Planck 2015 results

Astronomy & Astrophysics

Special feature

Planck 2015 results

XVI. Isotropy and statistics of the CMB

Planck Collaboration: P. A. R. Ade[90], N. Aghanim[60], Y. Akrami[65,104], P. K. Aluri[55], M. Arnaud[76], M. Ashdown[72,6], J. Aumont[60], C. Baccigalupi[89], A. J. Banday[101,9,★], R. B. Barreiro[67], N. Bartolo[32,68], S. Basak[89], E. Battaner[102,103], K. Benabed[61,100], A. Benoît[58], A. Benoit-Lévy[26,61,100], J.-P. Bernard[101,9], M. Bersanelli[35,49], P. Bielewicz[86,9,89], J. J. Bock[69,11], A. Bonaldi[70], L. Bonavera[67], J. R. Bond[8], J. Borrill[14,95], F. R. Bouchet[61,93], F. Boulanger[60], M. Bucher[1], C. Burigana[48,33,50], R. C. Butler[48], E. Calabrese[98], J.-F. Cardoso[77,1,61], B. Casaponsa[67], A. Catalano[78,75], A. Challinor[64,72,12], A. Chamballu[76,16,60], H. C. Chiang[29,7], P. R. Christensen[87,38], S. Church[97], D. L. Clements[56], S. Colombi[61,100], L. P. L. Colombo[25,69], C. Combet[78], D. Contreras[24], F. Couchot[74], A. Coulais[75], B. P. Crill[69,11], M. Cruz[21], A. Curto[67,6,72], F. Cuttaia[48], L. Danese[89], R. D. Davies[70], R. J. Davis[70], P. de Bernardis[34], A. de Rosa[48], G. de Zotti[45,89], J. Delabrouille[1], F.-X. Désert[54], J. M. Diego[67], H. Dole[60,59], S. Donzelli[49], O. Doré[69,11], M. Douspis[60], A. Ducout[61,56], X. Dupac[39], G. Efstathiou[64], F. Elsner[26,61,100], T. A. Enßlin[83], H. K. Eriksen[65], Y. Fantaye[37], J. Fergusson[12], R. Fernandez-Cobos[67], F. Finelli[48,50], O. Forni[101,9], M. Frailis[47], A. A. Fraisse[29], E. Franceschi[48], A. Frejsel[87], A. Frolov[92], S. Galeotta[47], S. Galli[71], K. Ganga[1], C. Gauthier[1,82], T. Ghosh[60], M. Giard[101,9], Y. Giraud-Héraud[1], E. Gjerløw[65], J. González-Nuevo[20,67], K. M. Górski[69,105], S. Gratton[72,64], A. Gregorio[36,47,53], A. Gruppuso[48], J. E. Gudmundsson[29], F. K. Hansen[65], D. Hanson[84,69,8], D. L. Harrison[64,72], S. Henrot-Versillé[74], C. Hernández-Monteagudo[13,83], D. Herranz[67], S. R. Hildebrandt[69,11], E. Hivon[61,100], M. Hobson[6], W. A. Holmes[69], A. Hornstrup[17], W. Hovest[83], Z. Huang[8], K. M. Huffenberger[27], G. Hurier[60], A. H. Jaffe[56], T. R. Jaffe[101,9], W. C. Jones[29], M. Juvela[28], E. Keihänen[28], R. Keskitalo[14], J. Kim[83], T. S. Kisner[80], J. Knoche[83], M. Kunz[18,60,3], H. Kurki-Suonio[28,44], G. Lagache[5,60], A. Lähteenmäki[2,44], J.-M. Lamarre[75], A. Lasenby[6,72], M. Lattanzi[33], C. R. Lawrence[69], R. Leonardi[39], J. Lesgourgues[62,99], F. Levrier[75], M. Liguori[32,68], P. B. Lilje[65], M. Linden-Vørnle[17], H. Liu[87,38], M. López-Caniego[39,67], P. M. Lubin[30], J. F. Macías-Pérez[78], G. Maggio[47], D. Maino[35,49], N. Mandolesi[48,33], A. Mangilli[60,74], D. Marinucci[37], M. Maris[47], P. G. Martin[8], E. Martínez-González[67], S. Masi[34], S. Matarrese[32,68,42], P. McGehee[57], P. R. Meinhold[30], A. Melchiorri[34,51], L. Mendes[39], A. Mennella[35,49], M. Migliaccio[64,72], K. Mikkelsen[65], S. Mitra[55,69], M.-A. Miville-Deschênes[60,8], D. Molinari[67,48], A. Moneti[61], L. Montier[101,9], G. Morgante[48], D. Mortlock[56], A. Moss[91], D. Munshi[90], J. A. Murphy[85], P. Naselsky[87,38], F. Nati[29], P. Natoli[33,4,48], C. B. Netterfield[22], H. U. Nørgaard-Nielsen[17], F. Noviello[70], D. Novikov[81], I. Novikov[87,81], C. A. Oxborrow[17], F. Paci[89], L. Pagano[34,51], F. Pajot[60], N. Pant[55], D. Paoletti[48,50], F. Pasian[47], G. Patanchon[1], T. J. Pearson[11,57], O. Perdereau[74], L. Perotto[78], F. Perrotta[89], V. Pettorino[43], F. Piacentini[34], M. Piat[1], E. Pierpaoli[25], D. Pietrobon[69], S. Plaszczynski[74], E. Pointecouteau[101,9], G. Polenta[4,46], L. Popa[63], G. W. Pratt[76], G. Prézeau[11,69], S. Prunet[61,100], J.-L. Puget[60], J. P. Rachen[23,83], R. Rebolo[66,15,19], M. Reinecke[83], M. Remazeilles[70,60,1], C. Renault[78], A. Renzi[37,52], I. Ristorcelli[101,9], G. Rocha[69,11], C. Rosset[1], M. Rossetti[35,49], A. Rotti[55], G. Roudier[1,75,69], J. A. Rubiño-Martín[66,19], B. Rusholme[57], M. Sandri[48], D. Santos[78], M. Savelainen[28,44], G. Savini[88], D. Scott[24], M. D. Seiffert[69,11], E. P. S. Shellard[12], T. Souradeep[55], L. D. Spencer[90], V. Stolyarov[6,96,73], R. Stompor[1], R. Sudiwala[90], R. Sunyaev[83,94], D. Sutton[64,72], A.-S. Suur-Uski[28,44], J.-F. Sygnet[61], J. A. Tauber[40], L. Terenzi[41,48], L. Toffolatti[20,67,48], M. Tomasi[35,49], M. Tristram[74], T. Trombetti[48], M. Tucci[18], J. Tuovinen[10], L. Valenziano[48], J. Valiviita[28,44], B. Van Tent[79], P. Vielva[67], F. Villa[48], L. A. Wade[69], B. D. Wandelt[61,100,31], I. K. Wehus[69], D. Yvon[16], A. Zacchei[47], J. P. Zibin[24], and A. Zonca[30]

(Affiliations can be found after the references)

ABSTRACT

We test the statistical isotropy and Gaussianity of the cosmic microwave background (CMB) anisotropies using observations made by the *Planck* satellite. Our results are based mainly on the full *Planck* mission for temperature, but also include some polarization measurements. In particular, we consider the CMB anisotropy maps derived from the multi-frequency *Planck* data by several component-separation methods. For the temperature anisotropies, we find excellent agreement between results based on these sky maps over both a very large fraction of the sky and a broad range of angular scales, establishing that potential foreground residuals do not affect our studies. Tests of skewness, kurtosis, multi-normality, N-point functions, and Minkowski functionals indicate consistency with Gaussianity, while a power deficit at large angular scales is manifested in several ways, for example low map variance. The results of a peak statistics analysis are consistent with the expectations of a Gaussian random field. The "Cold Spot" is detected with several methods, including map kurtosis, peak statistics, and mean temperature profile. We thoroughly probe the large-scale dipolar power asymmetry, detecting it with several independent tests, and address the subject of a posteriori correction. Tests of directionality suggest the presence of angular clustering from large to small scales, but at a significance that is dependent on the details of the approach. We perform the first examination of polarization data, finding the morphology of stacked peaks to be consistent with the expectations of statistically isotropic simulations. Where they overlap, these results are consistent with the *Planck* 2013 analysis based on the nominal mission data and provide our most thorough view of the statistics of the CMB fluctuations to date.

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

★ Corresponding author: A. J. Banday, e-mail: anthony.banday@irap.omp.eu

1. Introduction

This paper, one of a set associated with the 2015 release of data from the *Planck*[1] mission (Planck Collaboration I 2016), describes a set of studies undertaken to determine the statistical properties of both the temperature and polarization anisotropies of the cosmic microwave background (CMB).

The standard cosmological model is described well by the Friedmann-Lemaître-Robertson-Walker solution of the Einstein field equations. This model is characterized by a homogeneous and isotropic background metric and a scale factor of the expanding Universe. It is hypothesized that at very early times the Universe went through a period of accelerated expansion, the so-called "cosmological inflation", driven by a hypothetical scalar field, the "inflaton". During inflation the Universe behaves approximately as a de Sitter space, providing the conditions by which some of its present properties can be realized and specifically relaxing the problem of initial conditions. In particular, the seeds that gave rise to the present large-scale matter distribution via gravitational instability originated as quantum fluctuations of the inflaton about its vacuum state. These fluctuations in the inflaton produce energy density perturbations that are distributed as a statistically homogeneous and isotropic Gaussian random field. Linear theory relates those perturbations to the temperature and polarization anisotropies of the CMB, implying a distribution for the anisotropies very close to that of a statistically isotropic Gaussian random field.

The aim of this paper is to use the full mission *Planck* data to test the Gaussianity and isotropy of the CMB as measured in both intensity and, in a more limited capacity, polarization. Testing these fundamental properties is crucial for the validation of the standard cosmological scenario, and has profound implications for our understanding of the physical nature of the Universe and the initial conditions of structure formation. Moreover, the confirmation of the statistically isotropic and Gaussian nature of the CMB is essential for justifying the corresponding assumptions usually made when estimating the CMB power spectra and other quantities to be obtained from the *Planck* data. Indeed, the isotropy and Gaussianity of the CMB anisotropies are implicitly assumed in critical science papers from the 2015 release, in particular those describing the likelihood and the derivation of cosmological parameter constraints (Planck Collaboration XI 2016; Planck Collaboration XIII 2016). Conversely, if the detection of significant deviations from these assumptions cannot be traced to known systematic effects or foreground residuals, the presence of which should be diagnosed by the statistical tests set forth in this paper, this would necessitate a major revision of the current methodological approaches adopted in deriving the mission's many science results.

Well-understood physical processes due to the integrated Sachs-Wolfe (ISW) effect (Planck Collaboration XVII 2014; Planck Collaboration XXI 2016) and gravitational lensing (Planck Collaboration XIX 2014; Planck Collaboration XV 2016) lead to secondary anisotropies that exhibit marked deviation from Gaussianity. In addition, Doppler boosting, due to our motion with respect to the CMB rest frame, induces

both a dipolar modulation of the temperature anisotropies and an aberration that corresponds to a change in the apparent arrival directions of the CMB photons (Challinor & van Leeuwen 2002). Both of these effects are aligned with the CMB dipole, and were detected at a statistically significant level on small angular scales in Planck Collaboration XXVII (2014). Beyond these, Planck Collaboration XXIII (2014, hereafter PCIS13) established that the *Planck* 2013 data set showed little evidence for non-Gaussianity, with the exception of a number of CMB temperature anisotropy anomalies on large angular scales that confirmed earlier claims based on WMAP data. Moreover, given that the broader frequency coverage of the *Planck* instruments allowed improved component separation methods to be applied in the derivation of foreground-cleaned CMB maps, it was generally considered that the case for anomalous features in the CMB had been strengthened. Hence, such anomalies have attracted considerable attention in the community, since they could be the visible traces of fundamental physical processes occurring in the early Universe.

However, the literature also supports an ongoing debate about the significance of these anomalies. The central issue in this discussion is connected with the role of a posteriori choices – whether interesting features in the data bias the choice of statistical tests, or if arbitrary choices in the subsequent data analysis enhance the significance of the features. Indeed, the WMAP team (Bennett et al. 2011) base their rejection of the presence of anomalies in the CMB on such arguments. Of course, one should attempt to correct for any choices that were made in the process of detecting an anomaly. However, in the absence of an alternative model for comparison to the standard Gaussian, statistically isotropic one adopted to quantify significance, this is often simply not possible. In this work, whilst it is recognized that care must be taken in the assessment of significance, we proceed on the basis that allowing a posteriori reasoning permits us to challenge the limits of our existing knowledge (Pontzen & Peiris 2010). That is, by focusing on specific properties of the observed data that are shown to be empirically interesting, we may open up new paths to a better theoretical understanding of the Universe. We will clearly describe the methodology applied to the data, and attempt to study possible links among the anomalies in order to search for a physical interpretation.

The analysis of polarization data introduces a new opportunity to explore the statistical properties of the CMB sky, including the possibility of improvement of the significance of detection of large-scale anomalies. However, this cannot be fully included in the current data assessment, since the component-separation products in polarization are high-pass filtered to remove large angular scales (Planck Collaboration IX 2016), owing to the persistence of significant systematic artefacts originating in the High Frequency Instrument (HFI) data (Planck Collaboration VII 2016; Planck Collaboration VIII 2016). In addition, limitations of the simulations with which the data are to be compared (Planck Collaboration XII 2016), in particular a significant mismatch in noise properties, limit the extent to which any polarization results can be included. Therefore, we only present a stacking analysis of the polarized data, although this is a significant extension of previous approaches found in the literature.

With future *Planck* data releases, it will be important to determine in more detail whether there are any pecularities in the CMB polarization, and if so, whether they are related to existing features in the CMB temperature field. Conversely, the absence of any corresponding features in polarization might imply that

---

[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).

the temperature anomalies (if they are not simply flukes) could be due to a secondary effect such as the ISW effect, or alternative scenarios in which the anomalies arise from physical processes that do not correlate with the temperature, e.g., texture or defect models. Either one of these possible outcomes could yield a breakthrough in understanding the nature of the CMB anomalies. Of course, there also remains the possibility that anomalies may be found in the polarization data that are unrelated to existing features in the temperature measurements.

Following the approach established in Planck Collaboration XXIII (2014), throughout this paper we quantify the significance of a test statistic in terms of the *p*-value. This is the probability of obtaining a test statistic at least as extreme as the observed one, under the assumption that the null hypothesis (i.e., primordial Gaussianity and isotropy of the CMB) is true. In some tests, where it is clearly justified to only use a one-tailed probability, the *p*-value is replaced by the corresponding upper- or lower-tail probability.

This paper covers all relevant aspects related to the phenomenological study of the statistical isotropy and Gaussian nature of the CMB measured by the *Planck* satellite. Specific theoretically-motivated model constraints on isotropy or non-Gaussianity, as might arise from non-standard inflationary models, the geometry and topology of the Universe, and primordial magnetic fields are provided in the companion papers (Planck Collaboration XVII 2016; Planck Collaboration XX 2016; Planck Collaboration XVIII 2016; Planck Collaboration XIX 2016). The paper is organized as follows. Section 2 summarizes the *Planck* full mission data used for the analyses, and important limitations of the polarization maps that are studied. Section 3 describes the characteristics of the simulations that constitute our reference set of Gaussian sky maps representative of the null hypothesis. In Sect. 4 the null hypothesis is tested with a number of standard tests that probe different aspects of non-Gaussianity. Several important anomalous features of the CMB sky, originally detected with the WMAP data and subsequently confirmed in PCIS13, are reassessed in Sect. 5. Aspects of the CMB fluctuations specifically related to dipolar asymmetry are examined in Sect. 6. The sensitivity of the results for a number of statistical tests to the sky fraction is examined in Sect. 7. Section 8 presents tests of the statistical nature of the polarization signal observed by *Planck* using a local analysis of stacked patches of the sky. Finally, Sect. 9 provides the main conclusions of the paper.

2. Data description

In this paper, we use data from the *Planck*-2015 full mission data release. This contains approximately 29 months of data for the HFI and 50 months for the Low Frequency Instrument (LFI). The release includes sky maps at nine frequencies in intensity (seven in polarization), with corresponding "half-mission" maps that are generated by splitting the full-mission data sets in various ways. The maps are provided in `HEALPix` format (Górski et al. 2005)[2], with a pixel size defined by the N_{side} parameter. This set of maps allows a variety of consistency checks to be made, together with estimates of the instrumental noise contributions to our analyses and limits on time-varying systematic artefacts. Full details are provided in a series of companion papers (Planck Collaboration II 2016; Planck Collaboration III 2016; Planck Collaboration IV 2016; Planck Collaboration V 2016; Planck Collaboration VI 2016; Planck Collaboration VII 2016; Planck Collaboration VIII 2016).

---

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

Our main results are based on estimates of the CMB generated by four distinct component-separation algorithms – `Commander`, `NILC`, `SEVEM`, and `SMICA` – as described in Planck Collaboration IX (2016). These effectively combine the raw *Planck* frequency maps in such a way as to minimize foreground residuals from diffuse Galactic emission. Note that the additional information in the full mission data set allows us to improve the reconstruction noise levels by roughly a factor of 2 (in temperature) as compared to the *Planck*-2013 nominal mission data release. The CMB intensity maps were derived using all channels, from 30 to 857 GHz, and provided at a common angular resolution of 5′ FWHM and $N_{\text{side}} = 2048$. The intensity maps are only partially corrected for the second order temperature quadrupole (Kamionkowski & Knox 2003). Therefore, where appropriate, the component-separated maps should be corrected for the residual contribution (Notari & Quartin 2015), specifically as described in Planck Collaboration IX (2016). The polarization solutions include all channels sensitive to polarization, from 30 to 353 GHz, at a resolution of 10′ FWHM and $N_{\text{side}} = 1024$. Possible residual emission is then mitigated in our analyses by the use of sky-coverage masks, provided for both intensity and polarization.

Since in some cases it is important to study the frequency dependence of the cosmological signal, either to establish its primordial origin or to test for the frequency dependence associated with specific effects such as Doppler boosting (see Sect. 6.4), we also consider the foreground-cleaned versions of the 100, 143, and 217 GHz sky maps generated by the `SEVEM` algorithm (Planck Collaboration IX 2016), hereafter referred to as `SEVEM-100`, `SEVEM-143`, and `SEVEM-217`, respectively.

For the present release, a post-processing high-pass-filtering has been applied to the CMB polarization maps in order to mitigate residual large-scale systematic errors in the HFI channels (Planck Collaboration VII 2016). The filter results in the elimination of structure in the maps on angular scales larger than about 10°, and a weighted suppression of power down to scales of 5°, below which the maps remain unprocessed.

Lower-resolution versions of these data sets are also used in the analyses presented in this paper. The downgrading procedure for maps is to decompose them into spherical harmonics on the full sky at the input `HEALPix` resolution. The spherical harmonic coefficients, $a_{\ell m}$, are then convolved to the new resolution using

$$a_{\ell m}^{\text{out}} = \frac{b_{\ell}^{\text{out}} p_{\ell}^{\text{out}}}{b_{\ell}^{\text{in}} p_{\ell}^{\text{in}}} a_{\ell m}^{\text{in}}, \tag{1}$$

where b_{ℓ} is the beam transfer function, p_{ℓ} is the `HEALPix` pixel window function, and the "in" and "out" superscripts denote the input and output resolutions. They are then synthesized into a map directly at the output `HEALPix` resolution. Masks are downgraded in a similar way. The binary mask at the starting resolution is first downgraded like a temperature map. The smooth downgraded mask is then thresholded by setting pixels where the value is less than 0.9 to zero and all others to unity in order to make a binary mask. Table 1 lists the N_{side} and FWHM values defining the resolution of these maps, together with the different masks and their sky coverages that accompany the signal maps. In general, we make use of standardized masks that are the union of those associated with the individual component-separation methods.

As recommended in Planck Collaboration IX (2016), the mask UT78 is adopted for all high-resolution analyses of temperature data. UTA76 is an extended version of this mask more suitable for some non-Gaussianity studies. The mask preferred

Table 1. Standardized data sets used in this paper.

N_{side}	FWHM [arcmin]	Mask	Unmasked pixels [%]
2048	5	UT78	77.6
2048	5	UTA76	76.1
1024	10	$UT_{1024}76$	75.6
512	20	$UT_{512}74$	73.7
256	40	$UT_{256}73$	72.5
128	80	$UT_{128}70$	69.7
64	160	$UT_{64}67$	67.0
32	320	$UT_{32}64$	63.8
16	640	$UT_{16}58$	58.4
1024	10	UPB77	77.4

Notes. The resolutions of the sky maps used are defined in terms of the N_{side} parameter and corresponding FWHM of the Gaussian beam with which they are convolved. The corresponding common masks and the fraction of unmasked pixels used for analysis are also specified.

for polarization studies, UPB77, is again the union of those determined for each component separation method, but in addition the polarized point sources detected at each frequency channel are excluded. These masks are then downgraded for lower-resolution studies. As a consequence of the common scheme applied in order to generate such low-resolution masks, they are generally more conservative than the corresponding ones used in the 2013 analyses.

In what follows, we will undertake analyses of the data at a given resolution denoted by a specific N_{side} value. Unless otherwise stated, this implies that the data have been smoothed to a corresponding FWHM as described above, and a standardized mask employed. Irrespective of the resolution in question, we will then often simply refer to the latter as the "common mask".

3. Simulations

The results presented in this paper are derived using the extensive full focal plane (FFP8) simulations described in Planck Collaboration XII (2016). Of most importance are the Monte Carlo (MC) simulations that provide the reference set of Gaussian sky maps used for the null tests employed here. They also form the basis of any debiasing in the analysis of the real data as required by certain statistical methods.

The simulations include both CMB signal and instrumental noise realizations that capture important characteristics of the *Planck* scanning strategy, telescope, detector responses, and data reduction pipeline over the full mission period. In particular, the signal realizations include FEBeCoP (Mitra et al. 2011) beam convolution at each of the *Planck* frequencies, and are propagated through the various component-separation pipelines using the same weights as derived from the *Planck* full mission data analysis.

The FFP8 fiducial CMB power spectrum has been adopted from our best estimate of the cosmological parameters from the first *Planck* data release (Planck Collaboration I 2014). This corresponds to a cosmology with baryon density given by $\omega_b = \Omega_b h^2 = 0.0222$, cold dark matter (CDM) density $\omega_c = \Omega_c h^2 = 0.1203$, neutrino energy density $\omega_\nu = \Omega_\nu h^2 = 0.00064$, density parameter for the cosmological constant $\Omega_\Lambda = 0.6823$, Hubble parameter $H_0 = 100h \, km \, s^{-1} \, Mpc^{-1}$ with $h = 0.6712$, spectral index of the power spectrum of the primordial curvature perturbation $n_s = 0.96$, and amplitude of the primordial power spectrum (at $k = 0.05 \, Mpc^{-1}$) $A_s = 2.09 \times 10^{-9}$, and with

the Thomson optical depth through reionization defined to be $\tau = 0.065$. Each realization of the CMB sky is generated including lensing, Rayleigh scattering, and Doppler boosting effects, the latter two of which are frequency-dependent. Unfortunately, the aberration contribution to the Doppler boost was erroneously omitted from the simulations, but, with possible exceptions described in Sect. 6, this does not lead to any significant impact on the results in this paper. A second order temperature quadrupole (Kamionkowski & Knox 2003) is added to each simulation with an amplitude corresponding to the residual uncorrected contribution present in the observed data, as described in Planck Collaboration XII (2016).

However, the *Planck* maps were effectively renormalized by approximately 2% to 3% in power in the time between the generation of the FFP8 simulations and the final maps. As discussed in Planck Collaboration XII (2016), correction for this calibration effect should have no significant impact on cosmological parameters. As recommended, in this paper the CMB component of the simulations is simply rescaled by a factor of 1.0134 before analysis.

Of somewhat more importance is an observed noise mismatch between the simulations and the data. Whilst this has essentially no impact on studies of temperature anisotropy, it imposes important limitations on the statistical studies of polarization sky maps that can be included here. Conversely, analyses based on 1-point statistics, such as the variance, and the N-point correlation functions have played important roles in establishing the nature of this mismatch, which seems to be scale-dependent with an amplitude around 20% at lower resolutions but falling to a few per cent at higher resolution. As a consequence, this paper only includes results from a stacking analysis of the polarized data, in which the stacking of the data themselves necessarily acts to lower the effect of the noise mismatch. Polarization studies that do not rely on auto-statistics can still yield interesting new results, as found in Planck Collaboration XIII (2016); Planck Collaboration XVII (2016); Planck Collaboration XVIII (2016).

4. Tests of non-Gaussianity

There is no unique signature of non-Gaussianity, but the application of a variety of tests over a range of angular scales allows us to probe the data for departures from theoretically motivated Gaussian statistics. One of the more important tests in the context of inflationary cosmology is related to the analysis of the bispectrum. This is discussed thoroughly in Planck Collaboration XVII (2016), and is therefore not discussed further in this paper. In this section, we present the results from a variety of statistical tools. Unless otherwise specified, the analyses are applied to all four component separation products (Commander, NILC, SEVEM, and SMICA) at a given resolution with the accompanying common mask, and significance levels are determined by comparison with the corresponding results derived from the FFP8 simulations, with typically 1000 being used for this purpose. Establishing the consistency of the results derived from the different component-separation techniques is essential in order to be able to make robust claims about the statistical nature of the observed temperature fluctuations, and potential deviations from Gaussianity.

4.1. One-dimensional moments

In this section we consider simple tests of Gaussianity based on the variance, skewness, and kurtosis of the CMB temperature
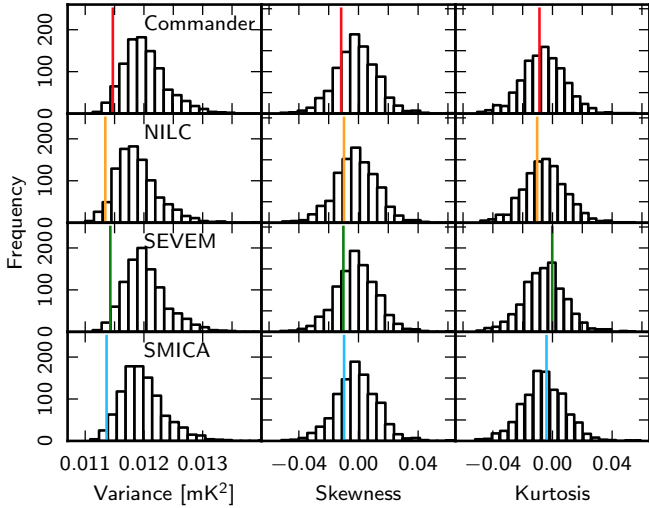
Fig. 1. Variance, skewness, and kurtosis for the four different component-separation methods – `Commander` (red), `NILC` (orange), `SEVEM` (green), and `SMICA` (blue) – compared to the distributions derived from 1000 Monte Carlo simulations.

maps. Previous analyses found an anomalously low variance in the WMAP sky maps (Monteserín et al. 2008; Cruz et al. 2011), which was subsequently confirmed in an analysis of the *Planck* 2013 data (PCIS13).

Cruz et al. (2011) developed the unit variance estimator to determine the variance, σ_0^2, of the CMB signal on the sky in the presence of noise. The normalized CMB map, u^X, is given by

$$u_i^X(\sigma_{X,0}^2) = \frac{X_i}{\sqrt{\sigma_{X,0}^2 + \sigma_{i,\mathrm{noise}}^2}}, \qquad (2)$$

where X_i is the observed temperature at pixel i and $\sigma_{i,\mathrm{noise}}^2$ is the noise variance for that pixel. Although this estimator is not optimal, Cruz et al. (2011) and Monteserín et al. (2008) have demonstrated that it is unbiased and sufficiently accurate for our purposes. The noise variance is estimated from the noise simulations for each component-separation algorithm. The CMB variance is then estimated by requiring that the variance of the normalized map u^X is unity. The skewness and kurtosis can then be obtained from the appropriately normalized map.

Figure 1 presents results for the variance, skewness, and kurtosis determined from the data at a resolution of 5′, $N_{\mathrm{side}} = 2048$. Good agreement between the component separation techniques is found, with small discrepancies likely due to sensitivity to the noise properties and their variation between methods.

Table 2 summarizes the lower-tail probabilities, defined as the percentage of MC simulations that show a lower variance, skewness, or kurtosis than the observed map, for these analyses. The results are in good agreement with PCIS13; the skewness and kurtosis are compatible with simulations, but the variance is marginally lower than in the simulations.

Although the variance is observed to be low, the results could still be affected by the presence of residual foregrounds at small scales in these maps, so that the true variance would be lower still. We assess this by application of the estimator to the cleaned frequency maps `SEVEM-100`, `SEVEM-143`, and `SEVEM-217`. The results, also presented in Table 2, are similar to those found for the combined map, although slightly less significant, which is most likely attributable to higher noise in the cleaned frequency maps.

Table 2. Lower-tail probabilities for the variance, skewness, and kurtosis of the component-separated maps.

Method	Probability [%]		
	Variance	Skewness	Kurtosis
`Commander`	3.2	17.2	35.3
`NILC`	3.3	20.9	30.9
`SEVEM`	1.9	20.5	56.8
`SMICA`	1.4	21.1	48.2
`SEVEM-100`	3.4	13.4	67.5
`SEVEM-143`	2.4	16.9	61.2
`SEVEM-217`	3.4	11.4	58.3

In conclusion, a simple statistical assessment of the *Planck* 2015 data using skewness and kurtosis shows no evidence for non-Gaussianity, although a low variance is found, which we will readdress in Sect. 5.1.

4.2. Testing the multi-normality of the CMB

Under the assumption of Gaussianity, the probability density function (PDF) of the N-dimensional pixelized temperature map is given by a multivariate Gaussian function:

$$f(\boldsymbol{T}) = \frac{1}{(2\pi)^{N_{\mathrm{pix}}/2} \det \mathsf{C}^{1/2}} \exp\left[-\frac{1}{2}\left(\boldsymbol{T}\mathsf{C}^{-1}\boldsymbol{T}^{\mathsf{T}}\right) \right], \qquad (3)$$

where $\boldsymbol{T}$ is a vector formed from the measured temperatures $T(\boldsymbol{x})$ over all positions allowed by the applied mask, N_{pix} is the number of pixels in the vector, and C is the covariance of the Gaussian field (of size $N_{\mathrm{pix}} \times N_{\mathrm{pix}}$).

Although the calculation of $\boldsymbol{T}\mathsf{C}^{-1}\boldsymbol{T}^{\mathsf{T}}$ can be achieved by conjugate gradient methods, the evaluation of $\det \mathsf{C}$ remains computationally difficult for the full *Planck* resolution at `HEALPix` $N_{\mathrm{side}} = 2048$. At a lower resolution, the problem is tractable, and the noise level can also be considered negligible compared to the CMB signal. That implies that under the assumption of isotropy the covariance matrix C is fully defined by the *Planck* angular power spectrum (C_ℓ):

$$\mathsf{C}_{ij} = \sum_{\ell=2}^{\ell_{\mathrm{max}}} \frac{2\ell+1}{4\pi} C_\ell b_\ell^2 P_\ell\left(\cos\theta_{ij}\right), \qquad (4)$$

where C_{ij} is the covariance between pixels i and j, θ_{ij} is the angle between them, P_ℓ are the Legendre polynomials, b_ℓ is an effective window function describing the combined effects of the instrumental beam and pixel window at resolution N_{side}, and ℓ_{max} is the maximum multipole probed.

Under the multivariate Gaussian hypothesis, the argument of the exponential in Eq. (3) should follow a χ^2 distribution with N_{pix} degrees of freedom, or, equivalently (for $N_{\mathrm{pix}} \gg 1$) a normal distribution $\mathcal{N}\left(N_{\mathrm{pix}}, \sqrt{2N_{\mathrm{pix}}}\right)$.

These χ^2 statistics are computed for the *Planck* 2015 component-separated CMB maps at $N_{\mathrm{side}} = 16$ and 32, then compared with the equivalent quantities derived from the corresponding FFP8 simulations. For those cases in which the covariance matrix is ill-conditioned, we use a principal component analysis approach to remove the lowest degenerate eigenvalues of the covariance matrix (see, e.g., Curto et al. 2011). This process is equivalent to adding uncorrelated regularization noise of amplitude $\approx 1\,\mu$K to the data before inversion. The results of the

Table 3. Lower-tail probabilities for the N-pdf χ^2 statistics derived from the *Planck* 2015 component-separated maps at $N_{\text{side}} = 16$ and 32.

	Probability [%]			
N_{side}	Comm.	NILC	SEVEM	SMICA
16	24.7	26.2	25.4	24.5
32	11.9	20.8	10.6	10.8

analysis are presented in Table 3 and indicate that the data are consistent with Gaussianity. We note that the lower-tail probabilities for the N-pdf decrease when the resolution of the data is increased from $N_{\text{side}} = 16$ to 32. However, this behaviour is consistent with that seen for simulations, and should not be considered to be significant.

4.3. N-point correlation functions

In this section, we present tests of the non-Gaussianity of the *Planck* 2015 temperature CMB maps using real-space N-point correlation functions. While harmonic-space methods are often preferred over real-space methods for studying primordial fluctuations, real-space methods have an advantage with respect to systematic errors and foregrounds, since such effects are usually localized in real space. It is therefore important to analyse the data in both spaces in order to highlight different features.

An N-point correlation function is defined as the average product of N temperatures, measured in a fixed relative orientation on the sky,

$$C_N(\theta_1, \ldots, \theta_{2N-3}) = \langle T(\hat{\boldsymbol{n}}_1) \cdots T(\hat{\boldsymbol{n}}_N) \rangle, \tag{5}$$

where the unit vectors $\hat{\boldsymbol{n}}_1, \ldots, \hat{\boldsymbol{n}}_N$ span an N-point polygon. Under the assumption of statistical isotropy, these functions depend only on the shape and size of the N-point polygon, and not on its particular position or orientation on the sky. Hence, the smallest number of parameters that uniquely determines the shape and size of the N-point polygon is $2N - 3$.

The correlation functions are estimated by simple product averages over all sets of N pixels fulfilling the geometric requirements set by $\theta_1, \ldots, \theta_{2N-3}$ characterizing the shape and size of the polygon,

$$\hat{C}_N(\theta_1, \ldots, \theta_{2N-3}) = \frac{\sum_i \left(w_1^i \cdots w_N^i\right)\left(T_1^i \cdots T_N^i\right)}{\sum_i \left(w_1^i \cdots w_N^i\right)}. \tag{6}$$

Pixel weights $w_1^i, \ldots, w_N^i$ can be introduced in order to reduce noise or mask boundary effects. Here they represent masking by being set to 1 for included pixels and to 0 for excluded pixels.

The shapes of the polygons selected for the analysis are the pseudo-collapsed and equilateral configurations for the 3-point function, and the rhombic configuration for the 4-point function, composed of two equilateral triangles that share a common side. We use the same definition of pseudo-collapsed as in Eriksen et al. (2005), i.e., an isosceles triangle where the length of the baseline falls within the second bin of the separation angles. The length of the longer edge of the triangle, θ, parameterizes its size. Analogously, in the case of the equilateral triangle and rhombus, the size of the polygon is parameterized by the length of the edge, θ. Note that these functions are chosen for ease of implementation, not because they are better suited for testing Gaussianity than other configurations. For a Gaussian field, Wick's theorem (Wick 1950) means that the ensemble average of

Table 4. Probabilities of obtaining values for the χ^2 statistic of the N-point functions for the *Planck* fiducial ΛCDM model at least as large as the observed values of the statistic for the *Planck* 2015 temperature CMB maps with resolution parameter $N_{\text{side}} = 64$, estimated using the `Commander`, `NILC`, `SEVEM`, and `SMICA` methods.

	Probability [%]			
Function	Comm.	NILC	SEVEM	SMICA
2-pt.	97.2	98.9	97.4	98.1
Pseudo-coll. 3-pt.	92.1	94.7	91.8	92.2
Equil. 3-pt.	74.0	80.4	75.8	79.0
Rhombic 4-pt.	64.6	70.9	65.6	65.9

the 4-point function may be written in terms of the 2-point function. In the following, all results refer to the connected 4-point function, i.e., are corrected for this Gaussian contribution.

We use a simple χ^2 statistic to quantify the agreement between the observed data and simulations, defined by

$$\chi^2 = \sum_{i,j=1}^{N_{\text{bin}}} \left(\hat{C}_N(\theta_i) - \langle C_N(\theta_i) \rangle\right) \mathsf{M}_{ij}^{-1} \left(\hat{C}_N(\theta_j) - \langle C_N(\theta_j) \rangle\right). \tag{7}$$

Here, $\hat{C}_N(\theta_i)$ is the N-point correlation function for the bin with separation angle θ_i, $\langle C_N(\theta_i) \rangle$ is the corresponding average from the MC simulation ensemble, and N_{bin} is the number of bins used for the analysis. If $\hat{C}_N^k(\theta_i)$ is the kth simulated N-point correlation function and N_{sim} is the number of simulations, then the covariance matrix M_{ij} is given by

$$\mathsf{M}_{ij} = \frac{1}{N'_{\text{sim}}} \sum_{k=1}^{N_{\text{sim}}} \left(\hat{C}_N^{(k)}(\theta_i) - \langle C_N(\theta_i) \rangle\right)\left(\hat{C}_N^{(k)}(\theta_j) - \langle C_N(\theta_j) \rangle\right), \tag{8}$$

where $N'_{\text{sim}} = N_{\text{sim}} - 1$. Following Hartlap et al. (2007), we then correct for bias in the inverse covariance matrix by multiplying it by the factor $(N'_{\text{sim}} - N_{\text{bin}} - 1)/N'_{\text{sim}}$. Below, we quote the significance level in terms of the fraction of simulations with a larger χ^2 value than the observed map.

We analyse the CMB estimates at a resolution of $N_{\text{side}} = 64$, this being constrained by computational limitations. The results are presented in Fig. 2, where we compare the N-point functions for the data and the mean values estimated from 1000 MC simulations. The probabilities of obtaining values of the χ^2 statistic for the *Planck* fiducial ΛCDM model at least as large as the observed values are given in Table 4.

It is worth noting that the values of the N-point functions for different angular separations are strongly correlated, and for this reason the figures show only one profile of each function in multi-dimensional space. Since the estimated probabilities take into account the correlations, they provide more reliable information on the goodness of fit between the data and a given model than simple inspection of the figures.

The results show excellent consistency between the CMB maps estimated using the different component-separation methods. No statistically significant deviations of the CMB maps from Gaussianity are found. Indeed, the slight preference for super-Gaussianity of the equilateral 3-point and 4-point functions observed for the 2013 data is now less pronounced. That may be caused by differences between the masks used for the analysis. Interestingly, the 2-point function shows clear evidence of a lack of structure for large separation angles. Such behaviour was originally noted for the WMAP first-year data by Bennett et al. (2003), and has subsequently been discussed at length in the literature (Efstathiou 2004; Copi et al. 2007, 2015). We will return to this issue in Sect. 5.2.
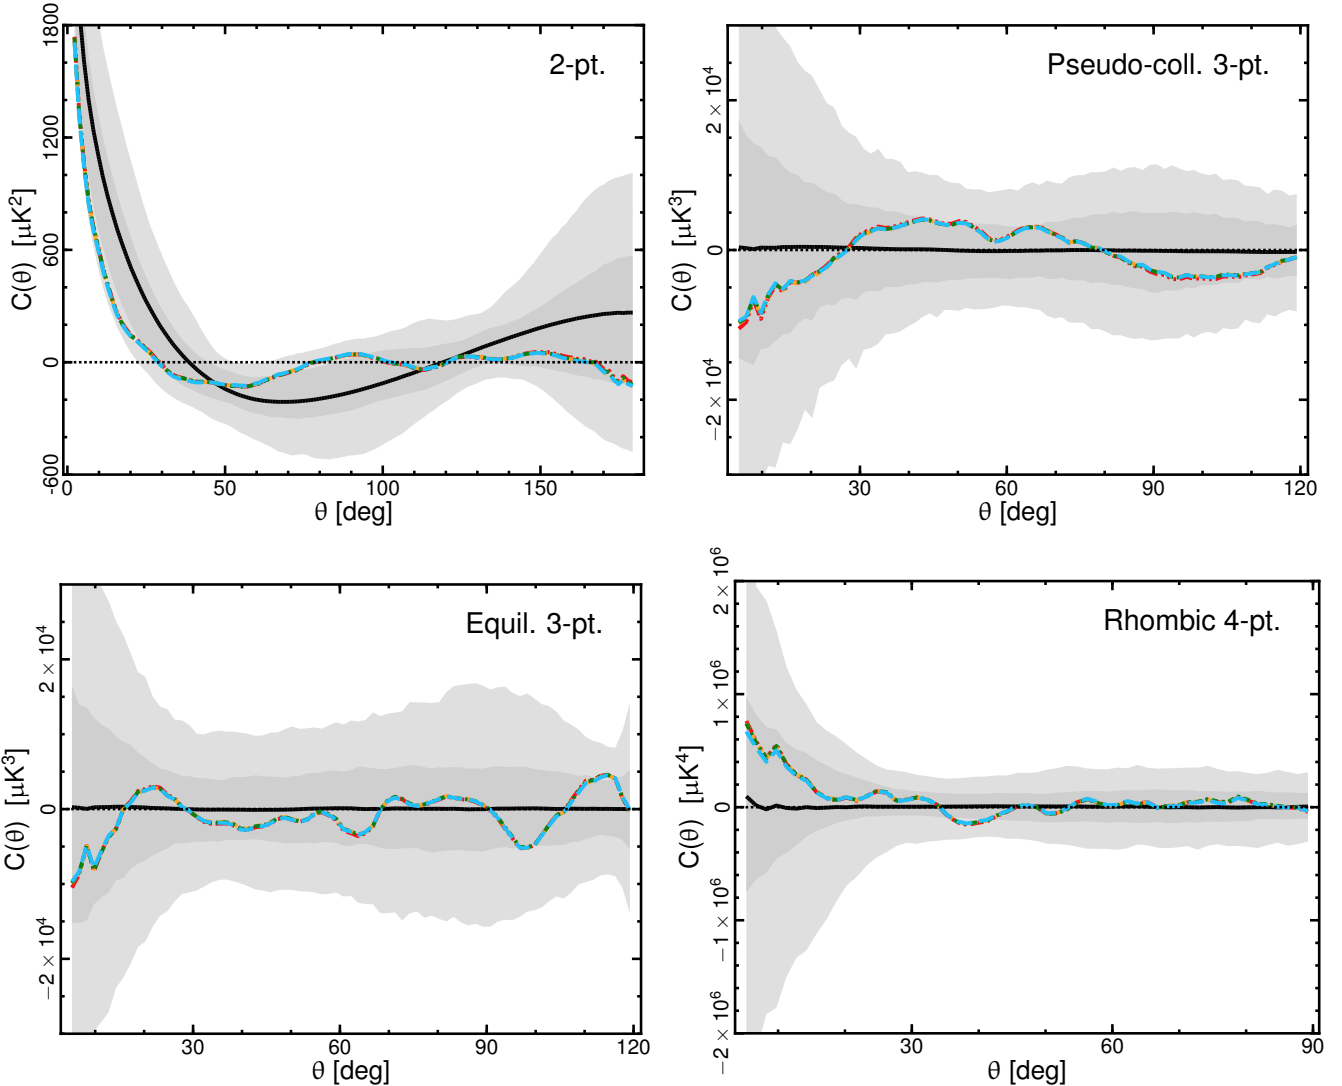
Fig. 2. *N*-point correlation functions determined from the $N_{side} = 64$ *Planck* CMB 2015 temperature maps. Results are shown for the 2-point, pseudo-collapsed 3-point (*upper left and right panels*, respectively), equilateral 3-point, and connected rhombic 4-point functions (*lower left and right panels*, respectively). The red dot-dot-dot-dashed, orange dashed, green dot-dashed, and blue long dashed lines correspond to the `Commander`, `NILC`, `SEVEM`, and `SMICA` maps, respectively. Note that the lines lie on top of each other. The black solid line indicates the mean determined from 1000 `SMICA` simulations. The shaded dark and light grey regions indicate the corresponding 68% and 95% confidence regions, respectively. See Sect. 4.3 for the definition of the separation angle θ.

4.4. Minkowski functionals

The Minkowski functionals (hereafter MFs) describe the morphology of fields in any dimension and have long been used as estimators of non-Gaussianity and anisotropy in the CMB (see e.g., Mecke et al. 1994; Schmalzing & Buchert 1997; Schmalzing & Gorski 1998; Komatsu et al. 2003; Eriksen et al. 2004b; Curto et al. 2007; De Troia et al. 2007; Spergel et al. 2007; Curto et al. 2008; Hikage et al. 2008; Komatsu et al. 2009; Planck Collaboration XXIII 2014). They are additive for disjoint regions of the sky and invariant under rotations and translations. In the literature, the contours are traditionally defined by a threshold ν, usually given in units of the sky standard deviation (σ_0).

We compute MFs for the regions colder and hotter than a given threshold ν. Thus, the three MFs, namely the area $V_0(\nu) = A(\nu)$, the perimeter $V_1(\nu) = C(\nu)$, and the genus $V_2(\nu) = G(\nu)$, are defined respectively as

$$V_0(\nu) = A(\nu) = \frac{N_\nu}{N_{pix}}, \tag{9}$$

$$V_1(\nu) = C(\nu) = \frac{1}{4A_{tot}} \sum_i S_i, \tag{10}$$

$$V_2(\nu) = G(\nu) = \frac{1}{2\pi A_{tot}}(N_{hot} - N_{cold}), \tag{11}$$

where N_ν is the number of pixels where $\Delta T/\sigma_0 > \nu$, N_{pix} is the total number of available pixels, A_{tot} is the total area of the available sky, N_{hot} is the number of compact hot spots, N_{cold} is the number of compact cold spots, and S_i is the contour length of each hot spot.

For a Gaussian random field in pixel space, the MFs can be written in terms of two functions: A_k, which depends only on the power spectrum, and v_k, which is a function only of the threshold ν (see, e.g., Vanmarcke 1983; Pogosyan et al. 2009; Gay et al. 2012; Matsubara 2010; Fantaye et al. 2015). The analytical expressions are

$$V_k(\nu) = A_k v_k(\nu), \tag{12}$$

Table 5. Probability $P\left(\chi^2 > \chi^2_{\rm Planck}\right)$ as a function of resolution for the unnormalized, classical Minkowski functionals.

$N_{\rm side}$	Probability [%]			
	Comm.	NILC	SEVEM	SMICA
1024	91.4	90.7	95.5	95.8
512	95.4	90.9	62.6	92.6
256	55.8	34.5	55.9	55.9
128	43.6	56.4	19.9	19.2
64	59.3	37.8	22.7	80.0
32	62.0	16.2	29.9	67.0
16	43.4	45.8	47.7	31.0

Table 6. Probability $P\left(\chi^2 > \chi^2_{\rm Planck}\right)$ as a function of resolution determined using normalized MFs.

$N_{\rm side}$	Probability [%]			
	Comm.	NILC	SEVEM	SMICA
2048	97.2	77.7	99.0	93.0
1024	93.1	98.0	90.2	92.6
512	53.7	36.7	30.4	77.6
256	89.0	85.9	96.8	58.1
128	93.0	63.5	94.1	37.1
64	37.1	70.4	54.1	62.5
32	28.9	77.4	75.5	46.7
16	33.1	39.4	44.1	38.8

with

$$v_k(\nu) = \exp(-\nu^2/2) H_{k-1}(\nu), \quad k \le 2, \tag{13}$$

$$v_3(\nu) = \frac{e^{-\nu^2}}{\mathrm{erfc}\left(\nu/\sqrt{2}\right)}, \tag{14}$$

and

$$H_n(\nu) = e^{\nu^2/2} \left(-\frac{\mathrm{d}}{\mathrm{d}\nu}\right)^n e^{-\nu^2/2}. \tag{15}$$

The amplitude A_k depends only on the shape of the power spectrum C_ℓ through the rms of the field σ_0 and its first derivative σ_1:

$$A_k = \frac{1}{(2\pi)^{(k+1)/2}} \frac{\omega_2}{\omega_{2-k}\omega_k} \left(\frac{\sigma_1}{\sqrt{2}\sigma_0}\right)^k, \quad k \le 2, \tag{16}$$

$$A_3 = \frac{2}{\pi} \left(\frac{\sigma_1}{\sqrt{2}\sigma_0}\right)^2, \tag{17}$$

where $\omega_k \equiv \pi^{k/2}/\Gamma(k/2 + 1)$.

Since this factorization is still valid in the weakly non-Gaussian case, we can use the normalized MFs, v_k, to focus on deviations from Gaussianity, with a reduced sensitivity to cosmic variance.

Apart from the characterization of the MFs using full-resolution temperature sky maps, we also consider results at different angular scales. In this paper, two different approaches are considered to study these degrees of freedom: in real space via a standard Gaussian smoothing and degradation of the maps; and, for the first time, in harmonic space by using needlets. Such a complete investigation provides an insight regarding the harmonic and spatial nature of possible non-Gaussian features detected with the MFs.

First, we apply scale-dependent analyses in real space by considering the sky maps at different resolutions. The three classical MFs – area, contour length, and genus – are evaluated over the threshold range $-3 \le \nu \le 3$ in σ units, with a step of 0.5. This provides a total of 39 different statistics. The values of these statistics for the *Planck* data are all within the 95% confidence region when compared with Gaussian simulations for all of the resolutions considered. A χ^2 value is computed for each component-separation method by combining the 39 statistics and taking into account their correlations (see e.g., Curto et al. 2007, 2008). The corresponding covariance matrix is computed using 1000 simulations. The *p*-value of this χ^2 test is presented in Table 5 for each component separation technique and for map resolutions between $N_{\rm side} = 1024$ and $N_{\rm side} = 16$. We find no significant deviations from Gaussianity for any of the resolutions considered.

Then we consider the four normalized functionals described above. For every scale we used 26 thresholds ranging between -3.5 and 3.5 in σ units, except for $\theta = 640'$ where 13 thresholds between -3.0 and 3.0 in σ units were more appropriate. Table 6 indicates that no significant deviation from Gaussianity is found.

