



Quijote-PNG: Quasi-maximum Likelihood Estimation of Primordial Non-Gaussianity in the Nonlinear Halo 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

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

⁶Istituto 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 boulevard Arago, F-75014 Paris, France

Received 2023 January 2; revised 2023 March 10; accepted 2023 March 14; published 2023 May 17

Abstract

We study primordial non-Gaussian signatures in the redshift-space halo field on nonlinear scales, using a quasi-maximum likelihood estimator based on optimally compressed power spectrum and modal bispectrum statistics. We train and validate the estimator on a suite of halo catalogs constructed from the QUIJOTE-PNG N -body simulations, which we release to accompany this paper. We verify its unbiasedness and near-optimality for the three main types of primordial non-Gaussianity (PNG): local, equilateral, and orthogonal. We compare the modal bispectrum expansion with a k -binning approach, showing that the former allows for faster convergence of numerical derivatives in the computation of the score function, thus leading to better final constraints. We find, in agreement with previous studies, that the local PNG signal in the halo field is dominated by the scale-dependent bias signature on large scales and saturates at $k \sim 0.2 h \text{ Mpc}^{-1}$, whereas the small-scale bispectrum is the main source of information for equilateral and orthogonal PNG. Combining the power spectrum and bispectrum on nonlinear scales plays an important role in breaking degeneracies between cosmological and PNG parameters; such degeneracies, however, remain strong for equilateral PNG. We forecast that PNG parameters can be constrained with $\Delta f_{\text{NL}}^{\text{local}} = 45$, $\Delta f_{\text{NL}}^{\text{equil}} = 570$, and $\Delta f_{\text{NL}}^{\text{ortho}} = 110$ on a cubic volume of $1(\text{Gpc } h^{-1})^3$ at $z = 1$, considering scales up to $k_{\text{max}} = 0.5 h \text{ Mpc}^{-1}$.

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

1. Introduction

The coming generation of spectroscopic and photometric galaxy surveys—e.g., Euclid, DESI, Sphex, Rubin Observatory, and Roman (LSST Science Collaboration et al. 2009; Laureijs et al. 2011; Doré et al. 2014; DESI Collaboration et al. 2016)—will allow us to study galaxy clustering with an unprecedented level of accuracy and precision, shedding further light on many open questions in cosmology. Among the many exciting possibilities, an interesting prospect, which we mainly focus on in this work, is that of improving our understanding of early universe physics, via high-precision tests of primordial non-Gaussianity (PNG).

Cosmic microwave background (CMB) measurements (Akrami et al. 2020), in agreement with theoretical expectations, have constrained the primordial cosmological perturbation field to be at most weakly non-Gaussian. This implies, for a large majority of early universe scenarios, that most of the PNG information is contained in the primordial bispectrum. For this reason, the bispectrum of dark matter tracers (e.g.,

galaxies) in the large-scale structure (LSS) can be a powerful probe of PNG. Crucially, the 3D galaxy bispectrum also gives us access, in principle, to a larger number of modes with respect to the 2D (angular) bispectrum of CMB anisotropies. Therefore, LSS bispectrum analyses can potentially lead to significant improvements in PNG constraints over current, CMB-based results. Achieving such improvements, however, will require including nonlinear scales in the analysis, carrying strong non-Gaussian (NG) signatures that are not primordial, but arise from the late-time, nonlinear evolution of cosmic structures. Disentangling the NG late-time component from the subdominant primordial one is therefore a crucial challenge in this kind of studies. It can be addressed either by analytical modeling of the bispectrum—via a suitable perturbative approach at mildly nonlinear scales (Cabass et al. 2022a, 2022b; D’Amico et al. 2022)—or by relying on fully numerical approaches, which evaluate the bispectrum (and/or other summary statistics; see, e.g., Friedrich et al. 2020; Biagetti et al. 2021; Valogiannis & Dvorkin 2022) using large mock data sets; field-level inference on large scales, not relying on specific statistical summaries, has also been recently considered (see Andrews et al. 2023). In this work, which is the fourth in a series of papers, following Jung et al. (2022) and Coulton et al. (2022a, 2022b), we base our study on the

QUIJOTE-PNG simulation suite, recently presented in Coulton et al. (2022a). Our main goal is to quantify the accuracy with which f_{NL} can be constrained using both the power spectrum and bispectrum of the dark matter halo field, up to strongly nonlinear scales ($k_{\text{max}} = 0.5 h \text{ Mpc}^{-1}$), using simulations with different kinds of PNG. More precisely, we consider the three main PNG bispectrum shapes, namely, the local, equilateral, and orthogonal shapes, which are predicted in a large variety of inflationary scenarios.

This work extends our initial analysis presented in Jung et al. (2022), where we worked at the level of the matter field. As in the previous analysis, we derive forecasts for PNG and standard cosmological parameters, by combining power spectrum and bispectrum measurements at nonlinear scales; our main focus is then on building an unbiased and nearly optimal quasi-maximum likelihood estimator, based on applying a MOPED-like compression algorithm to a modal decomposition of the data bispectrum. In our companion paper (Coulton et al. 2022b) we independently performed a similar analysis at a different redshift ($z = 0$, versus $z = 1$ in this work), but we employed a binned decomposition in k -space instead of the modal approach adopted here; in that work we studied in detail the PNG information content in the halo field while focusing on important numerical convergence issues. Therefore, the two analyses are complementary for studying the robustness of our approach and covering a full range of crucial issues, from numerical stability to the precise quantification of the information gain obtained from different observables at different scales and the demonstration of nearly optimal and unbiased bispectrum data compression for parameter estimation. Taken together, we think these works represent an important development in the effort to build a data analysis pipeline to be applied to observations. We release the halo catalogs of the QUIJOTE-PNG¹³ suite used in these works.

Since in the current work we consider tracers of the underlying density field, an additional signature of PNG arises—in comparison to those of the previous matter field analysis—in the form of a scale dependency in the tracer bias. Such a feature has a power-law behavior, with a degree determined by the squeezed limit of the PNG bispectrum shape under study: it is most prominent for local NG shapes, with a $\sim 1/k^2$ behavior, and absent in the equilateral case. Scale-dependent bias has been the object of significant study in the literature (see, e.g., Afshordi & Tolley 2008; Dalal et al. 2008; Matarrese & Verde 2008; Slosar et al. 2008; Seljak 2009; Desjacques & Seljak 2010; Castorina et al. 2018; Chan et al. 2019; Giri et al. 2022) and has been used to extract local PNG constraints from BOSS data (Slosar et al. 2008; Ross et al. 2013; Leistedt et al. 2014; Mueller et al. 2021; Cabass et al. 2022a; D’Amico et al. 2022). Recently, however, it has been pointed out that accurate modeling of scale-dependent bias from PNG also depends on the details of galaxy formation, making its use as a tool for measuring the PNG amplitude f_{NL} significantly more challenging than previously thought (Barreira 2020, 2022a, 2022b). The effect of scale-dependent bias is automatically incorporated in our analysis, where we are mostly concerned with assessing its relative constraining power on different NG shapes, as compared to the bispectrum, and verifying agreement with both our analysis in Coulton et al.

(2022b) and theoretical expectations (see, e.g., de Putter 2018; Karagiannis et al. 2018).

The paper is structured as follows. In Section 2 we briefly review the NG models considered in the analysis. In Section 3 we illustrate our methodology for data compression and parameter estimation. In Section 4 we discuss our Fisher matrix analysis, showing expected parameter constraints on different scales, and describe the application of our quasi-maximum likelihood, joint power spectrum and bispectrum estimator to simulated data. In Section 5 we summarize our main results and draw our final conclusions. In Appendix A we provide more details about the implementation of shot-noise modes in the bispectrum estimator. In Appendix B we show a comparison between the modal and binned approaches to bispectrum estimation, and in Appendix C we discuss the results of a preliminary study aimed at the application of the CARPool technique to the evaluation of covariances and numerical derivatives.

2. Bispectrum Shapes

Violating any condition of the standard inflationary model induces a deviation from the perfect Gaussian initial conditions, which leads to nonzero high-order correlators. The largest of them, in most inflationary models, is the bispectrum, i.e., the three-point correlation function of Fourier modes, defined as

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

The primordial bispectrum is generally written as

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

where f_{NL} is the dimensionless amplitude parameter corresponding to a given primordial bispectrum shape $F(k_1, k_2, k_3)$, which encompasses the dependence of the bispectrum on triplets of Fourier space modes.

In this work, we focus on building estimators to measure the f_{NL} of three of the most common primordial shapes, namely local, equilateral, and orthogonal¹⁴ bispectra (see Coulton et al. 2022a and references therein for a complete description of these templates).

For dark matter tracers (e.g., halos), the presence of PNG has a significant impact, due to the introduced coupling between large- and small-scale modes, with the most widely known example being that of the local type. In this case, the halo overdensity on large scales will no longer depend only on the matter overdensity, but also on the primordial gravitational potential (see Desjacques et al. 2018 for a review). This results in a scale-dependent term that introduces an important PNG signature on large scales of a correlator. This is of particular importance to the power spectrum of observed dark matter tracers, since it enhances the PNG signal within the two-point correlation function, which otherwise would have been very limited, e.g., in the case of dark matter field analysis (see, e.g., Coulton et al. 2022a; Jung et al. 2022).

The effect of the scale-dependent term on the power spectrum has been extensively studied in the literature, especially for the local PNG type. However, recent developments have made the measurement of f_{NL} , by such a term in the power spectrum, challenging, due to the perfect degeneracy

¹³ <https://quijote-simulations.readthedocs.io/en/latest/png.html>

¹⁴ We use the orthogonal-LSS shape described in Coulton et al. (2022a), which is a better approximation of the nonseparable orthogonal bispectrum than the standard template of CMB analyses for the 3D matter field.

between f_{NL} and the scale-dependent bias coefficient b_ϕ (Barreira 2020, 2022b). For the halo bispectrum, a significant amount of the PNG signal is located within the primordial part (Equation (2)), while the scale-dependent terms (studied at the theoretical level, e.g., in Karagiannis et al. 2018), which could carry a notable amount of information on local PNG, suffer from the same limitations as the power spectrum (Barreira 2022a).

