Planck 2015 results

Astronomy & Astrophysics

Special feature

Planck 2015 results

XI. CMB power spectra, likelihoods, and robustness of parameters

Planck Collaboration: N. Aghanim[63], M. Arnaud[78], M. Ashdown[74, 6], J. Aumont[63], C. Baccigalupi[91], A. J. Banday[103, 10], R. B. Barreiro[69], J. G. Bartlett[1, 71], N. Bartolo[32, 70], E. Battaner[105, 106], K. Benabed[64, 102], A. Benoît[61], A. Benoit-Lévy[24, 64, 102], J.-P. Bernard[103, 10], M. Bersanelli[35, 51], P. Bielewicz[86, 10, 91], J. J. Bock[71, 12], A. Bonaldi[72], L. Bonavera[20], J. R. Bond[9], J. Borrill[15, 96], F. R. Bouchet[64, 94, ★], F. Boulanger[63], M. Bucher[1], C. Burigana[50, 33, 52], R. C. Butler[50], E. Calabrese[98], J.-F. Cardoso[79, 1, 64], A. Catalano[80, 77], A. Challinor[66, 74, 13], H. C. Chiang[28, 7], P. R. Christensen[87, 38], D. L. Clements[59], L. P. L. Colombo[23, 71], C. Combet[80], A. Coulais[77], B. P. Crill[71, 12], A. Curto[69, 6, 74], F. Cuttaia[50], L. Danese[91], R. D. Davies[72], R. J. Davis[72], P. de Bernardis[34], A. de Rosa[50], G. de Zotti[47, 91], J. Delabrouille[1], F.-X. Désert[57], E. Di Valentino[64, 94], C. Dickinson[72], J. M. Diego[69], K. Dolag[104, 84], H. Dole[63, 62], S. Donzelli[51], O. Doré[71, 12], M. Douspis[63], A. Ducout[64, 59], J. Dunkley[98], X. Dupac[40], G. Efstathiou[66], F. Elsner[24, 64, 102], T. A. Enßlin[84], H. K. Eriksen[67], J. Fergusson[13], F. Finelli[50, 52], O. Forni[103, 10], M. Frailis[49], A. A. Fraisse[28], E. Franceschi[50], A. Frejsel[87], S. Galeotta[49], S. Galli[73], K. Ganga[1], C. Gauthier[1, 83], M. Gerbino[100, 89, 34], M. Giard[103, 10], E. Gjerløw[67], J. González-Nuevo[20, 69], K. M. Górski[71, 107], S. Gratton[74, 66], A. Gregorio[36, 49, 56], A. Gruppuso[50, 52], J. E. Gudmundsson[100, 89, 28], J. Hamann[101, 99], F. K. Hansen[67], D. L. Harrison[66, 74], G. Helou[12], S. Henrot-Versillé[76], C. Hernández-Monteagudo[14, 84], D. Herranz[69], S. R. Hildebrandt[71, 12], E. Hivon[64, 102], W. A. Holmes[71], A. Hornstrup[17], K. M. Huffenberger[26], G. Hurier[63], A. H. Jaffe[59], W. C. Jones[28], M. Juvela[27], E. Keihänen[27], R. Keskitalo[15], K. Kiiveri[27, 45], J. Knoche[84], L. Knox[29], M. Kunz[18, 63, 3], H. Kurki-Suonio[27, 45], G. Lagache[5, 63], A. Lähteenmäki[2, 45], J.-M. Lamarre[77], A. Lasenby[6, 74], M. Lattanzi[33, 53], C. R. Lawrence[71], M. Le Jeune[1], R. Leonardi[8], J. Lesgourgues[65, 101], F. Levrier[77], A. Lewis[25], M. Liguori[32, 70], P. B. Lilje[67], M. Lilley[64, 94], V. Lindholm[27, 45], M. Linden-Vørnle[17], V. Lindholm[27, 45], M. López-Caniego[40], J. F. Macías-Pérez[80], B. Maffei[72], G. Maggio[49], D. Maino[35, 51], N. Mandolesi[50, 33], A. Mangilli[63, 76], M. Maris[49], P. G. Martin[9], E. Martínez-González[69], S. Masi[34], S. Matarrese[32, 70, 42], P. R. Meinhold[30], A. Melchiorri[34, 54], M. Migliaccio[66, 74], M. Millea[29], S. Mitra[58, 71], M.-A. Miville-Deschênes[63, 9], A. Moneti[64], L. Montier[103, 10], G. Morgante[50], D. Mortlock[59], S. Mottet[64, 94], D. Munshi[93], J. A. Murphy[85], A. Narimani[22], P. Naselsky[88, 39], F. Nati[28], P. Natoli[33, 4, 53], F. Noviello[72], D. Novikov[82], I. Novikov[87, 82], C. A. Oxborrow[17], F. Paci[91], L. Pagano[34, 54], F. Pajot[63], D. Paoletti[50, 52], B. Partridge[44], F. Pasian[49], G. Patanchon[1], T. J. Pearson[12, 60], O. Perdereau[76], L. Perotto[80], V. Pettorino[43], F. Piacentini[34], M. Piat[1], E. Pierpaoli[23], D. Pietrobon[71], S. Plaszczynski[76], E. Pointecouteau[103, 10], G. Polenta[4, 48], N. Ponthieu[63, 57], G. W. Pratt[78], S. Prunet[64, 102], J.-L. Puget[63], J. P. Rachen[21, 84], M. Reinecke[84], M. Remazeilles[72, 63, 1], C. Renault[80], A. Renzi[37, 55], I. Ristorcelli[103, 10], G. Rocha[71, 12], M. Rossetti[35, 51], G. Roudier[1, 77, 71], B. Rouillé d'Orfeuil[76], J. A. Rubiño-Martín[68, 19], B. Rusholme[60], L. Salvati[34], M. Sandri[50], D. Santos[80], M. Savelainen[27, 45], G. Savini[90], D. Scott[22], P. Serra[63], L. D. Spencer[93], M. Spinelli[76], V. Stolyarov[6, 97, 75], R. Stompor[1], R. Sunyaev[84, 95], D. Sutton[66, 74], A.-S. Suur-Uski[27, 45], J.-F. Sygnet[64], J. A. Tauber[41], L. Terenzi[92, 50], L. Toffolatti[20, 69, 50], M. Tomasi[35, 51], M. Tristram[76], T. Trombetti[50, 33], M. Tucci[18], J. Tuovinen[11], G. Umana[46], L. Valenziano[50], J. Valiviita[27, 45], F. Van Tent[81], P. Vielva[69], F. Villa[50], L. A. Wade[71], B. D. Wandelt[64, 102, 31], I. K. Wehus[71, 67], D. Yvon[16], A. Zacchei[49], and A. Zonca[30]

(Affiliations can be found after the references)

ABSTRACT

This paper presents the *Planck* 2015 likelihoods, statistical descriptions of the 2-point correlation functions of the cosmic microwave background (CMB) temperature and polarization fluctuations that account for relevant uncertainties, both instrumental and astrophysical in nature. They are based on the same hybrid approach used for the previous release, i.e., a pixel-based likelihood at low multipoles ($\ell < 30$) and a Gaussian approximation to the distribution of cross-power spectra at higher multipoles. The main improvements are the use of more and better processed data and of *Planck* polarization information, along with more detailed models of foregrounds and instrumental uncertainties. The increased redundancy brought by more than doubling the amount of data analysed enables further consistency checks and enhanced immunity to systematic effects. It also improves the constraining power of *Planck*, in particular with regard to small-scale foreground properties. Progress in the modelling of foreground emission enables the retention of a larger fraction of the sky to determine the properties of the CMB, which also contributes to the enhanced precision of the spectra. Improvements in data processing and instrumental modelling further reduce uncertainties. Extensive tests establish the robustness and accuracy of the likelihood results, from temperature alone, from polarization alone, and from their combination. For temperature, we also perform a full likelihood analysis of realistic end-to-end simulations of the instrumental response to the sky, which were fed into the actual data processing pipeline; this does not reveal biases from residual low-level instrumental systematics. Even with the increase in precision and robustness, the ΛCDM cosmological model continues to offer a very good fit to the *Planck* data. The slope of the primordial scalar fluctuations, n_s, is confirmed smaller than unity at more than 5σ from *Planck* alone. We further validate the robustness of the likelihood results against specific extensions to the baseline cosmology, which are particularly sensitive to data at high multipoles. For instance, the effective number of neutrino species remains compatible with the canonical value of 3.046. For this first detailed analysis of *Planck* polarization spectra, we concentrate at high multipoles on the E modes, leaving the analysis of the weaker B modes to future work. At low multipoles we use temperature maps at all *Planck* frequencies along with a subset of polarization data. These data take advantage of *Planck*'s wide frequency coverage to improve the separation of CMB and foreground emission. Within the baseline ΛCDM cosmology this requires $\tau = 0.078 \pm 0.019$ for the reionization optical depth, which is significantly lower than estimates without the use of high-frequency data for explicit monitoring of dust emission. At high multipoles we detect residual systematic errors in E polarization, typically at the μK^2 level; we therefore choose to retain temperature information alone for high multipoles as the recommended baseline, in particular for testing non-minimal models. Nevertheless, the high-multipole polarization spectra from *Planck* are already good enough to enable a separate high-precision determination of the parameters of the ΛCDM model, showing consistency with those established independently from temperature information alone.

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

1. Introduction

This paper presents the angular power spectra of the cosmic microwave background (CMB) and the related likelihood functions, calculated from *Planck*[1] 2015 data, which consists of intensity maps from the full mission, along with a subset of the polarization data.

The CMB power spectra contain all of the information available if the CMB is statistically isotropic and distributed as a multivariate Gaussian. For realistic data, these must be augmented with models of instrumental noise, of other instrumental systematic effects, and of contamination from astrophysical foregrounds.

The power spectra are, in turn, uniquely determined by the underlying cosmological model and its parameters. In temperature, the power spectrum has been measured over large fractions of the sky by the Cosmic Background Explorer (COBE; Wright et al. 1996) and the *Wilkinson* Microwave Anistropy Probe (WMAP; Bennett et al. 2013), and in smaller regions by a host of balloon- and ground-based telescopes (e.g., Netterfield et al. 1997; Hanany et al. 2000; Grainge et al. 2003; Pearson et al. 2003; Tristram et al. 2005b; Jones et al. 2006; Reichardt et al. 2009; Fowler et al. 2010; Das et al. 2011, 2014; Keisler et al. 2011; Story et al. 2013). The *Planck* 2013 power spectrum and likelihood were discussed in Planck Collaboration XV (2014, hereafter Like13).

The distribution of temperature and polarization on the sky is further affected by gravitational lensing by the inhomogeneous mass distribution along the line of sight between the last scattering surface and the observer. This introduces correlations between large and small scales, which can be estimated by computing the expected contribution of lensing to the 4-point function (i.e., the trispectrum). This can in turn be used to determine the power spectrum of the lensing potential, as is done in Planck Collaboration XV (2016) for this *Planck* release, and to further constrain the cosmological parameters via a separate likelihood function (Planck Collaboration XIII 2016).

Over the last decade, CMB intensity (temperature) has been augmented by linear polarization data (e.g., Kovac et al. 2002; Kogut et al. 2003; Sievers et al. 2007; Dunkley et al. 2009; Pryke et al. 2009; Araujo et al. 2012; Polarbear Collaboration 2014). Because linear polarization is given by both an amplitude and direction, it can, in turn, be decomposed into two coordinate-independent quantities, each with a different dependence on the cosmology (e.g., Seljak 1997; Kamionkowski et al. 1997; Zaldarriaga & Seljak 1997). One, the so-called *E* mode, is determined by much the same physics as the intensity, and therefore enables an independent measurement of the background cosmology, as well as a determination of some new parameters (e.g., the reionization optical depth). The other polarization observable, the *B* mode, is only sourced at early times by gravitational radiation, as produced, for example, during an inflationary epoch. The *E* and *B* components are also conventionally taken to be isotropic Gaussian random fields, with only *E* expected to be correlated with intensity. Thus we expect to be able to measure four independent power spectra, namely the three auto-spectra C_ℓ^{TT}, C_ℓ^{EE}, and C_ℓ^{BB}, along with the cross-spectrum C_ℓ^{TE}.

Estimating these spectra from the likelihood requires cleaned and calibrated maps for all *Planck* detectors, along with a quantitative description of their noise properties. The required data processing is discussed in Planck Collaboration II (2016), Planck Collaboration III (2016), Planck Collaboration IV (2016), Planck Collaboration V (2016), and Planck Collaboration VIII (2016) for the low-frequency instrument (LFI; 30, 44, and 70 GHz) and Planck Collaboration VII (2016) and Planck Collaboration VIII (2016) for the high-frequency instrument (HFI; 100, 143, 217, 353, 585, and 857 GHz). Although the CMB is brightest over 70–217 GHz, the full range of *Planck* frequencies is crucial to distinguish between the cosmological component and sources of astrophysical foreground emission, present in even the cleanest regions of sky. We therefore use measurements from those *Planck* bands dominated by such emission as a template to model the foreground in the bands where the CMB is most significant.

This paper presents the C_ℓ^{TT}, C_ℓ^{EE}, and C_ℓ^{TE} spectra, likelihood functions, and basic cosmological parameters from the *Planck* 2015 release. A complete analysis in the context of an extended ΛCDM cosmology of these and other results from *Planck* regarding the lensing power spectrum results, as well as constraints from other observations, is given in Planck Collaboration XIII (2016). Wider extensions to the set of models are discussed in other *Planck* 2015 papers; for example, Planck Collaboration XIV (2016) examines specific models for the dark energy component and extensions to general relativity, and Planck Collaboration XX (2016) discusses inflationary models.

This paper shows that the contribution of high-ℓ systematic errors to the polarization spectra are at quite a low level (of the order of a few μK^2), therefore enabling an interesting comparison of the polarization-based cosmological results with those derived from C_ℓ^{TT} alone. We therefore discuss the results for C_ℓ^{TE} and C_ℓ^{EE} at high multipoles. However, the technical difficulties involved with polarization measurements and subsequent data analysis, along with the inherently lower signal-to-noise ratio (especially for *B* modes), thus require a careful understanding of the random noise and instrumental and astrophysical systematic effects. For this reason, at large angular scales (i.e., low multipoles ℓ) the baseline results use only a subset of *Planck* polarization data.

Because of these different sensitivities to systematic errors at different angular scales, as well as the increasingly Gaussian behaviour of the likelihood function at smaller angular scales, we adopt a hybrid approach to the likelihood calculation (Efstathiou 2004, 2006), splitting between a direct calculation of the likelihood on large scales and the use of pseudo-spectral estimates at smaller scales, as we did for the previous release.

The plan of the paper reflects this hybrid approach along with the importance of internal tests and cross-validation. In Sect. 2, we present the low-multipole ($\ell < 30$) likelihood and its validation. At these large scales, we compute the likelihood function directly in pixel space; the temperature map is obtained by a Gibbs sampling approach in the context of a parameterized foreground model, while the polarized maps are cleaned of foregrounds by a template removal technique.

In Sect. 3, we introduce the high-multipole ($\ell \geq 30$) likelihood and present its main results. At these smaller scales, we employ a pseudo-C_ℓ approach, beginning with a numerical spherical harmonic transform of the full-sky map, debiased and deconvolved to account for the mask and noise.

Section 4 is devoted to the detailed assessment of this high-ℓ likelihood. One technical difference between Like13

Table 1. Likelihood codes and datasets.

Name	Description
PlanckTT	Full *Planck* temperature-only C_ℓ^{TT} likelihood
PlanckTT, TE, EE ...	PlanckTT combined with high-ℓ $C_\ell^{TE}+C_\ell^{EE}$ likelihood
lowP	Low-ℓ polarization $C_\ell^{TE} + C_\ell^{EE} + C_\ell^{BB}$ likelihood
lowTEB	Low-ℓ temperature-plus-polarization likelihood
PlikTT...........	High-ℓ C_ℓ^{TT}-only likelihood
PlikEE...........	High-ℓ C_ℓ^{EE}-only likelihood
PlikTE...........	High-ℓ C_ℓ^{TE}-only likelihood
PlikTT, TE, EE	High-ℓ $C_\ell^{TT} + C_\ell^{TE} + C_\ell^{EE}$ likelihood
Plik_lite	High-ℓ $C_b^{TT} + C_b^{TE} + C_b^{EE}$, foreground-marginalized bandpower likelihood
tauprior	Gaussian prior, $\tau = 0.07 \pm 0.02$
highL	ACT+SPT high-ℓ likelihood
WP	WMAP low-ℓ polarization likelihood[a]

Notes. We use these designations throughout the text to refer to specific likelihood codes and implementations that use different input data. A sum of spectra in the description column designates the joint likelihood of these spectra. [a] "Low-ℓ" refers to $\ell < 23$ for WP, but $\ell < 30$ for the *Planck* likelihoods.

and the present work is the move from the `CamSpec` code to `Plik` for high-ℓ results as well as the released software (Planck Collaboration 2015). The main reason for this change is that the structure of `Plik` allows more fine-grained tests on the polarization spectra for individual detectors or subsets of detectors. We are able to compare the effect of different cuts on *Planck* and external data, as well as using methods that take different approaches to estimate the maximum-likelihood spectra from the input maps; these illustrate the small impact of differences in methodology and data preparation, which are difficult to assess otherwise.

We then combine the low- and high-ℓ algorithms to form the full *Planck* likelihood in Sect. 5, assessing there the choice of $\ell = 30$ for the hybridization scale and establishing the basic cosmological results from *Planck* 2015 data alone.

Finally, in Sect. 6 we conclude. A series of Appendices discusses sky masks and gives more detail on the individual likelihood codes, both the released version and a series of other codes used to validate the overall methodology.

To help distinguish the many different likelihood codes, which are functions of different parameters and use different input data, Table 1 summarizes the designations used throughout the text.

2. Low-multipole likelihood

At low multipoles, the current *Planck* release implements a standard joint pixel-based likelihood including both temperature and polarization for multipoles $\ell \leq 29$. Throughout this paper, we denote this likelihood "lowTEB", while "lowP" denotes the polarization part of this likelihood. For temperature, the formalism uses the CMB maps cleaned with `Commander` (Eriksen et al. 2004, 2008) maps, while for polarization we use the 70 GHz LFI maps and explicitly marginalize over the 30 GHz and 353 GHz maps taken as tracers of synchrotron and dust emission, respectively (see Sect. 2.3), accounting in both cases for the induced noise covariance in the likelihood.

This approach is somewhat different from the *Planck* 2013 low-ℓ likelihood. As described in Like13, this comprised two nearly independent components, covering temperature and polarization information, respectively. The temperature likelihood employed a Blackwell-Rao estimator (Chu et al. 2005) at $\ell \leq$ 49, averaging over Monte Carlo samples drawn from the exact

power spectrum posterior using `Commander`. For polarization, we had adopted the pixel-based 9-year WMAP polarization likelihood, covering multipoles $\ell \leq 23$ (Bennett et al. 2013).

The main advantage of the exact joint approach now employed is mathematical rigour and consistency to higher ℓ, while the main disadvantage is a slightly higher computational expense due to the higher pixel resolution required to extend the calculation to $\ell = 29$ in polarization. However, after implementation of the Sherman-Morrison-Woodbury formula to reduce computational costs (see Appendix B.1), the two approaches perform similarly, both with respect to speed and accuracy, and our choice is primarily a matter of implementational convenience and flexibility, rather than actual results or performance.

2.1. Statistical description and algorithm

We start by reviewing the general CMB likelihood formalism for the analysis of temperature and polarization at low ℓ, as described for instance by Tegmark & de Oliveira-Costa (2001), Page et al. (2007), and in Like13. We begin with maps of the three Stokes parameters $\{T, Q, U\}$ for the observed CMB intensity and linear polarization in some set of `HEALPix`[2] (Górski et al. 2005) pixels on the sky. In order to use multipoles $\ell \leq \ell_{cut} = 29$ in the likelihood, we adopt a `HEALPix` resolution of $N_{side} = 16$ which has 3072 pixels (of area 13.6 deg^2) per map; this accommodates multipoles up to $\ell_{max} = 3N_{side} - 1 = 47$, and, considering separate maps of T, Q, and U, corresponds to a maximum of $N_{pix} = 3 \times 3072 = 9216$ pixels in any given calculation, not accounting for any masking.

After component separation, the data vector may be modelled as a sum of cosmological CMB signal and instrumental noise, $\boldsymbol{m}^X = \boldsymbol{s}^X + \boldsymbol{n}^X$, where $\boldsymbol{s}$ is assumed to be a set of statistically isotropic and Gaussian-distributed random fields on the sky, indexed by pixel or spherical-harmonic indices (ℓm), with $X = \{T, E, B\}$ selecting the appropriate intensity or polarization component. The signal fields $\boldsymbol{s}^X$ have auto- and cross-power spectra C_ℓ^{XY} and a pixel-space covariance matrix

$$\mathsf{S}(C_\ell) = \sum_{\ell=2}^{\ell_{max}} \sum_{XY} C_\ell^{XY} \mathsf{P}_\ell^{XY}. \qquad (1)$$

---

[2] http://healpix.sourceforge.org

Here we restrict the spectra to $XY = \{TT, EE, BB, TE\}$, with $N_{\text{side}} = 16$ pixelization, and P_ℓ^{XY} is a beam-weighted sum over (associated) Legendre polynomials. For temperature, the explicit expression is

$$(\mathsf{P}_\ell^{TT})_{i,j} = \frac{2\ell + 1}{4\pi} B_\ell^2 P_\ell(\hat{\boldsymbol{n}}_i \cdot \hat{\boldsymbol{n}}_j), \tag{2}$$

where $\hat{\boldsymbol{n}}_i$ is a unit vector pointing towards pixel i, B_ℓ is the product of the instrumental beam Legendre transform and the HEALPix pixel window, and P_ℓ is the Legendre polynomial of order ℓ; for corresponding polarization components, see, e.g., Tegmark & de Oliveira-Costa (2001). The instrumental noise is also assumed to be Gaussian distributed, with a covariance matrix N that depends on the *Planck* detector sensitivity and scanning strategy, and the full data covariance is therefore $\mathsf{M} = \mathsf{S} + \mathsf{N}$. With these definitions, the full likelihood expression reads

$$\mathcal{L}(C_\ell) = \mathcal{P}(\boldsymbol{m}|C_\ell) = \frac{1}{2\pi|\mathsf{M}|^{1/2}} \exp\left(-\frac{1}{2}\boldsymbol{m}^\mathsf{T} \mathsf{M}^{-1} \boldsymbol{m}\right), \tag{3}$$

where the conditional probability $\mathcal{P}(\boldsymbol{m}|C_\ell)$ defines the likelihood $\mathcal{L}(C_\ell)$.

The computational cost of this expression is driven by the presence of the matrix inverse and determinant operations, both of which scale computationally as $O(N_{\text{pix}}^3)$. For this reason, the direct approach is only computationally feasible at large angular scales, where the number of pixels is low. In practice, we only analyse multipoles below or equal to $\ell_{\text{cut}} = 29$ with this formalism, requiring maps with $N_{\text{side}} = 16$. Multipoles between $\ell_{\text{cut}} + 1$ and ℓ_{max} are fixed to the best-fit ΛCDM spectrum when calculating S. This division between varying and fixed multipoles speeds up the evaluation of Eq. (3) through the Sherman-Morrison-Woodbury formula and the related matrix determinant lemma, as described in Appendix B.1. This results in an order-of-magnitude speed-up compared to the brute-force computation.

2.2. Low-ℓ temperature map and mask

Next, we consider the various data inputs that are required to evaluate the likelihood in Eq. (3), and we start our discussion with the temperature component. As in 2013, we employ the Commander algorithm for component separation. This is a Bayesian Monte Carlo method that either samples from or maximizes a global posterior defined by some explicit parametric data model and a set of priors. The data model adopted for the *Planck* 2015 analysis is described in detail in Planck Collaboration X (2016), and reads

$$\boldsymbol{s}_\nu(\theta) = g_\nu \sum_{i=1}^{N_{\text{comp}}} \mathsf{F}_\nu^i(\beta_i, \Delta_\nu) \, \boldsymbol{a}_i + \sum_{j=1}^{N_{\text{template}}} T_\nu^j \boldsymbol{b}_j^\nu, \tag{4}$$

where θ denotes the full set of unknown parameters determining the signal at frequency ν. The first sum runs over N_{comp} independent astrophysical components including the CMB itself; $\boldsymbol{a}_i$ is the corresponding amplitude map for each component at some given reference frequency; β_i is a general set of spectral parameters for the same component; g_ν is a multiplicative calibration factor for frequency ν; Δ_ν is a linear correction of the bandpass central frequency; and the function $\mathsf{F}_\nu^i(\beta_i, \Delta_\nu)$ gives the frequency dependence for component i (which can vary pixel-by-pixel and is hence most generally an $N_{\text{pix}} \times N_{\text{pix}}$ matrix). In the second sum,

T_ν^j is one of a set of N_{template} correction template amplitudes, accounting for known effects such as monopole, dipole, or zodiacal light, with template maps $\boldsymbol{b}_j^\nu$.

In 2013, only *Planck* observations between 30 and 353 GHz were employed in the corresponding fit. In the updated analysis, we broaden the frequency range considerably, by including the *Planck* 545 and 857 GHz channels, the 9-year WMAP observations between 23 and 94 GHz (Bennett et al. 2013), and the Haslam et al. (1982) 408 MHz survey. We can then separate the low-frequency foregrounds into separate synchrotron, free-free, and spinning-dust components, as well as to constrain the thermal dust temperature pixel-by-pixel. In addition, in the updated analysis we employ individual detector and detector-set maps rather than co-added frequency maps, and this gives stronger constraints on both line emission (primarily CO) processes and bandpass measurement uncertainties. For a comprehensive discussion of all these results, we refer the interested reader to Planck Collaboration X (2016).

For the purposes of the present paper, the critical output from this process is the maximum-posterior CMB temperature sky map, shown in the top panel of Fig. 1. This map is natively produced at an angular resolution of 1° FWHM, determined by the instrumental beams of the WMAP 23 GHz and 408 MHz frequency channels. In addition, the Commander analysis provides a direct goodness-of-fit measure per pixel in the form of the χ^2 map shown in Planck Collaboration X (2016, Fig. 22). Thresholding this χ^2 map results in a confidence mask that may be used for likelihood analysis, and the corresponding masked region is indicated in the top panel of Fig. 1 by a gray boundary. Both the map and mask are downgraded from their native HEALPix $N_{\text{side}} = 256$ pixel resolution to $N_{\text{side}} = 16$ before insertion into the likelihood code, and the map is additionally smoothed to an effective angular resolution of 440′ FWHM.

The middle panel of Fig. 1 shows the difference between the *Planck* 2015 and 2013 Commander maximum-posterior maps, where the gray region now corresponds to the 2013 confidence mask. Overall, there are large-scale differences at the 10 μK level at high Galactic latitudes, while at low Galactic latitudes there are a non-negligible number of pixels that saturate the colour scale of $\pm 25\,\mu$K. These differences are well understood. First, the most striking red and blue large-scale features at high latitudes are dominated by destriping errors in our 2013 analysis, due to bandpass mismatch in a few frequency channels effectively behaving as correlated noise during map making. As discussed in section 3 of Planck Collaboration X (2016) and illustrated in Fig. 2 therein, the most significant outliers have been removed from the updated 2015 analysis, and, consequently, the pattern is clearly visible from the difference map in Fig. 1. Second, the differences near the Galactic plane and close to the mask boundary are dominated by negative CO residuals near the Fan region, at Galactic coordinates $(l, b) \approx (110°, 20°)$; by negative free-free residuals near the Gum nebula at $(l, b) \approx (260°, 15°)$; and by thermal dust residuals along the plane. Such differences are expected because of the wider frequency coverage and improved foreground model in the new fit. In addition, the updated model also includes the thermal Sunyaev-Zeldovich (SZ) effect near the Coma and Virgo clusters in the northern hemisphere, and this may be seen as a roughly circular patch near the Galactic north pole.

Overall, the additional frequency range provided by the WMAP and 408 MHz observations improves the component separation, and combining these data sets makes more sky effectively available for CMB analysis. The bottom panel of Fig. 1 compares the two χ^2-based confidence masks. In total, 7% of
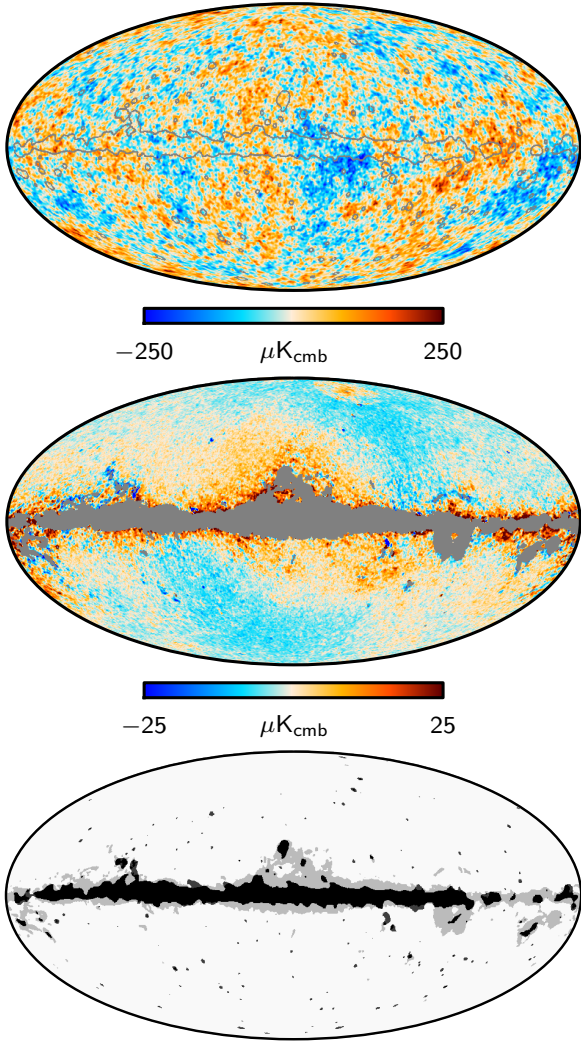
Fig. 1. *Top*: `Commander` CMB temperature map derived from the *Planck* 2015, 9-year WMAP, and 408 MHz Haslam et al. observations, as described in Planck Collaboration X (2016). The gray boundary indicates the 2015 likelihood temperature mask, covering a total of 7% of the sky. The masked area has been filled with a constrained Gaussian realization. *Middle*: difference between the 2015 and 2013 `Commander` temperature maps. The masked region indicates the 2013 likelihood mask, removing 13% of the sky. *Bottom*: comparison of the 2013 (gray) and 2015 (black) temperature likelihood masks.

the sky is removed by the 2015 confidence mask, compared with 13% in the 2013 version.

The top panel in Fig. 2 compares the marginal posterior low-ℓ power spectrum, $D_\ell \equiv C_\ell \ell(\ell + 1)/(2\pi)$, derived from the updated map and mask using the Blackwell-Rao estimator (Chu et al. 2005) with the corresponding 2013 spectrum (Like13). The middle panel shows their difference. The dotted lines indicate the expected variation between the two spectra, σ_ℓ, accounting only for their different sky fractions[3]. From this, we can compute

$$\chi^2 = \sum_{\ell=2}^{29} \left(\frac{D_\ell^{2015} - D_\ell^{2013}}{\sigma_\ell} \right)^2, \tag{5}$$

---

[3] These rms estimates were computed with the `PolSpice` power-spectrum estimator (Chon et al. 2004) by averaging over 1000 noiseless simulations.
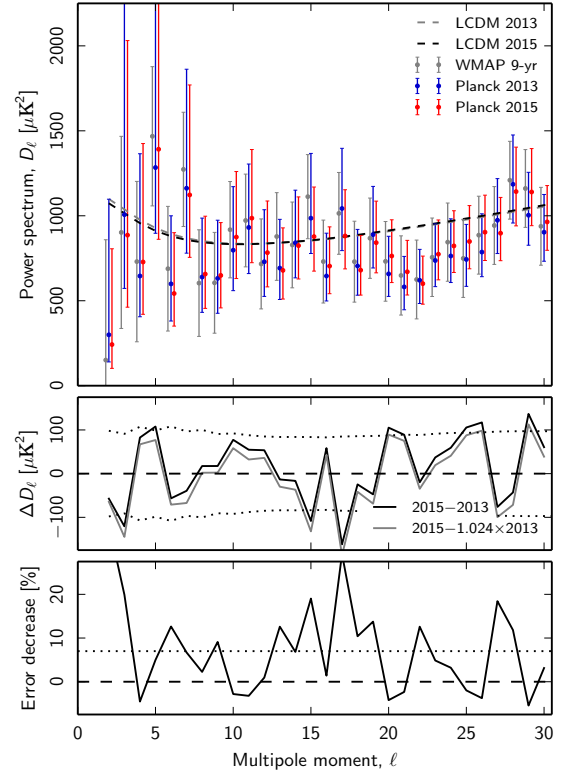

Fig. 2. *Top*: comparison of the *Planck* 2013 (blue points) and 2015 (red points) posterior-maximum low-ℓ temperature power spectra, as derived with `Commander`. Error bars indicate asymmetric marginal posterior 68% confidence regions. For reference, we also show the final 9-year WMAP temperature spectrum in light gray points, as presented by Bennett et al. (2013); note that the error bars indicate symmetric Fisher uncertainties in this case. The dashed lines show the best-fit ΛCDM spectra derived from the respective data sets, including high-multipole and polarization information. *Middle*: difference between the 2015 and 2013 maximum-posterior power spectra (solid black line). The gray shows the same difference after scaling the 2013 spectrum up by 2.4%. Dotted lines indicate the expected $\pm 1\sigma$ confidence region, accounting only for the sky fraction difference. *Bottom*: reduction in marginal error bars between the 2013 and 2015 temperature spectra; see main text for explicit definition. The dotted line shows the reduction expected from increased sky fraction alone.

and we find this to be 21.2 for the current data set. With 28 degrees of freedom, and assuming both Gaussianity and statistical independence between multipoles, this corresponds formally to a probability-to-exceed (PTE) of 82%. According to these tests, the observed differences are consistent with random fluctuations due to increased sky fraction alone.

As discussed in Planck Collaboration I (2016), the absolute calibration of the *Planck* sky maps has been critically reassessed in the new release. The net outcome of this process was an effective recalibration of +1.2% in map domain, or +2.4% in terms of power spectra. The gray line in the middle panel of Fig. 2 shows the same difference as discussed above, but after rescaling the 2013 spectrum up by 2.4%. At the precision offered by these large-scale observations, the difference is small, and either calibration factor is consistent with expectations.

Finally, the bottom panel compares the size of the statistical error bars of the two spectra, in the form of

$$r_\ell \equiv \frac{\left(\sigma_\ell^{\rm l} + \sigma_\ell^{\rm u}\right)\big|_{2013}}{\left(\sigma_\ell^{\rm l} + \sigma_\ell^{\rm u}\right)\big|_{2015}} - 1, \tag{6}$$

where $\sigma_\ell^{\rm u}$ and $\sigma_\ell^{\rm l}$ denote upper and lower asymmetric 68% error bars, respectively. Thus, this quantity measures the decrease in error bars between the 2013 and 2015 spectra, averaged over the upper and lower uncertainties. Averaging over 1000 ideal simulations and multipoles between $\ell = 2$ and 29, we find that the expected change in the error bar due to sky fraction alone is 7%, in good agreement with the real data. Note that because the net uncertainty of a given multipole is dominated by cosmic variance, its magnitude depends on the actual power spectrum value. Thus, multipoles with a positive power difference between 2015 and 2013 tend to have a smaller uncertainty reduction than points with a negative power difference. Indeed, some multipoles have a negative uncertainty reduction because of this effect.

For detailed discussions and higher-order statistical analyses of the new `Commander` CMB temperature map, we refer the interested reader to Planck Collaboration X (2016) and Planck Collaboration XVI (2016).

2.3. 70 GHz polarization low-resolution solution

The likelihood in polarization uses only a subset of the full *Planck* polarization data, chosen to have well-characterized noise properties and negligible contribution from foreground contamination and unaccounted-for systematic errors. Specifically, we use data from the 70 GHz channel of the LFI instrument, for the full mission except for Surveys 2 and 4, which are conservatively removed because they stand as 3σ outliers in survey-based null tests (Planck Collaboration II 2016). While the reason for this behaviour is not completely understood, it is likely related to the fact that these two surveys exhibit the deepest minimum in the dipole modulation amplitude (Planck Collaboration II 2016; Planck Collaboration IV 2016), leading to an increased vulnerability to gain uncertainties and to contamination from diffuse polarized foregrounds.

To account for foreground contamination, the *Planck Q* and *U* 70 GHz maps are cleaned using 30 GHz maps to generate a template for low-frequency foreground contamination, and 353 GHz maps to generate a template for polarized dust emission (Planck Collaboration Int. XIX 2015; Planck Collaboration Int. XXX 2016; Planck Collaboration IX 2016). Linear polarization maps are downgraded from high resolution to $N_{\rm side} = 16$ employing an inverse-noise-weighted averaging procedure, without applying any smoothing (Planck Collaboration VI 2016).

The final cleaned *Q* and *U* maps, shown in Fig. 3, retain a fraction $f_{\rm sky} = 0.46$ of the sky, masking out the Galactic plane and the "spur regions" to the north and south of the Galactic centre.

At multipoles $\ell < 30$, we model the likelihood assuming that the maps follow a Gaussian distribution with known covariance, as in Eq. (3). For polarization, however, we use foreground-cleaned maps, explicitly taking into account the induced increase in variance through an effective noise correlation matrix.

To clean the 70 GHz *Q* and *U* maps we use a template-fitting procedure. Restricting $\boldsymbol{m}$ to the *Q* and *U* maps (i.e., $\boldsymbol{m} \equiv [Q, U]$) we write

$$\boldsymbol{m} = \frac{1}{1 - \alpha - \beta} \left(\boldsymbol{m}_{70} - \alpha \boldsymbol{m}_{30} - \beta \boldsymbol{m}_{353} \right), \qquad (7)$$

where $\boldsymbol{m}_{70}$, $\boldsymbol{m}_{30}$, and $\boldsymbol{m}_{353}$ are bandpass-corrected versions of the 70, 30, and 353 GHz maps (Planck Collaboration III 2016; Planck Collaboration VII 2016), and α and β are the scaling coefficients for synchrotron and dust emission, respectively. The
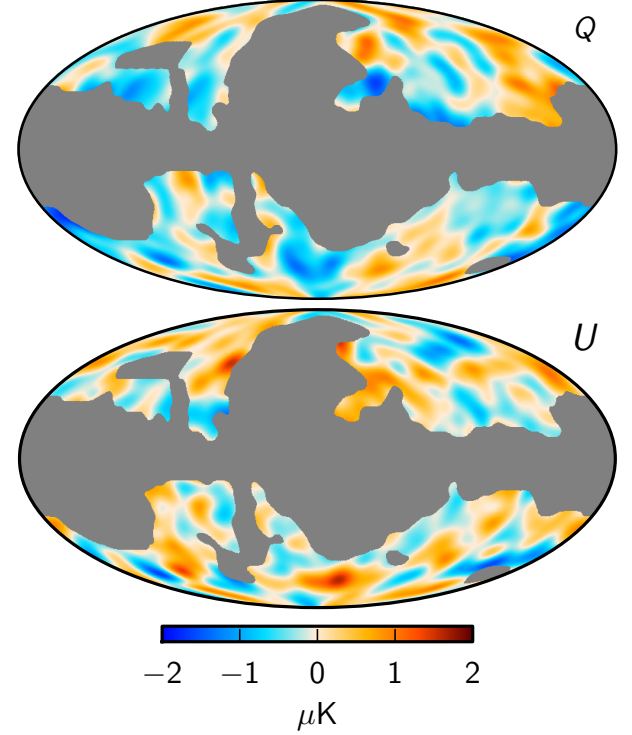

Fig. 3. Foreground-cleaned, 70 GHz *Q* (*top*) and *U* (*bottom*) maps used for the low-ℓ polarization part of the likelihood. Each of the maps covers 46% of the sky.

latter can be estimated by minimizing the quantity

$$\chi^2 = (1 - \alpha - \beta)^2 \boldsymbol{m}^{\sf T} {\sf C}_{\rm S+N}^{-1} \boldsymbol{m}, \qquad (8)$$

where

$${\sf C}_{\rm S+N} \equiv (1 - \alpha - \beta)^2 \langle \boldsymbol{m}\boldsymbol{m}^{\sf T} \rangle = (1 - \alpha - \beta)^2 {\sf S}(C_\ell) + {\sf N}_{70}. \qquad (9)$$

Here ${\sf N}_{70}$ is the pure polarization part of the 70 GHz noise covariance matrix[4] (Planck Collaboration VI 2016), and C_ℓ is taken as the *Planck* 2015 fiducial model (Planck Collaboration XIII 2016). We have verified that using the *Planck* 2013 model has negligible impact on the results described below. Minimization of the quantity in Eq. (8) using the form of the covariance matrix given in Eq. (9) is numerically demanding, since it would require inversion of the covariance matrix at every step of the minimization procedure. However, the signal-to-noise ratio in the 70 GHz maps is relatively low, and we may neglect the dependence on the α and β of the covariance matrix in Eq. (8) using instead:

$${\sf C}_{\rm S+N} = {\sf S}(C_\ell) + {\sf N}_{70}, \qquad (10)$$

so that the matrix needs to be inverted only once. We have verified for a test case that accounting for the dependence on the scaling parameters in the covariance matrix yields consistent results. We find $\alpha = 0.063$ and $\beta = 0.0077$, with 3σ uncertainties $\delta_\alpha \equiv 3\sigma_\alpha = 0.025$ and $\delta_\beta \equiv 3\sigma_\beta = 0.0022$. The best-fit values quoted correspond to a polarization mask using 46% of the sky and correspond to spectral indexes (with 2σ errors) $n_{\rm synch} = -3.16 \pm 0.40$ and $n_{\rm dust} = 1.50 \pm 0.16$, for synchrotron and dust emission respectively (see Planck Collaboration X 2016, for a definition of the foreground spectral indexes). To select

---

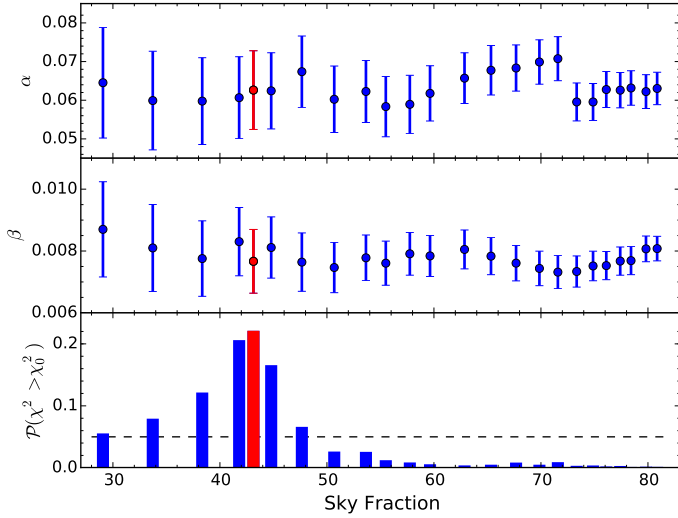[4] We assume here, and have checked in the data, that the noise-induced *TQ* and *TU* correlations are negligible.

Fig. 4. *Upper panels*: estimated best-fit scaling coefficients for synchrotron (α) and dust (β), for several masks, whose sky fractions are displayed along the bottom horizontal axis (see text). *Lower panel*: the probability to exceed, $\mathcal{P}(\chi^2 > \chi_0^2)$. The red symbols identify the mask from which the final scalings are estimated, but note how the latter are roughly stable over the range of sky fractions. Choosing such a large "processing" mask ensures that the associated errors are conservative.

the cosmological analysis mask, the following scheme is employed. We scale to 70 GHz both $\boldsymbol{m}_{30}$ and $\boldsymbol{m}_{353}$, assuming fiducial spectral indexes $n_{synch} = -3.2$ and $n_{dust} = 1.6$, respectively. In this process, we do not include bandpass correction templates. From either rescaled template we compute the polarized intensities $P = \sqrt{Q^2 + U^2}$ and sum them. We clip the resulting template at equally spaced thresholds to generate a set of 24 masks, with unmasked fractions in the range from 30% to 80% of the sky. Finally, for each mask, we estimate the best-fit scalings and evaluate the probability to exceed, $\mathcal{P}(\chi^2 > \chi_0^2)$, where χ_0^2 is the value achieved by minimizing Eq. (8). The $f_{sky} = 43\%$ processing mask is chosen as the tightest mask (i.e., the one with the greatest f_{sky}) satisfying the requirement $\mathcal{P} > 5\%$ (see Fig. 4). We use a slightly smaller mask ($f_{sky} = 46\%$) for the cosmological analysis, which is referred to as the R1.50 mask in what follows.

We define the final polarization noise covariance matrix used in Eq. (3) as

$$N = \frac{1}{(1 - \alpha - \beta)^2} \left(N_{70} + \delta_\alpha^2 \boldsymbol{m}_{30} \boldsymbol{m}_{30}^{\mathsf{T}} + \delta_\beta^2 \boldsymbol{m}_{353} \boldsymbol{m}_{353}^{\mathsf{T}} \right). \quad (11)$$

We use 3σ uncertainties, δ_α and δ_β, to define the covariance matrix, conservatively increasing the errors due to foreground estimation. We have verified that the external (column to row) products involving the foreground templates are sub-dominant corrections. We do not include further correction terms arising from the bandpass leakage error budget since they are completely negligible. Intrinsic noise from the templates also proved negligible.

2.4. Low-ℓ Planck power spectra and parameters

We use the foreground-cleaned Q and U maps derived in the previous section along with the `Commander` temperature map to derive angular power spectra. For the polarization part, we use the noise covariance matrix given in Eq. (11), while assuming only 1 μK^2 diagonal regularization noise for temperature. Consistently, a white noise realization of the corresponding variance
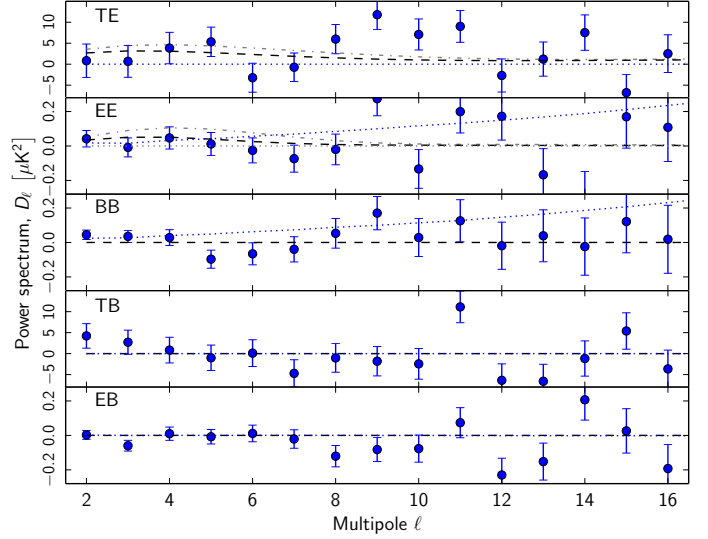

Fig. 5. Polarized QML spectra from foreground-cleaned maps. Shown are the 2013 *Planck* best-fit model ($\tau = 0.089$, dot-dashed) and the 2015 model ($\tau = 0.067$, dashed), as well as the 70 GHz noise bias computed from Eq. (11) (blue dotted).

is added to the `Commander` map. By adding regularization noise, we ensure that the noise covariance matrix is numerically well conditioned.

For power spectra, we employ the `BolPol` code (Gruppuso et al. 2009), an implementation of the quadratic maximum likelihood (QML) power spectrum estimator (Tegmark 1997; Tegmark & de Oliveira-Costa 2001). Figure 5 presents all five polarized power spectra. The errors shown in the plot are derived from the Fisher matrix. In the case of EE and TE we plot the *Planck* 2013 best-fit power spectrum model, which has an optical depth $\tau = 0.089$, as derived from low-ℓ WMAP-9 polarization maps, along with the *Planck* 2015 best model, which has $\tau = 0.067$ as discussed below[5]. Since the EE power spectral amplitude scales with τ as τ^2 (and TE as τ), the 2015 model exhibits a markedly lower reionization bump, which is a better description of *Planck* data. There is a 2.7σ outlier in the EE spectrum at $\ell = 9$, not unexpected given the number of low-ℓ multipole estimates involved.

To estimate cosmological parameters, we couple the machinery described in Sect. 2.1 to `cosmomc`[6] (Lewis & Bridle 2002). We fix all parameters that are not sampled to their *Planck* 2015 ΛCDM best-fit value (Planck Collaboration XIII 2016) and concentrate on those that have the greatest effect at low ℓ: the reionization optical depth τ, the scalar amplitude A_s, and the tensor-to-scalar ratio r. Results are shown in Table 2 for the combinations (τ, A_s) and (τ, A_s, r).

It is interesting to disentangle the cosmological information provided by low-ℓ polarization from that derived from temperature. Low-ℓ temperature mainly contains information on the combination $A_s e^{-2\tau}$, at least at multipoles corresponding to angular scales smaller than the scale subtended by the horizon at reionization (which itself depends on τ). The lowest temperature multipoles, however, are directly sensitive to A_s. On the other hand, large-scale polarization is sensitive to the combination $A_s \tau^2$. Thus, neither low-ℓ temperature nor polarization can separately constrain τ and A_s. Combining temperature and

---

[5] The models considered have been derived by fixing all parameters except τ and A_s to their full multipole range 2015 best-fit values

[6] http://cosmologist.info/cosmomc/

Table 2. Parameters estimated from the low-ℓ likelihood.

Parameter	ΛCDM	ΛCDM+r
τ	0.067 ± 0.023	0.064 ± 0.022
$\log[10^{10}A_{\rm s}]$	2.952 ± 0.055	$2.788^{+0.19}_{-0.09}$
r	0	[0, 0.90]
$z_{\rm re}$	$8.9^{+2.5}_{-2.0}$	$8.5^{+2.5}_{-2.1}$
$10^9 A_{\rm s}$	$1.92^{+0.10}_{-0.12}$	$1.64^{+0.29}_{-0.17}$
$A_{\rm s}{\rm e}^{-2\tau}$	$1.675^{+0.082}_{-0.093}$	$1.45^{+0.24}_{-0.14}$

Notes. For the centre column the set of parameters $(\tau, A_{\rm s})$ was sampled, while it was the set $(\tau, A_{\rm s}, r)$ for the right column. Unsampled parameters are fixed to their ΛCDM 2015 best-fit fiducial values. All errors are 68% CL (confidence level), while the upper limit on r is 95%. The bottom portion of the table shows a few additional derived parameters for information.

polarization breaks the degeneracies and puts tighter constraints on these parameters.

In order to disentangle the temperature and polarization contributions to the constraints, we consider four versions of the low-resolution likelihood.

1. The standard version described above, which considers the full set of T, Q, and U maps, along with their covariance matrix, and is sensitive to the TT, TE, EE, and BB spectra.
2. A temperature-only version, which considers the temperature map and its regularization noise covariance matrix. It is only sensitive to TT.
3. A polarization-only version, considering only the Q and U maps and the QQ, QU, and UU blocks of the covariance matrix. This is sensitive to the EE and BB spectra.
4. A mixed temperature-polarization version, which uses the previous polarization-only likelihood but multiplies it by the temperature-only likelihood. This is different from the standard T, Q, U version in that it assumes vanishing temperature-polarization correlations.

The posteriors derived from these four likelihood versions are displayed in Fig. 6. These plots show how temperature and polarization nicely combine to break the degeneracies and provide joint constraints on the two parameters. The degeneracy directions for cases (2) and (3) are as expected from the discussion above; the degeneracy in case (2) flattens for increasing values of τ because for such values the scale corresponding to the horizon at reionization is pulled forward to $\ell > 30$. By construction, the posterior for case 4 must be equal to the product of the temperature-only (2) and polarization-only (3) posteriors. This is indeed the case at the level of the two-dimensional posterior (see lower right panel of Fig. 6). It is not immediately evident in the one-dimensional distributions because this property does not survive the final marginalization over the non-Gaussian shape of the temperature-only posterior. It is also apparent from Fig. 6 that EE and BB alone do not constrain τ. This is to be expected, and is due to the inverse degeneracy of τ with $A_{\rm s}$, which is almost completely unconstrained without temperature information, and not to the lack of EE signal. By assuming a sharp prior $10^9 A_{\rm s}{\rm e}^{-2\tau} = 1.88$, corresponding to the best estimate obtained when also folding in the high-ℓ temperature information (Planck Collaboration XIII 2016), the polarization-only analysis yields $\tau = 0.051^{+0.022}_{-0.020}$ (red dashed curve in Fig. 6). The latter
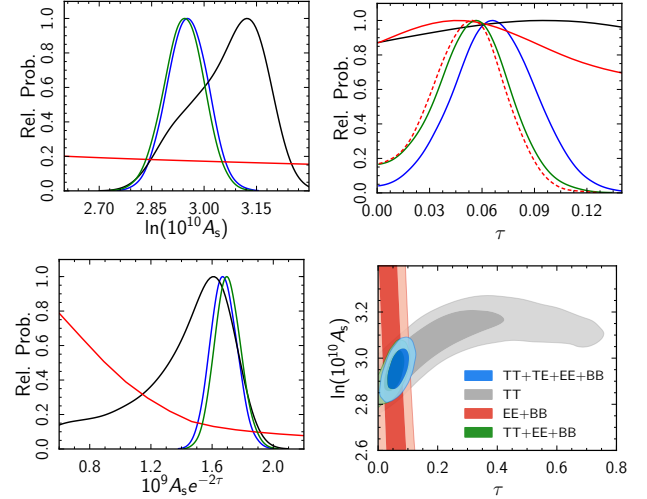
Fig. 6. Likelihoods for parameters from low-ℓ data. *Panels 1–3*: one-dimensional posteriors for $\log[10^{10}A_{\rm s}]$, τ, and $A_{\rm s}{\rm e}^{-2\tau}$ for the several sub-blocks of the likelihood, for cases 1 (blue), 2 (black), 3 (red), and 4 (green) – see text for definitions; dashed red is the same as case 3 but imposes a sharp prior $10^9 A_{\rm s}{\rm e}^{-2\tau} = 1.88$. *Panel 4*: two-dimensional posterior for $\log[10^{10}A_{\rm s}]$ and τ for the same data combinations; shading indicates the 68% and 95% confidence regions.

bound does not differ much from having $A_{\rm s}$ constrained by including TT in the analysis, which yields $\tau = 0.054^{+0.023}_{-0.021}$ (green curves). Finally, the inclusion of non-vanishing temperature-polarization correlations (blue curves) increases the significance of the τ detection at $\tau = 0.067 \pm 0.023$. We have also performed a three-parameter fit, considering τ, $A_{\rm s}$, and r for all four likelihood versions described above, finding consistent results.

2.5. Consistency analysis

Several tests have been carried out to validate the 2015 low-ℓ likelihood. Map-based validation and simple spectral tests are discussed extensively in Planck Collaboration IX (2016) for temperature, and in Planck Collaboration II (2016) for *Planck* 70 GHz polarization. We focus here on tests based on QML and likelihood analyses, respectively employing spectral estimates and cosmological parameters as benchmarks.

We first consider QML spectral estimates C_ℓ derived using `BolPol`. To test their consistency, we consider the following quantity:

$$\chi_{\rm h}^2 = \sum_{\ell=2}^{\ell_{\max}} (C_\ell - C_\ell^{\rm th}) \, {\sf M}_{\ell\ell'}^{-1} \, (C_\ell - C_\ell^{\rm th}), \qquad (12)$$

where ${\sf M}_{\ell\ell'} = \langle (C_\ell - C_\ell^{\rm th})(C_{\ell'} - C_{\ell'}^{\rm th}) \rangle$, $C_\ell^{\rm th}$ represents the fiducial *Planck* 2015 ΛCDM model, and the average is taken over 1000 signal and noise simulations. The latter were generated using the noise covariance matrix given in Eq. (11). We also use the simulations to sample the empirical distribution for $\chi_{\rm h}^2$, considering both $\ell_{\max} = 12$ (shown in Fig. 7, along with the corresponding values obtained from the data) and $\ell_{\max} = 30$, for each of the six CMB polarized spectra. We report in Table 3 the empirical probability of observing a value of $\chi_{\rm h}^2$ greater than for the data (hereafter, PTE). This test supports the hypothesis that the observed polarized spectra are consistent with *Planck*'s best-fit cosmological model and the propagated instrumental uncertainties. We verified that the low PTE values obtained for TE are
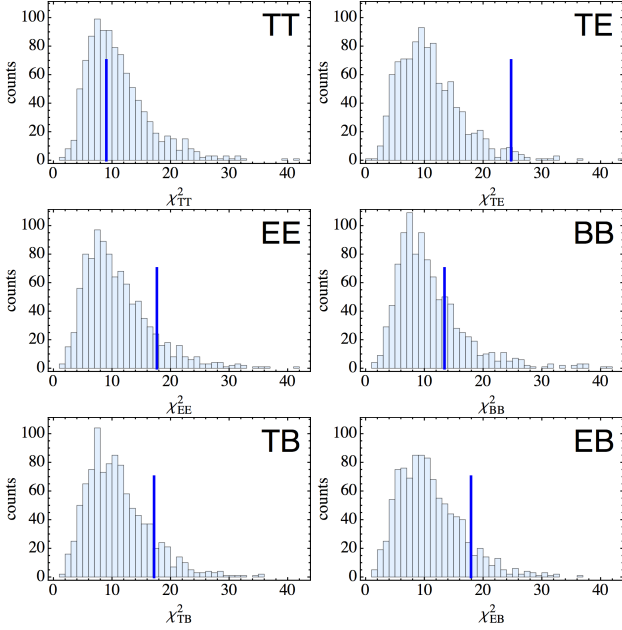
Fig. 7. Empirical distribution of χ_h^2 derived from 1000 simulations, for the case $\ell_{max} = 12$ (see text). Vertical bars reindicate the observed values.

Table 3. Empirical probability of observing a value of χ_h^2 greater than that calculated from the data.

Spectrum	PTE [%]	
	$\ell_{max} = 12$	$\ell_{max} = 30$
TT	57.6	94.2
EE	12.0	50.8
TE	2.2	2.3
BB	24.7	20.6
TB	12.3	35.2
EB	10.2	4.5

related to the unusually high (but not intrinsically anomalous) estimates $9 \leq \ell \leq 11$, a range that does not contribute significantly to constraining τ. For spectra involving *B*, the fiducial model is null, making this, in fact, a null test, probing instrumental characteristics and data processing independent of any cosmological assumptions.

In order to test the likelihood module, we first perform a 45° rotation of the reference frame. This leaves the *T* map unaltered, while sending $Q \rightarrow -U$ and $U \rightarrow Q$ (and, hence, $E \rightarrow -B$ and $B \rightarrow E$). The sub-blocks of the noise covariance matrix are rotated accordingly. We should not be able to detect a τ signal under these circumstances. Results are shown in Fig. 8 for all the full *TQU* and the *TT+EE+BB* sub-block likelihoods presented in the previous section. Indeed, rotating polarization reduces only slightly the constraining power in τ for the *TT+EE+BB* case, suggesting the presence of comparable power in the latter two. On the other hand, τ is not detected at all when rotating the full *T, Q, U* set, which includes *TE* and *TB*. We interpret these results as further evidence that the *TE* signal is relevant for constraining τ, a result that cannot be reproduced by substituting *TB* for *TE*. These findings appear consistent with the visual impression of the low-ℓ spectra of Fig. 5. We have also verified that our results stand when *r* is sampled.
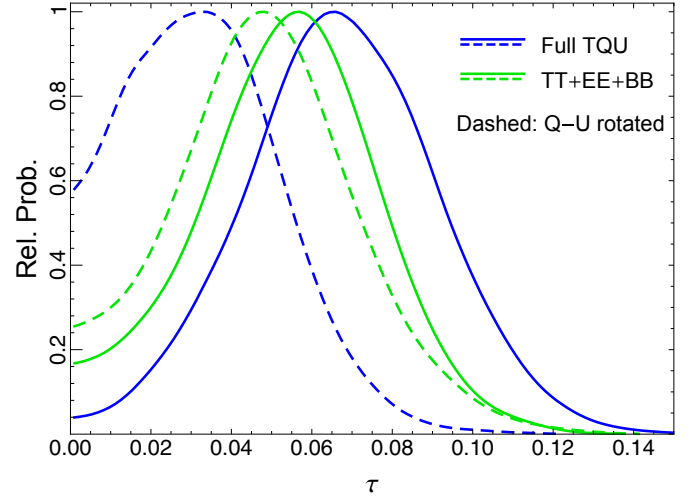

Fig. 8. Posterior for τ for both rotated and unrotated likelihoods. The definition and colour convention of the datasets shown are the same as in the previous section (see Fig. 6), while solid and dashed lines distinguish the unrotated and rotated likelihood, respectively.

As a final test of the 2015 *Planck* low-ℓ likelihood, we perform a full end-to-end Monte Carlo validation of its polarization part. For this, we use 1000 signal and noise full focal plane (FFP8) simulated maps (Planck Collaboration XII 2016), whose resolution has been downgraded to $N_{side} = 16$ using the same procedure as that applied to the data. We make use of a custom-made simulation set for the *Planck* 70 GHz channel, which does not include Surveys 2 and 4. For each simulation, we perform the foreground-subtraction procedure described in Sect. 2.3 above, deriving foreground-cleaned maps and covariance matrices, which we use to feed the low-ℓ likelihood. As above, we sample only $\log[10^{10} A_s]$ and τ, with all other parameters kept to their *Planck* best-fit fiducial values. We consider two sets of polarized foreground simulations, with and without the instrumental bandpass mismatch at 30 and 70 GHz. To emphasize the impact of bandpass mismatch, we do not attempt to correct the polarization maps for bandpass leakage. This choice marks a difference from what is done to real data, where the correction is performed (Planck Collaboration II 2016); thus, the simulations that include the bandpass mismatch effect should be considered as a worst-case scenario. This notwithstanding, the impact of bandpass mismatch on estimated parameters is very small, as shown in Fig. 9 and detailed in Table 4. Even without accounting for bandpass mismatch, the bias is at most 1/10 of the final 1σ error estimated from real data posteriors. The Monte Carlo analysis also enables us to validate the (Bayesian) confidence intervals estimated by cosmomc on data by comparing their empirical counterparts observed from the simulations. We find excellent agreement (see Table 4).

The validation described above only addresses the limited number of instrumental systematic effects that are modelled in the FFP8 simulations, i.e., the bandpass mismatch. Other systematics may in principle affect the measurement of polarization at large angular scales. To address this issue, we have carried out a detailed analysis to quantify the possible impact of LFI-specific instrumental effects in the 70 GHz map (see Planck Collaboration III 2016, for details). Here we just report the main conclusion of that analysis, which estimates the final bias on τ due to all known instrumental systematics to be at most 0.005, i.e., about 0.25σ, well below the final error budget.

Table 4. Statistics for the empirical distribution of estimated cosmological parameters from the FFP8 simulations.

Parameter	Cosmomc best-fit			Cosmomc mean			Standard deviation	
	mean	σ	Δ	mean	σ	Δ	mean	σ
τ	0.0641 ± 0.0007	0.0227	-4.1%	0.0650 ± 0.0006	0.0190	-0.1%	0.0186 ± 0.0001	0.0030
τ^*	0.0665 ± 0.0007	0.0226	$+6.4\%$	0.0672 ± 0.0006	0.0189	$+11.0\%$	0.0185 ± 0.0001	0.0031
$\log[10^{10}A_s]$	3.035 ± 0.002	0.059	-9.4%	3.036 ± 0.002	0.055	-8.0%	0.0535 ± 0.0001	0.0032
$\log[10^{10}A_s^*]$	3.039 ± 0.002	0.059	-1.7%	3.040 ± 0.002	0.056	-0.3%	0.0533 ± 0.0001	0.0033

Notes. Mean and standard deviation for cosmological parameters, computed over the empirical distributions for the estimated best-fit (*left columns*) and mean (*centre columns*) values, as obtained from the FFP8 simulation set. Asterisked parameters flag the presence of (untreated) bandpass mismatch in the simulated maps. The columns labeled Δ give the bias from the input values in units of the empirical standard deviation. This bias always remains small, being at most 0.1σ. Also, note how the empirical standard deviations for the estimated parameters measured from the simulations are very close to the standard errors inferred from `cosmomc` posteriors on real data. The *rightmost columns* show statistics of the standard errors for parameter posteriors, estimated from each `cosmomc` run. The input FFP8 values are $\tau_{input} = 0.0650$ and $\log[10^{10}A_s]_{input} = 3.040$.
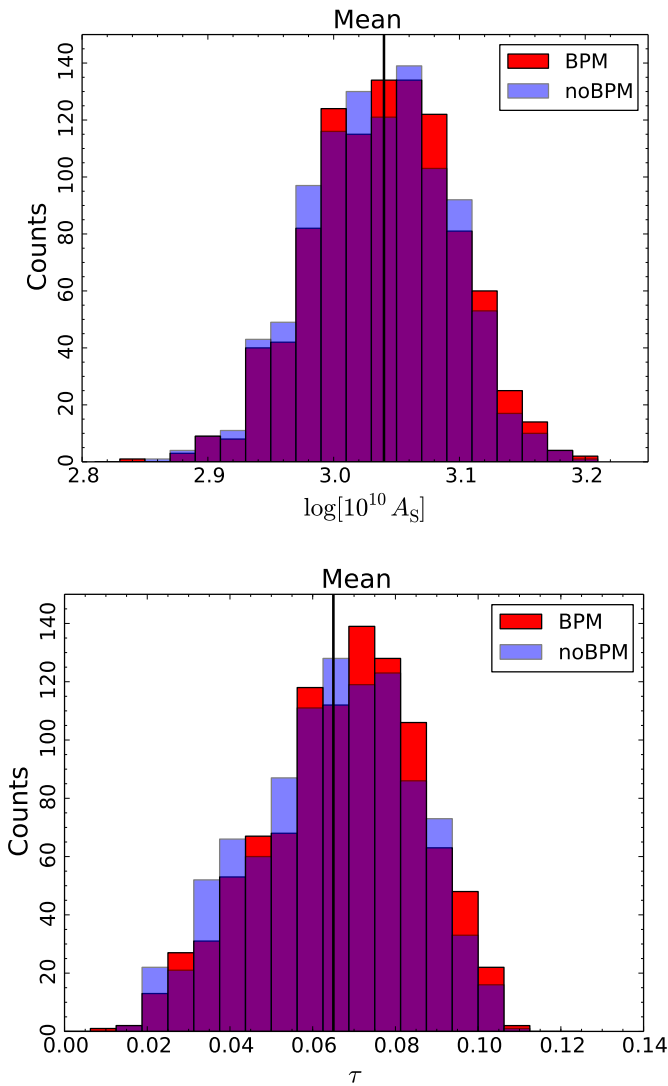
Fig. 9. Empirical distribution of the mean estimated values for $\log[10^{10}A_s]$ (*top*) and τ (*bottom*), derived from 1000 FFP8 simulations (see text). For each simulation, we perform a full end-to-end run, including foreground cleaning and parameter estimation. Blue bars refer to simulations that do not include the instrumental bandpass mismatch, while red bars do. The violet bars flag the overlapping area, while the vertical black lines show the input parameters. We note that the (uncorrected) bandpass mismatch effect hardly changes the estimated parameters.

Table 5. Scalings for synchrotron (α) and dust (β) obtained for WMAP, when WMAP K band and *Planck* 353 GHz data are used as templates.

Band	α	β
Ka	0.3170 ± 0.0016	0.0030 ± 0.0002
Q	0.1684 ± 0.0014	0.0031 ± 0.0003
V	0.0436 ± 0.0017	0.0079 ± 0.0003

2.6. Comparison with WMAP-9 polarization cleaned with Planck 353 GHz

In Like13, we attempted to clean the WMAP-9 low resolution maps using a preliminary version of *Planck* 353 GHz polarization. This resulted in an approximately 1σ shift towards lower values of τ, providing the first evidence based on CMB observations that the WMAP best-fit value for the optical depth may have been biased high. We repeat the analysis here with the 2015 *Planck* products. We employ the procedure described in Bennett et al. (2013), which is similar to that described above for *Planck* 2015. However, in contrast to the *Planck* 70 GHz foreground cleaning, we do not attempt to optimize the foreground mask based on a goodness-of-fit analysis, but stick to the processing and analysis masks made available by the WMAP team. WMAP's P06 mask is significantly smaller than the 70 GHz mask used in the *Planck* likelihood, leaving 73.4% of the sky. Specifically we minimize the quadratic form of Eq. (8), separately for the Ka, Q, and V channels from the WMAP-9 release, but using WMAP-9's own K channel as a synchrotron tracer rather than *Planck* 30 GHz[7]. The purpose of the latter choice is to minimize the differences with respect to WMAP's own analysis. However, unlike the WMAP-9 native likelihood products, which operate at $N_{side} = 8$ in polarization, we use $N_{side} = 16$ in Q and U, for consistency with the *Planck* analysis. The scalings we find are consistent with those from WMAP (Bennett et al. 2013) for α in both Ka and Q. However, we find less good agreement for the higher-frequency V channel, where our scaling is roughly 25% lower than that reported in WMAP's own analysis[8]. We

---

[7] To exactly mimic the procedure followed by the WMAP team, we exclude the signal correlation matrix from the noise component of the χ^2 form. We have checked, however, that the impact of this choice is negligible for WMAP.

