



Quijote-PNG: Quasi-maximum Likelihood Estimation of Primordial Non-Gaussianity in the Nonlinear Dark Matter Density Field

Gabriel Jung^{1,2}, Dionysios Karagiannis³, Michele Liguori^{1,2}, Marco Baldi^{4,5,6}, William R. Coulton⁷, Drew Jamieson⁸, Licia Verde^{9,10}, Francisco Villaescusa-Navarro^{7,11}, and Benjamin D. Wandelt^{7,12}

¹Dipartimento di Fisica e Astronomia “G. Galilei,”Università degli Studi di Padova, via Marzolo 8, I-35131, Padova, Italy; gabriel.jung@protonmail.com

²INFN, Sezione di Padova, via Marzolo 8, I-35131, Padova, Italy

³Department of Physics and Astronomy, University of the Western Cape, Cape Town 7535, South Africa

⁴Dipartimento di Fisica e Astronomia, Alma Mater Studiorum—University of Bologna, Via Piero Gobetti 93/2, I-40129 Bologna BO, Italy

⁵INAF—Osservatorio Astronomico di Bologna, Via Piero Gobetti 93/3, I-40129 Bologna BO, Italy

⁶INFN—Istituto Nazionale di Fisica Nucleare, Sezione di Bologna, Viale Bertini Pichat 6/2, I-40127 Bologna BO, Italy

⁷Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA

⁸Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Straße 1, D-85748 Garching, Germany

⁹ICREA, Pg. Lluís Companys 23, Barcelona, E-08010, Spain

¹⁰ICC, University of Barcelona, IEEC-UB, Martí i Franquès, 1, E-08028 Barcelona, Spain

¹¹Department of Astrophysical Sciences, Princeton University, 4 Ivy Lane, Princeton, NJ 08544, USA

¹²Sorbonne Université, CNRS, UMR 7095, Institut d’Astrophysique de Paris, 98 bis bd Arago, F-75014 Paris, France

Received 2022 June 15; revised 2022 September 28; accepted 2022 October 5; published 2022 November 22

Abstract

Future large-scale structure surveys are expected to improve current bounds on primordial non-Gaussianity (PNG), with a significant impact on our understanding of early universe physics. The level of such improvements will however strongly depend on the extent to which late-time nonlinearities erase the PNG signal on small scales. In this work, we show how much primordial information remains in the bispectrum of the nonlinear dark matter density field by implementing a new, simulation-based methodology for joint estimation of PNG amplitudes (f_{NL}) and standard Lambda cold dark matter parameters. The estimator is based on optimally compressed statistics, which, for a given input density field, combine power spectrum and modal bispectrum measurements, and numerically evaluate their covariance and their response to changes in cosmological parameters. In this first analysis, we focus on the matter density field, and we train and validate the estimator using a large suite of N -body simulations (QUIJOTE-PNG), including different types of PNG (local, equilateral, orthogonal). We explicitly test the estimator’s unbiasedness, optimality, and stability with respect to changes in the total number of input realizations. While the dark matter power spectrum itself contains negligible PNG information, as expected, including it as an ancillary statistic increases the PNG information content extracted from the bispectrum by a factor of order 2. As a result, we prove the capability of our approach to optimally extract PNG information on nonlinear scales beyond the perturbative regime, up to $k_{\text{max}} = 0.5 h \text{ Mpc}^{-1}$. At the same time, we discuss the significant information on cosmological parameters contained on these scales.

Unified Astronomy Thesaurus concepts: [Non-Gaussianity \(1116\)](#); [Cosmological parameters from large-scale structure \(340\)](#); [Fisher’s Information \(1922\)](#)

1. Introduction

The best cosmological probe of primordial non-Gaussianity (PNG) has been, up to now, the cosmic microwave background (CMB; Akrami et al. 2020). The available information in the angular bispectrum (i.e., the three-point function of harmonic multipoles) of CMB primary anisotropies has been, however, nearly completely extracted. If we want to obtain significant improvements over current PNG constraints for most models, we thus have to look at different observables. For this purpose, galaxy clustering carries a significant potential (Alvarez et al. 2014; Karagiannis et al. 2018; Meerburg et al. 2019), since the 3D galaxy density field contains many more bispectrum modes than the 2D CMB fluctuation fields. Most of these modes are however in the gravitational nonlinear regime. The important challenge to face is therefore to separate, up to the smallest possible scales, the

PNG signal from the late-time NG component, induced by evolution of cosmic structures.

To tackle this issue, significant efforts have been devoted to obtain theoretical predictions of the galaxy bispectrum, via a suitable perturbative treatment valid on the largest scales (for the current state of the art, see Ivanov et al. 2022, and references therein). This approach has even recently led to the first constraints on PNG using both the galaxy power spectrum and bispectrum, by analyzing data from the BOSS survey (Dawson et al. 2013; Cabass et al. 2022a, 2022b; D’Amico et al. 2022).¹³ Recent works based on field-level inference (Baumann & Green 2021; Andrews et al. 2022), neural networks (Giri et al. 2022), or reconstruction methods (Shirasaki et al. 2021), have however shown that much more information should be present in the data. To extract this extra information, alternative promising observables, such as the density probability density function (Mao et al. 2014; Uhlemann et al. 2018; Friedrich et al. 2020), persistent homology (Biagetti et al. 2021, 2022a),

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

¹³ Previous analyses based on the BOSS galaxy bispectrum, in different cosmological contexts, can be found, for example, in Gil-Marín (2015a, 2015b, 2017) and Slepian et al. (2017a, 2017b).

or higher-order correlation functions in Fourier space (Gualdi et al. 2021a, 2021b), have been considered. It has also been confirmed that probing nonlinear scales, using the power spectrum and the bispectrum, improves significantly the constraints on cosmological parameters (Hahn et al. 2020; Hahn & Villaescusa-Navarro 2021). Checking if this is also the case for PNG is then an important question. It is however unclear if analytical approaches can go significantly beyond the mildly nonlinear regime.

In this work, we thus consider instead the alternative, simulation-based route using a large number of realizations of the matter density field. This paper, and its companion paper Coulton et al. (2022), are the first of a series of studies, in which we plan to take a step-by-step approach and address the problem in increasing level of complexity. We start here by discussing the general implementation of our statistical estimation and data compression algorithms based on the modal bispectrum estimator, while in Coulton et al. (2022) we describe in more detail our input simulations and use them to perform an independent Fisher matrix analysis using a binned bispectrum estimator. Moreover, for additional complementarity, the two analyses are performed at different redshifts. In Coulton et al. (2022), the main focus is on $z=0$, where the nonlinear effects are maximal. Here, we study the redshift $z=1$, more suited for comparison with upcoming surveys such as Euclid (Laureijs et al. 2011; Amendola et al. 2018). At this initial stage, we focus our analysis just on the dark matter field, with a twofold aim: to set up the general pipeline that we will also use in the following analyses and to address the first crucial question of how much PNG information can be in principle extracted from the matter field, by pushing the analysis up to small scales, which would be hard to model analytically. In a follow-up paper, we will extend this analysis to halos, whereas in later studies we will finally consider the galaxy density field and account for additional important effects, such as redshift-space distortions and incomplete sky coverage.

We extract the power spectrum and bispectrum of a large number of realizations of the matter density field, allowing us to precisely describe the contribution of PNG down to nonlinear scales ($k_{\max} = 0.5 h \text{ Mpc}^{-1}$), as well as measuring the corresponding covariance matrix, including all non-Gaussian and off-diagonal terms (see Biagetti et al. 2022b for a discussion of their relative importance). We then combine these measurements into optimally compressed statistics, a procedure that has been shown to be extremely efficient in dealing with the galaxy bispectrum in Gualdi et al. (2018, 2019). Finally, following the general scheme developed in Alsing & Wandelt (2018), we derive an estimator to jointly measure Lambda cold dark matter (Λ CDM) and PNG parameters.

This paper is organized as follows. In Section 2, we describe the PNG models that we consider in this work and define the parameters in our analysis; in Section 3, we describe our estimators of the power spectrum and bispectrum and the optimal data compression procedure applied to build a quasi-maximum likelihood estimator of PNG and cosmological parameters; in Section 4, we discuss the application of these estimators to our input mock data sets and show the main results of our analysis; in Section 5, we summarize our conclusions and discuss future prospects.

2. Primordial Bispectrum Shapes

For most inflationary models, the main PNG signature is a nonvanishing bispectrum in the primordial curvature perturbation

field, which is defined as

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

where Φ is related to the comoving perturbation variable on super-horizon scales by $\Phi \equiv \frac{3}{5}\zeta$ and the Dirac delta function enforces translation invariance. The quantity $F(k_1, k_2, k_3)$ defines instead the functional dependence of the bispectrum on different Fourier space triplets—the so-called bispectrum shape—and it depends only on the length of the wavevectors in virtue of rotation invariance. Finally, the dimensionless PNG amplitude parameter f_{NL} measures the strength of the PNG signal, for a given shape. The main goal of our analysis is to build an efficient f_{NL} estimator for the three main bispectrum shapes, namely, the local, equilateral, and orthogonal¹⁴ shapes, defined by the following templates (see Akrami et al. 2020, and references therein):

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

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

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

In the present study we directly work with the matter density field; in such case, for small PNG, nearly all the information on $f_{\text{NL}}^{\text{local}}$, $f_{\text{NL}}^{\text{equil}}$, and $f_{\text{NL}}^{\text{ortho}}$ is contained in the field bispectrum. When dealing instead with observed dark matter tracers, it is well known that an important additional signature of PNG arises in the form of a scale-dependent bias term, sourced by the squeezed configurations of the primordial bispectrum (Afshordi & Tolley 2008; Dalal et al. 2008; Matarrese & Verde 2008; Slosar et al. 2008; Verde & Matarrese 2009; Desjacques & Seljak 2010; hence, especially important for PNG of the local type). In such case, the large-scale power spectrum carries significant extra information on the PNG amplitude. Since at this stage we perform joint estimation of f_{NL} and standard Λ CDM cosmological parameters, the power spectrum is already included in our analysis. Therefore, in follow-up studies that will consider dark matter halos and galaxies, no essential modification of our current data analysis pipeline will be needed, except in the input fields to analyze.

¹⁴ In the companion paper (Coulton et al. 2022), two different orthogonal shapes are discussed: the standard CMB one, which we use here, and a corrected template more accurate for LSS. While this may change the numerical results presented in this work, as they have different behaviors in the squeezed limit, it has no impact on the general conclusions.

3. Estimators

3.1. Summary Statistics

As discussed, the aim of our work is to efficiently combine power spectrum and bispectrum measurements of the matter field in the input simulations, to maximize the PNG information that can be extracted, up to the smallest possible scales. This will be achieved by resorting to a suitable compression scheme, taking the power spectrum and bispectrum as the starting summary statistics and requiring numerical evaluation of their covariances, to build the standard quasi-maximum likelihood estimator described in Alsing & Wandelt (2018). This data compression step is also a key ingredient in performing density-estimation likelihood-free inference (see, e.g., Alsing et al. 2019), an alternative approach we aim to pursue in a future work. Forecasts based on data compression also have the advantage of being conservative, whereas a standard Fisher matrix analysis using the starting set of modes could in principle underestimate errors; this would happen in the case of a lack of statistical convergence in the Monte Carlo averages, due to a low number of input simulations. In the current analysis, we however explicitly show that we can achieve optimality, thanks to fast and efficient estimators to measure the power spectra and bispectra, which are applied to a large enough set of input realizations.

