

Squeezing f_{NL} out of the matter bispectrum with consistency relations

Samuel Goldstein^{1,2,*}, Angelo Esposito^{3,4}, Oliver H. E. Philcox^{1,5,6,3}, Lam Hui,¹
J. Colin Hill^{1,7}, Roman Scoccimarro,⁸ and Maximilian H. Abitbol⁹

¹*Center for Theoretical Physics, Department of Physics, Columbia University,
New York, New York 10027, USA*

²*Max-Planck-Institute für Astrophysik, Karl-Schwarzschild-Straße 1, 85748 Garching, Germany*

³*School of Natural Sciences, Institute for Advanced Study,
1 Einstein Drive, Princeton, New Jersey 08540, USA*

⁴*Dipartimento di Fisica, Sapienza Università di Roma, Piazzale Aldo Moro 2, I-00185 Rome, Italy*

⁵*Simons Society of Fellows, Simons Foundation, New York, New York 10010, USA*

⁶*Department of Astrophysical Sciences, Princeton University, Princeton, New Jersey 08540, USA*

⁷*Center for Computational Astrophysics, Flatiron Institute, New York, New York 10010, USA*

⁸*Center for Cosmology and Particle Physics, Department of Physics, New York University,
New York, New York 10003, USA*

⁹*University of Oxford, Department of Physics, Denys Wilkinson Building,
Keble Road, Oxford OX1 4LS, United Kingdom*



(Received 16 September 2022; accepted 2 December 2022; published 29 December 2022)

We show how consistency relations can be used to robustly extract the amplitude of local primordial non-Gaussianity (f_{NL}) from the squeezed limit of the matter bispectrum, well into the nonlinear regime. First, we derive a nonperturbative relation between primordial non-Gaussianity and the leading term in the squeezed bispectrum, revising some results present in the literature. This relation is then used to successfully measure f_{NL} from N -body simulations. We discuss the dependence of our results on different scale cuts and redshifts. Specifically, the analysis is strongly dependent on the choice of the smallest soft momentum, q_{min} , which is the most sensitive to primordial bispectrum contributions, but is largely independent of the choice of the largest hard momentum, k_{max} , due to the non-Gaussian nature of the covariance. We also show how the constraints on f_{NL} improve at higher redshift, due to a reduced off-diagonal covariance. In particular, for a simulation with $f_{\text{NL}} = 100$ and a volume of $(2.4 \text{ Gpc}/h)^3$, we measure $f_{\text{NL}} = 98 \pm 12$ at redshift $z = 0$ and $f_{\text{NL}} = 97 \pm 8$ at $z = 0.97$. Finally, we compare our results with a Fisher forecast, showing that the current version of the analysis is satisfactorily close to the Fisher error. We regard this as a first step towards the realistic application of consistency relations to constrain primordial non-Gaussianity using observations.

DOI: [10.1103/PhysRevD.106.123525](https://doi.org/10.1103/PhysRevD.106.123525)

I. INTRODUCTION

One of the longest standing questions in modern cosmology concerns the mechanism that set the initial conditions for the classical evolution of matter in the Universe. In particular, one asks whether an initial inflationary period has taken place and, if so, what were the underlying physics. In this respect, a crucial aspect is to understand if the primordial fluctuations are consistent with those

predicted by single-field inflation, or whether additional light degrees of freedom played a role. To answer this, one typically relies on data for fluctuations that are small enough to lie within the quasilinear regime, whereupon perturbation theory provides reliable predictions. Such a theory can then be used to robustly extract the physical parameters. In contrast, modes at larger momentum (e.g. $k \gtrsim 0.2 h/\text{Mpc}$ for large-scale structures at redshift zero) fall in the nonlinear regime and are usually discarded due to modeling difficulties, despite being abundant and precisely measured. It would be ideal to have a tool able to take advantage of this data.

One possibility arises from the “consistency relations,” which are identities connecting the squeezed limit of $(N + 1)$ -point correlators (i.e. those with one vanishing momentum leg) and N -point correlators. Importantly, they follow solely from symmetry arguments and thus remain

*sjg2215@columbia.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by the Max Planck Society.

valid in the nonlinear regime, where conventional perturbation theory fails. After being first pointed out by Maldacena [1], intensive work has been done in discovering additional relations arising from new symmetries as well as in clarifying the assumptions underlying their formulation [2–14]. In the context of large-scale structure, consistency relations can be traced back to the Galilean invariance of the evolution equations, as first identified by Kehagias and Riotto [15] and Peloso and Pietroni [16], and further investigated in [17–22].

Crucially, the consistency relations for a large-scale structure are modified depending on whether the initial conditions for the matter fluctuations are Gaussian, as verified nonperturbatively in [23]. More precisely, if the initial conditions correspond to those of single-field slow-roll inflation, we expect $B(\mathbf{q}, \mathbf{k})/P(q)$ to have no inverse-powers-of- q pole. Here, B and P are the bispectrum and power spectrum respectively, and \mathbf{q}, \mathbf{k} are the momenta with the understanding that $q \ll k$ (the squeezed limit). Moreover, the symmetries from which the consistency relations originate characterize not only the evolution of matter distribution, but also that of astrophysically complex objects, such as galaxies (which can merge, form, annihilate, and so on) [22, 24–26]. Finally, consistency relations are expected to hold true in redshift space [18, 22, 27], making them amenable for application to galaxy surveys, which include strong anisotropy around the line-of-sight direction due to redshift-space distortions.

In this work, we take an important step towards the practical implementation of consistency relations as a tool to obtain nonperturbative information on the initial conditions of the Universe, by studying the squeezed limit of the matter bispectrum. Using a suite of N -body simulations of matter in real space, we show that the consistency relations can be used to precisely measure the local primordial non-Gaussianity parameter, f_{NL} . The idea is that if the initial conditions were *not* that of single-field slow-roll inflation (the curvaton [28, 29] model being an example, characterized by a nonvanishing f_{NL}), the squeezed B/P is no longer protected from having poles in the soft momentum q . In what follows, we consider only local-type non-Gaussianities which correspond to integer order poles as we will show in Sec. II. Alternative inflation models, such as quasisingle field inflation [7, 30–32], can lead to noninteger order poles or oscillatory behavior in the squeezed B/P . We leave a discussion of such models to a future work.

We first derive the nonperturbative relation between the squeezed limit of the bispectrum and f_{NL} in Sec. II, correcting and clarifying some results present in the literature. Section III then highlights and addresses a number of subtleties regarding the procedure of measuring and analyzing the bispectrum in this context, the union of which allows for robust constraints on primordial non-Gaussianity, as presented in Sec. IV. Finally, we conclude

and discuss the ingredients necessary for a future, realistic application of consistency relations to large-scale structure data in Sec. V.

The Appendices give useful theoretical results and methodology validation. Appendix A includes tree-level and one-loop standard Eulerian perturbation theory calculations of the squeezed bispectrum in the presence of local primordial non-Gaussianity. Appendix B validates our bispectrum parametrization. We derive our likelihood and Fisher matrix in Appendix C 1 and present a quasi-optimal weighting scheme in Appendix D. In Appendix E we show the full posterior of our main analysis.

Conventions.—Throughout this paper we work in natural units, $c = 1$. Our fiducial cosmological parameters are $\Omega_{m,0} = 0.25$ (of which $\Omega_b = 0.04$), $\Omega_\Lambda = 0.75$, $h = 0.7$, $n_s = 1$, and $\sigma_8 = 0.8$.

II. SQUEEZED BISPECTRUM IN THE PRESENCE OF LOCAL PRIMORDIAL NON-GAUSSIANITY

A. Derivation

Let us begin by deriving the squeezed limit of the bispectrum in the presence of local primordial non-Gaussianity. We will work nonperturbatively in gravitational evolution, and follow a similar approach to the derivation of scale-dependent bias found in [33]. Additionally, we assume a single redshift, z , and generally suppress this argument for clarity. In the presence of a background (soft) potential, Φ_L , the locally measured small-scale power spectrum is modulated, taking the form¹

$$P(k|\Phi_L) = P(k) + \sum_{q'} \frac{\partial P(k)}{\partial \Phi_{L,q'}} \Phi_{L,q'} + \dots, \quad (1)$$

where we are assuming that $q' \equiv |\mathbf{q}'| \ll k$, such that we can treat the long mode as a background, in the presence of which the power spectrum is evaluated. This induces a coupling between the hard mode power spectrum, $P(k)$, and the soft mode density, δ_q . The bispectrum is then obtained from the $\langle \delta_q P(k|\Phi_L) \rangle$ correlator, and is given by

$$\begin{aligned} B(\mathbf{q}, \mathbf{k}) &= V \sum_{q'} \frac{\partial P(k)}{\partial \Phi_{L,q'}} \langle \delta_q \Phi_{L,q'} \rangle + \dots \\ &= V \frac{\partial P(k)}{\partial \Phi_{L,-q}} \frac{P(q)}{\alpha(q)} + \dots, \end{aligned} \quad (2)$$

where V is the volume. Here $P(q)$ is soft and thus in the linear regime. This result uses the Poisson equation to relate the long-wavelength density and potential, $\delta_q = \alpha(q)\Phi_{L,q}$, with

¹In our notation, the power spectrum and bispectrum are defined as $\langle \delta_q \delta_k \rangle \equiv V \delta_{q+k}^K P(q)$ and $\langle \delta_q \delta_k \delta_p \rangle \equiv V \delta_{q+k+p}^K B(\mathbf{q}, \mathbf{k})$ respectively, with V the volume and δ^K the Kronecker delta.

$$\alpha(q) \equiv \frac{2}{3} \frac{q^2 T(q) D_{\text{md}}(z)}{\Omega_{m,0} H_0^2}, \quad (3)$$

where $\Omega_{m,0}$ and H_0 are the matter density and expansion rate today, $T(q)$ is the transfer function (normalized to unity for $q \rightarrow 0$), and $D_{\text{md}}(z)$ is the growth rate, normalized to $a(z) = 1/(1+z)$ in the matter-dominated era.² We are interested in the form of Eq. (2) in the presence of local primordial non-Gaussianity, which can be parametrized by a primordial gravitational potential at subhorizon scales given by (e.g. [34,35])

$$\Phi(\mathbf{x}) = \phi(\mathbf{x}) + f_{\text{NL}}(\phi^2(\mathbf{x}) - \langle \phi^2 \rangle), \quad (4)$$

where $\phi(\mathbf{x})$ is a Gaussian random field with expectation value $\langle \phi \rangle$.

In the separate Universe picture, the soft mode acts as a background rescaling, equivalent to varying the locally measured clustering amplitude, σ_8 , i.e.

$$B(\mathbf{q}, \mathbf{k}) = V \frac{\partial P(k)}{\partial \log \sigma_8} \frac{\partial \log \sigma_8}{\partial \Phi_{L,-q}} \frac{P(q)}{\alpha(q)} + \dots \quad (5)$$

To compute the potential derivative, we utilize the definition of σ_8 and f_{NL} (e.g. [36–38]):

$$\sigma_8(\Phi_L) = \sigma_8(0) \left[1 + \frac{2f_{\text{NL}}}{V} \sum_{\mathbf{q}'} \Phi_{L,\mathbf{q}'} + \mathcal{O}(f_{\text{NL}}^2) \right], \quad (6)$$

where we have noted that $\Phi_L(\mathbf{x})$ is a slowly varying field, thus we can neglect its spatial dependence. Combining Eqs. (3), (5), and (6), we find the leading term in the squeezed bispectrum to be

$$B(\mathbf{q}, \mathbf{k}) = \frac{6f_{\text{NL}}\Omega_{m,0}H_0^2}{D_{\text{md}}(z)} \frac{\partial P(k)}{\partial \log \sigma_8^2} \frac{P(q)}{q^2 T(q)} + \mathcal{O}(f_{\text{NL}}^2). \quad (7)$$

The result above differs from some former works in important ways. First, it may appear to differ from, e.g. [25,39], which present a rigorous tree-level definition and find a k dependence of the form $P(k)$ rather than $\partial P(k)/\partial \log \sigma_8^2$. This discrepancy is resolved by noting that, at first (tree-level) order, $P(k) \propto \sigma_8^2$, and thus the two coincide. More importantly, a nonperturbative derivation of the squeezed bispectrum is also presented in [16], but neither the logarithmic derivative nor the $D_{\text{md}}(z)$ factor appear in the final result. The presence of these factors is also verified by a direct one-loop calculation, as we report in Appendix A. In what follows, we further confirm their necessity in order to recover the correct f_{NL} when the latter

²Note that this differs from the normalization used in many cosmological codes, e.g. CLASS.

is extracted from simulations including primordial non-Gaussianity.

