

Nonperturbative test of consistency relations and their violation

Angelo Esposito¹, Lam Hui², and Roman Scoccimarro³

¹*Theoretical Particle Physics Laboratory (LPTP), Institute of Physics, EPFL, 1015 Lausanne, Switzerland*

²*Department of Physics, Center for Theoretical Physics, Columbia University,
538W 120th Street, New York, New York 10027, USA*

³*Center for Cosmology and Particle Physics, Department of Physics, New York University, New York,
New York 10003, USA*



(Received 11 June 2019; published 28 August 2019)

In this paper, we verify the large scale structure consistency relations using N -body simulations, including modes in the highly nonlinear regime. These relations (pointed out by Kehagias & Riotto and Peloso & Pietroni) follow from the symmetry of the dynamics under a shift of the Newtonian potential by a constant and a linear gradient, and predict the absence of certain poles in the ratio between the (equal time) squeezed bispectrum and power spectrum. The consistency relations, as symmetry statements, are exact, but have not been previously checked beyond the perturbative regime. Our test using N -body simulations not only offers a nonperturbative check, but also serves as a warm-up exercise for applications to observational data. A number of subtleties arise when taking the squeezed limit of the bispectrum—we show how to circumvent or address them. An interesting by-product of our investigation is an explicit demonstration that the linear-gradient symmetry is unaffected by the periodic boundary condition of the simulations. Lastly, we verify using simulations that the consistency relations are violated when the initial conditions are non-Gaussian (of the local f_{NL} type). The methodology developed here paves the way for constraining primordial non-Gaussianity using large scale structure data, including (numerous) highly nonlinear modes that are otherwise hard to interpret and utilize.

DOI: [10.1103/PhysRevD.100.043536](https://doi.org/10.1103/PhysRevD.100.043536)

I. INTRODUCTION

One of the key questions in modern cosmology concerns the initial condition of the Universe. Are the primordial fluctuations consistent with what one would expect from single-field inflation? Or do they arise from a scenario in which additional light fields, besides the inflaton, play an important role? Or more radically, is some mechanism other than inflation at work?

The standard approach to answering these questions is to work with data in the linear or quasilinear regime where perturbation theory can be relied upon to give reliable predictions. Modes in the nonlinear regime (for instance, with momentum $k \gtrsim 0.2h/\text{Mpc}$ in large scale structure data) are not utilized, even though they are abundant and measured with high precision.

The consistency relations offer an interesting alternative, where some of the information hidden in the nonlinear regime can be brought to light. First pointed out by Maldacena [1], consistency relations connect a squeezed $(N + 1)$ -point correlation function (squeezing means one of the momentum legs is soft) to an N -point function (see also [2,3]). More recent work pointed out additional consistency relations coming from new symmetries, clarified the assumptions behind consistency relations and emphasized their exact, nonperturbative nature, analogous to soft

theorems in high energy physics [4–10]. The nonperturbative nature of consistency relations is a mere curiosity for the microwave background since its fluctuations are small and linear, but becomes very interesting for a large scale structure. Kehagias & Riotto and Peloso & Pietroni [11,12] pointed out the relevant large scale structure consistency relations. It can be shown that of the infinite tower of general relativistic consistency relations [8], two has non-trivial Newtonian, sub-Hubble limits [13,14].

The study of large scale structure concerns, at a minimum, the following quantities: the mass fluctuation δ , the peculiar velocity \vec{v} and the gravitational potential Φ . (One can further expand this list to include the galaxy count fluctuation δ^g and the galaxy peculiar velocity \vec{v}^g . The symmetries discussed below apply to them as well, where δ^g and \vec{v}^g transform in the same way as δ and \vec{v} do, see e.g., [15,16].) The dynamics of (sub-Hubble) fluctuations exhibits two nonlinearly realized symmetries in a matter + cosmological constant universe.¹ One is a constant shift in the gravitational potential:

$$\Phi \rightarrow \Phi + c, \quad (1)$$

¹The split into two separate symmetries here follows the discussion of [17].

where c is independent of space but possibly a function of time. The other involves adding a linear gradient to the gravitational potential, together with a transformation of the spatial coordinates and the velocity field [14]:

$$\vec{x} \rightarrow \vec{x} + \vec{n}, \quad \Phi \rightarrow \Phi - (\mathcal{H}\vec{n}' + \vec{n}'') \cdot \vec{x}, \quad \vec{v} \rightarrow \vec{v} + \vec{n}', \quad (2)$$

where \vec{n} is independent of space but a function of time. Here $' \equiv \partial/\partial\eta$ is the derivative with respect to the conformal time η , and $\mathcal{H} \equiv a'/a$ is the comoving Hubble parameter, with a being the scale factor. The above is a symmetry of the large scale structure dynamics for \vec{n} having any time dependence, but the adiabatic mode condition [18] dictates that \vec{n} must match the time dependence of the linear growth factor, and likewise c should match the corresponding time dependence of the gravitational potential (see discussions in [10,17] and point 2 below).

The consistency relations corresponding to a shift of the gravitational potential by a constant and by a linear gradient are, respectively,

$$\lim_{\vec{q} \rightarrow 0} q^2 \frac{\langle \delta_{\vec{q}} \delta_{\vec{k}_1}^- \dots \delta_{\vec{k}_N}^- \rangle^{c'}}{P_\delta(q)} = 0, \quad (3)$$

and

$$\lim_{\vec{q} \rightarrow 0} \vec{\nabla}_q \left[q^2 \frac{\langle \delta_{\vec{q}} \delta_{\vec{k}_1}^- \dots \delta_{\vec{k}_N}^- \rangle^{c'}}{P_\delta(q)} \right] = - \sum_{a=1}^N \frac{D(\eta_a)}{D(\eta)} \vec{k}_a \langle \delta_{\vec{k}_1}^- \dots \delta_{\vec{k}_N}^- \rangle^{c'}, \quad (4)$$

where $P_\delta(q)$ is the mass power spectrum, D is the linear growth factor and c' denotes the connected correlator with the overall δ -function removed. The time dependence is as follows: the soft mode \vec{q} is at time η [likewise for $P_\delta(q)$] while the hard mode \vec{k}_a is at time η_a . Several comments are called for on these two consistency relations.

