Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **517,** 4620–4636 (2022) https://doi.org/10.1093/mnras/stac2744
Advance Access publication 2022 September 26

CMB power spectra and cosmological parameters from *Planck* PR4 with CamSpec

Erik Rosenberg ,★ Steven Gratton and George Efstathiou
Institute of Astronomy and Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge CB3 0HA, UK

Accepted 2022 September 20. Received 2022 September 20; in original form 2022 July 11

ABSTRACT

We present angular power spectra and cosmological parameter constraints derived from the *Planck* PR4 (NPIPE) maps of the cosmic microwave background. NPIPE, released by the Planck Collaboration in 2020, is a new processing pipeline for producing calibrated frequency maps from *Planck* data. We have created new versions of the CamSpec likelihood using these maps and applied them to constrain the lambda cold dark matter (ΛCDM) model and single-parameter extensions. We find excellent consistency between NPIPE and the *Planck* 2018 maps at the parameter level, showing that the *Planck* cosmology is robust to substantial changes in the mapmaking. The lower noise of NPIPE leads to $\sim$10 per cent tighter constraints, and we see both smaller error bars and a shift towards the ΛCDM values for beyond-ΛCDM parameters including Ω_K and A_L.

Key words: methods: data analysis – cosmic background radiation – cosmological parameters.

1 INTRODUCTION

The *Planck* survey of the anisotropies of the cosmic microwave background (CMB) has been invaluable to cosmology. It has tightly constrained the lambda cold dark matter (ΛCDM) model parameters, ruled out many plausible extensions, and as an important legacy data set is now frequently combined with ground-based CMB data and other cosmological probes. *Planck* analysis has evolved significantly in this time, from the initial 2013 data release (Planck Collaboration I 2014) to the addition of polarization in 2015 (PR2) (Planck Collaboration I 2016), and substantial improvement of systematic corrections in 2018 (PR3) (Planck Collaboration I 2020; Planck Collaboration VI 2020). Further work on the data processing and mapmaking led to the 2019 publication of SRoll2 (Delouis et al. 2019), an extension of 2018's SRoll destriping algorithm (Planck Collaboration XLVI 2016), and later the NPIPE release in 2020 (Planck Collaboration LVII 2020). NPIPE is a new and independent pipeline to produce frequency maps from the raw time-ordered data (TOD), with substantial differences in detector calibration and systematic corrections compared to previous releases. In this paper, we present a new version of the CamSpec likelihood, developed in previous *Planck* likelihood papers (Planck Collaboration XV 2014; Planck Collaboration XI 2016; Planck Collaboration V 2020) and described in detail in Efstathiou & Gratton (2021). We use NPIPE 100, 143, and 217 GHz maps to calculate pseudo-power spectra and covariances and so build a high-ℓ NPIPE likelihood. The 353 and 545 GHz channels are also used as Galactic dust templates. We compare the power spectra and parameter constraints obtained to those found with the 2018 maps, allowing us to evaluate the robustness of *Planck* cosmology to quite substantial changes in the data processing.

We begin by summarizing the NPIPE mapmaking pipeline (Section 2), the data and masks we use (Sections 3.1 and 3.2), and the dust-cleaning methodology (Section 3.3). We show details of the NPIPE temperature-to-polarization (TP) leakage and effective polarization efficiencies in Section 3.4. Our noise estimation method and NPIPE noise spectra are presented in Sections 3.5 and 3.6, with the latter focusing on noise correlations between data splits. NPIPE power spectra are shown in Section 4, along with tests of internal consistency. We present parameter constraints from the new NPIPE likelihoods in Section 5 and explore extensions to ΛCDM in Section 6. We compare to an alternative NPIPE likelihood as well as to ACT and SPT in Section 7 and conclude in Section 8. For conciseness, we refer to the *Planck* 2018 high-frequency instrument (HFI) mapmaking and likelihood papers (Planck Collaboration III 2020; Planck Collaboration V 2020) as HFI18 and PPL18. Similarly, we call the CamSpec paper (Efstathiou & Gratton 2021) EG21 and the NPIPE paper (Planck Collaboration LVII 2020) NP20.

2 INTRODUCTION TO NPIPE

The NPIPE pipeline is described in detail in NP20, including a comparison of the processing and resulting maps to PR3, a description of the simulations, and an estimate of the optical depth to reionization τ. Here we summarize the main aspects of the pipeline relevant to this work, focusing on HFI.

2.1 NPIPE pre-processing

The first stage of the NPIPE pipeline is pre-processing, performed on single detectors and mostly on single pointing periods. This is analogous to the construction of cleaned TOD in previous *Planck*

★ E-mail: er510@cam.ac.uk

releases (Planck Collaboration VII 2016; Planck Collaboration III 2020, section 2.1).

The first processing stage involves estimation of the sky signal from raw bolometer data, for use in subsequent steps. *Planck* scanned a given circle on the sky ∼39–65 times before moving on to the next. Within one such pointing period the data from each revolution may be binned together to estimate the sky signal on the circle. These estimates, referred to as (HEALPix[1]) ring maps, were used extensively in PR3 and earlier releases. NPIPE instead prefers a 'global signal estimate' in all cases except glitch removal. The global signal estimate involves sampling full-mission maps from an earlier iteration of NPIPE into the time domain, then adding an estimate of the orbital dipole and convolving by the bolometer transfer function. This signal estimate is intended to be less noisy than that derived from rings, and also allows for the use of 'repointing' data, taken between the pointing periods.

Apart from the signal estimation, the most significant differences between NPIPE pre-processing and that done previously are in the handling of 4K lines and glitches. As previously (Planck Collaboration VII 2016), 4K lines are fitted for each pointing period with sine and cosine templates at selected harmonic frequencies of the 4K cooler electronics in the time domain. NPIPE fits for six more harmonics seen in some detectors. For the six strongest (highest S/N) lines NPIPE also chooses to use a shorter fitting period than the full pointing period so as to reduce line residuals due to the time-varying amplitudes and phases of the lines. When removing glitches – large residuals in the data due primarily to cosmic ray hits – NPIPE uses a phase-binned estimate of the signal for each pointing period, similar to the PR3 rings. However, NPIPE employs wider bins, and uses as the signal estimate a polynomial fit to the data in a given bin plus its two neighbours instead of a simple mean. The additional smoothing leads to a less noisy signal estimate, reducing half-ring correlations and allowing the glitch threshold to be increased, reducing the number of flagged glitches by about half. Finally, the bolometer time response deconvolution step has been modified to minimize the effects of the gaps left by glitch removal on the science data, reducing small-scale noise and half-ring correlations.

2.2 NPIPE reprocessing and destriping

For the reprocessing stage NPIPE uses a generalized destriper (Keihänen et al. 2004; Keihänen, Kurki-Suonio & Poutanen 2005; Keihänen et al. 2010) similar to SRoll (Planck Collaboration XLVI 2016) to perform calibration and systematic template fitting. Writing the data model $\mathbf{d} = \mathsf{P}\mathbf{m} + \mathsf{F}\mathbf{a} + \mathbf{n}$ we have TOD $\mathbf{d}$, pointing matrix P, sky map $\mathbf{m}$, and a white noise term $\mathbf{n}$. F is a matrix of time-domain templates representing systematic and astrophysical effects to be removed, with $\mathbf{a}$ the amplitude of each template. NPIPE solves for template amplitudes $\mathbf{a}$, using them to make a template-cleaned TOD $\mathbf{d}_{\rm clean} = \mathbf{d} - \mathsf{F}\mathbf{a}$. Since the template for gain fluctuations is derived from the estimated sky map, the cleaned TOD is used to update estimates of F and $\mathbf{m}$ and the process iterated until the template amplitudes converge to zero.

While using a similar principle to HFI18, NPIPE differs significantly in the templates F used. One major change is in the imposition of a prior on the polarization maps during calibration. This is done to resolve a degeneracy created between gain and polarization residuals due to the *Planck* scanning strategy. In effect fluctuations are allowed in the gain which are compensated by spurious polarization signal. To address this problem NPIPE follows the methods of LFI (Planck Collaboration II 2020) and uses the 30, 217, and 353 GHz polarization maps as templates in F. This allows the gains to be solved on an effectively unpolarized sky map, removing the degeneracy. Because the polarization templates are foreground dominated, the destriper will try to fit the CMB signal with the other templates in F. This suppresses the CMB polarization signal and leads to a large-scale polarization transfer function in NPIPE.

Another significant change in NPIPE is the use of much shorter baselines for low-frequency noise offsets. F models $1/f$ noise as a constant offset for each baseline period; HFI18 used a single such offset per ring, while in NPIPE each corresponds to 1° on the sky. A single pointing period would therefore include thousands of these shorter baselines. NP20 finds that this method better captures low-frequency instrumental noise, leading to lower noise at $\ell > 20$.

Finally, NPIPE introduces two time domain polarization templates to measure corrections to polarization efficiency and polarization angle, and also adds a template to better account for systematics from the analogue-to-digital converter non-linearity (ADCNL). In addition to allowing for ADCNL-induced effective gain fluctuations as in SRoll, NPIPE includes a 'distortion' template corresponding to effective gains that change as a function of input signal level. This is done to further reduce TP leakage from the ADCNL.

2.3 Other differences between NPIPE and PR3

In making maps from independent subsets of the data from which to calculate cross-spectra, NPIPE chooses to split into sets of detectors, referred to as A and B, rather than the half-mission split used in EG21 and PPL18. The NPIPE detector sets are similar to the detector-set splits explored extensively in EG21 but also group one or two spiderweb bolometers with each set of polarization sensitive bolometers. More importantly however the NPIPE reprocessing steps – destriping and systematic template fitting – are performed independently on each set. This is in contrast to PR3 in which the full-mission data set at each frequency was fit for systematics (gains and offsets, bandpass corrections, transfer function, far side lobes) together and used for all map products. Indeed, SRoll explicitly minimizes the χ^2 difference in signal between one bolometer and the average of all bolometers in a frequency band when solving for the template amplitudes; this can introduce correlations in the subset maps. This suggests NPIPE detector sets should be more independent than those of PR3, as we test in Section 3.6.

Finally, we note that NPIPE reports an 8 per cent increase in integration time from including repointing manoeuvre data, taken between the stable pointing periods as the satellite reoriented from ring to ring. This is reflected in lower noise levels as discussed in Section 3.5, although it does not account for the entire difference.

3 METHODS AND POWER SPECTRA

To create an NPIPE likelihood we begin by following the methods of EG21, generating analogues of the 12.5HMcl likelihood described there. As in previous CamSpec releases we use the pseudo-C_ℓ method (Hivon et al. 2002; Kogut et al. 2003) in which power spectra are computed on masked skies and a coupling matrix is applied to deconvolve the effect of the mask. In this section, we discuss foreground cleaning, calibrations, and systematic corrections

[1] http://healpix.sourceforge.net

subsequently applied to the deconvolved spectra. We also introduce the noise modelling and investigate correlated noise in NPIPE.

3.1 Data

We use as our starting point *Planck* HFI maps at 100, 143, 217, 353, and 545 GHz. For our PR3-based analysis we use the 2018 half-mission maps with associated pixel covariances from the *Planck* Legacy Archive (PLA).[2] As in EG21 and previous releases the main likelihoods presented in this paper use data-split differences to estimate the noise power, and we compare to results from end-to-end simulations. For PR3 we estimate noise with the FFP10 simulations, available on the PLA, and half-mission odd-even ring differences.

For NPIPE we use the A/B detector-set maps at the same frequencies, available at NERSC,[3] with the associated wcov_mcscaled files for pixel covariances. We use the NPIPE simulations and half-ring A/B maps from NERSC for noise estimation. We subtract the PR2 solar dipole before using the frequency maps, noting that imperfections in this subtraction should not affect our high-ℓ likelihoods. The NPIPE low-ℓ EE transfer function is neglected as we do not use $\ell < 30$.

To deconvolve the instrument beams we employ transfer functions W_ℓ calculated with QuickPol (Hivon, Mottet & Ponthieu 2017) and provided with each data release.

3.2 Sky masks

The sky masks used here are identical to those used in EG21 and described in Planck Collaboration XI (2016), up to corrections for 'missing pixels' which we address momentarily. As previously we use diffuse foreground masks to remove strong Galactic dust emission and also mask regions with significant emission from point sources, CO, and extended objects by applying additional 'point source' masks. In all our likelihoods we use a diffuse mask that retains 80 per cent sky area before apodization (mask80) at all frequencies in temperature and polarization. Each frequency has a different point source mask in temperature, and the 143 GHz point source mask is used for all frequencies in polarization.

Planck half-mission maps contain a number of pixels (from 2500 for 100 GHz half-mission 1 to 115 000 in 217 GHz half-mission 2) marked as missing due to poor determination of the polarization in those pixels. NPIPE maps do not contain any such missing pixels, owing to the inclusion of data from repointing manoeuvres as well as changes to the glitch detection. Accordingly no missing pixels are masked in our NPIPE likelihoods.

