

General CMB and primordial bispectrum estimation: Mode expansion, map making, and measures of F_{NL}

J. R. Fergusson, M. Liguori, and E. P. S. Shellard

Centre for Theoretical Cosmology, Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, United Kingdom

(Received 3 May 2010; published 1 July 2010)

We present a detailed implementation of two parallel bispectrum estimation methods which can be applied to general nonseparable primordial and cosmic microwave background (CMB) bispectra. The method exploits bispectrum mode decompositions on the tetrahedral domain of allowed wave number or multipole values, using both separable basis functions and related orthonormal modes. We provide concrete examples of such modes constructed from symmetrized tetrahedral polynomials, demonstrating the rapid convergence of expansions describing nonseparable bispectra. We use these modes to create rapid and robust pipelines for generating simulated CMB maps of high resolution ($l > 2000$) given an arbitrary primordial power spectrum and bispectrum or an arbitrary late-time CMB angular power spectrum and bispectrum. By extracting coefficients for the same separable basis functions from an observational map, we are able to present an efficient f_{NL} estimator for a given theoretical model with a nonseparable bispectrum. The estimator has two manifestations, comparing the theoretical and observed coefficients at either primordial or late times, thus encompassing a wider range of models, including secondary anisotropies and lensing as well as active models, such as cosmic strings. We provide examples and validation of both f_{NL} estimation methods by direct comparison with simulations in a WMAP-realistic context. In addition, we demonstrate how the full primordial and CMB bispectrum can be extracted from observational maps using these mode expansions, irrespective of the theoretical model under study. We also propose a universal definition of the bispectrum parameter F_{NL} , so that the integrated bispectrum on the observational domain can be more consistently compared between theoretical models. We obtain WMAP5 estimates of f_{NL} for the equilateral model from both our primordial and late-time estimators which are consistent with each other, as well as results already published in the literature. These general bispectrum estimation methods should prove useful for non-Gaussianity analysis with the Planck satellite data, as well as in other contexts.

DOI: 10.1103/PhysRevD.82.023502

PACS numbers: 98.80.Es, 98.80.Cq

I. INTRODUCTION

Standard inflationary scenarios predict the Universe to be close to flat with primordial curvature perturbations which are nearly scale invariant and Gaussian. All these predictions are in very good accord with cosmic microwave background (CMB) and large-scale structure measurements, such as those provided by WMAP and SDSS. Despite this remarkable agreement, present observations are not able to completely rule out alternatives to inflation, nor to effectively discriminate among the vast number of different inflationary models that have been proposed. However, almost all such quantitative comparisons derive from inferred measurements of the primordial two-point correlator or power spectrum $P(k)$ from $\langle \zeta \zeta \rangle$, where ζ is the curvature perturbation. If we wish to subject inflation to more stringent tests and to distinguish between competing models then perhaps the best prospects are offered by studying non-Gaussianity, that is, the higher order correlators beyond the power spectrum. The three-point correlator of the CMB or bispectrum $B_{l_1 l_2 l_3}$ is a projection on the sky of the evolved primordial bispectrum $B(k_1, k_2, k_3)$ arising from $\langle \zeta \zeta \zeta \rangle$, consisting of contributions from triangle con-

figurations with side-lengths given by the wave numbers k_1, k_2, k_3 . The bispectrum has attracted most attention in the literature to date and its study is usually simplified to the characterization of a single nonlinearity parameter f_{NL} , which schematically is given by the ratio $f_{\text{NL}} \approx B(k, k, k)/P(k)^2$.

Standard inflation, that is, single-field slow-roll inflation, predicts a very small bispectrum with $f_{\text{NL}} \sim 0.01$ [1,2], possessing a characteristic scale-invariant local shape. (This local shape is dominated by squeezed triangle configurations, that is, those for which one side is much smaller than the others, e.g. $k_1 \ll k_2, k_3$.) In fact, such a low signal would be undetectable even by an ideal noiseless CMB experiment, because it is below the level of NG contamination expected from secondary anisotropies $f_{\text{NL}} \approx \mathcal{O}(1)$. However, measurement of a significantly larger primordial $f_{\text{NL}} \gtrsim 1$ would have profound consequences because it would signal the need for new physics during inflation or even a paradigm shift away from it. Present measurements of this local f_{NL} are equivocal with the WMAP team reporting [3]

$$f_{\text{NL}} = 51 \pm 60 \quad (95\%) \quad (1)$$

and with other teams obtaining higher [4] (WMAP3) or equivalent values [5,6], while with improved WMAP5 noise analysis a lower value was found $f_{\text{NL}} = 38 \pm 42$, but at a similar 2σ significance [7]. The Planck satellite experiment is expected to markedly improve precision measurements with $\Delta f_{\text{NL}} = 5$ or better [8].

Further motivation for the study of the bispectrum comes from the prospect of distinguishing alternative more complex models of inflation which can produce non-Gaussianity with potentially observable amplitudes $f_{\text{NL}} \gtrsim 1$, but also in a variety of different bispectrum shapes, that is, with the non-Gaussian signal peaked for different triangle configurations of wavevectors. To date only special separable bispectrum shapes have been constrained by CMB data, that is, those that can be expressed (schematically) in the form $B(k_1, k_2, k_3) = X(k_1)Y(k_2)Z(k_3)$, or else can be accurately approximated in this manner. All CMB analysis, such as those quoted above for the local shape (1), exploits this separability to reduce the dimensionality of the required integrations and summations to bring them to a tractable form. The separable approach reduces the problem from one of $\mathcal{O}(l_{\text{max}}^5)$ operations to a manageable $\mathcal{O}(l_{\text{max}}^3)$ [9]. Other examples of meaningful constraints on separable bispectrum shapes using WMAP5 data include those for the equilateral shape [3] and another shape “orthogonal” to both equilateral and local [10]. Despite these three shapes being a good approximation to non-Gaussianity from a number of classes of inflation models, they are not exhaustive in their coverage of known primordial models [11], nor other types of late-time non-Gaussianity, such as that from cosmic strings [12,13]; they cannot be expected to be, given the functional degrees of freedom available. Bringing observations to bear on this much broader class of cosmological models, therefore, is the primary motivation for this paper.

In a previous paper [14], we described a general approach to the estimation of nonseparable CMB bispectra. The method has developed out of the first direct calculations of the reduced CMB bispectrum $b_{l_1 l_2 l_3}$ which surveyed a wide variety of nonseparable primordial models, revealing smooth coherent patterns of acoustic peaks on the tetrahedral domain of allowed multipole values. Since the $b_{l_1 l_2 l_3}$ could be well represented using a limited number of bins, we could similarly decompose them into orthogonal mode functions which converged in relatively short mode expansions [11]. Here, we describe the detailed implementation of these methods in a comprehensive dual approach to estimating bispectrum parameters which is illustrated in Fig. 1. We present concrete examples of separable basis functions \mathcal{Q}_n (symmetrized tetrahedral polynomials) and corresponding orthonormal modes \mathcal{R}_n on the domain of allowed wave numbers k_1, k_2, k_3 ; these are then deployed within a more general mode expansion methodology. In the first primordial implementation, we decompose an arbitrary nonseparable shape S using separable basis functions with coefficients α_n . This expansion can be used for a fast calculation of the full CMB bispectrum $B_{l_1 l_2 l_3}$ (Sec. III), as well as leading to a robust method for generating simulated maps from a given power spectrum $P(k)$ and bispectrum $B(k_1, k_2, k_3)$ (Sec. IV). Our main emphasis here, however, is on a primordial estimator for f_{NL} which is achieved by a confrontation between theory, represented by the α_n coefficients, and a set of observational coefficients β_n found by extracting the same modes from the observed CMB map (Sec. III). Examples of simulated maps and recovery of the input f_{NL} are given in Sec. V in a WMAP-realistic context.

In the second and parallel late-time implementation (see Fig. 1), we assume the theoretical CMB bispectrum $B_{l_1 l_2 l_3}$ is calculated already from the primordial shape [11] or

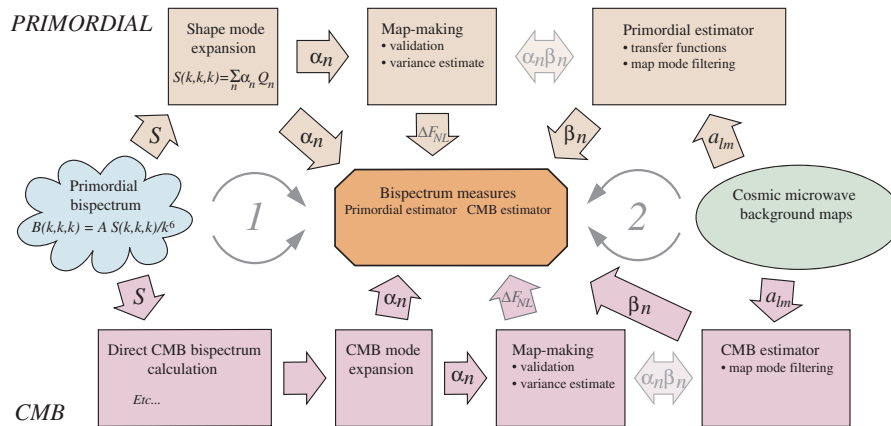


FIG. 1 (color online). Flow chart for the two general estimator methodologies described and implemented in this article using complete separable mode expansions. Note the overall redundancy which assists estimator validation and the independence of the extraction of expansion coefficients from theory α_n (cycle 1) and data β_n (cycle 2). Explanations for the schematic equations can be found in the main text.

because it is a late-time effect ranging from secondary anisotropies through to fluctuations induced by cosmic strings. A separable mode expansion of $B_{l_1 l_2 l_3}$ allows for a simpler and more direct approach to f_{NL} estimation, as well as simulated map generation, in a wider variety of scenarios. Here, as well as primordial models we consider the antithetical example of cosmic strings. These two estimator methods are complementary with each having distinct advantages depending on the properties and generation mechanism of the non-Gaussianity under investigation. They provide independent validation in situations where both are applicable.

It remains to point out recent and related developments, especially those by colleagues in Planck Working Group 4 (Non-Gaussianity). To date most primordial shapes have been assumed to be scale invariant, but in Ref. [15] some deviations from the local shape were considered in developing a more general approach. Spherical Mexican wavelets, using a limited number of scales, were employed in Refs. [5,16] to estimate f_{NL} for the local shape with WMAP3 data, providing a constraint consistent with (1). Similar work has been achieved for needlets with corresponding constraints [6], again essentially tailoring the method to the local template using local shape map simulations. Another approach to a late-time estimator has also exploited the smoothness of the reduced CMB bispectrum by using a limited number of multipole bins [17]. The method was tested for the local shape using map simulations, and emphasized Planck forecasts investigating the pattern of acoustic peaks in the local model. We shall discuss here how these late-time approaches—whether wavelets, bins or other alternatives—fall within the general mode expansion methodology outlined previously [14] and can be applied, in principle, to explore nonseparable primordial models beyond local non-Gaussianity. We point out in the implementation presented here, however, that direct estimation of the CMB bispectrum can be achieved without reference to the calculated bispectrum for a particular model and without relying on corresponding CMB map simulations.

II. THE CMB BISPECTRUM AND f_{NL} ESTIMATION

A. Relation between primordial and CMB bispectra

In this section we will review some basic definitions and mathematical formulas that will be used throughout the rest of the paper. Our work will be concerned with the analysis of the three-point function induced by a NG primordial gravitational potential $\Phi(\mathbf{k})$ in the CMB temperature fluctuation field. Temperature anisotropies are represented using the a_{lm} coefficients of a spherical harmonic decomposition of the cosmic microwave sky,

$$\frac{\Delta T}{T}(\hat{\mathbf{n}}) = \sum_{lm} a_{lm} Y_{lm}(\hat{\mathbf{n}}).$$

The primordial potential Φ is imprinted on the CMB multipoles a_{lm} by a convolution with transfer functions $\Delta_l(k)$ representing the linear perturbation evolution, through the integral

$$a_{lm} = 4\pi(-i)^l \int \frac{d^3 k}{(2\pi)^3} \Delta_l(k) \Phi(\mathbf{k}) Y_{lm}(\hat{\mathbf{k}}). \quad (2)$$

The CMB bispectrum is the three-point correlator of the a_{lm} , so substituting we obtain

$$B_{m_1 m_2 m_3}^{l_1 l_2 l_3} = \langle a_{l_1 m_1} a_{l_2 m_2} a_{l_3 m_3} \rangle \quad (3)$$

$$= (4\pi)^3 (-i)^{l_1 + l_2 + l_3} \int \frac{d^3 k_1}{(2\pi)^3} \frac{d^3 k_2}{(2\pi)^3} \times \frac{d^3 k_3}{(2\pi)^3} \Delta_{l_1}(k_1) \Delta_{l_2}(k_2) \Delta_{l_3}(k_3) \times \quad (4)$$

$$\langle \Phi(\mathbf{k}_1) \Phi(\mathbf{k}_2) \Phi(\mathbf{k}_3) \rangle Y_{l_1 m_1}(\hat{\mathbf{k}}_1) Y_{l_2 m_2}(\hat{\mathbf{k}}_2) Y_{l_3 m_3}(\hat{\mathbf{k}}_3), \quad (5)$$

where $k_1 = |\mathbf{k}_1|$, $k_2 = |\mathbf{k}_2|$ and $k_3 = |\mathbf{k}_3|$. Here, we define the primordial bispectrum as

$$\langle \Phi(\mathbf{k}_1) \Phi(\mathbf{k}_2) \Phi(\mathbf{k}_3) \rangle = (2\pi)^3 B_{\Phi}(k_1, k_2, k_3) \times \delta(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3), \quad (6)$$

where the delta function enforces the triangle condition, that is, the constraint imposed by translational invariance that wave vectors in Fourier space must close to form a triangle, $\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3 = 0$. We replace the delta function in (6) with its exponential integral form, substitute this into Eq. (3) and integrate out the angular parts of the three \mathbf{k}_i integrals in the usual manner to yield

$$B_{m_1 m_2 m_3}^{l_1 l_2 l_3} = \left(\frac{2}{\pi}\right)^3 \int x^2 dx \int dk_1 dk_2 dk_3 (k_1 k_2 k_3)^2 \times B_{\Phi}(k_1, k_2, k_3) \Delta_{l_1}(k_1) \Delta_{l_2}(k_2) \times \Delta_{l_3}(k_3) j_{l_1}(k_1 x) j_{l_2}(k_2 x) j_{l_3}(k_3 x) \times \int d\Omega_x Y_{l_1 m_1}(\hat{\mathbf{x}}) Y_{l_2 m_2}(\hat{\mathbf{x}}) Y_{l_3 m_3}(\hat{\mathbf{x}}). \quad (7)$$

The last integral over the angular part of x is known as the Gaunt integral which can be expressed in terms of Wigner-3j symbols as

$$\mathcal{G}_{m_1 m_2 m_3}^{l_1 l_2 l_3} \equiv \int d\Omega_x Y_{l_1 m_1}(\hat{\mathbf{x}}) Y_{l_2 m_2}(\hat{\mathbf{x}}) Y_{l_3 m_3}(\hat{\mathbf{x}}) = \sqrt{\frac{(2l_1 + 1)(2l_2 + 1)(2l_3 + 1)}{4\pi}} \times \begin{pmatrix} l_1 & l_2 & l_3 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} l_1 & l_2 & l_3 \\ m_1 & m_2 & m_3 \end{pmatrix}. \quad (8)$$

Given that most theories we shall consider are assumed to be isotropic, it is usual to work with the angle-averaged bispectrum,

$$B_{l_1 l_2 l_3} = \sum_{m_i} \binom{l_1 \quad l_2 \quad l_3}{m_1 \quad m_2 \quad m_3} \langle a_{l_1 m_1} a_{l_2 m_2} a_{l_3 m_3} \rangle. \quad (9)$$

or the even more convenient reduced bispectrum which removes the geometric factors associated with the Gaunt integral,

$$B_{m_1 m_2 m_3}^{l_1 l_2 l_3} = \mathcal{G}_{m_1 m_2 m_3}^{l_1 l_2 l_3} b_{l_1 l_2 l_3}. \quad (10)$$

The reduced bispectrum from (3) then takes the much simpler form

$$b_{l_1 l_2 l_3} = \left(\frac{2}{\pi}\right)^3 \int x^2 dx \int dk_1 dk_2 dk_3 (k_1 k_2 k_3)^2 B_\Phi(k_1, k_2, k_3) \times \Delta_{l_1}(k_1) \Delta_{l_2}(k_2) \Delta_{l_3}(k_3) j_{l_1}(k_1 x) j_{l_2}(k_2 x) j_{l_3}(k_3 x). \quad (11)$$

Here, it is important to note that the Gaunt integral in (10) encodes several constraints on the angle-averaged bispectrum $B_{l_1 l_2 l_3}$ which are no longer transparent in the reduced bispectrum $b_{l_1 l_2 l_3}$. These are, first, that the sum of the three multipoles l_i must be even and, secondly, that the l_i 's satisfy the triangle condition, analogously to the wave numbers k_i . For wave numbers, the triangle condition is enforced through the x -integral over the three spherical Bessel functions $j_l(k_i x)$ which evaluates to zero if the k_i 's cannot form a triangle, whereas in multipole space it is enforced by the angular integration $d\Omega_x$ over the spherical harmonics $Y_{l_i m_i}$ in (8). Appreciating the origin of these constraints is important when we later consider the separability of the reduced bispectrum expression (7).

