arXiv:1303.5075v2 [astro-ph.CO] 25 Mar 2013

Astronomy & Astrophysics manuscript no. 0_Like_top
26th March 2013 © ESO 2013

Planck 2013 results. XV. CMB power spectra and likelihood

Planck Collaboration: P. A. R. Ade[89], N. Aghanim[62], C. Armitage-Caplan[94], M. Arnaud[76], M. Ashdown[73,6], F. Atrio-Barandela[19], J. Aumont[62],
C. Baccigalupi[88], A. J. Banday[97,10], R. B. Barreiro[70], J. G. Bartlett[1,71], E. Battaner[98], K. Benabed[63,96], A. Benoît[60], A. Benoit-Lévy[26,63,96],
J.-P. Bernard[10], M. Bersanelli[36,52], P. Bielewicz[97,10,88], J. Bobin[76], J. J. Bock[71,11], A. Bonaldi[72], L. Bonavera[70], J. R. Bond[9], J. Borrill[14,91],
F. R. Bouchet[63,96]*, F. Boulanger[62], M. Bridges[73,6,66], M. Bucher[1], C. Burigana[51,34], R. C. Butler[51], E. Calabrese[94], J.-F. Cardoso[77,1,63],
A. Catalano[78,75], A. Challinor[66,73,12], A. Chamballu[76,16,62], L.-Y Chiang[65], H. C. Chiang[28,7], P. R. Christensen[84,40], S. Church[93], D. L. Clements[58],
S. Colombi[63,96], L. P. L. Colombo[25,71], C. Combet[78], F. Couchot[74], A. Coulais[75], B. P. Crill[71,85], A. Curto[6,70], F. Cuttaia[51], L. Danese[88],
R. D. Davies[72], R. J. Davis[72], P. de Bernardis[35], A. de Rosa[51], G. de Zotti[48,88], J. Delabrouille[1], J.-M. Delouis[63,96], F.-X. Désert[55], C. Dickinson[72],
J. M. Diego[70], H. Dole[62,61], S. Donzelli[52], O. Doré[71,11], M. Douspis[62], J. Dunkley[94], X. Dupac[43], G. Efstathiou[66], F. Elsner[63,96], T. A. Enßlin[81],
H. K. Eriksen[68], F. Finelli[51,53], O. Forni[97,10], M. Frailis[50], A. A. Fraisse[28], E. Franceschi[51], T. C. Gaier[71], S. Galeotta[50], S. Galli[63], K. Ganga[1],
M. Giard[97,10], G. Giardino[44], Y. Giraud-Héraud[1], E. Gjerløw[68], J. González-Nuevo[70,88], K. M. Górski[71,100], S. Gratton[73,66], A. Gregorio[37,50],
A. Gruppuso[51], J. E. Gudmundsson[28], F. K. Hansen[68], D. Hanson[82,71,9], D. Harrison[66,73], G. Helou[11], S. Henrot-Versillé[74],
C. Hernández-Monteagudo[13,81], D. Herranz[70], S. R. Hildebrandt[11], E. Hivon[63,96], M. Hobson[6], W. A. Holmes[71], A. Hornstrup[17], W. Hovest[81],
K. M. Huffenberger[99], G. Hurier[62,78], T. R. Jaffe[97,10], A. H. Jaffe[58], J. Jewell[71], W. C. Jones[28], M. Juvela[27], E. Keihänen[27], R. Keskitalo[23,14],
K. Kiiveri[27,47], T. S. Kisner[80], R. Kneissl[42,8], J. Knoche[81], L. Knox[30], M. Kunz[18,62,3], H. Kurki-Suonio[27,47], G. Lagache[62], A. Lähteenmäki[2,47],
J.-M. Lamarre[75], A. Lasenby[6,73], M. Lattanzi[34], R. J. Laureijs[44], C. R. Lawrence[71], M. Le Jeune[1], S. Leach[88], J. P. Leahy[72], R. Leonardi[43],
J. León-Tavares[45,2], J. Lesgourgues[95,87], M. Liguori[33], P. B. Lilje[68], V. Lindholm[27,47], M. Linden-Vørnle[17], M. López-Caniego[70], P. M. Lubin[31],
J. F. Macías-Pérez[78], B. Maffei[72], D. Maino[36,52], N. Mandolesi[51,5,34], D. Marinucci[39], M. Maris[50], D. J. Marshall[76], P. G. Martin[9],
E. Martínez-González[70], S. Masi[35], S. Matarrese[33], F. Matthai[81], P. Mazzotta[38], P. R. Meinhold[31], A. Melchiorri[35,54], L. Mendes[43], E. Menegoni[35],
A. Mennella[36,52], M. Migliaccio[66,73], M. Millea[30], S. Mitra[57,71], M.-A. Miville-Deschênes[62,9], D. Molinari[51], A. Moneti[63], L. Montier[97,10],
G. Morgante[51], D. Mortlock[58], A. Moss[90], D. Munshi[89], P. Naselsky[84,40], F. Nati[35], P. Natoli[34,4,51], C. B. Netterfield[21], H. U. Nørgaard-Nielsen[17],
F. Noviello[72], D. Novikov[58], I. Novikov[84], I. J. O'Dwyer[71], F. Orieux[63], S. Osborne[93], C. A. Oxborrow[17], F. Paci[88], L. Pagano[35,54], F. Pajot[62],
R. Paladini[59], D. Paoletti[51,53], B. Partridge[46], F. Pasian[50], G. Patanchon[1], P. Paykari[76], O. Perdereau[74], L. Perotto[78], F. Perrotta[88], F. Piacentini[35],
M. Piat[1], E. Pierpaoli[25], D. Pietrobon[71], S. Plaszczynski[74], E. Pointecouteau[97,10], G. Polenta[4,49], N. Ponthieu[62,55], L. Popa[64], T. Poutanen[47,27,2],
G. W. Pratt[76], G. Prézeau[11,71], S. Prunet[63,96], J.-L. Puget[62], J. P. Rachen[22,81], A. Rahlin[28], R. Rebolo[69,15,41], M. Reinecke[81], M. Remazeilles[62,1],
C. Renault[78], S. Ricciardi[51], T. Riller[81], C. Ringeval[67,63,96], I. Ristorcelli[97,10], G. Rocha[71,11], C. Rosset[1], G. Roudier[1,75,71], M. Rowan-Robinson[58],
J. A. Rubiño-Martín[69,41], B. Rusholme[59], M. Sandri[51], L. Sanselme[78], D. Santos[78], G. Savini[86], D. Scott[24], M. D. Seiffert[71,11], E. P. S. Shellard[12],
L. D. Spencer[89], J.-L. Starck[76], V. Stolyarov[6,73,92], R. Stompor[1], R. Sudiwala[89], F. Sureau[76], D. Sutton[66,73], A.-S. Suur-Uski[27,47], J.-F. Sygnet[63],
J. A. Tauber[44], D. Tavagnacco[50,37], L. Terenzi[51], L. Toffolatti[20,70], M. Tomasi[52], M. Tristram[74], M. Tucci[18,74], J. Tuovinen[83], M. Türler[56],
L. Valenziano[51], J. Valiviita[47,27,68], B. Van Tent[79], J. Varis[83], P. Vielva[70], F. Villa[51], N. Vittorio[38], L. A. Wade[71], B. D. Wandelt[63,96,32],
I. K. Wehus[71], M. White[29], S. D. M. White[81], D. Yvon[16], A. Zacchei[50], and A. Zonca[31]

(Affiliations can be found after the references)

Preprint online version: 26th March 2013

Abstract

This paper presents the *Planck* likelihood, a complete statistical description of the two-point correlation function of the CMB temperature fluctuations that accounts for all known relevant uncertainties, both instrumental and astrophysical in nature. We use this likelihood to derive our best estimate of the CMB angular power spectrum from *Planck* over three decades in multipole moment, ℓ, covering $2 \leq \ell \leq 2500$. The main source of error at $\ell \lesssim 1500$ is cosmic variance. Uncertainties in small-scale foreground modelling and instrumental noise dominate the error budget at higher ℓs. For $\ell < 50$, our likelihood exploits all *Planck* frequency channels from 30 to 353 GHz, separating the cosmological CMB signal from diffuse Galactic foregrounds through a physically motivated Bayesian component separation technique. At $\ell \geq 50$, we employ a correlated Gaussian likelihood approximation based on a fine-grained set of angular cross-spectra derived from multiple detector combinations between the 100, 143, and 217 GHz frequency channels, marginalizing over power spectrum foreground templates. We validate our likelihood through an extensive suite of consistency tests, and assess the impact of residual foreground and instrumental uncertainties on the final cosmological parameters. We find good internal agreement among the high-ℓ cross-spectra with residuals below a few $\mu\mathrm{K}^2$ at $\ell \lesssim 1000$, in agreement with estimated calibration uncertainties. We compare our results with foreground-cleaned CMB maps derived from all *Planck* frequencies, as well as with cross-spectra derived from the 70 GHz *Planck* map, and find broad agreement in terms of spectrum residuals and cosmological parameters. We further show that the best-fit ΛCDM cosmology is in excellent agreement with preliminary *Planck* EE and TE polarisation spectra. We find that the standard ΛCDM cosmology is well constrained by *Planck* from the measurements at $\ell \lesssim 1500$. One specific example is the spectral index of scalar perturbations, for which we report a $5.4\,\sigma$ deviation from scale invariance, $n_\mathrm{s} \neq 1$. Increasing the multipole range beyond $\ell \simeq 1500$ does not increase our accuracy for the ΛCDM parameters, but instead allows us to study extensions beyond the standard model. We find no indication of significant departures from the ΛCDM framework. Finally, we report a tension between the *Planck* best-fit ΛCDM model and the low-ℓ spectrum in the form of a power deficit of 5–10% at $\ell \lesssim 40$, with a statistical significance of 2.5–3 σ. Without a theoretically motivated model for this power deficit, we do not elaborate further on its cosmological implications, but note that this is our most puzzling finding in an otherwise remarkably consistent dataset.