3.3 Dust cleaning

The likelihoods presented here all include 'dust cleaning', using higher frequency *Planck* maps to remove the effects of Galactic dust at the power spectrum level as in equation (1):

$$\begin{aligned} C^{X_{\nu_1}Y_{\nu_2}\,\mathrm{clean}} = &\left(1+\alpha^{X_{\nu_1}}\right)\left(1+\alpha^{Y_{\nu_2}}\right) C^{X_{\nu_1}Y_{\nu_2}} \\ &-(1+\alpha^{X_{\nu_1}})\alpha^{Y_{\nu_2}} C^{X_{\nu_1}Y_{\nu_T}} \\ &-(1+\alpha^{Y_{\nu_2}})\alpha^{X_{\nu_1}} C^{Y_{\nu_2}X_{\nu_T}} + \alpha^{X_{\nu_1}}\alpha^{Y_{\nu_2}} C^{X_{\nu_T}Y_{\nu_T}}. \end{aligned} \quad (1)$$

The α^{X_ν} are cleaning coefficients at frequency ν; C^{XY} is a mask-deconvolved, beam-corrected power spectrum where $X, Y \in \{T, E, B\}$; ν_T is the template frequency, here 353 GHz for TE and EE spectra and 545 GHz for TT spectra.

The cleaning coefficients α^{X_ν} are calculated by minimizing the function $\sum_{\ell_{\min}}^{\ell_{\max}} C_\ell^{X_\nu X_\nu\,\mathrm{clean}}$. This minimization gives a biased coefficient α' and so we instead use the corrected value $\alpha = \alpha'/(1+\alpha')$ as the cleaning coefficient. In temperature we minimize over the ℓ-range $100 \le \ell \le 500$, where Galactic dust dominates over the cosmic infrared background (CIB). For polarization spectra we use $30 \le \ell \le 300$. To limit the impact of noise from the 353 GHz maps in polarization, we fit the dust contribution with a power law and at $\ell > 150$ subtract this fit from the uncleaned spectra rather than using the cleaned spectra.

As demonstrated in EG21 this method is very effective at removing dust emission in both temperature and polarization. This allows all TE and EE spectra to be coadded, and no foreground nuisance parameters are included for polarization in the likelihood. For temperature, we include nuisance parameters describing power-law residuals for each frequency combination, accounting for power from residual CIB, point sources, and other foregrounds. At 100 GHz large-scale foregrounds are not removed as accurately as at 143 GHz, as 100 GHz maps contain CO emission and non-negligible synchrotron emission that are difficult to capture accurately in even a more complicated foreground model. 100 GHz is also noisier at high multipoles and the primary CMB temperature fluctuations are already excellently determined by 143 and 217 GHz. Therefore, we omit 100 GHz from the temperature likelihoods.

3.4 Calibration, polarization efficiencies, and TP leakage

3.4.1 Temperature calibration

We apply multiplicative calibration factors to each temperature spectrum entering the likelihood, interpreting these differences as small transfer function errors due to mismodelling of the beam outside the main beam. The calibration factor $c_{\nu_1\nu_2}$ for a temperature spectrum $D_\ell^{\nu_1\nu_2}$ is calculated by minimizing the χ^2 of equation (2):

$$\chi^2 = \sum_k \frac{\left(c_{\nu_1\nu_2} D_k^{\nu_1\nu_2} - D_k^{143\times143}\right)^2}{\sigma_{\nu_1\nu_2}^2}. \quad (2)$$

The sum is over bandpowers k in the multipole range $50 \le \ell \le 500$, and σ is the scatter of those bandpowers over the fitted multipole range. The 143×143 spectrum is subtracted so as to eliminate cosmic variance. In order to remove foreground effects before minimization we also apply a conservative Galactic mask (mask30), the 217 GHz point source mask, and 545 GHz cleaning to all frequencies. The resulting temperature calibration factors $c_{\nu_1\nu_2}$ are applied to each TT spectrum in the likelihood. NPIPE TT calibrations relative to 143×143 are given in Table A1.

3.4.2 Effective polarization efficiencies

Calibration factors are also computed for the polarization spectra and interpreted as effective polarization efficiencies. In reality these may

[2] https://pla.esac.esa.int/pla

[3] National Energy Research Scientific Computing Center, see https://crd.lbl.gov/divisions/scidata/c3/c3-research/cosmic-microwave-background/cmb-data-analysis-at-nersc/

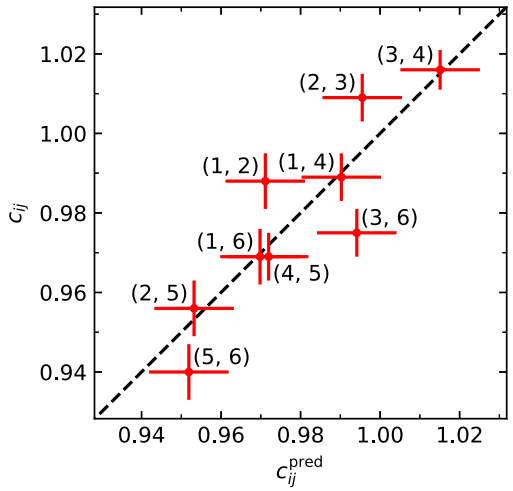

Figure 1. A comparison of NPIPE 12.6 calibration factors c_{ij} measured from the EE spectra and those predicted from triplets of TE spectra. The labels refer to the EE index given in Table A1.

be due to a combination of real polarization efficiencies, polarization angle errors, and beam effects as in temperature. Similarly to temperature we calculate the calibration factor c_k for a given TE or EE cross-spectrum k by minimizing equation (3):

$$\chi^2 = \sum_{\ell_1\ell_2} \left(C^k_{\ell_1} - c_k C^{\rm theory}_{\ell_1}\right) \left(M^k\right)^{-1}_{\ell_1\ell_2} \left(C^k_{\ell_2} - c_k C^{\rm theory}_{\ell_2}\right). \quad (3)$$

Here, $C_\ell^{\rm theory}$ is a theoretical power spectrum[4] and M^k is the covariance matrix for spectrum k. We minimize with respect to $C_\ell^{\rm theory}$ to avoid instability in the estimate from the high noise in polarization; we expect any bias introduced favouring the theoretical model used for calibration to appear as an overall multiplicative factor that can be absorbed in the $c_{\rm TE}$ and $c_{\rm EE}$ factors of the likelihood (see Section 5). We see no indications for significant pulls outside the prior range of these calibration parameters or strong correlations with them in any ΛCDM extensions we test. The calibrations c_k, given in Table A1, are divided out from each spectrum before coaddition. As an additional check on these calibration factors, we can use the TE calibrations to predict those for EE. We group together TE spectra sharing the same polarization map and average their calibration factors to calculate an effective polarization efficiency $\bar\rho$ for each map. These can be combined into a predicted calibration $c_{ij}^{\rm pred} = \bar\rho_i \bar\rho_j$ for each EE spectrum. We compare these predicted values to the measured EE calibrations in Fig. 1. We see a strong correlation, suggesting that our interpretation of these calibration factors as effective polarization efficiencies is a good one.

3.4.3 *Temperature-to-polarization leakage*

In PR3 as well as NPIPE great care was taken to eliminate TP leakage due for example to bandpass mismatch or ADCNL. Another source of TP leakage is beam mismatch; since polarization maps are created using pairs of bolometers, any mismatch of the beams of the individual bolometers in a pair will create TP leakage in the map. We calculate a correction for this effect using the off-diagonal terms of the beam matrix W_ℓ computed using QuickPol, as shown in equation (4):

$$\bar{\tilde C}_\ell^{TE} = W_\ell^{TETE}\tilde C_\ell^{TE} + W_\ell^{TETT}\tilde C_\ell^{TT} \quad (4a)$$

$$\bar{\tilde C}_\ell^{EE} = W_\ell^{EEEE}\tilde C_\ell^{EE} + W_\ell^{EETT}\tilde C_\ell^{TT}. \quad (4b)$$

Here, $\bar{\tilde C}_\ell$ is the beam-convolved spectrum, while $\tilde C_\ell$ is beam-corrected; in practice, we use theoretical spectra C_ℓ^{TT} rather than $\tilde C_\ell^{TT}$ to estimate the leakage corrections. The QuickPol correction of equation (4) is applied for TE but found to be negligible for EE.

To test the effectiveness of the leakage correction, we follow the method laid out in EG21 and organize the uncorrected TE spectra into 'triplets'. Each triplet consists of three spectra all sharing the same polarization map but using different frequency maps in temperature. We expect TP leakage in the polarization map to correlate identically to the dominant CMB component of each temperature map, hence we look for correlated residuals in these triplets as signatures of TP leakage and compare to the correction factor of equation (4). The TE spectra are mask and beam deconvolved, corrected for the polarization efficiencies of Table A1, and dust cleaned with 353 GHz. Residuals are calculated with respect to the mean of all TE spectra rather than to a theoretical model to reduce the impact of cosmic variance.

This test is shown for selected NPIPE spectra in Fig. 2; the solid lines are a fit to the data, while the dotted lines show the leakage correction derived from QuickPol. For 100 GHz, where leakage is most significant, we see excellent agreement between the three spectra using different temperature maps, and between the data and the QuickPol prediction. The same is largely true of the other spectra as well, although there the leakage is quite small. Although the agreement is not perfect, complete neglect of this leakage is found in section 6.4 of EG21 to cause up to 1σ shifts in cosmological parameters; therefore, we expect any error from an imperfect leakage model to be significantly less than 1σ. We also note that both the measured and predicted TP leakage is quite different between NPIPE detector sets and the PR3 half-missions presented in EG21, being for example of a similar magnitude but opposite sign in 100 GHz. Since the leakage correction is calculated from QuickPol in the same way in each case, we attribute this difference to the different split of the data in the two analyses.

3.5 Noise estimation and covariances

We use analytic covariance matrices as given in the appendix A.2 of EG21. To model the noise in *Planck* we estimate a noise power spectrum for each map and fit it with an analytic function of the form

$$N_\ell = A\left(\frac{100}{\ell}\right)^\alpha + B\frac{(\ell/1000)^\beta}{(1+(\ell/\ell_c)^\gamma)^\delta}. \quad (5)$$

Here, A, α, B, β, ℓ_c, γ, and δ are free parameters. Weight functions $\psi_\ell = N_\ell/N_\ell^{\rm white}$ are computed for each T, Q, U component, and the pixel noise estimates $(\sigma^X)_i^2$ in the original formulae for the covariances are replaced by $\sqrt{\psi_\ell^X\psi_{\ell'}^X}(\sigma^X)_i^2$.

Accurate modelling of the noise is especially important for the noise-dominated polarization spectra and has been a persistent challenge. Traditionally, we have estimated noise spectra using differences of data splits. The Plik likelihood in PPL18 uses half-ring difference (HRD) maps for this purpose – half-ring maps are made using the first or second halves of each pointing period and their difference gives an estimate of the noise. The PR3 half-ring maps are known to be correlated due to the glitch processing, so the HRD spectra are an underestimate of the true noise level. In PPL18 this effect was corrected for with an empirical bias

[4] $C_\ell^{\rm theory}$ here is the ΛCDM best fit to CamSpec 12.1HM TT from EG21.

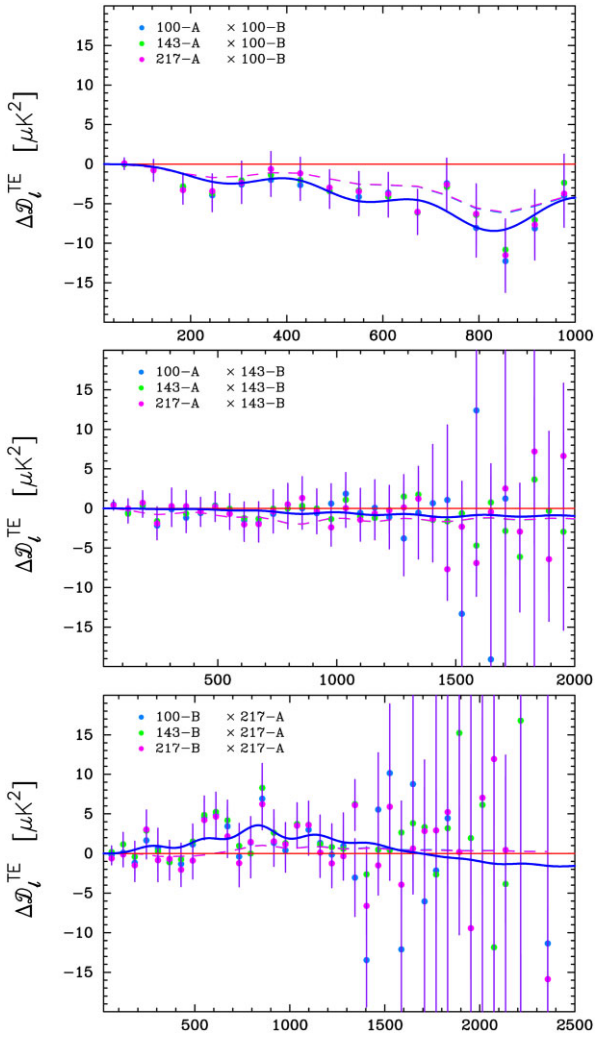

Figure 2. Residuals of NPIPE TE spectra with respect to the mean of all TE spectra. The spectra are mask and beam deconvolved, 353 GHz cleaned, and corrected for effective polarization efficiencies. We see a clear correlation between spectra with the same polarization maps, suggesting the residuals are due to TP leakage. The dashed lines show the beam-induced leakage model (from equation 4) corrected for in the likelihood. The solid lines are a fit to the residuals.