3.1.1. Power Spectrum

To measure the power spectrum, we use the standard estimator (e.g., Feldman et al. 1994)

$$\hat{P}(k_i) = \frac{1}{VN_i} \sum_{\mathbf{k} \in \Delta_i} \delta(\mathbf{k}) \delta^*(\mathbf{k}), \quad (5)$$

where $\delta(\mathbf{k})$ is the density field in Fourier space defined on a grid, the k range has been divided into bins Δ_i of width the fundamental mode, V is the survey volume, and N_i is the number of vectors \mathbf{k} in each bin.

3.1.2. Modal Bispectrum

A well-suited choice to measure the bispectrum statistics is the so-called modal bispectrum estimator, originally developed for CMB NG analysis in Fergusson et al. (2010, 2012a) and then extended to LSS in Fergusson et al. (2012b), Regan et al. (2012), and Schmittfull et al. (2013; see also Lazanu et al. 2016, 2017; Hung et al. 2019a, 2019b; Byun et al. 2021; Byun & Krause 2022).

The main idea of modal estimation is to expand a general weighted bispectrum shape on a factorizable basis:

$$w(k_1, k_2, k_3) B(k_1, k_2, k_3) \simeq \sum_{m,n=0}^{N-1} \gamma_{mn}^{-1} \beta_m Q_n(k_1, k_2, k_3), \quad (6)$$

where $Q_n \equiv q_{1p}(k_1) q_{r}(k_2) q_{s_3}(k_3)$ corresponds to the symmetrized product of one-dimensional basis functions $q_p(k)$ and a partial ordering $n \leftrightarrow (p, q, r)$ has been introduced (see, e.g., Fergusson et al. 2012a), w is a weight function and β are the expansion coefficients. If the input bispectrum is a smooth function of k_1 , k_2 , and k_3 , it can be well approximated by truncating the sum over a number N of modes that is much smaller than the total number of Fourier triangles. $\gamma_{mn} \equiv \langle Q_m | Q_n \rangle$ is the inner product between mode functions, where the following scalar product (see

Appendix A.2 for details on the derivation)

$$\langle f(k_1, k_2, k_3) | g(k_1, k_2, k_3) \rangle \equiv \frac{1}{8\pi^4} \int_{\mathcal{V}_T} dk_1 dk_2 dk_3 \times f(k_1, k_2, k_3) g(k_1, k_2, k_3), \quad (7)$$

is computed over the bispectrum *tetrapyd* domain \mathcal{V}_T , i.e., the set of all the Fourier triplets that form a closed triangle and such that each wavenumber has a length in the chosen range of $[k_{\min}, k_{\max}]$.

Each coefficient β_n can be estimated by fitting the corresponding mode Q_n to the data bispectrum. Making the customary choice of a separable weight function using the total power spectrum $P(k)$:

$$w(k_1, k_2, k_3) \equiv \frac{\sqrt{k_1 k_2 k_3}}{\sqrt{P(k_1) P(k_2) P(k_3)}}, \quad (8)$$

one can show (e.g., Fergusson et al. 2012b) that the modal estimator takes the form:

$$\hat{\beta}_n = \frac{1}{V} \int d^3x M_p(\mathbf{x}) M_q(\mathbf{x}) M_r(\mathbf{x}), \quad (9)$$

where

$$M_p(\mathbf{x}) \equiv \int \frac{d^3k}{(2\pi)^3} \frac{q_p(k) \delta(\mathbf{k})}{\sqrt{k P(k)}} e^{i\mathbf{k} \cdot \mathbf{x}}. \quad (10)$$

This inverse Fourier transform can be computed in an efficient way using fast Fourier transformation (FFT) routines (e.g., FFTW¹⁵), making $\hat{\beta}_n$ extremely fast to extract from data. In this paper, we work directly with the observed modal coefficients β_n to estimate PNG amplitudes f_{NL} and other cosmological parameters. We therefore do not need to evaluate the inner product matrix γ , a rather technical step of the numerical pipeline. The only part in this work, for which the calculation of γ is needed is when we reconstruct the bispectrum from the modal basis, presented in Figures 2 and 14, obtained by substituting the measured β_n into Equation (6). We describe our method to compute γ in detail in Appendix A.2.

The basis of one-dimensional functions $q_p(k)$ we use here is mainly composed of well-behaved polynomials normalized within the *tetrapyd* that fits inside a unit cube (see Fergusson et al. 2012a for details). This means that in the above expressions $q_p(k) \rightarrow q_p((k - k_{\min}) / (k_{\max} - k_{\min}))$. We also include custom modes based on the separable matter bispectrum shapes, i.e., the gravity-induced tree-level bispectrum and the local PNG template Equation (2) to further improve the bispectrum decomposition in Equation (6). This was introduced in Hung et al. (2019a) and further developed in Byun et al. (2021); see Appendix A.1 for details. A general useful feature of the modal estimator is that, for a proper choice of basis, it allows for a preliminary data compression step, by reconstructing the signal in a small number of modes, compared to the starting amount of triplets in Fourier space. This speeds up the numerical evaluation of covariances, which is in turn a key ingredient for the final compression step described in the next section.

¹⁵ <http://www.fftw.org>

Table 1
Cosmological Parameters and PNG Amplitudes of the QUIJOTE and QUIJOTE-PNG Simulations

	σ_8	Ω_m	Ω_b	n_s	h	$f_{\text{NL}}^{\text{local}}$	$f_{\text{NL}}^{\text{equil}}$	$f_{\text{NL}}^{\text{ortho}}$
Fiducial	0.834	0.3175	0.049	0.9624	0.6711	0	0	0
Steps	± 0.015	± 0.01	± 0.002	± 0.02	± 0.02	± 100	± 100	± 100

3.2. Data Compression and Quasi-maximum Likelihood Estimator

To build an optimal estimator of PNG amplitude parameter f_{NL} 's, we follow the method developed in Alsing & Wandelt (2018), based on a two-step data compression scheme. The first step consists of extracting a set of summary statistics from a full data set, which in our case are the power spectrum $P(k)$ and bispectrum modes β_n , described in the previous paragraphs. The second step is an additional compression, down to only n numbers where n is the number of parameters we aim to infer from data. The optimal compressed quantities, in the sense that no information about the wanted parameters is lost, are the components of the score function (the parameter gradient of the log-likelihood), which can be computed explicitly by assuming an approximate form for the likelihood of the summary statistics. When working with the power spectrum and bispectrum, a Gaussian likelihood is a reasonable approximate choice, by virtue of the central limit theorem.¹⁶ The compressed statistics, called \mathbf{t} , then take the following form:

$$\mathbf{t} = \nabla_{\theta} \mu_*^T \mathbf{C}_*^{-1} (\mathbf{d} - \mu_*), \quad (11)$$

where \mathbf{d} corresponds to the summary statistics (a vector of modal coefficients β_n and possibly power spectrum bins $P(k)$), μ and \mathbf{C} are, respectively, the mean and the covariance of \mathbf{d} , θ denotes the parameters of interest and the subscript $*$ indicates that the quantity is evaluated at the chosen fiducial cosmology. Note that in Equation (11), we made the further assumption that the covariance \mathbf{C} does not depend on the parameters θ , and in that case, the compressed statistics \mathbf{t} is equivalent to the MOPED data compression of Heavens et al. (2000). By construction, the covariance of the compressed statistic \mathbf{t} at the fiducial point θ_* is equal to the Fisher matrix \mathbf{F}_* , where

$$\mathbf{F} = \nabla_{\theta} \mu^T \mathbf{C}^{-1} \nabla_{\theta}^T \mu. \quad (12)$$

One can also write the following quasi-maximum likelihood estimator for the parameters θ :

$$\hat{\theta} = \theta_* + \mathbf{F}_*^{-1} \mathbf{t}, \quad (13)$$

which has an expected covariance given by the Fisher matrix evaluated at the fiducial point \mathbf{F}_* .

In this work, we compute both the derivatives and the covariance matrix from N -body simulations, which makes it possible to include nonlinear scales in the analysis and estimate their constraining power on PNG.

¹⁶ The validity of this assumption is tested a posteriori with the different analyses presented in Section 4, where we verify that the estimator developed yields unbiased and quasi-optimal results.

4. Analysis

4.1. Simulations

The analysis presented in this work is mainly based on the publicly available QUIJOTE suite of N -body simulations;¹⁷ see Villaescusa-Navarro et al. (2020) for details. These simulations are cubic boxes of length $1 h^{-1}$ Gpc, containing 512^3 particles, generated with the GADGET-III code (Springel 2005), with input transfer functions computed by CAMB (Lewis et al. 2000). Initial conditions are generated at $z_i = 127$ using the code 2LPTIC (Crocce et al. 2006). We use a set of 8000 simulations at fiducial cosmology to evaluate the covariance matrix, together with smaller sets of 500 realizations, in which parameters are one by one slightly displaced from their fiducial values, to numerically compute the derivatives in Equation (12). Fiducial parameter values and steps used for numerical differentiation are reported in Table 1. In Appendix B, we illustrate the effects of small parameter variations on the power spectrum and bispectrum. Throughout the analysis, we work at redshift $z = 1$.

To study the effects of PNG, we use the QUIJOTE-PNG set presented in detail in the companion paper (Coulton et al. 2022). Non-Gaussian initial conditions are generated using the method developed in Scoccimarro et al. (2012), and evolved following the exact same procedure as the standard QUIJOTE N -body simulations described above. For each of the three standard primordial shapes—local (Equation (2)), equilateral (Equation (3)), and orthogonal (Equation (4))—we analyze two sets of 500 realizations, with either $f_{\text{NL}} = +100$ or $f_{\text{NL}} = -100$.

For further testing, we also consider three additional sets of 10 simulations, with $f_{\text{NL}}^{\text{local}} = -40, 0, \text{ and } 40$, all generated from the same Gaussian seeds, independently of the QUIJOTE fiducial realizations, but with the same numerical specifications and input cosmological parameters. For this suite of simulations, initial conditions have been generated with a modified version of the PNGRun code (Wagner et al. 2010) that imprints local NG contributions on top of the primordial Gaussian potential field, starting from a prescribed shape of the primordial matter bispectrum (see Wagner et al. 2010 for a detailed description of the PNGRun algorithm). The simulations have been evolved with GADGET-III by including the contribution of relativistic species in the background cosmic expansion, which results in a different normalization at the starting redshift of the runs, compared to QUIJOTE to match the amplitude of the same perturbations at $z = 0$.

4.2. Results

The first step of our analysis consists of the extraction of power spectrum and bispectrum modes from all available simulations, using the methodology outlined in Section 3.1.

Modal bispectrum estimation requires us to fix k_{max} beforehand. In this work, we consider three different values: $k_{\text{max}} = 0.07 h \text{ Mpc}^{-1}$ (linear), $0.2 h \text{ Mpc}^{-1}$ (mildly nonlinear), and $0.5 h \text{ Mpc}^{-1}$ (nonlinear), in order to study the amount of

¹⁷ <https://quijote-simulations.readthedocs.io>




Figure 1. Impact of PNG on the matter power spectrum at $z = 1$ (ratio NG/G, averaged from 500 N -body simulations). We show the local (green), equilateral (orange), and orthogonal (purple) shapes. Solid and dashed lines correspond, respectively, to positive and negative f_{NL} values (± 100).

information as a function of scales. We begin by depositing particle positions into a regular grid using the public `Pylians3`¹⁸ code, with a fourth-order interpolation scheme and a resolution $N_{\text{grid}} = 360$ in the nonlinear case, as in Hahn et al. (2020), and the faster cloud-in-cell interpolation at the lower resolutions $N_{\text{grid}} = 128$ and $N_{\text{grid}} = 64$ ¹⁹ for the $k_{\text{max}} = 0.2$ and $0.07 h \text{ Mpc}^{-1}$ cases, respectively.

