

Planck 2015 results

XVI. Isotropy and statistics of the CMB

Planck Collaboration: P. A. R. Ade⁹⁰, N. Aghanim⁶⁰, Y. Akrami^{65,104}, P. K. Aluri⁵⁵, M. Arnaud⁷⁶, M. Ashdown^{72,6}, J. Aumont⁶⁰, C. Baccigalupi⁸⁹, A. J. Banday^{101,9,*}, R. B. Barreiro⁶⁷, N. Bartolo^{32,68}, S. Basak⁸⁹, E. Battaner^{102,103}, K. Benabed^{61,100}, A. Benoît⁵⁸, 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. Bonaldi⁷⁰, L. Bonavera⁶⁷, J. R. Bond⁸, J. Borrill^{14,95}, F. R. Bouchet^{61,93}, F. Boulanger⁶⁰, M. Bucher¹, C. Burigana^{48,33,50}, R. C. Butler⁴⁸, E. Calabrese⁹⁸, J.-F. Cardoso^{77,1,61}, B. Casaponsa⁶⁷, 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. Church⁹⁷, D. L. Clements⁵⁶, S. Colombi^{61,100}, L. P. L. Colombo^{25,69}, C. Combet⁷⁸, D. Contreras²⁴, F. Couchot⁷⁴, A. Coullais⁷⁵, B. P. Crill^{69,11}, M. Cruz²¹, A. Curto^{67,6,72}, F. Cuttaia⁴⁸, L. Danese⁸⁹, R. D. Davies⁷⁰, R. J. Davis⁷⁰, P. de Bernardis³⁴, A. de Rosa⁴⁸, G. de Zotti^{45,89}, J. Delabrouille¹, F.-X. Désert⁵⁴, J. M. Diego⁶⁷, H. Dole^{60,59}, S. Donzelli⁴⁹, O. Doré^{69,11}, M. Douspis⁶⁰, A. Ducout^{61,56}, X. Dupac³⁹, G. Efstathiou⁶⁴, F. Elsner^{26,61,100}, T. A. Enßlin⁸³, H. K. Eriksen⁶⁵, Y. Fantaye³⁷, J. Fergusson¹², R. Fernandez-Cobos⁶⁷, F. Finelli^{48,50}, O. Forni^{101,9}, M. Frailis⁴⁷, A. A. Fraisse²⁹, E. Franceschi⁴⁸, A. Frejsel⁸⁷, A. Frolov⁹², S. Galeotta⁴⁷, S. Galli⁷¹, K. Ganga¹, C. Gauthier^{1,82}, T. Ghosh⁶⁰, M. Giard^{101,9}, Y. Giraud-Héraud¹, E. Gjerløw⁶⁵, J. González-Nuevo^{20,67}, K. M. Górski^{69,105}, S. Gratton^{72,64}, A. Gregorio^{36,47,53}, A. Gruppuso⁴⁸, J. E. Gudmundsson²⁹, F. K. Hansen⁶⁵, D. Hanson^{84,69,8}, D. L. Harrison^{64,72}, S. Henrot-Versillé⁷⁴, C. Hernández-Monteagudo^{13,83}, D. Herranz⁶⁷, S. R. Hildebrandt^{69,11}, E. Hivon^{61,100}, M. Hobson⁶, W. A. Holmes⁶⁹, A. Hornstrup¹⁷, W. Hovest⁸³, Z. Huang⁸, K. M. Huffenberger²⁷, G. Hurier⁶⁰, A. H. Jaffe⁵⁶, T. R. Jaffe^{101,9}, W. C. Jones²⁹, M. Juvela²⁸, E. Keihänen²⁸, R. Keskitalo¹⁴, J. Kim⁸³, T. S. Kisner⁸⁰, J. Knoche⁸³, M. Kunz^{18,60,3}, H. Kurki-Suonio^{28,44}, G. Lagache^{5,60}, A. Lähteenmäki^{2,44}, J.-M. Lamarre⁷⁵, A. Lasenby^{6,72}, M. Lattanzi³³, C. R. Lawrence⁶⁹, R. Leonardi³⁹, J. Lesgourgues^{62,99}, F. Levrier⁷⁵, M. Liguori^{32,68}, P. B. Lilje⁶⁵, M. Linden-Vørnle¹⁷, H. Liu^{87,38}, M. López-Cañiego^{39,67}, P. M. Lubin³⁰, J. F. Macías-Pérez⁷⁸, G. Maggio⁴⁷, D. Maino^{35,49}, N. Mandolesi^{48,33}, A. Mangilli^{60,74}, D. Marinucci³⁷, M. Maris⁴⁷, P. G. Martin⁸, E. Martínez-González⁶⁷, S. Masi³⁴, S. Matarrese^{32,68,42}, P. McGehee⁵⁷, P. R. Meinhold³⁰, A. Melchiorri^{34,51}, L. Mendes³⁹, A. Mennella^{35,49}, M. Migliaccio^{64,72}, K. Mikkelsen⁶⁵, S. Mitra^{55,69}, M.-A. Miville-Deschênes^{60,8}, D. Molinari^{67,48}, A. Moneti⁶¹, L. Montier^{101,9}, G. Morgante⁴⁸, D. Mortlock⁵⁶, A. Moss⁹¹, D. Munshi⁹⁰, J. A. Murphy⁸⁵, P. Naselsky^{87,38}, F. Nati²⁹, P. Natoli^{33,4,48}, C. B. Netterfield²², H. U. Nørgaard-Nielsen¹⁷, F. Novello⁷⁰, D. Novikov⁸¹, I. Novikov^{87,81}, C. A. Oxborrow¹⁷, F. Paci⁸⁹, L. Pagano^{34,51}, F. Pajot⁶⁰, N. Pant⁵⁵, D. Paoletti^{48,50}, F. Pasian⁴⁷, G. Patanchon¹, T. J. Pearson^{11,57}, O. Perdereau⁷⁴, L. Perotto⁷⁸, F. Perrotta⁸⁹, V. Pettorino⁴³, F. Piacentini³⁴, M. Piat¹, E. Pierpaoli²⁵, D. Pietrobon⁶⁹, S. Plaszczynski⁷⁴, E. Pointecouteau^{101,9}, G. Polenta^{4,46}, L. Popa⁶³, G. W. Pratt⁷⁶, G. Prézeau^{11,69}, S. Prunet^{61,100}, J.-L. Puget⁶⁰, J. P. Rachen^{23,83}, R. Rebolo^{66,15,19}, M. Reinecke⁸³, M. Remazeilles^{70,60,1}, C. Renault⁷⁸, A. Renzi^{37,52}, I. Ristorcelli^{101,9}, G. Rocha^{69,11}, C. Rosset¹, M. Rossetti^{35,49}, A. Rotti⁵⁵, G. Roudier^{1,75,69}, J. A. Rubiño-Martín^{66,19}, B. Rusholme⁵⁷, M. Sandri⁴⁸, D. Santos⁷⁸, M. Savelainen^{28,44}, G. Savini⁸⁸, D. Scott²⁴, M. D. Seiffert^{69,11}, E. P. S. Shellard¹², T. Souradeep⁵⁵, L. D. Spencer⁹⁰, V. Stolyarov^{6,96,73}, R. Stompor¹, R. Sudiwala⁹⁰, R. Sunyaev^{83,94}, D. Sutton^{64,72}, A.-S. Suur-Uski^{28,44}, J.-F. Sygnet⁶¹, J. A. Tauber⁴⁰, L. Terenzi^{41,48}, L. Toffolatti^{20,67,48}, M. Tomasi^{35,49}, M. Tristram⁷⁴, T. Trombetti⁴⁸, M. Tucci¹⁸, J. Tuovinen¹⁰, L. Valenziano⁴⁸, J. Valiviita^{28,44}, B. Van Tent⁷⁹, P. Vielva⁶⁷, F. Villa⁴⁸, L. A. Wade⁶⁹, B. D. Wandelt^{61,100,31}, I. K. Wehus⁶⁹, D. Yvon¹⁶, A. Zacchei⁴⁷, J. P. Zibin²⁴, and A. Zonca³⁰

