Draft version December 28, 2021
Typeset using LaTeX **twocolumn** style in AASTeX631

The Simons Observatory: a new open-source power spectrum pipeline applied to the Planck legacy data

Zack Li,[1,2] Thibaut Louis,[3] Erminia Calabrese,[4] Hidde Jense,[4] David Alonso,[5] J. Richard Bond,[2] Steve K. Choi,[6,7] Jo Dunkley,[1,8] Giulio Fabbian,[9,4] Xavier Garrido,[3] Andrew H. Jaffe,[10] Mathew S. Madhavacheril,[11,12] P. Daniel Meerburg,[13] Umberto Natale,[4] and Frank J. Qu[14]

[1] *Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA*
[2] *Canadian Institute for Theoretical Astrophysics, 60 St. George Street, University of Toronto, Toronto, ON, M5S 3H8, Canada*
[3] *Université Paris-Saclay, CNRS/IN2P3, IJCLab, 91405 Orsay, France*
[4] *School of Physics and Astronomy, Cardiff University, The Parade, Cardiff, Wales CF24 3AA, UK*
[5] *Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, United Kingdom*
[6] *Department of Physics, Cornell University, Ithaca, NY 14853, USA*
[7] *Department of Astronomy, Cornell University, Ithaca, NY 14853, USA*
[8] *Joseph Henry Laboratories of Physics, Jadwin Hall, Princeton University, Princeton, NJ, USA 08544*
[9] *Center for Computational Astrophysics, Flatiron Institute, 162 5th Ave, New York, NY 10010, USA*
[10] *Blackett Laboratory, Prince Consort Road, Imperial College, London SW7 2AZ UK*
[11] *Perimeter Institute for Theoretical Physics, 31 Caroline Street N, Waterloo ON N2L 2Y5 Canada*
[12] *Department of Physics and Astronomy, University of Southern California, Los Angeles, CA, 90007, USA*
[13] *Van Swinderen Institute for Particle Physics and Gravity, University of Groningen, Nijenborgh 4, 9747 AG Groningen, The Netherlands*
[14] *DAMTP, Centre for Mathematical Sciences, Wilberforce Road, Cambridge CB3 0WA, UK*

ABSTRACT

We present a reproduction of the *Planck* 2018 angular power spectra at $\ell > 30$, and associated covariance matrices, for intensity and polarization maps at 100, 143 and 217 GHz. This uses a new, publicly available, pipeline that is part of the `PSpipe` package. As a test case we use the same input maps, ancillary products, and analysis choices as in the *Planck* 2018 analysis, and find that we can reproduce the spectra to 0.1σ precision, and the covariance matrices to 10%. We show that cosmological parameters estimated from our re-derived products agree with the public *Planck* products to 0.1σ, providing an independent cross-check of the *Planck* team's analysis. Going forward, the publicly-available code can be easily adapted to use alternative input maps, data selections and analysis choices, for future optimal analysis of *Planck* data with new ground-based Cosmic Microwave Background data.

1. INTRODUCTION

The Cosmic Microwave Background (CMB) has helped build a precisely-measured Standard Model of Cosmology. The current state-of-the-art is dominated by legacy data from the *Planck* satellite mission, which obtained cosmological constraints from analysis of the CMB anisotropies in temperature and polarization. Results from the final collaboration data release were presented in Planck Collaboration V (2020) and Planck Collaboration VI (2020) (hereafter, PL20 and PC20) for the power spectra and likelihood, and derived cosmology. A description of a re-analysis using the *Planck* team's alternative `CamSpec` code was also presented in Efstathiou & Gratton (2019). We expect to match and then surpass the statistical power of *Planck* from the ground in the coming decade, through ongoing and future ground-based surveys of the CMB, such as from the Atacama Cosmology Telescope (ACT, Thornton et al. 2016), the South Pole Telescope (SPT, Benson et al. 2014), the upcoming Simons Observatory (SO, Simons Observatory Collaboration 2019), and CMB-S4 (Abazajian et al. 2019) at the end of the decade.

Planck will provide a critical part of new cosmological constraints coming from ground-based surveys at the South Pole and Chile, for at least the next decade. *Planck* measured the temperature power spectrum to the cosmic variance limit up to multipoles of $\ell \sim 1800$ with excellent control of systematic effects, and such measurements at low to moderate scales in temperature are challenging to make robustly from the ground due to the atmosphere and ground pick-up. Polarization measurements from *Planck* were also made across the full sky, containing cosmological information currently inaccessible from the ground. Until the planned LiteBIRD satellite mission delivers science results near the end of this decade (Hazumi et al. 2020), the *Planck*

legacy dataset will be a part of all cosmological constraints from the CMB at large scales in both temperature and polarization. The scales measured by *Planck* anchor cosmology in important ways, particularly for constraining the primordial power spectrum parameters n_s and A_s, and the optical depth to reionization, $\tau_{\rm reio}$, as well as properties such as extra light relic particles, isocurvature fluctuations, and dark matter interactions with baryons.

The set of power spectra and the associated covariance matrix is the core product for many analyses of the CMB. Until now, joint analysis of *Planck* with high-resolution experiments like ACT (Aiola et al. 2020; Choi et al. 2020) and SPT (Dutcher et al. 2021) has usually been done by combining power spectrum likelihoods, avoiding calculating the covariance between overlapping maps from the different experiments. With the sky coverage of ground-based surveys now increasing considerably (De Bernardis et al. 2016; Naess et al. 2020), future state-of-the-art cosmological model constraints will require estimating new cross-spectra and covariance matrices for the *Planck* maps in combination with the new datasets. One of the goals of this paper is to take a step towards these future data combinations by demonstrating a new, publicly-available, pipeline for estimating the *Planck* spectra and covariances. The modern cosmological model derived from the *Planck* data also exhibits a number of possible tensions (Planck Collaboration V 2020), either internally between different subsets of the data, or with external data (for parameters including the Hubble constant, H_0, the amplitude of structure formation, σ_8, and the lensing amplitude, A_L) that could be hints of new components in the model. Further scrutiny of the *Planck* maps with an independent pipeline therefore serves as an important check for these glimpses of possible discoveries.

In this paper, we provide a detailed description of the construction of cross-spectra, covariance matrices, and the resulting multi-frequency likelihood obtained from analysis of the high-ℓ *Planck* 2018 data release (PR3), made using the power spectrum pipeline developed for the Simons Observatory. While this paper focuses on reproducing and testing the approach implemented by the *Planck* collaboration in PL20, the new machinery is flexible for modified use, for example for computing statistics over different sky regions and in cross-correlation with new ground-based data. We chose to implement the analysis choices in PL20 in this paper, to provide a concrete comparison point from which to make future changes. The *Planck* analysis is complicated, and an open-source pipeline will allow different groups to explore optimum analysis and map choices, for example following those made for the `CamSpec` likelihood in Efstathiou & Gratton (2019), or for the use of new `NPIPE` maps in Planck Collaboration LVII (2020). We anticipate making our own new analysis choices for the combination of *Planck* with Simons Observatory data in future work. The pipeline for this particular reproduction is released as part of the Simons Observatory power spectrum package, `PSPipe`[1]. We present detailed documentation to accompany this code[2].

The paper is organized as follows. In Sec. 2 we describe the data products which enter into the analysis. In Sec. 3 we present the power spectrum methods, results, and comparisons to the public *Planck* data products. In Sec. 4 we describe the methods used to estimate analytic covariance matrices. In Sec. 5 we summarize the modeling of the spectra and in Sec. 6 we show cosmological parameters estimated using our power spectrum products. We conclude in Sec. 7.

2. DATA PRODUCTS

We use the maps and ancillary data products as described in PL20, and summarized below. Most of the products require some manipulation in order to closely reproduce the *Planck* PR3 analysis, and we summarize these modifications here.

2.1. *Maps*

The primary inputs for this analysis are the publicly available *Planck* HFI half-mission maps from the 2018 data release. We expect to provide an analogous analysis of `NPIPE` maps (Planck Collaboration LVII 2020) in future work. We follow previous *Planck* cosmology analyses and focus on the 100, 143, and 217 GHz half-mission frequency maps, shown in Figure 1. All of these maps are in HEALPix pixelization with $N_{\rm side} = 2048$, and have missing pixels marked with the HEALPix `UNSEEN` pixel value. As in the PL20 analysis, we restrict our analysis to computing power spectra between different half-mission splits, to avoid both noise bias and time-dependent systematic effects. We use the maps in I, Q, U Stokes vectors, as well as the pixel variance II, QQ, and UU. We do not use the QU intra-pixel covariances. We use these maps to estimate the spectra as well as the noise in the construction of the covariance matrix, described in Sec. 4. This differs from what is done in PL20, in which 'half-ring half-mission maps', that were internally available to the *Planck* team, are used for the noise estimation.

[1] https://github.com/simonsobs/PSpipe/blob/master/project/Planck_cov

[2] https://simonsobs.github.io/planck-pr3-web/

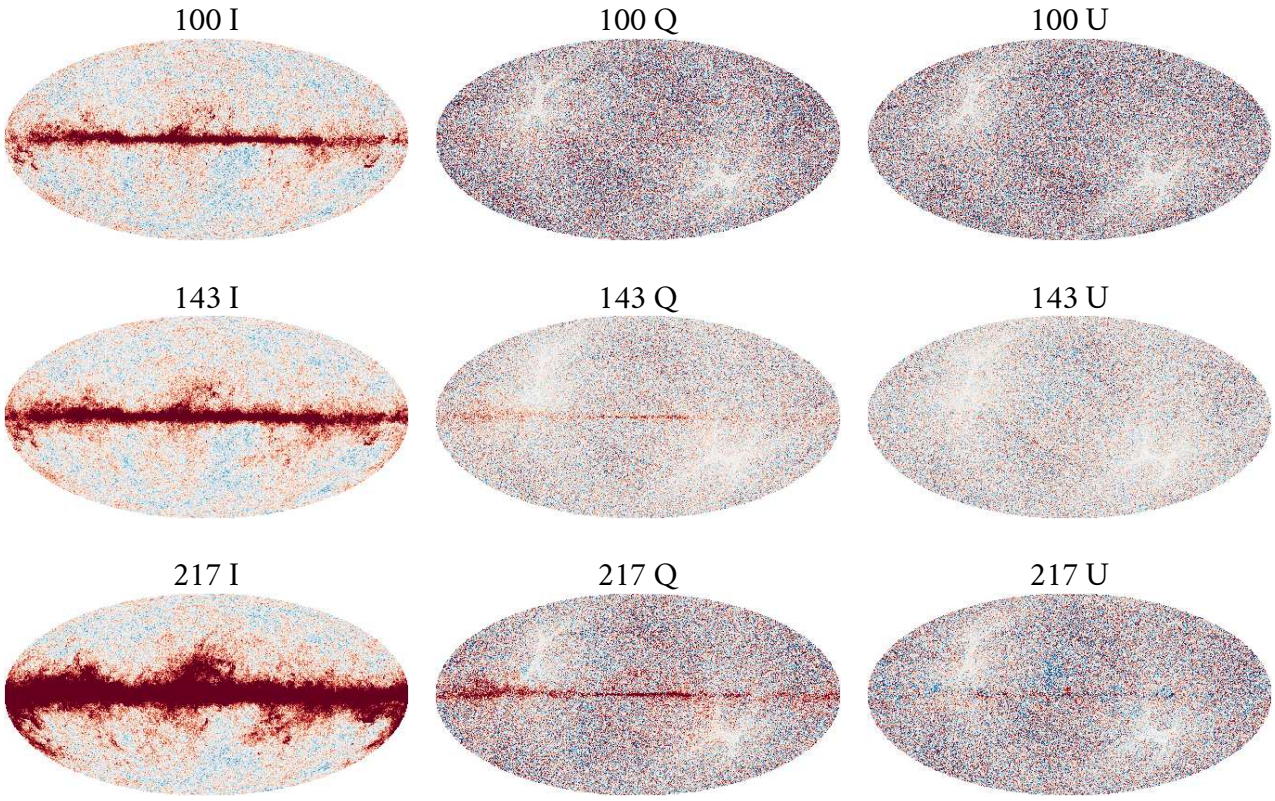

Figure 1. *Planck* 2018 maps from the first half-mission at 100, 143 and 217 GHz (Planck Collaboration VI 2020), with I, Q, and U Stokes vectors, as produced and presented in Planck Collaboration VI (2020) and used here. The second half-mission maps look similar, and we use the two half-mission map splits to compute spectra, as done in PL20. Maps are Mollweide projections with a color scale range of ± 500 μK for Stokes I and ± 200 μK in Stokes Q/U.

Polarization calibration: The calibration of the *Planck* polarization maps relative to the temperature maps is uncertain, as they lack the bright dipole used to calibrate the temperature maps. The relative calibration is limited by the accuracy with which their polarization efficiencies could be determined either on the ground or in flight. To best match the spectra in the public `Plik` likelihood we determined that calibration factors of 0.9995 for 100 GHz, 0.999 for 143 GHz, and 0.999 for 217 GHz should be multiplied with the publicly released maps. We note that the uncertainty in polarization calibration ($\sim 1\%$) dwarfs these ad-hoc calibration factors.

2.2. *Masks*

As in PL20, we use the publicly released `plik` sky masks shown in Figure 2 to avoid bright contaminated regions caused by the Galaxy and by extragalactic point sources. These comprise the temperature and polarization masks at 100, 143 and 217 GHz since both the point source and Galactic masks are different at each frequency. Polarized point source emission is negligible at *Planck* noise levels, so the polarization masks do not include point source holes. The Galactic masks cover 66%, 57%, and 47% of the sky at 100, 143, and 217 GHz (Planck Collaboration XI 2016b; Planck Collaboration V 2020) and are apodized with a 4.71 degree FWHM ($\sigma = 2^\circ$) Gaussian window function, whereas the point source mask is apodized with a 30 arcminute FWHM. The masks used for estimating the *Planck* power spectrum do not weight for noise or hits.

These masks also include a sharp mask on the 'missing pixels' in the maps by setting the mask values at the locations of the missing pixels to zero[3]. This represents a small fraction of the total pixels at $N_{\rm side} = 2048$ (corresponding to 50,331,648 pixels), with the most impacted frequency map (217 GHz hm2) missing 0.2% of the pixels. Although the mask is not band-limited, simulations have shown that the estimated pseudo-C_ℓ spec-

[3] There are 2,528 missing pixels for 100 GHz hm1, 50,917 pixels for 100 GHz hm2, 6,531 pixels for 143 GHz hm1, 64,307 pixels for 143 GHz hm2, 13,358 pixels for 217 GHz hm1, and 115,439 pixels for 217 GHz hm2.

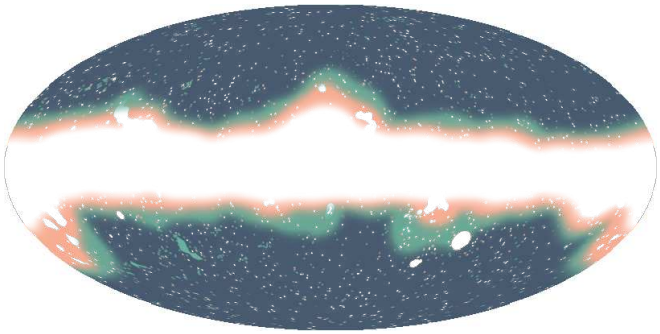

Figure 2. The *Planck* masks which combine Galactic and point source masking. We show the 100 GHz (orange), 143 GHz (green), and 217 GHz (dark blue) masks in temperature overlaid. The corresponding polarization mask for each frequency does not have the point source holes. Point sources detected with $S/N > 5$ are masked, with holes of radius $\sigma = FWHM/\sqrt{\ln 8}$ of the effective Gaussian beam at that frequency (Planck Collaboration XI 2016a). The colors correspond to 66%, 57% and 47% retained sky fraction at each frequency respectively, corresponding to the average of the square of the weight function. For each frequency, the parts of the mask that are colored refer to areas which are not masked at that frequency.

