



Publication Year	2021
Acceptance in OA	2022-03-17T16:30:27Z
Title	Towards cosmological constraints from the compressed modal bispectrum: A robust comparison of real-space bispectrum estimators
Authors	Byun, Joyce, Oddo, Andrea, PORCIANI, CRISTIANO, SEFUSATTI, Emiliano
Publisher's version (DOI)	10.1088/1475-7516/2021/03/105
Handle	http://hdl.handle.net/20.500.12386/31689
Journal	JOURNAL OF COSMOLOGY AND ASTROPARTICLE PHYSICS
Volume	2021

Towards cosmological constraints from the compressed modal bispectrum: a robust comparison of real-space bispectrum estimators

Joyce Byun,^{a,1} Andrea Oddo,^{b,c,d} Cristiano Porciani,^e and Emiliano Sefusatti^{f,c,d}

^aDépartement de Physique Théorique and Center for Astroparticle Physics (CAP), University of Geneva, 24 quai Ernest Ansermet, CH-1211 Geneva, Switzerland

^bSISSA - International School for Advanced Studies, Via Bonomea 265, 34136 Trieste, Italy

^cInstitute for Fundamental Physics of the Universe, Via Beirut 2, 34151 Trieste, Italy

^dIstituto Nazionale di Fisica Nucleare, Sezione di Trieste, via Valerio 2, 34127 Trieste, Italy

^eArgelander Institut für Astronomie der Universität Bonn, Auf dem Hügel 71, 53121 Bonn, Germany

^fIstituto Nazionale di Astrofisica, Osservatorio Astronomico di Trieste, via Tiepolo 11, 34143 Trieste, Italy

E-mail: joyce.byun@unige.ch, andrea.oddo@sissa.it, porciani@astro.uni-bonn.de, emiliano.sefusatti@inaf.it

Abstract. Higher-order clustering statistics, like the galaxy bispectrum, can add complementary cosmological information to what is accessible with two-point statistics, like the power spectrum. While the standard way of measuring the bispectrum involves estimating a bispectrum value in a large number of Fourier triangle bins, the compressed modal bispectrum approximates the bispectrum as a linear combination of basis functions and estimates the expansion coefficients on the chosen basis. In this work, we compare the two estimators by using parallel pipelines to analyze the real-space halo bispectrum measured in a suite of N -body simulations corresponding to a total volume of $\sim 1,000 h^{-3} \text{Gpc}^3$, with covariance matrices estimated from 10,000 mock halo catalogs. We find that the modal bispectrum yields constraints that are consistent and competitive with the standard bispectrum analysis: for the halo bias and shot noise parameters within the tree-level halo bispectrum model up to $k_{\text{max}} \approx 0.06 (0.10) h \text{Mpc}^{-1}$, only 6 (10) modal expansion coefficients are necessary to obtain constraints equivalent to the standard bispectrum estimator using ~ 20 to 1,600 triangle bins, depending on the bin width. For this work, we have implemented a modal estimator pipeline using Markov Chain Monte Carlo simulations for the first time, and we discuss in detail how the parameter posteriors and modal expansion are robust to, or sensitive to, several user settings within the modal bispectrum pipeline. The combination of the highly efficient compression that is achieved and the large number of mock catalogs available allows us to quantify how our modal bispectrum constraints depend on the number of mocks that are used to estimate covariance matrices and the functional form of the likelihood.

¹Corresponding author.

Contents

1	Introduction	1
2	Modal decomposition method	3
2.1	Modal estimator	5
2.2	Inner product methods	7
2.3	Bispectrum model and custom modes	9
2.4	Orthonormal basis	11
2.5	Summary	14
3	Data and analysis	15
3.1	Simulations and mock halo catalogs	15
3.2	Likelihood and MCMC	15
4	Results	17
4.1	Mean and covariance of modal coefficients	17
4.2	Standard bispectrum vs reconstructed bispectrum	20
4.3	Benchmark comparison: modal bispectrum vs standard bispectrum	22
4.4	Robustness checks in the modal implementation	25
4.5	Modal expansion correlators	33
4.6	Covariances	38
4.7	Gaussian vs Sellentin-Heavens likelihood	38
5	Conclusions	42
A	1-dimensional basis functions	49
B	Calculating the inner product using 1-dimensional FFT	49

1 Introduction

Gravitational clustering and nonlinear bias induce a non-Gaussianity in the large-scale structure (LSS) of the Universe that can be quantified by higher-order correlation functions, like the 3-point correlation function (3PCF) and its Fourier counterpart, the bispectrum. The measurement of the galaxy bispectrum in redshift surveys can contribute additional constraining power towards a wide range of science goals, improving on what is achievable using only 2-point correlations, like the power spectrum.

To date, the most precise measurements of the galaxy bispectrum and 3PCF are from the SDSS Baryon Oscillation Spectroscopic Survey [1–7]. In the near future, spectroscopic galaxy surveys like DESI,¹ Euclid,² SPHEREx,³ and the Roman Space Telescope⁴ [8–11] will map galaxy distributions over larger areas of the sky, to higher redshifts, and with more precision

¹<https://www.desi.lbl.gov>

²<https://www.euclid-ec.org>

³<https://spherex.caltech.edu>

⁴<https://wfirst.ipac.caltech.edu>

than before, opening up new opportunities for higher-order galaxy clustering statistics to be used as stronger probes of Λ CDM (the standard cosmological model dominated by a cosmological constant called Λ and cold dark matter), dark energy, modified gravity theories, primordial non-Gaussianity, and massive neutrinos (e.g. [12–25]).

Bispectrum data sets are naturally much larger than for power spectra, because they typically capture the correlation amplitudes for a large number of triangle bins, $B(k_1, k_2, k_3)$, rather than being a function of only one wavenumber, like $P(k)$. This presents practical challenges that can potentially limit the full exploitation of the data that will be available from upcoming LSS surveys. One particularly acute challenge for bispectrum analyses is the large number of mock catalogs that are typically necessary to accurately estimate the large data covariance matrices for galaxy clustering analyses. One remedy to this problem has been to explore the accuracy of fast, approximate simulation codes that reduce the computational resources needed for generating mocks [26, 27]. Alternatively, it may be possible to obtain equivalent covariance matrices using fewer or smaller volume mocks (e.g. [28–32]) or even no mocks, if an accurate theoretical model of the covariance matrix is available (e.g. [33–36]). A third strategy that has been pursued in tandem is to develop methods that compress the information contained in the bispectrum into smaller, more manageable data sets.

Much work to date falls into this last category of bispectrum compression methods. Using the standard bispectrum estimator, choosing to use wider wavenumber bins reduces the total number of triangle bins, but at the same time erases some of the triangle-dependence that encodes cosmological information. Other compression methods and compressed bispectrum observables include: Karhunen-Loève compression of the standard bispectrum estimator [37, 38] (which is similar to the MOPED algorithm [39, 40]), subspace projection of the standard bispectrum estimator [41], binning triangles based on their geometrical properties [42, 43], skew-spectra [44–46], position-dependent power spectra (also called integrated bispectra) [47–49], line correlation functions [50–56], and the modal bispectrum. The focus of this work is to implement and explore the last of these, and to compare it with a standard bispectrum analysis.

The modal bispectrum describes the bispectrum as a linear combination of smooth 3-dimensional basis functions, such that the observable data are the expansion coefficients over a chosen basis. If the bispectrum is relatively smooth, and the chosen basis is suitable for describing changes in the bispectrum induced by the model parameters we wish to constrain, we expect that the modal expansion coefficients will provide an efficient compression of the cosmological information that is typically distributed over a large number of triangle bins. The modal expansion method was originally developed in the context of primordial non-Gaussianity in the cosmic microwave background [57–59] before it was adapted for the LSS bispectrum [60–62]. It has been used to test and develop theoretical models of the matter bispectrum [62–64] and compare the matter bispectra measured from different dark matter simulation codes [65] and mock-making prescriptions [66]. These previous works have been in real space, and an extension of the modal expansion method to redshift space was outlined in [67].

The first direct comparison between the modal bispectrum and the standard bispectrum estimators was in [13], where Fisher forecasted constraints on Λ CDM cosmological parameters and galaxy bias using the real-space matter modal bispectrum and standard bispectrum estimators were equivalent. For this work, we have implemented a new modal bispectrum analysis pipeline that uses the real-space halo bispectrum to constrain galaxy bias and shot noise parameters. This builds on previous work by implementing the modal

bispectrum method in a Markov chain Monte Carlo (MCMC) analysis pipeline for the first time. Where possible, we have adhered closely to the standard bispectrum analysis in [68], so the two estimators can be rigorously compared. In the process of implementing the new modal bispectrum pipeline, we have explored several technical details of the modal method’s implementation that have not previously been presented in the literature, and we discuss how the modal estimator pipeline is sensitive (or not) to these details.

The main message of this work is that the modal bispectrum estimator provides an extremely efficient compression of the information contained within the halo bispectrum, resulting in parameter constraints that are at least as strong as the standard bispectrum estimator. Depending on the specific settings within the modal bispectrum pipeline, we find that as few as 10 modal expansion coefficients are necessary for galaxy bias and shot noise parameter constraints to converge when the smallest scale is given by $k_{\text{max}} \approx 0.10 h \text{Mpc}^{-1}$. While the modal decomposition method does require some additional calculations and machinery, we show that this overhead is small, especially when compared to the benefits of having a highly compressed data set. For example, in this work we use the modal bispectrum pipeline to show how bispectrum constraints can depend on the number of mock catalogs that are used to estimate the covariance matrix—a result that has not been attempted using the standard bispectrum estimator due to the extremely large number of mocks that would be required.

The outline of the paper is as follows. In Section 2 we review the modal decomposition method. In Section 3 we describe the data that we use, our likelihood modeling, and MCMC simulation details. We discuss our new results in Section 4 and summarize our main conclusions in Section 5.

2 Modal decomposition method

The basic premise of the modal decomposition method is that the bispectrum is a relatively smooth function of Fourier-space triangles, so the bispectrum that is normally measured in a large number of triangle bins, $N_{\text{triangles}}$, is very well approximated by a linear combination of a smaller number, N_{modes} , of basis mode functions. We write this as

$$w(k_1, k_2, k_3)B(k_1, k_2, k_3) \approx \sum_{n=0}^{N_{\text{modes}}-1} \beta_n^Q Q_n(k_1, k_2, k_3), \quad (2.1)$$

where $w(k_1, k_2, k_3)$ is a weighting function, $B(k_1, k_2, k_3)$ is the bispectrum, β_n^Q are a set of modal expansion coefficients, and $Q_n(k_1, k_2, k_3)$ are the modal basis functions. The modal coefficients therefore correspond to the amplitudes of template bispectra in the data.

The optimal estimator for the amplitude of a single bispectrum template for an isotropic and statistically homogeneous density field in the limit of weak non-Gaussianity is [69]

$$\hat{\varepsilon} = \frac{1}{N_\varepsilon} \int_{\mathbf{k}_1} \int_{\mathbf{k}_2} \int_{\mathbf{k}_3} (2\pi)^3 \delta_D(\mathbf{k}_{123}) B^{\text{template}}(k_1, k_2, k_3) \frac{\delta_{\mathbf{k}_1} \delta_{\mathbf{k}_2} \delta_{\mathbf{k}_3}}{P(k_1)P(k_2)P(k_3)}, \quad (2.2)$$

where we have introduced the shorthand notations $\int_{\mathbf{k}} \equiv \int \frac{d^3k}{(2\pi)^3}$ and $\mathbf{k}_{123} \equiv \mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3$ to make our expressions more compact. $\delta_{\mathbf{k}}$ is the observed density field on a discretized Fourier-space grid, $P(k_i)$ is the total power spectrum (including the shot noise), and N_ε is a normalization

constant.⁵ After defining the quantity

$$\hat{\mathcal{B}}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) \equiv \frac{\delta_{\mathbf{k}_1} \delta_{\mathbf{k}_2} \delta_{\mathbf{k}_3}}{V} \mathbf{1}_{\mathbf{k}_{123}}, \quad (2.3)$$

where V is the survey volume and $\mathbf{1}_{\mathbf{k}_{123}}$ is a Kronecker symbol that is unity if $\mathbf{k}_{123} = 0$ and zero otherwise, we see that the amplitude estimator is effectively a normalized weighted inner product between B^{template} and $\hat{\mathcal{B}}$,

$$\hat{\varepsilon} = \frac{V}{N_\varepsilon} \langle\langle wB^{\text{template}} | w\hat{\mathcal{B}} \rangle\rangle, \quad (2.4)$$

where the definition of the inner product is

$$\langle\langle wB^{\text{template}} | w\hat{\mathcal{B}} \rangle\rangle \equiv \int_{\mathbf{k}_1} \int_{\mathbf{k}_2} \int_{\mathbf{k}_3} (2\pi)^3 \delta_D(\mathbf{k}_{123}) \frac{wB^{\text{template}}(k_1, k_2, k_3) w\hat{\mathcal{B}}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)}{k_1 k_2 k_3} \quad (2.5)$$

and the weighting function is

$$w(k_1, k_2, k_3) \equiv \frac{\sqrt{k_1 k_2 k_3}}{\sqrt{P(k_1)P(k_2)P(k_3)}}. \quad (2.6)$$

Then the normalization constant must be $N_\varepsilon \equiv V \langle\langle wB^{\text{template}} | wB^{\text{template}} \rangle\rangle$ so that the ensemble average of the estimated amplitude $\langle\hat{\varepsilon}\rangle \rightarrow 1$ if $\langle\hat{\mathcal{B}}\rangle \rightarrow B^{\text{template}}$.

If the ensemble average $\langle\hat{\mathcal{B}}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)\rangle = B^{\text{obs}}(k_1, k_2, k_3)$, we can perform the angular integrals in the inner product analytically, following steps detailed in [60], which we also review here. We first use

$$\delta_D(\mathbf{k}_{123}) = \frac{1}{(2\pi)^3} \int d^3x e^{i\mathbf{k}_{123} \cdot \mathbf{x}} \quad (2.7)$$

and rewrite the exponential part as

$$e^{i\mathbf{k} \cdot \mathbf{x}} = 4\pi \sum_{\ell m} i^\ell j_\ell(kx) Y_{\ell m}(\hat{\mathbf{k}}) Y_{\ell m}^*(\hat{\mathbf{x}}) \quad (2.8)$$

to get

$$\begin{aligned} \langle\langle wB^{\text{template}} | wB^{\text{obs}} \rangle\rangle &= \int d^3x (4\pi)^3 \left[\int_{\mathbf{k}_1} \sum_{\ell_1 m_1} i^{\ell_1} j_{\ell_1}(k_1 x) Y_{\ell_1 m_1}(\hat{\mathbf{k}}_1) Y_{\ell_1 m_1}^*(\hat{\mathbf{x}}) \right] \\ &\quad \times \left[\int_{\mathbf{k}_2} \sum_{\ell_2 m_2} i^{\ell_2} j_{\ell_2}(k_2 x) Y_{\ell_2 m_2}(\hat{\mathbf{k}}_2) Y_{\ell_2 m_2}^*(\hat{\mathbf{x}}) \right] \\ &\quad \times \left[\int_{\mathbf{k}_3} \sum_{\ell_3 m_3} i^{\ell_3} j_{\ell_3}(k_3 x) Y_{\ell_3 m_3}(\hat{\mathbf{k}}_3) Y_{\ell_3 m_3}^*(\hat{\mathbf{x}}) \right] \\ &\quad \times \frac{wB^{\text{template}}(k_1, k_2, k_3) wB^{\text{obs}}(k_1, k_2, k_3)}{k_1 k_2 k_3}. \end{aligned} \quad (2.9)$$

The integral over $\hat{\mathbf{k}}_i$ inside each pair of square brackets is⁶

$$\int d\Omega_{\mathbf{k}_i} Y_{\ell_i m_i}(\hat{\mathbf{k}}_i) = \sqrt{4\pi} \delta_{\ell_i 0} \delta_{m_i 0}, \quad (2.10)$$

⁵We choose our Fourier transform convention so that the forward and backward transforms are $\delta(\mathbf{k}) = \int d^3x \delta(\mathbf{x}) e^{i\mathbf{k} \cdot \mathbf{x}}$ and $\delta(\mathbf{x}) = (2\pi)^{-3} \int d^3k \delta(\mathbf{k}) e^{-i\mathbf{k} \cdot \mathbf{x}}$.

⁶Our spherical harmonics are normalized such that $\int d\Omega_{\mathbf{k}} Y_{\ell m}(\hat{\mathbf{k}})^2 = 1$ and $Y_{00} = 1/\sqrt{4\pi}$.

which forces all ℓ_i and m_i in eq. (2.9) to be zero, giving

$$\begin{aligned} \langle\langle wB^{\text{template}}|wB^{\text{obs}}\rangle\rangle &= \int d^3x \frac{(4\pi)^{9/2}}{(2\pi)^9} \left[\int dk_1 k_1^2 j_0(k_1x) Y_{00}(\hat{\mathbf{x}}) \right] \\ &\quad \times \left[\int dk_2 k_2^2 j_0(k_2x) Y_{00}(\hat{\mathbf{x}}) \right] \\ &\quad \times \left[\int dk_3 k_3^2 j_0(k_3x) Y_{00}(\hat{\mathbf{x}}) \right] \\ &\quad \times \frac{wB^{\text{template}}(k_1, k_2, k_3) wB^{\text{obs}}(k_1, k_2, k_3)}{k_1 k_2 k_3}. \end{aligned} \quad (2.11)$$

In the final step, integration over \mathbf{x} using

$$\int dx x^2 j_0(k_1x) j_0(k_2x) j_0(k_3x) = \frac{\pi}{8k_1 k_2 k_3} \quad (2.12)$$

$$\int d\Omega_{\mathbf{x}} Y_{00}(\hat{\mathbf{x}})^3 = \frac{1}{\sqrt{2\pi}} \quad (2.13)$$

shows that the inner product is

$$\langle\langle wB^{\text{template}}|wB^{\text{obs}}\rangle\rangle = \frac{1}{8\pi^4} \int_{\mathcal{V}_T} dk_1 dk_2 dk_3 wB^{\text{template}}(k_1, k_2, k_3) wB^{\text{obs}}(k_1, k_2, k_3), \quad (2.14)$$

where the subscript \mathcal{V}_T signifies that the 3-dimensional integral must only cover the volume, sometimes called a *tetrapyd*, where (k_1, k_2, k_3) can form a closed triangle. Therefore, estimating the amplitude of a given template bispectrum is closely related to calculating a weighted inner product between the template and observed bispectrum over the tetrapyd space.

In previous literature on the modal bispectrum, the inner product in eq. (2.14) is sometimes written in terms of (x_1, x_2, x_3) instead of (k_1, k_2, k_3) , where $x_i \equiv (k_i - k_{\min}) / (k_{\max} - k_{\min})$, such that the allowed (x_1, x_2, x_3) form a tetrapyd that fits inside of a unit cube. We will sometimes use a different notation to define the inner product over this unit tetrapyd,

$$\langle f|g\rangle \equiv \int_{\mathcal{V}_T} dx_1 dx_2 dx_3 f(x_1, x_2, x_3) g(x_1, x_2, x_3). \quad (2.15)$$

2.1 Modal estimator

In this work, it is not the amplitude of one template that we are interested in, but the expansion coefficients of a general bispectrum on a chosen set of basis functions. In this case, we make the replacement $wB \rightarrow w\hat{B}$ in eq. (2.1) and take the inner product of both sides with Q_m to obtain

$$\langle\langle Q_m|w\hat{B}\rangle\rangle = \sum_{n=0}^{N_{\text{modes}}-1} \hat{\beta}_n^Q \gamma_{nm}, \quad (2.16)$$

where we have defined the positive-definite symmetric matrix $\gamma_{nm} \equiv \langle\langle Q_n|Q_m\rangle\rangle$.⁷ To estimate the modal coefficients, $\hat{\beta}_n^Q$, we first measure the N_{modes} inner products, $\langle\langle Q_n|w\hat{B}\rangle\rangle$ on the left-hand side, and solve the linear matrix equation in eq. (2.16).