Third, we tested MFs on needlet components. The needlet components of the CMB field as defined by Marinucci et al. (2008) and Baldi et al. (2009) are given by:

$$\beta_j(\hat{\boldsymbol{n}}) = \sum_{\ell=B^{j-1}}^{B^{j+1}} b^2\left(\frac{\ell}{B^j}\right) \sum_m a_{\ell m} Y_{\ell m}(\hat{\boldsymbol{n}})$$

$$= \sum_{\ell=B^{j-1}}^{B^{j+1}} b^2\left(\frac{\ell}{B^j}\right) T_\ell(\hat{\boldsymbol{n}}). \tag{18}$$

Here, $T_\ell(\hat{\boldsymbol{n}})$ denotes the component at multipole ℓ of the CMB map $T(\hat{\boldsymbol{n}})$, i.e.,

$$T(\hat{\boldsymbol{n}}) = \sum_\ell T_\ell(\hat{\boldsymbol{n}}), \tag{19}$$

where $\hat{\boldsymbol{n}} \in S^2$ denotes the pointing direction, B is a fixed parameter (usually taken to be between 1 and 2) and $b(.)$ is a smooth function such that $\sum_j b^2(\ell/B^j) = 1$ for all ℓ. Fantaye et al. (2015) show in a rigorous way that a general analytical expression for MFs at a given needlet scale j, which deals with an arbitrary mask and takes into account the spherical geometry of the sky, can be written as

$$V_k^j = \sum_{i=0}^k t_{(2-i)} A_i^j v_i, \tag{20}$$

where $t_0 = 2$, $t_1 = 0$, and $t_2 = 4\pi$ are respectively the Euler-Poincaré characteristic, boundary length, and area of the full sphere. The quantities v_k are the normalized MFs given in Eq. (13), while the needlet scale amplitudes A_k^j have a similar form as A_k but with the variances of the map and its first derivative given by

$$\sigma_0^2 = \sum_\ell b^4\left(\frac{\ell}{B^j}\right) C_\ell \frac{2\ell+1}{4\pi}, \tag{21}$$

$$\sigma_1^2 = \sum_\ell b^4\left(\frac{\ell}{B^j}\right) C_\ell \frac{2\ell+1}{4\pi} \frac{\ell(\ell+1)}{2}. \tag{22}$$

Implementing the MFs in needlet space has several advantages: the needlet filter is localized in pixel space, hence the needlet component maps are minimally affected by masked regions, especially at high-frequency j; and the double-localization properties of needlets (in real and harmonic space) allow a much more
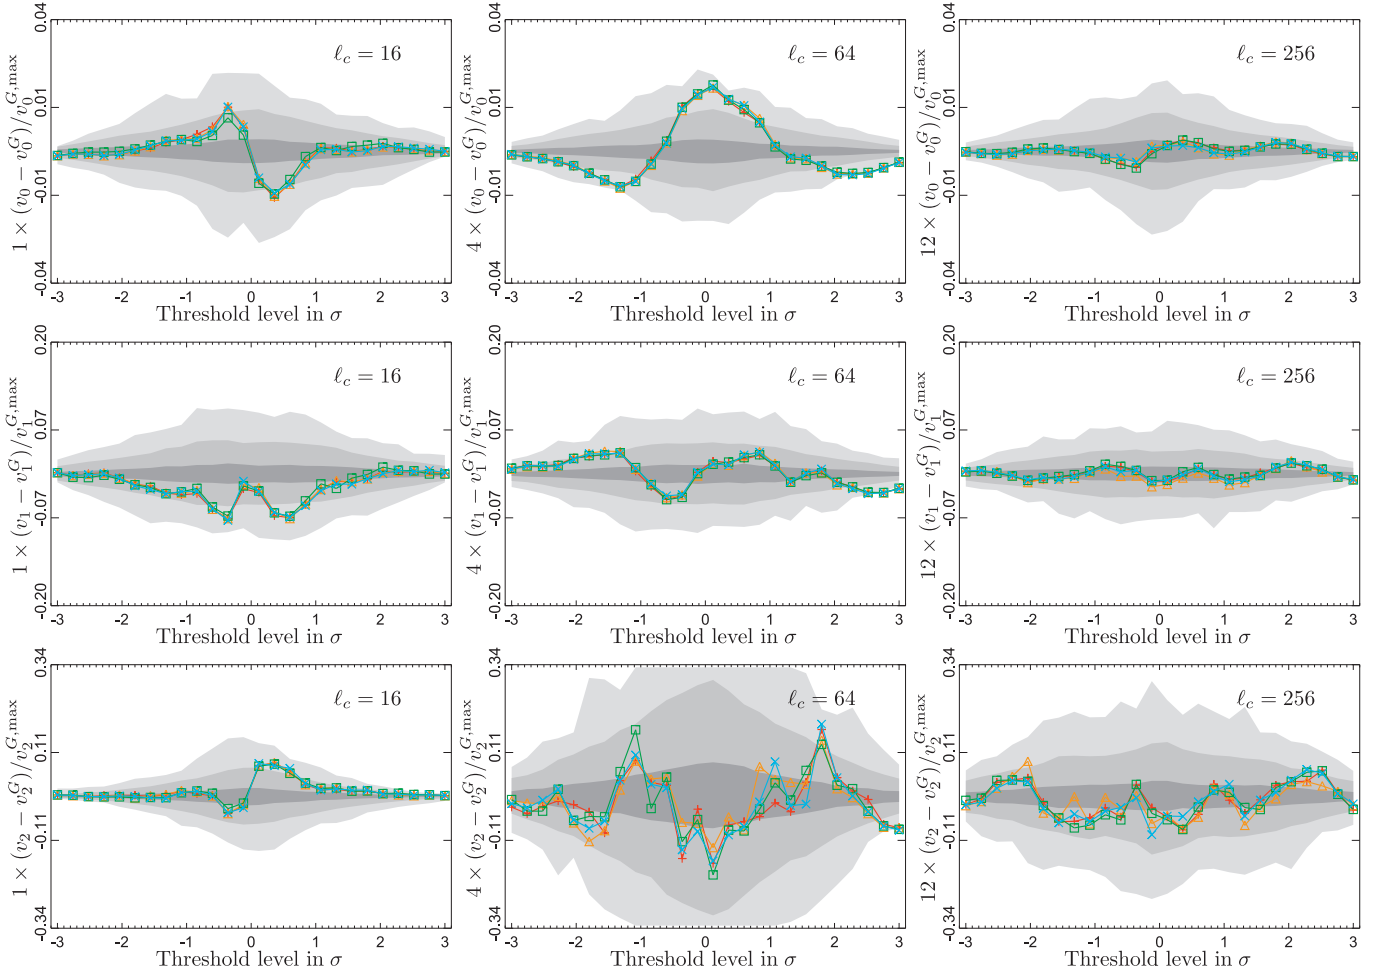
Fig. 3. Needlet space MFs for *Planck* 2015 data using the four component-separated maps, `Commander` (red), `NILC` (orange), `SEVEM` (green), and `SMICA` (blue); the grey regions, from dark to light, correspond, respectively, to 1, 2, and 3σ confidence regions estimated from the 1000 FFP8 simulations processed by the `Commander` method. The columns *from left to right* correspond to the needlet parameters $j = 4, 6$, and 8, respectively; the jth needlet parameter has compact support over multipole ranges $[2^{j-1}, 2^{j+1}]$. The $\ell_c = 2^j$ value indicates the central multipole of the corresponding needlet map. Note that to have the same range at all the needlet scales, the vertical axis has been multiplied by a factor that takes into account the steady decrease of the variance of the MFs as a function of scale.

precise, scale-by-scale, interpetation of any possible anomalies. While the behaviour of standard all-scale MFs is contaminated by the large cosmic variance of the low multipoles, this is no longer the case for MFs evaluated at the highest needlet scales; in such circumstances, the variance of normalized components may be shown to decrease steadily, entailing a much greater detection power in the presence of anomalies. Finally, and most importantly, the needlet MFs are more sensitive to the shape of the power spectrum than the corresponding all-scale MFs.

The needlet parameters we use in this analysis are $B = 2$, $j = 3, 4, 5, 6, 7, 8$. Since the masks in pixel space are map-resolution dependent, we also use different masks for each needlet scale. These new masks are constructed by multiplying the high-resolution common mask with the upgraded version of the appropriate low-resolution common mask. For needlet scales $j = 2$ and $j = 3$, we use the common mask defined at $N_{\text{side}} = 16$, and upgraded to $N_{\text{side}} = 2048$. Similarly, for the higher needlet scales, $j = 2^n$, where $n = 4, 5, 6, 7, 8$, we use upgraded versions of the common masks defined at $N_{\text{side}} = 2^n$.

The results concerning needlet MFs from the `Commander`, `NILC`, `SEVEM`, and `SMICA` foreground-cleaned temperature maps for needlet scales $B = 2$, $j = 4, 6, 8$ are shown in Fig. 3. All cases are computed using 26 thresholds ranging between -3.5

and 3.5 in σ units. The figure shows the fractional difference be-

Table 7. Probability $P\left(\chi^2 > \chi^2_{\text{Planck}}\right)$ as a function of needlet scale.

Needlet scale (ℓ range)	Probability [%]			
	Comm.	NILC	SEVEM	SMICA
3 (4,16)	32.1	36.1	40.4	39.8
4 (8,32)	84.0	57.9	79.4	59.4
5 (16,64)	23.8	11.2	29.1	43.8
6 (32,128)	14.8	38.9	33.5	34.1
7 (64,256)	11.9	7.5	15.4	1.1
8 (128,512)	46.1	55.2	67.7	52.2

tween the *Planck* data and the FFP8 simulations in area (top panels), boundary length (middle panels), and genus (bottom panels) for different needlet scales. The jth needlet scale has compact support over the multipole ranges $[2^{j-1}, 2^{j+1}]$. All the scales we considered are consistent with the Gaussian FFP8 simulations. This can be seen in Fig. 4, where we compare the data and simulation χ^2 values, which are computed by combining the three MFs with an appropriate covariance matrix. The vertical lines in
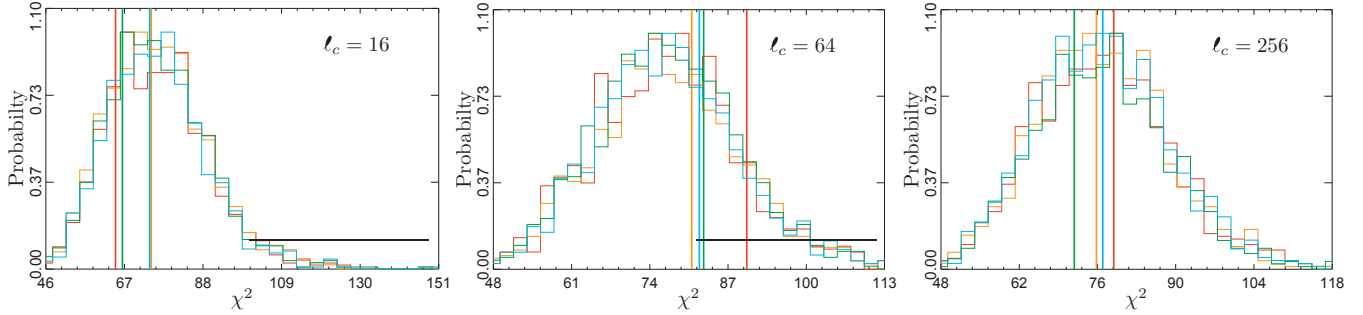
Fig. 4. Histograms of χ^2 for the *Planck* 2015 `Commander` (red), `NILC` (orange), `SEVEM` (green), and `SMICA` (blue) foreground-cleaned maps analysed with the common mask. The χ^2 is obtained by combining the three MFs in needlet space with an appropiate covariance matrix. The histograms are for the FFP8 simulations, while the vertical lines are for the data. The figures from *left to right* are for the needlet scales $j = 4, 6$, and 8, with the central multipoles $\ell_c = 2^j$ shown in each panel.

these figures represent the data, while the histogram shows the results for the 1000 FFP8 simulations. We also show in Table 7 the *p*-values for the four component-separation methods, as well as all needlet scales we considered. Despite the relatively small *p*-values for some scales, the *Planck* temperature maps show no significant deviation from the Gaussian simulations up to $\ell_{max} = 512$, which corresponds to the maximum multipole of our highest-frequency needlet map.

4.5. Multiscale analyses

Multiscale data analysis is a powerful approach for probing the fundamental hypotheses of the isotropy and Gaussianity of the CMB. The exploration of different scales (in an almost independent manner) not only helps to test the specific predictions of a given scenario for the origin and evolution of the fluctuations, but also is an important check on the impact of systematic errors or other contaminants on the cosmological signal.

There are several ways of performing a multiscale analysis, the simplest being to smooth/degrade the CMB map to different resolutions. However, in this section, we will focus on image processing techniques related to the application of wavelets and more general band-pass filtering kernels to the original CMB fluctuations. The advantage of wavelet-like analyses over scale degradation is clear: they allow the exploration of characteristics of the data that are related to specific angular scales. Wavelets have already been extensively used in the study of the Gaussianity and isotropy of the CMB (e.g., McEwen et al. 2007; Vielva 2007). Indeed, a wavelet-based (needlet) analysis of the *Planck* 2015 data has already been presented in Sect. 4.4.

We recall that in the 2013 analysis, some of the applied estimators deviated from the null hypothesis. In particular, it was determined that the cold area of the spherical Mexican hat wavelet (SMHW, Martinez-González et al. 2002) coefficients at scales of around 5° yielded a *p*-value of 0.3%. In addition, we also found an excess in the kurtosis of the wavelet coefficients on the same scales. Previous analyses (for a review, see Vielva 2010) have suggested that the "Cold Spot" (see Sect. 5.7) was the major contributor to these statistical outliers.

In what follows, we will consider the application of the SMHW, together with matched filters for a 2D-Gaussian profile (GAUSS), and for generalized spherical Savitzky-Golay kernels (SSG, Savitzky & Golay 1964, see Appendix A).

The application of a filter $\psi(R, p)$ to a signal on the sky $S(p)$ can be written as

$$\omega_S(R, p) = \sum_{\ell=0}^{\ell_{max}} \sum_{m=-\ell}^{m=\ell} s_{\ell m} W_\ell^\psi(R) Y_{\ell m}(p),$$ (23)

where p represents a given position/pixel, R parameterizes a characteristic scale for the filter (e.g., a wavelet scale), $W_\ell^\psi(R)$ is the window function associated with the filter $\psi(R, p)$, ℓ_{max} is the maximum multipole allowed by the corresponding `HEALPix` pixelization, and $Y_{\ell m}(p)$ is the spherical harmonic basis. Here, $s_{\ell m}$, the spherical harmonic coefficients of the analysed map, are given by

$$s_{\ell m} = \int d\Omega\, Y_{\ell m}^*(p) S(p),$$ (24)

where $d\Omega = d\theta \sin\theta d\phi$ and the asterisk denotes complex conjugation. Note that the filtered map (or the wavelet coefficient map, if $\psi(R, p)$ is a continous wavelet) conserves the statistical properties of the original map, since the convolution is a linear operation. In particular, if $S(p)$ is a Gaussian and statistically isotropic random signal, $\omega_S(R, p)$ is also Gaussian and statistically isotropic.

In the present work, the signal $S(p)$ corresponds to a temperature map $T(p)$. Several statistics can then be computed from the derived filtered map as a function of the filter scale, in particular, the first moments (the dispersion σ_R, the skewness S_R, and the kurtosis K_R), the total area above/below a given threshold, and the peak distribution. These statistics are compared to the corresponding results determined from the FFP8 simulations to establish the degree of compatibility with the null hypothesis.

4.5.1. First-order moments of the multiscale maps

For the three filters considered (SMHW, GAUSS, and SSG84[3]) the variance, skewness, and kurtosis are computed at 18 scales, $R(\text{arcmin}) = \{2, 4, 7, 14, 25, 50, 75, 100, 150, 200, 250, 300, 400, 500, 600, 750, 900, 1050\}$. These scales are chosen to be consistent with previous analyses. They cover a wide angular range, and are selected so that the intervals between them increase with scale. Notice that, for a given scale, the three filters do not cover exactly the same multipole range, since that depends on the specific filter definition. This can be seen in

---

[3] The digits 8 and 4 denote the order of the spherical Savitzky-Golay kernel and the smoothing weight, described in Appendix A.
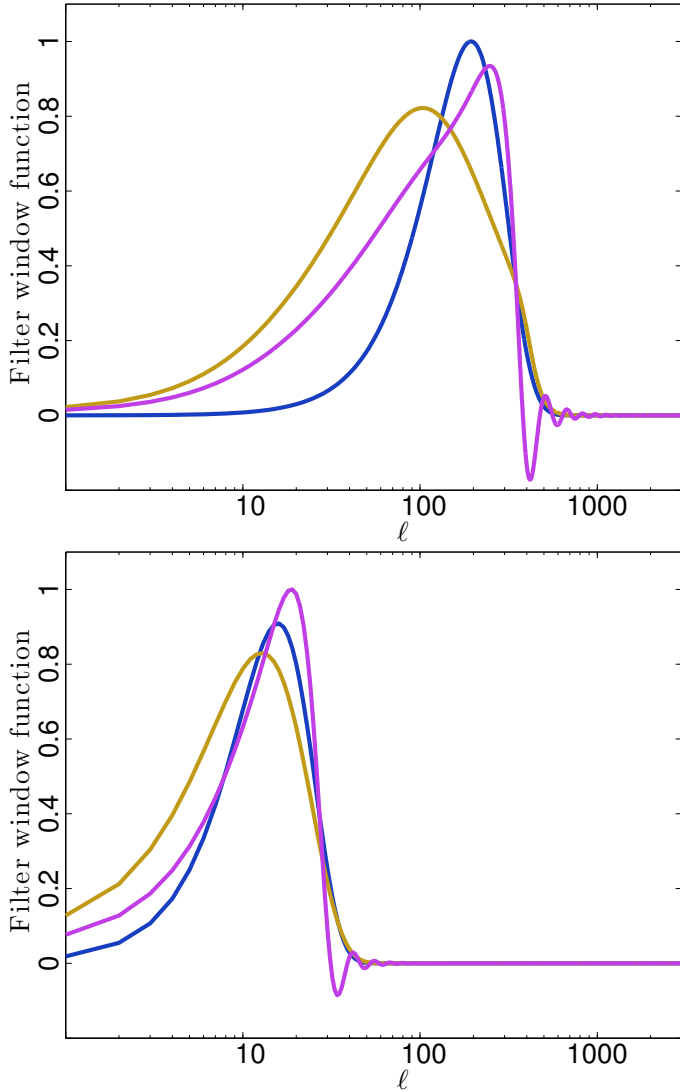
Fig. 5. Comparison of the window functions (normalized to have equal area) for the SMHW (blue), GAUSS (yellow), and SSG84 (magenta) filters. The scales shown are 25′ (*top*) and 250′ (*bottom*).

Fig. 5: the SMHW is the narrowest filter, followed by SSG84, then GAUSS. The three filters have an equivalent effective ℓ_{max}, but differ in the effective ℓ_{min}. Overall, the differences between the filters become smaller with increasing effective scale. In this paper, we refer to both the scale, R, and FWHM as parameters defining the size of the filters. For the SMHW, these are related by $FWHM = R\sqrt{8\ln 2}$, whereas for the GAUSS and SSG84 filters, the scale is defined to be half the FWHM. The latter definition is appropriate for filters that include pre-whitening since it is simple yet matches the ℓ-space bandwidth reasonably well.

Following the procedure explained in PCIS13, after convolution with a given filter, the common mask is extended to omit pixels from the analysis that could be contaminated by the mask. These pixels introduce an extra correlation between the data and the simulations, degrading the statistical power of the comparison with the null hypothesis (see, e.g., Vielva et al. 2004). For a given scale R, the exclusion mask is defined by extending an auxiliary mask by a distance $2R$ from its border, where the auxiliary mask is that part of the common mask related to residual

diffuse Galactic emission (i.e., the auxiliary mask does not mask point sources).

The following figures represent the upper-tail probability (UTP) for a given statistic, i.e., the fraction of simulations that yield a value equal to or greater than that obtained for the data. In fact, as explained in PCIS13, if a given UTP is larger than 0.5, a new quantity is defined as mUTP = 1 − UTP. Therefore, mUTP is constrained to lie between $1/N$ and 0.5, where N is the number of simulations used for each statistic.

Figure 6 presents the comparison of the CMB temperature maps with the corresponding simulations for the SMHW, GAUSS, and SSG84 filters. The full mission *Planck* data confirm the results already obtained with the 2013 release for temperature. In particular, for the SMHW, we find (i) an excess of kurtosis (≈0.8%) at scales of around 300′; (ii) that the dispersion of the wavelet coefficients at these scales and at around 700′ is relatively low (≈1%); and (iii) that the dispersion of the wavelet coefficients at scales below 5′ is significantly high (≲0.1%).

The excess of kurtosis has been previously associated with the Cold Spot (e.g., Vielva et al. 2004), and the low value of the standard deviation of the coefficients on large scales could be related to the low variance discussed in Sect. 5.1. Regarding the large dispersion of the coefficients on the smallest scales, this can be understood either by the presence of residual foreground contributions (extragalactic point sources) or by incomplete characterization of the true instrumental noise properties by the FFP8 simulations. We explore these possibilities with two additional tests undertaken with the SMHW.

Figure 7 shows the significance of the statistics derived from the SEVEM-100, SEVEM-143, and SEVEM-217 maps. The three cleaned maps yield very consistent values of the mUTP for the standard deviation, skewness, and kurtosis of the wavelet coefficients, with only small differences seen at small scales. This frequency-independence of the results argues against the foreground residuals hypothesis. Figure 8 presents the same statistics as applied to an estimator of the noise properties of the CMB maps. This is derived from the half-difference of the half-ring data sets, which provides the best estimate of the noise properties of the full mission data set. However, since there is still a known mismatch in noise properties, any conclusions will be more qualitative than quantitative. Nevertheless, the noise study reveals that, at the smallest scales, there are some discrepancies with the FFP8 simulations, and in particular the estimated dispersion of the SMHW noise coefficients is higher than predicted.

4.5.2. The area above/below a threshold

In the context of the study of the Cold Spot, the area above/below a given threshold, as a function of the SMHW wavelet scale, has been demonstrated to provide a useful and robust statistic (e.g., Cruz et al. 2005), since it is rather independent of any masking required. Our previous analysis (PCIS13) confirmed that the CMB temperature fluctuations exhibit an anomalously large cold area on scales of around 10°, which can be mostly associated with the Cold Spot. Here, we extend the analysis by including results derived using the GAUSS and SSG84 filters.

At a given scale R and threshold ν, the cold ($A_R^{-\nu}$) and hot ($A_R^{+\nu}$) areas of a filtered map are defined as

$$A_R^{-\nu} \equiv \#\{\omega_S(R, p) < -\nu\}, \qquad (25)$$
$$A_R^{+\nu} \equiv \#\{\omega_S(R, p) < +\nu\}, \qquad (26)$$

where the operator # represents the number of pixels p in which the condition defined between the braces is satisfied.
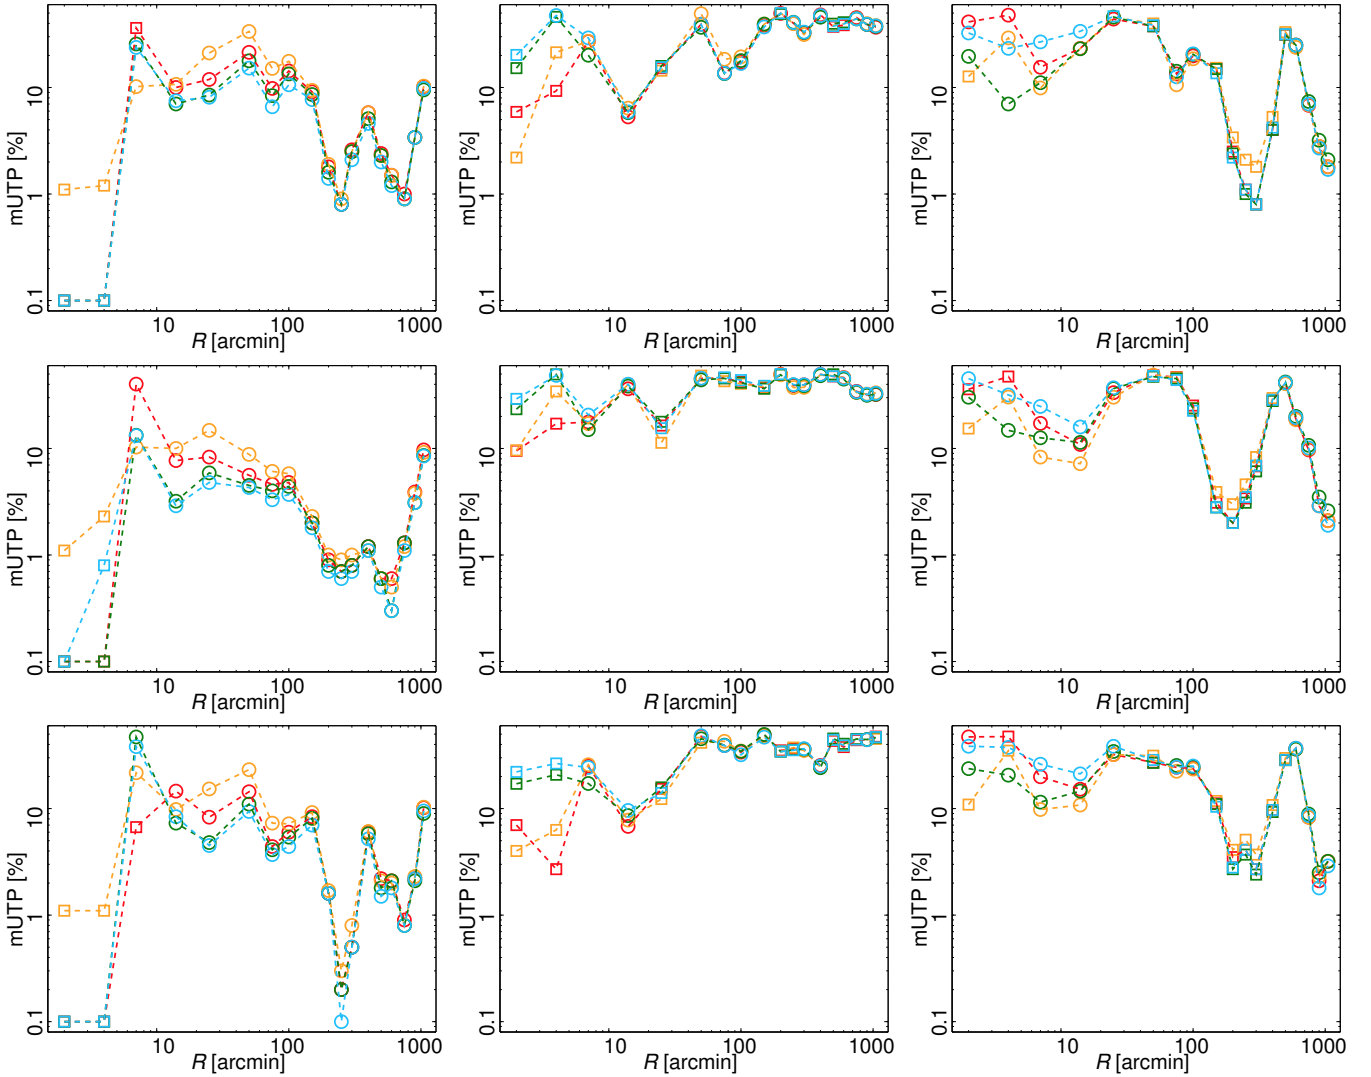
Fig. 6. Modified upper tail probabilities (mUTP) obtained from the analyses of the filter coefficients as a function of the filter scale R for the Commander (red), NILC (orange), SEVEM (green), and SMICA (blue) sky maps. *From left to right*, the panels correspond to standard deviation, skewness, and kurtosis results, when determined using the SMHW (*top*), GAUSS (*middle*), and SSG84 (*bottom*) filters. The squares represent UTP values above 0.5, whereas circles represent UTP values below 0.5.
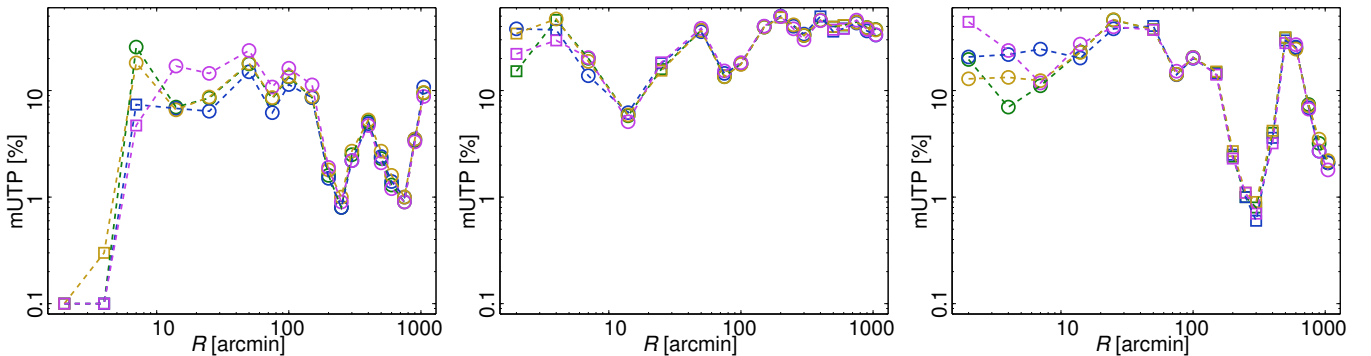

Fig. 7. Modified upper tail probabilities (mUTP) obtained from the analyses of the SMHW coefficients as a function of the wavelet scale R for the SEVEM-100 (blue), SEVEM-143 (yellow), SEVEM-217 (magenta), and SEVEM (green) maps. *From left to right*, the panels correspond to the standard deviation, skewness, and kurtosis.

Table 8 summarizes the results for the hot and cold areas determined for the four CMB temperature maps analysed with the common mask (and its associated exclusion masks). The results are similar to those obtained in 2013, with some small differences on those scales related to the Cold Spot (between 200′ and 400′). Specifically, the cold area is slightly less significant for smaller values of R, whereas the anomalous behaviour remains for larger filter scales. The three filters are in
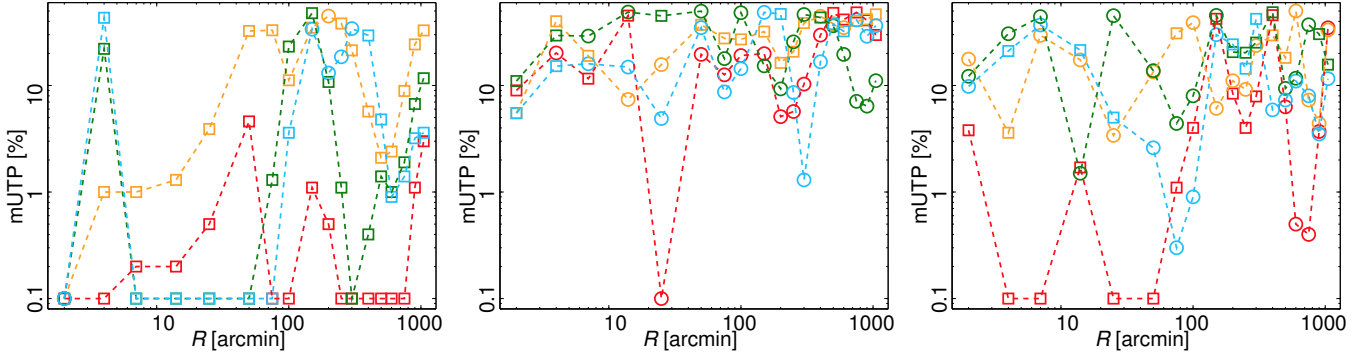
Fig. 8. Modified upper tail probabilities (mUTP) obtained from the analyses of the SMHW coefficients as a function of the wavelet scale R for the `Commander` (red), `NILC` (orange), `SEVEM` (green), and `SMICA` (blue) half-ring half-difference noise estimates. From *left to right*, the panels correspond to the standard deviation, skewness, and kurtosis.

reasonable agreement, but, as expected from Fig. 6, the SMHW yields higher significance levels than the SSG84 and GAUSS filters. However, it is worth recalling that, for a given scale, the three filters are not probing exactly the same multipole range and therefore some differences should be expected.

In Fig. 9 we plot the areas for thresholds $\nu > 3.0$, where the threshold is defined in units of σ_R, as determined from the `SEVEM` temperature map. The results for `Commander`, `NILC`, and `SMICA` are in good agreement with these. The panels refer to SMHW scales of $R = 200', 250', 300'$, and $400'$. The most extreme value (in terms of σ_R) for each area is indicated.

The coldest area corresponds to the Cold Spot with the minimum value of the wavelet coefficient at the position $(209°, -57°)$ in Galactic coordinates. The hottest area has already been identified in the WMAP data (e.g., Vielva et al. 2007). The results are insensitive to the choice of CMB temperature map that is adopted. It is clear that the southern Galactic hemisphere yields more anomalous signatures than the northern one. These results confirm the importance of the Cold Spot as the most extreme feature in the analysed sky. More insights about its nature are provided in Sect. 5.7.

4.5.3. Peak statistics

The statistical properties of local extrema (both minima and maxima, which we refer to collectively as "peaks") provide an alternative approach to search for evidence of non-Gaussianity in the data. Such peaks, defined as pixels whose amplitudes are higher or lower than the corresponding values for all of their nearest neighbours, trace topological properties of the data. Peak locations and amplitudes, and various derived quantities, such as their correlation functions, have previously been used to characterize the WMAP sky maps by Larson & Wandelt (2004, 2005) and Hou et al. (2009).

The statistical properties of peaks for a statistically isotropic Gaussian random field were derived by Bond & Efstathiou (1987). The integrated number density of peaks, n_{pk} (composed of maxima and minima with corresponding densities n_{max} and n_{min}), with amplitudes x above a certain threshold $\nu = x/\sigma$ is given by

$$\frac{n_{max} + n_{min}}{n_{pk}} \left(\frac{x}{\sigma} > \nu \right) = \sqrt{\frac{3}{2\pi}} \gamma^2 \nu \exp\left(-\frac{\nu^2}{2}\right) \tag{27}$$
$$+ \frac{1}{2} \operatorname{erfc}\left[\frac{\nu}{\sqrt{2 - \frac{4}{3}\gamma^2}} \right],$$

Table 8. Modified upper tail probability (mUTP) for the cold (top) and hot (bottom) areas.

Area	Scale [arcmin]	Probability [%]			
		Comm.	NILC	SEVEM	SMICA
SMHW					
Cold	200	3.8	5.1	3.7	3.8
	250	1.4	2.4	1.4	1.4
	300	0.4	1.5	0.4	0.4
	400	0.9	0.9	0.9	0.9
Hot	200	2.0	2.6	1.7	1.5
	250	2.4	3.0	2.1	2.0
	300	4.2	5.0	4.1	3.9
	400	...	...	...	...
GAUSS					
Cold	200	1.7	2.7	1.7	1.7
	250	1.2	1.2	1.2	1.2
	300	1.6	1.8	1.2	1.8
	400	...	...	...	...
Hot	200	2.9	3.5	2.8	2.6
	250	5.7	6.4	5.6	5.4
	300	...	...	...	...
	400	...	...	...	...
SSG84					
Cold	200	9.4	11.0	9.4	9.0
	250	12.3	13.4	10.8	12.3
	300	1.4	2.6	1.4	1.5
	400	0.9	1.9	0.8	0.9
Hot	200	1.1	1.8	1.0	0.9
	250	4.8	5.1	4.5	4.3
	300	...	...	...	...
	400	...	...	...	...

Notes. Results are given for the $\nu > 4\sigma_R$ threshold of the SMHW, GAUSS, and SSG84 coefficients. The four most significant scales related to the Cold Spot feature are shown. An ellipsis (...) indicates that no area above that threshold was found in the data.

where σ is the rms fluctuation amplitude measured on the sky, and γ is the spectral shape parameter of the underlying field. Uncharacteristically cold and hot spots are then manifested as extreme outliers in the peak values, and can constitute evidence for non-Gaussianity or deviation from isotropy.
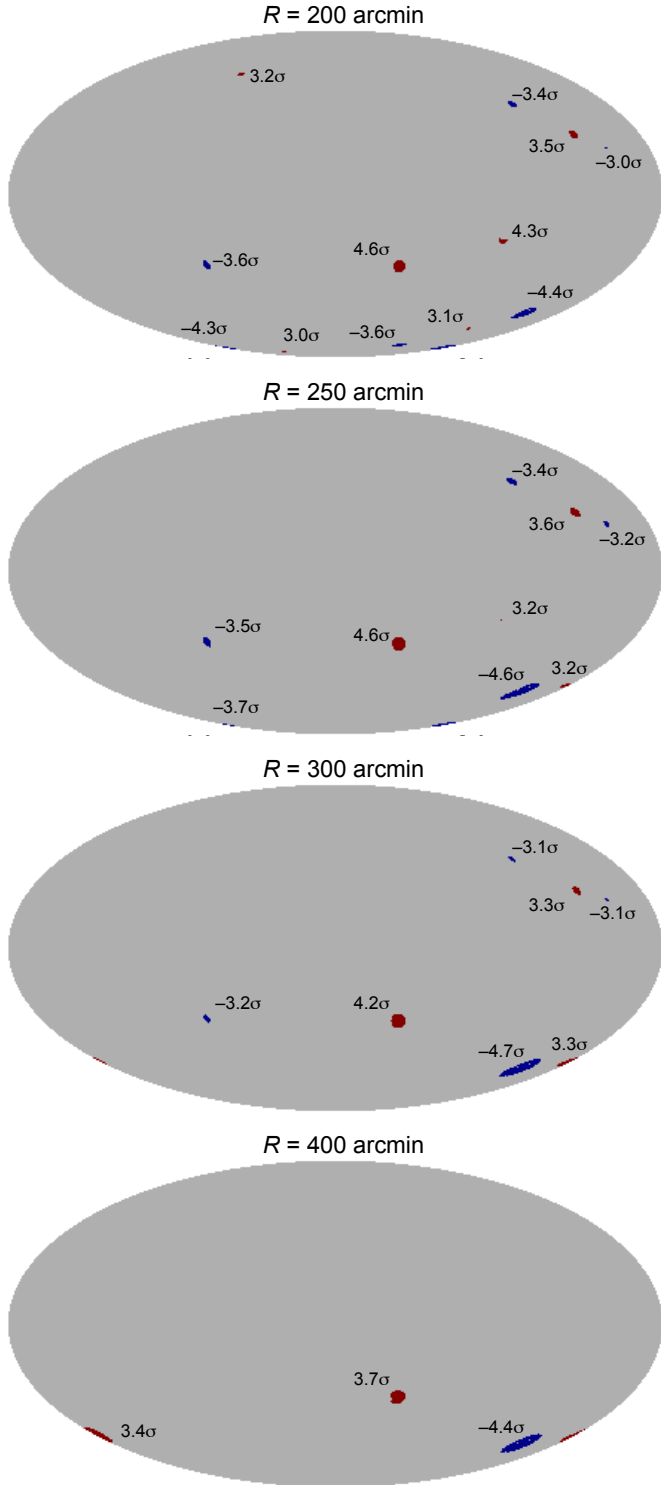
Fig. 9. Cold and hot areas for thresholds $\nu > 3.0$ as determined from the SEVEM temperature map. From *top to bottom*, the maps are for SMHW scales of $R = 200'$, $R = 250'$, $R = 300'$, and $R = 400'$.

Table 9. Peak counts in maps filtered to different scales.

Filter scale [arcmin]	Number of minima/maxima				
	Comm.	NILC	SEVEM	SMICA	Match
		SMHW			
200	176/187	170/178	173/182	169/182	161/169
250	105/105	104/103	107/123	105/107	97/ 99
300	70/ 70	71/ 70	70/ 72	68/ 71	66/ 66
400	43/ 32	46/ 32	44/ 31	43/ 33	37/ 30
		GAUSS			
200	152/170	152/166	157/179	156/165	142/155
250	102/ 93	104/ 95	108/ 99	99/101	92/ 85
300	60/ 63	57/ 62	63/ 64	56/ 62	50/ 53
400	33/ 28	29/ 29	31/ 33	29/ 28	24/ 27
		SSG84			
200	180/187	178/183	180/185	183/183	167/175
250	131/119	118/114	122/123	121/110	109/103
300	68/ 69	73/ 68	73/ 73	70/ 68	56/ 61
400	29/ 35	29/ 36	29/ 32	30/ 38	27/ 27

order to avoid potential contamination by boundary effects, the mask is extended by rejecting pixels with an effective convolution weight that differs from unity by more than 12%. Peaks are extracted from the filtered map (removing any that are adjacent to masked pixels), their positions and values are recorded for further analysis, and their cumulative density function (CDF) is constructed by sorting peak values. Table 9 presents peak counts for the component-separated sky maps for several different kernels and representative filtering scales, together with the number of peaks that are common to all maps. There is excellent agreement between the various CMB estimates. All statistical inference is then performed by comparison of the peak distributions derived from the data with equivalently processed simulations. As an internal consistency check, the properties of the FFP8 simulations are found to be in agreement with the predictions of Eq. (27).

Figure 10 presents the distributions of peaks for the SMICA CMB map filtered with two representative kernels on scales of 40′ and 800′ FWHM. The lower panels show the empirical peak CDFs as a function of peak value x, defined for a set of n peaks at values $\{X_i\}$ as

$$F_n(x) = \frac{1}{n} \sum_{i=1}^{n} I_{X_i \leq x}, \quad I_{X_i \leq x} \equiv \begin{cases} 1, & \text{if } X_i \leq x \\ 0, & \text{otherwise.} \end{cases} \quad (28)$$

For plotting purposes alone, the horizontal axis is scaled in units of σ defined by Eq. (27) and derived from the underlying median CDF, $\bar{F}(x)$, of the simulations. The upper panels show the difference between the observed and median simulated CDF values, $\sqrt{n}\,[F_n(x) - \bar{F}(x)]$, with the grey bands representing the 68.3%, 95.4%, and 99.7% regions of the simulated CDF distributions. The maximal value of this difference defines a Kolmogorov-Smirnov (KS) deviation estimator:

$$K_n \equiv \sqrt{n} \, \sup_x \left| F_n(x) - \bar{F}(x) \right|. \quad (29)$$

This forms the basis of a standard KS test of consistency between the two distributions. Although the KS deviation has a known limiting distribution, we also derive its CDF directly from the simulations.
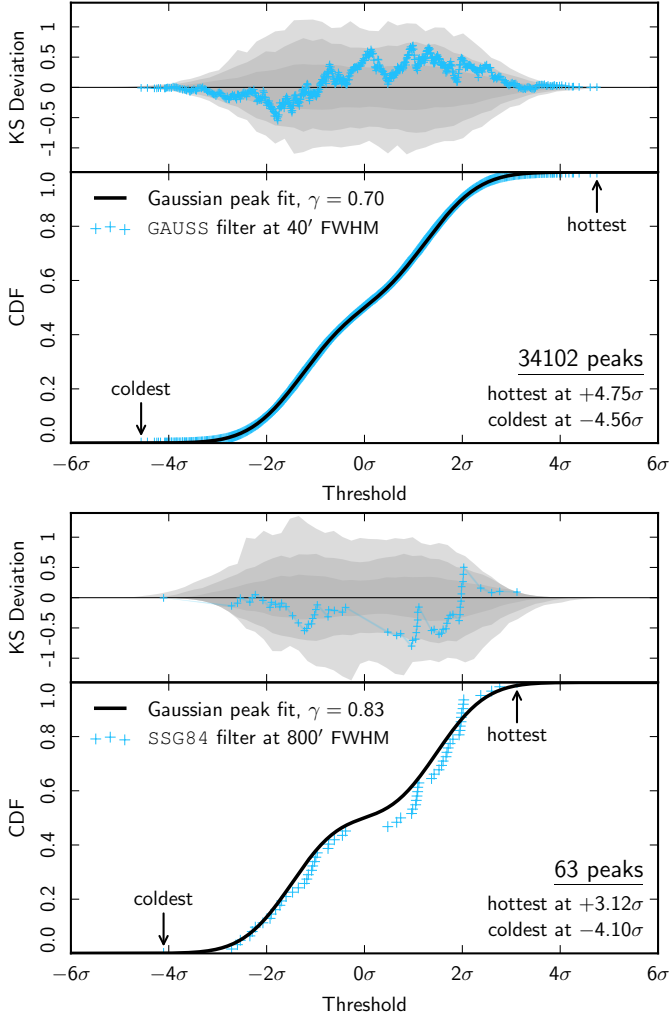
Here, we consider the peak statistics of the *Planck* component-separated temperature maps at $N_{\text{side}} = 2048$. The maps are pre-whitened as described in Appendix A. This step allows the construction of an estimator that is nearly optimal with respect to the fiducial CMB properties. After application of the common mask, weighted convolutions of the data are performed with either SSG or GAUSS kernels of variable scale. In

Fig. 10. Cumulative density function of the peak distribution for the SMICA CMB temperature map. The *top row* shows the peak CDF for maps filtered with a GAUSS kernel of 40′ FWHM. The *bottom row* shows the corresponding peak CDF for an SSG84 kernel of 800′ FWHM. The spectral shape parameter γ (see Eq. (27)) is the best-fit value for the simulated ensemble, as indicated by the cyan circle in Fig. 11. Similar results are obtained for the other component-separation methods.

The temperature peak distributions in Fig. 10 are consistent with Gaussian peak statistics, apart from a single anomalously cold peak on scales around 800′ FWHM. This corresponds to the previously reported Cold Spot. Although this exercise confirms that the Cold Spot is a rare cold feature, as already noted by Cruz et al. (2005) and confirmed in this paper, the most peculiar characteristic of the Cold Spot is not its coldness, but rather its size. A more detailed analysis of its nature is presented in Sect. 5.7.

The probability that the observed sky exceeds the value of the KS deviation for the adopted fiducial cosmology can be determined by counting the number of simulations with $K_{n'} > K_n^{(\text{sky})}$. The *p*-values for the KS test comparing the CDF of the observed sky with the median peak CDF derived from simulations for several different kernels and representative scales are summarized in Table 10. The similarly derived *p*-values for the total peak counts are summarized in Table 11. Most of the results indicate that the two distributions are highly consistent, with the exception of results for the SSG84 filter on scales of about
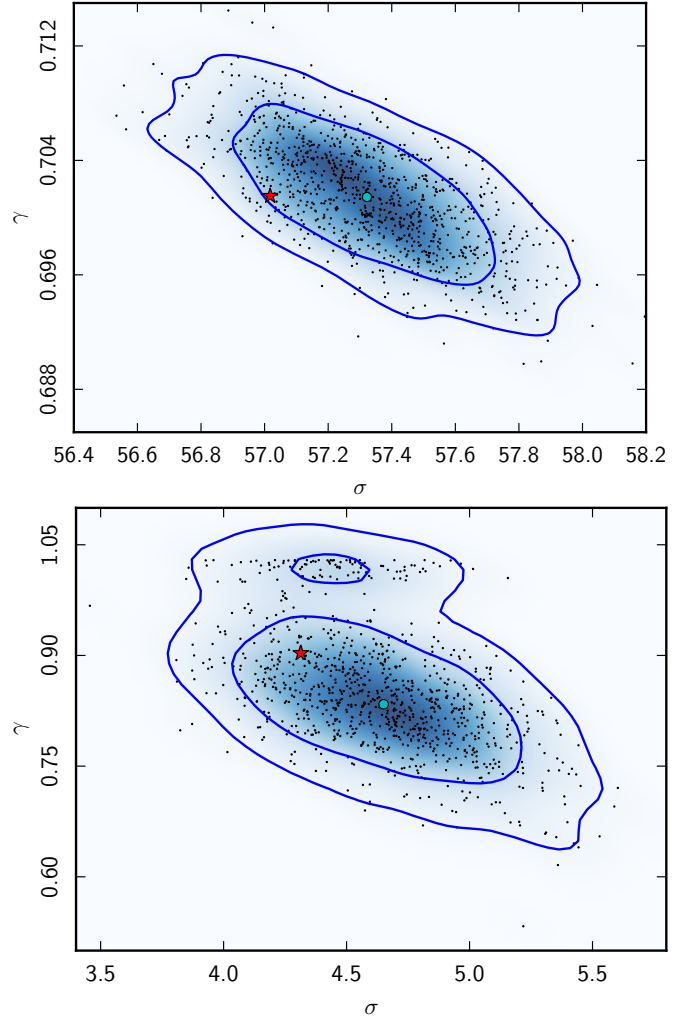

Fig. 11. Distribution of best-fit Gaussian peak CDF spectral shape parameters, σ and γ (as defined in Eq. (27)), recovered from 1000 simulations, as indicated by the black dots and the smoothed density map and compared to those derived for the observed sky (shown by the red star). The blue contours enclose 68% and 95% of the parameter distribution, and the cyan circle represents the best-fit parameters for the median peak CDF determined from simulations. The *upper panel* shows the peak CDF parameters for the SMICA map filtered with a GAUSS kernel of 40′ FWHM. The *lower panel* shows the corresponding peak CDF for an SSG84 kernel of 800′ FWHM. Similar results are obtained for the other component-separation methods.

500′ FWHM. This deviation appears to be related to a hemispherical asymmetry in the peak CDFs, and will be discussed further in Sect. 5.6.

One can also test whether the observed values of the parameters, σ and γ as defined in Eq. (27), are consistent with the simulation ensemble, under the assumption that the peak distributions in the *Planck* data are described by a Gaussian peak CDF. Figure 11 demonstrates the consistency of the best-fit values of these parameters, corresponding to the peak distributions in Fig. 10, with equivalent values derived from the simulations.

Inspired by the analysis of the WMAP first-year data in Larson & Wandelt (2004) which found fewer extreme peaks than expected, we additionally evaluate whether the distributions of maxima and minima are separately consistent with simulations. The mean of all maxima, and the negative of the mean of all minima, are calculated for the filtered map, and the observed
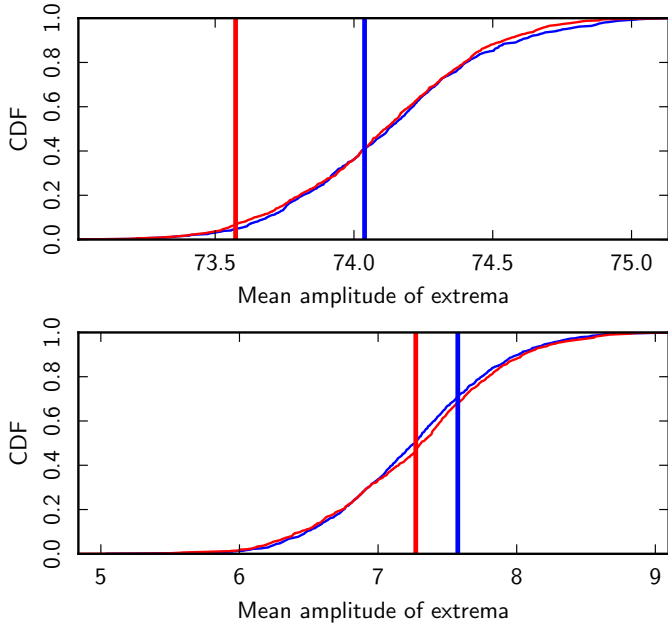
Fig. 12. Cumulative density function of the mean amplitude of all extrema, maxima (red) and minima (blue), derived from simulations, compared to the equivalent values observed for the SMICA CMB temperature map. The *upper panel* shows the peak mean amplitudes for maps filtered with a GAUSS kernel of 40′ FWHM. The *lower panel* shows the corresponding peak CDF for an SSG84 kernel of 800′ FWHM. Similar results are obtained for the other component separation methods. Since the filter kernel normalization is free, and the pre-whitened map to which the filter is applied is dimensionless, the plots are essentially in arbitrary units.

Table 10. Modified upper tail probability (mUTP) for the KS test, comparing the data with the median peak CDF derived from simulations.

Filter scale [arcmin]	Probability [%]			
	Comm.	NILC	SEVEM	SMICA
SMHW				
200	22.0	42.8	45.9	40.5
250	11.3	17.6	3.1	11.4
300	49.4	38.5	38.4	32.1
400	32.6	24.7	35.3	24.7
GAUSS				
200	41.3	46.6	14.4	47.2
250	43.7	34.8	7.6	48.4
300	46.3	9.9	28.0	7.7
400	30.7	5.6	35.8	6.6
SSG84				
200	37.1	36.7	24.0	37.5
250	0.5	1.7	0.8	5.4
300	17.5	12.2	0.3	9.3
400	47.4	44.6	47.5	47.8

values are compared to the simulated distributions in Fig. 12. The observed minima/maxima means are found to be in good agreement with the fiducial values.

The probability that the coldest peak seen on the sky is consistent with the adopted fiducial cosmology is evaluated as a function of both filter shape and size by counting the number of simulations with $x_{\mathrm{coldest}} < x_{\mathrm{coldest}}^{(\mathrm{sky})}$. The results obtained for the SMHW filter are summarized in Fig. 13. Consistent behaviour is seen when the GAUSS and SSG84 filters are applied. The error

Table 11. Modified upper tail probability (mUTP) for the total peak count, comparing the data with the peak count CDF derived from simulations.

Filter scale [arcmin]	Probability [%]			
	Comm.	NILC	SEVEM	SMICA
SMHW				
200	6.1	36.9	16.2	27.2
250	32.9	47.5	1.0	25.6
300	48.8	51.7	44.7	44.3
400	33.8	16.2	34.6	26.4
GAUSS				
200	7.1	11.2	0.7	8.7
250	18.2	11.2	2.1	8.1
300	29.0	12.8	48.2	10.0
400	11.9	3.0	26.6	2.8
SSG84				
200	0.2	3.0	1.0	1.7
250	0.1	1.7	0.1	2.1
300	9.3	22.6	50.7	12.2
400	0.3	0.4	0.1	2.3

bars represent the sampling uncertainty due to the finite number of realizations, and are determined using a bootstrap method. As the filters overlap substantially, different points are highly correlated. The *Planck* CMB maps are consistent with the expectations of a statistically isotropic Gaussian model. The most significant deviation, found at an effective filter bandwidth given by $\ell = 20$, is attributable to a single region on the sky – the Cold Spot.

4.5.4. Peak locations as a function of scale

The application of a filter kernel of variable size to a map extends it into what can be considered a "multiscale space", such that features on different scales are represented by a one-parameter family of smoothed maps. This technique is often used for feature detection and mathematical morphology analysis. Here, we introduce a morphological description of temperature maps based on the peak connectedness graph in multiscale space, and apply this technique to a statistical analysis of the *Planck* CMB data. Like most morphological analyses, it is equally applicable to intrinsically non-Gaussian maps, but here we focus on the Gaussian random field statistics and attempt to understand what features of the CMB temperature map are responsible for the Cold Spot.

