Aalto University

Ade, P.A.R.; Aghanim, N.; Armitage-Caplan, C.; Arnaud, M.; Ashdown, M.; Atrio-Barandela, F.; Aumont, J.; Baccigalupi, C.; Banday, A.J.; Barreiro, R.B.; Bartlett, J.G.; Bartolo, N.; Battaner, E.; Benabed, K.; Benoît, A.; Benoit-Lévy, A.; Bernard, J.-P.; Bersanelli, M.; Bielewicz, P.; Bobin, J.; Bock, J.J.; Bonaldi, A.; Bonavera, L.; Bond, J.R.; Borrill, J.; Bouchet, F.R.; Bridges, M.; Bucher, M.; Burigana, C.; Butler, R.C.; Cardoso, J.-F.; Catalano, A.; Challinor, A.; Chamballu, A.; Chiang, H.C.; Chiang, L.-Y.; Christensen, P.R.; Church, S.; Clements, D.L.; Colombi, S.; Colombo, L.P.L.; Couchot, F.; Coulais, A.; Crill, B.P.; Curto, A.; Cuttaia, F.; Danese, L.; Davies, R.D.; Davis, R.J.; de Bernardis, P.

Planck 2013 results. XXIV. Constraints on primordial non-Gaussianity

Planck 2013 results

Astronomy & Astrophysics

Special feature

Planck 2013 results

XXIV. Constraints on primordial non-Gaussianity

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

(Affiliations can be found after the references)

ABSTRACT

The *Planck* nominal mission cosmic microwave background (CMB) maps yield unprecedented constraints on primordial non-Gaussianity (NG). Using three 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 $f_{NL}^{local} = 2.7 \pm 5.8$, $f_{NL}^{equil} = -42 \pm 75$, and $f_{NL}^{ortho} = -25 \pm 39$ (68% CL statistical). Non-Gaussianity is detected in the data; using skew-C_ℓ statistics we find a nonzero bispectrum from residual point sources, and the integrated-Sachs-Wolfe-lensing bispectrum at a level expected in the ΛCDM scenario. 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 confirmed by skew-C_ℓ, wavelet bispectrum and Minkowski functional estimators. Beyond estimates of individual shape amplitudes, we present model-independent, three-dimensional reconstructions of the *Planck* CMB bispectrum and thus derive constraints on early-Universe scenarios that generate primordial NG, including general single-field models of inflation, excited initial states (non-Bunch-Davies vacua), and directionally-dependent vector models. We provide an initial survey of scale-dependent feature and resonance models. These results bound both general single-field and multi-field model parameter ranges, such as the speed of sound, $c_s \geq 0.02$ (95% CL), in an effective field theory parametrization, and the curvaton decay fraction $r_D \geq 0.15$ (95% CL). The *Planck* data significantly limit the viable parameter space of the ekpyrotic/cyclic scenarios. The amplitude of the four-point function in the local model $\tau_{NL} < 2800$ (95% CL). Taken together, these constraints represent the highest precision tests to date of physical mechanisms for the origin of cosmic structure.

★ Corresponding author: N. Bartolo, e-mail: `nicola.bartolo@pd.infn.it`

1. Introduction

This paper, one of a set associated with the 2013 release of data from the *Planck*[1] mission (Planck Collaboration I 2014), describes the constraints on primordial non-Gaussianity (NG) obtained using the cosmic microwave background (CMB) maps derived from the data acquired by *Planck* during its nominal operations period, i.e., between 12 August 2009 and 27 November 2010.

Primordial NG is one of the most informative fingerprints of the origin of structure in the Universe, probing physics at extremely high energy scales inaccessible to laboratory experiments. Possible departures from a purely Gaussian distribution of the CMB anisotropies provide powerful observational access to this extreme physics (Allen et al. 1987; Salopek & Bond 1990; Falk et al. 1993; Gangui et al. 1994; Verde et al. 2000; Gangui & Martin 2000b; Wang & Kamionkowski 2000; Komatsu & Spergel 2001; Acquaviva et al. 2003; Maldacena 2003; Babich et al. 2004; for recent reviews Bartolo et al. 2004a; Liguori et al. 2010; Chen 2010b; Komatsu 2010; Yadav & Wandelt 2010). A robust detection of primordial NG – or a strong constraint on it – discriminates among competing mechanisms for the generation of the cosmological perturbations in the early Universe. Different inflationary models, firmly rooted in modern theoretical particle physics, predict different amplitudes, shapes, and scale dependence of NG. As a result, primordial NG is complementary to the scalar-spectral index of curvature perturbations and the tensor-to-scalar amplitude ratio, distinguishing between inflationary models that are degenerate on the basis of their power spectra alone. Even in the simplest models of inflation, consisting of a single slowly-rolling scalar field, a small (but calculable) level of NG is predicted (Acquaviva et al. 2003; Maldacena 2003); this is undetectable in present-quality CMB and large-scale structure measurements. However, as demonstrated by a large body of work in recent years, extending this simplest paradigm will generically lead to detectable levels of NG in CMB anisotropies. Critically, a robust detection of primordial NG would rule out all canonical single-field slow-roll models of inflation, pointing to physics beyond the simplest "textbook" picture of inflation. Conversely, significant improvements in the constraints on primordial NG strongly limit extensions to the simplest paradigm, thus providing powerful clues to the physical mechanism that generated cosmic structure.

If the primordial fluctuations are Gaussian-distributed, then they are completely characterised by their two-point correlation function, or equivalently, their power spectrum. If they are non-Gaussian, there is additional statistical information in the higher-order correlation functions, which is not captured by the two-point correlation function. In particular, the 3-point correlation function, or its Fourier counterpart, the bispectrum, is important because it is the lowest-order statistic that can distinguish between Gaussian and non-Gaussian perturbations. One of the main goals of this paper is to constrain the amplitude and shape of primordial NG using the angular bispectrum of the CMB anisotropies. The CMB angular bispectrum is related to the primordial bispectrum defined by

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

Here we define the potential Φ in terms of the comoving curvature perturbation ζ on super-horizon scales by $\Phi \equiv (3/5)\zeta$. In matter domination, on super-horizon scales, Φ is equivalent to Bardeen's gauge-invariant gravitational potential (Bardeen 1980), and we adopt this notation for historical consistency. The bispectrum $B_\Phi(k_1, k_2, k_3)$ measures the correlation among three perturbation modes. Assuming translational and rotational invariance, it depends only on the magnitudes of the three wavevectors. In general the bispectrum can be written as

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

Here, $f_{\rm NL}$ is the so-called "nonlinearity parameter" (Gangui et al. 1994; Wang & Kamionkowski 2000; Komatsu & Spergel 2001; Babich et al. 2004), a dimensionless parameter measuring the amplitude of NG. The bispectrum is measured 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 of the bispectrum (Babich et al. 2004), which encodes much physical information. It can also encode the scale dependence, i.e., the running, of the bispectrum (Chen 2005c)[2]. Different NG shapes are linked to distinctive physical mechanisms that can generate such non-Gaussian fingerprints in the early Universe. For example, the so-called "local" NG (Gangui et al. 1994; Verde et al. 2000; Wang & Kamionkowski 2000; Komatsu & Spergel 2001) is characterized by a signal that is maximal for "squeezed" triangles with $k_1 \ll k_2 \simeq k_3$ (or permutations; Maldacena 2003) which occurs, in general, when the primordial NG is generated on super-horizon scales. Conversely, "equilateral" NG (Babich et al. 2004) peaks for equilateral configurations $k_1 \approx k_2 \approx k_3$, due to correlations between fluctuation modes that are of comparable wavelengths, which can occur if the three perturbation modes mostly interact when they cross the horizon approximately at the same time. Other relevant shapes include the so-called "folded" (or flattened) NG (Chen et al. 2007b), which is due to correlations between perturbation modes that are enhanced for $k_1 + k_2 \approx k_3$, or the "orthogonal" NG (Senatore et al. 2010) that generates a signal with a positive peak at the equilateral configuration and a negative peak at the folded configuration.

We now sketch how non-Gaussian information in the initial conditions is transferred to observable quantities (in this instance, the CMB anisotropies) in the context of inflation. Primordial perturbations in the inflaton field(s) $\phi(\boldsymbol{x}, t) = \phi_0(t) + \delta\phi(\boldsymbol{x}, t)$ (where $\delta\phi$ denotes quantum fluctuations about the background value $\phi_0(t)$) can be characterized by the comoving curvature perturbation ζ, since this is conserved on super-horizon scales for adiabatic perturbations. The inflaton fluctuations $\delta\phi$ (in the flat gauge) induce a curvature perturbation[3] $\zeta = -(H/\dot\phi_0)\,\delta\phi$ at linear order; however, nonlinearities induce corrections to this relation. The primordial NG in the curvature perturbation ζ is intrinsically nonlinear, so that its contribution to the CMB anisotropies is transferred linearly at leading

---

[2] Specifically, one can define the shape of the bispectrum as the dependence of $F(k_1, k_2, k_3)(k_1 k_2 k_3)^2$ on the ratios of momenta, e.g., (k_2/k_1) and (k_3/k_1), once the overall scale of the triangle $K = k_1 + k_2 + k_3$ is fixed. The scale dependence of the bispectrum can be characterized by the dependence of $F(k_1, k_2, k_3)(k_1 k_2 k_3)^2$ on the overall scale K, once the ratios (k_2/k_1) and (k_3/k_1) are fixed (see, e.g., Chen 2010b).

[3] For the curvature perturbation, we follow the notation and sign conventions of Komatsu et al. (2011). ζ is also sometimes denoted $\mathcal{R}$ (see e.g., Lidsey et al. 1997; Lyth & Riotto 1999, and references therein), while the comoving curvature perturbation $\mathcal{R}$ as defined, e.g., in Malik & Wands (2009) is such that $\mathcal{R} = -\zeta$.

order. In particular, at the linear level, the curvature perturbation ζ is related to Bardeen's gravitational potential Φ during the matter-dominated epoch by $\Phi = (3/5)\zeta$ and $\Delta T/T \sim g\,\zeta$, where g is the linear radiation transfer function; thus, any primordial NG will be transferred to the CMB even at linear order. For example, in the large-angular scale limit, the linear Sachs-Wolfe effect reads $\Delta T/T = -\Phi/3 = -\zeta/5$. Further, any other field excited during the inflationary phase which develops quantum fluctuations contributing to the primordial curvature perturbation – whether or not it is driving inflation – can leave its non-Gaussian imprint in the CMB anisotropies.

Thus the bispectrum of Eq. (1) measures the fundamental (self-) interactions of the scalar field(s) involved in the inflationary phase and/or generating the primordial curvature perturbation, as well as measuring nonlinear processes occurring during or after inflation. It therefore brings insights into the fundamental physics behind inflation, possibly allowing for the first time a reconstruction of the inflationary Lagrangian itself. For example, in a large class of inflationary models which involve additional light field(s) different from the inflaton, the super-horizon evolution of the fluctuations in the additional field(s) and their transfer to the adiabatic curvature perturbations can generate a large primordial NG of the local type. This is the case for curvaton-type models (Linde & Mukhanov 1997; Lyth & Wands 2002; Lyth et al. 2003) where the late-time decay of a scalar field, belonging to the non-inflationary sector of the theory, induces curvature perturbations; models where the curvature perturbation is generated by the local fluctuations of the inflaton's coupling to matter during the reheating phase (Kofman 2003; Dvali et al. 2004a); and multi-field models of inflation (see, e.g., Bartolo et al. 2002; Bernardeau & Uzan 2002; Vernizzi & Wands 2006; Rigopoulos et al. 2006, 2007; Lyth & Rodriguez 2005; Byrnes & Choi 2010). Since the nonlinear processes take place on super-horizon scales, the form of NG is local in real space and thus, in Fourier space, the bispectrum correlates large and small Fourier modes. "Equilateral" NG (Babich et al. 2004) is a generic feature of single-field models with a non-canonical kinetic term, which can also generate the "orthogonal" type of NG (Senatore et al. 2010). In general, these models are characterized by higher-derivative interactions of the inflaton field. The correlation between the fluctuation modes is suppressed when one of the modes is on super-horizon scales, because the derivative terms are redshifted away, so that the correlation is maximal for three modes of comparable wavelengths that cross the horizon at the same time. An example of "folded" NG is the one generated in a class of single-field models with non-Bunch-Davies vacuum (Chen et al. 2007b; Holman & Tolley 2008). Indeed, these and other types of primordial NG can also be produced in other models, and we refer to Sect. 2 for more details. All these models can easily yield primordial NG with an amplitude much bigger than the one predicted in the standard models of single-field slow-roll inflation, for which the NG amplitude turns out to be proportional to the usual slow-roll parameters $f_{NL} \sim O(\epsilon, \eta)$ (Acquaviva et al. 2003; Maldacena 2003).

Given that a robust detection of primordial NG would represent a breakthrough in the understanding of the physics governing the Universe during its very first stages, it is crucial that all sources of contamination are sufficiently understood to firmly control their effects. In particular, any nonlinearity in the post-inflationary Universe can introduce NG into perturbations that were initially Gaussian. Therefore, one must ensure that a primordial origin is not ascribed to a non-primordial contaminant; however, estimators of (primordial) NG from CMB data will also typically be sensitive to such contaminating signals. Potential

non-primordial sources of NG can be classified into four broad categories: instrumental systematic effects (see e.g., Donzelli et al. 2009); residual foregrounds and point sources; secondary CMB anisotropies, such as the Sunyaev-Zeldovich (SZ) effect (Zeldovich & Sunyaev 1969), gravitational lensing (see Lewis & Challinor 2006, for a review), the integrated Sachs-Wolfe (ISW) effect (Sachs & Wolfe 1967) or the Rees-Sciama effect (Rees & Sciama 1968); and effects arising from nonlinear (second-order) perturbations in the Boltzmann equations (due to the nonlinear nature of General Relativity and the nonlinear dynamics of the photon-baryon fluid at recombination). Among the secondary anisotropies, the cross-correlation of the ISW/Rees-Sciama and lensing (Goldberg & Spergel 1999) produces the dominant contamination to the (local) primordial NG. The impact is mainly on the local type of primordial NG, because the ISW-lensing correlation couples the large-scale gravitational potential fluctuations sourcing the ISW effect with the small-scale lensing effects of the CMB, thus producing a bispectrum which peaks on the squeezed configurations, as for the local shape. Detailed analyses have shown that the ISW-lensing bispectrum can introduce a bias to local primordial NG, while the bias to equilateral primordial NG is negligible (see Serra & Cooray 2008; Smith & Zaldarriaga 2011; Hanson et al. 2009b; Lewis et al. 2011, 2012; Mangilli & Verde 2009; Junk & Komatsu 2012; Mangilli et al. 2013). In our analysis we have carefully accounted for this effect (we report the values of the ISW-lensing bias in Sect. 5.2, and demonstrate the detection of the effect with skew-C_ℓs), as well as validating our results through an extensive suite of simulations and null tests in order to quantify the effects of systematic effects and diffuse and point-source foregrounds. Finally, a consistent treatment of weak NG in the CMB must account for additional contributions that arise at the nonlinear (second-order) level both in the gravitational perturbations after inflation ends, and for the evolution of the CMB anisotropies at second-order in perturbation theory at large and small angular scales. It has been shown that these second-order CMB effects yield negligible contamination to primordial NG for *Planck*-quality data (Bartolo et al. 2004b, 2005, 2010c, 2012; Creminelli & Zaldarriaga 2004b; Boubekeur et al. 2009; Nitta et al. 2009; Senatore et al. 2009; Khatri & Wandelt 2009; Bartolo & Riotto 2009; Khatri & Wandelt 2010; Bartolo et al. 2010c; Creminelli et al. 2011b; Bartolo et al. 2012; Huang & Vernizzi 2013; Su et al. 2012; Pettinari et al. 2013).

Previous constraints on various shapes of primordial NG come from the WMAP-9 data (Bennett et al. 2013). For the local shape they find $f_{NL}^{local} = 37 \pm 20$ (68% CL). For equilateral-type NG, they obtain $f_{NL}^{equil} = 51 \pm 136$ (68% CL), while for the orthogonal shape $f_{NL}^{ortho} = -245 \pm 100$ (68% CL). Other analyses employing different estimators give compatible constraints. Limits on other shapes, such as e.g. flattened and feature models, have also been obtained (Fergusson et al. 2012).

Before concluding this section let us point out the connection between the analyses presented here and in the companion paper Planck Collaboration XXIII (2014) on the statistical and isotropy properties of the CMB. Statistical anisotropy and NG are essentially two alternative descriptions of the same phenomenon on the sky (Ferreira & Magueijo 1997). Specifically any Gaussian but statistically anisotropic model becomes, after averaging over the possible (a priori unknown) orientations of the anisotropy, a statistically isotropic non-Gaussian model. For example local NG can be generated by large-scale field fluctuations that couple to the small-scale power. For the given fixed realization of large-scale modes that we see, the small-scale anisotropies

look anisotropic on the sky, and it is equally valid to describe this as a Gaussian anisotropic model (assuming the large-scale modes are Gaussian). In this paper we mostly focus on the non-Gaussian interpretation of various physically motivated models, although it is useful to bear both perspectives in mind, in particular when considering what forms of non-primordial signal might cause contamination. Planck Collaboration XXIII (2014) consider a broad class of more general phenomenological forms of anisotropy, which are complementary to the analysis presented here.

This paper is organized as follows. In Sect. 2, we present models generating primordial NG that have been tested in this paper. Section 3 summarizes the statistical estimators used to constrain the CMB bispectrum from Planck data and the methods for the reconstruction of the CMB bispectrum. Section 4 summarizes the statistical estimator used to constrain the CMB trispectrum. In Sect. 5, 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. Section 6 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. Section 7 contains our main results: we present constraints on f_{NL} for the local, equilateral, and orthogonal bispectra, and a selected set of other bispectrum shapes; we show a reconstruction of the CMB bispectrum, and give limits on the CMB trispectrum. In Sect. 8 we validate these results by performing a series of null tests on the data to assess the robustness of our results. We also evaluate the impact of the Planck data processing on the primordial NG signal. In Sect. 9, we discuss the main implications of Planck's constraints on primordial NG for early Universe models. We conclude in Sect. 10. The realistic Planck simulations used in various steps of the analysis and validation tests are described in Appendix A. Appendix B contains a derivation of the expected scatter between f_{NL} results on the same map from different estimators used in the validation tests of Sect. 6, while Appendix C presents a comparison of constraints on some selected non-standard bispectrum shapes using different foreground-cleaned maps.

2. Inflationary models for primordial non-Gaussianity

There is a simple reason why standard single-field models of slow-roll inflation predict a tiny level of NG, of the order of the usual slow-roll parameters $f_{NL} \sim O(\epsilon, \eta)$[4]: in order to achieve an accelerated period of expansion, the inflaton potential must be very flat, thus suppressing the inflaton (self-)interactions and any sources of nonlinearity, and leaving only its weak gravitational interactions as the main source of NG. This fact leads to

---

[4] This has been shown in the pioneering research which demonstrated that perturbations produced in single-field models of slow-roll inflation are characterized by a low-amplitude NG (Salopek & Bond 1990; Falk et al. 1993; Gangui et al. 1994). Later Acquaviva et al. (2003) and Maldacena (2003) obtained a complete quantitative prediction for the nonlinearity parameter in single-field slow-roll inflation models, also showing that the predicted NG is characterized by a shape dependence which is more complex than suggested by previous results expressed in terms of the simple parameterization $\Phi(\boldsymbol{x}) = \Phi_L(\boldsymbol{x}) + f_{NL}\Phi_L^2(\boldsymbol{x})$ (Gangui et al. 1994; Verde et al. 2000; Wang & Kamionkowski 2000; Komatsu & Spergel 2001), where Φ_L is the linear gravitational potential.

a clear distinction between the simplest models of inflation, and scenarios where a significant amplitude of NG can be generated (e.g., Komatsu 2010), as follows. The simplest inflationary models are based on a set of minimal conditions: (i) a single weakly-coupled neutral single scalar field (the inflaton, which drives inflation and generates the curvature perturbations); (ii) with a canonical kinetic term; (iii) slowly rolling down its (featureless) potential; (iv) initially lying in a Bunch-Davies (ground) vacuum state. In the last few years, an important theoretical realization has taken place: a detectable amplitude of NG with specific triangular configurations (corresponding broadly to well-motivated classes of physical models) can be generated if any one of the above conditions is violated (Bartolo et al. 2004a; Liguori et al. 2010; Chen 2010b; Komatsu 2010; Yadav & Wandelt 2010):

- "local" NG, where the signal peaks in "squeezed" triangles ($k_1 \ll k_2 \simeq k_3$) (e.g., multi-field models of inflation);
- "equilateral" NG, peaking for $k_1 \approx k_2 \approx k_3$. Examples of this class include single-field models with non-canonical kinetic term (Chen et al. 2007b), such as k-inflation (Armendariz-Picon et al. 1999; Chen et al. 2007b) or Dirac-Born-Infield (DBI) inflation (Silverstein & Tong 2004; Alishahiha et al. 2004), models characterized by more general higher-derivative interactions of the inflaton field, such as ghost inflation (Arkani-Hamed et al. 2004), and models arising from effective field theories (Cheung et al. 2008);
- "folded" (or flattened) NG. Examples of this class include: single-field models with non-Bunch-Davies vacuum (Chen et al. 2007b; Holman & Tolley 2008) and models with general higher-derivative interactions (Senatore et al. 2010; Bartolo et al. 2010a);
- "orthogonal" NG which is generated, e.g., in single-field models of inflation with a non-canonical kinetic term (Senatore et al. 2010), or with general higher-derivative interactions.

All these models naturally predict values of $|f_{NL}| \gg 1$. A detection of such a signal would rule out the simplest models of single-field inflation, which, obeying all the conditions above, are characterized by weak gravitational interactions with $|f_{NL}| \ll 1$.

The above scheme provides a general classification of inflationary models in terms of the corresponding NG shapes, which we adopt for the data analysis presented in this paper:

1. "general" single-field inflationary models (tested using the equilateral, orthogonal and folded shapes);
2. multi-field models of inflation (tested using the local shape).

In each class, there exist specific realizations of inflationary models which are characterized by the same underlying physical mechanism, generating a specific NG shape. We will investigate these classes of inflationary models by constraining the corresponding NG content, focusing on amplitudes and shapes. We also perform a survey of non-standard models giving rise to alternative specific shapes of NG. Different NG shapes are observationally distinguishable if their cross-correlation is sufficiently low; almost all of the shapes analysed in this paper are highly orthogonal to each other (e.g., Babich et al. 2004; Fergusson & Shellard 2007).

There are exceptional cases which evade this classification: for example, some exotic non-local single-field theories of inflation produce local NG (Barnaby & Cline 2008), while some multi-field models can produce equilateral NG, e.g., if some particle production mechanism is present (examples include trapped

inflation Green et al. 2009, and some models of axion inflation Barnaby & Peloso 2011; Barnaby et al. 2011, 2012b). Another example arises in a class of multi-field models where the second scalar field is not light, but has a mass $m \approx H$, of the order of the Hubble rate during inflation. Then NG with an intermediate shape, interpolating between local and equilateral, can be produced – "quasi-single field" models of inflation (Chen & Wang 2010a,b) – for which the NG shape is similar to the so-called constant NG of Fergusson & Shellard (2007). Furthermore, there is the possibility of a superposition of shapes (and/or running of NG), generated if different mechanisms sourcing NG act simultaneously during the inflationary evolution. For example, in multi-field DBI inflation, equilateral NG is generated at horizon crossing from the higher-derivative interactions of the scalar fields, and it adds to the local NG arising from the super-horizon nonlinear evolution (e.g., Langlois et al. 2008a,b; Arroja et al. 2008; Renaux-Petel 2009).

In the following subsections, we discuss each of these possibilities in turn. The reader already familiar with this background material may skip to Sect. 3.

2.1. General single-field models of inflation

Typically in models with a non-standard kinetic term (or more general higher-derivative interactions), inflaton perturbations propagate with an effective sound speed c_s which can be smaller than the speed of light, and this results in a contribution to the NG amplitude $f_{\mathrm{NL}} \sim c_{\mathrm{s}}^{-2}$ in the limit $c_{\mathrm{s}} \ll 1$. 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 $c_{\mathrm{s}}^2 = (\partial P / \partial X)/(\partial P / \partial X + 2X(\partial^2 P / \partial X^2))$.

In this case, two interaction terms give the dominant contribution to primordial NG, one of the type $(\dot{\delta\phi})^3$ and the other of the type $\dot{\delta\phi}(\nabla \delta\phi)^2$, which arise from expanding the $P(X, \phi)$ Lagrangian. Each of these two interaction terms generates a bispectrum with a shape similar to the equilateral type, with the second inflaton interaction yielding a nonlinearity parameter $f_{\mathrm{NL}} \approx c_{\mathrm{s}}^{-2}$, independent of the amplitude of the other bispectrum. Equilateral NG is usually generated by derivative interactions of the inflaton field; derivative terms are suppressed when one perturbation mode is frozen on super-horizon scales during inflation, and the other two are still crossing the horizon, so that the correlation between the three perturbation modes will be suppressed, while it is maximal when all the three modes cross the horizon at the same time, which happens for $k_1 \approx k_2 \approx k_3$.

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

$$B_\Phi^{\mathrm{equil}}(k_1, k_2, k_3) = 6A^2 f_{\mathrm{NL}}^{\mathrm{equil}}$$

$$\times \left\{ -\frac{1}{k_1^{4-n_{\mathrm{s}}} k_2^{4-n_{\mathrm{s}}}} - \frac{1}{k_2^{4-n_{\mathrm{s}}} k_3^{4-n_{\mathrm{s}}}} - \frac{1}{k_3^{4-n_{\mathrm{s}}} k_1^{4-n_{\mathrm{s}}}} - \frac{2}{(k_1 k_2 k_3)^{2(4-n_{\mathrm{s}})/3}} \right.$$
$$\left. + \left[\frac{1}{k_1^{(4-n_{\mathrm{s}})/3} k_2^{2(4-n_{\mathrm{s}})/3} k_3^{4-n_{\mathrm{s}}}} + (5 \text{ permutations}) \right] \right\}, \quad (3)$$

where $P_\Phi(k) = A/k^{4-n_s}$ is the power spectrum of Bardeen's gravitational potential with normalization A^2 and scalar spectral index n_s. For example, the models introduced in the string theory framework based on the DBI action (Silverstein & Tong 2004; Alishahiha et al. 2004) can be described within the $P(X, \phi)$-class, and they give rise to an equilateral NG with an overall

amplitude $f_{\mathrm{NL}}^{\mathrm{equil}} = -(35/108)c_{\mathrm{s}}^{-2}$ for $c_{\mathrm{s}} \ll 1$, which turns out typically to be $f_{\mathrm{NL}}^{\mathrm{equil}} < -5$[5].

The equilateral shape emerges also in models characterized by more general higher-derivative interactions, such as ghost inflation (Arkani-Hamed et al. 2004) or models within effective field theories of inflation (Cheung et al. 2008; Senatore et al. 2010; Bartolo et al. 2010a).

Taken individually, each higher-derivative interaction of the inflaton field generically gives rise to a bispectrum with a shape which is similar – but not identical to – the equilateral form (an example is provided by the two interaction terms discussed above for an inflaton with a non-standard kinetic term). Therefore it has been shown, using an effective field theory approach to inflationary perturbations, that it is possible to build a combination of the corresponding similar equilateral shapes to generate a bispectrum that is orthogonal to the equilateral one, the so-called "orthogonal" shape. This can be approximated by the template (Senatore et al. 2010)

$$B_\Phi^{\mathrm{ortho}}(k_1, k_2, k_3) = 6A^2 f_{\mathrm{NL}}^{\mathrm{ortho}}$$

$$\times \left\{ -\frac{3}{k_1^{4-n_{\mathrm{s}}} k_2^{4-n_{\mathrm{s}}}} - \frac{3}{k_2^{4-n_{\mathrm{s}}} k_3^{4-n_{\mathrm{s}}}} - \frac{3}{k_3^{4-n_{\mathrm{s}}} k_1^{4-n_{\mathrm{s}}}} - \frac{8}{(k_1 k_2 k_3)^{2(4-n_{\mathrm{s}})/3}} \right.$$
$$\left. + \left[\frac{3}{k_1^{(4-n_{\mathrm{s}})/3} k_2^{2(4-n_{\mathrm{s}})/3} k_3^{4-n_{\mathrm{s}}}} + (5 \text{ perm.}) \right] \right\}. \quad (4)$$

The orthogonal bispectrum can also arise as the predominant shape in some inflationary realizations of Galileon inflation (Renaux-Petel et al. 2011).

Non-separable single-field bispectrum shapes: while most single-field inflation bispectra can be well-characterized by the equilateral and orthogonal shapes, we note that these are separable ansätze which only approximate the contributions from two leading order terms in the cubic Lagrangian. In an effective field theory approach these correspond to two shapes which can be associated directly with the inflaton field interactions $\dot{\pi}(\partial_i \pi)^2$ and $\dot{\pi}^3$ (in the language of the effective field theory of inflation the inflaton scalar degree of freedom π is related to the comoving curvature perturbation as $\zeta = -H\pi$). They are, respectively (Senatore et al. 2010; see also Chen et al. 2007b[6])

$$B_\Phi^{\mathrm{EFT1}}(k_1, k_2, k_3) = \frac{6A^2 f_{\mathrm{NL}}^{\mathrm{EFT1}}}{(k_1 k_2 k_3)^3} \frac{(-9/17)}{(k_1 + k_2 + k_3)^3}$$

$$\times \left\{ \sum_i k_i^6 + \sum_{i \neq j} \left[3k_i^5 k_j - k_i^4 k_j^2 - 3k_i^3 k_j^3 \right] \right. \quad (5)$$

$$\left. + \sum_{i \neq j \neq l} \left[3k_i^4 k_j k_l - 9k_i^3 k_j^2 k_l - 2k_i^2 k_j^2 k_l^2 \right] \right\},$$

$$B_\Phi^{\mathrm{EFT2}}(k_1, k_2, k_3) = \frac{6A^2 f_{\mathrm{NL}}^{\mathrm{EFT2}}}{k_1 k_2 k_3} \frac{27}{(k_1 + k_2 + k_3)^3}. \quad (6)$$

---

[5] An effectively single-field model with a non-standard kinetic term and a reduced sound speed for the adiabatic perturbation modes might also arise in coupled multi-field systems, where the heavy fields are integrated out: see discussions in, e.g., Tolley & Wyman (2010), Achúcarro et al. (2011), Shiu & Xu (2011).

[6] Notice that the two shapes (5) and (6) correspond to a linear combination of the two shapes found in Chen et al. (2007b).

These shapes differ from equilateral in the flattened or collinear limit. DBI inflation gives a closely related shape of particular interest phenomenologically (Alishahiha et al. 2004),

$$B_\Phi^{\mathrm{DBI}}(k_1, k_2, k_3) = \frac{6A^2 f_{\mathrm{NL}}^{\mathrm{DBI}}}{(k_1 k_2 k_3)^3} \frac{(-3/7)}{(k_1 + k_2 + k_3)^2} \tag{7}$$

$$\times \left\{ \sum_i k_i^5 + \sum_{i \neq j} \left[2k_i^4 k_j - 3k_i^3 k_j^2 \right] + \sum_{i \neq j \neq l} \left[k_i^3 k_j k_l - 4k_i^2 k_j^2 k_l \right] \right\}.$$

For brevity, we have given the scale-invariant form of the shape functions, without the mild power spectrum running. There are also sub-leading order terms which give rise to additional non-separable shapes, but these are expected to be much smaller without special fine-tuning.

2.2. Multi-field models

This class of models generally includes an additional light scalar field (or more fields) during inflation, which can be different from the inflaton, and whose fluctuations contribute to the final primordial curvature perturbation of the gravitational potential. It could be the case of inflation driven by several scalar fields – "multiple-field inflation" – or the one where the inflaton drives the accelerated expansion, while other scalar fields remain subdominant during inflation. This encompasses, for instance, a large class of multi-field models which leads to non-Gaussian isocurvature perturbations (for earlier works, see e.g., Linde & Mukhanov 1997; Peebles 1997; Bucher & Zhu 1997). More importantly, such models can also lead to cross-correlated and non-Gaussian adiabatic and isocurvature modes, where NG is first generated by large nonlinearities in some scalar (possibly non-inflatonic) sector of the theory, and then efficiently transferred to the inflaton adiabatic sector(s) through the cross-correlation of adiabatic and isocurvature perturbations[7] (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). Another interesting possibility is the curvaton model (Mollerach 1990; Enqvist & Sloth 2002; Lyth & Wands 2002; Moroi & Takahashi 2001), where a second light scalar field, subdominant during inflation, decays after inflation ends, producing the primordial density perturbations which can be characterized by a high NG level (e.g., Lyth & Wands 2002; Lyth et al. 2003; Bartolo et al. 2004d). NG in the curvature perturbation can be generated at the end of inflation, e.g., due to the nonlinear dynamics of (p)reheating (e.g., Enqvist et al. 2005; Chambers & Rajantie 2008; Barnaby & Cline 2006; see also Bond et al. 2009) or, as in modulated (p)reheating and modulated hybrid inflation, due to local fluctuations in the decay rate/interactions of the inflaton field (Kofman 2003; Dvali et al. 2004a,b; Bernardeau et al. 2004; Zaldarriaga 2004; Lyth 2005; Salem 2005; Lyth & Riotto 2006; Kolb et al. 2006; Cicoli et al. 2012). The common feature of all these models is that a large NG in the curvature perturbation can be produced via both a transfer of super-horizon non-Gaussian isocurvature perturbations in the second field (not necessarily the inflaton) to

the adiabatic density perturbations, and via additional nonlinearities in the transfer mechanism. Since, typically, this process occurs on super-horizon scales, the form of NG is local in real space. Being local in real space, the bispectrum correlates large and small scale Fourier modes. The local bispectrum is given by (Falk et al. 1993; Gangui et al. 1994; Verde et al. 2000; Wang & Kamionkowski 2000; Komatsu & Spergel 2001)

$$B_\Phi^{\mathrm{local}}(k_1, k_2, k_3) = 2 f_{\mathrm{NL}}^{\mathrm{local}} \Big[P_\Phi(k_1) P_\Phi(k_2) + P_\Phi(k_1) P_\Phi(k_3)$$

$$+ P_\Phi(k_2) P_\Phi(k_3) \Big]$$

$$= 2 A^2 f_{\mathrm{NL}}^{\mathrm{local}} \left[\frac{1}{k_1^{4-n_s} k_2^{4-n_s}} + \mathrm{cycl.} \right]. \tag{8}$$

Most of the signal-to-noise ratio in fact peaks in the squeezed configurations ($k_1 \ll k_2 \simeq k_3$)

$$B_\Phi^{\mathrm{local}}(k_1 \to 0, k_2, k_3) \to 4 f_{\mathrm{NL}}^{\mathrm{local}} P_\Phi(k_1) P_\Phi(k_2). \tag{9}$$

The typical example of a curvature perturbation that generates the bispectrum of Eq. (8) is the standard local form for the gravitational potential (Hodges et al. 1990; Kofman et al. 1991; Salopek & Bond 1990; Gangui et al. 1994; Verde et al. 2000; Wang & Kamionkowski 2000; Komatsu & Spergel 2001)

$$\Phi(\boldsymbol{x}) = \Phi_L(\boldsymbol{x}) + f_{\mathrm{NL}}^{\mathrm{local}} (\Phi_L^2(\boldsymbol{x}) - \langle \Phi_L^2(\boldsymbol{x}) \rangle), \tag{10}$$

where $\Phi_L(\boldsymbol{x})$ is the linear Gaussian gravitational potential and $f_{\mathrm{NL}}^{\mathrm{local}}$ is the amplitude of a quadratic nonlinear correction (though this is not the only possibility: e.g., the gravitational potential produced in multiple-field inflation models generally cannot be reduced to the Eq. (10)). For example, in the (simplest) adiabatic curvaton models, the NG amplitude turns out to be (Bartolo et al. 2004d,c) $f_{\mathrm{NL}}^{\mathrm{local}} = (5/4r_{\mathrm{D}}) - 5r_{\mathrm{D}}/6 - 5/3$, for a quadratic potential of the curvaton field (Lyth & Wands 2002; Lyth et al. 2003; Lyth & Rodriguez 2005; Malik & Lyth 2006; Sasaki et al. 2006), where $r_{\mathrm{D}} = [3\rho_{\mathrm{curvaton}}/(3\rho_{\mathrm{curvaton}} + 4\rho_{\mathrm{radiation}})]_{\mathrm{D}}$ is the "curvaton decay fraction" evaluated at the epoch of the curvaton decay in the sudden decay approximation. Therefore, for $r_{\mathrm{D}} \ll 1$, a high level of NG is imprinted.

There exists a clear distinction between multi-field and single-field models of inflation that can be probed via a consistency condition (Maldacena 2003; Creminelli & Zaldarriaga 2004a; Chen et al. 2007b; Chen 2010b): in the squeezed limit, single-field models predict a bispectrum

$$B_\Phi^{\mathrm{single-field}}(k_1 \to 0, k_2, k_3 = k_2) \to \frac{5}{3}(1 - n_s) P_\Phi(k_1) P_\Phi(k_2), \tag{11}$$

and thus $f_{\mathrm{NL}} \sim O(n_s - 1)$ in the squeezed limit, in a model-independent sense (i.e., not only for standard single-field models). This means that a significant detection of local NG (in the squeezed limit) would rule out a very large class of single-field models of inflation (not just the simplest ones). Although based on very general conditions, the consistency condition of Eq. (11) can be violated in some well-motivated inflationary settings (we refer the reader to Chen 2010b; Chen et al. 2013 and references therein for more details).

Quasi-single field inflation: quasi-single field inflation has an extra field (or fields) with mass m close to the Hubble parameter H during inflation; these models evolve quiescently, producing a calculable non-Gaussian signature (Chen & Wang 2010b). The resulting one-parameter bispectrum smoothly interpolates

---

[7] This may happen, for instance, if the inflaton field is coupled to the other scalar degrees of freedom, as expected on particle physics grounds. These scalar degrees of freedom may have large self-interactions, so that their quantum fluctuations are intrinsically non-Gaussian, because, unlike the inflaton case, the self-interaction strength in such an extra scalar sector does not suffer from the usual slow-roll conditions.

between local and equilateral models, though in a non-trivial manner:

$$B_\Phi^{\mathrm{QSI}}(k_1, k_2, k_3) = \frac{\sqrt{972} A^2 f_{\mathrm{NL}}^{\mathrm{QSI}}}{(k_1 k_2 k_3)^{3/2}} \frac{N_\nu[8 k_1 k_2 k_3/(k_1 + k_2 + k_3)^3]}{N_\nu[8/27](k_1 + k_2 + k_3)^{3/2}} \quad (12)$$

where $\nu = (9/4 - m^2/H^2)^{1/2}$ and N_ν is the Neumann function of order ν. Quasi-single field models can also produce an essentially "constant" bispectrum defined by $B^{\mathrm{const}}(k_1, k_2, k_3) = 6 A^2 f_{\mathrm{NL}}^{\mathrm{const}}/(k_1 k_2 k_3)^2$. The constant model is the simplest possible non-zero primordial shape, with all its late-time CMB structure simply reflecting the behaviour of the transfer functions.

Alternatives to inflation: local NG can also be generated in some alternative scenarios to inflation, for instance in cyclic/ekpyrotic models (for a review, see Lehners 2010), due to the same basic curvaton mechanism described above. In this case, typical values of the nonlinearity parameter can easily reach $|f_{\mathrm{NL}}^{\mathrm{local}}| > 10$.

2.3. Non-standard models giving rise to alternative specific forms of NG

Non-Bunch-Davies vacuum and higher-derivative interactions: another interesting bispectrum shape is the folded one, which peaks in flattened configurations. To facilitate data analyses, the flat shape has been usually parametrized by the template (Meerburg et al. 2009)

$$B_\Phi^{\mathrm{flat}}(k_1, k_2, k_3) = 6 A^2 f_{\mathrm{NL}}^{\mathrm{flat}}$$
$$\times \left\{ \frac{1}{k_1^{4-n_s} k_2^{4-n_s}} + \frac{1}{k_2^{4-n_s} k_3^{4-n_s}} + \frac{1}{k_3^{4-n_s} k_1^{4-n_s}} + \frac{3}{(k_1 k_2 k_3)^{2(4-n_s)/3}} \right.$$
$$\left. - \left[\frac{1}{k_1^{(4-n_s)/3} k_2^{2(4-n_s)/3} k_3^{4-n_s}} + (5 \text{ perm.}) \right] \right\}. \quad (13)$$

The initial quantum state of the inflaton is usually specified by requiring that, at asymptotically early times and short distances, its fluctuations behave as in flat space. Deviations from this standard "Bunch-Davies" vacuum can result in interesting features in the bispectrum. Models with an initial non-Bunch-Davies vacuum state (Chen et al. 2007b; Holman & Tolley 2008; Meerburg et al. 2009; Ashoorioon & Shiu 2011) can generate sizeable NG similar to this type. NG highly correlated with such a template can be produced in single-field models of inflation from higher-derivative interactions (Bartolo et al. 2010a), and in models where a "Galilean" symmetry is imposed (Creminelli et al. 2011a). In both cases, cubic inflaton interactions with two derivatives of the inflaton field arise. Single-field inflation models with a small sound speed, studied in Senatore et al. (2010), can generate the flat shape, as a result of a linear combination of the orthogonal and equilateral shapes. In fact, from a simple parametrization point of view, the flat shape can be always written as $F_{\mathrm{flat}}(k_1, k_2, k_3) = [F_{\mathrm{equil}}(k_1, k_2, k_3) - F_{\mathrm{ortho}}(k_1, k_2, k_3)]/2$ (Senatore et al. 2010). Despite this, we provide constraints also on the amplitude of the flat bispectrum shape of Eq. (13).

For models with excited (i.e., non-Bunch-Davies) initial states, the resulting NG shapes are model-dependent, but they are usually characterized by the importance of flattened or collinear triangles, with $k_3 \approx k_1 + k_2$ along the edges of the tetrapyd. We will denote the original flattened bispectrum shape, given in Eqs. (6.2) and (6.3) of Chen et al. (2007b), by B_Φ^{NBD}; it is generically much more flattened than the "flat" model of Eq. (13). Although this shape was derived specifically for power-law k-inflation, it encapsulates several different shapes, with amplitudes which can vary between different phenomenological

models. These shapes are also typically oscillatory, being regularized by a cutoff scale k_c giving the oscillation period; this cutoff $k_c \approx (c_s \tau_c)^{-1}$ is determined by the (finite) time τ_c in the past when the non-Bunch-Davies component was initially excited. For excited canonical single-field inflation, the two leading order shapes can be described (Agullo & Parker 2011) by the ansatz

$$B_\Phi^{\mathrm{NBD}i} = \frac{2 A^2 f_{\mathrm{NL}}^{\mathrm{NBD}i}}{(k_1 k_2 k_3)^3} \left\{ f_i(k_1, k_2, k_3) \right. \quad (14)$$
$$\left. \times \frac{1 - \cos[(k_2 + k_3 - k_1)/k_c]}{k_2 + k_3 - k_1} + 2 \text{ perm.} \right\},$$

where $f_1(k_1, k_2, k_3) = k_1^2(k_2^2 + k_3^2)/2$ is dominated by squeezed configurations, $f_2(k_1, k_2, k_3) = k_2^2 k_3^2$ has a flattened shape, and $i = 1, 2$. Note that for all oscillatory shapes, the relevant bispectrum equation defines the normalisation of f_{NL}. The flattened signal is most easily enhanced in the limit of small sound speed c_s, for which a regularized ansatz is given by (Chen 2010b)

$$B_\Phi^{\mathrm{NBD}3} = \frac{2 A^2 f_{\mathrm{NL}}^{\mathrm{NBD}3}}{k_1 k_2 k_3} \left[\frac{k_1 + k_2 - k_3}{(k_c + k_1 + k_2 - k_3)^4} + 2 \text{ perm.} \right]. \quad (15)$$

Scale-dependent feature and resonant models: oscillating bispectra can be generated from violation of a smooth slow-roll evolution ("feature" or "resonant" NG). These models have the distinctive property of a strong running NG, which breaks approximate scale-invariance. A sharp feature in the inflaton potential forces the inflaton field away from the attractor solution, and causes oscillations as it relaxes back; these oscillations can appear in the bispectrum (Wang & Kamionkowski 2000; Chen et al. 2007a, 2008), as well as the power spectrum and other correlators. An analytic form for the oscillatory bispectrum for these feature models is (Chen et al. 2008)

$$B_\Phi^{\mathrm{feat}}(k_1, k_2, k_3) = \frac{6 A^2 f_{\mathrm{NL}}^{\mathrm{feat}}}{(k_1 k_2 k_3)^2} \sin \left[\frac{2\pi(k_1 + k_2 + k_3)}{3 k_c} + \phi \right], \quad (16)$$

where ϕ is a phase factor and k_c is a scale associated with the feature, which is linked in turn to an effective multipole periodicity ℓ_c of the CMB bispectrum. Typically, these oscillations will decay with an envelope of the form $\exp[-(k_1 + k_2 + k_3)/m k_c]$ for a model-dependent parameter m.

Closely related "resonant" bispectra can be created by periodic features superimposed on a smooth inflation potential (Chen et al. 2008; Flauger & Pajer 2011); these induce small periodic features in the background evolution, with which the quantum inflaton fluctuations can resonate while still inside the horizon. Resonant models are particularly relevant in the context of axion inflation models (e.g., Flauger et al. 2010; Flauger & Pajer 2011; Barnaby et al. 2012b). These mechanisms also create oscillatory behaviour in the bispectrum, but with a more constant amplitude and a wavelength that becomes logarithmically stretched. Here, the resonant oscillations for most models can be represented in the form

$$B_\Phi^{\mathrm{res}}(k_1, k_2, k_3) = \frac{6 A^2 f_{\mathrm{NL}}^{\mathrm{res}}}{(k_1 k_2 k_3)^2} \sin \left[C \ln(k_1 + k_2 + k_3) + \phi \right], \quad (17)$$

where the constant $C = 1/\ln(3 k_c)$ and ϕ is a phase.

Finally, we note that periodic features in the inflationary potential can excite the vacuum state, as well as perturbing

the background inflation trajectory (Chen 2010a). Such models offer the intriguing possibility of combining the flattened non-Bunch-Davies shape with periodic oscillations:

$$B_\Phi^{\text{resNBD}}(k_1, k_2, k_3) = \frac{2A^2 f_{\text{NL}}^{\text{resNBD}}}{(k_1 k_2 k_3)^2} \{\exp[-k_c^{3/5}(k_2 + k_3 - k_1)/2k_1]$$

$$\times \sin[k_c((k_2 + k_3 - k_1)/2k_1 + \ln k_1) + \phi] + 2 \text{ perm.}\}. \quad (18)$$

This ansatz represents the dominant folded resonant contribution in inflationary models with non-canonical kinetic terms, which competes with resonant (Eq. (17)) and equilateral (Eq. (3)) contributions; however, for slow-roll single-field inflation, there are additional terms.

Directional dependence motivated by gauge fields: additional variations of the bispectrum shape have been proposed for models with vector fields, which can have an additional directional dependence through the parameter $\mu_{12} = \hat{k}_1 \cdot \hat{k}_2$ where $\hat{k} = k/k$. For example, primordial magnetic fields sourcing curvature perturbations can cause a dependence on both μ and μ^2 (Shiraishi et al. 2012), and a coupling between the inflaton ϕ and the gauge field strength F^2 can yield a μ^2 dependence (Barnaby et al. 2012a; Bartolo et al. 2013). We can parameterize these shapes as variations on the local shape, following Shiraishi et al. (2013), as

$$B_\Phi(k_1, k_2, k_3) = \sum_L c_L[P_L(\mu_{12})P_\Phi(k_1)P_\Phi(k_2) + 2 \text{ perm}], \quad (19)$$

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

$$B_\Phi^{L=1}(k_1, k_2, k_3) = \frac{2A^2 f_{\text{NL}}^{L=1}}{(k_1 k_2 k_3)^2} \left[\frac{k_3^2}{k_1^2 k_2^2}(k_1^2 + k_2^2 - k_3^2) + 2 \text{ perm.}\right]. \quad (20)$$

Also the recently introduced "solid inflation" model (Endlich et al. 2013) generates bispectra similar to Eq. (19). 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 e.g., $k_1 + k_2 \approx k_3$, while in the squeezed limit, $L = 1$ is suppressed whereas $L = 2$ grows like the local bispectrum shape (i.e., the $L = 0$ case). Whether or not the underlying gauge field models prove robust, this directional dependence on the wave vectors is a generic feature which yields distinct bispectrum families, deserving closer study.

Warm inflation: in warm inflation (Berera 1995), where dissipative effects are important, a non-Gaussian signal can be generated (e.g., Moss & Xiong 2007) that peaks in the squeezed limit – but with a more complex shape than the local one – and exhibiting a low cross-correlation with the other shapes (see references in Liguori et al. 2010).

2.4. Higher-order non-Gaussianity: the trispectrum

The connected four-point functions of CMB anisotropies (or the harmonic counterpart, the so-called trispectrum) can also provide crucial information about the mechanism that gave rise to the primordial curvature perturbations (Okamoto & Hu 2002). The primordial trispectrum is usually characterised by two amplitudes τ_{NL} and g_{NL}: τ_{NL} is most often related to f_{NL}^2-type contributions, while g_{NL} is the amplitude of intrinsic cubic nonlinearities in the primordial gravitational potential (corresponding, in terms of field interactions, to a scalar-exchange and to a contact interaction term, respectively). They correspond to "soft"

limits of the full four-point function, with respectively the diagonal and one side of the general wavevector trapezoid being much smaller than the others. In the CMB maps they appear respectively approximately as a spatial variation in amplitude of the small-scale fluctuations, and a spatial variations in the value of f_{NL} correlated with the large-scale temperature. In addition to possible primordial signals that are the focus of this paper there is also expected to be a large lensing trispectrum (of very different shape), discussed in detail in Planck Collaboration XVII (2014).

