

Multi-variate joint PDF for non-Gaussianities: exact formulation and generic approximations

Licia Verde,^{a,b} Raul Jimenez,^{a,b} Luis Alvarez-Gaume,^b Alan F. Heavens,^c Sabino Matarrese^{d,e}

^aICREA & ICC, Instituto de Ciencias del Cosmos, University of Barcelona (IEEC-UB), Martí i Franques 1, Barcelona 08028, Spain.

^bTheory Group, Physics Department, CERN, CH-1211, Geneva 23, Switzerland.

^cImperial Centre for Inference and Cosmology, Imperial College, Blackett Laboratory, Prince Consort Road, London SW7 2AZ U.K.

^dDipartimento di Fisica e Astronomia G. Galilei, Università degli Studi di Padova, I-35131 Padova, Italy.

^eINFN, Sezione di Padova, I-35131 Padova, Italy.

E-mail: liciaverde@icc.ub.edu, raul.jimenez@icc.ub.edu, luis.alvarez-gaume@cern.ch, a.heavens@imperial.ac.uk, sabino.matarrese@pd.infn.it

Abstract. We provide an exact expression for the multi-variate joint probability distribution function of non-Gaussian fields primordially arising from local transformations of a Gaussian field. This kind of non-Gaussianity is generated in many models of inflation. We apply our expression to the non-Gaussianity estimation from Cosmic Microwave Background maps and the halo mass function where we obtain analytical expressions. We also provide analytic approximations and their range of validity. For the Cosmic Microwave Background we give a fast way to compute the PDF which is valid up to more than 7σ for f_{NL} values (both true and sampled) not ruled out by current observations, which consists of expressing the PDF as a combination of bispectrum and trispectrum of the temperature maps. The resulting expression is valid for any kind of non-Gaussianity and is not limited to the local type. The above results may serve as the basis for a fully Bayesian analysis of the non-Gaussianity parameter.

Contents

1	Introduction	1
2	Set-up	2
3	Gaussian field $\Phi = \phi$	4
3.1	\mathcal{M} invertible, $n = L = M$	5
3.2	\mathcal{M} identity matrix or invertible, $n < L$, $L = M$	5
3.3	\mathcal{M} non invertible, $n < L$, $L < M$	6
3.4	Adding Gaussian noise	8
4	Quadratic local non-Gaussianity, $f_{\text{NL}} \neq 0$	8
4.1	\mathcal{M} identity matrix	9
4.2	Approximations	11
4.3	Estimators vs. full PDF	13
4.4	If \mathcal{M} is not the identity matrix	19
4.5	\mathcal{M} not invertible, $n < M$	19
4.6	Approximations	20
4.7	Adding Gaussian (or near-Gaussian) noise	22
5	Application I: measuring f_{nl} from a CMB map	22
5.1	Signal-only regime	22
5.2	Adding noise	22
5.3	Second-order PDF expansion	23
6	g_{NL} case	24
7	Application II: Mass function of dark matter halos	25
7.1	Application III: three-dimensional matter overdensity (e.g. 21cm surveys)	25
8	Conclusions	25

1 Introduction

One of the most interesting predictions of slow-roll single field inflation is that it should generate primordial perturbations that are nearly Gaussian distributed, with deviations from non-Gaussianity present at the level of the value of the slow-roll parameters [1–10]. Current constraints on the value of the slow-roll parameters e.g., [11] indicate that non-Gaussianity from single field slow-roll inflation will not be measured from the sky, even if one could access all modes in the sky with futuristic 21cm experiments [12, 13]. However, any deviation from the simplest slow roll (with standard Bunch-Davis vacuum and canonical kinetic term) implies non-Gaussianity, possibly at a detectable level. In particular many multi-field models of inflation will produce detectable non-Gaussianity that current or future surveys could measure: see the review by Refs.[10, 14] for a detailed explanation of the physics behind it. Although in principle the primordial non-Gaussian signal should be present both in the Cosmic Microwave background and in the late large-scale structure of the Universe, most of the effort has so far been devoted to the cosmic microwave background where the gravitationally-induced non-Gaussianity is small e.g.,[5, 15–18] and refs therein. However, if the non-Gaussian signal is at a clearly detectable level, current estimators (which are exact in the Gaussian limit) become sub-optimal and in particular can underestimate the size of confidence regions [19, 20]. Therefore it is important to understand exquisitely the distribution of errors of any estimator, not only its variance (see [36]): the full probability distribution function (PDF) is needed.

Further, a possible detection of non-Gaussianity from the sky would be a direct probe of the properties of de Sitter space, as a study of the different correlators of non-Gaussianity would amount to performing transition experiments on de Sitter; something that will never be accessible from high-energy experiments in the imaginable future. This could open a new window to understand the properties of gravity in de Sitter. One could, for example, answer the question if fluctuations could be treated semi-classically as we do now.

Work so far has concentrated in finding the one-point PDF for primordial non-Gaussian initial conditions as this has a direct application on the calculation of the halo mass function and rare event abundance. This area has seen an exponential expansion starting with the works of [23, 24]. However for applications to e.g., the Cosmic Microwave Background (CMB) an expression for the multi-variate joint PDF is needed.

It is common practice to describe non-Gaussianity as a Taylor expansion of the potential of the form $\Phi = \phi + f_{\text{NL}}(\phi^2 - \langle \phi^2 \rangle)$ [5, 7], where ϕ is a Gaussian random field. This can be extended to the cubic power if the quadratic term is sub-dominant or zero, in which case the free parameter is g_{NL} . This is the so-called local type non-Gaussianity because it is a local transformation in real space. More general kinds of non-Gaussianity beyond the local type can be described by their bispectrum (although this description is not exhaustive as fields with the same bispectrum and power spectrum can have different higher-order correlations) e.g., [14] and references therein.

In addition to the above, it is the multi-variate joint PDF of the non-Gaussian primordial field which is predicted by theory and which has important clues about the physics of the perturbations of de Sitter space; therefore having an exact estimate of it goes beyond the simple experimental merit of constraining one parameter, as it may yield clues on fundamental physics.

Previous work on estimating the PDF for the CMB bispectrum or the non-Gaussianity parameter f_{NL} has focussed on performing Monte Carlo simulations, but these are extremely expensive and may not fully sample the tails of the distribution (i.e. beyond 3σ) [20, 25–27]. Further complications arise from the fact that the PDF will depend at some level on the value of f_{NL} itself. Furthermore, the general techniques developed here may find useful applications in non-Gaussian late-time fields, such as the galaxy density field in the presence of nonlinear bias or nonlinearity from gravitational instability.

In this paper we circumvent these limitations and provide an exact solution based on analytical methods. In particular, we provide an exact analytical formula for the one-point PDF in Eq. (4.12) that can be used for any computation of the mass function of dark matter halos.

Furthermore, an exact formula for the joint multi-variate PDF is provided in Eq. (4.20) and (4.26), which can be numerically integrated. For the CMB we provide a very accurate formula for the PDF (out to more than 7σ fluctuations in the field for f_{NL} values not ruled out by present data) expressed in terms of the observed bispectra and trispectra. This serves as the basis of a fully Bayesian analysis of f_{NL} . Previous work in this direction can be found in [21, 22, 28, 29]. The approach, approximations and applicability regimes considered here are fully complementary.

The paper is organised in a pedagogical way, first, as a warm up, solving the straightforward Gaussian case in § 3 and then the non-Gaussian one in § 4, where we also discuss comparison with previous estimators and present an analytic approximation which is valid out to more than 7σ fluctuations in the field for f_{NL} values not ruled out by present data and applies to non-Gaussianities beyond the local type. We conclude in § 5 with an application of the formalism to the mass function of dark matter halos and the CMB. The reader interested only in practical applications can find our relevant formulas in Eqs. (4.26, 4.28, 5.2).

2 Set-up

Let us define a field A , which is the observable field and can be the temperature of the CMB or the a_ℓ^m or the matter overdensity field etc. The value of the field at any spatial point i is related to the underlying potential Φ via

$$A(\mathbf{x}) = A_i = \int d^3y \mathcal{M}(\mathbf{x}, \mathbf{y}) \Phi(\mathbf{y}) \equiv (\mathcal{M}, \Phi) \quad (2.1)$$

where in the second equality we have introduced a short hand notation. For many applications $\mathcal{M}(\mathbf{x}, \mathbf{y}) = \mathcal{M}(|\mathbf{x} - \mathbf{y}|) = \mathcal{M}(r)$. For example if $A(\mathbf{x})$ is the (linearly evolved) density field $\delta(\mathbf{x})$ smoothed on a given radius R then we could express this in Fourier space:

$$\widetilde{\mathcal{M}}(k) \equiv \int d^3y e^{-i\mathbf{k} \cdot \mathbf{y}} \mathcal{M}(|\mathbf{y}|) = \widetilde{W}_R(k) T(k) g(k) \quad (2.2)$$

where $\widetilde{W}_R(k)$ is the Fourier transform of the assumed filter (typically a top hat for the mass function), $T(k)$ is the transfer function and $g(k) = -2/3(k/H_0)^2 \Omega_{m,0}^{-1}$. The fact that Fourier quantities are complex should not be a problem; one works with a complex field rather than a real one. If the A_i denote the spherical harmonic coefficients of the CMB temperature field, a_{lm} , then the equation becomes:

$$a_{lm} = 4\pi(-i)^l \int \frac{d^3k}{(2\pi)^3} \Phi(\mathbf{k}) g_l(k) Y_{lm}^*(\hat{k}) \quad (2.3)$$

where $g_l(k)$ denotes the radiation transfer function and $\Phi(\mathbf{k}) = \int d^3x \exp[-i\mathbf{k} \cdot \mathbf{x}] \Phi(\mathbf{x})$. Effectively \mathcal{M} acquires extra parameters: $\mathcal{M}^{lm}(k)$. If we observe a Fourier space quantity then we can relate it to the real space Φ by

$$A_j = \int d^3y \exp(-i\mathbf{k} \cdot \mathbf{y}) \Phi(\mathbf{y}) \quad (2.4)$$

i.e. $\mathcal{M}(\mathbf{k}, \mathbf{y}) = \exp(-i\mathbf{k} \cdot \mathbf{y})$, and j labels the wavevector \mathbf{k} .

Several physically-motivated models for primordial non-Gaussianity are expressed as local transformations of an underlying Gaussian field. For example, the non-Gaussianity of the so-called *local* type [5, 7, 30] is given by

$$\Phi = \phi + f_{\text{NL}}(\phi^2 - \langle \phi^2 \rangle) = \phi + f_{\text{NL}}(\phi^2 - \sigma_\phi^2), \quad (2.5)$$

where ϕ is a Gaussian random field and f_{NL} is the non-Gaussianity parameter which is assumed to be small. Here by *small* we mean that this expression is a truncation to linear order in f_{NL} of a physical model that includes e.g., the self coupling of the field during inflation. In the following we would like not to make any further approximations if possible, thus deriving whenever possible expressions that are valid even if $|f_{\text{NL}}\sigma_\phi|$ is not $\ll 1$. Other local models include the so called g_{NL} parameter [24, 31, 32]:

$$\Phi = \varphi + g_{\text{NL}}(\varphi^3 - 3\varphi\sigma_\varphi^2) \equiv \varphi + g'_{\text{NL}}\phi^3 \quad (2.6)$$

where φ, ϕ are Gaussian random fields.

In both cases the observable field A can be written as a sum of a Gaussian term A^G , corresponding to ϕ plus a non-Gaussian correction A^{NG} , corresponding to the term $\propto \phi^2$ or $\propto \phi^3$.

The full information about the Gaussian field ϕ is given by the Gaussian generating functional $\mathcal{P}[\phi]$ (see e.g., Eq 11 of Ref. [23]) which is the basis of several calculations in this paper. For example, the joint probability of A_1, \dots, A_n is given by:

$$\mathcal{P}(A_1, A_2, \dots, A_n) = \int [\mathcal{D}\phi] \mathcal{P}(\phi) \prod_{i=1}^n \delta^D[A_i - (\mathcal{M}, \Phi(\phi))] \quad (2.7)$$

where we have explicitly written that the field Φ is a function of the Gaussian random field ϕ . Here,

$$\mathcal{P}(\phi) = \frac{\exp[-1/2(\phi, K_0, \phi)]}{\int D[\phi] \exp[-1/2(\phi, K_0, \phi)]} \quad (2.8)$$

and we have used the shorthand notation

$$\int d^3y d^3z \phi(\mathbf{y}) K(\mathbf{y} - \mathbf{z}) \phi(\mathbf{z}) \equiv (\phi, K, \phi), \quad (2.9)$$

and with K_0 is defined by the relation to the 2-point correlation function of ϕ , ξ^ϕ :

$$\int d^3y K_0(|\mathbf{x} - \mathbf{y}|) \xi^\phi(|\mathbf{y} - \mathbf{z}|) = \delta(|\mathbf{x} - \mathbf{z}|). \quad (2.10)$$

Index	range	meaning
$a:g$	$1:M$	number of elements of the full discretized space of Φ, ϕ
$i:z$	$1:n$	number of elements of the discretized A array for which we want the joint probability
ν or greek	$n+1:M$	the remaining elements of the full discretized space

Table 1. Summary of the index convention. The number of elements of the observable field A is L . Of course $n \leq L$, $L \leq M$, and $n \leq M$. Note that for \mathcal{M} to be invertible a necessary but not sufficient condition is that the maximum possible size of A to be M i.e. $M = L$.

Then to compute, for example, the probability distribution of a product of n A 's we use

$$\mathcal{P}(B \equiv A_1 \dots A_n) = \int dA_1 dA_2 \dots dA_n \mathcal{P}(A_1, \dots, A_n) \delta^D(B - A_1 \dots A_n). \quad (2.11)$$

Here n is 3 for the three-point function (in Fourier space it is related to the bispectrum), $n = 2$ for the correlation function (in Fourier space it is related to the power spectrum) of A and $n = 1$ for the (one-point) PDF of A .

In these types of integrals, the Fourier representation of the Dirac delta function involves highly oscillatory functions and complicates the integrals. On the other hand, observations are often discretized/pixelized, so one possibility is to discretize these integrals to simplify their analytic calculation, then if needed take the continuous limit at the end. Let us pixelize our field A (say A is known at each pixel of a map) and we have L pixels, let us also pixelize the Φ field for which we have in total M pixels ($L \leq M$). As we will see this simplifies a lot the spatial integrals and the integrals in $\mathcal{D}[\phi]$. The integrals in A_i are however continuous integrals over all possible values that the field can take at position i .

