A&A 594, A11 (2016)
DOI: 10.1051/0004-6361/201526926

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], 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)

Received 9 July 2015 / Accepted 18 May 2016

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, $n_{\rm s}$, 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

---

[*] Corresponding author: F. R. Bouchet, e-mail: `bouchet@iap.fr`

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 $\mu\mathrm{K}^2$), 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

---

[1] *Planck* (`http://www.esa.int/Planck`) is a project of the European Space Agency (ESA) with instruments provided by two scientific consortia funded by ESA member states and led by Principal Investigators from France and Italy, telescope reflectors provided through a collaboration between ESA and a scientific consortium led and funded by Denmark, and additional contributions from NASA (USA).

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_ℓ^{TE}+C_ℓ^{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 $C_\mathrm{b}^{TT} + C_\mathrm{b}^{TE} + C_\mathrm{b}^{EE}$, 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 $\ell \leq \ell_\mathrm{cut} = 29$ in the likelihood, we adopt a HEALPix resolution of $N_\mathrm{side} = 16$ which has 3072 pixels (of area 13.6 deg^2) per map; this accommodates multipoles up to $\ell_\mathrm{max} = 3N_\mathrm{side} - 1 = 47$, and, considering separate maps of T, Q, and U, corresponds to a maximum of $N_\mathrm{pix} = 3 \times 3072 = 9216$ 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 $X = \{\mathrm{T, E, B}\}$ 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

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

---

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

Here we restrict the spectra to $XY = \{\mathrm{TT}, \mathrm{EE}, \mathrm{BB}, \mathrm{TE}\}$, with $N_{\mathrm{side}} = 16$ 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

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

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 $O(N_{\mathrm{pix}}^3)$. 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 $\ell_{\mathrm{cut}} = 29$ with this formalism, requiring maps with $N_{\mathrm{side}} = 16$. Multipoles between $\ell_{\mathrm{cut}} + 1$ 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

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

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 $N_{\mathrm{pix}} \times N_{\mathrm{pix}}$ 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` $N_{\mathrm{side}} = 256$ pixel resolution to $N_{\mathrm{side}} = 16$ 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 ±25 μK. 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

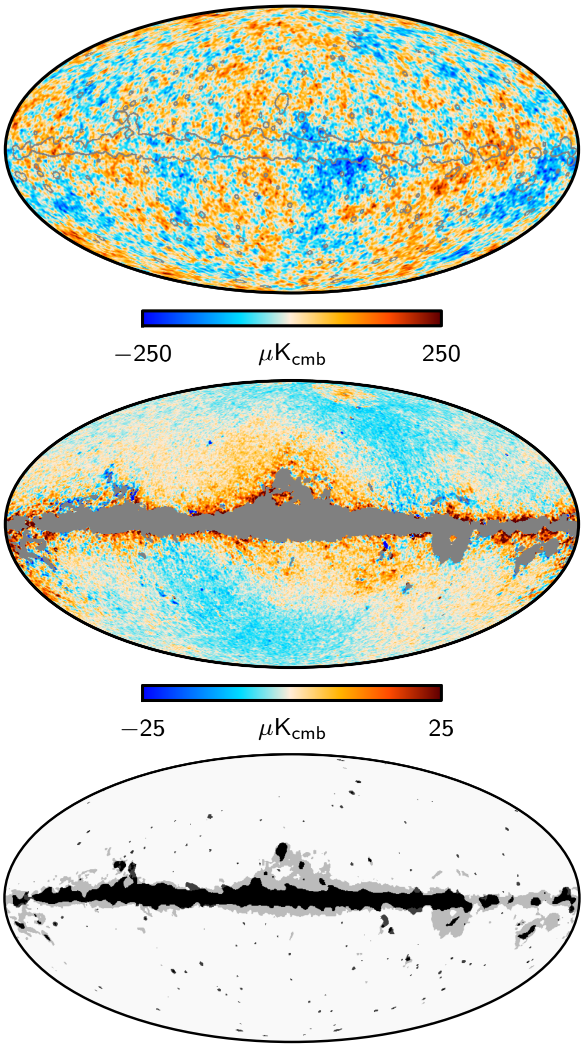

Fig. 1. *Top*: `Commander` CMB temperature map derived from the *Planck* 2015, 9-year WMAP, and 408 MHz Haslam et al. observations, as described in Planck Collaboration X (2016). The gray boundary indicates the 2015 likelihood temperature mask, covering a total of 7% of the sky. The masked area has been 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

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

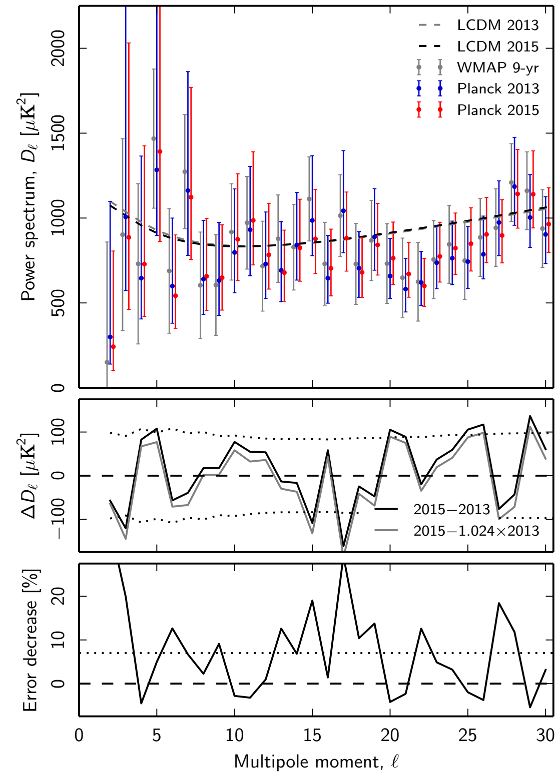

Fig. 2. *Top*: comparison of the *Planck* 2013 (blue points) and 2015 (red points) posterior-maximum low-ℓ temperature power spectra, as derived with `Commander`. Error bars indicate asymmetric marginal posterior 68% confidence regions. For reference, we also show the final 9-year WMAP temperature spectrum 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

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

---

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

where σ_ℓ^{u} and σ_ℓ^{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 $N_{\mathrm{side}} = 16$ 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 $f_{\mathrm{sky}} = 0.46$ 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

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

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 latter can be estimated by minimizing the quantity

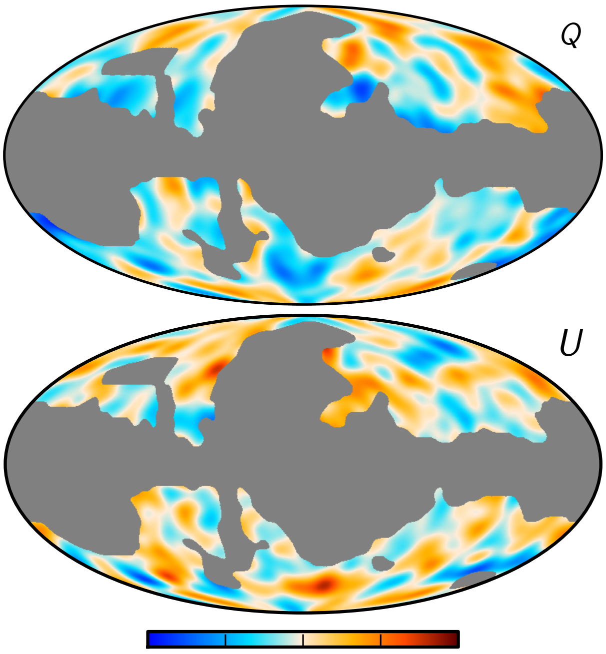

Fig. 3. Foreground-cleaned, 70 GHz Q (*top*) and U (*bottom*) maps used for the low-ℓ polarization part of the likelihood. Each of the maps covers 46% of the sky.