As an initial visual check, in Figure 1 we show the impact on the power spectrum of including PNG in the simulations at redshift $z = 1$. As expected, such impact is essentially negligible on linear scales and rises to the percent level, at most, in the nonlinear regime, despite the large NG amplitudes in input ($f_{\text{NL}} = \pm 100$). Moreover, the effect of local and equilateral NG is very degenerate. In Figure 2, we show the same comparison for bispectrum configurations, reconstructed from our modal basis, by using Equation (6), where the details in calculating the γ matrix, needed for this step, can be found in Appendix A.2. While the effect remains at the percent level for a single triangle, we have now of course a much larger number of available configurations. Moreover, we can clearly see how the effect of the local primordial template peaks in the squeezed limit—where one k is much smaller than the other two—whereas the equilateral shape produces the largest signal when three k are of the same order. The orthogonal bispectrum has both squeezed and equilateral components.

In Figures 3 and 4, we show some modal bispectrum coefficients β_n of the QUIJOTE simulations, evaluated using Equation (9). Figure 3 is the direct equivalent of Figure 2 at the modal level, and also highlights the very distinct behaviors of the three primordial shapes. In Figure 4, we compare the modal coefficients of the fiducial Gaussian simulations to their

theoretical prediction, computed at the tree level. The first five modes correspond to the special functions based on the tree level and local separable templates (see Appendix A.2 for more details), while the rest are polynomials of increasing degree, from left to right. As expected, the difference between measured and predicted signals increases significantly with k_{max} . Note also that in this analysis we work with the Q_n separable modes defined in Equation (6). These, by definition, are not orthogonal with respect to the inner product defined in Equation (7) and always display nontrivial correlations. For this reason, in modal bispectrum analyses a new basis is often introduced, by orthogonalizing the Q_n templates. This procedure, however, only removes the correlations in the weakly non-Gaussian case, and after verification, does not bring any improvement to the results presented in this paper. On the contrary, it can add some numerical instability when including high-order polynomial modes; thus, it is not included here.

After extracting summary statistics from the data and checking that they behave according to expectation, the next step of the analysis is the optimal compression to the score function, following the procedure described in Section 3.2. This requires the evaluation of the mean and the inverse covariance matrix of the estimated power spectrum bins and bispectrum modal coefficients, as well as their derivatives with respect to each considered parameter (cosmological + NG amplitudes). We compute the data covariance \hat{C} from 8000 QUIJOTE N -body simulations at fiducial cosmology and we apply the Hartlap–Anderson correction factor (Hartlap et al. 2007) to obtain an unbiased estimate of the precision matrix:

$$C^{-1} = \frac{n_r - n_d - 2}{n_r - 1} \hat{C}^{-1}, \quad (14)$$

where n_r is the number of realizations (8000 here) and n_d represents the size of the data vector (a few hundred at most here).²⁰ In Figure 5, we show the full correlation matrix given by $C_{ij} / \sqrt{C_{ii} C_{jj}}$, including both power spectrum bins $P(k)$ and mode amplitudes β_n . As expected, the power spectrum becomes more correlated on smaller scales. Bispectrum modal coefficients however have the opposite behavior, becoming less correlated at higher k_{max} . This is due to the fact that each mode is, by definition, the fit of a given template (polynomial, or simple separable bispectrum shape) to the data, including all scales in the range of $[k_{\text{min}}, k_{\text{max}}]$. Thus, an increased k_{max} helps to better distinguish between them, leading to lower correlations. Note also the nonzero correlation between the power spectrum and the bispectrum, which will play an important role in the analyses presented in this paper. We do not include super-sample covariance terms in our analyses, which have a negligible impact on our results as shown in our companion paper (Coulton et al. 2022). We also use these 8000 simulations to compute the mean value of the modal coefficient vector, which fixes the score function to zero at fiducial cosmology with no PNG.

The derivatives can be calculated quickly from the sets of 500 simulations with one adjusted parameter and matching

¹⁸ <https://github.com/franciscovillaescusa/Pylians3>

¹⁹ It has been argued in Sefusatti et al. (2016) that the FFT-based standard bispectrum estimator can probe scales up to $k_{\text{max}} = 2k_{\text{Ny}}/3$, where the Nyquist frequency is $k_{\text{Ny}} = k_f N_{\text{grid}}/2$ and $k_f = 2\pi/L$ is the fundamental frequency of the box. This is due to the factor $e^{i(k_1+k_2+k_3)x}$, present in the standard estimator, which is invariant under a translation $k_i \rightarrow k_i + k_f N_{\text{grid}}/3$. This indicates that any estimator that has this exponent factor will have the above momentum cutoff scale, rather than $k_{\text{max}} = k_{\text{Ny}}$. In Byun et al. (2021), they test this and find indications that the modal estimator falls under the same category, since it takes a form similar to the standard bispectrum estimator (see Equation (A6)). Taking this into account, gives the largest values for the wavenumbers k to be $k_{\text{max}} \simeq 0.75, 0.26$ and $0.13 h \text{ Mpc}^{-1}$ for $N_{\text{grid}} = 360, 128$, and 64 , respectively, well above our chosen k_{max} values.

²⁰ This correction factor is never more than a few percent here, a regime where it has been verified (see, e.g., Gualdi et al. 2021a; Gil-Marín 2022) that it provides results similar to the more advanced method developed in Sellentin & Heavens (2016) when estimating Fisher bounds.




Figure 2. Impact of PNG on the matter bispectrum at $z = 1$ (ratio NG/G, averaged from 500 N -body simulations). We show the local (top rows), equilateral (middle rows), and orthogonal (bottom rows) shapes.

random seeds,²¹ using the standard formula:

$$\frac{\partial d_n}{\partial \theta} \simeq \frac{d_n(\theta_{\text{fid}} + \Delta\theta) - d_n(\theta_{\text{fid}} - \Delta\theta)}{2\Delta\theta}, \quad (15)$$

²¹ In addition to using matching random seeds, different step values between simulations have been tested in order to reduce the numerical noise when evaluating the derivatives.

where θ can be any of the cosmological or NG amplitude parameters considered in this paper and d is either power spectrum or bispectrum mode estimates.

4.3. Fisher Constraints

With the numerically evaluated, joint covariance matrix of power spectrum bins and bispectrum modes, described in the

Table 2
Joint Constraints on Cosmological Parameters and PNG from the Power Spectrum and the Modal Bispectrum at $z = 1$, for Different k_{\max}

Fiducial	k_{\max} ($h \text{ Mpc}^{-1}$)	σ_8 0.834	Ω_m 0.3175	Ω_b 0.049	n_s 0.9624	h 0.6711	$f_{\text{NL}}^{\text{local}}$ 0	$f_{\text{NL}}^{\text{equil}}$ 0	$f_{\text{NL}}^{\text{ortho}}$ 0
$P(k)$	0.07	± 0.17	± 0.32	± 0.32	± 3.9	± 4.4			
	0.2	± 0.012	± 0.039	± 0.018	± 0.24	± 0.24			
	0.5	± 0.0045	± 0.011	± 0.0062	± 0.042	± 0.062			
β_n	0.07	± 0.59	± 0.95	$\pm 1.$	± 12	± 14			
	0.2	± 0.014	± 0.051	± 0.023	± 0.32	± 0.31			
	0.5	± 0.0063	± 0.016	± 0.006	± 0.066	± 0.071			
$P(k) + \beta_n$	0.07	± 0.17	± 0.31	± 0.31	± 3.7	± 4.2			
	0.2	± 0.011	± 0.035	± 0.015	± 0.21	± 0.21			
	0.5	± 0.0038	± 0.0091	± 0.0048	± 0.031	± 0.048			
$P(k)$	0.07	± 0.55	± 0.41	± 0.36	± 4.5	± 4.9	$\pm 200,000$	$\pm 500,000$	$\pm 180,000$
	0.2	± 0.13	± 0.07	± 0.031	± 0.49	± 0.39	$\pm 31,000$	$\pm 85,000$	$\pm 36,000$
	0.5	± 0.069	± 0.035	± 0.013	± 0.24	± 0.18	± 9600	$\pm 29,000$	$\pm 11,000$
β_n	0.07	± 1.1	± 1.3	± 1.4	± 18	± 20	± 670	± 2300	± 1500
	0.2	± 0.016	± 0.059	± 0.027	± 0.37	± 0.36	± 91	± 390	± 300
	0.5	± 0.0068	± 0.018	± 0.0064	± 0.078	± 0.079	± 39	± 150	± 110
$P(k) + \beta_n$	0.07	± 0.17	± 0.31	± 0.31	± 3.7	± 4.2	± 350	± 930	± 610
	0.2	± 0.011	± 0.035	± 0.015	± 0.21	± 0.21	± 52	± 170	± 120
	0.5	± 0.0038	± 0.0093	± 0.0048	± 0.033	± 0.048	± 22	± 94	± 58

Note. We analyzed 8000 QUIJOTE N -body simulations of $1 \text{ (Gpc}/h)^3$ volume at fiducial cosmology, and sets of 500 N -body simulations with one adjusted parameter.

previous section, we can derive expected error bars for the quasi-maximum likelihood estimator (Equation (13)), using the Fisher matrix defined in Equation (12).

In Table 2, we present the resulting constraints on cosmological parameters σ_8 , Ω_m , Ω_b , n_s , and h , and PNG amplitude parameters $f_{\text{NL}}^{\text{local}}$, $f_{\text{NL}}^{\text{equil}}$, and $f_{\text{NL}}^{\text{ortho}}$, for the three different k_{\max} studied. We consider the power spectrum and the bispectrum (modal coefficients) separately and then analyze them jointly, in order to gain insight into their relative constraining power at different scales.

In the upper part of the table, we assume Gaussian initial conditions. As pointed out in Hahn et al. (2020), we see in this case that adding bispectrum information does help improve power spectrum constraints on cosmological parameters. While the gain is marginal on linear scales, it becomes significant in the nonlinear regime, where the bispectrum alone is actually a more powerful probe than the power spectrum alone, for all considered parameters, with the exception of σ_8 .

In the lower part of the table, we include both cosmological parameters and PNG of the three standard shapes. In this case, if we look at bounds from the power spectrum alone, we not only see that PNG amplitudes are essentially unconstrained, as expected, but also that errors on other parameters significantly degrade in comparison with the primordial Gaussian case. This is due to degeneracies arising between f_{NL} and cosmological parameters, which are even more pronounced at nonlinear scales. The most affected parameter is σ_8 , the error of which increases by more than a factor 3. The bispectrum now plays a crucial role not only in constraining PNG at a significant level, but also in breaking such degeneracies and allowing for large improvements in predicted errors for cosmological parameters.

Note also that in Table 2, we assume that all PNG shapes are simultaneously present in the data. However, we verified that similar degeneracies are introduced even when we consider one NG shape at a time. This is not surprising, remembering that introducing NG of different types produces a similar impact on the power spectrum (see Figure 1).