Key words. Cosmology: cosmic background radiation – Surveys – Methods: data analysis

---

[*] Corresponding author: F. R. Bouchet, bouchet@iap.fr.

1. Introduction

This paper, one of a set associated with the 2013 release of data from the *Planck*[1] mission (Planck Collaboration I 2013), describes the CMB power spectra and the likelihood that we derive from the *Planck* data.

The power spectrum of the Cosmic Microwave Background (CMB) is a unique signature of the underlying cosmological model (e.g., Spergel et al. 2003; Hinshaw et al. 2012). It has been measured over the whole sky by *COBE* and *WMAP*, and over smaller regions by ground-based and sub-orbital experiments (e.g., Tristram et al. 2005; Jones et al. 2006; Reichardt et al. 2009; Fowler et al. 2010; Das et al. 2011; Keisler et al. 2011; Story et al. 2012; Das et al. 2013). By mapping the whole sky to scales of a few arcminutes, *Planck* now measures the power spectrum over an unprecedented range of scales from a single experiment. To estimate cosmological parameters from the power spectrum requires a likelihood function that propagates uncertainties.

In this paper we describe the power spectra obtained from the *Planck* temperature data, as well as the associated likelihood function. Since the probability distribution of the power spectrum is non-Gaussian at large scales, we follow a hybrid approach to construct the likelihood (Efstathiou 2004, 2006), using a Gibbs sampling based approach at low multipoles, ℓ, and a pseudo-C_ℓ technique at high multipoles (Hivon et al. 2002).

The high-ℓ part of the *Planck* likelihood is based on power spectra estimated from each *Planck* detector in the frequency range 100 to 217 GHz, allowing careful assessment of each detector's response to the sky emission. We implement two independent likelihood methods. The first, used in the distributed likelihood code, estimates the power spectrum at every multipole, together with the associated covariance matrix. The second takes a simplified form, binning the spectra, and is used to explore the stability of the results with respect to different instrumental and astrophysical systematic effects. The methods give consistent results.

Unresolved extragalactic foregrounds make a significant contribution to the power spectra at high multipoles. We develop a model for these foregrounds, designed to allow the *Planck* likelihood to be combined with high resolution data from the Atacama Cosmology Telescope (ACT) and the South Pole Telescope (SPT). We combine frequencies and model unresolved foregrounds in a physical way, as in e.g., Shirokoff et al. (2010); Dunkley et al. (2011); Reichardt et al. (2012), performing component separation at small scales at the power spectrum level. On large scales, $\ell < 50$, Galactic contamination is more significant. We use the *Planck* temperature maps in the range $30 \leq \nu \leq 353$ GHz to separate Galactic foregrounds in the maps, and then estimate the full probability distribution of the CMB power spectrum.

This paper is structured as follows. In Sect. 2 we describe the pseudo-C_ℓ likelihoods, and in Sect. 3 set up the foreground model. The power spectra and derived cosmological parameters are presented in Sects. 4 and 5, and an assessment of their accuracy and robustness is made in Sects. 6 and 7.

In Sect. 8 we describe the low-ℓ likelihood, and conclude by presenting the complete *Planck* likelihood in Sect. 9.

Table 1. Detectors used to make the maps for this analysis. Spider Web Bolometers (SWB) are used individually; Polarised Sensitive Bolometer pairs (PSBs, denoted a and b) are used in pairs, and we consider only the maps estimated from two pairs of PSBs. The relevant effective beams, and their uncertainties, are given in Planck Collaboration VII (2013).

Set name	Frequency [GHz]	Type	Detectors	FWHM^a [arcmin]
100-ds0	100	PSB	All 8 detectors	9.65
100-ds1	100	PSB	1a+1b + 4a+4b	
100-ds2	100	PSB	2a+2b + 3a+3b	
143-ds0	143	MIX	11 detectors	7.25
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	4.99
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	

2. High-ℓ likelihoods

The *Planck* maps consist of the order 5×10^7 pixels for each detector, so a likelihood described directly at the pixel level would be too time consuming. A significant compression of data can be achieved with minimal information loss using pseudo-C_ℓ power spectra, even in the case of incomplete sky coverage. Here we describe the form of the likelihood function of the compressed data, given a sky signal and instrumental model.

Following Hamimeche & Lewis (2008), we assume a Gaussian form of the likelihood based on pseudo-spectra that have been corrected to account for partial sky masking. We use a 'fine-grained' data description, computing spectra of maps from individual detectors or detector sets. Table 1 describes the 13 maps used in the analysis, spanning 100 to 217 GHz. We compute the spectra at these multiple frequencies to simultaneously constrain the CMB and foreground contributions. We choose these frequencies as a trade-off between adding further information, and adding further complexity to the foreground model, which would be needed to including the adjacent 70 and 353 GHz channels (see Sect. 3 for further discussion). In our baseline analysis the spectra are computed at each multipole, together with an estimate of the full covariance matrix with off-diagonal errors between different spectra and multipoles. As in the *WMAP* analysis, we use only cross-spectra between detectors, alleviating the need to accurately model the mean noise contribution.

In this section, we begin with a reminder of the pseudo-spectrum approach, and describe our baseline likelihood distribution, hereafter referred to as the `CamSpec` likelihood. We then show how a compression of spectra within a given frequency can be achieved with negligible loss of information. We describe the signal and instrument model, including detector noise properties, calibration, and beam uncertainties.

Next, we describe an alternative, simpler, form of the likelihood, hereafter referred to as `Plik`, based on binned power spectra with an inverse-Wishart distribution. This does not require the pre-computation of large covariance matrices, so changing the sky or instrument modeling is straightforward. This simpler form of the likelihood will be used to assess the robustness of our likelihood methodology with respect to technical choices and astrophysical foreground modeling.

---

[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 (in particular the lead countries France and Italy), with contributions from NASA (USA) and telescope reflectors provided by a collaboration between ESA and a scientific consortium led and funded by Denmark.

In Sect. 7.5 we also compare these likelihoods to pseudo-spectra computed directly from CMB maps estimated by multi-frequency component separation (Planck Collaboration XII 2013).

2.1. *The `CamSpec` likelihood*

We define $\tilde{T}^i_{\ell m}$ as the spherical harmonic coefficients of the weighted temperature map of detector i. The pseudo-spectrum at multipole ℓ, for the detector pair (i, j), is then given by

$$\tilde{\mathbf{C}}^{ij}_\ell = \frac{1}{2\ell+1}\sum_m \tilde{\mathbf{T}}^i_{\ell m}\tilde{\mathbf{T}}^{j\dagger}_{\ell m}, \tag{1}$$

where the dagger, †, denotes the Hermitian transpose. This is related to the 'deconvolved' spectrum, $\hat{C}^{T_{ij}}$, by a coupling matrix,

$$\tilde{C}^{T_{ij}} = M^{TT}_{ij}\hat{C}^{T_{ij}}. \tag{2}$$

For an isotropic signal on the sky, the ensemble average of these deconvolved spectra are equal to the spectra of the theoretical models (including CMB and isotropic unresolved foregrounds) that we wish to test. For completeness, the coupling matrices are given explicitly in Appendix A.1.

In the first method, `CamSpec`, we form the deconvolved spectra $\hat{C}_\ell$ without any prior smoothing of the pseudo-spectra $\tilde{C}_\ell$. Even for the largest sky masks used in our analysis (see Sect. 3), the coupling matrices are non-singular. The deconvolution requires the evaluation of $\sim N^2_{\rm map}$ coupling matrices for a data set with $N_{\rm map}$ sky maps, which takes a moderate, but not excessive, amount of computer time.

A more challenging computational task is to compute the covariances of the pseudo-spectra, i.e., $\mathrm{Cov}(\tilde{C}^{T_{ij}}\tilde{C}^{T_{pq}})$. Here we need to compute $N^4_{\rm map}$ coupling matrices, and the problem rapidly becomes computationally intractable even for relatively low values of $N_{\rm map}$. For the moment we will assume that these covariance matrices are available and describe their computation in Appendix A.4. We will use the notation $\tilde{\mathbf{X}} = \mathrm{Vec}(\tilde{\mathbf{C}})$ to denote a column vector for which the index p of a single element X_p denotes the map combination (i, j) and multipole ℓ. We denote the covariance matrix of this vector as

$$\tilde{\mathcal{M}} = (\tilde{\mathbf{X}} - \langle\tilde{\mathbf{X}}\rangle)(\tilde{\mathbf{X}} - \langle\tilde{\mathbf{X}}\rangle)^T. \tag{3}$$

As explained later, the deconvolved detector set cross-spectra given by Eq. 2 can be efficiently combined within a given frequency pair after a small effective recalibration, taking into account their respective isotropised beam transfer function and noise levels (see Appendix A.3 for the detailed procedure). Covariance estimates of these combined spectra can be deduced from those of the detector set cross-spectra. The covariance matrix is computed for a fixed fiducial model, and we approximate the likelihood as a Gaussian, described in Appendix A.5. The likelihood thus takes the form $p = \mathrm{e}^{-S}$ with

$$S = \tfrac{1}{2}\left(\hat{\mathbf{X}}-\mathbf{X}\right)^T \hat{\mathcal{M}}^{-1}\left(\hat{\mathbf{X}}-\mathbf{X}\right).$$

For the current analysis we include the following (deconvolved) spectrum combinations,

$$\hat{\mathbf{X}} = (\hat{C}^{100\times100}_\ell, \hat{C}^{143\times143}_\ell, \hat{C}^{217\times217}_\ell, \hat{C}^{143\times217}_\ell), \tag{4}$$

coupled to a parametric model of the CMB and foreground power spectra. The multipole ranges we select depend on frequency, as described in Sect. 5. We do not include the 100 × 143 and 100 × 217 spectra since these spectra carry little additional information about the primary CMB anisotropies, but would require us to solve for additional unresolved foreground parameters. This tradeoff of information versus complexity was also considered for the use of the 70 GHz and 353 GHz data, which we choose not to include except for cross-checks.

