

The hemispherical asymmetry from a scale-dependent inflationary bispectrum

Article (Accepted Version)

Byrnes, Christian T, Regan, Donough, Seery, David and Tarrant, Ewan R M (2016) The hemispherical asymmetry from a scale-dependent inflationary bispectrum. *Journal of Cosmology and Astroparticle Physics*, 2016 (06). 025. ISSN 1475-7516

This version is available from Sussex Research Online: <http://sro.sussex.ac.uk/61311/>

This document is made available in accordance with publisher policies and may differ from the published version or from the version of record. If you wish to cite this item you are advised to consult the publisher's version. Please see the URL above for details on accessing the published version.

Copyright and reuse:

Sussex Research Online is a digital repository of the research output of the University.

Copyright and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable, the material made available in SRO has been checked for eligibility before being made available.

Copies of full text items generally can be reproduced, displayed or performed and given to third parties in any format or medium for personal research or study, educational, or not-for-profit purposes without prior permission or charge, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

The hemispherical asymmetry from a scale-dependent inflationary bispectrum

Christian T. Byrnes, Donough Regan, David Seery
and Ewan R. M. Tarrant

Astronomy Centre, University of Sussex, Falmer, Brighton BN1 9QH, UK

E-mail: C.Byrnes@sussex.ac.uk, D.Regan@sussex.ac.uk, D.Seery@sussex.ac.uk

Abstract. If the primordial bispectrum is sufficiently large then the CMB hemispherical power asymmetry may be explained by a large-scale mode of exceptional amplitude which perturbs the ζ two-point function. We extend previous calculations, which were restricted to one- or two-source scenarios, by providing a method to compute the response of the two-point function in any model yielding a ‘local-like’ bispectrum. In general, this shows that it is not the reduced bispectrum $f_{\text{NL}}(k_1, k_2, k_3)$ which sources the amplitude and scale-dependence of the mode coupling but rather a combination of ‘response functions’. We discuss why it is difficult to construct successful scenarios and enumerate the fine-tunings which seem to be required. Finally, we exhibit a concrete model which can be contrived to match the observational constraints and show that to a Planck-like experiment it would appear to have $|\hat{f}_{\text{NL}}^{\text{local}}| \sim |\hat{f}_{\text{NL}}^{\text{equi}}| \sim |\hat{f}_{\text{NL}}^{\text{ortho}}| \lesssim 1$. Therefore, contrary to previous analyses, we conclude that it is possible to generate the asymmetry while respecting observational constraints on the bispectrum and low- ℓ multipoles even without tuning our location on the long-wavelength mode.

Contents

1	Introduction	1
2	Biasing the two-point function by a long-wavelength mode	4
2.1	The operator product expansion	4
2.2	The perturbed two-point function	9
2.3	Single-source models	12
3	Why it is difficult to produce a suitable response	14
3.1	Scale-dependence of the bispectrum	14
3.2	Models with constant η_σ are not viable	18
4	A working but contrived model	23
4.1	A step in η	24
4.2	The Grischuk–Zel’dovich effect	27
4.3	The shape and amplitude of the bispectrum	30
5	Conclusions	34

1 Introduction

Observations of the cosmic microwave background anisotropy on very large scales have accumulated modest evidence for a small number of anomalies in tension with our simplest picture of the early universe [1]. One of these is a roughly dipolar modulation of power which enhances the temperature fluctuations in one hemisphere. It is as if the power spectrum $\mathcal{P}(k)$ of modes contributing to the CMB anisotropy took the form

$$\mathcal{P}^{\text{obs}}(k) \approx \frac{k^3 P(k)}{2\pi^2} \left(1 + 2A(k) \hat{\mathbf{p}} \cdot \hat{\mathbf{n}} + \dots \right), \quad (1.1)$$

where $P(k)$ is the power spectrum for a statistically homogeneous and isotropic curvature perturbation synthesized by the early universe. The function $A(k)$ represents the scale-dependent amplitude of modulation, and $\hat{\mathbf{p}}$ and $\hat{\mathbf{n}}$ are unit vectors in the direction of maximum asymmetry and the line of sight, respectively, measured from Earth.

This effect has been observed with amplitude $A \approx 0.07$ on the largest scales in the WMAP [2–5] and Planck [6, 7] microwave background surveys. (See also Refs. [8, 9].) Numerically this should be interpreted as an average amplitude over a range of k contributing to the low multipoles of the CMB. On smaller scales Flender & Hotchkiss obtained the constraint $A \lesssim 0.0045$ suggesting that the modulation must exhibit a strong scale dependence [10], and improving an earlier constraint due to Hirata [11]. More recently Aiola et al. succeeded in estimating A as a function of scale, finding an approximate fit to a power law

k^n with spectral index n of roughly -0.5 [12].¹ These measurements are expected to improve in the near future owing to the arrival of new data from polarization and CMB lensing [14] which will extend the number of independent measurements of the longest visible modes.

If this effect is real, its implications for early universe models are not yet clear. Erickcek, Kamionkowski & Carroll suggested that a modulation of this kind could be produced by a fluctuation with wavelength longer than the scale of the last-scattering surface and an anomalously large amplitude [15, 16]. The linear gradient generated by this mode could produce a suitable asymmetry if the small-scale ζ two-point function responds to its presence. Any such response implies a nontrivial correlation between the long wavelength fluctuation $\delta\sigma(\mathbf{k}_L)$ and two curvature modes $\zeta(\mathbf{k})$ with $k \gg k_L$, and therefore should be controlled by ‘squeezed’ momentum configurations of the three-point function $\langle\delta\sigma(\mathbf{k}_3)\zeta(\mathbf{k}_1)\zeta(\mathbf{k}_2)\rangle$ with $k_3 \ll k_1, k_2$. But this three-point function will also contribute to the observable three-point function $\langle\zeta(\mathbf{k}_3)\zeta(\mathbf{k}_1)\zeta(\mathbf{k}_2)\rangle$ at a level which depends on the way in which $\delta\sigma$ contributes to ζ . Therefore, in this scenario, we should expect a modulation of the form (1.1) to be accompanied by some level of non-Gaussianity—although its precise amplitude cannot be predicted without further information.

Efforts have been made to quantify this effect. Working with a ‘single-source’ model in which the fluctuations of a single field dominate ζ , Lyth used the separate universe picture to compute how short-scale ζ modes would be biased by a long-wavelength perturbation [17, 18]. This calculation suggested that the response of the ζ two-point function would be characterized by the amplitude of the reduced bispectrum, f_{NL} , and that compatibility with constraints on the CMB quadrupole would probably require $|f_{\text{NL}}| \gtrsim 10$ or more.

If this estimate were applicable to the amplitude of three-point correlations in squeezed configurations then it would be incompatible with the simplest slow-roll, single-field inflationary models. Therefore achieving sufficient modulation would entail more complex dynamics, either during or after inflation. Even in these models, an amplitude of order $|f_{\text{NL}}| \gtrsim 10$ would place the scenario in tension with recent Planck constraints on the amplitude of three-point correlations in squeezed configurations.

Similar conclusions were reached by Abolhasani, Bagram, Namjoo & Firouzjahi [19–21], Kanno, Sasaki & Tanaka [22] and later Kobayashi, Cortês & Liddle [23]. A variant of the scenario, based on the response being generated by a large trispectrum rather than bispectrum, was suggested by Kenton, Mulryne & Thomas [24]. But none of these analyses explicitly accounted for the strong scale-dependence of three-point correlations which is presumably entailed by the strong scaling of $A(k)$. At a minimum this would require specification of the scale at which any amplitude constraint was intended to apply. Moreover the separate universe approach used in these investigations does not make manifest which momentum configurations of the three-point function are being invoked. This is not merely a pedantic point. If strong scaling is present the amplitude of correlation on squeezed configurations can differ by orders of magnitude compared to the amplitude on equilateral configurations—

¹Aiola et al. reported their constraint as an ℓ -dependent modulation $A(\ell)$ of the angular power spectrum C_ℓ . A power-law primordial spectrum of the form $P(k) \sim k^{-(3+n)}$ induces an angular power spectrum of the form $C_\ell \sim \ell^{-(2+n)}$. Therefore the observed modulation $A(\ell)$ implies a primordial power spectrum modulated by the same power law. This result was noted recently by Adhikari, Shandera & Erickcek [13].

even for configurations of the same overall scale.² For example, in Ref. [17] it was suggested that the amplitude relevant for determining the response of the ζ two-point function would come from equilateral configurations. The amplitude in squeezed configurations might then be very different. Since it is the squeezed configurations which mostly contribute to present observational constraints, for example through their contribution to CMB estimators for the amplitude $f_{\text{NL}}^{\text{local}}$ of the ‘local’ template, this could dramatically change our conclusions regarding the viability of the model or even the requirement for dynamics beyond slow-roll, single-field inflation.

