



Publication Year	2016
Acceptance in OA @INAF	2020-05-27T13:58:47Z
Title	Planck 2015 results. XVII. Constraints on primordial non-Gaussianity
Authors	Planck Collaboration; Ade, P. A. R.; Aghanim, N.; Arnaud, M.; Arroja, F.; et al.
DOI	10.1051/0004-6361/201525836
Handle	http://hdl.handle.net/20.500.12386/25238
Journal	ASTRONOMY & ASTROPHYSICS
Number	594

Planck 2015 results. XVII. Constraints on primordial non-Gaussianity

Planck Collaboration: P. A. R. Ade⁹⁷, N. Aghanim⁶³, M. Arnaud⁷⁹, F. Arroja^{71,85}, M. Ashdown^{75,6}, J. Aumont⁶³, C. Baccigalupi⁹⁵, M. Ballardini^{51,53,34}, A. J. Banday^{109,10}, R. B. Barreiro⁷⁰, N. Bartolo^{33,71}, S. Basak⁹⁵, E. Battaner^{110,111}, K. Benabed^{64,108}, A. Benoît⁶¹, A. Benoit-Lévy^{26,64,108}, J.-P. Bernard^{109,10}, M. Bersanelli^{37,52}, P. Bielewicz^{89,10,95}, J. J. Bock^{72,12}, A. Bonaldi⁷³, L. Bonavera⁷⁰, J. R. Bond⁹, J. Borrill^{15,101}, F. R. Bouchet^{64,99}, F. Boulanger⁶³, M. Bucher¹, C. Burigana^{51,35,53}, R. C. Butler⁵¹, E. Calabrese¹⁰⁴, J.-F. Cardoso^{80,1,64}, A. Catalano^{81,78}, A. Challinor^{67,75,13}, A. Chamballu^{79,17,63}, H. C. Chiang^{30,7}, P. R. Christensen^{90,40}, S. Church¹⁰³, D. L. Clements⁵⁹, S. Colombi^{64,108}, L. P. L. Colombo^{25,72}, C. Combet⁸¹, F. Couchot⁷⁷, A. Coullais⁷⁸, B. P. Crill^{72,12}, A. Curto^{70,6,75}, F. Cuttaia⁵¹, L. Danese⁹⁵, R. D. Davies⁷³, R. J. Davis⁷³, P. de Bernardis³⁶, A. de Rosa⁵¹, G. de Zotti^{48,95}, J. Delabrouille¹, F.-X. Désert⁵⁷, J. M. Diego⁷⁰, H. Dole^{63,62}, S. Donzelli⁵², O. Doré^{72,12}, M. Douspis⁶³, A. Ducout^{64,59}, X. Dupac⁴², G. Efstathiou⁶⁷, F. Elsner^{26,64,108}, T. A. Enßlin⁸⁶, H. K. Eriksen⁶⁸, J. Fergusson¹³, F. Finelli^{51,53}, O. Forni^{109,10}, M. Frailis⁵⁰, A. A. Fraisse³⁰, E. Franceschi⁵¹, A. Frejse⁹⁰, S. Galeotta⁵⁰, S. Galli⁷⁴, K. Ganga¹, C. Gauthier^{1,85}, T. Ghosh⁶³, M. Giard^{109,10}, Y. Giraud-Héraud¹, E. Gjerløw⁶⁸, J. González-Nuevo^{21,70}, K. M. Górski^{72,112}, S. Gratton^{75,67}, A. Gregorio^{38,50,56}, A. Gruppuso⁵¹, J. E. Gudmundsson^{106,92,30}, J. Hamann^{107,105}, F. K. Hansen⁶⁸, D. Hanson^{87,72,9}, D. L. Harrison^{67,75}, A. Heavens⁵⁹, G. Helou¹², S. Henrot-Versillé⁷⁷, C. Hernández-Monteagudo^{14,86}, D. Herranz⁷⁰, S. R. Hildebrandt^{72,12}, E. Hivon^{64,108}, M. Hobson⁶, W. A. Holmes⁷², A. Hornstrup¹⁸, W. Hovest⁸⁶, Z. Huang⁹, K. M. Huffenberger²⁸, G. Hurier⁶³, A. H. Jaffe⁵⁹, T. R. Jaffe^{109,10}, W. C. Jones³⁰, M. Juvela²⁹, E. Keihänen²⁹, R. Keskitalo¹⁵, J. Kim⁸⁶, T. S. Kisner⁸³, J. Knoche⁸⁶, M. Kunz^{19,63,3}, H. Kurki-Suonio^{29,46}, F. Lacasa^{63,47}, G. Lagache^{5,63}, A. Lähteenmäki^{2,46}, J.-M. Lamarre⁷⁸, A. Lasenby^{6,75}, M. Lattanzi³⁵, C. R. Lawrence⁷², R. Leonardi⁸, J. Lesgourgues^{65,107}, F. Levrier⁷⁸, A. Lewis²⁷, M. Liguori^{33,71}, P. B. Lilje⁶⁸, M. Linden-Vørnle¹⁸, M. López-Cañiego^{42,70}, P. M. Lubin³¹, J. F. Macías-Pérez⁸¹, G. Maggio⁵⁰, D. Maino^{37,52}, N. Mandolesi^{51,35}, A. Mangilli^{63,77}, D. Marinucci³⁹, M. Maris⁵⁰, P. G. Martin⁹, E. Martínez-González⁷⁰, S. Masi³⁶, M.-A. Miville-Deschênes^{63,9}, A. Moneti⁶⁴, L. Montier^{109,10}, G. Morgante⁵¹, D. Mortlock⁵⁹, A. Moss⁹⁸, M. Münchmeyer⁶⁴, D. Munshi⁹⁷, J. A. Murphy⁸⁸, P. Naselsky^{91,41}, F. Nati³⁰, P. Natoli^{35,4,51}, C. B. Netterfield²², H. U. Nørgaard-Nielsen¹⁸, F. Noviello⁷³, D. Novikov⁸⁴, I. Novikov^{90,84}, C. A. Oxborrow¹⁸, F. Paci⁹⁵, L. Pagano^{36,54}, F. Pajot⁶³, D. Paoletti^{51,53}, F. Pasian⁵⁰, G. Patanchon¹, H. V. Peiris²⁶, O. Perdereau⁷⁷, L. Perotto⁸¹, F. Perrotta⁹⁵, V. Pettorino⁴⁵, F. Piacentini³⁶, M. Piat¹, E. Pierpaoli²⁵, D. Pietrobon⁷², S. Plaszczynski⁷⁷, E. Pointecouteau^{109,10}, G. Polenta^{4,49}, L. Popa⁶⁶, G. W. Pratt⁷⁹, G. Prézeau^{12,72}, S. Prunet^{64,108}, J.-L. Puget⁶³, J. P. Rachen^{23,86}, B. Racine¹, R. Rebolo^{69,16,20}, M. Reinecke⁸⁶, M. Remazeilles^{73,63,1}, C. Renault⁸¹, A. Renzi^{39,55}, I. Ristorcelli^{109,10}, G. Rocha^{72,12}, C. Rosset¹, M. Rossetti^{37,52}, G. Roudier^{1,78,72}, J. A. Rubiño-Martín^{69,20}, B. Rusholme⁶⁰, M. Sandri⁵¹, D. Santos⁸¹, M. Savelainen^{29,46}, G. Savini⁹³, D. Scott²⁴, M. D. Seiffert^{72,12}, E. P. S. Shellard¹³, M. Shiraiishi^{33,71}, K. Smith⁹⁴, L. D. Spencer⁹⁷, V. Stolyarov^{6,102,76}, R. Stompor¹, R. Sudiwala⁹⁷, R. Sunyaev^{86,100}, P. Sutter⁶⁴, D. Sutton^{67,75}, A.-S. Suur-Uski^{29,46}, J.-F. Sygnet⁶⁴, J. A. Tauber⁴³, L. Terenzi^{96,51}, L. Toffolatti^{21,70,51}, M. Tomasi^{37,52}, M. Tristram⁷⁷, A. Troja³⁷, M. Tucci¹⁹, J. Tuovinen¹¹, L. Valenziano⁵¹, J. Valiviita^{29,46}, B. Van Tent⁸², P. Vielva⁷⁰, F. Villa⁵¹, L. A. Wade⁷², B. D. Wandelt^{64,108,32}, I. K. Wehus^{72,68}, D. Yvon¹⁷, A. Zacchei⁵⁰, and A. Zonca³¹

(Affiliations can be found after the references)

Received xxxx, Accepted xxxxx

ABSTRACT

The *Planck* full mission cosmic microwave background (CMB) temperature and *E*-mode polarization maps are analysed to obtain constraints on primordial non-Gaussianity (NG). Using three classes of optimal bispectrum estimators — separable template-fitting (KSW), binned, and modal — we obtain consistent values for the primordial local, equilateral, and orthogonal bispectrum amplitudes, quoting as our final result from temperature alone $f_{\text{NL}}^{\text{local}} = 2.5 \pm 5.7$, $f_{\text{NL}}^{\text{equil}} = -16 \pm 70$, and $f_{\text{NL}}^{\text{ortho}} = -34 \pm 33$ (68 % CL, statistical). Combining temperature and polarization data we obtain $f_{\text{NL}}^{\text{local}} = 0.8 \pm 5.0$, $f_{\text{NL}}^{\text{equil}} = -4 \pm 43$, and $f_{\text{NL}}^{\text{ortho}} = -26 \pm 21$ (68 % CL, statistical). The results are based on comprehensive cross-validation of these estimators on Gaussian and non-Gaussian simulations, are stable across component separation techniques, pass an extensive suite of tests, and are consistent with estimators based on measuring the Minkowski functionals of the CMB. The effect of time-domain de-glitching systematics on the bispectrum is negligible. In spite of these test outcomes we conservatively label the results including polarization data as preliminary, owing to a known mismatch of the noise model in simulations and the data. Beyond estimates of individual shape amplitudes, we present model-independent, three-dimensional reconstructions of the *Planck* CMB bispectrum and derive constraints on early universe scenarios that generate primordial NG, including general single-field models of inflation, axion inflation, initial state modifications, models producing parity-violating tensor bispectra, and directionally dependent vector models. We present a wide survey of scale-dependent feature and resonance models, accounting for the “look elsewhere” effect in estimating the statistical significance of features. We also look for isocurvature NG, and find no signal, but we obtain constraints that improve significantly with the inclusion of polarization. The primordial trispectrum amplitude in the local model is constrained to be $g_{\text{NL}}^{\text{local}} = (-9.0 \pm 7.7) \times 10^4$ (68 % CL statistical), and we perform an analysis of trispectrum shapes beyond the local case. The global picture that emerges is one of consistency with the premises of the Λ CDM cosmology, namely that the structure we observe today was sourced by adiabatic, passive, Gaussian, and primordial seed perturbations.

Key words. cosmology: cosmic background radiation – cosmology: observations – cosmology: theory – cosmology: early Universe – cosmology: inflation – methods: data analysis

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 the constraints on primordial non-Gaussianity (NG) obtained using the cosmic microwave background (CMB) maps from the full *Planck* mission, including a first analysis of some of the *Planck* polarization data.

Primordial NG is one of the most powerful tests of inflation, and more generally of high-energy early Universe physics (for some reviews, see Bartolo et al. 2004a, Liguori et al. 2010, Chen 2010b, Komatsu 2010, Yadav & Wandelt 2010). In fact, the simplest models of inflation (characterized by a single scalar field slowly rolling along a smooth potential) predict the generation of primordial fluctuations that are almost Gaussian distributed, with a tiny deviation from Gaussianity of the order of the slow-roll parameters (Acquaviva et al. 2003; Maldacena 2003). The 2013 *Planck* results on primordial NG are consistent with such a prediction, being compatible with Gaussian primordial fluctuations: the standard scenario of single-field, slow-roll inflation has survived its most stringent test to date. For example, in Planck Collaboration XXIV (2014) we obtained $f_{\text{NL}}^{\text{local}} = 2.7 \pm 5.8$, $f_{\text{NL}}^{\text{equil}} = -42 \pm 75$, and $f_{\text{NL}}^{\text{ortho}} = -25 \pm 39$ for the amplitudes of three of the most well-studied shapes of primordial NG. On the other hand, it is well known that any deviations from the standard picture of inflation have the potential to produce distinctive NG signatures at a detectable level in the CMB anisotropies.² Therefore, as already shown in Planck Collaboration XXIV (2014) (see also Planck Collaboration XXII (2014)) improved NG constraints allow severe limits to be placed on various classes of inflationary models that extend the simplest paradigm, in a way that is strongly complementary to the power-spectrum constraints (i.e., scalar spectral index of curvature perturbations and tensor-to-scalar amplitude ratio).

One of the main goals of this paper is to improve NG constraints using mainly the angular bispectrum of CMB anisotropies, i.e., the harmonic transform of the 3-point angular correlation function. We also investigate higher-order NG correlators like the trispectrum. We follow the same notation as Planck Collaboration XXIV (2014). The CMB angular bispectrum is related to the primordial bispectrum

$$\langle \Phi(\mathbf{k}_1)\Phi(\mathbf{k}_2)\Phi(\mathbf{k}_3) \rangle = (2\pi)^3 \delta^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B_{\Phi}(k_1, k_2, k_3), \quad (1)$$

where the field Φ , related to the comoving curvature perturbation ζ on super-horizon scales by $\Phi \equiv (3/5)\zeta$, is such that in the matter era, and on super-horizon scales, it reduces to Bardeen’s gauge-invariant gravitational potential (Bardeen 1980). The bispectrum $B_{\Phi}(k_1, k_2, k_3)$ measures the correlation among three perturbation modes. If translational and rotational invariance are assumed, it depends only on the magnitude of the three wavevectors. In general the bispectrum can be written as

$$B_{\Phi}(k_1, k_2, k_3) = f_{\text{NL}} F(k_1, k_2, k_3), \quad (2)$$

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).

² We refer the reader to Planck Collaboration XXIV (2014) and references therein for a detailed summary of the models and underlying physical mechanisms generating various types of primordial NG.

where we have introduced the dimensionless “nonlinearity parameter” f_{NL} (Gangui et al. 1994; Wang & Kamionkowski 2000; Komatsu & Spergel 2001; Babich et al. 2004), measuring the NG amplitude. The bispectrum is obtained by sampling triangles in Fourier space. The dependence of the function $F(k_1, k_2, k_3)$ on the type of triangle (i.e., the configuration) formed by the three wavevectors describes the *shape* (and the scale dependence) of the bispectrum (Babich et al. 2004), which encodes much physical information. Different NG shapes are linked to distinctive physical mechanisms that can generate such NG fingerprints in the early Universe.

In this paper the limits on primordial NG are mainly improved through the use of the full mission data, as well as by exploiting the polarization information.

Planck results on primordial NG also provide a reconstruction of the full CMB bispectrum through different techniques (see Sect. 6.2). This complements (and adds to) the extraction of single amplitudes f_{NL} for specific bispectrum shapes. Such a reconstruction can point to interesting features in the bispectrum signal that go beyond the usual standard scale-invariant shapes (such as the well known “local” and “equilateral” configurations).

As we have seen, the *Planck* 2013 NG paper (Planck Collaboration XXIV 2014) significantly improved constraints on the standard primordial NG models with scale-invariant local, equilateral or orthogonal shapes. The *Planck* NG paper also included constraints from the modal estimator on a variety of other primordial models, including DBI inflation, non-Bunch-Davies models (excited initial states), directionally-dependent vector inflation models, warm inflation, and scale-dependent feature and resonance models. All scale-invariant bispectra were strongly constrained, with the possible exception of highly flattened non-Bunch-Davies models. On the other hand, the preliminary investigation of primordial oscillatory models seemed to be more promising, in that two specific feature models appeared to produce fits of some significance. One aim of the present work is to expand the detail and scope of investigations of feature and resonant models and to examine the significance of these results with a more careful analysis of the “look elsewhere” effect, through exploring multi-parameter results using large ensembles of Gaussian simulations. Also we will thoroughly analyse or re-analyse other primordial NG signals that are theoretically well-motivated and those which have appeared in the literature since the first data release. These include primordial NG arising in the context of inflation models where vector fields play a non-negligible role or primordial NG generated in the tensor (gravitational waves) perturbations. Each of these primordial NG signals carry distinctive signatures that may have been imprinted at the inflationary epoch, thus opening up a new window into the detailed physics of inflation.

The paper is organized as follows. In Sect. 2 we briefly discuss the primordial NG models that we test in this paper. Section 3 summarizes the optimal statistical estimators used to constrain the CMB bispectrum and trispectrum from *Planck* temperature and polarization data. In Sect. 4 we discuss the non-primordial contributions to the CMB bispectrum and trispectrum, including foreground residuals after component separation and focusing on the f_{NL} bias induced by the ISW-lensing bispectrum. We also analyse the impact on primordial NG estimation from the residuals of the deglitching processing. Section 5 describes an extensive suite of tests performed on realistic simulations to validate the different estimator pipelines, and compare their performance. Using simulations, we also quantify the impact on f_{NL} of using a variety of component-separation tech-

niques. In Sect. 6 we derive constraints on f_{NL} for the local, equilateral, and orthogonal bispectra and present a reconstruction of the CMB bispectrum. We also present a reconstruction of the primordial curvature fluctuations. In Sect. 7 we validate these results by performing a series of null tests on the data to assess the robustness of our results. Section 8 investigates scale-dependent NG models and other selected bispectrum shapes. Section 9 presents the *Planck* limits on the CMB trispectrum. In Sect. 10 we provide constraints on CMB local bispectrum and trispectrum from Minkowski functionals. In Sect. 11 we discuss the main implications of *Planck*'s constraints on primordial NG for early Universe models. We conclude in Sect. 12. Appendix A contains the derivation of the statistical estimator for the amplitudes characterizing a “direction-dependent” primordial non-Gaussianity. Appendix B contains some details about Minkowski functionals.

2. Models

In this Section we briefly highlight the *classes* of inflationary models investigated in this paper, and describe the distinctive NG they generate. Within each class a common underlying physical process gives rise to the corresponding NG shape, illustrated by concrete realizations of inflationary models. For each class we therefore provide the explicit form of the bispectrum shapes chosen for the data analysis, emphasizing extensions with variants and distinctly new shapes beyond those already described in [Planck Collaboration XXIV \(2014\)](#).

2.1. General single-field models of inflation

This class of models includes inflationary models with a non-standard kinetic term (or more general higher-derivative interactions), in which the inflaton fluctuations propagate with an effective sound speed c_s which can be smaller than the speed of light. For example, models with a non-standard kinetic term are described by an inflaton Lagrangian $\mathcal{L} = P(X, \phi)$, where $X = g^{\mu\nu} \partial_\mu \phi \partial_\nu \phi$, with at most one derivative on ϕ , and the sound speed is given by $c_s^2 = (\partial P / \partial X) / (\partial P / \partial X + 2X \partial^2 P / \partial X^2)$.

The NG parameter space of this class of models is generically well described by two NG shapes — “equilateral” and “orthogonal” ([Senatore et al. 2010](#)) — since usually there are two dominant interaction terms of the inflaton field giving rise to the overall NG signal. One of these typically produces a bispectrum very close to the equilateral type with $f_{\text{NL}} \sim c_s^{-2}$ in the limit $c_s \ll 1$ ([Chen et al. 2007b](#); [Senatore et al. 2010](#)).

The equilateral-type NG is well approximated by the template ([Creminelli et al. 2006](#))

$$B_\Phi^{\text{equil}}(k_1, k_2, k_3) = 6A^2 f_{\text{NL}}^{\text{equil}} \times \left\{ -\frac{1}{k_1^{4-n_s} k_2^{4-n_s}} - \frac{1}{k_2^{4-n_s} k_3^{4-n_s}} - \frac{1}{k_3^{4-n_s} k_1^{4-n_s}} - \frac{2}{(k_1 k_2 k_3)^{2(4-n_s)/3}} + \left[\frac{1}{k_1^{(4-n_s)/3} k_2^{2(4-n_s)/3} k_3^{4-n_s}} + (5 \text{ permutations}) \right] \right\}. \quad (3)$$

Here $P_\Phi(k) = A/k^{4-n_s}$ represents Bardeen's gravitational potential power spectrum, A^2 being the normalization and n_s the scalar spectral index. DBI inflationary models based on string theory ([Silverstein & Tong 2004](#); [Alishahiha et al. 2004](#)) provide physically well-motivated examples of the $P(X, \phi)$ -model. They are characterized by an almost equilateral NG with $f_{\text{NL}}^{\text{equil}} = -(35/108)c_s^{-2}$ for $c_s \ll 1$, which typically is $f_{\text{NL}}^{\text{equil}} < -5$.

The “orthogonal” shape template is ([Senatore et al. 2010](#))

$$B_\Phi^{\text{ortho}}(k_1, k_2, k_3) = 6A^2 f_{\text{NL}}^{\text{ortho}} \times \left\{ -\frac{3}{k_1^{4-n_s} k_2^{4-n_s}} - \frac{3}{k_2^{4-n_s} k_3^{4-n_s}} - \frac{3}{k_3^{4-n_s} k_1^{4-n_s}} - \frac{8}{(k_1 k_2 k_3)^{2(4-n_s)/3}} + \left[\frac{3}{k_1^{(4-n_s)/3} k_2^{2(4-n_s)/3} k_3^{4-n_s}} + (5 \text{ perm.}) \right] \right\}. \quad (4)$$

Equilateral and orthogonal shapes emerge also in models characterized by more general higher-derivative interactions, such as ghost inflation ([Arkani-Hamed et al. 2004](#)), effective field theories of inflation ([Cheung et al. 2008](#); [Senatore et al. 2010](#); [Bartolo et al. 2010a](#)), or the so “Galileon-like” models of inflation (see, e.g., [Burrage et al. 2011](#)). The latter model is constructed starting from some specific underlying symmetry for the inflaton field, and is characterized by strongly constrained derivative interactions.

2.2. Multi-field models

This class of models is characterized by the presence of additional light scalar degrees of freedom besides the inflaton, whose fluctuations give rise, or contribute, to the final primordial curvature perturbation at the end of inflation. This includes the case of “multiple-field inflation”, where inflation is driven by more than one scalar field, as well as scenarios in which additional scalar fields remain subdominant during the inflationary expansion. From the point of view of primordial NG, the element in common to all these models is that a potentially detectable level of NG in the curvature perturbation is generated via a transfer of super-horizon non-Gaussian isocurvature perturbations in the second field (not necessarily the inflaton) to the adiabatic (curvature) density perturbations, accompanied by nonlinearities in the transfer mechanism. This process typically takes place on super-horizon scales, thus implying a local form of NG in real space. When going to Fourier space, this leads to a correlation between large and small scale modes. The bispectrum for this class of models is indeed largest on so-called “squeezed” triangles ($k_1 \ll k_2 \simeq k_3$). The local bispectrum is ([Falk et al. 1993](#); [Gangui et al. 1994](#); [Gangui & Martin 2000](#); [Verde et al. 2000](#); [Wang & Kamionkowski 2000](#); [Komatsu & Spergel 2001](#))

$$B_\Phi^{\text{local}}(k_1, k_2, k_3) = 2f_{\text{NL}}^{\text{local}} \left[P_\Phi(k_1)P_\Phi(k_2) + P_\Phi(k_1)P_\Phi(k_3) + P_\Phi(k_2)P_\Phi(k_3) \right] = 2A^2 f_{\text{NL}}^{\text{local}} \left[\frac{1}{k_1^{4-n_s} k_2^{4-n_s}} + \text{cycl.} \right]. \quad (5)$$

There is a broad literature on examples and specific realizations of this transfer mechanism from isocurvature to adiabatic perturbations ([Bartolo et al. 2002](#); [Bernardeau & Uzan 2002](#); [Vernizzi & Wands 2006](#); [Rigopoulos et al. 2006, 2007](#); [Lyth & Rodriguez 2005](#); [Tzavara & van Tent 2011](#); for a review on NG from multiple-field inflation models, see [Byrnes & Choi 2010](#)). An alternative, important possibility is the curvaton model ([Mollerach 1990](#); [Linde & Mukhanov 1997](#); [Enqvist & Sloth 2002](#); [Lyth & Wands 2002](#); [Moroi & Takahashi 2001](#)). In this type of scenario, a second light scalar field, subdominant during inflation, decays after inflation, generating primordial density perturbations with a potentially high level of NG (e.g., [Lyth & Wands 2002](#); [Lyth et al. 2003](#); [Bartolo et al.](#)

2004c). In the (simplest) adiabatic curvaton models, the local f_{NL} parameter was found to be (Bartolo et al. 2004c,b) $f_{\text{NL}}^{\text{local}} = (5/4r_{\text{D}}) - 5r_{\text{D}}/6 - 5/3$, when the curvaton field has a quadratic potential (Lyth & Wands 2002; Lyth et al. 2003; Lyth & Rodriguez 2005; Malik & Lyth 2006; Sasaki et al. 2006). In the previous formula, $r_{\text{D}} = [3\rho_{\text{curvaton}}/(3\rho_{\text{curvaton}} + 4\rho_{\text{radiation}})]_{\text{D}}$ is the ‘‘curvaton decay fraction’’ evaluated at the epoch of the curvaton decay in the sudden decay approximation. It is then easy to see that, for low values of r_{D} , a high level of NG can be generated.³

2.3. Isocurvature non-Gaussianity

Isocurvature NG, which was only sketched from the purely theoretical point of view in the 2013 paper, can now be analysed thanks to the polarization information.

In most of the models mentioned above, the main focus is on the level of primordial NG in the final curvature perturbation ζ . However, in inflationary scenarios where different scalar fields play a non-negligible role, residual isocurvature perturbation modes can remain after inflation. Isocurvature modes are usually investigated by considering their contribution to the power spectrum. However, if present, they would also contribute to the bispectrum, producing in general both a pure isocurvature bispectrum and mixed bispectra because of the cross-correlation between isocurvature and adiabatic perturbations (Komatsu 2002; Bartolo et al. 2002; Komatsu et al. 2005; Kawasaki et al. 2008; Langlois et al. 2008; Kawasaki et al. 2009; Hikage et al. 2009; Langlois & Lepidi 2011; Langlois & van Tent 2011; Kawakami et al. 2012; Langlois & van Tent 2012). While one might expect isocurvature NG to be negligible, since both (linear) isocurvature modes and (adiabatic) NG appear to be very small, and searches for isocurvature NG using WMAP data did not lead to any detections (Hikage et al. 2013a,b), this expectation can be tested at significantly higher precision by *Planck*. Moreover, there exist inflation models (Langlois & Lepidi 2011) where isocurvature modes, while remaining a small fraction in the power spectrum, would dominate the bispectrum.

At the time of recombination there are in principle four possible distinct isocurvature modes (in addition to the adiabatic mode): cold dark matter (CDM); baryon; neutrino density; and neutrino velocity isocurvature modes (Bucher et al. 2000). In this paper we will only consider isocurvature NG of the local type and always limit ourselves to considering the adiabatic mode together with just one type of isocurvature mode (considering each of the four types separately). Otherwise the number of free parameters becomes so large that no meaningful limits can be derived. Moreover, we assume the same spectral index for the primordial isocurvature power spectrum and the isocurvature-adiabatic cross-power spectrum as for the adiabatic power spectrum. Under those assumptions, as shown by Langlois & van Tent (2011), we have in principle six independent f_{NL} parameters: the usual purely adiabatic one; a purely isocurvature one; and four correlated ones.

³ NG perturbations can arise also at the end of inflation, e.g., from nonlinearities during the (pre)heating phase (e.g., Enqvist et al. 2005; Chambers & Rajantie 2008; Barnaby & Cline 2006; see also Bond et al. 2009) or from fluctuations in the inflaton decay rate or interactions, as found in modulated (pre)heating and modulated hybrid inflation (Kofman 2003; Dvali et al. 2004a,b; Bernardeau et al. 2004; Zaldarriaga 2004; Lyth 2005; Salem 2005; Lyth & Riotto 2006; Kolb et al. 2006; Cicoli et al. 2012).

The primordial shape templates are a generalization of Eq. (5), see Langlois & van Tent (2011, 2012):

$$B^{IJK}(k_1, k_2, k_3) = 2f_{\text{NL}}^{IJK} P_{\Phi}(k_2)P_{\Phi}(k_3) + 2f_{\text{NL}}^{JKI} P_{\Phi}(k_1)P_{\Phi}(k_3) + 2f_{\text{NL}}^{KIJ} P_{\Phi}(k_1)P_{\Phi}(k_2), \quad (6)$$

where I, J, K label the different modes (adiabatic and isocurvature). The invariance under the simultaneous exchange of two of these indices and the corresponding momenta means that $f_{\text{NL}}^{IJK} = f_{\text{NL}}^{IKJ}$, hence reducing the number of independent parameters from eight to six, in the case of two modes. The different bispectra vary most importantly through the fact that different types of radiation transfer functions $g_{\ell}^I(k)$ are used to project the primordial template onto the CMB: the reduced bispectra are of the form

$$b_{\ell_1 \ell_2 \ell_3}^{IJK} = 6 \int_0^{\infty} r^2 dr \alpha_{\ell_1}^I(r) \beta_{\ell_2}^J(r) \beta_{\ell_3}^K(r), \quad (7)$$

with

$$\alpha_{\ell}^I(r) \equiv \frac{2}{\pi} \int k^2 dk j_{\ell}(kr) g_{\ell}^I(k), \quad (8)$$

$$\beta_{\ell}^I(r) \equiv \frac{2}{\pi} \int k^2 dk j_{\ell}(kr) g_{\ell}^I(k) P_{\Phi}(k). \quad (9)$$

Here j_{ℓ} is the spherical bessel function and we use the notation $(\ell_1 \ell_2 \ell_3) \equiv [\ell_1 \ell_2 \ell_3 + 5 \text{ perm.}]/3!$. In addition to the isocurvature index, each transfer function carries a polarization index that we do not show here. It is important to note that, unlike the case of the purely adiabatic mode, the inclusion of polarization improves the constraints on the isocurvature NG significantly, as predicted by Langlois & van Tent (2011, 2012).

2.4. Resonance and axion monodromy models

Oscillatory models for NG are physically well-motivated. Large-field inflation faces an inherent UV completion problem because the inflaton field is required to move over large distances in field space relative to the Planck mass m_{Pl} . An effective shift symmetry can enforce potential flatness and this can be naturally implemented in a string theory context with axions and a periodically modulated potential, so-called ‘‘axion monodromy’’ models. This periodicity can generate resonances in the inflationary fluctuations with logarithmically-spaced oscillations, creating imprints in the power spectrum, the bispectrum and trispectrum (Chen et al. 2008; Flauger et al. 2010; Hannestad et al. 2010; Flauger & Pajer 2011). On the other hand, sharp features or corners in an inflationary potential can temporarily drive the inflaton away from slow-roll; these large changes in the field and derivatives can create evenly-spaced oscillations, to be discussed in the next subsection. However, in multifield models residual oscillations after corner-turning can also lead to log-spaced oscillations, just as in the resonance models (Chen 2011; Achúcarro et al. 2011; Battefeld et al. 2013; Chen et al. 2015). A preliminary search for bispectrum resonance signals was performed in the first *Planck* analysis (Planck Collaboration XXIV 2014) and our purpose here is to substantially increase the frequency range and number of models investigated.

Simple resonance model: Periodic features in the inflationary potential can induce oscillations with frequency ω that can resonate through any interactions with the inflationary fluctuations, contributing to the bispectrum. Provided that $\omega > H$, this mode starts inside the horizon but its frequency decreases as it is

stretched by inflation, until frozen when $\omega \simeq H$. Thus periodic features introduce a driving force which can scan across a wide range of frequencies. The simplest basic behaviour of such resonant models yields logarithmic stretching and can be described by the non-scale-invariant shape (see [Chen et al. 2008](#); [Chen 2010b](#))

$$B_{\Phi}^{\text{res}}(k_1, k_2, k_3) = \frac{6A^2 f_{\text{NL}}^{\text{res}}}{(k_1 k_2 k_3)^2} \sin [C \ln(k_1 + k_2 + k_3) + \phi], \quad (10)$$

where the constant $C = 1/\ln(3k_*)$, k_* is a wavenumber associated with the periodicity, and ϕ is a phase. These oscillations constructively and destructively interfere with the oscillations created by the CMB transfer functions, introducing additional nodal points in the CMB bispectrum.

Generalized resonance models: In a more general context, it is possible to have more complicated resonant shapes and envelopes. Resonant single-field models with varying sound speed c_s generate three leading-order bispectrum terms ([Chen 2010a](#)):

$$\begin{aligned} B^{\text{res-cs}}(k_1, k_2, k_3) = & \frac{6A^2}{(k_1 k_2 k_3)^2} \left\{ f_{\text{NL}}^{\text{res1}} \sin [C \ln(k_1 + k_2 + k_3) + \phi] \right. \\ & + 3 f_{\text{NL}}^{\text{res2}} \frac{k_1^2 + k_2^2 + k_3^2}{(k_1 + k_2 + k_3)^2} \cos [C \ln(k_1 + k_2 + k_3) + \phi] \\ & \left. + 27 f_{\text{NL}}^{\text{res3}} \frac{k_1 k_2 k_3}{(k_1 + k_2 + k_3)^3} \sin [C \ln(k_1 + k_2 + k_3) + \phi] \right\}. \quad (11) \end{aligned}$$

The first term on the right-hand side of Eq. (11) is the basic resonant shape given in Eq. (10), while the second and third terms have the same oscillatory behaviour, but modulated by a (mildly) flattened shape, and an equilateral shape respectively. The third term is in fact the second generic shape arising in effective field theory and correlates well with the equilateral shape in Eq. (12). The second term in Eq. (11) weakly favours flattened triangles, but there are regimes for resonant models that can generate much stronger flat shapes. If the resonance begins very deep inside the horizon, then the second (negative energy) mode can also make a significant contribution that is associated with enfolded or flat bispectra; this is similar to having an excited initial state or non-Bunch-Davies (NBD) vacuum.

With these two physical motivations in mind we also investigate classes of models with resonant oscillations modulated by both the equilateral and flattened shapes, defined by

$$S^{\text{eq}}(k_1, k_2, k_3) = \frac{\tilde{k}_1 \tilde{k}_2 \tilde{k}_3}{k_1 k_2 k_3}, \quad S^{\text{flat}} = 1 - S^{\text{eq}}, \quad (12)$$

where $\tilde{k}_1 \equiv k_2 + k_3 - k_1$ (here, for simplicity we ignore the spectral index dependence of the equilateral shape in Eq. 3). The corresponding equilateral and flattened resonant bispectra ansätze are then

$$\begin{aligned} B^{\text{res-eq}}(k_1, k_2, k_3) & \equiv S^{\text{eq}}(k_1, k_2, k_3) \times B^{\text{res}}(k_1, k_2, k_3) \\ & = \frac{6A^2 f_{\text{NL}}^{\text{res-eq}}}{(k_1 k_2 k_3)^2} \frac{\tilde{k}_1 \tilde{k}_2 \tilde{k}_3}{k_1 k_2 k_3} \sin [C \ln(k_1 + k_2 + k_3) + \phi], \quad (13) \end{aligned}$$

$$B^{\text{res-flat}}(k_1, k_2, k_3) \equiv S^{\text{flat}}(k_1, k_2, k_3) \times B^{\text{res}}(k_1, k_2, k_3). \quad (14)$$

We note that typically non-Bunch-Davies bispectra can be much more sharply peaked in the flattened or squeezed limits than Eq. (14), but our purpose here is to determine if this type of resonant model is favoured by the *Planck* data; that is, whether Eq. (14) warrants further investigation with other flattened profiles.

2.5. Scale-dependent oscillatory feature models

Temporary violations of slow-roll inflation can occur if there are sharp features in the inflationary potential ([Chen et al. 2007a](#)), as well as changes in the sound speed c_s or sharp turns in field space in multifield inflation. The inflaton field makes temporary departures from the attractor solution, which typically have a strong scale-dependent running modulated by a sinusoidal oscillation; there are model-dependent counterparts in the power spectrum, bispectrum, and trispectrum. For example, sharper or narrower features induce a relatively larger signal in the bispectrum (see e.g., [Chen 2010b](#)). An example is the analytic envelope solutions predicted for both the power spectrum and bispectrum for the single field models with a specific inflaton feature shape ([Adshead et al. 2012](#)); a search for these was presented previously in the *Planck* Inflation paper ([Planck Collaboration XXII 2014](#)) and likewise no significant signal was found using the corresponding bispectrum envelopes at the available modal resolution ([Planck Collaboration XXIV 2014](#)). In this new analysis, we will emphasize the search for generic oscillatory behaviour in the data over a larger range in modal resolution, although we will also look for the shapes predicted for simple features in single field models.

Constant feature model: In the previous investigation of *Planck* data using a coarse parameter grid ([Planck Collaboration XXIV 2014](#)), we searched for the simplest ansatz for an oscillatory bispectrum signal ([Chen et al. 2007a](#)):

$$B^{\text{feat}}(k_1, k_2, k_3) = \frac{6A^2 f_{\text{NL}}^{\text{feat}}}{(k_1 k_2 k_3)^2} \sin [\omega(k_1 + k_2 + k_3) + \phi], \quad (15)$$

where ϕ is a phase factor and ω is a frequency associated with the specific shape of the feature in the potential that disrupts the slow-roll evolution. In the earlier analysis, we also considered a damping envelope, which slightly increased the apparent significance of the best-fit feature models, though at the cost of an additional parameter (see single-field solutions below).

Generalized feature models: Here, we again search for oscillatory signals in a model-independent manner. We will modulate the bispectrum cross-sections with the physically motivated equilateral and flattened shapes, reflecting the physical contexts in which they could have been generated, as for the resonant models described above in Eq. (11). If there are potential features in a model with a varying sound speed, then we can expect there to be oscillatory contributions to the bispectrum signal with a dominant equilateral shape. Motivated by the equilateral resonance model in Eq. (11), we will search for the following *equilateral feature ansatz*:

$$B^{\text{feat-eq}}(k_1, k_2, k_3) \equiv S^{\text{eq}}(k_1, k_2, k_3) \times B^{\text{feat}}(k_1, k_2, k_3) \quad (16)$$

$$= \frac{6A^2 f_{\text{NL}}^{\text{feat-eq}}}{(k_1 k_2 k_3)^2} \frac{\tilde{k}_1 \tilde{k}_2 \tilde{k}_3}{k_1 k_2 k_3} \sin [\omega(k_1 + k_2 + k_3) + \phi]. \quad (17)$$

For extremely sharp features, it is possible to excite the inflationary fluctuations as if there were a non-Bunch-Davies vacuum: the oscillatory signal becomes modulated with a flattened shape ([Chen et al. 2007a](#)). Again, motivated by the enfolded resonance model in Eq. (14), we take the following simple flattened ansatz:

$$B^{\text{feat-flat}}(k_1, k_2, k_3) \equiv S^{\text{flat}}(k_1, k_2, k_3) \times B^{\text{feat}}(k_1, k_2, k_3). \quad (18)$$

Although the exact profile of the flattened shape can be much more highly peaked on the faces in these NBD models, this ansatz should be adequate for testing whether these models are

favoured. We note that while the power spectrum is insensitive to the underlying scenario creating the features, the bispectrum shape will reveal whether features arise from varying sound speed or highly excited features in the potential.

Single field feature solutions: Here we use the full analytic bispectrum solution given by [Adshead et al. \(2012\)](#), but the dominant leading-order behaviour takes the form

$$B^{K^2 \cos}(k_1, k_2, k_3) = \frac{6A^2 f_{\text{NL}}^{K^2 \cos}}{(k_1 k_2 k_3)^2} K^2 D(\alpha \omega K) \cos(\omega K), \quad (19)$$

where $K = k_1 + k_2 + k_3$ and $D(\alpha \omega K) = \alpha \omega / (K \sinh(\alpha \omega K))$ is an envelope function, with parameter α setting an overall cut-off for the bispectrum at large wavenumbers or multipoles. This envelope and the overall K^2 scaling distinguishes this realistic case from the simple separable constant feature ansatz of Eq. (15). We shall allow the envelope parameter α to vary from $\alpha = 0$, with no envelope (the infinitely thin limit for a feature in the potential) through to large α , with a narrow domain for the bispectrum. Alternative analytic solutions where the bispectrum is created by a variation in the sound speed c_s are dominated by the $K \sin(\omega K)$ term, as in

$$B^{K \sin}(k_1, k_2, k_3) = \frac{6A^2 f_{\text{NL}}^{K \sin}}{(k_1 k_2 k_3)^2} K D(\alpha \omega K) \sin(\omega K). \quad (20)$$

For the simplest models there is a predicted relationship between the power spectrum and bispectrum amplitude (e.g., see also [Achucarro et al. 2013](#) for a two-field model). We note that typically the power spectrum has larger signal-to-noise at low frequency (i.e., below $\omega \simeq 1000$) while the bispectrum dominates at higher frequency.

2.6. Non-Gaussianity from excited initial states

It is well known that if the initial vacuum state for inflation is excited and deviates from the standard Bunch-Davies vacuum, then measurable non-Gaussianities can be produced ([Chen et al. 2007b](#); [Holman & Tolley 2008](#); [Meerburg et al. 2009](#); [Ashoorioon & Shiu 2011](#)). These models generically lead to non-Gaussianity that peaks in the flattened limit, where $k_1 + k_2 \approx k_3$, and also often has oscillatory behaviour. Here we constrain the same selection of templates found in the 2013 *Planck* analysis, namely the flat model in Eq. (12), Non-Bunch-Davies (NBD) ([Chen et al. 2007b](#)), NBD1 and NBD2 models ([Agullo & Parker 2011](#)) (now called ‘‘NBD1 cos’’ and ‘‘NBD2 cos’’) and NBD3 ([Chen 2010b](#)). We also introduce three new templates, NBD sin which is motivated by ([Chen 2010a](#)) and takes the form

$$B^{\text{NBD-sin}}(k_1, k_2, k_3) = \frac{2A^2 f_{\text{NL}}^{\text{NBD-sin}}}{(k_1 k_2 k_3)^2} \left(e^{-\omega \tilde{k}_1} + e^{-\omega \tilde{k}_2} + e^{-\omega \tilde{k}_3} \right) \times \sin(\omega K + \phi), \quad (21)$$

where again $K = k_1 + k_2 + k_3$ and $\tilde{k}_i = K - 2k_i$. The other two templates are extensions of the NBD1 cos and NBD2 cos models found in [Agullo & Parker \(2011\)](#) and take the form

$$B^{\text{NBDi-sin}}(k_1, k_2, k_3) = \frac{2A^2 f_{\text{NL}}^{\text{NBDi-sin}}}{(k_1 k_2 k_3)^2} [f_i(k_1, k_2, k_3) \times \sin(\omega \tilde{k}_1) / \tilde{k}_1 + 2 \text{ perm.}], \quad (22)$$

where $f_1(k_1, k_2, k_3) = k_1^2(k_2^2 + k_3^2)/2$, which is dominated by squeezed configurations, and $f_2(k_1, k_2, k_3) = k_2^2 k_3^2$, which has a flattened shape.

2.7. Directional-dependence motivated by gauge fields

Some models where primordial vector fields are present during inflation predict interesting NG signatures. This is the case of a coupling of the inflaton field φ to the kinetic term of a gauge field A^μ , \mathcal{L} contains $-I^2(\varphi)F^2$, where $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ and the coupling $I^2(\varphi)F^2$ is chosen so that scale invariant vector perturbations are produced on superhorizon scales ([Barnaby et al. 2012b](#); [Bartolo et al. 2013a](#)). The bispectrum turns out to be the sum of two contributions: one of the local shape; and another that is also enhanced in the squeezed limit ($k_1 \ll k_2 \simeq k_3$), but featuring a non-trivial dependence on the angle between the small and the large wave vectors through the parameter $\mu_{12} = \hat{\mathbf{k}}_1 \cdot \hat{\mathbf{k}}_2$ (where $\hat{\mathbf{k}} = \mathbf{k}/k$) as μ_{12}^2 . Also, primordial magnetic fields sourcing curvature perturbations can cause a dependence on both μ and μ^2 ([Shiraishi et al. 2012](#)).

We can parametrize these shapes as variations on the local shape ([Shiraishi et al. 2013a](#)), as

$$B_\Phi(k_1, k_2, k_3) = \sum_L c_L [P_L(\mu_{12}) P_\Phi(k_1) P_\Phi(k_2) + 2 \text{ perm.}], \quad (23)$$

where $P_L(\mu)$ is the Legendre polynomial with $P_0 = 1$, $P_1 = \mu$, and $P_2 = \frac{1}{2}(3\mu^2 - 1)$. For example, for $L = 1$ we have the shape

$$B_\Phi^{L=1}(k_1, k_2, k_3) = \frac{2A^2 f_{\text{NL}}^{L=1}}{(k_1 k_2 k_3)^2} \left[\frac{k_3^2}{k_1^2 k_2^2} (k_1^2 + k_2^2 - k_3^2) + 2 \text{ perm.} \right]. \quad (24)$$

The local template corresponds to $c_i = 2f_{\text{NL}} \delta_{i0}$. Here and in the following the nonlinearity parameters f_{NL}^L are related to the c_L coefficients by $c_0 = 2f_{\text{NL}}^{L=0}$, $c_1 = -4f_{\text{NL}}^{L=1}$, and $c_2 = -16f_{\text{NL}}^{L=2}$. The $L = 1, 2$ shapes exhibit sharp variations in the flattened limit, for example for $k_1 + k_2 \approx k_3$, while in the squeezed limit, $L = 1$ is suppressed whereas $L = 2$ grows like the local bispectrum shape (i.e., the $L = 0$ case). The $I^2(\varphi)F^2$ models predict $c_2 = c_0/2$, while primordial curvature perturbations sourced by large-scale magnetic fields generate non-vanishing c_0 , c_1 , and c_2 . Quite interestingly, in the proposed ‘‘solid inflation’’ scenario ([Endlich et al. 2013](#); see also [Bartolo et al. 2013b](#); [Endlich et al. 2014](#); [Sitwell & Sigurdson 2014](#); [Bartolo et al. 2014](#)) bispectra similar to Eq. (23) can be generated, in this case with $c_2 \gg c_0$ ([Endlich et al. 2013, 2014](#)). Therefore, measurements of the c_i coefficients can be an efficient probe of some detailed aspects of the inflationary mechanism, such as the existence of primordial vector fields during inflation (or a non-trivial symmetry structure of the inflaton fields, as in solid inflation).

2.8. Non-Gaussianity from gauge-field production during axion inflation

The same shift symmetry that leads to axion (monodromy) models of inflation (Sect. 2.4) naturally allows (from an effective field theory point of view) for a coupling between a pseudoscalar axion inflaton field and a gauge field of the type $\mathcal{L} \supset -(\alpha/4f)\phi F^{\mu\nu} \tilde{F}_{\mu\nu}$, where the parameter α is dimensionless and f is the axion decay constant ($\tilde{F}^{\mu\nu} = \epsilon^{\mu\nu\gamma\beta} F_{\gamma\beta}/2$). This scenario has a rich and interesting phenomenology both for scalar and tensor primordial fluctuations (see, e.g., [Barnaby & Peloso 2011](#); [Sorbo 2011](#); [Barnaby et al. 2011, 2012c](#); [Linde et al. 2013](#); [Meerburg & Pajer 2013](#); [Ferreira & Sloth 2014](#)). Gauge field quanta are produced by the background motion of the inflaton field, and these in turn source curvature perturbations through an inverse decay process of the gauge field. A bispectrum of curvature fluctuations is generated as ([Barnaby et al.](#)

2011; Meerburg & Pajer 2013)⁴

$$B^{\text{inv.dec}} = 6A^2 f_{\text{NL}}^{\text{inv.dec}} \frac{\sum_i k_i^3 f_3(\xi_*, \frac{k_2}{k_1}, \frac{k_3}{k_1})}{\prod_i k_i^3 f_3(\xi_*, 1, 1)}, \quad (25)$$

where the exact expression for the function f_3 can be found in equation (3.29) of Barnaby et al. (2011) (see also Meerburg & Pajer 2013). Here ξ characterizes the coupling strength of the axion to the gauge field $\xi = a|\dot{\phi}|/(2fH)$. The inverse decay bispectrum peaks for equilateral configuration, since $\delta\varphi$ is mostly sourced by the inverse decay ($\delta A + \delta A \rightarrow \delta\varphi$), when two modes of the vector fields are of comparable magnitude (the correlation with the equilateral template is 94 % and with the orthogonal one is 4 %). We do however constrain the exact shape in Eq. (25), without resorting to the equilateral template. Another interesting observational signature that can shed light on the role played by pseudo-scalars in the early Universe is provided by tensor NG, to which we turn next.

2.9. Parity-violating tensor non-Gaussianity motivated by pseudo-scalars

While the majority of the studies on primordial and CMB NG focus on the scalar mode, tensor-mode NG has been attracting attention as a probe of high-energy theories of gravity (e.g., Maldacena & Pimentel 2011; McFadden & Skenderis 2011; Soda et al. 2011; Shiraishi et al. 2011; Gao et al. 2011) or primordial magnetic fields (Brown & Crittenden 2005; Shiraishi et al. 2012; Shiraishi 2012).⁵

Recently, the possibility of observable tensor bispectra has been vigorously discussed in a model where the inflaton couples to a pseudoscalar field (Barnaby et al. 2012a; Cook & Sorbo 2013; Ferreira & Sloth 2014). In this model, through the gravitational coupling to the U(1) gauge field, gravitational waves ($h_{ij} \equiv \delta g_{ij}^{TT}/a^2 = \sum_{s=\pm} h^{(s)} e_{ij}^{(s)}$) receive NG corrections, where only one of the two spin states is enhanced. The bispectrum, generally formed as

$$\left\langle \prod_{i=1}^3 h^{(s_i)}(\mathbf{k}_i) \right\rangle = (2\pi)^3 \delta^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B_h^{s_1 s_2 s_3}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3), \quad (26)$$

is accordingly polarized, with $B_h^{+++} \gg B_h^{++-}, B_h^{+-+}, B_h^{---}$. This NG enhancement is a sub-horizon effect and therefore B_h^{+++} is maximized at the equilateral limit ($k_1 \simeq k_2 \simeq k_3$) (Cook & Sorbo 2013).