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

where

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

Here 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:

$$\mathsf{C}_{\mathrm{S+N}} = \mathsf{S}(C_\ell) + \mathsf{N}_{70}, \tag{10}$$

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) $n_{\mathrm{synch}} = -3.16 \pm 0.40$ and $n_{\mathrm{dust}} = 1.50 \pm 0.16$, 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.

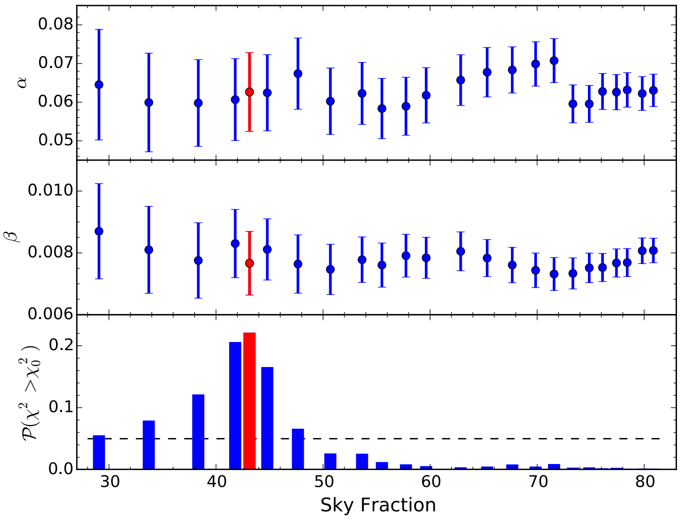

Fig. 4. *Upper panels*: estimated best-fit scaling coefficients for synchrotron (α) and dust (β), for several masks, whose sky fractions are displayed along the bottom horizontal axis (see text). *Lower panel*: the probability to exceed, $\mathcal{P}(\chi^2 > \chi_0^2)$. The red symbols identify the 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 $n_{\rm synch} = -3.2$ and $n_{\rm dust} = 1.6$, 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 $f_{\rm sky} = 43\%$ processing mask is chosen as the tightest mask (i.e., the one with the greatest $f_{\rm sky}$) satisfying the requirement $\mathcal{P} > 5\%$ (see Fig. 4). We use a slightly smaller mask ($f_{\rm sky} = 46\%$) 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

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

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

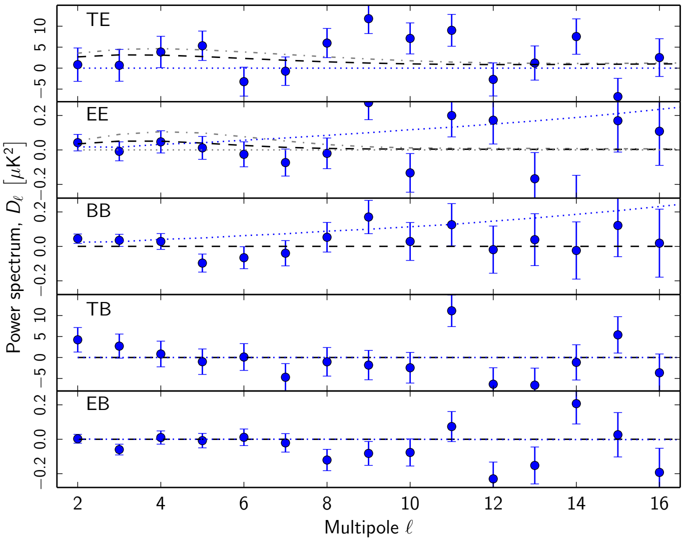

Fig. 5. Polarized QML spectra from foreground-cleaned maps. Shown are the 2013 *Planck* best-fit model ($\tau = 0.089$, dot-dashed) and the 2015 model ($\tau = 0.067$, dashed), as well as the 70 GHz noise bias computed from Eq. (11) (blue dotted).

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 $A_{\rm s}$, and the tensor-to-scalar ratio r. Results are shown in Table 2 for the combinations $(\tau, A_{\rm s})$ and $(\tau, A_{\rm s}, r)$.

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 $A_{\rm s}{\rm e}^{-2\tau}$, 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 $A_{\rm s}$. On the other hand, large-scale polarization is sensitive to the combination $A_{\rm s}\tau^2$. Thus, neither low-ℓ temperature nor polarization can separately constrain τ and $A_{\rm s}$. Combining temperature and

---

[5] The models considered have been derived by fixing all parameters except τ and $A_{\rm s}$ 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
$\log[10^{10}A_\mathrm{s}]$	2.952 ± 0.055	$2.788^{+0.19}_{-0.09}$
r	0	[0, 0.90]
z_re	$8.9^{+2.5}_{-2.0}$	$8.5^{+2.5}_{-2.1}$
$10^9 A_\mathrm{s}$	$1.92^{+0.10}_{-0.12}$	$1.64^{+0.29}_{-0.17}$
$A_\mathrm{s}\mathrm{e}^{-2\tau}$	$1.675^{+0.082}_{-0.093}$	$1.45^{+0.24}_{-0.14}$

Notes. For the centre column the set of parameters (τ, A_s) was sampled, while it was the set (τ, A_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_s, which is almost completely unconstrained without temperature information, and not to the lack of EE signal. By assuming a sharp prior $10^9 A_\mathrm{s}\mathrm{e}^{-2\tau} = 1.88$, 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 bound does not differ much from having A_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_s, and r for all four likelihood versions described above, finding consistent results.