⁷In [13], the inner product over the unit tetrapyd, $\bar{\gamma} \equiv \langle Q|Q\rangle$, was also defined and used to convert between γ matrices calculated over different k -ranges, given by k_{\min} and k_{\max} . This was motivated by the assumption that $\bar{\gamma}$ could be computed once, and thereafter γ over a general k -range could be computed as $\gamma = (k_{\max} - k_{\min})^3 / (8\pi^4) \bar{\gamma}$. However, we note that this was not quite correct; this rescaling can only be done when $k_{\min} = 0$, because in general γ depends on k_{\min} and k_{\max} in a way that cannot be factored out. We thank Dionysios Karagiannis for noticing this.

To make the measurement of

$$\langle\langle Q_n | w \hat{\mathcal{B}} \rangle\rangle = \frac{1}{V} \int_{\mathbf{k}_1} \int_{\mathbf{k}_2} \int_{\mathbf{k}_3} (2\pi)^3 \delta_D(\mathbf{k}_{123}) \frac{Q_n(k_1, k_2, k_3) \delta_{\mathbf{k}_1} \delta_{\mathbf{k}_2} \delta_{\mathbf{k}_3}}{\sqrt{k_1 k_2 k_3} \sqrt{P(k_1) P(k_2) P(k_3)}} \quad (2.17)$$

computationally tractable, we require that the Q_n basis functions can be written in separable form as a product of three 1-dimensional functions,

$$Q_n(k_1, k_2, k_3) = q_{\{p\}}(k_1) q_{\{r\}}(k_2) q_{\{s\}}(k_3). \quad (2.18)$$

The p , r , and s subscripts on the right side index the different 1-dimensional functions that we have chosen, and the curly brackets require that the Q_n functions are invariant to permutations of k_1 , k_2 , and k_3 . In Appendix A, we give additional details on how we compute the $q_n(k)$ from either normal or Legendre polynomials and how we have chosen the mapping between $\{prs\} \leftrightarrow n$.

Taking advantage of the separability of Q_n and using eq. (2.7) to rewrite the delta function in its exponential form, the final expression for $\langle\langle Q_n | w \hat{\mathcal{B}} \rangle\rangle$ simplifies into a computationally tractable expression written concisely as

$$\langle\langle Q_n | w \hat{\mathcal{B}} \rangle\rangle = \frac{1}{V} \int d^3 x M_{\{p\}}(\mathbf{x}) M_{\{r\}}(\mathbf{x}) M_{\{s\}}(\mathbf{x}), \quad (2.19)$$

where

$$M_r(\mathbf{x}) \equiv \int \frac{d^3 k}{(2\pi)^3} \frac{e^{i\mathbf{k}\cdot\mathbf{x}}}{\sqrt{k P(k)}} q_r(k) \delta_{\mathbf{k}}. \quad (2.20)$$

Therefore the inner product can be computed very efficiently using fast Fourier transform (FFT) routines (such as FFTW⁸ or Intel MKL⁹ libraries), if the basis of Q_n functions are multiplicative separable.

We note that this estimator requires minimal modifications to the standard bispectrum estimator, which takes the form

$$\begin{aligned} \hat{B}(k_1, k_2, k_3) &= \frac{V}{N_\Delta(k_1, k_2, k_3)} \int_{\mathbf{q}_1} \int_{\mathbf{q}_2} \int_{\mathbf{q}_3} (2\pi)^3 \delta_D(\mathbf{q}_{123}) \tilde{\Pi}_{k_1}(\mathbf{q}_1) \tilde{\Pi}_{k_2}(\mathbf{q}_2) \tilde{\Pi}_{k_3}(\mathbf{q}_3) \delta_{\mathbf{q}_1} \delta_{\mathbf{q}_2} \delta_{\mathbf{q}_3} \\ &= \frac{V}{N_\Delta(k_1, k_2, k_3)} \int d^3 x \mathcal{D}_{k_1}(\mathbf{x}) \mathcal{D}_{k_2}(\mathbf{x}) \mathcal{D}_{k_3}(\mathbf{x}), \end{aligned} \quad (2.21)$$

where

$$\mathcal{D}_k(\mathbf{x}) \equiv \int \frac{d^3 q}{(2\pi)^3} e^{i\mathbf{q}\cdot\mathbf{x}} \tilde{\Pi}_k(\mathbf{q}) \delta_{\mathbf{q}}, \quad (2.22)$$

and $\tilde{\Pi}_k(\mathbf{q})$ is a binning function that is 1 if $|\mathbf{q}| \in [k - \Delta k/2, k + \Delta k/2]$ and zero otherwise. N_Δ is the number of $(\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3)$ triangles that are averaged inside a (k_1, k_2, k_3) triangle bin,

$$\begin{aligned} N_\Delta(k_1, k_2, k_3) &\equiv V^2 \int_{\mathbf{q}_1} \int_{\mathbf{q}_2} \int_{\mathbf{q}_3} \delta_D(\mathbf{q}_{123}) \tilde{\Pi}_{k_1}(\mathbf{q}_1) \tilde{\Pi}_{k_2}(\mathbf{q}_2) \tilde{\Pi}_{k_3}(\mathbf{q}_3) \\ &= \frac{V^2}{(2\pi)^3} \int d^3 x \Pi_{k_1}(\mathbf{x}) \Pi_{k_2}(\mathbf{x}) \Pi_{k_3}(\mathbf{x}), \end{aligned} \quad (2.23)$$

⁸<http://www.fftw.org>

⁹<https://software.intel.com/content/www/us/en/develop/tools/math-kernel-library.html>

where $\Pi_k(\mathbf{x})$ is the inverse Fourier transform of $\tilde{\Pi}_k(\mathbf{q})$. Comparing the standard bispectrum estimator with the estimator for $\langle\langle Q_n | w \hat{\mathcal{B}} \rangle\rangle$, we see that they both require very similar computational steps, and the modal estimator recovers the bispectrum estimator by making the replacement $q_r(k)/\sqrt{kP(k)} \rightarrow \tilde{\Pi}_k(\mathbf{q})$. The critical difference, however, is that for a single realization, while the bispectrum estimator is computed once per (k_1, k_2, k_3) triangle bin, the modal estimator is computed once for each Q_n . Also, the memory requirement for the bispectrum estimator is such that each k_i bin requires a full Fourier grid to store the corresponding $\mathcal{D}_{k_i}(\mathbf{x})$, while for the modal estimator each 1-dimensional basis function q_r requires its own grid to store $M_r(\mathbf{x})$. Therefore the modal estimator is typically more computationally efficient compared to the standard bispectrum estimator.

Once the $\langle\langle Q_n | w \hat{\mathcal{B}} \rangle\rangle$ have been measured, we estimate $\hat{\beta}_n^Q$ by numerically solving the linear equation in eq. (2.16). Using the $\hat{\beta}_n^Q$ coefficients, we can also calculate a reconstructed bispectrum as

$$B_{\text{rec}} = \frac{1}{w} \sum_n \hat{\beta}_n^Q Q_n. \quad (2.24)$$

Later, in Section 4.2, we compare this bispectrum, B_{rec} , with the bispectrum measured using the standard estimator, \hat{B} .

We note that the γ matrix only needs to be computed once for a desired wavenumber range (k_{\min}, k_{\max}) and choice of Q_n basis functions. Different methods for calculating the inner products in γ have been discussed to date in the literature, and one of the goals of this work is to compare these methods. In the next section, we summarize the inner product methods that we implement and compare in this work.

2.2 Inner product methods

Here, we briefly describe the four different methods we have implemented in this work for calculating the inner product matrix

$$\gamma \equiv \langle\langle Q_n | Q_m \rangle\rangle = \frac{1}{8\pi^4} \int_{\mathcal{V}_T} dk_1 dk_2 dk_3 Q_n(k_1, k_2, k_3) Q_m(k_1, k_2, k_3). \quad (2.25)$$

Monte Carlo integration

This method uses the Monte Carlo algorithm called Vegas included in the public Cuba library for multidimensional numerical integration [70, 71] to calculate the inner product via random sampling of the 3-dimensional tetrapyd space. The free parameters for this method are the convergence tolerance and the maximum number of samples. In addition to being very slow to converge, we find it quite challenging, despite different choices in the integration parameters, to avoid a non-positive definite γ , which cannot be used for the modal analysis pipeline.

Voxels

This method divides the cubic volume of (k_1, k_2, k_3) from k_{\min} up to k_{\max} into a grid of smaller cubes, called *voxels*, and calculates eq. (2.25) by integrating over each voxel using tri-linear interpolation of the integrand within each voxel (as described in Appendix A2 of [13]). Because of the shape of the tetrapyd volume, some care must be taken to properly integrate over voxels that intersect with the tetrapyd boundary. The only free parameter of this method is the grid resolution set by the number of individual voxels, N_v , spanning the chosen k -range in each of the three dimensions.

3D FFT

This method calculates the inner product in the same way that the modal estimator does: we take the expression for $\langle\langle Q_n | w\hat{\mathcal{B}} \rangle\rangle$ in eq. (2.17) and make the replacement $w\hat{\mathcal{B}} \rightarrow Q_m$ to find

$$\langle\langle Q_n | Q_m \rangle\rangle = \int_{\mathbf{k}_1} \int_{\mathbf{k}_2} \int_{\mathbf{k}_3} (2\pi)^3 \delta_D(\mathbf{k}_{123}) \frac{Q_n(k_1, k_2, k_3) Q_m(k_1, k_2, k_3)}{k_1 k_2 k_3} \quad (2.26)$$

$$= \int d^3x \int_{\mathbf{k}_1} \int_{\mathbf{k}_2} \int_{\mathbf{k}_3} e^{i\mathbf{k}_{123} \cdot \mathbf{x}} \frac{q_{\{p\}}(k_1) q_{\{r\}}(k_2) q_{\{s\}}(k_3) q_{\{a\}}(k_1) q_{\{b\}}(k_2) q_{\{c\}}(k_3)}{k_1 k_2 k_3}. \quad (2.27)$$

Then, similarly to eq. (2.19), we write this in a compact form as [65]

$$\begin{aligned} \langle\langle Q_n | Q_m \rangle\rangle = \frac{1}{6} \int d^3x \{ & M_{pa}(\mathbf{x}) [M_{rb}(\mathbf{x}) M_{sc}(\mathbf{x}) + M_{rc}(\mathbf{x}) M_{sb}(\mathbf{x})] \\ & + M_{pb}(\mathbf{x}) [M_{rc}(\mathbf{x}) M_{sa}(\mathbf{x}) + M_{ra}(\mathbf{x}) M_{sb}(\mathbf{x})] \\ & + M_{pc}(\mathbf{x}) [M_{ra}(\mathbf{x}) M_{sb}(\mathbf{x}) + M_{rb}(\mathbf{x}) M_{sa}(\mathbf{x})] \}, \end{aligned} \quad (2.28)$$

where we have defined

$$M_{pa}(\mathbf{x}) \equiv \int_{\mathbf{k}} e^{i\mathbf{k} \cdot \mathbf{x}} \frac{q_p(k) q_a(k)}{k}. \quad (2.29)$$

Like the modal estimator that we have already discussed, this expression for $\langle\langle Q_n | Q_m \rangle\rangle$ can be computed quickly using existing FFT software. The free parameters of this method are, as with any discrete Fourier transform, the real-space volume and the FFT grid resolution.

1D FFT¹⁰

This method computes the inner product using 1-dimensional FFTs by evaluating the expression in eq. (2.11) after the replacements $wB^{\text{template}} \rightarrow Q_n$ and $wB^{\text{obs}} \rightarrow Q_m$. In this case, using $j_0(k_i x) = \sin(k_i x)/k_i x$ and $Y_{00} = 1/\sqrt{4\pi}$, the inner product becomes

$$\begin{aligned} \langle\langle Q_n | Q_m \rangle\rangle = \frac{1}{2\pi^5} \int dx \frac{1}{x} \{ & F_{pa}(x) [F_{rb}(x) F_{sc}(x) + F_{rc}(x) F_{sb}(x)] \\ & + F_{pb}(x) [F_{rc}(x) F_{sa}(x) + F_{ra}(x) F_{sb}(x)] \\ & + F_{pc}(x) [F_{ra}(x) F_{sb}(x) + F_{rb}(x) F_{sa}(x)] \}, \end{aligned} \quad (2.30)$$

where

$$F_{pa}(x) \equiv \int dk q_p(k) q_a(k) \sin(kx). \quad (2.31)$$

The integral over k in $F_{pa}(x)$ can be performed using 1-dimensional FFTs, as described in Chapter 13.9 of *Numerical Recipes* [72], while the outermost 1-dimensional integral over x can be done using standard numerical integration methods (in our case, the Cuhre routine included in the Cuba library). We have deferred the numerical details of this calculation to Appendix B. This method has two free parameters corresponding to the grid resolutions in k and x , and we find that the resulting γ is positive-definite only when these resolutions are sufficiently high.

¹⁰We thank Dionysios Karagiannis for suggesting the 1D FFT method.

Summary

We have described four different methods for computing the inner product matrix, $\gamma \equiv \langle\langle Q|Q \rangle\rangle$. The Monte Carlo routine, called Vegas in the Cuba library, fails to converge to positive-definite γ , so we do not use it subsequently in this work. The remaining three methods we have implemented are calculated independently and give numerically different results for γ , stemming from the fact that each method makes different assumptions and approximations about the inner product. The voxel and 1D FFT methods assume that the \mathbf{k}_i wavevectors are sampled very finely, such that the inner product is effectively a continuous integral. These two methods are still completely different in their numerical implementation. In contrast, the 3D FFT method assumes each \mathbf{k}_i is discretely sampled in three dimensions, and so it is the only method that accounts explicitly for the discrete sampling of Fourier space, treating this sampling in the same way that the modal estimator is applied to the data through $\langle\langle Q_n|w\mathcal{B} \rangle\rangle$ in eq. (2.19). Our benchmark constraints use the 3D FFT method, and part of Section 4.4 investigates whether the methods described here have an impact on the resulting parameter constraints.

2.3 Bispectrum model and custom modes

We use a tree-level standard perturbation theory (SPT) model for the real-space halo bispectrum that matches the modeling in [68]:

$$B_h(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) = b_1^3 B_m(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) + b_2 b_1^2 \Sigma(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) + 2\gamma_2 b_1^2 K(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) + \frac{1 + \alpha_1}{\bar{n}} b_1^2 [P_L(k_1) + P_L(k_2) + P_L(k_3)] + \frac{1 + \alpha_2}{\bar{n}^2}, \quad (2.32)$$

where the tree-level matter bispectrum is

$$B_m(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) = 2[F_2(\mathbf{k}_1, \mathbf{k}_2)P_L(k_1)P_L(k_2) + F_2(\mathbf{k}_1, \mathbf{k}_3)P_L(k_1)P_L(k_3) + F_2(\mathbf{k}_2, \mathbf{k}_3)P_L(k_2)P_L(k_3)] \quad (2.33)$$

and $P_L(k)$ is the linear matter power spectrum. b_1 and b_2 are the linear and quadratic bias parameters, while γ_2 is the tidal bias. The shot noise terms in the second row of eq. (2.32) are parametrized by α_1 and α_2 , such that $\alpha_1 = \alpha_2 = 0$ correspond to Poissonian shot noise. The kernel definitions are

$$F_2(\mathbf{k}_1, \mathbf{k}_2) \equiv \frac{5}{7} + \frac{1}{2} \frac{\mathbf{k}_1 \cdot \mathbf{k}_2}{k_1 k_2} \left(\frac{k_1}{k_2} + \frac{k_2}{k_1} \right) + \frac{2}{7} \left(\frac{\mathbf{k}_1 \cdot \mathbf{k}_2}{k_1 k_2} \right)^2 \quad (2.34)$$

$$\Sigma(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) \equiv P_L(k_1)P_L(k_2) + P_L(k_1)P_L(k_3) + P_L(k_2)P_L(k_3) \quad (2.35)$$

$$K(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) \equiv [(\hat{\mathbf{k}}_1 \cdot \hat{\mathbf{k}}_2)^2 - 1] P_L(k_1)P_L(k_2) + [(\hat{\mathbf{k}}_1 \cdot \hat{\mathbf{k}}_3)^2 - 1] P_L(k_1)P_L(k_3) + [(\hat{\mathbf{k}}_2 \cdot \hat{\mathbf{k}}_3)^2 - 1] P_L(k_2)P_L(k_3). \quad (2.36)$$

In this work, we consider the M5 model in [68], where the cosmological parameters are fixed and we only vary the bias and shot noise parameters, $(b_1, b_2, \gamma_2, \alpha_1, \alpha_2)$. This allows the theory predictions for the halo bispectrum to be computed very quickly, as a linear combination of precomputed terms corresponding to the B_m , Σ , K , and P_L terms in eq. (2.32) after accounting for binning effects. Similarly, in this work we require fast

theoretical predictions for the modal coefficients, β_n^Q , and we achieve this by taking advantage of *custom modes*.

First proposed in [65], custom modes are a set of four separable basis functions that by design reproduce exactly the tree-level matter bispectrum. We can see that this should be possible by rewriting the F_2 perturbation theory kernel in eq. (2.34) as

$$F_2(k_1, k_2, k_3) = \frac{5}{7} + \frac{1}{2} \left(\frac{k_3^2 - k_1^2 - k_2^2}{2k_1 k_2} \right) \left(\frac{k_1}{k_2} + \frac{k_2}{k_1} \right) + \frac{2}{7} \left(\frac{k_3^2 - k_1^2 - k_2^2}{2k_1 k_2} \right)^2, \quad (2.37)$$

where, after expanding this expression, we see that each term will be separable in k_1 , k_2 , and k_3 . More specifically, the weighted bispectrum, wB_m , can be written as a linear combination of four modes,

$$Q_0^{\text{tree}}(k_1, k_2, k_3) = q_{\{0\}}^{\text{tree}}(k_1) q_1^{\text{tree}}(k_2) q_{\{1\}}^{\text{tree}}(k_3) \quad (2.38)$$

$$Q_1^{\text{tree}}(k_1, k_2, k_3) = q_{\{0\}}^{\text{tree}}(k_1) q_2^{\text{tree}}(k_2) q_{\{3\}}^{\text{tree}}(k_3) \quad (2.39)$$

$$Q_2^{\text{tree}}(k_1, k_2, k_3) = q_{\{1\}}^{\text{tree}}(k_1) q_3^{\text{tree}}(k_2) q_{\{4\}}^{\text{tree}}(k_3) \quad (2.40)$$

$$Q_3^{\text{tree}}(k_1, k_2, k_3) = q_{\{3\}}^{\text{tree}}(k_1) q_3^{\text{tree}}(k_2) q_{\{5\}}^{\text{tree}}(k_3), \quad (2.41)$$

where the custom 1-dimensional basis functions are

$$q_0^{\text{tree}}(k) = \sqrt{\frac{k}{P(k)}} \frac{5}{14} \quad (2.42)$$

$$q_1^{\text{tree}}(k) = \sqrt{\frac{k}{P(k)}} P_L(k) \quad (2.43)$$

$$q_2^{\text{tree}}(k) = -\sqrt{\frac{k}{P(k)}} P_L(k) k^2 \quad (2.44)$$

$$q_3^{\text{tree}}(k) = \sqrt{\frac{k}{P(k)}} \frac{P_L(k)}{k^2} \quad (2.45)$$

$$q_4^{\text{tree}}(k) = \sqrt{\frac{k}{P(k)}} \frac{3}{14} k^2 \quad (2.46)$$

$$q_5^{\text{tree}}(k) = \sqrt{\frac{k}{P(k)}} \frac{1}{14} k^4. \quad (2.47)$$

We note that it is important to distinguish between $P_L(k)$, which is the linear power spectrum appearing in the tree-level matter bispectrum model, from $P(k)$ (without a subscript) which is the power spectrum appearing in the definition of the weighting function, $w(k_1, k_2, k_3)$. In addition to the four custom modes above, in this work we add two more,

$$Q_4^{\text{tree}}(k_1, k_2, k_3) = q_{\{0\}}^{\text{tree}}(k_1) q_0^{\text{tree}}(k_2) q_{\{1\}}^{\text{tree}}(k_3) \quad (2.48)$$

$$Q_5^{\text{tree}}(k_1, k_2, k_3) = q_{\{0\}}^{\text{tree}}(k_1) q_0^{\text{tree}}(k_2) q_{\{0\}}^{\text{tree}}(k_3), \quad (2.49)$$

to model the shot noise terms.