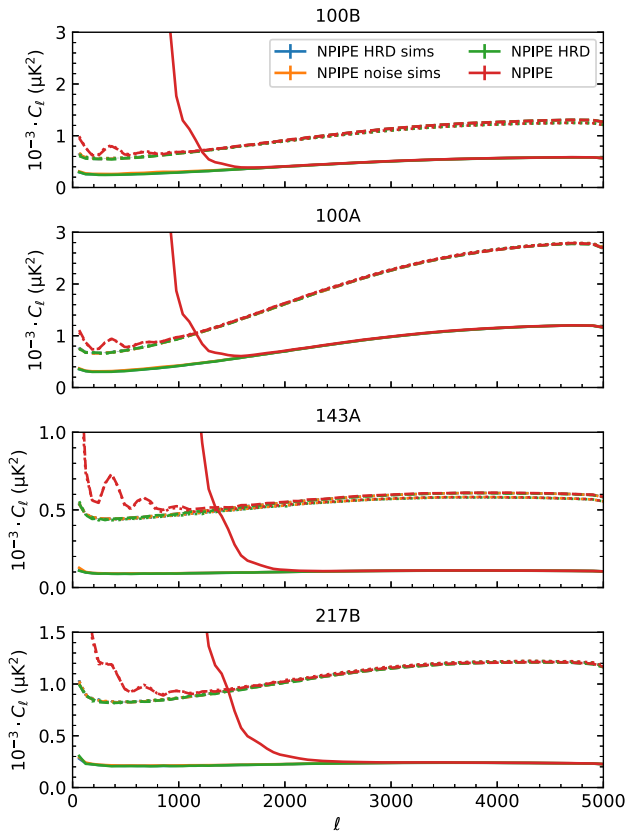

Figure 3. NPIPE noise spectra before mask and beam deconvolution. The solid, dashed, and dotted lines show the autospectra of T, Q, and U maps, respectively. For each frequency and split we show the autospectrum, HRD spectrum, and the simulated version of each. For the simulations we show the mean of 100 simulations. 143B and 217A are qualitatively similar to 143A/217B and are omitted.

correction; EG21 chose instead to use odd-even ring differences (OED), calculated similarly to HRDs but using odd and even rings. There is evidence that these maps overestimate the noise, particularly for 100 × 100 GHz and at large scales, but in this paper we continue to use them for PR3.

In Fig. 3 we present noise power spectra for NPIPE, comparing noise estimates derived from simulations to autospectra of the detector-set maps and to HRDs. From all panels of Fig. 3 we see no evidence of underestimation of the noise by NPIPE HRDs at a level that would affect our analysis. NP20 attributes this reduction in half-ring correlations to the changes in the glitch processing. We observe excellent agreement between the NPIPE simulations, HRDs, and autospectra at high multipoles. Below ℓ of 2000 however the HRD spectra can be below the simulations by a few per cent in polarization and up to 5 per cent in temperature, especially at 100 GHz. NP20 also introduces a small-scale 'noise alignment' correction to the simulated noise spectra in order to match to A − B difference spectra. This further boosts the simulated noise power by 1–2 per cent in both temperature and polarization. These differences are of similar order to the accuracy of the fitting function equation (5), which is able to match the spectra to the 1–2 per cent level in polarization and 5 per cent in temperature. The current version of the NPIPE likelihood also neglects to weight the HRD map with per-pixel variances $w_i = \sigma_1\sigma_2/(\sigma_1^2 + \sigma_2^2)$ (see equation 5.3 of EG21), instead using a constant weighting of 0.5. This leads to a 10 per cent overestimate of the QQ noise for 100B and 143B and smaller (1–2 per cent) errors in polarization for other spectra. The effect on parameters should be small, but will be corrected in future work.

In Fig. 4, we see that NPIPE has somewhat lower noise levels than the PR3 maps, potential overestimation from odd-even maps aside. This is expected to be due primarily to the addition of repointing data and the use of short baselines in destriping, and results in a 10–30 per cent noise reduction at the scales we use, dependent on frequency and multipole.

3.6 Correlated noise

We use cross-spectra of the NPIPE HRD maps to check for correlated noise, with examples from 143 × 217 shown in Fig. 5 and others in Fig. A1. These spectra are consistent with zero, indicating there is negligible correlated noise; notably, this is also true for A × A and B × B spectra. In particular, we find that unlike for the PR3

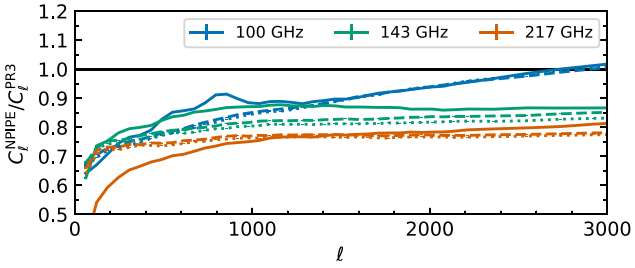

Figure 4. NPIPE undeconvolved full-mission simulated noise spectra divided by their PR3 counterparts. The solid, dashed, and dotted lines show autospectra of T, Q, and U maps, respectively.

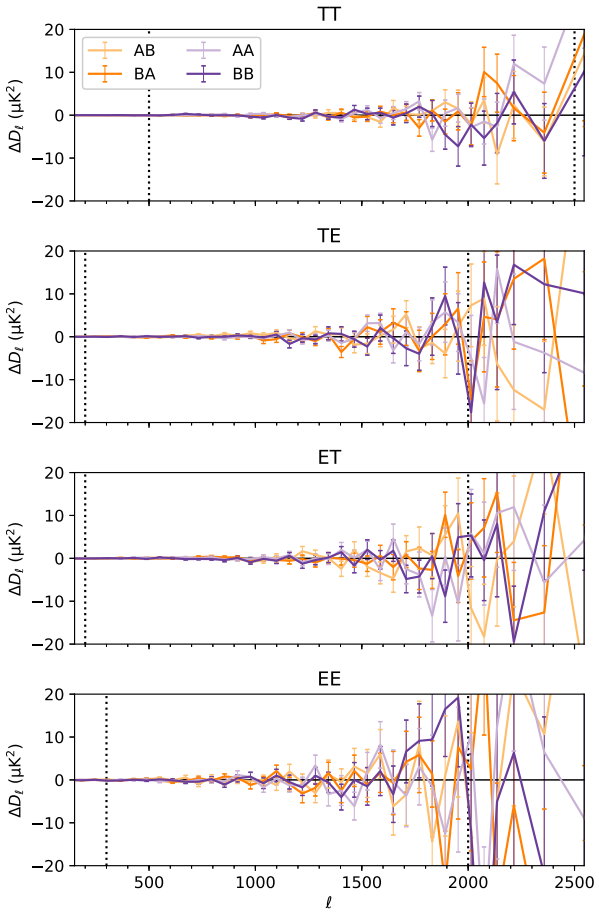

Figure 5. NPIPE HRD cross-spectra for 143 × 217 GHz. The error bars show the standard deviation $\sigma/\sqrt{\Delta\ell}$ within each bandpower of bin-width $\Delta\ell$. The vertical lines indicate the ℓ-range used in the likelihood.

detector sets the 143 × 143 and 217 × 217 TT spectra show no evidence of correlated noise. The most obvious deviations from zero are at high-ℓ in the 143B × 217B EE spectrum, where the noise is nevertheless very high. Therefore, we hope to add all AA and BB spectra in TE and EE to a future version of the likelihood as the addition of 6 spectra – AA and BB for each of 100 × 143, 100 × 217, and 143 × 217 – would significantly improve the constraining power of EE. TE would improve by a smaller amount as each new spectrum is equivalent to one already in the likelihood up to A-B exchange of the already signal-dominated temperature map.

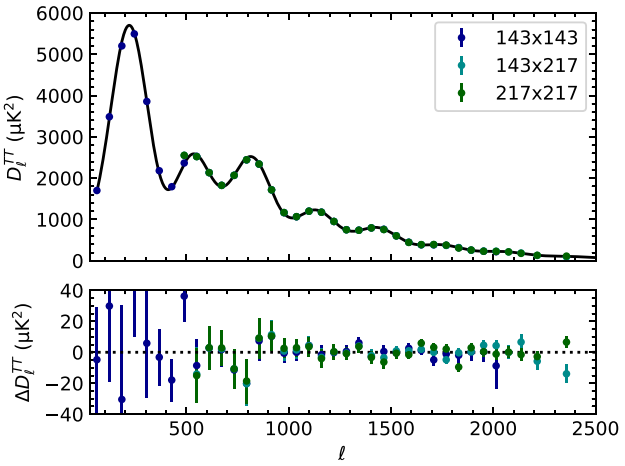

Figure 6. NPIPE 12.6 TT power spectra with best-fitting foregrounds removed. Residuals are shown with respect to the TTTEEE best-fitting theory.

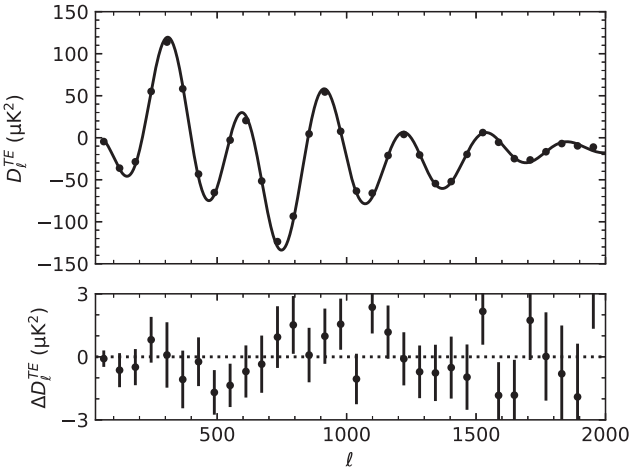

Figure 7. NPIPE 12.6 TE power spectrum and residual with respect to the TTTEEE best-fitting theory.

3.7 New likelihoods and nomenclature

Before discussing the spectra and parameter constraints in more detail we pause to introduce the nomenclature surrounding our likelihoods. As discussed previously our new likelihoods build on CamSpec12.5HMcl, presented in EG21. The covariance matrices of 12.5HMcl included contributions from foregrounds despite being a foreground-cleaned likelihood; this led to some overestimation of the error bars at high multipoles. This had negligible effect on cosmology but affected χ^2 values in TT and has been corrected in the new 12.6HMcl, which in this paper we will call PR3_12.6. Using NPIPE A and B maps in place of PR3 half-mission maps and half-ring instead of odd-even noise estimates we have an analogous NPIPE likelihood, PR4_12.6.

4 SPECTRA AND CONSISTENCY CHECKS

As in CamSpec 12.5HMcl we use co-added spectra at 143 × 143, 143 × 217, and 217 × 217 GHz in TT, while co-adding all A × B cross-spectra of 100, 143, and 217 GHz in TE and EE. The multipole ranges used are given in Table A2. Co-added NPIPE 12.6 power spectra are shown in Figs 6 (TT), 7 (TE), and 8 (EE), with residuals to the ΛCDM best fit to TTTEEE (discussed in Section 5). Note

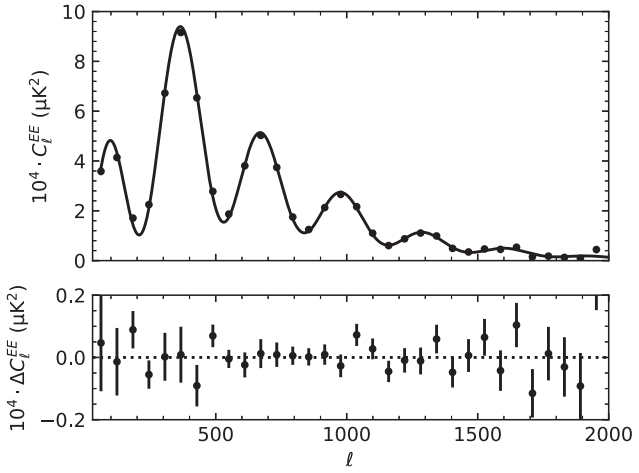

Figure 8. NPIPE 12.6 EE power spectrum and residual with respect to the TTTEEE best-fitting theory.

Table 1. χ^2 of the different components of the PR4_12.6 likelihood with respect to the TTTEEE best-fitting model. N_D is the size of the data vector. $\hat{\chi}^2 = \chi^2/N_D$ is the reduced χ^2. The last column gives the number of standard deviations of $\hat{\chi}^2$ from unity.

PR4_12.6	ℓ range	N_D	$\hat{\chi}^2$	$(\hat{\chi}^2 - 1)/\sqrt{2/N_D}$
TT 143 × 143	30–2000	1971	1.021	0.67
TT 143 × 217	500–2500	2001	0.985	−0.47
TT 217 × 217	500–2500	2001	1.002	0.05
TT Coadded	30–2500	2471	1.025	0.87
TT All	30–2500	5973	1.074	4.07
TE	30–2000	1971	1.055	1.73
EE	30–2000	1971	1.026	0.82
TEEE	30–2000	3942	1.046	2.02
TTTEEE	30–2500	9915	1.063	4.46

that we plot binned data here for clarity, but the actual likelihood is unbinned.