The effect of the scale-dependent bias will be taken into account within the framework of the forward modeling. In a simulation-based approach we assume tight priors on the scale-dependent bias coefficient b_ϕ , in order to focus on the f_{NL} constraints (see also Coulton et al. 2022b for an extensive discussion).

3. Method

In this section we review the main aspects of our methodology for data compression and quasi-maximum likelihood estimation of cosmological and PNG parameters, starting from the evaluation of power spectrum and modal bispectrum summary statistics.

3.1. Quasi-maximum Likelihood Estimator

Starting from a given data vector \mathbf{d} (a given set of summary statistics, like the power spectrum and/or the bispectrum) that depends on some parameters of interest denoted by $\boldsymbol{\theta}$ (e.g., f_{NL}), one can write the following quasi-maximum likelihood estimator for the value of the parameters (see Alsing & Wandelt 2018 for details):

$$\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}_* + \mathbf{F}_*^{-1} [\nabla_{\boldsymbol{\theta}} \boldsymbol{\mu}_*^T \mathbf{C}_*^{-1} (\mathbf{d} - \boldsymbol{\mu}_*)] \equiv \boldsymbol{\theta}_* + \mathbf{F}_*^{-1} \mathbf{t}, \quad (3)$$

where the subscript $*$ denotes that the quantities are evaluated at some chosen fiducial point, and $\boldsymbol{\mu}$ and \mathbf{C} are, respectively, the mean and the covariance of \mathbf{d} . The two key ingredients of this estimator, which are the Fisher information \mathbf{F} and the compressed score statistic \mathbf{t} , will be detailed below. Note also that in this expression we assume a Gaussian likelihood and a dependence on parameters through the mean only, a reasonable assumption as verified in Jung et al. (2022).

The Fisher matrix, a standard method to evaluate the information content of some observables, is given by

$$\mathbf{F} = \nabla_{\boldsymbol{\theta}} \boldsymbol{\mu}^T \mathbf{C}^{-1} \nabla_{\boldsymbol{\theta}}^T \boldsymbol{\mu}. \quad (4)$$

This requires knowledge of the derivatives $\nabla_{\boldsymbol{\theta}} \boldsymbol{\mu}$ and the covariance \mathbf{C} , which can be both evaluated from a large set of simulations, as we do in this work (see Section 4). However, reaching numerical convergence for the joint analysis of multiple parameters may be very challenging and therefore prone to wrong results. In Jung et al. (2022) we checked, by studying the matter field, that if the covariance matrix is not converged, it would typically induce suboptimal error bars, and that nonconverged derivatives could bias the estimated parameters. In Coulton et al. (2022b) we showed that noisy derivatives could lead to overconfident error bars when working with the halo field.

To tackle this problem, the alternative compressed Fisher method described in Coulton et al. (2022b; see also W. Coulton & B. Wandelt 2023, in preparation, for details) can provide conservative bounds. It consists of two steps, the first of which

is to compress the data to the score function using (see Alsing & Wandelt 2018)

$$\mathbf{t} = \nabla_{\boldsymbol{\theta}} \boldsymbol{\mu}_*^T \mathbf{C}_*^{-1} (\mathbf{d} - \boldsymbol{\mu}_*), \quad (5)$$

which is equivalent to the MOPED compression scheme of Heavens et al. (2000). This operation reduces the data vector \mathbf{d} of size n to only p , which is the number of parameters of interest, while keeping all relevant information about these parameters. Then, to compute the compressed Fisher matrix one has only to apply the standard expression of Equation (4) to the compressed data. An important subtlety of this scheme is that it requires one to use two separate sets of simulations for the two different steps. The first part is used for the compression step, to build new summary statistics, which will be suboptimal if derivatives are noisy. The second part is then compressed and used to estimate the Fisher matrix from the compressed statistics, which is suboptimal if the compression step is suboptimal, but is also a lot less noisy due to the much lower dimensionality of the compressed statistics.¹⁵

3.2. Summary Statistics

In this work, we use the same observables as in Jung et al. (2022), based on the power spectrum and bispectrum statistics as they contain significant and complementary information about both the Λ CDM cosmological parameters and the PNG amplitudes f_{NL} .

The standard power spectrum estimator of a field $\delta(\mathbf{k})$ defined on a grid of fundamental mode k_f is given by

$$\hat{P}(k_i) = \frac{1}{V N_i} \sum_{\mathbf{k} \in \Delta_i} \delta(\mathbf{k}) \delta^*(\mathbf{k}), \quad (6)$$

where V is the survey volume, and a binning of the k range has been introduced with each bin Δ_i having a width k_f and containing N_i independent vectors of \mathbf{k} .

As was initially shown for CMB NG analysis in Fergusson et al. (2010, 2012a), and extended later to the 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 bispectrum information can be efficiently extracted from data by measuring the following modal coefficients:

$$\hat{\beta}_n = \frac{1}{V} \int d^3x M_p(\mathbf{x}) M_q(\mathbf{x}) M_r(\mathbf{x}), \quad (7)$$

where

$$M_p(\mathbf{x}) \equiv \int \frac{d^3k}{(2\pi)^3} \frac{q_p(k) \delta(\mathbf{k})}{\sqrt{kP(k)}} e^{i\mathbf{k}\cdot\mathbf{x}}, \quad (8)$$

for a well-chosen basis of one-dimensional functions $q_p(k)$ and mode triplets $n \leftrightarrow (p, q, r)$. We refer the reader to Jung et al. (2022) for the details of the exact setup we use for the analyses presented in Section 4, as they are almost identical (the only change is the addition of two special modes, introduced in Byun et al. 2021 and recalled in Appendix A, describing the shot-noise component of the bispectrum expected from halos).

¹⁵ The variance of the procedure can be significantly decreased by repeating the procedure many times (generating different splittings each time) and computing the Monte Carlo average of the compressed Fisher matrix.

Table 1
The Fiducial Values of the Cosmological Parameters and PNG Amplitudes, together with Their Variations, Used in the Analysis

	σ_8	Ω_m	n_s	h	$M_{\min} (M_{\odot} h^{-1})$	$f_{\text{NL}}^{\text{local}}$	$f_{\text{NL}}^{\text{equil}}$	$f_{\text{NL}}^{\text{ortho}}$
Fiducial	0.834	0.3175	0.9624	0.6711	3.2×10^{13}	0	0	0
Steps	± 0.015	± 0.01	± 0.02	± 0.02	$\pm 0.1 \times 10^{13}$	± 100	± 100	± 100

4. Analyses

4.1. Specifications

For our analysis we use the publicly available QUIJOTE¹⁶ and QUIJOTE-PNG¹³ suites of N -body simulations (Villaescusa-Navarro et al. 2020; Coulton et al. 2022a). Each simulation represents a periodic cubic box of length $1 h^{-1}$ Gpc that contains 512^3 particles, run with the GADGET-III code (Springel 2005). Initial conditions are generated at $z_i = 127$ with the code 2LPTIC (Crocco et al. 2006) in the Gaussian case and the code 2LPTPNG¹⁷ in the NG case (Scoccimarro et al. 2012; Coulton et al. 2022a); linear matter power spectra and transfer functions are obtained from CAMB (Lewis et al. 2000). Finally, dark matter halos are identified using the friends-of-friends algorithm (Davis et al. 1985) with the value of the linking length equal to $b = 0.2$; we select those with a mass larger than $M_{\min} = 3.2 \times 10^{13} M_{\odot} h^{-1}$, corresponding to a number density $\bar{n} \sim 5.1 \times 10^{-5} h^3 \text{Mpc}^{-3}$ at $z = 1$ (see Hahn et al. 2020 for a power spectrum and bispectrum analysis of these halo catalogs focused on cosmological parameters).

We construct the halo density field in redshift space at $z = 1$ by depositing the halo positions, displaced radially by the velocity, on a grid of size $N_{\text{grid}} = 256$, using a fourth-order interpolation scheme implemented in the PYLIANS3 code¹⁸ (Villaescusa-Navarro 2018). We then measure the power spectrum and modal bispectrum monopoles using the estimators in Equations (6) and (7) including modes up to $k_{\max} = 0.5 h \text{Mpc}^{-1}$.

For the numerical computation of the covariance matrix, we have 15,000 simulations at fiducial cosmology, whereas smaller sets of 500 realizations, with varying input parameters, are used to evaluate the derivatives in Equation (5). In particular the analysis is focused on the PNG amplitudes of the three shapes considered here, $f_{\text{NL}}^{\text{local}}$, $f_{\text{NL}}^{\text{equil}}$, and $f_{\text{NL}}^{\text{ortho}}$; four cosmological parameters Ω_m , n_s , σ_8 , and h ; and one parameter related to the halo bias, M_{\min} . The variation of the minimum halo mass generates distinct catalogs with $M_{\min}^{\text{fid}} \pm \Delta M_{\min}$ (see Table 1), which consequently leads to a variation of the halo number density. This is roughly equivalent to a variation of the linear bias contribution (see, e.g., Desjacques et al. 2018 for details), though it propagates, to a minor extent, to higher-order terms. Although this bias model is quite simplistic, it is still useful within the framework of a first-order analysis presented in this work. A thorough investigation on the impact of the bias parameters, within a simulation-based approach, requires a population of halos with a halo occupation distribution (HOD) and the variation of the HOD parameters, which is left for future work. More details about the specifications of these simulations, concerning all the parameters considered in our analyses, can be found in Table 1.

¹⁶ <https://quijote-simulations.readthedocs.io>

¹⁷ <https://github.com/dsjamieson/2LPTPNG>

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

4.2. Fisher Constraints

We aim to evaluate the information on Λ CDM parameters and PNG amplitudes contained in the power spectrum and bispectrum of the halo field at redshift $z = 1$ using a Fisher matrix formalism. This analysis complements the work of Coulton et al. (2022a), as we focus on a different redshift and make use of a different bispectrum estimator. We explore the dependence on the number of simulations used and the chosen k_{\max} and the role of the different summary statistics. We show the results in Figures 1–3 and Table 2.