1. The consistency relations are in general of an unequal time form. In this paper, we focus on the equal time limit, in which case the right-hand side of Eq. (4) vanishes. Thus, the content of the consistency relations is simple, that the *equal time* correlator

$$\frac{\langle \delta_{\vec{q}} \delta_{\vec{k}_1}^- \dots \delta_{\vec{k}_N}^- \rangle^{c'}}{P_\delta(q)} \text{ has no } q^{-2} \text{ pole and no } q^{-1} \text{ pole} \quad (5)$$

in the $\vec{q} \rightarrow 0$ limit. The lack of a $1/q^2$ pole follows from the shift symmetry, and the lack of a $1/q$ pole follows from the linear gradient symmetry. That this statement is correct (for Gaussian initial conditions) is easy to check in perturbation theory (see e.g., [11,12,14]). But the consistency relations, as symmetry statements, are expected to be stronger than this. What we wish to accomplish in this paper is to test this statement in the nonperturbative regime using N -body

simulations (i.e., with the hard momenta \vec{k}_a 's on nonlinear scales).²

2. It should be emphasized that the consistency relations are not statements merely about a strictly vanishing \vec{q} . Indeed, an exact $\vec{q} = 0$ mode is not even observable. Rather, the consistency relations are statements about the absence of certain divergences as \vec{q} is taken to be smaller and smaller, such as (5). This is why the so-called adiabatic mode condition is crucial [18,19]. This condition ensures that the symmetry in question, which in general originates as a gauge redundancy, generates a $\vec{q} = 0$ mode that is smoothly connected to a physical mode of a small but finite \vec{q} . For more discussions on this point, see [10].

3. The consistency relations, Eqs. (3) and (4), take a particularly simple form in Lagrangian space where the corresponding ‘‘right-hand side’’ vanishes even if the hard modes are at unequal times. See [17] for a discussion.

4. The consistency relations take essentially the same form even in redshift space, as pointed out by [20]. This means they can be profitably applied in galaxy surveys where the line-of-sight direction is almost always in redshift space.

5. There is the question of how galaxy biasing affects the consistency relations. As mentioned above, the relevant symmetries remain good symmetries even for the dynamics of galaxies (which can form, merge and so on).³ Thus, the

²We focus on the equal time correlator largely for simplicity. There is also a practical reason for doing so: that the $1/q$ pole associated with the unequal time contributions [i.e., the right-hand side of Eq. (4)] is naturally suppressed in observational data. Recall that the unequal times refer to the times of the hard modes (η_1 for \vec{k}_1 , η_2 for \vec{k}_2 and so on); the hard modes are by definition short wavelength perturbations which also means their separation in time cannot be too big—keep in mind that observational data are confined to the light cone. One can see from Eq. (4) that if the η_a 's are close to each other, one is almost summing the \vec{k}_a 's which yields zero. Nonetheless, it is worth asking how big of a $1/q$ pole one might inadvertently generate by measuring a correlator averaged over some survey volume, which inevitably spans a range of redshifts. Some care in defining the average might be useful to ensure it is negligible. It is also worth noting that the unequal-time contributions do not generate a $1/q^2$ pole. A $1/q^2$ pole can only appear with certain primordial non-Gaussianities (see point 6 below).

³Galaxy dynamics is of course different from mass dynamics: mass conservation is replaced by galaxy number density evolution that has a source (or sink) term; galaxies are subject to forces beyond gravity. The key observation is that as long as these new terms/forces depend only on mass/galaxy density, velocity gradients (or velocity difference between different species) and second derivatives of the gravitational potential (tidal forces), the symmetries espoused in Eqs. (1) and (2) hold. For instance, it is crucial the new forces on a galaxy do not depend on the absolute velocity, i.e., some form of equivalence principle (see point 6 below). There is an additional requirement: that the squeezed momentum \vec{q} must be sufficiently soft, that on that scale, gravity dominates (even though for the hard momenta \vec{k} 's, the dynamics can be complicated). See [14,21] for further discussions.

consistency relations, Eqs. (3) and (4), remain valid even if the hard modes $\delta_{\vec{k}_1}^g, \dots, \delta_{\vec{k}_N}^g$ are replaced by galaxy density fluctuations $\delta_{\vec{k}_1}^g, \dots, \delta_{\vec{k}_N}^g$. The soft mode $\delta_{\vec{q}}^g$ can be replaced by $\delta_{\vec{q}}^g/b^g$ where b^g is the galaxy bias; likewise $P_\delta(q)$ can be replaced by $P_{\delta^g}(q)/b^{2g}$. In the soft limit, b^g is expected to be a constant,⁴ and thus the consistency relations Eqs. (3) and (4) are modified in a simple way. The equal time version (5) in fact takes the same form, i.e., the equal time correlator:

$$\frac{\langle \delta_{\vec{q}}^g \delta_{\vec{k}_1}^g \dots \delta_{\vec{k}_N}^g \rangle^{c'}}{P_{\delta^g}(q)} \text{ has no } q^{-2} \text{ pole and no } q^{-1} \text{ pole} \quad (6)$$

in the $\vec{q} \rightarrow 0$ limit.

6. Two important assumptions go into deriving the consistency relations. One is the equivalence principle, that on sufficiently large scales—i.e., $\vec{q} \rightarrow 0$ —all objects fall at the same rate (whereas on small scales, different objects can be subject to different forces, such as pressure forces, etc.). See [14,22] for a discussion. The other important assumption, which we focus on in this paper, is Gaussian initial conditions. More precisely, it is the assumption that in the squeezed limit, the primordial connected N -point function vanishes for $N > 2$, something that follows from single-clock inflation.⁵ From the point of view of initiating cosmological N -body simulations, imposing Gaussian initial conditions is sufficient to guarantee the validity of the consistency relations stated above, and this is what we adopt in this paper. It is not surprising that the consistency relations, or the precise form they take, are sensitive to initial conditions, since the symmetries underlying them are nonlinearly realized or spontaneously broken—in other words, exactly how the initial conditions, or the “vacuum,” breaks the symmetries in question dictates the form of the consistency relations [10]. Examples that violate the stated consistency relations generally involve extra light fields during inflation, for instance the curvaton, a spectator scalar that dominates the curvature fluctuations [27–29].⁶ The curvaton (or modulated-reheating) model motivates initial conditions of the local f_{NL} type (see Sec. IID), and we will examine how the consistency relations are violated in such a case. The ultimate goal

⁴This holds if the initial conditions were Gaussian, an assumption that goes into the derivation of the consistency relations themselves. Or more precisely, this assumes single-field or single-clock initial conditions. See the discussion in point 6 below.