[8] There is little point in comparing the scalings obtained for dust, as WMAP employs a model which is not calibrated to physical units.
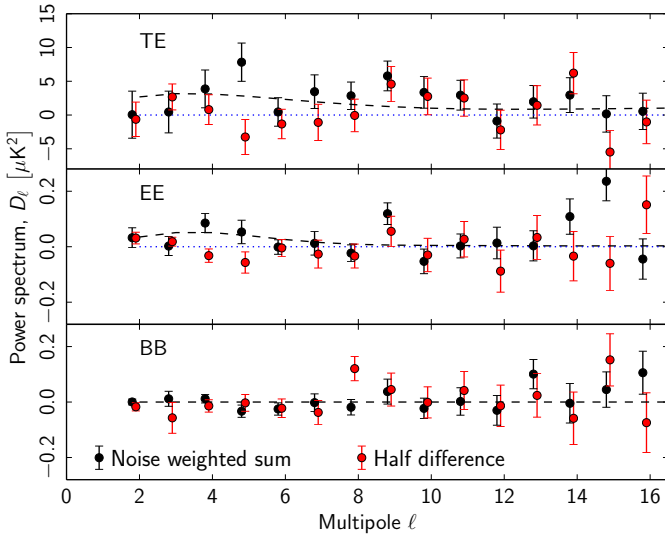
Fig. 10. `BolPol` spectra for the noise-weighted sum (black) and half-difference (red) WMAP and *Planck* combinations. The temperature map employed is always the `Commander` map described in Sect. 2.2 above. The fiducial model shown has $\tau = 0.065$.

combine the three cleaned channels in a noise-weighted average to obtain a three-band map and an associated covariance matrix.

We evaluate the consistency of the low-frequency WMAP and *Planck* 70 GHz low-ℓ maps. Restricting the analysis to the intersection of the WMAP P06 and *Planck* R1.50 masks ($f_{sky} = 45.3\%$), we evaluate half-sum and half-difference Q and U maps. We then compute the quantity $\chi^2_{sd} = \boldsymbol{m}^T N^{-1} \boldsymbol{m}$ where $\boldsymbol{m}$ is either the half-sum or the half-difference $[Q, U]$ combination and N is the corresponding noise covariance matrix. Assuming that χ^2_{sd} is χ^2 distributed with 2786 degrees of freedom we find a PTE($\chi^2 > \chi^2_{sd}$) = 1.3×10^{-5} (reduced $\chi^2 = 1.116$) for the half-sum, and PTE = 0.84 (reduced $\chi^2 = 0.973$) for the half-difference. This strongly suggests that the latter is consistent with the assumed noise, and that the common signal present in the half-sum map is wiped out in the difference.

We also produce noise-weighted sums of the low-frequency WMAP and *Planck* 70 GHz low-resolution Q and U maps, evaluated in the union of the WMAP P06 and *Planck* R1.50 masks ($f_{sky} = 73.8\%$). We compute `BolPol` spectra for the noise-weighted sum and half-difference combinations. These *EE*, *TE*, and *BB* spectra are shown in Fig. 10 and are evaluated in the intersection of the P06 and R1.50 masks. The spectra also support the hypothesis that there is a common signal between the two experiments in the typical multipole range of the reionization bump. In fact, considering multipoles up to $\ell_{max} = 12$ we find an empirical PTE for the spectra of the half-difference map of 6.8% for *EE* and 9.5% for *TE*, derived from the analysis of 10000 simulated noise maps. Under the same hypothesis, but considering the noise-weighted sum, the PTE for *EE* drops to 0.8%, while that for *TE* is below the resolution allowed by the simulation set (PTE < 0.1%). The *BB* spectrum, on the other hand, is compatible with a null signal in both the noise-weighted sum map (PTE = 47.5%) and the half-difference map (PTE = 36.6%).

We use the *Planck* and WMAP map combinations to perform parameter estimates from low-ℓ data only. We show here results from sampling $\log[10^{10}A_s]$, τ, and the tensor-to-scalar ratio r, with all other parameters kept to the *Planck* 2015 best fit (the case with $r = 0$ produces similar results). Figure 11 shows the posterior probability for τ for several *Planck* and WMAP
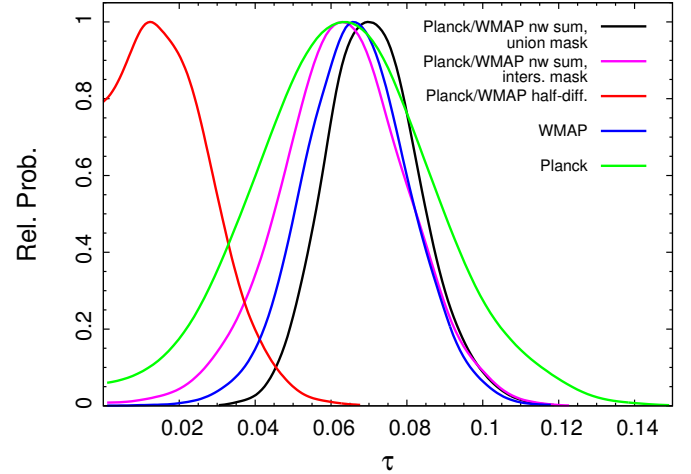

Fig. 11. Posterior probabilities for τ from the WMAP (cleaned with *Planck* 353 GHz as a dust template) and *Planck* combinations listed in the legend. Results are presented for the noise-weighted sum both in the union and the intersection of the two analysis masks. The half-difference map is consistent with a null detection, as expected.

Table 6. Selected parameters estimated from the low-ℓ likelihood, for *Planck*, WMAP and their noise-weighted combination.

Parameter	*Planck*	WMAP	*Planck*/WMAP
τ	$0.064^{+0.022}_{-0.023}$	$0.067^{+0.013}_{-0.013}$	$0.071^{+0.011}_{-0.013}$
z_{re}	$8.5^{+2.5}_{-2.1}$	$8.9^{+1.3}_{-1.3}$	$9.3^{+1.1}_{-1.1}$
$\log[10^{10}A_s]$	$2.79^{+0.19}_{-0.09}$	$2.87^{+0.11}_{-0.06}$	$2.88^{+0.10}_{-0.06}$
r	$[0, 0.90]$	$[0, 0.52]$	$[0, 0.48]$
$A_s e^{-2\tau}$	$1.45^{+0.24}_{-0.14}$	$1.55^{+0.16}_{-0.10}$	$1.55^{+0.14}_{-0.11}$

Notes. The temperature map used is always *Planck* `Commander`. Only $\log[10^{10}A_s]$, τ, and r are sampled. The other ΛCDM parameters are kept fixed to the *Planck* 2015 fiducial. The likelihood for the noise-weighted combination is evaluated in the union of the WMAP P06 and Planck R1.50 masks.

combinations. They are all consistent, except the *Planck* and WMAP half-difference case, which yields a null detection for τ – as it should. As above, we always employ the `Commander` map in temperature. Table 6 gives the mean values for the sampled parameters, and for the derived parameters z_{re} (mean redshift of reionization) and $A_s e^{-2\tau}$. Results from a joint analysis of the WMAP-based low-ℓ polarization likelihoods presented here and the *Planck* high-ℓ likelihood are discussed in Sect. 5.7.1.

3. High-multipole likelihood

At high multipoles ($\ell > 29$), as in Like13, we use a likelihood function based on pseudo-C_ℓs calculated from *Planck* HFI data, as well as further parameters describing the contribution of foreground astrophysical emission and instrumental effects (e.g., calibration, beams). Aside from the data themselves, the main advances over 2013 include the use of high-ℓ polarization information along with more detailed models of foregrounds and instrumental effects.

Section 3.1 introduces the high-ℓ statistical description, Sect. 3.2 describes the data we use, Sects. 3.3 and 3.4 describe foreground and instrumental modelling, and Sect. 3.5 describes

the covariance matrix between multipoles and spectra. Section 3.6 validates the overall approach on realistic simulations, while Sect. 3.7 addresses the question of the potential impact of low-level instrumental systematics imperfectly corrected by the DPC processing. The reference results generated with the high multipole likelihood are described in Sect. 3.8. A detailed assessment of these results is presented in Sect. 4.

3.1. Statistical description

Assuming a Gaussian distribution for the CMB temperature anisotropies and polarization, all of the statistical information contained in the *Planck* maps can be compressed into the likelihood of the temperature and polarization auto- and cross-power spectra. In the case of a perfect CMB observation of the full sky (with spatially uniform noise and isotropic beam-smearing), we know the joint distribution of the empirical temperature and polarization power spectra and can build an exact likelihood, which takes the simple form of an inverse Wishart distribution, uncorrelated between multipoles. For a single power spectrum (i.e., ignoring polarization and temperature cross-spectra between detectors) the likelihood for each multipole ℓ simplifies to an inverse χ^2 distribution with $2\ell + 1$ degrees of freedom. At high enough ℓ, the central limit theorem ensures that the shape of the likelihood is very close to that of a Gaussian distributed variable. This remains true for the inverse Wishart generalization to multiple spectra, where, for each ℓ, the shape of the joint spectra and cross-spectra likelihood approaches that of a correlated Gaussian (Hamimeche & Lewis 2008; Elsner & Wandelt 2012). In the simple full-sky case, the correlations are easy to compute (Hamimeche & Lewis 2008), and only depend on the theoretical CMB *TT*, *TE*, and *EE* spectra. For small excursions around a fiducial cosmology, as is the case here given the constraining power of the *Planck* data, one can show that computing the covariance matrix at a fiducial model is sufficient (Hamimeche & Lewis 2008).

The data, however, differ from the idealized case. In particular, foreground astrophysical processes contribute to the temperature and polarization maps. As we see in Sect. 3.3, the main foregrounds in the frequency range we use are emission from dust in our Galaxy, the clustered and Poisson contributions from the cosmic infrared background (CIB), and radio point sources. Depending on the scale and frequency, foreground emission can be a significant contribution to the data, or even exceed the CMB. This is particularly true for dust near the Galactic plane, and for the strongest point sources. We excise the most contaminated regions of the sky (see Sect. 3.2.2). The remaining foreground contamination is taken into account in our model, using the fact that CMB and foregrounds have different emission laws; this enables them to be separated while estimating parameters.

Foregrounds also violate the Gaussian approximation assumed above. The dust distribution, in particular, is clearly non-Gaussian. Following Like13, however, we assume that outside the masked regions we can neglect non-Gaussian features and assume that, as for the CMB, all the relevant statistical information about the foregrounds is encoded in the spatial power spectra. This assumption is verified to be sufficient for our purposes in Sect. 3.6, where we assess the accuracy of the cosmological parameter constraints in realistic Monte Carlo simulations that include data-based (non-Gaussian) foregrounds.

Cutting out the foreground-contaminated regions from our maps biases the empirical power spectrum estimates. We de-bias

them using the `PolSpice`[9] algorithm (Chon et al. 2004) and, following Like13, we take the correlation between multipoles induced by the mask and de-biasing into account when computing our covariance matrix. The masked-sky covariance matrix is computed using the equations in Like13, which are extended to the case of polarization in Appendix C.1.1. Those equations also take into account the inhomogeneous distribution of coloured noise on the sky using a heuristic approach. The approximation of the covariance matrix that can be obtained from those equations is only valid for some specific mask properties, and for high enough multipoles. In particular, as discussed in Appendix C.1.4, correlations induced by point sources cannot be faithfully described in our approximation. Similarly, Monte Carlo simulations have shown that our analytic approximation loses accuracy around $\ell = 30$. We correct for both of those effects using empirical estimates from Monte Carlo simulations. The computation of the covariance matrix requires knowledge of both the CMB and foreground power spectra, as well as the map characteristics (beams, noise, sky coverage). The CMB and foreground power spectra are obtained iteratively from previous, less accurate versions of the likelihood.

At this stage, we would thus construct our likelihood approximation by compressing all of the individual *Planck* detector data into mask-corrected (pseudo-) cross-spectra, and build a grand likelihood using these spectra and the corresponding analytical covariance matrix:

$$-\ln \mathcal{L}(\hat{\boldsymbol{C}}|\boldsymbol{C}(\theta)) = \frac{1}{2}\left[\hat{\boldsymbol{C}} - \boldsymbol{C}(\theta)\right]^{\mathsf{T}} \mathsf{C}^{-1}\left[\hat{\boldsymbol{C}} - \boldsymbol{C}(\theta)\right] + \text{const.}, \quad (13)$$

where $\hat{\boldsymbol{C}}$ is the data vector, $\boldsymbol{C}(\theta)$ is the model with parameters θ, and C is the covariance matrix. This formalism enables us to separately marginalize over or condition upon different components of the model vector, separately treating cases such as individual frequency-dependent spectra, or temperature and polarization spectra. Obviously, *Planck* maps at different frequencies have different constraining powers on the underlying CMB, and following Like13 we use this to impose and assess various cuts to keep only the most relevant data.

We therefore consider only the three best CMB *Planck* channels, i.e., 100 GHz, 143 GHz, and 217 GHz, in the multipole range where they have significant CMB contributions and low enough foreground contamination after masking; we therefore did not directly include the adjacent channels at 70 GHz and 350 GHz in the analysis. In particular, including the 70 GHz data would not bring much at large scales where the results are already cosmic variance limited, and would entail additional complexity in foreground modelling (synchrotron at large scales, additional radio sources excisions at small scales). The cuts in multipole ranges is be described in detail in Sect. 3.2.4. Further, in order to achieve a significant reduction in the covariance matrix size (and computation time), we compress the data vector (and accordingly the covariance matrix), both by co-adding the individual detectors for each frequency and by binning the combined power spectra. We also co-add the two different *TE* and *ET* inter-frequency cross-spectra into a single *TE* spectrum for each pair of frequencies. This compression is lossless in the case without foregrounds. The exact content of the data vector is discussed in Sect. 3.2.

---

[9] http://www2.iap.fr/users/hivon/software/PolSpice/

The model vector $\boldsymbol{C}(\theta)$ must represent the content of the data vector. It can be written schematically as

$$C^{XY}_{\nu\times\nu'}\big|_\ell(\theta) = M^{XY}_{ZW,\nu\times\nu'}\big|_\ell(\theta_{\rm inst})\ C^{ZW,\rm sky}_{\nu\times\nu'}\big|_\ell(\theta) + N^{XY}_{\nu\times\nu'}\big|_\ell(\theta_{\rm inst}),$$

$$C^{ZW,\rm sky}_{\nu\times\nu'}\big|_\ell(\theta) = C^{ZW,\rm cmb}\big|_\ell(\theta) + C^{ZW,\rm fg}_{\nu\times\nu'}\big|_\ell(\theta), \tag{14}$$

where $C^{XY}_{\nu\times\nu'}\big|_\ell(\theta)$ is the element of the model vector corresponding to the multipole ℓ of the XY cross-spectra (X and Y being either T or E) between the pair of frequencies ν and ν'. This element of the model originates from the sum of the microwave emission of the sky, i.e., the CMB ($C^{ZW,\rm cmb}\big|_\ell(\theta)$) which does not depend of the pair of frequencies (all maps are in units of $K_{\rm cmb}$), and foreground ($C^{ZW,\rm fg}_{\nu\times\nu'}\big|_\ell(\theta)$). Section 3.3 describes the foreground modelling. The mixing matrix $M^{XY}_{ZW,\nu\times\nu'}\big|_\ell(\theta_{\rm inst})$ accounts for imperfect calibration, imperfect beam correction, and possible leakage between temperature and polarization. It does depend on the pair of frequencies and can depend on the multipole[10] when accounting for imperfect beams and leakages. Finally, the noise term $N^{XY}_{\nu\times\nu'}\big|_\ell(\theta_{\rm inst})$ accounts for the possible correlated noise in the XY cross-spectra for the pair of frequencies $\nu\times\nu'$. Sections 3.2.3 and 3.4 describe our instrument model.

3.2. Data

The data vector $\hat{C}$ in the likelihood equation (Eq. (13)) is constructed from concatenated temperature and polarization components,

$$\hat{C} = \left(\hat{C}^{TT}, \hat{C}^{EE}, \hat{C}^{TE}\right), \tag{15}$$

which, in turn, comprise the following frequency-averaged spectra:

$$\hat{C}^{TT} = \left(\hat{C}^{TT}_{100\times100}, \hat{C}^{TT}_{143\times143}, \hat{C}^{TT}_{143\times217}, \hat{C}^{TT}_{217\times217}\right) \tag{16}$$

$$\hat{C}^{EE} = \left(\hat{C}^{EE}_{100\times100}, \hat{C}^{EE}_{100\times143}, \hat{C}^{EE}_{100\times217}, \hat{C}^{EE}_{143\times143}, \hat{C}^{EE}_{143\times217}, \hat{C}^{EE}_{217\times217}\right) \tag{17}$$

$$\hat{C}^{TE} = \left(\hat{C}^{TE}_{100\times100}, \hat{C}^{TE}_{100\times143}, \hat{C}^{TE}_{100\times217}, \hat{C}^{TE}_{143\times143}, \hat{C}^{TE}_{143\times217}, \hat{C}^{TE}_{217\times217}\right). \tag{18}$$

The TT data selection is very similar to Like13. We still discard the 100×143 and 100×217 cross-spectra in their entirety. They contain little extra information about the CMB, as they are strongly correlated with the high S/N maps at 143 and 217 GHz. Including them, in fact, would only give information about the foreground contributions in these cross-spectra, at the expense of a larger covariance matrix with increased condition number. In TE and EE, however, the situation is different since the overall S/N is significantly lower for all spectra, so a foreground model of comparatively low complexity can be used and it is beneficial to retain all the available cross-spectra.

We obtain cross power spectra at the frequencies $\nu\times\nu'$ using weighted averages of the individual beam-deconvolved, mask-corrected half-mission (HM) map power spectra,

$$\hat{C}^{XY}_{\nu\times\nu'}\big|_\ell = \sum_{(i,j)\in(\nu,\nu')} w^{XY}_{i,j}\big|_\ell \times \hat{C}^{XY}_{i,j}\big|_\ell, \tag{19}$$

---

[10] We assume an ℓ-diagonal mixing matrix here. This is not necessarily the case, as sub-pixel beam effects, for example, can induce mode couplings. As discussed in Sect. 3.4.3, those were estimated in Like13 and found to be negligible for temperature. They are not investigated further in this paper.

Table 7. Detector sets used to make the maps for this analysis.

Set	ν [GHz]	Type	Detectors	*FWHM*
100-ds0	100	PSB	8 detectors	9.′68
100-ds1	100	PSB	1a+1b + 4a+4b	
100-ds2	100	PSB	2a+2b + 3a+3b	
143-ds0	143	MIX	11 detectors	7.′30
143-ds1	143	PSB	1a+1b + 3a+3b	
143-ds2	143	PSB	2a+2b + 4a+4b	
143-ds3	143	SWB	143-5	
143-ds4	143	SWB	143-6	
143-ds5	143	SWB	143-7	
217-ds0	217	MIX	12 detectors	5.′02
217-ds1	217	PSB	5a+5b + 7a+7b	
217-ds2	217	PSB	6a+6b + 8a+8b	
217-ds3	217	SWB	217-1	
217-ds4	217	SWB	217-2	
217-ds5	217	SWB	217-3	
217-ds6	217	SWB	217-4	
353-ds0	353	MIX	12 detectors	4.′94
545-ds0	545	SWB	3 detectors	4.′83

Notes. SWBs may be used individually; PSBs are used in pairs (denoted a and b), and we consider only the maps estimated from two pairs of PSBs. The FWHM quoted here correspond to a Gaussian whose solid angle is equivalent to that of the effective beam; see Planck Collaboration VIII (2016) for details.

where $XY \in \{TT, TE, EE\}$, and $w^{XY}_{i,j}\big|_\ell$ is the multipole-dependent inverse-variance weight for the detector-set map combination (i, j), derived from its covariance matrix (see Sect. 3.5). For $XY = TE$, we further add the ET power spectra of the same frequency combination to the sum of Eq. (19); i.e., the average includes the correlation of temperature information from detector-set i and polarization information of detector-set j and vice versa.

We construct the *Planck* high-multipole likelihood solely from the HFI channels at 100, 143, and 217 GHz. These perform best as they have high S/N combined with manageably low foreground contamination. As in Like13, we only employ 70 GHz LFI data for cross-checks (in the high-ℓ regime), while the HFI 353 GHz and 545 GHz maps are used to determine the dust model.

3.2.1. Detector combinations

Table 7 summarizes the main characteristics of individual HFI detector sets used in the construction of the likelihood function. As discussed in Sect. 3.1, the likelihood does not use the cross-spectra from individual detector-set maps; instead, we first combine all those contributing at each frequency to form weighted averages. As in 2013, we disregard all auto-power-spectra as the precision required to remove their noise bias is difficult to attain and even small residuals may hamper a robust inference of cosmological parameters (Like13).

In 2015, the additional data available from full-mission observations enables us to construct nearly independent full-sky maps from the first and the second halves of the mission duration. We constructed cross-spectra by cross-correlating the two half-mission maps, ignoring the half-mission auto-spectra at the

Table 8. Masks used for the high-ℓ analysis.

Frequency [GHz]	Mask	
	Temperature	Polarization
100	T66	P70
143	T57	P50
217	T47	P41

Notes. Temperature and polarization masks used in the likelihood are identified by T and P, followed by two digits that specify the retained sky fraction (percent). As discussed in Appendix A, T masks are derived by merging apodized Galactic, CO, and extragalactic sources masks. P masks, instead, are simply given by apodized Galactic masks.

expense of a very small increase in the uncertainties. This differs from the procedure used in 2013, when we estimated cross-spectra between detectors or detector-sets, and has the advantage of minimizing possible contributions from systematic effects that are correlated in the time domain.

The main motivation for this change from 2013 is that the correlated noise between detectors (at the same or different frequencies) is no longer small enough to be neglected (see Sect. 3.4.4). And while the correction for the "feature" around $\ell = 1800$, which was (correctly) attributed to residual ^{4}He-JT cooler lines in 2013 (Planck Collaboration VI 2014), has been improved in the 2015 TOI processing pipeline (Planck Collaboration VII 2016), cross-spectra between the two half-mission periods can help to suppress time-dependent systematics, as argued by Spergel et al. (2015). Still, in order to enable further consistency checks, we also build a likelihood based on cross-spectra between full-mission detector-set maps, applying a correction for the effect of correlated noise. The result illustrates that not much sensitivity is lost with half-mission cross-spectra (see the whisker labelled "DS" in Figs. 35, 36, and C.10).

3.2.2. Masks

Temperature and polarization masks are used to discard areas of the sky that are strongly contaminated by foreground emission. The choice of masks is a trade-off between maximizing the sky coverage to minimize sample variance, and the complexity and potentially insufficient accuracy of the foreground model needed in order to deal with regions of stronger foreground emission. The masks combine a Galactic mask, excluding mostly low Galactic-latitude regions, and a point-source mask. We aim to maximize the sky fraction with demonstrably robust results (see Sect. 4.1.2 for such a test).

Temperature masks are obtained by merging the apodized Galactic, CO, and point-source masks described in Appendix A. In polarization, as discussed in Planck Collaboration Int. XXX (2016), even at 100 GHz foregrounds are dominated by the dust emission, so for polarization analysis we employ the same apodized Galactic masks as we use for temperature, because they are also effective in reducing fluctuations in polarized dust emission at the relatively small scales covered by the high-ℓ likelihood (contrary to the large Galactic scales), but we do not include a compact-source mask because polarized emission from extragalactic foregrounds is negligible at the frequencies of interest (Naess et al. 2014; Crites et al. 2015).

Table 8 lists the masks used in the likelihood at each frequency channel. We refer throughout to the masks by explicitly indicating the percentage of the sky they retain: T66, T57, T47 for temperature and P70, P50, P41 for polarization. G70, G60,
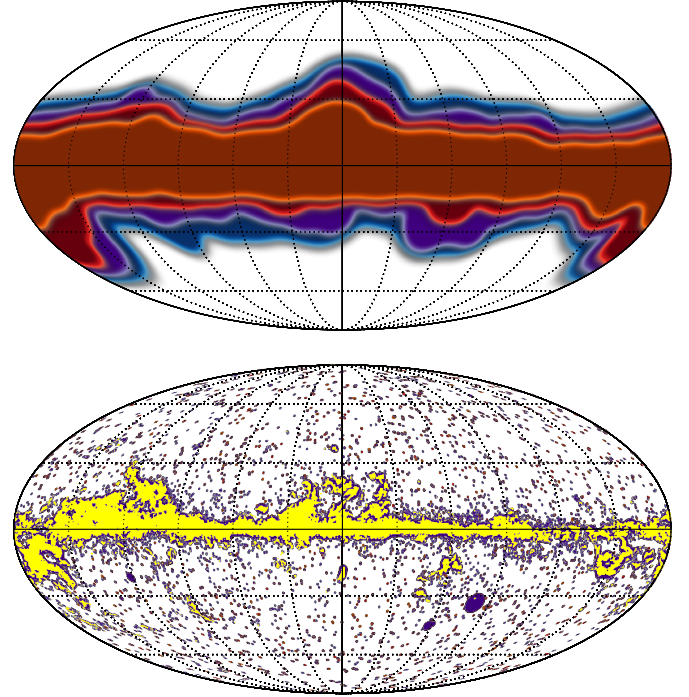
Fig. 12. *Top*: apodized Galactic masks: G41 (blue), G50 (purple), G60 (red), and G70 (orange); these are identical to the polarization masks P41 (used at 217 GHz), P50 (143 GHz), P70 (100 GHz). *Bottom*: extragalactic-object masks for 217 GHz (purple), 143 GHz (red), and 100 GHz (orange); the CO mask is shown in yellow.

G50, and G41 denote the apodized Galactic masks. As noted above, the apodized P70, P50, and P41 polarization masks are identical to the G70, G50, and G41 Galactic masks.

The Galactic masks are obtained by thresholding the smoothed, CMB-cleaned 353 GHz map at different levels to obtain different sky coverage. All of the Galactic masks are apodized with a $4°71$ FWHM ($\sigma = 2°$) Gaussian window function to localize the mask power in multipole space. In order to adapt to the different relative strengths of signal, noise, and foregrounds, we use different sky coverage for temperature and polarization, ranging in effective sky fraction from 41% to 70% depending on the frequency. The Galactic masks are shown in Fig. 12.

For temperature we use the G70, G60, and G50 Galactic masks at (respectively) 100 GHz, 143 GHz, and 217 GHz. For the first release of *Planck* cosmological data (Planck Collaboration XI 2016) we made more conservative choices of masks than in this paper ($f_{sky} = 49\%, 31\%$, and 31% at, respectively, 100, 143, and 217 GHz, to be compared to $f_{sky} = 66\%, 57\%$, and 47%). Admitting more sky into the analysis requires a thorough assessment of the robustness of the foreground modelling, and in particular of the Galactic dust model (see Sect. 3.3). When retaining more sky close to the Galactic plane at 100 GHz, maps start to show contamination by CO emission that also needs to be masked. This was not the case in the *Planck* 2013 analysis. We therefore build a CO mask as described in Appendix A. Once we apply this mask, the residual foreground at 100 GHz is consistent with dust and there is no evidence for other anisotropic foreground components, as shown by the double-difference spectra between the 100 GHz band and the 143 GHz band where there is no CO line (Sect. 3.3.1). We also use the CO mask at 217 GHz, although we expect it to have a smaller impact since at this frequency CO emission is fainter
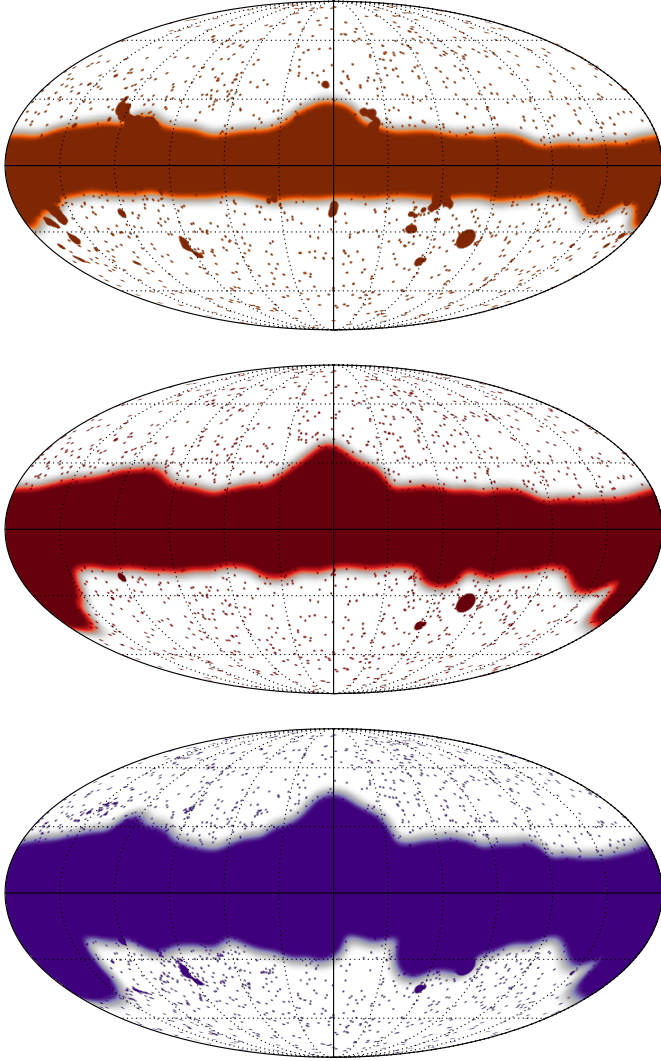
Fig. 13. *Top to bottom*: temperature masks for 100 GHz (T66), 143 GHz (T57) and 217 GHz (T47). The colour scheme is the same as in Fig. 12.

and the applied Galactic cut wider. The extragalactic "point" source masks in fact include both point sources and extended objects; they are used only with the temperature maps. Unlike in 2013, we use a different source mask for each frequency, taking into account different source selection and beam sizes (see Appendix A). Both the CO and the extragalactic object masks are apodized with a 30′ FWHM Gaussian window function. The different extragalactic masks, as well as the CO mask, are shown in Fig. 12. The resulting mask combinations for temperature are shown in Fig. 13.

3.2.3. Beam and transfer functions

The response to a point source is given by the combination of the optical response of the *Planck* telescope and feed-horns (the optical beam) with the detector time response and electronic transfer function (whose effects are partially removed during the TOI processing). This response pattern is referred to as the "scanning beam". It is measured on planet transits (Planck Collaboration VIII 2016). However, the value in any pixel resulting from the map-making operation comes from a sum over many different elements of the timeline, each of which has hit the pixel in a different location and

from a different direction. Furthermore, combined maps are weighted sums of individual detectors. All of these result in an "effective beam" window function encoding the multiplicative effect on the angular power spectrum. We note that beam non-circularity and the non-uniform scanning of the sky create differences between auto- and cross-detector beam window functions (Planck Collaboration VII 2014).

In the likelihood analysis, we correct for this by using the effective beam window function corresponding to each specific spectrum; the window functions are calculated with the QuickBeam pipeline, except for one of the alternative analyses (Xfaster) which relied on the FEBeCoP window functions (see Planck Collaboration VII 2016; Planck Collaboration VII 2014, and references therein for details of these two codes). In Sect. 3.4.3 we discuss the model of their uncertainties.

3.2.4. Multipole range

Following the approach taken in Like13, we use specifically tailored multipole ranges for each frequency-pair spectrum. In general, we exclude multipoles where either the S/N is too low for the data to contribute significant constraints on the CMB, or the level of foreground contamination is so high that the foreground contribution to the power spectra cannot be modelled sufficiently accurately; high foreground contamination would also require us to consider possible non-Gaussian terms in the estimation of the likelihood covariance matrix. We impose the same ℓ cuts for the detector-set and half-mission likelihoods for comparison, and we exclude the $\ell > 1200$ range for the 100×100 spectra, where the correlated noise correction is rather uncertain.

Figure 14 shows the unbinned S/N per frequency for TT, EE, and TE, where the signal is given by the frequency-dependent CMB and foreground power spectra, while the noise term contains contributions from cosmic variance and instrumental noise and is given by the diagonal elements of the power-spectrum covariance matrix. The figure also shows the S/N assuming only cosmic variance (CV) in the noise term, obtained either by a full calculation of the covariance matrix with instrumental noise set to zero, or using the approximation

$$\sigma_{CV}^{\{TT,EE\}} = \sqrt{\left(\frac{2}{(2\ell+1)f_{\text{sky}}}\right)\left(C_\ell^{\{TT,EE\}}\right)^2}$$

$$\sigma_{CV}^{TE} = \sqrt{\left(\frac{2}{(2\ell+1)f_{\text{sky}}}\right)\frac{\left(C_\ell^{TE}\right)^2 + C_\ell^{TT}C_\ell^{EE}}{2}}. \tag{20}$$

(see e.g. Percival & Brown 2006).

This figure illustrates that the multipole cuts we apply ensure that the $|S/N| \gtrsim 1$. The TT multipole cuts are similar to those adopted in Like13. While otherwise similar to the 2013 likelihood, the revised treatment of dust in the foreground model enables the retention of multipoles $\ell < 500$ of the 143×217 and 217×217 GHz TT spectra. As discussed in detail in Sect. 3.3.1, we are now marginalizing over a free amplitude parameter of the dust template, which was held constant for the 2013 release. Furthermore, the greater sky coverage at 100 GHz maximizes its weight at low ℓ, so that the best estimate of the CMB signal on large scales is dominated by 100 GHz data. We do not detect noticeable parameter shifts when removing or including multipoles at $\ell < 500$. See Sect. 4.1 for an in-depth analysis of the impact of different choices of multipole ranges on cosmological parameters.
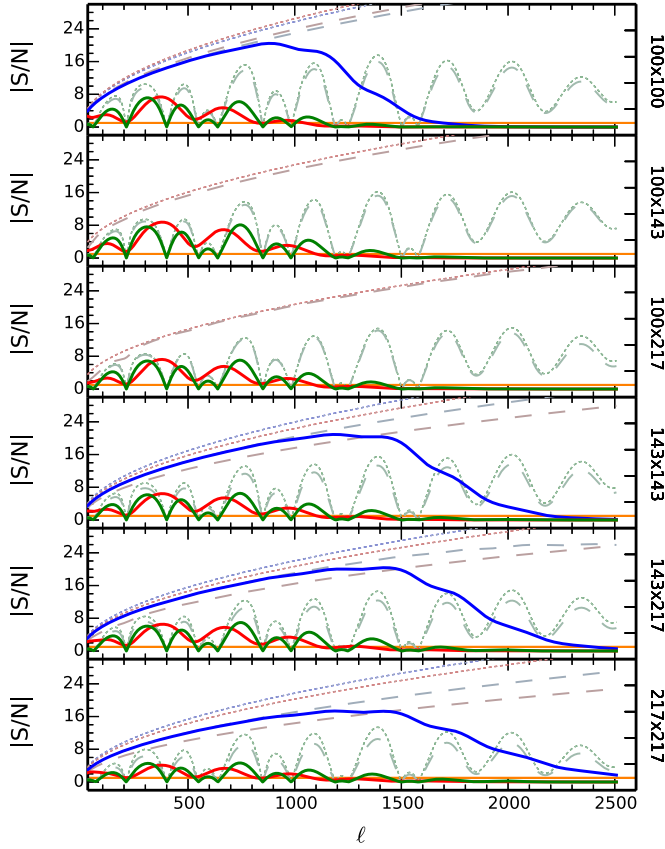
Fig. 14. Unbinned S/N per frequency for *TT* (solid blue, for those detector combinations used in the estimate of the *TT* spectrum), *EE* (solid red), and *TE* (solid green). The horizontal orange line corresponds to $S/N = 1$. The dashed lines indicate the S/N in a cosmic-variance-limited case, obtained by forcing the instrumental noise terms to zero when calculating the power spectrum covariance matrix. The dotted lines indicate the cosmic-variance-limited case computed with the approximate formula of Eq. (20).

For *TE* and *EE* we are more conservative, and cut the low S/N 100 GHz data at small scales ($\ell > 1000$), and the possibly dust-contaminated 217 GHz at large scales ($\ell < 500$). Only the 143×143 *TE* and *EE* spectra cover the full multipole range, restricted to $\ell < 2000$. Retaining more multipoles would require more in-depth modelling of residual systematic effects, which is left to future work. All the cuts are summarized in Table 9 and shown in Fig. 15.

Figure 14 also shows that each of the *TT* frequency power spectra is cosmic-variance dominated in a wide interval of multipoles. In particular, if we define as cosmic-variance dominated the ranges of multipoles where cosmic variance contributes more than half of the total variance, we find that the 100×100 GHz spectrum is cosmic-variance dominated at $\ell \lesssim 1156$, the 143×143 GHz at $\ell \lesssim 1528$, the 143×217 GHz at $\ell \lesssim 1607$, and the 217×217 GHz at $\ell \lesssim 1566$. To determine these ranges, we calculated the ratio of cosmic to total variance, where the cosmic variance is obtained from the diagonal elements of the covariance matrix after setting the instrumental noise to zero. Furthermore, we find that each of the *TE* frequency power spectra is cosmic-variance limited in some limited ranges of multipoles, below $\ell \lesssim 150$ ($\ell \lesssim 50$ for the 100×100)[11], in the range

---

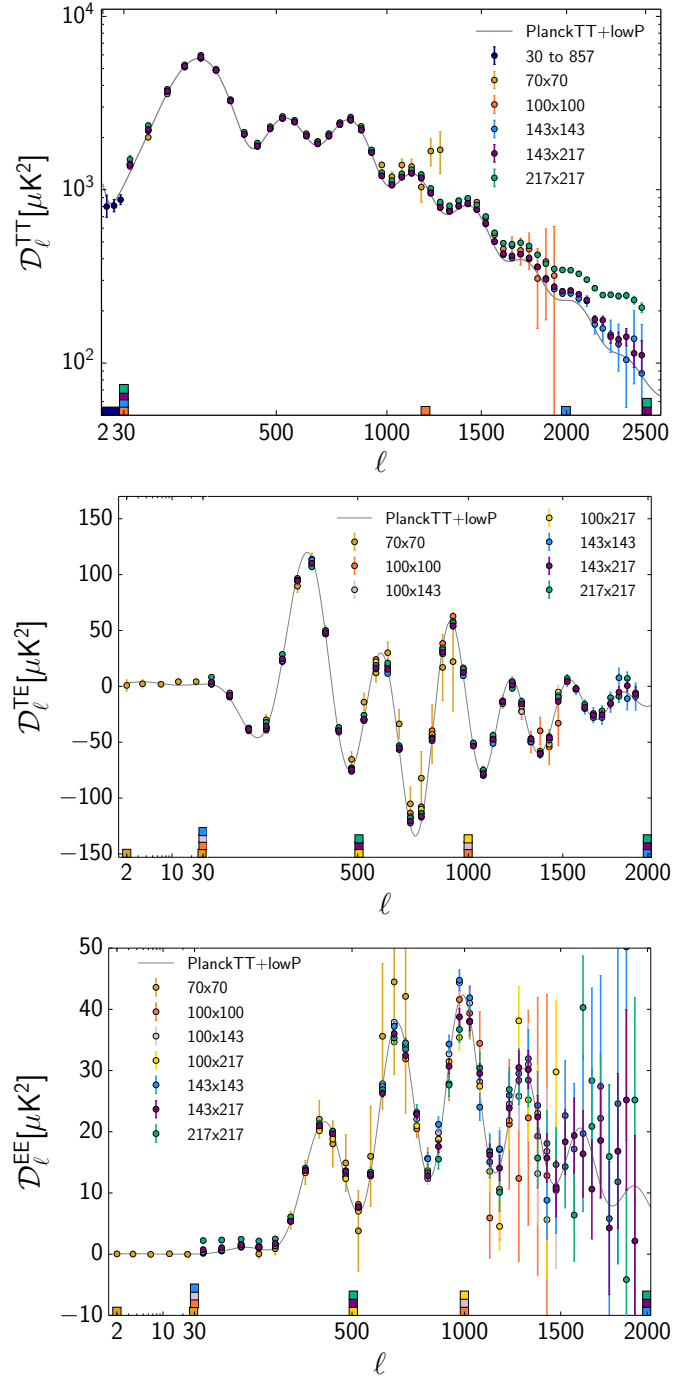[11] Recall that these statements refer to the high-ℓ likelihood ($\ell \geq 30$).

Fig. 15. *Planck* power spectra (not yet corrected for foregrounds) and data selection. The coloured tick marks indicate the ℓ-range of the cross-spectra included in the *Planck* likelihood. Although not used in the high-ℓ likelihood, the 70 GHz spectra at $\ell > 29$ illustrate the consistency of the data. The grey line indicates the best-fit *Planck* 2015 spectrum. The *TE* and *EE* plots have a logarithmic horizontal scale for $\ell < 30$.

$\ell \approx 250{-}450$ and additionally in the range $\ell \approx 650{-}700$ only for the 100×143 GHz and the 143×217 GHz power spectra.

Finally, when we co-add the foreground-cleaned frequency spectra to provide the CMB spectra (see Appendix C.4), we find that the CMB *TT* power spectrum is cosmic-variance dominated at $\ell \lesssim 1586$, while *TE* is cosmic-variance dominated at $\ell \lesssim 158$ and $\ell \approx 257{-}464$.

Table 9. Multipole cuts for the `Plik` temperature and polarization spectra at high ℓ.

Frequency [GHz]	Multipole range
TT	
100×100 ..	30–1197
143×143 ..	30–1996
143×217 ..	30–2508
217×217 ..	30–2508
TE	
100×100 ..	30–999
100×143 ..	30–999
100×217 ..	505–999
143×143 ..	30–1996
143×217 ..	505–1996
217×217 ..	505–1996
EE	
100×100 ..	30–999
100×143 ..	30–999
100×217 ..	505–999
143×143 ..	30–1996
143×217 ..	505–1996
217×217 ..	505–1996

Fig. 16. The relative weights of each frequency cross-spectrum in the *TT* (*top*), *TE* (*middle*) and *EE* (*bottom*) best-fit solution. Sharp jumps are due to the multipole selection. Weights are normalized to sum to one.

Due to the different masks, multipole ranges, noise levels, and to a lesser extent differing foreground contamination, each cross-spectrum ends up contributing differently as a function of scale to the best CMB solution. The determination of the mixing weights is described in Appendix C.4. Figure 16 presents the resulting (relative) weights of each cross-spectra. In temperature, the 100×100 spectrum dominates the solution until $\ell \approx 800$, when the solution becomes driven by the 143×143 up to $\ell \approx 1400$. The 143×217 and 217×217 provide the solution for the higher multipoles. In polarization, the 100×143 dominates the solution until $\ell \approx 800$ (with an equal contribution from 100×100 until $\ell \approx 400$ in *TE* only) while the higher ℓ range is dominated by the 143×217 contribution. Not surprisingly, the weights of the higher frequencies tend to increase with ℓ.

3.2.5. Binning

The 2013 baseline likelihood used unbinned temperature power spectra. For this release, we include polarization, which substantially increases the size of the numerical task. The 2015 likelihood therefore uses binned power spectra by default, downsizing the covariance matrix and speeding up likelihood computations. Indeed, even with the multipole-range cut just described, the unbinned data vector has around 23 000 elements, two thirds of which correspond to *TE* and *EE*. For some specific purposes (e.g., searching for oscillatory features in the *TT* spectrum or testing χ^2 statistics) we also produce an unbinned likelihood.

The spectra are binned into bins of width $\Delta\ell = 5$ for $30 \leq \ell \leq 99$, $\Delta\ell = 9$ for $100 \leq \ell \leq 1503$, $\Delta\ell = 17$ for $1504 \leq \ell \leq 2013$, and $\Delta\ell = 33$ for $2014 \leq \ell \leq 2508$, with a weighting of the C_ℓ proportional to $\ell(\ell + 1)$ over the bin widths,

$$C_{\rm b} = \sum_{\ell=\ell_b^{\rm min}}^{\ell_b^{\rm max}} w_b^\ell C_\ell, \text{ with } w_b^\ell = \frac{\ell(\ell + 1)}{\sum_{\ell=\ell_b^{\rm min}}^{\ell_b^{\rm max}} \ell(\ell + 1)}. \qquad (21)$$

The bin-widths are odd numbers, since for approximately azimuthal masks we expect a nearly symmetrical correlation function around the central multipole. It is shown explicitly in Sect. 4.1 that the binning does not affect the determination of cosmological parameters in ΛCDM-type models, which have smooth power spectra.
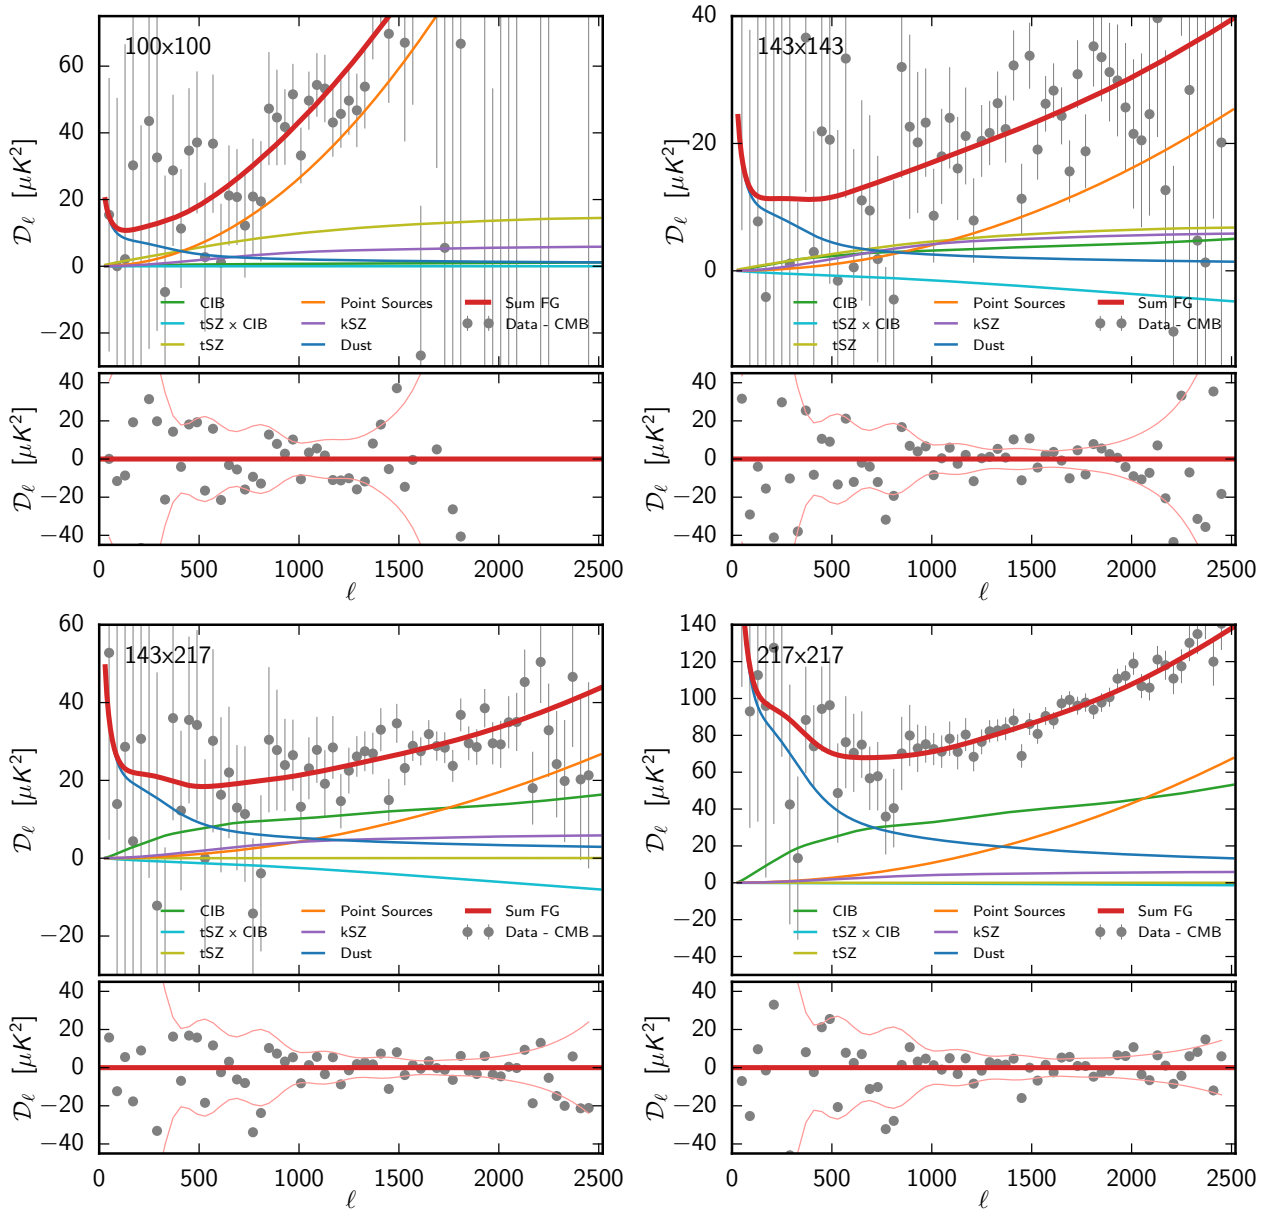
Fig. 17. Best foreground model in each of the cross-spectra used for the temperature high-ℓ likelihood. The data corrected by the best theoretical CMB C_ℓ are shown in grey. The *bottom panel* of each plot shows the residual after foreground correction. The pink line shows the 1σ value from the diagonal of the covariance matrix (32% of the unbinned points are out of this range).

3.3. Foreground modelling

Most of the foreground elements in the model parameter vector are similar to those in Like13. The main differences are in the dust templates, which have changed to accommodate the new masks. The *TE* and *EE* foreground model only takes into account the dust contribution and neglects any other Galactic polarized emission, in particular the synchrotron contamination. Nor do we mask out any extragalactic polarized foregrounds, as they have been found to be negligible by ground-based, small-scale experiments (Naess et al. 2014; Crites et al. 2015).

Figure 17 shows the foreground decomposition in temperature for each of the cross-spectra combinations we use in the likelihood. The figure also shows the CMB-corrected data (i.e., data minus the best-fit ΛCDM CMB model) as well as the residuals after foreground correction. In each spectrum, dust dominates the low-ℓ modes, while point sources dominate the smallest scales. For 217×217 and 143×217, the intermediate range

has a significant CIB contribution. We note that for 100×100, even when including 66% of the sky, the dust contribution is almost negligible and the point-source term is dominant well below $\ell = 500$. The least foreground-contaminated spectrum is 143×143. For comparison, Fig. 18 shows the full model, including the CMB. The foreground contribution is a small fraction of the total power at large scales.

Table 10 summarizes the parameters used for astrophysical foreground modelling and their associated priors.

3.3.1. Galactic dust emission

Galactic dust is the main foreground contribution at large scales and thus deserves close attention. This section describes how we model its power spectra. We express the dust contribution to the power spectrum calculated from map X at frequency ν and map
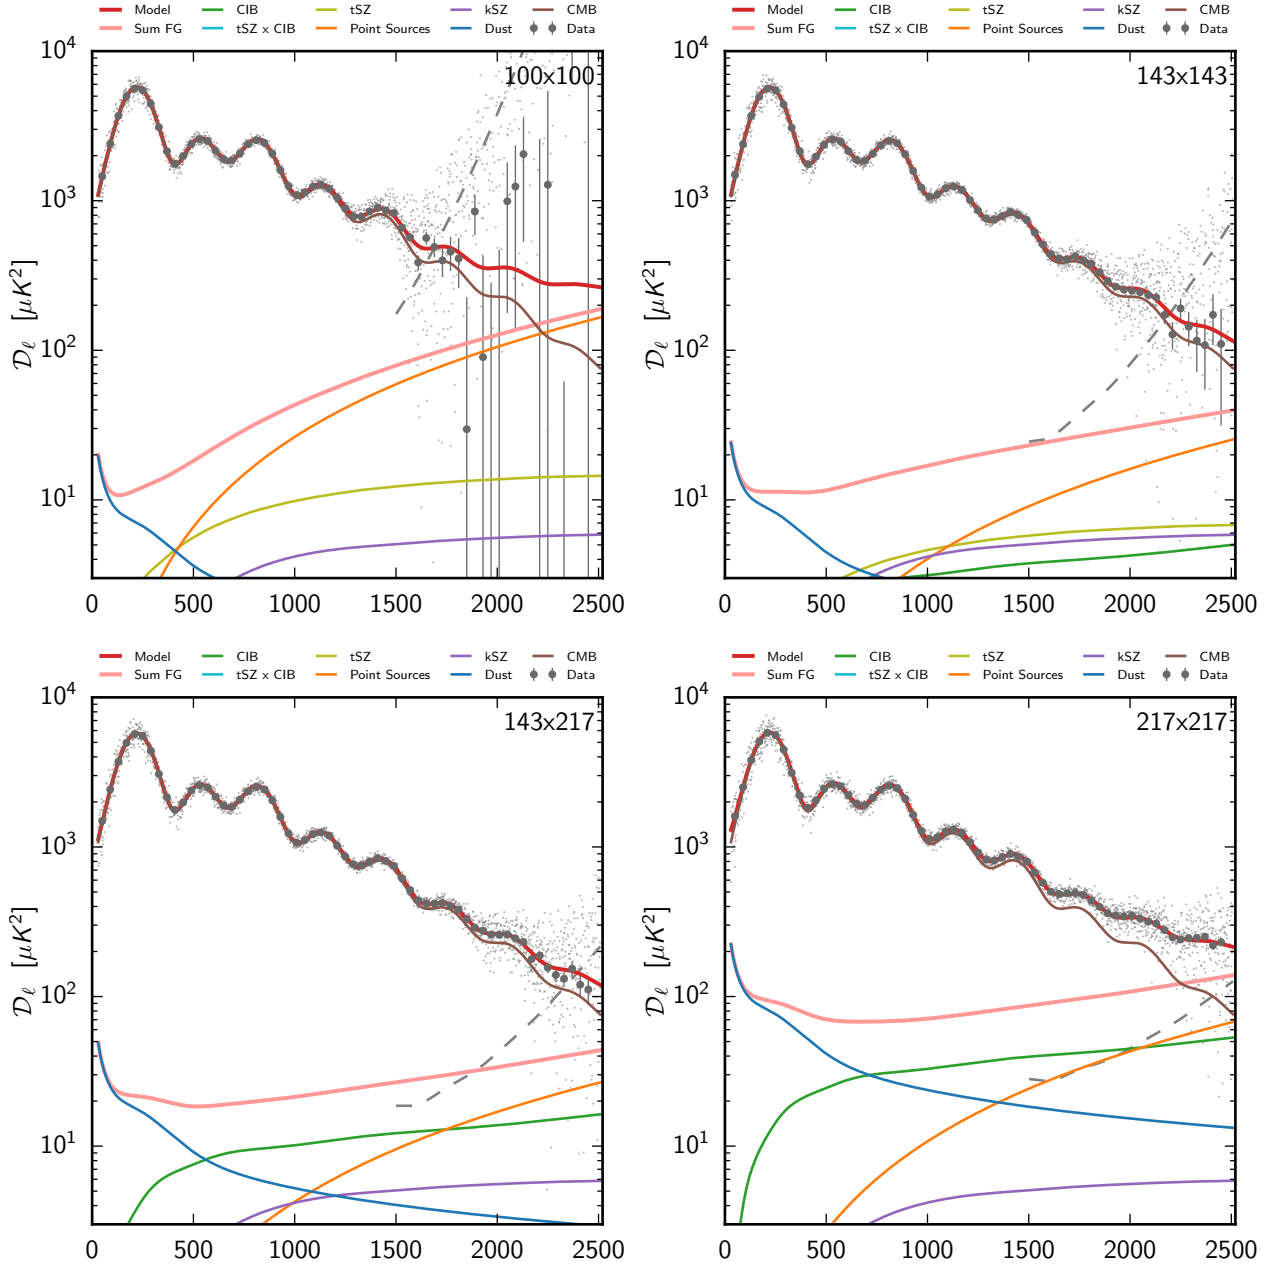
Fig. 18. Best model (CMB and foreground) in each of the cross-spectra used for the temperature high-ℓ likelihood. The small light grey points show the unbinned data point, and the dashed grey line show the square root of the noise contribution to the diagonal of the unbinned covariance matrix.

Y at frequency ν' as

$$\left(C_{\nu\times\nu'}^{XY,\text{dust}}\right)_\ell = A_{\nu\times\nu'}^{XY,\text{dust}} \times C_\ell^{XY,\text{dust}}, \tag{22}$$

where XY is one of TT, EE, or TE, and $C_\ell^{XY,\text{dust}}$ is the template dust power spectrum, with corresponding amplitude $A_{\nu\times\nu'}^{XY,\text{dust}}$. We assume that the dust power spectra have the same spatial dependence across frequencies and masks, so the dependence on sky fraction and frequency is entirely encoded in the amplitude parameter A. We do not try to enforce any a priori scaling with frequency, since using different masks at different frequencies makes determination of this scaling difficult. When both frequency maps ν and ν' are used in the likelihood with the same mask, we simply assume that the amplitude parameter can be written as

$$A_{\nu\times\nu'}^{XY,\text{dust}} = a_\nu^{XY,\text{dust}} \times a_{\nu'}^{XY,\text{dust}}. \tag{23}$$

This is clearly not exact when $XY = TE$ and $\nu \neq \nu'$. Similarly the multipole-dependent weight used to combine TE and ET for different frequencies breaks the assumption of an invariant dust template. These approximations do not appear to be the limiting factor of the current analysis.

In contrast to the choice we made in 2013, when all Galactic contributions were fixed and a dust template had been explicitly subtracted from the data, we now fit for the amplitude of the dust contribution in each cross-spectrum, in both temperature and polarization. This enables exploration of the possible degeneracy between the dust amplitude and cosmological parameters. A comparison of the two approaches is given in Sect. D.1 and Fig. D.2.

In the following, we describe how we build our template dust power spectrum from high-frequency data and evaluate the

Table 10. Parameters used for astrophysical foregrounds and instrumental modelling.

Parameter	Prior range	Definition
A_{100}^{PS}	[0, 400]	Contribution of Poisson point-source power to $\mathcal{D}_{3000}^{100\times100}$ for *Planck* (in $\mu\mathrm{K}^2$)
A_{143}^{PS}	[0, 400]	As for A_{100}^{PS} but at 143 GHz
A_{217}^{PS}	[0, 400]	As for A_{100}^{PS} but at 217 GHz
$A_{143\times217}^{\mathrm{PS}}$	[0, 400]	As for A_{100}^{PS} but at 143×217 GHz
A_{217}^{CIB}	[0, 200]	Contribution of CIB power to $\mathcal{D}_{3000}^{217}$ at the *Planck* CMB frequency for 217 GHz (in $\mu\mathrm{K}^2$)
A^{tSZ}	[0, 10]	Contribution of tSZ to $\mathcal{D}_{3000}^{143\times143}$ at 143 GHz (in $\mu\mathrm{K}^2$)
A^{kSZ}	[0, 10]	Contribution of kSZ to $\mathcal{D}_{3000}$ (in $\mu\mathrm{K}^2$)
$\xi^{\mathrm{tSZ}\times\mathrm{CIB}}$	[0, 1]	Correlation coefficient between the CIB and tSZ
$A_{100}^{\mathrm{dust}TT}$	[0, 50] (7 ± 2)	Amplitude of Galactic dust power at $\ell = 200$ at 100 GHz (in $\mu\mathrm{K}^2$)
$A_{143}^{\mathrm{dust}TT}$	[0, 50] (9 ± 2)	As for $A_{100}^{\mathrm{dust}TT}$ but at 143 GHz
$A_{143\times217}^{\mathrm{dust}TT}$	[0, 100] (21 ± 8.5)	As for $A_{100}^{\mathrm{dust}TT}$ but at 143×217 GHz
$A_{217}^{\mathrm{dust}TT}$	[0, 400] (80 ± 20)	As for $A_{100}^{\mathrm{dust}TT}$ but at 217 GHz
c_{100}	[0, 3] (0.9990004 ± 0.001)	Power spectrum calibration for the 100 GHz
c_{217}	[0, 3] (0.99501 ± 0.002)	Power spectrum calibration for the 217 GHz
y_{cal}	[0.9, 1.1] (1 ± 0.0025)	Absolute map calibration for *Planck*
$A_{100}^{\mathrm{dust}EE}$	[0, 10] (0.06 ± 0.012)	Amplitude of Galactic dust power at $\ell = 500$ at 100 GHz (in $\mu\mathrm{K}^2$)
$A_{100\times143}^{\mathrm{dust}EE}$	[0, 10] (0.05 ± 0.015)	As for $A_{100}^{\mathrm{dust}EE}$ but at 100×143 GHz
$A_{100\times217}^{\mathrm{dust}EE}$	[0, 10] (0.11 ± 0.033)	As for $A_{100}^{\mathrm{dust}EE}$ but at 100×217 GHz
$A_{143}^{\mathrm{dust}EE}$	[0, 10] (0.1 ± 0.02)	As for $A_{100}^{\mathrm{dust}EE}$ but at 143 GHz
$A_{143\times217}^{\mathrm{dust}EE}$	[0, 10] (0.24 ± 0.048)	As for $A_{100}^{\mathrm{dust}EE}$ but at 143×217 GHz
$A_{217}^{\mathrm{dust}EE}$	[0, 10] (0.72 ± 0.14)	As for $A_{100}^{\mathrm{dust}EE}$ but at 217 GHz
$A_{100}^{\mathrm{dust}TE}$	[0, 10] (0.14 ± 0.042)	Amplitude of Galactic dust power at $\ell = 500$ at 100 GHz (in $\mu\mathrm{K}^2$)
$A_{100\times143}^{\mathrm{dust}TE}$	[0, 10] (0.12 ± 0.036)	As for $A_{100}^{\mathrm{dust}TE}$ but at 100×143 GHz
$A_{100\times217}^{\mathrm{dust}TE}$	[0, 10] (0.3 ± 0.09)	As for $A_{100}^{\mathrm{dust}TE}$ but at 100×217 GHz
$A_{143}^{\mathrm{dust}TE}$	[0, 10] (0.24 ± 0.072)	As for $A_{100}^{\mathrm{dust}TE}$ but at 143 GHz
$A_{143\times217}^{\mathrm{dust}TE}$	[0, 10] (0.6 ± 0.18)	As for $A_{100}^{\mathrm{dust}TE}$ but at 143×217 GHz
$A_{217}^{\mathrm{dust}TE}$	[0, 10] (1.8 ± 0.54)	As for $A_{100}^{\mathrm{dust}TE}$ but at 217 GHz

Notes. The columns indicate the symbol for each parameter, the prior used for exploration (square brackets denote uniform priors, parentheses indicate Gaussian priors), and definitions. Beam eigenmode amplitudes require a correlation matrix to fully describe their joint prior and so do not appear in the table; they are internally marginalized over rather than explicitly sampled. This table only lists the instrumental parameters that are explored in the released version, but we do consider more parameters to assess the effects of beam uncertainties and beam leakage; see Sect. 3.4.3.

amplitude of the dust contamination at each frequency and for each mask.

As we shall see later in Sect. 4.1.2, the cosmological values recovered from TT likelihood explorations do not depend on the dust amplitude priors, as shown by the case "No gal. priors" in Fig. 35 and discussed in Sect. 4.1.2. The polarization case is discussed in Sect. C.3.5. Section 5.3 and Figs. 44 and 45 show the correlation between the dust and the cosmological or other foreground parameters. The dust amplitudes are found to be nearly

uncorrelated with the cosmological parameters except for *TE*. However, the priors do help to break the degeneracies between foreground parameters, which are found to be much more correlated with the dust. In Appendix E we further show that our results are insensitive to broader changes in the dust model.

Galactic TT dust emission. We use the 545 GHz power spectra as templates for Galactic dust spatial fluctuations. The 353 GHz

detectors also have some sensitivity to dust, along with a significant contribution from the CMB, and hence any error in removing the CMB contribution at 353 GHz data translates into biases on our dust template. This is much less of an issue at 545 GHz, to the point where entirely ignoring the CMB contribution does not change our estimate of the template. Furthermore, estimates using 545 GHz maps tend to be more stable over a wider range of multipoles than those obtained from 353 GHz or 857 GHz maps.

We aggressively mask the contribution from point sources in order to minimize their residual, the approximately white spectrum of which is substantially correlated with the value of some cosmological parameters (see the discussion of parameter correlations in Sect. 5.3). The downside of this is that the point-source masks remove some of the brightest Galactic regions that lie in regions not covered by our Galactic masks. This means that we cannot use the well-established power-law modelling advocated in Planck Collaboration XI (2014) and must instead compute an effective dust (residual) template.

All of the masks that we use in this section are combinations of the joint point-source, extended-object, and CO masks used for 100 GHz, 143 GHz, and 217 GHz with Galactic masks of various sizes. In the following discussion we refer only to the Galactic masks, but in all cases the masks contain the other components as well. The half-mission cross-spectra at 545 GHz provide us with a good estimate of the large-scale behaviour of the dust. Small angular scales, however, are sensitive to the CIB, with the intermediate range of scales dominated by the clustered part and the smallest scales by the Poisson distribution of infrared point sources. These last two terms are statistically isotropic, while the dust amplitude depends on the sky fraction. Assuming that the shapes of the dust power spectra outside the masks do not vary substantially as the sky fraction changes, we rely on mask differences to build a CIB-cleaned template of the dust.

Figure 19 shows that this assumption is valid when changing the Galactic mask from G60 to G41. It shows that the 545 GHz cross-half-mission power spectrum can be well represented by the sum of a Galactic template, a CIB contribution, and a point source contribution. The Galactic template is obtained by computing the difference between the spectra obtained in the G60 and the G41 masks. This difference is fit to a simple analytic model

$$C_\ell^{TT,\mathrm{dust}} \propto (1 + h\,\ell^k\,\mathrm{e}^{-\ell/t}) \times (\ell/\ell_p)^n, \qquad (24)$$

with $h = 2.3 \times 10^{-11}$, $k = 5.05$, $t = 56$, $n = -2.63$, and fixing $\ell_p = 200$. The model behaves like a $C_{\ell,\mathrm{dust}}^{TT} \propto \ell^{-2.63}$ power law at small scales, and has a bump around $\ell = 200$. The CIB model we use is described in Sect. 3.3.2.

We can compare this template model with the dust content in each of the power spectra we use for the likelihood. Of course those power spectra are strongly dominated by the CMB, so, to reveal the dust content, one has to rely on the same trick that was used for 545 GHz. This however is not enough, since the CMB cosmic variance itself is significant compared to the dust contamination. We can build an estimate of the CMB cosmic variance by assuming that at 100 GHz the dust contamination is small enough that a mask difference gives us a good variance estimate.

