The Predictive Power of Computational Astrophysics as a Discovery Tool
Proceedings IAU Symposium No. 362, 2023
D. Bisikalo, D. Wiebe & C. Boily, eds.
doi:10.1017/S1743921322001788

The Effect of Non-Gaussian Primordial Perturbations on Large-Scale Structure

G. A. Peña and G. N. Candlish

Instituto de Física y Astronomía, Universidad de Valparaíso, Gran Bretaña 1111, Valparaíso, Chile
email: greco.pena@postgrado.uv.cl

Abstract. The late-time effect of primordial non-Gaussianity offers a window into the physics of inflation and the very early Universe. In this work we study the consequences of a particular class of primordial non-Gaussianity that is fully characterized by initial density fluctuations drawn from a non-Gaussian probability density function, rather than by construction of a particular form for the primordial bispectrum. We numerically generate multiple realisations of cosmological structure and use the late-time matter polyspectra to determine the effect of these modified initial conditions. In this non-Gaussianity has only a small imprint on the first polyspectra, when compared to a standard Gaussian cosmology. Furthermore, some of our models present an interesting scale-dependent deviation from the Gaussian case in the bispectrum and trispectrum, although the signal is at most at the percent level.

Keywords. Large-scale structure of Universe – inflation – methods: numerical

1. Introduction

One of the current challenges in cosmology is to understand the physical processes that gave rise to the primordial inhomogeneities of the Universe. These inhomogeneities are generated during cosmic inflation in the standard model of cosmology. There is a wide variety of inflationary models that spontaneously give rise to the initial conditions of the Universe. While the simplest models, assuming a single inflationary scalar field, lead to a nearly Gaussian distribution of primordial perturbations, more complex models with larger numbers of degrees of freedom can produce measureable levels of non-Gaussianity (here a comprehnsive review: Celoria & Matarrese (2018)).

If primordial perturbations are the seeds that gave rise to the structures of the Universe, we can assume that there is an underlying signal corresponding to the primordial non-Gaussianities. In this work we consider a simple approach to test this signal on the formation of large-scale structure. We use a recently-proposed characterization of this primordial non-Gaussianity in terms of a non-Gaussian PDF describing the primordial curvature perturbations Chen et al. (2018a,b)). Using this approach will allow us to generate initial conditions for our simulations by simply changing the PDF from which we draw our initial sample of density perturbations. A significant difference of this approach compared to previous approaches to studying non-Gaussianity in cosmological simulations is that we do not limit ourselves to specifying a particular form of the primordial bispectrum, as we have the full PDF available to us.

In this work we aim to perform a preliminary exploration of the parameter space of this novel characterization of the primordial non-Gaussianity. In order to extract the signal from that induced through the non-linear process of gravitational collapse we must average over several realisations of cosmological evolutions. Thus we would, in principle,

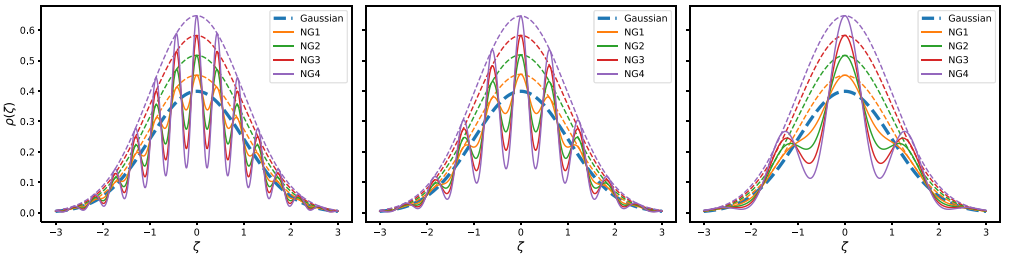

Figure 1. The non-Gaussian probability density distributions for the primordial curvature perturbations considered in this work (see Eq. (108) in Chen et al. (2018a)). The thick blue dashed line is the fiducial Gaussian distribution. All of our PDFs have unit variance. The PDF cut in the horizontal range between $\zeta_{min} = -4$, and $\zeta_{max} = 4$.

require multiple N-body simulations where we vary the initial sample of random numbers as well as the parameters of the PDF. To facilitate the generation of these realisations we have used the L-PICOLA code (Howlett et al. (2015)), which has been verified to be sufficiently accurate for our purposes by comparing with a smaller number of full N-body simulations run using the RAMSES code (Teyssier (2002)).