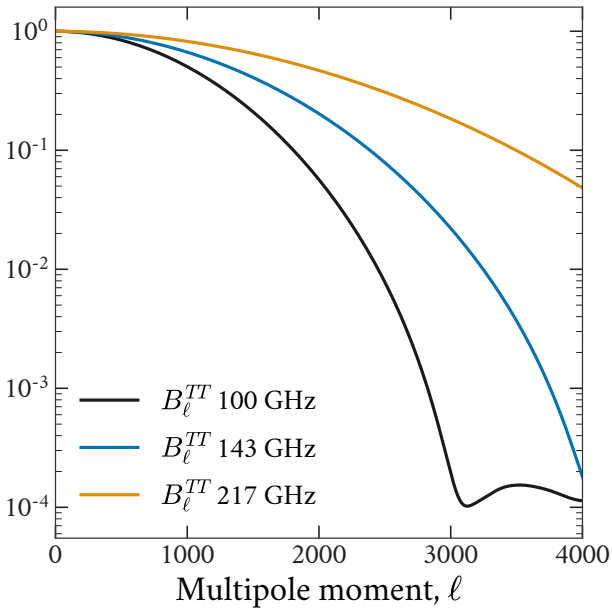

Figure 3. The temperature beam window $B_\ell^{TT} = (W_\ell^{TT,TT})^{1/2}$ for three half-mission-1 frequency maps, as released in the Planck Legacy Archive `RIMO R3.01`. Polarization beams differ by a few percent at the smallest scales that *Planck* can measure ($\ell \sim 2500$).

tra are not significantly biased Planck Collaboration XI (2016a).

The PR3 maps also feature a few thousand pixels at moderate Galactic latitude with extraordinarily high estimated pixel variance, σ_p, particularly in polarization. These are understood to arise from pixels with very few hits, which sometimes led to the inversion of a nearly-degenerate matrix during map processing. These are most noticeable at 100 GHz, where these pixels can have $\sim 10^6\times$ more variance than the typical pixel. In this analysis we apply an additional mask which imposes a sharp pixel cut on pixels with QQ and UU variance above $10^6\ \mu\mathrm{K}^2_{\mathrm{CMB}}$. These pixels do not have an effect on spectra estimation, regardless of mask – however, their huge estimated variance would contaminate the covariance estimation, where the noise-like term involves a power spectrum of the noise variance map. We use the union of these missing pixels masks across the I, Q, and U channels for all as the missing pixel mask for all channels.

2.3. *Beams*

The beam window function relates the signal angular power spectrum $C_{\mathrm{map}}(\ell)$ of a map to the true underlying sky angular power spectrum $C_{\mathrm{sky}}(\ell)$. Consider spectrum channels A, $B \in \{T, E, B\}$. We use the beams generated from the `QuickPol` formalism (Hivon et al. 2017a), which computes a beam matrix $W_{AB,A'B'}(\ell)$, such that

$$C^{AB}_{\mathrm{map}}(\ell) = \sum_{A'B'} W_{AB,A'B'}(\ell)\, w^2_{\mathrm{pix}}(\ell)\, C^{A'B'}_{\mathrm{sky}}(\ell). \qquad (1)$$

Here $W_{AB,A'B'}(\ell)$ is the instrumental beam transfer function for the contribution of spectrum $A'B'$ to spectrum AB. The binning into pixels also induces $w_{\mathrm{pix}}(\ell)$, the pixel window. For example, $W_{EE,TT}(\ell)$ is the mixing of the true TT spectrum into the observed EE spectrum. The *Planck* beam is azimuthally asymmetric, so the average of the scanning beams for a given pixel depends on the location of the pixel and the scanning strategy. The choice of mask therefore defines an effective, azimuthally and spatially averaged, beam. Since we use the likelihood mask following the PL20 analysis, these beams were already computed and released in the `RIMO R3.01` for multipoles up to 4000, and are shown in Figure 3. We smoothly extend the beam in a power law for multipoles from 4000 up to the maximum multipole $\sim 3n_{\mathrm{side}}$ induced by the Nyquist frequency. We also multiply this beam correction by the corresponding temperature or polarization HEALPix pixel window function for $N_{\mathrm{side}} = 2048$. These beams are sufficient for this re-analysis of the *Planck* maps, but future cross-

correlations with ground-based surveys will require recomputing the *Planck* beam for each new mask using `QuickPol` (Hivon et al. 2017b). We also separately provide a `QuickPol` implementation in the software package developed for this analysis, `PowerSpectra.jl`.

2.4. *Binning*

We first compute unbinned, mask-deconvolved spectra in $2 \leq \ell \leq 2508$. To pass these new inputs to the *Planck* likelihood for cosmological analysis and also make comparisons with the *Planck* spectra, we then bin these spectra with the same scheme as PL20. We use an $\ell(\ell+1)$ weighting for C_ℓ inside each bin, over bins of $\Delta\ell = 5$ for $30 \leq \ell \leq 99$, $\Delta\ell = 9$ for $100 \leq \ell \leq 1503$, $\Delta\ell = 17$ for $1504 \leq \ell \leq 2013$, and $\Delta\ell = 33$ for $2014 \leq \ell \leq 2508$. Typically the mode-coupling matrix is binned before it is inverted. The binning operation regularizes the matrix, making it more tractable to invert, particularly for ground-based experiments with partial sky coverage. However, *Planck* is a full-sky experiment and the mode-coupling matrix of these masks is well-conditioned. Inverting before binning allows the *Planck* bandpower window function to be a top-hat with a $\ell(\ell+1)$ weighting within the bins.

3. POWER SPECTRA

The analysis in this paper uses `PSPipe`[4], a general purpose public pipeline for computation and analysis of CMB power spectra, and in this paper we expand it to deliver the full *Planck* high-ℓ power spectrum analysis pipeline.

In the nominal analysis shown in this paper we use a `PSPipe` pipeline implementation in the Julia language, `PowerSpectra.jl`, as this high-level language allowed for the easy construction of the approximate covariance matrix computed in this work. In particular, the summations involved in these expressions map easily to for-loops that are inefficient in Python. We also perform a number of cross-checks of the power spectra computed with the other two codes in `PSPipe`: `NaMaster` (Alonso et al. 2019) and `pspy` (Louis et al. 2020). The three codes within the pipeline have been tested against each other, and agree to numerical precision. The multiple backends within `PSPipe` accelerates the analysis process, by providing independent checks and reference implementations, complementing existing testing suites for the individual codes.

As part of `PSPipe`, we provide a full suite of scripts which reproduce every step of the power spectrum analysis presented in the paper, from automatic download of the data to cosmological parameter estimation. We expect this to be an important resource to be used both directly and as a reference implementation, for future large-sky-area surveys in the millimeter band.

[4] https://github.com/simonsobs/PSpipe

3.1. *Power spectrum estimation*

We compute the power spectra at 100, 143 and 217 GHz, using the products described in Sec. 2. We begin by producing the masked maps, multiplying the half-mission maps with the masks described in Sec. 2.2.

Monopole and Dipole Subtraction: We then estimate and subtract from each frequency map a monopole and dipole, using the unmasked pixels. The monopole, in particular, is extremely large in comparison to the scalar fluctuations even at large scales. The mask couples the monopole to other scales, and we encounter issues in accurately treating this monopole mode coupling, as the monopole is so much larger relative to the power on other scales. Small errors in estimating the mode-coupling matrix elements related to the monopole are multiplied by the large monopole, leading to significant bias on the resulting spectra. The power spectrum analysis itself does not use information from the monopole and dipole modes. We fit both the monopole and dipole simultaneously using the `healpy` package, and conservatively restrict this fit to map pixels where the corresponding mask pixel has value 1, excluding regions where the mask is 0 or apodized to a value less than 1. We confirmed that this choice does not substantially affect the determination of the monopole and dipole, compared to simply using the mask. Fitting with the mask avoids bias in the monopole and dipole from the point sources, the Galaxy, and missing pixels.

Mode-Coupling: We then compute the mode-coupled spectrum, $\tilde{C}_\ell$, between pairs of fields,

$$\tilde{C}_\ell^{ABij} = \frac{1}{2\ell+1} \sum_m \mathsf{m}_{\ell m}^{i,A} \mathsf{m}_{\ell m}^{*j,B}, \tag{2}$$

where $\mathsf{m}_{\ell m}^{i,A}$ are the spherical harmonic coefficients of the masked map i of channel $A \in \{T, E\}$. The effect of the mask is then decoupled, by solving for the decoupled spectrum $\hat{C}_\ell$ following

$$\tilde{C}_\ell^{ABij} = \sum_{\ell'} M_{\ell\ell'}^{ABij} \hat{C}_{\ell'}^{ABij}, \tag{3}$$

where $M_{\ell\ell'}^{ABij}$is the mode-coupling matrix between fields i and j for the cross-spectrum $AB \in \{TT, TE, EE, BB\}$ (Hivon et al. 2002; Alonso et al. 2019). A field i refers to some combination of frequency and data split. Note that although we compute the BB pseudo-spectrum, we only use it in this analysis in order to remove its contribution

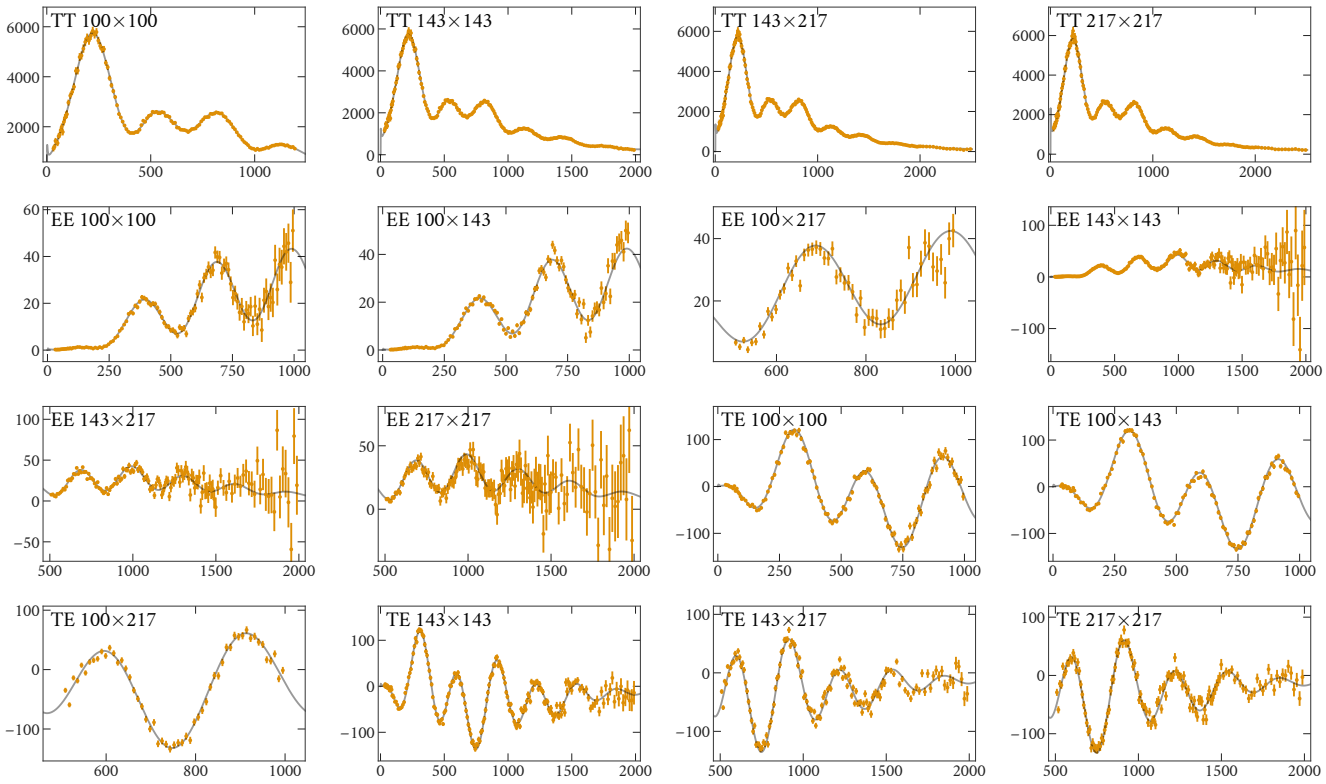

Figure 4. Re-estimated *Planck* spectra and errors at 100, 143 and 217 GHz, and the cross-frequency spectra, in terms of $D_\ell = \ell(\ell+1)C_\ell/2\pi$ in units of $\mu\mathrm{K}^2$. We show just the angular range and suite of cross-spectra used in the likelihood, following the multipole cuts in PL20. We show the PR3 best-fit theory in the solid black line.

to the EE power spectrum through mode coupling. This is important primarily for the noise power spectrum, where EE and BB contributions are comparable. The PL20 analysis used `PolSpice` for this calculation (Chon et al. 2004), which performs this same operation using correlation functions.

We compute mode-decoupled spectra for every cross spectrum $C_\ell(\mathsf{m}^{i,A}, \mathsf{m}^{j,B})$ for frequency-split combinations with frequencies $\in \{100, 143, 217\}$ and splits $\in \{\mathrm{hm1}, \mathrm{hm2}\}$. Cross-spectra within the same half-mission are not used for signal calculations. We combine TE and ET spectra with a flat average, i.e. $(C_\ell^{TE} + C_\ell^{ET})/2$. To estimate the data cross-spectra used in the likelihood, we then correct for the beam and pixel window by applying their inverse to the decoupled spectra. Note that this differs from PL20, which used inverse-variance weighting in order to combine spectra.

3.2. *Comparison to PL20 spectra*

The spectra at each of the three frequencies, for TT, TE and EE, are shown in Figure 4. For spectra which are combined, such as TT from different frequencies and half-missions or TE and ET, are combined together using uniform weighting. The residuals compared to the PL20 spectra for 143 GHz TT, TE and EE are in Figure 5, computed here using the errors from the PL20 products. The residuals compared to the full suite of spectra are shown in the Appendix in Figure 15, including cross-frequencies. We find general agreement to generally at the 0.05σ level, with a notable deviation in EE at large scales up to 0.3σ, and TE at small scales to 0.1σ. Considering the same angular ranges as in PL20, the full set of binned spectra can be concatenated to form a single data vector of total length 2290. The ordering of this vector is shown in Table 3.2, together with the multipole ranges.

4. COVARIANCE MATRIX

We now review the method for hybrid estimation of covariance matrices, building on Hansen et al. (2002); Hinshaw et al. (2003); Efstathiou (2004); Brown et al. (2005); Challinor & Chon (2005); Efstathiou (2006); Planck Collaboration XV (2014); Planck Collaboration XI (2016b), used in PL20 to produce the *Planck* 2018 likelihood. We describe the details of our implementation, and note where this differs from PL20 due to

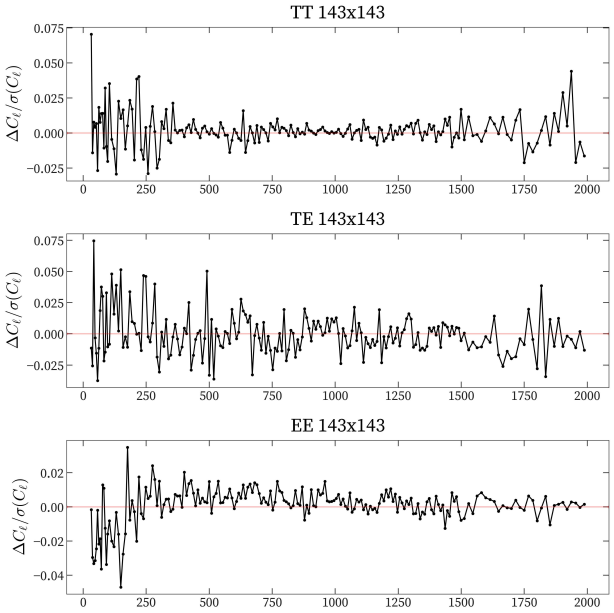

Figure 5. Residuals between our re-estimated power spectra and the public *Planck* PL20 spectra at 143 GHz, shown as a fraction of a standard deviation. The agreement is within 0.08σ over the full angular range; the EE agrees to within 0.05σ, with the largest deviation at large angular scales. The residuals for the full suite of spectra are in Figure 15 in the Appendix.

the available products, and where the approach differs from the analytic covariance method used in `NaMaster` (García-García et al. 2019) and `pspy`.

Following PL20, the model describing the time-ordered data is

$$d_t = (\mathbf{P}_{tp}\mathbf{B}_t\mathbf{w}_p)\, s_p + n_t, \tag{4}$$

for data d_t, pointing matrix $\mathbf{P}_{tp}$, beam operator $\mathbf{B}_{tp}$, pixel window operator $\mathbf{w}_p$, true underlying signal map s_p, and noise n_t (Hivon et al. 2002). Note that the noise is not affected by any of the beam-related operators. The likelihood assumes Gaussian variance about a fiducial signal power spectrum defined by

$$C_\ell^{AB,ij} \equiv W_\ell^{AB,ij}\, p_\ell^2 \left(C_\ell^{\rm AB,CMB} + C_\ell^{\rm AB,FG}(f_i, f_j)\right), \tag{5}$$

where f_k is the frequency of the field k, $W_\ell^{AB,ij}$ is the beam window function, and p_ℓ the pixel window function. In units of $\mu{\rm K}_{\rm CMB}$, the CMB power spectrum stays constant between frequencies, but the foreground signal varies. As in PL20, the fiducial CMB spectrum is given by the *Planck* best-fitting ΛCDM model, with foreground model given in Sec. 5.