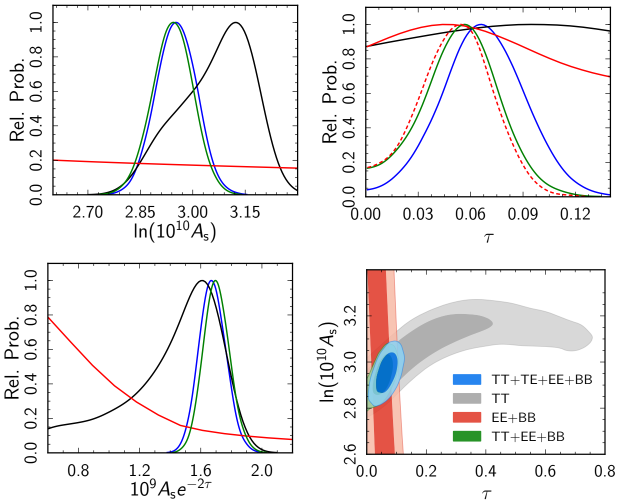

Fig. 6. Likelihoods for parameters from low-ℓ data. *Panels 1–3*: one-dimensional posteriors for $\log[10^{10}A_\mathrm{s}]$, τ, and $A_\mathrm{s}\mathrm{e}^{-2\tau}$ for the several sub-blocks of the likelihood, for cases 1 (blue), 2 (black), 3 (red), and 4 (green) – see text for definitions; dashed red is the same as case 3 but imposes a sharp prior $10^9 A_\mathrm{s}\mathrm{e}^{-2\tau} = 1.88$. *Panel 4*: two-dimensional posterior for $\log[10^{10}A_\mathrm{s}]$ and τ for the same data combinations; shading indicates the 68% and 95% confidence regions.

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:

$$\chi^2_\mathrm{h} = \sum_{\ell=2}^{\ell_\mathrm{max}} (C_\ell - C_\ell^\mathrm{th})\, \mathsf{M}^{-1}_{\ell\ell'}\, (C_\ell - C_\ell^\mathrm{th}), \tag{12}$$

where $\mathsf{M}_{\ell\ell'} = \langle (C_\ell - C_\ell^\mathrm{th})(C_{\ell'} - C_{\ell'}^\mathrm{th}) \rangle$, C_ℓ^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 χ^2_h, considering both $\ell_\mathrm{max} = 12$ (shown in Fig. 7, along with the corresponding values obtained from the data) and $\ell_\mathrm{max} = 30$, for each of the six CMB polarized spectra. We report in Table 3 the empirical probability of observing a value of χ^2_h 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

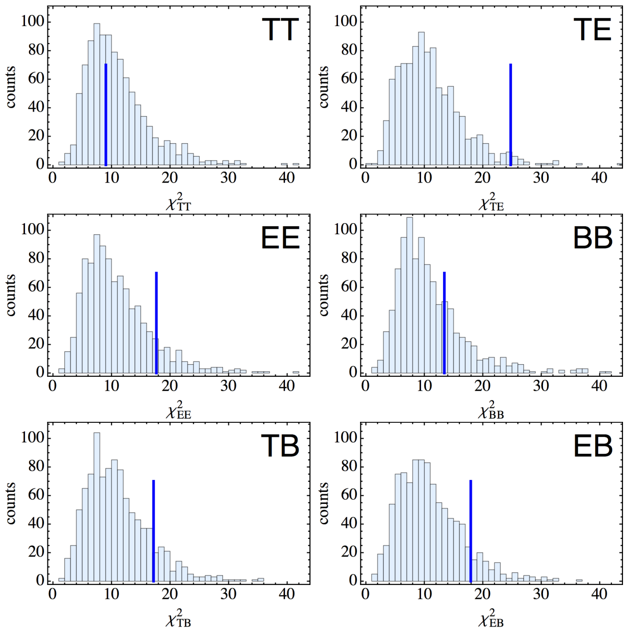

Fig. 7. Empirical distribution of χ^2_{h} derived from 1000 simulations, for the case $\ell_{\max} = 12$ (see text). Vertical bars reindicate the observed values.

Table 3. Empirical probability of observing a value of χ^2_h 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.

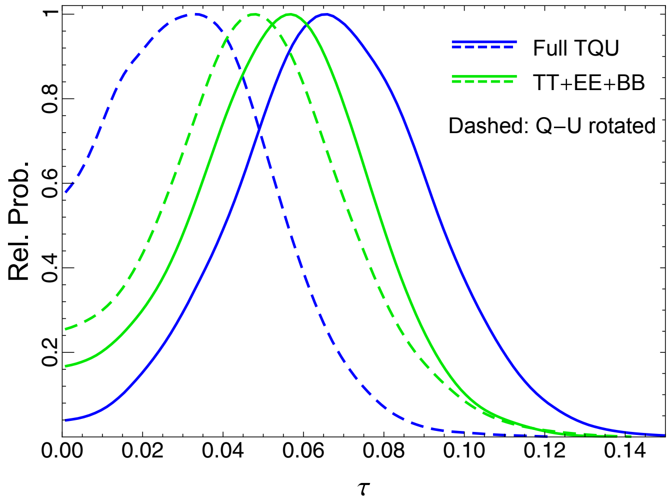

Fig. 8. Posterior for τ for both rotated and unrotated likelihoods. The definition and colour convention of the datasets shown are the same as in the previous section (see Fig. 6), while solid and dashed lines distinguish the unrotated and rotated likelihood, respectively.

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 $N_{\mathrm{side}} = 16$ 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 $\log[10^{10}A_{\mathrm{s}}]$ 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.

	`Cosmomc` best-fit			`Cosmomc` mean			Standard deviation	
Parameter	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
$\log[10^{10}A_\mathrm{s}]$	3.035 ± 0.002	0.059	−9.4%	3.036 ± 0.002	0.055	−8.0%	0.0535 ± 0.0001	0.0032
$\log[10^{10}A_\mathrm{s}^*]$	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 $\tau_\mathrm{input} = 0.0650$ and $\log[10^{10}A_\mathrm{s}]_\mathrm{input} = 3.040$.