As highlighted in Coulton et al. (2022a), a main difficulty of this simulation-based approach is to accurately compute numerical derivatives of both the power spectrum and bispectrum with respect to the different parameters considered. This is illustrated in Figure 1, where we show that using smaller subsets of the 500 available pairs of simulations per parameter leads to spurious smaller 1σ uncertainties when jointly analyzing $\{\sigma_8, \Omega_m, h, n_s, M_{\min}, f_{\text{NL}}^{\text{local}}, f_{\text{NL}}^{\text{equil}}, f_{\text{NL}}^{\text{ortho}}\}$ using the traditional Fisher uncertainties (i.e., the dashed lines), indicating a lack of numerical convergence.

Instead of the computationally intensive possibility of producing many more simulations, we use here an alternative method, described briefly in Section 3.1, of computing conservative constraints from a lower number of simulations. As expected, the resulting 1σ error bars decrease when we use more simulations to calculate numerical derivatives, and using the full set they are only between 5% and 25% larger than the unconverged standard Fisher constraints. The results are stable when using 250 pairs of simulations or more to compute each derivative.

In the simpler situation where we consider only the three parameters $\{\sigma_8, \Omega_m, f_{\text{NL}}^{\text{local}}\}$ in the analysis, the two methods give very similar results when numerical convergence is reached (using at least 100 pairs of simulations per derivative). This is why in the rest of this work we always use the conservative approach, knowing it is equivalent to the standard Fisher approach in cases where numerical accuracy can be reached with the available simulations, and otherwise only leads to a reasonable overestimation of order 10% as verified. Note that we manage to keep this overestimation small here due to the use of the modal bispectrum, rather than a standard “binned” approach, because it compresses the original data more efficiently leading to more stable numerical derivatives (we need less than 50 modes to extract the full information of the bispectrum up to $k_{\max} = 0.5 h \text{Mpc}^{-1}$). This is shown explicitly in Appendix B.

In Figure 2 we study the dependence of the constraints on k_{\max} , considering values from 0.1 to $0.5 h \text{Mpc}^{-1}$. The largest improvement (for both Λ CDM cosmological and PNG parameters) is obtained between $k_{\max} = 0.1$ and $0.2 h \text{Mpc}^{-1}$, at which point the error bars on $f_{\text{NL}}^{\text{local}}$ become saturated (as well as those on h). However, for the equilateral and orthogonal shapes considering smaller scales yields better constraints (a few percent for each additional increase of $0.1 h \text{Mpc}^{-1}$). For other parameters, the gain can even be larger, justifying the

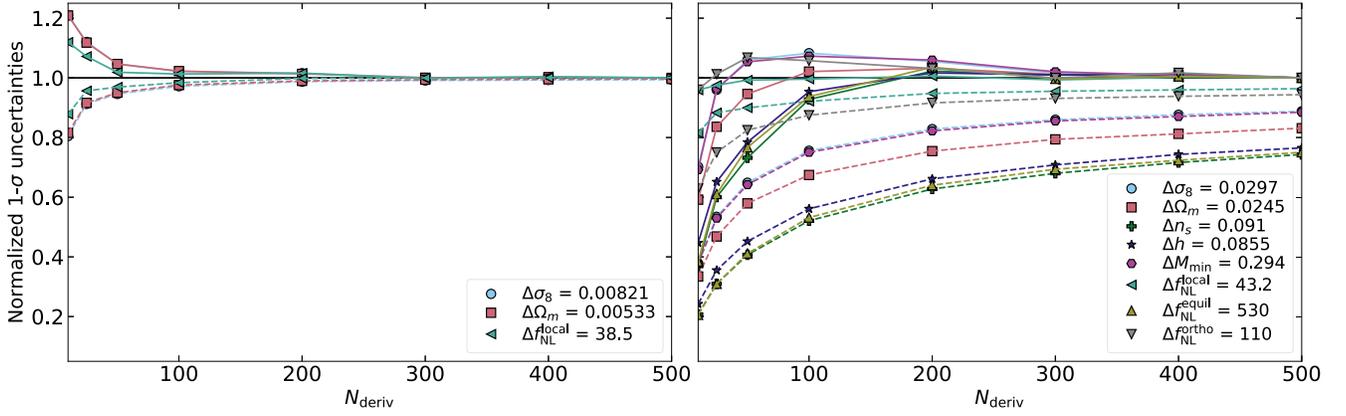


Figure 1. The stability of the Fisher 1σ uncertainties under variations of the number of simulations used to compute derivatives. The analysis includes both the power spectrum and the bispectrum information of the halo field at $z = 1$, with scales up to $k_{\max} = 0.5 h \text{ Mpc}^{-1}$. In the left panel, three parameters $\{\sigma_8, \Omega_m, f_{\text{NL}}^{\text{local}}\}$ are analyzed jointly, while in the right panel $\{h, n_s, M_{\min}, f_{\text{NL}}^{\text{equil}}, f_{\text{NL}}^{\text{ortho}}\}$ are also included. The dashed lines correspond to the standard Fisher error bars and the solid lines are computed from the compressed summary statistics (see Section 3.1).

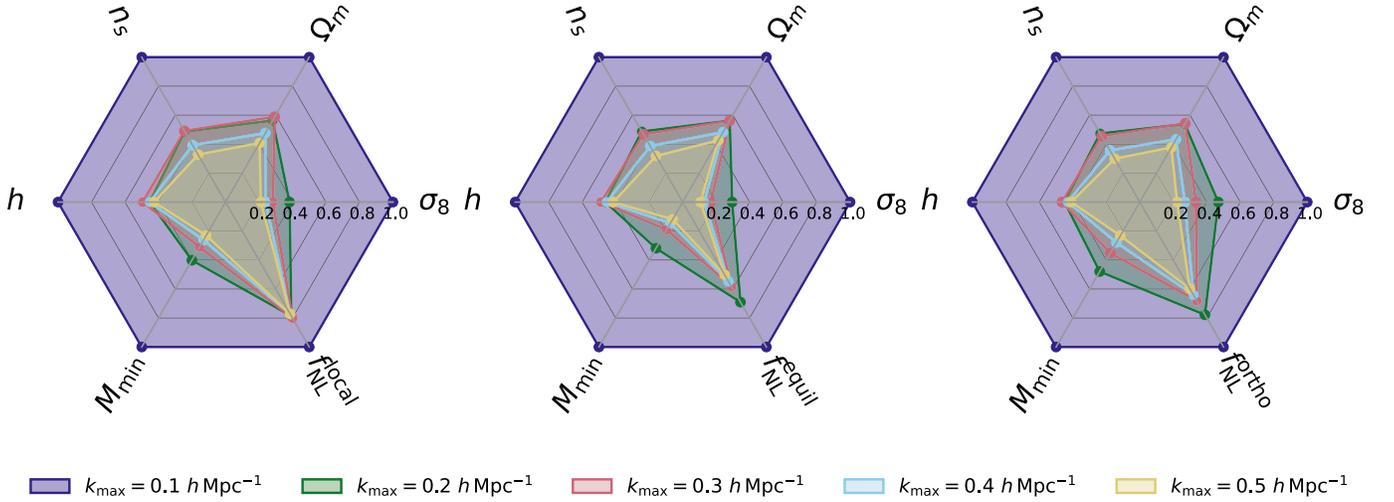


Figure 2. Normalized 1σ Fisher error bars for the joint analysis of cosmological parameters and one PNG shape at a time, using both the power spectrum and modal bispectrum information of the halo field at $z = 1$ for different k_{\max} from $0.1 h \text{ Mpc}^{-1}$ to $0.5 h \text{ Mpc}^{-1}$. All error bars are computed from the compressed summary statistics (see Section 3.1).

need to probe these nonlinear scales. Note also that all these improvements are computed using the conservative error bars obtained from the compressed summary statistics. Including smaller scales in the analysis typically leads to less converged numerical derivatives, and the less converged these derivatives are, the more suboptimal the conservative approach becomes. This means that we may be underestimating slightly the constraining power of small scales (in any case, this effect should not be large, as for $k_{\max} = 0.5 h \text{ Mpc}^{-1}$ this overestimation is $\sim 10\%$ as can be seen in Figure 1).

In Figure 3 we show the information content of the halo power spectrum, the halo bispectrum, and their combination. The bispectrum is a much more efficient probe of the equilateral and orthogonal shapes than the power spectrum, while for other parameters they yield constraints of the same order separately. Their combination always helps to reduce degeneracies, although to a lesser extent than that for the matter field studied previously in Coulton et al. (2022a) and Jung et al. (2022).

In Table 2, we present the 1σ conservative constraints on the ΛCDM parameters and PNG amplitudes using jointly the power spectrum and bispectrum and including small scales up to $k_{\max} = 0.5 h \text{ Mpc}^{-1}$. Unlike the matter field case discussed in Coulton et al. (2022a) and Jung et al. (2022), including PNG shapes in the analysis slightly increases the error bars on the ΛCDM cosmological parameters. The different PNG shapes are also less correlated, as analyzing them jointly increases only slightly their own error bars.

4.3. Parameter Estimation

As was shown in Jung et al. (2022) for the matter field, the simple quasi-maximum likelihood estimator (see Equation (3)) built from the Fisher matrix at some chosen fiducial cosmology is very efficient at measuring ΛCDM cosmological parameters and PNG amplitudes using power spectrum and bispectrum information. Here we extend this conclusion to the halo field.