The covariance matrix for the suite of cross-spectra is estimated for the combination of signal and noise on the masked sky. It combines the fiducial power spectra and

AB	f_1	f_2	$\ell_{\rm min}$	$\ell_{\rm max}$
TT	100	100	30	1197
TT	143	143	30	1996
TT	143	217	30	2508
TT	217	217	30	2508
EE	100	100	30	999
EE	100	143	30	999
EE	100	217	505	999
EE	143	143	30	1996
EE	143	217	505	1996
EE	217	217	505	1996
TE	100	100	30	999
TE	100	143	30	999
TE	100	217	505	999
TE	143	143	30	1996
TE	143	217	505	1996
TE	217	217	505	1996

Table 1. Multipole ranges for the data vector for each cross-spectrum between frequency f_1 and f_2, using the same choices as in PL20. The TE is an average of TE and ET for cross-frequency spectra. All cross-spectra that arises from different half-missions are included in an average, for any combination of frequencies and channel $AB \in \{TT, TE, EE\}$. The ordering of the rows matches the ordering in the data vector used in the likelihood.

projector functions (the product of the two masks) that account for the mode-coupling from the mask and non-uniform noise, as described in Planck Collaboration XI (2016a) (PL15). We include the expression here for the covariance matrix between a TT cross-spectrum from maps i, j and a TT cross-spectrum from maps p, q, to demonstrate the overall structure of the calculation:

$$\begin{aligned}
&{\rm Cov}(\hat{C}_\ell^{TT\,i,j}, \hat{C}_\ell^{TT\,p,q}) \\
&\quad\approx \sqrt{C_\ell^{TT\,i,p} C_{\ell'}^{TT\,i,p} C_\ell^{TT\,j,q} C_{\ell'}^{TT\,j,q}}\, \Xi_{TT}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,p)^{TT},(j,q)^{TT}\right]_{\ell\ell'} \\
&\quad+ \sqrt{C_\ell^{TT\,i,q} C_{\ell'}^{TT\,i,q} C_\ell^{TT\,j,p} C_{\ell'}^{TT\,j,p}}\, \Xi_{TT}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,q)^{TT},(j,p)^{TT}\right]_{\ell\ell'} \\
&\quad+ \sqrt{C_\ell^{TT\,i,p} C_{\ell'}^{TT\,i,p}}\, \Xi_{TT}^{\emptyset\emptyset,TT}\left[(i,p)^{TT},(j,q)^{TT}\right]_{\ell\ell'} \\
&\quad+ \sqrt{C_\ell^{TT\,j,q} C_{\ell'}^{TT\,j,q}}\, \Xi_{TT}^{\emptyset\emptyset,TT}\left[(j,q)^{TT},(i,p)^{TT}\right]_{\ell\ell'} \\
&\quad+ \sqrt{C_\ell^{TT\,i,q} C_{\ell'}^{TT\,i,q}}\, \Xi_{TT}^{\emptyset\emptyset,TT}\left[(i,q)^{TT},(j,p)^{TT}\right]_{\ell\ell'} \\
&\quad+ \sqrt{C_\ell^{TT\,j,p} C_{\ell'}^{TT\,j,p}}\, \Xi_{TT}^{\emptyset\emptyset,TT}\left[(j,p)^{TT},(i,q)^{TT}\right]_{\ell\ell'} \\
&\quad+ \Xi_{TT}^{TT,TT}\left[(i,p)^{TT},(j,q)^{TT}\right]_{\ell\ell'} + \Xi_{TT}^{TT,TT}\left[(i,q)^{TT},(j,p)^{TT}\right]_{\ell\ell'}
\end{aligned} \tag{6}$$

The $\Xi_{AB}^{X,Y}$ terms are projector functions describing mode-coupling from the mask for channel $AB \in \{TT, TE, EE\}$, and X, Y refer to the kind of noise-

variance weighting applied to the mask. $X = \emptyset\emptyset$ refers to the mask alone, $X = TT$ refers to weighting by the II variance, and PP refers to weighting by the QQ and UU variance. The suite of projector functions are given in the Appendix, and follow those in Appendix C1 of PL16.

The noise power spectrum is an important input for the covariance matrix. If the noise was diagonal in pixel space, leading to a white noise power spectrum, then Equation 6 would be a good approximation for the spectrum covariance. However, we see from Figure 6 that this is not the case (we describe the estimation of these noise terms in the next section), with many spectrum channels deviating significantly from the white noise prediction for the noise power spectrum. Following PL16, PL20, we define the ratio of the measured noise power spectrum with the predicted power spectrum from the measured white noise level,

$$\mathcal{R}^{i,X}_{\ell,\ell'} = \sqrt{\frac{N^{i,X}_{\ell,\rm data} N^{i,X}_{\ell',\rm data}}{N^{i,X}_{\ell,\rm white} N^{i,X}_{\ell',\rm white}}}. \tag{7}$$

In this case $N_{\ell,\rm data}$ refers to the effective white noise level presented in Table 4.1, constant with ℓ. A factor of $\mathcal{R}^{i,X}_\ell \mathcal{R}^{i,Y}_\ell$ is multiplied with each $\Xi^{X,Y}_{\ell,\ell'}$ for each factor of $X, Y \in \{TT, PP\}$, where $N^{PP}_\ell = N^{EE}_\ell$. For example, Equation 6 transforms to

$$\begin{aligned}
&\mathrm{Var}(\hat{C}^{TT\,i,j}_\ell, \hat{C}^{TT\,p,q}_\ell)_{\rm nonwhite} \\
&\quad\approx \sqrt{C^{TT\,i,p}_\ell C^{TT\,i,p}_{\ell'} C^{TT\,j,q}_\ell C^{TT\,j,q}_{\ell'}}\, \Xi^{\emptyset\emptyset,\emptyset\emptyset}_{TT} \left[(i,p)^{TT}, (j,q)^{TT}\right]_{\ell\ell'} \\
&\quad+ \sqrt{C^{TT\,i,q}_\ell C^{TT\,i,q}_{\ell'} C^{TT\,j,p}_\ell C^{TT\,j,p}_{\ell'}}\, \Xi^{\emptyset\emptyset,\emptyset\emptyset}_{TT} \left[(i,q)^{TT}, (j,p)^{TT}\right]_{\ell\ell'} \\
&\quad+ \sqrt{C^{TT\,i,p}_\ell C^{TT\,i,p}_{\ell'}}\, \Xi^{\emptyset\emptyset,TT}_{TT} \left[(i,p)^{TT}, (j,q)^{TT}\right]_{\ell\ell'} \mathcal{R}^{j,TT}_{\ell,\ell'} \mathcal{R}^{q,TT}_{\ell,\ell'} \\
&\quad+ \sqrt{C^{TT\,j,q}_\ell C^{TT\,j,q}_{\ell'}}\, \Xi^{\emptyset\emptyset,TT}_{TT} \left[(j,q)^{TT}, (i,p)^{TT}\right]_{\ell\ell'} \mathcal{R}^{i,TT}_{\ell,\ell'} \mathcal{R}^{p,TT}_{\ell,\ell'} \\
&\quad+ \sqrt{C^{TT\,i,q}_\ell C^{TT\,i,q}_{\ell'}}\, \Xi^{\emptyset\emptyset,TT}_{TT} \left[(i,q)^{TT}, (j,p)^{TT}\right]_{\ell\ell'} \mathcal{R}^{j,TT}_{\ell,\ell'} \mathcal{R}^{p,TT}_{\ell,\ell'} \\
&\quad+ \sqrt{C^{TT\,j,p}_\ell C^{TT\,j,p}_{\ell'}}\, \Xi^{\emptyset\emptyset,TT}_{TT} \left[(j,p)^{TT}, (i,q)^{TT}\right]_{\ell\ell'} \mathcal{R}^{i,TT}_{\ell,\ell'} \mathcal{R}^{q,TT}_{\ell,\ell'} \\
&\quad+ \Xi^{TT,TT}_{TT} \left[(i,p)^{TT}, (j,q)^{TT}\right]_{\ell\ell'} \mathcal{R}^{i,TT}_{\ell,\ell'} \mathcal{R}^{p,TT}_{\ell,\ell'} \mathcal{R}^{j,TT}_{\ell,\ell'} \mathcal{R}^{q,TT}_{\ell,\ell'} \\
&\quad+ \Xi^{TT,TT}_{TT} \left[(i,q)^{TT}, (j,p)^{TT}\right]_{\ell\ell'} \mathcal{R}^{i,TT}_{\ell,\ell'} \mathcal{R}^{q,TT}_{\ell,\ell'} \mathcal{R}^{j,TT}_{\ell,\ell'} \mathcal{R}^{p,TT}_{\ell,\ell'}.
\end{aligned} \tag{8}$$

The full suite of the blocks for TE, EE, and the covariances between these three sets of spectra is given in PL16 and repeated in the Appendix for completeness.

This covariance matrix expression is not quite the same as the Gaussian, uniform-noise approximations typically used by `NaMaster` and `pspy`. There one defines the total power spectrum

$$P^{i,A}_\ell \equiv C^{i,A}_\ell + N^{i,A}_\ell, \tag{9}$$

with covariance, for the TT block for example, given by

$$\begin{aligned}
&\mathrm{Var}(\hat{C}^{TT\,i,j}_\ell, \hat{C}^{TT\,p,q}_\ell)_{\rm uniform} \\
&\quad\approx \sqrt{P^{TT\,i,p}_\ell P^{TT\,i,p}_{\ell'} P^{TT\,j,q}_\ell P^{TT\,j,q}_{\ell'}}\, \Xi^{\emptyset\emptyset,\emptyset\emptyset}_{TT} \left[(i,p)^{TT}, (j,q)^{TT}\right]_{\ell\ell'} \\
&\quad+ \sqrt{P^{TT\,i,q}_\ell P^{TT\,i,q}_{\ell'} P^{TT\,j,p}_\ell P^{TT\,j,p}_{\ell'}}\, \Xi^{\emptyset\emptyset,\emptyset\emptyset}_{TT} \left[(i,q)^{TT}, (j,p)^{TT}\right]_{\ell\ell'}.
\end{aligned} \tag{10}$$

We note that `NaMaster` uses the arithmetic and not geometric mean for these covariances. In the limit of uniform noise, the projector functions Ξ^{TTTT}_{TT} are proportional to $N^2_\ell \Xi^{\emptyset\emptyset,\emptyset\emptyset}_{TT}$ and Equation 8 is equivalent to Equation 10. The *Planck* maps are sufficiently non-uniform in noise that approximating them as uniform, as implemented in `NaMaster` and `pspy`, can cause order 10% inaccuracies in the covariance matrix. While this may not be enough to substantially affect cosmological parameters, mis-estimating the covariance matrices to such an extent can have consequences for null tests. As a rough rule of thumb, Efstathiou & Gratton (2019) requires that for a data vector of length N, the noise estimates need to satisfy

$$\frac{\Delta\sigma^2}{\sigma^2} \lessapprox \sqrt{\frac{2}{N}}, \tag{11}$$

for accurate χ^2, which for *Planck* requires $\sim 1\%$ estimates of the covariance matrices. We found that failing to treat the noise anisotropy leads to covariance matrix differences on the order of 10%.

4.1. *Noise Model*

We require sufficiently accurate estimation of the TT, TE and EE noise power spectra to compute the contribution of noise to the covariance matrices. Here we use the model described in Planck Collaboration XV (2014); Efstathiou & Gratton (2019) for N_ℓ, assuming the *Planck* noise power spectrum follows a smooth power-law behavior at small and large scales:

$$N^{\rm model}_\ell = A \left(\frac{100}{\ell}\right)^\alpha + B \frac{(\ell/1000)^\beta}{(1 + (\ell/\ell_c)^\gamma)^\delta}. \tag{12}$$

This differs from the `plik` model used in PL20, which uses a more general polynomial expression. The *Planck* PL20 analysis relied on jointly fitting the large-scale term for both half-mission maps at a given frequency by using the half-mission half-ring difference maps, which can be used to estimate noise power spectra for each half-mission map. This two-stage fitting process makes it easier to use the general expression used in PL20.

These half-mission half-ring difference maps are not publicly available, so we use the half-mission maps instead, together with this simpler model. We find that

freq	split	A	B	α	β	ℓ_c	γ	δ
100	1	0.000849	0.705	0.117	0.654	$5 \cdot 10^5$	$8.51 \cdot 10^{-11}$	7.82
100	2	0.000849	0.705	0.11	0.654	$5 \cdot 10^5$	$8.51 \cdot 10^{-11}$	7.82
143	1	0.000255	0.15	$1.75 \cdot 10^{-5}$	1.96	$8.23 \cdot 10^3$	2.13	8
143	2	0.00024	0.289	0.00848	$2.13 \cdot 10^{-11}$	$5.49 \cdot 10^5$	−0.286	2.52
217	1	0.000734	1.91	0.000454	0	$1 \cdot 10^4$	2.12	−0.135
217	2	0.000579	3.41	0.000434	0	$7.3 \cdot 10^3$	6	6
100	1	0.00103	0.0985	$6.29 \cdot 10^{-5}$	−0.805	309	1.46	−1.28
100	2	0.00138	0.268	0.000703	0.93	$5 \cdot 10^5$	$-7.73 \cdot 10^{-8}$	$-7.86 \cdot 10^{-8}$
143	1	0.000407	0.877	0.000823	0.0725	$1.16 \cdot 10^3$	6	−0.0188
143	2	0.000891	0.313	0.000357	0.486	$1.14 \cdot 10^3$	−6	0.0132
217	1	$1.24 \cdot 10^{-8}$	8	0.000852	0.541	813	−0.799	−1.25
217	2	$4.26 \cdot 10^{-6}$	6.38	$7.27 \cdot 10^{-9}$	2.38	$3.43 \cdot 10^4$	−0.365	−8

Table 2. Noise model parameters for the TT (top) and EE (bottom) power spectra, for each frequency and half-mission map. We emphasize that the estimation of *Planck* noise spectra is a challenging problem, particularly at large scales, and this noise model we present here is primarily an intermediary product used for this power spectrum analysis. They are representative of, but do not fully characterize, the *Planck* noise properties. These model parameters provide a smooth fit to the estimated noise spectra (see Appendix Figure 17 and 18). This full table is available in the PSpipe repository.

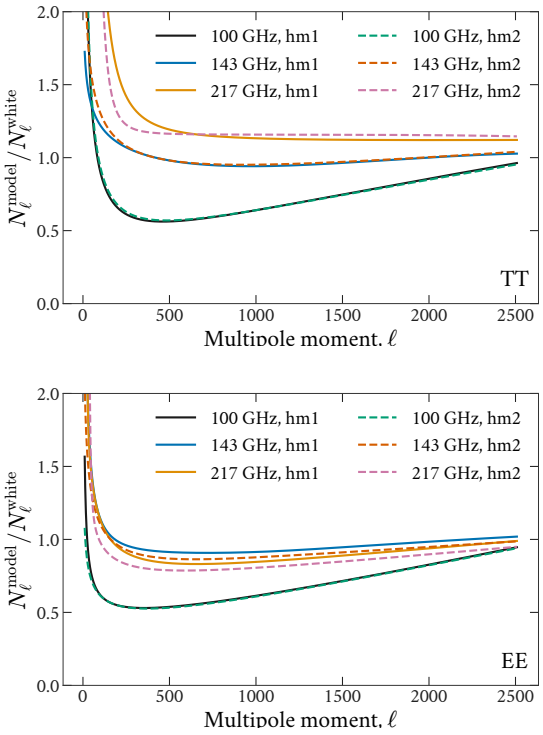

Figure 6. The ratio between the white noise and measured noise model in temperature (top) and polarization (bottom), estimated by differencing the auto and cross-spectra of the half-mission maps. This is a key quantity in the covariance matrix calculation. The noise model parameters estimated for each 100, 143 and 217 GHz half-mission map are reported in Table 4.1. The white noise levels $N_{\ell,\mathrm{data}}$ are presented in Table 4.1, and are constant with ℓ. The temperature noise spectra (top panel) are poorly determined for $\ell \lesssim 500$. However, the covariance matrices are not sensitive to the noise power spectrum at these signal-dominated scales.