⁵The primordial consistency relations can be expressed as the vanishing of the squeezed N -point function if one accounts for the fact that the metric fluctuations enter into the definition of physical momenta. See [23–26] for a discussion.

⁶Ultra-slow-roll inflation, while strictly a single field model, has essentially an extra clock due to the importance of what normally would be discarded as the decaying mode. See [10,30–32].

would be to check consistency relations in observational data, and put a bound on local f_{NL} for instance. The robustness of the consistency relations means we can freely employ data in the highly nonlinear regime (the high momentum \vec{k} modes), involving astrophysically realistic fluctuations, e.g., galaxies.

7. One might worry that the consistency relation could be violated by the finite size of the simulation box, especially for the symmetry transformation that involves shifting the gravitational potential by a term linear in \vec{x} (Eq. (2), which seems naïvely inconsistent with the periodic boundary condition of the simulations. However, from the point of view of the particles, all they see is the gradient of the potential, and the symmetry in question simply shifts this gradient by a constant, which does respect the periodic boundary condition. The fact that, as we will see, the consistency relations hold in the N -body data indeed confirms this expectation.

8. Lastly, it should be kept in mind that in the presence of features in the power spectrum (e.g., acoustic peaks), the bispectrum could present a $1/q$ behavior for mildly squeezed triangles, albeit recovering the behavior expected from Eq. (5) in the strict $q \rightarrow 0$ limit. When such features are present a simple power series in q will not suffice to describe the squeezed bispectrum, and the complete dependence should be taken into account [33,34]. For the range of k 's we are considering here, this is a negligible effect anyway.

To summarize, the goal of this paper is twofold. First, we test the consistency relations (5) at equal time using the results of N -body simulations with Gaussian initial conditions, focusing on the three-point function or bispectrum. To the best of our knowledge, this is the first time that the consistency relations have been verified for scales that are well within the nonperturbative regime.⁷ Second, we show that when the initial conditions for the primordial fields are non-Gaussian of the local f_{NL} type, deviations from (5) are observed, as expected from theoretical arguments [12,14,38].

II. CHECKING THE CONSISTENCY RELATIONS IN N -BODY SIMULATIONS

We describe in Sec. IIA our methodology, focusing in particular on how to obtain the bispectrum in the squeezed limit. This is followed by a discussion in Sec. IIB of how we fit the bispectrum with a power series in the squeezed momentum q . The results of the fit are presented in

⁷A different kind of consistency relation has been tested in [35] using N -body simulations as well. That interesting relation concerns the higher order coefficients of the low q expansion of the bispectrum [36,37]. More specifically, it concerns the q^0 behavior in the context of (5), and its derivation crucially relies on the hard observables being mass fluctuations as opposed to galaxy fluctuations. The consistency relations we focus on are instead more robust and valid even for galaxy observables.




FIG. 1. The bispectrum is defined for three Fourier modes whose momenta sum to zero. The resulting triangle can be uniquely labeled by $q \equiv |\vec{q}|$ (the softest momentum), $k \equiv |\vec{k}|$ (the hardest momentum), and the angle between them θ . By virtue of the fact that q is the softest and k is the hardest, θ is between $\pi/2$ and π .

Sec. II C, for N -body simulations with Gaussian initial conditions. We verify that the consistency relations are indeed satisfied, even though the high momentum modes are in the nonlinear regime. We draw attention to, and comment on, the fact that the linear-gradient consistency relation [i.e., the lack of $1/q$ pole in (5)] is satisfied, despite the periodic boundary conditions of the simulations—which one might naively expect to invalidate the linear gradient symmetry of Eq. (2). We demonstrate in Sec. II D that the consistency relations are violated for simulations with non-Gaussian initial conditions of the local f_{NL} type.

A. Setup and details of the measurement

We use a suite of N -body simulations consisting of $N_r = 40$ realizations with Gaussian initial conditions. The box size is $L = 2.4 \text{ Gpc}/h$ comoving, with 1280^3 particles. The cosmological parameters are $\Omega_\Lambda = 0.75$, $\Omega_m = 0.25$ (of which $\Omega_b = 0.04$), $h = 0.7$, $n_s = 1$ and $\sigma_8 = 0.8$. We analyze the simulation outputs at redshift $z = 0$. For further details on the simulations, see [39].

A prime observable of focus is the bispectrum, in the so-called squeezed limit, i.e., when one of the legs (in momentum space) is soft. We are particularly interested in what happens when that leg, labeled by the momentum q , becomes softer and softer as other quantities that label the relevant momentum-space triangle are kept fixed. A convenient parametrization is to take them to be the highest momentum leg, labeled by k , and its angular separation from the soft leg, labeled by θ . With this choice, θ is between $\pi/2$ and π (see Fig. 1).⁸ We will have more to say about the choice of parametrization below.

The bispectrum in the squeezed limit can then be expressed as a power series in the soft mode,

$$B_\delta(q, k, \theta) = \sum_{n=-2}^{\infty} a_n(k, \theta) P_\delta(q) q^n. \quad (7)$$

We will truncate this power series at some finite n , with the understanding that this is a good approximation for small

⁸By restricting ourselves to k being the highest momentum and θ between $\pi/2$ and π , we are implicitly assuming parity invariance: that two triangles related to each other by a reflection have the same bispectrum.

values of q —the precise n at which we truncate will be determined by the goodness-of-fit to the data. The consistency relations (5) tell us that, for equal time correlators, one has $a_{-2}(k, \theta) = a_{-1}(k, \theta) = 0$. The goal of this paper is to check this prediction. We wish to do it in a way that does not assume any knowledge of the coefficients a_n . They are known robustly only within perturbation theory, that is, if k is not too large. For large k 's, nonlinearity, or baryonic physics in the case of galaxy observables (in anticipation of applications to observational data), makes it difficult to robustly predict a_n . Thus we carry out the analysis without prior assumptions on them.