B. Modeling of the bispectrum

To extract the primordial signal given in Eq. (7), we require a model for the squeezed bispectrum that also encodes the contributions coming from the nonprimordial gravitational evolution. Beyond perturbative scales, the functional form of these is not well understood. Nonetheless, consistency relations tell us that the dependence on the soft mode is known, and that these late time contributions to the squeezed bispectrum scale with positive powers of the soft mode, q . In light of this, we write the squeezed bispectrum in the following form:

$$B(q, k, \theta) = a_{-2}(k) \frac{P(q)}{q^2 T(q)} + a_{-1}(k, \theta) \frac{P(q)}{q T(q)} + \sum_{n=0}^{\infty} a_n(k, \theta) q^n P(q), \quad (8)$$

where we choose a parametrization in terms of the magnitude of the softest momentum, q , the magnitude of one of the hard modes, k , and the angle between them, θ . Due to the consistency relations, a_{-2} and a_{-1} can only be sourced by primordial physics beyond single-field slow-roll inflation.³ Our parametrization includes only integer order poles since we are assuming local primordial non-Gaussianity. For the range of soft momenta we consider here the transfer function cannot be represented as a simple power series, thus we must include it explicitly. Inspired by perturbation theory, we only include it in the first two terms, since the remaining pieces are dominated by the bispectrum due to gravitational evolution, rather than by the primordial one, see Eq. (7). We will validate this parametrization in Sec. IV.

A comparison with Eq. (7) tells us that, at leading order in f_{NL} , local primordial non-Gaussianity leads to a nonzero a_{-2} coefficient,

$$a_{-2}(k) = \frac{6f_{\text{NL}}\Omega_{m,0}H_0^2}{D_{\text{md}}(z)} P(k) \frac{\partial \log P(k)}{\partial \log \sigma_8^2}, \quad (9)$$

while the other nonperturbative coefficients are *a priori* unknown.⁴ Nonetheless, one can still use invariance under parity to argue that [42]

³This is no longer true for unequal time correlators, which are expected to feature a (computable) $1/q^2$ term even in the absence of primordial non-Gaussianity [40].

⁴A prediction for a_0 can actually be made, as shown and verified in [41]. Nonetheless, it relies on the nonlinear observable being mass fluctuations, as opposed to, for example, galaxy ones. Given the scope of our broader program, we will remain agnostic about the nature of all nonperturbative coefficients, except for a_{-2} .




FIG. 1. Triangles related by the parity transformation reported in Eq. (10). When all pairs of this sort are added together, the odd terms in the expansion of the bispectrum vanish.

$$a_n(k, \pi - \theta) = (-1)^n a_n(k, \theta). \quad (10)$$

This means that, if one properly averages the bispectrum over all possible relative angles, the contributions to the odd coefficients coming from triangles related by the above parity transformation cancel (see Fig. 1). This reduces the set of unknown parameters to the even coefficients. We will return to this aspect in Sec. III A and Appendix B.

Our plan of action is the following: use N -body simulations to measure the a_n coefficients (properly averaged over angles and hard momenta, as explained in Sec. III A), then extract f_{NL} from a_{-2} using Eq. (9).

III. METHODOLOGY

To validate our measurements of f_{NL} , we use a suite of N -body simulations consisting of 40 realizations with Gaussian initial conditions and 12 realizations with local primordial non-Gaussianity ($f_{\text{NL}} = 100$). The simulations were run with 1280^3 particles in a comoving cubic volume of $L^3 = (2.4 \text{ Gpc}/h)^3$. The cosmological parameters are $\Omega_{m,0} = 0.25$ (of which $\Omega_b = 0.04$), $\Omega_\Lambda = 0.75$, $h = 0.7$, $n_s = 1$, and $\sigma_8 = 0.8$. We analyze the $z = 0$ and $z = 0.97$ snapshots. For further details on the simulations, we refer the reader to [35].

A. Measuring the power spectrum and bispectrum

To obtain an estimate of f_{NL} using the consistency relations, we will require both the matter power spectrum and bispectrum. These are measured using fast Fourier transforms, first assigning the simulation particles to a 1024^3 mesh using Triangular Shaped Cloud interpolation. The soft momentum is collected in bins of width $2k_f$, with $k_f \equiv 2\pi/L \simeq 2.6 \times 10^{-3} h/\text{Mpc}$ being the fundamental mode, starting from modes of magnitude $q > k_f$ up to $q_{\text{max}} = 25k_f \simeq 0.065 h/\text{Mpc}$ —i.e. 12 soft momentum bins. The hard modes entering the bispectrum will, instead, be taken within a range that is variable, depending on our analysis choices (as discussed below).

Throughout our analysis we pay particular attention to avoid any artifacts induced by binning and the discrete nature of the momenta. Such effects are more prominent for the softest momenta which, at the same time, are the most sensitive to the signal arising from primordial non-Gaussianity. For this reason, as far as the power spectrum and its related quantities are concerned, we refrain from any

modeling (i.e. using linear theory) and directly measure the following binned quantities:

$$\hat{M}_{i,n} \equiv \begin{cases} \langle q^n \delta_q \delta_{-q} / T(q) \rangle_i, & \text{if } n < 0 \\ \langle q^n \delta_q \delta_{-q} \rangle_i, & \text{otherwise,} \end{cases} \quad (11)$$

where by $\langle \dots \rangle_i$ we represent the average over q modes whose magnitude falls within the i th bin. By directly measuring $\hat{M}_{i,n}$ from the simulations, we also take advantage of the sample variance cancellation arising from the fact that bispectrum and power spectrum are measured from the same density field.

For the bispectrum, we measure the following quantities:

$$\hat{B}_i \equiv \langle \text{Re} \delta_q \delta_k \delta_{-q-k} \rangle_i, \quad \text{with } k \in [k_{\text{min}}, k_{\text{max}}], \quad (12)$$

where k_{min} and k_{max} define the hard momentum scales of our analysis. We sum over all k s with magnitude between these values. More precisely, if the relative angle between q and k is allowed to cover the full range between 0 and 2π , then the hard modes contributing to the bispectrum in the soft momenta bin q_i effectively range from $k_{\text{min}} - q_i$ to $k_{\text{max}} + q_i$, and hence the hard momenta range varies with q . In order to avoid finite resolution effects from the simulation or our particle assignment scheme, we impose a maximum momentum in our analysis of $K = 270k_f \simeq 0.71 h/\text{Mpc}$. This means that we must always make sure that $k_{\text{max}} + q_{\text{max}} \leq K$. Once this is done, we are guaranteed that the average in Eq. (12) is done over all possible relative angles, θ , hence ensuring that the odd coefficients in the expansion in Eq. (8) average to zero, as discussed below Eq. (10). Moreover, we simply average over all k s between k_{min} and k_{max} with uniform weight, instead of treating different k bins separately. This way we can remain agnostic about the k dependence of all coefficients in Eq. (8) except for $a_{-2}(k)$, otherwise we would need to introduce a new set of (unknown) coefficients for each k bin.⁵

In light of the discussion above, our theory model for the bispectrum is given by

$$\hat{B} = \sum_{n=-2,0,2} \bar{a}_n \hat{M}_n, \quad (13)$$

where vectors run over the soft mode index, i , and the \bar{a}_n coefficients correspond to the a_n s in Eq. (8) averaged over angles and over hard modes with magnitude between k_{min} and k_{max} . In Appendix B we show that repeating the

⁵Alternatively, one could average the bispectrum over different k modes using an optimal weighting scheme. We tested several quasioptimal weighting strategies (derived in Appendix D) and found negligible improvement on our constraints, thus we defer a discussion of weighting to a future work. The Fisher forecast presented in Sec. IV B suggests that our analysis is fairly close to optimal even without including k -dependent weights.

analysis with the inclusion of the odd coefficients (\bar{a}_{-1}, \bar{a}_1) leads to values consistent with zero. This confirms that the average over triangles is done properly, and it is therefore sensible to only consider the subset of even coefficients. We also show that \bar{a}_4 is consistent with zero so that we can truncate the series at \bar{a}_2 .

Finally, in order to extract f_{NL} from \bar{a}_{-2} , we measure $\bar{F} \equiv \langle \delta_k \delta_{-k} \partial \log P(k) / \partial \log \sigma_8^2 \rangle$, averaged over $k \in [k_{\text{min}}, k_{\text{max}}]$, using the same bins as the bispectrum. A couple of subtleties arise when estimating this quantity. First, in principle, $P(k)$ should be computed from a theoretical model as it is an average over realizations; however, since we do not *a priori* know $P(k)$ in the nonlinear regime, we instead measure it from the simulations. This does not significantly impact our results since there is little variance in $P(k)$ at high k . Second, $\partial \log P(k) / \partial \log \sigma_8^2$ should be determined from simulations, for example, via a separate Universe prescription; however, this is a computationally expensive and time consuming process. For our fiducial analysis it suffices to estimate the function using HALOFIT [43]. This represents the main source of theoretical uncertainty when constraining f_{NL} from consistency relations using the real-space matter distribution.⁶ For a measured \bar{F} and \bar{a}_{-2} (averaged over the same hard momenta), we can compute f_{NL} using the relationship

$$f_{\text{NL}} = \frac{\bar{a}_{-2}}{\bar{F}} \frac{D_{\text{md}}(z)}{6\Omega_{m,0}H_0^2}. \quad (14)$$

B. Likelihood and covariance

To formulate our likelihood for the nonperturbative coefficients, \bar{a}_n , we must account for the stochastic nature of both the measured bispectrum $\hat{\mathbf{B}}$ and the theory model, the latter of which depends on the measured binned power spectra, $\hat{\mathbf{M}}_n$, which we also consider as a stochastic variable. As shown in Appendix C 1, this can be taken into account by a likelihood that, for a single realization, r , is given by

$$\mathcal{L} \propto \frac{1}{\sqrt{\det \mathcal{C}(\bar{a}_n)}} \left[1 + \frac{\delta \mathbf{B}(\bar{a}_n) \cdot \mathcal{C}^{-1}(\bar{a}_n) \cdot \delta \mathbf{B}(\bar{a}_n)}{N_R - 1} \right]^{-\frac{N_R}{2}}, \quad (15)$$

where $\delta \mathbf{B}(\bar{a}_n) \equiv \hat{\mathbf{B}}^r - \sum_n \bar{a}_n \hat{\mathbf{M}}_n^r$ is the difference between the measured bispectrum and our theory model [Eq. (13)], $\mathcal{C}_{ij}(\bar{a}_n) \equiv \langle \delta B_i(\bar{a}_n) \delta B_j(\bar{a}_n) \rangle$ is the *parameter dependent* covariance of $\delta \mathbf{B}$, and N_R is the number of realizations used

⁶To elucidate the importance of this theoretical error, we have repeated the analysis instead taking the $\partial \log P(k) / \partial \log \sigma_8^2$ derivative from the QUIOTE simulations [44] at their fiducial cosmology. We find that the final value of f_{NL} changes by less than 5%, which is smaller than our statistical uncertainty (see Sec. IV).

to estimate it. Note that we use a multivariate t distribution instead of a Gaussian likelihood to account for the fact that our covariance is estimated from a finite number of realizations [45] as discussed below.