Summary.—In this paper we have three major goals.

First, we explain how to compute the response of the ζ two-point function to long-wavelength perturbations in scenarios more general than the single-source model. Our approach shows explicitly how the response depends on information embedded in the squeezed limit of the three-point function $\langle \delta\sigma(\mathbf{k}_3)\zeta(\mathbf{k}_1)\zeta(\mathbf{k}_2) \rangle$, and clarifies which momentum configurations are relevant. It does not rely on the separate universe method, although in some circumstances that could be used to compute the required correlation functions.

Second, despite the effort which has been invested in studying the Erickcek et al. scenario, it is still unclear what is entailed in an inflationary model (perhaps extended by a later curvaton phase) that achieves a suitable asymmetry by this mechanism. Explicit models were studied by McDonald [25] and Kanno et al. [22]; but if observation forces us to consider early-universe scenarios for the origin of the asymmetry—and this is not yet clear—we should like qualitative guidance concerning the features to be expected. We explain in general terms why it is difficult to manufacture a scenario which produces a bispectrum with suitable scale-dependence without simultaneously producing other undesirable features, such as an unacceptably tilted power spectrum or a large trispectrum.

Third, we address the issue of observational constraints. As explained above, the strong running with scale entailed by the scale-dependence of $A(k)$ makes it unclear how large a bispectrum amplitude is acceptable. The estimates made in Refs. [17, 18, 22, 23] suggest that the required amplitude may be too large, but these cannot be compared directly to constraints reported by Planck or WMAP. To make a comparison we must determine how the amplitude of the bispectrum varies with scale and squeezing. We exhibit a concrete (but contrived) model which satisfies current observational constraints. The construction of this model exemplifies the general difficulties encountered in building a successful scenario—but having done so, we use it to demonstrate the shape and magnitude of the three-point correlations which it produces. We use these to estimate how a Planck-like experiment would view the bispectrum through the response of the estimators for $\hat{f}_{\text{NL}}^{\text{local}}$, $\hat{f}_{\text{NL}}^{\text{equi}}$ and $\hat{f}_{\text{NL}}^{\text{ortho}}$. It is then possible to address the viability of the scenario.

Structure.—This paper is structured as follows. In §2 we develop a formalism which can be used to compute how an arbitrary long-wavelength mode biases the ζ two-point function.

²Here (and in the remainder of this paper) we are taking the *scale* of a momentum configuration specified by the triangle $\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3 = 0$ to mean its perimeter $k_t = k_1 + k_2 + k_3$. Its *shape* is specified by the ratios k_i/k_t of its sides to k_t . For squeezed configurations one side becomes much smaller than the other two. Taking this side to be k_3 that gives $k_3/k_t \ll 1$ while the other two sides have roughly $k_1/k_t \sim k_2/k_t \sim 0.5$.

The key results are collected in §§2.1–2.2. Our formalism applies to an arbitrary inflationary model provided the squeezed limit of the bispectrum is not suppressed in a sense to be made precise below. It does not make use of the slow-roll approximation. In §2.3 we explain how to relate our approach to the existing literature, including Refs. [17, 18, 22–24].

In §3 we give a heuristic argument explaining why it is difficult to build inflationary models with the correct response, even if we allow for effects subsequent to inflation such as a curvaton era. In §3.1 we explain what is required to produce a bispectrum with the correct scaling properties. In §3.2 we show that the simplest of these scenarios, where the scale dependence is generated by a large negative η -parameter, has difficulty generating a bispectrum of sufficient amplitude during the inflationary epoch. We sketch the problems encountered if we attempt to go beyond the simplest scenario.

§4 describes a working model which produces a suitable response during inflation by introducing a sharp feature in the potential. This evades the constraints discussed in §3, but also exemplifies the tunings which seem required to construct a viable model. We outline the model in §4.1 and give numerical results for the response of the two-point function to biasing. In §4.2 we compute constraints on the amplitude of the large-scale mode from the Grischuk–Zel’dovich effect, and in §4.3 we discuss constraints from the bispectrum including estimates for the response of the Planck estimators $\hat{f}_{\text{NL}}^{\text{local}}$, $\hat{f}_{\text{NL}}^{\text{equi}}$ and $\hat{f}_{\text{NL}}^{\text{ortho}}$. Finally, we summarize our conclusions in §5.

Notation.—Throughout this paper we adopt units in which $c = \hbar = 1$. The reduced Planck mass is defined by $M_{\text{P}}^2 = (8\pi G)^{-1/2}$, where G is Newton’s gravitational constant. We work with a collection of light scalar fields and their momenta, initially indexed by Greek labels $\{\alpha, \beta, \dots\}$. In order to work with compact formulae we adapt this summation convention when multiple times of evaluation are under discussion, as described in §3.1.

We collect our notational conventions in Table 1.

2 Biasing the two-point function by a long-wavelength mode

In this section we obtain a formula for the response of a short-wavelength two-point function to the presence of an underlying perturbation with much longer wavelength. Our result will be valid, up to gradient-suppressed corrections, to linear order in the amplitude of the long-wavelength mode but non-perturbatively in the short-wavelength modes. We work with the ζ two-point function because this is the case to which we will eventually apply our result, but the method is general and can be used to study the biasing of any n -point function.

2.1 The operator product expansion

Consider a region of spacetime with comoving spatial extent M , within which we wish to predict the ζ two-point function; see Fig. 1. We imagine that this region is enclosed within a uniform larger patch of extent $L \gg M$, and we suppose that an early-universe mechanism such as inflation has seeded a set of statistically isotropic and homogeneous fluctuations within this large patch. In the scenario of Erickcek et al., the particular realization within the L -patch contains a rare large-amplitude mode with wavenumber $k_{\text{L}} \ll 1/M$.

notation	meaning	definition
$P(k)$	power spectrum for $\langle\zeta\zeta\rangle$	Eq. (2.4)
$\mathcal{P}(k)$	dimensionless power spectrum of $\langle\zeta\zeta\rangle$	Eq. (1.1)
$B(k_1, k_2, k_3)$	bispectrum for $\langle\zeta\zeta\zeta\rangle$	Eq. (2.20)
$T(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4)$	connected trispectrum for $\langle\zeta\zeta\zeta\zeta\rangle$	Eq. (2.23)
$f_{\text{NL}}(k_1, k_2, k_3)$	reduced bispectrum for $\langle\zeta\zeta\zeta\rangle$	Eq. (2.19)
$\tau_{\text{NL}}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4)$	amplitude of τ_{NL} -mode in T	Eq. (2.22)
$\mathcal{P}_\sigma(k)$	dimensionless power spectrum for $\langle\sigma\sigma\rangle$	Eq. (2.13)
$\mathcal{P}_{\zeta_\sigma}(k)$	dimensionless power spectrum for $\langle\zeta_\sigma\zeta_\sigma\rangle$, where $\zeta_\sigma = N_\sigma\delta\sigma$	Eq. (2.25)
$\Sigma^{\alpha\beta}$	power spectrum for $\langle\delta\phi^\alpha\delta\phi^\beta\rangle$	Eq. (2.10)
$B^\lambda(k_1, k_2, k_3)$	bispectrum for $\langle\delta\phi^\lambda\zeta\zeta\rangle$	Eq. (2.8)
$A(k)$	amplitude of asymmetry	Eq. (1.1)
$\rho_\mu(k)$	linear response function for $\delta\phi^\mu$	Eq. (2.6)
$\rho_\zeta(k)$	linear response function for ζ	Eq. (2.28)
E	enhancement or exceptionality of long-wavelength mode	Eq. (2.14)
α, \mathbf{k}_L	scale of long-wavelength mode	Eq. (2.15)
x_{ls}	comoving distance to last-scattering surface	p. 10
$\hat{\mathbf{n}}$	orientation of line-of-sight from Earth	p. 10
$\hat{\mathbf{p}}$	orientation of exceptional mode	Eq. (2.15)
R	relative contribution of σ to $\mathcal{P}(k)$	Eq. (2.25)
$N_\alpha, N_{\alpha\beta}$	gauge transformation from field fluctuations to ζ	Eq. (2.27)
$\Gamma_a^\alpha, \Gamma_{ab}^\alpha$	separate-universe coefficients, depending on <i>two</i> times	Eq. (3.1)
$u^\alpha_\beta, u^\alpha_{\beta\gamma}$	transport equation coefficients	Eqs. (3.3a), (3.3b)
η_σ	η -parameter $V_{\sigma\sigma}/3H^2$ for σ	
ξ_σ	parameter $M_{\text{P}}V_{\sigma\sigma\sigma}/3H^2$ for σ	Eq. (3.8b)
k_t	perimeter of momentum n -gon	p. 15
$k_{\ell=1}$	roughly corresponds to $\ell = 1$, $k_{\ell=1} \approx 1/14,000 \text{ Mpc}^{-1}$	Eq. (4.2)

Table 1: Notation used in this paper. Time-dependent quantities such as n -point functions, power spectra and bispectra are always evaluated at the time of interest unless otherwise specified. Note that this includes the gauge-transformation coefficients $N_\alpha, N_{\alpha\beta}$.