B. Separable primordial shapes and CMB bispectrum solutions

Given that the primordial power spectrum is very nearly scale invariant, it is expected that the bispectrum will behave similarly. In order to bring the bispectrum to a scale-invariant form we have to appropriately eliminate a k^6 scaling which naturally arises in (6). This is usually achieved by multiplying through by the factor $(k_1 k_2 k_3)^2$ appearing in (11) and defining a primordial shape function as

$$S(k_1, k_2, k_3) \equiv \frac{1}{N} (k_1 k_2 k_3)^2 B_\Phi(k_1, k_2, k_3), \quad (12)$$

where N is a normalization factor which is often taken such that for equal k_i the shape function has unit value $S(k, k, k) = 1$. (This normalization is also used for f_{NL} , but it only strictly applies for scale invariance and, in any case, leads to inconsistent comparisons between different models, as we shall discuss in Sec. IV.) We thus characterize scale-invariant models in terms of an overall amplitude, parametrized by f_{NL} , and their transverse shape, described by $S(k_1, k_2, k_3)$ on a triangular slice with $k_1 + k_2 + k_3 = \text{const}$ [18]. This leaves a two-dimensional space on which it is most elegant to use the two independent variables $\tilde{\alpha}, \tilde{\beta}$

[14,19]

$$\tilde{\alpha} = (k_2 - k_3)/\tilde{k},$$

$$\tilde{\beta} = (\tilde{k} - k_1)/\tilde{k},$$

$$\text{where } \tilde{k} = \frac{1}{2}(k_1 + k_2 + k_3) = \text{const}, \quad (13)$$

with the following domains $0 \leq \tilde{\beta} \leq 1$ and $-(1 - \tilde{\beta}) \leq \tilde{\alpha} \leq 1 - \tilde{\beta}$. For scale-dependent models with a nontrivial variation in \tilde{k} , the full three-dimensional dependence on the k_i must be retained. In terms of the shape function (12), the reduced bispectrum (11) can be rewritten as

$$b_{l_1 l_2 l_3} = \frac{1}{N} \left(\frac{2}{\pi}\right)^3 \int x^2 dx \int dk_1 dk_2 dk_3 S(k_1 k_2 k_3) \Delta_{l_1}(k_1) \times \Delta_{l_2}(k_2) \Delta_{l_3}(k_3) j_{l_1}(k_1 x) j_{l_2}(k_2 x) j_{l_3}(k_3 x). \quad (14)$$

The simplest possible shape function is the constant model

$$S(k_1, k_2, k_3) = 1, \quad (15)$$

for which a large-angle analytic solution for the reduced bispectrum was presented in Ref. [11],

$$b_{l_1 l_2 l_3}^{\text{const}} = \frac{\Delta_\Phi^2}{27N} \frac{1}{(2l_1 + 1)(2l_2 + 1)(2l_3 + 1)} \times \left[\frac{1}{l_1 + l_2 + l_3 + 3} + \frac{1}{l_1 + l_2 + l_3} \right], \quad (l \ll 200). \quad (16)$$

Here, we take the Sachs-Wolfe approximation that $\Delta_l(k) = \frac{1}{3} j_l((\tau_o - \tau_{\text{dec}})k)$ for $l \ll 200$ and exploit the manifest separability of the expression (14) to perform the one-dimensional k_i integrations individually. The more general constant solution does not have an analytic solution for $l \geq 200$, for the reason that the transfer functions cannot be expressed in a simple form, but it can be evaluated numerically from the expression

$$b_{l_1 l_2 l_3}^{\text{const}} = \frac{\Delta_\Phi^2}{N} \int x^2 dx I_{l_1}(x) I_{l_2}(x) I_{l_3}(x),$$

$$\text{where } I_l(x) = \frac{2}{\pi} \int dk \Delta_l(k) j_l(kx). \quad (17)$$

The large-angle solution (16) is an important benchmark with which to compare the shape of late-time CMB bispectra from other models $b_{l_1 l_2 l_3}$ (note the l^{-4} scaling) and, additionally, it has some further recent physical motivation [20].

The most studied scale-invariant shape function is the local model,

$$\begin{aligned}
S(k_1, k_2, k_3) &= \frac{1}{3} \left(\frac{k_1^2}{k_2 k_3} + \frac{k_2^2}{k_1 k_3} + \frac{k_3^2}{k_1 k_2} \right) \\
&\approx \frac{(k_1 k_2 k_3)^2}{3 \Delta_{\Phi}^2} [P(k_1)P(k_2) + P(k_2)P(k_3) \\
&\quad + P(k_3)P(k_1)], \tag{18}
\end{aligned}$$

where, in the second line, we also allow for power spectra which are nearly scale invariant, defined by $\langle \Phi(\mathbf{k})\Phi^*(\mathbf{k}') \rangle = (2\pi)^3 P(k)\delta(\mathbf{k} - \mathbf{k}')$ with $P(k) \sim k^{-3}$. Using the Sachs-Wolfe approximation again, this has the corresponding large-angle analytic solutions

$$\begin{aligned}
b_{l_1 l_2 l_3}^{\text{local}} &= \frac{2\Delta_{\Phi}^2}{27\pi^2} \left(\frac{1}{l_1(l_1+1)l_2(l_2+1)} + \frac{1}{l_2(l_2+1)l_3(l_3+1)} \right. \\
&\quad \left. + \frac{1}{l_3(l_3+1)l_1(l_1+1)} \right) \tag{19}
\end{aligned}$$

Here, we see that the divergences for the squeezed triangles ($k_1 \ll k_2, k_3, \dots$) in the primordial shape (18) are also reflected in $b_{l_1 l_2 l_3}^{\text{local}}$, making it a much less useful for relative comparison than the constant model (16). It is straightforward, in principle, to calculate the full bispectrum from the separable expressions arising from (18),

$$b_{l_1 l_2 l_3}^{\text{local}} = \int x^2 dx [\alpha_{l_1}(x)\beta_{l_2}(x)\beta_{l_3}(x) + (2 \text{ perms})], \tag{20}$$

where the separated integrals analogous to (17) become

$$\begin{aligned}
\alpha_{l_1}(x) &= \frac{2}{\pi} \int dk k^2 \Delta_l(k) j_l(kx), \\
\beta_{l_1}(x) &= \frac{2}{\pi} \int dk k^2 P(k) \Delta_l(k) j_l(kx). \tag{21}
\end{aligned}$$

We note that these highly oscillatory integrals must be evaluated numerically with considerable care. The separable equilateral shape has also received a great deal of attention with [18]

$$\begin{aligned}
S(k_1, k_2, k_3) &= \frac{(k_1 + k_2 - k_3)(k_2 + k_3 - k_1)(k_3 + k_1 - k_2)}{k_1 k_2 k_3} \\
&= -2 - \left[\frac{k_1^2}{k_2 k_3} + (2 \text{ perms}) \right] \\
&\quad + \left[\frac{k_1}{k_2} + (5 \text{ perms}) \right]. \tag{22}
\end{aligned}$$

This is a much more regular shape than local (18) with the signal dominated by equilateral triangle configurations $k_1 \approx k_2 \approx k_3$ (the apparent divergence of the local shape in the second term cancels against the third). There is no simple large-angle analytic solution known for the equilateral model, unlike (17) and (20). In order to calculate the full equilateral bispectrum we evaluate the simplified expression

$$\begin{aligned}
b_{l_1 l_2 l_3}^{\text{equil}} &= \int x^2 dx \{ 2\delta_{l_1} \delta_{l_2} \delta_{l_3} + [\alpha_{l_1} \beta_{l_2} \beta_{l_3} + (2 \text{ perms})] \\
&\quad + [\beta_{l_1} \gamma_{l_2} \delta_{l_3} + (5 \text{ perms})] \}, \tag{23}
\end{aligned}$$

where α_l, β_l are given in (21) and γ_l, δ_l are defined by

$$\begin{aligned}
\gamma_l(x) &= \frac{2}{\pi} \int dk k^2 P(k)^{1/3} \Delta_l(k) j_l(kx), \\
\delta_l(x) &= \frac{2}{\pi} \int dk k^2 P(k)^{2/3} \Delta_l(k) j_l(kx). \tag{24}
\end{aligned}$$

The equilateral shape is not derived directly from a physical model, but was chosen phenomenologically as a good separable approximation to specific models including the nonlocal part of Maldacena's original shape [1], as well as noncanonical cases such as higher derivative models [21] and DBI inflation [22] (for a review of single-field inflation shapes, see e.g. Ref. [23]). These shapes are, in general, nonseparable from the perspective of the integral (14). Here, we give a specific shape example for a model with higher derivative operators (which is also identical to DBI inflation):

$$\begin{aligned}
S(k_1, k_2, k_3) &= \frac{1}{k_1 k_2 k_3 (k_1 + k_2 + k_3)^2} \\
&\quad \times \left(\sum_i k_i^5 + \sum_{i \neq j} (2k_i^4 k_j - 3k_i^3 k_j^2) \right. \\
&\quad \left. + \sum_{i \neq j \neq l} (k_i^3 k_j k_l - 4k_i^2 k_j^2 k_l) \right). \tag{25}
\end{aligned}$$

Not only is the equilateral shape (22) an excellent approximation to (25), a full Fisher matrix analysis of the respective CMB bispectra has shown they are 99% correlated out to $l_{\text{max}} \leq 2000$ [11]. However, a simple separable approximation is not necessarily available for arbitrary primordial shapes, nor is a particular separable representation necessarily convenient from a calculational perspective (as we shall discuss in Sec. V for the equilateral case above). In Ref. [11], we reviewed models currently proposed in the literature showing that families of CMB bispectra arising from nonseparable shapes, such as feature and flattened models, are largely independent of the separable models currently constrained observationally (see also discussion of a ‘‘cosine’’ shape correlator in Ref. [18]). The independence of two shapes S and S' can be calculated from the integral [11]

$$F_{\epsilon}(S, S') = \int_{\mathcal{V}_k} S(k_1, k_2, k_3) S'(k_1, k_2, k_3) \omega_{\epsilon}(k_1, k_2, k_3) d\mathcal{V}_k, \tag{26}$$

where we choose the weight to be

$$w(k_1, k_2, k_3) = \frac{1}{k_1 + k_2 + k_3}, \tag{27}$$

reflecting the scaling we see in the CMB correlator we

meet in the next section. The shape correlator is then defined by

$$\tilde{C}(S, S') = \frac{F(S, S')}{\sqrt{F(S, S)F(S', S')}}. \quad (28)$$

By way of further illustration of the need to move beyond simple separable primordial shape functions, we present the late-time CMB bispectrum predicted analytically for cosmic strings [13]

$$b_{l_1 l_2 l_3}^{\text{string}} = \frac{A}{(\xi l_1 l_2 l_3)^2} \left[(l_3^2 - l_1^2 - l_2^2) \left(\frac{L}{2l_3} + \frac{l_3}{50L} \right) \times \sqrt{\frac{l_*}{500}} \text{erf}(0.3\xi l_3) + 2 \text{ perms} \right], \quad (l \leq 2000), \quad (29)$$

where $l_{\min} = \min(l_1, l_2, l_3)$, $l_* = \min(500, l_{\min})$, $\xi = \min(1/500, 1/l_{\min})$ and

$$L = \xi \sqrt{\frac{1}{2}(l_1^2 l_2^2 + l_2^2 l_3^2 + l_3^2 l_1^2) - \frac{1}{4}(l_1^4 + l_2^4 + l_3^4)}. \quad (30)$$

Here, $A \sim (8\pi G\mu)^3$ is a model-dependent amplitude with $G\mu = \mu/m_{\text{Pl}}^2$ measuring the string tension μ relative to the Planck scale. The cutoffs around $l \approx 500$ in (29) are associated with the string correlation length at decoupling (perturbations with $l \gtrsim 500$ can only be causally seeded after last scattering). (For the original small angle solution valid for $l \gg 2000$, see Refs. [12,13].) Here, the nonseparable nature and very different scaling of the string CMB bispectrum are clear from a comparison with (19). Moreover, given the late-time origin of this signal from string metric perturbations, the modulating effect of acoustic peaks from the transfer functions is absent.

C. Estimators for f_{NL} and related correlators

The main purpose of this non-Gaussian CMB analysis is to measure the CMB bispectrum induced by non-Gaussianities in the primordial gravitational potential, the link being given by Eq. (11). Unfortunately, the bispectrum signal is too weak to measure individual multipoles directly, so to compare theory with observation we must use an estimator which sums over the available multipoles. An estimator can be thought of as performing a least squares fit of the bispectrum predicted by theory $\langle a_{l_1 m_1} a_{l_2 m_2} a_{l_3 m_3} \rangle$ to the bispectrum actually obtained from observations $a_{l_1 m_1}^{\text{obs}} a_{l_2 m_2}^{\text{obs}} a_{l_3 m_3}^{\text{obs}}$. Ignoring sky cuts and inhomogeneous noise, the estimator is weighted with the expected signal variance from C_l and written in the simple form

$$\begin{aligned} \mathcal{E} &= \frac{1}{N} \sum_{l_i m_i} \frac{\langle a_{l_1 m_1} a_{l_2 m_2} a_{l_3 m_3} \rangle a_{l_1 m_1}^{\text{obs}} a_{l_2 m_2}^{\text{obs}} a_{l_3 m_3}^{\text{obs}}}{C_{l_1} C_{l_2} C_{l_3}} \\ &= \frac{1}{N} \sum_{l_i m_i} \mathcal{G}_{m_1 m_2 m_3}^{l_1 l_2 l_3} b_{l_1 l_2 l_3} \frac{a_{l_1 m_1}^{\text{obs}} a_{l_2 m_2}^{\text{obs}} a_{l_3 m_3}^{\text{obs}}}{C_{l_1} C_{l_2} C_{l_3}}. \end{aligned} \quad (31)$$

where we have used (10) and the Gaunt integral is given in (8) and N is the usual normalization factor,

$$N = \sum_{l_i} \frac{B_{l_1 l_2 l_3} B_{l_1 l_2 l_3}}{C_{l_1} C_{l_2} C_{l_3}}. \quad (32)$$

We note from the second line of (31) that, for a given theoretical model, we need only calculate the reduced bispectrum $b_{l_1 l_2 l_3}$ rather than the much more challenging full bispectrum, $\langle a_{l_1 m_1} a_{l_2 m_2} a_{l_3 m_3} \rangle$.

The above estimator has been shown to be optimal [24] for general bispectra in the limit where the non-Gaussianity is small and the observed map is free of instrument noise and foreground contamination. Of course, this is an idealized case and we need to consider taking into account the effect of sky cuts and inhomogeneous noise, which was considered in some detail in Refs. [25,26]. In the more general case the optimal estimator takes the form:

$$\begin{aligned} \mathcal{E} &= \frac{1}{N} \sum_{l_i m_i} \begin{pmatrix} l_1 & l_2 & l_3 \\ m_1 & m_2 & m_3 \end{pmatrix} B_{l_1 l_2 l_3} [(C^{-1} a^{\text{obs}})_{l_1 m_1} \\ &\times (C^{-1} a^{\text{obs}})_{l_2 m_2} (C^{-1} a^{\text{obs}})_{l_3 m_3} \\ &+ (C_{l_1 m_1, l_2 m_2}^{-1} (C^{-1} a^{\text{obs}})_{l_3 m_3})], \end{aligned} \quad (33)$$

where the covariance matrix C is now nondiagonal due to mode-mode coupling introduced by the mask and anisotropic noise. Moreover, due to the breaking of isotropy, an additional term linear in the a_{lm} has now to be added in order to maintain the optimality of the estimator [24]. In the ideal case one can easily see that the linear term is proportional to a monopole, while the covariance matrix is diagonal and equal to $1/C_l$, thus reproducing the initial formula (31).

In this paper we will follow the approach of [27] and approximate the estimator (33) as

$$\mathcal{E} = \frac{1}{\tilde{N}} \sum_{l_i m_i} \frac{\mathcal{G}_{m_1 m_2 m_3}^{l_1 l_2 l_3} \tilde{b}_{l_1 l_2 l_3}}{\tilde{C}_{l_1} \tilde{C}_{l_2} \tilde{C}_{l_3}} (a_{l_1 m_1}^{\text{obs}} a_{l_2 m_2}^{\text{obs}} - 6C_{l_1 m_1, l_2 m_2}^{\text{sim}}) a_{l_3 m_3}^{\text{obs}}, \quad (34)$$

where the tilde denotes modification to include experimental effects. The normalization becomes

$$\tilde{N} = f_{\text{sky}} \sum_{l_i} \frac{\tilde{B}_{l_1 l_2 l_3}^2}{\tilde{C}_{l_1} \tilde{C}_{l_2} \tilde{C}_{l_3}}, \quad (35)$$

with the C_l 's and $b_{l_1 l_2 l_3}$ now incorporating beam and noise effects through

$$\tilde{C}_l = b_l^2 C_l + N_l \quad \text{and} \quad \tilde{b}_{l_1 l_2 l_3} = b_{l_1} b_{l_2} b_{l_3} b_{l_1 l_2 l_3}. \quad (36)$$

Here, b_l is the beam transfer function, N_l the noise power spectrum, f_{sky} the fraction of the sky remaining after application of the mask and $C_{l_1 m_1, l_2 m_2}^{\text{sim}}$ is the covariance matrix calculated from Gaussian simulations. In what follows, it will be clear from the context whether beams, noise

and masks are being incorporated in the analysis, so for simplicity we shall continue with the original estimator notation (31).

The estimator (31) also naturally defines a correlator for testing whether two competing bispectra could be differentiated by an ideal experiment. Replacing the observed bispectrum with one calculated from a competing theory we have,

$$\mathcal{C}(B, B') = \frac{1}{N} \sum_{l_i} \frac{B_{l_1 l_2 l_3} B'_{l_1 l_2 l_3}}{C_{l_1} C_{l_2} C_{l_3}}, \quad (37)$$

where now the normalization N is defined as follows,

$$N = \sqrt{\sum_{l_i} \frac{B_{l_1 l_2 l_3}^2}{C_{l_1} C_{l_2} C_{l_3}}} \sqrt{\sum_{l_i} \frac{B'_{l_1 l_2 l_3}{}^2}{C_{l_1} C_{l_2} C_{l_3}}}. \quad (38)$$

While this late-time correlator is the best measure of whether two CMB bispectra are truly independent, it requires a full calculation of the CMB bispectrum which is time consuming in general. In [11] we determined that for the majority of models the shape correlator (28) introduced earlier is sufficient to determine independence.

An inspection of Eqs. (31) and (33) shows that a brute force numerical implementation of the optimal estimator above would take $\mathcal{O}(l_{\max}^5)$ operations. This means an implementation is not feasible for the angular resolutions achieved by present and forthcoming data sets (e.g. in the signal dominated regime we have $l_{\max} \lesssim 500$ for WMAP and $l_{\max} \lesssim 2000$ for Planck). However, as initially shown in Ref. [9], if a specific theoretical bispectrum can be written in separable form as $B(k_1, k_2, k_3) = X(k_1)Y(k_2)Z(k_3)$ then the computational cost of the algorithm can be reduced to $\mathcal{O}(l_{\max}^3)$ operations, making the estimation tractable even at very high angular resolutions. This establishes the fact that separability is a crucial property for realistic data analysis, even though it is not generic for well-motivated inflationary and other models. As we have seen, the usual solution adopted has been to approximate the primordial nonseparable shape under study using a separable form that is highly correlated with the original. This kind of approach requires a case-by-case analysis of all nonseparable bispectra arising from different models and an educated ‘‘guess’’ of a good separable approximation, the close correlation of which must be verified numerically before moving on to the real analysis. Besides being impractical, this can also prove to be extremely difficult in specific cases. The aim of this work is then to find a completely general mathematical framework to ‘‘separate’’ shapes, both primordial and late-time, and thus build a general pipeline for f_{NL} estimation and simulation of non-Gaussian CMB maps, that can be applied to any shape of interest.