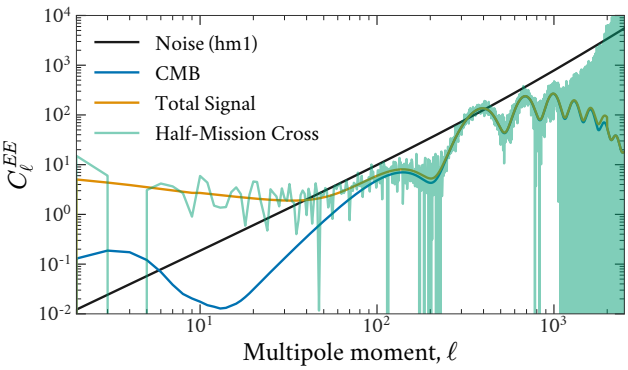

Figure 7. The EE noise model (black) for half-mission hm1 at 143 GHz, compared to the total signal (the sum of theory and foreground components) and the estimated half-mission cross-spectrum.

our analysis, which uses only the publicly available products, still reproduces the intermediate products and cosmology to high fidelity.

We compute the noise power spectrum for channel $X \in \{TT, EE\}$ as the difference between the mask-deconvolved auto and cross-spectra,

$$\begin{aligned} N_\ell^{X,\mathrm{hm1}\times\mathrm{hm1}} &= C_\ell^{X,\mathrm{hm1}\times\mathrm{hm1}} - C_\ell^{X,\mathrm{hm1}\times\mathrm{hm2}}, \\ N_\ell^{X,\mathrm{hm2}\times\mathrm{hm2}} &= C_\ell^{X,\mathrm{hm2}\times\mathrm{hm2}} - C_\ell^{X,\mathrm{hm1}\times\mathrm{hm2}}. \end{aligned} \tag{13}$$

We note that the noise in our convention is not modified by the beam. Since the half-mission maps have very slightly different beams (their difference is $\sim 10^{-4}$ the value of the beam transfer function), we compute these half-mission differences by correcting for the beam within the auto- and cross-spectra, differencing the spec-

freq	split	TT $[10^{-3}\,\mu\mathrm{K}^2]$	EE $[10^{-3}\,\mu\mathrm{K}^2]$
100	1	1.06583	2.47488
100	2	1.01859	2.38146
143	1	0.208448	0.965957
143	2	0.194149	0.893841
217	1	0.409548	2.16995
217	2	0.374725	1.91914

Table 3. The amplitude of the white noise power spectrum N_ℓ estimated for the *Planck* frequency maps. We include more significant figures than necessary to aid in reproducibility.

tra, and then re-applying the beam appropriately for that half-mission. We also check that we reproduce the qualitative differences reported in PL20 between noise power spectra estimated from either half-mission or half-ring maps.

In order to estimate the parameters of the noise model in Equation 12, we optimize a Gaussian likelihood with a diagonal covariance matrix that simply weights modes as $N_\ell/\sqrt{2\ell+1}$ when obtaining the noise model fit. This likelihood is then optimized through an adaptive differential evolution optimizer (Wang et al. 2014) implemented in `BlackBoxOptim.jl` (Feldt & Stukalov 2018). We report the noise model parameters for each of the *Planck* half-mission frequency maps in Table 4.1, and plot the resulting noise curve ratios in Figure 6. The total noise for EE at 143 GHz is shown in Figure 7, where we also show various contributions to the signal spectrum in comparison to the noise. We show the estimated noise spectra and our smooth fits to those spectra in Figure 17 and Figure 18.

The *Planck* noise power spectra are challenging to determine, even in polarization. The *Planck* polarization maps are individually noise-dominated at most scales, allowing for a more accurate determination of the noise power spectrum. In contrast, the temperature maps contain substantial signal, so the estimation of the noise power spectrum is more challenging. There is therefore uncertainty in the large-scale noise power spectra shown in Figure 6, that is difficult to characterize using the available public data products from the *Planck* collaboration. We note however that, at these scales, the noise contribution to the covariance matrix is subdominant and thus does not affect cosmological constraints; these noise curves are therefore a sufficient intermediate product for the power spectrum estimation. We show this behavior in Figure 16, where we plot the signal-to-noise ratios from some of the *Planck* frequency maps.

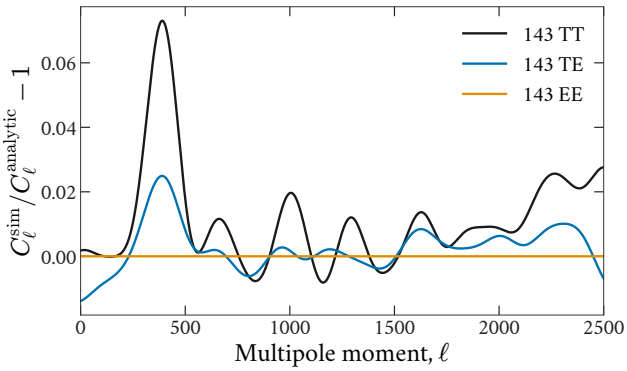

Figure 8. The ratio of the covariance estimated from signal-only simulations compared to the analytic estimate. This ratio is non-unity due to the presence of sharply-masked point sources in the simulations. This correction ratio is applied to the fiducial power spectrum used in the data covariance matrix.

4.2. *Point source treatment*

Point sources violate the approximation in the covariance matrix that the spectra of the masks are declining rapidly in power towards small scales. The apodization around each point source required to satisfy this assumption would result in a substantial reduction in signal. To accommodate the low apodization, we instead follow PL20 by computing signal-only simulations of each cross-spectrum used in the analysis. Using only the fiducial signal power spectrum C_ℓ, we obtain a correction ratio by dividing the predicted analytic covariance diagonal with that from the simulations. This correction ratio is applied to the fiducial power spectrum in the covariance estimation. This is the same approach as in PL20, but we generate our own simulations for this purpose.

In Figure 8 we show the correction ratio for each multipole estimated using this method at 143 GHz, based on a smooth interpolation of the ratio estimates from 1000 simulations using a Gaussian Process with an exponential kernel. PL20 used a spline fit, but the two different non-parametric models should produce nearly identical results. We distribute these correction ratios and code to generate these correction factors within `PSPipe`.

4.3. *Comparison to PL20 covariances*

A comparison of the *Planck* PL20 and our covariance matrix for the 100 GHz TT and E-mode covariance subblocks are shown in Figures 9 and 10. The length of the data vector in both cases is the number of multipole bins at this frequency, with $\ell_{\rm max} = 1250$ at 100 GHz. The middle panels show the diagonal elements of the matrices as a function of multipole, with the ratio of the

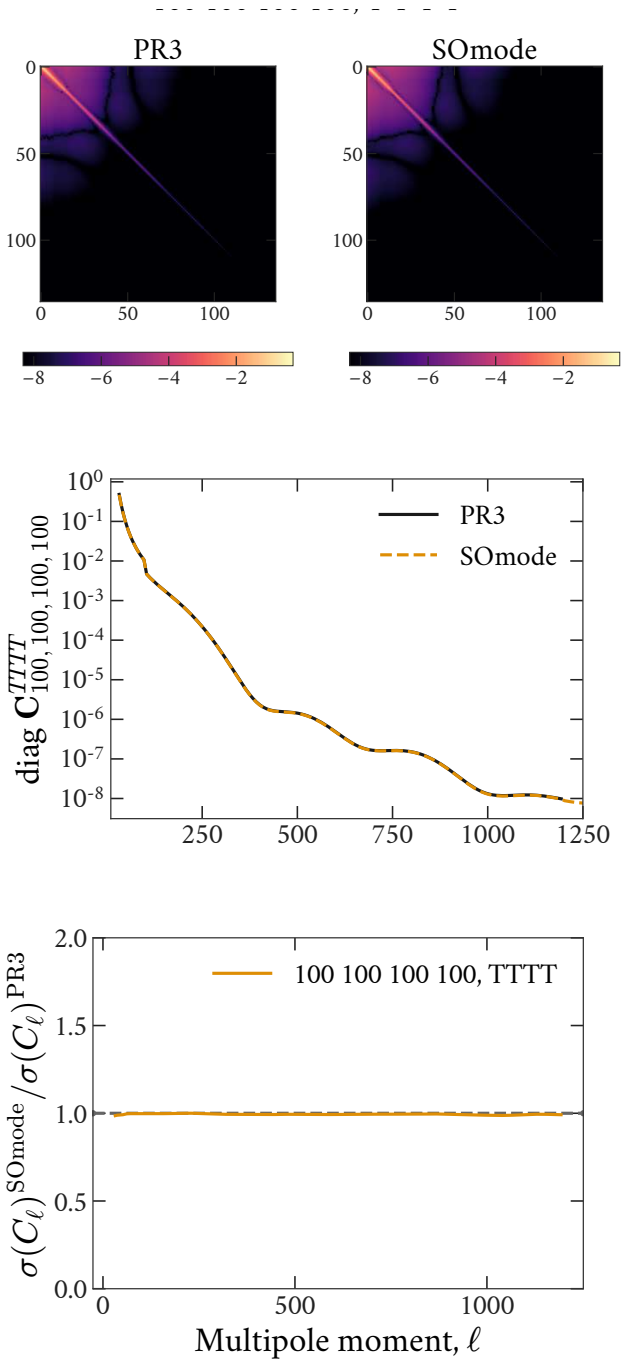

Figure 9. A comparison of the temperature covariance sub-blocks for cross-spectra between the half-missions. We note the excellent agreement in both the diagonal and off-diagonal structure. In the heatmaps above, we show the logarithm of the absolute value of the covariance matrix elements.

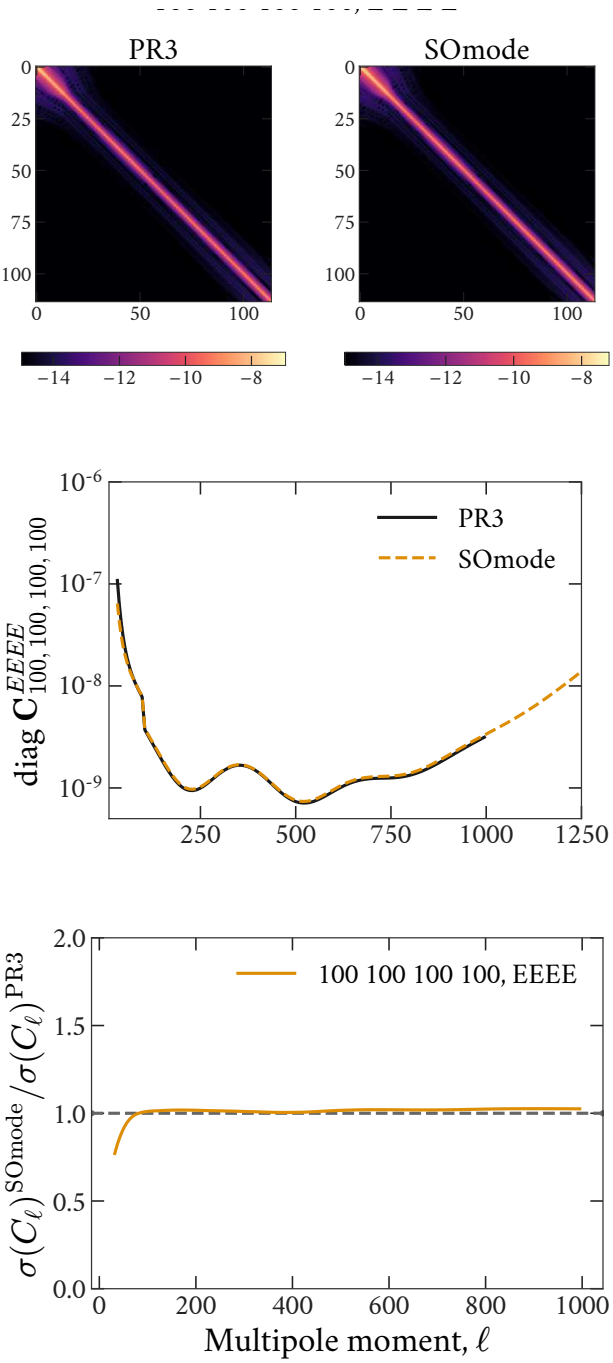

Figure 10. A comparison of the E-mode square root of the covariance sub-block diagonals for cross-spectra between the half-missions. Our best candidate for the origin of this low-ℓ discrepancy is the noise power spectrum estimate. We used the difference of the auto- and cross-spectra from the half-mission frequency maps, rather than the half-mission half-ring maps used in the PL20 analysis.

two estimates in the lower panels. The *TTTT* covariance matrices agree on the diagonal to better than 6% for all frequencies. The *EEEE* covariance also shows good agreement, with a deviation at the largest scales; our method has a lower estimate at $\ell < 100$ by $\sim 30\%$. The off-diagonal structure also shows consistent behav-

ior. The full covariance matrix, and comparison of the diagonals for the full data vector, is shown in Figure 11.

5. LIKELIHOOD

The Simons Observatory collaboration have also developed likelihood software to return the likelihood of the power spectra given some theoretical model, and we use it here to analyse these new inputs.

The modeling of the *Planck* high-ℓ multi-frequency power spectra is described in detail in PL20 and Planck Collaboration XV (2014); Planck Collaboration XI (2016b), with PL20 presenting the modeling of the PR3 spectra and differences/improvements with respect to the previous releases. Here, our goal is to test the re-estimated spectra and covariance matrices, so we use the same modeling assumptions used in the PL20 reference likelihood `Plik`, as summarized below. We have implemented this in `MFLikePlik`[5], a branch of the SO multi-frequency likelihood for the power spectrum analysis of the Large Aperture Telescope `MFLike`[6].

The data vector of the spectra, $\mathbf{C_d}$, and its covariance matrix $\mathbf{\Sigma}$ from Sections 3 and 4 are used to form a Gaussian likelihood:

$$-2\ln\mathcal{L}(\mathbf{C_d}|\mathbf{C_t}(\theta)) = [\mathbf{C_d}-\mathbf{C_t}(\theta)]^T\,\mathbf{\Sigma}^{-1}\,[\mathbf{C_d}-\mathbf{C_t}(\theta)] + \text{const.}\,, \tag{14}$$

where $\mathbf{C_t}(\theta)$ is the theoretical prediction as function of cosmological, foreground and nuisance parameters θ. For each cross-frequency binned spectrum (between frequencies i and j) in temperature or polarization, this can be written as:

$$\mathbf{C_t^{ij}}(\theta) = \mathbf{C_t^{CMB}}(\theta_1) + \mathbf{C_t^{sec,ij}}(\theta_2) + \mathbf{C_t^{sys,ij}}, \tag{15}$$

where $\mathbf{C_t^{CMB}}$ is a binned CMB theory prediction for a given set of cosmological parameters θ_1, and the other two terms described in more detail below account for frequency-dependent foregrounds and residual systematics, respectively. The spectra in PL20 are only binned at the last step of processing, resulting in a top-hat bandpower window function, followed by a $\ell(\ell+1)$ weighting within each bin. This binning operation is performed on the sum of the CMB theory and foreground components, to compute the likelihood.

Foregrounds: $\mathbf{C_t^{sec,ij}}(\theta_2)$ is the secondary signal from Galactic and extragalactic emission and depends on 20 foreground parameters θ_2. In temperature this term includes thermal and kinetic Sunyaev-Zel'dovich (SZ) effects, dusty star-forming and radio galaxies appearing as point sources, a clustering term for the cosmic infrared background (CIB), a correlation between SZ and CIB, and thermal dust emission from our Galaxy. In polarization the Galactic dust emission is modeled for both TE and EE. We refer the reader to PL20 for the modeling of all these components in the likelihood, and in particular to Table 16 of PL20 for the definition of the θ_2 parameters, their range of variation and the priors imposed on them. For our implementation of these components in `MFLikePlik` we adopt `fgspectra`[7], an SO library that builds cross-spectra predictions for foreground components at different frequencies for a given set of model parameters.

Systematics: $\mathbf{C_t^{sys,ij}}$ is a term correcting the model for residual levels of systematic effects. A module in `MFLikePlik` includes the PL20 frequency-dependent templates for beam leakage, sub-pixel noise corrections, and correlated noise corrections to the model vector with fixed amplitudes. We do not re-estimate these templates in this analysis.

Calibration and polarization efficiencies: The full theory vector is then calibrated and corrected for polarization efficiencies before comparing to the data. As in PL20, the model is calibrated assuming a fixed temperature calibration equal to 1 at 143 GHz, and polarization efficiencies also fixed to 1.021, 0.966, and 1.04 for 100, 143, and 217 GHz, respectively. The 100 and 217 GHz temperature calibration factors are added as nuisance parameters as a third component of the θ vector and varied with the same Gaussian priors used in PL20.

All the data products including the new spectra and covariance inputs presented here, the binning and weighting schemes, and the PL20 templates for systematic effects, are stored as plain text files in the `MFLikePlik` data folder. Additional templates needed to model foregrounds are stored in the `fgspectra` library.