The key feature we exploit is that Eq. (7) takes a *factorized* form: for each n , the dependence on the soft momentum q is factorized from the dependence on the hard momentum k (and θ). The coefficients that contain the k and θ dependence, $a_n(k, \theta)$, can be treated as free parameters when fitting the bispectrum. As a simplifying procedure, since we are not ultimately interested in the k and θ dependence of the bispectrum or a_n , we average over all possible values of k and θ when we measure the bispectrum for a given q .⁹

At this point, a subtlety occurs because of the discrete nature of the Fourier modes in a finite volume. Let us focus on the coefficient a_n for a particular n . Our procedure is effectively to compute some averaged version of a_n by summing over all possible k 's and θ 's at a fixed q , i.e., summing over all triangles which has one momentum leg of magnitude q . The issue is this: within our set of discrete Fourier or momentum modes, for a given q , not all possible k 's and θ 's are actually allowed—in fact, the span of possible k 's and θ 's would depend on the value of q in a subtle way; this means the averaged a_n would end up inheriting a subtle q dependence. *This q dependence cannot be predicted without prior knowledge or assumption of how a_n depends on k and θ .* It is useful to concretely see how this comes about by dividing the k 's and θ 's into bins, labeled by i . For instance, a bin centered at (k_i, θ_i) might have contributions from $N_{i,q}$ triangles. Note how the q dependence is “sneaked” in through the fact that $N_{i,q}$ depends on q . In this language, averaging over all possible triangles for a given q amounts to computing the following:

$$\begin{aligned} \bar{B}_\delta(q) &= \frac{\sum_i B_\delta(q, k_i, \theta_i) N_{i,q}}{\sum_i N_{i,q}} \\ &= \sum_{n=-2}^{\infty} \frac{\sum_i a_n(k_i, \theta_i) N_{i,q}}{\sum_i N_{i,q}} P_\delta(q) q^n \\ &\equiv \sum_{n=-2}^{\infty} \bar{a}_n(q) P_\delta(q) q^n. \end{aligned} \quad (8)$$

⁹We will later check this procedure by varying the range of θ over which we average.

We are thus left with an averaged a_n , which we call \bar{a}_n , which has an unwanted q dependence which cannot be predicted without making assumptions about how a_n behaves for high k 's. Thus, imagine we fit the N -body data with n up to, e.g., 1. Even if one puts aside the possible q dependence of \bar{a}_{-2} and \bar{a}_{-1} (which for Gaussian initial conditions are expected to vanish), the unknown q dependence of \bar{a}_0 and \bar{a}_1 is problematic.

This way of spelling out the problem also suggests its cure. The above averaging weighs each i th bin by the number of triangles in it, $N_{i,q}$. We can instead weigh each bin equally (or for that matter, use any other weights as long as they do not depend on q)¹⁰:

$$\bar{B}_\delta(q) = \frac{\sum_i B_\delta(k_i, \theta_i, q)}{\sum_i} \equiv \sum_{n=-2}^{\infty} \bar{a}_n P_\delta(q) q^n. \quad (9)$$

The coefficients \bar{a}_n are now given by

$$\bar{a}_n \equiv \frac{\sum_i a_n(k_i, \theta_i)}{\sum_i}, \quad (10)$$

and are truly independent of the soft momentum. They are treated as free parameters in our fit of the data.

To simplify the analysis, we also bin in q . In particular, if $\{\bar{q}\}$ is a bin with average soft momentum \bar{q} , the binned version of the bispectrum (9) is

$$\bar{B}_\delta(\bar{q}) = \sum_{q \in \{\bar{q}\}} \bar{B}_\delta(q) = \sum_{n=-2}^{n_{\max}} \bar{a}_n M_n(\bar{q}), \quad (11)$$

where $M_n(\bar{q}) = \sum_{q \in \{\bar{q}\}} P_\delta(q) q^n$ is also measured from the data to avoid any theoretical bias. The value of n_{\max} is something we have to experiment with: qualitatively, the more squeezed our triangles are (smaller q 's), the lower is the n_{\max} we need. In the next section we will rely on data to find how many \bar{a}_n 's we need to account for higher order corrections to the small q expansion.

Lastly, let us comment on our bispectrum triangle parametrization, described in Fig. 1. One alternative [40] is to parametrize the triangle in terms of \vec{q} , $\vec{k} + \vec{q}/2$ (where \vec{k} is one of the two high momentum legs), and the angle between them, say β . Assuming invariance under parity (a reflection of the triangle), the bispectrum should be unchanged under $\cos \beta \rightarrow \cos(\pi - \beta)$. Thus, if q enters into the bispectrum only through $\cos \beta$, the squeezed bispectrum should contain only even powers of q , as suggested by [40]. This appears to be true in some cases but not in others—for instance, it can be checked in perturbation theory that if the initial conditions were of the local f_{NL} type, the squeezed mass bispectrum depends

¹⁰In practice, this means when we loop through the triangles for a given q , we weigh them by $1/N_{i,q}$.




FIG. 2. Measured value of $\Delta^2(k) = 4\pi k^3 P_\delta(k)$. For the hard modes under consideration, $k \in [0.52, 0.66]h/\text{Mpc}$, one notices that $\Delta^2(k) \gtrsim 5$, i.e., we are far away from the linear regime.

on the transfer function at the soft-momentum q (which equals 1 when $q = 0$, but has corrections with both even and odd powers of q) [12].¹¹ We thus do not find a significant advantage for using the alternative parametrization, although the bispectrum can be analyzed that way if one wishes.

B. Details of the fit

Our analysis is performed averaging the bispectrum as in Eq. (9) over hard modes ranging from $k = 0.52h/\text{Mpc}$ to $k = 0.65h/\text{Mpc}$ (corresponding to $k = 199k_f$ to $k = 251k_f$, with $k_f = 2\pi/L \simeq 2.6 \times 10^{-3}h/\text{Mpc}$ being the fundamental mode), and over all the relative angles, $\theta \in [\pi/2, \pi]$.¹² As one can see from Fig. 2 the hard modes we are considering are well within the nonlinear regime.

We also measure the bispectrum for soft momenta ranging from $q = 3k_f$ to $q = 19k_f$, with a bin size $\Delta k = 2k_f$. The choice of this window for q 's requires some explanations. The high end $q = 19k_f$ is chosen to include as many modes as possible (thus minimizing error bars on \bar{a}_{-2} and \bar{a}_{-1}), while still staying within the squeezed limit such that the expansion in Eq. (9), truncated at n of a few, is a good approximation.

