



Quijote-PNG: The Information Content of the Halo Power Spectrum and Bispectrum

William R Coulton¹ , Francisco Villaescusa-Navarro^{1,2} , Drew Jamieson³ , Marco Baldi^{4,5,6} , Gabriel Jung^{7,8} ,
Dionysios Karagiannis⁹, Michele Liguori^{7,8}, Licia Verde^{10,11} , and Benjamin D. Wandelt^{1,12}

¹ Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA; wrcoulton@flatironinstitute.org

² Department of Astrophysical Sciences, Princeton University, 4 Ivy Lane, Princeton, NJ 08544, USA

³ Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Straße 1, D-85748 Garching, Germany

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

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

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

⁷ Dipartimento di Fisica e Astronomia “G. Galilei,” Università degli Studi di Padova, via Marzolo 8, I-35131 Padova, Italy

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

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

¹⁰ ICREA, Pg. Lluís Companys 23, E-08010 Barcelona, Spain

¹¹ ICC, University of Barcelona, IEEC-UB, Martí i Franquès, 1, E-08028 Barcelona, Spain

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

Received 2022 August 23; revised 2022 November 23; accepted 2022 November 29; published 2023 February 7

Abstract

We investigate how much can be learnt about four types of primordial non-Gaussianity (PNG) from small-scale measurements of the halo field. Using the QUIJOTE-PNG simulations, we quantify the information content accessible with measurements of the halo power spectrum monopole and quadrupole, the matter power spectrum, the halo–matter cross spectrum, and the halo bispectrum monopole. This analysis is the first to include small, nonlinear scales, up to $k_{\text{max}} = 0.5 \text{ h Mpc}^{-1}$, and to explore whether these scales can break degeneracies with cosmological and nuisance parameters making use of thousands of N -body simulations. We perform all the halo measurements in redshift space with a single sample comprised of all halos with mass $> 3.2 \times 10^{13} h^{-1} M_{\odot}$. For local PNG, measurements of the scale-dependent bias effect from the power spectrum using sample variance cancellation provide significantly tighter constraints than measurements of the halo bispectrum. In this case measurements of the small scales add minimal additional constraining power. In contrast, the information on equilateral and orthogonal PNG is primarily accessible through the bispectrum. For these shapes, small-scale measurements increase the constraining power of the halo bispectrum by up to $4\times$, though the addition of scales beyond $k \approx 0.3 \text{ h Mpc}^{-1}$ improves constraints largely through reducing degeneracies between PNG and the other parameters. These degeneracies are even more powerfully mitigated through combining power spectrum and bispectrum measurements. However, even with combined measurements and small-scale information, equilateral non-Gaussianity remains highly degenerate with σ_8 and our bias model.

Unified Astronomy Thesaurus concepts: [Cosmology \(343\)](#); [Non-Gaussianity \(1116\)](#); [Observational cosmology \(1146\)](#); [Large-scale structure of the universe \(902\)](#)

1. Introduction

In the coming decade, a range of deep and wide photometric and spectroscopic galaxy surveys, e.g., the Rubin Observatory, DESI, SPHEREX, Euclid, and Roman (LSST Science Collaboration et al. 2009; Laureijs et al. 2011; Doré et al. 2014; Spergel et al. 2015; DESI Collaboration et al. 2016), will produce catalogs detailing the properties of tens of billions of galaxies. Through analyzing the spatial distribution of these galaxies, and how this distribution evolves over time, we will advance our understanding in many areas of physics, from constraining the properties of neutrinos to characterizing dark energy to clarifying a range of tensions hinted at in current cosmological observations (see, e.g., Dvorkin et al. 2019; Green et al. 2019; Slosar et al. 2019; Di Valentino et al. 2021, for overviews).

A particularly exciting prospect is what can be learnt about the primordial universe. Information on the primordial physics

can be extracted by studying the statistical properties of the primordial perturbations. An especially powerful avenue is through primordial non-Gaussianity (PNG), which characterizes how the distribution of the primordial perturbations deviates from Gaussianity. Through constraining PNG we help characterize the physics of the early universe, for example by learning about the field content, strength of interactions, and more (Maldacena 2003; Alishahiha et al. 2004; Creminelli & Zaldarriaga 2004; Senatore et al. 2010); see, e.g., Chen et al. (2007) and Meerburg et al. (2019) for overviews. This paper is the third in a series of papers (Coulton et al. 2023; Jung et al. 2022) focused on examining what we can learn about PNG from measurements of the large-scale structure of the universe (LSS) from nonlinear scales. This work focuses on an examination of assessing how much information is gained by combining binned bispectrum measurements of the halo field with power spectrum measurements. In our companion paper, G. Jung et al. (2022, in preparation), we complement this work with a modal bispectrum analysis at higher redshift, demonstrating our conclusions apply more broadly, and we develop a quasi-maximum-likelihood estimator, demonstrating near-optimal and unbiased compression of the bispectrum. These pair of



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

papers bring our analysis one step closer to modeling observations.

As cosmic microwave background (CMB) measurements (Komatsu & Spergel 2001; Hinshaw et al. 2013; Planck Collaboration IX 2020) have limited the primordial universe to be weakly non-Gaussian, it is instructive to search for primordial bispectra, the harmonic equivalent of three-point functions. The bispectrum is the lowest-order deviation from a Gaussian distribution (Taylor & Watts 2001; Babich 2005), and for many classes of theoretical models the bispectrum will be the most sensitive probe of PNG. Searches for primordial bispectra typically focus on searching for theoretically motivated bispectra, known as templates or shapes, as this allows for the most statistically powerful and physically informative inferences (Babich 2005; Komatsu et al. 2005; Yadav et al. 2007). In this work, we focus on four shapes of the primordial bispectrum: *local*, *equilateral*, *orthogonal-LSS* and *orthogonal-CMB* (the last two are two approximations for a bispectrum that arises in the effective field theory of inflation; Cheung et al. 2008a, 2008b; Senatore et al. 2010). We refer the reader to Achúcarro et al. (2022) for a recent overview of the theoretical and observational status of PNG.

Biased tracers of the LSS, such as galaxies or halos, are potentially very powerful probes of PNG as they exhibit a novel signature of PNG, a large-scale scale dependence in the bias, that is neither seen in the matter field nor generated by many other physical processes.¹³ Specifically, Dalal et al. (2008) found that in the presence of *local* PNG the bias of halos on large scales, which otherwise is expected to be approximately constant, scales as $\sim 1/k^2$. These results were subsequently expanded to show that other types of PNG can also induce scale-dependent biases, with the degree of the power-law behavior determined by the squeezed limit of the primordial bispectrum, $B(k_1, k, k)$, with $k \gg k_1$ (Matarrese & Verde 2008; Schmidt & Kamionkowski 2010). This feature provides new observational avenues for constraining PNG as measurements of the power spectrum can now be used to constrain PNG. The power of this can be dramatically enhanced through the use of sample variance cancellation techniques, which in principle allow types of PNG with scale-dependent biases to be measured to a level limited only by the shot-noise level (Seljak 2009; Castorina et al. 2018; Chan et al. 2019).

Given the prominent $\sim 1/k^2$ feature of *local* non-Gaussianity, several groups have attempted to measure this using measurements of quasars and from spectroscopic galaxy surveys, such as the Baryon Oscillation Spectroscopic Survey (BOSS; Slosar et al. 2008; Ross et al. 2013; Leistedt et al. 2014; Mueller et al. 2021; Cabass et al. 2022a; D’Amico et al. 2022b).¹⁴ While the best of these constraints (Cabass et al. 2022a; D’Amico et al. 2022b) is not yet competitive with CMB-based constraints (Planck Collaboration IX 2020), constraints from a broad range of surveys are forecast to provide significantly tighter constraints on local PNG that should surpass CMB-based constraints (see Figure 5 of Achúcarro et al. 2022, for a compendium of forecasts). For local PNG, future possible detections will likely not be limited by sample

variance but rather the shot noise of the samples and the difficulty of understanding large-scale systematic effects (Ross et al. 2011; Pullen & Hirata 2013; Rezaie et al. 2021). Note, however, that as has been highlighted in Barreira (2020, 2022a, 2022b), converting any detection of PNG to a specific amplitude of the primordial bispectrum may be challenging without advances in our understanding of galaxy formation. The importance of this in ruling out single field inflationary models depends on the exact value predicted by single field inflation; whether it is zero or slow-roll suppressed continues to be subject to a theoretical discussion (Pajer et al. 2013; Matarrese et al. 2021).

On the other hand, measuring other interesting shapes of PNG, such as *equilateral* or *orthogonal*, which do not generate strong scale-dependent biases, is generally significantly more challenging for LSS measurements than equivalent CMB measurements. These constraints are driven by measurements of the LSS bispectrum and are highly challenging as the late-time nonlinear evolution generates similar bispectra to the primordial ones (see, e.g., Gil-Marín et al. 2015a, 2015b, 2017; Slepian et al. 2017a, 2017b, for details of the measurements of nonlinear bispectrum with BOSS). Recently there have been several new developments. First Cabass et al. (2022b) and D’Amico et al. (2022b) have used effective field theory methods to robustly produce constraints on these types of PNG from measurements of the BOSS bispectrum. However, just applying these methods to larger volumes will likely be insufficient to vastly improve upon CMB constraints and, as is discussed in Philcox et al. (2022), pushing these theoretical models deeper into the nonlinear regime is highly complex. A second avenue, presented in Baumann & Green (2022), notes that, given the inherently different physics of the processes generating these effects, it may be possible to use field-level analyses or higher-order correlation functions to overcome these degeneracies.

Our work seeks to address two questions. First, we aim to quantify how much information on *local*, *equilateral*, and *orthogonal* non-Gaussianity can be gained by extending redshift-space measurements of the halo power spectrum and bispectrum to small, nonlinear scales. To do this we use the suite of QUIJOTE-PNG simulations presented in Coulton et al. (2023). By using simulations instead of perturbative methods we have an accurate (dark-matter-only) theoretical model for the nonlinear regime, and so can estimate the information content available in smaller scales. While our approach imposes several assumptions, for example on the form of the halo bias, it allows us to assess an upper limit on the information content to guide future works. Second, we are interested in assessing whether there is more information about local PNG beyond that accessible to power spectrum measurements of the scale-dependent bias. de Putter (2018) first investigated this question analytically, finding that for samples with moderate or low sample number densities there is minimal extra information beyond scale-dependent bias. Our simulations allow us to test this result and relax the perturbative modeling assumptions. This work builds on the first papers in this series (Coulton et al. 2023; Jung et al. 2022), where we generated a large suite of simulations containing PNG, developed and validated optimal compressed estimators, and used a range of methods to assess the information accessible with power spectrum and bispectrum measurements of the matter field. Our work is part of a large community interest in

¹³ Neutrinos also lead to a scale-dependent bias, e.g., Chiang et al. (2018). However, this can be distinguished due to the different k scaling.

¹⁴ Note that other complementary methods that use the mass function (Lucchin & Matarrese 1988; Matarrese et al. 2000; Scoccimarro et al. 2004; Shandera et al. 2013; Sabti et al. 2021) have also been used to constrain *local* PNG.

using small scales to constrain PNG (e.g., Uhlemann et al. 2018; Friedrich et al. 2020; Biagetti et al. 2021, 2022; Andrews et al. 2022; Giri et al. 2022).

This paper is structured as follows. In Section 2 we outline the primordial bispectra considered in this work and briefly review scale-dependent bias effect, and in Section 3 we describe the simulations used in this work. Then, in Section 4 we overview the implementations of our power spectrum and bispectrum measurements, presenting the resulting measurements subsequently in Section 5. Next, in Section 6 we describe the details of our Fisher forecast setup, including two methods we use to enhance the robustness of our results. Finally, we explore the information content of our measurements in Section 7, before presenting our conclusions in Section 8. In Appendix A we examine the stability of the derivatives estimates and, finding the standard methods to be unconverged, implement two methods that mitigate this issue, thereby allowing robust constraints. In Appendix B we examine the stability of our results with respect to variations in the number of simulations used to estimate the covariance matrix.

2. Primordial Non-Gaussianity and Scale-dependent Bias

The focus of this paper is primordial bispectra, defined as

$$\begin{aligned} &\langle \Phi(\mathbf{k}_1)\Phi(\mathbf{k}_2)\Phi(\mathbf{k}_3) \rangle \\ &= (2\pi)^3 \delta^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B^\Phi(k_1, k_2, k_3), \end{aligned} \quad (1)$$

where $\Phi(\mathbf{k})$ is the primordial potential. In this work we consider the same four types of primordial bispectrum simulated and discussed in Coulton et al. (2023). In the following, we define the relevant primordial bispectra and refer the reader to Chen's (2010), Meerburg et al.'s (2019), and Achúcarro et al.'s (2022) reviews on the physics of PNG and to Coulton et al. (2023) for more details on the bispectra considered here.

We consider four shapes: the *local* shape, with bispectrum

$$B_\Phi^{\text{Local}}(k_1, k_2, k_3) = 2f_{\text{NL}}^{\text{Local}} P_\Phi(k_1)P_\Phi(k_2) + 2 \text{ perm.}; \quad (2)$$

the *equilateral* shape, with bispectrum

$$\begin{aligned} B_\Phi^{\text{Equil.}}(k_1, k_2, k_3) &= 6f_{\text{NL}}^{\text{Equil.}} [-P_\Phi(k_1)P_\Phi(k_2) + 2 \text{ perm.} \\ &- 2(P_\Phi(k_1)P_\Phi(k_2)P_\Phi(k_3))^{\frac{2}{3}} + P_\Phi(k_1)^{\frac{1}{3}}P_\Phi(k_2)^{\frac{2}{3}}P_\Phi(k_3) \\ &+ 5 \text{ perm.}]; \end{aligned} \quad (3)$$

and two approximations to the *orthogonal* bispectrum, the *orthogonal-LSS*, with bispectrum

$$\begin{aligned} B_\Phi^{\text{Orth-LSS}}(k_1, k_2, k_3) &= 6f_{\text{NL}}^{\text{Orth-LSS}} (P_\Phi(k_1)P_\Phi(k_2)P_\Phi(k_3))^{\frac{2}{3}} \\ &\times \left[-\left(1 + \frac{9p}{27}\right) \frac{k_3^2}{k_1 k_2} + 2 \text{ perms} + \left(1 + \frac{15p}{27}\right) \frac{k_1}{k_3} \right. \\ &+ 5 \text{ perms} - \left(2 + \frac{60p}{27}\right) \\ &+ \frac{p}{27} \frac{k_1^4}{k_2^2 k_3^2} + 2 \text{ perms} - \frac{20p}{27} \frac{k_1 k_2}{k_3^2} + 2 \text{ perms} \\ &\left. - \frac{6p}{27} \frac{k_1^3}{k_2 k_3^2} + 5 \text{ perms} + \frac{15p}{27} \frac{k_1^2}{k_3^2} + 5 \text{ perms} \right], \end{aligned} \quad (4)$$

and the *orthogonal-CMB*, with bispectrum

$$\begin{aligned} B_\Phi^{\text{Orth-CMB}}(k_1, k_2, k_3) &= 6f_{\text{NL}}^{\text{Orth-CMB}} [-3P_\Phi(k_1)P_\Phi(k_2) \\ &+ 2 \text{ perm.} - 8(P_\Phi(k_1)P_\Phi(k_2)P_\Phi(k_3))^{\frac{2}{3}} \\ &+ 3P_\Phi(k_1)^{\frac{1}{3}}P_\Phi(k_2)^{\frac{2}{3}}P_\Phi(k_3) + 5 \text{ perm.}], \end{aligned} \quad (5)$$

where

$$p = \frac{27}{-21 + \frac{743}{7(20\pi^2 - 193)}}, \quad (6)$$

where $P_\Phi(k_i)$ is the primordial power spectrum and f_{NL}^X is the amplitude of non-Gaussianity for shape X . The *orthogonal-LSS* bispectrum is the more physically motivated shape, and so we primarily discuss this orthogonal shape; see Senatore et al. (2010) for a detailed discussion. However, the *orthogonal-CMB* shape exhibits interesting novel features that we also discuss.

