

Biasing and the search for primordial non-Gaussianity beyond the local type

Jérôme Gleyzes^{▼,★}, Roland de Putter[▼], Daniel Green[◆], Olivier Doré^{▼,★}

[▼] *California Institute of Technology, Pasadena, CA 91125*

[★] *Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109*

[◆] *Department of Physics, University of California, Berkeley, CA 94720, USA*

Abstract

Primordial non-Gaussianity encodes valuable information about the physics of inflation, including the spectrum of particles and interactions. Significant improvements in our understanding of non-Gaussianity beyond Planck require information from large-scale structure. The most promising approach to utilize this information comes from the scale-dependent bias of halos. For local non-Gaussianity, the improvements available are well studied but the potential for non-Gaussianity beyond the local type, including equilateral and quasi-single field inflation, is much less well understood. In this paper, we forecast the capabilities of large-scale structure surveys to detect general non-Gaussianity through galaxy/halo power spectra. We study how non-Gaussianity can be distinguished from a general biasing model and where the information is encoded. For quasi-single field inflation, significant improvements over Planck are possible in some regions of parameter space. We also show that the multi-tracer technique can significantly improve the sensitivity for all non-Gaussianity types, providing up to an order of magnitude improvement for equilateral non-Gaussianity over the single-tracer measurement.

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 2 | Modeling the halo power spectrum | 3 |
| 2.1 | Scale-dependent bias from primordial non-Gaussianity | 3 |
| 2.2 | Non-linear and non-local halo bias | 4 |
| 2.2.1 | The ingredients | 5 |
| 2.2.2 | Tree level | 6 |
| 2.2.3 | 1-loop | 6 |
| 2.2.4 | The kernels | 8 |
| 2.2.5 | The halo-halo power spectrum | 9 |
| 3 | Forecasting methodology: f_{NL} constraints from the halo power spectrum | 9 |
| 4 | Equilateral non-Gaussianity | 12 |
| 4.1 | Optimistic forecast, b_δ and f_{NL} only | 12 |
| 4.2 | Including non-local bias | 13 |
| 4.3 | Including non-linearities | 14 |
| 4.4 | Summary | 15 |
| 5 | Beyond single-field inflation | 15 |
| 5.1 | Particle physics and the squeezed limit | 15 |
| 5.2 | Constraints on $f_{\text{NL}}^{(\Delta)}$ ($\Delta = 0 - 2$) | 16 |
| 5.3 | Constraints on quasi-single field inflation | 17 |
| 6 | Multi-tracer technique | 20 |
| 6.1 | The set-up | 20 |
| 6.2 | The results | 22 |
| 7 | Summary and Conclusions | 25 |
| A | Calculating the halo power spectrum | 27 |

1 Introduction

Understanding the physics of the very early universe is one of the central goals of modern cosmology. Of particular interest is physics responsible for the creation of the initial fluctuations that seeded structure formation. Primordial non-Gaussianity (PNG) is a powerful probe as it is sensitive to non-linear evolution of the fluctuations back to the time when the modes were created. Different models of the early universe can produce qualitatively different predictions for the deviations from Gaussianity, meaning a detection or new upper limits will help determine the mechanism responsible. Constraints on primordial non-Gaussianity from the CMB [1] currently provide the best limits on all types of non-Gaussianity, but fall short of some well-motivated targets. Significant improvements are possible with large-scale structure (LSS) and can potentially reach some of these thresholds [2].

Scale-dependent halo bias is a powerful probe of primordial non-Gaussianity [3–6]. In particular, scale-dependent bias is sensitive to the squeezed-limit bispectrum of primordial fluctuations. The single-field consistency relations [7, 8] state that in any single-field model of inflation, the dependence of this bispectrum on the wave vector of the long mode must be of the form (see also [9–13] for more detailed discussions on consistency relations in the large-scale structure)

$$\langle \varphi(\mathbf{k}_L)\varphi(\mathbf{q}_1)\varphi(\mathbf{q}_2) \rangle' \propto k_L^2 P_p(k_L) P_p(|\mathbf{q}_1 + \mathbf{q}_2|/2) \cdots + \mathcal{O}(k_L^4), \quad (1.1)$$

where φ is the primordial Newtonian potential with power spectrum P_p . This primordial potential is related in linear theory to the Newtonian potential during matter domination by $\Phi(q) = D(z)T(k)\varphi(q)$ where $D(z)$ is the linear growth factor of matter fluctuations, normalized such that $D(z) = (1+z)^{-1}$ during matter domination and $T(q)$ is the transfer function, which goes to unity as $q \rightarrow 0$. The important feature is that the small-scale power is modulated by the primordial long-wavelength Newtonian potential. As a result, any scale-dependent bias $b_{\text{NG}}(k)$ for the halo density field due to primordial non-Gaussianity, $\delta_h \supset b_{\text{NG}}(k)\delta$, will be inversely proportional to the transfer function, $b_{\text{NG}}(k) \propto 1/T(k)$. This is realized for equilateral type non-Gaussianity, since (see [14])

$$b_{\text{NG}}(q) = 9(b_\delta - 1)f_{\text{NL}}^{\text{Eq}}\Omega_m\delta_c \frac{H_0^2 R_*^2}{D(z)T(q)}, \quad (1.2)$$

where b_δ is the linear bias, Ω_m is the density parameter of matter, δ_c is a critical density that typically appears in peak-background split calculations (see for example [15]), H_0 is the Hubble parameter today and R_* is the Lagrangian radius of the objects considered. On the other hand, this consistency relation can be violated if multiple fields contributed to inflation [16–19]. For example, local type non-Gaussianity, which arises when there are multiple light fields, gives

$$b_{\text{NG}}(q) = 3(b_\delta - 1)f_{\text{NL}}^{\text{Loc}}\Omega_m\delta_c \frac{H_0^2}{D(z)T(q)q^2} \quad (1.3)$$

(in the single-field case, the consistency relation implies that any bias contribution with this $\propto q^{-2}$ scale dependence has to be exactly zero modulo projection effects [20, 21]).

More generally, models where one of the fields is not light also violate the consistency relation (see for example quasi-single field inflation [22]), and give

$$b_{\text{NG}}(q) \propto f_{\text{NL}}^{\Delta} \Omega_m \delta_c \frac{H_0^2}{D(z)T(q)q^2} (qR_*)^{\Delta} \quad (1.4)$$

with $\Delta = 0 - 3/2$.

Constraints from scale-dependent halo bias on these types of primordial non-Gaussianity have been extensively studied in the literature, especially in the case of local non-Gaussianity (e.g. [23, 24]), but also for equilateral [25–27] and quasi-single field non-Gaussianity [28–30]. One potential concern is that the scale-dependent bias effect may in principle be degenerate with contributions to halo bias from other sources than primordial non-Gaussianity, specifically non-linear bias and/or non-local¹ bias (evidence for both of those has been observed in simulations [31, 32]). While based on fundamental principles it is hard to mimic the k_L^{-2} scale-dependence of local-type scale-dependent bias by such contributions, it may be an important effect in the other cases mentioned above. This can potentially severely weaken the constraining power of halo clustering on primordial non-Gaussianity. In particular, for equilateral non-Gaussianity, the scale-dependent bias can be expanded as $b_{\text{NG}}(k) \propto 1/T(k) = 1 + (k/k_{\text{eq}})^2 + \dots$, which at first sight is fully degenerate with the gradient bias expansion $b_{\delta} + b_{k^2} (R_*k)^2 + b_{k^4} (R_*k)^4 + \dots$. In fact, considering only the k^2 terms, and assuming b_{k^2} to be of order unity, suggests an error floor on $f_{\text{NL}}^{\text{Eq}}$ of order 1000 [25, 33]. On the other hand, the effects of equilateral non-Gaussianity and non-local bias should not be exactly the same, as they come in with a different characteristic scale. For primordial non-Gaussianity, the typical scale is k_{eq} ($\sim 1.6 \times 10^{-2} h/\text{Mpc}$), while it is the Lagrangian size of the halos of interest for the gradient bias ($R_*^{-1} \gtrsim 10^{-1} h/\text{Mpc}$). In other words, one expects the degeneracy to not be exact when taking into account higher order terms.

In this paper, we study in more detail the constraints that can be obtained on primordial non-Gaussianity from scale-dependent bias in the presence of non-linear biasing and non-local (gradient) bias. We follow the approach of [33–35] to write the most general bias expansion based on principles of symmetry and the equivalence principle. We then forecast constraints on primordial non-Gaussianity, marginalizing over the effects of the various bias contributions, where for non-linearities, we include all terms up to 1-loop order.

Of particular interest is how information about equilateral and quasi-single field-type non-Gaussianity is encoded in galaxy power spectra. We study dependence on survey volume, number density, use of multiple tracers (as proposed in [36]) to understand the reach and limitations of scale-dependent bias for non-local non-Gaussianity. After marginalizing over bias, only quasi-single field inflation with $\Delta < 0.3$ and local non-Gaussianity show the potential for significant improvements over Planck [1] with realistic configurations. Nevertheless, we show that the multi-tracer technique allows for significant improvements over single tracers and could extend the potential reach of this large-scale structure and scale-dependent bias.

The paper is outlined as follows: In Section 2, we detail the model for the halo power spectra in terms of a biasing model. In Section 3, we explain our forecasting methodology.

¹By this we mean that δ_h can depend on spatial derivatives of δ .

In Section 4, we focus on equilateral non-Gaussianity, because there one might expect degeneracy with other bias terms to be the strongest. We then turn to non-Gaussianity of the types that violate the consistency relations in Section 5 (i.e. require multiple fields in the early universe) and quantify the degradation to parameter constraints due to non-linear and non-local biasing. In Section 6, we discuss the multi-tracer technique and demonstrate the potential for improving constraints on equilateral non-Gaussianity. We conclude in Section 7.

The reader that is familiar with non-local, non-linear bias and/or MCMC forecasting may want to go directly to the results, starting in Section 4.

2 Modeling the halo power spectrum

In this section, we describe the model of halo clustering. Our discussion here follows the systematic, general approach of [33, 34, 37–42] (see also the review [15] and references therein). We will consider the halo power spectrum in configuration space, i.e. we will not include redshift space distortions, which would make the power spectrum anisotropic, and would significantly complicate the analysis.

2.1 Scale-dependent bias from primordial non-Gaussianity

We start with the leading order, local contributions to the halo overdensity δ_h , including the scale-dependent bias from primordial non-Gaussianity, which is our signal,

$$\delta_h = b_\delta \delta + f_{\text{NL}} b_\Psi \Psi + \epsilon_0. \quad (2.1)$$

Here, δ is the matter overdensity and b_δ the Gaussian, linear bias. For non-Gaussian initial conditions, there can be a modulation of the small-scale variance of initial perturbations by a long-wavelength fluctuation,

$$\Psi(\mathbf{q}) \equiv \left(\frac{q}{\mu}\right)^\Delta \varphi(\mathbf{q}), \quad (2.2)$$

where φ is the primordial metric perturbation, related to the primordial curvature perturbation, and μ an arbitrary energy scale. The term $b_\Psi \Psi$ above describes the simplest form of bias due to this primordial non-Gaussianity (see [33, 43] for additional refinements) and has been tested in simulations [44–46]. Finally, ϵ_0 is a stochastic white noise contribution, which is assumed to be uncorrelated with the other terms.

The primordial metric fluctuation can be related to the present-day matter fluctuation by,

$$\delta(\mathbf{q}) = \mathcal{M}(q) \varphi(\mathbf{q}), \quad \text{with} \quad \mathcal{M}(q) \equiv \frac{2q^2 T(q) D(z)}{3\Omega_m H_0^2}. \quad (2.3)$$

We can then write the non-Gaussian contribution as a scale-dependent bias,

$$\delta_h(\mathbf{q}) = (b_\delta + b_{\text{NG}}(q)) \delta(\mathbf{q}) + \epsilon_0, \quad b_{\text{NG}}(q) = b_\Psi \mathcal{M}^{-1}(q) \left(\frac{q}{\mu}\right)^\Delta. \quad (2.4)$$

Parametrizing the primordial non-Gaussianity (specifically, the squeezed-limit bispectrum) in the conventional way, and applying a simple halo model, the scale-dependent bias can

be written in terms of a non-Gaussianity parameter² [4, 26, 49–51],

$$b_{\text{NG}}(q) = 2 f_{\text{NL}}^{\text{Loc}} (b_\delta - 1) \delta_c \mathcal{M}^{-1}(q) \quad (\text{local}), \quad (2.5)$$

$$b_{\text{NG}}(q) = 6 f_{\text{NL}}^{\text{Eq}} (b_\delta - 1) \delta_c (q R_*)^2 \mathcal{M}^{-1}(q) \quad (\text{equilateral}), \quad (2.6)$$

$$b_{\text{NG}}(q) = 6 f_{\text{NL}}^{(\Delta)} (b_\delta - 1) \delta_c (q R_*)^\Delta \mathcal{M}^{-1}(q) \quad (\text{generic exponent } \Delta \in [0, 2]), \quad (2.7)$$

where we recall that $\delta_c (= 1.686)$ is the critical overdensity for spherical collapse, and R_* is the Lagrangian radius of the halos of interest.³ The amplitudes of non-Gaussianity, f_{NL} , defined in eqs. (2.5)–(2.7) are determined by the amplitudes of the squeezed limit of the bispectrum. Constraints from the CMB often normalize f_{NL} by the amplitude of the bispectrum in the equilateral configuration, which can be non-trivially related to the amplitude in the squeezed limit. This can be important for comparing limits from scale-dependent bias with limits from Planck, as we will see in the case of quasi-single field inflation (Section 5.3).

The above expressions for scale-dependent bias are based on the squeezed-limit behavior of the primordial bispectrum. This means their range of validity is technically limited to $q/k_S \sim q R_* \ll 1$, where q is the long mode, and k_S the wave vector of the short modes. For a given halo type, this places a requirement on the maximum wave vector included in the analysis, k_{max} . In order to maximize the information from scale-dependent bias, we will push to relatively large values, $k_{\text{max}} R_* \lesssim 1$, but one has to keep in mind that for modes that do not satisfy $q R_* \ll 1$, the actual form of the scale-dependent bias may be modified. We will come back to the optimization of k_{max} in Section 4.2.

2.2 Non-linear and non-local halo bias

In this section, we will explicitly go step by step through the calculations of the halo-halo power spectrum in the presence of PNG. The reader not interested in those technical details is encouraged to skip to either the forecasting method Section 3, or even directly to the results in Section 4.

Specifically, we want to compute

$$\langle \delta_h(\mathbf{k}) \delta_h(\mathbf{q}) \rangle = (2\pi)^3 \delta^{(D)}(\mathbf{k} + \mathbf{q}) P_{hh}(k), \quad (2.8)$$

using the methodology of standard perturbation theory (SPT). We refer the reader to [53] for definitions and notations.