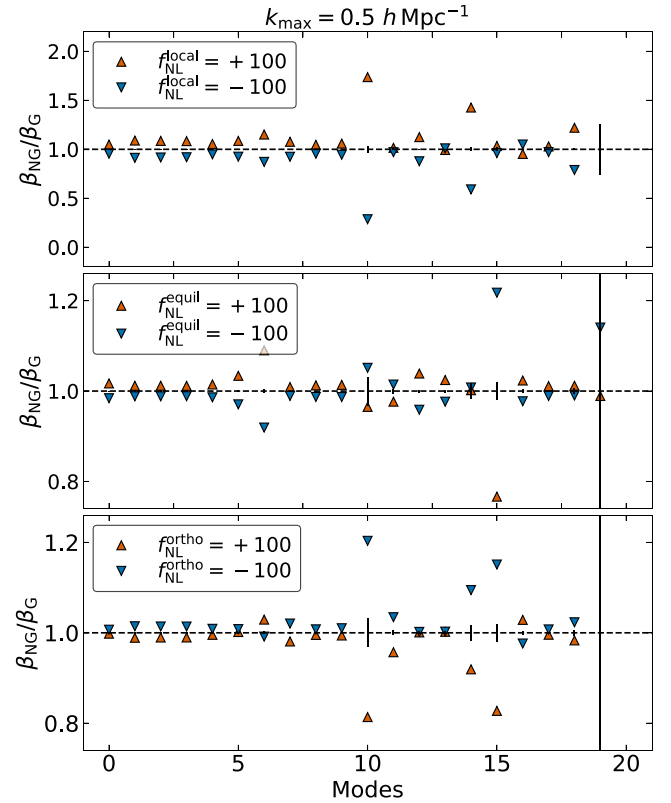


Figure 3. Impact of PNG on the modal bispectrum coefficients (ratio NG/G, averaged from 500 N -body simulations) at $z = 1$. The vertical black lines show the standard errors for 500 Gaussian fiducial simulations and for most modes these are unobservably small.

When combining the power spectrum and bispectrum information, we recover very similar error bars on cosmological parameters as in the purely Gaussian case, meaning that degeneracies are




Figure 4. The difference between the modal bispectrum coefficients of the QUIJOTE N -body simulations (average from 8000 simulations at fiducial cosmology) and their theoretical prediction computed at the tree level, divided by their respective standard deviation. Note the different vertical scales of each plot.

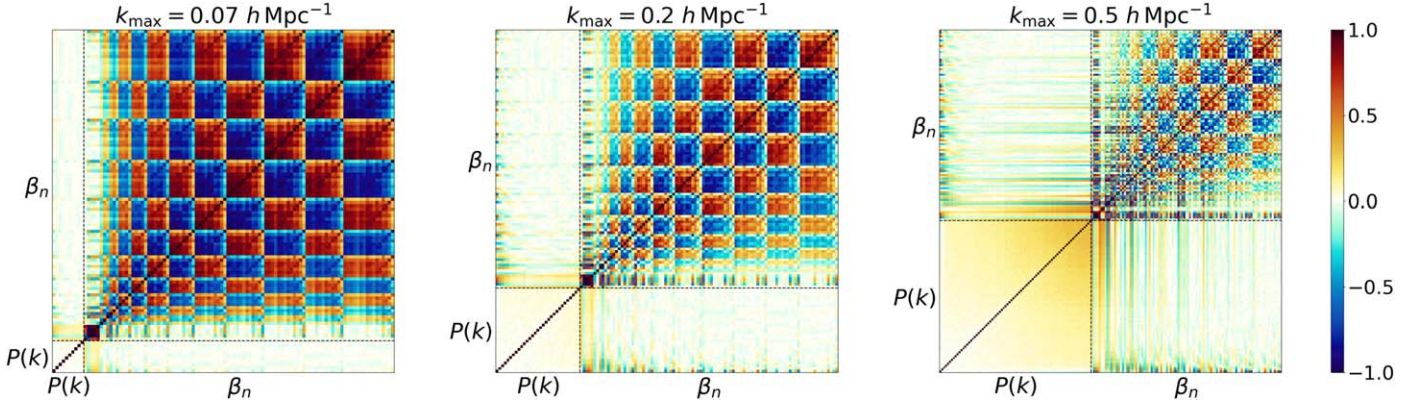


Figure 5. The correlation matrix of the power spectrum $P(k)$ and modal bispectrum coefficients β_n of the QUIJOTE simulations for different k_{\max} . We include a similar number of bispectrum modes in each case, using polynomial basis functions up to a degree of 12.

Table 3

Fisher 1σ Constraints on the Three PNG Shapes from the Power Spectrum and the Modal Bispectrum at $k_{\max} = 0.5 h \text{ Mpc}^{-1}$, for a Cubic Volume of $1 (\text{Gpc}/h)^3$ and at $z = 1$, after Marginalization Over Cosmological Parameters

$f_{\text{NL}}^{\text{local}}$	$f_{\text{NL}}^{\text{equil}}$	$f_{\text{NL}}^{\text{ortho}}$
± 16	± 77	± 40

Note. Each shape is analyzed independently.

strongly broken. For the same reason, we also see that PNG bounds from the bispectrum alone are improved by a factor ~ 2 when we add power spectrum measurements, despite the fact that the power spectrum yields almost no PNG constraining power by itself. Interestingly, even if we completely fix cosmological parameters in the analysis, combining power spectrum and bispectrum still produces some improvement in PNG constraints. In this case, degeneracies clearly cannot play a role. It turns out that the improvements are now due to correlations between the power spectrum and the bispectrum, displayed in Figure 5. Intuitively, at the tree level, the gravitational bispectrum is linked to the power spectrum squared, via a convolution kernel; hence, measuring the latter produces an information gain also on the former (in each simulation, fluctuations of the power spectrum and bispectrum around their respective averages, which are fixed by fixing cosmological parameters, are correlated). A better knowledge of the gravitational bispectrum allows in turn for a better separation of this component from the primordial signal.²²

²² An additional role is played by PNG contributions to the power spectrum, at the loop level, if we perform a single-shape analysis. Such contributions are however degenerate among different shapes.

In Table 3, we report the independent Fisher bounds for the three PNG shapes, after marginalizing over cosmological parameters.

Another important result is that probing nonlinear scales significantly improves the constraints on every considered parameter. This is further illustrated in Figure 6, where we show the contours obtained using jointly the power spectrum and the bispectrum for different k_{\max} . A more detailed analysis of the information content of different scales is in our companion paper (Coulton et al. 2022), where we show that errors tend to saturate at $k_{\max} = 0.3 h \text{ Mpc}^{-1}$, at $z = 0$, due to correlations between Fourier modes at strongly nonlinear scales.

The accuracy of the modal expansion of the bispectrum has a strong impact on expected error bars and thus it is important to verify that the modal results presented in this section are fully converged. There are several solutions to improve the modal reconstruction, in order to achieve higher accuracy with as small as possible number of modes. Some of them involve using custom separable templates, corresponding to, or highly correlated with, expected signals in the data. In this work, we already pursue this approach by including the tree-level matter bispectrum and the primordial local shape in the modal basis. In the future, we plan to explore it further, by considering, for example, separable templates that are strongly correlated to the predicted bispectrum at one loop in perturbation theory, or describing the other primordial shapes.

The obvious, although less economical, alternative approach to improve reconstruction accuracy simply consists of including additional, higher-order polynomials in the basis. We consider here polynomials up to degree 20, which are combined in triplets to obtain several hundred bispectrum modes. In Figure 7, we

