

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^{164,108}, 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. Gruppiso⁵¹, 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, S. Matarrese^{33,71,44}, P. McGehee⁶⁰, P. R. Meinhold³¹, A. Melchiorri^{36,54}, L. Mendes⁴², A. Mennella^{37,52}, M. Migliaccio^{67,75}, S. Mitra^{58,72}, 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. Shiraishi^{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-Uksi^{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 6 February 2015 / Accepted 27 January 2016

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

* Corresponding author: Nicola Bartolo, e-mail: bartolo@pd.infn.it

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 anisotropies². 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 techniques. 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 generated³.

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, 2009; Langlois et al. 2008; Hikage et al. 2009; Langlois & Lepidi 2011; Langlois & van Tent 2011, 2012; Kawakami et al. 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 (p)reheating 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 (p)reheating and modulated hybrid inflation (Kofman 2003; Dvali et al. 2004a,b; Bernardeau et al. 2004; Zaldarriaga 2004; Lyth 2005; Saleh 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 $f_{\text{NL}}^{IJK} = f_{\text{NL}}^{I,KJ}$, 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](#)):

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

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

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

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

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 [Achúcarro 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

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

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, 2014](#); see also [Bartolo et al. 2013b, 2014](#); [Sitwell & Sigurdson 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

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

where the exact expression for the function f_3 can be found in Eq. (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 $\xi = \alpha|\phi|/(2fH)$. 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

⁴ For simplicity we assume a scale-invariant bispectrum.

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

($F_{\zeta}^{\text{equil}} \equiv (5/3)^3 F_{\Phi}^{\text{equil}} = (5/3)^3 B_{\Phi}^{\text{equil}}/f_{\text{NL}}^{\text{equil}}$) 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:

$$\frac{\Delta T}{T}(\hat{\mathbf{n}}) = \sum_{\ell m} a_{\ell m}^T Y_{\ell m}(\hat{\mathbf{n}}),$$

$$E(\hat{\mathbf{n}}) = \sum_{\ell m} a_{\ell m}^E Y_{\ell m}(\hat{\mathbf{n}}). \quad (29)$$

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.

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,

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

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)

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

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)

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

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

$$\begin{aligned} \langle b^A, b^B \rangle = & \sum_{X_i, X'_i} \sum_{\ell_i} b_{\ell_1 \ell_2 \ell_3}^{X_1 X_2 X_3, A} h_{\ell_1 \ell_2 \ell_3}^2 \left(\mathbf{C}^{-1} \right)_{\ell_1}^{X_1 X'_1} \left(\mathbf{C}^{-1} \right)_{\ell_2}^{X_2 X'_2} \left(\mathbf{C}^{-1} \right)_{\ell_3}^{X_3 X'_3} b_{\ell_1 \ell_2 \ell_3}^{X'_1 X'_2 X'_3, B} \end{aligned} \quad (39)$$

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 Eqs. (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- \mathbf{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,

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).

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, in prep.), 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,

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

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

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 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, 2015), 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 bin⁸, and the same is true for

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.

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 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

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).

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 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 $(S/N) \propto \ell \sqrt{f_{\text{sky}}}$ 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 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

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.

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]).

impact on our f_{NL} results. We devote particular attention to assessing potential biases that these NG signals might induce on the 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

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

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

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

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).

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 usual¹². 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 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

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. 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.

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.

Method	Lensing-ISW amplitude			
	SMICA	SEVEM	NILC	Commander
<i>T</i>				
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
<i>T+E</i>				
Binned . . .	0.82 ± 0.27	0.75 ± 0.28	0.85 ± 0.26	0.84 ± 0.27

Notes. Error bars are 68% CL; see the main text for how they have been determined.

and 7. The error bars will also improve when measured with the other bispectrum estimators. As already seen in 2013, the values for $f_{\text{NL}}^{\text{ISW}}$ 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




Fig. 2. 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.

and late-type infrared galaxies (see e.g., [Toffolatti et al. 1998](#); [González-Nuevo et al. 2005](#)), while the latter are dusty star-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

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

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 $\ell \simeq 200\text{--}700$ ([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,

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

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.

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

Notes. The error bars have been determined using FFP8 simulations.

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 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 $A_{\text{CIB}} \approx 6 \times 10^{-27}$ 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. 2015).

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

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.

	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

Notes. 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.

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.