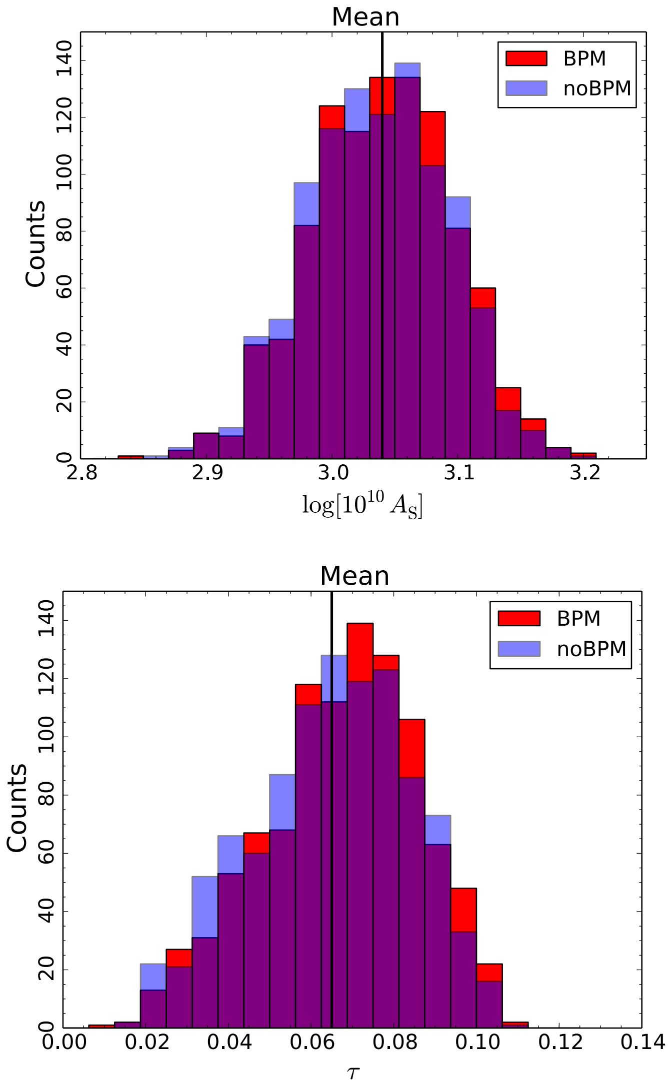

Fig. 9. Empirical distribution of the mean estimated values for $\log[10^{10}A_\mathrm{s}]$ (*top*) and τ (*bottom*), derived from 1000 FFP8 simulations (see text). For each simulation, we perform a full end-to-end run, including foreground cleaning and parameter estimation. Blue bars refer to simulations 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 $N_\mathrm{side} = 8$ in polarization, we use $N_\mathrm{side} = 16$ 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.

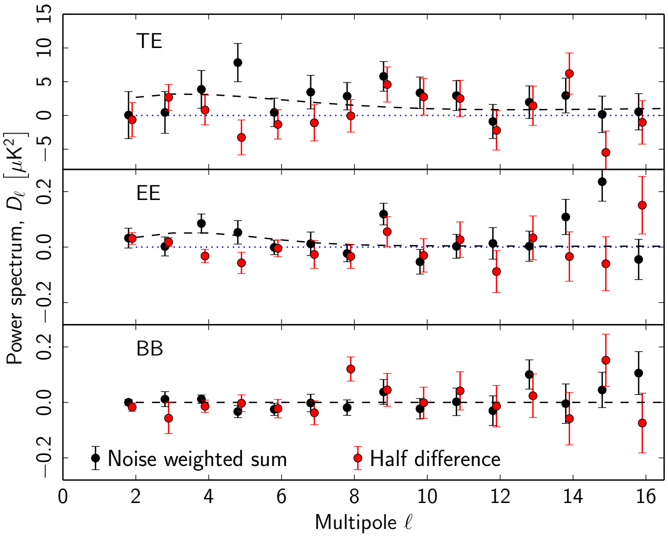

Fig. 10. `BolPol` spectra for the noise-weighted sum (black) and half-difference (red) WMAP and *Planck* combinations. The temperature map employed is always the `Commander` map described in Sect. 2.2 above. The fiducial model shown has $\tau = 0.065$.

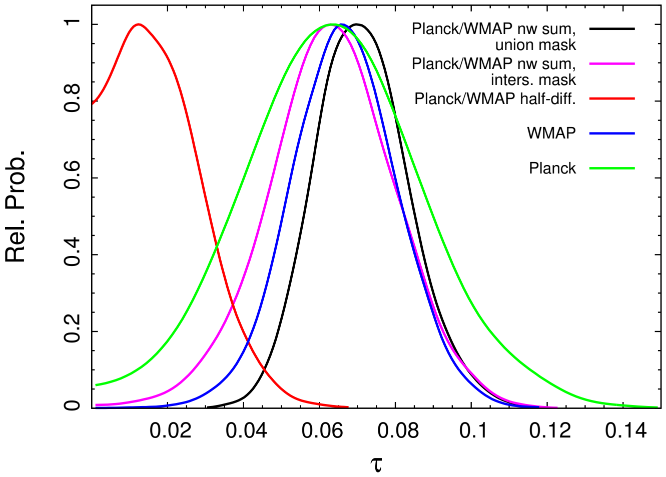

Fig. 11. Posterior probabilities for τ from the WMAP (cleaned with *Planck* 353 GHz as a dust template) and *Planck* combinations listed in the legend. Results are presented for the noise-weighted sum both in the union and the intersection of the two analysis masks. The half-difference map is consistent 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	$0.064^{+0.022}_{-0.023}$	$0.067^{+0.013}_{-0.013}$	$0.071^{+0.011}_{-0.013}$
$z_{\rm re}$	$8.5^{+2.5}_{-2.1}$	$8.9^{+1.3}_{-1.3}$	$9.3^{+1.1}_{-1.1}$
$\log[10^{10}A_{\rm s}]$	$2.79^{+0.19}_{-0.09}$	$2.87^{+0.11}_{-0.06}$	$2.88^{+0.10}_{-0.06}$
r	[0, 0.90]	[0, 0.52]	[0, 0.48]
$A_{\rm s}{\rm e}^{-2\tau}$	$1.45^{+0.24}_{-0.14}$	$1.55^{+0.16}_{-0.10}$	$1.55^{+0.14}_{-0.11}$

Notes. The temperature map used is always *Planck* `Commander`. Only $\log[10^{10}A_{\rm s}]$, 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.

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 ($f_{\rm sky} = 45.3\%$), we evaluate half-sum and half-difference Q and U maps. We then compute the quantity $\chi^2_{\rm sd} = \boldsymbol{m}^{\rm T}{\rm N}^{-1}\boldsymbol{m}$ 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 $\chi^2_{\rm sd}$ is χ^2 distributed with 2786 degrees of freedom we find a PTE($\chi^2 > \chi^2_{\rm sd}$) $= 1.3 \times 10^{-5}$ (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 ($f_{\rm sky} = 73.8\%$). 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 $\ell_{\rm max} = 12$ 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 $\log[10^{10}A_{\rm s}]$, 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 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 $z_{\rm re}$ (mean redshift of reionization) and $A_{\rm s}{\rm e}^{-2\tau}$. 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:

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

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,

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

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

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

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

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

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,

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

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.

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.

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

---

[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 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, 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.