(Affiliations can be found after the references)

Received 5 June 2015 / Accepted 7 September 2015

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 peculiarities 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

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

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 $H_0 = 100h \text{ km s}^{-1} \text{ Mpc}^{-1}$ 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 $k = 0.05 \text{ Mpc}^{-1}$) $A_s = 2.09 \times 10^{-9}$, 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

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

where X_i is the observed temperature at pixel i and $\sigma_{i,\text{noise}}^2$ 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 $5'$, $N_{\text{side}} = 2048$. 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:

$$f(\mathbf{T}) = \frac{1}{(2\pi)^{N_{\text{pix}}/2} \det \mathbf{C}^{1/2}} \exp\left[-\frac{1}{2}(\mathbf{T}\mathbf{C}^{-1}\mathbf{T}^T)\right], \quad (3)$$

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

Although the calculation of $\mathbf{T}\mathbf{C}^{-1}\mathbf{T}^T$ can be achieved by conjugate gradient methods, the evaluation of $\det \mathbf{C}$ remains computationally difficult for the full *Planck* resolution at HEALPix $N_{\text{side}} = 2048$. 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 \mathbf{C} is fully defined by the *Planck* angular power spectrum (C_ℓ):

$$C_{ij} = \sum_{\ell=2}^{\ell_{\text{max}}} \frac{2\ell+1}{4\pi} C_\ell b_\ell^2 P_\ell(\cos \theta_{ij}), \quad (4)$$

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 $N_{\text{pix}} \gg 1$) a normal distribution $\mathcal{N}(N_{\text{pix}}, \sqrt{2N_{\text{pix}}})$.

These χ^2 statistics are computed for the *Planck* 2015 component-separated CMB maps at $N_{\text{side}} = 16$ 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 $\approx 1 \mu\text{K}$ 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.

N_{side}	Probability [%]			
	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,

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

where the unit vectors $\hat{n}_1, \dots, \hat{n}_N$ 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 $\theta_1, \dots, \theta_{2N-3}$ characterizing the shape and size of the polygon,

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

Pixel weights w_1^i, \dots, 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 Λ CDM 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.

Function	Probability [%]			
	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

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

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 k th simulated N -point correlation function and N_{sim} is the number of simulations, then the covariance matrix \mathbf{M}_{ij} is given by

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

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 Λ CDM 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.

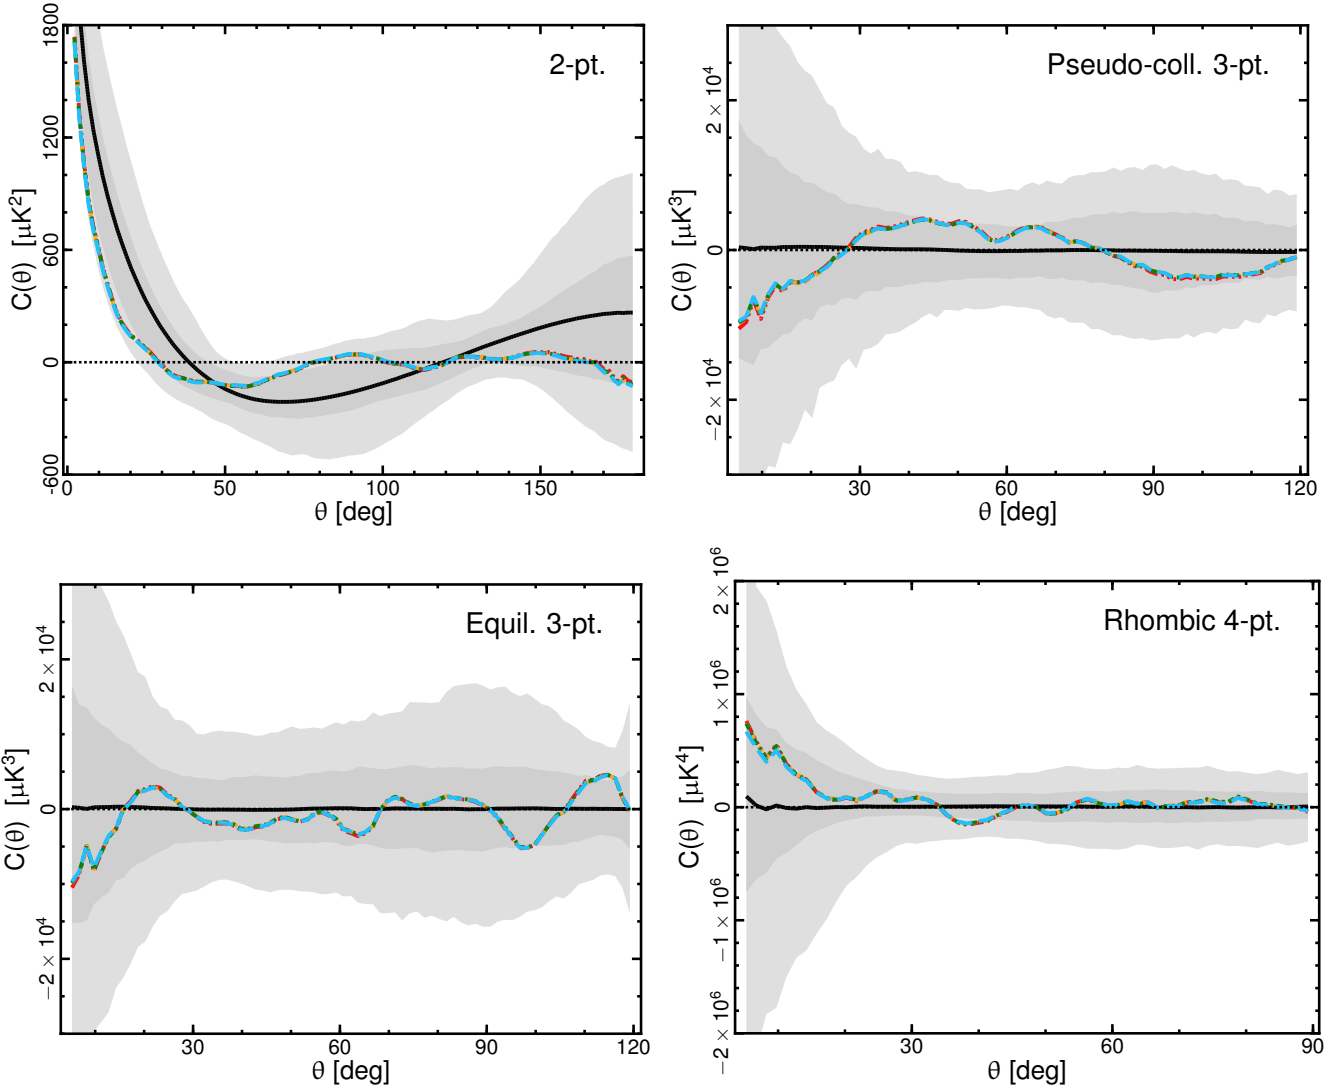


Fig. 2. N -point correlation functions determined from the $N_{\text{side}} = 64$ *Planck* CMB 2015 temperature maps. Results are shown for the 2-point, pseudo-collapsed 3-point (*upper left and right panels*, respectively), equilateral 3-point, and connected rhombic 4-point functions (*lower left and 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

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

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

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

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

$$V_k(\nu) = A_k v_k(\nu), \quad (12)$$

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

N_{side}	Probability [%]			
	Comm.	NILC	SEVM	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

with

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

$$v_3(\nu) = \frac{e^{-\nu^2}}{\text{erfc}(\nu/\sqrt{2})}, \quad (14)$$

and

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

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 :

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

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

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 $-3 \leq \nu \leq 3$ 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 $N_{\text{side}} = 1024$ and $N_{\text{side}} = 16$. We find no significant deviations from Gaussianity for any of the resolutions considered.