Figure 20 shows the mask difference (corrected for cosmic variance) between G60 and G41 for the 217 GHz and 143 GHz half-mission cross-spectra, as well as the dust model from Eq. (24). The dust model has been rescaled to the expected mask difference dust residual for the 217 GHz. The 143 GHz
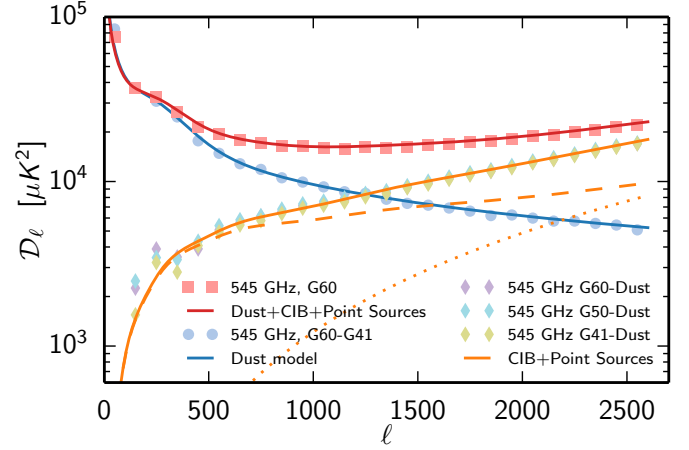

Fig. 19. Dust model at 545 GHz. The dust template is based on the G60–G41 mask difference of the 545 GHz half-mission cross-spectrum (blue line and circles, rescaled to the dust level in mask G60). Coloured diamonds display the difference between this model (rescaled in each case) and the cross half-mission spectra in the G41, G50, and G60 masks. The residuals are all in good agreement (less so at low ℓ, because of sample variance) and are well described by the CIB+point source prediction (orange line). Individual CIB and point sources contributions are shown as dashed and dotted orange lines. The red line is the sum of the dust model, CIB, and point sources for the G60 mask, and is in excellent agreement with the 545 GHz cross half-mission spectrum in G60 (red squares). In all cases, the spectra were computed by using different Galactic masks supplemented by the single combination of the 100 GHz, 143 GHz, and 217 GHz point sources, extended objects and CO masks.

mask-difference has also been rescaled in a similar way. The ratio between the two is about 14. Rescaling factors are obtained from Table 11. Error bars are estimated based on the scatter in each bin. The agreement with the model is very good at 217 GHz, but less good at 143 GHz where the greater scatter is probably dominated at large scales by the chance correlation between CMB and dust (which, as we see in Eq. (25), varies as the square root of the dust contribution to the spectra), and at small scale by noise. We also tested these double differences for other masks, namely G50–G41 and G60–G50, and verified that the results are similar (i.e., general agreement although with substantial scatter).

Finally, we can estimate the level of the dust contamination in each of our frequency maps used for CMB analysis by computing their cross-spectra with the 545 GHz half-mission maps. Assuming that all our maps m^ν have in common only the CMB and a variable amount of dust, and assuming that $m^{545} = m^{\mathrm{cmb}} + a^{545}m^{\mathrm{dust}}$, the cross-spectra between each of our CMB frequencies maps and the 545 GHz map is

$$\left(C_{545\times\nu}^{TT}\right)_\ell = C_\ell^{TT,\mathrm{cmb}} + a_{545}^{TT,\mathrm{dust}} a_\nu^{TT,\mathrm{dust}} C_\ell^{TT,\mathrm{dust}}$$
$$+ (a_{545}^{TT,\mathrm{dust}} + a_\nu^{TT,\mathrm{dust}}) C_\ell^{\mathrm{chance}}, \qquad (25)$$

where C_ℓ^{chance} is the chance correlation between the CMB and dust distribution (which would vanish on average over many sky realizations). By using the 100 GHz spectrum as our CMB estimate and assuming that the chance correlation is small enough, one can measure the amount of dust in each frequency map by fitting the rescaling factor between the (CMB cleaned) 545 GHz spectrum and the cross frequency spectra. This approach is limited by the presence of CIB which has a slightly different emission law than the dust. We thus limit our fits to the multipoles

Table 11. Contamination level in each frequency, $D_{\ell=200}$.

Frequency [GHz]	Contamination Level $[\mu K^2]$				
	Mask				
	G41	G50	G60	G70	CIB
100	1.6 ± 0.8	1.6 ± 0.8	3.2 ± 1.2	7.1 ± 1.6	0.24 ± 0.04
143	6.0 ± 1.2	6.6 ± 1.4	10 ± 1.8	23.5 ± 4	1.0 ± 0.2
217	84 ± 16	91 ± 18	150 ± 20	312 ± 35	10 ± 2

Notes. The levels reported in this table correspond to the amplitude of the contamination, $\mathcal{D}_\ell$, at $\ell = 200$ in μK^2. They are obtained at each frequency by fitting the 545 GHz cross half-mission spectra against the CMB-corrected 545×100, 545×143 and 545×217 spectra over a range of multipoles. The CMB correction is obtained using the 100 GHz cross half-mission spectra. This contamination is dominated by dust, with a small CIB contribution. The columns labelled with a Galactic mask name (G41, G50, G60, and G70) correspond to the results when combining those masks with the same CO, extended object, and frequency-combined point-source masks. The CIB contribution is shown in the last column. The errors quoted here include the variation when changing the range of multipoles used from $30 \leq \ell \leq 1000$ to $30 \leq \ell \leq 500$.
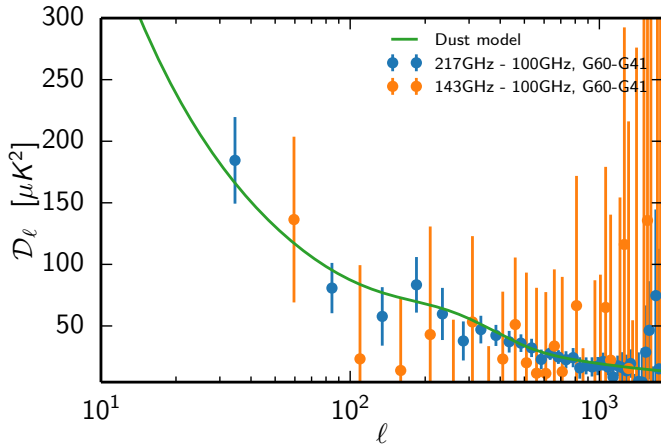

Fig. 20. Dust model versus data. In blue, the power spectrum of the double mask difference between 217 GHz and 100 GHz half-mission cross-spectra in masks G60 and G41 (complemented by the joint masks for CO, extended objects, and point sources). In orange, the equivalent spectrum for 143 and 100 GHz. The mask difference enables us to remove the contribution from all the isotropic components (CMB, CIB, and point sources) in the mean. But simple mask differences are still affected by the difference of the CMB in the two masks due to cosmic variance. Removing the 100 GHz mask difference, which is dominated by the CMB, reduces the scatter significantly. The error bars are computed as the scatter in bins of size $\Delta\ell = 50$. The dust model (green) based on the 545 GHz data has been rescaled to the expected dust contamination in the 217 GHz mask difference using values from Table 11. The 143 GHz double mask difference is also rescaled to the level of the 217 GHz difference; i.e., it is multiplied by approximately 14. Different multipole bins are used for the 217 GHz and 143 GHz data to improve readability.

$\ell < 1000$ where the CIB is small compared to the dust and we ignore the emission-law differences.

Table 11 reports the results of those fits at each frequency, for each Galactic mask. The error range quoted corresponds to the error of the fits, taking into account the variations when changing the multipole range of the fit from $30 \leq \ell \leq 1000$ to $30 \leq \ell \leq 500$. The values reported correspond to the sum of the CIB and the dust contamination at $\ell = 200$. The last column gives the estimate of the CIB contamination at the same multipole from the joint cosmology and foreground fit. From this table, the ratio of the dust contamination at map level between the 217 GHz and

100 GHz is around 7, while the ratio between the 217 GHz and 143 GHz is close to 3.7.

We derive our priors on the foreground amplitudes from this table, combining the 545 GHz fit with the estimated residual CIB contamination, to obtain the following values: $(7\pm2)\,\mu K^2$ for the 100×100 spectrum (G70); $(9\pm2)\,\mu K^2$ for 143×143 (G60); and $(80\pm20)\,\mu K^2$ for 217×217 (G50). Finally the 143×217 value is obtained by computing the geometrical average between the two auto spectra under the worst mask (G60), yielding $(21\pm8.5)\,\mu K^2$.

Galactic TE and EE dust emission. We evaluate the dust contribution in the *TE* and *EE* power spectra using the same method as for the temperature. However, instead of the 545 GHz data we use the maps at 353 GHz, our highest frequency with polarization information. At sufficently high sky fractions, the 353 GHz *TE* and *EE* power spectra are dominated by dust. As estimated in Planck Collaboration Int. XXX (2016), there is no other significant contribution from the Galaxy, even at 100 GHz. Following Planck Collaboration Int. XXX (2016), and since we do not mask any "point-source-like" region of strong emission, we can use a power-law model as a template for the polarized Galactic dust contribution. Enforcing a single power law for *TE* and *EE* and our different masks, we obtain an index of $n = -2.4$. We use the same cross-spectra-based method to estimate the dust contamination. The dust contribution being smaller in polarization, removing the CMB from the 353×353 and the $353 \times \nu$ (with ν being one of 100, 143 or 217) is particularly important. Our two best CMB estimates in *EE* and *TE* being 100 and the 143 GHz, we checked that using any of 100×100, 143×143, or 100×143 does not change the estimates significantly. Table 12 gives the resulting values. As for the *TT* case, the cross-frequency, cross-masks estimates are obtained by computing the geometric average of the auto-frequency contaminations under the smallest mask.

3.3.2. Extragalactic foregrounds

The extragalactic foreground model is similar to that of 2013 and in the following we describe the differences. Since we are neglecting any possible contribution in polarization from extragalactic foregrounds, we omit the *TT* index in the following descriptions of the foreground models. The amplitudes are expressed as $\mathcal{D}_\ell$ at $\ell = 3000$ so that, for any component, the template, C_{3000}^{FG}, satisfies $C_{3000}^{FG}\,\mathcal{A}_{3000} = 1$ with $\mathcal{A}_\ell = \ell(\ell+1)/(2\pi)$.

Table 12. *TE* and *EE* dust contamination levels, $D_{\ell=500}$.

Spectrum	Contamination level [μK^2]		
	100 GHz (G70)	143 GHz (G50)	217 GHz (G41)
$D^{TE}_{\ell=500}$			
100 GHz (G70)	0.14 ± 0.042	0.12 ± 0.036	0.3 ± 0.09
143 GHz (G50)		0.24 ± 0.072	0.6 ± 0.018
217 GHz (G41)			1.8 ± 0.54
$D^{EE}_{\ell=500}$			
100 GHz (G70)	0.06 ± 0.012	0.05 ± 0.015	0.11 ± 0.033
143 GHz (G50)		$0.1 \ \ \pm 0.02$	0.24 ± 0.048
217 GHz (G41)			0.72 ± 0.14

Notes. Values reported in the table correspond to the evaluation of the contamination level in each frequency by fitting the 353 GHz cross half-mission spectra against the CMB-corrected 353×100, 353×143 and 353×217 spectra over a range of multipoles. The CMB correction is obtained using the 100 GHz cross half-mission spectra (we have similar results at 143 GHz). Level reported here correspond to the amplitude of the contamination $\mathcal{D}_\ell$ at $\ell = 500$ in μK^2.

The cosmic infrared background. The CIB model has a number of differences from that used in Like13. First of all, it is now entirely parameterized by a single amplitude $\mathcal{D}^{CIB}_{217}$ and a template C^{CIB}_ℓ:

$$\left(C^{CIB}_{\nu \times \nu'}\right)_\ell = a^{CIB}_\nu a^{CIB}_{\nu'} C^{CIB}_\ell \times \mathcal{D}^{CIB}_{217}, \qquad (26)$$

where the spectral coefficients a^{CIB}_ν represent the CIB emission law normalized at $\nu = 217$ GHz.

In 2013, the template was an effective power-law model with a variable index with expected value $n = -1.37$ (when including the "highL" data from ACT and SPT). We did not assume any emission law and fitted the 143 GHz and 217 GHz amplitude, along with their correlation coefficient. The Planck Collaboration has studied the CIB in detail in Planck Collaboration XXX (2014) and now proposes a one-plus-two-halo model, which provides an accurate description of the *Planck* and IRAS CIB spectra from 3000 GHz down to 217 GHz. We extrapolate this model here, assuming it remains appropriate in describing the 143 GHz and 100 GHz data. The CIB emission law and template are computed following Planck Collaboration XXX (2014). The template power spectrum provided by this work has a very small frequency dependence that we ignore.

At small scales, $\ell > 2500$, the slope of the template is similar to the power law used in Like13. At larger scales, however, the slope is much shallower. This is in line with the variation we observed in 2013 on the power-law index of our simple CIB model when changing the maximum multipole. The current template is shown as the green line in the *TT* foreground component plots in Fig. 17.

In 2013, the correlation between the 143 GHz and 217 GHz CIB spectra was fitted, favouring a high correlation, greater than 90% (when including the "highL" data). The present model yields a fully correlated CIB between 143 GHz and 217 GHz.

We now include the the CIB contribution at 100 GHz, which was ignored in 2013. Another difference with the 2013 model is that the parameter controlling the amplitude at 217 GHz now directly gives the amplitude in the actual 217 GHz *Planck* band at $\ell = 3000$, i.e., it includes the colour correction. The ratio between the two is 1.33. The 2013 amplitude of the CIB contribution at $\ell = 3000$ (including the highL data) was $66 \pm 6.7 \, \mu K^2$,

while our best estimate for the present analysis is $63.9 \pm 6.6 \, \mu K^2$ (PlanckTT+lowP).

Point sources. At the likelihood level, we cannot differentiate between the radio- and IR-point sources. We thus describe their combined contribution by their total emissivity per frequency pair,

$$\left(C^{PS}_{\nu \times \nu'}\right)_\ell = \mathcal{D}^{PS}_{\nu \times \nu'} / \mathcal{A}_{3000}, \qquad (27)$$

where $\mathcal{D}_{\nu \times \nu'}$ is the amplitude of the point-source contribution in $\mathcal{D}_\ell$ at $\ell = 3000$. Contrary to 2013, we do not use a correlation parameter to represent the 143×217 point-source contribution; instead we use a free amplitude parameter. This has the disadvantage of not preventing a possible unphysical solution. However, it simplifies the parameter optimization, and it is easier to understand in terms of contamination amplitude.

Kinetic SZ (kSZ). We use the same model as in 2013. The kSZ emission is parameterized with a single amplitude and a fixed template from Trac et al. (2011),

$$\left(C^{kSZ}_{\nu \times \nu'}\right)_\ell = C^{kSZ}_\ell \times \mathcal{D}^{kSZ}, \qquad (28)$$

where $\mathcal{D}^{kSZ}$ is the kSZ contribution at $\ell = 3000$.

Thermal SZ (tSZ). Here again, we use the same model as in 2013. The tSZ emission is also parameterized by a single amplitude and a fixed template using the $\epsilon = 0.5$ model from Efstathiou & Migliaccio (2012),

$$\left(C^{tSZ}_{\nu \times \nu'}\right)_\ell = a^{tSZ}_\nu a^{tSZ}_{\nu'} C^{tSZ}_\ell \times \mathcal{D}^{tSZ}_{143}, \qquad (29)$$

where a^{tSZ}_ν is the thermal Sunyaev-Zeldovich spectrum, normalized to $\nu_0 = 143$ GHz and corrected for the *Planck* bandpass colour corrections. Ignoring the bandpass correction, we recall that the tSZ spectrum is given by

$$a^{tSZ}_\nu = \frac{f(\nu)}{f(\nu_0)}, \ f(\nu) = \left(x \coth\left(\frac{x}{2}\right) - 4\right), \ x = \frac{h\nu}{k_B T_{cmb}}. \qquad (30)$$

Thermal SZ × CIB correlation. Following [Like13] the cross-correlation between the thermal SZ and the CIB, tSZ× CIB, is parameterized by a single correlation parameter, ξ, and a fixed template from Addison et al. (2012),

$$\left(C^{\mathrm{tSZ \times CIB}}_{\nu \times \nu'}\right)_\ell = \xi \sqrt{\mathcal{D}^{\mathrm{tSZ}}_{143} \, \mathcal{D}^{\mathrm{CIB}}_{217}}$$
$$\times \left(a^{\mathrm{tSZ}}_\nu \, a^{\mathrm{CIB}}_{\nu'} + a^{\mathrm{tSZ}}_{\nu'} \, a^{\mathrm{CIB}}_\nu\right)$$
$$\times C^{\mathrm{tSZ \times CIB}}_\ell, \tag{31}$$

where a^{tSZ}_ν is the thermal Sunyaev-Zeldovich spectrum, corrected for the *Planck* bandpass colour corrections and a^{CIB}_ν is the CIB spectrum, rescaled at $\nu = 217\,\mathrm{GHz}$ as in the previous paragraphs.

SZ prior. The kinetic SZ, the thermal SZ, and its correlation with the CIB are not constrained accurately by the *Planck* data alone. Besides, the tSZ×CIB level is highly correlated with the amplitude of the tSZ. In 2013, we reduced the degeneracy between those parameters and improved their determination by adding the ACT and SPT data. In 2015, we instead impose a Gaussian prior on the tSZ and kSZ amplitudes, inspired by the constraints set by these experiments. From a joint analysis of the *Planck* 2013 data with those from ACT and SPT, we obtain

$$\mathcal{D}^{\mathrm{kSZ}} + 1.6\mathcal{D}^{\mathrm{tSZ}} = (9.5 \pm 3)\,\mu\mathrm{K}^2, \tag{32}$$

in excellent agreement with the estimates from Reichardt et al. (2012), once they are rescaled to the *Planck* frequencies (see Planck Collaboration XIII 2016, for a detailed discussion).

As can be seen in Fig. 17, the kSZ, tSZ, and tSZ×CIB correlations are always dominated by the dust, CIB, and point-source contributions.

3.4. Instrumental modelling

The following sections describe the instrument modelling elements of the model vector, addressing the issues of calibration and beam uncertainties in Sects. 3.4.1–3.4.3, and describing the noise properties in Sect. 3.4.4. For convenience, Table 10 defines the symbol used for the calibration parameters and the priors later used for exploring them.

3.4.1. Power spectra calibration uncertainties

As in 2013, we allow for a small recalibration of the different frequency power spectra, in order to account for residual uncertainties in the map calibration process. The mixing matrix in the model vector from Eq. (14) can be rewritten as

$$\left(M^{XY}_{ZW,\nu \times \nu'}\right)_\ell (\theta_{\mathrm{inst}}) = G^{XY}_{\nu \times \nu'}(\theta_{\mathrm{calib}}) \left(M^{XY,\mathrm{other}}_{ZW,\nu \times \nu'}\right)_\ell (\theta_{\mathrm{other}}),$$

$$G^{XY}_{\nu \times \nu'}(\theta_{\mathrm{calib}}) = \frac{1}{y_{\mathrm{P}}^2} \left(\frac{1}{2\sqrt{c^{XX}_\nu c^{YY}_{\nu'}}} + \frac{1}{2\sqrt{c^{XX}_{\nu'} c^{YY}_\nu}} \right), \tag{33}$$

where c^{XX}_ν is the calibration parameter for the XX power spectrum at frequency ν, X being either T or E, and y_{P} is the overall *Planck* calibration. We ignore the ℓ-dependency of the weighting function between the TE and ET spectra at different frequencies that are added to form an effective cross-frequency TE cross-spectrum. As in 2013, we use the TT at 143 GHz as our inter-calibration reference, so that $c^{TT}_{143} = 1$.

We further allow for an overall *Planck* calibration uncertainty, whose variation is constrained by a tight Gaussian prior,

$$y_{\mathrm{P}} = 1 \pm 0.0025. \tag{34}$$

This prior corresponds to the estimated overall uncertainty, which is discussed in depth in Planck Collaboration I (2016).

The calibration parameters can be degenerate with the foreground parameters, in particular the point sources at high ℓ (for TT) and the Galaxy for 217 GHz at low ℓ. We thus proceed as in 2013, and measure the calibration refinement parameters on the large scales and on small sky fractions near the Galactic poles. We perform the same estimates on a range of Galactic masks (G20, G30, and G41) restricted to different maximum multipoles (up to $\ell = 1500$). The fits are performed either by minimizing the scatter between the different frequency spectra, or by using the SMICA algorithm (see Planck Collaboration VI 2014, Sect. 7.3) with a freely varying CMB and generic foreground contribution. For the TT spectra, we obtained in both cases very similar recalibration estimates, from which we extracted the conservative Gaussian priors on recalibration factors,

$$c^{TT}_{100} = 0.999 \pm 0.001, \tag{35}$$
$$c^{TT}_{217} = 0.995 \pm 0.002. \tag{36}$$

These are compatible with estimates made at the map level, but on the whole sky; see Planck Collaboration VIII (2016).

3.4.2. Polarization efficiency and angular uncertainty

We now turn to the polarization recalibration case. The signal measured by an imperfect PSB is given by

$$d = G(1+\gamma)\left[I + \rho(1+\eta)\left(Q\cos 2(\phi+\omega) + U\sin 2(\phi+\omega)\right)\right]+n, \tag{37}$$

where I, Q, and U are the Stokes parameters; n is the instrumental noise; G, ρ, and ϕ are the nominal photometric calibration factor, polar efficiency, and direction of polarization of the PSB; and γ, η, and ω are the (small) errors made on each of them (see, e.g., Jones et al. 2007). Due to these errors, the measured cross-power spectra of maps a and b are then contaminated by a spurious signal given by

$$\Delta C^{TT}_\ell = (\gamma_a + \gamma_b)\, C^{TT}_\ell, \tag{38a}$$
$$\Delta C^{TE}_\ell = \left(\gamma_a + \gamma_b + \eta_b - 2\omega_b^2\right) C^{TE}_\ell, \tag{38b}$$
$$\Delta C^{EE}_\ell = \left(\gamma_a + \gamma_b + \eta_a + \eta_b - 2\omega_a^2 - 2\omega_b^2\right) C^{EE}_\ell$$
$$+ 2\left(\omega_a^2 + \omega_b^2\right) C^{BB}_\ell, \tag{38c}$$

where γ_x, η_x, and ω_x, for $x = a, b$, are the effective instrumental errors for each of the two frequency-averaged maps. Pre-flight measurements of the HFI polarization efficiencies, ρ, had uncertainties $|\eta_x| \approx 0.3\%$, while the polarization angle of each PSB is known to $|\omega_x| \approx 1°$ (Rosset et al. 2010). Analysis of the 2015 maps shows the relative photometric calibration of each detector at 100 to 217 GHz to be known to about $|\gamma_x| = 0.16\%$ at worst, with an absolute orbital dipole calibration of about 0.2%, while analysis of the Crab Nebula observations showed the polarization uncertainties to be consistent with the pre-flight measurements (Planck Collaboration VIII 2016).

Assuming C^{BB}_ℓ to be negligible, and ignoring $\omega^2 \ll |\eta|$ in Eq. (38), the Gaussian priors on γ and η for each frequency-averaged polarized map would have rms of $\sigma_\gamma = 2 \times 10^{-3}$ and

$\sigma_\eta = 3 \times 10^{-3}$. Adding those uncertainties in quadrature, the auto-power spectrum recalibration c_ν^{EE} introduced in Eq. (33) would be given, for an equal-weight combination of $n_d = 8$ polarized detectors, by

$$c_\nu^{EE} = 1 \pm 2 \sqrt{\frac{\sigma_\gamma^2 + \sigma_\eta^2}{n_d}} = 1 \pm 0.0025. \qquad (39)$$

The most accurate recalibration factors for *TE* and *EE* could therefore be somewhat different from *TT*. We found, though, that setting the *EE* recalibration parameter to unity or implementing those priors makes no difference with respect to cosmology; i.e., we recover the same cosmological parameters, with the same uncertainties. Thus, for the baseline explorations, we fixed the *EE* recalibration parameter to unity,

$$c_\nu^{EE} = 1, \qquad (40)$$

and the uncertainty on *TE* comes only from the *TT* calibration parameter through Eq. (33).

We also explored the case of much looser priors, and found that best-fit calibration parameters deviate very significantly, and reach values of several percent (between 3% and 12% depending on the frequencies and on whether we fit the *EE* or *TE* case). This cannot be due to the instrumental uncertainties embodied in the prior. In the absence of an informative prior, this degree of freedom is used to minimize the differences between frequencies that stem from other effects, not included in the baseline modelling.

The next section introduces one such effect, the temperature-to-polarization leakage, which is due to combining detectors with different beams without accounting for it at the map-making stage (see Sect. 3.4.3). But anticipating the results of the analysis described in Appendix C.3.5, we note that when the calibration and leakage parameters are explored simultaneously without priors, they remain in clear tension with the priors (even if the level of recalibration decreases slightly, by typically 2%, showing the partial degeneracy between the two). In other words, when calibration and leakage parameters are both explored with their respective priors, there is evidence of residual unmodelled systematic effects in polarization – to which we will return.

3.4.3. Beam and transfer function uncertainties

The power spectra from map pairs are corrected by the corresponding effective beam window functions before being confronted with the data model. However, these window functions are not perfectly known, and we now discuss various related sources of errors and uncertainties, the impact of which on the reconstructed C_ℓs is shown in Fig. 21.

Sub-pixel effects. The first source of error, the so-called "sub-pixel" effect, discussed in detail in Like13, is a result of the *Planck* scanning strategy and map-making procedure. Scanning along rings with very low nutation levels can result in the centroid of the samples being slightly shifted from the pixel centres; however, the map-making algorithm assigns the mean value of samples in the pixel to the centre of the pixel. This effect, similar to the gravitational lensing of the CMB, has a non-diagonal influence on the power spectra, but the correction can be computed given the estimated power spectra for a given data selection, and recast into an additive, fixed component. We showed in Like13 that including this effect had little impact on the cosmological parameters measured by *Planck*.

Masking effect. A second source of error is the variation, from one sky pixel to another, of the effective beam width, which is averaged over all samples falling in that pixel. While all the `HEALPix` pixels have the same surface area, their shape – and therefore their moment of inertia (which drives the pixel window function) – depends on location, as shown in Fig. 22, and therefore makes the effective beam window function depend on the pixel mask considered. Of course the actual sampling of the pixels by *Planck* leads to individual moments of inertia slightly different from the intrinsic values shown here, but spot-check comparisons of this semi-analytical approach used by `QuickBeam` with numerical simulations of the actual scanning by `FEBeCoP` showed agreement at the 10^{-3} level for $\ell < 2500$ on the resulting pixel window functions for sky coverage varying from 40 to 100%.

In the various Galactic masks used here (Figs. 12–13) the contribution of the unmasked pixels to the total effective window function departs from the full-sky average (which is not included in the effective beam window functions), and we therefore expect a different effective transfer function for each mask. We ignored this dependence and mitigated its effect by using transfer functions computed with the Galactic mask G60 which retains an effective sky fraction (including the mask apodization) of $f_{sky} = 60\%$, not too different from the sky fractions f_{sky} between 41 and 70% (see Sect. 3.2.2) used for computing the power spectra.

Figure 21 compares the impact of these two sources of uncertainty on the stated *Planck* statistical error bars for $\Delta\ell = 30$. It shows that, for $\ell < 1800$ where most of the information on ΛCDM lies, the error on the *TT* power spectra introduced by the sub-pixel effect and by the sky-coverage dependence are less than about 0.1%, and well below the statistical error bars of the binned C_ℓ. In the range $1800 \leq \ell \leq 2500$, which helps constrain one-parameter extensions to base ΛCDM (such as N_{eff}), the relative error can reach 0.4% (note as a comparison that the high-ℓ ACT experiment states a statistical error of about 3% on the bin $2340 \leq \ell \leq 2540$, Das et al. 2014). The bottom panel shows the Monte Carlo error model of the beam window functions, which provides negligible (ℓ-coupled) uncertainties. Even if this model is somewhat optimistic, since it does not include the effect of the ADC non-linearities and the colour-correction effect of beam measurements on planets (Planck Collaboration VII 2016), we note that even expanding them by a factor of 10 keeps them within the statistical uncertainty of the power spectra.

Modelling the uncertainties. As in the 2013 analysis, the beam uncertainty eigenmodes were determined from 100 (improved) Monte Carlo (MC) simulations of each planet observation used to measure the scanning beams, then processed through the same `QuickBeam` pipeline as the nominal beam to determine their effective angular transfer function $B(\ell)$. Thanks to the use of Saturn and Jupiter transits instead of the dimmer Mars used in 2013, the resulting uncertainties are now significantly smaller (Planck Collaboration VII 2016).

For each pair of frequency maps (and frequency-averaged beams) used in the present analysis, a singular-value decomposition (SVD) of the correlation matrix of 100 Monte Carlo based $B(\ell)$ realizations was performed over the ranges $[0, \ell_{max}]$ with $\ell_{max} = (2000, 3000, 3000)$ at (100, 143, 217 GHz), and the five leading modes were kept, as well as their covariance matrix (since the error modes do exhibit Gaussian statistics). We therefore have, for each pair of beams, five ℓ-dependent templates,
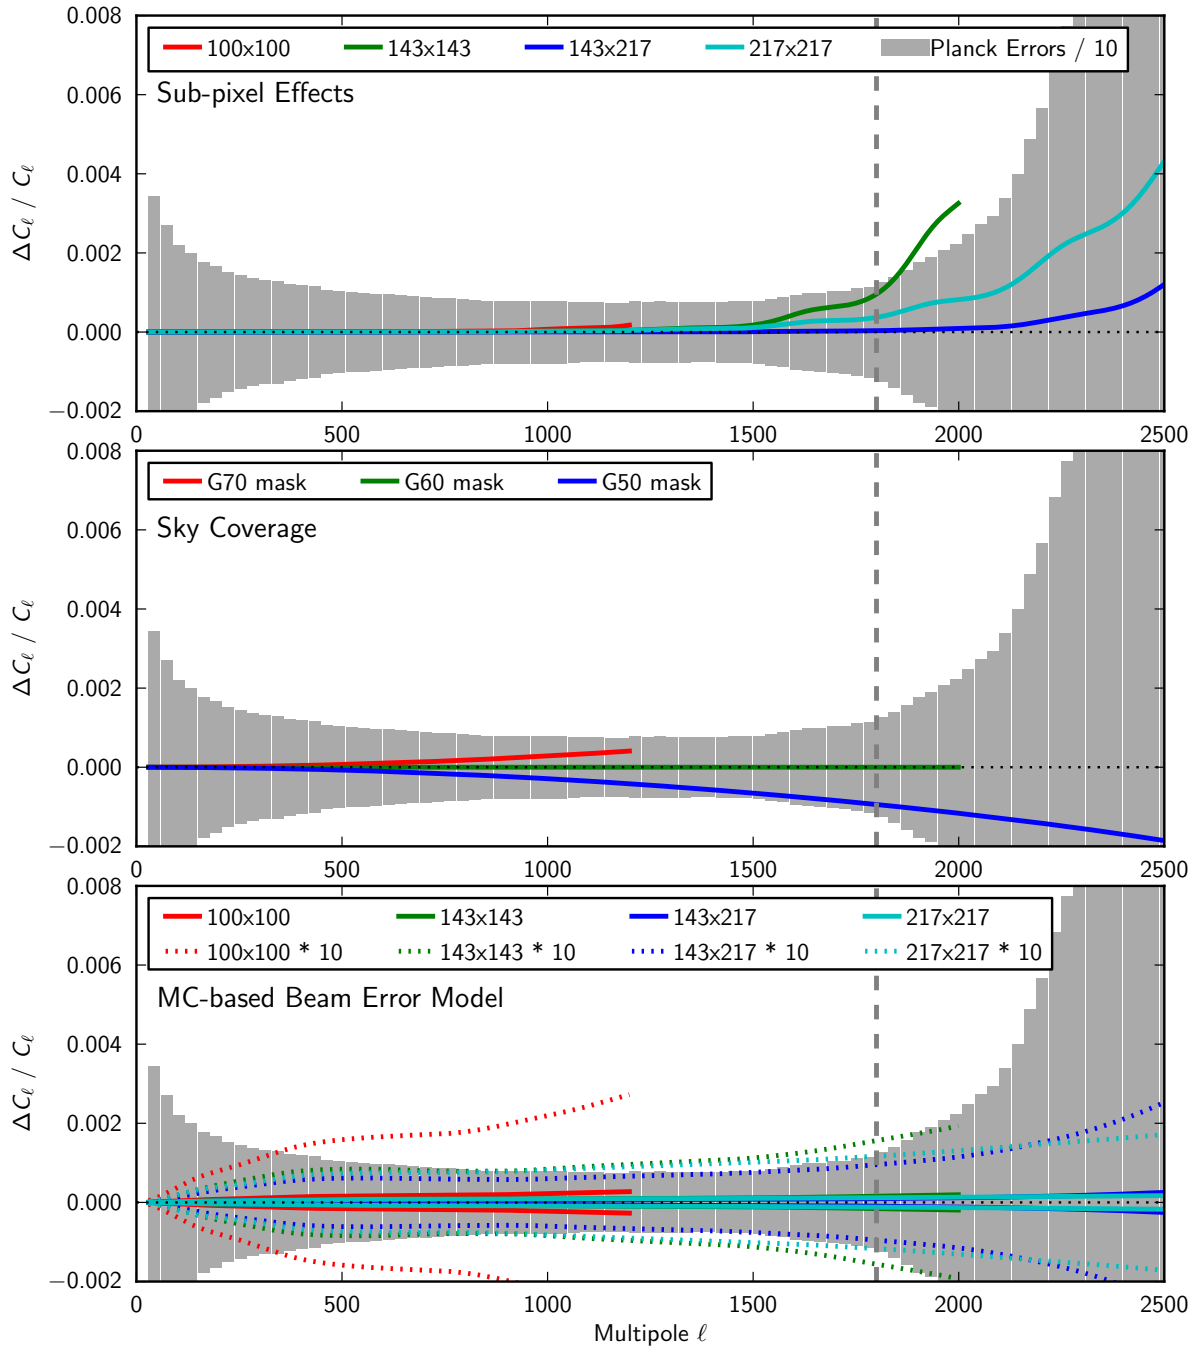
Fig. 21. Contribution of various beam-window-function-related errors and uncertainties to the C_ℓ relative error. In each panel, the grey histogram shows the relative statistical error on the *Planck* CMB *TT* binned power spectrum (for a bin width $\Delta\ell = 30$) *divided by 10*, while the vertical grey dashes delineate the range $\ell < 1800$ that is most informative for base ΛCDM. *Top*: estimation of the error made by ignoring the sub-pixel effects for a fiducial C_ℓ including the CMB and CIB contributions. *Middle*: error due to the sky mask, for the Galactic masks used in the *TT* analysis. *Bottom*: current beam window function error model, shown at 1σ (solid lines) and 10σ (dotted lines).

each associated with a Gaussian amplitude centred on 0, and a covariance matrix coupling all of them.

Including the beam uncertainties in the mixing matrix of Eq. (14) gives

$$\left(M_{ZW,\nu\times\nu'}^{XY}\right)_\ell(\theta_{\mathrm{inst}}) = \left(M_{ZW,\nu\times\nu'}^{XY,\mathrm{other}}\right)_\ell(\theta_{\mathrm{other}})\left(\Delta W_{\nu\times\nu'}^{ZW}\right)_\ell(\theta_{\mathrm{beam}}),$$

$$\left(\Delta W_{\nu\times\nu'}^{ZW}\right)_\ell(\theta_{\mathrm{beam}}) = \exp\sum_{i=1}^{5} 2\,\theta_{\nu\times\nu'}^{ZW,i}\left(E_{\nu\times\nu'}^{ZW,i}\right)_\ell, \qquad (41)$$

where $\left(\Delta W_{\nu\times\nu'}^{ZW}\right)_\ell(\theta_{\mathrm{beam}})$ stands for the beam error built from the eigenmodes $\left(E_{\nu\times\nu'}^{ZW,i}\right)_\ell$. The quadratic sum of the beam eigenmodes is shown in Fig. 21. This is much smaller (less than a percent) than the combined *TT* spectrum error bars. This contrasts with the 2013 case where the beam uncertainties were greater; for instance, for the 100, 143, and 217 GHz channel maps, the rms of the $W(\ell) = B(\ell)^2$ uncertainties at $\ell = 1000$ dropped from $(61, 23, 20) \times 10^{-4}$ to $(2.2, 0.84, 0.81) \times 10^{-4}$, respectively. The fact that beam uncertainties are sub-dominant in the total error
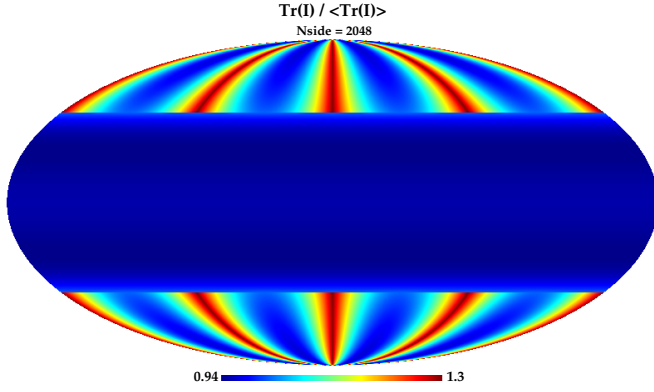
Tr(I) / <Tr(I)>
Nside = 2048

0.94 ▬▬▬ 1.3

Fig. 22. Map of the relative variations of the trace of the `HEALPix` pixel moment of inertia tensor at $N_{\text{side}} = 2048$ in Galactic coordinates.

budget is even more pronounced in polarization, where noise is higher. We use the beam modes computed from temperature data, combined with appropriate weights when used as parameters affecting the *TE* and *EE* spectra.

As in 2013, instead of including the beam error in the vector model, we include its contribution to the covariance matrix, linearizing the vector model so that

$$\left(C_{\nu\times\nu'}^{XY}\right)_\ell (\theta) = \left(C_{\nu\times\nu'}^{XY}\right)_\ell (\theta, \theta_{\text{beam}} = 0) + \left(\Delta W_{\nu\times\nu'}^{ZW}\right)_\ell (\theta_{\text{beam}}) \left(C_{\nu\times\nu'}^{XY}\right)_\ell^*, \tag{42}$$

where $\left(C_{\nu\times\nu'}^{XY}\right)_\ell^*$ is the fiducial spectrum *XY* for the pair of frequencies $\nu \times \nu'$ obtained using the best cosmological and foreground model. We can then marginalize over the beam uncertainty, enlarging the covariance matrix to obtain

$$\mathsf{C}_{\text{beam marg.}} = \mathsf{C} + \boldsymbol{C}^* \left\langle \Delta \boldsymbol{W} \Delta \boldsymbol{W}^\mathsf{T} \right\rangle \boldsymbol{C}^{*\mathsf{T}}, \tag{43}$$

where $\left\langle \Delta \boldsymbol{W} \Delta \boldsymbol{W}^\mathsf{T} \right\rangle$ is the Monte Carlo based covariance matrix, restricted to its first five eigenmodes.

In 2013, beam errors were marginalized for all the modes except the two greatest of the 100×100 spectrum. In the present release we instead marginalize over all modes in *TT*, *TE*, and *EE*. We also performed a test in which we estimated the amplitudes for all of the first five beam eigenmodes in *TT*, *TE*, and *EE*, and found no indication of any beam error contribution (see Sect. 4.1.3 and Fig. 35).

Temperature-to-polarization leakage. Polarization measurements are differential by nature. Therefore any unaccounted discrepancy in combining polarized detectors can create some leakage from temperature to polarization (Hu et al. 2003). Sources of such discrepancies in the current HFI processing include, but are not limited to: differences in the scanning beams that are ignored during the map-making; differences in the noise level, because of the individual inverse noise weighting used in HFI; and differences in the number of valid samples.

For this release, we did not attempt to model and remove a priori the form and amplitude of this coupling between the measured *TT*, *TE*, and *EE* spectra; we rather estimate the residual effect by fitting *a posteriori* in the likelihood some flexible template of this coupling, parameterized by some new nuisance parameters that we now describe.

The temperature-to-polarization leakage due to beam mismatch is *assumed* to affect the spherical harmonic

coefficients via

$$a_{\ell m}^T \quad \longrightarrow \quad a_{\ell m}^T, \tag{44a}$$

$$a_{\ell m}^E \quad \longrightarrow \quad a_{\ell m}^E + \varepsilon(\ell) a_{\ell m}^T, \tag{44b}$$

and, for each map, the spurious polarization power spectrum $C_\ell^{XY} \equiv \sum_m a_{\ell m}^X a_{\ell m}^{Y*}/(2\ell + 1)$ is modelled as

$$\Delta C_\ell^{TE} = \varepsilon(\ell) C_\ell^{TT}, \tag{45a}$$

$$\Delta C_\ell^{EE} = \varepsilon^2(\ell) C_\ell^{TT} + 2\varepsilon(\ell) C_\ell^{TE}. \tag{45b}$$

Here ε_ℓ is a polynomial in multipole ℓ determined by the effective beam of the detector-assembly measuring the polarized signal. Considering an effective beam map $b(\hat{\boldsymbol{n}})$ (rotated so that it is centred on the north pole), its spherical harmonic coefficients are defined as $b_{\ell m} \equiv \int \mathrm{d}\hat{\boldsymbol{n}}\, b(\hat{\boldsymbol{n}})\, Y_{\ell m}^*(\hat{\boldsymbol{n}})$. As a consequence of the *Planck* scanning strategy, pixels are visited approximately every six months, with a rotation of the focal plane by 180°, and we expect $b_{\ell m}$ to be dominated by even values of m, and especially the modes $m = 2$ and 4, which describe the beam ellipticity. As noted by, e.g., Souradeep & Ratra (2001) for elliptical Gaussian beams, the *Planck*-HFI beams for a detector d obey

$$b_{\ell m}^{(d)} \simeq \beta_m^{(d)} \ell^m b_{\ell 0}^{(d)}. \tag{46}$$

We therefore fit the spectra using a fourth-order polynomial

$$\varepsilon(\ell) = \varepsilon_0 + \varepsilon_2 \ell^2 + \varepsilon_4 \ell^4, \tag{47}$$

treating the coefficients ε_0, ε_2, and ε_4 as nuisance parameters in the MCMC analysis. Tests performed on detailed simulations of *Planck* observations with known mismatched beams have shown that Eqs. (45) and (47) describe the power leakage due to beam mismatch with an accuracy of about 20% in the ℓ range $100-2000$.

The equations above suggest that the same polynomial ε can describe the contamination of the *TE* and *EE* spectra for a given pair of detector sets. But in the current `Plik` analysis, the *TE* cross-spectrum of two different maps a and b is the inverse-variance-weighted average of the cross-spectra $T_a E_b$ and $T_b E_a$, while *EE* is simply $E_a E_b$. In addition, the temperature maps include the signal from SWBs, which is obviously not the case for the *E* maps. We therefore describe the *TE* and *EE* corrections by different ε parameters. Similarly, we treated the parameters for the *EE* cross-frequency spectra as being uncorrelated with the parameters for the auto-frequency ones.

The leakage is driven by the discrepancy between the individual effective beams $b_{\ell m}^{(d)}$ making up a detector assembly, coupled with the details of the scanning strategy and relative weight of each detector. If we assumed a perfect knowledge of the beams, precise – but not necessarily accurate – numerical predictions of the leakage would be possible. However, we preferred to adopt a more conservative approach in which the leakage was free to vary over a range wide enough to enclose the true value. On the other hand, in order to limit the unphysical range of variations permitted by so many nuisance parameters, we need priors on the ε_m terms used in the Monte Carlo explorations. We assume Gaussian distributions of zero mean with a standard deviation σ_m representative of the dispersion found in simulations of the effect with realistic instrumental parameters. We found $\sigma_0 = 1 \times 10^{-5}$, $\sigma_2 = 1.25 \times 10^{-8}$, and $\sigma_4 = 2.7 \times 10^{-15}$. This procedure ignores correlations between terms of different m, and is therefore likely substantially too permissive.
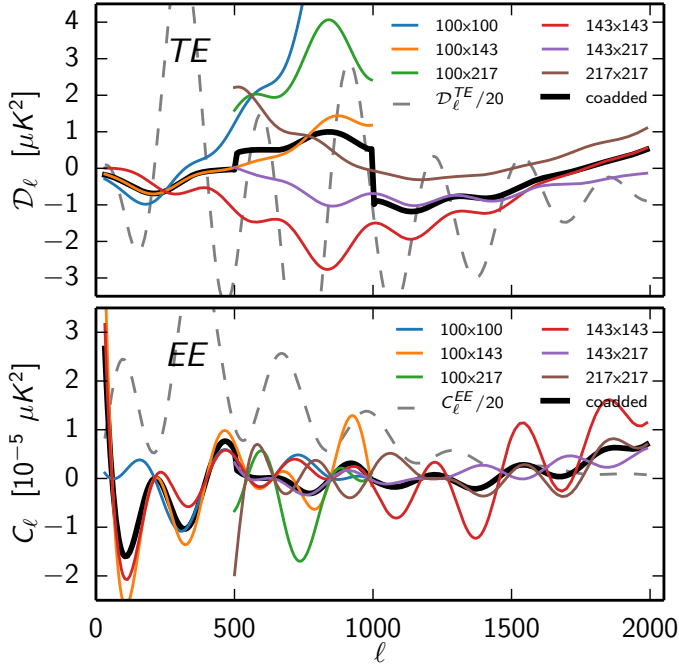
Fig. 23. Best fit of the power spectrum leakage due to the beam mismatch for *TE* (Eq. (45a), *upper panel*) and *EE* (Eq. (45b), *lower panel*). In each case, we show the correction for individual cross-spectra (coloured thin lines) and the co-added correction (black line). The individual cross-spectra corrections are only shown in the range of multipoles where the data from each particular pair is used. The individual correction can be much higher than the co-added correction. The co-added correction is dominated by the best S/N pair for each multipole. For example, up to $\ell = 500$, the *TE* co-added correction is dominated by the 100×143 contribution. The grey dashed lines show the *TE* and *EE* best-fit spectra rescaled by a factor of 20, to give an idea of the location of the model peaks.

Another way of deriving the beam leakage would be to use a cosmological prior, i.e., by finding the best fit when holding the cosmological parameters fixed at their best-fit values for base ΛCDM. Figure 23 shows the result of this procedure for the cross-frequency pairs. The figure also shows the implied correction for the co-added spectra. This correction is dominated by the pair with the highest S/N at each multipole. The fact that different sets are used in different ℓ-ranges leads to discontinuities in the correction template of the co-added spectrum. As can be seen in the figure, the co-added beam-leakage correction, of order μK^2, is much smaller than the individual corrections, which partially compensate each other on average (but improve the agreement between the individual polarized cross-frequency spectra).

It is shown in Appendix C.3.5 that neither procedure is fully satisfactory. The cosmological prior leads to nuisance parameters that vastly exceed the values allowed by the physical priors, and the physical priors are clearly overly permissive (leaving the cosmological parameters unchanged but with doubled error bars for some parameters). In any case, the agreement between the different cross-spectra remains much poorer in polarization than in temperature (see Sect. 4.4, Fig. 40, and Appendix C.3.5); they present oscillatory features similar to the ones produced by our beam leakage model, but the model is clearly not sufficient. For lack of a completely satisfactory global instrumental model, this correction is only illustrative and it is not used in the baseline likelihood.
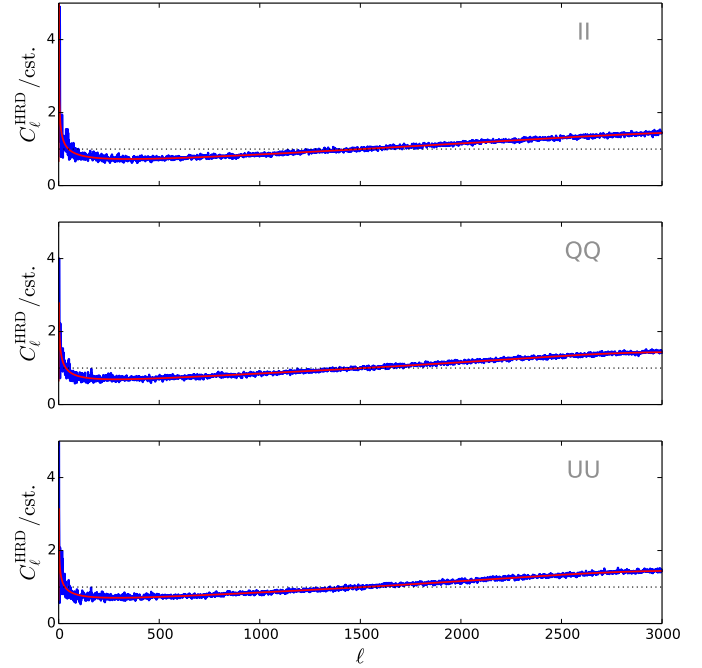
Fig. 24. Deviations from a white noise power spectrum induced by noise correlations. We show half-ring difference power spectra for 100 GHz half-mission 1 maps (blue lines) of Stokes parameters *I* (*top panel*), *Q* (*middle panel*), and *U* (*bottom panel*). The best-fitting analytical model of the form Eq. (48) is over-plotted in red.

3.4.4. Noise modelling

To predict the variance of the empirical power spectra, we need to model the noise properties of all maps used in the construction of the likelihood. As described in detail in Planck Collaboration VII (2016) and Planck Collaboration VIII (2016), the *Planck* HFI maps have complicated noise properties, with noise levels varying spatially and with correlations between neighbouring pixels along the scanning direction.

For each channel, full-resolution noise variance maps are constructed during the map-making process (Planck Collaboration VIII 2016). They provide an approximation to the diagonal elements of the true $n_{\mathrm{pix}} \times n_{\mathrm{pix}}$ noise covariance matrix for Stokes parameters *I* (temperature only), or *I*, *Q*, and *U* (temperature and polarization). While it is possible to capture the anisotropic nature of the noise variance with these objects, noise correlations between pixels remain unmodelled. To include deviations from a white-noise power spectrum, we therefore make use of half-ring difference maps. Choosing the 100 GHz map of the first half-mission as an example, we show the scalar (spin-0) power spectra of the three temperature and polarization maps in Fig. 24, rescaled by arbitrary constants. We find that the logarithm of the HFI noise power spectra as given by the half-ring difference maps can be accurately parameterized using a fourth-order polynomial with an additional logarithmic term,

$$\log(C_\ell^{\mathrm{HRD}}) = \sum_{i=0}^{4} \alpha_i\, \ell^i + \alpha_5 \log(\ell + \alpha_6). \qquad (48)$$

Since low-frequency noise and processing steps like deglitching leave residual correlations between both half-ring maps, noise estimates derived from their difference are biased low, at the percent level at high-ℓ (where it was first detected and understood, see Planck Collaboration VI 2014). We correct for this
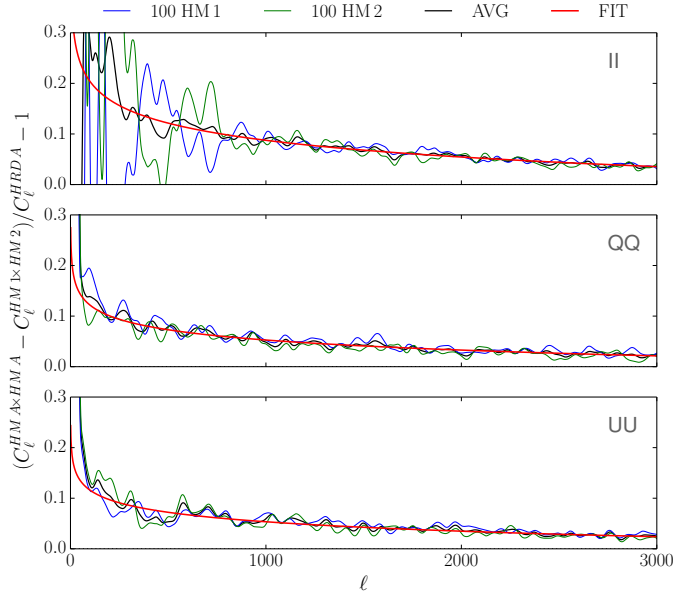
Fig. 25. Difference between auto and cross-spectra for the 100 GHz half-mission maps, divided by the noise estimate from half-ring difference maps (blue and green lines). Noise estimates derived from half-ring difference maps are biased low. We fit the average of both half-mission curves (black line) with a power law model (red line). The analysis procedure is applied to the Stokes parameter maps *I*, *Q*, and *U* (*top to bottom*). All data power spectra are smoothed.

effect by comparing the difference of auto-power-spectra and cross-spectra (assumed to be free of noise bias) at a given frequency with the noise estimates obtained from half-ring difference maps. As shown in Fig. 25, we use a power-law model with free spectral index to fit the average of the ratios of the first and second half-mission results to the half-ring difference spectrum, using the average to nullify chance correlations between signal and noise:

$$C_\ell^{\mathrm{bias}} = \alpha_0 \, \ell^{\alpha_1} + \alpha_2. \tag{49}$$

At a multipole moment of $\ell = 1000$, we obtain correction factors for the temperature noise estimate obtained from half-ring difference maps of 9%, 10%, and 9% at 100, 143, and 217 GHz, respectively.

In summary, our HFI noise model is obtained as follows. For each map, we capture the anisotropic nature of the noise amplitude by using the diagonal elements of the pixel-space noise covariance matrix. The corresponding white-noise power spectrum is then modulated in harmonic space using the product of the two smooth fitting functions given in Eqs. (48) and (49).

Correlated noise between detectors. If there is some correlation between the noise in the different cuts in our data, the trick of only forming effective frequency-pair power spectra from cross-spectra to avoid the noise biases fails. In 2013, we evaluated the amplitude of such correlated noise between different detsets. The correlation, if any, was found to be small, and we estimated its effect on the cosmological parameter fits to be negligible. As stated in Sect. 3.2.1, the situation is different for the 2015 data. Indeed, we now detect a small but significant correlated noise contribution between the detsets. This is the reason we change our choice of data to estimate the cross-spectra, from detsets to half-mission maps. The correlated noise appears to be much less significant in the latter.
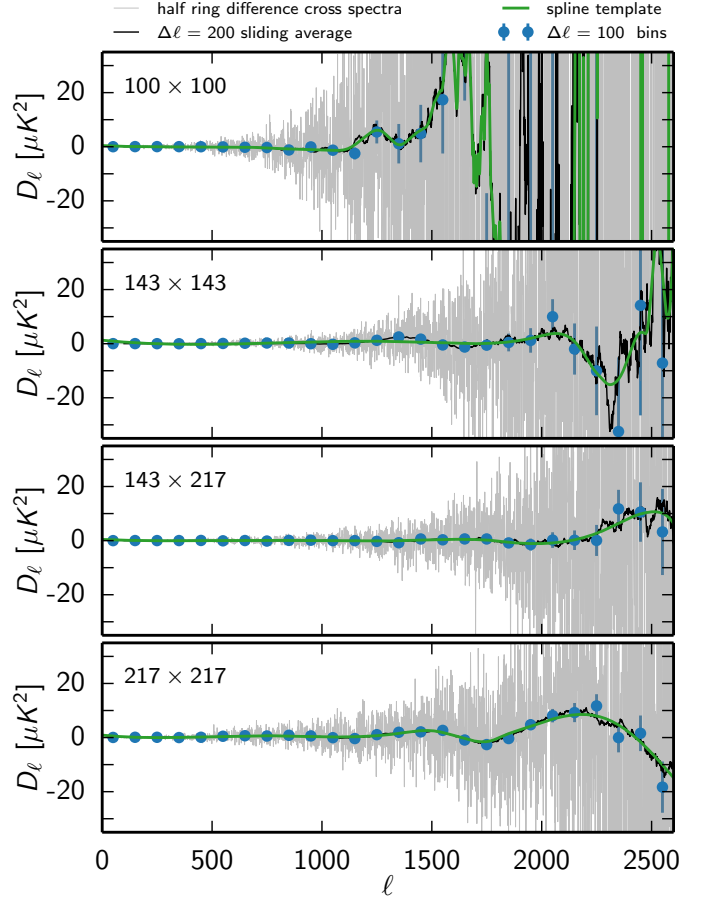

Fig. 26. Correlated noise model. In grey are shown the cross-detector TT spectra of the half-ring difference maps. The black line show the same, smoothed by a $\Delta\ell = 200$ sliding average, while the blue data points are a $\Delta\ell = 100$ binned version of the grey line. Error bars simply reflect the scatter in each bin. The green line is the spline-smoothed version of the data that we use as our correlated noise template.

To estimate the amount of correlated noise in the data, we measured the cross-spectra between the half-ring difference maps of all the individual detsets. The cross-spectra are then summed using the same inverse-variance weighting that we used in 2013 to form the effective frequency-pair spectra. Figure 26 shows the spectra for each frequency pair. All of these deviate significantly from zero. We build an effective correlated noise template by fitting a smoothing spline on a $\Delta\ell = 200$ sliding average of the data. Given the noise level in polarization, we did not investigate the possible contribution of correlated noise in EE and TE.

Section 4.1.1 shows that when these correlated noise templates are used, the results of the detsets likelihood are in excellent agreement with those based on the baseline, half-mission one.

3.5. Covariance matrix structure

The construction of a Gaussian approximation to the likelihood function requires building covariance matrices for the pseudo-power spectra. Mathematically exact expressions exist, but they are prohibitively expensive to calculate numerically at *Planck* resolution (Wandelt et al. 2001); we thus follow the approach taken in Like13 and make use of analytical approximations

(Hansen et al. 2002; Hinshaw et al. 2003; Efstathiou 2004; Challinor & Chon 2005).

For our baseline likelihood, we calculate covariance matrices for all 45 unique detector combinations that can be formed out of the six frequency-averaged half-mission maps at 100, 143, and 217 GHz. To do so, we assume a fiducial power spectrum that includes the data variance induced by the CMB and all foreground components described in Sect. 3.3; this variance is computed assuming these components are Gaussian-distributed. The effect of this approximation regarding Galactic foregrounds is tested by means of simulations in Sect. 3.6. The fiducial model is taken from the best-fit cosmological and foreground parameters; since they only become available after a full exploration of the likelihood, we iteratively refine our initial guess. As discussed in Sect. 3.1, the data vector used in the likelihood function of Eq. (13) is constructed from frequency-averaged power spectra. Following Like13, for each polarization combination, we therefore build averaged covariance matrices for the four frequencies $\nu_1, \nu_2, \nu_3, \nu_4$,

$$\mathrm{Var}(\hat{C}_\ell^{XY\,\nu_1,\nu_2}, \hat{C}_{\ell'}^{ZW\,\nu_3,\nu_4}) = \sum_{\substack{(i,j)\in(\nu_1,\nu_2)\\(p,q)\in(\nu_3,\nu_4)}} w_\ell^{XY\,i,j} w_{\ell'}^{ZW\,p,q}$$
$$\times \mathrm{Var}\left(\hat{C}_\ell^{XY\,i,j}, \hat{C}_{\ell'}^{ZW\,p,q}\right), \tag{50}$$

where $X, Y, Z, W \in \{T, E\}$, and $w^{XY\,i,j}$ is the inverse-variance weight for the combination (i, j), computed from

$$w_\ell^{XY,j} \propto 1/\mathrm{Var}\left(\hat{C}_\ell^{XY\,i,j}, \hat{C}_\ell^{XY\,i,j}\right), \tag{51}$$

and normalized to unity. For the averaged $XY = TE$ covariance (and likewise for $ZW = TE$), the sum in Eq. (50) must be taken over the additional permutation $XY = ET$. That is, the two cases where the temperature map of channel i is correlated with the polarization map of channel j and vice versa are combined into a single frequency-averaged covariance matrix. These matrices are then combined to form the full covariance used in the likelihood,

$$\mathsf{C} = \begin{pmatrix} C^{TTTT} & C^{TTEE} & C^{TTTE} \\ C^{EETT} & C^{EEEE} & C^{EETE} \\ C^{TETT} & C^{TEEE} & C^{TETE} \end{pmatrix}, \tag{52}$$

where the individual polarization blocks are constructed from the frequency-averaged covariance matrices of Eq. (50) (Like13).

Appendix C.1.1 provides a summary of the equations used to compute temperature and polarization covariance matrices and presents a validation of the implementation through direct simulations. Let us note that, for the approximations used in the analytical computation of the covariance matrix to be precise, the mask power spectra have to decrease quickly with multipole moment ℓ; this requirement gives rise to the apodization scheme discussed in Sect. 3.2.2. In the presence of a point-source mask, however, the condition may no longer be fulfilled, reducing the accuracy of the approximations assumed in the calculation of covariance matrices. We discuss in Appendix C.1.4 the heuristic correction we developed to restore the accuracy, which is based on direct simulations of the effect.

3.6. FFP8 simulations

In order to validate the overall implementation and our approximations, we generated 300 simulated HFI half-mission map sets in the frequency range 100 to 217 GHz, which we analysed like the real data. For the CMB, we created realizations of the ΛCDM

Table 13. Shifts of parameters over 300 TT simulations.

Parameter	300 sims	r_{A}^{30}	r_{A}^{65}	r_{A}^{100}
$\Omega_{\mathrm{b}}h^2$	0.27	0.62	0.50	0.52
$\Omega_{\mathrm{c}}h^2$	-0.71	-0.65	-0.44	0.00
θ	1.48	1.67	1.67	1.29
τ	-0.57	-0.38	-0.56	-0.38
$\ln\left(10^{10}A_{\mathrm{s}}\right)$	-0.70	-0.52	-0.65	-0.35
n_{s}	1.86	1.87	1.46	0.78
A_{CIB}^{217}	-0.99	-1.09	-1.44	-1.34
gal_{545}^{100}	0.31	0.13	-0.09	0.04
gal_{545}^{143}	0.40	-0.21	-0.23	-0.19
$\mathrm{gal}_{545}^{143-217}$	-0.22	-0.35	0.36	0.22
gal_{545}^{217}	1.61	1.48	2.19	2.04

Notes. Shifts are given in units of the posterior width rescaled by $1/\sqrt{300}$. If the parameters were uncorrelated, 68% of the shifts would be expected to lie within $\pm 1\sigma$. The effect of varying the value of $\ell_{\min}$ is measured on the likelihood of the average spectra over 300 realizations, labelled $r_{\mathrm{A}}^{\ell_{\min}}$. A significant decrease of the bias on n_{s} is obtained by not including low-ℓ multipoles, at the cost, however, of a degradation in the determination of the foreground amplitudes A_{CIB}^{217}, $\mathrm{gal}_{545}^{143-217}$, and gal_{545}^{217}.

model with the best-fit parameters obtained in this paper. After convolving the CMB maps with beam and pixel window functions, we superimposed CIB, dust, and noise realizations from the FFP8 simulations (Planck Collaboration XII 2016) that capture both the correlation structure and anisotropy of foregrounds and noise. We then computed power spectra using the set of frequency-dependent masks described in Sect. 3.2.2 and created the corresponding `Plik` TT likelihood. We modified the shape of the foreground spectra to fit the FFP8 simulations, but kept the parameterization used on the data. In the case of dust, we used priors similar to those used on data. Furthermore, in the following the dust amplitude parameter is named $\mathrm{gal}_{545}^{\nu \times \nu'}$. We then ran an MCMC sampler to derive the cosmological and foreground parameters posterior distributions for all dataset realizations.

For each simulation, we computed the shift of the derived posterior mean parameters with respect to the input cosmology, normalized by their posterior widths σ_{post}. When a Gaussian prior with standard deviation σ_{prior} is used, we rescale σ_{post} by $[1 - \sigma_{\mathrm{post}}^2/\sigma_{\mathrm{prior}}^2]^{1/2}$; this is the case for τ and for the Galactic dust amplitudes gal_{545}^{ν} in the four cross-frequency channels used. In Fig. 27, we show histograms of the shifts we found for all 300 simulations for the six baseline cosmological parameters, as well as the FFP8 CIB and galactic dust amplitudes. As shown in the figure, we recover the input parameters with little bias and a scatter of the normalized parameter shifts around unity. The p-values of the Kolmogorov-Smirnov test that we ran are given in the legend and we do not detect significant departures from normality. The average reduced χ^2 for the histograms of Fig. 27 is equal to 1.02.

Table 13 (second column) compiles the average shifts of Fig. 27, but in order to gauge whether they are as small as expected for this number of simulations (assuming no bias), the shifts are expressed in units of the posterior width rescaled by $1/\sqrt{300}$. We note that the shift of the average is above one (scaled)σ in three cases out of a total of 11 parameters (68% of the Δs would be expected to lie within 1σ if the parameters
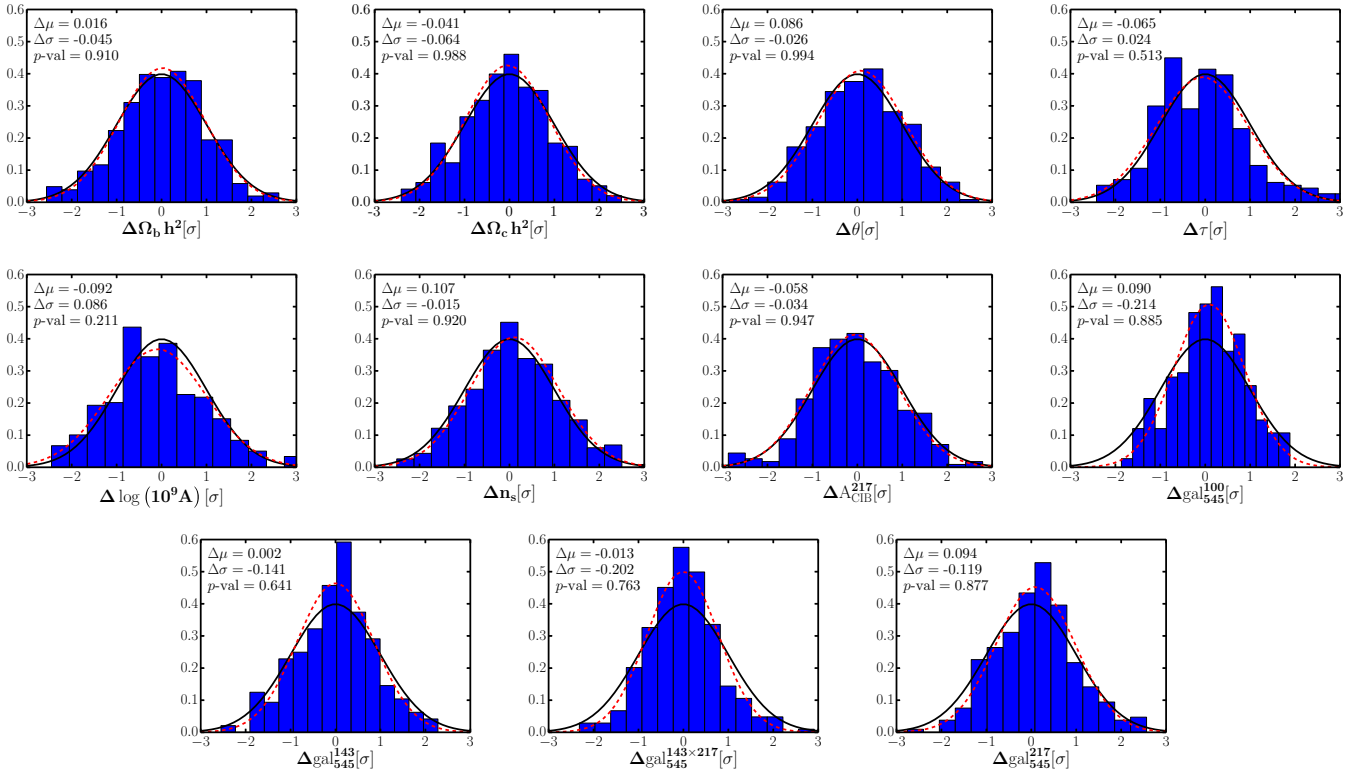
Fig. 27. `Plik` parameter results on 300 simulations for the six baseline cosmological parameters, as well as the FFP8 CIB and Galactic dust amplitudes. The simulations include quite realistic CMB, noise, and foregrounds (see text). The distributions of inferred posterior mean parameters are centred around their input values with the expected scatter. Indeed the dotted red lines show the best-fit Gaussian for each distribution, with a mean shift, $\Delta\mu$, and a departure $\Delta\sigma$ from unit standard deviation given in the legend; both are close to zero. These best fits are thus very close to Gaussian distributions with zero shift and unit variance, which are displayed for reference as black lines. The legend gives the numerical value of $\Delta\mu$ and $\Delta\sigma$, as well as the p-values of a Kolmogorov-Smirnov test of the histograms against a Gaussian distribution shifted from zero by $\Delta\mu$ and with standard deviation shifted from unity by $\Delta\sigma$. This confirms that the distributions are consistent with Gaussian distributions with zero mean and unit standard deviation, with a small offset of the mean.

were uncorrelated), with θ, $n_{\rm s}$, and gal$_{545}^{217}$ at the 1.7, 2.0, and 1.5 (scaled)σ level, respectively.

Before proceeding, let us note that an estimate (third column) of these shifts is obtained by simply computing the shift from a single likelihood using as input the average spectra of the 300 simulations. This effectively reduces cosmic variance and noise amplitude by a factor $\sqrt{300}$ and, more importantly, it decreases the cost and length of the overall computation, enabling additional tests. These shift estimates are noted $r_{\rm A}$. The table shows that significant improvement in the determination of $n_{\rm s}$ is obtained by removing low-ℓ multipoles. Indeed, Cols. 4 and 5 of Table 13 show the variation of the shift when the $\ell_{\rm min}$ of the high-ℓ likelihood is increased from 30 to 65 and 100. The shift in $n_{\rm s}$ is decreased by a factor two, while the decrease in the number of bins per cross-frequency spectrum is only reduced from 199 to 185 (having little impact on the size of the covariance matrix of cosmological parameters).