The simplest local trispectrum is given by

$$\langle \Phi(k_1)\Phi(k_2)\Phi(k_3)\Phi(k_4) \rangle = (2\pi)^3 \delta^{(3)}(k_1 + k_2 + k_3 + k_4)$$

$$\times \left\{ \frac{25}{9} \tau_{\text{NL}} [P_\Phi(k_1)P_\Phi(k_2)P_\Phi(k_{13}) + (11 \text{ perm.})] \right.$$

$$\left. + 6g_{\text{NL}} [P_\Phi(k_1)P_\Phi(k_2)P_\Phi(k_3) + (3 \text{ perm.})] \right\}, \quad (21)$$

where $k_{ij} \equiv |k_i + k_j|$. Previous constraints on τ_{NL} and g_{NL} have been derived, e.g., by Smidt et al. (2010) who obtained $-7.4 \times 10^5 < g_{\text{NL}} < 8.2 \times 10^5$ and $-0.6 \times 10^4 < \tau_{\text{NL}} < 3.3 \times 10^4$ (at 95% CL) analysing WMAP-5 data; for the same datasets Fergusson et al. (2010b) obtained $-5.4 \times 10^5 < g_{\text{NL}} < 8.6 \times 10^5$ (68% CL). This kind of trispectrum typically arises in multi-field inflationary models where large NG arise from the conversion of isocurvature perturbations on superhorizon scales. If the curvature perturbation is the standard local form, in real space one has $\Phi(x) = \Phi_L(x) + f_{\text{NL}}^{\text{local}}(\Phi_L^2(x) - \langle\Phi_L^2\rangle) + g_{\text{NL}}\Phi_L^3(x)$. In this case, $\tau_{\text{NL}} = (6f_{\text{NL}}^{\text{local}}/5)^2$; however, in general the trispectrum amplitude can be larger.

The trispectrum is a complementary observable to the CMB bispectrum as it can further distinguish different inflationary scenarios. This is because the same interactions that lead to the bispectrum might be responsible also for a large trispectrum, so that the different NG parameters can be related to each other in a well-defined way within specific models. If there is a non-zero squeezed-shape bispectrum there must necessarily be a trispectrum, with $\tau_{\text{NL}} \geq (6f_{\text{NL}}^{\text{local}}/5)^2$ (Suyama & Yamaguchi 2008; Suyama et al. 2010; Sugiyama et al. 2011; Sugiyama 2012; Lewis 2011; Smith et al. 2011; Assassi et al. 2012; Kehagias & Riotto 2012). In the simplest inflationary scenarios the prediction would be $\tau_{\text{NL}} = (6f_{\text{NL}}^{\text{local}}/5)^2$, but larger values would indicate more complicated dynamics. Several inflationary scenarios have been found in which the bispectrum is suppressed, thus leaving the trispectrum as the largest higher-order correlator in the data. A detection of a large trispectrum and a negligible bispectrum would be a smoking gun for these models. This is the case, for example, of certain curvaton and multi-field field models of inflation (Byrnes et al. 2006; Sasaki et al. 2006; Byrnes & Choi 2010), which for particular parameter choice can produce a significant τ_{NL} and g_{NL} and small f_{NL}. Large trispectra are also possible in single-field models of inflation with higher-derivative interactions (see, e.g., Chen et al. 2009; Arroja et al. 2009; Senatore & Zaldarriaga 2011; Bartolo et al. 2010b), but these would be suppressed in the squeezed limit since they are generated by derivative interactions at horizon-crossing, and hence only project weakly onto the local shapes. These equilateral trispectra arise can be well-described by some template forms (Fergusson et al. 2010b). Naturally, higher-order correlations could also be considered, but are not directly studied in this paper.

3. Statistical estimation of the CMB bispectrum

In this section, we review the statistical techniques that we use to estimate the nonlinearity parameter $f_{\rm NL}$. We begin by fixing some notation and describing the CMB angular bispectrum in Sect. 3.1. We then introduce in Sect. 3.2 the optimal $f_{\rm NL}$ bispectrum estimator. From Sect. 3.2.1 onwards we describe in detail the different implementations of the optimal estimator that were developed and applied to *Planck* data.

3.1. The CMB angular bispectrum

Temperature anisotropies are represented using the $a_{\ell m}$ coefficients of a spherical harmonic decomposition of the CMB map,

$$\frac{\Delta T}{T}(\hat{\boldsymbol{n}}) = \sum_{\ell m} a_{\ell m} Y_{\ell m}(\hat{\boldsymbol{n}}); \tag{22}$$

we write $C_\ell = \langle |a_{\ell m}|^2 \rangle$ for the angular power spectrum and $\hat{C}_\ell = (2\ell+1)^{-1} \sum_m |a_{\ell m}|^2$ for the corresponding (ideal) estimator; hats "^" denote estimated quantities. The CMB angular bispectrum is the three-point correlator of the $a_{\ell m}$:

$$B_{\ell_1 \ell_2 \ell_3}^{m_1 m_2 m_3} \equiv \langle a_{\ell_1 m_1} a_{\ell_2 m_2} a_{\ell_3 m_3} \rangle. \tag{23}$$

If the CMB sky is rotationally invariant, the angular bispectrum can be factorized as follows:

$$\langle a_{\ell_1 m_1} a_{\ell_2 m_2} a_{\ell_3 m_3} \rangle = \mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} b_{\ell_1 \ell_2 \ell_3}, \tag{24}$$

where $b_{\ell_1 \ell_2 \ell_3}$ 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:

$$\begin{aligned} \mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} &\equiv \int Y_{\ell_1 m_1}(\hat{\boldsymbol{n}}) Y_{\ell_2 m_2}(\hat{\boldsymbol{n}}) Y_{\ell_3 m_3}(\hat{\boldsymbol{n}}) \, {\rm d}^2 \hat{\boldsymbol{n}} \\ &= h_{\ell_1 \ell_2 \ell_3} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix}, \end{aligned} \tag{25}$$

where $h_{\ell_1 \ell_2 \ell_3}$ is a geometrical factor,

$$h_{\ell_1 \ell_2 \ell_3} = \sqrt{\frac{(2\ell_1+1)(2\ell_2+1)(2\ell_3+1)}{4\pi}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix}. \tag{26}$$

The Wigner-3j symbol in parentheses enforces rotational symmetry, and allows us to restrict 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 $\ell_{\rm max}$ (the latter being defined by the finite angular resolution of the experiment under study). This three-dimensional domain $\mathcal{V}_\mathcal{T}$ of allowed multipoles, sometimes referred to in the following as a "tetrapyd", is illustrated in Fig. 1 and it is explicitly defined by

Triangle condition: $\ell_1 \le \ell_2 + \ell_3$ for $\ell_1 \ge \ell_2, \ell_3$, +perms.,
Parity condition: $\ell_1 + \ell_2 + \ell_3 = 2n, \quad n \in \mathbb{N},$ (27)
Resolution: $\ell_1, \ell_2, \ell_3 \le \ell_{\rm max}, \quad \ell_1, \ell_2, \ell_3 \in \mathbb{N}.$

Here, $\mathcal{V}_\mathcal{T}$ is the isotropic subset of the full space of bispectra, denoted by $\mathcal{V}$.

One can also define an alternative rotationally-invariant reduced bispectrum $B_{\ell_1 \ell_2 \ell_3}$ in the following way:

$$B_{\ell_1 \ell_2 \ell_3} \equiv h_{\ell_1 \ell_2 \ell_3} \sum_{m_1 m_2 m_3} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix} B_{\ell_1 \ell_2 \ell_3}^{m_1 m_2 m_3}. \tag{28}$$
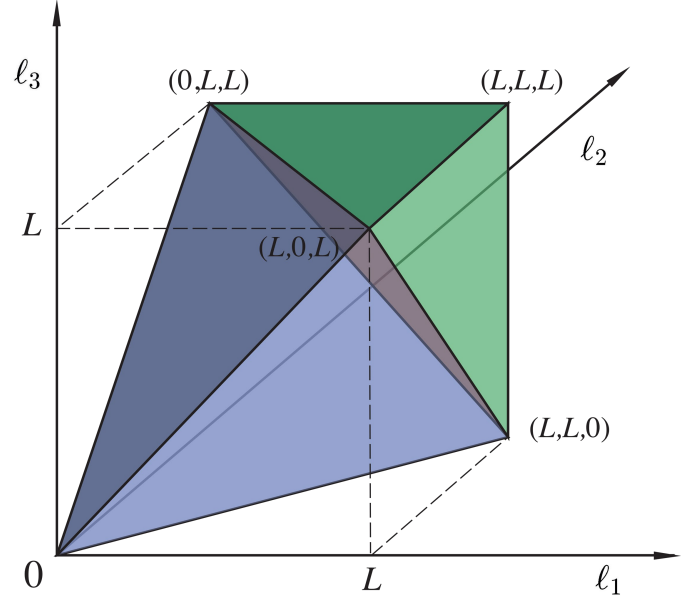

Fig. 1. Permitted observational domain of Eq. (27) for the CMB bispectrum $b_{\ell_1 \ell_2 \ell_3}$. Allowed multipole values (ℓ_1, ℓ_2, ℓ_3) lie inside the shaded "tetrapyd" region (tetrahedron+pyramid), satisfying both the triangle condition and the experimental resolution $\ell < L \equiv \ell_{\rm max}$.

Note that this $B_{\ell_1 \ell_2 \ell_3}$ is equal to $h_{\ell_1 \ell_2 \ell_3}$ times the angle-averaged bispectrum as defined in the literature. From Eqs. (24) and (25), and the fact that the sum over all m_i of the Wigner-3j symbol squared is equal to 1, it is easy to see that $B_{\ell_1 \ell_2 \ell_3}$ is related to the reduced bispectrum by

$$B_{\ell_1 \ell_2 \ell_3} = h_{\ell_1 \ell_2 \ell_3}^2 b_{\ell_1 \ell_2 \ell_3}. \tag{29}$$

The interest in this bispectrum $B_{\ell_1 \ell_2 \ell_3}$ is that it can be estimated directly from maximally-filtered maps of the data:

$$\hat{B}_{\ell_1 \ell_2 \ell_3} = \int {\rm d}^2 \hat{\boldsymbol{n}} T_{\ell_1}(\hat{\boldsymbol{n}}) T_{\ell_2}(\hat{\boldsymbol{n}}) T_{\ell_3}(\hat{\boldsymbol{n}}), \tag{30}$$

where the filtered maps $T_\ell(\hat{\boldsymbol{n}})$ are defined as:

$$T_\ell(\hat{\boldsymbol{n}}) \equiv \sum_m a_{\ell m} Y_{\ell m}(\hat{\boldsymbol{n}}). \tag{31}$$

This can be seen by replacing the $B_{\ell_1 \ell_2 \ell_3}^{m_1 m_2 m_3}$ in Eq. (28) by its estimate $a_{\ell_1 m_1} a_{\ell_2 m_2} a_{\ell_3 m_3}$ and then using Eq. (25) to rewrite the Wigner symbol in terms of a Gaunt integral, which in its turn is expressed as an integral over the product of three spherical harmonics.

3.2. CMB bispectrum estimators

The full bispectrum for a high-resolution map cannot be evaluated explicitly because of the sheer number of operations involved, $O(\ell_{\rm max}^5)$, as well as the fact that the signal will be too weak to measure in individual multipoles with any significance. Instead, we essentially use a least-squares fit to compare the bispectrum of the observed CMB multipoles with a particular theoretical bispectrum $b_{\ell_1 \ell_2 \ell_3}$. We then extract an overall "amplitude parameter" $f_{\rm NL}$ for that specific template, after defining a suitable normalization convention so that we can write $b_{\ell_1 \ell_2 \ell_3} = f_{\rm NL} b_{\ell_1 \ell_2 \ell_3}^{\rm th}$, where $b_{\ell_1 \ell_2 \ell_3}^{\rm th}$ is defined as the value of the theoretical bispectrum ansatz for $f_{\rm NL} = 1$.

Optimal 3-point estimators (introduced by Heavens 1998; see also Gangui & Martin 2000a), are those which saturate the Cramér-Rao bound. Taking into account the fact that instrument noise and masking can break rotational invariance, it has been shown that the general optimal f_{NL} estimator can be written as (Babich 2005; Creminelli et al. 2006; Senatore et al. 2010; Verde et al. 2013):

$$
\begin{aligned}
\hat{f}_{\mathrm{NL}} = \frac{1}{N} \sum_{\ell_i, m_i} & \mathcal{G}^{\ell_1 \ell_2 \ell_3}_{m_1 m_2 m_3} b^{\mathrm{th}}_{\ell_1 \ell_2 \ell_3} \\
& \times [C^{-1}_{\ell_1 m_1, \ell'_1 m'_1} a_{\ell'_1 m'_1} C^{-1}_{\ell_2 m_2, \ell'_2 m'_2} a_{\ell'_2 m'_2} C^{-1}_{\ell_3 m_3, \ell'_3 m'_3} a_{\ell'_3 m'_3} \\
& - 3 C^{-1}_{\ell_1 m_1, \ell_2 m_2} C^{-1}_{\ell_3 m_3, \ell'_3 m'_3} a_{\ell'_3 m'_3}],
\end{aligned}
\tag{32}
$$

where $\mathbf{C}^{-1}$ is the inverse of the covariance matrix $C_{\ell_1 m_1, \ell_2 m_2} \equiv \langle a_{\ell_1 m_1} a_{\ell_2 m_2} \rangle$ and N is a suitable normalization chosen to produce unit response to $b^{\mathrm{th}}_{\ell_1 \ell_2 \ell_3}$.

In the expression of the optimal estimator above we note the presence of two contributions, one (hereafter defined the "cubic term" of the estimator) is cubic in the observed $a_{\ell m}$ and correlates the bispectrum of the data to the theoretical fitting template $b^{\mathrm{th}}_{\ell_1 \ell_2 \ell_3}$, while the other is linear in the observed $a_{\ell m}$ (hereafter, the "linear term"), which is zero on average. In the rotationally-invariant case the linear term is proportional to the monopole in the map, which has been set to zero, so in this case the estimator simply reduces to the cubic term. However, when rotational invariance is broken by realistic experimental features such as a Galactic mask or an anisotropic noise distribution, the linear term has an important effect on the estimator variance. In this case, the coupling between different ℓ would in fact produce a spurious increase in the error bars (coupling of Fourier modes due to statistical anisotropy can be "misinterpreted" by the estimator as NG). The linear term correlates the observed $a_{\ell m}$ to the power spectrum anisotropies and removes this effect, thus restoring optimality (Creminelli et al. 2006; Yadav et al. 2008, 2007).

The actual problem with Eq. (32) is that its direct implementation to get an optimal f_{NL} estimator would require measurement of all the bispectrum configurations from the data. As already mentioned at the beginning of this section, the computational cost of this would scale like $\ell_{\max}^5$ and be totally prohibitive for high-resolution CMB experiments. Even taking into account the constraints imposed by isotropy, the number of multipole triples $\{\ell_1, \ell_2, \ell_3\}$ is of the order of 10^9 at *Planck* resolution, and the number of different observed bispectrum configurations $\hat{b}^{\ell_1 \ell_2 \ell_3}_{m_1 m_2 m_3}$ is of the order of 10^{15}. For each of them, costly numerical evaluation of the Wigner symbol is also required. This is completely out of reach of existing supercomputers. It is then necessary to find numerical solutions that circumvent this problem and in the following subsections we will show how the different estimators used for the f_{NL} *Planck* data analysis address this challenge. Before entering into a more accurate description of these different methods, we would like however to stress again that they are all going to be different implementations of the optimal f_{NL} estimator defined by Eq. (32); therefore they are conceptually equivalent and expected to produce f_{NL} results that are in very tight agreement. This will later on allow for stringent validation tests based on comparing different pipelines. On the other hand, it will soon become clear that the different approaches that we are going to discuss also open up a range of additional applications beyond simple f_{NL} estimation for standard bispectra. Such applications include, for example, full bispectrum reconstruction (in a suitably smoothed domain), tests of directional

dependence of f_{NL}, and other ways to reduce the amount of data, going beyond simple single-number f_{NL} estimation. So different methods will also provide a vast range of complementary information.

Another important preliminary point, to notice before discussing different techniques, is that none of the estimators in the following sections implement exactly Eq. (32), but a slightly modified version of it. In Eq. (32) the CMB multipoles always appear weighted by the inverse of the full covariance matrix. Inverse covariance filtering of CMB data at the high angular resolutions achieved by experiments like WMAP and *Planck* is another very challenging numerical issue, which was fully addressed only recently (Smith et al. 2009; Komatsu et al. 2011; Elsner & Wandelt 2013). For our analyses we developed two independent inverse-covariance filtering pipelines. The former is based on an extension to *Planck* resolution of the algorithm used for WMAP analysis (Smith et al. 2009; Komatsu et al. 2011); the latter is based on the algorithm described in Elsner & Wandelt (2013). However, detailed comparisons interestingly showed that our estimators perform equally well (i.e., they saturate the Cramér-Rao bound) if we approximate the covariance matrix as diagonal in the filtering procedure and we apply a simple diffusive inpainting procedure to the masked areas of the input CMB maps. A more detailed description of our inpainting and Wiener filtering algorithms can be found in Sect. 3.3.

In the diagonal covariance approximation, the minimum variance estimator is obtained by making the replacement $(C^{-1}a)_{\ell m} \rightarrow a_{\ell m}/C_\ell$ in the cubic term and then including the linear term that minimizes the variance for this class of cubic estimator (Creminelli et al. 2006). This procedure leads to the following expression:

$$
\begin{aligned}
\hat{f}_{\mathrm{NL}} = \frac{1}{N} \sum_{\ell_i, m_i} & \mathcal{G}^{\ell_1 \ell_2 \ell_3}_{m_1 m_2 m_3} \tilde{b}^{\mathrm{th}}_{\ell_1 \ell_2 \ell_3} \\
& \times \left[\frac{\tilde{a}_{\ell_1 m_1}}{\tilde{C}_{\ell_1}} \frac{\tilde{a}_{\ell_2 m_2}}{\tilde{C}_{\ell_2}} \frac{\tilde{a}_{\ell_3 m_3}}{\tilde{C}_{\ell_3}} - 6 \frac{\tilde{C}_{\ell_1 m_1, \ell_2 m_2}}{\tilde{C}_{\ell_1} \tilde{C}_{\ell_2}} \frac{\tilde{a}_{\ell_3 m_3}}{\tilde{C}_{\ell_3 m_3}} \right],
\end{aligned}
\tag{33}
$$

where the tilde denotes the modification of C_ℓ and $b_{\ell_1 \ell_2 \ell_3}$ to incorporate instrument beam and noise effects, and indicates that the multipoles are obtained from a map that was masked and preprocessed through the inpainting procedure detailed in Sect. 3.3. This means that

$$
\tilde{b}_{\ell_1 \ell_2 \ell_3} \equiv b_{\ell_1} b_{\ell_2} b_{\ell_3} b_{\ell_1 \ell_2 \ell_3}, \quad \tilde{C}_\ell \equiv b_\ell^2 C_\ell + N_\ell,
\tag{34}
$$

where b_ℓ denotes the experimental beam, and N_ℓ is the noise power spectrum. For simplicity of notation, in the following we will drop the tilde and always assume that beam, noise and inpainting effects are properly included.

Using Eqs. (28) and (29) we can rewrite Eq. (33) in terms of the bispectrum $B_{\ell_1 \ell_2 \ell_3}$:

$$
\hat{f}_{\mathrm{NL}} = \frac{6}{N} \sum_{\ell_1 \le \ell_2 \le \ell_3} \frac{B^{\mathrm{th}}_{\ell_1 \ell_2 \ell_3} \left(B^{\mathrm{obs}}_{\ell_1 \ell_2 \ell_3} - B^{\mathrm{lin}}_{\ell_1 \ell_2 \ell_3} \right)}{V_{\ell_1 \ell_2 \ell_3}}.
\tag{35}
$$

In the above expression, B^{th} is the theoretical template for B (with $f_{\mathrm{NL}} = 1$) and B^{obs} denotes the observed bispectrum (the cubic term), extracted from the (inpainted) data using Eq. (30). B^{lin} is the linear correction, also computed using Eq. (30) by replacing two of the filtered temperature maps by simulated Gaussian ones and averaging over a large number of them (three permutations). The variance V in the inverse-variance weights is given by $V_{\ell_1 \ell_2 \ell_3} = g_{\ell_1 \ell_2 \ell_3} h^2_{\ell_1 \ell_2 \ell_3} C_{\ell_1} C_{\ell_2} C_{\ell_3}$ (remember that these should be viewed as being the quantities with tildes, having

beam and noise effects included) with $g_{\ell_1\ell_2\ell_3}$ a permutation factor ($g_{\ell_1\ell_2\ell_3} = 6$ when all ℓ are equal, $g_{\ell_1\ell_2\ell_3} = 2$ when two ℓ are equal, and $g_{\ell_1\ell_2\ell_3} = 1$ otherwise). Both Eqs. (33) and (35) will be used in the following. Equation (33) will provide the starting point for the KSW, skew-C_ℓ and modal estimators, while Eq. (35) will be the basis for the binned and wavelets estimators.

Next, we will describe in detail the different methods, and show how they address the numerical challenge posed by the necessity to evaluate a huge number of bispectrum configurations. To summarize loosely: the KSW estimator, the skew-C_ℓ approach and the separable modal methodology achieve massive reductions in computational costs by exploiting structural properties of b^{th}, e.g., separability. On the other hand, the binned bispectrum and the wavelet approaches achieve computational gains by data compression of B^{obs}.

3.2.1. The KSW estimator

To understand the rationale behind the KSW estimator (Komatsu et al. 2005, 2003; Senatore et al. 2010; Creminelli et al. 2006; Yadav et al. 2008, 2007; Yadav & Wandelt 2008; Smith & Zaldarriaga 2011), assume that the theoretical reduced bispectrum $b^{\mathrm{th}}_{\ell_1\ell_2\ell_3}$ can be exactly decomposed into a separable structure, e.g., there exist some sequences of functions $\alpha(\ell, r)$, $\beta(\ell, r)$ such that we can approximate $b_{\ell_1\ell_2\ell_3}$ as

$$b_{\ell_1\ell_2\ell_3} \simeq \int \left[\beta(\ell_1, r)\beta(\ell_2, r)\alpha(\ell_3, r) + \beta(\ell_1, r)\beta(\ell_3, r)\alpha(\ell_2, r) \right.$$
$$\left. + \beta(\ell_2, r)\beta(\ell_3, r)\alpha(\ell_1, r)\right] \, r^2 \, \mathrm{d}r, \quad (36)$$

where r is a radial coordinate. This assumption is fulfilled in particular by the local shape (Komatsu & Spergel 2001; Babich et al. 2004), with $\alpha(\ell, r)$ and $\beta(\ell, r)$ involving integrals of products of spherical Bessel functions and CMB radiation transfer functions. Let us consider the optimal estimator of Eq. (33) and neglect for the moment the linear part. Exploiting Eq. (36) and the factorizability property of the Gaunt integral (Eq. (25)), the cubic term of the estimator can be written as:

$$S_{\mathrm{cub}} = \int \mathrm{d}r\, r^2 \int \mathrm{d}^2\hat{\boldsymbol{n}}\, A(\hat{\boldsymbol{n}}, r)B^2(\hat{\boldsymbol{n}}, r) \quad (37)$$

where

$$A(\hat{\boldsymbol{n}}, r) = \sum_{\ell m} \frac{\alpha(\ell, r)\, a_{\ell m} Y_{\ell m}(\hat{\boldsymbol{n}})}{C_\ell}, \quad (38)$$

and

$$B(\hat{\boldsymbol{n}}, r) = \sum_{\ell m} \frac{\beta(\ell, r)\, a_{\ell m} Y_{\ell m}(\hat{\boldsymbol{n}})}{C_\ell}. \quad (39)$$

From the formulae above we see that the overall triple integral over all the configurations ℓ_1, ℓ_2, ℓ_3 has been factorized into a product of three separate sums over different ℓ. This produces a massive reduction in computational time, as the problem now scales like $\ell^3_{\max}$ instead of the original $\ell^5_{\max}$. Moreover, the bispectrum can be evaluated in terms of a cubic statistic in pixel space from Eq. (37), and the functions $A(\hat{\boldsymbol{n}}, r)$, $B(\hat{\boldsymbol{n}}, r)$ are obtained from the observed $a_{\ell m}$ by means of Fast Harmonic Transforms.

It is easy to see that the linear term can be factorized in analogous fashion. Again considering the local shape type of decomposition of Eq. (36), it is possible to find:

$$S_{\mathrm{lin}} = \frac{-6}{N} \int \mathrm{d}r\, r^2 \int \mathrm{d}^2\hat{\boldsymbol{n}} \left[2 \langle A(r, \hat{\boldsymbol{n}})B(r, \hat{\boldsymbol{n}})\rangle_{\mathrm{MC}} \right.$$
$$\left. \times B(r, \hat{\boldsymbol{n}}) + \langle B(r, \hat{\boldsymbol{n}})B(r, \hat{\boldsymbol{n}})\rangle_{\mathrm{MC}}\, A(r, \hat{\boldsymbol{n}})\right], \quad (40)$$

where $\langle\cdot\rangle_{\mathrm{MC}}$ denotes a Monte Carlo (MC) average over simulations accurately reproducing the properties of the actual data set (basically we are taking a MC approach to estimate the product between the theoretical bispectrum and the $a_{\ell m}$ covariance matrix appearing in the linear term expression).

The estimator can be finally expressed as a function of S_{cub} and S_{lin}:

$$\hat{f}_{\mathrm{NL}} = \frac{S_{\mathrm{cub}} + S_{\mathrm{lin}}}{N}. \quad (41)$$

Whenever it can be applied, the KSW approach makes the problem of f_{NL} estimation computationally feasible, even at the high angular resolution of the *Planck* satellite. One important caveat is that factorizability of the shape, which is the starting point of the method, is not a general property of theoretical bispectrum templates. Strictly speaking, only the local shape is manifestly separable. However, a large class of inflationary models can be extremely well approximated by separable equilateral and orthogonal templates (Babich et al. 2004; Creminelli et al. 2006; Senatore et al. 2010). The specific expressions of cubic and linear terms are of course template-dependent, but as long as the template itself is separable their structure is analogous to the example shown in this section, i.e., they can be written as pixel space integrals of cubic products of suitably-filtered CMB maps (involving MC approximations of the $a_{\ell m}$ covariance for the linear term). For a complete and compact summary of KSW implementations for local, equilateral and orthogonal bispectra see Komatsu et al. (2009, Appendix).

3.2.2. The skew-C_ℓ extension

The skew-C_ℓ statistics were introduced by Munshi & Heavens (2010) to address an issue with estimators such as KSW which reduce the map to an estimator of f_{NL} for a given type of NG. This level of data compression, to a single number, has the disadvantage that it does not allow verification that a NG signal is of the type which has been estimated. KSW on its own cannot tell if a measurement of f_{NL} of given type is actually caused by NG of that type, or by contamination from some other source or sources. The skew-C_ℓ statistics perform a less radical data compression than KSW (to a function of ℓ), and thus retain enough information to distinguish different NG signals. The desire to find a statistic which is able to fulfil this rôle, but which is still optimal, drives one to a statistic which is closely related to KSW, and indeed reduces to it when the scale-dependent information is not used. A further advantage of the skew-C_ℓ is that it allows joint estimation of the level of many types of NG simultaneously. This requires a large number of simulations for accurate estimation of their covariance matrix, and they are not used in this role in this paper. However, they do play an important part in identifying which sources of NG are clearly detected in the data, and which are not.

We define the skew-C_ℓ statistics by extending from KSW, as follows: from Eq. (37), the numerator $\mathcal{E}$ can be rewritten as

$$\mathcal{E} = \sum_\ell \left[C^{A,B^2}_\ell + 2C^{AB,B}_\ell\right] \quad (42)$$

where

$$C^{A,B^2}_\ell = \int \int_{S^2} \sum_{\ell_1,\ell_2} \sum_{m_1,m_2,m} \left[\frac{\beta(\ell_1; r)a_{\ell_1 m_1} Y_{\ell_1 m_1}(\hat{\boldsymbol{n}})}{C_{\ell_1}}\right.$$
$$\left. \times \frac{\beta(\ell_2; r)a_{\ell_2 m_2} Y_{\ell_2 m_2}(\hat{\boldsymbol{n}})}{C_{\ell_2}} \frac{\alpha(\ell; r)a_{\ell m} Y_{\ell m}(\hat{\boldsymbol{n}})}{C_\ell}\right] r^2 \mathrm{d}^2\hat{\boldsymbol{n}}\mathrm{d}r \quad (43)$$

and

$$
\begin{aligned}
C_\ell^{AB,B} = \int \int_{S^2} \sum_{\ell_1,\ell_2} \sum_{m_1,m_2,m} & \left[\frac{\beta(\ell_1;r)a_{\ell_1 m_1}Y_{\ell_1 m_1}(\hat{\boldsymbol{n}})}{C_{\ell_1}} \right. \\
& \times \left. \frac{\alpha(\ell_2;r)a_{\ell_2 m_2}Y_{\ell_2,m_2}(\hat{\boldsymbol{n}})}{C_{\ell_2}} \frac{\beta(\ell;r)a_{\ell m}Y_{\ell m}(\hat{\boldsymbol{n}})}{C_\ell} \right] r^2 \mathrm{d}^2\hat{\boldsymbol{n}}\mathrm{d}r. \quad (44)
\end{aligned}
$$

The skew-C_ℓ approach allows for the full implementation of the KSW procedure, when the sum in Eq. (42) is fully evaluated; furthermore, it allows for extra degrees of flexibility, e.g., by restricting the sum to subsets of the multipole space, which may highlight specific features of the NG signal. Each form of NG considered has its own α, β, hence its own set of skew-C_ℓ, denoted $S_\ell \equiv C_\ell^{A,B^2} + 2C_\ell^{AB,B}$, and we have chosen to illustrate here with the local form, but as with KSW the method can be extended to other separable shapes, and some skew-C_ℓ do not involve integrals, such as the ISW-lensing skew statistic. Note that in this paper we do not fit the S_ℓ directly, but instead we estimate the NG using KSW, and then verify (or not) the nature of the NG by comparing the skew-C_ℓ with the theoretical expectation. No further free parameters are introduced at this stage. This procedure allows investigation of KSW detections of NG of a given type, assessing whether or not they are actually due to NG of that type.

3.2.3. Separable modal methodology

Primordial bispectra need not be manifestly separable (like the local bispectrum), or be easily approximated by separable ad hoc templates (equilateral and orthogonal), so the direct KSW approach above cannot be applied generically (nor to late-time bispectra). However, we can employ a highly-efficient generalization by considering a complete basis of separable modes describing any late-time bispectrum (see Fergusson & Shellard 2007; Fergusson et al. 2010a), as applied to WMAP-7 data for a wide variety of separable and non-separable bispectrum models (Fergusson et al. 2012). See also Planck Collaboration XXIII (2014) and Planck Collaboration XXV (2014). We can achieve this by expanding the signal-to-noise-weighted bispectrum as

$$
\frac{b_{\ell_1\ell_2\ell_3}}{\sqrt{C_{\ell_1}C_{\ell_2}C_{\ell_3}}} = \sum_{i,j,k} \alpha_{ijk}Q_{ijk}(\ell_1,\ell_2,\ell_3), \quad (45)
$$

where the (non-orthogonal) separable modes Q_n are defined by

$$
\begin{aligned}
Q_{ijk}(\ell_1,\ell_2,\ell_3) = \frac{1}{6} & \left[q_i(\ell_1)q_j(\ell_2)q_k(\ell_3) + q_j(\ell_1)q_i(\ell_2)q_k(\ell_3) \right. \\
& \left. + \text{cyclic perms in } i,j,k \right]. \quad (46)
\end{aligned}
$$

It is more efficient to label the basis as Q_n, with the subscript n representing an ordering of the $\{i,j,k\}$ products (e.g., by distance $i^2 + j^2 + k^2$). The $\bar{q}_i(\ell)$ are any complete basis functions up to a given resolution of interest and they can be augmented with other special functions adapted to target particular bispectra. The modal coefficients in the bispectrum of Eq. (45) are given by the inner product of the weighted bispectrum with each mode

$$
\alpha_n = \sum_p \gamma_{np}^{-1} \left\langle \frac{b_{\ell_1\ell_2\ell_3}}{\sqrt{C_{\ell_1}C_{\ell_2}C_{\ell_3}}}, Q_p(\ell_1,\ell_2,\ell_3) \right\rangle \quad (47)
$$

where the modal transformation matrix is

$$
\begin{aligned}
\gamma_{np} &= \langle Q_n, Q_p \rangle \\
&\equiv \sum_{\ell_1,\ell_2,\ell_3} h_{\ell_1\ell_2\ell_3}^2 Q_{ijk}(\ell_1,\ell_2,\ell_3) Q_{i'j'k'}(\ell_1,\ell_2,\ell_3). \quad (48)
\end{aligned}
$$

In the following, the specific basis functions $\bar{q}_i(\ell)$ we employ include either weighted Legendre-like polynomials or trigonometric functions. These are combined with the Sachs-Wolfe local shape and the separable ISW shape in order to obtain high correlations to all known bispectrum shapes (usually in excess of 99%).

Substituting the separable mode expansion of Eq. (45) into the estimator and exploiting the separability of the Gaunt integral (Eq. (25)), yields

$$
\begin{aligned}
\mathcal{E} = \frac{1}{N^2} \sum_{n\leftrightarrow prs} \alpha_n \int \mathrm{d}^2\hat{\boldsymbol{n}} & \left[\bar{M}_{\{p}(\hat{\boldsymbol{n}})\bar{M}_r(\hat{\boldsymbol{n}})\bar{M}_{s\}}(\hat{\boldsymbol{n}}) \right. \\
& \left. - 6 \left\langle \bar{M}_{\{p}^G(\hat{\boldsymbol{n}})\bar{M}_r^G(\hat{\boldsymbol{n}}) \right\rangle \bar{M}_{s\}}(\hat{\boldsymbol{n}}) \right], \quad (49)
\end{aligned}
$$

where the $\bar{M}_p(\hat{\boldsymbol{n}})$ represent versions of the original CMB map filtered with the basis function q_p (and the weights $(\sqrt{C_\ell})^{-1}$), that is,

$$
\bar{M}_p(\hat{\boldsymbol{n}}) = \sum_{\ell m} q_p(\ell) \frac{a_{\ell m}}{\sqrt{C_\ell}} Y_{\ell m}(\hat{\boldsymbol{n}}). \quad (50)
$$

The maps $\bar{M}_p^G(\hat{\boldsymbol{n}})$ incorporate the same mask and a realistic model of the inhomogeneous instrument noise; a large ensemble of these maps, calculated from Gaussian simulations, is used in the averaged linear term in the estimator of Eq. (49), allowing for the subtraction of these important effects. Defining the integral over these convolved product maps as cubic and linear terms respectively,

$$
\beta_n^{\text{cub}} = \int \mathrm{d}^2\hat{\boldsymbol{n}} \, \bar{M}_{\{p}(\hat{\boldsymbol{n}})\bar{M}_r(\hat{\boldsymbol{n}})\bar{M}_{s\}}(\hat{\boldsymbol{n}}), \quad (51)
$$

$$
\beta_n^{\text{lin}} = \int \mathrm{d}^2\hat{\boldsymbol{n}} \, \left\langle \bar{M}_{\{p}^G(\hat{\boldsymbol{n}})\bar{M}_r^G(\hat{\boldsymbol{n}}) \right\rangle \bar{M}_{s\}}(\hat{\boldsymbol{n}}), \quad (52)
$$

the estimator reduces to a simple sum over the mode coefficients

$$
\mathcal{E} = \frac{1}{N^2} \sum_n \alpha_n \bar{\beta}_n, \quad (53)
$$

where $\bar{\beta}_n^Q \equiv \bar{\beta}_n^{Q\text{cub}} - \bar{\beta}_n^{Q\text{lin}}$. The estimator sum in Eq. (53) is now straightforward to evaluate because of separability, since it has been reduced to a product of three sums over the observational maps (Eq. (49)), followed by a single 2D integral over all directions (Eq. (51)). The number of operations in evaluating the estimator sum is only $O(\ell^2)$.

For the purposes of testing a wide range inflationary models, we can also define a set of primordial basis functions $\overline{Q}_{ink}(k_1,k_2,k_3) = \bar{q}_i(k_1)\bar{q}_j(k_2)\bar{q}_k(k_3) + \text{perms.}$ for wavenumbers satisfying the triangle condition (again we will order the $\{i,j,k\}$ with n). This provides a separable expansion for an arbitrary primordial shape function $S(k_1,k_2,k_3) = B(k_1,k_2,k_3)/(k_1k_2k_3)^2$, that is,

$$
S(k_1,k_2,k_3) = \sum_n \bar{\alpha}_n \overline{Q}_n(k_1,k_2,k_3). \quad (54)
$$

Using the same transfer functions as in the KSW integral (37), we can efficiently project forward each separable primordial mode $\overline{Q}_n(k_1,k_2,k_3)$ to a corresponding late-time solution $\widetilde{Q}_n(l_1 l_2 l_3)$ (essentially the projected CMB bispectrum for $\overline{Q}_n(k_1,k_2,k_3)$). By finding the inner product between these projected modes $\widetilde{Q}_n(\ell_1,\ell_2,\ell_3)$ and the CMB basis functions $Q_n(\ell_1,\ell_2,\ell_3)$, we can obtain the transformation matrix (Fergusson et al. 2010a,b)

$$
\Gamma_{np} = \sum_r \bar{\gamma}_{np}^{-1} \langle \widetilde{Q}_r(\ell_1,\ell_2,\ell_3), Q(\ell_1,\ell_2,\ell_3) \rangle, \quad (55)
$$

which projects the primordial expansion coefficients α_n^Q to late-time:

$$\alpha_n = \sum_p \Gamma_{np} \bar{\alpha}_p. \tag{56}$$

When Γ_{np} is calculated once we can efficiently convert any given primordial bispectrum $B(k_1, k_2, k_3)$ into its late-time CMB bispectrum counterpart using Eq. (45). We can use this to extend the KSW methodology and to search for the much broader range of primordial models beyond local, equilateral and orthogonal, having validated on these standard shapes.

3.2.4. Binned bispectrum

In the binned bispectrum approach (Bucher et al. 2010), the computational gains are achieved by data compression of the observed $\hat{B}$. This is quite feasible, because like the power spectrum the bispectrum is a rather smooth function, with features on the scale of the acoustic peaks. In this way one obtains an enormous reduction of the computational resources needed at the cost of only a tiny increase in variance (typically 1%).

More precisely, the following statistic is considered,

$$T_i(\hat{\boldsymbol{n}}) = \sum_{\ell \in \Delta_i} \sum_{m=-\ell}^{+\ell} a_{\ell m} Y_{\ell m}(\hat{\boldsymbol{n}}), \tag{57}$$

where the Δ_i are intervals (bins) of multipole values $[\ell_i, \ell_{i+1}-1]$, for $i = 0, \ldots, (N_{\text{bins}} - 1)$, with $\ell_0 = \ell_{\text{min}}$ and $\ell_{N_{\text{bins}}} = \ell_{\text{max}} + 1$, and the other bin boundaries chosen in such a way as to minimize the variance of $\hat{f}_{\text{NL}}$. The binned bispectrum is then obtained by using T_i instead of T_ℓ in the expression for the sample bispectrum of Eq. (30), to obtain:

$$B_{i_1 i_2 i_3}^{\text{bin}} = \int T_{i_1}(\hat{\boldsymbol{n}}) T_{i_2}(\hat{\boldsymbol{n}}) T_{i_3}(\hat{\boldsymbol{n}}) \mathrm{d}^2 \hat{\boldsymbol{n}}. \tag{58}$$

The linear term B^{lin} is binned in an analogous way, and the theoretical bispectrum template B^{th} and variance V are also binned by summing them over the values of ℓ inside the bin. Finally f_{NL} is determined using the binned version of Eq. (35), i.e., by replacing all quantities by their binned equivalent and replacing the sum over ℓ by a sum over bin indices i. An important point is that the binned bispectrum estimator does not mix the observed bispectrum and the theoretical template weights until the very last step of the computation of $\hat{f}_{\text{NL}}$ (the final sum over bin indices). Thus, the (binned) bispectrum of the map is also a direct output of the code. Moreover, one can easily study the ℓ-dependence of the results by omitting bins from this final sum.

The full binned bispectrum allows one to explore the bispectral properties of maps independent of a theoretical model. Binned bispectra have been used to compare component separation maps and single-frequency maps dominated by foregrounds, as presented in Sects. 3.4.2 and 7. In its simplest implementation, which is used in this paper, the binned estimator uses top-hat filters in harmonic space, which makes the Gaussian noise independent between different bins; however, slightly overlapping bins could be used to provide better localization properties in pixel space. In this sense the binned estimator is related to the wavelet estimators, which we discuss below.

3.2.5. Wavelet f_{NL} estimator

Wavelet methods are well-established in the CMB literature and have been applied to virtually all areas of the

data analysis pipeline, including map-making and component separation, point source detection, search for anomalies and anisotropies, cross-correlation with large scale structure and the ISW effect, and many others (see for instance Antoine & Vandergheynst 1998; Martínez-González et al. 2002; Cayon et al. 2003; McEwen et al. 2007a,b; Pietrobon et al. 2006; Starck et al. 2006; Cruz et al. 2007; Faÿ et al. 2008; Feeney et al. 2011; Geller & Mayeli 2009a,b; Starck et al. 2010; Scodeller et al. 2011; Fernández-Cobos et al. 2012). These methods have the advantage of possessing localization properties both in real and harmonic space, allowing the effects of masked regions and anisotropic noise to be dealt with efficiently.

In terms of the current discussion, wavelets can be viewed as a way to compress the sample bispectrum vector by a careful binning scheme in the harmonic domain. See also Planck Collaboration XXIII (2014). In particular, the wavelet bispectrum can be rewritten as

$$q_{ijk} = \frac{1}{4\pi} \frac{1}{\sigma_i \sigma_j \sigma_k} \int \mathrm{d}^2 \hat{\boldsymbol{n}}\, w(R_i, \hat{\boldsymbol{n}}) w(R_j, \hat{\boldsymbol{n}}) w(R_k, \hat{\boldsymbol{n}}), \tag{59}$$

where

$$w(R; \boldsymbol{b}) = \int \mathrm{d}^2 \hat{\boldsymbol{n}}\, f(\hat{\boldsymbol{n}}) \Psi(\hat{\boldsymbol{n}}, \boldsymbol{b}; R) = \sum_{\ell m} a_{\ell m} \omega_\ell(R) Y_{\ell m}(\hat{\boldsymbol{n}}). \tag{60}$$

Here $\boldsymbol{b}$ is the position on the sky at which the wavelet coefficient is evaluated and σ is the dispersion of the wavelet coefficient map $w(R; \boldsymbol{b})$ for the scale R. The filters $\omega_\ell(R)$ can be seen as the coefficients of the expansion into spherical harmonic of the Spherical Mexican Hat Wavelet (SMHW) filter

$$\Psi(\boldsymbol{x}, \boldsymbol{n}; R) = \frac{1}{\sqrt{2\pi}} \frac{1}{N(R)} \left[1 + \left(\frac{y}{2}\right)^2 \right]^2 \left[2 - \left(\frac{y}{R}\right)^2 \right] e^{-y^2/2R^2}. \tag{61}$$

Here $N(R) = R\sqrt{1 + R^2/2 + R^4/4}$ is a normalizing constant and $y = 2\tan(\theta/2)$ represents the distance between $\boldsymbol{x}$ and $\boldsymbol{n}$, evaluated on the stereographic projection on the tangent plane at $\boldsymbol{n}$; θ is the corresponding angular distance, evaluated on the spherical surface.

The implementation of the linear-term correction can proceed in analogy with the earlier cases. However, note that, in view of the real-space localization properties of the wavelet filters, the linear term here is smaller than for KSW and related approaches, although not negligible. Moreover, it can be well-approximated by a term-by-term sample-mean subtraction for the wavelet coefficients, which allows for a further reduction of computational costs. Further details can be found in Curto et al. (2011a,b, 2012); Regan et al. (2013; see also Lan & Marinucci 2008; Rudjord et al. 2009; Pietrobon et al. 2009, 2010; Donzelli et al. 2012, for related needlet-based procedures).

3.3. Wiener filtering

As discussed in Sect. 3.2, the f_{NL} bispectrum estimator requires, in principle, inverse covariance filtering of the data to achieve optimality (equivalent to Wiener filtering up to a trivial multiplication by the inverse of the signal power spectrum).

We have used the iterative method of Elsner & Wandelt (2013) for Wiener filtering simulations and data. The algorithm makes use of a messenger field, introduced to mediate between the pixel space and harmonic space representation, where noise and signal properties can be specified most directly. Since the Wiener filter is the maximum a posteriori solution, we monitor the χ^2 of the current solution as a convergence diagnostics.

We terminate the algorithm as soon as the improvement in the posterior between two consecutive steps has dropped below a threshold of $\Delta\chi^2 \leq 10^{-4}\sigma_{\chi^2}$, where σ_{χ^2} is the standard deviation of χ^2-distributed variables with a number of degrees of freedom given by the number of observed pixels. Results of this method have been cross-validated using an independent conjugate gradient inversion algorithm with multi-grid preconditioning, originally developed for the analysis of WMAP data in Smith et al. (2009). Applying this estimator to simulations pre-processed using the above mentioned algorithms yielded optimal error bars as expected.

However, we found that optimal error bars could also be achieved for all shapes using a much simpler diffusive inpainting pre-filtering procedure that can be described as follows: all masked areas of the sky (both Galactic and point sources) are filled in with an iterative scheme. Each pixel in the mask is filled with the average of all surrounding pixels, and this is repeated 2000 times over all masked pixels (we checked on simulations that convergence of all f_{NL} estimates was achieved with 2000 iterations). Note that the effect of this inpainting procedure, especially visible for the Galactic mask, is effectively to apodize the mask, reproducing small-scale structure near the edges and only large-scale modes in the interior. This helps to prevent propagating any sharp-edge effects or lack of large-scale power in the interior of the mask to the unmasked regions during harmonic transforms.

Any bias and/or excess variance arising from the inpainting procedure were assessed through MC validation (see Sect. 6) and found to be negligible. Since the inpainting procedure is particularly simple to implement, easily allows inclusion of realistic correlated-noise models in the simulations, and retains optimality, we chose inpainting as our data filtering procedure for the f_{NL} analysis.

3.4. CMB bispectrum reconstruction

3.4.1. Modal bispectrum reconstruction

Modal and related estimators can be used to reconstruct the full bispectrum from the modal coefficients β_{ijk} obtained from map filtering with the separable basis functions $Q_{ijk}(\ell_1, \ell_2, \ell_3)$ (Eq. (51)) (Fergusson et al. 2010a). It is easy to show that the expectation value for β_{ijk} (or equivalently β_n) for an ensemble of non-Gaussian maps generated with a CMB bispectrum shape given by α_n (Eq. (47)) (with amplitude f_{NL}) is

$$\langle\beta_n\rangle = f_{\mathrm{NL}}\sum_p \gamma_{np}\alpha_p, \tag{62}$$

where γ_{np} is defined in Eq. (48). Hence, we expect the best fit coefficients for a particular α_n realization to be given by the β_ns themselves (for a sufficiently large signal). Assuming that we can extract the β_n coefficients accurately from a particular experiment, we can directly reconstruct the CMB bispectrum using the expansion of Eq. (46), that is,

$$
\begin{aligned}
b_{\ell_1\ell_2\ell_3} &= \sqrt{C_{\ell_1}C_{\ell_2}C_{\ell_3}}\sum_{n,p}\gamma_{np}^{-1}\beta_p\, Q_n(\ell_1,\ell_2,\ell_3)\\
&= \sqrt{C_{\ell_1}C_{\ell_2}C_{\ell_3}}\sum_n\beta_n^R R_n(\ell_1,\ell_2,\ell_3),
\end{aligned}
\tag{63}
$$

where, for convenience, we have also defined orthonormalized basis functions $R_n(\ell_1, \ell_2, \ell_3)$ with coefficients α_n^R and β_n^R such that $\langle R_n, R_p\rangle = \delta_{np}$. This method has been validated for simulated maps, showing the accurate recovery of CMB bispectra, and it

has been applied to the WMAP-7 data to reconstruct the full 3D CMB bispectrum (Fergusson et al. 2012).