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

N_{side}	Probability [%]			
	Comm.	NILC	SEVM	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

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:

$$\begin{aligned} \beta_j(\hat{\mathbf{n}}) &= \sum_{\ell=B^{j-1}}^{B^{j+1}} b^2 \left(\frac{\ell}{B^j} \right) \sum_m a_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}) \\ &= \sum_{\ell=B^{j-1}}^{B^{j+1}} b^2 \left(\frac{\ell}{B^j} \right) T_\ell(\hat{\mathbf{n}}). \end{aligned} \quad (18)$$

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

$$T(\hat{\mathbf{n}}) = \sum_\ell T_\ell(\hat{\mathbf{n}}), \quad (19)$$

where $\hat{\mathbf{n}} \in S^2$ denotes the pointing direction, B is a fixed parameter (usually taken to be between 1 and 2) and $b(\cdot)$ 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

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

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

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

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

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

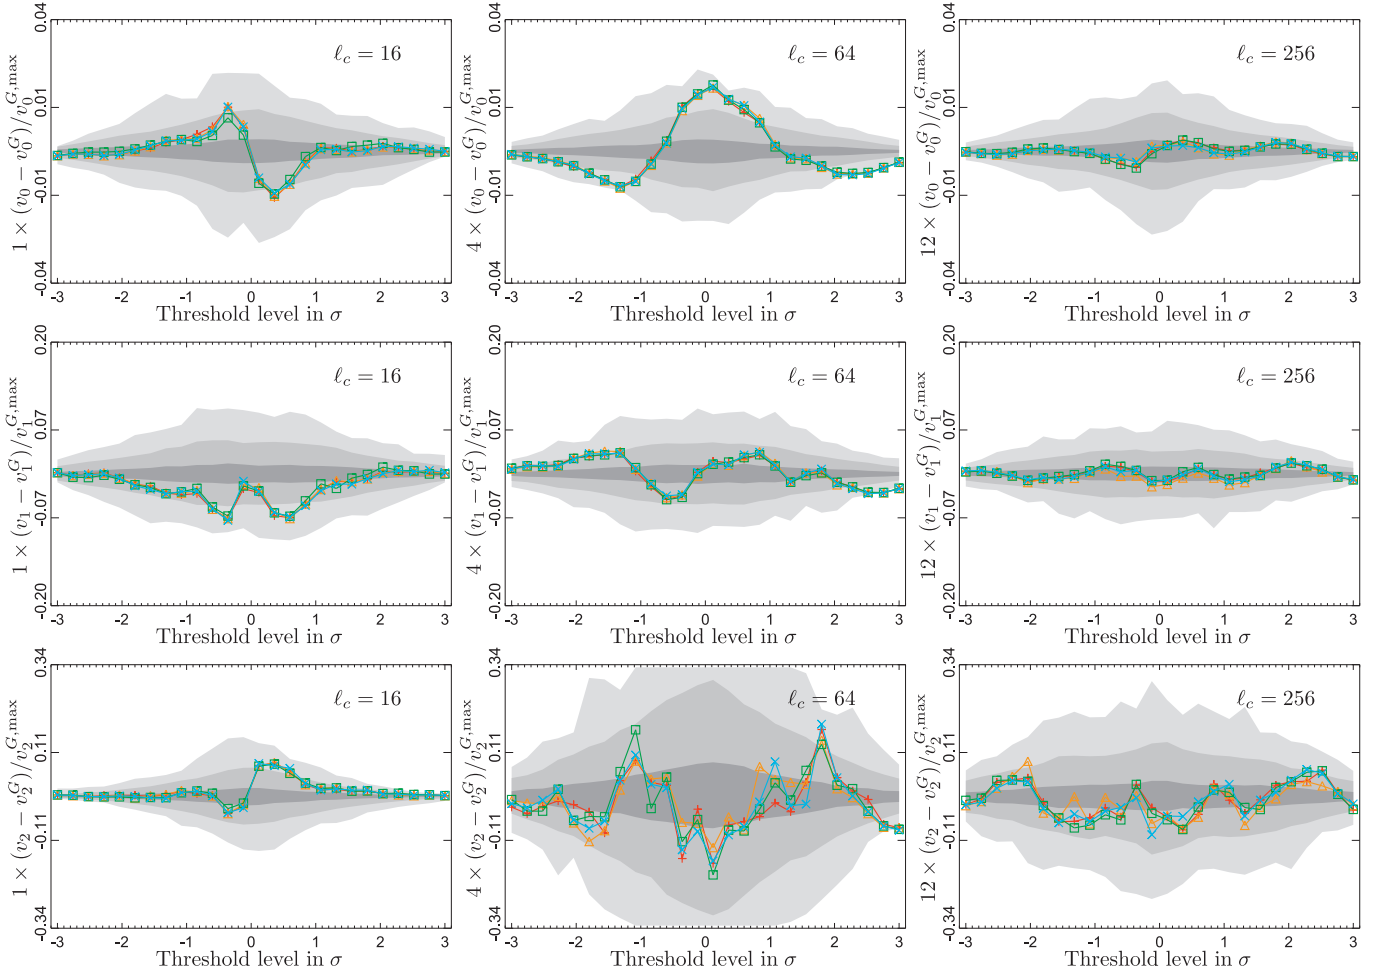


Fig. 3. Needle space MFs for *Planck* 2015 data using the four component-separated maps, Commander (red), NILC (orange), SEVEM (green), and SMICA (blue); the grey regions, from dark to light, correspond, respectively, to 1, 2, and 3 σ confidence regions estimated from the 1000 FFP8 simulations processed by the Commander method. The columns *from left to right* correspond to the needle parameters $j = 4, 6,$ and $8,$ respectively; the j th needle 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 needlelet map. Note that to have the same range at all the needlelet 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, interpretation 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 needlelet 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 needlelet MFs are more sensitive to the shape of the power spectrum than the corresponding all-scale MFs.

The needlelet 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 needlelet 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 needlelet 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 needlelet 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 needlelet MFs from the Commander, NILC, SEVEM, and SMICA foreground-cleaned temperature maps for needlelet 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 $P(\chi^2 > \chi^2_{\text{Planck}})$ as a function of needlelet 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 needlelet scales. The j th needlelet 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 appropriate 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

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

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

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

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 continuous 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 \mathcal{S}_R , and the kurtosis \mathcal{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 SSG84³) 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 ($\approx 0.8\%$) at scales of around $300'$; (ii) that the dispersion of the wavelet coefficients at these scales and at around $700'$ is relatively low ($\approx 1\%$); and (iii) that the dispersion of the wavelet coefficients at scales below $5'$ is significantly high ($\leq 0.1\%$).