Since a real analysis would assume a fiducial value of $f_{\text{NL}} = 0$ for the covariance and we only have twelve realizations with primordial non-Gaussianity, we always compute our covariance using the $N_R = 40$ realizations with Gaussian initial conditions. Furthermore, to ease the computational burden associated with sampling a likelihood with a parameter dependent covariance, we evaluate the latter at a fixed set of fiducial \bar{a}_n . Evaluating the covariance at a fiducial set of parameters is also necessary so that we can use covariance from simulations without primordial non-Gaussianity when fitting measurements from simulations with $f_{\text{NL}} = 100$. To determine the fiducial set of coefficients we use an iterative procedure (similar to a Newton-Raphson algorithm) to maximize the likelihood for $\hat{\mathbf{B}}$ and $\hat{\mathbf{M}}_n$, which are the mean measurements from the set of simulations with Gaussian initial conditions.

Specifically, we first maximize the likelihood with a covariance evaluated at $\bar{a}_n = 0$. Then, we use the new coefficients to recompute the covariance, which is used to remaximize the likelihood, and so on. We stop the procedure once all parameters have converged with a fractional accuracy of 10^{-9} , which takes around ten iterations. We have checked that this procedure does not affect our results in an appreciable way.⁷ Once the set of fiducial coefficients has been determined, and thus the covariance, \mathcal{C} , we sample from the parameter posterior distribution. For our analysis, we vary the \bar{a}_n coefficients with uniform wide priors between -10^9 and $10^9 (\text{Mpc}/h)^{3+n}$. The sampling is done using the MULTINEST package [46] with 500 live points and an evidence tolerance of 0.005.

IV. RESULTS

We first present a series of fiducial results carried out on measurements from a single realization of our simulations at two redshifts ($z = 0$ and $z = 0.97$) with soft momenta $0.005 < q < 0.06 h/\text{Mpc}$ and hard momenta $0.2 < k < 0.6 h/\text{Mpc}$. The full posteriors for this analysis can be found in Appendix E. We will later validate this choice of scale cuts, and investigate the corresponding sensitivity of our f_{NL} constraints.

In Fig. 2 we plot the squeezed bispectrum and our maximum *a posteriori* model prediction for a single realization at redshift $z = 0$. The error bars are computed from the fiducial covariance introduced in Sec. III B. As expected, in the presence of local primordial non-Gaussianity, there is a sizable contribution from the $1/q^2$

⁷Our fiducial results for the simulations with Gaussian initial conditions change from $f_{\text{NL}} = -0.6 \pm 12$ to $f_{\text{NL}} = -0.1 \pm 12$ when including the parameter dependent covariance.




FIG. 2. Top panel: measured hard-mode-averaged squeezed bispectrum (red points) and maximum *a posteriori* theory prediction (black lines) for a single realization of the $z = 0$ snapshots without (left) and with (right) local primordial non-Gaussianity assuming $q_{\text{max}} = 23k_f \simeq 0.06h/\text{Mpc}$, $k_{\text{min}} = 75k_f \simeq 0.2h/\text{Mpc}$, and $k_{\text{max}} = 231k_f \simeq 0.6h/\text{Mpc}$. Gray lines indicate the individual contributions of each component of the theory model. In the presence of local primordial non-Gaussianity, there is a sizable contribution from the $1/q^2$ term in the squeezed bispectrum. The error bars are computed from the fiducial covariance described in Sec. III B, scaled to match the volume of a single realization. Bottom panel: fractional residuals with respect to the maximum *a posteriori* prediction.

term in the squeezed bispectrum. Moreover, the \bar{a}_2 contribution becomes significant at scales $q \gtrsim 0.035h/\text{Mpc}$, matching the conclusions of Appendix B. For this particular realization and choice of scale cuts, we find $f_{\text{NL}} = 0 \pm 12$ and $f_{\text{NL}} = 98 \pm 12$, for the measurement without and with local primordial non-Gaussianity, respectively. Our constraints improve considerably at redshift $z = 0.97$ for which we find $f_{\text{NL}} = -1 \pm 8$ and $f_{\text{NL}} = 97 \pm 8$.

The improvement at higher redshift can be understood as follows. On the one hand, the primordial contribution to the squeezed bispectrum, Eq. (7), scales as $B \sim P^2/D(z) \sim D^3(z)$ for growth factor $D(z)$. On the other hand, the bispectrum covariance scales as $C_B \sim P^3 \sim D^6(z)$. It follows that the signal-to-noise ratio has the scaling $B/\sqrt{C_B} \sim D^0(z)$, which is redshift independent. Nonetheless, at higher redshift, gravitational non-Gaussianity is suppressed by a factor of $D(z)$, which (a) reduces the \bar{a}_n signal for $n > 0$ (which must be disentangled from f_{NL}), and (b) reduces the non-Gaussian covariance. These effects increase the signal to noise with which primordial non-Gaussianity can be measured at high redshift. This indicates a transfer of information on f_{NL} into higher-point functions at late times.

A. Squeezed criterion

In this section we fit the squeezed bispectra for different ranges of minimum hard momenta, k_{min} , and maximum soft momenta, q_{max} , to address how squeezed the bispectrum must be in order for consistency relations to be applied reliably. Since our focus in this section is on method validation, all results use the mean bispectrum measurement obtained from 12 realizations with Gaussian initial conditions, and 12 realizations with local primordial non-Gaussianity. The covariance is, however, always evaluated using the 40 Gaussian simulations, but scaled to match the volume of 12 realizations.

In the top panel of Fig. 3 we show the mean value for f_{NL} and its 68% confidence error, for both Gaussian and non-Gaussian initial conditions, as well as for redshifts $z = 0$ and $z = 0.97$. We do this for different values of k_{min} , but fixing $k_{\text{max}} = k_{\text{min}} + \Delta k$, with $\Delta k = 5k_f \simeq 0.013h/\text{Mpc}$. We fix the minimum and maximum soft momenta to be $q_{\text{min}} = k_f \simeq 0.002h/\text{Mpc}$ and $q_{\text{max}} = 23k_f \simeq 0.06h/\text{Mpc}$. For sufficiently high k_{min} , our fit closely matches the expected value of f_{NL} . However, for $k_{\text{min}} \lesssim 0.15h/\text{Mpc}$ we obtain biased results, a signal of the fact that the triangles are not sufficiently squeezed for the relation derived in Eq. (7) to

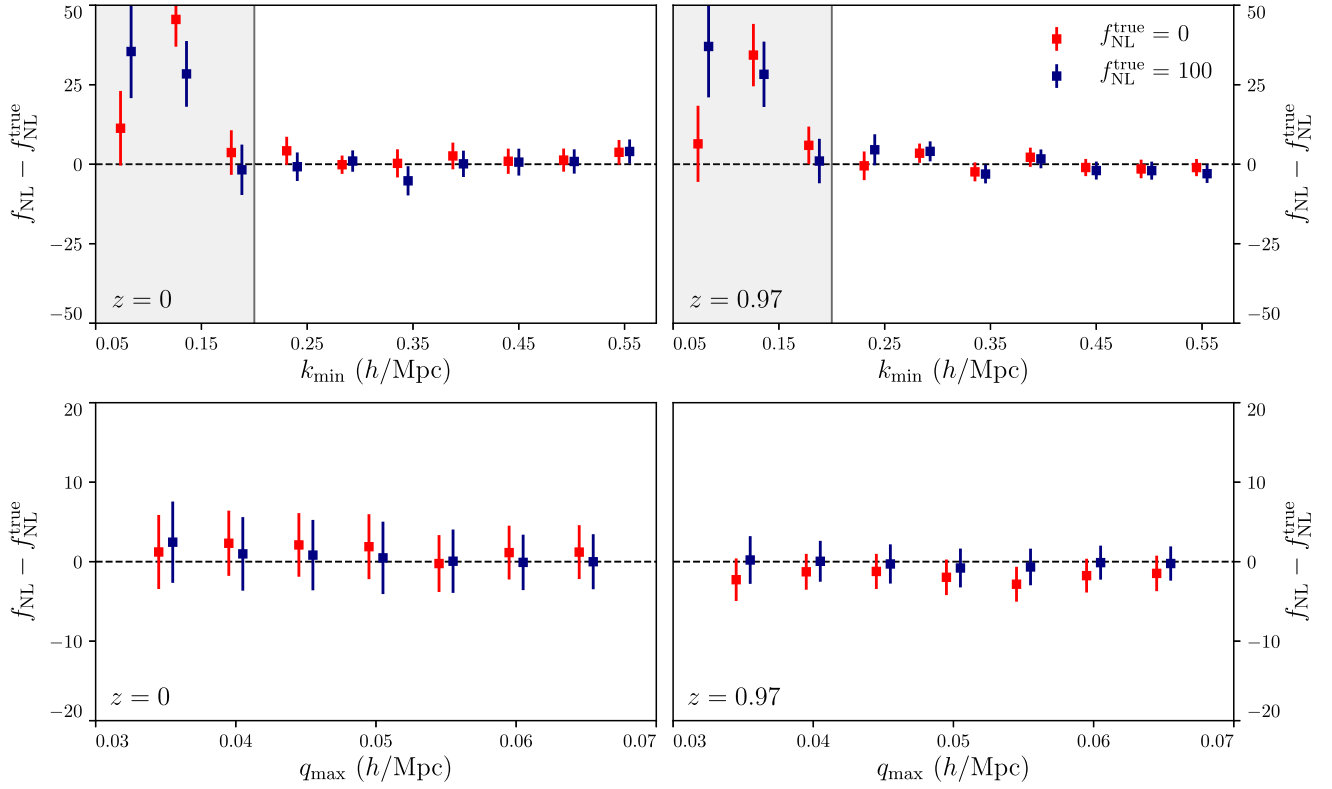


FIG. 3. Top panel: mean value for f_{NL} and its 68% error, for different sets of hard modes at redshift $z = 0$ (left) and $z = 0.97$ (right), and for both Gaussian (red) and non-Gaussian (blue) initial conditions. The maximum soft momentum has been fixed to $q_{\text{max}} = 23k_f \simeq 0.06h/\text{Mpc}$. The hard modes have been collected in a thin bin of size $\Delta k = 2k_f \simeq 0.013h/\text{Mpc}$, with varying k_{min} . The shaded gray region denotes modes with $k_{\text{min}} < 75k_f \simeq 0.2h/\text{Mpc}$, which are excluded in the remainder of our analyses. In this regime, triangles are not sufficiently squeezed, thus our estimates of f_{NL} become biased. Bottom panel: as above, but fixing the hard momentum range to be $k_{\text{min}} = 75k_f \simeq 0.2h/\text{Mpc}$ and $k_{\text{max}} = 231k_f \simeq 0.6h/\text{Mpc}$ and varying the maximum soft momentum, q_{max} . All the error bars refer to a volume corresponding to 12 realizations. These results confirm the expectation that consistency relations are only applicable when q/k is sufficiently small, and that most of the signal is contained in the lowest q modes.

apply. The scale at which we recover unbiased results appears to be consistent between the two redshifts. In light of the above results, unless otherwise stated, for the remainder of the analyses we fix the minimum hard momentum to $k_{\text{min}} \simeq 0.2h/\text{Mpc}$.

