










Quijote-PNG: The Information Content of the Halo Mass Function

Gabriel Jung¹ , Andrea Ravenni^{2,3}, Marco Baldi^{4,5,6} , William R Coulton⁷, Drew Jamieson⁸ , Dionysios Karagiannis⁹, Michele Liguori^{2,3,10}, Helen Shao¹¹ , Licia Verde^{12,13} , Francisco Villaescusa-Navarro^{7,11} , and Benjamin D. Wandelt^{7,14} 

¹ Université Paris-Saclay, CNRS, Institut d’Astrophysique Spatiale, F-91405 Orsay France

² Dipartimento di Fisica e Astronomia “G. Galilei,” Università degli Studi di Padova, via Marzolo 8, I-35131 Padova, Italy

³ INFN, Sezione di Padova, via Marzolo 8, I-35131 Padova, Italy

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

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

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

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

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

⁹ Department of Physics and Astronomy, University of the Western Cape, Cape Town 7535, South Africa

¹⁰ Dipartimento di Fisica, Università degli Studi di Trento, Via Sommarive 14, I-38123 Povo (TN), Italy

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

¹² ICREA, Pg. Lluís Companys 23, Barcelona, E-08010, Spain

¹³ ICC, University of Barcelona, IEEC-UB, Martí i Franquès, 1, E-08028 Barcelona, Spain

¹⁴ Sorbonne Université, CNRS, UMR 7095, Institut d’Astrophysique de Paris, 98 bis bd Arago, F-75014 Paris, France

Received 2023 June 1; revised 2023 September 22; accepted 2023 September 27; published 2023 October 27

Abstract

We study signatures of primordial non-Gaussianity (PNG) in the redshift-space halo field on nonlinear scales using a combination of three summary statistics, namely, the halo mass function (HMF), power spectrum, and bispectrum. The choice of adding the HMF to our previous joint analysis of the power spectrum and bispectrum is driven by a preliminary field-level analysis, in which we train graph neural networks on halo catalogs to infer the PNG f_{NL} parameter. The covariance matrix and the responses of our summaries to changes in model parameters are extracted from a suite of halo catalogs constructed from the QUIJOTE-PNG N -body simulations. We consider the three main types of PNG: local, equilateral, and orthogonal. Adding the HMF to our previous joint analysis of the power spectrum and bispectrum produces two main effects. First, it reduces the equilateral f_{NL} predicted errors by roughly a factor of 2 while also producing notable, although smaller, improvements for orthogonal PNG. Second, it helps break the degeneracy between the local PNG amplitude, $f_{\text{NL}}^{\text{local}}$, and assembly bias, b_{ϕ} , without relying on any external prior assumption. Our final forecasts for the PNG parameters are $\Delta f_{\text{NL}}^{\text{local}} = 40$, $\Delta f_{\text{NL}}^{\text{equil}} = 200$, $\Delta f_{\text{NL}}^{\text{ortho}} = 85$, on a cubic volume of $1(\text{Gpc}/h)^3$, with a halo number density of $\bar{n} \sim 5.1 \times 10^{-5} h^3 \text{Mpc}^{-3}$, at $z = 1$, and considering scales up to $k_{\text{max}} = 0.5 h \text{Mpc}^{-1}$.

Unified Astronomy Thesaurus concepts: [Large-scale structure of the universe \(902\)](#); [Non-Gaussianity \(1116\)](#); [Fisher’s Information \(1922\)](#)

1. Introduction

The presence of a certain degree of non-Gaussianity (NG) in the primordial cosmological perturbation field is a general prediction of both inflationary and other early Universe scenarios. In addition, both the level of the predicted NG signal and the shape of the expected NG signatures are significantly model-dependent. This makes primordial non-Gaussianity (PNG) a powerful tool to constrain inflation or alternative primordial models and provide clues about physics at very high energy scales.

From an observational point of view, the challenging aspect of any PNG analysis is that the expected NG signatures are very small, and the optimal statistic that maximizes their signal-to-noise ratio is unknown from low-redshift observables. Indeed, to date, there has been no experimental detection of a PNG signal, although significant constraints have been placed using cosmic microwave background (CMB) data; the CMB is an ideal observable for PNG studies, since it formed at early

times, when cosmological perturbations were still in the linear regime, hence preserving the statistical features of the primordial fluctuation field. The most precise results currently come from the analysis of Planck CMB data, which produced an upper bound on the level of PNG at roughly less than 0.1% of the amplitude of the Gaussian component of the field (Akrami et al. 2020).

The open question is whether and how we can obtain more stringent PNG constraints—or achieve a detection—with future cosmological observations. In this respect, it is known that, after Planck, CMB data have nearly saturated their PNG constraining power, with possible improvements of, at most, a factor of ~ 2 for relevant parameters in a majority of scenarios (Finelli et al. 2018; Abazajian et al. 2019). It is therefore necessary to explore different observables. Galaxy clustering is a natural candidate for two main reasons. First of all, in the limit of weak PNG, the bispectrum (i.e., the three-point function of the Fourier/harmonic modes) of primordial cosmological perturbations contains most of the non-Gaussian information, and the three-dimensional galaxy density field contains more bispectrum modes for NG analysis than the two-dimensional CMB map. Furthermore, some models—notably, those producing a “local-type” bispectrum, where the signal

peaks on squeezed Fourier mode triangles—generate a characteristic scale-dependent signature in the galaxy power spectrum on very large scales (Dalal et al. 2008; Matarrese & Verde 2008; McDonald 2008; Slosar et al. 2008; Desjacques & Seljak 2010a; Giannantonio & Porciani 2010), which can be used to constrain NG.

In both cases, however, there are some important complications to consider. As far as bispectrum analysis is concerned, the big caveat is that the additional modes in the large-scale structure (LSS) bispectrum are in the nonlinear regime. Hence, they present a “late-time” component generated by the nonlinear gravitational evolution of structures, which is hard to disentangle and much larger than the primordial one. Of course, this late-time three-point signal is interesting in itself, since it carries a lot of information about cosmological parameters and structure evolution (Hahn et al. 2020; Hahn & Villaescusa-Navarro 2021); however, as long as we are focused on PNG, it is a massive source of contamination, with an amplitude ~ 1000 times larger than the primordial signal of interest. The scale-dependent power spectrum signature on large scales clearly does not present this problem and was considered for a long time to be a cleaner LSS probe of PNG, although limited to a subset of all possible PNG scenarios. However, a significant issue has recently been pointed out again in this area (Reid et al. 2010; Barreira 2020, 2022), namely, the degeneracy produced by the breaking of the universality relation that was generally used to link the NG galaxy bias parameter b_ϕ to the linear bias parameter b_1 . This is due to halo/galaxy assembly bias effects, and, if not addressed in any way, it allows us only to constrain the $b_\phi f_{\text{NL}}$ combination.

A key objective in cosmological PNG studies is thus developing optimal data analysis strategies to overcome, at least partially, the aforementioned issues. As far as the $b_\phi(b_1)$ relation is concerned, an active effort is being put into characterizing it as well as possible via numerical studies of N -body simulations (Barreira 2020, 2022; Lazeyras et al. 2023; Sullivan et al. 2023) in order to produce accurate priors. Another logical line of attack, which we start exploring in this work, is that of going beyond a power spectrum + bispectrum analysis and including extra summary statistics, which could help disentangle the PNG signal from late-time evolution effects. The open question, with no straightforward answer, is, of course, which summary statistics are best suited to this purpose? In this paper, we explore the halo mass function (HMF) as an interesting candidate. This choice was not casual but was driven by training graph neural networks (GNNs) to perform field-level likelihood-free inference on halo catalogs from QUIJOTE-PNG simulations. The analysis of the outcome of those calculations led us to the conclusion that the model was extracting information from the abundance of halos, as we explain in Section 3.1. Therefore, the HMF can be seen as a machine learning–driven statistic that stands ahead of others.

Furthermore, our choice is also justified at a theoretical level, since the HMF has been known for a long time to be sensitive to non-Gaussian initial conditions (ICs), which are able to skew its distribution by changing the abundance of massive halos, and it was proposed as an interesting complementary PNG probe to the bispectrum in a number of papers (Matarrese et al. 2000; Sefusatti et al. 2007; Grossi et al. 2007, 2009; Pillepich et al. 2010; Desjacques et al. 2009; Desjacques & Seljak 2010b; LoVerde & Smith 2011; Palma et al. 2020).

On top of this, a major advantage of the HMF is that it directly depends on the PNG amplitude parameter f_{NL} . Therefore, it does not exhibit the b_ϕ - f_{NL} degeneracy that affects the scale-dependent power spectrum signature.