Shape	f_{NL}						
	KSW	Binned	Modal 1	Modal 2	B – KSW	M1 – KSW	M2 – KSW
<i>T</i> 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
<i>T</i> Equilateral	7 ± 53	5 ± 58	6 ± 53	8 ± 56	-2 ± 12	-1.0 ± 8.4	0 ± 17
<i>T</i> Orthogonal	-22 ± 27	-22 ± 28	-22 ± 27	-17 ± 30	0.5 ± 9.4	-0.2 ± 4.2	5 ± 11
<i>E</i> Local	-0.9 ± 4.1	-1.3 ± 3.4	-0.9 ± 3.7	...	-0.3 ± 2.9	0.1 ± 0.5	...
<i>E</i> Equilateral	-9 ± 42	-10 ± 42	-10 ± 40	...	-1 ± 11	-0.7 ± 9.2	...
<i>E</i> Orthogonal	4 ± 13	5 ± 13	4 ± 12	...	0.1 ± 3.8	-0.3 ± 2.7	...
<i>T+E</i> 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
<i>T+E</i> Equilateral	0 ± 20	2 ± 22	3 ± 21	0 ± 23	1.4 ± 5.8	2.3 ± 7.3	0 ± 12
<i>T+E</i> Orthogonal	-4 ± 10	-4 ± 9	-6 ± 9	-5 ± 12	0.3 ± 2.2	-1.1 ± 3.1	-1.0 ± 7.1

Notes. 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.

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.

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

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
<i>T</i> 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
<i>T</i> Equilateral	11 ± 61	12 ± 65	9 ± 63	12 ± 62	1 ± 15	-1.9 ± 9.6	1 ± 12
<i>T</i> Orthogonal	-19 ± 31	-18 ± 34	-20 ± 32	-18 ± 35	1 ± 12	-1.3 ± 5.1	0.8 ± 8.8
<i>E</i> Local	-2 ± 29	-4 ± 28	-1 ± 29	...	-2 ± 12	0.4 ± 5.5	...
<i>E</i> Equilateral	1 ± 191	-18 ± 195	-6 ± 200	...	-19 ± 47	-7 ± 23	...
<i>E</i> Orthogonal	-6 ± 101	0 ± 107	-6 ± 102	...	6 ± 25	-0.3 ± 10	...
<i>T+E</i> 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
<i>T+E</i> Equilateral	13 ± 46	11 ± 49	9 ± 48	13 ± 47	-2 ± 10	-4 ± 13	-0.3 ± 7.0
<i>T+E</i> 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
<i>T</i> 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
<i>T</i> Equilateral	11 ± 73	9 ± 75	6 ± 76	11 ± 70	-2 ± 19	-5 ± 14	0 ± 12
<i>T</i> Orthogonal	-22 ± 37	-21 ± 37	-23 ± 36	-20 ± 37	2 ± 14	-0.9 ± 6.1	2.6 ± 9.2
<i>E</i> Local	4 ± 36	0 ± 35	5 ± 37	...	-4 ± 16	1 ± 13	...
<i>E</i> Equilateral	-32 ± 242	-49 ± 209	-38 ± 246	...	-17 ± 88	-6 ± 34	...
<i>E</i> Orthogonal	-9 ± 138	-7 ± 139	-7 ± 142	...	2 ± 45	2 ± 19	...
<i>T+E</i> 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
<i>T+E</i> Equilateral	19 ± 50	16 ± 50	15 ± 53	16 ± 45	-3 ± 14	-4 ± 19	-3.2 ± 9.8
<i>T+E</i> Orthogonal	-12 ± 25	-11 ± 26	-13 ± 25	-11 ± 23	1.9 ± 8.7	-1.0 ± 9.9	1.4 ± 5.9

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 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 $f_{\text{NL}}^{\text{local}} = 11.5 \pm 6.4$, $f_{\text{NL}}^{\text{equil}} = -7.5 \pm 51$, and $f_{\text{NL}}^{\text{ortho}} = -0.4 \pm 29$. 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 signal being significantly different from zero there (the Fisher error bars