2. Simulations

2.1. *Initial conditions*

The inflationary scenario that gives rise to the generation of the non-Gaussian initial conditions comprises perturbations in an axion-like isocurvature field (with sinusoidal potential) which are coupled to the primordial curvature field. From this scenario rise non-Gaussian initial conditions characterized by the probability density distribution of the primordial curvature perturbation (see Eq. (104) of Chen et al. (2018a,b)). As PDF is characterized by an oscillatory modulation, we generate 4 different amplitudes (referred to as non-Gaussianity levels), and 3 different frequencies, as shown in Fig. 1. The thick blue dashed line indicates a Gaussian PDF with a standard deviation of unity. The orange, green, red, and purple lines correspond to the four levels of non-Gaussianity we consider. The dashed lines correspond to Gaussian envelopes that touch the upper peaks of the oscillations, and each panel represents a different modulated frequency. Note that we are considering large deviations from the amplitude of Gaussianity to explore the feasibility of detecting this type of NG in n-point late-time statistics; observational constraints such as those of the CMB could probably exclude our non-Gaussian primordial PDF.

We generate a sample of each of our non-Gaussian PDFs by using a simple accept-reject technique, whereby we generate uniformly distributed random values within a bounding box that includes the highest peak of the oscillatory PDF and is cut in the horizontal range between $\zeta_{min} = -4$, and $\zeta_{max} = 4$. This technique consists in accepting the values that belong to the distribution while rejecting those points that do not. The procedure ends when we generate N values for all distributions, where N is the total number of particles in our simulations, chosen to be equal to the number of points used in our discretized density field mesh. The final step is to sort the non-Gaussian distributions according to the ordering of the Gaussian one. In other words, we generate a "white noise field" using the PDFs at all points.

This procedure was repeated for all 4 levels of non-Gaussianity, all 3 frequencies, and 5 different values of the initial seed of the uniform random number generator, to generate 5 different statistical realizations. Therefore, in total, we produced initial conditions for 65 models, 5 of them being fiducial Gaussian models. To provide clarity when indicating a model, we have adopted the following naming:

https://doi.org/10.1017/S1743921322001788 Published online by Cambridge University Press

• Initial condition type: G for Gaussian, NG for non-Gaussian.
• Level of non-Gaussianity: 1, 2, 3, or 4. This only applies for NG models. The meaning of these levels will be clarified shortly.
• Frequency of non-Gaussianity: f1, f2, or f3. This only applies for NG models. The meaning of these levels will be clarified shortly.
• Random realisations: r1, r2, r3, r4 or r5. This will only be necessary when we refer to one specific realization. For most results, we average over all realizations.

For our simulations, we have used the N-body code, RAMSES, and the mock catalog, L-PICOLA. To generate the initial conditions for RAMSES we have passed the white noise to MUSIC (Hahn & Abel (2011)), where it is rearranged in a discretized grid. Then it is transformed to Fourier space and multiplied with a k-space transfer function generated by CAMB:

$$\delta(\vec{k}) = ck^{n_s/2}T(k)\mu(\vec{k}), \tag{2.1}$$

where $\mu(\vec{k})$ is the Fourier-transformed set of white noise, the transfer function $T(k)$ corresponds to a standard LCDM cosmology, and c is a normlisation constant. The real space over-density field $\delta(r)$ is then obtained by inverse Fourier transformation. Finally the 2LPT method is applied to generate the initial conditions.

In the case of L-PICOLA, we have modified the initial condition generation to read and use our random numbers, where the 2LPT method is also used. Therefore, the generation of the initial conditions of L-PICOLA is the same as MUSIC. We use RAMSES for just one realization with frequency f1 (5 models). All other models are generated with L-PICOLA.

For all our simulations, we have assumed a standard ΛCDM cosmology with the following parameters: $\Omega_m = 0.3$, $\Omega_b = 0.04$, $\Omega_\Lambda = 0.7$, $\sigma_8 = 0.88$, and $n_s = 0.96$. In addition, for all realisations we have used a 500 Mpc box, and a particle number of 256^3. In the case of RAMSES, which employs the AMR method, we have set a coarser grid resolution of 256^3 points, with 6 levels of refinement. In the case of L-PICOLA we have a fixed grid resolution of 256^3 points.