Let us define the index convention we will use. L denotes the size of the observable array A , of which we want to compute the n point distribution $n \leq L$, M denotes the size of the pixelized array ϕ (or Φ). Note that $n \leq L \leq M$. Let us call $M - n = N$. The i indices (or latin indices of the last part of the alphabet: $i-z$) run from 1 to n and ν (or greek indices) runs from $n+1$ to M . From now on latin indices of the first part of the alphabet ($a-h$) runs from 1 to M . Note that if A and Φ have the same dimensions then \mathcal{M} becomes a square matrix ($M \times M$) and everything below is relatively simple. However this does not need to be, as it is clear in the case of the a_{lm} where the i index runs over the discrete- l, m but the index j represent the integral in k which is in principle a continuous variable (but is discretized in most practical applications).

Discretizing we obtain for Equations such as Eq. (2.10) and Eq. (2.1): $\sum_b K_{0ab} \xi_{bc}^\phi = \mathcal{I}_{ac}$ and $\int d^3y \mathcal{M}(x_j, \mathbf{y}) \Phi(\mathbf{y}) \rightarrow \sum_a \mathcal{M}_{ja} \Phi_a$ etc. Also we can express $D[\phi] = d\phi_i d\phi_\nu$.

The Gaussian integral in the denominator of Eq. (2.8) gives

$$\sqrt{\frac{(2\pi)^M}{\det \mathbf{K}_0}}. \quad (2.12)$$

The Dirac delta function becomes $\delta(A_i - \sum_j \mathcal{M}_{ij} \phi_j)$.

Finally, let us report here a result for Gaussian integrals with we will use repeatedly below. The Gaussian integral is

$$\int d\varphi \exp \left[-\frac{1}{2} \varphi^T \mathbf{K} \varphi + \mathbf{B}^T \varphi \right] = \sqrt{\frac{(2\pi)^m}{\det \mathbf{K}}} \exp \left[\frac{1}{2} \mathbf{B}^T \mathbf{K}^{-1} \mathbf{B} \right] \quad (2.13)$$

with \mathbf{K} a square matrix $m \times m$ and \mathbf{B} and φ are vectors of size m .

3 Gaussian field $\Phi = \phi$

To illustrate the approach and familiarize ourselves with some of the expressions, let us start in the simplest case where $f_{\text{NL}} = 0$, $g_{\text{NL}} = 0$. We begin by considering the case where \mathcal{M} is a square,

invertible matrix. We then progressively consider more realistic/complicated cases: \mathcal{M} not invertible and the addition of (Gaussian) noise. In what follows, we denote with the over script arrow the vector composed of the value of the field in all the pixels, e.g., $\vec{A} = \{A_1, \dots, A_n\}$.

3.1 \mathcal{M} invertible, $n = L = M$

In this case we can do the following change of variables: $\phi \rightarrow \mathcal{M}\phi$ or $\hat{\phi}_a = \sum_{ab} \mathcal{M}_{ab} \phi_b$ and $\hat{K}_0 = (\mathcal{M}^{-1})^T K_0 \mathcal{M}^{-1}$ (matrix multiplication); $\prod_a d\phi_a = \frac{\prod_a d\hat{\phi}_a}{\det(\mathcal{M})}$. We thus obtain:

$$\begin{aligned} \mathcal{P}(\vec{A}) &= \sqrt{\frac{\det K_0}{(2\pi)^M}} \int \frac{\prod_i d\hat{\phi}_i}{\det(\mathcal{M})} \prod_i \delta^D(A_i - \hat{\phi}_i) \exp \left[-\frac{1}{2} \sum_{ij} \hat{\phi}_i \hat{\phi}_j \hat{K}_{0ij} \right] \\ &= \sqrt{\frac{\det \hat{K}_0}{(2\pi)^M}} \exp \left[-\frac{1}{2} \sum_{ab} A_a A_b \hat{K}_{0ab} \right], \end{aligned} \quad (3.1)$$

which is not surprising given that if the ϕ field is a multivariate Gaussian, the A field properties are also those of a multi-variate Gaussian, albeit with a different covariance.

Note that \hat{K}_0 is defined so that

$$\int d^3y \hat{K}_0(|\mathbf{x} - \mathbf{y}|) \xi_A(|\mathbf{y} - \mathbf{z}|) = \delta(|\mathbf{x} - \mathbf{z}|). \quad (3.2)$$

In other words, \hat{K}_0 is the equivalent quantity of K_0 , not for the underlying field ϕ , but for the *observable* field A .

3.2 \mathcal{M} identity matrix or invertible, $n < L$, $L = M$

Let us start with the simpler, special case where \mathcal{M} is the identity matrix, for example, when looking at the distribution of the product of n pixel values of an unsmoothed field. The values taken by the field ϕ will be denoted by φ (here φ plays the role of A but we have changed notation to make clear that one is really seeing the same field).

We start from:

$$\begin{aligned} \mathcal{P}(\varphi_i, \dots, \varphi_n) &= \sqrt{\frac{\det K_0}{(2\pi)^M}} \int \prod_{i,\nu} d\phi_i d\phi_\nu \prod_i \delta^D(\varphi_i - \phi_i) \times \\ &\exp \left[-\frac{1}{2} \sum_{ij} \phi_i \phi_j K_{0ij} - \sum_{i,\nu} \phi_i \phi_\nu K_{0i\nu} - \frac{1}{2} \sum_{\mu,\nu} \phi_\mu \phi_\nu K_{0\mu\nu} \right]. \end{aligned} \quad (3.3)$$

The integral over $d\phi_i$ can be done trivially exploiting the Dirac delta function yielding:

$$\begin{aligned} \mathcal{P}(\varphi_i, \dots, \varphi_n) &= \sqrt{\frac{\det K_0}{(2\pi)^M}} \int \prod_\nu d\phi_\nu \times \\ &\exp \left[-\frac{1}{2} \sum_{ij} \varphi_i \varphi_j K_{0ij} - \sum_{i,\nu} \varphi_i \phi_\nu K_{0i\nu} - \frac{1}{2} \sum_{\mu,\nu} \phi_\mu \phi_\nu K_{0\mu\nu} \right]. \end{aligned} \quad (3.4)$$

Here we recognize the Gaussian integral of Eq. (2.13) where in particular $B_\nu = \sum_i \varphi_i K_{0i\nu}$ and \mathbf{K} is the matrix made of the elements $K_{0\mu\nu}$.

Note that the matrix K_0 can be split in block matrices as

$$\mathbf{K}_0 = \begin{pmatrix} K & \mathbb{k} \\ \mathbb{k}^T & \mathcal{K} \end{pmatrix} \quad (3.5)$$

where K is made of the K_{0ij} elements (it is a square sub matrix of size $n \times n$), \mathcal{K} is made by the $K_{0\mu\nu}$ elements (it is a square matrix of size $N \times N$), and \mathbb{k} is an $n \times N$ matrix.

Finally we obtain that:

$$\mathcal{P}(\vec{\varphi}) = \sqrt{\frac{\det(\mathcal{G})}{(2\pi)^n}} \exp \left[-\frac{1}{2} \sum_{ij=1}^n \varphi_i \varphi_j \mathcal{G}_{ij} \right]. \quad (3.6)$$

Here

$$\mathcal{G}_{ij} = K_{0ij} - \sum_{\mu\nu} K_{0i\mu} K_{0j\mu} \mathcal{K}^{-1} = K - \mathbb{k} \mathcal{K}^{-1} \mathbb{k}^T, \quad (3.7)$$

and we have used the fact that

$$\det \mathbf{K}_0 = \det \begin{pmatrix} K \\ \mathbb{k} \end{pmatrix} = \det(K - \mathbb{k} \mathcal{K}^{-1} \mathbb{k}^T) \det(\mathcal{K}) \quad (3.8)$$

implying $\det(K_0)/\det(\mathcal{K}) = \det(\mathcal{G})$.

Eq. (3.6), not surprisingly, is a Gaussian but with a different covariance from that of ϕ . It is not difficult to see, using the Woodbury formula

$$(A + UCV)^{-1} = A^{-1} - A^{-1}U(C^{-1} + VA^{-1}U)^{-1}A^{-1} \quad (3.9)$$

(where A , U , C , and V denote generic matrices) that

$$[\mathcal{G}^{-1}]_{ij} = [K_0^{-1}]_{ij} = \xi_{ij} \quad (3.10)$$

as expected.

In case \mathcal{M} is not the identity matrix but it is still an invertible matrix, as long as the array A has been pixelized in as many elements as ϕ , it is straightforward to compute $\mathcal{P}(\vec{A})$ even if $n < M$. One just does the substitution: $\phi_i \rightarrow \sum_a \mathcal{M}_{ia} \phi_a$ etc. as above. In total one gets,

$$\mathcal{P}(\vec{A}) = \sqrt{\frac{\det(\hat{\mathcal{G}})}{(2\pi)^n}} \exp \left[-\frac{1}{2} \sum_{ij=1}^n A_i A_j \hat{\mathcal{G}}_{ij} \right], \quad (3.11)$$

where $\hat{K}_0 = (\mathcal{M}^{-1})^T K_0 \mathcal{M}^{-1}$ (matrix multiplication) and $\hat{\mathcal{G}}$ is the equivalent of \mathcal{G} (see Eq. (3.7)) but made from \hat{K}_0 instead of K_0 , $[\hat{\mathcal{G}}^{-1}]_{ij} = [\hat{K}_0^{-1}]_{ij} = \xi_{Aij}$. That is we have a Gaussian with covariance given by that of the observed field.

3.3 \mathcal{M} non invertible, $n < L$, $L < M$

This is only a slightly more complicated case for a Gaussian field (we will see below that for a non-Gaussian field, making \mathcal{M} non invertible complicates the resulting expressions considerably). In this case the observable field A is pixelized in L elements, but the dimension of the field ϕ (and the field Φ) is larger, for example $M = L \times H$ (i.e. there is one (or more) extra dimensions which gets pixelized to have H elements). To keep in mind a possible relevant application, this is the case of a CMB temperature map: the primordial field Φ is three-dimensional, the observed temperature fluctuation field around a given observer is two-dimensional, therefore \mathcal{M} is non invertible.

It is easy to see that one can rewrite Eq. (2.7) considering that if ϕ is Gaussian then its projection, e.g., to a_{lm} , will also be Gaussian (\mathcal{M} is a linear operator). So we only have to reinterpret Eq. (3.1) in light of Eq. (2.1) to have $\hat{\phi} = A^G$ and \hat{K}_0 to be the equivalent quantity for the Gaussian A^G , in the case of the a_{lm} it will be connected to the C_ℓ . In this case there is no need to invert \mathcal{M} . It is however instructive in view of the non-Gaussian case, to derive the same result going through the actual calculations.

We start from

$$\begin{aligned} \mathcal{P}(\vec{A}) &= \sqrt{\frac{\det K_0}{(2\pi)^M}} \int \prod_{i,\nu} d\phi_i d\phi_\nu \prod_i \delta^D(A_i - \sum_{ia} \mathcal{M}_{ia} \phi_a) \times \\ &\exp \left[-\frac{1}{2} \sum_{ij} \phi_i \phi_j K_{0ij} - \sum_{i,\nu} \phi_i \phi_\nu K_{0i\nu} - \frac{1}{2} \sum_{\mu,\nu} \phi_\mu \phi_\nu K_{0\mu\nu} \right], \end{aligned} \quad (3.12)$$

and split \mathcal{M} into block matrices as:

$$\mathcal{M} = \begin{pmatrix} \aleph & \mathcal{N} \end{pmatrix}, \quad (3.13)$$

where \aleph is a $n \times n$ sub matrix and \mathcal{N} is $n \times N$. We also split the vector $\phi = \{\phi_i, \phi_\nu\} = \{\phi', \phi''\}$. We start by re-writing the argument of the Dirac delta as:

$$A_i - \sum_j \mathcal{M}_{ij} \phi_j - \sum_\nu \mathcal{M}_{i\nu} \phi_\nu = A_i - (\aleph \phi')_i - (\mathcal{N} \phi'')_i \quad (3.14)$$

and make the following change of variables

$$A \longrightarrow \aleph^{-1} A \equiv \tilde{A}. \quad (3.15)$$

\aleph is invertible if the rank of \mathcal{M} is maximal. Note that there is some freedom in splitting the sum in Eq.(3.14), but we argue that the rank of \mathcal{M} is always maximal if the information in A is not redundant (e.g., in the absence of noise if pixels are not counted more than once). This change of variables gives a multiplicative overall factor $|\det \aleph|$ and which changes the argument of the Dirac delta to

$$\tilde{A}_i - \phi_i - (\aleph^{-1} \mathcal{N} \phi'')_i \equiv \tilde{A}_i - \phi_i - \tilde{\phi}''_i. \quad (3.16)$$

Now we can integrate over the $d\phi_i$ and eliminate the Dirac deltas giving the following substitution $\phi_i \longrightarrow \tilde{A}_i - \tilde{\phi}''_i$. We get:

$$\mathcal{P}(\vec{A}) = \sqrt{\frac{\det K_0}{(2\pi)^M}} |\det \aleph| \int d\phi'' \exp \left[-\frac{1}{2} \left(\tilde{A}^T K \tilde{A} + 2\tilde{A}^T \mathcal{R} \phi'' + \phi''^T \mathcal{Q} \phi'' \right) \right] \quad (3.17)$$

$$= \sqrt{\frac{\det K_0}{(2\pi)^M}} |\det \aleph| \int d\phi'' \exp \left[-\frac{1}{2} \{ \tilde{A}, \phi'' \}^T H \{ \tilde{A}, \phi'' \} \right] \quad (3.18)$$

where

$$H = \begin{pmatrix} K & \mathcal{R} \\ \mathcal{R}^T & \mathcal{Q} \end{pmatrix} \quad (3.19)$$

$$\mathcal{R} = \mathbb{K} - K \aleph^{-1} \mathcal{N} \quad (3.20)$$

and

$$\mathcal{Q} = K + \mathcal{N}^T \aleph^{-1 T} K \aleph^{-1} \mathcal{N} - \mathcal{N}^T \aleph^{-1 T} \mathbb{K} - \mathbb{K}^T \aleph^{-1} \mathcal{N}, \quad (3.21)$$

which is an $N \times N$ matrix.