Having determined a range of suitable hard momentum scale cuts, we now analyze the dependence of our results on the soft momenta scale cuts. To do this, we fix the maximum hard momentum to $k_{\text{max}} \simeq 0.6h/\text{Mpc}$ and the minimum soft momentum again to $q_{\text{min}} = k_f \simeq 0.002h/\text{Mpc}$, but vary the maximum soft momentum, q_{max} . The scale cuts are such that each successive value of q_{max} contains one more bin than the previous one. We report the results in the bottom panel of Fig. 3. As one can see, we recover the correct value of f_{NL} for both sets of initial conditions, essentially independently of the value chosen for q_{max} . The error on f_{NL} also decreases with increasing maximum soft momentum, albeit mildly. Both observations are explained by the fact that most of the signal is contained in the first few q bins. For this reason, from now on we fix $q_{\text{max}} = 23k_f \simeq 0.06h/\text{Mpc}$.

In summary, our constraints on f_{NL} from the matter bispectrum are unbiased for scale cuts with $k_{\text{min}}/q_{\text{max}} \gtrsim 3$ at redshifts $z = 0$ and $z = 0.97$. Since most of the information is contained within the softest momentum bin, the more relevant scale is likely set by k_{min} and q_{min} . For our analysis, the associated ratio is then $k_{\text{min}}/q_{\text{min}} \approx 30$ (100) for $q_{\text{min}} = 0.007h/\text{Mpc}$ ($0.002h/\text{Mpc}$). Further discussion on the parametrization of the bispectrum and on the scale cuts can be found in Appendix B.

B. Information content and Fisher forecasts

Having demonstrated that the consistency relations can be used to extract unbiased constraints on f_{NL} in the nonlinear regime, we now analyze the information content on f_{NL} in the consistency relations. Of particular interest is how our constraints on f_{NL} change as we vary q_{min} and k_{max} , as well as how our results are limited by naively averaging over all hard momenta with uniform weight. To this end, we fit the squeezed bispectra for varying

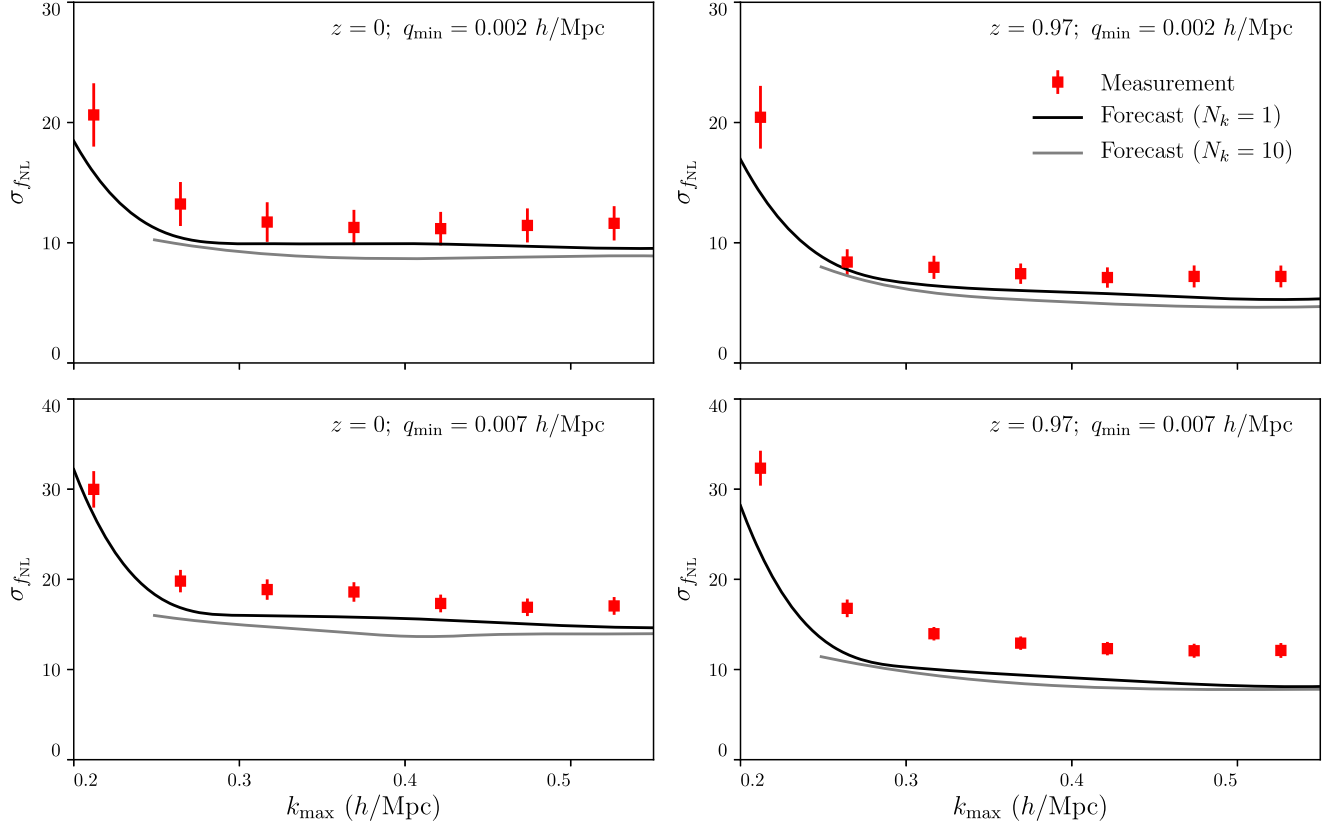


FIG. 4. Comparison of the uncertainty on f_{NL} from our analysis of 40 simulations with a volume of $(2.4 \text{ Gpc}/h)^3$ and $f_{\text{NL}} = 0$ (points) with forecasted errors for the same volume (lines). The error is shown as a function of k_{max} where $k_{\text{min}} = 75k_f \simeq 0.2h/\text{Mpc}$ has been fixed for all analyses. Constraints are shown with (top) and without (bottom) the softest momentum bin for redshift $z = 0$ (left) and $z = 0.97$ (right). The Fisher forecast is performed assuming all modes have been averaged in a single k bin (black), matching our analysis, as well as splitting the data into ten k bins (gray). The improvement from treating the k bins separately is less than 10%, which suggests that there is little to gain from using a more complex hard momentum weighting scheme and/or introducing different values of \bar{a}_0 and \bar{a}_2 for each k bin. The error bars quickly saturate as a function of k_{max} .

maximum hard momenta, k_{max} , and minimum soft momenta, q_{\min} , and additionally perform a series of analytic Fisher forecasts. The derivation of the Fisher matrix used in these forecasts is presented in Appendix C 2.

In Fig. 4 we plot the error bar on f_{NL} as a function of k_{max} for redshift $z = 0$ and $z = 0.97$. We show the mean and the 68% confidence limit for the error on f_{NL} computed by fitting each of the 40 realizations with Gaussian initial conditions individually. All results assume a maximum soft momentum $q_{\text{max}} \simeq 0.06h/\text{Mpc}$ and minimum hard momentum $k_{\min} \simeq 0.2h/\text{Mpc}$. The bottom panels exclude the first soft momenta bin. Removing the softest momenta bin increases the error on f_{NL} from 12 ± 1 (7 ± 1) to 17 ± 1 (12 ± 1) at $k_{\text{max}} = 0.54h/\text{Mpc}$ and $z = 0$ ($z = 0.97$), confirming that most of the information is contained within the softest momenta bin.

Turning our attention to the small-scale information, we find that the error on f_{NL} quickly saturates with increasing k_{max} . This is due to the non-Gaussian contributions to the covariance and is consistent with recent results analyzing local primordial non-Gaussianity in the nonlinear

regime [47]. Thus, even though the consistency relations can be used to constrain f_{NL} at scales well beyond $0.35h/\text{Mpc}$, we find essentially no benefit when including those modes. The situation may be less dire at high redshifts as was found in [48]. In particular, when we increase k_{max} from $121k_f \simeq 0.32h/\text{Mpc}$ to $221k_f \simeq 0.58h/\text{Mpc}$ we find that $\sigma_{f_{\text{NL}}}$ remains essentially constant at $z = 0$ (11.7 ± 1.6 to 11.6 ± 1.4), but decreases by 10% at $z = 0.97$ (8.0 ± 1.0 to 7.2 ± 0.9). Given that the non-Gaussian contributions to the covariance are particularly important for local shapes, even at high redshifts [48,49], the gains at higher redshift may be more modest than those found for other types of non-Gaussianities.

To ensure that the lack of small-scale information is not an artifact of our suboptimal weighting scheme (i.e. summing all k modes with uniform weights), we perform a series of Fisher forecasts. The black curves show the forecasts assuming that all hard momenta are averaged into a single k bin, as is the procedure used in our analysis. We find that our errors are within 20% of the Fisher error for all scale cuts, with the exception of the $z = 0.97$ snapshot with

the first bin excluded, for which our error bars are 50% larger. Note that for simplicity the Fisher forecasts assume $\bar{a}_2 = 0$, which leads to an underestimate of $\sigma_{f_{\text{NL}}}$ due to the degeneracy between \bar{a}_2 and f_{NL} , as discussed in Appendix B.⁸

We also perform the Fisher forecast assuming the hard momenta between k_{min} and k_{max} are subdivided into $N_k = 10$ linearly spaced hard momenta bins. This is equivalent to the improvement that one would expect if we accounted for the different values of \bar{a}_0 over ten bins, as opposed to simply averaging over all modes. Interestingly, there is only a very modest (<10%) improvement when treating the hard momenta bins separately, hence there is likely little to gain by implementing an optimal weighting scheme. We found that the error on f_{NL} does start to decrease if we consider $N_k > 20$; however, given that the covariance is estimated from only 40 realizations, we do not include those results and defer this study to a future work in which we have a more robust covariance.

Finally, using our forecasting pipeline, we present an *extremely idealistic* forecast for the error on local f_{NL} that we can hope to achieve from the Dark Energy Spectroscopic Instrument (DESI) [50] using our method. Assuming DESI will measure a combined survey volume of $170 \text{ Gpc}^3 \simeq 58.4 (\text{Gpc}/h)^3$ (Table 3 of [51]) and neglecting the (likely relevant) complications due to galaxy bias and redshift-space distortions, we find $\sigma_{f_{\text{NL}}} \simeq 4.9$ and 3.3 using the covariance at redshift 0 and 0.97, respectively, and scale cuts $q_{\text{min}} \simeq 0.002$, $q_{\text{max}} \simeq 0.06$, $k_{\text{min}} \simeq 0.2$, $k_{\text{max}} \simeq 0.3h/\text{Mpc}$. Although these forecasts are highly optimistic, combining our analysis with existing techniques to constrain f_{NL} from galaxy clustering, particularly those that use the scale-dependent bias, may considerably improve our ability to constrain local primordial non-Gaussianity.