III. BISPECTRUM MODE DECOMPOSITION

Our goal is to represent arbitrary nonseparable primordial bispectra $B(k_1, k_2, k_3)$ or CMB bispectra $b_{l_1 l_2 l_3}$ on their respective wavenumber or multipole domains using a rapidly convergent mode expansion [14]. Moreover, we need to achieve this in a separable manner, making tractable the three-dimensional integrals required for bispectrum estimation (14) by breaking them down into products of one-dimensional integrals. In particular, this means that we wish to expand an arbitrary nonseparable primordial shape function as

$$S(k_1, k_2, k_3) = \sum_p \sum_r \sum_s \alpha_{prs} q_p(k_1) q_r(k_2) q_s(k_3), \quad (39)$$

where the q_p are appropriate basis mode functions which are convergent and complete, that is, they span the space of all functions on the bispectrum wave number (or multipole) domain. In what follows below, we present one pathway for efficiently achieving this objective in stages. First, we create examples of one-dimensional mode functions $q_p(k_1)$ in the k_1 -direction which are orthogonal and well-behaved over the full wave number (or multipole) domain. We then construct three-dimensional products of these mode functions $q_p(k_1) q_r(k_2) q_s(k_3) \rightarrow \mathcal{Q}_n$ creating a complete basis for all possible bispectra on the given domain. Finally, by orthonormalizing these product basis functions $\mathcal{Q}_n \rightarrow \mathcal{R}_n$, we obtain a rapid and convenient method for calculating the relevant expansion coefficients α_{prs} in (39). The subsequent discussion and implementation of general primordial and CMB bispectrum estimators, as well as mapmaking methods, is then built around these mode functions q_p , \mathcal{Q}_n , and \mathcal{R}_n . Here, we use bounded symmetric polynomials as a concrete and working implementation of this methodology, and we defer discussion about other possible basis mode functions which have been investigated to the end of the section.

A. Tetrahedral domain and weight functions

In Fourier space, the primordial bispectrum $B(k_1, k_2, k_3)$ is defined when the three wave vectors $\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3$ close to form a triangle $\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3 = 0$. Since each such triangle is uniquely defined by the lengths of its sides $k_1 = |\mathbf{k}_1|, k_2 = |\mathbf{k}_2|, k_3 = |\mathbf{k}_3|$, we only require wave numbers in the bispectrum argument. In terms of these three wave numbers, the triangle condition restricts the allowed combinations into a tetrahedral region defined by

$$\begin{aligned} k_1 &\leq k_2 + k_3 \quad \text{for } k_1 \geq k_2, k_3, \quad \text{or } k_2 \leq k_1 + k_3 \\ \text{for } k_2 &\geq k_1, k_3, \quad \text{or } k_3 \leq k_1 + k_2 \quad \text{for } k_3 \geq k_1, k_2. \end{aligned} \quad (40)$$

This region forms a regular tetrahedron if we impose the restriction that $k_1 + k_2 + k_3 < 2k_{\max}$, however, it is more natural to extend the domain out to values given by a

maximum wave number in each direction $k_1, k_2, k_3 \leq k_{\max}$. This extension is motivated by issues both of separability and observation. The allowed domain $\mathcal{V}_{\mathcal{T}}$ is then a hexahedron formed by the intersection of a tetrahedron and a cube. It can be obtained from a regular tetrahedron (two-thirds of the total volume) by gluing on top a regular triangular pyramid constructed from the corner of the cube (as illustrated in Fig. 2). For brevity, let us denote this asymmetric triangular bipyramid as a *tetrapyd*, from the merger of a tetrahedron and a pyramid. Of course, bispectrum symmetries are such that it is only necessary to use one sixth of this domain, but aesthetics and intuition are helped by keeping the full domain while making a restriction to symmetrized functions.

We will frequently need to integrate functions $f(k_1, k_2, k_3)$ over the tetrapyd domain (40), which for brevity we will denote as $\mathcal{V}_{\mathcal{T}}$ with the integration given explicitly by

$$\begin{aligned} \mathcal{T}[f] &\equiv \int_{\mathcal{V}_{\mathcal{T}}} f(k_1, k_2, k_3) w(k_1, k_2, k_3) d\mathcal{V}_{\mathcal{T}} \\ &= K^3 \left\{ \int_0^{1/2} \int_y^{1-y} \int_{x-y}^{x+y} F W dz dx dy \right. \\ &\quad + \int_0^{1/2} \int_x^{1-x} \int_{y-x}^{x+y} F W dz dy dx \\ &\quad + \int_{1/2}^1 \int_{1-x}^x \int_{x-y}^1 F W dz dy dx \\ &\quad \left. + \int_{1/2}^1 \int_{1-y}^y \int_{y-x}^1 F W dz dx dy \right\}. \end{aligned} \quad (41)$$

where $K = k_{\max}$, $w(k_1, k_2, k_3)$ is an appropriate weight function, and we have made the transformation $x = k_1/K$, $y = k_2/K$, $z = k_3/K$ with $F(x, y, z) = f(Kx, Ky, Kz)$ and $W(x, y, z) = w(Kx, Ky, Kz)$. For inte-




FIG. 2 (color online). Tetrahedral domain (“tetrapyd”) for allowed multipole values l for the CMB bispectrum $b_{l_1 l_2 l_3}$ or, with wave numbers k for the primordial bispectrum $B(k_1, k_2, k_3)$. The regular tetrahedral region defined up to the equilateral slice $l_1 + l_2 + l_3 \leq 2l_{\max} \equiv 2L$ (shaded brown) contains two-thirds of the overall volume. The rest of the domain is given by the regular triangular pyramid on top which fills the volume to the corner of the encompassing cube defined by $l_1, l_2, l_3 \leq L$. An origami tetrapyd is also shown (right) with folding instructions.

grals over the product of two functions f and g we can define their inner product $\langle f, g \rangle \equiv \mathcal{T}[fg]$, essentially defining a Hilbert space of possible shape functions in the domain (40). The total volume of the tetrapyd domain is given by $\mathcal{T}[1] = K^3/2$. Initially, for the sake of simplicity, on the primordial wavenumber domain we will restrict attention to unit side-length $K = 1$ and weight $w = 1$.

We note that it is important to incorporate a weight function for a variety of reasons. For example, the primordial shape function $S(k_1, k_2, k_3)$ can be shown to possess a nearly linear scaling with respect to the CMB bispectrum estimator; on the multipole domain $w_{l_1 l_2 l_3}$ is nontrivial. A fairly close correspondence between the two can be obtained using $w(k_1, k_2, k_3) \approx 1/(k_1 + k_2 + k_3)$ [11] which explains its choice in the shape correlator (28). The choice of weight function also affects mode expansion convergence and for certain shapes it may be convenient to eliminate dependencies by rescaling with a separable function. For the shapes we consider here however, this is not necessary.

When analysing the CMB bispectrum it is particularly important to extend the tetrahedral domain to include multipoles in the top pyramidal region shown in Fig. 2. In principle, this pyramid contains 33% of the triple $l_1 l_2 l_3$ combinations available in the observational data, e.g. with Planck out to $l_1, l_2, l_3 \leq 2000$. The tetrapyd domain for the reduced bispectrum $b_{l_1 l_2 l_3}$ becomes the discrete $\{l_1, l_2, l_3\}$ combinations satisfying

$$\begin{aligned} l_1, l_2, l_3 &\leq l_{\max}, & l_1, l_2, l_3 &\in \mathbb{N}, \\ l_1 &\leq l_2 + l_3 & \text{for } l_1 \geq l_2, l_3, & + \text{cyclic perms}, \\ l_1 + l_2 + l_3 &= 2n, & n &\in \mathbb{N}. \end{aligned} \quad (42)$$

In Fig. 3 we illustrate contrasting bispectra on this domain for the equilateral and local models (here with $l_{\max} = 2000$).

In multipole space, we will be primarily dealing with a summation over all possible $\{l_1, l_2, l_3\}$ combinations in the estimator (31) or the closely related correlator (37). The appropriate weight function in the sum is then

$$w_{l_1 l_2 l_3} = \frac{1}{4\pi} (2l_1 + 1)(2l_2 + 1)(2l_3 + 1) \begin{pmatrix} l_1 & l_2 & l_3 \\ 0 & 0 & 0 \end{pmatrix}, \quad (43)$$

where we note that the third condition in (42) arises as a selection rule from the Wigner-3j symbol. Despite the discrete origin of the function $w_{l_1 l_2 l_3}$, like the reduced bispectrum $b_{l_1 l_2 l_3}$, it varies smoothly. It is particularly uniform on cross-sectional slices $l_1 + l_2 + l_3 = 2L$, except for a finite rise very close to the boundaries. While the Wigner-3j symbols are easily calculable (especially in the $m_i = 0$ case when performed in advance for a lookup table), it is more convenient to work in the continuum limit $w_{l_1 l_2 l_3} \rightarrow w(l_1, l_2, l_3)$ when considering domains with large l_{\max} . To achieve this we take the exact expression in terms




FIG. 3 (color online). The reduced CMB bispectra for the equilateral model (left) and the local model (right) plotted on the tetrahedral region shown in Fig. 42 (from [11]). Several density contours are illustrated (light blue positive and magenta negative) and $b_{l_1 l_2 l_3}$ is normalized by scaling relative to the constant Sachs-Wolfe solution (16) $b_{l_1 l_2 l_3}^{\text{model}}/b_{l_1 l_2 l_3}^{\text{const}}$. Note the acoustic peaks induced by the transfer functions and the center weighting for the equilateral model, contrasting with the corner-weighting for the local case [14].

of factorials (for even combinations with $l_1 + l_2 + l_3 = 2l$, $l \in \mathbb{N}$),

$$\binom{l_1 \ l_2 \ l_3}{0 \ 0 \ 0} = (-1)^l \sqrt{\frac{(2l - 2l_1)!(2l - 2l_2)!(2l - 2l_3)!}{(2l + 1)!}} \times \frac{l!}{(l - l_1)!(l - l_2)!(l - l_3)!}, \quad (44)$$

and then we substitute the Gosper approximation for all these factorials, that is,

$$l! \approx \sqrt{\left(2l + \frac{1}{3}\right) \pi} l^l e^{-l}. \quad (45)$$

The discrete multipole weight function (43) then reduces to a straightforward continuum version

$$w(l_1, l_2, l_3) = \frac{1}{2\pi^2} \frac{(2l_1 + 1)(2l_2 + 1)(2l_3 + 1)(2l + \frac{1}{3})}{(2l - 2l_1 + \frac{1}{3})(2l - 2l_2 + \frac{1}{3})(2l - 2l_3 + \frac{1}{3})} \times \sqrt{\frac{(2l - 2l_1 + \frac{1}{6})(2l - 2l_2 + \frac{1}{6})(2l - 2l_3 + \frac{1}{6})}{(2l + \frac{1}{6})}}. \quad (46)$$

This is a remarkably accurate representation for the exact discrete $w_{l_1 l_2 l_3}$ with the difference between the weight functions being less than 0.01% (0.1%) for about 95% (99%) of the allowed triples $l_1 l_2 l_3$ on the domain (42) with $2 \leq l_1, l_2, l_3 \leq 2000$. The worst approximation by $w(l_1, l_2, l_3)$ never differs by more than 2.5% and such

points are exclusively located very near the boundaries, leaving an overall integrated error over the entire domain (42) of less than 0.01%. Nevertheless, care must be exercised using this approximation for edge- or corner-weighted models. With this caveat in mind, we can define the multipole sum equivalent to the wave number tetrapyd integration (41) as

$$\begin{aligned} \mathcal{T}[f] &= \sum_{\{l_1 l_2 l_3\} \in \mathcal{V}_{\mathcal{T}}} w_{l_1 l_2 l_3} f_{l_1 l_2 l_3} \\ &= \frac{1}{2} \int_{\mathcal{V}_{\mathcal{T}}} w(l_1, l_2, l_3) f(l_1, l_2, l_3) d\mathcal{V}_{\mathcal{T}}, \end{aligned} \quad (47)$$

with the inner product again defined by $\langle f, g \rangle = \mathcal{T}[fg]$. It will be clear from the context whether we are dealing with multipole or wave number integrations.

The weight function $w(l_1, l_2, l_3)$ (or $w_{l_1 l_2 l_3}$) in (46) possesses an overall scaling which grows linearly with l , as illustrated in Fig. 4. It can be convenient to eliminate this scaling, so that the weight function becomes very nearly constant. We can achieve this by dividing $w(l_1, l_2, l_3)$ by a separable function as

$$w_s(l_1, l_2, l_3) = \frac{w(l_1, l_2, l_3)}{(2l_1 + 1)^{1/3} (2l_2 + 1)^{1/3} (2l_3 + 1)^{1/3}}. \quad (48)$$

The result is shown in Fig. 4 where it is evident that $w_s \approx \text{const}$ everywhere except on the boundaries. For uniform or center-weighted bispectrum models, such as the equilateral




FIG. 4 (color online). Scaling comparison of the multipole domain weight function $w(l_1, l_2, l_3)$ (or $w_{l_1 l_2 l_3}$) given in (46) and the modified weight function $w_s(l_1, l_2, l_3)$ given in (48), which is rescaled by a separable function. On the left, the equal- l values are shown with the linear scaling of w (dashed) contrasting with the constant w_s (solid). On the right, a density plot of w_s is shown on the $l_1 + l_2 + l_3 = 2L$ slice with $L = 2000$. Note the uniformity $w_s \approx \text{const}$, except very close to the edges where there is about a factor of 4 rise to the maximum value on the perimeter.

model, the multipole domain with weight $w_s(l_1, l_2, l_3)$ becomes essentially identical to that for the primordial wave numbers (40) with $w(k_1, k_2, k_3) = 1$, so that it is a good approximation to proceed with the same polynomial expansions.

Finally, we comment on the freedom to absorb an arbitrary separable function v_l into the weight functions $w(k_1, k_2, k_3)$ or $w_{l_1 l_2 l_3}$, such as in the example (48) above. If we define a new weight \bar{w} in the estimator as

$$\bar{w}_{l_1 l_2 l_3} = w_{l_1 l_2 l_3} / (v_{l_1} v_{l_2} v_{l_3})^2, \quad (49)$$

then we must similarly rescale the estimator functions as $(b_{l_1 l_2 l_3} / \Delta) = v_{l_1} v_{l_2} v_{l_3} (b_{l_1 l_2 l_3} / \Delta)$. This rescaling should be separable, otherwise it would compromise the separability of the methods we outline here, undermining their efficiency. As we have seen it can prove convenient to make the weight functions scale invariant for practical purposes, thus facilitating better convergence of mode expansions for typical bispectra. However, in principle, we can also exploit this separability in order to remove pathologies from singular shapes, such as the local model, using a mode expansion to describe the more regular deviations away from these shapes. The important point is to consistently use both the new weight \bar{w} and the estimator rescaling throughout the analysis pipeline, including the generation of appropriate orthonormal mode functions.

B. Orthogonal polynomials on a tetrahedral domain

We next construct some concrete realizations of mode functions which are orthogonal on the tetrahedral domain $\mathcal{V}_{\mathcal{T}}$ and which have the form required for a separable expansion (39). First, we will generate one-dimensional

orthogonal polynomials $q_p(x)$ for unit weight $w = 1$, before discussing their promotion to three-dimensions and alternative weights. These tetrahedral polynomials are analogues of the more familiar Legendre polynomials $P_n(x)$ on the unit interval. Considering functions $q_p(x)$ depending only on the x -coordinate, we integrate over the y - and z -directions to yield the reduced weight function $\tilde{w}(x)$ for $x \in [0, 1]$ (we take $K = 1$):

$$\tilde{w}(x) = \frac{1}{2} x(4 - 3x), \quad \text{with} \quad \mathcal{T}[f] = \int_0^1 f(x) \tilde{w}(x) dx. \quad (50)$$

This simplifies our domain integration (41) for functions of only x , and the moments for each power of x become simply

$$w_n \equiv \mathcal{T}[x^n] = \frac{n + 6}{2(n + 3)(n + 2)}. \quad (51)$$

From these we can create orthogonal polynomials using the generating function,

$$q_n(x) = \frac{1}{\mathcal{N}} \begin{vmatrix} 1/2 & 7/24 & 1/5 & \dots & w_n \\ 7/24 & 1/5 & 3/20 & \dots & w_{n+1} \\ \dots & \dots & \dots & \dots & \dots \\ w_{n-1} & w_n & w_{n+1} & \dots & w_{2n-1} \\ 1 & x & x^2 & \dots & x^n \end{vmatrix}, \quad (52)$$

where we choose the normalization factor \mathcal{N} such that $\mathcal{T}[q_n] = 1$ for all $n \in \mathbb{N}$, that is, so that the $q_n(x)$ are orthonormal

$$\langle q_n, q_p \rangle \equiv \mathcal{T}[q_n q_p] = \int_{\mathcal{V}_{\mathcal{T}}} q_n(x) q_p(x) d\mathcal{V}_{\mathcal{T}} = \delta_{np}. \quad (53)$$

The first few orthonormal polynomials on the tetrahedral domain (40) are explicitly

$$\begin{aligned}
 q_0(x) &= \sqrt{2}, & q_1(x) &= 5.787(-\frac{7}{12} + x), \\
 q_2(x) &= 23.32(\frac{54}{215} - \frac{48}{43}x + x^2), \\
 q_3(x) &= 93.83(-0.09337 + 0.7642x - 1.631x^2 + x^3), \\
 q_4(x) &= 376.9(0.03192 - 0.4126x \\
 &\quad + 1.531x^2 - 2.139x^3 + x^4), \\
 q_5(x) &= 1512(-0.01033 + 0.1929x - 1.084x^2 \\
 &\quad + 2.549x^3 - 2.644x^4 + x^5), \dots
 \end{aligned} \tag{54}$$

These can be obtained easily from the generating determinant (52) in MATHEMATICA or similar applications.

We note that the q_n 's are only orthogonal in one dimension (e.g. we have $\mathcal{T}[q_n(x)q_p(y)] \neq \delta_{np}$ in general). However, as product functions of x , y and z they form an independent and well-behaved basis which we will use to construct orthonormal three-dimensional eigenfunctions. In practice, these q_n 's will remain the primary calculation tools throughout, notably when performing separable integrations. Where they differ from the separable functions used to represent bispectra in the literature, they generally have a number of distinct advantages, as we shall detail at the end of this section. Finally, we point out that for a regular tetrahedron [in contrast to the tetrapyd domain (40)], the volume weight function is $\tilde{w}(x) = 2x(1-x)$ and so the behavior is different at $x = 1$ where the weight vanishes, unlike (50). The first orthonormal polynomials in this case are $q_0(x) = \sqrt{3}$, $q_1(x) = 0.387(2y - 1)$, $q_2 = 32.4(y - 0.724)(y - 0.276)$, ...