These changes with $\ell_{\rm min}$ therefore trace the small biases back to the lowest-ℓ bins. It suggests that the Gaussian approximation used in the high-ℓ likelihood starts to become mildly inaccurate at $\ell = 30$. Indeed, even if noticeable, this effect would contribute at most a 0.11σ bias on $n_{\rm s}$. This is further confirmed by the lack of a detectable effect found in Sect. 5.1 when varying the hybridization scale in TT between `Commander` and `Plik`. However, the exclusion of low-ℓ information degrades our ability to accurately reconstruct the foreground amplitudes $A_{\rm CIB}^{217}$, gal$_{545}^{143-217}$, and gal$_{545}^{217}$. Indeed, the dust spectral amplitudes in the

143×217 and 217×217 channels are highest at low multipoles, and the CIB spectrum in the range $30 \le \ell \le 100$ also adds substantial information.

In spite of this low-ℓ trade-off between an accurate determination of $n_{\rm s}$ on the one hand and $A_{\rm CIB}^{217}$, gal$_{545}^{143-217}$, and gal$_{545}^{217}$ on the other, we can conclude that the `Plik` implementation is behaving as expected and can be used for actual data analysis.

Appendix C.2 extends this conclusion to the joint `PlikTT`, EE, TE likelihood case.

3.7. End-to-end simulations

While the previous section validated our methodology, our approximations, and the overall implementation, this does not yet give the sensitivity to residual systematic uncertainties undetected by data consistency checks. These are by their very nature very much more difficult to address realistically, since, when an effect is detected and sufficiently well understood, it can be modelled and is corrected for, in general at the TOI-processing stage; only the uncertainty of the correction needs to be addressed. Still, HFI has developed a complete model of the instrument which contains all identified systematic effects and enables realistic simulation of the instrumental response. We have therefore generated a number of full-mission time streams which we have then processed with the DPC TOI processing pipeline in order to create map datasets as close to instrumental reality as we can in order to assess the possible impact of low-level residual

Table 14. End-to-end parameter shifts for a single realization of CMB and foregrounds, along with five different noise realizations. Shifts are computed with respect to those obtained without noise and with instrumental effects turned off.

Parameter	1	2	3	4	5	Mean	Median	σ_{FFP8}
$\Omega_{\text{b}}h^2$	0.32	0.09	−0.3	−0.02	−0.17	−0.01	−0.02	0.42
$\Omega_{\text{c}}h^2$	−0.07	−0.20	−0.22	−0.22	0.30	−0.08	−0.20	0.35
θ	0.18	0.08	0.24	−0.12	0.15	0.10	0.14	0.45
τ	0.08	0.19	−0.02	0.006	−0.21	0.01	0.005	0.25
$\ln\left(10^{10}A_{\text{s}}\right)$	0.04	0.10	−0.11	−0.11	−0.21	−0.05	−0.11	0.25
n_{s}	0.11	0.06	0.22	−0.13	−0.64	−0.07	0.05	0.40

Notes. In columns labelled 1 to 5: shifts of the cosmological parameters of the five noise realizations of the end-to-end simulations with respect to those obtained for the simulations without noise and with instrumental effects turned off, normalized by the end-to-end simulations' posterior widths. In columns labelled "Mean" and "Median": the corresponding mean and median. In column labelled "σ_{FFP8}": the standard deviation of the distribution obtained from the cosmological parameter shifts of 100 FFP8 simulations, varying the noise only with respect to the cosmological parameters of the CMB only.

Table 15. End-to-end parameter shifts for four different CMB realizations but comprising four pairs of realizations with the same noise realization with respect to those obtained without noise and with instrumental effects turned off.

Parameter	4	5	6	7	8	Δ_{5-6}	Δ_{4-7}	Δ_{4-8}	Δ_{7-8}	σ_{FFP8}
$\Omega_{\text{b}}h^2$	−0.02	−0.17	−0.25	−0.53	0.29	0.08	0.51	0.31	0.83	0.34
$\Omega_{\text{c}}h^2$	−0.22	0.30	−0.20	0.52	−0.59	0.50	0.74	0.37	1.11	0.30
θ	−0.12	0.15	−0.34	−1.06	0.41	0.48	0.94	0.52	1.47	0.35
τ	0.006	−0.21	−0.16	−0.17	−0.16	0.04	0.18	0.16	0.01	0.175
$\ln\left(10^{10}A_{\text{s}}\right)$	−0.11	−0.21	−0.29	−0.15	−0.34	0.07	0.04	0.24	0.20	0.19
n_{s}	−0.13	−0.64	−0.39	−0.86	0.24	0.24	0.72	0.37	1.03	0.35

Notes. In columns labelled 4 to 8: shifts of the cosmological parameters of the end-to-end simulations with respect to those obtained for the simulations without noise and with instrumental effects turned off, normalized by the end-to-end simulations' posterior widths. We point out that realizations numbered 4 and 5 are common to the sets of Tables 14 and 15. In columns labelled "Δ_{5-6}" to "Δ_{7-8}", the absolute differences in the shifts within pairs of realizations having different CMB but the same foregrounds and noise realizations. In column labelled "σ_{FFP8}": the standard deviation of the distribution obtained from the cosmological parameter shifts of 100 FFP8 simulations, varying the CMB only but keeping the same FFP8 noise realization, with respect to the cosmological parameters of the CMB only.

instrumental systematics, the effects of which might have remained undetected otherwise.

In this section, we report on the shifts in the values of the cosmological and foreground parameters induced by these specific residual systematic effects, comparing the results of a TT likelihood analysis for two overlapping sets of five simulations:

1. five simulations of maps at 100 GHz, 143 GHz and 217 GHz, for a single realization of the CMB and of the foregrounds but for five different realizations of the noise,
2. five simulations of maps at 100 GHz, 143 GHz and 217 GHz, composed of four CMB realizations, two noise realizations, a single realization of the foregrounds, but forming four pairs of realizations having the same noise but different CMB.

These simulations sum to a total of eight distinct simulations and are numbered from 1 to 8 in Tables 14 and 15. To be more explicit, in the second set, among simulations numbered 4 to 8, simulations 4, 7 and 8 have different CMB, but the same noise as each other. Simulations 5 and 6 have different CMB, and the same noise as each other, but different from simulations 4, 7 and 8. Realizations 4 and 5, having the same CMB but different noise, are common to the two sets of five realizations.

Each of these have been performed twice, with the end-to-end (instrument plus TOI processing) pipeline and noise contribution switched either on or off. End-to-end simulations are computationally very costly, (typically a week for each

simulated mission dataset) and hence only a few realizations were generated).

As explained in Sect. 5.4 of Planck Collaboration VII (2016), the end-to-end simulations are created by feeding the TOI processing pipeline with simulated data to evaluate and characterize the overall transfer function and the respective contribution of each individual effect on the determination of the cosmological parameters. Simulated TOIs are produced by applying the real mission scanning strategy to a realistic input sky specified by the *Planck* Sky Model (PSM; Delabrouille et al. 2013) containing a lensed CMB realization, galactic diffuse foregrounds, and the dipole components. To this sky-scanned TOI, we add a white-noise component, representing the phonon and photon noises. The very low-temporal-frequency thermal drift seen in the real data is also added to the TOI. The noisy sky TOI is then convolved with the appropriate bolometer transfer functions. Another white-noise component, representing Johnson noise and read-out noise, is also added. Simulated cosmic rays using the measured glitch rates, amplitudes, and shapes are added to the TOI. This TOI is interpolated to the electronic HFI fast-sampling frequency. It is then converted from analogue to digital using a simulated non-linear analogue-to-digital converter (ADC). Identified 4 K cooler spectral lines are added to the TOI. Both effects (ADC and 4 K lines) are derived from the measured in-flight behaviour. The TOI finally goes through the data compression/decompression algorithm used for communication between the *Planck* satellite and Earth. The simulated TOI

is then processed in the same way as the real mission data for cleaning and systematic error removal, calibration, destriping, and map-making.

Some limitations of the current end-to-end approach follow. No pointing error is included, although previous (dedicated) simulations suggest that this has negligible effect. In addition, this effect was included in the dedicated simulations performed to assess the precision of the beam recovery procedure. The first step of the TOI processing is to correct the ADC non-linearity (ADC NL). For the flight data, the ADC NL was determined by using HFI's measured signal at the end of the HFI mission, with the instrument's cooling system switched off and an instrument temperature equal to 4 K. This determination relied on supposing the signal to be perfect white noise and therefore to correspond to the distortions brought in by the ADC. In the current implementation, we assume perfect knowledge of ADC NL and 4 K lines. This is of course not true for the real data and future end-to-simulations, accompanying *Planck*'s next data release, will improve our model of this effect. After ADC NL correction, the signal is converted to volts. Deglitching is then performed by flagging glitch heads and by using glitch tail time-lines. This enables the creation of the thermal baseline which is used for signal demodulation. The thermal baseline and glitch tails are subtracted, the signal is converted to watts, and the 4 K lines are removed. The resulting signal is then deconvolved by the bolometer transfer functions. We do not include uncertainties in the glitch tail shape used in the deglitching procedure, i.e., the templates are the same for the simulations and the processing; but here again, previous studies suggest any difference is a small effect.

The analysis of these sets of end-to-end simulations, and of their counterparts for which all instrumental effects are turned off, is performed similarly to that of the simulations described in Sect. 3.6. Angular power spectra for all cross-half-missions and for all frequency combinations are computed using the *Planck* masks described in Sect. 3.2.2 and with the appropriate beam functions. Noise levels are evaluated as described in Sect. 3.4.4. Templates for galactic foregrounds (CO, free-free, synchrotron, thermal and spinning dust), the kinetic and thermal SZ effects, the cosmic infrared background, and radio and IR point sources are constructed based on the PSM input foreground maps. The covariance matrix is computed with the method outlined in Appendix C.1 with the aforementioned input CMB power spectrum, input foreground spectra, noise levels, beam functions and masks.

All sets of power spectra and the inverse covariance matrix are then binned and used in the likelihood analysis performed using an MCMC sampler together with Plik and PICO in order to determine the best fit cosmological parameters. The shifts in cosmological parameter values induced by the imperfect correction of instrumental effects by the TOI processing pipeline are then computed for the end-to-end simulations with respect to those obtained for the simulations without noise and with instrumental effects turned off, normalized by the end-to-end simulations' posterior widths. Comparing shifts computed in this way cancels out cosmic variance and chance correlations between the CMB and the foregrounds and are thus fully attributable to the instrument and to the noise, which cannot be disentangled, as well as to CMB-noise chance correlations. That is, those shifts probe directly the scatter and possible biases induced by residual systematics effects.

The mean and median shifts for the five simulations with a single CMB realization and a single foreground realization, but different noise realizations, are given in Table 14. In order

to verify that these shifts are within expectations, we computed the shifts in cosmological parameters for 100 FFP8 simulations, each with identical CMB signal but different FFP8 noise, with respect to the cosmological parameters obtained for the CMB only, normalized by their posterior widths. The standard deviations of the resulting distributions are given in the column labelled "σ_{FFP8}" of Table 14 and can be compared with the shifts obtained for the five end-to-end simulations. All shifts are within 1σ of the shifts expected from FFP8. In addition, there is no indication of any detectable bias. All shifts are thus compatible with scatter introduced by noise.

The shifts for the five realizations with four different CMB realizations, the same foregrounds, but comprising four pairs with the same noise realization, are given in Cols. 4 to 8 of Table 15. As mentioned at the beginning of this section, realizations numbered 4 and 5 are common to the sets of Tables 14 and 15. In the columns labelled "Δ_{5-6}" to "Δ_{7-8}", we computed the absolute differences in the shifts within pairs of realizations having different CMB but the same foreground and noise realizations. We compare these differences to the standard deviations of the distributions of cosmological parameter shifts of 100 FFP8 simulations, varying the CMB but keeping the same FFP8 noise realization, with respect to the cosmological parameters of the corresponding CMB but without noise (column labelled "σ_{FFP8}"). These distributions quantify the impact of CMB-noise correlations on the determination of the cosmological parameters. The Table shows that among all Δ's, 11 are within $1\sigma_{\text{FFP8}}$, 7 are within 1 to $2\sigma_{\text{FFP8}}$, 3 are within 2 to $3\sigma_{\text{FFP8}}$, 2 are within 3 to $4\sigma_{\text{FFP8}}$. 50% of the differences are within 1σ and 78% within 2σ. At the very worst, taking the example of Δ_{7-8} a cosmological parameter (θ) moves a total of 4σ, from -3σ to 1σ in units of σ_{FFP8} when the CMB is changed but the noise is left the same. This is rare but can be expected in a few percent of simulations. As in the case of the shifts listed in Table 14, there is thus no detectable bias, with all shifts compatible with those expected from FFP8.

In summary, we have detected no sign as yet of systematic biases of the cosmological parameters due to known low-level instrumental effects as corrected by the current HFI TOI processing pipeline. An increase in the significance of these tests is left for further work once the simulation chain is further optimized for more massive numerical work.

3.8. High-multipole reference results

This section describes the results obtained using the baseline Plik likelihood, in combination with a prior on the optical depth to reionization, $\tau = 0.07 \pm 0.02$ (referred to, in TT, as PlikTT+tauprior). The robustness and validation of these results (presented in Sect. 4) can therefore be assessed independently of any potential low-ℓ anomaly, or hybridization issues. The full low-ℓ + high-ℓ likelihood will be discussed in Sect. 5.

Figure 28 shows the high-ℓ co-added CMB spectra in TT, TE, and EE, and their residuals with respect to the best-fit ΛCDM model in TT (red line), both ℓ-by-ℓ (grey points) and binned (blue circles). The blue error bars per bin are derived from the diagonal of the covariance matrix computed with the best-fit CMB as fiducial model. The bottom sub-panels with residuals also show (yellow lines) the diagonal of the ℓ-by-ℓ covariance matrix, which may be compared to the dispersion of the individual ℓ determinations. Parenthetically, it provides graphical evidence that TT is dominated by cosmic variance through $\ell \approx 1600$, while TE is cosmic-variance dominated at $\ell \lesssim 160$ and $\ell \approx 260-460$. The jumps in the polarization diagonal-covariance
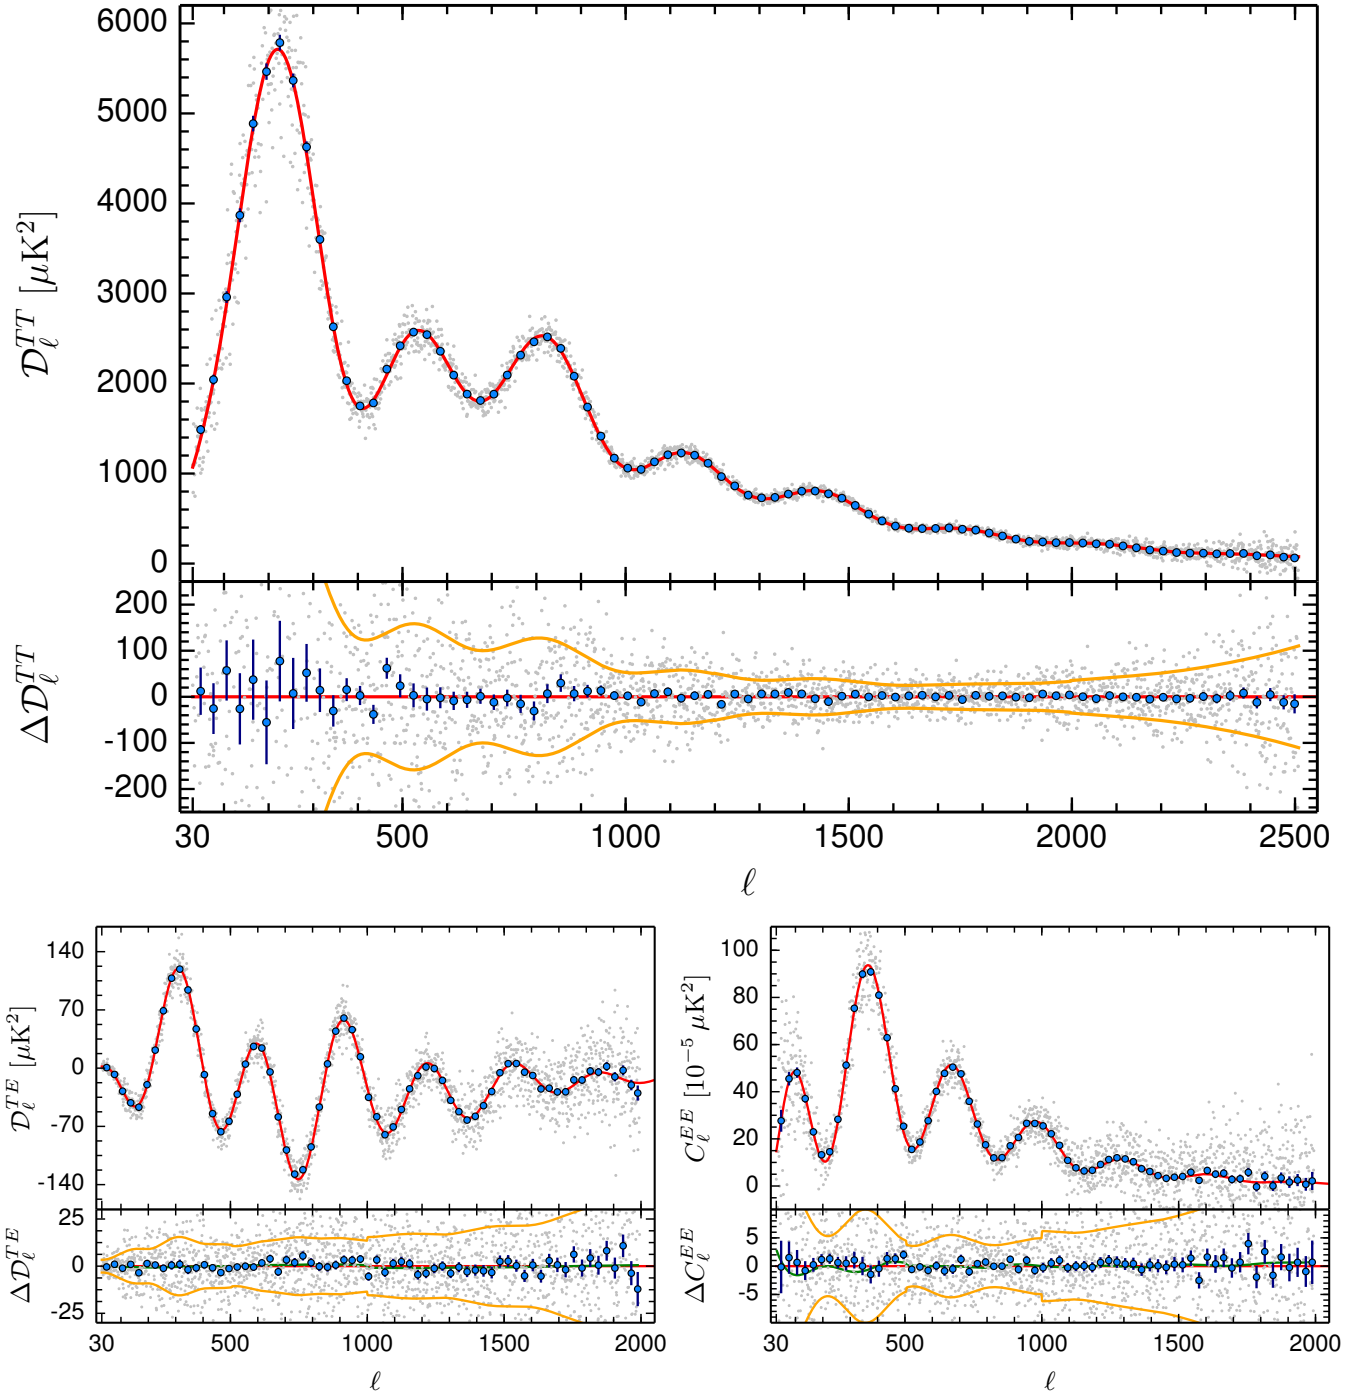
Fig. 28. Plik 2015 co-added *TT*, *TE*, and *EE* spectra. The blue points are for bins of $\Delta\ell = 30$, while the grey points are unbinned. The *lower panels* show the residuals with respect to the best fit PlikTT+tauprior ΛCDM model. The yellow lines show the 68% unbinned error bars. For *TE* and *EE*, we also show the best-fit beam-leakage correction (green line; see text and Fig. 23).

error-bars come from the variable ℓ ranges retained at different frequencies, which therefore vary the amount of data included discontinuously with ℓ. Figure 29 zooms in to five adjacent ℓ ranges on the co-added spectra to allow close inspection of the data distribution around the model.

More quantitatively, Table 16 shows the χ^2 values with respect to the ΛCDM best fit to the PlikTT+tauprior data combination for the unbinned CMB co-added power spectra (obtained as described in Appendix C.4). The *TT* spectrum has a reduced χ^2 of 1.03 for 2479 degrees of freedom, corresponding to a probability to exceed (PTE) of 17.2%; the base ΛCDM model

is therefore in agreement with the co-added data. The best-fit ΛCDM model in *TT* also provides an excellent description of the co-added polarized spectra, with a PTE of 12.8% in *TE* and 34.6% in *EE*. This already suggests that extensions with, e.g., isocurvature modes can be severely constrained.

Despite this overall agreement, we note that the PTEs are not uniformly good for all cross-frequency spectra (see in particular the 100×100 and 100×217 in *TE*). This shows that the baseline instrumental model needs to include further effects to describe all of the data in detail, even if the averages over frequencies appear less affected. The green line in Fig. 28 (mostly visible in
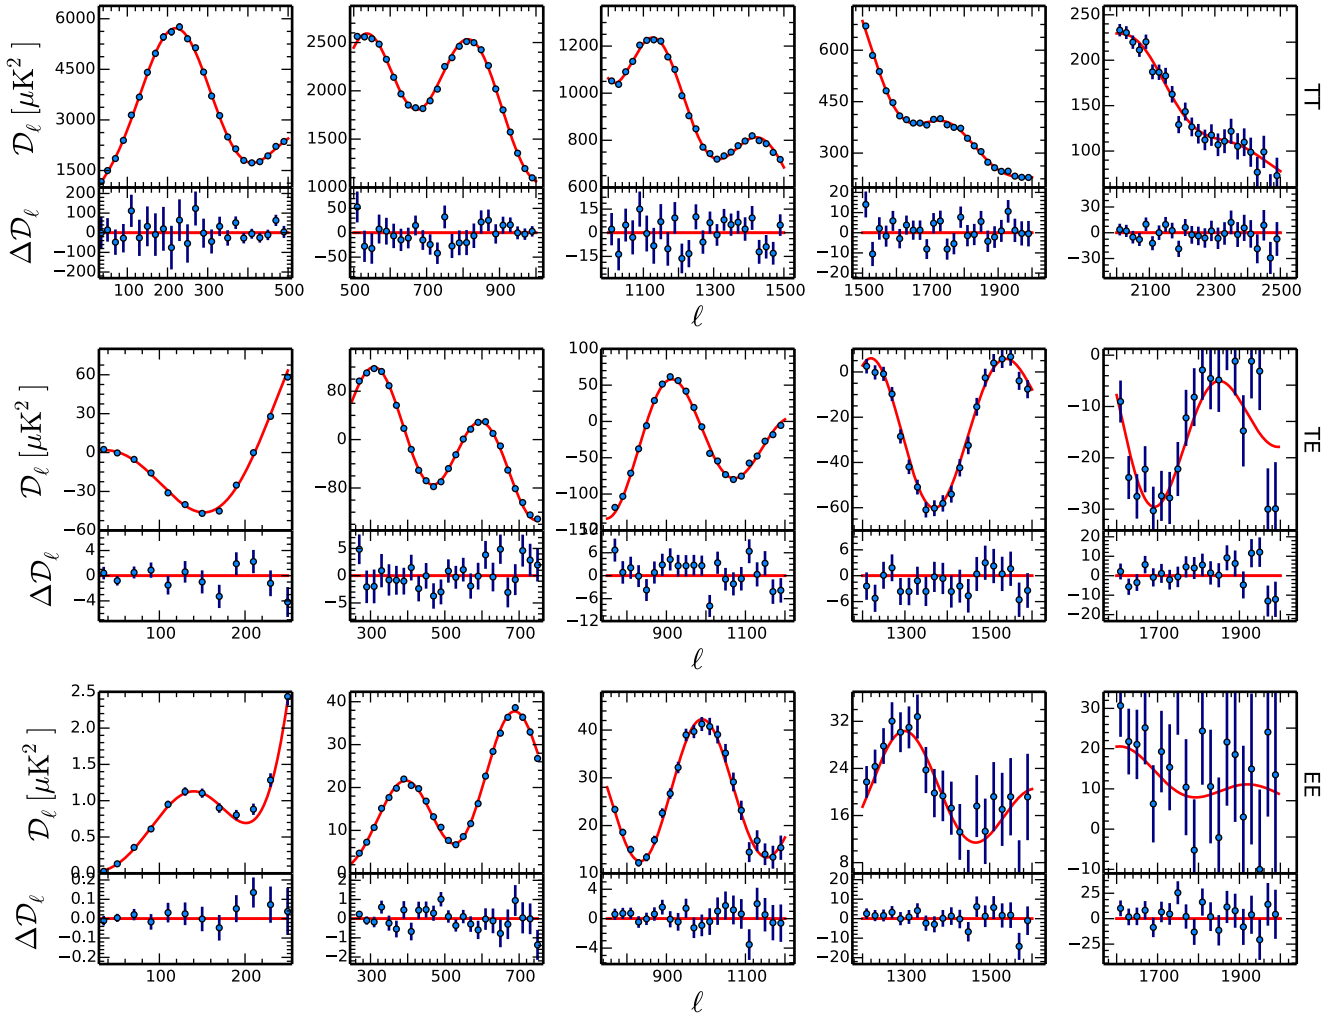
Fig. 29. Zoom in to various ℓ ranges of the HM co-added power spectra, together with the `PlikTT`+tauprior ΛCDM best-fit model (red line). We show the *TT* (*top*), *TE* (*centre*) and *EE* (*bottom*) power spectra. The *lower panels* in each plot show the residuals with respect to the best-fit model.

the ΔC_ℓ^{EE} plot) shows the best-fit leakage correction (shown on its own in Fig. 23), which is obtained when fixing the cosmology to the *TT*-based model. Let us recall, though, that this correction is for illustrative purposes only, and it is set to zero for all actual parameter searches. Indeed, we shall see that these leakage effects are not enough to bring all the data into full concordance with the model.

In more quantitative detail, Fig. 30 shows the binned ($\Delta\ell = 100$) residuals for the co-added CMB spectra in units of the standard deviation of each data point, (data − model)/error. For *TT*, we find the greatest deviations at $\ell \approx 434$ (-1.8σ), 464 (2.7σ), 1214 (-2.1σ), and 1450 ($-1.8,\sigma$). At $\ell = 1754$, where we previously reported a deficit due to the imperfect removal of the ^{4}He-JT cooler line (see Planck Collaboration XIII 2016, Sect. 3), there is now a less significant fluctuation, at the level of -1.4σ. The residuals in polarization show similar levels of discrepancy.

In order to assess whether these deviations are specific to one particular frequency channel or appear as a common signal in all the spectra, Fig. 31 shows foreground-cleaned *TT* power spectra differences across all frequencies, in units of standard deviations (details on how this is derived can be found in Appendix C.3.2). The agreement between *TT* spectra is clearly quite good. Figure 32 then shows the residuals per frequency for the *TT* power spectra with respect to the ΛCDM

`PlikTT`+tauprior best-fit model (see also the zoomed-in residual plots in Fig. C.5). The $\ell \approx 434$, 464, and 1214 deviations from the model appear to be common to all frequency channels, with differences between the frequencies smaller than 2σ. However, the deviation at $\ell \approx 1450$ is higher at 217×217 than in the other channels. In particular, the inter-frequency differences (Fig. 31) between the 217×217 power spectrum and the 100×100, 143×143, and 143×217 ones show deviations at $\ell \approx 1450$ at the roughly 1.7, 2.6, and 3.4σ levels, respectively.

This inter-frequency difference is due to a deficit in the residuals of the 217×217 channel of about -3.4σ in the bin centred at 1454 in Fig. 32. To better quantify this deviation, we also fit for a feature of the type $\cos^2((\pi/2)(\ell - \ell_p)/(\Delta\ell))$, with maximum amplitude centred at $\ell_p = 1460$, width $\Delta\ell = 25$ (we impose the feature to be zero at $|\ell - \ell_p| > \Delta\ell$) and with an independent amplitude in each frequency channel. At 217×217, we find an amplitude of $(-37.44 \pm 9.5)\mu$K^2, while in the other channels we find $(-15.0 \pm 7.8)\mu$K^2 at 143×143 and $(-19.7 \pm 7.9)\mu$K^2 at 143×217. This outlier seems to be at least in part due to chance correlation between the CMB and dust. Indeed, the amplitude of the feature in the different spectra is in rough agreement with the dust emission law. Moreover, the feature can also be found when varying the retained sky fraction in the galactic mask, again with an amplitude scaling compatible with a dust
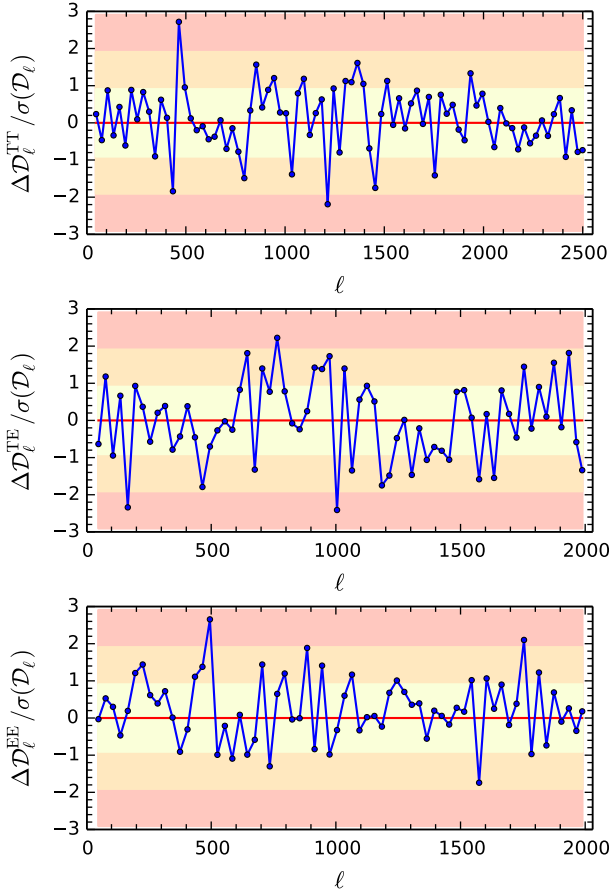
Fig. 30. Residuals of the co-added CMB *TT* power spectra, with respect to the `PlikTT`+tauprior best-fit model, in units of standard deviation. The three coloured bands (from the centre, yellow, orange, and red) represent the ± 1, ± 2, and $\pm 3\sigma$ regions.

Fig. 31. Inter-frequency foreground-cleaned *TT* power spectra differences, in μK^2. Each of the sub-panels shows the difference, after foreground subtraction, between pairs of frequency power spectra (the spectrum named on the vertical axis minus the one named on the horizontal axis), in units of standard deviation. The coloured bands identify deviations that are smaller than one (yellow), two (orange), or three (red) standard deviations. We show the differences for both the HM power spectra (blue points) and the DS power spectra (light blue points) after correlated noise correction. Figure 41 displays the same quantities for the *TE* and *EE* spectra.

origin. We discuss below the impact on cosmological parameters, see the case "CUT $\ell = 1404-1504$" in Fig. 35.

Finally we note that there is a deficit in the $\ell = 500-800$ region (in particular between $\ell = 700$ and 800) in the residuals of all the frequency spectra, roughly in correspondence with the position of the second and third peaks. Section 4.1 is dedicated to the study of these deviations and their impact on cosmological parameters. In spite of these marginally significant deviations from the model, the χ^2 values shown in Table 16 indicate that the ΛCDM model is an acceptable fit to each of the unbinned individual frequency power spectra, with PTEs always $\mathcal{P} \gtrsim 10\%$ in *TT*. We therefore proceed to examine the parameters of the best-fit model.

The cosmological parameters of interest are summarized in Table 17. Let us note that the cosmological parameters inferred here are obtained using the same codes, priors, and assumptions as in Planck Collaboration XIII (2016), except for the fact that we use the much faster `PICO` (Fendt & Wandelt 2007a) code instead of `CAMB` when estimating cosmological parameters[12] from *TT*, *TE* or *TT*, *TE*, *EE* using high-ℓ *Planck* data. Appendix C.5 establishes that the results obtained with the two codes only differ by small fractions of a standard deviation (less than 15% for most parameters, with a few more extreme deviations). However, we still use the `CAMB` code for results from *EE* alone, since in this case the parameter space explored is so wide that it includes
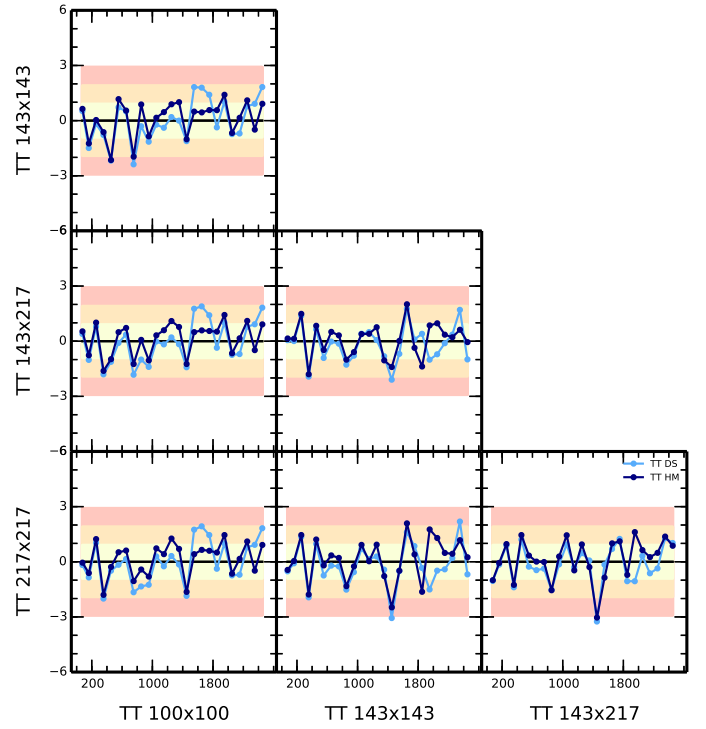
regions outside the `PICO` interpolation region (see Appendix C.5 for further details).

Figure 33 shows the posterior distributions of each pair of parameters of the base ΛCDM model from `PlikTT`+tauprior. The upper-right triangle compares the 1σ and 2σ contours for the full likelihood with those derived from only the $\ell < 1000$ or the $\ell \geq 1000$ data. Section 4.1.6 addresses the question of whether the results from these different cases are consistent with what can be expected statistically. The lower-left triangle further shows that the results are not driven by the data from a specific channel, i.e., dropping any of the 100, 143, or 217 GHz map data from the analysis does not lead to much change. The next section provides a quantitative analysis of this and other jack-knife tests.

We now turn to polarization results. Inter-frequency comparisons and residuals for *TE* and *EE* spectra are analysed in detail in Sect. 4.4. Suffice it to say here that the results are less satisfactory than in *TT*, both in the consistency between frequency spectra and in the detailed χ^2 results. This shows that the instrumental data model for polarization is less complete than for temperature, with residual effects at the μK^2 level. The model thus needs to be further developed to take full advantage of the HFI data in polarization, given the level of noise achieved. We thus consider the high-ℓ polarized likelihood as a "beta" version. Despite these limitations, we include it in the product delivery, to allow external reproduction of the results, even though the tests that we show indicate that it should not be used when searching for weak deviations (at the μK^2 level) from the baseline model.

---

[12] The definition of A_L differs in `PICO` and `CAMB`; see Appendix C.5.
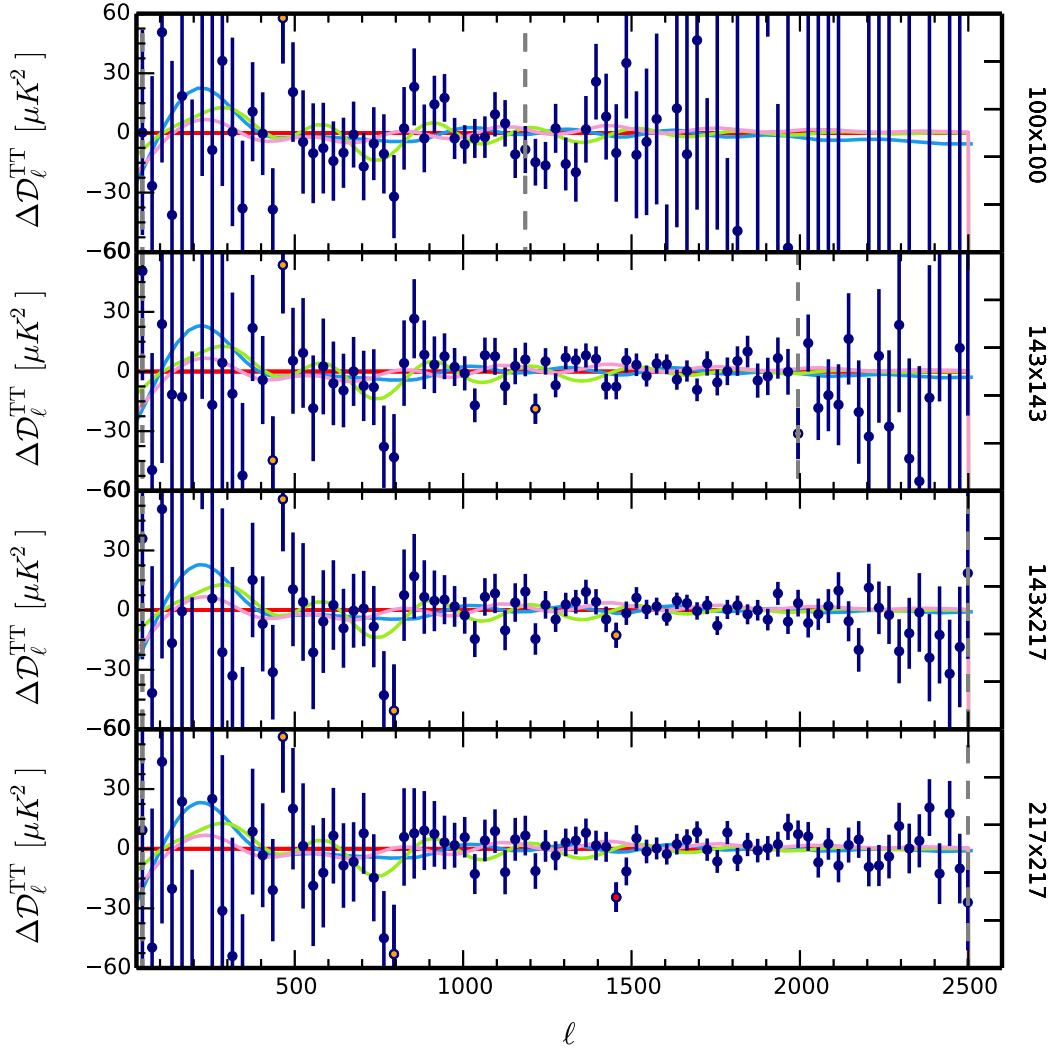
Fig. 32. Residuals in the half-mission *TT* power spectra after subtracting the `PlikTT+tauprior` ΛCDM best-fit model (blue points, except for those which differ by at least 2 or 3σ, which are coloured in orange or red, respectively). The light blue line shows the difference between the best-fit model obtained assuming a ΛCDM+A_L model and the ΛCDM best-fit baseline; the green line shows the difference of best-fit models using the ℓ_{max} = 999 likelihood (fixing the foregrounds to the baseline solution) minus the baseline best-fit (both in the ΛCDM framework); while the pink line is the same as the green one but for ℓ_{max} = 1404 instead of ℓ_{max} = 999; see text in Sect. 4.1. For the *TE* and *EE* spectra, see Fig. 40.

Nevertheless, we generally find agreement between the *TT*, *TE*, and *EE* spectra. Figure 34 shows the *TE*, and *EE* residual spectra conditioned on *TT*, which are close to zero. This is particularly the case for *TE* below ℓ = 1000, which gives some confidence in the polarization model. Most of the data points for *TE* and *EE* lie in the ±2σ range. As for all χ^2-based evaluations, the interpretation of this result depends crucially on the quality of the error estimates, i.e., on the quality of our noise model (see Sect. 3.4.4). We further note that the agreement is consistent with the finding that unmodelled instrumental effects in polarization are at the μK^2 level.

4. Assessment of the high-multipole likelihood

This section describes tests that we performed to assess the accuracy and robustness of the reference results of the high-ℓ likelihood that were presented above. First we establish the robustness of the *TT* results using `Plik` alone in Sect. 4.1 and with other likelihoods in Sect. 4.2. We verify in Sect. 4.3 that the amplitudes of the compact-source contributions derived at various frequencies are consistent with our current knowledge of source counts.

We then summarize in Sect. 4.4 the results of the detailed tests of the robustness of the polarization results, which are expanded upon in Appendix C.3.5. The paper Planck Collaboration XVI (2016) examines the dependence of the power spectrum on angular direction.

4.1. TT robustness tests

Figure 35 shows the marginal mean and the 68% CL error bars for cosmological parameters calculated assuming different data choices, likelihoods, parameter combinations, and data combinations. The 31 cases shown assume a base-ΛCDM framework, except when otherwise specified. The reference case uses the `PlikTT+tauprior` data combination. Figure 36 adds the specific results for the lensing parameter A_L (left) in a ΛCDM+A_L framework and for the effective number of relativistic species N_{eff} (right) in a ΛCDM+N_{eff} extended framework.

In both figures, the grey bands show the standard deviation of the parameter shifts relative to the baseline likelihood expected when using a sub-sample of the data (e.g., excising ℓ-ranges
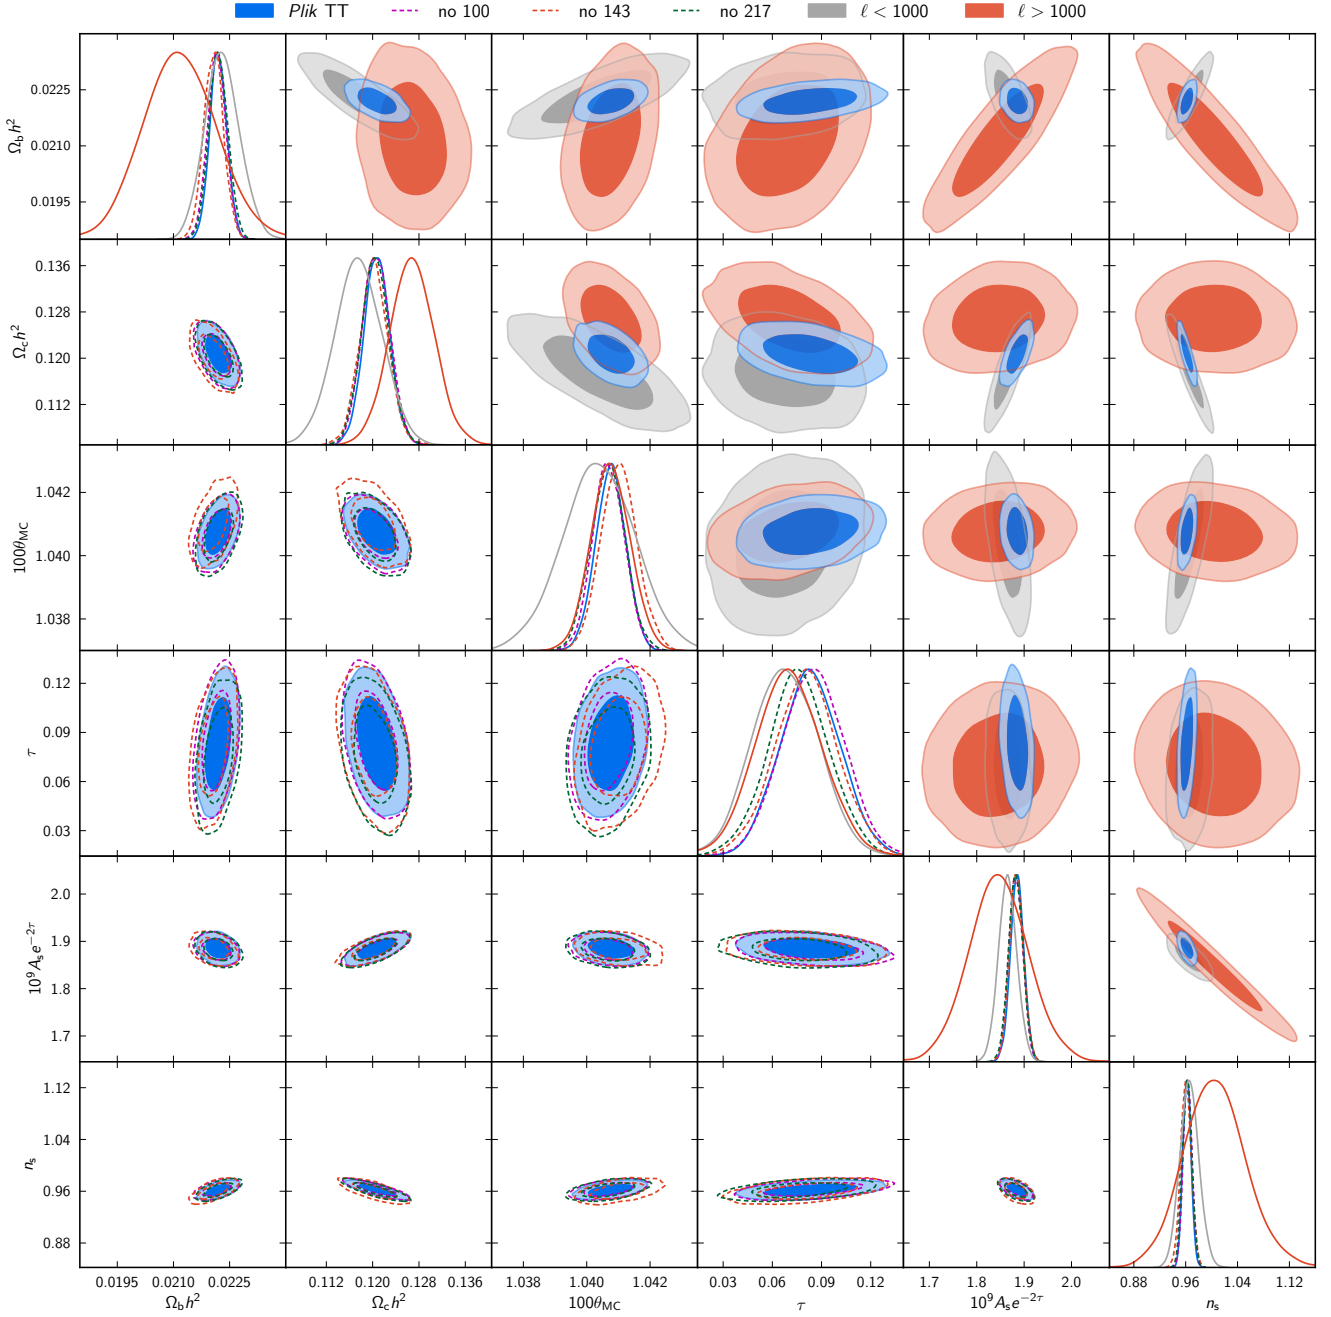
Fig. 33. ΛCDM parameters posterior distribution for `PlikTT+tauprior`. The lower left triangle of the matrix displays how the constraints are modified when the information from one of the frequency channels is dropped. The upper right triangle displays how the constraints are modified when the information from multipoles ℓ greater or less than 1000 is dropped. All the results shown in this figure were obtained using the `CAMB` code.

or frequencies). Because the data sets used to make inferences about a model are changed, one would naturally expect the inferences themselves to change, simply because of the effects of noise and cosmic variance. The inferences could also be influenced by inadequacies in the model, deficiencies in the likelihood estimate, and systematic effects in the data. Indeed, one may compare posterior distributions from different data subsets with each other and with those from the full data set, in order to assess the overall plausibility of the analysis.

To this end it is useful to have some idea about the typical variation in posteriors that one would expect to see even in the ideal case of an appropriate model being used to fit data sets with correct likelihoods and no systematic errors. It can be shown (Gratton & Challinor, in prep.) that if Y is a subset of a data set X, and $\boldsymbol{P}_X$ and $\boldsymbol{P}_Y$ are vectors of the maximum-likelihood parameter values for the two data sets, then the sampling distribution of the differences of the parameter values is given by

$$\overline{(\boldsymbol{P}_Y - \boldsymbol{P}_X)(\boldsymbol{P}_Y - \boldsymbol{P}_X)^{\mathsf{T}}} = \mathrm{cov}(\boldsymbol{P}_Y) - \mathrm{cov}(\boldsymbol{P}_X), \qquad (53)$$

i.e., the covariance of the differences is simply the difference of their covariances. Here the covariances are approximated by the
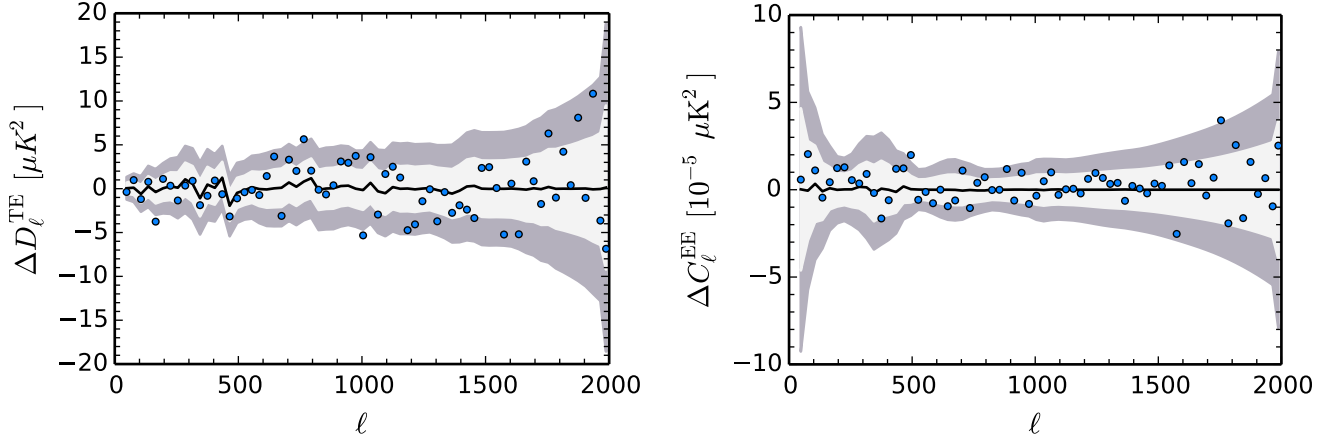
Fig. 34. *TE* (*left*) and *EE* (*right*) residuals conditioned on the *TT* spectrum (black line) with 1 and 2σ error bands. The blue points are the actual *TE* and *EE* residuals. We do not include any beam-leakage correction here.

Table 16. Goodness-of-fit tests for the `Plik` temperature and polarization spectra at high ℓ.

Frequency [GHz]	$f_{sky}[\%]^a$	Multipole range	χ^2	χ^2/N_ℓ	N_ℓ	$\Delta\chi^2\sqrt{2N_\ell}^b$	PTE [%]^c	χ_{norm}^d	PTE$_\chi[\%]^e$
TT									
100×100	66	30–1197	1234.91	1.06	1168	1.38	8.50	-0.30	76.44
143×143	57	30–1996	2034.59	1.03	1967	1.08	14.09	-0.39	69.91
143×217	49	30–2508	2567.11	1.04	2479	1.25	10.63	-1.07	28.25
217×217	47	30–2508	2549.40	1.03	2479	1.00	15.87	-0.17	86.72
Co-added		30–2508	2545.50	1.03	2479	0.94	17.22	-0.16	87.17
TE									
100×100	67	30–999	1089.75	1.12	970	2.72	0.43	3.70	0.02
100×143	50	30–999	1033.38	1.07	970	1.44	7.72	0.92	35.66
100×217	41	505–999	527.85	1.07	495	1.04	14.85	5.05	0.00
143×143	50	30–1996	2028.18	1.03	1967	0.98	16.45	-2.21	2.69
143×217	41	505–1996	1606.06	1.08	1492	2.09	2.02	-0.75	45.19
217×217	41	505–1996	1431.65	0.96	1492	-1.10	86.60	1.33	18.20
Co-added		30–1996	2038.54	1.04	1967	1.14	12.76	0.09	93.09
EE									
100×100	70	30–999	1027.14	1.06	970	1.30	9.89	1.13	25.88
100×143	52	30–999	1048.77	1.08	970	1.79	3.94	1.77	7.72
100×217	43	505–999	479.49	0.97	495	-0.49	68.33	-3.01	0.26
143×143	50	30–1996	2001.48	1.02	1967	0.55	28.87	3.74	0.02
143×217	43	505–1996	1430.95	0.96	1492	-1.12	86.89	-0.71	47.70
217×217	41	505–1996	1409.48	0.94	1492	-1.51	93.66	-1.39	16.45
Co-added		30–1996	1991.37	1.01	1967	0.39	34.55	1.88	6.00

Notes. $^{(a)}$ Effective fraction of the sky retained in the analysis. For the *TE* cross-spectra between two different frequencies, we show the smaller f_{sky} of the *TE* or *ET* combinations. $^{(b)}$ $\Delta\chi^2 = \chi^2 - N_\ell$ is the difference from the mean, assuming the best-fit *TT* base-ΛCDM model is correct, here expressed in units of the expected dispersion, $\sqrt{2N_\ell}$. $^{(c)}$ Probability to exceed the tabulated value of χ^2. $^{(d)}$ Weighted linear sum of deviations, scaled by the standard deviation, as defined in Eq. (60). $^{(e)}$ Probability to exceed the absolute value $|\chi_{norm}|$.

inverses of the appropriate Fisher information matrices evaluated for the true model. One might thus expect the scatter in the modes of the posteriors to follow similarly, and to be able, if the parameters are well-constrained by the data, to use covariances of the appropriate posteriors on the right-hand side.

4.1.1. Detset likelihood

We have verified (case "DS") that the results obtained using the half-mission cross-spectra likelihood are in agreement with those obtained using the detset (DS) cross-spectra likelihood. As

explained in Sect. 3.4.4, the main difficulty in using the DS likelihood is that the results might depend on the accuracy of the correlated noise correction. Reassuringly, we find that the results from the HM and DS likelihoods agree within 0.2σ. This is an important cross-check, since we expect the two likelihoods to be sensitive to different kinds of temporal systematics. Direct differences of half-mission versus detset-based *TT* cross-frequency spectra are compared in Fig. 31 (Fig. 41 shows similar plots for the *TE* and *EE* spectra.).

When using the detsets, we fit the calibration coefficients of the various detector sets with respect to a reference. The

Table 17. Cosmological parameters used in this analysis.

Parameter	Prior range	Baseline	Definition
$\omega_b \equiv \Omega_b h^2$	$[0.005, 0.1]$	. . .	Baryon density today
$\omega_c \equiv \Omega_c h^2$	$[0.001, 0.99]$	. . .	Cold dark matter density today
$\theta \equiv 100\theta_{MC}$	$[0.5, 10.0]$	. . .	$100 \times$ approximation to r_*/D_A (used in CosmoMC)
τ	$[0.01, 0.8]$	. . .	Thomson scattering optical depth due to reionization
τ	(0.07 ± 0.02)		
N_{eff}	$[0.05, 10.0]$	3.046	Effective number of neutrino-like relativistic degrees of freedom (see text)
Y_P	$[0.1, 0.5]$	BBN	Fraction of baryonic mass in helium
A_L	$[0.0, 10]$	1	Amplitude of the lensing power relative to the physical value
n_s	$[0.8, 1.2]$	. . .	Scalar spectrum power-law index ($k_0 = 0.05\ \mathrm{Mpc}^{-1}$)
$\ln(10^{10}A_s)$	$[2, 4.0]$	. . .	Log power of the primordial curvature perturbations ($k_0 = 0.05\ \mathrm{Mpc}^{-1}$)
Ω_Λ	. . .	. . .	Dark energy density divided by the critical density today
Age	. . .	. . .	Age of the Universe today (in Gyr)
Ω_m	. . .	. . .	Matter density (inc. massive neutrinos) today divided by the critical density
z_{re}	. . .	. . .	Redshift at which Universe is half reionized
H_0	$[20, 100]$	. . .	Current expansion rate in $\mathrm{km\,s}^{-1}\,\mathrm{Mpc}^{-1}$
$100\theta_D$	. . .	. . .	$100 \times$ angular extent of photon diffusion at last scattering
$100\theta_{eq}$	. . .	. . .	$100 \times$ angular size of the comoving horizon at matter-radiation equality

Notes. The columns indicate the cosmological parameter symbol, their uniform prior ranges in square brackets, or between parenthesis for a Gaussian prior, the baseline values if fixed for the standard ΛCDM model, and their definition. These parameters are the same as for the previous release. The top block lists the estimated parameters, while the lower block lists derived parameters.

resulting best-fit values are very close to one[13], with the greatest calibration refinement being less than 0.2%, in line with the accuracy expected from the description of the data processing in Planck Collaboration VIII (2016). This verifies that the maps produced by the HFI DPC and used for the half-mission-based likelihood come from the aggregation of well-calibrated and consistent data.

4.1.2. Impact of Galactic mask and dust modelling

We have tested the robustness of our results with respect to our model of the Galactic dust contribution in various ways.

Galactic masks. We have examined the impact of retaining a smaller fraction of the sky, less contaminated by Galactic emission. The baseline TT likelihood uses the G70, G60, and G50 masks (see Appendix A) at 100, 143, and 217 GHz, respectively. We have tested the effects of using G50, G41, and G41 (corresponding to $f_{sky}^{noap} = 0.60, 0.50,$ and 0.50 before apodization, case "M605050" in Fig. 35), and of the priors on the Galactic dust amplitudes relative to these masks described in Table 11. We find stable results as we vary these sky cuts, with the greatest shift in θ_{MC} of 0.5σ, compatible with the expected shift of 0.57σ calculated using Eq. (53). Going to higher sky fraction is more difficult. Indeed, the improvement in the parameter determination from increasing the sky fraction at 143 GHz and 217 GHz would be modest, as we would only gain information in the small-scale regime, which is not probed by 100 GHz. Increasing the sky fraction at 100 GHz is also more difficult because our estimates have shown that adding as little as 5% of the sky closer to the Galactic plane requires a change in the dust template and more than doubles the dust contamination at 100 GHz.

Amplitude priors. We have tested the impact of not using any prior (i.e., using arbitrarily wide, uniform priors) on the Galactic dust amplitudes (case "No gal. priors" in Fig. 35). Again, cosmological results are stable, with the greatest shifts in $\ln(10^{10}A_s)$ of 0.23σ and in n_s of 0.20σ. The values of the dust amplitude parameters, however, do change, and their best-fit values increase by about $15\,\mu\mathrm{K}^2$ for all pairs of frequencies, while at the same time the error bars of the dust amplitude parameters increase very significantly. All of the amplitude levels obtained from the 545 GHz cross-correlation are within 1σ of this result. The dust levels from this experiment are clearly unphysically high, requiring $22\,\mu\mathrm{K}^2$ ($\mathcal{D}_\ell, \ell = 200$) for the 100×100 pair. This level of dust contamination is clearly not allowed by the 545×100 cross-correlation, demonstrating that the prior deduced from it is informative. Nevertheless, the fact that cosmological parameters are barely modified in this test indicates that the values of the dust amplitudes are only weakly correlated with those of the cosmological parameters, consistent with the results of Figs. 44 and 45 below, which show the parameter correlations quantitatively.

Galactic dust template slope. We have allowed for a variation of the Galactic dust index n, defined in Eq. (24), from its default value $n = -2.63$, imposing a Gaussian prior of -2.63 ± 0.05 ("GALINDEX" case in Fig. 35). We find no shift in cosmological parameters (smaller than $\sim 0.1\sigma$) and recover a value for the index of $n = -2.572 \pm 0.038$, consistent with our default choice.

Impact of $\ell \lesssim 500$ at 217 GHz. We have analysed the impact of excising the first 500 multipoles ("LMIN=505 at 217 GHz" in Fig. 35) in the 143×217 and 217×217 spectra, where the Galactic dust contamination is the strongest. We find very good stability in the cosmological parameters, with the greatest change being a 0.16σ increase in n_s. This is compatible with the expectations estimated from Eq. (53) of 0.14σ. The inclusion of the first 500 multipoles at 217 GHz in the baseline Plik likelihood is one of the sources of the roughly 0.45σ difference

---

[13] The fitted values are 1.0000, 0.9999, 1.0000, 1.0000, 0.9987, 0.9986, 0.9992, 0.9989, 0.9989, 0.9981, 0.9989, 1.0000, and 0.9999 for detsets 100-ds1, 100-ds2, 143-ds1, 143-ds2, 143-5, 143-6, 143-7, 217-1, 217-2, 217-3, 217-4, 217-ds1, and 217-ds2, respectively.
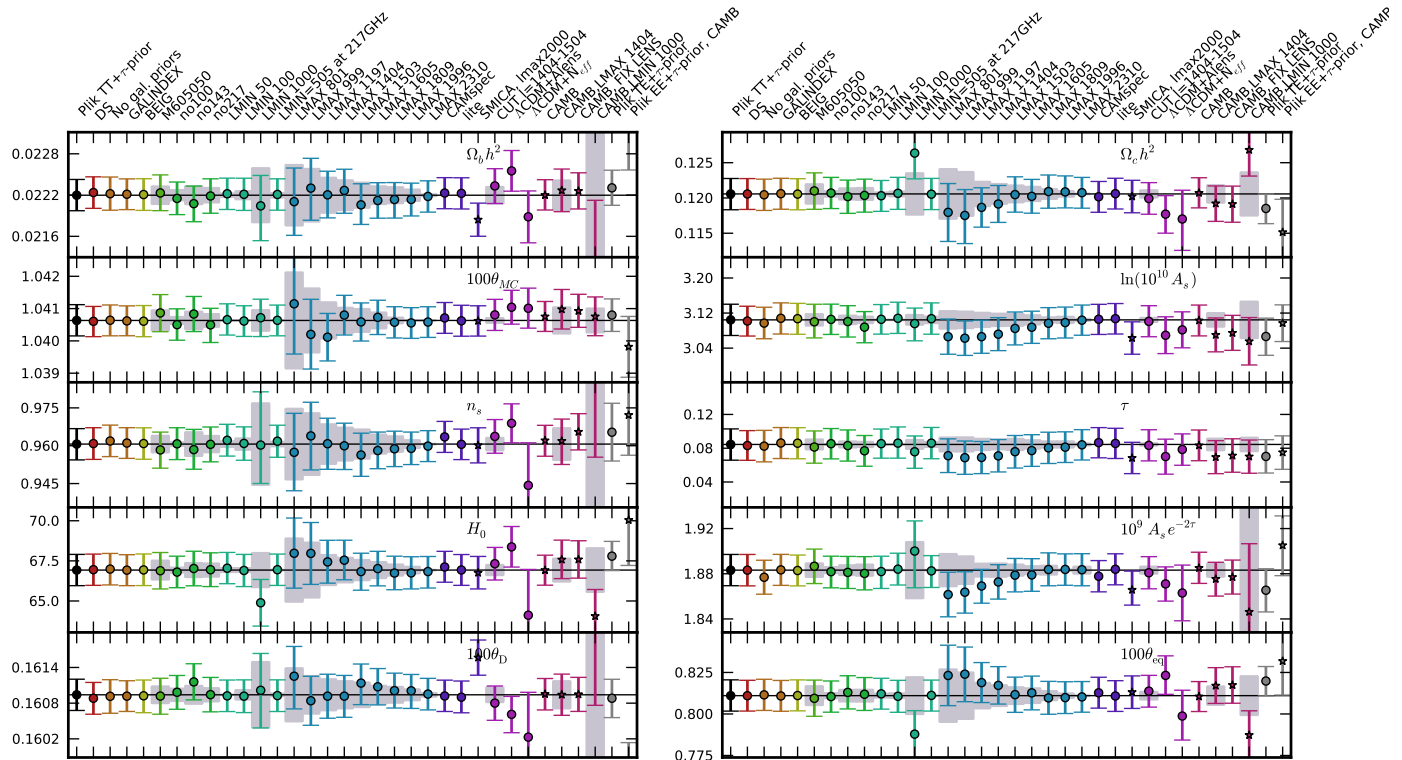
Fig. 35. Marginal mean and 68% CL error bars on cosmological parameters estimated with different data choices for the `Plik` likelihood, in comparison with results from alternate approaches or model. We assume a ΛCDM model and use variations of the `PlikTT` likelihood in most of the cases, in combination with a prior $\tau = 0.07 \pm 0.02$ (using neither low-ℓ temperature nor polarization data). The "`PlikTT`+tauprior" case (black dot and thin horizontal black line) indicates the baseline (HM, $\ell_{min} = 30$, $\ell_{max} = 2508$), while the other cases are described in Sect. 4.1 (and 4.2, 5.6, E.4). The grey bands show the standard deviation of the expected parameter shift, for those cases where the data used is a sub-sample of the baseline likelihood (see Eq. (53)). All the results were run with `PICO` except for few ones that were run with `CAMB`, as indicated in the labels.

in n_s observed when using the `CamSpec` code, since the latter excises that range of multipoles; for further discussion see Planck Collaboration XIII (2016, Table 1 and Sect. 3.1), as well as Sect. 4.2.

4.1.3. Impact of beam uncertainties

The case labelled "BEIG" in Fig. 35 corresponds to the exploration of beam eigenvalues with priors 10 times higher than indicated by the analysis of our MC simulation of beam uncertainties (which indicated by dotted lines in Fig. 21). This demonstrates that these beam uncertainties are so small in this data release that they do not contribute to the parameter posterior widths. They are therefore not enabled by default.

4.1.4. Inter-frequency consistency and redundancy

We have tested the effect of estimating parameters while excluding one frequency channel at a time. In Figs. 33 and 35, the "no100" case shows the effect of excluding the 100×100 frequency spectrum, the "no143" of excluding the 143×143 and 143×217 spectra, and the "no217" of excluding the 143×217 and 217×217 spectra.

We obtain the greatest deviations in the "no217" case for $\ln(10^{10}A_s)$ and τ, which shift to lower values by 0.53σ and 0.47σ, about twice the expected shift calculated using Eq. (53), 0.25σ and 0.23σ respectively (in units of standard deviations of the "no217" case). The value of $\Omega_c h^2$ decreases by only -0.1σ. Figure 37 further shows the 217×217 spectrum conditioned

on the 100×100 and 143×143 ones. This conditional deviates significantly in two places, at $\ell = 200$ and $\ell = 1450$. The $\ell = 1450$ case was already discussed in Sect. 3.8 and is further analysed in Sect. 4.1.6. Around $\ell = 200$, we see some excess scatter (both positive and negative) in the data around a jump between two consecutive bins of the conditional. This corresponds to the two bins around the first peak (one right before and the other almost at the location of the first peak), as can be seen in Fig. 28. All of the frequencies exhibit a similar behaviour (see Fig. 32); however, it is most pronounced in the 217 GHz case. This multipole region is also near the location of the bump in the effective dust model. The magnitude of this excess power in the model is not big enough or sharp enough to explain this excess scatter (see Fig. 17). Finally, note that the best-fit CMB solution at large scales is dominated by the 100×100 data, which are measured on a greater sky fraction (see Fig. 14).

This test shows that the parameters of the ΛCDM model do not rely on any specific frequency map, except for a weak pull of the higher resolution 217 GHz data towards higher values of both A_s and τ (but keeping $A_s \exp(-2\tau)$ almost constant).