On top of the SPT contributions up to 1-loop, we will also include the most general bias model for δ_h . To do so, we will allow δ_h to depend on any term that is permitted by the symmetries of the system, e.g. rotational symmetry and the equivalence principle. This means that we allow δ_h to depend on derivatives of the matter overdensity δ , leading to terms like $\nabla^2 \delta$, known as non-local bias. We will also include non-linear bias, i.e. we will allow δ_h to depend on terms like δ^2 . Our derivation will follow closely the work of

²Interestingly, these f_{NL} parameters in a given volume may be biased relative to the true statistical f_{NL} [47] but the parameter Δ is robust to such effects [48].

³For the quasi-single field case, the conventional prefactor is not 6, but a more complicated function of $\Delta \in [0, 3/2]$, see [52]. When comparing our results to the existing QSF constraints in Section 5.2, we will use the full normalization. However, when drawing a link from local to equilateral configurations, we will use the normalization (2.7).

[34] (and also [39]), which computed $P_{hm} \propto \langle \delta(\mathbf{k})\delta_h(\mathbf{q}) \rangle$ for gaussian initial conditions, to which we will add results of [33] regarding the PNG.

In the expansion of δ_h , the only term in Ψ that we will keep is the linear one (i.e. the first line of eq (2.34) in [33]), the other (loop) terms being much smaller for the scales of interest. Our expansion is going to be in $\delta_G(\mathbf{k}) = \mathcal{M}(k)\varphi_G(\mathbf{k})$, where φ_G is a purely gaussian variable⁴. Let's first discuss the structure of the terms that we will get. As in [33], we will conduct the expansion with diagrams for an easier representation.

2.2.1 The ingredients

The goal of this section is to present what are the terms that appear in the bias when considering that it is in principle non-local and non-linear. We will go through order by order (in δ_G).

- **Linear in δ_G**

Those terms will be represented by a line. External lines should be weighted by a linear δ_h , given by

$$\delta_h(q) = \left[b_\delta + f_{\text{NL}} \frac{b_\Psi}{\mathcal{M}(q)} \left(\frac{q}{\mu} \right)^\Delta + b_{q^2} q^2 R_*^2 + b_{q^4} q^4 R_*^4 + \dots \right] \delta(\vec{q}). \quad (2.9)$$

The $b_{q^{2n}} q^{2n} R_*^{2n}$ terms come from Fourier transforming the dependence of δ_H on $(\nabla^2)^{(n)}\delta$. The dots signify that in principle there are other terms in this gradient expansion. Internal lines on the other hand are necessarily δ_G . Contracting two lines gives a power spectrum P_G (represented by a black dot), with the appropriate bias weights. There is also a stochastic term, ϵ_0 , that is allowed by the symmetries. It does not correlate with δ_G and gives a noise contribution to the power spectrum.

- **Quadratic in δ_G**

These will be represented by a striped circle with one line going in and two going out. The circle will have a weighting function $G(\mathbf{k} - \mathbf{p}, \mathbf{p})$ (where \mathbf{k} is the momentum of the line going in) that will depend on the exact nature of the quadratic term. Since our operators are expressed in terms of δ and not δ_h , internal lines can only carry the standard perturbation theory kernel of second order for δ , F_2 . The restriction to the kernel F_2 will be shown by a circle with a grid. If the line is an external line, then two additional types of terms are possible: the quadratic overdensity δ^2 and the square of tidal tensor $s_{ij}^2 = (\nabla_i \nabla_j \Phi_\delta)^2 - \frac{\delta^2}{3}$, where $\Phi_\delta \equiv \nabla^{-2}\delta$. Similarly to the linear case, there is a stochastic term $\epsilon_\delta \delta$ that appears. We will see that this term can be absorbed in the definition of the linear bias (see App. A).

- **Cubic in δ_G**

These will be represented by a striped square with one line going in and three going out. The square will have a weighting function $R(\mathbf{q} - \mathbf{p}_1 - \mathbf{p}_2, \mathbf{p}_1, \mathbf{p}_2)$ (where \mathbf{q} is the momentum of the line going in) that will depend on the exact nature of the cubic

⁴In principle, the non-linear relation between the matter overdensity δ and δ_G due to the PNG will show up as a modification to the kernels of standard perturbation theory. However we have checked that those modifications are negligible compared to the effect of the scale-dependent bias.

term. At one loop, on top of the standard perturbation theory kernel of third order for δ , F_3 , we have four additional contributions: δ^3 , δs_{ij}^2 , $\nabla_i \nabla_j \Phi_\delta \nabla_j \nabla_k \Phi_\delta \nabla_k \nabla_i \Phi_\delta$. The last one, denoted $\text{Tr} [\Pi^{(1)} \Pi^{(2)}]$ in [33] (it is related to the variable st in [34] although it is not exactly the same). This term is more subtle. At first order, δ and the velocity divergence θ are the same up to a normalization factor, which means that only a single variable is needed to describe them. At second order they are not the same, but their difference (called $\Pi_{ij}^{(2)}$ in [33]) is not a new variable, i.e. it can be expressed in terms of the expressions in the paragraph right above (see [33]). Only when multiplied with a first order quantity such as $\Pi_{ij}^{(1)} = \nabla_i \nabla_j \Phi_\delta$ does it become a new, independent variable, which is why it only starts appearing at cubic order. We will give more explicit expressions when looking directly at the kernels. There are also mixed stochastic terms, $\epsilon_{\delta^2} \delta^2$ and $\epsilon_{s^2} s_{ij}^2$, which give divergent contributions and therefore will not appear in the final (renormalized) result.

Now that we have the ingredients, it is a matter of getting all the possible terms in the expansion of P_{hh} up to 1-loop. That will be represented by diagrams, similarly to standard perturbation theory [53]. The terms in P_{hh} are obtained by contracting the “outgoing” lines from two ingredients in the list above. The contractions, that we will denote with a black dot, come with a linear power spectrum P_G and imposes the sum of the momenta of the lines to be zero. We will illustrate this below. Note that in principle, we should be working with renormalized quantities [39] from the start. Instead, we will work with the bare quantities, and keep only the non-divergent parts in our final result.

2.2.2 Tree level

Only the terms that are at most linear in δ_G can appear here. Thus, we have to contract two external lines as in eq. (2.9), as well as two stochastic terms ϵ_0 . This gives

$$P_{hh}^{\text{Tree}}(q) \equiv b_{\text{full}}^2 P_G(q) + P_{\epsilon_0}, \quad (2.10)$$

where we have defined

$$b_{\text{full}}(q) \equiv b_\delta + f_{\text{NL}} \frac{b_\Psi}{\mathcal{M}(q)} \left(\frac{q}{\mu} \right)^\Delta + b_{q^2} q^2 R_*^2 + b_{q^4} q^4 R_*^4 + \dots \quad (2.11)$$

and we denoted the variance of the ϵ_0 by P_{ϵ_0} .

2.2.3 1-loop

Here, there are three types of diagrams that contribute. 1) a second-order vertex contracted with another second-order vertex (which is similar to the usual P_{22} of SPT). 2) a third order vertex with one loop and one contraction with a linear δ_h (the equivalent of P_{13}). There is a third type, that does not appear in SPT. This is because our quadratic operators are constructed from the full δ , and not the first order $\delta_G^{(1)}$ as in SPT. Therefore, we can have a second-order vertex contracted with the SPT second-order vertex F_2 and a linear δ_h , which we call type 3. The structures of the diagrams are shown in Fig. 1

Let us now compute the contribution of the different types of diagrams.

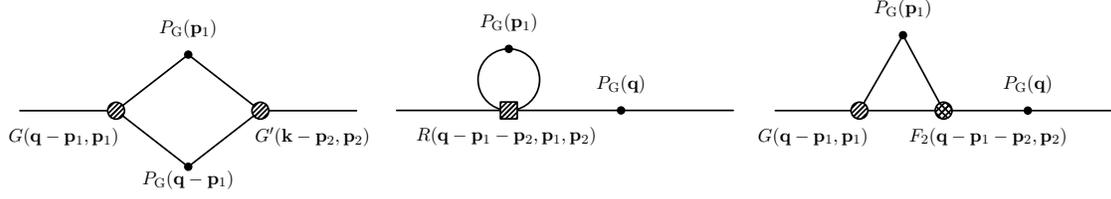


Figure 1: *Left:* Type 1) two second-order vertices. *Middle:* Type 2) a cubic vertex and a linear δ_h . *Right:* Type 3) a second-order vertex and one SPT kernel F_2 contracted with δ_h . The black dots represent the gaussian spectrum P_G , and each external line carries a factor b_{full} (eq. (2.11)).

- **Type 1)**

The contractions impose $\mathbf{p}_1 = -\mathbf{p}_2$, so that there is only one integral over momentum (which is expected because we are looking at 1-loop). The diagram and its symmetrical partner therefore give

$$P_{hh}(q) \supset \int_{\mathbf{p}} G(\mathbf{q} - \mathbf{p}, \mathbf{p}) [G'(\mathbf{p} - \mathbf{q}, -\mathbf{p}) + G'(-\mathbf{p}, \mathbf{p} - \mathbf{q})] P_G(|\mathbf{q} - \mathbf{p}|) P_G(p). \quad (2.12)$$

We use the shorthand notation $\int_{\mathbf{p}} = \int d^3p$. If $G = G' = F_2$, we recover P_{22} , the 1-loop term in SPT obtained when contracting two second-order $\delta_G^{(2)}$.

- **Type 2)**

The contractions impose $\mathbf{p}_2 = -\mathbf{q}$. The diagram and its symmetrical partners (2 others, depending on which leg of the cubic vertex is contracted with the second external line) therefore give

$$P_{hh}(q) \supset P_G(q) \int_{\mathbf{p}} P_G(p) [R(-\mathbf{p}, \mathbf{p}, -\mathbf{q}) + R(-\mathbf{p}, -\mathbf{q}, \mathbf{p}) + R(-\mathbf{q}, -\mathbf{p}, \mathbf{p})]. \quad (2.13)$$

Again, if $R = F_3$, this is P_{13} , the 1-loop term in SPT obtained when contracting a first-order $\delta_G^{(1)}$ with a third-order $\delta_G^{(3)}$.

- **Type 3)**

This type is slightly more involved. However, following the basic rules above, the diagram and its symmetrical partners (3 others, depending on where the second vertex is connected and where the second external line connects to the second vertex) give

$$P_{hh}(q) \supset 2P_G(q) \int_{\mathbf{p}} P_G(p) F_2(-\mathbf{p}, \mathbf{q}) [G(\mathbf{q} - \mathbf{p}, \mathbf{p}) + G(\mathbf{p}, \mathbf{q} - \mathbf{p})], \quad (2.14)$$

where we have used that $F_2(\mathbf{k}_1, \mathbf{k}_2) = F_2(-\mathbf{k}_1, -\mathbf{k}_2) = F_2(\mathbf{k}_2, \mathbf{k}_1)$.⁵

⁵In principle, there is another diagram where the external δ_h is connect to one branch of the first kernel G while the two branches of the second kernel F_2 are connected together. However, this leads to a vertex $F_2(\mathbf{p}, -\mathbf{p})$ which is zero.

We therefore have the ingredients as well as the structure of the bias expansion. The only thing left is to assign the proper kernel to each ingredient, and then we will be able to compute P_{hh} at 1-loop.

2.2.4 The kernels

Let us now translate the non-linear bias terms to quadratic $G(\mathbf{k}, \mathbf{p})$ and cubic $R(\mathbf{k}, \mathbf{p}, \mathbf{q})$ kernels. In order to compactly write the expressions, it is convenient to introduce the cosine between \mathbf{k}_i and \mathbf{k}_j

$$\mu_{ij} \equiv \frac{\mathbf{k}_i \cdot \mathbf{k}_j}{k_i k_j}, \quad (2.15)$$

as well as another cosine, between $\mathbf{q} - \mathbf{p}$ and \mathbf{p}

$$\mu_- \equiv \frac{(\mathbf{q} - \mathbf{p}) \cdot \mathbf{p}}{|\mathbf{q} - \mathbf{p}| p}. \quad (2.16)$$

- **Quadratic operators:** We have two kernels, δ^2 and $s_{ij}^2 = (\nabla_i \nabla_j \Phi_\delta)^2 - \frac{\delta^2}{3}$. We do not consider $\delta\Psi$ and $\epsilon_\Psi\Psi$ as they are much smaller (suppressed by small non-Gaussianity and loop order).

We can therefore associate the kernels: $\delta^2 \rightarrow 1$ and $s_{ij}^2 \rightarrow G(\mathbf{q} - \mathbf{p}, \mathbf{p}) = \mu_-^2 - 1/3$, where μ_- is defined in eq. (2.16).

On top of it, there is the standard SPT kernel, F_2 (see App. A for details).

- **Cubic operators** We have the following kernels (they are symmetrized):

$$\begin{aligned} \delta^3 \rightarrow 1, \quad \delta s_{ij}^2 \rightarrow R(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) &= \frac{1}{3} (\mu_{12}^2 + \mu_{23}^2 + \mu_{13}^2 - 1), \\ \nabla_i \nabla_j \Phi_\delta \nabla_j \nabla_k \Phi_\delta \nabla_k \nabla_i \Phi_\delta &\rightarrow \mu_{12} \mu_{23} \mu_{13}, \end{aligned} \quad (2.17)$$

with μ_{ij} is defined in eq. (2.15). The kernel for $\text{Tr} [\Pi^{(1)} \Pi^{(2)}]$ reads

$$\text{Tr} [\Pi^{(1)} \Pi^{(2)}] \rightarrow R(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) = \frac{1}{3} \left[\left(\frac{\mathbf{k}_1 \cdot (\mathbf{k}_2 + \mathbf{k}_3)}{k_1 |\mathbf{k}_2 + \mathbf{k}_3|} \right)^2 [G_2(\mathbf{k}_2, \mathbf{k}_3) - F_2(\mathbf{k}_2, \mathbf{k}_3)] + \text{perms} \right], \quad (2.18)$$

with G_2 the second-order kernel for the velocity (see [53]). Plugging the expressions for F_2 and G_2 , one finds

$$\text{Tr} [\Pi^{(1)} \Pi^{(2)}] \rightarrow R(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) = \frac{2}{21} \left[\left(\frac{\mathbf{k}_1 \cdot (\mathbf{k}_2 + \mathbf{k}_3)}{k_1 |\mathbf{k}_2 + \mathbf{k}_3|} \right)^2 [\mu_{23}^2 - 1] + \text{perms} \right]. \quad (2.19)$$

2.2.5 The halo-halo power spectrum

In this section, we will combine everything to compute P_{hh} at first order in loops and f_{NL} . We have already computed the tree level contribution

$$P_{hh}^{\text{Tree}}(q) = b_{\text{full}}^2 P_{\text{G}}(q) + P_{\epsilon_0}, \quad (2.20)$$

with b_{full} defined in eq. (2.11).