Now let us turn to the polynomials $\bar{q}(x)$ which are orthonormal on the multipole domain (42), using the weight functions w given in (46) and w_s given in (48). For definiteness we take $L \equiv l_{\max} = 2000$, so that $x = l_1/L$, $y = l_2/L$ and $z = l_3/L$. The generating function (52) can be obtained as above but now using the moments $w_n \equiv \mathcal{T}[x^n] = \int w(x, y, z) d\mathcal{V}_{\mathcal{T}}$ (or by undertaking Gram-Schmidt orthogonalization from $\bar{q}_0 = \text{const}$). The resulting first few polynomials for the multipole domain are then

$$\begin{aligned}
 \bar{q}_0(x) &= 0.07378, & \bar{q}_1(x) &= 0.3017(-0.6110 + x), \\
 \bar{q}_2(x) &= 1.223(0.2665 - 1.145x + x^2), \\
 \bar{q}_3(x) &= 4.933(-0.1000 + 0.7951x - 1.659x^2 + x^3), \\
 \bar{q}_4(x) &= 19.85(0.0345 - 0.4342x \\
 &\quad + 1.578x^2 - 2.169x^3 + x^4), \\
 \bar{q}_5(x) &= 79.55(-0.0106 + 0.1975x - 1.103x^2 \\
 &\quad + 2.576x^3 - 2.657x^4 + x^5), \dots
 \end{aligned} \tag{55}$$

A cursory comparison with q_n given above for the flat wave number domain will show that these polynomials




FIG. 5 (color online). The orthonormal one-dimensional tetrahedral $q_n(x)$ plotted on the unit interval for $n = 0-5$. The behavior is smooth and bounded across the domain even for high n , except where the weight function $w(x)$ vanishes at $x = 0$. Also plotted for comparison are the rescaled Legendre polynomials $P_n(2x - 1)$ (dashed lines). Despite q_n and P_n sharing qualitative features such as n nodal points, their properties and orthogonality on $\mathcal{V}_{\mathcal{T}}$ are very different.

are very similar for low n , despite the linear scaling behavior of w . However, if we remove this scaling as in the flatter weight w_s in (48), the polynomials become near identical as illustrated in Fig. 6. It is clear that each of these polynomial sets would suffice as independent basis functions on the multipole domain. However, using the correctly weighted versions leads to improvements in the immediate orthogonality of the three-dimensional polynomials we shall construct in the following discussion.

C. Bispectrum symmetries and three-dimensional basis functions

We can represent arbitrary bispectra on the tetrahedral domain (40) using a suitable set of independent basis




FIG. 6 (color online). Orthonormal polynomials $\bar{q}_n(x)$ for the multipole domain (42) with weight functions w given in (46) [solid line] and w_s given in (48) [dashed line], as well as the previous $q_n(x)$ for unit weight [dotted line] (shown already in Fig. 5). Despite the different scaling of w , these tetrahedral polynomials are very similar, particularly the latter two with flattened weight functions.

functions formed from products $q_p(x)q_r(y)q_s(z)$ of the orthogonal polynomials (54) (or with different weight functions, such as (55)). (Here, we again take $x = k_1/k_{\max}$, $y = k_2/k_{\max}$, $z = k_3/k_{\max}$ or $x = l_1/l_{\max}$, etc.) Both primordial bispectra $B(k_1, k_2, k_3)$ and CMB bispectra $b_{l_1 l_2 l_3}$ on (40) possess six symmetries made from combinations of discrete $\pi/3$ rotations around the line $x = y = z$ and/or reflections which interchange the axes. We can impose these six symmetries on our products by summing the relevant permutations and defining the 3D basis function

$$\begin{aligned} \mathcal{Q}_n(x, y, z) &= \frac{1}{6}[q_p(x)q_r(y)q_s(z) + q_r(x)q_s(y)q_p(z) \\ &\quad + q_s(x)q_p(y)q_r(z) + q_p(x)q_s(y)q_r(z) \\ &\quad + q_s(x)q_r(y)q_p(z) + q_r(x)q_p(y)q_s(z)] \\ &\equiv q_{\{pqr\}} \quad \text{with } n \leftrightarrow \{prs\}, \end{aligned} \quad (56)$$

where we use the notation $\{prs\}$ to denote the six permutations of prs . Here, for convenience, we have specified a one-to-one mapping $n \leftrightarrow \{prs\}$ ordering the permuted indices into a list labeled by n (see below). Alternatively, we could directly represent bispectra in a power series using sums of monomial symmetric polynomials which like (56) are also separable; that is, we could identify our set of basis functions with the following

$$\begin{aligned} 1, \quad x + y + z, \quad xy + yz + zx, \quad x^2 + y^2 + z^2, \\ xyz, \quad x^3 + y^3 + z^3, \quad \text{etc.} \end{aligned} \quad (57)$$

The $\mathcal{Q}_n(x, y, z)$ we defined in (56) are themselves ultimately constructed from these through the q_p products. However, the \mathcal{Q}_n have two distinct advantages which are, first, they already have partial orthogonality built in which improves their convenience and convergence and, secondly, unlike the elements of (57), the q_p polynomials remain bounded and well-behaved when convolved with transfer functions, as we shall emphasize in the mapmaking discussion.

Since we will be dealing with relatively small numbers of basis functions, it is convenient to order the symmetric products $\mathcal{Q}_n = q_{\{pqr\}}$ linearly with a single index n ; here we offer two comparable alternatives for achieving this. The first is by ‘‘slicing’’ such that triples are ordered by the sum $p + r + s$ and the second is by ‘‘distance’’ from the origin, that is, $p^2 + r^2 + s^2$.

Slicing the prs naturally groups the \mathcal{Q}_n by the overall order of the polynomials from which they are made. The subscript n , with a specific choice of subordering, relates to the prs via

$$\begin{array}{cccc} \underline{0 \rightarrow 000} & 4 \rightarrow 111 & 8 \rightarrow 022 & 12 \rightarrow 113 \\ \underline{1 \rightarrow 001} & 5 \rightarrow 012 & 9 \rightarrow 013 & 13 \rightarrow 023 \\ 2 \rightarrow 011 & \underline{6 \rightarrow 003} & \underline{10 \rightarrow 004} & 14 \rightarrow 014 \\ \underline{3 \rightarrow 002} & 7 \rightarrow 112 & 11 \rightarrow 122 & \underline{15 \rightarrow 005} \cdots, \end{array} \quad (58)$$

where we have underlined the transitions between polynomial order. The number d_N of independent symmetric polynomial products $\mathcal{Q}_n \mathcal{Q}_p \mathcal{Q}_r$ which can be formed at each polynomial order N is a combinatorial problem but the sequence begins as follows and we give a recurrence relation for any further elements:

$$\begin{aligned} \{d_N\} &= \{1, 1, 2, 3, 4, 5, 7, 8, 10, 12, \dots\}, \\ d_N &= 1 + d_{N-2} + d_{N-3} - d_{N-5}. \end{aligned} \quad (59)$$

For consistency when using slicing, we will usually decompose functions with polynomials up to a specific order N .

The distance ordering of the \mathcal{Q}_n is more straightforward with

$$\begin{array}{cccc} 0 \rightarrow 000 & 2 \rightarrow 011 & 4 \rightarrow 002 & 6 \rightarrow 112 \\ 8 \rightarrow 122 & 1 \rightarrow 001 & 3 \rightarrow 111 & 5 \rightarrow 012 \\ 7 \rightarrow 022 & 9 \rightarrow 003 \cdots \end{array} \quad (60)$$

This approach is the analogue of state counting over spherical shells in the continuum limit and the basis functions can be grouped accordingly. Distance ordering has some advantage by reshuffling to higher n the pure states $00p$ which turn out to be most affected by masking.

While the \mathcal{Q}_n 's by construction are an independent set of three-dimensional basis functions on the domain (40), they are not in general orthogonal. In Fig. 7, we illustrate the inner product matrix $\gamma_{np} = \langle \mathcal{Q}_n, \mathcal{Q}_p \rangle$, showing partial orthogonality (nearly diagonal γ_{np}) because of their origin as products of orthogonal q_r 's. However, this is not sufficient because we need the convenience of a fully orthonormal basis to efficiently decompose arbitrary bispectra. For this reason, we undertake an iterative Gram-Schmidt orthogonalization process to construct an orthonormal set \mathcal{R}_n from the \mathcal{Q}_n , that is, satisfying

$$\langle \mathcal{R}_n, \mathcal{R}_p \rangle = \delta_{np}. \quad (61)$$

Formally, we have a Gram matrix $\Gamma = (\langle \mathcal{Q}_n, \mathcal{Q}_p \rangle)$ made from the independent functions \mathcal{Q}_n , and therefore positive definite, which needs to be factorized as $\Gamma = \Lambda^T \Lambda$ where $\Lambda = (\langle \mathcal{Q}_n, \mathcal{R}_p \rangle)$ is triangular (i.e. an LU or Cholesky decomposition). As we require explicit relationships between \mathcal{Q}_n and \mathcal{R}_n , we run through the main steps in the Gram-Schmidt process.

Let us assume that we have achieved this orthonormalization up to n , that is, such that $\langle \mathcal{R}_n, \mathcal{R}_m \rangle = \delta_{nm}, \forall m \leq n$.




FIG. 7 (color online). Partial orthogonality of the symmetric product polynomials \mathcal{Q}_n illustrated through the inner product matrix $\langle \mathcal{Q}_n, \mathcal{Q}_p \rangle$ for $0 \leq n, p < 10$ (left panel). Lower triangular matrix λ_{np} in (62) illustrating the decomposition of the orthonormal \mathcal{R}_n into the \mathcal{Q}_p arising through the Gram-Schmidt process (right panel); this is the inverse of $\langle \mathcal{Q}_n, \mathcal{R}_p \rangle$. To improve comparison, the \mathcal{Q}_n 's have been unit normalized.

This means we can represent any basis function \mathcal{Q}_p in terms of the \mathcal{R}_m and vice versa by inversion, so we can write

$$\mathcal{R}_m = \sum_{p=0}^m \lambda_{mp} \mathcal{Q}_p \quad \text{for } m, p \leq n, \quad (62)$$

where λ_{mp} is a lower triangular matrix with $(\lambda^{-1})_{np}^\top = \langle \mathcal{Q}_n, \mathcal{R}_p \rangle$. We wish by induction to construct the next orthonormal polynomial \mathcal{R}_{n+1} and infer from this the sum over basis functions up to \mathcal{Q}_{n+1} . We achieve this by taking the next independent basis function, \mathcal{Q}_{n+1} , as a first approximation to an unnormalized \mathcal{R}'_{n+1} and then we project out all components dependent on the \mathcal{R}_m ($m \leq n$),

$$\begin{aligned} \mathcal{R}'_{n+1} &\equiv \sum_{p=0}^{n+1} \lambda'_{n+1p} \mathcal{Q}_p \\ &= \mathcal{Q}_{n+1} - \sum_{m=0}^n \mathcal{R}_m \int_{\mathcal{V}_T} \mathcal{Q}_{n+1} \mathcal{R}_m w d\mathcal{V}_T \\ &= \mathcal{Q}_{n+1} - \sum_{m=0}^n \sum_{r=0}^m \sum_{s=0}^m \lambda_{mr} \lambda_{ms} \gamma_{n+1s} \mathcal{Q}_r \end{aligned} \quad (63)$$

where in the second line we have substituted (62) and the γ_{n+1s} are determined from the relative orthogonality of the \mathcal{Q}_n 's,

$$\gamma_{n+1s} = \langle \mathcal{Q}_{n+1}, \mathcal{Q}_s \rangle = \int_{\mathcal{V}_T} \mathcal{Q}_{n+1} \mathcal{Q}_s w d\mathcal{V}_T. \quad (64)$$

By equating coefficients in the expression (63) we can determine that

$$l'_{n+1p} = \delta_{n+1p} - \sum_{r=p}^n \sum_{s=0}^r l_{rp} l_{rs} \gamma_{n+1s}. \quad (65)$$

Unit normalizing appropriately, we obtain the coefficients λ_{n+1p} which define the new orthonormal \mathcal{R}_{n+1} we are

seeking, that is, we have

$$l_{n+1p} = l'_{n+1p} / \left(\sum_{r,s=0}^{n+1} l'_{n+1r} l'_{n+1s} \gamma_{rs} \right). \quad (66)$$

In Fig. 7, we see the orthogonalization process at work for the first 10 modes by plotting the matrix coefficients for $\langle \mathcal{Q}_n, \mathcal{Q}_p \rangle$ and $\langle \mathcal{Q}_n, \mathcal{R}_p \rangle$. At each order n , the independent component in \mathcal{R}_n is provided by \mathcal{Q}_n , as indicated by the dominant diagonal term. This is a good approximation at low order, but the mixing increases with n . We also illustrate several of the orthogonal polynomials \mathcal{R}_n on the tetrapyd domain Fig. 8 for the slicing ordering (58). These are primarily the lowest modes and demonstrate the build up of the number of nodal points and lines as the order increases. As an aside, we note Gram-Schmidt orthogonalization in the form given above is inherently unstable numerically, though this can be easily corrected by using the modified Gram-Schmidt process. However, we do not iterate to sufficiently high n to notice any significant degradation in accuracy, as verified by determining orthogonality.

D. Mode decomposition of the bispectrum

We have constructed examples of an orthonormal basis $\{\mathcal{R}_n\}$ out of monomial symmetric polynomials (57) which span the set of symmetric functions on the tetrahedral domain (40). The \mathcal{R}_n polynomials will possess the properties of more familiar orthonormal eigenmodes in other contexts, notably completeness and the convergence of mode expansions for well-behaved functions. We proceed by considering an arbitrary primordial bispectrum (12) described by the shape function $S(k_1, k_2, k_3)$ and decomposing it as follows




FIG. 8 (color online). Three-dimensional orthonormal polynomials \mathcal{R}_n on the tetrahedral domain (40). Taken from top left (and moving across and then down) these are \mathcal{R}_0 , \mathcal{R}_1 , \mathcal{R}_2 , \mathcal{R}_3 , \mathcal{R}_4 , and \mathcal{R}_{41} (bottom right).

$$S(k_1, k_2, k_3) = \sum_{n=0}^{\infty} \alpha_n^{\mathcal{R}} \mathcal{R}_n(x, y, z), \quad (67)$$

where the expansion coefficients $\alpha_n^{\mathcal{R}}$ are given by

$$\alpha_n^{\mathcal{R}} = \langle \mathcal{R}_n, S \rangle = \int_{\mathcal{V}_{\mathcal{T}}} \mathcal{R}_n S \omega d\mathcal{V}_{\mathcal{T}}, \quad (68)$$

and $K = k_{\max}$ and $k_1 = Kx$ etc. on the domain $\mathcal{V}_{\mathcal{T}}$ de-

fined in (40). For practical purposes, we shall always work with partial sums up to a given $N = n_{\max}$ with

$$S_N = \sum_{n=0}^N \alpha_n^{\mathcal{R}} \mathcal{R}_n(x, y, z), \quad S = \lim_{N \rightarrow \infty} S_N. \quad (69)$$

We shall assume that the expansion (69) is the best fit mode expansion of degree N (for this particular mode ordering). Given the complete orthonormal basis \mathcal{R}_n , Parseval's

theorem for the integrated product of two functions implies

$$\langle S, S' \rangle = \int_{\mathcal{V}_{\mathcal{T}}} SS' wd\mathcal{V}_{\mathcal{T}} = \lim_{N \rightarrow \infty} \sum_{n=0}^N \alpha_n^{\mathcal{R}} \alpha_n^{\mathcal{R}'}, \quad (70)$$

which, for the square of a function S , yields the sum of the squares of the expansion coefficients, $\mathcal{T}[S^2] = \sum_n \alpha_n^{\mathcal{R}2}$.

In order to accomplish our original goal of a general separable expansion (39), we must now transform backwards from the orthonormal \mathcal{R}_n sum (69) into an expansion over the separable product functions $\mathcal{Q}_n = q_{\{p,q,r,s\}}$ through

$$S_N = \sum_{n=0}^N \alpha_n^{\mathcal{Q}} \mathcal{Q}_n(x, y, z), \quad (71)$$

where the $\alpha_n^{\mathcal{Q}}$ can be obtained from the $\alpha_n^{\mathcal{R}}$ as

$$\alpha_n^{\mathcal{Q}} = \sum_{p=0}^N (\lambda^{\top})_{np} \alpha_p^{\mathcal{R}}, \quad (72)$$

with the transformation matrix λ_{np} defined in (62) (this is triangular and not orthogonal in general). Note the complication that $\alpha_n^{\mathcal{Q}}$ also contains contributions from \mathcal{R}_p components with $n < p \leq N$, since $(\lambda^{\top})_{np}$ is upper triangular. The inverse transformation

$$\alpha_n^{\mathcal{R}} = \sum_p^N (\lambda^{-1})_{np}^T \alpha_p^{\mathcal{Q}}, \quad (73)$$

has coefficients given by $(\lambda^{-1})_{np} = \langle \mathcal{Q}_n, \mathcal{R}_p \rangle$. We have already noted that the degree of nonorthogonality of the \mathcal{Q}_n basis is described by $\gamma_{np} = \langle \mathcal{Q}_n, \mathcal{Q}_p \rangle$ in (64) which is in turn related to λ_{np} through

$$(\gamma^{-1})_{np} = \sum_r^N (\lambda^{\top})_{nr} \lambda_{rp}. \quad (74)$$

When substituted into Parseval's theorem (70) in the \mathcal{Q}_n basis, we see that the coefficients of different degrees become mixed as

$$\langle S_N, S_N \rangle = \sum_n^N \alpha_n^{\mathcal{R}2} = \sum_n^N \sum_p^N \alpha_n^{\mathcal{Q}} \gamma_{np} \alpha_p^{\mathcal{Q}} \quad (75)$$

The separable \mathcal{Q}_n expansion (71) is important for most practical calculational purposes but its coefficients are constructed at the outset using the orthonormal \mathcal{R}_n . For interpreting results from the estimator it is helpful to transform back to the \mathcal{R}_n basis in order to understand the normalized spectrum $\alpha_n^{\mathcal{R}}$ using Parseval's theorem (70). We finally note that all the transformation matrices, λ_{np} and γ_{np} in (64), need only be calculated once, at the same time as the \mathcal{R}_n polynomials are generated, and then stored for later reference.

In Fig. 9, we demonstrate polynomial convergence for the DBI model and its separable equilateral approximation by showing the cross-correlation between the shape function and the partial sum (69). We also provide the actual expansion coefficients $\alpha_n^{\mathcal{R}}$ for the primordial shape functions in Fig. 10 (along with the $\bar{\alpha}_n^{\mathcal{R}}$ for the CMB bispectra). Using only 6 three-dimensional \mathcal{R}_n polynomials we achieve a better than 98% cross-correlation with the original analytic expressions in both cases [i.e. using symmetric products of at most quadratic q_p polynomials from (54)]. Here, we undertake the full Fisher matrix analysis between the theoretical CMB bispectrum and its approximation using the methods described in Ref. [11]. More generally, we note that for all well-behaved bispectra the polynomial expansion has proved to be rapidly convergent. The decomposition of the primordial bispectra is also numerically




FIG. 9 (color online). Correlation of the reconstructed bispectra to the original for partial sums of the decomposition up to a given mode n . The plot includes the primordial bispectra for the equilateral and DBI models, the CMB bispectrum for the equilateral and DBI models and the CMB bispectrum produced at late times by cosmic strings. In all cases, we find that with 15 three-dimensional modes we have a correlation greater than 98%, thus demonstrating very rapid convergence. For the CMB bispectra, convergence is limited by matching the acoustic peaks introduced by the transfer functions, whereas the primordial models converge at 98% accuracy with only 6 modes.