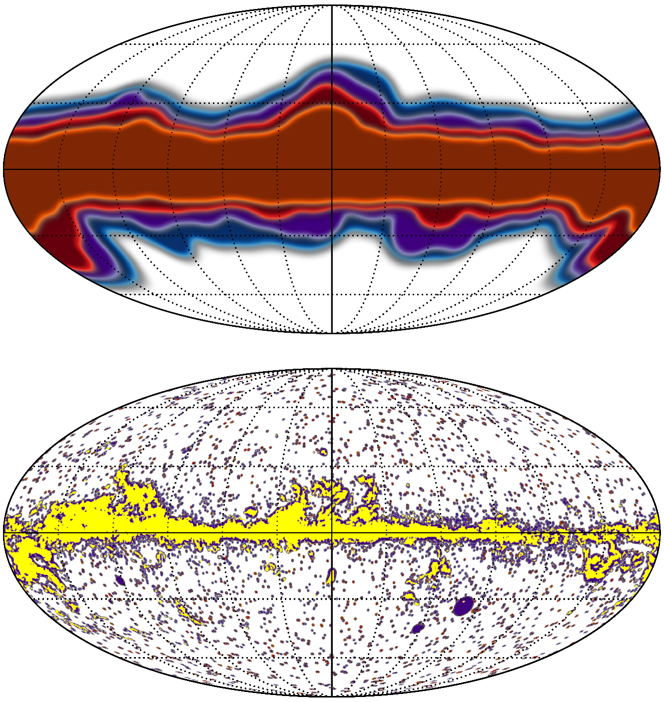

Fig. 12. *Top*: apodized Galactic masks: G41 (blue), G50 (purple), G60 (red), and G70 (orange); these are identical to the polarization masks P41 (used at 217 GHz), P50 (143 GHz), P70 (100 GHz). *Bottom*: extragalactic-object masks for 217 GHz (purple), 143 GHz (red), and 100 GHz (orange); the CO mask is shown in yellow.

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 ($f_{\rm sky} = 49\%, 31\%$, and 31% at, respectively, 100, 143, and 217 GHz, to be compared to $f_{\rm sky} = 66\%, 57\%$, 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

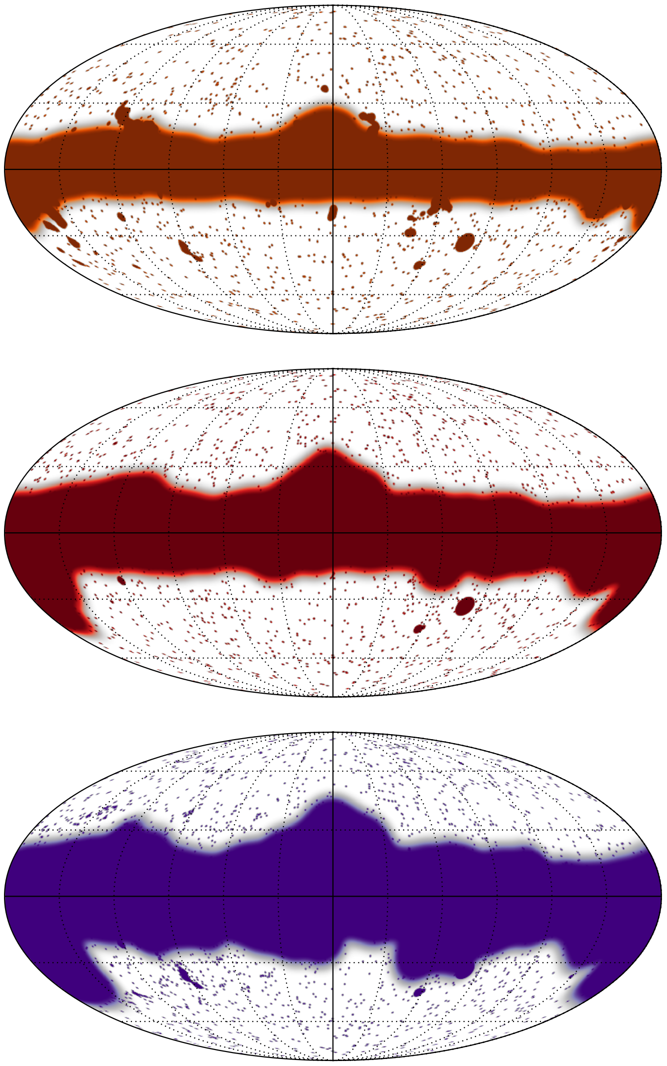

Fig. 13. *Top to bottom*: temperature masks for 100 GHz (T66), 143 GHz (T57) and 217 GHz (T47). The colour scheme is the same as in Fig. 12.

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

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

(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.

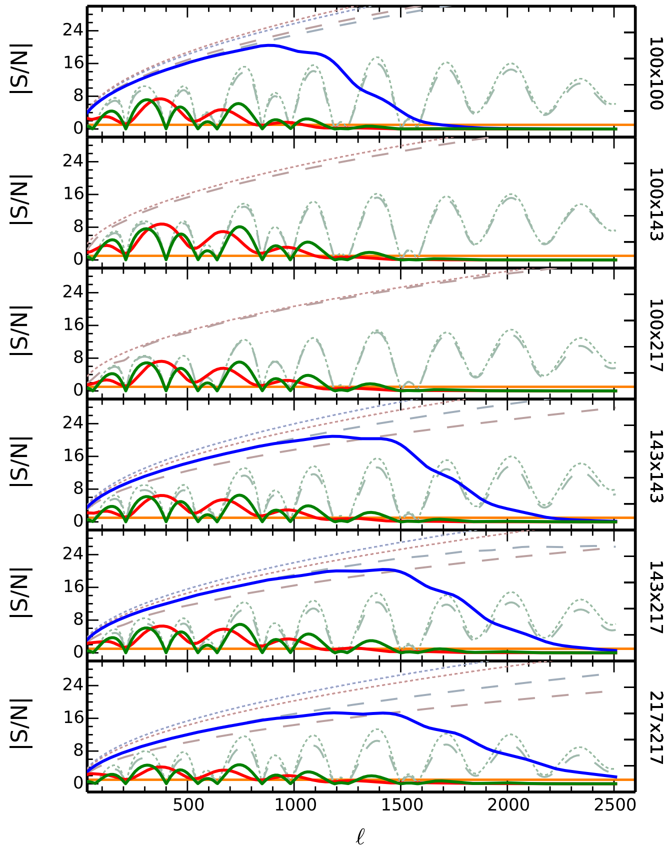

Fig. 14. Unbinned S/N per frequency for TT (solid blue, for those detector combinations used in the estimate of the TT spectrum), EE (solid red), and TE (solid green). The horizontal orange line corresponds to $S/N = 1$. The dashed lines indicate the S/N in a cosmic-variance-limited case, obtained by 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 $\ell \approx 250$–450 and additionally in the range $\ell \approx 650$–700 only for the 100×143 GHz and the 143×217 GHz power spectra.