The fiducial covariance matrix is composed of the blocks shown in Fig. 1. The off-diagonal blocks in this matrix accurately account for the correlations between the power spectra at different frequencies.

This description would be sufficient for perfectly known calibrations and beam transfer functions of each detector sets' cross-spectra. Planck Collaboration VII (2013) describes in detail these uncertainties, and shows that for each detector set pair, (i, j), the effective beam transfer function can be expressed as

$$B^{ij}(\ell) = B^{ij}_{\rm mean}(\ell)\exp\left(\sum_{k=1}^{n_{\rm modes}} g^{ij}_k E^{ij}_k(\ell)\right), \tag{5}$$

described further in Appendix A.6, with $n_{\rm modes}$ beam error eigenmodes $E^{ij}_k(\ell)$, and their covariance matrix. These modes are then combined into generalised beam eigenmodes corresponding to the spectra $\hat{\mathbf{X}}$. The associated covariance matrix is used to construct a Gaussian posterior distribution of the eigenmodes, which allows marginalization over the uncertainties.

Finally, in the construction of the covariance matrix, one needs to accurately specify the contribution of the instrumental noise. Even if there is no bias on the spectra due to instrumental noise, having removed auto-spectra, the latter dominates the covariance matrix on small scales. Fortunately, the *Planck* scanning strategy at the ring level allows us to make estimates of the noise pseudo-spectra from half-ring difference maps (see Planck Collaboration XII 2013). These half-ring difference maps, together with the knowledge of the noise variance per pixel for each detector set, can be used to derive the noise contribution to the covariance matrix with good accuracy (see Appendix A.8 for details).

2.2. *The `Plik` likelihood*

We now describe the alternative form of the likelihood, inspired by Cardoso et al. (2008), used for cross-checks and robustness tests. We start from the full-sky exact likelihood for a Gaussian signal, which for $N_{\rm map}$ detector maps is given by

$$p(\mathrm{maps}|\theta) \propto \exp -\{\sum_\ell (2\ell+1)\mathcal{K}(\hat{C}_\ell, \mathbf{C}_\ell(\theta))\},$$

where θ is a vector containing the parameters of the signal model, and $\hat{C}_\ell$ are the empirical angular spectra. $\mathcal{K}(\mathbf{A}, \mathbf{B})$ denotes the Kullback divergence between two n-variate zero-mean Gaussian distributions with covariance matrices $\mathbf{A}$ and $\mathbf{B}$, and is given by

$$\mathcal{K}(\mathbf{A}, \mathbf{B}) = \frac{1}{2}\Big[\mathrm{tr}(\mathbf{A}\mathbf{B}^{-1}) - \log\det(\mathbf{A}\mathbf{B}^{-1}) - n\Big].$$

As already noted, a sky cut introduces off-diagonal couplings between different multipoles. In this method we bin the power spectra in such a way that these off-diagonal terms of the covariance are negligible. This is adequate to model sources with

$$\tilde{\mathcal{M}} = \left(\begin{array}{c|c|c|c}
(100\times 100)\times(100\times 100) & (100\times 100)\times(143\times 143) & (100\times 100)\times(217\times 217) & (100\times 100)\times(143\times 217) \\ \hline
(143\times 143)\times(100\times 100) & (143\times 143)\times(143\times 143) & (143\times 143)\times(217\times 217) & (143\times 143)\times(143\times 217) \\ \hline
(217\times 217)\times(100\times 100) & (217\times 217)\times(143\times 143) & (217\times 217)\times(217\times 217) & (217\times 217)\times(143\times 217) \\ \hline
(143\times 217)\times(100\times 100) & (143\times 217)\times(143\times 143) & (143\times 217)\times(217\times 217) & (143\times 217)\times(143\times 217)
\end{array}\right).$$

Figure 1. The covariance matrix blocks used in the likelihood, accounting for the correlations between cross-spectra estimated from the 100, 143, and 217 GHz channels.

slowly varying spectra, such as foregrounds, and the CMB anisotropies for standard cosmologies. In this case, the likelihood takes the form

$$p(\text{maps}|\theta) \propto \exp -\mathcal{L}(\theta), \quad \text{with} \quad \mathcal{L}(\theta) = \sum_{\mathrm{q}=1}^{\mathrm{Q}} \mathrm{n_q}\, \mathcal{K}(\hat{\mathrm{C}}_\mathrm{q}, \mathbf{C}_\mathrm{q}), \tag{6}$$

where the angular spectra $\hat{C}_\ell$ for each cross-frequency spectrum have been averaged into Q spectral bins using spectral windows $w_q(\ell)$ $(q = 1, \dots, Q)$, with

$$\hat{C}_q = \sum_\ell w_q(\ell)\, _\ell\hat{C}_\ell, \qquad \mathbf{C}_q = \sum_\ell w_q(\ell)\, \mathbf{C}_\ell.$$

Here $w_q(\ell)$ denotes the window function for the q-th bin, and the same symbol, $\mathbf{C}$, is used to denote binned or unbinned spectra. The effective number of modes in the q-th bin is

$$n_q = f_\mathrm{sky} \cdot \frac{(\sum_\ell w_q(\ell)^2)^2}{\sum_\ell w_q(\ell)^4/(2\ell+1)}.$$

We adopt a spectral binning defined by

$$w_q(\ell) = \begin{cases} \frac{\ell(\ell+1)(2\ell+1)}{\sum_{\ell^q_\mathrm{min}}^{\ell^q_\mathrm{max}} \ell(\ell+1)(2\ell+1)} & \ell^q_\mathrm{min} \le \ell \le \ell^q_\mathrm{max}, \\ 0 & \text{otherwise.} \end{cases}$$

The `Plik` bin width is $\Delta\ell = 9$ from $\ell = 100$ to $\ell = 1503$, then $\Delta\ell = 17$ to $\ell = 2013$, and finally $\Delta\ell = 33$ to $\ell_\mathrm{max} = 2508$. This ensures that correlations between any two bins are smaller than 10 %.

While this binned likelihood approximation does not fully capture all couplings between different multipoles, it has a notable advantage in computational speed, and it agrees well with the primary likelihood. It is therefore very well suited for performing an extensive suite of robustness tests, as many more parameters can be considered in a short time. Further, instrumental effects can be investigated quickly to assess the agreement between pairs of detectors within a frequency channel, such as individual detector calibrations and beam errors.

A specific example is the impact of beam uncertainty parameters on the likelihood. This can be investigated by re-expressing the model covariance matrices as

$$\mathbf{C}_\ell = \mathbf{B}_\ell(\gamma)\, \mathbf{C}_\ell(\theta)\, \mathbf{B}_\ell(\gamma)^\mathrm{T}, \tag{7}$$

where $\mathbf{C}_\ell(\theta)$ is the model covariance including both signal and noise, and $\mathbf{B}_\ell(\gamma)$ is a diagonal matrix encoding the beam (and calibration) errors with elements given by [2]

$$B^i_\ell(\gamma) = \exp\left(\sum_{k=1}^{n_\mathrm{modes}} \delta^i_k\, E^{ii}_k(\ell)\right). \tag{8}$$

Here, $E^{ii}_k(\ell)$ are the eigenmodes of the (auto-)spectra, similar to Eq. 5. Note that Eq. 7 does not contain the *mean* beam transfer function, since it is already included in the empirical spectra. Thus, using Eq. 7 `Plik` approximates the cross-spectrum beam errors as the harmonic mean of the corresponding auto-spectrum beam errors, under the assumption that $\mathbf{B}_\ell$ is diagonal between detectors. This approximate factorisation is intrinsically linked to the assumed Kullback shape of the `Plik` likelihood, and is later demonstrated to work well for both simulations and data.

The `Plik` likelihood method also provides a direct estimate of the detector noise power spectra as it can include the empirical auto-spectra, and we find that these noise estimates are in good agreement with the noise spectra used to construct the `CamSpec` likelihood covariance matrix. The method can also produce a binned CMB power spectrum independent of the underlying cosmological model, providing a direct quality assessment of the foreground model parametrisation. In practice, we proceed in two steps. First, we jointly estimate the noise together with all other parameters using both auto and cross-spectra. Then we fix the noise estimates, and use the fiducial Gaussian approximation to explore the remaining free parameters excluding the auto-spectra, optionally including only specific data combinations.

3. Foreground emission model and sky masks

3.1. Sky masks

The Galactic emission varies strongly in both complexity and strength across the sky. It is therefore necessary to find a balance between maximizing the sky coverage to reduce statistical uncertainties, and establishing a simple yet efficient foreground model. In this paper, we threshold the 353 GHz temperature map to define a basic set of diffuse Galactic masks (shown in Appendix B), which form a sequence of increasing sky fraction, to minimize the contribution from diffuse dust emission. The sky fractions retained by these masks is summarized in Table 2.

For *Planck*, we need to estimate the covariance matrices to percent level precision. For temperature spectra, and in the

---

[2] From Eqs. 5, 7, and 8, we have $\delta^i_k = g^{ii}_k/2$ at first order.

Table 2. Area of sky retained by combining diffuse foreground and point source masks, once apodised.