4.1.5. Changes of parameters with ℓ_{min}

We have checked the stability of the results when changing ℓ_{min} from the baseline value of $\ell_{min} = 30$ to $\ell_{min} = 50$ and 100 (and $\ell_{min} = 1000$, which is discussed in Sect. 4.1.6). These correspond to the cases labelled "LMIN 50" and "LMIN 100" in Fig. 35 (to be compared to the reference case "`PlikTT`+tauprior"). This check is important, since the
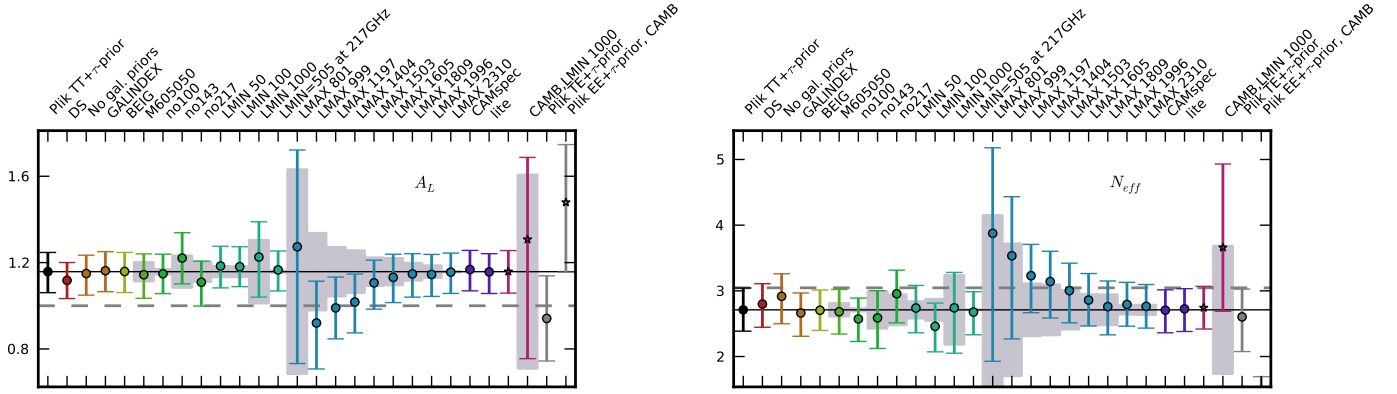
Fig. 36. Marginal mean and 68% CL error bars on the parameters A_L (*left*) and N_{eff} (*right*) in ΛCDM extensions, estimated with different data choices for the `PlikTT` likelihood in comparison with results from alternate approaches or model, combined with a Gaussian prior on $\tau = 0.07 \pm 0.02$ (i.e., neither low-ℓ temperature nor polarization data). The "PlikTT+tauprior" case indicates the baseline (HM, $\ell_{min} = 30$, $\ell_{max} = 2508$), while the other cases are described in subsections of Sect. 4.1. The thin horizontal black line shows the baseline result and the thick dashed grey line displays the ΛCDM value ($A_L = 1$ and $N_{eff} = 3.04$). The grey bands show the standard deviation of the expected parameter shift, for those cases where the data used is a sub-sample of the baseline likelihood (see Eq. (53)).
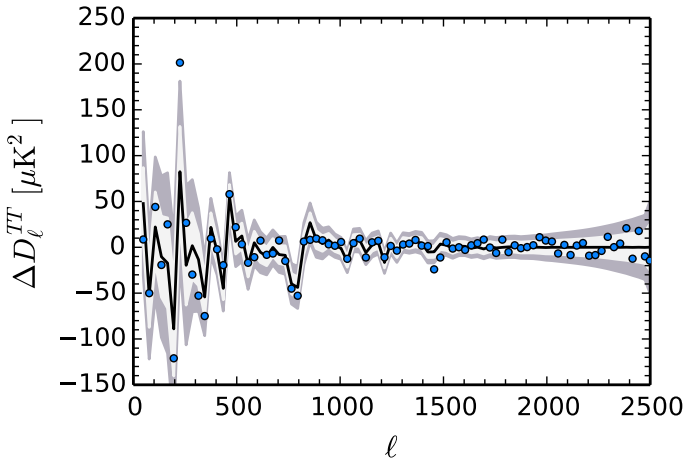

Fig. 37. 217×217 spectrum conditioned on the joint result from the 100×100 and 143×143 spectra. The most extreme outliers are at $\ell = 200$ and $\ell = 1450$.

Gaussian approximation assumed in the likelihood is bound to fail at very low ℓ (for further discussion, see Sect. 3.6).

The results are in good agreement, with shifts in parameters smaller than 0.2σ, well within expectations calculated from Eq. (53). This is also confirmed in Fig. 42, where the TT hybridization scale of the full likelihood is varied (i.e., the multipole where the low-ℓ and high-ℓ likelihoods are joined).

4.1.6. Changes of parameters with ℓ_{max}

We have tested the stability of our results against changes in the maximum multipole ℓ_{max} considered in the analysis. We test the restriction to ℓ_{max} in the range $\ell_{max} = 999-2310$, with the baseline likelihood having $\ell_{max} = 2508$. For each frequency power spectrum we choose $\ell_{max}^{freq} = \min(\ell_{max}, \ell_{max}^{freq, base})$, where $\ell_{max}^{freq, base}$ is the baseline ℓ_{max} at each frequency as reported in Table 16. The results shown in Fig. 35 use the same settings as the baseline likelihood (in particular, we leave the same nuisance parameters free to vary) and always use a prior on τ.

The results in Fig. 35 suggest there is a shift in the mean values of the parameters when using low ℓ_{max}; e.g., for $\ell_{max} = 999$,

$\ln(10^{10}A_s)$, τ, and $\Omega_c h^2$ are lower by 1.0, 0.8, and 0.8σ with respect to the baseline parameters. These parameters then converge to the baseline values for $\ell_{max} \gtrsim 1500$. Following the arguments given earlier (Eq. (53)), when using these nested sub-samples of the baseline data we expect shifts of the order of 0.5, 0.4, and 0.8σ respectively, in units of the standard deviation of the $\ell_{max} = 999$ results. We further note that the value of θ for $\ell_{max} \lesssim 1197$ is lower compared to the baseline value. In particular, at $\ell_{max} = 1197$, its value is 0.8σ low, while the expected shift is of the order of 0.7σ, in units of the standard deviation of the $\ell_{max} = 1197$ results. The value of θ then rapidly converges to the baseline for $\ell_{max} \gtrsim 1300$. Figure C.8 in Appendix C.3.3 also shows that these shifts are related to a change in the amplitude of the foreground parameters. In particular, the overall level of foregrounds at each frequency decreases with increasing ℓ_{max}, partially compensating for the increase in $\ln(10^{10}A_s)$ and $\Omega_c h^2$. Although all these shifts are compatible with expectations within a factor of 2, we performed some further investigations in order to understand the origin of these changes. In the following, we provide a tentative explanation.

Table 18 shows the difference in χ^2 between the best-fit model obtained using $\ell_{max} = 999$ (or $\ell_{max} = 1404$) and the baseline `PlikTT+tauprior` best-fit solution in different multipole intervals. For this test, we ran the ℓ_{max} cases fixing the nuisance parameters to the baseline best-fit solution. This is required in order to be able to "predict" the power spectra at multipoles higher than ℓ_{max}, since otherwise the foreground parameters, which are only weakly constrained by the low-ℓ likelihood, can converge to unreasonable values. We note that fixing the foregrounds has an impact on cosmological parameters, which can differ from the ones shown in Fig. 35 (see Appendix C.3.4 for a direct comparison). Nevertheless, since the overall behaviour with ℓ_{max} is similar, we use this simplified scenario to study the origin of the shifts.

The χ^2 differences in Table 18 indicate that the cosmology obtained using $\ell_{max} = 999$ is a better fit in the region between $\ell = 630$ and 829. In particular, the low value of θ preferred by the $\ell_{max} = 999$ data set shifts the position of the third peak to smaller scales. This enables a better fit to the low points at $\ell \approx 700-850$ (before the third peak), followed by the high points at $\ell \approx 850-950$ (after the third peak). This is also clear from the residuals and the green solid line in Fig. 32, which shows

Table 18. Difference of χ^2 values between pairs of best-fit models in different ℓ−ranges for the co-added TT power spectrum.

Multipole range	$\Delta_{\ell_{max}=999}$	$\Delta_{\ell_{max}=1404}$	Δ_{A_L}
30–129	0.1	0.31	0.4
130–229	0.07	0.05	0.3
230–329	−0.4	−0.22	−0.45
330–429	0.34	−0.09	0.22
430–529	−0.01	0.17	0.26
530–629	0.61	−0.26	−0.2
630–729	−1.66	−0.8	−0.8
730–829	−1.15	−0.13	−0.79
830–929	−0.45	0.01	0.91
930–1029	−0.87	0.41	0.58
1030–1129	2.17	−0.94	−0.24
1130–1229	1.65	1.47	−0.17
1230–1329	0.87	0.17	−0.08
1330–1429	6.21	−1.46	−0.64
1430–1529	−0.2	3.35	−0.62
1530–1629	0.78	0.27	−0.44
1630–1729	0.73	0.9	0.06
1730–1829	0	1.18	−0.01
1830–1929	0.59	−0.08	−0.31
1930–2029	0.21	0.04	−0.04
2030–2129	0	0.57	−0.12
2130–2229	0.11	0.19	−0.18
2230–2329	−0.17	0.25	−0.2
2330–2429	0.06	−0.16	0.09
2430–2508	2.63	2.66	−0.19

Notes. The first column shows the ℓ-range, the second shows the difference $\Delta_{\ell_{max}=999}$ between the χ^2 values for a ΛCDM best-fit model obtained using either a likelihood with $\ell_{max} = 999$ or the baseline, i.e., $\Delta_{999} \equiv \left(\chi^2_{\ell_{max}=999} - \chi^2_{BASE}\right)_{\Lambda CDM}$. The $\ell_{max} = 999$ case was run fixing the foreground parameters to the best fit of the baseline case. The third column is the same as the second, but for $\ell_{max} = 1404$. The fourth column shows the difference Δ_{A_L} between the χ^2 values obtained in the ΛCDM+A_L and the ΛCDM frameworks. In this case, all the foreground and nuisance parameters were free to vary in the same way as in the baseline case.

the difference in best-fit models between the $\ell_{max} = 999$ case and the reference case. However, the values in Table 18 also show that the $\ell_{max} = 999$ cosmology is disfavoured by the multipole region between $\ell \approx 1330−1430$, before the fifth peak. The $\ell_{max} = 999$ model predicts too little power in this multipole range, which can be better fit if the position of the fifth peak moves to lower multipoles. As a consequence, θ shifts to higher values when including $\ell_{max} \gtrsim 1400$.

Concerning the shifts in $\Omega_c h^2$, A_s and τ, Fig. 35 shows that these parameters converge to the full baseline solution between $\ell_{max} = 1404$ and $\ell_{max} = 1505$. The $\Delta\chi^2$ values in Table 18 between the best-fit $\ell_{max} = 1404$ case and the baseline suggest that the $\ell_{max} = 1404$ cosmology is disfavoured by the multipole region $\ell = 1430−1530$ (fifth peak), and – at somewhat lower significance – by the regions close to the fourth peak ($\ell \approx 1130−1230$) and the sixth peak ($\ell \approx 1730−1829$). The pink line in Fig. 32 shows the differences between the $\ell_{max} = 1404$ best-fit model and the baseline, and it suggests that the $\ell_{max} = 1404$ cosmology predicts an amplitude of the high-ℓ peaks that is too large.

This effect can be compensated by more lensing, which can be obtained with greater values of $\Omega_c h^2$ and $\ln(10^{10}A_s)$, as well as a greater value of τ to compensate for the increase in A_s in the normalization of the spectra, as observed when considering

$\ell_{max} \gtrsim 1500$. This also explains why the baseline ($\ell_{max} = 2508$) best-fit solution prefers a value of the optical depth which is 0.8σ higher than the mean value of the Gaussian prior ($\tau = 0.07 \pm 0.02$), $\tau = 0.085 \pm 0.018$. In order to verify this interpretation, we performed the following test (using the CAMB code instead of PICO). We fixed the theoretical lensing power spectrum to the best-fit parameters preferred by the $\ell_{max} = 1404$ cosmology, and estimated cosmological parameters using the baseline likelihood. This is the "CAMB, FIX LENS" case in Fig. 35, which shows that cosmological parameters shift back to the values preferred at $\ell_{max} = 1404$ ("CAMB, $\ell_{max} = 1404$") if they cannot alter the amount of lensing in the model.

Since the $\ell \approx 1400−1500$ region is also affected by the deficit at $\ell = 1450$ (described in Sect. 3.8), we tested whether excising this multipole region from the baseline likelihood (with $\ell_{max} = 2508$) has an impact on the determination of cosmological parameters. The results in Fig. 35 (case "CUT $\ell = 1404−1504$") show that the parameter shifts are at the level of 0.47, −0.29, 0.38, and 0.45σ on $\Omega_b h^2$, $\Omega_c h^2$, θ, and n_s, respectively (0.39, 0.09, 0.24, and 0.29σ expected from Eq. (53)), confirming that this multipole region has some impact on the parameters, although it cannot completely account for the shift between the $\ell_{max} \approx 1400$ case and the baseline.

We also estimated cosmological parameters including only multipoles $\ell > 1000$ ("LMIN 1000" case), and compared them to the "LMAX 999" case[14] (see also Appendix C.5). The two-dimensional posterior distributions in Fig. 33 show the complementarity of the information from $\ell \le 999$ and $\ell \ge 1000$, with degeneracy directions between pairs of parameters changing in these two multipole regimes. The $\ell_{min} = 1000$ likelihood sets constraints on the amplitude of the spectra $A_s e^{-2\tau}$ and on n_s that are almost a factor of 2 weaker than the ones obtained with the baseline likelihood, and somewhat higher than the ones obtained with $\ell_{max} = 999$. The value of τ is thus more effectively determined by its prior and shifts downward by 0.59σ with respect to the baseline. The value of $\Omega_c h^2$ shifts upward by 1.7σ (cf. 0.8σ expected from Eq. (53)). Whether this change is just due to a statistical fluctuation is still a matter of investigation.

However, since parameter shifts are correlated, we evaluated whether the ensemble of the shifts in all cosmological parameters between the $\ell_{max} = 999$ and $\ell_{min} = 1000$ cases are compatible with statistical expectations. In order to do so, we computed the χ^2_Δ statistic of the shift as

$$\chi^2_\Delta = \sum_{ij} \Delta_i \Sigma^{-1}_{ij} \Delta_j, \tag{54}$$

where Δ_i is the difference in best-fit value of the ith parameter between the $\ell_{max} = 999$ and $\ell_{min} = 1000$ cases and Σ is the covariance matrix of the expected shifts, calculated as the sum

---

[14] During the revision of this paper, we noticed that the $\ell > 1000$ case explores regions of parameter space that are outside the optimal PICO interpolation region, as also remarked by Addison et al. (2016). This inaccuracy mainly affected this particular test for constraints on n_s and $\Omega_b h^2$: the error bars for these parameters were underestimated by a factor of about 2 while the mean values were misestimated by about 0.8σ with respect to runs performed with CAMB. Nevertheless, we found that for all other parameters, and in all other likelihood tests presented in this section, this problem did not arise, since the explored parameter space was entirely contained in the PICO interpolation region so as to guarantee accurate results, as also detailed in Sect. C.5. Furthermore, this inaccuracy does not change any of the conclusions of this paper. We therefore decided to keep in Fig. 35 the results obtained with PICO but we have added results for the $\ell > 1000$ case obtained with CAMB (case "CAMB, $l_{min} = 1000$").

of the parameter covariance matrices obtained in each of the two cases, ignoring correlations between the two datasets. We include in this calculation the ΛCDM parameters ($\Omega_b h^2$, $\Omega_c h^2$, θ, n_s, $A_s \exp(-2\tau)$), excluding τ, since the constraints on this parameter are dominated by the same prior in both cases, and using $A_s \exp(-2\tau)$ instead of $\ln(10^{10} A_s)$, since the latter is very correlated with τ and the TT power spectrum is mostly sensitive to the combination $A_s \exp(-2\tau)$. Finally, we estimate the χ_Δ^2 both in the case where we leave the foregrounds free to vary or in the case where we fix them to the best fit of the baseline PlikTT + tauprior solution. Assuming that χ_Δ^2 has a χ^2 distribution for 5 degrees of freedom, we find that the shifts observed in the data are consistent with simulations at the 1.2σ (1.1σ with fixed foregrounds) level for the case where we do not include the low-ℓ TT likelihood at $\ell < 30$ to the $\ell_{max} = 999$ case, and at the 1.5σ (1.4σ with fixed foregrounds) level for the case where we include the low-ℓ TT likelihood. We also find that the use of $A_s \exp(-2\tau)$ instead of $\ln(10^{10} A_s)$ changes these significances only in the case where we include the low-ℓ TT likelihood to the $\ell_{max} = 999$ case and leave the foregrounds free to vary, in which case we find consistency at the level of 1.8σ, in agreement with the findings of Addison et al. (2016; although in this case the use of $\ln(10^{10} A_s)$ and the exclusion of τ makes this test less indicative of the true significance of the shifts). In all cases, we do not find evidence for a discrepancy between the two datasets. A more precise and extended evaluation and discussions of these shifts, based on numerical simulations, will be presented in a future publication.

4.1.7. Impact of varying A_L

Figure 36 (left) displays the impact of various choices on the value of the lensing parameter A_L in the ΛCDM+A_L framework. The baseline likelihood prefers a value of A_L that is about 2σ greater than the physical value, $A_L = 1$. It is clear that this preference only arises when data with $\ell_{max} \gtrsim 1400$ are included, and it is caused by the same effects as we proposed in Sect. 4.1.6 to explain the shifts in parameters at $\ell_{max} \gtrsim 1400$ in the ΛCDM case. More lensing helps to fit the data in the $\ell \approx 1300-1500$ region, as indicated by the χ^2 differences between the ΛCDM+A_L best-fit and the ΛCDM one in Table 18. This drives the value of A_L to 1.159 ± 0.090 with PlikTT+tauprior, 1.8σ higher than expected. The case "ΛCDM+A_L" of Fig. 35 also shows that opening up this unphysical degree of freedom shifts the other cosmological parameters at the 1σ level; e.g., $\Omega_c h^2$ and A_s shift closer to the values preferred in the ΛCDM case when using $\ell_{max} \lesssim 1400$. While in the ΛCDM case high values of these parameters allow increasing lensing, in the ΛCDM+A_L case this is already ensured by a high value of A_L, so $\Omega_c h^2$ and A_s can adopt values that better fit the $\ell \lesssim 1400$ range. When using PlikTT in combination with the lowTEB likelihood, the deviation increases to 2.4σ, $A_L = 1.204 \pm 0.086$,[15] due to the fact that more lensing allows smaller values of $\Omega_c h^2$ and A_s and a greater value of n_s, better fitting the deficit at $\ell \approx 20$ in the temperature power spectrum (see Planck Collaboration XIII 2016, Sect. 5.1.2 and Fig. 13).

4.1.8. Impact of varying N_{eff}

We have investigated the effect of opening up the N_{eff} degree of freedom in order to assess the robustness of the constraints

on the ΛCDM extensions, which rely heavily on the high-ℓ tail of the data. Figure 36 (right) shows that N_{eff} departs from the standard 3.04 value by about 1σ when using PlikTT+tauprior, $N_{eff} = 2.7 \pm 0.33$. The χ^2 improvement for this model over ΛCDM is only $\Delta\chi^2 = 1.5$. We note that when the lowTEB likelihood (or alternatively, the low-ℓ TT likelihood plus the prior on τ) is used in combination with PlikTT, the value of N_{eff} shifts higher by about 1σ, $N_{eff} = 3.09 \pm 0.29$. This shift is about a factor 2 more than the one expected from Eq. (53), 0.5σ, between the PlikTT+tauprior and PlikTT+tauprior+low-ℓ TT cases. This shift is due to the fact that the deficit at $\ell \approx 20$ is better fit by higher n_s and, as a consequence, an increase in N_{eff} helps decreasing the enhanced power at high ℓ.

Figure 36 also shows that, not surprisingly, the most extreme variations as compared to the reference case (less than 1σ) arise when the high-resolution data are dropped (by reducing ℓ_{max} or by removing the 217 GHz channel), owing to the strong dependence of the N_{eff} constraints on the damping tail.

Having opened up this degree of freedom, the standard parameters are now about 1σ away (see case "ΛCDM+N_{eff}" of Fig. 35), and such a model would prefer quite a low value of H_0, which would then be at odds with priors derived from direct measurements (see Planck Collaboration XIII 2016, for an in-depth analysis).

4.2. Intercomparison of likelihoods

In addition to the baseline high-ℓ Plik likelihood, we have developed four other high-ℓ codes, CamSpec, Hillipop, Mspec, and Xfaster. CamSpec and Xfaster have been described in separate papers (Planck Collaboration XV 2014; Rocha et al. 2011), and brief descriptions of Mspec and Hillipop are given in Appendix D. These codes have been used to perform data consistency tests, to examine various analysis choices, and to cross-check each other by comparing their results and ensuring that they are the same. In general, we find good agreement between the codes, with only minor differences in cosmological parameters.

The CamSpec, Hillipop, and Mspec codes are, like Plik, based on pseudo-C_ℓ estimators and an analytic calculation of the covariance (Efstathiou 2004, 2006), with some differences in the approximations used to calculate this covariance. The Xfaster code (Rocha et al. 2011) is an an approximation to the iterative, maximum likelihood, quadratic band-power estimator based on a diagonal approximation to the quadratic Fisher matrix estimator (Rocha et al. 2011, 2010), with noise bias estimated using difference maps, as described in Planck Collaboration IX (2016). For temperature, all of the codes use the same Galactic masks, but they differ in point-source masking: Hillipop uses a mask based on a combination of $S/N > 7$ and cuts based on flux, while the others use the baseline $S/N > 5$ mask described in Appendix A. The codes also differ in foreground modelling, in the choice of data combinations, and in the ℓ-range. For the comparison presented here, all make use of half-mission maps.

Figure 38 shows a comparison of the power spectra residuals and error bars from each code, while Fig. E.5 in Appendix E.4 compares the combined spectra with the best-fit model. In temperature, the main feature visible in these plots is an overall nearly constant shift, up to $10 \mu K^2$ in some cases. This represents a real difference in the best-fit power each code attributes to foregrounds. For context, it is useful to note the statistical uncertainty on the foregrounds; for example, the 1σ error on the total foreground power at 217 GHz at $\ell = 1500$ is $2.5 \mu K^2$ (calculated here with Mspec, but similar for the other codes). Shifts

---

[15] These results were obtained with the PICO code, and are thus close to but not identical to those obtained with CAMB and reported in Planck Collaboration XIII (2016).
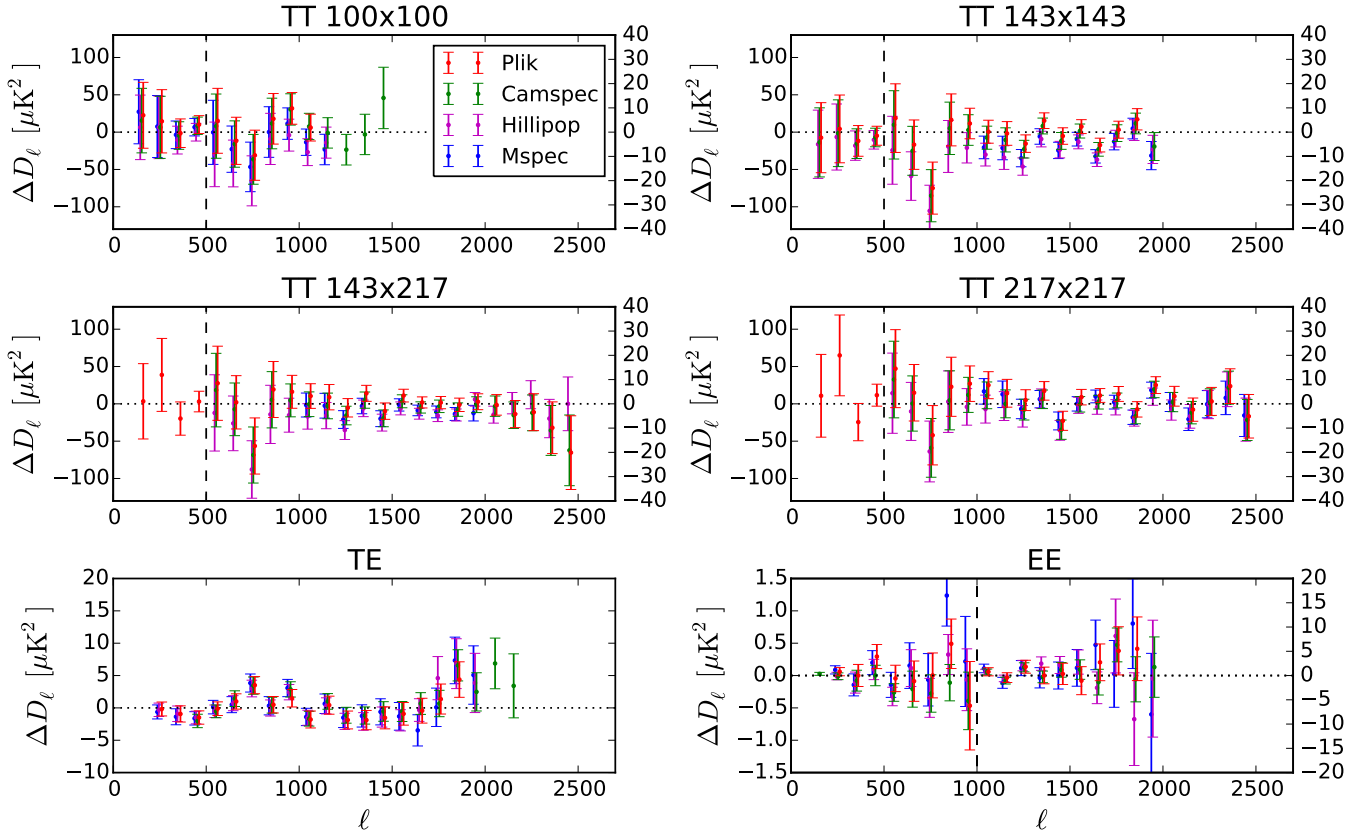
Fig. 38. Comparison of power spectra residuals from different high-ℓ likelihood codes. The figure shows "data/calib − FG − $\texttt{Plik}_{\mathrm{CMB}}$", where "data" stands for the empirical cross-frequency spectra, "FG" and "calib" are the best-fit foreground model and recalibration parameter for each individual code at that frequency, and the best-fit model $\texttt{Plik}_{\mathrm{CMB}}$ is subtracted for visual presentation. These plots thus show the difference in the amount of power each code attributes to the CMB. The power spectra are binned in bins of width $\Delta\ell = 100$. The y-axis scale changes at $\ell = 500$ for *TT* and $\ell = 1000$ for *EE* (vertical dashes).

of this level do not lead to very large differences in cosmological parameters except in a few cases that we discuss.

For easier visual comparison of error bars, we show in Fig. 39 the ratios of each code's error bars to those from $\texttt{Plik}$. These have been binned in bins of width $\Delta\ell = 100$, and are thus sensitive to the correlation structure of each code's covariance matrix, up to 100 multipoles into the off-diagonal. For all the codes and for both temperature and polarization, the correlation between multipoles separated by more than $\Delta\ell = 100$ is less than 3%, so Fig. 39 contains the majority of the relevant information about each code's covariance.

A few differences are visible, mostly at high frequency, when the 217 GHz data are used. First, the $\texttt{Hillipop}$ error bars in *TT* for 143×217 become increasingly tighter than the other codes at $\ell > 1700$. This is because $\texttt{Hillipop}$, unlike the other codes, gives non-zero weight to 143×217 spectra when both the 143 and the 217 GHz maps come from the same half-mission. This leads to a slight increase in power at high ℓ compared to $\texttt{Plik}$, as can be seen in Fig. 38. Conversely, the $\texttt{Hillipop}$ error bars are slightly larger by a few percent at $\ell < 1700$; however the source of this difference is not understood. Second, the $\texttt{Mspec}$ error bars in temperature are increasingly tighter towards higher frequency, as compared to other codes; for 217×217, $\texttt{Mspec}$ uncertainties are smaller by 5−10% for ℓ between 1000 and 2000. This arises from the $\texttt{Mspec}$ map-based Galactic cleaning procedure, which removes excess variance due to CMB–foreground correlations by subtracting a scaled 545 GHz map. However, for polarization, where one must necessarily clean with the noisier

353 GHz maps, the $\texttt{Mspec}$ error bars for *TE* and *EE* become larger. $\texttt{CamSpec}$, which also performs a map cleaning for low-ℓ polarization, switches to a power-spectrum cleaning at higher ℓ to mitigate this effect.

The differences in ΛCDM parameters from *TT* are shown in Table 19. Generally, parameters agree to within a fraction of σ, but with some differences we discuss. One thing to keep in mind in interpreting this comparison is that these differences are not necessarily indicative of systematic errors. Some of the differences are expected due to statistical fluctuations because different codes weight the data differently.

One of the biggest differences with respect to the baseline code is in n_{s}, which is higher by about 0.45σ for $\texttt{CamSpec}$, with a related downward shift of $A_{\mathrm{s}}\mathrm{e}^{-2\tau}$. To put these shifts into perspective, we refer to the whisker plots of Figs. 35 and 36 which compare $\texttt{CamSpec}$ *TT* results with $\texttt{Plik}$ in the ΛCDM case (base and extended). A difference in n_{s} of about 0.16σ between $\texttt{Plik}$ and $\texttt{CamSpec}$ can be attributed to the inclusion in $\texttt{Plik}$ of the first 500 multipoles for 143×217 and 217×217; these multipoles are excluded in $\texttt{CamSpec}$ (see also Sect. 4.1.2). Indeed, cutting out those multipoles in $\texttt{Plik}$ brings n_{s} closer by 0.16σ to the $\texttt{CamSpec}$ value and slightly degrades the constraint on n_{s} compared to the full $\texttt{Plik}$ result. Using Eq. (53), we see that the shift and degradation in constraining power are consistent with expectations. A similar 0.16σ shift can be attributed to different dust templates. $\texttt{CamSpec}$ uses a steeper power law index (-2.7). Using the $\texttt{CamSpec}$ template in $\texttt{Plik}$ brings n_{s} closer to the $\texttt{CamSpec}$ value. Allowing the power law index of the galactic
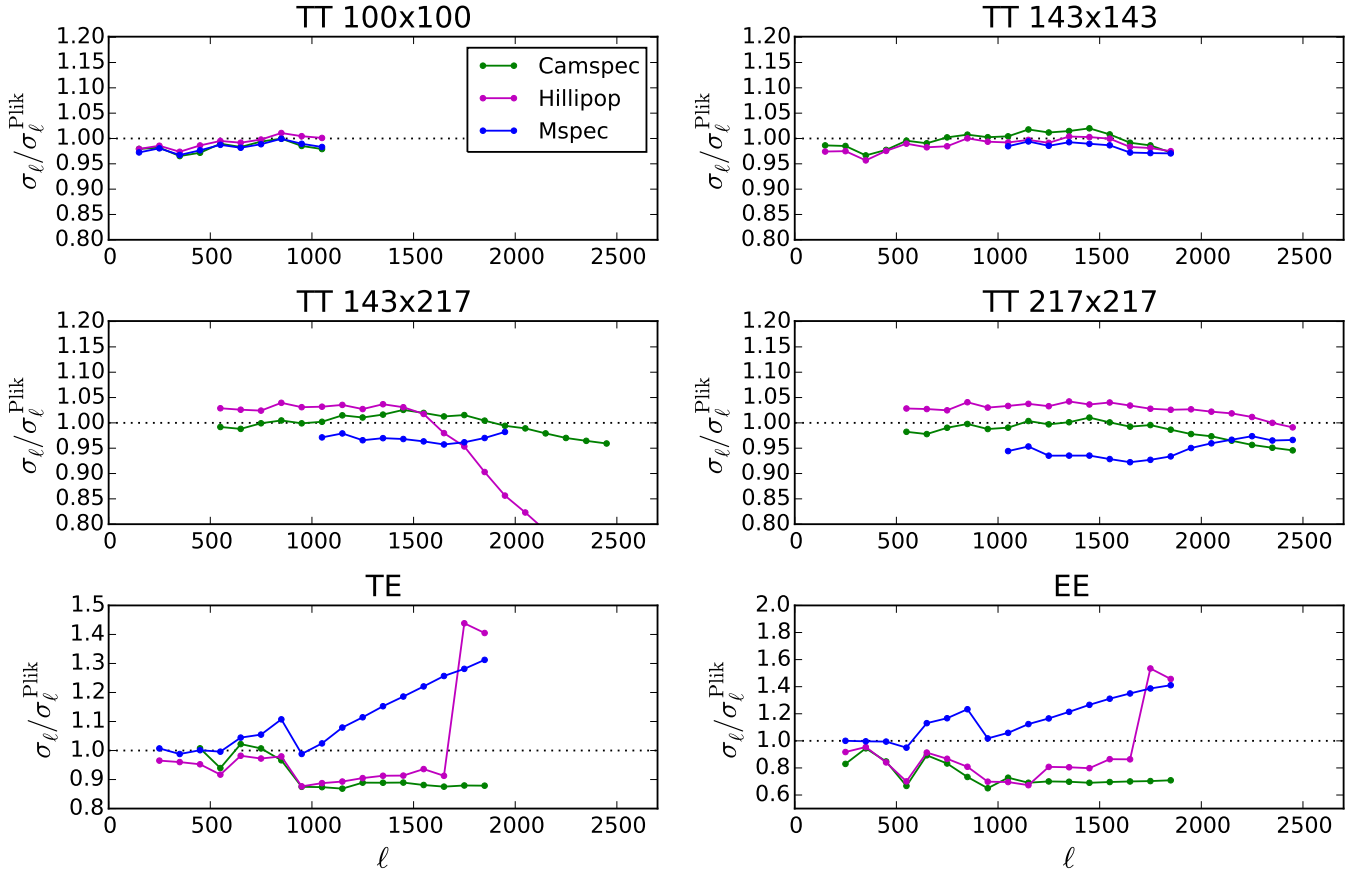
Fig. 39. Comparison of error bars from the different high-ℓ likelihood codes. The quantities plotted are the ratios of each code's error bars to those from `Plik`, and are for bins of width $\Delta\ell = 100$. Results are shown only in the ℓ range common to `Plik` and the code being compared.

template to vary when exploring cosmological parameters yields a slightly shallower slope (see Sect. 4.1.2). The slope of the dust template is mainly determined at relatively high ℓ, i.e., in the regime where it is hardest to determine the template accurately since the dust contribution is only a small fraction of the CIB and point-source contributions (see the $\ell \gtrsim 1000$ parts of Figs. 19 and 20). The remaining difference of 0.13σ arises from differences in data preparation (maps, calibration, binning) and covariance estimates. We therefore believe that a 0.2σ is a conservative upper bound of the systematic error in n_s associated with the uncertainties in the modelling of foregrounds, which is the biggest systematic uncertainty in TT.

A shift that is less well understood is the $\approx 1\sigma$ shift in $A_s e^{-2\tau}$ between `Plik` and `Hillipop`. The preference for a lower amplitude from `Hillipop` is sourced by the lower power attributed to the CMB, seen in Fig. 38. With τ partially fixed by the prior, this implies lower A_s and hence a smaller lensing potential envelope, explaining the somewhat lower value of A_L found by `Hillipop`. Tests performed with the same code suggest that 1σ is too great a shift to be explained simply by the different foreground models, so some part of it must be due to the different data weighting; as can be seen in Fig. 39, `Hillipop` gives less weight to $500 \lesssim \ell \lesssim 1500$, and slightly more outside of this region.

This comparison also shows the stability of the results with respect to the Galactic cleaning procedure. `Mspec` and `Plik` use different procedures, yet their parameter estimates agree to better than 0.5σ (see Appendix D.1). But we note that the `Plik`–`CamSpec` differences are higher in the polarization case, and can reach 1σ, as can be judged from the whisker plot in polarization of Fig. C.10.

4.3. Consistency of Poisson amplitudes with source counts

The Poisson component of the foreground model is sourced by shot-noise from astrophysical sources. In this section we discuss the consistency between the measured Poisson amplitudes and other probes and models of the source populations from which they arise. The Poisson amplitude priors that we calculate are not used in the main analysis, because they improve uncertainties on the cosmological parameters by at most 10%, and only for a few extensions; instead they serve as a self-consistency check.

This type of check was also performed in Like13, which we update here by:

1. developing a new method for calculating these priors that is accurate enough to give realistic uncertainties on Poisson predictions (for the first time);
2. including a comparison of more theoretical models;
3. taking into account the 2015 point-source masks.

In Like13 the Poisson power predictions were calculated via

$$C_\ell = \int_0^{S_{\rm cut}} \mathrm{d}S\, S^2 \frac{\mathrm{d}N}{\mathrm{d}S}, \tag{55}$$

where $\mathrm{d}N/\mathrm{d}S$ is the differential number count, $S_{\rm cut}$ is an effective flux-density cut above which sources are masked, and the integral was evaluated independently at each frequency. Although it is adequate for rough consistency checks, Eq. (55) ignores the facts that the 2013 point-source mask was built from a union of sources detected at different frequencies, and that the *Planck* flux-density cut varies across the sky, and it also ignores the effect of Eddington bias. In order to accurately account for all of

Table 19. Comparison between the parameter estimates from different high-ℓ codes.

Parameter	Plik	CamSpec	Hillipop	Mspec	Xfaster (SMICA)
$\Omega_{\mathrm{b}}h^2$	0.02221 ± 0.00023	0.02224 ± 0.00023	0.02218 ± 0.00023	0.02218 ± 0.00024	0.02184 ± 0.00024
$\Omega_{\mathrm{c}}h^2$	0.1203 ± 0.0023	0.1201 ± 0.0023	0.1201 ± 0.0022	0.1204 ± 0.0024	0.1202 ± 0.0023
$100\theta_{\mathrm{MC}}$	1.0406 ± 0.00047	1.0407 ± 0.00048	1.0407 ± 0.00046	1.0409 ± 0.00050	1.041 ± 0.0005
τ	0.085 ± 0.018	0.087 ± 0.018	0.075 ± 0.019	0.075 ± 0.018	0.069 ± 0.019
$10^9 A_{\mathrm{s}}\mathrm{e}^{-2\tau}$	1.888 ± 0.014	1.877 ± 0.014	1.870 ± 0.011	1.878 ± 0.012	1.866 ± 0.015
n_{s}	0.962 ± 0.0063	0.965 ± 0.0066	0.961 ± 0.0072	0.959 ± 0.0072	0.960 ± 0.0071
Ω_{m}	0.3190 ± 0.014	0.3178 ± 0.014	0.3164 ± 0.014	0.3174 ± 0.015	0.3206 ± 0.015
H_0	67.0 ± 1.0	67.1 ± 1.0	67.1 ± 1.0	67.1 ± 1.1	66.8 ± 1.0

Notes. Each column gives the results for various high-ℓ *TT* likelihoods at $\ell > 50$ when combined with a prior of $\tau = 0.07 \pm 0.02$. The SMICA parameters were obtained for $\ell_{\mathrm{max}} = 2000$.

these effects, we now calculate the Poisson power as

$$C_\ell^{ij} = \int_0^\infty \mathrm{d}S_1 \ldots \mathrm{d}S_n \, S_i S_j \frac{\mathrm{d}N(S_1,\ldots,S_n)}{\mathrm{d}S_1 \ldots \mathrm{d}S_n} \, I(S_1,\ldots,S_n), \quad (56)$$

where the frequencies are labelled $1\ldots n$, the differential source count model, $\mathrm{d}N/\mathrm{d}S$, is now a function of the flux densities at *each* frequency, and $I(S_1,\ldots,S_n)$ is the joint "incompleteness" of our catalogue for the particular cut that was used to build the point-source mask.

The joint incompleteness was determined by injecting simulated point sources into the *Planck* sky maps, using the procedure described in Planck Collaboration XXVI (2016). The same point-source detection pipelines that were used to produce the Second *Planck* Catalogue of Compact Sources (PCCS2) were run on the injected maps, producing an ensemble of simulated *Planck* sky catalogues with realistic detection characteristics. The joint incompleteness is defined as the probability that a source would not be included in the mask as a function of the source flux density, given the specific masking thresholds being considered. The raw incompleteness is a function of sky location, because the *Planck* noise varies across the sky. The incompleteness that appears in Eq. (56) is integrated over the region of the sky used in the analysis; the injection pipeline estimates this quantity by injecting sources only into these regions.

Equation (56) can be applied to any theoretical model which makes a prediction for the multi-frequency $\mathrm{d}N/\mathrm{d}S$. We have adopted the following models.

1. For radio galaxies we have two models. The first is the Tucci et al. (2011) model, updated to include new source-count measurements from Mocanu et al. (2013). We also consider a phenomenological model that is a power law in flux density and frequency, and assumes that the sources' spectral indices are Gaussian-distributed with mean $\bar{\alpha}$ and standard deviation σ_α; we use different values for $\bar{\alpha}$ and σ_α above and below 143 GHz. We shall refer to this second model as the "power-law" model, and the differential source counts are given by

$$\frac{\mathrm{d}N(S_1,S_2,S_3)}{\mathrm{d}S_1\mathrm{d}S_2\mathrm{d}S_3} = \frac{A(S_1 S_2 S_3)^{\gamma-1}}{2\pi\sigma_{12}\sigma_{23}} \quad (57)$$
$$\times \exp\left[-\frac{(\alpha(S_1,S_2)-\bar{\alpha}_{12})^2}{2\sigma_{12}^2} - \frac{(\alpha(S_2,S_3)-\bar{\alpha}_{32})^2}{2\sigma_{32}^2}\right],$$

where labels $1-3$ refer to *Planck* 100, 143, and 217 GHz and $\alpha(S_i,S_j) = \ln(S_j/S_i)/\ln(\nu_j/\nu_i)$. Both radio models are excellent fits to the available source-count data, and we take

the difference between them as an estimate of model uncertainty. With the power-law model we are additionally able to propagate uncertainties in the source count data to the final Poisson estimate via MCMC.

2. For dusty galaxies we use the Béthermin et al. (2012) model, as in Planck Collaboration XXX (2014). The model is in good agreement with the number counts measured with te *Spitzer* Space Telescope and the *Herschel* Space Observatory. It also gives a reasonable CIB redshift distribution, which is important for cross-spectra, and is a very good fit to CIB power spectra (see Béthermin et al. 2013). In contrast to the radio-source case, the major contribution to the dusty galaxy Poisson power arises from sources with flux densities well below the cuts; for example, we note that decreasing the flux-density cuts by a factor of 2 decreases the Poisson power by less than 1% at the relevant frequencies. In this case, Eq. (55) is a sufficient and more convenient approximation, and we make use of it when calculating Poisson levels for dusty galaxies.

We give predictions for Poisson levels for three different masks: (1) the 2013 point-source mask, which was defined for sources detected at $S/N > 5$ at any frequency between 100 and 353 GHz; (2) the 2015 point-source mask, which is frequency-dependent and includes $S/N > 5$ sources detected only at each individual frequency (used by Plik, CamSpec, and Mspec in this work); and (3) the Hillipop mask, which is also frequency-dependent and involves both a S/N cut and a flux-density cut[16].

Table 20 summarizes the main results of this section. Generally, we find good agreement between the priors from source counts and the posteriors from chains, with the priors being much more constraining. The exception to the good agreement is at 100 GHz where the prediction is lower than the measured value by around 4σ for the baseline 2015 mask and 6σ for the Hillipop mask. This is a sign either of a foreground modelling error or (perhaps more likely) of a residual unmodelled systematic in the data. We note that this disagreement was not present in Like13, where the Poisson amplitude at 100 GHz was found to be smaller. We also note that removing the relative calibration prior (Eq. (35)) or increasing the ℓ_{max} at 100 GHz by a few hundred reduces the tension in the Mspec results. In any case, it is unlikely to affect parameter estimates at all, since very little

---

[16] We note that the Hillipop mask was constructed partly so that Eq. (55) would be an accurate approximation. We find that for the radio contribution is is accurate to 2%, or 1σ, and for the dust contribution it is essentially exact.

Table 20. Priors on the Poisson amplitudes given a number of different point-source masks and models.

Mask	Type	Model	Power spectrum			
			100×100	143×143	143×217	217×217
Baseline 2013	Radio	Power-law	84 ± 3	29 ± 1	16 ± 1	9 ± 1
	Dusty	Bethermin	4 ± 1	13 ± 3	41 ± 8	129 ± 25
Baseline 2015	Radio	Power-law	148 ± 7	40 ± 1	16 ± 1	10 ± 1
		Tucci	139	40	16	11
	Dusty	Bethermin	4 ± 1	13 ± 3	41 ± 8	129 ± 25
		Plik	260 ± 28	44 ± 8	39 ± 10	97 ± 11
		Mspec	317 ± 46	22 ± 13	12 ± 7	21 ± 9
Hillipop 2015	Radio	Power-law	150 ± 7	47 ± 2	18 ± 1	11 ± 1
		Tucci	141	47	18	12
	Dusty	Bethermin	4 ± 1	13 ± 3	41 ± 8	129 ± 25
		Hillipop	372 ± 38	58 ± 21	53 ± 24	105 ± 18

Notes. Entries are $\mathcal{D}_\ell$ at $\ell = 3000$ in μK^2 and are given at the effective band centre for each component. Uncertainties on the "power-law" model are statistical errors propagated from uncertainties in the Mocanu et al. (2013) source-count data. Priors on the dust component have formally been calculated only for the Hillipop mask, but they are repeated for the other masks, for which they are accurate to better than 1%. The results from different codes, to which these predictions should be compared, use TT $\ell > 50$ data with a prior of $\tau = 0.07 \pm 0.02$. For Mspec about 90% of the dusty contribution is cleaned out at the map level, hence the measured values above are in some cases far less than prior value.

cosmological information comes from the multipole range at 100 GHz that constrains the Poisson amplitude.

4.4. TE and EE test results

4.4.1. Residuals per frequency and inter-frequency differences

Figure 40 shows the residuals for each frequency and Fig. 41 shows the differences between frequencies of the *TE* and *EE* power spectra (the procedure is explained in Appendix C.3.2). The residuals are calculated with respect to the best-fit cosmology as preferred by PlikTT+tauprior, although we use the best-fit solution of the PlikTT, TE, EE+tauprior run to subtract the polarized Galactic dust contribution.

The binned inter-frequency residuals show deviations at the level of a few μK^2 from the best-fit model. These deviations do not necessarily correspond to high values of the χ^2 calculated on the unbinned data (see Table 16). This is because some of the deviations are relatively small for the unbinned data and correctly follow the expected χ^2 distribution. However, if the deviations are biased (e.g., have the same sign) in some ℓ range, they can result in larger deviations (and large χ^2) after binning. Thus, the χ^2 calculated on unbinned data is not always sufficient to identify these type of biases. We therefore also use a second quantity, χ, defined as the weighted linear sum of residuals, to diagnose biased multipole regions or frequency spectra:

$$\chi = \boldsymbol{w}^\mathsf{T}(\hat{\boldsymbol{C}} - \boldsymbol{C}) \quad \text{with} \quad \boldsymbol{w} = (\mathrm{diag}\,\mathsf{C})^{-1/2}, \tag{58}$$

where $\hat{\boldsymbol{C}}$ is the unbinned vector of data in the multipole region or frequency spectrum of interest, $\boldsymbol{C}$) is the corresponding model, and $\boldsymbol{w}$ is a vector of weights, equal to the inverse standard deviation evaluated from the diagonal of the corresponding covariance matrix C. The χ statistic is distributed as a Gaussian with zero mean and standard deviation equal to

$$\sigma_\chi = \sqrt{\boldsymbol{w}^\mathsf{T}\mathsf{C}\boldsymbol{w}}. \tag{59}$$

We then define the normalized χ_{norm} as the χ in units of standard deviation,

$$\chi_{\mathrm{norm}} = \chi/\sigma_\chi. \tag{60}$$

The χ_{norm} values that we obtain for different frequency power spectra are given in Table 16.

For *EE*, the worst-behaved spectra from the χ_{norm} point of view are 143×143 (3.7σ deviation) and 100×217 (-3.0σ), while from the χ^2 point of view, the worst is 100×143 (PTE = 3.9%). For *TE*, the worst from the χ_{norm} point of view are 100×217 (5σ), 100×100 (3.7σ), and 143×143 (-2.2σ), while from the χ^2 point of view the worst is 100×100 (PTE = 0.43%). The extreme deviations from the expected distributions show that the frequency spectra are not described very accurately by our data model. This is also clear from Fig. 41, which shows that there are differences of up to 5σ between pairs of foreground-cleaned spectra.

However, as the co-added residuals in Fig. 29 show, systematic effects in the different frequency spectra appear to average out, leaving relatively small residuals with respect to the PlikTT+tauprior best-fit cosmology. In other words, these effects appear not to be dominated by common modes between detector sets or across frequencies. This is also borne out by the good agreement between the data and the expected polarization power spectra conditioned on the temperature ones, as shown in the conditional plots of Fig. 34.

4.4.2. TE and EE robustness tests

For *TE* and *EE*, we ran tests of robustness similar to those applied earlier to TT. These are presented in Appendix C.3.5, and the main conclusions are the following. We find that the Plik cosmological results are affected by less than 1σ when using detset cross-spectra instead of half-mission ones. This is also the case when we relax the dust amplitude priors, when we marginalize over beam uncertainties, or when we change ℓ_{min}
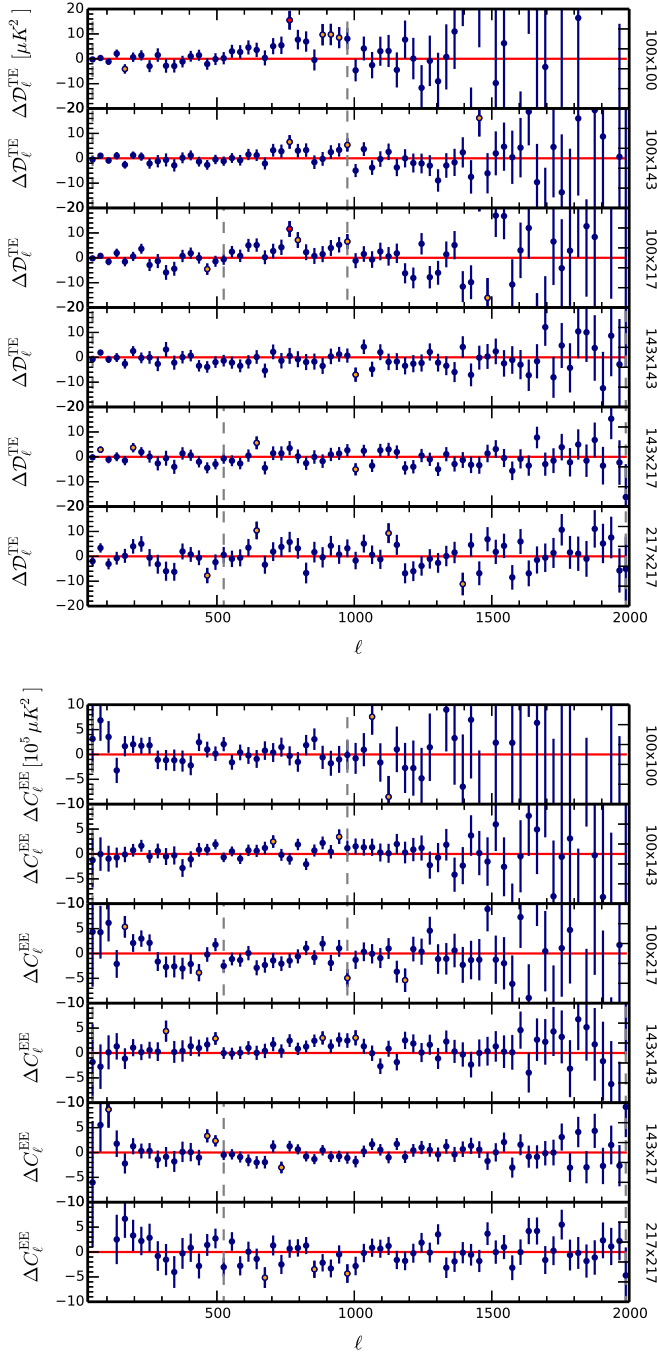
Fig. 40. Residual frequency power spectra after subtraction of the `PlikTT`+tauprior best-fit model. We clean Galactic dust from the spectra from using the best-fit solution of `PlikTT`, TE, EE+tauprior. The residuals are relative to the baseline HM power spectra (blue points, except for those that deviate by at least 2 or 3σ, which are shown in orange or red, respectively). The vertical dashed lines delimit the ℓ ranges retained in the likelihood. *Upper*: *TE* power spectra. *Lower*: *EE* power spectra.

Fig. 41. Inter-frequency foreground-cleaned power-spectra differences. Each panel shows the difference of two frequency power spectra, that indicated on the left axis minus that on the bottom axis, after subtracting foregrounds using the best-fit PlanckTT+lowP foreground solutions. Differences are shown for both the HM power spectra (dark blue) and the DS power spectra (light blue).

or ℓ_{max}. The alternative `CamSpec` likelihood has larger shifts, but still smaller than 1σ in *TE* and 0.5σ in *EE*. However, we also see larger shifts (more than 2σ in *TE*) with `Plik` when some frequency channels are dropped; and, when they are varied, the beam leakage parameters adopt much higher values than expected from the prior, while still leaving some small discrepancies between individual cross-spectra that have yet to be explained.

These results shows that our data model leaves residual instrumental systematic errors and is not yet sufficient to take advantage of the full potential of the HFI polarization information. Indeed, the current data model and likelihood code do not account satisfactorily for deviations at the μK^2 level, even if they can be captured in part by our beam leakage modelling. Nevertheless, the results for the ΛCDM model obtained from the `PlikTE`+tauprior and `PlikEE`+tauprior runs are in

good agreement with the results from `PlikTT+tauprior` (see Appendix C.3.6). This agreement between temperature and polarization results within ΛCDM is not a proof of the accuracy of the co-added polarization spectra and their data model, but rather a check of consistency at the μK^2 level. This consistency is, of course, a very interesting result in itself. But this comparison of probes cannot yet be pushed further to check for the potential presence of a physical inconsistency within the base model that the data could in principle detect or constrain.

5. The full *Planck* spectra and likelihoods

This section discusses the results that are obtained by using the full *Planck* likelihood. Section 5.1 first addresses the question of robustness with respect to the choice of the hybridization scale (the multipole at which we transition from the low-ℓ likelihood to the high-ℓ likelihood). Sections 5.2 and 5.3 then present the full results for the power spectra and the baseline cosmological parameters. Section 5.4 summarizes the full systematic error budget. Section 5.5 concentrates on the significance of the possibly anomalous structure around $\ell \approx 20$ in this new release. We then introduce in Sect. 5.6 a useful compressed *Planck* high-ℓ temperature and polarization CMB-only likelihood, `Plik_lite`, which, when applicable, enables faster parameter exploration. Finally, in Sect. 5.7, we compare the *Planck* 2015 results with the previous results from WMAP, ACT, and SPT.

5.1. Insensitivity to hybridization scale

Before we use the low-ℓ and high-ℓ likelihoods together, we address the question of the hybridization scale, ℓ_{hyb}, at which we switch from one to the other (neglecting correlations between the two regimes, as we did and checked in Like13). To that end, we focus on the TT case and use a likelihood based on the Blackwell-Rao estimator and the `Commander` algorithm (Chu et al. 2005; Rudjord et al. 2009) as described in Sect. 2.2, since this likelihood can be used to much higher ℓ_{max} than the full pixel-based T, E, B one. For this test without polarization data, we assume the same $\tau = 0.07 \pm 0.02$ prior as before.

The whisker plot of Fig. 42 shows the marginal mean and the 68% CL error bars for base-ΛCDM cosmological parameters when ℓ_{hyb} is varied from the baseline value of 30 (case "LOWL 30") to ℓ_{hyb} =50, 100, 150, 200, and 250, and compared to the `PlikTT+tauprior` case. The difference between the "LOWL 30" and "PlikTT+tauprior" values shows the effect of the low-ℓ dip at $\ell \approx 20$, which reaches 0.5σ on n_s. The plot shows that the effect of varying ℓ_{hyb} from 30 to 150 is a shift in n_s by less than 0.1σ. This is the result of the Gaussian approximation pushed to $\ell_{\text{min}} = 30$, already discussed in the simulation section (Sect. 3.6). It would have been much too slow to run the full low-ℓ TEB likelihood with ℓ_{max} substantially greater than 30, and we decided against the only other option, to leave a gap in polarization between $\ell = 30$ and the hybridization scale chosen in TT.

5.2. The Planck 2015 CMB spectra

The visual appearance of *Planck* 2015 CMB co-added spectra in TT, TE, and EE can be seen in Fig. 50. Goodness-of-fit values can be found in Table E.1 of Appendix E. These differ somewhat from those given previously in Table 16 for `Plik` alone, because the inclusion of low ℓ in temperature brings in the $\ell \approx 20$ feature (see Sect. 5.5). Still, they remain acceptable, with PTEs all above 10% (16.8% for TT).

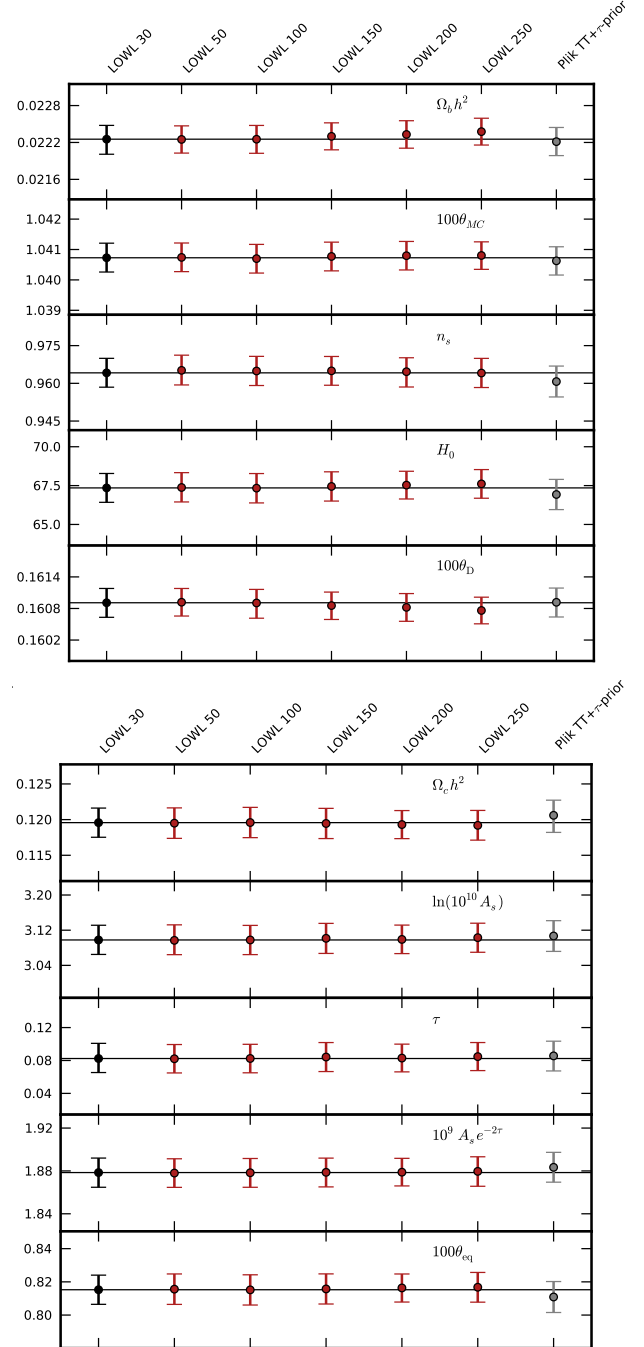**Fig. 42.** Marginal mean and 68% CL error bars on cosmological parameters estimated with different multipoles for the transition between the low-ℓ and the high-ℓ likelihood. Here we use only the TT power spectra and a Gaussian prior on the optical depth $\tau = 0.07 \pm 0.02$, within the base-ΛCDM model. "PlikTT+tauprior" refers to the case where we use the `Plik` high-ℓ likelihood only.

With this release, *Planck* now detects 36 extrema in total, consisting of 19 peaks and 17 troughs. Numerical values for the positions and amplitudes of these extrema may be found in Table E.2 of Appendix E.2, which also provides details of the steps taken to derive them. We provide in Appendix E.3 an alternate display of the correlation between temperature and (*E*-mode) polarization by showing their Pearson correlation coefficient and their decorrelation angle versus scale (Figs. E.2 and E.3).

5.3. Planck 2015 model parameters

Figure 43 compares constraints on pairs of parameters as well as their individual marginals for the base-ΛCDM model. The grey contours and lines correspond to the results of the 2013 release (Like13), which was based on TT and WMAP polarization at low ℓ (denoted by WP), using only the data from the nominal mission. The blue contours and lines are derived from the 2015 baseline likelihood, `PlikTT+lowTEB` ("PlanckTT+lowP" in the plot), while the red contours and line are obtained from the full `PlikTT, EE, TE+lowTEB` likelihood ("PlanckTT, TE, EE+lowP" in the plot, see Appendix E.1 for the relevant robustness tests). In most cases the 2015 constraints are in quite good agreement with the earlier constraints, with the exception of the normalization A_s, which is higher by about 2%, reflecting the 2015 correction of the *Planck* calibration which was indeed revised upward by about 2% in power. The figure also illustrates the consistency and further tightening of the parameter constraints brought by adding the E-mode polarization at high ℓ. The numerical values of the *Planck* 2015 cosmological parameters for base ΛCDM are given in Table 21.

As shown in Fig. 44, the degeneracies between foreground and calibration parameters generally do not affect the determination of the cosmological parameters. In the `PlikTT+lowTEB` case (top panel), the dust amplitudes appear to be nearly uncorrelated with the basic ΛCDM parameters. Similarly, the 100 and 217 GHz channel calibration is only relevant for the level of foreground emission. Cosmological parameters are, however, mildly correlated with the point-source and kinetic SZ amplitudes. Correlations are strongest (up to 30%) for the baryon density ($\Omega_b h^2$) and spectral index (n_s). We do not show correlations with the *Planck* calibration parameter (y_P), which is uncorrelated with all the other parameters except the amplitude of scalar fluctuations (A_s). The bottom panel shows the correlation for the `PlikTE+lowTEB` and `PlikEE+lowTEB` cases, which do not affect the cosmological parameters, except for 20% correlations in EE between the spectral index (n_s) and the dust contamination amplitude in the 100 and 143 GHz maps.

We also display in Fig. 45 the correlations between the foreground parameters and the cosmological parameters in the `PlikTT+lowTEB` case when exploring classical extensions to the ΛCDM model. While n_{run} seems reasonably insensitive to the foreground parameters, some extensions do exhibit a noticeable correlation, up to 40% in the case of Y_{He} and the point-source level at 143 GHz.

Finally, we note that power spectra and parameters derived from CMB maps obtained by the component-separation methods described in Planck Collaboration IX (2016) are generally consistent with those obtained here, at least when restricted to the $\ell < 2000$ range in TT; this is detailed in Sect. E.4.

5.4. Overall systematic error budget assessment

The tests presented throughout this paper and its appendices documented our numerous tests of the *Planck* likelihood code and its outputs. Here, we summarize those results and attempt to isolate the dominant sources of systematic uncertainty. This assessment is of course a difficult task. Indeed, all known systematics are normally corrected for, and when relevant, the uncertainty on the correction is included in the error budget and thus in the error bar we report. In that sense, except for a very few cases where we decided to leave a known uncertainty in the data, this section tries to deal with the more difficult task of evaluating the unknown uncertainty!

Table 21. Constraints on the basic six-parameter ΛCDM model using *Planck* angular power spectra.

Parameter	PlanckTT+lowP 68% limits	PlanckTT, TE, EE+lowP 68% limits
$\Omega_b h^2$	0.02222 ± 0.00023	0.02225 ± 0.00016
$\Omega_c h^2$	0.1197 ± 0.0022	0.1198 ± 0.0015
$100\theta_{MC}$	1.04085 ± 0.00047	1.04077 ± 0.00032
τ	0.078 ± 0.019	0.079 ± 0.017
$\ln(10^{10} A_s)$	3.089 ± 0.036	3.094 ± 0.034
n_s	0.9655 ± 0.0062	0.9645 ± 0.0049
H_0	67.31 ± 0.96	67.27 ± 0.66
Ω_Λ	0.685 ± 0.013	0.6844 ± 0.0091
Ω_m	0.315 ± 0.013	0.3156 ± 0.0091
$\Omega_m h^2$	0.1426 ± 0.0020	0.1427 ± 0.0014
$\Omega_m h^3$	0.09597 ± 0.00045	0.09601 ± 0.00029
σ_8	0.829 ± 0.014	0.831 ± 0.013
$\sigma_8\Omega_m^{0.5}$	0.466 ± 0.013	0.4668 ± 0.0098
$\sigma_8\Omega_m^{0.25}$	0.621 ± 0.013	0.623 ± 0.011
z_{re}	$9.9^{+1.8}_{-1.6}$	$10.0^{+1.7}_{-1.5}$
$10^9 A_s$	$2.198^{+0.076}_{-0.085}$	2.207 ± 0.074
$10^9 A_s e^{-2\tau}$	1.880 ± 0.014	1.882 ± 0.012
Age/Gyr	13.813 ± 0.038	13.813 ± 0.026
z_*	1090.09 ± 0.42	1090.06 ± 0.30
r_*	144.61 ± 0.49	144.57 ± 0.32
$100\theta_*$	1.04105 ± 0.00046	1.04096 ± 0.00032
z_{drag}	1059.57 ± 0.46	1059.65 ± 0.31
r_{drag}	147.33 ± 0.49	147.27 ± 0.31
k_D	0.14050 ± 0.00052	0.14059 ± 0.00032
z_{eq}	3393 ± 49	3395 ± 33
k_{eq}	0.01035 ± 0.00015	0.01036 ± 0.00010
$100\theta_{s,eq}$	0.4502 ± 0.0047	0.4499 ± 0.0032
f_{2000}^{143}	29.9 ± 2.9	29.5 ± 2.7
$f_{2000}^{143\times217}$	32.4 ± 2.1	32.2 ± 1.9
f_{2000}^{217}	106.0 ± 2.0	105.8 ± 1.9

Notes. The top group contains constraints on the six primary parameters included directly in the estimation process. The middle group contains constraints on derived parameters. The last group gives a measure of the total foreground amplitude (in μK^2) at $\ell = 2000$ for the three high-ℓ temperature spectra used by the likelihood. These results were obtained using the `CAMB` code, and are identical to the ones reported in Table 3 in Planck Collaboration XIII (2016).

This section summarizes the contribution of the known systematic uncertainties along with these potential unknown unknowns, specifically highlighting both internal consistency tests based on comparing subsets of the data, along with those using end-to-end instrumental simulations.

5.4.1. Low-ℓ budget

The low-ℓ likelihood has been validated using both internal consistency tests and simulation-based, tests. Here we summarize only the main result of the analysis, which has been set forth in Sect. 2 above.