A model-independent template of the equilateral-type polarized tensor bispectrum is given by (Shiraishi et al. 2013b, 2015)

$$B_h^{+++}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) = f_{\text{NL}}^{\text{tens}} F_{\zeta}^{\text{equil}}(k_1, k_2, k_3) \times \frac{16\sqrt{2}}{27} e_{ij}^{(+)*}(\hat{\mathbf{k}}_1) e_{jk}^{(+)*}(\hat{\mathbf{k}}_2) e_{ki}^{(+)*}(\hat{\mathbf{k}}_3), \quad (27)$$

with the polarization tensor $e_{ij}^{(s)}$ obeying $e_{ij}^{(s)}(\hat{\mathbf{k}}) e_{ij}^{(s')}(\hat{\mathbf{k}}) = 2\delta_{s,-s'}$ and $e_{ij}^{(s)*}(\hat{\mathbf{k}}) = e_{ij}^{(-s)}(\hat{\mathbf{k}}) = e_{ij}^{(s)}(-\hat{\mathbf{k}})$. We here have introduced a tensor nonlinearity parameter, by normalizing with the equilateral bispectrum template of curvature perturbations ($F_{\zeta}^{\text{equil}} \equiv$

$(5/3)^3 F_{\Phi}^{\text{equil}} = (5/3)^3 B_{\Phi}^{\text{equil}}/f_{\text{NL}}^{\text{equil}}$) in the equilateral limit, yielding

$$f_{\text{NL}}^{\text{tens}} \equiv \lim_{k_i \rightarrow k} \frac{B_h^{+++}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)}{F_{\zeta}^{\text{equil}}(k_1, k_2, k_3)}. \quad (28)$$

The template Eq. (27) can adequately reconstruct the tensor bispectra created in the pseudoscalar inflation models⁶ (Shiraishi et al. 2013b), and thus the amplitude $f_{\text{NL}}^{\text{tens}}$ is directly connected with the model parameters, e.g., the coupling strength of the pseudoscalar field to the gauge field ξ (for details see Sect. 11).

The CMB temperature and E -mode bispectra sourced by the parity-violating tensor NG have not only the usual parity-even ($\ell_1 + \ell_2 + \ell_3 = \text{even}$) signals but also parity-odd ($\ell_1 + \ell_2 + \ell_3 = \text{odd}$) contributions, which cannot be sourced by known scalar bispectra (Kamionkowski & Souradeep 2011; Shiraishi et al. 2011). Moreover, their shapes are mostly distinct from the scalar templates, due the different radiation transfer functions; hence they can be measured essentially independently of the scalar NG (Shiraishi et al. 2013b). The analysis of the WMAP temperature data distributed in $\ell_1 + \ell_2 + \ell_3 = \text{odd}$ configurations leads to an observational limit $f_{\text{NL}}^{\text{tens}} = (0.8 \pm 1.1) \times 10^4$ (Shiraishi et al. 2015). This paper updates the limit, by analysing both parity-even and parity-odd signals in the *Planck* temperature and E -mode polarization data.

3. Statistical estimation of the CMB bispectrum for polarized maps

We now provide a brief overview of the main statistical techniques that we use to estimate the nonlinearity parameter f_{NL} from temperature and polarization CMB data, followed by a description of the data set that will be used in our analysis.

The CMB temperature and polarization fields are characterized using the multipoles of a spherical harmonic decomposition of the CMB maps:

$$\begin{aligned} \frac{\Delta T}{T}(\hat{\mathbf{n}}) &= \sum_{\ell m} a_{\ell m}^T Y_{\ell m}(\hat{\mathbf{n}}), \\ E(\hat{\mathbf{n}}) &= \sum_{\ell m} a_{\ell m}^E Y_{\ell m}(\hat{\mathbf{n}}). \end{aligned} \quad (29)$$

At linear order, the relation between the primordial perturbation field and the CMB multipoles is (e.g., Komatsu & Spergel 2001)

$$a_{\ell m}^X = 4\pi(-i)^{\ell} \int \frac{d^3 k}{(2\pi)^3} \Phi(\mathbf{k}) Y_{\ell m}(\hat{\mathbf{k}}) \Delta_{\ell}^X(k), \quad (30)$$

where $X = \{T, E\}$ denotes either temperature or E -mode polarization, Φ is the primordial gravitational potential, and Δ_{ℓ}^X represents the linear CMB radiation transfer function.

The CMB angular bispectrum is the three-point correlator of the $a_{\ell m}$ s

$$B_{\ell_1 \ell_2 \ell_3}^{m_1 m_2 m_3, X_1 X_2 X_3} \equiv \langle a_{\ell_1 m_1}^{X_1} a_{\ell_2 m_2}^{X_2} a_{\ell_3 m_3}^{X_3} \rangle, \quad (31)$$

where $X_i = \{T, E\}$. If the CMB sky is rotationally invariant, and the bispectra we are considering have even parity (which is true for combinations of T and E), then the angular bispectrum can be factorized as

$$\langle a_{\ell_1 m_1}^{X_1} a_{\ell_2 m_2}^{X_2} a_{\ell_3 m_3}^{X_3} \rangle = \mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} b_{\ell_1 \ell_2 \ell_3}^{X_1 X_2 X_3}, \quad (32)$$

⁶ The form of the tensor bispectrum is the same whether the inflaton field is identified with the pseudoscalar field or not.

⁴ For simplicity we assume a scale-invariant bispectrum.

⁵ See Planck Collaboration XIX (2016) for the *Planck* constraints on magnetically-induced NG.

where $b_{\ell_1 \ell_2 \ell_3}^{X_1 X_2 X_3}$ is the so-called *reduced bispectrum*, and $\mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3}$ is the Gaunt integral, defined as the integral over the solid angle of the product of three spherical harmonics,

$$\mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} \equiv \int Y_{\ell_1 m_1}(\hat{\mathbf{n}}) Y_{\ell_2 m_2}(\hat{\mathbf{n}}) Y_{\ell_3 m_3}(\hat{\mathbf{n}}) d^2 \hat{\mathbf{n}}. \quad (33)$$

The Gaunt integral (often written in terms of Wigner 3j-symbols) enforces rotational symmetry, and restricts attention to a tetrahedral domain of multipole triplets $\{\ell_1, \ell_2, \ell_3\}$, satisfying both a triangle condition and a limit given by some maximum resolution ℓ_{\max} (the latter being defined by the finite angular resolution of the experiment under study).

Our goal is to extract the nonlinearity parameter f_{NL} from the data, for different primordial shapes. To achieve this, we essentially fit a theoretical CMB bispectrum ansatz $b_{\ell_1 \ell_2 \ell_3}$ to the observed 3-point function. Theoretical predictions for CMB angular bispectra arising from early Universe primordial models can be obtained by applying Eq. (30) to the primordial bispectra of Sect. 2, (see e.g., Komatsu & Spergel 2001). Optimized cubic bispectrum estimators were introduced by Heavens (1998), and it has been shown that for small NG the general optimal polarized f_{NL} estimator can be written as (Creminelli et al. 2006)

$$\begin{aligned} \hat{f}_{\text{NL}} = \frac{1}{N} \sum_{X_i X'_i} \sum_{\ell_i m_i} \sum_{\ell'_i m'_i} \mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} b_{\ell_1 \ell_2 \ell_3}^{X_1 X_2 X_3, \text{th}} \\ \left\{ \left[(\mathbf{C}_{\ell_1 m_1, \ell'_1 m'_1}^{-1})^{X_1 X'_1} a_{\ell'_1 m'_1}^{X_1} (\mathbf{C}_{\ell_2 m_2, \ell'_2 m'_2}^{-1})^{X_2 X'_2} a_{\ell'_2 m'_2}^{X_2} \right. \right. \\ \left. \left. (\mathbf{C}_{\ell_3 m_3, \ell'_3 m'_3}^{-1})^{X_3 X'_3} a_{\ell'_3 m'_3}^{X_3} \right] - \right. \\ \left. \left[(\mathbf{C}_{\ell_1 m_1, \ell_2 m_2}^{-1})^{X_1 X_2} (\mathbf{C}_{\ell_3 m_3, \ell'_3 m'_3}^{-1})^{X_3 X'_3} a_{\ell'_3 m'_3}^{X_3} + \text{cyclic} \right] \right\}, \quad (34) \end{aligned}$$

where N is a suitable normalization chosen to produce unit response to $b_{\ell_1 \ell_2 \ell_3}^{\text{th}}$. Note that we are implicitly defining a suitable normalization convention so that $b_{\ell_1 \ell_2 \ell_3} = f_{\text{NL}} b_{\ell_1 \ell_2 \ell_3}^{\text{th}}$, and $b_{\ell_1 \ell_2 \ell_3}^{\text{th}}$ is the value of the theoretical template when $f_{\text{NL}} = 1$. \mathbf{C}^{-1} is the inverse of the block matrix:

$$\mathbf{C} = \begin{pmatrix} \mathbf{C}^{TT} & \mathbf{C}^{TE} \\ \mathbf{C}^{ET} & \mathbf{C}^{EE} \end{pmatrix}, \quad (35)$$

and the blocks represent the full TT , TE , and EE covariance matrices, with \mathbf{C}^{ET} being the transpose of \mathbf{C}^{TE} . All quantities in the previous equation (i.e., CMB multipoles, bispectrum template and covariances matrices) are assumed to properly incorporate instrumental beam and noise.

As standard for these estimators, we note in square brackets (below) the presence of two contributions. One is cubic in the observed $a_{\ell m}$ s, and correlates the bispectrum of the data to the theoretical fitting template $b_{\ell_1 \ell_2 \ell_3}^{\text{th}}$. This is generally called the ‘‘cubic term’’ of the estimator. The other contribution is linear in the observed $a_{\ell m}$ s (‘‘linear term’’). This part corrects for mean-field contributions to the error bars, introduced by rotational invariance-breaking features, such as a mask or anisotropic/correlated instrumental noise (Creminelli et al. 2006; Yadav et al. 2008).

The inverse covariance filtering operation implied by Eq. (34) is a challenging numerical task, which has been successfully performed only recently (Smith et al. 2009; Elsner & Wandelt 2012). This step can be avoided by working in the ‘‘diagonal covariance approximation’’. In this approach, the estimator is built by neglecting off-diagonal entries of the covariance matrix in the cubic term in Eq. (34), and then finding

the linear term that minimizes the variance for this specific cubic statistic. Applying such a procedure yields (Yadav et al. 2007)

$$\begin{aligned} \hat{f}_{\text{NL}} = \frac{1}{N} \sum_{X_i X'_i} \sum_{\ell_i m_i} \mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} (\mathbf{C}^{-1})_{\ell_1}^{X_1 X'_1} (\mathbf{C}^{-1})_{\ell_2}^{X_2 X'_2} (\mathbf{C}^{-1})_{\ell_3}^{X_3 X'_3} b_{\ell_1 \ell_2 \ell_3}^{X_1 X_2 X_3, \text{th}} \\ \left[a_{\ell_1 m_1}^{X'_1} a_{\ell_2 m_2}^{X'_2} a_{\ell_3 m_3}^{X'_3} - \mathbf{C}_{\ell_1 m_1, \ell_2 m_2}^{X'_1 X'_2} a_{\ell_3 m_3}^{X'_3} - \mathbf{C}_{\ell_1 m_1, \ell_3 m_3}^{X'_1 X'_3} a_{\ell_2 m_2}^{X'_2} \right. \\ \left. - \mathbf{C}_{\ell_2 m_2, \ell_3 m_3}^{X'_2 X'_3} a_{\ell_1 m_1}^{X'_1} \right], \quad (36) \end{aligned}$$

where \mathbf{C}_ℓ^{-1} is the inverse of the 2×2 matrix

$$\mathbf{C}_\ell = \begin{pmatrix} \mathbf{C}_{\ell}^{TT} & \mathbf{C}_{\ell}^{TE} \\ \mathbf{C}_{\ell}^{ET} & \mathbf{C}_{\ell}^{EE} \end{pmatrix}. \quad (37)$$

This expression can also be written as

$$\hat{f}_{\text{NL}} = \frac{\langle b^{\text{th}}, b^{\text{obs}} \rangle}{\langle b^{\text{th}}, b^{\text{th}} \rangle}, \quad (38)$$

where the observed (reduced) bispectrum includes the linear correction term and the inner product is defined as

$$\langle b^A, b^B \rangle = \quad (39)$$

$$\sum_{X_i X'_i} \sum_{\ell_i} b_{\ell_1 \ell_2 \ell_3}^{X_1 X_2 X_3, A} h_{\ell_1 \ell_2 \ell_3}^2 (\mathbf{C}^{-1})_{\ell_1}^{X_1 X'_1} (\mathbf{C}^{-1})_{\ell_2}^{X_2 X'_2} (\mathbf{C}^{-1})_{\ell_3}^{X_3 X'_3} b_{\ell_1 \ell_2 \ell_3}^{X'_1 X'_2 X'_3, B}$$

with

$$h_{\ell_1 \ell_2 \ell_3} = \sqrt{\frac{(2\ell_1 + 1)(2\ell_2 + 1)(2\ell_3 + 1)}{4\pi}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix}, \quad (40)$$

where the last term is the Wigner 3j symbol. The denominator in Eq. (38), $\langle b^{\text{th}}, b^{\text{th}} \rangle$ is the normalization constant N .

The price to pay for the simplification obtained in Eq. (36) is, in principle, loss of optimality. However, in practice we found in our previous temperature analysis (Planck Collaboration XXIV 2014) that error bars obtained with this simplified procedure are very close to optimal, provided the $a_{\ell m}$ s are pre-filtered with a simple diffusive inpainting technique (see Planck Collaboration XXIV 2014 for details). We find that this still holds true when we include polarization and pre-inpaint the T, Q, U input maps. Given its practical advantages in terms of speed and simplicity, we adopt this method in the following analysis.

A well-known, major issue with both Eq. (34) and (36) is that their direct implementation would require evaluation of all the bispectrum configurations from the data. The computational cost of this would scale like ℓ_{\max}^5 and be totally prohibitive for high-resolution CMB experiments like *Planck*. The different bispectrum estimation techniques applied to our analysis are essentially defined by the approach adopted to circumvent this problem. The advantage of having multiple independent implementations of the optimal bispectrum estimator is twofold. First, by cross-validating and comparing outputs of different pipelines, it strongly improves the robustness of the results. Second, different methods are complementary, in the sense that they have specific capabilities which go beyond simple f_{NL} estimation. For example, the skew- C_ℓ method defined below facilitates the monitoring of NG foreground contamination, while the binned and modal estimators allow model-independent reconstruction of the data bispectrum, and so on. The skew- C_ℓ method enables the nature of any detected NG to be determined. Thus, the simultaneous application of all these techniques also allows us to increase the range and scope of our analysis.

In the following, we briefly outline the main features of the three optimal bispectrum estimation pipelines that are used for *Planck* measurements of f_{NL} . We will only provide a short summary here, focused on the extension to polarization data, referring the reader who is interested in more technical aspects to our previous analysis of temperature data (Planck Collaboration XXIV 2014).

3.1. KSW and skew- C_ℓ estimators

KSW and skew- C_ℓ estimators (Komatsu et al. 2005; Munshi & Heavens 2010) can be used for bispectrum templates that are written in *factorizable* (separable) form, i.e., as a linear combination of separate products of functions⁷ of k_1 , k_2 , and k_3 . This allows reduction of the three-dimensional integration over the bispectrum configurations into a product of three separate one-dimensional sums over ℓ_1 , ℓ_2 , ℓ_3 . This leads to a massive reduction in computational time ($O(N_{\text{pix}})$, where N_{pix} is the number of pixels in the map). The main difference between the KSW and skew- C_ℓ pipelines is that the former estimates the f_{NL} amplitude directly, whereas the latter initially estimates the so called “bispectrum-related power spectrum” (in short, “skew- C_ℓ ”) function. Roughly speaking, the skew- C_ℓ associates, with each angular wavenumber ℓ , the contribution to the amplitude f_{NL} (for each given shape) extracted from all triangles with one fixed side of size ℓ . After resummation over the contributions from each ℓ -bin, the final point-like f_{NL} estimate is obtained exactly as KSW. Equipping the KSW estimator with a skew- C_ℓ extension can be particularly useful in the presence of (expected) spurious NG contaminants in the data. The slope of the skew- C_ℓ statistic is in fact shape-dependent and can be used to separate multiple NG components in the map.

3.2. Modal estimators

Modal estimators (Fergusson et al. 2010a, 2012) are based on decomposing the bispectrum (both from theory and from data) into a sum of uncorrelated separable templates, forming a *complete basis* in bispectrum space, and measuring the amplitude of each. The evaluation of the amplitude for each template can be sped up by using a KSW approach (since the templates themselves are separable by construction). All amplitudes form a vector, also referred to as the “mode spectrum”. It is then possible to measure the correlation of the observed data mode spectrum with the theoretical mode spectra for different primordial shapes, in order to obtain estimates of the primordial f_{NL} . Note also that the observed mode spectrum from data is theory-independent, and contains *all* the information from the data. Correlating the observed mode spectrum to theoretical mode vectors then allows the extraction of all the f_{NL} amplitudes simultaneously. This makes modal estimators naturally suited for NG analyses, both when there are a large number of competing models to analyse, or when a model has free parameters through which we wish to scan (more than 500 shapes were analysed when applying this technique to *Planck* data). Another advantage is that by expanding into separable basis templates, the modal estimator does not

require separability of the starting theoretical shape in order to be applicable. Finally, after obtaining the data mode spectrum, it is possible to build a linear combination of the basis templates, using the measured amplitudes as coefficients, thus obtaining a model-independent full reconstruction of the bispectrum of the data. Of course the reconstructed bispectrum will be smoothed, as the estimator must use a finite number of basis templates.

For this analysis, the modal method is implemented in two ways. One of them generalizes our previous temperature modal pipeline by expanding, for each shape, the corresponding TTT , EEE , TTE and EET bispectra. We then exploit separability to build the covariance matrix of these expanded bispectra (Liguori et al. 2015), and to measure f_{NL} efficiently using Eq. (36). This modal pipeline will be referred to throughout the paper as the “Modal 1” pipeline.

The other implementation, which we will refer to as “Modal 2”, utilizes a novel approach where the $a_{\ell m}^T$ and $a_{\ell m}^E$ are first orthogonalized to produce new uncorrelated unit variance $\hat{a}_{\ell m}$ coefficients,

$$\hat{a}_{\ell m}^T = \frac{a_{\ell m}^T}{\sqrt{C_\ell^{TT}}} \quad (41)$$

$$\hat{a}_{\ell m}^E = \frac{C_\ell^{TT} a_{\ell m}^E - C_\ell^{TE} a_{\ell m}^T}{\sqrt{C_\ell^{TT} C_\ell^{EE} - C_\ell^{TE2}}}. \quad (42)$$

We then decompose the new bispectra as

$$\hat{b}_{\ell_1 \ell_2 \ell_3}^{X_1 X_2 X_3} = \sum_{m_i} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix} \langle \hat{a}_{\ell_1 m_1}^{X_1} \hat{a}_{\ell_2 m_2}^{X_2} \hat{a}_{\ell_3 m_3}^{X_3} \rangle, \quad (43)$$

which can be constrained independently, since they are uncorrelated. In this case the estimator then takes on a particularly simple form (Fergusson 2014). This new form is mathematically equivalent to the previous modal method, but involves significantly fewer terms in the estimator. However, due to the orthogonalization procedure we cannot constrain the full EEE bispectrum without further processing, just the additional part which is orthogonal to temperature. For this reason, although the “Modal 2” $T+E$ results incorporate all the polarization information, the EEE results alone are not presented here.

In our analysis, both modal techniques (together with all the other estimators described in this section) were used to measure f_{NL} for the three main shapes i.e., local, equilateral, and orthogonal. Besides this, we optimized the two pipelines for different purposes. The “Modal 1” estimator was adopted to perform a large number of robustness tests of our results, especially in relation to the local, equilateral and orthogonal measurements. The “Modal 2” pipeline was instead mostly used to study a large number of “non-standard” primordial shapes (e.g., oscillatory bispectra). For this reason, each pipeline uses a different set of basis templates. The “Modal 1” estimator starts from a polynomial basis with 600 modes, and includes nine more modes that are the contributions from last scattering to the exact radial KSW expansion of the local, equilateral and orthogonal templates. The “Modal 2” expansion uses a high-resolution basis with 2000 polynomial modes, augmented with a Sachs-Wolfe local bispectrum template, in order to improve convergence efficiency in the squeezed limit. In this way, the high resolution estimator provides the ability to scan across a wide variety of non-separable and oscillatory shapes, while the lower resolution pipeline gives efficient convergence in the f_{NL} measurements for the standard

⁷ We note that the local, equilateral, and orthogonal templates of Sect. 2 are separable. In fact, while the theoretical local NG models are manifestly separable, the equilateral and orthogonal templates of Eqs. (3) and (4) are factorizable approximations of the original non-separable shapes, that were derived exactly with the purpose of allowing the application of this type of estimator (Creminelli et al. 2006; Senatore et al. 2010).

local, equilateral, and orthogonal shapes, offering rapid analysis for validation purposes. The “Modal 1” pipeline can also be generalized for the estimation of parity-odd bispectra, which is included in our analysis of non-standard shapes.

3.3. Binned bispectrum estimator

One can also use a binned estimator (Bucher et al. 2010, 2016), exploiting the fact that the theoretical bispectra of interest are generally smooth functions in ℓ -space. As a result, data and templates can be binned in ℓ with minimal loss of information, but with large computational gains from data compression. The data bispectrum in the binning grid is then computed and compared to the binned primordial shapes to obtain f_{NL} . No KSW-like approach, which requires separability and mixing of theoretical and observational bispectra in the computation, is required. Instead, the binned data bispectrum and the binned theoretical bispectrum and covariance are computed and stored completely independently, and only combined at the very last stage in a sum over the bins to obtain f_{NL} . This means that it is very easy to test additional shapes or different cosmologies, and the data bispectrum can also be studied on its own in a non-parametric approach. In particular the smoothed binned bispectrum approach, also used in this paper, investigates the (smoothed) binned bispectrum of the map divided by its expected standard deviation, to test if there is a significant bispectral NG of any type in the map. Another advantage of the binned bispectrum estimator is that the dependence of f_{NL} on ℓ can be investigated for free, simply by leaving out bins from the final sum.

In more detail, the computation for the binned bispectrum estimator is based on Eqs. (38) and (39). However, instead of using the reduced bispectrum $b_{\ell_1\ell_2\ell_3}^{X_1X_2X_3}$, all expressions start from the alternative rotationally-invariant reduced bispectrum $B_{\ell_1\ell_2\ell_3}^{X_1X_2X_3} = h_{\ell_1\ell_2\ell_3}^2 b_{\ell_1\ell_2\ell_3}^{X_1X_2X_3}$, where h is defined in Eq. (40). The expression in Eq. (39) for the inner product remains the same when replacing b by B , except that the h^2 becomes h^{-2} . The importance of $B_{\ell_1\ell_2\ell_3}^{X_1X_2X_3}$ is that it can be determined directly from maximally-filtered maps;

$$B_{\ell_1\ell_2\ell_3}^{X_1X_2X_3} = \int d^2\hat{\mathbf{n}} M_{\ell_1}^{X_1}(\hat{\mathbf{n}}) M_{\ell_2}^{X_2}(\hat{\mathbf{n}}) M_{\ell_3}^{X_3}(\hat{\mathbf{n}}), \quad (44)$$

where

$$M_{\ell}^X(\hat{\mathbf{n}}) = \sum_m a_{\ell m}^X Y_{\ell m}(\hat{\mathbf{n}}). \quad (45)$$

Binning is then implemented by adding a sum over all ℓ inside a bin to the expression for the filtered map given in Eq. (45), thus obtaining the observed binned bispectrum of the map $B_{i_1i_2i_3}^{X_1X_2X_3, \text{obs}}$, with bin indices i_1, i_2, i_3 . The linear correction term is obtained in a similar way (and subtracted from the cubic term), but with two of the maps in Eq. (44) replaced by Gaussian simulations, and taking the average over a large number of those.

The theoretical templates are binned simply by summing the exact expression over the ℓ s inside a bin,⁸ and the same is true for the covariance matrix $h_{\ell_1\ell_2\ell_3}^2 C_{\ell_1}^{X_1X_1} C_{\ell_2}^{X_2X_2} C_{\ell_3}^{X_3X_3}$. The binning is optimized in such a way as to maximize the overlap, defined using the inner product of Eq. (39) between the binned and the exact template for all shapes under consideration. Finally the estimate

⁸ We note that the enormous computational gain of the binned bispectrum estimator comes from the binned determination of the observed bispectrum; determining the theoretical bispectrum templates is fast, even when done exactly.

of f_{NL} is computed using Eq. (38), where the inner product now contains a sum over bin indices i instead of multipoles ℓ , and the bispectra and covariance matrix are replaced by their binned versions.

3.4. Data set and simulations

In the following, we will apply our bispectrum estimation pipelines both to simulations and data, and consider a large number of shapes, either primordial or non-primordial in origin. Simulations will be used for a wide range of purposes, from comparisons of the outcomes of different estimators, to tests of instrumental systematics and foreground contamination, as well as Monte Carlo evaluation of error bars. For this reason, many different sets of simulated maps will be used, with features that will vary, depending on the specific application, and will be described case by case throughout the paper. Most of the time, however, we will use the FFP8 simulation data set described in Planck Collaboration XII (2016), or mock data sets obtained by processing FFP8 maps in various ways. These are the most realistic simulations available, modelling the CMB sky and the instrumental effects of *Planck* to the best of our current knowledge. They have passed through the same steps of the component separation pipelines as the real sky map and are the same maps as used for the final validation of the estimators in Sect. 5.3.

As far as actual data are concerned (Planck Collaboration I 2016; 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) the maps analysed in this work are the *Planck* 2015 sky map, both in temperature and in E polarization, as cleaned with the four component separation methods SMICA, SEVEM, NILC, and Commander (Planck Collaboration IX 2016). As explained in Planck Collaboration VII (2016), the polarization map has had a high-pass filter applied to it, since the characterization of systematics and foregrounds in low- ℓ polarization is not yet satisfactory. This filter removes the scales below $\ell = 20$ completely, and those between $\ell = 20$ and $\ell = 40$ partially. In all our analyses we use $\ell_{\text{min}} = 40$ for polarization, in order to be independent of the details of this filter. For temperature, we use $\ell_{\text{min}} = 2$. All the final cleaned maps are smoothed with a 5' Gaussian beam in temperature, and a 10' Gaussian beam in polarization.

The maps are masked to remove the brightest parts of the Galaxy as well as significant point sources. The masks used are the common masks of the *Planck* 2015 release in temperature and polarization, which are the union of the confidence masks for the four component separation methods⁹ (Planck Collaboration IX 2016). The sky coverages are respectively $f_{\text{sky}} = 0.76$ in temperature and $f_{\text{sky}} = 0.74$ in polarization. The stability of our results as a function of the mask is investigated in Sect. 7.2, where we show that our temperature and joint temperature plus polarization results do not change significantly when we consider a larger sky coverage.

In Sects. 7.1 and 7.3 we also compare the performance of different component separation methods, and conclude that, with

⁹ We note that the *Planck* collaboration produced two slightly different sets of union masks (see Planck Collaboration IX 2016 for details). We choose to adopt the more conservative set in this paper, as we found that the agreement between different component separation methods significantly increases with these masks when we measure f_{NL} of shapes that peak in the squeezed limit (while the differences are very small in other cases).

respect to bispectrum estimation, the most accurate results are obtained using the SMICA data set. As already done for the 2013 release, we will thus consider SMICA as our main data set, using the other methods for important cross-checking purposes.

If we consider only temperature, current SMICA data become noise dominated at $\ell \simeq 2000$, while previous nominal mission data were noise dominated at $\ell \simeq 1700$. The mask used for the 2013 release was also slightly larger than the current one ($f_{\text{sky}} = 0.73$ in 2013 vs. $f_{\text{sky}} = 0.76$ in 2015). Since the f_{NL} signal-to-noise ratio, as quantified by the Fisher Information Matrix, scales as $(\text{S/N}) \propto \ell \sqrt{f_{\text{sky}}}$ in the signal dominated regime and saturates in the noise-dominated regime, we expect an improvement in our f_{NL} temperature constraints of about 20% when going from the 2013 nominal mission release to the current results. Adding polarization and accounting for all possible *TTE*, *EET*, and *EEE* bispectra produces further improvements. Since we are neglecting the first 40 polarization multipoles, such improvements are expected to be fairly small for shapes peaking in the squeezed limit, and more pronounced for equilateral type bispectra. A Fisher matrix approach shows that error bars are expected to improve by about 10% for the local shape and about 40% for the equilateral shape. This is in good agreement with our actual measurements, as can be seen from the results presented in Sect. 6 onwards.

3.5. Data analysis settings

Now we detail the general setup adopted for the analysis of *Planck* 2015 data by the four different optimal bispectrum estimation pipelines, described in previous sections.

As already explained, inpainting of the masked regions of the sky is a preliminary data filtering operation that all pipelines must perform, in order to retain optimality. We found that the inpainting method used in 2013 (Planck Collaboration XXIV 2014) for temperature maps still works well when polarization is included (note that it is the original *T*, *Q*, and *U* maps that should be inpainted, not the derived *E* map). We adopt a simple diffusive inpainting method. First the masked regions of the map are filled with the average value of the rest of the map. Then the value of each masked pixel is replaced by the average value of its (generally eight) direct neighbour pixels. The latter step is repeated a fixed number of times (2000).¹⁰ Relevant, final computations in map space (see e.g., Eq. 44) are always done after remasking, so that the inpainted regions of the map are not used directly. The relevance of the inpainting procedure is that it reduces the effect of the sharp edges and the lack of large-scale power inside the mask leaking into the rest of the map during harmonic transforms.

For the linear correction term and to determine error bars, we use the FFP8 simulations (see Planck Collaboration XII 2016, and Sect. 3.4), filtered through the different component separation pipelines, using the same weights as used for the actual data when co-adding frequency channels. To compute all theoretical quantities (like the bispectrum templates and the ISW-lensing bias) we use the *Planck* 2015 best-fit cosmological parameters

¹⁰ For bispectrum purposes we found no difference between the results when performing the procedure without a buffer (the so-called ‘Gauss-Seidel’ method, where amongst the neighbours will be pixels both at the current and at the previous iteration) and with a buffer (the so-called ‘Jacobi’ method, where all neighbour pixels will be at the previous iteration), except that the former converges faster. We found 2000 iterations to work well in the ‘Gauss-Seidel’ case.

as our fiducial cosmology. However, results are quite insensitive to small changes in these parameters.

As pointed out in Sect. 3.4, low- ℓ multipoles are filtered out of the input polarization data set, so that all of our analyses will use $\ell_{\text{min}} = 40$ in polarization, and $\ell_{\text{min}} = 2$ in temperature. The choice of ℓ_{max} is dictated by the angular resolution of the cleaned maps, which is $5'$ in temperature, and $10'$ in polarization, and by the fact that the temperature data become noise-dominated at $\ell \simeq 2000$, while the polarization information saturates around $\ell \simeq 1000$. The KSW and binned estimators use $\ell_{\text{max}} = 2500$ for temperature, while the modal estimators use $\ell_{\text{max}} = 2000$. As shown explicitly in Sect. 7.4, results are completely stable between $\ell = 2000$ and $\ell = 2500$, so that this has no impact on f_{NL} . Similarly the binned estimator uses $\ell_{\text{max}} = 2000$ for polarization, while the other estimators use $\ell_{\text{max}} = 1500$, but again Sect. 7.4 shows that this difference is unimportant. The estimators also differ in the number of maps used to compute the linear correction term and the error bars, but generally it is of the order of 200. This difference is due to the different convergence properties of the estimators, some converging faster than others.

The binned bispectrum estimator uses 57 bins¹¹ for the analysis, which were determined by optimizing the correlation between the exact and the binned templates for the different shapes in temperature and polarization, as well as the full combined case. This is equivalent to minimizing the variance of the different f_{NL} parameters, where we focus in particular on the primordial shapes.

As previously explained, we use two different versions of the polarized modal pipelines, called ‘Modal 1’ and ‘Modal 2’ in the paper. Besides technical and conceptual implementation differences, the two modal estimators also use different sets of basis templates. The ‘Modal 1’ pipeline uses 600 polynomial modes, plus nine ‘KSW radial modes’, computed at last scattering, while ‘Modal 2’ has a basis formed by 2000 polynomial modes, augmented with a Sachs-Wolfe local bispectrum template. Due to the way polarization is implemented in the ‘Modal 2’ pipeline, it cannot determine results for *E*-only. More details and explanations of the different choices are provided in Sect. 3.2.

As already stressed, the use of several independent bispectrum estimators, and several completely independent component separation methods allows a remarkable level of cross-validation of our results in order to establish their robustness. The fact that the bispectrum estimators are statistically equivalent and produce practically optimal results will be established in Sect. 5. The validation of the component separation methods is described in Planck Collaboration IX (2016) and Sect 7.

4. Non-primordial contributions to the CMB bispectrum

Here we investigate several bispectra of non-primordial origin that are expected to be present in the data, and quantify their impact on our f_{NL} results. We devote particular attention to assessing potential biases that these NG signals might induce on the

¹¹ The boundary values of the bins are: 2, 4, 10, 18, 30, 40, 53, 71, 99, 126, 154, 211, 243, 281, 309, 343, 378, 420, 445, 476, 518, 549, 591, 619, 659, 700, 742, 771, 800, 849, 899, 931, 966, 1001, 1035, 1092, 1150, 1184, 1230, 1257, 1291, 1346, 1400, 1460, 1501, 1520, 1540, 1575, 1610, 1665, 1725, 1795, 1846, 1897, 2001, 2091, 2240, and 2500 (i.e., the first bin is [2, 3], the second [4, 9], etc., while the last one is [2240, 2500]).

primordial bispectra. When forecasting such biases, we assume the data analysis settings discussed in Sects. 3.4 and 3.5.

4.1. Non-Gaussianity from the lensing-ISW bispectrum

The correlation between the gravitational lensing of the CMB anisotropies and the integrated Sachs-Wolfe (ISW) effect gives rise to a secondary CMB bispectrum — characterized by an oscillatory behaviour and peaked on squeezed configurations — that is a well-known contaminant to the primordial NG signal (Hanson & Lewis 2009; Mangilli & Verde 2009; Lewis et al. 2011; Mangilli et al. 2013). The temperature-only 2013 *Planck* results (Planck Collaboration XXIV 2014; Planck Collaboration XIX 2014; Planck Collaboration XVII 2014) showed evidence for the first time for the lensing-ISW CMB bispectrum and associated bias. Based on the same methodology used for the 2013 *Planck* data analysis (Planck Collaboration XXIV 2014), here we update the computation of the lensing-ISW bispectrum and its bias to include the full mission temperature and polarization data.

As shown by Cooray & Melchiorri (2006), the direct lensing-ISW correlation in E -polarization due to rescattering of the temperature quadrupole generated by the ISW effect is negligible. However, as explained in Lewis et al. (2011), there is an important correlation between the lensing potential and the large-scale E -polarization generated by scattering at reionization. Because the lensing potential is highly correlated with the ISW signal, this also leads to a non-zero lensing-ISW bispectrum in polarization.

To determine $f_{\text{NL}}^{\text{LISW}}$, the amplitude parameter of the lensing-ISW bispectrum, one simply inserts the theoretical template for this shape into the general expression of Eq. (38). The template is given by (Hu 2000; Lewis et al. 2011)

$$\begin{aligned}
 b_{\ell_1 \ell_2 \ell_3}^{X_1 X_2 X_3, \text{LISW}} = & C_{\ell_2}^{X_2 \phi} \tilde{C}_{\ell_3}^{X_1 X_3} f_{\ell_1 \ell_2 \ell_3}^{X_1} + C_{\ell_3}^{X_3 \phi} \tilde{C}_{\ell_2}^{X_1 X_2} f_{\ell_1 \ell_3 \ell_2}^{X_1} \\
 & + C_{\ell_1}^{X_1 \phi} \tilde{C}_{\ell_3}^{X_2 X_3} f_{\ell_2 \ell_1 \ell_3}^{X_2} + C_{\ell_3}^{X_3 \phi} \tilde{C}_{\ell_1}^{X_1 X_2} f_{\ell_2 \ell_3 \ell_1}^{X_2} \\
 & + C_{\ell_1}^{X_1 \phi} \tilde{C}_{\ell_2}^{X_2 X_3} f_{\ell_3 \ell_1 \ell_2}^{X_3} + C_{\ell_2}^{X_2 \phi} \tilde{C}_{\ell_1}^{X_1 X_3} f_{\ell_3 \ell_2 \ell_1}^{X_3}. \quad (46)
 \end{aligned}$$

Here $C_{\ell}^{T\phi}$ and $C_{\ell}^{E\phi}$ are the temperature/polarization-lensing potential cross power spectra, and the tilde on \tilde{C}_{ℓ}^{TT} , \tilde{C}_{ℓ}^{TE} , and \tilde{C}_{ℓ}^{EE} indicates that it is the lensed TT , TE , or EE power spectrum. The functions $f_{\ell_1 \ell_2 \ell_3}^{T,E}$ are defined by

$$\begin{aligned}
 f_{\ell_1 \ell_2 \ell_3}^T &= \frac{1}{2} [\ell_2(\ell_2 + 1) + \ell_3(\ell_3 + 1) - \ell_1(\ell_1 + 1)], \\
 f_{\ell_1 \ell_2 \ell_3}^E &= \frac{1}{2} [\ell_2(\ell_2 + 1) + \ell_3(\ell_3 + 1) - \ell_1(\ell_1 + 1)] \\
 &\quad \times \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 2 & 0 & -2 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix}^{-1}, \quad (47)
 \end{aligned}$$

if $\ell_1 + \ell_2 + \ell_3$ is even and ℓ_1, ℓ_2, ℓ_3 satisfy the triangle inequality, and zero otherwise.

In this paper our main concern with the lensing-ISW bispectrum is not so much to determine its amplitude (although that is also of great interest), but to compute its influence on the primordial shapes. The bias $\Delta f_{\text{NL}}^{\text{P}}$ due to the lensing-ISW bispectrum on the estimation of a given primordial amplitude f_{NL}^{P} is given by

$$\Delta f_{\text{NL}}^{\text{P}} = \frac{\langle b^{\text{LISW}}, b^{\text{P}} \rangle}{\langle b^{\text{P}}, b^{\text{P}} \rangle}, \quad (48)$$

where the inner product is defined in Eq. (39).

Table 1. Bias in the three primordial f_{NL} parameters due to the lensing-ISW signal for the four component separation methods.

Shape	lensing-ISW f_{NL} bias			
	SMICA	SEVEM	NILC	Commander
T Local	7.5	7.5	7.3	7.0
T Equilateral	1.1	1.2	1.3	1.8
T Orthogonal	-27	-27	-26	-26
E Local	1.0	1.1	1.0	1.1
E Equilateral	2.6	2.7	2.5	2.9
E Orthogonal	-1.3	-1.3	-1.2	-1.5
$T+E$ Local	5.2	5.5	5.1	4.9
$T+E$ Equilateral	3.4	3.4	3.4	3.6
$T+E$ Orthogonal	-10	-11	-10	-10

Table 2. Results for the amplitude of the lensing-ISW bispectrum from the SMICA, SEVEM, NILC, and Commander foreground-cleaned maps, for different bispectrum estimators. Error bars are 68% CL; see the main text for how they have been determined.

Method	lensing-ISW amplitude			
	SMICA	SEVEM	NILC	Commander
T				
KSW	0.79 ± 0.28	0.78 ± 0.28	0.78 ± 0.28	0.84 ± 0.28
Binned	0.59 ± 0.33	0.60 ± 0.33	0.68 ± 0.33	0.65 ± 0.36
Modal2	0.72 ± 0.26	0.73 ± 0.26	0.73 ± 0.26	0.78 ± 0.27
$T+E$				
Binned	0.82 ± 0.27	0.75 ± 0.28	0.85 ± 0.26	0.84 ± 0.27

The values for the bias are given in Table 1. It should be noted that these are the results as computed exactly with Eq. (48). They can differ slightly from the ones used in e.g., Table 10, where each estimator adopts values computed using the approximations appropriate to the method. However, these differences are completely insignificant. As seen already in Planck Collaboration XXIV (2014), for T -only the bias is very significant for local and to a lesser extent for orthogonal NG. For local NG the bias is larger than the error bars on f_{NL} . We see that for E -only the effect is non-zero but not significant. For the full $T+E$ case, the bias is smaller than for T -only, but large enough that it is important to take into account.

The results for $f_{\text{NL}}^{\text{LISW}}$ can be found in Table 2. The polarized version of the template has only been implemented in the binned bispectrum estimator. Error bars have been determined based on FFP8 simulations as usual.¹² The KSW estimator implements the lensing-ISW template exactly, while the binned and modal estimators use approximations, as explained in Sect. 3. In particular for the binned estimator the correlation between the binned

¹² The average value of the lensing-ISW amplitude determined from the FFP8 simulations is around 0.85 of the expected value. This value is very consistent across bispectrum estimators and component separation methods, which provides a useful consistency test in its own right. Except for this effect, all other tests on the temperature FFP8 maps show them to be very robust and to behave as expected, for example in the determination of the lensing-ISW bias on the local shape. We took this effect into account by increasing all error bars in the table by the appropriate factor (i.e., dividing them by ≈ 0.85).




Fig. 1. The skew- C_ℓ spectrum for the lensing-ISW effect (red line with data points), from the temperature map. The blue curve is the theoretically-expected spectrum. Note that the points beyond $\ell = 1500$ are significantly correlated.

and exact lensing-ISW template is relatively low, since it is a difficult template to bin (unlike all the other templates considered in this paper), which is reflected in the larger error bars. Tests performed on FFP8, as well as other tests, demonstrate that the lower correlation does not lead to a bias compared to the other estimators. We will use the KSW results to draw our conclusions.

We see that temperature results from the full mission are consistent with the 2013 nominal mission (Planck Collaboration XXIV 2014). Including polarization yields results that also appear consistent and decrease the error bars. However, for now the $T+E$ conclusions should be considered preliminary, for the reasons related to polarization data discussed in detail in Sects. 6 and 7. The error bars will also improve when measured with the other bispectrum estimators. As already seen in 2013, the values for $J_{\text{NL}}^{\text{LISW}}$ are slightly low compared to the expected value of 1, but not significantly so. On the other hand, the detection of the lensing-ISW bispectrum is significant, even with our conservative rescaling of the error bars. The hypothesis of having no lensing-ISW bispectrum is excluded at 2.8σ using temperature alone, and improves to 3.0σ with the current preliminary result when including polarization. As mentioned above, the latter result is likely to improve with further analysis of the *Planck* data. In Fig. 1 we present the results of the skew- C_ℓ analysis for lensing-ISW NG for the T map, which illustrates that the instrument and data processing are not removing this expected NG signal from the data.

4.2. Non-Gaussianity from extragalactic point sources

The auto-bispectra of extragalactic point sources are a potential contaminant to primordial NG estimates at *Planck* frequencies. The basic modelling and methodology of this section follows the corresponding section in Planck Collaboration XXIV (2014).

Extragalactic point sources are divided into populations of unclustered and clustered sources. The former are radio and late-type infrared galaxies (see e.g., Toffolatti et al. 1998; González-Nuevo et al. 2005), while the latter are dusty star-




Fig. 2. The skew- C_ℓ spectrum for unclustered point sources (red line with data points), from the temperature map. The blue curve is the theoretical spectrum, given the amplitude determined with the KSW estimator.

forming galaxies constituting the cosmic infrared background (CIB; Lagache et al. 2005). The contamination due to both types of sources in NG estimators is handled via dedicated bispectrum templates which are fitted jointly with the primordial NG templates.

The unclustered sources have a white noise distribution, and hence constant polyspectra. Their reduced angular bispectrum template is thus

$$b_{\ell_1 \ell_2 \ell_3}^{\text{unclust}} = \text{constant}. \quad (49)$$

This constant is usually noted b_{PS} or b_{src} in the literature (e.g., Komatsu & Spergel 2001). This constant template is valid in polarization as well as temperature, since the polarization angles of point sources are less clustered than the source density. However, since not all these point sources are polarized, we do not measure the same sources in temperature and in polarization. In fact, there is no detection of the bispectrum of unclustered point sources in the cleaned *Planck* polarization map, unlike in the temperature map, where Table 3 (binned bispectrum estimator) and Fig. 2 (skew- C_ℓ s) show a clear detection.

The clustered sources (CIB) have a more complex bispectrum in temperature, reflecting the distribution of the large-scale structure and the clustering of galaxies in dark matter halos (Argüeso et al. 2003; Lacasa et al. 2012; Crawford et al. 2014). The *Planck* results have allowed the measurement of the CIB bispectrum at frequencies 217, 353, 545 GHz in the range $\ell \approx 200 - 700$ (Planck Collaboration XXX 2014). In this multipole range, a power law was found to fit the measurement, with an exponent consistent between frequencies. However, at lower multipoles theoretical models for the CIB power spectrum (e.g., Planck Collaboration XXX 2014) and bispectrum (Lacasa et al. 2014; Pénin et al. 2014) predict a flattening of the CIB power. We thus take the TTT CIB bispectrum template to be a broken power law,

$$b_{\ell_1 \ell_2 \ell_3}^{\text{CIB}} \propto \left[\frac{(1 + \ell_1/\ell_{\text{break}})(1 + \ell_2/\ell_{\text{break}})(1 + \ell_3/\ell_{\text{break}})}{(1 + \ell_0/\ell_{\text{break}})^3} \right]^q, \quad (50)$$

where the index is $q = 0.85$, the break is located at $\ell_{\text{break}} = 70$, and $\ell_0 = 320$ is the pivot scale for normalization. Dusty

Table 3. Joint estimates of the bispectrum amplitudes of unclustered and clustered point sources in the cleaned *Planck* temperature map, determined with the binned bispectrum estimator. The error bars have been determined using FFP8 simulations.

Map	$b_{\text{PS}}/(10^{-29})$	$A_{\text{CIB}}/(10^{-27})$
SMICA	5.6 ± 2.7	0.4 ± 1.4
SEVEM	7.9 ± 2.8	0.8 ± 1.4
NILC	9.3 ± 2.7	-0.3 ± 1.4
Commander	5.9 ± 3.2	1.4 ± 1.6

star-forming galaxies emit with a low polarization fraction, and polarization correlates only over the smallest scales, so that the CIB is negligibly polarized. We thus take vanishing templates for its polarized bispectra

$$b_{\ell_1 \ell_2 \ell_3}^{\text{CIB}, TTE} = b_{\ell_1 \ell_2 \ell_3}^{\text{CIB}, TEE} = b_{\ell_1 \ell_2 \ell_3}^{\text{CIB}, EEE} = 0. \quad (51)$$

Both point source templates, Eqs. (49) and (50), have been implemented in the binned bispectrum estimator described in Sect. 3. The results for these two templates applied to the *Planck* temperature map, cleaned with the four component separation methods, can be found in Table 3. Since the two templates are highly correlated, the results have been determined in a combined analysis. The results have also been determined jointly with the primordial local, equilateral, and orthogonal templates, and the lensing-ISW bias has been subtracted, but all of this makes a negligible difference. Contamination from unclustered sources is detected in all component-separated maps. However, A_{CIB} is not detected.

The order of magnitude of the bispectrum amplitudes found in Table 3 is consistent with expectations. Indeed, for radio sources at 217 GHz and with a flux cut based on the Planck ERCSC (Planck Collaboration VII 2011), Lacasa & Aghanim (2014), the forecasted b_{PS} is approximately 2×10^{-28} . For the CIB, the *Planck* 2013 measurement (Planck Collaboration XXX 2014) at 217 GHz gives $A_{\text{CIB}} \simeq 6 \times 10^{-27}$ when translated into dimensionless units. The results reported in Table 3 are consistent at the order-of-magnitude level with these estimates, although they are lower, because we are analysing cleaned maps.

The unclustered point source and CIB templates are highly correlated, at 93%. For this reason it was not deemed a priority for the other bispectrum estimators to implement the CIB template as well. Moreover, both point source templates are negligibly correlated with the primordial NG templates and the lensing-ISW template (the maximum being the correlation between equilateral and CIB templates at 2.7%, while correlations with the unclustered point source template are well below 1%). For this reason, and despite the detection of point sources in the cleaned maps, it makes no difference for the primordial results if point sources are included in a joint analysis or completely neglected.

An additional contaminant to the cosmological CMB bispectrum arises from the correlation between the gravitational lensing of the CMB anisotropies and the CIB anisotropies. This correlation was detected in Planck Collaboration XVIII (2014) using an optimal cross-spectrum estimator. The CIB-lensing bispectrum might couple with any of the primordial shapes. However, the amplitude of the CIB bispectrum is predicted to be small in the cleaned *Planck* maps and it has actually not been detected (see Table 3). The CIB-lensing bispectrum signal is frequency dependent, and it is mostly dominant in the very high *Planck* frequencies (see e.g., Curto et al. 2014).

4.3. Non-Gaussianity from residuals of the deglitching processing

Cosmic rays interacting with the cryogenic detectors induce spikes in timelines. These high-amplitude, fast-rising signals are followed by a decay tail. We observe three families of glitches, characterized by their temporal shape. The amplitude and time constants of the decays depend on which part of the satellite is hit (Catalano et al. 2014; Planck Collaboration X 2014). These random events are Poisson-distributed in time and produce highly non-Gaussian systematics.

A method has been developed to remove them directly at the time-ordered information (TOI) level. This process is performed iteratively, and is described in detail in Planck Collaboration X (2014). The short glitches are just flagged in the data, whereas for the long ones only the fast part is flagged, and the long tail is subtracted from the timeline. This procedure is not perfect, and there are residuals from the potentially biased errors in the fit, and the undetected glitches under the threshold of 3.2σ of the TOI noise rms. They could in principle produce a non-Gaussian signal in the final map. In addition, these residuals could interact with the mapmaking procedure at the destriping level, since the error on the offset determination could be non-Gaussian due to undetected glitches or a possible bias in the errors of the removal of tails. This is important, because in more than 95% of the TOI data, tails have been subtracted.

To estimate the effect of these residuals on the determination of NG, we created two sets of simulations (one including glitches and the other not) for every bolometer of the 143 GHz channel. We generated Gaussian CMB maps, and applied the full TOI processing with a realistic instrumental noise (Planck Collaboration VII 2016). In the simulations *with glitches*, we added glitches at the TOI level, following the properties measured in the data, and cleaned them with the procedure applied to the data. For the simulations *without glitches*, we have the same CMB and noise realization, but no glitches added at the TOI level.