Here we briefly give an overview of the scale-dependent bias effect; we refer the reader to Slosar et al. (2008), Matarrese & Verde (2008), Desjacques et al. (2011), and Schmidt & Kamionkowski (2010) for more detailed descriptions. The halo density contrast, $\delta_h(\mathbf{k}|M, z)$, is defined as

$$\delta_h(\mathbf{x}|M, z) = \frac{n_h(\mathbf{x}|M, z)}{\bar{n}_h(M, z)} - 1, \quad (7)$$

where $n_h(\mathbf{x}|M, z)$ is the halo counts of halos with mass M at redshift z and $\bar{n}_h(M, z)$ is its mean. In the absence of PNG and on the largest scales, the halo overdensity is related to the matter overdensity, $\delta_m(\mathbf{k})$, by

$$\delta_h(\mathbf{k}|M, z) = b_1(M, z)\delta_m(\mathbf{k}, z) + \epsilon(z), \quad (8)$$

where $b_1 = \frac{1}{\bar{n}_h} \frac{\partial n_h}{\partial \delta_m}$ characterizes the response of the halo counts to the large-scale matter fluctuation and $\epsilon(z)$ is a shot-noise term from the discrete nature of the halos.

PNG can introduce a coupling between large- and small-scale modes. The most well-known example is *local* PNG, in which small-scale modes, k_s , of the primordial potential, are modulated by long-wavelength modes, k_l (e.g., Komatsu & Spergel 2001):

$$\Phi(\mathbf{k}_s) \approx \Phi(\mathbf{k}_s)(1 + f_{\text{NL}}^{\text{Local}} \Phi(\mathbf{k}_l)). \quad (9)$$

This coupling means that the halo overdensity is no longer just a function of the large-scale matter overdensity, but also the primordial potential at that point. Hence,

$$\begin{aligned} \delta_h(\mathbf{k}|M, z) &= b_1(M, z)\delta_m(\mathbf{k}, z) + b_\phi(M, k, z)\alpha(k, z)\Phi(\mathbf{k}) + n(\mathbf{k}, z) \\ &= [b_1(M, z) + b_\phi(M, k, z)]\delta_m(\mathbf{k}, z) + \epsilon(z), \end{aligned} \quad (10)$$

having used in the second line the relation between the late-time matter overdensity and the primordial potential, $\delta_m(k, z) = \alpha(k, z)\Phi(k)$, where $\alpha(k, z)$ is the matter transfer function. (Note the transfer function in the first line is to simplify the definition of b_ϕ .) Unlike b_1 , the new bias term, b_ϕ , depends on scale.

For halos we can use the peak-background split method, assuming a universal mass function (Press & Schechter 1974; Sheth et al. 2001) to compute the relation between b_Φ , b_1 , and the primordial bispectrum $B_\Phi(k_1, k_2, k_3)$. In Schmidt &

Table 1
The Key Properties of the 15,000 Simulations Used to Compute Our Covariance Matrices

Ω_m	Ω_Λ	Ω_b	σ_8	h	n_s	Box Size (Mpc h ⁻¹)	$N_{\text{particles}}^{\frac{1}{3}}$	ICs	Mass Resolution ($M_\odot h^{-1}$)
0.3175	0.6825	0.049	0.834	0.6711	0.9624	1000	512	2LPTPNG	6.56×10^{11}

Note. Note these simulations have Gaussian primordial initial conditions, i.e., $f_{\text{NL}}^X = 0$ for all shapes, X . See Villaescusa-Navarro et al. (2020) for more details on the simulations.

Kamionkowski (2010) this was found to be

$$b_\Phi(M, z) = \frac{2f_{\text{NL}}\delta_c}{\alpha(k, z)}(b_1 - 1) \frac{\sigma_W^2(k)}{\sigma_M^2}, \quad (11)$$

where $\delta_c \approx 1.686$ is the spherical collapse threshold, and σ_M^2 is the variance of the top-hat smooth matter field:

$$\sigma_M^2 = \int \frac{d^3k'}{(2\pi)^3} W^2(k') P_\delta(k'), \quad (12)$$

and σ_W^2 is the smoothed variance in the presence of PNG:

$$\sigma_W^2 = \int \frac{d^3k'}{(2\pi)^3} \frac{W^2(k') P_\delta(k') B_\Phi(\mathbf{k}, \mathbf{k}', \mathbf{k} - \mathbf{k}')}{2f_{\text{NL}}(P_\Phi(k) P_\Phi(k') + 2 \text{ perms.})}. \quad (13)$$

In these last equations, $W_M(k)$ is the Fourier transform of the top hat, the radius of which is set by the halo mass such that $R^3 = \frac{3\bar{p}}{4\pi} M$ and $P_\delta(k)$ is the matter power spectrum.

For the case of *local* PNG $\sigma_W^2(k) \sim \text{constant}$ and, given that on large scales $\alpha(k, z) \sim k^2$, this means that $b_\Phi \propto 1/k^2$. For *orthogonal-CMB* $\sigma_W^2(k) \sim k$, which leads to $b_\Phi \propto 1/k$, while for *equilateral* and *orthogonal-LSS* we have $\sigma_W^2(k) \sim k^2$ and so $b_\Phi \propto 1$. Note that as the $f_{\text{NL}}^{\text{local}}$ correction to the power spectrum, on currently observable scales, is constrained to be small, the dominant $f_{\text{NL}}^{\text{local}}$ correction to the halo power spectrum is the b_Φ (i.e., the $1/k^2$ term) and not the b_Φ^2 ($1/k^4$) term. In fact, the b_Φ^2 term only becomes important once a detection of $f_{\text{NL}}^{\text{local}}$ is made.

A particularly powerful technique to boost the constraints using the scale-dependent bias effect is sample variance cancellation (Seljak 2009). The intuition underpinning this method is as follows: measurements of scale-dependent bias are most powerful on the largest scales, however on these scales only a very small number of modes are measurable. This means that the sample variance on such measurements is very large and limits the constraining power on PNG. However, one can utilize the fact that the bias coefficient relating the halo/galaxy field, δ_h , to the underlying matter field, δ_m , is deterministic, i.e., mode-by-mode $\delta_h(\mathbf{k}) = b(k)\delta_m(\mathbf{k}) + n(\mathbf{k})$ (as in Equation (10)). Thus, if we can access the information contained in the ratio of halo and matter fields, we can perform a measurement in which the stochasticity from the matter field cancels: this is powerful as this means the sample variance terms cancel, leaving only the shot-noise terms $n(\mathbf{k})$. In the idealized zero shot-noise case, this means that the bias coefficient could be measured arbitrarily well and thus PNG could be detected with arbitrarily high significance. In practice this technique will likely be performed using multiple tracers with different biases (e.g., Seljak 2009) as the three-dimensional (3D) matter field is not directly observable. However, a range of works have proposed using CMB lensing (Schmittfull & Seljak 2018), kinetic Sunyaev-Zel'dovich measurements

(Münchmeyer et al. 2019), and other probes to perform sample variance cancellation with proxies for the 3D matter field.

There is one complication with using measurements of scale-dependent bias to constrain PNG: the analytic form for the scale-dependent bias coefficient, b_Φ in Equation (11), is only valid for biased tracers with universal mass functions. This is not true for observables such as quasars or galaxies, as has been explicitly explored in, e.g., Slosar et al. (2008), Reid et al. (2010), and Barreira (2020, 2022b). A conservative approach to this issue would be to marginalize over this parameter, as is discussed in Cabass et al. (2022a) and Barreira (2022b), however the prior is very important and is currently poorly understood. Note this marginalization is necessary as breaking this degeneracy from the data itself is very challenging, as is discussed in, e.g., Barreira (2022b). In applying a simulation-based approach in a self-consistent way, we implicitly assume that we know the bias relation, Equation (11), perfectly. Our results should therefore be understood as theoretical limits. The multiplicative uncertainty in the bias relation for realistic tracers can be folded in after the fact. We discuss this further in the conclusions, but leave this extension to follow-up studies.

3. Simulations

In this work we use the QUIJOTE-PNG N -body simulations presented in Villaescusa-Navarro et al. (2020) and Coulton et al. (2023), and refer the reader to these papers for a detailed description. The simulations were run with GADGET-III (Springel 2005) and the halos were identified using the friends-of-friends algorithm (Davis et al. 1985). In our analysis we consider a single population of halos that is comprised of all halos with $M > 3.2 \times 10^{13} h^{-1} M_\odot$, which corresponds to a mean tracer density $\bar{n}(z=0) = 1.55 \times 10^{-4} h^3 \text{Mpc}^{-3}$. We use the 15,000 simulations run at the fiducial cosmology, with the parameters summarized in Table 1, and the simulations designed for estimating derivatives, summarized in Table 2, with each parameter perturbed above and below its fiducial value. There are 500 simulations per parameter for the positive perturbation and 500 for the negative perturbation. Note that in this work we restrict our analysis to the parameters in Table 2, rather than including all the parameters varied in the QUIJOTE suite, which include w , Ω_b , $\sum m_\nu$, due to the convergence challenges discussed in Section 6. This does not significantly limit our results as the other parameters are, with the exception of Ω_b , extensions to Lambda cold dark matter (Λ CDM) that would likely be constrained separately, while Ω_b would be far better measured by the CMB (Planck Collaboration VI 2020).

The statistical probes, described in Section 4, are all Fourier space statistics, and so to measure these we need to preprocess the particle data and halo catalogs. First we assign them to a grid with 512^3 voxels using the ‘‘cloud-in-cell’’ (CIC) mass-assignment scheme (Hockney & Eastwood 1981). To compute the overdensity field (δ_m and δ_h for the matter and halos,

Table 2

The Parameter Values of the Simulations Used to Compute the Numerical Derivatives

Parameter	Lower Value	Upper Value
$\Omega_{m,0}$	0.3075	3275
h	0.6511	0.6911
n_s	0.9424	0.9824
σ_8	0.819	0.849
$f_{\text{NL}}^{\text{Local}}$	-100	+100
$f_{\text{NL}}^{\text{Equil}}$	-100	+100
$f_{\text{NL}}^{\text{Orth-LSS}}$	-100	+100
$f_{\text{NL}}^{\text{Orth-CMB}}$	-100	+100
δ_b	-0.035	0.035
$M_{\text{min}}(M_{\odot} h^{-1})$	3.1×10^{13}	3.3×10^{13}

Note. For each desired derivative, we use 500 simulations with the parameter perturbed above and 500 perturbed below its fiducial value. Note δ_b is used to compute the super sample covariance and M_{min} is used to compute the derivative with respect to the halo bias.

respectively) we subtract the mean value from the field and then normalize by the mean. Note this mean is computed for each realization, an important subtlety for the halo field. The Nyquist frequency of the grid is 1.61 h Mpc^{-1} , which is much larger than the maximum scale used in our analysis, $k = 0.5 \text{ h Mpc}^{-1}$, and thus dramatically reduces the effects of aliasing; see, e.g., Sefusatti et al. (2016) or Appendix A of Foreman et al. (2020), for a discussion of aliasing effects on the bispectrum.

For the halo field we move from real space to redshift space by displacing the halo positions using the line-of-sight component of the halo velocity, which is computed as a mass weighting of the halo's particles. We are free to choose the direction of the line of sight and we make three grids by projecting along the x , y , and z axes. While these three grids share the same real-space configuration, they are not perfectly correlated in redshift space, especially on small scales, and so we can use them to further reduce the Monte Carlo uncertainty. We then Fourier transform the grid and deconvolve the CIC window function (see Jing 2005; Sefusatti et al. 2016) and then measure our statistical probes.

The halo catalogs used in this work will be available at <https://quijote-simulations.readthedocs.io/en/latest/png.html> on acceptance of the paper.

4. Statistical Probes

In this work we combine four statistics: the matter power spectrum, the halo power spectrum monopole and quadrupole, the halo–matter cross power spectrum and the bispectrum monopole. The combined analysis of the matter, halo–matter, and halo power spectrum will allow us to exploit the sample variance cancellation method described in Section 3. In this section we describe how we compute these statistics. The matter field statistics are computed as in Coulton et al. (2023) and so here we described the computation of the halo and the halo–matter cross terms.

4.1. Halo and Matter Power Spectrum

The halo redshift-space power spectrum depends both on the magnitude of the wavevector and the angle to the line of sight,

$\mu = \mathbf{k} \cdot \mathbf{n}$, hence $P^{hh}(k, \mu)$. To extract information available in redshift space we compute the power spectrum monopole, $P_0^{hh}(k)$, and quadrupole, $P_2^{hh}(k)$, defined as

$$P_\ell^{hh}(k) = \frac{2\ell + 1}{2} \int d\mu \mathcal{L}_\ell(\mu) P^{hh}(k, \mu), \quad (14)$$

where $\mathcal{L}_\ell(\mu)$ are the Legendre polynomials.