To construct a multiscale representation, we trace the location of the peaks in the smoothed, whitened CMB map as the smoothing scale is varied. As the smoothing scale increases, peaks merge and the total peak count decreases. Linking closest peak neighbours in position-scale space, from the finest to the coarsest resolution, produces an acyclic graph that encapsulates the peak "merger tree" history as the scale is varied. A summary of all the peak positions and CDF ranks for the SSG84 filter kernel on scales ranging from 120′ to 1200′ FWHM is shown in Fig. 14. The peaks are represented by discs of varying size (reflecting the filter scale) and colour (reflecting the peak temperature rank), with peaks at all scales projected onto a single map. The lower panel shows the peak linkage graph on the coarser scales; for the statistical analysis 81 filter scales are used, log-spaced from 120′ to 1200′. Peaks of the same type (i.e., maxima to maxima and minima to minima) are linked to
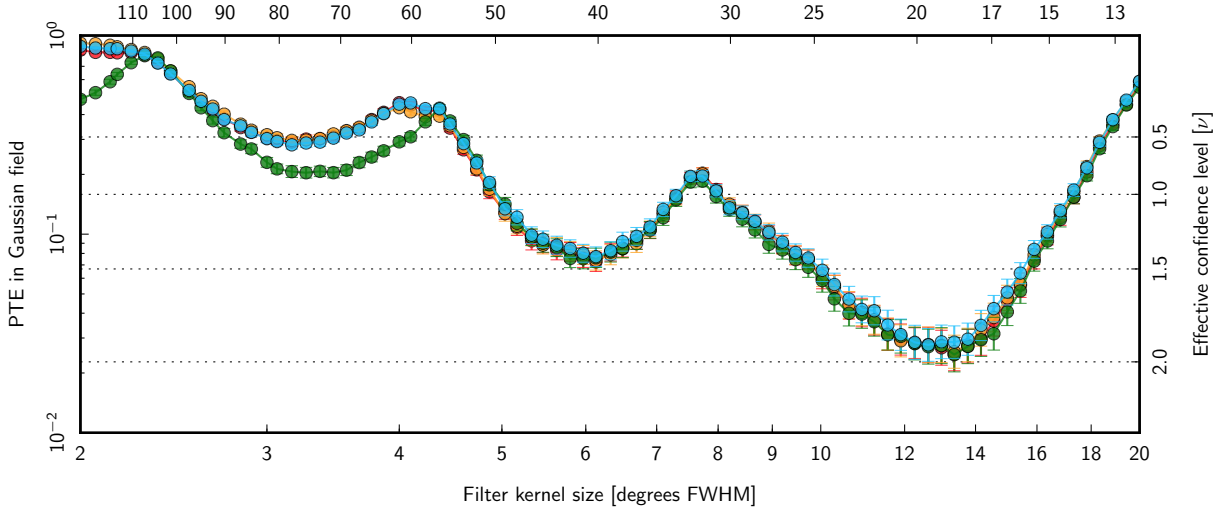
Fig. 13. Fraction of the Gaussian random field realizations in which the coldest peak is as cold as or colder than that observed, as a function of SMHW filter scale for `Commander` (red), `NILC` (orange), `SEVEM` (green), and `SMICA` (blue).
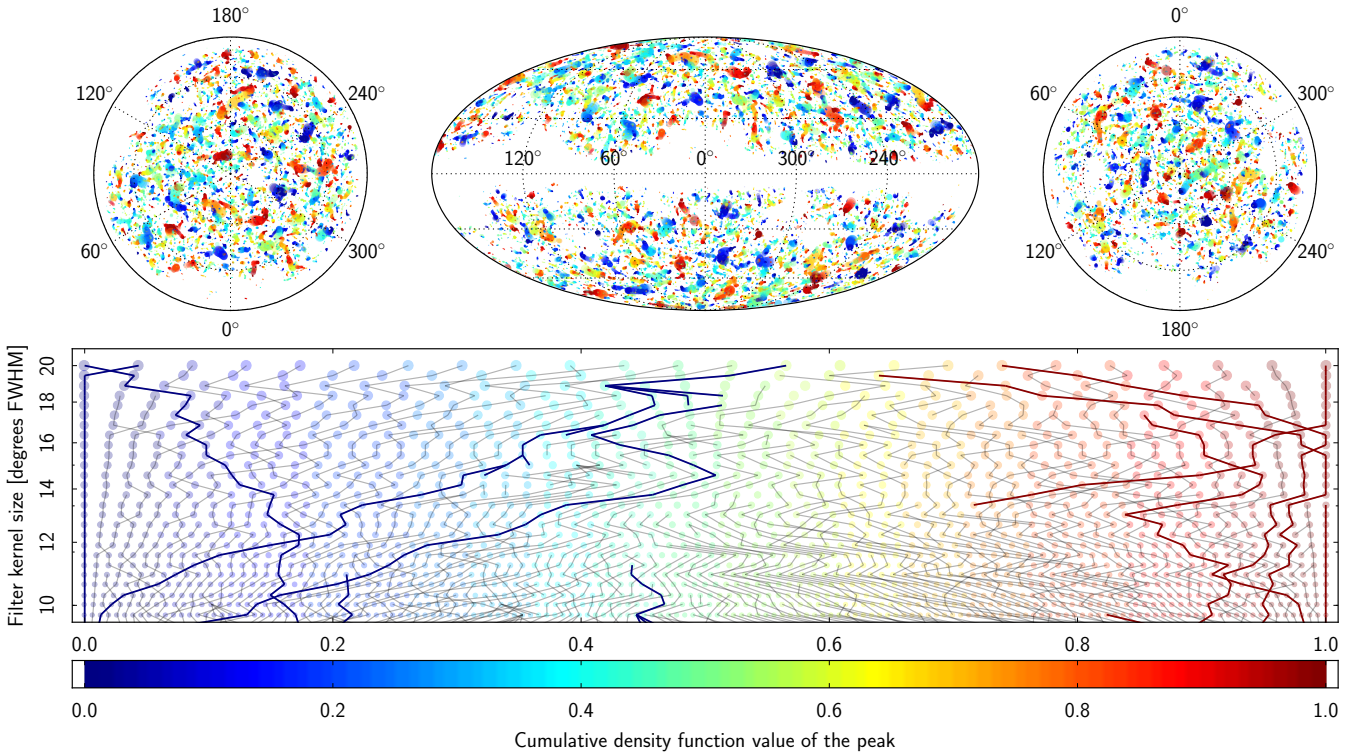

Fig. 14. Peak positions and CDF rank summarized for all filtering scales. The three sky-view panels in the *top row* show a Lambert projection of the north pole, the usual full sky Mollweide projection, and a Lambert projection of the south pole. The *lower panel* shows the peak heights (in percentile of the peak distribution on the horizontal axis) as a function of filter scale (on the vertical axis, in logarithmic scale), truncated to larger scales for clarity. Circles represent peaks (nodes of the graph) coloured according to their percentile level, and scaled according to kernel size. Black lines represent edges connecting peaks at different scales (according to a minimal distance measure). The components connected to the coldest and hottest peaks at any scale are highlighted by thicker edges, and are navy blue and dark red in colour. Note that there are thick lines that do not touch the 0 and 1 percentiles in the plot view. Those edges are connected to extreme percentile values, but at scales smaller than those shown in the plot. The Cold Spot is represented by the connected nodes that have the smallest percentiles except for the coarsest scale in the plot view.

the closest peak on the coarser scale according to a distance measure, $ds^2 + df^2$, where ds^2 is the metric on the unit sphere, and df^2 is the difference of peak temperature ranks (but only if that distance is within a predetermined fraction of the filter scale FWHM).

The resulting peak linkage graph is then analysed for connectedness. The simplest quantifiable measure is the node-degree distribution, shown in Fig. 15 for `SMICA`. The node-degree distribution is highly peaked at 2; this population corresponds to a single peak being traced across multiple scales. Pre-whitening effectively decorrelates the Gaussian map across different scales, so that the resulting node distribution shows a sizeable population of degree 0 and 1 nodes. When compared to the linkage graphs derived from the simulation ensemble, the
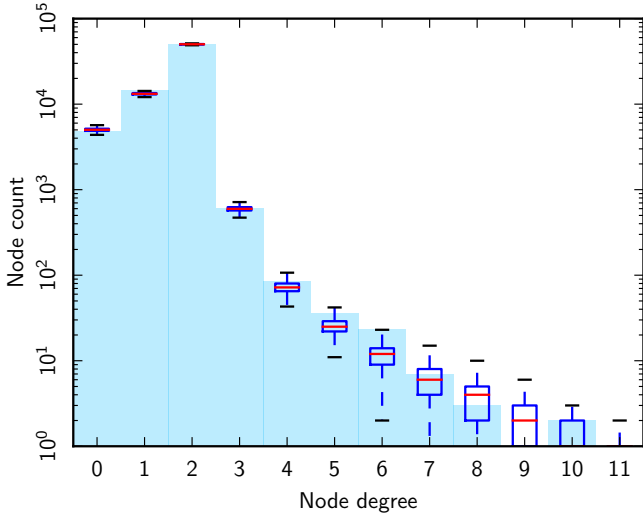
Fig. 15. Distribution of node degrees in the multiscale peak linkage graph determined for the SMICA map (cyan), compared with the median (red line), first to third quartile (blue box), and 95% (whiskers) derived from 1000 FFP8 simulations.

node-degree distribution of the peak linkage graph derived from *Planck* CMB data is consistent, with a slight excess in node counts of degrees 5 and 6.

5. Anomalies in the microwave sky

The previous section established the lack of evidence for significant non-Gaussianity in the *Planck* data. Here we consider several important anomalies that were originally detected in the WMAP sky maps, and later confirmed in the analyses described in PCIS13. Many of these are connected to evidence for a violation of isotropy, or to a preferred direction, in the CMB. Tests that involve dipolar power asymmetry, either directly or via measures of directionality, are collected together in Sect. 6. In this section we consider only those anomalies not directly related to dipolar power asymmetry.

The microwave sky is intrinsically statistically anisotropic due to our motion with respect to the CMB rest frame. The resulting Doppler boosting effect, introduced in Sect. 1, was detected in the 2013 *Planck* data (Planck Collaboration XXVII 2014). For completeness, Appendix B repeats the analysis with the *Planck* full mission data set, though based only on the full velocity estimator (β), which is the sum of the modulation and the aberration contributions. However, since the effects of Doppler boosting are now included in the simulations used for that analysis, this constitutes a consistency check for this release. More importantly, since both the data and simulations now include the effect, it is not necessary to consider deboosted data in many of the studies reported here, unlike in PCIS13 (although one exception in Sect. 6.4 makes use of unboosted simulations to search for the frequency-dependent signature of the effect in the SEVEM-100, SEVEM-143, and SEVEM-217 sky maps). However, we note that some care must be taken due to the absence of the aberration contribution in the simulations. Indeed, this leads to the slightly, but not alarmingly, low PTE for $\beta_\parallel$ in Appendix B. However, we not expect any impact on the results presented in this section.

Before presenting our results, we return to the issue of a posteriori correction, which particle physicists refer to as correcting for the "look-elsewhere effect" (LEE). Since there are many

tests that can be performed on the data to look for a violation of statistical isotropy, we expect some to indicate detections at, for example, roughly 3σ levels, since even a statistically isotropic CMB sky is a realization of an underlying statistical process corresponding to many independent random variables. However, in the absence of an existing theoretical framework (i.e., a physical model) to predict such anomalies, it is difficult to interpret their significance. It is then necessary, and equally challenging, to address the question of how often such detections would be found for statistically isotropic Gaussian skies. Unfortunately, it is not always clear how to answer this question.

There will always be a degree of subjectivity when deciding exactly how to assess the significance of these types of features in the data. As an example, one could argue that the large-scale dipole modulation signal we see is coming specifically from super-Hubble modes, in which case performing an LEE correction for dipole modulation that could have been seen on small scales ($\ell \gtrsim 100$) would not make sense. Models for such a super-Hubble modulation exist and an example was examined in Planck Collaboration XX (2016), the conclusion being that the model could only explain part of the dipole modulation and that the allowed part was perfectly consistent with cosmic variance.

In this paper, we adopt a pragmatic approach. When there is a clear mechanism for doing so, we attempt to correct for the "multiplicity of tests", or the possible ways in which an anomalous signal might have been detected but was not, as a consequence of any a posteriori (data-driven) choices made in searching for it. In such cases, a strong dependence of the significance on the correction would indicate that we should be cautious about the uncorrected result. When such an obvious correction is not possible, we clearly describe the methodology applied to the data and its limitations. With this approach, we also recognize that any statistical assessment is partially subjective, including those that purport to correct for the LEE.

Although many of the observed effects described in this and the next section may elude theoretical prediction today, we continue to highlight them since there is a real possibility that the significance of one or more might increase at a later date, perhaps when polarization data are included in the analysis, and lead to new insights into early Universe physics. Alternatively, such observations may directly motivate the construction of models that can make predictions for features that can be sought in new data sets. This is particularly the case for anomalies on the largest angular scales, which may have a specific connection to inflation.

5.1. Variance, skewness, kurtosis

Previous analyses of the WMAP data (Monteserín et al. 2008; Cruz et al. 2011; Gruppuso et al. 2013) have reported that the variance of the CMB sky is lower than that of simulations based on the ΛCDM model. PCIS13 confirmed this, and proposed a possible explanation of the apparent incompatibility of the observed variance with a fiducial cosmological model that has been determined from the same data set. Specifically, whilst the map-based variance is dominated by contributions from large angular scales on the sky, the cosmological parameter fits are relatively insensitive to these low-order ℓ-modes, and are instead largely dominated by scales corresponding to $\ell > 50$. Therefore the variance of the map appears to be anomalous, since there is a dearth of large-angular-scale power compared to the model.

In Sect. 4.1, we again confirmed the presence of low variance in the data. Here, we extend the analysis to investigate which angular scales are responsible for the low variance by applying the
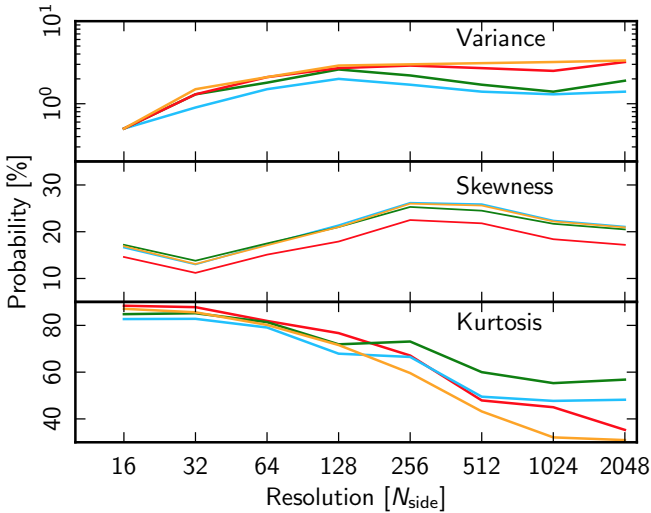
Fig. 16. Lower tail probability of the variance (*top panel*), skewness (*centre panel*), and kurtosis (*bottom panel*) obtained at different resolutions from the Commander (red), NILC (orange), SEVEM (green), and SMICA (blue) sky maps.

Table 12. Lower-tail probability for the variance, skewness, and kurtosis of the low resolution $N_{\rm side} = 16$ component-separated maps obtained with the common mask and two extended versions thereof.

Method	Probability [%]		
	Variance	Skewness	Kurtosis
Common mask ($f_{\rm sky} = 58\%$)			
Commander	0.5	14.6	88.4
NILC	0.5	16.9	87.1
SEVEM	0.5	17.2	84.8
SMICA	0.5	16.6	82.7
$f_{\rm sky} = 48\%$			
Commander	0.1	29.4	65.0
NILC	0.1	29.6	60.8
SEVEM	0.1	29.4	62.4
SMICA	0.1	29.4	57.3
$f_{\rm sky} = 40\%$			
Commander	0.4	35.2	32.4
NILC	0.4	34.4	28.7
SEVEM	0.4	34.3	30.2
SMICA	0.4	33.8	25.5

unit variance estimator to lower resolution component-separated maps, specifically those from $N_{\rm side} = 1024$ to $N_{\rm side} = 16$, with the corresponding common masks, and then comparing the results with those determined from 1000 MC simulations. The results are shown in Fig. 16.

All of the component-separation methods that we consider yield very consistent results which indicate an increasingly anomalous low variance at lower resolutions, with the lower-tail probability reaching a minimum value of 0.5% at $N_{\rm side} = 16$. We then consider the impact of a possible look-elsewhere effect by evaluating the minimum lower tail probability of each simulation irrespective of the $N_{\rm side}$ resolution at which it occurs. By comparing the distribution of these values with that of the data, we infer that the probability is slightly weakened to a value of about 1%. These results are compatible with a lack of power on large angular scales. Since the variance estimator is heavily weighted towards low-ℓ modes, this has an increasing impact on the observed variance when going from high to low resolution sky maps. Conversely, the skewness and kurtosis are consistent with the simulations, although there is some indication of a weak scale-dependence (albeit at low significance).

We also investigate the stability of the results at $N_{\rm side} = 16$ with respect to the possible presence of residual foregrounds by considering two additional masks obtained by extending the edge of the $N_{\rm side} = 16$ common mask by 5° and 9°, reducing the usable sky fraction from 58% to 48% and 40%, respectively. We then re-apply the unit variance estimator to the low resolution component-separated maps with these masks and determine the variance, skewness, and kurtosis values (see Table 12).

The results from 48% of the sky reveal that only 1 simulation in 1000 is found to be more anomalous (i.e., exhibit lower variance) than the observed map. In addition, both the skewness and kurtosis become more compatible with the ΛCDM model. With the more aggressive mask, the lower-tail probability slightly increases again. However, given the limited number of pixels involved in the analysis, this shift may be related to the effects of sample variance.

Overall, our results may be explained by the presence of a low-variance anomaly in the primordial CMB signal – the stability of the low-variance significance argues against foreground

contamination being responsible for the lack of observed power. This is reinforced by the decrease in variance when regions close to the common mask borders, where foreground residuals are most likely to be observed, are omitted from the analysis.

5.2. N-point correlation function anomalies

5.2.1. Lack of large-angle correlations

We first reassess the lack of correlation seen in the 2-point angular correlation function at large angular separations as reported in Sect. 4.3, and previously noted for both WMAP and the 2013 *Planck* temperature maps (Bennett et al. 2003; Copi et al. 2015). We attempt to quantify this lack of structure using the statistic proposed by Spergel et al. (2003):

$$S(x) = \int_{-1}^{x} \left[\hat{C}_2(\theta) \right]^2 \, {\rm d}(\cos\theta) , \qquad (30)$$

where $\hat{C}_2(\theta)$ is our estimate of the 2-point correlation function. Generally, the upper limit on the integral has been taken to correspond to a separation angle of 60°, possibly (as noted by Copi et al. 2009) motivated by the COBE-DMR 4-year results (Hinshaw et al. 1996). Inspection of the top panel of Fig. 2 suggests that the *Planck* 2-point function lies close to zero between 80° and 170°, but for consistency with previous work we compute the statistic $S_{1/2}$, for $x \equiv \cos 60° = \frac{1}{2}$. The results are presented in Table 13. We find that the data indeed show a lack of correlations on large angular scales, with a significance consistent with that found by Copi et al. (2015) (although note that the sense of the p-values differs between the papers).

Possible criticisms of the $S_{1/2}$ statistic include that it has been designed a posteriori to test for a lack of large-angle correlations, and that it does not account for the high degree of correlation between bins at different angular scales. We can address these concerns, at least in part, by considering a modified version of the commonly used and well understood χ^2 statistic used in previous studies. In order to test the same hypothesis as the $S_{1/2}$ statistic – that there are no correlations on scales larger than some angular cut-off – we do not subtract an averaged 2-point

Table 13. Probabilities of obtaining values for the $S_{1/2}$ and χ_0^2 statistics for the *Planck* fiducial ΛCDM model at least as large as the observed values of the statistic for the *Planck* 2015 temperature CMB maps with resolution parameter $N_{side} = 64$, estimated using the `Commander`, `NILC`, `SEVEM`, and `SMICA` maps.

	Probability [%]			
Statistic	Comm.	NILC	SEVEM	SMICA
$S_{1/2}$	99.5	99.6	99.5	99.6
$S(x)$ (global)	97.7	97.8	97.8	97.9
$\chi_0^2(\theta > 60°)$	98.1	98.8	98.1	98.4

Notes. We show also the corresponding estimation of the global *p*-value for the $S(x)$ statistic.

function when computing the χ^2, i.e., we use a statistic defined as

$$\chi_0^2(\theta_{min}, \theta_{max}) = \sum_{i,j=i_{min}}^{i_{max}} \hat{C}_2(\theta_i) M_{ij}^{-1} \hat{C}_2(\theta_j), \qquad (31)$$

where i_{min}, i_{max} denote the index of the bins corresponding to the minimum and maximum value of the separation angles θ_{min} and θ_{max}, respectively. In this analysis, we adopt $\theta_{min} = 60°$ and $\theta_{max} = 180°$. M_{ij} is the covariance matrix given by Eq. (8), estimated using MC simulations corresponding to the fiducial ΛCDM model. The results are shown in Table 13. The significance level of the anomaly is slightly smaller for the χ_0^2 statistic than that derived with $S_{1/2}$. We note that this statistic is closely related to the $A(x)$ measure proposed by Hajian (2007).

A further potential criticism of the $S_{1/2}$ statistic relates to the a posteriori choice of 60° for the separation angle that delineates the interesting region of behaviour of the correlation function. We therefore consider the generalized statistic $S(x)$ and compute its value for all values of x, both for the data and for the simulations. Then, for each value of x, we determine the number of simulations with a higher value of $S(x)$, and hence infer the most significant value of the statistic and the separation angle that it corresponds to. However, since such an analysis is sensitive to the LEE, we define a global statistic to evaluate the true significance of the result. Specifically, we repeat the procedure for each simulation, and search for the largest probability irrespective of the value of x at which it occurs. The fraction of these probabilities higher than the maximum probability found for the data defines a global *p*-value. As seen in Table 13, this corresponds to values of order 98% for all of the CMB estimates.

The previous analyses essentially test how consistent the observed 2-point correlation function data is with a lack of correlations on large angular scales, in particular for separation angles $\theta > 60°$. A conventional χ^2 statistic allows us to test the consistency of this quantity with the predictions of the ΛCDM model. In this case, the statistic is defined as in Eq. (7), except that we constrain the computations to those bins that correspond to the intervals defined by $\theta < 60°$ and $\theta > 60°$. The results of these studies are shown in Table 14.

The analysis for $\theta < 60°$ indicates that the observed 2-point function is a good match to the mean 2-point function predicted by the ΛCDM model. Moreover, for $\theta > 60°$ the results suggest that the problem is that the fit of the data to the model is too good, and this is even more pronounced for an analysis in the full separation angle range.

Overall, the tests indicate an unusually good fit of the observed 2-point function both to zero and to the predictions of the ΛCDM model for angles above 60°. This problem may be

Table 14. Probabilities of obtaining values for the χ^2 statistic for the *Planck* fiducial ΛCDM model at least as large as the observed values of the statistic for the *Planck* 2015 temperature CMB maps with resolution parameter $N_{side} = 64$, estimated using the `Commander`, `NILC`, `SEVEM`, and `SMICA` maps.

	Probability [%]			
Statistic	Comm.	NILC	SEVEM	SMICA
$\chi^2(\theta < 60°)$	91.5	93.3	91.6	91.7
$\chi^2(\theta > 60°)$	96.8	98.3	96.9	98.1

related to the fact that the theoretical variance for the best-fit model is larger than the observed value at large scales, so that the simulations based on this model that have been used in all of the statistical tests may overestimate the variance of the 2-point function.

5.2.2. Hemispherical asymmetry

We now turn to a reassessment of the asymmetry between the real-space *N*-point correlation functions computed on hemispheres reported previously for the WMAP (Eriksen et al. 2005) and *Planck* 2013 temperature maps (PCIS13). We initially focus the analysis on the hemispheres determined in the ecliptic coordinate frame for which the largest asymmetry was observed. However, we also carry out the corresponding calculations in other relevant reference frames, such as those defined by the Doppler boost (DB, see Sect. 6.4, Appendix B, and Planck Collaboration XXVII 2014) and the dipole modulation (DM, see Sect. 6.2) directions. We use the same configurations of the *N*-point functions as described in Sect. 4.3. However, here the functions are not averaged over the full sky and depend on a choice of specific direction, so they constitute tools for studying statistical isotropy rather than non-Gaussianity (Ferreira & Magueijo 1997).

As in Sect. 4.3, we analyse the CMB estimates at a resolution of $N_{side} = 64$ and quantify their agreement with the fiducial cosmological model using a χ^2 statistic. The results determined from the *Planck* 2015 temperature data for the ecliptic hemispheres are shown in Fig. 17. If we consider that the χ^2 statistic itself can act as a measure of fluctuation level, then asymmetry between the two measured hemispheres can be quantified by the ratio of the corresponding χ^2 values. The probabilities of obtaining values of the χ^2 statistic or ratio for the *Planck* fiducial ΛCDM model at least as large as the observed values are given in Table 15. Since we do not have any predictions concerning the behaviour of a given hemisphere, in the case of the χ^2 ratios we provide the complementary probabilities of the 2-tailed statistic.

The significance levels of the 3- and 4-point functions in the northern hemisphere are nominally very high, exceeding 99.9% for the pseudo-collapsed 3-point function. However, proper interpretation requires that one recognize that the analysis is affected by a posteriori choices for the smoothing scale and reference frame defining the hemispheres. This typically leads to an overestimation of the significance of the results. Accounting for such effects requires the repetition of the analysis for all possible reference directions and also for data at other resolutions. Unfortunately, because of computational limitations, such an extended analysis is not possible for these higher-order statistics. Nevertheless, the observed properties of the *Planck* data are consistent with a clear lack of fluctuations in a direction towards the north ecliptic pole. However, the χ^2-ratio statistic indicates

Table 15. Probabilities of obtaining values for the χ^2 statistic and ratio of χ^2 of the N-point functions shown in Fig. 17 for the *Planck* fiducial ΛCDM model at least as large as the observed values of the statistic for the *Planck* 2015 CMB maps estimated on northern and southern ecliptic hemispheres.

	Probability [%]			
Hemisphere	Comm.	NILC	SEVEM	SMICA
2-point function				
Northern	89.7	90.6	89.8	88.0
Southern	80.5	82.7	82.9	77.6
χ^2-ratio	22.6	21.0	19.7	22.3
Pseudo-collapsed 3-point function				
Northern	>99.9	>99.9	>99.9	99.7
Southern	35.1	34.9	35.8	31.4
χ^2-ratio	98.8	98.5	98.5	98.4
Equilateral 3-point function				
Northern	98.6	98.6	98.8	98.4
Southern	45.7	45.7	47.8	42.6
χ^2-ratio	86.6	86.7	86.6	86.7
Rhombic 4-point function				
Northern	99.7	99.7	99.7	99.6
Southern	22.8	22.5	23.2	20.1
χ^2-ratio	97.3	97.1	97.2	97.0

Table 16. Probabilities of obtaining values for the χ^2 statistic and ratio of χ^2 of the N-point functions shown in Fig. 18 for the *Planck* fiducial ΛCDM model at least as large as the observed values of the statistic for the SMICA map on hemispheres defined by the ecliptic (*first column*), Doppler boost (DB, *second column*), and dipolar modulation (DM, *third column*) reference frames.

	Probability [%]		
Hemisphere	Ecl.	DB	DM
2-point function			
Negative	88.0	86.9	61.8
Positive	77.6	91.1	59.9
χ^2-ratio	22.3	5.1	7.7
Pseudo-collapsed 3-point function			
Negative	99.7	64.1	95.9
Positive	31.4	79.3	48.3
χ^2-ratio	98.4	23.3	78.6
Equilateral 3-point function			
Negative	98.4	54.8	>99.9
Positive	42.6	95.0	78.4
χ^2-ratio	86.7	67.7	88.2
Rhombic 4-point function			
Negative	99.6	46.4	97.5
Positive	20.1	86.3	23.2
χ^2-ratio	97.0	57.9	92.5

a slightly smaller significance level for the asymmetry, not exceeding 99% for any of the N-point functions.

The results for the N-point correlation functions determined in the DB and DM reference frames for the SMICA map are shown in Fig. 18 and the probabilities are presented in Table 16. Note that the positive hemisphere for the ecliptic reference frame corresponds to the southern hemisphere in the previous table. Whilst the largest asymmetry is seen in ecliptic coordinates, a substantial asymmetry is present also for the DM direction. This can be explained by the fact that the DM direction is more closely aligned with the south ecliptic pole (with a separation of around 47°) than the DB direction is. For the DB direction we do not find any significant asymmetry. The equivalent results for Commander, NILC, and SEVEM are consistent with those shown here.

In conclusion, the correlation functions for the *Planck* 2015 temperature data are consistent with the results presented in PCIS13. Specifically, we observe that the northern hemisphere correlation functions are relatively featureless (both the 3- and 4-point functions lie very close to zero), whereas the southern hemisphere functions exhibit a level of structure consistent with Gaussian simulations.

5.3. Constraints on quadrupolar modulation

The most natural extension of the class of statistically anisotropic models that we have considered previously involves the quadrupolar modulation of an initially statistically isotropic CMB sky map. No detection of a corresponding quadrupolar power asymmetry is currently claimed. An initial BipoSH analysis of the WMAP 7-year data (Bennett et al. 2011) found evidence of corresponding non-zero spectra, $A_{\ell\ell}^{20}$ and $A_{\ell\ell+2}^{20}$, in ecliptic coordinates. However, Hanson et al. (2010) demonstrated that the signal could be attributed to an incomplete treatment of beam asymmetries in the data, and this was subsequently confirmed in Bennett et al. (2013). The corresponding analysis of the *Planck* 2013 data indicated consistency with statistical isotropy (Planck Collaboration XXIII 2014).

Here, we proceed further and consider the quadrupolar modulation of the primordial power spectrum as suggested by Ackerman et al. (2007):

$$P(\boldsymbol{k}) = P(k)\left[1 + \sum_M g_{2M}\, Y_{2M}(\hat{\boldsymbol{k}})\right]. \tag{32}$$

Given such a spectrum, the CMB temperature field is expected to exhibit a correlation between $a_{\ell m}$ and $a_{\ell+\Delta\, m'}^*$ with $\Delta = 0, 2$. Therefore, the BipoSH coefficients $A_{\ell\ell}^{2M}$ and $A_{\ell\ell+2}^{2M}$ are sensitive to g_{2M}. In the limit of weak anisotropy, Kim & Komatsu (2013) proposed an optimal estimator for g_{2M}:

$$\hat{g}_{2M} = \frac{1}{2}\sum_{M'}\left(\mathsf{F}^{-1}\right)_{MM'}\sum_{\ell m}\sum_{\ell' m'}\left.\frac{\partial \mathsf{C}_{\ell m, \ell' m'}}{\partial g_{2M'}}\right|_{g_{2M}=0} \tag{33}$$
$$\times\left[(\mathsf{C}^{-1}\boldsymbol{a}^*)_{\ell m}(\mathsf{C}^{-1}\boldsymbol{a})_{\ell' m'}\right.$$
$$\left.-\langle(\mathsf{C}^{-1}\boldsymbol{a}^*)_{\ell m}(\mathsf{C}^{-1}\boldsymbol{a})_{\ell' m'}\rangle\right]_{g_{2M}=0},$$

where $\boldsymbol{a}$ is the CMB data vector in harmonic space and C is its covariance matrix, and

$$\mathsf{F}_{MM'} \equiv \frac{1}{2}\sum_{\ell m}\sum_{\ell' m'}\left[(\mathsf{C}^{-1})_{\ell m}\frac{\partial \mathsf{C}_{\ell m, \ell' m'}}{\partial g_{2M}}(\mathsf{C}^{-1})_{\ell' m'}\frac{\partial \mathsf{C}_{\ell' m', \ell m}}{\partial g_{2M'}}\right]_{g_{2M}=0}. \tag{34}$$

Here, $\langle(\mathsf{C}^{-1}\boldsymbol{a}^*)_{\ell m}(\mathsf{C}^{-1}\boldsymbol{a})_{\ell' m'}\rangle_{g_{2M}=0}$ is the mean field in the absence of the quadrupolar modulation. Observation-specific issues such as incomplete sky coverage, inhomogeneous noise, and asymmetric beams will result in a non-zero mean field, which can be estimated for the *Planck* data using simulations. Due to the otherwise prohibitive computational cost, we adopt a diagonal approximation for the inverse of the covariance matrix:

$$(\mathsf{C}^{-1})_{\ell m, \ell' m'} \approx 1/(C_\ell + N_\ell)\,\delta_{\ell\ell'}\delta_{mm'}, \tag{35}$$

where C_ℓ and N_ℓ are the signal and noise power spectra respectively. Uncertainties are computed by applying the estimator to simulations.
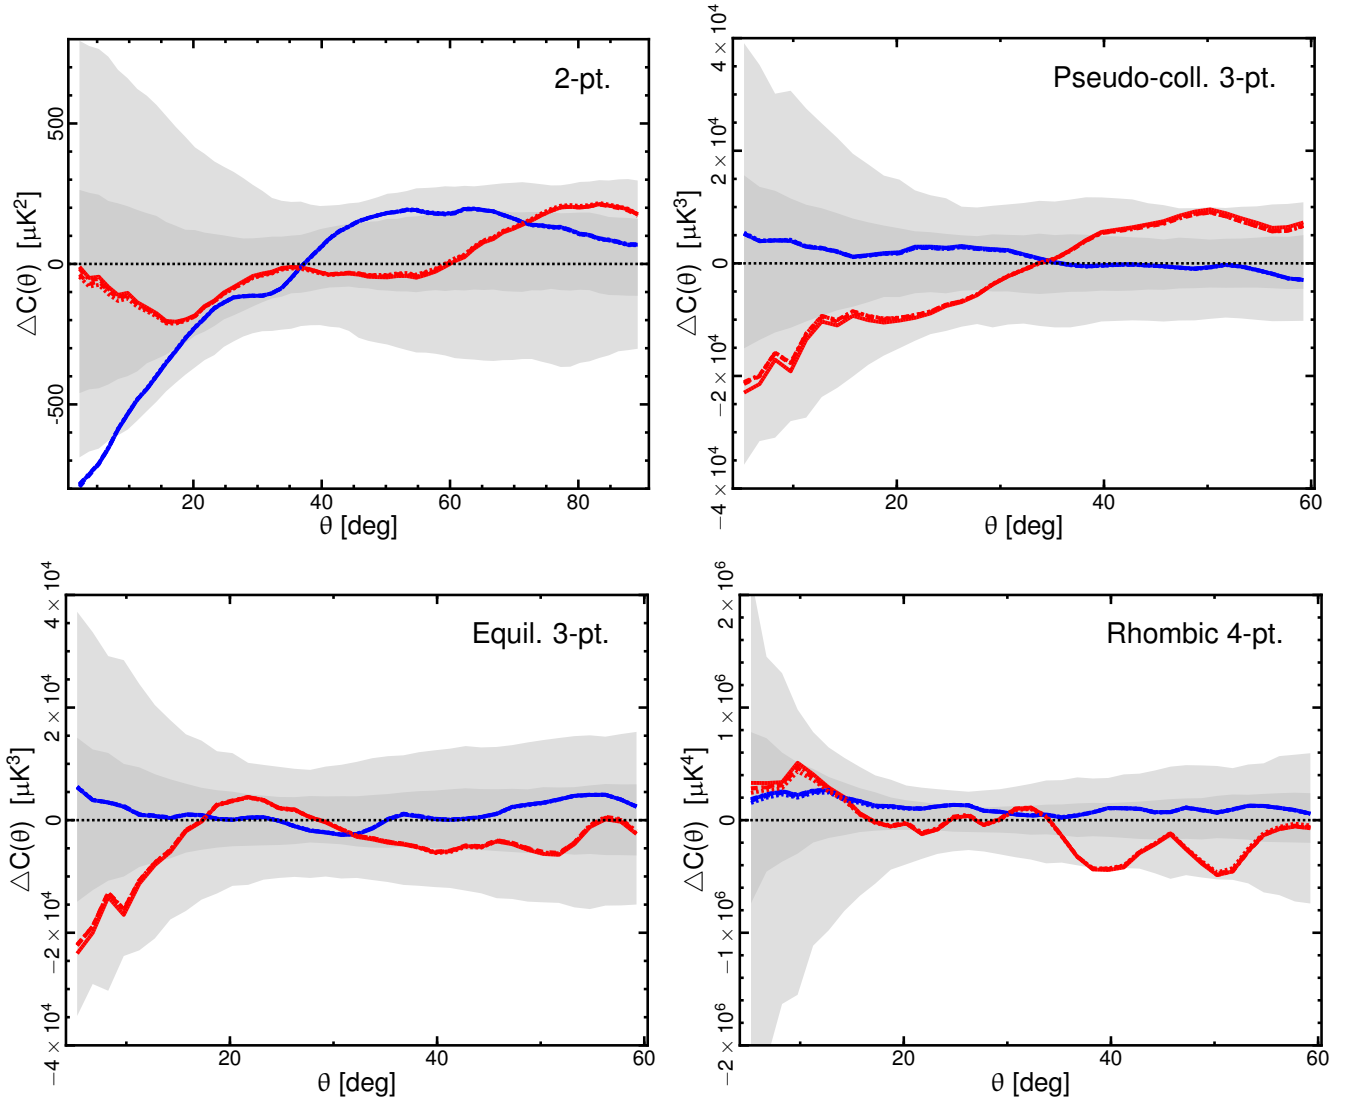
Fig. 17. Difference of the *N*-point correlation functions determined from the N_{side} = 64 *Planck* CMB 2015 temperature estimates and the corresponding means estimated from 1000 simulations. Results are shown for the 2-point, pseudo-collapsed 3-point (*upper left and right panels*, respectively), equilateral 3-point, and connected rhombic 4-point functions (*lower left and right panels*, respectively). Correlation functions are shown for the analysis performed on northern (blue) and southern (red) hemispheres determined in the ecliptic coordinate frame. The solid, dashed, dot-dashed, and dotted lines correspond to the Commander, NILC, SEVEM, and SMICA maps, respectively. Note that the lines lie on top of each other. The shaded dark and light grey regions indicate, for reference, the 68% and 95% confidence regions, respectively, determined from the SMICA simulations.

Table 17. Constraints on the quadrupolar modulation, determined from the Commander, NILC, SEVEM, and SMICA foreground-cleaned maps.

| Method | $g_{2M} \times 10^2$ | | | $g_2 \times 10^2$ | |
	$M = 0$	$M = 1$	$M = 2$	Data	Simulation
Commander	1.31 ± 1.22	$(0.43 \pm 0.86) + i(-0.01 \pm 0.68)$	$(1.08 \pm 0.89) + i(-0.38 \pm 0.86)$	0.97	1.12 ± 0.37
NILC	0.88 ± 1.21	$(0.37 \pm 0.85) + i(\ 0.33 \pm 0.67)$	$(0.87 \pm 0.88) + i(-0.26 \pm 0.86)$	0.77	1.11 ± 0.37
SEVEM	0.85 ± 1.22	$(0.35 \pm 0.85) + i(\ 0.34 \pm 0.67)$	$(1.00 \pm 0.88) + i(-0.25 \pm 0.86)$	0.81	1.11 ± 0.37
SMICA	1.10 ± 1.10	$(0.46 \pm 0.81) + i(\ 0.26 \pm 0.64)$	$(0.93 \pm 0.83) + i(-0.26 \pm 0.82)$	0.85	1.05 ± 0.34

Notes. The *first three columns* correspond to the five independent parts of the quadrupolar modulation, which we have chosen to present using a complex notation for g_{2M}. The quoted error bars are at the 68% confidence level. The quadrupolar modulation amplitude is given in the *fourth column*, while the mean and standard deviation of g_2, estimated from simulations, are provided in the fifth column.

Table 17 presents results from an analysis of the *Planck* data using the extended common mask, UTA76, and limiting the range of multipoles to $2 \leq \ell \leq 1200$. When including data at higher ℓ-values, simulations show evidence for large statistical uncertainties in the recovered g_{2M} values that are a consequence of the many holes in the mask related to point
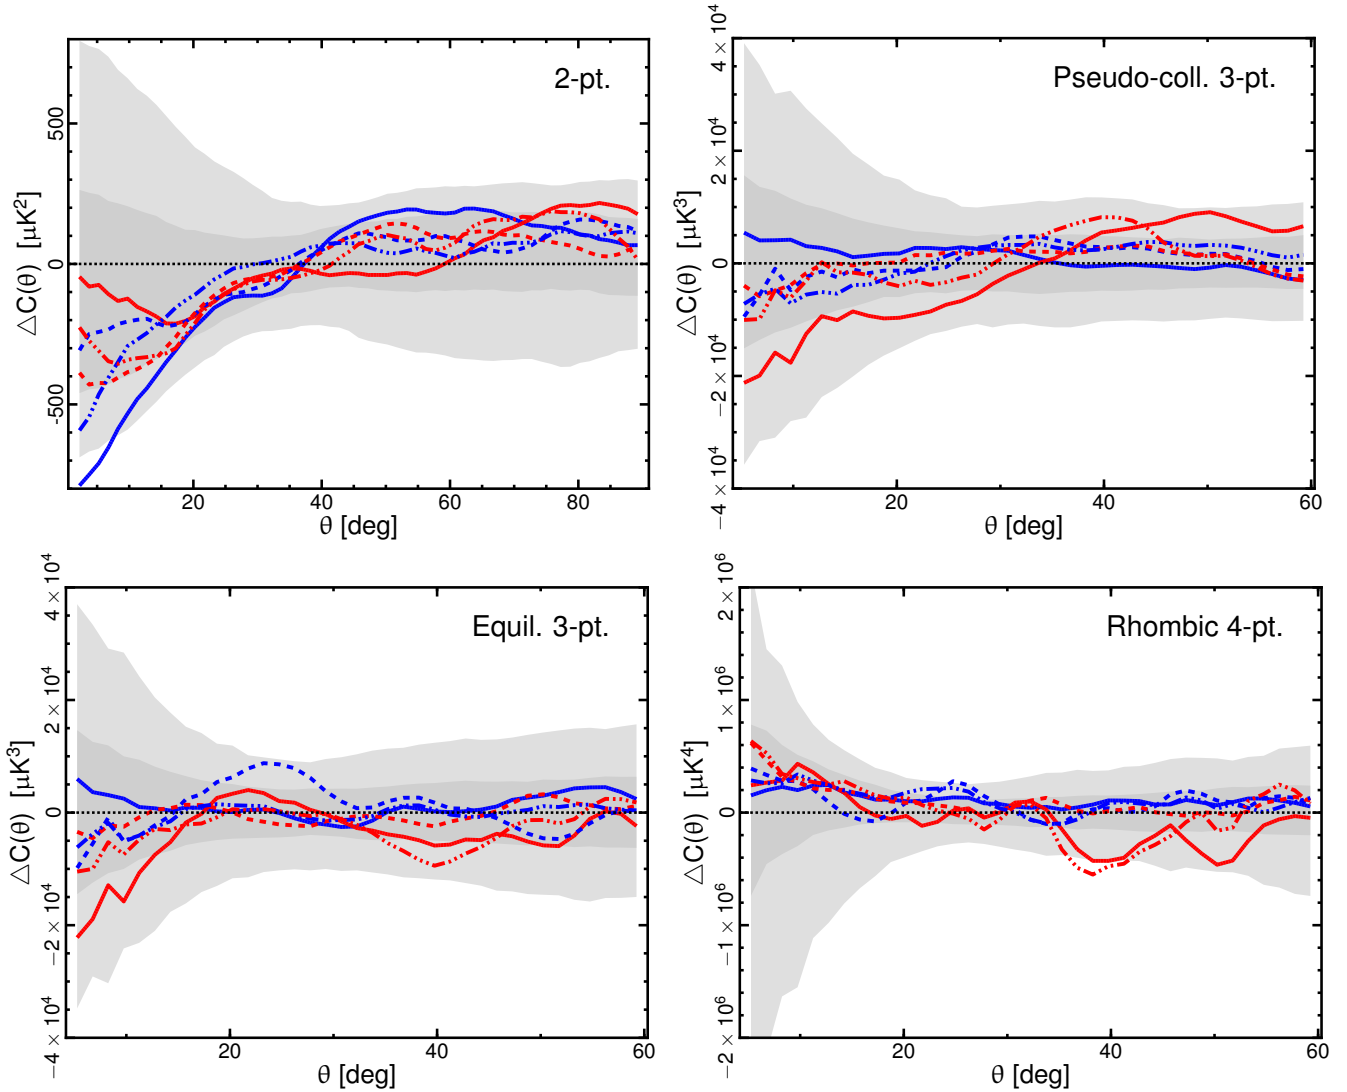
Fig. 18. Difference of the *N*-point correlation functions determined from the $N_{\text{side}} = 64$ *Planck* `SMICA` CMB 2015 temperature estimates and the corresponding means estimated from 1000 simulations. Results are shown for the 2-point, pseudo-collapsed 3-point (*upper left and right panels*, respectively), equilateral 3-point, and connected rhombic 4-point functions (*lower left and right panels*, respectively). Correlation functions are shown for the analysis performed on negative (blue) and positive (red) hemispheres determined in the ecliptic (solid lines), Doppler boost (DB, dashed lines), and dipole modulation (DM, dot-dashed lines) coordinate frames. The shaded dark and light grey regions indicate the 68% and 95% confidence regions, respectively.

sources. Therefore, imposing this limit $\ell \leq 1200$ does not significantly affect the constraining power of the analysis. We then estimate the amplitude of the quadrupolar modulation using the relation $g_2 = \left(1/5 \sum_M |g_{2M}|^2\right)^{1/2}$. Due to the nature of the estimator, which is necessarily positive, the estimation is biased. For an unbiased assessment, we estimate the mean and standard deviation of g_2 from simulations. We find no evidence for quadrupolar modulation of the primordial power spectrum. However, the derived limits allow us to impose tight constraints on statistically anisotropic inflationary models, such as those including vector fields during inflation. A companion paper, Planck Collaboration XX (2016), contains a more complete discussion on the theoretical implications of this constraint.

5.4. Point-parity asymmetry

The CMB anisotropy field defined on the sky, $T(\hat{\boldsymbol{n}})$, may be divided into symmetric, $T^+(\hat{\boldsymbol{n}})$, and antisymmetric, $T^-(\hat{\boldsymbol{n}})$, functions with respect to the centre of the sphere, as previously described in PCIS13. These functions have even and odd parity,

and thus correspond to spherical harmonics with even and odd ℓ-modes, respectively. On the very large scales corresponding to the Sachs-Wolfe plateau of the temperature power spectrum ($2 \leq \ell \leq 30$), the Universe should be parity neutral with no particular parity preference exhibited by the CMB fluctuations. However, an odd point-parity preference has previously been observed in the WMAP data releases (Land & Magueijo 2005a,b; Kim & Naselsky 2010a,b; Gruppuso et al. 2011) and the *Planck* 2013 results. Here, we investigate the parity asymmetry in the 2015 temperature maps at $N_{\text{side}} = 32$. We consider the following estimator:

$$R^{\text{TT}}(\ell_{\text{max}}) = \frac{D_+^{\text{TT}}(\ell_{\text{max}})}{D_-^{\text{TT}}(\ell_{\text{max}})}, \tag{36}$$

where $D_+(\ell_{\text{max}})$ and $D_-(\ell_{\text{max}})$ are given by

$$D_{+,-}^{\text{TT}} = \frac{1}{\ell_{\text{tot}}^{+,-}} \sum_{\ell=2,\ell_{\text{max}}}^{+,-} \frac{\ell(\ell+1)}{2\pi} C_\ell^{\text{TT}}, \tag{37}$$
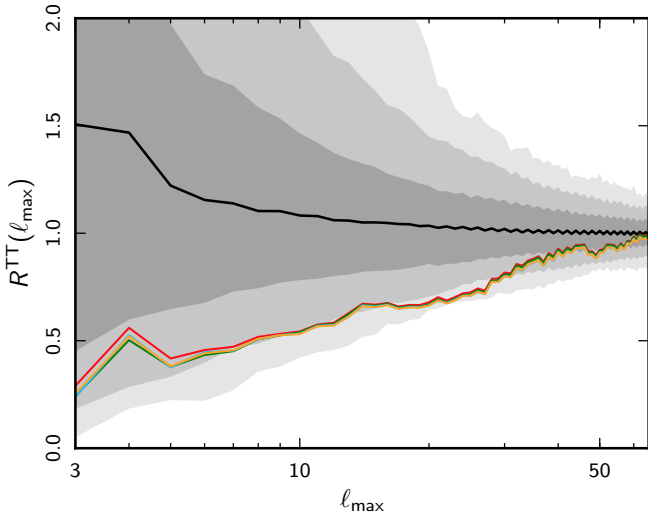
Fig. 19. Ratio $R^{TT}(\ell_{max})$ for `Commander` (red), `NILC` (orange), `SEVEM` (green), and `SMICA` (blue) determined at $N_{side} = 32$. The shaded grey regions indicate the distribution of the statistic derived from the `SMICA` MC simulations, with the dark, lighter, and light grey bands corresponding to the 1, 2, and 3σ confidence levels.

Fig. 20. Lower-tail probability of the point-parity estimator for `Commander` (red), `NILC` (orange), `SEVEM`, (green), and `SMICA` (blue).

$\ell_{tot}^{+,-}$ is the total number of even (+) or odd (−) multipoles included in the sum up to ℓ_{max}, and D_ℓ^{TT} is the temperature angular power spectrum computed using a quadratic maximum likelihood (QML) estimator (Gruppuso et al. 2011). The $\ell(\ell + 1)/(2\pi)$ factor in Eq. (37) effectively flattens the spectrum across the ℓ-range of the Sachs-Wolfe plateau (up to $\ell = 50$) in a ΛCDM model.

Figure 19 presents the ratio, $R^{TT}(\ell_{max})$, for the 2015 component-separated maps, together with the distribution determined from the `SMICA` MC simulations which serves as a reference for the expected behaviour of the statistic in a parity-neutral Universe. The distributions for the other CMB maps are very similar. The four component-separation products are in good agreement, indicating an odd-parity preference at very large scales for the multipole range considered in this test.

Figure 20 shows the lower-tail probability for the data as compared to simulations as a function of ℓ_{max}. The results are in good agreement with those in PCIS13. The cleaned CMB maps yield generally consistent profiles which signify an anomalous odd-parity preference in the multipole region $\ell_{max} = 20$–30. The minimum in the lower-tail probability occurs at $\ell = 28$ corresponding to a value of 0.2% for `NILC`, `SEVEM`, and `SMICA`, and 0.3% for `Commander`[4].

As a first attempt to quantify any a posteriori effects in the significance levels, we consider how many MC simulations appear in the lower tail of the MC distribution with a probability equal to, or lower than, 0.2%, for at least one ℓ_{max} value over a specific range. For ℓ_{max} in the range 3–50, the total number of simulated maps with this property is less than 20 over 1000 MC maps, implying that, even considering the LEE, an odd-parity preference is observed with a lower-tail probability of less than 2%.
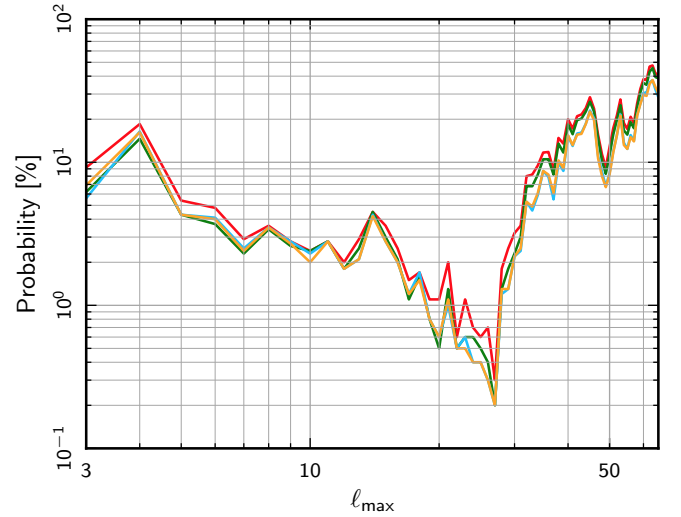
5.5. Mirror-parity asymmetry

For the *Planck* 2013 data release, we studied the properties of the temperature data at a resolution of $N_{side} = 16$ under reflection with respect to a plane to search for mirror symmetries. Such a symmetry might be connected to non-trivial topologies (Starobinsky 1993; Stevens et al. 1993; de Oliveira-Costa et al. 1996). In Planck Collaboration XXIII (2014), we reported evidence for an antisymmetry plane, with a perpendicular direction given by $(l, b) = (264°, -17°)$, However, the probability of the results was slightly dependent on the method of foreground cleaning, with a *p*-value ranging from 0.5% for `Commander-Ruler` to 8.9% for `SMICA`. The same direction was also found in the WMAP 7-year data (Finelli et al. 2012), and is close to that determined for the dipole modulation in the *Planck* 2013 data release (PCIS13), suggesting possible connections between the two directional anomalies.