To quantify whether or not there is a model-independent deviation from Gaussianity, we can consider the total integrated bispectrum. By summing over all multipoles, we can define a total integrated nonlinearity parameter $\bar{F}_{\mathrm{NL}}$ which, with the orthonormal modal decomposition of Eq. (63) becomes (Fergusson et al. 2012)

$$\bar{F}_{\mathrm{NL}}^2 = \frac{1}{N_{\mathrm{loc}}^2}\sum_{l_i}\frac{h_{\ell_1\ell_2\ell_3}^2 b_{\ell_1\ell_2\ell_3}^2}{C_{\ell_1}C_{\ell_2}C_{\ell_3}} = \frac{\sum_n \beta_n^{R2}}{\sum_n \alpha_n^{R2}}, \tag{64}$$

where N_{loc} is the normalization for the local $f_{\mathrm{NL}} = 1$ model. The expectation value $\langle\bar{F}_{\mathrm{NL}}^2\rangle$ contains more than just the three-point correlator, with contributions from products of the two-point correlators and even higher-order contributions,

$$\bar{F}_{\mathrm{NL}}^2 \approx \frac{1}{N_{\mathrm{loc}}^2}\left[6n_{\max} + \sum_{n=1}^{n_{\max}}\left(F_{\mathrm{NL}}^2\alpha_n^{R2} + \langle\beta_n^{R2}\rangle_{6\mathrm{pt}} + ...\right)\right]. \tag{65}$$

Here $\bar{F}_{\mathrm{NL}}^2$ is the full 6-point function over the unconnected Gaussian part, i.e., the product of $3C_\ell$. So, for a Gaussian input this recovers an average of 1 per mode. In the non-Gaussian case the leading-order contributions after the Gaussian part are the bispectrum squared and the product of the power spectrum and the trispectrum, which enter at the same order (for additional explanations see Fergusson et al. 2012).

3.4.2. Smoothed binned bispectrum reconstruction

As explained in Sect. 3.2.4, the full binned bispectrum of the maps under study is one of the products of the binned estimator code.

Given the relatively fine binning (about 50 bins up to $\ell = 2000$ or 2500), most of the measurement in any single bin-triplet is noise. If combinations of maps are chosen so that the CMB primordial signal dominates, most of this noise is Gaussian, reflecting the fact that even when $\langle xyz\rangle = 0$, for a particular statistical realization, xyz is almost certainly non-zero. If our goal is to test whether there is any statistically significant signal, then it makes sense to normalize by defining a new field $\mathcal{B}_{i_1 i_2 i_3}$, which is the binned bispectrum divided by its expected standard deviation (computed in the standard way assuming Gaussian statistics). The distribution of $\mathcal{B}_{i_1 i_2 i_3}$ in any one bin-triplet is very nearly Gaussian as a result of the central limit theorem, and there is almost no correlation between different bins.

We could study the significance of the extreme values at this fine resolution, but it also makes sense to smooth in order to detect features coherent over a wider range of ℓ. If the three-dimensional domain over which the bispectrum is defined (consisting of those triplets in the range $[\ell_{\min}, \ell_{\max}]$ satisfying the triangle inequality and parity condition) did not have boundaries, the smoothing would be more straightforward. We smooth with a Gaussian kernel of varying width in $\Delta(\ln\ell)$, normalized so that the smoothed function in each pixel would be normal-distributed if the input map were Gaussian. In this way, based on the extreme values, it is possible to decide on whether there is a statistically significant NG signal in the map in a "blind" (non-parametric) way.

4. Statistical estimation of the CMB trispectrum

4.1. The squeezed-diagonal trispectrum: τ_{NL}

The 4-point function, or equivalently the trispectrum, of the CMB, can also place interesting constraints on inflationary physics. There are several physically interesting "shapes" of the trispectrum (e.g., Huang & Shiu 2006; Byrnes et al. 2006; Fergusson et al. 2010b; Izumi et al. 2012), in analogy to the bispectrum case. In the simplest non-Gaussian models, the CMB bispectrum has larger signal-to-noise than the trispectrum, but there are examples of technically natural models in which the trispectrum has larger signal-to-noise (e.g., Senatore & Zaldarriaga 2011; Baumann & Green 2012; see also Bartolo et al. 2010b). This can happen in models in which the field modulating the fluctuation amplitude is only weakly correlated to the observed large-scale curvature perturbation.

Analysis of the trispectrum is more challenging than that of the bispectrum, due to the increased range of systematic effects and secondary signals which can contribute. For example, gravitational lensing of the CMB generates a many-sigma contribution to the trispectrum, though it has a distinctive anisotropic shape that differs from primordial NG modulated by scalar fields. As an instrumental example, any mismatch between the true covariance of the observed CMB plus noise and the covariance which is assumed in the analysis (due, for example, to mischaracterisation of the pointing, beams, or noise properties) will generally lead to biases in the estimated trispectrum. Due to these challenges, we have deferred a full analysis of the primordial trispectrum to a future paper, and here focus on the simplest squeezed shape that can provide useful constraints on primordial models, τ_{NL}.

τ_{NL} is most easily understood as measuring the large-scale modulation of small-scale power. The constraints on f_{NL} show that such a modulation must be small if correlated with the temperature. However it is possible for multi-field inflation models to produce squeezed-shape modulations which are uncorrelated with the large-scale curvature perturbations. Such models can be constrained by the trispectrum, conventionally parameterized by τ_{NL} in the squeezed-diagonal shape.

For example, consider the case where a small-scale Gaussian curvature perturbation ζ_0 is modulated by another field ϕ so that the primordial perturbation is given by

$$\zeta(\boldsymbol{x}) = \zeta_0(\boldsymbol{x})[1 + \phi(\boldsymbol{x})], \tag{66}$$

where $\phi(\boldsymbol{x})$ is a large-scale modulating field (with amplitude $\ll 1$). The large-scale modes of ϕ can be measured from the modulation they induce in the small-scale ζ power spectrum. If ϕ has a nearly scale-invariant spectrum, the nearly-white cosmic variance noise on the reconstruction dominates on small-scales, so only the very largest modes can be reconstructed (Kogo & Komatsu 2006). A reconstruction of ϕ is going to be limited to only very large-scale variations, in which case the scale of the variation is very large compared to the width of the last-scattering surface; i.e., in any particular direction a large-scale modulating field will modulate all small-scale perturbations through the last-scattering surface by approximately the same amount. This approximation is good at the percent level, and can readily be related to the full trispectrum estimator, as discussed in more detail in Pearson et al. (2012, Sect. IV; see also Okamoto & Hu 2002; Munshi et al. 2011; Smidt et al. 2010).

A large-scale power modulation therefore translates directly into a large-scale modulation of the small-scale CMB temperature:

$$T(\hat{\boldsymbol{n}}) \approx T_{\mathrm{g}}(\hat{\boldsymbol{n}})[1 + \phi(\hat{\boldsymbol{n}}, r_*)] \equiv T_{\mathrm{g}}(\hat{\boldsymbol{n}})[1 + f(\hat{\boldsymbol{n}})], \tag{67}$$

where T_{g} are the usual small-scale Gaussian CMB temperature anisotropies and r_* is the radial distance to the last-scattering surface. We can quantify the trispectrum as a function of modulation scale by using the power spectrum of the modulation,

$$\tau_{\mathrm{NL}}(L) \equiv \frac{C_L^f}{C_L^{\zeta_*}}. \tag{68}$$

As is conventional, we normalize relative to $C_L^{\zeta_*}$, the power spectrum of the primordial curvature perturbation at the location of the recombination surface. The field f is directly observable, but $C_L^{\zeta_*}$ is not, since the curvature perturbation can only be constrained very indirectly on very large scales. We shall therefore give constraints on f, which is directly constrained by *Planck*, but also on τ_{NL} for comparison with the inflation literature. Note that $\tau_{\mathrm{NL}} \sim 500$ corresponds to an $f = O(10^{-3})$ modulation.

A general quadratic estimator methodology for reconstructing f was developed in Hanson & Lewis (2009), which we broadly follow here. The structure is essentially identical to that for lensing reconstruction (Planck Collaboration XVII 2014), where here instead of reconstructing a lensing potential (or deflection angle), we are reconstructing a scalar modulation field. The quadratic maximum likelihood estimator for the large-scale modulation field f (assuming it is small) is given by

$$\hat{f}_{LM} = \mathcal{F}_{LML'M'}^{-1} \left[\bar{f}_{L'M'} - \langle \bar{f}_{L'M'} \rangle \right], \tag{69}$$

where $\bar{f}$ is a quadratic function of the filtered data that can be calculated quickly in real space:

$$\bar{f}_{LM} = \int \mathrm{d}^2 \hat{\boldsymbol{n}} Y_{LM}^* \left[\sum_{\ell_1 m_1} \bar{T}_{\ell_1 m_1}^i Y_{\ell_1 m_1} \right] \left[\sum_{\ell_2 m_2} \tilde{C}_{\ell_2} \bar{T}_{\ell_2 m_2}^j Y_{\ell_2 m_2} \right]. \tag{70}$$

Here $\bar{T}^i = C^{-1} \tilde{T}^i$ is an inverse-variance filtering sky map (which accounts for sky cuts and inhomogeneous noise), and $\tilde{C}_{\ell_2}$ is the lensed C_ℓ. The "mean field" $\bar{f}^{\mathrm{MF}} \equiv \langle \bar{f} \rangle$ can be estimated from simulations, along with the Fisher normalization $\mathcal{F}$ that is given by the covariance of $\bar{f} - \bar{f}^{\mathrm{MF}}$. The i, j indices are included here, since we shall be using different sky maps with independent noise to avoid noise biases at the level of modulation field reconstruction. For low L and high $\ell_{\max}$ the reconstruction noise is very nearly constant (white, because each small patch of sky gives a nearly-uncorrelated but noisy estimate of the small-scale power), and the reconstruction is very local.

In practice the inverse-variance filtering is imperfect, the noise cannot be modelled exactly, and the normalizing Fisher matrix $\mathcal{F}_{\ell m \ell' m'}$ evaluated from simulations would be inaccurate. Instead we focus on $\bar{f}$ directly, which is approximately an inverse-variance weighed reconstruction of the modulation, and is manifestly very local in real space (and hence zero in the cut part of the sky). Since the reconstruction noise, which also approximately determines the normalization, is nearly white (constant in L), $\bar{f} - \bar{f}^{\mathrm{MF}}$ in real space has an expectation nearly proportional to the underlying modulation outside the mask.

We then define an estimator of the modulation power spectrum

$$\hat{C}_L^f = k_L \left(\frac{A_L}{2L+1} \sum_M |\bar{f}_{LM} - \bar{f}_{LM}^{\mathrm{MF}}|^2 - N_L^{(0)} \right) \tag{71}$$

where

$$N_L^{(0)} = \frac{A_L}{2L+1} \sum_M \langle |\bar{f}_{LM} - \bar{f}_{LM}^{\text{MF}}|^2 \rangle_0 \tag{72}$$

is a noise bias for zero signal estimated from simulations. The normalization A_L is the analytic ideal full-sky normalization which is very close to a constant, and k_L is a calibration factor determined from with-signal simulations. On small scales $k_L \propto f_{\text{sky}}^{-1}$, but has some scale dependence: it increases towards $k_L \propto f_{\text{sky}}^{-2}$ at $L = 0$ at very low L. We shall sometimes plot $\bar{C}_L^f \equiv k_L^{-1} \hat{C}_L^f$, corresponding to the uncalibrated reconstruction of the power modulation, which is very local in real space.

For each value of the modulation scale L, Eq. (71) defines a separate estimator for τ_{NL}

$$\hat{\tau}_{\text{NL}}(L) \equiv \frac{\hat{C}_L^f}{C_L^{\zeta_\star}}. \tag{73}$$

We can combine estimators from all ℓ by constructing

$$\hat{\tau}_{\text{NL},1} \approx N^{-1} \sum_{L'=L_{\min}}^{L_{\max}} C_{L'}^{\zeta_\star} \text{cov}_{L'L}^{-1} \hat{C}_L^f, \tag{74}$$

where $N = \sum_{L'=L_{\min}}^{L_{\max}} C_{L'}^{\zeta_\star} \text{cov}_{L'L}^{-1} C_L^{\zeta_\star}$ and $\text{cov}_{LL'}$ is the covariance of $\hat{C}_L^f$ from simulations with $\tau_{\text{NL}} = 0$. On the full-sky the estimators from each L would be independent, but the mask introduces significant coupling between the very low multipoles and this form of the estimator allows us to account for this. In the full-sky uncorrelated approximation, with a nearly scale-invariant primordial spectrum and using the whiteness of the reconstruction noise, the estimator for τ_{NL} simplifies to (Pearson et al. 2012)

$$\hat{\tau}_{\text{NL}} \approx N^{-1} \sum_{L=L_{\min}}^{L_{\max}} \frac{2L+1}{L^2(L+1)^2} \frac{\hat{C}_L^f}{C_L^{\zeta_\star}}, \tag{75}$$

where $N \equiv L_{\min}^{-2} - (L_{\max} + 1)^{-2}$. This result does not require many simulations to estimate the covariance accurately for inversion, and is typically expected to give very similar results with an error bar that is less than 10% larger. We calculate both as a cross-check, but report results for $\hat{\tau}_{\text{NL}}$ because it is more robust, and in our simulation results actually has slightly lower tails (though larger variance). Mean fields and the $N_L^{(0)}$ bias are estimated in all cases from 1000 zero-τ_{NL} simulations, and the mask used retains about 70% of the sky.

If there is a nearly scale-invariant signal, so $C_L^f \propto C_L^\zeta$ as expected in most multi-field inflation models, the contributions fall rapidly $\propto 1/L^3$, as expected when measuring a scale-invariant signal that has large white reconstruction noise. The signal is therefore on very large scales, with typically half the Fisher signal in the dipole modulation and 95% of the signal at $L \leq 4$, justifying the squeezed approximations used. We use $L_{\max} = 10$ for the estimators, which includes almost all of the signal-to-noise but avoids excessive contamination with the "blue" spectrum of lensing contributions. However, due to the small number of modes involved, the posterior distributions of τ_{NL} can have quite broad tails corresponding to the finite probability that all the largest-scale modulation modes just happen to be near zero. To improve constraints on large values it can help to include a larger range of L, and we consider L up to $L_{\max} = 50$, which is about the limit of where the approximations are valid.

5. Non-primordial contributions to the CMB bispectrum and trispectrum

In this subsection we present the steps followed to account for and remove the main non-primordial contributions to CMB NG.

5.1. Foreground subtraction

Foreground emission signals in the microwave bands have a strong non-Gaussian signature. Therefore any residual emission in the CMB data can give a spurious apparent primordial NG detection. In our analysis we considered *Planck* CMB foreground-cleaned maps created using several independent techniques, as described in Planck Collaboration XII (2014): explicit parametrization and fitting of foregrounds in real space (Commander-Ruler, C-R) (Eriksen et al. 2006, 2008); Spectral Matching of foregrounds implementing Independent Component Analysis (SMICA) (Delabrouille et al. 2003; Cardoso et al. 2008); Internal Linear Combination (Needlet Internal Linear Combination, NILC) (Delabrouille et al. 2009); and Internal Template Fitting (SEVEM) (Fernández-Cobos et al. 2012). These and other techniques underwent a pre-launch testing phase (Leach et al. 2008). Each method provides a *Planck* CMB foreground-cleaned map with a confidence mask, which defines the trusted cleaned region of the sky; an estimate of the noise in the output CMB map obtained from half-ring difference maps; and an estimate of the beam transfer function of the processed map. The resolution reaches 5 arcminutes. In addition a union of all the confidence masks, denoted as U73, is provided. Channels from both the Low Frequency Instrument (LFI, Planck Collaboration II 2014) and the High Frequency Instrument (HFI, Planck Collaboration VI 2014) of *Planck* are used to achieve each of the reconstructed CMB templates. The validation of CMB reconstruction through component separation is based on the inspection of several observables, as explained in detail in Planck Collaboration XII (2014): the two-point correlation function and derived cosmological parameters; indicators of NG including the f_{NL} results presented in the present paper; and cross-correlation with known foreground templates. Based on various figures-of-merit, the foreground cleaning techniques performed comparably well (Planck Collaboration XII 2014).

In order to test foreground residuals a battery of simulations is required. In the simulations the foreground emission was modelled with the pre-launch version of the *Planck* Sky Model (PSM), based on observations of the emission from our own Galaxy and known extra-Galactic sources, largely in the radio and infrared bands. The PSM is described in Delabrouille et al. (2013) and includes models of CMB (including a dipole), diffuse Galactic emissions (synchrotron, free-free, thermal dust, Anomalous Microwave Emission and CO molecular lines), emission from compact objects (thermal SZ effect, kinetic SZ effect, radio sources, infrared sources, correlated far-infrared background and ultra-compact H II regions). The sky model includes total intensity as well as polarization, which was not used in this paper.

The PSM has been used to create the sixth round of full focal plane (FFP6) simulations, a set of simulations for the 2013 data release based on detailed models of the sky and instrument (e.g., noise properties, beams, satellite pointing and map-making process), consisting of both Gaussian and non-Gaussian CMB realizations. A description of the FFP6 simulations can be found in Appendix A (see Planck Collaboration 2013, for details). The FFP6 set has been used to test and validate the component separation algorithms employed in *Planck* and to establish

uncertainties on the outputs (Planck Collaboration XII 2014). We will also use FFP6 simulations in the next sections in order to validate our analysis (see Sects. 6–8).

5.2. The integrated Sachs-Wolfe-lensing bispectrum

One of the most relevant mechanisms that can generate NG from secondary CMB anisotropies is the coupling between weak lensing and the ISW (Sachs & Wolfe 1967) effect. This is in fact the leading contribution to the CMB secondary bispectrum with a blackbody frequency dependence (Goldberg & Spergel 1999; Verde & Spergel 2002; Giovi et al. 2005).

Weak lensing of the CMB is caused by gradients in the matter gravitational potential that distorts the CMB photon geodesics. The ISW on the other hand arise because of time-varying gravitational potentials due to the linear and nonlinear growth of structure in the evolving Universe. Both the lensing and the ISW effect are then related to the matter gravitational potential and thus are correlated phenomena. This gives rise to a non-vanishing 3-point correlation function. Furthermore, lensing is related to nonlinear processes which are therefore non-Gaussian. A detailed description of the signal, which accounts also for the contribution from the early-ISW effect, can be found in Lewis (2012).

The ISW-lensing bispectrum takes the form:

$$B_{\ell_1\ell_2\ell_3}^{m_1m_2m_3} \equiv \langle a_{\ell_1m_1}a_{\ell_2m_2}a_{\ell_3m_3}\rangle = \langle a_{\ell_1m_1}^{\mathrm{P}}a_{\ell_2m_2}^{\mathrm{L}}a_{\ell_3m_3}^{\mathrm{ISW}}\rangle + 5 \text{ perm.}, \quad (76)$$

where P, L, and ISW indicate primordial, lensing and ISW contributions respectively. This becomes

$$B_{\ell_1\ell_2\ell_3}^{m_1m_2m_3\,(\mathrm{ISW-L})} = \mathcal{G}_{\ell_1\ell_2\ell_3}^{m_1m_2m_3}b_{\ell_1\ell_2\ell_3}^{\mathrm{ISW-L}}, \quad (77)$$

where $\mathcal{G}_{\ell_1\ell_2\ell_3}^{m_1m_2m_3}$ is the Gaunt integral and $b_{\ell_1\ell_2\ell_3}^{\mathrm{ISW-L}}$ is the reduced bispectrum given by

$$b_{\ell_1\ell_2\ell_3}^{\mathrm{ISW-L}} = \frac{\ell_1(\ell_1+1) - \ell_2(\ell_2+1) + \ell_3(\ell_3+1)}{2}$$
$$\times \tilde{C}_{\ell_1}^{\mathrm{TT}}C_{\ell_3}^{\mathrm{T}\phi} + (5 \text{ perm.}). \quad (78)$$

Here $\tilde{C}_\ell^{\mathrm{TT}}$ is the lensed CMB power spectrum and $C_\ell^{\mathrm{T}\phi}$ is the ISW-lensing cross-power spectrum (Lewis 2012; Goldberg & Spergel 1999; Verde & Spergel 2002; Cooray & Hu 2000) that expresses the statistical expectation of the correlation between the lensing and the ISW effect.

As shown in Hanson et al. (2009b), Mangilli & Verde (2009), and Lewis et al. (2011), the ISW-lensing bispectrum can introduce a contamination in the constraints on primordial local NG from the CMB bispectrum. Both bispectra are maximal for squeezed or nearly squeezed configurations. The bias on a primordial f_{NL} (e.g., local) due to the presence of the ISW-lensing cross correlation signal is

$$\Delta f_{\mathrm{NL}}^{\mathrm{local}} = \frac{\hat{S}}{N}, \quad (79)$$

with

$$\hat{S} = \sum_{2\leqslant\ell_1\ell_2\ell_3} \frac{B_{\ell_1\ell_2\ell_3}^{\mathrm{ISW-L}}B_{\ell_1\ell_2\ell_3}^{\mathrm{P}}}{V_{\ell_1\ell_2\ell_3}}, \quad N = \sum_{2\leqslant\ell_1\ell_2\ell_3} \frac{\left(B_{\ell_1\ell_2\ell_3}^{\mathrm{P}}\right)^2}{V_{\ell_1\ell_2\ell_3}}, \quad (80)$$

where $B^{\mathrm{ISW-L}}$ and B^{P} refer respectively to the ISW-lensing and the primordial bispectrum, and V is defined below Eq. (35).

The bias in the estimation of the three primordial f_{NL} from *Planck* is given in Table 1. As one can see, taking into account

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

	ISW-lensing f_{NL} bias			
Shape	SMICA	NILC	SEVEM	C-R
Local	7.1	7.0	7.1	6.0
Equilateral	0.4	0.5	0.4	1.4
Orthogonal	−22	−21	−21	−19

Table 2. Results for the amplitude of the ISW-lensing bispectrum from the SMICA, NILC, SEVEM, and C-R foreground-cleaned maps, for the KSW, binned, and modal (polynomial) estimators; error bars are 68% CL.

	ISW-lensing amplitude			
Method	SMICA	NILC	SEVEM	C-R
KSW	0.81 ± 0.31	0.85 ± 0.32	0.68 ± 0.32	0.75 ± 0.32
Binned	0.91 ± 0.37	1.03 ± 0.37	0.83 ± 0.39	0.80 ± 0.40
Modal	0.77 ± 0.37	0.93 ± 0.37	0.60 ± 0.37	0.68 ± 0.39

the f_{NL} statistical error bars shown, e.g., in Table 8, the local shape is most affected by this bias (at the level of more than $1\sigma_{\mathrm{local}}$), followed by the orthogonal shape (at the level of about $0.5\sigma_{\mathrm{ortho}}$), while the equilateral shape is hardly affected. In this paper we have taken into account the bias reported in Table 1 by subtracting it from the measured f_{NL}[8].

The results for the amplitude of the ISW-lensing bispectrum from the different foreground-cleaned maps are given in Table 2. It should be noted that the binned and modal estimators are less correlated to the exact template for the ISW-lensing shape than they are for the primordial shapes, hence their larger error bars compared to KSW (which uses the exact template by construction Mangilli et al. 2013). The conclusion is that we detect the ISW-lensing bispectrum at a value consistent with the fiducial value of 1, at a significance level of 2.6σ (taking the SMICA-KSW value as reference). For details about comparisons between different estimators and analysis of the data regarding primordial shapes we refer the reader to Sects. 6 and 7.

We show for the SMICA map in the top figure of Fig. 2 the measured skew-C_ℓ spectrum (see Sect. 3.2.2) for optimal detection of the ISW-lensing bispectrum, along with the best-fitting estimates of f_{NL} from the KSW method for different values of ℓ. It should be noted that the skew-C_ℓ spectrum is *not* a fit to the KSW data points; its shape is fully fixed by the template under consideration, with only the overall amplitude as a free parameter. Hence the agreement between the curve and the points is good evidence that KSW is really detecting the ISW-lensing effect and not some other source of NG. Note that point sources, at the level determined by their own skew-spectrum, do not contribute significantly to the ISW-lensing statistic). See Planck Collaboration XVII (2014) and Planck Collaboration XIX (2014) for further information about the detection by *Planck* of the ISW-lensing signal.

5.3. Point-sources bispectrum

Extra-Galactic point sources at *Planck* frequencies are divided into two broad categories: radio sources with synchrotron and/or

---

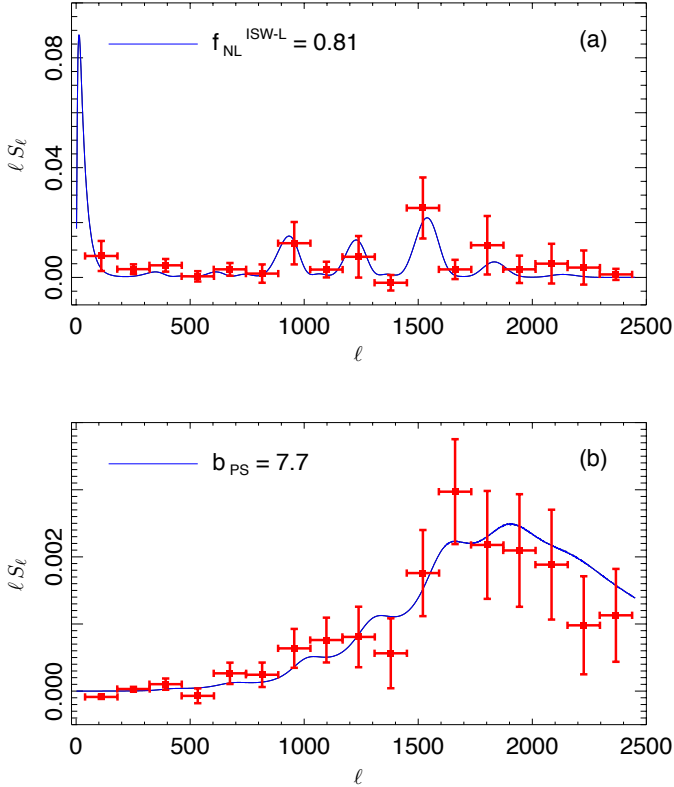[8] See Kim et al. (2013) for other debiasing techniques.

Fig. 2. Binned skew-C_ℓ statistics from the SMICA map for **a)** ISW-lensing and **b)** Poisson point sources. Theoretical curves are not fitted to the data shown, but are plotted with the amplitude (the only free parameter) determined from the KSW technique. The Poisson point-source foreground is clearly detected, and the ISW-lensing skew-spectrum is evident and consistent with the overall 2.6σ detection. b_{ps} is the Poisson point-source amplitude in dimensionless units of 10^{-29}, and f_{NL}^{ISW-L} is the ISW-lensing amplitude in units of that expected from the *Planck* best-fit cosmology. Error bars come from covariance estimates from 1000 simulated maps, and the points are mildly correlated.

Table 3. Results for the amplitude of the point source (Poisson) bispectrum (in dimensionless units of 10^{-29}) from the SMICA, NILC, SEVEM, and C-R foreground-cleaned maps, for the KSW, binned, and modal (polynomial) estimators; error bars are 68% CL.

	Point source bispectrum amplitude/10^{-29}			
Method	SMICA	NILC	SEVEM	C-R
KSW	7.7 ± 1.5	9.2 ± 1.7	7.6 ± 1.7	1.1 ± 5.1
Binned	7.7 ± 1.6	8.2 ± 1.6	7.5 ± 1.7	0.9 ± 4.8
Modal	$10 \quad \pm 3$	$11 \quad \pm 3$	$10 \quad \pm 3$	0.5 ± 6

Notes. Note that the KSW and binned estimators use $\ell_{max} = 2500$, while the modal estimator has $\ell_{max} = 2000$.

sources by jointly fitting primordial and Poisson templates to the data.

Our final measured point-source bispectrum amplitudes from the data are reported in Table 3. The amplitude is expressed in dimensionless units, i.e., it has been divided by the appropriate power of the monopole temperature T_0, and has been multiplied by 10^{29}. As shown in Sect. 8.1, the Poisson template is the only one that still evolves significantly between $\ell = 2000$ and $\ell = 2500$. This explains the differences between the values of the KSW and binned (that use $\ell_{max} = 2500$) and the modal (that uses $\ell_{max} = 2000$) estimators. It has been shown that for the same value of ℓ_{max} all three estimators agree very well.

We finally conclude from Table 3 that we detect the point-source bispectrum with high significance in the SMICA, NILC, and SEVEM cleaned maps, while it is absent from the C-R cleaned map. The measured skew-C_ℓ spectrum of the SMICA map in the bottom figure of Fig. 2 gives further evidence that the NG from foreground point sources is convincingly detected. The only degree of freedom in this plot is the amplitude, which is not set by a direct fit to the skew-C_ℓ, but rather is estimated by KSW. As a result, the good agreement with the shape of this skew-C_ℓ spectrum is powerful evidence that there is NG from point sources. However, this still turns out to be a negligible contaminant for primordial f_{NL} studies, due to the very low correlation between the Poisson bispectrum and the primordial shapes.

5.4. Non-primordial contributions to the trispectrum

The main non-instrumental source of non-primordial signal is the kinematic modulation dipole due to the peculiar velocity of the earth, $\boldsymbol{v}$, whose magnitude is $O(10^{-3})$ (Challinor & van Leeuwen 2002; Kosowsky & Kahniashvili 2011; Amendola et al. 2011). If data are used to constrain τ_{NL} using the dipole modulation (which shrinks the Fisher error by a factor of two relative to starting at $L = 2$), the dipole-induced signal must be subtracted, since its modulation reconstruction has signal-to-noise larger than unity at *Planck* resolution. Confirmation that this signal is detected with the expected magnitude and direction is a good test of our methodology. The dipole signal seen by *Planck* is studied in detail in Planck Collaboration XXVII (2014), so we only summarize the key points here.

The local Doppler effect modulates the observed CMB temperature $T_0[1 + \Delta T(\hat{\boldsymbol{n}})]$ by $1 + \hat{\boldsymbol{n}} \cdot \boldsymbol{v}$ at leading order, so that $T(\hat{\boldsymbol{n}}) = T_0[1 + \Delta T(\hat{\boldsymbol{n}})](1 + \hat{\boldsymbol{n}} \cdot \boldsymbol{v})$. The spectrum in each direction remains a blackbody, but the relative response in the intensity $I_\nu(\nu, \hat{\boldsymbol{n}})$ at the observed *Planck* frequencies is however frequency dependent. The effective thermodynamic fractional temperature anisotropy $\Delta\Theta$ at each frequency for zero peculiar

free-free emission; and infrared galaxies with thermal emission from dust heated by young stars. Radio sources are dominant at central CMB frequencies up to 143 GHz, and can be considered unclustered (Toffolatti et al. 1998; González-Nuevo et al. 2005). Hence their bispectrum is constant and is related to their number counts as

$$b_{ps} = k_\nu^3 \int_0^{S_c} S^3 \frac{dn}{dS} dS, \tag{81}$$

with S the flux density, dn/dS the number counts per steradian, S_c the flux cut and k_ν the conversion factor from flux to relative temperature elevation, depending on the frequency and instrumental bandpass.

Infrared galaxies become important at higher frequencies, 217 GHz and above, and are highly clustered in dark matter halos, which enhances their bispectrum on large angular scales (Lacasa et al. 2012; Curto et al. 2013). However, in the *Planck* context it was shown by Lacasa & Aghanim (2014) that the IR bispectrum is more than 90% correlated with the Poissonian template of the radio sources. So a joint estimation of f_{NL} with a Poissonian bispectrum template will essentially account for the IR signal, and provide quasi-identical values compared to an analysis accounting for the IR bispectrum template. Indeed, in our final optimal bispectrum constraints for primordial shapes, we will account for the potential contamination from point

velocity is defined by

$$I_\nu(\nu, \hat{\boldsymbol{n}}) = I_\nu(\nu) \left[1 + \left. \frac{\text{d}\ln I_\nu}{\text{d}\ln T} \right|_\nu \Delta\Theta(\hat{\boldsymbol{n}}) \right]. \qquad (82)$$

With peculiar velocity the temperature T depends on the second order term $\Delta T \hat{\boldsymbol{n}} \cdot \boldsymbol{v}$, so expanding the Planck function to second order then gives a change in the effective small-scale temperature anisotropy from both first and second order terms in the expansion of I_ν:

$$\begin{aligned} \Delta\Theta(\hat{\boldsymbol{n}}) &\rightarrow \left[1 + \hat{\boldsymbol{n}} \cdot \boldsymbol{v} + T \frac{\text{d}^2 I_\nu/\text{d}T^2}{\text{d}I_\nu/\text{d}T} \hat{\boldsymbol{n}} \cdot \boldsymbol{v} \right] \Delta\Theta(\hat{\boldsymbol{n}}) \\ &= (1 + [x \coth(x/2) - 1] \, \hat{\boldsymbol{n}} \cdot \boldsymbol{v}) \, \Delta\Theta(\hat{\boldsymbol{n}}), \end{aligned} \qquad (83)$$

where $x \equiv h\nu/k_B T_0$ (and we neglect small second-order nonmodulation terms). Thus the anisotropies in the *Planck* maps have a dipolar modulation given by $1 + b_\nu \hat{\boldsymbol{n}} \cdot \boldsymbol{v}$, where for the frequency bands we use $b_{143} \approx 2$, $b_{217} \approx 3$, and $\beta \equiv |\boldsymbol{v}| = 1.23 \times 10^{-3}$ in the direction of CMB dipole. In addition our peculiar velocity induces kinetic aberration, which looks at leading order exactly like a dipole lensing convergence and only projects weakly into the power anisotropy estimator. For constraining τ_{NL} both of the expected kinematic signals can be included in the simulations, and hence subtracted in the mean field of the modulation reconstruction.

Secondary effects are dominated by the significant and very blue lensing signal. However unlike for the f_{NL} bispectrum lensing only overlaps with τ_{NL} at a small fraction of the error bar as long as only low modulation multipoles $L \lesssim 10$ are used, where the τ_{NL} signal peaks (Pearson et al. 2012). We include lensing in the simulations, so lensing is straightforwardly accounted for in our analysis by its inclusion in the $N_L^{(0)}$ noise bias (Eq. (72)) and mean field.

A variety of instrumental effects can also give a spurious modulation signal if not modelled accurately. In particular the mean field due to anisotropic noise is very large (Hanson et al. 2009a). On the ultra-large scales of interest for τ_{NL}, our understanding of the noise is not adequate to calculate accurately and subtract this large signal. Instead, as for the power spectrum estimation, we use cross-map estimators that have no noise mean field on average. Both the noise and most other instrumental effects such as gain variations are expected to produce a signal with approximate symmetry about the ecliptic plane. Our modulation reconstruction methodology is especially useful here, since we can easily inspect the orientation of any signal found; for example a naive treatment of the noise not using cross-maps would give a large apparent quadrupolar modulation signal aligned with the ecliptic, corresponding to percent-level misestimation of the noise mean field from inaccurate noise simulation.

Beam asymmetries are included in the simulation, as described in Planck Collaboration XVII (2014), but their effect is very small, since the modulation we are reconstructing is isotropic.

Since we are reconstructing a modulation of small-scale power, the estimator is totally insensitive to smooth large-scale foregrounds. However large-scale variation in small-scale foreground power can mimic a trispectrum modulation. We project out 857 GHz as a dust template in our inverse-variance filtering procedure, as described in Appendix A of Planck Collaboration XVII (2014), but do not include any other foreground model in the trispectrum analysis. Any unmodelled foreground power variation would increase the τ_{NL} signal, so our modelling is sufficient to place a robust upper limit.

6. Validation tests

The f_{NL} results quoted in this paper have all been cross-validated using multiple bispectrum-based estimators from different groups. Having multiple estimators was extremely useful for the entire analysis, for two main reasons. First, it allowed great improvement in the robustness of the final results. In the early stages of the work the comparison between different independent techniques helped to resolve bugs and other technical issues in the various computer codes, while during the later stages it was very useful to understand the data and find the optimal way of extracting information about the various bispectrum templates. Secondly, besides these cross-checking purposes, different estimators provide also interesting complementary information, going beyond simple f_{NL} estimation. For example, the binned and modal estimators provide a reconstruction of the full bispectrum of the data (smoothed in different domains), the skew-C_ℓ estimator allows monitoring of the contribution to f_{NL} from different sources of NG, the wavelets reconstruction allows f_{NL} directionality tests, and so on.

In this section we are concerned with the first point above, that is, the use of multiple bispectrum-based pipelines as a way to improve the robustness of the results. For this purpose, a large amount of work was dedicated to the development and analysis of various test maps, in order to validate the estimators. This means not only checking that the various estimators recover the input f_{NL} within the expected errors, but also that the results agree on a map-by-map basis.

The section is split into two parts. Section 6.1 shows results on a set of initially full-sky, noiseless, Gaussian CMB simulations, to which we add, in several steps, realistic complications, including primordial NG, anisotropic coloured noise, and a mask, showing the impact on the results at each step. In Sect. 6.2 we show our results on a set of simulations that mimic the real data as closely as possible (except for the presence of foreground residuals, which will be studied in Sect. 8.4): no primordial NG, but NG due to the ISW-lensing effect; simulated instrumental effects and realistic noise; and simulations passed through the component separation pipelines. In fact these are the FFP6 simulations (see Appendix A) that are used to determine the error bars for the final *Planck* results.

We present here only a small subset of the large number of validation tests that were performed. For example, we also had a number of "blind f_{NL} challenges", in which the different groups received a simulated data set with an unknown value of input f_{NL} for a given shape and they had to report their estimated values. In addition different noise models were tested (white vs. coloured and isotropic vs. anisotropic), leading to the conclusion that it is important to make the noise in the estimator calibration as realistic as possible (coloured and anisotropic). We also tested different Galactic and point source masks, with and without inpainting, concluding that it is best to fill in both the point sources and the Galactic mask, using a sufficient number of iterations in our diffusive procedure to entirely fill in the point source gaps, while at the same time only effectively apodizing the Galactic mask (no small-scale structure in its interior). There were also various tests on FFP simulations of *Planck* data with Gaussian or non-Gaussian CMB and all foregrounds, provided by the PSM (see Appendix A.4). These simulations were tested both before and after they had passed through the component separation pipelines. In all comparison tests the results were consistent with input f_{NL} values and differences between estimators were consistent with theoretical expectations.

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

The aim of the first set of validation tests is threefold. First, we want to study the level of agreement from different estimators in ideal conditions (i.e., full-sky noiseless data). The expected scatter between measurements is, in this case, entirely due to the slightly imperfect correlation between weights of estimators that adopt different schemes to approximate the primordial shape templates. For this case the scatter can be computed analytically (see Appendix B for details). We can then verify that our results in ideal conditions match theoretical expectations. This is done in Sect. 6.1.1. Second, we want to make sure that the estimators are unbiased and correctly recover $f_{\rm NL}$ in input for local, equilateral, and orthogonal shapes. This is done in Sects. 6.1.2 and 6.1.3, where a superposition of local, equilateral and orthogonal bispectra is included in the simulations and the three $f_{\rm NL}$ values are estimated both independently and jointly. Finally we want to understand how much the agreement between pipelines in ideal conditions is degraded when we include a realistic correlated noise component and a sky cut, thus requiring the introduction of a linear term in the estimators in order to account for off-diagonal covariance terms introduced by the breaking of rotational invariance. Since we want to study the impact of adding noise and masking separately, we will first work on a set of full-sky maps with noise in Sect. 6.1.2, and then add a mask in Sect. 6.1.3.

The tests were applied to the KSW, binned and modal estimators. These are the bispectrum pipelines used to analyse *Planck* data in Sect. 7. Our goal for this set of tests is not so much to attain the tightest possible agreement between methods, as it is to address the points summarized in the above paragraph. For this reason the estimator implementations used in this specific Section were slightly less accurate but faster to compute than those adopted for the final data analysis of Sect. 7. The primary difference with respect to the main analysis is that a smaller number of simulations was used to calibrate the linear term (80–100 in these tests, as against 200 or more for the full analysis). For the modal estimator we also use a faster expansion with a smaller number of modes: 300 here versus 600 in the high accuracy version of the pipeline[9] used in Sect. 7. Even with many fewer modes, the modal estimator is still quite accurate: the correlation coefficient for the modal expansion of the local template is 0.95, while for the equilateral and orthogonal shapes it is 0.98.

6.1.1. Ideal Gaussian simulations

As a basis for the other tests we start with the ideal case, a set of 96 simulations of a full-sky Gaussian CMB, with a Gaussian beam with FWHM 5 arcmin and without any noise, cut off at $\ell_{\rm max} = 2000$ in our analyses. The independent Fisher matrix error bars in that case are 4.2 for local NG, 56 for equilateral, and 28 for orthogonal.

Note that this test does not make sense for all estimators, and hence results are not included for all of them. For example, for the binned estimator the optimal binning depends on the noise. While this dependence is not very strong, the difference between no noise and *Planck* noise is sufficiently large that a completely

different binning would have to be used just for this test, going against the purpose of this section to validate the estimators as used for the data analysis[10].

The purpose here is mostly aimed at checking consistency with the following formula (derived in Appendix B) for the expected scatter (standard deviation) between $f_{\rm NL}$ results of the same map from an exact and an approximate estimator:

$$\sigma_{\delta f_{\rm NL}} = \Delta_{\rm th} \frac{\sqrt{1 - r^2}}{r}. \tag{84}$$

Here $\Delta_{\rm th}$ is the standard deviation of the exact estimator and r is the correlation coefficient that gives the correlation of the approximate bispectrum template with the exact one, defined as

$$r \equiv \frac{\sum_{\ell_1 \le \ell_2 \le \ell_3} \frac{B^{\rm th}_{\ell_1 \ell_2 \ell_3} B^{\rm exp}_{\ell_1 \ell_2 \ell_3}}{g_{\ell_1 \ell_2 \ell_3} C_{\ell_1} C_{\ell_2} C_{\ell_3}}}{\sqrt{\sum_{\ell_1 \le \ell_2 \le \ell_3} \frac{(B^{\rm th}_{\ell_1 \ell_2 \ell_3})^2}{g_{\ell_1 \ell_2 \ell_3} C_{\ell_1} C_{\ell_2} C_{\ell_3}} \sum_{\ell_1 \le \ell_2 \le \ell_3} \frac{(B^{\rm exp}_{\ell_1 \ell_2 \ell_3})^2}{g_{\ell_1 \ell_2 \ell_3} C_{\ell_1} C_{\ell_2} C_{\ell_3}}}}, \tag{85}$$

where the label "th" denotes the initial bispectrum shape to fit to the data, and "exp" is the approximate expanded one. Note that this formula has been obtained under the simplifying assumptions of Gaussianity, full-sky coverage and homogeneous noise. For applications dealing with more realistic cases we might expect the scatter to become larger, while remaining qualitatively consistent.

The results averaged over the whole set of maps are given in Table 4 for the KSW and modal estimators individually, as well as for their difference. The plane wave modal expansion implemented here achieves about 98% correlation with the separable shapes used by KSW. According to the formula above we then expect a standard deviation of map-by-map differences of order $0.2\Delta_{f_{\rm NL}}$ for a given shape, where $\Delta_{f_{\rm NL}}$ is the corresponding $f_{\rm NL}$ error bar. Looking at the left-hand side of Table 4, we see that the error bars are 4 for local NG, 50 for equilateral, and 30 for orthogonal. So we predict a standard deviation of map-by-map differences of 0.8, 10 and 6 for local, equilateral, and orthogonal NG, respectively. As one can see from the "Modal-KSW" column, the measurements are in excellent agreement with the theoretical expectation.

6.1.2. Non-Gaussian simulations with realistic noise

A set of 96 full-sky non-Gaussian CMB simulations was created according to the process described by Fergusson et al. (2010a), with local $f_{\rm NL}^{\rm local} = 12$, equilateral $f_{\rm NL}^{\rm equil} = 35$, and orthogonal $f_{\rm NL}^{\rm ortho} = -22$. The effect of a 5 arcmin beam was added, as well as realistic coloured and anisotropic noise according to the specifications of the SMICA cleaned map. The independent Fisher matrix error bars in that case are 5.3 for local, 63 for equilateral, and 33 for orthogonal NG, while the joint ones are respectively 6.0, 64, and 37.

The results averaged over the whole set are given in Table 5 for the various estimators individually, as well as for the differences with respect to KSW. Compared to the previous case we now deviate from the exact theoretical expectation for two reasons: we include a realistic correlated noise component; and we

---

[9]   While most of the modal results in this paper come from the most accurate 600 modes pipeline, a few computationally intensive data validation tests of Sect. 8 also use the fast 300 modes version; therefore the results in this section also provide a direct validation of the fast modal pipeline.

[10]   While the binning with 48 bins and $\ell_{\rm max} = 2000$ used in the validation tests of Sects. 6.1.2 and 6.1.3 is also slightly different from the binning used for the data analysis with 51 bins and $\ell_{\rm max} = 2500$, these differences are very small and the binnings have very similar correlation coefficients of 0.99 or more for local and equilateral shapes, and 0.95 for orthogonal.

Table 4. Results for f_{NL} for the set of ideal Gaussian simulations described in Sect. 6.1.1 for the KSW and modal estimators and for their difference, assuming all shapes to be independent.

Shape	f_{NL}		
	KSW	Modal	Modal − KSW
Local	-0.5 ± 4.1	-0.5 ± 4.1	-0.05 ± 0.63
Equilateral	2.2 ± 48	1.3 ± 48	-0.9 ± 8.9
Orthogonal	-1.1 ± 29	-1.0 ± 30	0.1 ± 6.5

have NG in the maps. The presence of NG in the input maps will lower the agreement between estimators with respect to the Gaussian case if the correlation between weights is not exactly 100%. This is even more true in this specific case, where NG of three different kinds is present in the input maps and also cross-correlation terms between different expanded shapes are involved (and propagated over in the joint analysis). Moreover, when noise is included the specific modal expansion used for this test is 95% correlated to the separable KSW local shape (so there is a 3% reduction of the correlation compared to the ideal case for the modal local shape); we thus expect a further degradation of the level of agreement for this specific case. Finally, in order to correct for noise effects, a linear term has to be added to the estimators. Since the linear term is obtained by MC averaging over just 80 or 96 simulations in this test (depending on the estimator), MC errors are also adding to the measured differences. Of course the MC error can be reduced by increasing the number of simulations in the linear term sample. We do this for the analysis of the real data and in Sect. 6.2, but it was computationally too expensive for this set of preliminary validation tests, so we decided here to just account for it in the final interpretation of the results.

As a consequence of the above, we can no longer expect the map-by-map f_{NL} differences to follow perfectly the theoretical expectation, obtained in the previous section in idealized conditions (full-sky, no noise, and Gaussianity). With these caveats in mind, the agreement between different pipelines remains very good, being about 0.3σ in most cases and about 0.5σ for the modal-KSW difference in the local case, which can be easily explained by the fact that this is the set of weights with the lowest correlation (95%, as mentioned above). All estimators are unbiased and recover the correct input values.

6.1.3. Impact of the mask

To the simulations of Sect. 6.1.2 we now apply the *Planck* union mask – denoted U73 – masking both the Galaxy and the brightest point sources and leaving 73% of the sky unmasked (Planck Collaboration XII 2014). This is the same mask used to analyse *Planck* data in Sect. 7. The independent Fisher matrix error bars in that case (taking into account the f_{sky} correction) are 6.2 for local NG, 74 for equilateral, and 39 for orthogonal, while the joint ones are respectively 7.1, 76, and 44.

All masked areas of the sky (both Galactic and point sources) are filled in with a simple iterative method. In this simple inpainting method each pixel in the mask is filled with the average of all eight surrounding pixels, and this is repeated 2000 times over all masked pixels. The filling-in helps to avoid propagating the effect of a sharp edge and the lack of large-scale power inside the mask to the unmasked regions during harmonic transforms.

This inpainting method is the one that was used to produce all NG results in this paper for methods that need it (KSW, binned and modal).

The results averaged over the whole set of simulations are given in Table 6 for the various estimators individually, as well as for the differences with respect to KSW. The map-by-map results are shown in Fig. 3.

This is the most realistic case we consider in this set of tests. Besides noise, we also include a sky cut and our usual mask inpainting procedure. All the caveats mentioned for the previous case are still valid, and possibly emphasized by the inclusion of mask and inpainting. In the light of this, the agreement is still very good, worsening a bit with respect to the "full-sky + noise" case only for the local measurement, where the mask is indeed expected to have the biggest impact. In the joint analysis all estimators recover the correct input values for the local and orthogonal cases, but all estimators find a value for equilateral NG that is somewhat too low. It is unclear whether this is an effect of masking and inpainting on the equilateral measurement or just a statistical fluctuation for this set of simulations. In any case, this potential bias is small compared to the statistical uncertainty, so that it would not have a significant impact on the final results.

To summarize the results of this Sect. 6.1, we performed an extensive set of validation tests between different f_{NL} estimators using strongly, but not perfectly, correlated primordial NG templates in their weights. The test consisted in comparing the f_{NL} measured by the different estimators for different sets of simulations, on a map-by-map basis. We started from ideal conditions: full-sky Gaussian noiseless maps. In this case we computed a theoretical formula providing the expected standard deviation of the f_{NL} differences, as a function of the correlations between the input NG templates in the different estimators. Our results match this formula very well. In the other two simulation sets we added realistic features (noise, mask and inpainting) and we included a linear combination of local, equilateral and orthogonal NG. First of all we verified that all the pipelines correctly recover the three f_{NL} input values, hence they are unbiased. Moreover, we observed that adding such features produces an expected slight degradation of the level of agreement between different pipelines, that nevertheless remains very good: about 0.3–0.4σ for equilateral and orthogonal NG, and about 0.5–0.6σ for local NG, which is the shape most affected by mask and noise contamination.

6.2. Validation of estimators on realistic *Planck* simulations

In the tests of the previous subsection we checked the bias of the estimators and studied their level of agreement, given the correlation between their weights, in the presence of noise and a sky cut. To speed up the computation while still retaining enough accuracy for the purposes of that analysis, we used a relatively small number of maps for linear term calibrations (80–100) and used a smaller number of modes than usual in the modal estimator. In the present subsection we instead try to simulate as accurately as possible real data analysis conditions. Our goal is to obtain an accurate MC-based expectation of the scatter between different f_{NL} measurements when the pipelines are run on actual *Planck* maps.