This work belongs to the QUIJOTE-PNG series (Coulton et al. 2023a, 2023b; Jung et al. 2022a, 2022b), where we aim to build a simulation-based pipeline to optimally extract NG information, pushing our analysis to smaller, nonlinear scales. This kind of approach is complementary to a perturbation theory–based likelihood analysis of the power spectrum and bispectrum (Moradinezhad Dizgah et al. 2021; Cabass et al. 2022a, 2022b; D’Amico et al. 2022). See also Giri et al. (2023) for an alternative simulation-based approach that uses large-scale modulation of small-scale power.

The paper is structured as follows. In Section 2, we briefly describe the simulation data set used in our analysis; in Section 3.1, we describe our preliminary field-level analysis; in Section 3.2, we recall and summarize the main methodological aspects of our data analysis pipeline to extract relevant summary statistics and compute the corresponding Fisher matrix; Section 3.3 is devoted to a specific discussion of the HMF, the main new ingredient with respect to our previous analyses, and how we extract it from simulations; our numerical Fisher forecasts are described in Section 4, where we also discuss the improvements coming from complementing the initial power spectrum + bispectrum analysis with HMF estimates; finally, we draw our conclusions in Section 5.

2. Simulations

In this work, we use the publicly available halo catalogs derived from the QUIJOTE suite of N -body simulations (Villaescusa-Navarro et al. 2020).¹⁵ These simulations have been produced using the codes 2LPTIC (Crocce et al. 2006) and 2LPTPNG (Scoccimarro et al. 2012; Coulton et al. 2023a)¹⁶ to generate ICs at $z = 127$, GADGET-III (Springel 2005) to follow their evolution up to $z = 0$, and the friends-of-friends (FOF) algorithm to identify the halos in each simulation (Davis et al. 1985).

We report the cosmological parameters of these simulations in Table 1. As described in Section 3.2, we use 15,000 simulations of the fiducial cosmology to evaluate covariance matrices and paired sets of 500 catalogs where one parameter is displaced by a small step from its fiducial value to compute derivatives with respect to all parameters considered in the analyses. As in Coulton et al. (2023b) and Jung et al. (2022b), we focus on the cosmological parameters $\{\sigma_8, \Omega_m, n_s, h\}$ ¹⁷ and PNG amplitudes $\{f_{\text{NL}}^{\text{local}}, f_{\text{NL}}^{\text{equil}}, f_{\text{NL}}^{\text{ortho}}\}$, including a simplified bias parameter M_{min} (the minimum mass of the halos included in the analysis). To ensure that the IC generation method has not generated unphysical higher-order N -point functions, which could impact the results presented here, we performed further validation of the ICs by examining the primordial trispectrum. As discussed in Appendix A, we find no evidence of large, unphysical trispectra in the ICs.

We focus our analyses on redshift $z = 1$, for which all power spectra and (modal) bispectra have been computed in

¹⁵ <https://quijote-simulations.readthedocs.io>

¹⁶ <https://github.com/dsjamieson/2LPTPNG>

¹⁷ We do not include Ω_b in the analyses presented here, as it is the parameter that is most affected by the numerical convergence issue mentioned in Section 3.2, and it does not significantly impact the results. Moreover, Ω_b is better constrained by CMB observations.

Table 1
The Parameters of the QUIJOTE and QUIJOTE-PNG Halo Catalogs Used in This Work

	N_{sims}	σ_8	Ω_m	Ω_b	n_s	h	$f_{\text{NL}}^{\text{local}}$	$f_{\text{NL}}^{\text{equil}}$	$f_{\text{NL}}^{\text{ortho}}$	$M_{\text{min}}(M_{\odot}/h)$
Fiducial	15,000	0.834	0.3175	0.049	0.9624	0.6711	0	0	0	3.2×10^{13}
σ_8^+	500	0.849	0.3175	0.049	0.9624	0.6711	0	0	0	3.2×10^{13}
σ_8^-	500	0.819	0.3175	0.049	0.9624	0.6711	0	0	0	3.2×10^{13}
Ω_m^+	500	0.834	0.3275	0.049	0.9624	0.6711	0	0	0	3.2×10^{13}
Ω_m^-	500	0.834	0.3075	0.049	0.9624	0.6711	0	0	0	3.2×10^{13}
n_s^+	500	0.834	0.3175	0.049	0.9824	0.6711	0	0	0	3.2×10^{13}
n_s^-	500	0.834	0.3175	0.049	0.9424	0.6711	0	0	0	3.2×10^{13}
h^+	500	0.834	0.3175	0.049	0.9624	0.6911	0	0	0	3.2×10^{13}
h^-	500	0.834	0.3175	0.049	0.9624	0.6511	0	0	0	3.2×10^{13}
$f_{\text{NL}}^{\text{local},+}$	500	0.834	0.3175	0.049	0.9624	0.6711	+100	0	0	3.2×10^{13}
$f_{\text{NL}}^{\text{local},-}$	500	0.834	0.3175	0.049	0.9624	0.6711	-100	0	0	3.2×10^{13}
$f_{\text{NL}}^{\text{local},+}$	500	0.834	0.3175	0.049	0.9624	0.6711	0	+100	0	3.2×10^{13}
$f_{\text{NL}}^{\text{local},-}$	500	0.834	0.3175	0.049	0.9624	0.6711	0	-100	0	3.2×10^{13}
$f_{\text{NL}}^{\text{ortho},+}$	500	0.834	0.3175	0.049	0.9624	0.6711	0	0	+100	3.2×10^{13}
$f_{\text{NL}}^{\text{ortho},-}$	500	0.834	0.3175	0.049	0.9624	0.6711	0	0	-100	3.2×10^{13}
M_{min}^+	500	0.834	0.3175	0.049	0.9624	0.6711	0	0	0	3.3×10^{13}
M_{min}^-	500	0.834	0.3175	0.049	0.9624	0.6711	0	0	0	3.1×10^{13}

Jung et al. (2022b). Results at lower redshifts, $z = 0.5$ and zero, are also shown in Appendix B.

3. Methods

3.1. Field-level Analysis

As we discussed in the Introduction, the problem of finding an optimal summary statistic that minimizes the error bars on a given cosmological or PNG parameter is unsolved. An alternative to using summary statistics is to perform field-level analysis. The goal with this kind of analysis is to maximize the amount of information that can be extracted without relying on summary statistics. While there are many types of methods to perform such analysis, in our case, we made use of GNNs (Battaglia et al. 2018). The advantages of GNNs over other methods are that (1) they do not impose a cut on scales, (2) symmetries (e.g., rotational and translational invariance) can be easily implemented, and (3) they can be more interpretable than other methods. Because of this, we decided to train the GNNs to perform field-level likelihood-free inference.

As a starting point, we run 1000 simulations, each containing 512^3 particles in a periodic box of size $1 h^{-1}$ Gpc. Each of those simulations has a different initial random seed but also a different value of $f_{\text{NL}}^{\text{local}}$ in the range $-300, +300$. The value of the cosmological parameters was the same in all simulations. We then trained a GNN to perform field-level likelihood-free inference on the value of $f_{\text{NL}}^{\text{local}}$. The architecture and training procedure are the same as those outlined in de Santi et al. (2023), Shao et al. (2022), and Villanueva-Domingo & Villaescusa-Navarro (2022).

From this exercise, we found that our model was able to infer the value of $f_{\text{NL}}^{\text{local}}$ with an error of $\sigma(f_{\text{NL}}^{\text{local}}) \sim 35$ at $z = 0$. In an attempt to understand the behavior of the network, we trained a deep set model (Zaheer et al. 2017), where the only information we made use of for the halos was their masses, not their spatial positions. By training such a model, we found that its performance was almost identical to that of the GNNs. We thus concluded that the network was likely not using the

clustering of the halos to perform the inference. Therefore, the network should be using the abundance of halos to infer $f_{\text{NL}}^{\text{local}}$.

To verify this, we trained a simple model consisting of fully connected layers on the HMF of the halo catalogs from the simulations. We found that this model performed almost as well as the GNN. From this exercise, we reached the conclusion that the HMF is a summary statistic that contains lots of information, likely more than clustering-based statistics, as the GNN did not use those to perform the inference. We emphasize that we trained the GNN using halo catalogs from simulations that only vary $f_{\text{NL}}^{\text{local}}$. Therefore, our results did not account for degeneracies with cosmological parameters that could degrade the constraints, as we shall see below.

This motivated a further analysis, illustrated in the following sections, in which we explicitly extract the power spectrum, bispectrum, and HMF from the QUIJOTE data set, as well as their covariance and response to variations in both cosmological and PNG parameters, in order to perform a full Fisher matrix forecast on nonlinear scales.

3.2. Fisher Information

In this section, we recall the main ingredients of our Fisher analysis pipeline, which was previously used in Jung et al. (2022b).

The Fisher information matrix, defined as

$$F_{ij} = \left(\frac{\partial \bar{s}}{\partial \theta_i} \right)^T \mathbf{C}^{-1} \left(\frac{\partial \bar{s}}{\partial \theta_j} \right), \quad (1)$$

allows us to estimate the variance, $\sigma^2(\theta_i) = \sqrt{(F^{-1})_{ii}}$, of the optimal unbiased estimator of a given summary statistic s with covariance \mathbf{C} assuming the statistic is Gaussian-distributed¹⁸ and neglecting the dependence of \mathbf{C} itself on the parameters (Carron 2013).