We then recognize a Gaussian integral yielding

$$\mathcal{P}(\vec{A}) = \sqrt{\frac{\det K_0}{(2\pi)^n}} |\det \aleph| \frac{1}{\sqrt{|\det \mathcal{Q}|}} \exp \left[-\frac{1}{2} A^T \tilde{\mathcal{F}} A \right] \quad (3.22)$$

where

$$\mathcal{F} = K - \mathcal{R} \mathcal{Q}^{-1} \mathcal{R}^T \quad (3.23)$$

and therefore, using Woodbury again, $[\mathcal{F}^{-1}]_{ij} = [H^{-1}]_{ij}$, and $\tilde{\mathcal{F}} = \aleph^{-1 T} \mathcal{F} \aleph^{-1}$. Of course, in this simple example we are using the general framework to compute the PDF of another Gaussian field, \tilde{A} ,

whose covariance is simply $\mathcal{M}^T \xi \mathcal{M}$. It is straightforward using the Woodbury formula to demonstrate that we recover this result. For the Gaussian case there is therefore no need to invert \mathcal{Q} or \aleph .

Equivalently, one could change the basis for ϕ to ψ so that the matrix \mathcal{M} is $\mathcal{M}' = (\aleph', 0)$. This being a linear operation means that ψ is still Gaussian. Note that this means: take the rectangular matrix given by $\{\aleph, \mathcal{N}\}$ and SVD decompose it as $U\Sigma V^T$ (or $U\Sigma V^*$ (conjugate transpose)). Then do the following change of variables: $\phi \rightarrow U\phi = \psi$ and therefore $\mathcal{M} \rightarrow U^{-1}\mathcal{M} = \mathcal{M}'$. The Jacobian of the transformation is therefore $\det U$.

Thus the argument of the Dirac delta becomes $A_i - \aleph'_i \psi'$ and as before we make the change of variable $A \rightarrow \aleph'^{-1}A \equiv \tilde{A}$. This now simplifies the equations above as $\mathcal{Q}' = \mathcal{K}'$ and $\mathcal{R}' = \mathbb{K}'$. (Of course the Jacobian of the transformation must be computed; alternatively one makes sure, a posteriori, that the probability is suitably normalized). We find:

$$\mathcal{P}(\tilde{A}) = \sqrt{\frac{\det K_0}{(2\pi)^n}} \frac{|\det \aleph'|}{|\det U|} \frac{1}{\sqrt{|\det \mathcal{Q}'|}} \exp \left[-\frac{1}{2} \tilde{A}^T \tilde{\mathcal{G}}' \tilde{A} \right] \quad (3.24)$$

where $\tilde{\mathcal{G}}' = \aleph'^{-1T} \mathcal{G}' \aleph'^{-1}$ and $\mathcal{G}' = K' - \mathcal{R}' \mathcal{Q}'^{-1} \mathcal{R}'^T$ (where prime denotes the quantity in the ψ basis).

Of course Eqs. (3.22) and (3.24) are equivalent but depending on the context one or the other might be more suitable for practical implementation. Note that $\tilde{\mathcal{G}}' = (\mathcal{M}'^T \xi' \mathcal{M}')^{-1}$ where we use the prime to indicate the quantity for the ψ field. Since $\xi' = U \xi U^T$ and $\mathcal{M}' = U^{-1} \mathcal{M}$ we obtain $\mathcal{M}'^T \xi' \mathcal{M}' = \mathcal{M}^T \xi \mathcal{M} = \xi_A$ and therefore $\tilde{\mathcal{G}}' = \tilde{\mathcal{G}}$, $[\tilde{\mathcal{G}}'^{-1}]_{ij} = \xi_A{}_{ij}$.

To interpret $\tilde{\mathcal{G}}$ e.g., in Eq. (3.22) we can readily write down the 1-point PDF ($n = 1$), where $\tilde{\mathcal{G}}$ is a scalar, $\tilde{\mathcal{G}}_{11}$:

$$\mathcal{P}(A_1) = \frac{\tilde{\mathcal{G}}}{\sqrt{2\pi}} \exp \left[-\frac{1}{2} A_1^2 \tilde{\mathcal{G}}_{11} \right]. \quad (3.25)$$

If A is the field in real space then $\tilde{\mathcal{G}}(R)$ is $1/\sigma_R^2$, the inverse variance of the field, with R indicating the smoothing scale. If A is in Fourier space then $\tilde{\mathcal{G}}(k)$ is $1/\mathcal{P}(k)$.

3.4 Adding Gaussian noise

Having Gaussian noise in a map means that the observed field \mathcal{A} is given by the superposition of the field A as before and an independent Gaussian field ϵ , with covariance Σ_ϵ , $\mathcal{A} = A + \epsilon$. Thus in Eqs. (2.7, 3.1, 3.12) etc. the Dirac delta gets substituted by a Gaussian with covariance Σ_ϵ . It is easy to see that if $\Sigma_\epsilon^{-1} \equiv \mathcal{E}$ then the second line of Eq. (3.1) remains the same provided that $\hat{K}_0 \rightarrow \mathcal{G}$ where $\mathcal{G} = [\hat{K}_0^{-1} + \mathcal{E}^{-1}]^{-1}$. In doing so we have used the result for the Gaussian integral Eq. (2.13) and that $\mathcal{E} - \mathcal{E}(\mathcal{E} + \hat{K}_0)^{-1} \mathcal{E} = [\hat{K}_0^{-1} + \mathcal{E}^{-1}]^{-1}$.

Therefore for all the subsequent cases above one must just make the substitution $\mathcal{G} \rightarrow [\mathcal{G}^{-1} + \mathcal{E}]^{-1}$. Note that we could have obtained this result more directly as follows: the PDF of the superposition of two independent processes (A and ϵ) is the convolution of the individual PDFs. If the two PDFs are Gaussian the resulting PDF will be a Gaussian with inverse covariance given by the sum of the individual inverse covariances.

4 Quadratic local non-Gaussianity, $f_{\text{NL}} \neq 0$

After having used our approach for a Gaussian field to reproduce known results and familiarize ourselves with the manipulation of some of the expressions, we go on to tackle the case of a quadratic, local non-Gaussianity. As we will see this apparently simple, non-linear transformation of a Gaussian field, in some cases –those of most practical interest– complicates the final expressions quite a lot because it spoils the analyticity of some of the final steps we were able to use in the Gaussian case. Here we will proceed somewhat in parallel to Sec. 3, by first considering the simplest case and then introducing complications progressively.

We begin by noting that f_{NL} appears only in the argument of the Dirac delta function (see Eq. (2.7)). It is much easier to derive the final result in the case where \mathcal{M} is the identity matrix. While this case might not be of great practical interest, although some applications to future data are possible and are mentioned below, it is a useful preliminary step to then tackle more complicated and realistic cases.

4.1 \mathcal{M} identity matrix

If \mathcal{M} is identity matrix, one possible way to proceed to evaluate Eq. (2.7) is to use the property of the Dirac delta:

$$\delta(f(x)) = \sum_r \frac{\delta(x - x_r)}{|f'(x_r)|} \quad (4.1)$$

where r here denotes the roots of the function f . In fact if \mathcal{M} is the identity matrix, we have to solve $A_i + C - \phi - f_{\text{NL}}\phi^2 = 0$, where C is a constant, to make the field Φ to have zero average ($C = f_{\text{NL}}\langle\phi^2\rangle$). The roots of this are:

$$\frac{-1 \pm \sqrt{1 + 4f_{\text{NL}}(A_i + C)}}{2f_{\text{NL}}} \quad (4.2)$$

and

$$f' = 1 + 2f_{\text{NL}}\phi. \quad (4.3)$$

Our Dirac delta becomes

$$\prod_i \left[\frac{\delta\left(\phi_i - \frac{-1 + \sqrt{1 + 4f_{\text{NL}}(A_i + C)}}{2f_{\text{NL}}}\right)}{\sqrt{1 + 4f_{\text{NL}}(A_i + C)}} + \frac{\delta\left(\phi_i - \frac{-1 - \sqrt{1 + 4f_{\text{NL}}(A_i + C)}}{2f_{\text{NL}}}\right)}{\sqrt{1 + 4f_{\text{NL}}(A_i + C)}} \right]. \quad (4.4)$$

In the case $n = 1$ there is no product, only the two deltas. Let us call

$$\begin{aligned} f_1(A_i) &= \frac{-1 + \sqrt{1 + 4f_{\text{NL}}(A_i + C)}}{2f_{\text{NL}}} \\ f_2(A_i) &= \frac{-1 - \sqrt{1 + 4f_{\text{NL}}(A_i + C)}}{2f_{\text{NL}}}. \end{aligned} \quad (4.5)$$

Then Eq.(2.7) becomes:

$$\begin{aligned} \mathcal{P}(\vec{A}|f_{\text{NL}}) &= \sqrt{\frac{\det K_0}{(2\pi)^M}} \prod_i \frac{1}{\sqrt{1 + 4f_{\text{NL}}(A_i + C)}} \int \prod_\nu d\phi_\nu \exp \left[-\frac{1}{2} \sum_{\mu,\nu} \phi_\mu \phi_\nu K_{0\mu\nu} \right] \\ &\times \sum_{rq} \exp \left[-\frac{1}{2} \sum_{ij} f_r(A_i) f_q(A_j) K_{0ij} - \sum_{i,\nu} f_r(A_i) \phi_\nu K_{0i\nu} \right]. \end{aligned} \quad (4.6)$$

As is well known, from the pioneering work of [33] to the more recent [22], for small f_{NL} the root with the minus sign will be exponentially suppressed. For illustration, Fig. (1) shows the effect of neglecting the root on the 1 point PDF (which can be easily computed) for several choices for $f_{\text{NL}}\sigma$. For simplicity we have taken ϕ to have unit variance; where the distribution is not plotted is because it becomes imaginary (i.e. there is no solution). The figure shows that neglecting the root $r = 2$ is a very good approximation since deviations are small and non-zero only close to the limit where the PDF is not defined. Neglecting the $r = 2$ root we obtain:

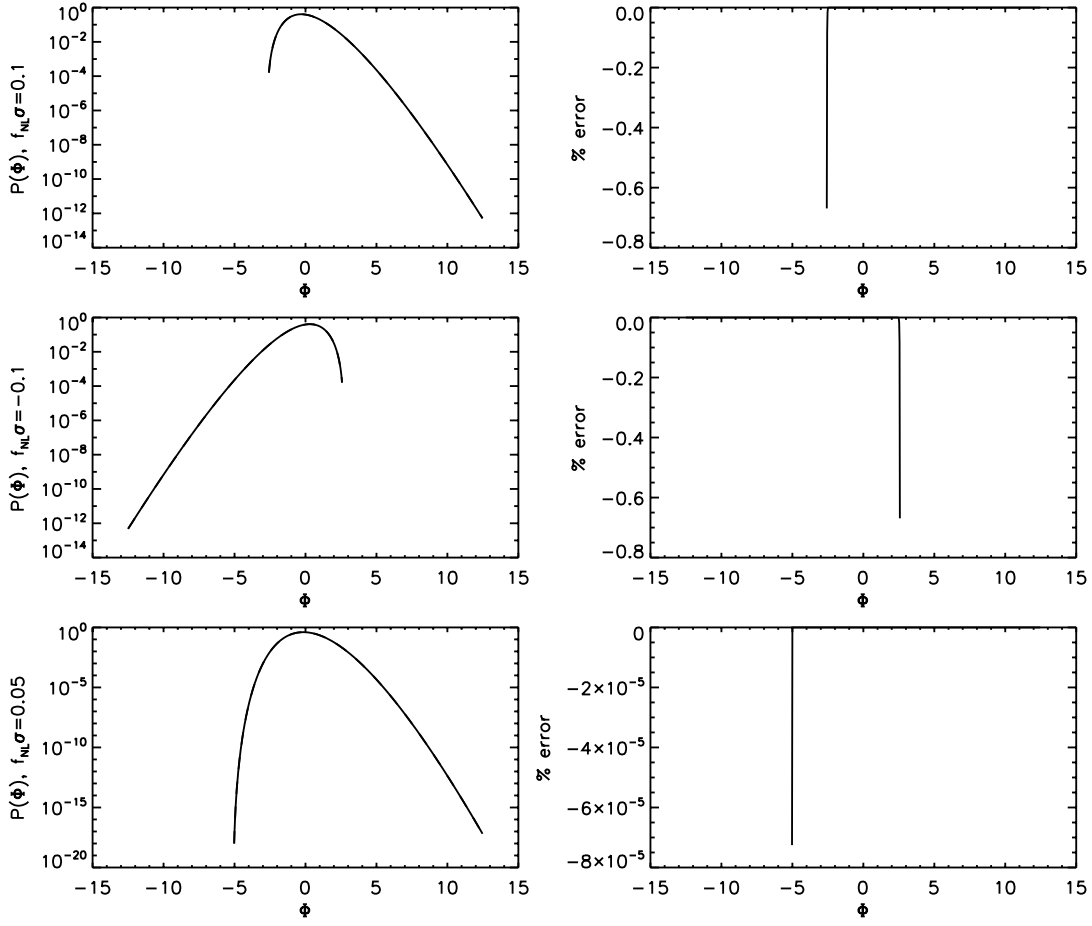


Figure 1. Exact one-point PDF for Φ (normalized to have unit variance) on the left panel and the effect of neglecting the exponentially suppressed root for “small” f_{NL} on the right panel. Several choices for $f_{\text{NL}}\sigma_\phi$ are shown. Where the distribution is not plotted is because it becomes imaginary (i.e. there is no solution). Neglecting the root $r = 1$ is a very good approximation since deviations are small and non-zero only close to the limit where the PDF is not defined.

$$\mathcal{P}(\vec{A}|f_{\text{NL}}) = \sqrt{\frac{\det K_0}{(2\pi)^M}} \prod_i \frac{1}{\sqrt{1 + 4f_{\text{NL}}(A_i + C)}} \int \prod_\nu d\phi_\nu \quad (4.7)$$

$$\times \exp \left[-\frac{1}{2} \sum_{ij} f_1(A_i) f_1(A_j) K_{0ij} - \sum_{i,\nu} f_1(A_i) \phi_\nu K_{0i\nu} - \frac{1}{2} \sum_{\mu,\nu} \phi_\mu \phi_\nu K_{0\mu\nu} \right]. \quad (4.8)$$

If $n < M$ we can recognize the Gaussian integral again and simplify the above equation:

$$\mathcal{P}(\vec{A}|f_{\text{NL}}) = \sqrt{\frac{\det \mathcal{G}}{(2\pi)^n}} \prod_i \frac{1}{\sqrt{1 + 4f_{\text{NL}}(A_i + C)}} \left\{ \exp \left[-\frac{1}{2} \sum_{ij} f_1(A_i) f_1(A_j) \mathcal{G}_{ij} \right] \right\}. \quad (4.9)$$