FIG. 10 (color online). Decomposition into orthonormal polynomials \mathcal{R}_n for both the primordial shape function (67) (above) and the CMB bispectrum estimator (89) (below) for the equilateral (red) and DBI (blue) models. In both cases, these results are for “slicing” polynomial ordering given in (58). The peak in the CMB bispectrum estimator modes (here, for $l_{\max} = 500$ at $n \approx 5$) arises because of the power shifted into the coherent acoustic peaks observed in Fig. 3; this is a distinguishing feature of the CMB bispectrum for most primordial models [14].

efficient using the orthogonal \mathcal{R}_n modes with each $\alpha_n^{\mathcal{R}}$ coefficient taking an average of 7 seconds to calculate.

We can equally well expand the CMB bispectrum $b_{l_1 l_2 l_3}$ at late times, using the same polynomials $b_{l_1 l_2 l_3} = \sum_n \bar{\alpha}_n^{\mathcal{R}} \mathcal{R}_n(x, y, z)$. However, our aim is to represent the bispectrum estimator \mathcal{E} given in (102), rather than $b_{l_1 l_2 l_3}$ itself. We, therefore, consider expanding a separable product which approximates \mathcal{E} with the same weight and scaling (schematically, $\sqrt{l} b_{l_1 l_2 l_3} / C_l^{3/2}$). We discuss this in the next section, but in the lower half of Fig. 10 we show the corresponding late-time expansion coefficients $\bar{\alpha}_n^{\mathcal{R}}$ for the equilateral and DBI CMB bispectra. Once again, convergence is rapid, see Fig. 9, with a 95% correlation achieved with only 12 $\bar{\mathcal{R}}_n$ polynomials at $l_{\max} = 500$, improving to 98% with 15 polynomials for both CMB bispectra. This is despite the fact that the expansion must incorporate additional features induced by the transfer functions. We emphasize the power shift from the low modes in the primordial bispectrum to a peak at higher modes $n \approx 5$ in the CMB bispectrum (for this slicing and l_{\max}). This is a common characteristic of the polynomial expansion for almost all bispectra of primordial origin and is a manifestation of the pattern of coherent acoustic peaks observed in Ref. [14]. The decomposition of the CMB bispectra is also

numerically efficient using the orthogonal $\bar{\mathcal{R}}_n$ modes with each $\bar{\alpha}_n^{\mathcal{R}}$ coefficient taking an average of 8 seconds to calculate.

E. Utility of the tetrahedral polynomials \mathcal{Q}_n and other alternatives

The three-dimensional polynomials \mathcal{Q}_n we have presented are just one possible set of basis functions which can be used as bispectrum eigenmodes for the methodology we present in the next section. They are built from products of the one-dimensional q_p 's which are orthonormal on the tetrahedral region (40) with given weight functions. These are the analogues of Legendre polynomials P_n . Unfortunately, unlike the P_n 's on a cube, they do not retain full orthogonality as separable products on the tetrahedral domain, though there is a substantial remnant. There are significant advantages to using the q_p 's, rather than the monomial symmetric polynomials in (57), in the same way that Legendre polynomials are more efficient than simple power series representations. As we shall discuss subsequently, there are further important benefits which arise when the \mathcal{Q}_n 's are decomposed into separable integrals over the q_p 's. Given the bounded and well-behaved nature of the q_p 's on their domain, these integrals reflect these

properties, eliminating diverging artifacts which are known for other separable approximations to bispectra in the literature (including difficulties for simple powers x^n).

There are other alternatives to expansions using the tetrahedral polynomials \mathcal{Q}_n and \mathcal{R}_n which we have considered. It is possible, for example, to expand an arbitrary bispectrum using separable products of more familiar orthonormal functions such as Legendre the P_n and Chebyshev T_n polynomials, as discussed in Ref. [14]. This entails using shifted polynomials on the full cubic domain $l_1, l_2, l_3 \leq l_{\max}$. The shortcoming of this approach is that the bispectrum is only defined on the tetrahedral region (40), so it has to be zero elsewhere or arbitrarily extended in some manner to fill the cube. This leads to generic overshooting of the expansion near the boundaries (the analogue of the Gibbs phenomena for Fourier series). Extensive experiments yielded very poor convergence with Legendre and Chebyshev polynomial expansions, as well as Fourier series, especially relative to that achieved with the tetrahedral \mathcal{Q}_n and \mathcal{R}_n polynomials. A further simple alternative is to transform the tetrahedral region into a cube (see Ref. [11]). This allows the bispectrum to be defined everywhere on the standard domain using the more familiar eigenmodes and thus yielding more rapid convergence. However, this compromises separability which is essential for the estimators we discuss below.

Were the rate of convergence to become a primary issue when representing the bispectrum, then there are further alternatives to polynomials. There is a significant literature

on eigenmodes on the regular tetrahedron or simplex because of its importance in crystallography and other contexts. For example, it is possible to define generalized sine and cosine functions on the simplex, as well as Koornwilder and generalized Chebyshev polynomials of the first and second kind (see, for example, ref. [28]). Such generalized eigenfunctions could, in principle, improve convergence, however, two significant developments are required. First, it is more natural to define the observational data on the tetrahedral domain with $l_1, l_2, l_3 \leq l_{\max}$ (the tetrapyd), rather than the simplex $l_1 + l_2 + l_3 \leq 2l_{\max}$, so generalized eigenfunctions must be derived explicitly for this domain (42). Second, these should be able to conveniently represent functions in separable form. The present tetrahedral polynomials \mathcal{R}_n and \mathcal{Q}_n do converge satisfactorily for all the primordial models studied to date, but more efficient mode expansions will continue to be investigated [29].

IV. MEASURES OF F_{NL}

A. Primordial F_{NL} estimator

We have obtained two related mode expansions for a general primordial shape function (12), one for an orthonormal basis \mathcal{R}_n (69) and the other for separable basis functions \mathcal{Q}_n (71). Substitution of the separable form into the expression for the reduced bispectrum (14) offers an efficient route to its direct calculation through

$$\begin{aligned}
 b_{l_1 l_2 l_3} &= \left(\frac{2}{\pi}\right)^3 \Delta_{\Phi}^2 f_{\text{NL}} \int x^2 dx dk_1 dk_2 dk_3 6 \sum_n \alpha_n^{\mathcal{Q}} \mathcal{Q}_n(k_1, k_2, k_3) \Delta_{l_1}(k_1) \Delta_{l_2}(k_2) \Delta_{l_3}(k_3) j_{l_1}(k_1 x) j_{l_2}(k_2 x) j_{l_3}(k_3 x) \\
 &= \Delta_{\Phi}^2 f_{\text{NL}} \sum_n \alpha_n^{\mathcal{Q}} \int x^2 dx \left[\left(\frac{2}{\pi} \int dk_1 q_p(k_1) \Delta_{l_1}(k_1) j_{l_1}(k_1 x)\right) \right. \\
 &\quad \times \left. \left(\frac{2}{\pi} \int dk_2 q_r(k_2) \Delta_{l_2}(k_2) j_{l_2}(k_2 x)\right) \left(\frac{2}{\pi} \int dk_3 q_s(k_3) \Delta_{l_3}(k_3) j_{l_3}(k_3 x)\right) + 5 \text{ permutations} \right] \\
 &= \Delta_{\Phi}^2 f_{\text{NL}} \sum_n \alpha_n^{\mathcal{Q}} \int x^2 dx q_{\{p}^1 q_r^2 q_s^3 \\
 &= \Delta_{\Phi}^2 f_{\text{NL}} \sum_n \alpha_n^{\mathcal{Q}} \int x^2 dx \mathcal{Q}_n^{l_1 l_2 l_3},
 \end{aligned} \tag{76}$$

where here we implicitly assume the mapping $n \leftrightarrow prs$ between indices for the \mathcal{Q}_n and the product basis functions from which they are formed, that is, $\mathcal{Q}_n = q_{\{p} q_r q_s$ [e.g. see the ordering in (58)]. For brevity we have also denoted as q_p^l the convolution of the basis function $q_p(k)$ with the transfer functions

$$q_p^l(x) = \frac{2}{\pi} \int dk q_p(k) \Delta_l(k) j_l(kx),$$

with $\mathcal{Q}_n^{l_1 l_2 l_3}(x) = q_{\{p}^{l_1}(x) q_r^{l_2}(x) q_s^{l_3}(x)$. (77)

[These q_p^l are the primordial counterparts of the \bar{q} defined in multipole space (55).] Here, in (76), the previously

intractable three-dimensional wave number integral separates into the product of three one-dimensional integrals which are relatively easy to evaluate. This has been achieved because the triangle condition has been enforced through the product of Bessel functions, giving a manifestly separable form and allowing us to interchange the orders of integration with x ; it is the basis for the analytic local (19) and constant (16) solutions on large angles, as well as all the analysis of separable shape functions to date (see, for example, Ref. [9]). With this mode expansion, all nonseparable theoretical CMB bispectra $b_{l_1 l_2 l_3}$ become calculable provided there is a convergent expansion for the shape function. Accurate hierarchical schemes already

exist against which to benchmark this method [14] but, in principle, it is more efficient.

Now consider the implications of this mode expansion for f_{NL} by substituting the decomposed $b_{l_1 l_2 l_3}$ (71) into the estimator expression (31) to obtain

$$\begin{aligned} \mathcal{E} &= \Delta_{\Phi}^2 f_{\text{NL}} \sum_n \alpha_n^{\mathcal{Q}} \sum_{l_i, m_i} \int x^2 dx q_{l_p}^{l_1}(x) q_{l_r}^{l_2}(x) q_{l_s}^{l_3}(x) \\ &\quad \times \int d^2 \hat{\mathbf{n}} Y_{l_1 m_1}(\hat{\mathbf{n}}) Y_{l_2 m_2}(\hat{\mathbf{n}}) Y_{l_3 m_3}(\hat{\mathbf{n}}) \frac{a_{l_1 m_1} a_{l_2 m_2} a_{l_3 m_3}}{C_{l_1} C_{l_2} C_{l_3}} \end{aligned} \quad (78)$$

$$\begin{aligned} &= \Delta_{\Phi}^2 f_{\text{NL}} \sum_n \alpha_n^{\mathcal{Q}} \int d^2 \hat{\mathbf{n}} \int x^2 dx \left[\sum_{l_1, m_1} q_p^{l_1} \frac{a_{l_1 m_1} Y_{l_1 m_1}}{C_{l_1}} \right. \\ &\quad \left. \times \sum_{l_2, m_2} q_r^{l_2} \frac{a_{l_2 m_2} Y_{l_2 m_2}}{C_{l_2}} \sum_{l_3, m_3} q_s^{l_3} \frac{a_{l_3 m_3} Y_{l_3 m_3}}{C_{l_3}} \right]. \end{aligned} \quad (79)$$

The break up of the wave number integration now extends also to the separation of the sum over the multipoles l_1, l_2, l_3 . The summation between the a_{lm} 's and each q_p integral creates a filtered map of the original data, which we can define in the above as

$$\begin{aligned} M_p(\hat{\mathbf{n}}, x) &= \sum_{lm} q_p^l \frac{a_{lm} Y_{lm}}{C_l} \\ &= \sum_{lm} \left[\frac{2}{\pi} \int q_p(k) \Delta_l(k) j_l(kx) dk \right] \frac{a_{lm} Y_{lm}(\hat{\mathbf{n}})}{C_l}. \end{aligned} \quad (80)$$

From these we can efficiently calculate product maps which essentially extract the \mathcal{Q}_n basis function contribution from the observational data,

$$\mathcal{M}_n^{\mathcal{Q}}(\hat{\mathbf{n}}, x) = M_p(\hat{\mathbf{n}}, x) M_r(\hat{\mathbf{n}}, x) M_s(\hat{\mathbf{n}}, x) \quad (81)$$

where again we exploit the correspondence $n \leftrightarrow prs$. Note, that in this case, there is no need to symmetrize the product map because it is implicit in the estimator expression. Integrating over directions and shells we can now obtain for the observational maps, the analogue of the primordial mode expansion coefficients $\alpha_n^{\mathcal{Q}}$,

$$\beta_n^{\mathcal{Q}} = \int d^2 \hat{\mathbf{n}} \int x^2 dx \mathcal{M}_n^{\mathcal{Q}}(\hat{\mathbf{n}}, x). \quad (82)$$

In common with the analysis of simple separable shapes, the shell integral over x is where the most significant computational effort is required.

Substituting into (78), the bispectrum estimator then collapses into a compact diagonal form

$$\mathcal{E} = \frac{6\Delta_{\Phi}^2}{N} \sum_n \alpha_n^{\mathcal{Q}} \beta_n^{\mathcal{Q}}. \quad (83)$$

The estimator has been reduced entirely to tractable integrals and sums which can be performed rapidly even at

$l_{\text{max}} = 2000$. We will demonstrate how efficiently it can recover f_{NL} from simulated maps in subsequent sections.

The form of the estimator in (83) suggests that further information can be extracted about the observed bispectrum beyond the f_{NL} for one specific theoretical model. This is because, through the coefficients $\beta_n^{\mathcal{Q}}$, we have obtained some sort of mode decomposition of the bispectrum of the observational map. However, the nonorthogonal and primordial nature of the \mathcal{Q}_n basis functions means that these $\beta_n^{\mathcal{Q}}$ require some effort in their interpretation. Consider the expectation value of $\beta_n^{\mathcal{Q}}$ obtained from an ensemble of maps generated for a particular theoretical model with shape function $S = \sum_n \alpha_n^{\mathcal{Q}} \mathcal{Q}_n$. Noting that the relation $\langle a_{l_1 m_1} a_{l_2 m_2} a_{l_3 m_3} \rangle = \mathcal{G}_{m_1 m_2 m_3}^{l_1 l_2 l_3} b_{l_1 l_2 l_3}$, the average over the product maps (81) becomes

$$\begin{aligned} \langle \beta_n^{\mathcal{Q}} \rangle &= \int d^2 \hat{\mathbf{n}} \int x^2 dx \langle \mathcal{M}_n^{\mathcal{Q}}(\hat{\mathbf{n}}, x) \rangle \\ &= \sum_{l_i, m_i} \left(\int x^2 dx q_{l_p}^{l_1} q_{l_r}^{l_2} q_{l_s}^{l_3} \right) \left(\mathcal{G}_{m_1 m_2 m_3}^{l_1 l_2 l_3} \right)^2 \frac{b_{l_1 l_2 l_3}}{C_{l_1} C_{l_2} C_{l_3}} \quad (84) \\ &= \sum_{l_1 l_2 l_3} \frac{1}{4\pi} \frac{(2l_1 + 1)(2l_2 + 1)(2l_3 + 1)}{C_{l_1} C_{l_2} C_{l_3}} \begin{pmatrix} l_1 & l_2 & l_3 \\ 0 & 0 & 0 \end{pmatrix}^2 \\ &\quad \times \int x^2 dx \mathcal{Q}_n^{l_1 l_2 l_3} \sum_p \alpha_p^{\mathcal{Q}} \int x^2 dx \mathcal{Q}_p^{l_1 l_2 l_3} \quad (85) \end{aligned}$$

$$\begin{aligned} &= \sum_p \alpha_p^{\mathcal{Q}} \sum_{l_1 l_2 l_3} \frac{w_{l_1 l_2 l_3}}{C_{l_1} C_{l_2} C_{l_3}} \int x^2 dx \mathcal{Q}_n^{l_1 l_2 l_3} \int x^2 dx \mathcal{Q}_p^{l_1 l_2 l_3} \\ &\equiv \sum_p \Gamma_{np}^{\mathcal{Q}} \alpha_p^{\mathcal{Q}}, \end{aligned} \quad (86)$$

where we have substituted the expression (76) for the reduced bispectrum and the weight $w_{l_1 l_2 l_3}$ is described in (43). Here, the matrix $\Gamma_{np}^{\mathcal{Q}}$ represents a late-time inner product $\Gamma_{np}^{\mathcal{Q}} = \langle \langle \mathcal{Q}_n, \mathcal{Q}_p \rangle \rangle$ analogous to $\gamma_{np} = \langle \mathcal{Q}_n, \mathcal{Q}_p \rangle$ in (64) (but with a different weight so that $\langle \langle \mathcal{R}_n, \mathcal{R}_p \rangle \rangle \neq \delta_{np}$). Determining the transformation matrix $\Gamma_{np}^{\mathcal{Q}}$ relating the $\alpha_n^{\mathcal{Q}}$ and $\beta_n^{\mathcal{Q}}$ appears to be a complicated task but, in fact, it reduces to separable sums and integrals over the one-dimensional products $q_p q_r$ convolved with Bessel and transfer functions, together with the final sum over the multipole domain (42). The latter is straightforward, especially in the continuum limit (46). It need only be evaluated once, given a robust prior estimate for the power spectrum C_l 's.

This discussion demonstrates that we can recover spectral information about the primordial shape function from the observational data through the relation

$$\alpha_n^{\mathcal{Q}} = \sum_p (\Gamma^{\mathcal{Q}^{-1}})_{np} \langle \beta_p^{\mathcal{Q}} \rangle, \quad (87)$$

which extends to the orthonormal coefficients $\alpha_n^{\mathcal{R}}$ using

(73). If the decomposition coefficients $\beta_n^{\mathcal{Q}}$ are found with adequate significance, we can reconstruct the shape function from a single realization through the expansion

$$S(k_1, k_2, k_3) = \sum_{n,p} (\Gamma^{\mathcal{Q}^{-1}})_{np} \beta_p^{\mathcal{Q}} \mathcal{Q}_n. \quad (88)$$

We will discuss this prospect in more detail in the next section about the late-time CMB estimator where the relation between the $\alpha_n^{\mathcal{Q}}$ and $\beta_n^{\mathcal{Q}}$ is more transparent.