Visually the ΛCDM best fit shown in Figs 6–8 appears to match the spectra well; to more quantitatively check the agreement of the data and the model for PR4_12.6 we calculate χ^2 values with respect to the TTTEEE best-fitting model, shown in Table 1. For each of the individual TT spectra as well as co-added TT, TE, and EE we find acceptable values for χ^2. However, the χ^2 values for total TT and TTTEEE are somewhat large, being over 4σ high. The source of these high χ^2 is primarily at $500 \lesssim \ell \lesssim 1000$, and to a lesser extent $\ell < 500$. We see the same effect in PR3_12.6, and it was also reported in Planck Collaboration XV (2014). The effect is not unexpected as when all three temperature spectra are present and particularly when as now computed on the same Galactic mask, the likelihood may respond to the 'cosmic variance free' combination $C_\ell^{143\times143} + C_\ell^{217\times217} - 2C_\ell^{143\times217}$ rendering the χ^2 particularly sensitive to the accuracy of the noise modelling in these cases as the usually dominant cosmic variance is cancelled. Comparing this combination to the expected null spectrum in the multipole range 500–2000 we find a χ^2 high by 3.5σ, supporting this interpretation. If we subtract off a version of this spectrum smoothed on a scale of $\Delta\ell = 61$ we still find that χ^2 is high by 3.1σ. This suggests we are seeing the effects of slightly high scatter throughout the ℓ-range rather than a smooth trend that could be fit by an improved foreground model or a different cosmology.

At low multipoles (30–500) where we use only 143 × 143 the χ^2 is 2.7σ higher than expected, but as there is only one spectrum

Table 2. Chi-squared values for NPIPE polarization data. The rows give the chi-squared χ^2 and the number of standard deviations from unity $\Delta\sigma = (\chi^2/N_D - 1)/\sqrt{2/N_D}$ for the TE, EE, and combined TEEE data sets when compared to ΛCDM best-fitting models to the TT, TE, and EE data (columns).

	$\chi^2(\Delta\sigma)$		
Data	TT Theory	TE Theory	EE Theory
TE	2080.7 (1.75)	2078.3 (1.71)	2119.1 (2.36)
EE	2023.2 (0.83)	2027.7 (0.90)	2015.0 (0.70)
TEEE	4123.5 (2.04)	4124.9 (2.06)	4161.4 (2.47)

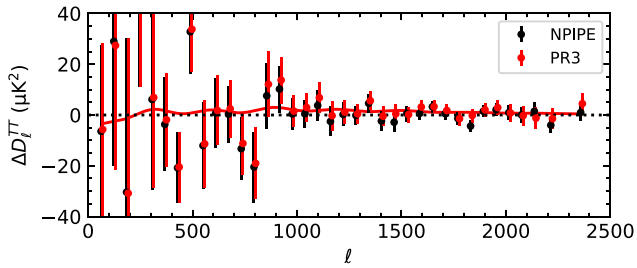

Figure 9. TT power spectrum residuals with respect to the PR4_12.6 TT best fit. The best-fitting foreground residuals have been subtracted from the data points, while calibrations have been applied to the theories. The best-fitting theory to PR3_12.6 TT is shown as a solid red line. The points are slightly offset for visual clarity.

this cannot be explained with the above argument. One possible explanation is the effect of excess variance introduced by the point-source masks that is not captured in the analytical covariance matrices. This effect is explained in depth in appendix C of Planck Collaboration XI (2016). A deviation of up to 10 per cent in the variance is found with the masks used there; from that we estimate that this effect could account for up to about 1σ of the high variance we see at $\ell < 500$. For this region the improvement in χ^2 is again marginal when a smoothed spectrum is subtracted, suggesting that relatively smooth cosmologies should not respond to this feature of the likelihood.

To further investigate the internal consistency of the NPIPE spectra, we can check how well the best ΛCDM fit to TT alone fits the TE and EE data. From Table 2 we find reduced chi-squareds $\hat{\chi}^2_{TE} = 1.055$ and $\hat{\chi}^2_{EE} = 1.026$ using the NPIPE TT best-fit solution. These values are 1.75σ and 0.83σ greater than 1, respectively, indicating that the TE and EE data are in acceptable agreement with the prediction of the TT model. We further compare these numbers to the best fits to the TE and EE data to understand how close the TT model is to the actual minima of the polarization likelihoods. If we assume the TT model represents the truth, then due to overfitting we expect the χ^2 of the TE and EE minima to improve over the 'true' (TT) model by a draw from a chi-squared distribution with $N_{\rm DOF}$ equal to the number of free parameters (see e.g. the appendix of Gratton & Challinor 2020). For TT-TE we find $\Delta\chi^2 = 2.4$, and for TT-EE $\Delta\chi^2 = 8.2$, roughly in line with the $\Delta\chi^2$ we expect for a seven-parameter model (six ΛCDM parameters and one free calibration) and therefore giving us confidence that the NPIPE polarization and temperature spectra are consistent.

Satisfied that our NPIPE spectra are internally consistent, we compare them directly to the PR3_12.6 spectra in Figs 9, 10, and 11. Beginning with temperature in Fig. 9, we find the co-added NPIPE TT spectrum to be up to 0.3 per cent low relative to PR3 at $200 \lesssim$

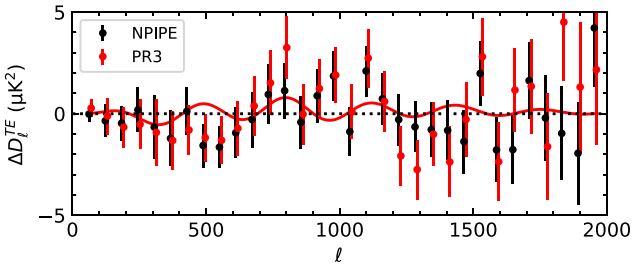

Figure 10. TE power spectrum residuals with respect to the PR4_12.6 TE best fit. The best-fitting theory to PR3_12.6 TE is shown as a solid red line.

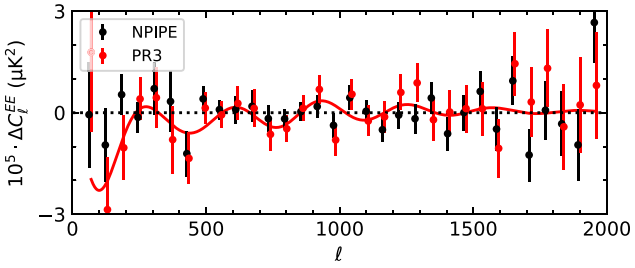

Figure 11. EE power spectrum residuals with respect to the PR4_12.6 EE best fit. The best-fitting theory to PR3_12.6 EE is shown as a solid red line.

$\ell \lesssim 1600$.[5] This is reflected in very small cosmological parameter shifts in TT (see Table A3). In TE (Fig. 10) we see good agreement between the spectra apart from a few exceptional points, and find similar ΛCDM best fits. The NPIPE residuals are slightly smaller, especially at $\ell > 1000$, but with a higher χ^2 due to the smaller error bars. There is somewhat more difference between the EE spectra (Fig. 11), with more large-scale power and a slight 'wavy' residual that leads to a preference for lower θ_* in NPIPE, as will be seen in Fig. 12 and Section 5.

To further quantify the level of agreement between the spectra we compare the PR4_12.6 ΛCDM best-fitting cosmology to the PR3_12.6 data and vice versa in Table 3. In doing so we leave the nuisance parameters in the likelihood free, as they are not necessarily expected to be the same for each likelihood. We find $\Delta\chi^2 = 2.5$ between the PR3 and PR4 TTTEEE cosmologies for the PR4 data and $\Delta\chi^2 = 1.4$ for the PR3 data, indicating similar best-fit cosmologies compatible with both data sets; the agreement is similarly good for TT and TE individually. For EE we have $\Delta\chi^2 = 8.8$ with PR4 data and $\Delta\chi^2 = 6.6$ for PR3. As with the visual comparison of Fig. 11, this again indicates some difference between PR3 and PR4 in EE.

5 BASE ΛCDM RESULTS

In this section, we describe our constraints on cosmological parameters in the base ΛCDM model using the NPIPE versions of the CamSpec likelihood. We use the standard set of six cosmological parameters for ΛCDM given in Table 4 with their priors. In addition to these cosmological parameters we have nuisance parameters $A_\nu^{\rm power}$ and $\gamma_\nu^{\rm power}$ for each TT spectrum, representing the amplitudes and spectral indices of power-law fits to remaining foreground residuals:

$$D_\ell^{\rm power} = A^{\rm power}\left(\frac{\ell}{1500}\right)^{\gamma^{\rm power}}. \quad (6)$$

These nuisance parameters are necessary to capture excesses from CIB and point sources in temperature that are not removed by 545 GHz cleaning. Any small foreground residuals remaining after power-law subtraction are ignored in the spectra and covariances. There are also three calibration parameters by which the theory is divided in the likelihood: the overall calibration $A^2_{\rm Planck}$ and individual polarization calibrations $C_{\rm TE}$ and $C_{\rm EE}$. Because all of our likelihoods are limited to $\ell \geq 30$ we supplement with low-ℓ ($2 \leq \ell \leq 29$) likelihoods from PR3; for all temperature likelihoods we include the Commander TT likelihood (Planck Collaboration IV 2020) and for polarization we add the SimAll EE likelihood (Planck Collaboration III 2020).

In Fig. 12, we compare parameter constraints from NPIPE TT, TE, and EE, again assessing agreement between the components. We find that the TT and TE constraints are highly consistent, and have similar constraining power; the EE constraints are broader but serve as a useful consistency check for *Planck*. Treating TT and EE as independent, we approximately quantify the consistency between their parameter constraints as follows. From our MCMC chains we compute a data vector of mean posterior values P and covariances C for our six sampled ΛCDM parameters. In this six-dimensional space we calculate $\Delta_P = P_{\rm TT} - P_{\rm EE}$ and $C_\Delta = C_{\rm TT} + C_{\rm EE}$, finding $\chi^2 = \Delta_P^{\rm T} C_\Delta^{-1} \Delta_P = 7.31$, indicating agreement within 0.5σ. Indeed, EE is very consistent for all parameters except the acoustic scale parameter θ_*, for which EE prefers a somewhat lower value (2.8σ lower than TT). This is reminiscent of the similar preference in PR3 EE for high n_s relative to TT and TE, which has disappeared in PR4_12.6. The θ_* tension is also slightly worse for extended models: for ΛCDM + A_L we have a 3.3σ difference, and for ΛCDM + Ω_K 3.2σ, due primarily to an increase in θ_* preferred by the TT spectrum in these models. Referring back to Fig. 12 however (see also Fig. 16), note that there is still substantial overlap of the θ_* contours; this is also true in the extended models. Given the consistency of NPIPE TT and TE on θ_*, which also agree very well with the values from PR3_12.6, we do not regard this shift as evidence for lower θ_*. Rather, it is more likely due to parameter degeneracies coupling to residual systematics in EE (the spectrum most likely to be affected by systematics), or perhaps a statistical fluctuation.

Moving on to the full TTTEEE likelihood, we present a comparison of ΛCDM parameter constraints from PR4_12.6 and PR3_12.6 in Table 5. Two-dimensional marginalized constraints, alongside those from the 2018 Plik likelihood and HiLLiPoP (see Section 7), are given in Fig. 13. We find consistency at the sigma level among all four likelihoods, for all cosmological parameters. Given the small weight carried by EE in a TTTEEE likelihood, note that any issues with EE or differences in their handling will have little impact on the overall cosmological constraints. The NPIPE likelihood contours are smaller than for PR3_12.6, and significantly more constraining than Plik. The widths of the one-dimensional 68 per cent constraints, as given in Table 5, shrink between 7 per cent and 14 per cent (14–21 per cent) compared to PR3_12.6 (Plik) TTTEEE for $\Omega_{\rm b}h^2$, $\Omega_{\rm c}h^2$, θ_*, $n_{\rm s}$, and H_0. Such decreases are consistent with the lower noise levels discussed in Section 3.5.

6 EXTENDED MODELS

We turn now to constraints of single-parameter extensions to the ΛCDM model. In particular we have varied curvature Ω_K, the

[5]The small offset is due to a smooth difference of 2–5 μK² in all three TT spectra in the given ℓ-range. Similar such differences are discussed in connection with effective calibrations in section 6.2 of EG21 for PR3 detector-set spectra.