Figure 1: Modulation of the power spectrum measured in an M -sized patch embedded within a larger L -sized patch. The L -patch is crossed by a long-wavelength mode (2.14) whose amplitude is enhanced above the typical amplitude by a factor E . In applications, the M -patch [located at the aggregate coordinate \mathbf{x}_+ ; see Eq. (2.2)] would be centred on the last-scattering surface.

Within the M -patch, the two-point function would respond to an infinitely long wavelength perturbation as if it were a shift in the zero-mode of the fields. Therefore to linear order in the amplitude of the perturbation,

$$\langle \zeta(\mathbf{x}_1)\zeta(\mathbf{x}_2) \rangle_M = \langle \zeta(\mathbf{x}_1)\zeta(\mathbf{x}_2) \rangle_L + \delta\phi^\mu \frac{\partial}{\partial\phi^\mu} \langle \zeta(\mathbf{x}_1)\zeta(\mathbf{x}_2) \rangle_L + \dots, \quad (2.1)$$

where ‘ \dots ’ denotes terms of second-order or higher in $\delta\phi$ which we have neglected. Except in §3 we are not using the slow-roll approximation, so $\delta\phi^\mu$ runs over the perturbations in the scalar fields and their momenta. The same is true for $\partial/\partial\phi^\mu$.

Our interest is in perturbations with large but finite wavelength. For such perturbations (2.1) represents the beginning of a series describing the response of $\langle \zeta\zeta \rangle$ to the position-dependent fluctuation $\delta\phi^\mu(\mathbf{x}_+)$, where \mathbf{x}_+ is the aggregate position of the M -patch surrounding \mathbf{x}_1 and \mathbf{x}_2 [26–29],

$$\mathbf{x}_+ = \frac{\mathbf{x}_1 + \mathbf{x}_2}{2}. \quad (2.2)$$

The two-point function will respond not only to the displacement $\delta\phi^\mu$ but also other local operators built from its gradients such as $\partial^2\delta\phi^\mu$. Therefore

$$\begin{aligned} \langle \zeta(\mathbf{x}_1)\zeta(\mathbf{x}_2) \rangle_M &= \langle \zeta(\mathbf{x}_1)\zeta(\mathbf{x}_2) \rangle_L + \delta\phi^\mu(\mathbf{x}_+) \frac{\partial}{\partial\phi^\mu} \langle \zeta(\mathbf{x}_1)\zeta(\mathbf{x}_2) \rangle_L \\ &+ \partial^2\delta\phi^\mu(\mathbf{x}_+) \frac{\partial}{\partial(\partial^2\phi^\mu)} \langle \zeta(\mathbf{x}_1)\zeta(\mathbf{x}_2) \rangle_L + \dots, \end{aligned} \quad (2.3)$$

where ‘ \dots ’ denotes contributions from other local operators which we have not written explicitly. In Eq. (2.3) all quantities are evaluated at the same time. A similar expansion was used by Mirbabayi & Simonović [30].

If $\delta\phi^\mu$ contains only long-wavelength contributions then its gradients will be suppressed. Where these suppressed terms can be neglected the second term in Eq. (2.3) will furnish the dominant response and the two-point function reacts to all long-wavelength modes in nearly the same way. But if $(\partial/\partial\phi^\mu)\langle\zeta\zeta\rangle$ is small or zero then the leading correction may come instead from a term such as $\partial^2\delta\phi^\mu$. In these cases the two-point function responds differently to perturbations of different wavelengths.

In this paper we focus on scenarios in which the dominant long-wavelength response comes from the $\delta\phi^\mu$ operator. This includes most models of interest because if the dominant response involves gradients it will be suppressed by powers of the ratio k_L/k , where k is a typical wavenumber contributing to $\langle\zeta\zeta\rangle$ and k_L is a long-wavelength mode in $\delta\phi^\mu$. We will see later that this suppression would make it difficult to generate a suitable asymmetry without a bispectrum of very large amplitude. In addition the scale-dependence may be incorrect; see Ref. [13].

The $\delta\phi^\mu$ operator in Eq. (2.3) was used by Schmidt & Hui to compute the linear response of the two-point function to immersion within a bath of long-wavelength modes [31].³ Their scenario was developed by Adhikari, Shandera & Erickcek [13]. Later Pajer, Schmidt & Zaldarriaga [27] demonstrated that the precise linear combination of \mathbf{x}_1 and \mathbf{x}_2 appearing in Eq. (2.2) is immaterial in the squeezed limit up to corrections of order $O(k^2)$. The discussion given here is closest to that of Namjoo et al. [19] and Kenton & Mulryne [32], both of whom effectively used (2.3) to study biasing of a three-point function by a long wavelength mode. A similar method has been proposed by Chiang et al. to measure the squeezed limit of the primordial bispectrum using large-scale structure [33].

Response function.—To use Eq. (2.3) we discard all gradient-suppressed terms and take its Fourier transform within the M -patch, leaving the location \mathbf{x}_+ of this patch fixed. If the power spectrum $P(k)$ within the L -patch satisfies

$$\langle\zeta(\mathbf{k}_1)\zeta(\mathbf{k}_2)\rangle_L = (2\pi)^3\delta(\mathbf{k}_1 + \mathbf{k}_2)P(k), \quad (2.4)$$

where $k = |\mathbf{k}_1| = |\mathbf{k}_2|$ is the common magnitude of \mathbf{k}_1 and \mathbf{k}_2 , the result can be written

$$\langle\zeta(\mathbf{k}_1)\zeta(\mathbf{k}_2)\rangle_M = (2\pi)^3\delta(\mathbf{k}_1 + \mathbf{k}_2)P(k)\left(1 + \delta\phi^\mu(\mathbf{x}_+)\rho_\mu(k) + \dots\right). \quad (2.5)$$

In Eq. (2.5) we have introduced the *linear response function* $\rho_\mu(k)$, defined to be a rescaled derivative of $P(k)$ with respect to the zero-modes within the L -patch,

$$\rho_\mu(k) \equiv \frac{1}{P(k)} \frac{\partial P(k)}{\partial\phi^\mu}. \quad (2.6)$$

³Schmidt & Hui’s scenario is very similar to the one considered in this paper, except that we will take a single long-wavelength mode to have an exceptional amplitude. In Schmidt & Hui’s scenario all the long-wavelength modes have a typical amplitude and the resulting anisotropy receives contributions from all of them. What is observed in a typical region with scale comparable to the horizon at the time of last scattering therefore depends on the distribution of large-scale modes, which was studied by Adhikari et al. [13]. They found the tail to be sufficiently broad that an exceptional amplitude might not be required if the bispectrum amplitude is sufficiently large on the largest observable scales. See also the discussion in §4.3 and §5.

It is a function only of k , and as we have explained it does not depend on the wavenumber of the source. If we require terms of higher-order in the amplitude of the long-wavelength mode then (2.5) could be extended by defining higher-order response functions $\rho_{\mu\nu} = P(k)^{-1}\partial_\mu\partial_\nu P(k)$, $\rho_{\mu\nu\lambda} = P(k)^{-1}\partial_\mu\partial_\nu\partial_\lambda P(k)$, and so on, which measure the quadratic, cubic or higher terms in the expansion.