B. CMB F_{NL} estimator

We turn now to the implementation details of the late-time CMB estimator originally proposed in [14]. Here, we presume that the CMB bispectrum $b_{l_1 l_2 l_3}$ for our nonseparable primordial model is precomputed using Eq. (76) or a robust hierarchical scheme [11,14]. In addition, this approach can accommodate any late-time source of non-Gaussianity in the CMB, including secondary anisotropies, gravitational lensing, active models such as cosmic strings, and even systematic experimental effects. For the late-time analysis we wish to expand the estimator functions using the orthonormal $\bar{\mathcal{R}}_n(l_1, l_2, l_3)$ and separable $\bar{\mathcal{Q}}_n(l_1, l_2, l_3)$ mode functions created out of products of the $\bar{q}_p(l)$ polynomials, for which we gave a concrete example (55). (Note that we denote the multipole modes with a bar, distinguishing them from the primordial q_p , \mathcal{Q}_n , \mathcal{R}_n which are functions of wave number k). Convergence of mode expansions on the multipole domain (42) has been found to be poor for quantities as scale dependent as $b_{l_1 l_2 l_3}$, so we choose to decompose the estimator functions directly as

$$\frac{v_{l_1} v_{l_2} v_{l_3}}{\sqrt{C_{l_1} C_{l_2} C_{l_3}}} b_{l_1 l_2 l_3} = \sum_n \bar{\alpha}_n^{\mathcal{Q}} \bar{\mathcal{Q}}_n, \quad (89)$$

where the separable v_l incorporates the freedom to make the weight function $w_{l_1 l_2 l_3}$ given in (43) even more scale invariant typically we shall use $v_l = (2l + 1)^{1/6}$ as defined in (48)]. The expression (89) means that we are effectively expanding in mode functions modulated by the C_l 's, that is, $\bar{\mathcal{Q}}_n \rightarrow \sqrt{C_l} \bar{\mathcal{Q}}_n / v_l$. These more closely mimic the acoustic peaks observed in the $b_{l_1 l_2 l_3}$ as illustrated in Fig. 11. We shall see that the estimator expansion with C_l in (89) is appropriate for primordial models, but different flatter choices will be more suitable for late-time anisotropy, such as that from cosmic strings.

We determine the implications for f_{NL} of our mode expansion (89) by substituting into the estimator (31),

$$\begin{aligned} \mathcal{E} = & \sum_{l_i, m_i} \sum_{n \rightarrow prs} \bar{\alpha}_n^{\mathcal{Q}} \bar{q}_{\{p} \bar{q}_r \bar{q}_{s\}} \int d^2 \hat{\mathbf{n}} Y_{l_2 m_2}(\hat{\mathbf{n}}) Y_{l_1 m_1}(\hat{\mathbf{n}}) \\ & \times Y_{l_3 m_3}(\hat{\mathbf{n}}) \frac{a_{l_1 m_1} a_{l_2 m_2} a_{l_3 m_3}}{v_{l_1} v_{l_2} v_{l_3} \sqrt{C_{l_1} C_{l_2} C_{l_3}}} \end{aligned} \quad (90)$$

$$\begin{aligned} = & \sum_{n \rightarrow prs} \bar{\alpha}_n^{\mathcal{Q}} \int d^2 \hat{\mathbf{n}} \left(\sum_{l_1, m_1} \bar{q}_{\{p} \frac{a_{l_1 m_1} Y_{l_1 m_1}}{v_{l_1} \sqrt{C_{l_1}}} \right) \left(\sum_{l_2, m_2} \bar{q}_r \frac{a_{l_2 m_2} Y_{l_2 m_2}}{v_{l_2} \sqrt{C_{l_2}}} \right) \\ & \times \left(\sum_{l_3, m_3} \bar{q}_{s\} \frac{a_{l_3 m_3} Y_{l_3 m_3}}{v_{l_3} \sqrt{C_{l_3}}} \right), \end{aligned} \quad (91)$$

where again we assume the correspondence between the label n and an ordered list of permuted triples $\{prs\}$, through $\bar{\mathcal{Q}}_n = \bar{q}_{\{p} \bar{q}_r \bar{q}_{s\}}$. As previously for the primordial estimator (78), we note that the sum between the $\bar{q}_p(l)$ and the a_{lm} creates filtered versions of the original CMB map defined by

$$\bar{M}_p(\hat{\mathbf{n}}) = \sum_{lm} q_p(l) \frac{a_{lm}}{v_l \sqrt{C_l}} Y_{lm}(\hat{\mathbf{n}}), \quad (92)$$

which are multiplied together in (90) to form the product map

$$\bar{\mathcal{M}}_n(\hat{\mathbf{n}}) = \bar{M}_p(\hat{\mathbf{n}}) \bar{M}_r(\hat{\mathbf{n}}) \bar{M}_s(\hat{\mathbf{n}}). \quad (93)$$

Integrating over directions, we can obtain the map mode expansion coefficient

$$\bar{\beta}_n^{\mathcal{Q}} = \int d^2 \hat{\mathbf{n}} \bar{\mathcal{M}}_n(\hat{\mathbf{n}}). \quad (94)$$

Thus the estimator reduces again to diagonal form

$$\mathcal{E} = \frac{1}{N} \sum_{n=0}^{n_{\text{max}}} \bar{\alpha}_n^{\mathcal{Q}} \bar{\beta}_n^{\mathcal{Q}}. \quad (95)$$

Like (83), it consists entirely of separable sums and tractable integrals which can be performed rapidly.

As before, the separation of the estimator into two $\bar{\alpha}_n^{\mathcal{Q}}$ and $\bar{\beta}_n^{\mathcal{Q}}$ halves indicates that this approach could offer more direct information about the bispectrum than just f_{NL} for one model. Consider the expectation value of $\bar{\beta}_n^{\mathcal{Q}}$ from an ensemble of maps with a given $b_{l_1 l_2 l_3}$ (and C_l) expanded as (89). Following the steps used to derive (84), we find a considerably simpler expression after substituting (89):

$$\begin{aligned} \langle \bar{\beta}_n^{\mathcal{Q}} \rangle = & \int d^2 \hat{\mathbf{n}} \int x^2 dx \langle \bar{\mathcal{M}}_n^{\mathcal{Q}}(\hat{\mathbf{n}}, x) \rangle \\ = & \sum_{l_i, m_i} \frac{\bar{q}_{\{p}(l_1) \bar{q}_r(l_2) \bar{q}_{s\}(l_3)}{v_{l_1} v_{l_2} v_{l_3} \sqrt{C_{l_1} C_{l_2} C_{l_3}}} (\mathcal{G}_{m_1 m_2 m_3}^{l_1 l_2 l_3})^2 b_{l_1 l_2 l_3} \\ = & \sum_{l_i, m_i} \frac{w_{l_1 l_2 l_3} \bar{\mathcal{Q}}_n(l_1, l_2, l_3)}{v_{l_1} v_{l_2} v_{l_3} \sqrt{C_{l_1} C_{l_2} C_{l_3}}} \sum_p \bar{\alpha}_p^{\mathcal{Q}} \bar{\mathcal{Q}}_p(l_1, l_2, l_3) \frac{\sqrt{C_{l_1} C_{l_2} C_{l_3}}}{v_{l_1} v_{l_2} v_{l_3}} \end{aligned} \quad (96)$$

$$= \sum_p \bar{\alpha}_p^{\mathcal{Q}} \sum_{l_1 l_2 l_3} \bar{w}_{l_1 l_2 l_3} \bar{\mathcal{Q}}_n \bar{\mathcal{Q}}_p = \sum_p \bar{\Gamma}_{np} \bar{\alpha}_p^{\mathcal{Q}}, \quad (98)$$

where the modified weight function $\bar{w}_{l_1 l_2 l_3}$ is given in (49) and $\bar{\Gamma}_{np} = \langle \bar{\mathcal{Q}}_n, \bar{\mathcal{Q}}_p \rangle$ as discussed previously. Hence, the




FIG. 11 (color online). Polynomials $\bar{\mathcal{R}}_n$ on the tetrahedral domain (42) used for representing modes in the CMB estimator multiplied by the weight function given in (89). These are ordered just as in Fig. 8 with $\bar{\mathcal{R}}_0$ (top left), $\bar{\mathcal{R}}_1$, $\bar{\mathcal{R}}_2$, $\bar{\mathcal{R}}_3$, $\bar{\mathcal{R}}_4$, and $\bar{\mathcal{R}}_{41}$ (bottom right). The last higher mode bears a superficial resemblance to the equilateral bispectrum in Fig. 3.

estimator, when applied to a map containing the bispectrum defined by $\bar{\alpha}_n^Q$, should have the expectation value

$$\langle \mathcal{E} \rangle = \frac{1}{N} \sum_n \sum_p \bar{\alpha}_n^Q \bar{\Gamma}_{np} \bar{\alpha}_p^Q. \quad (99)$$

Now rotating to our orthonormal basis $\bar{\mathcal{R}}_n$, we note that from the relation (75) we can deduce the simple and elegant form

$$\langle \mathcal{E} \rangle = \frac{1}{N} \sum_n \bar{\alpha}_n^{\mathcal{R}2}. \quad (100)$$

That is, we expect the best fit $\bar{\beta}_n^{\mathcal{R}}$'s for a particular realization to be the $\bar{\alpha}_n^{\mathcal{R}}$'s themselves. The simplicity of this result is not unexpected, since it would be obtained by correlating a bispectrum decomposed into the $\bar{\mathcal{R}}_n$ with itself. The advance here is that extracting the spectrum $\bar{\beta}_n^{\mathcal{R}}$ from the observed map would be intractable for large l_{\max} , were it not for the transformation made to a non-

orthogonal separable frame. Assuming the coefficients $\bar{\beta}_n^{\mathcal{R}}$ are measured with some significance from a particular experiment, we can go further and reconstruct the map bispectrum using (89)

$$b_{l_1 l_2 l_3} = \frac{\sqrt{C_{l_1} C_{l_2} C_{l_3}}}{v_{l_1} v_{l_2} v_{l_3}} \sum_n \bar{\beta}_n^{\mathcal{R}} \bar{\mathcal{R}}_n. \quad (101)$$

We reiterate that the viability of this fast and general reconstruction scheme [14] depends on two key factors, first, the smoothness of the reduced bispectrum $b_{l_1 l_2 l_3}$, requiring few modes to characterize it, and, second, on the completeness of the orthonormal basis from which the separable expansion was obtained. We note that this methodology can be applied using any complete mode expansions, beyond the polynomial examples given here, as well as with over-complete decompositions, such as wavelets, or with binning. In the next section, we will demonstrate the efficacy of this method with simulated maps (for a sufficiently large f_{NL}), recovering the expected $\bar{\alpha}_n^{\mathcal{R}}$ spectrum and the main distinguishing features of the bispectrum $b_{l_1 l_2 l_3}$.

C. Observable F_{NL} normalization

In previous work [11], we pointed out the shortcomings of normalizing the quantity f_{NL} using the conventions employed to date in the literature (see also [25]). At present, the central point in the primordial shape function defined in (12) is normalized to unity assuming scale invariance, that is, $S(k, k, k) = 1$ with no k -dependence. This produces inconsistent results between models peaking or dipping at this central point (actually along this line); contrast the factor of 7 between the quoted variances of the equilateral and local models for exactly this reason. Furthermore, the definition is not well defined for models which are not scale invariant, such as feature models, and it is simply not applicable to non-Gaussian signals created at late times, such as those induced by cosmic strings or secondary anisotropies.

We, therefore, propose a universally defined bispectrum non-Gaussianity parameter F_{NL} which (i) is a measure of the total observational signal expected for the bispectrum of the model in question and (ii) is normalized for direct comparison with the canonical local model (in particular, with $F_{\text{NL}}^{\text{loc}} = f_{\text{NL}}^{\text{loc}}$ for a given l_{max}). We presume that we have an unnormalized CMB bispectrum $b_{l_1 l_2 l_3}$ accurately calculated for a specific theoretical model over the whole observationally relevant domain $l \leq l_{\text{max}}$. This can be achieved for any model using the separable mode expansion (76) or hierarchical methods [11]. We then define F_{NL} from an adapted version of the estimator (31) with

$$F_{\text{NL}} = \frac{1}{N \bar{N}_{\text{loc}}} \sum_{l_i m_i} \mathcal{G}_{m_1 m_2 m_3}^{l_1 l_2 l_3} b_{l_1 l_2 l_3} \frac{a_{l_1 m_1} a_{l_2 m_2} a_{l_3 m_3}}{C_{l_1} C_{l_2} C_{l_3}}, \quad (102)$$

where N is the appropriate normalization factor for the

given model,

$$N^2 = \sum_{l_i} \frac{B_{l_1 l_2 l_3}^2}{C_{l_1} C_{l_2} C_{l_3}}, \quad (103)$$

and \bar{N}_{loc} is the normalization for the local model *with* $f_{\text{NL}} = 1$,

$$\bar{N}_{\text{loc}}^2 = \sum_{l_i} \frac{B_{l_1 l_2 l_3}^{\text{loc}(f_{\text{NL}}=1)^2}}{C_{l_1} C_{l_2} C_{l_3}}. \quad (104)$$

This F_{NL} estimator will certainly recover the usual f_{NL} for the local model, but it is also clear that it will also equitably compare the total integrated observational bispectrum with that obtained from the $f_{\text{NL}} = 1$ local model. Of course, these definitions presume a sum to a given $l = l_{\text{max}}$ (which should be quoted) but results for primordial models should not depend strongly on this cutoff, unless scale invariance is broken. In any case, diffusion from the transfer functions means that the primordial signal is dying out beyond $l \geq 2000$, so we propose a canonical cutoff at $l_{\text{max}} = 2000$ (which is also relevant in the medium term for the Planck experiment). Late-time anisotropies, such as cosmic strings, do not generically fall off exponentially for $l \geq 2000$, but meaningful comparisons to the local $f_{\text{NL}} = 1$ model can be made with the same definition (102) on this domain, and alternative measures can be proposed elsewhere. In principle, the normalized estimator (102) can also be adapted as a gross measure of the total bispectral signal over the given domain, irrespective of the possible underlying physical model. For example, using the reconstruction from Parseval's theorem (70), the estimator provides a measure of F_{NL}^2 which should then be normalized relative to the total expectation for the local model with $N = N_{\text{loc}}$ in (102).

If the CMB bispectrum $b_{l_1 l_2 l_3}$ is not known precisely for the primordial model under study, then the normalization factor N in (103) can still be estimated using the shape function $S(k_1, k_2, k_3)$. Primordial and CMB correlators are closely related, so one can obtain a fairly accurate approximation to the relative normalizations above (103) and (104) from [11]

$$\tilde{N}^2 = \int_{\mathcal{V}_k} S^2(k_1, k_2, k_3) w(k_1, k_2, k_3) d\mathcal{V}_k, \quad (105)$$

where the appropriate weight function was found to be $w(k_1, k_2, k_3) \approx 1/(k_1 + k_2 + k_3)$ and the domain \mathcal{V}_k is given by $k_1, k_2, k_3 \leq k_{\text{max}}(l_{\text{max}})$ [refer to the discussion before (28) in Sec. II]. Here, we note that $N/\bar{N}_{\text{loc}f_{\text{NL}}=1} \approx \tilde{N}/\tilde{N}_{\text{loc}f_{\text{NL}}=1}$. Using this primordial shape function normalization \tilde{N} in Ref. [11] led to a comparable definition of $\tilde{f}_{\text{NL}} \approx F_{\text{NL}}$, which can be useful for making fairly accurate projections of non-Gaussianity or for renormalizing f_{NL} constraints for different models into more compatible F_{NL} constraints.

Below we renormalize published and model-dependent constraints on f_{NL} into the integrated bispectral measure F_{NL} , using the expression (102) together with accurate calculations of $B_{l_1 l_2 l_3}$ for each model [7,10,30,10]:

$$-4 < f_{\text{NL}}^{\text{local}} < 80 \Rightarrow -4 < F_{\text{NL}}^{\text{local}} < 80 \quad (106)$$

$$-125 < f_{\text{NL}}^{\text{equil}} < 435 \Rightarrow -24 < F_{\text{NL}}^{\text{equil}} < 83 \quad (107)$$

$$-375 < f_{\text{NL}}^{\text{warm}} < 37 \Rightarrow -93 < F_{\text{NL}}^{\text{warm}} < 9 \quad (108)$$

$$-369 < f_{\text{NL}}^{\text{ortho}} < 71 \Rightarrow -114 < F_{\text{NL}}^{\text{ortho}} < 22. \quad (109)$$

Note the much more consistent variance found for the different models with F_{NL} , thus aiding direct comparison, as well as the exact correspondence for the local model to which it is normalized.

V. CMB MAP SIMULATIONS FOR GENERAL BISPECTRA

A. Mapmaking with separable shape functions and its limitations

In the limit of weak non-Gaussianity, an algorithm to produce non-Gaussian CMB simulations with a given power spectrum and bispectrum for separable primordial shapes was described in Ref. [31]. Here, we present it in a more transparent notation, generalizing the method to non-separable shapes using the mode decompositions of the previous sections. We note that while several other methods have been produced for simulating maps with non-Gaussianity of a specific form, such as for local non-Gaussianity [32–34] or with a particular equal- l bispectrum [35,36], the approach presented here is completely general. A by-product is that the generalized approach is more robust and reliable, because the polynomial mode functions are better behaved than the separable approximations which have been previously employed. In this algorithm the non-Gaussian components of the CMB multipoles are obtained using the following formula:

$$a_{lm}^{\text{NG}} = \frac{1}{6} \sum_{l_i m_i} B_{ll_2 l_3} \begin{pmatrix} l & l_2 & l_3 \\ m & m_2 & m_3 \end{pmatrix} \frac{a_{l_2 m_2}^{G*}}{C_{l_2}} \frac{a_{l_3 m_3}^{G*}}{C_{l_3}}, \quad (110)$$

where a_{lm}^G is the Gaussian part of the CMB multipoles, generated using the angular power spectrum C_l , while $B_{ll_2 l_3}$ is the given bispectrum of the theoretical model for which simulations are required. Although Eq. (110) is completely general, as before, its numerical evaluation is only computationally affordable for bispectra that can be written in separable form. We have emphasized already that separability results in a reduction of the computational cost of the estimator (31) from $O(l_{\text{max}}^5)$ to $O(l_{\text{max}}^3)$ operations; the

same argument applies here allowing a rewriting of (110) into an equivalent form in pixel space (see below).

The limitation dictated by separability is clearly overcome by using our eigenfunction representation for the bispectrum (71). The basic idea is to start by expanding an arbitrary bispectrum shape S using the separable polynomial decomposition S_N until a good level of convergence is achieved and then to substitute the mode decomposition into (110). The accuracy of convergence is parametrized in terms of the correlation $\tilde{C}(S, S_N)$ between the original nonseparable shape and the eigenmode expansion, as defined previously (28). Note that this convergence can also be checked more accurately using the full Fisher matrix correlation on the CMB bispectra $\mathcal{C}(b_{l_1 l_2 l_3}, b_{l_1 l_2 l_3}^N)$, calculated using the separable approach (76) or else accurate hierarchical approaches [11]. In previous sections (see Fig. 9), we have noted how rapid this convergence is for well-behaved nonseparable shapes, such as DBI inflation (or at late times with cosmic strings).

