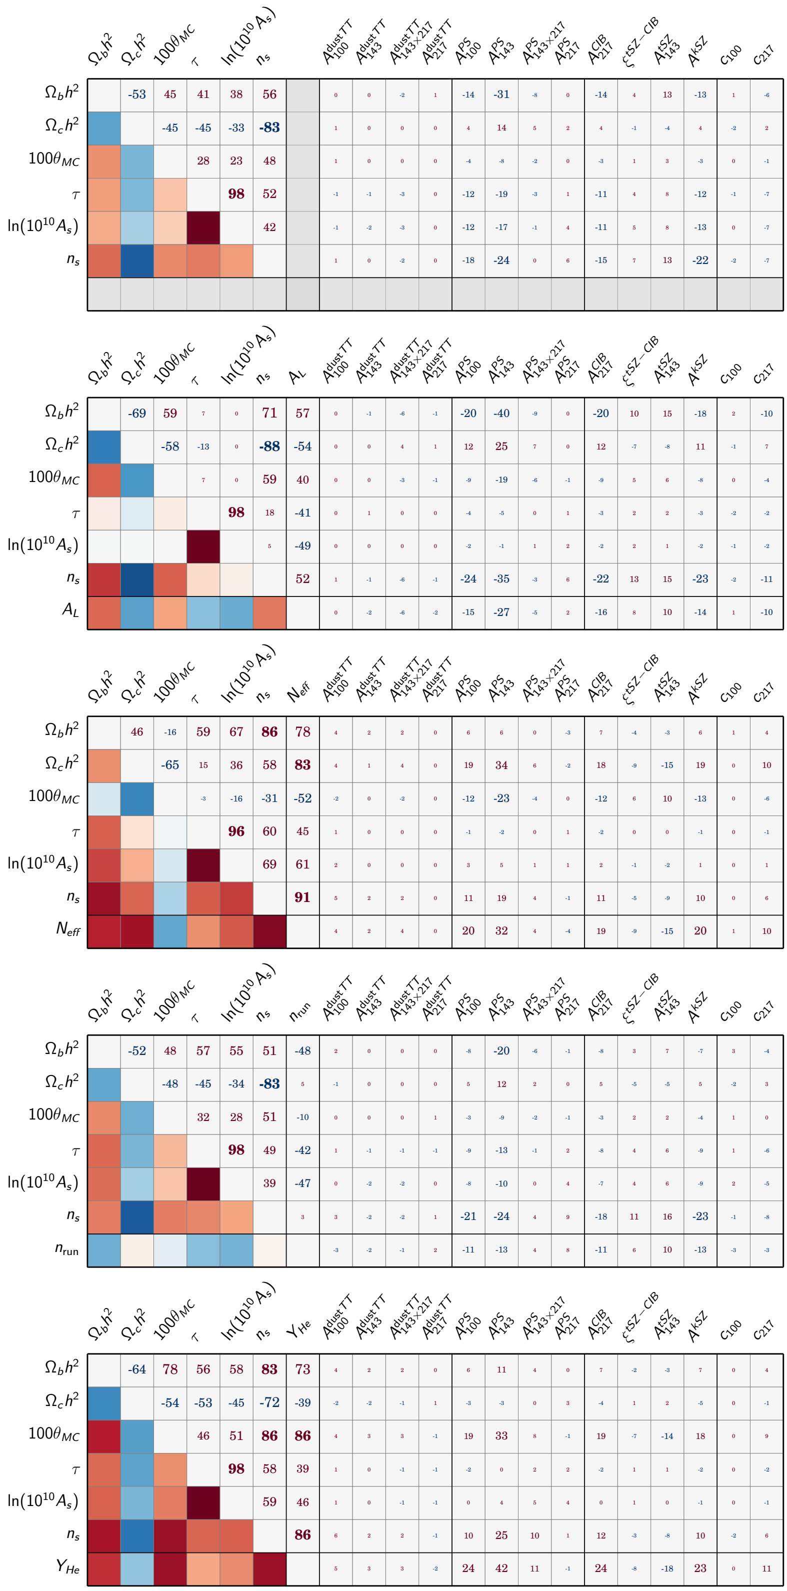

Fig. 45. Parameter correlations for PlikTT+lowTEB, including some ΛCDM extensions. The leftmost column is identical to Fig. 44 and is repeated here to ease comparison. Including extensions to the ΛCDM model changes the correlations between the cosmological parameters, sometimes dramatically, as can be seen in the case of $A_{\rm L}$. There is no correlation between the cosmological parameters (including the extensions) and the dust amplitude parameters. In most cases, the extensions are correlated with the remaining foreground parameters (and in particular with the point-source amplitudes at 100 and 143 GHz, and with the level of CIB fluctuations) with a strength similar to those of the other cosmological parameters (i.e., less than 30%). $Y_{\rm He}$ exhibits a stronger sensitivity to the point-source levels.

their combined noise and do not exhibit a reionization signal, as shown in Fig. 11.

We learn from these tests that if the EE and TE signal we measure at 70 GHz is due to systematics, then these systematics should affect only the above spectra in such a way to mimic a genuine reionization signal, and one that is fully compatible in the maps with that present in (cleaned) WMAP data. This is extremely unlikely and conclude that *Planck* 70 GHz is dominated by a genuine contribution from the sky, compatible with a signal from cosmic reionization.

The tests described so far do not let us accurately quantify the magnitude of a possible systematic contribution, nor to exclude artefacts arising from the data pipeline itself and, specifically, from the foreground cleaning procedure. These can be only controlled through detailed end-to-end tests, using the FFP8 simulations (Planck Collaboration XII 2016). As detailed in Sect. 2.5, we have performed end-to-end validation with 1000 simulated frequency maps containing signal, noise, and foreground contributions as well as specific systematics effects, mimicking all the steps in the actual data pipeline. Propagating to cosmological parameters (τ and A_s, which are most relevant at low ℓ in the ΛCDM model), we detect no bias within the simulation error budget. The total impact of any unknown systematics on the final τ estimate is at most 0.1σ. This effectively rules out any detectable systematic contribution from the data pipeline or or from the instrumental effects considered in the FFP8 simulations. A complementary analysis has been performed in Planck Collaboration III (2016), including further systematic contributions not incorporated into FFP8. This study, which should be taken as a worst-case scenario, limits the possible contribution to final τ of all known systematics at 0.005, i.e., about 0.25σ. We conclude that we were unable to detect any systematic contribution to the 2015 *Planck* τ measurement as driven by low ℓ, and have limited it to well within our final statistical error budget.

Finally, since the submission of this paper, dedicated work on HFI data at low ℓ leads to a higher-precision determination of τ (Planck Collaboration Int. XLVI 2016) which is consistent with the one described in this paper. This latest work paves the way towards a future release of improved *Planck* likelihoods.

5.4.2. High-ℓ budget

We now turn to the high-ℓ likelihood. The approximate statistical model from which we build the likelihood function may turn out to be an unfaithful representation of the data for three main reasons. First the equations describing the likelihood or the parameters of those equations can be inaccurate. They are, of course, since we are relying on approximations, but we expect that in the regime where they are used our approximations are good enough not to bias the best fit or strongly alter the estimation of error bars. We call such errors due to a breakdown of the approximations a “methodological systematic”. We may also lump into this any coding errors. Second, our data model must include a faithful description of the relation between the sky and the data analysed, i.e., one needs to describe the transfer function and/or additive biases due to the non-ideal instrument and data processing. Again, an error in this model or in its parameters translates into possible errors that we call “instrumental systematics”. Finally, to recover the properties of the CMB, the contribution of astrophysical foregrounds must be correctly modelled and accounted for. Errors in this model or its parameters is denoted “astrophysical systematics”. When propagating each of these systematics to cosmological parameters, this is always within the framework of the ΛCDM model, as systematic effects can project differently into parameters depending on the details of the model.

We investigated the possibility of methodological systematics with massive Monte-Carlo simulations. One of the main technical difficulties of the high-ℓ likelihood is the computation of the covariance of the band powers. Appendix C.1.3 describes how we validated the covariance matrix, through the use of Monte-Carlo simulations, to better than a percent accuracy. This includes a first-order correction for the excess scatter due to point source masks, which can induce a systematic error in the covariance reaching a maximum of around 10% near the first peak and the largest scales ($\ell < 50$), somewhat lower (about 5% or less) at other scales. In Sect. 3.6 we propagated the effect of those possible methodological systematics to the cosmological parameters and found a 0.1σ systematic shift on n_s, when using the temperature data, which decreases when cutting the largest and most non-Gaussian modes. This is further demonstrated on the data in Sect. 5.1 where we vary the hybridization scale. At this stage it is unclear whether this is a sign of a breakdown of the Gaussian approximation at those scales, or if it is the result of the limitations of our point source correction to the covariance matrix. We did not try to correct for this bias in the likelihood and we assess this 0.1σ effect on n_s to be the main contribution to the methodological systematics error budget.

Instrumental systematics are mainly assessed in three ways. First, given a foreground model, we estimate the consistency between frequencies and between the TT, EE and TE combinations at the spectrum and at the parameter level (removing some cross-spectra). For TT, the agreement is excellent, with shifts between parameters that are always compatible with the extra cosmic variance due to the removal of data when compared to the baseline solution (see Figs. 31 and 42). TE and EE inter-frequency tests reveal discrepancies between the different cross spectra that we assigned to leakage from temperature to polarization (see Fig. 40). In co-added spectra, these discrepancies tend to average out, leaving a few-μK^2-level residual in the difference between the co-added TE and EE spectra and their theoretical predictions based on the TT parameters. Section 3.4.3 describes an effective model that succeeds in capturing some of that mismatch, in particular in TE. But as argued in Sect. 3.4.3 and Appendix C.3.5 one cannot, at this stage, use this model as-is to correct for the leakage, or to infer the level of systematic it may induce on cosmological parameters, due to a lack of a good prior on the leakage model parameters. However, cosmological parameters deduced from the current polarization likelihoods are in perfect agreement with those calculated from the temperature, within the uncertainty allowed by our covariance. The second way we assess possible instrumental systematics is by comparing the detset (DS) and the half-mission (HM) results. As argued in Sect. 3.4.4, the DS cross spectra are known to be affected by a systematic noise correlation that we correct for. Ignoring any uncertainty in this correction (which is difficult to assess), the overall shift between the HM- and DS-based parameters is of the order of 0.2σ (on ω_b) at most on TT (Sect. 4.1.1 and Fig. 35), similar in TE and slightly worse in EE, particularily for n_s. Since the uncertainty on the correlated noise correction is not propagated, those shifts are only upper bounds on possible instrumental systematics (at least those which would manifest differently in these two data cuts which are completely different as regards temporal systematics). Finally, in Sect. 3.7, we evaluate the propagation of all known instrumental effects to parameters. Due to the cost of the required massive end-to-end simulations, this test can only reveal large deviations; no such instrumental systematic bias is detected in this test. To summarize, our instrumental systematics budget is at most 0.2σ in temperature, slightly higher in EE, and there is no sign of bias due to temperature-to-polarization leakage that would not be compatible with our covariance (within the ΛCDM framework).

Finally, we assess the contribution of astrophysical systematics. Given the prior findings on polarisation, we only discuss the case of temperature here. The uncertainty on the faithfulness of the astrophysical model is relatively high, and we know from the DS/HM comparison that our astrophysical components certainly absorb part of the correlated noise that is not entirely captured by our model. In that sense, the recovered astrophysical parameters may be a biased estimate of the real astrophysical foreground contribution (due to the flexibility of the model which may absorb residual instrumental systematics provided they are sufficiently small). At small scales, the dominant astrophysical component is the point source Poisson term. We checked in Sect. 4.3 that the recovered point source contributions are in general agreement with models of their expected level. This is much less the case at 100 GHz and we argued in Sect. 4.3 that, nonetheless, an error in the description of the Poisson term is unlikely to translate into a bias in the cosmological parameters, as the point source contribution is negligible at all scales where the 100 GHz spectrum dominates the CMB solution. At large scales, the dust is our strongest foreground. We checked in Fig. 35 the effect of either marginalizing out the slope of the dust spectrum or removing the amplitude priors (i.e., making them arbitrarily wide). When marginalizing over the slope, one recovers a value compatible with the one in our model (-2.57 ± 0.038 whereas our model uses -2.63) and the cosmological parameters do not change (Sect. 4.1.2). When comparing the baseline likelihood result to `CamSpec` which uses a slightly different template we find a 0.16σ systematic shift in n_s that can be attributed to the steeper dust template slope (-2.7) (Sect. 4.2). When ignoring the amplitude priors, a 0.2σ shift appears on n_s (and A_s, due to its correlation with n_s). However, in this case the level of dust contribution increases by about $20\,\mu K^2$ in all spectra, which corresponds to more than doubling the 100×100 dust contribution. This level is completely ruled out by the 100×545 cross spectrum, which enables estimation of the dust contribution in the 100 GHz channel. The parameter shift can hence be attributed to a degeneracy between the dust model and the cosmological model broken by the prior on the amplitude parameters. We also use the fact that the dust distribution is anisotropic on the sky and evaluate the cosmological parameters on a smaller sky fraction. On TT there is no shift in the parameters that cannot be attributed to the greater cosmic variance on the smaller sky fraction. We are also making a simplifying assumption by describing the dust as a Gaussian field with a specific power spectrum. The numerical simulations (FFP9 and End-to-end) that include a realistic, anisotropic template for the dust contribution do not uncover any systematic effect due to that approximation. In the end, we believe that 0.2σ on n_s is a conservative upper bound of our astrophysical systematic bias on the cosmological parameters. There is, however, a possibility of a residual instrumental bias affecting foreground parameters (but not cosmology), but we cannot, at this stage, provide quantitative estimates.

To summarize, our systematic error budget consists of a 0.1σ methodology bias on n_s for TT, at most a 0.2σ instrumental bias on TT (on ω_b), TE and possibly a slightly greater one on EE. The few-μK^2-level leakage residual in polarization does not appear to project onto biases on the ΛCDM parameters. We conservatively evalute our astrophysical bias to be 0.2σ on n_s. The astrophysical parameters might suffer from instrumental biases.

5.5. The low-ℓ "anomaly"

In Like13 we noted that the *Planck* 2013 low-ℓ temperature power spectrum exhibited a tension with the *Planck* best-fit model, which is mostly determined by high-ℓ information. In order to quantify such a tension, we performed a series of tests, concluding that the low-ℓ power anomaly was mainly driven by multipoles between $\ell = 20$ and 30, which happen to be systematically low with respect to the model. The effect was shown to be also present (although less pronounced) using WMAP data (again, see Like13 and Page et al. 2007). The statistical significance of this anomaly was found to be around 99%, with slight variations depending on the *Planck* CMB solution or the estimator considered. This anomaly has drawn significant attention as a potential tracer of new physics (e.g., Kitazawa & Sagnotti 2015, 2014; Dudas et al. 2012; see also Destri et al. 2008), so it is worth checking its status in the 2015 analysis.

We present here updated results from a selection of the tests performed in 2013. While in Like13 we only concentrated on temperature, we now also consider low-ℓ polarization, which was not available as a *Planck* product in 2013. We first perform an analysis through the Hausman test (Polenta et al. 2005), modified as in Like13 for the statistic $s_1 = \sup_r B(\ell_{\max}, r)$, with $\ell_{\max} = 29$ and

$$B(\ell_{\max}, r) = \frac{1}{\sqrt{\ell_{\max}}} \sum_{\ell=2}^{\mathrm{int}(\ell_{\max} r)} H_\ell, \quad r \in [0, 1]\,, \tag{61}$$

$$H_\ell = \frac{\hat{C}_\ell - C_\ell}{\sqrt{\mathrm{Var}\,\hat{C}_\ell}}, \tag{62}$$

where $\hat{C}_\ell$ and C_ℓ denote the observed and model power spectra, respectively. Intuitively, this statistic measures the relative bias between the observed spectrum and model, expressed in units of standard deviations, while taking the so-called "look-elsewhere effect" into account by maximizing s_1 over multipole ranges. We use the same simulations as described in Sect. 2.3, which are based on FFP8, for the likelihood validation. We plot in Fig. 46 the empirical distribution for s_1 in temperature and compare it to the value inferred from the *Planck* `Commander` 2015 map described in Sect. 2 above. The significance for the `Commander` map has weakened from 0.7% in 2013 to 2.8% in 2015. This appears consistent with the changes between the 2013 and 2015 `Commander` power spectra shown in Fig. 2, where we can see that the estimates in the range $20 < \ell < 30$ were generally shifted upwards (and closer to the *Planck* best-fit model) due to revised calibration and improved analysis on a larger portion of the sky. We also report in the lower panel of Fig. 46 the same test for the EE power spectrum, finding that the observed *Planck* low-ℓ polarization maps are anomalous only at the 7.7% level.

As a further test of the low-ℓ and high-ℓ *Planck* constraints, we compare the estimate of the primordial amplitude A_s and the optical depth τ, first separately for low and high multipoles, and then jointly. Results are displayed in Fig. 47, showing that the $\ell < 30$ and the $\ell \geq 30$ data posteriors in the primordial amplitude are separated by 2.6σ, where the standard deviation is computed as the square root of the sum of the variances of each posterior. We note that a similar separation exists for τ, but it is only significant at the 1.5σ level. Fixing the value of the high-ℓ parameters to the *Planck* 2013 best-fit model slightly increases the significance of the power anomaly, but has virtually no effect on τ. A joint analysis using all multipoles retrieves best-fit values in A_s and τ which are between the low and high-ℓ posteriors. This behaviour is confirmed when the *Planck* 2015 lensing

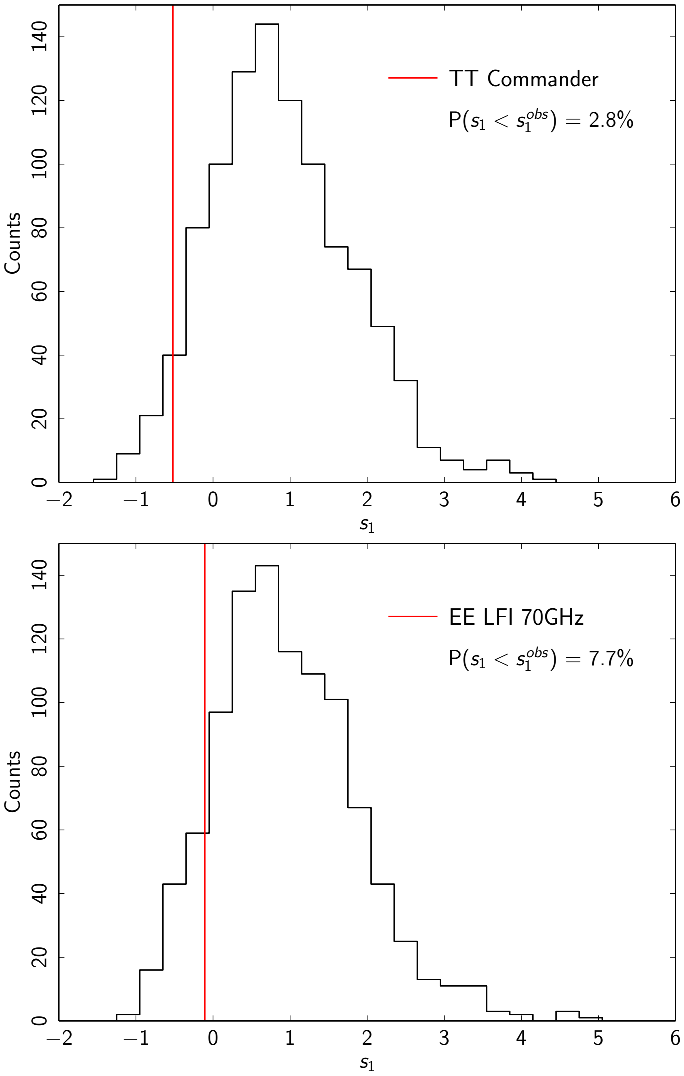

Fig. 46. *Top*: empirical distribution for the Hausman s_1 statistic for TT derived from simulations; the vertical bar is the observed value for the *Planck* `Commander` map. *Bottom*: the empirical distribution of s_1 for EE and the *Planck* 70 GHz polarization maps described in Sect. 2.

likelihood (Planck Collaboration XV 2016) is used in place of low-ℓ polarization.

Finally, we note a similar effect on $N_{\rm eff}$, which, in the high-ℓ analysis with a τ prior is about 1σ off the canonical value of 3.04, but is right on top of the canonical value once the lowP and its $\ell = 20$ dip is included.

5.6. Compressed CMB-only high-ℓ likelihood

We extend the Gibbs sampling scheme described in Dunkley et al. (2013) and Calabrese et al. (2013) to construct a compressed temperature and polarization *Planck* high-ℓ CMB-only likelihood, `Plik_lite`, estimating CMB band-powers and the associated covariances after marginalizing over foreground contributions. Instead of using the full multi-frequency likelihood to directly estimate cosmological parameters and nuisance parameters describing other foregrounds, we take the intermediate step of using the full likelihood to extract CMB temperature and polarization power spectra, marginalizing over possible Galactic and extragalactic contamination. In the process, a new covariance matrix is generated for the marginalized spectra, which therefore includes foreground uncertainty. We refer to Appendix C.6.2 for a description of the methodology and to Fig. C.12 for a comparison between the multi-frequency data and the extracted CMB-only band-powers for TT, TE, and EE.

By marginalizing over nuisance parameters in the spectrum-estimation step, we decouple the primary CMB from non-CMB information. We use the extracted marginalized spectra and covariance matrix in a compressed, high-ℓ, CMB-only likelihood. No additional nuisance parameters, except the overall *Planck* calibration $y_{\rm P}$, are then needed when estimating cosmology, so the convergence of the MCMC chains is significantly faster. To test the performance of this compressed likelihood, we compare results using both the full multi-frequency likelihood and the CMB-only version, for the ΛCDM six-parameter model and for a set of six ΛCDM extensions.

We show in Appendix C.6.2 that the agreement between the results of the full likelihood and its compressed version is excellent, with consistency to better than 0.1σ for all parameters. We have therefore included this compressed likelihood, `Plik_lite`, in the *Planck* likelihood package that is available in the *Planck* Legacy Archive[17].

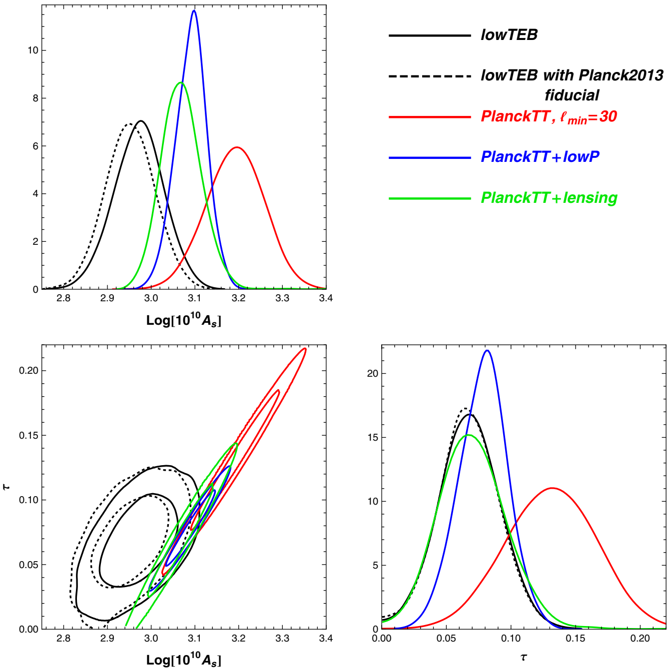

Fig. 47. Joint estimates of primordial amplitude $A_{\rm s}$ and τ for the data sets indicated in the legend. For low-ℓ estimates, all other parameters are fixed to the 2015 fiducial values, except for the dashed line, which uses the *Planck* 2013 fiducial. The PlanckTT+lowP estimates fall roughly half way between the low- and high-ℓ only ones.

5.7. Planck and other CMB experiments

5.7.1. WMAP-9

In Sect. 2.6 we presented the WMAP-9-based low-ℓ polarization likelihood, which uses the *Planck* 353 GHz map as a dust tracer, as well as the *Planck* and WMAP-9 combination. Results for these likelihoods are presented in Table 22, in conjunction with the *Planck* high-ℓ likelihood. Parameter results for the joint *Planck* and WMAP data set in the union mask

[17] http://pla.esac.esa.int/pla/

Table 22. Selected parameters estimated from *Planck*, WMAP, and their noise-weighted combination in low-ℓ polarization, assuming *Planck* in temperature at all multipoles.

Parameter	*Planck*	*WMAP*	*Planck+WMAP*
τ	$0.077^{+0.019}_{-0.018}$	$0.071^{+0.012}_{-0.012}$	$0.074^{+0.012}_{-0.012}$
$z_{\rm re}$	$9.8^{+1.8}_{-1.6}$	$9.3^{+1.1}_{-1.1}$	$9.63^{+1.1}_{-1.0}$
$\log[10^{10}A_{\rm s}]$	$3.087^{+0.036}_{-0.035}$	$3.076^{+0.022}_{-0.022}$	$3.082^{+0.021}_{-0.023}$
r	[0, 0.11]	[0, 0.096]	[0, 0.10]
$A_{\rm s}e^{-2\tau}$	$1.878^{+0.010}_{-0.010}$	$1.879^{+0.011}_{-0.010}$	$1.879^{+0.010}_{-0.010}$

Notes. The *Planck* `Commander` temperature map is always used at low ℓ, while the `Plik` *TT* likelihood is used at high ℓ. All the base-ΛCDM parameters and r are sampled.

are further discussed in Planck Collaboration XIII (2016) and Planck Collaboration XX (2016).

We now illustrate the state of agreement reached between the *Planck* 2015 data, in both the raw and likelihood processed form, and the final cosmological power spectra results from WMAP-9. In 2013 we noted that the difference between WMAP-9 and *Planck* data was mostly related to calibration, which is now resolved with the upward calibration shift in the *Planck* 2015 maps and spectra, as discussed in Planck Collaboration I (2016). This leads to the rather impressive agreement that has been reached between the two *Planck* instruments and WMAP-9.

Figure 48 (top panel) shows all the spectra after correction for the effects of sky masking, with different masks used in the three cases of the *Planck* frequency-map spectra, the spectrum computed from the *Planck* likelihood, and the WMAP-9 final spectrum. The *Planck* 70, 100, and 143 GHz spectra (which are shown as green, red, and blue points, respectively) were derived from the raw frequency maps (cross-spectra of the half-ring data splits for the signal, and spectra of the difference thereof for the noise estimates) on approximately 60% of the sky (with no apodization), where the sky cuts include the Galaxy mask, and a concatenation of the 70, 100, and 143 GHz point-source masks.

The spectrum computed from the *Planck* likelihood (shown in black as both individual and binned C_ℓ values in Fig. 48) was described earlier in the paper. We recall that it was derived with no use of the 70 GHz data, but including the 217 GHz data. Importantly, since it illustrates the likelihood output, this spectrum has been corrected (in the spectral domain) for the residual effects of diffuse foreground emission, mostly in the low-ℓ range, and for the collective effects of several components of discrete foreground emission (including tSZ, point sources, CIB, etc.). This spectrum effectively carries the information that drives the likelihood solution of the *Planck* 2015 best-fit CMB anisotropy, shown in brown. Our aim here is to show the conformity between this *Planck* 2015 solution and the raw *Planck* data (especially at 70 GHz) and the WMAP-9 legacy spectrum.

The WMAP-9 spectrum (shown in magenta as both individual and binned C_ℓ values) is the legacy product from the WMAP-9 mission, and it represents the final results of the WMAP team's efforts to clean the residual effects of foreground emission from the cosmological anisotropy spectrum.

All these spectra are binned the same way, starting at $\ell = 30$ with $\Delta\ell = 40$ bins, and the error-bars represent the error on the mean within each bin. In the low-ℓ range, especially near the first peak, the error calculation includes the cosmic variance contribution from the multipoles within each bin, which vastly exceeds any measurement errors (all the measurements shown here have high S/N over the first spectral peak), so we would expect good agreement between the errors derived for all the spectra in the completely signal-dominated range of the data.

The figure shows how WMAP-9 loses accuracy above $\ell \approx 800$ due to its inherent beam resolution and instrumental noise, and shows how the LFI 70 GHz data achieve improved fidelity over this range. HFI was designed to improve over both WMAP-9 and LFI in both noise performance and angular resolution, and the gains achieved are clearly visible, even over the relatively modest range of ℓ shown here, in the tiny spread of the individual C_ℓ values of the *Planck* 2015 power spectrum. While the overall agreement of the various spectra, especially in the low-ℓ range, is noticeable in this coarse plot, it is also clear that the *Planck* raw frequency-map spectra do show excess power over the *Planck* best-fit spectrum at the higher end of the ℓ-range shown – the highest level at 70 GHz and the lowest at 143 GHz. This illustrates the effect of uncorrected discrete foreground residuals in the raw spectra.

A better view of these effects is seen in the bottom panel of Fig. 48. Here we plot the binned values from the top panel as deviations from the best-fit model. Naturally, the black bins of the likelihood output fit well, since they were derived jointly with the best-fit spectrum, while correcting for foreground residuals. The WMAP-9 points show good agreement, given their errors, with the *Planck* 2015 best fit, and illustrate very tight control of the large-scale residual foregrounds (at the low-ℓ range of the figure); beyond $\ell \sim 600$ the WMAP-9 spectrum shows an increasing loss of fidelity. *Planck* raw 70, 100, and 143 GHz spectra show excess power in the lowest ℓ bin due to diffuse foreground residuals. The higher-ℓ range now shows more clearly the upward drift of power in the raw spectra, growing from 143 GHz to 70 GHz. This is consistent with the well-determined integrated discrete foreground contributions to those spectra. As previously shown in Planck Collaboration XXXI (2014, Fig. 8), the unresolved discrete foreground power (computed with the same sky masks as used here) can be represented in the bin near $\ell = 800$ as levels of approximately 40 μK^2 at 70 GHz, 15 μK^2 at 100 GHz, and 5 μK^2 at 143 GHz, in good agreement with the present figure.