¹¹Even in those cases where the squeezed bispectrum parametrized according to [40] appears as an even-power series in q , the information does not allow us to for instance infer \bar{a}_{-1} from \bar{a}_{-2} (these are parameters we are ultimately interested in, using our parametrization). In those cases, \bar{a}_{-1} is not directly related to \bar{a}_{-2} but is instead related to the average over k and θ of the derivatives of a_{-2} with respect to k and θ .

¹²As a simple check, we also repeated our analysis averaging over $\theta \in [\pi/2, 3\pi/4]$. The conclusions are largely unchanged, suggesting that adjusting the angular weighting does not have a significant impact on the outcome, and that the vanishing of \bar{a}_{-2} and \bar{a}_{-1} for Gaussian initial conditions is not the result of accidental cancellation when averaging over angles.

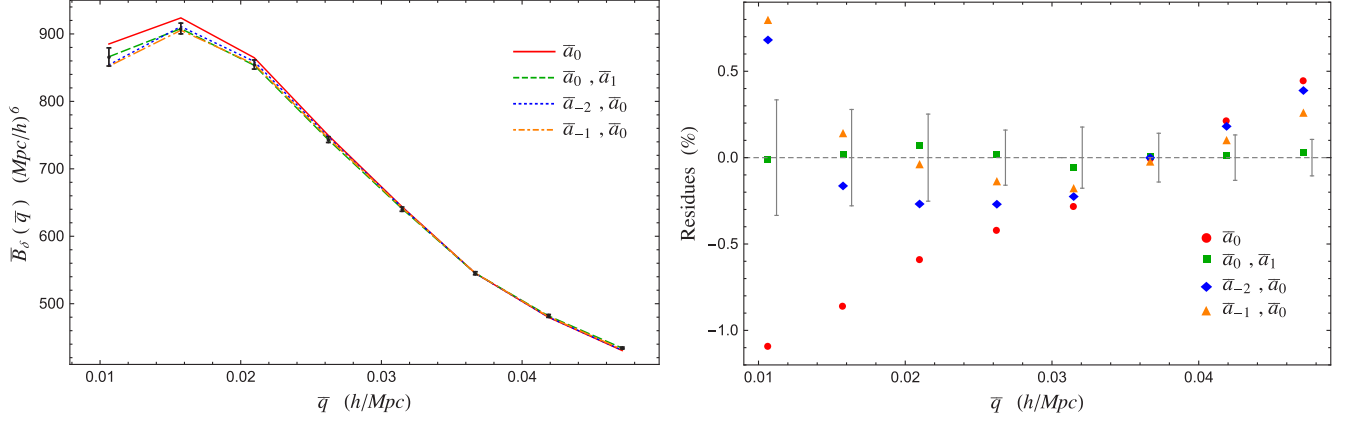


FIG. 3. Left panel: Comparison between N -body data (points with error bars) and some of the models (lines with different styles/colors), including different sets of parameters. The y -axis is the bispectrum \bar{B}_δ while the x -axis is soft momentum \bar{q} . Note that \bar{B}_δ is a function only of \bar{q} , because we have already summed over k 's and θ 's [see Eqs. (9) and (11)]. Here, the error bars reflect the statistical spread in \bar{B}_δ . Right panel: Percentage difference between data and model for the same set of models as the left panel, i.e., the y -axis is $[\bar{B}_\delta(\bar{q}) - \sum_n \bar{a}_n M_n(\bar{q})] / [\bar{B}_\delta(\bar{q}) + \sum_n \bar{a}_n M_n(\bar{q})]$, expressed in percentage. The error bars shown represent the statistical spread in this residue. The error bars are model dependent—the ones shown correspond to that of the (\bar{a}_0, \bar{a}_1) model. Note that the error bars are correlated across different \bar{q} 's, which partly explains why the (\bar{a}_0, \bar{a}_1) model appears well within the error bars at all momenta.

The low end $q = 3k_f$ is chosen because the procedure of eliminating the unwanted q dependence in \bar{a}_n , described in Sec. II A, is actually not perfect. Recall that we form bins, labeled as (k_i, θ_i) , and compute \bar{B}_δ and \bar{a}_n [Eqs. (9) and (10)] by giving these bins equal weights. In doing so, it is important that each bin is actually not empty, that there are triangles that fall into them. Thus, the bins have to be sufficiently wide. But there is some tension between using wide bins and using the bin-averaged a_n as a fair representation of how a_n varies with k and θ .¹³ In the Appendix, we show a test of our procedure for a particular model of $a_n(k, \theta)$ (one motivated by perturbation theory), and check to what extent our procedure yields \bar{a}_n that is truly q independent. We find that this works well as long as $q \geq 3k_f$. Hence we restrict our analysis only to soft modes such that $q \geq 3k_f$.

To determine the best-fit values for the parameters \bar{a}_n we maximize the following likelihood for each realization, $r = 1, \dots, N_r$:

$$\mathcal{L}_{(r)} \propto \frac{1}{\sqrt{\det \mathbf{C}}} \exp \left[-\frac{1}{2} \mathbf{\Delta}^{(r)} \cdot \mathbf{C}^{-1} \cdot \mathbf{\Delta}^{(r)} \right]. \quad (12)$$

We define the vector $\mathbf{\Delta}^{(r)} = \bar{\mathbf{B}}_\delta^{(r)} - \sum_n \bar{a}_n^{(r)} \bar{\mathbf{M}}_n^{(r)}$, and the covariance matrix $C_{ij} = \langle \Delta_i \Delta_j \rangle$. All vectors run over the soft momenta, $f_i \equiv f(\bar{q}_i)$, and the angular brackets stand for an average over the available realizations, i.e.,

¹³In other words, within a wide bin, the precise set of k 's and θ 's that fall into that bin would depend on q , and thus we are not achieving the goal of eliminating the unwanted q dependence.

$\langle f_i \rangle = \frac{1}{N_r - 1} \sum_r f_i^{(r)}$ (for instance, the covariance matrix is obtained by averaging over realizations).