With these definitions for six custom modes in total, Q_n^{tree} for $n = 0, \dots, 5$, we can reproduce the tree-level halo bispectrum model in eq. (2.32) corresponding to values of the

$\beta_0^{\text{tree}} =$	$6b_1^3 + \frac{42}{5}b_1^2b_2 - \frac{42}{5}b_1^2\gamma_2$
$\beta_1^{\text{tree}} =$	$6b_1^3 - \frac{42}{5}b_1^2\gamma_2$
$\beta_2^{\text{tree}} =$	$6b_1^3 - 28b_1^2\gamma_2$
$\beta_3^{\text{tree}} =$	$6b_1^3 + 21b_1^2\gamma_2$
$\beta_4^{\text{tree}} =$	$\frac{588}{25}b_1^2\left(\frac{1+\alpha_1}{n}\right)$
$\beta_5^{\text{tree}} =$	$\frac{2744}{125}\left(\frac{1+\alpha_2}{n^2}\right)$

Table 1: For an input model with free parameters $(b_1, b_2, \gamma_2, \alpha_1, \alpha_2)$, we list the modal coefficients for the six custom modes, Q_n^{tree} for $n = 0, \dots, 5$, that reproduce the tree-level halo bispectrum model in eq. (2.32).

parameters $(b_1, b_2, \gamma_2, \alpha_1, \alpha_2)$ by choosing the modal coefficients, β_n^{tree} , as shown in Table 1, i.e.

$$wB_h(k_1, k_2, k_3) = \sum_{n=0}^5 \beta_n^{\text{tree}} Q_n^{\text{tree}}(k_1, k_2, k_3). \quad (2.50)$$

The fact that we can write the model predictions for β_n^{tree} as trivial functions of $(b_1, b_2, \gamma_2, \alpha_1, \alpha_2)$ means that, as in the standard bispectrum analysis of [68], the calculation of the model predictions is very fast, allowing for MCMC simulations to run quickly.

We state for emphasis that this expansion of the halo bispectrum model is *exact*—it is not an approximation. In our subsequent analysis, unless otherwise mentioned explicitly, we always use a basis where the first six basis functions are these Q_n^{tree} , and starting with the seventh basis function, we use the Q_n that we have introduced earlier, which are either constructed from 1-dimensional normal or Legendre polynomials, which we call $q_n(k)$.

For a general model of the halo bispectrum or a general cosmological parameter set, it may not be possible to predict theoretical values of β_n^Q as quickly as what we use here. Finding a general strategy to manage this problem is outside the scope of this work, but it is an interesting challenge for future work. We note that this computational bottleneck has a counterpart in the standard bispectrum pipeline, where it is necessary to quickly calculate predictions for $B(k_1, k_2, k_3)$ in all triangle bins, accounting for the bin size, for a general bispectrum model and parameter set.

2.4 Orthonormal basis

Finally, in this section we introduce one more set of basis functions that are rotations of any general set of separable Q_n (which may include custom modes). We label this new basis R_n and call it the orthonormal basis because it satisfies

$$\langle\langle R_n | R_m \rangle\rangle = \frac{(k_{\text{max}} - k_{\text{min}})^3}{8\pi^4} \delta_{nm}. \quad (2.51)$$

We comment that the factor of $(k_{\text{max}} - k_{\text{min}})^3 / (8\pi^4)$ on the right hand side only changes the overall amplitude of all R_n by a constant factor, and it is only present here because we made the arbitrary choice to require that the R_n basis is orthonormal in the unit tetrapyd space, i.e. $\langle R_n | R_m \rangle = \delta_{nm}$.

Then we require that R_n is a linear combination of the Q_n basis as

$$R_n \equiv \sum_m \lambda_{nm}^{-1} Q_m, \quad (2.52)$$

and by deriving $\langle\langle R|R \rangle\rangle$ and setting it equal to the identity matrix, we see that γ and λ are related by

$$\gamma = \frac{(k_{\max} - k_{\min})^3}{8\pi^4} \lambda \cdot \lambda^T, \quad (2.53)$$

where λ is the lower triangular matrix resulting from the Cholesky decomposition.

Fig. 1 shows the six custom modes Q_n^{tree} and the first six R_n for $k_{\max} \approx 0.10 h \text{ Mpc}^{-1}$, with the 4-dimensional data represented in a plot similar in style to [57].¹¹ The three axes in each plot are the x_i defined by $(k_i - k_{\min})/(k_{\max} - k_{\min})$, where the origin is in the lower left corner. To focus on the overall triangle-dependence of each basis function, we have normalized each one to equal unity when $x_1 = x_2 = x_3 = 1$ in the upper right corner. We have also removed the half of the tetrapyd with $x_1 > x_2$, if the vertical axis is x_3 , to show the interior of the tetrapyd region. We notice that the Q_1^{tree} , Q_2^{tree} , and Q_3^{tree} modes are similar and peak (in red) at the edges of the tetrapyd corresponding to squeezed triangles, while the other three custom modes are largest at equilateral triangles with $k_1 = k_2 = k_3 = k_{\max}$. The plots showing Q_0^{tree} and R_0 are identical because $R_0 \propto Q_0^{\text{tree}}$ by construction. Unlike the Q_n^{tree} , however, all of the R_n look dissimilar because they have been defined to be orthogonal to each other, i.e. $\langle R_n | R_m \rangle = \delta_{nm}$.

Using the fact that the weighted bispectrum must be the same regardless of the basis,

$$\sum_n \beta_n^Q Q_n = \sum_m \beta_m^R R_m, \quad (2.54)$$

we see that the two sets of coefficients are related by

$$\beta^R = \lambda^T \cdot \beta^Q. \quad (2.55)$$

In the modal pipeline of this work, it is never necessary to calculate anything with R_n itself directly. We just use its definition to work with the β^R coefficients as our data, rather than the β^Q . Whether we do our MCMC analysis in terms of β^Q or β^R does not matter, but the β^R have the advantage that the numerical values of the β_n^R coefficients for $n < N_{\text{modes}}$ will not change if the size of the basis, given by N_{modes} , is increased. This is because λ^{-1} in eq. (2.52) is lower triangular, so R_n only depends on Q_m with $m \leq n$, and β^R can also be expressed as $\langle\langle R_n | w\mathcal{B} \rangle\rangle = (k_{\max} - k_{\min})^3 / (8\pi^4) \beta_n^R$. On the other hand, we can see from $\langle\langle Q | w\mathcal{B} \rangle\rangle = \gamma \cdot \beta^Q$ that all numerical values of β^Q will change as the basis set is increased.

The definition of R_n in eq. (2.51) corresponds to defining β^R that are orthogonal in the limit of Gaussian covariance. We note that the β_n^R can be written as

$$\beta_n^R = \frac{8\pi^4}{(k_{\max} - k_{\min})^3} \langle\langle R_n | w\mathcal{B} \rangle\rangle \quad (2.56)$$

$$= \frac{8\pi^4}{(k_{\max} - k_{\min})^3} \frac{1}{V} \int_{\mathbf{k}_1} \int_{\mathbf{k}_2} \int_{\mathbf{k}_3} (2\pi)^3 \delta_D(\mathbf{k}_{123}) \frac{R_n(k_1, k_2, k_3) \delta_{\mathbf{k}_1} \delta_{\mathbf{k}_2} \delta_{\mathbf{k}_3}}{\sqrt{k_1 k_2 k_3} \sqrt{P(k_1) P(k_2) P(k_3)}}, \quad (2.57)$$

¹¹The R_n shown in the figure were calculated the default modal settings described in detail at the beginning of Section 4.




Figure 1: Plots of the six custom Q_n basis functions (first two rows) and the first six R_n basis functions (bottom two rows) for $k_{\max} \approx 0.10 h \text{ Mpc}^{-1}$. Each function is plotted over the unit tetrapyd of allowed (x_1, x_2, x_3) combinations, with the origin in the lower left corner, and for easier readability is normalized to unity at $x_1 = x_2 = x_3 = 1$ in the upper right corner. These plots show only half of the tetrapyd to show more of the interior region.

such that the covariance $\langle \beta_n^R \beta_m^R \rangle$ requires evaluating $\langle \delta_{\mathbf{k}_1} \delta_{\mathbf{k}_2} \delta_{\mathbf{k}_3} \delta_{\mathbf{k}'_1} \delta_{\mathbf{k}'_2} \delta_{\mathbf{k}'_3} \rangle$. The leading-order Gaussian contribution to this is

$$\langle \delta_{\mathbf{k}_1} \delta_{\mathbf{k}_2} \delta_{\mathbf{k}_3} \delta_{\mathbf{k}'_1} \delta_{\mathbf{k}'_2} \delta_{\mathbf{k}'_3} \rangle_G = 6(2\pi)^9 \delta_D(\mathbf{k}_1 + \mathbf{k}'_1) \delta_D(\mathbf{k}_2 + \mathbf{k}'_2) \delta_D(\mathbf{k}_3 + \mathbf{k}'_3) P(k_1) P(k_2) P(k_3), \quad (2.58)$$

such that the Gaussian covariance for the modal coefficients is

$$\langle \beta_n^R \beta_m^R \rangle_G = \frac{6}{V} \left[\frac{8\pi^4}{(k_{\max} - k_{\min})^3} \right]^2 \langle R_n | R_m \rangle \quad (2.59)$$

$$= \frac{6}{V} \frac{8\pi^4}{(k_{\max} - k_{\min})^3} \delta_{nm}. \quad (2.60)$$

Non-Gaussian contributions to the covariance will generally couple orthonormal modal coefficients with different n and m . Later, in Section 4.6, we evaluate the impact of assuming the Gaussian covariance on the parameter constraints.

2.5 Summary

Here we put together the practical steps of the modal method necessary to implement it, and summarize the expressions necessary to estimate and model modal coefficients, β^Q and β^R , with or without custom modes included.

In the first step, we choose a k -range, bounded by k_{\min} and k_{\max} , and we choose the 1-dimensional basis functions $q_n(k)$ that are combined to get the separable basis of Q_n , which may or may not include custom modes, depending on the choice of the user and the bispectrum model. After the Q_n have been defined, γ is computed for this basis, using a chosen method—we have discussed four options in Section 2.2. Then we use the Cholesky decomposition to numerically calculate λ , which defines the orthonormal R_n basis.

In the second step, we obtain our measurements from simulations. This is done by measuring $\langle Q | w \mathcal{B} \rangle$, and then solving

$$\langle Q | w \mathcal{B} \rangle = \gamma \cdot \hat{\beta}^Q \quad (2.61)$$

$$\langle Q | w \mathcal{B} \rangle = \frac{(k_{\max} - k_{\min})^3}{8\pi^4} \lambda \cdot \hat{\beta}^R \quad (2.62)$$

to get the modal coefficients.

In the last step, we need a function for predicting $\beta^Q(\theta)$ and $\beta^R(\theta)$. When custom modes are used, we set the modal coefficients to be

$$\beta_n^Q(\theta) = \begin{cases} \beta_n^{\text{tree}}(\theta) & \text{if } 0 \leq n \leq 5 \\ 0 & \text{otherwise} \end{cases} \quad (2.63)$$

$$\beta^R(\theta) = \lambda^T \cdot \beta^Q(\theta). \quad (2.64)$$

When custom modes are *not* used, we have to estimate the β^Q in

$$\sum_{n=0}^5 \beta_n^{\text{tree}}(\theta) Q_n^{\text{tree}} \approx \sum_m \beta_m^Q Q_m, \quad (2.65)$$

where the sum over m on the right side does not include any custom modes. Taking the inner product of this with Q on both sides (where Q again does not include any custom modes), we find that we need to solve

$$\langle Q | Q^{\text{tree}} \rangle \cdot \beta^{\text{tree}}(\theta) = \gamma \cdot \beta^Q(\theta) \quad (2.66)$$

$$\langle Q | Q^{\text{tree}} \rangle \cdot \beta^{\text{tree}}(\theta) = \frac{(k_{\max} - k_{\min})^3}{8\pi^4} \lambda \cdot \beta^R(\theta) \quad (2.67)$$

for β^Q and/or β^R . Therefore, when custom modes are not included in the modal basis, we need to precompute the rectangular matrix $\langle Q | Q^{\text{tree}} \rangle$ as a means to obtaining theory predictions quickly.

3 Data and analysis

3.1 Simulations and mock halo catalogs

We use the same simulation and halo catalog data as in [68], since the aim of our work is to compare the modal bispectrum constraints with the results from the standard bispectrum in that work. The data consist of two sets of simulations. The first is a suite of 298 N -body simulations, called Minerva, created using the Gadget-2 code and first presented in [73]. Each realization is a $L = 1500 h^{-1} \text{Mpc}$ box evolved to $z = 1$ based on the same fiducial flat ΛCDM cosmology. Halos are identified using a friends-of-friends algorithm such that the minimum halo mass is $1.12 \times 10^{13} h^{-1} M_{\odot}$, and the mean number density is $\bar{n} = 2.13 \times 10^{-4} h^3 \text{Mpc}^{-3}$. The measurements from these Minerva simulations are the data that we fit in our analysis.

We also use a set of 10,000 mock halo catalogs generated using the approximate N -body code Pinocchio [74–76], which in [68] were used to obtain bispectrum covariance matrices. 298 of these realizations have initial conditions that match those of the Minerva realizations. The halos in the Pinocchio simulations were chosen with a different mass threshold (compared to the Minerva data) such that the large-scale amplitude of the total halo power spectrum matches what is measured in the Minerva N -body simulations, because in the Gaussian limit the bispectrum covariance depends directly on the total power spectrum [68].

For the modal estimator, as with the bispectrum measurements in [68], we map the halo positions to the grid of halo density values using the 4th-order interlacing method in [77] obtained with the public PowerI4¹² code and run the estimator using a FFT grid size of $N_g = 256$.

We note that by simultaneously fitting to 298 Minerva realizations, the results we present in Section 4 correspond to a total volume of $\approx 1,000 h^{-3} \text{Gpc}^3$, which is much larger than any real galaxy survey planned for the near future. Therefore the results we present should be interpreted as a proof of principle of the modal bispectrum method, while the exact numerical values of the parameter constraints will change for more realistic survey scenarios in smaller volumes.

3.2 Likelihood and MCMC

We implement two likelihood functions: a Gaussian likelihood, which we take to be our benchmark case, and the likelihood proposed by Sellentin and Heavens in [78] (SH in what follows). The two likelihoods differ in how they account for errors in the estimated covariance matrix, due to the fact that it is estimated from a finite number of mocks, but they both assume that the observable data is Gaussian distributed. We have checked that the probability distribution functions of the β_n^R modal coefficients measured in the N -body simulations and mock catalogs do not show any strong indications of non-Gaussianity. In Section 4.7, we compare results from the two likelihoods.

The Gaussian likelihood for our analysis is

$$\ln \mathcal{L} = -\frac{1}{2} \sum_n \sum_m \delta\beta_n^R \hat{C}_{nm}^{-1} \delta\beta_m^R \equiv -\frac{1}{2} (\delta\beta^R)^T \cdot \hat{C}^{-1} \cdot \delta\beta^R, \quad (3.1)$$

where $\delta\beta_n^R \equiv \beta_n^R(\theta) - \beta_n^{R,\text{sims}}$. The covariance matrix estimated from N_s mock catalogs is

$$\tilde{C}_{nm} \equiv \frac{1}{N_s - 1} \sum_i^{N_s} (\beta_n^{R(i)} - \bar{\beta}_n^R) (\beta_m^{R(i)} - \bar{\beta}_m^R). \quad (3.2)$$

¹²<https://github.com/sefusatti/PowerI4>

Though this is an unbiased estimator of the covariance matrix, taking the inverse of this \tilde{C}_{nm} will result in a biased estimate of the precision matrix, which can be statistically debiased by a multiplicative factor [79–81],

$$\hat{C}^{-1} = \Gamma \tilde{C}^{-1}, \quad (3.3)$$

where

$$\Gamma \equiv \frac{N_s - N_{\text{bins}} - 2}{N_s - 1} \quad (3.4)$$

and N_{bins} is the number of data bins. We note however that any one particular \tilde{C} will have statistical noise such that applying the Γ factor may actually bring the final parameter constraints closer to, or further away from, what we would have obtained using the true precision matrix. In other words, the Γ factor does not take into account the statistical nature of the estimated precision matrix—the fact that it is still a noisy estimate of an unknown quantity.

Instead, SH derives a likelihood that is the Gaussian likelihood marginalized over the unknown covariance matrix, conditioned on our estimate of it, to arrive at

$$\ln \mathcal{L} = -\frac{N_s}{2} \ln \left[1 + \frac{(\delta\beta^R)^T \cdot \tilde{C}^{-1} \cdot \delta\beta^R}{N_s - 1} \right] + \ln \left(\frac{\bar{c}_p}{\sqrt{\det \tilde{C}}} \right). \quad (3.5)$$

\bar{c}_p is a constant that depends only on N_s and N_{bins} , and since we assume the covariance matrix does not depend on the parameters, we drop the second $\ln(\dots)$ term. For one particular estimate of the covariance matrix, this SH likelihood will also have errors that are too large or too small, and be biased. However, on average the SH likelihood should yield parameter constraints that are closer to the true one.

The two likelihoods should equal each other, and approach the true answer, in the limit that the covariance matrix is well-estimated by a sufficiently large N_s . Conversely, they will show differences when N_s is small, e.g. $\Gamma \ll 1$.

As in [68], we simultaneously fit all 298 Minerva realizations, which means that our total likelihood is

$$\ln \mathcal{L}_{\text{total}} = \sum_i^{N_R} \ln \mathcal{L}_i, \quad (3.6)$$

where $\ln \mathcal{L}_i$ is the likelihood for one realization, and equal to eq. (3.1) or eq. (3.5).

We use wide, uniform priors for all parameters, and we compare our results with those from [68] using their ‘broad’ priors, which were $b_1 \in [0.5, 5]$, $b_2, \gamma \in [-5, 5]$, $\alpha_1 \in [-10, 10]$, and $\alpha_2 \in [-100, 100]$. The analysis in [68] also explored the effect of narrower priors on α_1 and α_2 , as well as fitting the simulation data with models with fewer than five parameters. However, in this work we only consider the five parameter model (called M5 in [68]) with broad priors, because we expect that a modal estimator pipeline fitting for more parameters with less informative priors will present a stronger test of the modal method.