V. CONCLUSIONS

Due to its ability to differentiate between single-field and multifield inflation, measuring local primordial non-Gaussianity is a key target for current and upcoming cosmological surveys. Here we present a method to measure local f_{NL} from squeezed configurations of the real-space matter bispectrum using the consistency relations. Crucially, our analysis is nonperturbative, allowing us to reliably use data deep into the nonlinear regime. To our knowledge, this is the first result using the consistency relations to directly measure local f_{NL} . Our main results are summarized as follows:

- (i) Using a suite of N -body simulations with local $f_{\text{NL}} = 100$ and with covariance scaled to match a volume of $(2.4 \text{ Mpc}/h)^3$, we measure $f_{\text{NL}} = 98 \pm 12$ at redshift $z = 0$ and $f_{\text{NL}} = 97 \pm 8$ at $z = 0.97$ from the real-space matter bispectrum using the consistency relations.
- (ii) Due to the non-Gaussian bispectrum covariance at late times in the short-scale regime, our error on f_{NL} is relatively insensitive to the choice of maximum hard momenta, k_{max} . In particular, we find little improvement when including modes beyond $k \gtrsim 0.35h/\text{Mpc}$ at redshifts 0 and 0.97, even though our model has been verified up to $k_{\text{max}} \simeq 0.7h/\text{Mpc}$.
- (iii) Constraints on f_{NL} from the consistency relations are extremely sensitive to the amount of large-scale information. The error on f_{NL} increases by 50% if we increase q_{min} from 0.002 to $0.007h/\text{Mpc}$ at redshifts $z = 0$ and $z = 0.97$.

In light of these findings, we conclude that even though the consistency relations provide a robust mechanism to constrain local f_{NL} using measurements in the nonlinear regime, the utility of this information is severely limited by the non-Gaussian covariance at small scales. We expect this effect to be less severe at high redshifts, although a more dedicated analysis is needed to verify this. Even if the consistency relations alone are unable to place competitive constraints on local f_{NL} , they provide a strong test of fundamental physics with minimal modeling assumptions and applying an analysis like the one presented in this work to observational data would be highly informative.

Given that the approach presented here used the real-space matter distribution, it would most easily be generalized to galaxy or CMB lensing surveys (which measure only the matter field), albeit with loss of radial information. Furthermore, recent and upcoming lensing surveys such as the Dark Energy Survey [52], the Legacy Survey of Space and Time [53], Atacama Cosmology Telescope [54], Simons Observatory [55], and CMB-S4 [56] will yield precise measurements of a wealth of nonlinear modes, making them particularly useful for consistency relation-based methods.

On the other hand, since the consistency relations are expected to hold for biased tracers in redshift space [18,22,24–27], one could extend our method to the bispectrum of photometric galaxy surveys. The main ingredient necessary to apply our method to 2D projected galaxy clustering measurements is a treatment of galaxy bias. This may be nontrivial given that the hard momenta in our analysis are in the nonlinear regime. Moreover, the relationship between the measured coefficient \bar{a}_{-2} and f_{NL} [Eq. (9)] will need to be modified to account for the scale-dependent bias and the response of the nonlinear galaxy power spectrum to a long-wavelength density perturbation, both of which could have significant theoretical uncertainties depending on the galaxy population [57].

⁸A more fair comparison between our measurement error and the Fisher error on f_{NL} can be made by fixing $\bar{a}_2 = 0$ and choosing $q_{\text{max}} \leq 0.035h/\text{Mpc}$ as was verified in Appendix C 2. In this case we find that our measured error on f_{NL} is 7% (14%) larger than the Fisher error for $N_k = 1$ ($N_k = 10$) at $z = 0$ assuming $q_{\text{min}} = 0.002h/\text{Mpc}$, $q_{\text{max}} = 0.025h/\text{Mpc}$, $k_{\text{min}} = 76k_f$, and $k_{\text{max}} = 231k_f$.

One could also expand the method presented in this work to 3D galaxy clustering measurements from spectroscopic galaxy surveys, such as BOSS [58], DESI [50], Euclid [59], and SPHEREx [60]. In addition to accounting for galaxy bias, this would require modeling redshift-space distortions and accounting for stochastic effects, i.e. shot noise. Such an analysis would complement and extend existing analyses which measured local f_{NL} from the redshift-space galaxy bispectrum in the perturbative regime [61,62] and reveal if the consistency relations contain information beyond the scale-dependent bias.

The consistency relations relate correlators of all orders, hence there is likely information from the consistency relations that we are not capturing in our analysis. This could be probed by considering higher-order correlators such as the trispectrum, but estimating such quantities and their covariances is both computationally and theoretically challenging. An alternative would be to constrain primordial non-Gaussianity at the field level which, by design, incorporates information from higher-order correlation functions [38,63]. A comparison with field level approaches would be extremely useful to quantify the information content of the consistency relations.

Finally, one could generalize the method presented in this work to constrain other types of non-Gaussianities that violate the consistency relations. The most direct extension would be to consider noninteger poles as are predicted in quasisingle field inflation scenarios [7,30–32]. This would require modifying our bispectrum parametrization to include such poles and their appropriate coefficients.

ACKNOWLEDGMENTS

We thank Kazuyuki Akitsu for sharing code used to compute the bispectrum. We are grateful to Matteo Biagetti, Will Coulton, Mikhail Ivanov, Kendrick Smith, Marko Simonović, and Beatriz Tucci for insightful discussions. We thank the referees for their helpful comments and suggestions. For most of the development of this work A. E. has been a Roger Dashen Member at the Institute for Advanced Study, whose work was also supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics under Award No. DE-SC0009988. O.H.E.P. is a Junior Fellow of the Simons Society of Fellows Junior Fellow and thanks the Institute for Advanced Study for their hospitality and support. S.G. acknowledges support from the Fulbright U.S. Student Program. L.H. acknowledges support by the DOE DE-SC0011941 and a Simons Fellowship in Theoretical Physics. J.C.H. acknowledges support from NSF Grant No. AST-2108536 and NASA Grant No. 21-ATP21-0129. The Flatiron Institute is supported by the Simons Foundation.

APPENDIX A: VERIFICATION OF THE SQUEEZED BISPECTRUM IN PERTURBATION THEORY

In this appendix, we verify the form of the squeezed bispectrum [Eq. (7)] using Eulerian perturbation theory, assuming $q \ll k$, with k in the quasilinear regime. We work in the infinite volume limit for clarity. Since we are interested in the leading $1/q^2$ term, which is not generated by the nonlinear gravitational evolution alone, we focus only on the contribution to the bispectrum coming from primordial non-Gaussianity.

1. Tree level

In the presence of local primordial non-Gaussianity, a nontrivial bispectrum is generated in addition to that from gravitational evolution. In particular, neglecting gravitational corrections, it is given by (e.g. [64])

$$\begin{aligned} B_{111}(\mathbf{q}, \mathbf{k}) &= \langle \delta_q^{(1)} \delta_k^{(1)} \delta_{-q-k}^{(1)} \rangle' \\ &= 2f_{\text{NL}} \frac{\alpha(|\mathbf{q} + \mathbf{k}|)}{\alpha(q)\alpha(k)} P_L(q)P_L(k) + 2 \text{ perms}, \end{aligned} \quad (\text{A1})$$

where P_L is the linear power spectrum, $\delta^{(n)}$ the n th order density field [65], and by $\langle \cdots \rangle'$ we indicate the ensemble-averaged correlator with the overall momentum conserving δ function removed. In the squeezed limit, recalling that $\alpha(q) \sim q^2$ [Eq. (3)], this takes the form

$$B_{111}(\mathbf{q}, \mathbf{k}) = \frac{4f_{\text{NL}}}{\alpha(q)} P_L(q)P_L(k) + \mathcal{O}((q/k)^{-1}). \quad (\text{A2})$$

This matches Eq. (5) with $\partial P(k)/\partial \log \sigma_8^2 \rightarrow P(k)$, which holds true at tree level.

2. One loop

When including nonlinear gravitational evolution, additional contributions are generated in the bispectrum. In particular, the one-loop contributions sourced by primordial non-Gaussianity are

$$\begin{aligned} B_{1\text{-loop}}(\mathbf{q}, \mathbf{k}) &= \langle \delta_q^{(3)} \delta_k^{(1)} \delta_{-q-k}^{(1)} \rangle' + 2 \text{ perms} \\ &\quad + \langle \delta_q^{(2)} \delta_k^{(2)} \delta_{-q-k}^{(1)} \rangle' + 2 \text{ perms} \\ &\quad + \langle \delta_q^{(\text{ct})} \delta_k^{(1)} \delta_{-q-k}^{(1)} \rangle' + 2 \text{ perms}, \end{aligned} \quad (\text{A3})$$

where the final term is an ultraviolet counterterm required by renormalization, with $\delta_q^{(\text{ct})} \equiv -c_s^2 q^2 \delta_q^{(1)}$, with c_s a constant corresponding to the speed of sound in the hydrodynamical description of matter evolution. Each of these terms involve five linear density fields, which, after

Wick contractions, involve the linear power spectrum, $\langle \delta^{(1)} \delta^{(1)} \rangle$, and the primordial bispectrum, $\langle \delta^{(1)} \delta^{(1)} \delta^{(1)} \rangle$ (except for the counterterm, which involves a primordial bispectrum and a factor of q^2). The squeezed bispectrum is

dominated by terms involving primordial bispectra of the form $B_{111}(\mathbf{q}, \mathbf{k}) \sim P_L(q)P_L(k)/\alpha(q)$, for some \mathbf{k} satisfying $q \ll k$; any other terms are suppressed by powers of q^2 . There are seven such contributions:

$$\begin{aligned} B_{1\text{-loop}}(\mathbf{q}, \mathbf{k}) \supset & 4 \int_p F_2(\mathbf{p}, \mathbf{k} - \mathbf{p}) F_2(\mathbf{p} + \mathbf{q}, \mathbf{k} - \mathbf{p}) B_{111}(\mathbf{q}, \mathbf{p}) P_L(|\mathbf{k} - \mathbf{p}|) + 3 B_{111}(\mathbf{q}, \mathbf{k}) \int_p F_3(\mathbf{k}, \mathbf{p}, -\mathbf{p}) P_L(p) \\ & + 3 B_{111}(\mathbf{q}, \mathbf{k}) \int_p F_3(-\mathbf{k} - \mathbf{q}, \mathbf{p}, -\mathbf{p}) P_L(p) + 3 P_L(|\mathbf{q} + \mathbf{k}|) \int_p F_3(\mathbf{k} + \mathbf{q}, \mathbf{p}, -\mathbf{p} - \mathbf{q}) B_{111}(\mathbf{q}, \mathbf{p}) \\ & + 3 P_L(k) \int_p F_3(-\mathbf{k}, \mathbf{p}, -\mathbf{p} - \mathbf{q}) B_{111}(\mathbf{q}, \mathbf{p}) - c_s^2 k^2 B_{111}(\mathbf{q}, \mathbf{k}) - c_s^2 |\mathbf{q} + \mathbf{k}|^2 B_{111}(\mathbf{q}, -\mathbf{q} - \mathbf{k}), \end{aligned} \quad (\text{A4})$$

utilizing momentum conservation, and denoting $\int_p \equiv \int d^3\mathbf{p}/(2\pi)^3$. The F_n functions are the standard perturbative kernels [65]. Terms three, five, and seven are simply terms two, four, and six with the hard momenta switched.