For $n = M$ this is instead:

$$\mathcal{P}(\vec{A}|f_{\text{NL}}) = \sqrt{\frac{\det K_0}{(2\pi)^n}} \prod_i \frac{1}{\sqrt{1 + 4f_{\text{NL}}(A_i + C)}} \left\{ \exp \left[-\frac{1}{2} \sum_{ij} f_1(A_i) f_1(A_j) K_{0ij} \right] \right\} \equiv \mathcal{P}(\Phi). \quad (4.10)$$

Note that this is the multivariate distribution of the Φ field (\mathcal{M} having been set to the identity). Note the similarity between this and Eq. 14 in [22]. Of course, to obtain the same quantity for \mathcal{M} not the identity matrix, one can then always write that

$$\mathcal{P}(\vec{A}|f_{\text{NL}}) = \int \mathcal{D}[\Phi] \mathcal{P}(\Phi) \delta^D(A - \int M\Phi). \quad (4.11)$$

This then must be integrated numerically via Monte-Carlo methods, as for example explored in [22, 28]. While Ref. [22, 28] proceeds by computing this (or a related) expression directly with a Monte Carlo approach, here we attempt to carry on analytically. If \mathcal{M} is the identity matrix, the one-point function can be obtained by setting, $i, j \equiv 1$ as for the Dirac delta of Eq. (4.9). One then recognizes the Gaussian integral in ϕ obtaining

$$\mathcal{P}(A_1|f_{\text{NL}}) = \sqrt{\frac{\det(\mathcal{G})}{2\pi}} \frac{1}{\sqrt{1 + 4f_{\text{NL}}(A_1 + C)}} \left(\exp \left[-\frac{1}{2} f_1^2(A_1) \mathcal{G}_{11} \right] + \exp \left[-\frac{1}{2} f_2^2(A_1) \mathcal{G}_{11} \right] \right).$$

(4.12)

4.2 Approximations

We now take a detour (subsections 4.2 and 4.3). We first explore possible approximations of this exact PDF expression, their performance and range of validity. For simplicity and clarity we do so in the case where \mathcal{M} is the identity matrix, i.e., where we can ignore the transfer function. Therefore the results will be valid only qualitatively and not strictly quantitatively in the (practically interesting case) when there is a transfer function, especially if it makes \mathcal{M} not invertible. Nevertheless the performance of the approximations considered will be directly useful in Sec.4.6. In addition they could be useful to find approximations to Eq. 4.26.

It is possible to further simplify our expressions in Sec. 4.1 by expanding to first (or higher) order in f_{NL} . For example expanding at first order the square root we obtain $(1 + 2f_{\text{NL}} \sum_i (A_i + C))$ both in Eqs. (4.9) and (4.10), in the same spirit then $f_1(A_i) \rightarrow A_i + C - A_i^2 f_{\text{NL}}$. Figure (2) shows how different approximations perform in the case of the one-point PDF (where \mathcal{M} is the identity matrix). We illustrate the performance of the different approximations by using the one-point PDF and taking ϕ to have unit variance ($\sigma = 1$). Note that due to this normalization the real expansion parameter is $f_{\text{NL}}\sigma$. For clarity below we will use the symbol f_{nl} to indicate the non-Gaussianity parameter for a field with unit variance, thus $f_{\text{nl}} = f_{\text{NL}}\sigma$.

The solid (black) line shows the first-order expansion of both the square root and the argument of the exponential:

$$\mathcal{P}^{1st}(\Phi|f_{\text{nl}}) = \frac{1}{\sqrt{2\pi}} (1 - 2f_{\text{nl}}\Phi) \exp \left[-\frac{1}{2} \Phi^2 + f_{\text{nl}}(\Phi^3 - \Phi) \right]. \quad (4.13)$$

Note that the Babich (2005) [33] (B05 from now on) approximation of the PDF (see their Eq. 20) would correspond to approximating $(1 - 2f_{\text{nl}}\Phi) \sim \exp(-2f_{\text{nl}}\Phi)$. On the other hand the first-order Edgeworth expansion (see e.g. [24] their Eq. (4.10)) corresponds to further expanding the non-Gaussian part in the exponential by $\exp[f_{\text{nl}}(\Phi^3 - \Phi)] \sim 1 + f_{\text{nl}}(\Phi^3 - \Phi)$ (thin solid (black) lines). The popular non-Gaussianity estimator [25, 33, 34] where the term $(\Phi^3 - 3\Phi)$ is considered stems from an expansion of the PDF (more on this below).

In order to make the path integrals in what follows analytic it is worth to note that an integrand made by the product of a polynomial with an exponential of a quadratic expression is integrable analytically but an exponential of an expression of power higher than quadratic is not. However expanding the $f_{\text{nl}}\Phi^3$ term in the exponent as $(1 + f_{\text{nl}}\Phi^3)$ performs very badly as Fig. (2) shows (dot-dashed).

Fig. (2) shows that the first-order expansion in f_{nl} does not provide a good approximation to the PDF, and indeed it is well known for example that a first-order Edgeworth expansion can give




Figure 2. Left: the one-point PDF (the underlying Gaussian field ϕ is normalized to have unit variance ($\sigma = 1$)). Solid line with diamonds: exact. Solid (black) line: keeping only terms proportional to $f_{\text{NL}}\sigma$ both in the argument of the exponential and in the factor in front. Solid thin (black) line: First-order expansion. Dotted line (purple): keeping terms proportional to $f_{\text{NL}}\sigma^2$ both in the argument of the exponential and in the factor in front. Short dashed line (blue): as above but expanding the terms $\propto f_{\text{NL}}\sigma^2$ in the argument of the exponential as $\exp(x) = (1 + x)$. Dot-dashed (light blue): expression where starting from an expansion that keeps terms proportional to $f_{\text{NL}}\sigma^2$ both in the argument of the exponential and in the factor in front, only terms $\propto \Phi$ and Φ^2 are kept in the argument of the exponential, the rest is expanded as $\exp(x) = 1 + x$. Solid thick (red): second-order Edgeworth expansion. Long dashed (red): Expansion to second order in $f_{\text{NL}}\sigma$. Dot-dot-dashed line (green): approximation of Babich (2005) Right: % error in the approximation (same linestyle as in the left panel). See text for more details.

negative values for the PDF (see discussion in [24]). Going to second order helps. Expanding to second order both the square-root and the argument of the exponential yields the dotted line in the figure. If the second-order part of the exponential is also expanded as $\exp(x) = 1 + x$ we obtain the short dashed line.

Expanding the full expression to second order in f_{NL} (thus expressing the PDF as a Gaussian multiplied by a correction) we obtain something similar (but not quite identical) to the second-order Edgeworth expansion:

$$\mathcal{P}^{\text{2nd}}(\Phi|f_{\text{NL}}) = \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{1}{2}\Phi^2\right] \left[1 + f_{\text{NL}}(\Phi^3 - 3\Phi) + \frac{f_{\text{NL}}^2}{2}(\Phi^6 - 11\Phi^4 + 23\Phi^2 - 5)\right] \quad (4.14)$$

where the second-order Edgeworth expansion would be:

$$\mathcal{P}^{\text{2ndEW}}(\Phi|f_{\text{NL}}) = \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{1}{2}\Phi^2\right] \left[1 + f_{\text{NL}}(\Phi^3 - 3\Phi) + \frac{f_{\text{NL}}^2}{2}(\Phi^6 - 13\Phi^4 + 33\Phi^2 - 9)\right]. \quad (4.15)$$

The fact that Eq. (4.14) and Eq. (4.15) are different at second order should not be too surprising. In fact while Eq. (4.14) is a polynomial expansion around $f_{\text{NL}} = 0$ using polynomials in their *natural* form, Eq. (4.15) is an expansion in Hermite polynomials, which, for the problem at hand, are an orthogonal basis. It is well known that expansions in orthogonal polynomials have smaller truncation errors, much better convergence and extrapolation properties than other polynomial expansions. Therefore a truncated Edgeworth expansion should be preferred to a Taylor expansion truncated at the same order.

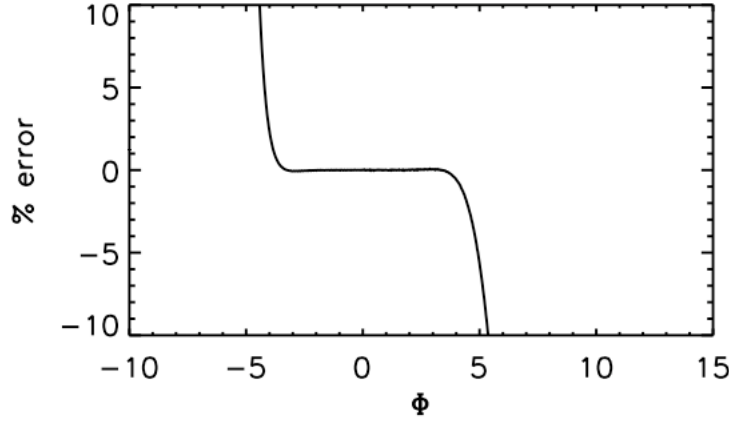


Figure 3. Hybrid approach: the exponential part preserves the Gaussian structure as it has only terms $\propto \phi$ and $\propto \phi^2$. Again ϕ is normalized to have unit variance.

Not unexpectedly, our numerical investigations have shown that the difference between Eq. (4.14) and Eq. (4.15), at least for relatively small f_{nl} , is small.

Note that the B05 [33] second-order expansion would correspond to

$$\mathcal{P}^{\text{2ndBabich}}(\Phi|f_{\text{nl}}) = \frac{1}{\sqrt{2\pi}} \exp \left[-\frac{1}{2}\Phi^2 + f_{\text{nl}}(\Phi^3 - 3\Phi) + \frac{f_{\text{nl}}^2}{2}(-5\Phi^4 + 14\Phi^2 - 5) \right]. \quad (4.16)$$

One “hybrid” approach is to look for simplicity of integration, first expand to second order both the square root and the argument of the exponential, then leave in the argument of the exponential only terms $\propto \phi^2$ and $\propto f_{\text{nl}}\phi$ the rest expand as $\exp(x) = 1 + x + x^2$. This is shown in Fig. (3).

Having different approximations used in the literature and the exact expression allows us to quantify their performance. In table (4.2) we show the range of Φ/σ where the different approximations work at better than 20%. We conclude that a second-order expansion (Edgeworth or plain second order or hybrid) reproduces well the PDF except in the extreme tails ($> 5\sigma$) for relatively high f_{nl} values and in addition offers the possibility of performing many of the relevant path integrals analytically. In what follows we will therefore use this approximation when needed.

Note that in the form of B05 the log of the PDF to second order is written as a Gaussian part plus a first-order term given by $I_1 \times f_{\text{NL}}$ and a second-order term $I_2 \times f_{\text{NL}}^2$. $\langle I_1 \rangle = 6f_{\text{NL}}$ and $\langle I_2 \rangle = 6$. If we write the second-order expansion or the second-order Edgeworth expansion in the B05 form we obtain the same I_1 expression but different expressions for I_2 . Note however that $\langle I_2 \rangle$ is always equal to 6.

While so far we have interpreted the PDF as a probability of Φ for a given f_{NL} , we can reinterpret it, assuming a uniform prior on f_{NL} , as the PDF for f_{NL} given Φ . Doing so, as we will show in details below, will uncover some properties of this PDF which we believe were missed so far. These stem from the fact that in $\mathcal{P}(\Phi|f_{\text{NL}})$, f_{NL} is fixed and it is therefore physically motivated to chose it to be small. Conversely, when considering $\mathcal{P}(f_{\text{NL}}|\Phi)$, the f_{NL} values are not bounded and can in principle range from $-\infty$ to ∞ . For large $|f_{\text{NL}}|$ (or combinations of relatively small $|f_{\text{NL}}|$ but large $|\Phi|$), $\mathcal{P}(f_{\text{NL}}|\Phi)$ has discontinuities which prevent any expansion in f_{NL} . It is only when working in a regime away from the discontinuities that expansions in f_{NL} are meaningful and moments of the distribution are well-defined so the the central limit theorem applies.

4.3 Estimators vs. full PDF

We continue our detour to discuss the choice of an estimator for f_{NL} given the exact and approximated expressions of the PDF in the simplified one-point case discussed above. We use this simplified case for clarity, as we will see below, it still captures all the interesting physics, features and limitations of

$f_{\text{NL}}\sigma$	1st order Babich(2005)	2nd order	second-order Edgeworth	Hybrid
0.1	3	4	5	4
0.01	4	6	6	6
0.001	7	10	10	8

Table 2. Number of sigmas where the different approximations fit the exact PDF to better than 20%. Recall that for simplicity throughout this section we have taken the field to have unit variance.

known estimators. As we will see below, this exercise is useful to clarify some common misconceptions. While all statistical information is enclosed in the full PDF, if the PDF is complicated or hard to compute it is often useful to work with estimators. An estimator represents a massive reduction of information: in this case for example the full PDF is substituted by a single number (the estimated f_{NL} value). It is worth recalling that the maximum likelihood estimator (MLE) is often desirable because, if it exists, under mild conditions, it is asymptotically the best unbiased estimator. To derive a MLE one interprets the PDF (probability of the data given the parameters) as the likelihood for the parameters given a set of data, and finds its maximum. This conditions give a combination of the data (Φ in our case at hand) which is a MLE of f_{NL} (or f_{nl}) i.e. \hat{f}_{NL} (or \hat{f}_{nl}). The drawback of a MLE is that in many problems the full PDF is not known or if it is known, its maximum cannot be found analytically. For our exact formulation of the PDF, it is clear that the maximum must be found numerically. In addition, in some regimes the PDF diverges making a MLE not well-defined. Fig. (4) –solid line, top panel– shows the locus of the maximum of the likelihood as a function of Φ for the simple one-point PDF investigated above (i.e., in this sub-section we keep considering ϕ to have unit variance). The dotted line is for guidance at $f_{\text{NL}} = 0$. The dashed line corresponds to the estimator of B05. In correspondence of the diamond symbols in the bottom panels we show the shape of the likelihood as a function of f_{nl} for fixed Φ values (0.1, 0.5, 1, 1.5). The likelihood is well-behaved as long as Φ is small. The location of the “cusp” in the curve corresponds to $|\Phi| = 1$ and for $\Phi = 1 + \epsilon$ (where ϵ is a small number) the PDF diverges, the discriminant becomes negative (and the PDF is zero) splitting the divergence through the middle. For $|\Phi| > 1$ there are two points where the PDF diverges, separated by a region where the discriminant is negative and the PDF is zero.