The key ingredient of the estimator is the Fisher matrix, which in this work is fully evaluated from a very large set of

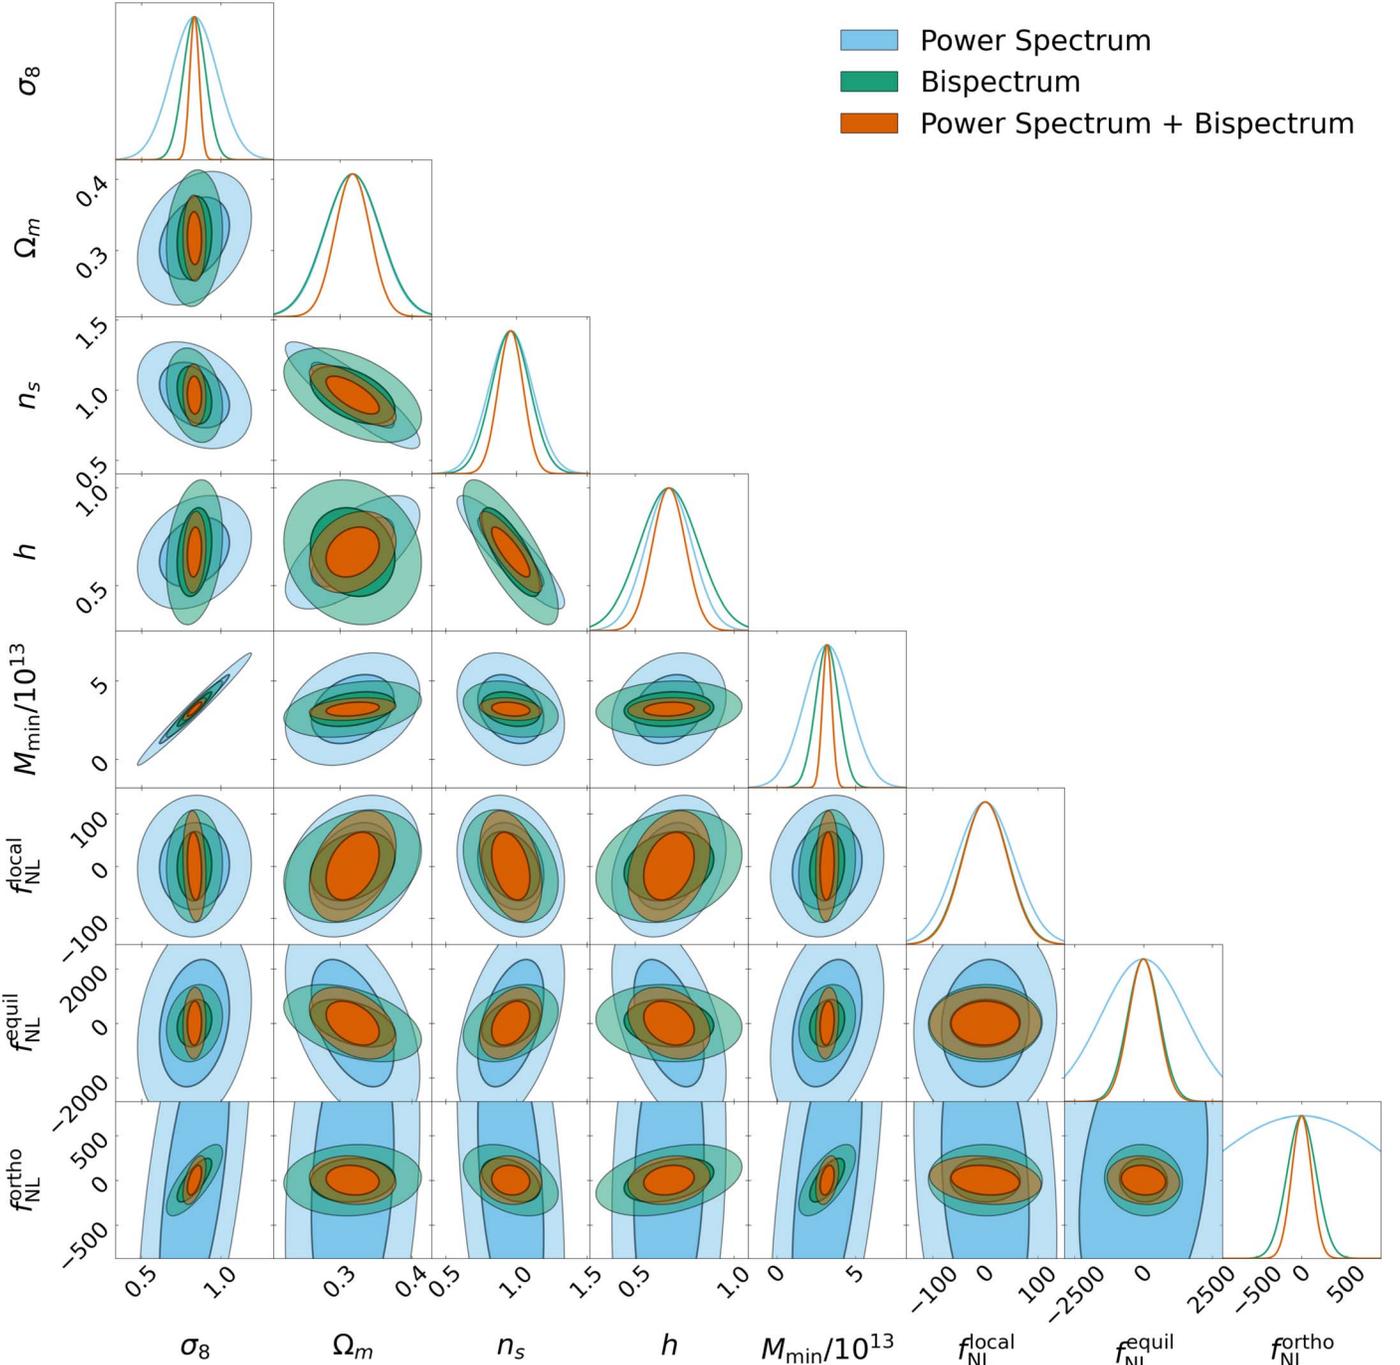


Figure 3. A comparison of the constraining power of the halo power spectrum and bispectrum at $z = 1$ for $k_{\max} 0.5 h \text{ Mpc}^{-1}$ on cosmological parameters and PNG amplitudes from the power spectrum and the modal bispectrum.

simulations. As discussed in the previous section, we use a two-step conservative approach for its computation leading to slightly suboptimal results, because numerical convergence is difficult to reach with the standard method. We verify that this leads nonetheless to unbiased and near-optimal measurements of parameters, by estimating jointly σ_8 , Ω_m , n_s , h , $f_{\text{NL}}^{\text{local}}$, $f_{\text{NL}}^{\text{equil}}$, and $f_{\text{NL}}^{\text{ortho}}$ in QUIJOTE simulations using both the power spectrum and the bispectrum.

In Figure 4, we study the effect of varying k_{\max} on the 1σ error bars of the quasi-maximum likelihood estimator. To compute these error bars, we use a set of 1000 simulations at fiducial cosmology and analyze it with the estimator calibrated using all

other simulations (using 14,000 simulations instead of 15,000 to calculate the covariance has been verified to have no impact on the results). We repeat the procedure for different sets of 1000 simulations, and compute the standard deviation of the results. As expected, this highlights a very similar behavior to that of the Fisher constraints discussed in the previous section. Concerning PNG, there is no improvement for $f_{\text{NL}}^{\text{local}}$ above $k_{\max} = 0.2 h \text{ Mpc}^{-1}$, while for the other two shapes there is no clear saturation yet (although the gain between $k_{\max} = 0.4$ and $0.5 h \text{ Mpc}^{-1}$ is only a few percent). For every other parameter considered (except h), the decrease of error bars is significant up to $k_{\max} = 0.5 h \text{ Mpc}^{-1}$. In Table 2, we report the corresponding

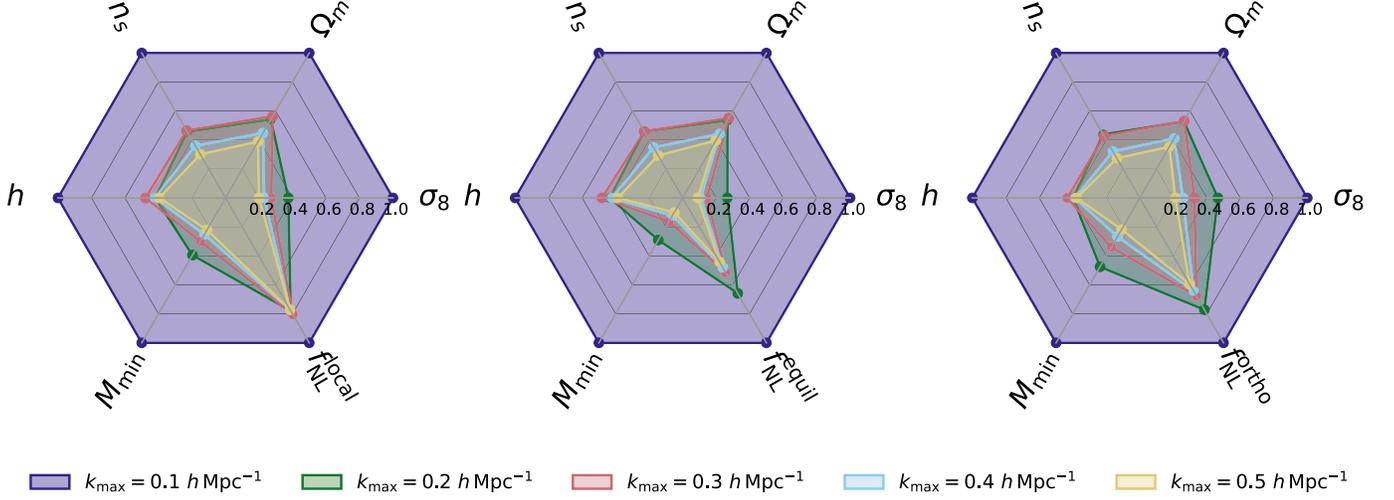


Figure 4. Same as Figure 2, but showing the 1σ error bars of the quasi-maximum likelihood estimator instead of the Fisher constraints.