In an analogous manner to the matter power we estimate these as

$$\hat{P}_\ell^{hh}(k_i) = \frac{2\ell + 1}{2N_i'} \sum_{k_i - \Delta/2 < |\mathbf{k}| \leq k_i + \Delta/2} \mathcal{L}_\ell(\mu) \delta^h(\mathbf{k}) \delta^{h*}(\mathbf{k}), \quad (15)$$

where N_i' is the estimator normalization and we use the same binning as for the matter power spectrum. We use the public Pylians3 code¹⁵ to compute these quantities. Additionally, we subtract the shot-noise term, $1/\bar{n}$, from the halo power spectrum monopole.

Likewise, we estimate the matter–halo cross power spectrum monopole as

$$\begin{aligned} \hat{P}_\ell^{mh}(k_i) &= \frac{1}{2N_i'} \\ &\times \sum_{k_i - \Delta/2 < |\mathbf{k}| \leq k_i + \Delta/2} (\delta^m(\mathbf{k}) \delta^{h*}(\mathbf{k}) + \delta^{m*}(\mathbf{k}) \delta^h(\mathbf{k})). \end{aligned} \quad (16)$$

As is discussed in Section 6, we limit the wavenumbers of the matter power spectrum and matter–halo cross power spectrum to $k < 0.1 \text{ h Mpc}^{-1}$.

4.2. Halo Bispectrum

The halo bispectrum is defined as

$$\begin{aligned} &\langle \delta^h(\mathbf{k}_1) \delta^h(\mathbf{k}_2) \delta^h(\mathbf{k}_3) \rangle \\ &= (2\pi)^3 \delta^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3). \end{aligned} \quad (17)$$

The redshift-space distortions mean that, like the halo power spectrum, the halo bispectrum also depends on the orientation of the modes relative to the line of sight. In this work we consider only the bispectrum monopole, which is defined simply averaging over all angles. Thus we use the same estimator as for the matter bispectrum:

$$\begin{aligned} \hat{B}_0^{hhh}(k_i, k_j, k_k) &= \frac{1}{N_{ijk}} \sum_{k_i - 1 < |\mathbf{k}_1| \leq k_i, k_j - 1 < |\mathbf{k}_2| \leq k_j, k_k - 1 < |\mathbf{k}_3| \leq k_k} \\ &\times \delta^h(\mathbf{k}_1) \delta^h(\mathbf{k}_2) \delta^h(\mathbf{k}_3) (2\pi)^3 \delta^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3). \end{aligned} \quad (18)$$

The halo bispectrum monopole contains a contribution from the discrete nature of the halo field. This term is known as the shot-noise term, and in this work we remove it. The shot-noise term is computed as

$$\begin{aligned} \hat{B}_0^{SN}(k_i, k_j, k_k) &= \frac{1}{N_{ijk}} \sum (2\pi)^3 \delta^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) \\ &\times \left[(P_0^{hh}(k_i) + P_0^{hh}(k_j) + P_0^{hh}(k_k)) \frac{1}{\bar{n}} + \frac{1}{\bar{n}^2} \right], \end{aligned} \quad (19)$$

where the implicit limits are the same as for the bispectrum estimator and \bar{n} is the mean halo number density. We use the halo power spectrum and mean halo number density as

¹⁵ <https://github.com/franciscovillaescusa/Pylians3>

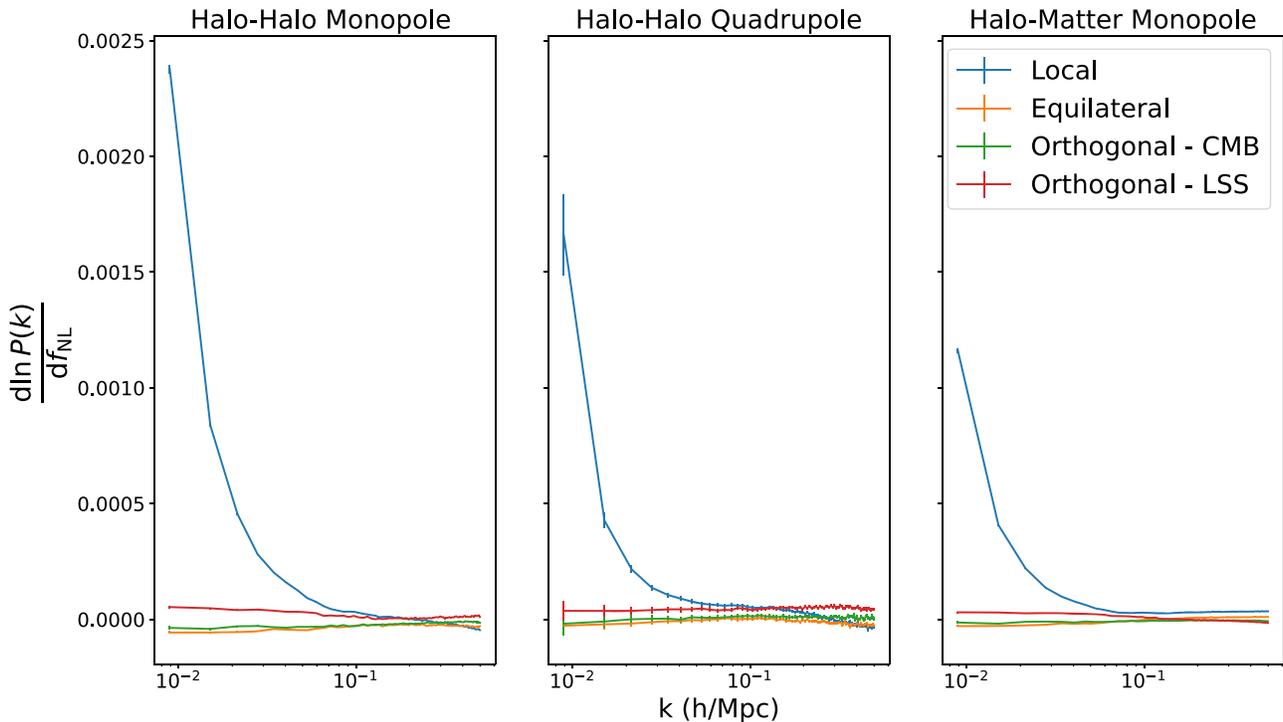


Figure 1. The derivatives of the halo monopole, halo quadrupole, and halo–matter monopole power spectra with respect to the amplitudes of the four types of primordial non-Gaussianity. These are computed at $z = 0.0$ and the halo field contains all halos with $M > 3.2 \times 10^{13} M_{\odot} h^{-1}$. The error bars denote the error on the mean. The large response of local non-Gaussianity arises due to the well-known, scale-dependent bias effect (see Section 2). The response of the power spectra to other types of non-Gaussianity is significantly weaker and largely featureless.

measured on that realization to compute this term. This removes the sample variance arising from the Poisson term; note, however, it does not remove the shot-noise contribution to the bispectrum variance.

5. Power Spectrum and Bispectrum Measurements

In this section, we explore how the power spectrum and bispectrum statistics are impacted by PNG using the simulations summarized in Table 2.

5.1. Power Spectrum

In Figure 1 we plot the derivative of the halo power spectrum with respect to the four shapes of PNG. For local, the scale-dependent bias effect, discussed Section 2, leads to an enhancement of the large-scale power spectrum, as expected on these scales as k^{-2} (see, e.g., Schmidt & Kamionkowski 2010, for an analytical description of this behavior). No scale-dependent bias effect is observed for the other shapes of PNG. For equilateral and orthogonal-LSS this is expected and is consistent with previous works and theory (e.g., Scoccimarro et al. 2012; Sefusatti et al. 2012; Wagner & Verde 2012). However, the orthogonal-CMB PNG is expected to exhibit a scale-dependent bias that scales as k^{-1} .

This apparent inconsistency with expectations for the orthogonal-CMB case arises as, for the halo sample used in this work, the coefficient for the scale-dependent bias term, Equation (11), is very small. When considering other samples the coefficient can be much larger and then the scale-dependent bias can be seen. In Figure 2, we show the ratio of the halo–matter power spectrum to the matter power spectrum

for our original sample, with a minimum mass $M_{\min} = 3.2 \times 10^{13} M_{\odot} h^{-1}$, and a second sample with $M_{\min} = 1 \times 10^{14} M_{\odot} h^{-1}$. On large scales this ratio isolates the linear bias term. For the second sample we see a scale-dependent PNG bias contribution for both the local and orthogonal-CMB cases. The small correction seen for our original sample is in fact consistent with previous work: Schmidt & Kamionkowski (2010) found that for a sample of halos at $z = 0.0$ and $M = 1 \times 10^{13} M_{\odot} h^{-1}$ the scale-dependent bias for orthogonal-CMB PNG is $\sim 1\%$ of the linear halo bias and more than an order of magnitude smaller than the equivalent local shape bias. For such small values of b_{ϕ} the scale-dependent bias, on the scales considered here, is swamped by the changes to b_1 induced by PNG.

5.2. Halo Bispectrum

In Figure 3 we investigate the signatures of PNG in the bispectrum at $z = 0.0$. On the largest scales the bispectra are all very well measured, and we have tested that these are well described by a simple linear bias times the PNG contribution to the matter bispectrum; see Coulton et al. (2023) and Jung et al. (2022) for more details of the matter bispectrum. Moving to smaller scales we see that the halo bispectrum derivatives tend to scale as $\sim k^{-3}$; interestingly, this scaling is significantly weaker than the one of the matter bispectrum, which scales as $\sim k^{-4}$. These results are qualitatively consistent with the theoretical models investigated in, e.g., Sefusatti (2009) and Baldauf et al. (2011), as well as previous simulations (Sefusatti et al. 2012; Tellarini et al. 2016), though we push to smaller scales in this work. Second, we notice that the bispectra are very noisy, much more so than the matter bispectra. The high

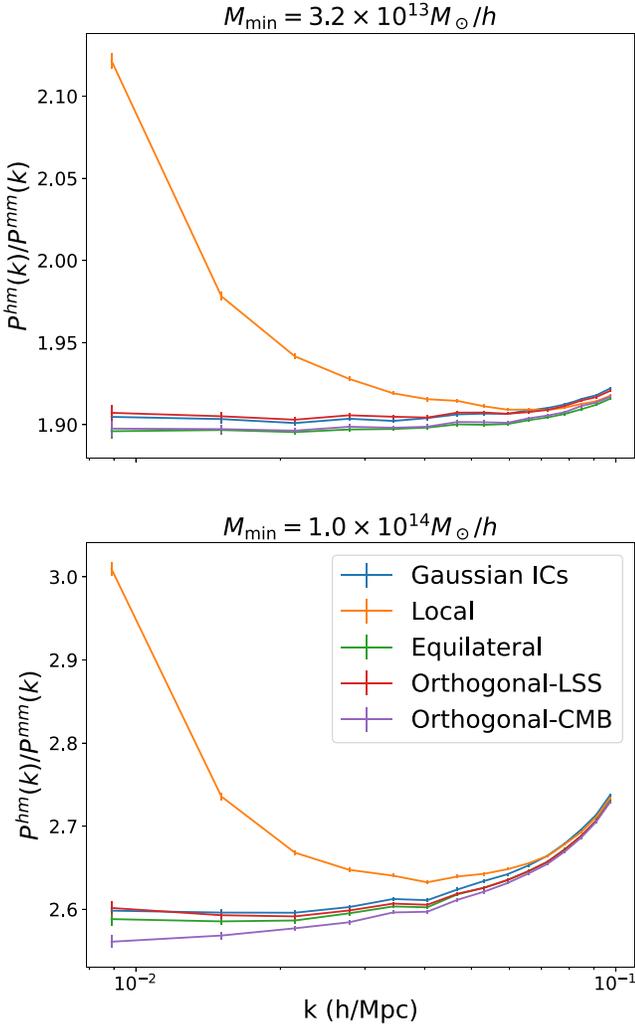


Figure 2. The average ratio of the halo–matter power spectrum to the matter–matter power spectrum for two halo mass limited samples at $z = 0.0$. The scale-dependent bias terms can be isolated through measurements of this ratio on large scales. These results show how the amplitude of the scale-dependent bias terms responds to changes in the halo sample. Interestingly, we only see evidence of the expected scale dependence in the orthogonal-CMB shape for the higher-mass sample; see the text for a detailed discussion.

levels of noise in the bispectrum derivatives, even after averaging over 500 simulations, leads to challenges when performing a Fisher forecast, and we explore the details of this in Section 6 and Appendix A.

To further understand the source of noise we investigate the properties of the bispectrum at the fiducial cosmology of QUIJOTE. In Figure 4 we compare the squeezed slice of the halo bispectrum with and without shot-noise subtraction to the matter bispectrum, scaled by the linear bias. Note that the linear bias was estimated by computing the ratio of the halo–matter cross power spectrum to the matter auto power spectrum. We see that on all scales the halo bispectrum without the shot-noise correction is significantly larger than the other bispectra, reaching a factor of ~ 10 on the smallest scales. This large shot noise acts as a source of effective noise in our analysis. When computing the derivatives shown in Figure 3 we use simulations with matched initial conditions: the amplitudes and phases of the initial conditions are identical and the only

difference is the sign of the PNG correction term; see Section IIIA in Coulton et al. (2023) for more details. For the matter field, using matched simulations means the effective noise terms cancel very effectively, leading to highly accurate, low-noise derivative measurements (note that this is further aided as noise in the matter field is inherently small on these scales). The halo field is more complicated as the slightly different evolution of the two simulations used to estimate the derivative ($f_{\text{NL}} = +100$ and $f_{\text{NL}} = -100$) leads to slight differences in the number and positions of the halos. These slight differences mean that the cancellation of the shot noise is not perfect and some of the large noise leaks in and masks the small bispectrum signals.

Note that once the shot noise has been subtracted, we find a good agreement between the scaled-matter bispectrum and the halo bispectrum on the largest scales. This is in agreement with perturbation theory results for the largest scales (Baldauf et al. 2011). On smaller scales we see large differences between the scaled-matter and shot-noise-subtracted bispectra: this is also expected, as moving to smaller scales higher-order bias terms become important, and beyond the nonlinear scale, $k \sim 0.08 \text{ Mpc } h^{-1}$, the perturbative description breaks down.

6. Fisher Forecasting and the Challenges of Convergence

The basic Fisher forecast methodology is very similar to Coulton et al. (2023) and so we outline the key steps and refer the reader to that work for more details.

We assume that the power spectra and bispectra likelihood can be approximated by a Gaussian distribution. Thus the Fisher information is given by

$$F_{IJ} = \frac{\partial \bar{O}(k_i)}{\partial \theta_I} C_{ij}^{-1} \frac{\partial \bar{O}(k_j)}{\partial \theta_J}, \quad (20)$$

where O is the data vector (composed of all or subsets of the power spectra and bispectrum measurements), θ_I are the parameters of interest, and C is the covariance matrix. The derivatives are estimated using central finite differences with a two-point stencil and the covariance matrix is estimated from the simulations.