Note that the vector $\mathbf{\Delta}$ depends on the fit parameters \bar{a}_n , and so the covariance matrix itself depends on the parameters. We use an iterative procedure (akin to the Newton-Raphson algorithm) to determine the optimal \bar{a}_n 's that maximize the likelihood. First, we determine C_{ij} with the \bar{a}_n 's set to zero. The maximization of the likelihood can thus be done analytically, because the remaining dependence on \bar{a}_n shows up only in the exponent of the likelihood in the standard χ^2 fashion (essentially equivalent to fitting the slope of a straight line). The resulting best-fit \bar{a}_n 's are plugged back into the definition of C_{ij} , and the whole procedure is repeated again to obtain a new set of best-fit \bar{a}_n 's; so on and so forth until convergence is achieved.

Once this is done, the final value of the ML estimators and their uncertainties is computed from the average and variance over realizations, i.e.,

$$\bar{a}_n = \frac{1}{N_r} \sum_{r=1}^{N_r} \bar{a}_n^{(r)}, \quad (13a)$$

$$\sigma_{\bar{a}_n}^2 = \frac{1}{N_r(N_r - 1)} \sum_{r=1}^{N_r} (\bar{a}_n^{(r)} - \bar{a}_n)^2. \quad (13b)$$

Note that the likelihood analysis itself, applied to each realization, does yield an error estimate, but we deem $\sigma_{\bar{a}_n}^2$ estimated from the spread between *independent* realizations as more reliable. For one thing, the likelihood analysis treats the data vector as Gaussian distributed, which is an approximation. The desire to have an accurate error

TABLE I. Detailed results of the likelihood fits with different sets of parameters (\bar{a}_n 's). The Bayesian information criterion clearly selects the (\bar{a}_0, \bar{a}_1) model.

\bar{a}_n included	$\bar{a}_{-2}(10^{-6} \text{ Mpc}/h)$	$\bar{a}_{-1}(10^{-2} \text{ Mpc}/h)^2$	BIC
\bar{a}_0	99.63
\bar{a}_0, \bar{a}_1	17.77
$\bar{a}_0, \bar{a}_1, \bar{a}_2$	19.82
\bar{a}_{-2}, \bar{a}_0	-30.4 ± 5.1	...	65.50
$\bar{a}_{-2}, \bar{a}_0, \bar{a}_1$	0.2 ± 6.7	...	19.84
\bar{a}_{-1}, \bar{a}_0	...	-42.3 ± 5.4	39.57
$\bar{a}_{-1}, \bar{a}_0, \bar{a}_1$...	0.6 ± 10.3	19.84
$\bar{a}_{-2}, \bar{a}_{-1}, \bar{a}_0$	69 ± 16	111 ± 17	22.35
$\bar{a}_{-2}, \bar{a}_{-1}, \bar{a}_0, \bar{a}_1$	16 ± 55	26 ± 87	21.83

estimate is why we analyze the realizations one at a time, as opposed to using all of them in one go.¹⁴

To determine the goodness of the fit, we rely on the Bayesian information criterion (BIC) [41,42]:

$$\text{BIC} = -2 \log \mathcal{L}_{\text{max}} + N_{\text{par}} \log N_q, \quad (14)$$

which has been shown to be dimensionally consistent, i.e., not to favor overfitted models [42]. Here \mathcal{L}_{max} is the maximum likelihood combining all realizations, N_{par} the number of parameters of the model and N_q the number of data points used. A model with the lowest BIC represents the best compromise between maximizing likelihood and minimizing the number of parameters.

C. Results

Let us now present the results of our analysis. In Fig. 3 we report some of the fits including different sets of parameters as well as the corresponding residues. In Table I we compare all the models we have tested.

Focus first on the first three models in the table, which do not involve \bar{a}_{-2} nor \bar{a}_{-1} . We see that the first model, involving \bar{a}_0 alone, is not a good fit to the N -body data (from both the BIC value in the table and from Fig. 3). Adding \bar{a}_1 greatly improves the fit, while further adding \bar{a}_2 does not lower the BIC score. Recall that in our power-series fit of the squeezed bispectrum as a function of a range of soft momenta [Eq. (11)], we do not know *a priori* how many higher order terms we need. This exercise tells us it is sufficient to stop at \bar{a}_1 (but necessary to include it), with the kind of precision and the range of soft momenta we have.

The rest of the models in the table involve \bar{a}_{-1} and/or \bar{a}_{-2} . In all cases, the BIC score worsens. The inferred values for \bar{a}_{-1} and \bar{a}_{-2} are consistent with zero, except for the $(\bar{a}_{-2}, \bar{a}_{-1}, \bar{a}_0)$ model. For this model, the fit prefers nonzero values for \bar{a}_{-2} and \bar{a}_{-1} to compensate for the lack of a \bar{a}_1 term. Note however this model has a worse BIC

¹⁴Analyzing the realizations all at once would give us essentially the same final best-fit \bar{a}_n , but would not let us reliably estimate the associated error bar.

score compared to the (\bar{a}_0, \bar{a}_1) model. It is also reassuring that the (\bar{a}_0, \bar{a}_1) fits the data well, with residues that have no clear trend with momenta (see the right panel of Fig. 3).

We conclude from this exercise that the N -body data, with Gaussian initial conditions, are consistent with a vanishing value for \bar{a}_{-2} and for \bar{a}_{-1} , confirming expectations from the consistency relations.

D. Violation of the consistency relations from non-Gaussian initial conditions

In this section we show that, when the initial conditions for the cosmological fields are non-Gaussian (of the local f_{NL} type), statistically significant deviations from the consistency relations in Eq. (5) are observed.¹⁵

We employ a smaller set of $N_r = 12$ realizations with the same cosmological parameters as before, but with an initial matter distribution characterized by a local non-Gaussian parameter $f_{\text{NL}} = 100$. The details of the measurement and the analysis are the same as in Secs. II A and II B. For the sake of checking whether or not deviations from consistency relations occur we limit our analysis to models that only include \bar{a}_0 and \bar{a}_1 (in addition to possibly \bar{a}_{-2} and \bar{a}_{-1}). We leave a more detailed study of the realizations with non-Gaussian initial condition for future work [43].