In addition to the bispectrum separability requirement, there is an important further caveat which can prevent the straightforward implementation of the algorithm (110). By construction, terms $\mathcal{O}(f_{\text{NL}}^2)$ and higher are not explicitly controlled. Following the discussion in [37] we can write the connected N-point functions as:

$$\langle a_{l_1 m_1}^* a_{l_2 m_2}^* \rangle = [C_{l_1} + f_{\text{NL}}^2 C_{l_1}^{\text{NG}}] \quad (111)$$

$$\langle a_{l_1 m_1}^* a_{l_2 m_2} a_{l_3 m_3} \rangle = [f_{\text{NL}} B_{l_1 l_2 l_3} + \mathcal{O}(f_{\text{NL}}^3)] \quad (112)$$

$$\langle a_{l_1 m_1} a_{l_2 m_2} a_{l_3 m_3} \dots a_{l_N m_N} \rangle = \mathcal{O}(f_{\text{NL}}^3). \quad (113)$$

Thus the condition that the map has the power spectrum C_l specified in the input will only be satisfied if the power spectrum of the non-Gaussian component in (111) remains small. Since this method does not control $\mathcal{O}(f_{\text{NL}}^2)$ terms, one has to ascertain that spuriously large $C_{l_1}^{\text{NG}}$ contributions do not affect the overall power spectrum significantly. It turns out that this effect plagues current map simulations if the standard separable expressions for the local and equilateral bispectra are directly substituted into (110), as we now demonstrate.

In Sec. II, we showed how the reduced bispectrum could be written explicitly in separable form for the local model $b_{l_1 l_2 l_3}^{\text{local}}$ in (20) and for the equilateral model $b_{l_1 l_2 l_3}^{\text{equil}}$ in (23). These were expressed in terms of one-dimensional convolution integrals between the transfer functions $\Delta_l(k)$ and powers of the power spectrum $P(k)$, with α_l , β_l , γ_l , δ_l corresponding to const, $P(k)$, $P(k)^{1/3}$, $P(k)^{2/3}$, respectively [refer to Eqs. (21) and (24)]. Just as we did with the f_{NL} estimator (80), we can sum the a_{lm}^G 's from the Gaussian maps with these functions to create filtered maps,

$$\begin{aligned}
M_\alpha(x, \hat{\mathbf{n}}) &\equiv \sum_{lm} \alpha_l(x) a_{lm}^G \frac{Y_{lm}(\hat{\mathbf{n}})}{C_l}, & M_\gamma(x, \hat{\mathbf{n}}) &\equiv \sum_{lm} \gamma_l(x) a_{lm}^G \frac{Y_{lm}(\hat{\mathbf{n}})}{C_l}, \\
M_\beta(x, \hat{\mathbf{n}}) &\equiv \sum_{lm} \beta_l(x) a_{lm}^G \frac{Y_{lm}(\hat{\mathbf{n}})}{C_l}, & M_\delta(x, \hat{\mathbf{n}}) &\equiv \sum_{lm} \delta_l(x) a_{lm}^G \frac{Y_{lm}(\hat{\mathbf{n}})}{C_l}
\end{aligned} \tag{114}$$

From products of these maps in pixel space, we can now obtain explicit expressions for the non-Gaussian a_{lm}^{NG} 's in these two separable cases [compare with the bispectrum expressions (20) and (23)]:

$$\begin{aligned}
a_{lm}^{\text{local}} &= \int dx x^2 \left[\frac{2}{3} \beta_l(x) \int d^2 \hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}}) M_\alpha(x, \hat{\mathbf{n}}) M_\beta(x, \hat{\mathbf{n}}) + \frac{1}{3} \alpha_l(x) \int d^2 \hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}}) M_\beta(x, \hat{\mathbf{n}}) M_\beta(x, \hat{\mathbf{n}}) \right], \\
a_{lm}^{\text{equil}} &= 2 \int dx x^2 \left[-2 \beta_l(x) \int d^2 \hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}}) M_\alpha(x, \hat{\mathbf{n}}) M_\beta(x, \hat{\mathbf{n}}) - \alpha_l(x) \int d^2 \hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}}) M_\beta(x, \hat{\mathbf{n}}) M_\beta(x, \hat{\mathbf{n}}) \right. \\
&\quad - 2 \delta_l(x) \int d^2 \hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}}) M_\delta(x, \hat{\mathbf{n}}) M_\delta(x, \hat{\mathbf{n}}) + 2 \gamma_l(x) \int d^2 \hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}}) M_\beta(x, \hat{\mathbf{n}}) M_\delta(x, \hat{\mathbf{n}}) \\
&\quad \left. + 2 \beta_l(x) \int d^2 \hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}}) M_\gamma(x, \hat{\mathbf{n}}) M_\delta(x, \hat{\mathbf{n}}) + 2 \delta_l(x) \int d^2 \hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}}) M_\beta(x, \hat{\mathbf{n}}) M_\gamma(x, \hat{\mathbf{n}}) \right]. \tag{115}
\end{aligned}$$

In the top panel of Fig. 12 we consider the contribution to the final C_l^{NG} from the various terms appearing in Eq. (115) taken separately. For example, we build a set of multipoles from the term $\int dx x^2 \frac{2}{3} \beta_l(x) \times \int d^2 \hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}}) M_\alpha(x, \hat{\mathbf{n}}) M_\beta(x, \hat{\mathbf{n}})$ and compute the resulting power spectrum, neglecting all other terms, and so forth. We then compare the power spectra of the non-Gaussian part to the input power spectrum of the Gaussian part for $f_{\text{NL}} = 100$. Our procedure underlines what was pointed out in [37]: some terms in the separable approximations to both the local and equilateral shapes produce spurious divergences at low l 's that are large enough to affect the final power spectrum of the map. More precisely, as can be seen in Fig. 12, the biggest problems come from the terms $\int dx x^2 \frac{2}{3} \beta_l(x) \int d^2 \hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}}) M_\alpha(x, \hat{\mathbf{n}}) M_\beta(x, \hat{\mathbf{n}})$ and $\int dx x^2 \beta_l(x) \int d^2 \hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}}) M_\gamma(x, \hat{\mathbf{n}}) M_\delta(x, \hat{\mathbf{n}})$. In Ref. [37], it was pointed out that the problem can be circumvented for the local model by modifying the expression (115) of a_{lm}^{local} so as to eliminate the pathological term, while leaving the final bispectrum of the map preserved with a change of weight for the remaining term. The same approach can also be applied to the equilateral case, leaving new tailored expressions for the non-Gaussian parts:

$$\begin{aligned}
a_{lm}^{\text{local}} &= \int dx x^2 \alpha_l(x) \int d^2 \hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}}) M_\beta(x, \hat{\mathbf{n}}) M_\beta(x, \hat{\mathbf{n}}), \\
a_{lm}^{\text{equil}} &= 2 \int dx x^2 \left[-3 \alpha_l(x) \int d^2 \hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}}) M_\beta(x, \hat{\mathbf{n}}) M_\beta(x, \hat{\mathbf{n}}) \right. \\
&\quad - 2 \delta_l(x) \int d^2 \hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}}) M_\delta(x, \hat{\mathbf{n}}) M_\delta(x, \hat{\mathbf{n}}) \\
&\quad \left. + 2 \gamma_l(x) \int d^2 \hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}}) M_\beta(x, \hat{\mathbf{n}}) M_\delta(x, \hat{\mathbf{n}}) \right]. \tag{116}
\end{aligned}$$

It is easy to verify that these modified expressions produce the correct bispectra in the final maps, they are numerically stable and so allow the simulation of non-Gaussian maps of

the local and equilateral type with given power spectrum and bispectrum. However, one can see how the necessity of looking at all the individual terms in the equations defining a_{lm}^{NG} , and the need to produce suitable modifications of the original formulae, means that the algorithm loses its generality. If additional shapes are considered then, in principle, different separation schemes could well encounter the problems outlined above. The good news is that the full generality of this approach is regained when the separation of the original shape is done using the eigenmode expansion introduced in this paper.

B. Mapmaking from arbitrary primordial shape functions

In order to see why this happens, it is useful to write down the equation for a_{lm}^{NG} in terms of our polynomial expression. Since we can decompose the three-point functions both at early and late times it is actually possible to generate a map in two different ways. The closest method to the ‘‘standard’’ one, just outlined above, is the one that start from the early time decomposition. In this case the primordial shape $S(k_1, k_2, k_3)$ is written as:

$$S(k_1, k_2, k_3) = \sum_n \alpha_n^{\mathcal{Q}} \mathcal{Q}_n = \sum_{pqr} \alpha_{pqr}^{\mathcal{Q}} q_p(k_2) q_q(k_1) q_r(k_3), \tag{117}$$

where the $\mathcal{Q}_n(k_1, k_2, k_3)$ are formed from products of the tetrahedral polynomials $q_p(k)$ given in (54) and the $\alpha_n^{\mathcal{Q}} \leftrightarrow \alpha_{pqr}^{\mathcal{Q}}$ are the coefficients of the eigenmode expansion for a given shape (recall the convenience of ordering the pqr with a single label n). The reduced angular bispectrum $b_{l_1 l_2 l_3}$ is obtained, as was shown in (76), by linearly projecting the primordial shape on the sphere using radiation transfer functions:

$$\begin{aligned}
 b_{l_1 l_2 l_3} &= \Delta_{\Phi}^2 f_{\text{NL}} \sum_n \alpha_n^{\mathcal{Q}} \int x^2 dx \left[\left(\frac{2}{\pi} \int dk_1 q_p(k_1) \Delta_{l_1}(k_1) j_{l_1}(k_1 x) \right) \left(\frac{2}{\pi} \int dk_2 q_r(k_2) \Delta_{l_2}(k_2) j_{l_2}(k_2 x) \right) \right. \\
 &\quad \left. \times \left(\frac{2}{\pi} \int dk_3 q_s(k_3) \Delta_{l_3}(k_3) j_{l_3}(k_3 x) \right) + 5 \text{ permutations} \right] \\
 &= \Delta_{\Phi}^2 f_{\text{NL}} \sum_n \alpha_n^{\mathcal{Q}} \int x^2 dx q_{\{p\}}^{l_1} q_{\{r\}}^{l_2} q_{\{s\}}^{l_3},
 \end{aligned}$$

$$\text{where } q_p^l(x) = \frac{2}{\pi} \int dk q_p(k) \Delta_l(k) j_l(kx),$$

(118)

Substituting Eq. (118) into (110), and using the standard technique of decomposing the integrals into tractable products of one-dimensional integrals, after some algebra, we obtain the general expression for a_{lm}^{NG} :

$$\begin{aligned}
 a_{lm}^{\text{NG}} &= \frac{1}{18} \sum_{pqr \rightarrow n} \alpha_n^{\mathcal{Q}} \int dx x^2 q_p^l(x) \\
 &\quad \times \int d^2 \hat{\mathbf{n}} Y_l^{m*}(\hat{\mathbf{n}}) M_q^G(\hat{\mathbf{n}}, x) M_r^G(\hat{\mathbf{n}}, x), \quad (119)
 \end{aligned}$$

where the $M_p^G(\hat{\mathbf{n}}, x)$ are filtered maps found by summing a set of Gaussian a_{lm}^G 's with the convolved tetrahedral polynomial q_p^l [refer to Eq. (80)]

$$\begin{aligned}
 M_p^G(\hat{\mathbf{n}}, x) &= \sum_{lm} q_p^l \frac{a_{lm}^G Y_{lm}}{C_l} \\
 &= \sum_{lm} \left[\frac{2}{\pi} \int q_p(k) \Delta_l(k) j_l(kx) dk \right] \frac{a_{lm}^G Y_{lm}(\hat{\mathbf{n}})}{C_l}. \quad (120)
 \end{aligned}$$

The $M_p^G(\hat{\mathbf{n}}, x)$ and $q_p^l(x)$ are now the analogues of the $M_\alpha(x, \hat{\mathbf{n}})$, $M_\beta(x, \hat{\mathbf{n}})$, $M_\gamma(x, \hat{\mathbf{n}})$, $M_\delta(x, \hat{\mathbf{n}})$, and α_l , β_l , γ_l , δ_l defined above. Note that it is not strictly necessary here to include cyclic permutations in (119) running over the indices $\{p, q, r\}$, as these are incorporated automatically. Further efficiencies can be achieved by exploiting the freedom to reorder terms in (119), taking out the polynomial q_p of highest order and convolving the maps with the two lower order polynomials; this is not necessitated by stability requirements (see below).

In principle, the numerical instabilities which cause problems for the standard separable approximations, could now affect the angular integrals over the polynomials q_p given in (119). However, as shown in Fig. 12, this is not the case. The key point is that all the functions $q_p^l(x)$ now scale as $\frac{1}{l(l+1)}$ (see Fig. 12), that is, in the same way as the nonpathological $\beta_l(x)$ term in the standard local and equilateral decompositions. For this reason the spherical harmonic projection of a product of two $M_p^G(\hat{\mathbf{n}}, x)$ maps is expected to have similar scaling properties as the term $\int d^2 n Y_l^{m*}(\hat{\mathbf{n}}) M_\beta(x, \hat{\mathbf{n}})^2$. This last integral was previously shown to be stable at low multipoles, as discussed for the local case in Ref. [37]). Thus all the integrals in Eq. (119) are going to be well-behaved at low l 's. Since the shape-dependent information is in the coefficients of the expansion

in $\alpha_n^{\mathcal{Q}}$ and not in the precomputed $q_p^l(x)$ modes, we are able to produce numerically stable results for any possible




FIG. 12 (color online). Convergence properties of the standard separable functions used in the literature to represent local and equilateral models (top panel). Here the functions have been convolved with the transfer functions Δ_l required in the bispectrum estimator or mapmaking algorithms. Note the poor scaling and divergence at low l for two of the separable combinations with the resulting power spectrum from non-Gaussianity rising to compete with the CMB power spectrum C_l 's ($f_{\text{NL}} = 1$). This poor scaling is contrasted with results for the tetrahedral polynomials $q_n(k)$ (lower panel). These remain bounded and roughly scale invariant over the full multipole range, even for very high order polynomials.

shape. Numerical tests were carried out for the local and equilateral case, confirming the previous statements. We suggest, therefore, that the eigenmode expansion provides a numerically stable and efficient means by which to generalize the algorithm in Ref. [31] to nonseparable bispectrum shapes.

C. Simulated maps from general CMB bispectra

It is useful to recap the discussion above by using separable mode expansions to create simulated maps at late times from a given CMB power spectrum C_l and




FIG. 13 (color online). Simulated non-Gaussian CMB maps from the equilateral model, created using the primordial mapmaking method (119). The upper panel shows a map simulation with $f_{\text{NL}} = 400$ (barely discernible from the underlying Gaussian template), the middle panel shows a map with a large NG signal with $f_{\text{NL}} = 4000$, while the lower panel shows the $f_{\text{NL}} = 400$ case above in a WMAP-realistic context using the KQ75 mask and with inhomogeneous noise added.

reduced bispectrum $b_{l_1 l_2 l_3}$. As before with the f_{NL} estimator, removal of the convolution with transfer functions, makes the late-time method much simpler and more transparent. We begin with the same expression

$$a_{lm} = a_{lm}^G + f_{\text{NL}} a_{lm}^{\text{NG}}, \quad (121)$$

where

$$a_{lm}^{\text{NG}} = \int d\hat{\mathbf{n}} \sum_{l_1, l_2, m_1, m_2} b_{l_1 l_2 l_3} Y_{l_1 m_1}(\hat{\mathbf{n}}) \frac{a_{l_2 m_2} Y_{l_2 m_2}(\hat{\mathbf{n}})}{C_{l_2}} \times \frac{a_{l_3 m_3} Y_{l_3 m_3}(\hat{\mathbf{n}})}{C_{l_3}}. \quad (122)$$

Now we expand the CMB bispectra using the eigenmode decomposition using weight functions motivated by the estimator [refer to (89)]

$$\frac{v_{l_1} v_{l_2} v_{l_3}}{\sqrt{C_{l_1} C_{l_2} C_{l_3}}} b_{l_1 l_2 l_3} = \sum \alpha_n \bar{Q}_n, \quad (123)$$

where v_l is a separable weight factor chosen to remove




FIG. 14 (color online). Simulated maps for non-Gaussian models using the late-time mapmaking method (124); this only includes the a_{lm}^{NG} contribution. The upper panel shows a non-Gaussian CMB map from cosmic strings obtained using the analytic expression for the string bispectrum (29). The lower panel shows a simulated non-Gaussian map for an equilateral model. When added to its Gaussian counterpart map from a_{lm}^G at an amplitude $f_{\text{NL}} = 600$, this equilateral map was used for the bispectrum recovery illustrated in Figs. 16 and 17. Note the red color cast from negative f_{NL} and blue from positive.

scaling from the CMB bispectrum, improving decomposition convergence. These weight factors are important for this late-time mapmaking method because they help remove the scaling of the $\sqrt{C_l}$ term in the $\bar{M}_p^G(\hat{\mathbf{n}})$ filtered maps, making their power spectrum flatter (the analogue of the problem discussed above for primordial map simulations). We can rewrite the non-Gaussian part as

$$a_{lm}^{\text{NG}} = \sum_{pqr \rightarrow n} \alpha_n^{\mathcal{Q}} \frac{q_{lp}(l)}{\sqrt{C_l} v_l} \int d\hat{\mathbf{n}} Y_{lm}(\hat{\mathbf{n}}) \bar{M}_q^G(\hat{\mathbf{n}}) \bar{M}_r^G(\hat{\mathbf{n}}), \quad (124)$$

where the $\bar{M}_p^G(\hat{\mathbf{n}})$ are defined in (92) and summed with Gaussian a_{lm}^G 's.

This method is straightforward to implement for a given theoretical $b_{l_1 l_2 l_3}$ and it is highly efficient. For example, it can produce simulated maps in 64 seconds for $l = 500$ with 16 eigenmodes (6 polynomials). Examples of maps produced with this method for the equilateral model can be seen in Fig. 13. It has the advantage that, as it depends only on the CMB bispectra, it can also be used to simulate maps for bispectra produced by late-time effects, like cosmic strings, gravitational lensing and secondary anisotropies. Plots of the non-Gaussian part of simulated maps can be seen in Fig. 14 for the nonseparable DBI inflation and cosmic string models.

VI. DIRECT COMPARISON OF BISPECTRUM ESTIMATORS

We have developed two complete numerical pipelines, implementing the eigenmode decomposition methods described in the previous sections. For a generic primordial shape $S(k_1, k_2, k_3)$ or a given CMB bispectrum $b_{l_1 l_2 l_3}$, an expansion in monomial symmetric polynomials \mathcal{Q}_n is performed followed by the generation of non-Gaussian map simulations. Bispectrum estimators are then applied to the map simulations in order to verify that the input f_{NL} can be properly recovered together with the expected variance. Both ‘‘early time decomposition’’ and ‘‘late-time decomposition’’ f_{NL} estimators have been fully implemented. The former starts from an expansion of the primordial shape $S(k_1, k_2, k_3)$ in Fourier space while the latter starts from an expansion of the reduced angular bispectrum $b_{l_1 l_2 l_3}$ in harmonic space, where in the second case radiation transfer functions have already been included in the expression for $b_{l_1 l_2 l_3}$. The redundancy provided by the two alternative pipelines provide a further check of the reliability of the final results.