To this aim we use FFP6 simulation maps described in Appendix A. The original FFP6 maps were lensed using the `Lenspix` algorithm, and processed through the `SMICA` component separation pipeline. They were then multiplied by the Galactic and point source mask U73 as in the actual f_{NL} analysis, and inpainted as usual. Since our final results show full

Table 5. Results from the different estimators for f_{NL} for the set of full-sky non-Gaussian simulations described in Sect. 6.1.2.

	f_{NL}				
Shape	KSW	Binned	Modal	Binned − KSW	Modal − KSW
Independent					
Local	13.8 ± 5.2	14.1 ± 5.2	14.1 ± 5.3	0.3 ± 2.1	0.4 ± 2.6
Equilateral	63 ± 57	62 ± 58	64 ± 57	−0.9 ± 20	1.0 ± 18
Orthogonal	−52 ± 37	−58 ± 40	−54 ± 37	−6.0 ± 13	−2.2 ± 12
Joint					
Local	11.7 ± 6.2	12.0 ± 6.6	12.0 ± 6.4	0.2 ± 2.7	0.2 ± 3.2
Equilateral	31 ± 59	29 ± 61	31 ± 59	−1.8 ± 21	−0.2 ± 19
Orthogonal	−20 ± 43	−22 ± 47	−21 ± 42	−2.1 ± 16	−0.6 ± 15

Notes. The simulations have $f_{NL} = 12, 35, −22$ respectively. Both the results for the estimators individually and for the differences with KSW are given.

Table 6. Results from the different estimators for f_{NL} for the set of masked non-Gaussian simulations described in Sect. 6.1.3.

	f_{NL}				
Shape	KSW	Binned	Modal	Binned − KSW	Modal − KSW
Independent					
Local	13.5 ± 7.1	13.1 ± 6.5	14.0 ± 6.8	−0.3 ± 3.5	0.5 ± 4.6
Equilateral	55 ± 64	50 ± 59	58 ± 63	−4.4 ± 24	3.3 ± 20
Orthogonal	−50 ± 45	−53 ± 46	−52 ± 45	−3.5 ± 16	−1.9 ± 15
Joint					
Local	11.7 ± 8.3	11.4 ± 7.9	12.2 ± 8.4	−0.3 ± 4.3	0.4 ± 5.7
Equilateral	23 ± 66	19 ± 59	24 ± 64	−3.8 ± 28	1.7 ± 25
Orthogonal	−18 ± 51	−20 ± 54	−18 ± 55	−1.3 ± 20	0.3 ± 20

Notes. Both the results for the estimators individually and for the differences with KSW are given.

consistency with Gaussianity for local, equilateral and orthogonal shapes, we do not include any primordial f_{NL} in these maps. We note that although the simulations were passed through SMICA in order to provide a realistic filtering of the data, they did not include any foreground components. The impact of foreground residuals will be studied separately in Sect. 8.4.

The configuration of all bispectrum pipelines was the same as used for the final data analysis, which implies a correlation of 99% or better between the weights of the KSW, binned and modal estimators. Linear terms were calibrated using 200 simulations, after verifying that this number allows accurate convergence for all the shapes. For this test we also included the wavelet bispectrum pipeline described in Sect. 3. Although this last estimator turns out to be about 30% suboptimal and, in its current implementation, less correlated with the primordial templates than the other estimators, it does provide an additional interesting cross-check of our results by introducing another decomposition basis. We thus used it to analyse SMICA data in Sect. 7, while the other three pipelines were used on all maps.

A comparison of the measured f_{NL} map-by-map for all shapes and estimators is shown in Fig. 4. As an overall figure of merit of the level of agreement achieved by different pipelines we take as usual the standard deviation of the map-by-map f_{NL} differences, $\sigma_{\delta_{f_{NL}}}$. Table 7 shows that the final agreement between the three optimal pipelines (KSW, binned, and modal) is close to saturating the ideal bound in Eq. (84) determined by the imperfect correlation of the weights, i.e., it varies from about

once to twice $\sigma_{\delta_{f_{NL}}} \simeq 0.15\,\Delta_{f_{NL}}$ for an $r = 0.99$ correlation. This is very consistent with the level of agreement that we find between estimators for the final results from the data, providing a good indication that no spurious NG features are present in the actual data set when compared to our simulations. It should be noted that we found a similarly good level of agreement between estimators for the non-primordial shapes of point sources and ISW-lensing, although we chose not to present those results here in order to focus on the primordial shapes. Finally, regarding the wavelet pipeline, the lower weight correlation and suboptimal error bars produce an expected larger scatter when compared to the other estimators. Nonetheless, the level of agreement is still of order 1σ, which is quite acceptable for consistency checks of the optimal results. Again, this MC expectation agrees with what we see in our results on the real data.

7. Results

For our analysis of *Planck* data we considered foreground-cleaned maps obtained with the four component separation methods SMICA, NILC, SEVEM, and C-R. For each map, f_{NL} amplitudes for the local, equilateral, and orthogonal primordial shapes have been measured using three (four for SMICA) bispectrum estimators described in Sect. 3. The results can be found in Sect. 7.1. These estimators, as explained earlier, basically use an expansion of the theoretical bispectrum templates in different domains, and truncate the expansion when a high level of
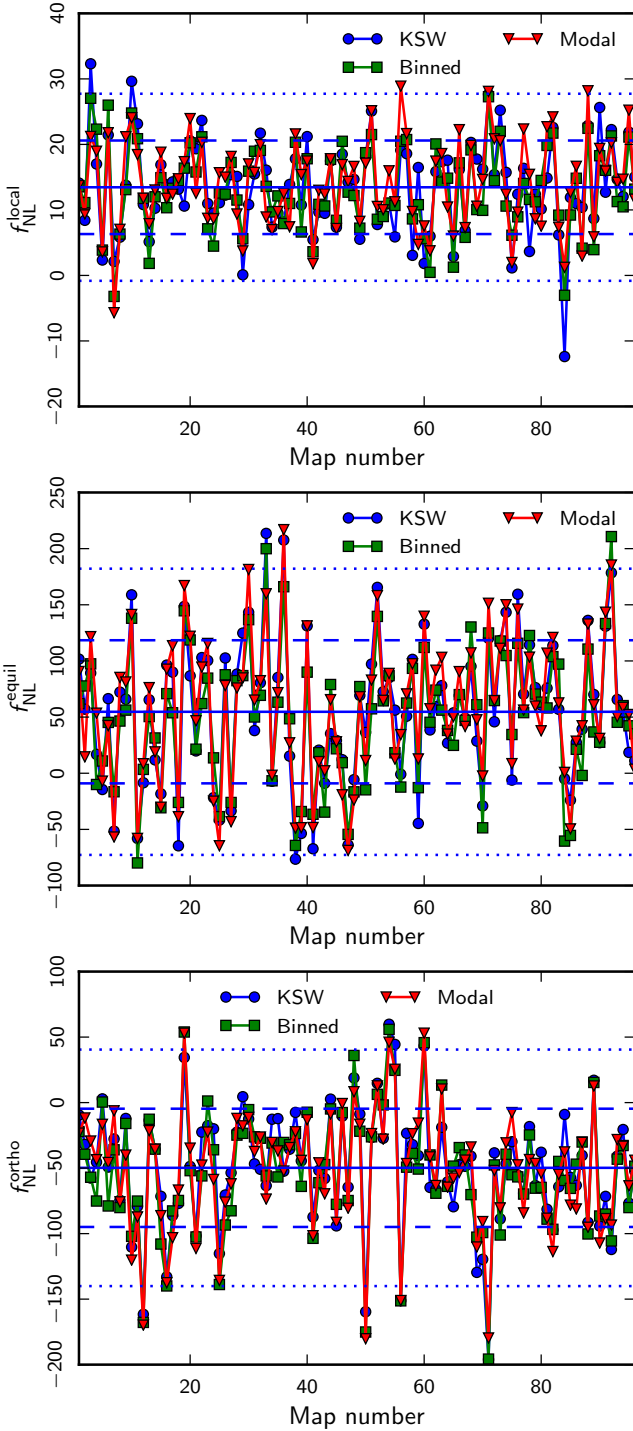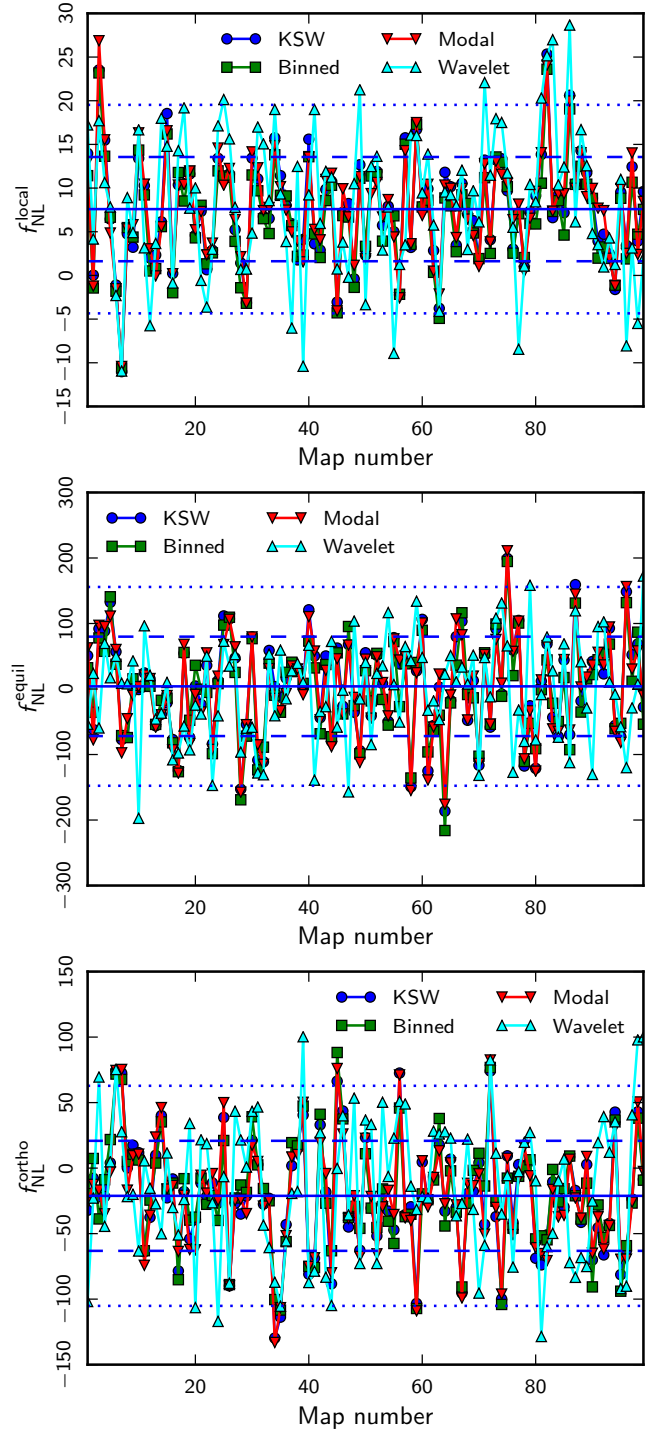
Fig. 3. Map-by-map comparison of the results from the different estimators for local (*top*), equilateral (*centre*), and orthogonal (*bottom*) f_{NL} for the set of masked non-Gaussian simulations described in Sect. 6.1.3, assuming the shapes to be independent. The horizontal solid line is the average value of all maps for KSW, and the dashed and dotted horizontal lines correspond to 1σ and 2σ deviations, respectively.

Fig. 4. Map-by-map comparison of the results from the different estimators for local (*top*), equilateral (*centre*), and orthogonal (*bottom*) f_{NL} for 99 maps from a set of realistic lensed simulations passed through the SMICA pipeline, described in Sect. 6.2, assuming the shapes to be independent. The horizontal solid line is the average value of the maps for KSW, and the dashed and dotted horizontal lines correspond to 1σ and 2σ deviations, respectively.

correlation with the primordial templates is achieved. These accurate decompositions, which are highly correlated with each other, are then matched to the data in order to extract f_{NL}. The different expansions are all different implementations of the maximum-likelihood estimator given in Eq. (32). So the final estimates are all expected to be optimal, and measure f_{NL} from nearly identical fitting templates. As discussed and tested

in detail on simulations in Sect. 6, central f_{NL} values from different methods are expected to be consistent with each other within about $0.3\sigma_{f_{NL}}$. It is then clear that comparing outputs from both different estimators and different component separation methods, as we do, allows for stringent internal consistency checks and improved robustness of the final f_{NL} results.

Table 7. Results from the different independent estimators for $f_{\rm NL}$ for 99 maps from a set of realistic lensed simulations passed through the SMICA pipeline, described in Sect. 6.2.

				$f_{\rm NL}$			
Parameter	KSW	Binned	Modal	Wavelet	Binned − KSW	Modal − KSW	Wavelet − KSW
Local	7.6 ± 6.0	6.8 ± 5.8	7.7 ± 5.9	8.1 ± 8.4	−0.8 ± 1.2	0.1 ± 1.4	0.5 ± 6.4
Equilateral . .	4 ± 76	−1 ± 72	2 ± 76	−3 ± 76	−5 ± 20	−2 ± 13	−7 ± 91
Orthogonal . .	−21 ± 42	−20 ± 41	−21 ± 42	−15 ± 53	1.6 ± 11	−0.1 ± 8	6.4 ± 48

Notes. Both the results for the estimators individually and for the differences with KSW are given.

In addition, the binned and modal techniques produce shape-independent full bispectrum reconstructions in their own different domains. These reconstructions, discussed in Sect. 7.2, complement the standard $f_{\rm NL}$ measurements in an important way, since they allow detection of possible NG features in the three-point function of the data that do not correlate significantly with the standard primordial shapes. This advantage is shared by the skew-C_ℓ method, also applied to the data. A detection of such features would either produce a warning that some residual spurious NG effects are still present in the data or provide an interesting hint of "non-standard" primordial NG that is not captured by the local, equilateral and orthogonal shapes. Additional constraints for a broad range of specific models are provided in Sect. 7.3 (see also Sect. 2.3). In Sect. 7.4 we present the constraints on local NG obtained with Minkowski Functionals. Finally, in Sect. 7.5 we present our CMB trispectrum results.

7.1. Constraints on local, equilateral and orthogonal $f_{\rm NL}$

Our goal here is to investigate the standard separable local, equilateral and orthogonal templates used e.g., in previous WMAP analyses (see e.g., Bennett et al. 2013). When using the modal, binned, or wavelet estimator, these theoretical templates are expanded approximately (albeit very accurately) using the relevant basis functions or bins. On the other hand, the KSW estimator by construction works with the *exact* templates and, for this reason, it is chosen as the baseline to provide the final $f_{\rm NL}$ results for the standard shapes (local, equilateral, orthogonal), see Table 8. However, both the binned and modal estimators achieve optimal performance and an extremely high correlation for the standard templates (~99%), so they are statistically equivalent to KSW, as demonstrated in the previous section. This means that we can achieve a remarkable level of cross-validation for our *Planck* NG results. We will be able to present consistent constraints for the local, equilateral and orthogonal models for all four *Planck* foreground-cleaned maps, using three independent optimal estimators (refer to Table 9). Regarding component separation methods, we adopt the SMICA map as the default for the final KSW results given its preferred status among foreground-separation techniques in Planck Collaboration XII (2014). The other component separation maps will be used for important cross-validation of our results and to evaluate potential sensitivity to foreground residuals.

All the results presented in this section were obtained using the union mask U73, which leaves 73% of the sky unmasked. The mask is the union of the confidence masks of the four different component separation methods, where each confidence mask defines the region where the corresponding CMB cleaning is trusted (see Planck Collaboration XII 2014). As will be shown
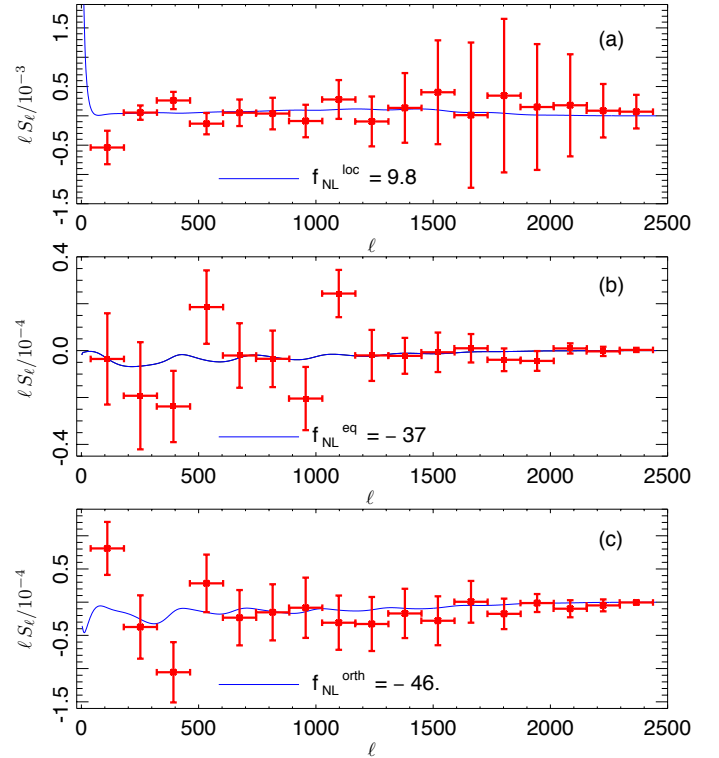

Fig. 5. Binned skew-C_ℓ statistics from the SMICA map for **a)** local; **b)** equilateral; and **c)** orthogonal. Theoretical curves are not fitted to the data shown, but are plotted with the amplitude (the only free parameter) determined from the KSW technique. There is no evidence for detection of primordial NG. Error bars are derived from the covariance of estimates from 1000 simulations. There are mild correlations between data points in all figures, but very strong correlations at high ℓ in the local case, reaching correlation coefficients $r > 0.99$ for $\ell > 1750$.

in Sect. 8.2, results are robust to changes that make the mask larger, but choosing a significantly smaller mask would leave some NG foreground contamination. For the linear term CMB and noise calibration, and error bar determination, we used sets of realistic FFP6 maps that include all steps of data processing, and have realistic noise and beam properties (see Appendix A). The simulations were also lensed using the Lenspix algorithm and filtered through the component separation pipelines.

In Table 8 we show results for the combination of the KSW estimator and the SMICA map, at a resolution of $\ell_{\rm max} = 2500$. We present both "independent" single-shape results and "ISW-lensing subtracted" ones. The former are obtained by directly fitting primordial templates to the data. For the latter, two additional operations have been performed. In the first place, as

Table 8. Results for the f_{NL} parameters of the primordial local, equilateral, and orthogonal shapes, determined by the KSW estimator from the SMICA foreground-cleaned map. Both independent single-shape results and results marginalized over the point source bispectrum and with the ISW-lensing bias subtracted are reported; error bars are 68% CL.

Shape and method	f_{NL}(KSW)	
	Independent	ISW-lensing subtracted
SMICA		
Local	9.8 ± 5.8	$\mathbf{2.7 \pm 5.8}$
Equilateral	-37 ± 75	$\mathbf{-42 \pm 75}$
Orthogonal	-46 ± 39	$\mathbf{-25 \pm 39}$

Notes. The final reported results of the paper are shown in bold face.

the name indicates, they have been corrected by subtracting the bias due to the correlation of the primordial bispectra to the late-time ISW-lensing contribution (Mangilli & Verde 2009; Junk & Komatsu 2012; Hanson et al. 2009b, see Sect. 5.2). In addition, a joint fit of the primordial shape with the (Poissonian) point-source bispectrum amplitude extracted from the data has been performed on the results marked "ISW-lensing subtracted"[11]. Since the ISW-lensing bispectrum is peaked on squeezed configurations, its impact is well known to be largest for the local shape. The ISW-lensing bias is also important for orthogonal measurements (there is a correlation coefficient $r \sim -0.5$ between the local and orthogonal CMB templates), while it is very small in the equilateral limit. The values of the ISW-lensing bias we subtract, summarized in Table 1, are calculated assuming the *Planck* best-fit cosmological model as our fiducial model. The same fiducial parameters were of course consistently used to compute the theoretical bispectrum templates and the estimator normalization. Regarding the point source contamination, we detect a Poissonian bispectrum at high significance in the SMICA map, see Sect. 5.3. However, marginalizing over point sources has negligible impact on the final primordial f_{NL} results, because the Poisson bispectrum template has very small correlations with all the other shapes.

In light of the discussion at the beginning of this section, we take the numbers from the KSW SMICA analysis in Table 8 as *the* final local, equilateral and orthogonal f_{NL} constraints for the current *Planck* data release. These results clearly show that no evidence of NG of the local, equilateral or orthogonal type is found in the data. After ISW-lensing subtraction, all f_{NL} for the three primordial shapes are consistent with 0 at 68% CL. Note that these numbers have been cross-checked using two completely independent KSW pipelines, one of which is an extension to *Planck* resolution of the pipeline used for the WMAP analysis (Bennett et al. 2013).

Unlike other methods, the KSW technique is not designed to provide a reconstruction of the full bispectrum of the data. However, the related skew-C_ℓ statistic described in Sect. 3.2.2 allows, for each given shape, visualization and study of the contribution to the measured f_{NL} from separate ℓ-bins. This is a useful tool to study potential spurious NG contamination in the data. We show for the SMICA map in Fig. 5 the measured skew-C_ℓ

spectrum for optimal detection of primordial local, equilateral and orthogonal NG, along with the best-fitting estimates of f_{NL} from the KSW method. Contrary to the case of the point source and ISW-lensing foregrounds (see Sect. 5), the skew-C_ℓ statistics do not show convincing evidence for detection of the primordial shapes, suggesting that these primordial effects are not significant sources of the NG in the map. Again, point sources contribute very little to this statistic; ISW-lensing contributes, but only a small fraction of the amplitude.

As mentioned before, our analysis went beyond the simple application of the KSW estimator to the SMICA map. All f_{NL} pipelines developed for *Planck* analysis were applied to all component-separated maps by SMICA, NILC, SEVEM, and C-R. We found from simulations in the previous sections that the KSW, binned, and modal pipelines saturate the Cramér-Rao bound, while the wavelet estimator in its current implementation provides slightly suboptimal results. Wavelets remain however a useful cross-check of the other methods, also given some technical complementarities, e.g., they are the only approach that does not require inpainting, as explained in Sect. 3. Hence we include wavelet results, but only for SMICA. The f_{NL} results for the optimal KSW, binned and modal bispectrum estimators, for the four component separation methods, are summarized in Table 9, one of the main products of our analysis of *Planck* data. The wavelet bispectrum analysis of SMICA is reported in Table 10. In the analysis, the KSW and binned bispectrum estimators considered multipoles up to $\ell_{max} = 2500$, while the modal estimator went to $\ell_{max} = 2000$. As shown in Sect. 8.1 and Table 16, error bars and central values for the three primordial shapes have converged at $\ell_{max} = 2000$, so the final primordial f_{NL} estimates from the three pipelines are directly comparable[12].

The binned bispectrum estimator used 51 bins, which were determined by optimizing the expected variance of the different f_{NL} parameters, focusing in particular on the primordial shapes.[13] The modal estimator employed a polynomial basis ($n_{max} = 600$) previously described in Fergusson et al. (2010a), but augmented with a local shape mode (approximating the SW large-angle local solution) to improve convergence in the squeezed limit. The above choices for the binned and modal methods produce a very high correlation (generally 99% or better) of the expanded/binned templates with the exact ones used by the KSW estimator. The wavelet estimator is based on third-order statistics generated by the different possible combinations of the wavelet coefficient maps of the SMHW evaluated at certain angular scales. See for example Antoine & Vandergheynst (1998) and Martínez-González et al. (2002) for detailed information about this wavelet. We considered a set of 15 scales logarithmically spaced between 1.3 and 956.3 arcmin and we also included the unconvolved map. The wavelet map $w(R_i; \boldsymbol{b})$ (Eq. (60)) for each angular scale R_i has an associated extended mask generated from the mask U73 following the procedure described and extensively used in Cutro et al. (2009b,a, 2011a,b, 2012), Donzelli et al. (2012); Regan et al. (2013). The wavelet coefficient maps are later combined into the third-order moments

---

[11] More precisely, in the subtracted ISW-lensing results the equilateral and orthogonal primordial shapes are also fitted jointly, although this has a nearly negligible impact on the final result because the two shapes are by construction nearly perfectly uncorrelated.

[12] The lower ℓ_{max} for the modal pipeline is also a conservative choice in view of the large survey of "non-standard" models that will be presented in Sect. 7.3.

[13] The boundary values of the bins are: 2, 4, 10, 18, 27, 39, 55, 75, 99, 130, 170, 224, 264, 321, 335, 390, 420, 450, 518, 560, 615, 644, 670, 700, 742, 800, 850, 909, 950, 979, 1005, 1050, 1110, 1150, 1200, 1230, 1260, 1303, 1346, 1400, 1460, 1510, 1550, 1610, 1665, 1725, 1795, 1871, 1955, 2091, 2240, and 2500 (i.e., the first bin is [2, 3], the second [4, 9], etc., while the last one is [2240, 2500]).

Table 9. Results for the $f_{\rm NL}$ parameters of the primordial local, equilateral, and orthogonal shapes, determined by the KSW, binned and modal estimators from the `SMICA`, `NILC`, `SEVEM`, and `C-R` foreground-cleaned maps.

	$f_{\rm NL}$					
	Independent			ISW-lensing subtracted		
Shape	KSW	Binned	Modal	KSW	Binned	Modal
`SMICA`						
Local	9.8 ± 5.8	9.2 ± 5.9	8.3 ± 5.9	$\mathbf{2.7 \pm 5.8}$	2.2 ± 5.9	1.6 ± 6.0
Equilateral	-37 ± 75	-20 ± 73	-20 ± 77	$\mathbf{-42 \pm 75}$	-25 ± 73	-20 ± 77
Orthogonal	-46 ± 39	-39 ± 41	-36 ± 41	$\mathbf{-25 \pm 39}$	-17 ± 41	-14 ± 42
`NILC`						
Local	11.6 ± 5.8	10.5 ± 5.8	9.4 ± 5.9	4.5 ± 5.8	3.6 ± 5.8	2.7 ± 6.0
Equilateral	-41 ± 76	-31 ± 73	-20 ± 76	-48 ± 76	-38 ± 73	-20 ± 78
Orthogonal	-74 ± 40	-62 ± 41	-60 ± 40	-53 ± 40	-41 ± 41	-37 ± 43
`SEVEM`						
Local	10.5 ± 5.9	10.1 ± 6.2	9.4 ± 6.0	3.4 ± 5.9	3.2 ± 6.2	2.6 ± 6.0
Equilateral	-32 ± 76	-21 ± 73	-13 ± 77	-36 ± 76	-25 ± 73	-13 ± 78
Orthogonal	-34 ± 40	-30 ± 42	-24 ± 42	-14 ± 40	-9 ± 42	-2 ± 42
`C-R`						
Local	12.4 ± 6.0	11.3 ± 5.9	10.9 ± 5.9	6.4 ± 6.0	5.5 ± 5.9	5.1 ± 5.9
Equilateral	-60 ± 79	-52 ± 74	-33 ± 78	-62 ± 79	-55 ± 74	-32 ± 78
Orthogonal	-76 ± 42	-60 ± 42	-63 ± 42	-57 ± 42	-41 ± 42	-42 ± 42

Notes. Both independent single-shape results and results marginalized over the point source bispectrum and with the ISW-lensing bias subtracted are reported; error bars are 68% CL. Final reported results of the paper are shown in bold.

Table 10. Results for the $f_{\rm NL}$ parameters of the primordial local, equilateral, and orthogonal shapes, determined by the suboptimal wavelet estimator from the `SMICA` foreground-cleaned map.

	$f_{\rm NL}$(wavelets)	
Shape	Independent	ISW-lensing subtracted
`SMICA`		
Local	10 ± 8.5	0.9 ± 8.5
Equilateral	89 ± 84	90 ± 84
Orthogonal	-73 ± 52	-45 ± 52

Notes. Both independent single-shape results and results marginalized over the point source bispectrum and with the ISW-lensing bias subtracted are reported; error bars are 68% CL. As explained in the text, our current wavelets pipeline performs slightly worse in terms of error bars and correlation to primordial templates than the other bispectrum estimators, but it still provides a useful independent cross-check of other techniques.

q_{ijk} (Eq. (59)), for a total 816 different statistics, and these statistics are used to constrain $f_{\rm NL}$ through a χ^2 test.

The high level of agreement between results from the KSW, binned and modal $f_{\rm NL}$ estimators, and from all the component separation pipelines, is representative of the robustness of our results with respect to residual foreground contamination, and is fully consistent with our preliminary MC analysis shown in Sect. 6. The scatter with wavelets is a bit larger, but this was expected due to the suboptimality of the wavelet estimator and is also in agreement with our MC expectations from the tests. Therefore wavelets do provide another successful cross-check.

7.2. Bispectrum reconstruction

As previously explained (see Sect. 3), in addition to looking in specific bispectrum-space directions and extracting the single number $f_{\rm NL}$ for given shapes, the binned and modal pipelines have the capability to generate a smoothed (i.e., either coarse-grained in ℓ-space, or truncated at a given expansion eigenmode) reconstruction of the full bispectrum of the data. See also Planck Collaboration XXIII (2014).

7.2.1. Modal bispectrum reconstruction

The modal pipeline was applied to the *Planck* temperature maps for the foreground-separation techniques `SMICA`, `NILC`, and `SEVEM` (Fergusson et al. 2010a). For this analysis we used two alternative sets of hybrid basis functions in order to cross-check results and identify particular signals. First, we employed trigonometric functions ($n_{\rm max} = 300$) augmented with the SW local mode, together with the three separable modes contributing to the CMB ISW-lensing signal. Secondly, we employed the same polynomial basis ($n_{\rm max} = 600$) with local SW mode as was used for $f_{\rm NL}$ estimation.

The modal coefficients $\beta_n^{\rm R}$ extracted from the *Planck* `SMICA`, `NILC`, and `SEVEM` maps are shown in Fig. 9. Here we have used the hybrid Fourier modes with local and ISW-lensing modes. These amplitudes show remarkable consistency between the different maps, demonstrating that the alternative foreground separation techniques do not appear to be introducing spurious NG. Note that here the $\beta_n^{\rm R}$ coefficients are for the orthonormalized modes R_n (Eq. (63)) and they have a roughly constant variance, so anomalously large modes can be easily identified. It is evident, for example, that among the low modes there are large signals, which include the ISW-lensing signal and point source contributions.
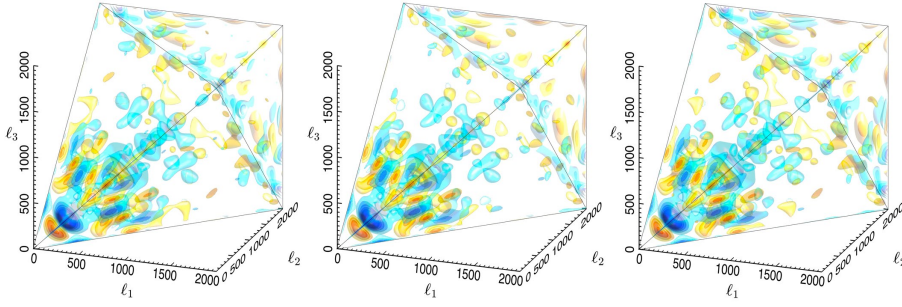
Fig. 6. Full 3D CMB bispectrum recovered from the *Planck* foreground-cleaned maps, including SMICA (*left*), NILC (*centre*) and SEVEM (*right*), using the hybrid Fourier mode coefficients illustrated in Fig. 9, These are plotted in three-dimensions with multipole coordinates $\{\ell_1, \ell_2, \ell_3\}$ on the tetrahedral domain shown in Fig. 1 out to $\ell_{\max} = 2000$. Several density contours are plotted with red positive and blue negative. The bispectra extracted from the different foreground-separated maps appear to be almost indistinguishable.
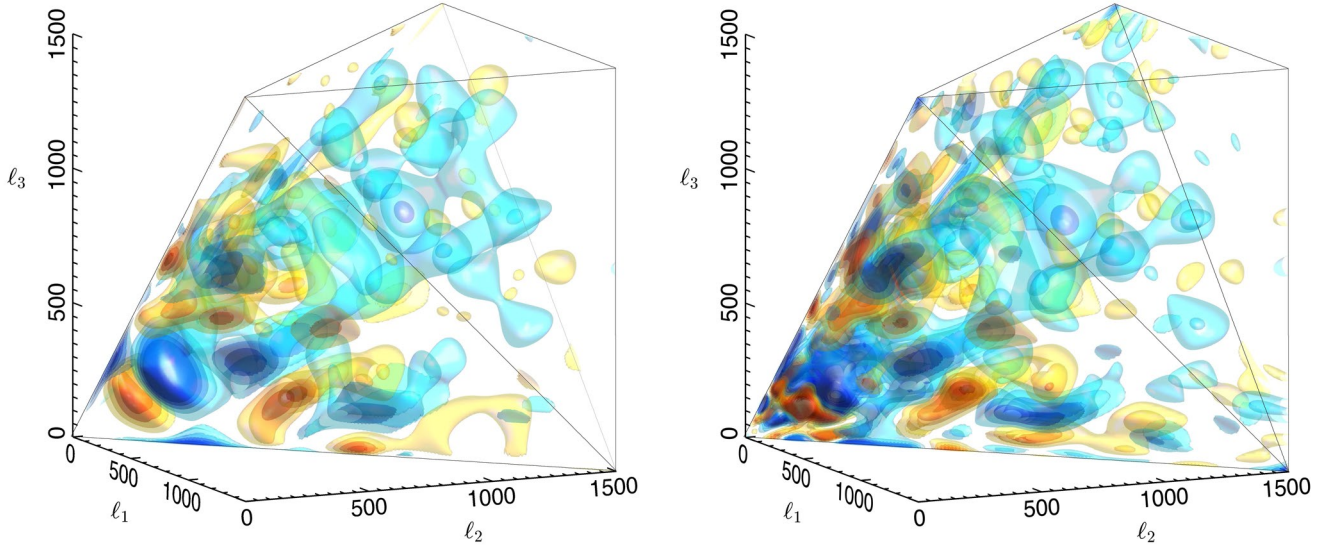
Fig. 7. *Planck* CMB bispectrum detail in the signal-dominated regime showing a comparison between full 3D reconstruction using hybrid Fourier modes (*left*) and hybrid polynomials (*right*). Note the consistency of the main bispectrum properties which include an apparently "oscillatory" central feature for low-ℓ together with a flattened signal beyond to $\ell \lesssim 1400$. Note also the periodic CMB ISW-lensing signal in the squeezed limit along the edges of the tetrapyd.

Using the modal expansion of Eq. (45) with Eq. (63), we have reconstructed the full 3D *Planck* bispectrum. This is illustrated in Fig. 6, where we show "tetrapyd" comparisons between different foreground cleaned maps. The tetrapyd (see Fig. 1) is the region defined by the multipoles that obey the triangle condition, with $\ell \le \ell_{\max}$. The 3D plots show the reduced bispectrum of the map, divided by a Sachs-Wolfe CMB bispectrum solution for a constant primordial shape, $S(k_1, k_2, k_3) = 1$. This constant primordial bispectrum template normalizaton is carried out in order to remove an $\sim\ell^4$ scaling from the starting bispectrum (it is analogous to multiplication of the power spectrum by $\ell(\ell + 1)$). To facilitate the interpretation of 3D bispectrum figures, note that squeezed configurations lie on the edges of the tetrapyd, flattened on the faces and equilateral in the interior, with $b_{\ell\ell\ell}$ on the diagonal. The colour levels are equally spaced with red denoting positive values, and blue denoting negative. Given the correspondence of the β_n^R coefficients for SMICA, NILC, and SEVEM, the reconstructed 3D signals also appear remarkably consistent, showing similar contours out to $\ell \lesssim 1500$. At large multipoles ℓ approaching $\ell_{\max} = 2000$, there is increased randomness in the reconstruction due to the rise in experimental noise and some evidence for a residual point source contribution.

There are some striking features evident in the 3D bispectrum reconstruction which appear in both Fourier and polynomial representations, as shown in more detail in Fig. 7. There is an apparent oscillation at low $\ell \lesssim 500$ already seen in WMAP-7

(Fergusson et al. 2012). Beyond out to $\ell \sim 1200$ there are further distinct features (mostly "flattened" on the walls of the tetrapyd), and an oscillating ISW-lensing contribution can be discerned in the squeezed limit. For comparison, the bispectrum reconstruction for one of the lensed Gaussian simulations described in Sect. 6.2 is shown in Fig. 8. ISW-lensing oscillating signatures in the squeezed limit are visible also in this case. When comparing results from a single simulation to the data, it is however always useful to keep in mind that the observed tetrapyd pattern is quite realization dependent, since full bispectrum reconstructions are noisy. Whatever its origin, Gaussian or otherwise, Fig. 7 reveals the CMB bispectrum of our Universe as observed by *Planck*.

The cumulative sum F_{NL}^2 over the squared orthonormal coefficients $\beta_n^{R\,2}$ from Eq. (64) for the *Planck* data is illustrated in Fig. 10 (upper panel). The *Planck* bispectrum contribution can be directly compared with Gaussian expectations averaged from 200 lensed Gaussian maps with simulated residual foregrounds. It is interesting to note that the integrated bispectrum signal fairly consistently exceeds the Gaussian mean by around 2σ over much of the domain. This includes the ISW and PS contributions for which subtraction only has a modest effect. Also shown (lower panel) is the corresponding cumulative F_{NL}^2 quantity as a function of multipole ℓ, for which features have visible counterparts at comparable ℓ in Fig. 7. Despite the high bispectrum signal, this χ^2-test over the orthonormal mode coefficients β_n^R is cumulatively consistent with Gaussianity for each of the three
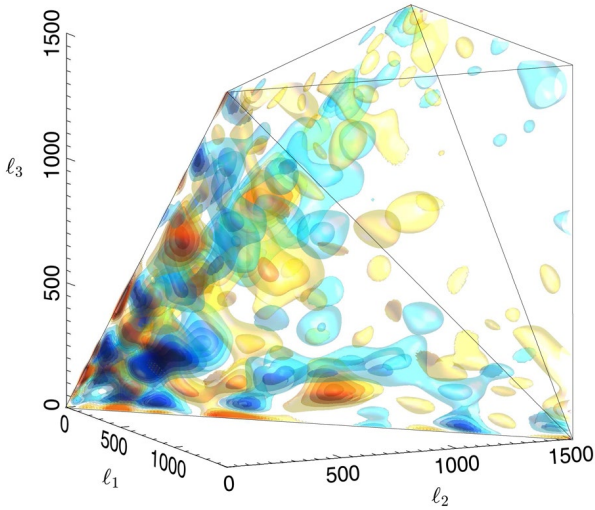
Fig. 8. CMB bispectrum from a Gaussian simulation including gravitational lensing. This reconstruction uses the same hybrid polynomials as for the *Planck* bispectrum in Fig. 7 (right), with which it can be compared. Note that an indication of a ISW-lensing bispectrum signal can be seen along the edges of the tetrahedron.
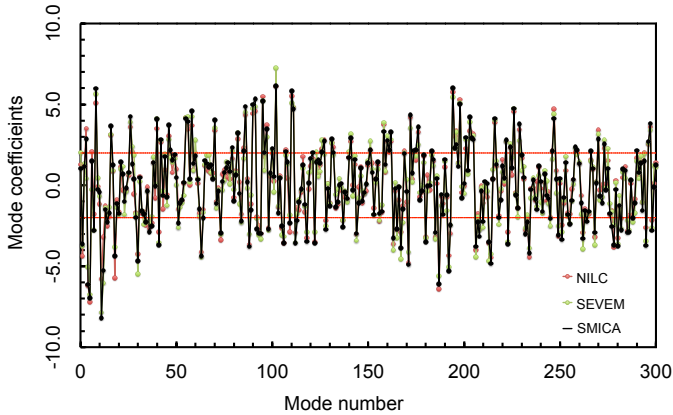

Fig. 9. Modal bispectrum coefficients β_n^R for the mode expansion (Eq. (63)) obtained from *Planck* foreground-cleaned maps using hybrid Fourier modes. The different component separation methods, SMICA, NILC and SEVEM exhibit remarkable agreement. The variance from 200 simulated noise maps was nearly constant for each of the 300 modes, with the average $\pm 1\sigma$ variation shown in red.

Fig. 10. Total integrated bispectrum F_{NL}^2 defined in Eq. (64) as a cumulative sum over orthonormal modal coefficients β_n^{R2} (*upper panel*) and over multipoles up to a given ℓ (*lower panel*). Above, the relative quantity $F_{NL}^2 \equiv \bar{F}_{NL}^2 - F_{NL}^{G\;2}$ is plotted, where $F_{NL}^{G\;2}$ is the mean obtained from 200 CMB Gaussian maps with the standard deviation shown as the red line. Below the square of the bispectrum is integrated over the tetrapyd out to ℓ and its significance plotted relative to the Gaussian standard deviation (1σ red line). A hybrid polynomial basis $n_{max} = 600$ is employed in the signal-dominated region $\ell \le 1500$.

component separation methods considered. It is however important to stress that this is a model independent integrated measurement of NG. Fitting the data with specific bispectrum templates can of course enhance the signal-to-noise ratio for given models, especially in light of the 1 to 2σ excess with respect to the Gaussian expectation, shown in the lower panel. This measurement is thus not in disagreement with detection of a residual point source (Poisson) bispectrum in the same maps (see Table 3), or with the results shown in our feature models survey of Sect. 7.3.3. We also note some differences in the high mode region between SEVEM and the other two methods, with the SEVEM results being closer to the Gaussian expectation. This is consistent with tests on simulations and with SEVEM measuring a slightly lower ISW-lensing amplitude than the other methods

(see Table 2). On the other hand the discrepancies are well within statistical bounds, as it can be seen by comparing them to the Gaussian standard deviation from simulations (red line). We thus conclude that, even accounting for these high modes deviations, the different component separation methods display a good level of internal consistency. It is also important to notice that this already good level of agreement becomes even stronger in "non-blind" f_{NL} measurements, when primordial bispectrum templates are fit to the data, and specific theoretically relevant regions of the tetrapyd are selected. This can be clearly seen from Table 9.

7.2.2. Binned bispectrum reconstruction

As explained in Sect. 3.4.2, it is interesting to study the smoothed observed bispectrum divided by its expected standard deviation, since this will indicate if there is a significant deviation from Gaussianity for certain regions of ℓ-space. This quantity is shown in Figs. 11 and 12 as a function of ℓ_1 and ℓ_2, for two different values (or rather, bins) of ℓ_3: the intermediate value [610, 654] in Fig. 11 and the high value [1330, 1374] in Fig. 12. Each figure shows the results for the SMICA, NILC, SEVEM, and C-R cleaned maps as well as for the raw 143 GHz channel map. For comparison, the result for one of the lensed Gaussian simulations described in Sect. 6.2 is also shown. The bispectra were
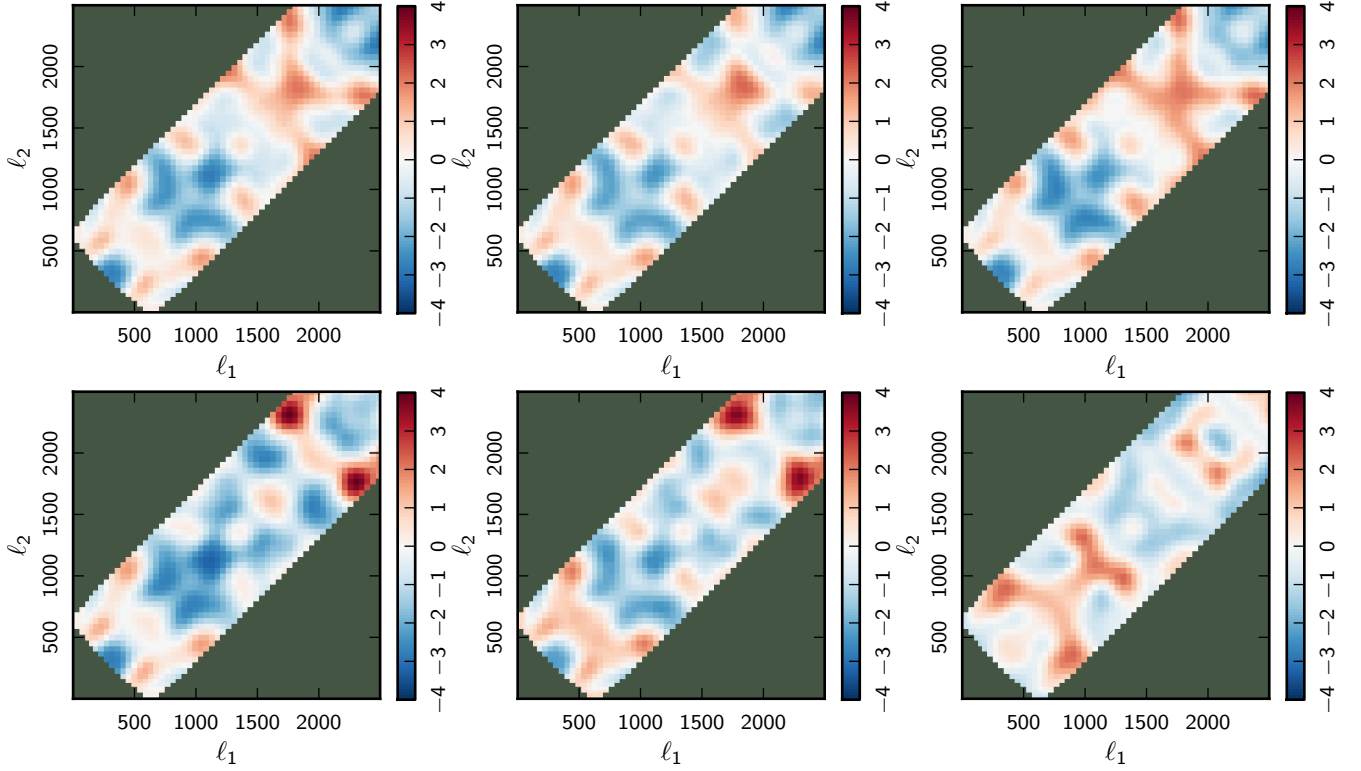
Fig. 11. Smoothed observed bispectrum as determined with the binned estimator divided by its expected standard deviation, as a function of ℓ_1 and ℓ_2, with ℓ_3 in the bin $[610, 654]$. *From left to right on the top row* are shown: SMICA, NILC, and SEVEM; and *on the bottom row*: C-R and the raw 143 GHz channel. For comparison purposes the *last figure on the bottom row* shows the same quantity for one of the lensed Gaussian simulations described in Sect. 6.2.
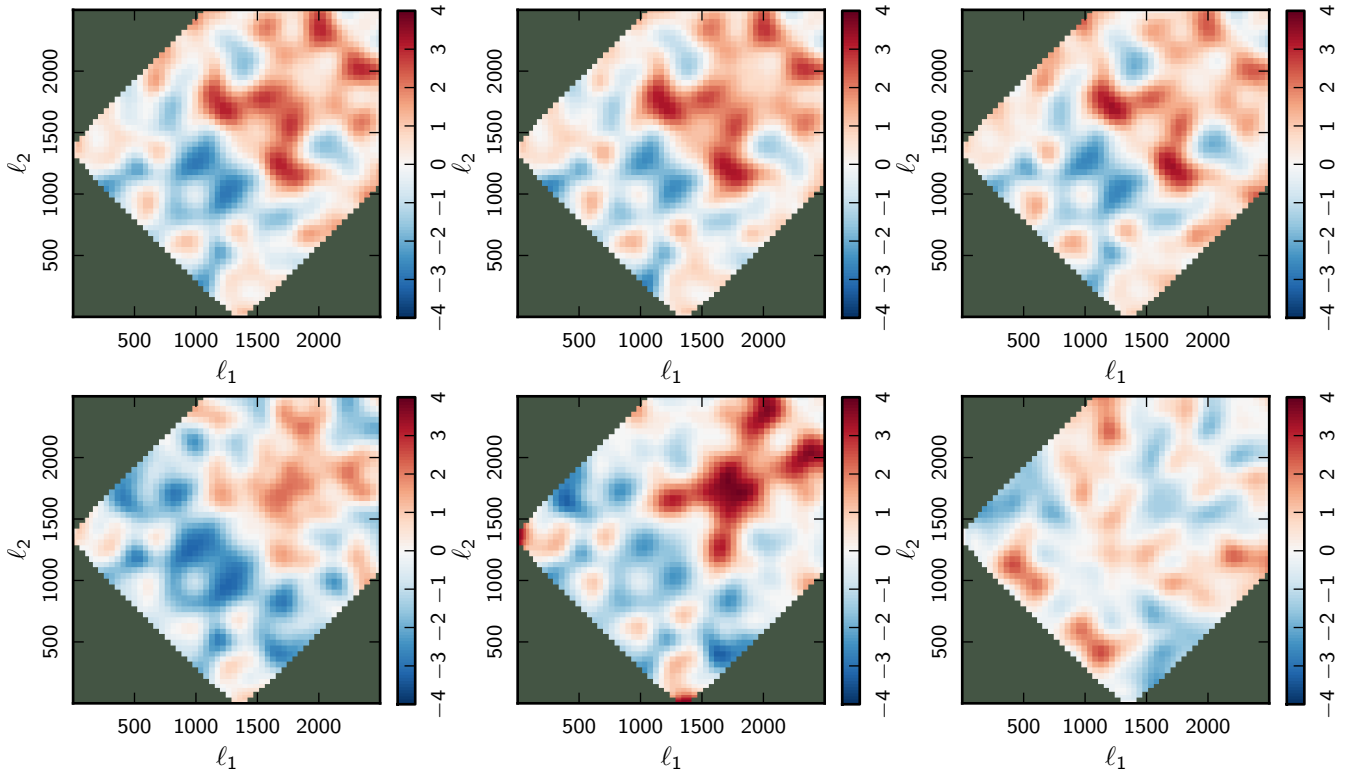

Fig. 12. Similar to Fig. 11, but with ℓ_3 in the bin $[1330, 1374]$.

obtained with the binned bispectrum estimator and smoothed
with a Gaussian kernel as explained in Sect. 3.4.2. Very blue
or red regions indicate significant NG, regions that are less
red or blue just represent expected fluctuations of a Gaussian
distribution.