5.7.2. ACT and SPT

Planck temperature observations are complemented at finer scales by measurements from the ground-based Atacama Cosmology Telescope (ACT) and South Pole Telescope (SPT). The ACT and SPT high-resolution data help *Planck* in separating the primordial cosmological signal from other Galactic and extragalactic emission, so as not to bias cosmological reconstructions in the damping-tail region of the spectrum. In 2013 we combined *Planck* with ACT (Das et al. 2014) and SPT (Reichardt et al. 2012) data in the multipole range $1000 < \ell < 10\,000$, defining a common foreground model and extracting cosmological parameters from all the data sets. Our updated "highL" temperature data include ACT power spectra at 148 and 218 GHz (Das et al. 2014) with a revised binning (Calabrese et al. 2013) and final beam estimates (Hasselfield et al. 2013), and SPT measurements in the range $2000 < \ell < 13\,000$ from the 2540 deg^2 SPT-SZ survey at 95, 150, and 220 GHz (George et al. 2015). However, in this new analysis, given the increased constraining power of the *Planck* full-mission data, we do not use ACT and SPT as primary data sets. Using the same ℓ cuts as the 2013 analysis (i.e., ACT data at $1000 < \ell < 10\,000$ and SPT at $\ell > 2000$) we only check for

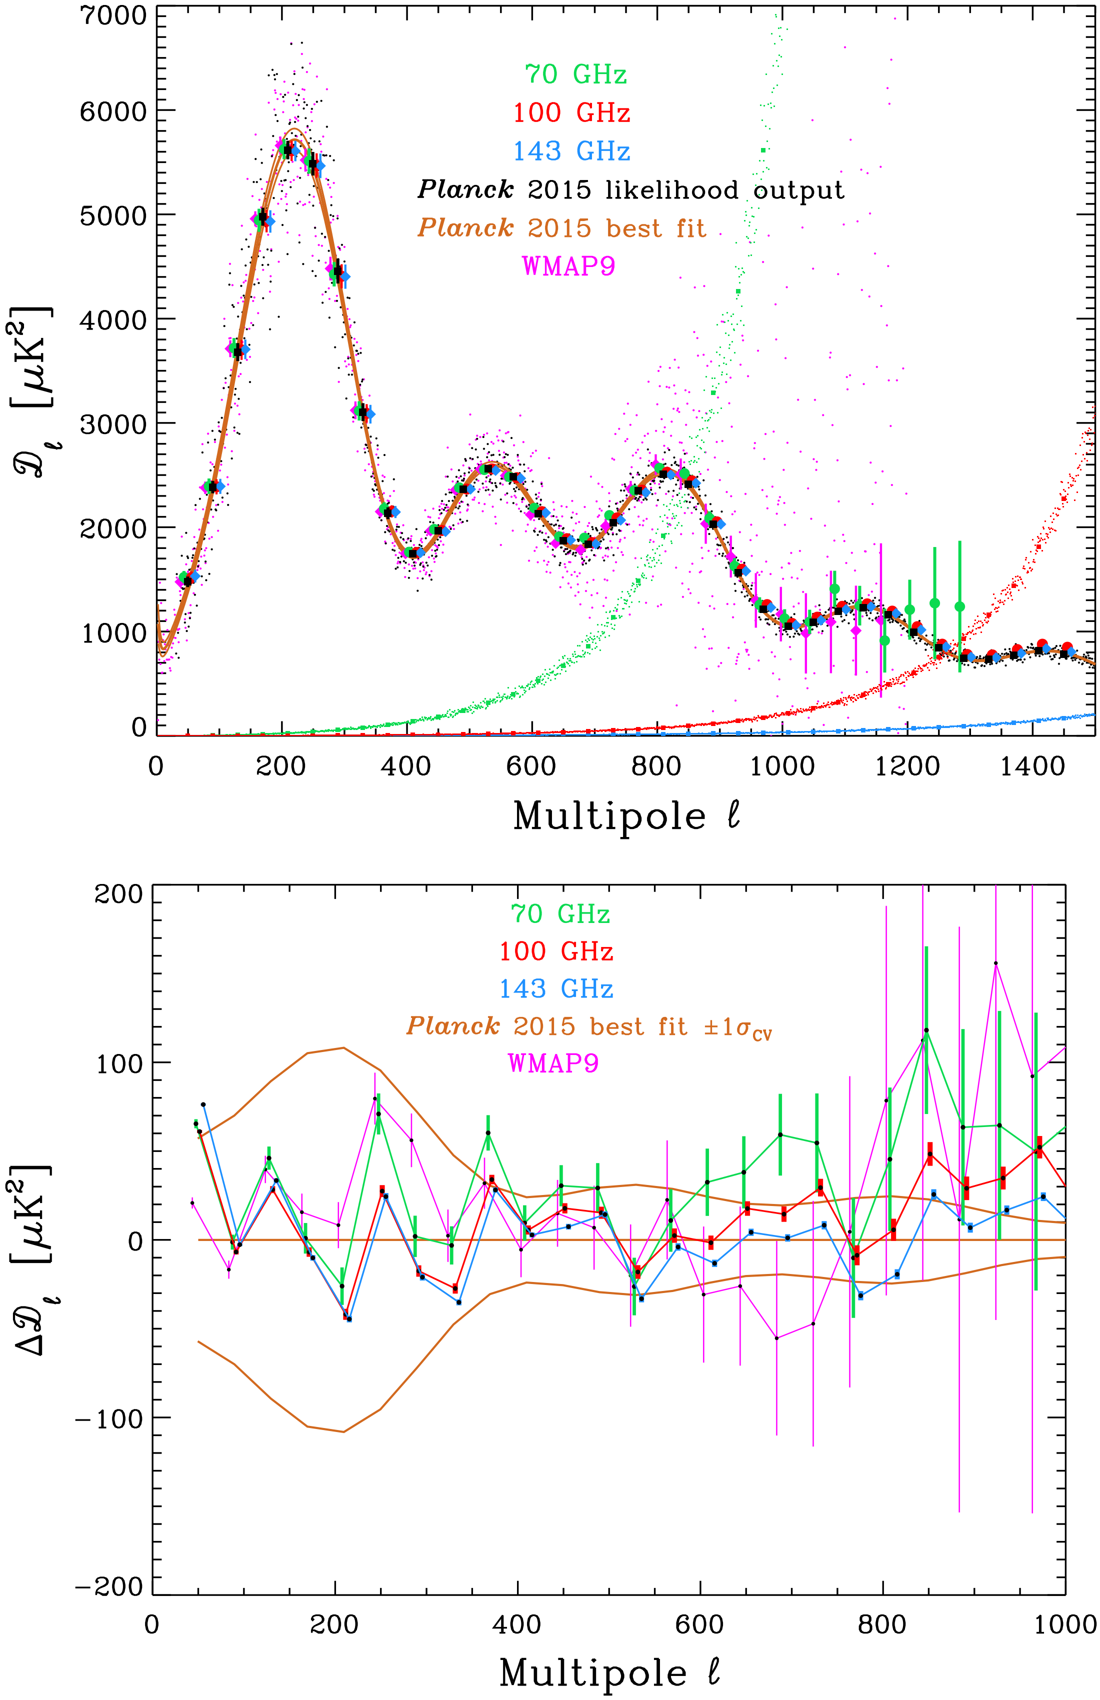

Fig. 48. Comparison of *Planck* and WMAP-9 CMB power spectra. *Top*: direct comparison. Noise spectra are derived from the half-ring difference maps. *Bottom*: residuals with respect to the *Planck* ΛCDM best-fit model. The error bars do not include the cosmic variance contribution (but the (brown) 1σ contour lines for the Likelihood best fit model do).

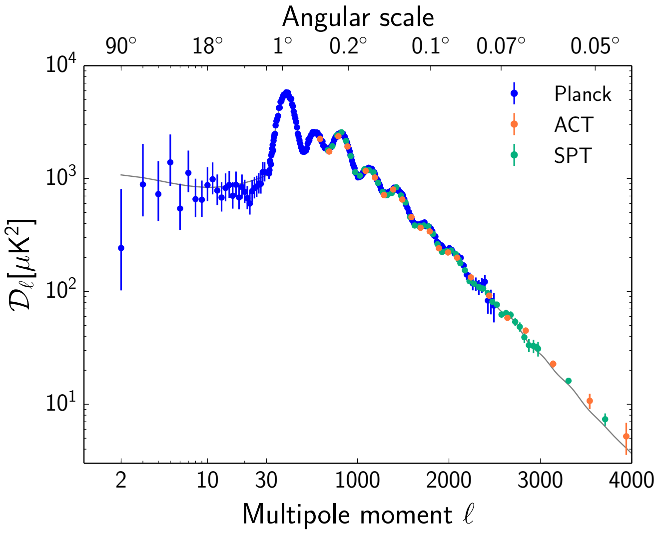

Fig. 49. CMB-only power spectra measured by *Planck* (blue), ACT (orange), and SPT (green). The best-fit PlanckTT+lowP ΛCDM model is shown by the grey solid line. ACT data at $\ell > 1000$ and SPT data at $\ell > 2000$ are marginalized CMB band-powers from multi-frequency spectra presented in Das et al. (2014) and George et al. (2015) as extracted in this work. Lower multipole ACT ($500 < \ell < 1000$) and SPT ($650 < \ell < 3000$) CMB power extracted by Calabrese et al. (2013) from multi-frequency spectra presented in Das et al. (2014) and Story et al. (2013) are also shown. The binned values in the range $3000 < \ell < 4000$ appear higher than the unbinned best-fit line because of the binning (this is numerically confirmed by the residual plot in Planck Collaboration XIII 2016, Fig. 9).

consistency and retain information on the nuisance foreground parameters that are not well constrained by *Planck* alone.

To assess the consistency between these data sets, we extend the *Planck* foreground model up to $\ell = 13\,000$ with additional nuisance parameters for ACT and SPT, as described in Planck Collaboration XIII (2016, Sect. 4). Fixing the cosmological parameters to the best-fit PlanckTT+lowP base-ΛCDM model and varying the ACT and SPT foreground and calibration parameters, we find a reduced $\chi^2 = 1.004$ (PTE = 0.46), showing very good agreement between *Planck* and the highL data.

As described in Planck Collaboration XIII (2016), we then take a further step and extend the Gibbs technique presented in Dunkley et al. (2013) and Calabrese et al. (2013; and applied to *Planck* alone in Sect. 5.6) to extract independent CMB-only band-powers from *Planck*, ACT, and SPT. The extracted CMB spectra are reported in Fig. 49. We also show ACT and SPT band-powers at lower multipoles as extracted by Calabrese et al. (2013). This figure shows the state of the art of current CMB observations, with *Planck* covering the low-to-high-multipole range and ACT and SPT extending into the damping region. We consider the CMB to be negligible at $\ell > 4000$ and note that these ACT and SPT band-powers have an overall calibration uncertainty (2% for ACT and 1.2% for SPT).

The inclusion of ACT and SPT improves the full-mission *Planck* spectrum extraction presented in Sect. 5.6 only marginally. The main contribution of ACT and SPT is to constrain small components (e.g., the tSZ, kSZ, and tSZ×CIB) that are not well determined by *Planck* alone. However, those components are sub-dominant for *Planck* and are well described by the prior based on the 2013 *Planck*+highL solutions imposed in the *Planck*-alone analysis. The CIB amplitude estimate improves by 40% when including ACT and SPT, but the CIB power is also reasonably well constrained by *Planck* alone. The main *Planck* contaminants are the Poisson sources, which are treated as independent and do not benefit from ACT and SPT. As a result, the errors on the extracted *Planck* spectrum are only slightly reduced, with little additional cosmological information added by including ACT and SPT for the baseline ΛCDM model (see also Planck Collaboration XIII 2016, Sect. 4).

6. Conclusions

The *Planck* 2015 angular power spectra of the cosmic microwave background derived in this paper are displayed in Fig. 50. These spectra in *TT* (top), *TE* (middle), and *EE* (bottom) are all quite consistent with the best-fit base-ΛCDM model obtained from *TT* data alone (red lines). The horizontal axis is logarithmic at $\ell < 30$, where the spectra are shown for individual multipoles, and linear at $\ell \geq 30$, where the data are binned. The error bars correspond to the diagonal elements of the covariance matrix. The lower panels display the residuals, the data being presented with different vertical axes, a larger one at left for the low-ℓ part and a zoomed-in axis at right for the high-ℓ part.

The 2015 *Planck* likelihood presented in this work is based on more temperature data than in the 2013 release, and on new polarization data. It benefits from several improvements in the processing of the raw data, and in the modelling of astrophysical foregrounds and instrumental noise. Apart from a revision of the overall calibration of the maps, discussed in Planck Collaboration I (2016), the most significant improvements are in the likelihood procedures:

(i) a joint temperature-polarization pixel-based likelihood at $\ell \leq 29$, with more high-frequency information used for foreground removal, and smaller sky masks (Sects. 2.1 and 2.2);
(ii) an improved Gaussian likelihood at $\ell \geq 30$ that includes a different strategy for estimating power spectra from data-subset cross-correlations, using half-mission data instead of detector sets (which enables us to reduce the effect of correlated noise between detectors, see Sects. 3.2.1 and 3.4.3), and better foreground templates, especially for Galactic dust (Sect. 3.3.1) that lets us mask a smaller fraction of the sky (Sect. 3.2.2) and to retain large-angle temperature information from the 217 GHz map that was neglected in the 2013 release (Sect. 3.2.4).

We performed several consistency checks of the robustness of our likelihood-making process, by introducing more or less freedom and nuisance parameters in the modelling of foregrounds and instrumental noise, and by including different assumptions about the relative calibration uncertainties across frequency channels and about the beam window functions.

For temperature, the reconstructed CMB spectrum and error bars are remarkably insensitive to all these different assumptions. Our final high-ℓ temperature likelihood, referred to as "PlanckTT" marginalizes over 15 nuisance parameters (12 modelling the foregrounds, and 3 for calibration uncertainties). Additional nuisance parameters (in particular, those associated with beam uncertainties) were found to have a negligible impact, and can be kept fixed in the baseline likelihood. Detailed end-to-end simulations of the instrumental response to the sky analysed like the real data did not uncover hidden low-level residual systematics.

For polarization, the situation is different. Variation of the assumptions leads to scattered results, with greater deviations than would be expected due to changes in the data subsets used,

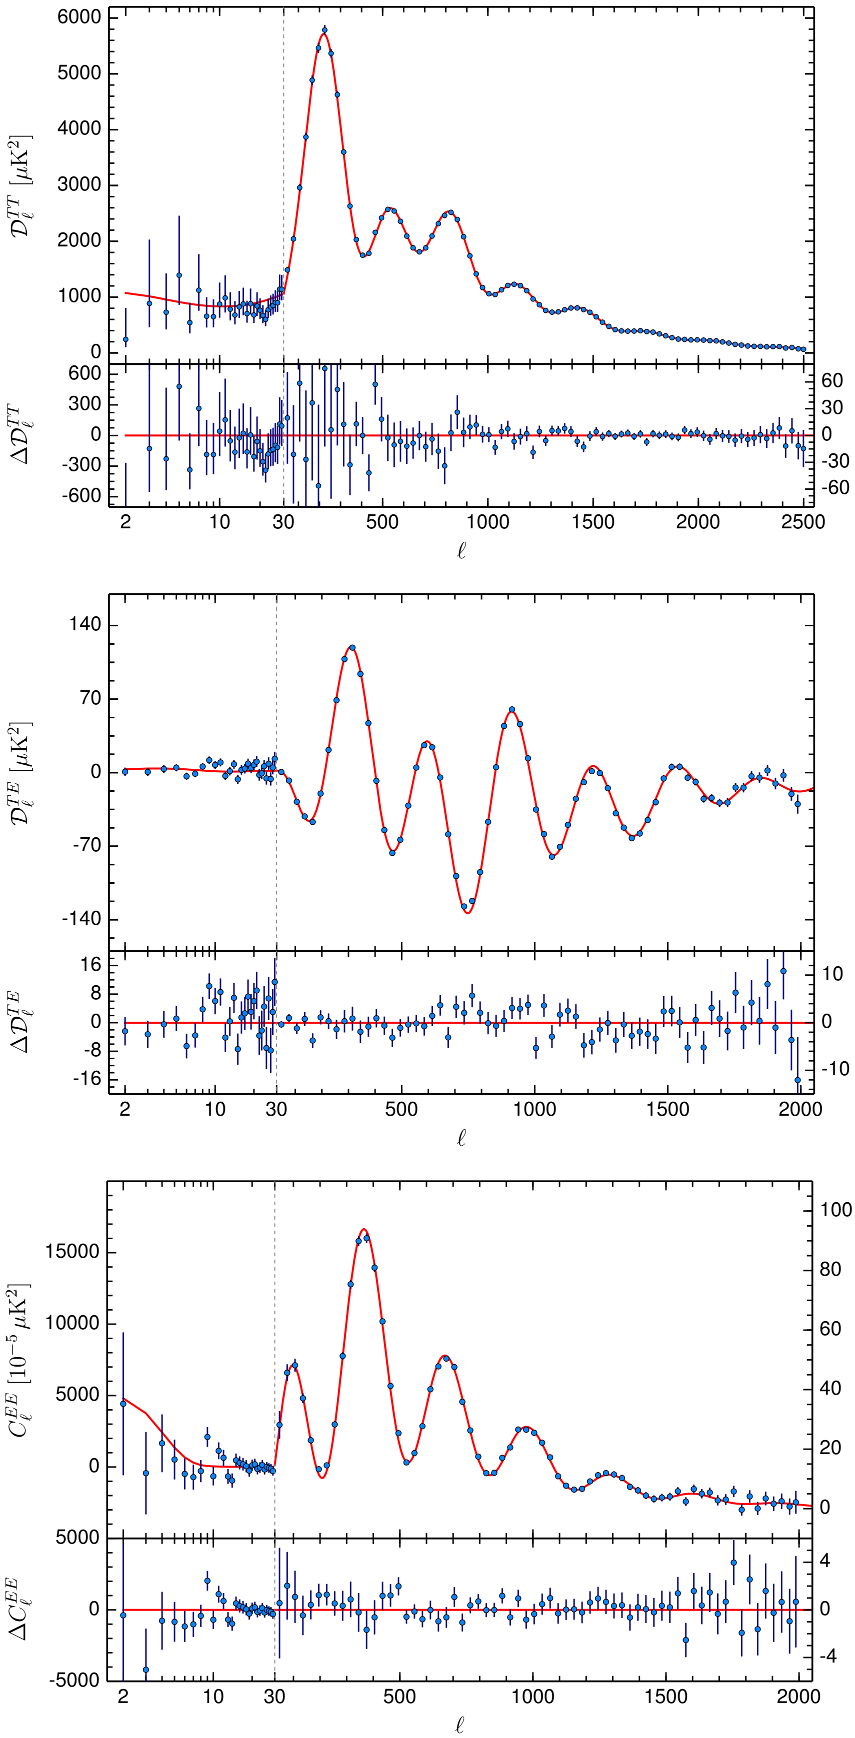

Fig. 50. *Planck* 2015 CMB spectra, compared with the base ΛCDM fit to PlanckTT+lowP data (red line). *The upper panels* show the spectra and the *lower panels* the residuals. In *all the panels*, the horizontal scale changes from logarithmic to linear at the "hybridization" scale, $\ell = 29$ (the division between the low-ℓ and high-ℓ likelihoods). For the residuals, the vertical axis scale changes as well, as shown by different left and right axes. We show $\mathcal{D}_\ell = \ell(\ell+1)C_\ell/(2\pi)$ for TT and TE, but C_ℓ for EE, which also has different vertical scales at low- and high-ℓ.

and at a level that is significant compared to the statistical error bars. This suggests that further systematic effects need to be either modelled or removed. In particular, our attempt to model calibration errors and temperature-to-polarization leakage suggests that the *TE* and *EE* power spectra are affected by systematics at a level of roughly 1 μK^2. Removal of polarization systematics at this level of precision requires further work, beyond the scope of this release. The 2015 high-ℓ polarized likelihoods, referred to as "`PlikTE`" and "`PlikEE`", or "`PlikTT`, EE, TE" for the combined version, ignore these uncertain corrections. They only include 12 additional nuisance parameters accounting for polarized foregrounds. Although these likelihoods are distributed in the *Planck* Legacy Archive[18], we stick to the PlanckTT+lowP choice in the baseline analysis of this paper and the companion papers such as Planck Collaboration XIII (2016), Planck Collaboration XIV (2016), and Planck Collaboration XX (2016).

We developed internally several likelihood codes, exploring not only different assumptions about foregrounds and instrumental noise, but also different algorithms for building an approximate Gaussian high-ℓ likelihood (Sect. 4.2). We compared these codes to check the robustness of the results, and decided to release:

(i) A baseline likelihood called `Plik` (available for TT, *TE*, *EE*, or combined observables), in which the data are binned in multipole space, with a bin-width increasing from $\Delta\ell = 5$ at $\ell \approx 30$ to $\Delta\ell = 33$ at $\ell \approx 2500$.

(ii) An unbinned version which, although slower, is preferable when investigating models with sharp features in the power spectra.

(iii) A simplified likelihood called `Plik_lite` in which the foreground templates and calibration errors are marginalized over, producing a marginalized spectrum and covariance matrix. This likelihood does not allow investigation of correlations between cosmological and foreground/instrumental parameters, but speeds up parameter extraction, having no nuisance parameters to marginalize over.

In this paper we have also presented an investigation of the measurement of cosmological parameters in the minimal six-parameter ΛCDM model and a few simple seven-parameter extensions, using both the new baseline *Planck* likelihood and several alternative likelihoods relying on different assumptions. The cosmological analysis of this paper does not replace the investigation of many extended cosmological models presented, e.g., in Planck Collaboration XIII (2016), Planck Collaboration XIV (2016), and Planck Collaboration XX (2016). However, the careful inspection of residuals presented here addresses two questions:

(i) a priori, is there any indication that an alternative model to ΛCDM could provide a significantly better fit?

(ii) if there is such an indication, could it come from caveats in the likelihood-building (imperfect data reduction, foreground templates or noise modelling) instead of new cosmological ingredients?

Since this work is entirely focused on the power-spectrum likelihood, it can only address these questions at the level of 2-point statistics; for a discussion of higher-order statistics, see Planck Collaboration XVI (2016) and Planck Collaboration XVII (2016).

[18] http://pla.esac.esa.int/pla/

The most striking result of this work is the impressive consistency of different cosmological parameter extractions, performed with different versions of the `PlikTT+tauprior` or PlanckTT+lowP likelihoods, with several assumptions concerning: data processing (half-mission versus detector set correlations); sky masks and foreground templates; beam window functions; the use of two frequency channels instead of three; different cuts at low ℓ or high ℓ; a different choice for the multipole value at which we switch from the pixel-based to the Gaussian likelihood; different codes and algorithms; the inclusion of external data sets like WMAP-9, ACT, or SPT; and the use of foreground-cleaned maps (instead of fitting the CMB+foreground map with a sum of different contributions). In all these cases, the best-fit parameter values drift by only a small amount, compatible with what one would expect on a statistical basis when some of the data are removed (with a few exceptions summarized below).

The cosmological results are stable when one uses the simplified `Plik_lite` likelihood. We checked this by comparing PlanckTT+lowP results from `Plik` and `Plik_lite` for ΛCDM, and for six examples of seven-parameter extended models.

Another striking result is that, despite evidence for small unsolved systematic effects in the high-ℓ polarization data, the cosmological parameters returned by the `PlikTT`, `PlikTE`, or `PlikEE` likelihoods (in combination with a τ prior or *Planck* lowP) are consistent with each other, and the residuals of the (frequency combined) *TE* and *EE* spectra after subtracting the temperature ΛCDM best-fit are consistent with zero. As has been emphasized in other *Planck* 2015 papers, this is a tremendous success for cosmology, and an additional proof of the predictive power of the standard cosmological model. It also suggests that the level of temperature-to-polarization leakage (and possibly other systematic effects) revealed by our consistency checks is low enough (on average over all frequencies) not to significantly bias parameter extraction, at least for the minimal cosmological model. We do not know yet whether this conclusion applies also to extended models, especially those in which the combination of temperature and polarization data has stronger constraining power than temperature data alone, e.g., dark matter annihilation (Planck Collaboration XIII 2016) or isocurvature modes (Planck Collaboration XX 2016). One should thus wait for a future *Planck* release before applying the *Planck* temperature-plus-polarization likelihood to such models. However, the fact that we observe a significant reduction in the error bars when including polarization data is very promising, since this reduction is expected to remain after the removal of systematic effects.

Careful inspection of residuals with respect to the best-fit ΛCDM model has revealed a list of anomalies in the *Planck* CMB power spectra, of which the most significant is still the low-ℓ temperature anomaly in the range $20 \leq \ell \leq 30$, already discussed at length in the 2013 release. In this 2015 release, with more data and with better calibration, foreground modelling, and sky masks, its significance has decreased from the 0.7% to the 2.8% level for the TT spectrum (Sect. 5.5). This probability is still small (although not very small), and the feature remains unexplained. We have also investigated the *EE* spectrum, where the anomaly, if any, is significant only at the 7.7% level.

Other "anomalies" revealed by inspection of residuals (and of their dependence on the assumptions underlying the likelihood) are much less significant. There are a few bins in which the power in the TT, *TE*, or *EE* spectrum lies 2–3σ away from the best-fit ΛCDM prediction, but this is not statistically unlikely and we find acceptable probability-to-exceed (PTE) levels. Nevertheless, in Sects. 3.8 and 4.1, we presented a careful investigation of these features, to see whether they could be caused by some imperfect modelling of the data. We noted that a deviation in the TT spectrum at $\ell \approx 1450$ is somewhat suspicious, since it is driven mostly by a single channel (217 GHz), and since it depends on the foreground-removal method. But this deviation is too small to be worrisome (1.8σ with the baseline `Plik` likelihood). As in the 2013 release, the data at intermediate ℓ would be fitted slightly better by a model with more lensing than in the best-fit ΛCDM model (to reduce the peak-to-trough contrast), but more lensing generically requires higher values of A_s and $\Omega_c h^2$ that are disfavoured by the rest of the data, in particular when high-ℓ information is included. This mild tension is illustrated by the preference for a value greater than unity for the unphysical parameter A_L, a conclusion that is stable against variations in the assumptions underlying the likelihoods. However, A_L is compatible with unity at the 1.8σ level when using the baseline PlanckTT likelihood with a conservative τ prior (to avoid the effect of the low-ℓ dip), so what we see here could be the result of statistical fluctuations.

This absence of large residuals in the *Planck* 2015 temperature and polarization spectra further establishes the robustness of the ΛCDM model, even with about twice as much data as in the *Planck* 2013 release. This conclusion is supported by several companion papers, in which many non-minimal cosmological models are investigated but no significant evidence for extra physical ingredients is found. The ability of the temperature results to pass several demanding consistency tests, and the evidence of excellent agreement down to the $\mu\mathrm{K}^2$ level between the temperature and polarization data, represent an important milestone set by the *Planck* satellite. The *Planck* 2015 likelihoods are the best illustration to date of the predictive power of the minimal cosmological model, and, at the same time, the best tool for constraining interesting, physically-motivated deviations from that model.

Acknowledgements. The Planck Collaboration acknowledges the support of: ESA; CNES, and CNRS/INSU-IN2P3-INP (France); ASI, CNR, and INAF (Italy); NASA and DoE (USA); STFC and UKSA (UK); CSIC, MINECO, JA and RES (Spain); Tekes, AoF, and CSC (Finland); DLR and MPG (Germany); CSA (Canada); DTU Space (Denmark); SER/SSO (Switzerland); RCN (Norway); SFI (Ireland); FCT/MCTES (Portugal); ERC and PRACE (EU). A description of the Planck Collaboration and a list of its members, indicating which technical or scientific activities they have been involved in, can be found at `http://www.cosmos.esa.int/web/planck/planck-collaboration`. We further acknowledge the use of the `CLASS` Boltzmann code (Lesgourgues 2011) and the `Monte Python` package (Audren et al. 2013) in earlier stages of this work. The likelihood code and some of the validation work was built on the library `pmclib` from the `CosmoPMC` package (Kilbinger et al. 2011). This research used resources of the IN2P3 Computer Center (`http://cc.in2p3.fr`) as well as of the Planck-HFI DPC infrastructure hosted at the Institut d'Astrophysique de Paris (France) and financially supported by CNES.