The next step is to remove the divergent quantities in the loop contributions (that would not have been there if we had worked with renormalized quantities from the beginning). Once this is taken into account (see App. A), the expression for the halo power spectrum reads

$$\begin{aligned} P_{hh}(q) &= P_{hh}^{\text{Tree}}(q) + b_{\text{full}}(q)^2 P^{\text{loop}}(q) \\ &+ 4b_{\text{full}}(q) \left\{ \int_{\mathbf{p}} \left[b_{\delta^2} + b_{s^2} \left(\mu_-^2 - \frac{1}{3} \right) \right] [P_{\text{G}}(|\mathbf{q} - \mathbf{p}|) - P_{\text{G}}(p)] P_{\text{G}}(p) \right\} \\ &+ 2 \int_{\mathbf{p}} \left[b_{\delta^2} + b_{s^2} \left(\mu_-^2 - \frac{1}{3} \right) \right]^2 [P_{\text{G}}(|\mathbf{q} - \mathbf{p}|) - P_{\text{G}}(p)] P_{\text{G}}(p) \\ &+ 6b_{\text{full}}(q) b_{\text{III}(2)} P_{\text{G}}(q) \int_{\mathbf{p}} \left(\frac{2}{7} \left[\left(\frac{\mathbf{q} \cdot \mathbf{p}}{qp} \right)^2 - 1 \right] \frac{2\mu_-^2}{3} + \frac{8}{63} \right) P_{\text{G}}(p), \end{aligned} \quad (2.21)$$

where $P^{\text{loop}} = P_{22} + P_{13}$ is the standard one loop contribution to the power spectrum [53]. This expression is similar to the results in [34], with the addition of the PNG. Notice that out of the type 2) and 3) diagrams, only one remains after renormalization (last line). Indeed, when the divergent parts are removed, all the non-zero cubic terms are proportional to each other, as already noted in [34]. This is why we regroup them into a single term, $b_{\text{III}(2)}$ (c.f. Appendix A). However, we cannot the same for the terms in second and third lines of eq. (2.21) because μ_- depends on \mathbf{q} and \mathbf{p} .

To better gauge the size of those different terms, we plot them in Fig. 2 at $z = 1.5$, choosing $f_{\text{NL}} = 1$, $b_{\delta} = 3.6$, $R_* = 3.7 \text{ Mpc}/h$ and all the other bias parameters equal to one. The cosmology is set to that of Planck 2015 [54], and the linear power spectrum is obtained from CAMB [55].

3 Forecasting methodology: f_{NL} constraints from the halo power spectrum

In order to estimate the constraining power of future surveys, we will resort to a Monte Carlo Markov Chain (MCMC), using the python package `emcee` [56]. To this end, we need the likelihood function, i.e. the probability to observe a halo power spectrum \hat{P} given a set of parameters $\boldsymbol{\theta}$. Since we want to focus on the bias expansion of eqs. (2.11) and (2.21), we will fix all the cosmological parameters. Therefore, our set of parameters is given by

$$\boldsymbol{\theta} = \{b_{\delta}, b_{q^2}, b_{q^4}, \dots, b_{s^2}, b_{\delta^2}, b_{\text{III}}, \Delta, P_{\epsilon_0}\}. \quad (3.1)$$

We approximate the likelihood by that of the power spectrum of a *Gaussian* galaxy density field, but using the full non-linear expression for the theory power spectrum, i.e. the

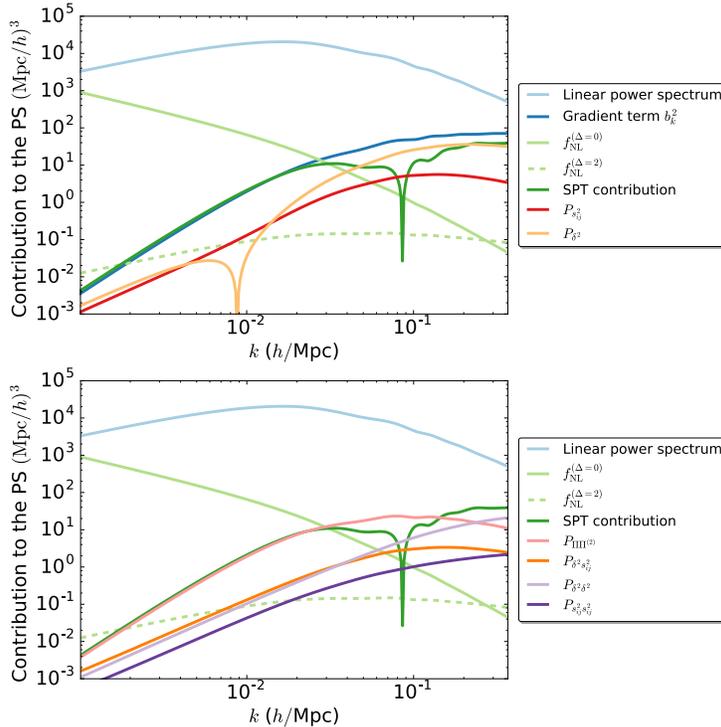


Figure 2: The contributions to P_{hh} compared to the linear power, $P_{\text{lin}} = b_\delta^2 P_G$. In the *top* panel, we show the terms linear in the extra bias coefficients, i.e. $b_{k2} k^2 R_*^2 P_G(k)$. For the *bottom* panel, we also include non-linear contributions, such as $b_{\delta^2}^2 P_{\delta^2 \delta^2}$. There are also terms of the form $(b_{k2} k^2 R_*^2)^2 P_G(k)$, but we choose not to show them here for clarity. The linear bias is chosen to be $b_\delta = 3.6$, the Lagrangian radius $R_* = 3.7 \text{ Mpc}/h$, the non-Gaussianity parameter $f_{\text{NL}} = 1$ and all the other bias parameters are one. We show the scale-dependent bias contribution (linear in f_{NL} only) for $\Delta = 0$ (local) and $\Delta = 2$ (equilateral). At low k , the local ($\Delta = 0$) contribution is the most important, and one can see by eye that no other bias terms can mimic its behavior. However, the $\Delta = 2$ curve is both much smaller than $\Delta = 0$ as well as more prone to degeneracies with non-local and non-linear contributions.

power spectrum $P_{\text{th}}(k, \theta)$ given in eq. (2.21). Given an observed (or in our case fiducial) spectrum $\hat{P}(k)$, this leads to the likelihood,

$$\log \mathcal{L}(\hat{P}(k) | \theta) = -\frac{N_k}{2} \left[\frac{\hat{P}(k)}{P_{\text{th}}(k, \theta)} - \log \left(\frac{\hat{P}(k)}{P_{\text{th}}(k, \theta)} \right) \right], \quad N_k \equiv \frac{k^2 V}{2\pi^2} \Delta k, \quad (3.2)$$

where N_k is the number of modes in a bin $[k, k + \Delta k]$ for a survey of volume V . We have ignored terms that do not vary with θ and therefore will not modify our exploration of the parameter space. The power spectrum thus is not a Gaussian variable, but follows a χ^2 distribution. This is particularly important for large scales, which are crucial for local non-Gaussianity, as the number of modes is low so that the central limit theorem does not apply, cf. [57].

The cosmology is given by a Λ CDM model with parameters given by the latest Planck release [54], and we will not vary them as we want to focus only on the potential degeneracy within the bias expansion. Our parameter estimations are for a survey of volume $V = 100 (\text{Gpc}/h)^3$, with a mean redshift $z_m = 1.5$. We will consider modes going

| Param. | Description | Fiducial | Prior |
|---------------------|--------------------|-------------------|------------------|
| b_δ | Linear bias | $b(M, z)$ | $[0, 10]$ |
| b_{k^n} | Gradient expansion | 0 | $[-4, 4]$ |
| b_{δ^2} | Quadratic bias | 0 | $[-4, 4]$ |
| b_{s^2} | Tidal bias | 0 | $[-4, 4]$ |
| $b_{\text{III}(2)}$ | Evolution bias | 0 | $[-4, 4]$ |
| f_{NL} | Non-Gaussianity | 0 | $[-10^5, 10^5]$ |
| Δ | Type of PNG | $[0, 2]$ | $[0, 2]$ |
| P_{ϵ_0} | Stochastic noise | $1/\bar{n}(M, z)$ | $[0, 2/\bar{n}]$ |

Table 1: The parameters in θ along with their fiducial values and priors. $b(M, z)$ and $\bar{n}(M, z)$ are respectively the bias and the mean density computed from the halo mass function at redshift z for a minimum halo mass M . b_{k^n} denotes generically all the terms in the bias expansions, whose truncation order will be specified later. For Δ , the fiducial value will vary according the type of PNG that we want to constrain.

from $k_{\text{min}} = 2\pi/V^{1/3}$ to $k_{\text{max}} = 1/(2R_*)^6$, with bin width $\Delta k = k_{\text{min}}$. The fiducial linear bias $b_{\delta, \text{fid}}$, the lagrangian radius R_* and galaxy density \bar{n} (which gives the shot-noise and fixes the fiducial value of $P_{\epsilon_0} = 1/\bar{n}$) are determined using the halo mass function and bias at redshift z_{m} [58, 59]. We will use a minimum halo mass of $M_{\text{h, min}} = 10^{13} M_\odot$ (see Section 6, and in particular Fig. 8 for a discussion on this choice), which implies,

$$b_{\delta, \text{fid}} = 3.6 \quad R_* = 3.7 \text{ Mpc}/h \quad \bar{n} = 1.2 \times 10^{-4} (h/\text{Mpc})^3. \quad (3.3)$$

This gives a $k_{\text{max}} = 0.135 h/\text{Mpc}$, well within the linear regime at redshift $z = 1.5$.

We will take flat priors for all the free parameters. For the bias expansion, in particular the gradient terms, we will impose a flat prior between $[-4, 4]$. Indeed, because we explicitly put the scaling with R_* in the spatial derivatives, the terms in front of them should be of order one. This is why the expansion is not completely degenerate with $1/T(k)$ in the scale-dependent bias, which evolves on a different spatial scale given by the matter radiation equality, k_{eq} . We will explore this in more details in the next section.

The fiducial parameter values and prior ranges are listed in Table 1. Although for realistic halos, the non-linear and non-local biases are expected to deviate from zero (see [32, 60–62] for assessment of a subset of these from simulations), we do not expect constraints on primordial non-Gaussianity to be very sensitive to the exact choice of fiducial values, so we set them to zero for simplicity.

The truncation order for the gradient expansion will be determined as the order when adding new gradient terms does not change the constraint on f_{NL} . As we will see, for $k_{\text{max}} = 1/(2R_*)$, this corresponds to including terms up to k^4 .

⁶ k_{max} is chosen so that the truncation of the gradient expansion $qR_* \ll 1$ is sensible.

4 Equilateral non-Gaussianity

We first consider constraints on scale-dependent bias due to equilateral non-Gaussianity, characterized by $f_{\text{NL}}^{\text{Eq}}$. According to eq. (2.6), this type of non-Gaussianity leads to a scale-dependence

$$b_{\text{NG}}(q) \propto \frac{1}{T(q)}, \quad (4.1)$$

where $T(q)$ is the transfer function of matter perturbations, normalized to one when $q \rightarrow 0$. Since the scale-independent part of $b_{\text{NG}}(q)$ is unobservable (because completely degenerate with the linear bias), this means that the signal is dominated by small scales, $q > k_{\text{eq}}$, where $k_{\text{eq}} \approx 0.016h/\text{Mpc}$ is the matter-radiation equality scale. This stands in strong contrast with the more commonly studied halo bias due to local non-Gaussianity, which is dominated by large scales. This can readily be seen in Fig. 2.

Since there is a large number of modes available on these small scales, in principle we expect scale-dependent bias to strongly constrain $f_{\text{NL}}^{\text{Eq}}$. On the other hand, if we expand the scale-dependent bias signal around $q = 0$,

$$b_{\text{NG}}(q) \propto 1 + c_2 \left(\frac{q}{k_{\text{eq}}} \right)^2 + \mathcal{O} \left(\left(\frac{q}{k_{\text{eq}}} \right)^4 \right), \quad (4.2)$$

we find a scale-dependence that will to a certain extent be degenerate with the non-local and even non-linear bias terms discussed in Section 2.2. However, the characteristic scale k_{eq} for the primordial non-Gaussianity signal is different than for the non-primordial terms R_*^{-1} which are determined by the size of the halos. Thus, the degeneracy may not be exact, and below we investigate quantitatively what the expected constraints on $f_{\text{NL}}^{\text{Eq}}$ are with and without the inclusion of non-linear and non-local biasing.

4.1 Optimistic forecast, b_δ and f_{NL} only

As a starting point, we consider constraints on $f_{\text{NL}}^{\text{Eq}}$ marginalized only over the linear, Gaussian bias b_δ . This corresponds to the case of the other bias parameters being known to vanish. It is also the standard approach to forecasting constraints on local non-Gaussianity, which relies on very large scales, where additional bias corrections are small. For the fiducial galaxy sample discussed in Section 3, Figure 3 shows in green the uncertainty in $f_{\text{NL}}^{\text{Eq}}$ as a function of the survey volume. The solid curves assume our default choice $k_{\text{max}} R_* = 1/2$ and, for comparison, the dashed curves show the constraints for $k_{\text{max}} R_* = 1/4, 1$ (we will discuss what is an appropriate cutoff choice in Section 4.2). Since there is a strong signal on small scales, the results are very sensitive to the choice of k_{max} .

We also show in dashed red the current value of $\sigma(f_{\text{NL}}^{\text{Eq}})$ from the bispectra of CMB fluctuations, measured by Planck [1]. For our preferred choice $k_{\text{max}} R_* = 1/2$, we see that a survey with volume $V \sim \mathcal{O}((100h^{-1}\text{Gpc})^3)$, comparable with next-generation galaxy surveys, appears to improve the $f_{\text{NL}}^{\text{Eq}}$ constraint significantly compared to the CMB measurement: $\sigma(f_{\text{NL}}^{\text{Eq}})_{\text{LSS}} \simeq 14$ vs $\sigma(f_{\text{NL}}^{\text{Eq}})_{\text{CMB}} = 44$. Thus, scale-dependent bias in theory appears as a promising way of improving our knowledge of equilateral non-Gaussianity.

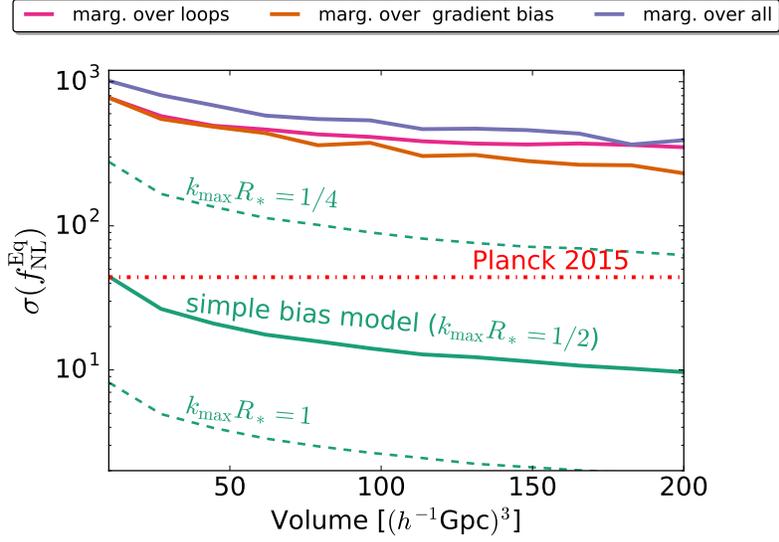


Figure 3: Projected uncertainty on equilateral non-Gaussianity from scale-dependent bias as a function of survey volume. The solid curves show the results for our default survey assumptions and the default cutoff $k_{\max} R_* = 1/2$. The simple bias model (green) corresponds to only marginalizing over the linear bias parameter b_δ . Applying a more general bias prescription leads to a severe degradation in constraining power, as seen with e.g. the purple curve (all bias terms marginalized), which is about a factor of 40 larger. The dashed green curves explore the effect of varying the small-scale cutoff while keeping the simple bias model, and the red dotted-dashed line indicates the current constraint from the CMB bispectra.