2.2. *Polyspectra Analysis*

An important feature of the nature of the non-Gaussian PDF is that only the even n-point functions are non-zero, while the odd n-point functions are zero just like a Gaussian PDF. On the other hand, possible evidence of primordial non-Gaussianities is mixed with non-Gaussianities coming from the gravitational collapse of the large-scale structure at all orders of the n-point function making it difficult to decode its signature. Therefore, as preliminary work, we decided to focus on the search for the possible signature of this primordial Gaussianity by considering the power spectrum, bispectrum, and trispectrum. As we did not find a major significance in the power spectrum, in this paper we only report the results for the bispectra and trispectra.

As a first step, we focused our analyses on symmetric configurations. For the bispectrum we choose $k_1 = k_2 = k_3$, and for the trispectrum we choose $k_1 = k_2 = k_3 = k_4$. For the latter case, we have considered all (possibly folded) quadrilaterals with equal side lengths.

We calculated the polyspectra using the Pylians code (Villaescusa-Navarro 2018). We have modified the code and added the trispectrum calculation for our purposes. We use 35 linearly-spaced bins in the range $2.2k_F \leq k \leq k_F N_k/3$, where the fundamental frequency is $k_F = 2\pi/L$, with $L = 500$ Mpc (box length), and N_k is the number of points used in our discretised Fourier space in each dimension, i.e, $N_k = 256$. We avoided small values of k corresponding to scales close to the box size as they are highly affected by the

https://doi.org/10.1017/S1743921322001788 Published online by Cambridge University Press

sample variance due to the small number of k configurations. At the other extreme we set the upper limit on k to $k_F N_k/3$ to avoid very high values of k where the estimator is expected to perform poorly (Sefusatti et al. (2016)). This also ensures that we are below the fundamental limit set by the Nyquist frequency $k_{Nyq} = k_F N_k/2$.

3. Results

3.1. *Variance due to differing realizations*

This section shows only our main results. We performed a normalized polyspectra analysis of variance on the results of the L-PICOLA simulations. We chose to work mainly with L-PICOLA because of its accuracy and speed. Our comparative analysis with RAMSES showed an excellent agreement of less than 1% at all scales. Therefore, we can have high confidence that the non-Gaussian polyspectra normalized by the Gaussian polyspectra are very well represented by L-PICOLA.

To remove the variance arising from differing realizations of the large-scale structure, we determine the average normalized polyspectra by averaging over all 5 realizations for each model. As discussed earlier, we consider only symmetrical configurations of the wavenumbers in this work, so we may treat all polyspectra as depending upon a single value of k. In this way, we can attempt to cancel out the contribution to these polyspectra arising from the non-linear structure formation, which will vary in each realization. We then use the various realizations to estimate the minimum and maximum values for the normalized polyspectra, which we will refer to as the variance around our averaged normalized polyspectra.

In the following plots (Fig. 2 and Fig. 3) of this section, we show the results of the analysis of variance in the polyspectra for models NG2 (left column), and NG4 (right column) for our frequencies f1, f2, and f3 (first, second and third row, respectively) at $z = 0$. In all plots, the red solid line represents the NG2 model, and the blue solid line represents the NG4 model, while the light grey strip will show the level of variance around these lines.

3.1.1. Bispectra

The results for the averaged normalized bispectra, at $z = 0$, are given in Fig. 2. The clearest thing to be seen in these results is that the variance increases as amplitude and frequency increase (see Fig. 1). The larger variance is given by the NG4f3 model, which shows a variance of $\sim 16\%$ (without considering the larger scales), and the rest of the models are below $\sim 10\%$. On the other hand, we can see a little scale dependence in the NG4f2 model, with a larger deviation from the equality line (black dashed line in the figure). We also can see a similar scale dependence, but smaller, in the NG2f2 model. These models presents a feature that we refer to as a "dip" between $k \sim 0.2$ h/Mpc and $k \sim 0.8$ h/Mpc. Apart from the presence of this "dip" in the models using frequency f2, there is little evidence of deviation from the Gaussian case in the other models.