Once we have obtained $\rho_\mu(k)$ it is sufficient (if nonzero) to characterize linear biasing within the M -patch. We can compute it by any convenient method—for example, by constructing a numerical derivative, or even by analytic differentiation if a closed-form expression can be found. However, as explained in §1, we expect that the response of the two-point function should be controlled by squeezed configurations of the three-point function $\langle\delta\phi^\mu\zeta\zeta\rangle$. To determine this relationship we return to (2.3), but now regarded as an operator product expansion for the bilinear $\zeta(\mathbf{x}_1)\zeta(\mathbf{x}_2)$ within the L -patch. Correlations between this bilinear and a distant point \mathbf{x}_3 are controlled by the expansion. If we arrange that $\langle\delta\phi^\mu(\mathbf{x})\rangle_L = 0$ then Eq. (2.3) asserts that the leading contribution comes from modulation of the two-point function,

$$\langle\delta\phi^\lambda(\mathbf{x}_3)\zeta(\mathbf{x}_1)\zeta(\mathbf{x}_2)\rangle_L \approx \langle\delta\phi^\lambda(\mathbf{x}_3)\delta\phi^\mu(\mathbf{x}_+)\rangle_L \frac{\partial}{\partial\phi^\mu} \langle\zeta(\mathbf{x}_1)\zeta(\mathbf{x}_2)\rangle_L \quad \text{if } |\mathbf{x}_3 - \mathbf{x}_+| \gg |\mathbf{x}_1 - \mathbf{x}_2|. \quad (2.7)$$

We now take the Fourier transform within the L -patch. For k_3 much less than k_1, k_2 the dominant contribution to the Fourier integral will come from spatial configurations for which Eq. (2.7) describes the behaviour of the three-point function. Defining the bispectrum of the mixed correlation function by

$$\langle\delta\phi^\lambda(\mathbf{k}_3)\zeta(\mathbf{k}_1)\zeta(\mathbf{k}_2)\rangle_L = (2\pi)^3\delta(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3)B^\lambda(k_1, k_2, k_3) \quad (2.8)$$

it can be shown that the Fourier transform yields

$$B^\lambda(k_1, k_2, k_3) \approx \Sigma^{\lambda\mu}(k_3)\partial_\mu P(|\mathbf{k}_1 - \mathbf{k}_3/2|) \quad \text{if } k_3 \ll k_1, k_2, \quad (2.9)$$

where $\Sigma^{\alpha\beta}$ is the power spectrum of the field fluctuations within the L -patch,

$$\langle\delta\phi^\alpha(\mathbf{k}_1)\delta\phi^\beta(\mathbf{k}_2)\rangle_L = (2\pi)^3\delta(\mathbf{k}_1 + \mathbf{k}_2)\Sigma^{\alpha\beta}(k_1). \quad (2.10)$$

Eq. (2.9) shows that the response function can be extracted from knowledge of $B^\lambda(k_1, k_2, k_3)$ and $\Sigma^{\alpha\beta}(k)$. To leading order in $k_3/k_1 \approx k_3/k_2$ the required combination is

$$\rho_\mu(k) \approx \frac{1}{P(k)}[\Sigma^{-1}(k_3)]_{\mu\lambda}B^\lambda(k, k, k_3) \quad \text{if } k_3 \ll k. \quad (2.11)$$

Provided the $\delta\phi^\mu$ operator dominates the long-wavelength response, the operator product expansion guarantees that the right-hand side becomes independent of k_3 for sufficiently squeezed configurations. In this case $\rho_\mu(k)$ can be determined approximately from any suitable configuration of this type. The wavenumber k represents the nearly equal magnitude of the short modes, $k \approx k_1 \approx k_2$.

Alternatively, it may happen that (2.11) does not approach a nonzero limit for small k_3 . This indicates that the $\delta\phi^\mu$ operator in (2.3) did not control the response to long-wavelength

perturbations, and a higher-order operator or one of the gradient-suppressed contributions is instead dominant. Examples of bispectra for which this occurs include the equilateral and orthogonal templates, because these diverge more slowly than the power-spectrum Σ as $k_3 \rightarrow 0$. In such cases it is possible to modify the discussion in this section to extract a corresponding response function, but we will not pursue this possibility; for the reasons explained above the response is suppressed, and a bispectrum of large amplitude is required to generate a suitable asymmetry. For an alternative approach to these templates see Adhikari, Shandera & Erickcek [13].

For the models to be studied in this paper Eq. (2.11) is sufficient. These have ‘local-like’ bispectra in the sense that correlations in the squeezed limit are not suppressed by powers of k_L/k , but there is no requirement that the bispectrum shape is a close match to the local template. Indeed, to generate a suitable scale-dependent asymmetry $A(k)$ we will require departures from the local shape. In these models Eq. (2.11) shows that the response of the two-point function depends on the squeezed limit of the mixed bispectrum $B^\lambda(k_1, k_2, k_3)$, although because $\rho_\mu(k)$ is itself independent of k_3 it cannot depend on ratios such as $k_3/k_t \sim k_3/k$ which measure the squeezing. We also note that no part of our analysis required the slow-roll approximation.

Finally, if we wished to include the quadratic or cubic response functions $\rho_{\mu\nu}, \rho_{\mu\nu\lambda}$ then these could be estimated in a similar way, by studying the double-soft limit of $\langle \delta\phi^\alpha \delta\phi^\beta \zeta \zeta \rangle$ or the triple-soft limit of $\langle \delta\phi^\alpha \delta\phi^\beta \delta\phi^\gamma \zeta \zeta \rangle$ —in each case as the momenta associated with the field fluctuations become much smaller than those carried by ζ .

OPE determines collapsed trispectrum.—The operator product expansion determines certain other n -point functions on a subset of configurations. In particular, an analysis similar to that of Eqs. (2.7) and (2.9) shows that Eq. (2.3) determines the four-point function on ‘collapsed’ configurations in terms of the response functions,

$$\langle \zeta(\mathbf{k}_1) \zeta(\mathbf{k}_2) \zeta(\mathbf{k}_3) \zeta(\mathbf{k}_4) \rangle_L^{\text{connected}} \approx (2\pi)^3 \delta\left(\sum_i \mathbf{k}_i\right) \Sigma^{\alpha\beta}(|\mathbf{k}_3 + \mathbf{k}_4|) \rho_\alpha(k_2) \rho_\beta(k_4) P(k_2) P(k_4) \quad (2.12)$$

if $|\mathbf{k}_3 + \mathbf{k}_4| \ll k_2, k_4$. Further formulas can be found for increasingly restrictive configurations of the higher n -point functions, but we do not study these in detail because there is no imminent prospect of restrictive observational constraints.

2.2 The perturbed two-point function

We now estimate the perturbation in the M -patch two-point function $\langle \zeta(\mathbf{k}_1) \zeta(\mathbf{k}_2) \rangle_M$ produced by a collection of long-wavelength modes crossing the L -patch. In this paper we are principally interested in the case where there is a rare fluctuation with enhanced amplitude associated with a single wavenumber \mathbf{k}_L , and for simplicity we will assume it is present in just one species which we denote σ . The case of multiple contributing species (or wavenumbers) can be handled by obvious modifications³ of our formulae.

Neglecting correlations between species, the typical variance of fluctuations in σ will be given by

$$\mathcal{P}_\sigma(k) = \frac{k^3}{2\pi^2} \Sigma^{\sigma\sigma}(k) \quad (\text{no sum}). \quad (2.13)$$

If this typical amplitude is enhanced by a factor E , the spatial variation of the long-wavelength mode will be described by

$$\delta\sigma(\mathbf{x}) \approx E\mathcal{P}_\sigma^{1/2}(k_L) \cos(\mathbf{k}_L \cdot \mathbf{x} + \vartheta) \quad (2.14)$$

where ϑ is a phase which will vary between realizations.

Contributions to the CMB anisotropy are generated by M -patches located on the last scattering surface at comoving distance x_{ls} , for which $\mathbf{x} = x_{\text{ls}}\hat{\mathbf{n}}$. (See Fig. 1.) The unit vector $\hat{\mathbf{n}}$ selects the line-of-sight from Earth. If \mathbf{k}_L corresponds to a spatial scale larger than x_{ls} then $\mathbf{k}_L \cdot \mathbf{x} \lesssim 1$, and in order that the linear gradient dominates we should require $\mathbf{k}_L \cdot \mathbf{x} \lesssim 10^{-1}$. In principle there is no upper bound on the spatial scale associated with k_L because a reduction in the gradient can always be compensated by adjusting the amplitude. However, unless \mathcal{P}_σ is rather red, moving k_L to larger scales will likely require an increase in the exceptionality E .

In the Erickcek et al. scenario we should prefer *some* exceptionality to justify our focus on a single wavenumber. The alternative, that all modes have nearly equal amplitude, is the scenario of Schmidt & Hui and Adhikari et al. [13, 31, 34]. But although some exceptionality is desirable the probability of obtaining an exceptional fluctuation decreases as E increases, and if we require E to be very large the scenario becomes unattractive. We will generally assume that the best arrangement is to set E as small as possible, but no lower than perhaps $O(10)$. Notice that in doing so we are electing to trade a $\sim 3\sigma$ discrepancy for *at least* a $\sim 10\sigma$ fluctuation. Therefore any realistic explanation of the hemispherical asymmetry which deploys this mechanism will likely require some way to manufacture a suitable exceptional amplitude without depending on Gaussian statistics. Possible examples include the proposals of Refs. [35, 36].