From Fig. 11 at an intermediate value of ℓ_3 we can conclude
that there is a very good agreement between SMICA, NILC, and
SEVEM for all values of ℓ_1 and ℓ_2, and with C-R up to about
$\ell_1, \ell_2 \sim 1500$. In fact, up to 1500 there is also a good agree-
ment with the raw 143 GHz channel. We also see no significant
non-Gaussian features in this figure (except maybe in the C-R
and raw maps at $\ell_1, \ell_2 > 2000$). The lensed Gaussian simula-
tion in the last panel looks quite Gaussian as well (as it should),
but very different from the others. This is not surprising, as all
we are seeing here are the small random fluctuations that ex-
ist in any Gaussian realization, and the simulation represents of
course another realization than the real data. (Note that the ISW-
lensing NG, while present in the simulation, is not really visible
in the particular slices shown here due to the linear scale of the
axes and the fact that the ISW-lensing NG peaks in the squeezed
configuration.)

Figure 12 at a high value of ℓ_3, on the contrary, shows signif-
icant non-Gaussian features in the raw map, but much less NG
in the cleaned maps. In particular one can see the point-source
bispectral signal at high-ℓ approximately at equilateral config-
urations in the data. This is absent in the the lensed Gaussian
simulation, which has no point sources. There is still an excel-
lent agreement between SMICA, NILC, and SEVEM. The C-R map
shows less NG than the other three cleaned maps, which is
consistent with the absence of a detection of the Poisson point
source bispectrum for C-R, see Table 3.

7.3. Constraints on specific targeted shapes

We have deployed the modal estimator to investigate a wide
range of the inflationary models described in Sect. 2. This is
the same validated estimator for which the standard f_{NL} re-
sults have been reported in the Sect. 7, but it is augmented with
the primordial modal decomposition and projection described in
Sect. 3.2.3. The resulting modal-projected local, equilateral and
orthogonal shapes are ~99% correlated with those found using
direct integration of Eq. (37) (as for the analysis above). Modal
correlations for the other models investigated were determined
for both the primordial shapes and the late-time projected de-
compositions and were all above 90%, unless stated otherwise.
This primordial modal estimator pipeline has been applied al-
ready extensively to the WMAP-7 data (Fergusson et al. 2012).

7.3.1. Nonseparable single-field bispectrum shape results

Having characterised single-field inflation bispectra using com-
binations of the separable equilateral and orthogonal ansätze,
we note that the actual leading-order non-separable contribu-
tions (Eqs. (6), (7)) exhibit significant differences in the collinear
(flattened) limit. For this reason we provide constraints on DBI
inflation (Eq. (7)) and the two effective field theory shapes
(Eqs. (5), (6)), as well as the ghost inflation bispectrum, which
is an exemplar of higher-order derivative theories (specifically
Eq. (3.8) in Arkani-Hamed et al. 2004). Using the primor-
dial modal estimator, with the SMICA foreground-cleaned data,

we find:

$$
\begin{aligned}
f_{NL}^{DBI} &= 11 \pm 69 & (F_{NL}^{DBI-eq} &= 10 \pm 77), \\
f_{NL}^{EFT1} &= 8 \pm 73 & (F_{NL}^{EFT1-eq} &= 8 \pm 77), \\
f_{NL}^{EFT2} &= 19 \pm 57 & (F_{NL}^{EFT2-eq} &= 27 \pm 79), \\
f_{NL}^{Ghost} &= -23 \pm 88 & (F_{NL}^{Ghost-eq} &= -20 \pm 75),
\end{aligned}
\tag{86}
$$

where we have normalized with the usual primordial f_{NL} con-
vention which is shape-dependent (i.e., the central value of the
shape function is taken such that $S(k,k,k) = 1$). In parenthe-
ses we also give a reweighted F_{NL}^{equil} constraint for easier com-
parison with the equilateral constraint from the same modal
estimator, i.e., we have rescaled using the Fisher variance for
the closely-related equilateral shape. Given the strong cross-
correlation (above 95%) between all these models, the equilat-
eral family results of (86) reveal larger differences around $\sigma/3$
than might be expected (and somewhat larger than observed pre-
viously in the WMAP data – Fergusson et al. 2012). The rea-
son for this variation between the equilateral shapes in *Planck*
appears to be the additional signal observed in the flattened
limit in the bispectrum reconstruction beyond the WMAP signal-
dominated range (see Fig. 6). There is also a contribution from
the small correlation difference between equilateral models from
primordial modal and KSW methods. The results for these mod-
els for all the SMICA, NILC and SEVEM foreground-separated
maps are given in Appendix C (Table C.3).

7.3.2. Non-Bunch-Davies vacuum results

We have investigated the non-separable shapes arising from ex-
cited initial states (non-Bunch-Davies vacuum models) which
usually peak in the flattened or collinear limit. In particular, we
have searched for the four non-separable bispectra described in
Eqs. (14) and (15), as well as the original flattened shape B_Φ^{NBD}
(Eqs. (6.2)–(3)) in Chen et al. 2007b). This entails choosing suit-
able cut-offs k_c to ensure that the signal is strongly flattened
(i.e., distinct from flat in Eq. (13)), while also accurately rep-
resented by the modal expansion at both early and late times
(Eqs. (54), (55)). For B_Φ^{NBD}, we adopted the same edge truncation
and mild Gaussian filter described in Fergusson et al. (2012),
while for B_Φ^{NBD1} and B_Φ^{NBD2}, which are described by Eq. (14),
we chose $k_c = 0.001$, and in Eq. (15) we take $k_c = 0.01$. The
shape correlations for most non-Bunch-Davies vacua were good
(above 90%), except for the strongly squeezed model with os-
cillations of Eq. (14) which was relatively poor (60%). Together
with the orthogonal (Eq. (4)), flat (Eq. (13)) and vector (Eq. (19))
shapes, these non-Bunch-Davies models explore a broad range
of flattened models, with a variety of different widths for picking
out signals around the faces of the tetrapyd (see Fig. 1).

The f_{NL} results obtained for the non-Bunch-Davies models
from the different foreground-cleaned map bispectra were con-
sistent and the constraints from SMICA (for brevity) are given in
Table 11. More comprehensive results from SMICA, NILC and
SEVEM can be found in Table C.3 in Appendix C. Both B_Φ^{NBD} and
B_Φ^{NBD2} (Eq. (14)) produced raw results above 2σ, in part pick-
ing out the flattened signal observed in the bispectrum recon-
struction in Fig. 6. However, these flattened squeezed signals are
also correlated with CMB ISW-lensing and so, after subtracting
the predicted ISW bias (as well as the measured point source
signal), most NBD f_{NL} results were reduced to 1σ or less (see
"Clean f_{NL}" column in Table 11). The exception was the most
flattened model B_Φ^{NBD} which remained higher $f_{NL}^{NDB} = 178 \pm 78$,

Table 11. Constraints on flattened or collinear bispectrum models (and related models) using the `SMICA` foreground-cleaned *Planck* map.

Flattened model (Eq. number)	Raw f_{NL}	Clean f_{NL}	Δf_{NL}	σ	Clean σ
Flat model (13)	70	37	77	0.9	0.5
Non-Bunch-Davies (NBD)	178	155	78	2.2	2.0
Single-field NBD1 flattened (14)	31	19	13	2.4	1.4
Single-field NBD2 squeezed (14)	0.8	0.2	0.4	1.8	0.5
Non-canonical NBD3 (15)	13	9.6	9.7	1.3	1.0
Vector model $L = 1$ (19)	−18	−4.6	47	−0.4	−0.1
Vector model $L = 2$ (19)	2.8	−0.4	2.9	1.0	−0.1

Notes. These bispectrum shapes, with equation numbers given, are described in detail in the text.

Table 12. *Planck* bispectrum estimation results for feature models compared to the `SMICA` foreground-cleaned maps.

	$f_{NL} \pm \Delta f_{NL}$ (σ)			
Wavenumber k_c	$\phi = 0$	$\phi = \pi/4$	$\phi = \pi/2$	$\phi = 3\pi/4$
0.01000	−110 ± 159 (−0.7)	−98 ± 167 (−0.6)	−17 ± 147 (−0.1)	56 ± 142 (0.4)
0.01125	434 ± 170 (2.6)	363 ± 185 (2.0)	57 ± 183 (0.3)	−262 ± 168 (−1.6)
0.01250	−70 ± 158 (−0.4)	130 ± 166 (0.8)	261 ± 167 (1.6)	233 ± 159 (1.5)
0.01375	35 ± 162 (0.2)	291 ± 145 (2.0)	345 ± 147 (2.3)	235 ± 162 (1.5)
0.01500	−313 ± 144 (−2.2)	−270 ± 137 (−2.0)	−95 ± 145 (−0.7)	179 ± 154 (1.2)
0.01625	81 ± 126 (0.6)	177 ± 141 (1.2)	165 ± 144 (1.1)	51 ± 129 (0.4)
0.01750	−335 ± 137 (−2.4)	−104 ± 128 (−0.8)	181 ± 117 (1.5)	366 ± 126 (2.9)
0.01875	−348 ± 118 (−3.0)	−323 ± 120 (−2.7)	−126 ± 119 (−1.1)	137 ± 117 (1.2)
0.02000	−155 ± 110 (−1.4)	−298 ± 119 (−2.5)	−241 ± 113 (−2.1)	−44 ± 105 (−0.4)
0.02125	−43 ± 96 (−0.4)	−186 ± 107 (−1.7)	−229 ± 115 (−2.0)	−125 ± 104 (−1.2)
0.02250	22 ± 95 (0.2)	−115 ± 92 (−1.2)	−194 ± 105 (−1.8)	−148 ± 107 (−1.4)
0.02375	70 ± 100 (0.7)	−56 ± 94 (−0.6)	−159 ± 93 (−1.7)	−164 ± 101 (−1.6)
0.02500	106 ± 93 (1.1)	6 ± 97 (0.1)	−103 ± 98 (−1.1)	−153 ± 94 (−1.6)

Notes. This preliminary survey on a coarse grid in the range $0.01 \leq k_c \leq 0.025$ and $0 \leq \phi < \pi$ finds specific models with significance up to 99.7%.

i.e., with signals at 2.0σ, 1.8σ and 2.1σ for `SMICA`, `NILC` and `SEVEM` respectively.

We emphasise that this has to be considered just as preliminary study of flattened NG in the *Planck* data using four exemplar models. In order to reach a complete statistical assessment of constraints regarding flattened models in forthcoming analyses, we will have to undertake a systematic search for best-fit *Planck* NBD models using the parameter freedom available.

7.3.3. Scale-dependent feature and resonant model results

We have investigated whether the *Planck* bispectrum reconstructions include oscillations expected in feature or resonant models (Eqs. (16), (17)). Although poorly correlated with scale-invariant shapes, the feature and resonant models have (at least) two free parameters – the period k_c and the phase ϕ – forming a model space which must be scanned to determine if there is any significant correlation (in the absence of any physical motivation for restricting attention to specific periodicities). We have undertaken an initial survey of these models with the wavelength range defined by the native resolution of the present modal estimator (hybrid local polynomials with 600 modes), similar to the feature model search in WMAP data in Fergusson et al. (2012). For feature models of Eq. (16) we can obtain high correlations (above 95%) for the predicted CMB bispectrum if we take $k_c > 0.01$,

that is, for an effective multipole periodicity $\ell_c > 140$ feature models are accurately represented.

The results of a first survey of feature models in the *Planck* data is shown in Table 12 for $0.01 \leq k_c \leq 0.1$ and phases $\phi = 0, \pi/4, \pi/2, 3\pi/4$ (for $\phi \geq \pi$ we will identify a correlation with the opposite sign). Again, there was good consistency between the different foreground-separation methods `SMICA`, `NILC` and `SEVEM` showing that the results are robust to potential residual foreground contamination in the data. For brevity we only give `SMICA` results here, while providing measurements from other component separation methods in Appendix C. Feature signals are typically largely uncorrelated with the ISW-lensing or point sources, but nevertheless we subtract these signals and give results for the cleaned f_{NL}. The Table 12 results show that there is a parameter region around $0.01 \leq k_c \leq 0.025$ for which signals well in excess of 2σ are possible (we undertook a broader search with $0.01 \leq k_c \leq 0.1$ but found only a low signal beyond $k > 0.3$). It appears that some feature models are able to match the low-ℓ "plus-minus" and other features in the *Planck* bispectrum reconstruction (see Fig. 6). The best fit model has $k_c = 0.0185$ ($\ell_c \approx 260$) and phase $\phi = 0$ with a signal -3σ. As a further validation step of our results, we also re-analysed the models with $>2.5\sigma$ significance using a different modal decomposition, namely an oscillating Fourier basis ($n_{max} = 300$) augmented with a local SW mode (the same used for the reconstruction plots in Sect. 7). The results from this basis are shown
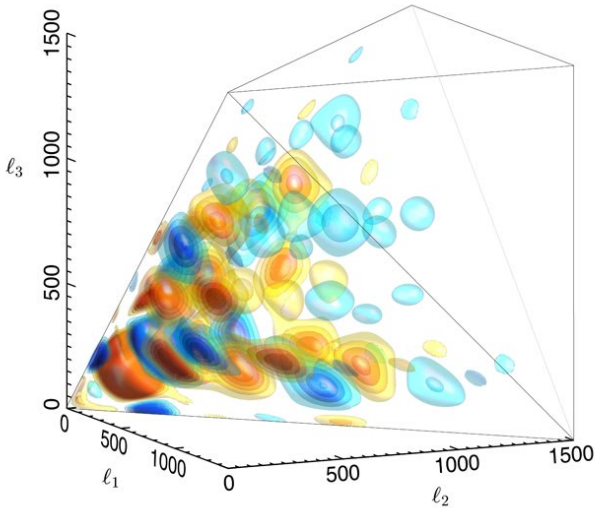
Fig. 13. CMB bispectrum shown for the best-fit feature model with an envelope with parameters $k = 0.01875$, phase $\phi = 0$ and $\Delta k = 0.045$ (see Table 13). Compare with the Planck bispectrum reconstruction, Fig. 7.

in Appendix C and they are fully consistent with the polynomial measurements presented here. The previous best-fit WMAP feature model, $k_{\rm c} = 0.014$ ($\ell_{\rm c} \approx 200$) and phase $\phi = 3\pi/4$, attained a 2.15σ signal with $\ell < 500$ (Fergusson et al. 2012), but it only remains at this level for *Planck*.

We note however that the apparently high statistical significance of these results is much lower if we consider this to be a blind survey of feature models, because we are seeking several uncorrelated models simultaneously. Following what we did for our study of impact of foregrounds in Sect. 8, we considered a set of 200 realistic lensed FFP6 simulations, processed through the `SMICA` pipeline, and including realistic foreground residuals. If we use this accurate MC sample to search for the same grid of 52 feature models as in Table 12, we find a typical maximum signal of $2.23(\pm0.56)\sigma$. Searching across all feature models (see below) studied here yields an expected maximum $2.37(\pm0.53)\sigma$ (whereas the survey for all 511 models from all paradigms investigated yielded $2.55(\pm0.52)\sigma$). This means that our best-fit model from data has a statistical significance below 1.5σ above the maximum signal expectation from simulations, so we conclude that we have *no significant detection of feature models* from *Planck* data.

Feature models typically have a damping envelope representing the decay of the oscillations as the inflaton returns to its background slow-roll evolution. Indeed, the feature envelope is a characteristic of the primordial mechanism producing the fluctuations, decaying as k increases for inflation while rising for contracting models like the ekpyrotic case (Chen 2011). We have made an initial survey to determine whether a decaying envelope improves the significance of any feature models. The envelope employed was a Gaussian centred at $k_{\rm c} = 0.045$ with a falloff $\Delta k = 0.015, 0.02, 0.025, 0.03, 0.035, 0.04, 0.045$ and results for specific parameters are given in Table 13. The best fit model remains $k = 0.01875$ ($\ell_{\rm c} = 265$) with phase $\phi = 0$ and the significance rises to 3.23σ, together with a second model $k = 0.02$ ($\ell_{\rm c} = 285$) $\phi = \pi/4$. However the caveats about blind survey statistics previously noted also do not allow a claim of any detection in this case. A plot of the best-fit feature model with a

decay envelope is shown in Fig. 13, for which the main features should be compared with those in Fig. 7. Non-Gaussian bispectrum signals from feature models typically produce counterparts in the power spectrum as will be described in Sect. 9. An improved statistical interpretation of the results presented in this section will be possible when this additional investigation will be completed.

We have also undertaken a survey of resonant models and the non-Bunch-Davies resonant models (or enfolded resonance models). With the modal estimator, we can achieve high accuracy for the predicted bispectrum for $k_{\rm c} > 0.001$ (note that this has a different logarithmic dependence to feature models and a varying effective $\ell_{\rm c}$). For the resonance model shape of Eq. (18), we have not undertaken an extensive survey, except selecting a likely range for a high signal with periodicity comparable to the feature model, that is, with $0.25 < k_{\rm c} < 0.5$ and phases $\phi = 0$, $\pi/4$, $\pi/2$, $3\pi/4$, π. However, no significant signal was found (all below 1σ), as can be verified in Table C.1 in Appendix C. For the enfolded resonance model shape of Eq. (18) , we have undertaken a preliminary search in the range $4 < k_{\rm c} < 12$ with the same phases. Again, no significant signal emerges from the *Planck* data, as shown in Table C.2 in Appendix C.

7.3.4. Directional dependence motivated by gauge fields

We have investigated whether there is significant NG from bispectrum shapes with non-trivial directional dependence (Eq. (19)), which are motivated by inflationary models with vector fields. Using the primordial modal estimator we obtained a good correlation with the $L = 1$ flattened-type model, but the squeezed $L = 2$ model produced a relatively poor correlation of only 60%, given the complexity of the dominant squeezed limit. Preliminary constraints on these models are given in the Table 11, showing no evidence of a significant signal.

7.3.5. Warm inflation

Warm inflation produces a related shape with a sign change in the squeezed limit. This also had a poor correlation, until smoothing (WarmS) was applied as described in Fergusson et al. (2012). The resulting bispectrum shows no evidence for significant correlation with *Planck* data (`SMICA`),

$$f_{\rm NL}^{\rm WarmS} = 4 \pm 33. \tag{87}$$

The full list of constraints for `SMICA`, `NILC` and `SEVEM` models can be found for warm inflation and vector models in Table C.3 in Appendix C.

7.3.6. Quasi-single-field inflation

Finally, quasi-single-field inflation has been analysed constraining the bispectrum shape of Q (Eq. (12)), that depends on two parameters, ν and $f_{\rm NL}^{\rm QSI}$. In order to constrain this model we have calculated modal coefficients for $0 \le \nu \le 1.5$ in steps of 0.01 (so 151 models in total). These were then applied to the data and the one with the greatest significance was selected. Results are shown in Fig. 26. The maximum signal occurred at $\nu = 1.5$, $f_{\rm NL}^{\rm QSI} = 4.79$ (0.31σ). To obtain error curves we performed a full likelihood using 2 billion simulations following the method described in Sefusatti et al. (2012). Such a large number of simulations was possible as they were generated from the modal β-covariance matrix which is calculated once from

Table 13. Feature model results with an envelope decay function of width Δk.

Wavenumber k_c; phase	$f_{NL} \pm \Delta f_{NL}$ (σ)			
	$\Delta k = 0.015$	$\Delta k = 0.03$	$\Delta k = 0.045$	Full
0.01125; $\phi = 0$	765 ± 275 (2.8)	703 ± 241 (2.9)	648 ± 218 (3.0)	434 ± 170 (2.6)
0.01750; $\phi = 0$	-661 ± 234 (-2.8)	-494 ± 192 (-2.6)	-425 ± 171 (-2.5)	-335 ± 137 (-2.4)
0.01750; $\phi = 3\pi/4$ ..	399 ± 207 (1.9)	438 ± 183 (2.4)	442 ± 165 (2.7)	366 ± 126 (2.9)
0.01875; $\phi = 0$	-562 ± 211 (-2.7)	-559 ± 180 (-3.1)	-515 ± 159 (-3.2)	-348 ± 118 (-3.0)
0.01875; $\phi = \pi/4$ ...	-646 ± 240 (-2.7)	-525 ± 189 (-2.8)	-468 ± 164 (-2.9)	-323 ± 120 (-2.7)
0.02000; $\phi = \pi/4$ ...	-665 ± 229 (-2.9)	-593 ± 185 (-3.2)	-500 ± 160 (-3.1)	-298 ± 119 (-2.5)

Notes. Results are only presented for feature models with better than 95% CL result on the full domain (see Table 12).

the 200 *Planck* realistic CMB simulations, rather than repeatedly from the CMB simulations themselves. The procedure is to take the 151×151 correlation matrix for the models (this is just the normalized dot product of the modal coefficients). This is then diagonalised using PCA, after which only the first 5 eigenvalues are kept as the remaining eigenvalues are $<10^{-10}$. The β-covariance matrix is projected into the same sub-basis where it is also diagonalised via PCA into 5 orthonormal modes, with the two leading modes closely correlated with local and equilateral. The procedure by which to produce a simulation is to generate five Gaussian random numbers and add the mean values obtained from the *Planck* data, rotating them to the sub-basis where we determine the ν with the greatest significance. The result is then projected back to the original space to determine the related f_{NL}. The two billion results from this MC analysis are then converted into confidence curves plotted in Fig. 26. The curve shows that there is no preferred value for ν with all values allowed at 3σ. This reflects the results obtained from data previously, where we found the least preferred value of $\nu = 0.86$ had only a marginally lower significance of 0.28σ (Sefusatti et al. 2012). Of course, these conclusions are directly related to the null results for both local and equilateral templates.

7.4. Constraints on local non-Gaussianity with Minkowski Functionals

In this subsection, we present constraints on local NG obtained with Minkowski Functionals (MFs). MFs describe the morphological properties of the CMB field and can be used as generic estimators of NG (Komatsu et al. 2003; Eriksen et al. 2004; De Troia et al. 2007; Hikage et al. 2008; Curto et al. 2008; Natoli et al. 2010; Hikage & Matsubara 2012; Modest et al. 2013). As they are sensitive to every order of NG, they can be used to constrain different bispectrum and trispectrum shapes (Hikage et al. 2006, 2008; Hikage & Matsubara 2012). They are therefore complementary to, and a useful validation of, optimal estimators. Their precise definition and analytic formulations are presented in Planck Collaboration XXIII (2014). The MF technique is also used in the companion paper Planck Collaboration XXV (2014).

We review here the properties of MFs, as a complementary tool to poly-spectrum based estimators.

First, they are defined in real space, which makes MFs robust to masking effects and no linear term is needed to take into account the anisotropy of the data model. Second, as MFs are sensitive to every non-Gaussian feature in the maps, they can be a useful probe of every potential bias in the bispectrum
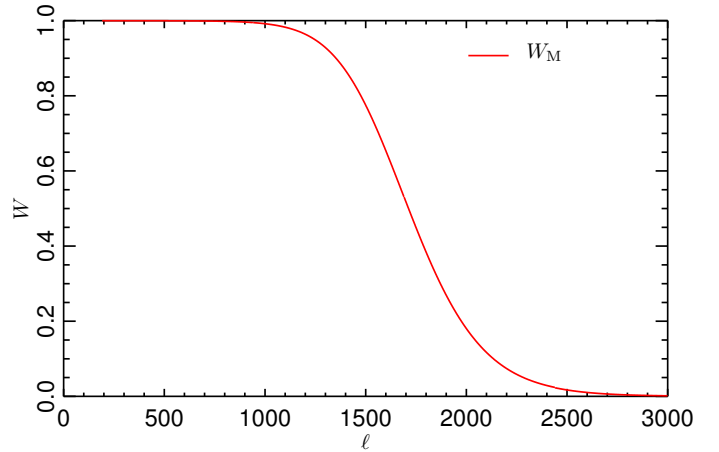

Fig. 14. The Wiener filter W_M used to constrain f_{NL}^{local} with MFs.

measurement, in particular the different astrophysical contaminations (foregrounds and secondaries).

There is a limitation to MF studies: they can be expressed in terms of weighted sums of the bispectrum (and trispectrum) in harmonic space (Matsubara 2010), hence the angle-dependence of the bispectrum is partially lost. This makes MFs suboptimal in two ways: increasing error bars for constraints on specific shapes and reducing the distinguishability of different bispectrum shapes. This lack of specificity can introduce biases, as MFs will partially confuse primordial and non-primordial sources of NG and can introduce degeneracies between different primordial shapes. Constraints on orthogonal and equilateral shapes are quite degenerate with MFs, we therefore chose here to focus on the local bispectrum shape. We also leave trispectrum analyses for future studies.

An attractive feature of MFs is their linearity for weak NG (f_{NL}) and weak signals (such as point sources, and Galactic residuals after masking and component separation) (Ducout et al. 2013). This property can be used to estimate different known non-primordial contributions.

7.4.1. Method

We constrain f_{NL}^{local} using the optimized procedure described in Ducout et al. (2013). To obtain constraints on f_{NL}^{local}, we apply a specific Wiener filter on the map (W_M), shown in Fig. 14. We do not use here the filter designed to enhance the information from the gradients of the map ($W_{D1} = \sqrt{\ell(\ell+1)}W_M$), because this filter is very sensitive to small-scale effects and may be biased by foreground residuals.

Table 14. Validation tests with MFs: results for $f_{\rm NL}^{\rm local}$ obtained using the filter $W_{\rm M}$, for $\ell_{\rm max} = 2000$ and $N_{\rm side} = 2048$.

Setup	$f_{\rm NL}^{\rm local}$(MFs)
Gaussian ideal	0.2 ± 10.7
Realistic noise, $f_{\rm NL}^{\rm local} = 12$	12 ± 13
Mask ($f_{\rm sky} = 0.73$), realistic noise $f_{\rm NL}^{\rm local} = 12$	12 ± 18

We use maps at HEALPix resolution $N_{\rm side} = 1024$ (Górski et al. 2005) and $\ell_{\rm max} = 2000$. Our results are based on the four normalized[14] functionals v_k ($k = 0, 3$) (respectively Area, Perimeter, Genus and $N_{\rm cluster}$), computed on $n_{\rm th} = 26$ thresholds ν, between $\nu_{\rm min} = -3.5$ and $\nu_{\rm max} = +3.5$ in units of the standard deviation of the map.

We combine all functionals into one vector y (of size $n = 104$). We then analyse this vector in a Bayesian way to obtain a posterior for the $f_{\rm NL}^{\rm local}$, and hence an estimate of this parameter. The principle is to compare the vector measured on the data $\hat{y}$ to the ones measured on non-Gaussian simulations with the same systematic effects (realistic noise, effective beam) and data processing (Wiener filtering) as the data, $\bar{y}(f_{\rm NL}^{\rm local})$. Modelling the MFs as multivariate Gaussians we obtain the posterior distribution for $f_{\rm NL}^{\rm local}$ with a χ^2 test:

$$P(f_{\rm NL}^{\rm local}|\hat{y}) \propto \exp\left[-\frac{\chi^2(\hat{y}, f_{\rm NL}^{\rm local})}{2}\right] \quad (88)$$

with

$$\chi^2(\hat{y}, f_{\rm NL}^{\rm local}) \equiv \left[\hat{y} - \bar{y}(f_{\rm NL}^{\rm local})\right]^{\rm T} \mathbf{C}^{-1} \left[\hat{y} - \bar{y}(f_{\rm NL}^{\rm local})\right]. \quad (89)$$

Since NG is weak, the covariance matrix C is computed with 10^4 Gaussian simulations, again reproducing effective beam, realistic noise and filtering of the data. The dependence of the MFs on $f_{\rm NL}^{\rm local}$, $\bar{y}(f_{\rm NL}^{\rm local})$, is obtained as an average of $\hat{y}$ measured on 100 simulations. The simulations used here are based on the WMAP-7 best-fit power spectrum (Komatsu et al. 2011), using the procedure described in Elsner & Wandelt (2009).

7.4.2. Validation tests

We report here validation of the MFs estimator on $f_{\rm NL}^{\rm local}$ in thoroughly realistic *Planck* simulations. This validation subsection is analogous to Sect. 6.1 concerning bispectrum-based estimators. The same tests (ideal Gaussian maps, full-sky non-Gaussian maps with noise and non-Gaussian maps with noise and mask) are performed, but different non-Gaussian simulations are used. Non-Gaussian CMB simulations as processed in Fergusson et al. (2010a) only guarantee the correctness of the 3-point correlations. Since the MFs are sensitive to higher-order n-point functions, they were validated with physical simulations (Elsner & Wandelt 2009).

The first test consists of 100 simulations of a full-sky Gaussian CMB, with a Gaussian beam with FWHM of 5 arcmin and without any noise, cut off at $\ell_{\rm max} = 2000$, with $N_{\rm side} = 2048$. Here validation tests were made at $N_{\rm side} = 2048$, but results (estimate and error bars) remain the same at $N_{\rm side} = 1024$ as we

keep the same $\ell_{\rm max}$. The second test includes non-Gaussian simulations with $f_{\rm NL}^{\rm local} = 12$ and realistic coloured and anisotropic noise, processed through the *Planck* simulation pipeline and the component-separation method SMICA. Finally, in the third test we add the union mask U73 to the previous simulations, masking both the Galaxy and the brightest point sources, and leaving 73% of the sky unmasked. Only the inpainting of the smallest holes in the point sources part of the mask was performed. For these three tests, the results are presented in Table 14. We give here the average estimate and error bar obtained on the 100 simulations, when we use the four functionals. The results show that the MF estimator is unbiased, robust, and a competitive alternative to bispectrum-based estimators. Moreover a map-by-map comparison of the results obtained on $f_{\rm NL}^{\rm local}$ with KSW and MFs estimators showed a fair agreement between the two methods.

7.4.3. Results

For our analysis we considered a foreground-cleaned map obtained with the component separation method SMICA at $N_{\rm side} = 1024$ and $\ell_{\rm max} = 2000$. As for the previous results in this section, we used the union mask U73, which leaves $f_{\rm sky} = 73\%$ of the sky after masking Galaxy and point sources. To take into account some instrumental effects (asymmetry of beams, component separation processing) and known non-Gaussian contributions (lensing), we used realistic unlensed and lensed simulations (10^3) of *Planck* data (FFP6 simulations, see Appendix A). First, MFs were applied to the unlensed simulations and the resulting curves served to calibrate the estimator, as the Gaussian part of the NG curves $\bar{y}(f_{\rm NL}^{\rm local})$[15]. This estimate is referred to as the "raw map". Secondly, MFs were applied to the lensed simulations, and the same procedure was applied, the result being referred as "lensing-subtracted". We summarize the procedure in the following equation:

$$\bar{y}(f_{\rm NL}^{\rm local}) = \bar{y}^{\rm G}_{Planck\ {\rm simulations,\ lensed}} + \Delta\bar{y}^{\rm NG}_{f_{\rm NL},\ {\rm NG\ simulations}}. \quad (90)$$

Here we assume that the MF respond linearly to lensing at first order and that primordial NG and lensing contributions are therefore additive.

Additionally, we tried to characterize other non-primordial contributions that one could expect in masked SMICA-cleaned maps covering 73% of the sky. To this end, we used simulations of extragalactic foregrounds and secondary anisotropies: uncorrelated (Poissonian) point sources; clustered CIB; and SZ clusters. These component simulations reproduce accurately the whole *Planck* data processing pipeline (beam asymmetry, component separation method). Using the linearity of MFs (Ducout et al. 2013), we could introduce these effects as a simple additive bias on the curves following

$$\hat{y} = \hat{y}^{\rm FGsubtracted} + \Delta\bar{y}^{\rm PS} + \Delta\bar{y}^{\rm CIB} + \Delta\bar{y}^{\rm SZ} \quad (91)$$

where $\Delta\bar{y}^{\rm PS,\cdots}$ is the average bias measured on 100 simulations. Note that the SZ simulation does not take into account the SZ-lensing correlation, which is expected to be negligible given the error bars.

Results are summarized in Table 15 and MFs curves are shown in Fig. 15, without including the lensing subtraction ("raw curves"). Considering the larger error bars of MF estimators, the constraints obtained are consistent with those from the

[14] Raw MFs V_k depend on the Gaussian part of fields through a normalization factor A_k that is a function only of the power spectrum shape. We therefore normalize functionals $v_k = V_k/A_k$ to focus on NG; see Planck Collaboration XXIII (2014) and references therein.

[15] The overall effect of data processing on the $f_{\rm NL}^{\rm local}$ constraint from MFs was evaluated as $f_{\rm NL}^{\rm local}$(process.) ~ 3.

Table 15. Estimates of $f_{\rm NL}^{\rm local}$ obtained with MFs on *Planck* data.

Map	$f_{\rm NL}^{\rm local}$	Source	Corresponding $\Delta f_{\rm NL}^{\rm local}$
Raw map	19.1 ± 19.3		
Lensing subtracted . .	8.5 ± 20.5	Lensing	$+10.6$
Lensing+PS subtracted	7.7 ± 20.3	Point sources	$+0.8$
Lensing+CIB subtracted	7.5 ± 20.5	CIB	$+1.0$
Lensing+SZ subtracted	6.0 ± 20.4	SZ	$+2.5$
All subtracted	4.2 ± 20.5	All	$+14.9$

Notes. Foreground and secondary effects are evaluated in terms of $f_{\rm NL}^{\rm local}$. Results are for SMICA at $N_{\rm side} = 1024$ and $\ell_{\rm max} = 2000$.
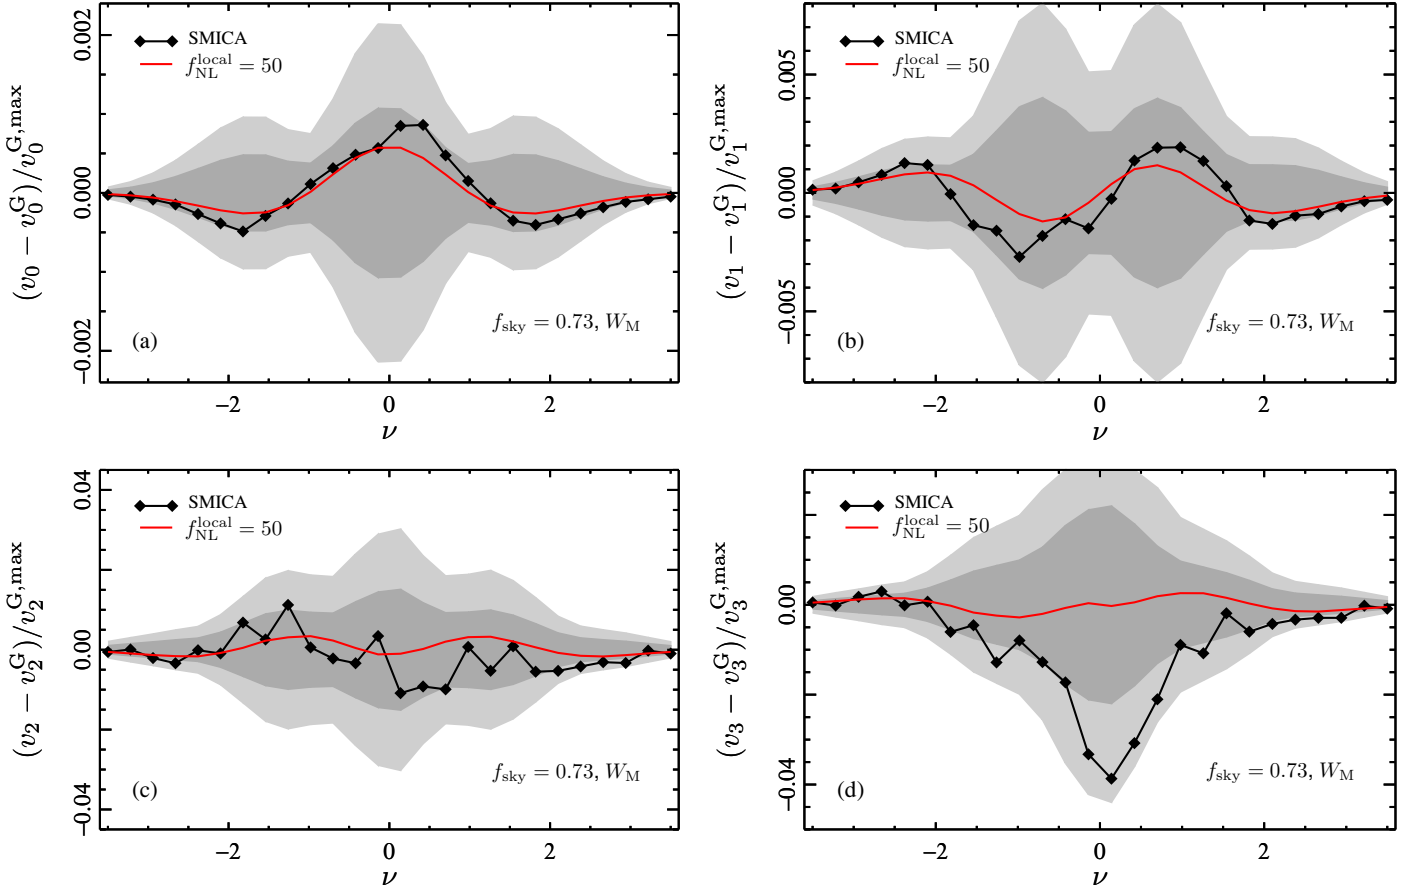
Fig. 15. MFs curves for SMICA at $N_{\rm side} = 1024$ and $\ell_{\rm max} = 2000$, for the four functionals v_k: **a)** area, **b)** perimeter, **c)** genus, and **d)** $N_{\rm cluster}$. The curves are the difference of each normalized MF, measured from the data, to the average from Gaussian *Planck* realistic simulations (not lensed). The difference curves are normalized by the maximum of the Gaussian curve. To compare the curves to the presence of primordial NG, the average difference curves for non-Gaussian simulations with $f_{\rm NL}^{\rm local} = 50$ is also represented (100 simulations).

bispectrum-based estimators, even without subtracting the expected non-primordial contributions. Moreover, results are quite robust to Galactic residuals: constraints obtained with other component separation methods (NILC and SEVEM), with different sky coverage, differ from the SMICA results presented here by less than $\Delta f_{\rm NL}^{\rm local} = 1$.

7.5. CMB trispectrum results

As shown in Fig. 16 the modulation reconstruction mean field has two large contributions, one from the mask and one from anisotropic noise, reflecting the fact that they both look like a large spatially-varying modulation of the fluctuation power. The

noise we use to estimate the mean field is taken from FFP6 simulations, adjusted with an additional $10\,\mu$K arcmin white noise component to match the power spectrum in the observed maps. However this is still only an approximate description of the instrumental noise present in the data. The mean field from non-independent maps (e.g., 143×143 GHz maps) shows a large noise anisotropy that is primarily quadrupolar before masking, and any mismatch between the simulated noise and reality would lead to a large error in the mean field subtraction. By instead using the modulation estimator for 143×217 GHz maps errors due to misestimation of the noise are avoided, and the mean field is then dominated by the shape of the Galactic cut, which is well known, and a smaller uncertainty from assumed simulation
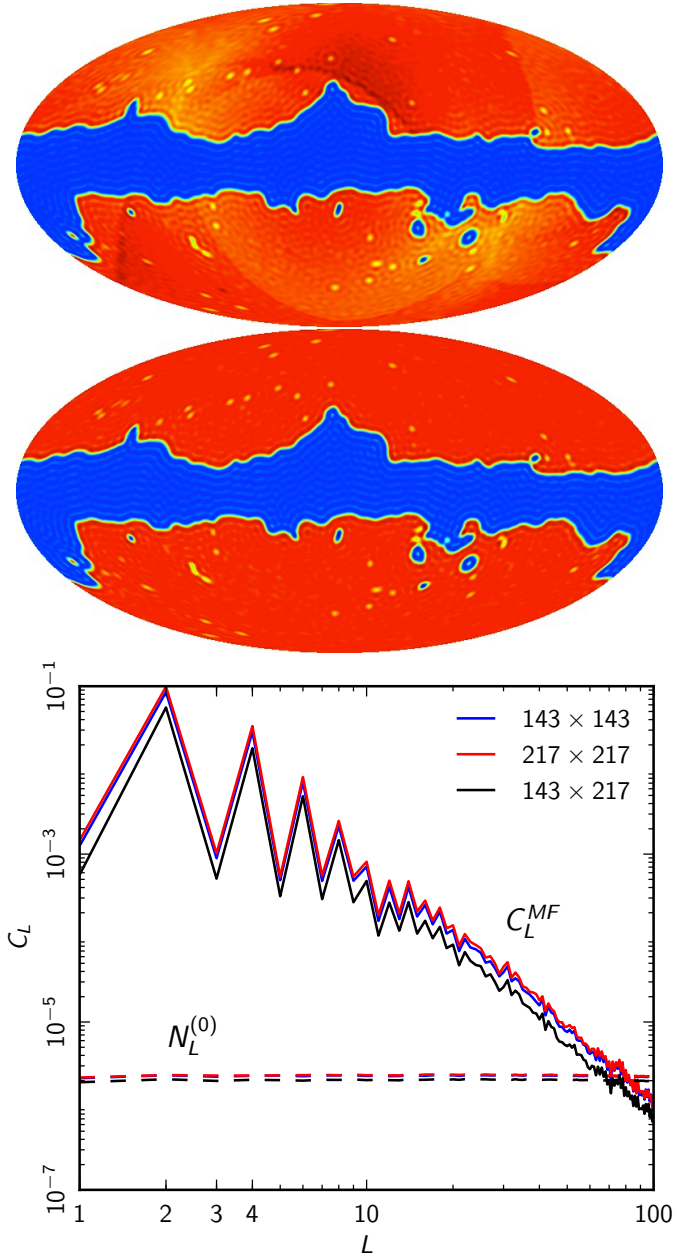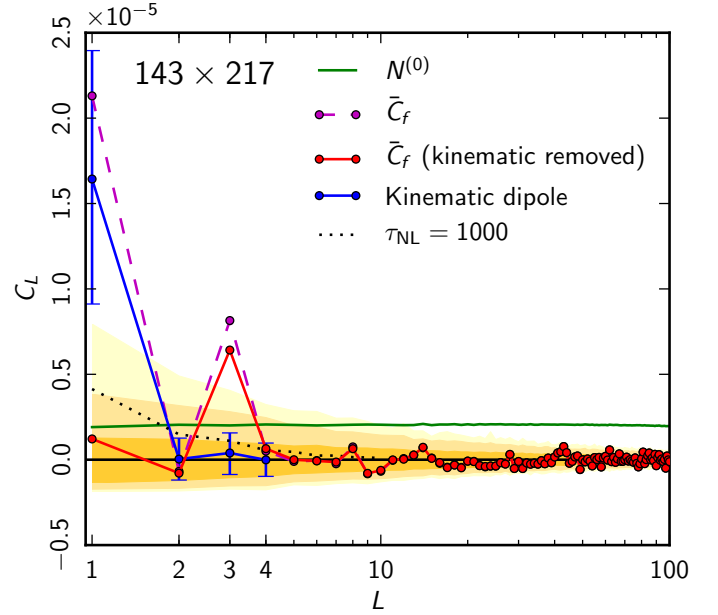
Fig. 17. Power spectrum of the power modulation reconstructed from 143×217 GHz maps. Shading shows the 68%, 95% and 99% CL intervals from simulations with no modulation or kinematic signal. The dashed lines are when the mean field simulations include no kinematic effects, showing a clear detection of a modulation dipole. The blue points show the expected kinematic modulation dipole signal from simulations, along with 1σ error bars (only first four points shown for clarity). The solid line subtracts the dipolar kinematic signal in the mean fields from simulations including the expected signal, and represents out best estimate of the non-kinematic signal (note this is not just a subtraction of the power spectra since the mean field takes out the fixed dipole anisotropy in real space before calculating the remaining modulation power). The dotted line shows the expected signal for $\tau_{\rm NL} = 1000$.

Fig. 16. The *two upper maps* show the modulation reconstruction mean field $\bar{f}^{\rm MF}(\hat{\boldsymbol{n}})$ at $L \le 100$, which is essentially a map of the expected total small-scale power on the masked map as a function of position (assuming there is no primordial power modulation). The *top mean field* map from the 143 GHz auto estimator has a large signal from both the cut (which can be calculated accurately), and from the noise anisotropy (aligned roughly with the ecliptic, which cannot be estimated very accurately from simulations). The *lower mean field* is the 143×217 GHz cross-estimator map, and does not have the contribution from the noise anisotropy (note the colour scale is adjusted). The *lower plot* shows the corresponding mean field power spectra compared to the reconstruction noise $N_L^{(0)}$ (connected part of the trispectrum); the reconstruction noise is much smaller than both the detector noise and mask contributions to the mean field. Since any $\tau_{\rm NL}$ signal is all on large scales we do not reconstruct modes above $L_{\rm max} = 100$.

power spectrum and beam errors (see Fig. 16). For this reason for our main result we work with modulation reconstructions generated from 143×217 GHz maps with independent noise, which removes the leading error due to noise mean field misestimation.

Figure 17 shows the reconstructed modulation power from 143×217 GHz maps that we use for our analysis. We show two results: one where we do not include the expected kinematic dipole signal in the mean field subtraction (see Sect. 5.4), and one were we do so that the reconstruction should then be dominated by the primordial and any unmodelled systematic effects. In the first case the 143×217 result gives a clear first detection of the dipolar kinematic modulation signal of roughly the expected magnitude (see Planck Collaboration XXVII (2014) for a more detailed discussion of this signal). Including the expected kinematic signal in the simulations (and hence the mean field) removes this signal, giving a cosmological modulation reconstruction that is broadly consistent with no modulation (statistical isotropy) except for the anomalous very significant signal in the modulation octopole.

Note that only the two-point reconstruction is free from noise bias, the four-point is still sensitive to noise modelling at the level of the subtraction of the $N_L^{(0)}$ reconstruction noise power spectrum (Eq. (72)). However as shown in the Fig. 17, $N_L^{(0)}$ is not that much larger than the reconstruction scatter at low multipoles where the $\tau_{\rm NL}$ signal peaks, so the sensitivity to noise misestimation is much less than in the mean field subtraction (where the large-scale noise anisotropy gives a large-scale mean field in the auto-estimators orders of magnitude larger, Fig. 16).

The $\tau_{\rm NL}$ estimator from the 143×217 GHz modulation reconstruction gives $\hat{\tau}_{\rm NL} = 442$, compared to a null hypothesis distribution $-452 < \hat{\tau}_{\rm NL} < 835$ at 95% CL ($\sigma_{\tau_{\rm NL}} \approx 335$). Our quoted error bar is assuming zero signal so that there is no signal cosmic

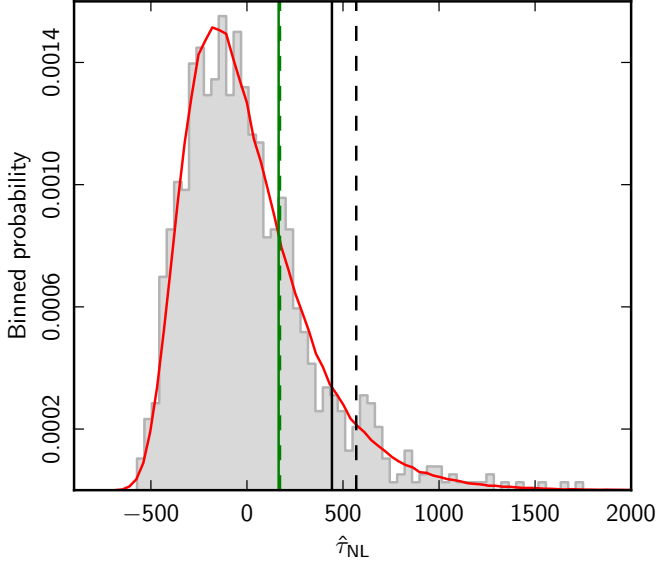**Fig. 18.** Distribution of $\hat{\tau}_{NL}$ estimators from Gaussian simulations ($L_{max} = 10$) compared to data estimates (vertical lines). The distribution is rather skewed because the main contributions are from $L \lesssim 4$ where the modulation power spectra have skewed χ^2 distributions with low degrees of freedom. The red line shows the predicted distribution for a weighted sum of $\hat{\tau}_{NL}(L)$ estimators assuming the reconstructed modulation modes are Gaussian with $2L + 1$ modes measured per L, which fits the full simulations well. The black vertical lines show the data estimates from $L_{max} = 10$, and should be compared to the green which have $L_{max} = 2$ and hence are insensitive to the anomalous octopole signal. The dashed lines are $\tau_{NL,1}$, the slightly more optimal variant of the estimator.

variance contribution, and the bulk of the apparent signal is coming from the high octopole seen in Fig. 17. The alternative estimator $\hat{\tau}_{NL,1}$ gives a slightly different weighting to the octopole, giving $\hat{\tau}_{NL,1} = 569$ with an expected null-hypothesis $\sigma_{\tau_{NL}} \approx 332$. The surprisingly large difference between the estimators can be explained as due to the large octopole signal, which has $\hat{\tau}_{NL}(L = 3) \approx 6000$. However the shape of the total signal would not be expected from a genuine τ_{NL} signal, since as shown in Fig. 17 on average the latter is expected to fall off approximately proportional to $1/L^2$ (i.e., a large primordial τ_{NL} would be expected in most realisations to give large dipole and quadrupole signals that we do not see). If we estimate τ_{NL} using $L_{max} = 2$ we obtain $\hat{\tau}_{NL} = 165$ with only a slightly larger null-hypothesis error $\sigma_{\tau_{NL}} = 349$, where in this case $\hat{\tau}_{NL,1} = 172$.