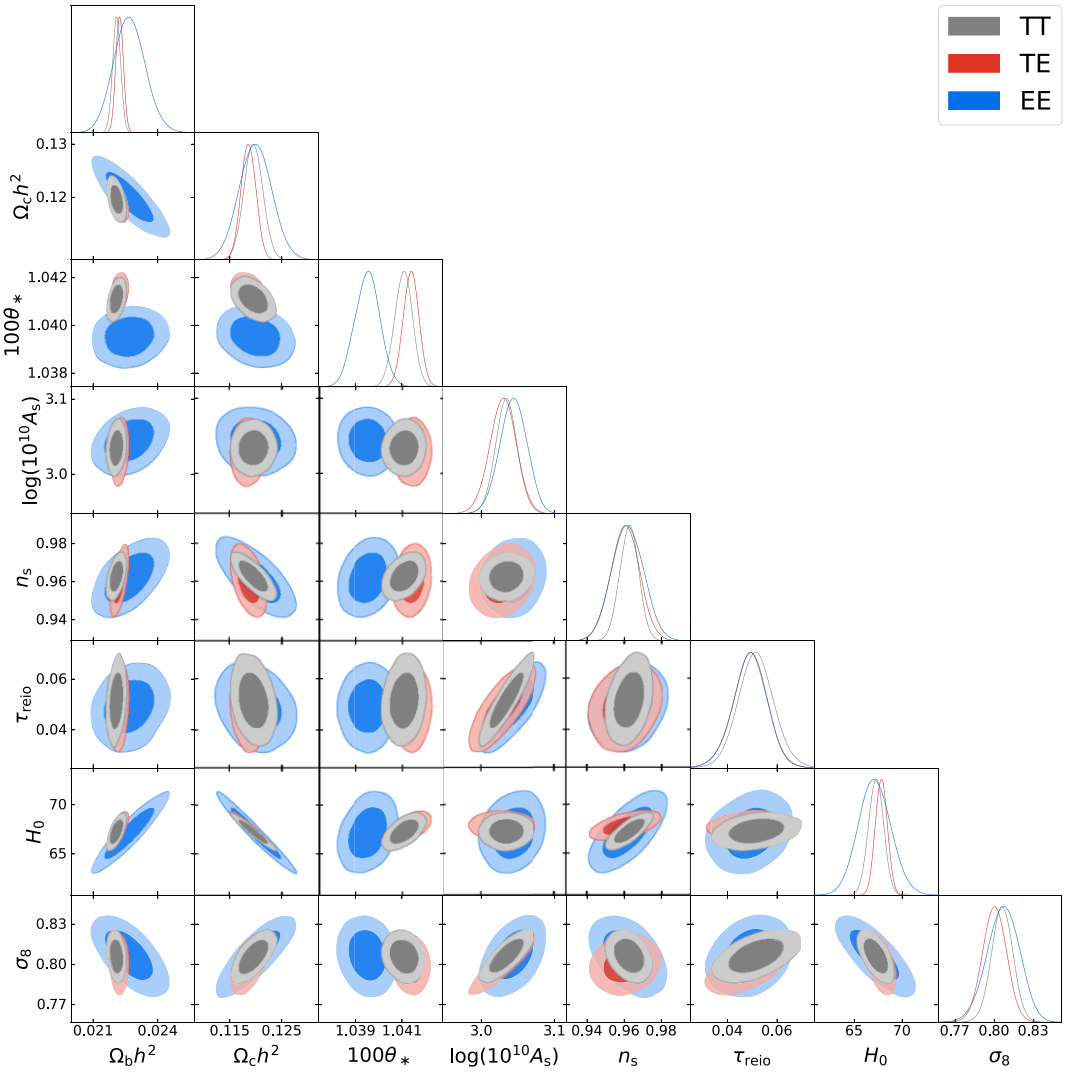

Figure 12. Comparison of 68 per cent and 95 per cent limits on cosmological parameters in ΛCDM from the NPIPE TT, TE, and EE likelihoods.

Table 3. χ^2 comparisons between PR4_12.6 (NP) and PR3_12.6. Each row gives the spectrum used for both the data and the best-fitting cosmology. The first column uses NPIPE for both data and theory, the second compares NPIPE data to the best-fitting theory from PR3, and so forth. Note that only cosmological parameters are kept fixed in the theory, while nuisance parameters are left free.

$\chi^2(\Delta\sigma)$	$NP_d - NP_{th}$	$NP_d - PR3_{th}$	$PR3_d - NP_{th}$	$PR3_d - PR3_{th}$
TT	6417.4 (4.066)	6417.8 (4.070)	6335.0 (3.312)	6334.9 (3.311)
TE	2078.3 (1.709)	2078.8 (1.717)	1965.0 (−0.096)	1964.4 (−0.105)
EE	2015.1 (0.702)	2023.9 (0.843)	1811.9 (−2.534)	1805.3 (−2.637)
TTTEEE	10543.1 (4.460)	10545.5 (4.477)	10116.1 (1.428)	10114.7 (1.418)

phenomenological A_L parameter, the effective number of light species $N_{\rm eff}$, and the sum of the neutrino masses m_ν one at a time. Marginalized constraints on each of these parameters are presented in Table 6, but we focus our discussion on Ω_K and A_L.

Starting with Ω_K we see a notable difference between NPIPE and PR3 polarization. The TT results hardly change, but in TE and especially EE we find with NPIPE a significant shrinking of the error bars, accompanied by a shift towards $\Omega_K = 0$, as illustrated in Fig. 14. In particular, the long tails allowing negative Ω_K are significantly reduced. Inspection of the 2D contours (not illustrated) do not suggest any obvious degeneracy between Ω_K and θ_*, so this shift in Ω_K appears independent from the low value of θ_* seen above in NPIPE EE. The maximum-likelihood point (best fit) of PR4_12.6 TTTEEE prefers $\Omega_K < 0$ at 1.2σ, similar to the 1.6σ preference in PR3_12.6 but now with a shift of the best fit from $\Omega_K = -0.026$ to -0.016 and a 20 per cent decrease in the size of the positive error bar. Similarly, the mean value of the posteriors prefers $\Omega_K < 0$ at 1.9σ in PR4_12.6 TTTEEE, down from 2.2σ in PR3 (see Table 6).

Turning to A_L, shown in Fig. 15, we see a similar trend to that for Ω_K. PR4_12.6 prefers a lower best-fit value of $A_L = 1.084$ compared

Table 4. Prior ranges on sampled parameters in both PR3_12.6 and PR4_12.6. Gaussian priors with mean μ and standard deviation σ are used for the calibrations, all other priors are flat. The TE and EE calibration parameters are not used in temperature-only likelihoods, and the foreground nuisance parameters $A_\nu^{\rm power}$ and $\gamma_\nu^{\rm power}$ are not used in polarization-only likelihoods.

	Prior limits
$\log(10^{10}A_{\rm s})$	1.61–3.91
$100\theta_{\rm MC}$	0.5–10
$n_{\rm s}$	0.8–1.2
$\Omega_{\rm b}h^2$	0.005–1
$\Omega_{\rm c}h^2$	0.001–0.99
$\tau_{\rm reio}$	0.01–0.8
$A_{\rm Planck}$	$\mu = 1; \sigma = 0.0025$
$c_{\rm TE}$	$\mu = 1; \sigma = 0.01$
$c_{\rm EE}$	$\mu = 1; \sigma = 0.01$
$A_{143}^{\rm power}$	0–50
$A_{217}^{\rm power}$	0–50
$A_{143{\rm x}217}^{\rm power}$	0–50
$\gamma_{143}^{\rm power}$	0–5
$\gamma_{217}^{\rm power}$	0–5
$\gamma_{143{\rm x}217}^{\rm power}$	0–5

Table 5. PR4_12.6 parameter constraints in ΛCDM, with a comparison to PR3_12.6. We report mean values and 68 per cent confidence intervals. Further entries are given in Table A3.

	PR4_12.6 TT	PR4_12.6 TTTEEE	PR3_12.6 TTTEEE
$\Omega_{\rm b}h^2$	0.02209 ± 0.00019	0.02218 ± 0.00013	0.02225 ± 0.00014
$\Omega_{\rm c}h^2$	0.1196 ± 0.0018	0.1197 ± 0.0011	0.1197 ± 0.0012
$100\,\theta_{\rm MC}$	1.04083 ± 0.00039	1.04075 ± 0.00024	1.04102 ± 0.00027
τ	0.0515 ± 0.0073	0.0517 ± 0.0072	0.0533 ± 0.0074
$\ln(10^{10}A_{\rm s})$	3.034 ± 0.015	3.035 ± 0.015	3.041 ± 0.015
$n_{\rm s}$	0.9626 ± 0.0052	0.9635 ± 0.0039	0.9667 ± 0.0042
H_0	67.26 ± 0.79	67.26 ± 0.49	67.40 ± 0.54
$100\,\theta_*$	1.04107 ± 0.00038	1.04098 ± 0.00024	1.04123 ± 0.00027
Ω_Λ	0.685 ± 0.011	0.6849 ± 0.0068	0.6859 ± 0.0076
$\Omega_{\rm m}$	0.315 ± 0.011	0.3150 ± 0.0068	0.3140 ± 0.0076
$\Omega_{\rm m}h^2$	0.1423 ± 0.0017	0.1425 ± 0.0011	0.1426 ± 0.0012
σ_8	0.8060 ± 0.0082	0.8067 ± 0.0067	0.8098 ± 0.0070
$\sigma_8\Omega_{\rm m}^{0.5}$	0.452 ± 0.011	0.4528 ± 0.0072	0.4538 ± 0.0079
$\sigma_8\Omega_{\rm m}^{0.25}$	0.604 ± 0.010	0.6043 ± 0.0070	0.6062 ± 0.0075
$z_{\rm re}$	7.42 ± 0.76	$7.43^{+0.75}_{-0.68}$	7.58 ± 0.75
$10^9A_{\rm s}$	2.079 ± 0.032	2.081 ± 0.031	2.092 ± 0.032
$10^9A_{\rm s}e^{-2\tau}$	1.875 ± 0.013	1.877 ± 0.011	1.880 ± 0.011
Age/Gyr	13.826 ± 0.031	13.820 ± 0.020	13.806 ± 0.022
$r_{\rm drag}$	147.52 ± 0.43	147.40 ± 0.25	147.31 ± 0.28
$A_{\rm Planck}$	1.0004 ± 0.0024	1.0004 ± 0.0024	1.0005 ± 0.0024
$c_{\rm TE}$	–	0.9975 ± 0.0036	1.0005 ± 0.0041
$c_{\rm EE}$	–	0.9975 ± 0.0037	1.0013 ± 0.0043
$A_{143}^{\rm power}$	$19.1^{+2.0}_{-2.8}$	$18.1^{+1.7}_{-2.4}$	$17.2^{+1.9}_{-2.4}$
$A_{143\times217}^{\rm power}$	$9.8^{+1.9}_{-2.7}$	$8.8^{+1.6}_{-2.3}$	$8.0^{+1.7}_{-2.4}$
$A_{217}^{\rm power}$	$12.9^{+1.9}_{-2.7}$	$11.9^{+1.6}_{-2.3}$	$12.8^{+1.7}_{-2.4}$
$\gamma_{143}^{\rm power}$	$0.97^{+0.23}_{-0.17}$	$1.01^{+0.20}_{-0.17}$	$1.12^{+0.28}_{-0.24}$
$\gamma_{143\times217}^{\rm power}$	$1.54^{+0.53}_{-0.61}$	$1.72^{+0.54}_{-0.61}$	$1.53^{+0.51}_{-0.74}$
$\gamma_{217}^{\rm power}$	1.39 ± 0.42	1.52 ± 0.41	1.40 ± 0.39

to $A_L = 1.141$ from PR3_12.6, and is now 1.5σ greater than $A_L = 1$ rather than 2.3σ. The means (Table 6) find $A_L > 1$ by 1.7σ (PR4) and 2.4σ (PR3).[6] The constraint from EE in particular has shifted to be centred at $A_L = 1$, while TT and TE show smaller downward shifts with slightly smaller error bars relative to PR3_12.6. Again, there does not appear to be a particular degeneracy between A_L and θ_*. To better test this we rerun cosmological parameter estimation with θ_* fixed to its PR3_12.6 best-fit value, and find that the NPIPE A_L constraints are effectively unchanged. This gives us confidence that the preference for excess smoothing of the power spectra represented by A_L really has decreased in the NPIPE maps, independent of other parameter shifts.

7 EXTERNAL COMPARISONS

7.1 HiLLiPoP

A limited analysis of the NPIPE maps using the HiLLiPoP[7] likelihood (Couchot et al. 2017) and focusing on the tensor-to-scalar ratio r has been published in Tristram et al. (2022). HiLLiPoP is also a pseudo-C_ℓ likelihood and uses NPIPE data products. However, the likelihood differs from ours in choice of masks, spectra, and foreground mitigation approach. Instead of cleaning spectra, HiLLiPoP attempts to fit parametric models to uncleaned temperature and polarization spectra, including those using 100 GHz data. Tristram et al. (2022) find high-ℓ TT to give essentially identical results to PR3; full high-ℓ TE and EE results are not given there due to that paper's focus on r. We use the public HiLLiPoP likelihood[8] with all default prior ranges on nuisance and foreground parameters together with our own priors on cosmological parameters to compare to our CamSpec likelihood. Note that we continue to use Commander and SimAll at low multipoles. These results are shown alongside Plik and our baseline CamSpec and NPIPE constraints in Fig. 13. We find that HiLLiPoP favours an even lower θ_* than PR4_12.6; this is driven by polarization as shown in Fig. 16, while the TT-only constraints (not shown) are similar to those of PR4_12.6. Otherwise the HiLLiPoP constraints do generally agree with our findings at the sigma level.

7.2 ACT DR4 and SPT-3G

We present a simple comparison of our results to the latest ground-based data, ACT DR4 and SPT-3G. More detailed comparisons to plik are presented in the ACT (Aiola et al. 2020; Choi et al. 2020) and SPT (Dutcher et al. 2021) papers, as well as in Handley & Lemos (2021), with a comparison to CamSpec in the appendix of EG21.