A powerful consistency test of the polarization data, described in Sect. 2.4, is derived by rotating some of the likelihood components by $\pi/4$. Specifically, the rotation is applied to the
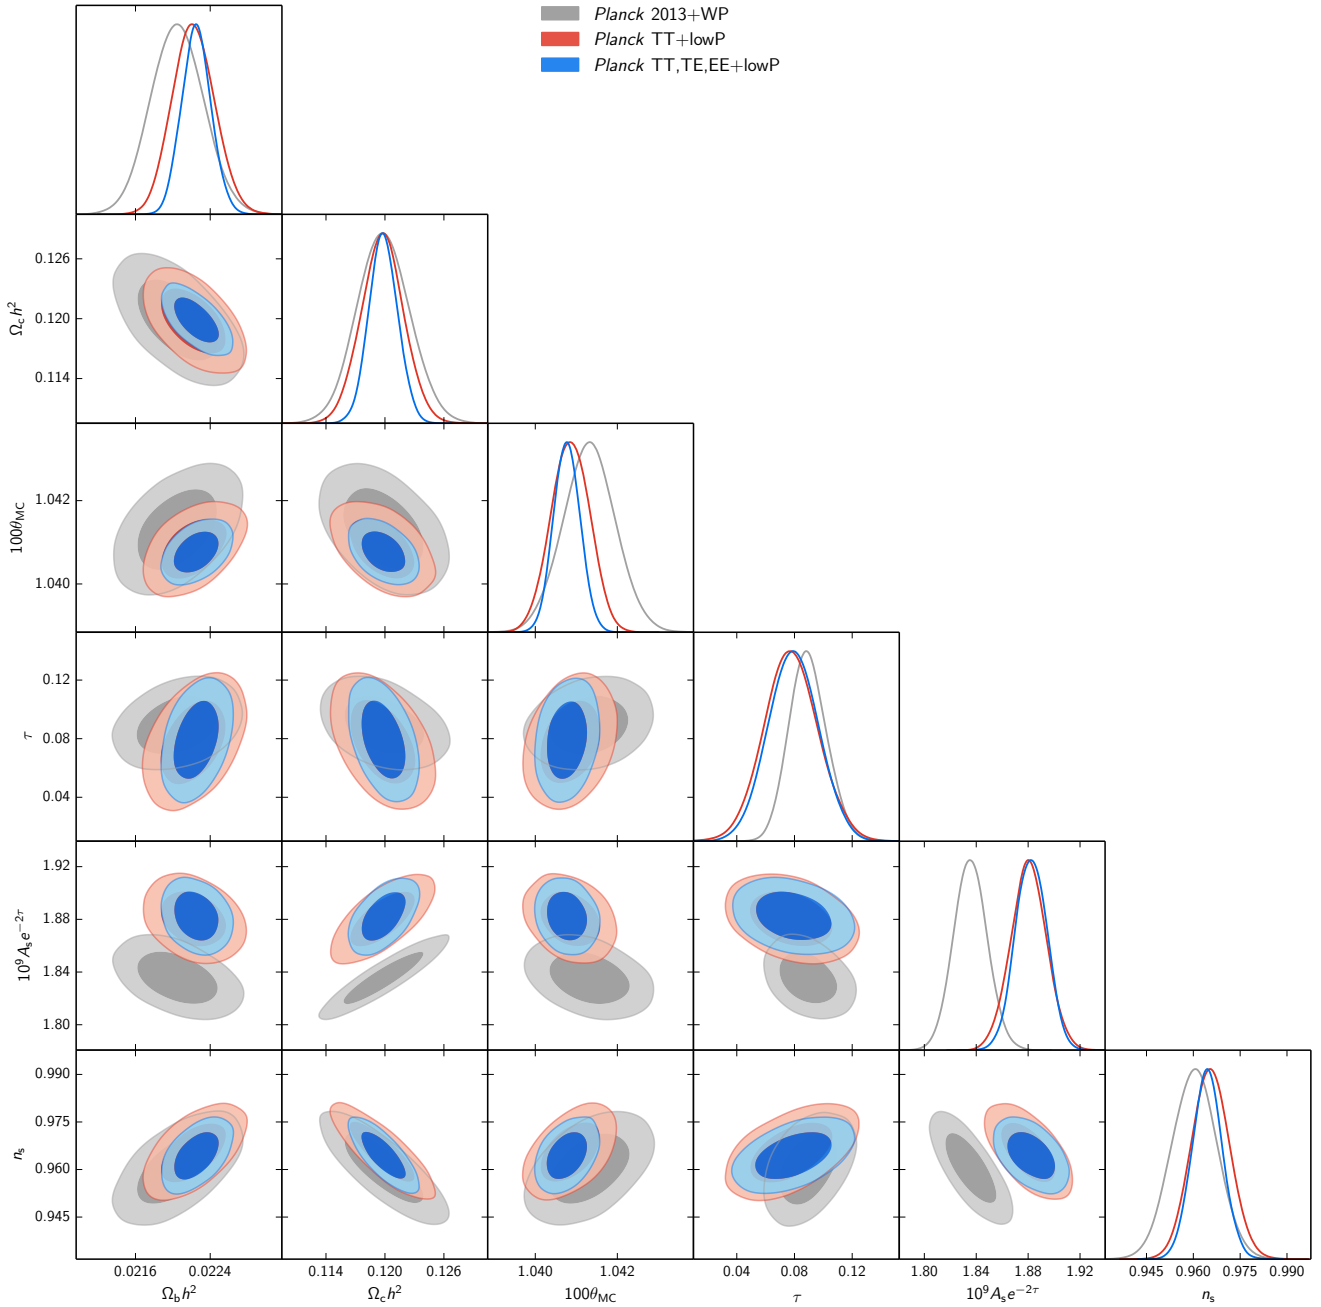
Fig. 43. ΛCDM parameter constraints. The grey contours show the 2013 constraints, which can be compared with the current ones, using either *TT* only at high ℓ (red) or the full likelihood (blue). Apart from further tightening, the main difference is in the amplitude, A_s, due to the overall calibration shift.

data maps and only to the noise covariance matrix (the likelihood being a scalar function, applying the same rotation to the signal matrix as well would be equivalent to not performing the rotation). The net effect is a conversion of $E \rightarrow -B$ and $B \rightarrow E$ for the signal, but leaving unaffected the Gaussian noise described in the covariance matrix. Under these circumstances, we do not expect to pick up any reionization signal, since it would then be present in *BB* or *TB*: the operation should result in a null τ detection. This is precisely what happens (see the blue dashed curve in Fig. 8). It is of course possible – though unlikely – that systematics are only showing up in the E channel, leaving B modes unaffected. Indeed, this possibility is further challenged by the fact that we do not detect anomalies in any of the six polarized power spectra; as detailed in Fig. 7, they are consistent with a

ΛCDM signal and noise as described by the final 70 GHz covariance matrix.

These tests are specific to *Planck* and aimed at validating the internal consistency of the datasets employed to build the likelihood. As a further measure of consistency, we have carried out a null test employing the WMAP data, detailed in Sect. 2.6. In brief, we have taken WMAP's K_a, Q and V channels and cleaned them from any polarized foreground contributions using a technique analogous to the one used to clean the LFI 70 GHz maps, employing the *Planck* 353 GHz map to minimize any dust contribution, but relying on WMAP's K channel to remove any synchrotron contribution. The resulting LFI 70 GHz and WMAP maps separately lead to compatible τ detections; their half-difference noise estimates are compatible with
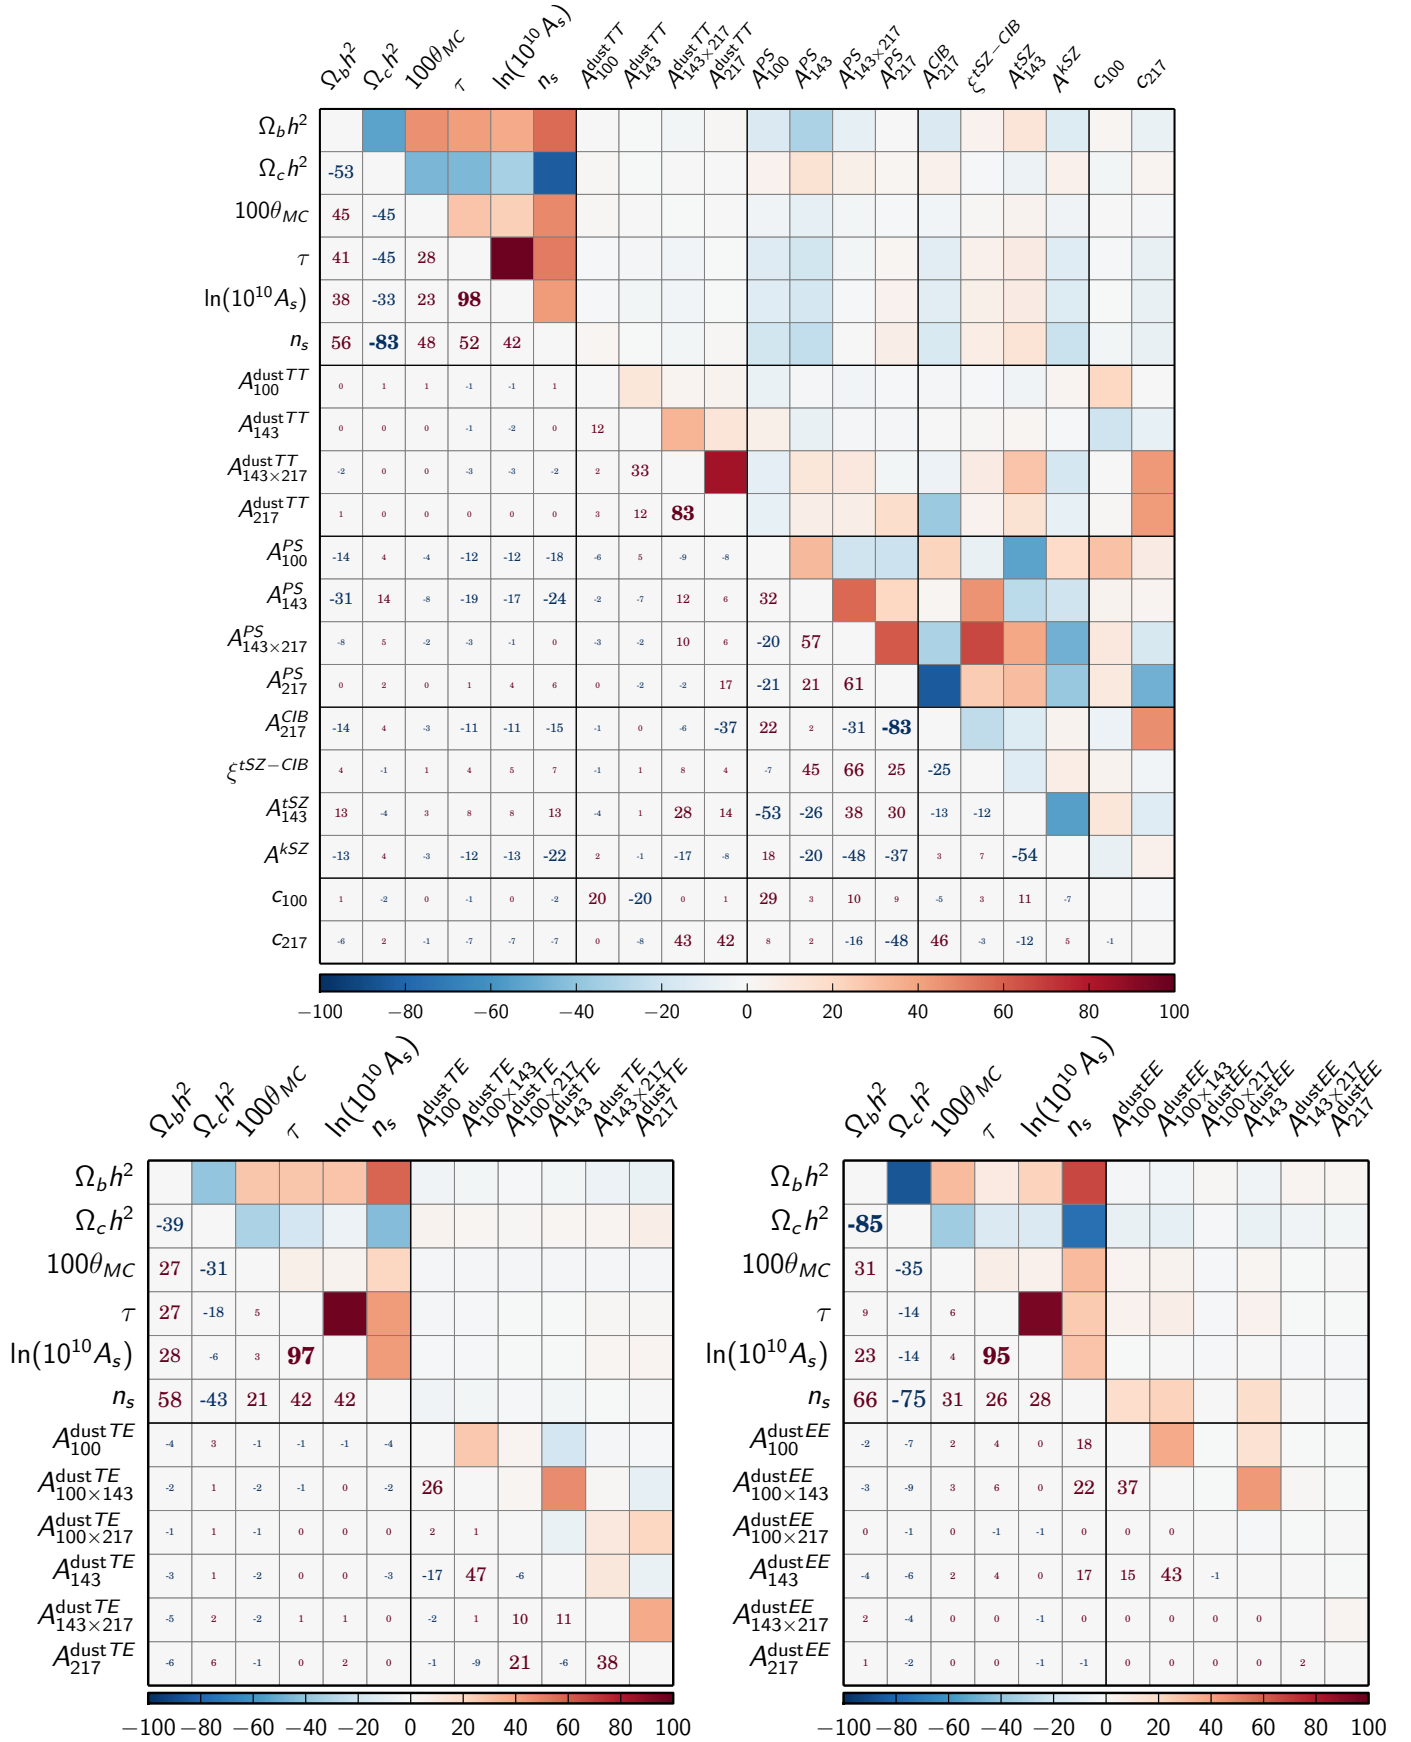
Fig. 44. Parameter correlations for `PlikTT+lowTEB` (*top*), `PlikTE+lowTEB` (*bottom left*), and `PlikEE+lowTEB` (*bottom right*). The degeneracies between foreground and calibration parameters do not strongly affect the determination of the cosmological parameters. In these figures the lower triangle gives the numerical values of the correlations in percent (with values below 10% printed at the smallest size), while the upper triangle represents the same values using a colour scale.
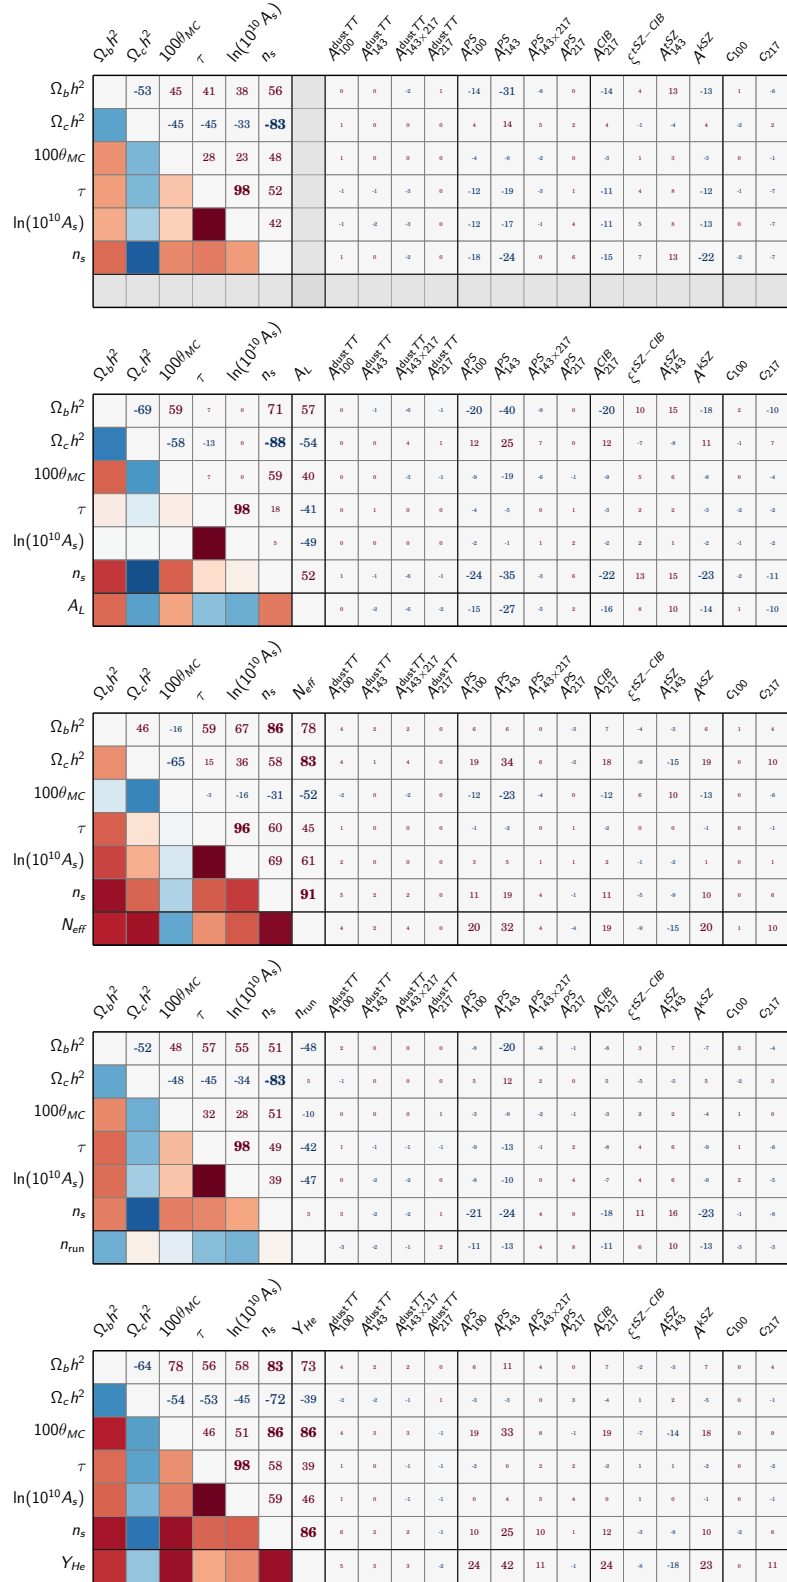
Fig. 45. Parameter correlations for `PlikTT+lowTEB`, including some ΛCDM extensions. The leftmost column is identical to Fig. 44 and is repeated here to ease comparison. Including extensions to the ΛCDM model changes the correlations between the cosmological parameters, sometimes dramatically, as can 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* interfrequency 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*, particularly 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 arbitrarly 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 evelute 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

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

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

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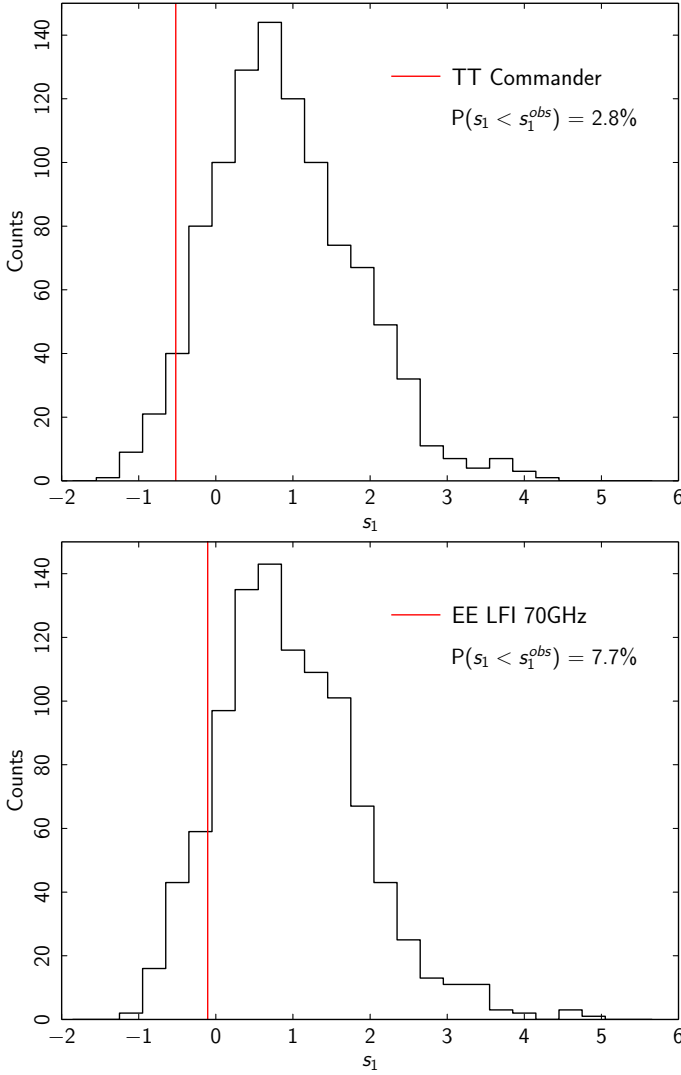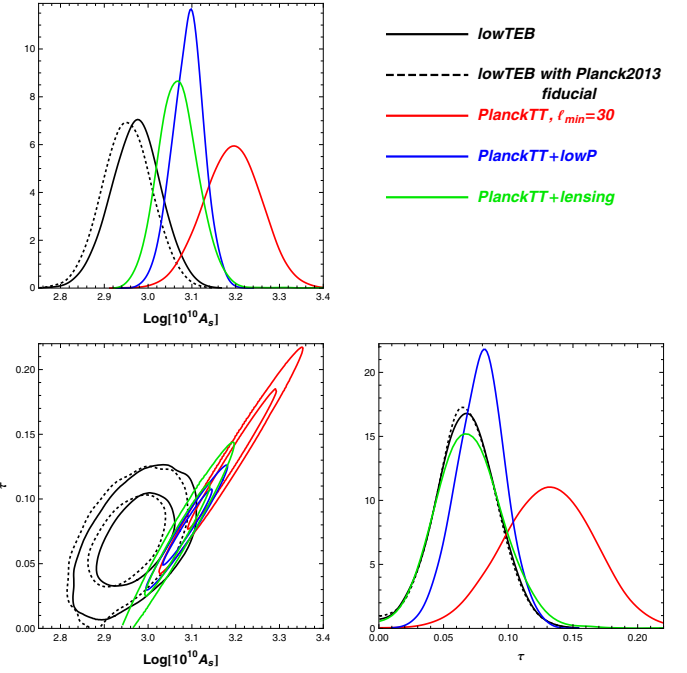
Fig. 47. Joint estimates of primordial amplitude A_s and τ for the data sets indicated in the legend. For low-ℓ estimates, all other parameters are fixed to the 2015 fiducial values, except for the dashed line, which uses the *Planck* 2013 fiducial. The PlanckTT+lowP estimates fall roughly half way between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

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_{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_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].

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	$[0, 0.11]$	$[0, 0.096]$	$[0, 0.10]$
$A_{\rm s}{\rm e}^{-2\tau}$	$1.878^{+0.010}_{-0.010}$	$1.879^{+0.011}_{-0.010}$	$1.879^{+0.010}_{-0.010}$

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 $40\,\mu{\rm K}^2$ at 70 GHz, $15\,\mu{\rm K}^2$ at 100 GHz, and $5\,\mu{\rm K}^2$ 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
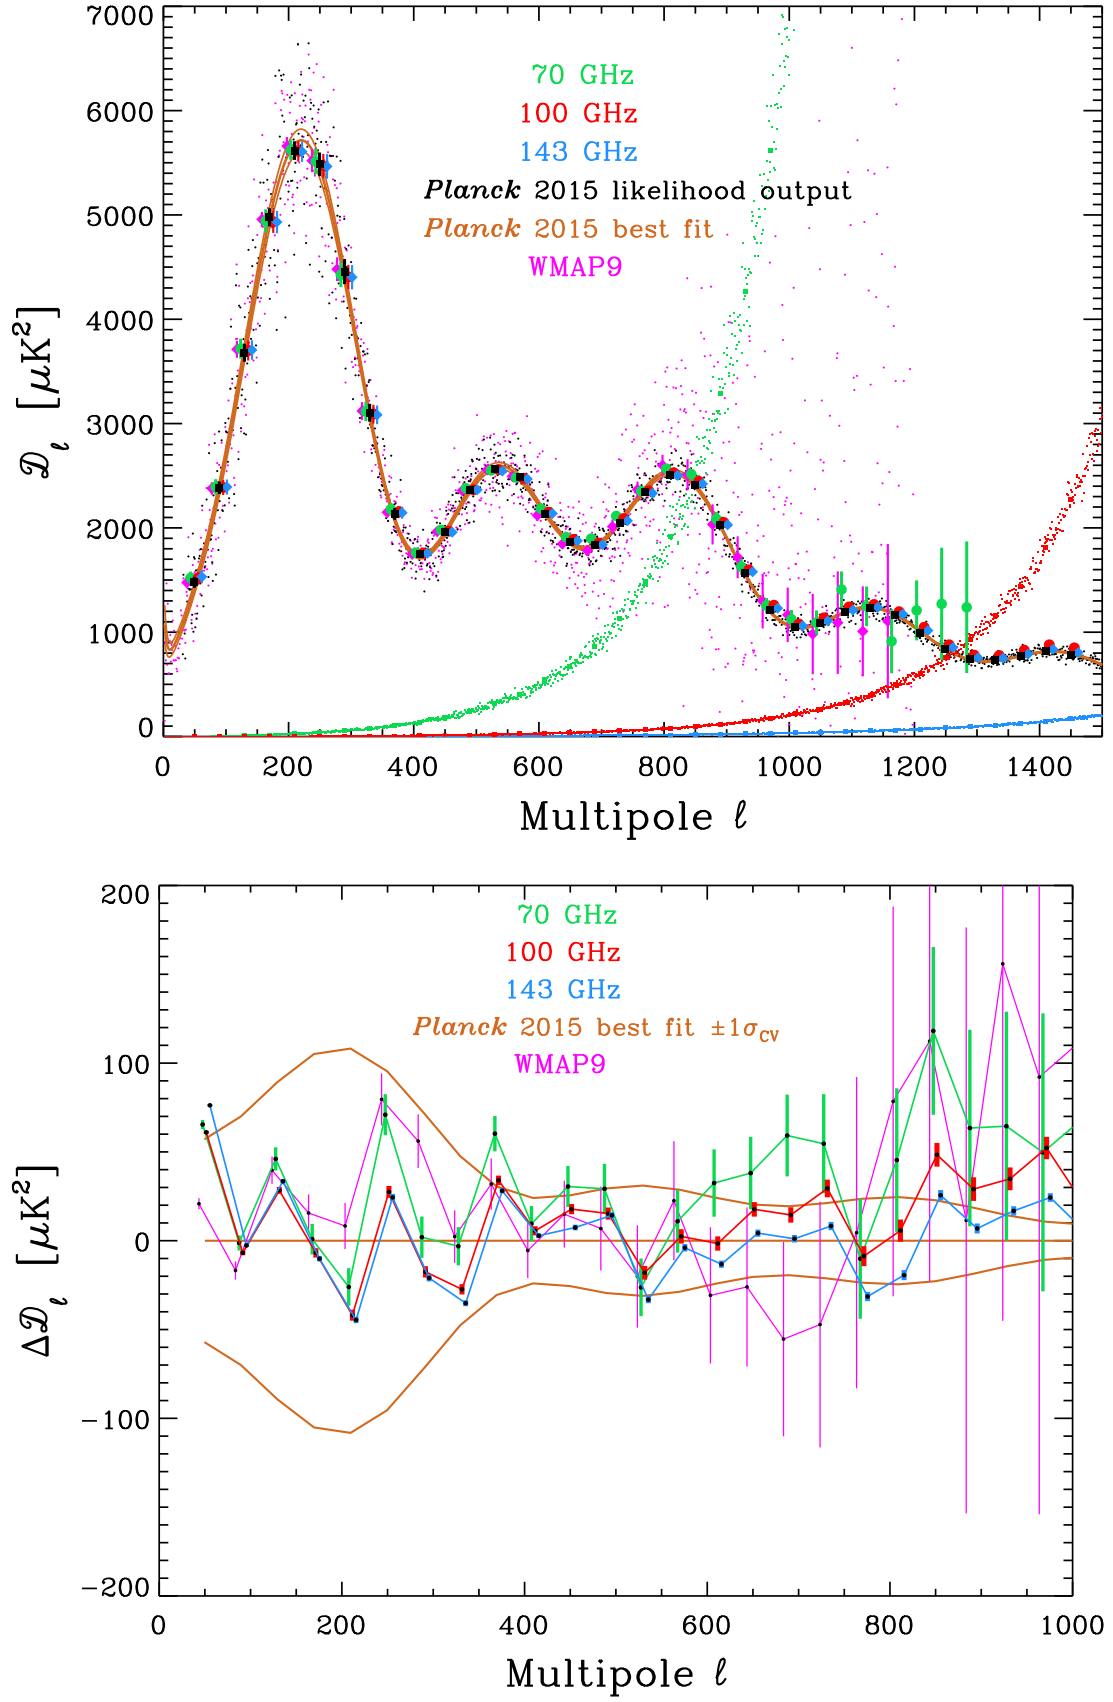
Fig. 48. Comparison of *Planck* and WMAP-9 CMB power spectra. *Top*: direct comparison. Noise spectra are derived from the half-ring difference maps. *Bottom*: residuals with respect to the *Planck* ΛCDM best-fit model. The error bars do not include the cosmic variance contribution (but the (brown) 1σ contour lines for the Likelihood best fit model do).
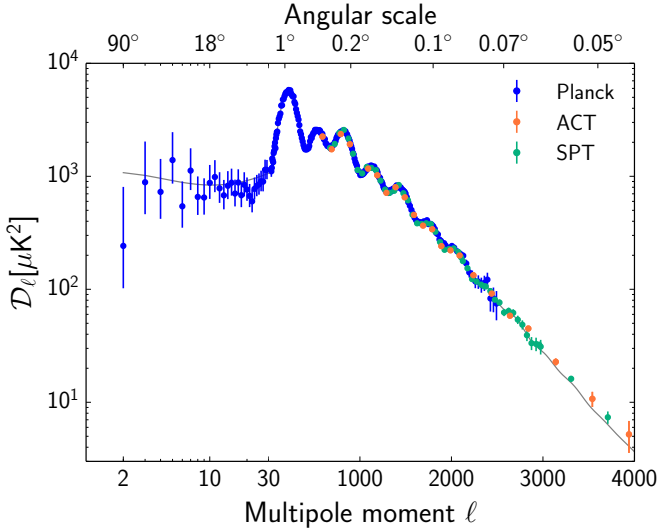
Fig. 49. CMB-only power spectra measured by *Planck* (blue), ACT (orange), and SPT (green). The best-fit PlanckTT+lowP ΛCDM model is shown by the grey solid line. ACT data at $\ell > 1000$ and SPT data at $\ell > 2000$ are marginalized CMB band-powers from multi-frequency spectra presented in Das et 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,
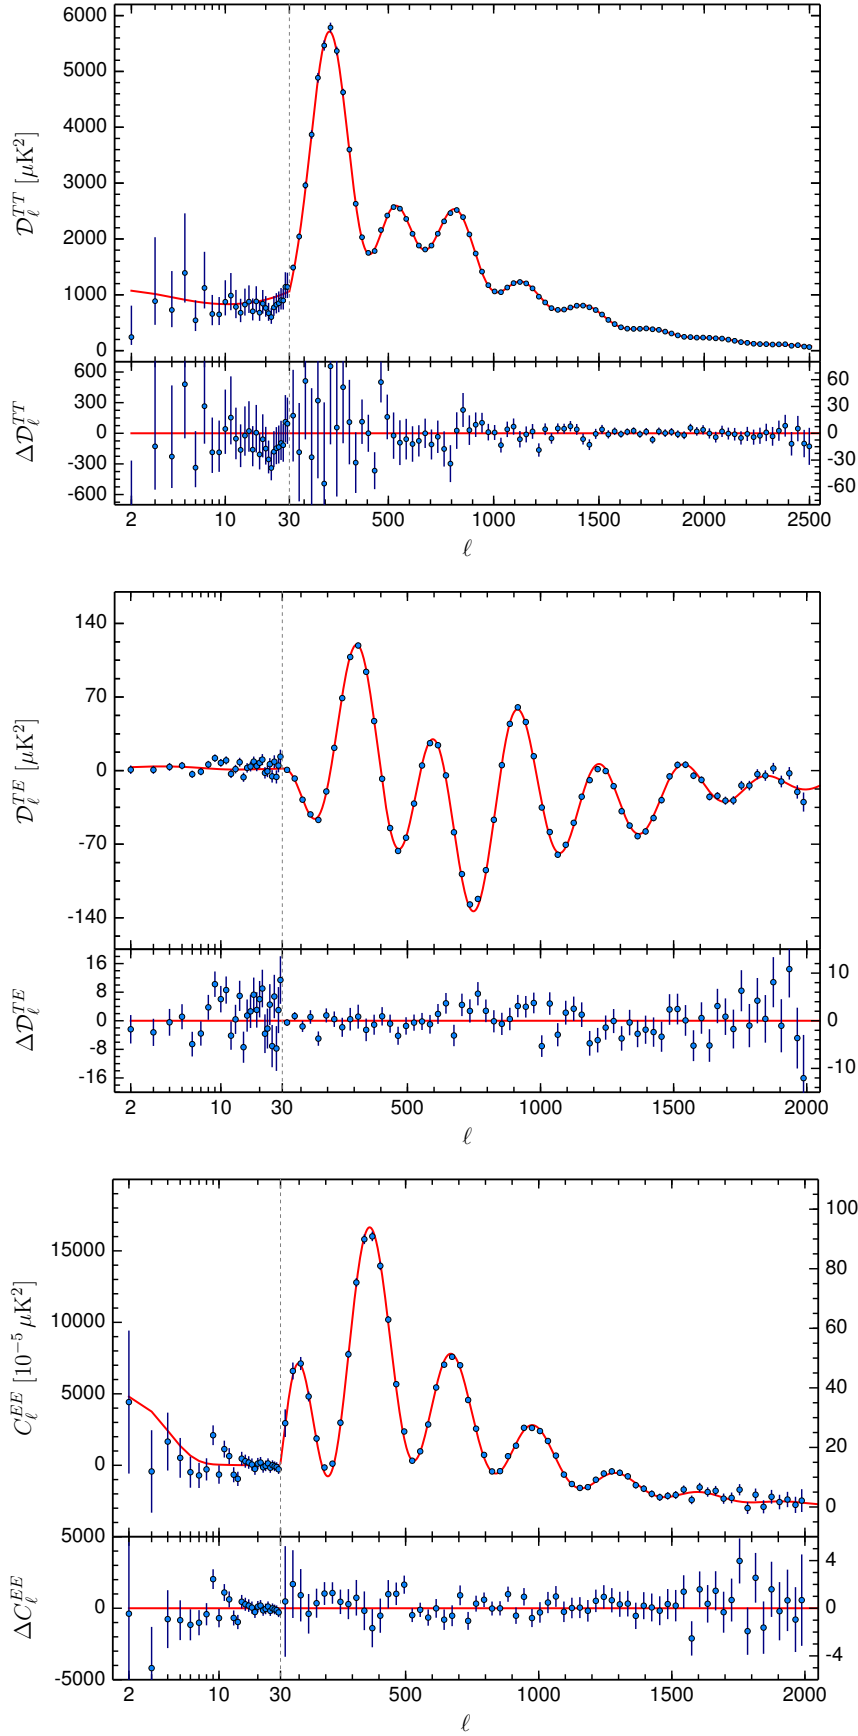
Fig. 50. *Planck* 2015 CMB spectra, compared with the base ΛCDM fit to PlanckTT+lowP data (red line). *The upper panels* show the spectra and the *lower panels* the residuals. In *all the panels*, the horizontal scale changes from logarithmic to linear at the "hybridization" scale, $\ell = 29$ (the division 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).

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 μK^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.

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_{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 $f_{sky} = \sum_i w_i^2 \Omega_i / (4\pi)$, 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_{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 $\sigma = FWHM/\sqrt{8}$ 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
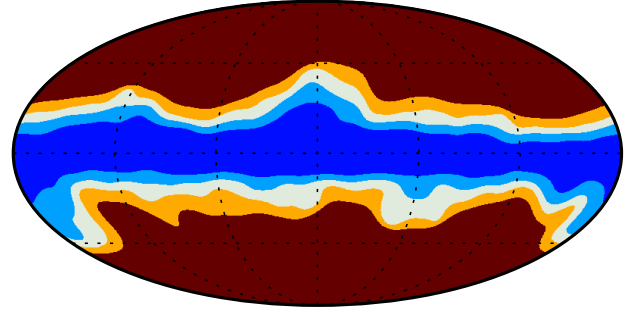
---

[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.

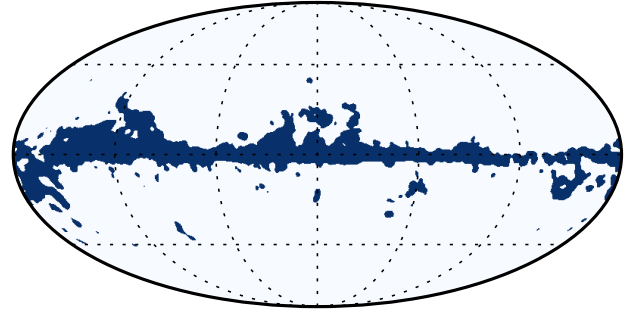**Fig. A.1.** Unapodized Galactic masks B50, B60, B70, and B80, from orange to dark blue.

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

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).

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

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 $(l, b) = (299°.85, −43°.6)$. [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:

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

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

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 $\mathsf{P}_\ell^{TT} = (\mathsf{V}_\ell^{TT})^{\mathsf{T}} \mathsf{A}_\ell^{TT} \mathsf{V}_\ell^{TT}$, 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

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

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

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

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

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

$$|\mathsf{M}| = |\mathsf{M}_0| |\mathsf{A}| |\mathsf{A}^{-1} + \mathsf{V} \mathsf{M}_0^{-1} \mathsf{V}^{\mathsf{T}}|. \qquad (B.6)$$

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

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

with the variable

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

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 $A_s e^{-2\tau}$) at $\ell = 200$. The resulting distribution is plotted in Fig. B.1. The best fit is at

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

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.
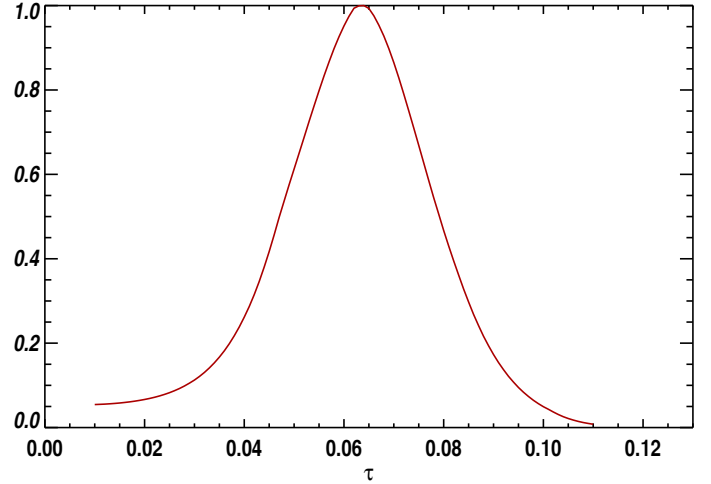

Fig. B.1. Distribution of the reionization optical depth τ using the `Lollipop` likelihood, based on the cross-correlation of the 100 and 143 GHz channels.

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*,

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

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:

$\mathrm{Var}(\hat{C}_\ell^{TT\,i,j}, \hat{C}_{\ell'}^{TT\,p,q})$

$$\begin{aligned}
&\approx \sqrt{C_\ell^{TT\,i,p} C_{\ell'}^{TT\,i,p} C_\ell^{TT\,j,q} C_{\ell'}^{TT\,j,q}} \; \Xi_{TT}^{00,00}\big[(i,p)^{TT}, (j,q)^{TT}\big]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{TT\,i,q} C_{\ell'}^{TT\,i,q} C_\ell^{TT\,j,p} C_{\ell'}^{TT\,j,p}} \; \Xi_{TT}^{00,00}\big[(i,q)^{TT}, (j,p)^{TT}\big]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{TT\,i,p} C_{\ell'}^{TT\,i,p}} \; \Xi_{TT}^{00,}\big[(i,p)^{TT}, (j,q)^{TT}\big]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{TT\,j,q} C_{\ell'}^{TT\,j,q}} \; \Xi_{TT}^{00,TT}\big[(j,q)^{TT}, (i,p)^{TT}\big]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{TT\,i,q} C_{\ell'}^{TT\,i,q}} \; \Xi_{TT}^{00,TT}\big[(i,q)^{TT}, (j,p)^{TT}\big]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{TT\,j,p} C_{\ell'}^{TT\,j,p}} \; \Xi_{TT}^{00,TT}\big[(j,p)^{TT}, (i,q)^{TT}\big]_{\ell\ell'} \\
&+ \Xi_{TT}^{TT,TT}\big[(i,p)^{TT}, (j,q)^{TT}\big]_{\ell\ell'} + \Xi_{TT}^{TT,TT}\big[(i,q)^{TT}, (j,p)^{TT}\big]_{\ell\ell'}.
\end{aligned} \tag{C.2}$$

TTTE block:

$\mathrm{Var}(\hat{C}_\ell^{TT\,i,j}, \hat{C}_{\ell'}^{TE\,p,q})$

$$\begin{aligned}
&\approx \frac{1}{2} \sqrt{C_\ell^{TT\,i,p} C_{\ell'}^{TT\,i,p}} \left(C_\ell^{TE\,j,q} + C_{\ell'}^{TE\,j,q} \right) \\
&\quad \times \Xi_{TT}^{00,00}\big[(i,p)^{TT}, (j,q)^{TP}\big]_{\ell\ell'} \\
&+ \frac{1}{2} \sqrt{C_\ell^{TT\,j,p} C_{\ell'}^{TT\,j,p}} \left(C_\ell^{TE\,i,q} + C_{\ell'}^{TE\,i,q} \right) \\
&\quad \times \Xi_{TT}^{00,00}\big[(i,q)^{TP}, (j,p)^{TT}\big]_{\ell\ell'} \\
&+ \frac{1}{2} \left(C_\ell^{TE\,j,q} + C_{\ell'}^{TE\,j,q} \right) \Xi_{TT}^{00,TT}\big[(j,q)^{TP}, (i,p)^{TT}\big]_{\ell\ell'} \\
&+ \frac{1}{2} \left(C_\ell^{TE\,i,q} + C_{\ell'}^{TE\,i,q} \right) \Xi_{TT}^{00,TT}\big[(i,q)^{TP}, (j,p)^{TT}\big]_{\ell\ell'}.
\end{aligned} \tag{C.3}$$

TETE block

$\mathrm{Var}(\hat{C}_\ell^{TE\,i,j}, \hat{C}_{\ell'}^{TE\,p,q})$

$$\begin{aligned}
&\approx \sqrt{C_\ell^{TT\,i,p} C_{\ell'}^{TT\,i,p} C_\ell^{EE\,j,q} C_{\ell'}^{EE\,j,q}} \; \Xi_{TE}^{00,00}\big[(i,p)^{TT}, (j,q)^{PP}\big]_{\ell\ell'} \\
&+ \frac{1}{2} \left(C_\ell^{TE\,i,q} C_{\ell'}^{TE\,j,p} + C_\ell^{TE\,j,p} C_{\ell'}^{TE\,i,q} \right) \Xi_{TT}^{00,00}\big[(i,q)^{TP}, (j,p)^{PT}\big]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{TT\,i,p} C_{\ell'}^{TT\,i,p}} \; \Xi_{TE}^{00,PP}\big[(i,p)^{TT}, (j,q)^{PP}\big]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{EE\,j,q} C_{\ell'}^{EE\,j,q}} \; \Xi_{TE}^{00,TT}\big[(j,q)^{PP}, (i,p)^{TT}\big]_{\ell\ell'} \\
&+ \Xi_{TE}^{TT,PP}\big[(i,p)^{TT}, (j,q)^{PP}\big]_{\ell\ell'}.
\end{aligned} \tag{C.4}$$

TTEE block:

$\mathrm{Var}(\hat{C}_\ell^{TT\,i,j}, \hat{C}_{\ell'}^{EE\,p,q})$

$$\begin{aligned}
&\approx \frac{1}{2} \left(C_\ell^{TE\,i,p} C_{\ell'}^{TE\,j,q} + C_\ell^{TE\,j,q} C_{\ell'}^{TE\,i,p} \right) \Xi_{TT}^{00,00}\big[(i,p)^{TP}, (j,q)^{TP}\big]_{\ell\ell'} \\
&+ \frac{1}{2} \left(C_\ell^{TE\,i,q} C_{\ell'}^{TE\,j,p} + C_\ell^{TE\,j,p} C_{\ell'}^{TE\,i,q} \right) \Xi_{TT}^{00,00}\big[(i,q)^{TP}, (j,p)^{TP}\big]_{\ell\ell'}.
\end{aligned} \tag{C.5}$$

TEEE block:

$\mathrm{Var}(\hat{C}_\ell^{TE\,i,j}, \hat{C}_{\ell'}^{EE\,p,q})$

$$\begin{aligned}
&\approx \frac{1}{2} \sqrt{C_\ell^{EE\,j,q} C_{\ell'}^{EE\,j,q}} \left(C_\ell^{TE\,i,p} + C_{\ell'}^{TE\,i,p} \right) \Xi_{EE}^{00,00}\big[(i,p)^{TP}, (j,q)^{PP}\big]_{\ell\ell'} \\
&+ \frac{1}{2} \sqrt{C_\ell^{EE\,j,p} C_{\ell'}^{EE\,j,p}} \left(C_\ell^{TE\,i,q} + C_{\ell'}^{TE\,i,q} \right) \Xi_{EE}^{00,00}\big[(i,q)^{TP}, (j,p)^{PP}\big]_{\ell\ell'} \\
&+ \frac{1}{2} \left(C_\ell^{TE\,i,p} + C_{\ell'}^{TE\,i,p} \right) \Xi_{EE}^{00,PP}\big[(i,p)^{TP}, (j,q)^{PP}\big]_{\ell\ell'} \\
&+ \frac{1}{2} \left(C_\ell^{TE\,i,q} + C_{\ell'}^{TE\,i,q} \right) \Xi_{EE}^{00,PP}\big[(i,q)^{TP}, (j,p)^{PP}\big]_{\ell\ell'}.
\end{aligned} \tag{C.6}$$

EEEE block:

$\mathrm{Var}(\hat{C}_\ell^{EE\,i,j}, \hat{C}_{\ell'}^{EE\,p,q})$

$$\begin{aligned}
&\approx \sqrt{C_\ell^{EE\,i,p} C_{\ell'}^{EE\,i,p} C_\ell^{EE\,j,q} C_{\ell'}^{EE\,j,q}} \; \Xi_{EE}^{00,00}\big[(i,p)^{PP}, (j,q)^{PP}\big]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{EE\,i,q} C_{\ell'}^{EE\,i,q} C_\ell^{EE\,j,p} C_{\ell'}^{EE\,j,p}} \; \Xi_{EE}^{00,00}\big[(i,q)^{PP}, (j,p)^{PP}\big]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{EE\,i,p} C_{\ell'}^{EE\,i,p}} \; \Xi_{EE}^{00,PP}\big[(i,p)^{PP}, (j,q)^{PP}\big]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{EE\,j,q} C_{\ell'}^{EE\,j,q}} \; \Xi_{EE}^{00,PP}\big[(j,q)^{PP}, (i,p)^{PP}\big]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{EE\,i,q} C_{\ell'}^{EE\,i,q}} \; \Xi_{EE}^{00,PP}\big[(i,q)^{PP}, (j,p)^{PP}\big]_{\ell\ell'} \\
&+ \sqrt{C_\ell^{EE\,j,p} C_{\ell'}^{EE\,j,p}} \; \Xi_{EE}^{00,PP}\big[(j,p)^{PP}, (i,q)^{PP}\big]_{\ell\ell'} \\
&+ \Xi_{EE}^{PP,PP}\big[(i,p)^{PP}, (j,q)^{PP}\big]_{\ell\ell'} + \Xi_{EE}^{PP,PP}\big[(i,q)^{PP}, (j,p)^{PP}\big]_{\ell\ell'}.
\end{aligned} \tag{C.7}$$

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,

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

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

and

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

where $X, Y \in \{00, TT, PP\}$, and $\alpha, \beta \in \{TT, TP, PT, PP\}$. They make use of window functions W,

$$W^{00,00}\big[(i,j)^\alpha, (p,q)^\beta\big]_\ell = \frac{1}{2\ell+1} \sum_m w_{\ell m}^{00}(i,j)^\alpha w_{\ell m}^{*\,00}(p,q)^\beta, \tag{C.11}$$

$$W^{00,TT}\big[(i,j)^\alpha, (p,q)^{TT}\big]_\ell = \frac{1}{2\ell+1} \sum_m w_{\ell m}^{00}(i,j)^\alpha w_{\ell m}^{*\,II}(p,q)^{TT}, \tag{C.12}$$

$$W^{00,PP}\big[(i,j)^{\alpha},(p,q)^{PP}\big]_{\ell} = \frac{1}{2\ell+1}\sum_{m}$$

$$\frac{1}{2}\left(w_{\ell m}^{00}(i,j)^{\alpha}w_{\ell m}^{*\,QQ}(p,q)^{PP} + w_{\ell m}^{00}(i,j)^{\alpha}w_{\ell m}^{*\,UU}(p,q)^{PP}\right), \quad \text{(C.13)}$$

$$W^{TT,TT}\big[(i,j)^{TT},(p,q)^{TT}\big]_{\ell} = \frac{1}{2\ell+1}\sum_{m}w_{\ell m}^{II}(i,j)^{TT}w_{\ell m}^{*\,II}(p,q)^{TT}, \tag{C.14}$$

$$W^{TT,PP}\big[(i,j)^{TT},(p,q)^{PP}\big]_{\ell} = \frac{1}{2\ell+1}\sum_{m}$$

$$\frac{1}{2}\left(w_{\ell m}^{II}(i,j)^{TT}w_{\ell m}^{*\,QQ}(p,q)^{PP} + w_{\ell m}^{II}(i,j)^{TT}w_{\ell m}^{*\,UU}(p,q)^{PP}\right), \quad \text{(C.15)}$$

and

$$W^{PP,PP}\big[(i,j)^{PP},(p,q)^{PP}\big]_{\ell} = \frac{1}{2\ell+1}\sum_{m}$$

$$\frac{1}{4}\left(w_{\ell m}^{QQ}(i,j)^{PP}w_{\ell m}^{*\,QQ}(p,q)^{PP} + w_{\ell m}^{UU}(i,j)^{PP}w_{\ell m}^{*\,UU}(p,q)^{PP}\right.$$
$$\left. + w_{\ell m}^{QQ}(i,j)^{PP}w_{\ell m}^{*\,UU}(p,q)^{PP} + w_{\ell m}^{UU}(i,j)^{PP}w_{\ell m}^{*\,QQ}(p,q)^{PP}\right).$$
$$\text{(C.16)}$$

In the above expressions, we defined the spherical harmonic coefficients of the effective weight maps w^{00},

$$w_{\ell m}^{00}(i,j)^{TT} = \sum_{p=1}^{N_{\text{pix}}} m_p^{i,T} m_p^{j,T} Y_{\ell m}^{*}(\hat{\boldsymbol{n}}_p)\Omega_p, \tag{C.17}$$

$$w_{\ell m}^{00}(i,j)^{PP} = \sum_{p=1}^{N_{\text{pix}}} m_p^{i,P} m_p^{j,P} Y_{\ell m}^{*}(\hat{\boldsymbol{n}}_p)\Omega_p, \tag{C.18}$$

$$w_{\ell m}^{00}(i,j)^{TP} = \sum_{p=1}^{N_{\text{pix}}} m_p^{i,T} m_p^{j,P} Y_{\ell m}^{*}(\hat{\boldsymbol{n}}_p)\Omega_p, \tag{C.19}$$

and

$$w_{\ell m}^{00}(i,j)^{PT} = \sum_{p=1}^{N_{\text{pix}}} m_p^{i,P} m_p^{j,T} Y_{\ell m}^{*}(\hat{\boldsymbol{n}}_p)\Omega_p, \tag{C.20}$$

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

$$w_{\ell m}^{II}(i,j)^{TT} = \delta_{i,j}\sum_{p=1}^{N_{\text{pix}}}\left(\sigma_p^{II}\right)^2 m_p^{i,T} m_p^{j,T} Y_{\ell m}^{*}(\hat{\boldsymbol{n}}_p)\Omega_p^2, \tag{C.21}$$

$$w_{\ell m}^{QQ}(i,j)^{PP} = \delta_{i,j}\sum_{p=1}^{N_{\text{pix}}}\left(\sigma_p^{QQ}\right)^2 m_p^{i,P} m_p^{j,P} Y_{\ell m}^{*}(\hat{\boldsymbol{n}}_p)\Omega_p^2, \tag{C.22}$$

and

$$w_{\ell m}^{UU}(i,j)^{PP} = \delta_{i,j}\sum_{p=1}^{N_{\text{pix}}}\left(\sigma_p^{UU}\right)^2 m_p^{i,P} m_p^{j,P} Y_{\ell m}^{*}(\hat{\boldsymbol{n}}_p)\Omega_p^2, \tag{C.23}$$

where the σ_p^2s 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 $\Xi_{\ell_1\ell_2}^{X,Y}$ for each factor of $X,Y \in \{TT, PP\}$ by an additional rescaling coefficient,

$$r_{\ell_1\ell_2} = \sqrt{\frac{N_{\ell_1}^{\text{data}} N_{\ell_2}^{\text{data}}}{N_{\ell_1}^{\text{white}} N_{\ell_2}^{\text{white}}}}. \tag{C.24}$$

Here, $N_{\ell}^{\text{data}}/N_{\ell}^{\text{white}}$ is the ratio of the observed noise power spectrum to the white-noise power spectrum predicted by the pixel noise variance values σ_p^2.

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),

$$M^{TT}(i,j)_{\ell_1\ell_2} = (2\ell_2+1)\sum_{\ell_3}\frac{2\ell_3+1}{4\pi}\begin{pmatrix}\ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0\end{pmatrix}^2 V^{TT}(i,j)_{\ell_3},$$
$$\text{(C.25)}$$

$$M^{EE}(i,j)_{\ell_1\ell_2} = (2\ell_2+1)\sum_{\ell_3}\frac{2\ell_3+1}{16\pi}\left(1+(-1)^{\ell_1+\ell_2+\ell_3}\right)^2$$
$$\times\begin{pmatrix}\ell_1 & \ell_2 & \ell_3 \\ -2 & 2 & 0\end{pmatrix}^2 V^{PP}(i,j)_{\ell_3}, \tag{C.26}$$

$$M^{TE}(i,j)_{\ell_1\ell_2} = (2\ell_2+1)\sum_{\ell_3}\frac{2\ell_3+1}{8\pi}\left(1+(-1)^{\ell_1+\ell_2+\ell_3}\right)$$
$$\times\begin{pmatrix}\ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0\end{pmatrix}\begin{pmatrix}\ell_1 & \ell_2 & \ell_3 \\ -2 & 2 & 0\end{pmatrix} V^{TP}(i,j)_{\ell_3}, \tag{C.27}$$

$$M^{ET}(i,j)_{\ell_1\ell_2} = (2\ell_2+1)\sum_{\ell_3}\frac{2\ell_3+1}{8\pi}\left(1+(-1)^{\ell_1+\ell_2+\ell_3}\right)$$
$$\times\begin{pmatrix}\ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0\end{pmatrix}\begin{pmatrix}\ell_1 & \ell_2 & \ell_3 \\ -2 & 2 & 0\end{pmatrix} V^{PT}(i,j)_{\ell_3}, \tag{C.28}$$

where

$$V^{TT}(i,j)_{\ell} = \frac{1}{2\ell+1}\sum_{m}m_{\ell m}^{i,T}m_{\ell m}^{*\,j,T}, \tag{C.29}$$

$$V^{PP}(i,j)_{\ell} = \frac{1}{2\ell+1}\sum_{m}m_{\ell m}^{i,P}m_{\ell m}^{*\,j,P}, \tag{C.30}$$

$$V^{TP}(i,j)_{\ell} = \frac{1}{2\ell+1}\sum_{m}m_{\ell m}^{i,T}m_{\ell m}^{*\,j,P}, \tag{C.31}$$

and

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

we obtain the final result for the deconvolved covariance matrix,

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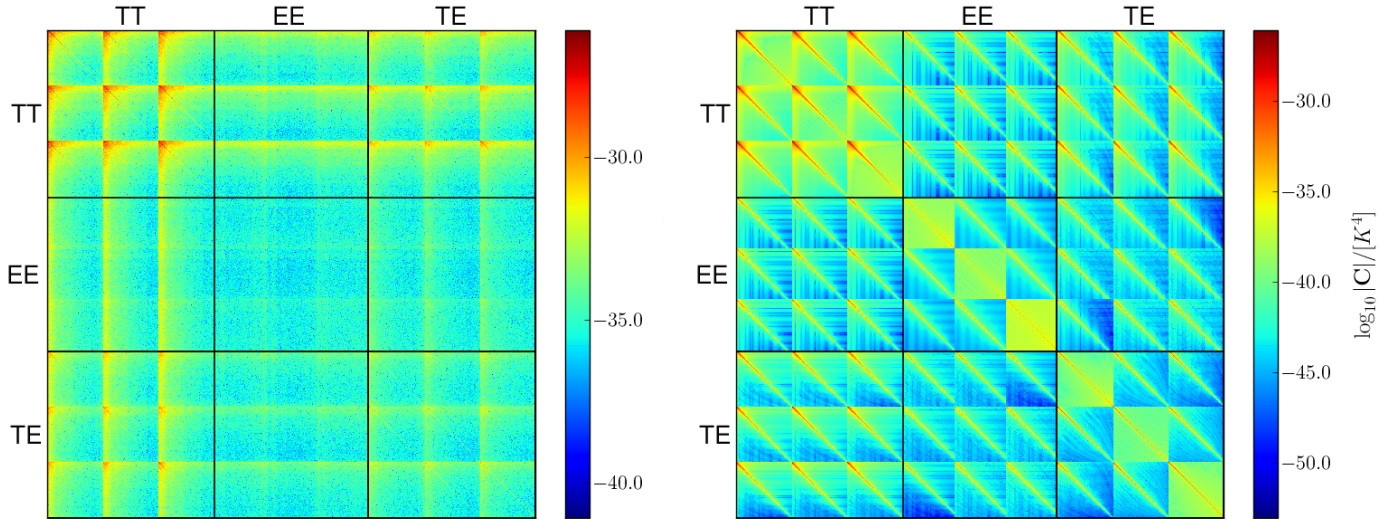$$\text{(C.33)}$$

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.

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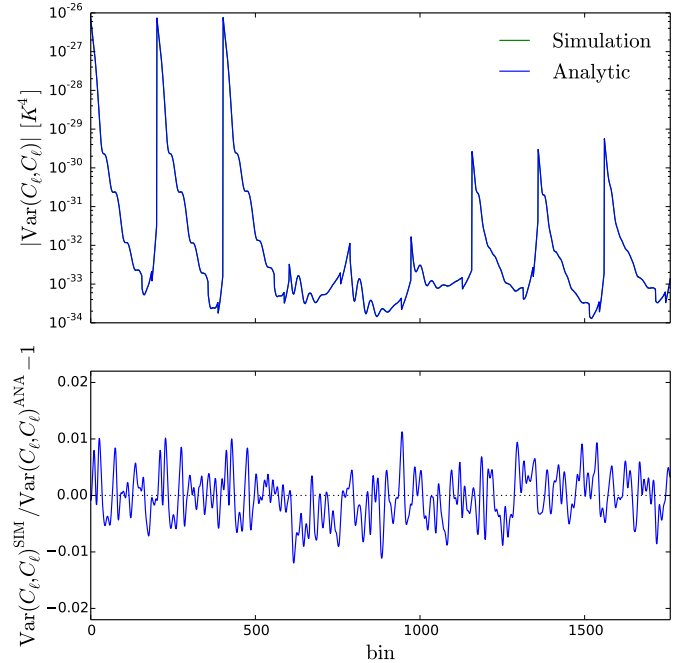
Fig. C.2. *Top*: diagonal elements of the empirical (green line) and analytic (blue line) covariance matrices; the two lines are indistinguishable. *Bottom*: ratio of the two estimates: the ratios differ from unity by <1% over the full multipole range for all frequency combinations and polarization blocks.

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.

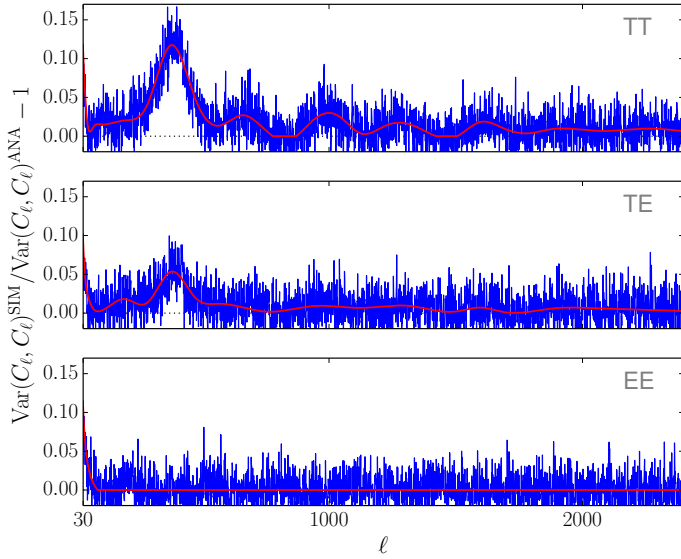**Fig. C.3.** Excess variance induced by the temperature point-source mask. The graphs compare the diagonal elements of the empirical and analytical power spectrum covariance matrices (blue lines) for *TT* (*upper panel*), *TE* (*middle panel*), and *EE* (*lower panel*), and show deviations at the 10% 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
$\Omega_{\rm b} h^2$	-1.09
$\Omega_{\rm c} h^2$	0.62
θ	-0.25
τ	-0.88
$\ln\left(10^{10} A_{\rm s}\right)$	-0.76
$n_{\rm s}$	0.25
$A_{\rm CIB}^{217}$	-0.75
${\rm gal}_{545}^{100}$	-0.03
${\rm gal}_{545}^{143}$	-0.05
${\rm gal}_{545}^{143-217}$	-0.28
${\rm gal}_{545}^{217}$	1.38
${\rm gal}_{EE}^{100}$	0.69
${\rm gal}_{EE}^{100-143}$	-0.80
${\rm gal}_{EE}^{100-217}$	0.02
${\rm gal}_{EE}^{143}$	-0.07
${\rm gal}_{EE}^{143-217}$	-1.21
${\rm gal}_{EE}^{217}$	1.08
${\rm gal}_{TE}^{100}$	-0.11
${\rm gal}_{TE}^{100-143}$	-0.39
${\rm gal}_{TE}^{100-217}$	0.32
${\rm gal}_{TE}^{143}$	0.55
${\rm gal}_{TE}^{143-217}$	-0.47
${\rm gal}_{TE}^{217}$	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 $n_{\rm s}$, θ, and τ. It also reduces the small bias in $n_{\rm s}$ 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
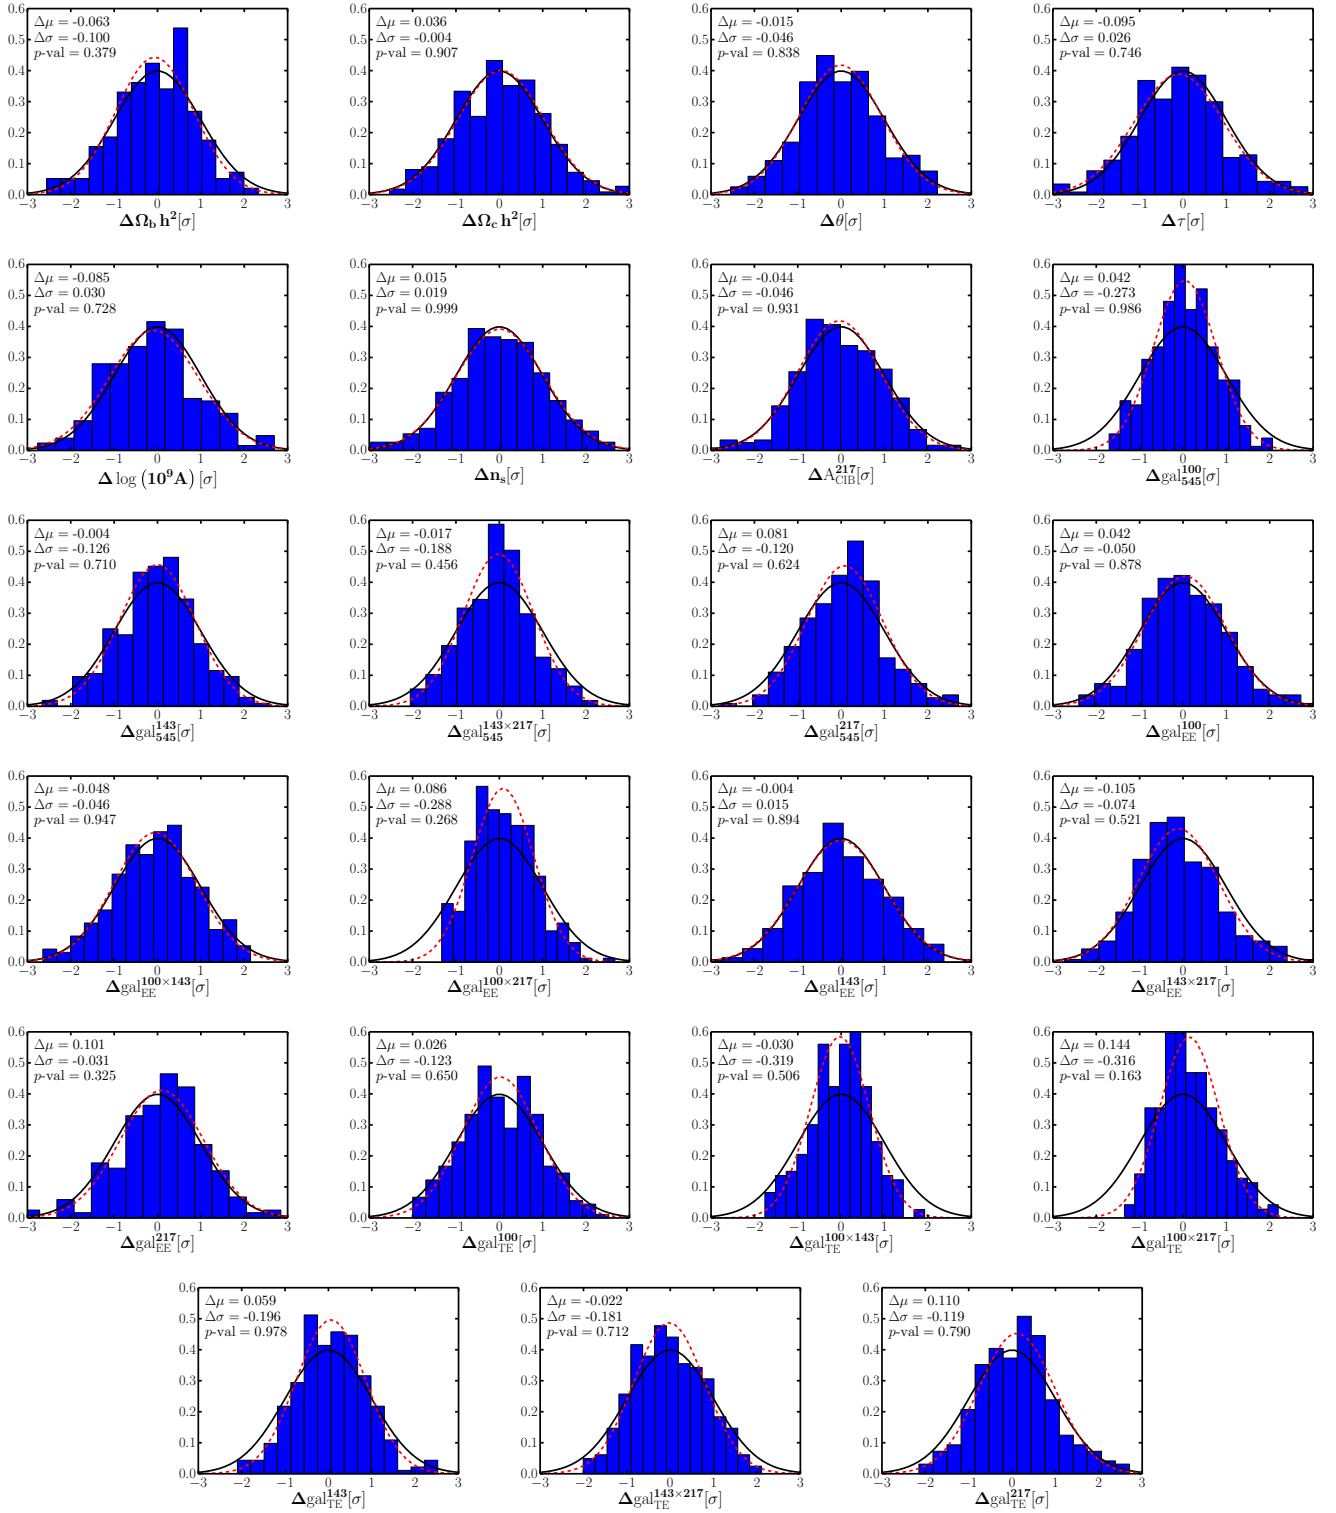
Fig. C.4. Plik parameter results from 300 simulations for the six baseline cosmological parameters, as well as the FFP8 CIB and Galactic dust amplitudes, as in Fig. 27, but for the joint PlikTT, EE, TE likelihood.