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

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

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

where the operator $\#$ represents the number of pixels p in which the condition defined between the braces is satisfied.

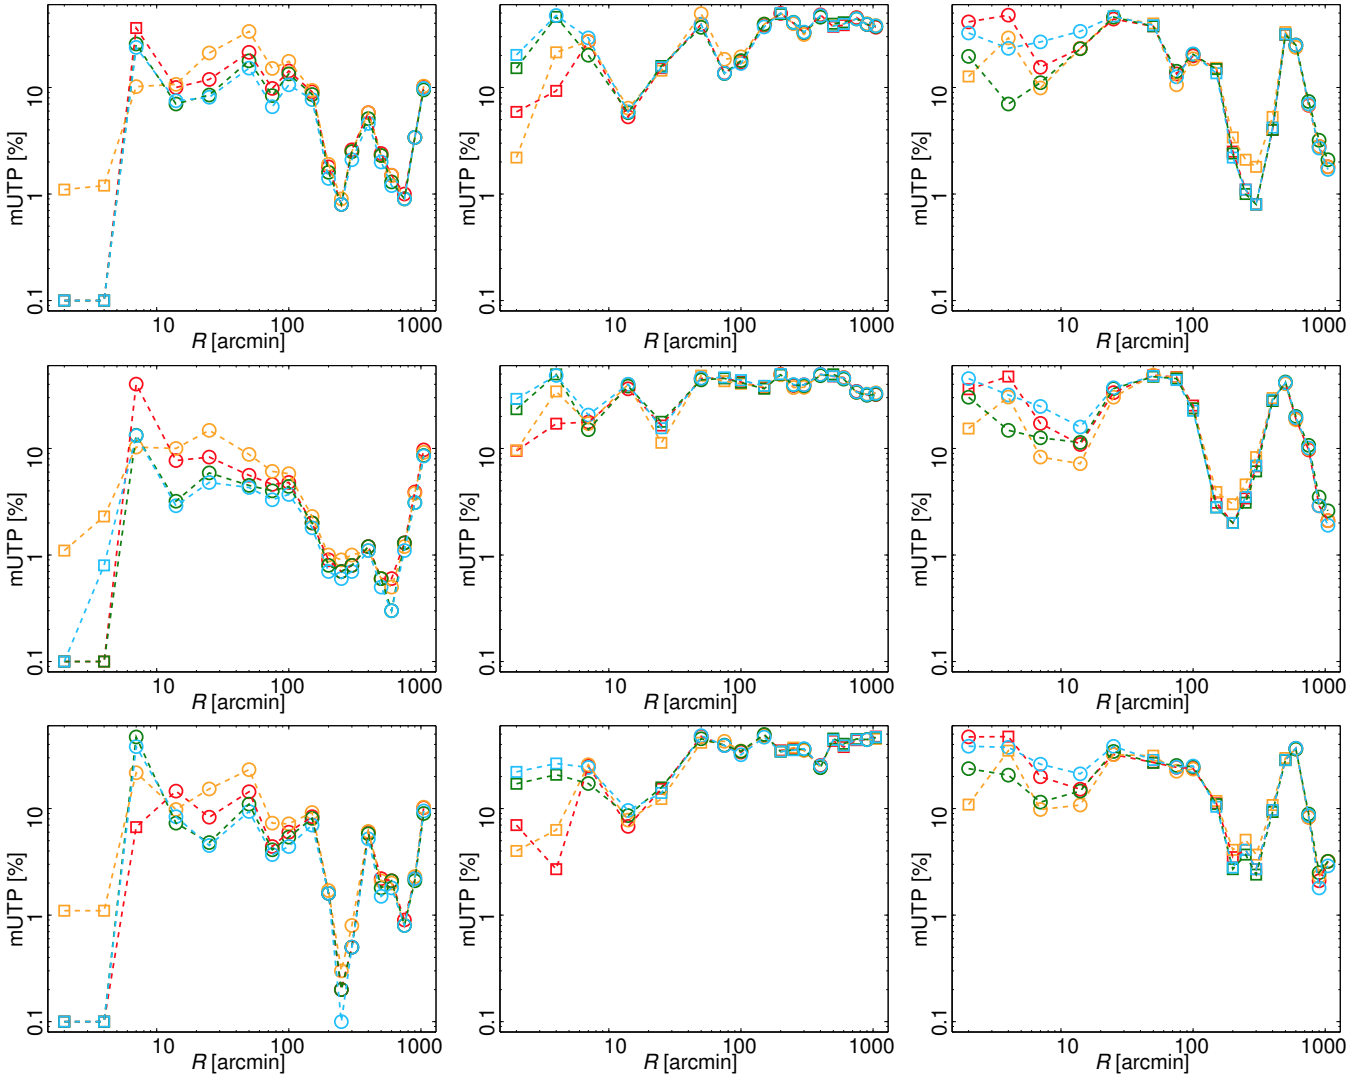


Fig. 6. Modified upper tail probabilities (mUTP) obtained from the analyses of the filter coefficients as a function of the filter scale R for the Commander (red), NILC (orange), SEVEM (green), and SMICA (blue) sky maps. From left to right, the panels correspond to standard deviation, skewness, and kurtosis 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

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

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 $R = 400'$.

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

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 $80'$ 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:

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

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.

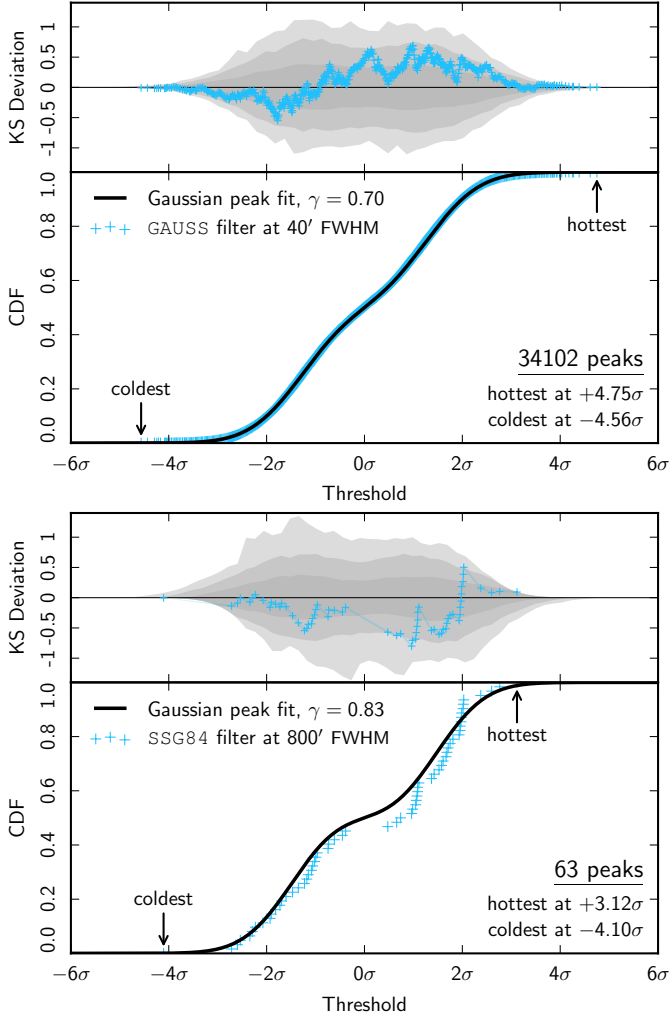


Fig. 10. Cumulative density function of the peak distribution for the SMICA CMB temperature map. The *top row* shows the peak CDF for maps filtered with a GAUSS kernel of 40' FWHM. The *bottom row* shows the corresponding peak CDF for an SSG84 kernel of 800' FWHM. The spectral shape parameter γ (see 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