Note that the distribution of $\hat{\tau}_{NL}$ is quite skewed because the signal is dominated by the very-low multipole modulation power spectra which have skewed χ^2-like distributions due to the large cosmic variance there (see Fig. 18; Hanson & Lewis 2009; Smith & Kamionkowski 2012). The reconstruction noise acts like nearly-uncorrelated Gaussian white noise, so each of the $\bar{C}_L^f$ comes from a sum of squares of $\sim 2L + 1$ modulation reconstruction modes; the shape of the $\hat{\tau}_{NL}$ distribution is consistent with what would come from calculating a weighted sum of χ^2-distributed random variables. If we assume that any primordial modulation modes giving rise to a physical τ_{NL} signal are also Gaussian, for any given physical τ_{NL} the $\hat{\tau}_{NL}(L)$ estimators would also have χ^2 distributions. This allows us to evaluate the posterior distribution of τ_{NL} given the observed $\hat{\tau}_{NL}$, in exactly the same way that one can do for the CMB temperature power spectrum. For each L the posterior distribution $P(\tau_{NL}(L)|\hat{\tau}_{NL}(L))$ on the full sky would have an inverse-gamma distribution. We
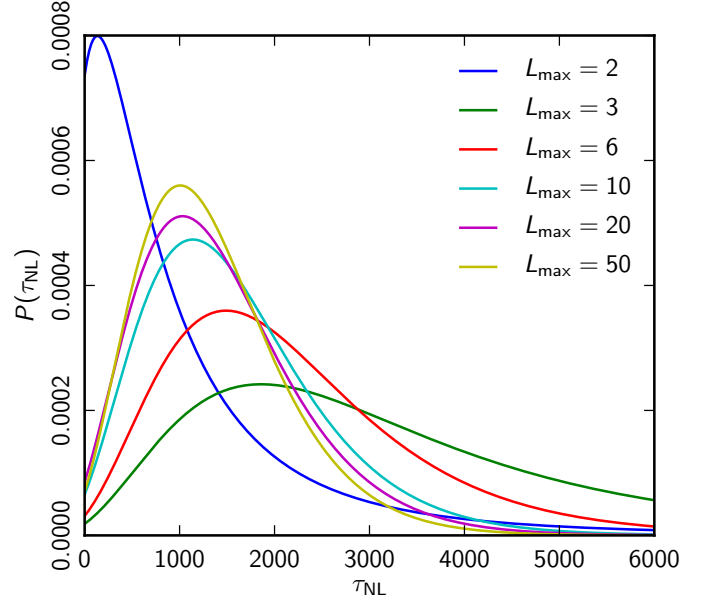

Fig. 19. Approximate posterior distributions $P(\tau_{NL}|\hat{\tau}_{NL}(L))$ for a range of L_{max}. The distributions have broad tails to high values because of the small number of large-scale modulation modes that are measured, and hence large cosmic variance. For $L_{max} \geq 3$ the distributions are pulled away from zero by the significant octopole modulation signal observed, and are gradually move back towards zero as we include more modulation modes that are inconsistent with large τ_{NL} values. As shown in Fig. 20 the octopole has significant frequency dependence and is therefore unlikely to be physical.

follow Hamimeche & Lewis (2008) by generalizing this to a cut-sky approximation for a range of multipoles:

$$-2 \ln P(\{\tau_{NL}(L)\}|\{\hat{\tau}_{NL}(L)\}) \approx$$
$$\sum_{LL'} \left[g(x(L)) N_{\tau_{NL}}^{(0)}(L) \right] \left[M^{-1} \right]_{LL'} \left[N_{\tau_{NL}}^{(0)}(L') g(x(L')) \right] \quad (92)$$

where $M_{LL'}$ is the covariance of the estimators calculated from null hypothesis simulations, $N_{\tau_{NL}}^{(0)}(L) = k_L N_L^{(0)}/C_L^{\zeta_*}$ is the estimator reconstruction noise,

$$x(L) \equiv \frac{\hat{\tau}_{NL}(L) + N_{\tau_{NL}}^{(0)}(L)}{\tau_{NL}(L) + N_{\tau_{NL}}^{(0)}(L)}, \quad (93)$$

and

$$g(x) \equiv \text{sign}(x - 1) \sqrt{2(x - \ln(x) - 1)}. \quad (94)$$

For uncorrelated χ^2-distributed estimators this distribution reduces to the exact distribution, a product of inverse-gamma distributions. This approximation to the posterior can be used to evaluate the probability of any scale dependent $\tau_{NL}(L)$, and does not rely on compression into a single $\hat{\tau}_{NL}$ estimator; it can therefore fully account for the observed distribution of modulation power between L. Here we focus on the main case of interest where $\tau_{NL}(L)$ is nearly-scale-invariant so that for all L we have $\tau_{NL}(L) = \tau_{NL}$.

The resulting posterior distributions of τ_{NL} are shown in Fig. 19 for a range of L_{max}. These are strongly skewed, in the same way as the posterior from the low quadrupole in the CMB temperature data. The high octopole is pulling the distributions up to higher τ_{NL} values, but increasing L_{max} can reduce the high-L tail because very large τ_{NL} values are inconsistent with the low modulation power seen at $L \neq 3$. With $L_{max} = 2$ the
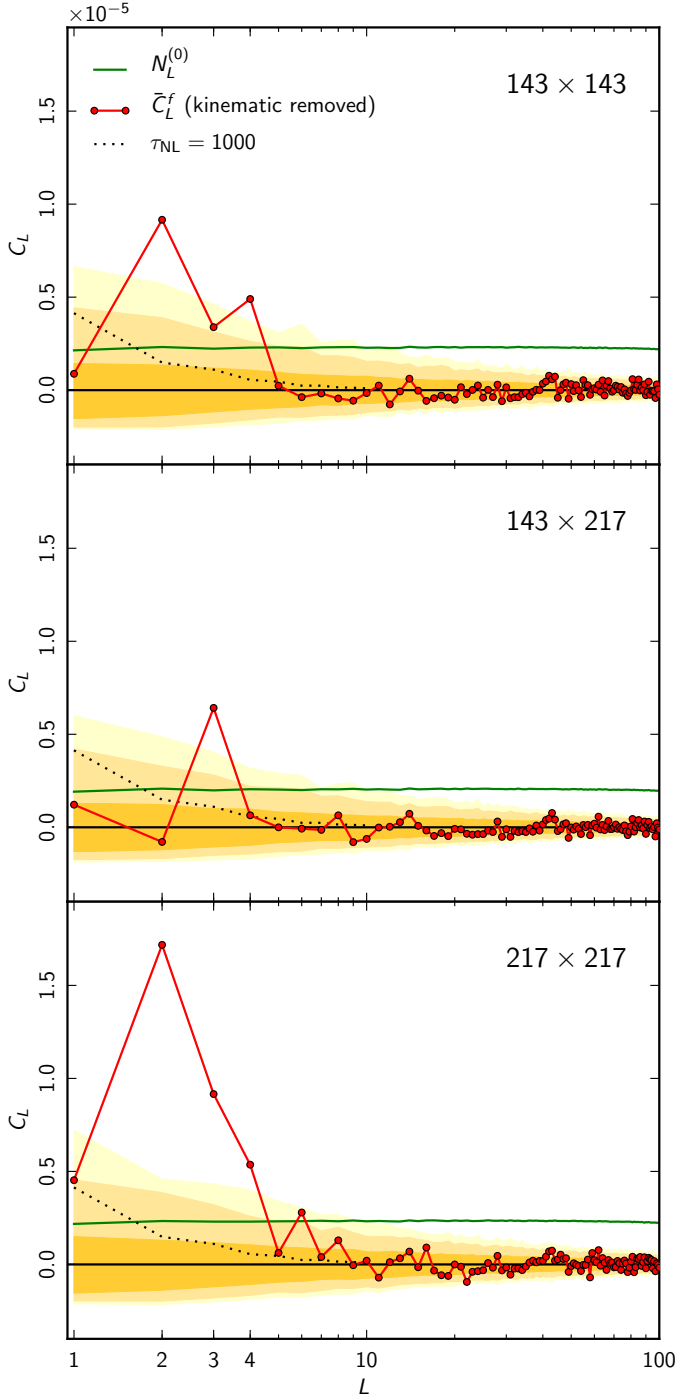
Fig. 20. Comparison of the un-normalized modulation power $\bar{C}_L^f$ with various combinations of frequencies. *The middle plot* shows the results used for our main $\tau_{\rm NL}$ result, since it removes the significant largely-quadrupolar signal from anisotropic noise misestimation seen in the two other plots. The noticeable difference in the odd octopole signal between channels indicates that the residual signal in 143×217 GHz is unlikely to be physical, but we cannot currently identify its origin.

Fig. 21. Dependence of the dipole modulation amplitude as a function of $\ell_{\rm max}$. *Upper panel*: the amplitude $|\hat{f}|$ of the dipole of the reconstructed power modulation from 143×217 GHz maps with the kinematic dipole subtracted; the dashed line shows that the observed signal decreases with $\ell_{\rm max}$ in the same way as the rms value expected from simulations. *Lower panel*: corresponding confidence values for the observed dipole of the modulation power spectrum $\bar{C}_f$ assuming it followed a chi-squared distribution; this shows the anomalous-looking results for $\ell_{\rm max} \sim 60$ and $\ell_{\rm max} \sim 600$ consistent with Planck Collaboration XXVII (2014), but that on smaller scales the observed signal is consistent with isotropy as expected from the scale-invariant $\bar{C}_f(L=1)$ constraint using $\ell_{\rm max} = 2000$ shown in Fig. 17.

posterior peaks near zero, but the distribution is then very broad because there are only about 8 modes, which therefore have large cosmic variance. For $L_{\rm max} = 50$ we find $\tau_{\rm NL} < 2800$ at 95% CL, which we take as our upper limit.

Figure 20 shows the modulation reconstructions for the 143 GHz and 217 GHz maps separately compared to the cross estimator. The picture is complicated here by the large signals

from noise misestimation seen in the 143×143 and 217×217 estimators, however the fact that the octopole in 143×143 is *lower* than in the cross-estimator indicates that the octopole signal is very unlikely to be mainly physical. Our measured $\tau_{\rm NL}$ limit in practice represents a strong upper limit on the level of primordial $\tau_{\rm NL}$ that could be present, since unmodelled varying small-scale foreground or non-constant gain/calibration would also only serve to increase the measured estimate compared to primordial on average. The octopole signal does vary slightly with Galactic mask, though at present we cannot clearly isolate its origin. If more extensive analysis (for example using cross-map estimators at the same frequency) can identify a non-physical origin and remove it, the quoted upper limit on $\tau_{\rm NL}$ would become significantly tighter. For a more extensive discussion of possible foreground and systematic effect issues that can affect 4-point estimators see Planck Collaboration XXVII (2014) and Planck Collaboration XVII (2014).

We have focussed on nearly scale invariant modulations here. As discussed in Planck Collaboration XXIII (2014) there are some potentially interesting "anomalies" in the distribution of power in the *Planck* data, especially the hemispherical power asymmetry at $\ell_{\rm max} \leq 600$. If these power asymmetries are physical rather than statistical flukes, they must be strongly scale-dependent, and would correspond to a scale-dependent $\tau_{\rm NL}$-like trispectrum. As we have shown here the dipole power asymmetry does not persist to small scales, except for the signal

aligned with the CMB dipole expected from kinematic effects: from our analysis using $\ell_{max} = 2000$ we find a primordial $\hat{\tau}_{NL}(L = 1)$ consistent with Gaussian simulations, corresponding to a dipole modulation amplitude $|f| \sim 0.2\%$. To be consistent with a $\sim7\%$ modulation at $\ell_{max} \sim 60$ and $\sim1\%$ modulation at ℓ_{max} 600, as seen in the lower-ℓ anisotropies of both WMAP and *Planck* the primordial trispectrum would have to be strongly scale-dependent on small-scale modes, so that larger small-scale modes are modulated more than the smallest ones (rather than just $\tau_{NL} = \tau_{NL}(L)$). Figure 21 shows how the allowed dipole modulation amplitude drops as ℓ_{max} increases (similar to Fig. 3 of Hanson & Lewis (2009) for WMAP but now extending to the higher ℓ_{max} available from *Planck*). This result is consistent with Planck Collaboration XXIII (2014), where different estimators and analysis cuts also find no evidence of primordial dipole-like power asymmetry on small scales, but confirm the large-scale distribution of power seen with WMAP and also a marginally-significant smaller-amplitude ($f \sim 1\%$) dipole-like modulation at $\ell_{max} \sim 600$.

8. Validation of *Planck* results

Here we perform a set of tests to check the robustness and stability of our f_{NL} measurements. As these are validation tests of *Planck* results, and not internal comparisons of bispectrum pipelines (already shown to be in tight agreement in Sects. 6 and 7) we will not employ all the bispectrum estimators on each test. In general we choose to use two estimators on each test, in order to have a cross-check of the outcomes without excessive redundancy.

8.1. Dependence on maximum multipole number

The dependence on the maximum multipole number ℓ_{max} of the SMICA results (assuming independent shapes) is shown in Fig. 22 (for the binned estimator) and Table 16 (for both the KSW and binned estimators). Testing the ℓ_{max} dependence is easiest for the binned estimator, where one can simply omit the highest bins in the final sum when computing f_{NL}. It is clear that we have reached convergence both for the values of f_{NL} and for their error bars at $\ell_{max} = 2500$, with the possible exception of the error bars of the diffuse point-source bispectrum. The diffuse point-source bispectrum template is dominated by equilateral configurations at high ℓ. Moreover, for all the shapes except point sources, results at $\ell_{max} = 2000$ are very close to those at $\ell_{max} = 2500$, taking into account the size of the error bars.

It is very interesting to see that at $\ell_{max} \sim 500$ we find a local f_{NL} result in very good agreement with the WMAP-9 value of 39.8 ± 19.9 (Bennett et al. 2013) (or 37.2 ± 19.9 after ISW-lensing bias subtraction). At these low ℓ_{max} values we also find negative values for orthogonal f_{NL}, although not as large or significant as the WMAP-9 value (which is -245 ± 100). One can clearly see the importance of the higher resolution of *Planck* both for the values of the different f_{NL} parameters and for their error bars.

It is also clear that the higher resolution of *Planck* is absolutely crucial for the ISW-lensing bispectrum; this is simply undetectable at WMAP resolution. On the other hand, the high sensitivity of *Planck* measurements also exposes us to a larger number of potentially spurious effects. For example we see that the bispectrum of point sources is also detected at high significance by *Planck* at $\ell_{max} \geq 2000$, while remaining undetectable at lower resolutions. The presence of this bispectrum in the data could in principle contaminate our primordial f_{NL} measurements. For this reason, the presence of a large point source

signal has been accounted for in previous sections by always including the Poisson bispectrum in a joint fit with primordial shapes. Fortunately, it turns out that the very low correlation between the primordial templates and the Poisson one makes the latter a negligible contaminant for f_{NL}, even when the residual point source amplitude is large.

8.2. Dependence on mask and consistency between frequency channels

To test the dependence on the mask, we have analysed the SMICA maps applying four different masks. Firstly the union mask U73 used for the final results in Sect. 7, which leaves 73% of the sky unmasked. Secondly we used the confidence mask CS-SMICA89 of the SMICA technique, which leaves 89% of the sky. Next, a bigger mask constructed by multiplying the union mask U73 with the *Planck* Galactic mask CG60, leading to a mask that leaves 56% of the sky. And finally a very large mask, leaving only 32% of the sky, which is the union of the mask CL31 – used for power spectrum estimation on the raw frequency maps – with the union mask U73 (for mask details see Planck Collaboration XII 2014 for U73, CS-SMICA89, and CG60; Planck Collaboration XV 2014 for CL31). The results of this analysis are presented in Table 17 for two different estimators: binned and modal. The f_{NL} are assumed independent here. In order to correctly interpret our results and conclusions, an important point to note is that binned results have been obtained choosing $\ell_{max} = 2500$, while modal results correspond to $\ell_{max} = 2000$. Primordial shape and ISW-lensing results and error bars saturate at $\ell_{max} = 2000$ (see Sect. 8.1), so the results from the two estimators are directly comparable in this case. The Poisson (point sources) bispectrum is however dominated by high-ℓ equilateral configurations and the signal for this specific template still changes from $\ell = 2000$ to $\ell = 2500$. The differences in central values and uncertainties between the two estimators for the Poisson shape are fully consistent with the different ℓ_{max} values. Direct comparisons on data and simulations between these two estimators and the KSW estimator showed that Poisson bispectrum results match each other very well when the same ℓ_{max} is used.

Results from the modal pipeline have uncertainties determined from MC simulations, while the results from the binned pipeline (in Table 17 and the next only) are given with Fisher error bars. It is very interesting to see that even with the large $f_{sky} = 0.32$ mask, the simple inpainting technique still allows us to saturate the (Gaussian) Cramér-Rao bound, except for the ISW-lensing shape where we have a significant detection of NG in a squeezed configuration (so that an error estimate assuming Gaussianity is not good enough). Finally we note that only for the tests in this and in the next paragraph we adopted a faster but slightly less accurate version of the modal estimator than the one used to obtain the final f_{NL} constraints in Sect. 7. In this faster implementation we use fewer modes in order to increase computational speed, and consequently we get a slight degrading of the level of correlation of our expanded templates with the initial primordial shapes. Note that the changes are small: we go from 99% correlation for local, equilateral and orthogonal shapes in the most accurate (and slower) implementation to 98% correlation for equilateral and orthogonal snapes, and 95% correlation for local shape in the faster version. This of course still allows for very stringent validation tests for all the primordial shapes, and produces results very close to the high-accuracy pipeline, while at the same time increasing overall
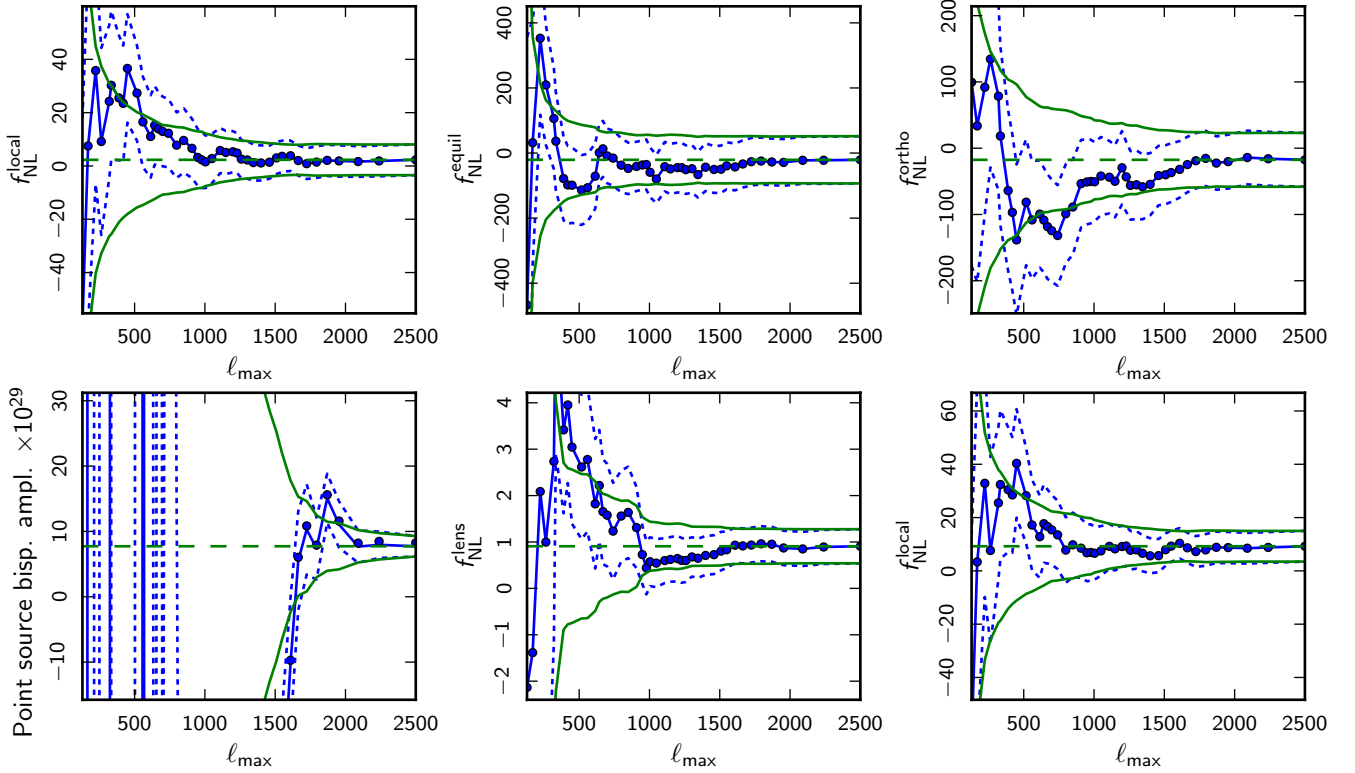
Fig. 22. Evolution of the f_{NL} parameters (solid blue line with data points) and their uncertainties (dashed lines) for the five bispectrum templates as a function of the maximum multipole number ℓ_{max} used in the analysis. *From left to right and top to bottom* the figures show respectively local, equilateral, orthogonal, diffuse point sources (all four with the ISW-lensing bias subtracted), ISW-lensing and local again (the last two without subtracting the ISW-lensing bias). To show better the evolution of the uncertainties, they are also plotted around the final value of f_{NL} (solid green lines without data points). The results are for SMICA, assume all shapes to be independent, and have been determined with the binned bispectrum estimator.

Table 16. Results for f_{NL} (assumed independent, without any correction for the ISW-lensing bias) of the SMICA cleaned map using different values of ℓ_{max}, for the KSW and binned estimators.

Shape	$\ell_{max} = 500$	$\ell_{max} = 1000$	$\ell_{max} = 1500$	$\ell_{max} = 2000$	$\ell_{max} = 2500$
			f_{NL}		
KSW					
Local	38 ± 18	6.4 ± 9.7	6.9 ± 6.2	9.1 ± 5.8	9.8 ± 5.8
Equilateral	-119 ± 121	-45 ± 88	-41 ± 75	-40 ± 75	-37 ± 75
Orthogonal	-163 ± 109	-89 ± 52	-57 ± 45	-45 ± 40	-46 ± 39
Diff.ps $/10^{-29}$ ...	$(-1.5 \pm 1.3) \times 10^4$	$(-7.9 \pm 3.1) \times 10^2$	-39 ± 18	10.0 ± 3.1	7.7 ± 1.5
ISW-lensing	3.2 ± 1.2	1.00 ± 0.43	1.00 ± 0.35	0.83 ± 0.31	0.81 ± 0.31
Binned					
Local	33 ± 18	6.6 ± 9.8	7.1 ± 6.1	8.5 ± 5.9	9.2 ± 5.9
Equilateral	-95 ± 107	-55 ± 77	-47 ± 72	-22 ± 73	-20 ± 73
Orthogonal	-102 ± 94	-69 ± 58	-60 ± 44	-35 ± 40	-39 ± 41
Diff.ps $/10^{-29}$ ...	$(-1.4 \pm 1.2) \times 10^4$	$(-8.2 \pm 2.9) \times 10^2$	-42 ± 17	9.9 ± 2.9	7.7 ± 1.6
ISW-lensing	2.6 ± 1.6	0.57 ± 0.52	0.80 ± 0.42	0.85 ± 0.38	0.91 ± 0.37

speed by almost a factor 2. Both versions of the modal pipeline were separately validated on simulations (see Sect. 6).

Besides confirming again the good level of agreement between the two estimators already discussed in Sects. 6 and 7, the main conclusion we draw from this analysis is that our measurements for all shapes are robust to changes in sky coverage, taking into account the error bars and significance levels, at least starting from a certain minimal mask. The $f_{sky} = 0.89$ mask is probably a bit too small, likely leaving foreground contamination around the edges of the mask, though even for this mask the

results are consistent within 1σ, except for point sources (which might suggest the presence of residual Galactic point source contamination for the small mask). The results from the $f_{sky} = 0.73$ and $f_{sky} = 0.56$ masks are highly consistent. This conclusion does not really change when going down to $f_{sky} = 0.32$, although uncertainties of course start increasing significantly for this large mask.

We also investigate if there is consistency between frequency channels when the largest mask with $f_{sky} = 0.32$ is used, and if these results agree with the SMICA results obtained with

Table 17. Results for f_{NL} (assumed independent) of the SMICA cleaned map using different masks as described in the main text (Sect. 8.2).

Shape	$f_{sky} = 0.89$	$f_{sky} = 0.73$	$f_{sky} = 0.56$	$f_{sky} = 0.32$
Binned				
Local	13.0 ± 5.4	9.2 ± 5.9	11 ± 6.8	6.1 ± 8.9
Equilateral	35 ± 66	−20 ± 73	−20 ± 83	39 ± 109
Orthogonal	−18 ± 36	−39 ± 39	−45 ± 45	−5 ± 59
Diff.ps /10^{-29} ...	14.0 ± 1.3	7.7 ± 1.4	9.0 ± 1.7	10.3 ± 2.2
ISW-lensing	0.69 ± 0.26	0.91 ± 0.29	0.84 ± 0.33	0.81 ± 0.43
Modal				
Local	12.1 ± 5.5	8.4 ± 6.0	12.3 ± 7.1	9.2 ± 8.7
Equilateral	52 ± 66	−56 ± 72	−31 ± 84	42 ± 104
Orthogonal	3.3 ± 35	−31 ± 40	−50 ± 47	−27 ± 59
Diff.ps /10^{-29} ...	20.6 ± 2.5	11.4 ± 2.8	10.7 ± 3.2	12.7 ± 3.9
ISW-lensing	0.42 ± 0.35	0.62 ± 0.40	1.1 ± 0.45	0.80 ± 0.48

Notes. Results are given for the binned and modal estimators. Uncertainties for the binned estimator in this table and the next are Fisher error bars. The modal estimator uses a faster and slightly less correlated expansion of the primordial templates for this test and the next than for other analyses (see Sect. 8.2 for more explanations). These caveats explain why the results shown in this table for the $f_{sky} = 0.73$ mask display small differences with respect to the corresponding numbers in the main results tables of Sect. 7, for both estimators. We also note that the binned estimator uses $\ell_{max} = 2500$ and the modal estimator $\ell_{max} = 2000$, which has an impact on the point source results as explained in the main text.

Table 18. Results for f_{NL} (assumed independent) for the raw frequency maps at 70, 100, 143, and 217 GHz with a very large mask ($f_{sky} = 0.32$) compared to the SMICA result with the union mask U73 ($f_{sky} = 0.73$), as determined by the binned (with $\ell_{max} = 2500$) and modal (with $\ell_{max} = 2000$) estimators.

Shape	SMICA	70 GHz	100 GHz	143 GHz	217 GHz
Binned					
Local	9.2 ± 5.9	19.7 ± 26.0	−2.5 ± 13.2	10.4 ± 9.8	−4.7 ± 9.6
Equilateral	−20 ± 73	159 ± 188	70 ± 132	48 ± 114	−9 ± 114
Orthogonal	−39 ± 39	−78 ± 139	−106 ± 81	−101 ± 64	−84 ± 63
Diff.ps /10^{-29} ...	7.7 ± 1.4	$(-1.4 \pm 2.3) \times 10^3$	−4.0 ± 64	8.7 ± 6.1	14.2 ± 3.0
ISW-lensing	0.91 ± 0.29	3.5 ± 2.2	0.35 ± 0.78	0.89 ± 0.50	0.87 ± 0.48
Modal					
Local	8.4 ± 6.0	36.5 ± 27.2	−6.6 ± 13.6	6.6 ± 9.4	−6.5 ± 8.9
Equilateral	−56 ± 72	74 ± 193	49 ± 123	81 ± 111	29 ± 110
Orthogonal	−31 ± 40	−225 ± 119	−75 ± 80	−133 ± 62	−112 ± 61
Diff.ps /10^{-29} ...	11.4 ± 2.8	$(-2.5 \pm 2.8) \times 10^3$	−45 ± 64	5.7 ± 7.0	25 ± 5.0
ISW-lensing	0.62 ± 0.40	2.6 ± 2.3	0.92 ± 0.80	0.78 ± 0.60	0.85 ± 0.56

Notes. The same caveats as for the previous table (Table 17) apply here as well.

the common mask. The results (assuming independent f_{NL}) are given both for the binned and the modal estimator in Table 18. As in the previous table, the full modal pipeline (faster but slightly less accurate version) has been run here, obtaining both central values from data and MC error bars from simulations, while the binned pipeline (which is slower in determining full error bars than the modal pipeline) is used to cross-check the modal measurements and has error bars given by simple Fisher matrix estimates. As one can see here and as was also checked explicitly in many other cases, the error bars from different estimators are perfectly consistent with each other and saturate the Cramér-Rao bound (except in the case of a significant non-Gaussian ISW-lensing detection).

A detailed analysis of Table 18 might actually suggest that the agreement between the two estimators employed for this test, although still clearly good, is slightly degraded when compared to their performance on clean maps from different component separation pipelines. If we compare e.g., SMICA results in Table 17 to raw data results in Table 18, we see that in the former case the discrepancy between the two estimators is at most of order $\sigma_{f_{NL}}/3$, and smaller in most cases. In the latter case, however, we notice several measurements displaying differences of order $\sigma_{f_{NL}}/2$ between the two pipelines, and the value of f_{NL}^{ortho} at 70 GHz being 1σ away. We explain these larger differences as follows. For SMICA runs we calibrated the estimator linear terms using FFP6 simulations, accurately reproducing noise properties and correlations (see Appendix A). On the other hand, for the present tests on raw frequency channels we adopted a simple noise model, based on generating uncorrelated noise multipoles with a power spectrum as extracted from the half-ring null

map, and remodulating the noise in pixel space according to the hit-count distribution. This approximation is expected to degrade the accuracy of the linear term calibration, and thus to produce a slightly lower agreement of different pipelines for shapes where the linear correction is most important. Those are the shapes that take significant contributions from squeezed triangles: local and ISW-lensing, and to a smaller but non-negligible extent orthogonal, i.e., exactly the shapes for which we find slightly larger differences.

We conclude that no significant fluctuations are observed when comparing measurements from different frequency channels (between themselves or with the clean and co-added SMICA map) or from different estimators on a given channel for the primordial shapes. The same is true for the ISW-lensing shape, although it should be noted that in particular the 70 GHz channel (like WMAP) does not have sufficient resolution to measure either the lensing or point source contributions. The uncertainties of the point source contribution vary significantly between frequency channels, although results remain consistent between channels given the error bars (when all multipoles up to $\ell_{max} = 2500$ are taken into account, as performed by the binned estimator). This is due to the fact that this shape is dominated by high-ℓ equilateral configurations, the signal-to-noise of which depends crucially on the beam and noise characteristics, which vary from channel to channel. In the SMICA map point sources are partially removed by foreground cleaning, explaining the significantly lower value than for 217 GHz. As explained before, differences between the binned and modal estimators regarding point sources are due to the different values of ℓ_{max} (2500 for binned and 2000 for modal), which particularly affects the 217 GHz channel and the SMICA cleaned map.

8.3. Null tests

To make sure there is no hidden NG in the instrumental noise, we performed a set of tests on null maps. These are noise-only maps obtained from differences between maps with the same sky signal. In the first place we constructed half-ring null maps, i.e., maps constructed by taking the difference between the first and second halves of each pointing period, divided by 2. Secondly, we constructed a survey difference map (Survey 2 minus Survey 1 divided by 2). A "survey" is defined as half a year of data, roughly the time needed to scan the full sky once; the nominal period of *Planck* data described by these papers contains two full surveys. Finally we constructed the detector set difference map ("detset 1" minus "detset 2" divided by 2). The four polarized detectors at each frequency from 100 to 353 GHz are split into two detector sets per frequency, in such a way that each set can measure all polarizations and the detectors in a set are aligned in the focal plane (see Planck Collaboration VI 2014; and Planck Collaboration XII 2014 for details on the null maps analysed in this section).

All these maps are analysed using the union mask U73 used for the final data results. However, in the case of the survey and detector set difference maps this mask needs to be increased by the unseen pixels. That effect only concerns a few additional pixels for the detector set null map, but is particularly important for the survey difference map, since a survey only approximately covers the full sky. The final f_{sky} of the mask used for the survey difference map is 64%.

The test consisted of extracting f_{NL} from the null maps described above, using only the cubic part of the bispectrum estimators (i.e., no linear term correction), and keeping the same weights as for the full "signal + noise" analysis. This means that the weights were not optimized for noise-only maps, as our aim was not to study the bispectrum of the noise per se but rather to check whether the noise alone produces a three-point function detectable by our estimators when they are run in the *same* configuration as for the actual CMB data analysis. For a similar reason it would have been pointless to introduce a linear term in this test. The purpose of the linear correction is in fact that of decreasing the error bars by accounting for off-diagonal covariance terms introduced by sky cuts and noise correlations when optimal weights are used, which is not the case here.

Our f_{NL} error bars for this test are obtained by running the estimators' cubic part on Gaussian noise simulations including realistic correlation properties. In the light of the above paragraph it is clear that such uncertainties have nothing to do with the actual uncertainties from CMB data, and cannot be compared to them.

Since SMICA was the main component-separation method for our final analysis of data, we present in Table 19 the results of our SMICA half-ring study using the KSW, binned and modal estimators, i.e., all the three main pipelines used in this paper. For the cleaned maps we do not have survey or detector set difference maps. Those are, however, available for single frequency channels. Thus we also studied all three types of null maps for the raw 143 GHz channel in Table 20, using the binned estimator. In both tables all f_{NL} shapes are assumed to be independent. The binned estimator is best suited for these specific tests as its cubic part is less sensitive to masking compared to other pipelines, especially modal. Therefore in this "cubic only" test, the binned results provide the most stringent constraints in terms of final error bars.

As one can see *Planck* passes these null tests without any problems: all values found for f_{NL} in these null maps are completely negligible compared to the final measured results on the data maps, and consistent with zero within the error bars.

8.4. Impact of foreground residuals

In Sect. 7 we applied our bispectrum estimators to *Planck* data filtered through four different component separation methods: SMICA, NILC, SEVEM and C-R (for a detailed description of component separation techniques used for *Planck* see Planck Collaboration XII 2014). The resulting set of f_{NL} measurements shows very good internal consistency both between different estimators (as expected from our MC validation tests of bispectrum pipelines described in Sect. 6) and between different foreground-cleaned maps. This already makes it clear that foreground residuals in the data are very well under control, and their impact on the final f_{NL} results is only at the level of a small fraction of the measured error bars. In this section we further investigate this issue, and validate our previous findings on data by running extensive tests in which we compare simulated data sets with and without foreground residuals from two different component separation pipelines, SMICA and NILC. The goal is to provide a MC-based assessment of the expected f_{NL} systematic error from residual foreground contamination.

For each component separation pipeline, we consider two sets of simulations. One set includes realistic *Planck* noise and beam, is masked and inpainted in the same way as we do for real data, and is processed through SMICA and NILC but it does *not* contain any foreground component. The other set is obtained by adding to the first one a number of diffuse foreground residuals: thermal and spinning dust components; free-free and synchrotron emission; kinetic and thermal SZ; CO lines

Table 19. Results for f_{NL} (assumed independent) of the SMICA half-ring null maps, determined by the KSW, binned and modal estimators.

	f_{NL}		
Shape	KSW	Binned	Modal
SMICA half-ring			
Local	-0.008 ± 0.18	-0.086 ± 0.20	$-0.13 \ \pm 0.35$
Equilateral	$-0.16 \ \pm 2.2$	$1.3 \ \ \pm 2.1$	$0.66 \ \pm 2.0$
Orthogonal	-0.035 ± 0.57	$0.51 \ \pm 0.57$	$0.14 \ \pm 0.60$
Diff.ps $/10^{-29}$ ...	$-0.05 \ \pm 0.60$	$0.03 \ \ \pm 0.68$	$0.05 \ \pm 0.65$
ISW-lensing	$(-0.06 \pm 2.0) \times 10^{-3}$	$(-2.2 \pm 4.7) \times 10^{-3}$	0.009 ± 0.030

Table 20. Results for f_{NL} (assumed independent) of several null maps determined by the binned estimator.

	f_{NL}			
Shape	SMICA half-ring	143 GHz half-ring	143 GHz survey	143 GHz detector set
Binned				
Local	-0.086 ± 0.20	-0.016 ± 0.073	0.43 ± 0.56	$1.9 \ \ \pm 1.7$
Equilateral	$1.3 \ \ \pm 2.1$	$3.2 \ \ \pm 1.8$	$-1.5 \ \pm 4.2$	$0.9 \ \pm 5.8$
Orthogonal	$0.51 \ \pm 0.57$	$1.2 \ \ \pm 0.6$	$-1.7 \ \pm 1.3$	$-1.3 \ \pm 1.8$
Diff.ps $/10^{-29}$ ...	$0.03 \ \pm 0.68$	$0.2 \ \ \pm 1.9$	$3.4 \ \ \pm 3.2$	$-1.0 \ \pm 4.3$
ISW-lensing	$(-2.2 \pm 4.7) \times 10^{-3}$	$(-0.5 \pm 1.7) \times 10^{-3}$	$(-0.6 \pm 11) \times 10^{-3}$	0.033 ± 0.026

Notes. We consider half-ring $(r1 - r2)/2$, survey $(s2 - s1)/2$, and detector set $(d1 - d2)/2$ difference maps for SMICA and the raw 143 GHz channel.

and correlated CIB. These residuals have been evaluated by applying the component separation pipelines to accurate synthetic *Planck* datasets including foreground emission according to the PSM (Delabrouille et al. 2013), and are of course dependent on the cleaning method adopted. The simple procedure of adding foreground residuals to the initially clean simulations is made possible because we consider only linear component separation methods for our analysis. Linearity is in general an important requirement for foreground cleaning algorithms aiming at producing maps suitable for NG analyses. All maps in both samples are lensed using the LensPix algorithm. We analyse both sets using different bispectrum estimators (modal, KSW, binned) for cross-validation purposes.

The presence of residual foreground components in the data can have two main effects on the measured f_{NL}. The first is to introduce a bias in the f_{NL} measurements due to the correlation between the foreground and primordial 3-point function for a given shape. The second is to increase the error bars while leaving the bispectrum estimator asymptotically unbiased. This is a consequence of accidental correlations between primordial CMB anisotropies and foreground emission. Of course these "CMB-CMB-foregrounds" bispectrum terms average to zero but they do not cancel in the bispectrum variance 6-point function. An interesting point to note is that a large foreground three-point function will tend to produce a *negative* bias in the local bispectrum measurements. That is because foreground emission produces a positive skewness of the CMB temperature distribution ("excess of photons"), and a positive skewness is in turn related to a negative $f_{\mathrm{NL}}^{\mathrm{local}}$[16]. A large negative $f_{\mathrm{NL}}^{\mathrm{local}}$ is thus a signature of

---

[16] While not rigorous, this argument captures the leading effect since Galactic foregrounds predominantly contaminate large scales. In principle, positively skewed, small scale foreground residuals ($\ell > 60$), or the negatively skewed SZ effect, can contribute positive bias. Our simulations with foreground residuals demonstrate that these contributions are subdominant.

significant foreground contamination in the map. This is indeed what we observe if we consider raw frequency maps with a small Galactic cut, which is why our frequency-by-frequency analysis in Sect. 8.2 was performed using a very large mask. For more details regarding effects of foreground contamination on primordial NG measurements in the context of the WMAP analysis see Yadav & Wandelt (2008); and Senatore et al. (2010).

In our test we built maps contaminated with foreground residuals by simply adding residual components to the clean maps. That means that the difference $(f_{\mathrm{NL}}^{\mathrm{residual}} - f_{\mathrm{NL}}^{\mathrm{clean}})_i$ for the i-th realization in the simulated sample exactly quantifies the change in f_{NL} due to foregrounds on that realization. In order to assess the level of residual foreground contamination on primordial and ISW-lensing shapes, first of all we consider the difference between average values and standard deviations of f_{NL} measured from the two map samples for each shape. As shown in Table 21 we find that neither the average nor the standard deviation shows a significant change between the two datasets. That means that foreground residuals are clearly subdominant, as they do not bias the estimator for any shape and they do not increase the variance through spurious correlations with the CMB primordial signal.

We also consider the difference $f_{\mathrm{NL}}^{\mathrm{residual}} - f_{\mathrm{NL}}^{\mathrm{clean}}$ on a map-by-map basis and compute its standard deviation. This is used as an estimate of the expected bias on a single realization due to the presence of residuals. As already expected from the negligible change in the standard deviations of the two samples, the variance of the map-by-map differences is also very small: Table 21 again shows that it is at most about $\sigma_{f_{\mathrm{NL}}}/6$ for any given shape, where $\sigma_{f_{\mathrm{NL}}}$ is the f_{NL} standard deviation for that shape. As an example, in Fig. 23 we show the measured values of $f_{\mathrm{NL}}^{\mathrm{local}}$ for the first 99 maps in both the SMICA and NILC samples, comparing results with and without residuals. It is evident also from this plot that the change in central value due to including residuals is very small. The very good agreement between the two component separation pipelines is also worth notice.

Table 21. Summary of our f_{NL} analysis of foreground residuals.

	f_{NL}					
Shape	SMICA clean	SMICA residual	NILC clean	NILC residual	SMICA residual − clean	NILC residual − clean
Modal						
Local	7.7 ± 5.9	7.8 ± 5.9	7.7 ± 5.8	7.4 ± 6.0	0.04 ± 1.0	−0.27 ± 1.1
Equilateral	−0.5 ± 77	−8.7 ± 79	−0.6 ± 78	−9.0 ± 79	−8.3 ± 8.2	−8.4 ± 8.3
Orthogonal	−23 ± 41	−25 ± 41	−24 ± 40	−26 ± 41	−2.0 ± 4.7	−2.4 ± 4.8
ISW-lensing	1.00 ± 0.38	1.01 ± 0.38	1.01 ± 0.38	1.02 ± 0.38	0.006 ± 0.052	0.013 ± 0.052

Notes. For realistic lensed FFP6 simulations processed through the SMICA and NILC component separation pipelines, we report: the average f_{NL} with and without foreground residuals added to the maps, the f_{NL} standard deviation in the same two cases, and the standard deviation of the map-by-map f_{NL} difference between the "clean" and "contaminated" sample. The impact of foreground residuals is clearly subdominant when compared to statistical error bars for all shapes. Results reported below have been obtained using the modal estimator.
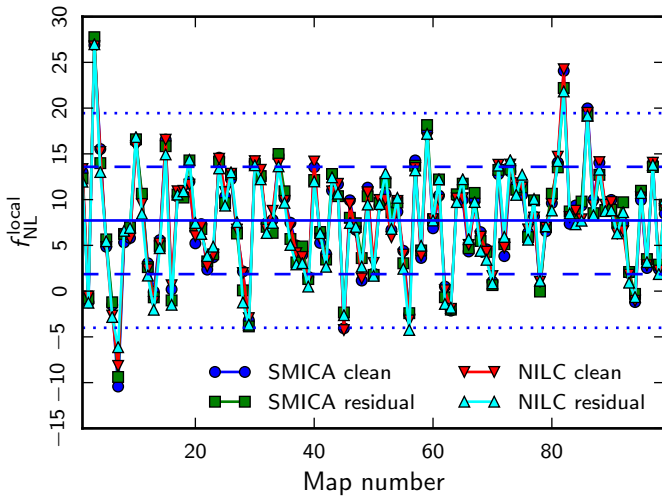
Fig. 23. The measured f_{NL}^{local} for the first 99 maps in the lensed FFP6 simulation sample used for the foreground studies presented in Sect. 8.4. We show measurements from SMICA and NILC processed maps both with and without foreground residuals. The horizontal solid line is the average value of the SMICA clean maps, and the dashed and dotted horizontal lines correspond to 1σ and 2σ deviations, respectively. The plot clearly shows the very small impact of including residuals, and the very good consistency between the two component separation pipelines.

From the study shown here and from the comparison between different component separation methods in Sect. 7, we can thus conclude that the combination of foreground-cleaned maps and $f_{sky} = 0.73$ sky coverage we adopt in this work provide a very robust choice for f_{NL} studies.

8.5. Impact of HFI time-ordered information processing on NG constraints

Our validation tests are designed to determine whether instrumental or data processing systematics could lead to a spurious detection of primordial NG. Our analysis protocol passes these tests which allows us to rule out this concern with confidence. These validation tests do not exclude the possibility that the extensive processing of HFI time-ordered information (TOI) could somehow *remove* non-Gaussian signals present on the sky and thus negatively impact *Planck*'s ability to detect them. If this were true it would weaken the bounds on primordial NG reported here. HFI TOI processing includes glitch fitting, gain drift

and correction (ADC non-linearity), 4K cooler line removal, telegraph noise step correction, and TOD filtering to correct for the bolometer time constant (Planck Collaboration VI 2014; Planck Collaboration VII 2014; Planck Collaboration X 2014).

The following facts constitute strong evidence that HFI TOI processing does not remove non-Gaussian signals:

1. *Planck*'s 2.6σ ISW-lensing bispectrum measurement is consistent with expectations from the LCDM model, and has the right skew-C_ℓ shape. Like NG of local type, this is a squeezed bispectrum template.
2. When frequency maps are combined to maximize the CIB signal, *Planck* finds a 25σ detection of the nearly squeezed CIB-lensing bispectrum, consistent with the CIB 2-point correlations inferred by *Planck* (Planck Collaboration XXX 2014).
3. *Planck* detects residual point sources in the SMICA maps at 5σ, seen as a bispectrum of equilateral shape at high ℓ consistent with the expected skew-spectrum shape (Fig. 2) on the angular scales where residual infrared sources and far infrared background are expected to be the dominant contaminants of the power spectrum according to the foreground residuals in FFP6 simulations of SMICA maps (Fig. E.3 in Planck Collaboration XII 2014).
4. The trispectrum signal imprinted by *Planck*'s motion with respect to the CMB rest frame is seen at a sensible level and in a plausible direction (Planck Collaboration XXVII 2014).
5. *Planck* detects the trispectrum signal due to lensing by large scale structure with high significance and in accordance with LCDM predictions based on the *Planck* parameter likelihood.
6. Figure 22 shows that we reproduce the WMAP-9 local NG results on those larger angular scales where the WMAP-9 data are cosmic variance limited.

We conclude that all the forms of NG that should have been seen by *Planck* have been seen, in quantitative agreement with the expected level across the entire range of angular scales probed. While it cannot be excluded with absolute certainty that some unknown aspect of HFI TOI processing could have affected some unknown form of NG, the presence of these non-Gaussian features in the resulting maps (in addition to signals such as the extracted map of y-distortion, and galactic and extragalactic foregrounds), gives us confidence in the force of the bounds on primordial NG described in this paper.

9. Implications for early Universe physics

The NG analyses performed in this paper show that *Planck* data are consistent with Gaussian primordial fluctuations. The standard models of single-field slow-roll inflation have therefore survived the most stringent tests of Gaussianity performed to date. On the other hand, the NG constraints obtained on different primordial bispectrum shapes (e.g., local, equilateral and orthogonal), after properly accounting for various contaminants, severely limit various classes of mechanisms for the generation of the primordial perturbations proposed as alternatives to the standard models of inflation.

In the following subsections, unless explicitly stated otherwise, we translate limits on NG to limits on parameters of the theories by constructing a posterior assuming the following: the sampling distribution is Gaussian (which is supported by Gaussian simulations); the likelihood is approximated by the sampling distribution but centred on the NG estimate (see Elsner & Wandelt 2009); uniform or Jeffreys' prior (where stated), over ranges which are physically meaningful, or as otherwise stated. Where two parameters are involved, the posterior is marginalized to give one-dimensional constraints on the parameter of interest.

9.1. General single-field models of inflation

DBI models: the constraints on f_{NL}^{equil} provide constraints on the sound speed with which the inflaton fluctuations can propagate during inflation. For example, for DBI models of inflation (Silverstein & Tong 2004; Alishahiha et al. 2004), where the inflaton field features a non-standard kinetic term, the predicted value of the nonlinearity parameter is $f_{NL}^{DBI} = -(35/108)(c_s^{-2} - 1)$ (Silverstein & Tong 2004; Alishahiha et al. 2004; Chen et al. 2007b). Although very similar to the equilateral shape, we have obtained constraints directly on the theoretical (nonseparable) predicted shape (Eq. 7)). The constraint $f_{NL}^{DBI} = 11 \pm 69$ at 68% CL (see Eq. (86)) implies

$$c_s^{DBI} \geq 0.07 \qquad 95\% \text{ CL}. \qquad (95)$$

The DBI class contains two possibilities based on string constructions. In ultraviolet (UV) DBI models, the inflaton field moves under a quadratic potential from the UV side of a warped background to the infrared side. It is known that this case is already at odds with observations, if theoretical internal consistency of the model and constraints on power spectra and primordial NG are taken into account (Baumann & McAllister 2007; Lidsey & Huston 2007; Bean et al. 2007; Peiris et al. 2007). Our results strongly limit the relativistic régime of these models even without applying the theoretical consistency constraints.