To parametrize the scale \mathbf{k}_L we write

$$\mathbf{k}_L = \frac{2\pi}{x_{\text{ls}}} \alpha \hat{\mathbf{p}}, \quad (2.15)$$

where $\hat{\mathbf{p}}$ is a unit vector and $\alpha < 1$ characterizes the ratio of x_{ls} to the spatial scale associated with k_L . After expanding in $\mathbf{k}_L \cdot \mathbf{x}$, Eq. (2.5) gives an expression for the modulated two-point function,

$$\langle \zeta(\mathbf{k}_1) \zeta(\mathbf{k}_2) \rangle_M \approx \langle \zeta(\mathbf{k}_1) \zeta(\mathbf{k}_2) \rangle_L \left\{ 1 - 2C(k) + 2A(k) \frac{\mathbf{x} \cdot \hat{\mathbf{p}}}{x_{\text{ls}}} + \dots \right\}, \quad (2.16)$$

where the quantities $A(k)$ and $C(k)$ are defined by

$$A(k) = \pi \alpha E \mathcal{P}_\sigma^{1/2}(k_L) \rho_\sigma(k) \sin \vartheta \quad (2.17a)$$

$$C(k) = -\frac{1}{2} E \mathcal{P}_\sigma^{1/2}(k_L) \rho_\sigma(k) \cos \vartheta. \quad (2.17b)$$

This is of the required form (1.1) with $\hat{\mathbf{n}} = \mathbf{x}/x_{\text{ls}}$. The scale-dependence of the modulation amplitude $A(k)$ is inherited from the scale-dependence of the response $\rho_\sigma(k)$.

Suppression of low multipoles.—Eq. (2.16) shows that in addition to the dipolar modulation there is an overall shift in amplitude due to $C(k)$ [18]. Assuming the spatial dependence of

the long-wavelength mode is described by (2.14), this is related to $A(k)$ via

$$C(k) = -\frac{A(k)}{2\pi\alpha} \cot \vartheta. \quad (2.18)$$

For more general dependence the coefficient of proportionality is altered, but the scaling $C(k) \sim A(k)/\alpha$ remains. In Refs. [17, 18, 37] it was suggested that $C(k)$ could be used to explain a second CMB anomaly—the observed low CMB quadrupole. Schwarz et al. proposed that a viable explanation of the hemispherical asymmetry should explain at least one other anomaly, so this outcome would be desirable [1]. Unfortunately, Eq. (2.18) will make $C(k)$ larger than required if α is small. Assuming the reported BICEP2 measurement of $r \sim 0.2$ (now known to have been confused by dust), Contaldi et al. estimated that $C(k)$ could be roughly of order 0.14 [38]. However, precise constraints do not seem to have been reported in more general circumstances. We will assume it should not be much larger than ~ 0.1 , and it should preferably not be negative.

This is an obstacle for construction of viable models. If we assume the spatial dependence in (2.14) then the amplitude $|C(k)|$ can be reduced by tuning ϑ . For more general spatial dependence it requires tuning the Taylor coefficient of order $(\mathbf{k}_L \cdot \mathbf{x})^0$ with respect to that of order $(\mathbf{k}_L \cdot \mathbf{x})^1$.

Bi- and trispectrum amplitudes.—If σ dominates the bi- and trispectrum of ζ then it is possible to go further and relate the amplitude of the asymmetry $A(k)$ to the degree of correlation in, respectively, squeezed and collapsed configurations of the three- and four-point functions. We will see in §§3–4 that this situation is realized in a large class of successful scenarios. To measure the amplitude of three-point correlations we define the reduced bispectrum $f_{\text{NL}}(k_1, k_2, k_3)$,

$$\frac{6}{5} f_{\text{NL}}(k_1, k_2, k_3) = \frac{B(k_1, k_2, k_3)}{P(k_1)P(k_2) + P(k_1)P(k_3) + P(k_2)P(k_3)}, \quad (2.19)$$

where $B(k_1, k_2, k_3)$ is the bispectrum for the three-point function of ζ ,

$$\langle \zeta(\mathbf{k}_1)\zeta(\mathbf{k}_2)\zeta(\mathbf{k}_3) \rangle_L = (2\pi)^3 \delta(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B(k_1, k_2, k_3). \quad (2.20)$$

In the squeezed limit $k_3 \ll k_1, k_2$ this becomes approximately

$$\frac{6}{5} f_{\text{NL}}(k, k, k_3) \approx \frac{B(k, k, k_3)}{2P(k_3)P(k)} \quad \text{if } k \ll k_3. \quad (2.21)$$

For four-point correlations we define

$$\tau_{\text{NL}}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4) \approx \frac{T(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4)}{4P(|\mathbf{k}_3 + \mathbf{k}_4|)P(k_2)P(k_4)} \quad \text{if } |\mathbf{k}_3 + \mathbf{k}_4| \ll k_2, k_4, \quad (2.22)$$

with $T(k_1, k_2, k_3, k_4)$ now the trispectrum defined by the connected four-point function of ζ ,

$$\langle \zeta(\mathbf{k}_1)\zeta(\mathbf{k}_2)\zeta(\mathbf{k}_3)\zeta(\mathbf{k}_4) \rangle_L^{\text{connected}} = (2\pi)^3 \delta(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3 + \mathbf{k}_4) T(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4). \quad (2.23)$$

Taking our position on the long-wavelength mode to be generic, so that $\sin \vartheta \sim \cos \vartheta \sim O(1)$, Eq. (2.9) predicts that for observably squeezed configurations—including those which

contribute to the CMB—we can write the reduced bispectrum in terms of the amplitude $A(k)$,

$$\frac{6}{5} f_{\text{NL}}(k, k, k_3) \approx \frac{A(k)}{2\pi\alpha E} R^{1/2}(k_3) \left(\frac{\mathcal{P}_\sigma(k_3)}{\mathcal{P}_\sigma(k_L)} \right)^{1/2} \frac{1}{\mathcal{P}^{1/2}(k_3)} \quad \text{if } k_L < k_3 \ll k, \quad (2.24)$$

where $\mathcal{P}(k)$ is the dimensionless version of the ζ power spectrum $P(k)$, and R measures the contribution of σ to the total power spectrum,

$$R(k) \equiv \frac{\mathcal{P}_{\zeta_\sigma}(k)}{\mathcal{P}(k)}. \quad (2.25)$$

Here ζ_σ is the linear contribution of σ to ζ . Despite its appearance (2.24) does not depend on $\mathcal{P}^{1/2}(k_L)$ since Eq. (2.17a) shows that $A(k)$ is proportional to it. Similarly, Eq. (2.12) predicts that for collapsed configurations

$$\tau_{\text{NL}}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4) \approx \frac{A(k_2)A(k_4)}{4\pi^2\alpha^2 E^2} \frac{\mathcal{P}_\sigma(|\mathbf{k}_3 + \mathbf{k}_4|)}{\mathcal{P}_\sigma(k_L)} \frac{1}{\mathcal{P}(|\mathbf{k}_3 + \mathbf{k}_4|)} \quad \text{if } k_L < |\mathbf{k}_3 + \mathbf{k}_4| \ll k_2, k_4. \quad (2.26)$$

We caution that these relations hold *only* if a single field σ dominates the bi- and trispectra of ζ . Note that this does not restrict Eqs. (2.24) and (2.26) to single-source models in the sense of Refs. [17, 18, 23], where the scale-dependent field σ dominates all the correlation functions of ζ , because there is no requirement for σ to dominate the two-point function. This more general scenario has already been studied by Kanno et al. [22]. They gave a formula equivalent to Eq. (2.24) (although derived by a different method and not including explicit scale-dependence), but we believe (2.26) is new.

Where Eqs. (2.24) and (2.26) apply, they predict that the amplitudes f_{NL} and τ_{NL} are enhanced by inverse powers of α and the ζ power-spectrum \mathcal{P} , but suppressed by powers of the asymmetry A and the exceptionality E .

2.3 Single-source models

Eqs. (2.16), (2.17a) and (2.17b) (and their obvious generalization to the case where the modulation is sourced by perturbations in several fields) explain how to calculate the asymmetry produced by the Erickcek et al. scenario in an arbitrary local-like model.

Refs. [17, 18, 23] studied biasing of the two-point function in the special case of single-source models in the sense defined above. These references reported results equivalent to (2.17a) with $\rho_\sigma(k)$ replaced by $12f_{\text{NL}}(k, k, k)/5$, and $\mathcal{P}_\sigma(k_L)$ replaced by the amplitude of ζ fluctuations, $\mathcal{P}_\zeta(k_L)$. Here $f_{\text{NL}}(k_1, k_2, k_3)$ is the reduced bispectrum defined in Eq. (2.19).