We estimated the bias caused by glitches on the measurement of f_{NL} using the binned bispectrum estimator. The bias on f_{NL} induced by the glitch residuals g on a map T , including noise and CMB is given by $\langle \hat{f}_{\text{NL}}(T + g) - \hat{f}_{\text{NL}}(T) \rangle$, where the noise in the weighting of the estimator is determined from the simulations *with glitches* (as it would be for the data). Results are shown in Table 4. For most shapes, we detect no significant bias. The higher signal and high dispersion for the local shape might be due to a mis-calibration of the linear correction. In any case, for all shapes the bias due to glitches is a negligible correction to the value of f_{NL} , given its error bars, and we will not take it into account in the remainder of the paper.

5. Validation tests

During the work for the 2013 release, culminating in the NG results of Planck Collaboration XXIV (2014), the advantage of having multiple independent bispectrum estimator implementations was amply demonstrated. This allows for very useful cross-checking of results, both during development and for the final analysis, thus greatly improving the robustness of and confidence in the final results. For this new release we followed the same procedure, with the same three principal bispectrum estimators: KSW; binned; and modal, all of which had their pipelines updated to handle polarization data in addition to temperature.

Table 4. Results on the impact of cosmic ray residuals on the estimation of f_{NL} at 143 GHz, determined using the binned bispectrum estimator. We produced 10 simulations. We report the mean of the bias defined in the text, and the error on this mean. We also show the Fisher error bars on f_{NL} for these simulations.

	Local	Equilateral	Orthogonal	Diffuse PS ($\times 10^{29}$)	Lens-ISW
<i>T</i> -only					
bias mean	1.1 ± 0.6	0.8 ± 1.6	-1.0 ± 0.7	0.5 ± 0.2	0.01 ± 0.01
$\sigma_{f_{\text{NL}}}$	5.2	64	34	2	0.2
<i>E</i> -only					
bias mean	2.4 ± 5.8	-9.6 ± 8.7	-7.1 ± 14.8	0.0 ± 0.1	-3.0 ± 1.4
$\sigma_{f_{\text{NL}}}$	38	157	90	0.6	7.8
<i>T+E</i>					
bias mean	1.8 ± 1.0	-5.1 ± 2.2	0.1 ± 1.5	0.01 ± 0.04	0.01 ± 0.01
$\sigma_{f_{\text{NL}}}$	4.4	43	22	0.3	0.2

Beyond the usefulness of cross-checking, the three estimators complement each other and have different strengths. The KSW estimator can treat separable bispectrum templates without approximation, but it is more work to add new templates and non-separable templates cannot be handled at all. The binned and modal estimators can reconstruct the full bispectrum (smoothed in different domains), while the skew- C_ℓ extension of the KSW estimator can be used to investigate the bispectrum beyond f_{NL} . The binned bispectrum estimator is the fastest on a single map or a set of unrelated maps, but becomes slower than the other two on a large set of realizations based on the same settings, because the linear correction term cannot be precomputed. The modal estimator can investigate a wide selection of oscillating or otherwise rapidly changing bispectrum templates that would be difficult to bin, while the binned bispectrum estimator can quickly implement and determine the f_{NL} of an additional template or the effect of a different cosmology if the binned bispectrum of the maps has already been computed. The binned estimator gets the dependence of f_{NL} on ℓ for free with its results, while the modal estimator allows for a statistical investigation of the mode coefficients.

In this Section we show some of the validation tests, in particular for polarization. In Sect. 5.1 we investigate the agreement between estimators, map-by-map, on sets of successively more realistic maps. In Sect. 5.2 we show that the estimators are unbiased in the presence of a non-zero f_{NL} . Finally, in Sect. 5.3 we show that the estimators are essentially optimal on a set of the most realistic *Planck* simulations available, which are those used to compute the error bars on our final results.

5.1. Agreement between estimators on a map-by-map basis

The maps used in this subsection are realistic simulations of the CMB (at resolution $N_{\text{side}} = 2048$) but without any foregrounds. They do not contain any primordial NG, but do include ISW-lensing. Since the final FFP8 simulations were not yet available, the main goal was to make sure that the estimators agreed with each other, not only on average, but also on a map-by-map basis. For this purpose it was enough to look at only 49 maps. Establishing optimality of the estimators requires a larger number of maps, and is shown on the FFP8 simulations in Sect. 5.3.

In our first test we include the effect of the 143 GHz beam, but in other respects the simulations are ideal (no noise, and no mask). The analysis used $\ell_{\text{max}} = 2000$ for both *T* and *E*. The results for the average over the maps for the KSW, binned, and

both modal estimators, as well as for the difference between each estimator and KSW, are shown in Table 5. The shapes are assumed to be independent in this analysis, which means that the bias on the local shape due to the ISW-lensing effect is clearly visible. Results are shown for *T*-only, *E*-only, and the full combined *T+E* analysis. Note that the second modal implementation cannot compute results for *E* alone. One clearly sees that the results agree very well. It is also interesting to note that in this ideal noiseless case, one can actually determine f_{NL} more accurately from polarization alone than from temperature alone. This is due to the narrower transfer function in polarization, so that the primordial bispectrum is less smoothed in its projection to two-dimensional harmonic space.

The second test is identical to the first, except that we add realistic anisotropic noise realizations to the full-sky maps, based on the 143 GHz channel. The estimators now require the use of the linear correction term, and results are shown in Table 6. The agreement is still very good, although slightly worse than in the ideal case, as expected. The fact that the error bars for the *T*-only local case here are actually a bit smaller than in the ideal case is an artefact of the small number of maps; i.e., the error bars have not completely converged yet. On the other hand, the fact that the error bars for *E*-only are much larger than in the ideal case is a real effect; the *Planck* single-frequency polarization maps are noise-dominated.

Finally, the third test is identical to the second, except that we now also add a mask. The mask chosen is realistic, based on the union of the confidence masks provided by the SMICA, NILC, SEVEM, and Commander methods for this particular set of simulations. It contains both a Galactic and a point source part. The temperature mask leaves 79% of the sky unmasked, while the polarization mask leaves 76%. The results are shown in Table 7, while the map-by-map comparison is given in Fig. 3. From the table we see that the agreement between the different bispectrum estimators is still very good and only slightly degraded when compared to the previous case. The typical discrepancy between the bispectrum estimators, even in this most realistic case, is less than about a third of the uncertainty on f_{NL} . This is apparent in the map-by-map comparison of Fig. 3.

5.2. Validation of estimators in the presence of primordial non-Gaussianity

After the map-by-map comparison of the previous section, we next want to make sure that the estimators are unbiased. For

Table 5. Results from the different f_{NL} estimators for the set of CMB simulations described in Sect. 5.1 in the ideal case without noise or mask. Both the results for the estimators individually and for the differences with KSW are given, for T -only, E -only, and the full combined $T+E$ analysis. The shapes are assumed independent.

Shape	f_{NL}						
	KSW	Binned	Modal 1	Modal 2	B – KSW	M1 – KSW	M2 – KSW
T Local	7.6 ± 5.4	7.4 ± 5.6	7.4 ± 5.1	7.2 ± 5.7	-0.3 ± 0.6	-0.2 ± 0.4	-0.5 ± 2.2
T Equilateral	7 ± 53	5 ± 58	6 ± 53	8 ± 56	-2 ± 12	-1.0 ± 8.4	0 ± 17
T Orthogonal	-22 ± 27	-22 ± 28	-22 ± 27	-17 ± 30	0.5 ± 9.4	-0.2 ± 4.2	5 ± 11
E Local	-0.9 ± 4.1	-1.3 ± 3.4	-0.9 ± 3.7	...	-0.3 ± 2.9	0.1 ± 0.5	...
E Equilateral	-9 ± 42	-10 ± 42	-10 ± 40	...	-1 ± 11	-0.7 ± 9.2	...
E Orthogonal	4 ± 13	5 ± 13	4 ± 12	...	0.1 ± 3.8	-0.3 ± 2.7	...
$T+E$ Local	2.2 ± 3.1	1.5 ± 2.5	2.1 ± 2.8	2.0 ± 3.3	-0.6 ± 1.0	0.0 ± 0.8	-0.2 ± 1.9
$T+E$ Equilateral	0 ± 20	2 ± 22	3 ± 21	0 ± 23	1.4 ± 5.8	2.3 ± 7.3	0 ± 12
$T+E$ Orthogonal	-4 ± 10	-4 ± 9	-6 ± 9	-5 ± 12	0.3 ± 2.2	-1.1 ± 3.1	-1.0 ± 7.1

Table 6. As Table 5, but with noise and no mask.

Shape	f_{NL}						
	KSW	Binned	Modal 1	Modal 2	B – KSW	M1 – KSW	M2 – KSW
T Local	6.7 ± 4.8	6.4 ± 5.2	6.7 ± 4.7	7.0 ± 5.3	-0.3 ± 1.0	0.1 ± 0.4	0.3 ± 1.2
T Equilateral	11 ± 61	12 ± 65	9 ± 63	12 ± 62	1 ± 15	-1.9 ± 9.6	1 ± 12
T Orthogonal	-19 ± 31	-18 ± 34	-20 ± 32	-18 ± 35	1 ± 12	-1.3 ± 5.1	0.8 ± 8.8
E Local	-2 ± 29	-4 ± 28	-1 ± 29	...	-2 ± 12	0.4 ± 5.5	...
E Equilateral	1 ± 191	-18 ± 195	-6 ± 200	...	-19 ± 47	-7 ± 23	...
E Orthogonal	-6 ± 101	0 ± 107	-6 ± 102	...	6 ± 25	-0.3 ± 10	...
$T+E$ Local	4.9 ± 4.2	4.5 ± 4.4	5.0 ± 4.2	4.9 ± 4.9	-0.4 ± 1.2	0.1 ± 1.5	-0.0 ± 1.2
$T+E$ Equilateral	13 ± 46	11 ± 49	9 ± 48	13 ± 47	-2 ± 10	-4 ± 13	-0.3 ± 7.0
$T+E$ Orthogonal	-11 ± 22	-11 ± 24	-13 ± 22	-11 ± 24	0.0 ± 7.3	-1.3 ± 7.1	0.7 ± 4.5

Table 7. As Table 5, but with noise and a mask.

Shape	f_{NL}						
	KSW	Binned	Modal 1	Modal 2	B – KSW	M1 – KSW	M2 – KSW
T Local	6.5 ± 5.1	6.1 ± 5.3	6.4 ± 5.0	6.0 ± 5.3	-0.4 ± 1.5	-0.1 ± 0.7	-0.5 ± 1.3
T Equilateral	11 ± 73	9 ± 75	6 ± 76	11 ± 70	-2 ± 19	-5 ± 14	0 ± 12
T Orthogonal	-22 ± 37	-21 ± 37	-23 ± 36	-20 ± 37	2 ± 14	-0.9 ± 6.1	2.6 ± 9.2
E Local	4 ± 36	0 ± 35	5 ± 37	...	-4 ± 16	1 ± 13	...
E Equilateral	-32 ± 242	-49 ± 209	-38 ± 246	...	-17 ± 88	-6 ± 34	...
E Orthogonal	-9 ± 138	-7 ± 139	-7 ± 142	...	2 ± 45	2 ± 19	...
$T+E$ Local	5.1 ± 5.3	4.2 ± 5.1	4.8 ± 5.0	4.5 ± 5.2	-1.0 ± 1.7	-0.3 ± 1.7	-0.6 ± 1.3
$T+E$ Equilateral	19 ± 50	16 ± 50	15 ± 53	16 ± 45	-3 ± 14	-4 ± 19	-3.2 ± 9.8
$T+E$ Orthogonal	-12 ± 25	-11 ± 26	-13 ± 25	-11 ± 23	1.9 ± 8.7	-1.0 ± 9.9	1.4 ± 5.9

this purpose we prepared a different set of 100 T and E CMB simulations, still with cosmological parameters as determined by *Planck*. This time ISW-lensing is not present, but there is a nonzero local $f_{\text{NL}} = 12$. To these maps we add the same beam, anisotropic noise, and mask as before. We again take $\ell_{\text{max}} = 2000$, and the results are given in Table 8.

We see that all the estimators correctly recover the input value, both in temperature and in polarization. The results for the equilateral and orthogonal shapes are consistent with the fact that those templates have a non-zero correlation with the local

shape (the Table gives the results for an analysis where all shapes are assumed independent). For example, a *joint* analysis of the $T+E$ binned estimator gives $f_{\text{NL}}^{\text{local}} = 11.5 \pm 6.4$, $f_{\text{NL}}^{\text{equil}} = -7.5 \pm 5.1$, and $f_{\text{NL}}^{\text{ortho}} = -0.4 \pm 2.9$. Except for the first modal estimator in E -only (due to an insufficient number of maps in the linear correction term), we also find that the error bars for the bispectrum-based estimators are very close to the Fisher errors. Note that a slight increase in the error bars compared to Fisher estimates is expected for the local shape in T -only and in $T+E$, due to the




Fig. 3. Map-by-map comparison of the results from the different estimators for local (left), equilateral (centre), and orthogonal (right) f_{NL} (taking the shapes to be independent), for the third set of simulations described in Sect. 5.1, including both noise and a mask. Results are shown for T -only (top), E -only (centre), and the full combined $T+E$ case (bottom). The legend for the estimators can be found in the top right figure. The horizontal solid line is the average value of all maps for KSW, and the dashed and dotted horizontal lines correspond to $\pm 1\sigma$ and $\pm 2\sigma$ deviations, respectively.

signal being significantly different from zero there (the Fisher error bars for the local case are 5.8 for T -only, 26 for E -only, and 5.0 for $T+E$). Hence the estimators are effectively optimal, as will be illustrated in more detail in the next section.

5.3. Validation of estimators on realistic *Planck* simulations

As a final validation test, we ran our estimators on a large set of the most realistic simulations available. These are the FFP8 simulations (Planck Collaboration XII 2016) using SMICA for foreground separation. They are the same simulations we use to de-

termine the error bars on our final SMICA results in Sect. 6. They contain the Collaboration’s best estimates of the CMB sky and of *Planck*’s noise and beam effects, and have been cleaned by SMICA in the same way as the real sky map. The mask used is the same common mask defined for the real data analysis. For this test the estimators were all processed with the same settings used for the final data analysis.

Here we take 159 of these maps, and process these using all the estimators. By contrast, for the final results in Sect. 6, the convergence of the error bars of each estimator was carefully checked, using more maps if required. This explains why there

Table 8. Results from the different estimators for f_{NL} for the set of CMB simulations with $f_{\text{NL}}^{\text{local}} = 12$ described in Sect. 5.2. Results are given for T -only, E -only, and the full combined $T+E$ analysis. The shapes are assumed independent (see the main text for a discussion of this point).

Shape	f_{NL}							
	KSW		Binned		Modal 1		Modal 2	
T								
Local ...	11.2 ± 6.7	10.9 ± 6.3	11.9 ± 6.6	11.6 ± 6.6				
Equilateral	26 ± 78	24 ± 77	31 ± 82	27 ± 76				
Orthogonal	-33 ± 34	-33 ± 35	-34 ± 36	-33 ± 36				
E								
Local ...	11 ± 29	12 ± 26	9 ± 36	...				
Equilateral	34 ± 182	32 ± 153	10 ± 241	...				
Orthogonal	-37 ± 110	-28 ± 115	-31 ± 143	...				
$T+E$								
Local ...	11.3 ± 5.5	11.2 ± 5.0	11.1 ± 5.8	11.0 ± 5.4				
Equilateral	29 ± 52	24 ± 50	28 ± 54	24 ± 50				
Orthogonal	-29 ± 26	-28 ± 23	-30 ± 28	-26 ± 23				

are some small differences between the error bars in Sect. 6 and the ones presented here.

The results are shown in Table 9. Note that these are the results from an independent analysis, without subtracting the ISW-lensing bias. We also show the results from Minkowski functionals (for the local case only).¹³ We see that there is very good agreement between the bispectrum estimators even on this most complex and realistic set of simulations. The standard deviation of the difference between bispectrum estimators generally stays below one third of the error bar on f_{NL} , the only exception being the T -only equilateral result for the Modal 1 pipeline, which is still smaller than one half of the error bar. We see that the results from Minkowski functionals are consistent, but clearly suboptimal for f_{NL} . They are however a valuable, independent check.

The exact Fisher error bars for the nine shapes considered in the table are, in the same order as the table: 5.4, 69, 35; 31, 131, 74; 4.7, 43, 21. Taking into account the relative error in the standard deviation of $1/\sqrt{2(n-1)}$, which is 5.6% for 159 maps, we see that all bispectrum estimators are effectively optimal on all shapes, except for the E -only equilateral case where they appear slightly suboptimal. The small suboptimality of the Modal 2 pipeline for the local shape seen here disappears once more maps are used (see the results in Sect. 6).

In conclusion, all these validation tests show that we have very good agreement between the results from the different bispectrum estimators, not just on average, but also on a map-by-map basis. In addition we see that, despite the approximations made in the pipelines, and the simple treatment of the masked part of the maps (diffusive inpainting method and f_{sky} factor), the bispectrum estimators we use are all essentially optimal.

¹³ Since the Minkowski-functional pipeline automatically subtracts the ISW-lensing bias, the theoretical value for the bias as computed from the Fisher matrix has been added to its results, to make a direct comparison possible.

6. Results

6.1. Constraints on local, equilateral, and orthogonal f_{NL}

In this section we investigate the local, equilateral, and orthogonal primordial templates. These are now established as the standard shapes to study first when investigating the bispectrum (see Sect. 2 for a theoretical motivation and description of these shapes). However, they represent only the tip of the bispectral iceberg, and many more shapes are investigated in Sect. 8, while full model-independent reconstructions of the bispectrum are presented in Sect. 6.2.

For a complete description of the *Planck* data set and the bispectrum estimator configurations we have used, we refer the reader in particular to Sect. 3.4 and Sect. 3.5. To summarize the overall analysis methodology, we have employed four independent bispectrum estimators on the full mission *Planck* temperature and polarization maps obtained from the four different component separation pipelines, SMICA, SEVEM, NILC and Commander. The bispectrum estimators are the KSW estimator with its skew- C_ℓ extension using exact separable templates (Sect. 3.1), the Binned estimator using fixed multipole bins (Sect. 3.3), and the Modal 1 and Modal 2 estimators, which both use separable eigenmode expansions (Sect. 3.2). Temperature is analysed over the multipole range $\ell_{\text{min}} = 2$ to $\ell_{\text{max}} = 2000$ or above and polarization is analysed from $\ell_{\text{min}} = 40$ to $\ell_{\text{max}} = 1500$ or above (Sect. 3.5). By employing inpainting and a linear term, all these estimators essentially achieve optimality (as shown by comparison with Fisher matrix forecasts). The linear term in Eq. (36) and the uncertainties are determined using the FFP8 simulations (Sect. 3.5), also processed through the different foreground-separation pipelines. Our thorough validation campaign for these estimators is presented in Sect. 5.

The results of the analysis of the four cleaned maps with the four estimators, for T -only, E -only, and full $T+E$, are shown in Table 10, which is one of the main results of this paper. Results are determined while assuming all shapes to be independent, and are shown both with and without subtraction of the ISW-lensing bias (see Sect. 4.1 for more details about ISW-lensing). This bias is most important (relative to the size of the error bars) for the local shape, but also non-negligible for the orthogonal shape. Results here have not been marginalized over the point source contributions. While Sect. 4.2 shows that there is still a significant contamination by unclustered point sources in the cleaned maps, the correlation with the primordial templates is so small that this has no impact on the results reported here (as checked explicitly).

While Table 10 is the main result of this Section, in order to simplify the use of the *Planck* results by the wider scientific community, we also present in Table 11 the results that can be considered the final *Planck* 2015 results for the local, equilateral, and orthogonal shapes. As in 2013, we select the combination of the KSW estimator and the SMICA map for this. The SMICA map consistently performs well in all data validation tests, which are discussed in detail in Sect. 7. The KSW estimator, while unable to deal with non-separable templates, treats separable templates exactly, and the local, equilateral, and orthogonal template are all separable. On the other hand, the binned and modal estimators can deal with non-separable shapes and have other advantages as well (like full bispectrum reconstruction), but at the price of using approximations for the templates. However, they have all been optimized in such a way that the correlation with the exact templates for the three primordial shapes is close to perfect, so that in the end the results by the different estimators are statistically equivalent. Compared to the corresponding values in

Table 9. Results from the different estimators for f_{NL} for the set of SMICA simulations based on FFP8 described in Sect. 5.3. Both the results for the estimators individually and for the differences with KSW are given, for T -only, E -only, and the full combined $T+E$ analysis. The shapes are assumed independent and the lensing-ISW bias has not been subtracted.

Shape	f_{NL}												
	KSW		Binned		Modal 1		Modal 2		Mink.F.	B – KSW	M1 – KSW	M2 – KSW	MF – KSW
<i>T</i>													
Local	7.1 ± 5.5	7.0 ± 5.4	6.2 ± 5.5	6.3 ± 6.2	7 ± 12	-0.1 ± 1.1	-0.9 ± 1.9	-0.8 ± 1.9	0 ± 11				
Equilateral	2 ± 67	4 ± 67	-4 ± 73	5 ± 66	...	2 ± 19	-6 ± 32	3 ± 18	...				
Orthogonal	-23 ± 32	-24 ± 33	-24 ± 33	-20 ± 36	...	-1 ± 11	-0.9 ± 9.1	3 ± 14	...				
<i>E</i>													
Local	0.5 ± 32	0 ± 35	1 ± 30	...	0 ± 49	-0.8 ± 8.3	0.7 ± 8.3	...	0 ± 37	2 ± 35	0 ± 37
Equilateral	7 ± 144	7 ± 143	9 ± 152	0 ± 37	2 ± 35	...	2 ± 22	-1 ± 17
Orthogonal	5 ± 72	7 ± 75	4 ± 73	2 ± 22	-1 ± 17
<i>T+E</i>													
Local	5.6 ± 5.1	5.0 ± 4.9	4.7 ± 4.8	4.3 ± 5.3	5 ± 11	-0.6 ± 1.2	-0.9 ± 1.5	-1.3 ± 1.7	-1 ± 11				
Equilateral	3 ± 46	5 ± 44	3 ± 46	4 ± 43	...	2 ± 14	0 ± 14	1.0 ± 9.7	...				
Orthogonal	-10 ± 22	-9 ± 22	-9 ± 21	-7 ± 22	...	0.8 ± 7.0	0.8 ± 7.3	2.7 ± 7.7	...				

Table 10, the difference in the numbers in the last column of Table 11 is due to the fact that in the latter the equilateral and orthogonal f_{NL} have been determined jointly.

Focusing on the results for temperature-only and the full temperature plus polarization ($T+E$) results, we see that there is no evidence for a non-zero bispectrum with any of these three primordial shapes (local, equilateral, and orthogonal). After ISW-lensing subtraction, all f_{NL} values are consistent with 0 at 68% CL. The temperature results are all very similar to the ones from the nominal mission data published in 2013 (Planck Collaboration XXIV 2014), with very minor improvements in the error bars due to the additional temperature data. We also see that results are quite consistent when including polarization, with error bars shrinking by about 15% for local, 35% for equilateral, and 40% for orthogonal.

Table 10 displays very good agreement between the results from the different estimators, at the level expected from the validation tests in Sect. 5. We also note how the error bars, which are determined using the FFP8 simulations, are statistically indistinguishable from the optimal Fisher expectation.

Different component separation methods also show a good level of agreement when looking at temperature-only and combined temperature plus polarization results. The accuracy of this statement will be shown and quantified in detail in Sect. 7. However, in the same section, we will also show how the agreement between f_{NL} extracted from different cleaned maps becomes significantly degraded when considering *polarization-only* results¹⁴. The reasons behind this loss of internal consistency are not fully understood at present. Polarization data are, however, much noisier than temperature data, implying that the EEE bispectra have a close to negligible weight in the final combined measurement, which is dominated by the TTT and TTE configurations. In fact, as just mentioned above, the combined measurement looks perfectly self-consistent: local, equilateral and orthogonal f_{NL} measurements in the $T+E$ column of Table 10 pass *all* our tests of robustness.

We can thus conclude that, while highly challenging from a technical point of view, the inclusion of polarization in our es-

imator pipelines has been a success, allowing for a significant tightening of the constraints on the three standard primordial bispectrum shapes. On the other hand, in light of the outstanding issues in E -only results, we present our results conservatively, and recommend the reader to consider all f_{NL} constraints that make use of polarization data throughout this paper as *preliminary* at the current stage. We stress again that this is a conservative choice, which is made despite the fact that *no test* to date shows any evidence of leakage of the issues in EEE bispectra into the $T+E$ measurements. A detailed description of all the data validation tests, which lead to the robustness-related assessments summarized here, can be found in Sect. 7 (for readers less interested in the technical details, the main results and conclusions of all these tests are summarized in Sect. 7.6).

6.2. Bispectrum reconstruction

6.2.1. Modal bispectrum reconstruction

The starting point for modal bispectrum estimation is the robust extraction of the modal coefficients β_n from each of the full mission foreground-separated maps, that is, SMICA, SEVEM, NILC, and Commander. The β_n -coefficients are obtained for each of the temperature, polarization, and mixed bispectrum components, TTT , TTE , TEE , and EEE . Their cross-correlation between cleaning methods is an important validation of their accuracy, as we shall discuss in the next section, with excellent correspondence for temperature and some differences remaining in polarization. The modal basis number $n_{\text{max}} = 2001$ for the full mission analysis has been substantially increased offering a higher effective resolution when compared to the 2013 *Planck* Data Release where $n_{\text{max}} = 601$ modes were used. Several different basis functions have been used, including trigonometric functions, sinlog basis functions, and polynomials (closely related to Legendre functions), with the latter chosen because of excellent convergence in the squeezed and flattened limits.

We can reconstruct the full three-dimensional *Planck* bispectrum, obtained using these basis functions, to visualize its main properties and to determine robustness. A comparison between the temperature-only bispectra from the Nominal Mission and full mission at the same $n_{\text{max}} = 601$ modal resolution is shown

¹⁴ The E -only f_{NL} agreement is still at a reasonable $1\sigma_{f_{\text{NL}}}$ level in most cases. However this is larger than expectations from simulations, as described in Sect. 7.

Table 10. Results for the f_{NL} parameters of the primordial local, equilateral, and orthogonal shapes, determined by the KSW, binned and modal estimators from the SMICA, NILC, SEVEM, and Commander foreground-cleaned maps. Results have been determined using an independent single-shape analysis and are reported both without and with subtraction of the ISW-lensing bias; error bars are 68 % CL.

Shape	f_{NL}														
	Independent						ISW-lensing subtracted								
	KSW		Binned		Modal 1		Modal 2		KSW		Binned		Modal 1		Modal 2
SMICA T															
Local	10.2 ± 5.7	8.7 ± 5.4	6.8 ± 5.5	7.8 ± 6.0	2.5 ± 5.7	1.3 ± 5.4	0.5 ± 5.5	1.7 ± 6.0							
Equilateral . .	-13 ± 70	-26 ± 66	-16 ± 67	-12 ± 68	-11 ± 70	-27 ± 66	-12 ± 67	-13 ± 68							
Orthogonal . .	-56 ± 33	-41 ± 33	-47 ± 33	-63 ± 36	-34 ± 33	-14 ± 33	-20 ± 33	-44 ± 36							
SMICA E															
Local	26 ± 32	35 ± 34	20 ± 30	...	26 ± 32	34 ± 34	20 ± 30	...							
Equilateral . .	144 ± 141	156 ± 143	147 ± 159	...	144 ± 141	155 ± 143	147 ± 159	...							
Orthogonal . .	-128 ± 72	-128 ± 75	-137 ± 73	...	-128 ± 72	-126 ± 75	-137 ± 73	...							
SMICA $T+E$															
Local	6.5 ± 5.0	5.8 ± 4.9	4.0 ± 4.8	4.8 ± 4.9	0.8 ± 5.0	0.7 ± 4.9	-0.6 ± 4.8	0.7 ± 4.9							
Equilateral . .	3 ± 43	12 ± 44	5 ± 48	6 ± 42	3 ± 43	9 ± 44	3 ± 48	5 ± 42							
Orthogonal . .	-36 ± 21	-34 ± 22	-30 ± 21	-37 ± 21	-25 ± 21	-24 ± 22	-21 ± 21	-30 ± 21							
SEVEM T															
Local	11.3 ± 5.7	9.7 ± 5.4	8.1 ± 5.8	9.3 ± 6.0	3.6 ± 5.7	2.3 ± 5.4	1.4 ± 5.8	3.1 ± 6.0							
Equilateral . .	-3 ± 69	-16 ± 66	-11 ± 75	-6 ± 68	-2 ± 69	-18 ± 66	-12 ± 75	-7 ± 68							
Orthogonal . .	-59 ± 33	-47 ± 33	-49 ± 34	-66 ± 36	-36 ± 33	-20 ± 33	-23 ± 34	-48 ± 36							
SEVEM E															
Local	60 ± 42	62 ± 42	44 ± 38	...	60 ± 42	61 ± 42	44 ± 38	...							
Equilateral . .	292 ± 167	320 ± 154	302 ± 183	...	292 ± 167	318 ± 154	302 ± 183	...							
Orthogonal . .	-184 ± 91	-156 ± 93	-172 ± 91	...	-183 ± 91	-154 ± 93	-172 ± 91	...							
SEVEM $T+E$															
Local	9.3 ± 5.2	8.3 ± 4.9	6.4 ± 5.0	7.9 ± 5.0	3.3 ± 5.2	2.8 ± 4.9	2.1 ± 5.0	3.5 ± 5.0							
Equilateral . .	9 ± 47	21 ± 48	15 ± 52	5 ± 45	8 ± 47	17 ± 48	14 ± 52	4 ± 45							
Orthogonal . .	-50 ± 23	-46 ± 23	-44 ± 23	-55 ± 22	-39 ± 23	-35 ± 23	-33 ± 23	-47 ± 22							
NILC T															
Local	10.5 ± 5.6	8.7 ± 5.4	6.4 ± 5.6	8.0 ± 6.2	3.0 ± 5.6	1.4 ± 5.4	0.3 ± 5.6	2.2 ± 6.2							
Equilateral . .	-28 ± 69	-45 ± 66	-31 ± 75	-15 ± 66	-28 ± 69	-47 ± 66	-30 ± 75	-17 ± 67							
Orthogonal . .	-67 ± 33	-48 ± 33	-50 ± 33	-63 ± 35	-45 ± 33	-22 ± 33	-28 ± 33	-44 ± 35							
NILC E															
Local	0 ± 33	18 ± 36	-1 ± 30	...	-1 ± 33	17 ± 36	-2 ± 30	...							
Equilateral . .	75 ± 140	97 ± 141	64 ± 162	...	75 ± 140	96 ± 141	64 ± 162	...							
Orthogonal . .	-79 ± 76	-96 ± 81	-78 ± 77	...	-78 ± 76	-94 ± 81	-78 ± 77	...							
NILC $T+E$															
Local	6.9 ± 5.1	6.1 ± 4.9	3.3 ± 4.9	5.3 ± 5.2	1.2 ± 5.1	0.9 ± 4.9	-2.4 ± 4.9	4.4 ± 5.2							
Equilateral . .	-9 ± 44	-4 ± 44	-15 ± 50	8 ± 42	-9 ± 44	-7 ± 44	-16 ± 50	4 ± 42							
Orthogonal . .	-35 ± 21	-31 ± 22	-27 ± 23	-32 ± 21	-25 ± 21	-21 ± 22	-16 ± 23	-26 ± 21							
Commander T															
Local	9.6 ± 6.1	9.4 ± 5.7	6.4 ± 6.6	7.9 ± 6.3	4.0 ± 6.1	2.4 ± 5.7	1.4 ± 6.6	3.3 ± 6.3							
Equilateral . .	-19 ± 71	-36 ± 68	-3 ± 77	-14 ± 70	-20 ± 71	-38 ± 68	-4 ± 77	-18 ± 70							
Orthogonal . .	-49 ± 35	-38 ± 34	-49 ± 36	-45 ± 37	-29 ± 35	-12 ± 34	-25 ± 38	-28 ± 37							
Commander E															
Local	33 ± 39	56 ± 40	28 ± 37	...	33 ± 39	55 ± 40	28 ± 37	...							
Equilateral . .	327 ± 165	369 ± 157	278 ± 178	...	327 ± 165	368 ± 157	278 ± 178	...							
Orthogonal . .	-52 ± 88	-70 ± 88	-56 ± 81	...	-52 ± 88	-67 ± 88	-56 ± 81	...							
Commander $T+E$															
Local	7.7 ± 5.2	7.9 ± 5.0	5.2 ± 5.4	6.8 ± 5.2	3.7 ± 5.2	3.0 ± 5.0	1.6 ± 5.4	3.7 ± 5.2							
Equilateral . .	16 ± 46	26 ± 45	30 ± 50	29 ± 46	14 ± 46	23 ± 45	28 ± 50	26 ± 46							
Orthogonal . .	-37 ± 22	-37 ± 23	-39 ± 23	-35 ± 22	-29 ± 22	-27 ± 23	-30 ± 23	-28 ± 22							

in Fig. 4. Note the excellent agreement with all the main features replicated in the new data. In Fig. 4 in the third bispectrum, we also demonstrate the much higher bispectrum resolution achieved with the full $n_{\text{max}} = 2001$ modes. The tetrapyd shape reflects the constraints on the wavenumbers ℓ_1 , ℓ_2 , and ℓ_3 , with the squeezed configuration appearing on the axes that lie

along one $\ell_i = 0$. The expected ISW-lensing bispectrum is an oscillating signal in the squeezed limit along the tetrapyd edges; it is now measured with a significance of 3.0σ (see Sect. 4.1). This ISW-lensing signal sets an interesting benchmark or threshold against which to compare the other strong features observed in the bispectrum and now defined with greater precision. The orig-




Fig. 4. Modal bispectrum reconstruction for *Planck* 2013 (top left) and 2015 (top right) temperature-only data, both using the SMICA maps. Here, we restrict the 2015 resolution to the same as 2013, using similar polynomials with $n_{\max} = 601$. The two bispectra are very close to being in complete agreement in the signal-dominated regime shown up to $\ell_{\max} = 1500$. In the lower panel, we show the *Planck* 2015 temperature bispectrum at high resolution using the full $n_{\max} = 2001$ polynomial modes. Large-scale features in the top panels become subdivided but the main 2013 signals remain, notably a stronger measurement of the ISW-lensing signal (the regular oscillations in the squeezed limit).

inal “plus-minus” feature, with a large positive red peak around $\ell \approx 150$ followed by a larger negative peak near $\ell \approx 250$, remains though with more substructure, together with a broad negative peak in the equilateral limit around $\ell \approx 900$, which can be associated with the third acoustic peak from the transfer functions. Oscillatory models, which can connect these three peaks,

achieve higher significance. The apparent signal observed in the flattened limit remains, with a distinct pattern of blue and red features on the surface of the tetrapyd.

We also include a comparison with WMAP-9 in Fig. 5, where we have restricted the reconstructions to $\ell_{\max} = 600$ for comparison with $n_{\max} = 601$ modes. These plots, using identical




Fig. 5. Modal reconstruction for the WMAP-9 bispectrum (left) and the Planck SMICA 2015 T -only bispectrum (right) plotted for the domain $\ell \leq 450$, using identical isosurface levels. Here, we employed the full 2001 eigenmodes for both the *Planck* analysis at $\ell_{\max} = 2000$ and for WMAP-9 analysis at $\ell_{\max} = 600$, but for comparison purposes we have only used the first 600 eigenmodes in order to obtain a comparable resolution. The main features in the WMAP-9 bispectrum have counterparts in the *Planck* version, revealing an oscillatory pattern in the central region, as well as features on the tetrapyd surface. The WMAP-9 bispectrum has a much larger noise signal beyond $\ell = 350$ than the more sensitive *Planck* experiment, leading to apparent residuals in this region.

Table 11. Results for the f_{NL} parameters of the primordial local, equilateral, and orthogonal shapes, determined by the KSW estimator from the SMICA foreground-cleaned map. Both independent single-shape results and results with the ISW-lensing bias subtracted are reported; error bars are 68% CL. The difference between the last column in this table and the corresponding values in the previous table is that in the second column here the equilateral and orthogonal shapes have been analysed jointly. The final reported results of the paper are shown in bold.

Shape and method	$f_{\text{NL}}(\text{KSW})$	
	Independent	ISW-lensing subtracted
SMICA (T)		
Local	10.2 ± 5.7	2.5 ± 5.7
Equilateral	-13 ± 70	-16 ± 70
Orthogonal	-56 ± 33	-34 ± 33
SMICA ($T+E$)		
Local	6.5 ± 5.0	0.8 ± 5.0
Equilateral	3 ± 43	-4 ± 43
Orthogonal	-36 ± 21	-26 ± 21

isosurfaces, show the same bispectrum structure including the “plus-minus” feature clearly bisecting the main $\ell = 200$ peak and the first oscillation of the ISW-lensing bispectrum visible along the lower tetrapyd edges. The WMAP-9 reconstruction only shows significant differences from *Planck* in the top right region, where the higher noise levels in WMAP-9 make its reconstruction less reliable.

All four components of the temperature and polarization bispectrum reconstruction obtained from SMICA are shown in Fig. 6. A direct comparison of the EEE polarization bispectrum for SEVEM, NILC and Commander, is shown in Fig. 7, where we

note that these are orthogonalized E -mode contributions (see the Modal 2 discussion in Sect. 3). It is interesting to observe patterns of features evident in the polarization bispectra from the different foreground-cleaned maps, which, although inherently noisier, have qualitative similarities. At a quantitative level, however, the polarization bispectra modes from different methods are less correlated in polarization than in temperature, as we discuss in Sect. 7.

6.2.2. Binned bispectrum reconstruction

The (reconstructed) binned bispectrum of a given map is a natural product of the binned bispectrum estimator code (see Sect. 3.3). To test if any bin has a significant NG signal, we study the binned bispectrum divided by its expected standard deviation, a quantity for which we will use the symbol $\mathcal{B}_{i_1 i_2 i_3}$. With the binning used in the estimator, the pixels are dominated by noise. We thus smooth in three dimensions with a Gaussian kernel of a certain width σ_{bin} . To avoid edge effects, due to the sharp boundaries of the domain of definition of the bispectrum, we renormalize the smoothed bispectrum, so that the pixel values would be normal-distributed for a Gaussian map.

In Figs. 8 and 9, we show slices of this smoothed binned signal-to-noise bispectrum $\mathcal{B}_{i_1 i_2 i_3}$ with a Gaussian smoothing of $\sigma_{\text{bin}} = 2$, as a function of ℓ_1 and ℓ_2 . Very red or very blue regions correspond to a significant NG of any type. The two figures only differ in the value chosen for the ℓ_3 -bin, which is [518, 548] for the first figure, and [1291, 1345] for the second. We have defined two cross-bispectra here: $\mathcal{B}_{i_1 i_2 i_3}^{T2E} \equiv \mathcal{B}_{i_1 i_2 i_3}^{TTE} + \mathcal{B}_{i_1 i_2 i_3}^{TET} + \mathcal{B}_{i_1 i_2 i_3}^{ETT}$; and $\mathcal{B}_{i_1 i_2 i_3}^{TE2} \equiv \mathcal{B}_{i_1 i_2 i_3}^{TEE} + \mathcal{B}_{i_1 i_2 i_3}^{ETE} + \mathcal{B}_{i_1 i_2 i_3}^{EET}$. These two cross-bispectra are then divided by their respective standard deviations (taking into account the covariance terms) to produce the corresponding $\mathcal{B}_{i_1 i_2 i_3}^{T2E}$ and $\mathcal{B}_{i_1 i_2 i_3}^{TE2}$. Those three different permutations are not equal a priori due to the condition $i_1 \leq i_2 \leq i_3$, which is implemented in the code to reduce computations by a factor of six.




Fig. 6. CMB temperature and polarization bispectrum reconstructions for *Planck* SMICA maps using the full set of polynomial modes with $n_{\max} = 2001$ and with signal-to-noise weighting. The top bispectra are the symmetric pure temperature TTT (left) plotted with $\ell \leq 1500$ and E -mode polarization EEE (right) shown for $30 \leq \ell \leq 1100$. Below are the mixed temperature/polarization bispectra with TTE on the left (with E multipoles in the z -direction) and TEE on the right (with T multipoles in the z -direction). All S/N thresholds are the same.




Fig. 7. Comparison of CMB polarization bispectrum EEE reconstructions for *Planck* NILC, SEVEM, and Commander foreground-separated maps with signal-to-noise weighting. Note that these results are not as internally consistent between the four methods, also comparing SMICA shown in Fig. 6, which is closest to NILC. We will compare the underlying modal coefficients below to demonstrate these differences quantitatively.




Fig. 8. Smoothed binned signal-to-noise bispectrum \mathcal{B} for the *Planck* 2015 cleaned sky map, as determined with the binned estimator, as a function of ℓ_1 and ℓ_2 for a fixed ℓ_3 -bin [518, 548]. From left to right results are shown for the four component separation methods SMICA, SEVEM, NILC, and Commander. From top to bottom are shown: TTT , TTT cleaned from radio and CIB point sources; $T2E$, $TE2$; and EEE . The colour range is in signal-to-noise from -4 to $+4$. The light grey regions are where the bispectrum is not defined, either because it is outside the triangle inequality or because of the cut $\ell_{\max}^E = 2000$.

However, part of the smoothing procedure involves adding the other five identical copies, so that in the end the plots are symmetric under interchange of ℓ_1 and ℓ_2 (and $\mathcal{B}_{i_1 i_2 i_3}$ is symmetric under interchange of all its indices). The grey areas in the plots are regions where the bispectrum is not defined, either because

it is outside of the triangle inequality, or because of the limitation $\ell_{\max}^E = 2000$. Given that in both plots ℓ_3 is fixed at less than 2000, $\mathcal{B}_{i_1 i_2 i_3}^{TE2}$ is not defined if both ℓ_1 and ℓ_2 are larger than 2000,

while $\mathcal{B}_{\ell_1 \ell_2 \ell_3}^{EEE}$ is undefined if either ℓ_1 or ℓ_2 (or both) are larger than 2000.

Results are shown for the four component separation methods SMICA, SEVEM, NILC, and Commander, and for *TTT*, *T2E*, *TE2*, and *EEE*. In addition we show on the second line of each figure the result for *TTT* with the radio (unclustered) and CIB (clustered) point source bispectra subtracted according to their jointly measured amplitudes. It is clear, in particular in the second figure, that at higher ℓ there is a very significant point source contamination in the cleaned *TTT* bispectra, in agreement with the results of Table 3. However, after removing it we do not see a clear signal of any other residual NG. Of course this is for the moment only a qualitative statement; more quantitative tools for studying the amount of NG in these smoothed bispectra are in development.

Looking at the polarized bispectra in the high- ℓ_3 slices, in particular *TE2* and *EEE*, we do see some bluer and redder regions that might indicate residual NG. This agrees with statements made earlier, and discussed in greater detail in the next section, that the *Planck* polarized bispectrum is for the moment not as clean and well-understood as the temperature one. We also see a very good qualitative agreement between the four component separation methods in temperature, which worsens somewhat when mixing in more and more polarization; in particular SMICA and NILC give very similar results.

6.3. Primordial Curvature Reconstruction

In this section, we compress the information in T and E maps into maps of projected primordial curvature fluctuations, ζ . To the extent that the primary CMB temperature and polarization are Gaussian fields and the early universe ζ fluctuations are presumed Gaussian, all of the information in T, E is encoded in the mean field (Wiener-filtered map) plus fluctuations, characterized by a covariance matrix; it is just a different compression/re-expression of the original statistics. Thus, one could look for non-Gaussian deviations by directly evaluating the 3-point and higher statistics of this re-expression. Since fluctuations as well as the mean are given, a full description of the errors follows. That a Gaussian assumption is made does not mean that the Wiener-filtered map is in fact Gaussian, since the constraining data may drive it from a map consistent with Gaussianity. Hence one can search in the statistics of the Wiener-filtered map for evidence for non-Gaussianity.

The weighting of the temperature and E -mode polarization associated with the Wiener filter is an optimal one (inverse-total-covariance weighted), and estimators for non-Gaussianity constructed using this expression are nearly optimal. However, since a (fiducial) primordial power spectrum is assumed for the Wiener filter (usually a uniform n_s ζ -spectrum, using the best fit Planck 2015 scalar spectral index and power amplitude A_s), this approach is best suited for perturbative non-Gaussianity of the sort we treat in this paper. In practice, the estimators used here for non-Gaussianity act directly on temperature and polarization data, rather than through the intermediary of optimal ζ -maps and their fluctuations.

The scalar fluctuations can be expressed in terms of the curvature variable $\zeta(\mathbf{x}) = \ln a(\mathbf{x})$ on uniform total density hypersurfaces, where a is the inhomogeneous expansion factor, or, equivalently, by its wavenumber transform $\zeta(\mathbf{k})$. In turn, $\zeta(\mathbf{k})$ can be expanded in multipoles, $\zeta_{LMk}(k)$. Instead of the magnitude of the wavenumber k , a mixed representation gives a multipole expansion at each comoving distance from our location, χ , $\zeta_{LMk}(\chi)$,

$M = 0, \dots, L$. Here $\kappa = c$ or s , with $\kappa = c$ referring to the real part of ζ_{LM} and $\kappa = s$ to the imaginary part. For $M = 0$ there is no s component, only c . The mean of the ζ field given the temperature field T and its covariance describing allowed fluctuations about that mean are

$$\langle \zeta_{LMk}(\chi) | a_{LMk}^T \rangle = C_L^{\zeta T}(\chi) [C_L^{TT}]^{-1} a_{LMk}^T \quad (52)$$

$$\langle \delta \zeta_{LMk}(\chi) \delta \zeta_{L'M'k'}(\chi') | T \rangle = \delta_{LL'} \delta_{MM'} \delta_{kk'} \times \left\{ C_L^{\zeta \zeta}(\chi, \chi') - C_L^{\zeta T}(\chi) [C_L^{TT}]^{-1} C_L^{T \zeta}(\chi') \right\}. \quad (53)$$

In these expressions, C^{TT} is the total covariance for the temperature, including both signal and noise variances. The specific forms assume the C^{TT} matrix is diagonal in multipole space, depending only on C_L^{TT} , but if the noise is inhomogeneous it will have off-diagonal components and the equations become matrix equations. In this section, we assume the noise is diagonal. Eq. (52) shows the unsurprising result that for each LMk only one mode is determined by T . Replacing T by E gives $\langle \zeta_{LMk}(\chi) | a_{LMk}^E \rangle$ and $\langle \delta \zeta \delta \zeta^\dagger | E \rangle$. Although T and E are correlated, the uncorrelated part of E delivers a different mode for ζ from the one given by T . Thus when ζ is constrained by both T and E , the two modes deliver substantially more information than for T and E alone, and the fluctuations about the mean are thereby diminished. This is further helped by the acoustic oscillations of polarization being out of phase with those for T , so when T is down, E is not, and vice versa. The interplay of the two modes is quantified by signal-to-noise in Fig. 10.

When both T and E are included as constraints, a two by two matrix appears which includes the C^{TE} correlation in the off-diagonal as well as the C^{TT} and C^{EE} along the diagonal:

$$\langle \zeta_{LMk}(\chi) | a_{LMk}^T, a_{LMk}^E \rangle = \left[C_L^{\zeta T}(\chi) \ C_L^{\zeta E}(\chi) \right]^\dagger \begin{pmatrix} C_L^{TT} & C_L^{TE} \\ C_L^{ET} & C_L^{EE} \end{pmatrix}^{-1} \begin{bmatrix} a_{LMk}^T \\ a_{LMk}^E \end{bmatrix}. \quad (54)$$

$$\langle \delta \zeta_{LMk}(\chi) \delta \zeta_{L'M'k'}(\chi') | T, E \rangle = \delta_{LL'} \delta_{MM'} \delta_{kk'} \left\{ C_L^{\zeta \zeta}(\chi, \chi') - \left[C_L^{\zeta T}(\chi) \ C_L^{\zeta E}(\chi) \right]^\dagger \begin{pmatrix} C_L^{TT} & C_L^{TE} \\ C_L^{ET} & C_L^{EE} \end{pmatrix}^{-1} \begin{bmatrix} C_L^{T \zeta}(\chi') \\ C_L^{E \zeta}(\chi') \end{bmatrix} \right\}. \quad (55)$$

Fig. 11 shows an all-sky mean field (Wiener-filtered) map of the curvature variable ζ constructed in (densely-packed) shells from the multipoles $\zeta_{LMk}(\chi)$. The figure actually shows ζ projected onto differential visibility, $w = de^{-\tau}/d\chi$. The map $\zeta_w(\theta, \phi)$ is the spherical transform of $\zeta_{LMkw} = \int w(\chi) d\chi \zeta_{LMk}(\chi)$. The top left map is $\langle \zeta_w | T \rangle$ using the all-sky maps SMICA DX11. The middle left map is $\langle \zeta_w | E \rangle$, based on just the polarization information. (The E -maps did not have high pass filtering imposed, so that an indication of the full results obtainable with E alone can be seen.) The bottom left panel shows $\langle \zeta_w | T, E \rangle$. Visually the combined two-mode constraint has more detail than either of the one-mode constrained maps, enhancing the estimation of ζ_w . In each of the three cases, two sample realizations are shown which include allowed fluctuations about the mean map. These illustrate the level of uncertainty in the patterns found in the mean maps, necessary ingredients to the full statistical description. Note that the fluctuations about the mean are larger for T or E alone than for the combined T, E .

In all cases, the fluctuations become very large indeed if we allow for arbitrary projections, not conditioned by the visibility. Another approach similar to the differential visibility projection




Fig. 9. Similar to Fig. 8, but with $\ell_3 \in [1291, 1345]$.

is to make a slice at $\chi = \chi_{\text{dec}}$; i.e., making the projection a delta-function in χ at the peak where $dw/d\ln\chi = 0$, which defines χ_{dec} . This single slice case is shown in Fig. 12. The figures look quite similar in structure to the projected differential visibility plots. Although $\chi \sim \chi_{\text{dec}}$ is where most information is, the fluctuations of a single slice are somewhat larger than for the differential visibility projection. The differential visibility includes a subdominant “reionization bump”, with a small weight focused at low L where sample variance is large; restricting the projec-

tion of ζ to be just over the reionization region results in a low L mean map that is largely swamped by the allowed fluctuations about it. The fluctuations about the mean become much larger if we project over all χ , since there are vast terrains in χ in which there is very little CMB temperature or polarization information; in particular between recombination and reionization. For that uncharted territory, a realization of the reconstructed ζ -map reverts to a realization of the fiducial ζ -power.