We now proceed to reanalyse the status of mirror symmetries using the *Planck* 2015 full mission temperature data at both $N_{side} = 16$ and $N_{side} = 32$. In order to avoid possible bias introduced by the use of the Galactic mask[5] the results are derived from the full-sky `Commander`, `NILC`, and `SMICA` maps described in Sect. 2. For `SEVEM`, a customized map is first produced by inpainting about 3% of the map along the Galactic plane using a diffusive inpainting technique. This is then smoothed to the appropriate lower resolutions for further analysis. Following Finelli et al. (2012), we consider the estimators in the pixel domain given by:

$$S^\pm(\hat{n}_i) = \frac{1}{N_{pix}} \sum_{j=1}^{N_{pix}} \left[\frac{1}{2} \left(\frac{\delta T}{T}(\hat{n}_j) \pm \frac{\delta T}{T}(\hat{n}_k) \right) \right]^2, \quad (38)$$

where the sum is over all N_{pix} `HEALPix` pixels, $(\delta T/T)(\hat{n}_j)$ is the CMB temperature anisotropy measured at the pixel defined by the unit vector $\hat{n}_j$, and $\hat{n}_k$ is the opposite direction with respect to the plane defined by $\hat{n}_i$, i.e.,

$$\hat{n}_k = \hat{n}_j - 2(\hat{n}_i \cdot \hat{n}_j)\hat{n}_i. \quad (39)$$

---

[4] In the case where we would like to test the probability of finding a Universe with either odd or even parity preference, the probability would be higher by a factor of about two.

[5] The Galactic mask induces a preferred direction in the analysis of the MC simulation ensemble, which affects the significance of the results determined from the data. See Ben-David & Kovetz (2014) for a discussion.

Table 18. Lower-tail probability for the $S^{\pm}$ statistics of the component-separated maps at $N_{side} = 16$ and $N_{side} = 32$.

Estimator	Probability [%]	Direction (l, b) [°]
$N_{side} = 16$		
Commander		
min(S^+)	2.9	(264.4, −17.0)
min(S^-)	12.0	(260.4, 48.1)
NILC		
min(S^+)	2.3	(264.4, −17.0)
min(S^-)	16.8	(260.4, 48.1)
SEVEM		
min(S^+)	1.6	(264.4, −17.0)
min(S^-)	13.5	(260.4, 48.1)
SMICA		
min(S^+)	2.7	(264.4, −17.0)
min(S^-)	19.1	(260.4, 48.1)
$N_{side} = 32$		
Commander		
min(S^+)	1.9	(264.4, −15.7)
min(S^-)	10.0	(265.3, 46.2)
NILC		
min(S^+)	1.2	(264.4, −15.7)
min(S^-)	10.3	(265.3, 46.2)
SEVEM		
min(S^+)	0.8	(264.4, −15.7)
min(S^-)	11.1	(265.3, 46.2)
SMICA		
min(S^+)	1.7	(264.4, −15.7)
min(S^-)	11.6	(265.3, 46.2)

Note that we expect S^+ to be small if the points on opposite sides of the mirror are negatives of each other, and S^- to be small when they are the same.

We compute these quantities for each of the 3072 (12288) directions defined at resolution $N_{side} = 16$ (32), and allow the j and k indices to run over all of the pixels of the low-resolution full-sky maps. We perform the same analysis on 1000 FFP8 simulations and store the minimum value of $S^{\pm}$ for each of these to compute probabilities. The results are summarized in Table 18 and Fig. 21.

We confirm that the full mission *Planck* temperature data at $N_{side} = 16$ exhibits the most anomalous mirror antisymmetry in the direction $(l, b) = (264°, −17°)$, consistent with the result from the 2013 nominal mission data, with a probability which ranges from 1.6% for SEVEM to 2.9% for Commander. This is within 40° of the preferred direction identified by the dipole modulation analysis in Sect. 6.2. The corresponding results at $N_{side} = 32$ yield approximately the same direction, $(l, b) = (264°, −16°)$, with a slightly increased probability, ranging from 0.8% for SEVEM to 1.9% for Commander.

We also note that the CMB pattern exhibits a mirror symmetry in the direction $(l, b) = (260°, 48°)$, consistent with that found in the WMAP 7-year data (Finelli et al. 2012), and close to that identified by the solar dipole (Planck Collaboration VIII 2016). However, the significance of the symmetry pattern is less than in the antisymmetric case.

This extension of the analysis to higher resolution than in our previous work shows that the antisymmetry property does not seem to be confined to the largest angular scales, although we have not attempted to correct for any a posteriori choices made in the analysis. The detailed connection of this antisymmetry property to the low-variance and hemispherical asymmetry observed on these scales remains an open issue.
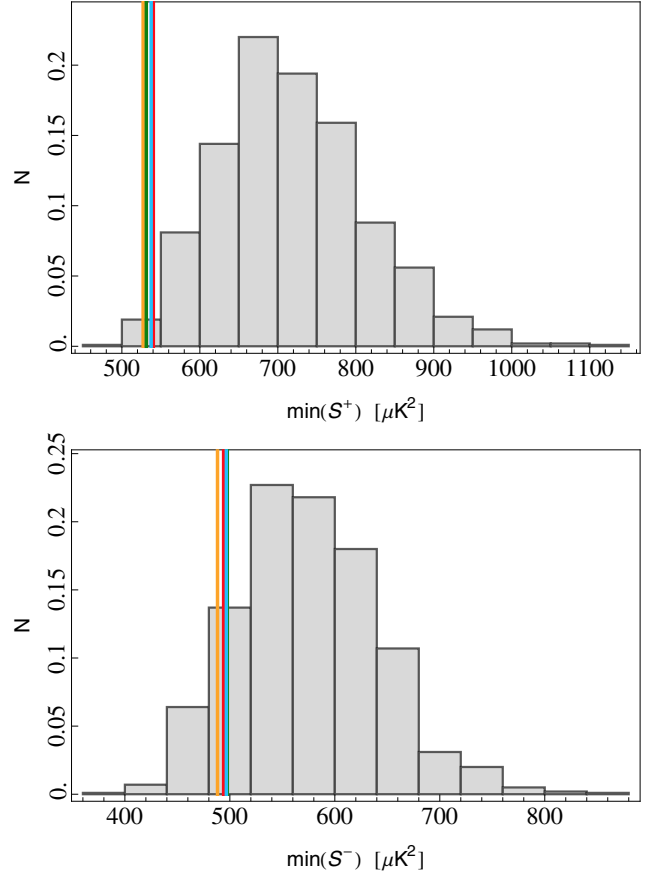
Fig. 21. Histograms of the S^+ (*top panel*) and S^- (*bottom panel*) statistic. The vertical lines show the minimum value for the estimator computed at $N_{side} = 32$ for Commander (red), NILC (orange), SEVEM (green), and SMICA (blue) maps. The grey area shows the same quantity computed from 1000 simulated SMICA maps.

5.6. Local peak statistics

Local extrema or peaks, as introduced in Sect. 4.5.3, can be employed to search for localized anomalies on the CMB sky by examining how their statistical properties vary in patches as a function of location.

Initially, we consider a further test for asymmetry by examining the differences in the peak distribution when divided according to orientation with respect to a previously specified asymmetry direction. In particular, we select the peaks both in a disc of radius 70° centred on $(l, b) = (225°, −18°)$ (the positive direction of the dipole defined in Sect. 6.2 for SMICA) and in the corresponding antipodal disc, then construct the empirical peak height CDFs to be compared with the full-sky median FFP8 distribution, as shown in Fig. 22. For maps filtered with a 40′ FWHM GAUSS filter the distribution of the peaks for the positive-direction disc is in general agreement with the full sky result, while that for the negative-direction is marginally different. Moreover, this pattern of behaviour is seen over a number of filtering scales, both for the KS deviation from the median full-sky simulated CDFs, and the spread of extremal values when comparing positive and negative regions. We also find that the properties of the negative disc affect the *p*-value results for a full sky KS test on data filtered with an SSG84 filter of 500′ FWHM, as seen in Sect. 4.5.3.

We can then extend the analysis for the 40′ GAUSS-filtered data by considering the variation in the peak statistical properties
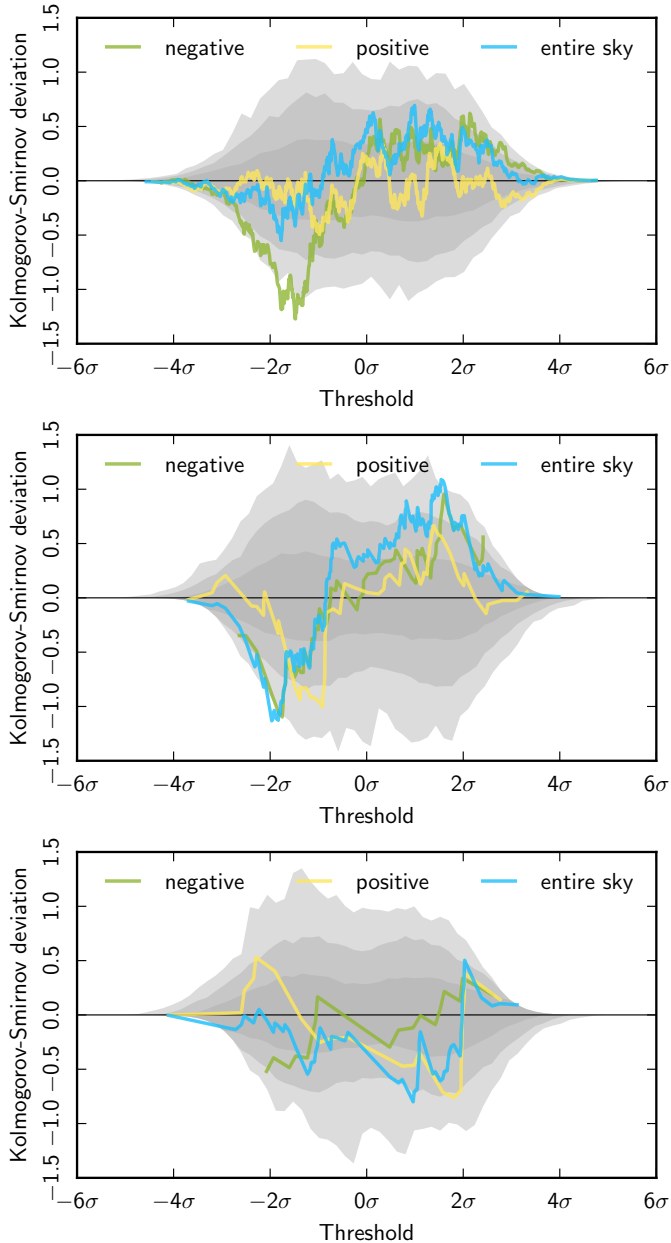
Fig. 22. KS-deviation of the peak distribution for 70° radius discs centred on the positive and negative asymmetry directions determined from the SMICA CMB temperature map in Sect. 6.2. From *top to bottom*, the plots correspond to maps filtered with a GAUSS kernel of 40′ FWHM, an SSG84 filter of 500′ FWHM, and an SSG84 filter of 800′ FWHM, respectively.

for a set of discs, each of which is centred on a pixel defined at $N_{side} = 256$. The simplest statistics to consider are the peak number counts. We therefore consider discs of 30° diameter and compute the peak counts for each disc. These are then compared to the corresponding peak count CDFs determined from simulations, and the upper- and lower-tail probabilities are assigned by counting the number of simulations above and below the observed counts at the same location. These quantities can then be visualized in the form of $N_{side} = 256$ sky maps. The derived $-\log_{10}(UTP)$ maps for each component-separation method are shown in Fig. 23. While we find that the total counts of peaks for the sky coverage defined by the common mask is consistent with simulations, significant regional variation is seen. Indeed,

the p-value for certain disc locations drops to 0.1% (i.e., the sky counts exceed anything seen in simulations). However, one needs to account for the a posteriori selection of significant regions in the determination of the true significance. It should also be noted that regional variations of the UTP are seen at similar levels when inspecting the peak-count statistics maps derived for randomly selected realizations of the simulations. Moreover, the significance of such peak-counting anomalies is degraded with larger disc diameters, and becomes insignificant for counts on the full sky. Thus, no significant anomalies can be claimed for the peak-count statistics of the *Planck* data.

A powerful non-parametric test of statistical isotropy is provided by the two-sample KS-deviation between the full sky empirical peak height CDF $F_n(x)$ (see Eq. (28)) and an empirical peak height CDF $F_{n'}(x)$ derived from a subsample of the distribution, again defined by the peaks within discs of 30° diameter as defined above. The two-sample KS-deviation

$$K_{nn'} \equiv \sqrt{\frac{nn'}{n+n'}} \sup_x |F_{n'}(x) - F_n(x)| \qquad (40)$$

for a partial sky region shares samples between the two CDFs, and can be calculated extremely efficiently using rank statistics according to

$$K_{nn'} \equiv \sqrt{\frac{nn'}{n+n'}} \max_i \left| \frac{r'(i)-1}{n'-1} - \frac{r(i)-1}{n-1} \right|, \qquad (41)$$

where r and r' denote the ranks of a value with index i in the full set of n and restricted set of n' samples, respectively. Maps of the upper tail probability are then determined by comparison with the equivalent quantities computed from simulations; $-\log_{10}(UTP)$ maps are shown in Fig. 24. The majority of the selected locations are consistent with the full-sky distribution, thus indicating the statistical isotropy of the *Planck* maps. The most prominent feature in each of the local KS-deviation maps appears south of the Galactic centre and may be associated with a cold region crossing the Galactic plane. However, as with the peak counts, it cannot be interpreted as statistically anomalous.

5.7. The Cold Spot

Since its discovery in the WMAP first-year data (Vielva et al. 2004), the Cold Spot, centred at Galactic coordinates $(l, b) = (210°, -57°)$ has been one of the most extensively studied large-scale CMB anomalies. In the 2013 release (Planck Collaboration XXIII 2014), *Planck* confirmed the apparently anomalous nature of this feature in temperature, in terms of the area of the SMHW coefficients on angular scales of $\approx 10°$ on the sky; the 2015 release has also confirmed this feature (see Sects. 4.5.2 and 4.5.3). The CMB temperature anisotropies around the Cold Spot as observed by *Planck* are shown in the top panel of Fig. 25. The peak merger tree within the Cold Spot region is presented in the lower panel of the figure and provides a multiscale view of its structure (see Sect. 4.5.4 for details).

The robustness of the detection of the anomalies discussed in this paper is a non-trivial issue. For the particular case of the Cold Spot, this has been reviewed by Vielva (2010), and addressed in detail by Cruz et al. (2006), paying specific attention to the impact of a posteriori choices. In particular, the latter study focused on the original test that indicated the presence of this feature on the sky, confirming a significance between 1% and 2%. An alternative analysis of the significance based on two statistical tests with different levels of conservativeness was made
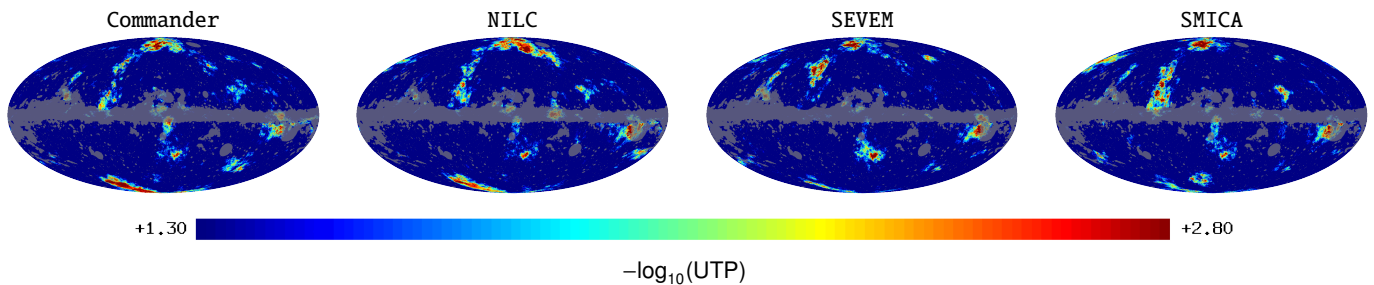
Fig. 23. Map of $-\log_{10}(UTP)$ for peak counts in the *Planck* 40′ GAUSS-filtered temperature data, where each pixel encodes the probability determined for a 30° diameter disc centred on it.
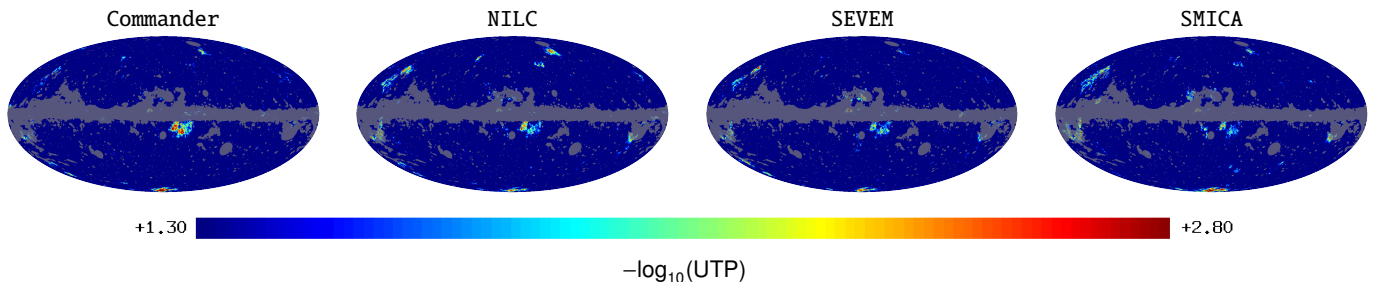

Fig. 24. Map of $-\log_{10}(UTP)$ for the two-sample KS-deviation where each pixel encodes the probability determined for a 30° diameter disc centred on it, as computed from the *Planck* 40′ GAUSS-filtered temperature data.

by McEwen et al. (2005), providing values of 0.1% and 4.7%, respectively. The statistical significance of the Cold Spot was questioned by Zhang & Huterer (2010) who found a low significance after performing a study based on different kernels. As discussed in more detail by Vielva (2010), this result can also be interpreted as evidence that not all kernels are necessarily suitable for the detection of arbitrary non-Gaussian features.

The possibility that the Cold Spot arises from instrumental systematics (Vielva et al. 2004) or foreground residuals (Liu & Zhang 2005; Cruz et al. 2006) has been largely rejected. However, several non-standard physical mechanisms have been proposed as possible explanations. These include the gravitational effect produced by a collapsing cosmic texture (Cruz et al. 2007), the linear and nonlinear ISW effect caused by a void in the large-scale structure (e.g., Tomita 2005; Inoue & Silk 2006; Rudnick et al. 2007; Tomita & Inoue 2008; Finelli et al. 2016), a cosmic bubble collision within the eternal inflation framework (Czech et al. 2010; Feeney et al. 2011; McEwen et al. 2012), and a localized version of the inhomogeneous reheating scenario within the inflationary paradigm (Bueno Sanchez 2014).

Since the other scenarios lack additional evidence, the void hypothesis would seem to be the most plausible, depending on the sizes, density contrasts, and profiles assumed in the computations, some of which are not in agreement with either observation (Cruz et al. 2008) or current *N*-body studies (Cai et al. 2010; Watson et al. 2014). However, Szapudi et al. (2015) have recently detected a large void in the WISE-2MASS galaxy catalogue aligned with the Cold Spot, with an estimated radius of around 200 h^{-1} Mpc, an averaged density contrast of $\bar{\delta} \approx -0.1$, and centred on a redshift of $z \approx 0.15$. Large voids with similar characteristics are not unusual in the standard ΛCDM model (Nadathur et al. 2014). In fact, *N*-body simulations predict about 20 such voids in the local Universe ($z < 0.5$). However, Zibin (2014) and Nadathur et al. (2014) indicate that the expected signal due to the linear and nonlinear ISW effects caused by this

structure is not large enough to explain the temperature decrement associated with the Cold Spot.

The new *Planck* data release allows us to further explore the statistical nature of the Cold Spot. Two previous studies (Zhao 2013; Gurzadyan et al. 2014) have claimed inconsistencies of the internal properties of the Cold Spot with the Gaussian hypothesis, which we re-address here. In particular, we consider the small-scale fluctuations within a disc-like region of radius ≈25°.

Several statistical quantities are computed from the full-resolution temperature maps within the Cold Spot region. This is divided into a central disc of diameter 1° surrounded by a set of 13 concentric annuli with central radii spaced in steps of about 2°, thus allowing us to build angular profiles for the mean, variance, skewness, and kurtosis. These are then compared to specialized CMB realizations, generated as follows. A set of Gaussian CMB skies is simulated using the FFP8 reference spectrum, and convolved with a Gaussian beam of 5′ FWHM. As for the FFP8 simulations themselves, these maps are rescaled, as discussed previously. Only those that contain a spot as extreme as the Cold Spot at a scale $R = 300′$ in SMHW space are retained, and these are rotated such that each simulated cold spot is relocated to the actual position of the Cold Spot (this ensures that the noise properties are identical for both data and simulations). This selection criterion corresponds to the characteristic that originally indicated the presence of the Cold Spot in the observed sky. As a final step, for each remaining CMB simulation a noise realization is added, consistent with each component-separation method.

The results are presented in Fig. 26. Focusing on the profile of the mean value, it is apparent that the largest deviations from the simulations appear on scales around 15°, which corresponds to a hot ring structure, as seen in Fig. 25 and previously discussed in Cayón et al. (2005) and Nadathur et al. (2014). Notice that on the smallest scales the mean profile is also somewhat deviant with respect to the simulations, but this may be connected to selection bias, since we are considering CMB
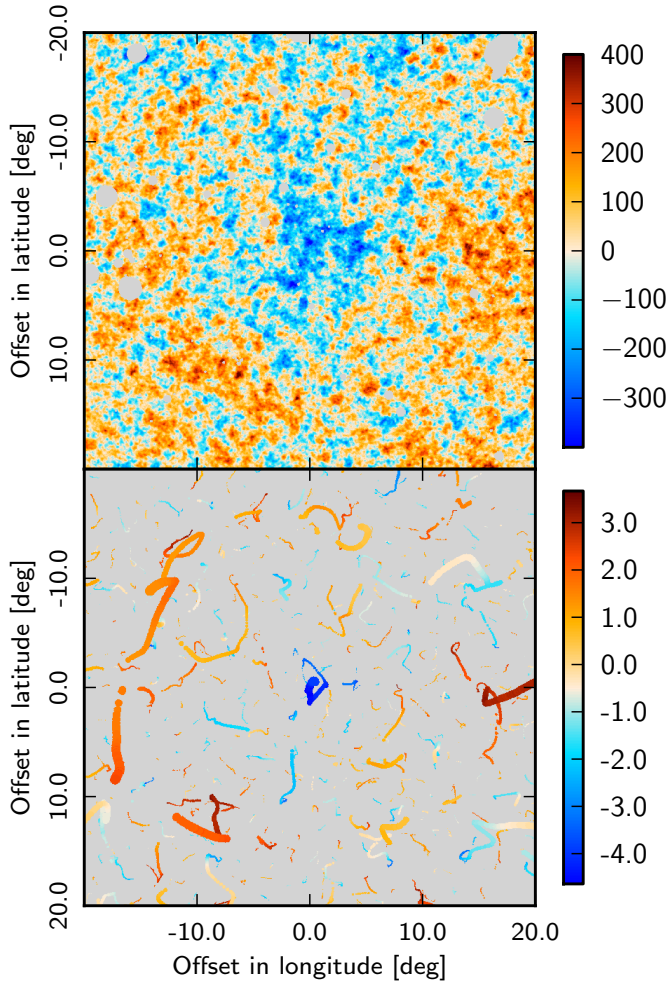
Fig. 25. *Top*: temperature patch centred on the Cold Spot. *Bottom*: peak merger tree within the Cold Spot region. The figure shows a region centred on the Cold Spot location in gnomonic projection, with all the peaks in SSG84-filtered maps with FWHM ranging from 80′ to 1200′ overlaid on the same plot. The size of the coloured circles is proportional to the filtering scale. The colour corresponds to the peak value, normalized in units of σ for a given filter scale. In both panels the data are from the SMICA CMB map at full resolution.

simulations containing a spot that is at least as cold as the Cold Spot. However, if we consider the distribution of the profiles corresponding to the coldest spots instead of the spots as extreme as the Cold Spot (removing the bias at the smallest scales) then the results do not change substantially (see below).

In order to quantify possible deviations from Gaussianity, we determine the probability of finding a χ^2 value larger than that of the data for each statistic, as summarized in Table 19. The χ^2 value for the data is computed using an estimate of the covariance matrix between different radial scales determined from the Cold Spot simulations (1000 for each component-separation method), and then compared to the theoretical χ^2 distribution with 13 degrees of freedom. The results indicate that the angular profile for the mean is poorly described by the simulations, of which less than 1% are found to have a higher χ^2 than the data (when considering the distribution corresponding to the coldest spot this probability becomes approximately 2%). We have checked that this deviation is not obviously associated with a particular sub-range of angular scales, implying that the mean profile is anomalous over the full range considered. Conversely,

Table 19. Probabilities of obtaining values for the χ^2 statistic of the angular profiles of the estimators shown in Fig. 26 larger than those determined from the data.

Angular profiles	Probability [%]			
	Comm.	NILC	SEVEM	SMICA
Mean	0.9	0.8	1.0	0.9
Variance	40.0	40.0	38.0	42.0
Skewness	79.0	82.0	85.0	80.0
Kurtosis	75.0	56.0	75.0	77.0

the radial profiles of the higher-order moments are compatible with the Gaussian simulations. The latter results are then in contradiction with a similar analysis (using discs instead of rings) by Zhao (2013) for the WMAP 9-year data. However, it appears that this may be related to the criteria applied for the selection of the Gaussian simulations used to define the null hypothesis. In particular, Zhao (2013) used the coldest pixel in real space as a means to identify those simulations that should be retained, as opposed to the existence of cold spots as extreme as the Cold Spot selected in the SMHW coefficient map at $R = 300'$. Since it is not implicit that such a temperature extremum is necessarily associated with an extended cold region, particularly one defined in wavelet space, the simulations used by Zhao (2013) did not contain features comparable to the nature of the Cold Spot. This explains why the Cold Spot seemed to be anomalous when looking at the small-scale fluctuations.

In conclusion, it appears that only the mean temperature profile of the Cold Spot should be considered anomalous when compared to CMB cold spots that are as statistically extreme. All other measures of its internal structure are consistent with expectations.

As a final remark, we note that the high-pass filtering currently applied to the *Planck* CMB polarization maps severely limits the possibility of conducting targeted analyses to discriminate between different possible origins of the Cold Spot. For example, no polarization signal would be expected in those models producing secondary anisotropies due to a gravitational effect, whereas a specific pattern might be expected in a bubble collision scenario (Czech et al. 2010). Appropriate tests will be pursued in future work, once the large-scale polarization data are available.

6. Dipole modulation and directionality

In this section, we examine isotropy violation related to dipolar asymmetry, various forms of which have been noted since the early WMAP releases (Eriksen et al. 2004a). We perform a non-exhaustive series of tests in an attempt to narrow down the nature of the asymmetry (on the assumption that it is not simply a statistical fluke). First, we will briefly describe some similarities and differences between the tests that are important for making a proper comparison of the results.

All the tests in this section share in common the fitting of a dipole. This is done either by fitting for a dipole explicitly in a map of power on the sky (Sects. 6.1 and 6.5), by employing Bayesian techniques in pixel space for a specific model (Sect. 6.2), or by measuring the coupling of ℓ to $\ell \pm 1$ modes in the CMB covariance matrix (Sects. 6.3, 6.4, and 6.6). The differences arise from how the fitted dipoles are combined, which determines the specific form of asymmetry that the test is sensitive to.
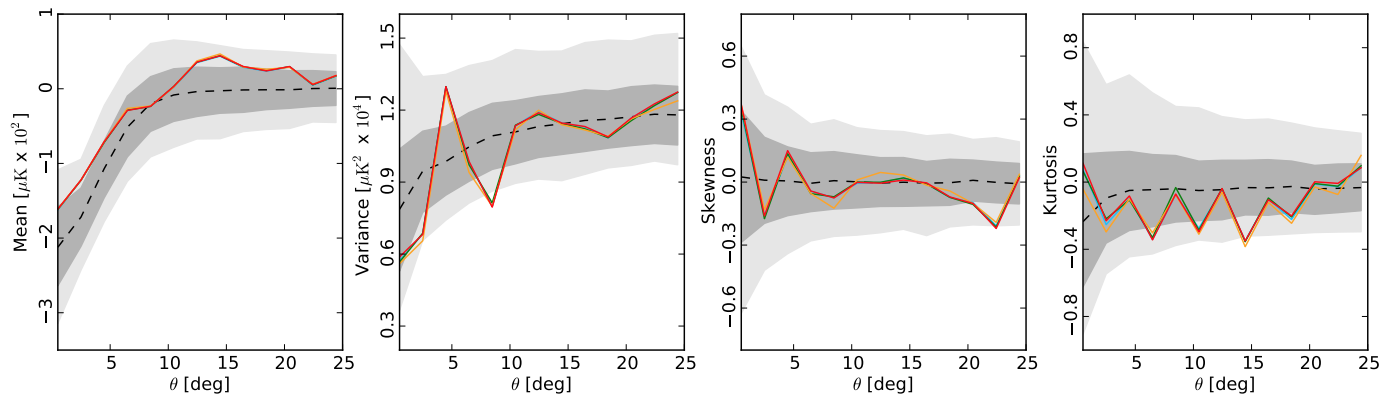
Fig. 26. From *left to right*: mean, variance, skewness, and kurtosis angular profiles computed for rings at radii θ centred on the Cold Spot position for `Commander` (red), `NILC` (orange), `SEVEM` (green), and `SMICA` (blue). The expected value obtained from the simulations is denoted by the black dashed line and the dark and light grey regions represent the 1σ and 2σ intervals.

The tests can be divided into two categories, amplitude-based and direction-based. Sections 6.1 to 6.4 are all sensitive to the *amplitude* of a dipole modulation. Specifically, Sect. 6.1 looks for dipole modulation in the pixel-to-pixel variance of the data, while Sects. 6.2–6.4 all search for dipole modulation of the angular power spectrum. The distinction between these two approaches is mainly one of ℓ weighting.

Sections 6.5 and 6.6 both examine aspects of *directionality* in the data, where the directions are extracted from dipole fits but combined in different ways. Section 6.5 fits for dipoles in band power (with similar results for variance) and only uses the direction information, while Sect. 6.6 weights each dipole equally across all scales and uses the amplitude information as well.

The differences between the approaches of these sections should be kept in mind when comparing their results. For example, although Sects. 6.5 and 6.6 both look for a directional signal in the data, they are optimized for different forms of deviations from statistical isotropy. It is therefore unsurprising that they arrive at different results. However, the signal found in Sect. 6.5, if not simply a statistical fluke, is constrained by the results of Sect. 6.6.

Regarding the impact on the dipolar modulation results of the lack of the aberration contribution to the simulations, we note the following. In general the analyses are either sensitive only to large angular scales, or only claim possible detections on such scales, where the effect of aberration will be negligible and hence the conclusions are unlikely to change. A possible exception is in relation to the results of Sect. 6.5, where claims are made about effects extending out to $\ell_{max} = 1500$. It is plausible that the effects of aberration could start to become important on these scales.

6.1. Variance asymmetry

The study of power asymmetry via the local variance of the CMB fluctuations was first performed by Akrami et al. (2014) for the *Planck* 2013 and WMAP 9-year temperature data. The approach was motivated by its conceptual and implementational simplicity, its directly intuitive interpretation, and by virtue of being defined in pixel space, a useful complementary to other mostly harmonic-based methods. The statistic was computed over patches of different sizes and positions on the sky, and compared with the values obtained from statistically isotropic simulations. It was found that none of the 1000 available simulations had a larger variance asymmetry than that estimated from the data. This suggested the presence of asymmetry at a

statistical significance of at least 3.3σ, with a preferred direction $(l, b) \approx (212°, -13°)$ in good agreement with other studies. In this section, we revisit the variance asymmetry and report the results of the analysis for the *Planck* 2015 temperature maps at full resolution, $N_{side} = 2048$.

The analysis proceeds as follows. We consider a set of discs of various sizes centred on the pixels of a `HEALPix` map defined by a specific N_{side} value. For each sky map, we first remove the monopole and dipole components from the masked sky and then compute the variance of the fluctuations on a given disc using only the unmasked pixels. This yields a local-variance map at the `HEALPix` resolution of interest. We also estimate the expected average and variance of the variance on each disc from the simulations and then subtract the resulting average variance map from both the observed and simulated local-variance maps. Finally, we define the amplitude and direction of the asymmetry by fitting a dipole to each of the local-variance maps, where each pixel is weighted by the inverse of the variance of the variances computed from the simulations at that pixel. At all stages, we use only the discs for which more than 10% of the area is unmasked, although our results are robust against the choice of this value. The computed local-variance amplitudes are then used to compare the data with statistically isotropic simulations. Note that we use only the dipole amplitudes of the local-variance maps to measure the significance of the asymmetry; the amplitudes of higher multipoles were shown by Akrami et al. (2014) to be consistent with statistically isotropic simulations and we therefore do not consider them in the present paper.

In Akrami et al. (2014), the sensitivity of the method to the disc size was assessed using both statistically isotropic and anisotropic simulations. The free parameters, i.e., the number and size of the discs, were then fixed by these simulations. It was found that for 3072 patches centred on the set of pixels defined at $N_{side} = 16$, the simulated asymmetry signals were not detected when either very small ($r_{disc} < 4°$) or very large ($r_{disc} > 16°$) discs were used.

The former effect is due to a combination of the low number of pixels per disc and an insufficient number of discs to cover the entire sky when $N_{side} = 16$ reference grids are used. However, it has recently been shown by Adhikari (2015) that using a larger number of small discs (by increasing N_{side} to 32, 64, 128, and 256, depending on the disc size) in order to cover the entire sky allows the local-variance method to detect the large-scale anomalous asymmetry as well as the Doppler boost signal from the *Planck* 2013 data, at a significance of $>3.3\sigma$. Fantaye (2014) has demonstrated that the Doppler boost signal can be
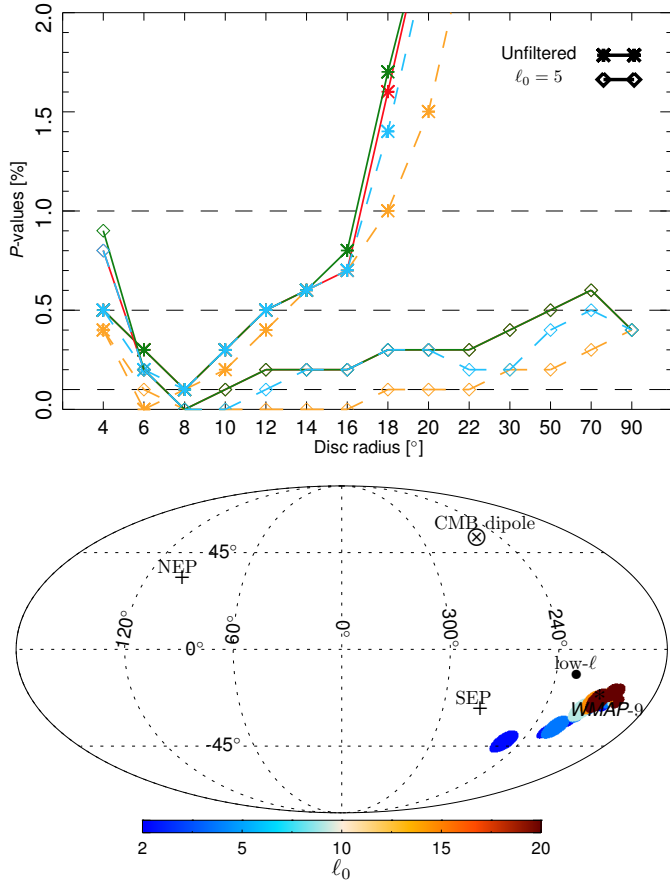
Fig. 27. *Upper panel*: *p*-values for variance asymmetry measured as the number of simulations with local-variance dipole amplitudes larger than those inferred from the data, as a function of disc radius for the four component-separated maps, `Commander` (red), `NILC` (orange), `SEVEM` (green), and `SMICA` (blue), and for unfiltered and high-pass-filtered cases. For the filtered case, the `Commander` curve is covered by the `SMICA` curve for small ($r_{disk} \leq 8$) disks, and by the `SEVEM` curve for large disks ($r_{disk} > 8$). *Lower panel*: local-variance dipole directions for the `SMICA` map. The colours, as indicated by the colourbar, correspond to different values of the high-pass filter central multipole ℓ_0. The size of a marker disc corresponds, from small to large, to the size of the disc used in the analysis, namely 4°, 12°, 20°, and 70°. The dipole directions from the `Commander`, `NILC`, and `SEVEM` component-separation methods are consistent with the case shown here. The low-ℓ and WMAP-9 directions are identical to those in Fig. 35.

detected at a similar level of significance using needlet bandpass filtering of the data, even with large discs, when simulations are deboosted. Here, in contrast to the 2013 analysis, we use maps which contain Doppler boosting, for both simulations and data, and therefore we do not detect any Doppler boost signal when using a large number of small discs.

The low observed significance levels when large discs are used is due to the cosmic variance associated with the largest-scale modes. Motivated by the analysis of Fantaye (2014), and in order to address this issue, we also perform analyses using a Butterworth high-pass filter,

$$H(\ell) = \frac{(\ell/\ell_0)^4}{1 + (\ell/\ell_0)^4},\qquad(42)$$

centred at multipoles $\ell_0 = 5, 10, 15, 20,$ and 30. In addition, the filtering of low multipoles allows us to establish the contribution of such modes to any detected asymmetry.

Table 20. *p*-values for the variance asymmetry measured by 8° discs for the four component-separated temperature maps and different high-pass filter scales.

	p-value [%]			
ℓ_0	Comm.	NILC	SEVEM	SMICA
Unfiltered ...	0.1	0.1	0.1	0.1
5	<0.1	<0.1	<0.1	<0.1
10	<0.1	<0.1	<0.1	<0.1
15	0.1	<0.1	0.1	<0.1
20	0.4	<0.1	0.3	0.2
30	1.8	0.8	1.8	1.7

Notes. The values represent the fraction of simulations with local-variance dipole amplitudes larger than those inferred from the data.

Here, based on the analysis of Akrami et al. (2014), we restrict our analysis to those disc sizes for which 3072 discs, corresponding to an $N_{side} = 16$ map, cover the entire sky, i.e., to the range 4°–90°. Consistent results can be obtained by choosing other values of N_{side} for a given disc size provided that the entire sky is covered by the discs. Here, for simplicity, we work with the same N_{side} (=16) for all disc sizes.

Our results for the measured amplitude of the variance asymmetry, compared to the values from the simulations, as well as the corresponding dipole directions, are shown in Fig. 27. The *p*-values are given for different disc sizes and in terms of the number of simulations with local-variance dipole amplitudes greater than the ones measured from the data. Note that since the discs with different sizes used in our analysis are correlated, the significance levels are also correlated. For this reason we choose to show the *p*-values as a function of disc size instead of combining them into a single number. Moreover, it should be noted that the significance values we present here do not incorporate any corrections to account for the choice of parameters adopted during method calibration, specifically the dipole amplitudes and directions for the anisotropic simulations that were used to fix the range of disc sizes and number of patches.

It can be seen from the upper panel of Fig. 27 that for the unfiltered map the significance of the power asymmetry drops quickly when we increase the disc size to radii greater than 16°. This is no longer the case, however, when the lowest multipoles are filtered out. For example, when the filter scale is set to $\ell_0 = 5$, i.e., when the very low multipoles which are affected most by cosmic variance are suppressed, the variance asymmetry is detected at the 3σ level for all disc sizes, as shown in Fig. 27. Table 20 presents the *p*-values of the variance asymmetry using 8° discs and for various values of ℓ_0. Our results show that variance asymmetry is detected with a remarkable significance for all disc sizes when very low multipoles are filtered out. In addition, the variance asymmetry amplitude slowly decreases with increasing ℓ_0, as seen in the upper panel of Fig. 28. For $\ell_0 \gtrsim 20$, the dipole amplitude becomes too small and we find no significant variance asymmetry. It is interesting to note, however, that the dipole directions found for large ℓ_0 are closely aligned with those found for $\ell_0 < 20$.

The lower panel of Fig. 27 shows the dipole directions we find using different disc sizes and different filter scales for `SMICA`. The dipole directions for the `Commander`, `NILC`, and `SEVEM` component-separated maps are very similar to those shown. The asymmetry directions found here are consistent with those determined by other analyses in this paper.
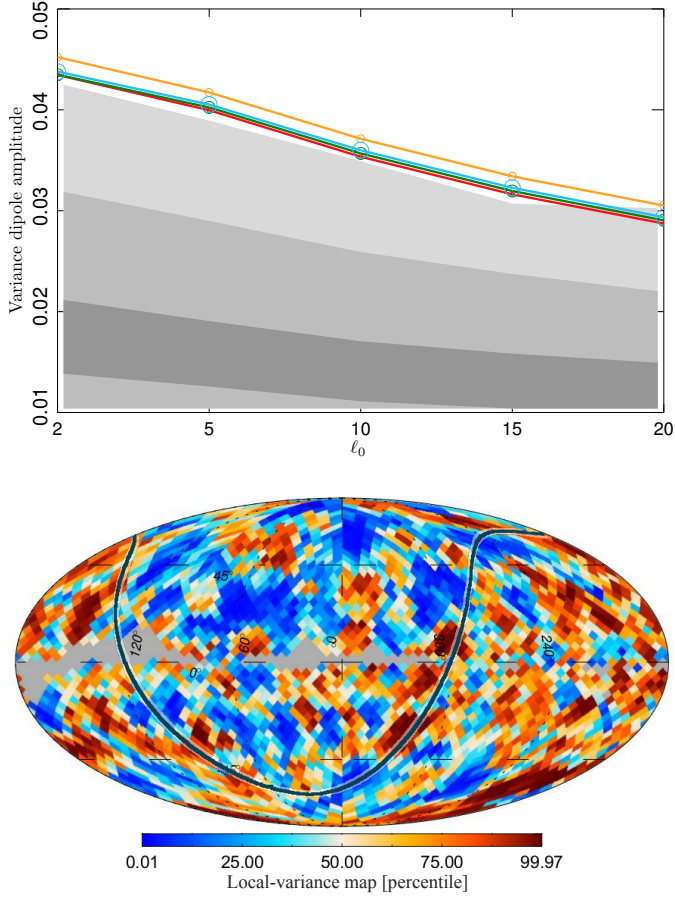
Table 21. Local-variance dipole amplitudes and directions.

ℓ_0	A^a	Direction (l, b) [°]
Unfiltered ..	0.052 ± 0.016	$(210, -26)$
5	0.046 ± 0.014	$(208, -24)$
10	0.040 ± 0.014	$(199, -16)$
15	0.038 ± 0.012	$(206, -16)$
20	0.028 ± 0.010	$(202, -18)$
30	0.025 ± 0.010	$(199, -19)$

Notes. All values quoted here are for 8° discs. This table is for the `Commander` component-separation method, but the results are similar for the other methods. $^{(a)}$ $A = 2(A_{\mathrm{Planck}} - \langle A_{\mathrm{FFP8}} \rangle)$, where A_{Planck} and A_{FFP8} are the local-variance dipole amplitudes of the data and the FFP8 simulations, respectively. The quoted errors are the dispersion of the simulation amplitudes. Assuming a pure dipole modulation model, A to first order would correspond to the modulation amplitude.

6.2. Dipole modulation: pixel-based likelihood

In PCIS13 we presented an analysis of the apparent anisotropic distribution of large-scale power in the *Planck* 2013 temperature data within the parametric framework defined by Gordon (2007) and Hoftuft et al. (2009), who introduced an explicit dipole modulation field to model potential hemispherical power asymmetry. The following is a direct update of that analysis using the *Planck* 2015 CMB data at $N_{\mathrm{side}} = 32$, retaining the 2013 common mask to explicitly test for consistency with the earlier study. All results are found to be in excellent agreement. In the following, we therefore only consider a smoothing scale of 5° FWHM as a representative example. This is the highest angular resolution accessible for an $N_{\mathrm{side}} = 32$ map.

Recall first the basic data model adopted in the dipole modulation approach: rather than assuming the CMB sky to be a statistically isotropic Gaussian field, we allow for an additional dipole modulation, resulting in a data model of the form $\boldsymbol{d} = \mathsf{BM}\boldsymbol{s} + \boldsymbol{n}$, where $\mathsf{M}_{ij} = (1 + \alpha\,\hat{\boldsymbol{p}} \cdot \hat{\boldsymbol{n}}_i)\delta_{ij}$ is an offset dipole field multiplying an intrinsically isotropic signal $\boldsymbol{s}$ with a dipole of amplitude α pointing towards some preferred direction $\hat{\boldsymbol{p}}$. B denotes convolution with an instrumental beam, and $\boldsymbol{n}$ denotes instrumental noise. Additionally, we model the power spectrum of the underlying statistically isotropic field in terms of a two-parameter amplitude–tilt model of the form $C_\ell(q, n) = q\,(\ell/30)^n\,C_\ell^{\Lambda\mathrm{CDM}}$, where $C_\ell^{\Lambda\mathrm{CDM}}$ is the best-fit *Planck* 2015 ΛCDM spectrum (Planck Collaboration XI 2016). The two parameters q and n can accommodate a deficit in power at low ℓ as compared to the best-fit cosmology that would otherwise create a tension with the underlying statistically isotropic model and result in the analysis measuring a combination of both asymmetry and power mismatch.

In the absence of any dipole modulation, $\alpha = 0$, the total data covariance matrix is given by $\mathsf{C} = \mathsf{B}\mathsf{S}_{\mathrm{iso}}\mathsf{B}^\mathsf{T} + \mathsf{N}$, where $\mathsf{S}_{\mathrm{iso}}$ is the standard statistically isotropic CMB covariance matrix given by the power spectrum, C_ℓ, N is the noise covariance matrix, and the corresponding likelihood is given by the usual expression for a multivariate Gaussian distribution. With dipole modulation, this generalizes straightforwardly to $\mathsf{C} = \mathsf{BMS}_{\mathrm{iso}}\mathsf{M}^\mathsf{T}\mathsf{B}^\mathsf{T} + \mathsf{N}$, with the likelihood given by

$$\mathcal{L}(\alpha, \hat{p}, q, n) \propto \frac{\exp\left[-\frac{1}{2}\boldsymbol{d}^t(\mathsf{BMSM}^\mathsf{T}\mathsf{B}^\mathsf{T} + \mathsf{N})^{-1}\boldsymbol{d}\right]}{\sqrt{|\mathsf{BMSM}^\mathsf{T}\mathsf{B}^\mathsf{T} + \mathsf{N}|}}. \quad (43)$$

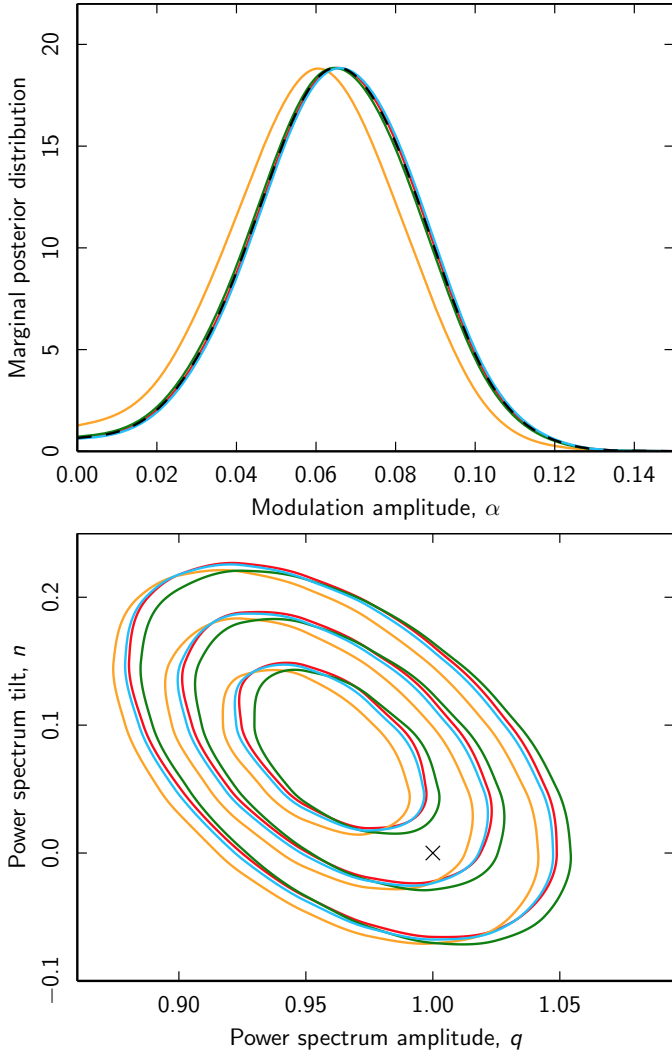Figure 29 and Table 22 summarize this five-dimensional likelihood in terms of marginal parameters for each of the four

Fig. 28. *Upper panel*: local-variance dipole amplitude for 8° discs as a function of the central multipole of the high-pass filter, ℓ_0, for the four component-separation methods, `Commander` (red), `NILC` (orange), `SEVEM` (green), and `SMICA` (blue). The grey regions, from dark to light, correspond, respectively, to 1σ, 2σ, and 3σ percentiles from the 1000 FFP8 simulations processed by the `Commander` method. *Lower panel*: mean-subtracted and inverse-variance-weighted local-variance map for the 8° discs and for the `Commander` component-separation method; each pixel is given in terms of the lower- and upper-tail probability of the measured value on that pixel compared to the values from the simulations. The pixels in grey correspond to the centres of the 8° discs on which the number of unmasked pixels in the full resolution map is lower than our threshold. The black curve superposed on the map indicates the boundary of the opposing hemispheres along the asymmetry axis. It is clear that the largest fraction of >95% outliers (red pixels) lie on the positive amplitude hemisphere of the local variance dipole, while the <5% outliers (blue pixels) are on the opposite hemisphere. The corresponding maps for `NILC`, `SEVEM`, and `SMICA` are very similar to the one shown here.

In the upper panel of Fig. 28, we show the local-variance dipole amplitudes for the 8° discs as a function of the central multipole of the high-pass filter, ℓ_0. In the lower panel of the same figure we show, as an example, the mean-subtracted and inverse-variance-weighted local-variance map using 8° discs for the `Commander` component-separation method. The pixels of the map are given in terms of the lower- and upper-tail probabilities of the values from the data compared to the values from the simulations. The maps for `NILC`, `SEVEM`, and `SMICA` are very similar. The numerical values of the local-variance dipole amplitudes and directions for the `Commander` method are given in Table 21; the values for the `NILC`, `SEVEM`, and `SMICA` methods are similar.

Fig. 29. *Top*: marginal constraints on the dipole modulation amplitude, as derived from *Planck* 2015 temperature observations at a smoothing scale of 5° FWHM for Commander (red), NILC (orange), SEVEM (green), and SMICA (blue). The plot corresponds directly to Fig. 32 of Planck Collaboration XXIII (2014). The Commander, SEVEM, and SMICA posteriors coincide almost perfectly both internally, and with the corresponding SMICA 2013 posterior, shown as a dashed black line. *Bottom*: corresponding marginal two-dimensional constraints on the low-ℓ power spectrum amplitude and tilt, (q, n), defined relative to the best-fit *Planck* 2015 ΛCDM model.