References

Aad, G. E. A. 2014, Phys. Rev. D, 90, 052004
Addison, G. E., Dunkley, J., & Spergel, D. N. 2012, MNRAS, 427, 1741
Addison, G. E., Huang, Y., Watts, D. J., et al. 2016, ApJ, 818, 132
Araujo, D., Bischoff, C., et al. QUIET Collaboration 2012, ApJ, 760, 145
Arnaud, M., Pratt, G. W., Piffaretti, R., et al. 2010, A&A, 517, A92
Audren, B., Lesgourgues, J., Benabed, K., & Prunet, S. 2013, J. Cosmol. Astropart. Phys., 2, 1
Battaglia, N., Natarajan, A., Trac, H., Cen, R., & Loeb, A. 2013, ApJ, 776, 83
Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, ApJS, 208, 20
Benoît, A., Ade, P., Amblard, A., et al. 2003, A&A, 399, L25
Béthermin, M., Daddi, E., Magdis, G., et al. 2012, ApJ, 757, L23
Béthermin, M., Wang, L., Doré, O., et al. 2013, A&A, 557, A66
Blas, D., Lesgourgues, J., & Tram, T. 2011, JCAP, 07, 034
Bond, J. R., Contaldi, C., & Pogosyan, D. 2003, Roy. Soc. London Philosoph. Trans. Ser. A, 361, 2435
Calabrese, E., Hlozek, R. A., Battaglia, N., et al. 2013, Phys. Rev. D, 87, 103012
Challinor, A., & Chon, G. 2005, MNRAS, 360, 509
Chon, G., Challinor, A., Prunet, S., Hivon, E., & Szapudi, I. 2004, MNRAS, 350, 914
Chu, M., Eriksen, H. K., Knox, L., et al. 2005, Phys. Rev. D, 71, 103002
Corasaniti, P. S., & Melchiorri, A. 2008, Phys. Rev. D, 77, 103507
Crites, A. T., Henning, J. W., Ade, P. A. R., et al. 2015, ApJ, 805, 36
Das, S., Marriage, T. A., Ade, P. A. R., et al. 2011, ApJ, 729, 62
Das, S., Louis, T., Nolta, M. R., et al. 2014, JCAP, 04, 014
de Bernardis, P., Ade, P. A. R., Bock, J. J., et al. 2002, ApJ, 564, 559
Delabrouille, J., Betoule, M., Melin, J.-B., et al. 2013, A&A, 553, A96
Destri, C., de Vega, H. J., & Sanchez, N. G. 2008, Phys. Rev. D, 78, 023013
Dudas, E., Kitazawa, N., Patil, S. P., & Sagnotti, A. 2012, JCAP, 05, 012
Dunkley, J., Komatsu, E., Nolta, M. R., et al. 2009, ApJS, 180, 306
Dunkley, J., Calabrese, E., Sievers, J., et al. 2013, JCAP, 07, 025
Durrer, R., Novosyadlyj, B., & Apunevych, S. 2003, ApJ, 583, 33
Efstathiou, G. 2004, MNRAS, 349, 603
Efstathiou, G. 2006, MNRAS, 370, 343
Efstathiou, G., & Migliaccio, M. 2012, MNRAS, 423, 2492
Elsner, F., & Wandelt, B. D. 2012, A&A, 542, A60
Eriksen, H. K., O'Dwyer, I. J., Jewell, J. B., et al. 2004, ApJS, 155, 227
Eriksen, H. K., Jewell, J. B., Dickinson, C., et al. 2008, ApJ, 676, 10
Fendt, W. A., & Wandelt, B. D. 2007a, ArXiv e-prints [arXiv:0712.0194]
Fendt, W. A., & Wandelt, B. D. 2007b, ApJ, 654, 2
Finelli, F., De Rosa, A., Gruppuso, A., & Paoletti, D. 2013, MNRAS, 431, 2961
Fowler, J. W., Acquaviva, V., Ade, P. A. R., et al. 2010, ApJ, 722, 1148
Galli, S., Benabed, K., Bouchet, F., et al. 2014, Phys. Rev. D, 90, 063504
George, E. M., Reichardt, C. L., Aird, K. A., et al. 2015, ApJ, 799, 177
Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759
Grainge, K., Carreira, P., Cleary, K., et al. 2003, MNRAS, 341, L23
Gruppuso, A., de Rosa, A., Cabella, P., et al. 2009, MNRAS, 400, 463
Hamimeche, S., & Lewis, A. 2008, Phys. Rev. D, 77, 103013
Hanany, S., Ade, P. A. R., Balbi, A., et al. 2000, ApJ, 545, L5
Hancock, S., & Rocha, G. 1997, in Microwave Background Anisotropies, eds. F. R. Bouchet, R. Gispert, B. Guiderdoni, & J. Trân Thanh Vân, 179–188
Hansen, F. K., Górski, K. M., & Hivon, E. 2002, MNRAS, 336, 1304
Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, A&AS, 47, 1
Hasselfield, M., Moodley, K., Bond, J. R., et al. 2013, ApJS, 209, 17
Hinshaw, G., Spergel, D. N., Verde, L., et al. 2003, ApJS, 148, 135
Hinshaw, G., Nolta, M. R., Bennett, C. L., et al. 2007, ApJS, 170, 288
Hivon, E., Górski, K. M., Netterfield, C. B., et al. 2002, ApJ, 567, 2
Hu, W., & White, M. 1997, Phys. Rev. D, 56, 596
Hu, W., Sugiyama, N., & Silk, J. 1997, Nature, 386, 37
Hu, W., Hedman, M. M., & Zaldarriaga, M. 2003, Phys. Rev. D, 67, 043004
James, F., & Roos, M. 1975, Computer Phys. Commun., 10, 343
Jones, W. C., Ade, P. A. R., Bock, J. J., et al. 2006, ApJ, 647, 823
Jones, W. C., Montroy, T. E., Crill, B. P., et al. 2007, A&A, 470, 771
Kamionkowski, M., Kosowsky, A., & Stebbins, A. 1997, Phys. Rev. D, 55, 7368
Keisler, R., Reichardt, C. L., Aird, K. A., et al. 2011, ApJ, 743, 28
Kilbinger, M., Benabed, K., Cappe, O., et al. 2011, ArXiv e-prints [arXiv:1101.0950]
Kitazawa, N., & Sagnotti, A. 2014, JCAP, 04, 017
Kitazawa, N., & Sagnotti, A. 2015, Modern Phys. Lett. A, 30, 1550137
Knox, L., & Page, L. 2000, Phys. Rev. Lett., 85, 1366
Kogut, A., Spergel, D. N., Barnes, C., et al. 2003, ApJS, 148, 161
Kovac, J. M., Leitch, E. M., Pryke, C., et al. 2002, Nature, 420, 772
Lesgourgues, J. 2011, ArXiv e-prints [arXiv:1104.2932]
Lewis, A., & Bridle, S. 2002, Phys. Rev., D66, 103511
Mangilli, A., Plaszczynski, S., & Tristram, M. 2015, MNRAS, 453, 3174
Mocanu, L. M., Crawford, T. M., Vieira, J. D., et al. 2013, ApJ, 779, 61
Naess, S., Hasselfield, M., McMahon, J., et al. 2014, JCAP, 10, 07
Netterfield, C. B., Devlin, M. J., Jarosik, N., Page, L., & Wollack, E. J. 1997, ApJ, 474, 47
Page, L., Nolta, M. R., Barnes, C., et al. 2003, ApJS, 148, 233
Page, L., Hinshaw, G., Komatsu, E., et al. 2007, ApJS, 170, 335
Pearson, T. J., Mason, B. S., Readhead, A. C. S., et al. 2003, ApJ, 591, 556
Peiris, H. V., Komatsu, E., Verde, L., et al. 2003, ApJS, 148, 213
Percival, W. J., & Brown, M. L. 2006, MNRAS, 372, 1104
Planck Collaboration. 2015, The Explanatory Supplement to the *Planck* 2015 results, http://wiki.cosmos.esa.int/planckpla/index.php/Main_Page (ESA)
Planck Collaboration VI. 2014, A&A, 571, A6

Planck Collaboration VII. 2014, A&A, 571, A7
Planck Collaboration XI. 2014, A&A, 571, A11
Planck Collaboration XII. 2014, A&A, 571, A12
Planck Collaboration XIII. 2014, A&A, 571, A13
Planck Collaboration XV. 2014, A&A, 571, A15
Planck Collaboration XVI. 2014, A&A, 571, A16
Planck Collaboration XXI. 2014, A&A, 571, A21
Planck Collaboration XXX. 2014, A&A, 571, A30
Planck Collaboration XXXI. 2014, A&A, 571, A31
Planck Collaboration I. 2016, A&A, 594, A1
Planck Collaboration II. 2016, A&A, 594, A2
Planck Collaboration III. 2016, A&A, 594, A3
Planck Collaboration IV. 2016, A&A, 594, A4
Planck Collaboration V. 2016, A&A, 594, A5
Planck Collaboration VI. 2016, A&A, 594, A6
Planck Collaboration VII. 2016, A&A, 594, A7
Planck Collaboration VIII. 2016, A&A, 594, A8
Planck Collaboration IX. 2016, A&A, 594, A9
Planck Collaboration X. 2016, A&A, 594, A10
Planck Collaboration XI. 2016, A&A, 594, A11
Planck Collaboration XII. 2016, A&A, 594, A12
Planck Collaboration XIII. 2016, A&A, 594, A13
Planck Collaboration XIV. 2016, A&A, 594, A14
Planck Collaboration XV. 2016, A&A, 594, A15
Planck Collaboration XVI. 2016, A&A, 594, A16
Planck Collaboration XVII. 2016, A&A, 594, A17
Planck Collaboration XVIII. 2016, A&A, 594, A18
Planck Collaboration XIX. 2016, A&A, 594, A19
Planck Collaboration XX. 2016, A&A, 594, A20
Planck Collaboration XXI. 2016, A&A, 594, A21
Planck Collaboration XXII. 2016, A&A, 594, A22
Planck Collaboration XXIII. 2016, A&A, 594, A23
Planck Collaboration XXIV. 2016, A&A, 594, A24
Planck Collaboration XXV. 2016, A&A, 594, A25
Planck Collaboration XXVI. 2016, A&A, 594, A26
Planck Collaboration XXVII. 2016, A&A, 594, A27
Planck Collaboration XXVIII. 2016, A&A, 594, A28
Planck Collaboration Int. XVI. 2014, A&A, 566, A54
Planck Collaboration Int. XIX. 2015, A&A, 576, A104
Planck Collaboration Int. XXX. 2016, A&A, 586, A133
Planck Collaboration Int. XLVI. 2016, A&A, submitted
Polarbear Collaboration, Ade, P. A. R., Akiba, Y., et al. 2014, ApJ, 794, 171
Polenta, G., Marinucci, D., Balbi, A., et al. 2005, JCAP, 11, 01
Pryke, C., Ade, P., Bock, J., et al. 2009, ApJ, 692, 1247
Readhead, A. C. S., Myers, S. T., Pearson, T. J., et al. 2004, Science, 306, 836
Reichardt, C. L., Ade, P. A. R., Bock, J. J., et al. 2009, ApJ, 694, 1200
Reichardt, C. L., Shaw, L., Zahn, O., et al. 2012, ApJ, 755, 70
Rocha, G., Contaldi, C. R., Bond, J. R., & Gorski, K. M. 2011, MNRAS, 414, 823
Rocha, G., Contaldi, C. R., Colombo, L. P. L., et al. 2010, arXiv e-print [arXiv:1008.4948]
Rosset, C., Tristram, M., Ponthieu, N., et al. 2010, A&A, 520, A13
Rudjord, Ø., Groeneboom, N. E., Eriksen, H. K., et al. 2009, ApJ, 692, 1669
Scott, D., & White, M. 1994, in CMB Anisotropies Two Years after COBE: Observations, Theory and the Future, ed. L. M. Krauss, 214
Seljak, U. 1997, ApJ, 482, 6
Shaw, L. D., Rudd, D. H., & Nagai, D. 2012, ApJ, 756, 15
Sievers, J. L., Achermann, C., Bond, J. R., et al. 2007, ApJ, 660, 976
Souradeep, T., & Ratra, B. 2001, ApJ, 560, 28
Spergel, D., Flauger, R., & Hlozek, R. 2015, Phys. Rev. D, 91, 023518
Story, K. T., Reichardt, C. L., Hou, Z., et al. 2013, ApJ, 779, 86
Taburet, N., Hernández-Monteagudo, C., Aghanim, N., Douspis, M., & Sunyaev, R. A. 2011, MNRAS, 418, 2207
Tegmark, M. 1997, Phys. Rev. D, 55, 5895
Tegmark, M., & de Oliveira-Costa, A. 2001, Phys. Rev. D, 64, 063001
Tinker, J., Kravtsov, A. V., Klypin, A., et al. 2008, ApJ, 688, 709
Trac, H., Bode, P., & Ostriker, J. P. 2011, ApJ, 727, 94
Tristram, M., Macías-Pérez, J. F., Renault, C., & Santos, D. 2005a, MNRAS, 358, 833
Tristram, M., Patanchon, G., Macías-Pérez, J. F., et al. 2005b, A&A, 436, 785
Tucci, M., Toffolatti, L., de Zotti, G., & Martínez-González, E. 2011, A&A, 533, A57
Wandelt, B. D., Hivon, E., & Górski, K. M. 2001, Phys. Rev. D, 64, 083003
Wright, E. L., Bennett, C. L., Górski, K., Hinshaw, G., & Smoot, G. F. 1996, ApJ, 464, L21
Zaldarriaga, M., & Seljak, U. 1997, Phys. Rev. D, 55, 1830

[1] APC, AstroParticule et Cosmologie, Université Paris Diderot, CNRS/IN2P3, CEA/lrfu, Observatoire de Paris, Sorbonne Paris Cité, 10 rue Alice Domon et Léonie Duquet, 75205 Paris Cedex 13, France
[2] Aalto University Metsähovi Radio Observatory and Dept of Radio Science and Engineering, PO Box 13000, 00076 Aalto, Finland
[3] African Institute for Mathematical Sciences, 6-8 Melrose Road, Muizenberg 7945, Cape Town, South Africa
[4] Agenzia Spaziale Italiana Science Data Center, via del Politecnico snc, 00133 Roma, Italy
[5] Aix-Marseille Université, CNRS, LAM (Laboratoire d'Astrophysique de Marseille) UMR 7326, 13388 Marseille, France
[6] Astrophysics Group, Cavendish Laboratory, University of Cambridge, J J Thomson Avenue, Cambridge CB3 0HE, UK
[7] Astrophysics & Cosmology Research Unit, School of Mathematics, Statistics & Computer Science, University of KwaZulu-Natal, Westville Campus, Private Bag X54001, 4000 Durban, South Africa
[8] CGEE, SCS Qd 9, Lote C, Torre C, 4° andar, Ed. Parque Cidade Corporate, CEP 70308-200, Brasília, DF, Brazil
[9] CITA, University of Toronto, 60 St. George St., Toronto, ON M5S 3H8, Canada
[10] CNRS, IRAP, 9 Av. colonel Roche, BP 44346, 31028 Toulouse Cedex 4, France
[11] CRANN, Trinity College, Dublin, Ireland
[12] California Institute of Technology, Pasadena, California, USA
[13] Centre for Theoretical Cosmology, DAMTP, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, UK
[14] Centro de Estudios de Física del Cosmos de Aragón (CEFCA), Plaza San Juan, 1, planta 2, 44001 Teruel, Spain
[15] Computational Cosmology Center, Lawrence Berkeley National Laboratory, Berkeley, California, USA
[16] DSM/Irfu/SPP, CEA-Saclay, 91191 Gif-sur-Yvette Cedex, France
[17] DTU Space, National Space Institute, Technical University of Denmark, Elektrovej 327, 2800 Kgs. Lyngby, Denmark
[18] Département de Physique Théorique, Université de Genève, 24 quai E. Ansermet, 1211 Genève 4, Switzerland
[19] Departamento de Astrofísica, Universidad de La Laguna (ULL), 38206 La Laguna, Tenerife, Spain
[20] Departamento de Física, Universidad de Oviedo, Avda. Calvo Sotelo s/n, 33007 Oviedo, Spain
[21] Department of Astrophysics/IMAPP, Radboud University Nijmegen, PO Box 9010, 6500 GL Nijmegen, The Netherlands
[22] Department of Physics & Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, British Columbia, Canada
[23] Department of Physics and Astronomy, Dana and David Dornsife College of Letter, Arts and Sciences, University of Southern California, Los Angeles, CA 90089, USA
[24] Department of Physics and Astronomy, University College London, London WC1E 6BT, UK
[25] Department of Physics and Astronomy, University of Sussex, Brighton BN1 9QH, UK
[26] Department of Physics, Florida State University, Keen Physics Building, 77 Chieftan Way, Tallahassee, Florida, USA
[27] Department of Physics, Gustaf Hällströmin katu 2a, University of Helsinki, 00014 Helsinki, Finland
[28] Department of Physics, Princeton University, Princeton, NJ 08544, USA
[29] Department of Physics, University of California, One Shields Avenue, Davis, CA 95616, USA
[30] Department of Physics, University of California, Santa Barbara, CA 93106, USA
[31] Department of Physics, University of Illinois at Urbana-Champaign, 1110 West Green Street, Urbana, Illinois, USA

[32] Dipartimento di Fisica e Astronomia G. Galilei, Università degli Studi di Padova, via Marzolo 8, 35131 Padova, Italy
[33] Dipartimento di Fisica e Scienze della Terra, Università di Ferrara, via Saragat 1, 44122 Ferrara, Italy
[34] Dipartimento di Fisica, Università La Sapienza, P.le A. Moro 2, 00185 Roma, Italy
[35] Dipartimento di Fisica, Università degli Studi di Milano, via Celoria, 16, 20133 Milano, Italy
[36] Dipartimento di Fisica, Università degli Studi di Trieste, via A. Valerio 2, 34127 Trieste, Italy
[37] Dipartimento di Matematica, Università di Roma Tor Vergata, via della Ricerca Scientifica 1, 00133 Roma, Italy
[38] Discovery Center, Niels Bohr Institute, Blegdamsvej 17, 2100 Copenhagen, Denmark
[39] Discovery Center, Niels Bohr Institute, Copenhagen University, Blegdamsvej 17, 2100 Copenhagen, Denmark
[40] European Space Agency, ESAC, Planck Science Office, Camino bajo del Castillo, s/n, Urbanización Villafranca del Castillo, 28691 Villanueva de la Cañada, Madrid, Spain
[41] European Space Agency, ESTEC, Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands
[42] Gran Sasso Science Institute, INFN, viale F. Crispi 7, 67100 L'Aquila, Italy
[43] HGSFP and University of Heidelberg, Theoretical Physics Department, Philosophenweg 16, 69120 Heidelberg, Germany
[44] Haverford College Astronomy Department, 370 Lancaster Avenue, Haverford, PA 19041, USA
[45] Helsinki Institute of Physics, Gustaf Hällströmin katu 2, University of Helsinki, 00014 Helsinki, Finland
[46] INAF–Osservatorio Astrofisico di Catania, via S. Sofia 78, Catania, Italy
[47] INAF–Osservatorio Astronomico di Padova, Vicolo dell'Osservatorio 5, 35122 Padova, Italy
[48] INAF– Osservatorio Astronomico di Roma, via di Frascati 33, 00040 Monte Porzio Catone, Italy
[49] INAF–Osservatorio Astronomico di Trieste, via G.B. Tiepolo 11, 40127 Trieste, Italy
[50] INAF/IASF Bologna, via Gobetti 101, 40129 Bologna, Italy
[51] INAF/IASF Milano, via E. Bassini 15, 20133 Milano, Italy
[52] INFN, Sezione di Bologna, viale Berti Pichat 6/2, 40127 Bologna, Italy
[53] INFN, Sezione di Ferrara, via Saragat 1, 44122 Ferrara, Italy
[54] INFN, Sezione di Roma 1, Università di Roma Sapienza, P.le Aldo Moro 2, 00185 Roma, Italy
[55] INFN, Sezione di Roma 2, Università di Roma Tor Vergata, via della Ricerca Scientifica 1, 00185 Roma, Italy
[56] INFN/National Institute for Nuclear Physics, via Valerio 2, 34127 Trieste, Italy
[57] IPAG: Institut de Planétologie et d'Astrophysique de Grenoble, Université Grenoble Alpes, IPAG; CNRS, IPAG, 38000 Grenoble, France
[58] IUCAA, Post Bag 4, Ganeshkhind, Pune University Campus, 411 007 Pune, India
[59] Imperial College London, Astrophysics group, Blackett Laboratory, Prince Consort Road, London, SW7 2AZ, UK
[60] Infrared Processing and Analysis Center, California Institute of Technology, Pasadena, CA 91125, USA
[61] Institut Néel, CNRS, Université Joseph Fourier Grenoble I, 25 rue des Martyrs, 38042 Grenoble, France
[62] Institut Universitaire de France, 103 bd Saint-Michel, 75005 Paris, France
[63] Institut d'Astrophysique Spatiale, CNRS, Univ. Paris-Sud, Université Paris-Saclay, Bât. 121, 91405 Orsay Cedex, France
[64] Institut d'Astrophysique de Paris, CNRS (UMR 7095), 98bis boulevard Arago, 75014 Paris, France
[65] Institut für Theoretische Teilchenphysik und Kosmologie, RWTH Aachen University, 52056 Aachen, Germany
[66] Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK
[67] Institute of Theoretical Astrophysics, University of Oslo, Blindern, 0371 Oslo, Norway
[68] Instituto de Astrofísica de Canarias, C/Vía Láctea s/n, La Laguna, 38205 Tenerife, Spain
[69] Instituto de Física de Cantabria (CSIC-Universidad de Cantabria), Avda. de los Castros s/n, 93005 Santander, Spain
[70] Istituto Nazionale di Fisica Nucleare, Sezione di Padova, via Marzolo 8, 35131 Padova, Italy
[71] Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 31109, USA
[72] Jodrell Bank Centre for Astrophysics, Alan Turing Building, School of Physics and Astronomy, The University of Manchester, Oxford Road, Manchester, M13 9PL, UK
[73] Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA
[74] Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge, CB3 0HA, UK
[75] Kazan Federal University, 18 Kremlyovskaya St., 420008 Kazan, Russia
[76] LAL, Université Paris-Sud, CNRS/IN2P3, 91898 Orsay, France
[77] LERMA, CNRS, Observatoire de Paris, 61 avenue de l'Observatoire, 75014 Paris, France
[78] Laboratoire AIM, IRFU/Service d'Astrophysique – CEA/DSM – CNRS – Université Paris Diderot, Bât. 709, CEA-Saclay, 91191 Gif-sur-Yvette Cedex, France
[79] Laboratoire Traitement et Communication de l'Information, CNRS (UMR 5141) and Télécom ParisTech, 46 rue Barrault 75634 Paris Cedex 13, France
[80] Laboratoire de Physique Subatomique et Cosmologie, Université Grenoble-Alpes, CNRS/IN2P3, 53 rue des Martyrs, 38026 Grenoble Cedex, France
[81] Laboratoire de Physique Théorique, Université Paris-Sud 11 & CNRS, Bâtiment 210, 91405 Orsay, France
[82] Lebedev Physical Institute of the Russian Academy of Sciences, Astro Space Centre, 84/32 Profsoyuznaya st., GSP-7, 117997 Moscow, Russia
[83] Leung Center for Cosmology and Particle Astrophysics, National Taiwan University, 10617 Taipei, Taiwan
[84] Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85741 Garching, Germany
[85] National University of Ireland, Department of Experimental Physics, Maynooth, Co. Kildare, Ireland
[86] Nicolaus Copernicus Astronomical Center, Bartycka 18, 00-716 Warsaw, Poland
[87] Niels Bohr Institute, Blegdamsvej 17, 2100 Copenhagen, Denmark
[88] Niels Bohr Institute, Copenhagen University, Blegdamsvej 17, 2100 Copenhagen, Denmark
[89] Nordita (Nordic Institute for Theoretical Physics), Roslagstullsbacken 23, 106 91 Stockholm, Sweden
[90] Optical Science Laboratory, University College London, Gower Street, London, UK
[91] SISSA, Astrophysics Sector, via Bonomea 265, 34136 Trieste, Italy
[92] SMARTEST Research Centre, Università degli Studi e-Campus, via Isimbardi 10, 22060 Novedrate (CO), Italy
[93] School of Physics and Astronomy, Cardiff University, Queens Buildings, The Parade, Cardiff, CF24 3AA, UK
[94] Sorbonne Université-UPMC, UMR 7095, Institut d'Astrophysique de Paris, 98bis boulevard Arago, 75014 Paris, France
[95] Space Research Institute (IKI), Russian Academy of Sciences, Profsoyuznaya Str, 84/32, 117997 Moscow, Russia
[96] Space Sciences Laboratory, University of California, Berkeley, CA 94720, USA
[97] Special Astrophysical Observatory, Russian Academy of Sciences, Nizhnij Arkhyz, Zelenchukskiy region, 369167 Karachai-Cherkessian Republic, Russia
[98] Sub-Department of Astrophysics, University of Oxford, Keble Road, Oxford OX1 3RH, UK
[99] Sydney Institute for Astronomy, School of Physics A28, University of Sydney, NSW 2006, Australia

[100] The Oskar Klein Centre for Cosmoparticle Physics, Department of Physics, Stockholm University, AlbaNova, 106 91 Stockholm, Sweden
[101] Theory Division, PH-TH, CERN, 1211 Geneva 23, Switzerland
[102] UPMC Univ. Paris 06, UMR 7095, 98bis boulevard Arago, 75014 Paris, France
[103] Université de Toulouse, UPS-OMP, IRAP, 31028 Toulouse Cedex 4, France
[104] University Observatory, Ludwig Maximilian University of Munich, Scheinerstrasse 1, 81679 Munich, Germany
[105] University of Granada, Departamento de Física Teórica y del Cosmos, Facultad de Ciencias, 18071 Granada, Spain
[106] University of Granada, Instituto Carlos I de Física Teórica y Computacional, 18071 Granada, Spain
[107] Warsaw University Observatory, Aleje Ujazdowskie 4, 00-478 Warszawa, Poland

Appendix A: Sky masks

This appendix provides details of the way we build sky masks for the high-ℓ likelihood. Since it is based on data at frequencies between 100 and 217 GHz, Galactic dust emission is the main diffuse foreground to minimize. We subtract the SMICA CMB temperature map (Planck Collaboration IX 2016) from the 353 GHz map and we adopt the resulting CMB-subtracted 353 GHz map as a tracer of dust. After smoothing the map with a 10° Gaussian kernel, we threshold it to generate a sequence of masks with different sky coverage. Galactic masks obtained in this way are named B80 to B50, where the number gives the retained sky fraction $f_{\rm sky}$ in percent (Fig. A.1).

For the likelihood analysis, we aim to find a trade-off between maximizing the sky coverage and having a simple, but reliable, foreground model of the data. The combination of masks and frequency channels retained is given in Table A.1. In order to get C_ℓ-covariance matrices for the cosmological analysis that are accurate at the few percent level (cf. Sect. 3.5), we actually use apodized versions of the Galactic masks. The apodization corresponds to a Gaussian taper of width $\sigma = 2°$ [19]. Apodized Galactic masks are also used for the polarization analysis. The effective sky fraction of an apodized mask is $f_{\rm sky} = \sum_i w_i^2 \Omega_i/(4\pi)$, where w_i is the value of the mask in pixel i and Ω_i is the solid angle of the pixel.

All the HFI frequency channels, except 143 GHz, are also contaminated by CO emission from rotational transition lines. Here we are concerned with emission around 100 and 217 GHz, associated with the CO $J = 1 \rightarrow 0$ and $J = 2 \rightarrow 1$ lines, respectively. Most of the emission is concentrated near the Galactic plane and is therefore masked out by the Galactic dust masks. However, there are some emission regions at intermediate and low latitudes that are outside the quite small B80 mask we use at 100 GHz. We therefore create a mask specifically targeted at eliminating CO emission. The Type 3 CO map, part of the *Planck* 2013 product delivery (Planck Collaboration XIII 2014), is sensitive to low-intensity diffuse CO emission over the whole sky. It is a multi-line map, derived using prior information on line ratios and a multi-frequency component separation method. Of the three types of *Planck* CO maps, this has the highest S/N. We smooth this map with a $\sigma = 120'$ Gaussian and mask the sky wherever the CO line brightness exceeds 1 $K_{\rm RJ}$ km s^{-1}. The mask is shown in Fig. A.2, before apodization with a Gaussian taper of $FWHM = 30'$.

Finally, we include extragalactic objects in our temperature masks, both point sources and nearby extended galaxies. The nearby galaxies that are masked are listed in Table A.2, together with the corresponding cut radii. For point sources, we build conservative masks for 100, 143, and 217 GHz separately. At each frequency, we mask sources that are detected above $S/N = 5$ in the 2015 point-source catalogue (Planck Collaboration XXVI 2016) with holes of radius three times the $\sigma = FWHM/\sqrt{\ln 8}$ of the effective Gaussian beam at that frequency. We take the FWHM values from the elliptic Gaussian fits to the effective beams (Planck Collaboration XXVI 2016), i.e., FWHM values of 9.′66, 7.′22, and 4.′90 at 100, 143, and 217 GHz, respectively. We apodize these masks with a Gaussian taper of $FWHM = 30'$. As already noted, these masks are designed to reduce the contribution of diffuse and discrete Galactic and extragalactic foreground emission in the "raw" (half-mission and detset) frequency maps used for the baseline high-ℓ likelihood.

The masks described in this appendix are used in the papers on cosmological parameters (Planck Collaboration XIII 2016), inflation (Planck Collaboration XX 2016), dark energy (Planck Collaboration XIV 2016), and primordial magnetic fields (Planck Collaboration XIX 2016), which are notable examples of the application of the high-ℓ likelihood. However, the masks differ from those adopted in some of the other *Planck* papers. For example, reconstructions of gravitational lensing (Planck Collaboration XV 2016) and integrated Sachs-Wolfe effect (Planck Collaboration XXI 2016), constraints on isotropy and statistics (Planck Collaboration XVI 2016), and searches for primordial non-Gaussianity (Planck Collaboration XVII 2016) mainly rely on the high-resolution foreground-reduced CMB maps presented in Planck Collaboration IX (2016). Those maps have been derived by four component-separation methods that combine data from different frequency channels to extract "cleaned" CMB maps. For each method, the corresponding confidence masks, for both temperature and polarization, remove regions of the sky where the CMB solution is not trusted. This is described in detail in Appendices A–D of Planck Collaboration IX (2016). The masks recommended for the analysis of foreground-reduced CMB maps are constructed

[19] We use the routine `process_mask` of the `HEALPix` package to obtain a map of the distance of each pixel of the mask from the closest null pixel. We then use a smoothed version of the distance map to build the Gaussian apodization. The smoothing of the distance map is needed to avoid sharp edges in the final mask.

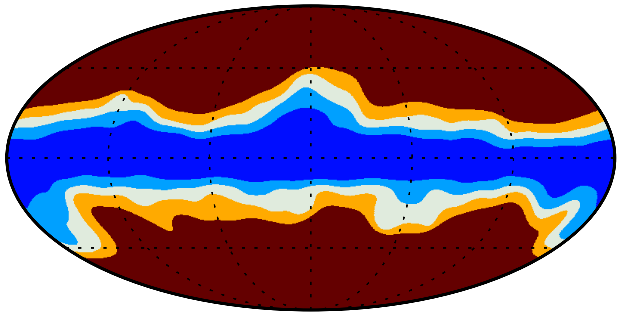

Fig. A.1. Unapodized Galactic masks B50, B60, B70, and B80, from orange to dark blue.