Table 2

Joint 1σ Error Bars on Cosmological Parameters and PNG from the Power Spectrum and the Modal Bispectrum of the Halo Field at $z = 1$, at $k_{\max} = 0.5 h \text{ Mpc}^{-1}$

	σ_8	Ω_m	n_s	h	$M_{\min}/10^{13}$	$f_{\text{NL}}^{\text{local}}$	$f_{\text{NL}}^{\text{equil}}$	$f_{\text{NL}}^{\text{ortho}}$
Fiducial	0.834	0.3175	0.9624	0.6711	3.2	0	0	0
Fisher	± 0.024	± 0.021	± 0.081	± 0.078	± 0.26
	± 0.025	± 0.022	± 0.086	± 0.081	± 0.27	± 42
	± 0.025	± 0.023	± 0.086	± 0.080	± 0.27	...	± 530	...
	± 0.029	± 0.021	± 0.080	± 0.078	± 0.29	± 110
	± 0.030	± 0.025	± 0.091	± 0.085	± 0.29	± 43	± 530	± 110
Estimator	± 0.025	± 0.022	± 0.089	± 0.084	± 0.27
	± 0.025	± 0.023	± 0.098	± 0.091	± 0.27	± 43
	± 0.026	± 0.024	± 0.092	± 0.084	± 0.28	...	± 570	...
	± 0.030	± 0.021	± 0.085	± 0.082	± 0.30	± 110
	± 0.031	± 0.025	± 0.094	± 0.087	± 0.31	± 45	± 570	± 110

Notes. In the first part, we report the Fisher constraints described in Section 4.2 and in the second part the corresponding error bars of the quasi-maximum likelihood estimator used in Section 4.3. We analyze 15,000 QUIOTE halo catalogs of $1 (\text{Gpc } h^{-1})^3$ volume at fiducial cosmology, and sets of 500 simulations with one adjusted parameter.

error bars at $k_{\max} = 0.5 h \text{ Mpc}^{-1}$, considering cosmological parameters only or jointly with the PNG shapes. For all parameters, the error bars of the quasi-maximum likelihood estimator are close to, or slightly larger than, the Fisher constraints reported in the same table (less than 10% difference).

In Figure 5, we compare the estimated parameters to the input values for different cases, focusing here on changes of the PNG amplitudes. We first study the mildly nonlinear regime ($k_{\max} = 0.2 h \text{ Mpc}^{-1}$) and then include also nonlinear scales ($k_{\max} = 0.5 h \text{ Mpc}^{-1}$). The measured parameters match their expected values for both ranges of scales with data sets at fiducial cosmology or with PNG of the equilateral or orthogonal type ($f_{\text{NL}}^{\text{equil}} = +100$ or $f_{\text{NL}}^{\text{ortho}} = +100$).

There are however large statistical deviations on several parameters for the simulations with a local NG shape (in particular, for $f_{\text{NL}}^{\text{local}}$, several data sets give a value more than 5σ away from the expected one). This difference in behavior between this specific set and the others can be explained by the fact that $f_{\text{NL}}^{\text{local}} = 100$ is more than 2σ away from the fiducial value of $f_{\text{NL}} = 0$ (based on the error bars given in Table 2) while $f_{\text{NL}}^{\text{ortho}} = 100$ and $f_{\text{NL}}^{\text{equil}} = 100$ have, respectively, an only slightly smaller and a few times smaller than 1σ deviation

from $f_{\text{NL}} = 0$. The NG simulations of the three shapes correspond to different regimes where a parameter is more or less displaced from the model we use to calibrate the estimator. This is confirmed in Figure 6, where we verify that simulations with $f_{\text{NL}}^{\text{local}} = 50$ (thus roughly a 1σ deviation) give this time the expected results.

These tests confirm the unbiasedness of the quasi-maximum likelihood estimator, with the caveat that the estimator must be calibrated relatively close to the actual parameter values. This, of course, is due to the fact that the entire method is based on a linear approximation of the likelihood around the fiducial parameters. For the same reason, however, it is clear that the issue can be immediately addressed—at the computational cost of producing new sets of simulations—by implementing a standard recursive procedure, in which the estimated parameters at the previous step generate the new fiducial model for the following step, until convergence. Note that this scenario is not bound to occur in practice, since current cosmological parameter constraints from, e.g., CMB data sets such as Planck, produce already quite narrow priors.

While it was shown in Section 4.2 that using a lower number of simulations to compute derivatives leads to more suboptimal

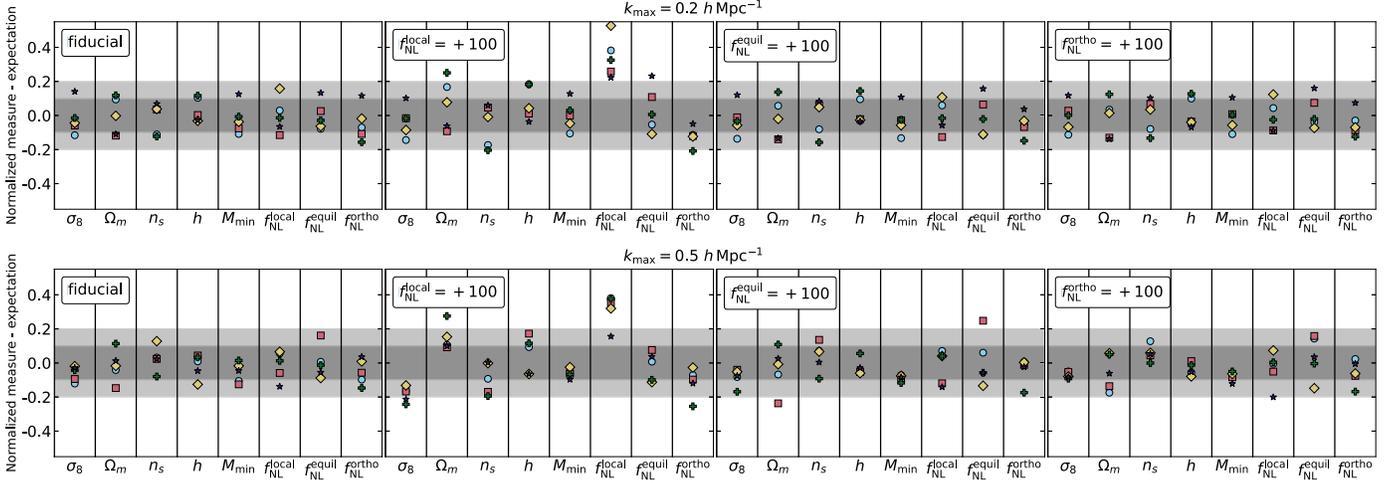


Figure 5. Relative difference of measured cosmological parameters and PNG amplitudes using the quasi-maximum likelihood estimator (Equation (3)) from their expected values. We use the power spectrum and the bispectrum of the halo field jointly, for $k_{\max} = 0.2 h \text{ Mpc}^{-1}$ in the top row and $k_{\max} = 0.5 h \text{ Mpc}^{-1}$ below. Each column corresponds to a given parameter (cosmological or PNG). Each panel corresponds to a different input cosmology of the data samples (i.e., one with the Gaussian initial conditions or either of the three types of PNG). For each input cosmology, we analyze five independent data sets of 100 realizations, each indicated by its own color and marker. The dark and light gray bands represent, respectively, the 2σ and 1σ intervals around the expected deviation (0).

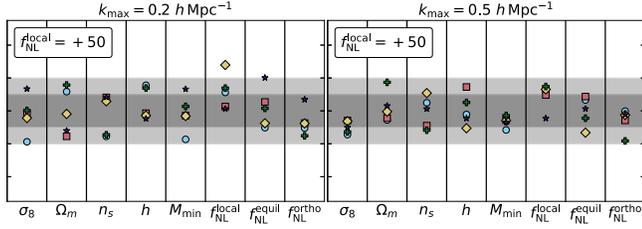


Figure 6. Same as Figure 5, but for simulations with $f_{\text{NL}}^{\text{local}} = +50$.

Fisher matrices, it is also important to verify the effect of changes in the number of simulations used to compute the covariance matrix. We explore this in Figure 7, where we show the increase of error bars due to using fewer simulations. Above 1000 simulations, the error bars of the quasi-maximum likelihood estimator are stable (variations at the percent level) and close to the Fisher estimates (10% differences at most).

5. Conclusions

In this paper, we have developed a joint power spectrum and bispectrum quasi-maximum likelihood estimator of cosmological and PNG parameters and applied it to the study of the halo field in the QUIJOTE-PNG simulation suite. The data analysis pipeline applies the optimal data compression methodology developed in Alsing & Wandelt (2018) and Heavens et al. (2000) to a set of power spectrum and modal bispectrum summary statistics, efficiently extracted from the input mock realizations. In this way, we have extended our previous analysis (Jung et al. 2022), which considered the matter field in the same data set.

The main technical complication that has arisen is related to the convergence of the numerical derivatives that are used to compute the Fisher information and to perform the final compression step. This turns out to be much slower now, with respect to the previous matter field analysis, leading to potential problems such as spurious “superoptimal” error bars in the final estimator. Interestingly, though, we have also found that our modal decomposition of the bispectrum makes derivative convergence much faster with respect to the binning approach we implemented

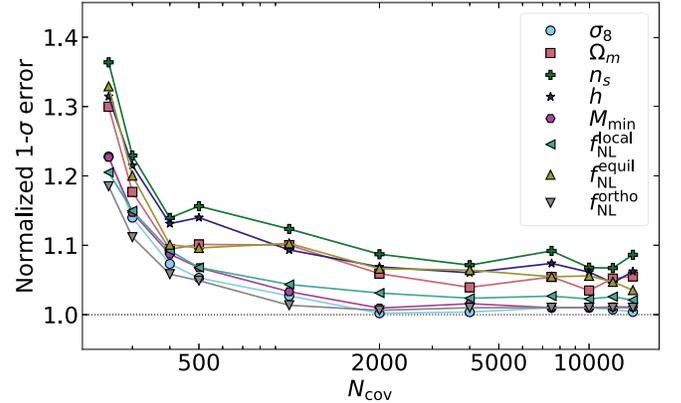


Figure 7. The impact of the number of simulations used to compute the covariance on the error bars of the quasi-maximum likelihood estimator (normalized by the Fisher constraints of Table 2). We use the quasi-maximum likelihood estimator calibrated using all the available QUIJOTE simulations at $z = 1$, including the power spectrum and bispectrum information up to $k_{\max} = 0.5 h \text{ Mpc}^{-1}$.