It is therefore interesting to look at infrared (IR) DBI models (Chen 2005b,a) where the inflaton field moves from the IR to the UV side, and the inflaton potential is $V(\phi) = V_0 - \frac{1}{2}\beta H^2 \phi^2$, with a wide range $0.1 < \beta < 10^9$ allowed in principle. In previous work (Bean et al. 2008) a 95% CL limit of $\beta < 3.7$ was obtained using WMAP. In a minimal version of IR DBI, where stringy effects are neglected and the usual field theory computation of the primordial curvature perturbation holds, one finds (Chen 2005c; Chen et al. 2007b) $c_s \simeq (\beta N/3)^{-1}$, $n_s - 1 = -4/N$, where N is the number of e-folds; further, primordial NG of the equilateral type is generated with an amplitude $f_{NL}^{DBI} = -(35/108)[(\beta^2 N^2/9) - 1]$. For this model, the range $N \geq 60$ is compatible with *Planck*'s 3σ limits on n_s (Planck Collaboration XXII 2014). Marginalizing over $60 \leq N \leq 90$, we

find

$$\beta \leq 0.7 \qquad 95\% \text{ CL}, \qquad (96)$$

dramatically restricting the allowed parameter space of this model.

Power-law k-inflation: these models (Armendariz-Picon et al. 1999; Chen et al. 2007b) predict $f_{NL}^{equil} = -170/(81\gamma)$, where $n_s - 1 = -3\gamma$, $c_s^2 \simeq \gamma/8$. Assuming a prior of $0 < \gamma < 2/3$, our constraint $f_{NL}^{equil} = -42 \pm 75$ at 68% CL (see Table 9) leads to the limit $\gamma \geq 0.05$ at 95% CL. On the other hand, *Planck*'s constraint on $n_s - 1$ yields the limit $0.01 \leq \gamma \leq 0.02$ (Planck Collaboration XXII 2014). These conflicting limits severely constrain this class of models.

Flat Models and higher derivative interactions: flat NG can characterize inflationary models which arise from independent interaction terms different from the $(\dot\pi)^3$ and $\dot\pi(\nabla\pi)^2$ discussed in Sect. 2 (see also Sect. 9.2). An example is given by models of inflation based on a Galilean symmetry (Creminelli et al. 2011a), where one of the inflaton cubic interaction terms allowed by the Galilean symmetry, $M_3[\dot\pi(\partial_i\partial_j\pi)^2/a^4 - 2H\dot\pi\ddot\pi^2 + 3H^3\dot\pi^3]$, contributes to the flat bispectrum with an amplitude $f_{NL}^{flat} = (35/256)(M_3 H)/(\epsilon M_{Pl}^2)$ (Creminelli et al. 2011a). Here, π is the relevant inflaton scalar degree of freedom, ϵ the usual slow-roll parameter and a the scale factor and H the Hubble parameter during inflation. Our constraint $f_{NL}^{flat} = 37 \pm 77$ at 68% CL (see Table 11) leads to $(M_3 H)/(\epsilon M_{Pl}^2) = 270 \pm 563$ at 68% CL. These interaction terms are similar to those arising in general inflaton field models that include extrinsic curvature terms, e.g., parameterized in the Effective Field Theory approach as $M^2\dot\pi(\partial_i\pi)^2/a^4$ (Bartolo et al. 2010a), which contribute to a flat bispectrum with an amplitude $f_{NL}^{flat} = (50/108) M^2/(c_s^2\epsilon M_{Pl}^2)$. In this case, we obtain $M^2/(c_s^2\epsilon M_{Pl}^2) = 80 \pm 166$ at 68% CL.

9.2. Implications for effective field theory of inflation

The effective field theory approach to inflation (Weinberg 2008; Cheung et al. 2008) provides a general way to scan the NG parameter space of inflationary perturbations. For example, one can expand the Lagrangian of the dynamically relevant degrees of freedom into the dominant operators satisfying some underlying symmetries. We will focus on general single-field models parametrized by the following operators (up to cubic order)

$$S = \int d^4x \sqrt{-g} \left[-\frac{M_{Pl}^2 \dot{H}}{c_s^2} \left(\dot\pi^2 - c_s^2 \frac{(\partial_i\pi)^2}{a^2} \right) \right. \qquad (97)$$
$$\left. - M_{Pl}^2 \dot{H}(1 - c_s^{-2})\dot\pi \frac{(\partial_i\pi)^2}{a^2} + \left(M_{Pl}^2 \dot{H}(1 - c_s^{-2}) - \frac{4}{3}M_3^4 \right)\dot\pi^3 \right]$$

where π is the scalar degree of freedom ($\zeta = -H\pi$). The measurements on f_{NL}^{equil} and f_{NL}^{ortho} can be used to constrain the magnitude of the inflaton interaction terms $\dot\pi(\partial_i\pi)^2$ and $(\dot\pi)^3$ which give respectively $f_{NL}^{EFT1} = -(85/324)(c_s^{-2} - 1)$ and $f_{NL}^{EFT2} = -(10/243)(c_s^{-2} - 1)[\tilde{c}_3 + (3/2)c_s^2]$ (Senatore et al. 2010; see also Chen et al. 2007b; Chen 2010b). These two operators give rise to shapes that peak in the equilateral configuration that are, nevertheless, slightly different, so that the total NG signal will be a linear combination of the two, possibly leading also to an orthogonal shape. There are two relevant NG parameters, c_s, the sound speed of the the inflaton fluctuations, and M_3 which characterizes the amplitude of the other operator $\dot\pi^3$. Following
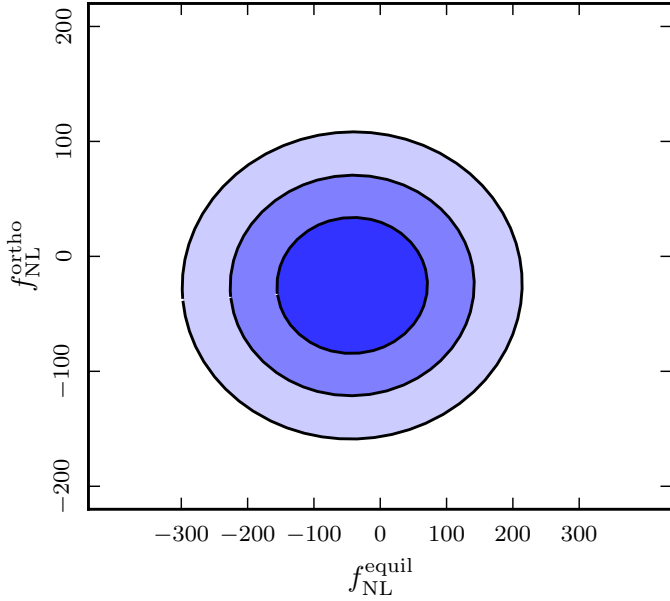
Fig. 24. 68%, 95%, and 99.7% confidence regions in the parameter space $(f_{\rm NL}^{\rm equil}, f_{\rm NL}^{\rm ortho})$, defined by thresholding χ^2 as described in the text.
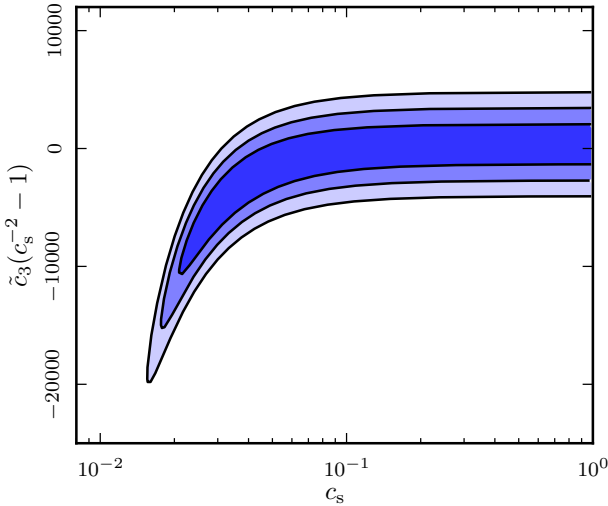

Fig. 25. 68%, 95%, and 99.7% confidence regions in the single-field inflation parameter space $(c_{\rm s}, \tilde{c}_3)$, obtained from Fig. 24 via the change of variables in Eq. (98).

Senatore et al. (2010) we will focus on the dimensionless parameter $\tilde{c}_3(c_{\rm s}^{-2} - 1) = 2M_3^4 c_{\rm s}^2/(\dot{H}M_{\rm Pl}^2)$. For example, DBI inflationary models corresponds to $\tilde{c}_3 = 3(1 - c_{\rm s}^2)/2$, while the non-interacting model (vanishing NG) correspond to $c_{\rm s} = 1$ and $M_3 = 0$ (or $\tilde{c}_3(c_{\rm s}^{-2} - 1) = 0$).

The mean values of the estimators for equilateral and orthogonal NG amplitudes are given in terms of $c_{\rm s}$ and $\tilde{c}_3$ by

$$f_{\rm NL}^{\rm equil} = \frac{1 - c_{\rm s}^2}{c_{\rm s}^2}(-0.275 + 0.0780A)$$

$$f_{\rm NL}^{\rm ortho} = \frac{1 - c_{\rm s}^2}{c_{\rm s}^2}(0.0159 - 0.0167A) \tag{98}$$

where $A = -(c_{\rm s}^2 + (2/3)\tilde{c}_3)$, and the coefficients are computed from the Fisher correlation matrix between the equilateral and orthogonal template bispectra and the theoretical bispectra arising from the two operators $\dot{\pi}(\nabla\pi)^2$ and $\dot{\pi}^3$. Given our constraints on $f_{\rm NL}^{\rm equil}$ and $f_{\rm NL}^{\rm ortho}$, and the covariance matrix **C** of the

joint estimators, we can define a χ^2 statistic given by $\chi^2(\tilde{c}_3, c_{\rm s}) = \boldsymbol{v}^{\rm T}(\tilde{c}_3, c_{\rm s})\mathsf{C}^{-1}\boldsymbol{v}(\tilde{c}_3, c_{\rm s})$, where the vector $\boldsymbol{v}$ is given by $v^i(\tilde{c}_3, c_{\rm s}) = f^i(\tilde{c}_3, c_{\rm s}) - f_P^i$. f_P^i, where $i = \{$equilateral, orthogonal$\}$, are the joint estimates of the equilateral and orthogonal $f_{\rm NL}$ measured by *Planck* and $f^i(\tilde{c}_3, c_{\rm s})$ is given by Eq. (98). Figure 24 shows the 68%, 95%, and 99.7% confidence regions for $f_{\rm NL}^{\rm equil}$ and $f_{\rm NL}^{\rm ortho}$, obtained by requiring $\chi^2 \leq 2.28, 5.99$, and 11.62 respectively, as appropriate for a χ^2 variable with two degrees of freedom. The corresponding confidence regions in the $(\tilde{c}_3, c_{\rm s})$ parameter space are shown in Fig. 25. After marginalizing over $\tilde{c}_3$ we find the following conservative bound on the inflaton sound-speed

$$c_{\rm s} \geq 0.02 \quad 95\% \text{ CL.} \tag{99}$$

Note that we have also looked explicitly for the non-separable shapes in Sect. 7.3.1, in particular the two effective field theory shapes and the DBI inflation shape (see Eqs. (5)–(7)).

9.3. Multi-field models

Curvaton models: Planck NG constraints have interesting implications for the simplest adiabatic curvaton models. They predict (Bartolo et al. 2004d,c)

$$f_{\rm NL}^{\rm local} = \frac{5}{4r_{\rm D}} - \frac{5r_{\rm D}}{6} - \frac{5}{3}, \tag{100}$$

for a quadratic potential of the curvaton field (Lyth & Wands 2002; Lyth et al. 2003; Lyth & Rodriguez 2005; Malik & Lyth 2006; Sasaki et al. 2006), where $r_{\rm D} = [3\rho_{\rm curvaton}/(3\rho_{\rm curvaton} + 4\rho_{\rm radiation})]_{\rm D}$ is the "curvaton decay fraction" evaluated at the epoch of the curvaton decay in the sudden decay approximation. Assuming a prior $0 < r_{\rm D} < 1$, given our constraint $f_{\rm NL}^{\rm local} = 2.7 \pm 5.8$ at 68% CL, we obtain

$$r_{\rm D} \geq 0.15 \quad 95\% \text{ CL.} \tag{101}$$

In Planck Collaboration XXII (2014) a limit on $r_{\rm D}$ is derived from the constraints on isocurvature perturbations under the assumption that there is some residual isocurvature fluctuations in the curvaton field. For this restricted case, they find $r_{\rm D} > 0.98$ (95% CL), compatible with the constraint obtained here.

Quasi-single field inflation: it is beyond the scope of this paper to perform a general multi-field analysis employing the local NG constraints. However, we have performed a detailed analysis for the quasi-single field models (see Eq. (12)). Quasi-single field (QSF) inflation models (Chen & Wang 2010b,a; Baumann & Green 2012) are a natural consequence of inflation model-building in string theory and supergravity (see Sect. 2.2). In addition to the inflaton field, these models have extra fields with masses of order the Hubble parameter, which are stabilized by supersymmetry. A distinctive observational signature of these massive fields is a one-parameter family of large NG whose squeezed limits interpolate between the local and the equilateral shape. Therefore, by measuring the precise momentum-dependence of the squeezed configurations in the NG, in principle, we are directly measuring the parameters of the theory naturally determined by the fundamental principle of supersymmetry. These models produce a bispectrum (Eq. (12)) depending on two parameters $\nu, f_{\rm NL}^{\rm QSI}$, with a shape that interpolates between the local shape, where $\nu = 1.5$ and the equilateral shape, where $\nu = 0$.

Results are shown in Fig. 26 (see Sect. 7.3.6 for details of the analyses). The best fit value corresponds to $\nu = 1.5$, $f_{\rm NL} = 4.79$ which would imply, within this scenario, that the extra field
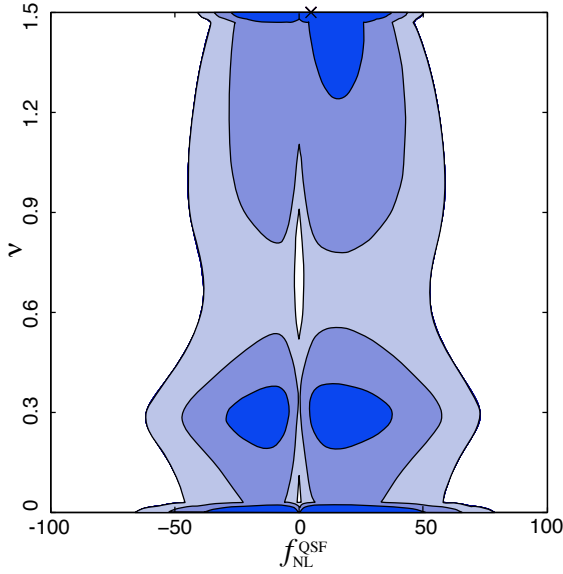
Fig. 26. 68%, 95%, and 99.7% confidence intervals for ν and f_{NL}^{QSI} for quasi-single field inflation. The best fit value of $\nu = 1.5$, $f_{NL}^{QSI} = 4.75$ is marked with an X. The contours were calculated using MC methods by creating 2×10^9 simulations using the β covariance matrix around this best fit model.

different from the inflaton has a mass $m \ll H$. However, the figure shows that there is no preferred value for ν with all values allowed at 3σ.

9.4. Non-standard inflation models

Constraints on excited initial states: results from Sect. 7.3 constrain a variety of models with flattened bispectra, often in combination with a non-trivial squeezed limit. The most notable examples are bispectra produced in excited initial state models (non-Bunch-Davies vacua), which can be generated by strong disturbances away from background slow-roll evolution or additional trans-Planckian physics (Chen et al. 2007b; Holman & Tolley 2008; Meerburg et al. 2009; Agullo & Parker 2011). The constraints we have obtained are summarized in Table 11, and cover four representative cases (see Eqs. (14), (15)) in the literature. We find no strong evidence for these flattened bispectra in the *Planck* data after subtraction of the ISW-lensing signal, with which all these models have some correlation. This is consistent with an earlier constraint on the NBD model from WMAP (Fergusson et al. 2012). However, this investigation is limited by the present resolution of the polynomial modal estimator ($n_{max} = 600$), so more strongly flattened models are not excluded by this analysis.

Directional dependence motivated by fields: directionally-dependent bispectra (Eq. (19)), motivated by inflation with gauge fields, have also been constrained (see Table 11). For example, models with a kinetic term of the vector field(s) $\mathcal{L} = -I^2(\phi)F^2/4$, where F^2 is the strength of the gauge field, and $I(\phi)$ is a function of the inflaton field which, with an appropriate time dependence (see, e.g., Ratra 1992), can generate vector fields during inflation. For these models the $L = 0, 2$ modes in the bispectrum are excited with $f_{NL}^L = X_L(|g_*|/0.1)(N_{k_3}/60)$, where $X_{L=0} = (80/3)$ and $X_{L=2} = -(10/6)$, respectively (Barnaby et al. 2012a; Bartolo et al. 2013; Shiraishi et al. 2013). Here g_* is the amplitude of a quadrupolar anisotropy in the power spectrum (see, e.g., Ackerman et al. 2007) and N is the number of e-folds before the end of inflation when the relevant

scales exit the horizon. These modes therefore relate the bispectrum amplitude to the level of statistical anisotropy in the power spectrum. Marginalizing over a prior $50 \leq N \leq 70$ assuming uniform priors on g_*, our constraints in Table 11 lead to the limits $-0.05 < g_* < 0.05$ and $-0.36 < g_* < 0.36$ from the $L = 0$, $L = 2$ modes respectively (95% CL). Note, however, that in the current analysis only a modest correlation was possible with the shape corresponding to the $L = 2$ mode. These results apply to all models where curvature perturbations are sourced by a $I^2(\phi)F^2$ term (see references in Shiraishi et al. 2013).

Feature models: non-scale-invariant oscillatory bispectrum shapes can be generated by sharp or periodic features in the inflaton potential, with particular recent interest on axion monodromy models motivated by string theory (see Sect. 2.3). We have undertaken a survey of simple feature models (Eq. (16)) which have a periodicity accessible to the polynomial modal estimator (wavenumbers $K = k_1 + k_2 + k_3 \gtrsim 0.01$), a two-parameter space spanned by K and a phase ϕ. There are interesting best fit models for the *Planck* CMB bispectrum around $K = 0.01875$, $\phi = 0$ (that is, with a large-ℓ bispectrum periodicity around $\Delta\ell = 260$), with results shown in Table 12. We note important caveats on the statistics of parameter surveys like this in Sect. 7.3.3; given the large numbers of feature models studied, we have to apply a higher threshold for statistical significance as shown for a survey of 200 Gaussian simulations. This feature survey takes forward earlier results for the WMAP data (Fergusson et al. 2012), with the apparent fit reflecting the signal observed in the *Planck* CMB reconstruction (see Fig. 7). Most attention on feature models has been motivated by the simplest single-field case for which there are correlated signals predicted in the bispectrum and power spectrum (see e.g., Chen et al. 2007a; Adshead et al. 2012). In this case, regions with small $k \sim 0.001$ are favoured for producing an observable bispectrum[17], given existing WMAP power spectrum constraints on these models. Periodicities $\Delta\ell \lesssim 20$ are anticipated (see Adshead et al. 2012) which are not accessible to the present modal bispectrum estimator analysis, but which are discussed in Planck Collaboration XXII (2014). Conversely, the *Planck* feature model survey using the power spectrum (Planck Collaboration XXII 2014) does cover bispectrum scales and parameters investigated in this paper. An analysis of the Bayesian posterior probability (Planck Collaboration XXII 2014) does not appear to provide evidence favoring parameters associated with the current best-fit bispectrum model. More detailed analysis using the specific bispectrum envelope for the single-field feature solution is being undertaken.

Resonance models: we have also investigated resonance models of Eq. (17) such as axion monodromy and enfolded resonance models of Eq. (18), in which non-Bunch-Davies vacua are excited by sharp features. This limited analysis focuses on periodicities associated with the best-fit feature model and the results are described in Tables C.1 and C.2. No significant signal was found in this domain for either of these two models. However, we note that the logarithmic dependence of the bispectrum creates challenges at low k, as we must ensure important features do not fall below the modal resolution. This restricts the present survey range, which will be extended in future. Again, we note that most attention on these models has focused on

---

[17] Planck Collaboration XVI (2014) confirms an anomaly in the power spectrum at $20 \lesssim \ell \lesssim 40$ first noted in WMAP, which leads to an improvement in likelihood when fitted with a feature in the inflationary potential (Peiris et al. 2003). Unfortunately, the best-fit parameters in this case do not lead to an observable bispectrum (Adshead et al. 2011).

higher ℓ-periodicities than those accessible to the present modal estimator (see e.g., Flauger & Pajer 2011; Peiris et al. 2013), for the same reason as the feature models. A resonance model survey using the *Planck* power spectrum has been undertaken and the results can be found in Planck Collaboration XXII (2014); however, this also currently excludes high frequencies that have received attention in the literature.

Warm inflation: this model, where dissipative effects are important, predicts $f_{NL}^{warm} = -15 \ln(1 + r_d/14) - 5/2$ (Moss & Xiong 2007) where the dissipation parameter $r_d = \Gamma/(3H)$ must be large for strong dissipation. The limit from WMAP is $r_d \leq 2.8 \times 10^4$ (Moss & Xiong 2007). Assuming a prior $0 \leq \log_{10} r_d \leq 4$, our constraint $f_{NL}^{warmS} = 4 \pm 33$ at 68% CL (see Sect. 7.3.5) yields a limit on the dissipation parameter of $\log_{10} r_d \leq 2.6$ (95% CL), improving the previous limit by nearly two orders of magnitude. The strongly-dissipative regime with $r_d \gtrsim 2.5$ still remains viable; however, the *Planck* constraint puts the model in a regime which can lead to an overproduction of gravitinos (see Hall & Peiris 2008, and references within).

9.5. Alternatives to inflation

Perhaps the most striking example is given by the ekpyrotic/cyclic models (for a review, see Lehners 2010) proposed as alternative to inflationary models. Typically they predict a local NG $|f_{NL}^{local}| > 10$. In particular, the so-called "ekpyrotic conversion" mechanism (in which isocurvature perturbations are converted into curvature perturbations during the ekpyrotic phase) yields $f_{NL}^{local} = -(5/12) c_1^2$, where c_1 is a parameter in the potential, requiring $10 \gtrsim c_1 \gtrsim 20$ for compatibility with power spectrum constraints. This case was $\sim 4\sigma$ discrepant with WMAP data, and with *Planck* it is decisively ruled out given our bounds $f_{NL}^{local} = 2.7 \pm 5.8$ at 68% CL (see Table 9) yielding $c_1 \leq 4.2$ at 95% CL. The predictions for the local bispectrum from other ekpyrotic models (based on the so called "kinetic conversion" taking place after the ekpyrotic phase) yield $f_{NL}^{local} = (3/2) \kappa_3 \sqrt{\epsilon} + 5$, where values $-1 < \kappa_3 < 5$ and $\epsilon \sim 100$ are natural (Lehners 2010). Taking $\epsilon \sim 100$, in this case we obtain $-0.9 < \kappa_3 < 0.6$ and $-0.2 < \kappa_3 < 1.3$ at 95% CL, for the plus and minus sign in f_{NL}^{local} respectively, significantly restricting the viable parameter space of this model.

9.6. Implications of CMB trispectrum results

The non-detection of local-type f_{NL} by *Planck* raises the immediate question of whether there might still be a large trispectrum. In this first analysis we have focused on the local shape τ_{NL}. Most inflation models predict $\tau_{NL} \sim O(f_{NL}^2)$ (Byrnes et al. 2006; Elliston et al. 2012), and hence given our tight f_{NL} limits, would be predicted to be very small. This is easily consistent with our conservative upper limit $\tau_{NL} < 2800$, and also with the significantly smaller signals found in the modulation dipole and quadrupole. Our upper limit also restricts the freedom of curvaton-like models with quadratic terms that are nearly uncorrelated with the curvature perturbation, which could in principle have $f_{NL} \sim 0$ but a large trispectrum (Byrnes & Choi 2010).

10. Conclusions

In this paper we have derived constraints on primordial NG, using the CMB maps derived from the *Planck* nominal mission data. Using three optimal bispectrum estimators – KSW, binned, and modal – we obtained consistent values for the primordial local, equilateral, and orthogonal bispectrum amplitudes, quoting as our final result $f_{NL}^{local} = 2.7 \pm 5.8$, $f_{NL}^{equil} = -42 \pm 75$, and $f_{NL}^{ortho} = -25 \pm 39$ (68% CL statistical). Hence there is no evidence for primordial NG of one of these shapes. We did, however, measure

the ISW-lensing bispectrum expected in the ΛCDM scenario, as well as a contribution from diffuse point sources, and these contributions are clearly seen in the form of the associated skew-C_ℓ. Indeed, the detection of ISW-lensing and Poisson point source skew-C_ℓ with the right shapes and at expected levels is good evidence that the data processing is not destroying NG in the maps, and gives confidence that the non-detections of primordial NG are robust. These results have been confirmed by measurements using the wavelet bispectrum, and Minkowski functional estimators, and demonstrated to be stable for the four different component separation techniques SMICA, NILC, SEVEM, and C-R, showing their robustness against foreground residuals. They have also passed an extensive suite of tests studying the dependence on the maximum multipole number and the mask, consistency checks between frequency channels, and several null tests. In addition, we have summarized in this paper an extensive validation campaign for the three optimal bispectrum estimators on Gaussian and non-Gaussian simulations.

Extending our analysis beyond estimates of individual shape amplitudes, we presented model-independent, three-dimensional reconstructions of the *Planck* CMB bispectrum using the modal and binned bispectrum estimators. These results were also shown to be fully consistent between the different component separation techniques even for the full bispectrum, and contained no significant NG signals of a type not captured by our standard templates at low multipole values. At high multipoles, some indications of unidentified NG signals were found, as also evidenced by the results from the skew-C_ℓ estimator. Further study will be required to ascertain whether these indications are due to foreground residuals, beams, data processing, or a more interesting signal.

Using the modal decomposition, we have presented constraints on key primordial NG scenarios, including general single-field models with non-separable shapes, excited initial states (non-Bunch-Davies vacua), and directionally-dependent vector models. We have also undertaken an initial survey of scale-dependent feature and resonance models.

Moving beyond three-point correlations, we have obtained limits from the *Planck* data on the amplitude of the local four-point function. Our trispectrum reconstruction yielded a signal consistent with zero except for an anomalously large octopole. Frequency dependence indicated that this was unlikely to be primordial, but allowing the signal to be primordial we placed an upper limit $\tau_{NL} < 2800$ (95% CL).

We have discussed the impact of these results on the inflationary model space, and derived bispectrum constraints on a selection of specific inflationary mechanisms, including both general single-field inflationary models (extensions to the standard single-field slow-roll case) and multi-field models. Our results led to a lower bound on the speed of sound, $c_s \geq 0.02$ (95% CL), in an effective field theory approach to inflationary models which includes models with non-standard kinetic terms. Strong constraints on other scenarios such as IR DBI, k-inflation, inflationary models involving gauge fields, and warm inflation have been obtained. Within the class of multi-field models, our measurements limit the curvaton decay fraction to $r_D \geq 0.15$ (95% CL). Ekpyrotic/cyclic scenarios were shown to be under pressure from the *Planck* data: we robustly ruled out the so-called "ekpyrotic conversion" mechanism, and found that the parameter space of the "kinetic conversion" mechanism is significantly limited.

With these results, the paradigm of standard single-field slow-roll inflation has survived its most stringent tests to-date.

Appendix A: Full focal plane simulations

The purpose of the full focal plane (FFP) simulations is to provide a complete, coherent realization of the full *Planck* mission (both HFI and LFI) using the best available estimates of its characteristics, including the satellite pointing, the individual detector beams, band passes and noise properties, and the various data flags. Each FFP data set consists of three parts: a fiducial observation of an input sky (including our best estimates of all of the foreground components and a realization of the CMB sky corresponding to a chosen cosmology), together with separate MC realization sets of the CMB and noise. The maps made from these simulated data are used both to validate and verify the analysis algorithms and their implementations and to quantify uncertainties in and remove biases from the analysis of the real data. The FFP simulations have been used for validating map-making, component separation (for both CMB and diffuse foreground recovery), power spectra estimation, cosmological parameter estimation and measurements of primordial NG and lensing. The complete specification of any FFP simulation consists of the definition of the inputs (mission characteristics and sky), the outputs (time-ordered data and maps), and the processes used to generate the latter from the former. For this data release we have used the sixth FFP simulation suite, henceforth referred to as FFP6.

A.1. Mission characteristics

The goal of FFP6 is to replicate the 2013 *Planck* data release (DR1) as closely as possible. The satellite pointing incorporates the PTCOR6 wobble corrections used by HFI (Planck Collaboration VI 2014), and the focal plane geometry of each instrument is provided by the corresponding DPC in the form of a reduced instrument model (RIMO) (Planck Collaboration II 2014, and Planck Collaboration VI 2014). The overall data flags are a combination of the pointing and detector flags. For both instruments, satellite repointing maneuvers and planet crossings are flagged. For pairs of detectors in the same horn, each sample flagged in one detector is flagged in the other member of the pair.

In addition to the geometric detector offsets, the RIMOs for both instruments provide a representation of the detector

bandpasses. For LFI, the RIMO also provides a parameterized noise model (white noise level, knee frequency and power law index) for each detector (Planck Collaboration II 2014). LFI beams were measured from Jupiter scans, fitted using Gaussian circular and elliptical approximations, converted into a polarized beam model, and smeared to account for the satellite motion (Planck Collaboration IV 2014). For HFI, the noise power spectral densities (PSDs) are averaged and smoothed versions of the raw ring-by-ring spectra determined by the HFI pipeline. This processing produced the per detector mean noise spectrum over the duration of the nominal mission, calibrated to convert from pre-processed time ordered information (TOI) units to thermodynamic K (Planck Collaboration VI 2014). HFI beam maps were synthesized from the elliptical beam parameters and Gauss-Hermite coefficients from the Mars observation (Planck Collaboration VII 2014).

A.2. Input sky

The FFP6 input sky is generated using the *Planck* Sky Model (PSM; Delabrouille et al. 2013) and includes the CMB together with the current best estimate foreground templates for the cosmic infrared background, CO line emission, free-free, synchrotron, spinning dust, thermal dust, kinetic and thermal SZ, and radio and infrared point sources. For simulation purposes, point sources are considered strong if their flux is more than 100 mJy at any of 30, 100, 300 or 1000 GHz; strong point sources are provided in catalogs, whilst weak sources are mapped. All maps are provided in thermodynamic K at HEALPix resolution 2048 (Górski et al. 2005) and smoothed with a 4 arcmin beam, while catalogues are in Jy. The CMB is generated as a lensed sum of a scalar, tensor and non-Gaussian realization with particular values of the tensor to scalar ratio (r) and NG parameter (f_{NL}). For every detector, each foreground component and catalogue of strong point sources also takes the appropriate band-passed from the RIMO and applies it.

A.3. FFP6 outputs

The FFP6 simulation consists of maps of the fiducial sky together with sets of MC realizations of the CMB and of the noise, each of which is produced using a distinct processing pipeline; full details on these can be found in the Explanatory Supplement (Planck Collaboration 2013).

The fiducial sky maps are made from explicit time-ordered data sets which are generated using the Level S software suite (Reinecke et al. 2006) for every detector and for every component; this results in 770 time-ordered datasets occupying 17TB of disk space. For every component we then make maps of the data at each frequency using the `Madam/TOAST` map-maker (Keihänen et al. 2010) for all the detectors, all of the detector quadruplets and all of the unpolarized detector singlets, and using both the full data and the two half-ring data subsets. In addition we make maps of the sum of all of the components' time streams with simulated noise added. HFI data are mapped at HEALPix resolution 2048 and LFI data at 1024; this results in 2370 maps occupying 0.4TB of disk space. We also produced hit maps, binned maps and the 3×3 block diagonal white-noise covariance matrices and their inverses.

The CMB MC provides a set of 1000 temperature-only maps of the full data combining all of the detectors at each frequency. These are made directly in the map domain by calculating the effective beam at each frequency using the `FEBeCoP`

software ([Mitra et al. 2011](#)) and applying it to each of 1000 realizations of the CMB sky drawn from the best-fit WMAP spectrum and generated using the HEALPix `synfast` tool ([Górski et al. 2005](#)); this results in 9000 maps occupying 1.2TB of disk space.

The MC noise maps are generated using `Madam/TOAST`, which has the ability to simulate the noise time stream on the fly (bypassing the otherwise cripplingly expensive writing and reading of time-ordered data objects). For every detector combination at each frequency, and for both the full data and half-ring data subsets, we generate 1000 noise realizations; this results in 222 000 maps occupying 33TB of disk space.

The bulk of the FFP simulations were performed at the US Department of Energy's National Energy Research Scientific Computing Center (NERSC); the LFI noise MC for the various detector quadruplets were produced at the Finnish Supercomputing Center (CSC); in total the generation of FFP6 required some 15 million CPU-hours and 50TB of disk space. All of the temperature data in the FFP6 simulation set are available for public use at NERSC.

A.4. Validation and exploitation

Specific tests for assessing our confidence in the claims were conducted for all the main results in the *Planck* 2013 release, namely CMB reconstruction through foreground cleaning, likelihood to cosmology, lensing and primordial NG; these are discussed in detail in the corresponding papers ([Planck Collaboration XII 2014](#); [Planck Collaboration XV 2014](#); [Planck Collaboration XVI 2014](#); [Planck Collaboration XVII 2014](#); [Planck Collaboration XXIII 2014](#) and this work). The validation was conducted with the necessary accuracy for measuring second-order cosmological effects like in the case of primordial NG as well as lensing. In the former case, blind tests on an FFP set with a non-zero and detectable $f_{\mathrm{NL}}^{\mathrm{local}} = 20.4075$ value in the CMB component were performed and the correct value (within the error bars) was detected after foreground cleaning. In the case of lensing, null detection tests were conducted successfully after the process of foreground cleaning on maps that did not contain this signal as input.

Residual systematics which were not taken into account belong to the low-ℓ and high-ℓ regimes for LFI and HFI, respectively. Although the latter were not translated directly into spurious NG effects, their effect was quantified. For LFI, sidelobe corrections have been quantified to a level which is comparable or below 0.05% in the beam transfer function on the range of ℓ between 200 and 1500, being increasingly sub-dominant with respect to noise at high ℓs. A number of HFI systematic effects and their treatment in the data processing are not included in the simulations. On large angular scales, the most important effects are the Analogue Digital Correction (ADC) non-linearities (which principally manifest themselves as an apparent gain drift), Zodiacal light emission, and far sidelobe pick-up of the Galaxy. On small angular scales the most important effects are cosmic ray hits (flagging of those from real data was taken into account into the FFP6 processing) and the 4He-JT (4K) cooler lines. The processing steps used to remove these effects and assessments of the residuals are described in [Planck Collaboration VI](#) ([2014](#)); [Planck Collaboration VIII](#) ([2014](#)); [Planck Collaboration X](#) ([2014](#)); [Planck Collaboration XIV](#) ([2014](#)).

Planck will be updating its suite of simulations for the forthcoming release including polarization. The instrumental characteristics described above will be updated, and particular care is being taken regarding the ability to simulate and control the main systematic effects affecting the polarised signal.

Appendix B: Expected level of agreement from bispectrum estimators with correlated weights

The estimator cross-validation work presented in Sect. [6.1](#) was based on comparing results from different estimators using sets (typically 50 to 100 simulations in size) of both Gaussian and non-Gaussian simulations. We started from idealized maps (e.g., full-sky, noiseless, Gaussian simulations) and then went on to include an increasing number of realistic features at each additional validation step. This allowed better testing and characterization of the response of different pipelines and bispectrum decompositions to various potential spurious effects in the data. As a preliminary step, we derived a general formula describing the expected level of agreement between estimators with different but strongly correlated weights, with the simplifying assumption of full-sky measurements and homogeneous noise. This theoretical expectation, summarized by Eq. ([84](#)), was then used as a benchmark against which to assess the quality of the results. The aim of this appendix is to describe in detail how we obtained Eq. ([84](#)).

Let us consider the idealized case of full-sky, noiseless CMB measurements (note that the following conclusions also work for homogeneous noise, because the pure cubic f_{NL} estimators without linear term corrections are still optimal). Under these assumptions, the f_{NL} estimator for a given CMB shape $B_{\ell_1 \ell_2 \ell_3}$ can be written simply as (see Sect. [3](#) for details):

$$\hat{f}_{\mathrm{NL}} = \frac{1}{F} \sum_{\ell_1 \leq \ell_2 \leq \ell_3} \sum_{m_1 m_2 m_3} \frac{B_{\ell_1 \ell_2 \ell_3}^{\mathrm{th}} a_{\ell_1 m_1} a_{\ell_2 m_2} a_{\ell_3 m_3}}{g_{\ell_1 \ell_2 \ell_3} C_{\ell_1} C_{\ell_2} C_{\ell_3}}, \tag{B.1}$$

where $B_{\ell_1 \ell_2 \ell_3}^{\mathrm{th}}$ is the angle-averaged bispectrum template for a given theoretical shape, $a_{\ell_1 m_1} a_{\ell_2 m_2} a_{\ell_3 m_3}$ is a bispectrum estimate constructed from the data, and F is the normalization of the estimator, provided by the Fisher matrix number

$$F = \sum_{\ell_1 \leq \ell_2 \leq \ell_3} \frac{(B_{\ell_1 \ell_2 \ell_3}^{\mathrm{th}})^2}{g_{\ell_1 \ell_2 \ell_3} C_{\ell_1} C_{\ell_2} C_{\ell_3}}. \tag{B.2}$$

As explained in Sect. [3](#), the different f_{NL} estimation techniques implemented in this work can be seen as separate implementations of the optimal estimator of Eq. ([B.1](#)). Each implementation is based on expanding the angular bispectrum as a sum of basis templates defined in different domains: for example in our analyses we built templates out of products of wavelets at different scales, cubic combinations of 1-dimensional polynomials and plane waves (what we call "modal estimator" in the main text), and ℓ-binning of the bispectrum (what we called "binned estimator" in the main text). Our initial theoretical templates in this work are the local, equilateral and orthogonal separable cases used in the KSW and skew-C_ℓ estimators. In this sense the KSW/skew-C_ℓ estimator provides an "exact" estimate of f_{NL} for this choice of shapes, while the other pipelines provide an approximate result that approaches KSW measurements as the expansions get more accurate. The differences between results from different pipelines became smaller and smaller as the approximate expanded templates converge to the starting one (e.g., by increasing the number of ℓ-bins or the number of wavelets

and polynomial triplets). The degree of convergence can be measured through the correlation coefficient r between the initial bispectrum and its reconstructed expanded version. The correlation coefficient is, as usual, defined as

$$r \equiv \frac{\sum_{\ell_1 \leq \ell_2 \leq \ell_3} \frac{B^{\rm th}_{\ell_1 \ell_2 \ell_3} B^{\rm exp}_{\ell_1 \ell_2 \ell_3}}{g_{\ell_1 \ell_2 \ell_3} C_{\ell_1} C_{\ell_2} C_{\ell_3}}}{\sqrt{\sum_{\ell_1 \leq \ell_2 \leq \ell_3} \frac{(B^{\rm th}_{\ell_1 \ell_2 \ell_3})^2}{g_{\ell_1 \ell_2 \ell_3} C_{\ell_1} C_{\ell_2} C_{\ell_3}} \sum_{\ell_1 \leq \ell_2 \leq \ell_3} \frac{(B^{\rm exp}_{\ell_1 \ell_2 \ell_3})^2}{g_{\ell_1 \ell_2 \ell_3} C_{\ell_1} C_{\ell_2} C_{\ell_3}}}} ,\qquad (B.3)$$

where the label "th" denotes the initial bispectrum shape to fit to the data, and "exp" is the approximate expanded one. The correlator between shapes naturally defines a scalar product in bispectrum space, and in the following the operation of correlating two shapes will be denoted by the symbol $\langle , \rangle$, so that the definition above would read $r = \langle B^{\rm th}, B^{\rm exp}\rangle / \sqrt{\langle B^{\rm exp}, B^{\rm exp}\rangle \langle B^{\rm th}, B^{\rm th}\rangle}$.

Whichever basis and separation scheme we have chosen, let us call $\mathcal{R}_n(\ell_1, \ell_2, \ell_3)$ the n-th template in the adopted bispectrum expansion, and α_n the coefficients of the expansion, so that we can write:

$$B^{\rm exp}_{\ell_1 \ell_2 \ell_3} = \sum_{n=0}^{N_{\rm exp}} \alpha_n \mathcal{R}_n(\ell_1, \ell_2, \ell_3) . \qquad (B.4)$$

From now on we will also call $\mathcal{R}_n(\ell_1, \ell_2, \ell_3)$ the "modes" of our expansion, $N_{\rm exp}$ is the number of modes we are using to approximate the starting template in order to obtain a correlation coefficient r. We will also assume that the modes form an orthonormal basis, that is:

$$\langle \mathcal{R}_{n_1}, \mathcal{R}_{n_2} \rangle = \delta^{n_2}_{n_1} , \qquad (B.5)$$

where $\delta^{n_2}_{n_1}$ is the Kronecker delta symbol. The orthogonality assumption does not imply loss of generality since any starting set of modes can always be rotated and orthogonalized. We now consider an expansion with a number of coefficients $N_{\rm th} > N_{\rm exp}$ that perfectly reproduces the initial bispectrum (i.e., $r = 1$), and is characterized by the same modes and coefficients as the previous one up to $N_{\rm exp}$. So we can write

$$B^{\rm th}_{\ell_1 \ell_2 \ell_3} = \sum_{n=0}^{N_{\rm th}} \alpha_n \mathcal{R}_n(\ell_1, \ell_2, \ell_3) . \qquad (B.6)$$

We now build two optimal estimators of $f_{\rm NL}$ for the shape $B_{\rm th}$: an "exact" estimator and an "approximate" one. In the exact estimator we fit the actual $B_{\rm th}$ shape to the data and obtain the estimate $\hat{f}_{\rm th}$, while in the approximate estimator we fit the expanded shape $B_{\rm exp}$ to obtain $\hat{f}_{\rm exp}$. We want to understand by how much the "exact" and "approximate" $f_{\rm NL}$ measurements are expected to differ, as a function of the correlation coefficient r between the weights $B_{\rm th}$ and $B_{\rm exp}$ that appear in the two estimators.

For each mode template $\mathcal{R}_n(\ell_1, \ell_2, \ell_3)$ we can build an optimal estimator (following Eq. (B.1)) in order to fit the mode to the data and get a corresponding amplitude estimate β_n. In other words, the observed bispectrum can then be reconstructed as in Eq. (B.4), but with coefficients β_n instead of α_n. Due to the orthonormality of the $\mathcal{R}_n$, the β coefficients have unit variance. It is then possible to show (Fergusson et al. 2010a) that the $f_{\rm NL}$ estimate for a given $B_{\rm th}$ with expansion parameters α_n is given by

$$\hat{f}_{\rm NL} = \frac{1}{F} \sum_n \alpha_n \beta_n. \qquad (B.7)$$

In the light of all the above, the exact and approximate estimators are:

$$\hat{f}^{\rm th}_{\rm NL} = \frac{1}{F_{\rm th}} \sum_{n=0}^{N_{\rm th}} \alpha_n \beta_n; \qquad (B.8)$$

$$\hat{f}^{\rm exp}_{\rm NL} = \frac{1}{F_{\rm exp}} \sum_{n=0}^{N_{\rm exp}} \alpha_n \beta_n. \qquad (B.9)$$

Thanks to the orthonormality properties of the modes, we can easily relate the Fisher matrix normalization F to the expansion coefficients α:

$$\begin{aligned}
F &\equiv \sum_{\ell_1 \leq \ell_2 \leq \ell_3} \frac{B^2_{\ell_1 \ell_2 \ell_3}}{g_{\ell_1 \ell_2 \ell_3} C_{\ell_1} C_{\ell_2} C_{\ell_3}} \\
&= \sum_{\ell_1 \leq \ell_2 \leq \ell_3} \frac{\sum_{n_1=0}^{N} \alpha_{n_1} \mathcal{R}_{n_1}(\ell_1, \ell_2, \ell_3) \sum_{n_2=0}^{N} \alpha_{n_2} \mathcal{R}_{n_2}(\ell_1, \ell_2, \ell_3)}{g_{\ell_1 \ell_2 \ell_3} C_{\ell_1} C_{\ell_2} C_{\ell_3}} \\
&= \sum_{n_1=0}^{N} \alpha_{n_1} \sum_{n_2=0}^{N} \alpha_{n_2} \langle \mathcal{R}_{n_1}, \mathcal{R}_{n_2} \rangle \\
&= \sum_{n_1=0}^{N} \sum_{n_2=0}^{N} \alpha_{n_1} \alpha_{n_2} \delta^{n_2}_{n_1} \\
&= \sum_{n=0}^{N} \alpha_n^2 . \qquad (B.10)
\end{aligned}$$

In analogous fashion we can derive an expression for the correlation coefficient r between the two estimators we are comparing. It is easy to check that the following holds:

$$r^2 = \frac{\sum_{n=0}^{N_{\rm exp}} \alpha_n^2}{\sum_{n=0}^{N_{\rm th}} \alpha_n^2}, \qquad (B.11)$$

and using the equation just derived above we can also write $r^2 = F_{\rm exp}/F_{\rm th}$, i.e., the square of the correlation coefficients between the estimators equals the ratio of the normalizations.

Armed with this preliminary notation we can now calculate the expected scatter between the exact and approximate $f_{\rm NL}$ measurement when we apply the two estimators to the same set of data. In order to quantify it we will consider the variable $\delta_{f_{\rm NL}} \equiv \hat{f}^{\rm th}_{\rm NL} - \hat{f}^{\rm exp}_{\rm NL}$, and calculate its standard deviation. We find

$$\begin{aligned}
\sigma^2_{\delta f_{\rm NL}} &\equiv \left\langle \left(\frac{1}{F_{\rm th}} \sum_{n=0}^{N_{\rm th}} \alpha_n \beta_n - \frac{1}{F_{\rm exp}} \sum_{n=0}^{N_{\rm exp}} \alpha_n \beta_n \right)^2 \right\rangle \\
&= \frac{1}{F^2_{\rm exp}} \left\langle \left(r^2 \sum_{n=0}^{N_{\rm th}} \alpha_n \beta_n - \sum_{n=0}^{N_{\rm exp}} \alpha_n \beta_n \right)^2 \right\rangle,
\end{aligned}$$
$$(B.12)$$

where we made use of Eq. (B.11). It is easy to show that orthonormality of the $\mathcal{R}$ templates implies no correlation of the amplitudes β, i.e., $\langle \beta_p \beta_q = \delta^q_p \rangle$. A straightforward calculation then yields:

$$\sigma^2_{\delta f_{\rm NL}} = \frac{1}{F^2_{\rm exp}} \sum_{n=0}^{N_{\rm exp}} \alpha_n^2 - 2r^2 \sum_{n=0}^{N_{\rm exp}} \alpha_n^2 + r^4 \sum_{n=0}^{N_{\rm th}} \alpha_n^2. \qquad (B.13)$$

This, together with Eqs. (B.10, B.11) finally gives:

$$\sigma_{\delta f_{\rm NL}} = \Delta_{\rm th} \frac{\sqrt{1 - r^2}}{r}, \qquad (B.14)$$

Table C.1. Results from a limited $f_{\rm NL}$ survey of resonance models of Eq. (17) with $0.25 \leq k_{\rm c} \leq 0.5$ using the SMICA component-separated map.

Wavenumber $k_{\rm c}$	$f_{\rm NL} \pm \Delta f_{\rm NL}$					
	$\phi = 0$	$\phi = \pi/5$	$\phi = 2\pi/5$	$\phi = 3\pi/5$	$\phi = 4\pi/5$	$\phi = \pi$
0.25	-16 ± 57	6 ± 63	19 ± 67	31 ± 69	38 ± 68	-6 ± 60
0.30	-66 ± 73	-57 ± 74	-44 ± 73	-26 ± 72	-7 ± 71	-65 ± 73
0.40	5 ± 57	40 ± 66	55 ± 71	63 ± 73	63 ± 71	22 ± 61
0.45	25 ± 56	34 ± 59	36 ± 62	34 ± 67	27 ± 69	30 ± 56
0.50	-2 ± 65	-13 ± 72	-16 ± 69	-16 ± 60	-14 ± 55	-7 ± 71

Notes. These models have a large-ℓ periodicity similar to the feature models in Table 12.

Table C.2. Results from a limited $f_{\rm NL}$ survey of non-Bunch-Davies feature models (or enfolded resonance models) of Eq. (18) with $4 \leq k_{\rm c} \leq 12$, again overlapping in periodicity with the feature model survey.