Our MCMC simulations are run using the code `emcee`¹³ [82]. Each chain uses 100 walkers that are started within a small sphere around an approximate maximum likelihood point. Our convergence criteria are that the integrated autocorrelation time τ is stable to within 1% and that the chain is at least 50τ long. After the chains have converged, we use the `getdist`¹⁴ package to analyze the chains and produce contour plots of the posteriors [83].

¹³<https://emcee.readthedocs.io>

¹⁴<https://getdist.readthedocs.io>

4 Results

Unless otherwise stated, the results we present use following default settings for the modal estimator pipeline. The first six Q_n basis functions are the six custom modes defined in Section 2.3, and the rest of the Q_n basis functions are constructed from 1-dimensional q_n functions that are normal polynomials. The basis functions are defined over the wavenumber range with $k_{\min} = 1.5 k_f \approx 0.006 h \text{ Mpc}^{-1}$ and one of two choices of k_{\max} , $13.5 k_f \approx 0.06 h \text{ Mpc}^{-1}$ or $24.5 k_f \approx 0.10 h \text{ Mpc}^{-1}$, where $k_f = 2\pi/L$ is the fundamental wavenumber and $L = 1500 h^{-1} \text{ Mpc}$ is the size of a simulation box. We use the 3D FFT method to compute the inner product matrix γ , where the grid resolution is $N_g = 256$ and the size of the real-space Fourier volume matches that of the simulation boxes. The power spectrum that appears in the weighting which defines the basis is the average total halo power spectrum P_h measured in the Minerva simulations.

We present our data consisting of the measured modal coefficients in Section 4.1. Sections 4.2 and 4.3 compare the modal bispectrum and standard bispectrum estimators, first by considering the similarities and differences in the bispectra that are measured, and then by comparing their resulting parameter constraints. The subsequent sections then focus exclusively on the modal bispectrum. In Section 4.4, we explore how the parameter constraints from the modal bispectrum can depend on a variety of settings within the modal bispectrum pipeline, while in Section 4.5 we discuss other ways that the convergence of the modal expansion could be estimated. Sections 4.6 and 4.7 quantify how our results depend on the estimated covariance matrices and likelihood modeling that are used in our analysis.

4.1 Mean and covariance of modal coefficients

In Fig. 2, we show the means, $\bar{\beta}_n^R$, and errors, $\Delta\beta_n^R$, of modal coefficients that are measured in the 298 realizations of Minerva N -body simulations for $k_{\max} = 13.5 k_f$ and $24.5 k_f$, where $k_{\min} = 1.5 k_f$ in both cases. The dashed gray horizontal lines show the Gaussian predictions for the error in eq. (2.60), which depends on k_{\max} .

$\bar{\beta}_n^R$ for the first few n are more easily detected, and if k_{\max} is higher, more of the first few n have a higher signal-to-noise ratio, $\bar{\beta}_n^R/\Delta\beta_n^R$: for $k_{\max} = 13.5 k_f$ ($24.5 k_f$), two (five) modes have a signal-to-noise ratio greater than 1. The errors on the coefficients measured from simulations are in good agreement with the Gaussian predictions, but there are some deviations, particularly for higher k_{\max} and small n . These observations can be explained with the reasoning that, since clustering is more non-linear on smaller scales, the case with higher k_{\max} will have larger $\bar{\beta}_n^R$ and more non-Gaussian $\Delta\beta_n^R$. The fact that these effects are concentrated at the low n modes can be explained by how the orthonormal basis of R_n functions have been defined. Since each R_n is by definition a linear combination of Q_0, \dots, Q_n that is orthogonal to (i.e. in the Gaussian limit, has no covariance with) all previous R_m for $m \leq n$, we generally expect higher R_n modes to have smaller amplitudes in the data, making the Gaussian error approximation more accurate.

In Fig. 3, we compare the means and errors of the modal expansion coefficients from the 298 matched Pinocchio realizations to the N -body measurements. The errors measured from Pinocchio are always within 10% of the errors from Minerva. For the first few modes, we note that the mean β^R measured from Pinocchio can differ from the N -body result by more than $\sim 0.1 \sigma$, but this is not unexpected. We recall that the halos in the mock catalogs were selected such that the total power spectrum, or the bispectrum variance, matched that of Minerva. Thus, by construction, the bispectrum variance from the mocks agrees with that




Figure 2: Means $\bar{\beta}_n^R$ and errors $\Delta\beta_n^R$ of modal expansion coefficients measured from 298 Minerva N -body simulations. The two panels show the measurements for $k_{\max} = 13.5 k_f$ (top, blue circles) and $24.5 k_f$ (bottom, red circles). Filled (empty) circles correspond to positive (negative) values. The gray dashed lines correspond to the Gaussian predictions for the error given in eq. (2.60).

of the N -body simulations, while the mean bispectrum from the mocks is suppressed relative to the N -body mean [68]. This overall suppression is what causes the β_0^R from Pinocchio to be less than what is measured in Minerva.

In Fig. 4, we compare the correlation matrices from Minerva vs the 10,000 Pinocchio mocks. As in the bispectrum case [68], we find that the correlation coefficients are generally

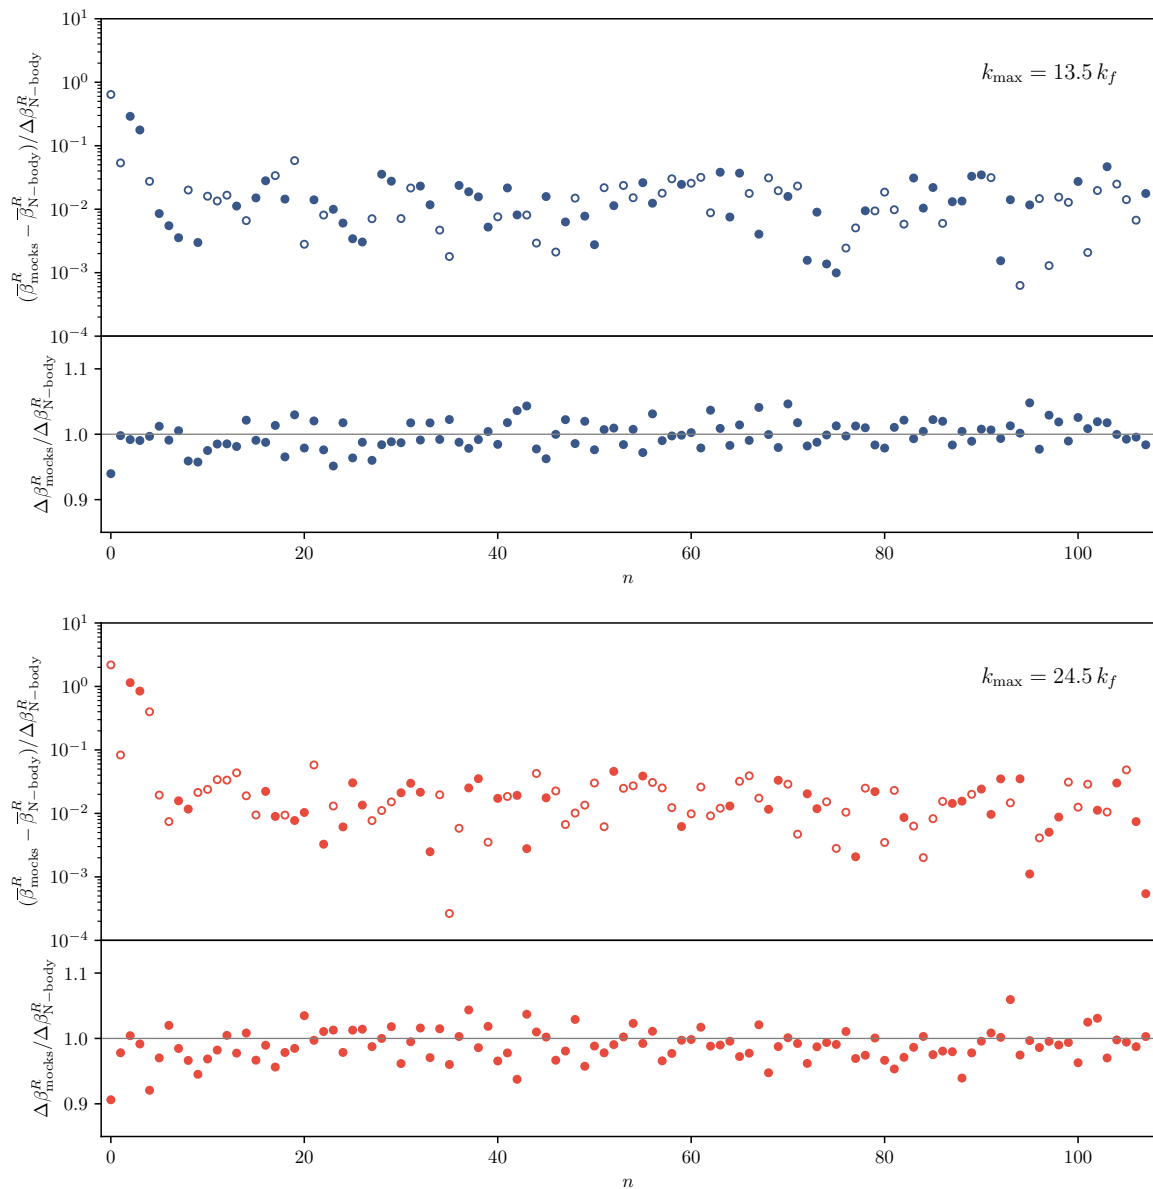


Figure 3: Means and errors of modal expansion coefficients measured from 298 realizations of Pinocchio mock halo catalogs are compared with the same measurements from the Minerva N -body simulations with matching initial conditions. The two panels show the measurements for $k_{\max} = 13.5 k_f$ (top, blue circles) and $24.5 k_f$ (bottom, red circles). Filled (empty) circles correspond to positive (negative) values. The bottom subplots show that the errors from the mocks are always within 10% of the N -body simulations.

closer to zero when 10,000 mocks are used, compared to 298 Minerva simulations. When only the 298 Pinocchio mocks with matched initial conditions are compared, the correlation coefficients agree very well (though we do not show this in a plot). Modal coefficients are more correlated when k_{\max} is higher. We compare how the different estimates of the covariance

matrix (using 298 Minerva simulations, 298 matched Pinocchio mocks, or all 10,000 Pinocchio mocks) impact the parameter constraints in Section 4.6.




Figure 4: Correlation matrices for 108 β^R modal coefficients at $k_{\max} = 13.5 k_f$ (left) and $24.5 k_f$ (right). Within each matrix, the lower triangular elements are from the 298 Minerva simulations while the upper triangular elements are from the full set of 10,000 Pinocchio realizations. Correlation coefficients tend to be closer to zero for lower k_{\max} and when more realizations are used.

4.2 Standard bispectrum vs reconstructed bispectrum

In this section, we look at the relationship between the standard bispectrum estimator in eq. (2.21) and the reconstructed bispectrum in eq. (2.24) at the level of an individual realization, the mean of many simulations, and the resulting covariance. Both estimators capture information about the bispectrum, but they are not equivalent, and in this section we examine which properties of the two estimators are the same or not.

For the comparisons in this section, we include scales up to $k_{\max} = 13.5 k_f$, and adopt triangle bins of width $\Delta k = s k_f$ with $s = 1$ for the standard bispectrum estimator. In this case, the standard bispectrum estimator measured the bispectrum in 294 triangle bins. Within each bin, we compare the standard estimator measurement with the bin-averaged value of B_{rec} in eq. (2.24). We have also done the comparison using B_{rec} evaluated on effective triangles with sorted side lengths, which was shown in [68] to provide a good approximation to the true bin average, and the results that follow are not changed.

Fig. 5 compares the standard bispectrum estimate B with B_{rec} for a single realization, showing that the two do not measure the same value in each bin. This is not unexpected, as the two estimators perform rather different operations on the same density grid in Fourier space in order to produce these bispectrum estimates.

Then, in Fig. 6, we extend the comparison to the suite of all 298 Minerva simulations. In this case, the mean B and B_{rec} are in much better agreement, on a bin-by-bin basis, where

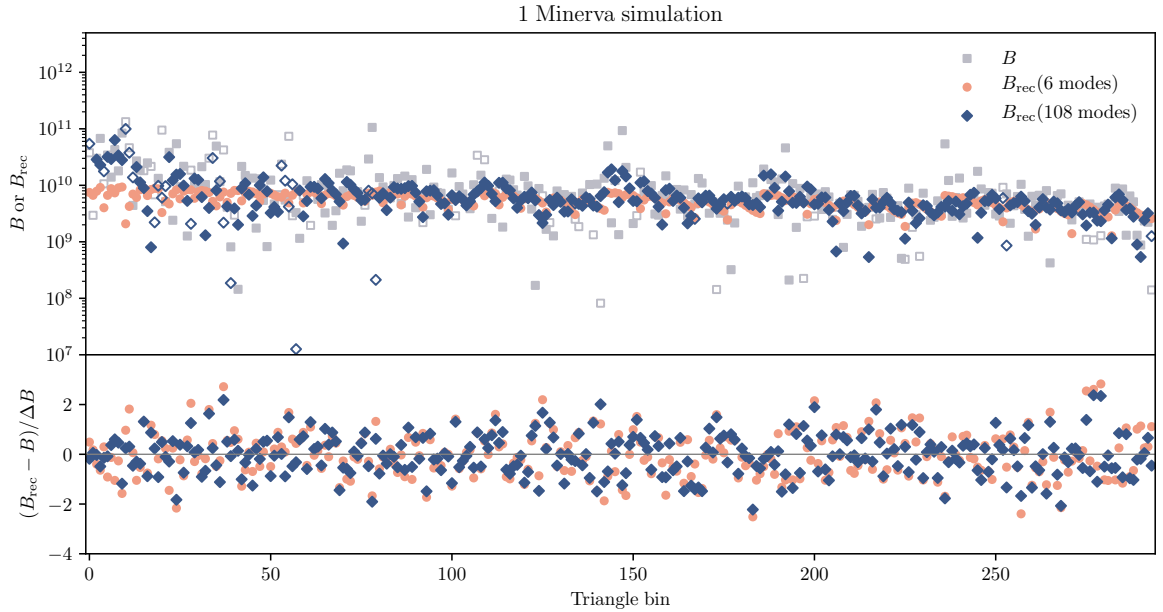


Figure 5: Comparison of the standard bispectrum vs reconstructed bispectrum on one Minerva simulation for $k_{\max} = 13.5 k_f$. In the top panel, filled (empty) markers signify positive (negative) values. The standard bispectrum measurement and the reconstructed bispectra do not obtain the same values in each bin, though the scatter between them is typically within two times the 1σ error of the standard estimator, ΔB .

the differences are typically $\lesssim 10\%$ of the error on the standard bispectrum estimate (middle panel). The difference is that the standard bispectrum estimates have more scatter, while the reconstructed bispectrum with fewer modes has less scatter. In the bottom panel of Fig. 6, we show one key difference between the standard bispectrum and modal bispectrum estimates: the error on the reconstructed bispectrum is typically suppressed relative to the error on the standard bispectrum, and it depends on the number of modes used in the reconstruction. This is shown also for the correlation matrices in Fig. 7, where we see that the triangle bins are much more correlated (and anti-correlated) when the reconstructed bispectrum is used, especially with fewer modes.

The underlying reason for this discrepancy is that the full N_{tri} -dimensional Gaussian distribution that is captured by the standard bispectrum covariance cannot be compressed into a N_{modes} -dimensional one, unless $N_{\text{modes}} = N_{\text{tri}}$, in which case there is no practical compression of the data. This can be illustrated by calculating the covariance of the reconstructed bispectrum as

$$\text{Cov}[B_{\text{rec}}(\Delta_i), B_{\text{rec}}(\Delta_j)] = \frac{1}{w(\Delta_i)w(\Delta_j)} \sum_n \sum_m \text{Cov}[\beta_n^R, \beta_m^R] R_n(\Delta_i) R_m(\Delta_j). \quad (4.1)$$

In the limit of Gaussian covariance for $\Delta_i = \Delta_j$,

$$\text{Var}[B_{\text{rec}}(\Delta_i)] = \frac{1}{w(\Delta_i)^2} \sum_n \text{Var}[\beta_n^R] R_n(\Delta_i)^2, \quad (4.2)$$

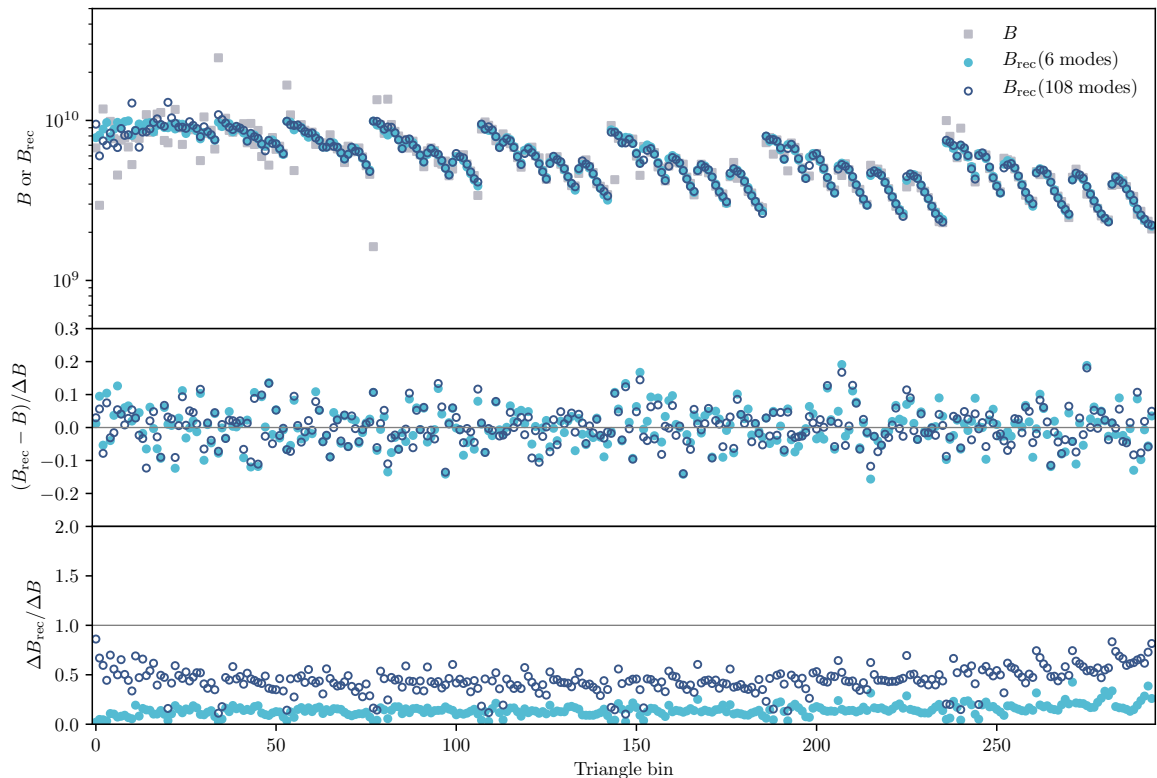


Figure 6: Comparison of the standard bispectrum vs reconstructed bispectrum means and errors over all 298 Minerva simulations for $k_{\max} = 13.5 k_f$. While the means agree to within $\sim 10\%$ of the error of the standard bispectrum estimator (middle panel), the error in the reconstructed bispectrum is generally suppressed relative to the error in the standard bispectrum measurement (bottom panel). This suppression is larger when fewer modes are used in the reconstruction.

such that adding more modes to the modal expansion will always increase the variance on B_{rec} , and conversely, using fewer modes will suppress the variance on B_{rec} (as seen in the bottom panel of Fig. 6).