Planck CMB maps, as evaluated over the common mask using the multi-dimensional grid-based Snake algorithm (Mikkelsen et al. 2013). All results correspond to a smoothing scale of 5° FWHM, the highest resolution supported by an $N_{side} = 32$ HEALPix grid, but, as in 2013, we consider all smoothing scales between 5° and 10° FWHM, reaching similar conclusions in each case: the dipole modulation results derived from the *Planck* 2015 temperature maps are essentially identical to the 2013 results, with improved internal consistency between the four CMB maps due to better mitigation of systematic errors. The best-fit dipole modulation amplitude at 5° FWHM is 6–7% whilst the low-ℓ power spectrum has an approximately 3–5% lower amplitude compared to the best-fit ΛCDM prediction. These results are fully consistent with expectations given that the *Planck* 2013 sky maps were already cosmic-variance-limited on these angular scales, and the 2015 maps differ from the 2013 maps at the level of only a few microkelvin (Planck Collaboration IX 2016).

Table 22. Summary of dipole modulation results at a smoothing scale of 5° for all *Planck* 2015 CMB temperature solutions, as derived by the brute-force likelihood given by Eq. (43).

Method	2013	2015
Dipole modulation amplitude, α		
Commander	0.078 ± 0.021	0.066 ± 0.021
NILC	0.069 ± 0.021	0.061 ± 0.022
SEVEM	0.066 ± 0.021	0.065 ± 0.021
SMICA	0.065 ± 0.021	0.066 ± 0.021
Dipole modulation direction, (l, b) [°]		
Commander	$(227, -15) \pm 19$	$(230, -16) \pm 24$
NILC	$(226, -16) \pm 22$	$(228, -19) \pm 29$
SEVEM	$(227, -16) \pm 24$	$(226, -17) \pm 25$
SMICA	$(226, -17) \pm 24$	$(225, -18) \pm 24$
Power spectrum amplitude, q		
Commander	$\cdots$	0.961 ± 0.025
NILC	$\cdots$	0.954 ± 0.024
SEVEM	$\cdots$	0.966 ± 0.025
SMICA	$\cdots$	0.960 ± 0.025
Power spectrum tilt, n		
Commander	$\cdots$	0.082 ± 0.043
NILC	$\cdots$	0.077 ± 0.043
SEVEM	$\cdots$	0.077 ± 0.043
SMICA	$\cdots$	0.081 ± 0.043

6.3. Dipole modulation: QML analysis

In this section we use the QML estimator introduced in Moss et al. (2011) and described in Appendix C to assess the level of dipole modulation in our estimates of the CMB sky at $N_{side} = 2048$. The specific implementation is essentially identical to that used in Hanson & Lewis (2009), Planck Collaboration XVII (2014), and Planck Collaboration XXVII (2014), and exploits the fact that dipole modulation of any cosmological parameter is equivalent to coupling of ℓ to ℓ ± 1 modes in the CMB covariance matrix to leading order (see Appendix C). Planck Collaboration XX (2016) presents an alternate analysis for a specific isocurvature model.

Since we are interested in dipole modulation there are three independent estimators. For our particular approach, these are a real-valued $m = 0$ and a complex-valued $m = 1$ estimator, and take the form

$$\tilde{X}_0 = \frac{6}{f_{10}} \frac{\sum_{\ell m} \delta C_{\ell \ell+1} A_{\ell m} \left(T^*_{\ell m} T_{\ell+1\,m} - \left\langle T^*_{\ell m} T_{\ell+1\,m} \right\rangle \right)}{\sum_{\ell} \delta C^2_{\ell \ell+1} (\ell + 1) F_\ell F_{\ell+1}}, \quad (44)$$

$$\tilde{X}_1 = \frac{6}{f_{11}} \frac{\sum_{\ell m} \delta C_{\ell \ell+1} B_{\ell m} \left(T^*_{\ell m} T_{\ell+1\,m+1} - \left\langle T^*_{\ell m} T_{\ell+1\,m+1} \right\rangle \right)}{\sum_{\ell} \delta C^2_{\ell \ell+1} (\ell + 1) F_\ell F_{\ell+1}}. \quad (45)$$

Here $T_{\ell m}$ are C-inverse filtered data and $F_\ell \equiv \left\langle T_{\ell m} T^*_{\ell m} \right\rangle$. We adopt the inverse-variance filter from Planck Collaboration XVII (2014), where the approximate filter functions are also specified. We define $\delta C_{\ell \ell+1} \equiv dC_\ell/dX + dC_{\ell+1}/dX$, where X is the parameter modulated, and $A_{\ell m}$ and $B_{\ell m}$ are numerical coefficients (details can be found in Appendix C). The factor f_{1m} corrects the normalization for errors introduced by masking:

$$f_{1m} \equiv \int d\Omega \, Y^*_{1m}(\Omega) M(\Omega), \quad (46)$$

where $M(\Omega)$ is the mask. Finally, we correct the direction for the effects of inhomogeneous noise which is not accounted for in the

Table 23. Amplitude (*A*) and direction of the low-ℓ dipole modulation signal determined from the QML analysis for the range $\ell \in [2, 64]$.

Method	A	Direction (l, b) [°]
Commander	$0.063^{+0.025}_{-0.013}$	$(213, -26) \pm 28$
NILC	$0.064^{+0.027}_{-0.013}$	$(209, -25) \pm 28$
SEVEM	$0.063^{+0.026}_{-0.013}$	$(211, -25) \pm 28$
SMICA	$0.062^{+0.026}_{-0.013}$	$(213, -26) \pm 28$

Notes. The errors are calculated from the cosmic variance expected for statistically isotropic CMB realizations.

filtering process, by weighting the $\tilde{X}_m$ by the inverse of the variance derived from filtered and mean-field corrected simulations.

The physics is readily accessible in this estimator: the ℓ-dependence in modulation determined by the parameter X is expressed in the $\delta C_{\ell\ell+1}$ factor, and the relevant scales appear directly in the limits of the sum. We consider the estimator over the range $\ell_{\min} = 2 \le \ell \le \ell_{\max}$. The modulation amplitude and direction are then given by

$$\tilde{A} = \sqrt{\tilde{X}_0^2 + 2|\tilde{X}_1|^2}, \tag{47}$$

$$\tilde{\theta} = \cos^{-1}\left(\frac{\tilde{X}_0}{\tilde{A}}\right), \tag{48}$$

$$\tilde{\phi} = -\tan^{-1}\left(\frac{\mathrm{Im}[\tilde{X}_1]}{\mathrm{Re}[\tilde{X}_1]}\right). \tag{49}$$

It is worth re-emphasizing that the quantities $\tilde{A}$, $\tilde{\theta}$, and $\tilde{\phi}$ are all dependent on the ℓ range considered.

As a consequence of the central limit theorem, for sufficiently large $\ell_{\max}$ the $\tilde{X}$s are Gaussian-distributed with mean zero, so that the amplitude parameter has a Maxwell-Boltzmann distribution. We fit to this distribution for $\ell_{\max} \ge 10$ when computing the *p*-value, so as not to be influenced by Poisson noise in the tails of the empirical distribution (and we have determined that this is a good fit to the simulations by applying a KS test). For the case of scalar amplitude modulation (i.e., $X = A_s$), and $\ell_{\min} = 2$, the cosmic-variance-limited expectation for the modulation amplitude from statistically isotropic skies is

$$\left\langle \frac{\Delta A_s}{A_s} \right\rangle \approx \sqrt{\frac{48}{\pi(\ell_{\max} + 4)(\ell_{\max} - 1)}}. \tag{50}$$

This is the cosmic variance for a scale-invariant dipole modulation, and gives a more explicit expression than the $\ell_{\max}^{-1}$ scaling discussed in Hanson & Lewis (2009).

The top panel of Fig. 30 presents results for the *p*-value of the fitted modulation amplitude as a function of $\ell_{\max}$. Note that there are several peaks, at $\ell \approx 40$ and $\ell \approx 67$ (the focus of most attention in the literature), and $\ell \approx 240$. The latter peak, while not previously emphasized, is also present in the WMAP results (see Fig. 15 in Bennett et al. 2011). It is also interesting to note that a modulation amplitude is observed at $\ell_{\max} \approx 800$ that is somewhat lower than what one would typically expect for a statistically isotropic sky. However, the significance is not at the level of the excess dipole modulation at low ℓ and will not be discussed further. The dip at $\ell_{\max} \approx 67$, with a *p*-value of
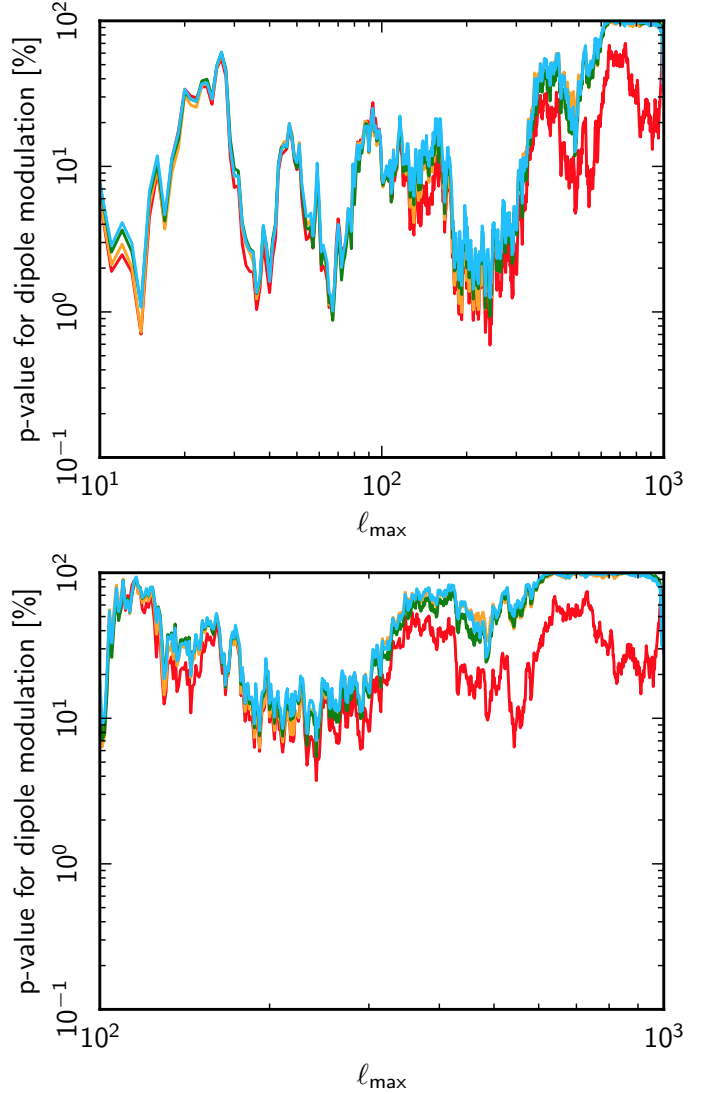

Fig. 30. Probability determined from the QML analysis for a Monte Carlo simulation to have a larger dipole modulation amplitude than the Commander (red), NILC (orange), SEVEM (green), or SMICA (blue) data sets, with (*top panel*) $\ell_{\min} = 2$ or (*bottom panel*) $\ell_{\min} = 100$. No significant modulation is found once the low-ℓ signal is removed. We emphasize that the statistic here is cumulative and apparent trends in the curves can be misleading.

0.9–1.0%, corresponds to the well-known low-ℓ dipole modulation[6]. Table 23 presents the corresponding dipole modulation parameters, which are seen to be consistent with previous studies. Note that the mean amplitude expected for a set of statistically isotropic simulations at this $\ell_{\max}$ is 2.9% (in close agreement with the expected value due to cosmic variance, Eq. (50)).

We have therefore determined a phenomenological signature of modulation for $\ell = 2$–67 with a *p*-value of 0.9–1.0%. If such a signal had been predicted by a specific model, then we could claim a significance of about 3σ. However, in the absence of such an a priori model, we can assess how often we might find a 3σ effect by chance, given that it could have occurred over any

---

[6] Actually only SEVEM and SMICA achieve their minimum at $\ell_{\max} = 67$, whereas NILC and Commander achieve theirs at $\ell_{\max} = 14$ and 240, respectively. Such scatter is expected when searching over a large number of possible ℓ ranges. The reconstructed amplitudes for each component-separation method are well within the error budgets of the estimator.
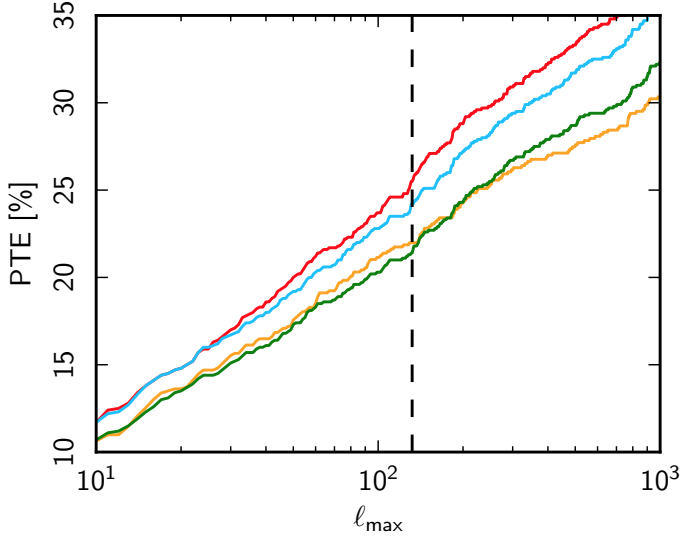
Fig. 31. Probability determined from the QML analysis for obtaining a dipole modulation amplitude at least as anomalous as the `Commander` (red), `NILC` (orange), `SEVEM` (green), and `SMICA` (blue) data sets, for the range $\ell \in [10, \ell_{\max}]$. The vertical line corresponds to $\ell_{\max} = 132$ which was used as the search limit in Bennett et al. (2011). The probability grows approximately logarithmically with $\ell_{\max}$. This means that the adopted probability to exceed is fortunately not very sensitive to $\ell_{\max}$, and for any reasonable choice is above 10%.

ℓ range. Since we are looking for a large-scale phenomenon, we assume that the analysis should include the corresponding low-ℓ modes and start at $\ell = 2$. In order to correct for a posteriori effects we then adopt the following scheme.

1. We calculate the modulation of each simulation on the scales 2–ℓ, where $\ell \in [10, \ell_{\max}]$. For each simulation we find the modulation that gives the smallest probability, η (in the same way that was done for the data).
2. With the distribution of ηs given by the simulations we then compare this to the data. That is, we calculate the probability that one would find oneself in a Hubble patch with a modulation amplitude up to $\ell \in [10, \ell_{\max}]$ that is as significant as (or more significant than) the modulation in the real data.

If $\ell_{\max} = 132$ (as chosen by Bennett et al. 2011), the probability of achieving a modulation as large as the *Planck* data in this range is higher than 10% (see Fig. 31). This is in agreement with the findings of the WMAP team (which found 10% and 13% in the same ℓ-range, using two different masks). Here, we do not quote a specific PTE for the dipole modulation since it depends on the choice of both $\ell_{\max}$ (albeit not so sensitively) and $\ell_{\min}$ (which we have decided not to marginalize over). However, it appears to be the case that the dipole modulation that we observe is quite unremarkable. That is, Gaussian fluctuations in a statistically isotropic Universe will reasonably often result in a dipole modulation with a comparable level of significance to that presented here.

Beyond this, evidence for dipole modulation is found at $\ell \approx 200$–300, with a smaller dip at $\ell \approx 500$. Given that the dipole modulation estimator is a cumulative quantity, it is possible that these features are statistically enhanced by the usual low-ℓ signal. To test this we analyse the dipole modulation as a function of $\ell_{\max}$ again, with the restriction $\ell_{\min} = 100$ applied in order to completely remove any low-ℓ influence. The outcome is presented in Fig. 30 (bottom). It is clear that even before introducing posterior corrections no significant modulation is found,

indicating that the p-values of the features at $\ell > 100$ were indeed exaggerated by the low-ℓ modulation.

6.4. Bipolar spherical harmonics

In the absence of the assumption of statistical isotropy, the CMB two-point correlation function $C(\hat{n}_1, \hat{n}_2) \neq C(\hat{n}_1 \cdot \hat{n}_2)$ can be most generally expanded in the bipolar spherical harmonic (BipoSH) basis representation as follows:

$$C(\hat{n}_1, \hat{n}_2) = \sum_{LM\ell_1\ell_2} \tilde{A}^{LM}_{\ell_1\ell_2} \{Y_{\ell_1}(\hat{n}_1) \otimes Y_{\ell_2}(\hat{n}_2)\}_{LM}. \qquad (51)$$

The BipoSH basis functions, $\{Y_{\ell_1}(\hat{n}_1) \otimes Y_{\ell_2}(\hat{n}_2)\}_{LM}$ are tensor products of ordinary spherical harmonic functions, and the corresponding expansion coefficients are termed BipoSH coefficients (Hajian & Souradeep 2003; Hajian & Souradeep 2006). The BipoSH basis provides a complete representation of any form of statistical isotropy violation with the key advantage of separating the angular scale-dependence of the signal in spherical harmonic multipoles, ℓ, from the nature of the violation indexed in the bipolar multipole space by L. Consequently, it is possible to simultaneously determine that such a signal is dipolar ($L = 1$), quadrupolar ($L = 2$), octopolar ($L = 3$), and so on, in nature and that the power is restricted to specific ranges of angular scales.

The estimation of BipoSH coefficients from CMB maps is a natural generalization of the more routinely undertaken estimation of the angular power spectrum C_l. To allow a direct connection to the angular power, we further introduce a set of BipoSH spectra at every bipolar harmonic moment, (L, M), labelled by a difference index d, defined as follows:

$$A^{LM}_{\ell\ell+d} = \tilde{A}^{LM}_{\ell\ell+d} \frac{\Pi_L}{\Pi_{\ell(\ell+d)} C^{L0}_{\ell0(\ell+d)0}}, \quad (0 \le d \le L), \qquad (52)$$

where $C^{LM}_{\ell_1m_1\ell_2m_2}$ are the Clebsch-Gordon coefficients and for brevity the notation $\Pi_{\ell_1\ell_2..\ell_n} = \prod_{i=1}^{n} \sqrt{(2\ell_i + 1)}$. BipoSH spectra, clearly, are then simply a generalized set of CMB angular power spectra, with the standard CMB angular power spectrum $C_\ell = A^{00}_{\ell\ell}$ being one of them[7]. While $A^{00}_{\ell\ell}$ quantifies the properties of the statistically isotropic part of the CMB fluctuations, the additional BipoSH coefficients quantify the statistically anisotropic part of the CMB two-point correlation function.

Thus BipoSH provides a mathematically complete description of all possible violations of statistical isotropy in a Gaussian CMB sky map. It is then always possible to translate any specific model for such a signal into the language of BipoSH and provide a common approach for the multiple specialized tests that have been implemented previously in this paper and elsewhere. However, improving on the analysis of the 2013 *Planck* data, a new formalism is developed in order to reliably analyse a masked sky, as concisely described in Appendix D. Aluri et al. (2015) provides a more detailed description of the approach and includes an explicit demonstration of its validity using simulations.

---

[7] The BipoSH spectra, as defined in Eq. (52), restrict us to working with only even-parity BipoSH coefficients ($L+d$ is even) due to the vanishing of $C^{L0}_{\ell0\ell+d0}$ otherwise. While most known isotropy-violating phenomena like weak lensing, Doppler boost, non-circular beams, etc., can only produce even-parity BipoSH spectra, measurement of odd-parity BipoSH spectra can be used to test for systematic effects, or to search for the signatures of exotic effects such as the lensing of CMB photons by tensor metric perturbations.

Table 24. Amplitude (*A*) and direction of the dipole modulation in Galactic coordinates as estimated for the multipole range $\ell \in [2, 64]$ using a BipoSH analysis.

Method	A	Direction (l, b) [°]
Commander ..	0.067 ± 0.023	$(230, -18) \pm 31$
NILC.......	0.069 ± 0.022	$(228, -17) \pm 30$
SEVEM......	0.067 ± 0.023	$(230, -17) \pm 31$
SMICA......	0.069 ± 0.022	$(228, -18) \pm 30$
SEVEM-100 ..	0.070 ± 0.023	$(231, -19) \pm 30$
SEVEM-143 ..	0.068 ± 0.023	$(230, -17) \pm 31$
SEVEM-217 ..	0.069 ± 0.023	$(229, -20) \pm 31$

Notes. The measured values of the dipole amplitude and direction are consistent for all maps.

Initially, we revisit the simple phenomenological model of dipole modulation of the CMB sky from Sect. 6.2,

$$T(\hat{\boldsymbol{n}}) = T_0(\hat{\boldsymbol{n}})\,(1 + \mathcal{M}(\hat{\boldsymbol{n}}))\,, \tag{53}$$

where $T(\hat{\boldsymbol{n}})$ represents the modulated CMB sky, $T_0(\hat{\boldsymbol{n}})$ is the underlying (statistically isotropic) random CMB sky, and $\mathcal{M}(\hat{\boldsymbol{n}})$ is a dipolar field. The BipoSH coefficients resulting from such a modulation are given by

$$A_{\ell\ell+1}^{1M} = \bar{A}_{\ell\ell+1}^{1M} + m_{1M}G_{\ell\ell+1}^{1}, \tag{54}$$

$$G_{\ell\ell+1}^{1} = \frac{C_\ell + C_{\ell+1}}{\sqrt{4\pi}} \sqrt{\frac{(2\ell+1)(2\ell+3)}{3}} C_{\ell0(\ell+1)0}^{10}. \tag{55}$$

Here $\bar{A}_{\ell\ell+1}^{1M}$ corresponds to the BipoSH coefficients of the unknown, but statistically isotropic, unmodulated CMB field, m_{1M} are the spherical harmonic coefficients of the modulation field, and C_ℓ is the best-fit CMB angular power spectrum.

The BipoSH representation further enables an estimate of the modulation field to be made over specific angular scales by windowing regions in multipole space in the sum over multipoles ℓ in Eq. (55). This additional information is important for identifying the origin of the isotropy-breaking signal, which could be either cosmological or due to systematic artefacts.

We perform the analysis for the $N_{\text{side}} = 2048$ component separated CMB maps with an apodized version of the common mask at that resolution and reconstruct the modulation signal in independent bins of width $\Delta\ell = 64$ up to $\ell_{\max} = 512$. The application of the common mask introduces a mean field bias in the BipoSH coefficients derived from the data. This bias is estimated from the FFP8 simulations and subtracted from the derived coefficients. The process of masking induces a coupling between the modulation field and the mask that results in a modification of the spectral shape of the modulation signal by the modified shape function (MSF; see Appendix D for details). Further, the covariance of the bias-subtracted BipoSH coefficients is not easy to derive analytically in this case. To overcome this problem, we consider the diagonal approximation to the covariance matrix and estimate it from simulations.

The results presented in the top panel of Fig. 32 indicate that the dipole modulation signal is most significant in the lowest multipole window $\ell \in [2, 64]$. Note that the power in the dipole modulation field $m_1 = (|m_{11}|^2 + |m_{10}|^2 + |m_{1-1}|^2)/3$ is related to the dipole amplitude by $A = 1.5\sqrt{m_1/\pi}$. The best-fit amplitude (*A*)
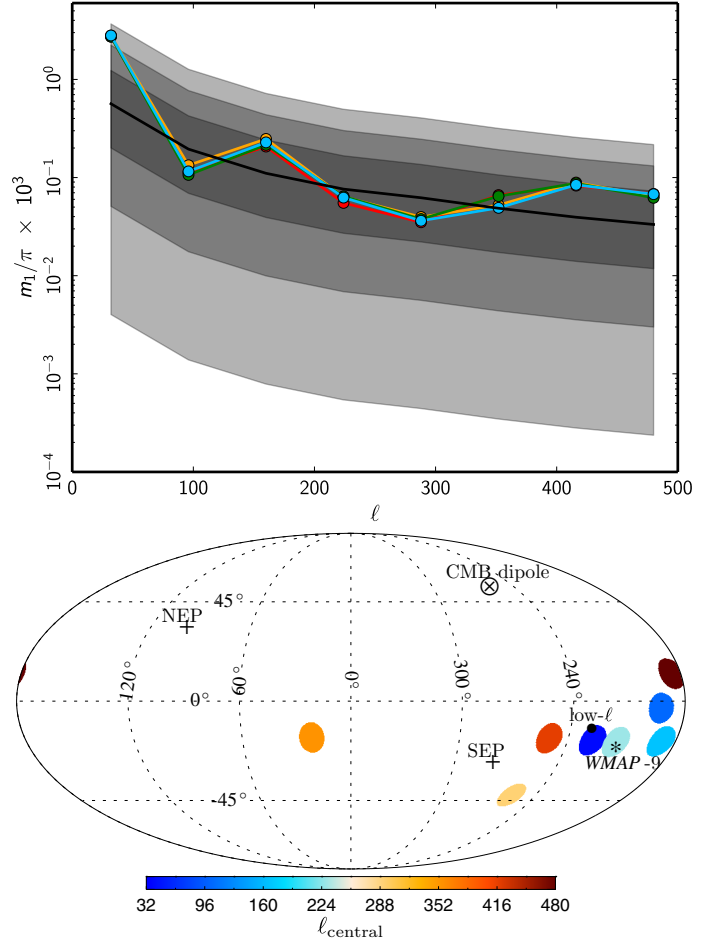

Fig. 32. *Top*: measured dipole modulation ($L = 1$) power in non-overlapping CMB multipole bins for Commander (red), NILC (orange), SEVEM (green), and SMICA (blue) as determined from a BipoSH analysis of the data. The power in the dipole of the modulation field is a χ^2-distributed variable with 3 degrees of freedom. The shaded regions in the plot depict, in dark-grey, grey, and light-grey respectively, the 1, 2, and 3σ equivalent intervals of the distribution function derived from simulations, while the solid black line denotes its median. Significant power in the dipole modulation is seen to be limited to $\ell = 2$–64 and does not extend to higher multipoles. *Bottom*: dipole modulation direction as determined from the SMICA map. The directions found from the other component separation maps are consistent with this analysis. The coloured circles denote the central value of the multipole bin used in the analysis, as specified in the colour bar. The low-ℓ and WMAP-9 directions are identical to those in Fig. 35.

and direction corresponding to the reconstructed dipole modulation field from this lowest multipole bin is quoted in Table 24 for each component-separation method. Also shown are the corresponding results for the cleaned frequency maps SEVEM-100, SEVEM-143, and SEVEM-217. As expected for signals with a cosmological origin, no evidence for frequency dependence is seen.

Since the amplitude of the dipole modulation field is consistent with zero within 2σ for all of the higher ℓ-bins considered, it is plausible that the simple modulation model in Eq. (53) is inadequate to describe the features seen in the BipoSH spectra and should minimally allow for the amplitude, $A(\ell)$, of the dipole to depend on CMB multipole, ℓ. Although this may appear to be a more complex model, it does not necessarily lack motivation. It is readily conceivable that physical mechanisms that cause a dipolar modulation of the random CMB sky would be
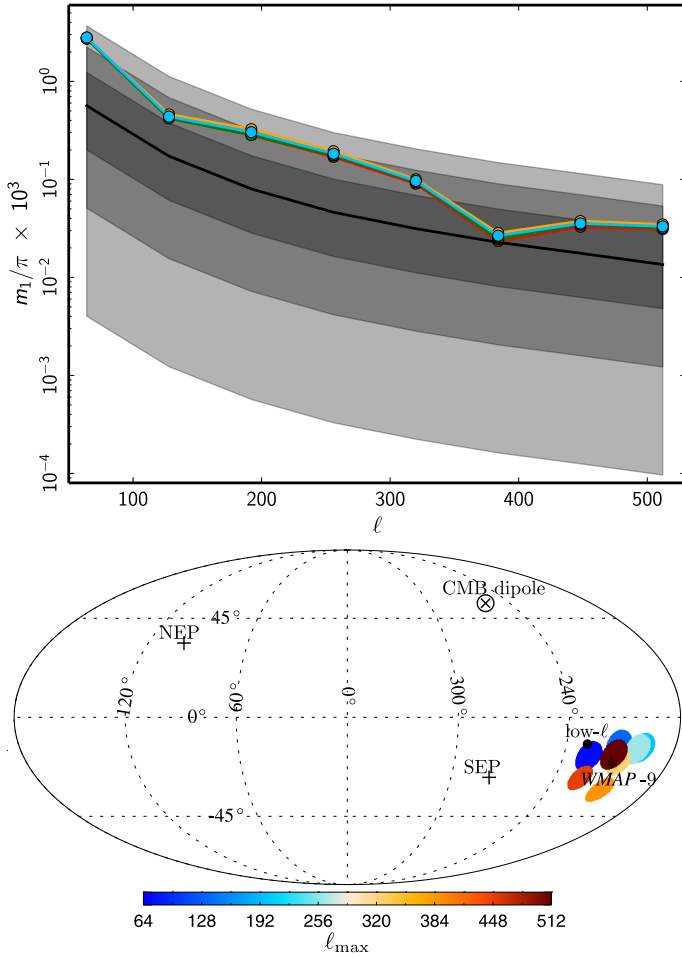
Table 25. Doppler boost amplitude ($|\boldsymbol{\beta}|$) and direction in Galactic coordinates derived over the multipole range $\ell \in [640, 1024]$ as evaluated from a BipoSH analysis.

Method	$|\boldsymbol{\beta}| \times 10^{-3}$	Direction (l, b) [°]
SEVEM-100	1.24 ± 0.66	$(277, 40) \pm 50$
SEVEM-143	1.35 ± 0.56	$(264, 39) \pm 39$
SEVEM-217	1.28 ± 0.45	$(257, 42) \pm 32$

Notes. The errors are estimated from an identical analysis of a set of 1000 Doppler boosted simulations for each frequency.

would not contaminate a dipole modulation signal up to $\ell_{max} \approx 700$. We now confirm the Doppler boost signal using the BipoSH methodology.

An equivalent description of the Doppler boost in terms of BipoSH coefficients is given by

$$A^{1M}_{\ell_1\ell_2} = \bar{A}^{1M}_{\ell_1\ell_2} + \beta_{1M} G^1_{\ell_1\ell_2}, \tag{56}$$

$$G^1_{\ell_1\ell_2} = \left\{ b_\nu [G^1_{\ell_1\ell_2}]^{\mathcal{M}} - [G^1_{\ell_1\ell_2}]^\phi \right\}$$

$$\times \sqrt{\frac{(2\ell_1 + 1)(2\ell_2 + 1)}{12\pi}} C^{10}_{\ell_10\ell_20}, \tag{57}$$

$$\left[G^1_{\ell_1\ell_2} \right]^{\mathcal{M}} = [C_{\ell_1} + C_{\ell_2}], \tag{58}$$

$$\left[G^1_{\ell_1\ell_2} \right]^\phi = [C_{\ell_1} + C_{\ell_2}] \tag{59}$$

$$+ [C_{\ell_1} - C_{\ell_2}] [\ell_1(\ell_1 + 1) - \ell_2(\ell_2 + 1)] / 2,$$

where $\beta_{1M} = \int dn Y_{1M}(\hat{\boldsymbol{n}}) \boldsymbol{\beta} \cdot \hat{\boldsymbol{n}}$, $\boldsymbol{\beta} = \boldsymbol{v}/c$ denotes the peculiar velocity of our local rest frame with respect to the CMB, and b_ν is the frequency-dependent boost factor, as discussed in more detail in Planck Collaboration XXVII (2014).

Since the Doppler boost signal has a frequency dependence, we perform our analysis on the SEVEM-100, SEVEM-143, and SEVEM-217 maps at $N_{side} = 2048$, and adopt values of $b_\nu = 1.51$, 1.96, and 3.07, respectively. A minimum variance estimator for β_{1M}, as discussed in Appendix D, is adopted with the shape function $G^L_{\ell_1\ell_2}$ replaced by the corresponding Doppler boost term given in Eq. (56). Corresponding unboosted CMB simulations were also used, in particular to correct for the mean field bias. However, we use a set of Doppler-boosted simulations in order to estimate the error on the reconstructed Doppler boost vector.

Since it is expected that the low multipole modes of the $A^{1M}_{l,l+1}$ spectrum are contaminated by the dipolar signal reported previously, in order to monitor the impact of this anomalous signal on the Doppler reconstruction we implement a cumulative analysis using multipoles with a varying ℓ_{min} from 2 to 640 in increments of $\Delta\ell_{min} = 128$ and a fixed $\ell_{max} = 1024$[8]. The recovered Doppler amplitudes from the three SEVEM frequency cleaned maps as a function of ℓ_{min} are shown in the top panel of Fig. 34, while the lower panel indicates the corresponding direction $\hat{\boldsymbol{\beta}}$ in Galactic coordinates determined from the SEVEM-217 data. Table 25 records the best-fit amplitudes and directions for $\ell \in [640, 1024]$.

Fig. 33. *Top*: measured dipole modulation power in cumulative CMB multipole bins for Commander (red), NILC (orange), SEVEM (green), and SMICA (blue) as determined from a BipoSH analysis of the data.. Colour coding as in Fig. 32. Note that the measurements in cumulative bins indicate a power in excess of 2σ up to multipole $\ell_{max} \sim 320$. The value on the horizontal axis denotes the maximum multipole used in the analysis, with $\ell_{min} = 2$. *Bottom*: modulation dipole direction as recovered from the SMICA map. The directions found from the other component-separation maps are consistent with these directions. The colour-coded points represent the directions recovered for the specific ℓ_{max} used in the analysis, with $\ell_{min} = 2$. The low-ℓ and WMAP-9 directions are identical to those in Fig. 35.

scale-dependent and possibly significant only at low wavenumbers. It is also intriguing to note that, although in most cases the amplitude of the modulation dipole is seen at low significance, the directions in the first four bins, $\ell_{32} \in [2, 64]$, $\ell_{96} \in [65, 128]$, $\ell_{160} \in [129, 192]$, and $\ell_{224} \in [193, 256]$, are seen to be clustered together, as shown in the bottom panel of Fig. 32. Note that the lower significance of the modulation for the multipole bins at $\ell > 64$ results in larger errors for their respective directions than the value quoted for the $\ell \in [2, 64]$ bin recorded in Table 24.

We extend our analysis to carry out the dipole modulation reconstruction in cumulative bins up to $\ell_{max} = 512$, making cumulative increments in the multipole in steps of $\Delta\ell = 64$. The results of this analysis are summarized in Fig. 33.

As noted previously, as a consequence of our motion with respect to the CMB rest frame, the observed CMB map is expected to be statistically anisotropic, as has been demonstrated in Planck Collaboration XXVII (2014) and Appendix B. Reassuringly, in PCIS13 it was established that such a signal

---

[8] We fix $\ell_{max} = 1024$ since at higher ℓ values the mismatch between the data and simulation power spectra becomes more important and is a concern for the bias subtraction applied when reconstructing the Doppler boost signal.
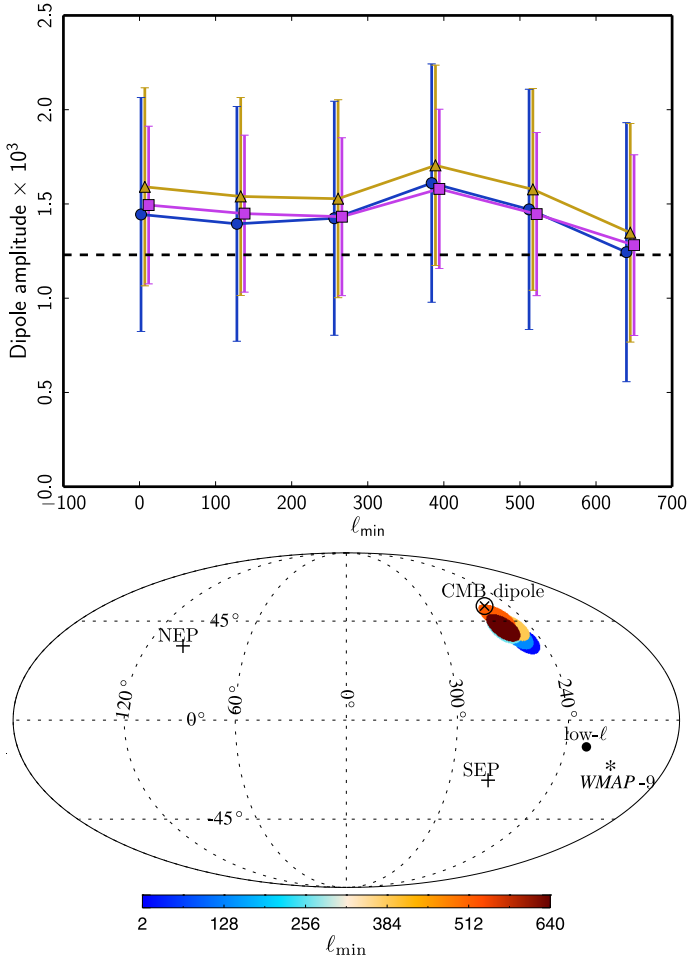
Fig. 34. *Top*: amplitude $|\beta|$ of the Doppler boost from the `SEVEM-100`, `SEVEM-143`, and `SEVEM-217` maps for different multipole bins determined using a BipoSH analysis. The maximum multipole of each bin is fixed at $\ell_{max} = 1024$, while ℓ_{min} is incremented from $\ell = 2$ to $\ell = 640$ in steps of $\Delta\ell = 128$. The dashed line corresponds to the actual dipole boost amplitude, $|\beta| = 1.23 \times 10^{-3}$. *Bottom*: Doppler boost direction $\hat{\beta}$ measured in Galactic coordinates from `SEVEM-217`. The coloured circles denote ℓ_{min} used in the analysis, while $\ell_{max} = 1024$ is held fixed. The low-ℓ and WMAP-9 directions are identical to those in Fig. 35.

6.5. Angular clustering of the power distribution

In the *Planck* 2013 data release we reported a possible deviation from statistical isotropy in the multipole range $\ell = 2$–600, thus confirming earlier findings based on the WMAP data (Hansen et al. 2009; Axelsson et al. 2013). This claim of asymmetry extending to higher multipoles was made only on the basis of the alignment of preferred directions as determined from maps of the power distribution on the sky for specific multipole ranges. In particular, it was found that the directions of the dipoles fitted to such maps in the multipole range $\ell = 2$–600 were significantly more aligned than in simulations. In addition, we showed that the ratio of the power spectra in the two opposite hemispheres defined by the asymmetry axis for $\ell = 2-600$ was not statistically anomalous (as later confirmed over the extended multipole range $\ell = 2-2000$ by Quartin & Notari 2015).

Here, we test for the alignment in the *Planck* 2015 data set. We adopt the approach for the estimation of the dipole alignment that was described in detail in PCIS13, a brief summary of which follows.

1. Local power spectra are estimated from the data at $N_{side} = 2048$ for 12 patches of the sky corresponding to the $N_{side} = 1$ `HEALPix` base pixels. Only those high-resolution pixels surviving the application of the common mask are included in the analysis[9]. As a consequence of this masking, when patches based on `HEALPix` pixels with $N_{side} > 1$ are used, the available sky fraction for those patches close to the Galactic plane is too small for power-spectrum estimation. For most of the analysis, we use the cross-spectra determined from half-mission data sets[10]. Due to a mismatch between the noise level in the data and the simulated maps, the results based on auto-spectra are less reliable and also more prone to other systematic effects than the cross-spectra. We therefore do not consider such results here. The spectra are binned over various bin sizes between $\Delta\ell = 8$ and $\Delta\ell = 32$.

2. For each power spectrum multipole bin, an $N_{side} = 1$ `HEALPix` map with the local power distribution is constructed.

3. The best-fit dipole amplitude and direction are estimated from this map using inverse-variance weighting, where the variance is determined from the local spectra computed from the simulations. We do not compute error bars for the direction, but expect this to be accounted for in part by the use of equivalently treated simulations in the clustering analysis.

4. A measure of the alignment of the different multipole blocks is then constructed. In PCIS13, we considered the mean angle between all possible pairs of dipole directions up to a given ℓ_{max}. Here, for greater consistency with Sect 6.6, we use the mean of the cosine of the angles, rather than of the angles themselves, between all pairs of dipoles. This effectively corresponds to the Rayleigh statistic (RS) introduced formally in Sect. 6.6, and we will refer to it as such, although it differs by ignoring all amplitude information. Clearly, smaller values of the RS correspond to less clustering.

5. The clustering as a function of ℓ_{max} is then assessed using p-values determined as follows. We first construct the RS using all multipoles up to ℓ_{max}. The p-value is then given by the fraction of simulations with a higher RS than for the data for this ℓ_{max}. A small p-value therefore means that there are few simulations that exhibit as strong clustering as the data. Note that the p-values are highly correlated as the RS is a cumulative function of ℓ_{max}.

6. We then define two measures of significance. To achieve this, it is necessary to reduce the 1499 different p-values determined for $\ell_{max} \in [2, 1500]$ to a single measure of clustering. We do this in two different ways, using the mean of these p-values, and by finding the minimum of the p-values, for both the data and for each available simulation. We then determine the percentage of simulations with (i) a lower mean p-value and (ii) a lower minimum p-value than the data. Note that these two measures of significance take into account different aspects of the data. Note further that since the RS is cumulative and the p-values therefore correlated, different scales are weighted unequally and a detection in the mean

---

[9] Departing from the analysis in PCIS13, we do not use an apodized version of the common mask. Simulations indicate that the error on the power spectrum for those multipoles in the range 300 to 500 where the significance is highest is up to 20% larger in this case, with the corresponding error on preferred direction being typically 8% larger.

[10] Note that simulated half-mission noise maps were generated by adjusting the properties of the existing 1000 (10 000 in the case of `SMICA`) noise simulations appropriately, thus explaining why only 500 (5000) simulations are used in this analysis.
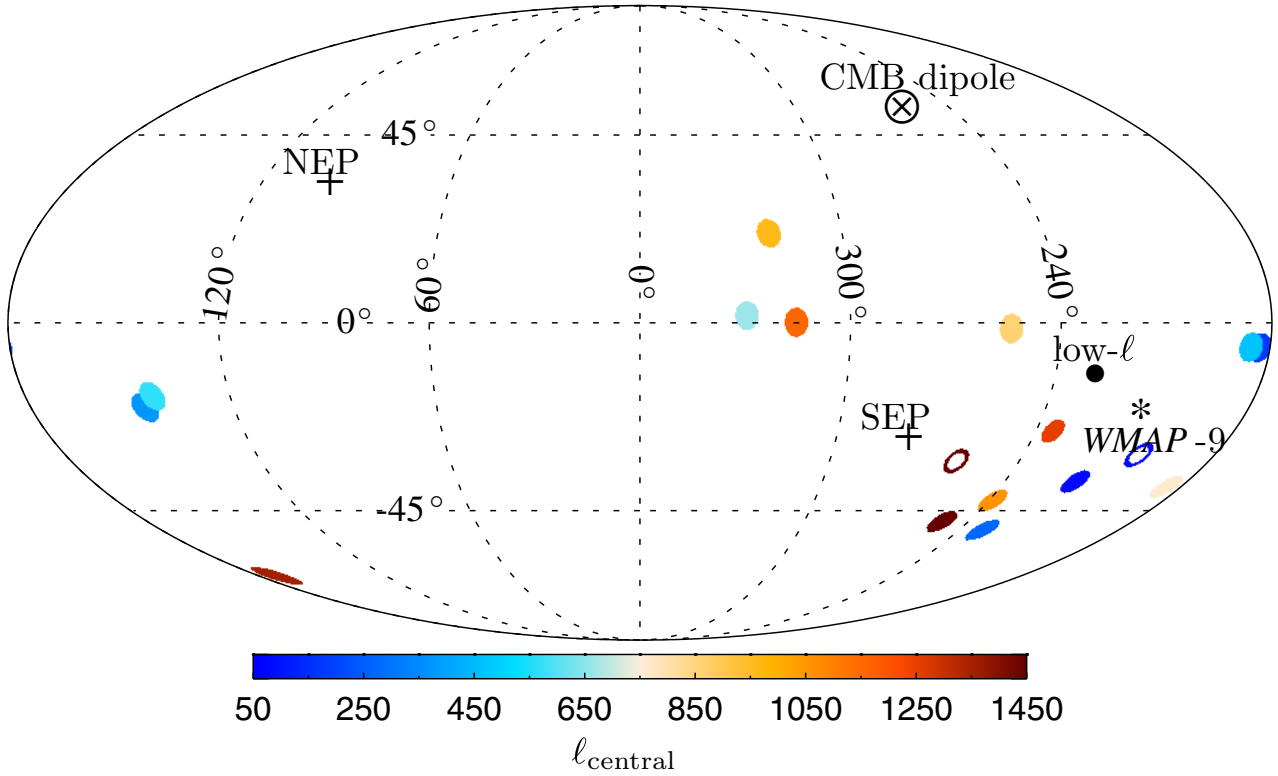
Fig. 35. Dipole directions for independent 100-multipole bins of the local power spectrum distribution from $\ell = 2$ to 1500 in the `SMICA` map with the common mask applied. We also show the preferred dipolar modulation axis (labelled as "low-ℓ") derived in Sect. 6.2, as well as the total direction for $\ell_{max} = 600$ determined from WMAP-9 (Axelsson et al. 2013). The average directions determined from the two multipole ranges $\ell \in [2, 300]$ and $\ell \in [750, 1500]$ are shown as blue and red rings, respectively. The error on the derived direction that results from masking the data is about $60°$, with only small variations related to bin size.

and/or minimum p-value may be difficult to interpret and to correct for the multiplicity of tests effect (LEE).

Note that the statistics defined in step 6 above correspond to two choices of what were referred to as "global statistics" in PCIS13 in order to assess the degree to which the significance of the results depends on a specific choice for ℓ_{max}. The mean p-value over all available ℓ_{max} measures the degree to which clustering is present over large multipole ranges independently of whether the clustering is strongly focused in one given direction. Clearly the p-values for different ℓ_{max} are strongly correlated, but if the clustering is present only over a small multipole range, the RS will drop and the corresponding p-values will eventually rise. By comparing this value to simulations, we test not only whether the dipole alignment in the data is stronger than in statistically isotropic random simulations, but also whether it is present over larger ranges of multipoles than expected. The minimum p-value will give strong detections if there is a strong asymmetry over a limited multipole range or weaker clustering over larger multipole ranges when the clustering is strongly focused in a given direction.

For `Commander`, `NILC`, and `SEVEM`, only 500 simulations are available. However, 5000 simulations are available for `SMICA`, which allows a better estimate of significance to be determined when the probabilities obtained are very low. In this case, we use half of the 5000 simulations to calibrate the statistic (obtain p-values following step 5 above) and the remaining half to determine significance levels (compute the mean and minimum over these p-values as a function of ℓ_{max} following step 6). When using 500 simulations, it is necessary to use the same set of simulations to calibrate as well as to obtain probabilities. A

related issue with these results is that this set of simulations (corresponding to the first 500 out of the 5000 available for `SMICA`) are observed to yield higher p-values for the clustering angle due to a statistical fluctuation. Another 9 sets of 500 simulations that can be obtained from partitioning the 5000 available `SMICA` simulations all result in lower p-values. As a consequence, we observe that results based on the larger number of simulations often give lower p-values than when only 500 simulations are used.

In Fig. 35 we show the dipole directions of the 15 lowest 100-multipole bins for the `SMICA` map. Here, the binning has been chosen for visualization purposes; in further analysis of the *Planck* data we use finer ℓ-intervals. The preferred low-ℓ modulation direction determined in Sect. 6.2 is also indicated, along with the WMAP-9 result determined over the range $\ell = 2$ to 600 (Axelsson et al. 2013). The observed clustering of the dipole directions is similar to that shown in figure 27 of PCIS13. Note that differences in masking, foreground subtraction, and residual systematic effects will displace the direction of a given dipole with respect to the previous analysis. Similar behaviour is seen for all of the *Planck* component-separated maps.

In PCIS13, we calculated the mean angle between all possible pairs of dipole directions determined from maps of the local power in multipole bins of size $\Delta\ell = 16$. Here we test the possible bias arising from such a choice by considering bin sizes between $\Delta\ell = 8$ and $\Delta\ell = 32$ in steps of 2. The lower limit avoids significant bin-to-bin coupling in the power spectra for smaller binnings, whilst the upper limit excludes cases where there are an insufficient number of derived dipoles from which the mean angle can be calculated, this leading to poor statistics. In addition to showing results for each bin size, we also calculate the
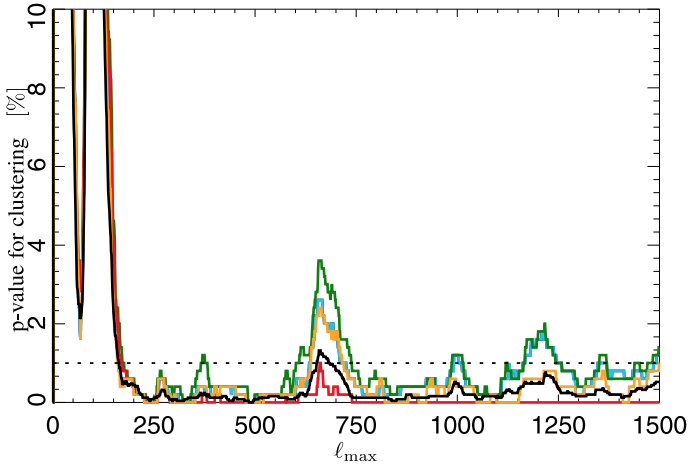
Fig. 36. Derived *p*-values for the angular clustering of the power distribution as a function of ℓ_{max}, determined for `Commander` (red), `NILC` (orange), `SEVEM` (green), and `SMICA` (blue), based on 500 simulations. For SMICA, the *p*-values based on 2500 simulations are also shown (black). The *p*-values are based on the fraction of simulations with a higher RS, determined over the ℓ-range up to the given ℓ_{max}, compared to the data. The results shown here have been marginalized over bin sizes in the range $\Delta\ell = 8$ to $\Delta\ell = 32$.

variance-weighted mean of the power spectra over all bin sizes (the C_ℓ for a given bin size is weighted by $1/\sqrt{N_b}$ where N_b is the bin size). In this way, we marginalize over bin sizes to obtain local power spectra and thereby the RS for each single multipole.

Figure 36 shows the *p*-values for the different component-separated maps, derived as described in step 5 above. We see that the results based on 500 simulations for `NILC`, `SEVEM`, and `SMICA` are in good agreement. The `Commander` results are less consistent, but this may be related to the fact that component separation was performed independently for the half-mission solutions, in contrast to the other methods, where component-separation solutions were obtained from the full mission data only. For `SMICA`, we also show *p*-values based on 2500 simulations. These more accurate results show lower *p*-values, and may indicate that those determined from only 500 simulations are not sufficiently stable. Note also that for $\ell < 100$ the *p*-values are not consistent with the detection of a low-ℓ asymmetry/modulation, as seen by other methods in this paper. However, for $\ell < 100$, there are very few bins and the variance of the RS might therefore be too high for this effect to be visible.

In agreement with the conclusions in PCIS13, a large degree of alignment is seen at least to $\ell_{max} \approx 600$. However, in contrast to the earlier results where the *p*-values started increasing systematically for $\ell_{max} > 1000$, the current *p*-values remain low for $\ell_{max} > 750$. The full component-separated maps which have higher resolution and sensitivity are used for the current analysis, instead of the single-frequency foreground-cleaned map (`SEVEM-143`) used in PCIS13. We note that the results for the updated `SEVEM-143` map are consistent with the earlier analysis, both with and without correction for the Doppler modulation. Note also that the `SMICA` results with improved statistics (based on 2500 simulations) generally show lower *p*-values than the corresponding results based on 500 simulations.

Table 26 presents the fraction of simulations with a lower mean/minimum *p*-value than in the data for a number of different cases. The table shows probabilities for `SMICA` with different bin sizes (showing only every second bin size since these are correlated), as well as for the results marginalized over bin sizes. We

Table 26. Significance of the angular clustering of the power distribution.