Taking the soft limit, $q \ll k, p$, and inserting Eq. (A2), these simplify to

$$\begin{aligned} B_{1\text{-loop}}(\mathbf{q}, \mathbf{k}) = & 4 f_{\text{NL}} \frac{P_L(q)}{\alpha(q)} \left[4 \int_p |F_2(\mathbf{p}, \mathbf{k} - \mathbf{p})|^2 P_L(p) P_L(|\mathbf{k} - \mathbf{p}|) + 12 P_L(k) \int_p F_3(\mathbf{k}, \mathbf{p}, -\mathbf{p}) P_L(p) \right. \\ & \left. - 2 c_s^2 k^2 P_L(k) \right] + \mathcal{O}((q/k)^{-1}). \end{aligned} \quad (\text{A5})$$

We recognize the term in parentheses as twice the one-loop power spectrum, $P_{1\text{-loop}}(k)$ [66,67]. In combination with Eq. (A2), we find the full f_{NL} contribution to the squeezed bispectrum up to one-loop order:

$$B(\mathbf{q}, \mathbf{k}) = \frac{4 f_{\text{NL}}}{\alpha(q)} P_L(q) [P_L(k) + 2 P_{1\text{-loop}}(k) + \dots] + \mathcal{O}((q/k)^{-1}), \quad (\text{A6})$$

which agrees with Eq. (7), noting that $P_L \propto \sigma_8^2$ and $P_{1\text{-loop}} \propto \sigma_8^4$. It is worth stressing that there is a factor of 2 multiplying the one-loop contribution to the power spectrum. Thus what fits inside the square brackets is *not* simply the nonlinear power spectrum.

APPENDIX B: VALIDATION OF THE PARAMETRIZATION

In this appendix, we perform a series of tests validating our fiducial analysis choice in which we vary f_{NL} , \bar{a}_0 , and \bar{a}_2 . In particular, we show that the odd coefficients in the power series expansion of the squeezed bispectrum are consistent with zero, that including the quadratic coefficient \bar{a}_2 is necessary to model the squeezed bispectrum at the scale cuts used in this work, and that coefficients higher than quadratic order can be neglected in our model. Unless otherwise stated, all results in this section use the mean measurements at redshift $z = 0$ from 12 realizations with covariance scaled to match the volume of 12 realizations assuming a maximum soft momentum $q_{\text{max}} = 0.06h/\text{Mpc}$ and hard momenta between 0.2 and $0.6h/\text{Mpc}$.

To show that \bar{a}_{-1} and \bar{a}_1 can be set to zero (matching parity considerations), we rerun our analysis by including either \bar{a}_{-1} or \bar{a}_1 as a free parameters.⁹ In Fig. 5 we plot the 2D marginalized posteriors of $f_{\text{NL}} - f_{\text{NL}}^{\text{true}}$ and the odd order coefficients. We include results with Gaussian initial conditions, as well as those assuming $f_{\text{NL}} = 100$. The posteriors on \bar{a}_{-1} and \bar{a}_1 are consistent with zero which suggests that our angular averaging is working as expected. Furthermore, including the odd coefficients as free parameters in our model does not bias the constraints on f_{NL} . We note, however, that the error on f_{NL} increases significantly when we vary \bar{a}_{-1} , suggesting that our measurements are unable to fully differentiate between a $1/q$ and $1/q^2$ pole in the squeezed bispectrum. On account of these tests, we conclude that we can set the odd parameters in our model to zero.

In Fig. 6 we instead plot the 68% confidence limits on f_{NL} as a function of the maximum soft momentum,

⁹We have also varied \bar{a}_{-1} and \bar{a}_1 simultaneously and our results remain unchanged; however, the constraints degrade significantly when letting all parameters vary, thus we focus on the results for which we vary only one coefficient at a time.




FIG. 5. Marginalized posteriors when including either \bar{a}_{-1} (left) or \bar{a}_1 (right) as free parameters in the model. The odd coefficients are consistent with zero as is expected from our angular averaging procedure. There is significant degeneracy between \bar{a}_{-1} and f_{NL} due to the difficulty in distinguishing between a $1/q$ and $1/q^2$ pole in our measurements.

q_{max} , assuming $\bar{a}_2 = 0$. As one can see, for $q_{\text{max}} \gtrsim 0.035h/\text{Mpc}$, this model gives biased constraints on f_{NL} . Indeed, at these scales the $\mathcal{O}((q/k)^2)$ correction to the squeezed bispectrum becomes comparable to the measurement error. Hence, we need to include higher-order terms if we wish to include these modes in our analysis without increasing k_{min} . Including \bar{a}_2 as a free parameter provides unbiased constraints on f_{NL} up to scales of at least $q_{\text{max}} \simeq 0.065h/\text{Mpc}$ as shown in Fig. 3. In practice one could

either include \bar{a}_2 as a free parameter, or impose a stricter soft momentum scale cut and fix $\bar{a}_2 = 0$. In our fiducial analysis we choose to vary \bar{a}_2 , since this leads to an error on f_{NL} that is approximately 20% smaller than what we find when fixing $\bar{a}_2 = 0$ and decreasing q_{max} . Alternatively, one could extend the validity of the two-parameter model to include more soft momenta bins by increasing k_{min} ; however, since the maximum accessible value of k is often limited both in observations and simulations




FIG. 6. Constraints on f_{NL} for varying soft momenta scale cuts assuming $\bar{a}_2 = 0$ and averaging over hard momenta between 0.2 and $0.6h/\text{Mpc}$. Results are shown for $z = 0$ for the simulations with Gaussian initial conditions (red), as well as those with primordial non-Gaussianity (blue). The horizontal offset between the two simulations is purely for visual purposes. We are unable to recover the true value of f_{NL} if we include soft momenta bins larger than $0.035h/\text{Mpc}$; therefore, we include \bar{a}_2 as a free parameter in our fiducial analysis.




FIG. 7. Constraints on f_{NL} and \bar{a}_4 when including \bar{a}_4 as a free parameter in our model. \bar{a}_4 is consistent with zero, suggesting that it is sufficient to only consider terms up to q^2 in our fiducial analysis.

(e.g. by shot noise or resolution), we do not utilize this method.

Finally, we rerun our analysis including terms up to quartic order in q by varying f_{NL} , \bar{a}_0 , \bar{a}_2 , and \bar{a}_4 . In Fig. 7 we show the 2D marginalized posteriors of $f_{\text{NL}} - f_{\text{NL}}^{\text{true}}$ and \bar{a}_4 . We find that \bar{a}_4 is completely consistent with zero and that we continue to recover unbiased constraints on f_{NL} . Including \bar{a}_4 as a free parameter increases the error on f_{NL} by 50%, thus we fix \bar{a}_4 and all higher-order coefficients to zero in our fiducial analysis.

APPENDIX C: LIKELIHOOD AND FISHER MATRIX

In this appendix we derive the likelihood and Fisher matrix used in our analysis. We first start on general grounds, and then introduce simplifying approximations to ease the computational burden.

Consider the angular averaged bispectrum, B_{iJ} , binned in both soft and hard momenta bins. Here i, j, \dots label the soft modes and I, J, \dots the hard ones. Consistency relations

imply that, in the squeezed limit, the angle-averaged bispectrum is a stochastic variable that fluctuates around its true value. For sufficiently thin bins, the latter is given by $\bar{B}_{iI} = s_{iI} \bar{P}_i$, where \bar{P}_i is the true (soft) power spectrum, and s_{iI} is defined as

$$s_{iI} \equiv f_{\text{NL}} P_{iI} + \sum_{n \geq 0} a_{n,I} q_i^n, \quad (\text{C1})$$

with

$$P_{iI} \equiv \frac{6\Omega_{m,0} H_0^2}{D_{\text{md}}(z) T_i q_i^2} \frac{\partial \bar{P}_I}{\partial \log \sigma_8^2}. \quad (\text{C2})$$

Similarly, the power spectrum P_i is a stochastic variable fluctuating around \bar{P}_i .

Assuming that both B_{iJ} and P_i are Gaussian distributed, we interpret their probability distribution as the likelihood for the unknown parameters $(f_{\text{NL}}, \mathbf{a}_n, \bar{P}_i)$, which reads

$$\mathcal{L}(f_{\text{NL}}, \mathbf{a}_n, \bar{P}_i) \propto \exp \left\{ -\frac{1}{2} \sum_i [(B_i - s_i \bar{P}_i) \cdot \Psi_i^{BB} \cdot (B_i - s_i \bar{P}_i) + 2(B_i - s_i \bar{P}_i) \cdot \Psi_i^{BP} (P_i - \bar{P}_i) + \Psi_i^{PP} (P_i - \bar{P}_i)^2] \right\}, \quad (\text{C3})$$

where, contrary to the main text, we explicitly let the vectors run over the hard mode indices, I, J . We have also used the fact that the covariance matrices are approximately diagonal in the soft mode, as verified in our simulations. Using the standard identities for the inverse of a block matrix, the Ψ 's can be related to the covariance matrices by

$$\Psi_i^{BB} = \left(\mathcal{C}_i^{BB} - \frac{\mathcal{C}_i^{BP} \otimes \mathcal{C}_i^{BP}}{\mathcal{C}_i^{PP}} \right)^{-1}, \quad (\text{C4a})$$

$$\Psi_i^{BP} = -\frac{(\mathcal{C}_i^{BB})^{-1} \cdot \mathcal{C}_i^{BP}}{\mathcal{C}_i^{PP} - \mathcal{C}_i^{BP} \cdot (\mathcal{C}_i^{BB})^{-1} \cdot \mathcal{C}_i^{BP}}, \quad (\text{C4b})$$

$$\Psi_i^{PP} = \frac{1}{\mathcal{C}_i^{PP} - \mathcal{C}_i^{BP} \cdot (\mathcal{C}_i^{BB})^{-1} \cdot \mathcal{C}_i^{BP}}, \quad (\text{C4c})$$

where $\mathcal{C}_{iIJ}^{BB} \equiv \langle \delta B_{iI} \delta B_{iJ} \rangle$, $\mathcal{C}_{iI}^{BP} \equiv \langle \delta B_{iI} \delta P_i \rangle$, $\mathcal{C}_i^{PP} \equiv \langle \delta P_i \delta P_i \rangle$, and $\langle \dots \rangle$ denotes an ensemble average, which will be estimated from simulations. In particular, in the space of the hard momenta bins, Ψ_i^{BB} is a matrix, Ψ_i^{BP} a vector, and Ψ_i^{PP} a scalar.

1. The likelihood for the analysis

As described in Sec. III A, in our main analysis we average over all hard modes collected in a single k bin. The

corresponding likelihood should then be obtained from Eq. (C3) when the vector index is trivial, e.g. $\mathbf{a}_n \rightarrow \bar{a}_n$. Since we do not fit for the true value of the power spectrum, we also marginalize over \bar{P}_i , obtaining

$$\mathcal{L}(f_{\text{NL}}, \bar{a}_n) \propto \prod_i \frac{1}{\sqrt{\mathcal{C}_i}} \exp \left[-\frac{1}{2} \frac{(B_i - s_i P_i)^2}{\mathcal{C}_i} \right], \quad (\text{C5})$$

with $\mathcal{C}_i \equiv \mathcal{C}_i^{BB} - 2s_i \mathcal{C}_i^{BP} + s_i^2 \mathcal{C}_i^{PP}$. Note that the likelihood above is simply a Gaussian one, but with a parameter dependent covariance.

In the analysis described in Sec. III B we start from the likelihood above, but introduce two modifications for practical purposes. First, since we are estimating the covariance matrix from simulations, we promote it to a nondiagonal one. We have checked that, even assuming a diagonal covariance in our analysis does not change the results appreciably. Second, in principle the only independent stochastic quantity in the theory model is the soft power spectrum, P_i . However, when accounting for finite size bins for which $\langle q^n \delta_q \delta_{-q} \rangle \neq q^n \langle \delta_q \delta_{-q} \rangle$, the different terms in s_i will weight P_i by different powers of q_i , with effects that are hard to predict. We therefore find it more convenient to treat the M_{in} —defined in Eq. (11)—as different stochastic variables. In summary, we employ the likelihood Eq. (C5), but promoting the covariance matrix to

$$\mathcal{C}_i \delta_{ij}^K \rightarrow \left\langle \left(B_i - \sum_n \bar{a}_n M_{in} \right) \left(B_j - \sum_n \bar{a}_n M_{jn} \right) \right\rangle, \quad (\text{C6})$$

which, for a finite number of realizations, results in Eq. (15).