For ACT each region, array, and season is independently calibrated to *Planck* PR3 temperature by cross-correlation. There is no separate polarization calibration; instead there is a free effective polarization efficiency in the likelihood. We compare to co-added CMB-only spectra with foregrounds and TP leakage corrections marginalized out. SPT also uses *Planck* PR3 cross-spectra to calibrate temperature of each sub-field before co-addition. They calibrate the final co-added spectra again to both TT and EE, with TE fixed. Foregrounds

[6]Considering the data entering PR3_12.6 as a subset of PR4_12.6 then one can apply the method of Gratton & Challinor (2020) to estimate the expected standard deviation of the shifts to be $\sqrt{0.061^2 - 0.056^2} = 0.024$, putting the difference of the means at 2.1σ.

[7]https://github.com/planck-npipe/hillipop

[8]Note that all references to HiLLiPoP in this paper use NPIPE data products.

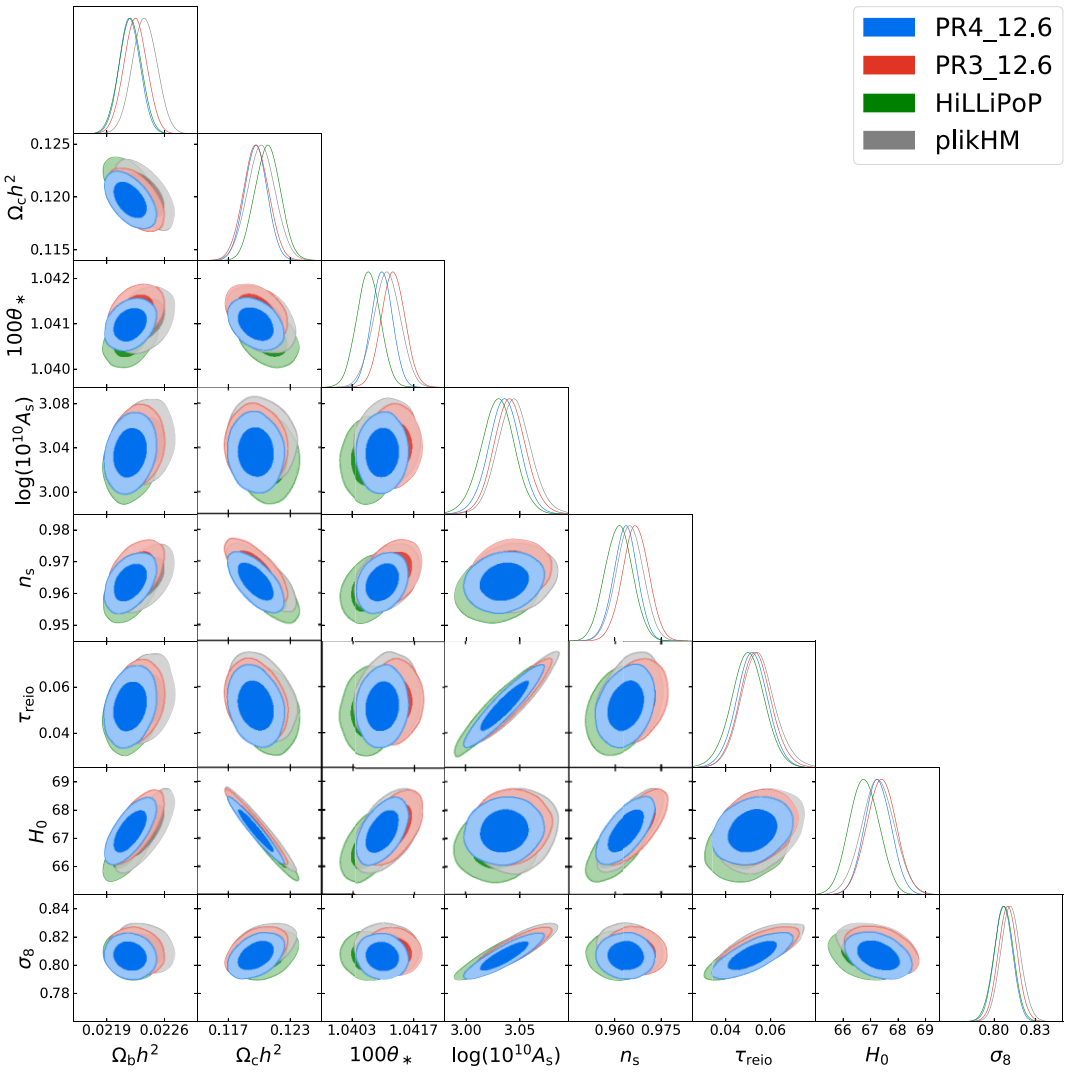

Figure 13. 2D 68 per cent and 95 per cent parameter constraints on cosmological parameters in ΛCDM using the PR4_12.6, PR3_12.6, `HiLLiPoP` (PR4), and `Plik` (PR3) TTTEEE likelihoods. Note that the scales are zoomed in relative to Fig. 12.

Table 6. Mean values and 68 per cent limits for beyond-ΛCDM parameters.

PR4_12.6	A_L	Ω_K	$N_{\rm eff}$	m_ν
TTTEEE	1.095 ± 0.056	$-0.025^{+0.013}_{-0.010}$	3.00 ± 0.21	<0.161
TT	1.198 ± 0.084	$-0.042^{+0.022}_{-0.016}$	$2.98^{+0.28}_{-0.35}$	<0.278
TE	0.96 ± 0.15	$-0.010^{+0.035}_{-0.015}$	$3.11^{+0.38}_{-0.42}$	<0.400
EE	0.995 ± 0.15	$-0.012^{+0.034}_{-0.017}$	4.6 ± 1.3	<2.37
PR3_12.6	A_L	Ω_K	$N_{\rm eff}$	m_ν
TTTEEE	1.146 ± 0.061	$-0.035^{+0.016}_{-0.012}$	$2.94^{+0.20}_{-0.23}$	<0.143
TT	1.215 ± 0.089	$-0.047^{+0.024}_{-0.017}$	$2.89^{+0.28}_{-0.32}$	<0.248
TE	0.96 ± 0.17	$-0.015^{+0.043}_{-0.015}$	$2.96^{+0.42}_{-0.49}$	<0.504
EE	1.15 ± 0.20	$-0.053^{+0.063}_{-0.029}$	$2.46^{+0.94}_{-1.7}$	–

are not removed from the polarization spectra, but are reported to be negligibly small (Dutcher et al. 2021).

We show power spectra and residuals comparing our `NPIPE` likelihood PR4_12.6 to ACT DR4 and SPT-3G in Figs 17, 18, and 19. Table 7 presents χ^2 values for the ACT and SPT data using the PR3_12.6 and PR4_12.6 best-fitting ΛCDM models.[9] In TT we find generally good agreement between ACT and *Planck* but see that ACT Wide is consistently high by about 10–20 μK^2 in D_ℓ from ℓ of about 1000 to 1700. Both ACT Deep and Wide are also a bit low relative to *Planck* at very high multipoles (above 3500). These differences are reflected in somewhat high χ^2 values in TT, slightly worse with the `NPIPE` best fit than with that of PR3. We interpret this as due to the calibration of ACT DR4 to the PR3 spectra. Indeed, allowing a free calibration parameter decreases the reduced χ^2 in TT to 1.30 (PR4) and 1.28 (PR3), confirming that calibration explains most of the difference.

In TE, like TT, the `NPIPE` spectra are so similar to those from PR3 that previous comparisons to ground-based experiments generally still apply. By eye ACT TE appears to agree well with `NPIPE`. However, we note that Choi et al. (2020) found significant

[9] χ^2 values are calculated using the public ACT likelihood (https://github.com/ACTCollaboration/pyactlike) and a `python` port (https://github.com/xgarrido/spt_likelihoods) of the public SPT-3G likelihood.

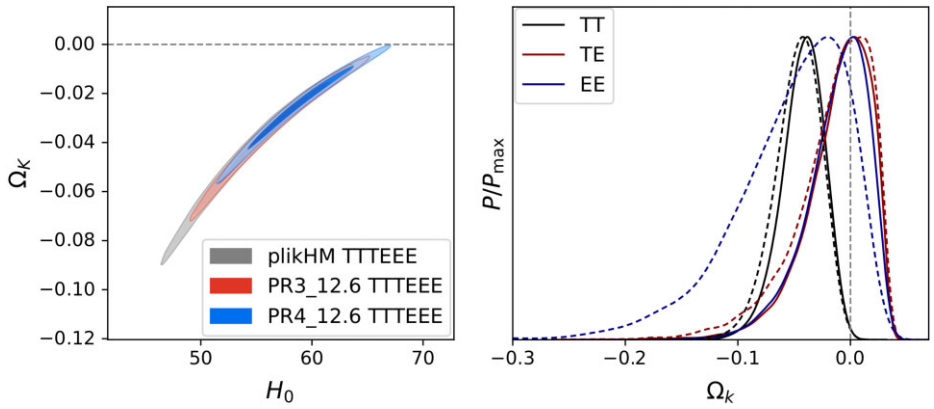

Figure 14. Left: 68 per cent and 95 per cent confidence limits in the $\Omega_K - H_0$ plane from the `Plik`, PR3_12.6, and PR4_12.6 TTTEEE likelihoods. Right: 1D marginalized densities on Ω_K from TT, TE, and EE. Constraints from `NPIPE` are shown with solid lines, and PR3_12.6 with dashed lines.

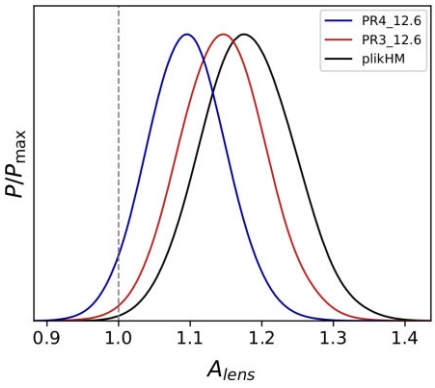

Figure 15. Marginalized constraints on A_L from the PR4_12.6, PR3_12.6, and `Plik` TTTEEE likelihoods.

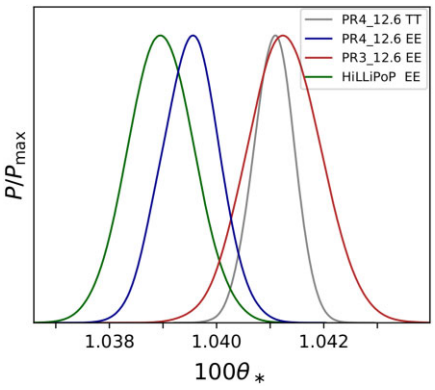

Figure 16. Marginalized constraints on θ_* from PR4_12.6 TT and the PR4_12.6, PR3_12.6, and `HiLLiPoP` (PR4) EE likelihoods. Both `NPIPE` EE likelihoods prefer lower values of θ_* relative to PR3 EE or TT.

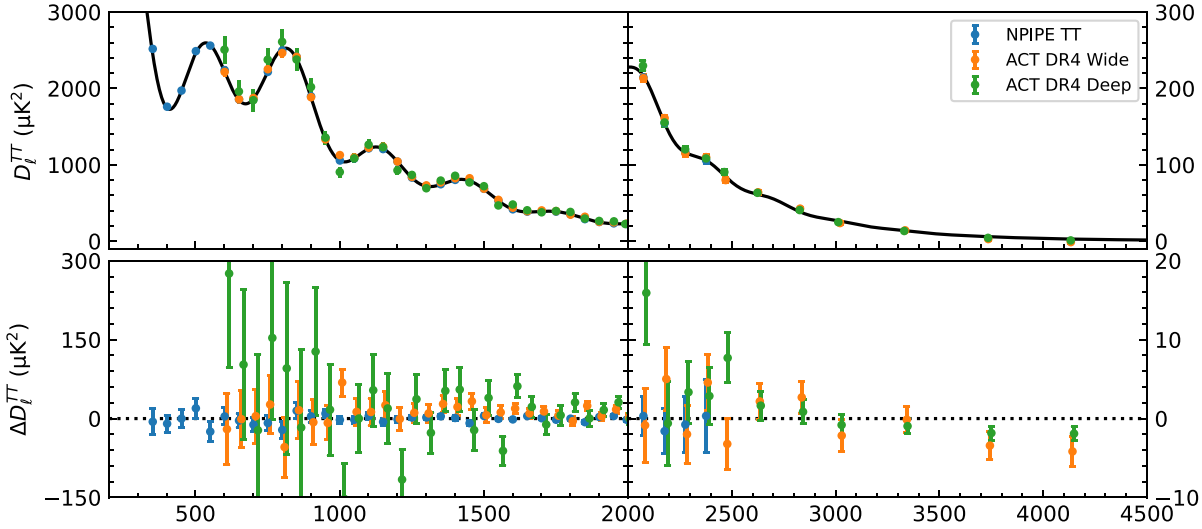

Figure 17. PR4_12.6 and ACT DR4 TT power spectra. `NPIPE` spectra are foreground-subtracted and binned using the ACT Wide window functions. The lower panels show residuals to the best-fitting theory spectrum for PR4_12.6 TT. At multipoles higher than 2000 we use an expanded scale to better show differences between the spectra.