The result of our analysis is unambiguous. The model with just \bar{a}_0 and \bar{a}_1 (i.e., no poles in the soft limit), which fits very well the bispectrum in the case of Gaussian initial conditions, is not a good description of the data obtained from $f_{\text{NL}} = 100$. From the fit we obtain $\text{BIC} = 74.12$, much larger than the one reported in Table I for the Gaussian case. Moreover, Fig. 4 shows that the likelihood fit for this model is not a good description of the data, which is also confirmed by the fact that the residues exhibit a parabolic pattern around zero. Indeed, introducing either \bar{a}_{-2} or \bar{a}_{-1} to the fit one obtains values that are statistically different from zero: $\bar{a}_{-2} = (11.3 \pm 1.8) \times 10^{-5} \text{ Mpc}/h$ with $\text{BIC} = 30.39$, or $\bar{a}_{-1} = (16.5 \pm 2.8) \times 10^{-3} (\text{Mpc}/h)^2$ with $\text{BIC} = 36.79$.¹⁶ The inferred values for \bar{a}_{-2} and \bar{a}_{-1} can in principle be turned into an estimate of f_{NL} , which we leave for future work.

This shows that, in the presence of a non-Gaussian distribution (of the local f_{NL} type) for the initial cosmological fields, the consistency relations in Eq. (5) are violated as expected [12,14,38].

¹⁵Specifically, the local f_{NL} model is this: the primordial Bardeen potential $\Phi_p(\vec{x}) = \phi(\vec{x}) + f_{\text{NL}}(\phi(\vec{x})^2 - \langle \phi^2 \rangle)$, where ϕ is a Gaussian random field. The Bardeen potential (after multiplication by the transfer function) gives the gravitational potential used in initializing N -body simulations (see [39] for details). A primordial non-Gaussianity of this type is motivated by the curvaton and modulated reheating models [27–29].

¹⁶For completeness, let us mention several additional models we investigated: the $(\bar{a}_0, \bar{a}_1, \bar{a}_2)$ model has a BIC of 56.5, and the $(\bar{a}_{-2}, \bar{a}_0, \bar{a}_1, \bar{a}_2)/(\bar{a}_{-1}, \bar{a}_0, \bar{a}_1, \bar{a}_2)/(\bar{a}_{-2}, \bar{a}_{-1}, \bar{a}_0, \bar{a}_1, \bar{a}_2)/(\bar{a}_{-2}, \bar{a}_{-1}, \bar{a}_0, \bar{a}_1)$ models have respectively a BIC score of 31.8, 34.3, 32.2, 31.2.




FIG. 4. Left panel: Comparison between N -body data and a (\bar{a}_0, \bar{a}_1) model fit for a case where the initial conditions are non-Gaussian (of the local f_{NL} type, with $f_{\text{NL}} = 100$). Right panel: The blue circles represent the residues for the (\bar{a}_0, \bar{a}_1) model, i.e., $[\bar{B}_\delta(\bar{q}) - \sum_n \bar{a}_n M_n(\bar{q})] / [\bar{B}_\delta(\bar{q}) + \sum_n \bar{a}_n M_n(\bar{q})]$ expressed in percentage, where \bar{B}_δ is the N -body bispectrum with initial conditions of local $f_{\text{NL}} = 100$. The error bars shown represent the statistical spread in this residue. Also shown as green squares are the same residues as the green squares in Fig. 3, i.e., residues for the Gaussian case ($f_{\text{NL}} = 0$). The larger error bars compared to those in Fig. 3 reflect the fact that fewer realizations are used in this analysis.

III. DISCUSSION

The search for primordial non-Gaussianities has been so far a challenging task. This is partly due to our lack of theoretical control over observables that are outside the linear regime. Consistency relations are nonperturbative statements that follow solely from symmetry arguments and, as such, might provide a key tool to overcome these difficulties.

In this paper we successfully test them, for the first time, in a regime well outside the domain of perturbation theory. In doing so, we highlight and solve a number of technical and conceptual subtleties associated with the analysis of the bispectrum in the squeezed regime from N -body simulations, whose systematic study has been lacking from the literature (see [35] for an exception, though see footnote 7).

Moreover, we show how in the presence of non-Gaussian initial conditions of the local f_{NL} type, significant deviations from the standard consistency relations are observed. This is the first step towards extracting constraints on f_{NL} from observational data, using the consistency relations (or violations thereof). The appeal is that with this method, (nonlinear) modes that are normally discarded can now be used. Several issues need to be investigated before this goal can be realized. They include: checking the consistency relations (1) for biased observables such as halos in N -body simulations or galaxies in hydrodynamic simulations, and (2) including redshift space distortions. As explained in Sec. I, the consistency relations are expected to be robust against these complications, but it would be useful to test the expectations against simulations—our simple exercise presented in this paper suggests there could well be subtleties that need to be understood and addressed.

ACKNOWLEDGMENTS

The authors are grateful to A. Joyce, A. Lewis, A. Nicolis, R. Penco, M. Pietroni, M. Simonović and S. Wong for

interesting and useful discussions, and to M. Abitbol for illuminating insights on the statistical analysis. The work done by A. E. is supported by the Swiss National Science Foundation under Contract No. 200020-169696 and through the National Center of Competence in Research SwissMAP. The work done by L. H. is supported in part by the NASA Grant No. NXX16AB27G and the DOE Grant No. DE-SC011941.

APPENDIX: CHECKING SOFT MOMENTUM FACTORIZATION

In this Appendix we show that the procedure outlined around Eqs. (9) and (10) might not eliminate the unwanted q dependence if q is extremely small.

As an explicit check let us consider the result obtained in perturbation theory. When the hard mode k is within the




FIG. 5. Comparison between $\bar{f}(\bar{q})$ and $\bar{a}_0^{\text{pt}} \bar{q}$. The procedure outlined around Eqs. (9) and (10), if it works, should make the two very close to each other. (Here, \bar{f} is the analog of \bar{B}_δ over there.) One can notice that, for very small soft momenta, the procedure does not work so well.

linear regime one can easily show that the squeezed limit of the bispectrum (see e.g., [14,44]) gives

$$a_0^{\text{pt}}(k, \theta) = \left(\frac{13}{7} + \frac{8}{7} \cos^2 \theta \right) P_\delta(k) - \cos^2 \theta k P'_\delta(k).$$

Let us then consider the toy model $f(q, k, \theta) = a_0^{\text{pt}}(k, \theta)q$ (i.e., f is our toy bispectrum for which a_0^{pt} as a function of k and θ is exactly known), and bin it as in Eqs. (9) and (11) over $k \in [79, 91]k_f$ and all relative angles. Let us call the result \bar{f} . Recall that the worry was that, with a discrete set of triangles, the average over k and θ of a_0^{pt} (we call this \bar{a}_0^{pt}) would secretly depend on q . In our toy example, since the k and θ dependence of a_0^{pt} are precisely known, we can compute this average exactly without reference to the

particular triangles we happen to have in our discrete grid. If our procedure to remove the unwanted q dependence works, then it should be $\bar{f}(\bar{q}) = \bar{a}_0^{\text{pt}}\bar{q}$.