2. The Fisher matrix

To determine the Fisher matrix we work with the original likelihood, Eq. (C3). To reduce the computational burden and reliably estimate the Fisher matrix from 40 mocks, we assume $\mathbf{a}_2 = \mathbf{0}$ and hence the model parameters are $(f_{\text{NL}}, \mathbf{a}_0, \bar{P}_i)$. The Fisher matrix is then given by the Hessian of $-\log \mathcal{L}$, and can be written in a block form as

$$\mathbf{F} = \begin{pmatrix} \mathbf{F}_1 & \mathbf{F}_2 \\ \mathbf{F}_2^T & \mathbf{F}_3 \end{pmatrix}. \quad (\text{C7})$$

If we have N_k hard mode bins, then \mathbf{F}_1 is a $(N_k + 1) \times (N_k + 1)$ matrix with entries,

$$F_1 \equiv \sum_i \bar{P}_i^2 \begin{pmatrix} \mathbf{p}_i \cdot \boldsymbol{\Psi}_i^{BB} \cdot \mathbf{p}_i & \mathbf{p}_i \cdot \boldsymbol{\Psi}_i^{BB} \\ \boldsymbol{\Psi}_i^{BB} \cdot \mathbf{p}_i & \boldsymbol{\Psi}_i^{BB} \end{pmatrix}. \quad (\text{C8})$$

If we have N_q soft bins, \mathbf{F}_2 is a $(N_k + 1) \times N_q$ matrix with entries,

$$(\mathbf{F}_2)_i \equiv \bar{P}_i \begin{pmatrix} \mathbf{p}_i \cdot \boldsymbol{\Psi}_i^{BB} \cdot \mathbf{s}_i + \mathbf{p}_i \cdot \boldsymbol{\Psi}_i^{BP} \\ \boldsymbol{\Psi}_i^{BB} \cdot \mathbf{s}_i + \boldsymbol{\Psi}_i^{BP} \end{pmatrix}. \quad (\text{C9})$$

Finally, \mathbf{F}_3 is a diagonal $N_q \times N_q$ matrix,

$$(\mathbf{F}_3)_{ij} = (\mathbf{s}_i \cdot \boldsymbol{\Psi}_i^{BB} \cdot \mathbf{s}_i + 2\mathbf{s}_i \cdot \boldsymbol{\Psi}_i^{BP} + \boldsymbol{\Psi}_i^{PP}) \delta_{ij}^K. \quad (\text{C10})$$

By the Cramér-Rao theorem, the variance on a measurement of f_{NL} will then satisfy the bound

$$\sigma_{f_{\text{NL}}}^2 \geq (\mathbf{F}^{-1})_{11}, \quad (\text{C11})$$

which is what we use in our analysis of Sec. IV B. For the forecasts in Sec. IV B we assume $f_{\text{NL}} = 0$ and that \bar{P}_i is equal the mean power spectrum as measured from the 40 realizations with Gaussian initial conditions. To determine a set of fiducial values for \mathbf{a}_0 we must account for its k dependence. To do this, we use the consistency relation from [41],

$$a_0(k) = \left[1 + \frac{13}{21} \frac{\partial}{\partial \ln D_+} - \frac{1}{3} \frac{\partial}{\partial \ln k} \right] P(k), \quad (\text{C12})$$

where D_+ is the linear growing mode. In particular, we estimate \mathbf{a}_0 by computing Eq. (C12) using HALOFIT. This prediction agrees within 5% of our measured value of \bar{a}_0 over all redshifts and scale cuts analyzed in this work, and is therefore sufficient for the forecasts.

APPENDIX D: OPTIMAL WEIGHTING SCHEMES

In this appendix, we consider possible optimal weighting schemes, $W(\mathbf{q}, \mathbf{k})$, for the compressed bispectrum, $B(\mathbf{q}) \equiv \int_{\mathbf{k}} W(\mathbf{q}, \mathbf{k}) B(\mathbf{q}, \mathbf{k})$. Note that in our main analysis we always fix $W(\mathbf{q}, \mathbf{k}) = 1$ as we find that the weights derived in this section do not noticeably improve our constraints on f_{NL} ; nevertheless, we present our main findings below since they could be relevant for future works measuring f_{NL} from consistency relations. To derive the weights, we will compute the maximum-likelihood estimator for the \bar{a}_n coefficients appearing in the bispectrum model of Eq. (13) for a given q bin using a simplified treatment of our likelihood and covariance.

For simplicity, we assume a Gaussian likelihood for the bispectrum and work in the continuous limit of \mathbf{q} and \mathbf{k} . The likelihood of the model parameters \bar{a}_n is then, up to a constant term,

$$\log \mathcal{L}(\bar{a}_n) \propto - \int_{\mathbf{q}\mathbf{q}'\mathbf{k}\mathbf{k}'} (\delta B(\bar{a}_n; \mathbf{q}, \mathbf{k}) \times \mathcal{C}_B^{-1}(\mathbf{q}, \mathbf{k}; \mathbf{q}', \mathbf{k}') \times \delta B(\bar{a}_n; \mathbf{q}', \mathbf{k}')), \quad (\text{D1})$$

where $\delta B(\bar{a}_n; \mathbf{q}, \mathbf{k}) \equiv B(\mathbf{q}, \mathbf{k}) - \sum_n \bar{a}_n \hat{M}_n(\mathbf{q})$ with $\hat{M}_n(\mathbf{q})$ defined via Eq. (11), \mathcal{C}_B is the squeezed bispectrum covariance, and the integral is carried out over all soft and hard modes allowed for a given choice of scale cuts. In the squeezed limit, the bispectrum covariance is approximately given by [49]

$$\begin{aligned} \mathcal{C}_B(\mathbf{q}, \mathbf{k}; \mathbf{k}', \mathbf{q}') &\simeq (2\pi)^3 \delta_D(\mathbf{q} - \mathbf{q}') \\ &\times [P(\mathbf{q}) \mathcal{C}_P(\mathbf{k}, \mathbf{k}') + 2B(\mathbf{q}', \mathbf{k}') B(\mathbf{q}, \mathbf{k})], \end{aligned} \quad (\text{D2})$$

where \mathcal{C}_P is the (nonlinear) power spectrum covariance.

Under null assumptions, i.e. $\bar{a}_n = 0$ for all n , the bispectrum vanishes and only the first term survives, and the likelihood for a single \mathbf{q} mode becomes

$$\log \mathcal{L}(\bar{a}_n, \mathbf{q}) \propto - \frac{1}{P(\mathbf{q})} \int_{\mathbf{k}\mathbf{k}'} (\delta B(\bar{a}_n; \mathbf{q}, \mathbf{k}) \times \mathcal{C}_P^{-1}(\mathbf{k}, \mathbf{k}') \times \delta B(\bar{a}_n; \mathbf{q}', \mathbf{k}')), \quad (\text{D3})$$

up to a constant term. Given Eq. (D3), we may derive an optimal estimator for \bar{a}_n (for a single \mathbf{q} mode) by maximizing $\log \mathcal{L}(\bar{a}_n, \mathbf{q})$ with respect to \bar{a}_n . This yields

$$\hat{\bar{a}}_n = \sum_m \mathcal{K}_{nm}^{-1} \int_{\mathbf{k}\mathbf{k}'} B(\mathbf{q}, \mathbf{k}) \mathcal{C}_P^{-1}(\mathbf{k}, \mathbf{k}') \hat{M}_m(\mathbf{q}), \quad (\text{D4})$$

with the coupling matrix

$$\mathcal{K}_{mn} = \int_{\mathbf{k}\mathbf{k}'} \hat{M}_m(\mathbf{q}) \mathcal{C}_P^{-1}(\mathbf{k}, \mathbf{k}') \hat{M}_n(\mathbf{q}). \quad (\text{D5})$$

The optimal weight is thus $W_n(\mathbf{q}, \mathbf{k}) \propto \int_{\mathbf{k}'} \mathcal{C}_P^{-1}(\mathbf{k}, \mathbf{k}') \hat{M}_n(\mathbf{q})$, with one weight per choice of n .

It is useful to consider two simplifying limits. First, if the covariance is dominated by its Gaussian contribution, $VC_P(\mathbf{k}, \mathbf{k}') = 2(2\pi)^3 \delta_D(\mathbf{k} - \mathbf{k}') P^2(k)$, yielding

$$\begin{aligned} \hat{a}_n &\rightarrow \sum_m \mathcal{K}_{nm}^{-1} \int_{\mathbf{k}} \frac{B(\mathbf{q}, \mathbf{k}) \hat{M}_m(\mathbf{q})}{P^2(k)}, \\ \mathcal{K}_{mn} &\rightarrow \int_{\mathbf{k}} \frac{\hat{M}_m(\mathbf{q}) \hat{M}_n(\mathbf{q})}{P^2(k)}, \end{aligned} \quad (\text{D6})$$

this allows us to remove the \mathbf{k}' integral. Second, if the coupling between \bar{a}_n coefficients is weak (i.e. \mathcal{K}_{mn} is close to diagonal), we may apply the optimal weighting only for the $m = -2$ coefficient. Inserting the relevant template [Eq. (9)], we find the quasioptimal weighting $W(\mathbf{q}, \mathbf{k}) = \partial \log P(k) / \partial \log \sigma_8^2 / P(k)$ (dropping \mathbf{k} -independent factors), and thus the compressed bispectrum

$$B(\mathbf{q}) \equiv \int_{\mathbf{k}} \frac{B(\mathbf{q}, \mathbf{k})}{P(k)} \frac{\partial \log P(k)}{\partial \log \sigma_8^2}. \quad (\text{D7})$$

Strictly speaking, this weighting could be far from optimal in the high- k limit, whereupon the power spectrum covariance is far from diagonal. In this case, invoking only the second assumption, one might use

$$B(\mathbf{q}) \equiv \int_{\mathbf{k}} B(\mathbf{q}, \mathbf{k}) \int_{\mathbf{k}'} \mathcal{C}_P^{-1}(\mathbf{k}, \mathbf{k}') \frac{\partial P(k')}{\partial \log \sigma_8^2}. \quad (\text{D8})$$

This may be of greater use in practice, and is straightforward to implement if the weighting is estimated from simulations, e.g. QUIJOTE. Finally, one may wish to retain the bispectra appearing in Eq. (D2), assuming nonzero fiducial values of \bar{a}_n ; this can be done similarly, though the covariance will become a function also of \mathbf{q} .

APPENDIX E: FULL POSTERIOR DISTRIBUTION OF FIDUCIAL ANALYSIS

In Fig. 8 we show the full marginalized posteriors of our fiducial analysis at $z = 0$ and $z = 0.97$. As seen in the text, we recover the true value of f_{NL} for both redshifts and both sets of initial conditions. There are noticeable degeneracies between all parameters, with a particularly strong correlation between \bar{a}_0 and \bar{a}_2 . We also find that \bar{a}_0 is larger for simulations with $f_{\text{NL}} = 100$. This can be understood by noting that $\bar{a}_0 \sim P(k)$ and a positive value of f_{NL} enhances the nonlinear power spectrum [68].