We focus our analysis on the four PNG bispectrum amplitudes, $f_{\text{NL}}^{\text{Local}}$, $f_{\text{NL}}^{\text{Equil.}}$, $f_{\text{NL}}^{\text{Orth.-LSS}}$, and $f_{\text{NL}}^{\text{Orth.-CMB}}$; four cosmological parameters, Ω_m , n_s , σ_8 , and h ; and one parameter controlling the halo bias, M_{min} . The bias parameter, M_{min} , corresponds to variations in the halo catalog mass threshold, and to compute derivatives of this we generate and process in an identical manner catalogs with a minimum halo mass $3.1 \times 10^{13} M_{\odot} h^{-1}$ and $3.3 \times 10^{13} M_{\odot} h^{-1}$ (see, e.g., Hahn et al. 2020, for a discussion of this parameterization). This bias parameter is in our setup roughly equivalent to the standard linear bias model. While the bias model presented here is overly simplistic it is still useful as a first estimate of the importance of the bias terms and for estimating the total accessible information. A more thorough analysis of the impact of bias parameters, which could be implemented in our simulations by populating the halos with a halo occupation distribution (HOD) and investigating the constraints on the HOD parameters, is left to future work. It is expected that including more bias parameters will lead to further degradation of the constraints, as is discussed in D’Amico et al. (2022a) and Philcox et al. (2022).

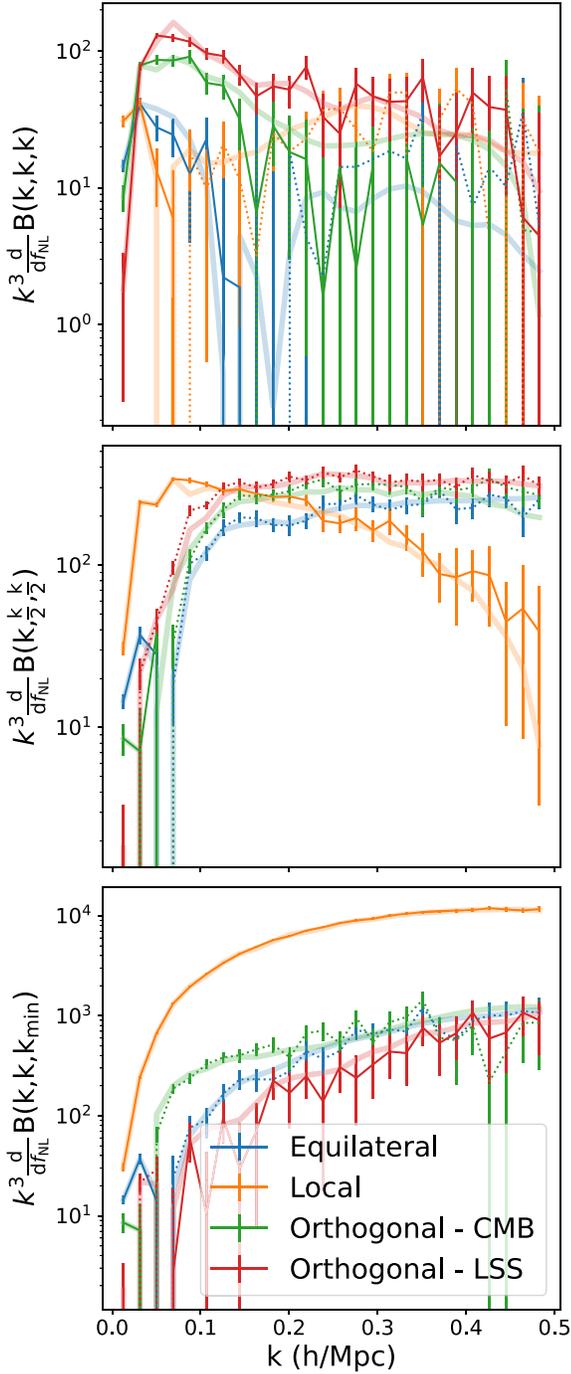


Figure 3. The response of the halo equilateral (top), folded (middle), and squeezed (bottom) bispectrum slices to the four different types of primordial non-Gaussianity at $z = 0$. The dotted lines denote regions where the bispectrum is negative and the error bars are the error on the mean of our simulations, which is equivalent to the error that would be measured with a volume of $\sim 500 \text{ Gpc}^3 h^{-3}$. At small scales, especially for the equilateral slice, the measurements are significantly impacted by noise, despite using 500 simulations with matched seeds. This large noise can bias our Fisher estimates, and in Appendix A we demonstrate two methods to mitigate this bias. The method used in the rest of the paper mitigates the noise by fitting a smooth function, via a Gaussian process, to the measurements. The thick shaded contour shows the resulting smoothed derivatives.

We use 14,500 simulations at the QUIJOTE fiducial cosmology to compute the covariance (the remaining 500 are used to compute the super sample covariance, SSC, terms). Unlike for the computation of the derivatives, we only include

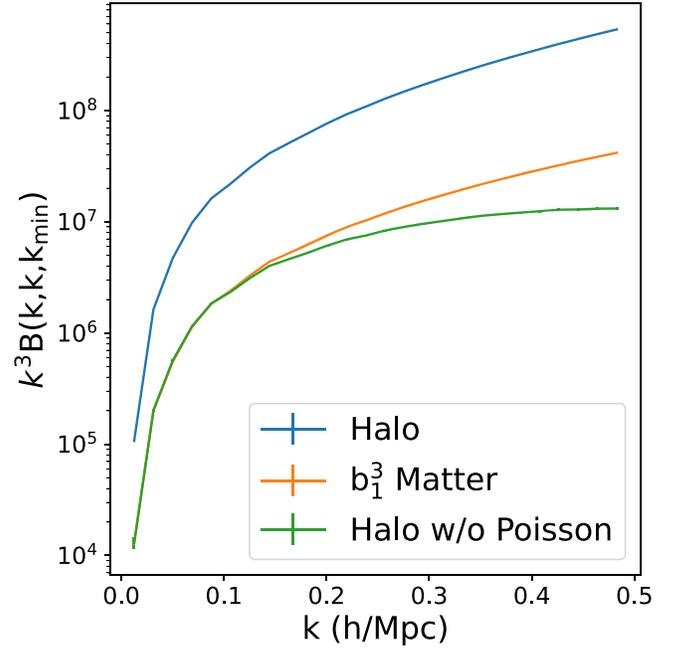


Figure 4. The squeezed bispectrum computed at the fiducial cosmology (Table 1). We compare the halo bispectrum to the matter bispectrum rescaled by the linear bias, b_1 , and to the halo bispectrum once we have removed the shot-noise contribution. The large shot-noise contribution to the bispectrum leads to the large noise seen in our bispectrum derivatives; see Figure 3.

projections along one line of sight in the simulations used to estimate the covariance matrix. This is to avoid biasing the covariance matrix with correlations between the correlated lines of sight. In Section 7.3 we investigate the importance of the different contributions to the covariance matrix, which is composed of Gaussian contributions, non-Gaussian terms from the connected three-, four-, and six-point functions, and the SSC (see, e.g., Kayo et al. 2013; Chan & Blot 2017; Galdi et al. 2018; Sugiyama et al. 2020).

In this work our primary focus is to investigate the information of the halo field. As a consequence, we limit the information included in the matter field to only the largest scale modes $k < 0.1 \text{ h Mpc}^{-1}$. Through this choice we can assess the benefits of sample variance cancellation (through the matter field on the largest scales) as will be used in future surveys through multitracer analyses (Slosar 2009), with CMB lensing (Schmittfull & Seljak 2018) or with the kinetic Sunyaev-Zel'dovich effect (Münchmeyer et al. 2019). The k cut used here is larger than would be accessible through some of these methods, however this compensates for the small size of our box. While using multiple tracers is a powerful method to employ sample variance cancellation, we do not use that here as, given our mass resolution, it is difficult to construct a second tracer field from our halo catalog with an observationally relevant tracer density.

An important part of Fisher forecasts with numerically estimated covariance matrices and derivatives is that the Monte Carlo noise in these components is sufficiently small so that it does not bias the forecasts. In Appendix A we examine the stability of derivative estimates and in Appendix B we investigate the stability of the covariance matrix. We find that our derivatives are not converged, leading to biased forecasts when using the standard way to compute the Fisher matrix. To mitigate this issue we implement two complementary methods

to obtain stable and robust results. This is an important, though technical, aspect of our forecast, and we provide a detailed description in Appendix A. The bispectrum results presented in the remainder of the paper use the “smoothed” (Gaussian process) method to ensure robust forecasts.

7. Constraining Power of the Halo Field

In this section we explore the importance of degeneracies, small-scale information, and the different contributions to the covariance matrix when quantifying the amount of information on PNG embedded into the halo power spectrum and bispectrum. Our analysis is performed at $z = 0.0$ and uses modes up to $k_{\max} = 0.5 \text{ h Mpc}^{-1}$ except for quantities involving the matter field, for which we use $k_{\max} = 0.1 \text{ h Mpc}^{-1}$. As discussed in Section 6, we use low k_{\max} measurements of the matter field to emulate observations that use the sample variance cancellation method and assess how much information they contain. The results in this section primarily focus on the local, equilateral, and orthogonal-LSS shapes as these most closely approximate physically generated bispectra.

7.1. Parameter Degeneracies

In Figure 5 we explore the parameter constraints obtainable from the bispectrum and power spectrum measurements. First, focusing on the power spectrum constraints, we see that, compared to the information available in measurements of the matter field power spectrum discussed in Coulton et al. (2023) and Jung et al. (2022), the halo power spectrum contains much more information on PNG. This is largely due to the scale-dependent bias effects. The information in scale-dependent bias can be more efficiently extracted with sample variance cancellation techniques, as originally proposed in Seljak (2009). In our analysis, we attempt to incorporate this effect by using both the monopole, matter monopole, and halo-matter cross power spectra. It can be seen that including these other power spectrum measurements vastly improves the constraints on PNG, particularly for local non-Gaussianity.

Moving to the bispectrum constraints, we see that the constraints on local PNG are worse for the bispectrum than for the power spectrum, highlighting the importance of scale-dependent bias measurements. As expected, for the equilateral and orthogonal shapes the bispectrum measurements are significantly more informative than the power spectrum measurements. We see that there are strong parameter degeneracies for the bispectrum measurements of PNG, especially with the effective bias parameter, M_{\min} , and the amplitude of clustering, σ_8 . The degradation of the constraints due to marginalization can be seen more clearly in Figure 6(b), where it can be seen that the bispectrum marginalization increases the constraints by more than 100%.

Finally, we see that combining the power spectrum and bispectrum measurements provides significant improvements for all the parameters except local PNG, where the gains are more modest. These improvements come from two related effects: the power spectrum and bispectrum exhibit different degeneracies between the PNG parameters and the Λ CDM parameters; and, second, through the complementary nature of the two probes the subsets of the parameters are better constrained by different probes, and when combined this improved constraining power propagates to the other parameters due to the strong degeneracies.

7.2. Are Small-scale Measurements Informative?

To further understand the information content, it is useful to examine which scales contribute to the constraining power. In Figure 6(a) we explore the constraining power in measurements of the halo field only as a function of scale. For the unmarginalized constraints, those obtained when measuring only the amplitude of one type of PNG, the information in the power spectrum changes very slowly with the inclusion of new modes, whereas the bispectrum constraints improve rapidly with k_{\max} up to $k_{\max} \sim 0.3 \text{ h Mpc}^{-1}$, after which the gains are more modest. This shows that, like the matter field (Coulton et al. 2023; Jung et al. 2022), there are minimal gains from pushing to very small scales—at least up to the maximum scale considered here. When considering marginalized constraints the situation is somewhat different: the power spectrum and bispectrum constraints both show significant improvements up to $k_{\max} \sim 0.4 \text{ h Mpc}^{-1}$. This improvement arises as the inclusion of these scales reduces the degeneracies and improves the constraints on the degenerate parameters. Thus, when we consider the joint analysis of the power spectrum and bispectrum, we see that the complementarity of the probes helps reduce the degradation of the constraints from marginalization. Interestingly, we find the joint marginalized constraints for local and orthogonal-LSS show minimal degradation from marginalization, and thus the joint case mirrors the unmarginalized case where there is minimal information in the small-scale modes $k \gtrsim 0.3 \text{ h Mpc}^{-1}$.

Next we repeat this analysis for the case where we include large-scale ($k < 0.1 \text{ h Mpc}^{-1}$) measurements of the matter auto and matter-halo cross power spectra. By comparing with Figure 6(a), we see the power of sample variance cancellation (which occurs on large scales through the inclusion of the matter field) as the constraints improve by a factor of ~ 5 . Note that, in the unmarginalized case, sample variance cancellation improves the power spectrum constraint for equilateral, despite the absence of a scale-dependent bias feature. This occurs as equilateral non-Gaussianity leads to a small shift in the linear bias on large scales; see Figure 2. As we cannot predict the linear bias coefficient, this information cannot be accessed in practice. This effect is completely degenerate with halo bias so when we marginalize over the bias model the constraints dramatically degrade. The evolution with k_{\max} otherwise mirrors the halo-only analysis: the unmarginalized power spectrum constraints improve weakly with the inclusion of small-scale modes, while the marginalized case shows significant improvements as the small-scale modes, up to $k_{\max} \sim 0.3 \text{ h Mpc}^{-1}$, are included. Likewise, the joint analysis of both probes also shows only modest improvements with k_{\max} .