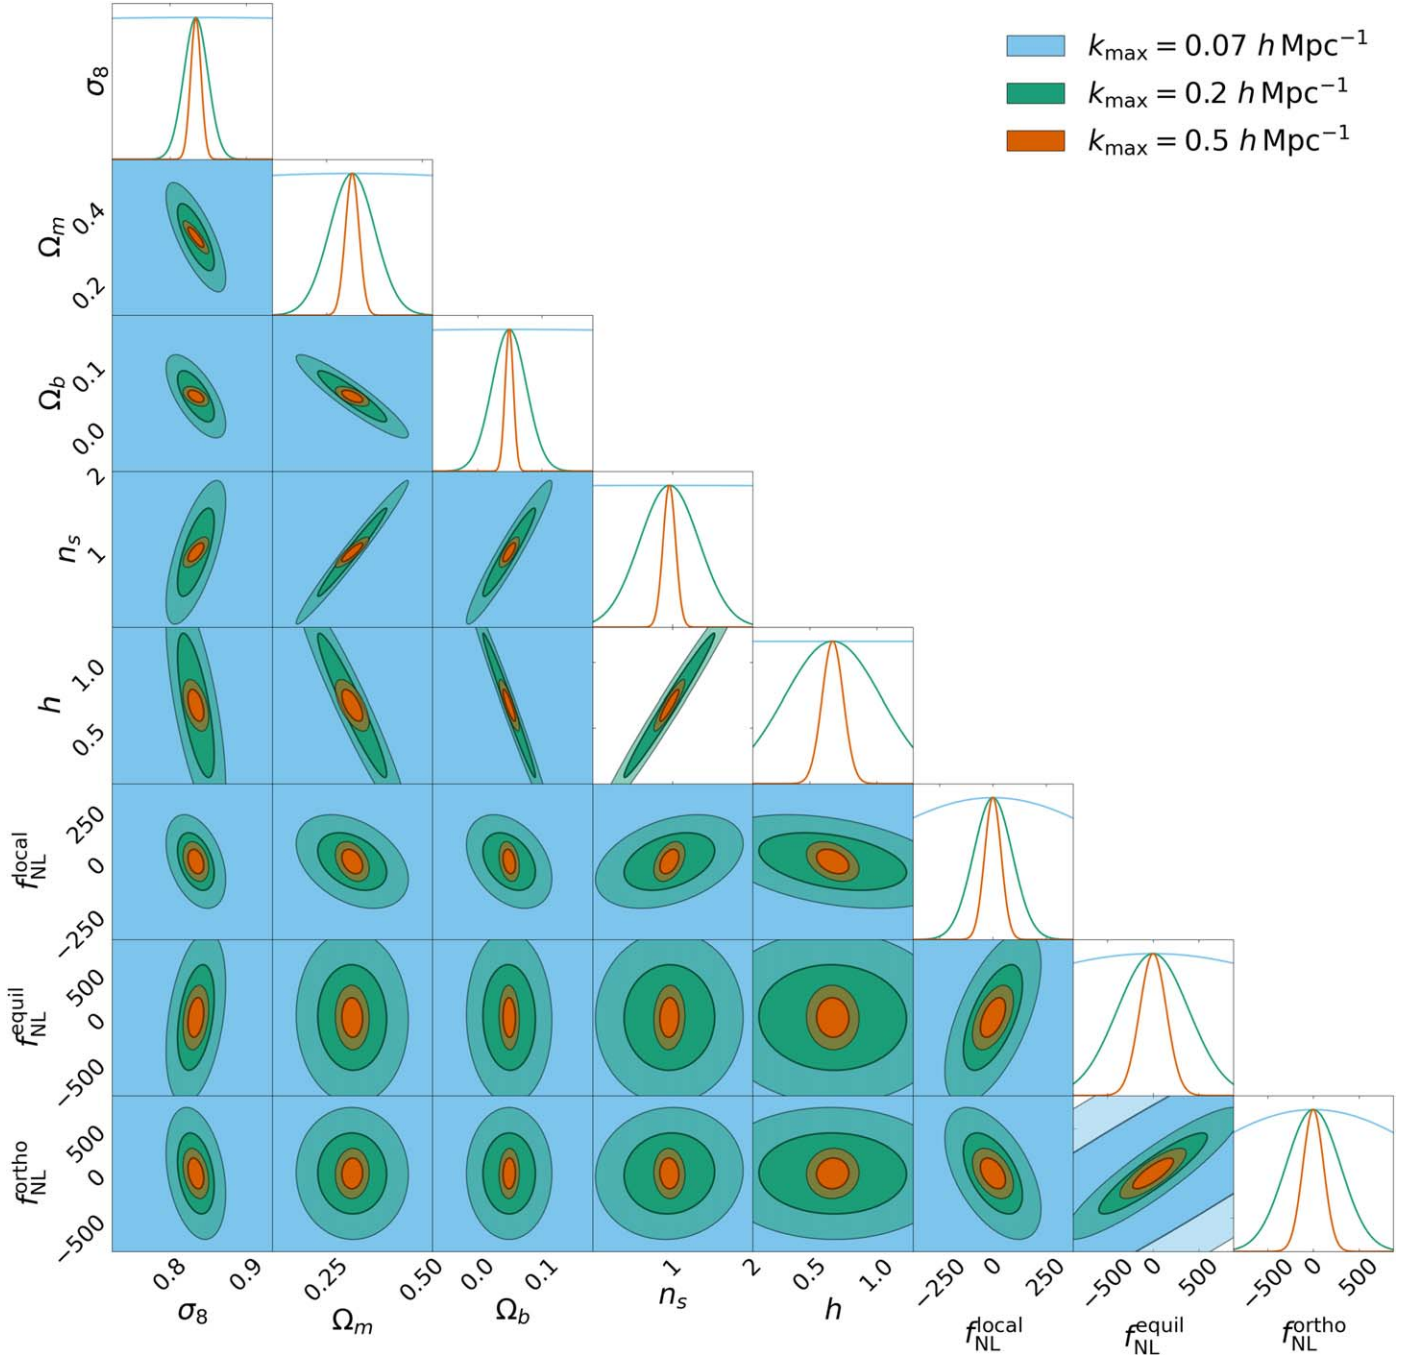


Figure 6. Joint constraints on cosmological parameters and PNG from the power spectrum and the modal bispectrum for different k_{\max} at $z = 1$.

study the convergence of the constraints as a function of the number of modes included in the modal expansion, each data point from left to right corresponding to adding polynomials of 1 degree higher. As could be expected, the accurate description of the bispectrum on small scales requires more modes. The constraints on some cosmological parameters are not even fully converged by taking polynomials up to degree 20, making our Fisher forecast in such case slightly overly conservative; we verify that the results on PNG are however always converged. The situation is significantly better with the joint power spectrum and bispectrum analysis. In the mildly nonlinear regime, we achieve convergence with ~ 100 modes, and using only 10 modes produces constraints that are already within 10% of their converged values. Even in the nonlinear regime, going beyond

~ 100 modes we improve the Fisher bounds only by $\sim 1\%$ (except for Ω_b and h for which it is of a few percent). This means that, for the combined power spectrum and bispectrum analysis, we can stop the polynomial expansion to a lower degree than 20, with a negligible loss of information. To be exact, the constraints given in Table 2 use polynomials up to the degree 12 for the joint power spectrum + bispectrum case, and 20 for the bispectrum-only analysis. Note that using higher-order polynomials can add numerical noise to the analysis, thus stopping the expansion at a reasonable degree, like we do here, is a more robust approach.

The numbers discussed above show that the modal representation is quite efficient. Compare, for example, the ~ 100 modal coefficients in the current analysis to the thousands of k triplets that are typically necessary to study the bispectrum at the

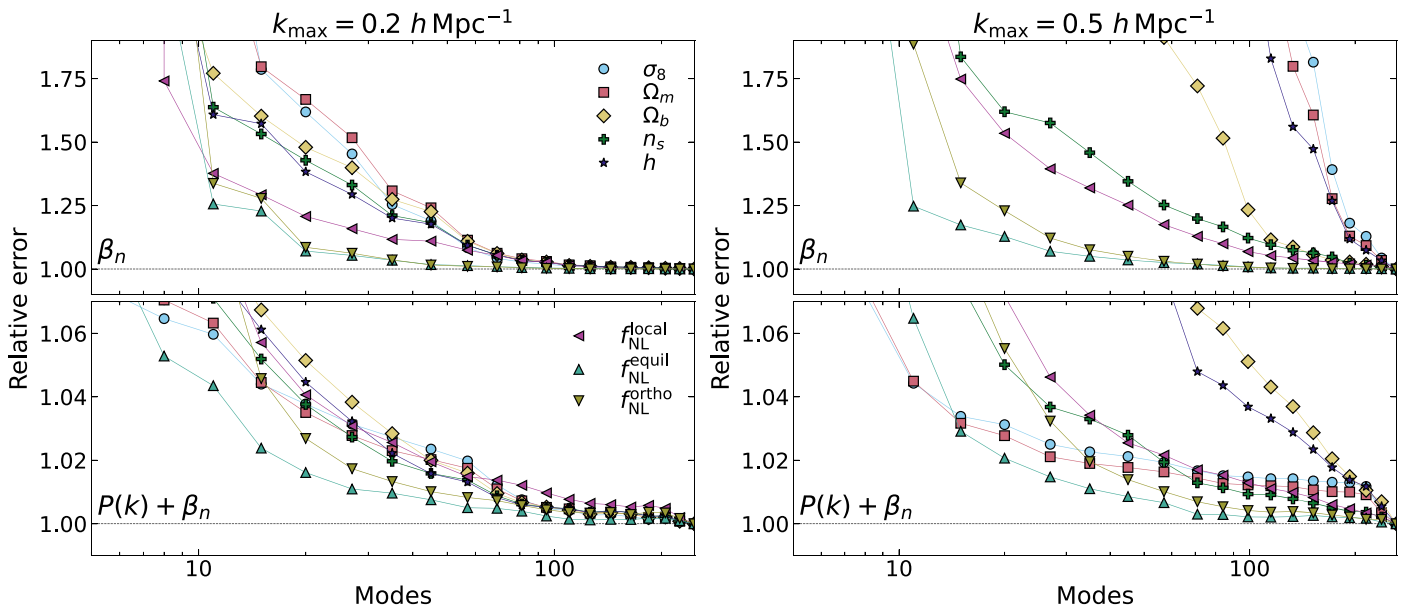


Figure 7. The convergence of constraints on cosmological parameters and PNG as a function of the number of modes used for the modal bispectrum estimation. The top row corresponds to the bispectrum-only case, while the bottom row also includes the power spectrum information. Note the difference in scale between the top and bottom rows. For each parameter, error bars are divided by the corresponding minimum; see Table 2 for the actual values. From left to right, each data point corresponds to include polynomials of 1 degree higher in the modal basis.

same nonlinear scales, using the standard, *binned* estimator (e.g., ~ 2000 triangle configurations with rather large bins of width $3k_F$). This more efficient compression is, for example, important to increase the accuracy of the numerical derivatives evaluated from simulations.

Convergence tests have to be performed not only for the modal basis, but also with respect to the total number of mock realizations used to evaluate the Fisher matrix. In Figure 8, we vary the number of simulations used to estimate derivatives and verify the impact of this change on the final forecast. We find that the constraints from the bispectrum and the joint power spectrum + bispectrum analyses remain very stable, even by using 25 simulations instead of the usual 500. At the same time, though, we see that the power spectrum-only forecast is not fully converged, even when we use the full 500 simulation sample, leading to somewhat overoptimistic constraints in Table 2, when cosmological parameters and PNG are studied jointly. As the power spectrum-only case has already been reported as essentially unconstraining for PNG, this does not change any of the main conclusions presented above (see the companion paper Coulton et al. 2022 for a more detailed study of this issue, which can for example be mitigated by assuming some prior on f_{NL}).

In Figure 8, we also verify that the constraints are highly stable when we change the number of simulations (at fiducial cosmology) used to estimate covariance matrices. Using only 1000 instead of 8000 simulations only leads to variations at the percent level. However, the requirement on covariance precision will become stricter in the next section, where we discuss the actual parameter inference.

4.4. Parameter Estimation

After deriving Fisher matrix forecasts, we now build the quasi-maximum likelihood estimator defined in Equation (13) and apply it to different data sets, to verify its efficiency to extract accurate cosmological information.

Our first validation test consists of estimating jointly cosmological parameters and PNG amplitudes (σ_8 , Ω_m , Ω_b , n_s , h , $f_{\text{NL}}^{\text{local}}$, $f_{\text{NL}}^{\text{equil}}$, $f_{\text{NL}}^{\text{ortho}}$) in the QUIJOTE simulations, for different cosmologies. We include nonlinear scales ($k_{\text{max}} = 0.5 h \text{ Mpc}^{-1}$) and use both the power spectrum and the bispectrum in this analysis, as their combination should lead to smaller error bars, as verified in the previous section.

In Figure 9, we show the corresponding results for fiducial cosmology data sets, as well as others having PNG ($f_{\text{NL}}^{\text{local}} = 100$, $f_{\text{NL}}^{\text{equil}} = 100$ or $f_{\text{NL}}^{\text{ortho}} = 100$).²³ We compute averages of the estimated parameters and compare them to their input values in the simulations. The correct values are recovered each time, confirming the unbiasedness of the estimator. Note also that we consider only cases in which these averages are extracted from simulations not overlapping with those used to evaluate derivatives and covariances. In this way, we avoid spurious correlations, which could arise from using overlapping sets of data both to calibrate the estimator and to measure parameters.

We also perform a similar analysis for the additional sets of 10 realizations, produced independently of the QUIJOTE simulations and having respectively $f_{\text{NL}}^{\text{local}} = -40$, 0, and 40 described in Section 4.1. The quasi-maximum likelihood estimator is built from all the available QUIJOTE simulations. The results are shown in Figure 10, where we focus on the combined power spectrum and modal bispectrum analysis at different k_{max} . All estimated central values are within or very close to the expected 1σ range. Moreover, the internal scatter between cosmological parameter estimates for the three data sets ($f_{\text{NL}}^{\text{local}} = -40$, 0, and 40) is extremely small.

²³ We focus here on changes in f_{NL} , rather than on other cosmological parameters, both because the main focus of this work is on the study of PNG and because, in the available QUIJOTE set, we can access realizations with $f_{\text{NL}}^{\text{local}} = 100$, which is more than a 4σ deviation from the fiducial value of $f_{\text{NL}} = 0$. This allows us to test an interesting regime, in which the data we analyze are generated from parameters that are significantly displaced from the model we use to build the estimator.




Figure 8. The stability of Fisher constraints under variations of the number of simulations. In the left panels, we vary the number of simulations used to calculate numerical derivatives, while on the right panels we vary the number of simulations used to measure covariances. All error bars are normalized by the result obtained using the full data set, see Table 7 for actual values. The top, middle, and bottom rows correspond respectively to power spectrum only, bispectrum only, and combined power spectrum and bispectrum analyses. Each column corresponds to the constraints on one parameter, and they are all analyzed jointly. We consider five different numbers of simulations to compute derivatives (from 25–500) and covariances (from 500–8000), each number corresponding to its own color and marker. We include scales down to $k_{\max} = 0.5 h \text{ Mpc}^{-1}$.