This leads to an interesting conclusion regarding the sensitivity of the bispectrum to this form of non-Gaussianity, which is by construction a symmetric oscillatory correction to an underlying Gaussian. This does not violate the property of Gaussian distributions that the odd-n moments are identically zero. Thus the 3-point correlation function (whose k-space analog is the bispectrum under consideration here) in perfect Gaussian conditions would vanish. Normally the presence of non-linear structure induces non-Gaussianities such that the bispectrum is not zero. What we apparently find, however, is that structure formation driven from this primordial non-Gaussianity imprints a marginally statistically significant signal in the bispectrum at the limit of detectability above the "noise" resulting from non-linearities, at least for this choice of parameters.

https://doi.org/10.1017/S1743921322001788 Published online by Cambridge University Press

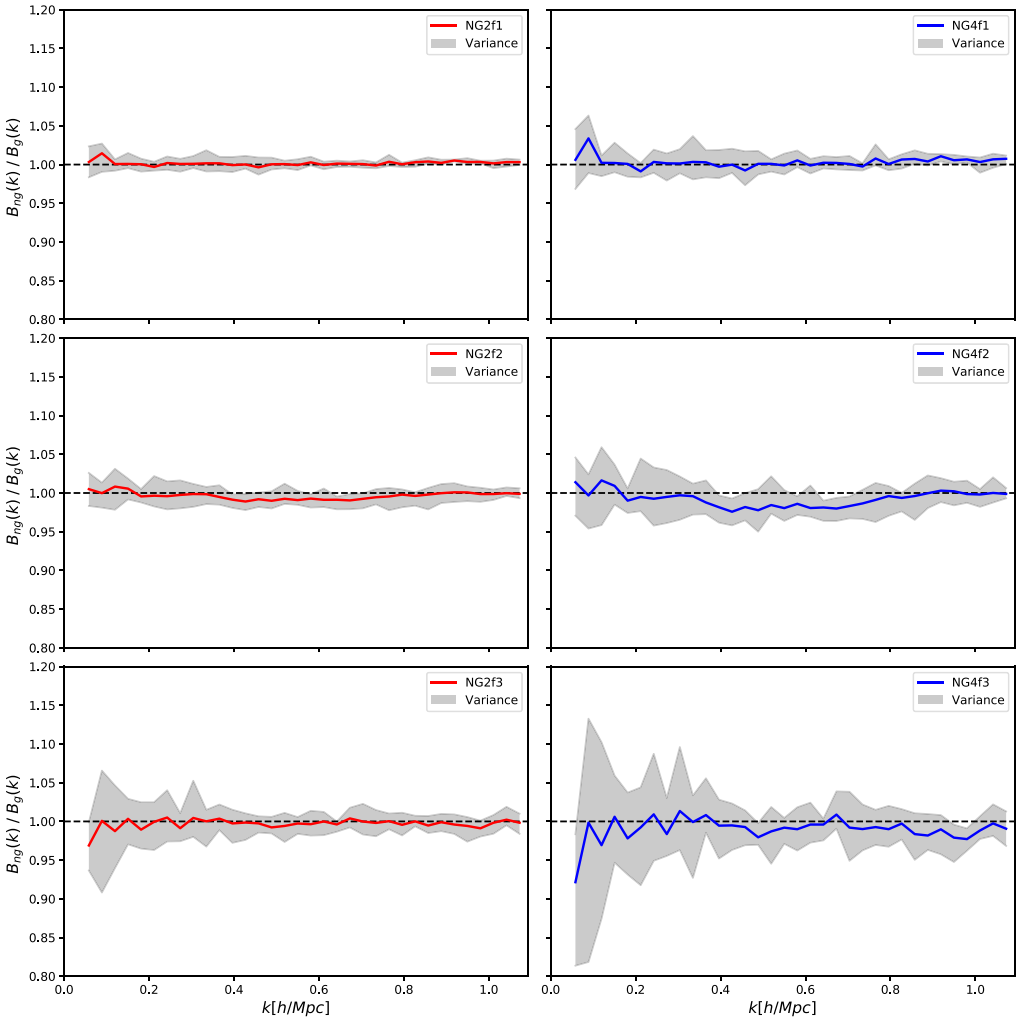

Figure 2. Non-Gaussian bispectra normalised with respect to the Gaussian bispectra, at $z = 0$. *Left column*: NG2 models, *right column*: NG4 models. *Top row*: frequency f1, *middle row*: frequency f2, *bottom row*: frequency f3. The red and blue lines indicate the average normalized bispectra (for the NG2 and NG4 models respectively) while the light grey strip shows the degree of variance around this average arising from the individual realizations.