Mask	Sky fraction [%]	Sky area [deg^2]
CL31	30.71	12 668
CL39	39.32	16 223
CL49	48.77	20 121

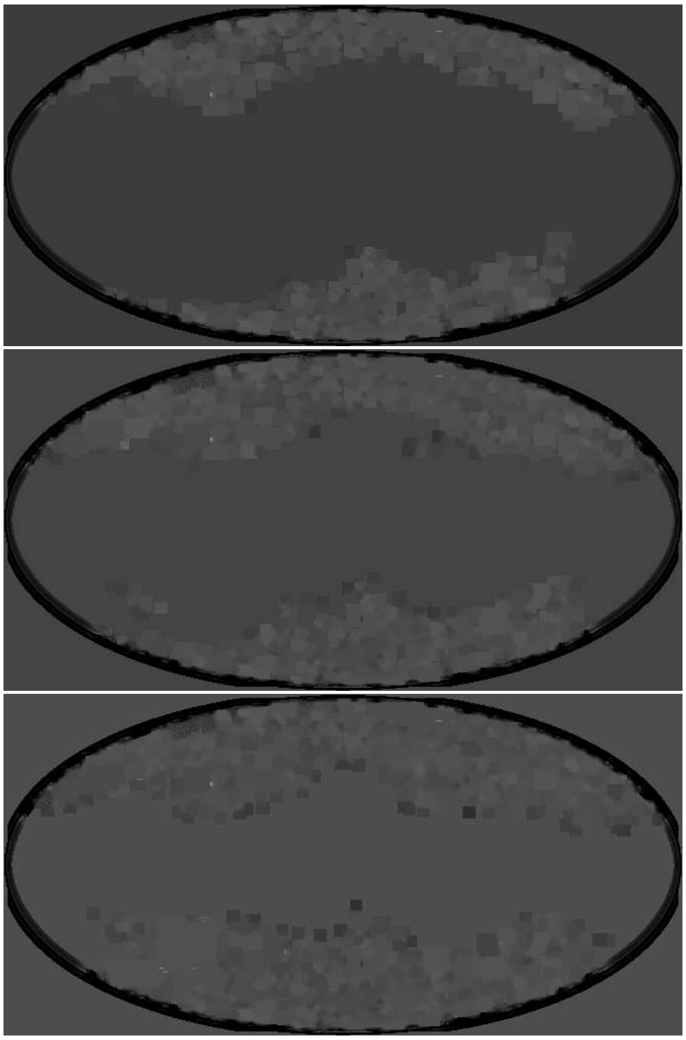

Figure 2. The set of masks (CL31, CL39, CL49) used for the likelihood analyses.

absence of point source holes, this precision can be achieved with sharp, non-apodised Galactic masks (Efstathiou 2004). However, the inclusion of point source holes introduces non-negligible low-ℓ power leakage, which in turn can generate errors of a few percent in the covariance matrices. We reduce this leakage by apodising the diffuse Galactic masks (see Appendix B for details). The point source mask is based on the union of the point sources detected between 100 and 353 GHz, and is also apodized. The point source flux cut is not critical, since the amplitudes of the Poisson contributions of unresolved sources are allowed to vary over a wide range in the likelihood analysis. Thus, we do not impose tight priors from source counts and other CMB experiments on the Poisson amplitudes. A set of the combined Galactic and point source masks, referred to as 'CLx', where 'x' is the percentage of sky retained, are shown in Fig. 2.

3.2. Galactic emission

The contamination from diffuse Galactic emission at low to intermediate multipoles can be reduced to low levels compared to CMB anisotropies by a suitable choice of masking. However, even with conservative masking, the remaining Galactic emission at high multipoles is non-negligible compared to other unresolved components, such as the Cosmic Infrared Background (CIB) anisotropies at 143 and 217 GHz. A clear way of demonstrating this is by differencing the power spectra computed with different masks, thereby highlighting the differences between the isotropic and non-isotropic unresolved components. Figure 3 shows (up to $\ell \leq 1400$) the 217 GHz power spectrum difference for the mask1 and mask0 masks[3], minus the corresponding difference for the 143 GHz frequency channel. Any isotropic contribution to the power spectrum (CMB, unresolved extragalactic sources, etc.) will cancel in such a double difference, leaving a non-isotropic signal of Galactic origin, free of the CMB induced cosmic variance scatter. Above $\ell > 1400$, Fig. 3 shows the mask differenced 217 GHz power spectrum, as the instrumental noise becomes significant at $\ell \gtrsim 1400$ for the 143 GHz channel.

In the same figure, these difference spectra are compared to the unbinned mask-differenced 857 GHz power spectrum, scaled to 217 GHz adopting a multiplicative factor[4] of $(9.93 \times 10^{-5})^2$; the dotted line shows a smooth fit to the unbinned spectrum. The agreement between this prediction and the actual dust emission at 217 GHz is excellent, and this demonstrates conclusively the existence of a small-scale dust emission component with an amplitude of $\sim 5-15\,\mu\mathrm{K}^2$ at 217 GHz if mask1 is used.

For cosmological parameter analysis this small-scale dust component must be taken into account, and several approaches may be considered:

1. Fit to a template shape, e.g., as shown by the dotted line in Fig. 3.
2. Reduce the amplitude by further masking of the sky.
3. Attempt a component separation by using higher frequencies.

The main disadvantage of the third approach is a potential signal-to-noise penalty, depending on which frequencies are used, as well as confusion with other unresolved foregrounds. This is particularly problematic with regards to the CIB, which has a spectrum very similar to that of Galactic dust. In the following we therefore adopt the two former solutions.

It is important to understand the nature of the small scale dust emission, and, as far as possible, to disentangle this emission from the CIB contribution at the HFI cosmological frequencies. We use the 857 GHz power spectrum for this purpose, noting that the dust emission at 857 GHz is so intense that this particular map provides an effectively noise-free dust emission map. In Fig. 4 we again show the 857 GHz mask power spectrum difference, but this time plotted on a log-log scale. The solid line shows the corresponding best-fit model defined by

$$\mathcal{D}_\ell = \frac{A\,(100/\ell)^\alpha}{[1 + (\ell/\ell_c)^2]^{\gamma/2}}, \quad (9)$$

---

[3] These are the combination of the non-apodised Galactic masks G35 and G22 with the apodised point source mask PSA82.

[4] The scaling coefficient for the 143 GHz spectrum is $(3.14 \times 10^{-5})^2$, derived from the 7-parameter fitting function of Eq. A.46.

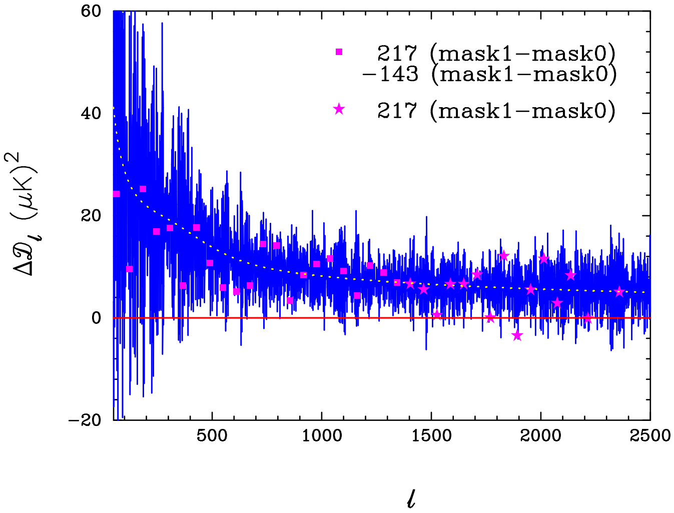

Figure 3. Differences between power spectra evaluated from masks1 and mask0, showing the presence of Galactic dust. For $\ell \leq 1400$ the spectra show the 217−143 "double-differenced" power spectrum, rescaled to correct for dust emission at 143 GHz. For $\ell > 1400$ the 217 mask differenced power spectrum is plotted. The blue line shows the 857 GHz mask-differenced power spectrum scaled to 217 GHz as described in the text, fit by the dotted line.

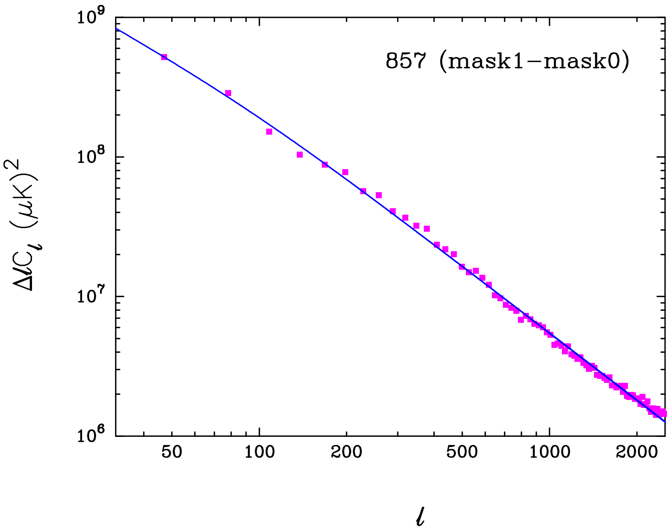

Figure 4. 857 GHz mask-differenced power spectrum (points), interpreted as Galactic dust emission. The solid line shows the best-fit model defined by Eq. 9.

with $A = 5.729 \times 10^8\,\mu\mathrm{K}^2$, $\alpha = 0.169$, $\ell_c = 905$, and $\gamma = 0.427$. At high multipoles this fit asymptotically approaches $C_\ell \propto \ell^{-2.6}$, which is compatible with previous knowledge about diffuse Galactic emission, i.e., a power-law behaviour with an index close to −3 extending to high multipoles (see e.g., Miville-Deschênes et al. 2007).

The upper panel in Fig. 5 shows the 857 GHz spectra for the four Galactic masks (G22, G35, G45 through to G56) with the point-source mask applied. They are compared to the 857 GHz CIB power spectrum from Planck Collaboration et al. (2011c), which for $\ell > 500$ can be described approximately as $\ell C_\ell \approx 7 \times 10^6\,\mu\mathrm{K}^2$. The best fit models of Eq. 9 are also shown, fitted to $\ell \leq 500$, where we expect diffuse emission to be dominant. The lower panel of Fig. 5 shows the same power spectra after subtracting the best-fit dust model. After subtracting the Galactic dust component, the recovered power spectra are consistent with the CIB measured in Planck Collaboration et al. (2011c) for all masks. The excess at high multipoles may be due to a combination of aliasing of large scale power through the point source masks at $\ell \gtrsim 3000$, Galactic point sources, and uncertainties in the 857 GHz beams.