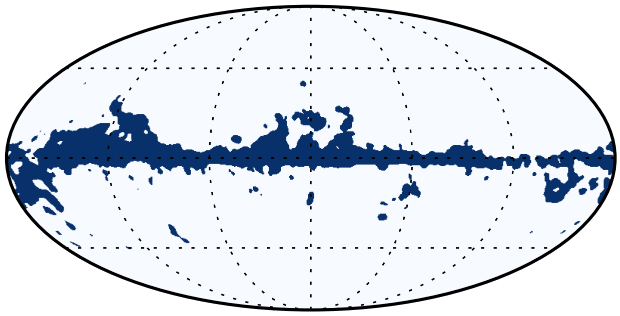

Fig. A.2. Unapodized CO mask ($f_{\rm sky} = 87\%$).

Table A.1. Galactic masks used for the high-ℓ analysis.

Frequency [GHz]	Temperature		Polarization	
100	B80	G70	B80	G70
143	B70	G60	B60	G50
217	B60	G50	B50	G41

Notes. For each frequency channel, the Galactic and apodized Galactic masks are labelled by their "B" and "G" prefixes, followed by the retained sky fraction (in percent).

Table A.2. Masked nearby galaxies and corresponding cut radii.

Galaxy	Radius [arcmin]
LMC	250
SMC	110
SMC ext[a]	50
M 31 F1[b]	80
M 31 F2	80
M 33	30
M 81	30
M 101	18
M 82	15
M 51	15
Cen A	15

Notes. [(a)] Inspection of the SMC at 857 GHz reveals an extra signal, localized in a small area near the border of the excised disk, which we mask with a disk centred at $(l, b) = (299.^\circ85, -43.^\circ6)$. [(b)] M 31 is elongated. Therefore, instead of cutting an unnecessarily large disk, we use two smaller disks centred at the focal points of an elliptical fit to the galaxy image (F1, F2).

as the unions of the confidence masks of all the four component separation methods. Their sky coverages are $f_{\rm sky} = 0.776$ in temperature and $f_{\rm sky} = 0.774$ in polarization. Since component separation mitigates the foreground contamination even at relatively low Galactic latitudes, those masks feature a thinner cut along the Galactic plane than the ones described in this appendix. Nevertheless, propagation of noise, beam, and extragalactic foreground uncertainties in foreground-cleaned CMB maps is more difficult, and this is the main reason why we do not employ them in the baseline high-ℓ likelihood. We also note that the recommended mask for temperature foreground-reduced maps has a greater number of compact object holes than the masks used here. This is due to the fact that some component separation techniques can introduce contamination of sources from a wider range of frequencies than the approach considered here for the high-ℓ power spectra. According to the tests provided in Sect. C.1.4, such masks would result in sub-optimal performance of the analytic C_ℓ-covariance matrices.

Appendix B: Low-ℓ likelihood supplement

B.1. Sherman-Morrison-Woodbury formula

In the *Planck* 2015 release we follow a pixel-based approach to the joint low-ℓ likelihood (up to $\ell = 29$) of T, Q, and U. This approach treats temperature and polarization maps consistently at `HEALPix` resolution $N_{\rm side} = 16$, as opposed to the WMAP low-ℓ likelihood, which incorporates polarization information from lower-resolution maps to save computational time (Page et al. 2007). The disadvantage of a consistent-resolution, brute-force approach lies in its computational cost (Like13), which may require massively parallel coding (and adequate hardware) in order to be competitive in execution time with the high-ℓ part of the CMB likelihood (see, e.g., Finelli et al. 2013 for one such implementation). Such a choice, however, would hamper the ease of code distribution across a community not necessarily specialized in massively parallel computing. Luckily, the Sherman-Morrison-Woodbury formula and the related matrix determinant lemma provide a means to achieve good timing without resorting to supercomputers. To see how this works, rewrite the covariance matrix from Eq. (3) in a form that explicitly separates the C_ℓ to be varied from those that stay fixed at the reference model:

$$\mathsf{M} = \sum_{XY} \sum_{\ell=2}^{\ell_{\rm cut}} C_\ell^{XY} \mathsf{P}_\ell^{XY} + \sum_{XY} \sum_{\ell=\ell_{\rm cut}+1}^{\ell_{\rm max}} C_\ell^{XY,{\rm ref}} \mathsf{P}_\ell^{XY} + \mathsf{N} \tag{B.1}$$

$$\equiv \sum_{XY} \sum_{\ell=2}^{\ell_{\rm cut}} C_\ell^{XY} \mathsf{P}_\ell^{XY} + \mathsf{M}_0, \tag{B.2}$$

where we have effectively redefined the fixed multipoles as "high-ℓ correlated noise", as far as the varying low-ℓ multipoles are concerned. Next, note that for fixed ℓ, P_ℓ^{TT} has rank[20] $\lambda = 2\ell + 1$, and this matrix may therefore be decomposed as $\mathsf{P}_\ell^{TT} = (\mathsf{V}_\ell^{TT})^{\sf T} \mathsf{A}_\ell^{TT} \mathsf{V}_\ell^{TT}$, where A_ℓ^{TT} and V_ℓ^{TT} are $(\lambda \times \lambda)$ and $(\lambda \times N_{\rm pix})$ matrices, respectively, which depend only upon the unmasked pixel locations. A similar decomposition holds for the $\mathsf{P}_\ell^{EE,BB}$ matrices, while P_ℓ^{TE} can be expanded in the $[\mathsf{V}_\ell^{TT}, \mathsf{V}_\ell^{EE}]$ basis for the corresponding ℓ. We can then write

$$\mathsf{M} = \mathsf{V}^{\sf T} \mathsf{A}(C_\ell) \mathsf{V} + \mathsf{M}_0, \tag{B.3}$$

where $\mathsf{V} = [\mathsf{V}_2^{TT}, \mathsf{V}_2^{EE}, \mathsf{V}_2^{BB}, \dots \mathsf{V}_{\ell_{\rm cut}}^{BB}]$ is an $(n_\lambda \times N_{\rm pix})$ matrix with $n_\lambda = 3[(\ell_{\rm cut} + 1)^2 - 4]$, and $\mathsf{A}(C_\ell)$ is an $(n_\lambda \times n_\lambda)$ block-diagonal matrix (accounting for four modes removed in monopole and dipole subtraction). Each ℓ-block in the latter matrix reads

$$\begin{bmatrix} C_\ell^{TT} \mathsf{A}_\ell^{TT} & C_\ell^{TE} \mathsf{A}_\ell^{TE} & 0 \\ C_\ell^{TE} \mathsf{A}_\ell^{TE} & C_\ell^{EE} \mathsf{A}_\ell^{EE} & 0 \\ 0 & 0 & C_\ell^{BB} \mathsf{A}_\ell^{BB} \end{bmatrix}. \tag{B.4}$$

Finally, using the Sherman-Morrison-Woodbury identity and the matrix determinant lemma, we can rewrite the inverse and determinant of M as

$$\mathsf{M}^{-1} = \mathsf{M}_0^{-1} - \mathsf{M}_0^{-1} \mathsf{V}^{\sf T} (\mathsf{A}^{-1} + \mathsf{V} \mathsf{M}_0^{-1} \mathsf{V}^{\sf T})^{-1} \mathsf{V} \mathsf{M}_0^{-1} \tag{B.5}$$

$$|\mathsf{M}| = |\mathsf{M}_0| \, |\mathsf{A}| \, |\mathsf{A}^{-1} + \mathsf{V} \mathsf{M}_0^{-1} \mathsf{V}^{\sf T}| \, . \tag{B.6}$$

Because neither V nor M_0 depends on C_ℓ, all terms involving only their inverses, determinants, and products may be precomputed and stored. Evaluating the likelihood for a new set of C_ℓ then requires only the inverse and determinant of an $(n_\lambda \times n_\lambda)$ matrix, not an $(N_{\rm pix} \times N_{\rm pix})$ matrix. For the current data selection, described in Sects. 2.2 and 2.3, we find $n_\lambda = 2688$, which is to be compared to $N_{\rm pix} = 6307$, resulting in an order-of-magnitude speed-up compared to the brute-force computation.

B.2. Lollipop

We performed a complementary analysis of low-ℓ polarization using the HFI data, in order to check the consistency with the LFI-based baseline result. The level of systematic residuals in the HFI maps at low ℓ is quite small, but comparable to the HFI noise (see Planck Collaboration VIII 2016), so these residuals should be either corrected, which is the goal of a future release, or accounted for by a complete analysis including parameters for all relevant systematic effects, which we cannot yet perform. Instead, we use `Lollipop`, a low-ℓ polarized likelihood function based on cross-power spectra. The idea behind this approach is that the systematics are considerably reduced in cross-correlation compared to auto-correlation.

At low multipoles and for incomplete sky coverage, the C_ℓ statistic is not simply distributed and is correlated between modes. `Lollipop` uses the approximation presented

[20] Masking can in principle reduce the effective rank, but for the high sky fractions used in the *Planck* analysis, this is not an issue.

in Hamimeche & Lewis (2008), modified as described in Mangilli et al. (2015) to apply to cross-power spectra. We restrict ourselves to the one-field approximation to derive a likelihood function based only on the EE power spectrum at very low multipoles. The likelihood function of the C_ℓ given the data $\tilde{C}_\ell$ is then

$$-2 \ln P(C_\ell|\tilde{C}_\ell) = \sum_{\ell\ell'} [X_g]^{\mathsf{T}}_\ell [M_f^{-1}]_{\ell\ell'} [X_g]_{\ell'}, \tag{B.7}$$

with the variable

$$\left[X_g\right]_\ell = \sqrt{C_\ell^f + O_\ell}\; g\!\left(\frac{\tilde{C}_\ell + O_\ell}{C_\ell + O_\ell}\right) \sqrt{C_\ell^{\mathrm{fid}} + O_\ell}, \tag{B.8}$$

where $g(x) = \sqrt{2(x - \ln x - 1)}$, C_ℓ^{fid} is a fiducial model and O_ℓ is the offset needed in the case of cross-spectra. This likelihood has been tested on Monte Carlo simulations including both realistic signal and noise. In order to extract cosmological information on τ from the EE spectrum alone, we restrict the analysis to the cross-correlation between the HFI 100 and 143 GHz maps, which exhibits the lowest variance.

At large angular scales, the HFI maps are contaminated by systematic residuals coming from temperature-to-polarization leakage (see Planck Collaboration VIII 2016). We used our best estimate of the Q and U maps at 100 and 143 GHz, which we correct for residual leakage coming from destriping uncertainties, calibration mismatch, and bandpass mismatch, using templates as described in Planck Collaboration VIII (2016). Even though the level of systematic effects is thereby significantly reduced, we still have residuals above the noise level in null tests at very low multipoles ($\ell \leqslant 4$). To mitigate the effect of this on the likelihood, we restrict the range of multipoles to $\ell = 5$–20.

Cross-power spectra are computed on the cleanest 50% of the sky by using a pseudo-C_ℓ estimate (`Xpol`, an extension to polarization of the code described in Tristram et al. 2005a). The mask corresponds to thresholding a map of the diffuse polarized Galactic dust at large scales. In addition, we also removed pixels where the intensity of diffuse Galactic dust and CO lines is strong. This ensures that bandpass leakage from dust and CO lines does not bias the polarization spectra (see Planck Collaboration VIII 2016).

We construct the C_ℓ correlation matrix using simulations including CMB signal and realistic inhomogeneous and correlated noise. In order to take into account the residual systematics, we derive the noise level from the estimated BB auto-spectrum where we neglect any possible cosmological signal. This overestimates the noise level and ensures conservative errors. However, this estimate assumes by construction a Gaussian noise contribution, which is not a full description of the residuals.

We then sample the reionization optical depth τ from the likelihood, with all other parameters fixed to the *Planck* 2015 best-fit values (Planck Collaboration XIII 2016). Without any other data, the degeneracy between A_s and τ is broken by fixing the amplitude of the first peak of the TT spectrum (directly related to $A_\mathrm{s}\mathrm{e}^{-2\tau}$) at $\ell = 200$. The resulting distribution is plotted in Fig. B.1. The best fit is at

$$\tau = 0.064^{+0.015}_{-0.016}, \qquad z_\mathrm{re} = 8.7^{+1.4}_{-1.6}, \tag{B.9}$$

in agreement with the current *Planck* low-ℓ baseline (see Table 2), even though this result only relies on the EE spectrum between $\ell = 5$ and 20.

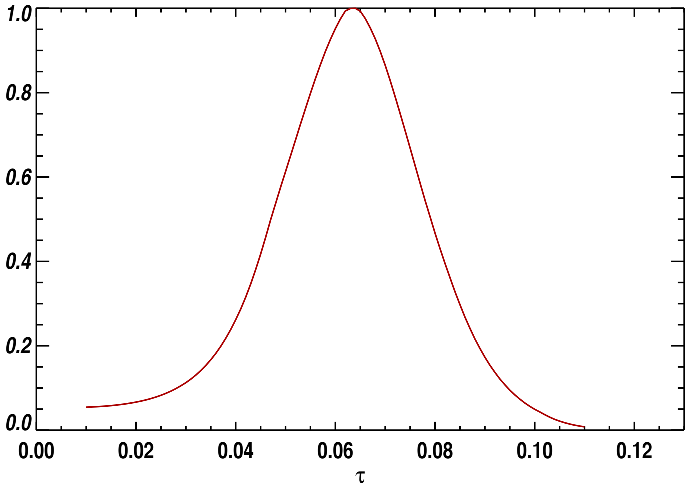

Fig. B.1. Distribution of the reionization optical depth τ using the `Lollipop` likelihood, based on the cross-correlation of the 100 and 143 GHz channels.

Appendix C: High-ℓ baseline likelihood: `Plik`

In this appendix, we provide detailed information on the `Plik` baseline likelihood used at high ℓ. First we describe in Sect. C.1 the `Plik` covariance matrix, by providing the equations we have implemented, by giving results from some of the numerical tests we carried out, and by describing our procedure to deal with the excess variance (as compared to the prediction of our approximate analytical model) due to the point source mask. Section C.2 validates the overall `Plik` implementation with Monte Carlo simulations of the full mission. For reference, Sect. C.3 gives the results of a large body of validation and stability tests on the actual data, including polarization in particular. We also discuss the numerical agreement of the temperature- and polarization-based results on base-ΛCDM parameters. Section C.4 describes how we calculate co-added CMB spectra from foreground-cleaned frequency power spectra. Section C.5 compares `Plik` cosmological results obtained using the `PICO` or `CAMB` codes. Finally, Sect. C.6 details how we marginalize over nuisance parameters to provide a fast but accurate CMB-only likelihood.

C.1. Covariance matrix

C.1.1. Structure of the covariance matrix

Here we summarize the mathematical formalism implemented to calculate the pseudo-power spectrum covariance matrices for temperature and polarization.

In the following, the fiducial power spectra C_ℓ are assumed to be the smooth theory spectra multiplied by beam (b) and pixel window function (p) for detectors i and j,

$$C_\ell^{i,j} = b_\ell^i\, b_\ell^j\, p_\ell^2 \left(C_\ell^{\mathrm{CMB}} + C_\ell^{\mathrm{FG}}(f_i, f_j)\right), \tag{C.1}$$

where the f_k denote the frequency dependence of the foreground contribution.

We now present the equations used to compute all the unique covariance matrix polarization blocks that can be formed from temperature and E-mode polarization maps (Hansen et al. 2002; Hinshaw et al. 2003; Efstathiou 2004; Challinor & Chon 2005; Like13). They approximate the variance of the biased pseudo-power spectrum coefficients, before correcting for the effects of pixel window function, beam, and mask.

TTTT block:

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

TTTE block:

$$
\begin{aligned}
&\mathrm{Var}(\hat{C}_\ell^{TT\,i,j},\hat{C}_{\ell'}^{TE\,p,q}) \\
&\approx \frac{1}{2}\sqrt{C_\ell^{TT\,i,p}C_{\ell'}^{TT\,i,p}}\left(C_\ell^{TE\,j,q}+C_{\ell'}^{TE\,j,q}\right) \\
&\times \Xi_{TT}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,p)^{TT},(j,q)^{TP}\right]_{\ell\ell'} \\
&+ \frac{1}{2}\sqrt{C_\ell^{TT\,j,p}C_{\ell'}^{TT\,j,p}}\left(C_\ell^{TE\,i,q}+C_{\ell'}^{TE\,i,q}\right) \\
&\times \Xi_{TT}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,q)^{TP},(j,p)^{TT}\right]_{\ell\ell'} \\
&+ \frac{1}{2}\left(C_\ell^{TE\,j,q}+C_{\ell'}^{TE\,j,q}\right)\Xi_{TT}^{\emptyset\emptyset,TT}\left[(j,q)^{TP},(i,p)^{TT}\right]_{\ell\ell'} \\
&+ \frac{1}{2}\left(C_\ell^{TE\,i,q}+C_{\ell'}^{TE\,i,q}\right)\Xi_{TT}^{\emptyset\emptyset,TT}\left[(i,q)^{TP},(j,p)^{TT}\right]_{\ell\ell'}.
\end{aligned}
\tag{C.3}
$$

TETE block

$$
\begin{aligned}
&\mathrm{Var}(\hat{C}_\ell^{TE\,i,j},\hat{C}_{\ell'}^{TE\,p,q}) \\
&\approx \sqrt{C_\ell^{TT\,i,p}C_{\ell'}^{TT\,i,p}C_\ell^{EE\,j,q}C_{\ell'}^{EE\,j,q}}\;\Xi_{TE}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,p)^{TT},(j,q)^{PP}\right]_{\ell\ell'} \\
&+ \frac{1}{2}\left(C_\ell^{TE\,i,q}C_{\ell'}^{TE\,j,p}+C_\ell^{TE\,j,p}C_{\ell'}^{TE\,i,q}\right)\Xi_{TT}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,q)^{TP},(j,p)^{PT}\right]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{TT\,i,p}C_{\ell'}^{TT\,i,p}}\;\Xi_{TE}^{\emptyset\emptyset,PP}\left[(i,p)^{TT},(j,q)^{PP}\right]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{EE\,j,q}C_{\ell'}^{EE\,j,q}}\;\Xi_{TE}^{\emptyset\emptyset,TT}\left[(j,q)^{PP},(i,p)^{TT}\right]_{\ell\ell'} \\
&+ \Xi_{TE}^{TT,PP}\left[(i,p)^{TT},(j,q)^{PP}\right]_{\ell\ell'}.
\end{aligned}
\tag{C.4}
$$

TTEE block:

$$
\begin{aligned}
&\mathrm{Var}(\hat{C}_\ell^{TT\,i,j},\hat{C}_{\ell'}^{EE\,p,q}) \\
&\approx \frac{1}{2}\left(C_\ell^{TE\,i,p}C_{\ell'}^{TE\,j,q}+C_\ell^{TE\,j,q}C_{\ell'}^{TE\,i,p}\right)\Xi_{TT}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,p)^{TP},(j,q)^{TP}\right]_{\ell\ell'} \\
&+ \frac{1}{2}\left(C_\ell^{TE\,i,q}C_{\ell'}^{TE\,j,p}+C_\ell^{TE\,j,p}C_{\ell'}^{TE\,i,q}\right)\Xi_{TT}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,q)^{TP},(j,p)^{TP}\right]_{\ell\ell'}
\end{aligned}
\tag{C.5}
$$

TEEE block:

$$
\begin{aligned}
&\mathrm{Var}(\hat{C}_\ell^{TE\,i,j},\hat{C}_{\ell'}^{EE\,p,q}) \\
&\approx \frac{1}{2}\sqrt{C_\ell^{EE\,j,q}C_{\ell'}^{EE\,j,q}}\left(C_\ell^{TE\,i,p}+C_{\ell'}^{TE\,i,p}\right)\Xi_{EE}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,p)^{TP},(j,q)^{PP}\right]_{\ell\ell'} \\
&+ \frac{1}{2}\sqrt{C_\ell^{EE\,j,p}C_{\ell'}^{EE\,j,p}}\left(C_\ell^{TE\,i,q}+C_{\ell'}^{TE\,i,q}\right)\Xi_{EE}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,q)^{TP},(j,p)^{PP}\right]_{\ell\ell'} \\
&+ \frac{1}{2}\left(C_\ell^{TE\,i,p}+C_{\ell'}^{TE\,i,p}\right)\Xi_{EE}^{\emptyset\emptyset,PP}\left[(i,p)^{TP},(j,q)^{PP}\right]_{\ell\ell'} \\
&+ \frac{1}{2}\left(C_\ell^{TE\,i,q}+C_{\ell'}^{TE\,i,q}\right)\Xi_{EE}^{\emptyset\emptyset,PP}\left[(i,q)^{TP},(j,p)^{PP}\right]_{\ell\ell'}.
\end{aligned}
\tag{C.6}
$$

EEEE block:

$$
\begin{aligned}
&\mathrm{Var}(\hat{C}_\ell^{EE\,i,j},\hat{C}_{\ell'}^{EE\,p,q}) \\
&\approx \sqrt{C_\ell^{EE\,i,p}C_{\ell'}^{EE\,i,p}C_\ell^{EE\,j,q}C_{\ell'}^{EE\,j,q}}\;\Xi_{EE}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,p)^{PP},(j,q)^{PP}\right]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{EE\,i,q}C_{\ell'}^{EE\,i,q}C_\ell^{EE\,j,p}C_{\ell'}^{EE\,j,p}}\;\Xi_{EE}^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,q)^{PP},(j,p)^{PP}\right]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{EE\,i,p}C_{\ell'}^{EE\,i,p}}\;\Xi_{EE}^{\emptyset\emptyset,PP}\left[(i,p)^{PP},(j,q)^{PP}\right]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{EE\,j,q}C_{\ell'}^{EE\,j,q}}\;\Xi_{EE}^{\emptyset\emptyset,PP}\left[(j,q)^{PP},(i,p)^{PP}\right]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{EE\,i,q}C_{\ell'}^{EE\,i,q}}\;\Xi_{EE}^{\emptyset\emptyset,PP}\left[(i,q)^{PP},(j,p)^{PP}\right]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{EE\,j,p}C_{\ell'}^{EE\,j,p}}\;\Xi_{EE}^{\emptyset\emptyset,PP}\left[(j,p)^{PP},(i,q)^{PP}\right]_{\ell\ell'} \\
&+ \Xi_{EE}^{PP,PP}\left[(i,p)^{PP},(j,q)^{PP}\right]_{\ell\ell'} + \Xi_{EE}^{PP,PP}\left[(i,q)^{PP},(j,p)^{PP}\right]_{\ell\ell'}.
\end{aligned}
\tag{C.7}
$$

In Eqs. (C.2)–(C.7), we have introduced the projector functions Ξ_{TT}, Ξ_{EE}, and Ξ_{TE} to describe the coupling between multipoles induced by the mask,

$$
\Xi_{TT}^{X,Y}\left[(i,j)^\alpha,(p,q)^\beta\right]_{\ell_1\ell_2} = \sum_{\ell_3}\frac{2\ell_3+1}{4\pi}\begin{pmatrix}\ell_1 & \ell_2 & \ell_3\\ 0 & 0 & 0\end{pmatrix}^2 \times W^{X,Y}\left[(i,j)^\alpha,(p,q)^\beta\right]_{\ell_3}, \tag{C.8}
$$

$$
\Xi_{EE}^{X,Y}\left[(i,j)^\alpha,(p,q)^\beta\right]_{\ell_1\ell_2} = \sum_{\ell_3}\frac{2\ell_3+1}{16\pi}\left(1+(-1)^{\ell_1+\ell_2+\ell_3}\right)^2 \times \begin{pmatrix}\ell_1 & \ell_2 & \ell_3\\ -2 & 2 & 0\end{pmatrix}^2 W^{X,Y}\left[(i,j)^\alpha,(p,q)^\beta\right]_{\ell_3}, \tag{C.9}
$$

and

$$
\Xi_{TE}^{X,Y}\left[(i,j)^\alpha,(p,q)^\beta\right]_{\ell_1\ell_2} = \sum_{\ell_3}\frac{2\ell_3+1}{8\pi}\left(1+(-1)^{\ell_1+\ell_2+\ell_3}\right) \times \begin{pmatrix}\ell_1 & \ell_2 & \ell_3\\ 0 & 0 & 0\end{pmatrix}\begin{pmatrix}\ell_1 & \ell_2 & \ell_3\\ -2 & 2 & 0\end{pmatrix} W^{X,Y}\left[(i,j)^\alpha(p,q)^\beta\right]_{\ell_3}, \tag{C.10}
$$

where $X,Y \in \{\emptyset\emptyset, TT, PP\}$, and $\alpha,\beta \in \{TT, TP, PT, PP\}$. They make use of window functions W,

$$
W^{\emptyset\emptyset,\emptyset\emptyset}\left[(i,j)^\alpha,(p,q)^\beta\right]_\ell = \frac{1}{2\ell+1}\sum_m w_{\ell m}^{\emptyset\emptyset}(i,j)^\alpha w_{\ell m}^{*\,\emptyset\emptyset}(p,q)^\beta, \tag{C.11}
$$

$$
W^{\emptyset\emptyset,TT}\left[(i,j)^\alpha,(p,q)^{TT}\right]_\ell = \frac{1}{2\ell+1}\sum_m w_{\ell m}^{\emptyset\emptyset}(i,j)^\alpha w_{\ell m}^{*\,II}(p,q)^{TT}, \tag{C.12}
$$

$$W^{\emptyset\emptyset,PP}\Big[(i,j)^{\alpha},(p,q)^{PP}\Big]_{\ell} = \frac{1}{2\ell+1}\sum_m$$
$$\frac{1}{2}\left(w^{\emptyset\emptyset}_{\ell m}(i,j)^{\alpha}w^{*\,QQ}_{\ell m}(p,q)^{PP} + w^{\emptyset\emptyset}_{\ell m}(i,j)^{\alpha}w^{*\,UU}_{\ell m}(p,q)^{PP}\right), \quad (C.13)$$
$$W^{TT,TT}\Big[(i,j)^{TT},(p,q)^{TT}\Big]_{\ell} = \frac{1}{2\ell+1}\sum_m w^{II}_{\ell m}(i,j)^{TT}w^{*\,II}_{\ell m}(p,q)^{TT}, \quad (C.14)$$
$$W^{TT,PP}\Big[(i,j)^{TT},(p,q)^{PP}\Big]_{\ell} = \frac{1}{2\ell+1}\sum_m$$
$$\frac{1}{2}\left(w^{II}_{\ell m}(i,j)^{TT}w^{*\,QQ}_{\ell m}(p,q)^{PP} + w^{II}_{\ell m}(i,j)^{TT}w^{*\,UU}_{\ell m}(p,q)^{PP}\right), \quad (C.15)$$

and

$$W^{PP,PP}\Big[(i,j)^{PP},(p,q)^{PP}\Big]_{\ell} = \frac{1}{2\ell+1}\sum_m$$
$$\frac{1}{4}\Big(w^{QQ}_{\ell m}(i,j)^{PP}w^{*\,QQ}_{\ell m}(p,q)^{PP} + w^{UU}_{\ell m}(i,j)^{PP}w^{*\,UU}_{\ell m}(p,q)^{PP}$$
$$+w^{QQ}_{\ell m}(i,j)^{PP}w^{*\,UU}_{\ell m}(p,q)^{PP} + w^{UU}_{\ell m}(i,j)^{PP}w^{*\,QQ}_{\ell m}(p,q)^{PP}\Big). \quad (C.16)$$

In the above expressions, we defined the spherical harmonic coefficients of the effective weight maps $w^{\emptyset\emptyset}$,

$$w^{\emptyset\emptyset}_{\ell m}(i,j)^{TT} = \sum_{p=1}^{N_{\rm pix}} m^{i,T}_p m^{j,T}_p Y^*_{\ell m}(\hat{\boldsymbol{n}}_p)\Omega_p, \quad (C.17)$$
$$w^{\emptyset\emptyset}_{\ell m}(i,j)^{PP} = \sum_{p=1}^{N_{\rm pix}} m^{i,P}_p m^{j,P}_p Y^*_{\ell m}(\hat{\boldsymbol{n}}_p)\Omega_p, \quad (C.18)$$
$$w^{\emptyset\emptyset}_{\ell m}(i,j)^{TP} = \sum_{p=1}^{N_{\rm pix}} m^{i,T}_p m^{j,P}_p Y^*_{\ell m}(\hat{\boldsymbol{n}}_p)\Omega_p, \quad (C.19)$$

and

$$w^{\emptyset\emptyset}_{\ell m}(i,j)^{PT} = \sum_{p=1}^{N_{\rm pix}} m^{i,P}_p m^{j,T}_p Y^*_{\ell m}(\hat{\boldsymbol{n}}_p)\Omega_p, \quad (C.20)$$

where m^T is the temperature mask (Stokes I), m^P the polarization mask (Stokes Q and U), and Ω_p the solid angle of pixel p.

Accordingly, the noise-variance-weighted maps w^{II}, w^{QQ}, and w^{UU} are

$$w^{II}_{\ell m}(i,j)^{TT} = \delta_{i,j}\sum_{p=1}^{N_{\rm pix}}\left(\sigma^{II}_p\right)^2 m^{i,T}_p m^{j,T}_p Y^*_{\ell m}(\hat{\boldsymbol{n}}_p)\Omega^2_p, \quad (C.21)$$
$$w^{QQ}_{\ell m}(i,j)^{PP} = \delta_{i,j}\sum_{p=1}^{N_{\rm pix}}\left(\sigma^{QQ}_p\right)^2 m^{i,P}_p m^{j,P}_p Y^*_{\ell m}(\hat{\boldsymbol{n}}_p)\Omega^2_p, \quad (C.22)$$

and

$$w^{UU}_{\ell m}(i,j)^{PP} = \delta_{i,j}\sum_{p=1}^{N_{\rm pix}}\left(\sigma^{UU}_p\right)^2 m^{i,P}_p m^{j,P}_p Y^*_{\ell m}(\hat{\boldsymbol{n}}_p)\Omega^2_p, \quad (C.23)$$

where the σ^2_ps are the noise variances in pixel p in the given Stokes map, and the Kronecker symbols $\delta_{i,j}$ ensure that there is only a noise contribution if the two detectors i and j are identical.