The modal method produces errors on B_{rec} that are suppressed relative to errors from the standard bispectrum estimator. This implies that measurements of the modal coefficients cannot give estimates of bispectrum errors that are directly relevant for standard bispectrum pipelines. For example, errors on B_{rec} should not be used to judge how well theoretical model predictions for the bispectrum would work in a pipeline using the standard bispectrum estimator.

4.3 Benchmark comparison: modal bispectrum vs standard bispectrum

Here we compare parameter constraints from the modal bispectrum and the standard bispectrum for the default modal pipeline settings detailed at the beginning of Section 4. In subsequent sections, we demonstrate how the modal bispectrum constraints can depend on these settings.




Figure 7: Comparison of correlation matrices from the standard bispectrum and reconstructed bispectra for 294 triangle bins and $k_{\max} = 13.5 k_f$. The noticeable differences in the correlation matrices for B and B_{rec} indicate that the covariance matrix for the standard bispectrum estimator cannot be recovered by the modal bispectrum. This discrepancy is especially pronounced when fewer modes are used in the reconstruction.

To put the bispectrum constraints in context, we also compare them with an independent constraint of $b_1 = 2.7081 \pm 0.0012$ from using chi-squared minimization to fit the ratio of cross-power spectra P_{hm}/P_{mm} from the Minerva simulations [84]. The fitting function is $P_{hm}(k)/P_{mm}(k) = b_1 + \text{coefficient} \times k^2$ up to $k_{\max} = 0.044 h \text{ Mpc}^{-1}$, where the k^2 term is used to account for scale-dependent loop corrections, which for the small error bars corresponding to the total volume of the Minerva simulations can be important even at scales as large as $k \approx 0.02 h \text{ Mpc}^{-1}$.

The benchmark constraints are shown in Fig. 8 for $k_{\max} = 13.5 k_f \approx 0.06 h \text{ Mpc}^{-1}$. This k_{\max} was chosen such that we conservatively consider scales over which the tree-level halo bispectrum model has been shown to accurately describe the data, and at this k_{\max} we are able to compare with standard bispectrum results from all three bin widths in [68], $s = 1, 2,$ and 3 .

The modal constraints using six modes and 59 modes have the same posteriors, with the contours overlapping to the point of making the 59 modes case almost invisible, indicating that the parameter constraints have fully converged with the six custom modes. This is fully consistent with the analysis in [68] which showed that the tree-level bispectrum model is a good fit to the data up to this k_{\max} .

We find that the modal bispectrum constraints are consistent with, though not identical to, the standard bispectrum constraints. We note that the modal bispectrum estimator, while it is a measure of the bispectrum, is an independent estimation of it, so we do not require (in the sense of a test) that the constraints be identical. Although all bin choices for the standard bispectrum estimator lead to consistent outcomes, the smallest bin width ($s = 1$) takes better advantage of the shape-dependence of the bispectrum leading to slightly narrower constraints. In comparison, the modal decomposition accounts for the shape-dependence of the bispectrum without loss of information due to the binning of wavenumbers.

In Fig. 9 we perform the same comparison except with the k -range extended to $k_{\max} = 24.5 k_f \approx 0.10 h \text{ Mpc}^{-1}$. In this case, by comparing the modal bispectrum constraints with




Figure 8: Benchmark comparison between modal bispectrum and standard bispectrum constraints for $k_{\max} = 13.5 k_f \approx 0.06 h \text{Mpc}^{-1}$. The vertical dashed black line represents the constraint on b_1 from the large-scale ratio of cross-power spectra P_{hm}/P_{mm} . Standard bispectrum constraints with three different bin sizes are shown: $s = 1$ (red), 2 (green), and 3 (gray). The modal bispectrum constraints with six modes (dark blue) and 59 modes (light blue) are identical, indicating that the modal bispectrum constraints have already converged with six modes.

different numbers of modes, we see that the six custom modes are no longer sufficient, but 10 modes has already converged, as there is no further benefit to using 59 modes. The fact that there is further information in four additional modes, on top of the six custom modes, is a sign that the halo bispectrum deviates from the tree-level prediction on scales within this

k -range, as studied in much more detail in [68]. Although we know that the tree-level halo bispectrum model is no longer a good model for the data to these scales, we can still compare the constraints with those from the standard bispectrum. At this k_{\max} , we can only compare the standard bispectrum constraints with the $s = 1$ binning. Fig. 9 shows that the results from the modal bispectrum and standard bispectrum agree very well, with only a small bias relative to each other (most visible in the 1-dimensional posterior for b_2 which is biased by $\sim 0.5\sigma$).

Finally, we comment on the amount of compression that has been achieved in these benchmark comparisons. At $k_{\max} = 13.5k_f$, the number of triangle bins that were used by the standard bispectrum estimator are 294 bins for $s = 1$, 49 bins for $s = 2$, and 19 bins for $s = 3$, and we found that the modal bispectrum constraints had converged using only six custom modes. At $k_{\max} = 24.5k_f$, the standard bispectrum used 1,585 bins with $s = 1$, yielding constraints that were very similar to the modal bispectrum using 10 modes. This shows that the modal bispectrum is able to efficiently compress the information contained in the bispectrum into a data set that is 3 to 160 times smaller, while preserving most of the important cosmological information that we are interested in.

4.4 Robustness checks in the modal implementation

In this section, we vary the settings in the modal analysis pipeline away from the benchmark settings to see how the modal constraints are sensitive to these choices.

Normal vs Legendre polynomials

The six custom modes are defined to fit the tree-level bispectrum, and so their form is independent of whether we choose normal or (shifted) Legendre polynomials as our $q_n(k)$. Therefore, we take $k_{\max} = 24.5k_f$ and 10 modes, and compare the constraints between choosing normal polynomials vs Legendre polynomials. The constraints from the Legendre polynomials are identical to the modal constraints with 10 modes and normal polynomials shown in Fig. 9, and plotted together only one set of posteriors would be visible, so to save space we do not show this comparison in a figure.

Custom modes

What is the impact of including the custom modes? To see this, we perform the same analysis as in the benchmark case for $k_{\max} = 13.5k_f$ and $24.5k_f$, but do not include the six custom modes, and see how the modal expansion convergence is affected by not including custom modes. The case for $k_{\max} = 13.5k_f$ is in Fig. 10, and for $k_{\max} = 24.5k_f$ is in Fig. 11. For lower k_{\max} , we find that 16 modes are sufficient for the parameter errors to agree to within 8% with the result from six custom modes, and the parameter means are shifted by $\lesssim 0.25\sigma$ compared to the six custom modes case. For higher k_{\max} , the expansion converges with 31 modes, compared to only 10 modes when six of these are the custom modes. (It appears to converge also with 16 modes, but we find that this is not stable, because the 16 modes case changes when compared with 23 modes.) The means are shifted by $\lesssim 0.3\sigma$ and the errors are consistent to within 1% compared to the benchmark case with custom modes. This shows that the inclusion of custom modes which are constructed based on an informative theoretical model for the bispectrum can help to compress the data into the most efficient basis.




Figure 9: Comparison of modal bispectrum and standard bispectrum constraints at $k_{\max} = 24.5 k_f \approx 0.10 h \text{Mpc}^{-1}$. At this k_{\max} , we can only compare with the $s = 1$ binning of the standard bispectrum estimator (red). By comparing the constraints using six custom modes (dark blue), 10 modes (gray), and 59 modes (light blue), we find that 10 modes are sufficient for the modal constraints to converge, as there is no further change when 59 modes are used.

Inner product methods

The benchmark results used a γ matrix computed with the 3D FFT method using $N_g = 256$ and $L = 1500 h^{-1} \text{Mpc}$. In this section, we compare constraints that have used three different numerical methods for computing γ : 3D FFT, voxels, and 1D FFT. (As mentioned in Section 2.2, we do not include in this comparison the γ computed using the Vegas routine in Cuba, as this results in matrices that are not positive-definite and therefore cannot be used.) All three




Figure 10: Comparison of modal bispectrum constraints at $k_{\max} = 13.5 k_f \approx 0.06 h \text{ Mpc}^{-1}$ without and with custom modes. We compare constraints without custom modes using 7 modes (gray), 11 modes (green), and 16 modes (red) with the converged benchmark constraints that used 6 custom modes (dark blue).

methods are unique and have their own free parameters, which roughly correspond to the resolution of the inner product integration in k -space. By comparing constraints obtained with γ matrices computed from each method, we find that both the method and how its free parameters are set can affect the resulting constraints.

Figs. 12 and 13 compare the methods for $k_{\max} = 13.5 k_f$ and $24.5 k_f$, respectively. The N_v that is shown for the voxel method is the number of voxel cells in each dimension, and the N and M resolution parameters for the 1D FFT method (defined in Appendix B), are

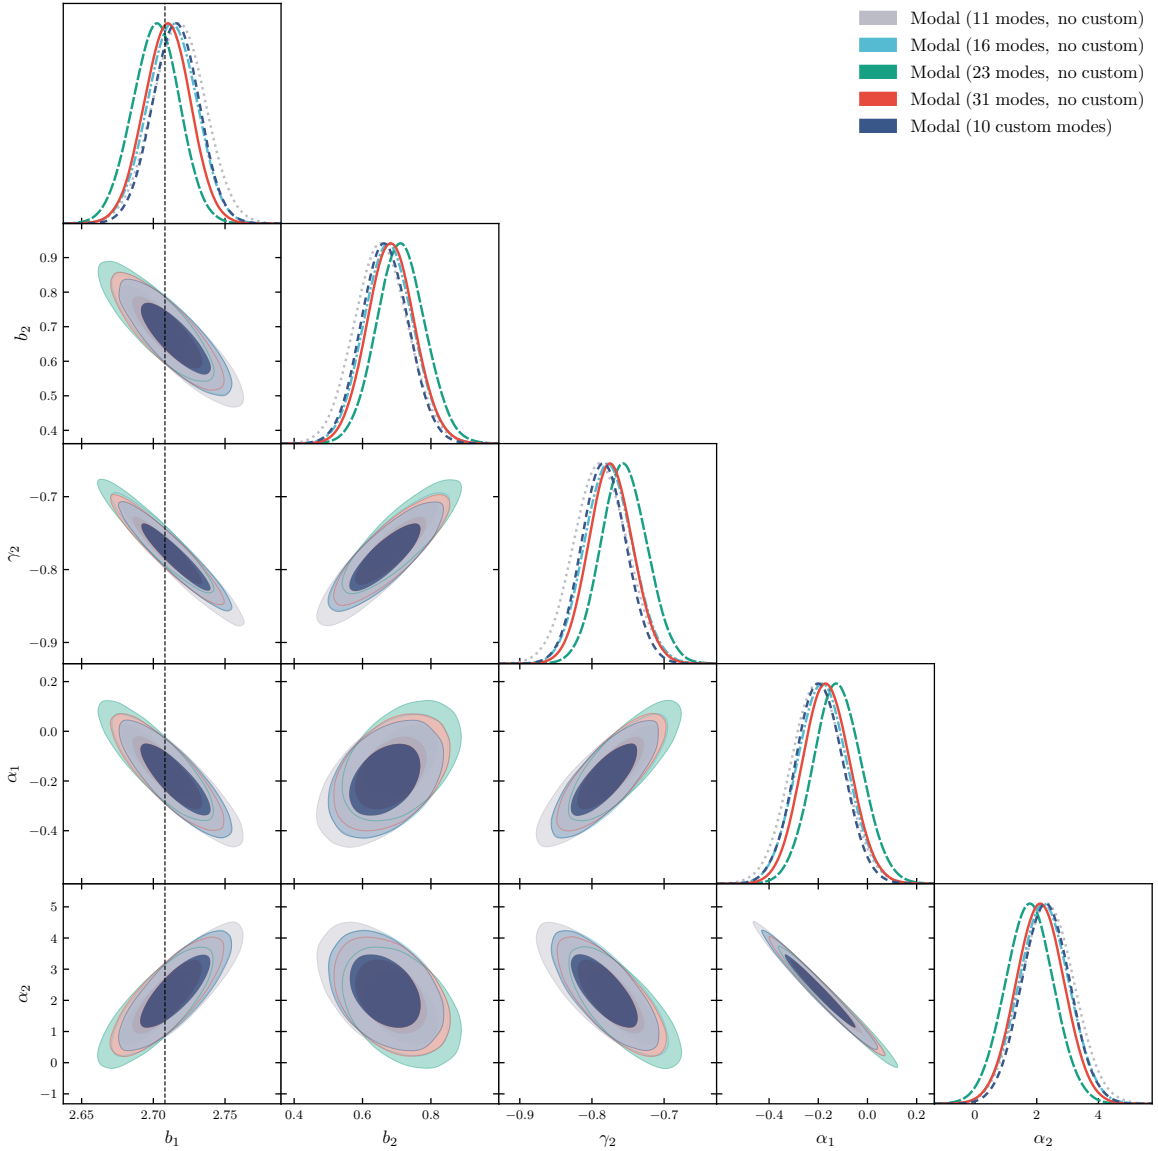


Figure 11: Comparison of modal bispectrum constraints at $k_{\max} = 24.5 k_f \approx 0.10 h \text{ Mpc}^{-1}$ without and with custom modes. We compare constraints without custom modes using 11 modes (gray), 16 modes (light blue), 23 modes (green), and 31 modes (red) with the converged benchmark constraints that used 10 custom modes (dark blue).

converged; doubling these values did not show any changes in the resulting constraints. When k_{\max} is small, we find that the voxel and 1D FFT methods converge on the same posterior for sufficiently high resolutions, but they strongly disagree with the 3D FFT result in both the position and size of the posterior. For higher k_{\max} , the different methods produce posterior contours that agree in their size, but with non-negligible shifts (biases) between them. The observation that the voxel and 1D FFT methods converge to each other for both k_{\max} , while disagreeing more strongly with the 3D FFT calculation at lower k_{\max} , reflects the fact that

the inner product calculation captured by γ must be treated using the same discretization scheme as the measurements in order to obtain correct constraints that are consistent with the standard bispectrum analysis. The voxel and 1D FFT methods are ways of calculating the inner product assuming it is a smooth continuous integral, and it is not straightforward to adapt these methods such that they account for the same discretization effects as the measurements, while the 3D FFT method is, by construction, computing the inner product in the same way that the measurements are taken.

When k_{\max} is low, the discretization of the Fourier grid during the estimation is more important to take into account, because fewer triangle configurations are being averaged (i.e. the tetrapyd is more sparsely sampled). This interpretation also be confirmed another way: within the 3D FFT method, the resolution is increased if the size of the FFT box, L , is larger. If we increase the resolution by increasing L to be twice as large as the simulation box, we find in both Figs. 12 and 13 that the resulting posterior becomes more similar to the voxel and 1D FFT case, as we would expect.

It may be the case that for a higher k_{\max} than what we have used in this work, the different methods for computing γ may yield results that are similar enough that the different methods can be interchangeable. This is expected because more cosmological information is contained in non-linear scales and the continuous integration of the voxel and 1D FFT methods becomes a better approximation to the discretized inner product when k_{\max} is higher. However, we emphasize that for the range of scales that we have used in this work, $k_{\max} \lesssim 0.10 h \text{Mpc}^{-1}$, the different methods for computing γ are *not* interchangeable, and the 3D FFT method is the only one which treats the inner product identically to how the measurements are performed, resulting in correct constraints.

The 3D FFT method is also preferable for the speed and ease of its calculation. The 3D FFTs are performed very quickly using the same FFT routines which are already necessary for the modal bispectrum (and standard bispectrum) measurements, while the other voxel and 1D FFT methods require different algorithms which must be coded independently and, in our implementation, are not as fast. As γ must only be computed once for a fixed k_{\max} , we do not anticipate that the computation of γ which is unique to the modal bispectrum analysis increases the computational cost of using the modal method by a significant amount.

Dependence on FFT grid resolution

Depending on k_{\max} , the FFT grid on which the 3D FFTs are calculated can have a configuration-space grid resolution much smaller than our default value of $N_g = 256$. Since L fixes the resolution of the Fourier grid to be $k_f \equiv 2\pi/L$, increasing N_g increases the k_{\max} that can be probed without too much aliasing contamination. [85] and [77] have suggested that the standard FFT bispectrum estimator, which takes a form very similar to the modal estimator, can probe up to $k_{\max} = k_f N_g/3 = 2k_{\text{Ny}}/3$, where k_{Ny} is the Nyquist wavenumber, unlike the power spectrum estimator which is valid up to $k_{\max} = k_f N_g/2 = k_{\text{Ny}}$. This is because the factor of $e^{i\mathbf{k}_{123}\cdot\mathbf{x}}$ in the estimator is invariant under shifts (in one dimension) of each k_i to $k_i + k_f N_g/3$ [77]. On the other hand, the opposite has been argued by [86] for the FFT bispectrum estimator and [65] for the modal estimator—that these estimators are valid up to $k_{\max} = k_{\text{Ny}}$.

Here we fix our k -range to have $k_{\max} = 13.5 k_f$ and test which criterion for N_g , either $N_g > 3k_{\max}/k_f \approx 41$ or $N_g > 2k_{\max}/k_f = 27$ is sufficient to return the same constraints from the modal estimator pipeline as the benchmark value of $N_g = 256$. We note that changing N_g requires the pipeline to be run from the beginning, starting with the construction of




Figure 12: Comparison of modal bispectrum constraints at $k_{\max} = 13.5 k_f \approx 0.06 h \text{ Mpc}^{-1}$ using different inner product methods. The voxel (green) and 1D FFT (red) methods have converged towards each other, and both disagree with the 3D FFT result (gray) that agreed with the standard bispectrum constraint in Fig. 8. The 3D FFT case with $L = 3000 h^{-1} \text{ Mpc}$ (dark blue) is also biased, illustrating that getting the discretization as in the measurements is important.

the configuration-space density grid, and including the calculation of γ with the 3D FFT method. Fig. 14 compares the constraints for different values of $N_g = 34, 42,$ and 256 . We find that using $N_g = 34$ leads to constraints that strongly disagree with the benchmark case of $N_g = 256$, while $N_g = 42$ gives identical constraints to the benchmark case, showing that the modal estimator is valid only up to $2k_{\text{Ny}}/3$. If this result is explained by the argument in




Figure 13: Comparison of modal bispectrum constraints at $k_{\max} = 24.5 k_f \approx 0.10 h \text{ Mpc}^{-1}$ using different inner product methods. As in Fig. 12, the voxel (green) and 1D FFT (red) methods have converged towards each other, and both are biased relative to the correct 3D FFT result (gray). The 3D FFT case with $L = 3000 h^{-1} \text{ Mpc}$ (dark blue) is also biased, illustrating that getting the discretization as in the measurements is important.

[77], then we would expect this result to also hold for the FFT-based standard bispectrum estimator, which also has a factor of $e^{i\mathbf{k}_{123}\cdot\mathbf{x}}$.

Weighting

In the benchmark analysis, the bispectrum was weighted by w in eq. (2.6), where the power spectrum $P(k_i)$ was the average total halo power spectrum measured from the Minerva




Figure 14: Comparison of modal bispectrum constraints at $k_{\max} = 13.5 k_f \approx 0.06 h \text{ Mpc}^{-1}$ using different FFT grid resolutions, N_g . The case with $N_g = 42$ (dark blue) yields identical results to the benchmark case with $N_g = 256$ (gray, hidden underneath the dark blue contours), because both satisfy $N_g > 3k_{\max}/k_f \approx 41$, while $N_g = 34$ (red) is insufficient and leads to incorrect constraints.

simulations. What is the effect of using a different weighting function?