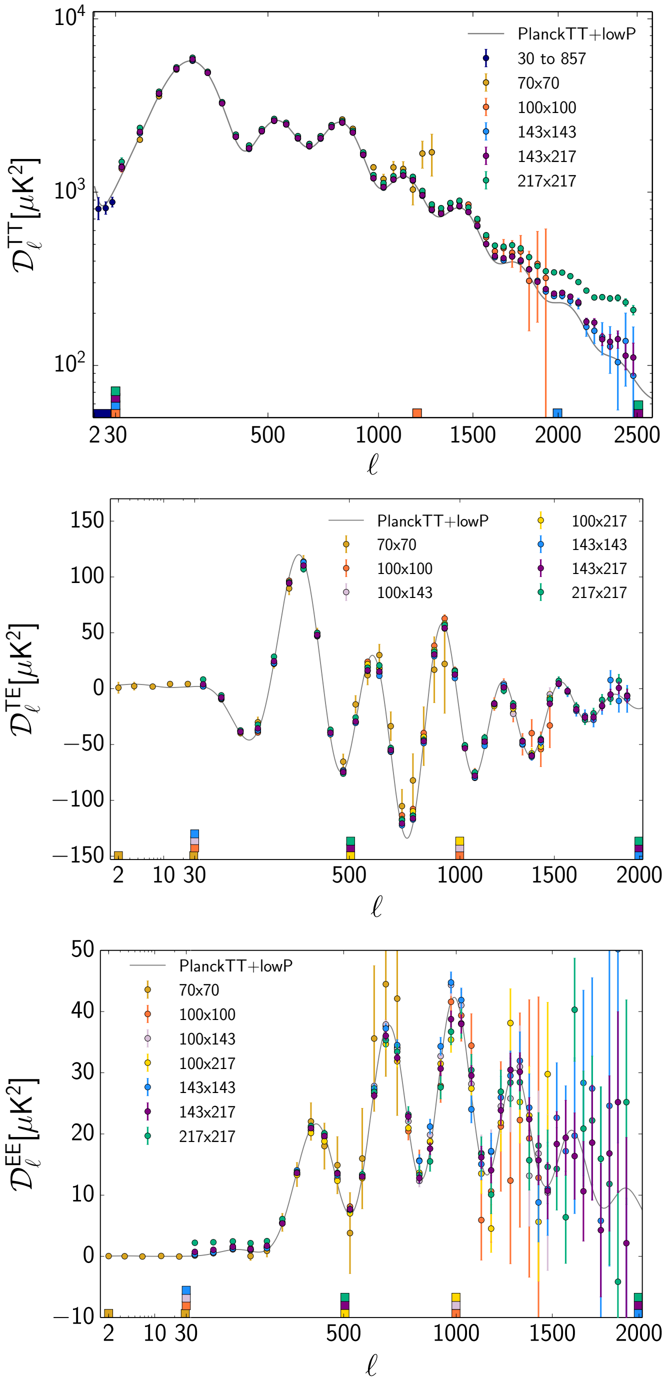

Fig. 15. *Planck* power spectra (not yet corrected for foregrounds) and data selection. The coloured tick marks indicate the ℓ-range of the cross-spectra included in the *Planck* likelihood. Although not used in the high-ℓ likelihood, the 70 GHz spectra at $\ell > 29$ illustrate the consistency of the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

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.

---

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

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

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 $30 \le \ell \le 99$, $\Delta\ell = 9$ for $100 \le \ell \le 1503$, $\Delta\ell = 17$ for $1504 \le \ell \le 2013$, and $\Delta\ell = 33$ for $2014 \le \ell \le 2508$, with a weighting of the C_ℓ proportional to $\ell(\ell+1)$ over the bin widths,

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

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.

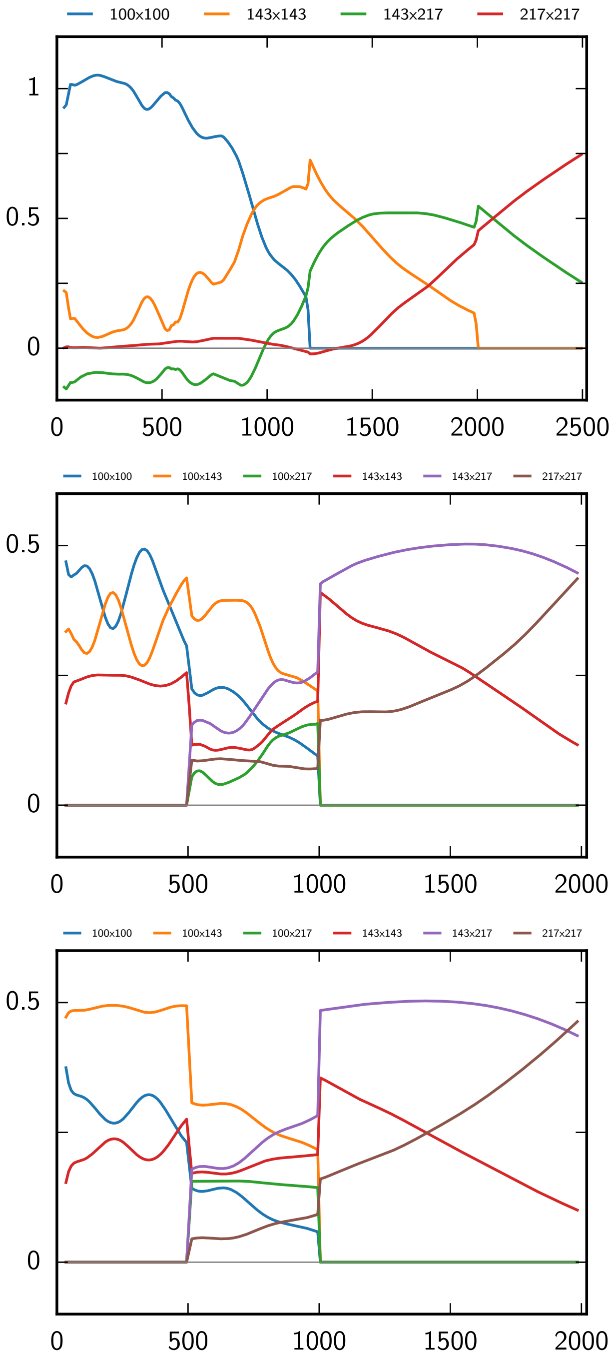

Fig. 16. The relative weights of each frequency cross-spectrum in the *TT* (*top*), *TE* (*middle*) and *EE* (*bottom*) best-fit solution. Sharp jumps are due to the multipole selection. Weights are normalized to sum to one.