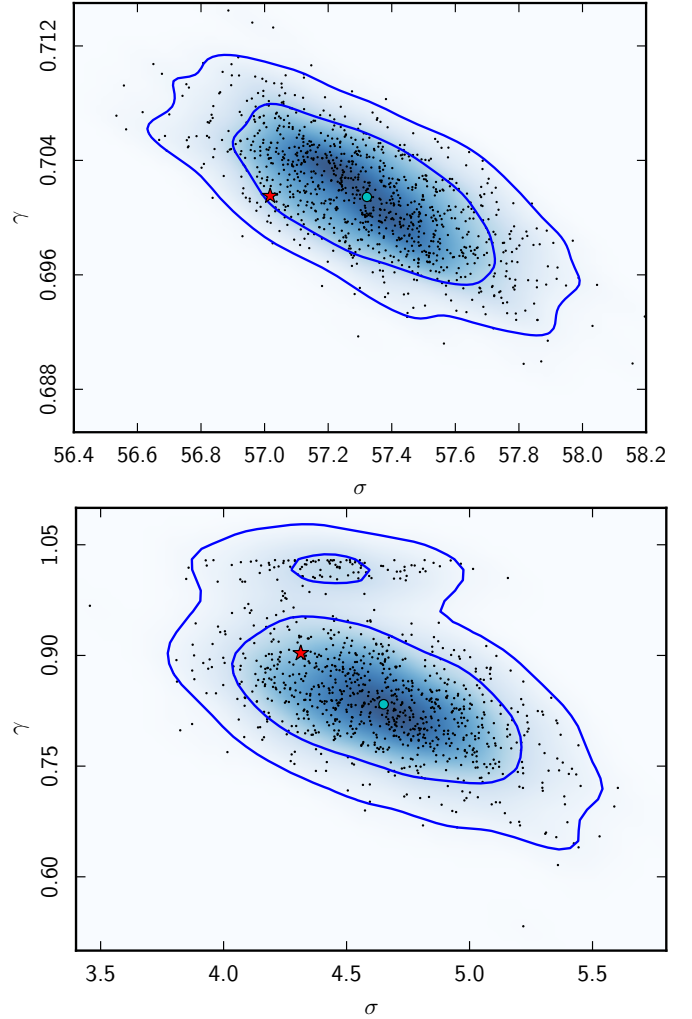


Fig. 11. Distribution of best-fit Gaussian peak CDF spectral shape parameters, σ and γ (as defined in Eq. (27)), recovered from 1000 simulations, as indicated by the black dots and the smoothed density map and compared to those derived for the observed sky (shown by the red star). The blue contours enclose 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

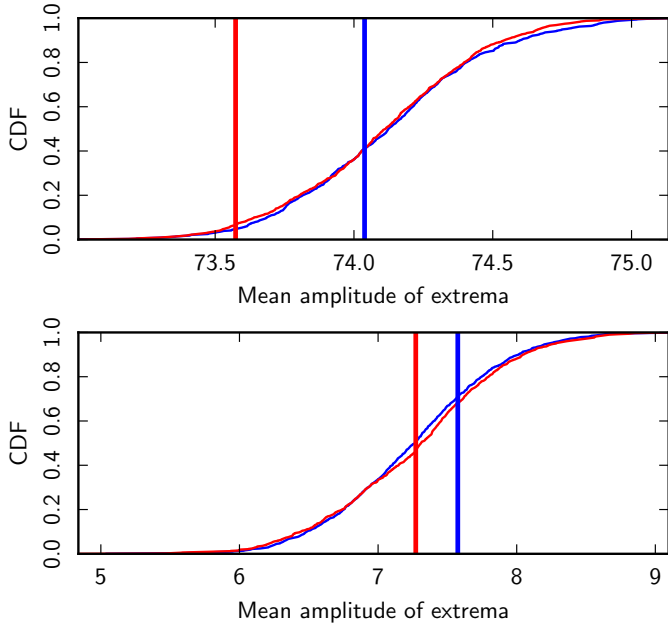


Fig. 12. Cumulative density function of the mean amplitude of all extrema, maxima (red) and minima (blue), derived from simulations, compared to the equivalent values observed for the SMICA CMB temperature map. The *upper panel* shows the peak mean amplitudes for maps filtered with a GAUSS kernel of 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 $x_{\text{coldest}} < x_{\text{coldest}}^{(\text{sky})}$. 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).

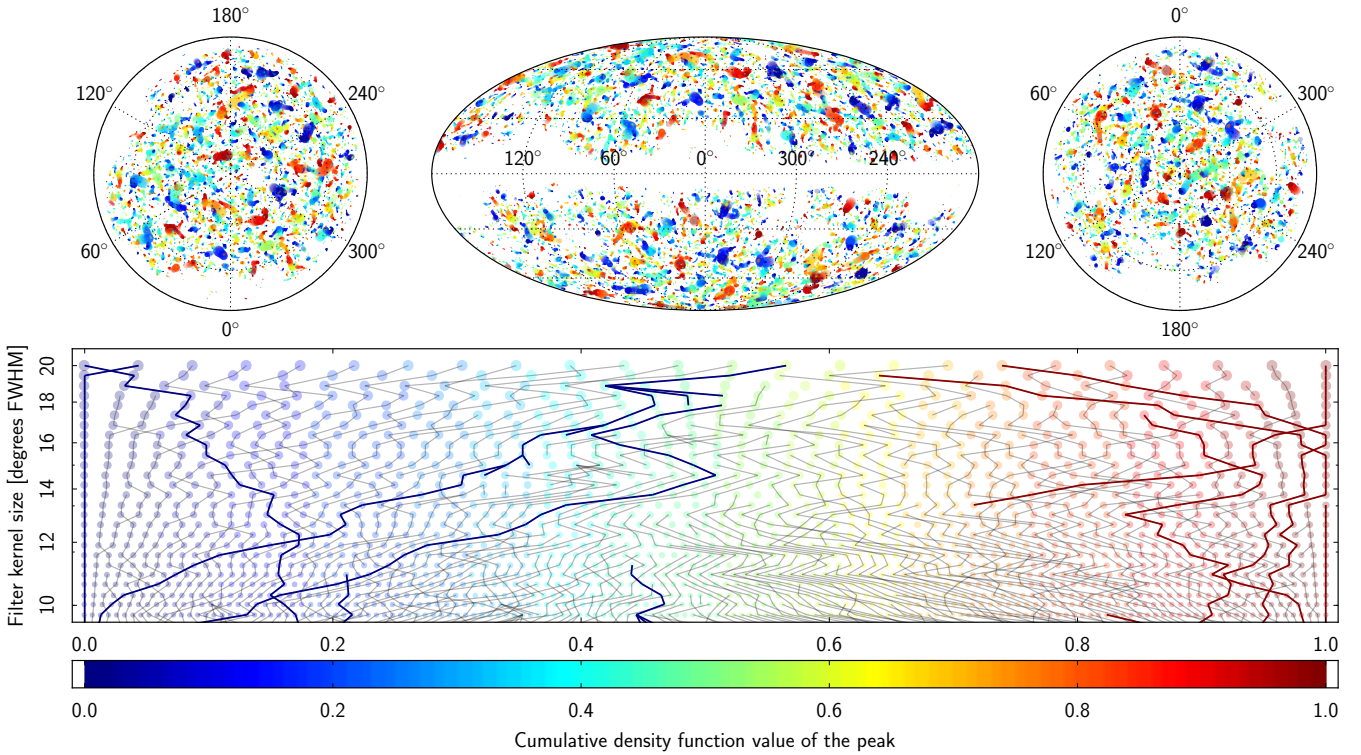


Fig. 14. Peak positions and CDF rank summarized for all filtering scales. The three sky-view panels in the *top row* show a Lambert projection of the north pole, the usual full sky Mollweide projection, and a Lambert projection of the south pole. The *lower panel* shows the peak heights (in percentile 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 $\beta_{||}$ 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 Λ CDM 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.