In the spherical harmonic representation of the noise-variance-weighted window functions that appear in Eqs. (C.8)–(C.10) it is possible to take into account noise correlations approximately. Following Like13 and given the characterization of the observed noise power spectra discussed in Sect. 3.4.4, we multiply the projector functions $\Xi^{X,Y}_{\ell_1\ell_2}$ for each factor of $X, Y \in \{TT, PP\}$ by an additional rescaling coefficient,

$$r_{\ell_1\ell_2} = \sqrt{\frac{N^{\rm data}_{\ell_1}N^{\rm data}_{\ell_2}}{N^{\rm white}_{\ell_1}N^{\rm white}_{\ell_2}}}. \quad (C.24)$$

Here, $N^{\rm data}_\ell/N^{\rm white}_\ell$ is the ratio of the observed noise power spectrum to the white-noise power spectrum predicted by the pixel noise variance values σ^2_p.

C.1.2. Mask deconvolution

In a last step, we correct the individual covariance matrix blocks for the effect of pixel window function, beam, and mask. Using the coupling matrices (Hivon et al. 2002; Kogut et al. 2003),

$$M^{TT}(i,j)_{\ell_1\ell_2} = (2\ell_2+1)\sum_{\ell_3}\frac{2\ell_3+1}{4\pi}\begin{pmatrix}\ell_1 & \ell_2 & \ell_3\\ 0 & 0 & 0\end{pmatrix}^2 V^{TT}(i,j)_{\ell_3}, \quad (C.25)$$
$$M^{EE}(i,j)_{\ell_1\ell_2} = (2\ell_2+1)\sum_{\ell_3}\frac{2\ell_3+1}{16\pi}\left(1+(-1)^{\ell_1+\ell_2+\ell_3}\right)^2 \times\begin{pmatrix}\ell_1 & \ell_2 & \ell_3\\ -2 & 2 & 0\end{pmatrix}^2 V^{PP}(i,j)_{\ell_3}, \quad (C.26)$$
$$M^{TE}(i,j)_{\ell_1\ell_2} = (2\ell_2+1)\sum_{\ell_3}\frac{2\ell_3+1}{8\pi}\left(1+(-1)^{\ell_1+\ell_2+\ell_3}\right) \times\begin{pmatrix}\ell_1 & \ell_2 & \ell_3\\ 0 & 0 & 0\end{pmatrix}\begin{pmatrix}\ell_1 & \ell_2 & \ell_3\\ -2 & 2 & 0\end{pmatrix} V^{TP}(i,j)_{\ell_3}, \quad (C.27)$$
$$M^{ET}(i,j)_{\ell_1\ell_2} = (2\ell_2+1)\sum_{\ell_3}\frac{2\ell_3+1}{8\pi}\left(1+(-1)^{\ell_1+\ell_2+\ell_3}\right) \times\begin{pmatrix}\ell_1 & \ell_2 & \ell_3\\ 0 & 0 & 0\end{pmatrix}\begin{pmatrix}\ell_1 & \ell_2 & \ell_3\\ -2 & 2 & 0\end{pmatrix} V^{PT}(i,j)_{\ell_3}, \quad (C.28)$$

where

$$V^{TT}(i,j)_\ell = \frac{1}{2\ell+1}\sum_m m^{i,T}_{\ell m}m^{*\,j,T}_{\ell m}, \quad (C.29)$$
$$V^{PP}(i,j)_\ell = \frac{1}{2\ell+1}\sum_m m^{i,P}_{\ell m}m^{*\,j,P}_{\ell m}, \quad (C.30)$$
$$V^{TP}(i,j)_\ell = \frac{1}{2\ell+1}\sum_m m^{i,T}_{\ell m}m^{*\,j,P}_{\ell m}, \quad (C.31)$$

and

$$V^{PT}(i,j)_\ell = \frac{1}{2\ell+1}\sum_m m^{i,P}_{\ell m}m^{*\,j,T}_{\ell m}, \quad (C.32)$$

we obtain the final result for the deconvolved covariance matrix,

$${\rm Var}(\hat{C}^{XY\ i,j}_\ell, \hat{C}^{ZW\ p,q}_{\ell'})^{\rm dec} = \Big[M^{XY}(i,j)^{-1}\ {\rm Var}(\hat{C}^{XY\ i,j}, \hat{C}^{ZW\ p,q}) \times\left(M^{ZW}(p,q)^{-1}\right)^{\dagger}\Big]_{\ell\ell'}\Big/\left(b^{X\ i}_\ell\, b^{Y\ j}_\ell\, b^{Z\ p}_{\ell'}\, b^{W\ q}_{\ell'}\, p^{2\ XY}_\ell\, p^{2\ ZW}_{\ell'}\right). \quad (C.33)$$

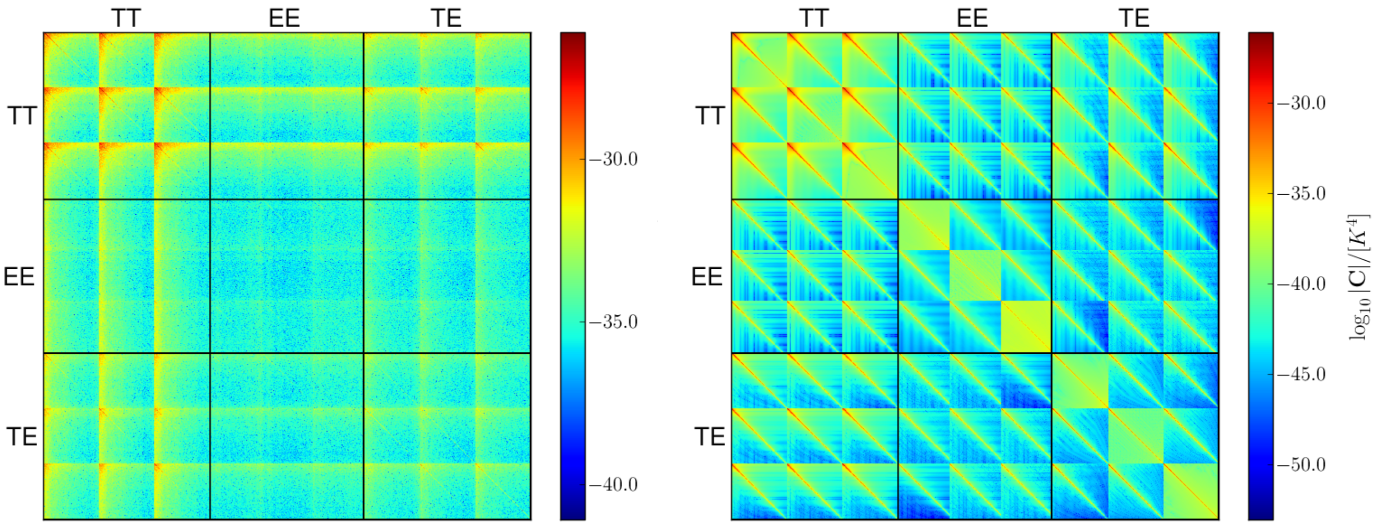

Fig. C.1. Combined C_ℓ-covariance matrices comprising the $TTTT$ (*upper left sub-block*), $EEEE$ (*middle sub-block*), and $TETE$ (*lower right sub-block*) covariances and their cross-correlations. *Left*: empirical covariance. *Right*: analytic covariance. We note the different scales; despite visual appearance, the diagonals are in good agreement.

C.1.3. Validation of the implementation

We verified the numerical implementation of the pipeline used to compute covariance matrices by means of Monte Carlo simulations. Specifically, we generated a set of 10 000 simulated maps for the four HFI detector sets 143-ds1, 143-ds2, 217-ds1, and 217-ds2. The simulations included CMB and an isotropic frequency-dependent foreground component, convolved with effective beam and pixel window functions. To each map, we added a realization of anisotropic, correlated noise.

In this test, we used a Galactic mask that leaves 40% of the sky for analysis at both frequencies and neglected the point source mask usually applied to temperature data. We then computed a total of 120 000 cross power spectra and constructed empirical covariance matrices for the 21 unique detector combinations that can be built from the four channels. Being based on at least 10 000 simulations each, the covariance matrix estimates reach an intrinsic relative precision of 1% or better.

We then compared the empirical covariance matrix to its approximate analytic counterpart computed using identical input parameters. To do so, we applied the standard post-processing procedure discussed in Sect. 3.5 to produce frequency averaged covariance matrices for all frequency combinations at 143 and 217 GHz. For the analysis, we adopted frequency-independent multipole ranges $100 \leq \ell \leq 2500$ for TT and TE, and $100 \leq \ell \leq 2000$ for EE. In a final step, we reduced the size of the matrices by binning. The temperature and polarization blocks were then combined into the single matrices shown in Fig. C.1. We note that, owing to the Monte Carlo noise floor, the colour scales are different, which may be misleading, since the diagonals appear to be fairly different, which is actually not the case. Indeed, Fig. C.2 compares the diagonal elements of the covariance matrix, and shows that for all polarization components and over the full multipole range, there is good agreement between the two covariance matrices, verifying the implementation of the equations summarized in the previous section, and their accuracy.

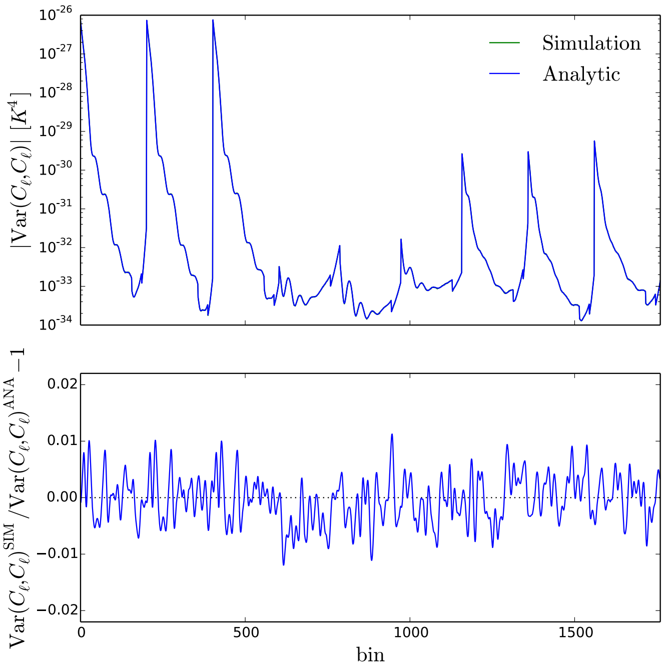

Fig. C.2. *Top*: diagonal elements of the empirical (green line) and analytic (blue line) covariance matrices; the two lines are indistinguishable. *Bottom*: ratio of the two estimates: the ratios differ from unity by <1% over the full multipole range for all frequency combinations and polarization blocks.

C.1.4. Excess variance induced by the point-source mask

The approximations used in the calculation of the covariance matrix assume that the power spectra of the masks decline rapidly, and therefore require a conservative apodization scheme at the expense of a reduction in the sky fraction available for analysis. The point-source masks used in the temperature analysis excise large numbers of sources with an approximately isotropic distribution. Owing to their high number, only a severely reduced apodization of individual holes is feasible in practice (cf. Sect. 3.2.2). As a consequence, the power spectrum of the combined Galactic and point-source mask flattens and the precision of the approximation deteriorates noticeably, leading to systematic errors in the calculated analytical covariance matrices.

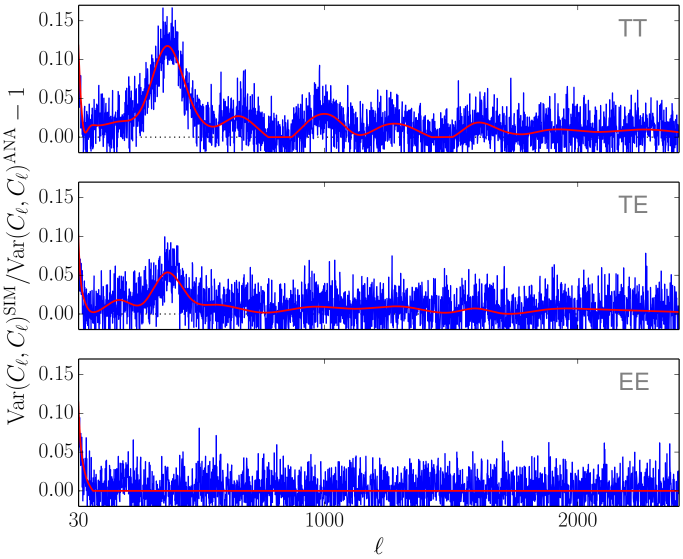

Fig. C.3. Excess variance induced by the temperature point-source mask. The graphs compare the diagonal elements of the empirical and analytical power spectrum covariance matrices (blue lines) for TT (*upper panel*), TE (*middle panel*), and EE (*lower panel*), and show deviations at the 10% level. The red lines are smooth fits based on cubic splines.

Here, we propose a heuristic approach to capture the variance modulations introduced by the point-source masks. In a first step, we use Monte Carlo simulations to quantify the level of mismatch between analytical and empirical power spectra variances. Since the point-source mask is frequency dependent, we simulate 5000 realizations of the six half-mission CMB and foreground maps, without noise contribution, at 100, 143, and 217 GHz. Using the reference Galactic and point-source masks in temperature, and Galactic masks in polarization (Sect. 3.2.2), we compute power spectra and construct empirical covariance matrices.

A comparison with the analytic covariance matrices reveals that the point-source mask has introduced excess variance that is not fully captured by the analytical approximation. In Fig. C.3 we plot results for the 217×217 GHz power spectrum variance, finding a deviation of up to about 10% at $\ell \approx 400$, with characteristic oscillating features in the TT and, to a lesser extent, in the TE power spectrum variance. Furthermore, on large scales ($\ell \lesssim 50$), the approximations start to break down in both temperature and polarization, a known feature of pseudo-power spectrum estimators (e.g., Efstathiou 2004).

In the signal-dominated regime, the analytical approximations of the covariance matrices are proportional to the square of the fiducial power spectrum C_ℓ (Eqs. (C.2)–(C.7)). Using spline fits to the variance ratios, we obtain correction factors that describe the excess scatter introduced by the point-source masks. We then multiply the fiducial power spectrum by the square-root of this ratio, cancelling the observed mismatch in the variance to first order.

C.2. Plik joint likelihood simulations

In Sect. 3.6 we discussed the 300 simulations performed to validate the overall implementation and our approximations for PlikTT. Here we complement that section with additional results for the full PlikTT, EE, TE joint likelihood.

Table C.1. Shifts of parameters for the joint PlikTT, EE, TE likelihood.

Parameter	300 sims
$\Omega_b h^2$	−1.09
$\Omega_c h^2$	0.62
θ	−0.25
τ	−0.88
$\ln(10^{10}A_s)$	−0.76
n_s	0.25
A^{217}_{CIB}	−0.75
gal^{100}_{545}	−0.03
gal^{143}_{545}	−0.05
$\mathrm{gal}^{143-217}_{545}$	−0.28
gal^{217}_{545}	1.38
gal^{100}_{EE}	0.69
$\mathrm{gal}^{100-143}_{EE}$	−0.80
$\mathrm{gal}^{100-217}_{EE}$	0.02
gal^{143}_{EE}	−0.07
$\mathrm{gal}^{143-217}_{EE}$	−1.21
gal^{217}_{EE}	1.08
gal^{100}_{TE}	−0.11
$\mathrm{gal}^{100-143}_{TE}$	−0.39
$\mathrm{gal}^{100-217}_{TE}$	0.32
gal^{143}_{TE}	0.55
$\mathrm{gal}^{143-217}_{TE}$	−0.47
gal^{217}_{TE}	1.20

Notes. The shifts are given in units of the posterior width rescaled by $300^{-1/2}$. If the parameters were uncorrelated, 68% of the shifts would be expected to lie within 1σ of their fiducial values. Of a total 23 parameters this would mean that 5 or 6 parameters are over 1σ away. As shown in the table, 3 parameters are in between 1 and 2σ, 2 parameters are marginally above 1σ and the remaining 18 parameters are well below 1σ.

Figure C.4 and Table C.1 show the full-likelihood parameter results; these are companions to Fig. 27 and Table 13 of Sect. 3.6, which were devoted to the TT case. The average reduced χ^2 corresponding to the histograms of Fig. C.4 is equal to 1.01. Compared to TT, the inclusion of EE and TE provides a significant improvement in the determination of several cosmological parameters, in particular n_s, θ, and τ. It also reduces the small bias in n_s already discussed in the main text, since the entire ℓ range is used in the joint analysis.

C.3. Plik validation and stability tests

This section complements the main text with detailed information on Plik results and tests on data, and how they are obtained. We start in Sect. C.3.1 with zooms in five adjacent ℓ-ranges of all the individual frequency cross-spectra, and their residuals with respect to the PlikTT+tauprior ΛCDM best-fit model, both in temperature and polarization. In order to facilitate the search for possible common features across frequency spectra, we compute inter-frequency power spectra differences, according to a

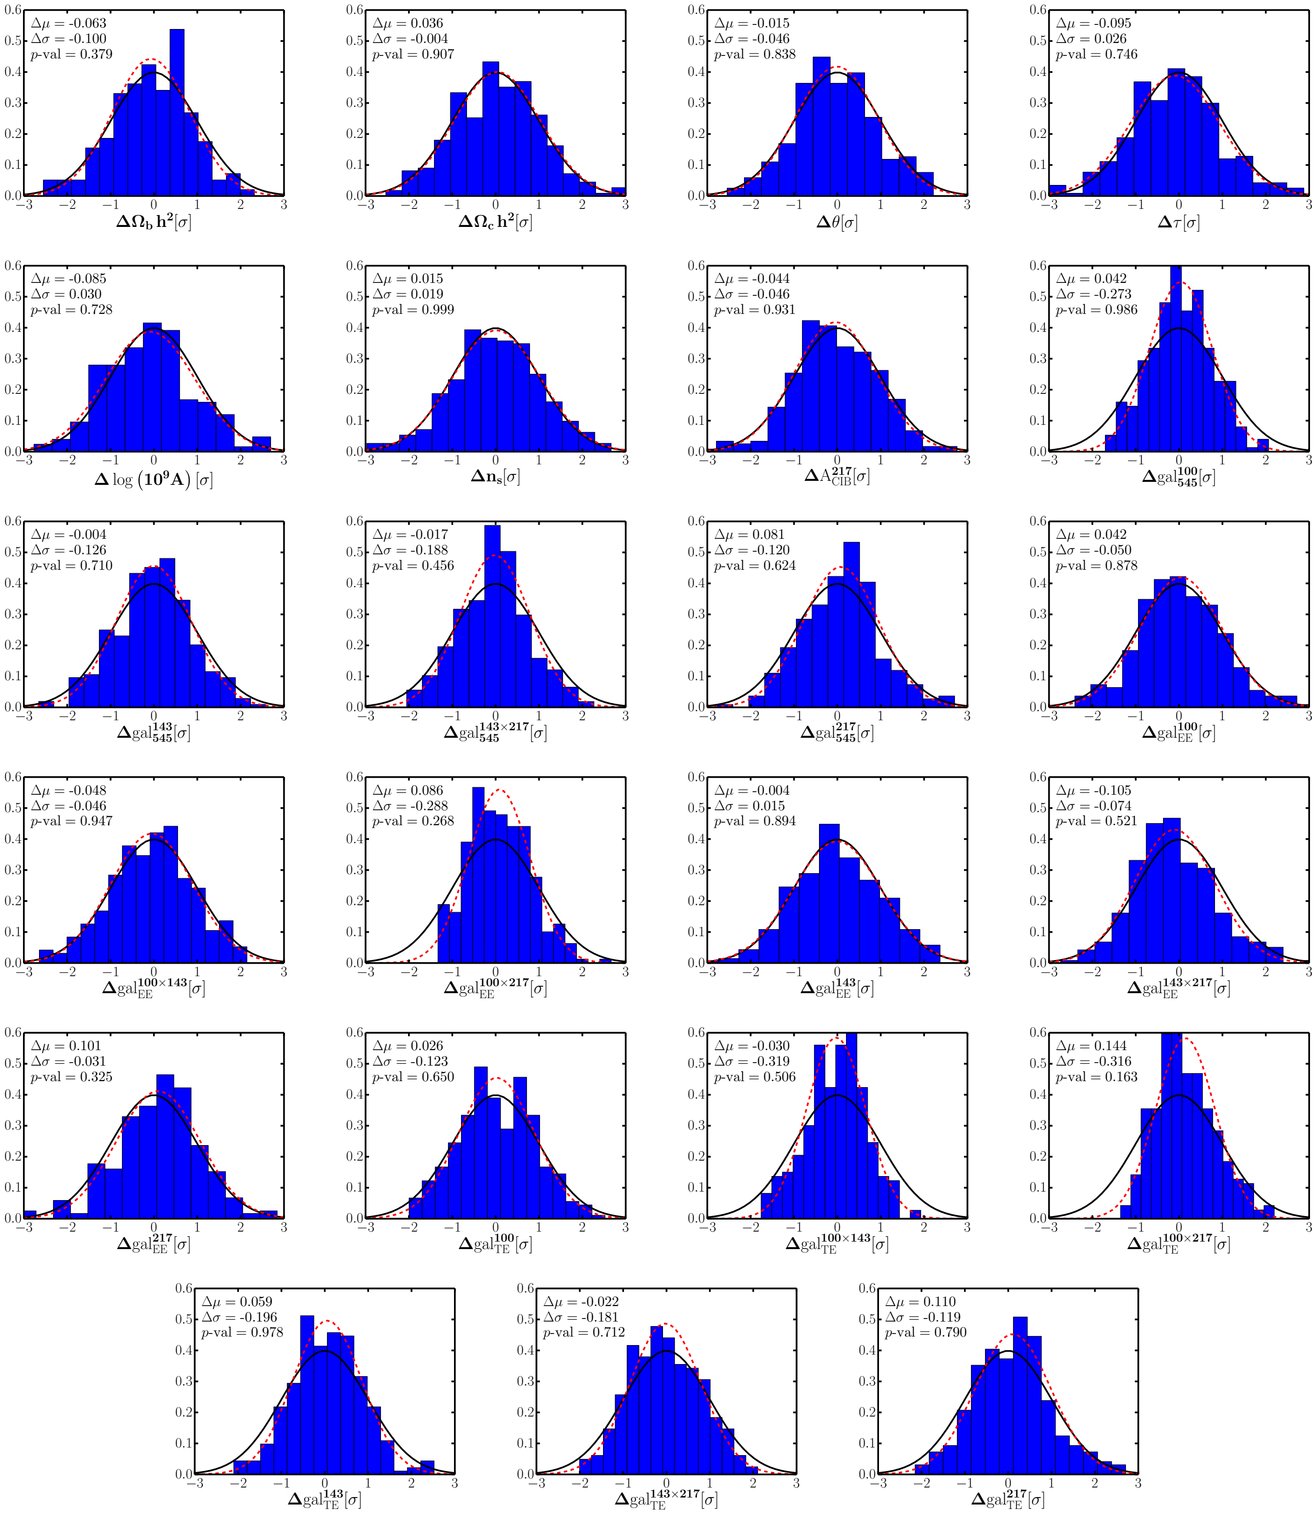

Fig. C.4. `Plik` parameter results from 300 simulations for the six baseline cosmological parameters, as well as the FFP8 CIB and Galactic dust amplitudes, as in Fig. 27, but for the joint `Plik`TT, EE, TE likelihood.

procedure discussed in Sect. C.3.2. Section 4.4.1 presents the corresponding results in polarization, which show that there are sizeable differences between pairs of foreground-cleaned spectra, much greater than those described in the main text for temperature. We proceed in Sect. C.3.5 to assess the robustness of the polarization results. Finally, we present in Sect. C.3.6 simulations to quantify whether the level of agreement between temperature- and polarization-based cosmological parameters is as expected.

C.3.1. Zoomed-in frequency power spectra and residuals

Figures C.5–C.7 show the frequency zoomed-in TT, EE and TE power spectra (respectively), in $\Delta\ell = 20$ bins. The red lines show

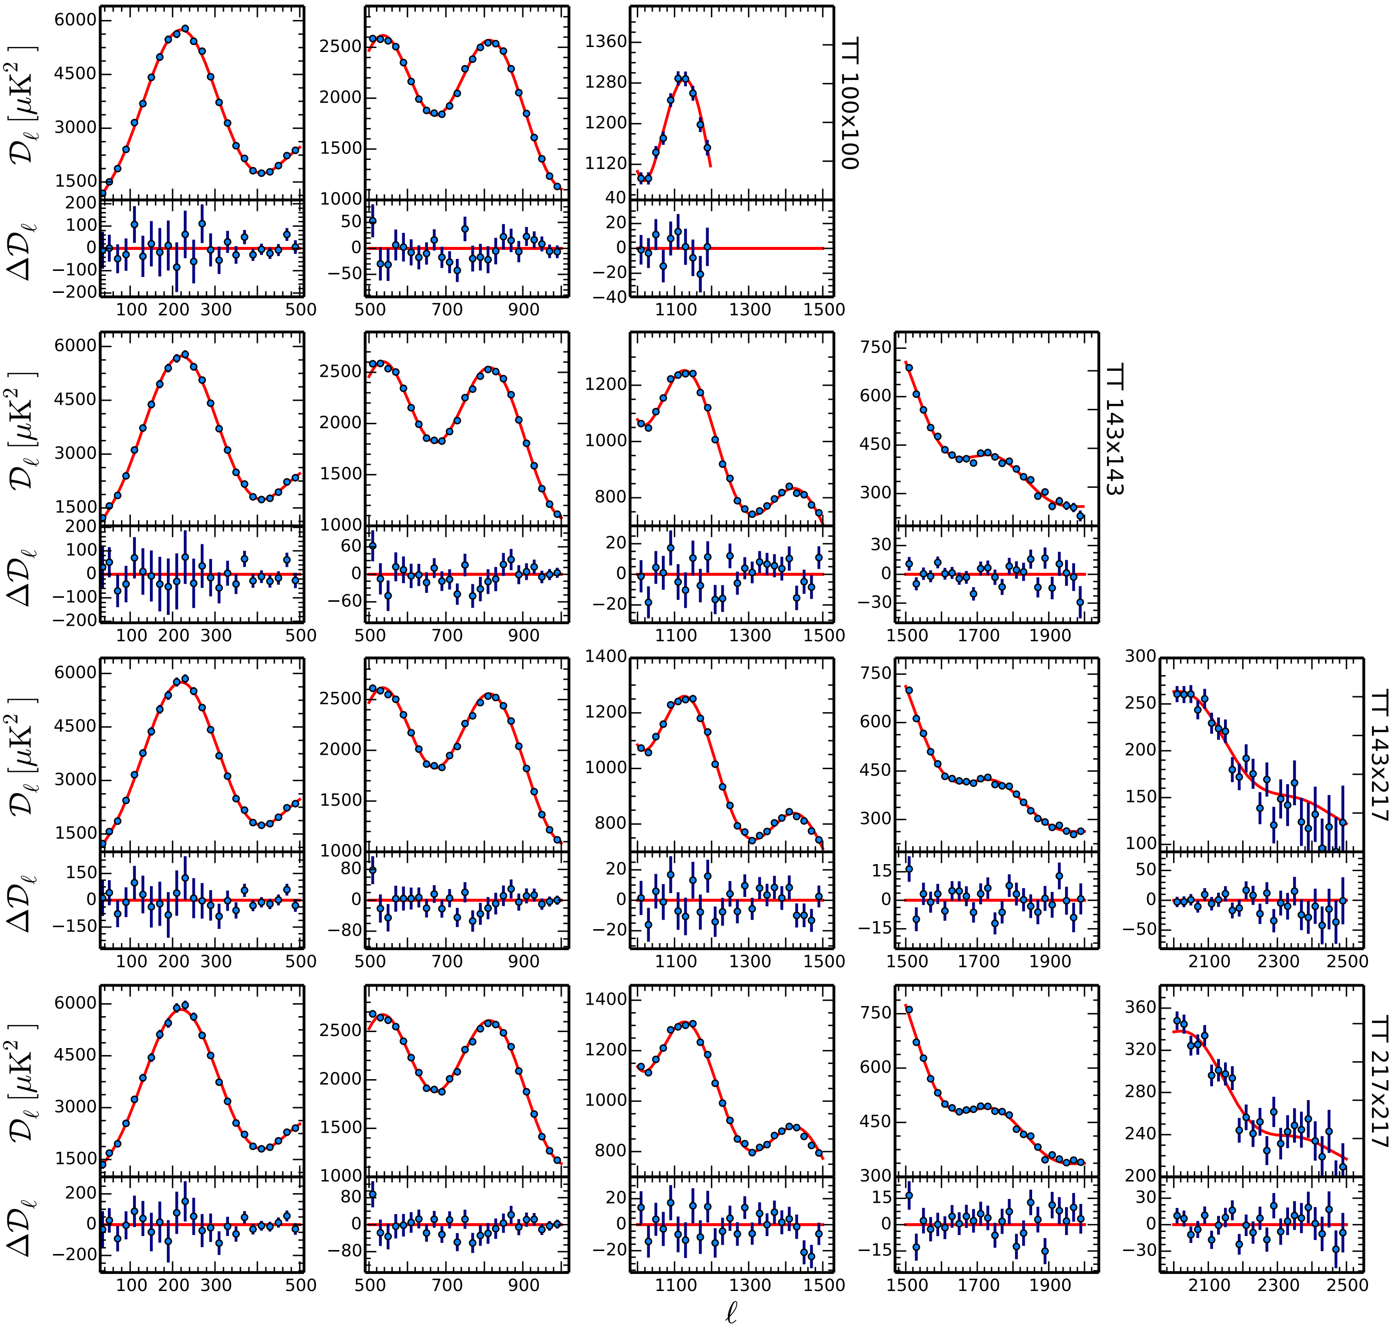

Fig. C.5. Per-frequency zoomed-in TT power spectra, in $\Delta\ell = 20$ bins. The red line shows the `PlikTT`+tauprior ΛCDM best-fit model. The *lower plots* show the residuals. We only show the ℓ ranges used in the baseline `Plik` likelihood.

the `PlikTT`+tauprior ΛCDM best-fit model. The lower plots show the residuals with respect to this best-fit model. We only show the multipole ranges that are included in the baseline analysis. These plots are meant to help the visual inspection of the residuals already shown in Fig. 32 and described in Sect. 3.8 for TT; and in Fig. 40 and described in Sect. 4.4 and Appendix 4.4.1 for TE and EE.