¹⁸ As verified in Jung et al. (2022a, 2022b), this is a good approximation for the power spectrum and bispectrum.

In this work, both the covariance and derivatives are computed from the simulations described in Section 2. The covariance matrix is evaluated using

$$\hat{\mathbf{C}} = \frac{1}{n_r - 1} (\mathbf{s} - \bar{\mathbf{s}})(\mathbf{s} - \bar{\mathbf{s}})^T, \quad (2)$$

where n_r is the number of realizations at fiducial cosmology (15,000 here). Then, to obtain an unbiased estimate of the precision matrix, we apply the Hartlap correction factor (Hartlap et al. 2007),

$$\mathbf{C}^{-1} = \frac{n_r - n_s - 2}{n_r - 1} \hat{\mathbf{C}}^{-1}, \quad (3)$$

where n_s is the length of the summary statistic vector \mathbf{s} (note, however, that this correction is very small here, $n_s \sim 10^2$, while $n_r = 15,000$).

The derivatives are calculated using finite difference,

$$\frac{\partial \bar{s}}{\partial \theta_i} = \frac{\bar{s}(\theta_i^{\text{fid}} + \delta\theta_i) - \bar{s}(\theta_i^{\text{fid}} - \delta\theta_i)}{2\delta\theta_i}, \quad (4)$$

where we use the sets of 500 simulations where one parameter θ_i is displaced by $\pm\delta\theta_i$ with respect to its fiducial value. However, it was noticed in Coulton et al. (2023b) and Jung et al. (2022b) that this number of realizations was not sufficient to obtain fully converged derivatives of the halo power spectrum and bispectrum, leading to spuriously low predictions when analyzing jointly cosmological parameters and PNG amplitudes. To overcome this issue, a conservative approach to Fisher matrix computations was developed in Coulton & Wandelt (2023) and Coulton et al. (2023b) that is based on computing the Fisher matrix from maximally compressed statistics instead of working with the summary statistics directly.

As shown in Heavens et al. (2000) and Alsing & Wandelt (2018), the compressed quantity defined by

$$\tilde{s}_i = \frac{\partial \bar{s}}{\partial \theta_i} \mathbf{C}^{-1} (\mathbf{s} - \bar{\mathbf{s}}) \quad (5)$$

conserves all of the statistical information about the parameter θ_i contained in the data vector \mathbf{s} , if \mathbf{s} follows a Gaussian likelihood (hence, the same assumption as for the Fisher matrix in Equation (1)). This compression uses the same ingredients as for the Fisher matrix computation (covariance and derivatives of \mathbf{s}), with the addition of the mean $\bar{\mathbf{s}}$ that is trivial to evaluate from the simulations at fiducial cosmology. Repeating the process for all parameters of interest in θ , one can then compute the Fisher matrix of the compressed statistics $\tilde{\mathbf{s}}$ by substituting it for \mathbf{s} in Equation (1). In practice, one has to separate the initial data set into two subsets. The first is used to perform the compression (i.e., compute the derivatives in Equation (5)), and the second is compressed (i.e., \mathbf{s} in Equation (5)) and then used to calculate the derivatives $\partial \tilde{s}_i / \partial \theta_j$ and covariance $\hat{\mathbf{C}}$ of the compressed statistics to obtain a conservative estimation of the Fisher matrix. In this work, we use 80% and 20% of the simulations for the two steps, respectively, which have been verified to give optimal and numerically stable results. We repeat the procedure for many random splits of the data




Figure 1. The HMF derivatives with respect to the parameters $\{\sigma_8, \Omega_m, f_{\text{NL}}^{\text{local}}, f_{\text{NL}}^{\text{equil}}, f_{\text{NL}}^{\text{ortho}}\}$ at $z = 0$ and 1 . For internal comparison, the derivative with respect a given parameter θ is multiplied by the finite difference $\Delta\theta$ used for its numerical estimation (see Table 1 for details). The vertical scale is logarithmic, except in the range $[-10^{-8}, 10^{-8}]$, where it is linear. Note that in some cases, we have a change of sign in the f_{NL} derivatives, implying an opposite effect of PNG on the abundance of high- and low-mass halos, respectively. This is consistent with previous findings in the literature, as pointed out in the main text. The decreasing behavior of all derivatives at high M is related to the exponential decay of the HMF in this mass range; note that a plot of the logarithmic derivatives would display clear differences between them, also at high M . The numerical results displayed here have all been cross-validated in the simulation-independent, halo model-based analysis that we describe in Section 4.4.

(between the two steps) and average the results to minimize the intrinsic variance of the method.

Finally, as shown in Coulton & Wandelt (2023), computing the following combination of the standard (overoptimistic) and compressed (conservative) Fisher matrices,

$$\mathbf{F}^{\text{combined}} = G(\mathbf{F}^{\text{standard}}, \mathbf{F}^{\text{compressed}}), \quad (6)$$

where G corresponds to the geometric mean defined by

$$G(A, B) = A^{\frac{1}{2}} (A^{-\frac{1}{2}} B A^{-\frac{1}{2}})^{\frac{1}{2}} A^{\frac{1}{2}}, \quad (7)$$

gives unbiased estimates of the Fisher error bars with a much smaller number of simulations. An illustration of the different convergences for the three methods is provided in Appendix C.

3.3. Halo Mass Function

In addition to the halo power spectrum and bispectrum, we consider the HMF, defined as the number of dark matter halos per unit of comoving volume per unit of logarithmic mass bins.

We measure it in the QUIJOTE simulations using 15 logarithmic bins corresponding to halo masses M between approximately 2.0×10^{13} and $4.6 \times 10^{15} M_{\odot} h^{-1}$ (note, however, that we do not use the first two bins in the analyses presented in Section 4). To be exact, we use the same binning as in Bayer et al. (2021), where the counted halos each contain between 30 and 7000 dark matter particles.¹⁹

In Figure 1, we show the impact of the three shapes of PNG on the HMF. Both the local and equilateral shapes increase the

¹⁹ The mass of a halo is given by $M = Nm_p$, where N is the number of dark matter particles it contains, and m_p is the mass of a dark matter particle. However, m_p depends on the cosmological parameter Ω_m , which requires inclusion of the correction term $-\frac{1}{\Omega_m} \frac{\partial \text{HMF}}{\partial \ln N}$ when computing the derivative $\frac{\partial \text{HMF}}{\partial \Omega_m}$ (see Bayer et al. 2021, for details). This derivative can also be evaluated by finite difference between bins of N .




Figure 2. The 1σ Fisher error bars on f_{NL} (local, equilateral, and orthogonal) from the HMF as a function of the maximum mass M_{max} of the halos considered ($M_{\text{min}} \sim 4.1 \times 10^{13} M_{\odot}/h$). These constraints are derived from the QUIJOTE suite of halo catalogs at $z = 0$ and 1, each having a $1 (\text{Gpc } h^{-1})^3$ volume. The solid lines (with triangles) are computed for each primordial shape independently, assuming a fixed cosmology (at fiducial values), while for the dashed–dotted lines, we marginalize over the cosmological parameters σ_8 and Ω_m . This highlights the large degeneracies between the parameters at the level of the HMF. For comparison, we also show the corresponding constraints from the power spectrum and bispectrum (horizontal solid and dashed–dotted lines for the independent and joint cases, respectively), as computed previously in Jung et al. (2022b; $M_{\text{min}} = 3.2 \times 10^{13} M_{\odot}/h$). If we consider the unmarginalized HMF results, we see that the f_{NL} constraining power is higher at $z = 1$ for the local and equilateral cases, despite the smaller number of halos at this redshift; this is clearly due to a stronger response of the HMF to variations in f_{NL} at higher redshift, consistent with previous findings (see, e.g., Figure 4 in LoVerde et al. 2008). The shape is due to the change of sign in the f_{NL} derivative at different masses, discussed in the main text and Figure 1.

number of massive halos for a positive f_{NL} value (and decrease it for a negative f_{NL}) and have very degenerate signatures, while for orthogonal PNG, it is the opposite. For less massive halos, the effect of PNG changes sign (with the switch occurring for higher masses for orthogonal PNG, which is the only one that appears in the mass range of the plot at $z = 1$). This effect was already present in early works on the HMF with PNG simulations (see, e.g., LoVerde et al. 2008) and is due to the fact that, at fixed Ω_m , more massive halos can only appear at the expense of less massive halos and matter in smaller structures.