4.2 Including non-local bias

Next, we add gradient bias terms,

$$\delta_h(\mathbf{q}) \supset \sum_{n \text{ even}}^{n_{\max}} b_{q^n} (R_* q)^n \delta^{(1)}(\mathbf{q}), \quad (4.3)$$

see Section 2.2, and marginalize over the bias parameters b_{q^n} . As discussed in the beginning of this section, these gradient bias terms come with a typically smaller characteristic length scale, R_* , than the characteristic scale for the onset of equilateral scale-dependent bias, k_{eq}^{-1} . However, without restrictions on the coefficients in eq. (4.3), if we simply expand both the signal, eq. (4.2), and the gradient bias in powers of q^2 , it is clear that they are exactly the same, i.e. we can mimic the effect of $f_{\text{NL}}^{\text{Eq}}$ exactly with the gradient bias expansion by absorbing the difference in characteristic scales in the coefficients, $b_{q^n} \sim c_n (R k_{\text{eq}})^{-n}$. This is a fine-tuning that does not reflect the physical difference between the two contributions. This physical difference is therefore imposed by our prior of order unity on the b_{q^n} parameters, see Table 1.

In general, the gradient bias expansion is expected to converge if the non-locality scale $\sim R_*$ associated with halo formation is small compared to the mode of interest, q^{-1} . Assuming this is satisfied, there must be some finite truncation of the expansion, n_{\max} , such that including terms beyond n_{\max} negligibly affects the results, and in particular $\sigma(f_{\text{NL}}^{\text{Eq}})$. We study this convergence for $k_{\max} R_* = 1/4, 1/2, 1$ in Figure 4. The dots show the constraints using the model with b_δ , $f_{\text{NL}}^{\text{Eq}}$ and the gradient bias expansion as a function of n_{\max} , assuming a survey volume $V = 100(h^{-1}\text{Gpc})^3$.

We see that for $k_{\max} R_* = 1/2$, the results have converged by $n_{\max} = 4$. For $k_{\max} R_* = 1$, while the $f_{\text{NL}}^{\text{Eq}}$ bound appears to converge by $n_{\max} = 8$, the cutoff k_{\max} is clearly outside the regime where the gradient expansion can be expected to make sense. For smaller cutoff, $k_{\max} R_* = 1/4$, convergence is reached early, at $n_{\max} = 2$, but since this cutoff means using a much smaller number of modes, the $f_{\text{NL}}^{\text{Eq}}$ constraint is significantly weakened. As a compromise between the gradient expansion being well behaved, and using as many modes as possible, we choose as our default value for the rest of this work the cutoff $k_{\max} R_* = 1/2$. As discussed in Section 2.1, our expressions for the scale-dependent bias assume the squeezed-limit bispectrum behavior, which technically is only valid for $q R_* \ll 1$. We thus note that for the smallest modes we include, there may be non-negligible corrections to the scale-dependent bias that we do not take into account.

Assuming the default cutoff k_{\max} motivated above, the orange curve in Figure 3 shows the constraint on $f_{\text{NL}}^{\text{Eq}}$ as a function of survey volume when including the gradient expansion. We see that marginalization over non-local bias severely weakens the bound on non-Gaussianity, by about a factor of 25.

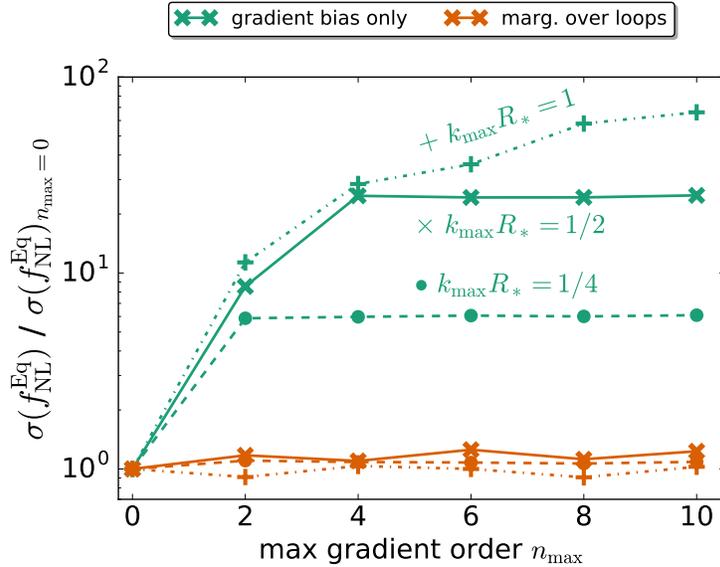


Figure 4: Constraint on $f_{\text{NL}}^{\text{Eq}}$ as a function of the order at which the gradient bias expansion is truncated, relative to the constraint in the absence of gradient bias (i.e. with only the zeroth order term included). The green curves show the case without marginalizing over the loop terms in the tracer power spectrum, assuming short wavelength cutoffs $k_{\max} R_* = 1/4, 1/2, 1$ (circle, cross and plus markers respectively). For the default choice, $k_{\max} R_* = 1/2$, convergence is reached at $n_{\max} = 4$, which we will use as the truncation for most of this work. The orange curves show the same quantities, but including loop contributions. While the dependence on n_{\max} is very different here, $n_{\max} = 4$ is still a reasonable choice.

4.3 Including non-linearities

Finally, we include the loop terms from non-linear bias and evolution, discussed in Section 2.2, marginalizing over the corresponding parameters. The orange crosses in Figure 4 confirm that even in this more general scenario, our truncation of the gradient expansion at $n_{\max} = 4$ is appropriate. The purple curve in Figure 3 shows the final constraints as

a function of survey volume. The degradation factor relative to the simple case without non-linear or non-local bias is about 40, slightly worse than with non-local biasing only. We note that it would be wrong to conclude from this that non-local bias is necessarily more degenerate with scale-dependent bias than the non-linear terms. If we add the non-linear terms first, and *then* include the non-local terms, we would again find that the first step gives the biggest deterioration (see the magenta curve in Fig. 3).

4.4 Summary

While a “naive” forecast suggests scale-dependent bias in galaxy surveys may improve the constraint on $f_{\text{NL}}^{\text{Eq}}$ beyond the current CMB error bar to $\sigma(f_{\text{NL}}^{\text{Eq}}) \simeq 20$, implementing a general treatment of bias and non-linear evolution leads to strong degeneracies, weakening the constraints by roughly a factor 40 leading to $\sigma(f_{\text{NL}}^{\text{Eq}}) \simeq 500$ for a survey volume $V = 100(h^{-1}\text{Gpc})^3$. On the bright side, we do find that the degeneracy is not exact, and that even after marginalization, information on $f_{\text{NL}}^{\text{Eq}}$ remains in scale-dependent bias. Unfortunately, the degraded uncertainties are at least an order of magnitude above the current Planck constraint, even assuming very large survey volumes of order a few $\times 100(h^{-1}\text{Gpc})^3$.

The forecasts in this section assume the use of a single halo sample. When multiple samples with different biases and with high number density are available, it is possible to evade the cosmic variance bound and to strongly improve the constraints [36, 63]. We will study in Section 6 if this method can make constraints on $f_{\text{NL}}^{\text{Eq}}$ from scale-dependent bias more competitive with the CMB. We also point out that for equilateral non-Gaussianity the halo bispectrum [64–66] is expected to do better than the halo power spectrum because it makes use of primordial bispectrum configurations beyond the squeezed limit, and the equilateral signal is dominated by those triangles (unlike local non-Gaussianity). Even for the halo bispectrum, however, degeneracies with non-linear evolution and biasing place limits on what can be achieved [25].

5 Beyond single-field inflation

5.1 Particle physics and the squeezed limit

Inflation is thought to have occurred at energies such that $H \lesssim 10^{14}$ GeV. During inflation, any particles with masses $m \lesssim H$ are excited from the vacuum and can potentially impact the evolution of the fluctuations we ultimately observe. These extra fields are particularly important in the squeezed limit as they can lead to violations of the single-field consistency relations [22, 67]. There are good reasons to think a plethora of new particles could appear at these energies and we should take this possibility seriously [22, 67–71].

The most well-studied possibility is local non-Gaussianity, which arises most commonly from scalar fields with $m \ll H$. In this case, the primordial perturbation of the gravitational potential φ can be written as

$$\varphi = \varphi_{\text{G}} + f_{\text{NL}}(\varphi_{\text{G}}^2 - \langle \varphi_{\text{G}} \rangle^2), \quad (5.1)$$

where φ_{G} is a gaussian variable of power spectrum P_{φ} . In the squeezed limit, local non-Gaussianity takes the form

$$\lim_{k_1 \rightarrow 0} B(k_1, k_2, k_3) = \langle \varphi(k_1)\varphi(k_2)\varphi(k_3) \rangle \rightarrow 4f_{\text{NL}}P_{\varphi}(k_1)P_{\varphi}(k_2) \quad (5.2)$$

where $k_1 \rightarrow 0$ forces $k_2 \simeq k_3$. For $m \simeq H$ (known as Quasi-Single Field, or QSF), the extra fields do not directly contribute to the dynamics at the end of inflation (often responsible for local non-Gaussianity) but lead to non-trivial mode coupling at horizon crossing. This leads to a squeezed⁷ bispectrum [22, 52]

$$\lim_{k_1 \rightarrow 0} B(k_1, k_2, k_3) \rightarrow -\frac{18\sqrt{3}}{\pi} \frac{\Gamma(3/2 - \Delta)}{2^\Delta N_{3/2-\Delta}(8/27)} f_{\text{NL}}^{(\Delta)} \left(\frac{k_1}{k_2}\right)^\Delta P_\varphi(k_1) P_\varphi(k_2) \quad (5.3)$$

where $\Delta = \frac{3}{2} - \sqrt{\frac{9}{4} - \frac{m^2}{H^2}}$ and N_ν is the second-kind Bessel function of order ν . For $m \leq 3H/2$, we find weaker power law scaling in the squeezed limit, with $0 < \Delta \leq \frac{3}{2}$. For $m > 3H/2$ this behavior becomes oscillatory and the amplitude is exponentially suppressed [70, 72]. In both cases, the overall power law is determined by the time evolution of the wavefunction for the massive fields outside the horizon. The suppression in the squeezed limit is ultimately due to the decay of the massive field between the time of horizon crossing of the short and long wavelength modes [67]. In these simple models, the power law never reaches the single-field limit $\Delta = 2$ because we must reproduce the a^{-3} dilution expected for very massive particles [70]. However, by using interactions, we can modify the time evolution to achieve any $0 \leq \Delta \leq 2$ [73].

5.2 Constraints on $f_{\text{NL}}^{(\Delta)}$ ($\Delta = 0 - 2$)

Here we study constraints on scale-dependent bias with arbitrary scale-dependence, $\Delta = 0 - 2$, where, in summary, $\Delta = 0$ corresponds to local non-Gaussianity, $\Delta = 2$ to equilateral non-Gaussianity, and the range $\Delta = 0 - \frac{3}{2}$ to quasi-single field inflation (QSFI) with a range of masses.

In order to understand the statistical power of scale-dependent bias to constrain these models, it is useful to consider the constraints where the amplitude in the squeezed limit is fixed as a function of Δ . Specifically, in this section we will assume

$$b_{\text{NG}}(q) = 6 f_{\text{NL}}^{(\Delta)} (b_\delta - 1) \delta_c (q R_*)^\Delta \mathcal{M}^{-1}(q) , \quad (5.4)$$

where we will take $f_{\text{NL}}^{(\Delta)}$ to vary independently of Δ . The case $\Delta = 0$ corresponds to local non-Gaussianity, although with a slightly unusual normalization $3f_{\text{NL}}^{(\Delta=0)} = f_{\text{NL}}^{\text{Loc}}$. In this limit, the signal-to-noise is dominated by the largest scales where they are robust to non-linear evolution and biasing. However, because of the non-trivial q -dependence of $\mathcal{M}^{-1}(q)$ there is also potentially information available at large q that is not degenerate with the bias parameters. As we increase $\Delta > 0$, the amount of information available at $q \rightarrow 0$ decreases and (naively) increases at $q \rightarrow k_{\text{max}}$.

The constraint on $f_{\text{NL}}^{(\Delta)}$ for varying $\Delta \in [0, 2]$ is shown in Figure 5. We see that when $\Delta < 1$, the results are unaffected by marginalization over bias parameters, which means most of the constraining power is coming from large scales $q \ll k_{\text{max}}$ (for example, one can see in Fig. 2 that the other contributions to P_{hh} are very different from that of $f_{\text{NL}}^{(\Delta=0)}$). However, for $\Delta > 1$, we find much weaker constraints when marginalizing over higher order biases.

⁷This expression is only valid for $\Delta \neq 3/2$. The expression for $\Delta = 3/2$ can be found in e.g. [52].

We can understand the forecasts qualitatively from the likelihood function in eq. (3.2). For a small deviation from the hypothetical model, $\hat{P} = P_{\text{th}} + \delta P$ then we have

$$\log \mathcal{L}(\hat{P}(k)|\boldsymbol{\theta}) \approx -\frac{k^3 V}{4\pi^2} \Delta \ln k \frac{1}{2} \left(\frac{\delta P(k)}{P_{\text{th}}(k, \boldsymbol{\theta})} \right)^2. \quad (5.5)$$

If we consider only the change from the primordial non-Gaussianity, we have $\delta P = 12b_\delta(b_\delta - 1)f_{\text{NL}}^{(\Delta)} \delta_c(k R_*)^\Delta \mathcal{M}^{-1}(k)P(k)$ and therefore the likelihood scales like

$$\log \mathcal{L}(\hat{P}(k)|\boldsymbol{\theta}) \approx -\frac{18k^3 V}{\pi^2} \Delta \ln k \frac{(b_\delta - 1)^2}{b_\delta^2} f_{\text{NL}}^{(\Delta)2} \delta_c^2(k R_*)^{2\Delta} \mathcal{M}^{-2}(k), \quad (5.6)$$

where we have assumed a sample-variance limited measurement. As $k \rightarrow 0$, $\mathcal{M} \propto k^2$ and we see therefore that for $\Delta \leq 1/2$, the signal-to-noise at low- k is dominated by the smallest k (largest scales). However, $\mathcal{M}(k)$ is not simply a power law and for $k > k_{\text{eq}}$ the $\mathcal{M}(k)$ grows more slowly than $k^{3/2}$ and therefore the signal-to-noise increases with increasing k . As a result, we expect there to be contributions to the signal both at largest and smallest scales in the survey for $\Delta < 1/2$. Increasing Δ makes the contribution to the signal increasingly dominated by the smallest scales, especially for $\Delta > 1/2$.