Furthermore, the support of the PDF depends on the parameter, since it is zero for $\Phi < -f_{\text{nl}} - 1/(4f_{\text{nl}})$, and this makes the Cramér-Rao bound invalid [35].

These conditions make the MLE not well-defined. Note that the B05 estimator is not a MLE except for $f_{\text{nl}} = 0$. Its error is shown to satisfy the Cramér-Rao bound, but that just corresponds to an expansion of the PDF i.e. it approximates the PDF as a function of f_{nl} as a Gaussian distribution. The plots in Fig. (2) show that this is not a good approximation: even when Φ is very small the likelihood has broad wings. While it is true that the central limit theorem should make the PDF close to Gaussian for practical applications where many independent measurements are considered, the full PDF shape carries much more information than its mean and the variance, as discussed at length in [36]. Since we have the exact expression for the PDF we can compute the possible bias of the estimator by computing (numerically) the expectation value of it.

Numerically we find that the estimator of B05 is unbiased even for large values of f_{nl} . In particular the bias of the estimator scales like f_{nl}^2 as follows:

$$\text{bias of estimator} = 0.001 \left(\frac{f_{\text{nl}}}{0.1} \right)^2. \quad (4.17)$$

Given that the maximum likelihood estimator is not well-defined and that the full likelihood is known, and that the PDF carries more information than just the mean and variance of an estimator, it is much more useful (although not necessarily simpler) to work directly with the full PDF. This is what we set up to do here.

The above discussion is valid for a single pixel. In any practical application many independent measurements will be combined, and a full analysis would include in-field correlations (e.g. [37]). For illustration, we consider first a full posterior analysis for f_{NL} based on the complete PDF in n pixels.

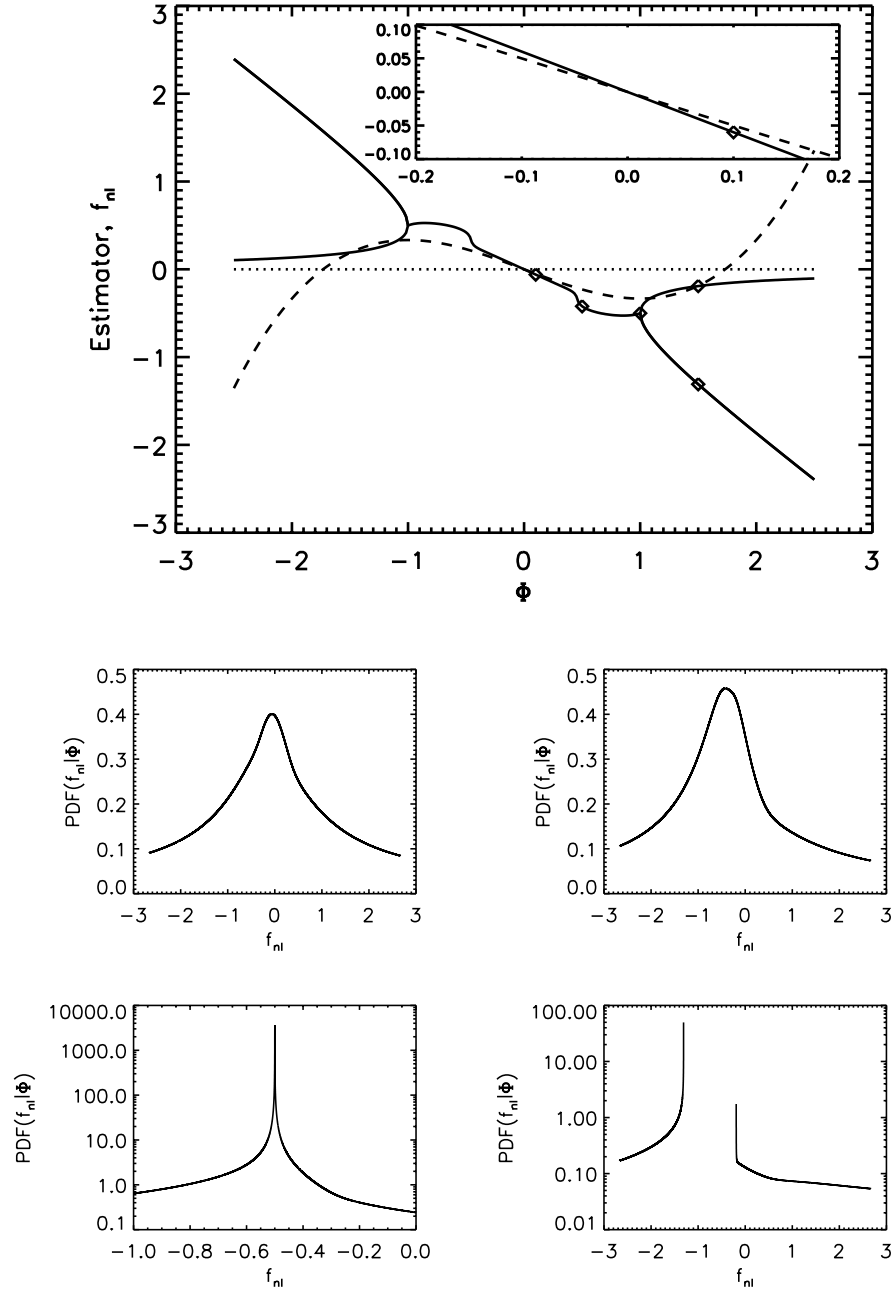


Figure 4. Top panel. Solid line: the numerically-determined maximum likelihood as a function of $f_{nl}(= f_{NL}\sigma)$ for several values of Φ . Here the underlying Gaussian field ϕ is taken to have unit variance ($\sigma = 1$). For $|\Phi| \geq 1$ the PDF diverges making a maximum likelihood estimator not well-defined. Dashed line: the Babich (2005) estimator. The dotted line is for reference at $f_{nl} = 0$. The inset shows a zoom in in the region around $f_{nl} = 0$. The Bottom 4 panels show slices through the f_{nl} likelihood for fixed values of Φ (0.1, 0.5, 1, 1.5 from top left to bottom right), which maximum correspond to the locations of the diamonds. Note that the likelihood for f_{nl} is always highly non-Gaussian. For $|\Phi| > 1$ there are two points where the PDF diverges that are separated by a region where PDF is zero. For $|\Phi| \geq 1$ the derivatives of the likelihood are non continuous.

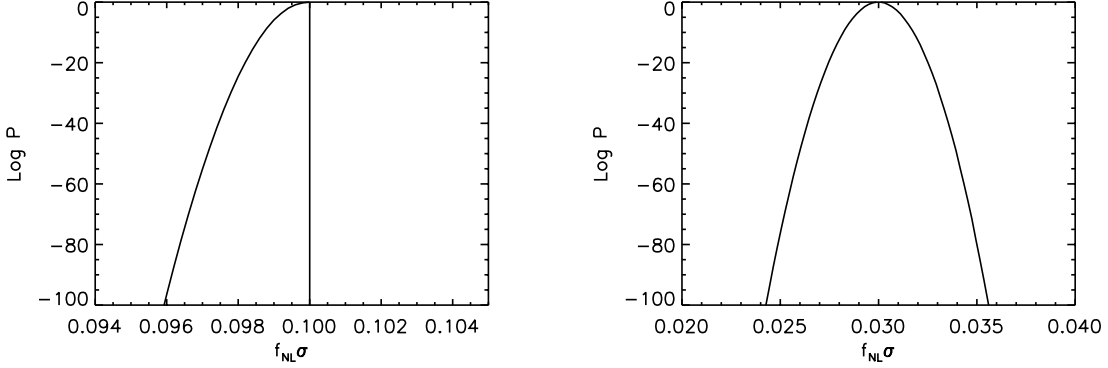


Figure 5. *Left panel:* $\text{Log}(\text{Posterior})$ for $f_{nl}(= f_{NL}\sigma)$, with true $f_{nl}^{\text{true}}=0.1$ and 10^6 pixels but ignoring sampling issues and noise. $f_{nl} > f_{nl}^{\text{true}}$ is excluded, as the true PDF is non-zero for some values of Φ where the trial PDF is zero. *Right panel:* $\text{Log}(\text{Posterior})$ for $f_{nl}(= f_{NL}\sigma)$ given 10^6 pixels, and $f_{nl}^{\text{true}} = 0.03$. A Gaussian noise of rms 0.01 added (recall that the underlying Gaussian field ϕ is normalized to have unit variance). The noise makes all values of Φ reachable in principle, and so the posterior is always non-zero. This removes the strict cutoff apparent in the left panel.

Using Bayes' theorem, the probability of the non-Gaussianity parameter having a value f_{nl} is $\mathcal{P}(f_{nl}|\Phi) \propto \mathcal{P}(\Phi|f_{nl})$, for a uniform prior on f_{nl} . For $i = 1, \dots, n$ pixels, assumed independent, the probability of f_{nl} given a set of pixel values $\{\Phi_i\} \equiv \vec{\Phi}$ (normalized so that the underlying Gaussian field ϕ has unit variance for the case at hand) is

$$\mathcal{P}(f_{nl}|\vec{\Phi}) = \prod_i \mathcal{P}(f_{nl}|\Phi_i) \propto \prod_i \mathcal{P}(\Phi_i|f_{nl})$$

so

$$\ln \mathcal{P}(f_{nl}|\vec{\Phi}) = \sum_i \ln \mathcal{P}(\Phi_i|f_{nl}) + \text{const.}$$

If the true f_{nl} is f_{nl}^{true} , then the probability of a pixel having a value between Φ and $\Phi + d\Phi$ is $\mathcal{P}(\Phi|f_{nl}^{\text{true}})d\Phi$. So the sum over the set of pixels becomes an integral over Φ , with a weight given by $\mathcal{P}(\Phi|f_{nl}^{\text{true}})$ (and a normalisation given by the number of pixels):

$$\ln \mathcal{P}(f_{nl}|\vec{\Phi}) = n \int d\Phi \mathcal{P}(\Phi|f_{nl}^{\text{true}}) \ln \mathcal{P}(\Phi|f_{nl}) + \text{const.}$$

Note that the for the purpose of this illustration the number of pixels enters only as a normalization factor; in fact the sum is evaluated as an integral and the limits of integration are (formally) from $-\infty$ to ∞ . In practice, given a finite number of pixels, the chances of sampling the tails of the distribution are small and thus effectively the limits of integration should depend on the number of pixels.

The above equation implies that $\mathcal{P}(f_{nl}|\vec{\Phi})$ is zero for $f_{nl} > f_{nl}^{\text{true}}$, since there is an excluded region which depends on f_{nl} : $\mathcal{P}(\Phi|f_{nl}) = 0$ for $\Phi < \Phi_{\min}$ where $\Phi_{\min} = -f_{nl} - 1/(4f_{nl})$ (Note that this limit is very small for small f_{nl} , e.g. for $f_{nl} = 0.01$, $\Phi > -25.01$). Hence, if $f_{nl} > f_{nl}^{\text{true}}$, there is a region of Φ where $\mathcal{P}(\Phi|f_{nl}^{\text{true}}) \neq 0$ and $\ln \mathcal{P}(\Phi|f_{nl}) \rightarrow -\infty$ and the integral diverges. This is borne out by doing the integration (see left panel of Fig. (5)) (Note that this argument applies to small f_{nl} , in which case Φ_{\min} increases with f_{nl}). This sharp discontinuity will also prevent any truncated expansion of the PDF being a good approximation.

However, in a practical case, the chances of reaching the 'excluded' region which is forbidden for a small (positive) f_{nl} are basically vanishingly small, so the *cliff* in the posterior is not going to happen in practice, because of sampling, or it could also happen because of noise. For example, if we introduce a small Gaussian noise into the measurements, then the probability of reaching the




Figure 6. Left panel: PDF simulated (3 million Φ 's, thick solid line), analytical exact expression (dashed line), first order (dotted) and second order (dot-dashed). Right panel: % difference compared with simulated PDF, analytical exact expression (solid), first order (dotted) and second order (dashed). Here the underlying Gaussian field ϕ has unit variance and $f_{nl}^{\text{true}} = 0.03$.

otherwise forbidden region of Φ is no longer zero, and this allows values of $f_{nl} > f_{nl}^{\text{true}}$ to have non-zero probability. The right panel of Fig. (5), shows the posterior for local non-Gaussianity with a Gaussian noise added, with rms only 0.01; the zero probabilities have gone away, and the posterior becomes quite Gaussian. This behaviour is quite robust for small smoothing; for large smoothing the posterior gets wider.

The cutoff disappears also if we now consider that for a sample of finite size the tails of the distribution are not well sampled. For example for $f_{nl} = 0.05(0.1)$, the region $\Phi < -5.05(-2.6)$ is not sampled if the sample is smaller than about 1.7×10^6 (80). The disappearance of the *cliff* makes truncated expansions in f_{NL} of the PDF viable. Note however that a truncation to first order will yield to a $\mathcal{P}^{1st}(f_{NL})$ that does not have a maximum, which therefore cannot be a good approximation. At the minimum to have a maximum the approximated PDF must be truncated at second order or higher in f_{NL} .

We next explore the performance of the estimator, the exact PDF and its approximations in the regime where sampling is finite. To do so we generated simulations (*realizations*) of uni-variate non-Gaussian fields for given values of f_{nl} (which we refer to as f_{nl}^{true}) and proceed to recover (measure) it from *samples* of the field (which we call f_{nl}^m , where m stands for *measured*). Before proceeding we verified that the distribution of the realization matches to analytic exact PDF for all f_{nl} values. This is shown in Fig. (6) for $f_{nl}^{\text{true}} = 0.03$.

We concentrate here on two regimes: one where there are few samples (only 50 Φ s; *small sample size*) and one where there are many more samples (1500 Φ s; *large sample size*). For simplicity we assume that the pixels (samples) are uncorrelated. Fig. (7) shows the distribution of the f_{nl}^m obtained using the exact PDF (solid line), the first and second-order PDF (dashed and dotted, almost indistinguishable) and the B05 estimator (Dot-dot-dot-dashed). To be explicit, since the pixels are uncorrelated we can multiply the individual likelihoods $\mathcal{P}(f_{nl}|\Phi_1 \dots \Phi_n) = \prod_i \mathcal{P}(f_{nl}|\Phi_i)$ and f_{nl}^m is the value that maximizes such product.

Clearly the distribution of f_{nl}^m is non-Gaussian. In particular we note that if f_{nl}^{true} is small (0.03 in this example) there is a non-negligible probability that the f_{nl} of the sample is ~ 0.15 if the sample is small. The shape of the PDF for small sample size reflects the behaviour seen in Fig. (4) bottom right panel.