4. Results

4.1. Constraints from the HMF

As a preliminary exercise, in Figure 2, we show the constraining power of the HMF on the PNG amplitudes f_{NL} of the three shapes, assuming the exactly known cosmological parameters. As expected, the HMF is, in this case, extremely sensitive to the presence of PNG, leading to even tighter constraints than the power spectrum and bispectrum. For example, our Fisher forecast on PNG of the local type is $\sigma(f_{\text{NL}}^{\text{local}}) \sim 30$ at $z = 0$, which is in very good agreement with the GNN and deep set results $\sigma(f_{\text{NL}}^{\text{local}}) \simeq 35$ (see Section 3.1) and more than twice as small as the equivalent power spectrum + bispectrum forecast error bar.

However, it is well known that there are large degeneracies between f_{NL} and several cosmological parameters, like σ_8 or Ω_m (Maturi et al. 2011), as can be verified in Figure 1.

When we jointly analyze all parameters, these degeneracies increase the errors significantly (by roughly 1 order of magnitude at $z = 1$ and slightly less at $z = 0$, where the change

of sign of the f_{NL} derivative, seen in Figure 1, helps distinguish it from the f_{NL} response to variations in other cosmological parameters), making them larger than those achievable from the power spectrum and bispectrum combination.

4.2. Joint Constraints with the Power Spectrum and Bispectrum

While, as expected, the HMF alone does not produce competitive f_{NL} constraints in comparison with the power spectrum and bispectrum, it does remain interesting to investigate whether a combined analysis of all three statistics can produce significant improvements; this is the main point of the present work. Complementing our previous power spectrum + bispectrum analysis with the HMF can, in principle, benefit us in two ways. First of all, it directly adds extra information about the f_{NL} parameter; also, it could be useful to help break the important degeneracy between f_{NL} and the so-called b_{ϕ} bias parameter.

Before presenting our results, let us review and discuss the latter point in more detail. In the presence of local PNG, the halo density fluctuation field $\delta_h(z)$ can be written to leading order as follows (Dalal et al. 2008; Matarrese & Verde 2008; Slosar et al. 2008; McDonald 2008; Giannantonio & Porciani 2010; Desjacques & Seljak 2010a):

$$\delta_h(z) = \left[b_1(z) + \frac{3\Omega_m H_0^2}{2D(z)k^2} b_{\phi} f_{\text{NL}} \right] \delta_m(z), \quad (8)$$

where δ_m is the matter density fluctuation, $D(z)$ is the growth factor, and b_1 and b_{ϕ} are the bias parameters, defined respectively as the response of δ_h to mass density δ_m and primordial potential ϕ . It is evident in this relation that the scale-dependent signature depends on both b_{ϕ} and f_{NL} and that the two parameters are completely degenerate. This issue can be avoided if one assumes, as was generally done, the universality relation between b_1 and b_{ϕ} , that is,

$$b_{\phi} = 2\delta_c(b_1 - 1), \quad (9)$$

where δ_c is the critical density for collapse. However, it was recently pointed out in Barreira (2020, 2022) that such a relation does not accurately describe the bias of either galaxies, selected by stellar mass, or halos, selected by concentration. Therefore, b_{ϕ} is not exactly determined anymore, and this reintroduces the b_{ϕ} - f_{NL} degeneracy problem. To overcome the issue, different studies have focused on using simulations to produce accurate priors on b_{ϕ} (Lazeyras et al. 2023) and exploiting the multitracer technique (Barreira & Krause 2023; Karagiannis et al. 2023; Sullivan et al. 2023). In the present context, the idea is instead to try and break the degeneracy by exploiting the information in the HMF—which selects all halos in each given mass bin—and its direct dependence on f_{NL} and not on b_{ϕ} .

For clarity, we split the discussion of our results into two parts. Initially, we assume universality in the $b_{\phi}(b_1)$ relation using Equation (9), and we measure the sheer extra information content in the HMF, in the absence of the b_{ϕ} - f_{NL} degeneracy.²⁰ Later on, we instead treat b_{ϕ} as a free parameter.

²⁰ Or, equivalently, we forecast the power spectrum + bispectrum + HMF constraining power on the $b_{\phi}f_{\text{NL}}$ parameter combination.




Figure 3. Ratio of 1σ Fisher error bars on the cosmological parameters and PNG amplitudes from the HMF, halo power spectrum, and halo bispectrum at $z = 1$, assuming b_ϕ is fixed. This illustrates how including the HMF tightens the constraints on several parameters (σ_8 and $f_{\text{NL}}^{\text{equil}}$ in particular). Note that the values of these error bars are given in Table 2 and Figure 11.

4.3. Fixing b_ϕ

The outcome of the first part of the analysis (assuming universality in $b_\phi(b_1)$) is illustrated in Figures 3 and 4 (see also Table 2). We see that by adding the HMF, the error bars on σ_8 and $f_{\text{NL}}^{\text{equil}}$ become roughly twice as small as the power spectrum + bispectrum result. Moreover, there is also a noticeable improvement for Ω_m and $f_{\text{NL}}^{\text{ortho}}$. For $f_{\text{NL}}^{\text{local}}$, there is instead no clear improvement; this seems to be due to the fact that in this case, the information content is totally dominated by the power spectrum contribution via scale-dependent bias. Such a contribution is instead smaller for the orthogonal shape and absent for the equilateral case, making the HMF inclusion more important for these scenarios, especially the equilateral one.

Note that we consider only halos with masses above $\sim 4 \times 10^{13} M_\odot h^{-1}$ in the HMF, which is larger than the fiducial $M_{\text{min}} = 3.2 \times 10^{13} M_\odot/h$ used to study the power spectrum and bispectrum. This means that the HMF is not sensitive at all to small variations of M_{min} around the fiducial value. However, through cross-correlated terms with the other summary statistics, the error bars on M_{min} are almost 2 orders of magnitude smaller.²¹ In Appendix C, we verify the numerical stability of our results by varying the number of simulations used.

²¹ An important caveat here is that it is important to verify whether this conclusion holds when considering a more complex bias model, which includes higher-order bias parameters; this will be done as part of a future work on mock galaxy catalogs by including numerical derivatives with respect to HOD parameters.

It is interesting to check which halo mass range gives the largest contribution to the observed improvements. To this purpose, we repeat the analysis by varying $M_{\text{min}}^{\text{HMF}}$, the lowest mass bin of the HMF used to evaluate Fisher matrices. Our results are displayed in Figure 5, which highlights different behaviors for the different parameters considered. Most importantly, for the two PNG parameters $f_{\text{NL}}^{\text{equil}}$ and $f_{\text{NL}}^{\text{ortho}}$, halos of intermediate masses ($\sim 2\text{--}6 \times 10^{14} M_\odot h^{-1}$ at $z = 0$ and slightly smaller at $z = 1$) play a significant role in the observed improvement of constraints, while less massive halos, despite being more numerous, have a much smaller effect. However, the situation is different for cosmological parameters like σ_8 and Ω_m , where those same less massive halos contain most of the information.

4.4. Breaking the b_ϕ - $f_{\text{NL}}^{\text{local}}$ Degeneracy with the HMF

Accounting for the effects of b_ϕ in our methodology is not straightforward, since b_ϕ cannot be explicitly included as an input parameter in our simulations, and this does not allow us to directly compute the numerical derivative $\partial s/\partial b_\phi$. To circumvent this issue in a simple way and be able to perform a first test of the ability of the HMF to remove degeneracies between b_ϕ and $f_{\text{NL}}^{\text{local}}$, we decide here to work under the conservative assumption that these two parameters are fully degenerate at the level of the halo power spectrum and bispectrum. In other words, we assume that $\partial s/\partial b_\phi \propto \partial s/\partial f_{\text{NL}}^{\text{local}}$, where s is either the power spectrum or the bispectrum. For the HMF, we instead set the derivative with respect to b_ϕ equal to zero, as it does not depend on this parameter, and compute the $f_{\text{NL}}^{\text{local}}$ derivative as usual.

In Figure 6, we show the 1σ Fisher constraints obtained in this assumption and compare them with the “ideal” (b_ϕ fixed) constraints derived in the previous section for different k_{max} (see also Table 2).

The most important result here is that the inclusion of the HMF makes it possible to break the b_ϕ - $f_{\text{NL}}^{\text{local}}$ degeneracy to a level that allows us to produce meaningful $f_{\text{NL}}^{\text{local}}$ constraints without resorting to any prior information on b_ϕ . The final $f_{\text{NL}}^{\text{local}}$ forecast is, however, degraded by a factor of ~ 2.5 with respect to the idealized, b_ϕ fixed case that was shown in Figure 11 in Appendix C. In order to achieve this constraining level, it is also crucial to include the information from the power spectrum and bispectrum at nonlinear scales (k between 0.2 and $0.5 h \text{ Mpc}^{-1}$), as it helps break degeneracies with several cosmological parameters (Ω_m in particular).