Thus, the best constraints are for Δ close to zero, where adding other bias terms does not significantly alter $\sigma(f_{\text{NL}})$. This is because for low Δ , the signal peaks at large scales which are linear and do not get large contributions from the rest of the bias expansion (see Fig. 2). By contrast, small scales are the most non-linear and degeneracies with biasing will appear as weakening constraints. Therefore, when approaching the equilateral case, the constraints notably worsen for the full bias expansion. The peak in the curve when marginalizing solely over the linear bias comes from the fact that for k close to k_{max} (where most of the signal is for these values of Δ), the transfer function $T(k)$ roughly behaves as k^{-1} , so that $\mathcal{M}(k) \propto k$. Thus, for $\Delta = 1$, there is enhanced degeneracy with the linear bias. As we increase $\Delta > 1$, it will become less degenerate with b_δ but more degenerate with the higher order biases, which is reflected in the differences in the forecasts for difference biasing models.

The scaling properties for these forecasts follow from eq. (5.5) and, in particular, the noise level set by the cosmic variance of the tracer power spectrum. However, mode-coupling occurs on a mode-by-mode basis and can, in principle, be measured without cosmic variance. We will discuss the multi-tracer approach to cosmic variance cancellation in Section 6, but we should anticipate that the forecasts will behave qualitatively differently because of the change of the noise properties.

5.3 Constraints on quasi-single field inflation

Having understood the scaling behavior of the signal-to-noise for generic scale-dependent bias, we are now in a position to understand the corresponding constraints on QSFI. We now must take into account the change to the normalization of the squeezed limit, given by eq. (5.3).

Planck has put constraints on non-Gaussianity coming from QSF inflation [1], using an expression for the bispectrum that interpolates between the local and equilateral shapes.

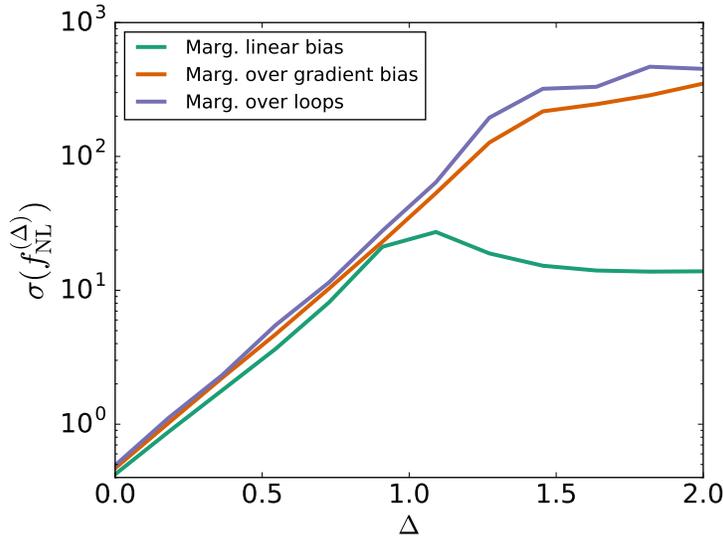


Figure 5: The evolution of $\sigma(f_{\text{NL}}^{(\Delta)})$ as a function of a fixed Δ , for different marginalization schemes. The green curve is the naive scenario, where marginalizing solely over the linear bias b_δ . For the orange curve, we also marginalized over gradient terms, but keeping the non-linear bias zero. The purple curve shows the full marginalization, over the every term in the expression (2.21). Up to $\Delta = 1$ the behavior of the three curves is similar, and the increase with Δ comes from the fact the scales that dominate the signal get away from the squeezed limit, which is what we probe best with scale-dependent bias. At $\Delta = 1$, the degeneracies start to play a more important role, as the signal is dominated by the small scales which are more non-linear and non-local. This explains the large degradation from the green to the purple curve.

The expression is given by [22]

$$B_\varphi(k_1, k_2, k_3) \equiv 6 C_\varphi^2 F(k_1, k_2, k_3),$$

$$F(k_1, k_2, k_3) \equiv \frac{3^{3/2}}{N_\nu(8/27)} f_{\text{NL}} \frac{N_\nu [8k_1 k_2 k_3 / (k_1 + k_2 + k_3)^3]}{(k_1 k_2 k_3)^{3/2} (k_1 + k_2 + k_3)^{3/2}}, \quad (5.7)$$

where $\nu = 3/2 - \Delta$. When looking at the squeezed limit⁸, one recovers eq. (5.3), which leads to a scale-dependent bias similar to eq. (2.7)

$$\lim_{k \rightarrow 0} b_{\text{NG}}^{\text{QSF}}(k) = -\frac{9\sqrt{3}}{\pi} \frac{\Gamma(3/2 - \Delta)}{2^\Delta N_{3/2-\Delta}(8/27)} f_{\text{NL}}^{(\Delta)} (b_\delta - 1) \delta_c (q R_*)^\Delta \mathcal{M}^{-1}(q), \quad (5.8)$$

which, as before, is only valid for $\Delta \neq 3/2$. It diverges as $\Delta \rightarrow 3/2$, because the limit $k \rightarrow 0$ does not commute with the limit $\Delta \rightarrow 3/2$. There is a different expression in this case, given in [52]. This expression is mainly for comparison with eq. (2.7). For our analysis however, we will not use it. Indeed, we expect that for large values of Δ , the signal peaks at scales where this approximation is not valid. This is why we will rather use the expression computed from integrating over the bispectrum (5.7) without taking the squeezed limit. This means that it is given by (see [52])

$$b_{\text{NG}}(k) = \frac{\delta_c [b_\delta - 1]}{2\mathcal{M}(k)} \frac{I_{21}(k)}{\sigma_m^2}, \quad (5.9)$$

⁸large-scale structure probes of quasi-single field inflation need not use the squeezed limit to place constraints on f_{NL} . See [74–76] for some examples.

where

$$\sigma_m^2 \equiv \int d^3q P_G(q) W_R(q)^2, \quad (5.10)$$

$$I_{21}(k) \equiv \frac{1}{P_\varphi(k)} \int d^3q \mathcal{M}(q) \mathcal{M}(|\mathbf{q} - \mathbf{k}|) W_R(q) W_R(|\mathbf{q} - \mathbf{k}|) B_\varphi(q, |\mathbf{q} - \mathbf{k}|, k), \quad (5.11)$$

and $W_R(q) = 3 [\sin(qR) - qR \cos(qR)] / (qR)^3$ is the Fourier transform of a top-hat filter.

We want to use this form of the scale-dependent bias to compare with Planck, which conducted an analysis of QSFI in [1]. Note that in their case, they used directly the bispectrum, not only the squeezed limit as we are doing (even with the bias given by eq. (5.9)). Therefore, we do not expect to improve their constraints for large values of Δ , where equilateral configurations receive a sizable part of the signal.

We show the results in Fig. 6, where we have let Δ be a free parameter, with flat prior between $[0, 3/2]$ and fiducial values $\Delta = 0$ and $f_{\text{NL}} = 0$.

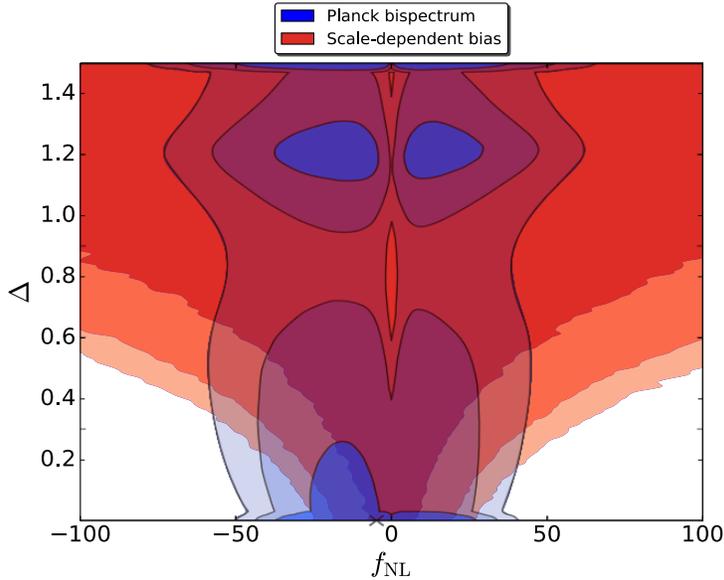


Figure 6: The red regions are the 68%, 95% and 99% contours from our MCMC runs. The blue are the 68%, 95% and 99% contours from Planck [1]. At low Δ , below ~ 0.3 , the scale-dependent bias does better than Planck. Indeed, in this case, the signal is dominated by the largest scales, where galaxy surveys have a decisive advantage. However, as Δ increase, our constraints keep on worsening because the signal is dominated by smaller scales, where the non-local and non-linear bias terms have significant contributions. These are also the scales at which one expects that the scale-dependent bias, that relies on squeezed configurations of the bispectrum, should perform poorly compared to the bispectrum. Indeed, for Planck, high Δ values just means that the signal is more dominated by equilateral shapes, which is well captured by the bispectrum. This is why the blue contours do not change significantly as a function of Δ .

As we saw in the previous section, as we increase Δ away from zero we slowly interpolate between being dominated by the largest scales to dominated by the smallest. It is therefore not surprising that Fig. 6 shows that we have better constraints than existing Planck limits on QSFI for $\Delta \lesssim 0.3$. However, as we increase Δ , our constraints keep on weakening, while those of Planck stay roughly the same. This again was expected from Fig. 5.

To be concrete, the transition in Fig. 5 occurs at $\Delta \sim 1$ and $\sigma(f_{\text{NL}}^{(\Delta)}) < 10$ for $\Delta < 0.5$. A significant difference between these figures is the normalization of the squeezed limit relative to the equilateral limit. For the purpose of a CMB experiment, all of the constraining power is coming from equilateral configurations. For $\Delta \sim 0.5$ we find that the amplitude of the QSFI normalization is suppressed by a factor of 4-5 compared to the expectation from the amplitude in the equilateral configurations. We compared in Fig. 7 the effect of the different normalizations on $\sigma(f_{\text{NL}}^{(\Delta)})$ for fixed values of Δ . One can see in Fig. 7 that at low Δ , where the signal gets a significant contribution from the local configuration, the approximation in eq. (5.8) performs well. However, as the Δ increases and the signal shifts toward higher k , the approximation breaks down. Moreover, as Δ increase, even if one can check that the signal increase faster with the normalization of eq. (2.7) than with the full QSF, the degeneracies are stronger with the former, resulting in a weakening of the constraint on $f_{\text{NL}}^{(\Delta)}$.

The broad lesson here is that this relative normalization is important when comparing limits from scale-dependent bias with probes that are not limited to squeezed configurations. This can be seen more easily if we imagine a model where the bispectrum is a sum of an equilateral and quasi-single field piece, $B = f_{\text{NL}}^{\text{comb},\Delta}(B_{\text{eq}} + \alpha B_{\Delta})$. Since scale-dependent bias cannot measure $f_{\text{NL}}^{\text{Eq}}$ better than Planck, as $\alpha \rightarrow 0$, we will find that Planck places stronger constraints on $f_{\text{NL}}^{\text{comb},\Delta}$ for all Δ . For QSFI, this suppression is significant and limits the improvement over Planck to $\Delta < 0.3$. Given that $f_{\text{NL}}^{(\Delta)} < 20$ for $\Delta < 1$ (and that $f_{\text{NL}}^{(\Delta=2)} = f_{\text{NL}}^{\text{Eq}}$) it is not unreasonable to expect larger improvements could be available in other models with non-trivial scaling in the squeezed limit.

6 Multi-tracer technique

6.1 The set-up

Since our constraints depend on measuring the power spectrum, *a priori* they suffer from a fundamental limit, the so-called cosmic variance limit. This is just to say that even if we were perfectly sampling the density field, because we only have access to a finite number of modes (especially for long modes), there is an error that cannot be beat. While this is certainly true when using a single tracer, it was proposed in [36] that by using more than one tracer, one can get a measurement of the bias (and therefore f_{NL}) that is not cosmic-variance limited. This has recently been applied to constraints on PNG from the galaxy bispectrum [77].

Let's take the example of two tracers, related to the density field δ through $\delta_1 = b^{(1)}\delta$ and $\delta_2 = b^{(2)}\delta$. While each separate measure of δ_i is limited by cosmic variance since one needs to know the stochastic variable δ , taking the ratio is independent of the realization of the underlying field δ . It depends only on the ratio of the biases, which we can thus measure without the cosmic-variance limit.

In order to see what type of improvement might be gained with this method, we also ran Monte Carlo Markov chains with two tracers keeping the same configuration (volume, redshift, etc) as before for the survey. Each one of them has its own independent set of bias parameters as in eq. (2.21), meaning that we have twice as many bias parameters as in the single-tracer case (but still a single f_{NL} parameter). In this case, the quantity of

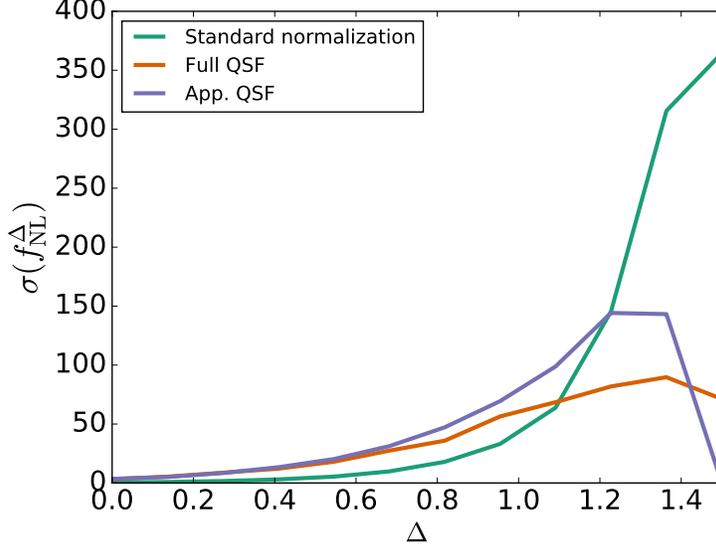


Figure 7: Evolution of the constraint on f_{NL} as a function of Δ , for different normalizations of the scale-dependent bias b_{NG} . The green curve corresponds to a scale-dependent bias given by eq. (2.7). The purple one is given by the local limit of the scale-dependent bias, Eq. (5.8), and is just a Δ dependent rescaling eq. (2.7) (which should not be trusted for $\Delta \rightarrow 3/2$). Finally, the orange curve shows the result with the full QSF normalization of eq. (5.9). At low Δ , where the PNG is dominated by the squeezed limit, the QSF normalization is well approximated by eq. (5.8) and the constraint with standard normalization is about 6 times better. As we go to higher value of Δ , the approximation eq. (5.8) starts to fail: $\sigma(f_{\text{NL}})$ goes to zero because the normalization diverges at $\Delta = 3/2$, and should not be trusted. Moreover, the QSF normalization gets much better constraints than with the normalization eq. (2.7). This is because when using eq. (5.9), the signal is much less degenerate at high k : one can check that for $\Delta = 3/2$ the signal is of the same order with eq. (2.7) and eq. (5.9). This difference in the constraints explains why we cannot interpret Fig. 6 with insights from Fig. 5.

interest is the matrix (the extension of the power spectrum in the single tracer case)

$$C = \begin{pmatrix} P_{11} & P_{12} \\ P_{12} & P_{22} \end{pmatrix}, \quad (6.1)$$

where the P_{ij} are the (cross) power spectra between the tracers i and j . Then, by analogy with the single tracer case, the likelihood is

$$\begin{aligned} \log \mathcal{L}_{\text{Double}}(\hat{C}(k)|\boldsymbol{\theta}) &= -\frac{N_k}{2} \left\{ \text{Tr} \left[\hat{C}(k) C_{\text{th}}^{-1}(k, \boldsymbol{\theta}) \right] - \text{Tr} \left[\log \left(\hat{C}(k) C_{\text{th}}^{-1}(k, \boldsymbol{\theta}) \right) \right] \right\}, \\ &= -\frac{N_k}{2} \left\{ \text{Tr} \left[\hat{C}(k) C_{\text{th}}^{-1}(k, \boldsymbol{\theta}) \right] - \log \left[\det \left(\hat{C}(k) C_{\text{th}}^{-1}(k, \boldsymbol{\theta}) \right) \right] \right\}. \end{aligned} \quad (6.2)$$

Here, \hat{C} is the matrix (6.1) using the fiducial (cross) power spectra $P_{ij} \rightarrow \hat{P}_{ij}$ and C_{th} using the theory ones $P_{ij} \rightarrow P_{\text{th},ij}$, given by the straightforward extension of eq. (2.21).