This behaviour disappears when the sample size is increased as shown in Fig.(8). Note that all approximations and the estimator have a variance that is larger than the one given by the exact PDF when f_{nl}^{true} is large. As $f_{nl}^{\text{true}} \rightarrow 0$ the PDF, its approximation and estimator distributions converge. The resulting distribution remains however non-Gaussian. This demonstrates the importance of working with a PDF (exact if f_{nl}^{true} is large, exact or approximated if f_{nl}^{true} is small): the confidence




Figure 7. Distribution of f_{nl}^m for the exact PDF (solid line), first order (dotted) second order (dashed), B05 estimator (Dot-dot-dot-dashed) in the small sample size limit. The thin vertical dotted line shown the location of f_{nl}^{true} (0.03 on the left panel and -0.07 on the right panel).




Figure 8. Distribution of f_{nl}^m for the exact PDF (solid line), first order (dotted) second order (dashed), B05 estimator (Dot-dot-dot-dashed) in the large sample size limit. The thin vertical dotted line shows the location of f_{nl}^{true} . Note that all approximations overestimate the width of the true distribution for large f_{nl}^{true} . As $f_{nl}^{\text{true}} \rightarrow 0$ the PDF, its approximation and estimator distributions converge.

intervals for $\gg 1\sigma$ or $\gg 68\%$ levels, cannot be estimated from the *rms*. In the absence of a PDF, many simulations of the experiment would be required to estimate these confidence intervals via Monte Carlo simulation. On the other hand knowing the PDF allows one to compute any desired confidence interval. Although in this section we have worked in the case where \mathcal{M} is the identity matrix, we expect that these findings hold qualitatively in the general case.

Our detour ends here; we now go back at deriving analytic expressions for the full PDF for increasingly complicated (realistic) cases.

4.4 If \mathcal{M} is not the identity matrix

If \mathcal{M} is not the identity matrix but it is an invertible square matrix then we need to go back to the expression for \mathcal{P} as:

$$\mathcal{P}(\vec{A}|f_{\text{NL}}) = \sqrt{\frac{\det K_0}{(2\pi)^M}} \int \prod_{i,\nu} d\phi_i d\phi_\nu \prod_i \delta^D(A_i - \sum_a \mathcal{M}_{ia} \Phi_a) \times \exp \left[-\frac{1}{2} \sum_{ij} \phi_i \phi_j K_{0ij} - \sum_{i,\nu} \phi_i \phi_\nu K_{0i\nu} - \frac{1}{2} \sum_{\mu,\nu} \phi_\mu \phi_\nu K_{0\mu\nu} \right]. \quad (4.18)$$

If $n = M$ then one could do

$$A_b \longrightarrow \sum_a (\mathcal{M}^{-1})_{ba} A_a \equiv \tilde{A}_b, \quad (4.19)$$

which also introduces a factor of $1/\det \mathcal{M}^{-1} = \det \mathcal{M}$ in front of the RHS. This complicates the Dirac delta function but leaves us with a Gaussian integral.

We obtain

$$\mathcal{P}(\vec{A}|f_{\text{NL}}) = \sqrt{\frac{\det K_0}{(2\pi)^M}} \det \mathcal{M} \prod_a \frac{1}{\sqrt{1 + 4f_{\text{NL}}(\tilde{A}_a + C)}} \times \sum_{rq} \left\{ \exp \left[-\frac{1}{2} \sum_{ab} f_r(\tilde{A}_a) f_q(\tilde{A}_b) K_{0ab} \right] \right\}; \quad (4.20)$$

if we keep only one root for small f , $r = q = 1$.

For a possible practical application \mathcal{M} is a square invertible matrix; for example when A corresponds to the matter overdensity field δ be it in real or Fourier space. If $n < M$ this route cannot be used (see below for the approach to use) but one could always write :

$$\mathcal{P}(\vec{A}|f_{\text{NL}}) = \int \mathcal{D}[\mathcal{A}] \mathcal{P}(\mathcal{A}) \delta^D(A - \mathcal{A}). \quad (4.21)$$

So for example for the one-point PDF one should do: $\mathcal{P}(A_1) = \int \mathcal{D}[\mathcal{A}] \mathcal{P}(\mathcal{A}) \delta^D(A_1 - \mathcal{A})$. By applying the Dirac delta one reduces the dimensionality of the path integral by one. The remaining integral might still have to be performed numerically.

4.5 \mathcal{M} not invertible, $n < M$

To obtain directly the probability of $\mathcal{P}(\vec{A})$ if $n < M$ and \mathcal{M} is not invertible let us go back to Eq. (4.18) and as before let us now split the Dirac delta function so its argument is given by the elements of:

$$A - \aleph \Phi' - \mathcal{N} \Phi'' = A - \aleph \phi' - \aleph f_{\text{NL}} \phi'^2 - \mathcal{N} \phi'' - f_{\text{NL}} \mathcal{N} \phi''^2 + f_{\text{NL}} \sigma_\phi^2 c \quad (4.22)$$

where the elements of the matrix c are given by $c_i = \sum_a \mathcal{M}_{ia}$. From now on let us redefine $c = f_{\text{NL}} \sigma_\phi^2 c$.

Now note that to impose the Dirac delta function one must first integrate in ϕ' as such it is convenient first to change variables so that the argument of the Dirac delta function is multiplied by \aleph^{-1} , which yields a $|\det \aleph|$ in front of the equation and then it is useful to define :

$$\tilde{A}(\phi''|f_{\text{NL}}) = \aleph^{-1} A - \aleph^{-1} \mathcal{N} \phi'' - f_{\text{NL}} \aleph^{-1} \mathcal{N} \phi''^2 + \aleph^{-1} c. \quad (4.23)$$

with this the roots of the quadratic equation are:

$$\phi_{r,i} = \frac{-1 \pm \sqrt{1 + 4f_{\text{NL}} \tilde{A}_i(\phi''|f_{\text{NL}})}}{2f_{\text{NL}}}, \quad (4.24)$$

and the Dirac delta function becomes

$$\frac{\delta^D(\phi_i - \phi_{r=1}) + \delta^D(\phi_i - \phi_{r=2})}{\sqrt{1 + 4f_{\text{NL}}\tilde{A}_i}}. \quad (4.25)$$

Thus we arrive at

$$\mathcal{P}(\vec{A}|f_{\text{NL}}) = \sqrt{\frac{\det K_0}{(2\pi)^M}} \det \mathbb{N} \int d\phi'' \prod_i \left(\sqrt{1 + 4f_{\text{NL}}\tilde{A}_i(\phi''|f_{\text{NL}})} \right)^{-1} \exp \left[-\frac{1}{2} \sum_{\nu\mu} \phi_\mu \phi_\nu K_{0\mu\nu} \right] \times \sum_{r=1,2} \exp \left[-\frac{1}{2} \left(\sum_{ij} \phi_{r,i} K_{0ij} \phi_{r,j} + 2 \sum_{i\nu} \phi_{r,i} K_{0i\nu} \phi''_\nu \right) \right]. \quad (4.26)$$

Again one can safely ignore the exponentially suppressed root. Note that even if only an integral in ϕ'' is left, there are complications to perform this analytically: ϕ'' appears not only in front of the exponential as an argument of the function \tilde{A} (under the square root), but also in the exponential both under square root in the ϕ_r and explicitly.

This equation is one of the main results of this paper as it can be interpreted as the probability of f_{NL} given the observable vector A . The expression is exact but the path integral must be performed numerically for different f_{NL} values. Note that the dimensionality of the integration has been reduced by a factor L compared to e.g., Eq. (4.11)– i.e. for a full sky CMB map the dimensionality of the integral has been reduced by a factor $\mathcal{O}(10^6)$. The integrand is close to a Gaussian so the numerical integration should converge easily but detailed implementation will depend on the particular application one is interested in. We will return on this in Sec. (5).

4.6 Approximations

In order to progress analytically approximations must be made. Taylor-like expansions around $f_{\text{NL}} = 0$ could be considered and they could be chosen so that the integral (Eq. 4.26) becomes analytic, in the same spirit as Sec. (4.2). Sec. (4.2) and (4.3) indicate that a second-order (Edgeworth) expansion is a good approximation for the full PDF if f_{NL} is small and the sample size is reasonably large. While we have quantified this numerically for the single pixel and Poisson, uncorrelated pixels case, we expect that qualitatively the same will hold in the more general, multi-variate case. While for the one-point PDF where \mathcal{M} is the identity matrix it is simple to expand directly the PDF, as done in Sec. (4.2), it becomes, still doable, but unnecessarily complicated if we want to work with Eq. (4.26). In fact the work leading up to Eq. (4.26) and the route we chose, are necessary if one wants to find a closed form for the exact PDF. If one instead wanted *a priori* to obtain an approximated expression, truncated at a given order, the simplest route to follow would have been different. In fact it is well known that a general expression for any (well behaved) PDF can be obtained from its generating functional by expanding the PDF in terms of all its (2nd and higher-order) correlations.

This expansion can be written as (see e.g., [39], the integrand of their Eq. 9, and [40]):

$$\mathcal{P}(\vec{A}|f_{\text{NL}}) = \frac{\det \mathbf{K}^{1/2}}{(2\pi)^{n/2}} \exp \left[\sum_{i=3}^{\infty} (-1)^i \sum_{[r_i]=1}^n \frac{\xi^{(i)}_{[r_i]}}{i!} \prod_{j=1}^i \frac{\partial}{\partial A_{r_j}} \right] \exp \left[-\frac{1}{2} (\vec{A}, \mathbf{K}, \vec{A}) \right] \quad (4.27)$$

where \mathbf{K} denotes the inverse covariance matrix for \vec{A} and $[r_i]$ is a compact notation to account for all relevant combinatorials for example for $i = 3$ the second sum runs over all triplets that can be combined out of the $A_1 \dots A_n$. Here $\xi^{(i)}$ denotes the connected correlation function (or corresponding polyspectra) of order i . This expression is exact, but it involves an infinite sum which must be truncated. In fact this expression simply says that any non-Gaussian distribution is specified by *all* its higher order correlations (only a Gaussian distribution is completely specified by its covariance).

In other words Eqs. (4.26) and (4.27) are mathematically equivalent for local non-Gaussianity once the ξ^n are the ones of the field A . For a practical application however, when the exact PDF is needed, Eq. (4.26) is preferable to Eq. (4.27) because of the infinite series in Eq. (4.27) which would require evaluation of an infinite sequence of correlation functions for A . Conversely, if for some applications an approximated expansion of the PDF is what is needed, it is overwhelmingly easier to work with Eq. (4.27), truncated at the appropriate order.

As an exercise it is not too algebraically cumbersome to see that at first order in f_{NL} Eq. (4.26) and Eq. (4.27) are truly equivalent. The univariate case for unit variance field was presented in Sec. (4.2). For the multi-variate case with a generic covariance, the most cumbersome part is to keep track of the various (symmetrized) permutations that go in $\xi^{(3)}$ for the local non-Gaussian case. Since in the simpler case analyzed in Sec. (4.2) and (4.3) we could compute the exact PDF, its first and second-order expansions, and compare the performance of the exact PDF with its approximations, we conclude that a second-order expansion of Eq. (4.27) is suitable for many interesting practical applications (f_{NL} in agreement with current constraints i.e., Ref. [38], sample size not too small). Here therefore we derive an explicit expression for the multi-variate non-Gaussian PDF to second-order using Eq. (4.27).

We expand to second order in f_{NL} by first truncating the first sum to $i = 4$ – i.e., the trispectrum – and then expanding the exponential in Taylor series to second order.

It is easy to see that if $A = \Phi$, $n = 1$ and Φ has unit variance, we recover Eq. (4.15). In doing so it is useful to keep in mind that for local non-Gaussianity $\xi^{(3)} = 6f_{\text{NL}}$ and $\xi^{(4)} = 24f_{\text{NL}}^2$.

In the general case by performing the derivatives we obtain:

$$\begin{aligned} \mathcal{P}(\vec{A}|f_{\text{NL}}) = \frac{\sqrt{\det \mathbf{K}}}{(2\pi)^{n/2}} \exp\left[-\frac{1}{2}(\vec{A}, \mathbf{K}, \vec{A})\right] \times \left\{ 1 + \frac{1}{6} \sum_{\text{all } l} \xi_{l_1 l_2 l_3}^{(3)} \left[(\mathbf{K}\vec{A})_{l_1} (\mathbf{K}\vec{A})_{l_2} (\mathbf{K}\vec{A})_{l_3} - 3\mathbf{K}_{l_1 l_2} (\mathbf{K}\vec{A})_{l_3} \right] + \right. \\ \frac{1}{24} \sum_{\text{all } l} \xi_{l_1 l_2 l_3 l_4}^{(4)} \left[3K_{l_1 l_2} K_{l_3 l_4} - 6\mathbf{K}_{l_1 l_2} (\mathbf{K}\vec{A})_{l_3} (\mathbf{K}\vec{A})_{l_4} + (\mathbf{K}\vec{A})_{l_1} (\mathbf{K}\vec{A})_{l_2} (\mathbf{K}\vec{A})_{l_3} (\mathbf{K}\vec{A})_{l_4} \right] + \\ \frac{1}{72} \sum_{l_1, \dots, l_6} \xi_{l_1 l_2 l_3}^{(3)} \xi_{l_4 l_5 l_6}^{(3)} \left[(\mathbf{K}\vec{A})_{l_1} (\mathbf{K}\vec{A})_{l_2} (\mathbf{K}\vec{A})_{l_3} (\mathbf{K}\vec{A})_{l_4} (\mathbf{K}\vec{A})_{l_5} (\mathbf{K}\vec{A})_{l_6} - 15\mathbf{K}_{l_1 l_2} \mathbf{K}_{l_3 l_4} \mathbf{K}_{l_5 l_6} \right. \\ \left. \left. - 15\mathbf{K}_{l_1 l_2} (\mathbf{K}\vec{A})_{l_3} (\mathbf{K}\vec{A})_{l_4} (\mathbf{K}\vec{A})_{l_5} (\mathbf{K}\vec{A})_{l_6} + 45\mathbf{K}_{l_1 l_2} \mathbf{K}_{l_3 l_4} (\mathbf{K}\vec{A})_{l_5} (\mathbf{K}\vec{A})_{l_6} \right] \right\} \end{aligned} \quad (4.28)$$

where l_i run from 1 to n . Here we recognize that in this approximation the PDF is given by a suitable combination of bispectrum and trispectrum. We will return to this point below in the application to the CMB in Sec. (5). The great advantage of this approach is that the resulting expression is valid for non-Gaussianity beyond the local type. While we have tested quantitatively the performance of this approximation for the local case, we expect that it will hold also e.g., for equilateral, orthogonal or flattened type. Of course, Eq. (4.28) is strictly second-order in f_{NL} for non-Gaussianity of the type where $\Phi^{NG} \sim f_{\text{NL}} \phi * \phi$ ($*$ denoting convolution or multiplication). Should other types of non-Gaussianity be considered, then the expression would not be necessarily second-order in the non-Gaussianity parameter. For example for the g_{NL} type where $f_{\text{NL}} = 0$, the terms with $\xi^{(3)}$ would drop out and the resulting expression would be first-order in g_{NL} .