We corroborate our findings with a simulation-independent analysis based on the halo model (for a review, see Cooray & Sheth 2002; Asgari et al. 2023). Within this framework, we describe the HMF and halo power spectrum following Takada & Spergel (2014) up to $k_{\text{max}} = 0.2 h \text{ Mpc}^{-1}$. We use the HMF and bias from Tinker et al. (2010) using $M_{200,m}$ directly as the mass definition in the mass integration. In the power spectrum analysis of the simulations, the halos are considered pointlike; thus, we use a Dirac delta as the halo profile. Thanks to the low k_{max} we use, the two-halo term dominates the signal, and this approximation is appropriate. The effect of PNG—here we only consider the local model—is included as a correction to the HMF parameterized according to LoVerde & Smith (2011) and through the scale-dependent halo bias shown in Equation (8). While we are aware that the $M_{200,m}$ mass does not match the FOF mass used in the rest of the paper, we still




Figure 4. Impact of the HMF on the 1σ constraints on the cosmological parameters and PNG amplitudes from the halo power spectrum and bispectrum at $z = 1$, assuming b_ϕ is fixed.

consider the HMF divided in 10 bins logarithmically spaced between 3.2×10^{13} and $3.2 \times 10^{15} M_\odot h^{-1}$ as observable. We bin the halo power spectrum in 30 bins logarithmically spaced between 6.3×10^{-3} and $0.2 h \text{ Mpc}^{-1}$. We choose a relatively low k_{max} to ensure that nonlinearities are negligible at this stage. In the HMF–halo power spectrum covariance, for which we again follow Takada & Spergel (2014), only the Gaussian terms are included at present. A more refined analysis, including a wider range of scales and masses, the complete covariance, the uncertainties on the parameterization of the HMF, and, crucially, the bispectrum, will be presented in a future work (Ravenni et al. 2023, in preparation).

The results are shown in Figure 7, which highlights a very good agreement between our preliminary theoretical computations and the purely simulation-based forecast. This result confirms that a joint analysis including the HMF is an interesting approach that deserves further investigation and could be adopted as a complementary strategy to those already implemented in the literature to address the b_ϕ – $f_{\text{NL}}^{\text{local}}$ degeneracy issue.

4.5. Removing Degeneracies with Planck Priors

As highlighted in Section 4.2, removing the degeneracies of the HMF using the information from the halo power spectrum and bispectrum significantly improves the constraints on PNG of the




Figure 5. Impact of varying the lowest mass bins of the HMF on the 1σ Fisher constraints on cosmological parameters and PNG amplitudes from the combination of HMF, power spectrum, and bispectrum at $z = 0$ and 1 , assuming b_ϕ is fixed. All errors are normalized by their equivalent using the power spectrum and bispectrum only. Note that we restrict only the mass range for the HMF.

Table 2

The 1σ Constraints on the Cosmological Parameters and PNG Amplitudes at $z = 1$ Obtained by Combining the Information of the Halo Power Spectrum, Bispectrum, and Mass Function, Each Measured from the QUIJOTE and QUIJOTE-PNG Simulations

	b_ϕ Fixed	No Prior on b_ϕ	Planck Priors
σ_8	0.012	0.013	0.005
Ω_m	0.018	0.017	0.002
n_s	0.075	0.075	0.003
h	0.072	0.071	0.017
$f_{\text{NL}}^{\text{local}}$	40	89	34
$f_{\text{NL}}^{\text{equil}}$	203		136
$f_{\text{NL}}^{\text{ortho}}$	85		79
$M_{\text{min}}/10^{13}$	0.019	0.045	0.009

equilateral type. In this section, we push the idea further by assuming strong but realistic priors on cosmological parameters based on CMB measurements from Planck.

We use the same Gaussian likelihood based on the Planck CMB data (Aghanim et al. 2020) as in Uhlemann et al. (2020) in Figure 8 in addition to our HMF, power spectrum, and bispectrum measurements to derive 1σ Fisher constraints (see also Table 2). For both $f_{\text{NL}}^{\text{local}}$ and $f_{\text{NL}}^{\text{equil}}$, it improves these constraints, while the effect is smaller for $f_{\text{NL}}^{\text{ortho}}$. Note also that the effect is strongest when the HMF is also considered in the analysis, meaning it removes degeneracies between the PNG and cosmological parameters at the level of the HMF. Concerning numerical convergence with the number of simulations used to compute the derivatives, including these Planck priors also improves it significantly, where only $f_{\text{NL}}^{\text{equil}}$ is not optimally constrained for the power spectrum + bispectrum case, and all parameters have converged when we add the HMF information.

5. Conclusion

In this work, we presented a combined analysis of the power spectrum, bispectrum, and mass function of dark matter halos

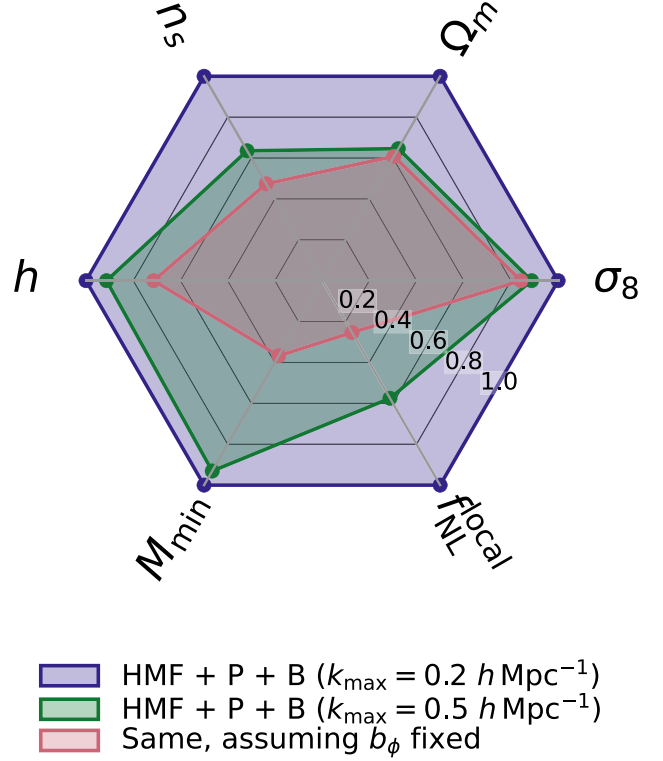


Figure 6. The HMF can break the b_ϕ - $f_{\text{NL}}^{\text{local}}$ degeneracy in the power spectrum and bispectrum. As in Figure 3, we show normalized 1σ Fisher error bars derived from the HMF, halo power spectrum, and bispectrum at $z = 1$. Here we assume that $f_{\text{NL}}^{\text{local}}$ and b_ϕ are fully degenerate at the power spectrum and bispectrum level, while the HMF does not depend on b_ϕ .

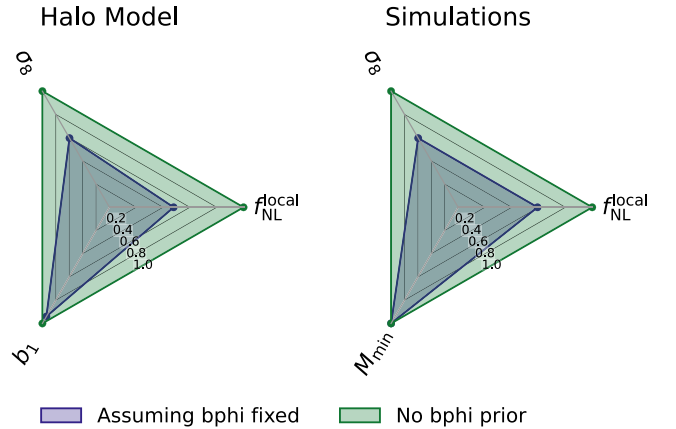


Figure 7. Similar to Figure 6, considering only $\{\sigma_8, f_{\text{NL}}^{\text{local}}\}$ and bias parameters. The 1σ Fisher constraints include the information contained in the HMF and the power spectrum information up to $k_{\text{max}} = 0.2 h \text{ Mpc}^{-1}$ computed using the halo model on the left and from simulations on the right. Note that both methods give $\sigma(f_{\text{NL}}^{\text{local}}) \sim 50$ and similar $\sigma(\sigma_8)$ (less than 20% difference).

in the QUIJOTE-PNG simulation suite. Our main goal was to verify whether adding the HMF to our previous joint power spectrum and bispectrum analyses (Coulton et al. 2023a, 2023b; Jung et al. 2022a, 2022b) could lead to improved constraints on primordial non-Gaussianity (PNG). The main underlying reason behind this analysis is that the HMF turned out to be the statistics used by a sophisticated graph neural network when carrying out a preliminary field-