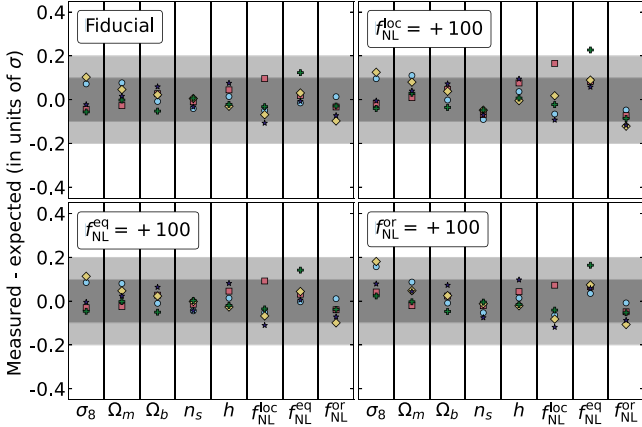


Figure 9. The unbiasedness of the quasi-maximum likelihood estimator (Equation (13)), when measuring cosmological parameters and PNG amplitudes. We use the power spectrum and the bispectrum jointly, up to $k_{\max} = 0.5 h \text{ Mpc}^{-1}$. Each individual column corresponds to a given parameter (cosmological or PNG). To be exact, we show the difference between the expected and measured values, divided by the respective Fisher error bar. Each panel corresponds to a different cosmology of the data samples (i.e., one with Gaussian initial conditions and the three types of PNG). For each of them, we analyze five independent data sets of 100 realizations, each being indicated by its own color and marker.

In general, we expect the estimator to yield slightly different results, depending on the exact set of simulations used for its construction. For this reason, we perform several stability tests, in which we study variations in the estimated parameters, by changing the number of simulations used to calculate covariances or derivatives. These tests highlight that less accurate derivatives and covariances, which can be obtained if we use a too small number of realizations in their evaluation, have a different impact on the final measurements.

A less precise covariance matrix leads to a suboptimal estimation of parameters, as shown in Figure 11. We compare the error bars of the quasi-maximum likelihood estimator applied to the QUIJOTE simulations at fiducial cosmology, as a function of the number of realizations used to compute covariances, to the Fisher forecasts described in Section 4.3. This confirms that our estimator yields close to optimal constraints, at the percent level for most parameters (5% for Ω_b and h), when using many simulations to evaluate the covariance. Using only 1000 simulations, the increase of error bars is still only of a reasonable order 10%, which is small compared to the gain due to including nonlinear scales in the analysis (a factor 2 or more for all parameters).

For the derivatives, the impact on optimality is negligible, but the results can be biased, as shown in Figure 12. To study the

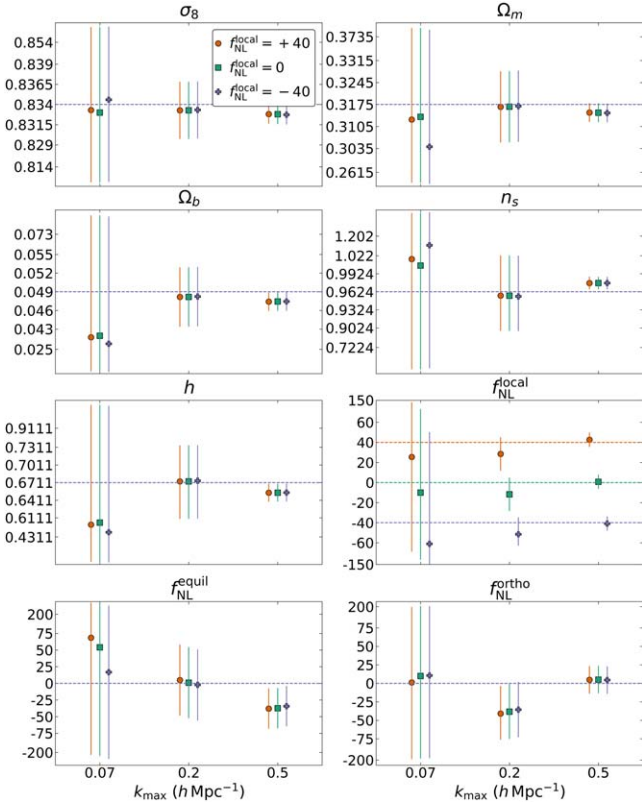


Figure 10. Cosmological and PNG amplitude parameters estimated from several sets of 10 N -body simulations for different k_{\max} with a different amount of local PNG in their initial conditions. The orange circles, green squares, and purple crosses correspond respectively to $f_{\text{NL}}^{\text{local}} = 40$, 0, and -40 , and the vertical dashed lines are the expected standard errors on the average from 10 simulations. The horizontal dashed lines are the expected values of the different data sets. The vertical scale is linear around 0, and logarithmic close to the edges.

typical size of the bias, we define less accurate estimators, using smaller subsets of the 500 simulations, which are initially used to compute derivatives. For example, using 25 simulations to compute derivatives gives 20 independent estimators. Then, we measure the difference between the measured parameters with every estimator and the benchmark results, obtained with an estimator calibrated using all available simulations. Finally, we compute the standard deviation of this difference as a function of the number of simulations used to calculate derivatives (from 1–400), which is what we show in Figure 12 after dividing it by the corresponding Fisher error bar.

We apply the same procedure for different sizes of data samples (10 and 100 simulations) to verify that the results are indeed biased, and not suboptimal, as, in the latter case, the spurious bias evaluated with this methodology would decrease with the number of simulations. As data samples, we consider three different choices. In the first, all parameters are set at their fiducial values, meaning in particular that all PNG amplitudes are set to 0. In the second, we set $f_{\text{NL}}^{\text{equil}} = 100$; we will define this case as *close to fiducial* since it corresponds to a 1σ deviation from $f_{\text{NL}}^{\text{equil}} = 0$. In the third, we take $f_{\text{NL}}^{\text{local}} = 100$; this is our *far from fiducial* case, since it is characterized by a 4σ deviation from $f_{\text{NL}}^{\text{local}} = 0$. We see that the impact of the number of simulations used to compute derivatives on the fiducial and close to fiducial cases is negligible. However, when analyzing the simulations with local PNG (far from




Figure 11. The error bars of the quasi-maximum likelihood estimator as a function of the number of realizations used to calculate the covariance matrix. We analyze the 8000 QUIJOTE simulations at $z = 1$, including both the power spectrum and the bispectrum, up to $k_{\max} = 0.5 h \text{ Mpc}^{-1}$. All error bars are divided by their corresponding Fisher prediction, as described in Section 4.3.

fiducial), the measured value of $f_{\text{NL}}^{\text{local}}$ displays a bias. However, this effect decreases rapidly and becomes already negligible when using more than 100 simulations to compute derivatives (let us recall that in our main analysis we use 500 realizations for this task). Even when using less than 10 simulations, it is still smaller than the typical 1σ error bar. The two other primordial shapes, both correlated with local PNG, are also affected in similar ways, whereas for cosmological parameters the effect is negligible.

We can thus conclude that the quasi-maximum likelihood estimator combining power spectrum and modal bispectrum is unbiased and can optimally extract information about cosmology and PNG up to significantly nonlinear scales, provided a sufficiently large number of simulations to calculate covariance and derivatives is used.

An important caveat is that to achieve optimality, it is also important to evaluate covariances and derivatives by using an input fiducial cosmology in the simulations, which is close enough to the actual value of parameters in the data. In our current test, this is true by construction. However, in general, it may be necessary to reach a good enough fiducial set of parameters by iteration. This poses the numerical challenge of generating thousands of N -body simulations at several different cosmologies. One efficient solution to overcome this issue is to use the CARPool method (Chartier & Wandelt 2021, 2022; Chartier et al. 2021), which combines a few high-fidelity simulations with many fast surrogates (typically 100–1000 times faster to produce) to obtain accurate estimates of the mean and covariance matrix of our observables. This approach is under investigation and will be further discussed in a forthcoming publication.

5. Conclusions

In this paper, we discussed the implementation of a new, simulation-based, joint power spectrum and bispectrum statistical estimator of PNG amplitudes and standard Λ CDM parameters in LSS data. The methodology follows the general approach developed in Alsing & Wandelt (2018) and Heavens et al. (2000), and it is based on extracting summary power spectrum and bispectrum statistics from the data and computing the score function, making the reasonable assumption that sampling distributions are Gaussian to a good approximation. This produces a vector of optimally compressed statistics, which is then used to build a quasi-maximum likelihood estimator of the parameters.




Figure 12. The typical size of the bias on estimated parameters, divided by their corresponding Fisher error bars, as a function of the number of simulations used to compute derivatives. We analyze QUIJOTE simulations for three different cosmologies, fiducial (left panel), $f_{\text{NL}}^{\text{local}} = 100$ (middle panel), and $f_{\text{NL}}^{\text{equil}} = 100$ (right panel), and include both the power spectrum and the bispectrum up to $k_{\text{max}} = 0.5 h \text{ Mpc}^{-1}$. The solid lines correspond to data sets of 10 simulations, while the dashed lines correspond to data sets of 100 simulations.

The required covariances and numerical derivatives with respect to parameters are evaluated from a large set of mock realizations with Gaussian and non-Gaussian initial conditions. To extract bispectrum summaries, we implement an efficient numerical pipeline for modal bispectrum estimation, allowing for a useful preliminary compression step, in which the information in all available Fourier triplets is described by a relatively small set of mode amplitudes.

Our pipeline was applied to a large set of QUIJOTE simulations at $z=1$, including newly produced sets of realizations with PNG of the local, equilateral, and orthogonal types. In the analysis presented here, we focused on the study of the dark matter field. Of course this is a vastly idealized scenario, compared to the analysis of the galaxy density field in redshift space, required in a real data set. However, a preliminary study of the matter field is a very useful starting point, as it both enables validating the entire methodology—developed here for the first time—and addresses at the same time the interesting question of how much information on PNG and cosmological parameters we can in principle extract, by pushing the analysis to small scales, where perturbative approaches fail, as we show, even in this simplified context. For this purpose, we derived Fisher bounds for the parameter set $\{f_{\text{NL}}^{\text{local}}, f_{\text{NL}}^{\text{equil}}, f_{\text{NL}}^{\text{ortho}}, \sigma_8, \Omega_m, \Omega_b, n_s, h\}$, using the covariance of the summary statistics—extracted from a training set of 8000 realizations—at three different cutoff scales, namely, $k_{\text{max}} = 0.07 h \text{ Mpc}^{-1}$, $k_{\text{max}} = 0.2 h \text{ Mpc}^{-1}$, and $k_{\text{max}} = 0.5 h \text{ Mpc}^{-1}$. We then applied the quasi-maximum likelihood estimator to our mock data and verified its unbiasedness.

The robustness of the results was thoroughly checked, for example, by reducing the number of training realizations used to calibrate the estimator, or to compute the Fisher forecasts. We confirmed that all our results are fully converged and remain stable (e.g., with results within the 1σ range, even when we used a few times less data to evaluate covariances). Moreover, we run the pipeline on a set of NG realizations, generated independently of the QUIJOTE suite and displaying a different input f_{NL} from the training set; also in this case, we recovered the unbiasedness of the estimator. In further consistency tests, we considered power spectrum and bispectrum estimates separately. Interestingly, starting from bispectrum results, we found out that PNG constraints significantly improve when we add power spectrum estimates. This effect persists when we fix all standard cosmological parameters to their known fiducial values. This might seem surprising at first, since there are no f_{NL} signatures at the tree level in the dark matter power spectrum, for any PNG shape. However, it turns out that the power spectrum works in this case as an ancillary

statistic, able to improve the separation between the primordial (signal) and gravitational (*noise*) bispectrum component, through its correlation with the latter.