Method	Bin size	Mean *p*-value	% (mean)	Min. *p*-value	% (min)
SMICA	8	261/2500	1.60	35/2500	16.2
SMICA	10	51/2500	0.08	3/2500	2.36
SMICA	12	75/2500	0.20	1/2500	0.96
SMICA	14	83/2500	0.16	2/2500	1.52
SMICA	16	78/2500	0.24	4/2500	2.00
SMICA	18	51/2500	0.04	1/2500	0.68
SMICA	20	21/2500	<0.04	1/2500	0.76
SMICA	22	60/2500	0.08	2/2500	1.24
SMICA	24	34/2500	0.08	2/2500	1.00
SMICA	26	38/2500	0.08	1/2500	0.96
SMICA	28	42/2500	0.20	0/2500	<0.52
SMICA	30	27/2500	0.20	0/2500	<0.60
SMICA	32	21/2500	0.04	0/2500	<0.52
SMICA	marg.	43/2500	<0.04	0/2500	<1.00
SMICA[a]	marg.	48/2500	<0.04	1/2500	1.70
SMICA[b]	marg.	47/2500	<0.04	0/2500	<1.16
SMICA[c]	marg.	50/2500	<0.04	0/2500	<0.76
SMICA[d]	marg.	254/2500	1.52	34/2500	20.1
Comm.	marg.	9/500	<0.20	0/500	<2.60
NILC	marg.	10/500	<0.20	0/500	<3.60
SEVEM	marg.	13/500	<0.20	0/500	<4.00
SMICA	marg.	11/500	<0.20	0/500	<3.60
Comm.[b]	marg.	11/500	<0.20	0/500	<3.00
NILC[b]	marg.	10/500	<0.20	0/500	<3.80
SEVEM[b]	marg.	12/500	<0.20	0/500	<3.40
SMICA[b]	marg.	11/500	<0.20	0/500	<3.80
Comm.[c]	marg.	8/500	0.20	0/500	<4.00
NILC[c]	marg.	14/500	0.20	1/500	7.20
SEVEM[c]	marg.	17/500	0.20	1/500	8.40
SMICA[c]	marg.	15/500	0.20	1/500	7.60

Notes. We indicate the actual mean/min *p*-value of the data, determined from Fig. 36 and written as a fraction of the number of simulations used to assess the values, together with the percentage of simulations with a lower mean/minimum *p*-value than the data. Unless otherwise specified, the numbers are determined from half-mission cross spectra $C_\ell\ell(\ell + 1)$, for all multipoles in the range $\ell = 2$–1500, and for the common mask. [a] Half-ring maps instead of half-mission maps. [b] $C_\ell(2\ell + 1)$ instead of $C_\ell\ell(\ell + 1)$. [c] Restricted to multipoles $\ell > 100$. [d] Restricted to multipoles $\ell > 200$.

also show results for the different component-separated maps, results based on half-ring cross-spectra instead of half-mission cross-spectra, and results using a different ℓ-weighting scheme, specifically $(2\ell + 1)C_\ell$ instead of $\ell(\ell + 1)C_\ell$, the former being a measure of the variance of the temperature fluctuations. The table indicates probabilities of approximately 0–2% for most of these cases, although results for the smallest bin size show much less significant results. This could be due to the strong anticorrelations between adjacent bins found for this bin size in those Galactic $N_{side} = 1$ patches with very small available sky fraction. For the other bin sizes, these correlations are much weaker. Note that many of the significances based on minimum *p*-value are only upper limits. This is due to the fact that the limited number of simulations in some cases results in the lowest minimum *p*-value being zero. When the minimum *p*-value in the data is zero, we show the percentage of simulations which also have zero as the minimum *p*-value. Clearly this fraction is only an upper limit on the real significance.

In order to further investigate the ℓ-dependence of the asymmetry, we follow two approaches from PCIS13. Firstly,
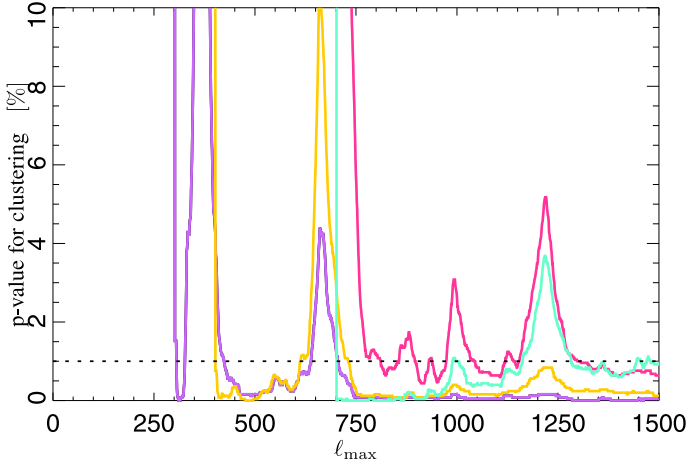
Fig. 37. Derived *p*-values for the angular clustering analysis as a function of $\ell_{\max}$, determined from SMICA based on 2500 simulations. The *p*-values are based on the fraction of simulations with a higher Rayleigh statistic up to the given $\ell_{\max}$ than in the data. The RS here is calculated over all pairs of dipole directions where one dipole in each pair is computed in the range $[\ell_{\lim}, \ell_{\max}]$, and the other is determined in the range $[2, \ell_{\lim}]$. The plot shows *p*-values for $\ell_{\lim} = 300$ (purple), $\ell_{\lim} = 400$ (yellow), $\ell_{\lim} = 500$ (pink), and $\ell_{\lim} = 700$ (cyan). The results have been marginalized over bin sizes in the range $\Delta\ell = 8$ to $\Delta\ell = 32$.

we restrict the analysis to multipoles above a minimum multipole $\ell_{\min}$. Table 26 indicates that clustering at the <1% significance level is still found when considering only those multipoles with $\ell_{\min}$ greater than 100. However, when this limit is increased to 200, no significant clustering is found. We then calculate the RS between pairs of dipoles where one dipole is determined from an ℓ-range above a certain limiting multipole $\ell_{\lim}$, and the other dipole below this limit. Figure 37 shows the RS as a function of $\ell_{\max}$ for some selected values of $\ell_{\lim}$. The $\ell_{\lim} = 300$ curve (purple) indicates that dipole directions for $\ell > 1000$ are significantly aligned with dipoles for $\ell < 300$. Similarly, the $\ell_{\lim} = 700$ curve (cyan) indicates that the dipole directions for $\ell = 700$–1000 are strongly correlated with the dipole directions for $\ell < 700$.

Combining these results, we note that when using only multipoles with (i) $\ell > 200$; or (ii) $\ell < 200$, no significant clustering is found. The strong clustering significance shown to persist to high multipoles in Fig. 36 must therefore be the result of clustering of the dipole directions between low and high multipoles as supported by Fig. 37. The low *p*-values can be explained by the alignment of dipole directions for multipoles extending all the way to $\ell = 1500$ correlated with directions for $\ell < 200$. The observed asymmetry is therefore not consistent with a model based on dipole modulation or power asymmetry located in one specific multipole range or for one given direction, but rather as a correlation of the dipole directions between $\ell < 200$ and $\ell > 200$. This correlation with lower multipoles is found to persist all the way to $\ell_{\max} = 1500$.

An advantage of the directional analysis performed here is that it focuses on a central issue for tests of deviation from isotropy – whether there is a preferred direction. Indeed, Bunn & Scott (2000) noted that the CMB may exhibit a pattern that cannot be identified from the power spectrum, but which would indicate some non-trivial large-scale structure. Evidence for the close correlation and alignment of directions on different angular scales may present a signature of broken statistical isotropy, since in the standard model, these directions should

all be independent random variables. In this context, we do not quote a specific direction for such asymmetry here since our results indicate a clustering of angles between different multipoles, but not necessarily that all multipoles are clustered about one specific direction. However, crucially we have shown that the measured clustering is driven by the correlations of directions between higher and lower multipoles.

Some of the analyses in other sections of the paper focus on dipolar modulation, a specific model for a dipolar power enhancement of the statistically isotropic CMB field towards a preferred direction of the sky, and use methods optimized for the detection of such a signal. While the results of Sect. 6.6 show no detection of the clustering of directions, there is no clear contradiction with the results presented here, since they are based on tests for $a_{\ell m}$ correlations between different multipoles as expected in the dipolar modulation model. The clustering analysis presented here is a model-independent test for deviations from statistical isotropy which could induce very different correlation structure. It is therefore sensitive to other forms of asymmetry, such as the addition of power in one part of the sky or more general phase correlations.

6.6. Rayleigh statistic: QML analysis

Results from Sect. 6.5 and in PCIS13 suggest that, beyond a dipole modulation of power on large angular scales, some form of directional asymmetry continues to small scales. There are also indications from Sect. 6.5 that the directions of dipolar asymmetry are correlated between large and small angular scales. Since the nature of the asymmetry is unknown we use the RS, a generic test for directionality that makes minimal assumptions about the nature of the asymmetry. This statistic has been used both in previous CMB studies (Stannard & Coles 2005) and other areas of cosmology (Scott 1991). In our context, for a statistically isotropic sky this statistic is identical to a three-dimensional random walk. The implementation here incorporates all information pertaining to modulation, not just the direction. The approach in this section differs from that of Sect. 6.5 in the method of reconstructing power, the choice of binning, and the choice of how to weight directions in each bin. Another important difference is that Sect. 6.5 only considers the direction of dipolar asymmetry and does not take into account its amplitude.

The statistic is cumulative and thus narrowing down the specific scales from which a signal may be originating is a non-trivial task. However, it is the case that all statistics that measure this form of asymmetry (dipole modulation or large-scale clustering of power) are in some way cumulative and so we will not worry about this issue any further. Another disadvantage of this approach is that it will generally be *less* powerful than a test that uses a specific model for the directionality. Again, this is a distinction shared when one compares any non-parametric versus parametric statistic.

The construction of the statistic is as follows.

1. Beginning with the estimator from Eqs. (44) and (45), compute the following binned quantities for the data and simulation:

$$\tilde{X}_{0,\,\ell} = \frac{6}{f_{10}} \frac{\sum_m A_{\ell m} \left(T^*_{\ell m} T_{\ell+1\,m} - \left\langle T^*_{\ell m} T_{\ell+1\,m} \right\rangle \right)}{\delta C_{\ell\ell+1} F_\ell F_{\ell+1} (\ell+1)}, \tag{60}$$

$$\tilde{X}_{1,\,\ell} = \frac{6}{f_{11}} \frac{\sum_m B_{\ell m} \left(T^*_{\ell m} T_{\ell+1\,m+1} - \left\langle T^*_{\ell m} T_{\ell+1\,m+1} \right\rangle \right)}{\delta C_{\ell\ell+1} F_\ell F_{\ell+1} (\ell+1)}. \tag{61}$$

For each ℓ this computes the coupling of ℓ to $\ell + 1$. We emphasize that this is a very natural choice of binning the estimator, since any parameter that is dipole modulated will lead to coupling of ℓ to $\ell \pm 1$ modes, albeit with different ℓ-weightings (below we describe why this is not an important issue).

2. Construct a three-dimensional vector out of the three estimators for both the data and the simulations[11], as defined by Eqs. (47)–(49).

3. Compute the mean amplitude from simulations and divide all vectors (data and simulations) by this amplitude. This choice ensures that each vector is treated equally, since we have no a priori reason to weight some scales more than others.

4. Add this new vector to the previous vector. If this is the first time going through this process the previous vector is the zero vector.

5. Repeat with $\ell \to \ell + 1$. Note that the statistics of this process are identical to a three dimensional random walk.

Given that a dipole modulation amplitude of roughly 3σ significance is known to exist at low ℓ (before a posteriori correction), one would expect a similar level of detection of asymmetry to be determined by the RS. Indeed, we find that asymmetry is present out to $\ell \approx 240$. Figure 38 (top) presents the p-values derived when the RS is computed as a function of ℓ_{max} from $\ell = 2$. The minimum p-value obtained by the data is 0.1–0.2%, to be compared to the value of 0.9–1.0% obtained for the dipole modulation amplitude at $\ell_{max} = 67$. The direction preferred by the data for $\ell_{max} \approx 240$ is $(l, b) = (208°, −29°)$, which is approximately $20°$ away from the dipole modulation direction determined to $\ell \approx 64$.

We correct for a posteriori statistics using the same procedure as in Sect. 6.3. Specifically, we count how often simulations find asymmetry in the range $10 \le \ell \le \ell_{max}$ that is more significant than that found for the data. From Fig. 39 it is clear that generic asymmetry at the significance level found in our CMB sky occurs about 6% or 8% of the time (depending on the range of ℓ one decides to search over).

While the PTE here is not very low, it is nevertheless somewhat lower than for the usual dipole modulation test. Hence, it seems worth exploring whether any of this signal comes from higher multipoles. Therefore we compute the RS starting at $\ell_{min} = 100$, to avoid the influence of asymmetry at lower ℓ. The lower panel of Fig. 38 presents the corresponding p-values as a function of ℓ_{max}. There is a striking similarity with the lower panel of Fig. 30. It is clear that, even in the absence of a posteriori correction, we find no significant asymmetry at larger ℓ. Hence most of the signal we are seeing in Fig. 38 (top) is due to the usual low-ℓ asymmetry.

We would like to stress that the results here are very similar to the results of the previous section. For each of the statistics used we are simply asking whether there is significant coupling of ℓ with $\ell \pm 1$ modes. The details of how to optimally combine these couplings for a given ℓ range depends on whether we are talking about dipole modulation or directionality (or some other related test, e.g., variance asymmetry). These details will change the range of scales over which the strongest signal in the data is found.

---

[11] Note that here we have not specified what $\delta C_{\ell\ell+1}$ is (it is fully specified by choosing a parameter X to modulate). This is because we have decided to weight each ℓ equally and thus any strictly positive choice for $\delta C_{\ell\ell+1}$ will be equivalent, since in step 3 we force the mean length of the vectors at each ℓ to be equal.
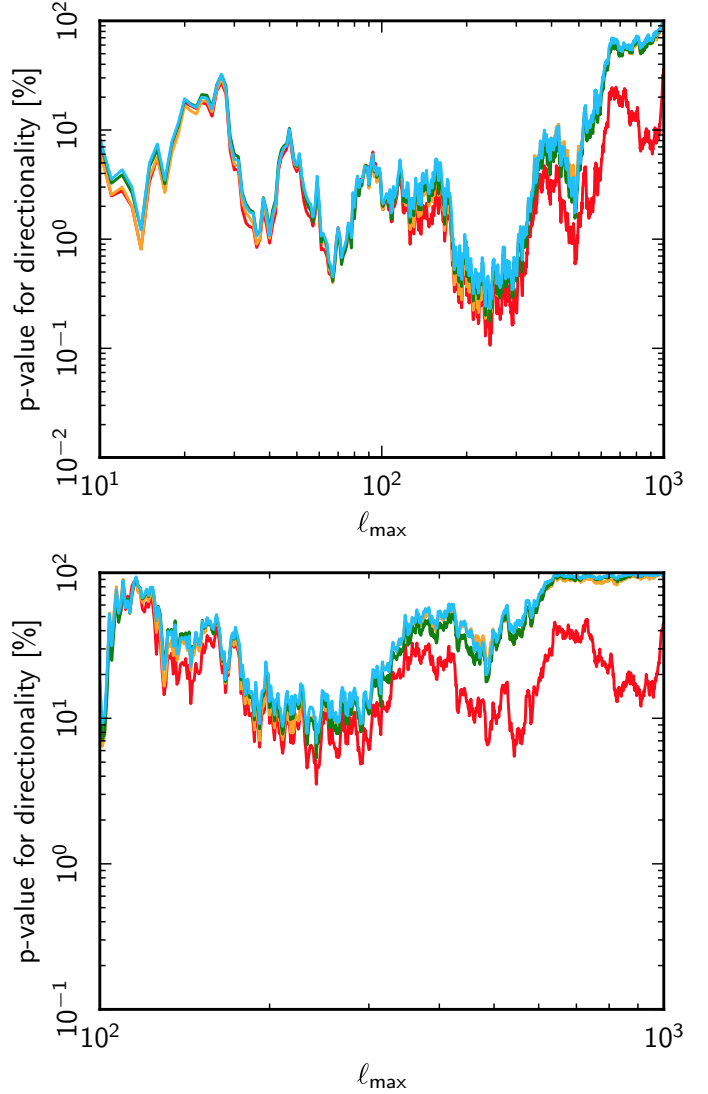

Fig. 38. Rayleigh statistic p-values determined from the QML analysis as a function of ℓ_{max} for the `Commander` (red), `NILC` (orange), `SEVEM` (green), and `SMICA` (blue) data sets, with (*top panel*) $\ell_{min} = 2$ and (*bottom panel*) $\ell_{min} = 100$. The general pattern of peaks is very similar to that in Fig. 30. We emphasize that the statistic here is cumulative and as such trends in the curves can be misleading.

7. Sensitivity of anomalies to enhanced sky coverage

One of the critical aspects in searching for anomalous features in sky maps is to ensure that the region being investigated constitutes a fair and unbiased sample. Since many of the claimed anomalies are on large angular scales, this implies that minimal masking should be applied to the data. However, residual foregrounds then become a significant consideration. The masks applied to the four component-separated maps studied in the bulk of this paper have been defined at high resolution, and then conservatively degraded for lower resolution studies. Such a procedure inevitably reduces the sky coverage available for analysis, and can be particularly problematic if significant structures are aligned by chance with the masked regions. Indeed, the WMAP team (Bennett et al. 2011) have drawn attention to several such features in their ILC reconstruction of the CMB sky, and these are clearly also present in the *Planck* `Commander`, `NILC`, `SEVEM`, and `SMICA` sky maps. A large cold spot is seen
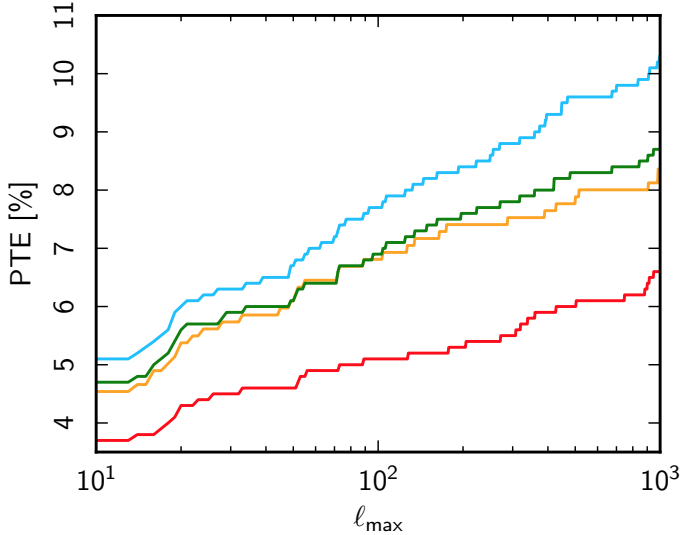
Fig. 39. Probability to exceed (PTE) the *p*-value of the signal from the `Commander` (red), `NILC` (orange), `SEVEM` (green), and `SMICA` (blue) data at ℓ = 230–240 (which is the multipole range with the most significant deviation) when searching over a range of multipoles up to ℓ_{max}, for the RS determined from the QML analysis. Much like the equivalent curve for dipole modulation, the PTE appears to grow approximately logarithmically with ℓ_{max}.

near to the Galactic centre, a significant fraction of which lies within the common mask at any resolution. However, despite its location and visual impression, the feature is neither likely to be attributable to residual foreground emission, nor is it inconsistent with the ΛCDM model (Gott et al. 2007). In addition, four elongated cold fingers stretching from near the Galactic equator to the south Galactic pole are seen, although no equivalent features are evident in the northern sky. Bennett et al. (2011) have noted that the alignment of the ℓ = 2 and ℓ = 3 multipoles (Tegmark et al. 2003) seems to be intimately connected with these large-scale cool fingers and the intervening warm regions. One of the latter also corresponds to the well-known "Bianchi VII$_h$" main lobe originally found in Jaffe et al. (2005).

Although we would ideally pursue full sky analyses, we prefer to remain mindful of the influence of residual foregrounds, but still seek to minimize the extent of any mask applied for analysis. In this context, and specifically for large-angular-scale studies, we consider the properties of an additional estimate of the CMB sky, also generated using the `Commander` component separation methodology. In particular, we note that the *Planck* low-ℓ likelihood analysis (Planck Collaboration XI 2016) uses the temperature solution from this study, degraded to a resolution of N_{side} = 16. The `Lkl-Commander` map, as we now refer to it, is initially derived from input data sets (32 bands) at 1° FWHM resolution and N_{side} = 256. This includes *Planck* individual detector and detector set maps from 30–857 GHz, the 9-year WMAP observations between 23 and 94 GHz, and the 408 MHz sky survey (Haslam et al. 1982), whereas the `Commander` map described in Planck Collaboration IX (2016) includes *Planck* data alone. It is believed that the 32-band solution is better (on large angular scales) than the *Planck*-only map, because the larger number of input frequencies allows more detailed foreground modelling, and in particular the separation of the low-frequency foregrounds into synchrotron, free-free, and spinning dust components. An associated confidence mask (hereafter $\text{LklT}_{256}93$) is then defined based on a goodness-of-fit measure per pixel,

Table 27. Lower-tail probability for the variance, skewness, and kurtosis of the `Lkl-Commander` map.

	Probability [%]		
Mask	Variance	Skewness	Kurtosis
$\text{LklT}_{16}94$	7.0	1.5	94.0
$\text{UT}_{16}58$	0.7	19.9	82.5

corresponding to a rejection of 7.3% of the pixels on the sky. A detailed discussion of these results can be found in Planck Collaboration X (2016).

We now consider the implications of using the `Lkl-Commander` map for studies of several large-angular-scale anomalies observed in previous sections, in particular since the larger sky coverage permitted by this data set should constitute a better sample of the Universe. Note that, at the resolutions of interest for the following analyses, the noise level is negligible (even accounting for the WMAP contribution) and should not have significant impact on the results. The exact details of the noise contribution to simulations is therefore unimportant.

7.1. Variance, skewness, and kurtosis

We begin by estimating the variance, skewness, and kurtosis of the CMB. We apply the unit variance estimator to the `Lkl-Commander` map, and specifically to the version used in the low-ℓ likelihood analysis, which is smoothed to 440' FWHM at a resolution of N_{side} = 16. A corresponding low-ℓ mask is generated by a simple degrading of the mask at N_{side} = 256, then setting those N_{side} = 16 pixels with a value less than 0.5 to zero and all others to unity. The resultant low-ℓ likelihood mask rejects only 6.4% of the sky. We compare the results for both this mask (also to be referred to as $\text{LklT}_{16}94$), and the standard common mask at this resolution ($\text{UT}_{16}58$). The results are summarized in Table 27 and show that, when using the low-ℓ likelihood mask, the lower tail probability for the variance is 7.0%. This value is higher than the corresponding values for the component separated maps as shown in Table 12. In addition the skewness and kurtosis are less consistent with Gaussianity than the component separated maps. However, when using the standard common mask at N_{side} = 16, the lower tail probability of the variance, skewness, and kurtosis become more compatible with those derived earlier.

There are two possible explanations for this behaviour. Either the variance of the CMB in the region close to the Galactic plane is intrinsically high, perhaps due to the presence of the various features noted above, or the presence of residual foregrounds increases the variance of the map. In order to attempt to distinguish between these options, we again apply the unit variance estimator to the standard component-separated maps[12], but this time utilising the low-ℓ mask. Although the component-separated maps are likely to contain

---

[12] Note that the `SEVEM` maps used in this section have been inpainted within 3% of the sky towards the Galactic centre using a simple diffusive inpainting algorithm. This prevents residual foreground contamination from propagating to neighbouring regions when downgrading the map. The other component-separated maps are not pre-processed in this way since some form of inpainting of the most contaminated regions was already implemented as part of the component separation algorithms.

Table 28. Lower-tail probability for the variance, skewness, and kurtosis of the Lkl-Commander map compared to the component separated maps, obtained using the low-ℓ likelihood mask LklT$_{16}$94.

	Probability [%]		
Map	Variance	Skewness	Kurtosis
Lkl-Commander	7.0	1.5	94.0
Commander	7.7	1.9	96.0
NILC	9.6	5.0	94.4
SEVEM	7.4	4.8	94.3
SMICA	7.7	3.7	93.7
SEVEM-100	8.4	0.4	97.9
SEVEM-143	7.7	3.7	95.5
SEVEM-217	8.3	0.7	95.2

Table 29. Probabilities to exceed the observed values of the χ^2 statistics for the Lkl-Commander and Commander maps at $N_{\mathrm{side}} = 64$.

	Probability [%]		
Function	Lkl-Comm.[a]	Lkl-Comm.[b]	Comm.[b]
2-pt.	84.3	97.1	97.2
Pseudo-coll. 3-pt.	76.8	92.1	92.1
Equil. 3-pt.	96.5	74.0	74.0
Rhombic 4-pt.	94.5	65.0	64.6

Notes. [a] low-ℓ mask, LklT$_{64}$92. [b] Common mask, UT$_{64}$67.

some foreground contamination in the regions omitted by application of the UT$_{16}$58 mask, it is appropriate to recall that this was constructed in a conservative way, and may also mask parts of the sky where the level of residual foregrounds can be considered negligible. In addition, we investigate the cleaned frequency maps produced by the SEVEM algorithm in order to test for the presence of frequency-dependent residual foregrounds. The results of the unit variance estimator analysis are summarized in Table 28.

All of the component separated maps show an increase in the lower tail probability from about 0.5% when the UT$_{16}$58 mask is applied to roughly 7% for the LklT$_{16}$94 mask. The small variations in results for the different maps may be attributable to the presence of residual foregrounds close to the Galactic plane. However, the increased probabilities can also be explained by the presence of CMB structures with higher variance within that region which is not rejected by the less conservative mask. Indeed, since the component-separated maps are affected by different residual foregrounds, if the source of the changes in probabilities is due only to the residual foregrounds, then we would expect a larger dispersion than what is observed. We also note that when we apply the low-ℓ likelihood mask the skewness and kurtosis values are shifted towards more extreme values. This implies that the sky signal is less Gaussian for the larger sky fraction, despite the results remaining compatible with the ΛCDM model assumed for the null tests. Both Commander maps are noteworthy in this regard.

An important issue is whether the changes in the statistics can simply be attributed to differences in the masks. We determine how many simulations show an increase in variance at least as large as that seen for the Lkl-Commander map when comparing the values derived for the UT$_{16}$58 and low-ℓ likelihood masks. Similarly, we determine how many simulations have increased skewness or kurtosis values with shifts at least as large as observed. When the three statistics are considered separately, the fraction of simulations that indicate such changes are 7.6%, 4.3%, and 13.9% for the variance, skewness, and kurtosis, respectively. Of course, such subsets of the simulations also include cases where a large shift in the statistic is observed, but the statistic would not be considered anomalous for either mask. If we also impose the requirement that the simulations have these shifts for all three quantities simultaneously, then only 2 maps from 1000 are found. Of course, such a requirement is rather strong, and at this stage we are likely to be approaching the limits of what can be said based on model-independent null tests. Indeed, in order to assess whether these results are sensitive to

a posteriori choices, we repeat the analysis but successively take each simulation as the reference. Thus, for each simulation the shift in the variance, skewness, and kurtosis is computed and then we determine how many times we find a case in which two or less of the remaining simulations simultaneously show larger shifts for the three moments. We find that 48 maps from 1000 satisfy these conditions. Given this, it is difficult to draw strong conclusions about the significance or otherwise of the mask-related changes in variance.

7.2. N-point correlation functions

The connection between sky coverage and the observed structure of the 2-point correlation function for large angular separations has previously been discussed in the literature, in particular in connection with the $S_{1/2}$ statistic discussed in Sect. 5.2. Bennett et al. (2011) consider that the use of a Galactic mask when computing these quantities is sub-optimal, and note that a full-sky computation of the 2-point correlation function from the 7-year WMAP ILC map lies within the 95% confidence region determined by simulations of their best-fit ΛCDM model over all angular separations. However, Copi et al. (2009) suggest that the origin of the inconsistencies between the full-sky and cut-sky large-scale angular correlations remains unknown, and that the observed discrepancies may indicate that the Universe is not statistically isotropic on these scales. We therefore consider the N-point correlation functions, and related statistics, of the Lkl-Commander map to contribute to this debate.

We compare results computed for both the Lkl-Commander and Commander maps at $N_{\mathrm{side}} = 64$ after smoothing to a FWHM of 160′. A mask is constructed for the Lkl-Commander map by degrading the LklT$_{256}$93 mask to $N_{\mathrm{side}} = 64$ and setting all resulting pixels with a value less than 0.5 to zero, with the remainder set to unity. The LklT$_{64}$92 mask retains 92% of the sky, to be compared to the 67% usable sky coverage allowed by the UT$_{64}$67 common mask at this resolution.

The results are presented in Fig. 40 where we compare the N-point functions for the data and the mean values estimated from 1000 Commander simulations. The probabilities for obtaining values of the χ^2 statistic for the *Planck* fiducial ΛCDM model at least as large as the observed values are provided in Table 29. For the estimation of the probabilities, we use the same set of 1000 Commander simulations for both versions of the Commander data. As noted previously, the details of the simulations for such highly smoothed data is essentially unimportant. We also provide an analysis of the Lkl-Commander map using the common mask to enable a direct comparison with the analysis of the Commander map. In this latter case, the results for both maps are in excellent agreement. However, the Lkl-Commander map is more consistent with simulations when the LklT$_{64}$92
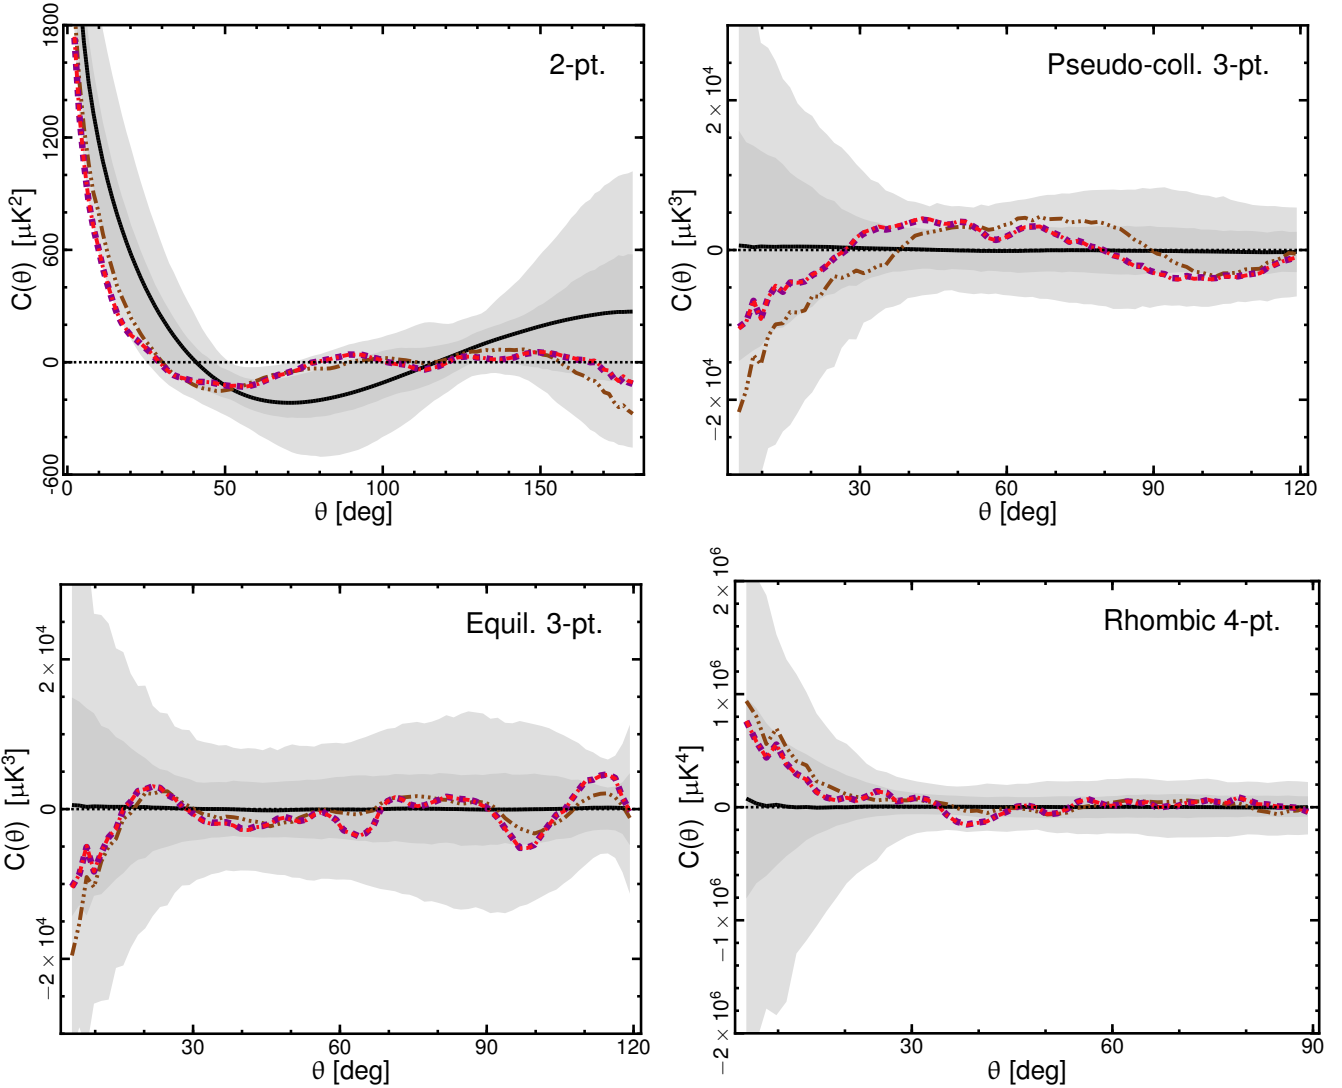
Fig. 40. N-point correlation functions determined from the $N_{side} = 64$ *Planck* CMB 2015 temperature estimates. Results are shown for the 2-point, pseudo-collapsed 3-point (*upper left and right panels*, respectively), equilateral 3-point, and connected rhombic 4-point functions (*lower left and right panels*, respectively). The brown three dot-dashed, purple dashed, and red dot-dashed lines correspond to the Lkl-Commander map analysed using the low-ℓ and common masks and the Commander map analysed using the common mask, respectively. Note that the dashed and dot-dashed lines lie on top of each other. The black solid line indicates the mean for 1000 MC simulations. The shaded dark and light grey regions indicate the 68% and 95% confidence regions, respectively, estimated using 1000 Commander simulations. See Sect. 4.3 for the definition of the separation angle θ.

mask is adopted for the 2-point and pseudo-collapsed 3-point functions, but less consistent for the equilateral 3-point and rhombic 4-point function results. Nevertheless, the results are generally in agreement with expectations for a Gaussian, statistically isotropic model of the CMB fluctuations.

The increased consistency of the 2-point function with simulations when analysing a larger sky fraction is consistent with the observations in Copi et al. (2009). We therefore quantify this further by determining the statistical quantities introduced in Sect. 5.2 for the Lkl-Commander map. In particular, we re-assess the lack of correlation determined previously for large angular scales. It is evident from Table 30 that the results for the $S_{1/2}$ and χ_0^2 statistics are less anomalous when the low-ℓ mask is applied. Moreover, the global p-value for the $S(x)$ statistic is substantially smaller.

We also repeat the conventional χ^2 analysis but constraining the computations to the two separate ranges defined by

Table 30. Probabilities for obtaining values of the $S_{1/2}$ and χ_0^2 statistics for the simulations at least as large as the observed values of the statistic estimated from the Lkl-Commander and Commander maps using the LklT$_{64}$92 and UT$_{64}$67 masks, respectively. We also show the corresponding estimation of the global p-value for the $S(x)$ statistic.

	Probability [%]	
Statistic	Lkl-Comm.	Comm.
$S_{1/2}$.....................	97.1	99.5
$S(x)$ (global)	90.9	97.7
$\chi_0^2(\theta > 60°)$	95.7	98.1

$\theta < 60°$ and $\theta > 60°$. The results of these studies are shown in Table 31. The analysis for separation angles $\theta > 60°$ indicates that the unusually good fit of the observed 2-point function to the mean 2-point function determined for the ΛCDM model is

Table 31. Probabilities for obtaining values of the χ^2 statistic for the simulations at least as large as the observed values of the statistic estimated from the Lkl-Commander and Commander maps using the $LklT_{64}92$ and $UT_{64}67$ masks, respectively.

Statistic	Probability [%]	
	Lkl-Comm.	Comm.
$\chi^2(\theta < 60°)$	52.9	91.5
$\chi^2(\theta > 60°)$	96.5	96.8

Table 32. Probabilities for obtaining values of the χ^2 statistic and ratio of χ^2 of the N-point functions for the *Planck* fiducial ΛCDM model at least as large as the observed values of the statistic on the northern and southern ecliptic hemispheres estimated from the Lkl-Commander and Commander maps using the $LklT_{64}92$ and $UT_{64}67$ masks, respectively.

Hemisphere	Probability [%]	
	Lkl-Comm.	Comm.
2-point function		
Northern	98.7	89.7
Southern	29.4	80.5
χ^2-ratio	93.3	22.6
Pseudo-collapsed 3-point function		
Northern	99.4	>99.9
Southern	14.9	35.1
χ^2-ratio	98.6	98.8
Equilateral 3-point function		
Northern	99.9	98.6
Southern	47.6	45.7
χ^2-ratio	94.4	86.6
Rhombic 4-point function		
Northern	99.8	99.7
Southern	24.5	22.8
χ^2-ratio	97.8	97.3

independent of the mask used in the analysis. Conversely, the results for the angles $\theta < 60°$ indicate a strong dependence on the mask. It appears that the decreased significance of the χ^2 statistic for the 2-point function of the Lkl-Commander map reported in Table 29 is related mainly to correlations in the data for separation angles smaller than 60°.

Our results do appear to indicate that computations made on larger sky fractions increase the consistency of the 2-point function with simulations. We therefore also test how the hemispherical asymmetry observed previously is affected. The results for the ecliptic frame are presented in Table 32. We find that the asymmetry is larger for the Lkl-Commander map than for the Commander map in the case of the 2-point function, but does not change substantially for the 3-point and 4-point functions.

7.3. Dipole modulation and directionality

7.3.1. Variance asymmetry

Here we apply the local-variance analysis of Sect. 6.1 to the Lkl-Commander map and compare the results with those of the Commander map. Contrary to the analysis of Sect. 6.1, where full-resolution ($N_{side} = 2048$) maps were used, here the Commander map is downgraded to $N_{side} = 256$ in order to consistently compare the results for both maps. The simulations used for estimating the significance levels are also downgraded to the same resolution, and convolved with the corresponding beam function. Otherwise, the procedure is identical to the one

Table 33. p-values for the variance asymmetry measured by different discs from the *Planck* 2015 Lkl-Commander and Commander temperature solutions using the $LklT_{256}93$ and $UT_{256}73$ masks, respectively.

Disc radius [°]	p-value [%]	
	Lkl-Comm.	Comm.
4	0.2	0.7
6	0.4	0.3
8	0.3	0.1
10	0.3	0.7
12	0.3	0.8
14	0.5	1.5
16	0.6	1.9
18	0.8	2.7
20	1.1	3.7

Notes. The values represent the fraction of simulations with local-variance dipole amplitudes larger than those inferred from the data. No high-pass filtering has been applied to the maps.

Table 34. Local-variance dipole directions measured by 8° discs for the *Planck* 2015 Lkl-Commander and Commander temperature solutions.

Method	(l, b) [°]
Lkl-Comm.	(225, −28)
Commander	(214, −24)

described in Sect. 6.1, e.g., the same number of discs has been used to construct the local-variance maps. Here we only present the results when no high-pass filtering has been applied to the maps; this is to avoid confusion as our objective in this section is only to compare the general properties of the Lkl-Commander map to those of the standard component-separated maps.

Table 33 summarizes the significance levels measured by our variance asymmetry analysis using discs of different radii, for the *Planck* 2015 Commander and Lkl-Commander temperature maps. The p-values represent the fraction of simulations with local-variance dipole amplitudes larger than those inferred from the data. We in addition present in Table 34 the preferred variance asymmetry directions for both maps using 8° discs.

Our results show consistency between the two maps. The small change in the preferred direction is expected from the change in the mask, and agrees specifically with the directions found by the analysis of the QML dipole modulation analysis in Sect. 7.3.3. One interesting observation is that the large variance asymmetry significance is now extended to cases where larger discs are used. Note that no high-pass filtering has been applied in the present analysis, and therefore p-values inferred from the Commander map increase with the disc size. As explained in Sect. 6.1, the low observed significance levels for larger discs is due to the cosmic variance associated with the largest-scale modes. The observed increase in the significance levels for the Lkl-Commander map is therefore interestingly consistent with this picture; the mask in this case is smaller and therefore a larger fraction of the sky is available. This in turn provides more data on the largest scales, and therefore lowers the impact of the cosmic variance.

7.3.2. Dipole modulation: pixel-based likelihood

Table 35 presents constraints on the dipole modulation model as derived from the Lkl-Commander map and the $LklT_{32}93$ mask that includes 93% of the sky, updating the results from Sect. 6.2

Table 35. Summary of dipole modulation results at a smoothing scale of 5° for the *Planck* 2015 `Lkl-Commander` and `Commander` temperature solutions, as derived by the brute-force likelihood given by Eq. (43).

Method	2013	2015
Dipole modulation amplitude, α		
`Lkl-Comm.`	$\cdots$	0.059 ± 0.020
`Commander`	0.078 ± 0.021	0.066 ± 0.021
Dipole modulation direction, (l, b) [°]		
`Lkl-Comm.`	$\cdots$	$(223, -17) \pm 23$
`Commander`	$(227, -15) \pm 19$	$(230, -16) \pm 24$
Power spectrum amplitude, q		
`Lkl-Comm.`	$\cdots$	0.970 ± 0.025
`Commander`	$\cdots$	0.961 ± 0.025
Power spectrum tilt, n		
`Lkl-Comm.`	$\cdots$	0.068 ± 0.045
`Commander`	$\cdots$	0.082 ± 0.043

Notes. The former results were derived using the $LklT_{32}93$ mask, whereas the latter are those determined previously in Sect. 6.2.

Table 36. Summary of the dipole modulation results for the range $\ell \in [2, 64]$ determined from the *Planck* 2015 `Lkl-Commander` and `Commander` temperature solutions, as derived by the QML estimator defined in Sect. 6.3 using the $LklT_{256}93$ and UT78 masks, respectively.

Method	A	Direction (l, b) [°]
`Lkl-Comm.`	$0.058^{+0.022}_{-0.012}$	$(227, -28) \pm 26$
`Commander`	$0.063^{+0.025}_{-0.013}$	$(213, -26) \pm 28$

for the `Commander` map. We find that all previously reported results are robust with respect to data selection and sky coverage. In particular, the best-fit dipole modulation amplitude at 5° FWHM is 5.9% in the `Lkl-Commander` map, and is thus stable to within about 0.3σ when increasing the sky fraction from 78% to 93%. Likewise, the marginal low-ℓ power spectrum amplitude, q, shifts upward by 0.4σ, and the power spectrum tilt, n, downward by 0.3σ, for the same sky fraction increases.

To assess the statistical significance of these shifts, we compare with Gaussian statistics, creating two Gaussian random vectors with 78 and 93 elements, respectively, where the first 78 elements of the latter vector are identical to the first vector. From these, we compute the difference between the two means, after normalizing each so that their individual errors in the mean are unity. Repeating this simple calculation 10^5 times, we find that 48% of all Gaussian realizations observe shifts larger than 0.3σ, and 34% observe shifts larger than 0.4σ. Thus, the parameter differences due to the different data selection and sky fractions reported above are consistent with expectations from random Gaussian statistics.

7.3.3. Dipole modulation: QML analysis

We also repeat the QML dipole modulation analysis of Sect. 6.3 for the `Lkl-Commander` map and corresponding mask. Table 36 summarizes the results of the low-ℓ dipole modulation for the `Lkl-Commander` temperature solution, compared with the `Commander` map.

The best-fit modulation amplitude for `Lkl-Commander` is 5.8% and the small 0.5% shift from the `Commander` best-fit amplitude corresponds to a decrease of approximately 0.4σ. These

Table 37. Amplitude (A) and direction of the dipole modulation in Galactic coordinates as estimated for the multipole range $\ell \in [2, 64]$ using the BipoSH analysis on `Lkl-Commander` and `Commander` maps. The former results were derived using the $LklT_{256}93$ mask; the latter are those determined previously in Sect. 6.4.

Method	A	Direction (l, b) [°]
`Lkl-Comm.` . .	0.063 ± 0.021	$(234, -27) \pm 31$
`Commander` . .	0.067 ± 0.023	$(230, -18) \pm 31$

results mirror very closely the results found above for the pixel-based likelihood approach to dipole modulation, as expected, and the observed shifts are perfectly consistent with those expected from the change in the mask.

7.3.4. Bipolar spherical harmonics

We next perform a dipole modulation analysis on the `Lkl-Commander` temperature map using the BipoSH formalism from Sect. 6.4. The dipole modulation amplitude inferred from the analysis is smaller that that deduced from analysing the `Commander` map as seen in Table 37. However, it should be noted that the probability for simulations to yield a dipole modulation amplitude equal to or greater than the amplitude inferred from data is 0.4%, which is smaller by a factor of approximately 2.4 as compared to the p-value inferred from analysis on `Commander`. The reduction in the dipole amplitude and the enhanced significance can both be attributed to the reduced power bias which is a result of the increased sky coverage.

7.4. Summary

Using a larger sky fraction in our analyses leads to small changes in the results related to large-angular-scale anomalies, but these are essentially consistent with expectations from random Gaussian statistics. In particular, the asymmetry in power on the sky, as parameterized by a dipole modulation model, is robust to mask changes.

8. Polarization analysis

As previously discussed in Sect. 2, large angular-scale CMB fluctuations in the *Planck* polarization data have been suppressed by a post-processing high-pass filter to minimize the impact of systematic artefacts. Therefore, no polarization results concerning CMB statistical anomalies on such scales are presented in this paper. In addition, a noise mismatch between simulations and data also limits our ability to study polarization more generally. Nevertheless, a local analysis of the polarization data for stacked patches of the sky can still be performed, in order to test the statistical properties of the CMB anisotropies. In this case, the stacking procedure mitigates the impact of the small-scale noise and potential systematic effects.

Traditionally, the Stokes parameters Q and U are used to describe the CMB polarization anisotropies (e.g., Zaldarriaga & Seljak 1997). Such quantities are not rotationally invariant, thus for the stacking analysis it is convenient to consider a local rotation of the Stokes parameters, resulting in quantities denoted by Q_r and U_r, as described in Sect. 8.1. Additionally, several other related quantities can be defined.

The polarization amplitude $P \equiv \sqrt{Q^2 + U^2}$ and polarization angle $\Psi \equiv \frac{1}{2}\arctan(U/Q)$, are commonly used quantities in, for example, Galactic astrophysics. However, unbiased estimators of these quantities in the presence of anisotropic and/or correlated noise are hard to define (Plaszczynski et al. 2014). Of course, a direct comparison of the observed (noise-biased) quantity to simulations analysed in the same manner is possible, but we elect here to defer the study of this representation of the polarization signal, using maps of the polarization amplitude only to define peaks around which stacking can be applied.

The rotationally invariant quantities referred to as E and B modes are commonly used for the global analysis of CMB data. Although the E-mode maps are not analysed in detail here, they are considered qualitatively, so that it is appropriate to recall their construction. Since the quantities $Q \pm iU$, defined relative to the direction vectors $\hat{n}$, transform as spin-2 variables under rotations around the $\hat{n}$ axis, they can be expanded as

$$(Q \pm iU)(\hat{n}) = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}^{(\pm 2)} \,_{\pm 2}Y_{\ell m}(\hat{n}), \tag{62}$$

where $_{\pm 2}Y_{\ell m}(\hat{n})$ are the spin-weighted spherical harmonics and $a_{\ell m}^{(\pm 2)}$ are the corresponding harmonic coefficients. If we define

$$a_{\ell m}^E = \frac{1}{2}\left(a_{\ell m}^{(2)} + a_{\ell m}^{(-2)}\right), \tag{63}$$

$$a_{\ell m}^B = \frac{-i}{2}\left(a_{\ell m}^{(2)} - a_{\ell m}^{(-2)}\right), \tag{64}$$

then the invariant quantities are given by

$$E(\hat{n}) = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}^E Y_{\ell m}(\hat{n}), \tag{65}$$

$$B(\hat{n}) = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}^B Y_{\ell m}(\hat{n}). \tag{66}$$

8.1. Stacking around temperature hot and cold spots

The stacking of CMB anisotropies around peaks (hot and cold spots) on the sky yields characteristic temperature and polarization patterns that contain valuable information about the physics of recombination (Komatsu et al. 2011). Statistical analysis of stacked images differs from the other tests in this paper in several respects. First, peak-related new physics may be revealed that is difficult to find in a global analysis, for example, the non-Gaussian CMB cold spots predicted by a modulated preheating model (Bond et al. 2009). Secondly, stacking is a local operation, which naturally avoids mask-induced complications. Thus stacking can be used as a transparent and intuitive method to test the robustness of anomalies found with other methods. Alternatively, it can be applied as a quality indicator of the data at the map level.

Our stacking procedure is as follows. Hot (or cold) peaks are selected in the temperature map as local extrema with negative (or positive) second derivatives, and classified relative to a given threshold ν (in rms units of the temperature map). Since the spinorial components Q and U are expressed in a local coordinate system, we employ a configuration in which the Stokes parameters around a peak at the direction $\hat{n}_0$ can be superposed (Kamionkowski et al. 1997). In particular, we use a locally defined rotation of the Stokes parameters that is written as:

$$Q_r(\hat{n}; \hat{n}_0) = -Q(\hat{n})\cos(2\phi) - U(\hat{n})\sin(2\phi), \tag{67}$$

$$U_r(\hat{n}; \hat{n}_0) = Q(\hat{n})\sin(2\phi) - U(\hat{n})\cos(2\phi), \tag{68}$$

where ϕ is the angle between the axis aligned along a meridian (pointing to the south by convention) in the local coordinate system centred on a peak at $\hat{n}_0$ and the great circle connecting this peak to a position $\hat{n}$. This definition decomposes the linear polarization into radial ($Q_r > 0$) and tangential ($Q_r < 0$) contributions around the peaks. This definition of Q_r is equivalent to the "tangential shear" used in weak lensing studies.

For visualization purposes, a flat patch around each peak is then extracted, and the average stacked image computed from the subset. A position on the sky at an angular distance θ from the central peak is labelled with the flat-sky coordinates

$$x = \varpi\cos\phi, \quad y = \varpi\sin\phi. \tag{69}$$

Here $\varpi = 2\sin(\theta/2) \approx \theta$ is the effective flat-sky radius. For the angular scales of a few degrees considered in the stacking analyses the difference between ϖ and θ is negligible. We use ϖ for analyses in the flat-sky approximation, and θ for analyses directly on the sphere.

The stacking process tends to provide an image with azimuthal symmetry about its centre, due to the almost uncorrelated orientations of the temperature peaks. The stacked images of temperature patches around hot spots selected above the null threshold for both the Commander data and a corresponding simulation are shown in the top row of Fig. 41. The observed patterns are in excellent agreement. Stacking around cold spots yields similar patterns but with flipped sign. Given the symmetry, it is often useful to consider the radial profile obtained by averaging the stacked image over the azimuthal angle ϕ. Figure 42 shows such a profile determined from the stacked temperature image.

At this point, it is useful to consider the underlying physics represented by the various patterns in the stacked images. During recombination, the sound horizon extends an angle $\theta_s = r_s/D_A \approx 0.011$ ($0.61°$), where $r_s \approx 0.15$ Gpc is the size of the sound horizon at recombination and $D_A \approx 14$ Gpc is the angular-diameter distance to the last scattering surface. To understand the ring patterns in the stacking image, projection effects must be taken into account. Firstly, all 3D modes with wavenumber $k \geq \ell/D_A$ contribute to a 2D ℓ-mode. More modes contribute to, and therefore enhance, the power at lower ℓ. For the first acoustic peak, the net effect is a $\pi/4$ phase shift towards lower ℓ, such that $\ell_s \approx (\pi - \pi/4)/\theta_s \approx 220$. The projected acoustic scale on the temperature map is of order $\theta_s^{2D} = \pi/\ell_s = 0.014$ ($0.81°$). Secondly, the stacked 2D modes around peaks interfere with each other. The first dark ring appears at $1.22\theta_s^{2D} \approx 0.017$ ($1.0°$). The factor 1.22 is the ratio of the first minimum of the projection kernel, the Bessel function J_0, to the first minimum of the unprojected cosine wave.