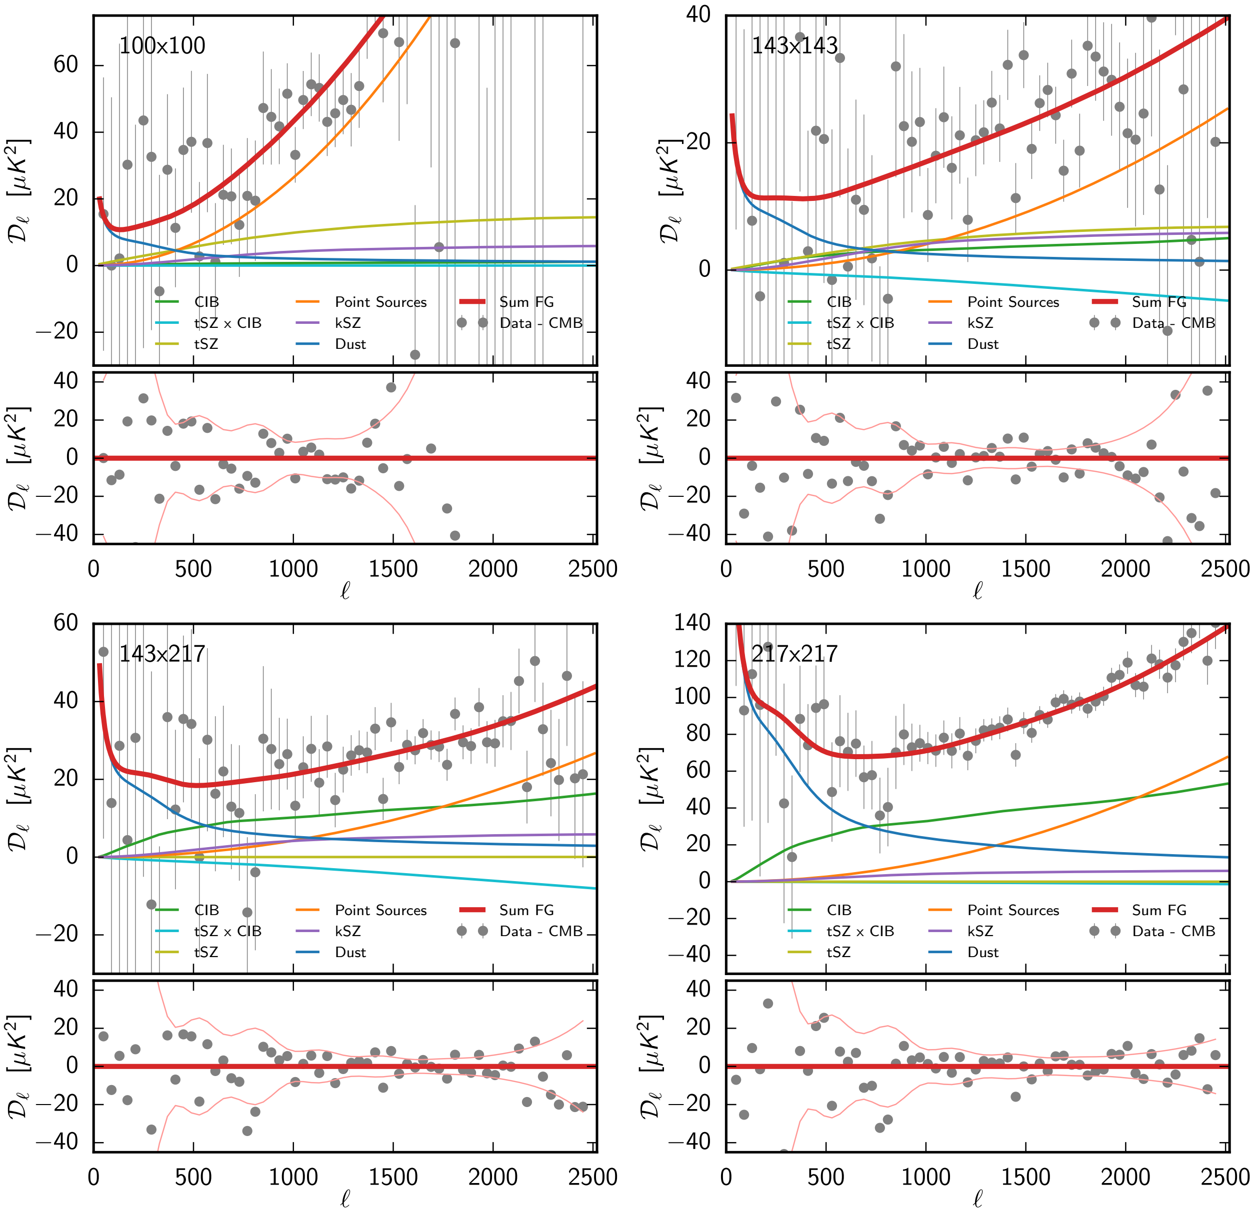

Fig. 17. Best foreground model in each of the cross-spectra used for the temperature high-ℓ likelihood. The data corrected by the best theoretical CMB C_ℓ are shown in grey. The *bottom panel* of each plot shows the residual after foreground correction. The pink line shows the 1σ value from the diagonal 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