6. TESTS OF PARAMETERS

Our re-estimated spectra and covariance products, which we will term `SOmode`, have been directly compared to the *Planck* PR3 products in previous sections. Here we look at the last step of the cosmological pipeline and test how cosmological parameters derived from our new inputs compare with those presented in PC20 for the basic ΛCDM model as well as a selection of extensions.

To do this we couple our likelihood to the publicly-available `Cobaya`[8] code (Torrado & Lewis 2021). We extend the `Cobaya` *Planck* yaml files to call `MFLikePlik`

[5] https://github.com/simonsobs/LAT_MFLike/tree/mflike-plik

[6] https://github.com/simonsobs/LAT_MFLike

[7] https://github.com/simonsobs/fgspectra

[8] https://cobaya.readthedocs.io/en/latest/

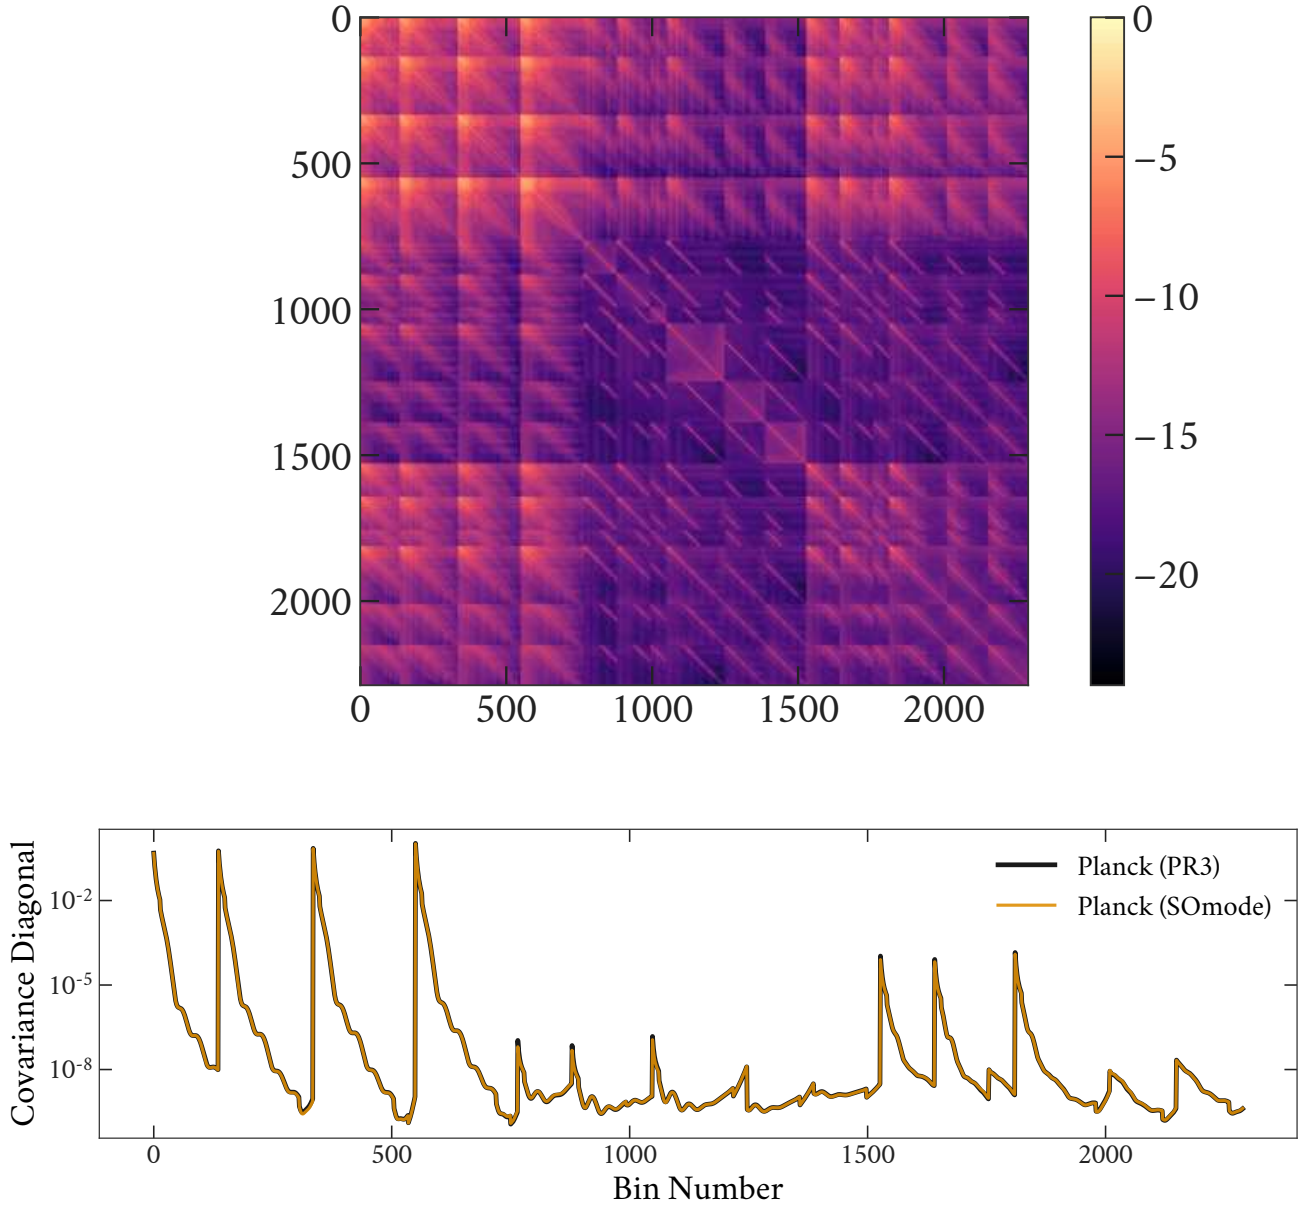

Figure 11. (Top) The full covariance matrix for the *Planck* 2018 spectra, with (bottom) a comparison of the covariance matrix diagonals for *Planck* and our estimate. The data vector has 2289 elements and is ordered overall by TT, EE, TE. The ordering by frequency is given in Table 3.2.

for the high-ℓ spectra and we also include the PR3 likelihood implementations, wrapped in `Cobaya`, of the *Planck* 'lowl' and 'lowE' likelihoods modeling the large-scale temperature and EE polarization data.

We find excellent agreement for the six basic ΛCDM parameters ($\Omega_b h^2$, $\Omega_c h^2$, A_s, n_s, $\theta_{\rm MC}$, τ) between our inputs and PC20, as shown in Figure 12, with all parameters agreeing to within 0.1σ. The distributions of the foreground parameters are shown in Figure 13. In this case, all the parameters are consistent with those reported in PC20 within 0.1σ, except for the amplitudes of thermal dust in TE at frequencies higher than 100 GHz. These exhibit some small shifts, at the level of $0.2-0.9\sigma$. As in PC20, the dust amplitudes are varied in TT and TE with a Gaussian prior imposed, and are fixed in EE. We attribute the small shifts in TE parameters to the small differences in the EE spectrum at low multipoles (see Sec. 3.1) where the dust contamination is

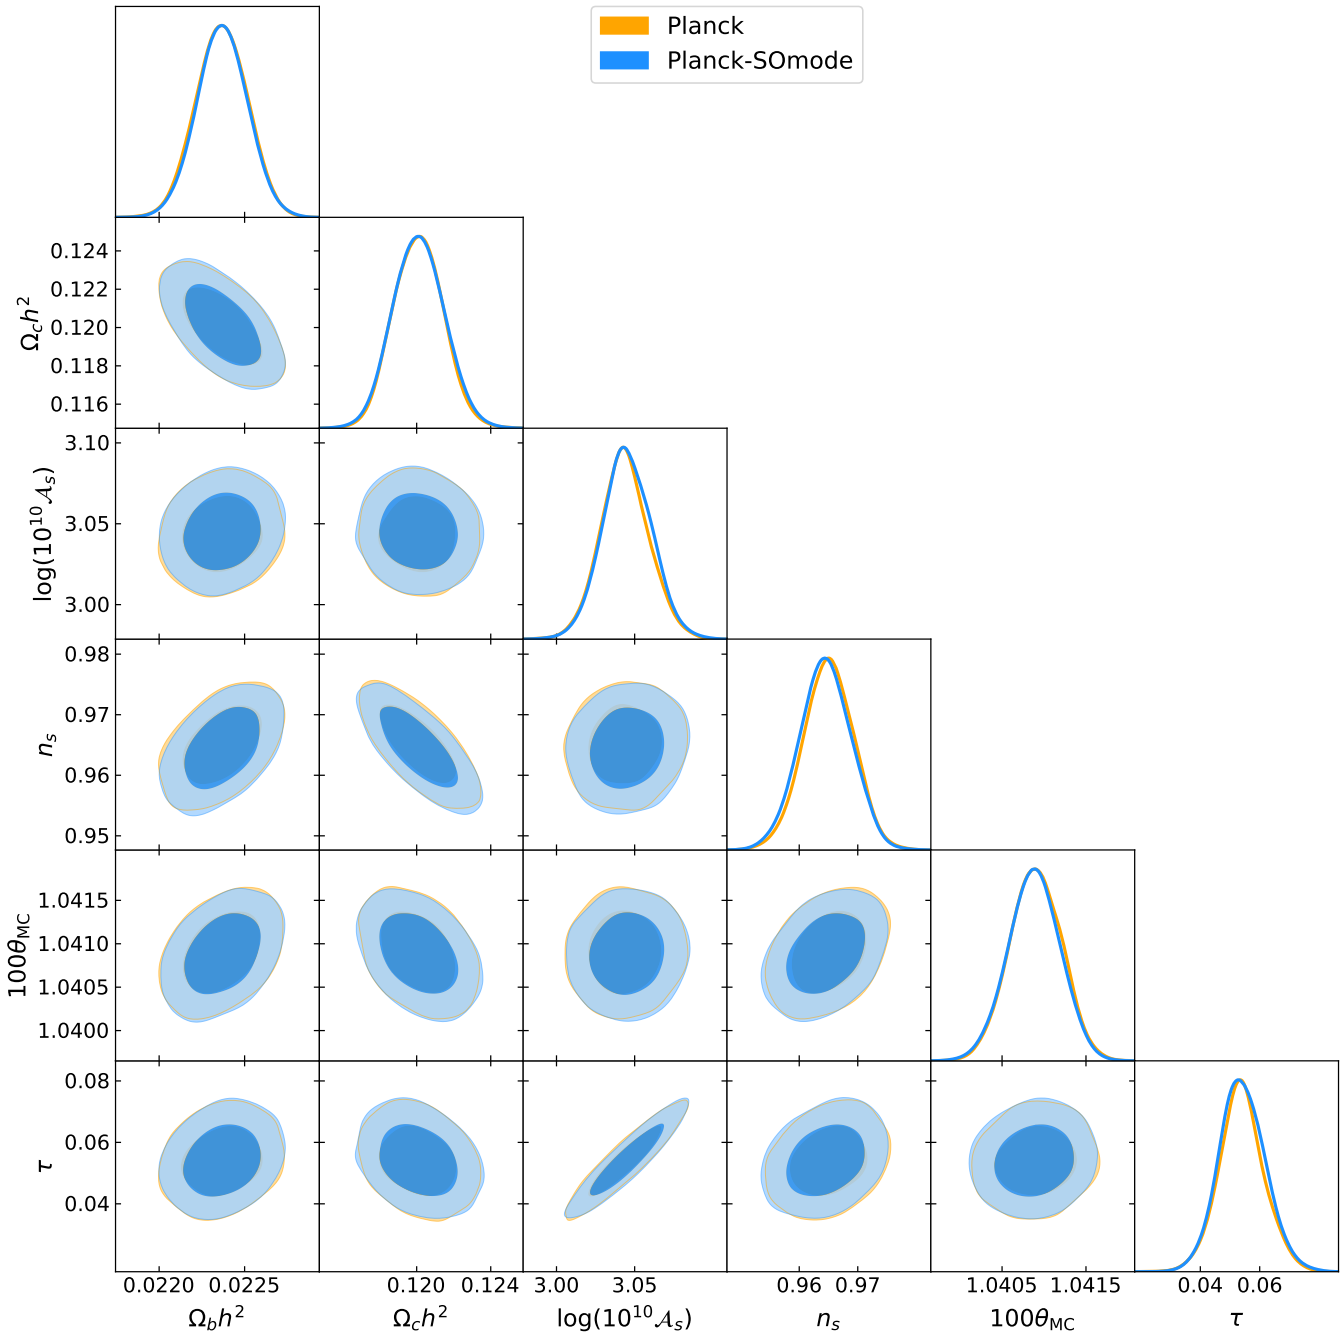

Figure 12. Constraints on the ΛCDM cosmological parameters derived from the spectra and covariance matrix estimated in this work (blue), compared to the *Planck* legacy results (orange). The diagonal panels show the 1-dimensional posterior distributions; the contour plots show the 2-dimensional 68% and 95% confidence levels. The two pipelines give constraints that agree to within 0.1σ.

significant, at 143 and 217 GHz. More precisely, since the dust amplitudes are fixed in EE, any difference at the EE spectrum level will be leveraged by the free-to-vary TE parameters. We tested that these shifts were indeed coming from spectrum differences, rather than differences in the covariance matrix, by estimating parameters using the PR3 spectra together with our new `SOmode` covariance matrix. This combination resulted in no parameter shifts in the foreground parameters.

The small differences in the covariance matrix described in Sec. 4 do manifest as a small difference in χ^2. For ΛCDM, the *Planck* best-fit χ^2 (as per the .likestats Getdist file) for `plikTTTEEE+lowl+lowE` is 2768.9, while for our combination of `MFLikePlik+lowl+lowE` we find 2698.7. This corresponds to a $\Delta\chi^2 = 70.2$ which is dominated by the `plik` χ^2 ($\Delta\chi^2_{\texttt{plik}} = 71.9$). Finding a lower χ^2 is consistent with the slightly increased noise and covariance elements we found with our products at $\ell > 400$. We confirmed this computing χ^2 for subsets of the data, the $\ell > 800$ range, and TE in particular, have the larger impact on the difference in χ^2. The difference in the TE spectra is dominated by the choice of computing the flat average of the various TE and ET cross-half-mission spectra that combine for the final result, rather than inverse-variance weighting.

We also explore a standard set of ΛCDM extensions to validate our products on multiple theories/parametrizations. We estimate constraints on the effective number of relativistic species, $N_{\rm eff}$, the running of the spectral index, $n_{\rm run}$, the amplitude of lensing smoothing in the power spectra, A_L, and spatial curvature, Ω_k. In all cases, we recover the PC20 results to within 0.1σ as shown in Figure 14.

7. CONCLUSION

The *Planck* legacy dataset will be a critical part of cosmological constraints for at least the next decade. We expect exciting new data from existing and upcoming ground-based experiments to shed light on the unknown features of the cosmological model. These data will have to be combined with *Planck* measurements of the large scale temperature modes in order to anchor parameters such as A_s, n_s, and τ, in addition to *Planck*'s clean measurements of the first few acoustic peaks.

Drawing upon wisdom accrued from a decade of *Planck* analysis, we have reproduced some of the main *Planck* cosmology products – the CMB power spectrum and associated covariances, tested with the cosmological likelihood. This paper presents a complete, validated, and publicly available spectrum analysis, and we also provide all associated intermediary data products and software tools. We have reproduced the spectra to extremely high fidelity using the same tools we expect to apply to new ground-based data, and we construct covariance matrix estimates that lead to a remarkably consistent reproduction of the main *Planck* cosmology. We expect these to be broadly useful for making robust cosmological claims from the ground, and advancement past the era of *Planck*.

In this paper, we have provided a detailed and complete description of the scientific choices and approximations used in this pipeline, but we expect the code and scripts to provide substantial additional value for ongoing and future power spectrum analysis. However, future analysis will need to perform various necessary extensions to our work in order to adapt this pipeline to the particulars of their dataset.

- We anticipate that spectra and covariances will require re-computation of the *Planck* spectra and covariances on new survey masks. We try to make substitution of masks in our pipeline straightforward.
- The azimuthally-averaged *Planck* beam varies based on sky patch due to the inhomoegenous scan pattern. We expect the future power spectrum analyst to also require new beams computed using `QuickPol` for *Planck* on their particular sky patch.
- The foreground model will change in future analysis, and that model is an input for the covariance matrix calculation as well as used in the likelihood.

We aim to make these tasks straightforward within the framework of `PSPipe` and `MFLike`, although re-estimating the Galactic contributions as inputs for the likelihood is beyond the scope of this current study.