unit variance estimator to lower resolution component-separated maps, specifically those from $N_{\text{side}} = 1024$ to $N_{\text{side}} = 16$, 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 $N_{\text{side}} = 16$. 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_{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 $N_{\text{side}} = 16$ with respect to the possible presence of residual foregrounds by considering two additional masks obtained by extending the edge of the $N_{\text{side}} = 16$ 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 Λ CDM 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

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

Method	Probability [%]		
	Variance	Skewness	Kurtosis
Common mask ($f_{\text{sky}} = 58\%$)			
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
$f_{\text{sky}} = 48\%$			
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
$f_{\text{sky}} = 40\%$			
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

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):

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

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 $x \equiv \cos 60^\circ = \frac{1}{2}$. 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 Λ CDM 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 maps.

Statistic	Probability [%]			
	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
$\chi_0^2(\theta > 60^\circ)$	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

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

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 $\theta_{\min} = 60^\circ$ and $\theta_{\max} = 180^\circ$. \mathbf{M}_{ij} is the covariance matrix given by Eq. (8), estimated using MC simulations corresponding to the fiducial Λ CDM 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 $\theta > 60^\circ$. A conventional χ^2 statistic allows us to test the consistency of this quantity with the predictions of the Λ CDM 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 $\theta < 60^\circ$ and $\theta > 60^\circ$. The results of these studies are shown in Table 14.

The analysis for $\theta < 60^\circ$ indicates that the observed 2-point function is a good match to the mean 2-point function predicted by the Λ CDM model. Moreover, for $\theta > 60^\circ$ 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 Λ CDM model for angles above 60° . This problem may be

Table 14. Probabilities of obtaining values for the χ^2 statistic for the *Planck* fiducial Λ CDM 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 maps.

Statistic	Probability [%]			
	Comm.	NILC	SEVEM	SMICA
$\chi^2(\theta < 60^\circ)$	91.5	93.3	91.6	91.7
$\chi^2(\theta > 60^\circ)$	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 $N_{\text{side}} = 64$ 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 Λ CDM 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 Λ CDM 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.

Hemisphere	Probability [%]			
	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

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).

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 Λ CDM 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.

Hemisphere	Probability [%]		
	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

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

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

Given such a spectrum, the CMB temperature field is expected to exhibit a correlation between $a_{\ell m}$ and $a_{\ell \pm \Delta m'}^*$ 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} :

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

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

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

Here, $\langle (\mathbf{C}^{-1} \mathbf{a}^*)_{\ell m} (\mathbf{C}^{-1} \mathbf{a})_{\ell' m'} \rangle_{g_{2M}=0}$ 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:

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

where C_ℓ and N_ℓ are the signal and noise power spectra respectively. Uncertainties are computed by applying the estimator to simulations.

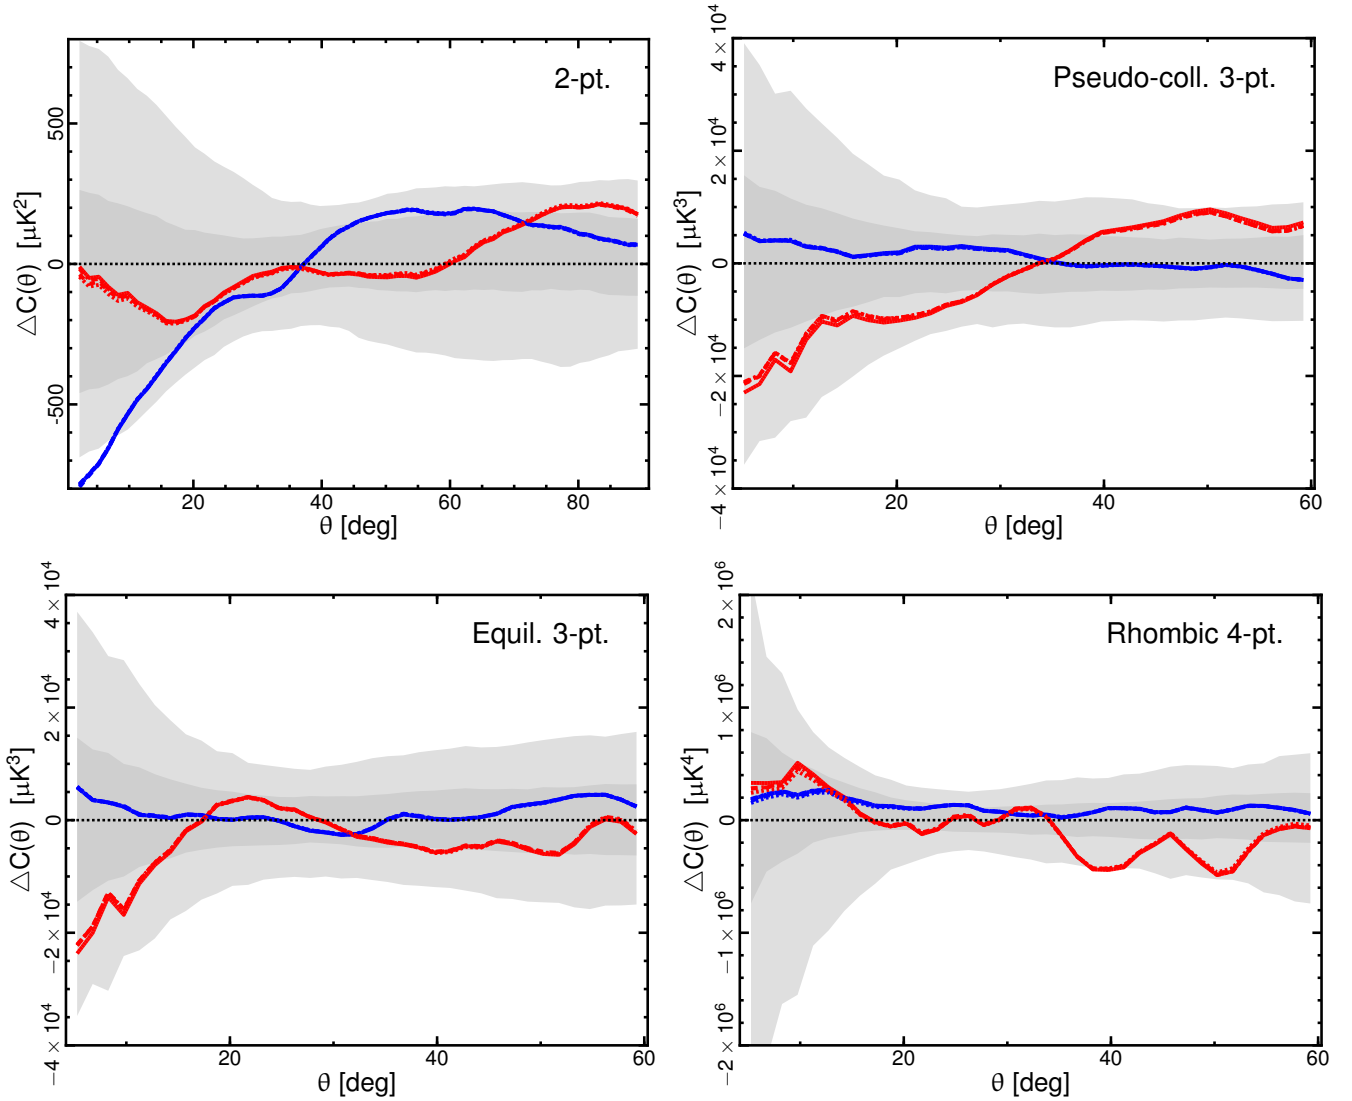


Fig. 17. Difference of the N -point correlation functions determined from the $N_{\text{side}} = 64$ *Planck* CMB 2015 temperature estimates and the corresponding means estimated from 1000 simulations. Results are shown for the 2-point, pseudo-collapsed 3-point (*upper left and right panels*, respectively), 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