In summary, the outcome of our analysis shows that our pipeline can extract information on f_{NL} and cosmological parameters up to small scales, deep into the nonlinear regime, which is a very promising starting point. The overall estimation procedure is unbiased and fast, once input simulations to train the estimator have been generated. A large set of input realizations, based on the QUIJOTE suite, was produced for this purpose and are publicly available.²⁴ These simulations are fully described in our companion paper (Coulton et al. 2022). In the same companion paper, we also study in more detail the information content of the power spectrum and bispectrum, at $z=0$, addressing the important points of quantifying how large the information contribution coming from nonlinear scales is and of establishing when such contribution saturates. This is done via a Fisher matrix analysis using bispectrum Fourier modes, which also includes a further battery of numerical convergence and validation tests.

In the work presented here, the training and validation sets are characterized in most cases by matching input cosmological parameters. In general, this would not be the case and too large a mismatch between the input fiducial cosmology and the actual value of parameters in the data would make the results suboptimal or biased. This can be solved by implementing a recursive procedure, in which the fiducial cosmology is updated by taking the best-fit values of the parameters at each step. Such an approach has the extra cost of having to generate new sets of simulations to recalibrate the estimator weights at each step. This issue can in principle be significantly alleviated by resorting to the CARPool method, which allows for accurate estimates of covariances by resorting to just a small set of high-fidelity realizations, combined with a large set of fast surrogates. This approach is under investigation and will be the object of a separate publication.

In forthcoming works, we will also gradually extend our analysis to include more and more realistic scenarios: first, we will analyze the actual density field of biased tracers (dark matter halos and galaxies), and later, we will account for redshift space, incomplete sky coverage, and other crucial observational effects, to finally reach the capability to analyze real data sets. Note that, given our simulation-based approach, this will not require any significant modification in the pipeline that was illustrated and validated here, but only changes in the input training sets and data.

²⁴ <https://quijote-simulations.readthedocs.io/en/latest/png.html>

G.J., M.L., and M.B. were supported by the project ‘‘Combining Cosmic Microwave Background and Large Scale Structure data: an Integrated Approach for Addressing Fundamental Questions in Cosmology,’’ funded by the MIUR Progetti di Ricerca di Rilevante Interesse Nazionale (PRIN) Bando 2017—grant 2017YJYZAH.

D.K. is supported by the South African Radio Astronomy Observatory (SARAO) and the National Research Foundation (grant No. 75415).

G.J. and M.L. also acknowledge support from the INDARK INFN Initiative (<https://web.infn.it/CSN4/IS/Linea5/InDark>), which provided access to CINECA supercomputing facilities (<https://www.cineca.it>).

M.B. acknowledges the use of computational resources from the parallel computing cluster of the Open Physics Hub (<https://site.unibo.it/openphysicshub/en>) at the Physics and Astronomy Department in Bologna.

L.V. acknowledges ERC (BePreSySe, grant agreement 725327), PGC2018-098866-B-I00 MCIN/AEI/10.13039/501100011033 y FEDER ‘‘Una manera de hacer Europa,’’ and the ‘‘Center of Excellence Maria de Maeztu 2020-2023’’ award to the ICCUB (CEX2019-000918-M funded by MCIN/AEI/10.13039/501100011033).

B.D.W. acknowledges the support of the ANR BIG4 project, grant ANR-16-CE23-0002 of the French Agence Nationale de la Recherche; and the Labex ILP (reference ANR-10-LABX-63) part of the Idex SUPER, and received financial state aid managed by the Agence Nationale de la Recherche, as part of the program Investissements d’avenir under the reference ANR-11-IDEX-0004-02. The Flatiron Institute is supported by the Simons Foundation.

Appendix A Modal Estimator

A.1. Modal Basis

The custom modes we use in this work are based on those developed in Hung et al. (2019a) and Byun et al. (2021). From the one-dimensional functions

$$\begin{aligned}
 q_0^{\text{tree}} &= \sqrt{\frac{k}{P(k)}} \frac{5}{14}, \\
 q_1^{\text{tree}} &= \sqrt{\frac{k}{P(k)}} P_L(k), \\
 q_2^{\text{tree}} &= -\sqrt{\frac{k}{P(k)}} P_L(k) k^2, \\
 q_3^{\text{tree}} &= \sqrt{\frac{k}{P(k)}} \frac{P_L(k)}{k^2}, \\
 q_4^{\text{tree}} &= \sqrt{\frac{k}{P(k)}} \frac{3}{14} k^2, \\
 q_5^{\text{tree}} &= \sqrt{\frac{k}{P(k)}} \frac{1}{14} k^4,
 \end{aligned} \tag{A1}$$

where $P_L(k)$ is the linear matter power spectrum, one can write the following four modes:

$$\begin{aligned}
 Q_0^{\text{tree}}(k_1, k_2, k_3) &= q_{\{0\}}^{\text{tree}}(k_1) q_{\{1\}}^{\text{tree}}(k_2) q_{\{1\}}^{\text{tree}}(k_3), \\
 Q_1^{\text{tree}}(k_1, k_2, k_3) &= q_{\{0\}}^{\text{tree}}(k_1) q_2^{\text{tree}}(k_2) q_3^{\text{tree}}(k_3), \\
 Q_2^{\text{tree}}(k_1, k_2, k_3) &= q_{\{1\}}^{\text{tree}}(k_1) q_3^{\text{tree}}(k_2) q_4^{\text{tree}}(k_3), \\
 Q_3^{\text{tree}}(k_1, k_2, k_3) &= q_{\{3\}}^{\text{tree}}(k_1) q_3^{\text{tree}}(k_2) q_5^{\text{tree}}(k_3),
 \end{aligned} \tag{A2}$$

which fit exactly the standard tree-level matter bispectrum.

Similarly, the local bispectrum template given in Equation (2) can be described by a single mode:

$$Q_0^{\text{local}}(k_1, k_2, k_3) = q_{\{0\}}^{\text{local}}(k_1) q_{\{1\}}^{\text{local}}(k_2) q_{\{1\}}^{\text{local}}(k_3), \tag{A3}$$

where the one-dimensional basis functions are

$$\begin{aligned}
 q_0^{\text{local}} &= \sqrt{\frac{k}{P(k)}} \sqrt{P_L(k)} k^{n_s/2-2}, \\
 q_1^{\text{local}} &= \sqrt{\frac{k}{P(k)}} \sqrt{P_L(k)} k^{2-n_s/2}.
 \end{aligned} \tag{A4}$$

A.2. Gamma Matrix

The gamma matrix is defined as the inner product (Equation (7)) of the mode functions Q_n , given by

$$\begin{aligned}
 \gamma_{mn} &\equiv \langle Q_m | Q_n \rangle \\
 &= \frac{1}{8\pi^4} \int_{\mathcal{V}_T} dk_1 dk_2 dk_3 Q_m(k_1, k_2, k_3) Q_n(k_1, k_2, k_3).
 \end{aligned} \tag{A5}$$

The above inner product matrix needs to be calculated only once for a chosen basis function Q_n and scale range $[k_{\min}, k_{\max}]$. Different ways exist to calculate the integral of the γ matrix, where these methods have been summarized and compared extensively in Byun et al. (2021). In this work, we use the method applied for the first time in Karagiannis (2018), where the inner product is calculated by utilizing 1D FFTs. In order to describe this method, we need to take a step back and use the definition of the inner product used to derive Equations (9) and (10), as well as Equation (A5), which is

$$\begin{aligned}
 \langle Q_m | Q_n \rangle &= \int_{k_1} \int_{k_2} \int_{k_3} (2\pi)^3 \delta_D(\mathbf{k}_{123}) \\
 &\quad \times \frac{Q_m(k_1, k_2, k_3) Q_n(k_1, k_2, k_3)}{k_1 k_2 k_3},
 \end{aligned} \tag{A6}$$

where for brevity we introduce the shorthand notations $\int_{\mathbf{k}} \equiv \int \frac{d^3k}{(2\pi)^3}$ and $\mathbf{k}_{123} \equiv \mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3$. Following the steps of Ferguson et al. (2012b), we use

$$\delta_D(\mathbf{k}_{123}) = \frac{1}{(2\pi)^3} \int d^3x e^{i\mathbf{k}_{123} \cdot \mathbf{x}}, \tag{A7}$$

where the exponent can be written as

$$e^{i\mathbf{k} \cdot \mathbf{x}} = 4\pi \sum_{\ell m} i^\ell j_\ell(kx) Y_{\ell m}(\hat{\mathbf{k}}) Y_{\ell m}^*(\hat{\mathbf{x}}). \tag{A8}$$

Using the above equation into Equation (A6) we get

$$\begin{aligned}
 \langle Q_m | Q_n \rangle &= \int d^3x (4\pi)^3 \int_{k_1} \sum_{\ell_1 m_1} i^{\ell_1} j_{\ell_1}(k_1 x) Y_{\ell_1 m_1}(\hat{\mathbf{k}}_1) Y_{\ell_1 m_1}^*(\hat{\mathbf{x}}) \\
 &\quad \times \int_{k_2} \sum_{\ell_2 m_2} i^{\ell_2} j_{\ell_2}(k_2 x) Y_{\ell_2 m_2}(\hat{\mathbf{k}}_2) Y_{\ell_2 m_2}^*(\hat{\mathbf{x}}) \\
 &\quad \times \int_{k_3} \sum_{\ell_3 m_3} i^{\ell_3} j_{\ell_3}(k_3 x) Y_{\ell_3 m_3}(\hat{\mathbf{k}}_3) Y_{\ell_3 m_3}^*(\hat{\mathbf{x}}) \\
 &\quad \times \frac{Q_m(k_1, k_2, k_3) Q_n(k_1, k_2, k_3)}{k_1 k_2 k_3}.
 \end{aligned} \tag{A9}$$

The integral over \mathbf{k}_i can be separated into a radial and angular part, where the latter is written as

$$\int d\Omega_{\mathbf{k}_i} Y_{\ell_i m_i}(\hat{\mathbf{k}}_i) = \sqrt{4\pi} \delta_{\ell_i 0} \delta_{m_i 0}, \quad (\text{A10})$$

where to derive the above we have used the normalization of spherical harmonics, i.e., $\int d\Omega_{\mathbf{k}_i} Y_{\ell_i m_i}(\hat{\mathbf{k}}_i)^2 = 1$ and $Y_{00} = 1/\sqrt{4\pi}$. The above equation forces all ℓ_i and m_i in Equation (A9) to zero, giving

$$\begin{aligned} \langle Q_m | Q_n \rangle &= \frac{(4\pi)^{9/2}}{(2\pi)^9} \int dx x^2 \int d\Omega_x \int dk_1 k_1^2 j_0(k_1 x) Y_{00}(\hat{\mathbf{x}}) \\ &\times \int dk_2 k_2^2 j_0(k_2 x) Y_{00}(\hat{\mathbf{x}}) \\ &\times \int dk_3 k_3^2 j_0(k_3 x) Y_{00}(\hat{\mathbf{x}}) \\ &\times \frac{Q_m(k_1, k_2, k_3) Q_n(k_1, k_2, k_3)}{k_1 k_2 k_3}. \end{aligned} \quad (\text{A11})$$