A secondary aim of our analysis has been to provide an independent check of the *Planck* analysis. The excellent agreement we find with the primary *Planck* power spectrum analysis adds confidence in the implementation. Although the pipeline we present in this paper has hewed closely to the PR3 release, we did experiment with choices such as methods for monopole/dipole subtraction, for masking missing pixels, using `PolSpice` versus other pseudo-C_ℓ codes, choices of noise modeling, and the choice of approximations in the estimation of the covariance matrix, all of which have order $\sim 1-10\%$ effect on the resulting products. The *Planck* dataset is complex, and details such as the computation of noise power spectra can have large effects on outputs like null tests. For example, the χ^2 depends quadratically on the noise power spectrum through the $\mathcal{R}_\ell$ terms.

We note that `NPIPE`, an alternative analysis of *Planck* time-ordered data, is now publicly available. We are in

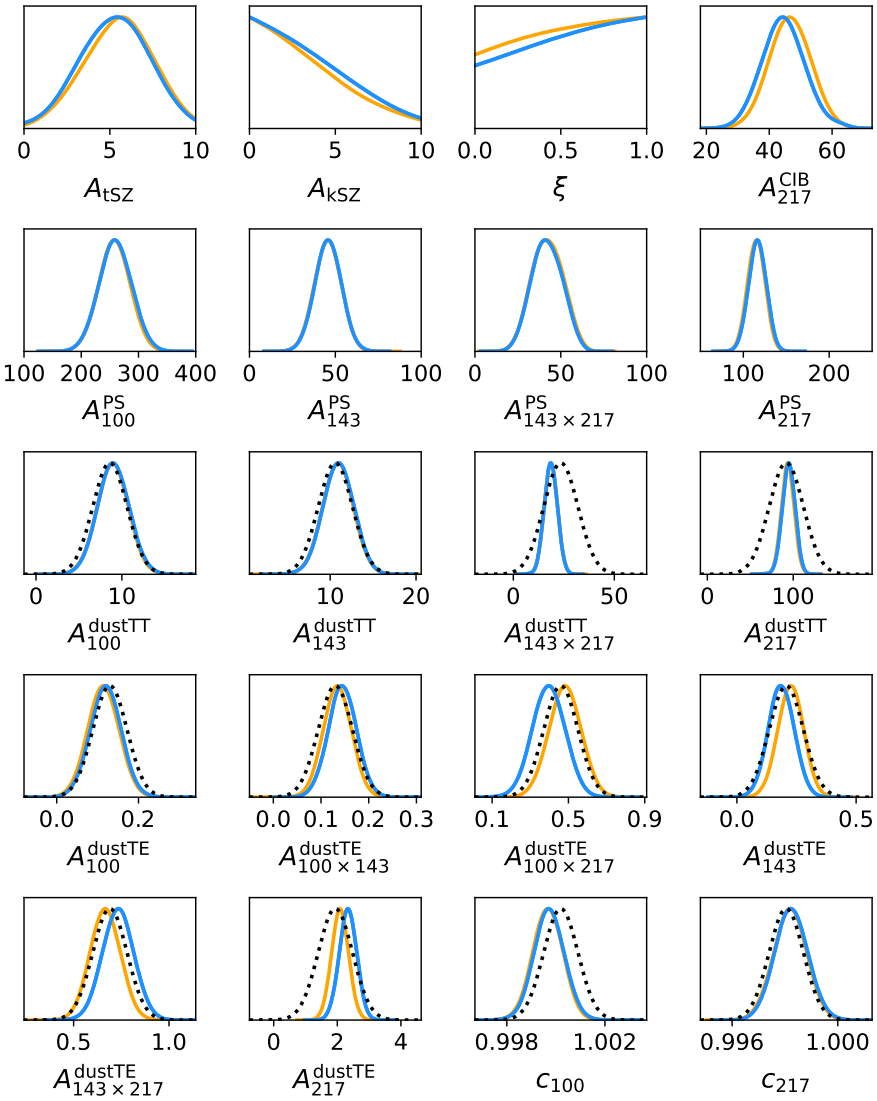

Figure 13. Foreground and calibration parameters derived from the spectra and covariances estimated in this work (blue), compared the *Planck* legacy results (orange) – from left to right and top to bottom these show: thermal and kinetic SZ amplitudes, the correlation between thermal SZ and CIB, the CIB amplitude, four point source amplitudes, four dust amplitudes in TT and six in TE, and two calibrations (see PL20 for more details). The dotted black lines show the Gaussian priors imposed on individual parameters. An additional joint prior is used for the thermal and kinetic SZ amplitudes and not shown here. The two pipelines give consistent results, with small shifts only in some of the TE dust amplitudes. This is attributed to small differences in the EE spectra at low multipoles.

the process of adapting our pipeline to these data, and leave a detailed analysis for future work. Future work will also incorporate adaptations for using the *Planck* data with current and future ground-based surveys.

8. ACKNOWLEDGMENTS

We are grateful to George Efstathiou and Steven Gratton for sharing useful information about the *Planck* analysis. We also thank Antony Lewis for helpful feedback. This work was performed using Princeton Research Computing resources, a consortium of groups including the Princeton Institute for Computational Science and Engineering (PICSciE) and the Office of Information Technology's High Performance Computing Center and Visualization Laboratory at Princeton University, as well as the Hawk high-performance computing cluster at the Advanced Research Computing at Cardiff (ARCCA). Research in Canada is supported by NSERC and CIFAR. JD gratefully acknowledges support from the Institute for Advanced

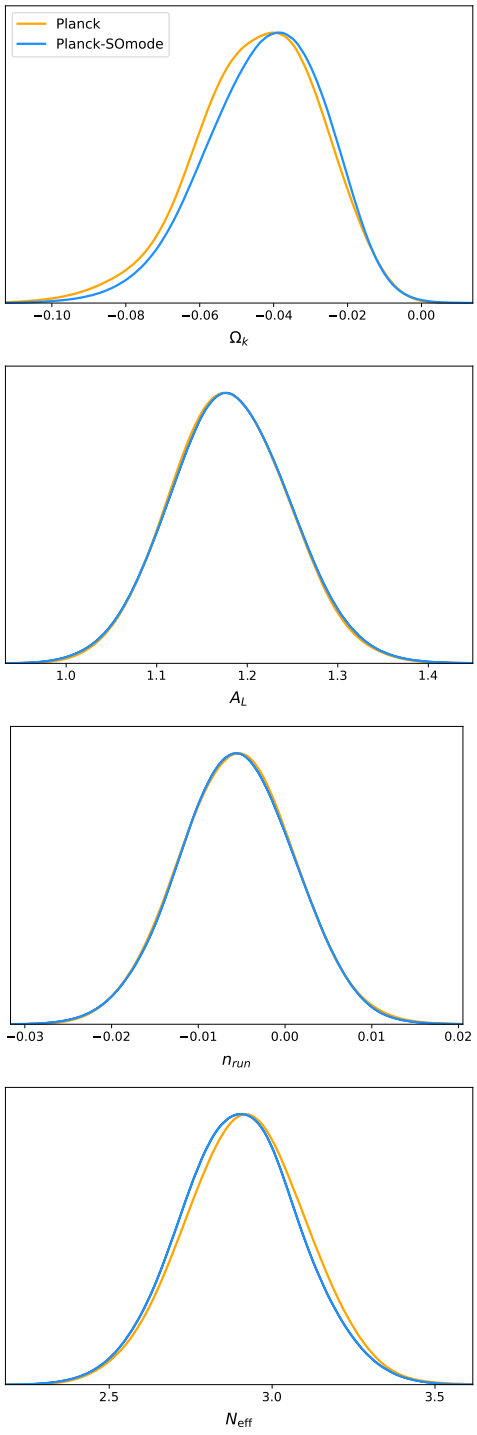

Figure 14. Posterior distributions for four single-parameter extensions to the ΛCDM model, derived using the spectra and covariance matrix estimated in this work (blue), compared to the *Planck* legacy results (orange). Only the extension parameter is shown for each model.

Study and from NSF grant AST-2108126. DA is supported by the Science and Technology Facilities Council through an Ernest Rutherford Fellowship, grant reference ST/P004474. EC acknowledges support from the STFC Ernest Rutherford Fellowship ST/M004856/2 and STFC Consolidated Grant ST/S00033X/1; EC, HJ and UN acknowledge additional support from the Horizon 2020 ERC Starting Grant (Grant agreement No 849169). GF acknowledges the support of the European Research Council under the Marie Skłodowska Curie actions through the Individual Global Fellowship No. 892401 PiCOGAMBAS. SKC acknowledges support from NSF award AST-2001866. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Industry Canada and by the Province of Ontario through the Ministry of Colleges and Universities. P.D.M acknowledges support from the Netherlands organization for scientific research (NWO) VIDI grant (dossier 639.042.730). Some of the results in this paper have been derived using the HEALPix (Górski et al. 2005), healpy (Zonca et al. 2019), and Healpix.jl (Tomasi & Li 2021) packages. We also acknowledge use of the matplotlib (Hunter 2007) package for producing plots in this paper.

REFERENCES

Abazajian, K., Addison, G., Adshead, P., et al. 2019, arXiv e-prints, arXiv:1907.04473. https://arxiv.org/abs/1907.04473

Aiola, S., Calabrese, E., Maurin, L., et al. 2020, JCAP, 2020, 047, doi: 10.1088/1475-7516/2020/12/047

Alonso, D., Sanchez, J., Slosar, A., & LSST Dark Energy Science Collaboration. 2019, MNRAS, 484, 4127, doi: 10.1093/mnras/stz093

Benson, B. A., Ade, P. A. R., Ahmed, Z., et al. 2014, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9153, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII, ed. W. S. Holland & J. Zmuidzinas, 91531P, doi: 10.1117/12.2057305

Brown, M. L., Castro, P. G., & Taylor, A. N. 2005, MNRAS, 360, 1262, doi: 10.1111/j.1365-2966.2005.09111.x

Challinor, A., & Chon, G. 2005, MNRAS, 360, 509, doi: 10.1111/j.1365-2966.2005.09076.x

Choi, S. K., Hasselfield, M., Ho, S.-P. P., et al. 2020, JCAP, 2020, 045, doi: 10.1088/1475-7516/2020/12/045

Chon, G., Challinor, A., Prunet, S., Hivon, E., & Szapudi, I. 2004, MNRAS, 350, 914, doi: 10.1111/j.1365-2966.2004.07737.x

De Bernardis, F., Stevens, J. R., Hasselfield, M., et al. 2016, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9910, Observatory Operations: Strategies, Processes, and Systems VI, ed. A. B. Peck, R. L. Seaman, & C. R. Benn, 991014, doi: 10.1117/12.2232824

Dutcher, D., Balkenhol, L., Ade, P. A. R., et al. 2021, arXiv e-prints, arXiv:2101.01684. https://arxiv.org/abs/2101.01684

Efstathiou, G. 2004, MNRAS, 349, 603, doi: 10.1111/j.1365-2966.2004.07530.x

—. 2006, MNRAS, 370, 343, doi: 10.1111/j.1365-2966.2006.10486.x

Efstathiou, G., & Gratton, S. 2019, arXiv e-prints, arXiv:1910.00483. https://arxiv.org/abs/1910.00483

Feldt, R., & Stukalov, A. 2018, BlackBoxOptim.jl, 0.6.1, https://github.com/robertfeldt/BlackBoxOptim.jl, GitHub

García-García, C., Alonso, D., & Bellini, E. 2019, JCAP, 2019, 043, doi: 10.1088/1475-7516/2019/11/043

Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759, doi: 10.1086/427976

Hansen, F. K., Górski, K. M., & Hivon, E. 2002, MNRAS, 336, 1304, doi: 10.1046/j.1365-8711.2002.05878.x

Hazumi, M., Ade, P. A. R., Adler, A., et al. 2020, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 11443, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 114432F, doi: 10.1117/12.2563050

Hinshaw, G., Spergel, D. N., Verde, L., et al. 2003, ApJS, 148, 135, doi: 10.1086/377225

Hivon, E., Górski, K. M., Netterfield, C. B., et al. 2002, ApJ, 567, 2, doi: 10.1086/338126

Hivon, E., Mottet, S., & Ponthieu, N. 2017a, A&A, 598, A25, doi: 10.1051/0004-6361/201629626

—. 2017b, A&A, 598, A25, doi: 10.1051/0004-6361/201629626

Hunter, J. D. 2007, Computing in Science & Engineering, 9, 90, doi: 10.1109/MCSE.2007.55

Louis, T., Naess, S., Garrido, X., & Challinor, A. 2020, PhRvD, 102, 123538, doi: 10.1103/PhysRevD.102.123538

Naess, S., Aiola, S., Austermann, J. E., et al. 2020, JCAP, 2020, 046, doi: 10.1088/1475-7516/2020/12/046

Planck Collaboration LVII. 2020, A&A, 643, A42, doi: 10.1051/0004-6361/202038073

Planck Collaboration V. 2020, A&A, 641, A5, doi: 10.1051/0004-6361/201936386

Planck Collaboration VI. 2020, A&A, 641, A6, doi: 10.1051/0004-6361/201833910

Planck Collaboration XI. 2016a, A&A, 594, A11, doi: 10.1051/0004-6361/201526926

—. 2016b, A&A, 594, A11, doi: 10.1051/0004-6361/201526926

Planck Collaboration XV. 2014, A&A, 571, A15, doi: 10.1051/0004-6361/201321573

Simons Observatory Collaboration. 2019, Journal of Cosmology and Astro-Particle Physics, 2019, 056, doi: 10.1088/1475-7516/2019/02/056

Thornton, R. J., Ade, P. A. R., Aiola, S., et al. 2016, The Astrophysical Journal Supplement Series, 227, 21, doi: 10.3847/1538-4365/227/2/21

Tomasi, M., & Li, Z. 2021, Healpix.jl: Julia-only port of the HEALPix library. http://ascl.net/2109.028

Torrado, J., & Lewis, A. 2021, JCAP, 2021, 057, doi: 10.1088/1475-7516/2021/05/057

Wang, Y., Li, H.-X., Huang, T., & Li, L. 2014, Appl. Soft Comput., 18, 232–247, doi: 10.1016/j.asoc.2014.01.038

Zonca, A., Singer, L., Lenz, D., et al. 2019, Journal of Open Source Software, 4, 1298, doi: 10.21105/joss.01298

APPENDIX

A. SPECTRUM RESIDUALS

We show the full suite of residuals of our spectra compared to the PL20 public spectra in Figure 15.

B. COVARIANCE MATRIX FUNCTIONS

Following PL16 and PL20 we use index names

$$\begin{aligned} \alpha, \beta &\in \{TT, TP, PT, PP\} \\ X, Y &\in \{\emptyset\emptyset, TT, PP\}. \end{aligned} \tag{B1}$$

Here, indices α, β refer to the kind of cross-spectrum (i.e., cross-spectra between spin 0 and 2, spin 2 and 2, etc.). In contrast, X, Y refer to the kind of noise-variance weighting applied to the mask. $X = \emptyset\emptyset$ refers to the mask alone, $X = TT$ refers to weighting by the II variance, and PP refers to weighting by the QQ and UU variance.