The decisive parameter here is the number density and the associated shot noise. When the density \bar{n} is low, the shot noise $1/\bar{n}$ is high. This means that the total noise, given by C_{th} , is dominated by the shot noise contribution – remember that P_{th} is given by eq. (2.21) – and not by the part proportional to P_G . In this case, the limiting factor

is thus not cosmic variance, and using the multi-tracer technique cannot improve the constraints. However, as soon as the shot noise is low enough for P_G to dominate the noise, using the multi-tracer technique can dramatically improve the constraints on f_{NL} . To illustrate this, we will take a toy model with two species, each one with a different value for its fiducial bias, which is a requirement for the multi-tracer approach to work (the improvement increases the more different the fiducial biases are [36]). In this case the only thing varying will be the (fiducial) number density, which means that we will go from shot-noise dominated to cosmic-variance limited.

We will also study a case closer to actual surveys, where the first tracer comes from a galaxy population with properties determined by a minimum halo mass, $M_{h,\text{min}}$, analogously to the discussion around eq. (3.3). We will not be free to change the number density, galaxy bias and Lagrangian radius independently, but instead they are jointly determined by $M_{h,\text{min}}$. For the second tracer, we will directly take the matter density field δ , that could in principle be obtained through weak lensing.

There is one small subtlety with k_{max} . The Lagrangian radius R_* depends on $M_{h,\text{min}}$ and can get small for low minimum halo mass. To ensure that we stay in linear theory for our forecast, we impose that the cut-off scale k_{max} be the minimum between $1/(2R_*)$ and the non linear scale, $\pi/(2R_{\text{NL}})$, where R_{NL} is such that the r.m.s. of linear fluctuations in spheres of radius R_{NL} is 1/2.

6.2 The results

To illustrate how the multi-tracer approach allows to beat cosmic variance, we plot on the top panels of Fig. 8 the expected precision on $f_{\text{NL}}^{\text{Loc}}$ (*left*) and $f_{\text{NL}}^{\text{Eq}}$ (*right*) as a function of the shot noise $1/\bar{n}$ (*top*), where \bar{n} is the number density of the population we consider. The two tracers are two different (idealized) population of galaxies, with different biases ($b_{\delta,1} = 3.6$ and $b_{\delta,2} = 2$) but same Lagrangian radius $R_* = 3.7 \text{ Mpc}/h$ and same number density. The single tracer is a combination of the two populations, with bias $(b_{\delta,1} + b_{\delta,2})/2$ and twice as many galaxies. While at low density (high noise) the behaviors are similar, when getting to larger densities, the single tracer sensitivity starts to plateau, limited by cosmic variance, while the double-tracer case keeps on improving.

To put this in a more realistic setting, we show in the bottom panels of Fig. 8 the expected precision on $f_{\text{NL}}^{\text{Loc}}$ (*left*) and $f_{\text{NL}}^{\text{Eq}}$ (*right*) as a function of the minimum halo mass defining one of the tracers, where the other tracer is taken to be a direct measurement of the dark matter density. The minimum halo mass determines the number density and mean bias of the sample in accordance with N-body simulations.

At low minimum halo masses, the density is higher, implying lower shot noise, and the multi-tracer approach is expected to improve the constraints. However, the bias gets closer to one, which both decreases the signal and the effect of the multi-tracer approach (which is more efficient the higher the difference in biases is [36]). On the other hand, at high minimum halo mass, the bias is very different from one, but the density is much lower, which means the signal is dominated by shot noise and therefore the multi-tracer approach is not effective.

Looking at the curves in the single tracer case (in purple), one can also see why we chose the default value of $M_{h,\text{min}} = 10^{13} M_{\odot}$, as it is the one that approximately minimizes $\sigma(f_{\text{NL}}^{\text{Eq}})$.

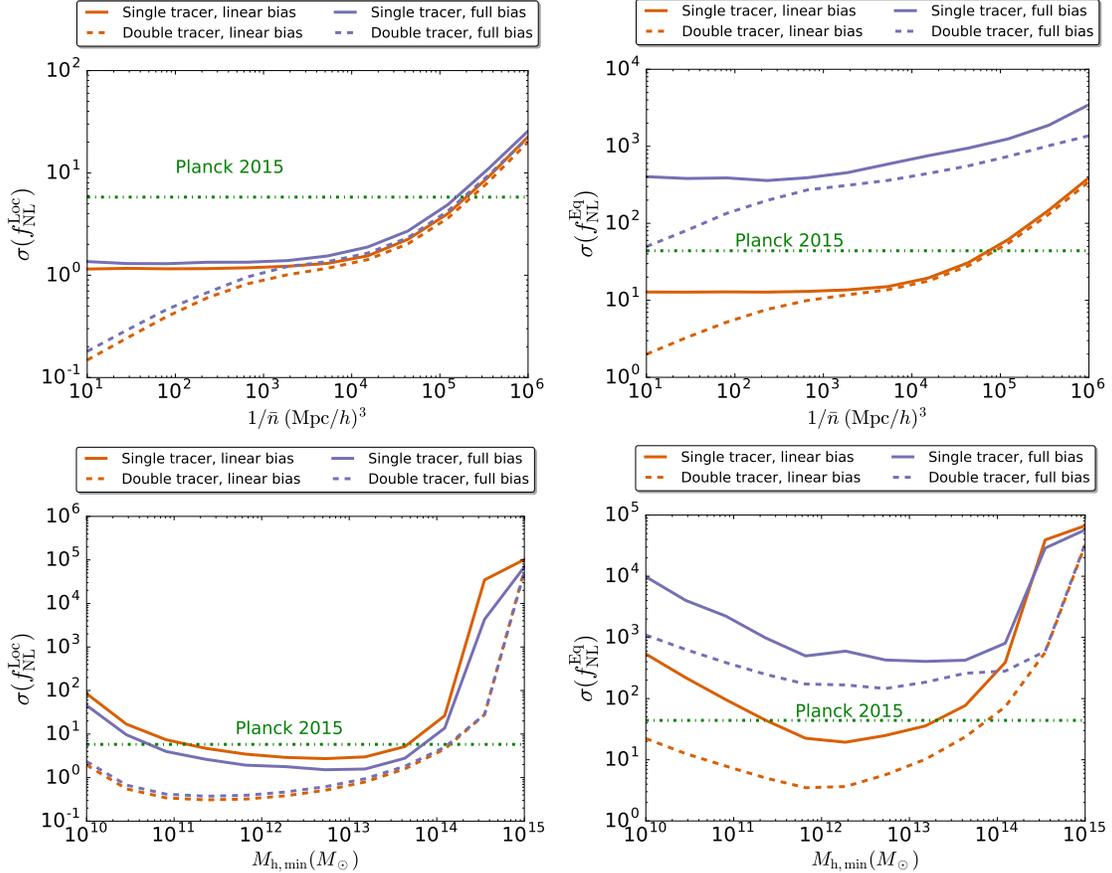


Figure 8: *Top:* Evolution of $\sigma(f_{\text{NL}}^{\text{Loc}})$ (left) and $\sigma(f_{\text{NL}}^{\text{Eq}})$ (right) as a function of the number density for a single tracer (full lines) and two tracers (dashed lines). For the double-tracer case, we have two galaxy populations with different biases ($b_{\delta,1} = 3.6$ and $b_{\delta,2} = 2$) and same number density. The single tracer is the sum of those two populations, with bias $(b_{\delta,1} + b_{\delta,2})/2$ and twice the number density of the double-tracer case. As the shot noise becomes negligible, the single-tracer case is limited by cosmic variance, while the double-tracer keeps on improving. *Bottom:* Evolution of $\sigma(f_{\text{NL}}^{\text{Loc}})$ (left) and $\sigma(f_{\text{NL}}^{\text{Eq}})$ (right) as a function of the minimum halo mass for a single tracer (solid lines) and two tracers (dashed lines). For the single tracer, the bias, Lagrangian radii and densities of galaxies \bar{n} are computed from the halo mass function. For the double-tracer, the additional tracer is taken to be the dark matter field (possibly obtained by weak lensing). Using two tracers always constitutes an improvement by roughly an order of magnitude over a single one (see main text for details). In purple, the results when marginalized over all bias parameters, in orange when the marginalization is only over the linear bias(es). The green lines are the current limit from Planck [1].

On each panel of Fig. 8, we have made a comparison between our MCMC forecasts (with a fiducial volume $V = 100 (\text{Gpc}/h)^3$) and the current constraints from Planck [1]. While for local PNG, we can (optimistically) improve by an order of magnitude on the Planck value, it is harder to imagine reaching those level for the equilateral PNG. Even at the minimum, our constraints are still a factor ~ 3 higher than those from Planck.

However, these plots show that there is no “error floor” due to the degeneracies with other bias parameters that would prevent measuring $f_{\text{NL}}^{\text{Eq}}$ using scale-dependent bias. Degeneracies are still an important factor in the poor performance, but this adds to the

fact that the signal is weak compared to local PNG, or in other words that its signal peaks for bispectrum configurations away from the squeezed limit domain. We can conclude that equilateral non-Gaussianity is not as ideal for scale-dependent bias as the local type for both reasons above. Nevertheless, for the case of quasi-single field inflation described in Section 5, the results are more promising, with potential improvement over Planck for most of the parameter space (see Figs. 9 and 10).

To gauge where one can expect an improvement in a setup similar to the bottom panels of Fig. 8, we plot in Fig. 9 the expected value of the 99% bound on $f_{\text{NL}}^{(\Delta)}$, for fixed values of Δ (using the normalization of eq. (5.9)). We compared it to an estimate of the bound from Planck, approximating the 3σ contours of Fig. 6 by a constant value of 50. $M_{\text{h,min}}$ is the minimum halo mass of the galaxy sample that constitutes the first tracer, the second being dark matter.

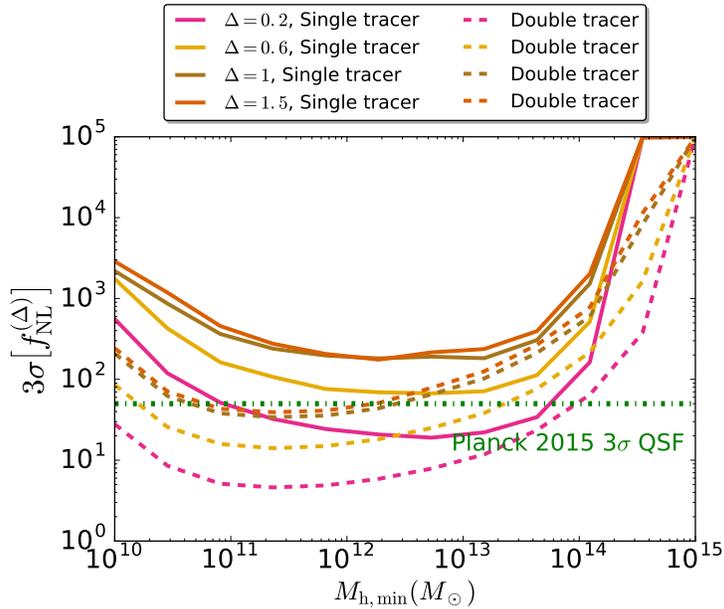


Figure 9: The evolution of 99% bound on $f_{\text{NL}}^{(\Delta)}$ from quasi-single field inflation (see Section 5), as a function of the minimum halo mass of the galaxy sample. The solid lines are for an analysis with this sample only, while the dashed one combine it with a second tracer, dark matter. For all values of Δ , there is an improvement over Planck when using two tracers and sufficiently small minimum halo masses (which means larger number density and smaller shot noise).

At first sight, it seems that for any value of Δ , one can find a range of minimum halo masses where double-tracer improves on Planck. However, the analysis in Fig. 9 assumes a fix value of Δ , while in Fig. 6 it was allowed to vary, bringing potential new degeneracies. This is why, in order to compare with Fig. 6, we look at the constraints in the plane (f_{NL}, Δ) using two tracers, dark matter and galaxies. To see how the improvement changes with minimum halo mass as in Fig. 8, we plot two set of contours for two different minimum halo masses. The results are shown in Fig. 10

When looking at doing the same analysis as before for galaxies (i.e. setting $M_{\text{h,min}} = 10^{13} M_{\odot}$) but using the information in the dark matter field as well, the region where scale dependent bias improves on Planck goes from $\Delta \lesssim 0.2$ to $\Delta \lesssim 0.6$. Moreover, as soon

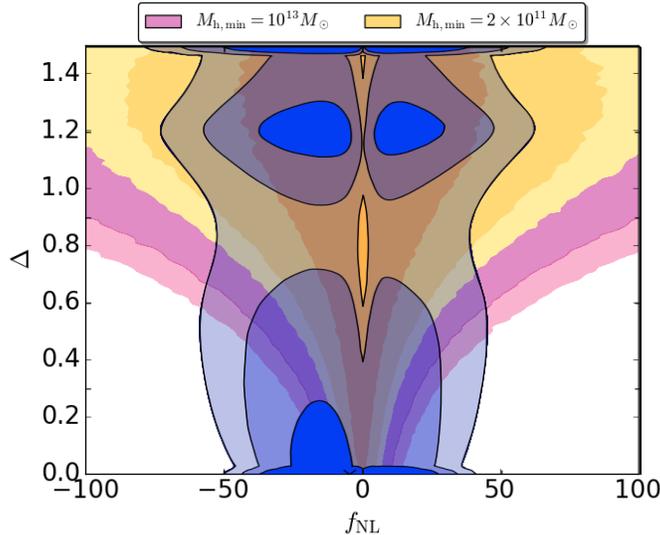


Figure 10: The green regions are the 68%, 95% and 99% contours from our MCMC with two tracers, dark matter and a galaxy sample with $M_{h,\min} = 10^{13} M_{\odot}$, similar to our analysis in the single tracer case. The red regions are the 68%, 95% and 99% contours when considering dark matter and a galaxy sample with $M_{h,\min} = 2 \times 10^{11} M_{\odot}$. The blue are the 68%, 95% and 99% contours from Planck [1]. While for the green contours, we recover only a minor advantage over the single tracer case of Fig. 6, for lower minimum halo mass (and therefore lower shot noise), the multitracers approach provides an improvement over Planck for almost all values of Δ .

as we use more galaxies by setting $M_{h,\min} = 2 \times 10^{11} M_{\odot}$, the shot noise reduces and the bounds are better than Planck’s for more than half of the range in Δ . In order to see that, one has to use the full scale-dependent bias eq. (5.9) and not the approximate eq. (5.8). In particular, it is not necessarily easy to compare those results to the equilateral case, we showed in Fig. 7 that the two normalization (2.7) and (5.9) give very different results for $\Delta \rightarrow 1.5$.