3.1.2. Trispectra

In the case of the trispectrum, we have considered working with quadrilaterals in equilateral configurations, which are generally folded in 3-dimensional space. Then we take $k_1 = k_2 = k_3 = k_4$, where we have not set a restriction on the other two additional degrees of freedom. On the other hand, it is worth remembering that the trispectrum is the analog in Fourier space of the connected 4-point correlation function, where the disconnected parts are given by-products of the power spectra, which correspond to the disconnected 2-point correlation functions (Verde & Heavens (2001)). Thus, including for the equilateral configuration considered here, the trispectrum is an independent statistical measure that goes beyond the power spectrum.

The results for the averaged normalized trispectra, at $z = 0$, are shown in Fig. 3. Here, we must take into consideration that we have a smaller range of values shown on the

https://doi.org/10.1017/S1743921322001788 Published online by Cambridge University Press

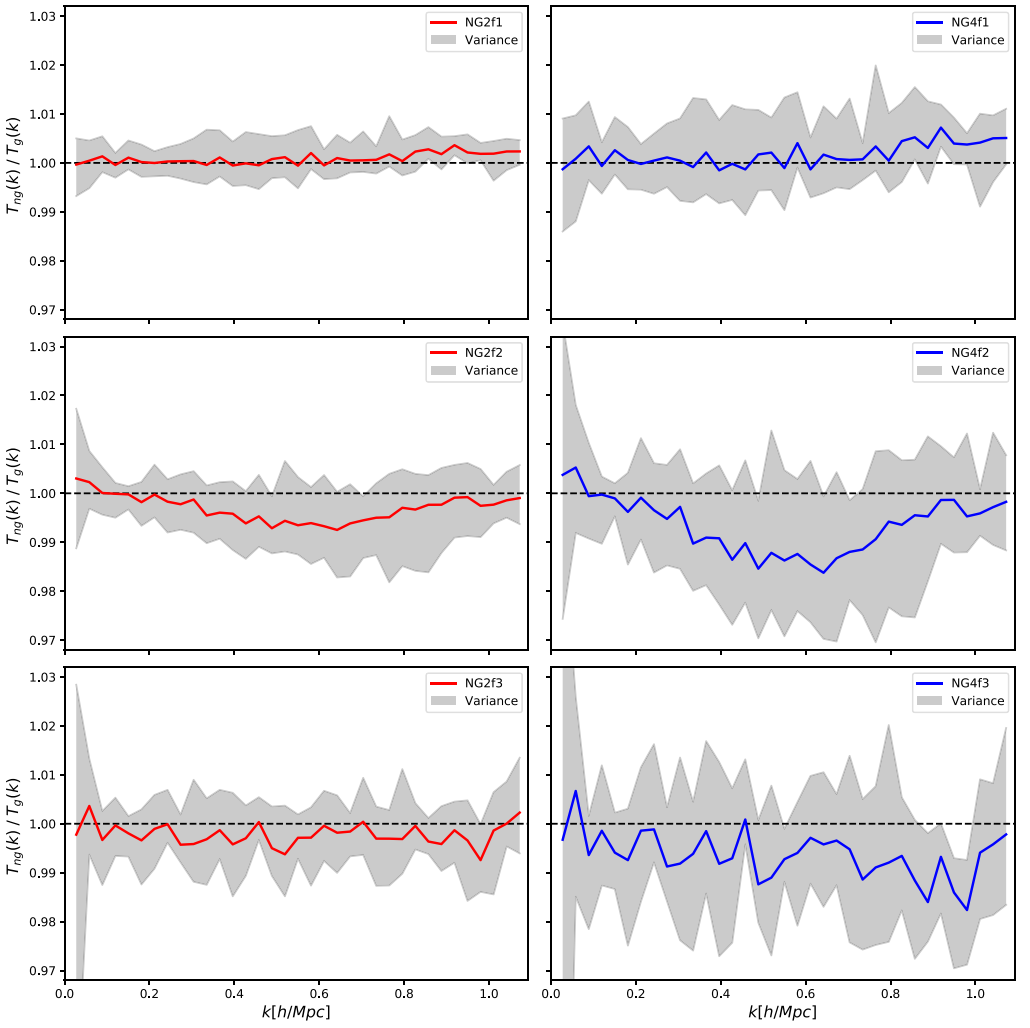

Figure 3. Non-Gaussian trispectra normalised with respect to the Gaussian trispectra, at $z = 0$. The panels are as for Fig. 2.