In Fig. 5 we report the result of our measurement. As one can see, the procedure works reasonably well only if the soft momentum is $q \gtrsim 3k_f$. In the analysis reported in Sec. II C we therefore exclude the lowest momentum bin. See Sec. II B for a discussion of why our procedure does not perfectly remove the unwanted q dependence.

Finally, it can be checked that, in the presence of local f_{NL} non-Gaussianities, the factorization in Eq. (11) holds better, even for low momenta. Indeed, the dominant $a_{-2}(k)$ term is expected to have no θ -dependence and only a mild k -dependence [12].

-
- [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] P. Creminelli, J. Noreña, and M. Simonović, *J. Cosmol. Astropart. Phys.* **07** (2012) 052.
[5] K. Hinterbichler, L. Hui, and J. Khoury, *J. Cosmol. Astropart. Phys.* **08** (2012) 017.
[6] V. Assassi, D. Baumann, and D. Green, *J. Cosmol. Astropart. Phys.* **11** (2012) 047.
[7] A. Kehagias and A. Riotto, *Nucl. Phys.* **B864**, 492 (2012).
[8] K. Hinterbichler, L. Hui, and J. Khoury, *J. Cosmol. Astropart. Phys.* **01** (2014) 039.
[9] W. D. Goldberger, L. Hui, and A. Nicolis, *Phys. Rev. D* **87**, 103520 (2013).
[10] L. Hui, A. Joyce, and S. S. C. Wong, *J. Cosmol. Astropart. Phys.* **02** (2019) 060.
[11] A. Kehagias and A. Riotto, *Nucl. Phys.* **B873**, 514 (2013).
[12] M. Peloso and M. Pietroni, *J. Cosmol. Astropart. Phys.* **05** (2013) 031.
[13] P. Creminelli, J. Noreña, M. Simonović, and F. Vernizzi, *J. Cosmol. Astropart. Phys.* **12** (2013) 025.
[14] B. Horn, L. Hui, and X. Xiao, *J. Cosmol. Astropart. Phys.* **09** (2014) 044.
[15] A. Kehagias, J. Norea, H. Perrier, and A. Riotto, *Nucl. Phys.* **B883**, 83 (2014).
[16] A. Kehagias, A. Moradinezhad Dizgah, J. Norea, H. Perrier, and A. Riotto, *J. Cosmol. Astropart. Phys.* **08** (2015) 018.
[17] B. Horn, L. Hui, and X. Xiao, *J. Cosmol. Astropart. Phys.* **09** (2015) 068.
[18] S. Weinberg, *Phys. Rev. D* **67**, 123504 (2003).
[19] M. Mirbabayi and M. Simonović, arXiv:1602.05196.
[20] P. Creminelli, J. Gleyzes, M. Simonović, and F. Vernizzi, *J. Cosmol. Astropart. Phys.* **02** (2014) 051.
[21] M. Peloso and M. Pietroni, *J. Cosmol. Astropart. Phys.* **04** (2014) 011.
[22] P. Creminelli, J. Gleyzes, L. Hui, M. Simonović, and F. Vernizzi, *J. Cosmol. Astropart. Phys.* **06** (2014) 009.
[23] T. Tanaka and Y. Urakawa, *J. Cosmol. Astropart. Phys.* **05** (2011) 014.
[24] E. Pajer, F. Schmidt, and M. Zaldarriaga, *Phys. Rev. D* **88**, 083502 (2013).
[25] R. Bravo, S. Mooij, G. A. Palma, and B. Pradenas, *J. Cosmol. Astropart. Phys.* **05** (2018) 025.
[26] R. Bravo, S. Mooij, G. A. Palma, and B. Pradenas, *J. Cosmol. Astropart. Phys.* **05** (2018) 024.
[27] D. H. Lyth, *J. Cosmol. Astropart. Phys.* **06** (2006) 015.
[28] G. Dvali, A. Gruzinov, and M. Zaldarriaga, *Phys. Rev. D* **69**, 023505 (2004).
[29] L. Kofman, arXiv:astro-ph/0303614.
[30] M. H. Namjoo, H. Firouzjahi, and M. Sasaki, *Europhys. Lett.* **101**, 39001 (2013).
[31] J. Martin, H. Motohashi, and T. Suyama, *Phys. Rev. D* **87**, 023514 (2013).
[32] B. Finelli, G. Goon, E. Pajer, and L. Santoni, *Phys. Rev. D* **97**, 063531 (2018).
[33] M. Mirbabayi, M. Simonović, and M. Zaldarriaga, arXiv:1412.3796.
[34] T. Baldauf, M. Mirbabayi, M. Simonović, and M. Zaldarriaga, *Phys. Rev. D* **92**, 043514 (2015).
[35] T. Nishimichi and P. Valageas, *Phys. Rev. D* **90**, 023546 (2014).
[36] A. Kehagias, H. Perrier, and A. Riotto, *Mod. Phys. Lett. A* **29**, 1450152 (2014).
[37] P. Valageas, *Phys. Rev. D* **89**, 123522 (2014).
[38] P. Valageas, A. Taruya, and T. Nishimichi, *Phys. Rev. D* **95**, 023504 (2017).
[39] R. Scoccimarro, L. Hui, M. Manera, and K. C. Chan, *Phys. Rev. D* **85**, 083002 (2012).
[40] A. Lewis, *J. Cosmol. Astropart. Phys.* **10** (2011) 026.
[41] G. Schwarz, *Ann. Stat.* **6**, 461 (1978).
[42] A. R. Liddle, *Mon. Not. R. Astron. Soc.* **351**, L49 (2004).
[43] A. Esposito, L. Hui, and R. Scoccimarro (to be published).
[44] F. Bernardeau, S. Colombi, E. Gaztanaga, and R. Scoccimarro, *Phys. Rep.* **367**, 1 (2002).