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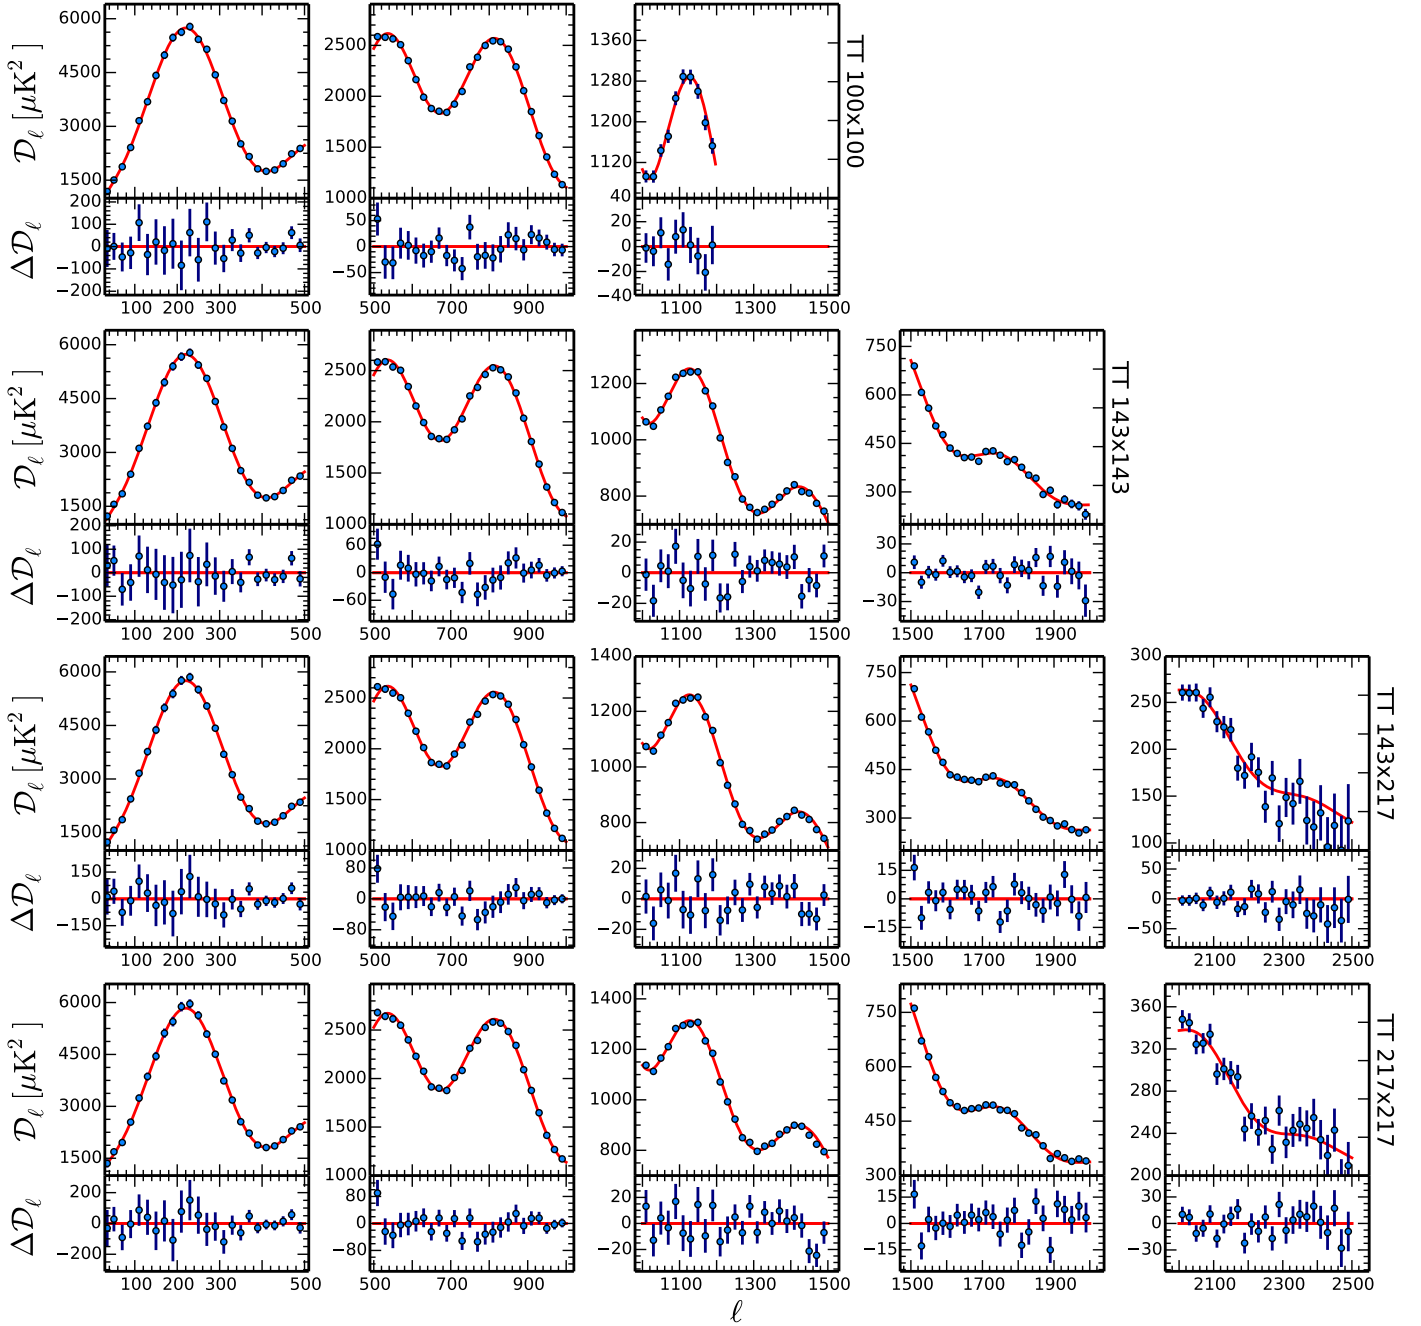
Fig. C.5. Per-frequency zoomed-in *TT* power spectra, in $\Delta\ell = 20$ bins. The red line shows the `PlikTT+tauprior` ΛCDM best-fit model. The *lower plots* show the residuals. We only show the ℓ ranges used in the baseline `Plik` likelihood.

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:

$$A = \left(\begin{array}{cc} \mathbb{1}^{XY} & -\mathbb{1}^{X'Y'} \end{array} \right), \tag{C.35}$$

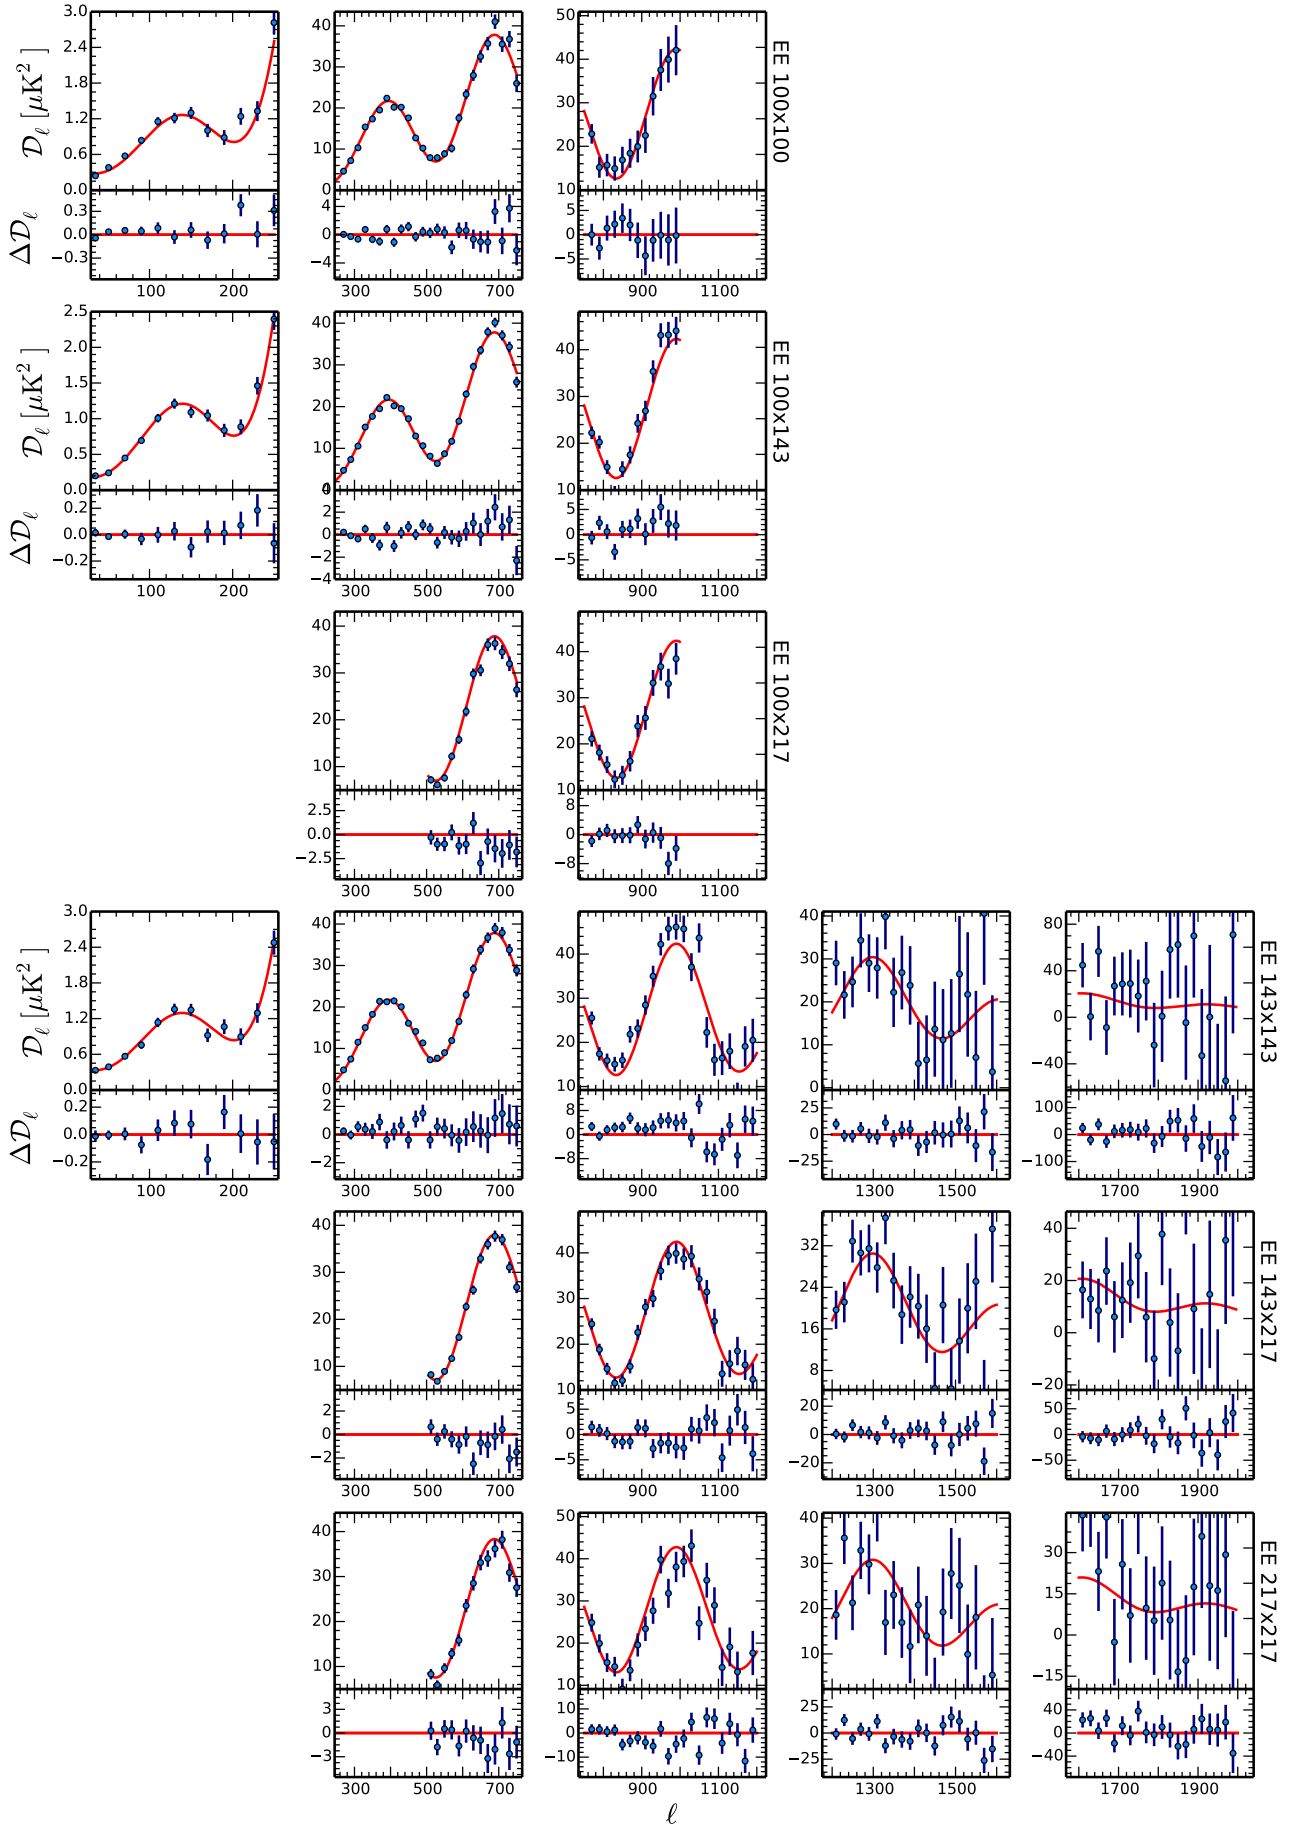where $\mathbb{1}^{XY}$ is the $n \times n$ identity matrix.

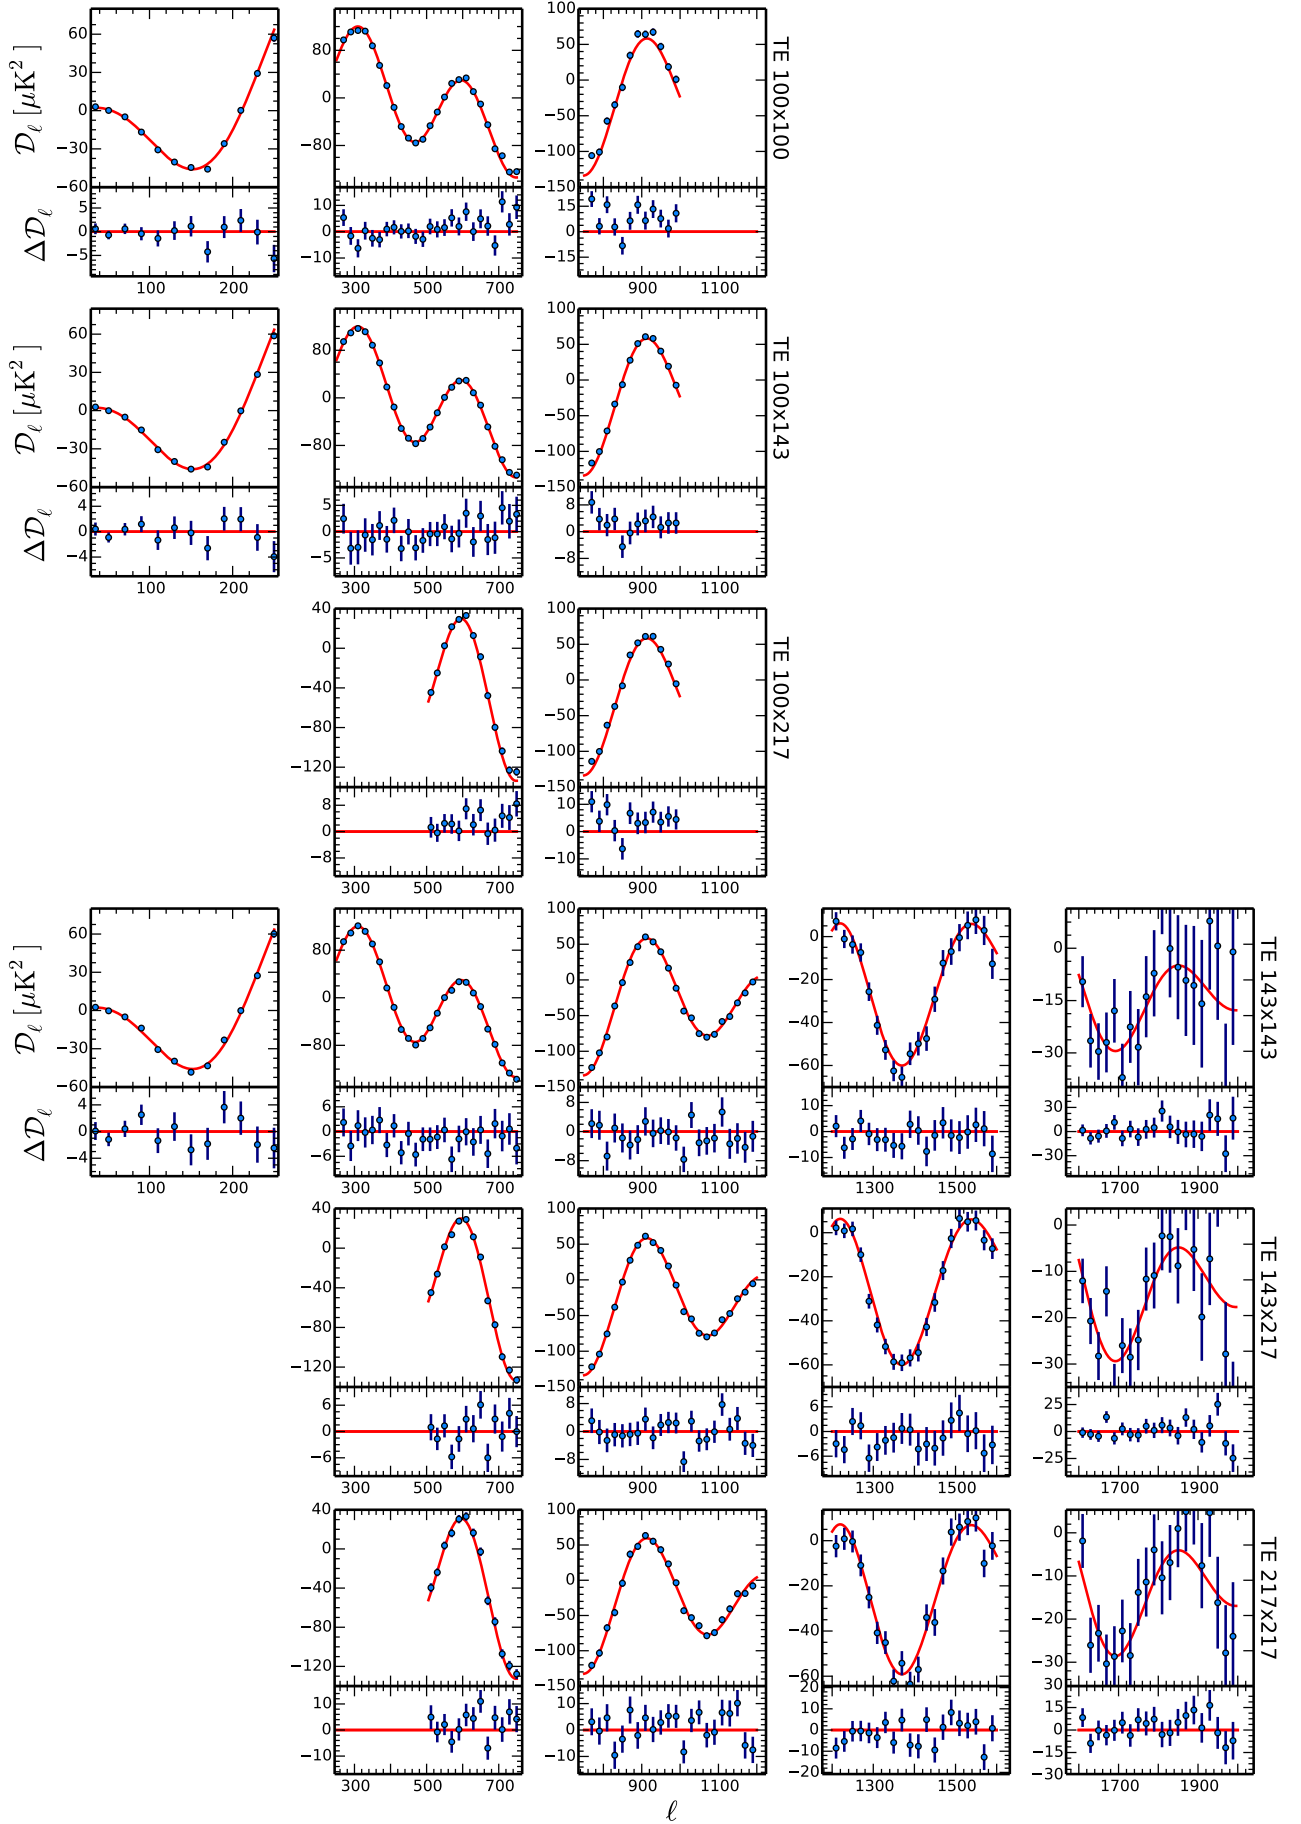**Fig. C.6.** Same as Fig. C.5, but for *EE*.

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 $100 \times 100\,\text{GHz}$ 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 $\ell_{\min}$. 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 $\ell_{\min} = 50, 100$ instead of the baseline $\ell_{\min} = 30$, 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 $\ell_{\max}$. We find that the overall amplitude of the foregrounds decreases when increasing the maximum multipole $\ell_{\max}$ included in the analysis[21]. This is related to the shift in cosmological parameters observed at different $\ell_{\max}$, which is described in Sect. 4.1.6. In Fig. C.8 the results for extragalactic foregrounds at $\ell_{\max} \lesssim 1200$ 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_{eff} case, the level of foregrounds is very similar to that in the base-ΛCDM case, while in the ΛCDM+A_{L} model it is few μK^2 lower at all frequencies.

CamSpec. The foreground contamination levels determined by the `CamSpec` and `Plik` codes differ by a few μK^2. This appears in Fig. C.8 as differences at the 1σ level in the sub-dominant (and ill-determined) foreground components (A^{kSZ}, A^{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_{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 $\ell_{\max}$

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 $\ell_{\max}$ 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 $\ln(10^{10}A_{\text{s}})$, $\Omega_{\text{c}}h^2$, and τ, i.e., they all increase with increasing $\ell_{\max}$. 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 $\ell_{\max} \approx 1200-1300$ already described in Sect. 4.1.6), fixing the foregrounds forces other parameters such as n_{s} and $\Omega_{\text{b}}h^2$ to shift as well. It is interesting to note that all the parameters tend to converge to the baseline solution between $\ell_{\max} = 1404$ 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_{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 $\ell_{\max}$), pushing $\Omega_{\text{c}}h^2$ to even lower values, following the $n_{\text{s}}-\Omega_{\text{c}}h^2$ 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 $\ell_{\max}^{\text{freq}} = \min(\ell_{\max}, \ell_{\max}^{\text{freq, baseline}})$, with $\ell_{\max}^{\text{freq, baseline}}$ the baseline $\ell_{\max}$ at each frequency as reported in Table 16 (e.g., in the $\ell_{\max} = 1404$ case, we still use the 100×100 power spectrum through $\ell = 1197$).
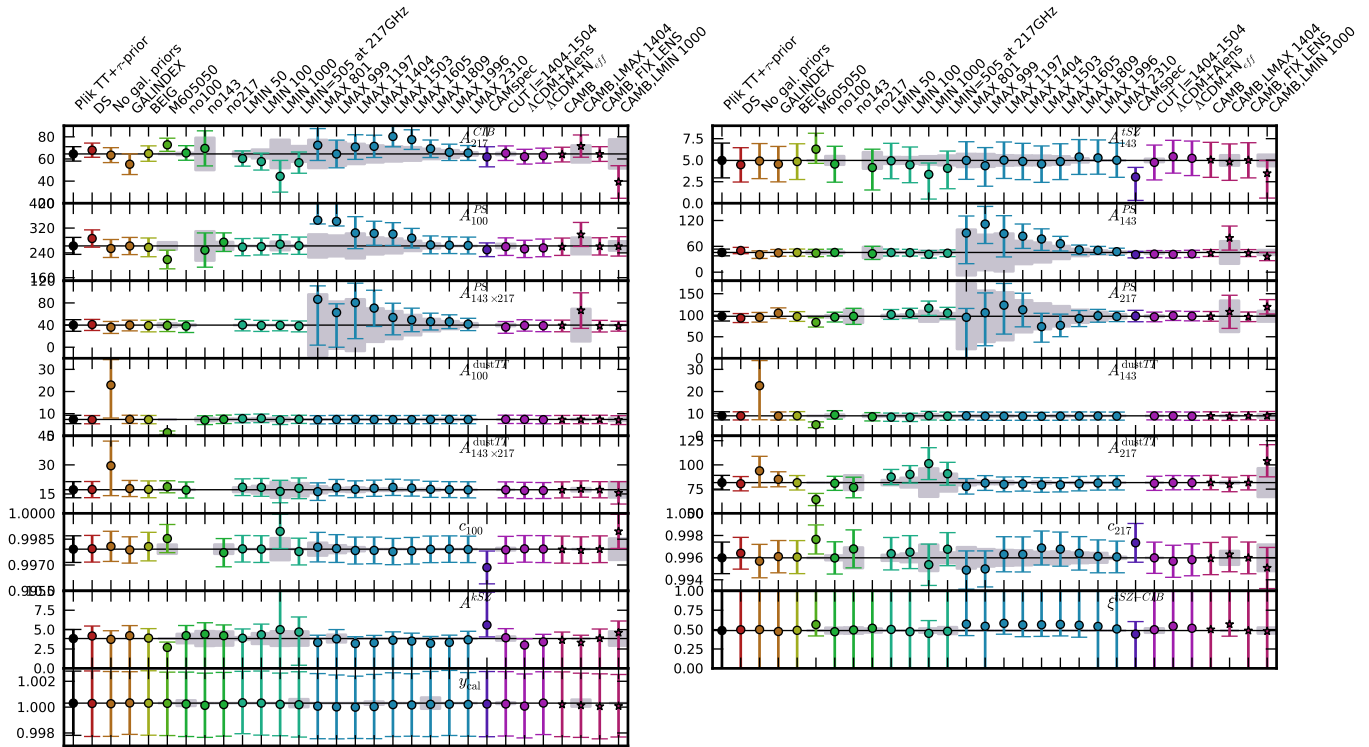
Fig. C.8. Marginal mean and 68% CL error bars on *TT* foreground parameters estimated when adopting different data choices for the `Plik` likelihood, in comparison with results from alternate approaches or models. We assume a ΛCDM model and always combine the `Plik` likelihood with a prior on $\tau = 0.07 \pm 0.02$ (we do not use low-ℓ temperature or polarization data here). "PlikTT+tauprior" indicates the baseline (HM, $\ell_{min} = 30$, $\ell_{max} = 2508$), 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 $\Omega_b h^2$ and θ higher than the baseline case by about 1σ, while n_s is lower by 1σ.

Larger Galactic mask. We examined the impact of using a larger Galactic mask (case "M605050") with f_{sky} = 0.50, 0.41, and 0.41 at 100, 143, and 217 GHz, respectively (corresponding to f_{sky}^{noap} = 0.60, 0.50, and 0.50 before apodization), instead of the baseline values f_{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 $-0.3\,\mu K^2$ 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 of σ, 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.
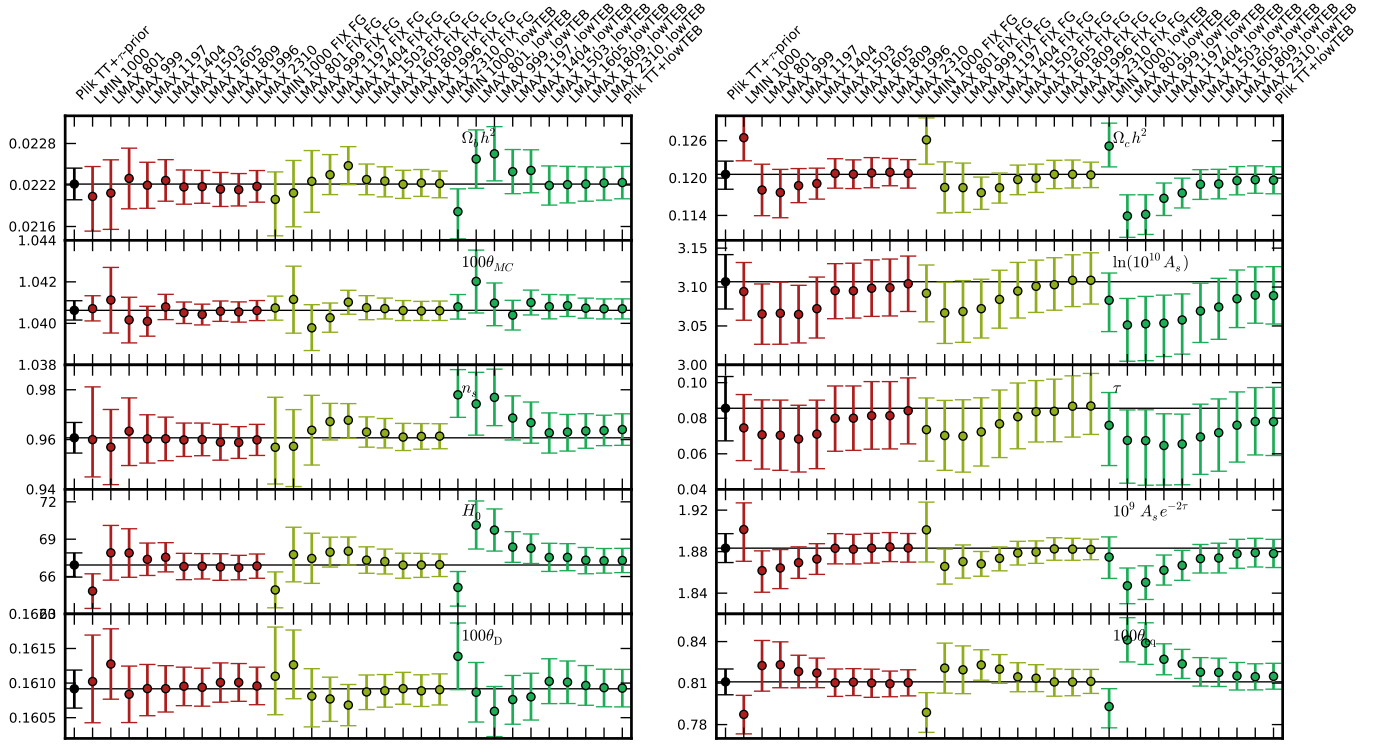
Fig. C.9. Marginal mean and 68% CL error bars on cosmological parameters estimated when adopting different data choices for the `Plik` likelihood. We assume a ΛCDM model and calculate parameters using different maximum multipole $\ell_{\max}$. The red points show the results for `PlikTT+tauprior`, with the points specifically labelled "PlikTT+tauprior" in black showing the baseline `PlikTT` likelihood at $\ell_{\max} = 2508$, 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 $1\,\mu\mathrm{K}^2$ ($\mathcal{D}_\ell$) in TE and $1 \times 10^{-5}\,\mu\mathrm{K}^2$ (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 $\ell_{\min}$. We find good stability in the results when changing the minimum multipole $\ell_{\min}$ considered in the analysis ("LMIN" case). The baseline likelihood has $\ell_{\min} = 30$, and we test the cases of $\ell_{\min} = 50$ and 100.

Changing $\ell_{\max}$. We observe small shifts when including maximum multipoles between $\ell_{\max} \sim 1000$ and 2000 ("LMAX" cases). This is not surprising, since even though the baseline has $\ell_{\max} = 2000$, most of the constraining power of our polarization spectra comes from $\ell < 1000$. When using $\ell_{\max} = 801$, 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.,
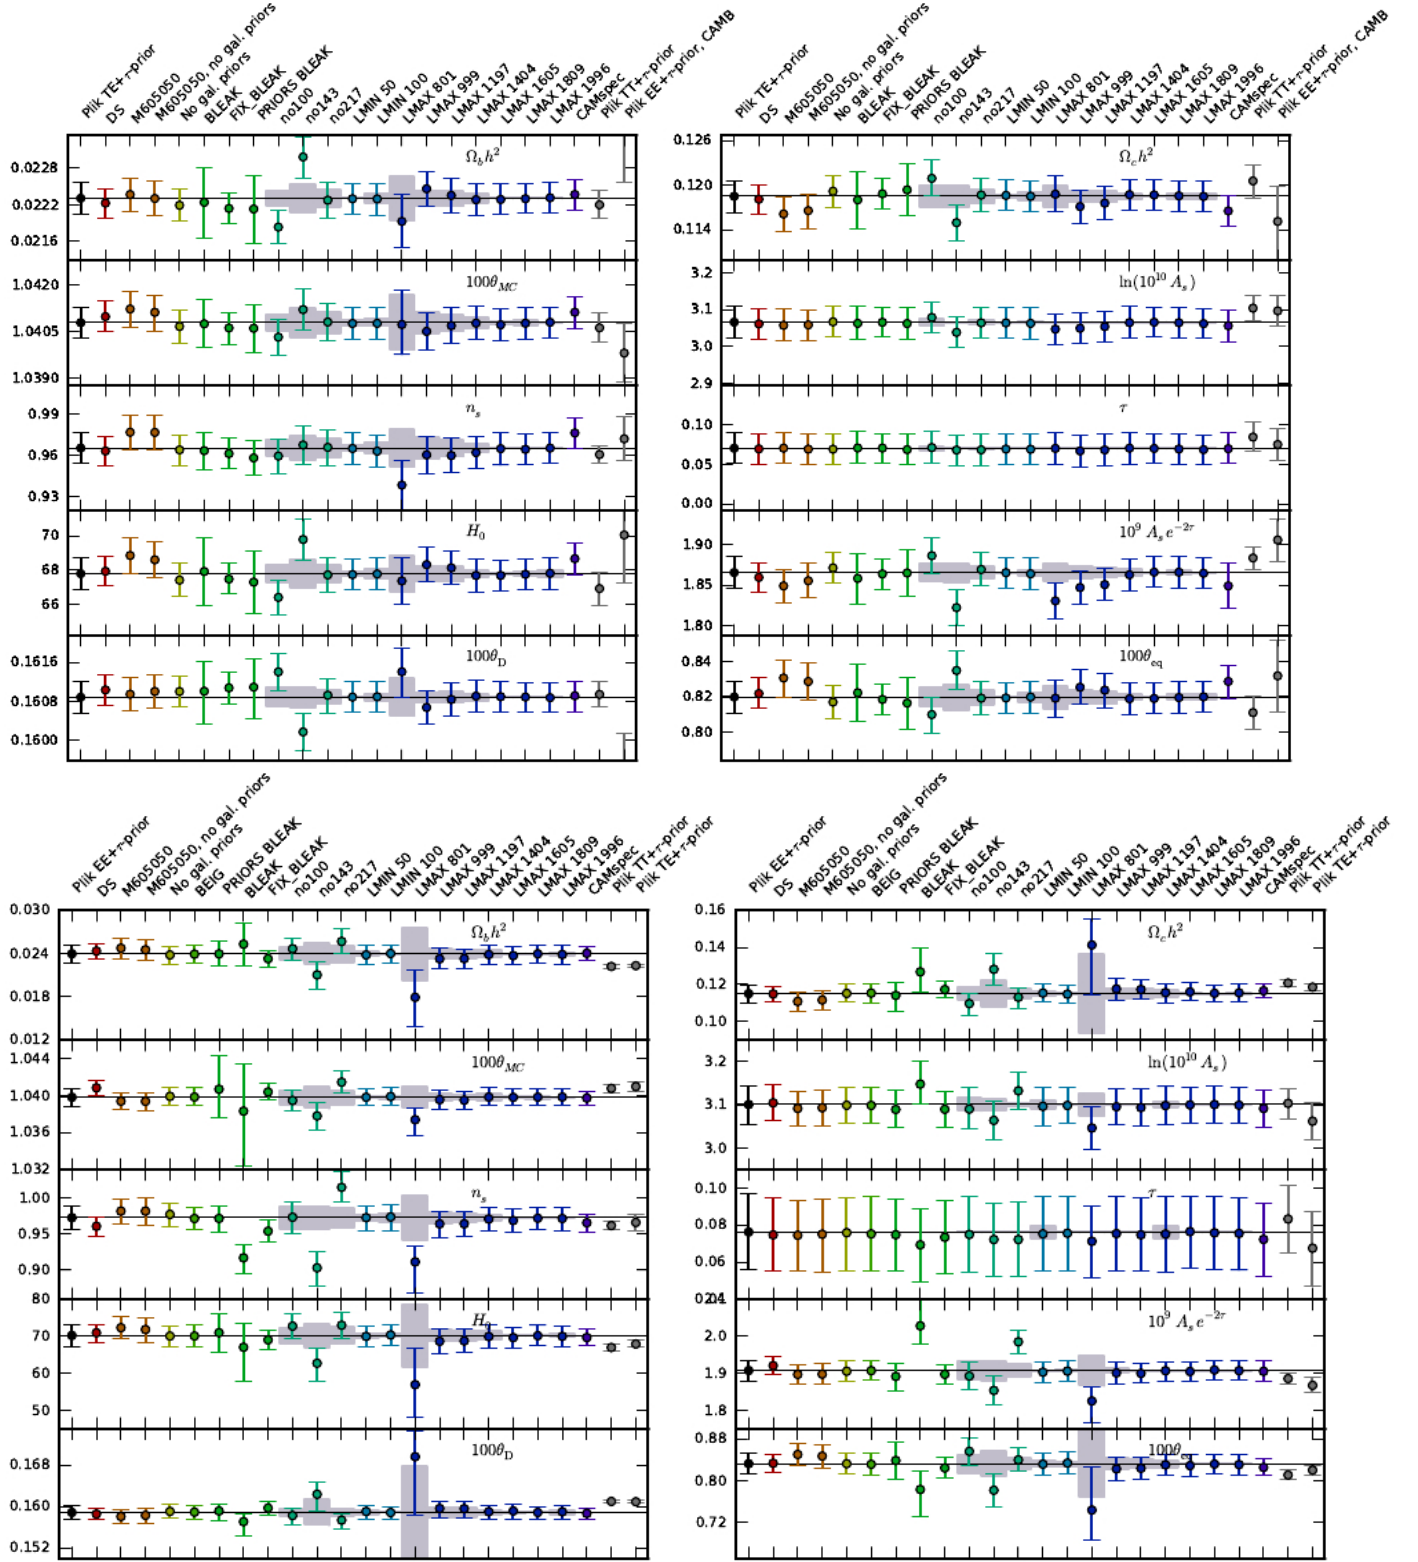
Fig. C.10. Marginal mean and 68% CL error bars on cosmological parameters estimated adopting different data choices for the `Plik` likelihood, in comparison with results from alternate approaches or model. *Top*: *TE* tests; we assume a ΛCDM model and use the `PlikTE+tauprior` likelihood in most of the 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, $\ell_{min} = 30$, $\ell_{max} = 1996$), 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_ℓ^{CMB} that minimizes:

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

where $\hat{\boldsymbol{C}}$ is the foreground-cleaned frequency data vector, $\boldsymbol{C}^{CMB}$ 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

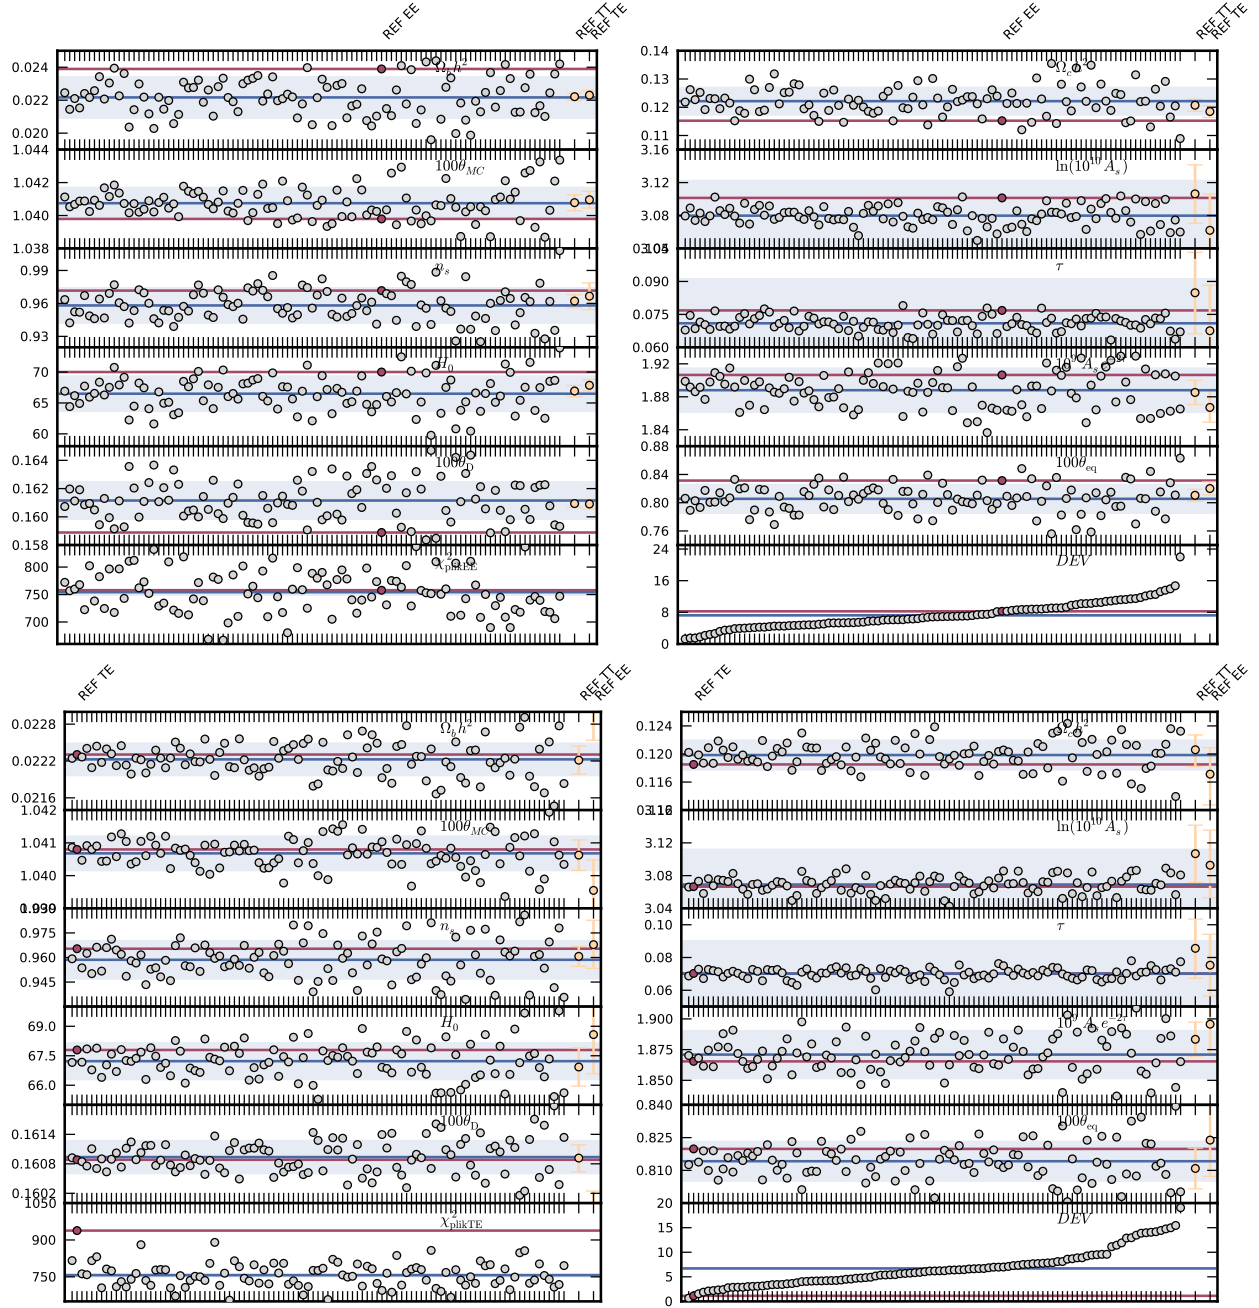**Fig. C.11.** Marginal mean and 68% CL error bars on cosmological parameters estimated from 100 *EE* (*left*) or *TE* (*right*) power-spectra simulations conditioned on the *TT* power spectrum, assuming as a fiducial cosmology the best-fit of the ΛCDM PlikTT+tauprior results (grey circles). The blue 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:

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

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

which we can rewrite

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

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

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

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

where we used the fact that

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

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

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

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

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

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 $\Omega_b h^2$ 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

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

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
$\Omega_b h^2$	0.00	0.00	0.63
$\Omega_c h^2$	0.05	−0.01	−0.40
θ	0.32	0.31	0.26
τ	−0.04	−0.14	0.08
$\ln(10^{10}A_s)$	−0.01	−0.11	0.21
n_s	0.23	0.10	0.24
H_0	0.01	0.07	0.52
$A_s \exp(-2\tau)$	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

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

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., $C_b^{\mathrm{th},ij} = \sum_\ell \mathcal{B}_{b\ell} C_\ell^{\mathrm{th},ij}$. We then write the model for the C_b parameters in vector form as

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

where C_b^{th} and C_b^{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_b^{CMB}, marginalized over the secondary parameters, θ. The posterior distribution for C_b^{CMB}, given the observed multi-frequency spectra C_b, can be written as

$$p(C_b^{\mathrm{CMB}}|C_b) = \int p(C_b^{\mathrm{CMB}},\theta|C_b)p(\theta)\mathrm{d}\theta. \qquad (C.48)$$

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

We write the multi-frequency *Planck* likelihood as

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

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

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

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

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

and covariance

$$\mathsf{Q} = \mathsf{A}^{\mathsf{T}}\Sigma^{-1}\mathsf{A}. \qquad (\mathrm{C.52})$$

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:

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

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_{b}s, and $C_{\mathrm{b}}^{\mathrm{th}}$ 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.
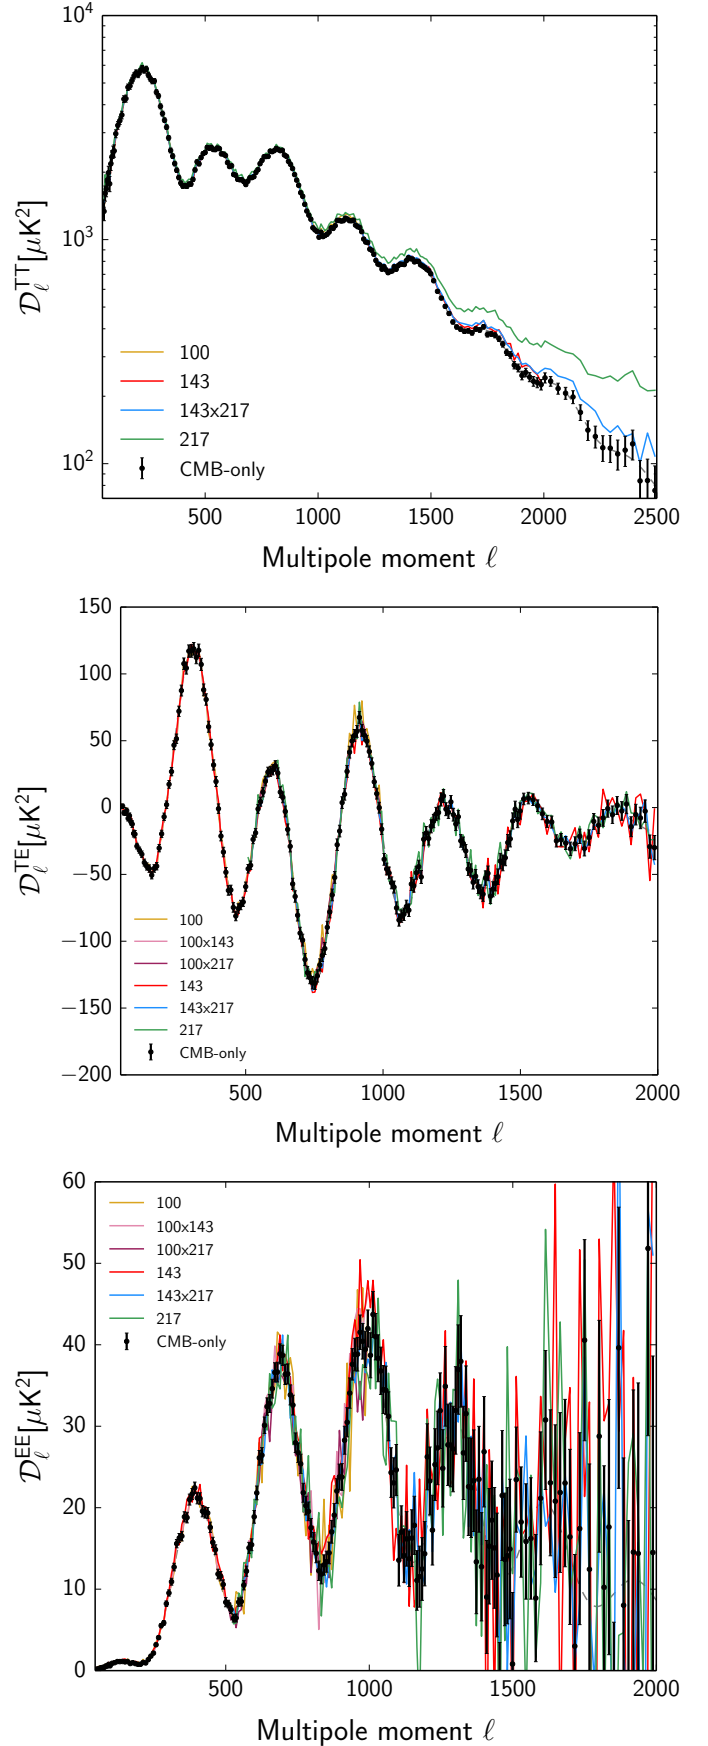

Fig. C.12. *Planck* multi-frequency power spectra (solid coloured lines) and extracted CMB-only spectra (black points).
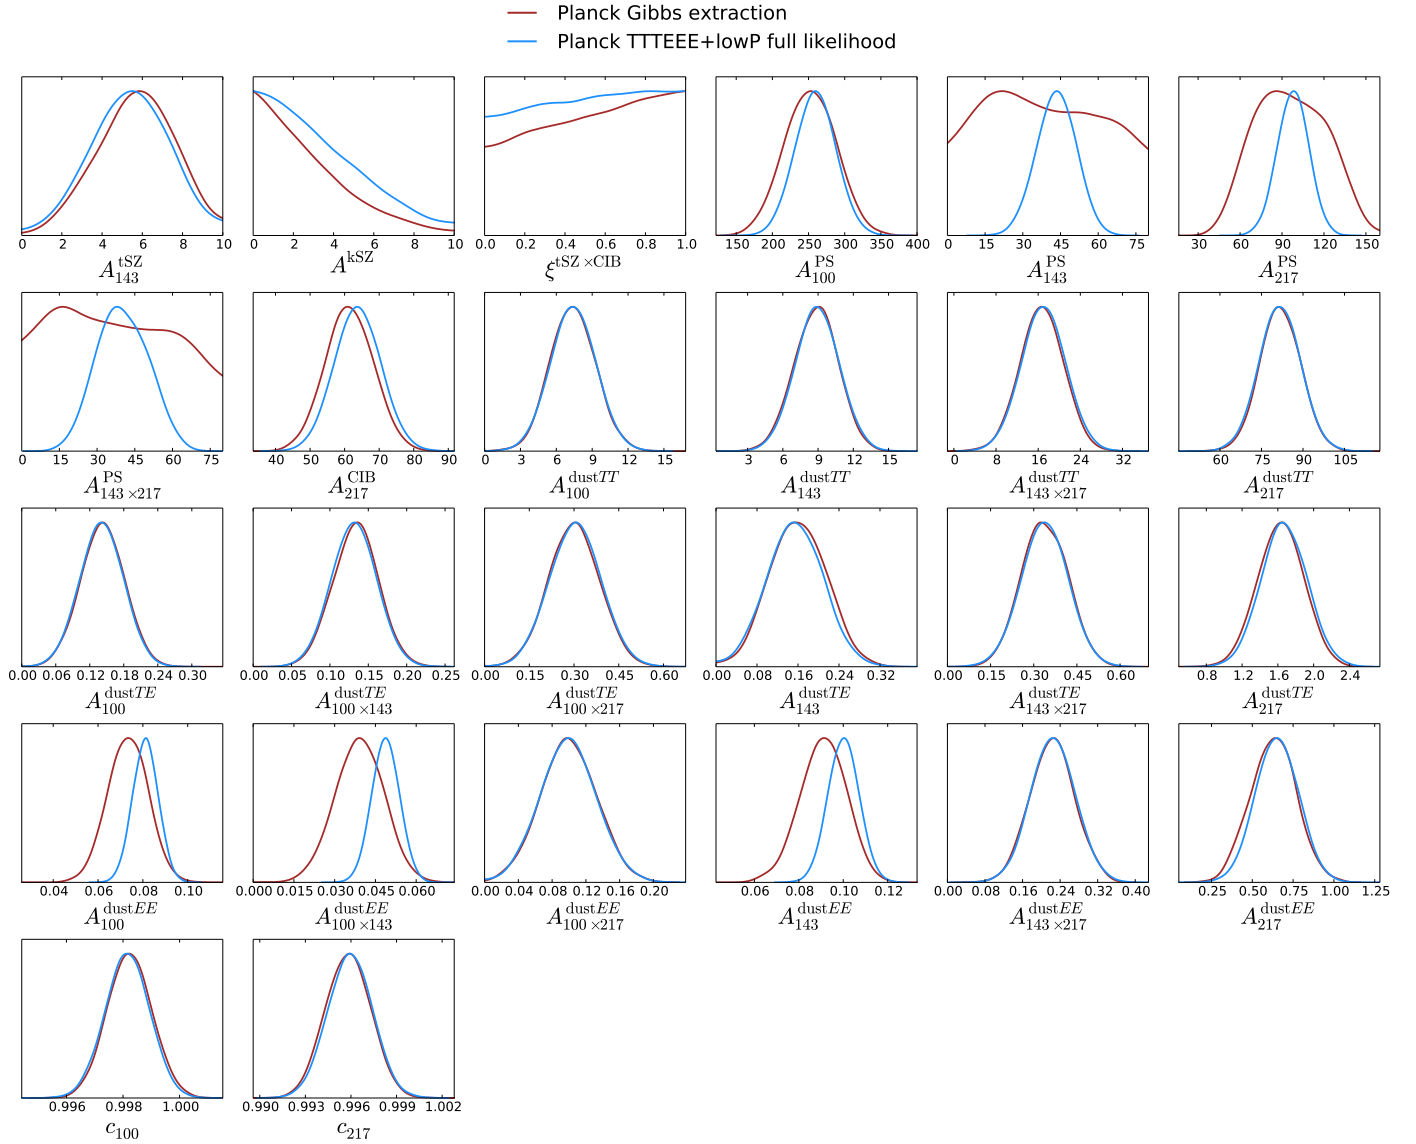
Fig. C.13. Comparison of the nuisance parameters estimated simultaneously with the CMB band-powers (red lines) and the results from the full multi-frequency likelihood (blue lines).

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 ${\rm d}n_{\rm s}/{\rm d}\ln k$, the tensor-to-scalar ratio r, the primordial helium fraction $Y_{\rm P}$, and the lensing amplitude $A_{\rm L}$. 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
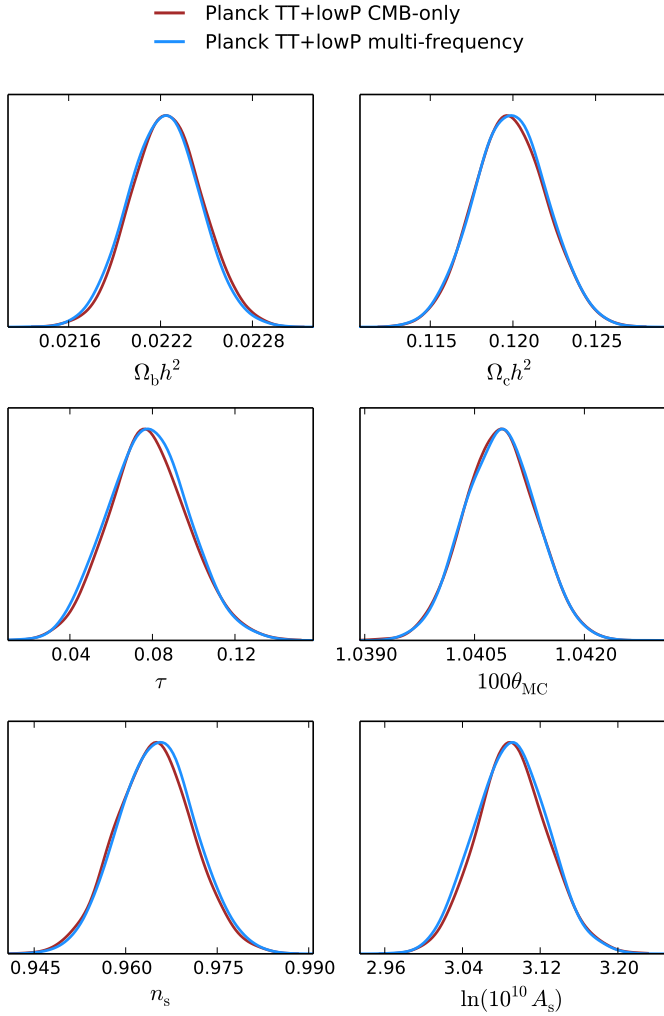
Fig. C.14. Comparison of the six base ΛCDM parameters estimated with the *Planck* compressed CMB-only likelihood (red lines) and the full multi-frequency likelihood (blue lines), in combination with *Planck* lowP data.
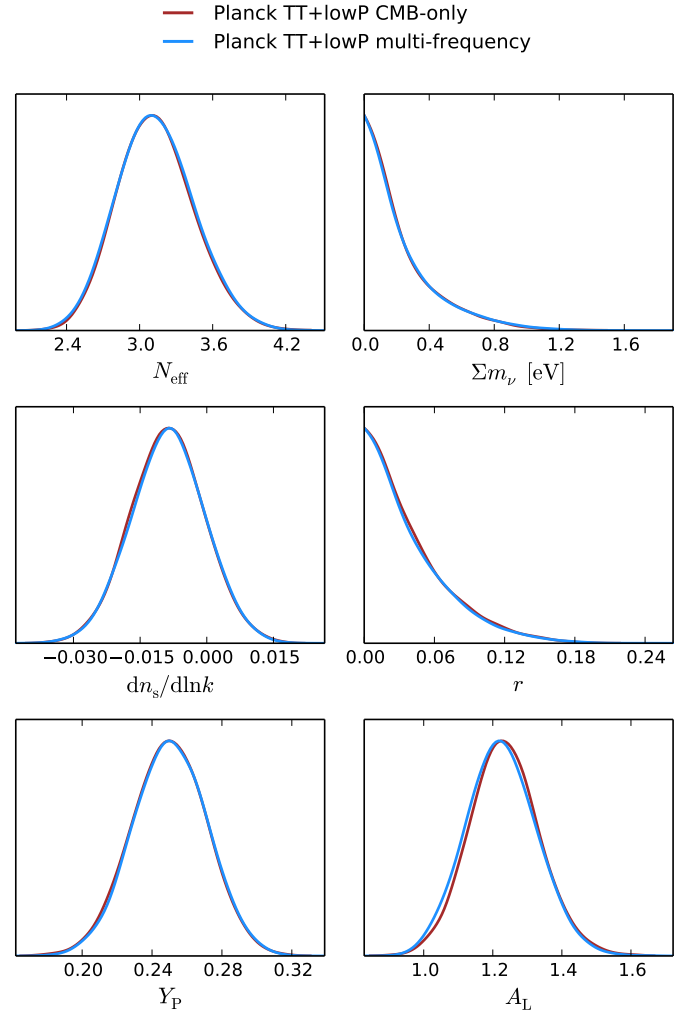
Fig. C.15. Comparison of extensions to the ΛCDM model from the CMB-only likelihood (red) and the multi-frequency likelihood (blue). There is excellent agreement between the two methods.

`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
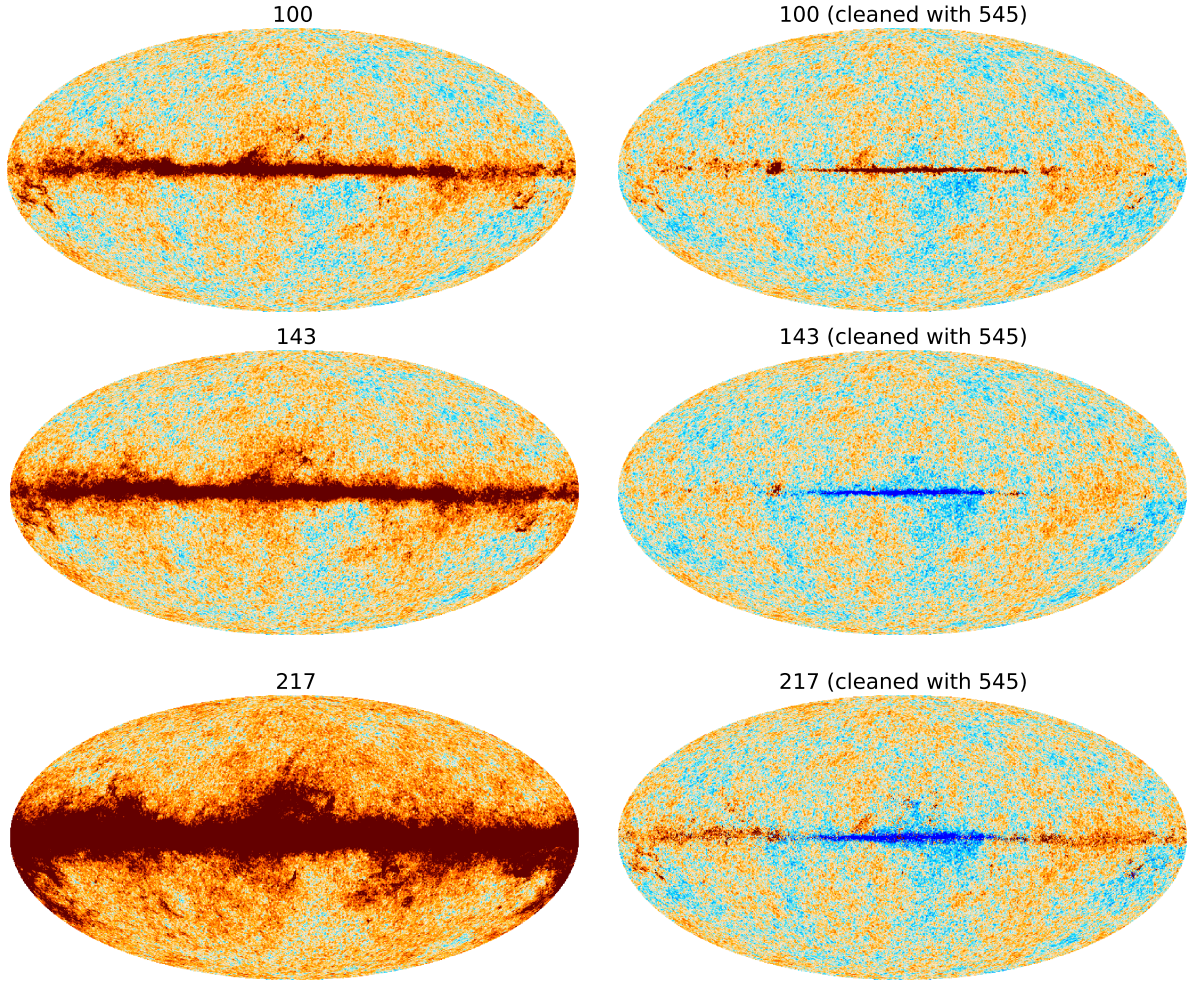
Fig. D.1. *Planck* temperature maps for the three CMB channels computed with `Mspec`. *Left*: raw frequency maps. *Right*: the same maps, after map cleaning with the 545 GHz map. The cleaning coefficients are given in Table D.1. Comparison- of power spectra from these maps with different levels of Galactic masking are shown in Fig. D.2.

correlated than 857 GHz. For polarization, only the 353 GHz detectors are polarization sensitive, and thus we use those.

Although we describe it as "map" subtraction, in practice it is done at the power-spectrum level by forming the equivalent linear combination of power spectra,

$$C_\ell^{\nu_1,\text{clean}} = (C_\ell^{\nu_1} - 2x C_\ell^{\nu_1 \times \nu_2} + x^2 C_\ell^{\nu_2})/(1-x)^2, \quad (D.1)$$

where the C_ℓ are mask- and beam-deconvolved power spectra, x is the cleaning coefficient, $\nu_1 \in \{100, 143, 217\}$ refers to one of the CMB channels, $\nu_2 \in \{353, 545\}$ is the cleaning frequency, and both have the same mask applied. We obtain the cleaning coefficient by maximizing the reduction in pixel variance due to cleaning, filtered to a given ℓ-range and on a particular mask. We choose the filter range to be $\ell = (50, 500)$ since that is where we expect dust to be dominant over extragalactic CIB. We use the $f_{\text{sky}} = 90\%$ mask to maximize the dust signal, fitting it well while avoiding bias from strong signals from the Galactic plane. In other words, we find x for each frequency ν by maximizing

$$\sum_{\ell=50}^{500} (2\ell+1)(C_\ell^\nu - C_\ell^{\nu,\text{clean}}). \quad (D.2)$$

The resulting cleaning coefficients are given in Table D.1. Temperature maps cleaned with these coefficients are shown in

Fig. D.1, visually demonstrating the remarkable effectiveness of this procedure. This can be seen more quantitatively at the power-spectrum level in Fig. D.2, which shows the power spectrum computed on different masks, differenced against the $f_{\text{sky}} = 30\%$ case to cancel out any isotropic components. The middle panel shows that the map-cleaning step leads to about a factor of ten reduction in Galactic power.

This suppression of power not only reduces sensitivity to errors in the Galactic power-spectrum modelling, but it also leads directly to a smaller covariance by reducing chance correlations between signal (CMB and extragalactic foregrounds) and Galactic dust. In general, if a map contains a sum of signal and dust, $T(\hat{n}) = S(\hat{n}) + D(\hat{n})$, its auto-spectrum contains signal, dust, and signal–dust terms, $C_\ell = C_\ell^{SS} + 2C_\ell^{SD} + C_\ell^{DD}$. The DD term can be modelled and subtracted as in the baseline `Plik` likelihood. The SD term has zero mean but non-zero variance, so it must be included in the covariance matrix. For high sky fractions this can become important, e.g., for $f_{\text{sky}} = 80\%$ at 217 GHz it is equal to the noise contribution at $\ell \approx 1500$. Conversely, the signal–dust term is not present if the dust is removed from the map initially as is done by the map-cleaning procedure.

For TT, there is a second step of `Mspec` cleaning that removes any remaining dust contamination left due to spatial variation of the dust spectral index (or equivalently the decorrelation

Table D.1. `Mspec` map cleaning coefficients.

Raw map	x	Template
T 100	0.0013	545
T 143	0.0024	545
T 217	0.0080	545
Q, U 100	0.019	353
Q, U 143	0.040	353
Q, U 217	0.128	353

Notes. The coefficients for the `Mspec` map cleaning procedure, with x defined as in Eq. (D.1)

between the CMB channels and the high-frequency cleaning channel). The level of this residual Galactic contamination can be seen in the single differences in the middle panel of Fig. D.2, however there is too much scatter due to sample variance to derive a template at all multipoles. To remedy this, we consider 353 GHz, where the dust intensity increases relative to the scatter and also subtract 143 GHz to remove the CMB contribution to the scatter. This gives the lower panel of Fig. D.2. The shape of the template is set by fitting a model to these residuals, which we take phenomenologically to be the sum of two power laws. During parameter estimation, this template is added to the foreground model with an amplitude parameter that is marginalized over. We place a tight prior on this amplitude parameter coming from fitting the single-differences directly. For *TE* and *EE*, there is no evidence of residual Galactic contamination for the sky fractions used, and thus we do not perform the second step.