The model explains by construction the 'double-difference' plot shown in Fig. 3. Specifically, this emission is consistent with a cirrus-like power spectrum, $C_\ell \propto \ell^{-2.6}$, extrapolated to high multipoles. Furthermore, the results of Fig. 5 demonstrate that over a wide area of sky, we can understand the 857 GHz power

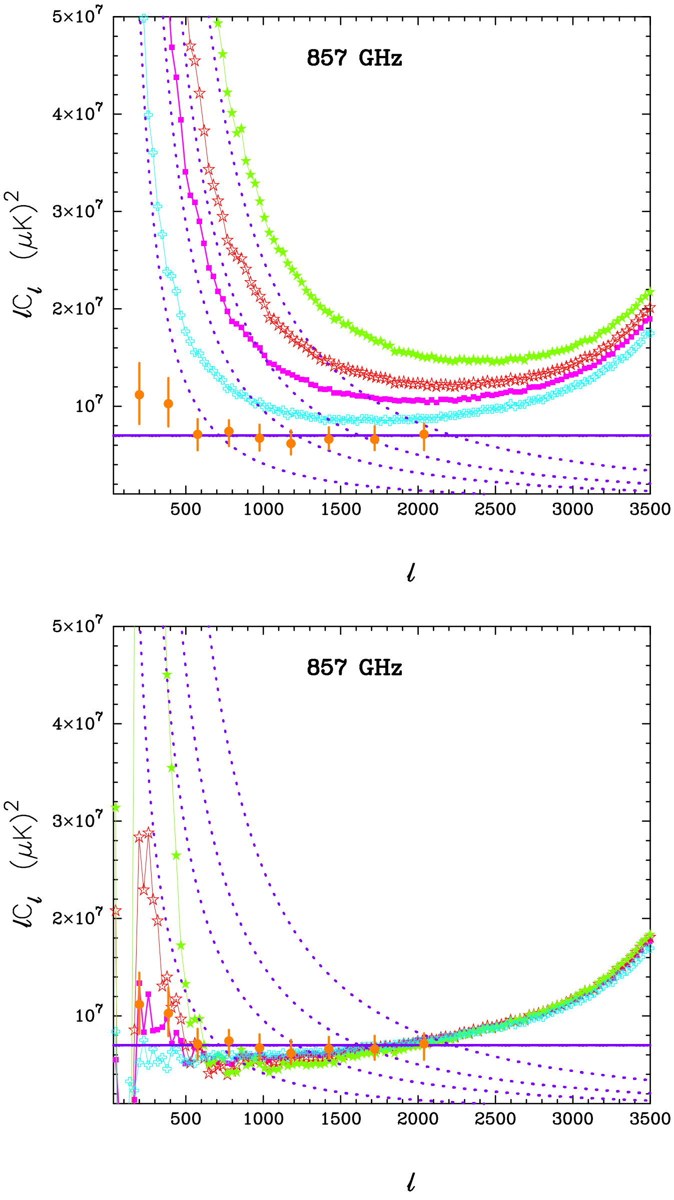

Figure 5. *Top*: 857 GHz power spectra for the four different masks defined in the text. The dotted lines show the best-fit model defined by Eq. 9 fit to $\ell \leq 500$, capturing the Galactic dust. An estimate of the CIB power spectrum is shown in orange points Planck Collaboration et al. (2011c). *Bottom*: Power spectra after subtracting the Galactic dust model.

spectrum in terms of a 'universal' cirrus spectrum together with an isotropic CIB component. These results provide strong evidence that an extragalactic CIB component dominates over the diffuse Galactic emission at multipoles $\ell \gtrsim 500$ over the full range of HFI frequencies outside the CL31 mask.

We take a different approach for the Galactic dust correction with the `Plik` likelihood. Rather than correcting the empirical spectra during a pre-processing step, the `Plik` likelihood implements an explicit one-parameter model that describes the dust contribution to the cross-spectrum between detectors i and j,

$$C_\ell^{\rm Dust}(i,\, j) = A^{\rm Dust}\, F(\nu_i, \nu_0) F(\nu_j, \nu_0) \left(\frac{\ell}{500}\right)^{-\gamma_d} g_i^{\rm Dust} g_j^{\rm Dust}. \qquad (10)$$

Here

$$F(\nu,\nu_0) = \frac{\nu^{\beta_d} B(T_d, \nu)}{\frac{\partial B(T_{\rm CMB},\nu)}{\partial T}} \Big/ \frac{\nu^{\beta_d} B(T_d, \nu_0)}{\frac{\partial B(T_{\rm CMB},\nu_0)}{\partial T}}, \qquad (11)$$

where the dust amplitude, $A^{\rm Dust}$, is measured in units of μK^2, ν_i is the reference frequency for map i, ν_0 is a reference frequency which is taken to be 143 GHz, $B(T, \nu)$ is the emission law of a blackbody with temperature T, and the dust color-correction terms, $g_i^{\rm Dust}$, are computed by integrating the dust spectrum within the spectral band of each detector (set). We fix the frequency and angular scaling parameters to $\gamma_d = 2.6$, $\beta_d = 1.6$ and $T_d = 18$K.

3.3. Poisson power from unresolved point sources

Unresolved galaxies contribute both shot noise and clustered power to the *Planck* maps. The Poisson contribution leads to a scale independent tem, $C_\ell =$ constant. We model this power with a single amplitude parameter for each auto-spectrum ($A^{\rm PS}_{100}$, $A^{\rm PS}_{143}$, and $A^{\rm PS}_{217}$) and a cross correlation coefficient for each cross spectrum ($A^{\rm PS}_{143\times217} = r_{\rm CIB}\sqrt{(A^{\rm CIB}_{143}A^{\rm CIB}_{217})}$). These quantities are not of primary interest for cosmological results, so to avoid modelling error we do not separate the power into that sourced by "dusty" or "radio" galaxies (i.e., with increasing or decreasing brightness with frequency, respectively) as is done in the analysis of the ACT and SPT power spectra (Dunkley et al. 2011; Reichardt et al. 2012; Dunkley et al. 2013). We also make no assumptions about their coherence between frequencies.

The Poisson power can be related to the flux density dN/dS via

$$C_\ell = \frac{1}{4\pi} \int d\hat{n} \int_0^{S_{\rm cut}(\hat{n})} dS\, S^2 \frac{dN}{dS}, \qquad (12)$$

where we have explicitly introduced the *Planck* flux cut $S_{\rm cut}(\hat{n})$. Since *Planck* utilizes a constant signal-to-noise cut, and the *Planck* noise varies significantly across the sky, this flux cut has a spatial dependence. Although this does not alter the shape of the Poisson term, extra care must be taken when comparing results[5] with models of dN/dS. In Sect. 7.3, we explore the consistency between the Poisson power recovered from the *Planck* power spectrum analysis and predictions from source count measurements.

3.4. Clustered power from unresolved point sources

Unresolved galaxies also contribute power because they trace large-scale structures. The mean flux from the radio galaxies is much smaller than that from the dusty galaxies, so only the dusty galaxies contribute a significant clustering term (Millea et al. 2012). The CIB clustering has been studied extensively, starting with Bond et al. (1986, 1991). Further theoretical investigation (Scott & White 1999; Haiman & Knox 2000) was stimulated by the detection of the infrared background in the *COBE* data (Puget et al. 1996; Fixsen et al. 1998), and the detection of bright "sub-millimeter" galaxies in SCUBA data (Hughes et al. 1998). Subsequently, the clustering has been detected at 160 microns (Lagache et al. 2007), at 250, 350 and 500 microns by the Balloon-borne Large Aperture Submillimeter Telescope (BLAST, Viero et al. 2009; Hajian et al. 2012) and at 217 GHz by SPT and ACT (Hall et al. 2010; Dunkley et al. 2011). Recent *Planck* measurements of the Cosmic Infrared Background (Planck Collaboration et al. 2011c) have extended the measurements at 217 GHz, 353 GHz, and 545 GHz to larger scales, and recent *Herschel* measurements (Amblard et al. 2011) have improved on the BLAST measurements and extended them to smaller angular scales.

Rather than attempt to establish a physical model of the CIB, we adopt in this analysis a phenomenological model that captures the CIB uncertainties for both Planck and high-ℓ experiments. Our baseline CIB model is a power-law spectrum with a free spectral index, $\mathcal{D}_\ell^{\rm CIB} \propto \ell^{\gamma^{\rm CIB}}$, with an amplitude at each frequency, $A^{\rm CIB}_{143}$ and $A^{\rm CIB}_{217}$, and a cross-correlation between frequencies, $A^{\rm CIB}_{143\times217} = r_{\rm CIB}\sqrt{(A^{\rm CIB}_{143}A^{\rm CIB}_{217})}$. We assume that the CIB clustering power at 100 GHz is negligible.

3.5. Unresolved Sunyaev-Zeldovich effects

Based on analysis of ACT and SPT data, the thermal Sunyaev-Zeldovich (tSZ) contribution is expected to contribute approximately $\mathcal{D}^{\rm tSZ}_{\ell=3000} \sim 9\,\mu{\rm K}^2$ at 100 GHz and $\mathcal{D}^{\rm tSZ}_{\ell=3000} \sim 4\,\mu{\rm K}^2$ at 143 GHz (Reichardt et al. 2012; Dunkley et al. 2013; Sievers et al. 2013). The kinetic Sunyaev-Zeldovich (kSZ) effect is expected to have a similar, or smaller, contribution, with $\mathcal{D}^{\rm kSZ}_{\ell=3000} \lesssim 7\,\mu{\rm K}^2$. In addition, theoretical arguments (Reichardt et al. 2012; Addison et al. 2012c) suggest that there should be a tSZ x CIB correlation that should contribute about the same order of magnitude as the kSZ term at 143 GHz.