C.3.2. Inter-frequency power spectra differences

We describe here the procedure followed to obtain the inter-frequency power spectra differences shown in Figs. 31 and 41. We first clean the frequency power spectra by subtracting from the data the best-fit foreground solution obtained using the `PlikTT`+tauprior (for TT) or `PlikTT`, TE, EE+tauprior (for TE and EE) data combinations, assuming a ΛCDM framework.

We then calculate the difference between a pair of cleaned spectra of length n as $\Delta_\ell^{XY-X'Y'} = C_\ell^{XY} - C_\ell^{X'Y'}$.

The covariance matrix C^Δ of the difference $\Delta_\ell^{XY-X'Y'}$ is then:

$$\mathsf{C}^\Delta = A\mathsf{C}^{XY,X'Y'}A^T, \tag{C.34}$$

where $\mathsf{C}^{XY,X'Y'}$ is the $2n\times 2n$ covariance matrix relative to the XY and $X'Y'$ spectra, and A is a $n \times 2n$ matrix with blocks:

$$A = \begin{pmatrix} \mathbb{1}^{XY} & -\mathbb{1}^{X'Y'} \end{pmatrix}, \tag{C.35}$$

where $\mathbb{1}^{XY}$ is the $n \times n$ identity matrix.

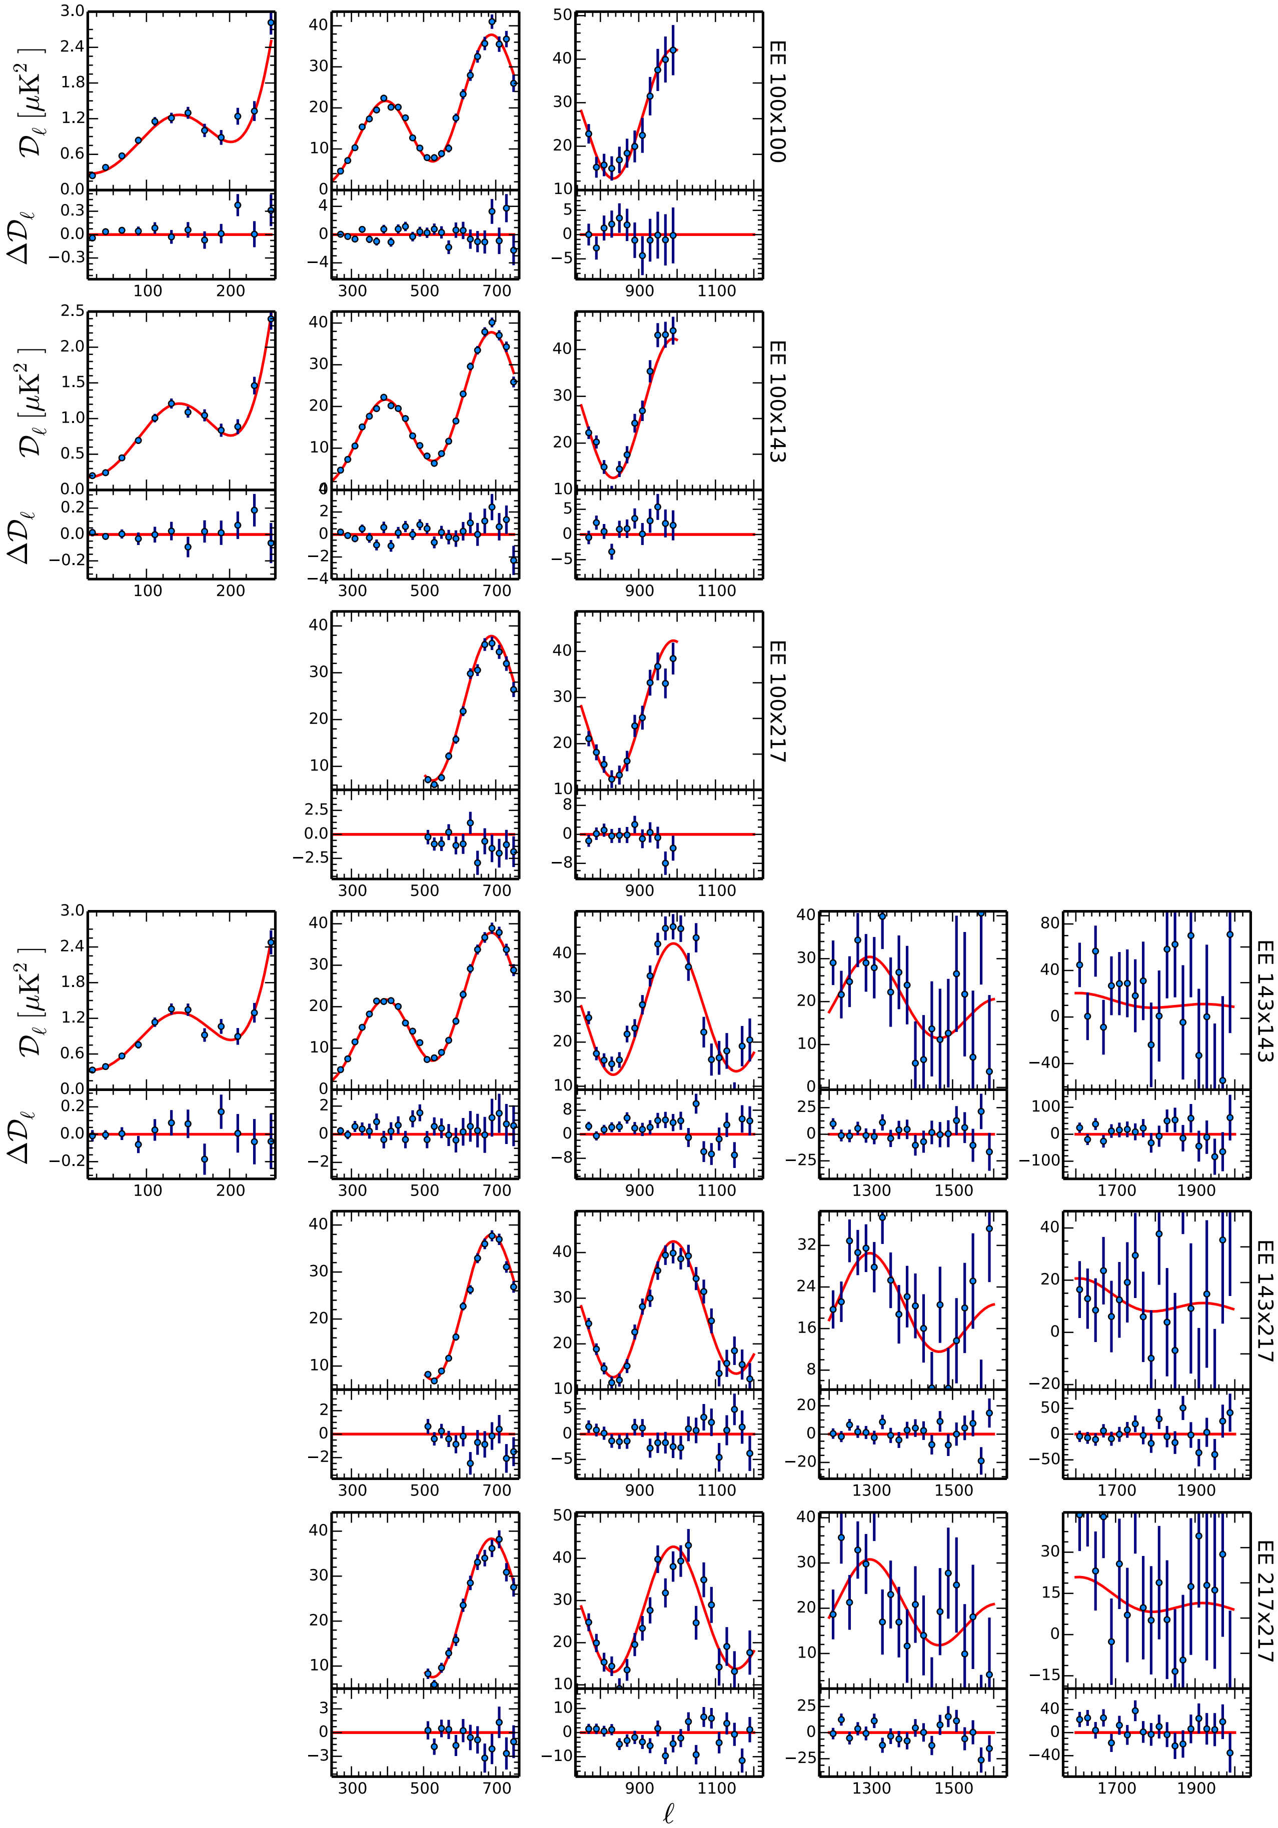

Fig. C.6. Same as Fig. C.5, but for *EE*.

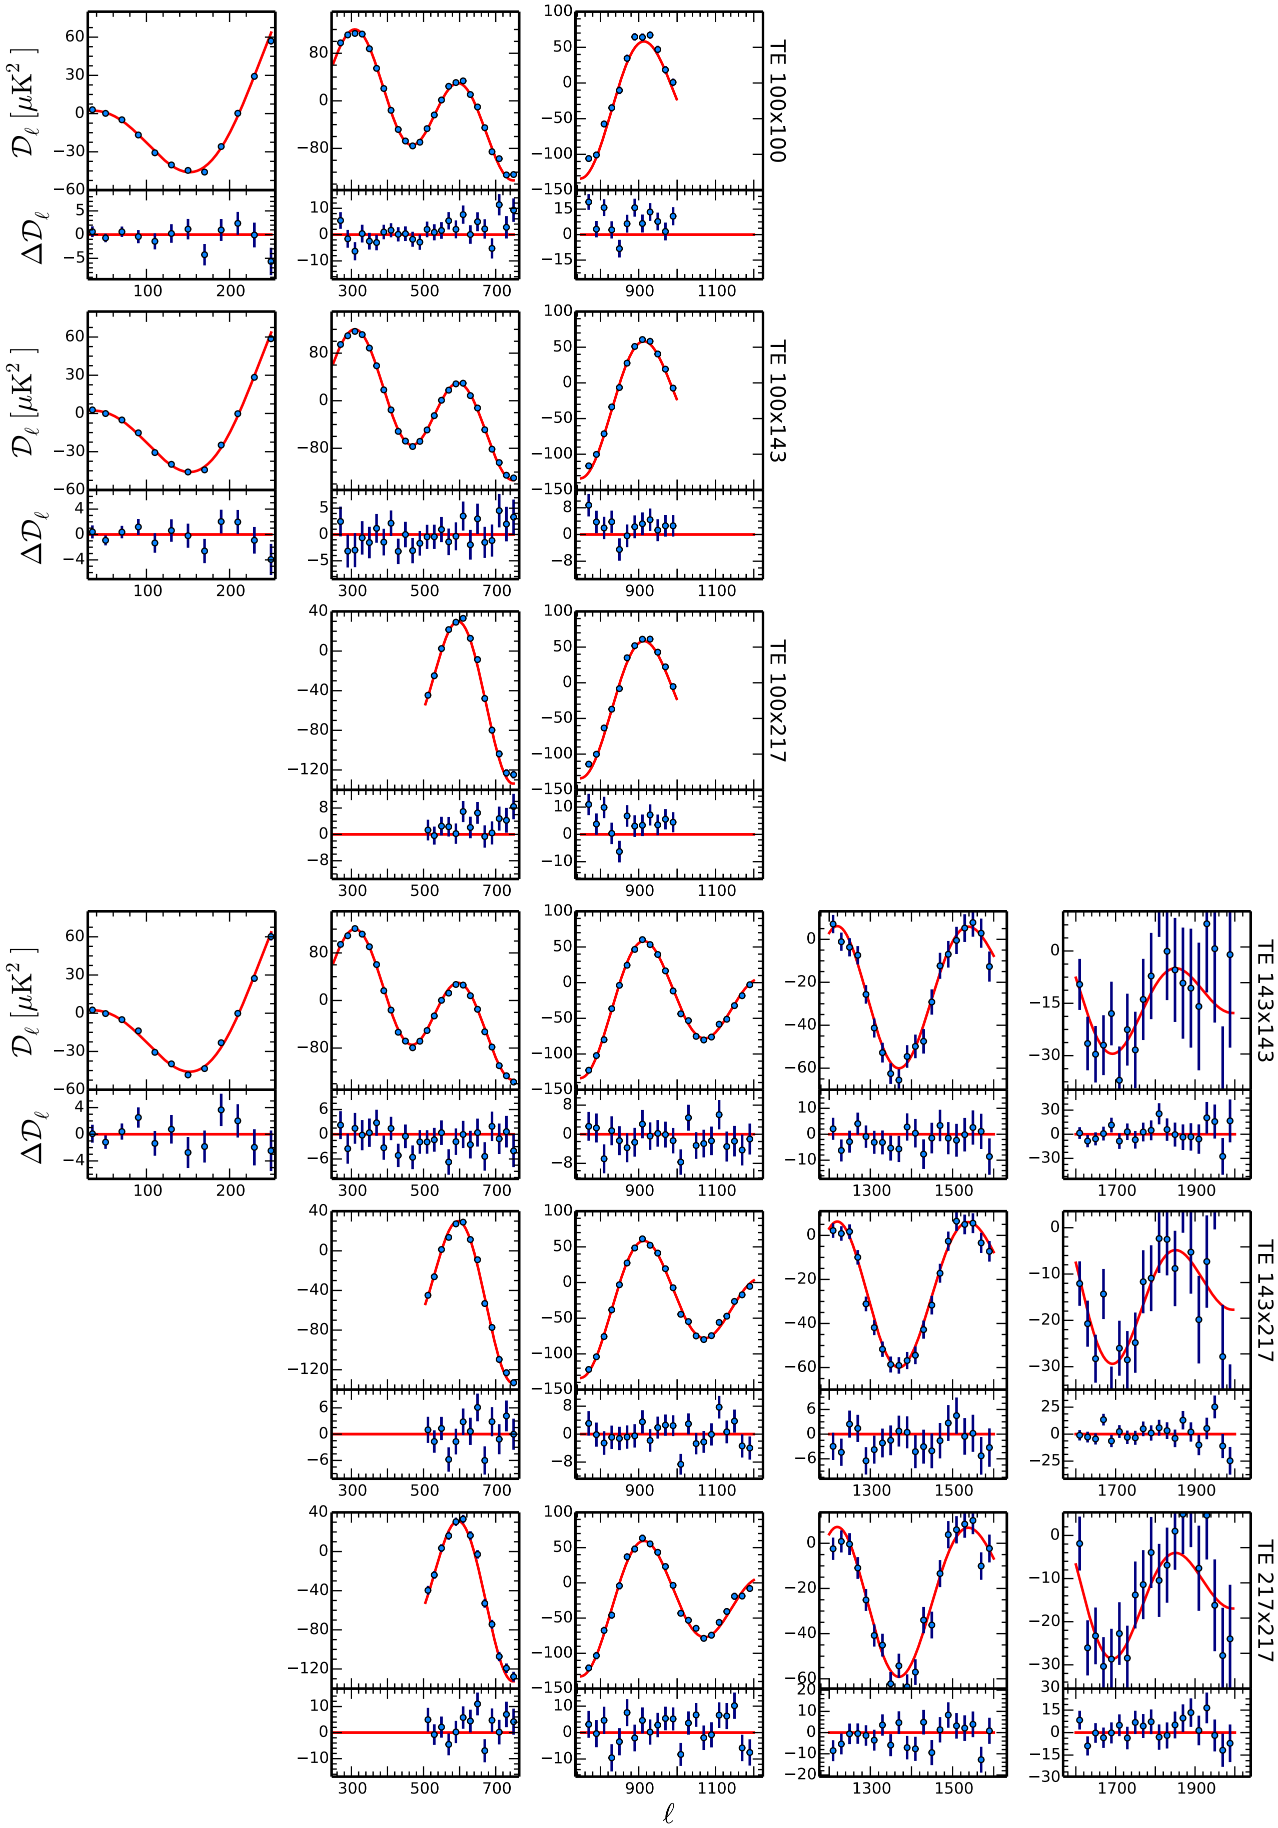

Fig. C.7. Same as Fig. C.5, but for *TE*.

C.3.3. Robustness tests on foreground parameters in *TT*

This section presents some further checks that we performed to validate the results from TT. Figure C.8 shows the marginal mean and the 68% confidence level error bars for the foreground parameters of the `Plik` TT high-ℓ likelihood under different assumptions about the data selection, foreground model, or treatment of the systematics. The cases considered are the same as those in Sect. 4.1, and the results for cosmological parameters can be found in Fig. 35. We now comment on them in turn.

Detset likelihood. In the detsets ("DS") case, the amplitude of the point sources at 100×100 GHz is higher than in the baseline case. This might indicate a residual correlated noise component in the DS spectra, not corrected by the procedure described in Sect. 3.4.4.

Impact of Galactic mask and dust modelling. We recover Galactic dust amplitudes within 1σ of the baseline values when we leave these parameters free to vary without any prior ("No gal priors") or when we leave the Galactic slope (described in Sect. 4.1.2) free to vary. The dust amplitudes for the "M605050" case (i.e., when we use more conservative Galactic masks, as detailed in Sect. 4.1.2) cannot be directly compared to the baseline values, since we expect smaller amplitudes when using reduced sky fractions.

Changes with $\ell_{\rm min}$. We observe variations by up to 1σ, as well as an increase in the error bars, in the level of dust contamination at 217×217 and of the CIB amplitude when we consider $\ell_{\rm min} = 50, 100$ instead of the baseline $\ell_{\rm min} = 30$, or when we excise the first 500 multipoles at 143×217 and 217×217. This is due to the fact that the lowest multipoles help in breaking the degeneracy between these two foreground components, giving tighter constraints when included in the analysis.

Changes with $\ell_{\rm max}$. We find that the overall amplitude of the foregrounds decreases when increasing the maximum multipole $\ell_{\rm max}$ included in the analysis[21]. This is related to the shift in cosmological parameters observed at different $\ell_{\rm max}$, which is described in Sect. 4.1.6. In Fig. C.8 the results for extragalactic foregrounds at $\ell_{\rm max} \lesssim 1200$ are not very meaningful, since these parameters are very weakly constrained in those multipole regions.

ΛCDM extensions. Figure C.8 also show the level of foregrounds obtained using the baseline likelihood in extensions of the ΛCDM model. In the ΛCDM+$N_{\rm eff}$ case, the level of foregrounds is very similar to that in the base-ΛCDM case, while in the ΛCDM+$A_{\rm L}$ model it is few $\mu{\rm K}^2$ lower at all frequencies.

`CamSpec`. The foreground contamination levels determined by the `CamSpec` and `Plik` codes differ by a few $\mu{\rm K}^2$. This appears in Fig. C.8 as differences at the 1σ level in the sub-dominant (and ill-determined) foreground components ($A^{\rm kSZ}$, $A^{\rm tSZ}_{143}$), together with different best-fit recalibration factors (c_{100}, c_{217}), a result of the different modelling choices made regarding the ℓ ranges retained, and small variations in the dust template (where it is least well determined by the data). As already mentioned earlier in the discussion of cosmological parameters, the strongest effect is in $n_{\rm s}$, resulting in our estimate of a 0.3σ systematic uncertainty on this parameter.

Other cases. The remaining cases shown in Fig. C.8 are described in Sect. 4.1. We find good agreement in the cases where we excise one frequency at a time, or when we use the `CAMB` code instead of `PICO`.

C.3.4. Further tests of the shift with $\ell_{\rm max}$

We have investigated whether different data combination choices have an impact on the shift in cosmological parameters we observe when we change the maximum multipole included in the analysis, as described in Sect. 4.1.6.

Figure C.9 shows the results for different $\ell_{\rm max}$ for three different settings. We show results for `PlikTT`+tauprior (red points), identical to the ones already shown in Fig. 35; for `PlikTT`+tauprior, but fixing the foregrounds to the best-fit of the baseline likelihood (yellow points); and for `PlikTT` combined with the low-ℓ likelihood in temperature and polarization (green points, `PlikTT`+lowTEB in the plot). This figure shows that in all these three cases we have similar behaviour for $\ln(10^{10}A_{\rm s})$, $\Omega_{\rm c}h^2$, and τ, i.e., they all increase with increasing $\ell_{\rm max}$. However, the evolution of the other parameters differs. While in the `PlikTT`+tauprior case the other parameters do not change significantly (apart from the shift in θ between $\ell_{\rm max} \approx 1200$–$1300$ already described in Sect. 4.1.6), fixing the foregrounds forces other parameters such as $n_{\rm s}$ and $\Omega_{\rm b}h^2$ to shift as well. It is interesting to note that all the parameters tend to converge to the baseline solution between $\ell_{\rm max} = 1404$ and 1505, confirming the impact of the fifth peak in determining the final solution, as already described in Sect. 4.1.6.

As far as the `PlikTT`+lowTEB combination is concerned, adding the low-ℓ multipoles in temperature pulls $n_{\rm s}$ to higher values in order to better fit the deficit at $\ell \sim 20$–30. This pull is more effective when excising the high-ℓ data (i.e., when using low $\ell_{\rm max}$), pushing $\Omega_{\rm c}h^2$ to even lower values, following the $n_{\rm s}$–$\Omega_{\rm c}h^2$ degeneracy.

C.3.5. Polarization robustness tests

We now present the results of the tests we conducted so far to assess the robustness and accuracy of the polarization results, with the same tools as used for TT (described in the main text). In the parameter domain, the results are summarized in Fig. C.10, which shows the marginal mean and the 68% confidence limit (CL) error bars for cosmological parameters using the `PlikTE` or `PlikEE` high-ℓ likelihoods under different assumptions about the data selection, foreground model, or treatment of the systematics. In the following, we comment, in turn, on each of the tests shown in this figure (from left to right). In most of the cases, we use the `Plik` likelihoods in combination with the usual Gaussian τ prior, $\tau = 0.07 \pm 0.02$. The reference `PlikTE`+tauprior and `PlikEE`+tauprior results for the ΛCDM model are denoted as "`PlikTE`+tauprior" and "`PlikEE`+tauprior". We note that all the TE tests are run with the `PICO` code, while the EE ones are run with the `CAMB` code, for the reasons given in Appendix C.5.

[21] We remind the reader that, in this test, at each frequency we always use $\ell_{\rm max}^{\rm freq} = \min(\ell_{\rm max}, \ell_{\rm max}^{\rm freq,\,baseline})$, with $\ell_{\rm max}^{\rm freq,\,baseline}$ the baseline $\ell_{\rm max}$ at each frequency as reported in Table 16 (e.g., in the $\ell_{\rm max} = 1404$ case, we still use the 100×100 power spectrum through $\ell = 1197$).

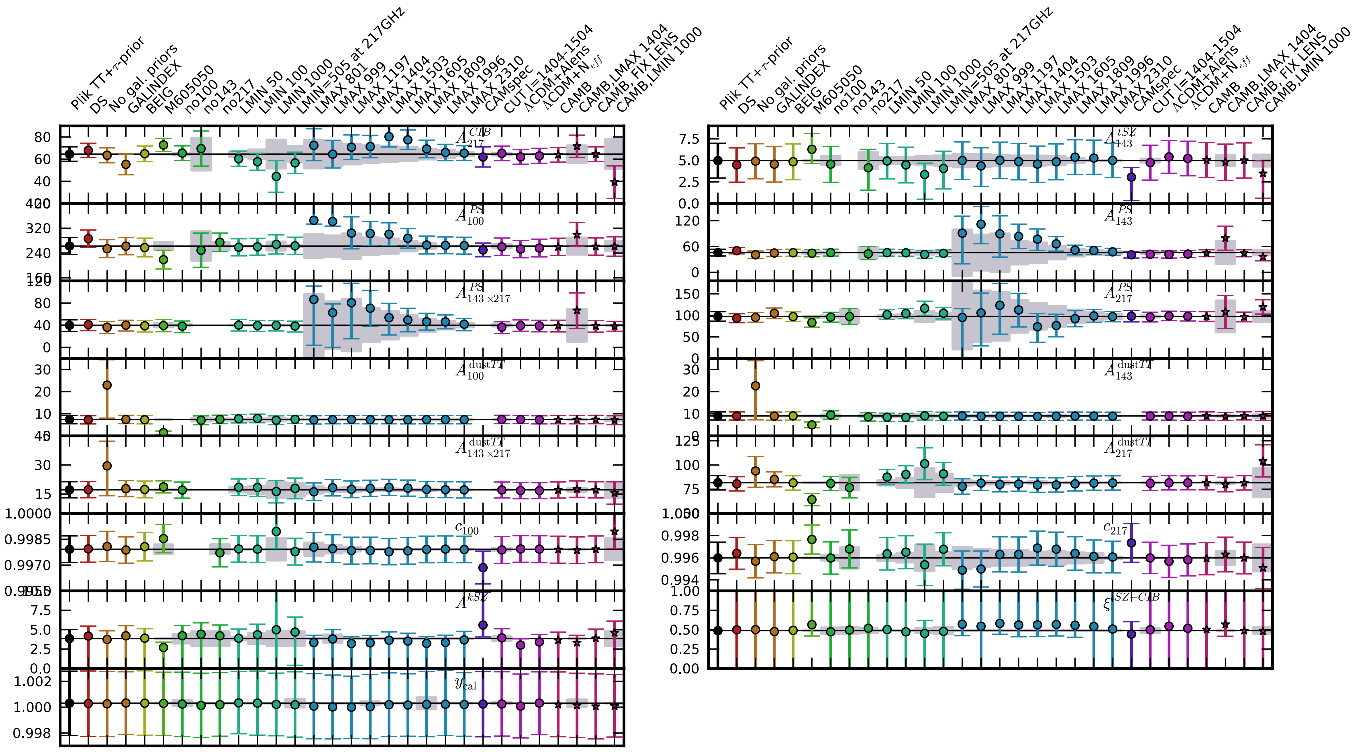

Fig. C.8. Marginal mean and 68% CL error bars on TT foreground parameters estimated when adopting different data choices for the `Plik` likelihood, in comparison with results from alternate approaches or models. We assume a ΛCDM model and always combine the `Plik` likelihood with a prior on $\tau = 0.07 \pm 0.02$ (we do not use low-ℓ temperature or polarization data here). "`PlikTT+tauprior`" indicates the baseline (HM, $\ell_{\rm min} = 30$, $\ell_{\rm max} = 2508$), while the other cases are described in Sect. 4.1. The grey bands show the standard deviation of the expected parameter shift, for those cases where the data used are a sub-sample of the baseline likelihood (see Eq. (53)).

Detsets. We find good agreement between the baseline cases based on half-mission spectra and those based on detset spectra (case "DS"). We find the greatest deviations in EE, where the DS case shows values of $\Omega_{\rm b}h^2$ and θ higher than the baseline case by about 1σ, while $n_{\rm s}$ is lower by 1σ.

Larger Galactic mask. We examined the impact of using a larger Galactic mask (case "M605050") with $f_{\rm sky} = 0.50$, 0.41, and 0.41 at 100, 143, and 217 GHz, respectively (corresponding to $f_{\rm sky}^{\rm noap} = 0.60$, 0.50, and 0.50 before apodization), instead of the baseline values $f_{\rm sky} = 0.70$, 0.50, and 0.41. In TE we observe substantial shifts in the parameters, at the level of $\lesssim 1\sigma$. We did not assess whether this is consistent with cosmic variance, but we note that the results remain compatible with "`PlikTT+tauprior`" at the 1σ level.

Galactic dust priors. We find that leaving the Galactic dust amplitudes completely free to vary ("No Gal. priors"), without applying the priors described in Sect. 3.3.1, does not have a significant impact on cosmological parameters. This suggest that our foreground model is satisfactory, despite its simplicity[22].

[22] We discovered late in the preparation of this paper that in some of the tests the prior for the 143 × 217 TE dust contamination was set inaccurately, with an offset of $-0.3\,\mu{\rm K}^2$ at $\ell = 500$. With our cuts, this spectrum contributes only at $\ell > 500$ where the dust contamination is already small compared to the signal. We verified that this has no impact on the cosmology and on our conclusions.

Beam eigenmodes. We have marginalized over the beam uncertainty eigenmodes (case "BEIG"), finding, as in TT, no impact on cosmological parameters.

Beam leakage. Section 3.4.3 presented a model for the polarization systematic error induced by assuming identical beams in detsets combined at the map-making stage (when the beams do in fact differ). Here we consider three cases for exploring the impact of the 18 amplitudes of the beam leakage model parameters, ε_m (for $m = 0$, 2, and 4; i.e., three parameters per cross-frequency spectrum): when we leave these amplitudes completely free to vary along with all other parameters (case "BLEAK"); when we apply the priors motivated in Sect. 3.4.3 (case "priors_BLEAK"); and when we use the best-fit values of these parameters ("FIX_BLEAK"). The amplitudes for "FIX_BLEAK" are obtained by a prior exploration while keeping all other parameters (TT cosmology and foregrounds) fixed. We find that this case has better goodness of fit without otherwise affecting the model.