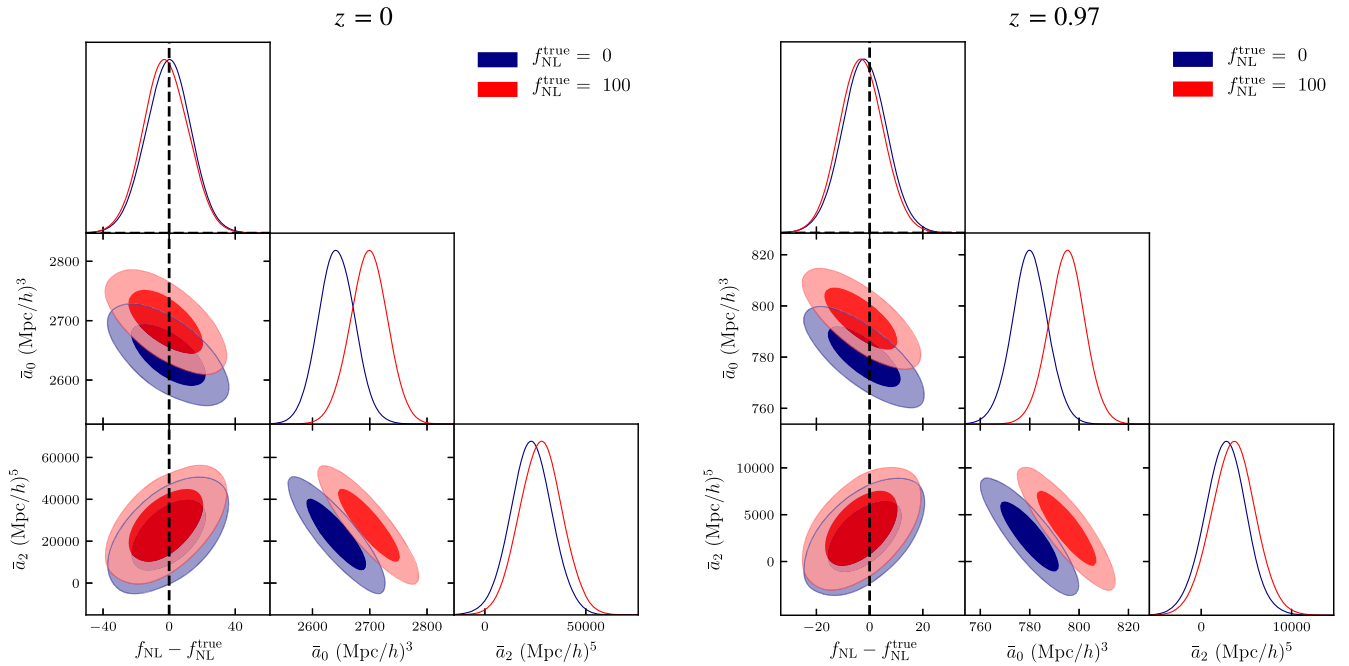


FIG. 8. Full marginalized posteriors of the fiducial results presented in Sec. IV at $z = 0$ (left) and $z = 0.97$ (right). We recover the true value of f_{NL} in all cases with considerable improvement at $z = 0.97$.

- [1] J. M. Maldacena, *J. High Energy Phys.* **05** (2003) 013.
- [2] P. Creminelli and M. Zaldarriaga, *J. Cosmol. Astropart. Phys.* **10** (2004) 006.
- [3] C. Cheung, A. L. Fitzpatrick, J. Kaplan, and L. Senatore, *J. Cosmol. Astropart. Phys.* **02** (2008) 021.
- [4] T. Tanaka and Y. Urakawa, *J. Cosmol. Astropart. Phys.* **05** (2011) 014.
- [5] P. Creminelli, J. Noreña, and M. Simonović, *J. Cosmol. Astropart. Phys.* **07** (2012) 052.
- [6] K. Hinterbichler, L. Hui, and J. Khoury, *J. Cosmol. Astropart. Phys.* **08** (2012) 017.
- [7] V. Assassi, D. Baumann, and D. Green, *J. Cosmol. Astropart. Phys.* **11** (2012) 047.
- [8] A. Kehagias and A. Riotto, *Nucl. Phys.* **B864**, 492 (2012).
- [9] E. Pajer, F. Schmidt, and M. Zaldarriaga, *Phys. Rev. D* **88**, 083502 (2013).
- [10] K. Hinterbichler, L. Hui, and J. Khoury, *J. Cosmol. Astropart. Phys.* **01** (2014) 039.
- [11] W. D. Goldberger, L. Hui, and A. Nicolis, *Phys. Rev. D* **87**, 103520 (2013).
- [12] T. Baldauf, M. Mirbabayi, M. Simonović, and M. Zaldarriaga, *Phys. Rev. D* **92**, 043514 (2015).
- [13] R. Bravo, S. Mooij, G. A. Palma, and B. Pradenas, *J. Cosmol. Astropart. Phys.* **05** (2018) 025.
- [14] L. Hui, A. Joyce, and S. S. C. Wong, *J. Cosmol. Astropart. Phys.* **02** (2019) 060.
- [15] A. Kehagias and A. Riotto, *Nucl. Phys.* **B873**, 514 (2013).
- [16] M. Peloso and M. Pietroni, *J. Cosmol. Astropart. Phys.* **05** (2013) 031.
- [17] P. Valageas, *Phys. Rev. D* **89**, 123522 (2014).
- [18] P. Creminelli, J. Gleyzes, M. Simonović, and F. Vernizzi, *J. Cosmol. Astropart. Phys.* **02** (2014) 051.
- [19] P. Valageas, A. Taruya, and T. Nishimichi, *Phys. Rev. D* **95**, 023504 (2017).
- [20] M. Marinucci, T. Nishimichi, and M. Pietroni, *Phys. Rev. D* **100**, 123537 (2019).
- [21] M. Marinucci, T. Nishimichi, and M. Pietroni, *J. Cosmol. Astropart. Phys.* **07** (2020) 054.
- [22] P. Creminelli, J. Noreña, M. Simonović, and F. Vernizzi, *J. Cosmol. Astropart. Phys.* **12** (2013) 025.
- [23] A. Esposito, L. Hui, and R. Scoccimarro, *Phys. Rev. D* **100**, 043536 (2019).
- [24] A. Kehagias, J. Noreña, H. Perrier, and A. Riotto, *Nucl. Phys.* **B883**, 83 (2014).
- [25] M. Simonović, Cosmological consistency relations, Ph.D. thesis, SISSA, Trieste, 2014.
- [26] A. Kehagias, A. Moradinezhad Dizgah, J. Noreña, H. Perrier, and A. Riotto, *J. Cosmol. Astropart. Phys.* **08** (2015) 018.
- [27] P. Creminelli, J. Gleyzes, L. Hui, M. Simonović, and F. Vernizzi, *J. Cosmol. Astropart. Phys.* **06** (2014) 009.
- [28] D. H. Lyth and D. Wands, *Phys. Lett. B* **524**, 5 (2002).
- [29] M. Zaldarriaga, *Phys. Rev. D* **69**, 043508 (2004).
- [30] X. Chen and Y. Wang, *J. Cosmol. Astropart. Phys.* **04** (2010) 027.
- [31] D. Baumann and D. Green, *Phys. Rev. D* **85**, 103520 (2012).
- [32] T. Noumi, M. Yamaguchi, and D. Yokoyama, *J. High Energy Phys.* **06** (2013) 051.
- [33] V. Desjacques, D. Jeong, and F. Schmidt, *Phys. Rep.* **733**, 1 (2018).
- [34] E. Komatsu and D. N. Spergel, *Phys. Rev. D* **63**, 063002 (2001).
- [35] R. Scoccimarro, L. Hui, M. Manera, and K. C. Chan, *Phys. Rev. D* **85**, 083002 (2012).
- [36] A. Slosar, C. Hirata, U. Seljak, S. Ho, and N. Padmanabhan, *J. Cosmol. Astropart. Phys.* **08** (2008) 031.
- [37] C.-T. Chiang, *Phys. Rev. D* **95**, 123517 (2017).
- [38] U. Giri, M. Münchmeyer, and K. M. Smith, *arXiv*: 2205.12964.
- [39] T. Baldauf, U. Seljak, and L. Senatore, *J. Cosmol. Astropart. Phys.* **04** (2011) 006.
- [40] B. Horn, L. Hui, and X. Xiao, *J. Cosmol. Astropart. Phys.* **09** (2014) 044.
- [41] T. Nishimichi and P. Valageas, *Phys. Rev. D* **90**, 023546 (2014).
- [42] A. Lewis, *J. Cosmol. Astropart. Phys.* **10** (2011) 026.
- [43] R. Takahashi, M. Sato, T. Nishimichi, A. Taruya, and M. Oguri, *Astrophys. J.* **761**, 152 (2012).
- [44] F. Villaescusa-Navarro *et al.*, *Astrophys. J. Suppl. Ser.* **250**, 2 (2020).
- [45] E. Sellentin and A. F. Heavens, *Mon. Not. R. Astron. Soc.* **456**, L132 (2016).
- [46] F. Feroz, M. P. Hobson, and M. Bridges, *Mon. Not. R. Astron. Soc.* **398**, 1601 (2009).
- [47] W. R. Coulton, F. Villaescusa-Navarro, D. Jamieson, M. Baldi, G. Jung, D. Karagiannis, M. Liguori, L. Verde, and B. D. Wandelt, *arXiv*:2206.15450.
- [48] T. Flöss, M. Biagetti, and P. D. Meerburg, *arXiv*: 2206.10458.
- [49] M. Biagetti, L. Castiblanco, J. Noreña, and E. Sefusatti, *J. Cosmol. Astropart. Phys.* **09** (2022) 009.
- [50] A. Aghamousa *et al.* (DESI Collaboration), *arXiv*: 1611.00036.
- [51] M. Levi *et al.* (DESI Collaboration), *arXiv*:1308.0847.
- [52] T. Abbott *et al.* (DES Collaboration), *arXiv*:astro-ph/0510346.
- [53] P. A. Abell *et al.* (LSST Science, LSST Project Collaborations), *arXiv*:0912.0201.
- [54] A. Kosowsky, *New Astron. Rev.* **47**, 939 (2003).
- [55] P. Ade *et al.* (Simons Observatory), *J. Cosmol. Astropart. Phys.* **02** (2019) 056.
- [56] K. N. Abazajian *et al.* (CMB-S4 Collaboration), *arXiv*: 1610.02743.
- [57] A. Barreira, *J. Cosmol. Astropart. Phys.* **11** (2022) 013.
- [58] B. Reid *et al.*, *Mon. Not. R. Astron. Soc.* **455**, 1553 (2016).
- [59] L. Amendola *et al.*, *Living Rev. Relativity* **21**, 2 (2018).
- [60] O. Doré *et al.*, *arXiv*:1412.4872.
- [61] G. Cabass, M. M. Ivanov, O. H. E. Philcox, M. Simonović, and M. Zaldarriaga, *Phys. Rev. D* **106**, 043506 (2022).
- [62] G. D'Amico, M. Lewandowski, L. Senatore, and P. Zhang, *arXiv*:2201.11518.
- [63] A. Andrews, J. Jasche, G. Lavaux, and F. Schmidt, *arXiv*:2203.08838.
- [64] Y. Akrami *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A9 (2020).

-
- [65] F. Bernardeau, S. Colombi, E. Gaztanaga, and R. Scoccimarro, [Phys. Rep. **367**, 1 \(2002\)](#).
- [66] D. Baumann, A. Nicolis, L. Senatore, and M. Zaldarriaga, [J. Cosmol. Astropart. Phys. 07 \(2012\) 051](#).
- [67] J. J. M. Carrasco, M. P. Hertzberg, and L. Senatore, [J. High Energy Phys. 09 \(2012\) 082](#).
- [68] A. Taruya, K. Koyama, and T. Matsubara, [Phys. Rev. D **78**, 123534 \(2008\)](#).