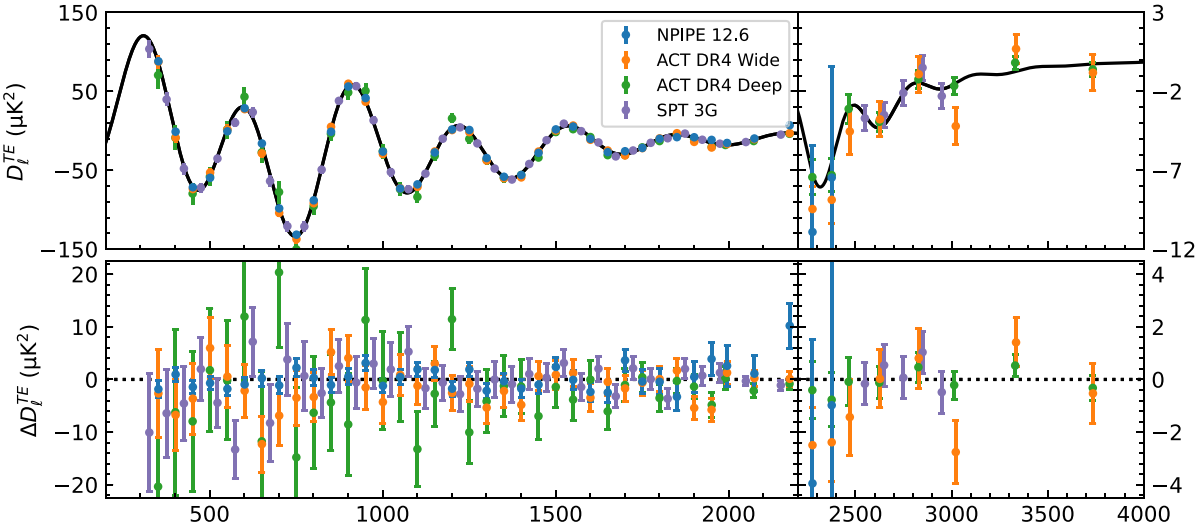

Figure 18. PR4_12.6, ACT DR4, and SPT-3G TE power spectra. The lower panels show residuals to the best-fitting theory spectrum for PR4_12.6 TE, binned using the appropriate window function. NPIPE is binned as ACT Wide. As in Fig. 17 the scale is expanded at high ℓ.

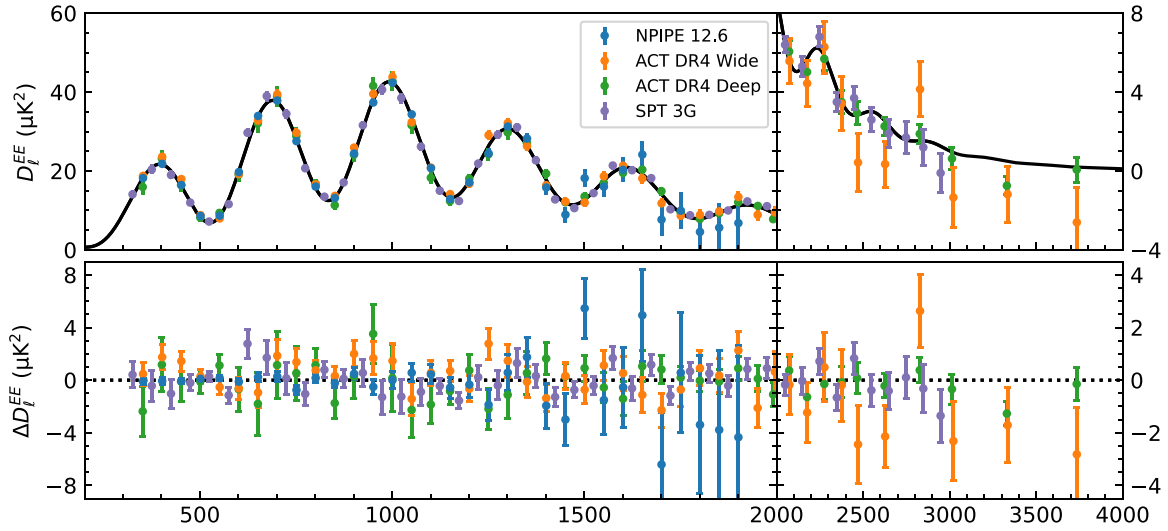

Figure 19. As Fig. 18 for EE.

Table 7. Reduced chi-squared $\hat{\chi}^2$ and uncertainties for ACT and SPT data (rows) with ΛCDM best fits from *Planck*, ACT, and SPT. The spectra used for each model match the data in all cases except the SPT best fits, which are always TEEE. Multipole ranges are as in the likelihood for each experiment.

Data	PR4_12.6	PR3_12.6	ACT DR4	SPT-3G TEEE
ACT TTTEEE	1.21 (2.43)	1.16 (1.87)	1.08 (0.88)	1.22 (2.51)
ACT TT	1.57 (3.62)	1.48 (3.01)	1.18 (1.15)	1.37 (2.33)
ACT TEEE	1.10 (0.92)	1.09 (0.82)	0.99 (−0.08)	1.19 (1.79)
SPT TEEE	0.98 (−0.34)	0.98 (−0.32)	1.02 (0.29)	0.97 (−0.53)

disagreement between PR3 and a subset of ACT TE spectra. There are no obvious features causing this disagreement, and further work is needed to identify its cause. On the SPT side, there is a run of low points at $\ell < 700$, while higher multipoles appear to agree well. In EE ACT DR4 Wide is consistently about 5 μK^2 low for $\ell > 2500$; otherwise all of the spectra agree well. The TEEE χ^2 is very good, being within 1σ for both ACT and SPT. Overall, we find, as for PR3, that the three experiments are broadly consistent but a more careful analysis with detailed map-based relative calibration would be necessary for a more quantitative comparison.

8 CONCLUSIONS

In this work, we have presented new high-ℓ CMB temperature and polarization power spectra and likelihoods using the *Planck* PR4 'NPIPE' maps. We find excellent consistency between the PR4 and PR3 power spectra, which translates to very good agreement on cosmological parameters as well. The most significant difference is in EE, for which NPIPE prefers a lower value of θ_*; however, we find that the PR3_12.6 best-fitting ΛCDM model is still a good fit to the NPIPE data.

The lower noise of the NPIPE maps leads to tighter parameter constraints, with a ∼10 per cent improvement in most ΛCDM

parameters in TTTEEE due primarily to improvements in polarization. For ΛCDM extensions we find that, relative to PR3, NPIPE polarization shrinks the error bars on Ω_K and A_L from EE by 40 per cent and 25 per cent, respectively, and by 15 per cent and 8 per cent in TTTEEE. The combination of smaller error bars and shifts towards the ΛCDM values continues the trend observed in EG21 of decreasing the Ω_K and A_L tensions as more data are used, as would be expected if these pulls were due to a statistical fluctuation. Overall, we conclude that NPIPE, despite substantial differences in the mapmaking, is completely consistent with PR3 at the power spectrum and parameter levels, offering yet more evidence of the robustness of *Planck* data and the cosmological results inferred from *Planck*.

ACKNOWLEDGEMENTS

We are grateful to the NPIPE team and the wider *Planck* Collaboration for the data products used in this work. We thank Reijo Keskitalo for helpful correspondence and thank Anthony Challinor and John Peacock for useful comments on earlier drafts of this paper. Thanks also to the referee for helpful suggestions on the submitted manuscript.

ER acknowledges support from an Isaac Newton Studentship. SG acknowledges the award of a Kavli Institute Fellowship at KICC.

We gratefully acknowledge use of the following software packages: PSPIPE (https://github.com/simonsobs/PSpipe), PSPY (https://github.com/simonsobs/pspy), CAMB (Lewis, Challinor & Lasenby 2000; Howlett et al. 2012), COBAYA (Powell 2009; Cartis et al. 2018a; Cartis, Roberts & Sheridan-Methven 2018b; Torrado & Lewis 2021), GETDIST (Lewis 2019), the MCMC sampler of Lewis & Bridle (2002) and Lewis (2013) using the fast-dragging procedure of Neal (2005), SCIPY (Virtanen et al. 2020), NUMPY (Harris et al. 2020), ASTROPY (Astropy Collaboration 2013, 2018), MATPLOTLIB (Hunter 2007), HEALPY (Zonca et al. 2019), and HEALPIX (Górski et al. 2005).

This work was performed in part using resources provided by the Cambridge Service for Data Driven Discovery (CSD3) operated by the University of Cambridge Research Computing Service (www.csd3.cam.ac.uk), provided by Dell EMC and Intel using Tier-2 funding from the Engineering and Physical Sciences Research Council (capital grant EP/T022159/1), and DiRAC funding from the Science and Technology Facilities Council (www.dirac.ac.uk). This research also used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231.

DATA AVAILABILITY

The *Planck* and NPIPE data products underlying the likelihoods presented here are publicly available at the PLA (http://pla.esac.esa.int/pla/) and at NERSC. The PR3_12.6 and PR4_12.6 likelihoods are both publicly available as internal likelihoods to Cobaya (https://github.com/CobayaSampler/cobaya), named planck_2018_highl_CamSpec2021 (PR3) and planck_NPIPE_highl_CamSpec (PR4).

REFERENCES

Aiola S. et al., 2020, J. Cosmol. Astropart. Phys., 2020, 047
Astropy Collaboration, 2013, A&A, 558, A33
Astropy Collaboration, 2018, AJ, 156, 123
Cartis C., Fiala J., Marteau B., Roberts L., 2018a, preprint (arXiv:1804.00154)
Cartis C., Roberts L., Sheridan-Methven O., 2018b, preprint (arXiv:1812.11343)
Choi S. K. et al., 2020, J. Cosmol. Astropart. Phys., 2020, 045
Couchot F., Henrot-Versillé S., Perdereau O., Plaszczynski S., Rouillé d'Orfeuil B., Spinelli M., Tristram M., 2017, A&A, 602, A41
Delouis J. M., Pagano L., Mottet S., Puget J. L., Vibert L., 2019, A&A, 629, A38
Dutcher D. et al., 2021, Phys. Rev. D, 104, 022003
Efstathiou G., Gratton S., 2021, Open J. Astrophys., 4, 8
Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759
Gratton S., Challinor A., 2020, MNRAS, 499, 3410
Handley W., Lemos P., 2021, Phys. Rev. D, 103, 063529
Harris C. R. et al., 2020, Nature, 585, 357
Hivon E., Górski K. M., Netterfield C. B., Crill B. P., Prunet S., Hansen F., 2002, ApJ, 567, 2
Hivon E., Mottet S., Ponthieu N., 2017, A&A, 598, A25
Howlett C., Lewis A., Hall A., Challinor A., 2012, J. Cosmol. Astropart. Phys., 1204, 027
Hunter J. D., 2007, Comput. Sci. Eng., 9, 90
Keihänen E., Kurki-Suonio H., Poutanen T., Maino D., Burigana C., 2004, A&A, 428, 287
Keihänen E., Kurki-Suonio H., Poutanen T., 2005, MNRAS, 360, 390
Keihänen E., Keskitalo R., Kurki-Suonio H., Poutanen T., Sirviö A., 2010, A&A, 510, A57
Kogut A. et al., 2003, ApJS, 148, 161
Lewis A., 2013, Phys. Rev. D, 87, 103529
Lewis A., 2019, preprint (arXiv:1910.13970)
Lewis A., Bridle S., 2002, Phys. Rev. D, 66, 103511
Lewis A., Challinor A., Lasenby A., 2000, ApJ, 538, 473
Neal R. M., 2005, preprint (arXiv:e-prints/math/0502099)
Planck Collaboration I, 2014, A&A, 571, A1
Planck Collaboration XV, 2014, A&A, 571, A15
Planck Collaboration I, 2016, A&A, 594, A1
Planck Collaboration VII, 2016, A&A, 594, A7
Planck Collaboration XI, 2016, A&A, 594, A11
Planck Collaboration I, 2020, A&A, 641, A1
Planck Collaboration II, 2020, A&A, 641, A2
Planck Collaboration III, 2020, A&A, 641, A3
Planck Collaboration IV, 2020, A&A, 641, A4
Planck Collaboration V, 2020, A&A, 641, A5
Planck Collaboration VI, 2020, A&A, 641, A6
Planck Collaboration LVII, 2020, A&A, 643, 42
Planck Collaboration XLVI, 2016, A&A, 596, A107
Powell J. D. M., 2009, Technical Report, Department of Applied Mathematics and Theoretical Physics, University of Cambridge.
Torrado J., Lewis A., 2021, J. Cosmol. Astropart. Phys., 2021, 057
Tristram M. et al., 2022, Phys. Rev. D, 105, 083524
Virtanen P. et al., 2020, Nat. Methods, 17, 261
Zonca A., Singer L., Lenz D., Reinecke M., Rosset C., Hivon E., Gorski K., 2019, J. Open Source Softw., 4, 1298

APPENDIX A: SUPPLEMENTARY TABLES AND FIGURES

This appendix contains the following supplementary information: polarization spectrum calibration factors (Table A1), description of

Table A1. Calibration factors measured from the NPIPE spectra. Each co-added TT spectrum is multiplied by the given factor upon incorporation into the likelihood. Each TE and EE spectrum is divided (not multiplied) by the relevant factor before coaddition into the single TE and EE spectra of the likelihoods. Indices 1–6 correspond to 100A, 100B, 143A, 143B, 217A, and 217B respectively.