When we leave the amplitudes completely free to vary, there is no significant impact on cosmology in TE, with shifts at the level of fractions ofσ, which is reassuring. For EE, though, we find large deviations in the "BLEAK" case, suggesting strong degeneracies between the cosmological and beam leakage parameters in EE. And for both TE and EE, we find that the beam leakage parameters adopt values in the "BLEAK" case that are much higher than the values expected from the priors. This shows that other residual systematic effects project substantially onto these template shapes, which is not surprising, given the additional degrees of freedom.

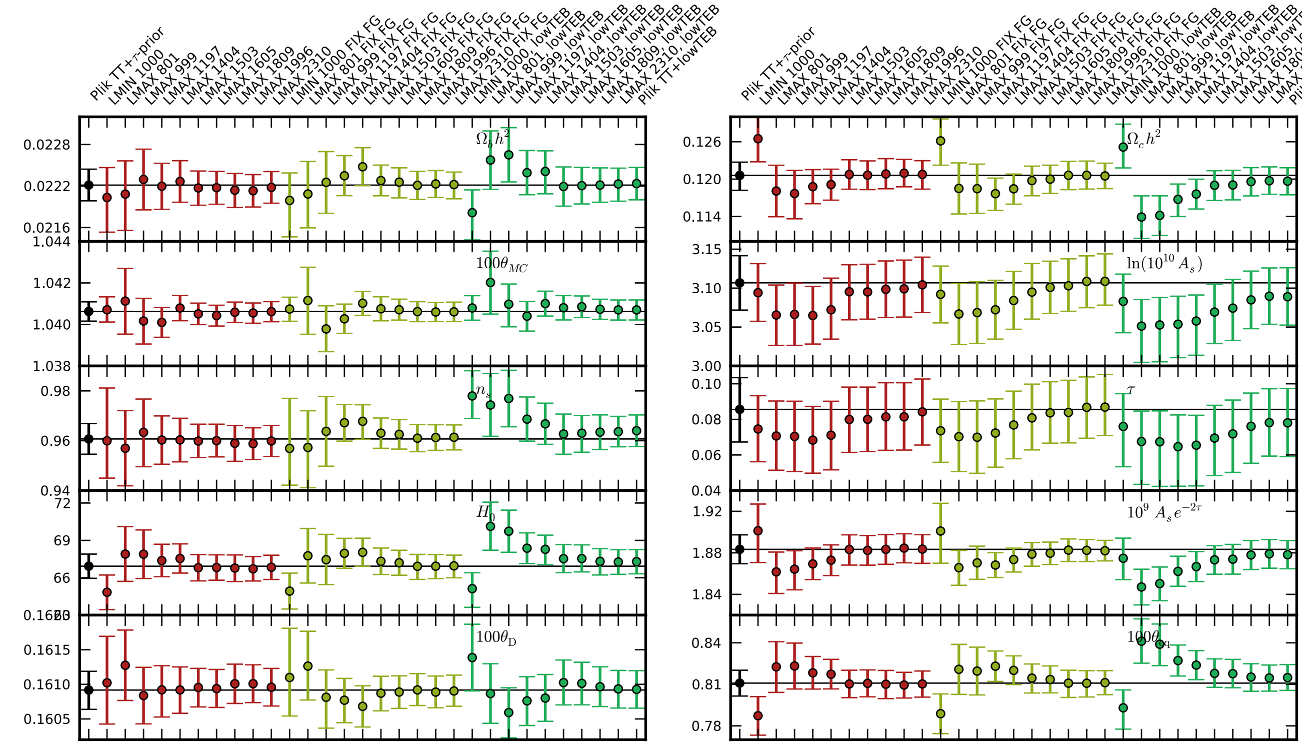

Fig. C.9. Marginal mean and 68% CL error bars on cosmological parameters estimated when adopting different data choices for the `Plik` likelihood. We assume a ΛCDM model and calculate parameters using different maximum multipole $\ell_{\rm max}$. The red points show the results for `PlikTT`+tauprior, with the points specifically labelled "`PlikTT`+tauprior" in black showing the baseline `PlikTT` likelihood at $\ell_{\rm max} = 2508$, the yellow points show results for `PlikTT`+tauprior but fixing the foregrounds to the best-fit of the baseline likelihood ("FIX FG"), and the green points show results for `PlikTT` combined with the low-ℓ likelihood in temperature and polarization ("`PlikTT`+lowTEB").

If we use our so-called cosmological prior (case "FIX_BLEAK"), i.e., when we fix leakage parameters to their best-fit values, in order to see how they improve the overall goodness of fit, the uncertainties remain close to the reference case (when the ε_m are set to zero) and of course the results shift slightly towards the "`PlikTT+tauprior`" result. By using this TT solution, the fit improves by $\Delta\chi^2 = 55$ in TE, and only $\Delta\chi^2 = 26$ in EE, while opening 18 new parameters (and TT has 765 bins, while TE and EE have 762 bins). For TE in particular, the corrections are not sufficient to significantly improve the χ^2, which is too large, and dominated by the disagreement between individual spectra. Furthermore, the beam-leakage parameter values that we recover are higher than what we expect from the physical priors.

If instead we apply the physical priors, the best-fit cosmological values are not strongly affected, except for a small shift towards the "`PlikTT+tauprior`" case, and the errors bars are increased substantially compared to the fixed-leakage-parameter cases. But we find that the χ^2 value of the fit does not improve significantly (i.e., barely any change in EE, and $\Delta\chi^2 \approx 20$ in TE). The discrepancy between frequencies remains. We also explored the simultaneous variation of the leakage and calibration parameters within their expected physical priors, and found results similar to the case of the variation of the leakage alone.

In any case, we cannot assign the origin of the frequency-spectra disagreement to beam leakage, alone or in combination with polarization recalibration. The surprisingly high values found for the leakage parameters when they are allowed to vary widely are indicative of the presence of other systematic effects that are absent from our model. We therefore do not include these corrections in the final baseline likelihood; we only use them to estimate the possible amount of residual beam leakage in the co-added spectra, which is around $1\,\mu{\rm K}^2$ ($\mathcal{D}_\ell$) in TE and $1 \times 10^{-5} \mu{\rm K}^2$ (C_ℓ) in EE.

Cutting out frequency channels. We have considered the cases where we eliminate all the power-spectra related to one particular frequency at a time, as in the TT analyses; e.g., the "no 100" case uses only the 143×143, 143×217, and 217×217 spectra. In TE, we see strong shifts (in opposite directions) when either the 100 or the 143 GHz data are removed, much more than one would expect due to the change of information (given by the grey bands in Fig. C.10). In EE, we instead see strong shifts in opposite directions when either the 143 or the 217 GHz data are dropped. Furthermore, we note in EE the rather big and similar change in EE parameters when the 143 GHz data are dropped and when the leakage parameters are varied.

Changing $\ell_{\rm min}$. We find good stability in the results when changing the minimum multipole $\ell_{\rm min}$ considered in the analysis ("LMIN" case). The baseline likelihood has $\ell_{\rm min} = 30$, and we test the cases of $\ell_{\rm min} = 50$ and 100.

Changing $\ell_{\rm max}$. We observe small shifts when including maximum multipoles between $\ell_{\rm max} \sim 1000$ and 2000 ("LMAX" cases). This is not surprising, since even though the baseline has $\ell_{\rm max} = 2000$, most of the constraining power of our polarization spectra comes from $\ell < 1000$. When using $\ell_{\rm max} = 801$, we find bigger shifts, non-Gaussian parameter posterior distributions (for EE), and a significant increase in the error bars. This increase is expected from Fisher-matrix forecasts (see, e.g.,

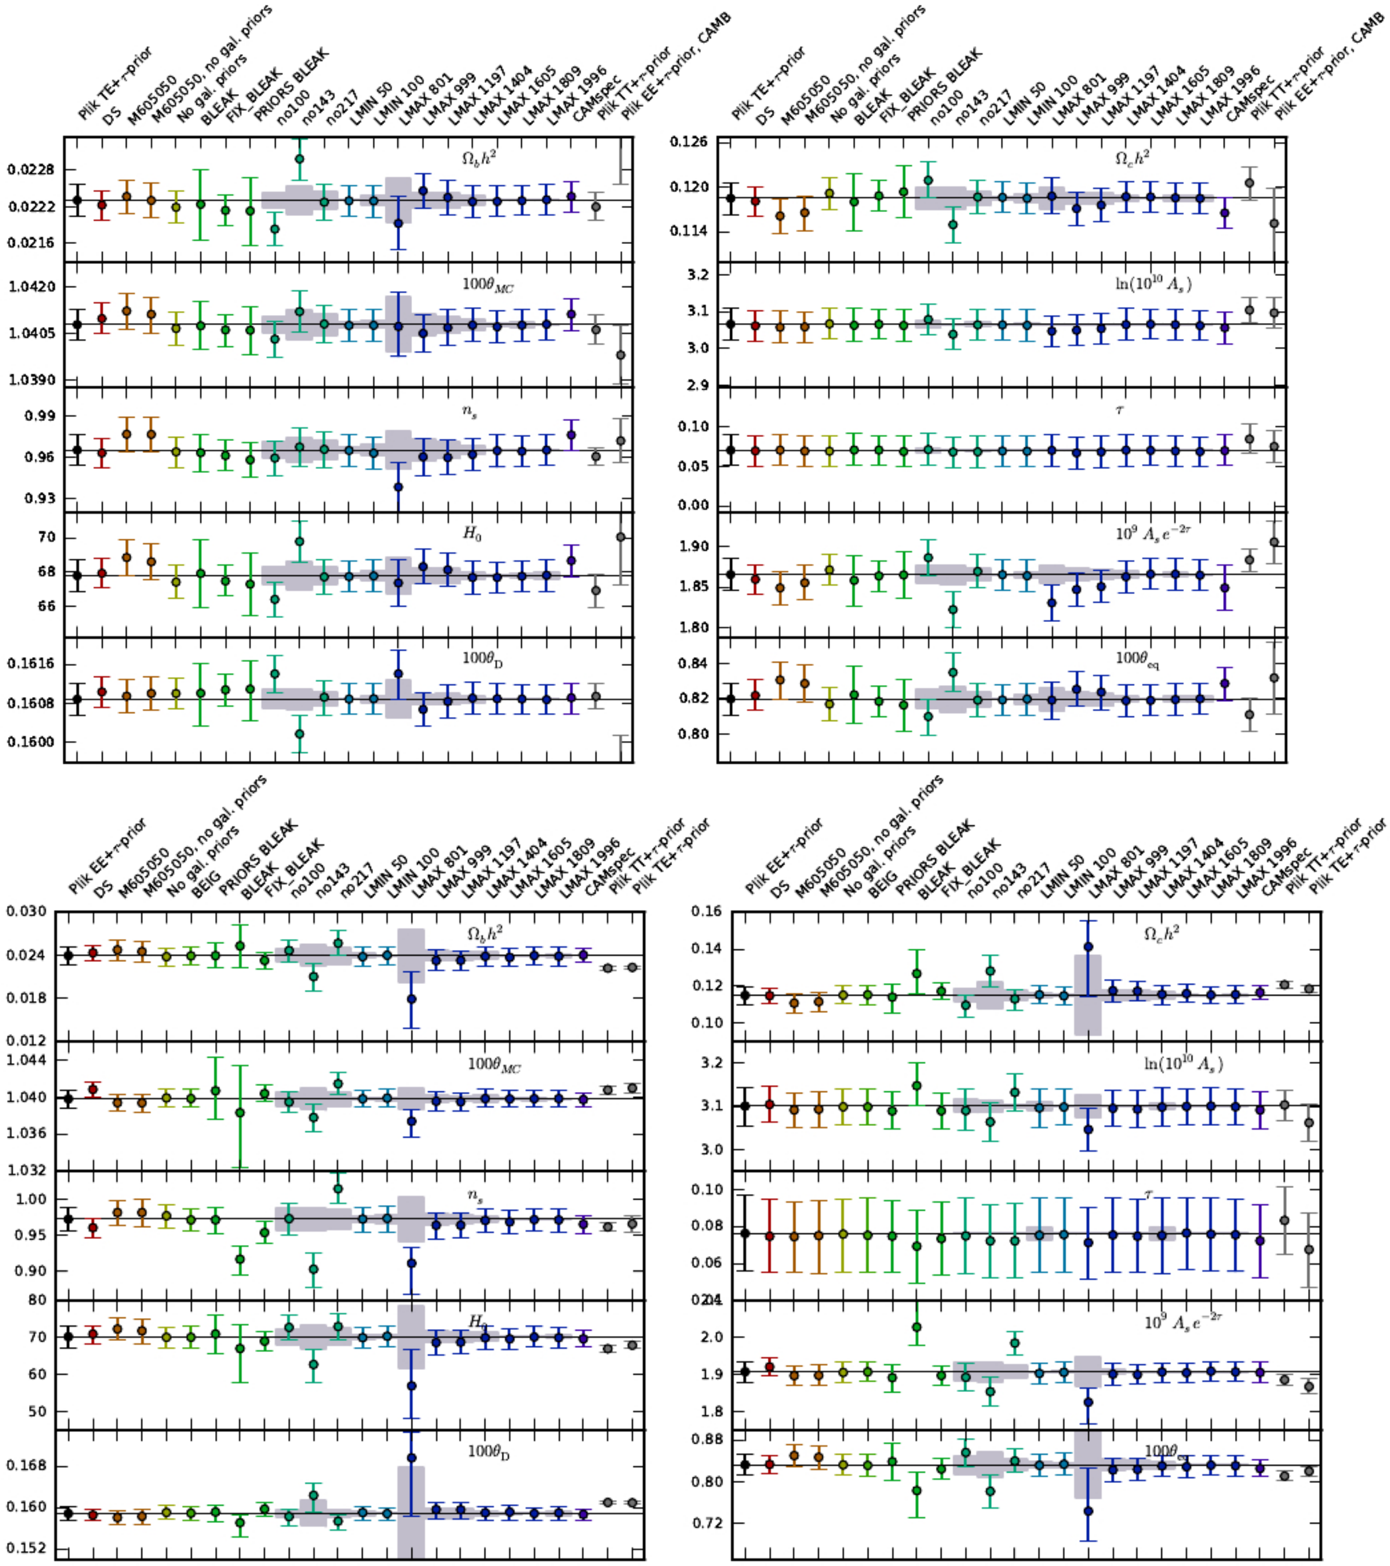

Fig. C.10. Marginal mean and 68% CL error bars on cosmological parameters estimated adopting different data choices for the `Plik` likelihood, in comparison with results from alternate approaches or model. *Top*: *TE* tests; we assume a ΛCDM model and use the `PlikTE`+tauprior likelihood in most of the cases, with a prior on $\tau = 0.07 \pm 0.02$ (we do not use low-ℓ temperature or polarization data here.). The "`PlikTE`+tauprior" case (black dot and thin horizontal black line) indicates the baseline (HM, $\ell_{\rm min} = 30$, $\ell_{\rm max} = 1996$), while the other cases are described in Appendix C.3.5. The grey bands show the standard deviation of the expected parameter shift, for those cases where the data used are a sub-sample of the baseline likelihood (see Eq. (53)). All the cases shown in these *TE* plots are run with `PICO`, except for the "`PlikEE`+tauprior, `CAMB`" case, which is run with `CAMB` (see Appendix C.5 for further details). *Bottom*: *EE* tests; the same as the top plots, but for the `PlikEE`+tauprior likelihood. For these *EE* plots we used `CAMB` instead of `PICO` to run all the cases (including `PlikTT`+tauprior and `PlikTE`+tauprior).

Fig. 8 of Galli et al. 2014), which show that the EE constraint on n_s is expected to be more than a factor of 2 weaker in the $\ell_{max} = 801$ case. This is confirmed by the tests presented here. Also, note that the grey bands in Fig. C.10, which indicate the standard deviation of the expected shifts, are calculated under the assumption of Gaussian parameter posterior distributions, and thus fail to properly describe non-Gaussian cases such as EE $\ell_{max} = 801$ considered here.

Comparison to `CamSpec`. We find relatively good consistency with the results of the `CamSpec` code, with shifts smaller than about 1σ in TE and 0.5σ in EE. Let us recall that the `CamSpec` and `Plik` codes adopt different choices of Galactic mask, Galactic dust treatment, and likelihood codes in polarization. Differences at this level therefore illustrate the good agreement reached for this release, and are useful to gauge the impact of quite different choices in the analysis procedures.

Remaining cases. As expected, the "lite" CMB-only likelihood is in agreement with the `Plik` code (see further discussion in Sect. C.6).

Finally, we note that in some of the cases discussed above, the calibration parameter for c_{217}^{TT} was wrongly set to unity instead of being to varied within its prior. We checked that this does not change our conclusions on the behaviour of the cosmological parameters and their uncertainties.

Summary. While a number of tests have been passed, the behaviour for masks, leakage parameters, and channel-data removal shows that systematic uncertainties are at least comparable to the statistical uncertainties. In the absence of a fully satisfactory data model, it is difficult to assess precisely the extent to which the extensive data averaging in the co-added TE or EE spectra effectively suppresses the residual systematic errors, many of which are detector-specific.

C.3.6. Agreement between temperature and polarization results

In order to assess the extent to which the cosmological parameters results that we obtain using the `PlikEE` or `PlikTE` data alone are compatible with those obtained from `PlikTT` alone, we performed the following test. We simulated 100 sets of TE or EE frequency power spectra conditioned on the TT power spectrum. As a fiducial model, we used the best-fit solution of the ΛCDM `PlikTT`+tauprior data combination. For all the polarization-related parameters (e.g., Galactic dust amplitudes) we used the best-fit solution of the `PlikTT`, TE, EE+tauprior data combination. We estimated cosmological parameters from each of these simulations, using the same assumptions as were adopted for the real data, and estimated the mean of the parameters obtained from the simulations. We then evaluated the deviation parameter $\mathcal{P}$ for each of the simulations as

$$\mathcal{P} = (\boldsymbol{P} - \langle \boldsymbol{P} \rangle)^{\mathsf{T}} \mathsf{P}^{-1} (\boldsymbol{P} - \langle \boldsymbol{P} \rangle), \tag{C.36}$$

where $\boldsymbol{P}$ is the vector of all varied parameters in the run (cosmological and foreground), P is the covariance matrix of the parameters, and $\langle \boldsymbol{P} \rangle$ is the mean of the parameters over the 100 simulations. The $\mathcal{P}$ parameter provides us a measure of how much all the parameters differ from their means, taking into account the correlations among them. We calculate the $\mathcal{P}$ parameter also for the results obtained from the real data, `PlikEE`+tauprior or `PlikTE`+tauprior, and compare these values to those obtained from the simulations. For EE, there are 36 simulations with a deviation $\mathcal{P}$ higher than the `PlikEE`+tauprior case, suggesting that the shifts in parameters we observe between `PlikEE`+tauprior and `PlikTT`+tauprior are in good agreement with expectations. For TE, there are 99 simulations with a deviation $\mathcal{P}$ higher than the `PlikTE`+tauprior case, suggesting that for TE the probability of obtaining parameters so close to the expected ones is only at the level of a few percent (although a more precise statement would require at least an order of magnitude more simulations). We note that this is not statistically very probable, but we could not identify any systematic reason why this should be so in all the tests conducted so far.

Figure C.11 shows the cosmological parameters obtained from the simulations (grey points), together with their mean (blue line). For clarity, we omit the error bars on the individual points (since they are all the same for each parameter), but show it instead as a light-blue band around the mean of the simulations. The cases shown in the figure are ordered by the $\mathcal{P}$ parameter from smallest to biggest (most "deviant").

It is interesting to note that the mean of the simulations, both for EE and TE, is very close to the cosmology obtained using the `PlikTT`+tauprior data, as expected. However, for A_s and τ, the mean of the simulations is almost 1σ lower than the value from `PlikTT`+tauprior. As explained in Sect. 4.1, the high value of A_s obtained from `PlikTT`+tauprior gives more lensing, better fitting the multipole region $\ell \approx 1400$–1500. This forces τ to adopt values about 1σ higher that those preferred by its Gaussian prior, in order to marginally compensate for the rise in A_s in the normalization of the power spectrum, $A_s \exp(-2\tau)$.

The high-ℓ TE and EE likelihoods, however, detect lensing at a much lower significance than in TT, and are thus sensitive only to the combination $A_s \exp(-2\tau)$. The individual constraints on A_s and τ are thus completely dominated by the prior on τ, centred on a value lower by about 1σ with respect to the value preferred by the `PlikTT`+tauprior data combination. As a consequence, the constraint on A_s from the simulated polarized spectra is lower that that obtained from the temperature data.

C.4. Co-added CMB spectra

This section illustrates the method we use to calculate the co-added CMB spectra. We first produce foreground-cleaned frequency power spectra using a fiducial model for the nuisance (e.g., foreground) parameters. The figures shown in Sects. 3–5 use the ΛCDM `PlikTT`+tauprior (PlanckTT+lowP) best-fit solution as a fiducial model for the temperature-related nuisance parameters, and `PlikTT`, TE, EE+tauprior (PlanckTT, TE, EE+lowP) for all the other polarization-specific nuisance parameters (e.g., polarized Galactic dust amplitudes). We then search for the maximum likelihood solution for the CMB power spectrum $C_\ell^{\rm CMB}$ that minimizes:

$$-\ln \mathcal{L}(\hat{\boldsymbol{C}}|\boldsymbol{C}^{\rm CMB}) \quad = \quad \frac{1}{2}\left[\hat{\boldsymbol{C}} - \boldsymbol{C}^{\rm CMB}\right]^{\mathsf{T}} \mathsf{C}^{-1} \left[\hat{\boldsymbol{C}} - \boldsymbol{C}^{\rm CMB}\right] + \text{const.}, \tag{C.37}$$

where $\hat{\boldsymbol{C}}$ is the foreground-cleaned frequency data vector, $\boldsymbol{C}^{\rm CMB}$ is the CMB vector we want to determine, and C is the covariance matrix. For instance, if we wanted to find the co-added CMB

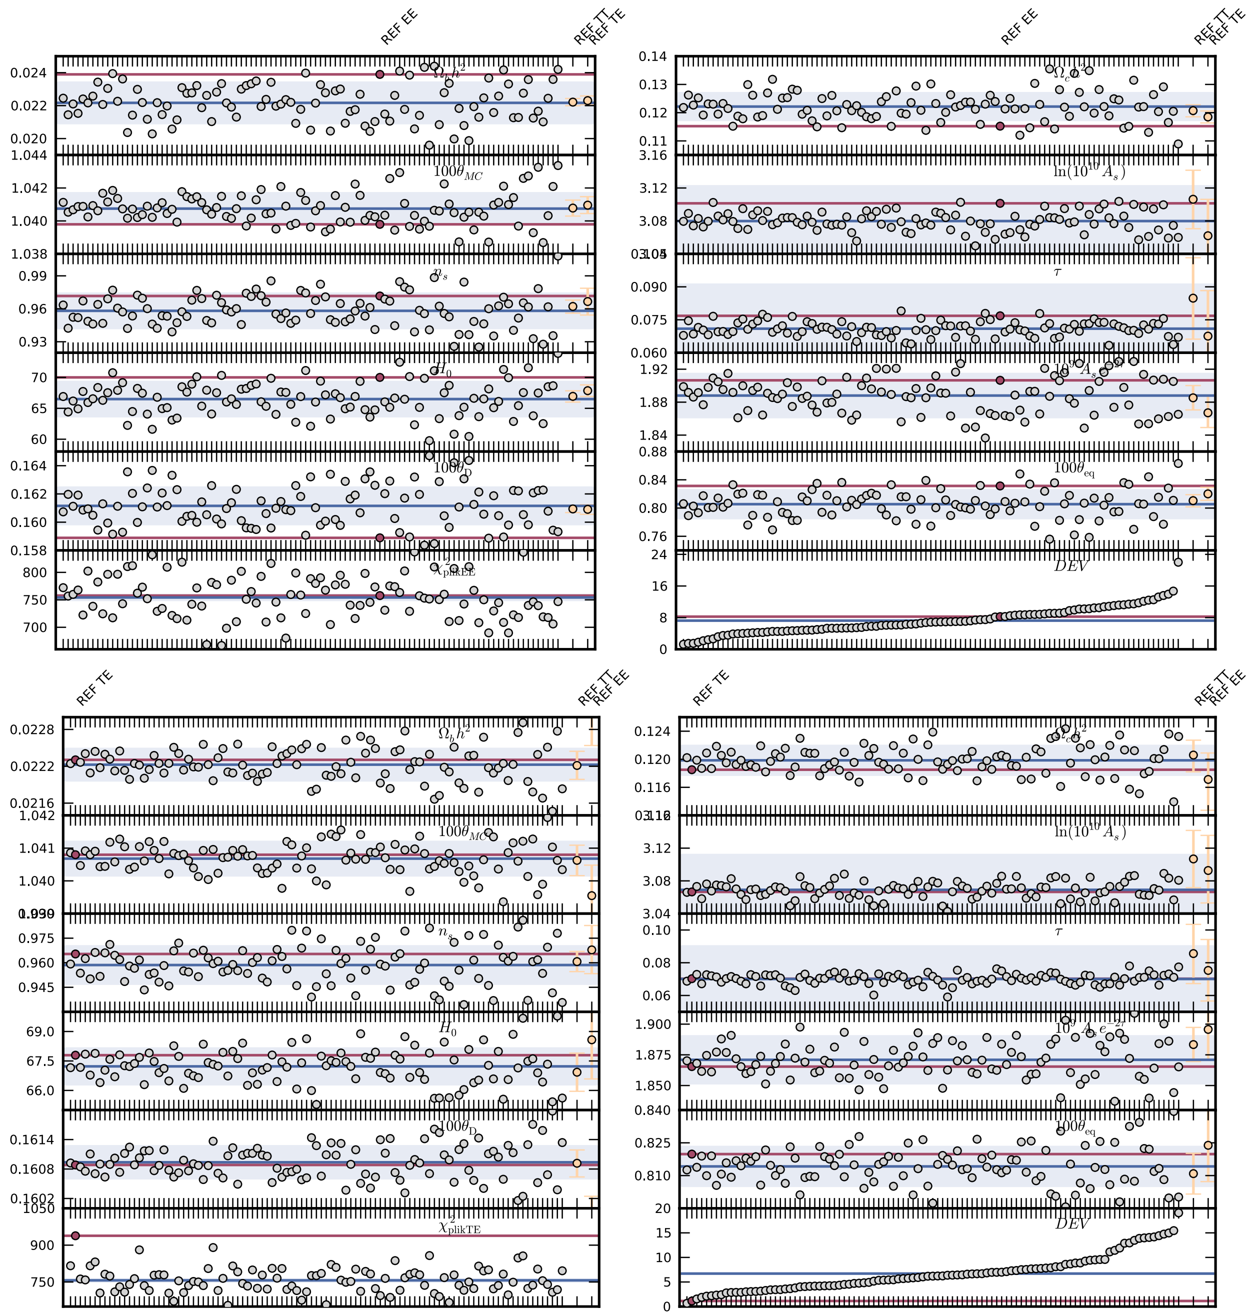

Fig. C.11. Marginal mean and 68% CL error bars on cosmological parameters estimated from 100 *EE* (*left*) or *TE* (*right*) power-spectra simulations conditioned on the TT power spectrum, assuming as a fiducial cosmology the best-fit of the ΛCDM `PlikTT`+tauprior results (grey circles). The blue line shows the mean of the simulations, i.e., the expected cosmology from the conditioned *EE* (or *TE*) spectra, while the blue band just shows the 68% CL error bar. The different cases are ordered from the least to the most "deviant" result according to the $\mathcal{P}$ parameter defined in Eq. (C.36) and called "DEV" in the plots. The `PlikEE`+tauprior, `PlikTE`+tauprior, and `PlikTT`+tauprior cases (in red or yellow) show the results from the real data. All the results in the *EE* plots were produced using the `CAMB` code, while those in the *TE* plots used the `PICO` code.

spectrum for TT alone, the vectors would be:

$$\hat{\boldsymbol{C}} = \left(\hat{C}^{TT}_{100\times100}, \hat{C}^{TT}_{143\times143}, \hat{C}^{TT}_{143\times217}, \hat{C}^{TT}_{217\times217}\right) \tag{C.38}$$

$$\boldsymbol{C}^{\rm CMB} = \left(\boldsymbol{C}^{TT,\rm CMB}, \boldsymbol{C}^{TT,\rm CMB}, \boldsymbol{C}^{TT,\rm CMB}, \boldsymbol{C}^{TT,\rm CMB}\right), \tag{C.39}$$

which we can rewrite

$$\boldsymbol{C}^{\rm CMB} = \mathsf{J}\, \boldsymbol{C}^{TT,\rm CMB}, \tag{C.40}$$

where J is a tall matrix which connects the power spectrum multipoles to the correct locations in the vector $\boldsymbol{C}^{\rm CMB}$; each column of the matrix contains only ones and zeros.

We minimize Eq. (C.37) by solving the linear system

$$\frac{\partial(-\ln \mathcal{L}(\hat{\boldsymbol{C}}))}{\partial \boldsymbol{C}^{\rm CMB}} = \frac{1}{2}\left(2\,\mathsf{J}^{\mathsf{T}}\mathsf{C}^{-1}\left[\hat{\boldsymbol{C}} - \mathsf{J}\,\boldsymbol{C}^{TT,\rm CMB}\right]\right) = 0, \tag{C.41}$$

where we used the fact that

$$\mathsf{J}^{\mathsf{T}}\mathsf{C}^{-1}\left[\hat{\boldsymbol{C}} - \boldsymbol{C}^{\rm CMB}\right] = \left(\mathsf{J}^{\mathsf{T}}\mathsf{C}^{-1}\left[\hat{\boldsymbol{C}} - \boldsymbol{C}^{\rm CMB}\right]\right)^{\mathsf{T}} = \left[\hat{\boldsymbol{C}} - \boldsymbol{C}^{\rm CMB}\right]^{\mathsf{T}}\mathsf{C}^{-1}\mathsf{J}, \tag{C.42}$$

since $\mathsf{C}^{-1} = (\mathsf{C}^{-1})^T$. The solution to Eq. (C.41) is just that of a generalized least-squares problem and is given by

$$\boldsymbol{C}^{TT,\mathrm{CMB}} = \left(\mathsf{J}^{\mathsf{T}}\mathsf{C}^{-1}\mathsf{J}\right)^{-1} \mathsf{J}^{\mathsf{T}}\mathsf{C}^{-1}\,\hat{\boldsymbol{C}}. \tag{C.43}$$

We then evaluate the covariance matrix $\mathsf{C}^{\mathrm{CMB}}$ of the co-added $\boldsymbol{C}^{TT,\mathrm{CMB}}$ spectrum as

$$\mathsf{C}^{\mathrm{CMB}} = \left(\mathsf{J}^{\mathsf{T}}\mathsf{C}^{-1}\mathsf{J}\right)^{-1}. \tag{C.44}$$

The matrix $\left(\mathsf{J}^{\mathsf{T}}\mathsf{C}^{-1}\mathsf{J}\right)^{-1}\ \mathsf{J}^{\mathsf{T}}\mathsf{C}^{-1}$ mixes the different frequency cross-spectra to compute the co-added solution. This matrix is flat and consists of the concatenation of blocks weighting each a particular cross-spectrum. Taking into account the different ℓ ranges for each, one can recast the blocks into diagonal-dominated square matrices. For a given multipole, ignoring the small out-of-band correlations, the relative weights of the cross-spectra in the co-added solution are given by the diagonals of those blocks. This is what we show Fig. 16.