Now, if we use for the integration over \mathbf{x} the fact that, $\int d\Omega_x Y_{00}(\mathbf{x})^3 = 1/\sqrt{4\pi}$ and $\int dx x^2 j_0(k_1 x) j_0(k_2 x) j_0(k_3 x) = \pi/(8k_1 k_2 k_3)$, where the latter is nonzero only for wavevectors that form a triangle, then we retrieve the expression in Equation (A6). However, at this point we will only use the expression for the angular part of the \mathbf{x} integral, while we substitute in Equation (A11) the basis functions $Q_m = q_{\{p_1\}}(k_1) q_{r_1}(k_2) q_{s_1}(k_3)$ and $Q_n = q_{\{p_2\}}(k_1) q_{r_2}(k_2) q_{s_2}(k_3)$, which is a separable product of one-dimensional functions $q_p(k)$. Furthermore, we use the expression for the zero-order Bessel function, i.e., $j_0(k_i x) = \sin(k_i x)/(k_i x)$, thus getting

$$\begin{aligned} \langle Q_m | Q_n \rangle &= \frac{1}{2\pi^5} \int dx \frac{1}{x} \frac{1}{6} \\ &(\xi_{p_1 p_2}(x) [\xi_{r_1 r_2}(x) \xi_{s_1 s_2}(x) + \xi_{r_1 s_2}(x) \xi_{s_1 r_2}(x)] \\ &\xi_{p_1 r_2}(x) [\xi_{r_1 s_2}(x) \xi_{s_1 p_2}(x) + \xi_{r_1 p_2}(x) \xi_{s_1 s_2}(x)] \\ &\xi_{p_1 s_2}(x) [\xi_{r_1 p_2}(x) \xi_{s_1 r_2}(x) + \xi_{r_1 r_2}(x) \xi_{s_1 p_2}(x)]), \end{aligned} \quad (\text{A12})$$

where

$$\begin{aligned} \xi_{pr}(x) &\equiv \int_{k_{\min}}^{k_{\max}} dk q_p \left(\frac{k - k_{\min}}{k_{\max} - k_{\min}} \right) \\ &\times q_r \left(\frac{k - k_{\min}}{k_{\max} - k_{\min}} \right) \sin(kx). \end{aligned} \quad (\text{A13})$$

The above integral over k can be calculated by using one-dimensional FFTs, as described in Chapter 13.9 of Numerical Recipes (Press et al. 1992), while the integration over x in Equation (A12) can be done by using conventional numerical integration methods, like the five-point Newton–Cotes quadrature rule used here. The accuracy of the outlined method, on calculating Equation (A12), depends on the grid resolution of k , which has to be sufficient, not only to accurately calculate the integral of Equation (A13), but the derived range and resolution of x generated by the FFT, should be adequate enough, as well, to retrieve accurate numerical result for the outermost integration over x . More details on this method can be found in Press et al. (1992).

Appendix B Impact of Cosmological Parameters

In Figures 13 and 14, we show the impact on the power spectrum and on the bispectrum of varying the cosmological parameters considered in this work, similar to what was done for PNG in Section 4.2. This highlights distinct behaviors for each parameter and each observable, in contrast with PNG, which is strongly degenerate at the power spectrum level.




Figure 13. Impact of cosmological parameters on the matter power spectrum (averaged from 500 N -body simulations). Solid and dashed lines correspond respectively to values perturbed above and below the fiducial.




Figure 14. Impact of cosmological parameters on the matter bispectrum (ratio bispectrum with a modified parameter to fiducial, averaged from 500 N -body simulations).

ORCID iDs

Gabriel Jung <https://orcid.org/0000-0002-2133-6254>
 Dionysios Karagiannis <https://orcid.org/0000-0002-4927-0816>
 Michele Liguori <https://orcid.org/0000-0002-0383-2254>
 Marco Baldi <https://orcid.org/0000-0003-4145-1943>
 William R. Coulton <https://orcid.org/0000-0002-1297-3673>
 Drew Jamieson <https://orcid.org/0000-0001-5044-7204>
 Licia Verde <https://orcid.org/0000-0003-2601-8770>
 Francisco Villaescusa-Navarro <https://orcid.org/0000-0002-4816-0455>
 Benjamin D. Wandelt <https://orcid.org/0000-0002-5854-8269>

References

- Afshordi, N., & Tolley, A. J. 2008, *PhRvD*, **78**, 123507
 Akrami, Y., Arroja, F., Ashdown, M., et al. 2020, *A&A*, **641**, A9
 Alsing, J., Charnock, T., Feeney, S., & Wandelt, B. 2019, *MNRAS*, **488**, 4440
 Alsing, J., & Wandelt, B. 2018, *MNRAS*, **476**, L60
 Alvarez, M., Baldauf, T., Bond, R., et al. 2014, arXiv:1412.4671
 Amendola, L., Appleby, S., Avgoustidis, A., et al. 2018, *LRR*, **21**, 2
 Andrews, A., Jasche, J., Lavaux, G., & Schmidt, F. 2022, arXiv:2203.08838
 Baumann, D., & Green, D. 2021, *JCAP*, **2022**, 061
 Biagetti, M., Calles, J., Castiblanco, L., Cole, A., & Noreña, J. 2022a, *JCAP*, **2022**, 002
 Biagetti, M., Castiblanco, L., Noreña, J., & Sefusatti, E. 2022b, *JCAP*, **2022**, 009
 Biagetti, M., Cole, A., & Shiu, G. 2021, *JCAP*, **04**, 061
 Byun, J., & Krause, E. 2022, *MNRAS*, Advance Access
 Byun, J., Oddo, A., Porciani, C., & Sefusatti, E. 2021, *JCAP*, **2021**, 105
 Cabass, G., Ivanov, M. M., Philcox, O. H. E., Simonović, M., & Zaldarriaga, M. 2022a, *PhRvD*, **106**, 043506
 Cabass, G., Ivanov, M. M., Philcox, O. H. E., Simonović, M., & Zaldarriaga, M. 2022b, *PhRvL*, **129**, 021301
 Chartier, N., Wandelt, B., Akrami, Y., & Villaescusa-Navarro, F. 2021, *MNRAS*, **503**, 1897
 Chartier, N., & Wandelt, B. D. 2021, *MNRAS*, **509**, 2220
 Chartier, N., & Wandelt, B. D. 2022, *MNRAS*, **515**, 1296
 Coulton, W. R., Villaescusa-Navarro, F., Jamieson, D., et al. 2022, arXiv:2206.01619
 Crocce, M., Pueblas, S., & Scoccimarro, R. 2006, *MNRAS*, **373**, 369
 Dalal, N., Dore, O., Huterer, D., & Shirokov, A. 2008, *PhRvD*, **77**, 123514
 D’Amico, G., Lewandowski, M., Senatore, L., & Zhang, P. 2022, arXiv:2201.11518
 Dawson, K. S., Schlegel, D. J., Ahn, C. P., et al. 2013, *AJ*, **145**, 10
 Desjacques, V., & Seljak, U. 2010, *AdAst*, **2010**, 908640
 Feldman, H. A., Kaiser, N., & Peacock, J. A. 1994, *ApJ*, **426**, 23
 Fergusson, J. R., Liguori, M., & Shellard, E. P. S. 2010, *PhRvD*, **82**, 023502
 Fergusson, J. R., Liguori, M., & Shellard, E. P. S. 2012a, *JCAP*, **12**, 032
 Fergusson, J. R., Regan, D. M., & Shellard, E. P. S. 2012b, *PhRvD*, **86**, 063511
 Friedrich, O., Uhlemann, C., Villaescusa-Navarro, F., et al. 2020, *MNRAS*, **498**, 464
 Gil-Marín, H. 2022, *JCAP*, **05**, 040

- Gil-Marín, H., Noreña, J., Verde, L., et al. 2015a, *MNRAS*, **451**, 539
- Gil-Marín, H., Percival, W. J., Verde, L., et al. 2017, *MNRAS*, **465**, 1757
- Gil-Marín, H., Verde, L., Noreña, J., et al. 2015b, *MNRAS*, **452**, 1914
- Giri, U., Münchmeyer, M., & Smith, K. M. 2022, arXiv:2205.12964
- Gualdi, D., Gil-Marín, H., Schuhmann, R. L., et al. 2019, *MNRAS*, **484**, 3713
- Gualdi, D., Gil-Marín, H., & Verde, L. 2021a, *JCAP*, **07**, 008
- Gualdi, D., Manera, M., Joachimi, B., & Lahav, O. 2018, *MNRAS*, **476**, 4045
- Gualdi, D., Novell, S., Gil-Marín, H., & Verde, L. 2021b, *JCAP*, **01**, 015
- Hahn, C., & Villaescusa-Navarro, F. 2021, *JCAP*, **2021**, 029
- Hahn, C., Villaescusa-Navarro, F., Castorina, E., & Scoccimarro, R. 2020, *JCAP*, **03**, 040
- Hartlap, J., Simon, P., & Schneider, P. 2007, *A&A*, **464**, 399
- Heavens, A. F., Jimenez, R., & Lahav, O. 2000, *MNRAS*, **317**, 965
- Hung, J., Fergusson, J. R., & Shellard, E. P. S. 2019a, arXiv:1902.01830
- Hung, J., Manera, M., & Shellard, E. P. S. 2019b, arXiv:1909.03248
- Ivanov, M. M., Philcox, O. H. E., Nishimichi, T., et al. 2022, *PhRvD*, **105**, 063512
- Karagiannis, D. 2018, PhD dissertation, University of Padova
- Karagiannis, D., Lazanu, A., Liguori, M., et al. 2018, *MNRAS*, **478**, 1341
- Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, arXiv:1110.3193
- Lazanu, A., Giannantonio, T., Schmittfull, M., & Shellard, E. P. S. 2016, *PhRvD*, **93**, 083517
- Lazanu, A., Giannantonio, T., Schmittfull, M., & Shellard, E. P. S. 2017, *PhRvD*, **95**, 083511
- Lewis, A., Challinor, A., & Lasenby, A. 2000, *ApJ*, **538**, 473
- Mao, Q., Berlind, A. A., McBride, C. K., et al. 2014, *MNRAS*, **443**, 1402
- Matarrese, S., & Verde, L. 2008, *ApJL*, **677**, L77
- Meerburg, P. D., Green, D., Abidi, M., et al. 2019, arXiv:1903.04409
- Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. 1992, *Numerical Recipes in FORTRAN 77: The Art of Scientific Computing* (2nd ed; Cambridge: Cambridge Univ. Press)
- Regan, D. M., Schmittfull, M. M., Shellard, E. P. S., & Fergusson, J. R. 2012, *PhRvD*, **86**, 123524
- Schmittfull, M. M., Regan, D. M., & Shellard, E. P. S. 2013, *PhRvD*, **88**, 063512
- Scoccimarro, R., Hui, L., Manera, M., & Chan, K. C. 2012, *PhRvD*, **85**, 083002
- Sefusatti, E., Crocce, M., Scoccimarro, R., & Couchman, H. M. P. 2016, *MNRAS*, **460**, 3624
- Sellentin, E., & Heavens, A. F. 2016, *MNRAS*, **456**, L132
- Shirasaki, M., Sugiyama, N. S., Takahashi, R., & Kitaura, F.-S. 2021, *PhRvD*, **103**, 023506
- Slepian, Z., Eisenstein, D. J., Beutler, F., et al. 2017a, *MNRAS*, **468**, 1070
- Slepian, Z., Eisenstein, D. J., Brownstein, J. R., et al. 2017b, *MNRAS*, **469**, 1738
- Slosar, A., Hirata, C., Seljak, U., Ho, S., & Padmanabhan, N. 2008, *JCAP*, **2008**, 031
- Springel, V. 2005, *MNRAS*, **364**, 1105
- Uhlemann, C., Pajer, E., Pichon, C., et al. 2018, *MNRAS*, **474**, 2853
- Verde, L., & Matarrese, S. 2009, *ApJL*, **706**, L91
- Villaescusa-Navarro, F., Hahn, C., Massara, E., et al. 2020, *ApJS*, **250**, 2
- Wagner, C., Verde, L., & Boubekur, L. 2010, *JCAP*, **10**, 022