PL16 give the projector functions $\Xi^{X,Y}_{AB}\left[(i,j)^\alpha, (p,q)^\beta\right]$, describing mode-coupling from the mask for maps i, j, p, q on channel $AB \in \{TT, TE, EE\}$. They are given in Appendix C1 of PL16 and we repeat them here for completeness:

$$\begin{aligned} \Xi^{X,Y}_{TT}\left[(i,j)^\alpha, (p,q)^\beta\right]_{\ell_1\ell_2} = \sum_{\ell_3} \frac{2\ell_3+1}{4\pi} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix} \\ \times W^{X,Y}\left[(i,j)^\alpha, (p,q)^\beta\right]_{\ell_3}, \end{aligned} \tag{B2}$$

$$\begin{aligned} \Xi^{X,Y}_{TE}\left[(i,j)^\alpha, (p,q)^\beta\right]_{\ell_1\ell_2} = \sum_{\ell_3} \frac{2\ell_3+1}{8\pi} \left(1 + (-1)^{\ell_1+\ell_2+\ell_3}\right) \\ \times \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ -2 & 2 & 0 \end{pmatrix} W^{X,Y}\left[(i,j)^\alpha, (p,q)^\beta\right]_{\ell_3}, \end{aligned} \tag{B3}$$

$$\begin{aligned} \Xi^{X,Y}_{EE}\left[(i,j)^\alpha, (p,q)^\beta\right]_{\ell_1\ell_2} = \sum_{\ell_3} \frac{2\ell_3+1}{16\pi} \left(1 + (-1)^{\ell_1+\ell_2+\ell_3}\right)^2 \\ \times \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ -2 & 2 & 0 \end{pmatrix}^2 W^{X,Y}\left[(i,j)^\alpha, (p,q)^\beta\right]_{\ell_3}. \end{aligned} \tag{B4}$$

The PL20 analysis ignores the effect of the BB spectrum in the covariance matrix estimates; we also state the corresponding BB projector here for completeness, which is also used in Alonso et al. (2019). Ignoring the conversion of EE into BB from coupling is only of order 5% in EE spectra at $\ell < 100$, and this is corrected from the diagonal through simulations, as we discuss later in this section.

$$\begin{aligned} \Xi^{X,Y}_{BB}\left[(i,j)^\alpha, (p,q)^\beta\right]_{\ell_1\ell_2} = \sum_{\ell_3} \frac{2\ell_3+1}{16\pi} \left(1 - (-1)^{\ell_1+\ell_2+\ell_3}\right)^2 \\ \times \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ -2 & 2 & 0 \end{pmatrix}^2 W^{X,Y}\left[(i,j)^\alpha, (p,q)^\beta\right]_{\ell_3}. \end{aligned} \tag{B5}$$

These projector functions depend on the cross-spectra W between the masks, and require a Wigner 3j symbol that comes from integrals over the spherical harmonic basis. Note that the 3j symbols[9] are subject to a selection rule, such that the nonzero terms in these summations are limited to $|\ell_1 - \ell_2| \leq \ell_3 \leq \ell_1 + \ell_2$.

[9] We implement these 3j symbols in a fast and stable set of recurrence relations within the publicly available package `WignerFamilies.jl`.

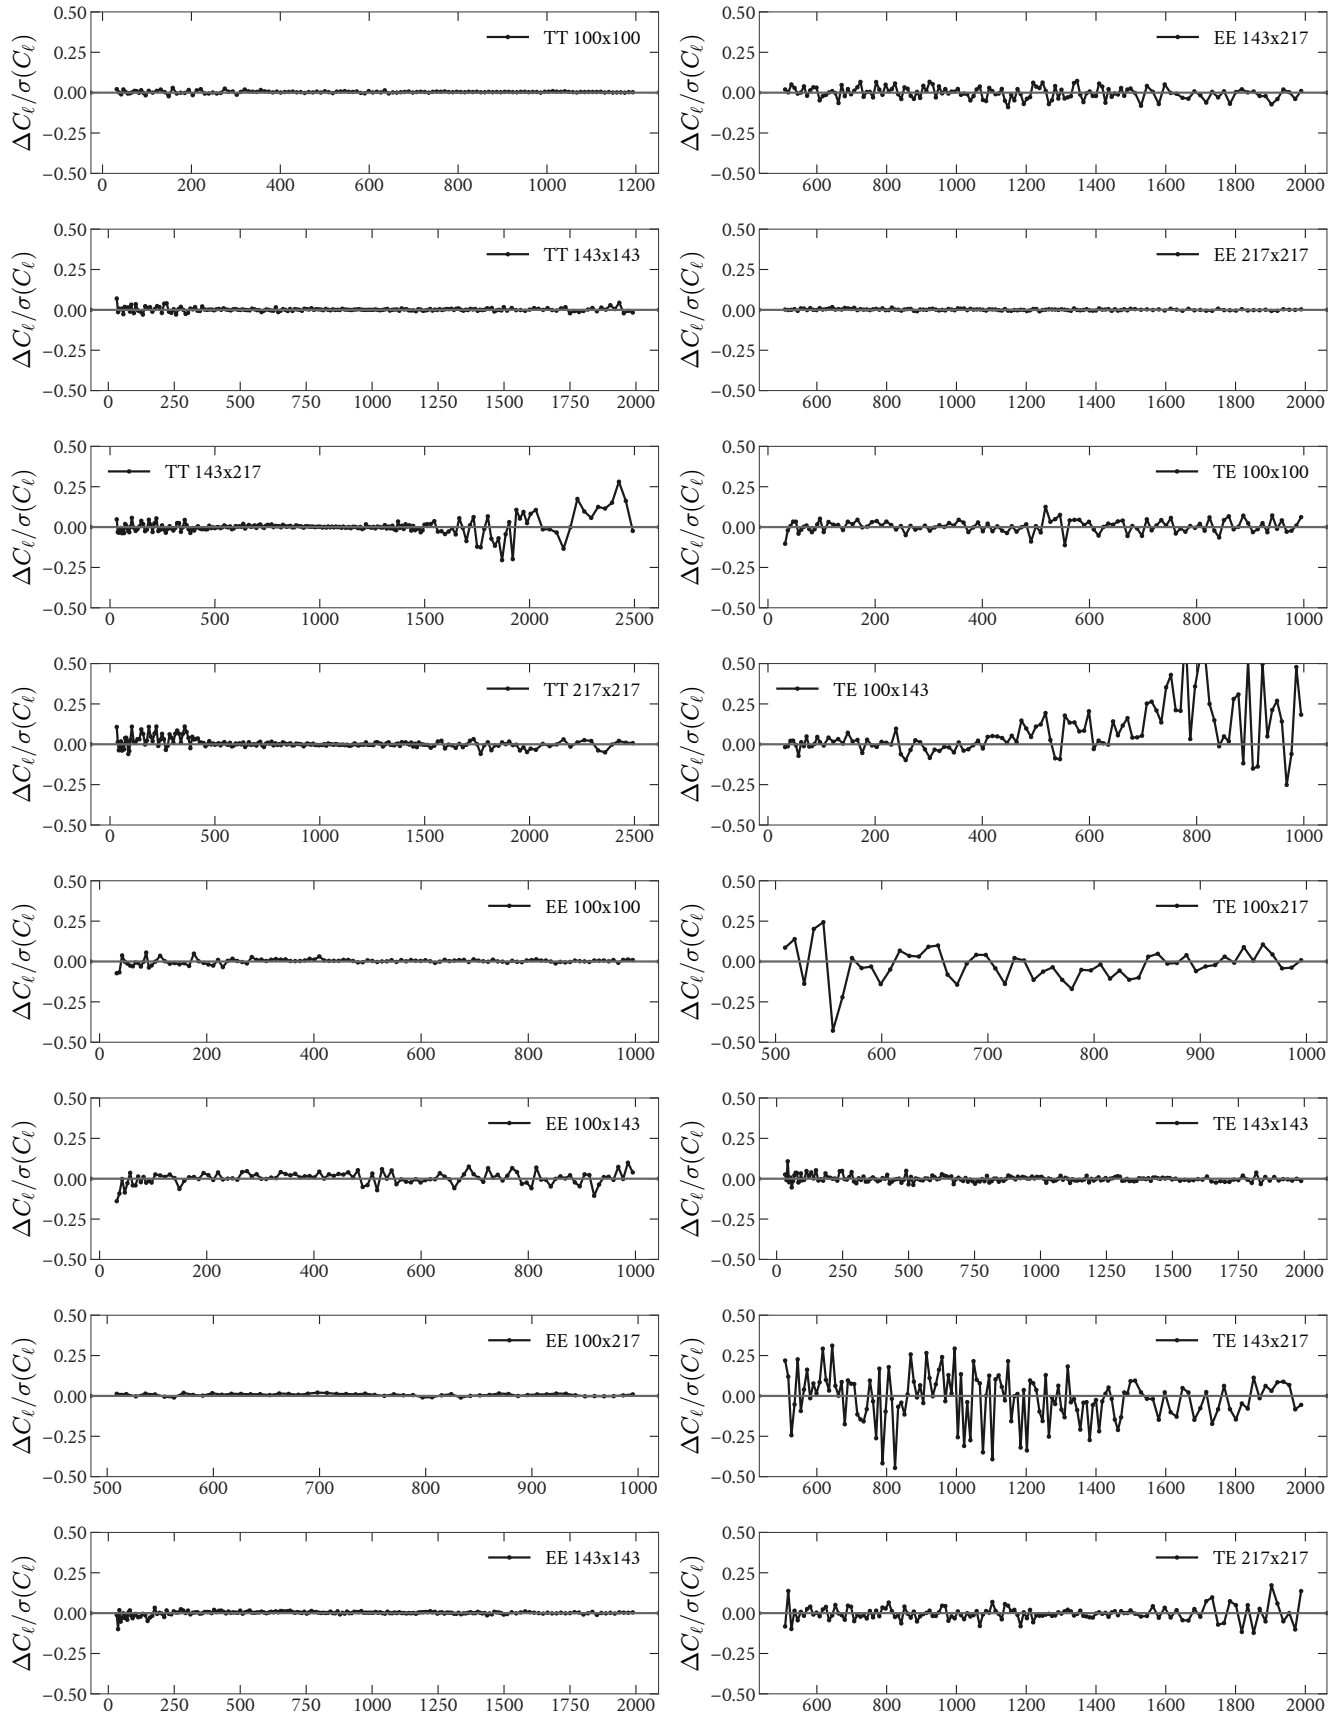

Figure 15. All cross-spectra used in the likelihood, showing the residual of our Planck-SOmode estimates compared to the spectra from the official PR3 analysis.

The window functions $W^{X,Y}\left[(i,j)^\alpha,(p,q)^\beta\right]$ in the projector functions, which are cross-spectra between the masks, are given in Appendix C2 of PL16; we repeat them here.

$$
\begin{aligned}
W^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,j)^\alpha,(p,q)^\beta\right]_\ell &= \frac{1}{2\ell+1}\sum_m w^{\emptyset\emptyset}_{\ell m}(i,j)^\alpha w^{*\emptyset\emptyset}_{\ell m}(p,q)^\beta \\
W^{\emptyset\emptyset,TT}\left[(i,j)^\alpha,(p,q)^{TT}\right]_\ell &= \frac{1}{2\ell+1}\sum_m w^{\emptyset\emptyset}_{\ell m}(i,j)^\alpha w^{*II}_{\ell m}(p,q)^{TT} \\
W^{TT,TT}\left[(i,j)^{TT},(p,q)^{TT}\right]_\ell &= \frac{1}{2\ell+1}\sum_m w^{II}_{\ell m}(i,j)^{TT} w^{*II}_{\ell m}(p,q)^{TT} \\
W^{\emptyset\emptyset,PP}\left[(i,j)^\alpha,(p,q)^{PP}\right]_\ell &= \frac{1}{2\ell+1}\sum_m \\
&\frac{1}{2}\left(w^{\emptyset\emptyset}_{\ell m}(i,j)^\alpha w^{*QQ}_{\ell m}(p,q)^{PP} + w^{\emptyset\emptyset}_{\ell m}(i,j)^\alpha w^{*UU}_{\ell m}(p,q)^{PP}\right) \\
W^{TT,PP}\left[(i,j)^{TT},(p,q)^{PP}\right]_\ell &= \frac{1}{2\ell+1}\sum_m \\
&\frac{1}{2}\left(w^{II}_{\ell m}(i,j)^{TT} w^{*QQ}_{\ell m}(p,q)^{PP} + w^{II}_{\ell m}(i,j)^{TT} w^{*UU}_{\ell m}(p,q)^{PP}\right) \\
W^{PP,PP}\left[(i,j)^{PP},(p,q)^{PP}\right]_\ell &= \frac{1}{2\ell+1}\sum_m \\
&\frac{1}{4}\big(w^{QQ}_{\ell m}(i,j)^{PP} w^{*QQ}_{\ell m}(p,q)^{PP} + w^{UU}_{\ell m}(i,j)^{PP} w^{*UU}_{\ell m}(p,q)^{PP} \\
&+ w^{QQ}_{\ell m}(i,j)^{PP} w^{*UU}_{\ell m}(p,q)^{PP} + w^{UU}_{\ell m}(i,j)^{PP} w^{*QQ}_{\ell m}(p,q)^{PP}\big)
\end{aligned}
\tag{B6}
$$

These $w_{\ell m}$ are spherical harmonic transforms of the effective weight maps. The $w^{\emptyset\emptyset}_{\ell m}(i,j)$ are simply the coefficients of the spherical harmonic transform of the i and j masks.

$$
\begin{aligned}
w^{\emptyset\emptyset}_{\ell m}(i,j)^{TT} &= \sum_i^{N_{\rm pix}} m_p^{i,T} m_p^{j,T} Y^*_{\ell m}(\hat{\mathbf{n}}_p)\Omega_p, \\
w^{\emptyset\emptyset}_{\ell m}(i,j)^{TP} &= \sum_i^{N_{\rm pix}} m_p^{i,T} m_p^{j,P} Y^*_{\ell m}(\hat{\mathbf{n}}_p)\Omega_p, \\
w^{\emptyset\emptyset}_{\ell m}(i,j)^{PT} &= \sum_i^{N_{\rm pix}} m_p^{i,P} m_p^{j,T} Y^*_{\ell m}(\hat{\mathbf{n}}_p)\Omega_p, \\
w^{\emptyset\emptyset}_{\ell m}(i,j)^{PP} &= \sum_i^{N_{\rm pix}} m_p^{i,P} m_p^{j,P} Y^*_{\ell m}(\hat{\mathbf{n}}_p)\Omega_p,
\end{aligned}
\tag{B7}
$$

The $w^{II}_{\ell m}(i,j)$, $w^{QQ}_{\ell m}(i,j)$, and $w^{UU}_{\ell m}(i,j)$ are the coefficients of noise-variance weighted maps from II, QQ, and UU respectively.

$$
\begin{aligned}
w^{II}_{\ell m}(i,j)^{TT} &= \delta_{i,j}\sum_i^{N_{\rm pix}} \left(\sigma_p^{II}\right)^2 m_p^{i,T} m_p^{j,T} Y^*_{\ell m}(\hat{\mathbf{n}}_p)\Omega_p, \\
w^{QQ}_{\ell m}(i,j)^{TT} &= \delta_{i,j}\sum_i^{N_{\rm pix}} \left(\sigma_p^{QQ}\right)^2 m_p^{i,P} m_p^{j,P} Y^*_{\ell m}(\hat{\mathbf{n}}_p)\Omega_p, \\
w^{UU}_{\ell m}(i,j)^{TT} &= \delta_{i,j}\sum_i^{N_{\rm pix}} \left(\sigma_p^{UU}\right)^2 m_p^{i,P} m_p^{j,P} Y^*_{\ell m}(\hat{\mathbf{n}}_p)\Omega_p,
\end{aligned}
\tag{B8}
$$

The σ_p^2 are pixel noise variances, and there is only a noise contribution if the two maps are the same.

The projector functions are linearly related to the mode-coupling matrices used in the pseudo-C$_\ell$ formalism to estimate unbiased spectra in the presence of a mask (Hivon et al. 2002). For a pseudo-spectrum $\tilde{C}_\ell^{AB}$ and decoupled spectrum $\hat{C}_\ell^{AB}$, we have $\tilde{C}_\ell^{TT} = (2\ell_2+1)\Xi_{\ell,\ell'}^{TT}\,\hat{C}_\ell^{TT}$ and $\tilde{C}_\ell^{TE} = (2\ell_2+1)\Xi_{\ell,\ell'}^{TE}\,\hat{C}_\ell^{TE}$. However, the spin 2 fields EE, EB, BE, and BB requires the solution of a linear system,

$$\begin{aligned}\tilde{C}_\ell^{EE} &= (2\ell_2+1)\Xi_{\ell,\ell'}^{EE}\,\hat{C}_\ell^{EE} + (2\ell_2+1)\Xi_{\ell,\ell'}^{BB}\,\hat{C}_\ell^{BB}\\ \tilde{C}_\ell^{BB} &= (2\ell_2+1)\Xi_{\ell,\ell'}^{BB}\,\hat{C}_\ell^{EE} + (2\ell_2+1)\Xi_{\ell,\ell'}^{EE}\,\hat{C}_\ell^{BB}\end{aligned} \tag{B9}$$

Thus we compute the C_ℓ^{BB} pseudo-spectrum, in order to get decoupled spectra. Note that we follow PL20 and do not use the contribution of C_ℓ^{BB} to the covariance, as it is subdominant to the variance induced by cosmic variance and noise, as suggested by the signal-only simulations presented in Section 4.2.

C. COVARIANCE MATRIX BLOCKS

We repeat here the expressions analogous to Equation 6 for TT, TE, and EE. The same $\mathcal{R}_\ell^{i,X}$ factors must be applied as in Section 4.