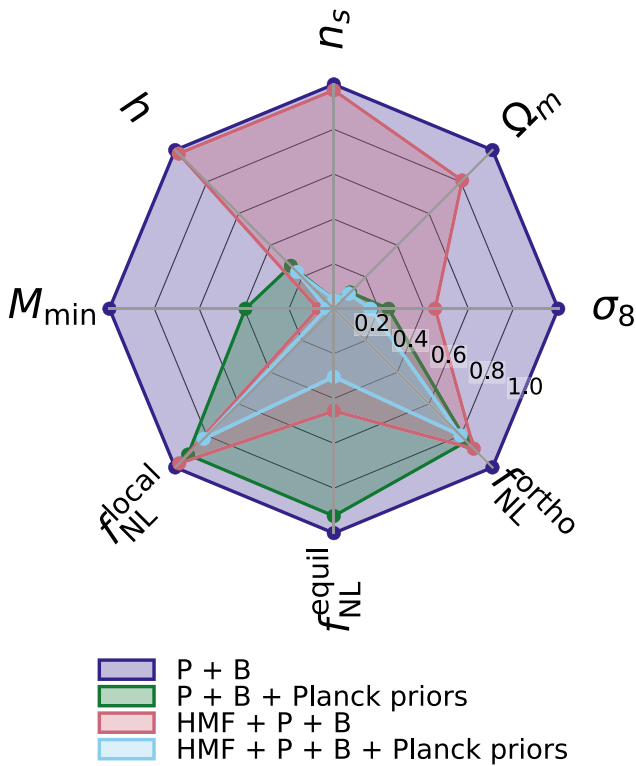


Figure 8. Similar to Figure 3, where we include Planck priors on the cosmological parameters $\{\sigma_8, \Omega_m, n_s, h\}$, and we assume b_ϕ is fixed.

level likelihood-free inference calculation. Furthermore, the HMF tail has been known for a long time to be strongly sensitive to PNG. Finally, the HMF not only carries complementary information to the power spectrum and bispectrum but also does not suffer from the b_ϕ - $f_{\text{NL}}^{\text{local}}$ assembly bias-PNG degeneracy that was recently pointed out in Barreira (2020, 2022) as an important issue in the analysis of local PNG.

Our results show that the HMF can indeed play a significant role in tightening the expected PNG bounds and breaking parameter degeneracies when its contribution is added to those of the power spectrum and bispectrum. In the first part of our analysis, we remove a priori the b_ϕ - $f_{\text{NL}}^{\text{local}}$ degeneracy by assuming universality in the $b_\phi(b_1)$ relation; i.e., we set $b_\phi = 2\delta_c(b_1 - 1)$. In this case, we see that the HMF is able to improve equilateral f_{NL} constraints by roughly a factor 2 and orthogonal f_{NL} constraints by 10%. Constraints on PNG of the local type are instead unchanged, since in this idealized scenario, the local PNG information is dominated by the large-scale power spectrum modes via scale-dependent bias.

In the second part of the analysis, we instead treat b_ϕ as a free parameter and assume that the responses of the halo power spectrum and bispectrum to the changes in b_ϕ and $f_{\text{NL}}^{\text{local}}$ are identical; that is, we assume that these two parameters are fully degenerate in a joint analysis of the power spectrum and bispectrum. Starting with this setup, we then see that the additional inclusion of the HMF is able to break the b_ϕ - $f_{\text{NL}}^{\text{local}}$ degeneracy at a significant level, without the need to rely on any prior on b_ϕ or any other external information. More precisely, our final $f_{\text{NL}}^{\text{local}}$ constraints after marginalizing over b_ϕ and other standard cosmological parameters are now degraded by a factor of ~ 2.5 compared to the ideal case, in which b_ϕ is fixed by the universality relation. We confirmed these results with a semianalytical, halo model-based evaluation of the Fisher matrix,

in which we restrict ourselves to the power spectrum and HMF, after verifying that for local PNG, these two observables give the dominant contributions to the final sensitivity. We note that to achieve the claimed level of precision on $f_{\text{NL}}^{\text{local}}$, it is important to include nonlinear scales in the analysis up to $k_{\text{max}} = 0.5 \text{ h Mpc}^{-1}$, since they help break additional important degeneracies that affect the HMF constraining power. We also stress that QUIJOTE-PNG simulations have a cosmological volume of 1 (h Gpc)^{-3} , making it not straightforward to generalize our forecasts to, e.g., a Euclid-like or other coming survey settings. For the same reason, a direct comparison with other forecasts—such as those based on the multitracer methodology and placing suitable priors on b_ϕ —is not easy to make at the moment. In a forthcoming publication, Ravenni et al. (2023, in preparation), we will produce more detailed semianalytical predictions for future surveys based on the halo model.

The results presented here have to be considered as preliminary also, as they rely on a simplified bias model for our tracers, and they do not account for systematic effects in the determination of the HMF from actual observations. Indeed, the dark matter mass of a halo is a quantity that is notoriously difficult to measure observationally, especially for high-redshift objects. Halos are complex and dynamic structures that are almost exclusively probed by the signal broadcast by the baryons they host. Dark mass measurements tend to require sophisticated and labor-intensive observations, which is unfeasible for a large number of objects, as needed for the HMF. Moreover, the sample completeness (for the host halo, not the tracers) needs to be known exquisitely well, which may constitute a formidable challenge. Among the most promising approaches are the clusters selected by the Sunyaev-Zel’dovich effect (signal at millimeter wavelengths; Mroczkowski et al. 2019), X-ray clusters (Pratt et al. 2019), and (optical) gravitational-lensing mass determination (e.g., Murray et al. 2022). For example, cluster catalogs will increase drastically with a suite of forthcoming experiments: eROSITA (Predehl et al. 2021), Simons Observatory (Ade et al. 2019), Euclid (Laureijs et al. 2011), Roman (Akeson et al. 2019), and Rubin (Ivezić et al. 2019). Cluster masses will not be measured directly but inferred through proxies; these proxies, however, will be provided as a product of these surveys and are expected to be or be made robust and reliable. An important ingredient for any HMF analysis would be to robustly quantify the probability distribution of the proxies as a function of the true halo mass. This can then be simply folded into the error budget and the uncertainty propagated through to the inferred parameters.

The results shown in this paper clearly show that a joint analysis of the HMF, power spectrum, and bispectrum of LSS tracers is a promising approach to constrain PNG, hence providing another motivation for further investigation in this direction and for addressing the aforementioned observational issues.

Acknowledgments

G.J. acknowledges support from the ANR LOCALIZATION project, grant ANR-21-CE31-0019/490702358 of the French Agence Nationale de la Recherche. A.R. acknowledges support from PRIN-MIUR 2020 METE under contract No. 2020KB33TP. The work of F.V.N. is supported by the Simons Foundation. D.K. is supported by the South African Radio Astronomy Observatory and the National Research Foundation (grant No. 75415). L.V. acknowledges the “Center of Excellence Maria de Maeztu 2020-2023” award to the ICCUB

(CEX2019-000918-M funded by MCIN/AEI/10.13039/501100011033).

Appendix A Examination of the ICs

The procedure used to generate the simulation ICs in Coulton et al. (2023a) is designed to produce a specific bispectrum. However, the method additionally modifies all other N -point functions. The most well-studied by-product of this procedure is modifications to the power spectrum. Scoccimarro et al. (2012) showed that it must be taken in when choosing how to generate the ICs to avoid having corrections that dominate the power spectrum. In Coulton et al. (2023a) and Jung et al. (2022a), the ICs were validated by examining the power spectrum and bispectrum. Those tests showed that the modifications to the power spectrum are small, and the correct bispectrum was generated. A concern for the results presented in this work and other studies of statistics beyond the two- and three-point functions is that the ICs may have unphysically large higher-order N -point functions that impact the results. The power spectrum and trispectrum are the leading-order unwanted by-products of the IC generation procedure. If we can show that the corrections to both are small, it is reasonable to assume that the impacts of the unphysical higher N -point functions of the ICs are negligible for studies of the HMF and other statistics of the simulations. Given that the power spectrum has already been validated, in this Appendix, we present an investigation into the properties of the trispectrum.