Fig. 10. Signal-to-noise for T and E as a function of multipole L and comoving distance χ for an ideal cosmic-variance limited experiment (upper panels) and for the *Planck* experiment, with noise determined from FFP8 simulations (lower panels). The complementary nature of the information provided by T and E is evident. The differential visibility is shown for comparison.

Fig. 10 quantifies where the information resides in L - χ space. Here, the signal variance is the ensemble average of the square of the mean field, assuming the ζ are drawn from a Gaussian with fiducial covariance,

$$\langle\langle\zeta_{LM\kappa}(\chi)|T\rangle\langle\zeta_{LM\kappa}(\chi')|T\rangle\rangle = C_L^{\zeta T}(\chi)[C_L^{TT}]^{-1}C_L^{T\zeta}(\chi'), \quad (56)$$

and the noise (i.e., fluctuation) variance is the covariance matrix, as given by Eq. (54) for T and the equivalent for E . Fig. 10 plots the signal-to-noise at each individual χ slice (hence ignores the

components off-diagonal in χ):

$$[S/N]_L^T(\chi) = \frac{\rho}{\sqrt{(1-\rho^2)}}; \quad \rho \equiv \frac{C_L^{\zeta T}(\chi)}{\sqrt{C_L^{TT}C_L^{\zeta\zeta}(\chi,\chi)}} \quad (57)$$

for T , with an equivalent expression for E . The signal-to-noise structure in χ can be contrasted with the differential visibility plotted in the bottom panel: Fig. 10 shows that the L - χ information has a reach beyond the differential visibility structure, especially for low multipoles, because the associated waves can straddle the last scattering surface. The “reionization bump” in signal-to-noise seen in Fig. 10 is not that prominent. The Integrated Sachs Wolfe impact on the S/N is evident, but is relatively low, with the consequence that we cannot draw out high significance results from the ISW effect for ζ reconstruction (or equivalently for gravitational potential reconstruction).

The top panels in Fig. 10 are for an ideal experiment, with no noise (apart from the cosmic variance “noise” in the Wiener map fluctuations). The bottom panels of Fig. 10 include realistic *Planck* noise, as estimated from the FFP8 simulations: the high signal to noise in E in the top panels is noticeably diminished around $L \sim 100$ over what a cosmic variance limited experiment would give.

Figs. 11 and 12 are Gaussian-filtered on a relatively large $L \sim 200$ scale (40' FWHM). To see what happens at higher resolution, Figs. 13 and 14 zoom in on a typical $20^\circ \times 20^\circ$ patch, with long waves removed using a filter $W_L = \sin^2(L-L_c)/\Delta L$ for $L_c < L < L_c + \Delta L$. The specific choices are $L_c = 20$ and $\Delta L = 20$, the values used for the *Planck* high pass for polarization maps. To allow direct comparison, we have just done the same high pass filtering for the T map. We also removed the means of the maps. The resolution is $L \sim 400$ (20' FWHM). Figs. 13 and 14 illustrate that the fluctuations play a larger role in this higher resolution regime. Note that the χ_{dec} slice has about the same fluctuation level as the differential visibility projection.

Lensing effects are not taken into account in these ζ -maps. In principle one could de-lens the temperature and polarization maps before forming the Wiener filter. In practice this would be a highly noisy operation with current *Planck* data. Hence the contaminating influence of lensing on these ζ -reconstructions would be treated by comparing simulated ζ -maps with and without lensing. Such corrections are expected to be a subdominant bias, as it has been modelled in the rest of the paper.

7. Validation of *Planck* results

In this section, we perform a battery of tests aimed at verifying the robustness of the results obtained in the previous section. Table 10 shows excellent agreement with our 2013 analysis of nominal mission data (Planck Collaboration XXIV 2014). The agreement using different component separation methods in temperature is also generally very good. Our focus here is thus on polarization bispectra. Redundancy is perhaps the most important element in our analysis, as far as robustness is concerned. We devote considerable attention to comparing the outcomes of different estimators and component separation pipelines, and assess their level of internal consistency. We also verify the stability of our results in the harmonic and pixel domains, by considering different sky cuts and multipole intervals. Given the large computational requirements of these tests, and since results from different optimal estimators agree very well, as shown in previous sections, we will variously use the KSW, binned, or modal pipeline for different tests. In doing this we will also exploit




Fig. 11. Mean field and fluctuations for differential visibility projected ζ , as described in the text. The filter used in these maps is a $40'$ FWHM Gaussian. The grey regions are masked out.

the complementarity of different decompositions, which might make some of them more suited for different tests than others (for example, the binned pipeline directly works with a harmonic space decomposition of the bispectrum, thus making it perfectly suited for tests of ℓ -dependence, the modal pipeline can perform quick model-independent tests by working on a relatively small subset of bispectrum modes, and so on).

7.1. Dependence on foreground cleaning method

7.1.1. Comparison between f_{NL} measurements

In Table 10 we show f_{NL} results for the local, equilateral and orthogonal shapes, using four different optimal estimators, and four different foreground cleaning pipelines. The agreement between different estimators, on a given map, is within a fraction of a standard deviation, in line with theoretical expectations and simulations, as reported in Sect. 5. This level of agreement applies to all of the $T+E$, (TTT), and (EEE) bispectra.

The overall picture becomes more complex when comparing outputs across different foreground cleaning methods and esti-

matoms. Whereas for TTT and $T+E$ results the agreement also seems quite good in this case (being at the level of half a sigma or better, for nearly all combinations of cleaned maps and shapes), larger discrepancies are present in the EEE bispectrum measurements. The most notable differences are found for the equilateral shape, where SMICA and NILC find values of f_{NL} consistent with 0 within 1σ , while SEVEM and Commander measure a roughly 2σ deviation from Gaussianity. The largest discrepancy is found for the pair Commander–NILC, using the binned pipeline (see table 10). This estimator recovers $f_{\text{NL}}^{\text{equil}} = 369 \pm 160$ using the Commander E -map and $f_{\text{NL}}^{\text{equil}} = 97 \pm 141$ using the NILC E -map. Other pipelines, and different choices of component separation methods, show slightly smaller but similar discrepancies, at a level of about 1.5σ . The same shape and estimator analysis of temperature maps shows good agreement: Commander recovers $f_{\text{NL}}^{\text{equil}} = -36 \pm 73$, while NILC gives $f_{\text{NL}}^{\text{equil}} = -45 \pm 71$. The combined $T+E$ measurement, still for the same modal pipeline and equilateral shape, yields $f_{\text{NL}}^{\text{equil}} = 26 \pm 50$ for Commander and $f_{\text{NL}}^{\text{equil}} = -4 \pm 46$ for NILC, corresponding to about half a standard




Fig. 12. Mean field and fluctuations for a χ_{dec} slice of ζ , as described in the text. The filter used in these maps is a 40' FWHM Gaussian. The grey regions are masked out.

deviation difference. This general trend is seen for other shapes and estimators.

Simulations were used to give insight into the expected level of disagreement. For each of the four component separation methods, we generated 100 FFP8-based Gaussian simulations with realistic beams and noise levels. These simulations start from the same initial single frequency realizations, and are processed through the four different foreground cleaning pipelines. The starting maps *do not* include any foreground component (the same map generation procedure is used in the Monte Carlo determination of error bars). The differences in final simulations are thus generated only by the different data filtering and coadding operations performed either in pixel, harmonic, or needlet domains by the various foreground cleaning methods, and by additional manipulations of the maps which are required for f_{NL} estimation, such as inpainting. Therefore, the average scattering in f_{NL} , measured from these realizations, provides us with a baseline assessment of the expected discrepancies between different methods when foreground residuals and other spurious sources of NG are negligible. We can then compare them with differ-

ences observed on the data to establish whether they are consistent with expectations, or are too large. The latter would raise the concern that foreground contamination, or other systematics, might be affecting the results.

Results are shown in Table 12, for EEE and $T+E$ and two different sky coverages. The scatter between f_{NL} values from simulations is about 0.5σ for both $T+E$ and EEE . This is smaller than the differences in the *Planck* f_{NL} values obtained from EEE analysis of different foreground-cleaned maps, especially for the equilateral shape. However, for the final combined $T+E$ measurement, observed differences are in good agreement with expectations from simulations for the majority of cases. Another important point is that the consistency shown in Table 12 for $T+E$ measurements is stable to the change of sky coverage (in polarization) from $f_{\text{sky}} = 0.74$ to $f_{\text{sky}} = 0.64$. This will be confirmed by additional tests later in this section. For the SMICA-SEVEM pair we also verified stability using an even larger mask with $f_{\text{sky}} = 0.52$.

Residual foregrounds may be responsible for at least some of the observed excess of scatter in EEE -derived f_{NL} between




Fig. 13. Mean field and fluctuations for differential visibility projected ζ , for a 20 degree by 20 degree patch, as described in the text. The filter used in these maps is a 20' FWHM Gaussian. Sources in grey are masked out.

different cleaning algorithms. This is supported by the fact that several EEE results for this test change significantly for different masks, and that discrepancies are alleviated by using a larger mask, especially for equilateral shapes (see e.g., SEVEM–NILC and SMICA–SEVEM in Table 12). However, modal coefficients and their correlations are stable to a change of mask (see below), as are values of f_{NL} for a given component separation method (see Table 13).

Another possible contributor is a mismatch between the noise model (used to build the estimator normalization, weights and linear term), and the actual noise in the data. Polarization data are very noisy, and it is a known problem that the model assumed underestimates the true noise. This means that the error bars for EEE f_{NL} results, quoted in Table 10, are somewhat underestimated, which does not seem to be a problem for the final $T+E$ results, since the weight of the EEE bispectrum in the final combined measurement is very low. This is confirmed by the results of this test. Indeed, we investigate EEE in detail because it is a useful and sensitive indicator of various systematics in the polarized maps (which could eventually leak into the TTE and TEE bispectra), rather than for its statistical weight in the final measurement. It is then fair to say that issues in the EEE bispectra, and related f_{NL} measurements are not yet fully understood and will require further investigation. Even though the $T+E$ are

consistent, we recommend that results that include polarization data are regarded as *preliminary* at this stage.

7.1.2. Comparison between reconstructed bispectra

It is important to stress that the conclusions reached at the end of the previous subsection refer to the three main bispectra in our analysis, defined by the standard scale-invariant local, equilateral, and orthogonal primordial shapes. These shapes select a specific subset of configurations in the overall bispectrum domain (essentially squeezed, equilateral, and flattened triangles). Therefore, testing consistency between methods for these shapes does not guarantee that results for the many other NG models considered in this work (such as, for example, the oscillatory bispectra of Sect. 8) will display the same level of agreement. For this reason we decided to perform a model-independent test of consistency between methods, based on comparisons between the β_n eigenmodes used for bispectrum reconstruction in Sect. 6.2. We also reconstruct the bispectrum starting from a binned ℓ -decomposition, and this will be used in Sect. 7.4 to study stability of the results in the harmonic domain. For the β_n study we consider a simple test based on measuring the correlation coefficient between modes extracted from different foreground-cleaned maps. The correlation is defined, as usual,




Fig. 14. Mean field and fluctuations for a χ_{dec} slice of ζ , for a 20 degree by 20 degree patch, as described in the text. The filter used in these maps is a 20' FWHM Gaussian. Sources in grey are masked out.

by

$$r_{ij}^2 = \frac{\text{cov}(\beta_n^i, \beta_n^j)^2}{(\sigma_n^2)_i (\sigma_n^2)_j}, \quad (58)$$

and we measure it for each combination of the SMICA, SEVEM, NILC, and Commander maps, labelled by the indices i, j . Results are given in Table 14 and 15 for the two modal pipelines and are illustrated in Fig. 15. These results show an excellent degree of correlation between different maps in temperature (especially for SMICA, SEVEM, and NILC), which is reduced when polarization is considered. In fact the correlation for polarization is not much lower than temperature for SMICA and NILC, while it reduces the correlation for the pairs SMICA–SEVEM, and NILC–SEVEM, and for Commander when paired with any other method. This is consistent with previous findings of our f_{NL} -based test. To test if these results are due to foreground residuals (or other effects that are not included in the simulations), we evaluate the same mode-mode correlations on the same sets of 100 realistic, foreground-free, Gaussian simulations as previously used, and processed through each of the different component separation pipelines. For this analysis we consider TTT and EEE bispectra, expanded via the low-resolution modal estimator. Our results are reported in Table 14, in the simulation column, and they clearly show that the trend in the simulations is consistent with what we see in the *Planck* data. In particular, EEE results

show a lower degree of correlation in simulated maps, for the same pairs of methods. The observed loss of correlation in polarization does not seem to come from unresolved foregrounds or other unaccounted systematics, but rather something intrinsic to the foreground-removal algorithms. They are substantially different, as SMICA and NILC both perform the cleaning in harmonic space, at the level of E and B multipoles, whereas SEVEM is essentially a pixel-space template fitting method, performing the subtraction on Q and U maps, or inpainted before f_{NL} estimation. These issues will be studied in greater detail in future work, using Wiener-filtered, as well as inpainted maps for f_{NL} estimation. However, we have already seen that the larger scatter between modes from different foreground cleaning methods does not have a serious impact on f_{NL} estimation, at least for the standard local, equilateral, and orthogonal shapes. The non-standard shapes need to be analysed separately to check robustness of NG polarization results. This is the approach we will take for the various non-standard NG models.

7.2. Dependence on sky coverage

For each of the four component separation methods, we have used two different polarization masks, namely the same polarization mask as in Sect. 6, with $f_{\text{sky}} = 0.74$ (defined as the polarization “common mask” in Sect. 3.4), and an extended mask

Table 12. Comparison between local, equilateral, and orthogonal f_{NL} results, obtained using the four different component separation pipelines. For each pair of cleaning methods, and for each NG model, we compute the difference in the measured f_{NL} . The quoted error bar is the standard deviation of the same difference, extracted from a set of 100 realistic Gaussian simulations per method, not including foregrounds. These results have been obtained using the low resolution modal pipeline. See main text for comments and further details.

Methods	$f_{\text{NL}}(\text{method}_1) - f_{\text{NL}}(\text{method}_2)$					
	$f_{\text{sky}} = 0.74$			$f_{\text{sky}} = 0.64$		
	Local	Equilateral	Orthogonal	Local	Equilateral	Orthogonal
SMICA–SEVEM						
<i>T</i>	-1.2 ± 0.9	-6.0 ± 8.7	1.5 ± 4.8
<i>E</i>	-19 ± 21	-155 ± 86	34 ± 57	5 ± 22	-82 ± 90	-11 ± 66
<i>T+E</i>	-2.4 ± 1.6	-10 ± 18	13.5 ± 9.4	-1.5 ± 1.7	-12 ± 18	13 ± 10
SMICA–NILC						
<i>T</i>	0.4 ± 1.0	14.5 ± 8.9	2.5 ± 4.7
<i>E</i>	26 ± 11	83 ± 52	-59 ± 27	26 ± 13	32 ± 56	-96 ± 28
<i>T+E</i>	-0.7 ± 0.9	20.0 ± 8.2	-3.3 ± 3.8	0.6 ± 0.9	18.4 ± 8.4	-4.5 ± 4.0
SMICA–Commander						
<i>T</i>	0.4 ± 3.5	-14 ± 23	1.7 ± 14
<i>E</i>	-3 ± 16	-130 ± 77	-81 ± 42	-13 ± 17	-117 ± 100	-59 ± 40
<i>T+E</i>	-1.3 ± 3.2	-25 ± 18	9 ± 10	-1.4 ± 3.3	-26 ± 18	13 ± 10
SEVEM–NILC						
<i>T</i>	1.6 ± 1.0	20 ± 12	1.0 ± 4.5
<i>E</i>	45 ± 26	239 ± 94	-94 ± 69	30 ± 29	114 ± 105	-86 ± 79
<i>T+E</i>	3.1 ± 1.8	30 ± 18	-17 ± 10	2.2 ± 1.9	30 ± 18	-18 ± 10
SEVEM–Commander						
<i>T</i>	1.6 ± 3.4	-8 ± 22	0 ± 14
<i>E</i>	16 ± 22	25 ± 112	-116 ± 59	-18 ± 25	-35 ± 121	-48 ± 64
<i>T+E</i>	1.2 ± 3.3	-14 ± 21	-5 ± 11	0.2 ± 3.4	-14 ± 20	0 ± 11
NILC–Commander						
<i>T</i>	0.0 ± 3.0	-28 ± 22	-1 ± 12
<i>E</i>	-29 ± 21	-213 ± 84	-22 ± 54	-39 ± 23	-149 ± 108	38 ± 55
<i>T+E</i>	-1.9 ± 3.1	-45 ± 18	12 ± 11	-2.0 ± 3.2	-44 ± 17	18 ± 11

with $f_{\text{sky}} = 0.64$. The temperature mask is kept unchanged in this test, and it covers a sky fraction $f_{\text{sky}} = 0.76$ (the temperature “common mask” of Sect. 3.4). We report the variation in f_{NL} for the three standard shapes in Table 13, which shows insensitivity to f_{sky} , in agreement with earlier results on *T+E*. In this case, however, the *EEE* results also seem quite stable, supporting the view that foreground residuals are not affecting our local, equilateral, and orthogonal f_{NL} results, especially for the final, combined *T+E* measurements. Tests on FFP8 simulations including foregrounds (see Sect. 7.3) suggest that f_{NL} measurements obtained from the SMICA and SEVEM maps are the most accurate under the current choice of mask. As a further check of these two methods we consider a third polarization mask, with $f_{\text{sky}} = 0.53$, and repeat the combined *T+E* f_{NL} measurement, also finding stable results. For SMICA we find $f_{\text{NL}}^{\text{local}} = 5.6 \pm 5.4$, $f_{\text{NL}}^{\text{equil}} = 65 \pm 58$, and $f_{\text{NL}}^{\text{ortho}} = -30 \pm 26$, while for SEVEM we obtain $f_{\text{NL}}^{\text{local}} = 9.4 \pm 5.4$, $f_{\text{NL}}^{\text{equil}} = 75 \pm 59$, $f_{\text{NL}}^{\text{ortho}} = -50 \pm 30$.

We also perform model-independent checks by looking at the correlation coefficient between different sets of bispectrum modes, in a similar way to Sect. 7.1.2, but now changing the polarization mask. Results are reported in Table 16, and confirm that the data and simulations behave similarly, and that polarization modes display a lower correlation level than temperature.

7.3. Tests on simulations

We consider two realistic data simulations, one of which is Gaussian, while the other includes local NG. We start with a foreground-free realization, add foregrounds according to the Planck Sky Model, and finally process through the four component separation pipelines. By estimating f_{NL} in the input foreground-free simulation, for each method, and comparing to f_{NL} recovered from the cleaned maps (or with the input local f_{NL} , for the NG case), we can assess both the impact of foregrounds on our measurement before subtraction and which method gives the highest accuracy. The necessity to clean is very apparent in the middle set of columns in Table 17, where no cleaning has been performed.

SMICA and SEVEM give the best results, both in the G and NG tests. In the G test, reported in Table 17, SMICA results show an agreement between the input and the cleaned map at the level $\sigma_{f_{\text{NL}}}/2$ for all shapes, and for all of *TTT*, *EEE*, and *T+E*. SEVEM displays a similar level of accuracy, except for $f_{\text{NL}}^{\text{local}}$, where the difference is larger, but within one standard deviation. NILC and Commander clearly perform worse for the local shape, with NILC showing a $2\sigma_{f_{\text{NL}}}$ difference, and Commander even larger than that. In the NG test, reported in Table 18, SEVEM gives the most accurate results, recovering the input with $\sigma_{f_{\text{NL}}}/2$ accuracy or better. Results for SMICA are accurate at the $1\sigma_{f_{\text{NL}}}$ level for the *TTT* constraint, and worse (about $2\sigma_{f_{\text{NL}}}$) in *EEE*. However, the combined *T+E* measurement is again very good ($\sigma_{f_{\text{NL}}}/2$). NILC

Table 13. For each of the four foreground cleaned maps, we compute f_{NL} for the local, equilateral, and orthogonal modes using two different polarization masks, one with $f_{\text{sky}} = 0.74$ and the other with $f_{\text{sky}} = 0.64$, while for temperature we use a single mask with $f_{\text{sky}} = 0.76$. We then calculate the difference between the two measurements and compare with expectations from simulations, obtained in the following way: firstly, we generate realistic Gaussian realizations for each component separation pipeline, not including foregrounds; then, for each simulated map and NG model, we measure f_{NL} using the two masks in turn; finally, we calculate the standard deviation on 100 realizations. See the main text for more details and a discussion of these results, which were obtained using the low resolution modal pipeline.

	$f_{\text{sky}} = 0.74$			$f_{\text{sky}} = 0.64$			Difference		
	Local	Equilateral	Orthogonal	Local	Equilateral	Orthogonal	Local	Equilateral	Orthogonal
SMICA									
T	6.8 ± 5.4	-17 ± 66	-48 ± 33
E	25 ± 30	147 ± 159	-137 ± 73	48 ± 31	220 ± 168	-180 ± 81	-23 ± 16	-73 ± 68	43 ± 34
$T+E$. . .	4.0 ± 4.8	5 ± 46	-30 ± 21	4.6 ± 5.2	19 ± 55	-37 ± 22	-0.7 ± 1.2	-14 ± 14	6.7 ± 7.7
SEVEM									
T	8.1 ± 5.8	-11 ± 75	-49 ± 34
E	44 ± 38	302 ± 183	-172 ± 91	43 ± 39	303 ± 191	-170 ± 96	1 ± 19	0 ± 76	-2 ± 49
$T+E$. . .	6.4 ± 5.0	15 ± 52	-44 ± 23	6.2 ± 5.3	31 ± 54	-50 ± 25	0.2 ± 1.3	-16 ± 15	6.3 ± 8.8
NILC									
T	6.4 ± 5.8	-31 ± 76	-50 ± 33
E	-1 ± 30	64 ± 162	-78 ± 77	22 ± 30	190 ± 162	-84 ± 77	-23 ± 16	-124 ± 67	6 ± 37
$T+E$. . .	3.3 ± 4.9	-15 ± 50	-27 ± 23	4.0 ± 5.3	1 ± 56	-33 ± 23	-0.7 ± 1.3	-16 ± 13	5.4 ± 7.5
Commander									
T	6.4 ± 6.6	-3 ± 77	-49 ± 36
E	28 ± 37	278 ± 178	-56 ± 81	61 ± 38	337 ± 188	-122 ± 91	-32 ± 20	-60 ± 92	66 ± 47
$T+E$. . .	5.2 ± 5.4	30 ± 50	-39 ± 23	6.0 ± 5.7	45 ± 55	-51 ± 25	-0.7 ± 1.5	-15 ± 14	11.5 ± 8.9

Table 14. Correlation coefficients between pairs of bispectrum modes, extracted using two different component-separated maps. For both TTT and EEE we compare correlations measured from data with averages over 100 Gaussian realizations. The simulations are processed through the different component separation pipelines in the same way as the data, but they do *not* include any foregrounds. The correlation is clearly lower for EEE bispectra than for TTT . However, this is seen not only in data but also in simulations, indicating that it is not due to foreground residual contamination or other unaccounted for systematics. The results presented in this table are obtained using the low resolution modal pipeline, with 610 modes; results on data have also been cross-checked with the high-resolution modal estimator, using 2001 modes, and they are stable.

Methods	$f_{\text{sky}} = 0.74$				$f_{\text{sky}} = 0.64$	
	TTT		EEE		EEE	
	Data	Simul.	Data	Simul.	Data	Simul.
SMICA–SEVEM	0.97	0.97	0.61	0.62	0.60	0.61
SMICA–NILC	0.97	0.97	0.95	0.95	0.95	0.95
SMICA–Commander . . .	0.78	0.81	0.70	0.70	0.73	0.73
SEVEM–NILC	0.96	0.97	0.54	0.55	0.54	0.54
SEVEM–Commander . . .	0.81	0.83	0.69	0.67	0.70	0.70
NILC–Commander	0.85	0.86	0.64	0.63	0.66	0.66

is also performing very well in TTT and $T+E$, while the EEE result is more than $2\sigma_{f_{\text{NL}}}$ off.

The test described here has several limitations, the main and most obvious one being that it has been performed on just two maps (due to lack of availability of a large sample of this type of simulation at this stage). Another clear issue is that some methods, in particular *Commander*, seem to perform much better on actual data than on these simulations. On the other hand some important trends, observed in the data in previous tests, are clearly reproduced here, like the good stability of *SMICA* and *SEVEM*, especially for the combined $T+E$ results and, most no-

tably, the fact that the clear degradation in the accuracy of the EEE measurement for some methods does *not* seem to propagate to $T+E$.

7.4. Dependence on multipole number

In this section we discuss another stability test of our results, namely the dependence of the results for f_{NL} on the maximum and minimum multipole number used in the analysis. This test is most easily performed with the binned bispectrum estimator, since it gets the dependence of f_{NL} on ℓ for free with its stan-




Fig. 15. Scatter plots showing correlations between bispectrum modes extracted from the different *Planck* foreground-cleaned maps, for all possible pairs of component separation methods. *Top:* TTT bispectrum modes. *Bottom:* EEE bispectrum modes. While temperature shows a strong correlation, the loss of correlation in polarization between the different methods is evident in these plots, as discussed in the text and quantified in Tables 14 and 15. Results here and in Table 15 have been obtained using the high resolution modal pipeline (2001 modes), while results in Table 14 have been obtained with the low resolution modal pipeline. By construction, the high resolution pipeline is measuring not the full EEE bispectrum of the map, but the component of EEE that is orthogonal to TTT . For this reason, r^2 measured by the two pipelines for EEE will not be identical. With this caveat in mind, the agreement between the two modal approaches is very good.

standard analysis, simply by leaving out bins in the final sum when computing f_{NL} (the binned equivalent of Eq. (39)).

The dependence on ℓ_{max} of the results for the three standard primordial shapes (local, equilateral, and orthogonal modes), is shown in Fig. 16, for T -only, E -only, and full $T+E$. As mentioned in Sect. 6, the KSW and binned estimators use $\ell_{\text{max}} = 2500$ for temperature, while the modal estimators use $\ell_{\text{max}} = 2000$. As can be seen in the figure, both the T -only and $T+E$ results are basically unchanged between $\ell = 2000$ and $\ell = 2500$ for

Table 15. The r^2 statistic, Eq.(58), showing the degree of correlation between measured bispectrum β coefficients for the component separation methods shown in Fig. 15 (upper three rows for temperature, lower for polarization). Correlation between SMICA, NILC, and SEVEM is excellent in temperature; however, it declines markedly for the latter in polarization. Note that results are from the high resolution Modal 2 pipeline using the orthogonal EEE component only.

	SEVEM	NILC	Commander
SMICA (T)	0.95	0.94	0.63
SEVEM (T)		0.92	0.66
NILC (T)			0.72
SMICA (E)	0.39	0.89	0.55
SEVEM (E)		0.30	0.50
NILC (E)			0.43

Table 16. Correlation coefficients between pairs of EEE bispectrum modes, extracted using two different masks for each of the four component-separated maps. We compare correlations measured from data with Monte Carlo averages over 100 Gaussian realizations. The simulations were processed through the different component separation pipelines in the same way as the data, but do *not* include any foreground component. According to this test, modal expansions are stable for a change of sky coverage, with measured correlations in full agreement with expectations from simulations.

Method	EEE	
	Data	Simul.
SMICA	0.87	0.87
SEVEM	0.87	0.87
NILC	0.87	0.87
Commander	0.88	0.87

all three shapes, showing that this difference has no impact on the results (as was to be expected from the excellent agreement between estimators in Table 10). In fact, the values of f_{NL} for T and $T+E$ are reasonably stable (given their error bars) down to much lower values of ℓ_{max} , of about 1000.

On the polarization side we can draw a similar conclusion. The binned estimator uses $\ell_{\text{max}} = 2000$ for polarization, while the other estimators use $\ell_{\text{max}} = 1500$, but we see that results for E remain basically unchanged between $\ell = 1500$ and $\ell = 2000$. Central values and error bars for E for all three shapes have clearly converged by $\ell = 1500$, and are in fact reasonably stable down to much lower values of about 700.

As we noted in the 2013 analysis (Planck Collaboration XXIV 2014), when going to the much lower WMAP resolution of $\ell_{\text{max}} \simeq 500$, we agree with the slightly high value of $f_{\text{NL}}^{\text{local}}$ that the WMAP team reported (Bennett et al. 2013). This value is also confirmed when including polarization. One can clearly see the value of the higher resolution of Planck.

The dependence on ℓ_{min} is shown in Fig. 17. Here all estimators used the same values, $\ell_{\text{min}}^T = 2$ and $\ell_{\text{min}}^E = 40$. As explained in Planck Collaboration VII (2016), not all systematic and foreground uncertainties in the low- ℓ HFI polarization data

Table 17. Comparison of component separation methods, using Gaussian FFP8 simulations. We firstly consider Gaussian, foreground-free simulations, with simulated noise for each frequency band, process them through each of the four foreground cleaning pipelines, and measure f_{NL} for the three standard shapes (columns labelled with “Input map”). We then include foregrounds and repeat the measurement, *before* applying the cleaning, and including realistic noise levels for each method (columns labelled with “Input map + foregrounds”); this step is performed in order to get an idea of the level of contamination introduced by foregrounds, before cleaning. Finally, we apply the different component separation methods, and again estimate f_{NL} from the final maps (columns labelled with “Cleaned map”). The discrepancies between f_{NL} measured on the input map, and f_{NL} extracted from the cleaned map, provide a figure of merit to assess how well foregrounds are subtracted by different methods. Results below have been obtained with the KSW estimator and the “cleaned map” results were also checked with the binned estimator.

	Input map			Input map + foregrounds			Cleaned map		
	Local	Equilateral	Orthogonal	Local	Equilateral	Orthogonal	Local	Equilateral	Orthogonal
SMICA									
T	5.2 ± 5.8	29 ± 71	-8 ± 34	-107.0 ± 5.8	-23 ± 71	27 ± 34	7.8 ± 5.8	38 ± 71	-20 ± 34
E	-39 ± 30	-99 ± 133	59 ± 69	-10 ± 30	-154 ± 133	-41 ± 69	-56 ± 30	-120 ± 133	65 ± 34
$T+E$. . .	5.9 ± 5.1	14 ± 45	-20 ± 22	-118.0 ± 5.1	-32 ± 45	8 ± 22	8.3 ± 5.2	14 ± 45	-22 ± 22
SEVEM									
T	5.6 ± 5.7	32 ± 69	-8 ± 32	-113.2 ± 5.7	-8 ± 69	34 ± 32	12.7 ± 5.7	35 ± 69	-25 ± 32
E	-17 ± 41	-149 ± 175	28 ± 95	-14 ± 41	-171 ± 175	-44 ± 95	-22 ± 41	-120 ± 175	41 ± 95
$T+E$. . .	7.7 ± 5.3	12 ± 49	-37 ± 24	-126.0 ± 5.3	-29 ± 49	-57 ± 25	13.0 ± 5.3	11 ± 49	-41 ± 24
NILC									
T	5.1 ± 5.7	32 ± 69	-5 ± 31	-102.0 ± 5.7	-14 ± 69	32 ± 31	17.8 ± 5.7	85 ± 69	-16 ± 31
E	-52 ± 33	-157 ± 156	72 ± 73	-6 ± 33	-155 ± 156	-47 ± 73	-76 ± 33	-179 ± 156	113 ± 73
$T+E$. . .	5.7 ± 5.0	7 ± 46	-15 ± 21	-117.0 ± 5.9	-27 ± 46	12 ± 21	15.8 ± 5.0	-20 ± 46	-7 ± 21
Commander									
T	0.5 ± 6.2	-5 ± 73	-14 ± 36	-127.0 ± 6.2	-25 ± 73	-137 ± 36	25.6 ± 6.2	67 ± 73	-17 ± 36
E	-51 ± 38	-64 ± 160	93 ± 86	-10 ± 38	-153 ± 160	-45 ± 86	-70 ± 38	-78 ± 159	138 ± 86
$T+E$. . .	1.6 ± 5.4	-2 ± 48	-21 ± 23	-137.0 ± 5.4	-29 ± 48	13 ± 23	20.4 ± 5.4	28 ± 48	-11 ± 23

Table 18. Same test as in Table 17, but with a NG map as input, with $f_{\text{NL}}^{\text{local}} = 8.8$. For this case, we only report the final value after foreground cleaning for each method. Results below have been obtained with the KSW estimator and double-checked with the binned estimator. ISW-lensing contributions are removed.

	Cleaned map. Input $f_{\text{NL}}^{\text{local}} = 8.8$		
	Local	Equilateral	Orthogonal
SMICA			
T	3.1 ± 5.8	47 ± 71	-6 ± 34
E	-53 ± 30	-113 ± 133	94 ± 69
$T+E$. . .	5.7 ± 5.1	22 ± 45	-19 ± 22
SEVEM			
T	8.0 ± 5.7	43 ± 69	-11 ± 32
E	-19 ± 41	-112 ± 175	35 ± 95
$T+E$. . .	10.2 ± 5.3	19 ± 49	-37 ± 24
NILC			
T	10.2 ± 5.7	84 ± 69	7 ± 31
E	-76 ± 33	-179 ± 156	113 ± 73
$T+E$. . .	10.1 ± 5.0	20 ± 46	4 ± 21
Commander			
T	22.2 ± 6.2	81 ± 73	-5 ± 36
E	-68 ± 38	-78 ± 160	132 ± 86
$T+E$. . .	18.3 ± 5.4	35 ± 48	-9 ± 23

have been fully characterized yet, and hence it was decided to filter out these data.

For equilateral and orthogonal shapes the values for f_{NL} and their error bars are quite stable as a function of ℓ_{min} up to about

$\ell = 100$ (and $\ell \simeq 300$ for E -only), which is not surprising, since these templates have little weight at low ℓ . The local template, on the other hand, depends very strongly on the lowest multipoles, which is reflected in the very rapid growth of the error bars when ℓ_{min} increases. Looking at T -only and $T+E$ we see a very similar pattern, with $f_{\text{NL}}^{\text{local}}$ being reasonably stable, although there are some jumps. The local result for E -only wanders a bit more outside of the $\pm 1 \sigma$ region, in agreement with the other tests in this section, which also indicate that E -only is not as stable as T -only and $T+E$. However, that is still quite acceptable, given the small weight of E -only in the full $T+E$ result.

One can work out quite generally that when Y is a subset of a data set X , and P_X and P_Y are the values of a parameter P determined from these two data sets, then the variance of the difference $P_Y - P_X$ is equal to $|\text{Var}(P_Y) - \text{Var}(P_X)|$. Hence we can determine how likely the jumps in f_{NL} as a function of ℓ_{min} are. It turns out that the jump in the T -only value of $f_{\text{NL}}^{\text{local}}$ between $\ell_{\text{min}} = 40$ and $\ell_{\text{min}} = 53$ is a 2.5σ effect (using the values of $f_{\text{NL}}^{\text{local}}$ before subtraction of the ISW-lensing bias, which also depends on ℓ). Similarly, the jump in the T -only value of $f_{\text{NL}}^{\text{equil}}$ between $\ell_{\text{min}} = 154$ and $\ell_{\text{min}} = 211$ is (by chance) also a 2.5σ effect. Given the fact that there are 57 bins, having such a jump appears to be consistent from a statistical point of view.

7.5. A directional analysis with a needlet-based modal estimator

The validation tests on simulations, described in Sect. 7.3, point to SMICA and SEVEM as the best foreground cleaning methods for f_{NL} estimation. Results in Table 10 show that SMICA also has




Fig. 16. Evolution of the f_{NL} parameters (solid blue line with data points) and their uncertainties (dashed lines) for the three primordial bispectrum templates as a function of the maximum multipole number ℓ_{max} used in the analysis. From left to right the figures show, respectively local, equilateral, and orthogonal, while the different rows from top to bottom show results for T -only, E -only, and full $T+E$. To indicate more clearly the evolution of the uncertainties, they are also plotted around the final value of f_{NL} (solid green lines without data points, showing the $\pm 1 \sigma$ range around the dashed green lines). The results are for SMICA, and assume all shapes to be independent, and that the ISW-lensing bias has been subtracted. They have been determined with the binned bispectrum estimator.

slightly smaller error bars, thus making it the method of choice for our final results.

As a further check of residual foreground contamination in the SMICA map, in this section we investigate the possible directional dependence of SMICA-derived third-order statistics by means of a needlet-based modal estimator (i.e., an estimator based on the decomposition described in Sect. 3.2, and references therein, where we use cubic combinations of needlets as our basis modes). In other words, we analyse the behaviour of the needlet bispectrum (see Lan & Marinucci 2008, Rudjord et al. 2010 and Donzelli et al. 2012) on separate patches of the sky, and we study the fluctuations of the corresponding residuals.

Rather than assuming a specific anisotropic model, we instead calculate the contribution to the local f_{NL} from different regions of the sky and look for evidence of anisotropy in the result.¹⁵

¹⁵ Even though we focus here on directional contributions to the local shape (which was typically found as one of the most sensitive to residual

Our modal needlet estimator has been validated with respect to the procedures considered in Sect. 5, showing excellent agreement. Since in this paper we use the needlet estimator only in this section, and as a diagnostic tool, we will not explicitly report the outcome of these validation tests here, for the sake of conciseness. The advantages of using needlet-based modal estimators have been advocated in Lan & Marinucci (2008), Rudjord et al. (2010), Fergusson et al. (2010a), and Fergusson et al. (2012); we refer to these papers for more discussion and details. In short, however, they can be summarized as follows:

1. it is possible to achieve a strong data compression, i.e., to investigate cubic statistics by means of a small number of modes (needlet frequencies) for many different bispectrum templates;

foreground contamination), this type of directional analysis can be done in a model-independent way, by looking separately at different needlet modes; we leave this for future work.




Fig. 17. Evolution of the f_{NL} parameters (solid blue line with data points) and their uncertainties (dashed lines) for the three primordial bispectrum templates as a function of the minimum multipole number ℓ_{min} used in the analysis. From left to right the figures show respectively, local, equilateral, and orthogonal, while the different rows from top to bottom show results for T -only, E -only, and full $T+E$. To indicate more clearly the evolution of the uncertainties, they are also plotted around the final value of f_{NL} (solid green lines without data points, with $\pm 1 \sigma$ range around the dashed green line). The results are for SMICA, and assume all shapes to be independent, as well as that the ISW-lensing bias has been subtracted. They have been determined with the binned bispectrum estimator.

2. needlet transforms have good correlation properties in pixel space, which allows study of the pixel contribution to the f_{NL} signal for the templates under investigation, by treating different directions independently.

In our analysis, we first divide the sky into several large “regions”, with boundaries defined by the pixels of a HEALPix grid (Górski et al. 2005). at lower resolution than the starting map (which is at $N_{\text{side}} = 2048$). For the low resolution grid we consider N_{side} values between 2 and 8. For each pixel in the coarse grid, we then compute the local f_{NL} using our modal needlet estimator, and neglect contributions from external regions.

The correlation matrices between f_{NL} measurements in different regions were computed via Monte Carlo simulations. $N_{\text{side}} = 4$ was chosen, as providing the best tradeoff between having a large number of regions for directional analysis (i.e., the total number of pixels in the low resolution grid), and a low correlation between different regions. This is shown in Fig. 18. It is readily seen that, at $N_{\text{side}} = 4$, the correlation is largely concentrated in one or two points near the main diagonals (where it

is still low, never exceeding 34%), and falls off rapidly for all other pixels. Note that the results here refer to temperature only. The EEE local polarization error bars are large, even for the full sky analysis, making this directional approach uninformative for *Planck* polarization data. We concentrate on the TTT bispectrum here, complementing other validation tests in this section, which are mostly focused on polarization.

Having obtained our correlation matrices, and having shown that different regions are essentially uncorrelated, we can then proceed to extract f_{NL} for each region in the actual SMICA map. As a test of directional-dependent contamination of the f_{NL} measurement, we can also compare our results, region by region, with the fluctuations expected by looking at the diagonal of our Monte Carlo correlation matrix.

The results of this analysis are shown in Fig. 19. In the left panel, we represent the directional local f_{NL} map, extracted with this method. In the right panel, we report the f_{NL} values, region by region, and compare them to expectations from simulations. The red line gives the expected standard deviation, while the blue




Fig. 18. Temperature only, local shape, pixel correlation matrix from Monte Carlo analysis, at $N_{\text{side}} = 4$. The most correlated pixels are those closest to the main diagonal; however, these values are always lower than 34 % in the chosen case.

one gives estimates on the component-separated maps with the Monte Carlo error bars. Our estimator is normalized in such a way that the sum of all these contributions would yield exactly the f_{NL} estimator for the full map, so these results can be viewed as a partition of the estimates along the different directions. It is readily seen that no significant fluctuation occurs, so that our results are consistent with the absence of directionally-dependent features (which could occur due to, for instance, residual foreground contamination). As an additional check, we have also investigated the possible presence of a dipole in these data, and found that our results are consistent with Gaussian isotropic simulations.

7.6. Summary of the main validation results

Throughout Sect. 7 we have shown a battery of tests aimed at evaluating the robustness of our data set, from the point of view of bispectrum estimators, focusing especially on the polarization part. We studied the stability of our results (local, equilateral, and orthogonal f_{NL} measurements, model-independent bispectrum reconstruction) under a change of sky coverage, multipole range, and choice of component separation methods. We also considered simulated data sets and studied the ability of different component separation methods to recover the input f_{NL} after foreground subtraction. Our main conclusions from these tests can be summarized as follows:

- TTT and $T+E$ results are stable both in the pixel and harmonic domains, for different component separation methods. For SMICA, we also checked that TTT temperature constraints on local f_{NL} show no evidence of a directional variation via a needlet-based analysis.
- SMICA and SEVEM perform better than NILC and Commander at recovering the original f_{NL} in foreground-cleaned simulations. At the same time, SMICA allows slightly better constraints on f_{NL} than SEVEM, due to a lower (by a small amount) noise level.
- EEE bispectra, and related f_{NL} measurements, have some problems, and do not pass all the tests. Different component separation methods show a low level of consistency (especially when comparing pixel-based cleaning methods to harmonic-based cleaning methods). This disagreement is

only partly alleviated by choosing a larger E -mask, so that residual foregrounds do not seem to fully explain all issues. An important caveat, already pointed out previously, is that the noise model in polarized FFP8 simulations is known to underestimate the actual noise level in the data, leading to some degree of underestimation of the error bars in our EEE results. We stress again, however, that this has little impact on the final $T+E$ constraints, due to the high noise level and consequent low statistical weight of EEE bispectra. This was verified in detail both on data and on simulations.

In light of the above analysis, we conclude that, as far as bispectrum estimation is concerned, our best cleaned map is the one produced by SMICA, in line with our previous 2013 analysis. We also conclude that our TTT -based f_{NL} constraints, summarized in Tables 10 and 11, are robust. Joint $T+E$ constraints pass all our validation tests. On the other hand, in the light of the remaining issues in the EEE bispectra and in the FFP8 polarized simulations, as we stressed at the end of Sect. 6.1, we suggest that all measurements that include polarization data in this paper should be regarded as *preliminary*.

8. Other non-Gaussianity shapes for f_{NL}

This section discusses new searches for NG beyond standard single-field inflation. The focus here is on extensions to the analysis undertaken in Planck Collaboration XXIV (2014) with new limits on isocurvature models, further oscillatory models over a broader frequency range, and parity-violating tensor NG. However, we also briefly revisit all the non-standard models constrained in Planck Collaboration XXIV (2014), including effective field theory, non-Bunch Davies, and directionally-dependent models, in particular noting the impact of the new (preliminary) polarization data on the previous constraints.

8.1. Isocurvature non-Gaussianity

Here we show the results obtained for a study of the isocurvature NG in the *Planck* 2015 SMICA map using the binned bispectrum estimator. As explained in Sect. 2.3, we only investigate isocurvature NG of the local type, and in addition always consider only one isocurvature mode (either cold dark matter, neutrino density, or neutrino velocity isocurvature) in addition to the adiabatic mode. Note that the baryon isocurvature mode behaves identically to the cold dark matter one, only rescaled by factors of Ω_b/Ω_c , so there is no need to consider it separately. In that case there are six different f_{NL} parameters: a purely adiabatic one (a,aa), which corresponds to the result from Sect. 6), a purely isocurvature one (i,ii); and four mixed ones (see Sect. 2.3 for an explanation of the notation).

The results are given in Table 19¹⁶ and show no clear signs of any isocurvature NG. There are a few values that deviate from zero by up to about 2.5σ , but such a small deviation, in particular given the large number of results, cannot be considered a detection. We do see that many constraints are tightened considerably when including polarization, by up to the predicted factor of about six for the cold dark matter a,ii, i,ai, and i,ii modes in

¹⁶ Compared to definitions in the literature based on ζ and S (see e.g., Langlois & van Tent 2012), here we adopt definitions based on $\Phi_{\text{adi}} = 3\zeta/5$ and $\Phi_{\text{iso}} = S/5$, in order to make the link with the standard adiabatic result more direct. Conversion factors to obtain results based on ζ and S are $6/5$, $2/5$, $2/15$, $18/5$, $6/5$, and $2/5$, for the six modes, respectively.




Fig. 19. Temperature only, local f_{NL} directional contributions from SMICA. As explained in the text, summing over all the pixel values would give the full sky f_{NL} needlet estimator result. The left panel displays the directional f_{NL} map. On the right, the blue points represent the f_{NL} contribution for each direction (i.e., for each pixel in the directional map), with Monte Carlo error bars. The red line is the average from simulations, which is consistent with zero.

Table 19. Results for local isocurvature NG, determined from the SMICA *Planck* 2015 map with the binned bispectrum estimator. In each case the adiabatic mode is considered together with one isocurvature mode (either cold dark matter, neutrino density, or neutrino velocity isocurvature). As explained in the text this gives six different f_{NL} parameters, indicated by the different combinations of the adiabatic (a) and isocurvature (i) modes. Results with two significant digits are shown for both an independent and a fully joint analysis, for T -only, E -only, and full $T+E$. In all cases the ISW-lensing bias has been subtracted.

Shape	f_{NL}					
	Independent			Joint		
	Cold dark matter	Neutrino density	Neutrino velocity	Cold dark matter	Neutrino density	Neutrino velocity
T a,aa	1.3 ± 5.4	1.3 ± 5.4	1.3 ± 5.4	21 ± 13	-27 ± 52	-32 ± 48
T a,ai	-2 ± 10	-4 ± 15	47 ± 29	-39 ± 26	140 ± 210	370 ± 350
T a,ii	59 ± 910	-130 ± 280	750 ± 360	17000 ± 8200	-4500 ± 4500	-1300 ± 3800
T i,aa	6 ± 50	3.0 ± 9.0	1.0 ± 4.7	96 ± 120	40 ± 99	-27 ± 51
T i,ai	3 ± 66	-5 ± 22	26 ± 21	-2100 ± 1000	220 ± 630	75 ± 170
T i,ii	76 ± 280	-100 ± 250	440 ± 230	4200 ± 2000	-750 ± 2400	-970 ± 1400
E a,aa	34 ± 34	34 ± 34	34 ± 34	66 ± 50	51 ± 120	-140 ± 150
E a,ai	-31 ± 200	70 ± 140	78 ± 93	-380 ± 310	-280 ± 640	1100 ± 620
E a,ii	-4200 ± 4000	-520 ± 2300	190 ± 940	-8800 ± 6100	-6400 ± 6200	-9400 ± 3900
E i,aa	-10 ± 87	42 ± 42	23 ± 27	27 ± 180	52 ± 170	54 ± 120
E i,ai	94 ± 250	83 ± 130	45 ± 62	910 ± 770	670 ± 850	-190 ± 420
E i,ii	690 ± 2200	390 ± 1400	260 ± 460	-6000 ± 5300	-4100 ± 5300	2200 ± 1600
$T+E$ a,aa	0.7 ± 4.9	0.7 ± 4.9	0.7 ± 4.9	5 ± 10	-35 ± 27	2 ± 24
$T+E$ a,ai	-2.6 ± 9.7	-5 ± 14	17 ± 22	-12 ± 20	74 ± 94	330 ± 130
$T+E$ a,ii	130 ± 450	-130 ± 240	130 ± 230	-1800 ± 1300	-3000 ± 1400	-3200 ± 1200
$T+E$ i,aa	30 ± 26	5.6 ± 7.7	-0.7 ± 4.1	53 ± 47	51 ± 45	-44 ± 24
$T+E$ i,ai	26 ± 38	2 ± 19	6 ± 15	140 ± 170	170 ± 210	20 ± 74
$T+E$ i,ii	38 ± 170	-26 ± 180	85 ± 130	-280 ± 390	-390 ± 860	480 ± 430

the joint analysis. As discussed in detail in Sect. 7, results including polarization data should be considered preliminary, and that is even more important here, since these results depend so strongly on the additional information from polarization.

In the results so far we allowed for a possible correlation between the isocurvature and adiabatic modes. However, if we assume that they are completely uncorrelated, with a zero cross power spectrum, then there are only two f_{NL} parameters, the a,aa and i,ii ones. In Table 20 we give the results for this uncorrelated

case. The independent results are the same as in the previous table, while in the joint results one can clearly see the difference between the neutrino density mode (the bispectrum template of which has a large overlap with the adiabatic one), and the cold dark matter and neutrino velocity modes (with templates that are very different from the adiabatic one). Again there is no evidence for any isocurvature NG; the almost 2σ result for the neutrino velocity isocurvature mode in the temperature-only case does not survive the addition of polarization.

Table 20. Similar to Table 19, except that we now assume that the adiabatic and isocurvature mode are completely uncorrelated. Hence there are only two f_{NL} parameters in this case, a purely adiabatic one and a purely isocurvature one.