We note that changing the weighting only changes two parts of the modal pipeline. First, the q_n^{tree} functions in eqs. (2.42)–(2.47) will change, such that the factor of $\sqrt{k/P(k)}$ in each one will be different. This will, however, not change the fact that the six custom modes, Q_n^{tree} , are able to reconstruct the tree-level halo bispectrum model exactly. The second change is that when $\langle\langle Q|w\hat{\mathcal{B}}\rangle\rangle$ in eq. (2.17) is estimated from simulations, the factor

of $[\sqrt{k_1 k_2 k_3} \sqrt{P(k_1)P(k_2)P(k_3)}]^{-1/2}$ in the integrand will change to $w/[k_1 k_2 k_3]$.

We have considered the case where w takes the same form as in eq. (2.6), but $P(k_i)$ is set to the linear *matter* power spectrum P_L . This power spectrum is different to the benchmark weighting in that the power spectrum does not have any halo bias, non-linearities, or shot noise. Therefore, we use this situation to reflect an analysis where the halo power spectrum in the weight is not perfectly modeled or measured. We find that this difference does not have any effect on the resulting parameter constraints, implying that the modal bispectrum constraints are not strongly affected by the particular power spectrum that is used for the weighting. In particular, it does not change how quickly the modal expansion converges.

Still, there is a reason to prefer the optimal weighting with the non-linear total halo power spectrum, which is that it is in this case that the covariance of the β^R is best approximated by the Gaussian covariance expression in eq. (2.60). If the linear power spectrum is used in the weighting, the $P(k_1)P(k_2)P(k_3)$ that appears in eq. (2.58) does not cancel out with the power spectra in the weight, such that the Gaussian covariance expression for β^R is not eq. (2.60).

This result does not necessarily mean that the parameter constraints are totally immune to especially sub-optimal choices for w . We have checked that when $w = 1$ is adopted, in other words, no weighting at all is used, the information in the bispectrum is less efficiently extracted. This is shown in Fig. 15, which compares the constraints in the ‘no weight’ case with the benchmark results for $k_{\max} = 13.5 k_f$. With no weighting, we find that the constraints using six custom modes is much weaker than, though still consistent with, the benchmark case of six custom modes with default weighting. This difference must originate from the choice of weighting, because in both cases the six custom modes, by construction, can exactly reproduce the tree-level bispectrum model that is a good description of the bispectrum up to this k_{\max} . The fact that the no weighting constraints are weaker is consistent with the fact that less optimal weighting of the (k_1, k_2, k_3) Fourier triangles should lead to less information being extracted. However, even in this case, the modal expansion method can compensate for the less-than-optimal weighting if a larger number of modes are included. For this k_{\max} , Fig. 15 shows that the benchmark constraints can be recovered if 21 modes are included.¹⁵

4.5 Modal expansion correlators

One way of measuring the accuracy of the modal expansion is to define and compute so-called *correlators* that quantify different aspects of the accuracy of the reconstruction. For example, the shape correlator \mathcal{S} , amplitude correlator \mathcal{A} , and total correlator \mathcal{T} are [63, 65]

$$\mathcal{S}(B_i, B_j) \equiv \frac{[B_i, B_j]}{\sqrt{[B_i, B_i][B_j, B_j]}} \quad (4.3)$$

$$\mathcal{A}(B_i, B_j) \equiv \frac{\sqrt{[B_i, B_i]}}{\sqrt{[B_j, B_j]}} \quad (4.4)$$

$$\mathcal{T}(B_i, B_j) \equiv 1 - \sqrt{1 - 2\mathcal{S}(B_i, B_j)\mathcal{A}(B_i, B_j) + \mathcal{A}(B_i, B_j)^2}, \quad (4.5)$$

¹⁵If $w = 1$, we note that both the last custom mode Q_5^{tree} and Q_0 will be constants that do not have any dependence on (k_1, k_2, k_3) , so a modal basis that includes both of them will correspond to having a non-positive definite γ matrix. For this reason, when $w = 1$ and custom modes are included, the basis sets with more than six modes leave out the Q_0 mode, but otherwise have the same ordering of Q_n functions as in the rest of this work.




Figure 15: Comparison of modal bispectrum constraints at $k_{\max} = 13.5 k_f \approx 0.06 h \text{Mpc}^{-1}$ using the benchmark weighting function in eq. (2.6) (green) vs no weight, where $w = 1$. When the six custom modes are used in both cases, $w = 1$ (red) leads to weaker parameter constraints, but this sub-optimal choice of weighting can be fully compensated for by using more modes. For this k_{\max} , 12 modes (dark blue) is not sufficient, but 21 modes (gray, overlapping exactly with the green contours) are able to recover the same constraints as the more optimal weighting.

where the square brackets notation above from [65] is

$$[B_i, B_j] \propto \sum_n \beta_n^{R(i)} \beta_n^{R(j)}, \quad (4.6)$$

such that $[B_i, B_j]$ is proportional to our $\langle\langle wB_i | wB_j \rangle\rangle$. The shape correlator \mathcal{S} takes values between -1 and 1 and is insensitive to constant multiplicative factors that change the bispectrum amplitude, while the amplitude correlator \mathcal{A} can take any positive value. The total correlator \mathcal{T} is sensitive to both the shape and the amplitude of the bispectrum, such that if both the shape and amplitude are perfectly matched, then $\mathcal{S} = \mathcal{A} = \mathcal{T} = 1$, and if either the shape or the amplitude are not perfectly reproduced then $\mathcal{T} < 1$. These correlators have an intuitive quantitative meaning in cases where an amplitude parameter (like f_{NL} , the amplitude of primordial non-Gaussianity, for example) is measured with Gaussian data covariances [65].

We show plots of \mathcal{S} and \mathcal{T} between the reconstructed bispectrum with n modes, $B_i = B_{\text{rec}}(N_{\text{modes}} = n)$, and the reconstructed bispectrum with 108 modes, $B_j = B_{\text{rec}}(N_{\text{modes}} = 108)$, in Fig. 16. The figure is for $k_{\text{max}} = 24.5 k_f \approx 0.10 h \text{Mpc}^{-1}$ and considers the mean modal bispectrum from 298 Minerva simulations. The vertical gray lines at $N_{\text{modes}} = 10$ and 31 mark the number of modes that we previously found in Sections 4.3 and 4.4 were sufficient for converged parameter constraints, with and without custom modes, respectively. Both panels show that the correlators are already $\mathcal{S}, \mathcal{T} \gtrsim 0.99$ with six custom modes, and further improvements are gained slowly as more modes are accumulated. The correlators approach unity more slowly in the absence of custom modes, but these too show small improvements after ~ 25 modes. This behavior implies that the \mathcal{S} and \mathcal{T} correlators are not good indicators for predicting how many modes would be sufficient to use the modal bispectrum to constrain cosmological parameters of interest. Given only the information in Fig. 16, it is not obvious how many modes will be needed for any specific purpose. This is partly because the correlators do not take into account other information that will influence the number of sufficient modes, such as which parameters the modal pipeline will be used to measure.

In this work, we determined a sufficient number of modes by checking that parameter posteriors had converged. However, this requires many steps, including measuring modal coefficients from a large number of simulations and running MCMC simulations multiple times for different N_{modes} to validate our results. In the absence of these, one may consider calculating Fisher forecasts to estimate the number of modes that would be needed, to check that the modal bispectrum still provides a good compression of the information in the bispectrum.

Commonly in Fisher matrix analyses, the fiducial parameter values are kept fixed, and parameter errors are forecasted. However, in this work we found that even though very few modes are needed to reproduce the size and degeneracy directions of the parameter contours, more modes are typically needed to reduce the bias in the positions of the contours in parameter space. (This implies that the modes are better at capturing the derivatives of the bispectrum with respect to our chosen parameters, than it is at capturing the mean bispectrum.) Therefore, it is also prudent to estimate the bias using the Fisher formalism [87, 88].

The Fisher matrix corresponding to the modal pipeline is

$$F_{ij} = \sum_{n,m}^{N_{\text{modes}}-1} \frac{\partial \beta_n^R}{\partial \theta_i} \hat{C}(N_{\text{modes}})_{nm}^{-1} \frac{\partial \beta_m^R}{\partial \theta_j}, \quad (4.7)$$

where the partial derivatives are evaluated for our tree-level model at a chosen fiducial. The parameter covariance matrix is then F^{-1} . The bias in the parameters due to an unaccounted




Figure 16: Shape correlator \mathcal{S} (top) and total correlator \mathcal{T} (bottom) for $k_{\max} = 24.5 k_f$. Smaller values of $1 - \mathcal{S}$ and $1 - \mathcal{T}$ correspond to smaller differences in the reconstructed bispectrum compared to the case where the maximum number of modes are used. The vertical gray lines mark $N_{\text{modes}} = 10$ and 31 , which we previously found was sufficient to obtain robust parameter constraints when custom modes are included or excluded, respectively.

for systematic error can be estimated as (e.g. [89])

$$b(\theta_i) = (F^{-1})_{ij} b_j \quad (4.8)$$

$$b_j = \sum_{n,m}^{\max N_{\text{modes}}-1} \beta_n^{R,\text{sys}} \hat{C}(\max N_{\text{modes}})^{-1}_{nm} \frac{\partial \beta_m^R}{\partial \theta_j} \quad (4.9)$$

where $\beta_n^{R,\text{sys}}$ is a source of residual systematic uncertainty. In our case, to calculate the bias due to truncating at a certain number of modes,

$$\beta_n^{R,\text{true}} = \beta_n^{R,\text{obs}} \quad \text{for } n < \max N_{\text{modes}} \quad (4.10)$$

$$\beta_n^{R,\text{truncated}} = \begin{cases} \beta_n^{R,\text{obs}} & \text{for } n < N_{\text{modes}} \\ 0 & \text{otherwise} \end{cases} \quad (4.11)$$

$$\begin{aligned} \beta_n^{R,\text{sys}} &= \beta_n^{R,\text{true}} - \beta_n^{R,\text{truncated}} \\ &= \begin{cases} 0 & \text{for } n < N_{\text{modes}} \\ \beta_n^{R,\text{obs}} & \text{for } n \geq N_{\text{modes}}, \end{cases} \end{aligned} \quad (4.12)$$

where $\beta_n^{R,\text{obs}}$ is the average measured from the Minerva simulations. We set $\max N_{\text{modes}} = 108$ if custom modes are included, and $\max N_{\text{modes}} = 102$ if they are not. Taking the case where $k_{\max} = 24.5 k_f$, we show the Fisher forecasted errors and bias in Fig. 17. Explicitly, we show

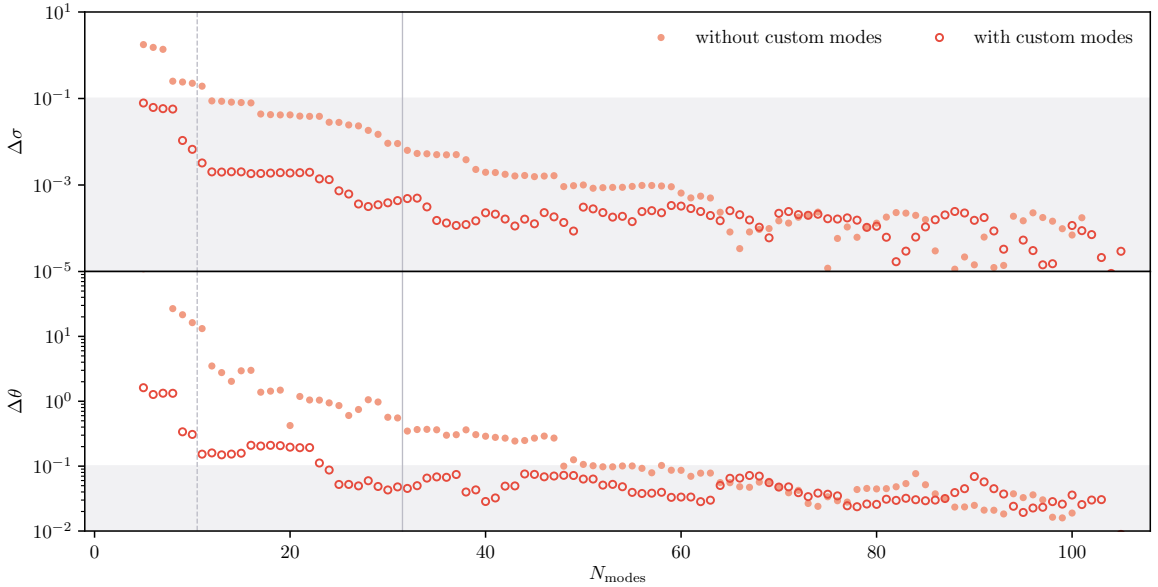


Figure 17: Results from Fisher forecasts on the parameter error (top) and bias (bottom), defined in eqs. (4.14) and (4.13), as a function of total modes used. The vertical gray lines mark $N_{\text{modes}} = 10$ and 31 , which we previously found was sufficient to obtain robust parameter constraints when custom modes are included or excluded, respectively. The shaded gray regions are where the error is within 10%, and the bias is within 0.1σ , of the result with $\max N_{\text{modes}}$.

$|\Delta\sigma|$ and $|\Delta\theta|$, where

$$\Delta\theta \equiv \frac{b(\theta)}{\sigma_\theta(\max N_{\text{modes}})} \quad (4.13)$$

$$\Delta\sigma \equiv \frac{\sigma_\theta(N_{\text{modes}})}{\sigma_\theta(\max N_{\text{modes}})} - 1, \quad (4.14)$$

and σ_θ is the Fisher forecasted error for parameter θ . To keep the plot simple, at each N_{modes} we have plotted the $|\Delta\sigma|$ and $|\Delta\theta|$ that is the largest among the five parameters. This illustrates a simple case where we already have MCMC simulations, and we are simply verifying, after the fact, that Fisher forecasts can provide similar indications of modal expansion convergence.

However, we note that the bias due to a truncation of the modal expansion is a purely non-Gaussian effect that cannot be estimated without the non-Gaussian covariance matrix. This is because the Gaussian covariance is diagonal, such that the b_j in eq. (4.9) would only be sensitive to the systematics in those $\beta_n^{R,\text{sys}}$ modes that also explicitly vary with the θ parameters of the modeling. Still, once a non-Gaussian covariance matrix is obtained, the Fisher formalism may allow for a forecast of how many modes are necessary for the errors and bias to converge to a desired level. This may be useful if a non-Gaussian covariance matrix is available, but one wants to estimate roughly how many modes may be needed without running potentially expensive MCMC simulations for multiple scenarios.

4.6 Covariances

All of the results we have discussed so far used covariance matrices estimated from the full set of 10,000 Pinocchio mocks. In this section, we explore how the modal bispectrum constraints are sensitive to the covariance matrix that is used. Unless otherwise mentioned explicitly, the results in this subsection use $k_{\max} = 24.5 k_f \approx 0.10 h \text{ Mpc}^{-1}$ and 10 modes.

In the limit of Gaussian covariance, the covariance matrix for the β_n^R modal coefficients is diagonal, with the same variance for each n . When $k_{\max} = 13.5 k_f$ with six modes, we find that the Gaussian covariance approximation is very accurate, giving the same constraints as the fully non-Gaussian covariance matrix estimated from 10,000 Pinocchio mocks. However, when $k_{\max} = 24.5 k_f$ with 10 modes, as shown in Fig. 18, the Gaussian covariance underestimates the parameter errors by up to 20%, and the constraints are biased by up to 2σ , depending on the parameter. Fig. 18 also shows that constraints using covariance matrices estimated from 298 Minerva simulations, 298 Pinocchio mocks with matched initial conditions, and the full set of 10,000 Pinocchio mocks are in good agreement: parameter errors agree to within 10% and biases are small, less than $\sim 0.4\sigma$.

In Fig. 19, we show how the parameter means and errors can depend on the number of mocks used to estimate the covariance matrix. From the 10,000 Pinocchio mocks, we take subsets of the mocks divided into groups of N_s and compare the resulting constraints. This is along the lines of [90], which considered the impact of covariance matrix errors on cosmological parameter constraints from the power spectrum. Similarly to eqs. (4.13) and (4.14), we define and show

$$\Delta\theta \equiv \frac{\bar{\theta}(N_s) - \bar{\theta}(N_s = 10^4)}{\sigma_\theta(N_s = 10^4)} \quad (4.15)$$

$$\Delta\sigma \equiv \frac{\sigma_\theta(N_s)}{\sigma_\theta(N_s = 10^4)} - 1. \quad (4.16)$$

Each subset of mocks corresponds to a single gray circle in each panel of Fig. 19, while the red points and error bars show the mean and standard deviation of the gray circles at one value of N_s . This comparison shows that the parameter errors are recovered to within 10% with only 300 mocks, but many more mocks are typically necessary to reduce the bias to the same level. For example, $N_s > 2000$ would be needed to reduce the bias to $\lesssim 0.1\sigma$. One caveat to this result, however, is that the red points in the plot are not independent, since they are dividing up the same realizations, just in different groups. This fact will tend to make the different N_s appear more consistent with the case we are comparing with, using all 10,000 mocks.

4.7 Gaussian vs Sellentin-Heavens likelihood

We compared the two likelihoods when all $N_s = 10^4$ mocks are used to estimate the covariance, and we find that they result in indistinguishable parameter constraints, which is the expected behavior when N_s is very large. For some smaller value of N_s , we expect that the two likelihoods will show different results. For $k_{\max} = 24.5 k_f$ with 10 modes, we show this comparison for two values of N_s , $N_s = 20$ and 300, in Fig. 20. For $N_s = 20$, we simulate 50 analyses, and for $N_s = 300$ we simulate 33. $\Delta\theta$ and $\Delta\sigma$ are defined in eqs. (4.15) and (4.16), and for each MCMC simulation, we plot five points for $\Delta\theta$ and $\Delta\sigma$, one point for each parameter.