in Coulton et al. (2022b). Although still not sufficient for a brute-force computation with the available realizations, such faster convergence suggests that more investigation should be done in the future to find the optimal bispectrum decomposition scheme for the best numerical stability. In the meantime, to circumvent the issue, we have implemented the method first described in Coulton et al. (2022b), which is based on computing the Fisher matrix of MOPED-compressed statistics extracted from an independent simulation set. This approach leads to stable, robust results, at the price of slight suboptimality in the final estimator. Despite such (small) suboptimality, we have verified that the forecasted errors significantly improve after including nonlinear scales up to $k_{\max} = 0.5 h \text{ Mpc}^{-1}$ (see Figure 2 as a summary of our main results), in agreement with our findings in Coulton et al. (2022b). Given the significant contribution provided by small-scale, shot-noise-dominated bispectrum triangles, further improvements could in principle be achieved in a future galaxy density analysis, by selecting higher-density tracers. In contrast to other parameters,

we have observed a saturation of the $f_{\text{NL}}^{\text{local}}$ error at a scale $k \sim 0.2 h \text{ Mpc}^{-1}$; this is again consistent with our previous findings and with other forecasts, such as those in Karagiannis et al. (2018), where it was shown that the $f_{\text{NL}}^{\text{local}}$ signal is dominated by the scale-dependent bias signature on large scales, both in the power spectrum and in squeezed bispectrum configurations.

After we investigated the power spectrum and bispectrum information content on nonlinear scales, we performed the final step of our analysis, which consisted in testing our quasi-maximum likelihood estimator on the simulated data set. We have verified that we can recover unbiased results, deep into the nonlinear regime, up to $k_{\text{max}} = 0.5 h \text{ Mpc}^{-1}$ (see Figures 5 and 6). Unbiasedness is of course verified only provided the starting fiducial parameter values in the estimator are close enough to the real ones. We have studied this in more detail by varying the input value of $f_{\text{NL}}^{\text{local}}$ in the analyzed simulations and verifying that biased results are obtained when the true $f_{\text{NL}}^{\text{local}}$ in the data is $\sim 2\sigma$ away from the fiducial choice in the estimator. In a realistic observational scenario, this issue can of course always be addressed by implementing a recursive estimation procedure, which however becomes more and more expensive, by requiring new mock realizations and recalibration of the estimator weights at each step. This suggests the need to investigate the possibility of reducing the overall computational cost of simulations. We have started a preliminary analysis in this direction, using the CARPool method (Chartier et al. 2021), which is further discussed in Appendix C. Another possibility is to use machine-learning-augmented simulations—see Kaushal et al. (2022), Jamieson et al. (2022), and Piras et al. (2023) for examples. Making use of these different techniques will play a key role in enabling simulation-based inference with the upcoming generation of galaxy surveys, which will have a much higher tracer density.

The recovered error bars are, as expected, slightly larger than the optimal Fisher bound. This is a direct consequence of the fact that, to secure unbiasedness and robustness of the results, we have calibrated the estimator weights using the stable yet conservative approximation of the Fisher matrix described above. In this case though, the slight suboptimality does not prevent us from obtaining large improvements in precision for the final parameter estimates when we include nonlinear scales in the analysis (see Figure 4). By extending our previous analysis to the halo field in redshift space, we have made a significant step forward toward the final development of an efficient, joint power spectrum and bispectrum estimation pipeline, able to extract cosmological and PNG parameters at strongly nonlinear scales from actual observations. In a follow-up work we will further extend the current analysis, by looking at the galaxy density field, simulated via a suitable HOD, following Hahn & Villaescusa-Navarro (2021). Marginalization over HOD parameters will also allow us to significantly improve the accuracy of our bias model, which is currently defined by a single parameter that describes the leading-order contribution and only to a minor extent captures higher-order effects.

Our conclusions are in full agreement with those in our companion work, Coulton et al. (2022b), where we performed an independent analysis at a different redshift ($z = 0$, versus $z = 1$ in this paper) and used a standard binning scheme for the bispectrum, rather than the modal approach developed here. Besides increasing the robustness of our conclusions via cross-validation of independent data analysis pipelines, the two works complement each other in several ways and together

cover a significant range of crucial issues: Coulton et al. (2022b) focused on addressing numerical stability issues, on assessing the information content of our observables at different scales, and on evaluating in detail all possible contributions to the error budget (such as shot noise and supersample covariance effects), whereas the present study, while cross-checking the previous Fisher matrix results, is more centered on optimal data compression and on the development and testing of related statistical estimators.

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 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 support from 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 support by the ANR BIG4 project, grant ANR-16-CE23-0002, of the French Agence Nationale de la Recherche and from Labex ILP (reference No. ANR-10-LABX-63) as part of Idex SUPER, as well as financial state aid managed by the ANR, as part of the program Investissements d’avenir under reference No. ANR-11-IDEX-0004-02. The Flatiron Institute is supported by the Simons Foundation.

Appendix A Shot-noise Modal Modes

The shot-noise contribution to the matter bispectrum at the tree level is given by

$$B^{\text{SN}}(k_1, k_2, k_3) = \frac{1}{\bar{n}} [P_L(k_1) + P_L(k_2) + P_L(k_3)] + \frac{1}{\bar{n}^2}, \quad (\text{A1})$$

where \bar{n} is the halo number density and $P_L(k)$ is the linear matter power spectrum. As introduced in Byun et al. (2021), this can be fully described in the modal way by using the two triplets (0, 0, 1) and (0, 0, 0) combining the following one-dimensional basis functions:

$$q_0 = \sqrt{\frac{k}{P(k)}} \frac{5}{14}, \quad q_1 = \sqrt{\frac{k}{P(k)}} P_L(k). \quad (\text{A2})$$

Appendix B Comparison with the Standard “Binned” Bispectrum Estimator

A key ingredient to computing the Fisher matrix (Equation (4)), which is used both for constraint forecasts

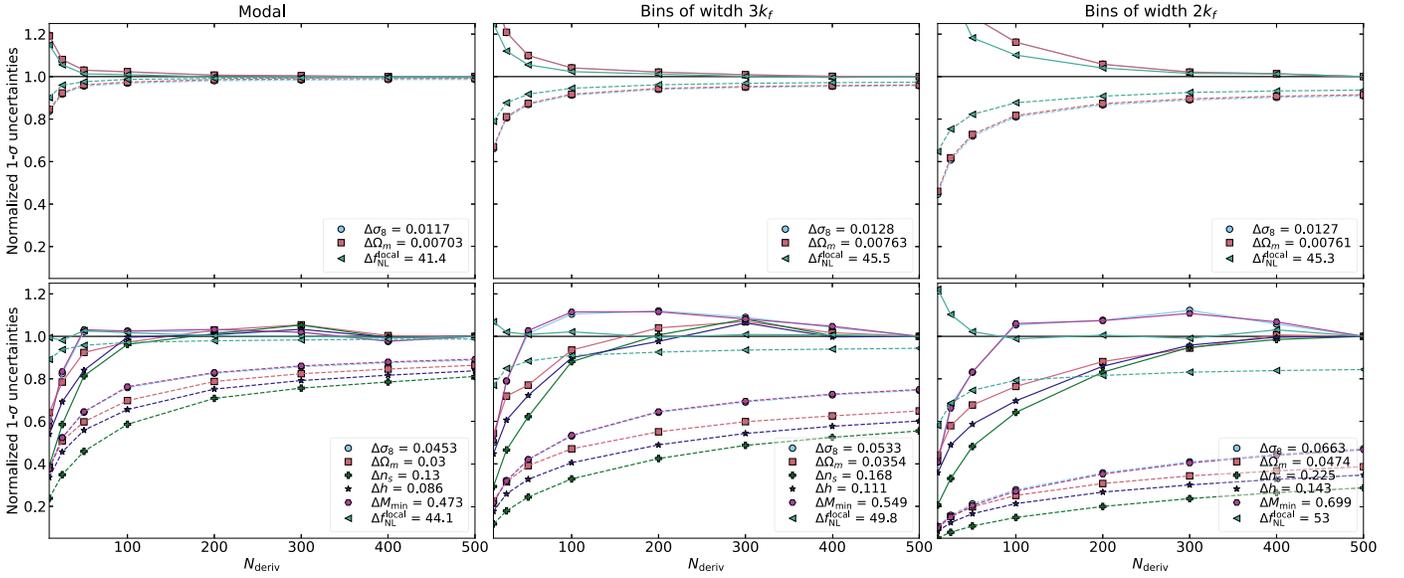


Figure 8. A comparison of 1σ uncertainties obtained with different bispectrum estimators. The leftmost column is obtained with the modal bispectrum estimator used throughout this paper, while the two others use a standard “binned” approach for different widths of bins ($3k_f$ in the middle panels and $2k_f$ in the right panels). Otherwise, this figure is similar to Figure 1, for $k_{\max} = 0.2 h \text{ Mpc}^{-1}$.

and for building estimators, is to have accurate derivatives of the summary statistics with respect to the different parameters considered. As discussed in Section 4.2, even the large sets of 500 paired simulations for each parameter of the QUIJOTE and QUIJOTE-PNG collections are not sufficient to reach the necessary numerical convergence for the power spectrum and bispectrum derivatives. This typically leads to an underestimation of 1σ error bars.

On the other hand, a two-step computation involving first optimally compressing the data and then computing the Fisher matrix from this compressed data (using different data sets for the two steps) yields slightly overestimated error bars. Combining the power spectrum and the bispectrum information of the halo field at $z=1$, we have verified that even when we include nonlinear scales up to $k_{\max} = 0.5 h \text{ Mpc}^{-1}$, the difference between the lower and upper bounds of constraints is at most of the order of 20% on the different parameters, a very reasonable difference. In this appendix, we show that the modal estimator, which by construction compresses the bispectrum information in the data, is a necessary ingredient for the efficiency of the method.