To contextualize these results, in Figure 7 we compare the information available in the joint power spectrum and bispectrum analysis to the information available in the primordial fields. There are several interesting features. First the constraints on local PNG on the largest scales are better in the late-time than in the primordial universe. This result is easily explained as the information encoded in the scale-dependent bias, which drives the constraints on the largest scales, actually encodes information from small scales: the formation of a halo is inherently controlled by these modes and the information of these small-scale modes is transferred to the halo bias. Note that we discuss the challenging issue of

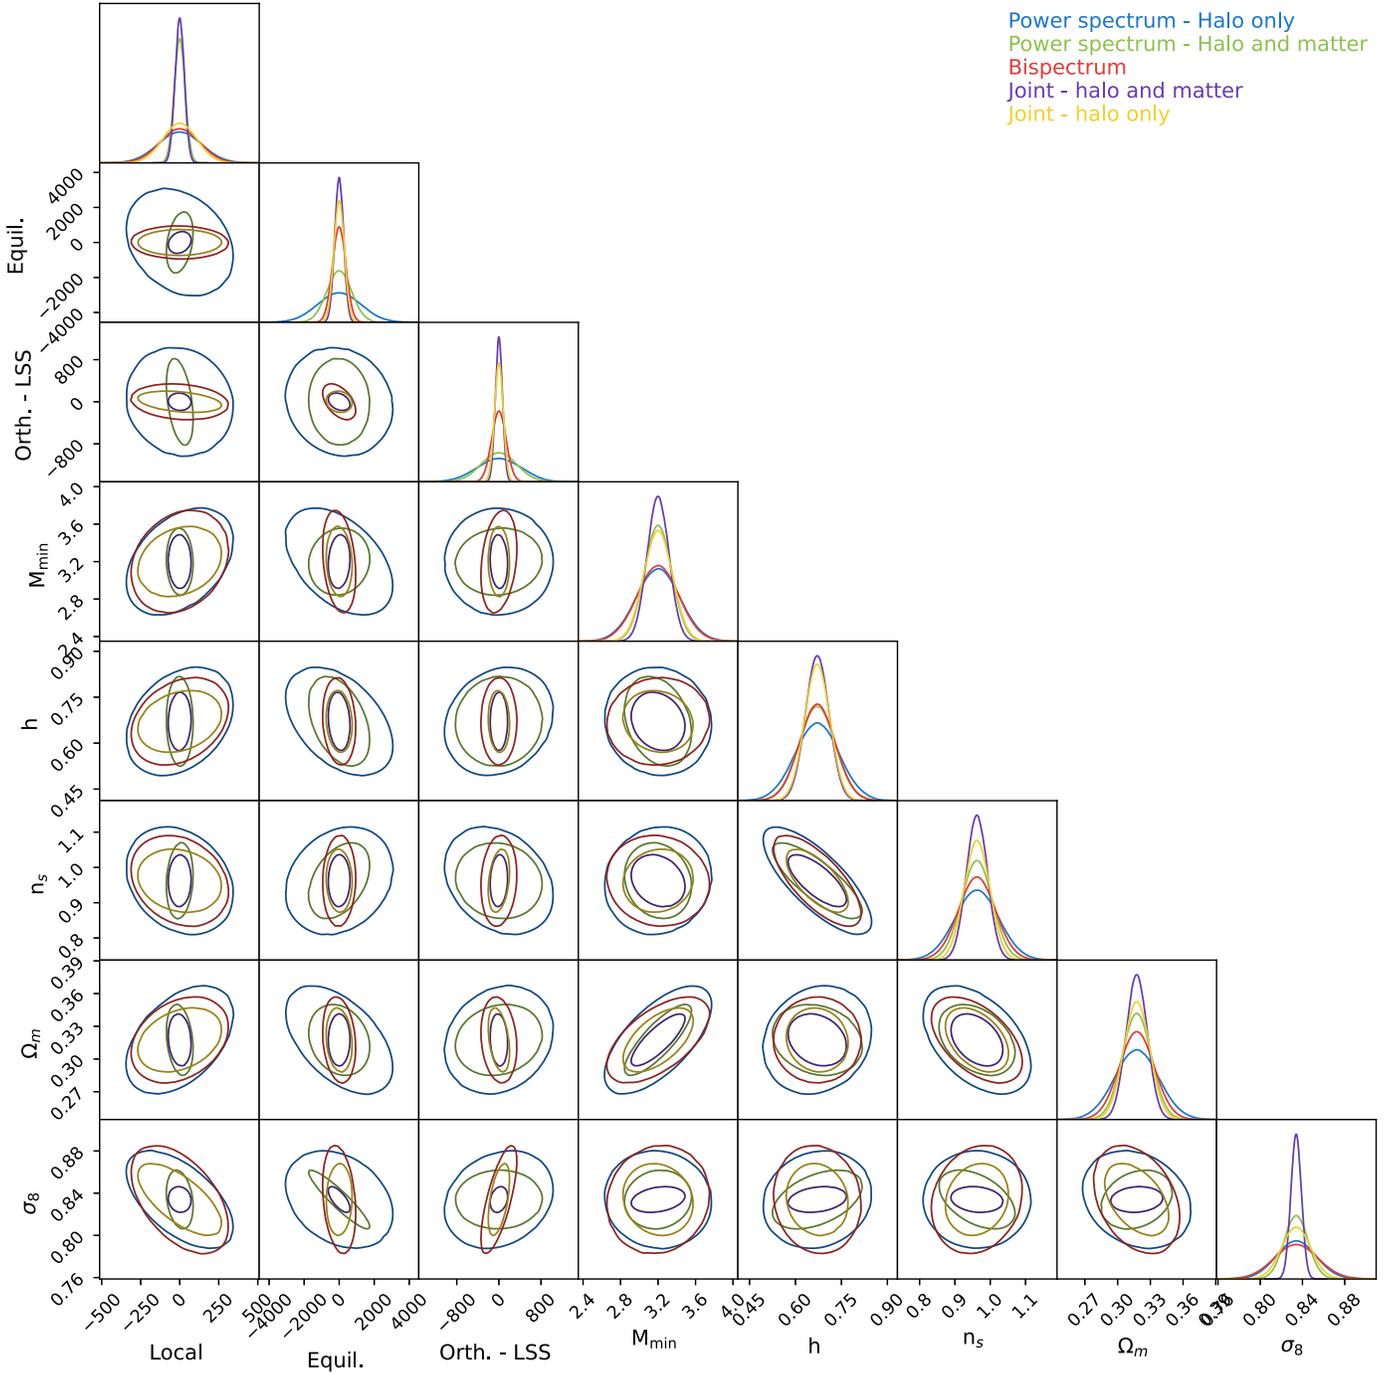


Figure 5. A comparison of the constraining power in the halo power spectrum monopole and quadrupole, the matter power spectrum, the halo–matter power spectrum monopole, and the halo bispectrum. This analysis is performed at $z = 0.0$ with a maximum scale of $k_{\max} = 0.5 \text{ h Mpc}^{-1}$, except for the matter field, for which we only include modes up to $k_{\max} = 0.1 \text{ h Mpc}^{-1}$. These results are for a $1 \text{ Gpc}^3 \text{ h}^{-3}$ volume at $z = 0.0$ for tracers with a number density $1.55 \times 10^{-4} \text{ h}^3 \text{ Mpc}^{-3}$. The contours show the 2σ constraints. These results demonstrate the sensitivities and different degeneracies’ directions for each of the probes.

marginalizing over b_ϕ in our concluding remarks (Section 8). Next, we can see a quantification of the saturation of the constraints: in the primordial space these improve at least as fast as k_{\max}^3 (Kalaja et al. 2021), whereas the late-time improvement is drastically reduced. This saturation was also seen in the matter field (Coulton et al. 2023) and in a similar manner arises as the signal-to-noise ratio of the probes, shown in Figure 8, improves slowly. We discuss the origin of this saturation in the next section.

7.3. Sensitivity to Modeling of the Covariance Matrix

The covariance matrix of our probes can be written as

$$C_{ij} = C_{ij}^{\text{Gaus.}} + C_{ij}^{\text{non-Gaus.}} + C_{ij}^{\text{SSC}}, \quad (21)$$

where the three terms are the Gaussian (or disconnected), the non-Gaussian (or connected), and super sample covariance (SSC) contributions. While the contribution of the Gaussian term (see Appendix B in Coulton et al. 2023, for explicit

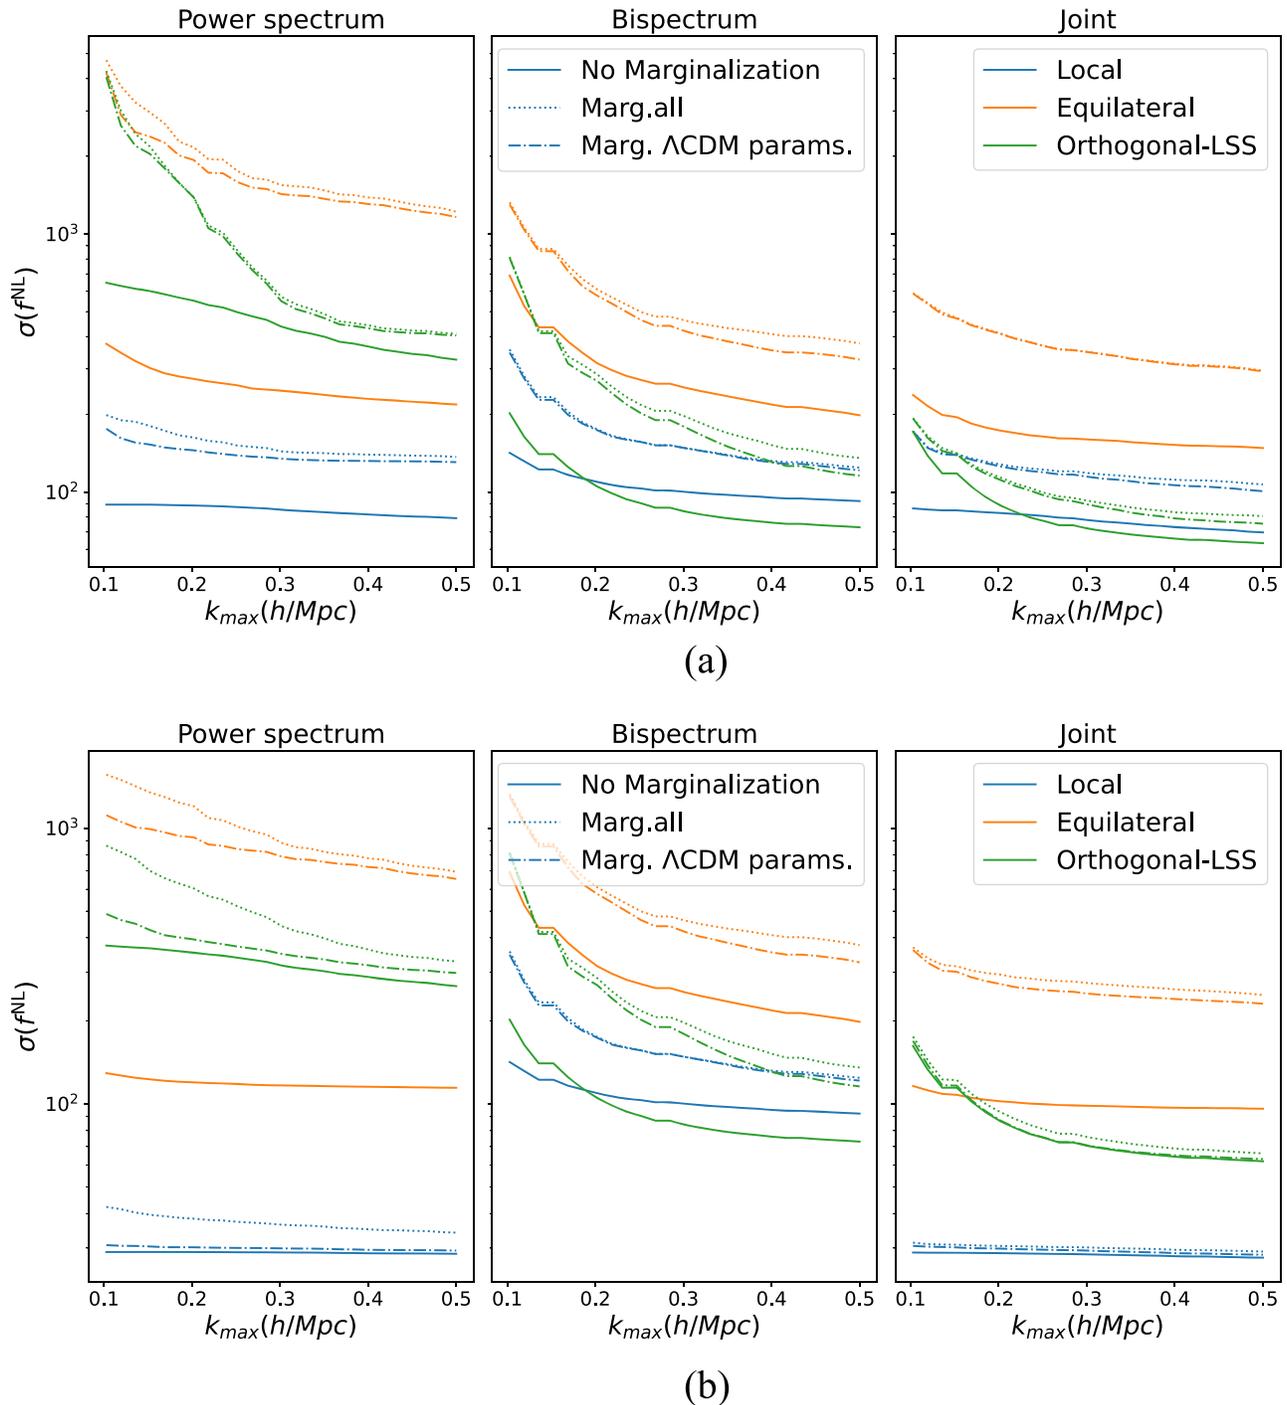


Figure 6. An examination of how the constraining power for power spectrum and bispectrum measurements vary as a function of the maximum scale used in the analysis. (a) Halo field only. (b) The halo field statistics combined with measurements of the large-scale ($k-1$) matter power spectrum and halo-matter cross power spectrum. The solid lines correspond to the 1σ constraining power when no marginalization is performed, the dashed-dotted line includes marginalization of the cosmological and bias parameters, and the dotted line corresponds to marginalization of cosmological and bias parameters and the other PNG shapes. These are computed at $z = 0$, with the same tracer as in Figure 5. These results demonstrate the value of including scales beyond the perturbative regime, $k \gtrsim 0.1 \text{ h Mpc}^{-1}$. The inclusion of the large-scale matter field, Figure 6(b), to the halo-field-only results, Figure 6(a), highlights the power of sample variance cancellation methods.

expressions) is straightforward to calculate analytically, the remaining two terms are complex and typically need to be estimated from simulations (see, e.g., Kayo et al. 2013; Chan & Blot 2017; Galdi et al. 2018; Sugiyama et al. 2020). Given the large size of our data vector, with 2370 elements for the joint bispectrum and halo-matter power spectrum analysis, we

require at least 2370 simulations in order for the estimated covariance matrix to be invertible. From Appendix B we see that we in fact need at least 3000 for our forecast constraints to be converged at the 10% level.