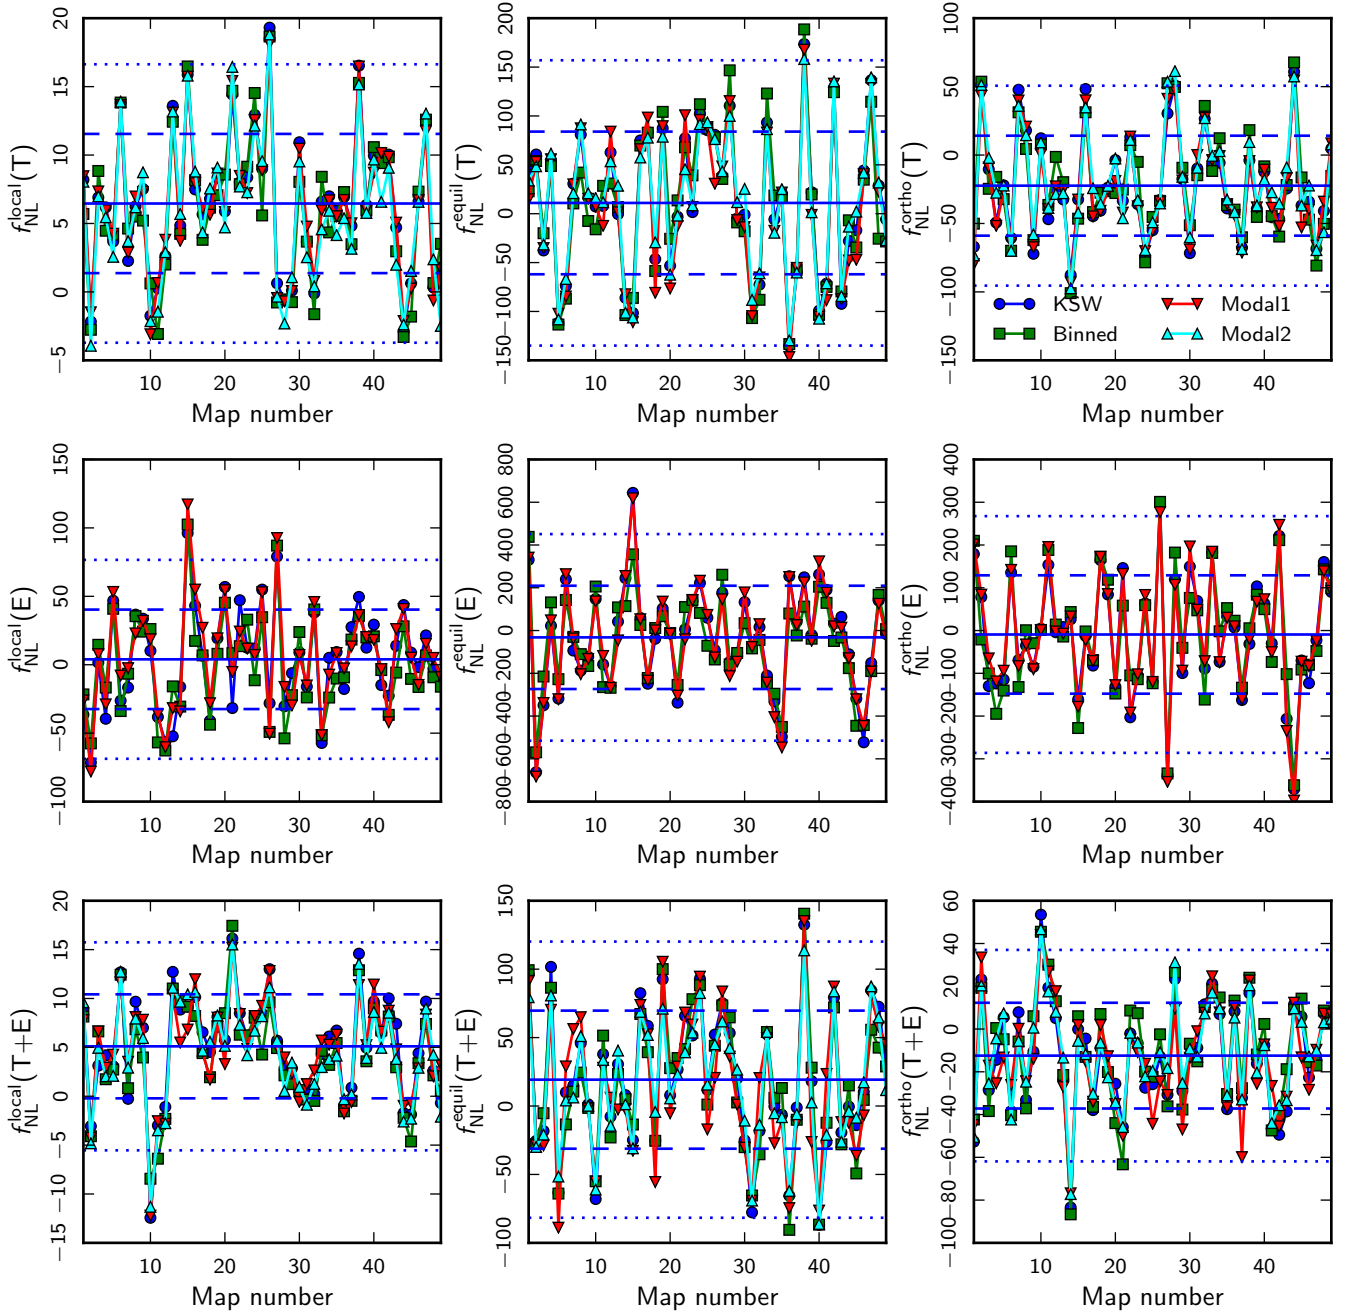


Fig. 3. Map-by-map comparison of the results from the different estimators for local (*left*), equilateral (*centre*), and orthogonal (*right*) f_{NL} (taking the shapes to be independent), for the third set of simulations described in Sect. 5.1, including both noise and a mask. Results are shown for 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.

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 determine 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 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

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.

Shape	f_{NL}			
	KSW	Binned	Modal 1	Modal 2
<i>T</i>				
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
<i>E</i>				
Local . . .	11 ± 29	12 ± 26	9 ± 36	...
Equilateral	34 ± 182	32 ± 153	10 ± 241	...
Orthogonal	-37 ± 110	-28 ± 115	-31 ± 143	...
<i>T+E</i>				
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

Notes. 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).

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.

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

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.

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 Sects. 3.4 and 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 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

Table 9. Results from the different estimators for f_{NL} for the set of SMICA simulations based on FFP8 described in Sect. 5.3.

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
Equilateral .	7 ± 144	7 ± 143	9 ± 152	0 ± 37	2 ± 35
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	...

Notes. 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.

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 estimator 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

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.

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 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

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.

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

Notes. 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.

benchmark or threshold against which to compare the other strong features observed in the bispectrum and now defined with greater precision. The original “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.