In Figure 8, we compare the convergence of the standard and conservative Fisher 1σ uncertainties obtained with the modal bispectrum (as in the rest of this paper) and the convergence of those obtained with a standard “binned” bispectrum estimator.¹⁹ A main result is that the constraints obtained with the modal estimator are the most stringent, with a difference of the order of 10% from the standard estimator with bins of width $3k_f$ and one of even higher order from that with smaller bins of width $2k_f$. Indeed the estimator using the smallest bins gives here the largest error bars, despite the fact

that in principle it should keep more information, due to the greater difficulty of computing sufficiently accurate numerical derivatives. This lack of convergence is also very clear when we compare the lower and upper bounds on error bars for all three methods. Using the full sets of simulations, the lower bounds are 10%–20% smaller than the upper limits for the modal estimator, 30% smaller for the bins of width $3k_f$, and as much as two times smaller for the bins of width $2k_f$. The modal estimator gives more stringent constraints, which are proven to be closer to the actual Fisher uncertainties, and should converge totally with a smaller number of simulations, as shown in the first row of Figure 8 in the simple situation, where the modal estimator has fully converged and the other two have not.

Appendix C Application of CARPool

As we have verified in this paper, the use of the quasi-maximum likelihood estimator is a powerful method to infer cosmological parameters and PNG amplitudes from halo catalogs using information beyond the mildly nonlinear regime, which, like other simulation-based methods, unfortunately can require a large number of costly forward simulations. Therefore, a key component of future applications will be to include the variance reduction CARPool technique, developed in Chartier et al. (2021) and Chartier & Wandelt (2021, 2022), in the full analysis pipeline.

The basic idea behind CARPool is to use a relatively small number of high-fidelity simulations combined with a large number of less accurate simulations, or surrogates, to measure some chosen summary statistics with much smaller error bars. In Chartier et al. (2021), these surrogates were computed using much faster, but less precise, N -body solvers like COLA (Tassev et al. 2013). This could for example be applied to the case of numerical derivatives, for which reaching numerical convergence typically requires thousands of costly simulations.

¹⁹ As was pointed out in Byun et al. (2021), the standard bispectrum estimator in Fourier space can be recovered in the modal fashion by using a simple basis with modes of the form $q_n(k) = \sqrt{kP(k)}\theta_i(k)$, where $\theta_i(k) = 1$ if $k \in \Delta_i$, after dividing the k range into bins Δ_i , and 0 otherwise. This is the implementation we use for this analysis.

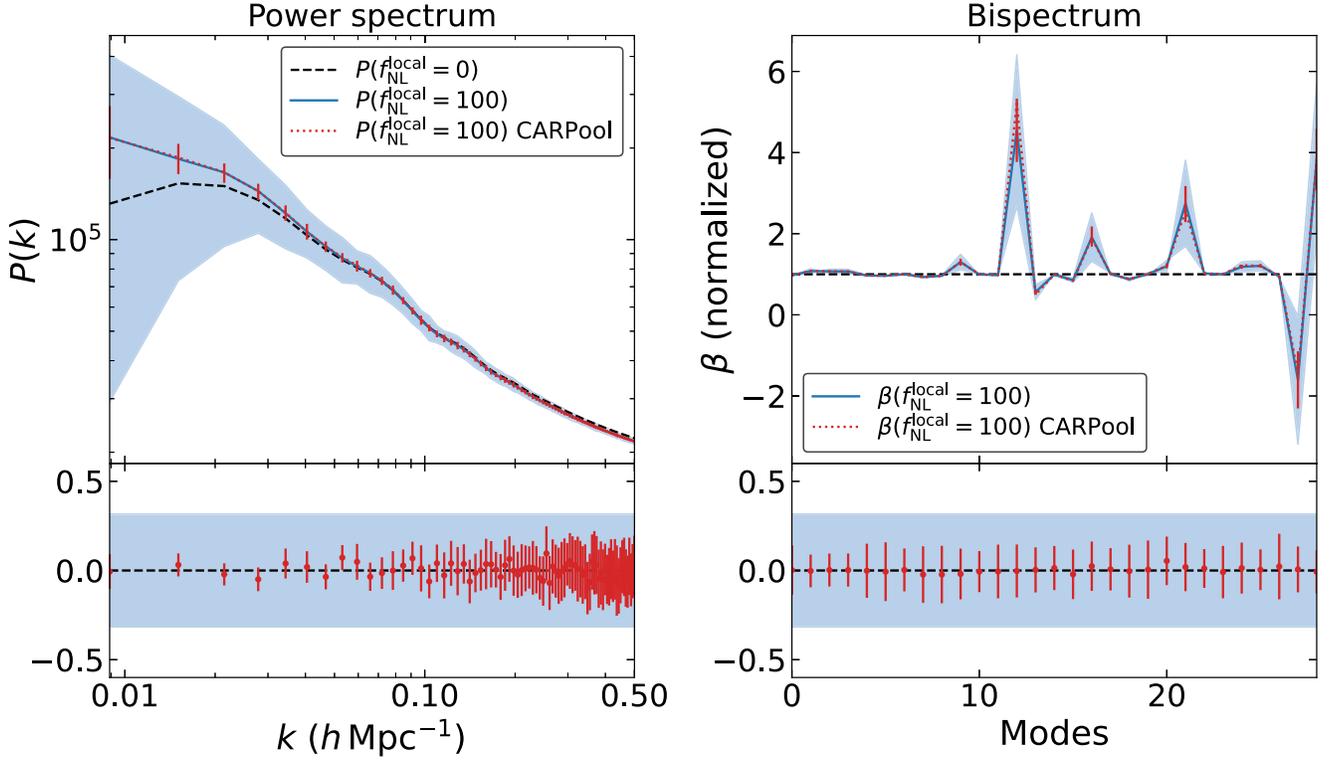


Figure 9. The CARPool method applied to the power spectrum (left column) and modal bispectrum (right column) of the halo field, at $z = 1$. In the top row, the black dashed lines correspond to the averages from the 15,000 QUIJOTE simulations at fiducial cosmology with $f_{\text{NL}}^{\text{local}} = 0$ (note that in the bispectrum case, all modal coefficients are normalized by dividing by the modes from these 15,000 simulations at fiducial cosmology). The blue lines correspond to the average from 500 simulations with $f_{\text{NL}}^{\text{local}} = +100$. The red dotted lines have been computed using the CARPool method (see Equation (C3)), using 10 simulations at $f_{\text{NL}}^{\text{local}} = +100$ as the high-fidelity simulations and the 15,000 simulations at fiducial cosmology as the surrogates. The blue areas and red vertical lines show the respective error bars from the two cases (they correspond to the standard errors for the sets of 10 simulations, and have been multiplied by a factor of 10 for visibility in the power spectrum case). In the bottom row, we show the difference between the CARPool estimates and the averages from 500 simulations, normalized by the standard deviation. The blue areas correspond to the standard error for 10 simulations and the error bars. Error bars on the CARPool estimates, shown in red, are computed by applying the CARPool method to many different sets of 10 simulations at $f_{\text{NL}}^{\text{local}} = +100$.

We leave this application for future work, and instead focus here on the use of CARPool to speed up the iteration process of quasi-maximum likelihood estimation.

To obtain unbiased estimates of cosmological parameters or f_{NL} , it is important that the fiducial cosmology where we evaluate the covariance and numerical derivatives is not too far from the actual parameter values. For example, in Section 4, we have seen that with a fiducial cosmology at $f_{\text{NL}}^{\text{local}} = 0$, the estimator is unbiased in measuring $f_{\text{NL}}^{\text{local}}$ in simulations with an input of $f_{\text{NL}}^{\text{local}} = 50$, but not with an input of $f_{\text{NL}}^{\text{local}} = 100$. Note that even if the measured bias were large when averaging from hundreds of simulations, it would still be smaller than the 1σ error bar, making it a good first estimate. Then, working by iteration and choosing a new fiducial cosmology at these roughly measured parameters should yield unbiased results. To avoid producing a completely new large set of simulations at the new fiducial cosmology, we can consider the original simulations as surrogates of the CARPool method (the idea of combining simulations at different cosmologies was also explored in Ding et al. 2022).

The main ingredients of the CARPool method are as follows:

1. A set of N paired high-fidelity simulations and surrogates sharing the same random seeds to produce their initial conditions, from which we measure some chosen summary statistic denoted by \mathbf{y} or \mathbf{c} (simulation or surrogate, respectively) and the corresponding sample

covariance given by

$$\hat{\Sigma}_{\mathbf{y}\mathbf{c}} = \frac{1}{N-1} \sum_{i=1}^N (\mathbf{y}_i - \bar{\mathbf{y}})(\mathbf{c}_i - \bar{\mathbf{c}})^T,$$

$$\hat{\Sigma}_{\mathbf{c}\mathbf{c}} = \frac{1}{N-1} \sum_{i=1}^N (\mathbf{c}_i - \bar{\mathbf{c}})(\mathbf{c}_i - \bar{\mathbf{c}})^T, \quad \bar{\mathbf{y}} = \frac{1}{N} \sum_{i=1}^N \mathbf{y}_i. \quad (\text{C1})$$

2. A separate set of M surrogates, for computing the mean of \mathbf{c} with the standard expression

$$\boldsymbol{\mu}_{\mathbf{c}} = \frac{1}{M} \sum_{i=1}^M \mathbf{c}_i. \quad (\text{C2})$$

Then, the key quantity to compute is

$$\mathbf{x} = \mathbf{y} - \hat{\beta}(\mathbf{c} - \boldsymbol{\mu}_{\mathbf{c}}), \quad (\text{C3})$$

which by construction has the same ensemble average as \mathbf{y} (i.e., $\bar{\mathbf{x}} = \bar{\mathbf{y}}$). The variance of \mathbf{x} is minimized when the control matrix $\hat{\beta}$ is given by

$$\hat{\beta} = \Sigma_{\mathbf{y}\mathbf{c}} \Sigma_{\mathbf{c}\mathbf{c}}^{-1}. \quad (\text{C4})$$