Since our purpose in this paper is to introduce the eigenmode expansion method, and to test its implementation, we will primarily apply our pipelines to map simulations, leaving detailed analysis of real data sets over a wider range of models for future publication [29]. As this is a proof-of-concept paper, we will mainly limit ourselves to the study of the simple equilateral family of models. This is because it is already well studied in the literature

(see e.g. [3,25,31]), which enables a useful comparison between the outcome of our numerical pipelines and previously published results for the equilateral shape. Moreover, the equilateral case does not require sophisticated noise analysis, unlike the local model. However, we will briefly consider other nonseparable models outlined earlier in the introduction, such as the related DBI model and the cosmic string bispectrum. We note that from the point of view of the eigenmode decomposition, the formal separability of the equilateral shape is irrelevant; it does not cause early termination of the expansion series which is nearly identical to the nonseparable DBI model (see Fig. 10). Having established the reliability of the eigenmode expansion method here, in a forthcoming publication [29] we will apply it to the study of families of nonseparable shapes using WMAP5 data.

A. Simulated observational maps

Using the algorithm described in Sec. V, we generated a set of 100 equilateral CMB maps with both the primordial and late-time decomposition pipelines. We worked at roughly WMAP resolution with $l_{\text{max}} = 500$ and HEALpix [38] $n_{\text{side}} = 512$, corresponding to a pixel number $N_{\text{pix}} \approx 10^6$. We then applied both our primordial and late-time estimators to both our primordial and late-time sets of simulated maps in all combinations. We found that in all cases the mapmaking methods gave consistent results, producing simulated maps from which the correct f_{NL} could be reliably recovered with the correct variance. Results for both primordial and late-time estimators on the same set of 50 equilateral maps (with and without the mask and inhomogeneous noise) with $f_{\text{NL}} = 300$ can be seen in Fig. 15. We observe that the two estimators produce consistent results on the same maps. Of course, there is some small variation between the results as the two estimators can be regarded to be independent but this proved always to be well within the variance.

In addition, we extracted the equilateral configurations B_{lll} of the bispectrum from the maps and compared the average over all the simulations to the semianalytic expectations obtained from the standard decomposition of the equilateral shape in terms of $\alpha_l(x)$, $\beta_l(x)$, $\gamma_l(x)$, $\delta_l(x)$ [refer to Eqs. (21) and (24)]. The recovered equilateral bispectrum values were in very good agreement between the semianalytic prediction from the standard α , β , γ , δ decomposition and the simulations, based on our eigenmode expansion, thus showing consistency with previous approaches.

Finally, we reiterate that this general approach to map simulation was highly efficient, producing Planck resolution maps for the equilateral model on short time scales. This made estimator validation through Monte Carlo simulations easily achievable with only modest resources. For other well-behaved bispectra, such as the cosmic string model, the general method proved robust. Examples of




FIG. 15 (color online). Recovery of f_{NL} from 50 simulated maps of the equilateral model, showing a direct map-by-map comparison between the primordial estimator (83) (blue) and the CMB estimator (95) (red). Ideal map recovery is shown in the top panel, while recovery for WMAP-realistic maps is shown below with beam, inhomogeneous noise and mask included (BNM). Both methods recovered the input $f_{\text{NL}} = 300$ with a variance of approximately $\Delta f_{\text{NL}} = 105$ (clean) and 150 (BNM). Note the overall consistency of the two independent estimators with a significantly lower variance evident between the methods $\Delta f_{\text{NL}} = 30$ (clean) and $\Delta f_{\text{NL}} = 103$ (BNM).

nonseparable maps already have been discussed and shown in Fig. 14.

B. Primordial and late-time f_{NL} estimators

Choosing an input value $f_{\text{NL}} = 300$ for the sets of equilateral map simulations described above, we compared results from both the primordial and late-time bispectrum estimators. In order to verify the consistency of the two methods we selected the late-time map sets and applied both estimators to it. The tests were performed starting from a noiseless full-sky map and then more realistic simulations were used, including partial sky-coverage and an anisotropic noise component. The rms noise was obtained by coadding WMAP V and W channel using the same scheme as the one adopted for non-Gaussian analysis by the WMAP team [3]. The sky-coverage was done using the KQ75 mask, also adopted by the WMAP5 team for their f_{NL} analysis. Only the approximate form (34) of the estimator is used, and not the full form (33) including the full covariance matrix and a linear term. Note however that this approximation has been demonstrated in several previous studies to work well for equilateral shapes. Moreover, for our purposes the approximate nearly-optimal estimator is all we need since it contains all the

dependence on the theoretical ansatz and thus all the dependence on our eigenmode expansion, which is the primary concern for this initial validation process.

We compared the f_{NL} recovered from each map using the two methods, as well as the final averages and variances. The variances were compared to expectations from Fisher matrix forecasts obtained both from our eigenmode expansion and from the “standard” $\alpha_l, \beta_l, \gamma_l, \delta_l$ decomposition of the equilateral shape used to date in other non-Gaussian analysis. In all cases the results were internally consistent and in agreement with Fisher matrix expectations, as summarized in Table I. This led us to conclude that the eigenmode expansion method appears to be a reliable way to produce non-Gaussian CMB simulations and f_{NL} estimators for primordial models, whether separable or otherwise.

Having verified the two estimator’s performance on simulated equilateral maps we then applied both of them to the WMAP5 data, coadding the V and W channels as discussed above. The primordial estimator obtained the result $-174 < f_{\text{NL}}^{\text{equil}} < 434$, which is consistent with the existing constraints obtained using standard separable primordial approach (given the caveat that a number of these results have now been superseded [10]). The constraint

TABLE I. Results obtained from the application of the primordial and late-time estimators as described in the text. In the first two columns, labeled by “Ideal simulations,” we consider ideal full-sky noiseless measurements, while in the last two columns, labeled by “WMAP5 simulations” we include noise and sky coverage in order to simulate a WMAP5-realistic experiment (see text for further explanation). We apply both estimators to a single set of maps, in this case created using the late-time mode expansion approach. In the last row, we calculate the difference between the f_{NL} recovered by the two techniques, map-by-map for 100 maps, and report the final internal standard deviation between the methods.

	Ideal simulations		WMAP5 simulations	
	Average	St. Dev.	Average	St. Dev.
Primordial estimator	292.9	104.8	297.7	152.1
Late-time estimator	300.6	104.9	278.7	160
Internal st. dev.	38.5		102.6	

using the late-time estimator was $-156 < f_{\text{NL}}^{\text{equil}} < 443$ which is consistent with the primordial result. As the late-time estimator can be regarded as independent of the primordial estimator, some small variation is to be expected as we have seen already in Fig. 15. The difference between the estimators is consistent with the internal variance of 103 noted in Table I for equilateral map simula-

tions in a WMAP-realistic context. We conclude that both the primordial and late-time estimators appear to be performing up to expectation.

Using the method described in Eq. (101), we can endeavour to recover the full bispectrum from a given map. To illustrate this capability, we created a single map realization from an equilateral bispectra with $f_{\text{NL}} = 600$, that is, a map with a 4σ non-Gaussian signal. We then used the late-time estimator to recover the $\tilde{\beta}_n^{\text{Q}}$ and $\tilde{\beta}_n^{\text{R}}$ mode coefficients described in (89). Recall that for results of sufficient significance, the $\tilde{\beta}_n^{\text{R}}$ should approximate the original theoretical model coefficients $\tilde{\alpha}_n^{\text{R}}$, that is, those used to generate the simulated map. We estimated the variance in each of the eigenmodes by applying the same method to 100 Gaussian simulations. The results for the orthonormal coefficients $\tilde{\alpha}_n^{\text{R}}$ are plotted in Fig. 16 for both ideal maps and for maps with inhomogeneous noise added and a mask applied. We see that we recover the first 7 modes well from the ideal map but the results from the map containing noise and mask are somewhat less encouraging. Clearly, more work is required to control noise and mask effects at higher mode numbers. By plotting the 3D bispectra from the reconstructions, see Fig. 17, we observe that it is possible to recover the main acoustic peak and some basic features of the CMB bispectra. We will address the challenging




FIG. 16 (color online). Recovered spectral coefficients $\tilde{\beta}_n^{\text{R}}$ from the late-time estimator (89) from a single map simulation for an equilateral model with $f_{\text{NL}} = 600$ (or normalized relative to the local model $F_{\text{NL}} \approx 110$); NG map simulation shown in the lower panel of Fig. 14. In both panels, the original $\tilde{\alpha}_n^{\text{R}}$ decomposition coefficients for the theoretical model are shown for comparison (blue). In the upper panel, the $\tilde{\beta}_n^{\text{R}}$ coefficients recovered from the single realization are shown, with error bars (2σ) estimated from 100 Gaussian maps. In the lower panel, the $\tilde{\beta}_n^{\text{R}}$ are recovered in a WMAP-realistic context using the KQ75 mask with inhomogeneous noise added. Note that the $\tilde{\beta}_n^{\text{R}}$ provide a remarkably good fit to the $\tilde{\alpha}_n^{\text{R}}$ given the significance of the non-Gaussian signal.




FIG. 17 (color online). Recovered 3D bispectrum using the late-time mode decomposition method (101) from a single map simulation for an equilateral model with $f_{\text{NL}} = 600$ (or relative to the local model, $F_{\text{NL}} \approx 110$); this figure shows the reconstruction of the bispectrum from the $\bar{\beta}_n^R$ expansion modes illustrated in Fig. 16. The left panel represents the original theoretical bispectrum used to construct a single realization of the map (just like those shown in Fig. 14). The middle panel represents the recovered bispectrum from the ideal map, while the right panel represents the recovery in a WMAP-realistic context using the KQ75 mask with inhomogeneous noise added. The main feature of the bispectrum, that is, the primary acoustic peak appears to be evident even in the noisy cut-sky case, given the significant non-Gaussian signal.

issues associated with bispectrum reconstruction in greater detail elsewhere [29].

VII. CONCLUSIONS

We have now implemented two comprehensive and independent pipelines for the analysis and estimation of general primordial or CMB bispectra. Both methods are based on dual mode expansions, exploiting a complete orthonormal eigenmode basis to efficiently decompose arbitrary bispectra into a separable polynomial expansion. These separable mode expansions, whether at late or early times, allow a reduction of the computational overhead to easily tractable levels, whether calculating the reduced bispectrum $b_{l_1 l_2 l_3}$, generating Planck resolution non-Gaussian map simulations, or directly estimating f_{NL} from simulations or real data sets. The method exploits the smoothness of the pattern of acoustic peaks observed in calculations reviewing all well-behaved primordial models, implying the rapid convergence of the corresponding mode expansions. While most calculations are performed in the separable basis, a final rotation of mode coefficients into the orthonormal frame allows for a simple interpretation of the contributions to f_{NL} using Parseval's theorem. In fact, the completeness of the orthonormal eigenmodes means, in principle, that it is straightforward to extract and reconstruct the full CMB bispectrum from the data, assuming the presence of a sufficiently significant non-Gaussian signal.

The main purpose of this paper has been to present a detailed theoretical framework for f_{NL} estimation using separable eigenmode expansions, irrespective of the spe-

cific polynomials or other basis functions employed. However, we have also presented some numerical results from the pipelines we have implemented, chiefly for the equilateral model where there are extensive published results for direct comparison. An important milestone for the validation of this approach has been the development of a robust and reliable mode expansion method for generating map simulations from arbitrary bispectra. While generalizing previous methods applied to specific separable cases, we noted that the scale invariance of the polynomial expansion modes eliminates numerical instabilities that previously had to be circumvented on a case-by-case basis. Given convergent mode expansions for well-behaved bispectra, high resolution map simulations for a wide variety of models can easily and efficiently be generated, with several examples illustrated here including late-time cosmic strings. The many map simulations created for the equilateral model with both primordial and late-time methods showed consistency in expected variance and f_{NL} recovery.

The primordial and late-time f_{NL} estimators using mode expansions were tested successfully on the simulated equilateral maps, matching expectations for semianalytic Fisher matrix forecasts and providing consistent unbiased results for f_{NL} . This was achieved for both ideal maps and in a WMAP-realistic context, incorporating beams, anisotropic noise and a mask. Application of the estimators to the WMAP5 data gave constraints on the equilateral model consistent with each other and previously published results. These encouraging results suggest that the approach will provide a robust and general framework for f_{NL} estimation for the wide variety of nonseparable models which remain to be constrained [11]. For single equilateral

map simulation with $f_{\text{NL}} = 600$, we were able to demonstrate a reasonable correspondence between the theoretical and recovered mode expansion coefficients, while also being able to recover key features of the full CMB bispectrum. However, a detailed discussion of such prospects has been left for a future publication [29]. We have also left aside for discussion elsewhere a more sophisticated treatment of sky cuts and inhomogeneous noise, which is more important for the analysis of the local model, as well as the potential for incorporating polarization data. Challenges remain for the full implementation of the primordial and late-time pipelines at Planck resolution, but the generality and robustness of this methodology suggests that it should prove to be a useful tool for exploring and constraining a much wider class of non-Gaussian models.

ACKNOWLEDGMENTS

We are very grateful for informative discussions with Xingang Chen and Kendrick Smith and we have also benefitted from useful conversations with Martin Bucher, Anthony Challinor, Olivier Forni, Enrique Martinez-Gonzalez, and Bartjan van Tent. The ongoing development of these methods has been regularly reported at Planck Working Group 4 (Non-Gaussianity), where we have been grateful for feedback and to learn of related work. We are also very grateful to Andrey Kaliazin for his help with developing and implementing the numerical code. Simulations were performed on the COSMOS supercomputer (an Altix 4700) which is funded by STFC, HEFCE, and SGI. J. R. F., M. L. and E. P. S. were supported by STFC Grant No. ST/F002998/1 and the Centre for Theoretical Cosmology.

-
- [1] Juan Martin Maldacena, *J. High Energy Phys.* **05** (2003) 013.
 - [2] Viviana Acquaviva, Nicola Bartolo, Sabino Matarrese, and Antonio Riotto, *Nucl. Phys.* **B667**, 119 (2003).
 - [3] E. Komatsu *et al.*, *Astrophys. J. Suppl. Ser.* **180**, 330 (2009).
 - [4] Amit P.S. Yadav and Benjamin D. Wandelt, *Phys. Rev. Lett.*, **100**, 181301 (2008).
 - [5] A. Curto *et al.*, *Mon. Not. R. Astron. Soc.* **393**, 615 (2009).
 - [6] Oystein Rudjord *et al.*, *Astrophys. J.* **701**, 369 (2009).
 - [7] Kendrick M. Smith, *J. Cosmol. Astropart. Phys.* **09** (2009) 006.
 - [8] (ESA Planck Consortium), “Planck: The Scientific Programme (Planck Blue Book)” (unpublished).
 - [9] Eiichiro Komatsu, David N. Spergel, and Benjamin D. Wandelt, *Astrophys. J.* **634**, 14 (2005).
 - [10] Leonardo Senatore, Kendrick M. Smith, and Matias Zaldarriaga, *J. Cosmol. Astropart. Phys.* **01** (2010) 028.
 - [11] J.R. Fergusson and E.P.S. Shellard, *Phys. Rev. D* **80**, 043510 (2009).
 - [12] Mark Hindmarsh, Christophe Ringeval, and Teruaki Suyama, *Phys. Rev. D* **80**, 083501 (2009).
 - [13] D. M. Regan and E. P. S. Shellard, “Cosmic String Power Spectrum, Bispectrum and Trispectrum” (unpublished).
 - [14] J.R. Fergusson and Edward P. S. Shellard, *Phys. Rev. D* **76**, 083523 (2007).
 - [15] Emiliano Sefusatti, Michele Liguori, Amit P.S. Yadav, Mark G. Jackson, and Enrico Pajer, *J. Cosmol. Astropart. Phys.* **12** (2009) 022.
 - [16] A. Curto, E. Martinez-Gonzalez, and R.B. Barreiro, *Astrophys. J.* **706**, 399 (2009).
 - [17] Martin Bucher, Bartjan Van Tent, and Carla Sofia Carvalho, “Detecting Bispectral Acoustic Oscillations from Inflation Using a New Flexible Estimator” (unpublished).
 - [18] Daniel Babich, Paolo Creminelli, and Matias Zaldarriaga, *J. Cosmol. Astropart. Phys.* **08** (2004) 009.
 - [19] G. I. Rigopoulos, E. P. S. Shellard, and B. J. W. van Tent, *Phys. Rev. D* **73**, 083522 (2006).
 - [20] Xingang Chen and Yi Wang, *Phys. Rev. D* **81**, 063511 (2010).
 - [21] Paolo Creminelli, *J. Cosmol. Astropart. Phys.* **10** (2003) 003.
 - [22] Mohsen Alishahiha, Eva Silverstein, and David Tong, *Phys. Rev. D* **70**, 123505 (2004).
 - [23] Xingang Chen, Min-xin Huang, Shamit Kachru, and Gary Shiu, *J. Cosmol. Astropart. Phys.* **01** (2007) 002.
 - [24] Daniel Babich, *Phys. Rev. D* **72**, 043003 (2005).
 - [25] Paolo Creminelli, Alberto Nicolis, Leonardo Senatore, Max Tegmark, and Matias Zaldarriaga, *J. Cosmol. Astropart. Phys.* **05** (2006) 004.
 - [26] Amit P.S. Yadav, *Astrophys. J.* **678**, 578 (2008).
 - [27] E. Komatsu *et al.*, *Astrophys. J. Suppl. Ser.* **180**, 330 (2009).
 - [28] Huiyuan Li and Yuan Xu, [arXiv:0803.0508](https://arxiv.org/abs/0803.0508).
 - [29] James R. Fergusson, Michele Liguori, and Edward P. S. Shellard, “CMB constraints on non-Gaussianity” (to be published).
 - [30] Ian G Moss and Chun Xiong, *J. Cosmol. Astropart. Phys.* **04** (2007) 007.
 - [31] Kendrick M. Smith and Matias Zaldarriaga, [arXiv:astro-ph/0612571](https://arxiv.org/abs/astro-ph/0612571).
 - [32] E. Komatsu *et al.*, *Astrophys. J. Suppl. Ser.*, **148**, 119 (2003).
 - [33] Michele Liguori, Sabino Matarrese, and Lauro Moscardini, *Astrophys. J.* **597**, 57 (2003).
 - [34] Franz Elsner and Benjamin D. Wandelt, *Astrophys. J. Suppl. Ser.* **184**, 264 (2009).
 - [35] Carlo R. Contaldi and Joao Magueijo, *Phys. Rev. D* **63**, 103512 (2001).
 - [36] Graca Rocha, M. P. Hobson, Sarah Smith, Pedro Ferreira, and Anthony Challinor, *Mon. Not. R. Astron. Soc.* **357**, 1 (2005).
 - [37] Duncan Hanson, Kendrick M. Smith, Anthony Challinor, and Michele Liguori, *Phys. Rev. D* **80**, 083004 (2009).
 - [38] K.M. Gorski, *Astrophys. J.* **622**, 759 (2005), <http://healpix.jpl.nasa.gov>.