The dark ring can also be regarded as a consequence of the correlation between T and $-\nabla^2 T$. At the temperature maxima $-\nabla^2 T$ is positive, with an amplitude of order $T_{\text{peak}}/(\theta_s^{2D})^2$. Thus, the quadratic terms in the local expansion of the temperature field have a negative contribution that grows as $-T_{\text{peak}}(\varpi/\theta_s^{2D})^2$. At $\varpi \gtrsim \theta_s^{2D}$ the quadratic terms dominate and the T-$(-\nabla^2 T)$ correlation becomes negative. Meanwhile, the T-$(-\nabla^2 T)$ correlation tends to zero on the scale $\varpi \gtrsim \theta_s^{2D}$, where the temperature autocorrelation becomes weak and the local quadratic expansion starts to fail. As shown in Fig. 42, the dark ring appears at the critical point where the T-$(-\nabla^2 T)$ correlation reaches its minimum and turns back toward zero.

We have discussed the projection effects that make the projected radial acoustic scale on a stacked T image larger than θ_s. For Q_r, the most striking patterns in the image have
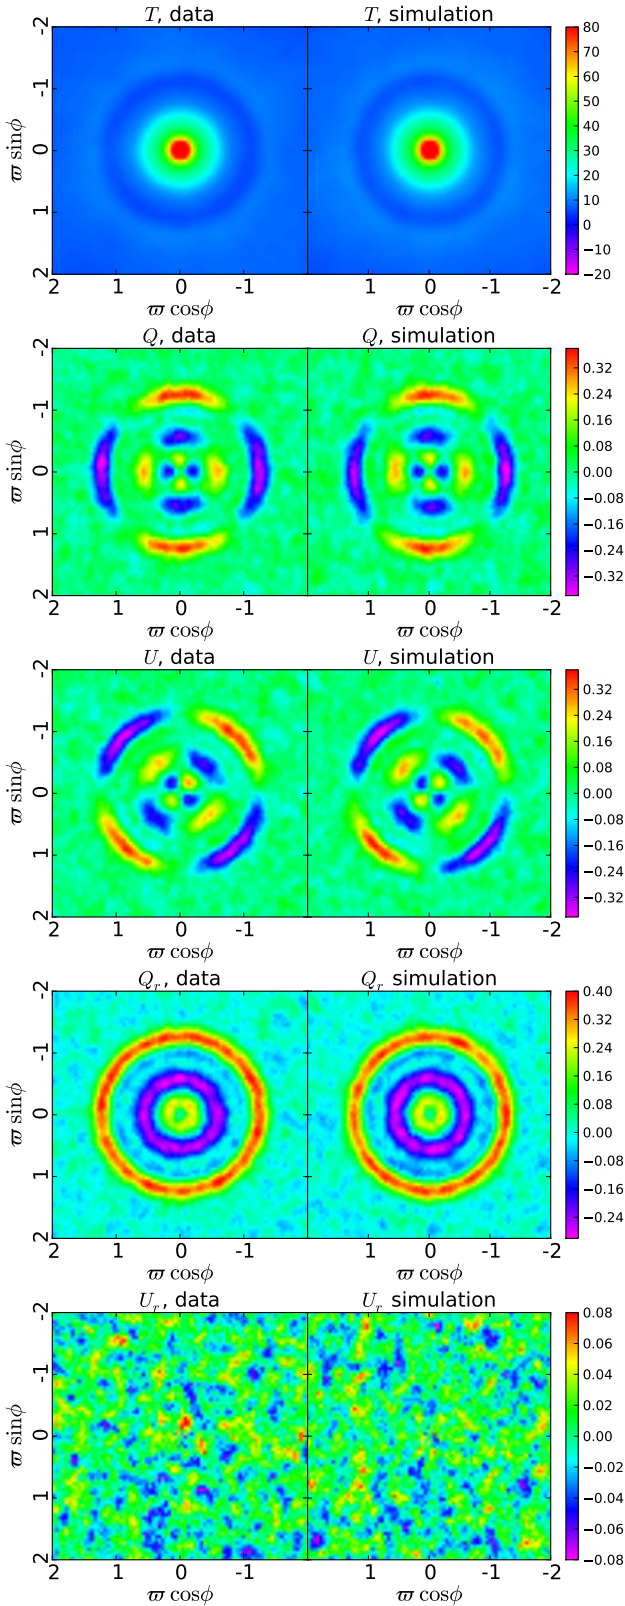
Fig. 41. From *top to bottom*, T, Q, U, Q_r, and U_r stacked images (in μK units) extracted around temperature hot spots selected above the null threshold ($\nu = 0$) in the `Commander` sky map for data (*left column*) and an equivalent simulation (*right column*). The horizontal and vertical axes of the flat-sky projection are labelled in degrees.

Fig. 42. Radial profile $\mu_T(\varpi)$ derived from the stacked temperature image (see Fig. 41 or 45). The denominators σ_0 and σ_2 are the theoretical rms values of CMB T and $\nabla^2 T$, respectively. The theoretical $\langle \mu_T(\varpi) \rangle$ is a linear combination of $\langle T(\varpi)(T(0)/\sigma_0) \rangle$ (green dash-dotted line) and $\langle T(\varpi)(-\nabla^2 T(0))/\sigma_2) \rangle$ (blue dotted line). For all four component-separated maps, the deviation of μ_T from the ensemble mean $\langle \mu_T \rangle$ of the fiducial model (here the *Planck* 2015 ΛCDM best fit) is consistent with cosmic variance, and can be related to the low-ℓ power deficit. The example power-deficit $\langle \mu_T \rangle$ (purple dashed line) is the theoretical prediction of $\langle \mu_T \rangle$ if the fiducial model C_ℓs are reduced by 10% in the range $2 \le \ell \le 50$.

more intuitive simple explanations, since the stacking is essentially the real-space equivalent of the temperature polarization

correlation. The projection function contains an extra ℓ^2 factor, which enhances the high-ℓ power and reduces the projected radial acoustic scale, coincidentally, back to $\approx \theta_s$. The quadrupole responsible for the polarization around peaks is induced by gravity on angular scales larger than twice the size of the horizon at decoupling. In the case of an overdensity, this causes a flow of photons towards the gravitational well on these scales, inducing a quadrupolar pattern (see, e.g., Coulson et al. 1994). The spherical symmetry of the gravitational interaction causes a rotation of the quadrupole in the vicinity of the well, resulting in a radial configuration in polarization. This radial polarization pattern implies $Q_r > 0$ and an overdensity implies $T < 0$ by the Sachs-Wolfe formulae, which leads to anticorrelation on these scales. Similarly, an underdensity leads to an outward flow and induces a tangential polarization pattern, once again leading to anticorrelation on these scales. At smaller scales, the polarized contribution is dominated by the dynamics of the photon fluid. The acoustic oscillations modulate the polarization pattern, leading to the different rings in the stacked images. The most noticeable rings in the stacked Q_r image are approximately at θ_s and $2\theta_s$. Thanks to the ℓ^2 enhancement, multiple acoustic peaks in the TE power spectrum may be captured and projected into ring patterns in the stacked polarization images. As photons flow towards the overdensity, they are compressed and the temperature increases, slowing the fluid descent into the well. Eventually, the radiation pressure becomes large enough to reverse the photon flow. This expansion cools the photons until they fall back towards the well. Note that the resulting inner ring was not observed in the WMAP analysis (Komatsu et al. 2011), since the resolution was too low.

Figure 41 clearly reveals all of the features described above. The two bright rings at $\theta_s \approx 0.011$ (0.6°) and $2\theta_s \approx 0.021$ (1.2°) are the predicted patterns associated with the first C_ℓ^{TE} acoustic
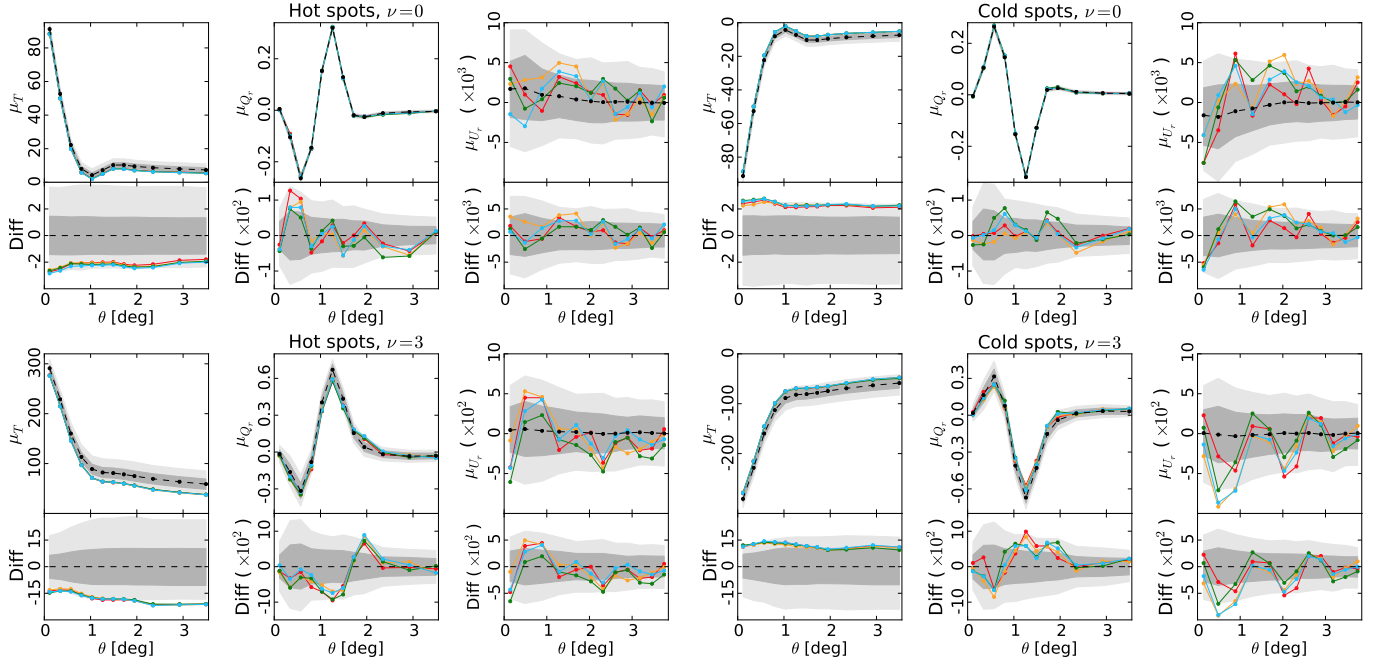
Fig. 43. Mean radial profiles of T, Q_r, and U_r in μK obtained for `Commander` (red), `NILC` (orange), `SEVEM` (green), and `SMICA` (blue). Each individual panel contains (*top*) the mean radial profiles and (*bottom*) the differences (denoted "Diff") between the mean profiles of the data and those computed from the ensemble mean of the simulations. Results based on stacks around temperature hot and cold spots are shown in the *left and right columns*, respectively. *Upper plots* present results for peaks selected above the null threshold, while *lower plots* show the equivalent results for peak amplitudes above (hot spots) or below (cold spots) 3 times the dispersion of the temperature map. The black dots (connected by dashed lines) depict the mean profiles and the shaded regions correspond to the 1σ (68%) and 2σ (95%) error bars. The mean profiles and error bars are determined from `SEVEM` simulations. Note that the Diff curves for each component-separation method are computed using the corresponding ensemble average, although only the ensemble average from `SEVEM` is shown here.

peak at $\ell \approx 310$, while the two faint rings are a striking illustration of the detection of multiple acoustic peaks in the TE power spectrum. The large-scale anticorrelation is suppressed due to the scale-dependent bias which results from the fact that peaks are defined by the second derivatives of the temperature field (e.g., Desjacques 2008).

We are now in a position to discuss the consistency of the *Planck* results with the predictions of a ΛCDM cosmology. For simplicity, further analysis is focused on the angular profiles, and specifically the mean, $\mu(\theta)$, estimated as the average of the angular profiles around all hot (cold) peaks above (below) a certain threshold ν. This analysis is performed directly on the sphere to avoid any repixelization error. Note that the expected value of the mean temperature angular profile is proportional to $\int \ell \, d\ell \, C_\ell^{TT} J_0(\ell\theta)$, whilst the expected values of the Q_r and U_r mean angular profiles are approximately proportional to $\int -\ell \, d\ell \, C_\ell^{TE} J_2(\ell\theta)$ and $\int -\ell \, d\ell \, C_\ell^{TB} J_2(\ell\theta)$, respectively. Since T has even parity and B has odd parity, the expectation value for C_ℓ^{TB} is zero, and the U_r mean angular profile is therefore expected to vanish.

A χ^2 estimator is used to quantify the differences between the profiles obtained from the data and the expected values estimated with simulations:

$$\chi^2 = [\mu(\theta) - \bar{\mu}(\theta)] \, \mathbf{C}^{-1} \, [\mu(\theta) - \bar{\mu}(\theta)]^\mathsf{T}, \tag{70}$$

with the covariance matrix defined as

$$C(i,j) = \frac{1}{N-1} \sum_{k=1}^{N} [\mu_k(\theta_i) - \bar{\mu}(\theta_i)] [\mu_k(\theta_j) - \bar{\mu}(\theta_j)], \tag{71}$$

where the sum is over the N simulations used to estimate this matrix and $\bar{\mu}(\theta)$ is the ensemble average. Note that although the profiles in Fig. 41 are derived from data at a resolution $N_{\text{side}} = 1024$,

faster convergence of the χ^2 statistic is achieved using maps at a lower resolution. We have verified that the results remain unchanged when adopting data with $N_{\text{side}} = 512$.

Figure 43 presents a comparison between the profiles obtained from the component-separated data and the mean value estimated from simulations processed through the `SEVEM` pipeline. Note that the error bars for the temperature profiles are asymmetric due to a bias in the selection of the peaks above a given threshold. Results for hot and cold spots are shown for two different thresholds, $\nu = 0$ and $\nu = 3$. There is generally excellent agreement between the different component-separation methods. A systematic deviation between the data and the simulations for the hot peaks in temperature ($\nu = 0$) is seen at a level greater than 1σ. This discrepancy increases at higher thresholds, reaching values of about 2σ for the $\nu = 3$ case. Similar behaviour is seen for the cold spots. For the Q_r angular profiles, the most striking differences appear around $\theta = 2°$ in the $\nu = 3$ case for hot peaks, and around $\theta = 1°5$ for the cold peaks. For the U_r angular profiles, where a null signal is expected (i.e., only noise is expected to be present), deviations at similar levels are seen.

Table 38 presents the corresponding p-values for this comparison. A theoretical χ^2 distribution is used to determine the probability that a sky drawn from the ΛCDM cosmology has a value larger than that derived from the data. We have verified this approach by comparing the empirical χ^2 distribution estimated from 100 simulations (in which the mean value and the covariance matrix are computed from a further 900 simulations) with the theoretical distribution with the corresponding degrees of freedom (see Fig. 44). The χ^2 value of the data is then estimated using the mean value and the covariance matrix determined from simulations. Although some differences are found among the
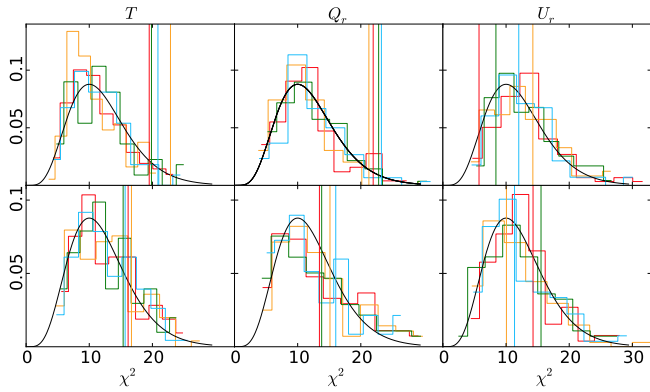
Fig. 44. χ^2 distributions obtained from the T (*left column*), Q_r (*middle column*), and U_r (*right column*) mean radial profiles centred on temperature hot spots selected above the null threshold (*upper row*) and three times the dispersion of the map (*bottom row*). The black lines correspond to the theoretical χ^2 distribution with 12 degrees of freedom, whilst the histograms show the distributions determined from 100 simulations computed through the `Commander` (red), `NILC` (orange), `SEVEM` (green), and `SMICA` (blue) pipelines. The vertical lines represent the χ^2 values obtained from the data.

Table 38. p-values of the T, Q_r, and U_r angular profiles computed from the stacking of hot and cold spots selected above the $\nu = 0$ and $\nu = 3$ thresholds.

	Probability [%]			
	Comm.	NILC	SEVEM	SMICA
$\nu = 0$ (hot spots)				
T	8	3	7	5
Q_r	4	5	3	3
U_r	93	28	75	44
$\nu = 3$ (hot spots)				
T	18	16	22	21
Q_r	34	23	31	19
U_r	28	61	21	50
$\nu = 0$ (cold spots)				
T	23	38	29	39
Q_r	86	85	63	78
U_r	14	11	39	34
$\nu = 3$ (cold spots)				
T	24	21	23	20
Q_r	21	51	29	52
U_r	30	13	30	8

component-separation methods, a general consistency between model and data is found.

Although the χ^2 test has the advantage of being sensitive to different types of deviations between model and data, does not assume prior knowledge about possible departures from the model, and can account for correlations between the various tests from which it is constructed, it can nevertheless be suboptimal under certain conditions. This appears to be the case when considering the systematic shift between data and simulations seen in the temperature profiles μ_T – the χ^2 statistic is not particularly sensitive to systematic deviations of constant sign. We therefore consider an alternative quantity, the integrated profile deviation, defined as

$$\Delta\mu_T(W) = \int_0^R \left[\mu_T(\theta) - \bar{\mu}_T(\theta)\right] W(\theta)\, d\theta, \qquad (72)$$

Table 39. p-values of $\Delta\mu_T$ computed from the stacking of hot and cold spots selected above the $\nu = 0$ and $\nu = 3$ thresholds from the `Commander`, `NILC`, `SEVEM`, and `SMICA` maps.

	Probability [%]			
	Comm.	NILC	SEVEM	SMICA
Hot spots				
T ($\nu = 0$)	96.0	95.8	96.2	97.1
T ($\nu = 3$)	98.6	98.2	98.3	98.7
Cold spots				
T ($\nu = 0$)	97.1	96.9	98.1	97.9
T ($\nu = 3$)	92.0	90.6	90.6	93.0

where R, the size of stacking patches, is taken to be $3°\!.5$ in this case. The weighting function is chosen to be proportional to the expected profile, but the results are robust for other choices, e.g., $W = 1$. The p-values obtained in this case are given in Table 39. These are consistent with what might be expected from visual inspection of the plots, i.e., the deviations are typically close to 2σ. These deviations are likely to be connected to the deficit in the observed power spectrum at low multipoles, as may be seen in Fig. 42. Here, the purple dashed line indicates the reduction in $\bar{\mu}_T$ if the theoretical C_ℓ values are reduced by 10% over the range $2 \leq \ell \leq 50$.

8.2. Generalized stacking

In this section, a much wider class of stacking methods is introduced, with particular emphasis on *oriented stacking*, a novel approach that has not previously been explored in the literature. We regard the stacking as *oriented* if the orientation of the local coordinate frame, and in particular the $\phi = 0$ axis, is correlated with the map that is being stacked. Thus, the stacking methodology in Sect. 8.1 is considered unoriented, because the orientation is defined relative to the local meridian pointing towards the Galactic south, rather than any property of the data themselves. Alternative approaches to unoriented stacking can also be considered. In this subsection, the orientation of each patch is chosen randomly from a uniform distribution in $[0, 2\pi)$. The unoriented T and Q_r images can then be directly compared with previous sections.

For unoriented stacking, the ensemble average of stacked fields cannot result in any intrinsic ϕ-dependence, as this would be averaged out by the uncorrelated orientation choices. The ϕ-dependence due to a specific choice of representation can always be removed via a local rotation. For example, the ensemble averages of $Q + iU$ around unoriented temperature peaks are proportional to $e^{2i\phi}$. A local rotation $(Q, U) \rightarrow (Q_r, U_r)$ (Kamionkowski et al. 1997) removes the $e^{2i\phi}$ factor and compresses the information into a single real map Q_r. For oriented stacking, the ϕ-dependence can be a mixture of a few Fourier modes ($e^{im\phi}$, for integer m). Each m mode corresponds to a radial (ϖ-dependent) function.

In what follows, we use the $N_{\rm side} = 1024$ component-separated maps at a resolution of $10'$ FWHM. The use of this higher resolution as compared to the $N_{\rm side} = 512$ data used in Sect. 8.1 is motivated by the smaller-scale features that are expected to result from the oriented stacking.

We also introduce the concept of the noise-free ensemble average (NFEA), which is defined as the ensemble average of stacked CMB-only maps for a fiducial cosmology. Recall that the fiducial model for the simulated sky maps, the *Planck* 2013

best-fit ΛCDM model (Planck Collaboration XVI 2014), differs from the updated *Planck* 2015 best-fit ΛCDM model (Planck Collaboration XIII 2016). In previous sections, this mismatch was partially accommodated by rescaling the CMB signal by a fixed scale factor. Here, we instead specifically adopt the 2015 best fit as a fiducial model for the data. When comparing the data to the simulations, we subtract the corresponding NFEA to minimize any bias resulting from cosmology dependence.

In the context of random Gaussian fields, the NFEA can be computed straightforwardly following Bardeen et al. (1986):

$$\langle M \rangle = \langle M\boldsymbol{w}^{\mathsf{T}} \rangle \langle \boldsymbol{w}\boldsymbol{w}^{\mathsf{T}} \rangle^{-1} \langle \boldsymbol{w} \rangle, \tag{73}$$

where M is the map (around the central peak) to be stacked, and $\boldsymbol{w}$ is the collection of Gaussian variables (on the central peak) that are related to peak selection and orientation determination. Equation (73) is only valid for Gaussian random variables. If the patch is rotated before stacking, the field value evaluated at a dynamic coordinate is, in general, not a random Gaussian variable. However, statistical isotropy guarantees that the rotation of patches is equivalent to an orientation constraint on the nonzero-spin field. For example, orienting each patch in the direction where $U = 0$ and $Q > 0$ is equivalent to the unoriented stacking case where only peaks satisfying the additional constraint $-\epsilon/2 < \arg(Q + iU) < \epsilon/2$ ($\epsilon \to 0^+$) are selected.

A further source of statistical bias can arise from noise mismatch between the simulations and the data. Since the effect of noise is to introduce random shifts in the peaks and hence suppress patterns in the stacked images, any noise mismatch can lead to pattern mismatch between the data and simulations. For the temperature data, the contribution due to noise mismatch is estimated to be at the sub-percent level, lower than the cosmic variance. For stacking on polarization peaks, the impact of the noise mismatch cannot be safely ignored. Thus, for quantitative comparisons in this paper, we only consider stacking on temperature peaks.

8.2.1. Oriented temperature stacking

The most straightforward way to orient a patch centred on a temperature peak is to align the horizontal axis with the major axis defined by a local quadratic expansion of the temperature field around the peak. The disadvantage of doing so is that the orientation is dominated by small-scale fluctuations that are noise-sensitive. A better choice is to use the major axis of the inverse Laplacian $\nabla^{-2}T$ that filters out the small-scale power. The inverse Laplacian is defined as:

$$\nabla^{-2}T(\hat{\boldsymbol{n}}) = -\sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} \frac{a_{\ell m}^{T}}{\ell(\ell+1)} Y_{\ell m}(\hat{\boldsymbol{n}}), \tag{74}$$

where $a_{\ell m}^{T}$ are the harmonic coefficients of the masked temperature map. Spin-2 maps Q_T, U_T are then defined by:

$$(Q_T \pm iU_T)(\hat{\boldsymbol{n}}) = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}^{T} [\,_{\pm 2}Y_{\ell m}(\hat{\boldsymbol{n}})]. \tag{75}$$

In the flat-sky limit, $Q_T \approx (\partial_x^2 - \partial_y^2)(\nabla^{-2}T)$ and $U_T \approx -2\partial_x\partial_y\nabla^{-2}T$. To align the $\nabla^{-2}T$ axes of the patches, we rotate each patch so that U_T vanishes and $Q_T \geq 0$ for the central peak.

Figure 45 presents the stacked images of SMICA temperature patches centred on temperature hot spots selected above the threshold $\nu = 0$, in both unoriented and oriented forms. These
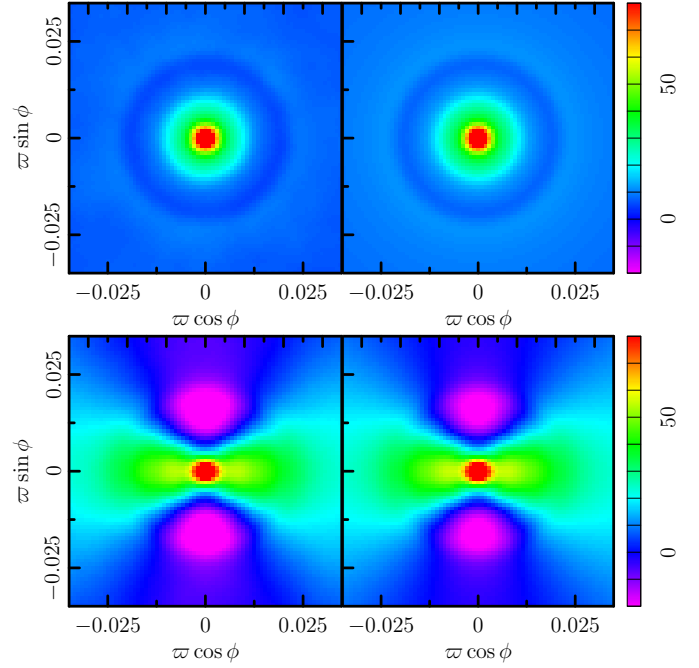

Fig. 45. Comparison between unoriented stacking (*upper panels*) and oriented stacking (*lower panels*) of temperature patches around temperature hot spots selected above the null threshold ($\nu = 0$). The *left panels* are the stacked SMICA maps, and the *right panels* their corresponding NFEAs. The image units are μK.

are seen to be in excellent agreement with their accompanying NFEAs, and, in the case of the unoriented stacks, with the results shown in Fig. 41, despite the different stacking methodologies adopted (and component separation method selected for visualization purposes).

The oriented T image is notably different from the unoriented one. The alignment between the major axis (of $\nabla^{-2}T$) and the horizontal axis results in an ellipse elongated along the horizontal axis, rather than a central disc. Moreover, the quadratic-term contribution is suppressed along the horizontal axis where the temperature profile is smoother, and enhanced along the vertical axis where the temperature profile is sharper. As a consequence, the dark ring visible in the upper panel at 1° splits into two cold blobs along the vertical axis.

To proceed with a quantitative analysis, we extract Fourier modes $T_m(\varpi)$ from the stacked map $T_{\text{stack}}(\varpi, \phi)$ as follows:

$$T_m(\varpi) = \frac{1}{(1 + \delta_{m0})\pi} \int_0^{2\pi} T_{\text{stack}}(\varpi, \phi) \cos m\phi \, d\phi, \tag{76}$$

where δ_{m0} is the Kronecker delta function. For odd m, the NFEA $\langle T_m \rangle$ vanishes due to statistical isotropy. For even m, a straightforward calculation shows that only $T_0(\varpi)$, which is equivalent to $\mu_T(\varpi)$, and $T_2(\varpi)$ have nonzero NFEAs.

As discussed previously in Sect. 8.1, there are some shortcomings of the standard χ^2 procedure that is generally used to assess the consistency of the data with simulations. The problem is simplified by studying the statistics of an integrated profile deviation:

$$\mathcal{T}_m(W) = \int_0^R [T_m(\varpi) - \langle T_m(\varpi) \rangle] W(\varpi) \, d\varpi, \tag{77}$$

where R, the size of the stacking patches, is taken to be 2° in our examples. The purpose of removing the NFEA, $\langle T_m(\varpi) \rangle$, which

Table 40. $\mathcal{T}_m$, as defined in Eqs. (77) and (78), for different thresholds ν.

Method	$\mathcal{T}_0$		$\mathcal{T}_2$	
	Hot spots	Cold spots	Hot spots	Cold spots
		Threshold $\nu = 0$		
Commander	$-0.03(\ 0.04^{+0.06+0.15}_{-0.04-0.07})$	$-0.04(\ 0.04^{+0.06+0.16}_{-0.04-0.07})$	$0.07(0.04^{+0.01+0.03}_{-0.02-0.04})$	$0.06(0.04^{+0.01+0.03}_{-0.02-0.04})$
NILC	$-0.04(\ 0.04^{+0.06+0.15}_{-0.04-0.07})$	$-0.05(\ 0.04^{+0.06+0.16}_{-0.04-0.07})$	$0.06(0.03^{+0.01+0.02}_{-0.02-0.04})$	$0.05(0.03^{+0.01+0.03}_{-0.02-0.04})$
SEVEM	$-0.03(\ 0.04^{+0.06+0.16}_{-0.04-0.07})$	$-0.04(\ 0.04^{+0.06+0.16}_{-0.04-0.07})$	$0.06(0.04^{+0.01+0.03}_{-0.02-0.04})$	$0.06(0.04^{+0.01+0.03}_{-0.02-0.04})$
SMICA	$-0.03(-0.01^{+0.03+0.07}_{-0.02-0.04})$	$-0.05(\ 0.00^{+0.03+0.07}_{-0.02-0.04})$	$0.06(0.04^{+0.01+0.03}_{-0.02-0.04})$	$0.06(0.04^{+0.01+0.03}_{-0.02-0.04})$
		Threshold $\nu = 1$		
Commander	$-0.06(\ 0.05^{+0.09+0.22}_{-0.05-0.10})$	$-0.06(\ 0.05^{+0.09+0.21}_{-0.06-0.09})$	$0.06(0.03^{+0.02+0.03}_{-0.02-0.04})$	$0.04(0.03^{+0.02+0.03}_{-0.02-0.04})$
NILC	$-0.06(\ 0.05^{+0.09+0.22}_{-0.05-0.10})$	$-0.07(\ 0.05^{+0.09+0.21}_{-0.05-0.09})$	$0.06(0.02^{+0.02+0.03}_{-0.02-0.04})$	$0.04(0.02^{+0.02+0.03}_{-0.02-0.04})$
SEVEM	$-0.06(\ 0.06^{+0.09+0.22}_{-0.05-0.10})$	$-0.06(\ 0.06^{+0.09+0.22}_{-0.05-0.10})$	$0.06(0.03^{+0.02+0.03}_{-0.02-0.04})$	$0.04(0.03^{+0.02+0.03}_{-0.02-0.04})$
SMICA	$-0.06(-0.01^{+0.04+0.10}_{-0.03-0.06})$	$-0.07(-0.01^{+0.04+0.10}_{-0.03-0.06})$	$0.06(0.03^{+0.02+0.03}_{-0.02-0.04})$	$0.04(0.03^{+0.02+0.03}_{-0.02-0.04})$

Notes. The expected values, together with the 1σ (68% CL) and 2σ (95% CL) ranges, determined from simulations are given in brackets.

Table 41. $\mathcal{T}_0$, as defined in Eqs. (77) and (78), for different thresholds ν and hemispheres.

Method	"North" $\mathcal{T}_0$		"South" $\mathcal{T}_0$	
	Hot spots	Cold spots	Hot spots	Cold spots
		Threshold $\nu = 0$		
Commander	$-0.02(\ 0.03^{+0.07+0.16}_{-0.04-0.07})$	$-0.03(\ 0.03^{+0.07+0.18}_{-0.04-0.07})$	$-0.05(\ 0.03^{+0.07+0.18}_{-0.05-0.07})$	$-0.06(\ 0.03^{+0.07+0.18}_{-0.04-0.07})$
NILC	$-0.02(\ 0.02^{+0.07+0.16}_{-0.04-0.07})$	$-0.03(\ 0.02^{+0.07+0.17}_{-0.04-0.07})$	$-0.05(\ 0.02^{+0.07+0.18}_{-0.04-0.07})$	$-0.06(\ 0.02^{+0.07+0.18}_{-0.04-0.07})$
SEVEM	$-0.02(\ 0.03^{+0.07+0.17}_{-0.04-0.07})$	$-0.03(\ 0.03^{+0.07+0.18}_{-0.04-0.07})$	$-0.05(\ 0.03^{+0.07+0.18}_{-0.05-0.07})$	$-0.06(\ 0.03^{+0.07+0.18}_{-0.04-0.07})$
SMICA	$-0.02(-0.01^{+0.04+0.09}_{-0.03-0.05})$	$-0.03(-0.01^{+0.04+0.09}_{-0.03-0.05})$	$-0.05(-0.01^{+0.04+0.08}_{-0.03-0.05})$	$-0.07(-0.01^{+0.04+0.08}_{-0.03-0.05})$
		Threshold $\nu = 1$		
Commander	$-0.04(\ 0.03^{+0.09+0.22}_{-0.06-0.10})$	$-0.05(\ 0.03^{+0.09+0.23}_{-0.06-0.10})$	$-0.08(\ 0.04^{+0.09+0.25}_{-0.06-0.11})$	$-0.08(\ 0.04^{+0.09+0.24}_{-0.06-0.10})$
NILC	$-0.05(\ 0.03^{+0.10+0.23}_{-0.06-0.10})$	$-0.06(\ 0.02^{+0.09+0.23}_{-0.06-0.10})$	$-0.08(\ 0.03^{+0.09+0.25}_{-0.06-0.11})$	$-0.08(\ 0.03^{+0.09+0.24}_{-0.06-0.10})$
SEVEM	$-0.04(\ 0.04^{+0.10+0.23}_{-0.06-0.10})$	$-0.05(\ 0.03^{+0.10+0.23}_{-0.06-0.10})$	$-0.08(\ 0.04^{+0.09+0.25}_{-0.07-0.11})$	$-0.08(\ 0.04^{+0.09+0.24}_{-0.06-0.11})$
SMICA	$-0.04(-0.02^{+0.05+0.13}_{-0.04-0.07})$	$-0.05(-0.02^{+0.05+0.13}_{-0.04-0.07})$	$-0.08(-0.02^{+0.05+0.11}_{-0.04-0.07})$	$-0.09(-0.02^{+0.05+0.12}_{-0.04-0.07})$

Notes. The "north" hemisphere is centred on the Galactic coordinate $(l, b) = (212°, -13°)$ and the "south" hemisphere in the opposite direction. The expected values, together with the 1σ (68% CL) and 2σ (95% CL) ranges, determined from simulations are given in brackets.

differs for the data and the simulations, is to minimize the impact of the cosmology dependence. A natural choice for the filter is $\langle T_m(\varpi)\rangle$ itself with a proper normalization:

$$W(\varpi) = \frac{\langle T_m(\varpi)\rangle}{\int_0^R \langle T_m(\varpi)\rangle^2 \mathrm{d}\varpi}. \qquad (78)$$

For the filter given by Eq. (78), the integrated profile deviation $\mathcal{T}_m$ describes the relative deviation from the NFEA. If ΛCDM is the correct model, the deviation is due to cosmic variance and noise. The distribution of $\mathcal{T}_m$ is obtained from simulations.

Table 40 presents a comparison of the $\mathcal{T}_m$ values derived from the *Planck* data and the FFP8 simulations. No inconsistencies in excess of the 3σ level have been found, although tensions around 2σ are seen.

The $m = 0$ projection kernel $J_0[(\ell + 1/2)\varpi]$ peaks at low ℓ. Thus $\mathcal{T}_0$ is cosmic-variance sensitive and the apparent discrepancy in it could be related to a low-ℓ power deficit. An example is shown in Fig. 42 for illustration. To test the robustness of this result, we have tried three additional filters: a top-hat filter $W = 1$, a linear filter $W = \varpi$, and a Gaussian filter $W = \exp(-\varpi^2/\sigma_g^2)$

with $\sigma_g = 1°$. In all cases, the power deficit remains at about the 2σ level.

Although the $\mathcal{T}_0$ deficit is not significant enough to falsify the ΛCDM model, further investigation of its properties may still be interesting and help us understand the other anomalies discussed in this paper. We consider two possibilities. Firstly the amplitude of the $\mathcal{T}_0$ deficit is of order 5–10%, which coincides with the level of hemispherical power asymmetry discussed in Sect. 6.1. To test whether the $\mathcal{T}_0$ deficit is localized on one hemisphere, we define the "north" direction to be aligned with the power asymmetry direction at $(l, b) = (212°, -13°)$ (Akrami et al. 2014) and compute $\mathcal{T}_0$ on the northern and southern hemispheres separately. The results are presented in Table 41. Although the $\mathcal{T}_0$ deficit is more significant for the southern hemisphere, it remains consistent with the ΛCDM prediction. Secondly, it is of interest to determine whether the $\mathcal{T}_0$ deficit is related to the Cold Spot discussed in Sect. 5.7. We therefore mask out the Cold Spot using a disc of radius 6° and repeat the calculation. The impact of this region on the $\mathcal{T}_0$ deficit is insignificant.
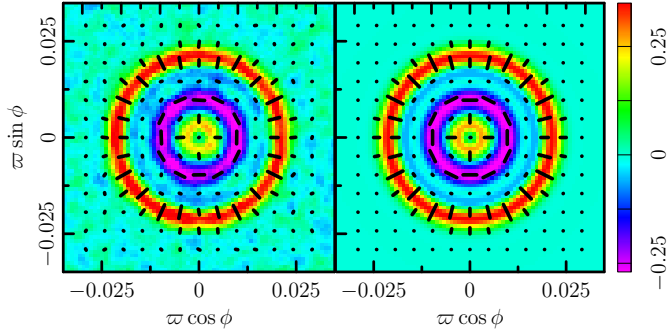
Fig. 46. Stacked Q_r image around temperature hot spots selected above the null threshold ($\nu = 0$) in the SMICA sky map. The *left panel* corresponds to the observed data and the *right panel* shows the NFEA. The image units are μK. The headless vectors (black solid lines) are the polarization directions for stacked Q_{stack}, U_{stack}. The lengths of the headless vectors are proportional to the polarization amplitude P_{stack}.

Tensions at the 2σ level are also seen for $\mathcal{T}_2$. However, due to the additional ℓ^2 factor in the projection kernel, the oriented ($m = 2$) component $\mathcal{T}_2$ is more sensitive to high-ℓ power where the cosmic variance is small, and an understanding of the noise properties of the data is more important. The former implies that the related uncertainty in $\mathcal{T}_2$ is, in general, smaller than that in $\mathcal{T}_0$. However, a mismatched cosmology, perhaps arising from a different primordial power amplitude A_s, can then lead to significant tension between the data and the simulations. Indeed, we find that without application of our cosmology calibration (i.e., the subtraction of the NFEA in Eq. (77)) the $\mathcal{T}_2$-tension between the data and simulations increases by about 0.5σ, whereas the variation of the $\mathcal{T}_0$-tension is $\lesssim 0.2\sigma$. The high-ℓ sensitivity of $\mathcal{T}_2$ also requires the use of an accurate noise model, and it is possible that the 1–2σ tension in $\mathcal{T}_2$ may be alleviated once improved noise simulations are available.

8.2.2. Oriented polarization stacking

The stacked Q and U images can be decomposed into Fourier modes, $Q + iU = \sum_{m=-\infty}^{\infty} P_m(\varpi) e^{im\phi}$. For unoriented $Q + iU$ stacking on temperature peaks, only $P_2(\varpi)$ has a non-zero NFEA, and it can be linked to the conventional Q_r stacking via $P_2 = -Q_r$. Figure 46 shows that the stacked Q_r image is in excellent agreement with its NFEA and the corresponding stacked image (fourth panel) in Fig. 41, despite the different stacking methodologies adopted (and component-separation method selected for visualization purposes). The length and orientation of the headless vectors represent the polarization amplitude, $P_{\text{stack}} \equiv \sqrt{Q_{\text{stack}}^2 + U_{\text{stack}}^2}$, and direction.

We next consider oriented stacking of the polarization maps, again using Q_T, U_T to define the orientation of the patches. The stacked polarization images around temperature peaks have $m = 0, 2, 4$ Fourier components. We can also choose to stack the polarization maps on P_T peaks, where $P_T = \sqrt{Q_T^2 + U_T^2}$. This picks up $m = 0, 4$ Fourier modes with no circularly symmetric (Q_r, $m = 2$) mode. In Fig. 47 we compare the (Q, U) images stacked centred either on T peaks (top panel) or on P_T peaks (bottom panel) with their corresponding NFEAs, and find excellent agreement.

For a quantitative comparison, we only consider stacking on temperature peaks and define the polarization integrated profile
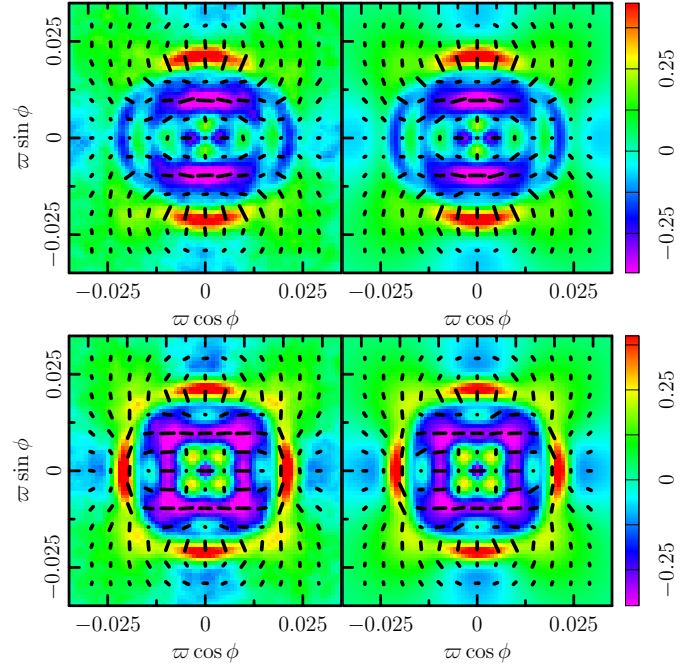

Fig. 47. Oriented stacking of polarization fields (Q, U) on temperature maxima (*upper panels*) and P_T maxima (*lower panels*). In both cases the threshold $\nu = 0$ is used and the orientation is chosen such that $U_T = 0$ and $Q_T \geq 0$ on the central peak. The image units are μK. The *left panels* are the stacked SMICA maps, and the *right panels* their NFEAs. See Fig. 46 for the meaning of the headless vectors (black dashed lines).

deviation

$$\mathcal{P}_m(W) = \int_0^R \left(P_m(\varpi) - \langle P_m(\varpi) \rangle \right) W(\varpi) \, d\varpi, \qquad (79)$$

where by default the filter is

$$W(\varpi) = \frac{\langle P_m(\varpi) \rangle}{\int_0^R \langle P_m(\varpi) \rangle^2 d\varpi}. \qquad (80)$$

The comparison of $\mathcal{P}_m$ ($m = 0, 2, 4$) between the data and the simulations is shown in Table 42, where the results are seen to be in excellent agreement.

Finally, we note that the peak selection does not have to be made from the temperature map. In Fig. 48 we show a few examples of stacking on polarization peaks using the $N_{\text{side}} = 512$ maps. The higher-resolution polarization data are too noisy for peak selection. In the upper panels, we compare stacked images of the E-mode map centred around E-mode peaks with the corresponding NFEA. We find that the noise impact is relatively minor for $20'$ FWHM maps and the plots are in qualitatively good agreement. Another possibility, shown in the lower panels, is to stack polarization maps centred on peaks determined from the corresponding polarization amplitude map. In this case the peaks are strongly biased by the quadratic noise contribution and quite visible deviation from the NFEA is observed in the stacked image.

9. Conclusions

In this paper, we have presented a study of the statistical isotropy and Gaussianity of the CMB using the *Planck* 2015 data, including the full mission for temperature. We do not claim that our results support or refute any particular physical model. Rather,

Table 42. $\mathcal{P}_m$, as defined in Eqs. (79) and (80), for different thresholds ν.

Method	$\mathcal{P}_0$	$\mathcal{P}_2$	$\mathcal{P}_4$
	Hot spots, threshold $\nu = 0$		
Commander	$0.06(0.02^{+0.03+0.05}_{-0.03-0.06})$	$-0.01(\ 0.01^{+0.01+0.02}_{-0.01-0.02})$	$0.04(0.01^{+0.02+0.05}_{-0.02-0.05})$
NILC	$0.05(0.02^{+0.03+0.05}_{-0.03-0.06})$	$-0.02(\ 0.00^{+0.01+0.02}_{-0.01-0.02})$	$0.03(0.01^{+0.02+0.05}_{-0.02-0.05})$
SEVEM	$0.05(0.02^{+0.03+0.06}_{-0.03-0.06})$	$0.01(\ 0.01^{+0.01+0.02}_{-0.01-0.02})$	$0.04(0.01^{+0.02+0.05}_{-0.02-0.05})$
SMICA	$0.05(0.03^{+0.03+0.05}_{-0.03-0.05})$	$-0.02(\ 0.00^{+0.01+0.02}_{-0.01-0.02})$	$0.03(0.01^{+0.02+0.05}_{-0.02-0.05})$
	Cold spots, threshold $\nu = 0$		
Commander	$0.06(0.02^{+0.03+0.05}_{-0.03-0.06})$	$-0.01(\ 0.01^{+0.01+0.02}_{-0.01-0.02})$	$0.03(0.01^{+0.03+0.05}_{-0.02-0.05})$
NILC	$0.06(0.02^{+0.03+0.05}_{-0.03-0.06})$	$-0.01(\ 0.00^{+0.01+0.02}_{-0.01-0.02})$	$0.04(0.01^{+0.02+0.04}_{-0.03-0.06})$
SEVEM	$0.06(0.03^{+0.03+0.05}_{-0.03-0.06})$	$0.01(\ 0.01^{+0.01+0.02}_{-0.01-0.02})$	$0.03(0.01^{+0.02+0.05}_{-0.03-0.05})$
SMICA	$0.05(0.03^{+0.03+0.05}_{-0.03-0.06})$	$-0.02(\ 0.00^{+0.01+0.02}_{-0.01-0.02})$	$0.03(0.02^{+0.02+0.04}_{-0.02-0.05})$
	Hot spots, threshold $\nu = 1$		
Commander	$0.04(0.02^{+0.03+0.06}_{-0.04-0.07})$	$-0.02(\ 0.00^{+0.02+0.03}_{-0.02-0.03})$	$0.05(0.01^{+0.03+0.06}_{-0.03-0.06})$
NILC	$0.06(0.02^{+0.03+0.07}_{-0.04-0.07})$	$-0.02(-0.01^{+0.01+0.03}_{-0.02-0.03})$	$0.05(0.01^{+0.03+0.06}_{-0.03-0.06})$
SEVEM	$0.05(0.02^{+0.04+0.07}_{-0.04-0.07})$	$-0.01(\ 0.00^{+0.02+0.03}_{-0.01-0.03})$	$0.05(0.01^{+0.03+0.06}_{-0.03-0.06})$
SMICA	$0.04(0.03^{+0.03+0.07}_{-0.03-0.07})$	$-0.02(\ 0.00^{+0.01+0.02}_{-0.01-0.03})$	$0.06(0.01^{+0.03+0.05}_{-0.03-0.06})$
	Cold spots, threshold $\nu = 1$		
Commander	$0.07(0.02^{+0.03+0.06}_{-0.03-0.07})$	$0.00(-0.01^{+0.02+0.03}_{-0.02-0.03})$	$0.01(0.01^{+0.03+0.06}_{-0.03-0.07})$
NILC	$0.08(0.02^{+0.03+0.06}_{-0.04-0.07})$	$-0.01(-0.01^{+0.01+0.03}_{-0.02-0.03})$	$0.01(0.01^{+0.03+0.06}_{-0.03-0.07})$
SEVEM	$0.09(0.02^{+0.03+0.07}_{-0.03-0.07})$	$0.00(\ 0.00^{+0.02+0.03}_{-0.02-0.03})$	$0.02(0.01^{+0.03+0.06}_{-0.03-0.06})$
SMICA	$0.06(0.03^{+0.03+0.06}_{-0.03-0.07})$	$-0.01(\ 0.00^{+0.01+0.02}_{-0.01-0.03})$	$0.02(0.01^{+0.03+0.06}_{-0.03-0.06})$

Notes. The expected values, together with 1σ (68% CL) and 2σ (95% CL) ranges, determined from simulations are given in brackets.
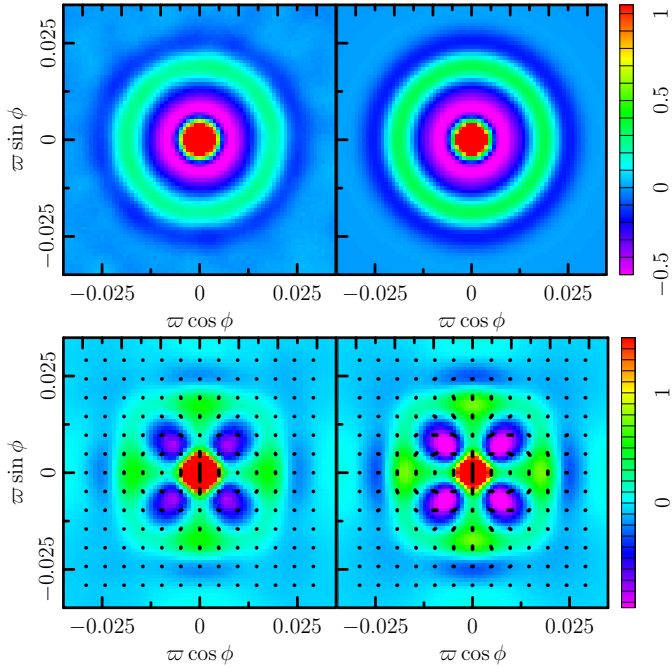

Fig. 48. *Top*: E-mode maps stacked on the unoriented E-mode maxima computed above the null threshold $\nu = 0$. *Bottom*: Q stacked around oriented polarization amplitude (P) maxima. In this case, no threshold is used and the orientation is chosen such that $U = 0$ and $Q \geq 0$ on the central peak. The *left panels* are the stacked SMICA maps, and the *right panels* their corresponding NFEAs. See Fig. 46 for the meaning of the headless vectors (black dashed lines). The image units are μK.

we focus on null-hypothesis testing: a number of tests are performed, then p-values are calculated and reported. It is in the

very nature of such a model-independent approach to leave the detailed interpretation to the reader. However, we do address the important subject of a posteriori correction where possible.

The statistical tests are performed on maps of the CMB anisotropy that result from the application of the four component-separation methods described in Planck Collaboration IX (2016). All of the results presented here are robust with respect to the choice of component-separated CMB map. This is important since it demonstrates the high quality and equivalence of the *Planck* component-separated data products rendered by different methodologies under varying assumptions.

We find that the CMB is largely consistent with statistical isotropy, although there are a few indications of anomalies with respect to the expectations of ΛCDM. Some of the tests we have performed are the same as those in PCIS13, in which case the results are consistent. Since many of these anomalies were also observed in the WMAP temperature data, we re-emphasize explicitly the statement we made in 2013 – that the agreement between the two independent experiments effectively rules out the possibility that the origin of these features can be found in residual systematic artefacts present in either data set. We have also performed a number of new tests, in order to try to narrow down the nature of the apparent violations of statistical isotropy. In addition, although the component-separated polarization maps contained in the *Planck* 2015 release are high-pass filtered, we have performed a stacking analysis that tests some aspects of the polarized sky while mitigating the impacts of noise and systematic effects.

In Sect. 4, we examined aspects of the Gaussianity of the CMB fluctuations. Tests of skewness, kurtosis, multi-normality, N-point functions, and Minkowski functionals yielded no indications of significant departures from Gaussianity, while the

variance of the CMB map was found to be low, in agreement with previous studies (PCIS13). First-order moments of filtered maps also exhibit the low-variance anomaly, as well as a kurtosis excess on certain scales associated with the Cold Spot. A new study of peak statistics finds results consistent with the expectations for a Gaussian random field, although the Cold Spot is again detected.