In Chartier et al. (2021), it was shown that a very efficient choice, using only the diagonal elements of $\Sigma_{\mathbf{y}\mathbf{c}}$ and $\Sigma_{\mathbf{c}\mathbf{c}}$, is the

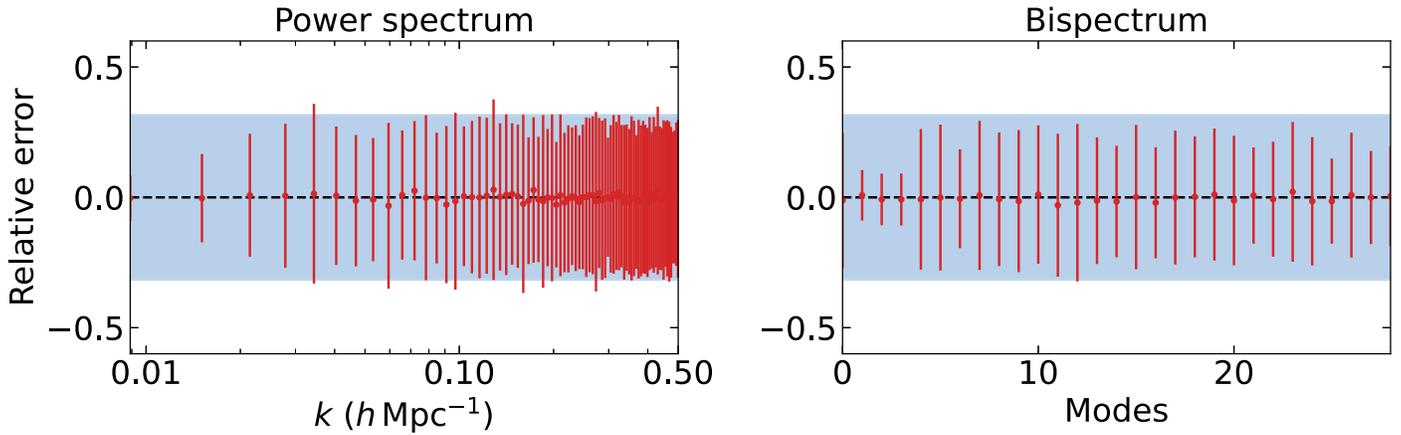


Figure 10. Similar to the bottom row of Figure 9, where the CARPool method is applied to compute the derivatives $\partial P(k)/\partial f_{\text{NL}}^{\text{local}}$ and $\partial \beta/\partial f_{\text{NL}}^{\text{local}}$ at the point $f_{\text{NL}}^{\text{local}} = +50$, using the previously computed derivatives at $f_{\text{NL}}^{\text{local}} = 0$ as surrogates.

following diagonal control matrix:

$$\beta^{\text{diag}} = \text{diag}\left(\frac{\text{cov}(y_1, c_1)}{\sigma(c_1)^2}, \frac{\text{cov}(y_2, c_2)}{\sigma(c_2)^2}, \dots, \frac{\text{cov}(y_n, c_n)}{\sigma(c_n)^2}\right), \quad (\text{C5})$$

where n is the size of the vectors \mathbf{y} and \mathbf{c} .

In Figure 9, we show the results obtained with the CARPool technique applied to the QUIJOTE-PNG set of halo catalogs. We use the large set of 15,000 QUIJOTE simulations at fiducial cosmology and with no PNG in their initial conditions as the surrogates, and a small set of 10 NG simulations ($f_{\text{NL}}^{\text{local}} = +100$) with the same Λ CDM cosmological parameters as the high-fidelity simulations, the goal being to predict the power spectrum and bispectrum more accurately outside of the fiducial cosmology. We compare the CARPool results to those of the 500 simulations with $f_{\text{NL}}^{\text{local}} = +100$ at our disposal and verify that they are indeed unbiased, as expected. We repeat the procedure using many subsets of 10 simulations among the 500 to check that the result is not spurious, and to derive error bars on the CARPool averages. For all power spectrum and bispectrum modes, the error bars are significantly smaller than the standard errors on the average from the 10 simulations alone. The effect is strongest on linear scales (small k for the power spectrum, and the first few bispectrum modes that describe the tree-level matter bispectrum), but is also present in the nonlinear regime.

In Figure 10, we follow a similar procedure to study the derivatives of the power spectrum and bispectrum with respect to $f_{\text{NL}}^{\text{local}}$. We use the derivatives evaluated at $f_{\text{NL}}^{\text{local}} = 0$ by finite difference applied to the 500 $f_{\text{NL}}^{\text{local}} = \pm 100$ simulations as surrogates, to compute the derivatives at $f_{\text{NL}}^{\text{local}} = 50$ using only a few simulations with $f_{\text{NL}}^{\text{local}} = 0$ or 100. For the power spectrum the improvement is small, even negligible in some cases, outside of the largest scales. For the bispectrum, there is a significant improvement of the first few modes describing the tree-level matter bispectrum. All error bars are reduced by the CARPool method, although the improvement is very small for some modal coefficients. One issue here is the small number of surrogates compared to that of the previous application (only 500 instead of 15,000), adding to the fact that we know that the surrogate derivatives are not even fully converged numerically (as discussed thoroughly in Section 4).

These examples illustrate briefly the possibilities of the CARPool technique. We leave its full implementation in the pipeline for future work, where we will also include the powerful ‘‘CARPool Bayes’’ technique introduced in Chartier & Wandelt (2022) for the fast and accurate estimation of the covariance matrix and its inverse.

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., & Wandelt, B. 2018, *MNRAS*, **476**, L60
 Andrews, A., Jasche, J., Lavaux, G., & Schmidt, F. 2023, *MNRAS*, **520**, 5746
 Barreira, A. 2020, *JCAP*, **12**, 031
 Barreira, A. 2022a, *JCAP*, **01**, 033
 Barreira, A. 2022b, *JCAP*, **11**, 013
 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*, **03**, 105
 Cabass, G., Ivanov, M. M., Philcox, O. H. E., Simonović, M., & Zaldarriaga, M. 2022a, *PhRvL*, **129**, 021301
 Cabass, G., Ivanov, M. M., Philcox, O. H. E., Simonović, M., & Zaldarriaga, M. 2022b, *PhRvD*, **106**, 043506
 Castorina, E., Feng, Y., Seljak, U., & Villaescusa-Navarro, F. 2018, *PhRvL*, **121**, 101301
 Chan, K. C., Hamaus, N., & Biagetti, M. 2019, *PhRvD*, **99**, 121304
 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. 2023a, *ApJ*, **943**, 64
 Coulton, W. R., Villaescusa-Navarro, F., Jamieson, D., et al. 2023b, *ApJ*, **943**, 178

- Crocce, M., Pueblas, S., & Scoccimarro, R. 2006, *MNRAS*, **373**, 369
- Dalal, N., Doré, O., Huterer, D., & Shirokov, A. 2008, *PhRvD*, **77**, 123514
- D'Amico, G., Lewandowski, M., Senatore, L., & Zhang, P. 2022, arXiv:2201.11518
- Davis, M., Efstathiou, G., Frenk, C. S., & White, S. D. M. 1985, *ApJ*, **292**, 371
- de Putter, R. 2018, arXiv:1802.06762
- DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016, arXiv:1611.00036
- Desjacques, V., Jeong, D., & Schmidt, F. 2018, *PhR*, **733**, 1
- Desjacques, V., & Seljak, U. 2010, *AdAst*, **2010**, 908640
- Ding, Z., Chuang, C., Yu, Y., et al. 2022, *MNRAS*, **514**, 3308
- Doré, O., Bock, J., Ashby, M., et al. 2014, arXiv:1412.4872
- 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
- Giri, U., Münchmeyer, M., & Smith, K. M. 2023, *PhRvD*, **107**, L061301
- Hahn, C., & Villaescusa-Navarro, F. 2021, *JCAP*, **2021**, 029
- Hahn, C., Villaescusa-Navarro, F., Castorina, E., & Scoccimarro, R. 2020, *JCAP*, **03**, 040
- Heavens, A., 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
- Jamieson, D., Li, Y., de Oliveira, R. A., et al. 2022, arXiv:2206.04594
- Jung, G., Karagiannis, D., Liguori, M., et al. 2022, *ApJ*, **940**, 71
- Karagiannis, D., Lazanu, A., Liguori, M., et al. 2018, *MNRAS*, **478**, 1341
- Kaushal, N., Villaescusa-Navarro, F., Giusarma, E., et al. 2022, *ApJ*, **930**, 115
- 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
- Leistedt, B., Peiris, H. V., & Roth, N. 2014, *PhRvL*, **113**, 221301
- Lewis, A., Challinor, A., & Lasenby, A. 2000, *ApJ*, **538**, 473
- LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009, arXiv:0912.0201
- Matarrese, S., & Verde, L. 2008, *ApJL*, **677**, L77
- Mueller, E.-M., Rezaie, M., Percival, W. J., et al. 2021, arXiv:2106.13725
- Piras, D., Joachimi, B., & Villaescusa-Navarro, F. 2023, *MNRAS*, **520**, 668
- Regan, D. M., Schmittfull, M. M., Shellard, E. P. S., & Fergusson, J. R. 2012, *PhRvD*, **86**, 123524
- Ross, A. J., Percival, W. J., Carnero, A., et al. 2013, *MNRAS*, **428**, 1116
- 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
- Seljak, U. 2009, *PhRvL*, **102**, 021302
- Slosar, A., Hirata, C., Seljak, U., Ho, S., & Padmanabhan, N. 2008, *JCAP*, **2008**, 031
- Springel, V. 2005, *MNRAS*, **364**, 1105
- Tassev, S., Zaldarriaga, M., & Eisenstein, D. 2013, *JCAP*, **06**, 036
- Valogiannis, G., & Dvorkin, C. 2022, *PhRvD*, **105**, 103534
- Villaescusa-Navarro, F. 2018, Pylians: Python libraries for the analysis of numerical simulations, Astrophysics Source Code Library, ascl:1811.008
- Villaescusa-Navarro, F., Hahn, C., Massara, E., et al. 2020, *ApJS*, **250**, 2