Shape	f_{NL}					
	Independent			Joint		
	Cold dark matter	Neutrino density	Neutrino velocity	Cold dark matter	Neutrino density	Neutrino velocity
T a,aa	1.3 ± 5.4	1.3 ± 5.4	1.3 ± 5.4	1.0 ± 5.3	19 ± 12	-0.2 ± 5.4
T i,ii	76 ± 280	-100 ± 250	440 ± 230	65 ± 280	-840 ± 540	440 ± 230
E a,aa	34 ± 34	34 ± 34	34 ± 34	33 ± 35	42 ± 40	35 ± 40
E i,ii	690 ± 2200	390 ± 1400	260 ± 460	210 ± 2200	-680 ± 1700	-31 ± 540
$T+E$ a,aa	0.7 ± 4.9	0.7 ± 4.9	0.7 ± 4.9	0.5 ± 5.0	3.0 ± 7.9	-0.3 ± 4.9
$T+E$ i,ii	38 ± 170	-26 ± 180	85 ± 130	35 ± 170	-120 ± 290	87 ± 130

8.2. Feature models

An important and well-motivated class of scale-dependent bispectra is the feature model, characterized by linear oscillations described by Eq. (15) and its variants in Eqs. (16) and (18). In Planck Collaboration XXIV (2014) we performed an initial search for a variety of feature models using the Modal estimator. This earlier search was limited to $\omega < 200$ by the native resolution of our implementation of the Modal estimator (using 600 modes), roughly the same range as the initial WMAP bispectrum feature model searches at lower precision (with only 50 eigenmodes (Fergusson et al. 2012)). Note, in the previous *Planck* release we used wavenumber k_c in line with the theory literature, but here we switch to frequency ω , in line with more recent observational power spectrum searches; the two are related by $\omega = 2\pi/3k_c$. With the improved estimator resolution (now using 2001 modes) we are able to achieve convergence over a broader range up to $\omega = 350$. We perform a frequency scan of 350 sampling points between $\omega = 10$ and $\omega = 350$, i.e., 35 independent frequencies and 10 phases. We also extend the number and variety of feature and resonance models that are investigated, essentially probing the resolution domain in which we have obtained a reliable Modal bispectrum reconstruction (see Fig. 4).

Constant feature ansatz: For the constant feature shape of Eq. (15), we can extend the frequency range much further with another approach. As the bispectrum in Eq. (15) is separable, it allows the construction of a KSW estimator (Münchmeyer et al. 2014) for direct bispectrum estimation at any given frequency. The bispectrum can be written as a sum of sine and cosine components which can be estimated separately (equivalent to measuring the amplitude and phase above) and this method was used to constrain frequencies up to $\omega = 3000$. The range where the two estimators overlap provides validation of the two methods and excellent agreement was seen (see Fig. 20).

Apart from cross-validation with two estimators, we have undertaken further tests to determine the robustness of the results to foreground and noise effects. In Fig. 21 (left panel), we show the effect on feature model results of the subtraction of the simple point source bispectrum, as well as the ISW-lensing bispectrum. This study was a major motivation for adopting the more conservative “extended” common mask, because the consistency between different component separation methods improved markedly for low frequencies, with the original common mask requiring much larger point-source subtractions (e.g., for NILC subtraction the maximum raw significance has reduced from $\sigma = 4.0$ to $\sigma_{\text{clean}} = 2.2$ at $\omega = 110$). After cleaning these signals, the SMICA, SEVEM and NILC results are in good

agreement, and also consistent with each other when polarization is included (while the Commander results generally have a larger variance and so are not included in the plotted averages). Fortunately, the effect of subtracting ISW-lensing and point source bispectra diminishes rapidly at higher frequencies $\omega > 200$ and should be negligible; subtraction was only undertaken in the Modal region $\ell < 350$. In Fig. 21 (right panel), we show the effect on the averaged significance of reducing the *Planck* domain from the usual $\ell_{\text{max}} = 2000$ to $\ell_{\text{max}} = 1500$ ($\ell_{\text{max}} = 1500$ to $\ell_{\text{max}} = 1000$ for E -modes). Despite the non-trivial change in overall signal-to-noise entailed, there is no strong evidence for an ℓ -dependent signal, as might be expected if there was substantial NG feature contamination in the noise-dominated region. Finally, we note that most peaks at low ω show some correlation between T -only and $T+E$, although there are notable exceptions, such as the peak at $\omega = 180$, which is removed after inclusion of polarization (see also the phase plots in Fig. 20 before marginalization). The temperature feature peaks observed in Fig. 21 at $\omega \approx 110, 150$, and 180 are consistent with the peaks identified previously in Planck Collaboration XXIV (2014).

In Fig. 20, the full set of frequency results, $0 < \omega < 3000$, for the constant feature model in Eq. (15) is shown for both the Modal and KSW estimators. There is good agreement for regions where the estimators overlap and PS and ISW lensing bispectra subtraction is not necessary ($\omega > 200$). Generally, there is tighter consistency between temperature-only results, than with polarization where there is additional scatter between foreground separation methods. Scanning across the full frequency range, there is no strong evidence for any large maximum that might indicate unequivocal evidence for a feature model signal. The maximum peak significance obtained with either T -only or $T+E$ is consistent with expectations for a Gaussian model over this frequency range. In particular, for the KSW estimator using the SMICA data, the highest significance found in the range $200 < \omega < 3000$ is 3.2σ in T -only and 2.9σ in $T+E$. To gauge the likelihood of these results occurring randomly, realistic Gaussian SMICA simulations were analysed with the KSW estimator and found to typically produce a highest peak with $3.1(\pm 0.3)\sigma$ over the same frequency range.

Generalized feature models: We have also deployed the Modal estimator to look at (non-separable) feature models with equilateral and flattened cross-sections, as motivated by varying sound speed scenarios and those with highly excited states, respectively. In the left panels in Fig. 22 we show results from the equilateral feature model of Eq. (16), including the frequency/phase




Fig. 20. Constant feature model results for both T -only and $T+E$ data across a wide frequency range. The upper four panels show the feature signal in the Modal range $0 < \omega < 350$. The two upper left panels show contours of the raw significance σ obtained from the SMICA map as a function of the frequency ω , for T -only and $T+E$, respectively. The upper right panels show the maximum signal after marginalizing over phase ϕ for both T -only and $T+E$ for all foreground separation models. The third and fourth panels show the maximum feature signal in both T -only and $T+E$ across the frequency range $0 < \omega < 1000$, plotting both Modal results (dashed lines) and KSW results (solid lines for $200 < \omega < 1000$); these show good agreement in the overlap. The lower two panels give the maximum KSW results for T -only and $T+E$ in the range $1000 < \omega < 3000$.




Fig. 21. Constant feature ansatz validation for the Modal estimator, showing the effect of ISW-lensing and point source subtraction for $\ell < 300$ (left panel) and the impact of a lower ℓ maximum cutoff on the average signal (right panel), i.e., lowering $\ell_{\max} = 2000$ to 1500 (T) and $\ell_{\max} = 1500$ to 1000 (E). All Modal 2 results in Sect. 8 have used the extended common mask, except the validation analysis at different resolutions (right panel) which for consistency employs the common mask.




Fig. 22. Generalized feature models analysed at $\ell_{\max} = 2000$ (E -modes $\ell_{\max} = 1500$) for the different *Planck* foreground separation methods, SMICA (blue), SEVEM (red), NILC (green), and Commander (yellow), together with the SSN (SMICA - SEVEM - NILC) average (black). The left three panels apply to the equilateral feature models, showing, respectively, in the top panel the full feature survey significance at each frequency and phase (temperature only), the maximum significance at each frequency for temperature only (middle), and with polarization (lower). The right three panels apply to the flattened feature models, showing the maximum significance at each frequency for temperature only (top right) and with polarization (middle right), along with significance at each frequency and phase for temperature and polarization (right lower).

contours before marginalization for the SMICA T -only map. Multiple peaks are apparent in the temperature signal across the Modal range up to an average maximum 3.3σ raw significance. However, from the lower panel it is clear that the polarization signal is not well correlated with the temperature in the equilateral case, generally reducing peak heights with the maximum now about 2.6σ (while eliminating the $\omega = 180$ peak altogether). This temperature peak remains present with the $\ell_{\max} = 1500$ cutoff, where the signal is slightly higher, but the polarization in this case is less well correlated (using $\ell_{\max}^E = 1000$). For Gaussian noise we would not expect polarization to reinforce a high temperature signal on average. It may also be that the equilateral temperature signal has some resid-

ual diffuse point-source contamination. The equilateral feature model is the most affected by the removal of point sources, so the presence of a more complex correlated PS bispectrum (not removed by the constant PS template subtraction) remains for future investigation.

Results for the flattened feature model in Eq. (18) are shown in Fig. 22 (right panels), displaying more coherence between temperature and polarization. The temperature signal with a 2.6σ peak between $50 < \omega < 150$ is reinforced by polarization and merges to make a broad 3σ peak around $\omega = 90$, together with another distinct peak at $\omega \approx 140$. Such broad frequency peaks are not expected, because neighbouring feature models should be nearly uncorrelated over a range $\Delta\omega_{\text{eff}} \approx 13$




Fig. 23. Single field feature model significance with a $K^2 \cos \omega K$ scaling dependence (Eq. 19) (left panels, T -only upper and $T+E$ lower) or with a $K \sin \omega K$ scaling (Eq. 20) (right panels). To find the maximum signal, these results have been marginalized over the α -dependent envelope function ranging from $\alpha \rightarrow 0$ (no envelope) to the maximum cutoff allowed by the Modal resolution $\alpha \omega = 90$.

Table 21. Peak statistics for the different feature models showing the *Raw* peak maximum significance (for the given Modal survey domain), the corrected significance of this *Single* maximum peak after accounting for the parameter survey size (the “look-elsewhere” effect) and the *Multi*-peak statistic which integrates across the adjusted significance of all peaks to determine consistency with Gaussianity. SMICA, SEVEM and NILC map analyses exhibit satisfactory bispectrum agreement for all the different models, whereas the Commander results produce some anomalously large results, especially for polarization. The significant signal for the equilateral features model in the T -only multi-peak statistic is reduced when polarization is added. The flattened feature model produces interesting results, which are reinforced with polarization to the 3σ -level, with a high multi-peak significance.

	SMICA			SEVEM			NILC			Commander		
	Raw	Single	Multi	Raw	Single	Multi	Raw	Single	Multi	Raw	Single	Multi
Features constant T -only	2.7	0.5	1.7	2.6	0.4	1.5	2.8	0.7	2.2	2.9	0.8	2.7
Features constant $T+E$	2.7	0.5	1.9	2.8	0.7	2.5	2.8	0.7	2.4	2.6	0.4	1.5
Features equilateral T -only	3.3	1.5	4.0	3.2	1.3	3.5	3.3	1.6	4.1	2.9	0.9	2.5
Features equilateral $T+E$	2.6	0.4	1.3	2.6	0.4	1.6	2.8	0.7	1.9	2.7	0.6	1.5
Features flattened T -only	2.5	0.3	1.4	2.6	0.4	1.6	2.7	0.5	2.1	2.8	0.8	2.7
Features flattened $T+E$	2.9	0.9	2.9	3.0	1.1	3.5	3.1	1.2	3.8	3.1	1.2	3.8
$K^2 \cos$ features T -only	2.5	0.7	1.9	2.3	0.6	1.6	2.7	1.0	2.5	2.2	0.3	1.1
$K^2 \cos$ features $T+E$	2.7	1.0	2.5	2.7	1.1	2.6	2.6	1.0	2.5	2.4	0.6	1.8
$K \sin$ feature T -only	2.8	1.2	2.8	2.7	1.1	2.7	3.0	1.5	3.4	2.6	0.9	2.3
$K \sin$ features $T+E$	2.1	0.3	1.0	2.9	1.4	3.1	2.4	0.6	1.7	2.3	0.5	1.6

(as we discuss below). As the phase plots in Fig. (22) indicate, this breadth in frequency ω may reflect the neighbouring feature models adjusting phase ϕ to match an underlying NG signal of a related, but different, nature. We note also that the frequency region for $\omega < 100$ is susceptible to some degeneracy with cosmological parameters. We shall consider a “look-elsewhere” statistical analysis of these results below.

Single-field feature solutions: We have also searched for the specific analytic solutions predicted for single-field inflation models with step-like potential features, as given in Eqs. (19) and (20), with results shown in Fig. 23. The highest peaks for the K^2 -cosine model occur around 2.5σ with temperature-only, then rises to 2.7σ when polarization data are included, again with peaks at other distinct frequencies apparent. The K -sine model shows a similar apparent signal level, with a maximum T -only 2.7σ peak, dropping to 2.4σ with polarization. One further difficulty with a positive interpretation of these bispectrum results in this low frequency range is that stronger S/N counterparts in the power spectrum are predicted for these simple mod-

els (Adshead et al. 2011), whereas no significant correlated oscillation signals are apparent at the relevant peak frequencies Planck Collaboration XX (2016).

Feature model peak statistics: In order to determine consistency with Gaussianity for these feature model results, we can apply a number of statistical tests developed for this specific purpose (Fergusson et al. 2015) and, if warranted, also apply these jointly in combination with power spectrum results, as for the WMAP polyspectra analysis (Fergusson et al. 2014). When scanning across the (ω, ϕ) parameter-dependent feature models, we are searching through independent models for which Gaussian noise, by chance, can lead to a large apparent signal. We must correct for this multiplicity of tests when determining the actual significance of results for a given model — this is a quantitative correction for any model with free parameters, distinct from the a posteriori choice of models to test. The simplest approach is to determine whether the maximum peak is consistent with Gaussian expectations, which can be determined from Monte Carlo simulations. However, in Fergusson et al. (2015)

it was recognized that these feature models can be accurately characterized analytically with a χ -distribution having two degrees of freedom¹⁷. Taking $\ell_{\max} = 2000$ and using this analysis, the frequency step size between models that are uncorrelated is approximately $\Delta\omega_{\text{eff}} = 13$, so we have an effective number of independent feature models $N_{\text{eff}} \approx 27$ for the Modal frequency range (with $N_{\text{eff}} \approx 230$ across the larger constant feature survey range). Accordingly appropriate look-elsewhere corrections have been applied to find an adjusted significance for the maximum peak signal found in all the feature model searches undertaken, which is shown in Table 21. Given that this feature model survey is over many independent frequency models and combinations of data, even the highest raw significances above 3σ (“Raw” column in Table 21) are reduced to a corrected significance below 2σ (“Single” column). Hence, there appears to be no evidence from maximum peak statistics for feature model deviations from Gaussianity.

Nevertheless, we also examine the possibility that multiple feature models are contributing to a NG signal, given the apparent emergence of several preferred frequency peaks. This integrated multi-peak statistic can also be accurately approximated analytically (Eq. 60) using a χ -distribution; essentially we sum over all independent frequencies using the single peak significance adjusted for the “look-elsewhere” effect (see Eq. (59)). Most of the signal surveys exhibit an unusual number of broad overlapping peaks within the accessible frequency domain, so the multi-peak statistic does yield a much higher significance, with many models above 2σ after “look-elsewhere” correction. Notable cases are the temperature-only signal for the equilateral feature model which yields an average significance of 3.4σ across the foreground-cleaned maps with concordant bispectrum results (i.e., SMICA, SEVEM and NILC); however, this interesting multi-peak significance drops to only 1.6σ when the polarization data are included (assuming the reliability of E results). On the other hand, the flattened feature model has an average multi-peak significance of 1.7σ in temperature only, which rises to 3.4σ with polarization included (higher at 3.7σ if Commander data were to be included in the average). In this case, beyond the number of peaks, it is also their width that contributes, with the main signal around $\omega \approx 90$, much broader than $\Delta\omega \approx 13$. Finally, after look-elsewhere effects are taken into account, the K^2 -cosine single-field solutions yield multi-peak statistics that rise with the inclusion of polarization data from 2.0σ to 2.5σ , while the K -sine falls from 3.0σ (T) to 1.9σ ($T+E$).

¹⁷ For the feature model with parameters (ω, ϕ) , the adjusted significance S for the raw significance σ after accounting for the “look-elsewhere” effect is given by (Fergusson et al. 2015)

$$S = \sqrt{2} \text{Erf}^{-1}[(F_{\chi,2}(\sigma))^{N_{\text{eff}}}], \quad (59)$$

where $F_{\chi,2}$ is the cumulative distribution function of the χ -distribution of degree two and N_{eff} is the effective number of independent feature models. This can also be used to investigate whether feature models are contributing at several frequencies. This multi-peak statistic integrates over all peak signals using the corrected significance S , i.e.,

$$S_I^2 = \frac{\Delta\omega}{\Delta\omega_{\text{eff}}} \sum_{\omega} 2 \text{Erf}^{-1}[F_{\chi,2}(\sigma(\omega))^{N_{\text{eff}}}]^2, \quad (60)$$

where $\Delta\omega$ is the sampling step-size and $\Delta\omega_{\text{eff}}$ is the effective correlation between independent models given by $\Delta\omega_{\text{eff}} = (\omega_{\max} - \omega_{\min})/(N_{\text{eff}} - 1)$. Essentially this sums up all significant peaks that have a non-zero adjusted significance, after accounting for the look-elsewhere effect in Eq. (59).

An interesting, but not entirely coherent, picture emerges from these searches for non-standard models in the new *Planck* temperature data, especially when combined with the additional (preliminary) polarization information. In Planck Collaboration XXIV (2014), we noted that the feature model searches provided interesting hints of NG. This more rigorous statistical analysis confirms this view, allowing for several alternative feature model explanations of the apparently high NG signal observed in the bispectrum reconstructions (see Fig. 4). However, there is no strong evidence for a single large feature model at a particular frequency; rather, the high multi-peak statistics indicate signal that is spread across several broad peaks. Given the variability between different feature models and polarization component-separation methods, we note the caveat that the integrated multi-peak statistic could be sensitive to calibration issues and foreground contamination. For this reason, we do not make strong claims for these non-standard signals at this stage, but we note that oscillatory models will continue to be investigated thoroughly over a wider frequency domain and using the more reliable polarization data available in the final *Planck* data release.

8.3. Resonance/axion monodromy models

Generalized resonance models: Using the Modal expansion, we have embarked on a survey of the simplest resonance model (Eq. 10), as well as the equilateral and flattened variants proposed in the literature, i.e., described by Eqs. (13) and (14), respectively. The raw significance for the resonance models for both temperature-only and temperature plus polarization data are shown in Fig. 24; these are the maximal results marginalized over the phase parameter ϕ . Previously, the resonance model was studied in Planck Collaboration XXIV (2014) using the Modal expansion over a narrower frequency range yielding no results above a raw significance of 1σ . In this extended analysis over a wider frequency range, the constant $\sin(\log)$ model (Eq. 10) produces 2.2σ peaks for T -only, and 2.6σ for $T+E$. The equilateral resonance model (Eq. 13) achieves a maximum 2.8σ in T -only at $\omega \approx 35$, rising to a more impressive average 3.2σ for $T+E$. For the flattened case (Eq. 14) we have 2.5σ and 3.0σ , respectively at $\omega \approx 12$. Qualitatively, the results shown in Fig. 24, exhibiting broad peaks, are similar to those for feature models.

Resonance model peak statistics: To determine the statistical significance of these results given the look-elsewhere effect of scanning across the parameters (ω, ϕ) , we have used the two peak statistics defined above in Eqs. (59) and (60) for feature models (Fergusson et al. 2015). In this case, the maximum peak statistic for the constant resonance model of 2.6σ ($T+E$) is readjusted to an unremarkable “look-elsewhere” single peak significance of 0.9σ . Likewise the apparently significant results above 3σ for the equilateral and flattened models now fall below 2.0σ with $T+E$. Using the single peak statistic alone, we would conclude that there is no strong evidence for any individual resonance model. Resonance models also generate oscillations in the power spectrum, and an analysis based on the 2015 temperature and polarization likelihood is presented in (Planck Collaboration XX 2016). A combined statistical treatment of resonance model power spectrum and bispectrum results will be reported in the future.

The multi-peak statistic in Eq. (60) integrates the resonance model signal across all frequencies to determine consistency with Gaussianity. The constant resonance model has a modest multi-peak signal but, like the feature models, the equilateral and




Fig. 24. Generalized resonance models analysed at $\ell_{\max} = 2000$ (E -modes $\ell_{\max} = 1500$) for the different *Planck* foreground separation methods, SMICA (blue), SEVEM (red), NILC (green), and Commander (yellow), together with the SSN (SMICA - SEVEM - NILC) average (black). The upper panels apply to the constant resonance model (Eq. 10), with T -only (left) and $T+E$ (right), the middle panels give results for the equilateral resonance model (Eq. 13), and the lower panels for the flattened resonance model (Eq. 14). Both the equilateral and flattened resonance models produce broad peaks, which are reinforced with polarization (middle and bottom right panels).

Table 22. Peak statistics for the resonance models showing the maximum *Raw* peak significance, the *Single* peak significance after accounting for the parameter survey “look-elsewhere” effect, and the *Multi*-peak statistic integrating across all peaks (also accounting for the “look-elsewhere” correction). There is some evidence for a high signal for both the equilateral and flattened resonance models, which increases when the polarization signal is added. This table does not include the results of the high frequency resonance model estimator, whose significance was assessed independently.

	SMICA			SEVEM			NILC			Commander		
	Raw	Single	Multi	Raw	Single	Multi	Raw	Single	Multi	Max.	Single	Multi
Sin(log) constant T -only	2.4	0.7	1.2	2.2	0.4	0.9	2.0	0.2	0.7	2.5	0.8	1.6
Sin(log) constant $T+E$	2.4	0.7	1.7	2.7	1.1	2.4	2.6	1.0	2.2	2.7	1.1	2.5
Sin(log) equilateral T -only	3.0	1.6	2.4	2.8	1.2	2.0	2.5	0.9	1.5	2.6	1.0	2.1
Sin(log) equilateral $T+E$	3.5	2.2	3.5	3.0	1.5	2.8	3.1	1.7	3.2	2.8	1.2	2.0
Sin(log) flattened T -only	2.5	0.7	1.8	2.5	0.8	1.9	2.6	0.9	2.1	3.0	1.6	3.2
Sin(log) flattened $T+E$	2.9	1.4	2.9	3.0	1.6	3.4	3.1	1.6	3.4	3.6	2.3	4.5

flattened resonance shapes offer stronger hints. The multi-peak equilateral signal rose from 1.9σ (T -only) to 3.1σ ($T+E$) after adjusting for the “look-elsewhere” effect, while the flattened signal went from 2.4σ (T -only) to 3.2σ ($T+E$). These interesting results, reflecting those obtained for feature models, suggests the fit to any underlying NG signal might await alternative, but related, oscillatory models for a more compelling explanation.

High frequency resonance model estimator: We have further surveyed the simple resonance model (Eq. 10) with a second approach. Using a model-specific expansion in terms of linear oscillation, proposed in Münchmeyer et al. (2015), it is possible to extend the frequency range of the analysis considerably. In this approach one exploits the fact that any bispectrum shape which

is a function of $k_1 + k_2 + k_3$ can be expanded in Fourier modes of $k_1 + k_2 + k_3$, resulting in an effectively one-dimensional expansion, as opposed to the general Modal expansion. In the present implementation we use 800 sine and cosine modes, which cover the full frequency space of the power spectrum search, i.e., $0 < \omega < 1100$. The significances found with this method are presented in Fig. 25. As was the case for the feature model, we find that SMICA, SEVEM and NILC results are in good agreement in both temperature and polarization, while Commander results generally have larger variance and so are not included in the plotted averages. We find that the largest peak of the average is 3.6σ in temperature and 3.1σ in temperature and polarization combined. The results of the Modal expansion discussed in the pre-

vious paragraph are in good agreement with the high frequency resonance model estimator in the overlapping frequency range.

Due to the high computational demands of this analysis, we have only exactly assessed the look-elsewhere effect for SMICA T data, for which we find an expected maximum peak of $3.5\sigma \pm 0.4\sigma$ in the case of Gaussian maps, to be compared to 3.7σ in the SMICA T data, demonstrating that the results are fully consistent with Gaussianity. The expected maximum peak for Gaussian maps was calculated from the Fisher matrix with the method described in Meerburg et al. (2015). The average over component separation methods as well as the $T+E$ data is even less significant. For the high frequency estimator we have assessed the significance of multiple peaks in the following way. Define the multi-peak amplitude A_M as the sum of squares of the M highest significances σ_i in the frequency range, i.e., $A_M = \left(\sum_{i=1}^M \sigma_i^2\right)^{1/2}$, where only approximately independent peaks with $\Delta\omega > 10$ are considered. One can then compare A_M to its distribution in the Gaussian case and get an individual significance σ_M for each number of peaks M , where we assume $M \leq 10$. The multi-peak statistic is then obtained by maximizing over M , leading to an additional look-elsewhere effect that we also accounted for. In this way we find that in the SMICA T data the raw peak maximum 3.7σ , is reduced to 0.5σ for the single-peak statistic and to 0.6σ in the multi-peak case. This large reduction in significance is due to the large number of independent frequencies as well as to the maximization over phases. One may argue that the frequency range $\omega < 1100$ is too large, as EFT arguments for resonance non-Gaussianities (Behbahani et al. 2012) limit the frequency range to $\mathcal{O}(10^2)$. As an example, we have therefore also calculated the look-elsewhere-corrected significances when we limit the analysis to $\omega < 250$. In this case we find a single-peak significance of 0.6σ and a multi-peak significance of 0.9σ . Clearly the results are fully consistent with Gaussianity.

8.4. Equilateral-type models and the effective field theory of inflation

There is considerable interest in equilateral-type models because they are physically well-motivated, through e.g., varying sound speed scenarios. There are generic predictions available from the effective field theory of inflation, notably the two specific effective field theory (EFT1 and EFT2) shapes that give rise to the equilateral and orthogonal approximations. These models were previously constrained in Planck Collaboration XXIV (2014) and the reader is referred to section 2 of that paper for analytic expressions for the specific shapes constrained here. In Table 23, we list the main equilateral-type models in the literature, giving constraints for T -only and $T+E$. All these models correlate well with the equilateral ansatz (Eq. 12) and likewise do not show a significant signal. However, despite this correlation, it is interesting to note the variation between models, largely due to the difference between these shapes in the flattened limit. The implications of these results are discussed in Sect. 11.

8.5. Models with excited initial states (non-Bunch-Davies vacua)

Non-Bunch-Davies (NBD) or excited initial states are models which produce flattened (or squeezed) bispectrum shapes. The wide variety of NBD models that have been proposed are briefly classified and labelled in Sect. 2, following a more exten-

sive overview in section 2 of Planck Collaboration XXIV (2014) where more analytic forms and the first constraints were presented. The latest *Planck* constraints for these models are listed in Table 24, obtained using the Modal 2 estimator with polarization. Despite the apparent “flattened” signal seen in the *Planck* bispectrum reconstructions (Fig. 4), this is generally not matched well by the specific modulation induced by the acoustic peaks for these scale-invariant NBD models. Tight constraints emerge for most models. The largest signal obtained is from the NBD sinusoidal shape which gives a 1.6σ T -only raw significance, rising to 2.1σ for $T+E$; this is hardly an impressive correspondence given the number of models surveyed and the parameter freedom used in maximizing the signal. However, an important caveat for NBD models is that the predicted shapes can be very narrow in the flattened limit, in which case solutions have been smoothed to match the current Modal resolution (though this has improved considerably since the *Planck* 2013 NG analysis). An improved match to the warm inflation shape means that the final constraint shown in Table 24 is more robust, with further implications discussed in Sect. 11.

8.6. Direction-dependent primordial non-Gaussianity

We impose observational limits on direction-dependent primordial NG parametrized by Eq. 23. Rather than using c_1 and c_2 we instead choose to work with the non-linearity parameters $f_{\text{NL}}^{L=1} = -c_1/4$ and $f_{\text{NL}}^{L=2} = -c_2/16$ (chosen to match a primordial bispectrum that is equal to the equilateral shape in the equilateral limit) keeping the notation from the 2013 results. We estimated the f_{NL}^L values from temperature data and high-pass filtered polarization data from the four foreground-cleaned CMB maps SMICA, NILC, SEVEM, and Commander, where we apply the common mask. The details of the KSW estimator and its derivation is presented in Appendix A. For temperature data, we use the common mask as adopted in Planck Collaboration XII (2014), which has more conservative foreground masking than the newly available mask. We choose the more conservative foreground masking, considering the fact that anisotropic NG is more sensitive to residual foregrounds. We set the maximum multipole to 2000 and 1000 for temperature and polarization data, respectively. Validating our analysis pipeline with realistic simulations, we find that the asymmetry of the *Planck* beam, coupled with the *Planck* scanning pattern, inflates the statistical fluctuations of the f_{NL}^L significantly. Noting the large angular scale of artificial anisotropy produced by the beam asymmetry, we set the minimum multipole to 101, and find that the statistical fluctuation of estimation from simulated data is close to the theoretical expectations.

These two shapes are also constrained using the Modal 2 estimator, which is not affected by the beam asymmetry and is used in the same form as elsewhere in the paper with multipoles from 2 to 2000 and 30 to 1500 being used for temperature and polarization, respectively. The present constraints are consistent with those found for T -only in Planck Collaboration XXIV (2014), but at higher resolution convergence has improved considerably, reflected in the lower variance.

We find that the ISW-lensing bispectrum and the unresolved point-sources bispectrum bias the estimation of the f_{NL}^L , in particular in the analysis of temperature data. For our final values, we subtract both these biases from our estimation. In Table 25, we report the estimated value of f_{NL}^L from the foreground-cleaned CMB maps. For $L = 1$ the effect of the differing ℓ -ranges between the two estimators is not so significant and the results are quite consistent. For $L = 2$, which has significant signal in




Fig. 25. Standard resonance model results for both T -only and $T+E$ across a wide frequency range using the high frequency resonance model estimator. The first and second panels show the signal in both T -only (first panel) and $T+E$ (second panel) across the frequency range $0 < \omega < 500$. The lower two panels give the results for T -only (third panel) and $T+E$ (fourth panel) in the range $500 < \omega < 1100$.

the squeezed configuration, the effect of removing small scales from the KSW estimator is more pronounced, resulting in significantly enlarged error bars. In light of this, the differences seen between the central values for the two methods is to be expected and does not indicate any inconsistencies between the two approaches. The slight differences between the results from different foreground-cleaned temperature maps are within the likely

range of statistical fluctuations, estimated from realistic simulations of CMB and noise propagated through the pipelines of foreground-cleaned map making. As seen in Table 25, we find that the estimated values of f_{NL}^L from *Planck* 2015 temperature plus polarization data are consistent with zero.

Table 23. Constraints on models with equilateral-type NG covering the shapes predicted by the effective field theory of inflation, together with constraints on specific non-canonical inflation models, such as DBI inflation. See [Planck Collaboration XXIV \(2014\)](#) (section 2) for further explanation of these specific models, with further implications discussed in Sect. 11.

Equilateral-type model	SMICA		SEVEM		NILC		Commander	
	$A \pm \sigma_A$	S/N						
Constant T -only	12 ± 38	0.3	16 ± 38	0.4	10 ± 37	0.3	1 ± 39	0.0
Constant $T+E$	18 ± 22	0.8	28 ± 24	1.2	12 ± 23	0.5	26 ± 24	1.1
Equilateral T -only	-15 ± 68	-0.2	-9 ± 68	-0.1	-19 ± 67	-0.3	-17 ± 69	-0.3
Equilateral $T+E$	5 ± 42	0.1	4 ± 45	0.1	-2 ± 42	-0.1	27 ± 45	0.6
EFT shape 1 T -only	-3 ± 65	0.0	3 ± 64	0.0	-7 ± 62	-0.1	-9 ± 66	-0.1
EFT shape 1 $T+E$	12 ± 39	0.3	15 ± 42	0.3	3 ± 39	0.1	32 ± 42	0.8
EFT shape 2 T -only	17 ± 50	0.3	22 ± 50	0.5	15 ± 47	0.3	8 ± 51	0.2
EFT shape 2 $T+E$	23 ± 29	0.8	31 ± 31	1.0	15 ± 29	0.5	36 ± 31	1.2
DBI inflation T -only	3 ± 62	0.0	9 ± 61	0.1	-1 ± 59	0.0	-4 ± 63	-0.1
DBI inflation $T+E$	15 ± 37	0.4	20 ± 39	0.5	7 ± 37	0.2	34 ± 40	0.9
Ghost inflation T -only	-50 ± 80	-0.6	-45 ± 80	-0.6	-54 ± 79	-0.7	-45 ± 82	-0.6
Ghost inflation $T+E$	-27 ± 50	-0.5	-37 ± 54	-0.7	-31 ± 51	-0.6	1 ± 55	0.0
Inverse decay T -only	17 ± 43	0.4	21 ± 43	0.5	14 ± 41	0.3	4 ± 44	0.1
Inverse decay $T+E$	23 ± 25	0.9	32 ± 27	1.2	15 ± 26	0.6	32 ± 27	1.2

Table 24. Constraints on models with excited initial states (non-Bunch-Davies models), as well as warm inflation. See Sect. 2 for further explanation and the labelling of these classes of NBD models. Note that the NBD, NBD1, and NBD2 models contain free parameters, so here we quote the maximum significance found over the available parameter range; the maximum for T and $T+E$ can occur at different parameter values (on which the error bars are also dependent).

Flattened-type model	SMICA		SEVEM		NILC		Commander	
	$A \pm \sigma_A$	S/N						
Flat model T -only	49 ± 65	0.8	57 ± 65	0.9	47 ± 65	0.7	19 ± 65	0.3
Flat model $T+E$	44 ± 37	1.2	70 ± 37	1.9	33 ± 37	0.9	47 ± 37	1.3
Non-Bunch-Davies T -only	42 ± 82	0.5	53 ± 82	0.6	26 ± 82	0.3	17 ± 82	0.2
Non-Bunch-Davies $T+E$	61 ± 47	1.3	76 ± 47	1.6	43 ± 47	0.9	58 ± 47	1.2
NBD sine T -only	567 ± 341	1.7	513 ± 341	1.5	588 ± 341	1.7	604 ± 341	1.8
NBD sine $T+E$	-387 ± 206	-1.9	-485 ± 218	-2.2	-425 ± 206	-2.1	-417 ± 210	-2.0
NBD1 cos flattened T -only	-10 ± 22	-0.5	-4 ± 22	-0.2	-8 ± 22	-0.4	-9 ± 22	-0.4
NBD1 cos flattened $T+E$	-20 ± 19	-1.1	-10 ± 19	-0.5	-19 ± 19	-1.0	-14 ± 19	-0.8
NBD2 cos squeezed T -only	10 ± 17	0.6	10 ± 17	0.6	8 ± 17	0.5	-2.5 ± 17	-0.1
NBD2 cos squeezed $T+E$	-3 ± 5	-0.5	-0.8 ± 5.5	-0.1	-4 ± 5	-0.8	-3.8 ± 5.5	-0.7
NBD1 sin flattened T -only	-25 ± 22	-1.1	-27 ± 22	-1.2	-18 ± 22	-0.8	-33 ± 23	-1.4
NBD1 sin flattened $T+E$	48 ± 30	1.6	49 ± 33	1.5	35 ± 31	1.1	26 ± 34	0.8
NBD2 sin squeezed T -only	-2.0 ± 1.4	-1.4	-1.4 ± 1.4	-1.0	-1.6 ± 1.4	-1.1	-1.3 ± 1.4	-0.9
NBD2 sin squeezed $T+E$	-0.8 ± 0.4	-1.9	-0.5 ± 0.4	-1.2	-0.6 ± 0.4	-1.4	-0.5 ± 0.4	-1.2
NBD3 non-canonical T -only ($\times 1000$)	-5.9 ± 6.7	-0.9	-6.0 ± 6.8	-0.9	-5.4 ± 6.8	-0.8	-5.5 ± 6.7	-0.8
NBD3 non-canonical $T+E$ ($\times 1000$)	-8.7 ± 5.0	-1.7	-6.2 ± 5.2	-1.2	-7.5 ± 5.2	-1.5	-9.4 ± 5.2	-1.8
WarmS inflation T -only	-23 ± 36	-0.6	-26 ± 36	-0.7	-32 ± 36	-0.9	-24 ± 36	-0.7
WarmS inflation $T+E$	-14 ± 23	-0.6	-28 ± 23	-1.2	-21 ± 23	-0.9	-17 ± 23	-0.7

8.7. Parity-violating tensor non-Gaussianity

We present observational limits on the parity-violating tensor nonlinearity parameter $f_{\text{NL}}^{\text{tens}}$ from the temperature and E -mode polarization data. Unlike the usual scalar-mode templates, the CMB bispectra sourced from the tensor NG (Eq. 27) are written in non-factorizable forms ([Shiraishi et al. 2013b](#)); hence, we use the separable Modal pipeline in our bispectrum estimations.

The parity-violating NG under examination induces non-vanishing signals not only in parity-even configurations ($\ell_1 + \ell_2 + \ell_3 = \text{even}$) but also in the parity-odd ones ($\ell_1 + \ell_2 + \ell_3 = \text{odd}$) in the temperature and E -mode polarization bispectra ([Shiraishi et al. 2013b](#)). The optimal estimator, including all (even + odd) bispectrum signals, is expressed by the linear combination of the parity-even and parity-odd estimators, reading

([Liguori et al. 2015](#))

$$\hat{f}_{\text{NL}}^{\text{all}} = \frac{N^{\text{even}} \hat{f}_{\text{NL}}^{\text{even}} + N^{\text{odd}} \hat{f}_{\text{NL}}^{\text{odd}}}{N^{\text{even}} + N^{\text{odd}}}, \quad (61)$$

where $N^{\text{even/odd}}$ is the normalization factor (related to the Fisher matrix as $N^{\text{even/odd}} = 6F^{\text{even/odd}}$) defined for $\ell_1 + \ell_2 + \ell_3 = \text{even/odd}$. The parity-even estimator $\hat{f}_{\text{NL}}^{\text{even}}$ can be computed using the original Modal methodology ([Fergusson et al. 2010a, 2012; Fergusson 2014; Liguori et al. 2015](#)), while in computations of the parity-odd estimator $\hat{f}_{\text{NL}}^{\text{odd}}$, we follow the extended spin-weighted pipeline ([Shiraishi et al. 2014, 2015; Liguori et al. 2015](#)).

Our $f_{\text{NL}}^{\text{tens}}$ estimations (with both temperature and polarization data) are based on the resolution of $\ell_{\text{max}} = 500$ and HEALPix $N_{\text{side}} = 512$, leading to feasible computational costs. This choice is not expected to change the results significantly, in compari-

Table 25. Direction-dependent NG results for both the $L = 1$ and $L = 2$ models. We present results from both the KSW and Modal 2 pipelines. The discrepancy between the central values for the $L = 2$ models is due to the differing ℓ ranges taken for the two estimators, the key difference being the KSW $\ell_{\min} = 101$. As this model peaks in the squeezed configuration, a significant portion of the signal is lost, which is reflected in the increased error bars.

	Commander		NILC		SEVEM		SMICA	
	$A \pm \sigma_A$	S/N						
$L = 1$								
Modal 2 T -only	-41 ± 43	-0.9	-58 ± 42	-1.4	-51 ± 43	-1.2	-49 ± 43	-1.1
KSW T -only	-8 ± 46	-0.2	-62 ± 46	-1.3	-34 ± 45	-0.8	-26 ± 45	-0.6
Modal 2 $T+E$	-28 ± 29	-1.0	-30 ± 27	-1.1	-49 ± 28	-1.7	-31 ± 26	-1.2
KSW $T+E$	-57 ± 33	-1.7	-62 ± 32	-1.9	-79 ± 32	-2.5	-54 ± 32	-1.7
$L = 2$								
Modal 2 T -only	0.7 ± 2.8	0.2	0.8 ± 2.8	0.4	1.1 ± 2.7	0.3	0.5 ± 2.7	0.2
KSW T -only	1.5 ± 5.1	0.3	-3.9 ± 5.1	-0.8	-0.4 ± 5.1	-0.1	0.1 ± 5.0	0.0
Modal 2 $T+E$	1.1 ± 2.4	0.5	0.5 ± 2.4	0.2	1.3 ± 2.4	0.6	-0.2 ± 2.3	-0.1
KSW $T+E$	-3.0 ± 4.1	-0.7	-3.6 ± 4.0	-0.9	-3.8 ± 4.0	-1.0	-1.3 ± 3.9	-0.3

son to the analysis at higher resolution, e.g., $\ell_{\max} = 2000$ and Healpix $N_{\text{side}} = 2048$, since the cosmic variance and instrumental noise are already far higher than the signals for $\ell \gtrsim 300$ (Shiraishi et al. 2013b). Only in the polarization data analysis is an effective ℓ_{\min} also adopted, which is motivated by the high-pass filtering process for $\ell \leq 40$ in component separation.

Within the above ℓ ranges, the theoretical bispectrum templates are decomposed with the eigenbasis composed of $\mathcal{O}(1 - 10)$ polynomials and some special functions reconstructing the tensor-mode features, e.g., temperature enhancement due to the ISW effect ($\ell \lesssim 100$), and two E -mode peaks created by reionization ($\ell \lesssim 10$) and recombination ($\ell \approx 100$). The resulting factorized templates are more than 95% correlated with the original ones. The validity of our numerical computations has been confirmed through the map-by-map comparisons of $\hat{f}_{\text{NL}}^{\text{even/odd}}$ at very low resolution, showing the consistency between the values from the Modal methodology and those obtained by the brute-force $\mathcal{O}(\ell^5)$ summations like Eq. (36). We have also checked that our parity-even estimator successfully leads to the constraints on $f_{\text{NL}}^{\text{local}}$, $f_{\text{NL}}^{\text{equil}}$, and $f_{\text{NL}}^{\text{ortho}}$ at $\ell_{\max} = 500$, compatible with the results from the binned estimator.

Our limits estimated from the foreground-cleaned temperature and high-pass filtered polarization data (SMICA, SEVEM, and NILC) are summarized in Table 26. The data and MC simulations used here, including all experimental features, i.e., beam, anisotropic noise levels and partial sky mask, have been inpainted using the identical recursive process adopted in the standard f_{NL} estimations (see Sect. 3.5). The sky fractions of the temperature and polarization masks adopted here are, respectively, $f_{\text{sky}} = 0.76$ and $f_{\text{sky}} = 0.74$. Although the error bars and the linear terms have been computed using 160 MC simulations, the resulting error bars are close to the expected values, $(f_{\text{sky}}F)^{-1/2}$.

We have confirmed the stability of the T -only constraints, and significant scatter of the E -only constraints both in the parity-even case and in the parity-odd one, when changing f_{sky} . Such E -mode instability has given insignificant effects on our $T+E$ constraints in the parity-even case, as they are determined almost exclusively by TTT , like the scalar NG analyses. In contrast, our parity-odd $T+E$ results vary a lot, due to the E -mode scatter (quantitatively speaking, only a few percent change of f_{sky} has shifted $f_{\text{NL}}^{\text{tens}}$ by more than 1σ), because TTE and TEE contribute significantly to the signal-to-noise ratio in the parity-odd case (Shiraishi et al. 2013b). Table 27 presents the correlations of the bispectra reconstructed from the component sep-

arated maps, also indicating the robustness of the T -only constraints and the instability of the E -only results. We report only stable results in Table 26 and conclude that there is no evidence at $> 2\sigma$ of $f_{\text{NL}}^{\text{tens}}$ in the parity-even, parity-odd or their whole domain.

The parity-odd components of the TTT and EEE bispectra extracted model-independently from the SMICA data are visually represented in Fig. 26. It is apparent from this figure that the *Planck* TTT bispectrum has similar features to the WMAP one (Shiraishi et al. 2015), e.g., distinctive signals distributed around $\ell_1 \approx \ell_2 \approx \ell_3$. As indicated by the roughly 70% correlation between the SMICA and WMAP bispectra (see Table 27), the *Planck* T -only limits in Table 26 are close to the WMAP ones (68% CL): $f_{\text{NL}}^{\text{tens}}/10^2 = 4 \pm 16$ for parity-even; and $f_{\text{NL}}^{\text{tens}}/10^2 = 80 \pm 110$ for parity-odd (Shiraishi et al. 2015).

Table 26. Results for the tensor nonlinearity parameter $f_{\text{NL}}^{\text{tens}}/10^2$, estimated from the SMICA, SEVEM, and NILC temperature and high-pass filtered polarization maps. We separately show the central values and the errors (68% CL) extracted from $\ell_1 + \ell_2 + \ell_3 = \text{even}$ (Even), $\ell_1 + \ell_2 + \ell_3 = \text{odd}$ (Odd) and their whole domain (All). The parity-odd constraints have also been obtained from the E -mode data, but they are still preliminary and not currently shown.

	Even	Odd	All
SMICA			
T	2 ± 15	120 ± 110	4 ± 15
$T+E$	0 ± 13		
SEVEM			
T	2 ± 15	120 ± 110	5 ± 15
$T+E$	4 ± 13		
NILC			
T	3 ± 15	110 ± 100	5 ± 15
$T+E$	1 ± 13		

9. Limits on the primordial trispectrum

So far, we have considered a variety of physically motivated possibilities for the inflationary 3-point function, or bispectrum. A similar phenomenology exists for the 4-point function, or

trispectrum. Our constraints on the trispectrum will use CMB temperature only; we do not use polarization in this section. We start by briefly reviewing the inflationary physics and classifying the signals we will search for.

First, some notation: a “primed” ζ -trispectrum $\langle \zeta_{k_1} \zeta_{k_2} \zeta_{k_3} \zeta_{k_4} \rangle'$ denotes the connected trispectrum without its momentum-conserving delta function, i.e.,

$$\langle \zeta_{k_1} \zeta_{k_2} \zeta_{k_3} \zeta_{k_4} \rangle = \langle \zeta_{k_1} \zeta_{k_2} \zeta_{k_3} \zeta_{k_4} \rangle' (2\pi)^3 \delta^{(3)}\left(\sum k_i\right) + \text{disc.}, \quad (62)$$

where “+ disc.” denotes the disconnected contributions to the 4-point function.

One possible signal is the “local” trispectrum $g_{\text{NL}}^{\text{local}}$, which arises if the non-Gaussian adiabatic curvature ζ is of cubic-type form (see, e.g., Okamoto & Hu (2002)), i.e.,

$$\zeta(\mathbf{x}) = \zeta_{\text{G}}(\mathbf{x}) + \frac{9}{25} g_{\text{NL}}^{\text{local}} \zeta_{\text{G}}(\mathbf{x})^3 \quad (63)$$

where $g_{\text{NL}}^{\text{local}}$ is a free parameter and ζ_{G} is a Gaussian field. In this model, the bispectrum is zero (since there is a $\zeta \rightarrow -\zeta$ symmetry) and the trispectrum is given by

$$\langle \zeta_{k_1} \zeta_{k_2} \zeta_{k_3} \zeta_{k_4} \rangle' = \frac{54}{25} g_{\text{NL}}^{\text{local}} \left[P_{\zeta}(k_1) P_{\zeta}(k_2) P_{\zeta}(k_3) + 3 \text{ perm.} \right]. \quad (64)$$

Analogously to the case of the local bispectrum, $f_{\text{NL}}^{\text{local}}$, the observational signal-to-noise for $g_{\text{NL}}^{\text{local}}$ is largest in the “squeezed” limit, $k_1 \ll \min(k_2, k_3, k_4)$, and there is a consistency relation which shows that in single-field inflation, the four-point function is always small in the squeezed limit (e.g., Senatore & Zaldarriaga 2012a). Thus $g_{\text{NL}}^{\text{local}}$ can only be detectably large in multi-field models. Conversely, there are multi-field models where $g_{\text{NL}}^{\text{local}}$ is detectable. The main obstacle here is technical naturalness, i.e., ensuring that radiative corrections do not generate an observationally larger bispectrum. This can be the case if a large bispectrum is forbidden by a Z_2 symmetry, or by supersymmetry (Senatore & Zaldarriaga 2012b).

A further category of four-point signals can be obtained by adding quartic interactions to the inflationary action. Following Smith et al. (2015), we will concentrate on the simplest possibility, by considering quartic operators consistent with the symmetries of inflation and having the lowest possible number of derivatives (Bartolo et al. 2010b; Senatore & Zaldarriaga 2011; Senatore & Zaldarriaga 2012b). There are three such operators, of the schematic form ∂^4 , $\partial^2(\partial_i \sigma)^2$, and $(\partial_i \sigma)^2(\partial_j \sigma)^2$. By a short calculation using the in-in formalism (Maldacena 2003),

Table 27. Correlation coefficients between pairs of bispectrum modes, extracted from two of the *Planck* component separated maps and the WMAP foreground-cleaned map at $\ell_{\text{max}} = 500$ resolution. We separately present the results estimated from $\ell_1 + \ell_2 + \ell_3 = \text{even}$ (Even) and $\ell_1 + \ell_2 + \ell_3 = \text{odd}$ (Odd) combinations. The loss of the correlations is confirmed in the *EEE* case, like Table 14.