A.1. Trispectrum Estimation

The trispectrum is defined as

$$\langle \delta(\mathbf{k}_1)\delta(\mathbf{k}_2)\delta(\mathbf{k}_3)\delta(\mathbf{k}_4) \rangle = T(k_1, k_2, k_3, k_4, K_a, K_b), \quad (\text{A1})$$

where $k_i = |\mathbf{k}_i|$, $K_a = |\mathbf{k}_1 + \mathbf{k}_2|$, and $K_b = |\mathbf{k}_1 + \mathbf{k}_3|$. Estimating the full trispectrum is computationally highly challenging; so, in this work, we measure trispectra averaged over K_b , i.e.,

$$T(k_1, k_2, k_3, k_4, K_a) \propto \sum_{K_b} T(k_1, k_2, k_3, k_4, K_a, K_b). \quad (\text{A2})$$

A binned version of this can be estimated as

$$\begin{aligned} \hat{T}(k_a, k_b, k_c, k_d, K_E) &= \frac{1}{N_{a,b,c,d,E}} \int \prod_{i=1,4} \frac{d^3 k_i}{(2\pi)^3} \int \frac{d^3 K_a}{(2\pi)^3} \\ &\times (2\pi)^3 \delta^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 - \mathbf{K}_a) (2\pi)^3 \delta^{(3)}(\mathbf{K}_a - \mathbf{k}_3 - \mathbf{k}_4) W_a(\mathbf{k}_1) \\ &\times W_b(\mathbf{k}_2) W_c(\mathbf{k}_3) W_d(\mathbf{k}_4) W_E(\mathbf{K}_a) \delta(\mathbf{k}_1)\delta(\mathbf{k}_2)\delta(\mathbf{k}_3)\delta(\mathbf{k}_4), \end{aligned} \quad (\text{A3})$$

where $W_a(k)$ selects modes that lie within binned a , and $N_{a,b,c,d,E}$ is the normalization. In this work, we use 14 equally spaced bins between $k = 0.0102$ and $0.193 h \text{ Mpc}^{-1}$. By utilizing

$$\delta^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) = \int d^3 x e^{i\mathbf{x}\cdot\mathbf{k}} \delta(\mathbf{k}), \quad (\text{A4})$$

we efficiently implement the estimator by first computing

$$\delta_{W_a}(\mathbf{x}) = \int \frac{d^3 k_a}{(2\pi)^3} e^{i\mathbf{x}\cdot\mathbf{k}} \delta(\mathbf{k}) W_a(\mathbf{k}), \quad (\text{A5})$$

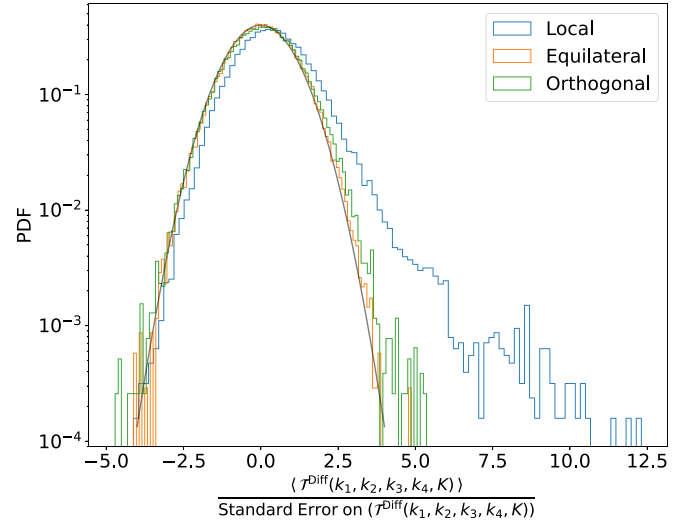


Figure 9. Significance of the detection of the trispectrum in the ICs for the three types of PNG. This is computed using 200 simulations of each type of PNG.

then computing

$$D_{ab}(\mathbf{K}) = \int d^3 x e^{-i\mathbf{x}\cdot\mathbf{K}} \delta_{W_a}(\mathbf{x}) \delta_{W_b}(\mathbf{x}), \quad (\text{A6})$$

and then the estimate is given by

$$\begin{aligned} \hat{T}(k_a, k_b, k_c, k_d, K_E) &= \frac{1}{N_{a,b,c,d,E}} \int \frac{d^3 K}{(2\pi)^3} D_{ab}(\mathbf{K}) D_{cd}(-\mathbf{K}) W_E(\mathbf{K}). \end{aligned} \quad (\text{A7})$$

The normalization is obtained by evaluating this estimator (without the $N_{a,b,c,d,E}$ term) on maps with $\delta(\mathbf{k}) = 1$.

A.2. Trispectrum of the ICs

To perform a stringent test of the ICs, we study the difference between the trispectrum of the ICs with $f_{\text{NL}} \neq 0$ and $f_{\text{NL}} = 0$, i.e.,

$$\begin{aligned} \hat{T}^{\text{diff}}(k_a, k_b, k_c, k_d, K_E) &= \hat{T}^{f_{\text{NL}}=0}(k_a, k_b, k_c, k_d, K_E) \\ &- \hat{T}^{f_{\text{NL}} \neq 0}(k_a, k_b, k_c, k_d, K_E). \end{aligned} \quad (\text{A8})$$

This cancels the leading noise contribution to the trispectrum measurement.

The results are shown in Figure 9. For equilateral, there is no detectable trispectrum. For orthogonal NG, there are small hints of a trispectrum signal. As this measurement uses 200 simulations and a method to cancel the cosmic variance, it is likely that this small trispectrum is negligible. However, the local case shows significant evidence of a trispectrum. This is not unexpected. Local PNG is generated in these simulations by

$$\Phi(\mathbf{x}) = \phi^G(\mathbf{x}) + f_{\text{NL}} (\phi^G(\mathbf{x})^2 - \langle \phi^G(\mathbf{x}) \rangle), \quad (\text{A9})$$

where $\Phi^G(\mathbf{x})$ is the Gaussian primordial potential. This generates a primordial trispectrum known as τ_{NL} (Kogo & Komatsu 2006). In many inflationary models, τ_{NL} is generated with local NG; thus, the trispectrum seen here is physical.

These trispectra measurements suggest that unphysical higher-order N -point functions are not significant in our simulations.

Appendix B
Analyses at Other Redshifts

We have performed a similar analysis using the QUIJOTE snapshots at $z = 0.5$ and zero to verify that our conclusions hold at other lower redshifts. As can be seen in Figure 10, this

is indeed the case. For all parameters, the relative improvements due to including the HMF in the Fisher analysis are of the same order (note, however, that the difference between the halo power spectrum and bispectrum results is more pronounced at lower redshifts).

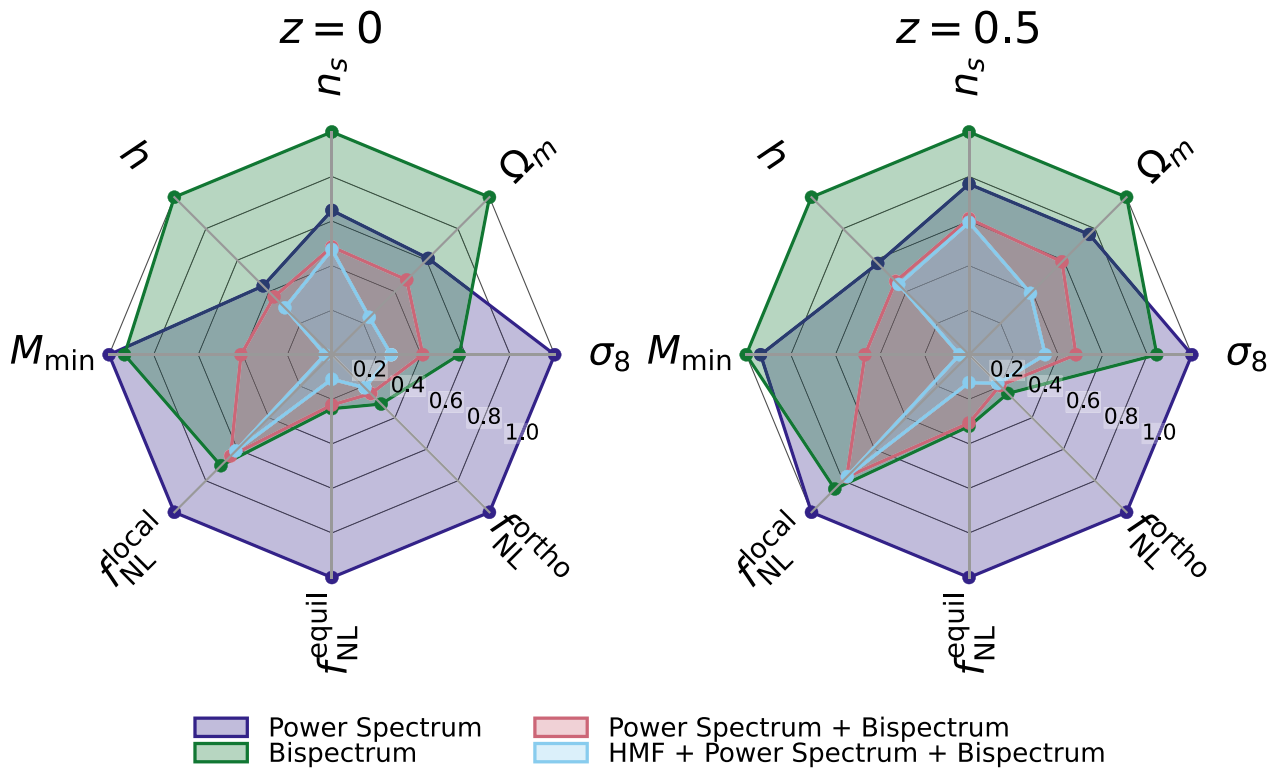


Figure 10. Similar to Figure 3 at redshifts $z = 0$ and 0.5 and with b_ϕ fixed.