Wavenumber $k_{\rm c}$	$f_{\rm NL} \pm \Delta f_{\rm NL}$ (σ)			
	$\phi = 0$	$\phi = \pi/4$	$\phi = \pi/2$	$\phi = 3\pi/4$
4	11 ± 146 (0.1)	2 ± 145 (0.0)	-7 ± 143 (0.0)	-15 ± 142 (-0.1)
6	52 ± 202 (0.3)	63 ± 203 (0.3)	72 ± 201 (0.4)	80 ± 197 (0.4)
8	100 ± 190 (0.5)	130 ± 189 (0.7)	158 ± 189 (0.8)	183 ± 190 (1.0)
10	188 ± 241 (0.8)	210 ± 242 (0.9)	230 ± 242 (1.0)	248 ± 243 (1.0)
12	180 ± 307 (0.6)	171 ± 310 (0.6)	158 ± 312 (0.5)	142 ± 314 (0.5)

where $\Delta_{\rm th}$ is the standard deviation of the exact estimator. Equation (B.14) is the starting point of our validation tests. It provides an estimate of the expected scatter between $f_{\rm NL}$ estimators with correlated weights, as a function of the $f_{\rm NL}$ error bars and of the correlation coefficient r. Note that this formula has been obtained under the simplifying assumptions of Gaussianity, full-sky coverage and homogeneous noise. For applications dealing with more realistic cases we might expect the scatter to become *larger* than this expectation, while remaining qualitatively consistent.

In our tests we started from sets of simulated maps and compared $f_{\rm NL}$ results from different pipelines and shapes on a map-by-map basis. In our validation tests the correlation levels between templates in different expansions were varying between $r \sim 0.95$ and $r \sim 0.99$, depending on the estimators and the shapes under study. Using the formula above, this corresponds to an expected scatter $0.15 \Delta \leq \sigma \leq 0.3 \Delta$, where Δ is the $f_{\rm NL}$ standard deviation for the shape under study.

In Sect. 6.1 we presented several applications to simulated data sets, showing that the recovered results are fully consistent with these expectations, and thus the different pipelines are fully consistent with each other.

Appendix C: Targeted bispectrum constraints

This Appendix provides further tables of results for primordial models, extending those given in Sect. 7.3, notably for resonance models, while it also gives additional validation checks for the modal bispectrum estimator, beyond the extensive tests described already in the paper for the standard bispectrum shapes. In particular, using each of the SMICA NILC and SEVEM maps, we will quote results for the main paradigms for non-standard bispectrum models, including comparisons for the feature model results. Remarkably consistent results are again obtained using the different foreground-separation methodologies and using different modal basis functions.

C.1. Additional results for resonance models and enfolded resonance models

As described in Sect. 7.3, we have surveyed resonance models (Eq. (17)) in the region of most interest for feature models, that is, with comparable periodicities to those with described in Table 12 near the best-fit feature model. The results from this initial survey for the SMICA component-separated map produced no significant signal, with the results Table C.1. We note that while the feature and resonance models proved similar for the WMAP analysis (Fergusson et al. 2012), wavelengths are much more differentiated for *Planck* and so it can be difficult to resolve the shortest wavelengths at low ℓ. This is the key limitation on surveys with the present estimator and will be circumvented in future, by using specific separable templates to improve the overall resolution. Just as local and ISW templates can be incorporated into the analysis, so can targeted feature modes. Note that consistent results were obtained using the NILC and SEVEM maps, though we only show results here for feature models with an envelope (see discussion below with Table C.4).

The enfolded resonant (or non-Bunch-Davies feature) models (Eq. (18)) were also surveyed for these periodicities and also yielded no significant signal; see Table C.2. These are interesting shapes which hold out the prospect of exhibiting both oscillatory and flattened features observed in the *Planck* bispectrum reconstruction, see Fig. 7. Due to similar resolution restrictions, only relatively large multipole periodicities ($\ell > 100$) have been surveyed in the present work, again searching around the periodicities exhibited for feature models.

C.2. Comparison of $f_{\rm NL}$ results from SMICA, NILC and SEVEM foreground-cleaned maps

Tables C.3 and C.4 compare bispectrum results extracted from *Planck* maps created using the three different component-separation techniques, SMICA, NILC and SEVEM. In addition

Table C.3. Summary of results from the modal estimator survey of primordial models for the main non-standard bispectrum shapes.

	F_{NL}	F_{NL}-Clean	StDev	Fisher	σ	σ-clean
SMICA						
Const (see text)	26	14	44	42	0.6	0.3
EFT1 shape (5)	13	8	73	70	0.2	0.1
EFT2 shape (5)	27	19	57	54	0.5	0.3
DBI inflation (7)	17	11	69	67	0.2	0.2
Ghost inflation (see text)	−27	−24	88	87	−0.3	−0.3
Flat model (13)	70	37	77	71	0.9	0.5
NBD (see text)	178	155	78	76	2.2	2.0
NBD1 flattened (14)	31	19	13	12	2.4	1.4
NBD2 squeezed (14)	0.8	0.2	0.4	0.5	1.8	0.5
NBD3 non-canonical (15)	13	9.6	9.7	9.0	1.3	1.0
Vector model $L = 1$ (19)	−18	−4.6	47	45	−0.4	−0.1
Vector model $L = 2$ (19)	2.8	−0.4	2.9	2.8	1.0	−0.1
WarmS inflation (see text)	−8	1	33	33	−0.2	0.04
SEVEM						
Const (see text)	23	11	44	42	0.5	0.2
EFT1 shape (5)	9	−1	72	71	0.1	−0.02
EFT2 shape (5)	21	14	56	54	0.4	0.2
DBI inflation (7)	13	7	68	67	0.2	0.1
Ghost inflation	−24	−21	88	88	−0.3	−0.2
Flat model (13)	63	31	76	72	0.8	0.4
NBD (see text)	159	137	78	76	2.0	1.8
NBD1 flattened (14)	30	18	12	12	2.4	1.4
NBD2 squeezed (14)	0.9	0.4	0.4	0.5	2.1	0.8
NBD3 non-canonical (15)	12	9	10	9	1.2	0.9
Vector model $L = 1$ (19)	−15	−2	47	45	−0.3	−0.04
Vector model $L = 2$ (19)	3.5	0.3	2.7	2.8	1.3	0.1
WarmS inflation	−11	−2	33	33	−0.3	−0.1
NILC						
Const (see text)	37	25	44	42	0.8	0.6
EFT1 shape (5)	20	−4	72	70	0.3	−0.05
EFT2 shape (5)	39	32	56	54	0.7	0.6
DBI inflation (7)	26	20	69	67	0.4	0.3
Ghost inflation	−36	−33	88	87	−0.4	−0.4
Flat model (13)	100	68	76	71	1.3	0.9
NBD (see text)	189	165	78	76	2.4	2.1
NBD1 flattened (14)	35	22	13	12	2.7	1.7
NBD2 squeezed (14)	0.8	0.3	0.4	0.5	1.9	0.6
NBD3 non-canonical (15)	17	13	9	9	1.8	1.4
Vector model $L = 1$ (19)	−41	−28	46	45	−0.9	−0.6
Vector model $L = 2$ (19)	3.8	0.6	2.6	2.8	1.4	0.2
WarmS inflation	−23	−14	32	33	−0.7	−0.4

Notes. This is an extended version of Table 11, but with results from SMICA, NILC and SEVEM.

Table C.4. Cross-validation of best fit feature model results for the SMICA, NILC and SEVEM foreground-cleaned maps.

Wavenumber Phase	$fnl \pm \Delta f_{NL}$ (σ)		
	NILC	SEVEM	SMICA
$k_c = 0.01125$ $\phi = 0$	458 ± 169 (2.7)	409 ± 169 (2.4)	434 ± 170 (2.6)
$k_c = 0.01750$ $\phi = 0$	−337 ± 131 (−2.6)	−328 ± 128 (−2.6)	−335 ± 137 (−2.4)
$k_c = 0.01750$ $\phi = 3\pi/4$ . . .	368 ± 124 (3.0)	348 ± 121 (2.9)	366 ± 126 (2.9)
$k_c = 0.01875$ $\phi = 0$	−359 ± 118 (−3.1)	−366 ± 115 (−3.2)	−348 ± 118 (−3.0)
$k_c = 0.01875$ $\phi = \pi/4$	−339 ± 117 (−2.9)	−328 ± 115 (−2.9)	−323 ± 120 (−2.7)
$k_c = 0.02000$ $\phi = \pi/4$	−305 ± 118 (−2.6)	−334 ± 118 (−2.8)	−298 ± 119 (−2.5)

Notes. Results are only presented for feature models with better than a 2.5σ result on the full domain (see Table 12).

Table C.5. Comparison of f_{NL} results for the hybrid polynomial and Fourier modes for a variety of non-separable and feature models.

	$f_{NL} \pm \Delta f_{NL}$ (σ)					
Model	Polynomial			Fourier ISW		
Const (see text)	14	± 44	(0.3)	10	± 44	(0.2)
EFT1 shape (5)	8	± 73	(0.1)	6	± 73	(0.1)
EFT2 shape (5)	19	± 57	(0.3)	13	± 57	(0.2)
DBI inflation (7)	12	± 69	(0.2)	−0.3	± 70	(0.0)
Ghost inflation (see text) ..	−24	± 88	(−0.3)	−48	± 89	(−0.5)
Flat model (13)	37	± 77	(0.5)	38	± 76	(0.5)
NBD (see text)	155	± 78	(2.0)	116	± 92	(1.3)
NBD1 flattened (14)	19	± 13	(1.4)	4	± 19	(0.2)
NBD2 squeezed (14)	0.25	± 0.45	(0.5)	−0.3	± 0.5	(−0.5)
NBD3 non-canonical (15) .	10	± 10	(1.0)	4	± 11	(0.3)
Vector model $L = 1$ (19) ..	−5	± 47	(−0.1)	−24	± 50	(−0.5)
Vector model $L = 2$ (19) ..	−0.4	± 2.8	(−0.1)	−1.0	± 3.2	(−0.3)
WarmS inflation (see text) .	1	± 33	(0.04)	−16	± 41	(−0.4)
Feature $k_c = 0.015$	−313	± 144	(2.1)	−264	± 161	(1.6)
Feature $k_c = 0.020$	−155	± 110	(−1.4)	−167	± 122	(−1.4)
Feature $k_c = 0.025$	106	± 93	(1.1)	110	± 98	(1.1)
Feature $k_c = 0.030$	56	± 89	(0.6)	78	± 96	(0.8)
Feature $k_c = 0.035$	22	± 82	(0.3)	15	± 84	(0.2)
Feature $k_c = 0.040$	−0.9	± 68	(−0.01)	4	± 69	(0.1)
Feature $k_c = 0.050$	−0.2	± 63	(0.00)	15	± 66	(0.2)
Feature $k_c = 0.080$	16	± 64	(0.2)	21	± 64	(0.3)

to the models discussed in Sect. 7.3, we have also included the constant model which is defined by $B^{const}(k_1, k_2, k_3) = 6A^2 f_{NL}^{const}/(k_1 k_2 k_3)^2$. The abbreviations EFT denotes the effective field theory single field shapes, NBD the non-Bunch-Davies (excited initial state) models, vector models are gauge field inflation with directional dependence, along with DBI, Ghost and Warm inflation, also described previously. The expression "cleaned" refers to removal of the predicted ISW-lensing signal and the measured point source signal, unless stated otherwise.

We note that there is good consistency between the different foreground-separation techniques for all models, whether equilateral, flattened, or squeezed as shown in Table C.3. For the non-scaling case, differences for the best-fit feature models were below $1/3\sigma$ confirming the interesting results discussed in Sect. 3.2.3, see Table C.4.

C.3. Comparison of f_{NL} results from ISW Fourier basis with hybrid polynomials

As described in Sect. 3.2.3, the modal bispectrum estimator can flexibly incorporate any suitable basis functions with which to expand the bispectrum and separably filter CMB maps (Fergusson et al. 2010a). For the *Planck* analysis, we have evolved two sets of basis functions to fulfil three basic criteria: first, to provide a complete basis for the bispectrum up to a given ℓ-resolution, secondly, to represent accurately primordial models of interest and, thirdly, to incorporate the CMB ISW-lensing signal, which with diffuse point sources provides a significant secondary signal which must be subtracted. The first basis functions are $n_{max} = 600$ polynomials (closely related to Legendre polynomials) which are supplemented with the Sachs-Wolfe local mode in order to represent more accurately the squeezed limit; enhanced orthogonality is preconditioned by choosing these from a larger set of polynomials. The second basis functions are $n_{max} = 300$ Fourier modes, augmented with the same

SW local mode, together with the the separable ISW-lensing modes. Both modal expansions proved useful, providing important validation and cross-checks, however, the twofold resolution improvement from the polynomials meant that most quantitative results employed this basis. This improved resolution was particularly important in probing flattened models on the edge of the tetrapyd.

In Fig. 7, we show a direct comparison between the modal reconstruction of the 3D bispectrum using the polynomial and Fourier mode expansions. The basic features of the two reconstructions are in good agreement, confirming a central feature which changes sign at low ℓ and a flattened signal beyond as discussed previously in Sect. 7.2. Notably the polynomial basis, with double the resolution, preserves the large-scale features observed in the Fourier basis.

In Table C.5, results from both basis expansions are shown for a variety of non-separable models. These demonstrate consistent results where the Fourier basis had sufficient resolution, as indicated by the ratio of the variance reflecting the Fisher ratio (i.e., above 90% correlation as indicated by the results in Appendix B). Note that we independently determine the estimator correlation between the exact solution and primordial decomposition and then at late times with the CMB decomposition; we use a polynomial basis as the overall benchmark here. This analysis also includes several feature models (phase $\phi = 0$) showing good agreement from the Fourier basis while the Fisher estimates remain accurate. Again, the hybrid Fourier basis degrades in accuracy towards $k = 0.02$ as it reaches its resolution limit, when the variance disparity rises towards 10%. With $n_{max} = 600$ modes and $\ell_{max} = 2000$, the polynomial basis retains a good correlation for all primordial feature models for $k > 0.01$. The accuracy and robustness of the feature model results have been verified using $\ell_{max} = 1500$ for the polynomial expansion, for example, obtaining 3.1σ with the SMICA map for the best-fit model ($K = 0.01875$, $\phi = 0$).

References

Achúcarro, A., Gong, J.-O., Hardeman, S., Palma, G. A., & Patil, S. P. 2011, J. Cosmol. Astropart. Phys., 1, 30

Ackerman, L., Carroll, S. M., & Wise, M. B. 2007, Phys. Rev. D, 75, 083502

Acquaviva, V., Bartolo, N., Matarrese, S., & Riotto, A. 2003, Nucl. Phys. B, 667, 119

Adshead, P., Hu, W., Dvorkin, C., & Peiris, H. V. 2011, Phys. Rev. D, 84, 3519

Adshead, P., Dvorkin, C., Hu, W., & Lim, E. A. 2012, Phys. Rev. D, 85, 3531

Agullo, I., & Parker, L. 2011, Gen. Relativ. Gravit., 43, 2541

Alishahiha, M., Silverstein, E., & Tong, D. 2004, Phys. Rev. D, 70, 3505

Allen, T. J., Grinstein, B., & Wise, M. B. 1987, Phys. Lett. B, 197, 66

Amendola, L., Catena, R., Masina, I., et al. 2011, J. Cosmol. Astropart. Phys., 7, 27

Antoine, J.-P., & Vandergheynst, P. 1998, J. Math. Phys., 39, 3987

Arkani-Hamed, N., Creminelli, P., Mukohyama, S., & Zaldarriaga, M. 2004, J. Cosmol. Astropart. Phys., 4, 1

Armendariz-Picon, C., Damour, T., & Mukhanov, V. 1999, Phys. Lett. B, 458, 209

Arroja, F., Mizuno, S., & Koyama, K. 2008, J. Cosmol. Astropart. Phys., 8, 15

Arroja, F., Mizuno, S., Koyama, K., & Tanaka, T. 2009, Phys. Rev. D, 80, 3527

Ashoorioon, A., & Shiu, G. 2011, J. Cosmol. Astropart. Phys., 3, 25

Assassi, V., Baumann, D., & Green, D. 2012, J. Cosmol. Astropart. Phys., 11, 47

Babich, D. 2005, Phys. Rev. D, 72, 3003

Babich, D., Creminelli, P., & Zaldarriaga, M. 2004, J. Cosmol. Astropart. Phys., 8, 9

Bardeen, J. M. 1980, Phys. Rev. D, 22, 1882

Barnaby, N., & Cline, J. M. 2006, Phys. Rev. D, 73, 106012

Barnaby, N., & Cline, J. M. 2008, J. Cosmol. Astropart. Phys., 6, 30

Barnaby, N., & Peloso, M. 2011, Phys. Rev. Lett., 106, 1301

Barnaby, N., Namba, R., & Peloso, M. 2011, J. Cosmol. Astropart. Phys., 4, 9

Barnaby, N., Namba, R., & Peloso, M. 2012a, Phys. Rev. D, 85, 3523

Barnaby, N., Pajer, E., & Peloso, M. 2012b, Phys. Rev. D, 85, 3525

Bartolo, N., & Riotto, A. 2009, J. Cosmol. Astropart. Phys., 3, 17

Bartolo, N., Matarrese, S., & Riotto, A. 2002, Phys. Rev. D, 65, 103505

Bartolo, N., Komatsu, E., Matarrese, S., & Riotto, A. 2004a, Phys. Rep., 402, 103

Bartolo, N., Matarrese, S., & Riotto, A. 2004b, J. Cosmol. Astropart. Phys., 1, 3

Bartolo, N., Matarrese, S., & Riotto, A. 2004c, Phys. Rev. Lett., 93, 231301

Bartolo, N., Matarrese, S., & Riotto, A. 2004d, Phys. Rev. D, 69, 043503

Bartolo, N., Matarrese, S., & Riotto, A. 2005, J. Cosmol. Astropart. Phys., 8, 10

Bartolo, N., Fasiello, M., Matarrese, S., & Riotto, A. 2010a, J. Cosmol. Astropart. Phys., 8, 8

Bartolo, N., Fasiello, M., Matarrese, S., & Riotto, A. 2010b, J. Cosmol. Astropart. Phys., 9, 35

Bartolo, N., Matarrese, S., & Riotto, A. 2010c, Adv. Astron., 2010

Bartolo, N., Matarrese, S., & Riotto, A. 2012, J. Cosmol. Astropart. Phys., 2, 17

Bartolo, N., Matarrese, S., Peloso, M., & Ricciardone, A. 2013, Phys. Rev. D, 87, 023504

Baumann, D., & Green, D. 2012, Phys. Rev. D, 85, 3520

Baumann, D., & McAllister, L. 2007, Phys. Rev. D, 75, 3508

Bean, R., Shandera, S. E., Tye, S.-H. H., & Xu, J. 2007, J. Cosmol. Astropart. Phys., 5, 4

Bean, R., Chen, X., Peiris, H., & Xu, J. 2008, Phys. Rev. D, 77, 023527

Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, ApJS, 208, 20

Berera, A. 1995, Phys. Rev. Lett., 75, 3218

Bernardeau, F., & Uzan, J.-P. 2002, Phys. Rev. D, 66, 103506

Bernardeau, F., Kofman, L., & Uzan, J.-P. 2004, Phys. Rev. D, 70, 083004

Bond, J. R., Frolov, A. V., Huang, Z., & Kofman, L. 2009, Phys. Rev. Lett., 103, 071301

Boubekeur, L., Creminelli, P., D'Amico, G., Noreña, J., & Vernizzi, F. 2009, J. Cosmol. Astropart. Phys., 8, 29

Bucher, M., & Zhu, Y. 1997, Phys. Rev. D, 55, 7415

Bucher, M., van Tent, B., & Carvalho, C. S. 2010, MNRAS, 407, 2193

Byrnes, C. T., & Choi, K.-Y. 2010, Adv. Astron., 2010 [arXiv:1002.3110]

Byrnes, C. T., Sasaki, M., & Wands, D. 2006, Phys. Rev. D, 74, 123519

Cardoso, J., Martin, M., Delabrouille, J., Betoule, M., & Patanchon, G. 2008, IEEE Select. Topics Signal Process., 2, 735

Cayon, L., Martínez-González, E., Argueso, F., Banday, A. J., & Gorski, K. M. 2003, MNRAS, 339, 1189

Challinor, A., & van Leeuwen, F. 2002, Phys. Rev. D, 65, 3001

Chambers, A., & Rajantie, A. 2008, Phys. Rev. Lett., 100, 1302

Chen, X. 2005a, J. High Energy Phys., 8, 45

Chen, X. 2005b, Phys. Rev. D, 71, 3506

Chen, X. 2005c, Phys. Rev. D, 72, 3518

Chen, X. 2010a, J. Cosmol. Astropart. Phys., 12, 3

Chen, X. 2010b, Adv. Astron. [arXiv:1002.1416]

Chen, X. 2011, Phys. Lett. B, 706, 111

Chen, X., & Wang, Y. 2010a, Phys. Rev. D, 81, 063511

Chen, X., & Wang, Y. 2010b, J. Cosmol. Astropart. Phys., 4, 27

Chen, X., Easther, R., & Lim, E. A. 2007a, J. Cosmol. Astropart. Phys., 6, 23

Chen, X., Huang, M.-X., Kachru, S., & Shiu, G. 2007b, J. Cosmol. Astropart. Phys., 1, 2

Chen, X., Easther, R., & Lim, E. A. 2008, J. Cosmol. Astropart. Phys., 4, 10

Chen, X., Hu, B., Huang, M.-X., Shiu, G., & Wang, Y. 2009, J. Cosmol. Astropart. Phys., 8, 8

Chen, X., Firouzjahi, H., Namjoo, M. H., & Sasaki, M. 2013, Europhys. Lett., 102, 59001

Cheung, C., Fitzpatrick, A. L., Kaplan, J., Senatore, L., & Creminelli, P. 2008, J. High Ener. Phys., 3, 14

Cicoli, M., Tasinato, G., Zavala, I., Burgess, C. P., & Quevedo, F. 2012, J. Cosmol. Astropart. Phys., 5, 39

Cooray, A., & Hu, W. 2000, ApJ, 534, 533

Creminelli, P., & Zaldarriaga, M. 2004a, J. Cosmol. Astropart. Phys., 10, 6

Creminelli, P., & Zaldarriaga, M. 2004b, Phys. Rev. D, 70, 083532

Creminelli, P., Nicolis, A., Senatore, L., Tegmark, M., & Zaldarriaga, M. 2006, J. Cosmol. Astropart. Phys., 5, 4

Creminelli, P., D'Amico, G., Musso, M., Noreña, J., & Trincherini, E. 2011a, J. Cosmol. Astropart. Phys., 2, 6

Creminelli, P., Pitrou, C., & Vernizzi, F. 2011b, J. Cosmol. Astropart. Phys., 11, 25

Cruz, M., Cayon, L., Martínez-González, E., Vielva, P., & Jin, J. 2007, ApJ, 655, 11

Curto, A., Macias-Perez, J. F., Martínez-González, E., et al. 2008, A&A, 486, 383

Curto, A., Martínez-González, E., & Barreiro, R. B. 2009a, ApJ, 706, 399

Curto, A., Martínez-González, E., Mukherjee, P., et al. 2009b, MNRAS, 393, 615

Curto, A., Martínez-González, E., & Barreiro, R. B. 2011a, MNRAS, 412, 1038

Curto, A., Martínez-González, E., Barreiro, R. B., & Hobson, M. P. 2011b, MNRAS, 417, 488

Curto, A., Martínez-González, E., & Barreiro, R. B. 2012, MNRAS, 426, 1361

Curto, A., Tucci, M., González-Nuevo, J., et al. 2013, MNRAS, 432, 728

De Troia, G., Ade, P. A. R., Bock, J. J., et al. 2007, ApJ, 670, L73

Delabrouille, J., Cardoso, J.-F., & Patanchon, G. 2003, MNRAS, 346, 1089

Delabrouille, J., Cardoso, J., Le Jeune, M., et al. 2009, A&A, 493, 835

Delabrouille, J., Betoule, M., Melin, J.-B., et al. 2013, A&A, 553, A96

Donzelli, S., Hansen, F. K., Liguori, M., & Maino, D. 2009, ApJ, 706, 1226

Donzelli, S., Hansen, F. K., Liguori, M., Marinucci, D., & Matarrese, S. 2012, ApJ, 755, 19

Ducout, A., Bouchet, F. R., Colombi, S., Pogosyan, D., & Prunet, S. 2013, MNRAS, 429, 2104

Dvali, G., Gruzinov, A., & Zaldarriaga, M. 2004a, Phys. Rev. D, 69, 083505

Dvali, G., Gruzinov, A., & Zaldarriaga, M. 2004b, Phys. Rev. D, 69, 023505

Elliston, J., Alabidi, L., Huston, I., Mulryne, D., & Tavakol, R. 2012, J. Cosmol. Astropart. Phys., 9, 1

Elsner, F., & Wandelt, B. D. 2009, ApJS, 184, 264

Elsner, F., & Wandelt, B. D. 2013, A&A, 549, A111

Endlich, S., Nicolis, A., & Wang, J. 2013, J. Cosmol. Astropart. Phys., 10, 11

Enqvist, K., & Sloth, M. S. 2002, Nucl. Phys. B, 626, 395

Enqvist, K., Jokinen, A., Mazumdar, A., Multamäki, T., & Väihkönen, A. 2005, Phys. Rev. Lett., 94, 161301

Eriksen, H. K., Novikov, D. I., Lilje, P. B., Banday, A. J., & Górski, K. M. 2004, ApJ, 612, 64

Eriksen, H. K., Dickinson, C., Lawrence, C. R., et al. 2006, ApJ, 641, 665

Eriksen, H. K., Jewell, J. B., Dickinson, C., et al. 2008, ApJ, 676, 10

Falk, T., Rangarajan, R., & Srednicki, M. 1993, ApJ, 403, L1

Faÿ, G., Guilloux, F., Betoule, M., et al. 2008, Phys. Rev. D, 78, 083013

Feeney, S. M., Johnson, M. C., Mortlock, D. J., & Peiris, H. V. 2011, Phys. Rev. D, 84, 043507

Fergusson, J. R., & Shellard, E. P. S. 2007, Phys. Rev. D, 76, 083523

Fergusson, J. R., Liguori, M., & Shellard, E. P. S. 2010a, Phys. Rev. D, 82, 023502

Fergusson, J. R., Regan, D. M., & Shellard, E. P. S. 2010b [arXiv:1012.6039]

Fergusson, J. R., Liguori, M., & Shellard, E. P. S. 2012, J. Cosmol. Astropart. Phys., 12, 32

Fernández-Cobos, R., Vielva, P., Barreiro, R. B., & Martínez-González, E. 2012, MNRAS, 420, 2162

Ferreira, P. G., & Magueijo, J. 1997, Phys. Rev. D, 56, 4578

Flauger, R., & Pajer, E. 2011, J. Cosmol. Astropart. Phys., 1, 17

Flauger, R., McAllister, L., Pajer, E., Westphal, A., & Xu, G. 2010, J. Cosmol. Astropart. Phys., 6, 9

Gangui, A., & Martin, J. 2000a, Phys. Rev. D, 62, 3004

Gangui, A., & Martin, J. 2000b, MNRAS, 313, 323

Gangui, A., Lucchin, F., Matarrese, S., & Mollerach, S. 1994, ApJ, 430, 447

Geller, D., & Mayeli, A. 2009a, Math. Z., 262, 895

Geller, D., & Mayeli, A. 2009b, Math. Z., 263, 235

Giovi, F., Baccigalupi, C., & Perrotta, F. 2005, Phys. Rev. D, 71, 3009

Goldberg, D. M., & Spergel, D. N. 1999, Phys. Rev. D, 59, 3002

González-Nuevo, J., Toffolatti, L., & Argüeso, F. 2005, ApJ, 621, 1

Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759

Green, D., Horn, B., Senatore, L., & Silverstein, E. 2009, Phys. Rev. D, 80, 3533

Hall, L. M. H., & Peiris, H. V. 2008, J. Cosmol. Astropart. Phys., 1, 27

Hamimeche, S., & Lewis, A. 2008, Phys. Rev. D, 77, 3013

Hanson, D., & Lewis, A. 2009, Phys. Rev. D, 80, 3004

Hanson, D., Rocha, G., & Górski, K. 2009a, MNRAS, 400, 2169

Hanson, D., Smith, K. M., Challinor, A., & Liguori, M. 2009b, Phys. Rev. D, 80, 083004

Heavens, A. F. 1998, MNRAS, 299, 805

Hikage, C., & Matsubara, T. 2012, MNRAS, 425, 2187

Hikage, C., Komatsu, E., & Matsubara, T. 2006, ApJ, 653, 11

Hikage, C., Matsubara, T., Coles, P., et al. 2008, MNRAS, 389, 1439

Hodges, H. M., Blumenthal, G. R., Kofman, L. A., & Primack, J. R. 1990, Nucl. Phys. B, 335, 197

Holman, R., & Tolley, A. J. 2008, J. Cosmol. Astropart. Phys., 5, 1

Huang, M. X., & Shiu, G. 2006, Phys. Rev. D, 74, 1301

Huang, Z., & Vernizzi, F. 2013, Phys. Rev. Lett., 110, 1303

Izumi, K., Mizuno, S., & Koyama, K. 2012, Phys. Rev. D, 85, 3521

Junk, V., & Komatsu, E. 2012, Phys. Rev. D, 85, 123524

Kehagias, A., & Riotto, A. 2012, Nucl. Phys. B, 864, 492

Keihänen, E., Keskitalo, R., Kurki-Suonio, H., Poutanen, T., & Sirviö, A. 2010, A&A, 510, A57

Khatri, R., & Wandelt, B. D. 2009, Phys. Rev. D, 79, 3501

Khatri, R., & Wandelt, B. D. 2010, Phys. Rev. D, 81, 3518

Kim, J., Rotti, A., & Komatsu, E. 2013, J. Cosmol. Astropart. Phys., 4, 21

Kofman, L. 2003 [arXiv:astro-ph/0303614]

Kofman, L., Blumenthal, G. R., Hodges, H., & Primack, J. R. 1991, in Large-scale Structures and Peculiar Motions in the Universe, eds. D. W. Latham, & L. A. N. da Costa, ASP Conf. Ser., 15, 339

Kogo, N., & Komatsu, E. 2006, Phys. Rev. D, 73, 3007

Kolb, E. W., Riotto, A., & Vallinotto, A. 2006, Phys. Rev. D, 73, 3522

Komatsu, E. 2010, Class. Quantum Grav., 27, 124010

Komatsu, E., & Spergel, D. N. 2001, Phys. Rev. D, 63, 3002

Komatsu, E., Kogut, A., Nolta, M. R., et al. 2003, ApJS, 148, 119

Komatsu, E., Spergel, D. N., & Wandelt, B. D. 2005, ApJ, 634, 14

Komatsu, E., Dunkley, J., Nolta, M. R., et al. 2009, ApJS, 180, 330

Komatsu, E., Smith, K. M., Dunkley, J., et al. 2011, ApJS, 192, 18

Kosowsky, A., & Kahniashvili, T. 2011, Phys. Rev. Lett., 106, 191301

Lacasa, F., & Aghanim, N. 2014, A&A, 569, A51

Lacasa, F., Aghanim, N., Kunz, M., & Frommert, M. 2012, MNRAS, 421, 1982

Lan, X., & Marinucci, D. 2008, Electr. J. Stat., 2, 332

Langlois, D., Renaux-Petel, S., Steer, D. A., & Tanaka, T. 2008a, Phys. Rev. Lett., 101, 061301

Langlois, D., Renaux-Petel, S., Steer, D. A., & Tanaka, T. 2008b, Phys. Rev. D, 78, 063523

Leach, S. M., Cardoso, J., Baccigalupi, C., et al. 2008, A&A, 491, 597

Lehners, J.-L. 2010, Adv. Astron., 2010, 67

Lewis, A. 2011, J. Cosmol. Astropart. Phys., 10, 26

Lewis, A. 2012, J. Cosmol. Astropart. Phys., 6, 23

Lewis, A., & Challinor, A. 2006, Phys. Rep., 429, 1

Lewis, A., Challinor, A., & Hanson, D. 2011, J. Cosmol. Astropart. Phys., 3, 18

Lidsey, J. E., & Huston, I. 2007, J. Cosmol. Astropart. Phys., 7, 2

Lidsey, J. E., Liddle, A. R., Kolb, E. W., et al. 1997, Rev. Mod. Phys., 69, 373

Liguori, M., Sefusatti, E., Fergusson, J. R., & Shellard, E. P. S. 2010, Adv. Astron., 2010, 73

Linde, A., & Mukhanov, V. 1997, Phys. Rev. D, 56, 535

Lyth, D. H. 2005, J. Cosmol. Astropart. Phys., 11, 6

Lyth, D. H., & Riotto, A. A. 1999, Phys. Rep., 314, 1

Lyth, D. H., & Riotto, A. 2006, Phys. Rev. Lett., 97, 121301

Lyth, D. H., & Rodriguez, Y. 2005, Phys. Rev. Lett., 95, 121302

Lyth, D. H. & Wands, D. 2002, Phys. Lett. B, 524, 5

Lyth, D. H., Ungarelli, C., & Wands, D. 2003, Phys. Rev. D, 67, 023503

Maldacena, J. 2003, J. High Ener. Phys., 5, 13

Malik, K. A., & Lyth, D. H. 2006, J. Cosmol. Astropart. Phys., 9, 8

Malik, K. A., & Wands, D. 2009, Phys. Rep., 475, 1

Mangilli, A., & Verde, L. 2009, Phys. Rev. D, 80, 3007

Mangilli, A., Wandelt, B., Elsner, F., & Liguori, M. 2013, A&A, 555, A82

Martínez-González, E., Gallegos, J. E., Argüeso, F., Cayon, L., & Sanz, J. L. 2002, MNRAS, 336, 22

Matsubara, T. 2010, Phys. Rev. D, 81, 3505

McEwen, J. D., Vielva, P., Hobson, M. P., Martínez-González, E., & Lasenby, A. N. 2007a, MNRAS, 376, 1211

McEwen, J. D., Vielva, P., Wiaux, Y., et al. 2007b, J. Fourier Anal. Appl., 13, 495

Meerburg, P. D., van der Schaar, J. P., & Stefano Corasaniti, P. 2009, J. Cosmol. Astropart. Phys., 5, 18

Mitra, S., Rocha, G., Górski, K. M., et al. 2011, ApJS, 193, 5

Modest, H. I., Räth, C., Banday, A. J., et al. 2013, MNRAS, 428, 551

Mollerach, S. 1990, Phys. Rev. D, 42, 313

Moroi, T., & Takahashi, T. 2001, Phys. Lett. B, 522, 215

Moss, I. G., & Xiong, C. 2007, J. Cosmol. Astropart. Phys., 4, 7

Munshi, D., & Heavens, A. 2010, MNRAS, 401, 2406

Munshi, D., Heavens, A., Cooray, A., et al. 2011, MNRAS, 412, 1993

Natoli, P., de Troia, G., Hikage, C., et al. 2010, MNRAS, 408, 1658

Nitta, D., Komatsu, E., Bartolo, N., Matarrese, S., & Riotto, A. 2009, J. Cosmol. Astropart. Phys., 5, 14

Okamoto, T., & Hu, W. 2002, Phys. Rev. D, 66, 3008

Pearson, R., Lewis, A., & Regan, D. 2012, J. Cosmol. Astropart. Phys., 3, 11

Peebles, P. J. E. 1997, ApJ, 483, L1

Peiris, H. V., Komatsu, E., Verde, L., et al. 2003, ApJS, 148, 213

Peiris, H., Baumann, D., Friedman, B., & Cooray, A. 2007, Phys. Rev. D, 76, 103517

Peiris, H. V., Easther, R., & Flauger, R. 2013, J. Cosmol. Astropart. Phys., 9, 18

Pettinari, G. W., Fidler, C., Crittenden, R., Koyama, K., & Wands, D. 2013, J. Cosmol. Astropart. Phys., 4, 3

Pietrobon, D., Balbi, A., & Marinucci, D. 2006, Phys. Rev. D, 74, 043524

Pietrobon, D., Cabella, P., Balbi, A., de Gasperis, G., & Vittorio, N. 2009, MNRAS, 396, 1682

Pietrobon, D., Cabella, P., Balbi, A., et al. 2010, MNRAS, 402, L34

Planck Collaboration 2013, The Explanatory Supplement to the Planck 2013 results, http://www.sciops.esa.int/wikiSI/planckpla/index.php?title=Main_Page (ESA)

Planck Collaboration I. 2014, A&A, 571, A1

Planck Collaboration II. 2014, A&A, 571, A2

Planck Collaboration III. 2014, A&A, 571, A3

Planck Collaboration IV. 2014, A&A, 571, A4

Planck Collaboration V. 2014, A&A, 571, A5

Planck Collaboration VI. 2014, A&A, 571, A6

Planck Collaboration VII. 2014, A&A, 571, A7

Planck Collaboration VIII. 2014, A&A, 571, A8

Planck Collaboration IX. 2014, A&A, 571, A9

Planck Collaboration X. 2014, A&A, 571, A10

Planck Collaboration XI. 2014, A&A, 571, A11

Planck Collaboration XII. 2014, A&A, 571, A12

Planck Collaboration XIII. 2014, A&A, 571, A13

Planck Collaboration XIV. 2014, A&A, 571, A14

Planck Collaboration XV. 2014, A&A, 571, A15

Planck Collaboration XVI. 2014, A&A, 571, A16

Planck Collaboration XVII. 2014, A&A, 571, A17

Planck Collaboration XVIII. 2014, A&A, 571, A18

Planck Collaboration XIX. 2014, A&A, 571, A19

Planck Collaboration XX. 2014, A&A, 571, A20

Planck Collaboration XXI. 2014, A&A, 571, A21

Planck Collaboration XXII. 2014, A&A, 571, A22

Planck Collaboration XXIII. 2014, A&A, 571, A23

Planck Collaboration XXIV. 2014, A&A, 571, A24

Planck Collaboration XXV. 2014, A&A, 571, A25

Planck Collaboration XXVI. 2014, A&A, 571, A26

Planck Collaboration XXVII. 2014, A&A, 571, A27

Planck Collaboration XXVIII. 2014, A&A, 571, A28

Planck Collaboration XXIX. 2014, A&A, 571, A29

Planck Collaboration XXX. 2014, A&A, 571, A30

Planck Collaboration XXXI. 2014, A&A, 571, A31

Ratra, B. 1992, ApJ, 391, L1

Rees, M. J., & Sciama, D. W. 1968, Nature, 217, 511

Regan, D., Mukherjee, P., & Seery, D. 2013, Phys. Rev. D, 88, 043512

Reinecke, M., Dolag, K., Hell, R., Bartelmann, M., & Enßlin, T. A. 2006, A&A, 445, 373

Renaux-Petel, S. 2009, J. Cosmol. Astropart. Phys., 10, 12

Renaux-Petel, S., Mizuno, S., & Koyama, K. 2011, J. Cosmol. Astropart. Phys., 11, 42

Rigopoulos, G. I., Shellard, E. P. S., & van Tent, B. J. W. 2006, Phys. Rev. D, 73, 083522

Rigopoulos, G. I., Shellard, E. P. S., & van Tent, B. J. W. 2007, Phys. Rev. D, 76, 083512

Rudjord, O., Hansen, F. K., Lan, X., et al. 2009, ApJ, 701, 369

Sachs, R. K., & Wolfe, A. M. 1967, ApJ, 147, 73

Salem, M. P. 2005, Phys. Rev. D, 72, 123516

Salopek, D. S., & Bond, J. R. 1990, Phys. Rev. D, 42, 3936

Sasaki, M., Väliviita, J., & Wands, D. 2006, Phys. Rev. D, 74, 3003

Scodeller, S., Rudjord, O., Hansen, F. K., et al. 2011, ApJ, 733, 121

Sefusatti, E., Fergusson, J. R., Chen, X., & Shellard, E. P. S. 2012, J. Cosmol. Astropart. Phys., 8, 33

Senatore, L., & Zaldarriaga, M. 2011, J. Cosmol. Astropart. Phys., 1, 3

Senatore, L., Tassev, S., & Zaldarriaga, M. 2009, J. Cosmol. Astropart. Phys., 9, 38

Senatore, L., Smith, K. M., & Zaldarriaga, M. 2010, J. Cosmol. Astropart. Phys., 1, 28

Serra, P., & Cooray, A. 2008, Phys. Rev. D, 77, 107305

Shiraishi, M., Nitta, D., Yokoyama, S., & Ichiki, K. 2012, J. Cosmol. Astropart. Phys., 3, 41

Shiraishi, M., Komatsu, E., Peloso, M., & Barnaby, N. 2013, J. Cosmol. Astropart. Phys., 5, 2

Shiu, G., & Xu, J. 2011, Phys. Rev. D, 84, 3509

Silverstein, E., & Tong, D. 2004, Phys. Rev. D, 70, 3505

Smidt, J., Amblard, A., Byrnes, C. T., et al. 2010, Phys. Rev. D, 81, 3007

Smith, K. M., & Zaldarriaga, M. 2011, MNRAS, 417, 2

Smith, K. M., Senatore, L., & Zaldarriaga, M. 2009, J. Cosmol. Astropart. Phys., 9, 6

Smith, K. M., Loverde, M., & Zaldarriaga, M. 2011, Phys. Rev. Lett., 107, 191301

Smith, T. L., & Kamionkowski, M. 2012, Phys. Rev. D, 86, 063009

Starck, J.-L., Moudden, Y., Abrial, P., & Nguyen, M. 2006, A&A, 446, 1191

Starck, J.-L., Murtagh, F., & Fadili, J. M. 2010, Sparse image and signal processing, wavelets, curvelets, morphological diversity (Cambridge: Cambridge University Press)

Su, S.-C., Lim, E. A., & Shellard, E. P. S. 2012 [arXiv:1212.6968]

Sugiyama, N. S. 2012, J. Cosmol. Astropart. Phys., 5, 32

Sugiyama, N. S., Komatsu, E., & Futamase, T. 2011, Phys. Rev. Lett., 106, 1301

Suyama, T., & Yamaguchi, M. 2008, Phys. Rev. D, 77, 023505

Suyama, T., Takahashi, T., Yamaguchi, M., & Yokoyama, S. 2010, J. Cosmol. Astropart. Phys., 12, 30

Toffolatti, L., Argueso Gomez, F., de Zotti, G., et al. 1998, MNRAS, 297, 117

Tolley, A. J., & Wyman, M. 2010, Phys. Rev. D, 81, 3502

Tzavara, E., & van Tent, B. 2011, J. Cosmol. Astropart. Phys., 6, 26

Verde, L., & Spergel, D. N. 2002, Phys. Rev. D, 65, 3007

Verde, L., Wang, L., Heavens, A. F., & Kamionkowski, M. 2000, MNRAS, 313, 141

Verde, L., Jimenez, R., Alvarez-Gaume, L., Heavens, A. F., & Matarrese, S. 2013, J. Cosmol. Astropart. Phys., 6, 23

Vernizzi, F., & Wands, D. 2006, J. Cosmol. Astropart. Phys., 5, 19

Wang, L., & Kamionkowski, M. 2000, Phys. Rev. D, 61, 063504

Weinberg, S. 2008, Phys. Rev. D, 77, 123541

Yadav, A. P. S., & Wandelt, B. D. 2008, Phys. Rev. Lett., 100, 181301

Yadav, A. P. S., & Wandelt, B. D. 2010, Adv. Astron., 2010, 71

Yadav, A. P. S., Komatsu, E., & Wandelt, B. D. 2007, ApJ, 664, 680

Yadav, A. P. S., Komatsu, E., Wandelt, B. D., et al. 2008, ApJ, 678, 578

Zaldarriaga, M. 2004, Phys. Rev. D, 69, 043508

Zeldovich, Y. B., & Sunyaev, R. A. 1969, Ap&SS, 4, 301

[1] APC, AstroParticule et Cosmologie, Université Paris Diderot, CNRS/IN2P3, CEA/lrfu, Observatoire de Paris, Sorbonne Paris Cité, 10 rue Alice Domon et Léonie Duquet, 75205 Paris Cedex 13, France

[2] Aalto University Metsähovi Radio Observatory, Metsähovintie 114, 02540 Kylmälä, Finland

[3] African Institute for Mathematical Sciences, 6-8 Melrose Road, 7945 Muizenberg, Cape Town, South Africa

[4] Agenzia Spaziale Italiana Science Data Center, via del Politecnico snc, 00133 Roma, Italy

[5] Agenzia Spaziale Italiana, viale Liegi 26, 00998 Roma, Italy

[6] Astrophysics Group, Cavendish Laboratory, University of Cambridge, J J Thomson Avenue, Cambridge CB3 0HE, UK

[7] Astrophysics & Cosmology Research Unit, School of Mathematics, Statistics & Computer Science, University of KwaZulu-Natal, Westville Campus, Private Bag X54001, 4000 Durban, South Africa

[8] CITA, University of Toronto, 60 St. George St., Toronto, ON M5S 3H8, Canada

[9] CNRS, IRAP, 9 Av. colonel Roche, BP 44346, 31028 Toulouse Cedex 4, France

[10] California Institute of Technology, Pasadena, California, USA

[11] Centre for Theoretical Cosmology, DAMTP, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, UK

[12] Centro de Estudios de Física del Cosmos de Aragón (CEFCA), Plaza San Juan, 1, planta 2, 44001 Teruel, Spain

[13] Computational Cosmology Center, Lawrence Berkeley National Laboratory, Berkeley, California, USA

[14] Consejo Superior de Investigaciones Científicas (CSIC), 2800 Madrid, Spain

[15] DSM/Irfu/SPP, CEA-Saclay, 91191 Gif-sur-Yvette Cedex, France

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

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

[18] Departamento de Física Fundamental, Facultad de Ciencias, Universidad de Salamanca, 37008 Salamanca, Spain

[19] Departamento de Física, Universidad de Oviedo, Avda. Calvo Sotelo s/n, 33007 Oviedo, Spain

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

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

[22] Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, California, USA

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

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

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

[26] Department of Physics and Astronomy, University of Sussex, Brighton BN1 9QH, UK

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

[28] Department of Physics, Gustaf Hällströmin katu 2a, University of Helsinki, 00014 Helsinki, Finland

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

[30] Department of Physics, University of California, Berkeley, California, USA

[31] Department of Physics, University of California, One Shields Avenue, Davis, California, USA

[32] Department of Physics, University of California, Santa Barbara, California, USA

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

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

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

[36] Dipartimento di Fisica, Università La Sapienza, P.le A. Moro 2, 00185 Roma, Italy

[37] Dipartimento di Fisica, Università degli Studi di Milano, via Celoria, 16, 20133 Milano, Italy

[38] Dipartimento di Fisica, Università degli Studi di Trieste, via A. Valerio 2, 34127 Trieste, Italy

[39] Dipartimento di Fisica, Università di Roma Tor Vergata, via della Ricerca Scientifica 1, 00133 Roma, Italy

[40] Dipartimento di Matematica, Università di Roma Tor Vergata, via della Ricerca Scientifica 1, 00133 Roma, Italy

[41] Discovery Center, Niels Bohr Institute, Blegdamsvej 17, 2100 Copenhagen, Denmark

[42] Dpto. Astrofísica, Universidad de La Laguna (ULL), 38206 La Laguna, Tenerife, Spain

[43] European Space Agency, ESAC, Planck Science Office, Camino bajo del Castillo s/n, Urbanización Villafranca del Castillo, 28691 Villanueva de la Cañada, Madrid, Spain

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

[45] Helsinki Institute of Physics, Gustaf Hällströmin katu 2, University of Helsinki, 00014 Helsinki, Finland

[46] INAF – Osservatorio Astronomico di Padova, Vicolo dell'Osservatorio 5, 35122 Padova, Italy

[47] INAF – Osservatorio Astronomico di Roma, via di Frascati 33, 00040 Monte Porzio Catone, Italy

[48] INAF – Osservatorio Astronomico di Trieste, via G.B. Tiepolo 11, 34143 Trieste, Italy

[49] INAF Istituto di Radioastronomia, via P. Gobetti 101, 40129 Bologna, Italy

[50] INAF/IASF Bologna, via Gobetti 101, 40129 Bologna, Italy

[51] INAF/IASF Milano, via E. Bassini 15, 20133 Milano, Italy

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

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

[54] IPAG: Institut de Planétologie et d'Astrophysique de Grenoble, Université Joseph Fourier, Grenoble 1/CNRS-INSU, UMR 5274, 38041 Grenoble, France

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

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

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

[58] Institut Néel, CNRS, Université Joseph Fourier Grenoble I, 25 rue des Martyrs, 38042 Grenoble, France

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

[60] Institut d'Astrophysique Spatiale, CNRS (UMR 8617) Université Paris-Sud 11, Bâtiment 121, 91405 Orsay, France

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

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

[63] Institute of Astronomy and Astrophysics, Academia Sinica, 106 Taipei, Taiwan

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

[65] Institute of Theoretical Astrophysics, University of Oslo, Blindern, 0315 Oslo, Norway

[66] Instituto de Astrofísica de Canarias, C/Vía Láctea s/n, 38200 La Laguna, Tenerife, Spain

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

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

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

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

[71] LAL, Université Paris-Sud, CNRS/IN2P3, 91383 Orsay, France

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

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

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

[75] Laboratoire de Physique Subatomique et de Cosmologie, Université Joseph Fourier Grenoble I, CNRS/IN2P3, Institut National Polytechnique de Grenoble, 53 rue des Martyrs, 38026 Grenoble Cedex, France

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

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

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

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

[80] MilliLab, VTT Technical Research Centre of Finland, Tietotie 3, 02044 Espoo, Finland

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

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

[83] Observational Cosmology, Mail Stop 367-17, California Institute of Technology, Pasadena CA, 91125, USA

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

[85] SB-ITP-LPPC, EPFL, 1015 Lausanne, Switzerland

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

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

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

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

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

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

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

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

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

[95] UPMC Univ. Paris 06, UMR7095, 98bis Boulevard Arago, 75014 Paris, France

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

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

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