To make a connexion with the results of Refs. [17, 18, 23] we reformulate our analysis in terms of the response of the ζ two-point function to a long-wavelength ζ fluctuation. In a single-source model this leads to no loss of generality since we need not distinguish between the field fluctuation which dominates ζ , and ζ itself. In any model (single-source or not) the response to a ζ fluctuation can be obtained by projecting $\rho_\mu(k)$ along the field-space unit vector corresponding to the orientation of ζ .⁴ To obtain this, note that ζ can always be

⁴By ‘field space’ we mean the space spanned by the field values and their momenta.

expressed as a composite of the field fluctuations defined on spatially flat hypersurfaces,

$$\zeta(\mathbf{k}) = N_\alpha \delta\phi^\alpha(\mathbf{k}) + \frac{1}{2} N_{\alpha\beta} \int \frac{d^3q}{(2\pi)^3} \delta\phi^\alpha(\mathbf{q}) \delta\phi^\beta(\mathbf{k} - \mathbf{q}) + \dots, \quad (2.27)$$

where ‘ \dots ’ denotes terms of higher order in the field fluctuations which have been omitted. The fluctuations ζ and $\delta\phi^\alpha$ are to be evaluated at the same time, and the gauge transformation coefficients N_α , $N_{\alpha\beta}$ can be found in the literature [39–41]. As above we caution that we are not using the slow-roll approximation and therefore the labels $\{\alpha, \beta, \dots\}$ run over both the scalar fields and their momenta.

Response to ζ .—At linear order $\zeta \approx N_\alpha \delta\phi^\alpha$. Therefore we can associate the curvature perturbation with a fluctuation oriented along the field-space unit vector $\hat{n}_\alpha = N_\alpha / (N_\lambda N_\lambda)^{1/2}$. In this expression and what follows, summation over repeated field-space indices is implied, irrespective of their ‘up’ or ‘down’ position. (We are working with a trivial field-space metric.) It follows that the response to a long-wavelength ζ fluctuation can be computed by the projection

$$\rho_\zeta(k) = \frac{\hat{n}_\alpha \rho_\alpha(k)}{(N_\lambda N_\lambda)^{1/2}} \approx \frac{N_\alpha [\Sigma^{-1}(k_3)]_{\alpha\mu} B^\mu(k, k, k_3)}{(N_\lambda N_\lambda) P(k)} \quad \text{if } k \ll k_3. \quad (2.28)$$

The linear formula for ζ shows that $P(k) = N_\alpha N_\beta \Sigma^{\alpha\beta}(k)$ up to corrections of higher order in the field fluctuations. Also, in this limit $B(k_1, k_2, k_3) \approx N_\mu B^\mu(k_1, k_2, k_3)$. Taken together, these relations suggest that Eq. (2.28) is *related* to the squeezed limit of the reduced bispectrum (2.21). Nevertheless it is not exactly the same. In Eq. (2.21) the combination $P(k_3) \approx N_\lambda N_\mu \Sigma^{\lambda\mu}(k_3)$ appears in the denominator, whereas in Eq. (2.28) only the contraction $N_\lambda N_\lambda$ appears there directly; the power spectrum factor appears in the numerator as a matrix inverse interposed between N_α and $B^\alpha(k_1, k_2, k_3)$.

In the single-source, slow-roll limit these relations simplify. The slow-roll approximation implies that ζ can be written using only the fluctuation in the single relevant field (without requiring the momentum fluctuation), so matrix multiplication and inversion reduce to ordinary multiplication and division. In this limit the distinction between Eqs. (2.28) and (2.21) disappears, and

$$\rho_\zeta(k) \approx \frac{B(k, k, k_3)}{P(k_3) P(k)} = \frac{12}{5} f_{\text{NL}}(k, k, k_3). \quad (2.29)$$

In a single-source model the right-hand side of (2.29) can be computed and shown to be independent of the soft mode k_3 . This observation follows from the formulae of Dias et al. [42], and was pointed out explicitly by Kenton & Mulryne [32]. The final result agrees with the separate-universe analyses (obtained by entirely different methods) presented in Refs. [17, 18, 23].

Eq. (2.29) and the property of independence from k_3 validate the statement made in Ref. [18], that biasing of the ζ two-point function in single-source models is controlled by the reduced bispectrum on equilateral configurations. In these models there is negligible difference between the reduced bispectrum in equilateral and squeezed configurations of the same scale.

More generally, Eq. (2.6) shows that the relevant configurations are squeezed rather than equilateral, although in local-like models the response cannot depend on the squeezing ratio

k_3/k_t . In an arbitrary model it need *not* happen that $f_{\text{NL}}(k_1, k_2, k_3)$ becomes independent of k_3 in the limit $k_3 \ll k_1, k_2$, which gives another demonstration that the response cannot equal the reduced bispectrum in general. Therefore to evaluate the amplitude of the asymmetry, and its compatibility with Planck constraints on the amplitude of three-point correlations in squeezed configurations, will generally require a model-dependent analysis. The same is true when checking the compatibility of a Grishchuk–Zel’dovich effect generated by an enhanced mode of wavelength longer than the scale of the last-scattering surface.

3 Why it is difficult to produce a suitable response

Up to this point, it remains an open question whether it is *possible* to manufacture a response function $\rho_\sigma(k)$ with suitable amplitude and scale-dependence without requiring an unacceptable exceptionality E or other undesirable features. In this section we discuss, in general terms, what properties appear required to yield a scale-dependent bispectrum, and explain why the simplest scenarios have difficulty in simultaneously producing an acceptable amplitude.

3.1 Scale-dependence of the bispectrum

Methods to compute the scale-dependence of a bispectrum generated during inflation have been discussed by several authors [43–45]. To describe the evolution of fluctuations on scales outside the horizon we use the separate universe picture to write an analogue of the gauge transformation (2.27) for the field fluctuations [46, 47]

$$\delta\phi^\alpha(\mathbf{k}) = \Gamma_a^\alpha \delta\phi^a(\mathbf{k}) + \frac{1}{2}\Gamma_{ab}^\alpha \int \frac{d^3q}{(2\pi)^3} \delta\phi^a(\mathbf{q})\delta\phi^b(\mathbf{k}-\mathbf{q}) + \dots, \quad (3.1)$$

which again does not invoke the slow-roll approximation. In this expression the fluctuation on the left-hand side is evaluated at the time of interest but the fluctuations on the right-hand side are evaluated at some earlier time. To prevent our formulae becoming cluttered by a proliferation of time labels we indicate evaluation at this earlier time by using the species labels $\{a, b, \dots\}$ instead of $\{\alpha, \beta, \dots\}$.

The mixed-index objects Γ_a^α and Γ_{ab}^α are derivatives of the background field configurations [47],

$$\Gamma_a^\alpha = \frac{\partial\phi^\alpha}{\partial\phi^a} \quad (3.2a)$$

$$\Gamma_{ab}^\alpha = \frac{\partial^2\phi^\alpha}{\partial\phi^a\partial\phi^b}. \quad (3.2b)$$

The background solution ϕ^α solves a system of differential equations $d\phi^\alpha/dN = u^\alpha$. Then, defining $u^\alpha{}_\beta = \partial_\beta u^\alpha$ and $u^\alpha{}_{\beta\gamma} = \partial_\gamma u^\alpha{}_\beta$ it is possible to write evolution equations for Γ_a^α and Γ_{ab}^α . As functions of the time N_+ defined by $\{\alpha, \beta, \dots\}$ they obey [47–51]

$$\frac{d}{dN_+}\Gamma_a^\alpha = u^\alpha{}_\beta\Gamma_a^\beta \quad (3.3a)$$

$$\frac{d}{dN_+}\Gamma_{ab}^\alpha = u^\alpha{}_\beta\Gamma_{ab}^\beta + u^\alpha{}_{\beta\gamma}\Gamma_a^\beta\Gamma_b^\gamma \quad (3.3b)$$

whereas as functions of the time N_- defined by $\{a, b, \dots\}$ they obey

$$\frac{d}{dN_-} \Gamma_a^\alpha = -\Gamma_b^\alpha u^b{}_a \quad (3.4a)$$

$$\frac{d}{dN_-} \Gamma_{ab}^\alpha = -\Gamma_c^\alpha u^c{}_{ab} - \Gamma_{cb}^\alpha u^c{}_a - \Gamma_{ac}^\alpha u^c{}_a. \quad (3.4b)$$

Variation with scale.—As we now explain, these equations for the time-dependence of Γ_a^α and Γ_{ab}^α enable us to determine the scale-dependence of correlation functions involving the $\delta\phi^\alpha$.⁵

Eq. (3.1) permits us to write the two- and three-point functions $\langle \delta\phi^\alpha \delta\phi^\beta \rangle$, $\langle \delta\phi^\alpha \delta\phi^\beta \delta\phi^\gamma \rangle$ evaluated at the late time N_+ in terms of $\langle \delta\phi^a \delta\phi^b \rangle$, $\langle \delta\phi^a \delta\phi^b \delta\phi^c \rangle$ evaluated at the early time N_- , giving

$$\langle \delta\phi^\alpha(\mathbf{k}_1) \delta\phi^\beta(\mathbf{k}_2) \rangle = \Gamma_a^\alpha \Gamma_b^\beta \langle \delta\phi^a(\mathbf{k}_1) \delta\phi^b(\mathbf{k}_2) \rangle \quad (3.5a)$$

and

$$\begin{aligned} \langle \delta\phi^\alpha(\mathbf{k}_1) \delta\phi^\beta(\mathbf{k}_2) \delta\phi^\gamma(\mathbf{k}_3) \rangle &= \Gamma_a^\alpha \Gamma_b^\beta \Gamma_c^\gamma \langle \delta\phi^a(\mathbf{k}_1) \delta\phi^b(\mathbf{k}_2) \delta\phi^c(\mathbf{k}_3) \rangle \\ &+ \Gamma_{mn}^\alpha \Gamma_b^\beta \Gamma_c^\gamma \int \frac{d^3q}{(2\pi)^3} \langle \delta\phi^m(\mathbf{q}) \delta\phi^b(\mathbf{k}_2) \rangle \langle \delta\phi^n(\mathbf{k}_1 - \mathbf{q}) \delta\phi^c(\mathbf{k}_3) \rangle \\ &+ \text{cyclic permutations.} \end{aligned} \quad (3.5b)$$

In what follows it is useful to make use of the scale k_t corresponding to the perimeter of the momentum n -gon characterizing each correlation function. For the two-point function $k_t = k_1 + k_2$ and for the three-point function $k_t = k_1 + k_2 + k_3$, with obvious generalizations to higher n -point functions.