For *Planck*, all of these SZ contributions are small in comparison to other unresolved foregrounds and are therefore poorly constrained by *Planck* data alone. Nevertheless, to eliminate biases in cosmological parameters (Millea et al. 2012; Zahn et al. 2005), we model their contributions, with appropriate constraints from higher resolution CMB experiments, using three templates.

First, for the thermal SZ effect we adopt the family of templates described by Efstathiou & Migliaccio (2012). These are based on the Komatsu & Seljak (2002) model, but use the 'universal' X-ray electron pressure profile, P_e, of Arnaud et al. (2010) extrapolated to high redshift via

$$P_e(z) \propto [(1-\Omega_\Lambda)(1+z)^3 + \Omega_\Lambda]^{4/3-\epsilon/2}. \qquad (13)$$

Here, ϵ describes departures from self-similar evolution, and a value of $\epsilon = 0.5$, which is adopted as the default for parameter estimation purposes, provides a good match to the results from recent hydrodynamical numerical simulations incorporating feedback processes (Battaglia et al. 2010, 2012). Figure 6

---

[5] One must also account for the fact that these numbers correspond to the amplitude for a suitably averaged spectral band, which is approximately that of the map, and is described in detail in Planck Collaboration IX (2013).

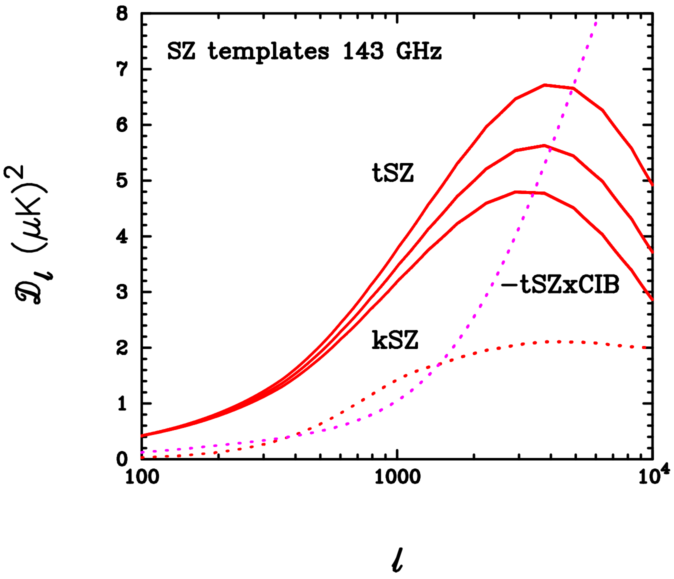

Figure 6. SZ templates at 143 GHz computed for a normalization of $\sigma_8 = 0.8$. The tSZ templates are from the model of Efstathiou & Migliaccio (2012) for three values of the evolution parameter ϵ, $\epsilon = 0$ (top), $\epsilon = 0.25$ (middle) and $\epsilon = 0.5$ (lower). The kinetic SZ template is from Trac et al. (2011). The tSZ x CIB cross correlation (143 × 143 GHz) is from the Addison et al. (2012c) template with parameters described in the text, and is negative for 143 × 143 GHz.

shows the tSZ templates for three values of ϵ; the template shape is not particularly sensitive to ϵ. We treat the (dimensionless) normalization of the tSZ template at 143 GHz as an adjustable parameter,

$$\mathcal{D}_\ell^{\mathrm{tSZ}} = A_{143}^{\mathrm{tSZ}} \mathcal{D}_\ell^{\mathrm{template}}. \tag{14}$$

This parameter fixes the amplitude at 100 GHz via the frequency dependence of the tSZ effect,

$$C^{\mathrm{tSZ}} \propto \left(x \frac{e^x + 1}{e^x - 1} - 4 \right)^2, \quad x = \frac{h\nu}{kT}. \tag{15}$$

We neglect the tSZ at 217 GHz.

Second, for the kinetic SZ effect we adopt the template described by Trac et al. (2011), and as in Eq. 14 we treat the dimensionless amplitude of the template, A^{kSZ}, as a free parameter,

$$\mathcal{D}_\ell^{\mathrm{kSZ}} = A^{\mathrm{kSZ}} \mathcal{D}_\ell^{\mathrm{template}}. \tag{16}$$

Third and finally, for the cross-correlation between the thermal SZ component and the CIB we adopt the template described by Addison et al. (2012b). In this case, the amplitude is parameterised in terms of a single correlation coefficient,

$$\left.\begin{aligned} \mathcal{D}_\ell^{\mathrm{tSZxCIB}} &= -2\xi \sqrt{\mathcal{D}_{3000}^{\mathrm{tSZ143}} \mathcal{D}_{3000}^{\mathrm{CIB143}}}\, \mathcal{D}_\ell^{\mathrm{template}} \quad (143 \times 143), \\ \mathcal{D}_\ell^{\mathrm{tSZxCIB}} &= -\xi \sqrt{\mathcal{D}_{3000}^{\mathrm{tSZ143}} \mathcal{D}_{3000}^{\mathrm{CIB217}}}\, \mathcal{D}_\ell^{\mathrm{template}} \quad (143 \times 217). \end{aligned}\right\} \tag{17}$$

These templates are plotted in Fig. 6, normalized to $\sigma_8 = 0.8$ and with $\xi = 1.0$ using a fiducial CIB amplitude. Note that with these parameters, the tSZ x CIB cross-spectrum approximately cancels the kSZ spectrum at 143 GHz.

As seen in Fig. 6, the SZ contributions are at the level of a few $\mu\mathrm{K}^2$, which, although small, must be taken into account to assess inter-frequency residuals. However, one can see that these templates have similar shapes at multipoles $\lesssim 2000$, and therefore they cannot be disentangled using *Planck* data alone. On the other hand, higher resolution experiments can break this degeneracy, and as shown in Planck Collaboration XVI (2013), the combination of *Planck*, ACT, and SPT, better constrains the amplitude of the thermal SZ effect. The ACT and SPT data at 150 GHz can be fitted to high (sub-$\mu\mathrm{K}^2$) accuracy *without* kSZ and tSZ x CIB templates, yet we expect a kSZ contribution of at least the amplitude shown in Fig. 6, and larger if we account for patchy reionization (see, e.g., Knox 2003, and references therein) and references therein). This implies a cancellation of the kSZ and tSZ x CIB contributions at 150 GHz (Addison et al. 2012c), as discussed in greater detail in Planck Collaboration XVI (2013).

4. Combined cross-spectra and consistency checks

The large number of cross spectra in a detector-by-detector power spectrum analysis allows for a number of internal consistency checks of the data. Within a frequency band, we expect to see exactly the same sky signals (primordial CMB, Galactic, and extra-galactic foregrounds), and so any intra-frequency residuals will reflect instrumental systematics, for example beam errors, 'gain' fluctuations, and band-pass mismatch. In contrast, inter-frequency residuals are harder to analyse because the sky signals vary with frequency. An accurate model of the unresolved foregrounds is therefore required to assess inter-frequency residuals. Furthermore, as we will show below, the scatter caused by chance CMB–foreground cross-correlations can dominate the inter-frequency residuals. For a precision experiment such as *Planck*, where the power spectra are expected to be signal dominated over a wide multipole range, intra- and inter-frequency residuals provide a powerful way of assessing possible systematic errors. It is essential that contributions of systematic errors to both types of residual are small enough that they have negligible impact on cosmological parameter analysis.

Figure 7 shows a selection of temperature cross-spectra and estimates of the analytic covariance matrices, together with the best-fit cosmological model described in Sect. 5. Unresolved foregrounds have been subtracted using the best-fit foreground parameters of the model described in Sect. 3. The scatter varies substantially between cross-spectra, reflecting differences in the instrument noise and effective resolution of different detector combinations. The analytic error model summarized in Appendix A.2 is indicated, modified by the non-white noise correction. This model provides an excellent description of the scatter seen in the data, over the full multipole range shown in plots, with an accuracy of a few percent or better.

4.1. *Intra-frequency residuals*

In this section we analyse the intra-frequency residuals at 143 and 217 GHz. There are $N_{\mathrm{spec}} = 10$ cross-spectra at 143 GHz and 15 cross-spectra at 217 GHz[6]. At each frequency, we solve for multiplicative ('effective' calibration) coefficients, y_i, that minimise

$$\chi^2 = \sum_\ell \sum_{ij,j>i} (y_i y_j \hat{C}_\ell^{ij} - \langle \hat{C}_\ell \rangle)^2, \tag{18}$$