Appendix C Convergence of Numerical Derivatives

In Figure 11, we study the impact of varying the number of simulations used to compute numerical derivatives on the 1σ Fisher constraints both with and without including the HMF in the analyses. This shows that the parameters for which the

improvement due to the HMF is the largest (i.e., σ_8 and $f_{\text{NL}}^{\text{equil}}$) also have a better numerical convergence with the number of simulations (a smaller difference between the standard and conservative compressed Fisher methods). Note also the stability of the combined Fisher results (variations at the % level) when using more than 200 simulations for the derivatives.











Figure 11. Stability of the Fisher 1σ error bars when varying the number of simulations used to compute derivatives for the three methods described in Section 3.2 (standard, compressed, and combined). In the left panels, the analysis includes the power spectrum (monopole + quadrupole) and bispectrum (monopole) information of the halo field at $z = 1$, with scales up to $k_{\text{max}} = 0.5 h \text{ Mpc}^{-1}$. In the right panels, we also consider the HMF (for halos with a mass larger than $4.1 \times 10^{13} M_{\odot} h^{-1}$). All error bars are normalized by their respective combined Fisher results, given explicitly in the legend for all parameters. They show that adding HMF can significantly reduce the error bars, in addition to improving the numerical convergence of the results (smaller relative differences between the compressed and standard methods) for several parameters, in particular σ_8 and $f_{\text{NL}}^{\text{equil}}$. Note that the lines corresponding to PNG parameters are in bold.

ORCID iDs

Gabriel Jung  <https://orcid.org/0000-0002-2133-6254>
 Marco Baldi  <https://orcid.org/0000-0003-4145-1943>
 Drew Jamieson  <https://orcid.org/0000-0001-5044-7204>
 Helen Shao  <https://orcid.org/0000-0002-0152-6747>
 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

- Abazajian, K., Addison, G., Adshead, P., et al. 2019, arXiv:1907.04473
 Ade, P., Aguirre, J., Ahmed, Z., et al. 2019, *JCAP*, **02**, 056
 Aghanim, N., Akrami, Y., Ashdown, M., et al. 2020, *A&A*, **641**, A5
 Akesson, R., Armus, L., Bachelet, E., et al. 2019, arXiv:1902.05569
 Akrami, Y., Arroja, F., Ashdown, M., et al. 2020, *A&A*, **641**, A9
 Alsing, J., & Wandelt, B. 2018, *MNRAS*, **476**, L60
 Asgari, M., Mead, A. J., & Heymans, C. 2023, arXiv:2303.08752
 Barreira, A. 2020, *JCAP*, **12**, 031
 Barreira, A. 2022, *JCAP*, **11**, 013
 Barreira, A., & Krause, E. 2023, arXiv:2302.09066
 Battaglia, P. W., Hamrick, J. B., Bapst, V., et al. 2018, arXiv:1806.01261
 Bayer, A. E., Villaescusa-Navarro, F., Massara, E., et al. 2021, *ApJ*, **919**, 24
 Cabass, G., Ivanov, M. M., Philcox, O. H. E., Simonović, M., & Zaldarriaga, M. 2022a, arXiv:2204.01781
 Cabass, G., Ivanov, M. M., Philcox, O. H. E., Simonović, M., & Zaldarriaga, M. 2022b, arXiv:2201.07238
 Carron, J. 2013, *A&A*, **551**, A88
 Cooray, A., & Sheth, R. K. 2002, *PhR*, **372**, 1
 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
 Coulton, W. R., & Wandelt, B. D. 2023, arXiv:2305.08994
 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
 Davis, M., Efstathiou, G., Frenk, C. S., & White, S. D. M. 1985, *ApJ*, **292**, 371
 de Santi, N. S. M., Shao, H., Villaescusa-Navarro, F., et al. 2023, *ApJ*, **952**, 69
 Desjacques, V., & Seljak, U. 2010a, *CQGra*, **27**, 124011
 Desjacques, V., & Seljak, U. 2010b, *AdAst*, **2010**, 908640
 Desjacques, V., Seljak, U., & Iliev, I. 2009, *MNRAS*, **396**, 85
 Finelli, F., Bucher, M., Achúcarro, A., et al. 2018, *JCAP*, **04**, 016
 Giannantonio, T., & Porciani, C. 2010, *PhRvD*, **81**, 063530
 Giri, U., Münchmeyer, M., & Smith, K. M. 2023, arXiv:2305.03070
 Grossi, M., Dolag, K., Branchini, E., Matarrese, S., & Moscardini, L. 2007, *MNRAS*, **382**, 1261
 Grossi, M., Verde, L., Carbone, C., et al. 2009, *MNRAS*, **398**, 321
 Hahn, C., & Villaescusa-Navarro, F. 2021, *JCAP*, **2021**, 029
 Hahn, C., Villaescusa-Navarro, F., Castorina, E., & Scoccimarro, R. 2020, *JCAP*, **2020**, 040
 Hartlap, J., Simon, P., & Schneider, P. 2007, *A&A*, **464**, 399
 Heavens, A., Jimenez, R., & Lahav, O. 2000, *MNRAS*, **317**, 965
 Ivezić, V., Kahn, S., Tyson, A., et al. 2019, *ApJ*, **873**, 111
 Jung, G., Karagiannis, D., Liguori, M., et al. 2022a, *ApJ*, **940**, 71
 Jung, G., Karagiannis, D., Liguori, M., et al. 2022b, arXiv:2211.07565
 Karagiannis, D., Maartens, R., Fonseca, J., Camera, S., & Clarkson, C. 2023, arXiv:2305.04028
 Kogo, N., & Komatsu, E. 2006, *PhRvD*, **73**, 083007
 Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, arXiv:1110.3193
 Lazeyras, T., Barreira, A., Schmidt, F., & Desjacques, V. 2023, *JCAP*, **01**, 023
 LoVerde, M., Miller, A., Shandera, S., & Verde, L. 2008, *JCAP*, **04**, 014
 LoVerde, M., Miller, A., Shandera, S., & Verde, L. 2008, *JCAP*, **2008**, 014
 LoVerde, M., & Smith, K. M. 2011, *JCAP*, **08**, 003
 Matarrese, S., & Verde, L. 2008, *ApJL*, **677**, L77
 Matarrese, S., Verde, L., & Jimenez, R. 2000, *ApJ*, **541**, 10
 Maturi, M., Fedeli, C., & Moscardini, L. 2011, *MNRAS*, **416**, 2527
 McDonald, P. 2008, *PhRvD*, **78**, 123519
 Moradinezhad Dizgah, A., Biagetti, M., Sefusatti, E., Desjacques, V., & Noreña, J. 2021, *JCAP*, **2021**, 015
 Mroczkowski, T., Nagai, D., Basu, K., et al. 2019, *SSRv*, **215**, 17
 Murray, C., Bartlett, J. G., Artis, E., & Melin, J.-B. 2022, *MNRAS*, **512**, 4785
 Palma, G. A., Scheiing Hirschfeld, B., & Sypsas, S. 2020, *JCAP*, **02**, 027
 Pillepich, A., Porciani, C., & Hahn, O. 2010, *MNRAS*, **402**, 191
 Pratt, G. W., Arnaud, M., Biviano, A., et al. 2019, *SSRv*, **215**, 25
 Predehl, P., Andriutschke, R., Arefiev, V., et al. 2021, *A&A*, **647**, A1
 Reid, B. A., Verde, L., Dolag, K., Matarrese, S., & Moscardini, L. 2010, *JCAP*, **07**, 013
 Scoccimarro, R., Hui, L., Manera, M., & Chan, K. C. 2012, *PhRvD*, **85**, 083002
 Sefusatti, E., Vale, C., Kadota, K., & Frieman, J. 2007, *ApJ*, **658**, 669
 Shao, H., Villaescusa-Navarro, F., Villanueva-Domingo, P., et al. 2022, *ApJ*, **944**, 27
 Slosar, A., Hirata, C., Seljak, U., Ho, S., & Padmanabhan, N. 2008, *JCAP*, **08**, 031
 Springel, V. 2005, *MNRAS*, **364**, 1105
 Sullivan, J. M., Prijon, T., & Seljak, U. 2023, *JCAP*, **08**, 004
 Takada, M., & Spergel, D. N. 2014, *MNRAS*, **441**, 2456
 Tinker, J. L., Robertson, B. E., Kravtsov, A. V., et al. 2010, *ApJ*, **724**, 878
 Uhlemann, C., Friedrich, O., Villaescusa-Navarro, F., Banerjee, A., & Codis, S. 2020, *MNRAS*, **495**, 4006
 Villaescusa-Navarro, F., Hahn, C., Massara, E., et al. 2020, *ApJS*, **250**, 2
 Villanueva-Domingo, P., & Villaescusa-Navarro, F. 2022, *ApJ*, **937**, 115
 Zaheer, M., Kottur, S., Ravanbakhsh, S., et al. 2017, arXiv:1703.06114