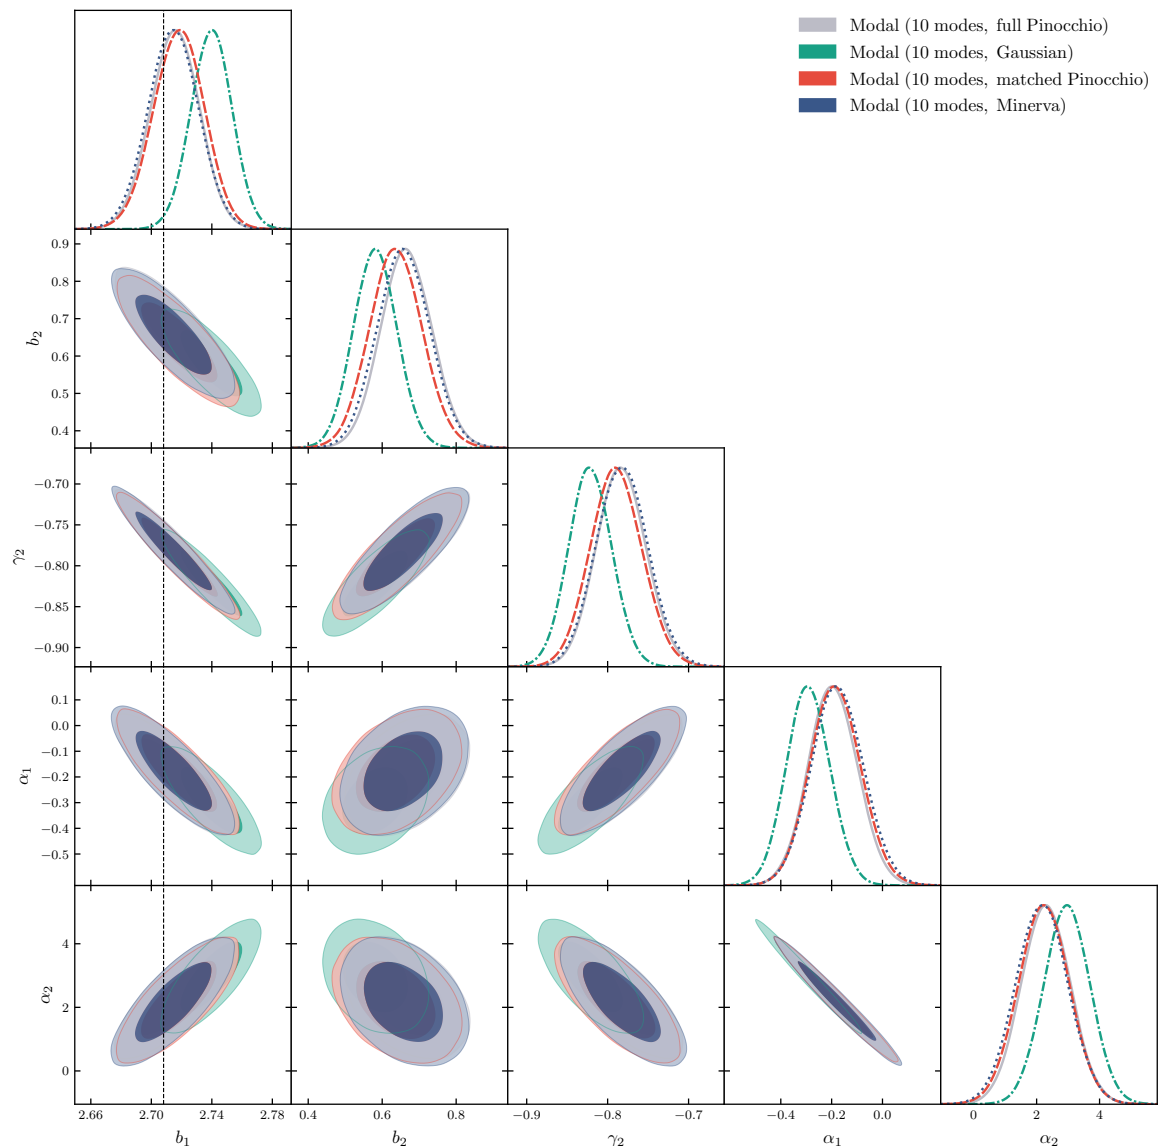


Figure 18: Comparison of parameter constraints when different covariance matrices are used to obtain constraints at $k_{\max} = 24.5 k_f \approx 0.10 h \text{ Mpc}^{-1}$ using 10 modes. The Gaussian covariance matrix (green) leads to biased and underestimated parameter constraints, but the covariance matrices estimated from 298 Minerva simulations (dark blue) and 298 Pinocchio simulations with matched initial conditions (red) are in good agreement with the benchmark modal constraints that used the full set of 10,000 Pinocchio simulations (gray, nearly identical to the dark blue and red contours).

The top row of Fig. 20 shows that $N_s = 300$ is sufficiently high to make differences between the two likelihoods negligible when constraints are compared between individual sets of 300 mocks. When $N_s = 20$ (middle row), which is much closer to the number of data bins, $N_{\text{modes}} = 10$, the two likelihoods can produce posteriors that are noticeably different. For any one analysis using 20 mocks, the two posteriors will be biased relative to each other and can produce parameter errors that are too big or too small relative to the true




Figure 19: Comparison of the bias and change in parameter errors, for $k_{\max} = 24.5 k_f$ with 10 modes, as a function of how many Pinocchio mocks, N_s , are used to estimate the covariance. $\Delta\theta$ and $\Delta\sigma$ are defined in eqs. (4.15) and (4.16). The gray circles mark the result from each set of N_s mocks, while the red error bars are the 1σ scatter of the gray circles. The gray shaded bands mark the regions where deviations are within 0.1σ of the benchmark result using $N_s = 10^4$.




Figure 20: Comparison of the bias (left column) and errors (right column) obtained with the Gaussian and SH likelihoods when covariance matrices are estimated using $N_s = 300$ (top row) or $N_s = 20$ (middle row) mocks. $\Delta\theta$ and $\Delta\sigma$ are defined in eqs. (4.15) and (4.16). The bottom row shows histograms of posterior biases and sizes for the likelihoods when $N_s = 20$. The solid lines are Gaussian fits to the histograms and show that the Sellentin-Heavens likelihood on average is slightly closer to the answer given by the full set of 10^4 mocks.

answer, which we assume is the result from 10^4 mocks. On average though, both likelihoods have posteriors that are unbiased relative to the true answer, but the parameter error from both likelihoods will be larger than the case with $N_s = 10^4$. This is expected, because the parameter constraints should be worse when the covariance is less well estimated from fewer mocks. However, with the SH likelihood, the result on average is closer to the truth: the scatter in the bias $|\Delta\theta|$ is smaller (bottom left panel), the average parameter errors are closer to the truth (bottom right panel), and the distribution of parameter errors are more tightly scattered around the true value (also bottom right panel).

Despite this, we find that in practice it does not matter which likelihood is implemented, because in either case enough mocks would have to be used to ensure that the parameter constraints are not dominated by the covariance matrix error. In this work, we find that once N_s is large enough for either likelihood to be stable to within a few tens of per cent (as shown in Fig. 19), the two likelihoods will produce identical results.

5 Conclusions

In this work, we have implemented an MCMC analysis using the compressed modal bispectrum for the first time. By using the same data, modeling, and analysis choices as [68], we are able to rigorously compare the constraints from the standard bispectrum estimator and the modal bispectrum estimator within a controlled setting. Specifically, we use the real-space tree-level halo bispectrum model to constrain the halo bias and shot noise parameters measured in the Minerva N -body simulations, which represents an idealized survey with volume $\approx 1,000 h^{-3} \text{Gpc}^3$.

Our key result is that the modal bispectrum provides a very efficient compression of the information in the bispectrum, while requiring minimal new calculations compared to the standard bispectrum analysis; the critical components of the pipeline are the modal estimator in eqs. (2.19) and (2.20) and the inner product matrix, γ , and both can be computed with minor modifications to the standard bispectrum estimator. We find that for $k_{\text{max}} = 13.5 k_f \approx 0.06 h \text{Mpc}^{-1}$ ($24.5 k_f \approx 0.10 h \text{Mpc}^{-1}$), the constraints on halo bias and shot noise parameters converge with only 6 (10) modal coefficients, yielding very similar constraints compared to the standard bispectrum analysis in [68] that used ~ 20 to 1,600 triangle bins. We showed that this convergence of the constraints with N_{modes} is only qualitatively reflected by the shape and total correlators (in Section 4.5), but Fisher forecasts can estimate the N_{modes} needed for parameter constraints to converge to a desired level.

We tested the robustness of the modal bispectrum constraints to different user choices within the modal pipeline implementation. We find that the choice between the normal polynomials or shifted Legendre polynomials for constructing the Q_n basis functions has no impact on the results, but using a near-optimal weighting function w to weight Fourier triangles and including some customized basis functions, like Q_n^{tree} , that are more tuned to the parameters being constrained can help minimize the number of modes needed for more efficient compression. We also compared different methods of computing the inner product matrix, $\gamma_{nm} \equiv \langle\langle Q_n | Q_m \rangle\rangle$, a critical piece of the pipeline, and find that only the 3D FFT method is always correct, though other methods appear to be approximately correct for higher k_{max} . This is because the 3D FFT method calculates the inner product between basis functions, $\langle\langle Q_n | Q_m \rangle\rangle$, on the Fourier-space grid in the same way that the modal estimator calculates the inner product between basis functions and the data, $\langle\langle Q | w\mathcal{B} \rangle\rangle$, treating data and theory in the most consistent way possible. The voxel and 1D FFT methods, on the other

hand, take the continuous limit of the inner product, which becomes a good approximation when the Fourier grid is very fine.

We also noted that, while the modal bispectrum and standard bispectrum estimators are both summary statistics of the true bispectrum, they are performing different operations on the density grid in Fourier space, $\delta(\mathbf{k})$. Thus they are not always interchangeable and care should be taken when comparing the two. To illustrate this, we have shown that they agree on the mean bispectrum averaged over many simulations (i.e. the mean standard bispectrum estimator vs the mean reconstructed bispectrum), but for one realization they give different answers for the triangle-dependence of the measured bispectrum. Additionally, the two estimators have different error properties, and the covariance of B_{rec} cannot be used as a substitute for the covariance of measurements made with the standard bispectrum estimator.

The highly efficient compression achieved by the modal bispectrum and the large number of simulations available have allowed us to explore how the modal estimator constraints depend on the the number of simulations used to estimate the covariance, N_s , and whether N -body simulations or Pinocchio approximate mocks are used. Such calculations can usually only be done in a limited way for bispectrum data sets using the standard estimator because of its much larger size. We find that the covariance matrices from 298 N -body simulations, 298 Pinocchio mocks with matched initial conditions, and the full set of 10,000 mocks lead to constraints that are biased by up to $\sim 0.4\sigma$ relative to each other, and to reduce this bias to $\lesssim 0.1\sigma$ would require $N_s > 2,000$ mocks. We also show that the Gaussian and Sellentin-Heavens likelihood functions only show different results when N_s is extremely low. However, since N_s should be large enough such that the error in the covariance matrix estimate is subdominant with either likelihood function (at which point the two likelihoods give identical parameter constraints), in practice either likelihood could be used, though in principle the Sellentin-Heavens likelihood is more theoretically motivated from a Bayesian perspective.

This work has shown that the modal method remains a promising avenue for accessing cosmological information in the bispectrum through a compressed data set, and we have developed a better understanding of how to implement and interpret the estimator and its results. However, the modal bispectrum pipeline presented here would require further work before it could be applied to a realistic galaxy catalog, including redshift-space distortions, and potentially probing a larger range of scales (higher k_{max}), where a theoretical model beyond the tree-level SPT bispectrum would be necessary. A larger k_{max} will most likely require more modes, ideally including more custom modes that would be theoretically motivated. If a theoretical model beyond the separable tree-level SPT one is used, computational methods in the pipeline will need to be adapted so that the inner product $\langle\langle Q|wB^{\text{theory}}\rangle\rangle$ could still be computed quickly. (The analogous problem in the standard bispectrum analysis is handled by evaluating the theoretical bispectrum at an effective triangle in each triangle bin [68].) It is likely that if k_{max} is sufficiently high, the calculation of the inner product could be well-approximated by another method that does not require evaluating the theory on the Fourier-space grid. Additionally, the modal bispectrum method presented here would need to be extended to capture anisotropies coming from redshift-space distortions [67], which are always present in real observations and also act as a source of more cosmological information. We plan to investigate these outstanding issues in future work.

Acknowledgments

We are grateful to Pierluigi Monaco for providing the Pinocchio mock halo catalogs and to Claudio Dalla Vecchia and Ariel Sánchez for providing the Minerva N -body simulations. We thank Dionysios Karagiannis for suggesting the use of the 1D FFT inner product method. We also wish to thank the Institute for Fundamental Physics of the Universe (IFPU) in Trieste, Italy for hosting the workshop of the Euclid Galaxy Clustering Higher-order Statistics Work Package where part of this work was done.

JB is supported by the Sinergia Grant No. 173716 from the Swiss National Science Foundation. ES acknowledges support from PRIN MIUR 2015 Cosmology and Fundamental Physics: illuminating the Dark Universe with Euclid.

The modal bispectrum analysis was performed on the Baobab cluster at the University of Geneva.

References

- [1] H. Gil-Marín, J. Noreña, L. Verde, W. J. Percival, C. Wagner, M. Manera et al., *The power spectrum and bispectrum of SDSS DR11 BOSS galaxies – I. Bias and gravity*, *Mon. Not. Roy. Astron. Soc.* **451** (2015) 539 [[1407.5668](#)].
- [2] H. Gil-Marín, L. Verde, J. Noreña, A. J. Cuesta, L. Samushia, W. J. Percival et al., *The power spectrum and bispectrum of SDSS DR11 BOSS galaxies – II. Cosmological interpretation*, *Mon. Not. Roy. Astron. Soc.* **452** (2015) 1914 [[1408.0027](#)].
- [3] H. Gil-Marín, W. J. Percival, L. Verde, J. R. Brownstein, C.-H. Chuang, F.-S. Kitaura et al., *The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: RSD measurement from the power spectrum and bispectrum of the DR12 BOSS galaxies*, *Mon. Not. Roy. Astron. Soc.* **465** (2017) 1757 [[1606.00439](#)].
- [4] Z. Slepian et al., *Detection of baryon acoustic oscillation features in the large-scale three-point correlation function of SDSS BOSS DR12 CMASS galaxies*, *Mon. Not. Roy. Astron. Soc.* **469** (2017) 1738 [[1607.06097](#)].
- [5] D. W. Pearson and L. Samushia, *A Detection of the Baryon Acoustic Oscillation features in the SDSS BOSS DR12 Galaxy Bispectrum*, *Mon. Not. Roy. Astron. Soc.* **478** (2018) 4500 [[1712.04970](#)].
- [6] D. W. Pearson and L. Samushia, *Erratum: A Detection of the Baryon Acoustic Oscillation Features in the SDSS BOSS DR12 Galaxy Bispectrum*, *Mon. Not. Roy. Astron. Soc.* **483** (2018) 915.
- [7] N. S. Sugiyama, S. Saito, F. Beutler and H.-J. Seo, *A complete FFT-based decomposition formalism for the redshift-space bispectrum*, *Mon. Not. Roy. Astron. Soc.* **484** (2019) 364 [[1803.02132](#)].
- [8] DESI collaboration, *The DESI Experiment, a whitepaper for Snowmass 2013*, [1308.0847](#).
- [9] EUCLID collaboration, *Euclid Definition Study Report*, [1110.3193](#).
- [10] O. Doré et al., *Cosmology with the SPHEREX All-Sky Spectral Survey*, [1412.4872](#).
- [11] D. Spergel et al., *Wide-Field Infrared Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report*, [1503.03757](#).
- [12] K. C. Chan and L. Blot, *Assessment of the Information Content of the Power Spectrum and Bispectrum*, *Phys. Rev. D* **96** (2017) 023528 [[1610.06585](#)].

- [13] J. Byun, A. Eggemeier, D. Regan, D. Seery and R. E. Smith, *Towards optimal cosmological parameter recovery from compressed bispectrum statistics*, *Mon. Not. Roy. Astron. Soc.* **471** (2017) 1581 [1705.04392].
- [14] Y.-S. Song, A. Taruya and A. Oka, *Cosmology with anisotropic galaxy clustering from the combination of power spectrum and bispectrum*, *JCAP* **08** (2015) 007 [1502.03099].
- [15] P. Gagrani and L. Samushia, *Information Content of the Angular Multipoles of Redshift-Space Galaxy Bispectrum*, *Mon. Not. Roy. Astron. Soc.* **467** (2017) 928 [1610.03488].
- [16] V. Yankelevich and C. Porciani, *Cosmological information in the redshift-space bispectrum*, *Mon. Not. Roy. Astron. Soc.* **483** (2019) 2078 [1807.07076].
- [17] D. Gualdi and L. Verde, *Galaxy redshift-space bispectrum: the Importance of Being Anisotropic*, *JCAP* **06** (2020) 041 [2003.12075].
- [18] N. Agarwal, V. Desjacques, D. Jeong and F. Schmidt, *Information content in the redshift-space galaxy power spectrum and bispectrum*, 2007.04340.
- [19] D. Yamauchi, S. Yokoyama and H. Tashiro, *Constraining modified theories of gravity with the galaxy bispectrum*, *Phys. Rev. D* **96** (2017) 123516 [1709.03243].
- [20] B. Bose and A. Taruya, *The one-loop matter bispectrum as a probe of gravity and dark energy*, *JCAP* **10** (2018) 019 [1808.01120].
- [21] B. Bose, J. Byun, F. Lacasa, A. Moradinezhad Dizgah and L. Lombriser, *Modelling the matter bispectrum at small scales in modified gravity*, *JCAP* **02** (2020) 025 [1909.02504].
- [22] M. Tellarini, A. J. Ross, G. Tasinato and D. Wands, *Galaxy bispectrum, primordial non-Gaussianity and redshift space distortions*, *JCAP* **06** (2016) 014 [1603.06814].
- [23] D. Karagiannis, A. Lazanu, M. Liguori, A. Raccanelli, N. Bartolo and L. Verde, *Constraining primordial non-Gaussianity with bispectrum and power spectrum from upcoming optical and radio surveys*, *Mon. Not. Roy. Astron. Soc.* **478** (2018) 1341 [1801.09280].
- [24] R. Ruggeri, E. Castorina, C. Carbone and E. Sefusatti, *DEMNUi: Massive neutrinos and the bispectrum of large scale structures*, *JCAP* **03** (2018) 003 [1712.02334].
- [25] C. Hahn, F. Villaescusa-Navarro, E. Castorina and R. Scoccimarro, *Constraining M_ν with the bispectrum. Part I. Breaking parameter degeneracies*, *JCAP* **03** (2020) 040 [1909.11107].
- [26] P. Monaco, *Approximate methods for the generation of dark matter halo catalogs in the age of precision cosmology*, *Galaxies* **4** (2016) 53 [1605.07752].
- [27] M. Colavincenzo et al., *Comparing approximate methods for mock catalogues and covariance matrices – III: bispectrum*, *Mon. Not. Roy. Astron. Soc.* **482** (2019) 4883 [1806.09499].
- [28] B. Joachimi, *Non-linear shrinkage estimation of large-scale structure covariance*, *Mon. Not. Roy. Astron. Soc.* **466** (2017) L83 [1612.00752].
- [29] O. Friedrich and T. Eifler, *Precision matrix expansion - efficient use of numerical simulations in estimating errors on cosmological parameters*, *Mon. Not. Roy. Astron. Soc.* **473** (2018) 4150 [1703.07786].
- [30] A. Hall and A. Taylor, *A Bayesian method for combining theoretical and simulated covariance matrices for large-scale structure surveys*, *Mon. Not. Roy. Astron. Soc.* **483** (2019) 189 [1807.06875].
- [31] D. W. Pearson and L. Samushia, *Estimating the power spectrum covariance matrix with fewer mock samples*, *Mon. Not. Roy. Astron. Soc.* **457** (2016) 993 [1509.00064].
- [32] C. Howlett and W. J. Percival, *Galaxy two-point covariance matrix estimation for next generation surveys*, *Mon. Not. Roy. Astron. Soc.* **472** (2017) 4935 [1709.03057].