Methods	<i>TTT</i>		<i>EEE</i>	
	Even	Odd	Even	Odd
SMICA–SEVEM . .	1.00	0.99	0.80	0.80
SMICA–NILC . . .	1.00	1.00	0.90	0.87
SEVEM–NILC . . .	0.99	1.00	0.70	0.60
SMICA–WMAP .	0.75	0.67




Fig. 26. Parity-odd signals ($\ell_1 + \ell_2 + \ell_3 = \text{odd}$) of the *TTT* (top) and *EEE* (bottom) bispectra ($\ell_i \leq 500$) recovered from the SMICA maps by means of the Modal decomposition with 101 simple polynomial-based eigenmodes, not including any special functions fitting the CMB tensor-mode features. In the panel for *EEE*, only the signals larger than $\ell = 40$ are shown. The *TTT* and *EEE* bispectra shown here are rescaled with a constant Sachs-Wolfe weighting and signal-to-noise weighting, respectively.

the associated four-point functions are (Smith et al. 2015; see also Huang & Shiu 2006):

$$\begin{aligned} \langle \zeta_{k_1} \zeta_{k_2} \zeta_{k_3} \zeta_{k_4} \rangle' &= \frac{9216}{25} g_{\text{NL}}^{\partial^4} A_{\zeta}^3 \int_{-\infty}^0 d\tau_E \tau_E^4 \left(\prod_{i=1}^4 \frac{e^{k_i \tau_E}}{k_i} \right) \\ &= \frac{221184}{25} g_{\text{NL}}^{\partial^4} A_{\zeta}^3 \frac{1}{k_1 k_2 k_3 k_4 K^5}; \end{aligned} \quad (65)$$

$$\begin{aligned}
 \langle \zeta_{k_1} \zeta_{k_2} \zeta_{k_3} \zeta_{k_4} \rangle' &= -\frac{13824}{325} g_{\text{NL}}^{\sigma^2(\partial\sigma)^2} A_\zeta^3 \int_{-\infty}^0 d\tau_E \tau_E^2 \\
 &\quad \times \left[\frac{(1-k_3\tau_E)(1-k_4\tau_E)}{k_1 k_2 k_3^3 k_4^3} (\mathbf{k}_3 \cdot \mathbf{k}_4) e^{\sum k_i \tau_E} \right. \\
 &\quad \left. + 5 \text{ perm.} \right] \\
 &= -\frac{27648}{325} g_{\text{NL}}^{\sigma^2(\partial\sigma)^2} A_\zeta^3 \\
 &\quad \times \left[\frac{K^2 + 3(k_3 + k_4)K + 12k_3 k_4}{k_1 k_2 k_3^3 k_4^3 K^5} (\mathbf{k}_3 \cdot \mathbf{k}_4) \right. \\
 &\quad \left. + 5 \text{ perm.} \right]; \quad (66)
 \end{aligned}$$

$$\begin{aligned}
 \langle \zeta_{k_1} \zeta_{k_2} \zeta_{k_3} \zeta_{k_4} \rangle' &= \frac{82944}{2575} g_{\text{NL}}^{(\partial\sigma)^4} A_\zeta^3 \int_{-\infty}^0 d\tau_E \\
 &\quad \times \left[\prod_{i=1}^4 \frac{(1-k_i\tau_E) e^{k_i\tau_E}}{k_i^3} \right] \\
 &\quad \times [(\mathbf{k}_1 \cdot \mathbf{k}_2)(\mathbf{k}_3 \cdot \mathbf{k}_4) + 2 \text{ perm.}] \\
 &= \frac{165888}{2575} g_{\text{NL}}^{(\partial\sigma)^4} A_\zeta^3 \\
 &\quad \times \left(\frac{2K^4 - 2K^2 \sum k_i^2 + K \sum k_i^3 + 12k_1 k_2 k_3 k_4}{k_1^3 k_2^3 k_3^3 k_4^3 K^5} \right) \\
 &\quad \times [(\mathbf{k}_1 \cdot \mathbf{k}_2)(\mathbf{k}_3 \cdot \mathbf{k}_4) + 2 \text{ perm.}]. \quad (67)
 \end{aligned}$$

Here $K = \sum k_i$ and we have introduced parameters $g_{\text{NL}}^{\sigma^4}$, $g_{\text{NL}}^{\sigma^2(\partial\sigma)^2}$, and $g_{\text{NL}}^{(\partial\sigma)^4}$ to parametrize the amplitude of each trispectrum.

The normalization of the g_{NL} -parameters is chosen so that $\langle \zeta_{k_1} \zeta_{k_2} \zeta_{k_3} \zeta_{k_4} \rangle = (216/25) g_{\text{NL}} A_\zeta^3 / k^9$ for ‘‘tetrahedral’’ configurations, with $|\mathbf{k}_i| = k$ and $(\mathbf{k}_i \cdot \mathbf{k}_j) = -k^3/3$. This is the analogue of the commonly-used normalization for the bispectrum, where f_{NL} parameters are defined so that $\langle \zeta_{k_1} \zeta_{k_2} \zeta_{k_3} \rangle = (18/5) f_{\text{NL}} A_\zeta^2 / k^6$ for equilateral configurations with $|\mathbf{k}_i| = k$.

For simplicity in Eqs. (65)–(67) we have assumed a scale-invariant initial power spectrum $P_\zeta(k) = A_\zeta/k^3$. In order to analyse *Planck* data, we must slightly generalize this to a power-law spectrum $P_\zeta(k) \propto k^{n_s-4}$. Our scheme for doing this follows Appendix C of Smith et al. (2015).

A Fisher matrix analysis shows that there is one large correlation among the three trispectra in Eqs. (65)–(67), so that to an excellent approximation we can treat only two of the trispectra as independent. To quantify this, in Smith et al. (2015) it is shown that the $\sigma^2(\partial\sigma)^2$ shape is 98.6% correlated to a linear combination of the shapes σ^4 and $(\partial\sigma)^4$. Therefore, we will only search for the parameters $g_{\text{NL}}^{\sigma^4}$ and $g_{\text{NL}}^{(\partial\sigma)^4}$.

We note that the analysis that leads to the trispectrum shapes $g_{\text{NL}}^{\sigma^4}$ and $g_{\text{NL}}^{(\partial\sigma)^4}$ is very similar to the analysis that leads to the ‘‘standard’’ bispectrum shapes $f_{\text{NL}}^{\text{equi}}$ and $f_{\text{NL}}^{\text{ortho}}$. However, there are some minor differences as follows. In the bispectrum case, one considers the cubic operators π^3 and $\pi(\partial\pi)^2$, but it is conventional to define observables $f_{\text{NL}}^{\text{equi}}$, $f_{\text{NL}}^{\text{ortho}}$ which are related to the operator coefficients by a linear transformation. This is done because the two cubic operators are about 90% correlated, so it is convenient to orthogonalize. In the trispectrum case the correlation is smaller (around 60% for *Planck*), and we have chosen to omit the orthogonalization step. Another reason to omit the orthogonalization step is that the trispectrum shape $(\partial\sigma)^4$ is a signature of multi-field inflation. In single field inflation, the $(\partial\sigma)^4$

trispectrum is not technically natural; radiative corrections generate cubic operators of the form π^3 or $\pi(\partial\pi)^2$, which generate a bispectrum with larger signal-to-noise.

There are more trispectrum shapes one might consider. For example, classifying Galilean invariant quartic operators leads to higher-derivative trispectra, which are not highly correlated to the trispectra considered above (Bartolo et al. 2013; Arroja et al. 2013). We have only considered ‘‘contact’’ diagrams arising from one power of a quartic operator, and it would be interesting to study ‘‘exchange’’ diagrams arising from two cubic operators and exchange of a light particle (e.g., Chen et al. 2009; Arroja et al. 2009; Chen & Wang 2010; Bartolo et al. 2010a; Baumann & Green 2012). We leave these as extensions for future work.

Summarizing, we will search for the following trispectrum signals:

$$\left\{ g_{\text{NL}}^{\text{local}}, g_{\text{NL}}^{\sigma^4}, g_{\text{NL}}^{(\partial\sigma)^4} \right\} \quad (68)$$

defined by Eqs. (64), (65), and (67) above.

9.1. Data analysis

Turning now to data analysis, we use the machinery from Smith et al. (2015). The first step is to represent each trispectrum as a small sum of factorizable terms as follows. The angular CMB trispectrum can be written either as an integral over comoving distance r (in the case of $g_{\text{NL}}^{\text{local}}$) or as a double integral over (τ, r) where τ is conformal time during inflation (in the case of the σ^4 or $(\partial\sigma)^4$ trispectra). We approximate the integral by a finite sum, which represents the CMB trispectrum as a sum of terms that are factorizable in a sense defined in Smith et al. (2015). A large number of sampling points are needed to obtain a good approximation to the integral, leading to a large number of terms in the factorizable representation. However, there exists an optimization algorithm, which takes as input a trispectrum that has been represented as a sum of many factorizable terms, and outputs a representation with fewer terms. The reduction can be quite dramatic, as shown in Table 28. The optimization algorithm guarantees that the output trispectrum accurately approximates the input trispectrum, in the sense that the two are nearly observationally indistinguishable.

Table 28. Number of factorizable terms N_{in} needed to represent each trispectrum by direct sampling of the integral, and number of terms N_{out} obtained after running the optimization algorithm from section VII of Smith et al. (2015).

Trispectrum	N_{in}	N_{out}
$g_{\text{NL}}^{\text{local}}$	436	17
σ^4	6955	73
$(\partial\sigma)^4$	20865	192

Armed with ‘‘small’’ factorizable representations for each trispectrum, the next step is to run an analysis pipeline that estimates the amplitude of each trispectrum from *Planck* data. We use the ‘‘pure MC’’ pipeline from Smith et al. (2015), which compares the trispectrum of the data to the mean trispectrum of an ensemble of simulations. This pipeline requires a filtering operation $d \rightarrow \tilde{a}_{\ell m}$ which processes the pixel-space CMB data d and generates a harmonic-space map $\tilde{a}_{\ell m}$. Our filtering operation is defined by the following steps:

1. Starting from the data d , we compute (with uniform pixel weighting) a best-fit monopole and dipole outside the Galactic mask. We use the temperature “common mask”, the union of the confidence masks for the SMICA, SEVEM, NILC, and Commander component separation methods (Planck Collaboration IX 2016).
2. The mask defines a few “islands”, i.e., isolated groups of pixels that are unmasked, but contained in a larger masked region. We slightly enlarge the mask so that it removes the islands.
3. We classify the components of the masked part of the sky into “small” masked regions with ≤ 1000 pixels (at HEALPix resolution $N_{\text{side}} = 2048$), and “large” regions with > 1000 pixels. Small regions usually correspond to point sources, and large regions typically correspond to areas of diffuse galactic emission. In small regions, we inpaint the CMB by assigning the unique map that agrees with the data on boundary pixels, and whose value in each interior pixel is the average of the neighboring pixels.
4. In large regions, we do not inpaint the CMB, but rather apodize the boundary of the large region using cosine apodization with 12’ radius.
5. We apply a spherical harmonic transform to the inpainted, apodized CMB map to obtain a harmonic-space map $a_{\ell m}$ with $\ell_{\text{max}} = 1600$. We then take the final filtered map $\tilde{a}_{\ell m}$ to be

$$\tilde{a}_{\ell m} = \frac{a_{\ell m}}{b_{\ell} C_{\ell} + b_{\ell}^{-1} N_{\ell}} \quad (69)$$

where b_{ℓ} is the beam, C_{ℓ} is the fiducial CMB power spectrum, and N_{ℓ} is the sky-averaged noise power spectrum (without beam deconvolution). To motivate this choice of ℓ -weighting, we note that for an ideal all-sky experiment with isotropic noise, we have $a_{\ell m} = b_{\ell} s_{\ell m} + n_{\ell m}$ where $s_{\ell m}, n_{\ell m}$ are signal and noise realizations. In this case, Eq. (69) weights the signal as $s_{\ell m}/(C_{\ell} + b_{\ell}^{-2} N_{\ell})$, which is optimal.

In our pipeline, we apply this filter to the component-separated SMICA maps (Planck Collaboration IX 2016), obtaining a harmonic-space map $\tilde{a}_{\ell m}$. We apply the same filter to 1000 Monte Carlo simulations to obtain an ensemble of harmonic-space maps. Our pipeline has the property that it always estimates the trispectrum of the data in excess of the trispectrum in the simulations. Since the simulations include lensing, this means that lensing bias will automatically be subtracted from our g_{NL} estimates.

Now that the filter, data realization, and Monte Carlo simulations have been fully specified, the details of the pipeline are described in section IX.B of Smith et al. (2015). For each trispectrum, the pipeline outputs an estimate of g_{NL} and an estimate of the statistical error. Our basic results are:

$$\begin{aligned} g_{\text{NL}}^{\text{local}} &= (-9.0 \pm 7.7) \times 10^4; \\ g_{\text{NL}}^{\sigma^4} &= (-0.2 \pm 1.7) \times 10^6; \\ g_{\text{NL}}^{(\partial\sigma)^4} &= (-0.1 \pm 3.8) \times 10^5. \end{aligned} \quad (70)$$

No deviation from Gaussian statistics is seen. These results significantly improve the previous best constraints on the trispectrum from WMAP (Vielva & Sanz 2010; Smidt et al. 2010; Fergusson et al. 2010b; Hikage & Matsubara 2012; Sekiguchi & Sugiyama 2013; Regan et al. 2013; Smith et al. 2015) and large-scale structure (Desjacques & Seljak 2010; Giannantonio et al. 2014; Leistedt et al. 2014).

A constraint on $g_{\text{NL}}^{\text{local}}$ from *Planck* 2013 data was recently reported by Feng et al. (2015), who find $g_{\text{NL}}^{\text{local}} = (-13 \pm 18) \times 10^4$.

Our central value in Eq. (70) agrees well with this result, but the statistical error is smaller by a factor of 2.3. This improvement is partly due to the lower noise levels in *Planck* 2015 data, and partly due to the use of a better estimator.

Each line in Eq. (70) is a “single- g_{NL} ” constraint; i.e., the constraint on one g_{NL} parameter with the other g_{NL} -parameters held fixed. For joint constraints, one needs to know the full covariance matrix. The correlation between $g_{\text{NL}}^{\text{local}}$ and the other two parameters is negligible, and the $g_{\text{NL}}^{\sigma^4}$ - $g_{\text{NL}}^{(\partial\sigma)^4}$ correlation is:

$$\text{Corr}(g_{\text{NL}}^{\sigma^4}, g_{\text{NL}}^{(\partial\sigma)^4}) = 0.61. \quad (71)$$

multi-field models of inflation will generally give a linear combination of σ^4 , $\sigma^2(\partial_i\sigma)^2$, and $(\partial_i\sigma)^2(\partial_j\sigma)^2$ trispectra. In this case we proceed as follows. First, if the $\sigma^2(\partial_i\sigma)^2$ coefficient is non-zero, we can use the near-degeneracy with a linear combination of the other two operators to absorb it into the effective values of $g_{\text{NL}}^{\sigma^4}$ and $g_{\text{NL}}^{(\partial\sigma)^4}$. A Fisher matrix analysis shows that the coefficients of this linear combination are

$$\begin{aligned} (g_{\text{NL}}^{\sigma^4})_{\text{eff}} &= 0.59 g_{\text{NL}}^{\sigma^2(\partial\sigma)^2} \\ (g_{\text{NL}}^{(\partial\sigma)^4})_{\text{eff}} &= 0.091 g_{\text{NL}}^{\sigma^2(\partial\sigma)^2} \end{aligned} \quad (72)$$

It is convenient to define the two-component parameter vector:

$$g_i = \begin{pmatrix} g_{\text{NL}}^{\sigma^4} \\ g_{\text{NL}}^{(\partial\sigma)^4} \end{pmatrix}. \quad (73)$$

We also compute a two-by-two Fisher matrix F_{ij} , whose diagonal is given by $F_{ii} = 1/\sigma_i^2$, where σ_i is the single- g_{NL} statistical error in Eq. (70), and whose off-diagonal is $F_{12} = r F_{11}^{1/2} F_{22}^{1/2}$, where r is the correlation in Eq. (71). This procedure gives:

$$F_{ij} = \begin{pmatrix} 3.3 & 9.2 \\ 9.2 & 68.7 \end{pmatrix} \times 10^{-13}. \quad (74)$$

For a given parameter vector g_i , we can define a trispectrum- χ^2 by

$$\chi^2(g) = [F_{ii}\hat{g}_i - (Fg)_i] F_{ij}^{-1} [F_{jj}\hat{g}_j - (Fg)_j] \quad (75)$$

where $\hat{g}_i = (-0.21 \times 10^6, -0.10 \times 10^5)$ is the vector of best-fit single- g_{NL} values from Eq. (70). This definition of χ^2 follows from the observation that $(F_{ii}\hat{g}_i)$ is an estimator with expectation value $(Fg)_i$ and covariance matrix $\text{Cov}(F_{ii}\hat{g}_i, F_{jj}\hat{g}_j) = F_{ij}$.

The inflationary implications of these trispectrum constraints are discussed in Sect. 11.5 below.

10. Minkowski functionals results

In this section, we present constraints on local NG at first and second order ($f_{\text{NL}}^{\text{local}}$ and $g_{\text{NL}}^{\text{local}}$) obtained with Minkowski functionals (MFs) on temperature and polarization E maps. MFs (Mecke et al. 1994; Schmalzing & Buchert 1997; Schmalzing & Gorski 1998; Winitzki & Kosowsky 1998) are a measure of fields’ local morphology used to constrain their stationarity, isotropy and Gaussianity. Mostly probing general NG in a frequentist fashion in two-dimensions on CMB maps (Eriksen et al. 2004; Komatsu et al. 2005; Modest et al. 2013; Natoli et al. 2010; Curto et al. 2008) or three-dimensions on LSS data (Park et al. 2005; Wiegand et al. 2014), they have also been used to measure specific NG targets with Bayesian methods, such as $f_{\text{NL}}^{\text{local}}$ (Hikage et al. 2006, 2008; Ducout et al. 2013; Planck Collaboration XXIV 2014), other bispectrum and

trispectrum shapes (Hikage & Matsubara 2012) and topological defects (Planck Collaboration XXV 2014). New developments have been made recently, using needlets (Fantaye et al. 2015), neural networks (Novaes et al. 2014) or allowing scale-dependent measurements (Munshi et al. 2013).

MF-based limits are well known to be suboptimal for $f_{\text{NL}}^{\text{local}}$ and $g_{\text{NL}}^{\text{local}}$, but they provide an independent cross-check of bispectrum and trispectrum-based estimators. They are complementary to optimal estimators: they are weighted integrals of the polyspectra and are sensitive to any source of NG, including foregrounds and secondaries. While MFs are not always able to distinguish between these different sources and systematics, they allow upper bounds to be put on them.

The most recent constraints (1σ) from MFs on $f_{\text{NL}}^{\text{local}}$ and $g_{\text{NL}}^{\text{local}}$ have been obtained with WMAP (Hikage & Matsubara 2012) and *Planck* (Planck Collaboration XXIV 2014):

$$f_{\text{NL}}^{\text{local}} = 4.2 \pm 20.5; \quad g_{\text{NL}}^{\text{local}} = (1.9 \pm 6.4) \times 10^5. \quad (76)$$

10.1. Method and definition of MFs

For a smoothed two-dimensional field δ of zero mean and of variance σ_0^2 , defined on the sphere, we consider an excursion set of height $\nu = \delta/\sigma_0$, i.e., the set of points where the field exceeds the threshold ν . We use four functionals denoted by $V_k(\nu)$ ($k = 0, 1, 2, 3$). The first three correspond to MFs: V_0 is the fractional Area of the regions above the threshold, V_1 is the Perimeter of these regions and V_2 is the Genus, defined as the total number of connected components of the excursion above the threshold minus the total number of connected components under the threshold. The fourth, V_3 , is the Number of clusters (also referred to as N_{clusters}). This is the number of connected regions above the threshold for positive thresholds and below the threshold for negative thresholds. Precise definitions and formulae for the quantities V_k as well as their expectation values for Gaussian fields are summarized in Appendix B.

We calculate the four normalized¹⁸ functionals $v_k(\nu)$ on $n_{\text{th}} = 26$ thresholds ν , between $\nu_{\text{min}} = -3.5$ and $\nu_{\text{max}} = +3.5$.

For this analysis, we used the same temperature and polarization E data, simulations and masks described in Sect. 3.4 for consistency with the bispectrum estimators. In addition, the maps are filtered to optimize constraints on local NG (Ducout et al. 2013), the filters used being similar to Wiener filters for T and E (W_{M}), and for the first (W_{D1}) and second (W_{D2}) derivatives of these fields (Fig. 27):

$$W_{\text{D1}} \propto \sqrt{\ell(\ell+1)} W_{\text{M}}; \quad (77)$$

$$W_{\text{D2}} \propto \ell(\ell+1) W_{\text{M}}. \quad (78)$$

For the temperature map, known point sources in the mask are inpainted.

We define the vector y as any combination $y = \{v_k^{A,W}\}$ with $k = \{0, 3\}$, $A = \{T, E\}$, $W = \{W_{\text{M}}, W_{\text{D1}}, W_{\text{D2}}\}$, \hat{y} being the vector measured on the data.

¹⁸ Raw Minkowski functionals V_k depend on the Gaussian part of fields through a normalization factor A_k , a function only of the shape of the power spectrum. We therefore normalize functionals $v_k = V_k/A_k$ to focus on NG, see Appendix B.

From these measurements, we then use a Bayesian method to jointly estimate $f_{\text{NL}}^{\text{local}}$ and $g_{\text{NL}}^{\text{local}}$,

$$P(f_{\text{NL}}^{\text{local}}, g_{\text{NL}}^{\text{local}} | \hat{y}) = \frac{P(\hat{y} | f_{\text{NL}}^{\text{local}}, g_{\text{NL}}^{\text{local}}) P(f_{\text{NL}}^{\text{local}}, g_{\text{NL}}^{\text{local}})}{\int P(\hat{y} | f_{\text{NL}}^{\text{local}}, g_{\text{NL}}^{\text{local}}) P(f_{\text{NL}}^{\text{local}}, g_{\text{NL}}^{\text{local}}) df_{\text{NL}}^{\text{local}} dg_{\text{NL}}^{\text{local}}}. \quad (79)$$

We take a uniform prior for $f_{\text{NL}}^{\text{local}}$ in the range -400 to 400 , and for $g_{\text{NL}}^{\text{local}}$ in the range -4×10^6 to $+4 \times 10^6$, while the evidence is just considered as a normalization.

Assuming MFs are multi-variate Gaussian-distributed we obtain the posterior distribution for $(f_{\text{NL}}^{\text{local}}, g_{\text{NL}}^{\text{local}})$ with a χ^2 test

$$P(f_{\text{NL}}^{\text{local}}, g_{\text{NL}}^{\text{local}} | \hat{y}) \propto \exp \left[-\frac{\chi^2(\hat{y}, f_{\text{NL}}^{\text{local}}, g_{\text{NL}}^{\text{local}})}{2} \right], \quad (80)$$

with

$$\chi^2(\hat{y}, f_{\text{NL}}^{\text{local}}, g_{\text{NL}}^{\text{local}}) \equiv \left[\hat{y} - \bar{y}_{\text{sim1}}(f_{\text{NL}}^{\text{local}}, g_{\text{NL}}^{\text{local}}) \right]^T C_{\text{sim2}}^{-1} \left[\hat{y} - \bar{y}_{\text{sim1}}(f_{\text{NL}}^{\text{local}}, g_{\text{NL}}^{\text{local}}) \right]. \quad (81)$$

For this test, we use two types of simulations to first construct a model including primordial NG $\bar{y}_{\text{sim1}}(f_{\text{NL}}^{\text{local}}, g_{\text{NL}}^{\text{local}})$ and secondly a covariance matrix

$$C_{\text{sim2}} \equiv \langle (y_{\text{sim2}} - \bar{y}_{\text{sim2}})(y_{\text{sim2}} - \bar{y}_{\text{sim2}})^T \rangle, \quad (82)$$

with $\bar{y}_{\text{sim}} \equiv \langle y_{\text{sim}} \rangle_{\text{sim}}$ averaged over the simulations. We now describe the details of these simulations.

– Simulations 1: Non-Gaussian model

For the first type of simulation, we included all possible sources of NG, assuming that the total and individual levels of NG are small enough that MFs are linear with respect to those NG levels (Ducout et al. 2013). The three kinds of NG we included are foreground residuals (Galactic residuals with scalable amplitude α , as well as radio sources, CIB anisotropies, secondaries (SZ, lensing and ISW-lensing, but not SZ-lensing) and primordial NG ($f_{\text{NL}}^{\text{local}}, g_{\text{NL}}^{\text{local}}$):

$$\text{sim1}^i = \text{map}_{\text{lensed}}^i(f_{\text{NL}}^{\text{local}}, g_{\text{NL}}^{\text{local}}) + \text{map}^{\text{fg}}(\text{radio sources, CIB, SZ}) + \alpha \times \text{map}^{\text{fg}}(\text{Galactic residuals}). \quad (83)$$

We tried to reproduce all instrumental effects, with realistic effective beams (isotropic window functions), noise from FFP8 simulations (Planck Collaboration XII 2016), filtered with component separation weights. We checked the accuracy of these simulations by comparing them to FFP8 simulations, using no foreground and no primordial NG. The astrophysical models are provided by the *Planck Sky Model* (PSM, Delabrouille et al. 2013), while the primordial NG simulations are computed as in Elsner & Wandelt (2009). The lensing uses LensPix¹⁹. The power spectrum used for these NG simulations and the lensing is the best-fit power spectrum from *Planck* 2013+ACT/SPT+BAO (Planck Collaboration XXII 2014). We created $n_1 = 200$ simulations i , using n_1 maps for the primordial NG, while we had only one astrophysical foreground simulation.

¹⁹ <http://cosmologist.info/lenspix>




Fig. 27. Filters used to optimize constraints on local NG, in harmonic space. The temperature filter W_M is a smoothed version of the true Wiener filter obtained with realistic models, while the E -mode W_M filter is adapted from the temperature one, with a cutoff value at $\ell \simeq 800$. The formulae for the derivative filters are given in Eq. (78).

– *Simulations 2: FFP8 (Planck Collaboration XII 2016) MC simulations*

Since NG is weak, the covariance matrix C is computed with $n_2 = 10^4$ simulations, including no primordial NG and no foregrounds. These simulations reproduce realistic instrumental effects (anisotropy of beams in particular), realistic noise and component separation filtering. The only NG still present in these simulations are lensing and the ISW-lensing correlation.

Validation of the estimator

Part of the validation for the MFs estimator is described in Sect. 5.3 for $f_{\text{NL}}^{\text{local}}$, to compare the results to bispectrum estimators on realistic simulations (FFP8 MC, second item above). In addition we present in Table 29 the results obtained on the same realistic simulations for $g_{\text{NL}}^{\text{local}}$, and on simulations containing primordial NG (first item above), with $f_{\text{NL}}^{\text{local}} = 10$ and $g_{\text{NL}}^{\text{local}} = 10^5$, with 200 simulations used in each case. These tests have been performed using the SMICA method with lensing bias removed.

10.2. Results

Results for $f_{\text{NL}}^{\text{local}}$ and $g_{\text{NL}}^{\text{local}}$ estimation with MFs on the four component separated maps in temperature and polarization are presented in Table 30. The results for polarization E -only maps are not quoted, since these results were not sufficiently stable (cf. Sect. 7.6). No deviation from Gaussianity is detected. $T+E$ analysis generally finds higher values for $f_{\text{NL}}^{\text{local}}$, but remains consistent with Gaussianity.

The posteriors for $f_{\text{NL}}^{\text{local}}$ and $g_{\text{NL}}^{\text{local}}$ from SMICA are shown in Fig. 28. One interesting point is that the estimates of $f_{\text{NL}}^{\text{local}}$ and $g_{\text{NL}}^{\text{local}}$ are almost uncorrelated ($r < 0.1$); this can be inferred when we consider the parity of MF deviations from Gaussianity, which is different for the two parameters (Matsubara 2010).

Foreground and secondary biases are removed from these estimates, since the NG model directly includes them. However, an estimation of their contribution in the map is reported in Table 31.

Table 29. Results for local NG parameters at first and second order, $f_{\text{NL}}^{\text{local}}$ and $g_{\text{NL}}^{\text{local}}$, obtained with Minkowski functionals on SMICA simulations in temperature and polarization. These results are corrected for the lensing and ISW-lensing biases unless stated otherwise. Parameters are estimated jointly, and we report marginalized results, quoting 1σ errors. The results are the average obtained from 200 simulations.

	$f_{\text{NL}}^{\text{local}}$	$g_{\text{NL}}^{\text{local}} (\times 10^4)$
$f_{\text{NL}}^{\text{local}} = 0, g_{\text{NL}}^{\text{local}} = 0$		
T	0 ± 13	-1 ± 19
E	1 ± 42	0 ± 23
$T + E$	1 ± 12	0 ± 13
$f_{\text{NL}}^{\text{local}} = 10, g_{\text{NL}}^{\text{local}} = 10^5$		
T	10 ± 13	9 ± 22
E	12 ± 42	10 ± 23
$T + E$	11 ± 12	10 ± 13

Table 30. Results for local NG parameters at first and second order, $f_{\text{NL}}^{\text{local}}$ and $g_{\text{NL}}^{\text{local}}$, obtained with Minkowski functionals on all four component separated maps in temperature and polarization. These results are corrected for the lensing and ISW-lensing biases unless stated otherwise. Parameters are estimated jointly, and we report marginalized results, quoting 1σ errors.

	$f_{\text{NL}}^{\text{local}}$	$g_{\text{NL}}^{\text{local}} (\times 10^4)$
SMICA		
T	2 ± 13	-17 ± 19
$T + E$	3 ± 12	-8 ± 13
SEVEM		
T	3 ± 13	-23 ± 20
$T + E$	7 ± 12	-9 ± 13
NILC		
T	10 ± 13	-23 ± 20
$T + E$	12 ± 12	-15 ± 13
Commander		
T	8 ± 13	-30 ± 19
$T + E$	10 ± 13	-18 ± 13

Foreground and secondary biases

Foreground residuals are generally negligible, in particular in the T analysis. This is different from the *Planck* 2013 results, where the residuals were more important; this can be explained by the beam correction applied to these previous estimates which exaggerated signals from small scales.

Lensing has a significant signature in MF estimation of $f_{\text{NL}}^{\text{local}}$, but is even stronger in $g_{\text{NL}}^{\text{local}}$ (the four-point correlation signature) and could be detected (and not treated just as a bias) with this estimator. The Wiener filters enhance the scales where lensing is dominant.




Fig. 28. Joint constraint on $f_{\text{NL}}^{\text{local}}$ and $g_{\text{NL}}^{\text{local}}$ obtained with MFs. The contour lines represent 1, 2 and 3 σ limits of a 2D-Gaussian distribution. Constraints were obtained with SMICA temperature and polarization E maps.

Table 31. Biases for local NG parameters at first and second order $f_{\text{NL}}^{\text{local}}$, $g_{\text{NL}}^{\text{local}}$ obtained with Minkowski Functionals on SMICA in temperature and polarization. Parameters are estimated jointly, and we report marginalized results. For the corresponding error (on one map), see Table 30.

	$\Delta f_{\text{NL}}^{\text{local}}$		$\Delta g_{\text{NL}}^{\text{local}} (\times 10^4)$	
	T	$T + E$	T	$T + E$
SMICA				
SZ	0.0	-0.3	2.3	1.1
CIB	0.7	0.5	-6.8	3.4
Galaxy	-0.1	-0.2	-1.2	3.1
PS	0.1	0.2	2.2	1.2
Lensing	16.5	10.0	63.1	40.4

11. Implications for early Universe physics

The NG constraints obtained in this paper show consistency of *Planck* data with Gaussian primordial fluctuations, thus confirming the results obtained in the 2013 release (Planck Collaboration XXIV 2014) and improving them through the inclusion of CMB polarization data. The standard single-field slow-roll models of inflation have therefore been confirmed as a viable scenario for inflation, passing one of their most stringent tests, based on lack of measurable deviations from Gaussianity. The constraints obtained on local, equilateral, and orthogonal NG, after accounting for various contaminants, strongly limit different mechanisms proposed as alternatives to the standard inflationary models to explain the seeds of cosmological perturbations. Measurements on deviations from Gaussianity for other primordial bispectral shapes help to shed light on more subtle effects about the detailed physics of inflation.

As in Planck Collaboration XXIV (2014), in the following we derive limits on parameters of the models from the NG con-

straints in the following way (unless explicitly stated otherwise): we construct a posterior based on the assumption that the sampling distribution is Gaussian (as supported by Gaussian simulations); the likelihood is approximated by the sampling distribution, but centred on the NG estimate (see Elsner & Wandelt 2009); we employ uniform or Jeffreys' priors, over intervals of the parameters values that are physically meaningful, or as otherwise stated; and in the cases when two or more parameters are involved, we marginalize the posterior to provide one-dimensional constraints on the parameter considered.

11.1. General single-field models of inflation

DBI models: DBI models of inflation (Silverstein & Tong 2004; Alishahiha et al. 2004), characterized by a non-standard kinetic term of the inflaton field, predict a non-linearity parameter $f_{\text{NL}}^{\text{DBI}} = -(35/108)(c_s^{-2} - 1)$, where c_s is the sound speed of the inflaton perturbations (Silverstein & Tong 2004; Alishahiha et al. 2004; Chen et al. 2007b). The corresponding bispectrum shape is very close to the equilateral shape. Nonetheless we have constrained the exact theoretical (non-separable) shape (see equation 7 of Planck Collaboration XXIV 2014). The constraint we obtain $f_{\text{NL}}^{\text{DBI}} = 2.6 \pm 61.6$ from temperature data ($f_{\text{NL}}^{\text{DBI}} = 15.6 \pm 37.3$ from temperature and polarization) at 68 % CL (with ISW-lensing and point sources subtracted, see Table 23) implies

$$c_s^{\text{DBI}} \geq 0.069 \quad 95 \% \text{ CL } (T\text{-only}), \quad (84)$$

and

$$c_s^{\text{DBI}} \geq 0.087 \quad 95 \% \text{ CL } (T+E). \quad (85)$$

In Planck Collaboration XXIV (2014) we constrained the so-called infrared (IR) DBI models (Chen 2005b,a), which arise in string frameworks. We focused on a minimal setup, considering a regime where stringy effects are negligible and predictions for primordial perturbations are built within standard field theory. In the companion paper Planck Collaboration XX (2016) we present an analysis of a more general class of IR DBI models which accounts for stringy signatures (see Bean et al. 2008) by combining *Planck* power spectrum and bispectrum constraints.

Implications for effective field theory of Inflation: In this subsection we use the effective field theory approach to inflation in order to translate the constraints on $f_{\text{NL}}^{\text{equil}}$ and $f_{\text{NL}}^{\text{ortho}}$ into limits on the parameters of the Lagrangian of general single-field models of inflation (of the type $P(X, \varphi)$ models). In particular we derive the most conservative bound on the sound speed of the inflaton perturbations for this class of models.

The effective field theory approach (Cheung et al. 2008; Weinberg 2008) provides an efficient way to constrain inflationary perturbations for various classes of models that incorporate deviations from the standard single-field slow-roll scenario. In this approach the Lagrangian of the system is expanded into the (lowest dimension) operators obeying the underlying symmetries. We consider general single-field models described by the following action

$$S = \int d^4x \sqrt{-g} \left[-\frac{M_{\text{Pl}}^2 \dot{H}}{c_s^2} \left(\dot{\pi}^2 - c_s^2 \frac{(\partial_i \pi)^2}{a^2} \right) - M_{\text{Pl}}^2 \dot{H} (1 - c_s^{-2}) \dot{\pi} \frac{(\partial_i \pi)^2}{a^2} + \left(M_{\text{Pl}}^2 \dot{H} (1 - c_s^{-2}) - \frac{4}{3} M_3^4 \right) \dot{\pi}^3 \right], \quad (86)$$

where the curvature perturbation is related to the scalar field π as $\zeta = -H\pi$. The inflaton interaction terms $\dot{\pi}(\partial_i \pi)^2$ and $(\dot{\pi})^3$




Fig. 29. 68 %, 95 %, and 99.7 % confidence regions in the parameter space $(f_{\text{NL}}^{\text{equil}}, f_{\text{NL}}^{\text{ortho}})$, defined by thresholding χ^2 as described in the text.

generate two kind of bispectra with amplitudes, respectively, $f_{\text{NL}}^{\text{EFT1}} = -(85/324)(c_s^{-2} - 1)$ and $f_{\text{NL}}^{\text{EFT2}} = -(10/243)(c_s^{-2} - 1)[\tilde{c}_3 + (3/2)c_s^2]$, where M_3 is the amplitude of the operator $\hat{\pi}^3$ (Senatore et al. 2010, see also Chen et al. 2007b; Chen 2010b). These two bispectra both peak for equilateral triangles in Fourier space. Nevertheless, they are sufficiently different, and the total NG signal turns out to be a linear combination of the two, leading also to an orthogonal shape. We can put constraints on c_s and the dimensionless parameter $\tilde{c}_3(c_s^{-2} - 1) = 2M_3^4 c_s^2 / (\dot{H} M_{\text{pl}}^2)$ (Senatore et al. 2010). Notice that DBI inflationary models corresponds to having $\tilde{c}_3 = 3(1 - c_s^2)/2$, while $c_s = 1$ and $M_3 = 0$ (or $\tilde{c}_3(c_s^{-2} - 1) = 0$) represent the non-interacting (vanishing NG) case.

The mean values of the estimators for $f_{\text{NL}}^{\text{equil}}$ and $f_{\text{NL}}^{\text{ortho}}$ are expressed in terms of c_s and \tilde{c}_3 by

$$\begin{aligned} f_{\text{NL}}^{\text{equil}} &= \frac{1 - c_s^2}{c_s^2} \left[-0.275 - 0.0780 c_s^2 - (2/3) \times 0.780 \tilde{c}_3 \right] \\ f_{\text{NL}}^{\text{ortho}} &= \frac{1 - c_s^2}{c_s^2} \left[0.0159 - 0.0167 c_s^2 - (2/3) \times 0.0167 \tilde{c}_3 \right]. \end{aligned} \quad (87)$$

Here the coefficients come from the Fisher matrix between the equilateral and orthogonal templates and the theoretical bispectra predicted by the two operators $\hat{\pi}(\nabla\pi)^2$ and $\hat{\pi}^3$. We use a χ^2 statistic given by $\chi^2(\tilde{c}_3, c_s) = \mathbf{v}^T(\tilde{c}_3, c_s) \mathbf{C}^{-1} \mathbf{v}(\tilde{c}_3, c_s)$, where $\mathbf{v}^i(\tilde{c}_3, c_s) = f^i(\tilde{c}_3, c_s) - f_p^i$ ($i = \{\text{equilateral, orthogonal}\}$), f_p^i being the joint estimates of equilateral and orthogonal f_{NL} (see Table 11), \mathbf{C} the covariance matrix of the joint estimators, and $f^i(\tilde{c}_3, c_s)$ is provided by Eq. (87). As an example in Fig. 29 we show the 68 %, 95 %, and 99.7 % confidence regions for $f_{\text{NL}}^{\text{equil}}$ and $f_{\text{NL}}^{\text{ortho}}$ obtained from the $T + E$ constraints, requiring $\chi^2 \leq 2.28, 5.99$, and 11.62 respectively (corresponding to a χ^2 variable with two degrees of freedom). In Fig. 30 we show the corresponding confidence regions in the (\tilde{c}_3, c_s) parameter space.




Fig. 30. 68 %, 95 %, and 99.7 % confidence regions in the single-field inflation parameter space (c_s, \tilde{c}_3) , obtained from Fig. 29 via the change of variables in Eq. (87).

Marginalizing over \tilde{c}_3 we find

$$c_s \geq 0.020 \quad 95 \% \text{ CL } (T\text{-only}), \quad (88)$$

and

$$c_s \geq 0.024 \quad 95 \% \text{ CL } (T+E). \quad (89)$$

The constraints improve by a few percent in T -only and by up to 25 % by including polarization, in comparison with those of Planck Collaboration XXIV (2014).

Galileon models of inflation: Galileon models of inflation (Burrage et al. 2011; Kobayashi et al. 2010; Mizuno & Koyama 2010; Ohashi & Tsujikawa 2012) are well motivated models based on the so called ‘‘Galilean symmetry’’ (Nicolis et al. 2009). They are characterized by stability properties that are quite well understood (ghost-free, and stable against quantum corrections) and can arise naturally within fundamental physics models (de Rham & Gabadadze 2010b,a). Moreover they are an interesting example of models where gravity is modified on large scales and we focus on them also as a typical example of a more general class of modified gravity theories that are ghost-free (the so called Horndeski theories, Horndeski 1974; see also Planck Collaboration XIV (2016) where these models are discussed in the context of dark energy/modified gravity scenarios for the late time evolution of the universe). The predictions for the primordial perturbations are very rich. Bispectra can be generated with the same shapes as the ‘‘EFT1’’ and ‘‘EFT2’’ bispectra (see also discussion in Creminelli et al. 2011), however the amplitude(s) scale with the fluctuation sound speed as c_s^{-4} , differently from the general single-field models of inflation considered in the above subsection. They can be written as (at the lowest-order in slow-roll

parameters)

$$f_{\text{NL}}^{\text{EFT1}} = \frac{17}{972} \left(-\frac{5}{c_s^4} + \frac{30}{c_s^2} - \frac{40}{c_s \bar{c}_s} + 15 \right) \quad (90)$$

$$f_{\text{NL}}^{\text{EFT2}} = \frac{1}{243} \left(\frac{5}{c_s^4} + \frac{30/A - 55}{c_s^2} + \frac{40}{c_s \bar{c}_s} - 320 \frac{c_s}{\bar{c}_s} - \frac{30}{A} + 275 - 225c_s^2 + 280 \frac{c_s^3}{\bar{c}_s} \right). \quad (91)$$

Here A , \bar{c}_s , and c_s are dimensionless parameters of the models. In particular c_s is the sound speed of the Galileon scalar field, while \bar{c}_s is a parameter that appears to break the standard consistency relation for the tensor-to-scalar perturbation ratio ($r = 16\epsilon\bar{c}_s = -8n_T\bar{c}_s$, n_T being the tensor spectral index)²⁰. Accordingly to Eq. (87) a linear combination of these two bispectra generate equilateral and orthogonal bispectra templates²¹. From the *Planck* constraints on $f_{\text{NL}}^{\text{equil}}$ and $f_{\text{NL}}^{\text{ortho}}$ (see Table 11), we derive constraints on these model parameters following the procedure described at the beginning of this section. We choose log-constant priors in the ranges $10^{-4} \leq A \leq 10^4$, and $10^{-4} \leq \bar{c}_s \leq 10^2$, together with a uniform prior $10^{-4} < c_s < 1$. These priors have been chosen essentially on the basis of perturbative regime validity of the theory and to allow for a quite wide range of parameter values. In Fig. 31, as an example, probability contours are shown in the parameter space (c_s, \bar{c}_s) from the $T + E$ constraints, after marginalizing over the parameter A . Marginalizing over both A and \bar{c}_s we find

$$c_s^{\text{Galileon}} \geq 0.21 \quad 95\% \text{ CL } (T\text{-only}), \quad (92)$$

and

$$c_s^{\text{Galileon}} \geq 0.23 \quad 95\% \text{ CL } (T+E). \quad (93)$$

Notice that interestingly enough the parameter \bar{c}_s can be negative in principle, corresponding to a blue spectral tilt of inflationary gravitational waves (without any kind of instability). We therefore also explore this branch, with a log-constant prior (for $-\bar{c}_s$), $-10^2 \leq \bar{c}_s \leq -10^{-4}$, and the same priors for the other parameters as above. Fig. 32 shows the probability contours in the (c_s, \bar{c}_s) plane, after marginalizing over the parameter A , for the $n_T < 0$ branch. Marginalizing over both A and \bar{c}_s gives

$$c_s^{\text{Galileon}} \geq 0.19 \quad 95\% \text{ CL } (T\text{-only}), \quad (94)$$

and

$$c_s^{\text{Galileon}} \geq 0.21 \quad 95\% \text{ CL } (T+E). \quad (95)$$

A combined analysis of the *Planck* bispectrum and power spectrum constraints on the Galileon models is presented in the companion *Planck* paper on inflation (Planck Collaboration XX 2016).

11.2. Multi-field models

Curvaton models: The simplest adiabatic curvaton models predict local NG with an amplitude (Bartolo et al. 2004c,b)

$$f_{\text{NL}}^{\text{local}} = \frac{5}{4r_D} - \frac{5r_D}{6} - \frac{5}{3}, \quad (96)$$

²⁰ For the explicit expressions of these parameters in terms of the coefficients of the Galileon Lagrangian see Planck Collaboration XX (2016).

²¹ We note that we are neglecting $O(\epsilon_1/c_s^4)$ corrections (where $O(\epsilon_1)$ means also $O(\eta_s, s, \dots)$); see (Burrage et al. 2011) and (Ribeiro & Seery 2011). These corrections will have a different shape associated with them and they are not necessarily small when compared with some of the terms displayed, e.g., the terms $O(1/c_s^2)$.




Fig. 31. 68 %, 95 %, and 99.7 % probability contours in the Galileon models for c_s and \bar{c}_s parameters for the $\bar{c}_s > 0$ branch (tensor spectral index $n_T < 0$).




Fig. 32. 68 %, 95 %, and 99.7 % probability contours in the Galileon models for c_s and \bar{c}_s parameters for the $\bar{c}_s < 0$ branch (blue tensor spectral index $n_T > 0$).

for a quadratic potential of the curvaton field (Lyth & Wands 2002; Lyth et al. 2003; Lyth & Rodriguez 2005; Malik & Lyth 2006; Sasaki et al. 2006), where $r_D = [3\rho_{\text{curvaton}}/(3\rho_{\text{curvaton}} + 4\rho_{\text{radiation}})]_D$ is the ‘‘curvaton decay fraction’’ at the time of the curvaton decay in the sudden decay approximation. Assuming a uniform prior, $0 < r_D < 1$, our constraint $f_{\text{NL}}^{\text{local}} = 2.5 \pm 5.7$ at 68 % CL (see Table 11) yields

$$r_D \geq 0.16 \quad 95\% \text{ CL } (T\text{-only}), \quad (97)$$

while accounting for temperature and polarization data ($f_{\text{NL}}^{\text{local}} = 0.8 \pm 5.0$ at 68 % CL) gives

$$r_D \geq 0.19 \quad 95\% \text{ CL } (T+E), \quad (98)$$

improving the previous *Planck* bound which was previously $r_D \geq 0.15$ (95% CL; [Planck Collaboration XXIV 2014](#)). In [Planck Collaboration XX \(2016\)](#), assuming there is some relic isocurvature fluctuations in the curvaton field, a limit on r_D is derived from the bounds on isocurvature fluctuations. In this restricted case, the limit $r_D > 0.98$ (95% CL) is derived, which is consistent with the constraint given here.

Notice that the above expression of $f_{\text{NL}}^{\text{local}}$ Eq. (96) is valid under the assumption that there is no significant decay of the inflaton into curvaton particles. In general one should account for such a possibility. For example, if the classical curvaton field survives and starts to dominate, then the curvaton particles produced during reheating (which have the same equation of state as the classical curvaton field) are expected to survive and dominate over other species at the epoch of their decay. The classical curvaton field and the curvaton particles decay at the same time, inevitably producing adiabatic perturbations (for a detailed discussion see [Linde & Mukhanov 2006](#)). A general formula for $f_{\text{NL}}^{\text{local}}$, accounting for the possibility that the inflaton field decays into curvaton particles, is provided in [Sasaki et al. \(2006\)](#):

$$f_{\text{NL}}^{\text{local}} = (1 + \Delta_s^2) \frac{5}{4r_D} - \frac{5r_D}{6} - \frac{5}{3}, \quad (99)$$

where $\Delta_s^2 = \rho_{\text{curv. particles}}/\rho_{\text{curv. field}}$ measures the ratio of the energy density of curvaton particles to the energy density of the classical curvaton field ([Linde & Mukhanov 2006](#); [Sasaki et al. 2006](#)) and ρ_{curvaton} in the expression for r_D is given by $\rho_{\text{curvaton}} = \rho_{\text{curv. particles}} + \rho_{\text{curv. field}}$. Using the uniform prior $0 < r_D < 1$ and $0 < \Delta_s^2 < 10^2$ our measurements of $f_{\text{NL}}^{\text{local}}$ constrain $\Delta_s^2 \leq 8.5$ at 95% CL (T) and $\Delta_s^2 \leq 6.9$ ($T+E$).

11.3. Non-standard inflation models

Directional dependence motivated by gauge fields: In [Table 25](#) we constrained directionally-dependent bispectra (Eq. 23). This kind of NG is generated by inflationary models characterized by the presence of gauge fields. An actual realization of this type of scenario can be obtained with a coupling between the inflaton and the gauge field(s), via the kinetic term of the field(s), i.e., $\mathcal{L} = -I^2(\phi)F^2/4$. In this formula, F^2 represents the strength of the gauge field, while $I(\phi)$ is a function of the inflaton field with an appropriate time dependence (see, e.g., [Ratra 1992](#)). In this type of scenario, vector fields can be generated during inflation, and this in turn determines the excitation of $L = 0$ and 2 modes in the bispectrum, with $f_{\text{NL}}^L = X_L(|g_*|/0.1)(N_{k_3}/60)$, where $X_{L=0} = (80/3)$ and $X_{L=2} = -(10/6)$ ([Barnaby et al. 2012b](#); [Bartolo et al. 2013a](#); [Shiraishi et al. 2013a](#)). The parameter g_* , appearing in the equations above, represents the amplitude of a quadrupolar anisotropy in the power spectrum (see, e.g., [Ackerman et al. 2007](#)), while N defines the number of e-folds, before the end of inflation, when the relevant scales exit the horizon. It is thus clear that these models predict both a degree of statistical anisotropy in the power spectrum, and a potentially non-negligible bispectrum, as well as a direct relation between the two.

Starting from our SMICA constraints from $T(T+E)$ in [Table 25](#), marginalizing over a uniform prior $50 \leq N \leq 70$, and assuming uniform priors on $-1 \leq g_* \leq 1$, we obtain the limits $-0.050 < g_* < 0.050$ ($-0.040 < g_* < 0.040$), and $-0.31 < g_* < 0.31$ ($-0.29 < g_* < 0.29$), from the $L = 0$, $L = 2$ modes, respectively (95% CL) (considering g_* as scale-independent).

We note that these constraints refer to all models in which curvature perturbations are sourced by an $I^2(\phi)F^2$ term (see references in [Shiraishi et al. 2013a](#)). The constraints we obtain are consistent with the tighter (model-independent) limits obtained in [Planck Collaboration XX \(2016\)](#) for the case of a scale-independent g_* , from an analysis of quadrupolar anisotropies in the CMB power spectrum.