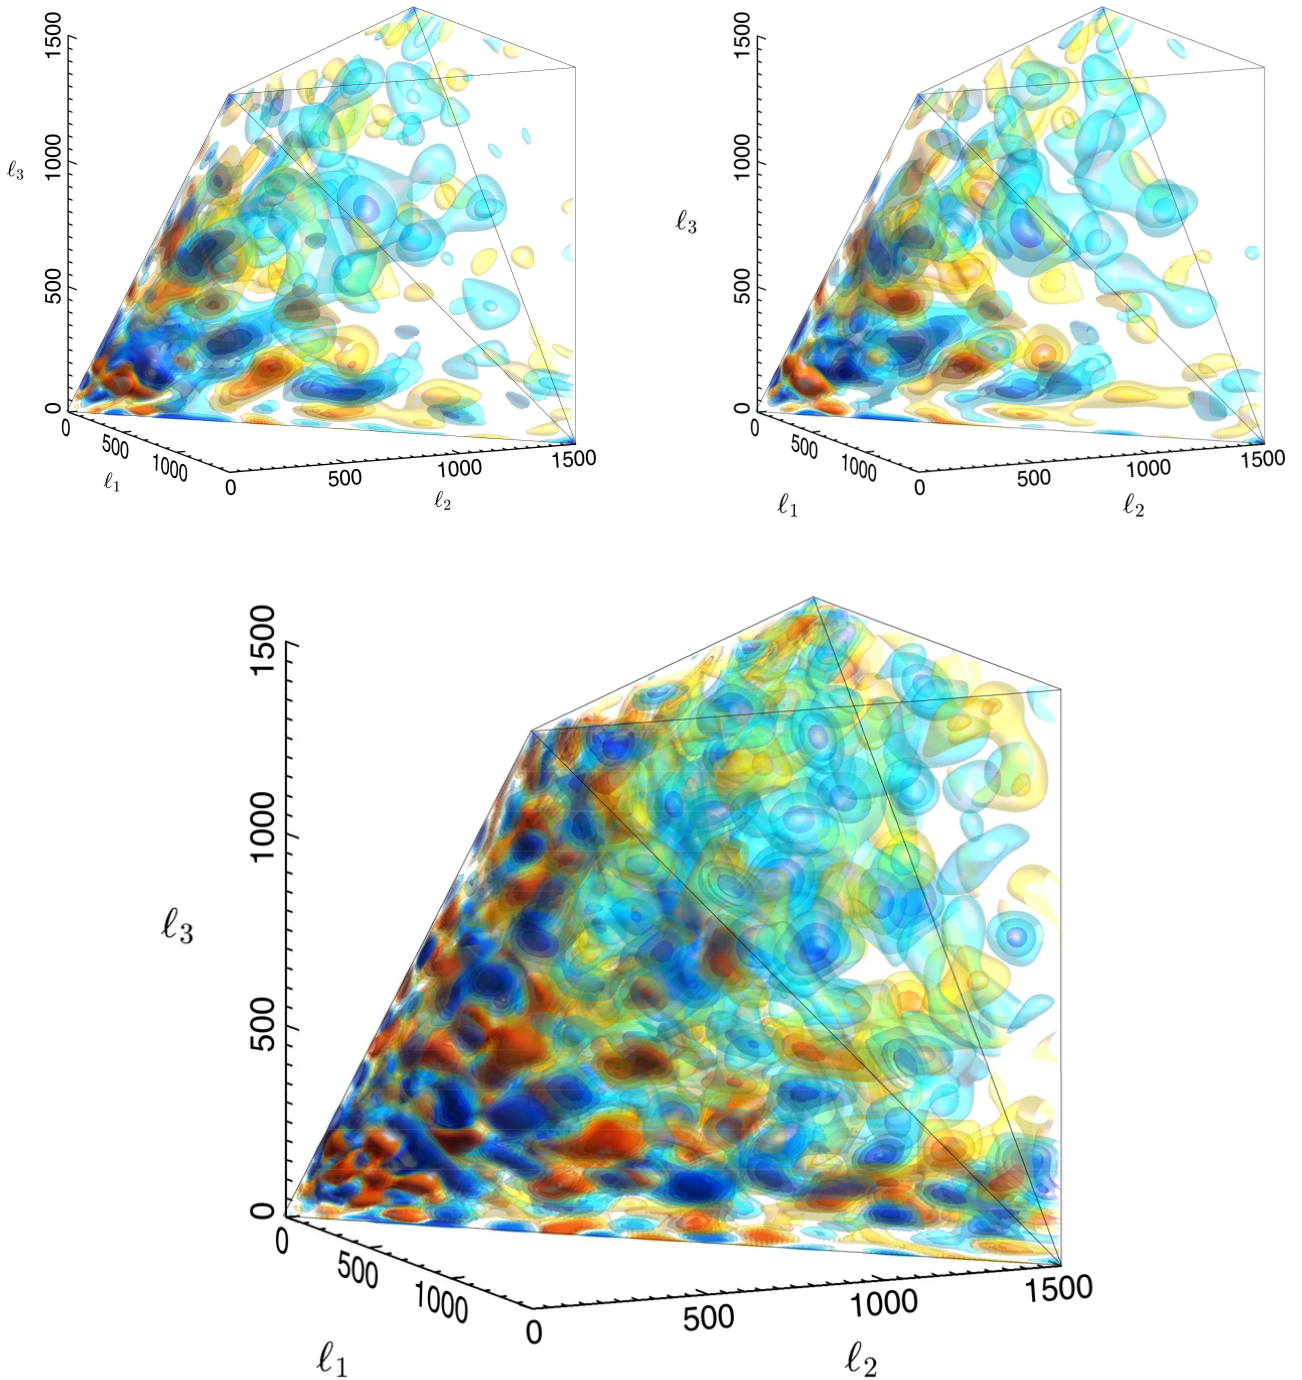


Fig. 4. Modal bispectrum reconstruction for *Planck* 2013 (*top left*) and 2015 (*top right*) temperature-only data, both using the SMICA maps. Here, we restrict the 2015 resolution to the same as 2013, using similar polynomials with $n_{\max} = 601$. The two bispectra are very close to being in complete 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).

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 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




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.

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

Notes. 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.

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. 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, $\mathcal{B}_{i_1 i_2 i_3}^{T2E}$ is not defined if both ℓ_1 and ℓ_2 are larger than 2000, while $\mathcal{B}_{i_1 i_2 i_3}^{TE2}$ 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




Fig. 6. CMB temperature and polarization bispectrum reconstructions for *Planck* SMICA maps using the full set of polynomial modes with $n_{\text{max}} = 2001$ 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.