Generating thousands of simulations is computationally very expensive, and hence we investigate the necessity of including

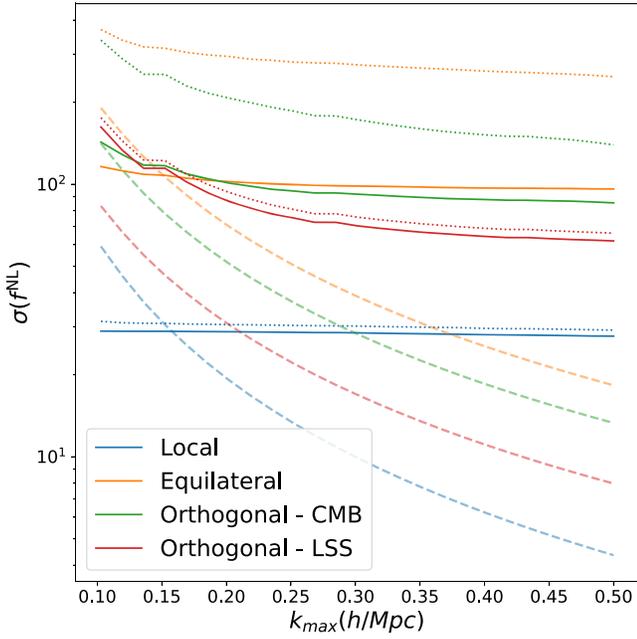


Figure 7. A comparison of the information in the primordial fields (dashed) to that accessible with joint measurements of the matter and halo power spectra and the halo bispectrum monopole (solid lines denote the unmarginalized constraints and the dotted are the marginalized constraints). The marginalization includes the marginalization of the Λ CDM parameters, the bias parameter, and the other PNG templates. Note, however, that we only marginalize over one of the orthogonal templates rather than both as these are two different approximations of the same physical bispectrum.

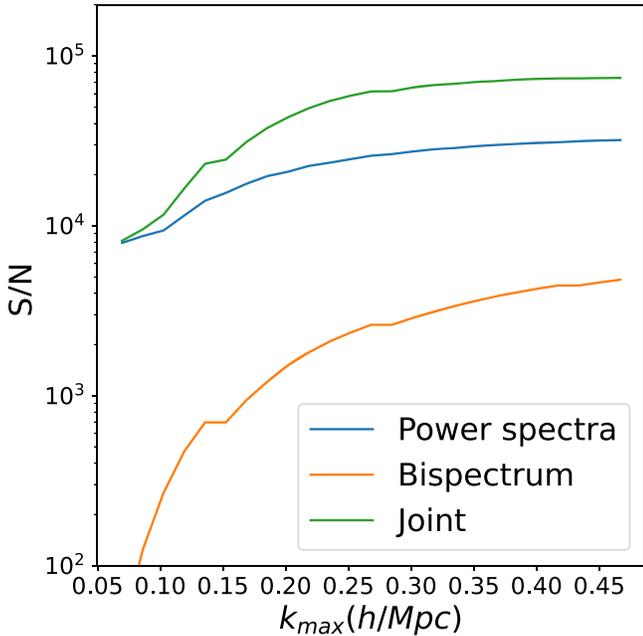


Figure 8. The cumulative signal-to-noise ratio (S/N) of the combined halo and matter power spectra, the bispectrum, and the combination of all the probes. At small scales the modes become increasingly correlated, meaning that the increase in S/N is significantly less than the naive scaling of $k^{\frac{3}{2}}$.

the non-Gaussian and SSC terms. We compute the SSC terms using the separate-universe method described in Li et al. (2014). We use 1000 simulations (500 with a large-scale overdensity mode and 500 with a large-scale underdensity mode; see Appendix B in Coulton et al. 2023 for explicit

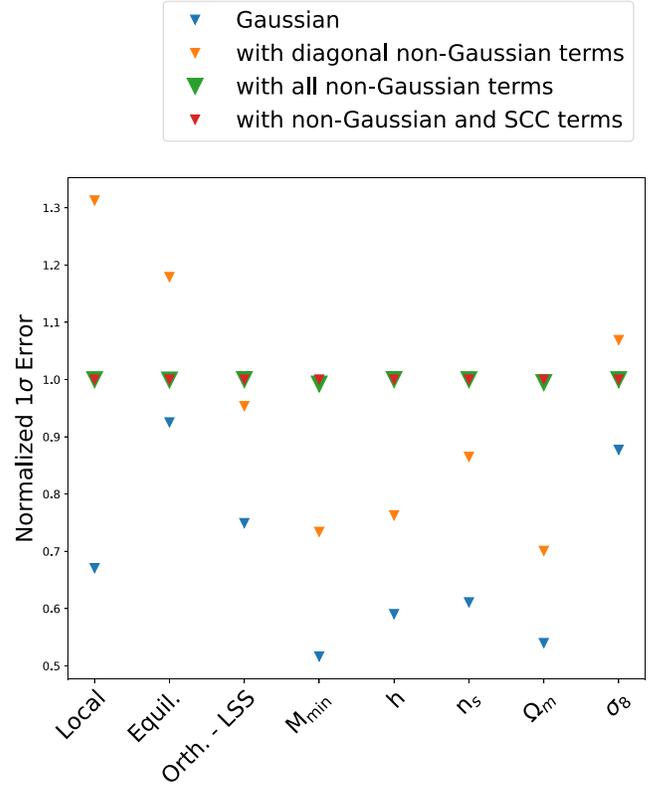


Figure 9. The impact of various components of the covariance matrix to the joint power spectrum and bispectrum parameter constraints. The constraints are normalized by the error obtained by the full covariance matrix (i.e., including all three terms in Equation (21)). While the non-Gaussian contributions to the covariance matrix are very important for parameter constraints, the super sample covariance terms are not; this is seen as the red and green symbols overlap entirely for most parameters (for clarity we have enlarged the green symbols).

formulae). This work builds on the results from many previous works (Chan & Blot 2017; Chan et al. 2018; Barreira 2019; Biagetti et al. 2022), which have found that the non-Gaussian and SSC terms can be the dominant contribution to the covariance matrix, by propagating the impact of these terms into the parameter constraints.

The parameter constraints obtained when only considering subsets of the covariance matrix are shown in Figure 9. As is expected given the large contribution of the non-Gaussian terms shown in Chan & Blot (2017) and Biagetti et al. (2022), we see that only accounting for the Gaussian terms can lead to a $\sim 100\%$ error on the resulting parameter constraints. Interestingly, our results are qualitatively similar to those reported in de Putter (2018) and Flöss et al. (2022), where a perturbative treatment of the leading non-Gaussian terms was performed. Note that this factor is significantly smaller than the equivalent factor for power spectrum and bispectrum measurements of the matter field (Coulton et al. 2023); this is because the shot-noise contribution to the halo field covariance matrix dominates on small scales and reduces the significance of the non-Gaussian terms. It can also be seen that the off-diagonal terms play an equally important role to the diagonal terms and lead to reductions, as well as increases, in the parameter constraints. These reductions occur as the power spectrum measurements act as an ancillary statistic and can remove effective noise in the bispectrum; see Biagetti et al. (2022) and Jung et al. (2022) for a more detailed description.

Finally, we see that, similar to the matter field, the SSC terms generally only lead to small changes in the parameter constraints. This arises as squeezed bispectrum configurations, which are important for our parameter constraints, are minimally impacted by the SSC terms (Chan et al. 2018; Barreira 2019). Given the scalings of the SSC term with volume (see, e.g., Li et al. 2014, for a discussion), the relative contribution of this term is not expected to become more important for larger volumes.

8. Conclusions

In this paper we have used numerical simulations to quantify the information on PNG accessible with measurements of the halo power spectrum, matter power spectrum, halo–matter cross spectrum, and halo bispectrum. For the first time we explore the information content of the halo field including nonlinear scales and the full, non-Gaussian covariance matrix.

First, for the sample considered here, we find that the scale-dependent bias feature contains almost all of the information on local PNG accessible with power spectrum and bispectrum measurements. Thus, when using sample variance cancellation techniques, only measurements of the large-scale power spectrum are needed—especially as there is minimal degeneracy with the other parameters considered here. These results reinforce the analytical results of de Putter (2018), who found that, for tracer number densities below $\bar{n} \sim \text{few} \times 10^{-3} \text{ h}^3 \text{ Mpc}^{-3}$, the bispectrum and trispectrum contains minimal additional information. Interestingly, when compared to the primordial field, the scale-dependent bias feature only contains information equivalent to $k \approx 0.15 \text{ h Mpc}^{-1}$, indicating that alternative late-time statistical probes (such as those discussed in, e.g., Uhlemann et al. 2018, Friedrich et al. 2020, Biagetti et al. 2021, 2022, Andrews et al. 2022, and Giri et al. 2022) may be able to provide additional information.

Second, for equilateral and orthogonal-LSS PNG, which have very weak scale-dependent bias features, we find that bispectrum measurements provide more constraining power than the power spectrum. However, both shapes exhibit strong degeneracies with the cosmological and bias parameters, especially equilateral. For the orthogonal-LSS shape, these degeneracies are largely broken as smaller scale information is included or through combining the bispectrum measurements with the power spectrum measurements. Yet this is not the case for the equilateral shape, highlighting the challenge of measuring this shape with bispectrum measurements.

Similar to the results found for the matter bispectrum (Coulton et al. 2023; Jung et al. 2022), we find that for the local and orthogonal-LSS shapes the information in the bispectrum seems to saturate at $k_{\text{max}} \approx 0.35 \text{ h Mpc}^{-1}$. For the matter bispectrum this occurred due the large non-Gaussian contributions to the covariance matrix. For the halo bispectrum it occurs due to the contribution from shot noise, which dominates on small scales. This suggests that more information may be accessible in the halo field, and therefore the galaxy field, if the shot noise were suppressed (e.g., for a higher density tracer). Additionally, we find that the SSC contribution to the covariance is negligible, generalizing the results of the squeezed limit bispectrum found in Barreira (2019).

The key limitation of our local PNG analysis is that we did not consider the impact of marginalizing over b_ϕ (Equation (10)), which is important as the constraints from scale-dependent bias are sensitive only to the product $b_\phi f_{\text{NL}}$. Our constraints are entirely driven by the scale-dependent bias effect and we would not be able to break the degeneracy between b_ϕ and f_{NL} with the bispectrum

measurements. Thus it would be conservative to interpret our results as constraints for $b_\phi f_{\text{NL}}$. While this degeneracy would not impact our ability to detect the existence of local PNG, it would complicate any interpretation of the specific value. The specific value would rely on the assumptions used in modeling or marginalizing b_ϕ , which depend upon the properties of the specific tracer. A trivial demonstration of this point was seen for orthogonal-CMB PNG: this type of PNG produces a $1/k$ scale-dependent bias but for the tracer considered here the associated bias coefficient is close to zero. This is a trivial case as for our halos the bias parameter can be computed with reasonable accuracy; however, it highlights how small bias coefficients can mask the PNG signals. Moradinezhad Dizgah et al. (2021) explored how the choice of priors in marginalizing over b_ϕ impact constraints and how the addition of bispectrum measurements affect this degeneracy. Informative b_ϕ priors allow strong f_{NL} inferences, but care must be taken to avoid biased results (Moradinezhad Dizgah et al. 2021). This issue is discussed in more depth in Barreira (2022b). In this work, we have also neglected relativistic effects, which can introduce similar features (Bruni et al. 2012; Jeong et al. 2012; Camera et al. 2015; Castorina & Dio 2022), though these can be differentiated via the different redshift evolution.

Similarly, an important and substantial limitation of our equilateral and orthogonal analysis is our bias model. Our bias model effectively only includes the leading bias term and neglects all higher-order biases, which is a gross simplification. We also investigated two variations on our bias models. First we considered replacing the halo mass cutoff, M_{min} , with a linear bias parameter b_1 . We found that using b_1 leads to very similar constraints to using the halo bias parameter, which arises as the two impact our statistics, in the ranges relevant for our bispectrum constraint, in a similar manner. The second test involved fitting constraints without any bias modeling. When we fit all the parameters jointly the contours only marginally widened when including the bias parameters, largely due to the large degeneracies with cosmological parameters discussed above. However, if we fix the cosmological parameters, including the bias parameter, it doubles the equilateral constraint, while leaving the local constraint largely unchanged. The constraints presented here would likely be further degraded when marginalizing over a set of realistic bias uncertainties, especially tidal biases; thus, we view our results as a “best case” estimate. In future work we will analyze this in more detail, by for example using a HOD to marginalize over more realistic biases, as was considered in Hahn & Villaescusa-Navarro (2021). Further, it would be highly interesting to consider different halo mass samples. Unfortunately, the small number of objects present in our simulations means that we cannot further divide our halo catalog without increasing the convergence issues faced in this analysis.

A technical challenge of this work arose due to the large shot noise found in the halo bispectrum, which makes it difficult to accurately measure the bispectrum derivatives. This issue is examined in detail in Appendix A, where we ensure that our results were robust by implementing two methods to mitigate the impact of this noise: first through removing the noise by smoothing with a Gaussian process, and second by using a newly developed method to estimate the Fisher information (W. Coulton & B. Wandelt 2022, in preparation). Without including these mitigation methods our forecast constraints would be biased too small by a factor of >2 ! This difficulty highlights the challenges that shot noise from halos (and

galaxies) poses for analyses exploiting small-scale information and, more generally, the need to carefully examine the convergence of numerical derivatives.

The authors are very grateful to Simone Ferraro, Oliver Philcox, Colin Hill, Daan Meerburg, Shirley Ho, David Spergel, and Yin Li for useful discussions throughout this work. G.J., M.L., and M.B. were supported by the project ‘‘Combining Cosmic Microwave Background and Large Scale Structure data: an Integrated Approach for Addressing Fundamental Questions in Cosmology,’’ funded by the MIUR Progetti di Ricerca di Rilevante Interesse Nazionale (PRIN) Bando 2017—grant No. 2017YJYZAH. D.K. is supported by the South African Radio Astronomy Observatory (SARAO) and the National Research Foundation (grant No. 75415). B.D.W. acknowledges support by the ANR BIG4 project, grant No. ANR-16-CE23-0002 of the French Agence Nationale de la Recherche, and the Labex ILP (reference ANR-10-LABX-63), part of the IDEX SUPER, and received financial state aid managed by the Agence Nationale de la Recherche as part of the program Investissements d’avenir under the reference ANR-11-IDEX-0004-02. The Flatiron Institute is supported by the Simons Foundation. L.V. acknowledges ERC (BePreSySe, grant agree-ment No. 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).

Appendix A Derivative Convergence Analysis

One method of verifying the convergence is to investigate whether the resulting Fisher information, and estimated parameter constraints, $\sigma_i^2 = F_{ii}^{-1}$, are stable to variations in the number of simulations used. In Figure 10 we explore the stability of the bispectrum constraints. We find that the constraints show large variations as the number of simulations is varied and the constraints scale roughly like the square root of the number of derivative simulations. This implies that the resulting constraints are not converged and not reliable.