vertical axes of the bispectra of Fig. 2. Then, we see a lower sensitivity in all models if we compare it to the bispectra.

For models of frequency f2, interestingly, we have again a "dip" in a similar range ($0.3 \lesssim k \lesssim 0.9$ h/Mpc) to that shown in the bispectra of frequency f2. In the case of the other frequencies, there is very little obvious scale dependence, as seen for the bispectra. The f3 models show some indications of suppression of the trispectrum across all accessible scales, with the smaller scales more suppressed in the NG4f3 model. It is also worth pointing out that the variance over realizations in the trispectra is lower than seen for the bispectra, being at most 4% for the NG4f3 model. However, this level of variance implies that these results are not statistically significant.

4. Conclusions

Our study has focussed on analyzing the low-n correlation functions (specifically their Fourier space analogs: the bispectrum and trispectrum) to search for possible signatures of this type of primordial non-Gaussianity at low redshift. We have found that using

https://doi.org/10.1017/S1743921322001788 Published online by Cambridge University Press

the averaged normalized polyspectra, there are some scale-dependent deviations from the Gaussian model, which may be at the limit of detectability. Interestingly, the most significant signal is seen in the bispectrum for the models with frequency f2. For all other models, even the most extreme case NG4f3 we see only sub-percent deviations from Gaussianity, well within the sample variance of differing non-linear realizations. Thus our best-case scenario, given by the model NG4f2 shows deviations from Gaussianity in the bispectrum at the level of 2%. The statistical significance of this deviation by considering sample variance is rather marginal but may be detectable if sufficient precision can be obtained.

It is also noteworthy that the signal present in the frequency f2 models is a scale-dependent suppression of the symmetric bispectrum and the symmetric trispectrum. Consideration of non-symmetric configurations in k-space may well lead to further insights, and a more extensive exploration of the parameter space may uncover models whose late-time n-point correlation functions show more significant deviations of a similar form. Such an exploration may help to disentangle the precise relationship between the non-Gaussian frequency and the scale dependence of the deviations from Gaussianity.

The authors acknowledge financial support from FONDECYT Regular No. 1181708. GP thanks the Postgrado en Astrofísica program of the Instituto de Física y Astronomía of the Universidad de Valparaíso for funding.

References

Celoria M., & Matarrese S. 2018, *J. Cosmology Astropart. Phus.*, 039
Chen X., Palma G. A., Scheihing H. B., & Sypsas S. 2018a, *Phys. Rev. D.*, 98, 083528
Chen X., Palma G. A., Scheihing H. B., & Sypsas S. 2018, *Phys. Rev. Lett.*, 121, 161302
Hahn, O. & Abel, T. 2011, *MNRAS*, 415, 3
Howlett C., Manera M., & Percival W. J. 2015, *Astronomy and Computing*, 12, 109
Sefusatti E., Crocce M., Scoccimarro R., & Couchman H. M. P. 2016, *MNRAS*, 460, 3624
Teyssier R. 2002, *A&A*, 382, 412
Verde L., & Heavens A. F. 2001, *ApJ*, 553, 14
Villaescusa-Navarro F. 2018, Pylians: Python libraries for the analysis of numerical simulations (ascl:1811.008)

Discussion

Anonymous: The use of L-PICOLA seems to be quite useful in the field of cosmological simulation, can you tell us a little more about it?

Greco: L-PICOLA is a is a dark matter halo catalog generator known as mocks catalogues. The main advantage of using such codes is that they can evolve a dark matter distribution from early times to the present day much faster than a full non-linear N-body simulation.

https://doi.org/10.1017/S1743921322001788 Published online by Cambridge University Press