---

[6] there is only one cross-spectrum at 100 GHz

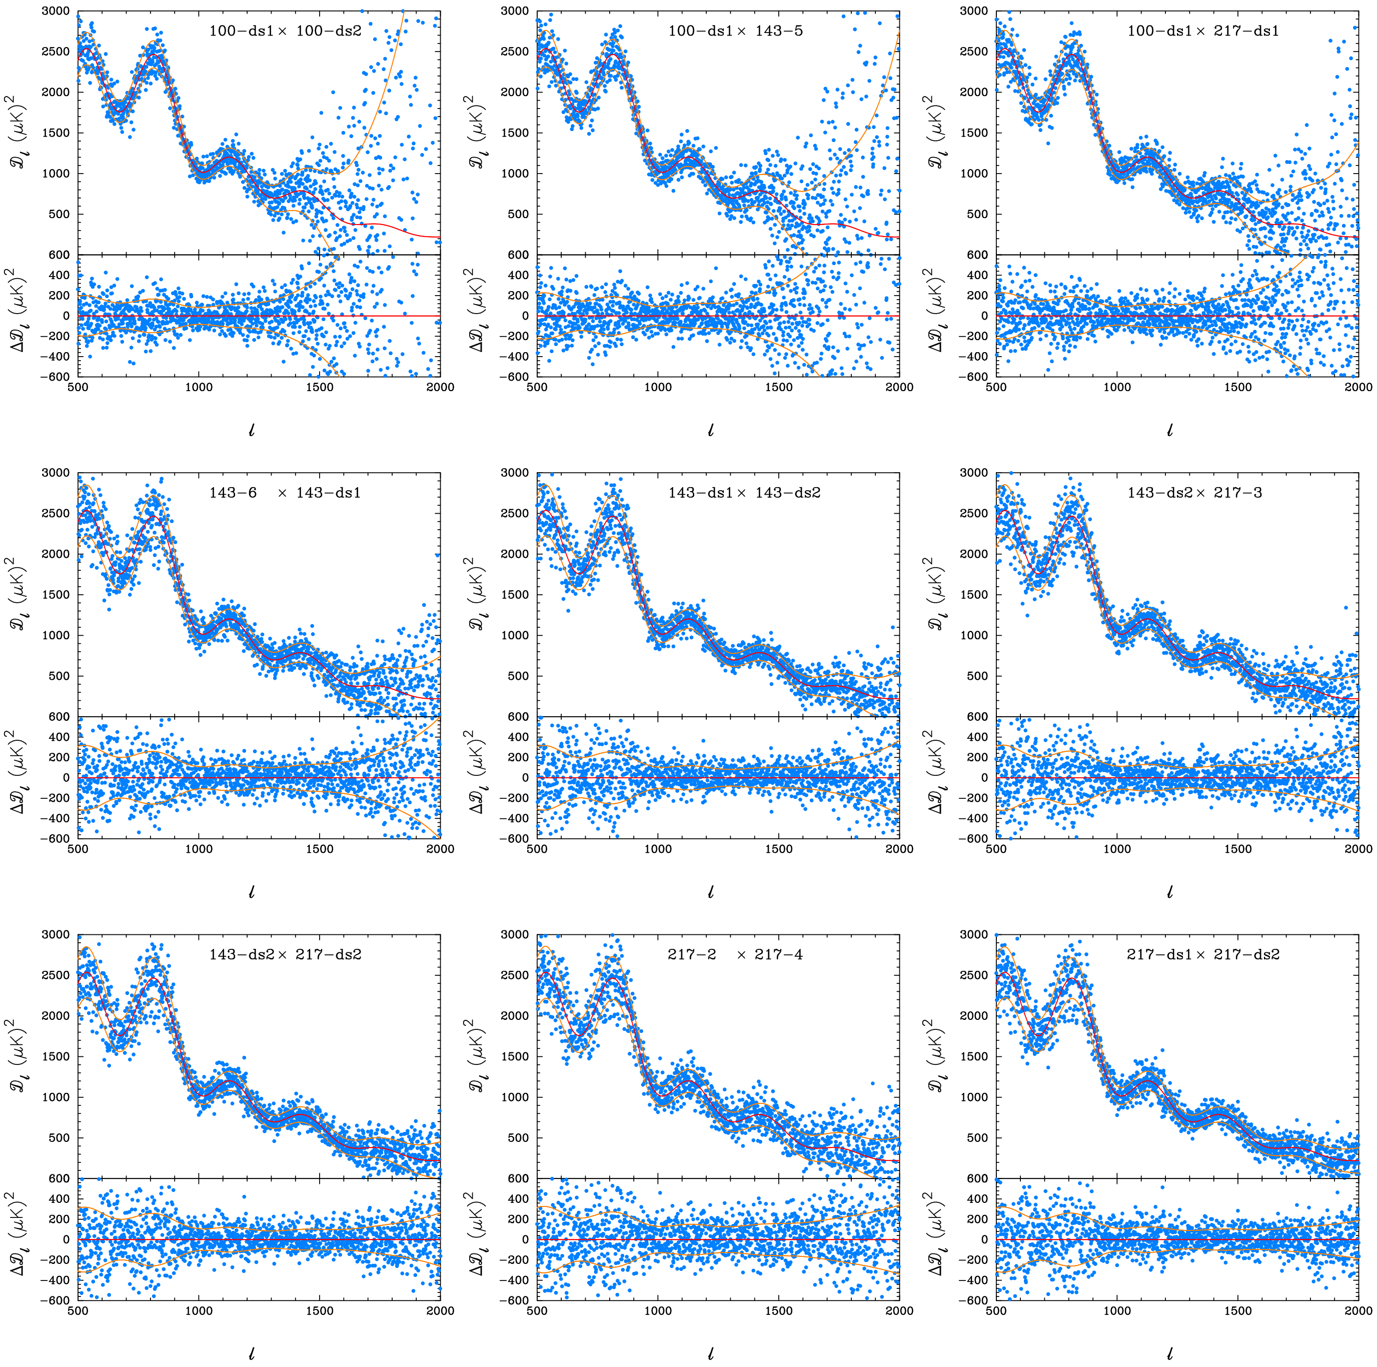

Figure 7. A selection of cross spectra from individual detectors, with the best fit unresolved foreground model subtracted. The best-fit six parameter ΛCDM model is shown, with $\pm 1\sigma$ errors determined from the diagonal components of the analytic covariance matrices. The lower panel in each plot shows the residuals with respect to the model.

where

$$\langle \hat{C}_\ell \rangle = \frac{1}{N_{\rm spec}} \sum_{ij,j>i} y_i y_j \hat{C}^{ij}_\ell, \qquad (19)$$

subject to the constraint that $y_1 = 1$ (where $i = 1$ corresponds to detector 5 at 143 GHz and detector 1 at 217 GHz). Note that the power spectra in Eq. 18 and 19 are corrected for beam transfer functions. To minimise the possible impact of beam errors and noise, we restrict the sum in Eq. 18 to the multipole range $50 \leq \ell \leq 500$ where the spectra are signal dominated. Numerical values for the calibration coefficients are given in Table 3, using mask CL31. The calibration factors are insensitive to the choice of mask or multipole range.

The results of Table 3 show that effective calibration factors of ∼0.2 % are quite typical for HFI maps, in the 100 – 217 GHz frequency range. These recalibrations are of the order of magnitude of the statistical errors of the calibrations on dipole (see Planck Collaboration VIII 2013, Table 2). Note that the data are corrected for individual bolometer time transfer functions (TTFs; Planck Collaboration VII 2013). For each detector, the TTF model is tuned to minimise survey differences and by construction normalised to unity at the spin frequency of the satellite

Table 3. Map calibration coefficients.

map	y_i	map	y_i
143-5	1.0000	217-1	1.0000
143-6	0.9988	217-2	0.9992
143-7	0.9980	217-3	0.9981
-	-	217-4	0.9985
143-ds1	0.9990	217-ds1	0.9982
143-ds2	0.9994	217-ds2	0.9975

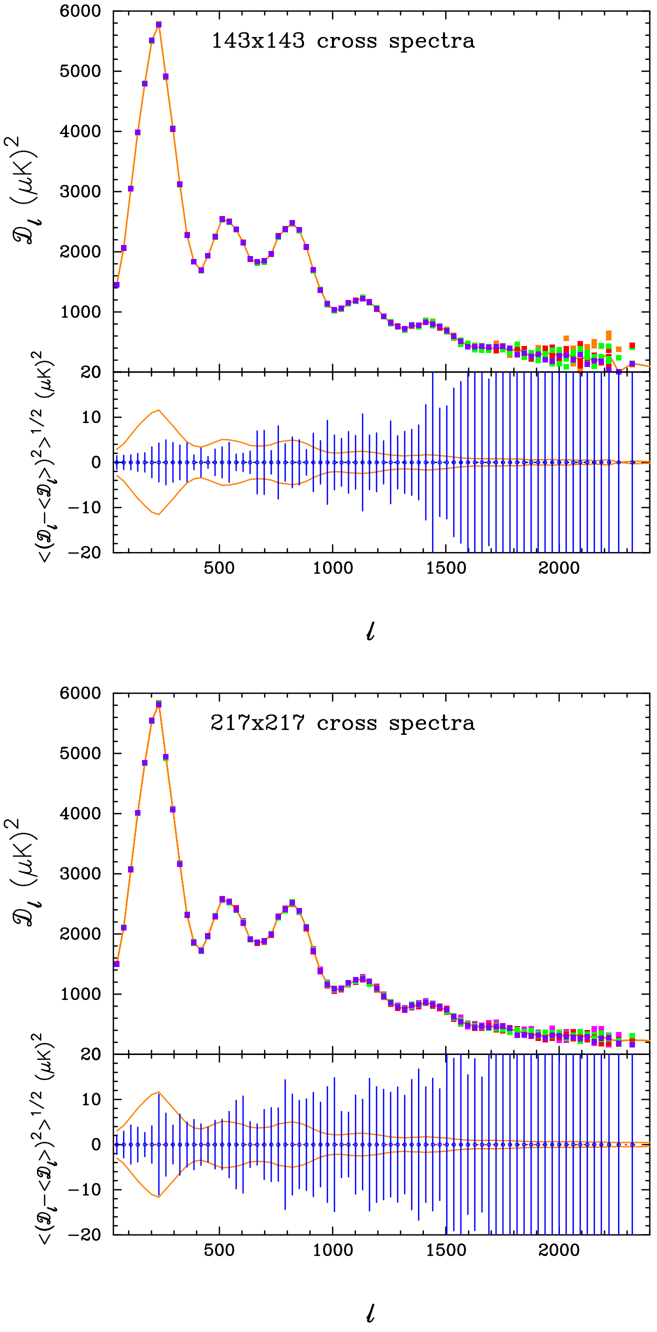

Figure 8. Cross spectra corrected for the beam and effective calibration, together with the mean power spectrum. *Top*: The 10 cross spectra at 143 GHz. *Bottom*: The 15 cross spectra at 217 GHz, with SWB×SWB spectra (magenta), SWB×ds1 (red), SWB×ds2 (green), ds1×ds2 (purple). The power spectra are distinguishable only at high multipoles where the data become noise dominated. The lower panels show the dispersion of the cross spectra around the mean, together with a ±0.2% calibration error.