C.5. PICO

We have used PICO to perform the extensive tests in this paper because it is much faster than CAMB, which is used in the *Planck* paper on parameters (Planck Collaboration XIII 2016). In this section we compare the results obtained using these two codes when evaluating cosmological parameters.

Table C.2 shows the parameter shifts (CAMB minus PICO), in units of standard deviations, assuming a ΛCDM model and using either code to evaluate cosmological parameters from the PlikTT+tauprior, PlikTE+tauprior, and PlikEE+tauprior data combinations. For the PlikTT+tauprior and the PlikTE+tauprior combinations, the biggest differences are in θ at about 0.3σ, and in n_{s} at about 0.2σ. These differences occur because (1) PICO was trained on the October 2012 version of CAMB whereas our CAMB runs use the January 2015 version (relevant differences include minor code changes and a slightly different default value of T_{CMB}); (2) PICO assumes three equal-mass neutrinos rather than one single massive one; and (3) a bug in the CosmoMC PICO wrapper caused a shift in N_{eff} of about 0.015. Despite these differences, the PICO results are sufficient for the inter-comparisons within this paper. While for PlikTT+tauprior and PlikTE+tauprior the PICO fitting error is negligible, for PlikEE+tauprior runs this is not the case, since the area of parameter space is much greater. For this reason, we actually use CAMB in these cases.

During the revision of this paper, we realized that this problem also affects the PlikTT likelihood test that excises the $\ell < 1000$ ($\ell_{\min} = 1000$ case) shown in Fig. 35. As mentioned in Sect. 4.1.6, this is due to the fact that this run explores regions of the parameter space that are wider than the PICO training region; this was also noticed by Addison et al. (2016). As a consequence, the results on n_{s} and $\Omega_{\mathrm{b}}h^2$ from this particular case have error bars underestimated by a factor of about two and mean values mis-estimated by about 0.8σ with respect to runs performed with CAMB. We therefore use CAMB rather than PICO to calculate results for this particular test.

Finally, we note that the definition of the A_{L} parameter used in CAMB is different from the one in PICO. The PICO A_{L} parameter is defined such that

$$C_\ell = A_{\mathrm{L}} C_\ell^{\mathrm{lensed}} + (1 - A_{\mathrm{L}}) C_\ell^{\mathrm{unlensed}}, \tag{C.45}$$

which is identical to CAMB's definition only to first order.

Table C.2. Differences between cosmological parameter estimates from CAMB and PICO.

Parameter	(CAMB−PICO)/σ(CAMB)) *TT*	*TE*	*EE*
$\Omega_{\mathrm{b}}h^2$	0.00	0.00	0.63
$\Omega_{\mathrm{c}}h^2$	0.05	−0.01	−0.40
θ	0.32	0.31	0.26
τ	−0.04	−0.14	0.08
$\ln(10^{10}A_{\mathrm{s}})$	−0.01	−0.11	0.21
n_{s}	0.23	0.10	0.24
H_0	0.01	0.07	0.52
$A_{\mathrm{s}}\exp(-2\tau)$	0.12	0.09	0.41

Notes. Parameter shifts, in standard deviations, obtained using PICO or CAMB. The results assume a ΛCDM model and the PlikTT+tauprior, PlikTE+tauprior, or PlikEE+tauprior data combinations.

C.6. Marginalized likelihood construction

C.6.1. Estimating temperature and polarization CMB-only spectra

The ℓ-range selection of the *Planck* high-ℓ likelihood defines $N_b = 613$ CMB band-powers, C_{b}. The C_{b} vector is structured in the following way: the first 215 elements describe the *Planck* *TT* CMB power spectrum, followed by 199 elements for the *EE* spectrum and 199 for *TE*.

The model for the theoretical power for a single cross-frequency spectrum (between frequencies i and j) in temperature or polarization, $C_\ell^{\mathrm{th},ij}$, is written as

$$C_\ell^{\mathrm{th},ij} = C_\ell^{\mathrm{CMB}} + C_\ell^{\mathrm{sec},ij}(\theta), \tag{C.46}$$

where $C_\ell^{\mathrm{sec},ij}(\theta)$ is the secondary signal given by thermal and kinetic SZ effects, clustered and Poisson point source emission, and Galactic emission, and is a function of secondary nuisance parameters θ. We convert $C_\ell^{\mathrm{th},ij}$ to band-powers by multiplying by the binning matrix $\mathcal{B}_{b\ell}$, i.e., $C_{\mathrm{b}}^{\mathrm{th},ij} = \sum_\ell \mathcal{B}_{b\ell} C_\ell^{\mathrm{th},ij}$. We then write the model for the C_{b} parameters in vector form as

$$C_{\mathrm{b}}^{\mathrm{th}} = \mathsf{A} C_{\mathrm{b}}^{\mathrm{CMB}} + C_{\mathrm{b}}^{\mathrm{sec}}(\theta), \tag{C.47}$$

where $C_{\mathrm{b}}^{\mathrm{th}}$ and $C_{\mathrm{b}}^{\mathrm{sec}}$ are multi-frequency spectra, and the mapping matrix A, with elements that are either 1 or 0, maps the CMB C_{b} vector (of length N_b), which is the same at all frequencies, onto the multi-frequency data. We calibrate the model as in the full multi-frequency likelihood, fixing the 143-GHz calibration factor to 1 and sampling the 100 and 217 calibration factors as nuisance parameters (i.e., as part of the θ vector).

We estimate $C_{\mathrm{b}}^{\mathrm{CMB}}$, marginalized over the secondary parameters, θ. The posterior distribution for $C_{\mathrm{b}}^{\mathrm{CMB}}$, given the observed multi-frequency spectra C_{b}, can be written as

$$p(C_{\mathrm{b}}^{\mathrm{CMB}}|C_{\mathrm{b}}) = \int p(C_{\mathrm{b}}^{\mathrm{CMB}}, \theta|C_{\mathrm{b}})p(\theta)\mathrm{d}\theta. \tag{C.48}$$

Rather than using, for example, Metropolis-Hastings, we use Gibbs sampling, which provides an efficient way to map out the joint distribution $p(C_{\mathrm{b}}^{\mathrm{CMB}}, \theta|C_{\mathrm{b}})$ and to extract the desired marginalized distribution $p(C_{\mathrm{b}}^{\mathrm{CMB}}|C_{\mathrm{b}})$. We do this by splitting the joint distribution into two conditional distributions: $p(C_{\mathrm{b}}^{\mathrm{CMB}}|\theta, C_{\mathrm{b}})$, and $p(\theta|C_{\mathrm{b}}^{\mathrm{CMB}}, C_{\mathrm{b}})$.

We write the multi-frequency *Planck* likelihood as

$$-2\ln\mathscr{L} = (\mathsf{A}C_\mathrm{b}^\mathrm{CMB} + C_\mathrm{b}^\mathrm{sec} - C_\mathrm{b})^\mathsf{T}\Sigma^{-1}(\mathsf{A}C_\mathrm{b}^\mathrm{CMB} + C_\mathrm{b}^\mathrm{sec} - C_\mathrm{b}) + \ln\det\Sigma, \tag{C.49}$$

which is a multivariate Gaussian. If $C_\mathrm{b}^\mathrm{sec}$ is held fixed, the conditional distribution for the CMB C_b parameters, $p(C_\mathrm{b}^\mathrm{CMB}|\theta, C_\mathrm{b})$, assuming a uniform prior for $p(C_\mathrm{b}^\mathrm{CMB})$, is then also a Gaussian. It has a distribution given by

$$-2\ln p(C_\mathrm{b}^\mathrm{CMB}|\theta, C_\mathrm{b}) = (C_\mathrm{b}^\mathrm{CMB} - \hat{C}_b)^\mathsf{T}\mathsf{Q}^{-1}(C_\mathrm{b}^\mathrm{CMB} - \hat{C}_b) + \ln\det\mathsf{Q}. \tag{C.50}$$

The mean, $\hat{C}_\mathrm{b}$, and covariance, Q, of this conditional distribution are obtained by taking the derivatives of the likelihood in Eq. (C.49) with respect to $C_\mathrm{b}^\mathrm{CMB}$. This gives mean

$$\hat{C}_b = \left[\mathsf{A}^T\Sigma^{-1}\mathsf{A}\right]^{-1}\left[\mathsf{A}^\mathsf{T}\Sigma^{-1}(C_\mathrm{b} - C_\mathrm{b}^\mathrm{sec})\right], \tag{C.51}$$

and covariance

$$\mathsf{Q} = \mathsf{A}^\mathsf{T}\Sigma^{-1}\mathsf{A}. \tag{C.52}$$

We draw a random sample from this Gaussian distribution by taking the Cholesky decomposition of the covariance matrix, $\mathsf{Q} = \mathsf{L}\mathsf{L}^\mathsf{T}$, and drawing a vector of Gaussian random variates G. The sample is then given by $C_\mathrm{b}^\mathrm{CMB} = \hat{C}_b + \mathsf{L}^{-1}G$.

If instead $C_\mathrm{b}^\mathrm{CMB}$ is held fixed, the conditional distribution for the secondary parameters, $p(\theta|C_\mathrm{b}^\mathrm{CMB}, C_\mathrm{b})$ can be sampled with the Metropolis algorithm in a simple MCMC code.

To map out the full joint distribution for θ and $C_\mathrm{b}^\mathrm{CMB}$ we alternate a Gibbs-sampling step, drawing a new vector $C_\mathrm{b}^\mathrm{CMB}$, with a Metropolis step, drawing a trial vector of the secondary parameters θ. About 700 000 steps are required for convergence of the joint distribution. The mean and covariance of the resulting marginalized CMB powers, $C_\mathrm{b}^\mathrm{CMB}$, are then estimated following the standard MCMC prescription.

Figure C.12 shows the multi-frequency data and the extracted CMB-only band-powers for TT, EE, and TE; the CMB is clearly separated out from foregrounds in both temperature and polarization.

Figure C.13 compares the nuisance parameters θ recovered in this model-independent sampling and the distributions obtained with the full likelihood. The parameters are consistent, with a broader distribution for the *Planck* Poisson sources. This degeneracy is observed because the sources can mimic blackbody emission and so are degenerate with the freely-varying CMB C_b parameters.

C.6.2. The `Plik_lite` CMB-only likelihood

We construct a CMB-only Gaussian likelihood from the extracted CMB C_b bandpowers in the following way:

$$-2\ln\mathscr{L}(\tilde{C}_b^\mathrm{CMB}|C_\mathrm{b}^\mathrm{th}) = \boldsymbol{x}^\mathsf{T}\tilde{\Sigma}^{-1}\boldsymbol{x}, \tag{C.53}$$

where $\boldsymbol{x} = \tilde{C}_b^\mathrm{CMB}/y_p^2 - C_\mathrm{b}^\mathrm{th}$, $\tilde{C}_b^\mathrm{CMB}$ and $\tilde{\Sigma}$ are the marginalized mean and covariance matrix for the C_bs, and C_b^th is the binned lensed CMB theory spectrum generated from `Plik`. The overall *Planck* calibration y_p is the only nuisance parameter left in this compressed likelihood. The Gaussianity assumption is a good approximation in the selected ℓ range, the extracted C_bs are well described by Gaussian distributions over the whole multiple range.

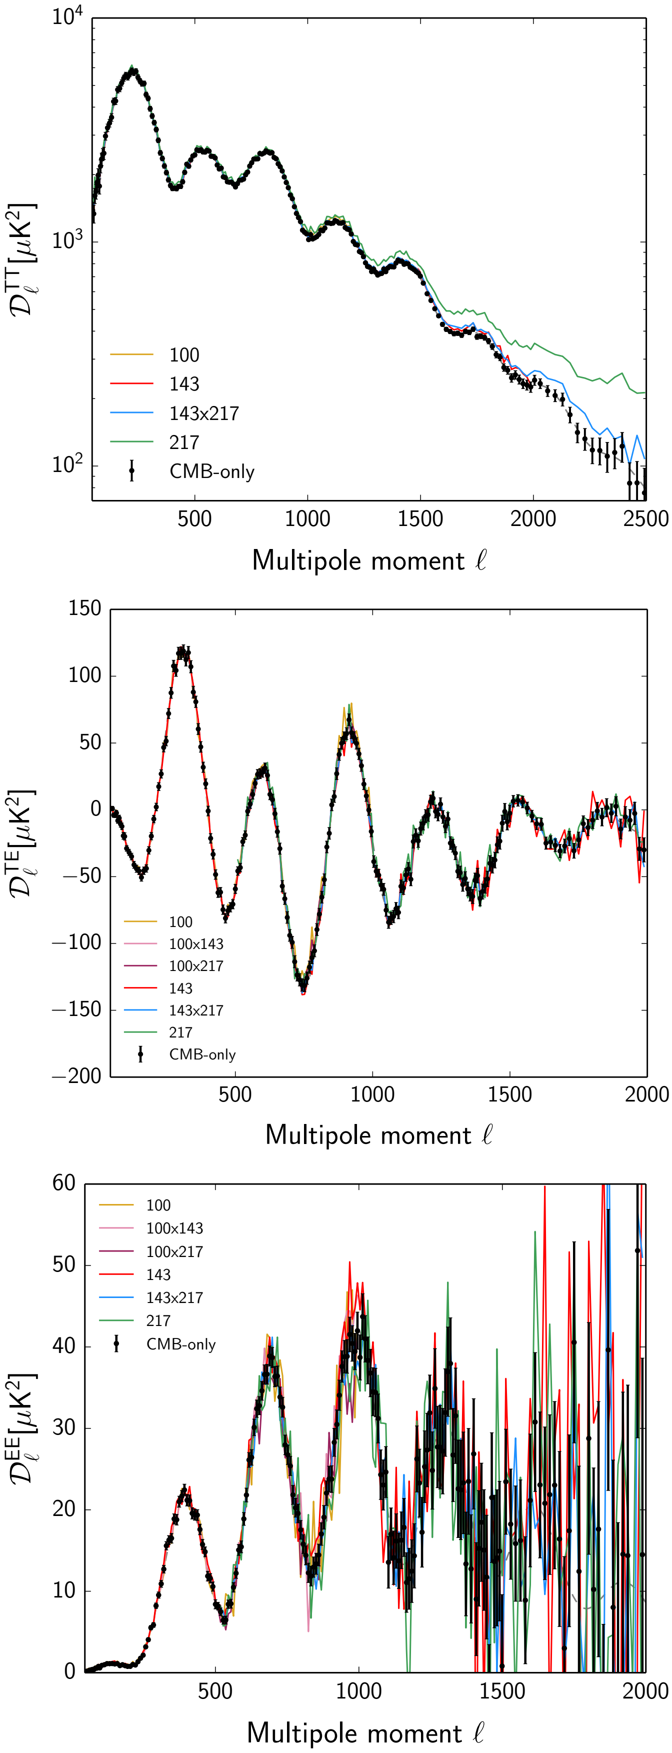

Fig. C.12. *Planck* multi-frequency power spectra (solid coloured lines) and extracted CMB-only spectra (black points).

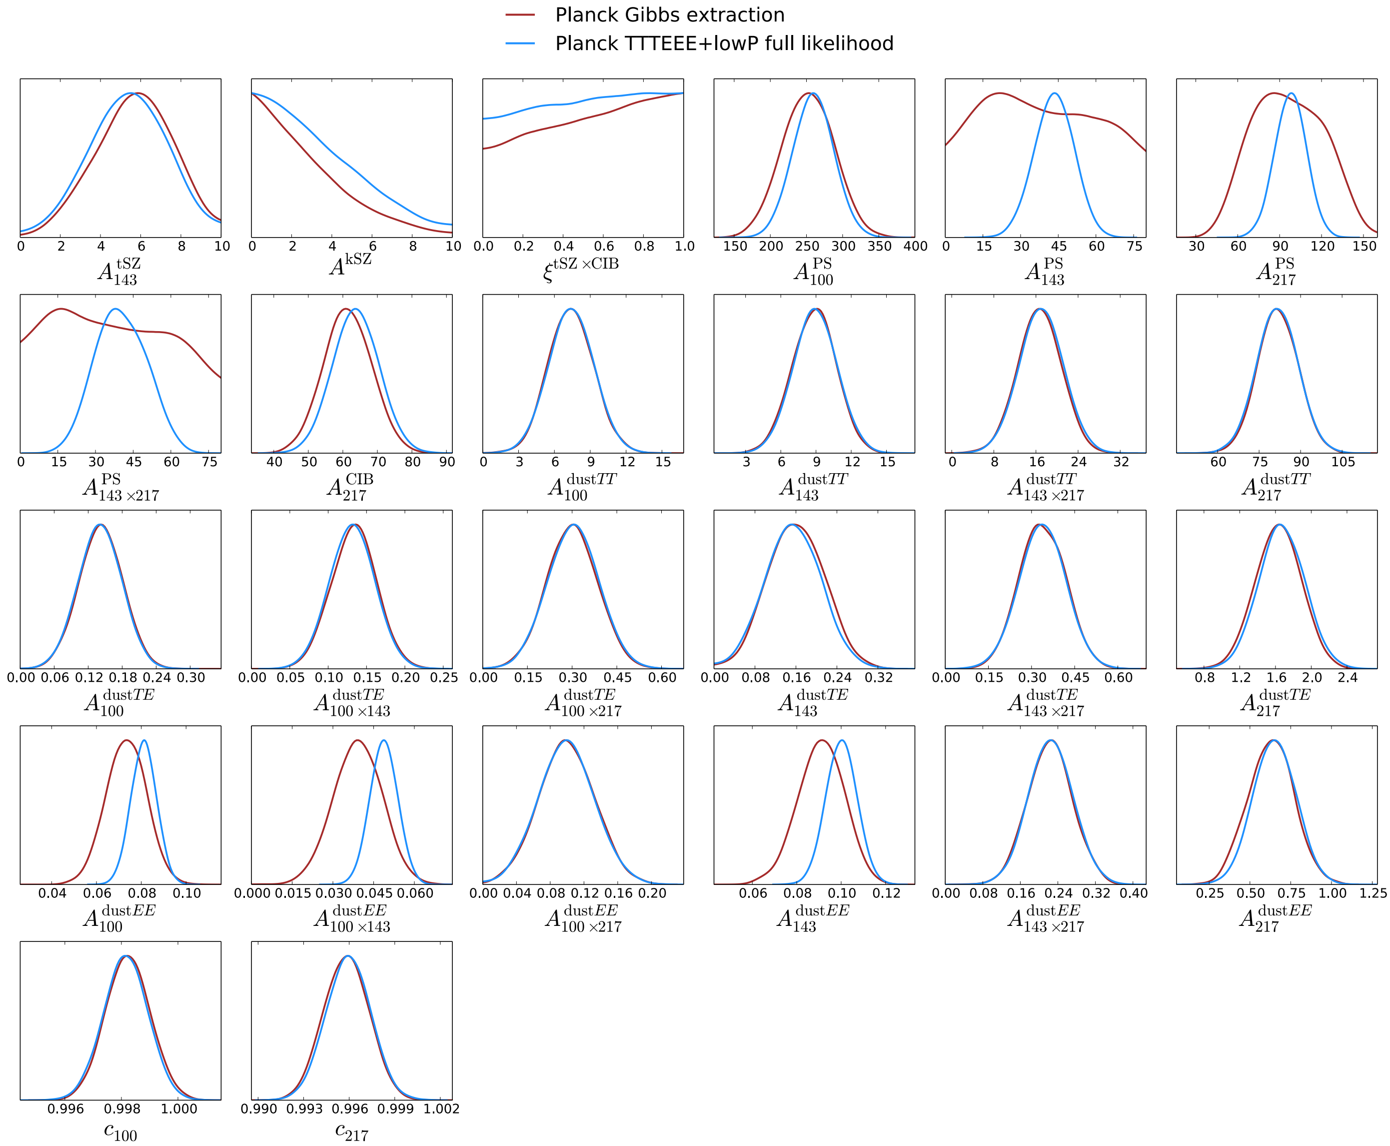

Fig. C.13. Comparison of the nuisance parameters estimated simultaneously with the CMB band-powers (red lines) and the results from the full multi-frequency likelihood (blue lines).

To test the performance of this compressed likelihood, we compare results using both the full multi-frequency likelihood and the CMB-only version. We report below examples for the baseline PlanckTT+lowP case. We first estimate cosmological parameters with `Plik_lite` for the restricted ΛCDM six-parameter model (see Fig. C.14) and compare them with the full-likelihood results. The agreement between the two methods is excellent, showing consistency to better than 0.1σ for all parameters.

We then extend the comparison to a set of six ΛCDM extensions, adding one parameter at a time to the base-ΛCDM model: the effective number of neutrino species $N_{\rm eff}$, the neutrino mass $\sum m_\nu$, the running of the spectral index $dn_s/d\ln k$, the tensor-to-scalar ratio r, the primordial helium fraction Y_P, and the lensing amplitude A_L. These parameters affect the damping tail more than the base set, and so are more correlated with the foreground parameters. Distributions for the added parameter in each of the six extensions are shown in Fig. C.15. Also in these cases we note that the agreement between the two methods is excellent, with all parameters differing by less than 0.1σ.

We find the same consistency when the polarization data are included in tests using the CMB-only high-ℓ TT, TE, and EE spectra in combination with lowP.

Appendix D: High-ℓ likelihood supplement

The *Planck* team have developed several independent approaches to the high-ℓ likelihood problem. These approaches and their implementations differ in several aspects, including the approximations, the foreground modelling, and the specific aspects that are checked. We have chosen `Plik`, for which the most supporting tests are available, as the baseline method. The comparison of the approaches given in the main text gives an indication of how well they agree, and the rather small differences give a feel for the remaining methodological uncertainties. In this appendix, we give a short description of two alternatives to `Plik`: `Mspec` and `Hillipop`. Another alternative, `CamSpec`, was the baseline for the previous *Planck* release, and has already been described in detail in Like13. Further comparison of `Plik` and

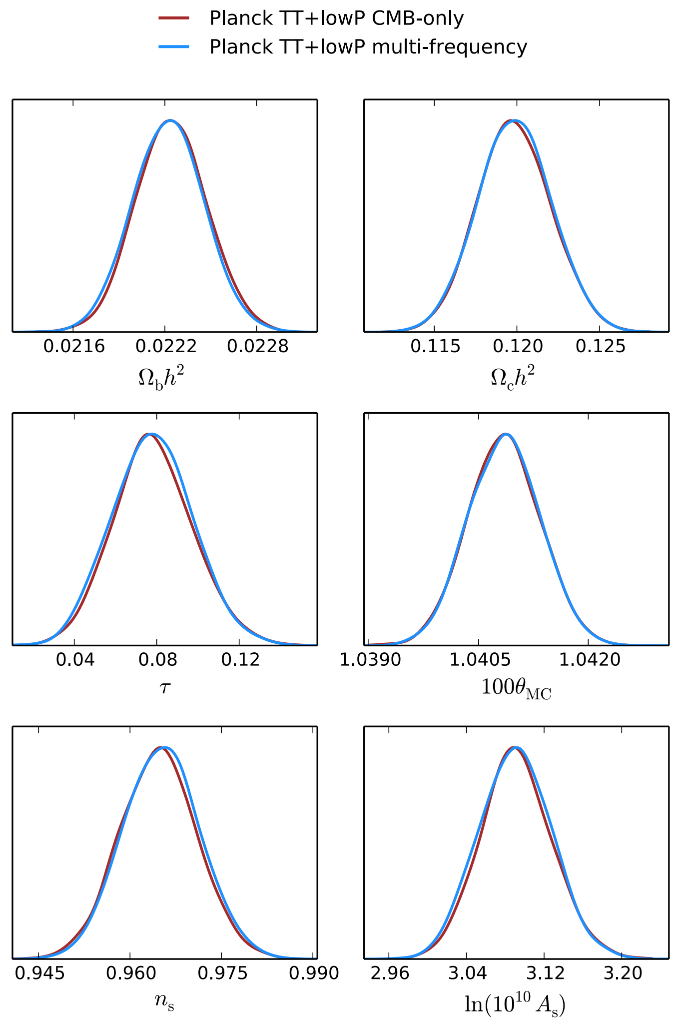

Fig. C.14. Comparison of the six base ΛCDM parameters estimated with the *Planck* compressed CMB-only likelihood (red lines) and the full multi-frequency likelihood (blue lines), in combination with *Planck* lowP data.

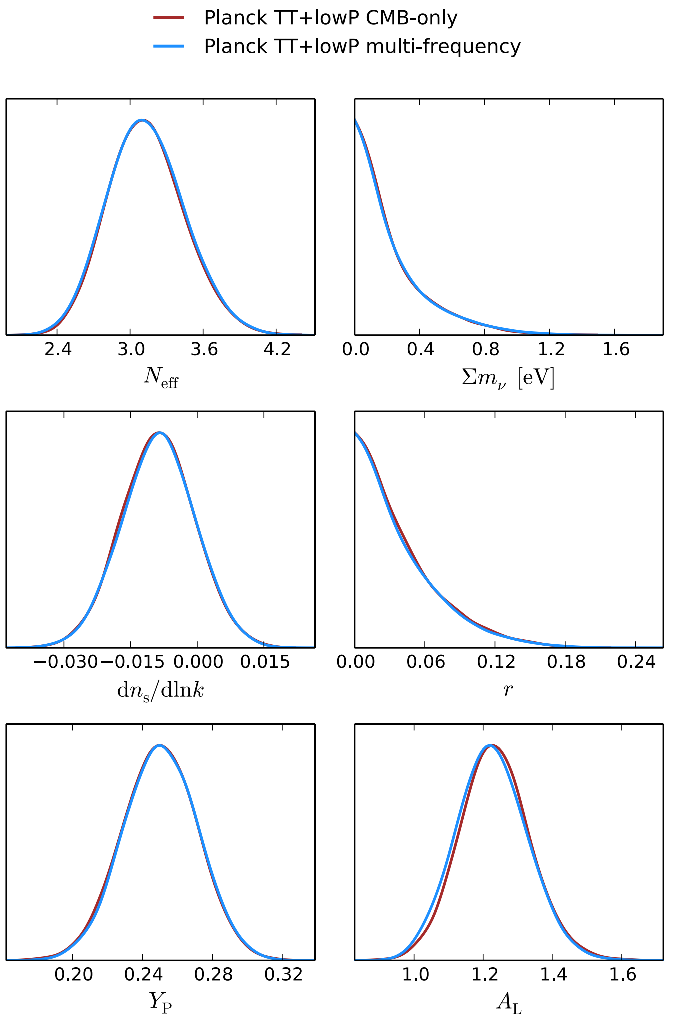

Fig. C.15. Comparison of extensions to the ΛCDM model from the CMB-only likelihood (red) and the multi-frequency likelihood (blue). There is excellent agreement between the two methods.

`CamSpec` is provided in the companion paper on cosmological parameters (Planck Collaboration XIII 2016).

D.1. `Mspec`

The `Mspec` likelihood differs from the baseline `Plik` likelihood mainly in the treatment of Galactic contamination in TT. `Mspec` results offer a cross-check of the baseline Galactic cleaning method, confirming that Galactic contamination does not have significant impact on the baseline parameters. A second smaller difference is the use of additional covariance approximations that reduce the computation cost while preserving satisfactory accuracy. We now describe these two aspects in more detail.

Galactic cleaning. Galactic dust cleaning in `Mspec` is a half-way point between some sophisticated component-separation methods (see Appendix E.4 and Planck Collaboration XII 2014) and the simple power-spectrum template subtraction or marginalization performed by `Plik`, `CamSpec`, and `Hillipop`. Component-separation methods are flexible and powerful, but propagation of beam and extragalactic-foreground uncertainties into the cleaned maps is difficult, and prohibitive in cost at high ℓ even when formally possible (e.g., for a Gibbs sampler). On the other hand, the power-spectrum template methods may be sensitive to errors in template shape and have bigger uncertainties due to signal-dust correlations.

`Mspec` cleaning is thus a two-step process. The first step is a simplified component-separation procedure that avoids the above shortcomings: we subtract a single scaled, high-frequency map from each CMB channel. This is very similar to the procedure used in Spergel et al. (2015), but it is targeted to remove Galactic as opposed to extragalactic contamination. It is also known as a "two-band ILC", and we refer to the procedure as "map cleaning" for short. The second step is to *then* subtract and marginalize a residual power-spectrum template model akin to the other likelihoods. We now describe each step in more detail.

In the map-cleaning step we subtract a scaled, higher-frequency *Planck* map from the lower-frequency CMB channels. This is a powerful method of cleaning, because the dust temperature is nearly uniform across the sky, and its intensity increases with frequency. High-frequency maps thus provide essentially noise-free dust maps that are highly correlated with the contamination at lower frequency. We choose to clean temperature maps with 545 GHz because it is less noisy than 353 GHz, but more