7 Summary and Conclusions

The ongoing search for primordial non-Gaussianity will continue to sharpen our understanding of the very early history of the universe and the origin of structure. Large-scale structure offers great potential for increasing our sensitivity over current limits from the CMB. However, large-scale structure is non-Gaussian in its own right and separating the two effects will be challenging.

In the case of local non-Gaussianity, scale-dependent bias is a particularly compelling target as it is distinguishable from effects of late time evolution and can be measured directly from the power spectra. The signal arises from squeezed configurations of the primordial correlation functions (especially the three-point function), which are strongly constrained by symmetries [8]. In the context of inflation, the squeezed limit is sensitive to the spectrum of particles and their interactions around the time of horizon crossing.

In this paper, we explored the use of scale-dependent bias for constraining non-Gaussianity beyond the local shape. In general, the scale-dependent bias signal can be expected to be partially degenerate with contributions to the halo power spectrum related

to non-linear evolution, in particular the non-linear evolution of matter fluctuations, non-linear biasing between halos and matter, and non-local halo biasing. A major focus of this work was to include these contributions in a comprehensive manner and to study primordial non-Gaussianity constraints after marginalization over this non-primordial physics.

We were particularly interested in the scale-dependence that arises from the transfer function between the *primordial* Newtonian potential (for which the mode-coupling due to primordial non-Gaussianity is specified) and the potential during matter domination. In principle, this makes any primordial mode-coupling distinguishable from mode-coupling due to non-linear structure formation. Here we quantified the degree to which equilateral and quasi-single field shapes can be measured in this way and studied the degeneracy with the parameters of a general biasing model.

For local and some quasi-single field shapes, significant improvements over Planck are achievable for a large volume large-scale structure survey: for example, $\sigma(f_{\text{NL}}^{\text{Loc}}) \simeq 1$ could be reached (similar to the forecasts of future surveys in [24, 78, 79]) with a single tracer of the matter density (e.g. a single galaxy sample), which could be improved by a factor of 5 with a multi-tracer approach. This is a level that is interesting from a theoretical point of view [80]. For equilateral non-Gaussianity, after marginalizing over a comprehensive halo biasing model, the sensitivity falls below the current Planck limits, reaching $\sigma(f_{\text{NL}}^{\text{Eq}}) \simeq 500$ with a single tracer, and about 145 with two tracers. The latter is about 3 times worse than current Planck sensitivity. Nevertheless, there is information encoded at large distances, where the transfer function is non-trivial, that can in principle be measured without cosmic variance. For this reason, forecasting using the multi-tracer technique shows significant improvements for equilateral non-Gaussianity. Moreover, the case of quasi-single field inflation is even more promising, with significant improvement possible over a good portion of the parameter space.

Ultimately, a large-scale structure survey should also include the information in the galaxy bispectrum as part of its analysis of primordial non-Gaussianity. For non-local shapes (with large signal in non-squeezed configurations), the galaxy bispectrum contains much more information than scale-dependent bias and could lead to significant improvements over Planck for all shapes because large-scale structure surveys have access to a much higher number of clustering modes than the CMB. However, a galaxy bispectrum analysis presents a number of challenges that are not present for the power spectrum. We have shown in this work that, optimistically, there is an interesting range of parameter space (beyond merely local-type non-Gaussianity) where improvements over Planck are achievable from the power spectrum alone.

Acknowledgements. We would like to thank Fabian Schmidt for helpful discussions. Part of the research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. This research is partially supported by NASA ROSES ATP 14-ATP14- 0093 grant. RdP and O.D. acknowledge support by the Heising-Simons foundation.

A Calculating the halo power spectrum

By summing the contribution from all the diagrams described in Section 2.2, one gets a halo-halo power spectrum of the following form

$$\begin{aligned}
P_{hh}(k) = & b_{\text{full}}^2 [P_G(k) + P_{22}(k) + P_{13}(k)] + P_{\epsilon_0} + 2P_{\epsilon_\delta} P_G(k) \\
& + 4b_{\text{full}} \left\{ \int_{\mathbf{p}} \left[b_{\delta^2} + b_{s^2} \left(\mu_-^2 - \frac{1}{3} \right) \right] [P_G(|\mathbf{q} - \mathbf{p}|) - P_G(p)] P_G(p) \right\} \\
& + 2 \int_{\mathbf{p}} \left[b_{\delta^2} + b_{s^2} \left(\mu_-^2 - \frac{1}{3} \right) \right]^2 P_G(|\mathbf{q} - \mathbf{p}|) P_G(p) \\
& + b_{\text{full}} P_G(k) \left\{ 6b_{\text{III}^{(2)}} \int_{\mathbf{p}} \frac{2}{7} \left[\left(\frac{\mathbf{k} \cdot \mathbf{p}}{k p} \right)^2 - 1 \right] \frac{2\mu^2}{3} P_G(p) \right. \\
& \left. + 8 b_{s^2} \int_{\mathbf{p}} \left(\mu_-^2 - \frac{1}{3} \right) F_2(-\mathbf{p}, \mathbf{k}) P_G(p) + 8 b_{\delta^2} \int_{\mathbf{p}} F_2(-\mathbf{p}, \mathbf{k}) P_G(p) \right\}.
\end{aligned} \tag{A.1}$$

We have used

$$\mu_- \equiv \frac{(\mathbf{k} - \mathbf{p}) \cdot \mathbf{p}}{|\mathbf{k} - \mathbf{p}| p}, \tag{A.2}$$

we recall that (see standard cosmological perturbation theory [53])

$$F_2(\mathbf{k}_1, \mathbf{k}_2) = \frac{5}{7} + \frac{2}{7} \left(\frac{\mathbf{k}_1 \cdot \mathbf{k}_2}{k_1 k_2} \right)^2 + \frac{1}{2} \frac{\mathbf{k}_1 \cdot \mathbf{k}_2}{k_1 k_2} \left[\frac{k_1}{k_2} + \frac{k_2}{k_1} \right], \tag{A.3}$$

and P_G is the power spectrum of δ_G . Let us go through the various terms in this lengthy expression

- The first line is the extension of SPT result including b_{full} plus the contribution from stochastic fields. The last one, coming from $\epsilon_\delta \delta$, can be reabsorbed in the definition of the linear bias b_δ .
- The second line corresponds to Type 1) diagrams of Fig. 1, where one of the kernel is F_2 . The factor 4 in front of the first curly parentheses comes for symmetry reasons: looking at the left of Fig. 1, there are two diagrams of the sort depending on how you connect the lines, and the extra factor of two comes for changing G and G' .
- The third line corresponds to Type 1) diagrams of Fig. 1, where both kernels are different from F_2 . The factor 2 comes again for symmetry reasons.
- The third line is the only contributing diagram of Type 2) that is not SPT. The factor of 6 comes from the symmetries of the diagram, the factors 2/3 and 2/7 from the definition of the kernel.
- In the last line we find the diagrams of the third type.

Once we have this expression in terms of “bare” quantities, we can remove the parts that lead to divergences⁹ as well as perform angular integration to simplify it. Moreover,

⁹Some of them only appear for $k \rightarrow 0$, which can be removed by replacing $P_G(|\mathbf{k} - \mathbf{p}|) \rightarrow P_G(|\mathbf{k} - \mathbf{p}|) - P_G(p)$.

the last term of the first line can be reabsorbed in a redefinition of b_δ , and will thus disappear from the expression. We then obtain

$$\begin{aligned}
P_{hh}(k) = & b_{\text{full}}^2 [P_G(k) + P_{22}(k) + P_{13}(k)] + P_{\epsilon_0} \\
& + 4b_{\text{full}} \left\{ \int_{\mathbf{p}} \left[b_{\delta^2} + b_{s^2} \left(\mu_-^2 - \frac{1}{3} \right) \right] [P_G(|\mathbf{q} - \mathbf{p}|) - P_G(p)] P_G(p) \right\} \\
& + 2 \int_{\mathbf{p}} \left[b_{\delta^2} + b_{s^2} \left(\mu_-^2 - \frac{1}{3} \right) \right]^2 [P_G(|\mathbf{k} - \mathbf{p}|) - P_G(p)] P_G(p) \\
& + b_{\text{full}} P_G(k) \left\{ 6b_{\text{III}^{(2)}} \int_{\mathbf{p}} \left(\frac{2}{7} \left[\left(\frac{\mathbf{k} \cdot \mathbf{p}}{kp} \right)^2 - 1 \right] \frac{2\mu_-^2}{3} + \frac{8}{63} \right) P_G(p) \right. \\
& \quad \left. + 8b_{s^2} \int_{\mathbf{p}} \left[\left(\mu_-^2 - \frac{1}{3} \right) F_2(-\mathbf{p}, \mathbf{k}) - \frac{34}{63} \right] P_G(p) \right\}.
\end{aligned} \tag{A.4}$$

In the last two terms, the integration over the angle can be done analytically, and one finds that the contributions are proportional. Therefore, we can regroup them into a single one, keeping only the $b_{\text{III}^{(2)}}$ term, leading to eq. (2.21).

References

- [1] **Planck** Collaboration, P. A. R. Ade *et al.*, “Planck 2015 results. XVII. Constraints on primordial non-Gaussianity,” [arXiv:1502.01592](#) [[astro-ph.CO](#)].
- [2] M. Alvarez *et al.*, “Testing Inflation with Large Scale Structure: Connecting Hopes with Reality,” [arXiv:1412.4671](#) [[astro-ph.CO](#)].
- [3] N. Dalal, O. Doré, D. Huterer, and A. Shirokov, “The imprints of primordial non-gaussianities on large-scale structure: scale dependent bias and abundance of virialized objects,” *Phys.Rev.* **D77** (2008) 123514, [arXiv:0710.4560](#) [[astro-ph](#)].
- [4] S. Matarrese and L. Verde, “The Effect of Primordial Non-Gaussianity on Halo Bias,” *Astrophys. J Lett.* **677** (Apr., 2008) L77, [arXiv:0801.4826](#).
- [5] A. Slosar, C. Hirata, U. Seljak, S. Ho, and N. Padmanabhan, “Constraints on local primordial non-Gaussianity from large scale structure,” *JCAP* **8** (Aug., 2008) 031, [arXiv:0805.3580](#).
- [6] V. Desjacques and U. Seljak, “Primordial non-Gaussianity from the large-scale structure,” *Classical and Quantum Gravity* **27** no. 12, (June, 2010) 124011, [arXiv:1003.5020](#).
- [7] J. M. Maldacena, “Non-Gaussian features of primordial fluctuations in single field inflationary models,” *JHEP* **05** (2003) 013, [arXiv:astro-ph/0210603](#) [[astro-ph](#)].
- [8] P. Creminelli and M. Zaldarriaga, “Single field consistency relation for the 3-point function,” *JCAP* **0410** (2004) 006, [arXiv:astro-ph/0407059](#) [[astro-ph](#)].
- [9] A. Kehagias and A. Riotto, “Symmetries and Consistency Relations in the Large Scale Structure of the Universe,” *Nucl. Phys.* **B873** (2013) 514–529, [arXiv:1302.0130](#) [[astro-ph.CO](#)].
- [10] M. Peloso and M. Pietroni, “Galilean invariance and the consistency relation for the nonlinear squeezed bispectrum of large scale structure,” *JCAP* **1305** (2013) 031, [arXiv:1302.0223](#) [[astro-ph.CO](#)].

- [11] P. Creminelli, J. Norena, M. Simonović, and F. Vernizzi, “Single-Field Consistency Relations of Large Scale Structure,” *JCAP* **1312** (2013) 025, [arXiv:1309.3557 \[astro-ph.CO\]](#).
- [12] P. Creminelli, J. Gleyzes, M. Simonović, and F. Vernizzi, “Single-Field Consistency Relations of Large Scale Structure. Part II: Resummation and Redshift Space,” *JCAP* **1402** (2014) 051, [arXiv:1311.0290 \[astro-ph.CO\]](#).
- [13] P. Creminelli, J. Gleyzes, L. Hui, M. Simonović, and F. Vernizzi, “Single-Field Consistency Relations of Large Scale Structure. Part III: Test of the Equivalence Principle,” *JCAP* **1406** (2014) 009, [arXiv:1312.6074 \[astro-ph.CO\]](#).
- [14] F. Schmidt and M. Kamionkowski, “Halo Clustering with Non-Local Non-Gaussianity,” *Phys. Rev.* **D82** (2010) 103002, [arXiv:1008.0638 \[astro-ph.CO\]](#).
- [15] V. Desjacques, D. Jeong, and F. Schmidt, “Large-Scale Galaxy Bias,” [arXiv:1611.09787 \[astro-ph.CO\]](#).
- [16] N. Bartolo, S. Matarrese, and A. Riotto, “Nongaussianity from inflation,” *Phys. Rev.* **D65** (2002) 103505, [arXiv:hep-ph/0112261 \[hep-ph\]](#).
- [17] N. Bartolo, E. Komatsu, S. Matarrese, and A. Riotto, “Non-Gaussianity from inflation: Theory and observations,” *Phys. Rept.* **402** (2004) 103–266, [arXiv:astro-ph/0406398 \[astro-ph\]](#).
- [18] D. H. Lyth and Y. Rodriguez, “The Inflationary prediction for primordial non-Gaussianity,” *Phys. Rev. Lett.* **95** (2005) 121302, [arXiv:astro-ph/0504045 \[astro-ph\]](#).
- [19] F. Vernizzi and D. Wands, “Non-gaussianities in two-field inflation,” *JCAP* **0605** (2006) 019, [arXiv:astro-ph/0603799 \[astro-ph\]](#).
- [20] R. de Putter, O. Doré, and D. Green, “Is There Scale-Dependent Bias in Single-Field Inflation?,” *JCAP* **1510** no. 10, (2015) 024, [arXiv:1504.05935 \[astro-ph.CO\]](#).
- [21] L. Dai, E. Pajer, and F. Schmidt, “On Separate Universes,” *JCAP* **1510** no. 10, (2015) 059, [arXiv:1504.00351 \[astro-ph.CO\]](#).
- [22] X. Chen and Y. Wang, “Quasi-Single Field Inflation and Non-Gaussianities,” *JCAP* **1004** (2010) 027, [arXiv:0911.3380 \[hep-th\]](#).
- [23] B. Leistedt, H. V. Peiris, and N. Roth, “Constraints on Primordial Non-Gaussianity from 800?000 Photometric Quasars,” *Phys. Rev. Lett.* **113** no. 22, (2014) 221301, [arXiv:1405.4315 \[astro-ph.CO\]](#).
- [24] O. Doré *et al.*, “Cosmology with the SPHEREX All-Sky Spectral Survey,” [arXiv:1412.4872 \[astro-ph.CO\]](#).
- [25] T. Baldauf, M. Mirbabayi, M. Simonović, and M. Zaldarriaga, “LSS constraints with controlled theoretical uncertainties,” [arXiv:1602.00674 \[astro-ph.CO\]](#).
- [26] T. Giannantonio, C. Porciani, J. Carron, A. Amara, and A. Pillepich, “Constraining primordial non-Gaussianity with future galaxy surveys,” *MNRAS* **422** (June, 2012) 2854–2877, [arXiv:1109.0958 \[astro-ph.CO\]](#).
- [27] J. Byun and R. Bean, “Non-Gaussian Shape Discrimination with Spectroscopic Galaxy Surveys,” *JCAP* **1503** no. 03, (2015) 019, [arXiv:1409.5440 \[astro-ph.CO\]](#).
- [28] J. Norena, L. Verde, G. Barenboim, and C. Bosch, “Prospects for constraining the shape of non-Gaussianity with the scale-dependent bias,” *JCAP* **1208** (2012) 019, [arXiv:1204.6324 \[astro-ph.CO\]](#).
- [29] E. Sefusatti, J. R. Fergusson, X. Chen, and E. P. S. Shellard, “Effects and Detectability of Quasi-Single Field Inflation in the Large-Scale Structure and Cosmic Microwave