(0.01666 Hz) to preserve the dipole calibration. The consistency of intra-frequency power spectrum residuals therefore provides a test of the consistency of the TTFs in addition to the beam transfer functions.

Figure 8 shows the remarkable consistency of the power spectra at each frequency. The upper panels show the spectra corrected for the beam and effective calibration, together with the mean cross spectra. The lower panels show the dispersion around the mean. In the signal dominated regime the cross spectra show an RMS dispersion of a few μK^2 (in bands of width $\Delta\ell = 31$), i.e., the band-averaged spectra are consistent to an accuracy of ~0.1–0.2 %. This excess scatter has negligible impact on cosmological parameter analysis.

The residuals of the cross spectra in band averages of width $\Delta\ell \sim 61$ are shown in Fig. 9, before and after correction for the effective intra-frequency calibrations. The reduction in scatter after correction is evident at $\ell \lesssim 500$, and the residual scatter is consistent with instrument noise and beam errors. At 217 GHz, beam errors dominate over noise at multipoles $\lesssim 1000$. There is no evidence that the excess scatter is caused by a small number of 'anomalous' detectors.

4.2. Inter-frequency residuals

The results of the previous section show that the intra-frequency cross-spectra between detector/detector sets are consistent to within a few μK^2 at multipoles $\ell \lesssim 1000$. In a likelihood analysis, there is therefore little loss of information in compressing the power spectra for each distinct frequency combination, as opposed to retaining the spectra for each map pair. This compression greatly reduces the size of the data vector and its covariance matrix, and speeds up the likelihood computation at high multipoles. In this section we inter-compare the residuals of these compressed power spectra.

One might naïvely expect that with accurate foreground modelling, the inter-frequency residuals in the signal dominated regime should be reduced to levels comparable to those seen in the intra-frequency comparisons described in the previous section. This is incorrect. Figure 10 shows power spectrum differences between the cosmologically significant spectra for *Planck* at high multipoles (143×143, 143×217, 217×217). In this figure, which is independent of the cosmological model, the best-fit unresolved foreground model has been subtracted from each spectrum, and relative calibration factors have been applied. Residual beam, calibration and unresolved foreground errors would show up in this figure as large-scale smooth residuals.

In fact, we see small-scale residuals at multipoles $\ell \lesssim 800$ which are considerably larger than expected from instrumental noise. This excess scatter arises from chance CMB–foreground cross-correlations. Even if the foreground contamination is much smaller than the CMB, chance cross-correlations can produce scatter in the inter-frequency power spectra that is large in the signal dominated regime. We demonstrate in Appendix C that the observed scatter can be predicted quantitatively.

At high enough multipoles, instrument noise, beam errors, and errors in foreground modelling dominate the inter-frequency residuals. A complete analysis of inter-frequency residuals therefore requires the full likelihood machinery and MCMC analysis to determine foreground, beam and calibration parameters. We will therefore revisit the inter-frequency residuals in the following sections.

5. Reference results of the high-ℓ likelihood

In this section we study the high-ℓ *Planck* likelihood, and present the power spectrum and parameters derived from the baseline likelihood for the basic six-parameter ΛCDM model. In order to break the well-known degeneracy between the optical depth, τ, and the scalar index of scalar perturbations, $n_{\rm s}$, we adopt a Gaussian prior on τ (inspired from WMAP7 data, i.e., 0.088 ± 0.015, see Komatsu et al. 2011) instead of the low-ℓ likelihood at $\ell < 50$. We return to the global *Planck* results after introducing the low-ℓ likelihood.

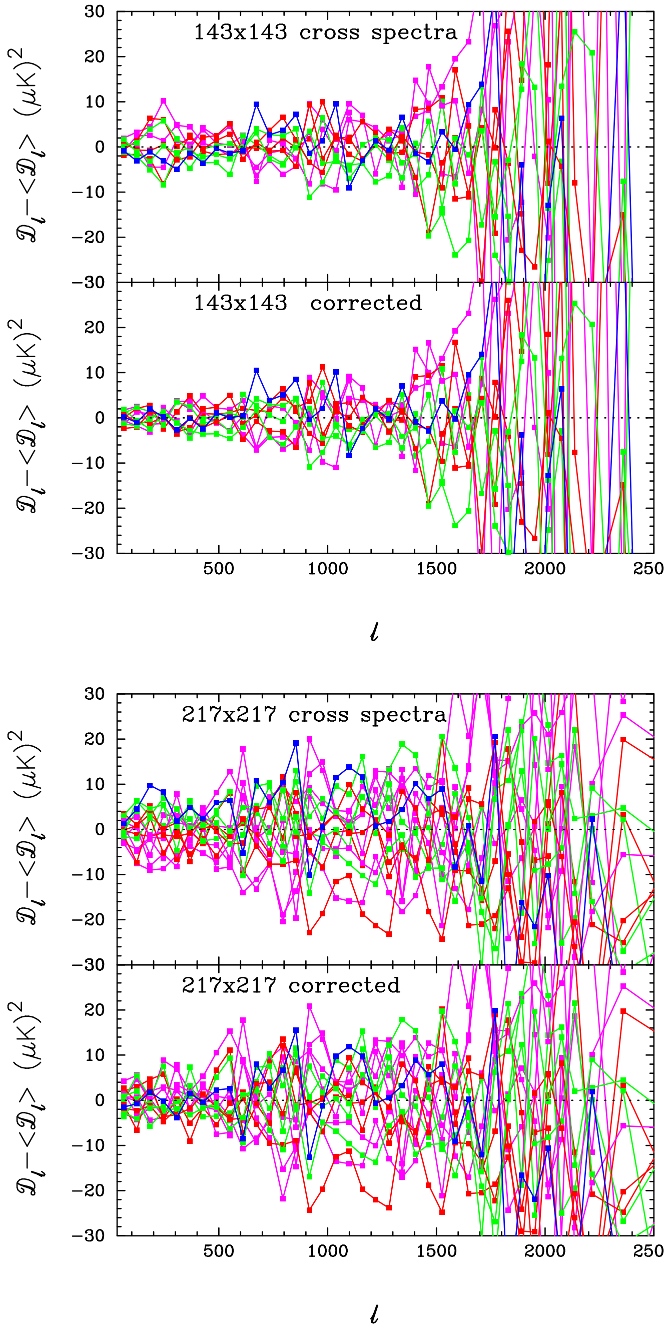

Figure 9. Cross spectra for the 143 GHz (*top*) and 217 GHz (*bottom*) channels, as in Fig. 8, before correction for multiplicative intra-frequency calibration coefficients (above), and after correction (below).

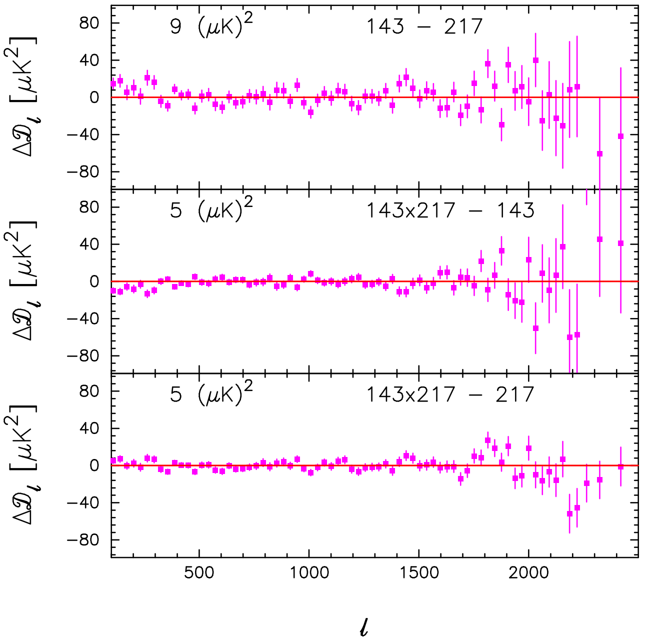

Figure 10. Differences between the 143 × 143, 143 × 217, and 217 × 217 cross spectra (plotted in bins of width $\delta\ell \approx 31$). The best-fit model for unresolved foregrounds has been subtracted from each spectrum. The numbers list the dispersions over the multipole range $800 \leq \ell \leq 1500$.

We choose separate masks for each frequency map to minimise Galactic foreground emission. First, since the HFI data are signal-dominated at $\lesssim 500$, and diffuse Galactic emission is low at 100 GHz outside the CL49 mask, there is little to be gained from analyzing the dustier 217 GHz with the same region; it contains no new information about the primordial CMB. At higher multipoles, $\ell \gtrsim 500$, we use the CL31 mask, optimizing the sky coverage while ensuring a low amplitude of small-scale Galactic emission relative to the isotropic unresolved foregrounds and the primordial CMB. In addition, we tune the multipole range for each frequency to mitigate Galactic foreground contamination and beam errors.

The choices of masks and angular ranges used in the high-ℓ likelihood are summarised in Table 4, together with basic χ^2 statistics with respect to the minimal ΛCDM model per cross-spectrum and combined. The 100 GHz cross-spectrum is computed over the largest sky fraction, a total of 49% of the sky, and measures the largest scales. On the other hand, it has lowest resolution, and it is therefore only used for $\ell \leq 1200$. The 143 GHz cross-spectrum has higher resolution, and is used for $\ell \leq 2000$. Finally, the 217 GHz cross-spectrum has the highest resolution, but also the most Galactic dust contamination, and is therefore evaluated from only 31% of the sky, but including an angular range of $500 \leq \ell \leq 2500$.

Given these masks and angular ranges, we compute the angular power spectra and covariance matrices, and construct the `CamSpec` likelihood. The angular power spectra for each fre-