To use Eqs. (3.5a) and (3.5b) we set the time N_- associated with $\{a, b, \dots\}$ to match the time at which the scale k_t exited the horizon in the sense $k_t/aH = 1$. Now suppose we vary the momentum configuration, first choosing to vary its scale k_t while leaving its shape (measured by the ratios k_i/k_t) fixed. This forces us to change the time N_- , giving two contributions to the variation of the two- and three-point functions: one arising from the change in evaluation time for Γ_a^α and Γ_{ab}^α , which can be computed using Eqs. (3.4a)–(3.4b); and the other from the change in the two- and three-point functions $\langle \delta\phi^a \delta\phi^b \rangle$, $\langle \delta\phi^a \delta\phi^b \delta\phi^c \rangle$.

For configurations which are close to equilateral in the sense that $k_i \sim k_t$ for each momentum k_i the change from $\langle \delta\phi^a \delta\phi^b \rangle$, $\langle \delta\phi^a \delta\phi^b \delta\phi^c \rangle$ can be obtained using known results for the n -point functions [53]. The dominant scaling is $1/k_t^3$ for the two-point function and $1/k_t^6$ for the three-point function, following directly from their engineering dimension but cancelling out in dimensionless combinations such as the reduced bispectrum. The remaining scaling comes mostly from the time dependence of the Hubble parameter H or the field momenta $d\phi^\alpha/dN$ evaluated at N_- . The contribution to (3.5b) from $\langle \delta\phi^a \delta\phi^b \delta\phi^c \rangle$ is known to be

⁵This method can be regarded as a refinement of the approach used in Refs. [43, 52]. These references effectively constructed linear approximations for Γ_a^α , Γ_{ab}^α as a function of N_- , invoking the slow-roll approximation to control the expansion. These approximations were valid over only a small range of k_t . Eqs. (3.4a) and (3.4b) replace these approximations. They can be used to determine the Γ -matrices at any scale and, as we have explained, they do not require the slow-roll approximation.

negligible unless the fields have nontrivial derivative interactions [54, 55], and we discard it in the following discussion.

With these assumptions we can estimate the scaling with k_t by temporarily invoking the slow-roll approximation to simplify our expressions. However, our conclusions will not depend on these simplifications. When we discuss a concrete model in §4 we will employ a numerical method which does not use the slow-roll approximation.

The slow-roll approximation makes the variation of H negligible when the scale k_t was leaving the horizon (even if slow-roll was subsequently violated), so the only contribution which needs to be kept is that generated by the Γ -matrices. Also, we do not need to retain fluctuations in the scalar field momenta and therefore we can restrict the indices on Γ_a^α , Γ_{ab}^α , $u^\alpha_{\beta\gamma}$ and $u^\alpha_{\beta\gamma}$ to the field fluctuations only. With this understanding we have

$$u^\alpha_{\beta\gamma} = -\frac{V_{\alpha\beta}}{3H^2} + \left(1 - \frac{\epsilon}{3}\right) \frac{\dot{\phi}^\alpha \dot{\phi}^\beta}{H^2 M_{\text{P}}^2} + \frac{2}{3} \frac{\dot{\phi}^{(\alpha} \ddot{\phi}^{\beta)}}{H^3 M_{\text{P}}^2} \quad (3.6a)$$

$$u^\alpha_{\beta\gamma} = -\frac{V_{\alpha\beta\gamma}}{3H^2} + \frac{\dot{\phi}^\alpha}{HM_{\text{P}}} u^\beta_{\gamma} + \frac{\dot{\phi}^\beta}{HM_{\text{P}}} u^\alpha_{\gamma} + \frac{\dot{\phi}^\gamma}{HM_{\text{P}}} u^\alpha_{\beta} - \frac{\dot{\phi}^\alpha \dot{\phi}^\beta \dot{\phi}^\gamma}{H^3 M_{\text{P}}^3} \quad (3.6b)$$

(in which the placement of indices should be regarded as immaterial, owing to our use of a trivial field-space metric). The quantity ϵ is the usual slow-roll parameter $\epsilon \equiv -\dot{H}/H^2$. If all fields are slowly rolling then $\dot{\phi}^\alpha/HM_{\text{P}} \ll 1$, making $u^\alpha_{\beta\gamma}$ principally sensitive to $V_{\alpha\beta}$ and $u^\alpha_{\beta\gamma}$ principally sensitive to $V_{\alpha\beta\gamma}$.

Obtaining power-law scaling.—In order to be concrete we restrict attention to the scenarios considered at the end of §2.2 in which only a single field σ dominates the ζ bispectrum. These allow the simplest possible statement. However we expect that our qualitative conclusions continue to apply in more general models.

First suppose that the second derivative of the potential in the direction σ is large compared to H , while the third derivative is small. Then $u^\sigma_{\sigma} \approx -\eta_\sigma$ while $u^\sigma_{\sigma\sigma} \approx 0$, where $\eta_\sigma = V_{\sigma\sigma}/3H^2$. It follows that

$$\frac{d}{d \ln k_t} \Gamma_\sigma^\sigma \approx \eta_\sigma \Gamma_\sigma^\sigma \quad (3.7a)$$

$$\frac{d}{d \ln k_t} \Gamma_{\sigma\sigma}^\sigma \approx 2\eta_\sigma \Gamma_{\sigma\sigma}^\sigma. \quad (3.7b)$$

If η_σ is approximately constant while some range of wavenumbers are leaving the horizon then we conclude that, over this range, $\Gamma_\sigma^\sigma \sim k_t^{\eta_\sigma}$ and $\Gamma_{\sigma\sigma}^\sigma \sim k_t^{2\eta_\sigma}$. This is the simplest mechanism by which one can generate significant scale dependence. The price we must pay to achieve any desired power-law scaling is a tuning of the mass, together with the necessity to keep this mass nearly constant over the desired range of scales. To obtain a red-tilted power law we require $\eta_\sigma < 0$, so during this epoch the field σ can be regarded as departing from a quadratic hilltop. The principal disadvantage of this mechanism is that it affects both Γ_σ^σ and $\Gamma_{\sigma\sigma}^\sigma$ and therefore both the two- and three-point functions will exhibit scale dependence. This makes model-building more complex because the σ contribution to the ζ two-point function must be kept sufficiently small that the resulting spectral index is acceptable.

Alternatively, in some scenarios it may be possible for $u^{\sigma}_{\sigma\sigma}$ to be large while u^{σ}_{σ} remains small. It is normally difficult to maintain this situation over many e -folds, because large contributions to $u^{\alpha}_{\beta\gamma}$ typically source large contributions to u^{α}_{β} . However, supposing it can be realized this scenario will generate scaling which satisfies

$$\frac{d}{d \ln k_t} \Gamma_{\sigma}^{\sigma} \approx 0 \quad (3.8a)$$

$$\frac{d}{d \ln k_t} \Gamma_{\sigma\sigma}^{\sigma} \approx \frac{\xi_{\sigma}}{M_{\text{P}}} \Gamma_{\sigma}^{\sigma}, \quad (3.8b)$$

where we have set $\xi_{\sigma} = M_{\text{P}} V_{\sigma\sigma\sigma} / 3H^2$. Compared to the large- η_{σ} case it is less simple to obtain a pure power-law, although by choosing ξ_{σ} appropriately it is still possible to generate scale dependence. The advantage of this scenario is that the two-point function does not acquire significant scale-dependence.