- [33] I. Mohammed, U. Seljak and Z. Vlah, *Perturbative approach to covariance matrix of the matter power spectrum*, *Mon. Not. Roy. Astron. Soc.* **466** (2017) 780 [[1607.00043](#)].
- [34] N. S. Sugiyama, S. Saito, F. Beutler and H.-J. Seo, *Perturbation theory approach to predict the covariance matrices of the galaxy power spectrum and bispectrum in redshift space*, *Mon. Not. Roy. Astron. Soc.* **497** (2020) 1684 [[1908.06234](#)].
- [35] D. Wadekar and R. Scoccimarro, *The Galaxy Power Spectrum Multipoles Covariance in Perturbation Theory*, [1910.02914](#).
- [36] A. Taruya, T. Nishimichi and D. Jeong, *The covariance of the matter power spectrum including the survey window function effect: N-body simulations vs. fifth-order perturbation theory on grid*, [2007.05504](#).
- [37] D. Gualdi, M. Manera, B. Joachimi and O. Lahav, *Maximal compression of the redshift space galaxy power spectrum and bispectrum*, *Mon. Not. Roy. Astron. Soc.* **476** (2018) 4045 [[1709.03600](#)].
- [38] D. Gualdi, H. Gil-Marín, R. L. Schuhmann, M. Manera, B. Joachimi and O. Lahav, *Enhancing BOSS bispectrum cosmological constraints with maximal compression*, *Mon. Not. Roy. Astron. Soc.* **484** (2019) 3713 [[1806.02853](#)].
- [39] A. Heavens, R. Jimenez and O. Lahav, *Massive lossless data compression and multiple parameter estimation from galaxy spectra*, *Mon. Not. Roy. Astron. Soc.* **317** (2000) 965 [[astro-ph/9911102](#)].
- [40] A. F. Heavens, E. Sellentin and A. H. Jaffe, *Extreme data compression while searching for new physics*, *Mon. Not. Roy. Astron. Soc.* **498** (2020) 3440 [[2006.06706](#)].
- [41] O. H. Philcox, M. M. Ivanov, M. Zaldarriaga, M. Simonovic and M. Schmittfull, *Fewer Mocks and Less Noise: Reducing the Dimensionality of Cosmological Observables with Subspace Projections*, [2009.03311](#).
- [42] D. Gualdi, H. Gil-Marín, M. Manera, B. Joachimi and O. Lahav, *Geometrical compression: a new method to enhance the BOSS galaxy bispectrum monopole constraints*, *Mon. Not. Roy. Astron. Soc.* **484** (2019) L29 [[1901.00987](#)].
- [43] D. Gualdi, H. Gil-Marín, M. Manera, B. Joachimi and O. Lahav, *GEOMAX: beyond linear compression for three-point galaxy clustering statistics*, *Mon. Not. Roy. Astron. Soc.* **497** (2020) 776 [[1912.01011](#)].
- [44] G. Pratten and D. Munshi, *Non-Gaussianity in Large Scale Structure and Minkowski Functionals*, *Mon. Not. Roy. Astron. Soc.* **423** (2012) 3209 [[1108.1985](#)].
- [45] M. Schmittfull, T. Baldauf and U. s. Seljak, *Near optimal bispectrum estimators for large-scale structure*, *Phys. Rev. D* **91** (2015) 043530 [[1411.6595](#)].
- [46] A. Moradinezhad Dizgah, H. Lee, M. Schmittfull and C. Dvorkin, *Capturing non-Gaussianity of the large-scale structure with weighted skew-spectra*, *JCAP* **04** (2020) 011 [[1911.05763](#)].
- [47] C.-T. Chiang, C. Wagner, F. Schmidt and E. Komatsu, *Position-dependent power spectrum of the large-scale structure: a novel method to measure the squeezed-limit bispectrum*, *JCAP* **05** (2014) 048 [[1403.3411](#)].
- [48] C.-T. Chiang, C. Wagner, A. G. Sánchez, F. Schmidt and E. Komatsu, *Position-dependent correlation function from the SDSS-III Baryon Oscillation Spectroscopic Survey Data Release 10 CMASS Sample*, *JCAP* **09** (2015) 028 [[1504.03322](#)].
- [49] C.-T. Chiang, *Position-dependent power spectrum: a new observable in the large-scale structure*, Ph.D. thesis, Munich U., 2015. [1508.03256](#).
- [50] D. Obreschkow, C. Power, M. Bruderer and C. Bonvin, *A Robust Measure of Cosmic Structure*

beyond the Power-Spectrum: Cosmic Filaments and the Temperature of Dark Matter, *Astrophys. J.* **762** (2013) 115 [1211.5213].

- [51] R. Wolstenhulme, C. Bonvin and D. Obreschkow, *Three-point Phase Correlations: a new Measure of Nonlinear Large-scale Structure*, *Astrophys. J.* **804** (2015) 132 [1409.3007].
- [52] A. Eggemeier, T. Battfeld, R. E. Smith and J. Niemeyer, *The Anisotropic Line Correlation Function as a Probe of Anisotropies in Galaxy Surveys*, *Mon. Not. Roy. Astron. Soc.* **453** (2015) 797 [1504.04036].
- [53] A. Eggemeier and R. E. Smith, *Cosmology with Phase Statistics: Parameter Forecasts and Detectability of BAO*, *Mon. Not. Roy. Astron. Soc.* **466** (2017) 2496 [1611.01160].
- [54] F. O. Franco, C. Bonvin, D. Obreschkow, K. Ali and J. Byun, *Probing redshift-space distortions with phase correlations*, *Phys. Rev.* **D99** (2019) 103530 [1805.10178].
- [55] K. Ali, D. Obreschkow, C. Howlett, C. Bonvin, C. Llinares, F. O. Franco et al., *Cosmological Constraints from Fourier Phase Statistics*, *Mon. Not. Roy. Astron. Soc.* **479** (2018) 2743 [1806.10276].
- [56] J. Byun, F. O. Franco, C. Howlett, C. Bonvin and D. Obreschkow, *Constraining the growth rate of structure with phase correlations*, *Mon. Not. Roy. Astron. Soc.* **497** (2020) 1765 [2005.06325].
- [57] J. Fergusson, M. Liguori and E. Shellard, *General CMB and Primordial Bispectrum Estimation I: Mode Expansion, Map-Making and Measures of f_{NL}* , *Phys. Rev. D* **82** (2010) 023502 [0912.5516].
- [58] J. Fergusson, M. Liguori and E. Shellard, *The CMB Bispectrum*, *JCAP* **12** (2012) 032 [1006.1642].
- [59] PLANCK collaboration, *Planck 2013 Results. XXIV. Constraints on primordial non-Gaussianity*, *Astron. Astrophys.* **571** (2014) A24 [1303.5084].
- [60] J. R. Fergusson, D. M. Regan and E. P. S. Shellard, *Rapid Separable Analysis of Higher Order Correlators in Large Scale Structure*, *Phys. Rev.* **D86** (2012) 063511 [1008.1730].
- [61] D. M. Regan, M. M. Schmittfull, E. P. S. Shellard and J. R. Fergusson, *Universal Non-Gaussian Initial Conditions for N-body Simulations*, *Phys. Rev.* **D86** (2012) 123524 [1108.3813].
- [62] M. M. Schmittfull, D. M. Regan and E. P. S. Shellard, *Fast Estimation of Gravitational and Primordial Bispectra in Large Scale Structures*, *Phys. Rev.* **D88** (2013) 063512 [1207.5678].
- [63] A. Lazanu, T. Giannantonio, M. Schmittfull and E. P. S. Shellard, *Matter bispectrum of large-scale structure: Three-dimensional comparison between theoretical models and numerical simulations*, *Phys. Rev.* **D93** (2016) 083517 [1510.04075].
- [64] A. Lazanu, T. Giannantonio, M. Schmittfull and E. Shellard, *Matter bispectrum of large-scale structure with Gaussian and non-Gaussian initial conditions: Halo models, perturbation theory, and a three-shape model*, *Phys. Rev. D* **95** (2017) 083511 [1511.02022].
- [65] J. Hung, J. R. Fergusson and E. P. S. Shellard, *Advancing the matter bispectrum estimation of large-scale structure: a comparison of dark matter codes*, **1902.01830**.
- [66] J. Hung, M. Manera and E. Shellard, *Advancing the matter bispectrum estimation of large-scale structure: fast prescriptions for galaxy mock catalogues*, **1909.03248**.
- [67] D. Regan, *An Inventory of Bispectrum Estimators for Redshift Space Distortions*, *JCAP* **12** (2017) 020 [1708.05303].
- [68] A. Oddo, E. Sefusatti, C. Porciani, P. Monaco and A. G. Sánchez, *Toward a robust inference method for the galaxy bispectrum: likelihood function and model selection*, *JCAP* **03** (2020) 056 [1908.01774].

- [69] D. Babich, *Optimal estimation of non-Gaussianity*, *Phys. Rev. D* **72** (2005) 043003 [[astro-ph/0503375](#)].
- [70] T. Hahn, *CUBA: A Library for multidimensional numerical integration*, *Comput. Phys. Commun.* **168** (2005) 78 [[hep-ph/0404043](#)].
- [71] T. Hahn, *Concurrent Cuba*, *J. Phys. Conf. Ser.* **608** (2015) 012066 [[1408.6373](#)].
- [72] W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery, *Numerical Recipes in Fortran 77: the Art of Scientific Computing. Second Edition*, vol. 1. 1996.
- [73] J. N. Grieb, A. G. Sánchez, S. Salazar-Albornoz and C. Dalla Vecchia, *Gaussian covariance matrices for anisotropic galaxy clustering measurements*, *Mon. Not. Roy. Astron. Soc.* **457** (2016) 1577 [[1509.04293](#)].
- [74] P. Monaco, T. Theuns and G. Taffoni, *Pinocchio: pinpointing orbit-crossing collapsed hierarchical objects in a linear density field*, *Mon. Not. Roy. Astron. Soc.* **331** (2002) 587 [[astro-ph/0109323](#)].
- [75] P. Monaco, E. Sefusatti, S. Borgani, M. Crocce, P. Fosalba, R. Sheth et al., *An accurate tool for the fast generation of dark matter halo catalogs*, *Mon. Not. Roy. Astron. Soc.* **433** (2013) 2389 [[1305.1505](#)].
- [76] E. Munari, P. Monaco, E. Sefusatti, E. Castorina, F. G. Mohammad, S. Anselmi et al., *Improving fast generation of halo catalogues with higher order Lagrangian perturbation theory*, *Mon. Not. Roy. Astron. Soc.* **465** (2017) 4658 [[1605.04788](#)].
- [77] E. Sefusatti, M. Crocce, R. Scoccimarro and H. Couchman, *Accurate Estimators of Correlation Functions in Fourier Space*, *Mon. Not. Roy. Astron. Soc.* **460** (2016) 3624 [[1512.07295](#)].
- [78] E. Sellentin and A. F. Heavens, *Parameter inference with estimated covariance matrices*, *Mon. Not. Roy. Astron. Soc.* **456** (2016) L132 [[1511.05969](#)].
- [79] G. M. Kaufmann, *Some Bayesian Moment Formulae*, Report No. 6710. Centre for Operations Research and Econometrics. Catholic University of Louvain, Heverlee, 1967.
- [80] T. W. Anderson, *An Introduction to Multivariate Statistical Analysis*. Wiley, 2003.
- [81] J. Hartlap, P. Simon and P. Schneider, *Why your model parameter confidences might be too optimistic: Unbiased estimation of the inverse covariance matrix*, *Astron. Astrophys.* (2006) [[astro-ph/0608064](#)].
- [82] D. Foreman-Mackey, D. W. Hogg, D. Lang and J. Goodman, *emcee: The MCMC Hammer*, *Publ. Astron. Soc. Pac.* **125** (2013) 306 [[1202.3665](#)].
- [83] A. Lewis, *GetDist: a Python package for analysing Monte Carlo samples*, [1910.13970](#).
- [84] A. Oddo et al., *in preparation*.
- [85] D. Jeong, *Cosmology with high ($z>1$) redshift galaxy surveys*, Ph.D. thesis, The University of Texas at Austin, 2010.
- [86] C. A. Watkinson, S. Majumdar, J. R. Pritchard and R. Mondal, *A fast estimator for the bispectrum and beyond – a practical method for measuring non-Gaussianity in 21-cm maps*, *Mon. Not. Roy. Astron. Soc.* **472** (2017) 2436 [[1705.06284](#)].
- [87] M. Tegmark, A. N. Taylor and A. F. Heavens, *Karhunen-Loève Eigenvalue Problems in Cosmology: How Should We Tackle Large Data Sets?*, *Astrophys. J.* **480** (1997) 22 [[astro-ph/9603021](#)].
- [88] L. Knox, R. Scoccimarro and S. Dodelson, *The Impact of inhomogeneous reionization on cosmic microwave background anisotropy*, *Phys. Rev. Lett.* **81** (1998) 2004 [[astro-ph/9805012](#)].
- [89] A. Amara and A. Réfrégier, *Systematic bias in cosmic shear: extending the Fisher matrix*, *Mon. Not. Roy. Astron. Soc.* **391** (2008) 228 [[0710.5171](#)].

- [90] L. Blot, P. S. Corasaniti, L. Amendola and T. D. Kitching, *Non-Linear Matter Power Spectrum Covariance Matrix Errors and Cosmological Parameter Uncertainties*, *Mon. Not. Roy. Astron. Soc.* **458** (2016) 4462 [1512.05383].

A 1-dimensional basis functions

The separable basis functions, $Q_n(k_1, k_2, k_3)$, that we use in this work are constructed out of a product of three 1-dimensional basis functions,

$$Q_n(k_1, k_2, k_3) = q_{\{p\}}(k_1)q_{\{r\}}(k_2)q_{\{s\}}(k_3). \quad (\text{A.1})$$

The p , r , and s subscripts on the right side index the different 1-dimensional functions that we have chosen, and the curly brackets require that the Q_n functions are invariant to permutations of k_1 , k_2 , and k_3 .

In principle, the q_n basis can be any function, but in this work we implement and compare two choices: normal polynomials and shifted Legendre polynomials. The normal q_n basis we use is constructed as described in [57], while the choice of shifted Legendre polynomials is

$$q_n(x) \equiv \tilde{P}_n(x) = P_n(2x - 1), \quad (\text{A.2})$$

where P_n are the usual non-shifted Legendre polynomials.

Given a definition of q_n , we still need to specify how we map between $\{prs\} \leftrightarrow n$ in eq. (A.1). Our convention is that we group the Q_n according to the maximum power of any term, $p + r + s$, and then within each group, we order the Q_n by sorting by increasing s , r , and finally p . Explicitly, our first 11 Q_n have prs as listed in Table 2.

In choosing the number of modes in a basis, we will always choose a number of modes such that all modes up to a given $p + r + s$ are included. For example, in determining how many modes are necessary for parameter constraints to converge when $k_{\text{max}} = 24.5 k_f$ and custom modes are not included, we compared results with $p + r + s \leq 3$ (7 modes), $p + r + s \leq 4$ (11 modes), $p + r + s \leq 5$ (16 modes), and so on, up to $p + r + s \leq 12$ (102 modes). This is a somewhat arbitrary, but simple, way of grouping modes together to simplify the analysis whenever we check how our results depend on how many modes are used.

We note that this choice of ordering is mostly arbitrary, though larger values of $p + r + s$ generally correspond to Q_n with smaller scale variations, such that the larger n modes are expected to have smaller amplitudes in the weighted bispectrum. This would also be true for the alternative option of ordering the modes according to their ‘distance’, $\sqrt{p^2 + r^2 + s^2}$, which was also described in [57].

B Calculating the inner product using 1-dimensional FFT

In this appendix, we detail the calculation of eqs. (2.30) and (2.31), copied here for convenience,

$$\begin{aligned} \langle\langle Q_n | Q_m \rangle\rangle = \frac{1}{2\pi^5} \int dx \frac{1}{x} & \{ F_{pa}(x) [F_{rb}(x)F_{sc}(x) + F_{rc}(x)F_{sb}(x)] \\ & + F_{pb}(x) [F_{rc}(x)F_{sa}(x) + F_{ra}(x)F_{sb}(x)] \\ & + F_{pc}(x) [F_{ra}(x)F_{sb}(x) + F_{rb}(x)F_{sa}(x)] \}, \end{aligned} \quad (\text{B.1})$$

n	p	r	s
0	0	0	0
1	0	0	1
2	0	1	1
3	0	0	2
4	1	1	1
5	0	1	2
6	0	0	3
7	1	1	2
8	0	2	2
9	0	1	3
10	0	0	4
etc.

Table 2: We show the prs corresponding to the first 11 Q_n basis functions, according to our convention of grouping functions first by $p+r+s$ (as shown by the horizontal lines) and within each group sorting by increasing s , r , and then p .

where

$$F_{pa}(x) \equiv \int dk q_p(k) q_a(k) \sin(kx). \quad (\text{B.2})$$

Specifically, we use 1-dimensional FFTs to compute eq. (B.2), and we require it for a range and resolution of x such that the outer integral over x in eq. (B.1) can be numerically calculated to sufficient accuracy. Doing this requires some care, and we follow the steps described in Numerical Recipes (Chapter 13.9, ‘‘Computing Fourier Integrals Using the FFT’’, hereafter NR). Here we summarize the key equations that we use in our case, and we refer the reader to NR for the step-by-step derivation of the method and more general expressions.

The aim is to numerically evaluate

$$\int_a^b h(t) e^{iwt} dt, \quad (\text{B.3})$$

and we begin by approximating $h(t)$ using an interpolation

$$h(t) \approx \sum_{j=0}^M h_j \psi\left(\frac{t-t_j}{\Delta}\right) + \sum_{j=\text{endpts}} h_j \phi_j\left(\frac{t-t_j}{\Delta}\right), \quad (\text{B.4})$$

where $\Delta \equiv (b-a)/M$, $h_j \equiv h(t_j)$, and $t_j \equiv a + j\Delta$ for $j = 0, \dots, M$. ψ and ϕ_j are kernel functions that depend on the interpolation scheme, which also determines which points at the boundary count as endpoints in the second sum in eq. (B.4).

We substitute this interpolating function into eq. (B.3), interchange the sum and integral in each of the two terms, and then make the change of variable $s \equiv (t-t_j)/\Delta$ in the first term and $s \equiv (t-a)/\Delta$ in the second term. This leads to

$$\int_a^b h(t) e^{iwt} dt \approx \Delta e^{iwa} \left[W(\theta) \sum_{j=0}^M h_j e^{i\theta j} + \sum_{j=\text{endpts}} h_j \alpha_j(\theta) \right], \quad (\text{B.5})$$

where $\theta \equiv w\Delta$ and the two functions are

$$W(\theta) \equiv \int_{-\infty}^{\infty} ds e^{i\theta s} \psi(s) \quad (\text{B.6})$$

$$\alpha_j(\theta) \equiv \int_{-\infty}^{\infty} ds e^{i\theta s} \phi_j(s-j). \quad (\text{B.7})$$

If the endpoint kernel is symmetric, then

$$\phi_{M-j}(s) = \phi_j(-s) \quad (\text{B.8})$$

$$\alpha_{M-j}(\theta) = e^{iw(b-a)} \alpha_j^*(\theta), \quad (\text{B.9})$$

so the second sum in eq. (B.5) can be more explicitly written as

$$\begin{aligned} \sum_{j=\text{endpts}} h_j \alpha_j(\theta) &= \alpha_0(\theta) h_0 + \alpha_1(\theta) h_1 + \alpha_2(\theta) h_2 + \alpha_3(\theta) h_3 + \dots \\ &+ e^{iw(b-a)} [\alpha_0^*(\theta) h_M + \alpha_1^*(\theta) h_{M-1} + \alpha_2^*(\theta) h_{M-2} + \alpha_3^*(\theta) h_{M-3} + \dots], \end{aligned} \quad (\text{B.10})$$

where the ellipses represent terms that are dropped when the interpolation kernels are cubic splines (or lower order).

We perform the first sum in eq. (B.5) using an FFT. The FFT grid size N must be $N \geq M + 1$, and it determines the values of w (and θ) that are sampled,

$$w_n \Delta \equiv \frac{2\pi n}{N}, \quad (\text{B.11})$$

for $n = 0, \dots, \frac{N}{2} - 1$. Therefore, we see that M and N are both free parameters of the calculation. Larger M allows the integration to be sensitive to higher frequency oscillations in t , and M must be higher if $h(t)$ is varying rapidly with t . N on the other hand is the grid resolution used for the FFT that we use to evaluate the sum in the equation above, so for a fixed M , a larger N produces finer sampling in w -space, which is especially important if the result of the integral is then going to be interpolated for different values of w , as we do when this is used to calculate the inner product.

In this work, we use the cubic order kernel functions in NR, which are implemented in the subroutine `dftcor` provided there, and use a modification of the code in `dftint` also provided to calculate the integral for only the $\int_a^b \sin(wt) h(t) dt$ part.