Spectrum	TT	c_k
143 × 217		0.9992 ± 0.0004
217 × 217		0.9988 ± 0.0005
Spectrum	**EE index**	c_k
100A × 100B	(1,2)	0.988 ± 0.007
100A × 143B	(1,4)	0.987 ± 0.006
100A × 217B	(1,6)	0.970 ± 0.007
100B × 143A	(2,3)	1.008 ± 0.006
100B × 217A	(2,5)	0.955 ± 0.007
143A × 143B	(3,4)	1.016 ± 0.005
143A × 217B	(3,6)	0.975 ± 0.006
143B × 217A	(4,5)	0.969 ± 0.006
217A × 217B	(5,6)	0.941 ± 0.007
Spectrum	**TE index**	c_k
100A × 100B	(1,2)	0.985 ± 0.009
100A × 143B	(1,4)	1.002 ± 0.008
100A × 217B	(1,6)	0.985 ± 0.009
100B × 143A	(2,3)	1.010 ± 0.008
100B × 217A	(2,5)	0.967 ± 0.009
143A × 143B	(3,4)	0.997 ± 0.008
143A × 217B	(3,6)	0.980 ± 0.009
143B × 217A	(4,5)	0.966 ± 0.009
217A × 217B	(5,6)	0.985 ± 0.009
100B × 100A	(2,1)	0.980 ± 0.010
143B × 100A	(4,1)	0.984 ± 0.010
217B × 100A	(6,1)	0.985 ± 0.010
143A × 100B	(3,2)	0.983 ± 0.009
217A × 100B	(5,2)	0.985 ± 0.009
143B × 143A	(4,3)	1.009 ± 0.008
217B × 143A	(6,3)	1.015 ± 0.008
217A × 143B	(5,4)	0.998 ± 0.008
217B × 217A	(6,5)	0.972 ± 0.009

the multipole ranges used in the CamSpec likelihoods (Table A2), HRD cross-spectra for NPIPE (Fig. A1), and ΛCDM parameter constraints for PR4_12.6 and PR3_12.6 TT, TE, and EE likelihoods (Table A3).

Table A2. Multipole ranges used in the likelihoods described in this paper.

	TT	TE	EE
100 × 100	–	30–1200	200–1200
100 × 143	–	30–1500	30–1500
100 × 217	–	200–1500	200–1200
143 × 143	30–2000	30–2000	200–2000
143 × 217	500–2500	200–2000	300–2000
217 × 217	500–2500	500–2000	500–2000
Total	30–2500	30–2000	30–2000

Table A3. PR4_12.6 and PR3_12.6 parameter constraints in ΛCDM. We report mean values and 68 per cent confidence intervals. TTTEEE is shown in Table 5.

	PR4_12.6 TT	PR4_12.6 TE	PR4_12.6 EE	PR3_12.6 TT	PR3_12.6 TE	PR3_12.6 EE
$\Omega_b h^2$	0.02209 ± 0.00019	0.02223 ± 0.00017	0.02269 ± 0.00073	0.02215 ± 0.00020	0.02221 ± 0.00019	0.02316 ± 0.00093
$\Omega_c h^2$	0.1196 ± 0.0018	0.1186 ± 0.0014	0.1200 ± 0.0031	0.1193 ± 0.0019	0.1191 ± 0.0016	0.1161 ± 0.0038
$100\,\theta_{MC}$	1.04083 ± 0.00039	1.04117 ± 0.00034	1.03934 ± 0.00055	1.04073 ± 0.00041	1.04142 ± 0.00040	1.04114 ± 0.00071
τ	0.0515 ± 0.0073	$0.0490^{+0.0071}_{-0.0064}$	0.0491 ± 0.0071	0.0523 ± 0.0075	0.0490 ± 0.0071	0.0525 ± 0.0074
$\ln(10^{10} A_s)$	3.034 ± 0.015	3.029 ± 0.019	3.044 ± 0.019	3.036 ± 0.015	3.025 ± 0.019	3.040 ± 0.020
n_s	0.9626 ± 0.0052	0.9606 ± 0.0078	0.9619 ± 0.0086	0.9654 ± 0.0054	0.9581 ± 0.0089	0.985 ± 0.012
H_0	67.26 ± 0.79	67.82 ± 0.62	67.1 ± 1.7	67.38 ± 0.83	67.71 ± 0.72	69.5 ± 2.2
$100\,\theta_*$	1.04107 ± 0.00038	1.04139 ± 0.00034	1.03951 ± 0.00054	1.04096 ± 0.00040	1.04164 ± 0.00040	1.04126 ± 0.00069
Ω_Λ	0.685 ± 0.011	0.6921 ± 0.0085	$0.681^{+0.023}_{-0.020}$	0.687 ± 0.011	0.6901 ± 0.0098	$0.709^{+0.026}_{-0.021}$
Ω_m	0.315 ± 0.011	0.3078 ± 0.0085	$0.319^{+0.020}_{-0.023}$	0.313 ± 0.011	0.3099 ± 0.0098	$0.291^{+0.021}_{-0.026}$
$\Omega_m h^2$	0.1423 ± 0.0017	0.1415 ± 0.0014	0.1433 ± 0.0025	0.1421 ± 0.0018	0.1420 ± 0.0016	0.1399 ± 0.0031
σ_8	0.8060 ± 0.0082	0.8001 ± 0.0093	0.807 ± 0.012	0.8064 ± 0.0083	0.7997 ± 0.0096	0.798 ± 0.015
$\sigma_8 \Omega_m^{0.5}$	0.452 ± 0.011	0.4438 ± 0.0096	0.455 ± 0.021	0.451 ± 0.012	0.445 ± 0.011	$0.430^{+0.024}_{-0.026}$
$\sigma_8 \Omega_m^{0.25}$	0.604 ± 0.010	0.5959 ± 0.0095	0.606 ± 0.018	0.603 ± 0.011	0.597 ± 0.010	0.586 ± 0.022
z_{re}	7.42 ± 0.76	$7.12^{+0.78}_{-0.63}$	$7.05^{+0.75}_{-0.63}$	7.49 ± 0.76	$7.13^{+0.78}_{-0.65}$	7.25 ± 0.72
$10^9 A_s$	2.079 ± 0.032	2.069 ± 0.038	2.099 ± 0.039	2.083 ± 0.032	2.061 ± 0.038	2.092 ± 0.042
$10^9 A_s e^{-2\tau}$	1.875 ± 0.013	1.876 ± 0.022	1.902 ± 0.024	1.876 ± 0.013	1.868 ± 0.023	1.883 ± 0.027
Age/Gyr	13.826 ± 0.031	13.798 ± 0.026	13.812 ± 0.091	13.822 ± 0.033	13.794 ± 0.029	13.69 ± 0.12
r_{drag}	147.52 ± 0.43	147.62 ± 0.36	146.76 ± 0.47	147.54 ± 0.43	147.50 ± 0.40	147.28 ± 0.67
A_{Planck}	1.0004 ± 0.0024	0.99998 ± 0.0025	1.0000 ± 0.0024	1.0003 ± 0.0024	1.0000 ± 0.0025	1.0000 ± 0.0024
c_{TE}	–	0.9997 ± 0.0092	–	–	0.9997 ± 0.0091	–
c_{EE}	–	–	1.0000 ± 0.0091	–	–	1.0004 ± 0.0091
A^{power}_{143}	$19.1^{+2.0}_{-2.8}$	–	–	$18.4^{+2.2}_{-2.9}$	–	–
$A^{power}_{143\times217}$	$9.8^{+1.9}_{-2.7}$	–	–	$9.2^{+2.1}_{-2.9}$	–	–
A^{power}_{217}	$12.9^{+1.9}_{-2.7}$	–	–	$14.0^{+2.1}_{-2.9}$	–	–
γ^{power}_{143}	$0.97^{+0.23}_{-0.17}$	–	–	$1.05^{+0.31}_{-0.25}$	–	–
$\gamma^{power}_{143\times217}$	$1.54^{+0.53}_{-0.61}$	–	–	$1.32^{+0.50}_{-0.72}$	–	–
γ^{power}_{217}	1.39 ± 0.42	–	–	1.27 ± 0.41	–	–

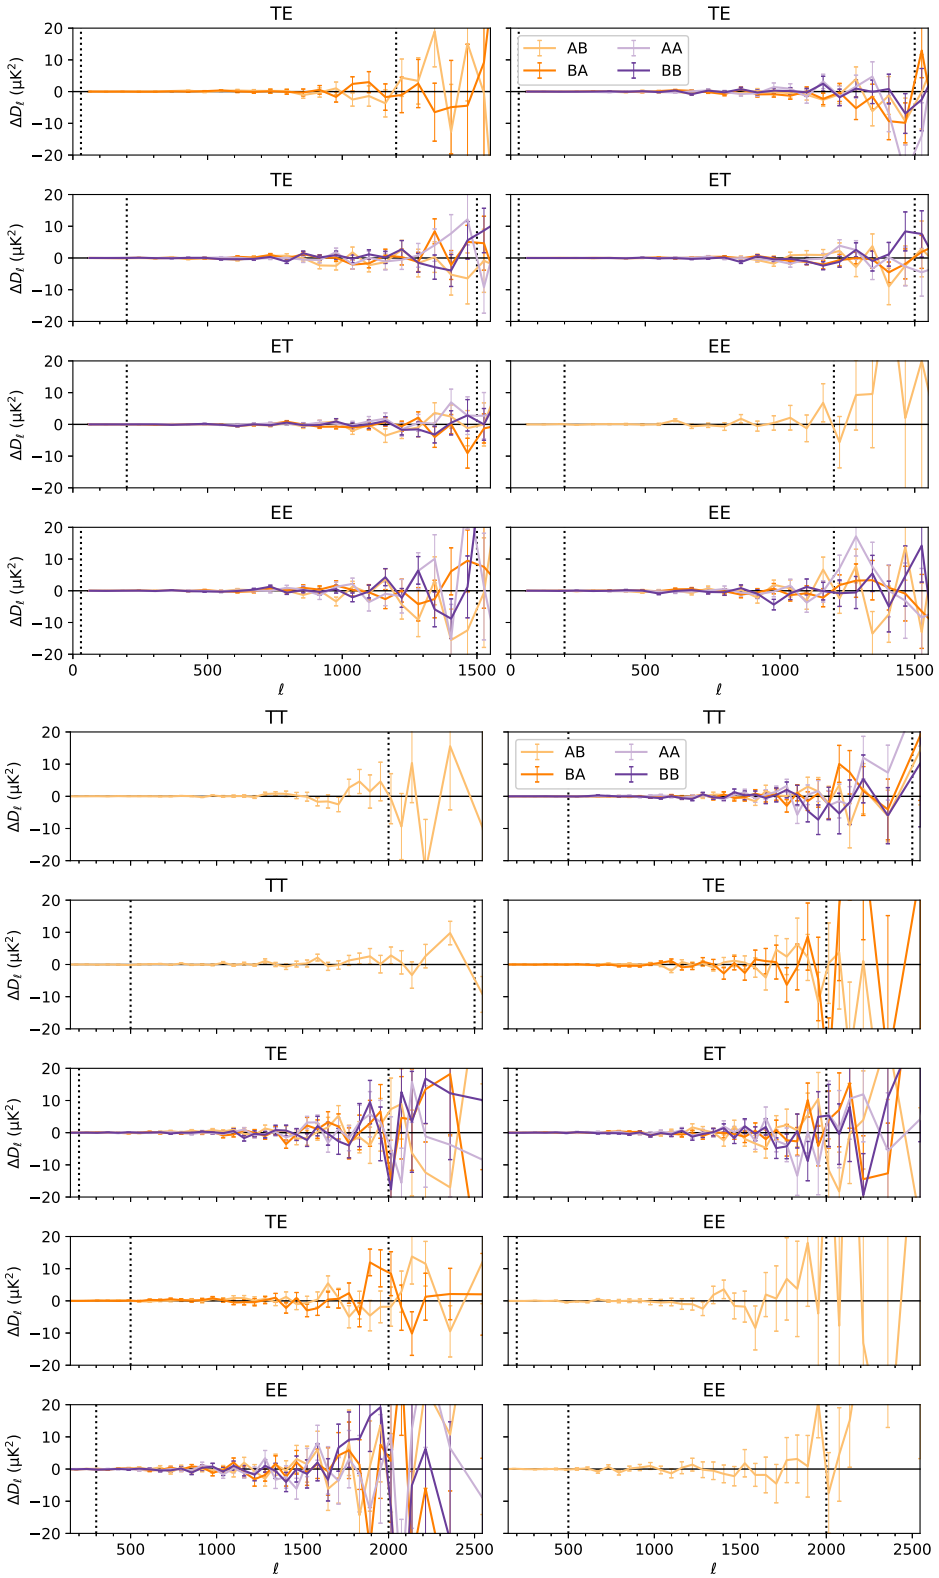

Figure A1. NPIPE half-ring difference cross-spectra. The error bars show the standard deviation $\sigma/\sqrt{\Delta\ell}$ within each bandpower of bin-width $\Delta\ell$.

This paper has been typeset from a TEX/LATEX file prepared by the author.