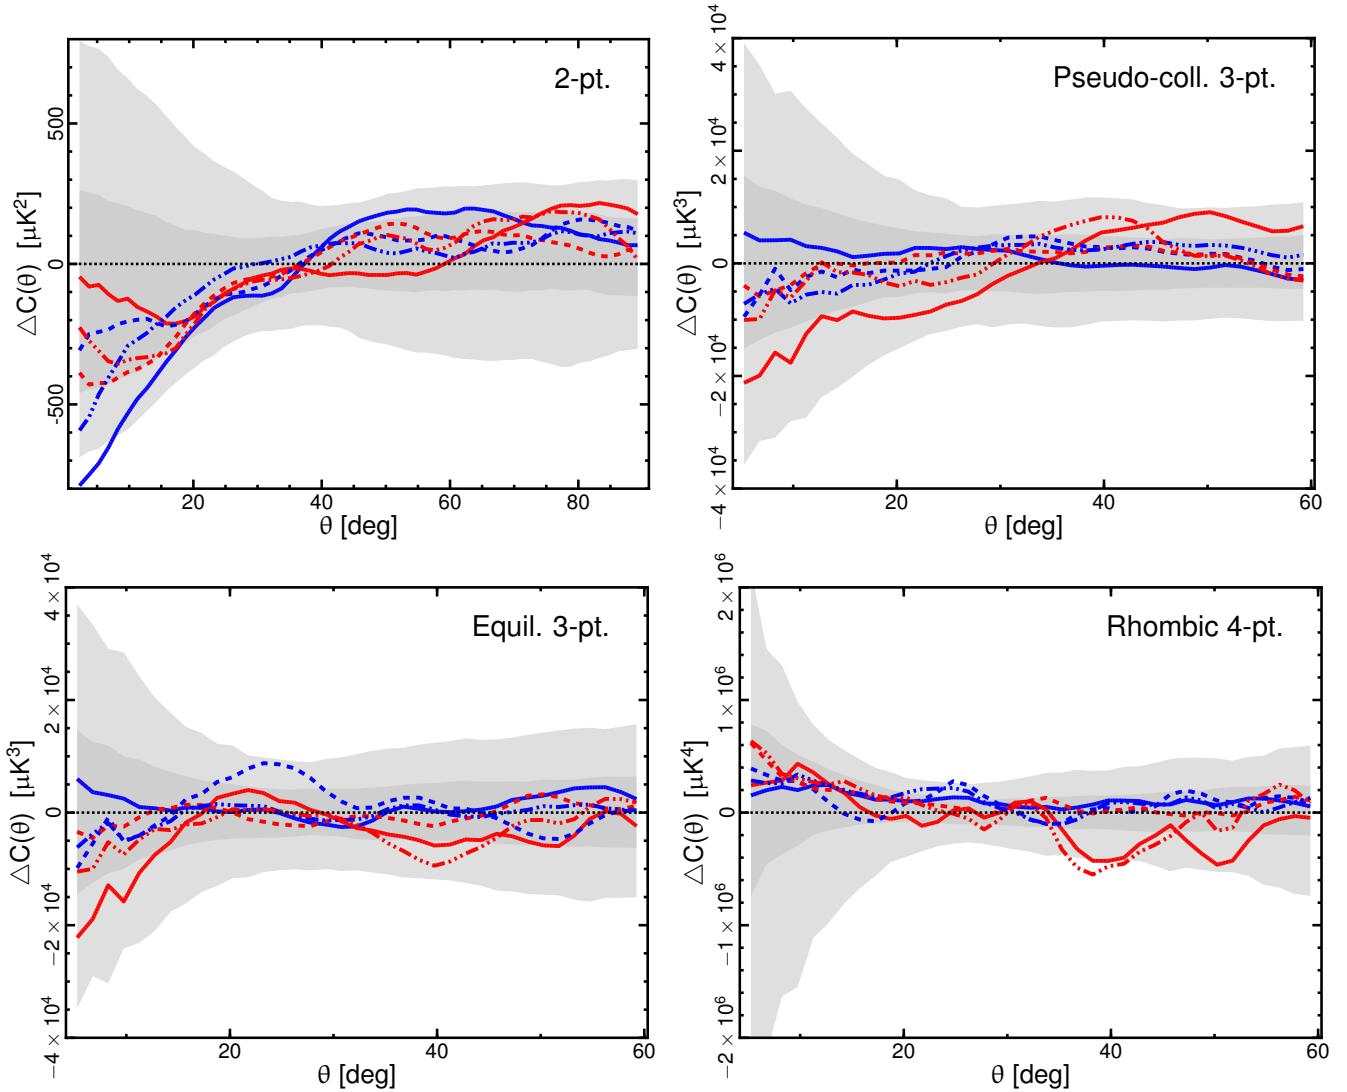


Fig. 18. Difference of the N -point correlation functions determined from the $N_{\text{side}} = 64$ *Planck* SMICA CMB 2015 temperature estimates and the corresponding means estimated from 1000 simulations. Results are shown for the 2-point, pseudo-collapsed 3-point (*upper left and right panels*, respectively), 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 $g_2 = (1/5 \sum_M |g_{2M}|^2)^{1/2}$. 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, $T(\hat{n})$, may be divided into symmetric, $T^+(\hat{n})$, and antisymmetric, $T^-(\hat{n})$, 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:

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

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

$$D_{+,-}^{\text{TT}} = \frac{1}{\ell_{\text{tot}}^{+,-}} \sum_{\ell=2, \ell_{\text{max}}}^{+,-} \frac{\ell(\ell+1)}{2\pi} C_{\ell}^{\text{TT}}, \quad (37)$$




Fig. 19. Ratio $R^{\text{TT}}(\ell_{\text{max}})$ for Commander (red), NILC (orange), SEVEM (green), and SMICA (blue) determined at $N_{\text{side}} = 32$. 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.

$\ell_{\text{tot}}^{+,-}$ 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 Λ CDM model.

Figure 19 presents the ratio, $R^{\text{TT}}(\ell_{\text{max}})$, 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 $\ell_{\text{max}} = 20\text{--}30$. 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 Commander⁴.

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%.

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.




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

5.5. Mirror-parity asymmetry

For the *Planck* 2013 data release, we studied the properties of the temperature data at a resolution of $N_{\text{side}} = 16$ 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 $(l, b) = (264^\circ, -17^\circ)$. 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 $N_{\text{side}} = 16$ and $N_{\text{side}} = 32$. In order to avoid possible bias introduced by the use of the Galactic mask⁵ 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:

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

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.,

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

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 $N_{\text{side}} = 16$ and $N_{\text{side}} = 32$.