NG from gauge field production during axion inflation: We have also constrained the inverse decay NG of Eq. (25) arising typically in models where the inflaton field is a pseudoscalar axion that couples to a gauge field. Using the modal estimator we get the following constraints (removing ISW-lensing bias):

$$f_{\text{NL}}^{\text{inv.dec}} = 17 \pm 43 \quad 68\% \text{ CL}, \quad (100)$$

for temperature only; and

$$f_{\text{NL}}^{\text{inv.dec}} = 23 \pm 26 \quad 68\% \text{ CL}, \quad (101)$$

from temperature+polarization (see [Table 23](#)). The NG amplitude is given by $f_{\text{NL}}^{\text{inv.dec}} = f_3(\xi_*, 1, 1)\mathcal{P}_*^3 e^{6\pi\xi_*}/\mathcal{P}_\zeta^2(k_*)$, where $\mathcal{P}^{1/2} = H^2/(2\pi|\dot{\phi}|)$ is the power spectrum of vacuum-mode curvature perturbations (i.e., the power spectrum predicted without the coupling to gauge fields), $\mathcal{P}_\zeta^2(k_*)$ is the dimensionless scalar power spectrum of curvature perturbations (a star denoting evaluation at the pivot scale). The NG parameter is exponentially sensitive to the strength of the coupling between the axion and the gauge field. From Eq. (101) we limit the strength of the coupling to $\xi \leq 2.5$ (95% CL). The details together with constraints on the axion decay constant can be found in [Planck Collaboration XX \(2016\)](#), where an overview of various observational limits on axion (monodromy) models of inflation is presented. This limit is in agreement with the one that can be derived from tensor non-Gaussianities (see below).

Tensor NG and pseudoscalars: In inflationary scenarios associated with a pseudoscalar coupling to a U(1) gauge field, the tensor bispectrum generated via the gravitational interaction with the gauge field is expressed by Eq. (27), and the amplitude $f_{\text{NL}}^{\text{tens}}$ depends on the following: the coupling strength of the pseudoscalar field to the gauge field (ξ); a slow-roll parameter for the inflaton (ϵ); and the power spectrum of vacuum-mode curvature perturbations (\mathcal{P}). The expression is $f_{\text{NL}}^{\text{tens}} \approx 6.4 \times 10^{11} \mathcal{P}^3 \epsilon^3 e^{6\pi\xi}/\xi^9$ ([Cook & Sorbo 2013](#); [Shiraishi et al. 2013b](#)). The constraints on $f_{\text{NL}}^{\text{tens}}$ presented in [Table 26](#) can then be used to constrain the model parameters. Clearly there are strong degeneracies, but if we marginalize over a uniform prior $1.5 \times 10^{-9} \leq \mathcal{P} \leq 3 \times 10^{-9}$ and set $\epsilon = 0.01$, then assuming a uniform prior $0.1 \leq \xi \leq 7$, from the SMICA (T -only or $T+E$) limit, we obtain $\xi < 3.3$ (95% CL).

Warm inflation: We update the constraints on warm inflation models in the strongly dissipative regime, when dissipative effects are relevant. In this regime $f_{\text{NL}}^{\text{warm}} = -15 \ln(1 + r_d/14) - 5/2$ ([Moss & Xiong 2007](#)) with a large dissipation parameter $r_d = \Gamma/(3H)$ (where Γ is a friction term for the inflaton evolution describing the energy transfer from the inflaton field to radiation). The limit from the 2013 *Planck* release is $\log_{10} r_d \leq 2.6$ (95% CL) ([Planck Collaboration XXIV 2014](#)). Assuming a constant prior $0 \leq \log_{10} r_d \leq 4$, the new SMICA constraint $f_{\text{NL}}^{\text{warmS}} = -23 \pm 36$ at 68% CL from T ($f_{\text{NL}}^{\text{warmS}} = -14 \pm 23$ from $T+E$), see [Table 24](#), yields a limit on the dissipation parameter of $\log_{10} r_d \leq 3.3$ ($\log_{10} r_d \leq 2.5$) at 95% CL, with the $T+E$ constraints (in brackets), slightly improving the 2013 *Planck* limits. Values of $r_d \gtrsim 2.5$ (strongly-dissipative regime)

are still allowed; however, the *Planck* constraint puts the model in a regime where there might be an overproduction of gravitinos (see [Hall & Peiris 2008](#) and references therein). Unlike the strong dissipative regime, in the intermediate and weak dissipative regimes ($r_D \leq 1$) the NG level strongly depends on the microscopic parameters (T/H and r_D), giving rise to a new additional bispectrum shape (for details see [Bastero-Gil et al. 2014](#)).

11.4. Alternatives to inflation

Ekpyrotic/cyclic models have been proposed as an alternative to inflation (for a review, see [Lehners 2010](#)). Local NG is generated from the conversion of “intrinsic” non-Gaussian entropy perturbation modes into curvature fluctuations. Models based on a conversion taking place during the ekpyrotic phase (the so called “ekpyrotic conversion mechanism”) are already ruled out ([Koyama et al. 2007](#); [Planck Collaboration XXIV 2014](#)). Ekpyrotic models where “kinetic conversion” occurs after the ekpyrotic phase predict a local bispectrum with $f_{\text{NL}}^{\text{local}} = (3/2)\kappa_3 \sqrt{\epsilon} \pm 5$, where the sign depends on the details of the conversion process ([Lehners & Steinhardt 2008](#); [Lehners 2010](#); [Lehners & Steinhardt 2013](#)), where $\epsilon \simeq 50$ or greater are typical values. If we take $\epsilon \simeq 100$ and a uniform prior on $-5 < \kappa_3 < 5$ the constraints on $f_{\text{NL}}^{\text{local}}$ from T -only (see [Table 11](#)), yield $-0.91 < \kappa_3 < 0.58$ and $-0.25 < \kappa_3 < 1.2$ at 95 % CL, for the plus and minus sign in $f_{\text{NL}}^{\text{local}}$ respectively. From the $T+E$ $f_{\text{NL}}^{\text{local}}$ constraints ([Table 11](#)) we obtain $-0.94 < \kappa_3 < 0.38$ and $-0.27 < \kappa_3 < 1.0$ at 95 % CL, for the plus and minus sign in $f_{\text{NL}}^{\text{local}}$ respectively. If we consider $\epsilon \simeq 50$ we derive the following limits: $-1.3 < \kappa_3 < 0.81$ and $-0.35 < \kappa_3 < 1.8$ at 95 % CL from T -only; $-1.3 < \kappa_3 < 0.53$ and $-0.38 < \kappa_3 < 1.5$ at 95 % CL from $T+E$. Another variant of the ekpyrotic models has been investigated in ([Qiu et al. 2013](#); [Li 2013](#); [Fertig et al. 2014](#)), where the intrinsic NG is zero and NG is generated only by non-linearities in the conversion mechanism, reaching a value of $f_{\text{NL}}^{\text{local}} \simeq \pm 5$.

11.5. Inflationary interpretation of CMB trispectrum results

We briefly interpret the trispectrum constraints in an inflationary context. First, consider the case of single field inflation. The action for the Goldstone boson π is highly constrained by residual diffeomorphism invariance (see, e.g., [Smith et al. \(2015\)](#)). To lowest order in the derivative expansion, the most general action is:

$$S_\pi = \int d^4x \sqrt{-g} \left\{ -M_{\text{Pl}}^2 \dot{H} (\partial_\mu \pi)^2 + 2M_2^4 \left[\dot{\pi}^2 + \dot{\pi}^3 - \dot{\pi} \frac{(\partial_i \pi)^2}{a^2} + (\partial_\mu \pi)^2 (\partial_\nu \pi)^2 \right] - \frac{M_3^4}{3!} \left[8 \dot{\pi}^3 + 12 \dot{\pi}^2 (\partial_\mu \pi)^2 + \dots \right] + \frac{M_4^4}{4!} \left[16 \dot{\pi}^4 + 32 \dot{\pi}^3 (\partial_\mu \pi)^2 + \dots \right] + \dots \right\}, \quad (102)$$

where the parameter M_4 is related to the trispectrum by

$$g_{\text{NL}}^{\dot{\sigma}^4} = \frac{25}{288} \frac{M_4^4}{H^4} A_\zeta c_s^3. \quad (103)$$

The $g_{\text{NL}}^{\dot{\sigma}^4}$ constraint in [Eq. \(70\)](#) translates to the following parameter constraint in single field inflation:

$$-9.70 \times 10^{14} < \frac{M_4^4}{H^4 c_s^3} < 8.59 \times 10^{14} \quad (95 \text{ \% CL}). \quad (104)$$




Fig. 33. 68 % and 95 % confidence regions in the $(g_{\text{NL}}^{\dot{\sigma}^4}, g_{\text{NL}}^{(\partial\sigma)^4})$ plane, with the Lorentz invariant model in [Eq. \(107\)](#) shown as the dashed line.

This constraint is a factor 1.8 better than WMAP.

Turning now to multifield inflation, we consider an action of the form

$$S_\sigma = \int d^4x \sqrt{-g} \left[\frac{1}{2} (\partial_\mu \sigma)^2 + \frac{1}{\Lambda_1^4} \dot{\sigma}^4 + \frac{1}{\Lambda_2^4} \dot{\sigma}^2 (\partial_i \sigma)^2 + \frac{1}{\Lambda_3^4} (\partial_i \sigma)^2 (\partial_j \sigma)^2 \right], \quad (105)$$

where σ is a spectator field that acquires quantum fluctuations with power spectrum $P_\sigma(k) = H^2/(2k^3)$ and converts to adiabatic curvature via $\zeta = (2A_\zeta)^{1/2} H^{-1} \sigma$. The trispectrum in this model is

$$\begin{aligned} g_{\text{NL}}^{\dot{\sigma}^4} A_\zeta &= \frac{25}{768} \frac{H^4}{\Lambda_1^4}, \\ g_{\text{NL}}^{\dot{\sigma}^2 (\partial\sigma)^2} A_\zeta &= -\frac{325}{6912} \frac{H^4}{\Lambda_2^4}, \\ g_{\text{NL}}^{(\partial\sigma)^4} A_\zeta &= \frac{2575}{20736} \frac{H^4}{\Lambda_3^4}, \end{aligned} \quad (106)$$

so we can constrain its parameters by thresholding the χ^2 defined in [Eq. \(75\)](#). For example, if we consider the Lorentz invariant model

$$S = \int d^4x \sqrt{-g} \left[\frac{1}{2} (\partial_\mu \sigma)^2 + \frac{1}{\Lambda^4} (\partial_\mu \sigma)^2 (\partial_\nu \sigma)^2 \right], \quad (107)$$

so that the parameters Λ_i of the more general action in [Eq. \(105\)](#) are given by $\Lambda_1^4 = -2\Lambda_2^4 = \Lambda_3^4 = \Lambda^4$, then by thresholding at $\Delta\chi^2 = 4$ (as appropriate for one degree of freedom), we obtain the following constraint on the parameter Λ :

$$-0.26 < \frac{H^4}{\Lambda^4} < 0.20 \quad (95 \text{ \% CL}). \quad (108)$$

Constraints in other parameter spaces can also be obtained by thresholding the χ^2 defined in [Eq. \(75\)](#). For example we show 68 % and 95 % confidence regions in the $(g_{\text{NL}}^{\dot{\sigma}^4}, g_{\text{NL}}^{(\partial\sigma)^4})$ -plane, obtained by thresholding at $\chi^2 = 2.28$ and $\chi^2 = 5.99$ as appropriate for a χ^2 random variable with two degrees of freedom.

DBI Trispectrum: The trispectrum constraint on the shape δ^4 in Eq. (70) can also be used to obtain a lower bound on the DBI model sound speed. This is because, in the small sound speed limit (Chen et al. 2009; Arroja et al. 2009), the dominant contribution to the contact interaction trispectrum (Huang & Shiu 2006) has this shape. The corresponding non-linearity parameter is $g_{\text{NL}}^{\delta^4} = -25/(768 c_s^4)$. We follow the same procedure as described at the beginning of this section 11 and, assuming a uniform prior in the range $0 \leq c_s \leq 1/5$, we can derive a constraint on c_s as

$$c_s^{\text{DBI}} \geq 0.021, \quad 95 \% \text{ CL.} \quad (109)$$

This constraint is consistent with the ones derived from the bispectrum measurements (see Eqs. 84 and 85)) and it is only a factor of about three worse. Notice, however, that in this case we are ignoring the scalar exchange contribution, which is of the same order in c_s .

Curvaton trispectrum: For the simplest curvaton scenario, the trispectrum non-linearity parameter $g_{\text{NL}}^{\text{local}}$ prediction is (Sasaki et al. 2006)

$$g_{\text{NL}}^{\text{local}} = \frac{25}{54} \left(-\frac{9}{r_{\text{D}}} + \frac{1}{2} + 10r_{\text{D}} + 3r_{\text{D}}^2 \right). \quad (110)$$

Following the procedure described at the beginning of Sect. 11, we use the observational constraint obtained in Sect. 9 (Eq. 70), and the same prior ($0 < r_{\text{D}} < 1$) as in 11.2, to obtain a lower bound on the curvaton decay fraction as

$$r_{\text{D}} \geq 0.05 \quad 95 \% \text{ CL.} \quad (111)$$

This limit is consistent with the previous ones derived using the bispectrum measurements and it is a factor of about 3 to 4 worse.

12. Conclusions

In this paper we have presented the constraints on primordial NG using the full *Planck* mission data. The results have improved compared to the *Planck* 2013 release (Planck Collaboration XXIV 2014) as a consequence of including data from the full mission and taking advantage of *Planck*'s polarization capability — the first time that maps of the CMB polarization anisotropies have been used to constrain primordial NG.

Using temperature data alone, the constraints on the local, equilateral, and orthogonal bispectrum templates are $f_{\text{NL}}^{\text{local}} = 2.5 \pm 5.7$, $f_{\text{NL}}^{\text{equil}} = -16 \pm 70$, and $f_{\text{NL}}^{\text{ortho}} = -34 \pm 33$. Moving from the nominal *Planck* 2013 data to the full mission data yielded modest improvements of up to 15% (for the orthogonal shape). After the inclusion of full mission polarization data, our final constraints become $f_{\text{NL}}^{\text{local}} = 0.8 \pm 5.0$, $f_{\text{NL}}^{\text{equil}} = -4 \pm 43$, and $f_{\text{NL}}^{\text{ortho}} = -26 \pm 21$, which represents a substantial step forward relative to *Planck* 2013, with error bars shrinking by 14% (local), 43% (equilateral), and 46% (orthogonal). These improved limits on the standard shapes enhance our understanding of different inflationary models that can potentially lead to subtle effects beyond the simplest models of inflation.

The reason that the polarization data provide such complementary constraints on primordial curvature perturbations is due to the phase shift of the CMB polarization transfer functions compared to the temperature transfer functions. So, despite the relatively much lower signal-to-noise in the polarization maps, their inclusion leads to appreciable improvements on limits on

NG parameters. Nevertheless, the full characterization of the noise properties in the polarized maps is still ongoing. In spite of the extensive testing and cross-checks validating the combined temperature and polarization results, we therefore conservatively recommend that *all* results that include polarization information, not just the polarization-only results, be taken as preliminary at this stage.

The complementary nature of the polarization information also represents an important cross-check on the analysis. The *Planck* results based on polarization alone are statistically consistent with the results based on temperature alone, with a precision comparable to that achievable in an optimal analysis of the WMAP 3-year temperature maps (Spergel et al. 2007; Creminelli et al. 2007; Yadav & Wandelt 2008).

Our analysis was subject to an extensive validation exercise. In addition to extensive simulation tests, including for the first time a detailed test of the impact of time-domain de-glitching, our results are supported by tests for robustness under change of estimator implementations (KSW, binned bispectrum, and two modal estimators), and variations in sky coverage as well as upper and lower multipole cutoffs. We also tested for possible directional dependence using a needlet estimator. These tests form the basis of our selection of SMICA as the main foreground cleaning method for our headline results.

The *Planck* 2015 analysis presented here provides constraints on a greatly extended range of template families. These extensions include a tenfold increase in the range of frequencies covered in feature models, giving rise to linearly oscillating bispectra, generalized shapes for oscillating models including for logarithmic oscillations, tests for deviations from the Bunch-Davies vacuum, models of equilateral type in the context of the effective field theory of inflation, and direction-dependent primordial NG. Beyond purely scalar mode templates we also tested for parity-violating tensor NG.

Using the full mission data with polarization, we have investigated the hints of NG reported in the *Planck* 2013 analysis of oscillatory features. While no individual feature or resonance model rises above our detection threshold of 3σ (after inclusion of the look-elsewhere effect), the results of integrated (multi-peak) statistical tests indicate that continued investigation of oscillatory and non-scaling models is warranted.

In addition to searches for specific NG templates, we present model-independent reconstructions of the temperature and polarization bispectra using the modal and binned bispectrum approaches. These full mission reconstructions can achieve twice the resolution of the *Planck* 2013 results, demonstrating excellent consistency in temperature, and good agreement with the WMAP9 reconstruction in regions where this earlier data set is signal-dominated.

The inclusion of polarization information leads to significantly improved constraints on NG in primordial isocurvature perturbations, providing complementary information to 2-point function results for models where the NG in isocurvature components is more easily detectable than its contribution to the power spectrum.

A significant addition to this year's analysis is the inclusion of detailed trispectrum results due to cubic NG. The local trispectrum is constrained by *Planck* temperature data to be $g_{\text{NL}}^{\text{local}} = (-9.0 \pm 7.7) \times 10^4$ and the other two shapes were also found to be consistent with Gaussianity. Both 3-point and 4-point constraints are consistent with the improved (though still suboptimal) constraints from Minkowski functionals, a very different estimation framework. This concordance adds confidence in our results.

We have discussed the implications of our results on the physics of the early Universe, showing that the n -point functions for $n > 2$ provide a significant window onto the primordial Universe beyond the power spectrum, constraining general-single field, multifield, and non-standard inflation models, as well as alternatives to inflation. Using bispectrum and trispectrum limits we updated results on the parameter space of the inflationary models (and alternatives) already tested in 2013, and constrained the parameter space of other well-motivated inflationary models (e.g., Galileon-like models of inflation, and models where axion/pseudoscalar fields are present during inflation).

The global picture that emerges is one of consistency with the premises of the Λ CDM cosmology, namely that the structure we observe today is the consequence of passive evolution of adiabatic, Gaussian, primordial seed perturbations. Nevertheless, NG at some level is expected in all inflationary models, and hence we should strive to find means to reduce errors on f_{NL} still further.

Acknowledgements. The Planck Collaboration acknowledges the support of: ESA; CNES and CNRS/INSU-IN2P3-INP (France); ASI, CNR, and INAF (Italy); NASA and DoE (USA); STFC and UKSA (UK); CSIC, MINECO, JA, and RES (Spain); Tekes, AoF, and CSC (Finland); DLR and MPG (Germany); CSA (Canada); DTU Space (Denmark); SER/SSO (Switzerland); RCN (Norway); SFI (Ireland); FCT/MCTES (Portugal); ERC and PRACE (EU). A description of the Planck Collaboration and a list of its members, indicating which technical or scientific activities they have been involved in, can be found at <http://www.cosmos.esa.int/web/planck/planck-collaboration>. Some of the results in this paper have been derived using the HEALPIX package. Part of this work was undertaken on the STFC COSMOS@DiRAC HPC Facility at the University of Cambridge, funded by UK BIS NEI grants. We gratefully acknowledge IN2P3 Computer Center (<http://cc.in2p3.fr>) for providing a significant amount of the computing resources and services needed for the analysis with the binned bispectrum estimator. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. We also acknowledge the IAP magique3 computer facilities. Some computations were performed on the GPC cluster at the SciNet HPC Consortium. SciNet is funded by the Canada Foundation for Innovation under the auspices of Compute Canada, the Government of Ontario, and the University of Toronto.

References

- Achúcarro, A., Gong, J.-O., Hardeman, S., Palma, G. A., & Patil, S. P., Features of heavy physics in the CMB power spectrum. 2011, *J. Cosmology Astropart. Phys.*, 1, 30
- Achúcarro, A., Gong, J.-O., Palma, G. A., & Patil, S. P., KachruCorrelating features in the primordial spectra. 2013, *Phys.Rev.*, D87, 121301, [arXiv:1211.5619](https://arxiv.org/abs/1211.5619)
- Ackerman, L., Carroll, S. M., & Wise, M. B., Imprints of a primordial preferred direction on the microwave background. 2007, *Phys. Rev. D*, 75, 083502
- Acquaviva, V., Bartolo, N., Matarrese, S., & Riotto, A., Gauge-invariant second-order perturbations and non-Gaussianity from inflation. 2003, *Nuclear Physics B*, 667, 119
- Adshead, P., Dvorkin, C., Hu, W., & Lim, E. A., Non-Gaussianity from step features in the inflationary potential. 2012, *Phys. Rev. D*, 85, 023531
- Adshead, P., Hu, W., Dvorkin, C., & Peiris, H. V., Fast Computation of Bispectrum Features with Generalized Slow Roll. 2011, *Phys. Rev. D*, 84, 043519
- Agullo, I. & Parker, L., Stimulated creation of quanta during inflation and the observable universe. 2011, *General Relativity and Gravitation*, 43, 2541
- Alishahiha, M., Silverstein, E., & Tong, D., DBI in the sky: Non-Gaussianity from inflation with a speed limit. 2004, *Phys. Rev. D*, 70, 123505
- Argüeso, F., González-Nuevo, J., & Toffolatti, L., Contributions of Point Extragalactic Sources to the Cosmic Microwave Background Bispectrum. 2003, *ApJ*, 598, 86, [arXiv:astro-ph/0307148](https://arxiv.org/abs/astro-ph/0307148)
- Arkani-Hamed, N., Creminelli, P., Mukohyama, S., & Zaldarriaga, M., Ghost inflation. 2004, *J. Cosmology Astropart. Phys.*, 4, 1
- Arroja, F., Bartolo, N., Dimastrogiovanni, E., & Fasiello, M., On the Trispectrum of Galileon Inflation. 2013, *JCAP*, 1311, 005, [arXiv:1307.5371](https://arxiv.org/abs/1307.5371)
- Arroja, F., Mizuno, S., Koyama, K., & Tanaka, T., Full trispectrum in single field DBI inflation. 2009, *Phys. Rev. D*, 80, 043527
- Ashoorioon, A. & Shiu, G., A Note on Calm Excited States of Inflation. 2011, *JCAP*, 1103, 025, [arXiv:1012.3392](https://arxiv.org/abs/1012.3392)
- Babich, D., Creminelli, P., & Zaldarriaga, M., The shape of non-Gaussianities. 2004, *J. Cosmology Astropart. Phys.*, 8, 9
- Bardeen, J. M., Gauge-invariant cosmological perturbations. 1980, *Phys. Rev. D*, 22, 1882
- Barnaby, N. & Cline, J. M., Non-Gaussian and nonscale-invariant perturbations from tachyonic preheating in hybrid inflation. 2006, *Phys. Rev. D*, 73, 106012
- Barnaby, N., Moxon, J., Namba, R., et al., Gravity waves and non-Gaussian features from particle production in a sector gravitationally coupled to the inflaton. 2012a, *Phys. Rev. D*, 86, 103508, [arXiv:1206.6117](https://arxiv.org/abs/1206.6117)
- Barnaby, N., Namba, R., & Peloso, M., Phenomenology of a pseudo-scalar inflaton: naturally large nongaussianity. 2011, *J. Cosmology Astropart. Phys.*, 4, 9
- Barnaby, N., Namba, R., & Peloso, M., Observable non-Gaussianity from gauge field production in slow roll inflation, and a challenging connection with magnetogenesis. 2012b, *Phys. Rev. D*, 85, 123523
- Barnaby, N., Pajer, E., & Peloso, M., Gauge field production in axion inflation: Consequences for monodromy, non-Gaussianity in the CMB, and gravitational waves at interferometers. 2012c, *Phys. Rev. D*, 85, 023525
- Barnaby, N. & Peloso, M., Large Non-Gaussianity in Axion Inflation. 2011, *Physical Review Letters*, 106, 181301
- Bartolo, N., Dimastrogiovanni, E., & Fasiello, M., The Trispectrum in the Effective Theory of Inflation with Galilean symmetry. 2013, *JCAP*, 1309, 037, [arXiv:1305.0812](https://arxiv.org/abs/1305.0812)
- Bartolo, N., Fasiello, M., Matarrese, S., & Riotto, A., Large non-Gaussianities in the effective field theory approach to single-field inflation: the bispectrum. 2010a, *J. Cosmology Astropart. Phys.*, 8, 8
- Bartolo, N., Fasiello, M., Matarrese, S., & Riotto, A., Large non-Gaussianities in the effective field theory approach to single-field inflation: the trispectrum. 2010b, *J. Cosmology Astropart. Phys.*, 9, 35
- Bartolo, N., Komatsu, E., Matarrese, S., & Riotto, A., Non-Gaussianity from inflation: theory and observations. 2004a, *Phys. Rep.*, 402, 103
- Bartolo, N., Matarrese, S., Peloso, M., & Ricciardone, A., Anisotropic power spectrum and bispectrum in the $f(\phi)F^2$ mechanism. 2013a, *Phys. Rev. D*, 87, 023504
- Bartolo, N., Matarrese, S., Peloso, M., & Ricciardone, A., Anisotropy in solid inflation. 2013b, *J. Cosmology Astropart. Phys.*, 8, 22, [arXiv:1306.4160](https://arxiv.org/abs/1306.4160)
- Bartolo, N., Matarrese, S., & Riotto, A., Non-Gaussianity from inflation. 2002, *Phys. Rev. D*, 65, 103505
- Bartolo, N., Matarrese, S., & Riotto, A., Gauge-Invariant Temperature Anisotropies and Primordial Non-Gaussianity. 2004b, *Physical Review Letters*, 93, 231301
- Bartolo, N., Matarrese, S., & Riotto, A., Non-Gaussianity in the curvaton scenario. 2004c, *Phys. Rev. D*, 69, 043503
- Bartolo, N., Peloso, M., Ricciardone, A., & Unal, C., The expected anisotropy in solid inflation. 2014, *J. Cosmology Astropart. Phys.*, 11, 9, [arXiv:1407.8053](https://arxiv.org/abs/1407.8053)
- Bastero-Gil, M., Berera, A., Moss, I. G., & Ramos, R. O., Theory of non-Gaussianity in warm inflation. 2014, *J. Cosmology Astropart. Phys.*, 12, 8, [arXiv:1408.4391](https://arxiv.org/abs/1408.4391)
- Battefeld, T., Niemeyer, J. C., & Vlaykov, D., Probing two-field open inflation by resonant signals in correlation functions. 2013, *J. Cosmology Astropart. Phys.*, 5, 6, [arXiv:1302.3877](https://arxiv.org/abs/1302.3877)
- Baumann, D. & Green, D., Signature of supersymmetry from the early universe. 2012, *Phys. Rev. D*, 85, 103520
- Bean, R., Chen, X., Peiris, H., & Xu, J., Comparing infrared Dirac-Born-Infeld brane inflation to observations. 2008, *Phys. Rev. D*, 77, 023527
- Behbahani, S. R., Dymarsky, A., Mirbabayi, M., & Senatore, L., (Small) Resonant non-Gaussianities: Signatures of a Discrete Shift Symmetry in the Effective Field Theory of Inflation. 2012, *JCAP*, 1212, 036, [arXiv:1111.3373](https://arxiv.org/abs/1111.3373)
- Bennett, C. L., Larson, D., Weiland, J. L., et al., Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results. 2013, *ApJS*, 208, 20
- Bernardeau, F., Kofman, L., & Uzan, J.-P., Modulated fluctuations from hybrid inflation. 2004, *Phys. Rev. D*, 70, 083004
- Bernardeau, F. & Uzan, J.-P., Non-Gaussianity in multifield inflation. 2002, *Phys. Rev. D*, 66, 103506
- Bond, J. R., Frolov, A. V., Huang, Z., & Kofman, L., Non-Gaussian Curvature Spikes from Chaotic Billiards in Inflation Preheating. 2009, *Physical Review Letters*, 103, 071301
- Brown, I. & Crittenden, R., Non-Gaussianity from cosmic magnetic fields. 2005, *Phys. Rev. D*, 72, 063002, [arXiv:astro-ph/0506570](https://arxiv.org/abs/astro-ph/0506570)
- Bucher, M., Moodley, K., & Turok, N., General primordial cosmic perturbation. 2000, *Phys. Rev. D*, 62, 083508, [arXiv:astro-ph/9904231](https://arxiv.org/abs/astro-ph/9904231)
- Bucher, M., Racine, B., & van Tent, B., The binned bispectrum estimator: template-based and non-parametric CMB non-Gaussianity searches. 2016, *J. Cosmology Astropart. Phys.*, 5, 055, [arXiv:1509.08107](https://arxiv.org/abs/1509.08107)

- Bucher, M., van Tent, B., & Carvalho, C. S., Detecting bispectral acoustic oscillations from inflation using a new flexible estimator. 2010, *MNRAS*, 407, 2193
- Burrage, C., de Rham, C., Seery, D., & Tolley, A. J., Galileon inflation. 2011, *J. Cosmology Astropart. Phys.*, 1, 14, [arXiv:1009.2497](#)
- Byrnes, C. T. & Choi, K.-Y., Review of Local Non-Gaussianity from Multifield Inflation. 2010, *Advances in Astronomy*, 2010
- Catalano, A., Ade, P., Atik, Y., et al., Impact of particles on the Planck HFI detectors: Ground-based measurements and physical interpretation. 2014, *ArXiv e-prints*, [arXiv:1403.6592](#)
- Chambers, A. & Rajantie, A., Lattice Calculation of Non-Gaussian Density Perturbations from the Massless Preheating Inflationary Model. 2008, *Physical Review Letters*, 100, 041302
- Chen, X., Inflation from warped space. 2005a, *Journal of High Energy Physics*, 8, 45
- Chen, X., Multithroat brane inflation. 2005b, *Phys. Rev. D*, 71, 063506
- Chen, X., Folded resonant non-Gaussianity in general single field inflation. 2010a, *J. Cosmology Astropart. Phys.*, 12, 3
- Chen, X., Primordial Non-Gaussianities from Inflation Models. 2010b, *Advances in Astronomy*, 2010
- Chen, X., Primordial Features as Evidence for Inflation. 2011, *ArXiv e-prints*, [arXiv:1104.1323](#)
- Chen, X., Easther, R., & Lim, E. A., Large non-Gaussianities in single-field inflation. 2007a, *J. Cosmology Astropart. Phys.*, 6, 23
- Chen, X., Easther, R., & Lim, E. A., Generation and characterization of large non-Gaussianities in single field inflation. 2008, *J. Cosmology Astropart. Phys.*, 4, 10
- Chen, X., Hu, B., Huang, M.-x., Shiu, G., & Wang, Y., Large primordial trispectra in general single field inflation. 2009, *J. Cosmology Astropart. Phys.*, 8, 8
- Chen, X., Huang, M.-x., Kachru, S., & Shiu, G., Observational signatures and non-Gaussianities of general single-field inflation. 2007b, *J. Cosmology Astropart. Phys.*, 1, 2
- Chen, X., Namjoo, M. H., & Wang, Y., Models of the Primordial Standard Clock. 2015, *J. Cosmology Astropart. Phys.*, 2, 27, [arXiv:1411.2349](#)
- Chen, X. & Wang, Y., Quasi-single field inflation and non-Gaussianities. 2010, *J. Cosmology Astropart. Phys.*, 4, 27
- Cheung, C., Fitzpatrick, A. L., Kaplan, J., Senatore, L., & Creminelli, P., The effective field theory of inflation. 2008, *Journal of High Energy Physics*, 3, 14
- Cicoli, M., Tasinato, G., Zavala, I., Burgess, C. P., & Quevedo, F., Modulated reheating and large non-gaussianity in string cosmology. 2012, *J. Cosmology Astropart. Phys.*, 5, 39
- Cook, J. L. & Sorbo, L., An inflationary model with small scalar and large tensor non-gaussianities. 2013, *J. Cosmology Astropart. Phys.*, 11, 47, [arXiv:1307.7077](#)
- Cooray, A. & Melchiorri, A., Searching for integrated Sachs Wolfe effect beyond temperature anisotropies: CMB E-mode polarization galaxy cross-correlation. 2006, *Journal of Cosmology and Astro-Particle Physics*, 1, 18, [arXiv:astro-ph/0511054](#)
- Crawford, T. M., Schaffer, K. K., Bhattacharya, S., et al., A Measurement of the Secondary-CMB and Millimeter-wave-foreground Bispectrum using 800 deg² of South Pole Telescope Data. 2014, *ApJ*, 784, 143, [arXiv:1303.3535](#)
- Creminelli, P., D'Amico, G., Musso, M., Noreña, J., & Trincherini, E., Galilean symmetry in the effective theory of inflation: new shapes of non-Gaussianity. 2011, *J. Cosmology Astropart. Phys.*, 2, 6, [arXiv:1011.3004](#)
- Creminelli, P., Nicolis, A., Senatore, L., Tegmark, M., & Zaldarriaga, M., Limits on non-Gaussianities from WMAP data. 2006, *J. Cosmology Astropart. Phys.*, 5, 4
- Creminelli, P., Senatore, L., Zaldarriaga, M., & Tegmark, M., Limits on f_{NL} parameters from Wilkinson Microwave Anisotropy Probe three-year data. 2007, *J. Cosmology Astropart. Phys.*, 3, 5, [arXiv:astro-ph/0610600](#)
- Curto, A., Macias-Perez, J. F., Martínez-González, E., et al., Constraints on the non-linear coupling parameter f_{nl} with Archeops data. 2008, *A&A*, 486, 383
- Curto, A., Tucci, M., Kunz, M., & Martinez-Gonzalez, E., The CIB-lensing bispectrum: impact on primordial non-Gaussianity and detectability for the Planck mission. 2014, *ArXiv e-prints*, [arXiv:1405.7029](#)
- de Rham, C. & Gabadadze, G., Generalization of the Fierz-Pauli action. 2010a, *Phys. Rev. D*, 82, 044020, [arXiv:1007.0443](#)
- de Rham, C. & Gabadadze, G., Selftuned massive spin-2. 2010b, *Physics Letters B*, 693, 334, [arXiv:1006.4367](#)
- Delabrouille, J., Betoule, M., Melin, J.-B., et al., The pre-launch Planck Sky Model: a model of sky emission at submillimetre to centimetre wavelengths. 2013, *A&A*, 553, A96, [arXiv:1207.3675](#)
- Desjacques, V. & Seljak, U., Signature of primordial non-Gaussianity of ϕ^3 -type in the mass function and bias of dark matter haloes. 2010, *Phys.Rev.*, D81, 023006, [arXiv:0907.2257](#)
- Donzelli, S., Hansen, F. K., Liguori, M., Marinucci, D., & Matarrese, S., On the Linear Term Correction for Needlet/Wavelet Non-Gaussianity Estimators. 2012, *ApJ*, 755, 19, [arXiv:1202.1478](#)
- Ducout, A., Bouchet, F. R., Colombi, S., Pogosyan, D., & Prunet, S., Non-Gaussianity and Minkowski functionals: forecasts for Planck. 2013, *MNRAS*, 429, 2104
- Dvali, G., Gruzinov, A., & Zaldarriaga, M., Cosmological perturbations from inhomogeneous reheating, freeze-out, and mass domination. 2004a, *Phys. Rev. D*, 69, 083505
- Dvali, G., Gruzinov, A., & Zaldarriaga, M., New mechanism for generating density perturbations from inflation. 2004b, *Phys. Rev. D*, 69, 023505
- Elsner, F. & Wandelt, B. D., Improved Simulation of Non-Gaussian Temperature and Polarization Cosmic Microwave Background Maps. 2009, *ApJS*, 184, 264
- Elsner, F. & Wandelt, B. D., Fast Wiener filtering of CMB maps. 2012, *ArXiv e-prints*, [arXiv:1211.0585](#)
- Endlich, S., Horn, B., Nicolis, A., & Wang, J., Squeezed limit of the solid inflation three-point function. 2014, *Phys. Rev. D*, 90, 063506, [arXiv:1307.8114](#)
- Endlich, S., Nicolis, A., & Wang, J., Solid inflation. 2013, *J. Cosmology Astropart. Phys.*, 10, 11
- Enqvist, K., Jokinen, A., Mazumdar, A., Multamäki, T., & Vähiköinen, A., Non-Gaussianity from preheating. 2005, *Physical Review Letters*, 94, 161301
- Enqvist, K. & Sloth, M. S., Adiabatic CMB perturbations in pre-Big-Bang string cosmology. 2002, *Nuclear Physics B*, 626, 395
- Eriksen, H. K., Novikov, D. I., Lilje, P. B., Banday, A. J., & Górski, K. M., Testing for Non-Gaussianity in the Wilkinson Microwave Anisotropy Probe Data: Minkowski Functionals and the Length of the Skeleton. 2004, *ApJ*, 612, 64
- Falk, T., Rangarajan, R., & Srednicki, M., The angular dependence of the three-point correlation function of the cosmic microwave background radiation as predicted by inflationary cosmologies. 1993, *ApJ*, 403, L1
- Fantaye, Y., Marinucci, D., Hansen, F., & Maino, D., Applications of the Gaussian kinematic formula to CMB data analysis. 2015, *Phys. Rev. D*, 91, 063501
- Feng, C., Cooray, A., Smidt, J., et al., Planck Trispectrum Constraints on Primordial Non-Gaussianity at Cubic Order. 2015, [arXiv:1502.00585](#)
- Fergusson, J., Efficient optimal non-Gaussian CMB estimators with polarisation. 2014, *Phys.Rev.*, D90, 043533, [arXiv:1403.7949](#)
- Fergusson, J., Gruetjen, H., Shellard, E., & Liguori, M., Combining power spectrum and bispectrum measurements to detect oscillatory features. 2015, *Phys.Rev.*, D91, 023502, [arXiv:1410.5114](#)
- Fergusson, J., Gruetjen, H., Shellard, E., & Wallisch, B., Polyspectra searches for sharp oscillatory features in cosmic microwave sky data. 2014, [arXiv:1412.6152](#)
- Fergusson, J. R., Liguori, M., & Shellard, E. P. S., General CMB and primordial bispectrum estimation: Mode expansion, map making, and measures of F_{NL} . 2010a, *Phys. Rev. D*, 82, 023502
- Fergusson, J. R., Liguori, M., & Shellard, E. P. S., The CMB bispectrum. 2012, *J. Cosmology Astropart. Phys.*, 12, 32, [arXiv:1006.1642](#)
- Fergusson, J. R., Regan, D. M., & Shellard, E. P. S., Optimal Trispectrum Estimators and WMAP Constraints. 2010b, *ArXiv e-prints*, [arXiv:1012.6039](#)
- Ferreira, R. Z. & Sloth, M. S., Universal Constraints on Axions from Inflation. 2014, *ArXiv e-prints*, [arXiv:1409.5799](#)
- Fertig, A., Lehnert, J.-L., & Mallwitz, E., Ekpyrotic Perturbations With Small Non-Gaussian Corrections. 2014, *Phys.Rev.*, D89, 103537, [arXiv:1310.8133](#)
- Flauger, R., McAllister, L., Pajer, E., Westphal, A., & Xu, G., Oscillations in the CMB from axion monodromy inflation. 2010, *J. Cosmology Astropart. Phys.*, 6, 9
- Flauger, R. & Pajer, E., Resonant non-gaussianity. 2011, *J. Cosmology Astropart. Phys.*, 1, 17
- Gangui, A., Lucchin, F., Matarrese, S., & Mollerach, S., The three-point correlation function of the cosmic microwave background in inflationary models. 1994, *ApJ*, 430, 447
- Gangui, A. & Martin, J., Cosmic microwave background bispectrum and slow-roll inflation. 2000, *MNRAS*, 313, 323
- Gao, X., Kobayashi, T., Yamaguchi, M., & Yokoyama, J., Primordial Non-Gaussianities of Gravitational Waves in the Most General Single-Field Inflation Model with Second-Order Field Equations. 2011, *Physical Review Letters*, 107, 211301, [arXiv:1108.3513](#)
- Giannantonio, T., Ross, A. J., Percival, W. J., et al., Improved Primordial Non-Gaussianity Constraints from Measurements of Galaxy Clustering and the Integrated Sachs-Wolfe Effect. 2014, *Phys.Rev.*, D89, 023511, [arXiv:1303.1349](#)
- González-Nuevo, J., Toffolatti, L., & Argüeso, F., Predictions of the Angular Power Spectrum of Clustered Extragalactic Point Sources at Cosmic Microwave Background Frequencies from Flat and All-Sky Two-dimensional Simulations. 2005, *ApJ*, 621, 1

- Górski, K. M., Hivon, E., Banday, A. J., et al., HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. 2005, *ApJ*, 622, 759, [arXiv:astro-ph/0409513](#)
- Hall, L. M. H. & Peiris, H. V., Cosmological constraints on dissipative models of inflation. 2008, *J. Cosmology Astropart. Phys.*, 1, 27
- Hannestad, S., Haugbolle, T., Jarnhus, P. R., & Sloth, M. S., Non-Gaussianity from Axion Monodromy Inflation. 2010, *JCAP*, 1006, 001, [arXiv:0912.3527](#)
- Hanson, D. & Lewis, A., Estimators for CMB statistical anisotropy. 2009, *Phys. Rev. D*, 80, 063004
- Heavens, A. F., Estimating non-Gaussianity in the microwave background. 1998, *MNRAS*, 299, 805
- Hikage, C., Kawasaki, M., Sekiguchi, T., & Takahashi, T., CMB constraint on non-Gaussianity in isocurvature perturbations. 2013a, *J. Cosmology Astropart. Phys.*, 7, 7, [arXiv:1211.1095](#)
- Hikage, C., Kawasaki, M., Sekiguchi, T., & Takahashi, T., Extended analysis of CMB constraints on non-gaussianity in isocurvature perturbations. 2013b, *J. Cosmology Astropart. Phys.*, 3, 20, [arXiv:1212.6001](#)
- Hikage, C., Komatsu, E., & Matsubara, T., Primordial Non-Gaussianity and Analytical Formula for Minkowski Functionals of the Cosmic Microwave Background and Large-Scale Structure. 2006, *ApJ*, 653, 11
- Hikage, C., Koyama, K., Matsubara, T., Takahashi, T., & Yamaguchi, M., Limits on isocurvature perturbations from non-Gaussianity in WMAP temperature anisotropy. 2009, *MNRAS*, 398, 2188, [arXiv:0812.3500](#)
- Hikage, C. & Matsubara, T., Limits on second-order non-Gaussianity from Minkowski functionals of WMAP 7-year data. 2012, *MNRAS*, 425, 2187
- Hikage, C., Matsubara, T., Coles, P., et al., Limits on primordial non-Gaussianity from Minkowski Functionals of the WMAP temperature anisotropies. 2008, *MNRAS*, 389, 1439
- Holman, R. & Tolley, A. J., Enhanced Non-Gaussianity from Excited Initial States. 2008, *JCAP*, 0805, 001, [arXiv:0710.1302](#)
- Horndeski, G. W., Second-Order Scalar-Tensor Field Equations in a Four-Dimensional Space. 1974, *International Journal of Theoretical Physics*, 10, 363
- Hu, W., Weak lensing of the CMB: A harmonic approach. 2000, *Phys. Rev. D*, 62, 043007, [arXiv:astro-ph/0001303](#)
- Huang, M.-X. & Shiu, G., Inflationary trispectrum for models with large non-Gaussianities. 2006, *Phys. Rev. D*, 74, 121301, [arXiv:hep-th/0610235](#)
- Kamionkowski, M. & Souradeep, T., Odd-parity cosmic microwave background bispectrum. 2011, *Phys. Rev. D*, 83, 027301, [arXiv:1010.4304](#)
- Kawakami, E., Kawasaki, M., Miyamoto, K., Nakayama, K., & Sekiguchi, T., Non-Gaussian isocurvature perturbations in dark radiation. 2012, *J. Cosmology Astropart. Phys.*, 7, 37, [arXiv:1202.4890](#)
- Kawasaki, M., Nakayama, K., Sekiguchi, T., Suyama, T., & Takahashi, F., Non-Gaussianity from isocurvature perturbations. 2008, *J. Cosmology Astropart. Phys.*, 11, 19, [arXiv:0808.0009](#)
- Kawasaki, M., Nakayama, K., Sekiguchi, T., Suyama, T., & Takahashi, F., A general analysis of non-gaussianity from isocurvature perturbations. 2009, *J. Cosmology Astropart. Phys.*, 1, 42, [arXiv:0810.0208](#)
- Kobayashi, T., Yamaguchi, M., & Yokoyama, J., Inflation Driven by the Galileon Field. 2010, *Physical Review Letters*, 105, 231302, [arXiv:1008.0603](#)
- Kofman, L., Probing String Theory with Modulated Cosmological Fluctuations. 2003, *ArXiv Astrophysics e-prints*, [arXiv:astro-ph/0303614](#)
- Kolb, E. W., Riotto, A., & Vallinotto, A., Non-Gaussianity from broken symmetries. 2006, *Phys. Rev. D*, 73, 023522
- Komatsu, E., The Pursuit of Non-Gaussian Fluctuations in the Cosmic Microwave Background. 2002, *ArXiv Astrophysics e-prints*, [arXiv:astro-ph/0206039](#)
- Komatsu, E., Hunting for primordial non-Gaussianity in the cosmic microwave background. 2010, *Classical and Quantum Gravity*, 27, 124010
- Komatsu, E. & Spergel, D. N., Acoustic signatures in the primary microwave background bispectrum. 2001, *Phys. Rev. D*, 63, 063002
- Komatsu, E., Spergel, D. N., & Wandelt, B. D., Measuring Primordial Non-Gaussianity in the Cosmic Microwave Background. 2005, *ApJ*, 634, 14
- Koyama, K., Mizuno, S., Vernizzi, F., & Wands, D., Non-Gaussianities from ekpyrotic collapse with multiple fields. 2007, *JCAP*, 0711, 024, [arXiv:0708.4321](#)
- Lacasa, F. & Aghanim, N., Optimal estimator for the amplitude of the bispectrum from infrared clustered sources. 2014, *A&A*, 569, A51
- Lacasa, F., Aghanim, N., Kunz, M., & Frommert, M., Characterization of the non-Gaussianity of radio and IR point sources at CMB frequencies. 2012, *MNRAS*, 421, 1982
- Lacasa, F., Pénin, A., & Aghanim, N., Non-Gaussianity of the cosmic infrared background anisotropies - I. Diagrammatic formalism and application to the angular bispectrum. 2014, *MNRAS*, 439, 123, [arXiv:1312.1251](#)
- Lagache, G., Puget, J.-L., & Dole, H., Dusty Infrared Galaxies: Sources of the Cosmic Infrared Background. 2005, *ARA&A*, 43, 727, [arXiv:astro-ph/0507298](#)
- Lan, X. & Marinucci, D., The needlets bispectrum. 2008, *Electronic Journal of Statistics*, 2, 332, [arXiv:0802.4020](#)
- Langlois, D. & Lepidi, A., General treatment of isocurvature perturbations and non-Gaussianities. 2011, *J. Cosmology Astropart. Phys.*, 1, 8, [arXiv:1007.5498](#)
- Langlois, D. & van Tent, B., Hunting for isocurvature modes in the cosmic microwave background non-Gaussianities. 2011, *Classical and Quantum Gravity*, 28, 222001, [arXiv:1104.2567](#)
- Langlois, D. & van Tent, B., Isocurvature modes in the CMB bispectrum. 2012, *J. Cosmology Astropart. Phys.*, 7, 40, [arXiv:1204.5042](#)
- Langlois, D., Vernizzi, F., & Wands, D., Non-linear isocurvature perturbations and non-Gaussianities. 2008, *J. Cosmology Astropart. Phys.*, 12, 4, [arXiv:0809.4646](#)
- Lehners, J.-L., Ekpyrotic Nongaussianity: A Review. 2010, *Advances in Astronomy*, 2010
- Lehners, J.-L. & Steinhardt, P. J., Non-Gaussian density fluctuations from entropically generated curvature perturbations in Ekpyrotic models. 2008, *Phys.Rev.*, D77, 063533, [arXiv:0712.3779](#)
- Lehners, J.-L. & Steinhardt, P. J., Planck 2013 results support the cyclic universe. 2013, *Phys.Rev.*, D87, 123533, [arXiv:1304.3122](#)
- Leistedt, B., Peiris, H. V., & Roth, N., Constraints on Primordial Non-Gaussianity from 800000 Photometric Quasars. 2014, *Phys.Rev.Lett.*, 113, 221301, [arXiv:1405.4315](#)
- Lewis, A., Challinor, A., & Hanson, D., The shape of the CMB lensing bispectrum. 2011, *J. Cosmology Astropart. Phys.*, 3, 18
- Li, M., Note on the production of scale-invariant entropy perturbation in the Ekpyrotic universe. 2013, *Phys.Lett.*, B724, 192, [arXiv:1306.0191](#)
- Liguori, M., Sefusatti, E., Fergusson, J. R., & Shellard, E. P. S., Primordial Non-Gaussianity and Bispectrum Measurements in the Cosmic Microwave Background and Large-Scale Structure. 2010, *Advances in Astronomy*, 2010
- Liguori, M., Shiraishi, M., Fergusson, J., & Shellard, E., . 2015, In prep.
- Linde, A., Mooij, S., & Pajer, E., Gauge field production in supergravity inflation: Local non-Gaussianity and primordial black holes. 2013, *Phys. Rev. D*, 87, 103506, [arXiv:1212.1693](#)
- Linde, A. & Mukhanov, V., Non-Gaussian isocurvature perturbations from inflation. 1997, *Phys. Rev. D*, 56, 535
- Linde, A. & Mukhanov, V., The curvaton web. 2006, *J. Cosmology Astropart. Phys.*, 4, 9, [arXiv:astro-ph/0511736](#)
- Lyth, D. H., Generating the curvature perturbation at the end of inflation. 2005, *J. Cosmology Astropart. Phys.*, 11, 6
- Lyth, D. H. & Riotto, A., Generating the Curvature Perturbation at the End of Inflation in String Theory. 2006, *Physical Review Letters*, 97, 121301
- Lyth, D. H. & Rodriguez, Y., Inflationary Prediction for Primordial Non-Gaussianity. 2005, *Physical Review Letters*, 95, 121302
- Lyth, D. H., Ungarelli, C., & Wands, D., Primordial density perturbation in the curvaton scenario. 2003, *Phys. Rev. D*, 67, 023503
- Lyth, D. H. & Wands, D., Generating the curvature perturbation without an inflaton. 2002, *Physics Letters B*, 524, 5
- Maldacena, J., Non-gaussian features of primordial fluctuations in single field inflationary models. 2003, *Journal of High Energy Physics*, 5, 13
- Maldacena, J. M. & Pimentel, G. L., On graviton non-gaussianities during inflation. 2011, *Journal of High Energy Physics*, 9, 45, [arXiv:1104.2846](#)
- Malik, K. A. & Lyth, D. H., A numerical study of non-Gaussianity in the curvaton scenario. 2006, *J. Cosmology Astropart. Phys.*, 9, 8
- Mangilli, A. & Verde, L., Non-Gaussianity and the CMB bispectrum: Confusion between primordial and lensing-Rees-Sciama contribution? 2009, *Phys. Rev. D*, 80, 123007, [arXiv:0906.2317](#)
- Mangilli, A., Wandelt, B., Elsner, F., & Liguori, M., Optimal bispectrum estimator and simulations of the CMB lensing-integrated Sachs Wolfe non-Gaussian signal. 2013, *A&A*, 555, A82
- Matsubara, T., Analytic Minkowski functionals of the cosmic microwave background: Second-order non-Gaussianity with bispectrum and trispectrum. 2010, *Phys. Rev. D*, 81, 083505
- McFadden, P. & Skenderis, K., Cosmological 3-point correlators from holography. 2011, *J. Cosmology Astropart. Phys.*, 6, 30, [arXiv:1104.3894](#)
- Mecke, K. R., Buchert, T., & Wagner, H., Robust morphological measures for large-scale structure in the Universe. 1994, *A&A*, 288, 697, [arXiv:astro-ph/9312028](#)
- Meerburg, P. D., Münchmeyer, M., & Wandelt, B., Joint resonant CMB power spectrum and bispectrum estimation. 2015, [arXiv:1510.01756](#)
- Meerburg, P. D. & Pajer, E., Observational constraints on gauge field production in axion inflation. 2013, *J. Cosmology Astropart. Phys.*, 2, 17, [arXiv:1203.6076](#)
- Meerburg, P. D., van der Schaar, J. P., & Stefano Corasaniti, P., Signatures of initial state modifications on bispectrum statistics. 2009, *J. Cosmology Astropart. Phys.*, 5, 18
- Mizuno, S. & Koyama, K., Primordial non-Gaussianity from the DBI Galileons. 2010, *Phys. Rev. D*, 82, 103518, [arXiv:1009.0677](#)