The end result of the `Mspec` cleaning procedure for *TT* is summarized in Fig. D.3, which shows a comparison of `Mspec` and `Plik` power spectra. The top row shows that the map-subtraction cleans out a significant amount of foreground power at 143×217 and 217×217, but makes a very small impact at lower frequencies, as expected. In the bottom row we see power spectra from both codes after total-foreground cleaning, and we find excellent point-by-point agreement. The agreement at the parameter level is also very good, as described in Sect. 4.2. Figure D.3 is similar to Fig. 2 of (Planck Collaboration XIII 2016), which compared `CamSpec` spectra with and without a map-cleaning step, but using a simpler model for the power spectrum of residuals left over after map cleaning. Here, by using the `Mspec` model instead, we confirm the consistency of our results with respect to Galactic cleaning using a more realistic foreground model.

Covariance approximations. `Mspec` computes the covariance of the pseudo-C_ℓs in the same way as the baseline likelihood, but with two additional approximations. First, we ignore the fact that the noise is non-uniform across the sky. Using the FFP8 simulations (Planck Collaboration XII 2016), which include realistic spatial variations of the noise, we find that this is an excellent approximation for *TT* and *TE*, where the noise is only important at high ℓ, but leads to about a 5% underestimation of the error bars for low-ℓ *EE*, which we correct for heuristically. Second, we always take the coupling kernel that appears in the pseudo-C_ℓ covariance to be the *TT* one. This leads to much less than 1% changes. The benefit is a huge simplification of the covariance expressions, allowing for all the temperature and polarization
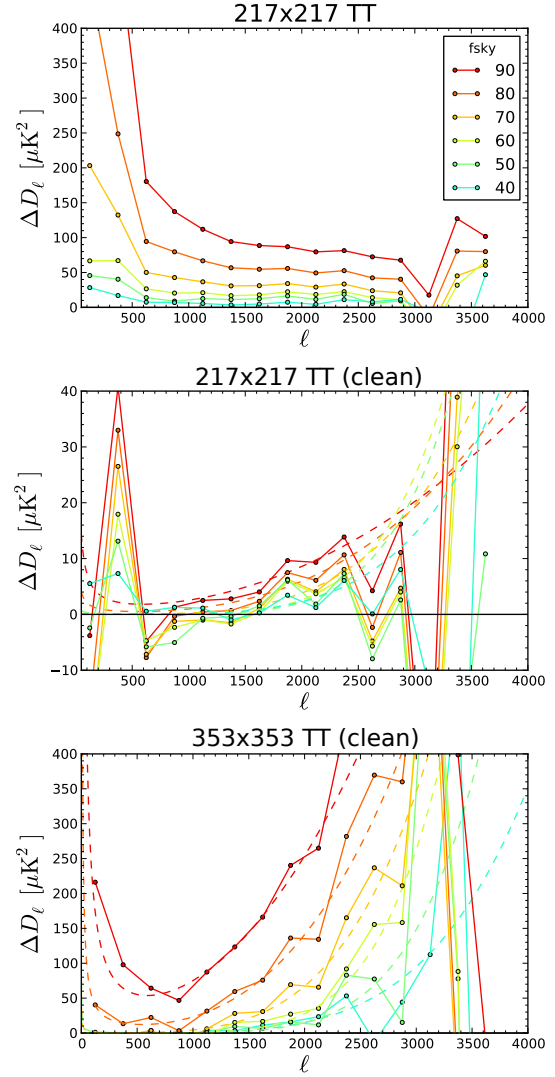

Fig. D.2. *Top*: power spectra of the 217 GHz raw temperature maps, with various Galactic masks, differenced against the $f_{\rm sky} = 30\%$ case. *Middle*: the same, but after map cleaning. The y-scale is now 10 times smaller. *Bottom*: the same as the middle panel, but for 353 GHz. Additionally, we first subtract the 143 GHz map to remove scatter due to CMB sample variance. A model consisting of the sum of two power laws is then fit through these points. This sets the shape of the residual Galactic template. In the middle panel this template is scaled by an amplitude to fit the residuals seen there, showing that the Galactic residuals after map cleaning do not change shape significantly between 353 and 217 GHz.

entries to be written concisely as

$$\langle \Delta \widetilde{C}_\ell^{ab,XY} \Delta \widetilde{C}_{\ell'}^{cd,ZW} \rangle = \frac{1}{2} \Big\{ C_{(\ell}^{ac,XZ} C_{\ell')}^{bd,YW} \Xi_{\ell\ell'}^{TT}[W^{ac,bd}]$$
$$+ C_{(\ell}^{ad,XW} C_{\ell')}^{bc,YZ} \Xi_{\ell\ell'}^{TT}[W^{ad,bc}] \Big\}, \quad (D.3)$$

where a, b, c, d label detectors, X, Y, Z, W each label one of T, E, B, and the C_ℓ's that appear on the right-hand side are the fiducial beam-convolved signal-plus-noise power-spectra.

D.2. `Hillipop`

`HiLLiPOP` is another high-ℓ likelihood procedure, based on a Gaussian approximation, to confront the *Planck* HFI data ℓ-by-ℓ with cosmological models. In this approach, the data consist
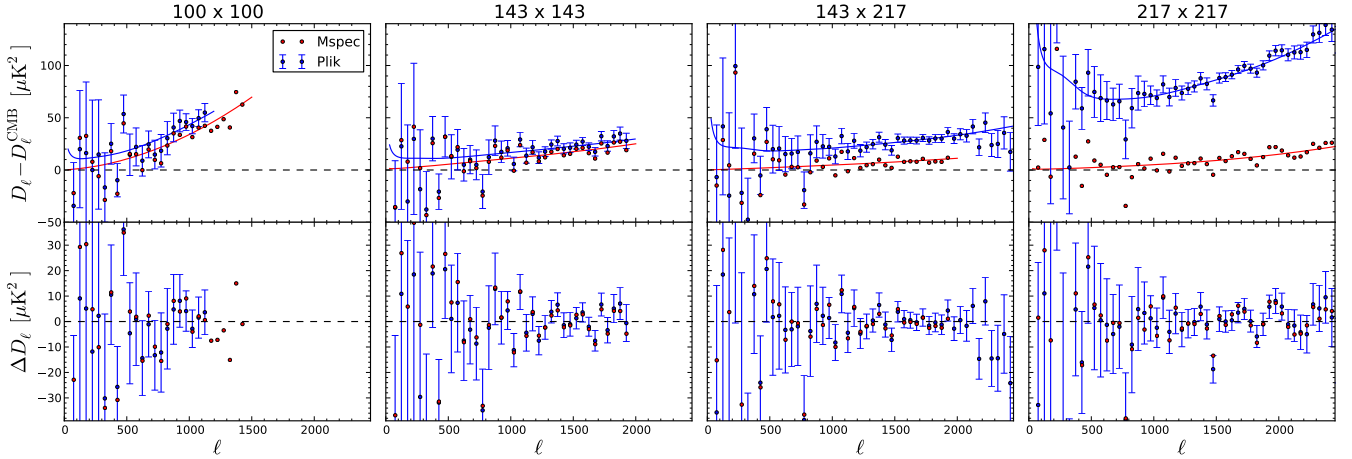
Fig. D.3. Comparison of power spectra on identical masks obtained from the `Plik` Galactic cleaning procedure and from the `Mspec` map cleaning. In the *top panels* we subtract the best-fit CMB spectrum as determined by each code; in the *bottom panel* we additionally subtract the Galactic and extragalactic foreground model.

of six maps: two sets of half-mission (I, Q, U) maps at 100, 143, and 217 GHz. Frequency-dependent apodized masks are applied to these maps in order to limit contamination from diffuse Galactic dust, Galactic CO lines, nearby galaxies, and extragalactic point sources. With regard to the latter, unlike the `Plik` masks, which are based on the point source catalogue with a flux-density cut (see Appendix A), the masks used here rely on a more refined procedure that preserves Galactic compact structures and ensures the completeness level at each frequency, but with a higher detection threshold (i.e., leaving more extragalactic sources unmasked). `Hillipop` retains 72, 62, and 48% of the sky at 100, 143, and 217 GHz, respectively, and uses the same set of masks for both temperature and polarization. Mask-deconvolved and beam-corrected cross-half-mission power spectra are computed using `Xpol`, an extension of the `Xspect` (Tristram et al. 2005a) code to polarization. From the six maps, we can derive 15 sets of power spectra in TT, EE, TE, and ET: one each for 100×100, 143×143, and 217×217, and four each for 100×143, 100×217, and 143×217.

The covariance matrix, which encompasses the correlations (ℓ-by-ℓ) between all these 60 power spectra, is estimated semi-analytically with `Xpol`. Unlike the `Plik` likelihood, which assumes a model for signal (of cosmological and astrophysical origin) and noise, here the calculation relies on data estimates only by using as input a smooth version of the estimated power spectra. Contributions from noise, sky emission, and the associated cosmic variance are automatically taken into account. Several approximations (as in Efstathiou 2006) are needed in this calculation and Monte Carlo simulations have been performed to test their accuracy. A precision better than a few percent is achieved.

In addition to the CMB component, the `HiLLiPOP` likelihood accounts for foreground residuals and differences in calibration between maps. A differential calibration coefficient, d_i, is defined per map, $\tilde{m}_i = (1 + d_i) m_i$, with the 143 half-mission-1 map calibration taken as reference.[23] We use different models for foregrounds in temperature and polarization. The temperature model includes contributions from cosmic infrared background (CIB), Galactic dust, thermal and kinetic Sunyaev-Zeldovich (tSZ and kSZ) effects, Poisson point sources (PS), and the cross-correlation between infrared galaxies and the tSZ effect (tSZ×CIB). The polarization model includes only Galactic dust.
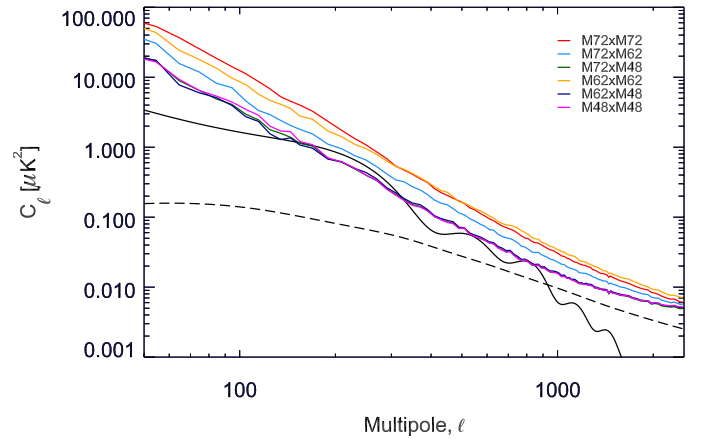
---

[23] Therefore, e.g., $\tilde{C}_\ell^{143h2 \times 217h1} = (1 + d_{143h2} + d_{217h1}) C_\ell^{143h2 \times 217h1}$.

Fig. D.4. Cross-power spectra of half-mission maps at 353 GHz in each combination of masks after subtraction of CMB (solid line) and CIB (dashed line) models.

The calibration coefficients are assigned Gaussian priors reflecting the uncertainties in the half-mission map calibration (from Planck Collaboration VIII 2016): $d_{100h1} = d_{100h2} = d_{143h1} = d_{143h2} = 0.000 \pm 0.002$ and $d_{217h1} = d_{217h2} = 0.004 \pm 0.002$.

`HiLLiPOP` uses physically motivated templates of foreground emission power spectra, in both photon frequency and spherical harmonic wave number, based on *Planck* measurements, as described below.

– Galactic dust: The TT, EE, and TE Galactic dust power-spectrum templates are derived following the methodology presented in Planck Collaboration Int. XXX (2016). We first estimate the TT, EE, and TE half-mission cross-power spectra at 353 GHz in the different combinations of masks (Fig. D.4) and then subtract the best-fit CMB power spectrum (see Planck Collaboration XVI 2014). In temperature, the CIB power spectrum (see Planck Collaboration XXX 2014) is also subtracted. A good fit is obtained on the resulting power spectra using the power-law model $A\ell^\alpha + B$, where B describes the Poisson contribution from unresolved point sources in temperature (the contribution of polarized sources is negligible in EE and TE). It is worth stressing that the shape of the Galactic dust component is found to be compatible with a power law. This is due to the choice

of point-source masks that have minimal effect on Galactic structures or bright cirrus. The masks used by Plik, which are based on lower flux-density cuts to further reduce the compact-source contribution, introduce a knee in the power spectra around $\ell \approx 300$.

– tSZ: The power-spectrum template is based on Tinker et al. (2008) for the mass function and Arnaud et al. (2010) for the "universal pressure profile". It contains both the one-halo and the two-halo terms (Taburet et al. 2011). A full description can be found in Planck Collaboration XXI (2014).

– kSZ: The power spectrum is taken from Battaglia et al. (2013) for the patchy reionization part and Shaw et al. (2012) for the Ostriker-Vishniac effect, both normalized to the *Planck* cosmological parameters.

– CIB and tSZ×CIB: The CIB power-spectrum templates are based on a halo model linking directly the galaxies' luminosities to their host dark-matter halo masses. It has been successfully applied in Planck Collaboration XXX (2014). For the tSZ×CIB power-spectrum templates, the tSZ power-spectrum template is based on Efstathiou & Migliaccio (2012).

– PS: we used a Poisson-like flat power spectrum for the unresolved point source contribution at each cross-frequency.

A single free parameter (A, scaling all frequencies equally) for each foreground template is used to adjust the amplitude of the power spectra. Each amplitude is assigned a uniform prior.

The model used to describe the *Planck TT* power spectra thus reads

$$\hat{C}_\ell^{T_i T_j} = (1 + c_i + c_j)\Big(C_\ell^{\mathrm{CMB},TT} + A_{\mathrm{dust}}^{TT} C_\ell^{\mathrm{dust},T_i T_j} + A_{\mathrm{PS}}^{T_i T_j}$$
$$+ A_{\mathrm{CIB}} C_\ell^{\mathrm{CIB},T_i T_j} + A_{\mathrm{tSZ}} C_\ell^{\mathrm{tSZ},T_i T_j}$$
$$+ A_{\mathrm{kSZ}} C_\ell^{\mathrm{kSZ}} + A_{\mathrm{tSZ} \times \mathrm{CIB}} C_\ell^{\mathrm{tSZ} \times \mathrm{CIB},T_i T_j}\Big),$$

while the *EE* and *TE* power spectrum models simply read

$$\hat{C}_\ell^{E_i E_j} = \left(1 + c_i + c_j\right)\left(C_\ell^{\mathrm{CMB},EE} + A_{\mathrm{dust}}^{EE} C_\ell^{\mathrm{dust},E_i E_j}\right),$$
$$\hat{C}_\ell^{T_i E_j} = \left(1 + c_i + c_j\right)\left(C_\ell^{\mathrm{CMB},TE} + A_{\mathrm{dust}}^{TE} C_\ell^{\mathrm{dust},T_i E_j}\right).$$

Among all the map cross-power spectra, we focus on just the six frequency cross-spectra (100×100, 100×143, 100×217, 143×143, 143×217, and 217×217) in *TT*, *EE*, and *TE* and compress the covariance matrix accordingly. The Hillipop likelihood then reads

$$-2 \ln \mathcal{L} = \sum_{\substack{X,Y \\ X',Y'}} \sum_{\substack{i \leqslant j \\ i' \leqslant j'}} \sum_{\substack{\ell = \ell_{\min}^{X_i Y_j} \\ \ell' = \ell_{\min}^{X'_{i'} Y'_{j'}}}}^{\substack{\ell_{\max}^{X_i Y_j} \\ \ell_{\max}^{X'_{i'} Y'_{j'}}}} \mathcal{R}_\ell^{X_i Y_j} \left[\Sigma_{\ell \ell'}^{X_i Y_j, X'_{i'} Y'_{j'}}\right]^{-1} \mathcal{R}_{\ell'}^{X'_{i'} Y'_{j'}}, \qquad (\mathrm{D.4})$$

where $\mathcal{R} = C_\ell - \hat{C}_\ell$ denotes the residual of the estimated power spectrum (C_ℓ) with respect to the model ($\hat{C}_\ell$), and Σ is the full covariance matrix, which is symmetric and positive-definite. The frequency band (100, 143, or 217 GHz) is given by the i, j indices and the CMB modes (T, E) by X, Y. The multipole ranges $[\ell_{\min}, \ell_{\max}]$ are chosen with the goal of limiting contamination in each power spectrum from diffuse Galactic dust emission at low ℓ and noise at high ℓ.

At the end, we have a total of 5 instrumental, 13 astrophysical, and 6 or more cosmological (ΛCDM and possible extensions) parameters, i.e., a total of 24 (or more) free parameters in
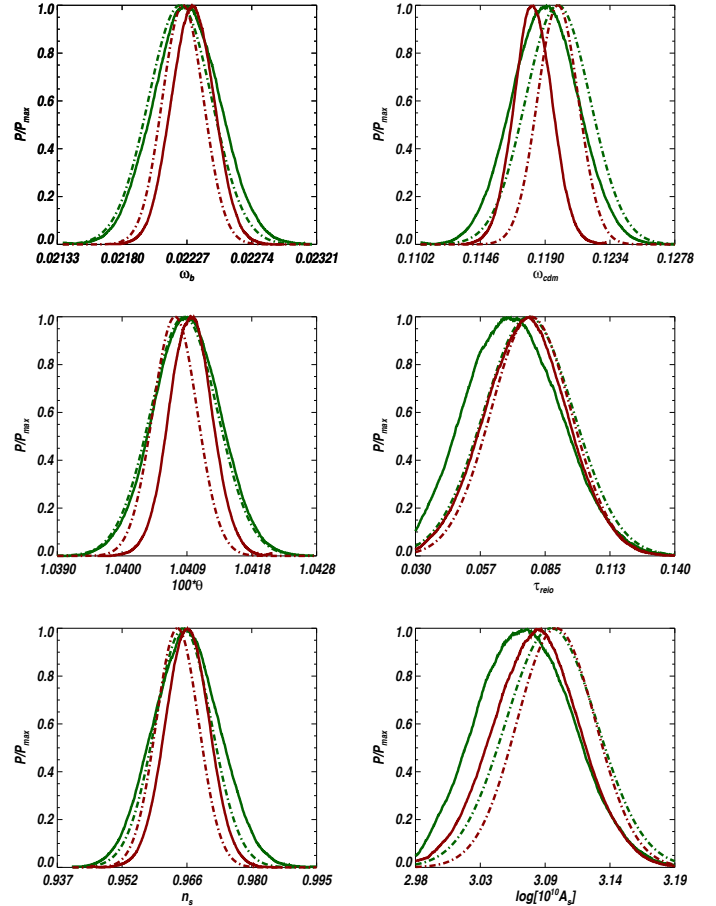

Fig. D.5. Marginalized constraints for the base-ΛCDM model obtained with the HiLLiPOP likelihood using *TT* only (green solid line) and the full temperature and polarization data (red solid line). For comparison Plik *TT* (dashed green) and Plik full data (dashed red) results are also shown. The high-ℓ information is always complemented with lowTEB.

Table D.2. HiLLiPOP goodness of fit.

CMB mode	χ^2	n_ℓ	$\Delta\chi^2/\sqrt{2n_\ell}$
TT	9949.7	9556	2.85
EE	7309.5	7256	0.44
TE	9322.5	8806	3.89

Notes. The $\Delta\chi^2 = \chi^2 - n_\ell$ is the difference from the mean (n_ℓ), assuming that the model follows a χ^2 distribution. The quantity $\Delta\chi^2/\sqrt{2n_\ell}$ is $\Delta\chi^2$ in units of the standard deviation. These numbers correspond to the global, all-frequency fits, and are therefore not directly comparable to the Plik numbers for the co-added CMB spectra.

the full (*TT*, *EE*, and *TE*) HiLLiPOP likelihood function. The theoretical CMB power spectra are generated with the CLASS Boltzmann solver (Lesgourgues 2011; Blas et al. 2011) or the PICO algorithm (Fendt & Wandelt 2007b). The HiLLiPOP likelihood function is explored using Minuit (James & Roos 1975). Table D.2 shows the best-fit χ^2 values for *TT*, *EE*, *TE*, and the full data set. The number of degrees of freedom (n_ℓ) is simply the total number of multipoles considered.

The posterior distributions of the six base ΛCDM parameters obtained with MCMC sampling are shown in Fig. D.5 in comparison with the baseline Plik results. Here the high$-\ell$ information is complemented with lowTEB. When considering

the *TT* data only, almost all parameters are compatible with the baseline within approximately 0.1σ, with the exception of $\Omega_c h^2$ where the shift is slightly higher (about 0.4σ). The approximately 0.5σ difference in τ and A_s can be understood as a mild preference of the `HiLLiPOP` likelihood for a lower A_L (1.20 ± 0.09 compared to 1.23 ± 0.10 for `Plik`). The shifted value for A_L comes in both cases from a tension between high ℓ and lowTEB, the value from `HilliPOP` alone for this parameter being compatible with unity at the 1σ level. Error bars from the baseline `Plik` and `HiLLiPOP` are nearly identical, with only a slightly bigger error bar for n_s in the latter.

When considering the full data set, the shifts with respect to the baseline `Plik` are more pronounced, but still remain within about 0.5σ. We observe the same trend for a lower $\Omega_c h^2$, but the difference in τ and A_s is instead alleviated by the compatibility of A_L for the full likelihoods. Again the error bars derived with both likelihood methods are nearly identical.

We checked the robustness of our results with respect to the foreground models. Recall that we adjust each foreground component to data using a single rescaling amplitude. In temperature, the dust amplitude is recovered almost perfectly ($A_{dust}^{TT} \approx 1$) while the A_{CIB} estimation lies 1.7σ away. Using the full likelihood, the dust *TT* remains perfectly compatible with $A_{dust}^{TT} = 1$ and the shift in the CIB amplitude is reduced to 0.5σ. The polarized dust amplitudes (A_{dust}^{TE} and A_{dust}^{EE}) are compatible with unity within about 2σ (this is true also for single *EE* and *TE* spectra). As described in Planck Collaboration XIII (2016), using *Planck*-only data, we are not very sensitive to SZ components. In any event, the marginalized posteriors on A_{tSZ}, A_{kSZ}, and $A_{tSZ\times CIB}$ are compatible with the expectations.

Compatibility with $A = 1$ for the foreground scaling parameters is a good indication of the consistency of the internal *Planck* templates. We have also tested the stability of the cosmological results with respect to the choice of priors on these scaling parameters, by considering a set of Gaussian priors, $A = 1.0 ± 0.2$. The χ^2 values remain unchanged, with almost no shift in cosmological parameters and a slight reduction of the error bar on n_s.

Appendix E: The *Planck* CMB likelihood supplement

In this appendix, we provide additional material to further characterize the CMB power spectra from *Planck*. Table E.1 provides goodness-of-fit values for the full likelihood; comparison with the equivalent Table 16 for `PlikTT+tauprior` shows the effect of replacing a τ prior by the full low-ℓ likelihood, which encompasses the $\ell \approx 20$ dip in *TT* power. In Sect. E.1 we extend our tests to combine low-ℓ temperature and polarization information with the baseline `Plik` likelihood at high ℓ. In the following Sect. E.2, we fit the shapes of the spectra as a series of (Gaussian) peaks and troughs, and in Sect. E.3 we provide an alternative display of the correlation between temperature and (*E*-mode) polarization. Section E.4 discusses the spectra and parameters one can derive by analysing the CMB maps cleaned by component-separation methods, and Sect. E.5 confirms that the profile likelihood approach (Planck Collaboration Int. XVI 2014) leads to parameter constraints consistent with those discussed in the main text. This offers a breakdown of the uncertainties on the cosmological parameters between those that arise from the finite sensitivity of the experiment and those coming from a lack of knowledge of the nuisance parameters related to the modelled foregrounds and instrumental effects.

E.1. TT, TE, EE robustness tests

This section presents the tests conducted so far to assess the robustness and accuracy of the results when using the high-ℓ `PlikTT, EE, TE` likelihood in combination with the low-ℓ likelihood in temperature and polarization.

The results are summarized in Fig. E.1, which shows the marginal mean and the 68% CL error bars for cosmological parameters using the `PlikTT, TE, EE+lowTEB` data combination and different assumptions about the data selection, foreground model, or treatment of the systematics for the `PlikTT, EE, TE` likelihood. In the following, we comment, in turn, on each of the tests shown in this figure (from left to right). The reference `PlikTT, TE, EE+lowTEB` results for the ΛCDM model are denoted as "`PlikTT, TE, EE+lowTEB`". All the tests are run with the `PICO` code. Most of the shifts we observe in the different cases are consequences of the shifts we observe in the *TT*, *TE*, and *EE* tests presented in Sect. 4 and Appendix C.3.5.

Detsets. We find good agreement between the baseline cases that use half-mission spectra and those that use detsets (case "DS"). The greatest deviation is an upward 0.5σ shift in θ, driven by the upward shifts observed in *TE* and *EE* in Fig. C.10.

Larger Galactic mask. We examined the impact of using larger Galactic masks, G50, G41, and G41 (case "M605050"), instead of the baseline ones at 100, 143, and 217 GHz, respectively (see also Sect. 3.2.2). We observe substantial shifts in the parameters, at the level of ≲1σ. These are mainly driven by the shifts already observed in the *TE* test in Fig. C.10.

Galactic dust priors. We find that leaving the Galactic dust amplitudes completely free to vary ("No gal. priors") in both temperature and polarization, without applying the priors described in Sect. 3.3.1, does not have a significant impact on cosmological parameters.

Cutting out frequency channels. We have considered the cases where we eliminate all the power-spectra related to one particular frequency at a time. We see shifts at the level of 1−2σ when eliminating all the power spectra containing the 100 GHz ("no100") or the 143 GHz ("no143") maps. These reflect the analogous shifts observed in the *EE* and *TE* tests.

Changing ℓ_{min} We find good stability in the results when changing the minimum multipole ℓ_{min} considered in the analysis for the `PlikTT, EE, TE` likelihood ("LMIN" case). The baseline likelihood has $\ell_{min} = 30$, and we test the cases of $\ell_{min} = 50$ and $\ell_{min} = 100$. For each frequency power spectrum in temperature or polarization we choose $\ell_{min}^{freq} = \max(\ell_{min}, \ell_{min}^{freq, base})$, with $\ell_{min}^{freq, base}$ being the baseline value of ℓ_{min} at each frequency, as reported in Table 16. For example, for the 217 × 217 spectrum in *EE* we use the baseline value $\ell_{min} = 505$ in all the cases we consider here.

Changing ℓ_{max}. We observe shifts when including maximum multipoles between $\ell_{max} \approx 800$ and 2500 ("LMAX" cases), consistent with the shifts already observed in *TT* in Fig. 35. For each frequency power spectrum in temperature or polarization

Table E.1. Goodness-of-fit tests for the 2015 *Planck* temperature and polarization spectra.

Frequency [GHz]	Multipole range[a]	χ^2 [b]	χ^2/N_ℓ	N_ℓ	$\Delta\chi^2\sqrt{2N_\ell}$[d]	PTE[%][d]
_TT_						
100×100	30–1197	1234.37	1.06	1168	1.37	8.7
143×143	30–1996	2034.45	1.03	1967	1.08	14.1
143×217	30–2508	2566.74	1.04	2479	1.25	10.7
217×217	30–2508	2549.66	1.03	2479	1.00	15.8
Combined _TT_	30–2508	2546.67	1.03	2479	0.96	16.8
_TE_						
100×100	30–999	1088.78	1.12	970	2.70	0.5
100×143	30–999	1032.84	1.06	970	1.43	7.9
100×217	505–999	526.56	1.06	495	1.00	15.8
143×143	30–1996	2028.43	1.03	1967	0.98	16.4
143×217	505–1996	1606.25	1.08	1492	2.09	2.0
217×217	505–1996	1431.52	0.96	1492	−1.11	86.7
Combined _TE_	30–1996	2046.11	1.04	1967	1.26	10.5
_EE_						
100×100	30–999	1027.89	1.06	970	1.31	9.6
100×143	30–999	1048.22	1.08	970	1.78	4.1
100×217	505–999	479.72	0.97	495	−0.49	68.1
143×143	30–1996	2000.90	1.02	1967	0.54	29.2
143×217	505–1996	1431.16	0.96	1492	−1.11	86.8
217×217	505–1996	1409.58	0.94	1492	−1.51	93.6
Combined _EE_	30–1996	1986.95	1.01	1967	0.32	37.2

Notes. [a] The ℓ range used in the high-ℓ likelihood. [b] The χ^2 are with respect to the best-fit from the `PlikTT+lowTEB` data combination in a ΛCDM framework. [c] $\Delta\chi^2 = \chi^2 - N_\ell$ is the difference from the mean assuming that the best-fit base _TT_ ΛCDM model is correct, expressed in units of the expected dispersion, $\sqrt{2N_\ell}$. [d] Probability to exceed the tabulated value of χ^2.

we choose $\ell_{\max}^{\mathrm{freq}} = \min(\ell_{\max}, \ell_{\max}^{\mathrm{freq,base}})$, with $\ell_{\max}^{\mathrm{freq,base}}$ being the baseline value of $\ell_{\max}$ at each frequency, as reported in Table 16.

Comparison to `CamSpec`. We find roughly 1σ shifts when comparing to the results obtained with the `CamSpec` code. Furthermore, we note that our "M605050" case is in better agreement with `CamSpec`. This is due to the fact that while `CamSpec` uses the same masks as `Plik` in temperature, it uses much larger masks in polarization, retaining about 50% of the sky (before apodization). We showed in Sect. 4.1.2 and Appendix C.3.5 that using a larger Galactic mask does not affect the _TT_ results, while it does change the _TE_ and _EE_ results. We therefore expect the results from the `CamSpec` code in the TT, TE, EE+lowTEB case to be in better agreement with our "M605050" test.

E.2. Peaks and troughs in Planck power spectra

The power spectrum of CMB temperature anisotropies has been measured by *Planck* to an exquisite level of precision, and the C_ℓ^{TT} spectrum has now been joined by C_ℓ^{TE} and C_ℓ^{EE}. This precision enables us to estimate the underlying cosmological parameters, but the C_ℓs are themselves a set of cosmological observables, whose properties can be described independently of any model. The peaks and troughs in the power spectra reflect their origin in oscillating sound waves. These features tell us that the Universe once contained a very hot, dense plasma, with the CMB anisotropies mostly originating from acoustic modes in the coupled photon-baryon fluid, driven by dark matter potential perturbations (e.g., Hu et al. 1997). The overall angular structure is determined by θ_*, the ratio of the sound horizon to the last-scattering surface distance, the statistical quantity that is best

constrained by *Planck*. However, the positions of the *individual* peaks and troughs are now well determined in their own right, and this information has become part of the canon of facts now known about our Universe.

Here we use the *Planck* data directly to fit for the multipoles of individual features in the measured _TT_, _TE_, and _EE_ power spectra. We specifically use the CMB-only bandpowers described in this paper and available in the *Planck* Legacy Archive[24], adopting the same weighting scheme within each bin. Fitting for the positions and amplitudes of features in the band-powers is a topic with a long history, with approaches becoming more sophisticated as the fidelity of the data improved (e.g., Scott & White 1994; Hancock & Rocha 1997; Knox & Page 2000; de Bernardis et al. 2002; Bond et al. 2003; Page et al. 2003; Benoît et al. 2003; Durrer et al. 2003; Readhead et al. 2004; Jones et al. 2006; Hinshaw et al. 2007; Corasaniti & Melchiorri 2008; Pryke et al. 2009). Following earlier approaches, we fit Gaussian functions to the peaks and troughs in C_ℓ^{TT} and C_ℓ^{EE}, but parabolas for the C_ℓ^{TE} features. We have to remove a featureless damping tail (which we do by using an extreme amount of lensing to wash out the structure) to fit the higher-ℓ C_ℓ^{TT} region and care has to be taken to treat the lowest-ℓ "recombination" peak in C_ℓ^{EE}. We do not try to fit the "reionization" bumps in the lowest few multipoles of C_ℓ^{TE} and C_ℓ^{EE} (even though these might technically be the "first peaks" in these power spectra). We explicitly focus on features in the conventional quantity $\mathcal{D}_\ell \equiv \ell(\ell+1)C_\ell/2\pi$; note that other quantities (e.g., C_ℓ) have maxima and minima at slightly different multipoles, and that the selection of which band-powers to use for fitting each peak is somewhat subjective.

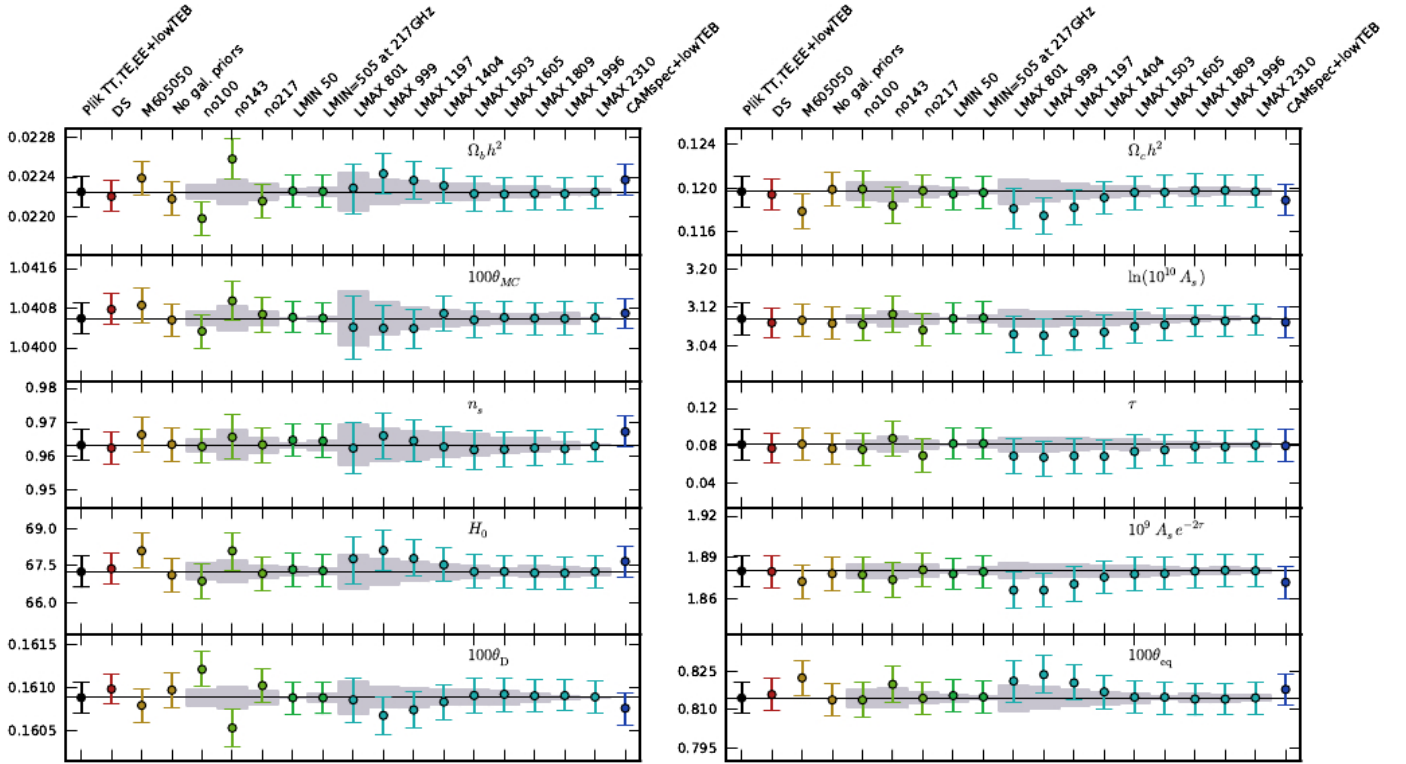
[24] http://pla.esac.esa.int/pla/

Fig. E.1. Marginal mean and 68% CL error bars on cosmological parameters estimated when adopting different data choices for the `Plik` likelihood, in comparison with results from alternate approaches or model. We assume a ΛCDM model and use variations of `PlikTT, EE, TE` in combination with the lowTEB likelihood for most of the cases. The "`PlikTT, EE, TE+lowTEB`" case (black dot and thin horizontal black line) indicates the baseline (HM, $\ell_{min} = 30$, $\ell_{max} = 2508$), 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 data (see Eq. (53)). All the cases are run with `PICO`.

Our numerical values, presented in Table E.2, are consistent with previous estimates, but with a dramatically increased number of features measured. *Planck* detects 36 extrema in total, consisting of 19 peaks and 17 troughs. The eighth C_ℓ^{TT} peak is only marginally detected using the released likelihood (although this cannot be seen from the values in the table, since the "height" includes the amplitude of the featureless spectrum). However, by digging further into the *Planck* data it would be possible to strengthen this detection and perhaps distinguish more features, since using the fact that the foreground power spectra have no structure on the relevant scales, one could be more liberal with foreground contamination for the purposes of feature detection.

E.3. T–E correlations in Planck power spectra

The information contained in the primary CMB anisotropies comes from both temperature and polarization. It is well known that the T and E-mode fields are correlated, and hence one has to measure C_ℓ^{TE}, as well as C_ℓ^{TT} and C_ℓ^{EE}, in order to extract all the information from T and E (see, e.g., Zaldarriaga & Seljak 1997; Kamionkowski et al. 1997; Hu & White 1997).

The C_ℓ^{TE} power spectrum has an amplitude that is roughly the geometric mean of the C_ℓ^{TT} and C_ℓ^{EE} power spectra and it oscillates in sign as a function of ℓ, depending on whether the fields are correlated or anticorrelated. Now that *Planck* has provided high quality measurements of all three power spectra over a wide range of multipoles, it is possible to plot the strength of

this correlation directly. In other words one can form the quantity

$$r_\ell \equiv \frac{C_\ell^{TE}}{\left(C_\ell^{TT} C_\ell^{EE}\right)^{1/2}} = \frac{\mathcal{D}_\ell^{TE}}{\left(\mathcal{D}_\ell^{TT} \mathcal{D}_\ell^{EE}\right)^{1/2}}, \tag{E.1}$$

which is the Pearson correlation coefficient for T and E in harmonic space.

We plot this in Fig. E.2, using the binned $\mathcal{D}_\ell^{TT}$, $\mathcal{D}_\ell^{TE}$, and $\mathcal{D}_\ell^{EE}$ values from the *Planck* 2015 data release. As expected the data fit the theoretical expectation (except for a slight bias at high multipoles, since we have to remove ℓ bins for which the EE power is measured to be quite small).

For the best-fit cosmology one can see that the maximum correlation is at $\ell \approx 300$ (actually $\ell = 282$), where T and E are 63% correlated. The most anti-correlated scale is at $\ell \approx 150$ (actually $\ell = 148$), where $r = -0.66$. When first observed (Kogut et al. 2003) the anti-correlation at relatively large scales, i.e., $r < 0$ for $50 \lesssim \ell \lesssim 250$, was seen as a confirmation of the adiabatic nature of the density perturbations on super-Hubble scales (Peiris et al. 2003). The oscillatory behaviour of all three power spectra of course confirms adiabaticity much more dramatically. But more quantitatively, looking at the r_ℓ plot enables us to directly gauge the strength of the $T-E$ correlation as a function of scale. The fact that the power spectrum is mostly negative (i.e., the troughs are typically deeper than the heights of the peaks) is the usual "baryon drag" effect (e.g., Hu et al. 1997).

Because $-1 \leq r \leq 1$, the Pearson correlation coefficient is sometimes interpreted geometrically as the cosine of a correlation angle. Figure E.2 could hence be plotted with this angle on the vertical axis, but such a plot would carry no additional

Table E.2. Positions and amplitudes of extrema in power spectra using *Planck* data.

Extremum	Multipole	Height [μK^2]
TT power spectrum		
Peak 1	220.0 ± 0.5	5717 ± 35
Trough 1	415.5 ± 0.8	1696 ± 13
Peak 2	537.5 ± 0.7	2582 ± 11
Trough 2	676.1 ± 0.8	1787 ± 12
Peak 3	810.8 ± 0.7	2523 ± 10
Trough 3	997.7 ± 1.4	1061 ± 5
Peak 4	1120.9 ± 1.0	1237 ± 4
Trough 4	1288.8 ± 1.6	737 ± 4
Peak 5	1444.2 ± 1.1	797.1 ± 3.1
Trough 5	1621.2 ± 2.3	400 ± 4
Peak 6	1776 ± 5	377.4 ± 2.9
Trough 6	1918 ± 7	245 ± 4
Peak 7	2081 ± 25	214 ± 4
Trough 7	2251 ± 8	119.5 ± 3.5
Peak 8	2395 ± 24	105 ± 4
TE power spectrum		
Trough 1	150.0 ± 0.8	-48.0 ± 0.8
Peak 1	308.5 ± 0.4	115.9 ± 1.1
Trough 2	471.2 ± 0.4	-74.4 ± 0.8
Peak 2	595.3 ± 0.7	28.6 ± 1.1
Trough 3	746.7 ± 0.6	-126.9 ± 1.1
Peak 3	916.9 ± 0.5	58.4 ± 1.0
Trough 4	1070.4 ± 1.0	-78.0 ± 1.1
Peak 4	1224 ± 1.0	0.7 ± 0.5
Trough 5	1371.7 ± 1.2	-60.9 ± 1.1
Peak 5	1536 ± 2.8	5.6 ± 1.3
Trough 6	1693.0 ± 3.3	-27.6 ± 1.3
Peak 6	1861 ± 4	1.2 ± 1.0
EE power spectrum		
Peak 1	137 ± 6	1.15 ± 0.07
Trough 1	197 ± 8	0.848 ± 0.034
Peak 2	397.2 ± 0.5	22.04 ± 0.14
Trough 2	525 ± 0.7	6.86 ± 0.16
Peak 3	690.8 ± 0.6	37.35 ± 0.25
Trough 3	832.8 ± 1.1	12.5 ± 0.4
Peak 4	992.1 ± 1.3	41.8 ± 0.5
Trough 4	1153.9 ± 2.7	12.3 ± 0.9
Peak 5	1296 ± 4	31.6 ± 1.0

Fig. E.2. Pearson correlation coefficient for T and E. The red curve is the theoretical best-fit to the PlanckTT, TE, EE+lowP data. The blue points are the binned *Planck* data described in this paper. Data points with negative values for *EE*, or low S/N in polarization (specifically $S/N < 1$) are omitted, which leads to an apparent bias between the theory curve and the data at the highest multipoles. The error bars here are calculated assuming a Gaussian distribution for the TT, TE, and EE power values, which should be a reasonable approximation at high ℓ. The Pearson coefficient is also normally distributed at high ℓ.

Fig. E.3. Decorrelation angle for T and E. This is the angle through which we have to rotate in $T-E$ space in order to make two uncorrelated quantities.

information. Nevertheless, there is an alternative geometrical interpretation of the $T-E$ correlation that is also worth examining. This involves considering the angle through which one has to rotate in $T-E$ space in order to decorrelate the two quantities. It is equivalent to a principal component analysis for two correlated variables, or the Jacobi rotation of a matrix with the rotation angle given by

$$\tan 2\theta = \frac{2\mathcal{D}_\ell^{TE}}{\mathcal{D}_\ell^{TT} - \mathcal{D}_\ell^{EE}}. \tag{E.2}$$

In Fig. E.3 we plot the quantity defined in Eq. (E.2). It can be regarded as the angle through which we have to rotate E and T to make the transformed quantity $C_\ell^{T'E'} = 0$. Or alternatively, it can be thought of in terms of T "leaking" into E. Since generally $C_\ell^{TT} \gg C_\ell^{EE}$, this angle is quite small. We can see from the figure that the required angle oscillates, becoming as high as $2°$ and as
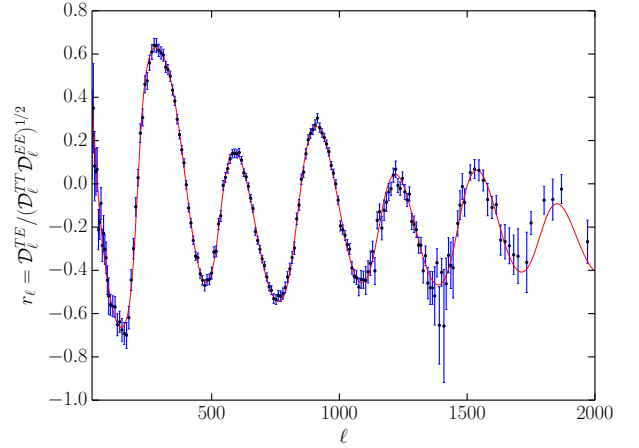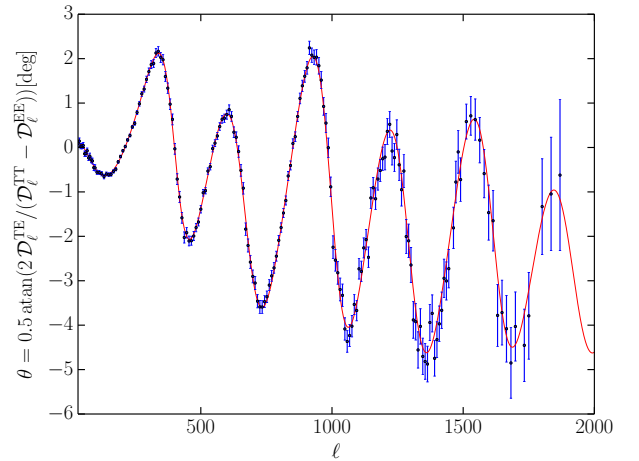
low as almost $-5°$; actually the greatest angle required is for $\ell = 1992$, where $\theta = -4°.63$. Again we see that for the best-fit cosmology the plot is mostly negative. Dimensionless quantities such as those plotted in Figs. E.2 and E.3 may have additional value in being independent of the overall normalization of the power spectra (i.e., A_s) or the calibration of the data.

E.4. Analysis of CMB maps derived by component-separation methods

The high-ℓ likelihoods considered in this paper perform component separation at the power-spectrum level, to fully exploit the signal at the smallest scales probed by *Planck* and to enable full error propagation. In this section, we describe the steps involved in the alternative approach of deriving CMB spectra and cosmological parameters from CMB maps obtained by component-separation techniques. These cleaning techniques, referred to as `Commander`, `SMICA`, `NILC`, and
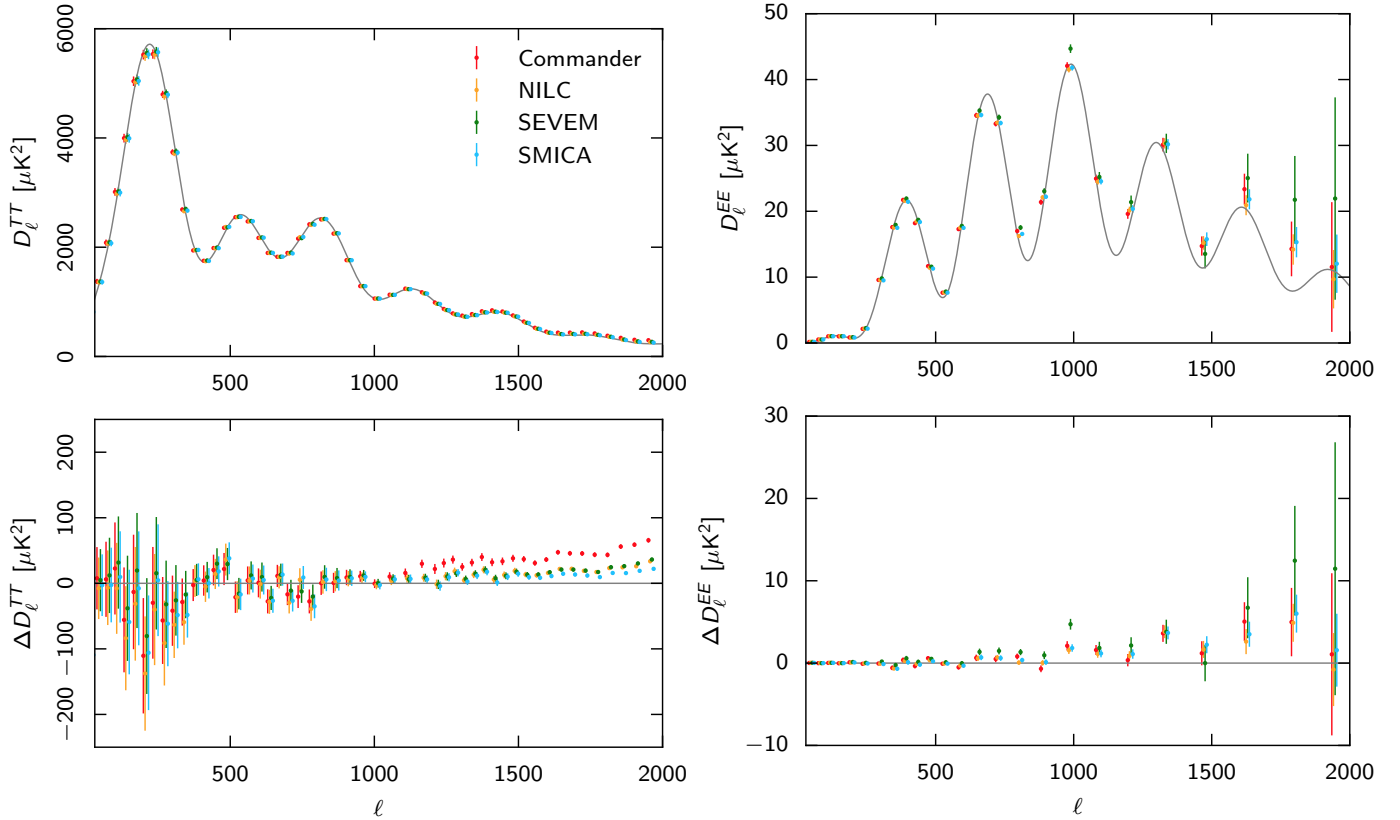
Fig. E.4. CMB *TT* (*left*) and *EE* (*right*) power spectra for each of the four foreground-cleaned CMB maps. *Upper panels* show raw band-powers; the grey lines show the best-fit ΛCDM model from the *Planck* 2015 likelihood. *Lower panels* show residual band-powers after subtracting the best-fit ΛCDM model, showing the residual extragalactic foreground contribution.

SEVEM, are described in Planck Collaboration IX (2016). The maps are weighted with the union of the confidence masks of the different component-separation methods, applying the `UT78` mask in temperature and the `UP78` mask in polarization (Planck Collaboration IX 2016), with a further cosine apodization of 10′.

In order to look at relatively small differences, we compare the angular power spectra of the four CMB maps and cosmological parameters with those from `Plik`. To estimate the power spectra we use the `XFaster` method, an approximation to the iterative, maximum likelihood, quadratic band-power estimator based on a diagonal approximation to the quadratic Fisher matrix estimator (Rocha et al. 2011, 2010). The noise bias is estimated using difference maps, as described in Planck Collaboration IX (2016).

We estimate the power spectra of the half-mission half-sum (HMHS) CMB maps. The HMHS spectra contain signal and noise. The noise bias is estimated from the half-mission half-difference (HMHD) data. The HMHD spectra contain only noise and potential systematic effects. The resulting spectra are shown in Fig. E.4. The top panels compare each of the four power spectra derived from the component-separation maps with the best-fit ΛCDM power spectrum derived from the *Planck* likelihood including multipoles up to $\ell = 2500$, with no subtraction of extragalactic foregrounds. The bottom panels show the spectrum differences between the component-separation methods and the best-fit spectrum, showing the residual extragalactic foreground contribution. We note that the `SMICA` map appears to be the least contaminated by foreground residuals at small scales. Still, it is obvious that the foreground residuals are too strong for a direct cosmological parameter analysis.

We therefore proceed, as in the baseline likelihoods, with component separation at the power-spectrum level (albeit with much smaller non-CMB contributions). One must nevertheless obtain templates for these residual foregrounds after component separation. To that effect, we propagate the simulated full-sky FFP8 foreground maps through the respective pipelines and estimate the resulting power spectra normalized to some pivotal multipole, under the hypothesis that the simulated sky is close enough to the real one for these residual templates to be accurate. We have indications that this is only marginally the case, and have used a dedicated foreground model for the validation of the `Plik` likelihood on the FFP8 simulations. This can be the source of the variation under $\ell_{\max}$ of some of the parameters obtained in this test.

We then estimate cosmological parameters using a Gaussian correlated likelihood derived from these `XFaster` bandpowers, coupled to `CosmoMC` (Lewis & Bridle 2002). Specifically, we include multipoles between $\ell_{\min} = 50$ and $\ell_{\max}$, where $\ell_{\max} = 1000$, 1500, or 2000 for temperature, and 1000 or 1500 for polarization. We explore the base six-parameter ΛCDM model, and, since low-ℓ data are not used in the likelihood, impose an informative Gaussian prior of $\tau = 0.070 \pm 0.006$.

The resulting spectra are shown in Fig. E.5, and the corresponding cosmological parameters are summarized in Fig. E.6 for both *TT* (filled symbols) and *EE* (unfilled symbols). The comparison of Fig. E.4 with Fig. E.5 enables assessment of the efficiency of this extra cleaning step, which brings the component-separation-based spectra into much greater agreement with those derived from the high-ℓ likelihood. To be more quantitative, we now turn to the corresponding cosmological parameters in Fig. E.6. Starting with the temperature cases, we
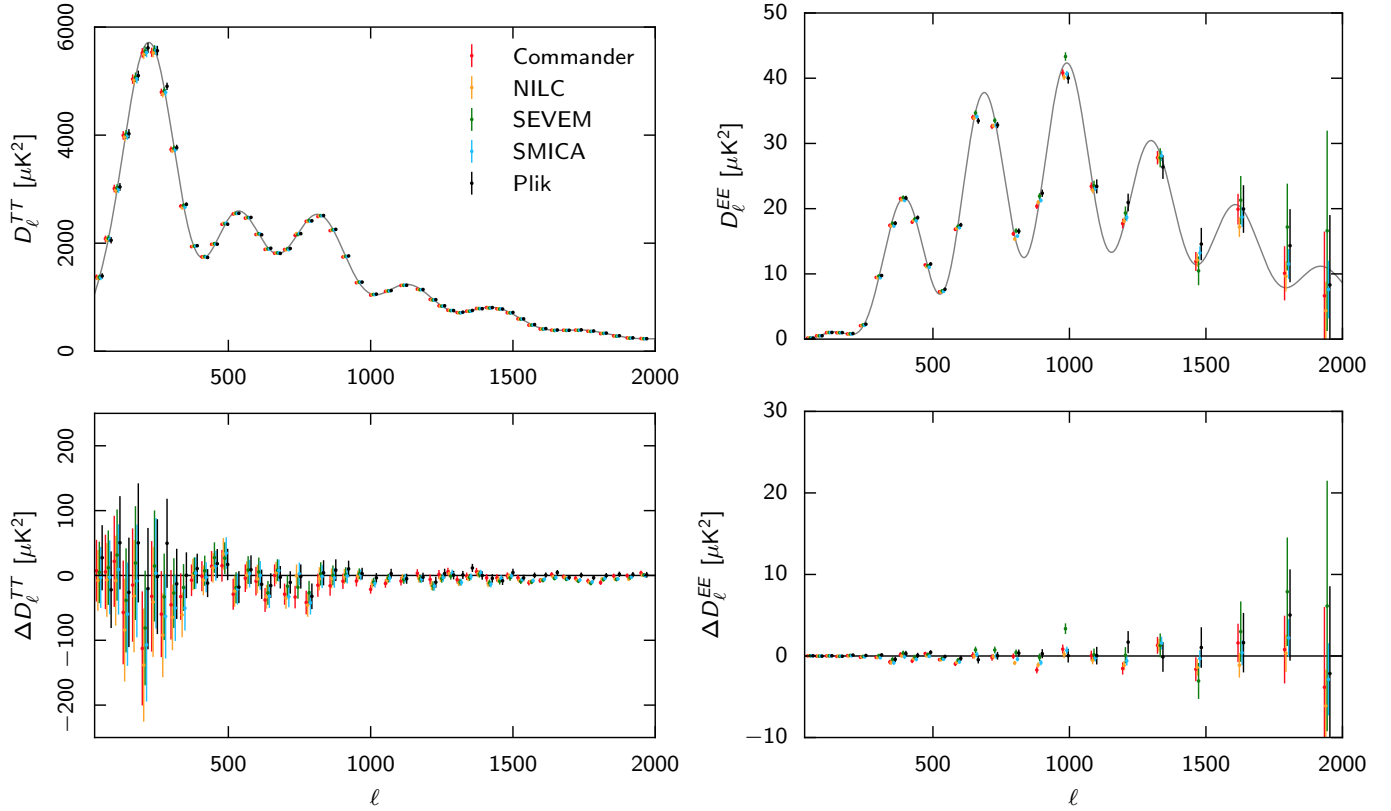
Fig. E.5. CMB *TT* (*left*) and *EE* (*right*) power spectra for each of the four foreground-cleaned CMB maps, as in Fig. E.4, after an extra cleaning out the extragalactic residuals at the spectrum level, along with the `Plik` results for comparison. *Upper panels* show band powers; the grey lines show the best-fit ΛCDM model from the *Planck* 2015 likelihood. *Lower panels* show residual band powers after subtracting the best-fit ΛCDM model.

Table E.3. Comparison between parameters based on various CMB temperature maps.

Parameter	Commander	NILC	SEVEM	SMICA
$\Omega_b h^2$	0.02207 ± 0.00025	0.02172 ± 0.00023	0.02185 ± 0.00023	0.02184 ± 0.00024
$\Omega_c h^2$	0.1178 ± 0.0024	0.1205 ± 0.0023	0.1214 ± 0.0023	0.1202 ± 0.0023
$100\theta_{MC}$	1.041 ± 0.0005	1.041 ± 0.0005	1.040 ± 0.0005	1.041 ± 0.0005
τ	0.075 ± 0.019	0.069 ± 0.018	0.069 ± 0.018	0.068 ± 0.019
$10^9 A_s e^{-2\tau}$	1.862 ± 0.011	1.868 ± 0.011	1.882 ± 0.012	1.866 ± 0.015
n_s	0.971 ± 0.007	0.964 ± 0.007	0.958 ± 0.007	0.960 ± 0.007
Ω_m	0.304 ± 0.014	0.322 ± 0.014	0.327 ± 0.014	0.321 ± 0.014
H_0	68.0 ± 1.1	66.6 ± 1.0	66.4 ± 1.0	66.8 ± 1.0

Notes. These constraints are obtained from *TT* in the range $50 < \ell < 2000$ combined with a prior of $\tau = 0.07 \pm 0.02$, allowing for residual foreground contributions (see text).

first observe reasonable overall internal agreement between the four component-separation methods, with almost all differences smaller than 1σ within each multipole band. Second, we also observe acceptable agreement with the best-fit *Planck* 2015 ΛCDM model derived from the likelihood, since most of the differences are within 1σ, at least for $\ell_{max} = 2000$. One notable exception to this agreement is the power spectrum amplitude, $A_s e^{-2\tau}$, which is systematically low by about 2σ for $\ell_{max} = 1000$ for all methods. We further note that otherwise all approaches find similar shifts of some parameters (ω_c) between $\ell_{max} = 1000$ and $\ell_{max} = 1400$. Table E.3 provides the numerical *TT* constraints at $\ell < 2000$.

One should keep in mind that the CMB cleaned maps used here have been obtained with a different objective than providing the best/safest estimation of the cosmological parameters.

Rather, they are meant to maximize the sky coverage, in particular for non-Gaussianity analyses; this entails a very different set of trade-offs. To reach this goal, the component separation methods use the LFI channels to clean regions close to the Galactic plane, where foreground sources other than dust need to be accounted for. The efficiency of the cleaning is dominated by the residual in the Galactic plane, where contamination is strong, possibly at the expense of regions far from the plane, where the residual compared to the CMB can then be higher than in a dedicated solution. The separation methods also require the use of a common mask across all frequencies, contrary to our spectrum-based approach where we can tailor the masks per frequency in order to minimize foreground contamination as well as the size and number of the point source holes. In the end, the
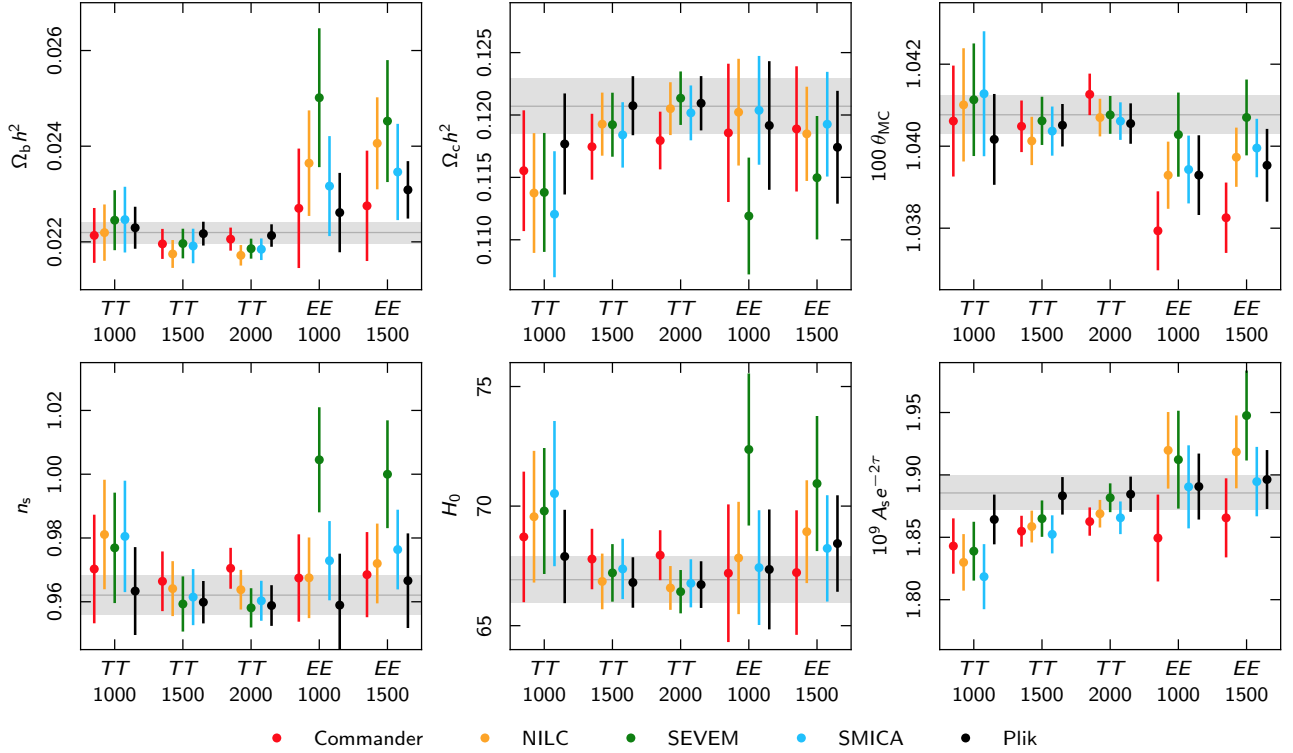
Fig. E.6. Comparison of cosmological parameters estimated from the *TT* and *EE* spectra computed from the CMB maps and further accounting for foreground residuals through method-tailored templates derived from the FFP8 simulations. The sets of points correspond to various values of ℓ_{max}, and lets the four methods and Plik to be compared when the same scales are used. For comparison, we also show the corresponding parameters obtained with the *Planck* 2015 likelihood including multipoles up to $\ell_{max} = 2500$ as the horizontal line surrounded by a grey band giving the uncertainties.

channel weights in the component-separated CMB map is different from the weighting in Plik, with 143 GHz dominating the CMB map at the first peak, whereas, thanks to the use of a higher sky fraction at this frequency, the 100×100 dominates the Plik spectrum (see Fig. 16). Further, the use of a larger point-source mask increases the excess scatter seen in Fig. C.1.4, which is not taken into account in XFaster. The propagation of instrumental uncertainties, like calibration and beam errors, are also handled very differently, if at all. Given all those caveats, it is remarkable to nevertheless find such a general agreement between the different component-separated maps created using the XFaster likelihood, between themselves and with the Plik likelihood.

The *EE*-based results are generally more discrepant, apart from SMICA. For instance Commander gives a quite significantly different value for θ_{MC}, NILC differs in Ω_b, and SEVEM in Ω_b, n_s, H_0, and $A_s e^{-2\tau}$.
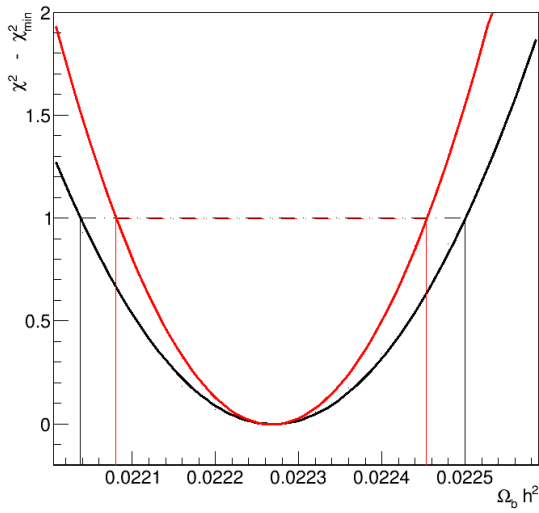
E.5. Profile likelihood

We have additionally made a comparison of the cosmological parameters determined using a "profile likelihood" method as in Planck Collaboration Int. XVI (2014), along with an independent Boltzmann code, CLASS (Lesgourgues 2011). The results of this analysis for the *Planck* lowTEB likelihood on the ΛCDM model are shown in the first column of Table E.4, which summarizes the values of the parameters and their errors at 68% CL for both Plik and CamSpec. This table should be compared with

the first column of Table 4 of Planck Collaboration XIII (2016), which shows the *Planck* 2015 baseline MCMC parameter results; there is excellent agreement with the profile-likelihood fits, the small difference observed for the mean value of θ_{MC} being attributed to a slightly different definition of this parameter in CLASS with respect to CAMB.

We have also estimated the contribution of the fit of the nuisance (foreground and instrumental) parameters to the error on the cosmological parameters; this analysis permits us to assess how much our lack of knowledge of the nuisance parameters impacts the cosmological error budget. First, a profile-likelihood fit is performed to estimate the full error on each parameter, giving, for instance, the black curve of Fig. E.7; this corresponds to Col. 3 in Table E.4. In a second step, another profile likelihood is built, fixing the nuisance parameters to their best-fit values; from this we obtain the "sensitivity-related" uncertainty (the red curve of Fig. E.7). This corresponds to the ultimate error one would obtain if we knew the nuisance parameters perfectly (and they had these best-fit values). Finally the "foreground and instrumental" error is deduced by quadratically subtracting the "sensitivity-related" uncertainty from the total error; this procedure has been used, for instance, in Aad (2014). The results are shown in Cols. 4 and 5 of Table E.4 for Plik and CamSpec, and are very similar for both likelihoods. Ω_b and n_s are more strongly impacted by the fit of the nuisance parameters, while $100\theta_{MC}$ is less so, which is consistent with the correlation matrix of the parameters.

Table E.4. Profile likelihood estimates of cosmological parameters within the ΛCDM model for `Plik` and `CamSpec`.

		Error [68% CL]		
Parameter	Estimate	Full	Sensitivity-related	Foreground and instrumental
`Plik` parameters				
$\Omega_b h^2$	0.02227	0.00023	0.00019	0.00014
$\Omega_c h^2$	0.1198	0.0022	0.0021	0.0007
$100\theta_{MC}$	1.04184	0.00044	0.00044	0.00003
τ	0.082	0.020	0.018	0.009
$\ln(10^{10} A_s)$	3.098	0.037	0.034	0.014
n_s	0.9663	0.0063	0.0051	0.0036
`CamSpec` parameters				
$\Omega_b h^2$	0.02229	0.00023	0.00018	0.00014
$\Omega_c h^2$	0.1194	0.0022	0.0021	0.0006
$100\theta_{MC}$	1.04102	0.00047	0.00047	0.00009
τ	0.080	0.020	0.018	0.009
$\ln(10^{10} A_s)$	3.091	0.037	0.034	0.016
n_s	0.9685	0.0062	0.0051	0.0036

Fig. E.7. Likelihood (specifically $\chi^2 - \chi^2_{min}$) profile of $\Omega_b h^2$ within the ΛCDM model for `Plik`. In black is shown the full profile-likelihood fit (to derive the full error of Table E.4) and in red the profile-likelihood fit when the nuisance are fixed to the values obtained when maximizing the likelihood function (to derive the "sensitivity related" error of Table E.4).