Given that in a practical application the A_i are measured (given) quantities, the PDF can be interpreted as the probability distribution for f_{NL} . Within a given model (like for example the local model) where the bispectrum and trispectrum coefficients are related (one being $\propto f_{\text{NL}}$ and the other $\propto f_{\text{NL}}^2$) the PDF can be used to constrain the parameter (i.e. find the maximum likelihood and the confidence intervals). In a more model-independent way, one could imagine finding joint constraints on the coefficients of the bispectrum and trispectrum and then comparing those with theory predictions. In light of these considerations and the fact that for f_{NL} values allowed by present data and for not too small sample sizes this expression gives and an extremely good approximation for the exact PDF, Eq. (4.28) represents one of the main results of this paper. For a CMB application of Eq. (4.28), see Sec. (5).

4.7 Adding Gaussian (or near-Gaussian) noise

As before (and using the same convention as of Sec. (3.4)) we can derive the result in two ways. The direct way is to recall that $\mathcal{P}(\mathcal{A}) = \mathcal{P}(A) \star \mathcal{P}(\epsilon)$ which will be further explored below (Sec. 5.2) in the context of measuring f_{nl} from a CMB map. In fact, although the approach is general, the implementation depends on the details of the concrete application. This approach might be best if the noise has complicated pixel-to-pixel correlations. The other way is to substitute the Dirac delta function by a Gaussian with width given by the noise covariance Σ_ϵ or by the noise distribution if it is non-Gaussian. Here we illustrate the approach in the Gaussian noise case.

In this case our starting equation becomes Eq. (2.7) but where the product of the Dirac delta function is instead $1/(2\pi)^{n/2} \exp[-1/2(A - \mathcal{M}\Phi, K_\epsilon, A - \mathcal{M}\Phi)]$ where K_ϵ denotes the inverse covariance of the noise and $\mathcal{M}\Phi$ should be interpreted as matrix times a vector. Unfortunately the resulting path integral cannot be done exactly analytically but can be performed numerically. However even in the signal-only regime the last path integral has to be performed numerically (see discussion around Eq. (4.26)) thus correctly including the effect of noise does not represent an additional computational burden.

On the other hand, if the noise contribution is sub-dominant compared to the signal (and yet non-negligible) a saddle point approximation performs very well: Eq. (4.26) then becomes the first term in a saddle point approximation. Note that Ref.[22] already noted that the presence of Gaussian noise has the effect of changing the Dirac delta functions for Gaussians. They approach the problem by implementing numerically the expressions corresponding with our Eq. (2.7) and Eq. (4.11).

5 Application I: measuring f_{nl} from a CMB map

5.1 Signal-only regime

It is possible to consider Eq. (4.26) and interpret it as the probability of f_{NL} given the observable vector \vec{A} (which could be $a_{\ell m}$ or T etc.). This expression is exact. One would need to still do a path integral numerically but only for a grid of values for f_{NL} i.e. only a finite, reduced number of times, to find the maximum likelihood and its confidence levels. The path integral can be computed using standard lattice QCD techniques. The integrand is close to a Gaussian so integration should converge easily. This makes possible to increase the size of the lattice, which sample ϕ'' in Eq. (4.26). In QCD applications the size of the lattice can be as large as billions, comparable with the size we estimate would be needed for a CMB application. Alternatively Monte-Carlo techniques such as those explored in [28] can be employed. We leave the numerical implementation for the evaluation of Eq. (4.26) for future work, but note that *a)* with respect to the formulation of Eq. (4.11) the dimensionality of the integral has been reduced by a factor $\mathcal{O}(10^6)$ and *b)* given a set of data (e.g., a CMB map) the integral for several f_{NL} values must be computed only once to obtain the best-fit value and confidence intervals.

This approach of using the PDF does not have the problem of the bispectrum that is that the number of triplets one can make grows much faster than the number of pixels in a map and therefore different bispectrum estimates are not independent (and thus the central limit theorem needs not to apply) [17]. One possible limitation is that the expression is exact only in the signal-dominated regime (i.e. no noise). This can be overcome as we discuss below.

5.2 Adding noise

As done in Sec. (3.4) consider that the observable field is \mathcal{A} and is given by the superposition of two independent random processes, the signal A and the noise ϵ , ($\mathcal{A}_i = A_i + \epsilon_i$). For the sum (superposition) of two independent random processes the probability distribution is the convolution of the two probabilities. Doing the convolution however might not be possible analytically. One thing to keep in mind is that if the noise is uncorrelated one need not do a N_{pix} -dimensional convolution, but one can do N_{pix} 1-dimensional convolutions. On the other hand, what we expect in the presence of noise is that the probability distribution of f_{NL} is broadened by the noise:

$$\mathcal{P}_{\mathcal{A}}(f_{\text{NL}}|\mathcal{A}) = \mathcal{P}_A(f_{\text{NL}}|A) \star \mathcal{P}_\epsilon(f_{\text{NL}}|\epsilon) \quad (5.1)$$

i.e. a convolution between the result of the signal-only regime with the distribution for the noise-only regime. Given that the noise has zero mean, $P_\epsilon(f_{\text{NL}}|\epsilon)$ must only depend on the noise covariance Σ_ϵ . The advantage of this approach is that one is left to do only a 1 dimensional convolution.

While $P_\epsilon(f_{\text{NL}}|\epsilon)$ has not been computed, if the noise is Gaussian or close to Gaussian, we know that it must be very well approximated by the distribution for the KSW estimator [34] applied to a purely Gaussian map (as it saturates the Cramér-Rao bound [36]). This distribution has been shown to be very well approximated by a Gaussian [20] –since for a large number of pixels the central limit theorem holds– with variance $\sigma_{f_{\text{NL}}}^2 = 1/(6\sigma_\epsilon^2 N_{\text{pix}})$. In summary, since the noise is not correlated with the signal, in the presence of noise the signal-only PDF, Eq. (4.26), must be convolved with the noise-only PDF for f_{NL} . If the noise is Gaussian, the noise-only PDF has a closed analytical expression; if not, it would need to be computed on a case-by-case basis. Having $P_{\mathcal{A}}(f_{\text{NL}}|\mathcal{A})$ already computed on a grid of f_{NL} values (and $P_\epsilon(f_{\text{NL}}|\epsilon)$ as an analytic expression or computed on the same grid values for f_{NL}) makes the convolution a discrete sum. Of course this is useful if the noise-only PDF is less broad than the signal-only PDF (i.e., significantly non zero over a smaller range of f_{NL} values than the signal-only PDF). If not, then one would be working in the noise-dominated regime where little useful information on the primordial f_{NL} can be extracted, or, alternatively, where the signal PDF is only a correction to the noise-only PDF. Thus rough approximations of the signal PDF can be safely used in this regime.

Another equivalent approach is the one outlined in sec. (4.7): substitute the Dirac delta function in Eq.(4.18) by the noise distribution. Then Monte-Carlo like approaches as adopted e.g. in [22] can be used to evaluate the multi-dimensional integral. Depending on the nature of the noise one or the other will be best suited.

5.3 Second-order PDF expansion

Eq.(4.28) as it is can be interpreted as the PDF given the temperature pixels in a CMB map (i.e. $A = T$ or $A = \Delta T/T$, n the number of pixels, l the pixel index and w^i the i -point correlation function). On the other hand it is useful to re-write it explicitly as a function of the spherical harmonic coefficients $a_{\ell m}$ also because it will make explicit the relation of our findings with previous work in the literature. In the following we denote by a the vector of a_ℓ^m and by n the number of harmonic coefficients considered. We obtain:

$$\begin{aligned} \mathcal{P}(a|f_{\text{NL}}) = & \frac{(\det C^{-1})^{1/2}}{(2\pi)^{n/2}} \exp \left[-\frac{1}{2} \sum_{\ell\ell'mm'} a_\ell^{*m} (C^{-1})_{\ell m \ell' m'} a_{\ell'}^{m'} \right] \times \\ & \left\{ 1 + \frac{1}{6} \sum_{\text{all } \ell_i m_j} \langle a_{\ell_1}^{m_1} a_{\ell_2}^{m_2} a_{\ell_3}^{m_3} \rangle \left[(C^{-1}a)_{\ell_1}^{m_1} (C^{-1}a)_{\ell_2}^{m_2} (C^{-1}a)_{\ell_3}^{m_3} - 3(C^{-1})_{\ell_1, \ell_2}^{m_1 m_2} (C^{-1}a)_{\ell_3}^{m_3} \right] + \right. \\ & \quad \frac{1}{24} \sum_{\text{all } \ell m} \langle a_{\ell_1}^{m_1} a_{\ell_2}^{m_2} a_{\ell_3}^{m_3} a_{\ell_4}^{m_4} \rangle \left[3(C^{-1})_{\ell_1 \ell_2}^{m_1 m_2} (C^{-1})_{\ell_3 \ell_4}^{m_3 m_4} \right. \\ & \quad \left. - 6(C^{-1})_{\ell_1, \ell_2}^{m_1 m_2} (C^{-1}a)_{\ell_3}^{m_3} (C^{-1}a)_{\ell_4}^{m_4} + (C^{-1}a)_{\ell_1}^{m_1} (C^{-1}a)_{\ell_2}^{m_2} (C^{-1}a)_{\ell_3}^{m_3} (C^{-1}a)_{\ell_4}^{m_4} \right] + \\ & \quad \frac{1}{72} \sum_{\ell_1, \dots, \ell_6} \langle a_{\ell_1}^{m_1} a_{\ell_2}^{m_2} a_{\ell_3}^{m_3} \rangle \langle a_{\ell_4}^{m_4} a_{\ell_5}^{m_5} a_{\ell_6}^{m_6} \rangle \left[(C^{-1}a)_{\ell_1}^{m_1} (C^{-1}a)_{\ell_2}^{m_2} (C^{-1}a)_{\ell_3}^{m_3} (C^{-1}a)_{\ell_4}^{m_4} (C^{-1}a)_{\ell_5}^{m_5} (C^{-1}a)_{\ell_6}^{m_6} \right. \\ & \quad \left. - 15(C^{-1})_{\ell_1 \ell_2}^{m_1 m_2} (C^{-1}a)_{\ell_3}^{m_3} (C^{-1}a)_{\ell_4}^{m_4} (C^{-1}a)_{\ell_5}^{m_5} (C^{-1}a)_{\ell_6}^{m_6} \right. \\ & \quad \left. - 15(C^{-1})_{\ell_1 m_2}^{m_1 m_2} (C^{-1})_{\ell_3 \ell_4}^{m_3 m_4} (C^{-1})_{\ell_5 \ell_6}^{m_5 m_6} + 45(C^{-1})_{\ell_1 \ell_2}^{m_1 m_2} (C^{-1})_{\ell_3 \ell_4}^{m_3 m_4} (C^{-1}a)_{\ell_5}^{m_5} (C^{-1}a)_{\ell_6}^{m_6} \right] \} \end{aligned} \quad (5.2)$$

Note that in the second line of Eq. (5.2) we recognize the standard formulation for approximating the PDF [33] which is the starting point to derive the standard KSW estimator [34]. In fact we recognize Eq. (5) of [41]. In the next two lines we recognize Eq. (32), but more specifically Eq. (158), of [42].

The last three lines are “new” terms that arise from expanding the exponential to second order in f_{NL} thus involving a term proportional to the bispectrum squared.

We argue here that interpreting this as a likelihood for f_{NL} enables one to combine optimally bispectrum and trispectrum measurements and obtain both best-fit value and confidence intervals for the non-Gaussianity parameter. As discussed above this expression is valid for non-Gaussianities more general than the local form as long as departures from Gaussianity are small as quantified in Sec. (4.3). This expression is strictly second order in f_{NL} for the type of non-Gaussianity where $\Phi^{NG} \sim f_{\text{NL}}\phi * \phi$ (* denoting convolution or multiplication) which includes local, equilateral, flattened and orthogonal shapes. For other types of non-Gaussianity, this expression will not necessarily be second order in the non-Gaussianity parameters. Within a given non-Gaussian model where the bispectrum and trispectrum amplitudes are specified by a single parameter, this PDF can be used to constrain such a parameter. Moreover, in a more model-independent approach one could use the above PDF to find joint constraints on the amplitude of the bispectrum and trispectrum for comparison with theory. Eq. (5.2) is one of the main results of this paper. Moreover note that this approximation to the PDF, being quadratic in f_{NL} , always has a maximum. In the standard formulation, where the PDF expansion is truncated to first order in f_{NL} , the PDF depends linearly on f_{NL} and therefore at face value has no maximum. Implicitly, the quadratic term is assumed to have a constant coefficient (the average of the true coefficient in the case of B05), and the estimator normalised later to be unbiased. With the complete PDF, it is easy to write down the maximum after expansion to second order. We rewrite Eq. (5.2) in a compact form

$$\mathcal{P}(a|f_{\text{NL}}) = G * [1 + f_{\text{NL}}b * w_1 + f_{\text{NL}}^2(t * w_2 + h * w_3)] \quad (5.3)$$

where G denotes the Gaussian factor, $b * w_1$ corresponds to the part on the second line of Eq. (5.2) that is proportional to f_{NL} evaluated for $f_{\text{NL}} = 1$, and similarly the terms $t * w_2$ and $h * w_3$ are defined by that equation. Then the formal maximum would be

$$\hat{f}_{\text{NL}} = -\frac{b * w_1}{2(t * w_2 + h * w_3)} \quad (5.4)$$

and its variance $1/(t * w_2 + h * w_3)$.