- Modest, H. I., Räth, C., Banday, A. J., et al., Scale-dependent non-Gaussianities in the CMB data identified with Minkowski functionals and scaling indices. 2013, MNRAS, 428, 551
- Mollerach, S., Isocurvature baryon perturbations and inflation. 1990, Phys. Rev. D, 42, 313
- Moroi, T. & Takahashi, T., Effects of cosmological moduli fields on cosmic microwave background. 2001, Physics Letters B, 522, 215
- Moss, I. G. & Xiong, C., Non-Gaussianity in fluctuations from warm inflation. 2007, J. Cosmology Astropart. Phys., 4, 7
- Münchmeyer, M., Bouchet, F., Jackson, M., & Wandelt, B., The Komatsu Spergel Wandelt estimator for oscillations in the cosmic microwave background bispectrum. 2014, Astron.Astrophys., 570, A94, [arXiv:1405.2550](#)
- Münchmeyer, M., Meerburg, P. D., & Wandelt, B. D., Optimal estimator for resonance bispectra in the CMB. 2015, Phys.Rev., D91, 043534, [arXiv:1412.3461](#)
- Munshi, D. & Heavens, A., A new approach to probing primordial non-Gaussianity. 2010, MNRAS, 401, 2406
- Munshi, D., Smidt, J., Cooray, A., et al., New approaches to probing Minkowski functionals. 2013, MNRAS, 434, 2830, [arXiv:1011.5224](#)
- Natoli, P., de Troia, G., Hikage, C., et al., BOOMERanG constraints on primordial non-Gaussianity from analytical Minkowski functionals. 2010, MNRAS, 408, 1658
- Nicolis, A., Rattazzi, R., & Trincherini, E., Galileon as a local modification of gravity. 2009, Phys. Rev. D, 79, 064036, [arXiv:0811.2197](#)
- Novaes, C. P., Bernui, A., Ferreira, I. S., & Wuensche, C. A., A Neural-Network based estimator to search for primordial non-Gaussianity in Planck CMB maps. 2014, ArXiv e-prints, [arXiv:1409.3876](#)
- Ohashi, J. & Tsujikawa, S., Potential-driven Galileon inflation. 2012, J. Cosmology Astropart. Phys., 10, 35, [arXiv:1207.4879](#)
- Okamoto, T. & Hu, W., Angular trispectra of CMB temperature and polarization. 2002, Phys. Rev. D, 66, 063008
- Park, C., Choi, Y.-Y., Vogeley, M. S., et al., Topology Analysis of the Sloan Digital Sky Survey. I. Scale and Luminosity Dependence. 2005, ApJ, 633, 11, [arXiv:astro-ph/0507059](#)
- Pénin, A., Lacasa, F., & Aghanim, N., Non-Gaussianity of the cosmic infrared background anisotropies - II. Predictions of the bispectrum and constraints forecast. 2014, MNRAS, 439, 143, [arXiv:1312.1252](#)
- Planck Collaboration VII, *Planck* early results. VII. The Early Release Compact Source Catalogue. 2011, A&A, 536, A7, [arXiv:1101.2041](#)
- Planck Collaboration I, *Planck* 2013 results. I. Overview of products and scientific results. 2014, A&A, 571, A1, [arXiv:1303.5062](#)
- Planck Collaboration II, *Planck* 2013 results. II. Low Frequency Instrument data processing. 2014, A&A, 571, A2, [arXiv:1303.5063](#)
- Planck Collaboration III, *Planck* 2013 results. III. LFI systematic uncertainties. 2014, A&A, 571, A3, [arXiv:1303.5064](#)
- Planck Collaboration IV, *Planck* 2013 results. IV. LFI Beams and window functions. 2014, A&A, 571, A4, [arXiv:1303.5065](#)
- Planck Collaboration V, *Planck* 2013 results. V. LFI Calibration. 2014, A&A, 571, A5, [arXiv:1303.5066](#)
- Planck Collaboration VI, *Planck* 2013 results. VI. High Frequency Instrument data processing. 2014, A&A, 571, A6, [arXiv:1303.5067](#)
- Planck Collaboration VII, *Planck* 2013 results. VII. HFI time response and beams. 2014, A&A, 571, A7, [arXiv:1303.5068](#)
- Planck Collaboration VIII, *Planck* 2013 results. VIII. HFI photometric calibration and mapmaking. 2014, A&A, 571, A8, [arXiv:1303.5069](#)
- Planck Collaboration IX, *Planck* 2013 results. IX. HFI spectral response. 2014, A&A, 571, A9, [arXiv:1303.5070](#)
- Planck Collaboration X, *Planck* 2013 results. X. HFI energetic particle effects: characterization, removal, and simulation. 2014, A&A, 571, A10, [arXiv:1303.5071](#)
- Planck Collaboration XI, *Planck* 2013 results. XI. All-sky model of thermal dust emission. 2014, A&A, 571, A11, [arXiv:1312.1300](#)
- Planck Collaboration XII, *Planck* 2013 results. XII. Diffuse component separation. 2014, A&A, 571, A12, [arXiv:1303.5072](#)
- Planck Collaboration XIII, *Planck* 2013 results. XIII. Galactic CO emission. 2014, A&A, 571, A13, [arXiv:1303.5073](#)
- Planck Collaboration XIV, *Planck* 2013 results. XIV. Zodiacal emission. 2014, A&A, 571, A14, [arXiv:1303.5074](#)
- Planck Collaboration XV, *Planck* 2013 results. XV. CMB power spectra and likelihood. 2014, A&A, 571, A15, [arXiv:1303.5075](#)
- Planck Collaboration XVI, *Planck* 2013 results. XVI. Cosmological parameters. 2014, A&A, 571, A16, [arXiv:1303.5076](#)
- Planck Collaboration XVII, *Planck* 2013 results. XVII. Gravitational lensing by large-scale structure. 2014, A&A, 571, A17, [arXiv:1303.5077](#)
- Planck Collaboration XVIII, *Planck* 2013 results. XVIII. The gravitational lensing-infrared background correlation. 2014, A&A, 571, A18, [arXiv:1303.5078](#)
- Planck Collaboration XIX, *Planck* 2013 results. XIX. The integrated Sachs-Wolfe effect. 2014, A&A, 571, A19, [arXiv:1303.5079](#)
- Planck Collaboration XX, *Planck* 2013 results. XX. Cosmology from Sunyaev-Zeldovich cluster counts. 2014, A&A, 571, A20, [arXiv:1303.5080](#)
- Planck Collaboration XXI, *Planck* 2013 results. XXI. Power spectrum and high-order statistics of the *Planck* all-sky Compton parameter map. 2014, A&A, 571, A21, [arXiv:1303.5081](#)
- Planck Collaboration XXII, *Planck* 2013 results. XXII. Constraints on inflation. 2014, A&A, 571, A22, [arXiv:1303.5082](#)
- Planck Collaboration XXIII, *Planck* 2013 results. XXIII. Isotropy and statistics of the CMB. 2014, A&A, 571, A23, [arXiv:1303.5083](#)
- Planck Collaboration XXIV, *Planck* 2013 results. XXIV. Constraints on primordial non-Gaussianity. 2014, A&A, 571, A24, [arXiv:1303.5084](#)
- Planck Collaboration XXV, *Planck* 2013 results. XXV. Searches for cosmic strings and other topological defects. 2014, A&A, 571, A25, [arXiv:1303.5085](#)
- Planck Collaboration XXVI, *Planck* 2013 results. XXVI. Background geometry and topology of the Universe. 2014, A&A, 571, A26, [arXiv:1303.5086](#)
- Planck Collaboration XXVII, *Planck* 2013 results. XXVII. Doppler boosting of the CMB: Eppur si muove. 2014, A&A, 571, A27, [arXiv:1303.5087](#)
- Planck Collaboration XXVIII, *Planck* 2013 results. XXVIII. The Planck Catalogue of Compact Sources. 2014, A&A, 571, A28, [arXiv:1303.5088](#)
- Planck Collaboration XXIX, *Planck* 2013 results. XXIX. The Planck catalogue of Sunyaev-Zeldovich sources. 2014, A&A, 571, A29, [arXiv:1303.5089](#)
- Planck Collaboration XXX, *Planck* 2013 results. XXX. Cosmic infrared background measurements and implications for star formation. 2014, A&A, 571, A30, [arXiv:1309.0382](#)
- Planck Collaboration XXXI, *Planck* 2013 results. XXXI. Consistency of the *Planck* data. 2014, A&A, 571, A31, [arXiv:1508.03375](#)
- Planck Collaboration I, *Planck* 2015 results. I. Overview of products and results. 2016, A&A, submitted, [arXiv:1502.01582](#)
- Planck Collaboration II, *Planck* 2015 results. II. Low Frequency Instrument data processing. 2016, A&A, submitted, [arXiv:1502.01583](#)
- Planck Collaboration III, *Planck* 2015 results. III. LFI systematic uncertainties. 2016, A&A, submitted, [arXiv:1507.08853](#)
- Planck Collaboration IV, *Planck* 2015 results. IV. LFI beams and window functions. 2016, A&A, in press, [arXiv:1502.01584](#)
- Planck Collaboration V, *Planck* 2015 results. V. LFI calibration. 2016, A&A, in press, [arXiv:1505.08022](#)
- Planck Collaboration VI, *Planck* 2015 results. VI. LFI maps. 2016, A&A, submitted, [arXiv:1502.01585](#)
- Planck Collaboration VII, *Planck* 2015 results. VII. High Frequency Instrument data processing: Time-ordered information and beam processing. 2016, A&A, in press, [arXiv:1502.01586](#)
- Planck Collaboration VIII, *Planck* 2015 results. VIII. High Frequency Instrument data processing: Calibration and maps. 2016, A&A, in press, [arXiv:1502.01587](#)
- Planck Collaboration IX, *Planck* 2015 results. IX. Diffuse component separation: CMB maps. 2016, A&A, submitted, [arXiv:1502.05956](#)
- Planck Collaboration XII, *Planck* 2015 results. XII. Full Focal Plane simulations. 2016, A&A, submitted, [arXiv:1509.06348](#)
- Planck Collaboration XIV, *Planck* 2015 results. XIV. Dark energy and modified gravity. 2016, A&A, submitted, [arXiv:1502.01590](#)
- Planck Collaboration XIX, *Planck* 2015 results. XIX. Constraints on primordial magnetic fields. 2016, A&A, submitted, [arXiv:1502.01594](#)
- Planck Collaboration XX, *Planck* 2015 results. XX. Constraints on inflation. 2016, A&A, submitted, [arXiv:1502.02114](#)
- Qiu, T., Gao, X., & Saridakis, E. N., Towards anisotropy-free and nonsingular bounce cosmology with scale-invariant perturbations. 2013, Phys.Rev., D88, 043525, [arXiv:1303.2372](#)
- Ratra, B., Cosmological 'seed' magnetic field from inflation. 1992, ApJ, 391, L1
- Regan, D., Gosenca, M., & Seery, D., Constraining the WMAP9 bispectrum and trispectrum with needlets. 2013, [arXiv:1310.8617](#)
- Ribeiro, R. H. & Seery, D., Decoding the bispectrum of single-field inflation. 2011, J. Cosmology Astropart. Phys., 10, 27, [arXiv:1108.3839](#)
- Rigopoulos, G. I., Shellard, E. P. S., & van Tent, B. J. W., Large non-Gaussianity in multiple-field inflation. 2006, Phys. Rev. D, 73, 083522
- Rigopoulos, G. I., Shellard, E. P. S., & van Tent, B. J. W., Quantitative bispectra from multifield inflation. 2007, Phys. Rev. D, 76, 083512
- Rudjord, Ø., Hansen, F. K., Lan, X., et al., Directional Variations of the Non-Gaussianity Parameter f_{NL} . 2010, ApJ, 708, 1321, [arXiv:0906.3232](#)
- Salem, M. P., Generation of density perturbations at the end of inflation. 2005, Phys. Rev. D, 72, 123516
- Sasaki, M., Valiviita, J., & Wands, D., Non-Gaussianity of the primordial perturbation in the curvaton model. 2006, Phys.Rev., D74, 103003, [arXiv:astro-ph/0607627](#)
- Schmalzing, J. & Buchert, T., Beyond Genus Statistics: A Unifying Approach to the Morphology of Cosmic Structure. 1997, ApJ, 482, L1, [arXiv:astro-ph/9702130](#)
- Schmalzing, J. & Gorski, K. M., Minkowski functionals used in the morphological analysis of cosmic microwave background anisotropy maps. 1998,

MNRAS, 297, 355, [arXiv:astro-ph/9710185](#)
 Sekiguchi, T. & Sugiyama, N., Optimal constraint on g_{NL} from CMB. 2013, JCAP, 1309, 002, [arXiv:1303.4626](#)
 Senatore, L., Smith, K. M., & Zaldarriaga, M., Non-Gaussianities in single field inflation and their optimal limits from the WMAP 5-year data. 2010, J. Cosmology Astropart. Phys., 1, 28
 Senatore, L. & Zaldarriaga, M., A naturally large four-point function in single field inflation. 2011, J. Cosmology Astropart. Phys., 1, 3
 Senatore, L. & Zaldarriaga, M., A Note on the Consistency Condition of Primordial Fluctuations. 2012a, JCAP, 1208, 001, [arXiv:1203.6884](#)
 Senatore, L. & Zaldarriaga, M., The Effective Field Theory of Multifield Inflation. 2012b, JHEP, 1204, 024, [arXiv:1009.2093](#)
 Shiraishi, M., Parity violation of primordial magnetic fields in the CMB bispectrum. 2012, J. Cosmology Astropart. Phys., 6, 15, [arXiv:1202.2847](#)
 Shiraishi, M., Komatsu, E., Peloso, M., & Barnaby, N., Signatures of anisotropic sources in the squeezed-limit bispectrum of the cosmic microwave background. 2013a, J. Cosmology Astropart. Phys., 5, 2
 Shiraishi, M., Liguori, M., & Fergusson, J. R., General parity-odd CMB bispectrum estimation. 2014, J. Cosmology Astropart. Phys., 5, 8, [arXiv:1403.4222](#)
 Shiraishi, M., Liguori, M., & Fergusson, J. R., Observed parity-odd CMB temperature bispectrum. 2015, J. Cosmology Astropart. Phys., 1, 7, [arXiv:1409.0265](#)
 Shiraishi, M., Nitta, D., & Yokoyama, S., Parity Violation of Gravitons in the CMB Bispectrum. 2011, Progress of Theoretical Physics, 126, 937, [arXiv:1108.0175](#)
 Shiraishi, M., Nitta, D., Yokoyama, S., & Ichiki, K., Optimal limits on primordial magnetic fields from CMB temperature bispectrum of passive modes. 2012, JCAP, 1203, 041, [arXiv:1201.0376](#)
 Shiraishi, M., Ricciardone, A., & Saga, S., Parity violation in the CMB bispectrum by a rolling pseudoscalar. 2013b, J. Cosmology Astropart. Phys., 11, 51, [arXiv:1308.6769](#)
 Silverstein, E. & Tong, D., Scalar speed limits and cosmology: Acceleration from D-acceleration. 2004, Phys. Rev. D, 70, 103505
 Sitwell, M. & Sigurdson, K., Quantization of perturbations in an inflating elastic solid. 2014, Phys. Rev. D, 89, 123509, [arXiv:1306.5762](#)
 Smidt, J., Amblard, A., Byrnes, C. T., et al., CMB constraints on primordial non-Gaussianity from the bispectrum (f_{NL}) and trispectrum (g_{NL} and τ_{NL}) and a new consistency test of single-field inflation. 2010, Phys. Rev. D, 81, 123007
 Smith, K. M., Senatore, L., & Zaldarriaga, M., Optimal limits on f_{NL}^{local} from WMAP 5-year data. 2009, J. Cosmology Astropart. Phys., 9, 6
 Smith, K. M., Senatore, L., & Zaldarriaga, M., Optimal analysis of the CMB trispectrum. 2015, [arXiv:1502.00635](#)
 Soda, J., Kodama, H., & Nozawa, M., Parity violation in graviton non-gaussianity. 2011, Journal of High Energy Physics, 8, 67, [arXiv:1106.3228](#)
 Sorbo, L., Parity violation in the Cosmic Microwave Background from a pseudoscalar inflaton. 2011, J. Cosmology Astropart. Phys., 6, 3, [arXiv:1101.1525](#)
 Spergel, D. N., Bean, R., Doré, O., et al., Three-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Implications for Cosmology. 2007, ApJS, 170, 377, [arXiv:astro-ph/0603449](#)
 Toffolatti, L., Argüeso Gomez, F., de Zotti, G., et al., Extragalactic source counts and contributions to the anisotropies of the cosmic microwave background: predictions for the Planck Surveyor mission. 1998, MNRAS, 297, 117
 Tzavara, E. & van Tent, B., Bispectra from two-field inflation using the long-wavelength formalism. 2011, J. Cosmology Astropart. Phys., 6, 26
 Vanmarcke, E. 1983, Random Fields, ed. Vanmarcke, E.
 Verde, L., Wang, L., Heavens, A. F., & Kamionkowski, M., Large-scale structure, the cosmic microwave background and primordial non-Gaussianity. 2000, MNRAS, 313, 141
 Vernizzi, F. & Wands, D., Non-Gaussianities in two-field inflation. 2006, J. Cosmology Astropart. Phys., 5, 19
 Vielva, P. & Sanz, J., Constraints on f_{NL} and g_{NL} from the analysis of the N-pdf of the CMB large scale anisotropies. 2010, Mon.Not.Roy.Astron.Soc., 404, 895, [arXiv:0910.3196](#)
 Wang, L. & Kamionkowski, M., Cosmic microwave background bispectrum and inflation. 2000, Phys. Rev. D, 61, 063504
 Weinberg, S., Effective field theory for inflation. 2008, Phys. Rev. D, 77, 123541
 Wiegand, A., Buchert, T., & Ostermann, M., Direct Minkowski Functional analysis of large redshift surveys: a new high-speed code tested on the luminous red galaxy Sloan Digital Sky Survey-DR7 catalogue. 2014, MNRAS, 443, 241, [arXiv:1311.3661](#)
 Winitzki, S. & Kosowsky, A., Minkowski functional description of microwave background Gaussianity. 1998, New A, 3, 75, [arXiv:astro-ph/9710164](#)
 Yadav, A. P. S., Komatsu, E., & Wandelt, B. D., Fast Estimator of Primordial Non-Gaussianity from Temperature and Polarization Anisotropies in the Cosmic Microwave Background. 2007, ApJ, 664, 680
 Yadav, A. P. S., Komatsu, E., Wandelt, B. D., et al., Fast Estimator of Primordial Non-Gaussianity from Temperature and Polarization Anisotropies in the Cosmic Microwave Background. II. Partial Sky Coverage and

Inhomogeneous Noise. 2008, ApJ, 678, 578
 Yadav, A. P. S. & Wandelt, B. D., Evidence of Primordial Non-Gaussianity (f_{NL}) in the Wilkinson Microwave Anisotropy Probe 3-Year Data at 2.8σ . 2008, Physical Review Letters, 100, 181301
 Yadav, A. P. S. & Wandelt, B. D., Primordial Non-Gaussianity in the Cosmic Microwave Background. 2010, Advances in Astronomy, 2010
 Zaldarriaga, M., Non-Gaussianities in models with a varying inflaton decay rate. 2004, Phys. Rev. D, 69, 043508

Appendix A: Derivation of an estimator for c_L

As parameterized by Eq. (23), we express a primordial bispectrum of direction-dependence:

$$B_{\Phi}(k_1, k_2, k_3) = \sum_{L \geq 1} c_L \left[P_L(\hat{\mathbf{k}}_1 \cdot \hat{\mathbf{k}}_2) P_{\Phi}(k_1) P_{\Phi}(k_2) + 2 \text{ perm.} \right] \quad (\text{A.1})$$

where $P_L(\hat{\mathbf{k}}_1 \cdot \hat{\mathbf{k}}_2)$ is a Legendre Polynomial of order L . It can be shown that such a primordial bispectrum leads to a CMB bispectrum

$$\begin{aligned} \langle a_{\ell_1 m_1}^{p_1} a_{\ell_2 m_2}^{p_2} a_{\ell_3 m_3}^{p_3} \rangle = & \quad (\text{A.2}) \\ (-i)^{\ell_1 + \ell_2 + \ell_3} (4\pi)^3 \int \frac{d^3 \mathbf{k}_1}{(2\pi)^3} g_{\ell_1}^{p_1}(k_1) Y_{\ell_1 m_1}^*(\hat{\mathbf{k}}_1) \times \\ & \int \frac{d^3 \mathbf{k}_2}{(2\pi)^3} g_{\ell_2}^{p_2}(k_2) Y_{\ell_2 m_2}^*(\hat{\mathbf{k}}_2) \int \frac{d^3 \mathbf{k}_3}{(2\pi)^3} g_{\ell_3}^{p_3}(k_3) Y_{\ell_3 m_3}^*(\hat{\mathbf{k}}_3) \times \\ & (2\pi)^3 \delta^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) c_L \left[P_L(\hat{\mathbf{k}}_1 \cdot \hat{\mathbf{k}}_2) P_{\Phi}(k_1) P_{\Phi}(k_2) + 2 \text{ perm.} \right] \\ = & \left(\begin{matrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{matrix} \right) \left(\sum_L c_L b_{\ell_1 \ell_2 \ell_3}^{(c_L), p_1 p_2 p_3} \right), \end{aligned}$$

where p denotes either temperature or E -mode polarization, $b_{\ell_1 \ell_2 \ell_3}^{(c_L), p_1 p_2 p_3}$ is the reduced CMB bispectrum, and the term with large parentheses denotes the Wigner-3j symbol. The reduced CMB bispectrum $b_{\ell_1 \ell_2 \ell_3}^{(c_L), p_1 p_2 p_3}$ is given by (Shiraishi et al. 2013a)

$$\begin{aligned} b_{\ell_1 \ell_2 \ell_3}^{(c_L), p_1 p_2 p_3} = & \frac{4\pi}{2L+1} \frac{w_{\ell_1} w_{\ell_2} w_{\ell_3} t^{\ell_1 + \ell_2 + \ell_3}}{h_{\ell_1 \ell_2 \ell_3}} \int_0^{\infty} r^2 dr \\ & \times \sum_{L_1 L_2 L_3} t^{L_1 + L_2 + L_3} h_{L_1 L_2 L_3} \left[h_{\ell_1 L_1 L} h_{\ell_2 L_2 L} (-1)^{L_1 + L_2 + L} \delta_{L_3, \ell_3} \right. \\ & \left. \times \beta_{\ell_1 L_1}^{p_1}(r) \beta_{\ell_2 L_2}^{p_2}(r) \alpha_{\ell_3}^{p_3}(r) \left\{ \begin{matrix} \ell_1 & \ell_3 & \ell_3 \\ L_2 & L_1 & L \end{matrix} \right\} + 2 \text{ perm.} \right], \quad (\text{A.3}) \end{aligned}$$

where $\delta_{L,l}$ denotes the Kronecker delta function, $\{ \dots \}$ the Wigner 6j symbol, ‘‘perm.’’ means permutations, w_{ℓ} is a beam window function, and

$$\alpha_{\ell}^p(r) = \frac{2}{\pi} \int k^3 d \ln k g_{\ell}^p(k) j_{\ell}(kr), \quad (\text{A.4})$$

$$\beta_{\ell, \ell'}^p(r) = \frac{2}{\pi} \int dk^3 \ln k P_{\Phi}(k) g_{\ell}^p(k) j_{\ell'}(kr). \quad (\text{A.5})$$

Here, $g_{\ell}^p(k)$ is the radiation transfer function for temperature or E -mode polarization, and $j_{\ell}(x)$ is a spherical Bessel function. The h symbol is

$$h_{\ell_1 \ell_2 \ell_3} = \sqrt{\frac{(2\ell_1 + 1)(2\ell_2 + 1)(2\ell_3 + 1)}{4\pi}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix}. \quad (\text{A.6})$$

By maximizing the likelihood with respect to c_L , we obtain the KSW estimator for c_L :

$$\hat{c}_L = \frac{1}{6 N_{c_L}} \sum_{p_i q_i} \sum_{\ell_i m_i} \mathcal{G}_{\ell_1 \ell_2 \ell_3}^{m_1 m_2 m_3} b_{\ell_1 \ell_2 \ell_3}^{(c_L), p_1 p_2 p_3} \times \left[(C^{-1} a)_{\ell_1 m_1}^{p_1} (C^{-1} a)_{\ell_2 m_2}^{p_2} (C^{-1} a)_{\ell_3 m_3}^{p_3} - 3 \langle (C^{-1} a)_{\ell_1 m_1}^{p_1} (C^{-1} a)_{\ell_2 m_2}^{p_2} \rangle_{MC} (C^{-1} a)_{\ell_3 m_3}^{p_3} \right], \quad (\text{A.7})$$

where MC denotes that the average is over Monte Carlo simulations and N_{c_L} is a normalization constant. The normalization constant N_{c_L} is given by

$$N_{c_L} = \frac{1}{6} \sum_{p_i q_i} \sum_{\ell_i} (h_{\ell_1 \ell_2 \ell_3})^2 \times b_{\ell_1 \ell_2 \ell_3}^{(c_L), p_1 p_2 p_3} (C^{-1})_{\ell_1}^{p_1 q_1} (C^{-1})_{\ell_2}^{p_2 q_2} (C^{-1})_{\ell_3}^{p_3 q_3} b_{\ell_1 \ell_2 \ell_3}^{(c_L), q_1 q_2 q_3}. \quad (\text{A.8})$$

Using Eq. (A.2), we find that

$$\sum_{p_i \ell_i m_i} \mathcal{G}_{\ell_1 \ell_2 \ell_3}^{m_1 m_2 m_3} b_{\ell_1 \ell_2 \ell_3}^{(c_L), p_1 p_2 p_3} (C^{-1} a)_{\ell_1 m_1}^{p_1} (C^{-1} a)_{\ell_2 m_2}^{p_2} (C^{-1} a)_{\ell_3 m_3}^{p_3} = \int r^2 dr \int d\Omega \frac{4\pi}{2L+1} \left(\sum_{q_3 \ell_3 m_3} \alpha_{\ell_3}^{p_3}(r) (C^{-1} a)_{\ell_3 m_3}^{p_3} Y_{\ell_3 m_3}(\hat{\mathbf{n}}) \right) \times \left(\sum_{M=-L}^L (-1)^M \left(\sum_{\ell'_1 m'_1} b_{\ell'_1 m'_1}^{LM}(r) Y_{\ell'_1 m'_1}(\hat{\mathbf{n}}) \right) \left(\sum_{\ell'_2 m'_2} b_{\ell'_2 m'_2}^{L-M}(r) Y_{\ell'_2 m'_2}(\hat{\mathbf{n}}) \right) + 2 \text{ perm.} \right) \quad (\text{A.9})$$

where

$$b_{\ell' m'}^{LM}(r) = \sum_{pq\ell m} (-1)^{\ell+m} i^{\ell+\ell'} h_{\ell' \ell L} \begin{pmatrix} \ell & \ell' & L \\ -m & m' & M \end{pmatrix} \beta_{\ell, \ell'}^p(r) w_\ell (C^{-1})_l^{pq} a_{\ell m}^q,$$

with $\beta_{\ell, \ell'}^p(r)$ being defined in Eq. (A.5). In the derivation above, we used the identities

$$P_L(\hat{\mathbf{k}}_1 \cdot \hat{\mathbf{k}}_2) = \frac{4\pi}{2L+1} \sum_M Y_{LM}(\hat{\mathbf{k}}_1) Y_{LM}^*(\hat{\mathbf{k}}_2),$$

and

$$\delta(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) = 8 \int dr dn^2 \sum_{\ell_1 m_1} i^{\ell_1} j_{\ell_1}(k_1 r) Y_{\ell_1 m_1}(\hat{\mathbf{k}}_1) Y_{\ell_1 m_1}^*(\hat{\mathbf{n}}) \times \sum_{\ell_2 m_2} i^{\ell_2} j_{\ell_2}(k_2 r) Y_{\ell_2 m_2}(\hat{\mathbf{k}}_2) Y_{\ell_2 m_2}^*(\hat{\mathbf{n}}) \sum_{\ell_3 m_3} i^{\ell_3} j_{\ell_3}(k_3 r) Y_{\ell_3 m_3}(\hat{\mathbf{k}}_3) Y_{\ell_3 m_3}^*(\hat{\mathbf{n}}),$$

Applying Eq. (A.9) to Eq. (A.7), we find

$$\hat{c}_L = \frac{1}{N_{c_L}} \frac{2\pi}{(2L+1)} \sum_{M=-L}^L (-1)^M \int r^2 dr \int d^2 \hat{\mathbf{n}} [A(\hat{\mathbf{n}}, r) B_{LM}(\hat{\mathbf{n}}, r) B_{L,-M}(\hat{\mathbf{n}}, r) - A(\hat{\mathbf{n}}, r) \langle B_{LM}(\hat{\mathbf{n}}, r) B_{L,-M}(\hat{\mathbf{n}}, r) \rangle_{MC} - 2 B_{L,-M}(\hat{\mathbf{n}}, r) \langle A(\hat{\mathbf{n}}, r) B_{LM}(\hat{\mathbf{n}}, r) \rangle_{MC}],$$

where

$$A(\hat{\mathbf{n}}, r) = \sum_{\ell m} \sum_{pq} \alpha_{\ell}^p(r) (C^{-1})_l^{pq} a_{\ell m}^q Y_{\ell m}(\hat{\mathbf{n}}),$$

$$B_{LM}(\hat{\mathbf{n}}, r) = \sum_{\ell m} b_{\ell m}^{LM}(r) Y_{\ell m}(\hat{\mathbf{n}}).$$

Since $b_{\ell' m'}^{LM}(r) \neq (-1)^{m'} [b_{\ell', -m'}^{LM}(r)]^*$, $B_{LM}(\hat{\mathbf{n}}, r)$ is not a real function, but a complex function. We estimate $B_{LM}(\hat{\mathbf{n}}, r)$ efficiently by computing the following with HEALPix (Górski et al. 2005):

$$\text{Re}[B_{LM}(\hat{\mathbf{n}}, r)] = \sum_{\ell m} \mathcal{R}_{\ell m}^{LM}(r) Y_{\ell m}(\hat{\mathbf{n}});$$

$$\text{Im}[B_{LM}(\hat{\mathbf{n}}, r)] = \sum_{\ell m} \mathcal{I}_{\ell m}^{LM}(r) Y_{\ell m}(\hat{\mathbf{n}}).$$

Here

$$\mathcal{R}_{\ell m}^{LM}(r) = \frac{b_{\ell m}^{LM}(r) + (-1)^m [b_{\ell, -m}^{LM}(r)]^*}{2},$$

$$\mathcal{I}_{\ell m}^{LM}(r) = \frac{b_{\ell m}^{LM}(r) - (-1)^m [b_{\ell, -m}^{LM}(r)]^*}{2i}.$$

Appendix B: Definition of Minkowski Functionals and theoretical expectations

For a field $f(x)$ of zero average and variance σ_0^2 defined on the two-dimensional sphere \mathbb{S}^2 , an overdense excursion set is defined as

$$\Sigma \equiv \{x \in \mathbb{S}^2 | f(x) > \nu \sigma_0\}. \quad (\text{B.1})$$

The boundary of the excursion is

$$\partial \Sigma \equiv \{x \in \mathbb{S}^2 | f(x) = \nu \sigma_0\}. \quad (\text{B.2})$$

Then the three Minkowski functionals on the sphere are

$$\text{Area} : V_0(\nu) = \frac{1}{4\pi} \int_{\Sigma} d\Omega, \quad (\text{B.3})$$

$$\text{Perimeter} : V_1(\nu) = \frac{1}{4\pi} \frac{1}{4} \int_{\partial \Sigma} dl, \quad (\text{B.4})$$

$$\text{Genus} : V_2(\nu) = \frac{1}{4\pi} \frac{1}{2\pi} \int_{\partial \Sigma} \kappa dl, \quad (\text{B.5})$$

where $d\Omega$ and dl are respectively elements of solid angles (surface) and of angle (distance), κ is the geodesic curvature. Note that the Genus can be also expressed as the number of components²² in the excursion minus the number of holes in the excursion.

The fourth functional $V_3(\nu)$, is defined, for $\nu > 0$, as the number of components in the excursion. Symmetrically, for $\nu < 0$, it is the number of underdense components (or the number of components in the excursion $\{x \in \mathbb{S}^2 | f(x) < \nu \sigma_0\}$).

In the Gaussian limit, the functionals can be expressed the following way (see, e.g., Matsubara 2010; Vanmarcke 1983):

$$V_k(\nu) = A_k v_k(\nu), \quad (\text{B.6})$$

with

$$v_k(\nu) = \exp(-\nu^2/2) H_{k-1}(\nu), \quad k \leq 2 \quad (\text{B.7})$$

$$v_3(\nu) = \frac{e^{-\nu^2}}{\text{erfc}(\nu/\sqrt{2})}, \quad (\text{B.8})$$

and

$$H_n(\nu) = e^{\nu^2/2} \left(-\frac{d}{d\nu} \right)^n e^{-\nu^2/2}. \quad (\text{B.9})$$

²² A component is a connected subset of the excursion.

The amplitude A_k depends only on the shape of the power spectrum C_ℓ :

$$A_k = \frac{1}{(2\pi)^{(k+1)/2}} \frac{\omega_2}{\omega_{2-k}\omega_k} \left(\frac{\sigma_1}{\sqrt{2}\sigma_0} \right)^k, \quad k \leq 2 \quad (\text{B.10})$$

$$A_3 = \frac{2}{\pi} \left(\frac{\sigma_1}{\sqrt{2}\sigma_0} \right)^2 \quad (\text{B.11})$$

where $\omega_k \equiv \pi^{k/2}/\Gamma(k/2+1)$, which gives $\omega_0 = 1$, $\omega_1 = 2$, $\omega_2 = \pi$ and σ_0 and σ_1 are respectively the rms of the field and its first derivatives.

-
- ¹ APC, AstroParticule et Cosmologie, Université Paris Diderot, CNRS/IN2P3, CEA/Irfu, Observatoire de Paris, Sorbonne Paris Cité, 10, rue Alice Domon et Léonie Duquet, 75205 Paris Cedex 13, France
- ² Aalto University Metsähovi Radio Observatory and Dept of Radio Science and Engineering, P.O. Box 13000, FI-00076 AALTO, Finland
- ³ African Institute for Mathematical Sciences, 6-8 Melrose Road, Muizenberg, Cape Town, South Africa
- ⁴ Agenzia Spaziale Italiana Science Data Center, Via del Politecnico snc, 00133, Roma, Italy
- ⁵ Aix Marseille Université, CNRS, LAM (Laboratoire d'Astrophysique de Marseille) UMR 7326, 13388, Marseille, France
- ⁶ Astrophysics Group, Cavendish Laboratory, University of Cambridge, J J Thomson Avenue, Cambridge CB3 0HE, U.K.
- ⁷ Astrophysics & Cosmology Research Unit, School of Mathematics, Statistics & Computer Science, University of KwaZulu-Natal, Westville Campus, Private Bag X54001, Durban 4000, South Africa
- ⁸ CGEE, SCS Qd 9, Lote C, Torre C, 4° andar, Ed. Parque Cidade Corporate, CEP 70308-200, Brasília, DF, Brazil
- ⁹ CITA, University of Toronto, 60 St. George St., Toronto, ON M5S 3H8, Canada
- ¹⁰ CNRS, IRAP, 9 Av. colonel Roche, BP 44346, F-31028 Toulouse cedex 4, France
- ¹¹ CRANN, Trinity College, Dublin, Ireland
- ¹² California Institute of Technology, Pasadena, California, U.S.A.
- ¹³ Centre for Theoretical Cosmology, DAMTP, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, U.K.
- ¹⁴ Centro de Estudios de Física del Cosmos de Aragón (CEFCA), Plaza San Juan, 1, planta 2, E-44001, Teruel, Spain
- ¹⁵ Computational Cosmology Center, Lawrence Berkeley National Laboratory, Berkeley, California, U.S.A.
- ¹⁶ Consejo Superior de Investigaciones Científicas (CSIC), Madrid, Spain
- ¹⁷ DSM/Irfu/SPP, CEA-Saclay, F-91191 Gif-sur-Yvette Cedex, France
- ¹⁸ DTU Space, National Space Institute, Technical University of Denmark, Elektrovej 327, DK-2800 Kgs. Lyngby, Denmark
- ¹⁹ Département de Physique Théorique, Université de Genève, 24, Quai E. Ansermet, 1211 Genève 4, Switzerland
- ²⁰ Departamento de Astrofísica, Universidad de La Laguna (ULL), E-38206 La Laguna, Tenerife, Spain
- ²¹ Departamento de Física, Universidad de Oviedo, Avda. Calvo Sotelo s/n, Oviedo, Spain
- ²² Department of Astronomy and Astrophysics, University of Toronto, 50 Saint George Street, Toronto, Ontario, Canada
- ²³ Department of Astrophysics/IMAPP, Radboud University Nijmegen, P.O. Box 9010, 6500 GL Nijmegen, The Netherlands
- ²⁴ Department of Physics & Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, British Columbia, Canada
- ²⁵ Department of Physics and Astronomy, Dana and David Dornsife College of Letter, Arts and Sciences, University of Southern California, Los Angeles, CA 90089, U.S.A.
- ²⁶ Department of Physics and Astronomy, University College London, London WC1E 6BT, U.K.
- ²⁷ Department of Physics and Astronomy, University of Sussex, Brighton BN1 9QH, U.K.
- ²⁸ Department of Physics, Florida State University, Keen Physics Building, 77 Chieftan Way, Tallahassee, Florida, U.S.A.
- ²⁹ Department of Physics, Gustaf Hällströmin katu 2a, University of Helsinki, Helsinki, Finland
- ³⁰ Department of Physics, Princeton University, Princeton, New Jersey, U.S.A.
- ³¹ Department of Physics, University of California, Santa Barbara, California, U.S.A.
- ³² Department of Physics, University of Illinois at Urbana-Champaign, 1110 West Green Street, Urbana, Illinois, U.S.A.
- ³³ Dipartimento di Fisica e Astronomia G. Galilei, Università degli Studi di Padova, via Marzolo 8, 35131 Padova, Italy
- ³⁴ Dipartimento di Fisica e Astronomia, ALMA MATER STUDIUM, Università degli Studi di Bologna, Viale Berti Pichat 6/2, I-40127, Bologna, Italy
- ³⁵ Dipartimento di Fisica e Scienze della Terra, Università di Ferrara, Via Saragat 1, 44122 Ferrara, Italy
- ³⁶ Dipartimento di Fisica, Università La Sapienza, P. le A. Moro 2, Roma, Italy
- ³⁷ Dipartimento di Fisica, Università degli Studi di Milano, Via Celoria, 16, Milano, Italy
- ³⁸ Dipartimento di Fisica, Università degli Studi di Trieste, via A. Valerio 2, Trieste, Italy
- ³⁹ Dipartimento di Matematica, Università di Roma Tor Vergata, Via della Ricerca Scientifica, 1, Roma, Italy
- ⁴⁰ Discovery Center, Niels Bohr Institute, Blegdamsvej 17, Copenhagen, Denmark
- ⁴¹ Discovery Center, Niels Bohr Institute, Copenhagen University, Blegdamsvej 17, Copenhagen, Denmark
- ⁴² European Space Agency, ESAC, Planck Science Office, Camino bajo del Castillo, s/n, Urbanización Villafranca del Castillo, Villanueva de la Cañada, Madrid, Spain
- ⁴³ European Space Agency, ESTEC, Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands
- ⁴⁴ Gran Sasso Science Institute, INFN, viale F. Crispi 7, 67100 L'Aquila, Italy
- ⁴⁵ HGSFP and University of Heidelberg, Theoretical Physics Department, Philosophenweg 16, 69120, Heidelberg, Germany
- ⁴⁶ Helsinki Institute of Physics, Gustaf Hällströmin katu 2, University of Helsinki, Helsinki, Finland
- ⁴⁷ ICTP South American Institute for Fundamental Research, Instituto de Física Teórica, Universidade Estadual Paulista, São Paulo, Brazil
- ⁴⁸ INAF - Osservatorio Astronomico di Padova, Vicolo dell'Osservatorio 5, Padova, Italy
- ⁴⁹ INAF - Osservatorio Astronomico di Roma, via di Frascati 33, Monte Porzio Catone, Italy
- ⁵⁰ INAF - Osservatorio Astronomico di Trieste, Via G.B. Tiepolo 11, Trieste, Italy
- ⁵¹ INAF/IASF Bologna, Via Gobetti 101, Bologna, Italy
- ⁵² INAF/IASF Milano, Via E. Bassini 15, Milano, Italy
- ⁵³ INFN, Sezione di Bologna, Via Imerio 46, I-40126, Bologna, Italy
- ⁵⁴ INFN, Sezione di Roma 1, Università di Roma Sapienza, Piazzale Aldo Moro 2, 00185, Roma, Italy
- ⁵⁵ INFN, Sezione di Roma 2, Università di Roma Tor Vergata, Via della Ricerca Scientifica, 1, Roma, Italy
- ⁵⁶ INFN/National Institute for Nuclear Physics, Via Valerio 2, I-34127 Trieste, Italy
- ⁵⁷ IPAG: Institut de Planétologie et d'Astrophysique de Grenoble, Université Grenoble Alpes, IPAG, F-38000 Grenoble, France, CNRS, IPAG, F-38000 Grenoble, France
- ⁵⁸ IUCAA, Post Bag 4, Ganeshkhind, Pune University Campus, Pune 411 007, India
- ⁵⁹ Imperial College London, Astrophysics group, Blackett Laboratory, Prince Consort Road, London, SW7 2AZ, U.K.

- ⁶⁰ Infrared Processing and Analysis Center, California Institute of Technology, Pasadena, CA 91125, U.S.A.
- ⁶¹ Institut Néel, CNRS, Université Joseph Fourier Grenoble I, 25 rue des Martyrs, Grenoble, France
- ⁶² Institut Universitaire de France, 103, bd Saint-Michel, 75005, Paris, France
- ⁶³ Institut d'Astrophysique Spatiale, CNRS, Univ. Paris-Sud, Université Paris-Saclay, Bât. 121, 91405 Orsay cedex, France
- ⁶⁴ Institut d'Astrophysique de Paris, CNRS (UMR7095), 98 bis Boulevard Arago, F-75014, Paris, France
- ⁶⁵ Institut für Theoretische Teilchenphysik und Kosmologie, RWTH Aachen University, D-52056 Aachen, Germany
- ⁶⁶ Institute for Space Sciences, Bucharest-Magurale, Romania
- ⁶⁷ Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, U.K.
- ⁶⁸ Institute of Theoretical Astrophysics, University of Oslo, Blindern, Oslo, Norway
- ⁶⁹ Instituto de Astrofísica de Canarias, C/Vía Láctea s/n, La Laguna, Tenerife, Spain
- ⁷⁰ Instituto de Física de Cantabria (CSIC-Universidad de Cantabria), Avda. de los Castros s/n, Santander, Spain
- ⁷¹ Istituto Nazionale di Fisica Nucleare, Sezione di Padova, via Marzolo 8, I-35131 Padova, Italy
- ⁷² Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, California, U.S.A.
- ⁷³ Jodrell Bank Centre for Astrophysics, Alan Turing Building, School of Physics and Astronomy, The University of Manchester, Oxford Road, Manchester, M13 9PL, U.K.
- ⁷⁴ Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA
- ⁷⁵ Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge, CB3 0HA, U.K.
- ⁷⁶ Kazan Federal University, 18 Kremlyovskaya St., Kazan, 420008, Russia
- ⁷⁷ LAL, Université Paris-Sud, CNRS/IN2P3, Orsay, France
- ⁷⁸ LERMA, CNRS, Observatoire de Paris, 61 Avenue de l'Observatoire, Paris, France
- ⁷⁹ Laboratoire AIM, IRFU/Service d'Astrophysique - CEA/DSM - CNRS - Université Paris Diderot, Bât. 709, CEA-Saclay, F-91191 Gif-sur-Yvette Cedex, France
- ⁸⁰ Laboratoire Traitement et Communication de l'Information, CNRS (UMR 5141) and Télécom ParisTech, 46 rue Barrault F-75634 Paris Cedex 13, France
- ⁸¹ Laboratoire de Physique Subatomique et Cosmologie, Université Grenoble-Alpes, CNRS/IN2P3, 53, rue des Martyrs, 38026 Grenoble Cedex, France
- ⁸² Laboratoire de Physique Théorique, Université Paris-Sud 11 & CNRS, Bâtiment 210, 91405 Orsay, France
- ⁸³ Lawrence Berkeley National Laboratory, Berkeley, California, U.S.A.
- ⁸⁴ Lebedev Physical Institute of the Russian Academy of Sciences, Astro Space Centre, 84/32 Profsoyuznaya st., Moscow, GSP-7, 117997, Russia
- ⁸⁵ Leung Center for Cosmology and Particle Astrophysics, National Taiwan University, Taipei 10617, Taiwan
- ⁸⁶ Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85741 Garching, Germany
- ⁸⁷ McGill Physics, Ernest Rutherford Physics Building, McGill University, 3600 rue University, Montréal, QC, H3A 2T8, Canada
- ⁸⁸ National University of Ireland, Department of Experimental Physics, Maynooth, Co. Kildare, Ireland
- ⁸⁹ Nicolaus Copernicus Astronomical Center, Bartycka 18, 00-716 Warsaw, Poland
- ⁹⁰ Niels Bohr Institute, Blegdamsvej 17, Copenhagen, Denmark
- ⁹¹ Niels Bohr Institute, Copenhagen University, Blegdamsvej 17, Copenhagen, Denmark
- ⁹² Nordita (Nordic Institute for Theoretical Physics), Roslagstullsbacken 23, SE-106 91 Stockholm, Sweden
- ⁹³ Optical Science Laboratory, University College London, Gower Street, London, U.K.
- ⁹⁴ Perimeter Institute for Theoretical Physics, Waterloo ON N2L 2Y5, Canada
- ⁹⁵ SISSA, Astrophysics Sector, via Bonomea 265, 34136, Trieste, Italy
- ⁹⁶ SMARTEST Research Centre, Università degli Studi e-Campus, Via Isimbardi 10, Novedrate (CO), 22060, Italy
- ⁹⁷ School of Physics and Astronomy, Cardiff University, Queens Buildings, The Parade, Cardiff, CF24 3AA, U.K.
- ⁹⁸ School of Physics and Astronomy, University of Nottingham, Nottingham NG7 2RD, U.K.
- ⁹⁹ Sorbonne Université-UPMC, UMR7095, Institut d'Astrophysique de Paris, 98 bis Boulevard Arago, F-75014, Paris, France
- ¹⁰⁰ Space Research Institute (IKI), Russian Academy of Sciences, Profsoyuznaya Str, 84/32, Moscow, 117997, Russia
- ¹⁰¹ Space Sciences Laboratory, University of California, Berkeley, California, U.S.A.
- ¹⁰² Special Astrophysical Observatory, Russian Academy of Sciences, Nizhnij Arkhyz, Zelenchukskiy region, Karachai-Cherkessian Republic, 369167, Russia
- ¹⁰³ Stanford University, Dept of Physics, Varian Physics Bldg, 382 Via Pueblo Mall, Stanford, California, U.S.A.
- ¹⁰⁴ Sub-Department of Astrophysics, University of Oxford, Keble Road, Oxford OX1 3RH, U.K.
- ¹⁰⁵ Sydney Institute for Astronomy, School of Physics A28, University of Sydney, NSW 2006, Australia
- ¹⁰⁶ The Oskar Klein Centre for Cosmoparticle Physics, Department of Physics, Stockholm University, AlbaNova, SE-106 91 Stockholm, Sweden
- ¹⁰⁷ Theory Division, PH-TH, CERN, CH-1211, Geneva 23, Switzerland
- ¹⁰⁸ UPMC Univ Paris 06, UMR7095, 98 bis Boulevard Arago, F-75014, Paris, France
- ¹⁰⁹ Université de Toulouse, UPS-OMP, IRAP, F-31028 Toulouse cedex 4, France
- ¹¹⁰ University of Granada, Departamento de Física Teórica y del Cosmos, Facultad de Ciencias, Granada, Spain
- ¹¹¹ University of Granada, Instituto Carlos I de Física Teórica y Computacional, Granada, Spain
- ¹¹² Warsaw University Observatory, Aleje Ujazdowskie 4, 00-478 Warszawa, Poland