- Background,” *JCAP* **1208** (2012) 033, [arXiv:1204.6318](#) [[astro-ph.CO](#)].
- [30] N. Agarwal, S. Ho, and S. Shandera, “Constraining the initial conditions of the Universe using large scale structure,” *JCAP* **1402** (2014) 038, [arXiv:1311.2606](#) [[astro-ph.CO](#)].
- [31] K. C. Chan, R. Scoccimarro, and R. K. Sheth, “Gravity and Large-Scale Non-local Bias,” *Phys. Rev.* **D85** (2012) 083509, [arXiv:1201.3614](#) [[astro-ph.CO](#)].
- [32] T. Baldauf, U. Seljak, V. Desjacques, and P. McDonald, “Evidence for Quadratic Tidal Tensor Bias from the Halo Bispectrum,” *Phys. Rev.* **D86** (2012) 083540, [arXiv:1201.4827](#) [[astro-ph.CO](#)].
- [33] V. Assassi, D. Baumann, and F. Schmidt, “Galaxy bias and primordial non-Gaussianity,” *JCAP* **12** (Dec., 2015) 043, [arXiv:1510.03723](#).
- [34] P. McDonald and A. Roy, “Clustering of dark matter tracers: generalizing bias for the coming era of precision LSS,” *JCAP* **8** (Aug., 2009) 20, [arXiv:0902.0991](#) [[astro-ph.CO](#)].
- [35] D. Baumann, S. Ferraro, D. Green, and K. M. Smith, “Stochastic Bias from Non-Gaussian Initial Conditions,” *JCAP* **1305** (2013) 001, [arXiv:1209.2173](#) [[astro-ph.CO](#)].
- [36] U. Seljak, “Extracting primordial non-gaussianity without cosmic variance,” *Phys. Rev. Lett.* **102** (2009) 021302, [arXiv:0807.1770](#) [[astro-ph](#)].
- [37] T. Matsubara, “Nonlinear Perturbation Theory Integrated with Nonlocal Bias, Redshift-space Distortions, and Primordial Non-Gaussianity,” *Phys. Rev.* **D83** (2011) 083518, [arXiv:1102.4619](#) [[astro-ph.CO](#)].
- [38] T. Matsubara, “Deriving an Accurate Formula of Scale-dependent Bias with Primordial Non-Gaussianity: An Application of the Integrated Perturbation Theory,” *Phys. Rev.* **D86** (2012) 063518, [arXiv:1206.0562](#) [[astro-ph.CO](#)].
- [39] V. Assassi, D. Baumann, D. Green, and M. Zaldarriaga, “Renormalized Halo Bias,” *JCAP* **1408** (2014) 056, [arXiv:1402.5916](#) [[astro-ph.CO](#)].
- [40] F. Schmidt, D. Jeong, and V. Desjacques, “Peak-background split, renormalization, and galaxy clustering,” *Phys. Rev. D* **88** no. 2, (July, 2013) 023515, [arXiv:1212.0868](#) [[astro-ph.CO](#)].
- [41] M. Mirbabayi, F. Schmidt, and M. Zaldarriaga, “Biased tracers and time evolution,” *JCAP* **7** (July, 2015) 30, [arXiv:1412.5169](#).
- [42] L. Senatore, “Bias in the effective field theory of large scale structures,” *JCAP* **11** (Nov., 2015) 007, [arXiv:1406.7843](#).
- [43] V. Assassi, D. Baumann, E. Pajer, Y. Welling, and D. van der Woude, “Effective theory of large-scale structure with primordial non-Gaussianity,” *JCAP* **11** (Nov., 2015) 024, [arXiv:1505.06668](#).
- [44] C. Wagner and L. Verde, “N-body simulations with generic non-Gaussian initial conditions II: halo bias,” *JCAP* **3** (Mar., 2012) 2, [arXiv:1102.3229](#).
- [45] R. Scoccimarro, L. Hui, M. Manera, and K. C. Chan, “Large-scale bias and efficient generation of initial conditions for nonlocal primordial non-Gaussianity,” *Phys. Rev. D* **85** no. 8, (Apr., 2012) 083002, [arXiv:1108.5512](#) [[astro-ph.CO](#)].
- [46] E. Sefusatti, M. Crocce, and V. Desjacques, “The halo bispectrum in N-body simulations with non-Gaussian initial conditions,” *MNRAS* **425** (Oct., 2012) 2903–2930, [arXiv:1111.6966](#).
- [47] B. Bayta?, A. Kesavan, E. Nelson, S. Park, and S. Shandera, “Nonlocal bispectra from super cosmic variance,” *Phys. Rev.* **D91** no. 8, (2015) 083518, [arXiv:1502.01009](#)

[astro-ph.CO].

- [48] B. Bonga, S. Brahma, A.-S. Deutsch, and S. Shandera, “Cosmic variance in inflation with two light scalars,” *JCAP* **1605** no. 05, (2016) 018, [arXiv:1512.05365](#) [astro-ph.CO].
- [49] F. Schmidt and M. Kamionkowski, “Halo clustering with nonlocal non-Gaussianity,” *Phys.Rev.D* **82** no. 10, (Nov., 2010) 103002, [arXiv:1008.0638](#) [astro-ph.CO].
- [50] V. Desjacques, D. Jeong, and F. Schmidt, “Non-Gaussian Halo Bias Re-examined: Mass-dependent Amplitude from the Peak-Background Split and Thresholding,” *Phys. Rev. D* **84** no. 6, (Sept., 2011) 063512, [arXiv:1105.3628](#) [astro-ph.CO].
- [51] V. Desjacques, D. Jeong, and F. Schmidt, “Accurate predictions for the scale-dependent galaxy bias from primordial non-Gaussianity,” *Phys. Rev. D* **84** no. 6, (Sept., 2011) 061301, [arXiv:1105.3476](#).
- [52] E. Sefusatti, J. R. Fergusson, X. Chen, and E. P. S. Shellard, “Effects and detectability of quasi-single field inflation in the large-scale structure and cosmic microwave background,” *JCAP* **8** (Aug., 2012) 033, [arXiv:1204.6318](#).
- [53] F. Bernardeau, S. Colombi, E. Gaztanaga, and R. Scoccimarro, “Large scale structure of the universe and cosmological perturbation theory,” *Phys. Rept.* **367** (2002) 1–248, [arXiv:astro-ph/0112551](#) [astro-ph].
- [54] **Planck** Collaboration, P. A. R. Ade *et al.*, “Planck 2015 results. XIII. Cosmological parameters,” [arXiv:1502.01589](#) [astro-ph.CO].
- [55] A. Lewis, A. Challinor, and A. Lasenby, “Efficient computation of CMB anisotropies in closed FRW models,” *Astrophys. J.* **538** (2000) 473–476, [arXiv:astro-ph/9911177](#) [astro-ph].
- [56] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, “emcee: The MCMC Hammer,” *Publ. Astron. Soc. Pac.* **125** (2013) 306–312, [arXiv:1202.3665](#) [astro-ph.IM].
- [57] B. Kalus, W. J. Percival, and L. Samushia, “Cosmological parameter inference from galaxy clustering: The effect of the posterior distribution of the power spectrum,” *Mon. Not. Roy. Astron. Soc.* **455** no. 3, (2016) 2573–2581, [arXiv:1504.03979](#) [astro-ph.CO].
- [58] J. Tinker, A. V. Kravtsov, A. Klypin, K. Abazajian, M. Warren, G. Yepes, S. Gottlöber, and D. E. Holz, “Toward a Halo Mass Function for Precision Cosmology: The Limits of Universality,” *Astrophys. J.* **688** (Dec., 2008) 709–728, [arXiv:0803.2706](#).
- [59] J. L. Tinker, B. E. Robertson, A. V. Kravtsov, A. Klypin, M. S. Warren, G. Yepes, and S. Gottlöber, “The Large-scale Bias of Dark Matter Halos: Numerical Calibration and Model Tests,” *Astrophys. J.* **724** (Dec., 2010) 878–886, [arXiv:1001.3162](#) [astro-ph.CO].
- [60] T. Lazeyras, C. Wagner, T. Baldauf, and F. Schmidt, “Precision measurement of the local bias of dark matter halos,” *JCAP* **2** (Feb., 2016) 018, [arXiv:1511.01096](#).
- [61] M. Musso, A. Paranjape, and R. K. Sheth, “Scale-dependent halo bias in the excursion set approach,” *MNRAS* **427** (Dec., 2012) 3145–3158, [arXiv:1205.3401](#).
- [62] M. Schmittfull, T. Baldauf, and U. Seljak, “Near optimal bispectrum estimators for large-scale structure,” *Phys. Rev. D* **91** no. 4, (Feb., 2015) 043530, [arXiv:1411.6595](#).
- [63] D. Yamauchi, K. Takahashi, and M. Oguri, “Constraining primordial non-Gaussianity via a multitracer technique with surveys by Euclid and the Square Kilometre Array,” *Phys. Rev. D* **90** no. 8, (2014) 083520, [arXiv:1407.5453](#) [astro-ph.CO].
- [64] E. Sefusatti and E. Komatsu, “The bispectrum of galaxies from high-redshift galaxy surveys: Primordial non-Gaussianity and non-linear galaxy bias,” *Phys. Rev. D* **76** (2007) 083004, [arXiv:0705.0343](#) [astro-ph].

- [65] S. Mizuno and S. Yokoyama, “Halo/Galaxy Bispectrum with Equilateral-type Primordial Trispectrum,” *Phys. Rev.* **D91** no. 12, (2015) 123521, [arXiv:1504.05505 \[astro-ph.CO\]](#).
- [66] I. Hashimoto, S. Mizuno, and S. Yokoyama, “Constraining equilateral-type primordial non-Gaussianities from imaging surveys,” *Phys. Rev.* **D94** no. 4, (2016) 043532, [arXiv:1605.07348 \[astro-ph.CO\]](#).
- [67] D. Baumann and D. Green, “Signatures of Supersymmetry from the Early Universe,” *Phys. Rev.* **D85** (2012) 103520, [arXiv:1109.0292 \[hep-th\]](#).
- [68] V. Assassi, D. Baumann, D. Green, and L. McAllister, “Planck-Suppressed Operators,” *JCAP* **1401** (2014) 033, [arXiv:1304.5226 \[hep-th\]](#).
- [69] N. Craig and D. Green, “Testing Split Supersymmetry with Inflation,” *JHEP* **07** (2014) 102, [arXiv:1403.7193 \[hep-ph\]](#).
- [70] N. Arkani-Hamed and J. Maldacena, “Cosmological Collider Physics,” [arXiv:1503.08043 \[hep-th\]](#).
- [71] H. Lee, D. Baumann, and G. L. Pimentel, “Non-Gaussianity as a Particle Detector,” [arXiv:1607.03735 \[hep-th\]](#).
- [72] T. Noumi, M. Yamaguchi, and D. Yokoyama, “Effective field theory approach to quasi-single field inflation and effects of heavy fields,” *JHEP* **06** (2013) 051, [arXiv:1211.1624 \[hep-th\]](#).
- [73] D. Green, M. Lewandowski, L. Senatore, E. Silverstein, and M. Zaldarriaga, “Anomalous Dimensions and Non-Gaussianity,” *JHEP* **10** (2013) 171, [arXiv:1301.2630 \[hep-th\]](#).
- [74] S. Shandera, A. Mantz, D. Rapetti, and S. W. Allen, “X-ray Cluster Constraints on Non-Gaussianity,” *JCAP* **1308** (2013) 004, [arXiv:1304.1216 \[astro-ph.CO\]](#).
- [75] E. Dimastrogiovanni, M. Fasiello, and M. Kamionkowski, “Imprints of Massive Primordial Fields on Large-Scale Structure,” *JCAP* **1602** (2016) 017, [arXiv:1504.05993 \[astro-ph.CO\]](#).
- [76] P. D. Meerburg, M. Münchmeyer, J. B. Muoz, and X. Chen, “Prospects for Cosmological Collider Physics,” [arXiv:1610.06559 \[astro-ph.CO\]](#).
- [77] D. Yamauchi, S. Yokoyama, and K. Takahashi, “Multitracer technique for galaxy bispectrum - An application to constraints on non-local primordial non-Gaussianities -,” [arXiv:1611.03590 \[astro-ph.CO\]](#).
- [78] C. Carbone, L. Verde, and S. Matarrese, “Non-Gaussian halo bias and future galaxy surveys,” *Astrophys. J.* **684** (2008) L1–L4, [arXiv:0806.1950 \[astro-ph\]](#).
- [79] R. de Putter and O. Dor, “Designing an Inflation Galaxy Survey: how to measure $\sigma(f_{\text{NL}}) \sim 1$ using scale-dependent galaxy bias,” [arXiv:1412.3854 \[astro-ph.CO\]](#).
- [80] R. de Putter, J. Gleyzes, and O. Doré, “The next non-Gaussianity frontier: what can a measurement with $\sigma(f_{\text{NL}}) \lesssim 1$ tell us about multifield inflation?,” [arXiv:1612.05248 \[astro-ph.CO\]](#).