The reason for this lack of convergence is the large noise in the derivatives, visible in Figure 3. We can understand this behavior by computing the expectation of the standard Fisher information estimator:

$$\langle \hat{F}_{ij} \rangle = \frac{d\mu}{d\theta_i} \Sigma^{-1} \frac{d\mu}{d\theta_j} + \Sigma^{-1} \text{Cov} \left[\frac{d\mu}{d\theta_i}, \frac{d\mu}{d\theta_j} \right], \quad (\text{A1})$$

where we have ignored any noise in the covariance matrix. We see that the Fisher information is biased high due to the covariance of the mean of the derivatives. Hence the parameter constraints, which depend on the inverse, are biased low. The covariance in the derivatives decreases as we utilize more simulations, and hence the bias decreases monotonically as the number of simulations increases.

To generate robust constraints we need to mitigate this noise, and in the following sections we discuss two such methods.

A.1. Smoothed Derivatives

One method of removing the impact of noise would be to fit a function through the derivatives and use this in the Fisher forecast. By eye it is suggestive that, for most of the configurations and

PNG shapes seen in Figure 3, the measured derivatives could be described by a smooth function. Unfortunately, it is challenging to write down a function that captures the full shape over all the scales (perturbative models are not able to capture the range of scales considered here).

As an alternative approach we explored using a Gaussian process (GP) to perform a nonparametric estimation of the underlying function; see Rasmussen & Williams (2006) for a review of GPs. In this approach we treat the measured derivatives as coming from a stochastic process, in this case a zero-mean GP, where we additionally use the measured variance of each point to assign a measurement uncertainty to each point. The intuition with this approach is the following: instead of specifying a full functional form, as in parametric approaches, we ask, given a set of data points, what can be inferred about the value of the function at another value. In this stochastic process framework the relation of the data at one point to any other point (observed or not) is given by the correlation structure of the process; thus, information is contained in the structure of the GP covariance matrix, i.e., the structure of the function is implicitly determined by the data points and the correlations among them. Given our data points (the derivatives from the simulations) and their noise, we fit for the properties of this Gaussian covariance matrix. The noise assigned to each point is vital as it allows the GP to smooth out the noise: it estimates what smooth functions are consistent with the data points, given their measurement noise. Note that we could have fit the GP directly to the simulations at θ , $\theta + \delta\theta$ and $\theta - \delta\theta$ including dependence on the parameters in the GP. Then, we can use the GP to estimate the derivatives. Empirically it was found that fitting the GPs to the derivatives directly worked better. We attribute this difference to sample variance cancellation: our simulations have matched initial seeds so when we compute the derivatives by central differencing the majority of the sample variance cancels. By fitting the GP directly to the simulations at θ , $\theta + \delta\theta$ and $\theta - \delta\theta$ it is difficult to incorporate this sample variance cancellation information, and without it we found the GP was not able to accurately model the derivatives.

We use the implementation provided in the scikit-learn library (Pedregosa et al. 2011). We use an isotropic radial basis function kernel, i.e., a Gaussian kernel, to describe the correlations between data points. The parameters characterizing the GP are estimated by maximizing the log-marginal likelihood from the derivatives. Before fitting the GP to the derivatives we divide the data by a smooth function to reduce the dynamic range of the problem. This provided significantly better performance. For the Λ CDM parameters we divide the derivatives by the mean of the covariance matrix simulations. For the PNG derivatives, we use a linear combination of the tree-level PNG bispectra and the mean of the covariance matrix (this combination helps capture the general behavior on both the large and small scales). The measurement error assigned to each point when fitting the GP is the error on the mean of our simulated derivatives. Once the derivatives have been smoothed by our GP we multiply back by the same base functions.

In Figure 3 we compare the resulting smoothed bispectra to the measurements. There are several notable features. First, where the bispectra are well measured the GP matches the measurements very well. In the highly noisy region the GP is a significantly smoother function, producing a curve that is

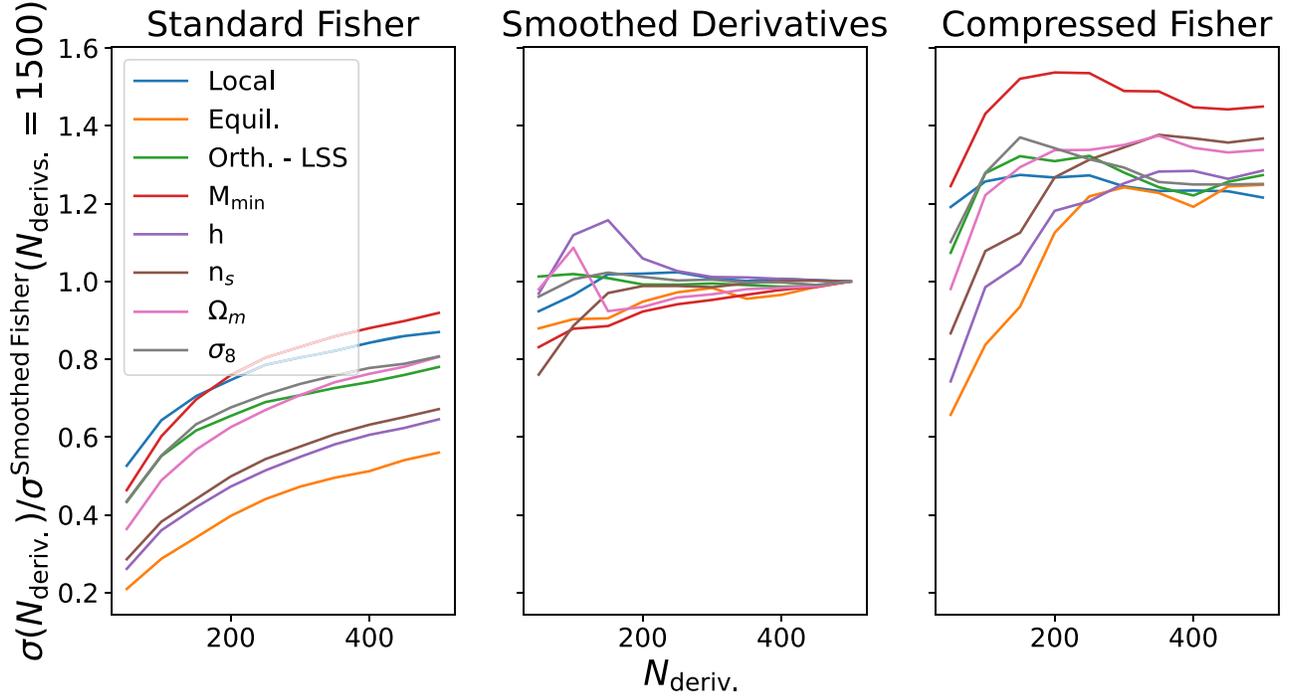


Figure 10. A examination of the stability of the Fisher forecast to changes in the number of simulations. In the left plot we plot the results of the standard Fisher forecast. In the center plot we show the results when we use a Gaussian process (GP) to mitigate noise in the derivatives. In the right plot we show an alternative Fisher forecast method using data compression. This last method, for $N_{\text{deriv.}} \gtrsim 250$, provides an approximate upper bound on the Fisher information. To easily compare across the different parameters, the Fisher estimates are normalized by the GP estimate for each parameter.

consistent with the simulation points, given their errors. In a few places the GP shows statistically significant deviations from the data points; this demonstrates the limitations of this method. These points typically occur when there is a steep gradient in one direction of the bispectrum, and thus these points represent an oversmoothing. Despite this caveat, the smoothed derivatives generally provide a good, smoothed approximation of the simulations.

Given this smoothing method, we can investigate how our constraining power changes as the number of simulations used to compute the smoothed derivatives is varied. This test allows us to assess the effectiveness of our smoothing procedure: if the derivatives are not sufficiently smoothed we expect to see constraints that increase as more simulations are added; on the other hand, if the smoothing process is too aggressive, and is removing important physical features, we may see constraints that decrease as the number of simulations are varied. As is seen in Figure 10, our constraints are now stable to variations in the number of simulations, thus providing a useful validation of our smoothing method.

A.2. Compressed Fisher Forecast

Given the potential caveats of the smoothed approach, it is worthwhile to explore alternative methods of validating our Fisher forecast. Recent work by W. Coulton & B. Wandelt (2022, in preparation) has demonstrated an alternative method of performing Fisher forecasts that, under certain conditions, leads to conservative forecast constraints. Thus this method allows another test of our smoothed derivatives and can be combined with the standard Fisher method to provide bounds on the constraints. Here we briefly give an overview of the method, hereafter referred to as the compressed Fisher forecast;

we refer the reader to W. Coulton & B. Wandelt (2022, in preparation) for more details.

In this approach instead of computing the Fisher information of the power spectrum and bispectrum, we compute the Fisher information of a set of compressed summary statistics. In principle, the information in the power spectrum and bispectrum can be losslessly compressed to a set of statistics with the same dimension as the number of parameters (this includes both cosmological parameters and nuisance parameters). One method of achieving this is to compress the data via (Alsing & Wandelt 2018)

$$\mathbf{t} = \nabla \mathcal{L}(\mathbf{d}, \theta), \quad (\text{A2})$$

where \mathbf{t} are the summary statistics, $\mathcal{L}(\mathbf{d}, \theta)$ is the likelihood for data, \mathbf{d} , given parameters, θ , and the derivative is with respect to the parameters of interest. In this compressed space, the Fisher information is

$$F_{ij}^t = \frac{\partial \mu^t}{\partial \theta_i} \Sigma^{t-1} \frac{\partial \mu^t}{\partial \theta_j}, \quad (\text{A3})$$

where μ^t and Σ^t are the mean and covariance of the compressed statistics. Performing the Fisher analysis in the compressed space has two advantages: first, the dimension of the data vectors are drastically reduced, so a smaller dimension covariance matrix is needed; second, the noise on the derivatives of the compressed statistic, which is required for the Fisher forecast, is expected to be lower as it is a weighted sum of all the data points and thereby averages down the noise in the uncompressed derivatives.

For the case of a Gaussian likelihood with parameter-independent covariance, as is used here and motivated by Scoccimarro (2000) and Carron (2013), the lossless

compression is (Tegmark et al. 1997; Heavens et al. 2000)

$$t = \frac{\partial \mu}{\partial \theta} C^{-1} (d - \mu), \quad (\text{A4})$$

where μ and C are the data mean and covariance. Compression has already been considered for bispectrum analysis (e.g., Gualdi et al. 2019; Philcox et al. 2021; Byun & Krause 2022) to aid working with very-high-dimensional data vectors, but here we are using it to improve the accuracy of the Fisher forecast (and we use a simulation-based rather than perturbative model). An issue that is immediately clear is that performing the optimal compression requires the same ingredients as performing the standard Fisher forecast. This issue can be resolved by splitting the simulations into two parts: the first part is used to estimate the compression and the second part is then compressed and used to estimate the quantities needed for the compressed Fisher forecast, i.e., Equation (A3). This splitting is necessary as reusing simulations will lead to biased Fisher estimates. In this work we use 75% of the data to estimate the compression and 25% to compute the compressed Fisher information. The use of estimated quantities in the compression results in a suboptimal compression, and so the resulting Fisher forecast errors will be larger than the truth (see, e.g., Lehmann & Casella 1998). Thus we have traded forecast errors that were biased too small for a set of conservative errors that are biased too large. In the limit of a sufficient number of simulations the standard Fisher and compressed Fisher estimates tend to the same values. Note that when the noise in the compressed

derivatives is large, which occurs for small numbers of simulations, this estimator will be biased low in the same manner as the standard estimator.

In Figure 10 we investigate the stability of our compressed analysis to the number of simulations used to compute the derivatives. Unlike the standard estimate, the compressed forecast reaches a regime where it is stable to changes in the number of simulations used: for most of the parameters the forecast errors are largely unchanged or they decrease as more than 250 simulations are used. The decrease occurs as, once convergence is reached, using more simulations only improves the effectiveness of the compression. Note that for very small numbers of simulations, $\lesssim 250$, the forecast errors based on the compressed Fisher method is biased low due to noisy estimates of the derivatives, in an identical effect to the standard forecast; see W. Coulton & B. Wandelt (2022, in preparation) for a detailed discussion of this estimator, as well as an improved estimator that accelerates convergence to the exact Fisher forecast.

Given that our compressed Fisher forecast is converged, we can compare these forecast constraints to those of the other methods. In Figure 10 we see that, as expected, the forecast constraints are larger than the GP constraints and the standard (and unconverged) forecast. This hierarchy provides a conservative estimate of the information available and validation that our smoothing method does not wash out the important features.