Section 5 provides an updated study of several previously known peculiarities. We study in detail the low variance anomaly, which appears to be associated with the known low-ℓ deficit in the angular power spectrum. We confirm the lack of large-scale angular correlations, relatively featureless northern ecliptic hemisphere 3- and 4-point functions, and indications of violations of point- and mirror-parity symmetry, although we make little or no attempt to correct these for a posteriori effects. We place tight constraints on a quadrupolar power modulation. The Cold Spot is examined further, and, while we find variance, skewness, and kurtosis angular profiles consistent with the expectations of statistically isotropic simulations, the mean temperature profile is anomalous at roughly the 1% level, apparently due to the surrounding hot ring – the feature that deviates most from the Gaussian model.

In Sect. 6 we perform a series of tests probing the well-known large-scale dipolar power asymmetry. We detect the asymmetry via pixel-to-pixel variance, as well as by measuring power explicitly or indirectly via ℓ to $\ell \pm 1$ mode coupling. The latter approach lends itself to a posteriori correction, which reduces the significance of the asymmetry substantially when no model for the anomaly is assumed. In addition, we perform two independent but related tests of directionality. One finds suggestions of anomalous clustering of directions out to relatively small scales while the other does not, evidently due to being optimized for slightly different forms of directionality.

Section 7 demonstrates that the significances of several large-angular-scale anomalies are robust to the use of larger sky coverage, with the observed small changes being consistent with expectations from random Gaussian statistics.

Finally, Sect. 8 presents the results of the stacking of temperature and polarization peaks. We find results that are largely consistent with statistically isotropic simulations, both for oriented and unoriented stacking. The exception is a low unoriented temperature profile, which seems to be yet another reflection of the large-scale power deficit.

With the *Planck* 2015 release, we are probably near the limit of our ability to probe the CMB anomalies with temperature fluctuations alone. The use of large-angular-scale polarization, expected for the final *Planck* release, should enable *independent* tests of these peculiar features. Importantly, this will reduce or eliminate the subjectivity and ambiguity in interpreting their statistical significance. It is a tantalizing possibility that some of the anomalies described in this paper will take us beyond the standard model of cosmology.

References

Ackerman, L., Carroll, S. M., & Wise, M. B. 2007, Phys. Rev. D, 75, 083502
Adhikari, S. 2015, MNRAS, 446, 4232
Akrami, Y., Fantaye, Y., Shafieloo, A., et al. 2014, ApJ, 784, L42
Aluri, P. K., Pant, N., Rotti, A., & Souradeep, T. 2015, Phys. Rev. D, 92, 083015
Axelsson, M., Fantaye, Y., Hansen, F. K., et al. 2013, ApJ, 773, L3
Baldi, P., Kerkyacharian, G., Marinucci, D., & Picard, D. 2009, Annals of Statistics, 37, 1150
Bardeen, J. M., Bond, J. R., Kaiser, N., & Szalay, A. S. 1986, ApJ, 304, 15
Ben-David, A., & Kovetz, E. D. 2014, MNRAS, 445, 2116
Bennett, C. L., Halpern, M., Hinshaw, G., et al. 2003, ApJS, 148, 1
Bennett, C. L., Hill, R. S., Hinshaw, G., et al. 2011, ApJS, 192, 17
Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, ApJS, 208, 20
Bond, J. R., & Efstathiou, G. 1987, MNRAS, 226, 655
Bond, J. R., Frolov, A. V., Huang, Z., & Kofman, L. 2009, Phys. Rev. Lett., 103, 071301
Bueno Sanchez, J. 2014, Phys. Lett. B, 739, 269
Bunn, E. F., & Scott, D. 2000, MNRAS, 313, 331
Cai, Y.-C., Cole, S., Jenkins, A., & Frenk, C. S. 2010, MNRAS, 407, 201
Cayón, L., Jin, J., & Treaster, A. 2005, MNRAS, 362, 826
Challinor, A., & van Leeuwen, F. 2002, Phys. Rev. D, 65, 103001
Copi, C. J., Huterer, D., Schwarz, D. J., & Starkman, G. D. 2007, Phys. Rev. D, 75, 023507
Copi, C. J., Huterer, D., Schwarz, D. J., & Starkman, G. D. 2009, MNRAS, 399, 295
Copi, C. J., Huterer, D., Schwarz, D. J., & Starkman, G. D. 2015, MNRAS, 451, 2978
Coulson, D., Crittenden, R. G., & Turok, N. 1994, Phys. Rev. Lett., 73, 2390
Cruz, M., Martínez-González, E., Vielva, P., & Cayón, L. 2005, MNRAS, 356, 29
Cruz, M., Tucci, M., Martínez-González, E., & Vielva, P. 2006, MNRAS, 369, 57
Cruz, M., Turok, N., Vielva, P., Martínez-González, E., & Hobson, M. 2007, Science, 318, 1612
Cruz, M., Martínez-González, E., Vielva, P., et al. 2008, MNRAS, 390, 913
Cruz, M., Vielva, P., Martínez-González, E., & Barreiro, R. B. 2011, MNRAS, 412, 2383
Curto, A., Aumont, J., Macías-Pérez, J. F., et al. 2007, A&A, 474, 23
Curto, A., Macías-Pérez, J. F., Martínez-González, E., et al. 2008, A&A, 486, 383
Curto, A., Martínez-González, E., & Barreiro, R. B. 2011, MNRAS, 412, 1038
Czech, B., Kleban, M., Larjo, K., Levi, T. S., & Sigurdson, K. 2010, J. Cosmol. Astropart. Phys., 12, 23
de Oliveira-Costa, A., Smoot, G. F., & Starobinsky, A. A. 1996, ApJ, 468, 457
De Troia, G., Ade, P. A. R., Bock, J. J., et al. 2007, ApJ, 670, L73
Desjacques, V. 2008, Phys. Rev. D, 78, 103503
Efstathiou, G. 2004, MNRAS, 348, 885
Eriksen, H. K., Hansen, F. K., Banday, A. J., Górski, K. M., & Lilje, P. B. 2004a, ApJ, 605, 14
Eriksen, H. K., Novikov, D. I., Lilje, P. B., Banday, A. J., & Górski, K. M. 2004b, ApJ, 612, 64
Eriksen, H. K., Banday, A. J., Górski, K. M., & Lilje, P. B. 2005, ApJ, 622, 58
Fantaye, Y. 2014, ArXiv e-prints [arXiv:1409.1114]
Fantaye, Y., Marinucci, D., Hansen, F., & Maino, D. 2015, Phys. Rev. D, 91, 063501
Feeney, S. M., Johnson, M. C., Mortlock, D. J., & Peiris, H. V. 2011, Phys. Rev. Lett., 107, 071301
Ferreira, P. G., & Magueijo, J. 1997, Phys. Rev. D, 56, 4578
Finelli, F., Gruppuso, A., Paci, F., & Starobinsky, A. 2012, J. Cosmol. Astropart. Phys., 1207, 049
Finelli, F., Garcia-Bellido, J., Kovacs, A., Paci, F., & Szapudi, I. 2016, MNRAS, 455, 1246
Fixsen, D. J., Cheng, E. S., Gales, J. M., et al. 1996, ApJ, 473, 576
Gay, C., Pichon, C., & Pogosyan, D. 2012, Phys. Rev. D, 85, 023011
Gordon, C. 2007, ApJ, 656, 636
Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759
Gott, J. R., Colley, W. N., Park, C.-G., Park, C., & Mugnolo, C. 2007, MNRAS, 377, 1668
Gruppuso, A., Finelli, F., Natoli, P., et al. 2011, MNRAS, 411, 1445
Gruppuso, A., Natoli, P., Paci, F., et al. 2013, J. Cosmol. Astropart. Phys., 7, 47
Gurzadyan, V. G., Kashin, A. L., Khachatryan, H., et al. 2014, A&A, 566, A135
Hajian, A. 2007, ArXiv e-prints [arXiv:astro-ph/0702723]
Hajian, A., & Souradeep, T. 2003, ApJ, 597, L5
Hajian, A., & Souradeep, T. 2006, Phys. Rev. D, 74, 123521
Hansen, F. K., Banday, A. J., Górski, K. M., Eriksen, H. K., & Lilje, P. B. 2009, ApJ, 704, 1448
Hanson, D., & Lewis, A. 2009, Phys. Rev. D, 80, 063004

Hanson, D., Lewis, A., & Challinor, A. 2010, Phys. Rev. D, 81, 103003
Hartlap, J., Simon, P., & Schneider, P. 2007, A&A, 464, 399
Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, A&AS, 47, 1
Hikage, C., Matsubara, T., Coles, P., et al. 2008, MNRAS, 389, 1439
Hinshaw, G., Banday, A. J., Bennett, C. L., et al. 1996, ApJ, 464, L25
Hinshaw, G., Weiland, J. L., Hill, R. S., et al. 2009, ApJS, 180, 225
Hoftuft, J., Eriksen, H. K., Banday, A. J., et al. 2009, ApJ, 699, 985
Hou, Z., Banday, A. J., & Górski, K. M. 2009, MNRAS, 396, 1273
Inoue, K. T., & Silk, J. 2006, ApJ, 648, 23
Jaffe, T. R., Banday, A. J., Eriksen, H. K., Górski, K. M., & Hansen, F. K. 2005, ApJ, 629, L1
Kamionkowski, M., & Knox, L. 2003, Phys. Rev. D, 67, 063001
Kamionkowski, M., Kosowsky, A., & Stebbins, A. 1997, Phys. Rev. D, 55, 7368
Kim, J., & Komatsu, E. 2013, Phys. Rev. D, 88, 101301
Kim, J., & Naselsky, P. 2010a, ApJ, 714, L265
Kim, J., & Naselsky, P. 2010b, Phys. Rev. D, 82, 063002
Knutsson, H., & Westin, C.-F. 1993, Conf. IEEE Comput. Soc., 515
Kogut, A., Spergel, D. N., Barnes, C., et al. 2003, ApJS, 148, 161
Komatsu, E., Kogut, A., Nolta, M. R., et al. 2003, ApJS, 148, 119
Komatsu, E., Dunkley, J., Nolta, M. R., et al. 2009, ApJS, 180, 330
Komatsu, E., Smith, K. M., Dunkley, J., et al. 2011, ApJS, 192, 18
Land, K., & Magueijo, J. 2005a, Phys. Rev. Lett., 95, 071301
Land, K., & Magueijo, J. 2005b, Phys. Rev. D, 72, 101302
Larson, D. L., & Wandelt, B. D. 2004, ApJ, 613, L85
Larson, D. L., & Wandelt, B. D. 2005, ArXiv e-prints [arXiv:astro-ph/0505046]
Liu, X., & Zhang, S. N. 2005, ApJ, 633, 542
Marinucci, D., Pietrobon, D., Balbi, A., et al. 2008, MNRAS, 383, 539
Martinez-González, E., Gallegos, J. E., Argüeso, F., Cayón, L., & Sanz, J. L. 2002, MNRAS, 336, 22
Matsubara, T. 2010, Phys. Rev. D, 81, 083505
McEwen, J. D., Hobson, M. P., Lasenby, A. N., & Mortlock, D. J. 2005, MNRAS, 359, 1583
McEwen, J. D., Vielva, P., Wiaux, Y., et al. 2007, J. Fourier Analysis and Applications, 13, 495
McEwen, J. D., Feeney, S. M., Johnson, M. C., & Peiris, H. V. 2012, Phys. Rev. D, 85, 103502
Mecke, K. R., Buchert, T., & Wagner, H. 1994, A&A, 288, 697
Mikkelsen, K., Næss, S. K., & Eriksen, H. K. 2013, ApJ, 777, 172
Mitra, S., Rocha, G., Górski, K. M., et al. 2011, ApJS, 193, 5
Monteserín, C., Barreiro, R. B., Vielva, P., et al. 2008, MNRAS, 387, 209
Moss, A., Scott, D., Zibin, J. P., & Battye, R. 2011, Phys. Rev. D, 84, 023014
Nadathur, S., Lavinto, M., Hotchkiss, S., & Räsänen, S. 2014, Phys. Rev. D., 90, 103510
Notari, A., & Quartin, M. 2015, J. Cosmol. Astropart. Phys., 6, 47
Planck Collaboration I. 2014, A&A, 571, A1
Planck Collaboration XV. 2014, A&A, 571, A15
Planck Collaboration XVI. 2014, A&A, 571, A16
Planck Collaboration XVII. 2014, A&A, 571, A17
Planck Collaboration XIX. 2014, A&A, 571, A19
Planck Collaboration XXIII. 2014, A&A, 571, A23
Planck Collaboration XXVII. 2014, A&A, 571, A27
Planck Collaboration I. 2016, A&A, 594, A1
Planck Collaboration II. 2016, A&A, 594, A2
Planck Collaboration III. 2016, A&A, 594, A3
Planck Collaboration IV. 2016, A&A, 594, A4
Planck Collaboration V. 2016, A&A, 594, A5
Planck Collaboration VI. 2016, A&A, 594, A6
Planck Collaboration VII. 2016, A&A, 594, A7
Planck Collaboration VIII. 2016, A&A, 594, A8
Planck Collaboration IX. 2016, A&A, 594, A9
Planck Collaboration X. 2016, A&A, 594, A10
Planck Collaboration XI. 2016, A&A, 594, A11
Planck Collaboration XII. 2016, A&A, 594, A12
Planck Collaboration XIII. 2016, A&A, 594, A13
Planck Collaboration XIV. 2016, A&A, 594, A14
Planck Collaboration XV. 2016, A&A, 594, A15
Planck Collaboration XVI. 2016, A&A, 594, A16
Planck Collaboration XVII. 2016, A&A, 594, A17
Planck Collaboration XVIII. 2016, A&A, 594, A18
Planck Collaboration XIX. 2016, A&A, 594, A19
Planck Collaboration XX. 2016, A&A, 594, A20
Planck Collaboration XXI. 2016, A&A, 594, A21
Planck Collaboration XXII. 2016, A&A, 594, A22
Planck Collaboration XXIII. 2016, A&A, 594, A23
Planck Collaboration XXIV. 2016, A&A, 594, A24
Planck Collaboration XXV. 2016, A&A, 594, A25

Planck Collaboration XXVI. 2016, A&A, 594, A26
Planck Collaboration XXVII. 2016, A&A, 594, A27
Planck Collaboration XXVIII. 2016, A&A, 594, A28
Plaszczynski, S., Montier, L., Levrier, F., & Tristram, M. 2014, MNRAS, 439, 4048
Pogosyan, D., Gay, C., & Pichon, C. 2009, Phys. Rev. D, 80, 081301
Pontzen, A., & Peiris, H. V. 2010, Phys. Rev. D, 81, 103008
Quartin, M., & Notari, A. 2015, J. Cosmol. Astropart. Phys., 1, 008
Rudnick, L., Brown, S., & Williams, L. R. 2007, ApJ, 671, 40
Savitzky, A., & Golay, M. J. E. 1964, J. Anal. Chem., 36, 1627
Schmalzing, J., & Buchert, T. 1997, ApJ, 482, L1
Schmalzing, J., & Gorski, K. M. 1998, MNRAS, 297, 355
Scott, D. 1991, A&A, 242, 1
Spergel, D. N., Verde, L., Peiris, H. V., et al. 2003, ApJS, 148, 175
Spergel, D. N., Bean, R., Doré, O., et al. 2007, ApJS, 170, 377
Stannard, A., & Coles, P. 2005, MNRAS, 364, 929
Starobinsky, A. A. 1993, JETP Lett., 57, 622
Stevens, D., Scott, D., & Silk, J. 1993, Phys. Rev. Lett., 71, 20
Szapudi, I., Kovács, A., Granett, B. R., et al. 2015, MNRAS, 450, 288
Tegmark, M., de Oliveira-Costa, A., & Hamilton, A. 2003, Phys. Rev. D, 68, 123523
Tomita, K. 2005, Phys. Rev. D, 72, 103506
Tomita, K., & Inoue, K. T. 2008, Phys. Rev. D, 77, 103522
Vanmarcke, E. 1983, in Random Fields (Cambridge, Massachusetts, USA: The MIT Press), March (Paper), ed. E. Vanmarcke, 372
Vielva, P. 2007, in SPIE Conf. Ser., 6701
Vielva, P. 2010, Adv. Astron., 2010, 592094
Vielva, P., Martínez-González, E., Barreiro, R. B., Sanz, J. L., & Cayón, L. 2004, ApJ, 609, 22
Vielva, P., Wiaux, Y., Martínez-González, E., & Vandergheynst, P. 2007, MNRAS, 381, 932
Watson, W. A., Diego, J. M., Gottlöber, S., et al. 2014, MNRAS, 438, 412
Wick, G. C. 1950, Phys. Rev., 80, 268
Zaldarriaga, M., & Seljak, U. 1997, Phys. Rev. D, 55, 1830
Zhang, R., & Huterer, D. 2010, Astropart. Phys., 33, 69
Zhao, W. 2013, MNRAS, 433, 3498
Zibin, J. P. 2014, ArXiv e-prints [arXiv:1408.4442]

[1] APC, AstroParticule et Cosmologie, Université Paris Diderot, CNRS/IN2P3, CEA/lrfu, Observatoire de Paris, Sorbonne Paris Cité, 10 rue Alice Domon et Léonie Duquet, 75205 Paris Cedex 13, France
[2] Aalto University Metsähovi Radio Observatory and Dept of Radio Science and Engineering, PO Box 13000, 00076 AALTO, Finland
[3] African Institute for Mathematical Sciences, 6–8 Melrose Road, Muizenberg, Cape Town, South Africa
[4] Agenzia Spaziale Italiana Science Data Center, via del Politecnico snc, 00133 Roma, Italy
[5] Aix Marseille Université, CNRS, LAM (Laboratoire d'Astrophysique de Marseille) UMR 7326, 13388 Marseille, France
[6] Astrophysics Group, Cavendish Laboratory, University of Cambridge, J J Thomson Avenue, Cambridge CB3 0HE, UK
[7] Astrophysics & Cosmology Research Unit, School of Mathematics, Statistics & Computer Science, University of KwaZulu-Natal, Westville Campus, Private Bag X54001, 4000 Durban, South Africa
[8] CITA, University of Toronto, 60 St. George St., Toronto, ON M5S 3H8, Canada
[9] CNRS, IRAP, 9 Av. colonel Roche, BP 44346, 31028 Toulouse Cedex 4, France
[10] CRANN, Trinity College, Dublin, Ireland
[11] California Institute of Technology, Pasadena, California CA 91125, USA
[12] Centre for Theoretical Cosmology, DAMTP, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, UK
[13] Centro de Estudios de Física del Cosmos de Aragón (CEFCA), Plaza San Juan, 1, planta 2, 44001 Teruel, Spain
[14] Computational Cosmology Center, Lawrence Berkeley National Laboratory, Berkeley, California, USA
[15] Consejo Superior de Investigaciones Científicas (CSIC), Madrid, Spain
[16] DSM/Irfu/SPP, CEA-Saclay, 91191 Gif-sur-Yvette Cedex, France

[17] DTU Space, National Space Institute, Technical University of Denmark, Elektrovej 327, 2800 Kgs. Lyngby, Denmark

[18] Département de Physique Théorique, Université de Genève, 24 Quai E. Ansermet,1211 Genève 4, Switzerland

[19] Departamento de Astrofísica, Universidad de La Laguna (ULL), 38206 La Laguna, Tenerife, Spain

[20] Departamento de Física, Universidad de Oviedo, Avda. Calvo Sotelo s/n, 33003 Oviedo, Spain

[21] Departamento de Matemáticas, Estadística y Computación, Universidad de Cantabria, Avda. de los Castros s/n, 39005 Santander, Spain

[22] Department of Astronomy and Astrophysics, University of Toronto, 50 Saint George Street, Toronto, Ontario, Canada

[23] Department of Astrophysics/IMAPP, Radboud University Nijmegen, PO Box 9010, 6500 GL Nijmegen, The Netherlands

[24] Department of Physics & Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, British Columbia, Canada

[25] Department of Physics and Astronomy, Dana and David Dornsife College of Letter, Arts and Sciences, University of Southern California, Los Angeles, CA 90089, USA

[26] Department of Physics and Astronomy, University College London, London WC1E 6BT, UK

[27] Department of Physics, Florida State University, Keen Physics Building, 77 Chieftan Way, Tallahassee, Florida, USA

[28] Department of Physics, Gustaf Hällströmin katu 2a, University of Helsinki, 00100 Helsinki, Finland

[29] Department of Physics, Princeton University, Princeton, New Jersey NJ 08544, USA

[30] Department of Physics, University of California, Santa Barbara, California CA 93106, USA

[31] Department of Physics, University of Illinois at Urbana-Champaign, 1110 West Green Street, Urbana, Illinois, USA

[32] Dipartimento di Fisica e Astronomia G. Galilei, Università degli Studi di Padova, via Marzolo 8, 35131 Padova, Italy

[33] Dipartimento di Fisica e Scienze della Terra, Università di Ferrara, via Saragat 1, 44122 Ferrara, Italy

[34] Dipartimento di Fisica, Università La Sapienza, P. le A. Moro 2, Roma, Italy

[35] Dipartimento di Fisica, Università degli Studi di Milano, via Celoria, 16 Milano, Italy

[36] Dipartimento di Fisica, Università degli Studi di Trieste, via A. Valerio 2, Trieste, Italy

[37] Dipartimento di Matematica, Università di Roma Tor Vergata, via della Ricerca Scientifica, 1 Roma, Italy

[38] Discovery Center, Niels Bohr Institute, 17 Blegdamsvej, Copenhagen, Denmark

[39] European Space Agency, ESAC, Planck Science Office, Camino bajo del Castillo, s/n, Urbanización Villafranca del Castillo, Villanueva de la Cañada, 28692 Madrid, Spain

[40] European Space Agency, ESTEC, Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands

[41] Facoltà di Ingegneria, Università degli Studi e-Campus, via Isimbardi 10, 22060 Novedrate (CO), Italy

[42] Gran Sasso Science Institute, INFN, viale F. Crispi 7, 67100 L'Aquila, Italy

[43] HGSFP and University of Heidelberg, Theoretical Physics Department, Philosophenweg 16, 69120 Heidelberg, Germany

[44] Helsinki Institute of Physics, Gustaf Hällströmin katu 2, University of Helsinki, Helsinki, Finland

[45] INAF–Osservatorio Astronomico di Padova, Vicolo dell'Osservatorio 5, Padova, Italy

[46] INAF–Osservatorio Astronomico di Roma, via di Frascati 33, Monte Porzio Catone, Italy

[47] INAF–Osservatorio Astronomico di Trieste, via G.B. Tiepolo 11, Trieste, Italy

[48] INAF/IASF Bologna, via Gobetti 101, Bologna, Italy

[49] INAF/IASF Milano, via E. Bassini 15, Milano, Italy

[50] INFN, Sezione di Bologna, via Irnerio 46, 40126 Bologna, Italy

[51] INFN, Sezione di Roma 1, Università di Roma Sapienza, Piazzale Aldo Moro 2, 00185 Roma, Italy

[52] INFN, Sezione di Roma 2, Università di Roma Tor Vergata, via della Ricerca Scientifica, 1 Roma, Italy

[53] INFN/National Institute for Nuclear Physics, via Valerio 2, 34127 Trieste, Italy

[54] IPAG: Institut de Planétologie et d'Astrophysique de Grenoble, Université Grenoble Alpes, IPAG; CNRS, IPAG, 38000 Grenoble, France

[55] IUCAA, Post Bag 4, Ganeshkhind, Pune University Campus, 411 007 Pune, India

[56] Imperial College London, Astrophysics group, Blackett Laboratory, Prince Consort Road, London, SW7 2AZ, UK

[57] Infrared Processing and Analysis Center, California Institute of Technology, Pasadena, CA 91125, USA

[58] Institut Néel, CNRS, Université Joseph Fourier Grenoble I, 25 rue des Martyrs, Grenoble, France

[59] Institut Universitaire de France, 103 bd Saint-Michel, 75005 Paris, France

[60] Institut d'Astrophysique Spatiale, CNRS (UMR 8617) Université Paris-Sud 11, Bâtiment 121, 91440 Orsay, France

[61] Institut d'Astrophysique de Paris, CNRS (UMR 7095), 98 bis Boulevard Arago, 75014 Paris, France

[62] Institut für Theoretische Teilchenphysik und Kosmologie, RWTH Aachen University, 52056 Aachen, Germany

[63] Institute for Space Sciences, Bucharest-Magurale, 077125 Romania

[64] Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

[65] Institute of Theoretical Astrophysics, University of Oslo, Blindern, 0371 Oslo, Norway

[66] Instituto de Astrofísica de Canarias, C/Vía Láctea s/n, La Laguna, 38205 Tenerife, Spain

[67] Instituto de Física de Cantabria (CSIC-Universidad de Cantabria), Avda. de los Castros s/n, 39005 Santander, Spain

[68] Istituto Nazionale di Fisica Nucleare, Sezione di Padova, via Marzolo 8, 35131 Padova, Italy

[69] Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, California, USA

[70] Jodrell Bank Centre for Astrophysics, Alan Turing Building, School of Physics and Astronomy, The University of Manchester, Oxford Road, Manchester, M13 9PL, UK

[71] Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA

[72] Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge, CB3 0HA, UK

[73] Kazan Federal University, 18 Kremlyovskaya St., 420008 Kazan, Russia

[74] LAL, Université Paris-Sud, CNRS/IN2P3, Orsay, France

[75] LERMA, CNRS, Observatoire de Paris, 61 Avenue de l'Observatoire, 75014 Paris, France

[76] Laboratoire AIM, IRFU/Service d'Astrophysique – CEA/DSM – CNRS – Université Paris Diderot, Bât. 709, CEA-Saclay, 91191 Gif-sur-Yvette Cedex, France

[77] Laboratoire Traitement et Communication de l'Information, CNRS (UMR 5141) and Télécom ParisTech, 46 rue Barrault, 75634 Paris Cedex 13, France

[78] Laboratoire de Physique Subatomique et Cosmologie, Université Grenoble-Alpes, CNRS/IN2P3, 53 rue des Martyrs, 38026 Grenoble Cedex, France

[79] Laboratoire de Physique Théorique, Université Paris-Sud 11 & CNRS, Bâtiment 210, 91405 Orsay, France

[80] Lawrence Berkeley National Laboratory, Berkeley, California, USA

[81] Lebedev Physical Institute of the Russian Academy of Sciences, Astro Space Centre, 84/32 Profsoyuznaya st., GSP-7, 117997 Moscow, Russia

[82] Leung Center for Cosmology and Particle Astrophysics, National Taiwan University, 10617 Taipei, Taiwan

[83] Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85741 Garching, Germany

[84] McGill Physics, Ernest Rutherford Physics Building, McGill University, 3600 rue University, Montréal, QC, H3A 2T8, Canada

[85] National University of Ireland, Department of Experimental Physics, Maynooth, Co. Kildare, Ireland

[86] Nicolaus Copernicus Astronomical Center, Bartycka 18, 00-716 Warsaw, Poland

[87] Niels Bohr Institute, Blegdamsvej 17, 2100 Copenhagen, Denmark

[88] Optical Science Laboratory, University College London, Gower Street, London, UK

[89] SISSA, Astrophysics Sector, via Bonomea 265, 34136 Trieste, Italy

[90] School of Physics and Astronomy, Cardiff University, Queens Buildings, The Parade, Cardiff, CF24 3AA, UK

[91] School of Physics and Astronomy, University of Nottingham, Nottingham NG7 2RD, UK

[92] Simon Fraser University, Department of Physics, 8888 University Drive, Burnaby BC, Canada

[93] Sorbonne Université-UPMC, UMR 7095, Institut d'Astrophysique de Paris, 98 bis Boulevard Arago, 75014 Paris, France

[94] Space Research Institute (IKI), Russian Academy of Sciences, Profsoyuznaya Str, 84/32, 117997 Moscow, Russia

[95] Space Sciences Laboratory, University of California, Berkeley, California, USA

[96] Special Astrophysical Observatory, Russian Academy of Sciences, Nizhnij Arkhyz, Zelenchukskiy region, 369167 Karachai-Cherkessian Republic, Russia

[97] Stanford University, Dept of Physics, Varian Physics Bldg, 382 via Pueblo Mall, Stanford, California, USA

[98] Sub-Department of Astrophysics, University of Oxford, Keble Road, Oxford OX1 3RH, UK

[99] Theory Division, PH-TH, CERN, 1211, Geneva 23, Switzerland

[100] UPMC Univ Paris 06, UMR 7095, 98 bis Boulevard Arago, 75014 Paris, France

[101] Université de Toulouse, UPS-OMP, IRAP, 31028 Toulouse Cedex 4, France

[102] University of Granada, Departamento de Física Teórica y del Cosmos, Facultad de Ciencias, 18071 Granada, Spain

[103] University of Granada, Instituto Carlos I de Física Teórica y Computacional, 18071 Granada, Spain

[104] University of Heidelberg, Institute for Theoretical Physics, Philosophenweg 16, 69120 Heidelberg, Germany

[105] Warsaw University Observatory, Aleje Ujazdowskie 4, 00-478 Warszawa, Poland

Appendix A: Generalized Savitzky-Golay polynomials

In the construction of optimal linear filters, one needs to combine information about the (statistically isotropic) CMB signal, anisotropic instrumental noise, masking to be applied for the elimination of foreground contributions, and a model for any non-Gaussian signal for matched filtering. These can be combined in a general framework of normalized convolutions (Knutsson & Westin 1993), where the filtered field is defined as

$$U = \frac{a\boldsymbol{B} \star w\boldsymbol{T}}{a\boldsymbol{B} \star \boldsymbol{B}^{\dagger}w}, \tag{A.1}$$

where $\boldsymbol{B}$ is the (multiscale) filtering beam function, $\boldsymbol{T}$ is the temperature, a and w their respective weights, and $\star$ denotes the usual convolution operation

$$\{a\boldsymbol{B} \star w\boldsymbol{T}\}(\boldsymbol{\xi}) = \sum_{\boldsymbol{x}} a(\boldsymbol{x})\boldsymbol{B}(\boldsymbol{x}) \cdot w(\boldsymbol{\xi}-\boldsymbol{x})\boldsymbol{T}(\boldsymbol{\xi}-\boldsymbol{x}). \tag{A.2}$$

In the absence of a specific model for the non-Gaussian signal, the beam functions can be taken to be orthogonal polynomials on a disc, weighted by some smoothing function, while the weights applied to the temperature maps are determined by the CMB and noise covariance.

In a simple approach, the information about the CMB signal can be utilized by pre-whitening the map by convolving it with an isotropic beam function $w_{\ell} = C_{\ell}^{-1/2}$ derived from the isotropic best-fit CMB power spectrum combined with a diagonal approximation to the instrumental noise covariance. After the component-separated CMB maps are pre-whitened, and the corresponding mask is applied to the resulting map, the multiscale filtering kernel b_{ℓ} is applied at various scales.

In this paper, the maps are pre-whitened with the 2013 best-fit cosmological parameter CMB spectrum (Planck Collaboration XV 2014), co-added to an isotropic noise power spectrum derived from the half-mission, half-difference noise maps appropriate for each component-separation method. No adjustment is made either for the recalibration of the 2015 data relative to the nominal results that the cosmological spectrum is derived from, or for the mismatch in noise level between the half-mission, half-difference and full-mission maps. This implies that the filtering is sub-optimal, but the data and simulations are treated consistently so there should be no significant impact on the results. The resulting pre-whitening beam function w_{ℓ} for the SMICA temperature map is shown in Fig. A.1.

The peak detector wavelets are taken to be Savitzky-Golay polynomials (Savitzky & Golay 1964), generalized to be defined on a disc with a polynomial smoothing weight function applied, as shown in Fig. A.1. A generalized spherical Savitzky-Golay kernel of order n and smoothing weight k (referred to as SSGnk in the text) is defined by a polynomial function of a radial variable $x = \sin(\theta/2)/\sin(\theta_{\max}/2)$,

$$F_{n,k}(x) = \left(\sum_{i=0}^{n/2} a_i x^{2i}\right)\left(1 - x^2\right)^k, \tag{A.3}$$

which is normalized to have unit mean on a disc and is orthogonal to all non-constant polynomials up to order n,

$$\int_0^1 x F_{n,k}(x)\, dx = 1, \quad \int_0^1 x^{i+1} F_{n,k}(x)\, dx = 0. \tag{A.4}$$

These are essentially high-order low-pass filters in harmonic space, but have compact support on the sphere. A few representative Savitzky-Golay polynomials are compared to a Gaussian kernel in Fig. A.1. Combined with pre-whitening, the total effect of the filters applied is described by the composite beam functions shown in Fig. A.1.

One should note a slight ℓ-space bandwidth mismatch between differently shaped kernels with the same FWHM value in real space, which is clear from the lower left panel of Fig. A.1. While not a problem in general, some care should be exercised when directly comparing results for different shape kernels. In particular, the ℓ value at which the filter kernel coefficient reaches $b_{\ell} = b_{\max}/2$ differs by a factor of 1.58 between the GAUSS and SSG84 kernels of the same FWHM.

Appendix B: Doppler boosting

The main effect of our relative motion with respect to the CMB rest frame is a dominant contribution to the CMB dipole (C_1); this is boosting of the monopole and has been detected previously (Kogut et al. 2003; Fixsen et al. 1996; Hinshaw et al. 2009). A subtler consequence of our motion is the boosting of all other multipoles. In fact, there are really two effects at work. The first is a modulation effect which increases power by approximately 0.25% in the direction of our motion and decreases it by the same amount in the opposite direction. This can equivalently be thought of as coupling between the multipoles ℓ and $\ell \pm 1$. The second is an aberration effect which shifts the apparent direction in which CMB photons arrive at our detectors toward the velocity direction.

Planck Collaboration XXVII (2014) reported a detection of this Doppler boosting, and an associated measurement of its velocity signature of 384 ± 78 (statistical) ± 115 (systematic) km s^{-1} in the known dipole direction, $(l, b) = (264°, 48°)$. Here, we demonstrate that the *Planck* 2015 data release remains in agreement with this result, by considering the angular scales $500 \leq \ell \leq 2000$. However, since the simulations employed in the analysis partially contain the effects of Doppler boosting (as noted in Sect. 3 the aberration contribution was erroneously omitted), we report a consistency check rather than a detection.

It is useful to perform a harmonic transform on the peculiar velocity vector,

$$\beta_{LM} = \int d\hat{\boldsymbol{n}}\, Y_{LM}^*(\hat{\boldsymbol{n}})\boldsymbol{\beta} \cdot \hat{\boldsymbol{n}}, \tag{B.1}$$

where only the $L = 1$ modes are non-zero. Following the convention in Planck Collaboration XXVII (2014), we rotate to an orthonormal basis, labelled $\boldsymbol{\beta}_{\parallel}$ (along the expected velocity direction), $\boldsymbol{\beta}_{\times}$ (parallel to the Galactic plane), and $\boldsymbol{\beta}_{\perp}$ (the remaining vector).

The peculiar velocity is detected using estimators that pick out the off-diagonal components of the CMB covariance matrix

$$\langle T_{\ell_1 m_1} T_{\ell_2 m_2}\rangle_{\mathrm{CMB}} = \sum_{LM} (-1)^M \begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & M \end{pmatrix}$$
$$\times \sqrt{\frac{(2\ell_1+1)(2\ell_2+1)(2L+1)}{4\pi}} W_{\ell_1\ell_2 L}^{\beta_v} \beta_{LM}. \tag{B.2}$$

The weight function W^{β_v} is a sum of the modulation ($b_v W^{\tau}$) and aberration (W^{ϕ}) effects. We quote results based on orthogonalized weight matrices,

$$W^{\hat{\phi}} = W^{\phi} - W^{\tau}\mathcal{R}^{\phi\tau}/\mathcal{R}^{\tau\tau} \tag{B.3}$$

$$W^{\hat{\tau}} = W^{\tau} - W^{\phi}\mathcal{R}^{\tau\phi}/\mathcal{R}^{\phi\phi}. \tag{B.4}$$
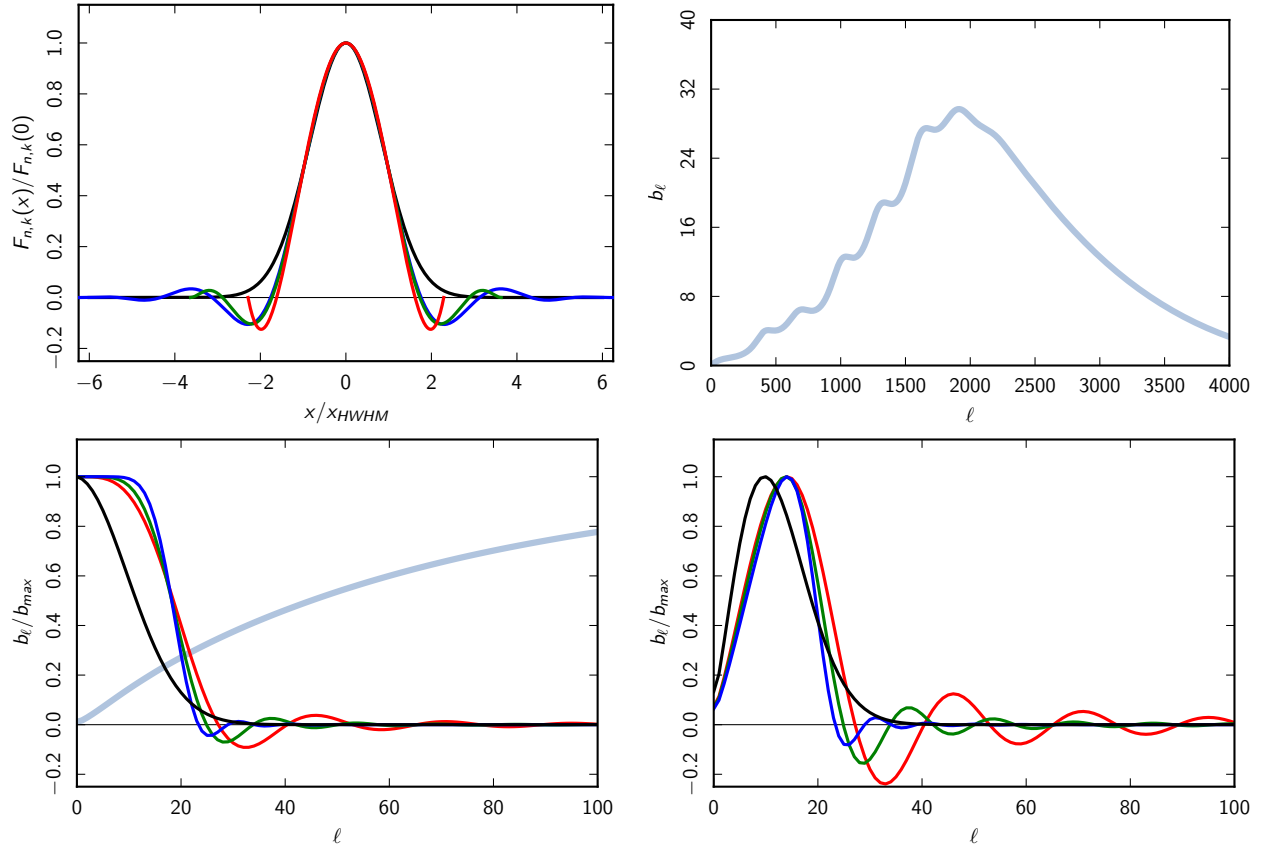
Fig. A.1. Generalized Savitzky-Golay polynomials are orthogonal to polynomials up to degree n on a disc, with smoothing weight applied. *Upper left panel* shows a few representative polynomial kernels (SSG21 in red, SSG42 in dark green, SSG84 in blue) and Gaussian (in black) as a function of radius (scaled to the same FWHM of 800′), *lower left* shows their harmonic space representation. *Right column* shows the pre-whitening kernel for the SMICA temperature map on the *top* (in light blue), and the corresponding composite kernels (WHITE*SSG21, etc.) on the *bottom* (in the same colours).

Table B.1. Significance measures for the β estimates for the 143×217 data set.

Estimator	χ^2	PTE [%]
$\hat{\beta}_{\parallel}$	3.28	7.01
$\hat{\beta}_{\perp}$	0.21	64.39
$\hat{\beta}_{\times}$	0.08	77.53
$\hat{\beta}$	3.38	33.70

Notes. χ^2 is formed from the three modes of β using the covariance matrix measured from Doppler boosted simulations.

Due to the clear connection between the velocity estimators and those used for the lensing analysis, we adopt the same data (143 GHz and 217 GHz sky maps, with dust foregrounds removed using the 857 GHz data as a template) and mask as used in Planck Collaboration XV (2016). The results summarized in Table B.1 show a slight excess of signal in the dipole direction of the data compared to simulations. This is due to the simulations used containing the modulation, but not aberration, part of the Doppler boost signal.

Appendix C: Generalized modulation estimator

Consider a parameter X that the (primary) CMB power spectrum is dependent on. Let X have a dipolar dependence of the form $X(\hat{\boldsymbol{n}}) = X_0 + \Delta X \hat{\boldsymbol{n}} \cdot \hat{\boldsymbol{m}}$ (this could correspond to a gradient in X across our observable volume), where X_0 is the average value,

$\hat{\boldsymbol{n}}$ is the direction to the last scattering surface, and $\hat{\boldsymbol{m}}$ is the gradient direction. To linear order in $\Delta X/X$, the measured spherical harmonics coefficients are given by

$$a_{\ell m} = a_{\ell m}^{\mathrm{iso}} + \sum_M \Delta X_M \sum_{\ell' m'} \frac{\mathrm{d} a_{\ell' m'}^{\mathrm{iso}}}{\mathrm{d} X} \xi_{\ell m \ell' m'}^M, \tag{C.1}$$

where the $a_{\ell m}^{\mathrm{iso}}$ are the unmodulated statistically isotropic modes. The $\xi_{\ell m \ell' m'}^M$ are coupling coefficients given by

$$\xi_{\ell m \ell' m'}^0 = \delta_{m'm} \left(\delta_{\ell' \ell - 1} A_{\ell - 1\, m} + \delta_{\ell' \ell + 1} A_{\ell m} \right), \tag{C.2}$$

$$\xi_{\ell m \ell' m'}^{\pm 1} = \delta_{m'm \mp 1} \left(\delta_{\ell' \ell - 1} B_{\ell - 1\, \pm m - 1} - \delta_{\ell' \ell + 1} B_{\ell\, \mp m} \right), \tag{C.3}$$

where

$$A_{\ell m} = \sqrt{\frac{(\ell + 1)^2 - m^2}{(2\ell + 1)(2\ell + 3)}}, \tag{C.4}$$

$$B_{\ell m} = \sqrt{\frac{(\ell + m + 1)(\ell + m + 2)}{2(2\ell + 1)(2\ell + 3)}}. \tag{C.5}$$

From Eq. (C.1) we can find the covariance matrix to first order in the components ΔX_M:

$$C_{\ell m \ell' m'} = C_\ell \delta_{\ell \ell'} \delta_{mm'} + \frac{\delta C_{\ell \ell'}}{2} \sum_M \Delta X_M \xi_{\ell m \ell' m'}^M, \tag{C.6}$$

where $\delta C_{\ell\ell+1} = dC_\ell/dX + dC_{\ell+1}/dX$. To determine the best-fit parameters, we proceed by maximizing the CMB likelihood function

$$\mathcal{L} = \frac{1}{\sqrt{2\pi|C|}} \exp(-\boldsymbol{d}^\dagger C^{-1} \boldsymbol{d}/2), \tag{C.7}$$

where $\boldsymbol{d}$ is the CMB temperature data. Equation (C.7) is maximized for the ΔX_M that satisfy

$$\boldsymbol{d}^\dagger C^{-1} \frac{dC}{d\Delta X_M} C^{-1} \boldsymbol{d} = \mathrm{Tr}\left[\left(C^{-1} \frac{dC}{d\Delta X_M}\right)\right]. \tag{C.8}$$

From Eq. (C.6) it is clear that the CMB covariance can be decomposed into an isotropic part (C_ℓ) and a small anisotropic part proportional to ΔX_M. By inverting Eq. (C.6) and using the orthogonality of the $\xi_{\ell m\ell'm'}^M$, we can determine the best-fit parameters

$$\Delta X_0 = \frac{6 \sum_{\ell m} \frac{\delta C_{\ell\ell+1}}{C_\ell C_{\ell+1}} A_{\ell m} a_{\ell m}^* a_{\ell+1\,m}}{\sum_\ell \frac{\delta C_{\ell\ell+1}^2}{C_\ell C_{\ell+1}}(\ell+1)}, \tag{C.9}$$

$$\Delta X_{+1} = \frac{6 \sum_{\ell m} \frac{\delta C_{\ell\ell+1}}{C_\ell C_{\ell+1}} B_{\ell m} a_{\ell m}^* a_{\ell+1\,m+1}}{\sum_\ell \frac{\delta C_{\ell\ell+1}^2}{C_\ell C_{\ell+1}}(\ell+1)}, \tag{C.10}$$

and $\Delta X_{-1} = -\Delta X_{+1}^*$, to first order in the anisotropy. These estimators are the full-sky, no-noise versions of Eqs. (44) and (45).

Errors can easily be found by expanding the log-likelihood about the best-fit parameters. The Fisher matrix is defined as

$$F_{MM'} \equiv \frac{1}{2}\mathrm{Tr}\left[\left(\frac{\partial C}{\partial \Delta X_M} C^{-1} \frac{\partial C}{\partial \Delta X_{M'}} C^{-1}\right)\right]. \tag{C.11}$$

Upon switching bases, we find

$$F_{0,0} = \frac{1}{4} \sum_{\ell m} \frac{\delta C_{\ell\ell+1}^2}{C_\ell C_{\ell+1}} A_{\ell m}^2, \tag{C.12}$$

$$F_{\Re(\Delta X_{+1}),\Re(\Delta X_{+1})} = \frac{1}{2} \sum_{\ell m} \frac{\delta C_{\ell\ell+1}^2}{C_\ell C_{\ell+1}} B_{\ell m}^2. \tag{C.13}$$

We can then assign the standard errors, $\sigma = \sqrt{F^{-1}}$.

Appendix D: Weighted-variance modified shape function estimator

The BipoSH representation characterizes the off-diagonal elements in the covariance matrix and is a generalization of the angular power spectrum, C_ℓ,

$$A_{\ell_1\ell_2}^{\mathrm{LM}} = \sum_{m_1 m_2} \langle a_{\ell_1 m_1} a_{\ell_2 m_2}\rangle C_{\ell_1 m_1 \ell_2 m_2}^{\mathrm{LM}}. \tag{D.1}$$

In general, it is not possible to analyse the full sky even for component-separated maps, due to the presence of residual contributions from diffuse Galactic emission and point sources. However, the application of a mask leads to coupling between the spherical harmonic modes. Hence, the correlation function is no longer described only by $C(\theta)$ or the power spectrum C_ℓ, and other quantities are required to completely quantify the statistical field.

We obtain an analytic expression for the observed BipoSH coefficients after the application of a mask in terms of the corresponding coefficients of the unmasked sky, and those of the mask itself,

$$\tilde{A}_{\ell_1\ell_2}^{\mathrm{LM}} = \sum_{\ell_3\ell_4} \frac{\Pi_{\ell_3\ell_4}}{\sqrt{4\pi}} \sum_{\ell_5\ell_6} \frac{\Pi_{\ell_5\ell_6}}{\sqrt{4\pi}} C_{\ell_3 0 \ell_5 0}^{\ell_1 0} C_{\ell_4 0 \ell_6 0}^{\ell_2 0}$$
$$\times \sum_{L_1 M_1 JK} \left\{ \begin{matrix} L & \ell_1 & \ell_2 \\ L_1 & \ell_3 & \ell_4 \\ J & \ell_5 & \ell_6 \end{matrix} \right\} \Pi_{L_1} \Pi_J A_{\ell_3\ell_4}^{L_1 M_1} W_{\ell_5\ell_6}^{JK} C_{L_1 M_1 JK}^{\mathrm{LM}} \tag{D.2}$$

where $\Pi_\ell = \sqrt{2\ell+1}$, $\tilde{A}_{\ell_1\ell_2}^{\mathrm{LM}}$ are the BipoSH coefficients of the masked sky map, $A_{\ell_1\ell_2}^{\mathrm{LM}}$ correspond to the BipoSH coefficients of the unmasked sky, $W_{\ell_1\ell_2}^{\mathrm{LM}}$ are the BipoSH coefficient of the mask itself, $C_{lml'm'}^{\mathrm{LM}}$ are the Clebsch-Gordon coefficients, and the term $\{\}$ in Eq. (D.2) is the $9j$−symbol. This quantifies the coupling between the BipoSH coefficients of the CMB sky map and those of the mask itself.

The underlying CMB sky may have deviations from statistical isotropy, as discussed in Sect. 6.4, due either to a dipole modulation ($L = 1$) of unknown origin, or to Doppler boosting ($L = 1$) of the temperature field. The BipoSH coefficients of such statistical isotropy-violating fields can be given by

$$A_{\ell_1\ell_2}^{\mathrm{LM}} = \bar{A}_{\ell_1\ell_2}^{\mathrm{LM}} + \phi_{\mathrm{LM}} G_{\ell_1\ell_2}^{\mathrm{L}}. \tag{D.3}$$

Here $\bar{A}_{\ell_1\ell_2}^{\mathrm{LM}}$ corresponds to the BipoSH coefficients of the unknown but statistically isotropic CMB field. This couples with BipoSH coefficients of the mask to introduce a mean field linear bias $\langle \mathcal{A}_{\ell_1\ell_2}^{\mathrm{LM}}\rangle_{\mathrm{mask}}$, which is estimated from simulations and subtracted from the BipoSH coefficients obtained from the masked sky. The ϕ_{LM} are the spherical harmonic coefficients of the field that breaks statistical isotropy, and $G_{\ell_1\ell_2}^{\mathrm{L}}$ is the shape function. Shape functions for dipole modulation and Doppler boosting are given in Eqs. (54) and (56), respectively.

Due to symmetries of the mask, which is largely defined by foreground residuals towards the Galactic plane, the dominant BipoSH modes of the mask correspond to $J = \{0, 2\}, K = 0$. Hence, for all practical purposes, signal is retained in the $L = 1$ mode itself, although masking modifies the shape function, now defined as the modified shape funtion in the rest of the text. A weighted variance modified shape function is defined as

$$\hat{\phi}_{\mathrm{LM}} = \sum_{\ell_1\ell_2} w_{\ell_1\ell_2}^{\mathrm{LM}} \frac{\hat{A}_{\ell_1\ell_2}^{\mathrm{LM}}}{K_{\ell_1\ell_2}^{\mathrm{LM}}}, \tag{D.4}$$

where $\hat{A}_{\ell_1\ell_2}^{\mathrm{LM}} = \tilde{A}_{\ell_1\ell_2}^{\mathrm{LM}} - \langle \mathcal{A}_{\ell_1\ell_2}^{\mathrm{LM}}\rangle_{\mathrm{mask}}$ and the weights are chosen such that $\sum_{\ell_1\ell_2} w_{\ell_1\ell_2}^{\mathrm{LM}} = 1$.

Here $K_{\ell_1\ell_2}^{\mathrm{LM}}$ is the MSF, which can be evaluated as

$$K_{\ell_1\ell_2}^{\mathrm{LM}} = \sum_{\ell_3\ell_4} \Pi_L G_{\ell_3\ell_4}^{\mathrm{L}} \frac{\Pi_{\ell_3\ell_4}}{\sqrt{4\pi}} \sum_{\ell_5\ell_6} \frac{\Pi_{\ell_5\ell_6}}{\sqrt{4\pi}} C_{\ell_3 0 \ell_5 0}^{\ell_1 0} C_{\ell_4 0 \ell_6 0}^{\ell_2 0}$$
$$\times \sum_{JK} \left\{ \begin{matrix} L & \ell_1 & \ell_2 \\ L & \ell_3 & \ell_4 \\ J & \ell_5 & \ell_6 \end{matrix} \right\} \Pi_J W_{\ell_5\ell_6}^{JK} C_{\mathrm{LM}JK}^{\mathrm{LM}}. \tag{D.5}$$

The weights are then given by

$$
w_{\ell_1\ell_2}^{\mathrm{LM}} = \frac{1}{\sum_M \left(\sigma_{\mathcal{A}_{\ell_1\ell_2}^{\mathrm{LM}}} / K_{\ell_1\ell_2}^{\mathrm{LM}}\right)^2}
$$
$$
\times \left[\sum_{\ell_1'\ell_2'} \frac{1}{\sum_M \left(\sigma_{\mathcal{A}_{\ell_1'\ell_2'}^{\mathrm{LM}}} / K_{\ell_1'\ell_2'}^{\mathrm{LM}}\right)^2}\right]^{-1}. \qquad (D.6)
$$