6 g_{NL} case

We can extend our formalism to the so-called g_{NL} model, in which

$$\Phi = \phi + g_{\text{NL}}(\phi^3 - 3\langle\phi^2\rangle\phi) \equiv \phi + g'_{\text{NL}}\phi^3, \quad (6.1)$$

which corresponds to the case when f_{NL} is zero. Having worked out in pedagogical detail the Gaussian case and the local f_{NL} one, here we will only sketch the derivation. In fact the steps are closely similar, only the resulting expressions (given the increasing number of roots involved) become quickly long and cumbersome. If \mathcal{M} is the identity matrix, we obtain

$$\begin{aligned} \mathcal{P}(\vec{A}) = & \sqrt{\frac{\det K_0}{(2\pi)^M}} \prod_i \frac{1}{|1 + 3g_{\text{NL}}\phi_r^2|} \int \prod_\nu d\phi_\nu \exp \left[-\frac{1}{2} \sum_{\mu,\nu} \phi_\mu \phi_\nu K_{0\mu\nu} \right] \\ & \times \exp \left[-\frac{1}{2} \sum_{ij} f_r(A_i) f_r(A_j) K_{0ij} - \sum_{i,\nu} f_r(A_i) \phi_\nu K_{0i\nu} \right] \end{aligned} \quad (6.2)$$

where ϕ_r is the real root of the third order equation $A - \phi - g_{\text{NL}}\phi^3 = 0$. The integral in ϕ can be performed analytically as done before. If \mathcal{M} is invertible but not the identity matrix, then $A \rightarrow \mathcal{M}^{-1}A$ and a factor of $\det \mathcal{M}$ should be included in front of Eq. (6.3). In the general case where \mathcal{M} is not invertible, the technique presented in Sec. (4.5) can be similarly applied here to obtain a closed expression. The resulting expression can be obtained straightforwardly, but it is long and cumbersome and therefore we do not write it explicitly here.

7 Application II: Mass function of dark matter halos

Note that in the case where we need the PDF of the linearly extrapolated δ^L , $\mathcal{P}(\delta^L)$, \mathcal{M} is invertible and we have obtained above an exact analytic expression:

$$\mathcal{P}(\delta^L|f_{\text{NL}}) = \frac{1}{(2\pi)^{M/2}} \frac{1}{\sigma_R} \frac{1}{\sqrt{1 + 4f_{\text{NL}}(\tilde{\delta}^L + f_{\text{NL}}\sigma_{\phi,R}^2)}} \sum_{r=1,2} \left\{ \exp \left[-\frac{1}{2} \frac{f_r^2(\tilde{\delta}^L)}{\sigma_{\phi,R}^2} \right] \right\}. \quad (7.1)$$

Here δ^L is the linearly extrapolated overdensity field smoothed on scale R , corresponding to the halo mass one is interested in; σ_R denotes the RMS of δ^L (smoothed on scale R); $\sigma_{\phi,R}$ is the corresponding quantity for the primordial field ϕ . The sum is over the two roots (recall that the exponentially suppressed root can be safely neglected). $\tilde{\delta} = \mathcal{M}^{-1}\delta$. Finally $f_{1,2}(\delta) = 1/(2f_{\text{NL}}) (-1 \pm \sqrt{1 + 4f_{\text{NL}}\delta})$. While in principle the inversion of \mathcal{M} here requires solving the Poisson equation, recall that in our approach the fields are pixelized and therefore techniques developed for setting the initial conditions for dark matter cosmological N-body simulations can be directly applied here. $\mathcal{P}(\delta^L)$ is the basic ingredient to compute the mass function. In the old-fashioned Press-Schechter approach one only had to integrate the one-point PDF for δ^L above the threshold for collapse δ_c (and introduce the fudge factor of 2). In this framework our approach gives an *exact* expression for the Press-Schechter mass function, which should not have the problem of negative PDF tails (as in the Edgeworth expansion, see [24, 43]) or being inaccurate for not-so-rare peaks [23]. The standard Press-Schechter approach fails in many respects but the one-point PDF remains the key ingredient for more sophisticated formulations of the mass function [43–46]. Ref. [47] showed the equivalence of simple barrier crossing to more modern and sophisticated approaches and ref. [45] showed how to compute the (Gaussian and non-Gaussian) mass functions starting from the 1-point PDF and using a suitably-defined threshold.

7.1 Application III: three-dimensional matter overdensity (e.g. 21cm surveys)

In the case of surveys that map directly the three-dimensional δ field, the size of the \mathcal{M} matrix is $M = n$ as in Sec. (4.4), which simplifies the computations over the CMB case. Although one still needs to compute the transfer function because non-Gaussianity is in the ϕ field, the equations simplify significantly, as the last path integral can be performed analytically, see Eq. (4.20).

8 Conclusions

We have presented a new formalism to write down a closed expression for the exact multi-variate joint probability distribution function (PDF) of a non-Gaussian field of the local type. We provide an exact formula for such PDF for the non-Gaussianity parameter f_{NL} in Eq. (4.26). This exact formulation has several applications ranging from the halo mass function to the analysis of CMB maps. In its more generic formulation (and this for applications that rely on the generic formulation such as application to CMB maps) the final integration must be performed numerically. The dimensionality of the integral is high but it has been reduced by a factor $\mathcal{O}(10^6)$ compared to the initial and intermediate formulation. While we estimate that the numerical calculation is doable it remains to be seen how computationally intensive it might be in practice for a realistic (Planck-size) CMB map. This is left for future work. However in special cases (e.g., as the fundamental step for the derivation of the halo mass function) the resulting expression is fully analytic. We present the exact analytic formula for the 1-point non-Gaussian PDF in Eq. (4.12) which provides the key ingredient for an exact formulation of the mass function of dark halos; a formula to do so is provided in Eq. (7.1).

These results serve as the basis for a fully Bayesian analysis of the non-Gaussianity parameter which has several advantages compared to the presently-used, state-of-the-art frequentists estimators. Most notably the fact that the PDF encloses the full information on the parameter of interest and therefore prohibitively expensive Monte Carlo estimation of the errors is not needed. While the analytic derivation is done in the signal-only regime, we also discuss how to include (analytically) the effect of instrumental noise and show that Gaussian noise can be included at little extra computational expense.

Having a full, exact PDF expression and formulating the problem in a Bayesian fashion, enabled us to uncover some features of the PDF for local non-Gaussianities that were previously ignored. Notably the fact that the PDF is neither smooth nor continuous, making any approximation based on expansions on the f_{NL} parameter of the PDF sometimes unsuitable (and applicability of the central limit theorem limited). We however identify a regime where a second-order expansion is a good approximation, this happens for a combination of f_{NL} not too large (i.e., not ruled out by current measurements) and a sample size which is not too small (few tens or equivalently for a full sky CMB analysis ($\ell < \ell_{\text{max}}$ where $\ell_{\text{max}} \sim \text{few or larger}$). That is, for a full sky CMB analysis that explores a wide ℓ range, the second-order expansion is a good approximation, however if one wanted to explore localized features or fine scale-dependence, the full PDF behaviour might have to be considered. We present explicit expressions for this second-order expansion approximation both in the general case, Eq. (4.28), and for the specific application to a CMB map, Eq. (5.2). In such expansions we can recognize previously known terms (such as the KSW estimator of Ref. [34], and the optimal trispectrum estimator of Ref. [42]) but also a new term $\propto f_{\text{NL}}^2$. This formulation enables one to combine optimally bispectrum and trispectrum measurements and obtain both best-fit value and confidence intervals for the non-Gaussianity parameter. As an added bonus, this truncated expression is valid for non-Gaussianity beyond the local type. While we have tested quantitatively the performance and regime of validity of this approximation for the local case, we expect that qualitatively it will hold also e.g. for equilateral, orthogonal or flattened type and in general for non-Gaussianity where $\Phi^{NG} \sim f_{\text{NL}}\phi * \phi$ (* denoting convolution or multiplication). Within a given model (like for example the local model) where the bispectrum and trispectrum coefficients are related, the PDF can be used to constrain the parameter (i.e. find the maximum likelihood and the confidence intervals). In a more model-independent way, one could imagine finding joint constraints on the coefficients of the bispectrum and trispectrum and then comparing those with theory predictions. Our formulas may also find application in large-scale galaxy surveys in searching for non-Gaussian features, compute the confidence intervals and in general do statistical inference also comparing and combining different data sets.

Acknowledgments

We thank Miguel Angel Vazquez-Mozo, A. Sabio Vera and Benjamin Wandelt for discussions and Michele Liguori for comments on the draft. LV acknowledges support from FP7-IDEAS-Phys.LSS 240117. LV and RJ are supported by FPA2011-29678- C02-02 and acknowledge hospitality of Galileo Galilei Institute for Theoretical Physics during some initial stages of this work. We thank Imperial College London for travel support which facilitated this collaboration. S.M. acknowledges partial financial support from ASI/INAF Agreement I/072/09/0 for the Planck LFI Activity of Phase E2 and PRIN 2009 project “La Ricerca di non-Gaussianità Primordiale”.

References

- [1] Allen, T. J., Grinstein, B., & Wise, M. B. 1987, Physics Letters B, 197, 66
- [2] Salopek, D. S., & Bond, J. R. 1990, PRD, 42, 3936
- [3] Falk, T., Rangarajan, R., & Srednicki, M. 1993, ApJL, 403, L1
- [4] Gangui, A., Lucchin, F., Matarrese, S., & Mollerach, S. 1994, ApJ, 430, 447
- [5] L. Verde, L. -M. Wang, A. Heavens and M. Kamionkowski, Mon. Not. Roy. Astron. Soc. **313** (2000) L141 [astro-ph/9906301].
- [6] Wang, L., & Kamionkowski, M. 2000, PRD, 61, 063504
- [7] E. Komatsu and D. N. Spergel, Phys. Rev. D **63** (2001) 063002 [astro-ph/0005036].
- [8] V. Acquaviva, N. Bartolo, S. Matarrese and A. Riotto, Nucl. Phys. B **667** (2003) 119 [astro-ph/0209156].
- [9] J. M. Maldacena, JHEP **0305** (2003) 013 [astro-ph/0210603].

- [10] N. Bartolo, E. Komatsu, S. Matarrese and A. Riotto, Phys. Rept. **402** (2004) 103 [astro-ph/0406398].
- [11] J. Norena, C. Wagner, L. Verde, H. V. Peiris and R. Easther, Phys. Rev. D **86** (2012) 023505 [arXiv:1202.0304 [astro-ph.CO]].
- [12] A. Pillepich, C. Porciani and S. Matarrese, Astrophys. J. **662**, 1 (2007) [astro-ph/0611126].
- [13] A. Cooray, Phys. Rev. Lett. **97** (2006) 261301 [astro-ph/0610257].
- [14] X. Chen, Adv. Astron. **2010** (2010) 638979 [arXiv:1002.1416 [astro-ph.CO]].
- [15] Komatsu, E. 2010, Classical and Quantum Gravity, 27, 124010
- [16] Liguori, M., Sefusatti, E., Fergusson, J. R., & Shellard, E. P. S. 2010, Advances in Astronomy, 980523.
- [17] Yadav, A. P. S., & Wandelt, B. D. 2010, Advances in Astronomy, 565248.
- [18] Planck Collaboration XXIV, 2013, [arXiv:1303.5084 [astro-ph.CO]].
- [19] Liguori, M., Yadav, A., Hansen, F. K., et al. 2007, Phys. Rev. D, 76, 105016
- [20] T. L. Smith, M. Kamionkowski and B. D. Wandelt, Phys. Rev. D **84** (2011) 063013 [arXiv:1104.0930 [astro-ph.CO]].
- [21] Vielva, P., & Sanz, J. L. 2009, MNRAS, 397, 837
- [22] F. Elsner, B. D. Wandelt and M. D. Schneider, Astron. Astrophys. **513**, A59 (2010) [arXiv:1002.1713 [astro-ph.CO]].
- [23] S. Matarrese, L. Verde and R. Jimenez, Astrophys. J. **541** (2000) 10 [astro-ph/0001366].
- [24] M. LoVerde, A. Miller, S. Shandera and L. Verde, JCAP **0804** (2008) 014 [arXiv:0711.4126 [astro-ph]].
- [25] P. Creminelli, L. Senatore and M. Zaldarriaga, JCAP **0703** (2007) 019 [astro-ph/0606001].
- [26] F. Elsner and B. D. Wandelt, Astrophys. J. Suppl. **184** (2009) 264 [arXiv:0909.0009 [astro-ph.CO]].
- [27] T. L. Smith, D. Grin and M. Kamionkowski, arXiv:1211.3417 [astro-ph.CO].
- [28] F. Elsner and B. D. Wandelt, Astrophys. J. **724**, 1262 (2010) [arXiv:1010.1254 [astro-ph.CO]].
- [29] T. A. Ensslin, M. Frommert and F. S. Kitaura, Phys. Rev. D **80** (2009) 105005 [arXiv:0806.3474 [astro-ph]].
- [30] X. -c. Luo, Phys. Rev. D **49** (1994) 3810 [astro-ph/9310039].
- [31] T. Okamoto and W. Hu, Phys. Rev. D **66** (2002) 063008 [astro-ph/0206155].
- [32] K. Enqvist and S. Nurmi, JCAP **0510** (2005) 013 [astro-ph/0508573].
- [33] D. Babich, Phys. Rev. D **72** (2005) 043003
- [34] E. Komatsu, D. N. Spergel and B. D. Wandelt, Astrophys. J. **634** (2005) 14 [astro-ph/0305189].
- [35] L. Tenorio, arXiv:061017 [astro-ph.CO]
- [36] M. Kamionkowski, T. L. Smith and A. Heavens, Phys. Rev. D **83** (2011) 023007 [arXiv:1010.0251 [astro-ph.CO]].
- [37] D. Marinucci, Annals of Statistics, **34**, (2006) 1
- [38] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5084
- [39] Matarrese, S., Lucchin, F., & Bonometto, S. A. 1986, ApJLett, 310, L21
- [40] Grinstein, B., & Wise, M. B. 1986, ApJ, 310, 19
- [41] Komatsu, E., Class.Quant.Grav.27: 124010,2010
- [42] Regan, D. M., Shellard, E. P. S., & Fergusson, J. R. 2010, Phys. Rev. D. , 82, 023520
- [43] M. Maggiore and A. Riotto, Astrophys. J. **711** (2010) 907 [arXiv:0903.1249 [astro-ph.CO]].
- [44] G. D'Amico, M. Musso, J. Norena and A. Paranjape, JCAP **1102** (2011) 001 [arXiv:1005.1203 [astro-ph.CO]].
- [45] A. Paranjape, R. K. Sheth and V. Desjacques, arXiv:1210.1483 [astro-ph.CO].

- [46] I. E. Aчитouv and P. S. Corasaniti, Phys. Rev. D **86** (2012) 083011 [arXiv:1207.4796 [astro-ph.CO]].
- [47] S. Ferraro, K. M. Smith, D. Green and D. Baumann, arXiv:1209.2175 [astro-ph.CO].