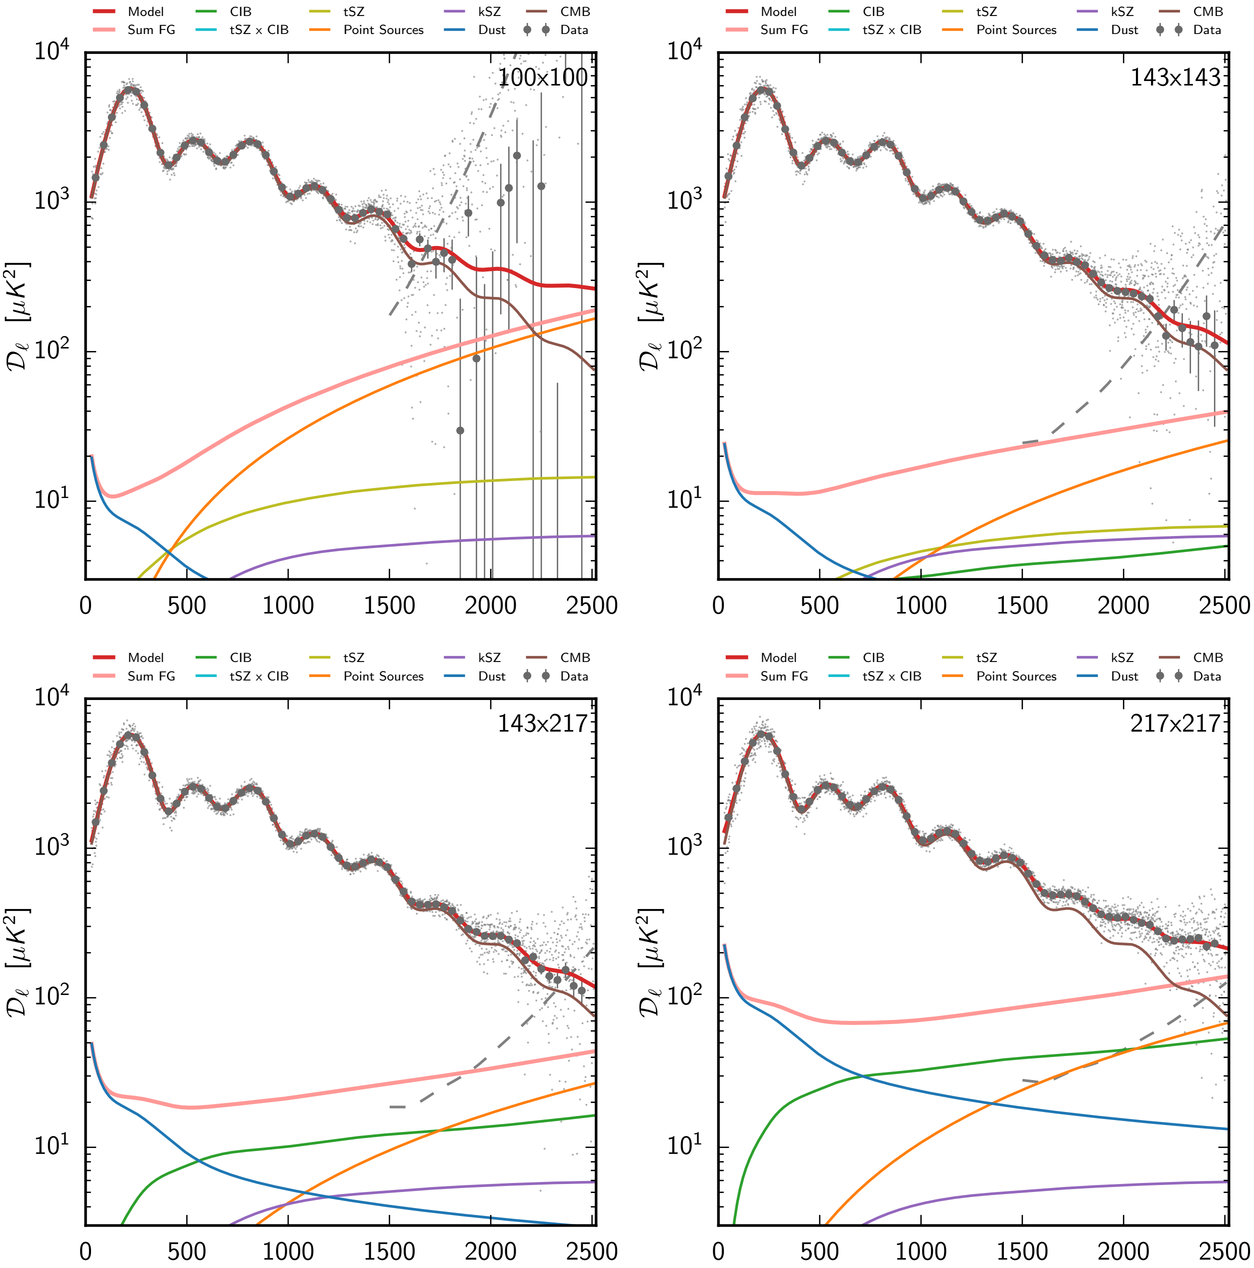

Fig. 18. Best model (CMB and foreground) in each of the cross-spectra used for the temperature high-ℓ likelihood. The small light grey points show the unbinned data point, and the dashed grey line show the square root of the noise contribution to the diagonal of the unbinned covariance matrix.

Y at frequency ν' as

$$\left(C^{XY,\mathrm{dust}}_{\nu\times\nu'}\right)_\ell = A^{XY,\mathrm{dust}}_{\nu\times\nu'} \times C^{XY,\mathrm{dust}}_\ell, \tag{22}$$

where XY is one of TT, EE, or TE, and $C^{XY,\mathrm{dust}}_\ell$ is the template dust power spectrum, with corresponding amplitude $A^{XY,\mathrm{dust}}_{\nu\times\nu'}$. 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

$$A^{XY,\mathrm{dust}}_{\nu\times\nu'} = a^{XY,\mathrm{dust}}_\nu \times a^{XY,\mathrm{dust}}_{\nu'}. \tag{23}$$

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

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

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 $C^{TT}_{\ell,\mathrm{dust}} \propto \ell^{-2.63}$ 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 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).

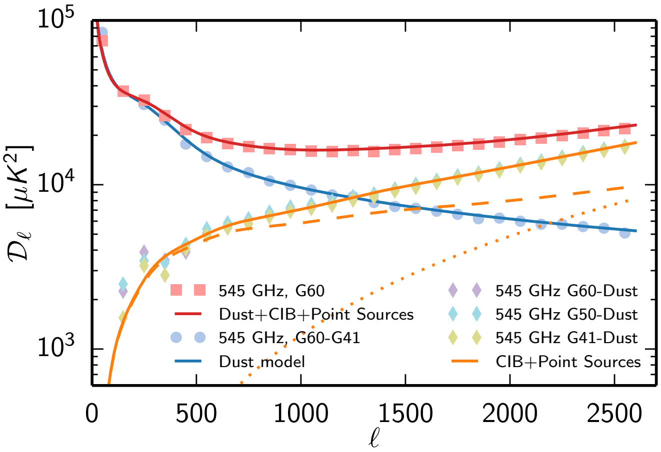

Fig. 19. Dust model at 545 GHz. The dust template is based on the G60–G41 mask difference of the 545 GHz half-mission cross-spectrum (blue line and circles, rescaled to the dust level in mask G60). Coloured diamonds display the difference between this model (rescaled in each case) and the cross half-mission 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.

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 $\boldsymbol{m}^\nu$ have in common only the CMB and a variable amount of dust, and assuming that $\boldsymbol{m}^{545} = \boldsymbol{m}^{\mathrm{cmb}} + a^{545}\boldsymbol{m}^{\mathrm{dust}}$, the cross-spectra between each of our CMB frequencies maps and the 545 GHz map is

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

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 [μK^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 $30 \leq \ell \leq 500$.

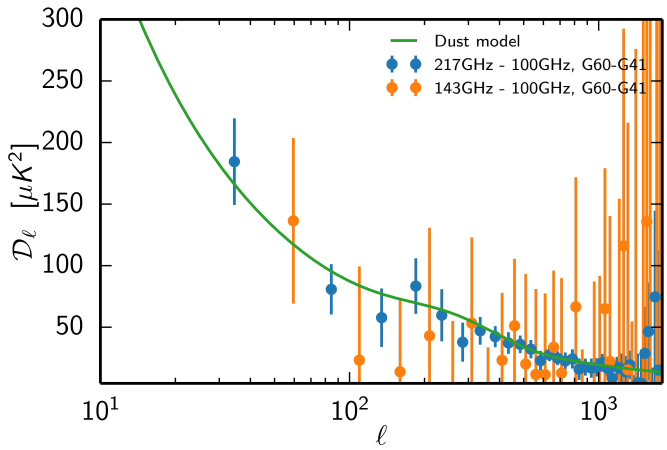

Fig. 20. Dust model versus data. In blue, the power spectrum of the double mask difference between 217 GHz and 100 GHz half-mission cross-spectra in masks G60 and G41 (complemented by the joint masks for CO, extended objects, and point sources). In orange, the equivalent spectrum for 143 and 100 GHz. 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 ± 2) μK^2 for the 100 × 100 spectrum (G70); (9 ± 2) μK^2 for 143 × 143 (G60); and (80 ± 20) μK^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 ± 8.5) μK^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 × ν (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, $C^{\rm FG}_{3000}$, satisfies $C^{\rm FG}_{3000}\mathcal{A}_{3000} = 1$ with $\mathcal{A}_\ell = \ell(\ell+1)/(2\pi)$.