TTTT block:—

$$\begin{aligned}&\mathrm{Var}(\hat{C}_\ell^{TT\,i,j}, \hat{C}_{\ell'}^{TT\,p,q})\\ &\quad\approx \sqrt{C_\ell^{TT\,i,p}C_{\ell'}^{TT\,i,p}C_\ell^{TT\,j,q}C_{\ell'}^{TT\,j,q}}\;\Xi_{TT}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,p)^{TT},(j,q)^{TT}\right]_{\ell\ell'}\\ &\quad+\sqrt{C_\ell^{TT\,i,q}C_{\ell'}^{TT\,i,q}C_\ell^{TT\,j,p}C_{\ell'}^{TT\,j,p}}\;\Xi_{TT}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,q)^{TT},(j,p)^{TT}\right]_{\ell\ell'}\\ &\quad+\sqrt{C_\ell^{TT\,i,p}C_{\ell'}^{TT\,i,p}}\;\Xi_{TT}^{\emptyset\emptyset,}\left[(i,p)^{TT},(j,q)^{TT}\right]_{\ell\ell'}\\ &\quad+\sqrt{C_\ell^{TT\,j,q}C_{\ell'}^{TT\,j,q}}\;\Xi_{TT}^{\emptyset\emptyset,TT}\left[(j,q)^{TT},(i,p)^{TT}\right]_{\ell\ell'}\\ &\quad+\sqrt{C_\ell^{TT\,i,q}C_{\ell'}^{TT\,i,q}}\;\Xi_{TT}^{\emptyset\emptyset,TT}\left[(i,q)^{TT},(j,p)^{TT}\right]_{\ell\ell'}\\ &\quad+\sqrt{C_\ell^{TT\,j,p}C_{\ell'}^{TT\,j,p}}\;\Xi_{TT}^{\emptyset\emptyset,TT}\left[(j,p)^{TT},(i,q)^{TT}\right]_{\ell\ell'}\\ &\quad+\Xi_{TT}^{TT,TT}\left[(i,p)^{TT},(j,q)^{TT}\right]_{\ell\ell'}+\Xi_{TT}^{TT,TT}\left[(i,q)^{TT},(j,p)^{TT}\right]_{\ell\ell'}\,.\end{aligned} \tag{C10}$$

TTTE block:—

$$\begin{aligned}&\mathrm{Var}(\hat{C}_\ell^{TT\,i,j}, \hat{C}_{\ell'}^{TE\,p,q})\\ &\quad\approx \frac{1}{2}\sqrt{C_\ell^{TT\,i,p}C_{\ell'}^{TT\,i,p}}\left(C_\ell^{TE\,j,q}+C_{\ell'}^{TE\,j,q}\right)\;\Xi_{TT}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,p)^{TT},(j,q)^{TP}\right]_{\ell\ell'}\\ &\quad+\frac{1}{2}\sqrt{C_\ell^{TT\,j,p}C_{\ell'}^{TT\,j,p}}\left(C_\ell^{TE\,i,q}+C_{\ell'}^{TE\,i,q}\right)\;\Xi_{TT}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,q)^{TP},(j,p)^{TT}\right]_{\ell\ell'}\\ &\quad+\frac{1}{2}\left(C_\ell^{TE\,j,q}+C_{\ell'}^{TE\,j,q}\right)\;\Xi_{TT}^{\emptyset\emptyset,TT}\left[(j,q)^{TP},(i,p)^{TT}\right]_{\ell\ell'}\\ &\quad+\frac{1}{2}\left(C_\ell^{TE\,i,q}+C_{\ell'}^{TE\,i,q}\right)\;\Xi_{TT}^{\emptyset\emptyset,TT}\left[(i,q)^{TP},(j,p)^{TT}\right]_{\ell\ell'}\,.\end{aligned} \tag{C11}$$

TETE block—

$$\begin{aligned}
&\mathrm{Var}(\hat{C}_\ell^{TE\,i,j}, \hat{C}_{\ell'}^{TE\,p,q}) \\
&\quad\approx \sqrt{C_\ell^{TT\,i,p} C_{\ell'}^{TT\,i,p} C_\ell^{EE\,j,q} C_{\ell'}^{EE\,j,q}}\; \Xi_{TE}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,p)^{TT},(j,q)^{PP}\right]_{\ell\ell'} \\
&\quad+\frac{1}{2}\left(C_\ell^{TE\,i,q} C_{\ell'}^{TE\,j,p} + C_\ell^{TE\,j,p} C_{\ell'}^{TE\,i,q}\right)\; \Xi_{TT}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,q)^{TP},(j,p)^{PT}\right]_{\ell\ell'} \\
&\quad+\sqrt{C_\ell^{TT\,i,p} C_{\ell'}^{TT\,i,p}}\; \Xi_{TE}^{\emptyset\emptyset,PP}\left[(i,p)^{TT},(j,q)^{PP}\right]_{\ell\ell'} \\
&\quad+\sqrt{C_\ell^{EE\,j,q} C_{\ell'}^{EE\,j,q}}\; \Xi_{TE}^{\emptyset\emptyset,TT}\left[(j,q)^{PP},(i,p)^{TT}\right]_{\ell\ell'} \\
&\quad+\Xi_{TE}^{TT,PP}\left[(i,p)^{TT},(j,q)^{PP}\right]_{\ell\ell'}\,.
\end{aligned} \tag{C12}$$

TTEE block:—

$$\begin{aligned}
&\mathrm{Var}(\hat{C}_\ell^{TT\,i,j}, \hat{C}_{\ell'}^{EE\,p,q}) \\
&\quad\approx \frac{1}{2}\left(C_\ell^{TE\,i,p} C_{\ell'}^{TE\,j,q} + C_\ell^{TE\,j,q} C_{\ell'}^{TE\,i,p}\right)\; \Xi_{TT}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,p)^{TP},(j,q)^{TP}\right]_{\ell\ell'} \\
&\quad+\frac{1}{2}\left(C_\ell^{TE\,i,q} C_{\ell'}^{TE\,j,p} + C_\ell^{TE\,j,p} C_{\ell'}^{TE\,i,q}\right)\; \Xi_{TT}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,q)^{TP},(j,p)^{TP}\right]_{\ell\ell'}
\end{aligned} \tag{C13}$$

TEEE block:—

$$\begin{aligned}
&\mathrm{Var}(\hat{C}_\ell^{TE\,i,j}, \hat{C}_{\ell'}^{EE\,p,q}) \\
&\quad\approx \frac{1}{2}\sqrt{C_\ell^{EE\,j,q} C_{\ell'}^{EE\,j,q}}\left(C_\ell^{TE\,i,p} + C_{\ell'}^{TE\,i,p}\right)\; \Xi_{EE}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,p)^{TP},(j,q)^{PP}\right]_{\ell\ell'} \\
&\quad+\frac{1}{2}\sqrt{C_\ell^{EE\,j,p} C_{\ell'}^{EE\,j,p}}\left(C_\ell^{TE\,i,q} + C_{\ell'}^{TE\,i,q}\right)\; \Xi_{EE}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,q)^{TP},(j,p)^{PP}\right]_{\ell\ell'} \\
&\quad+\frac{1}{2}\left(C_\ell^{TE\,i,p} + C_{\ell'}^{TE\,i,p}\right)\; \Xi_{EE}^{\emptyset\emptyset,PP}\left[(i,p)^{TP},(j,q)^{PP}\right]_{\ell\ell'} \\
&\quad+\frac{1}{2}\left(C_\ell^{TE\,i,q} + C_{\ell'}^{TE\,i,q}\right)\; \Xi_{EE}^{\emptyset\emptyset,PP}\left[(i,q)^{TP},(j,p)^{PP}\right]_{\ell\ell'}\,.
\end{aligned} \tag{C14}$$

EEEE block:—

$$\begin{aligned}
&\mathrm{Var}(\hat{C}_\ell^{EE\,i,j}, \hat{C}_{\ell'}^{EE\,p,q}) \\
&\quad\approx \sqrt{C_\ell^{EE\,i,p} C_{\ell'}^{EE\,i,p} C_\ell^{EE\,j,q} C_{\ell'}^{EE\,j,q}}\; \Xi_{EE}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,p)^{PP},(j,q)^{PP}\right]_{\ell\ell'} \\
&\quad+\sqrt{C_\ell^{EE\,i,q} C_{\ell'}^{EE\,i,q} C_\ell^{EE\,j,p} C_{\ell'}^{EE\,j,p}}\; \Xi_{EE}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,q)^{PP},(j,p)^{PP}\right]_{\ell\ell'} \\
&\quad+\sqrt{C_\ell^{EE\,i,p} C_{\ell'}^{EE\,i,p}}\; \Xi_{EE}^{\emptyset\emptyset,PP}\left[(i,p)^{PP},(j,q)^{PP}\right]_{\ell\ell'} \\
&\quad+\sqrt{C_\ell^{EE\,j,q} C_{\ell'}^{EE\,j,q}}\; \Xi_{EE}^{\emptyset\emptyset,PP}\left[(j,q)^{PP},(i,p)^{PP}\right]_{\ell\ell'} \\
&\quad+\sqrt{C_\ell^{EE\,i,q} C_{\ell'}^{EE\,i,q}}\; \Xi_{EE}^{\emptyset\emptyset,PP}\left[(i,q)^{PP},(j,p)^{PP}\right]_{\ell\ell'} \\
&\quad+\sqrt{C_\ell^{EE\,j,p} C_{\ell'}^{EE\,j,p}}\; \Xi_{EE}^{\emptyset\emptyset,PP}\left[(j,p)^{PP},(i,q)^{PP}\right]_{\ell\ell'} \\
&\quad+\Xi_{EE}^{PP,PP}\left[(i,p)^{PP},(j,q)^{PP}\right]_{\ell\ell'} + \Xi_{EE}^{PP,PP}\left[(i,q)^{PP},(j,p)^{PP}\right]_{\ell\ell'}\,.
\end{aligned} \tag{C15}$$

D. NOISE SPECTRUM FITS

In Figure 16 we show the ratio of signal power spectrum to noise power spectrum for the half-mission 1 maps, to illustrate how at large angular scale in temperature the noise spectrum is significantly smaller than the signal spectrum, and thus the covariance is dominated by cosmic variance. In polarization the noise spectrum is not sub-dominant at any scale. Figures 17 and 18 show the smooth noise models fit to the estimated noise power spectra, for temperature and polarization respectively.

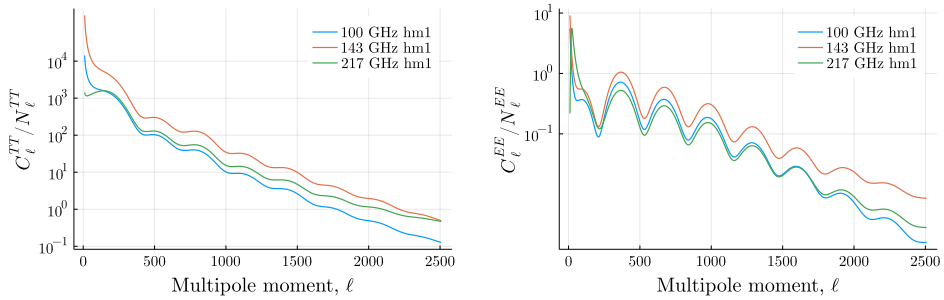

Figure 16. The ratio of signal power spectrum (including foregrounds) and noise power spectrum for the half-mission 1 frequency maps in temperature (above) and polarization (below).

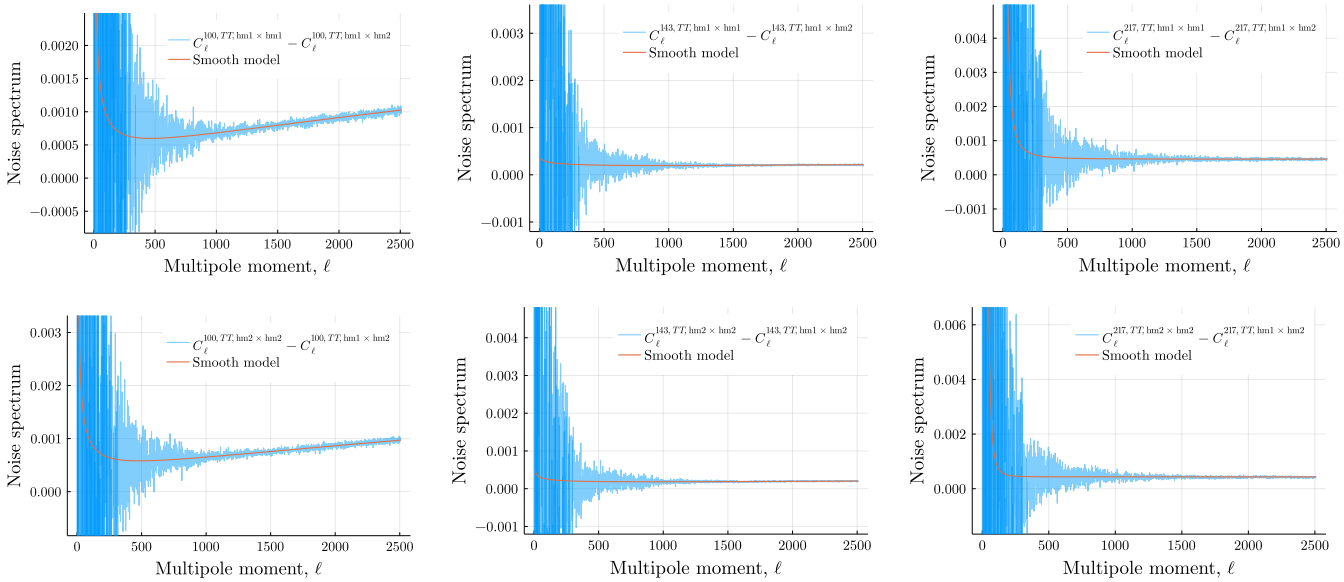

Figure 17. The estimated noise power spectrum (blue) and our smooth model fits (orange) for the *Planck* temperature maps. At large scales ($\ell < 500$), the spectra are signal-dominated. The estimate from the difference of auto- and cross-spectra from the half-mission splits is thus contaminated by noise × signal terms. We note that the covariance matrix is not sensitive to the noise power spectrum at such scales, since those scales are signal-dominated.

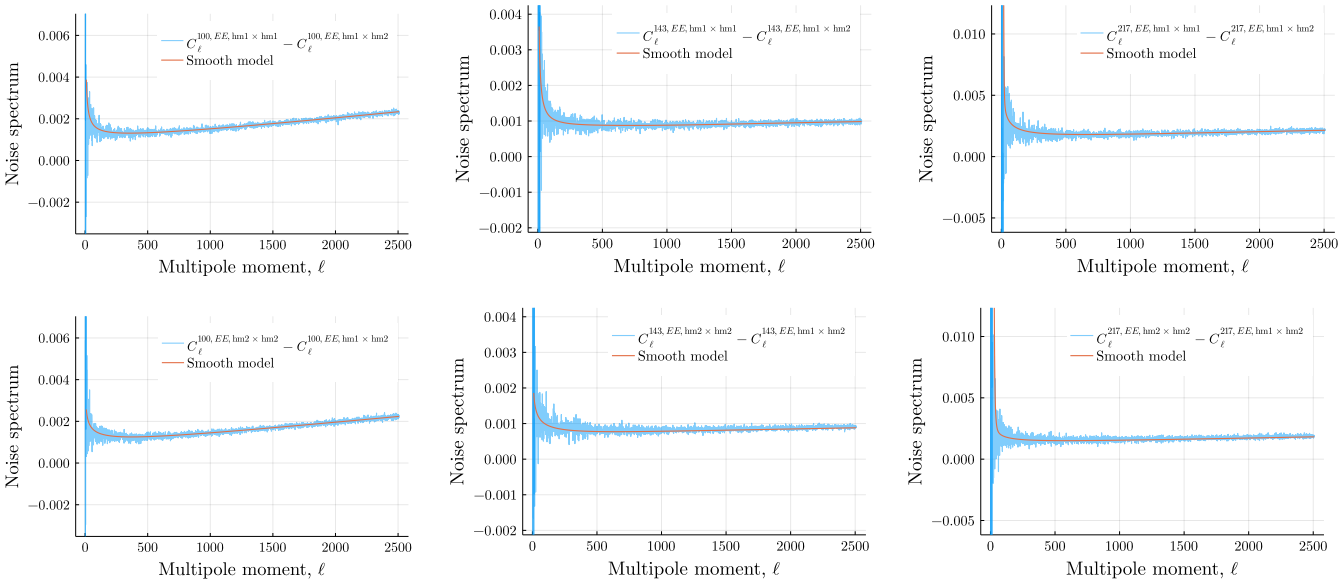

Figure 18. The estimated noise power spectrum (blue) and our smooth model fits (orange) for the *Planck* polarization maps. Unlike temperature, the polarization power spectra are always noise-dominated in every bin, leading to a well-determined noise power spectrum and resulting model fit.