Appendix B Convergence of the Covariance Matrix

In Figure 11 we demonstrate that our forecast constraints only exhibit a $\lesssim 2\%$ variation as the number of simulations is

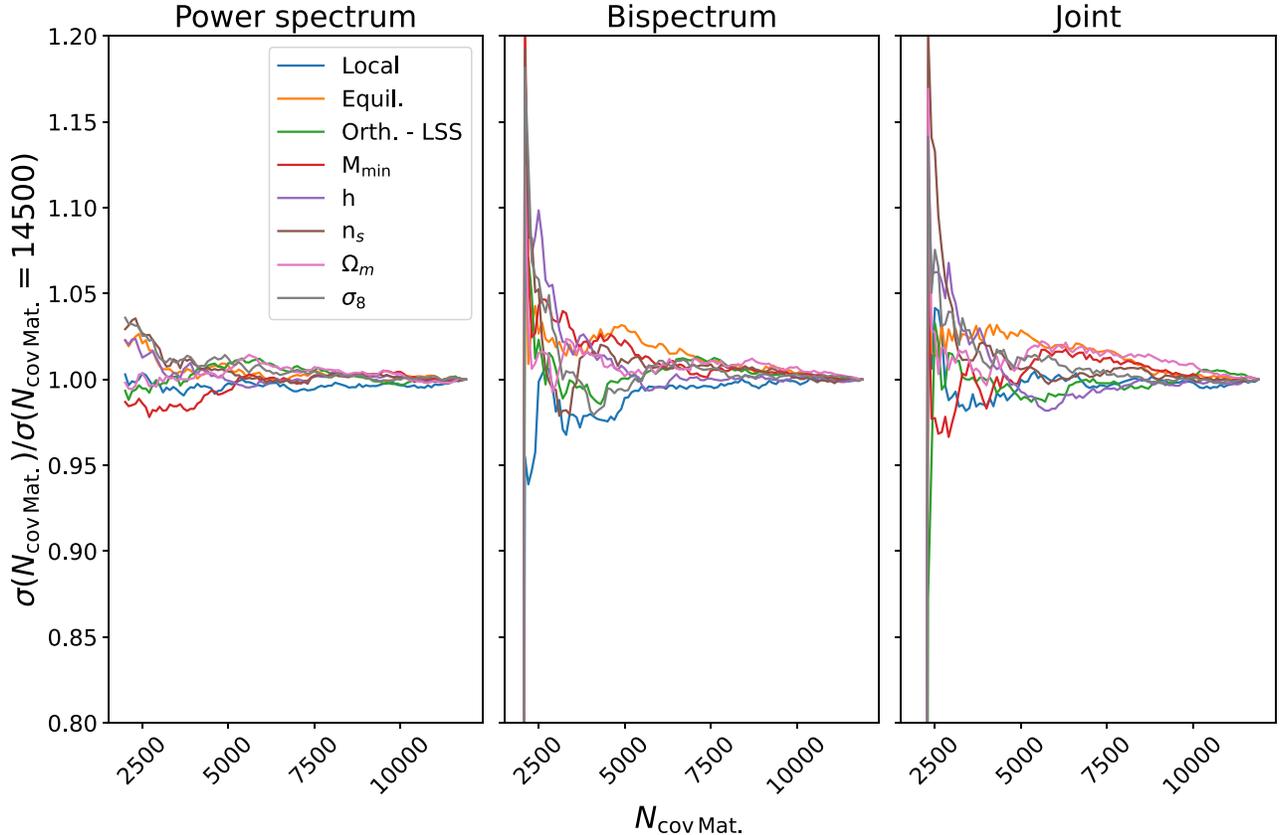


Figure 11. The stability of our joint constraints to variations in the number of simulations used to compute the covariance matrix. The small variations seen demonstrate that the covariance matrix used in our full analysis, computed with 14,500 simulations, is converged.

changed from 7000 to 14,500, and thus demonstrates that our analysis is likely converged with respect to the number of covariance matrix simulations.

ORCID iDs

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

References

- Achúcarro, A., Biagetti, M., Braglia, M., et al. 2022, arXiv:2203.08128
 Alishahiha, M., Silverstein, E., & Tong, D. 2004, *PhRvD*, **70**, 123505
 Alsing, J., & Wandelt, B. 2018, *MNRAS*, **476**, L60
 Andrews, A., Jasche, J., Lavaux, G., & Schmidt, F. 2022, arXiv:2203.08838
 Babich, D. 2005, *PhRvD*, **72**, 043003
 Baldauf, T., Seljak, U., & Senatore, L. 2011, *JCAP*, **2011**, 006
 Barreira, A. 2019, *JCAP*, **2019**, 008
 Barreira, A. 2020, *JCAP*, **2020**, 031
 Barreira, A. 2022a, *JCAP*, **2022**, 033
 Barreira, A. 2022b, *JCAP*, **2022**, 13
 Baumann, D., & Green, D. 2022, *JCAP*, **2022**, 61
 Biagetti, M., Calles, J., Castiblanco, L., Cole, A., & Noreña, J. 2022, *JCAP*, **2022**, 02
 Biagetti, M., Castiblanco, L., Noreña, J., & Sefusatti, E. 2022, *JCAP*, **2022**, 09
 Biagetti, M., Cole, A., & Shiu, G. 2021, *JCAP*, **04**, 061
 Bruni, M., Crittenden, R., Koyama, K., et al. 2012, *PhRvD*, **85**, 041301
 Byun, J., & Krause, E. 2022, *MNRAS*, in press
 Cabass, G., Ivanov, M. M., Philcox, O. H. E., Simonović, M., & Zaldarriaga, M. 2022a, *PhRvD*, **106**, 043506
 Cabass, G., Ivanov, M. M., Philcox, O. H. E., Simonović, M., & Zaldarriaga, M. 2022b, *PhRvL*, **129**, 021301
 Camera, S., Maartens, R., & Santos, M. G. 2015, *MNRAS*, **451**, L80
 Carron, J. 2013, *A&A*, **551**, A88
 Castorina, E., & Dio, E. Di. 2022, *JCAP*, **2022**, 061
 Castorina, E., Feng, Y., Seljak, U., & Villaescusa-Navarro, F. 2018, *PhRvL*, **121**, 101301
 Chan, K. C., & Blot, L. 2017, *PhRvD*, **96**, 023528
 Chan, K. C., Hamaus, N., & Biagetti, M. 2019, *PhRvD*, **99**, 121304
 Chan, K. C., Moradinezhad Dizgah, A., & Noreña, J. 2018, *PhRvD*, **97**, 043532
 Chen, X. 2010, *AdAst*, **2010**, 638979
 Chen, X., Huang, M.-X., Kachru, S., & Shiu, G. 2007, *JCAP*, **2007**, 002
 Cheung, C., Fitzpatrick, A. L., Kaplan, J., & Senatore, L. 2008a, *JCAP*, **2008**, 021
 Cheung, C., Fitzpatrick, A. L., Kaplan, J., Senatore, L., & Creminelli, P. 2008b, *JHEP*, **2008**, 014
 Chiang, C.-T., Hu, W., Li, Y., & LoVerde, M. 2018, *PhRvD*, **97**, 123526
 Coulton, W. R., Villaescusa-Navarro, F., Jamieson, D., et al. 2023, *ApJ*, **943**, 64
 Creminelli, P., & Zaldarriaga, M. 2004, *JCAP*, **2004**, 006
 Dalal, N., Doré, O., Huterer, D., & Shirokov, A. 2008, *PhRvD*, **77**, 123514
 D'Amico, G., Donath, Y., Lewandowski, M., Senatore, L., & Zhang, P. 2022a, arXiv:2206.08327
 D'Amico, G., Lewandowski, M., Senatore, L., & Zhang, P. 2022b, 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. 2011, *PhRvD*, **84**, 063512
 Di Valentino, E., Mena, O., Pan, S., et al. 2021, *CQGra*, **38**, 153001
 Doré, O., Bock, J., Ashby, M., et al. 2014, arXiv:1412.4872
 Dvorkin, C., Gerbino, M., Alonso, D., et al. 2019, *BAAS*, **51**, 64
 Flöss, T., Biagetti, M., & Meerburg, P. D. 2022, arXiv:2206.10458
 Foreman, S., Coulton, W., Villaescusa-Navarro, F., & Barreira, A. 2020, *MNRAS*, **498**, 2887
 Friedrich, O., Uhlemann, C., Villaescusa-Navarro, F., et al. 2020, *MNRAS*, **498**, 464
 Gil-Marín, H., Noreña, J., Verde, L., et al. 2015a, *MNRAS*, **451**, 539
 Gil-Marín, H., Percival, W. J., Verde, L., et al. 2017, *MNRAS*, **465**, 1757
 Gil-Marín, H., Verde, L., Noreña, J., et al. 2015b, *MNRAS*, **452**, 1914
 Giri, U., Münchmeyer, M., & Smith, K. M. 2022, arXiv:2205.12964
 Green, D., Amin, M. A., Meyers, J., et al. 2019, *BAAS*, **51**, 159
 Gualdi, D., Gil-Marín, H., Schuhmann, R. L., et al. 2019, *MNRAS*, **484**, 3713
 Gualdi, D., Manera, M., Joachimi, B., & Lahav, O. 2018, *MNRAS*, **476**, 4045
 Hahn, C., & Villaescusa-Navarro, F. 2021, *JCAP*, **2021**, 029
 Hahn, C., Villaescusa-Navarro, F., Castorina, E., & Scoccimarro, R. 2020, *JCAP*, **2020**, 040
 Heavens, A. F., Jimenez, R., & Lahav, O. 2000, *MNRAS*, **317**, 965
 Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, *ApJS*, **208**, 19
 Hockney, R. W., & Eastwood, J. W. 1981, *Computer Simulation Using Particles* (Boca Raton, FL: CRC Press)
 Jeong, D., Schmidt, F., & Hirata, C. M. 2012, *PhRvD*, **85**, 023504
 Jing, Y. P. 2005, *ApJ*, **620**, 559
 Jung, G., Karagiannis, D., Liguori, M., et al. 2022, *ApJ*, **940**, 71
 Kalaja, A., Meerburg, P. D., Pimentel, G. L., & Coulton, W. R. 2021, *JCAP*, **2021**, 050
 Kayo, I., Takada, M., & Jain, B. 2013, *MNRAS*, **429**, 344
 Komatsu, E., & Spergel, D. N. 2001, *PhRvD*, **63**, 063002
 Komatsu, E., Spergel, D. N., & Wandelt, B. D. 2005, *ApJ*, **634**, 14
 Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, arXiv:1110.3193
 Lehmann, E. L., & Casella, G. 1998, *Theory of Point Estimation* (2nd ed.; New York: Springer)
 Leistedt, B., Peiris, H. V., & Roth, N. 2014, *PhRvL*, **113**, 221301
 Li, Y., Hu, W., & Takada, M. 2014, *PhRvD*, **89**, 083519
 LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009, arXiv:0912.0201
 Lucchin, F., & Matarrese, S. 1988, *ApJ*, **330**, 535
 Maldacena, J. 2003, *JHEP*, **013**, 2003
 Matarrese, S., Pilo, L., & Rollo, R. 2021, *JCAP*, **2021**, 062
 Matarrese, S., & Verde, L. 2008, *ApJL*, **677**, L77
 Matarrese, S., Verde, L., & Jimenez, R. 2000, *ApJ*, **541**, 10
 Meerburg, P. D., Green, D., Flauger, R., et al. 2019, *BAAS*, **51**, 107
 Moradinezhad Dizgah, A., Biagetti, M., Sefusatti, E., Desjacques, V., & Noreña, J. 2021, *JCAP*, **2021**, 015
 Mueller, E.-M., Rezaie, M., Percival, W. J., et al. 2021, arXiv:2106.13725
 Münchmeyer, M., Madhavacheril, M. S., Ferraro, S., Johnson, M. C., & Smith, K. M. 2019, *PhRvD*, **100**, 083508
 Pajer, E., Schmidt, F., & Zaldarriaga, M. 2013, *PhRvD*, **88**, 083502
 Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *J Mach. Learn. Res.*, **12**, 2825, <https://www.jmlr.org/papers/volume12/pedregosa11a/pedregosa11a.pdf>
 Philcox, O. H. E., Ivanov, M. M., Cabass, G., et al. 2022, *PhRvD*, **106**, 043530
 Philcox, O. H. E., Ivanov, M. M., Zaldarriaga, M., Simonović, M., & Schmittfull, M. 2021, *PhRvD*, **103**, 043508
 Planck Collaboration VI 2020, *A&A*, **641**, A6
 Planck Collaboration IX 2020, *A&A*, **641**, A9
 Press, W. H., & Schechter, P. 1974, *ApJ*, **187**, 425
 Pullen, A. R., & Hirata, C. M. 2013, *PASP*, **125**, 705
 Rasmussen, C. E., & Williams, C. K. I. 2006, *Gaussian Processes for Machine Learning, Adaptive Computation and Machine Learning* (Cambridge, MA: MIT Press)
 Reid, B. A., Verde, L., Dolag, K., Matarrese, S., & Moscardini, L. 2010, *JCAP*, **2010**, 013
 Rezaie, M., Ross, A. J., Seo, H.-J., et al. 2021, *MNRAS*, **506**, 3439
 Ross, A. J., Ho, S., Cuesta, A. J., et al. 2011, *MNRAS*, **417**, 1350
 Ross, A. J., Percival, W. J., Carnero, A., et al. 2013, *MNRAS*, **428**, 1116
 Sabti, N., Muñoz, J. B., & Blas, D. 2021, *JCAP*, **2021**, 010
 Schmidt, F., & Kamionkowski, M. 2010, *PhRvD*, **82**, 103002
 Schmittfull, M., & Seljak, U. 2018, *PhRvD*, **97**, 123540
 Scoccimarro, R. 2000, *ApJ*, **544**, 597
 Scoccimarro, R., Hui, L., Manera, M., & Chan, K. C. 2012, *PhRvD*, **85**, 083002
 Scoccimarro, R., Sefusatti, E., & Zaldarriaga, M. 2004, *PhRvD*, **69**, 103513
 Sefusatti, E. 2009, *PhRvD*, **80**, 123002
 Sefusatti, E., Crocce, M., & Desjacques, V. 2012, *MNRAS*, **425**, 2903
 Sefusatti, E., Crocce, M., Scoccimarro, R., & Couchman, H. M. P. 2016, *MNRAS*, **460**, 3624
 Seljak, U. 2009, *PhRvL*, **102**, 021302
 Senatore, L., Smith, K. M., & Zaldarriaga, M. 2010, *JCAP*, **2010**, 028

- Shandera, S., Mantz, A., Rapetti, D., & Allen, S. W. 2013, *JCAP*, 2013, 004
- Sheth, R. K., Mo, H. J., & Tormen, G. 2001, *MNRAS*, 323, 1
- Slepian, Z., Eisenstein, D. J., Beutler, F., et al. 2017a, *MNRAS*, 468, 1070
- Slepian, Z., Eisenstein, D. J., Brownstein, J. R., et al. 2017b, *MNRAS*, 469, 1738
- Slosar, A. 2009, *JCAP*, 2009, 004
- Slosar, A., Hirata, C., Seljak, U., Ho, S., & Padmanabhan, N. 2008, *JCAP*, 2008, 031
- Slosar, A., Mandelbaum, R., & Eisenstein, D. 2019, *BAAS*, 51, 97
- Spergel, D., Gehrels, N., Baltay, C., et al. 2015, arXiv:1503.03757
- Springel, V. 2005, *MNRAS*, 364, 1105
- Sugiyama, N. S., Saito, S., Beutler, F., & Seo, H.-J. 2020, *MNRAS*, 497, 1684
- Taylor, A. N., & Watts, P. I. R. 2001, *MNRAS*, 328, 1027
- Tegmark, M., Taylor, A. N., & Heavens, A. F. 1997, *ApJ*, 480, 22
- Tellarini, M., Ross, A. J., Tasinato, G., & Wands, D. 2016, *JCAP*, 2016, 014
- Uhlemann, C., Pajer, E., Pichon, C., et al. 2018, *MNRAS*, 474, 2853
- Villaescusa-Navarro, F., Hahn, C., Massara, E., et al. 2020, *ApJS*, 250, 2
- Wagner, C., & Verde, L. 2012, *JCAP*, 2012, 002
- Yadav, A. P. S., Komatsu, E., & Wandelt, B. D. 2007, *ApJ*, 664, 680