Estimator	Probability [%]	Direction (l, b) [$^\circ$]
$N_{\text{side}} = 16$		
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)
$N_{\text{side}} = 32$		
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 $N_{\text{side}} = 16$ (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 $N_{\text{side}} = 16$ exhibits the most anomalous mirror antisymmetry in the direction $(l, b) = (264^\circ, -17^\circ)$, 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 $N_{\text{side}} = 32$ yield approximately the same direction, $(l, b) = (264^\circ, -16^\circ)$, 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 $(l, b) = (260^\circ, 48^\circ)$, 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) statistics. The vertical lines show the minimum value for the estimator computed at $N_{\text{side}} = 32$ 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 $(l, b) = (225^\circ, -18^\circ)$ (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

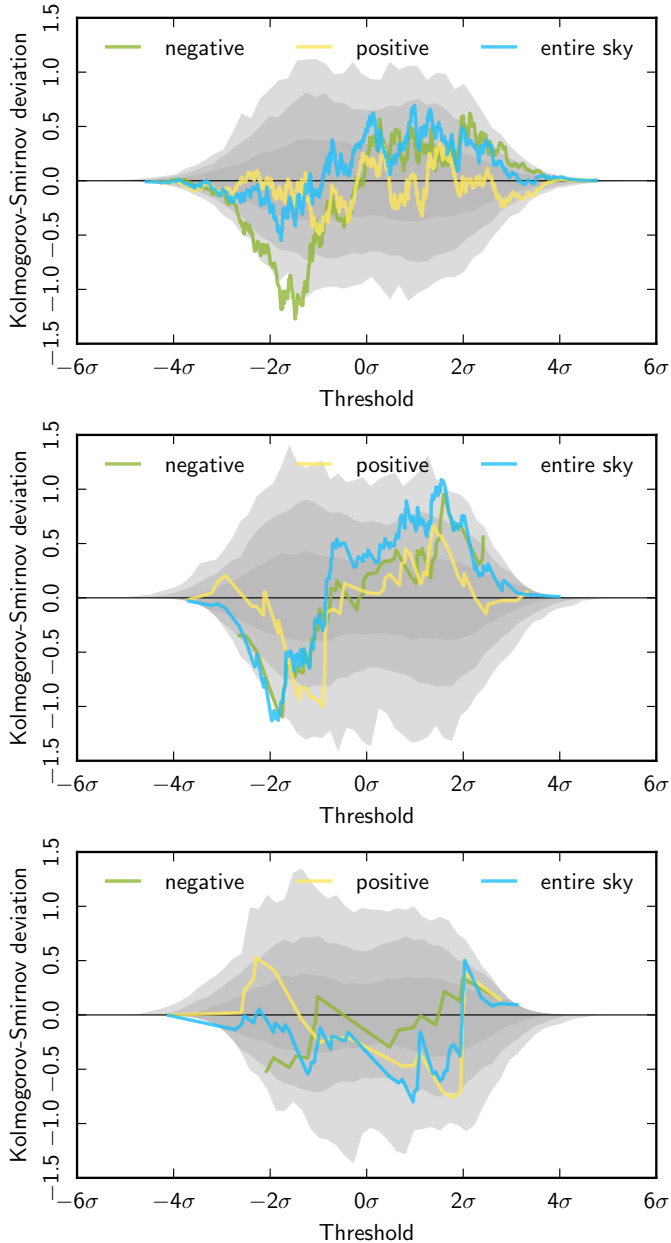


Fig. 22. KS-deviation of the peak distribution for 70° radius discs centred on the positive and negative asymmetry directions determined from the SMICA CMB temperature map in Sect. 6.2. From *top to bottom*, the plots correspond to maps filtered with a GAUSS kernel of $40'$ FWHM, an SSG84 filter of $500'$ FWHM, and an SSG84 filter of $800'$ FWHM, respectively.

for a set of discs, each of which is centred on a pixel defined at $N_{\text{side}} = 256$. 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 $N_{\text{side}} = 256$ sky maps. The derived $-\log_{10}(\text{UTP})$ 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

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

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

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

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; $-\log_{10}(\text{UTP})$ 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 $(l, b) = (210^\circ, -57^\circ)$ 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 $\approx 10^\circ$ 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 $-\log_{10}(\text{UTP})$ 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 $-\log_{10}(\text{UTP})$ 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 Λ CDM 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 $\approx 25^\circ$.

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 complementarity 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 $(l, b) \approx (212^\circ, -13^\circ)$ 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, $N_{\text{side}} = 2048$.

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 $N_{\text{side}} = 16$, the simulated asymmetry signals were not detected when either very small ($r_{\text{disc}} < 4^\circ$) or very large ($r_{\text{disc}} > 16^\circ$) 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 $N_{\text{side}} = 16$ 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

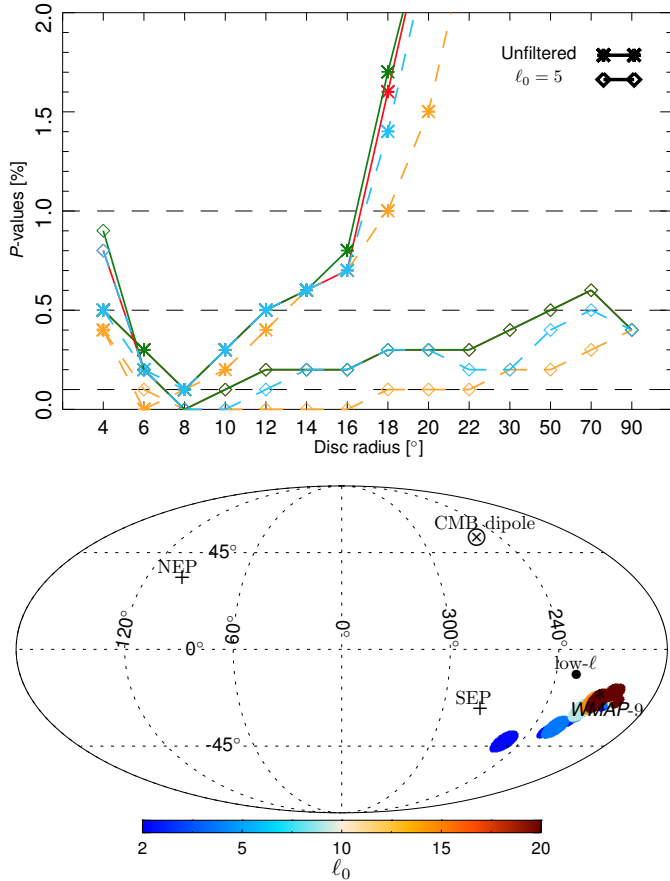


Fig. 27. *Upper panel:* p -values for variance asymmetry measured as the number of simulations with local-variance dipole amplitudes larger than those inferred from the data, as a function of disc radius for the four component-separated maps, Commander (red), NILC (orange), SEVEM (green), and SMICA (blue), and for unfiltered and high-pass-filtered cases. For the filtered case, the Commander curve is covered by the SMICA curve for small ($r_{\text{disk}} \leq 8$) discs, and by the SEVEM curve for large discs ($r_{\text{disk}} > 8$). *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,

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

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.

ℓ_0	p -value [%]			
	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 $N_{\text{side}} = 16$ 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 $N_{\text{side}} (=16)$ 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 $\ell_0 \geq 20$, 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.