimeter-wave sources in South Pole Telescope survey data: source counts, catalog and statistics for an 87 square-degree field*, *Astrophys. J.* **719** (2010) 763 [[arXiv:0912.2338](#)] [[IN SPIRE](#)].
- [17] L.M. Mocanu et al., *Extragalactic Millimeter-wave Point-source Catalog, Number Counts and Statistics from 771 deg² of the SPT-SZ Survey*, *Astrophys. J.* **779** (2013) 61 [[arXiv:1306.3470](#)] [[IN SPIRE](#)].
- [18] K.M. Górski et al., *HEALPix — A Framework for high resolution discretization and fast analysis of data distributed on the sphere*, *Astrophys. J.* **622** (2005) 759 [[astro-ph/0409513](#)] [[IN SPIRE](#)].
- [19] O.J. Dunn, *Multiple Comparisons among Means*, *J. Am. Statist. Assoc.* **56** (1961) 52 [[IN SPIRE](#)].