If we were to abandon the slow-roll approximation then similar conclusions would apply if, respectively, u^{σ}_{σ} or $u^{\sigma}_{\sigma\sigma}$ are larger than the other components of u^{α}_{β} and $u^{\alpha}_{\beta\gamma}$. In the case of multiple fields a similar discussion will apply to each field individually unless the u -matrices couple the scale-dependent species.

In either case, once a set of scaling behaviours have been generated by Eqs. (3.4a) and (3.4b), the subsequent evolution cannot generate new ones. Eqs. (3.3a) and (3.3b) *do* allow the evolution to vary the scaling observed in any particular correlation function by linearly mixing the available scalings with new amplitudes—but it is only these amplitudes which depend on the superhorizon epoch and not the scale-dependence itself. This conclusion is quite general and does not depend on the slow-roll approximation.

Variation with shape.—We now return to the alternative possibility of variations in the momentum configuration which leave k_t fixed but vary the side ratios k_i/k_t . We describe this as a change of shape. It is relevant only for the three- and higher n -point functions.

A variation in shape does not alter the evaluation time N_- for the coefficients Γ_a^{α} , Γ_{ab}^{α} in Eqs. (3.5b). It will change only the three-point function $\langle \delta\phi^a(\mathbf{k}_1) \delta\phi^b(\mathbf{k}_2) \delta\phi^c(\mathbf{k}_3) \rangle$ and a subset of the two-point functions appearing in the cyclically-permuted quadratic terms. The effect on the two-point functions is to change their evaluation time relative to the horizon-crossing time for their wavenumber, and Eq. (3.5a) shows that this can be expressed in terms of the 2-component Γ coefficient evaluated between suitable times. The effect on the three-point function is more difficult to extract but has recently been calculated by Kenton & Mulryne [32]. It can also be expressed purely in terms of the 2-component Γ coefficient.

The conclusion is that models which generate a significant scale-dependence through a large η_{σ} will almost always exhibit strong scaling as a function of the squeezing k_i/k_t . Together with the scaling with k_t , this will have implications for the degree to which we can interpret recent Planck measurements of $f_{\text{NL}}^{\text{local}}$ (or the amplitude of other templates) as measurements of the correlation amplitude in squeezed configurations. Conversely, models which generate scale-dependence through a large ξ_{σ} while keeping η_{σ} small will exhibit much smaller scaling as a function of squeezing because the 3-component Γ coefficient is not needed to describe scaling with shape at fixed k_t .

3.2 Models with constant η_σ are not viable

In this paper we focus on the simpler mechanism which generates significant k_t -scaling using a large second derivative, as in Eqs. (3.7a) and (3.7b). On the basis of what has been said above, we should expect this k_t -dependence to be accompanied by significant squeezing-dependence.

Scaling of f_{NL} .—How large an $|\eta_\sigma|$ is required? We focus on a simple model in which inflation is driven by a field ϕ which acquires a nearly scale-invariant spectrum while a second field σ acquires a scale-dependent fluctuation using Eqs. (3.5a) and (3.5b). The simplest arrangement occurs if the σ fluctuation contributes significantly to the ζ three-point function but not its two-point function, because then the two-point function can be insulated from any large scale-dependent contribution carried by $\delta\sigma$.

First consider variations of scale k_t . In the circumstances described above, $\langle\zeta\zeta\zeta\rangle$ will scale like $\langle\delta\sigma\delta\sigma\delta\sigma\rangle \sim k_t^{-6}(\Gamma_\sigma^\sigma)^2(\Gamma_{\sigma\sigma}^\sigma)^2 \sim k_t^{-6+4\eta_\sigma}$ whereas the scaling of $\langle\zeta\zeta\rangle$ will be independent, like $k_t^{-3+(n_s-1)}$. Therefore to lowest order in slow-roll parameters

$$f_{\text{NL}} \sim k_t^{4\eta_\sigma-2(n_s-1)} \quad (\text{variation of } k_t, \text{ shape fixed}). \quad (3.9)$$

Our approximations make this prediction independent of shape, although in principle shape-dependent contributions are present through the correlation functions $\langle\delta\phi^a\delta\phi^b\rangle$, $\langle\delta\phi^a\delta\phi^b\delta\phi^c\rangle$ which we have neglected. This property was noticed in Ref. [43], and we briefly reconsider it in §4.3 below.

For measurements of the response function we are instead interested in squeezed isosceles configurations as described in §2.1. On squeezed configurations the results of Kenton & Mulryne together with Eq. (2.21) suggest $\langle\zeta\zeta\zeta\rangle \sim \langle\delta\sigma\delta\sigma\delta\sigma\rangle \sim k_3^{-3}(\Gamma_\sigma^\sigma)^2 \sim k_3^{-3+2\eta_\sigma}$ (where now the time N_- appearing in the Γ -matrix is to be interpreted as the horizon exit time for k_3 [32]) while the scaling of the ζ two-point function is still $\langle\zeta\zeta\rangle \sim k_3^{-3+(n_s-1)}$. Therefore

$$f_{\text{NL}} \sim \left(\frac{k_3}{k_t}\right)^{2\eta_\sigma-(n_s-1)} \quad (\text{variation of shape, } k_t \text{ fixed}). \quad (3.10)$$

Putting Eqs. (3.9) and (3.10) together enables us to estimate the scaling on a sequence of isosceles triangles which fix k_3 but neither k_t nor k_3/k_t ,

$$f_{\text{NL}}(k, k, k_3) \sim k^{4\eta_\sigma-2(n_s-1)} \left(\frac{k_3}{k}\right)^{2\eta_\sigma-(n_s-1)} \sim k^{2\eta_\sigma-(n_s-1)} \quad \text{if } k \gg k_3. \quad (3.11)$$

Eq. (2.24) shows that, under our assumptions, the asymmetry amplitude $A(k)$ will scale like $f_{\text{NL}}(k, k, k_3)$ at fixed k_3 . Therefore we can estimate the η_σ required to generate a fixed power law. The scalar spectral index n_s is known from observation to be of order $n_s \approx 0.97$ [56], and to obtain a power-law for $f_{\text{NL}}(k, k, k_3)$ which is a little less steep than $k^{-0.5}$ we will take $\eta_\sigma \approx -0.2$.

Ridge models.—This choice for η_σ is attractive, because its relative largeness provides a means to synthesize three-point correlations with sufficient amplitude to modulate the ζ two-point function. Elliston et al. showed that an inflationary trajectory initially parallel to but slightly displaced from a quadratic ‘ridge’ in the inflationary potential generates an

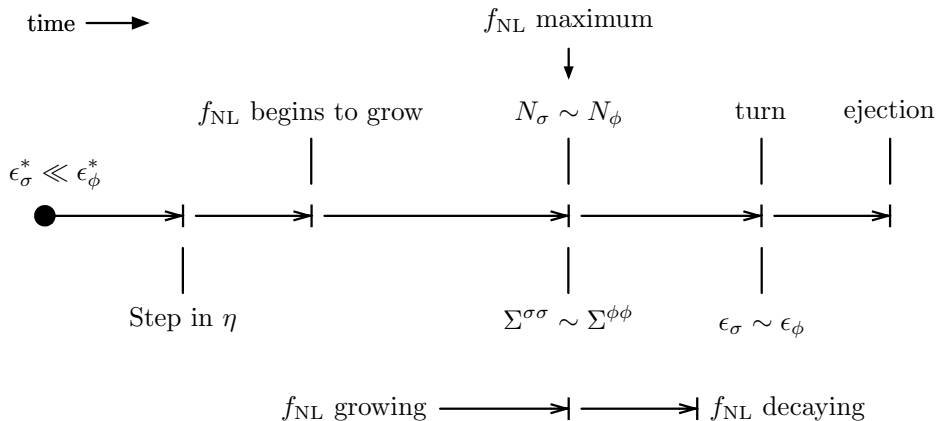


Figure 2: Evolution of f_{NL} generated by dispersion from a ridge with constant η .

enhanced bispectrum when the trajectory turns, eventually becoming nearly perpendicular to its original direction of travel [57]. The amplitude of the bispectrum at the point of maximum enhancement is proportional to the η parameter characterizing the ridge. Moreover, the ridge will naturally generate a *negative* η and therefore a red-tilted power-law.

The evolution of the bispectrum amplitude during this process is depicted in Fig. 2. The direction parallel to the ridge is the inflaton direction ϕ and the perpendicular direction is σ . We measure the relative kinetic energies by writing $\delta = \dot{\sigma}/\dot{\phi} = \sqrt{\epsilon_\sigma/\epsilon_\phi}$, where the ϵ -parameters are defined by

$$\epsilon_\phi \equiv \frac{M_{\text{P}}^2}{2} \left(\frac{V_\phi}{V} \right)^2 \quad (3.12a)$$

$$\epsilon_\sigma \equiv \frac{M_{\text{P}}^2}{2} \left(\frac{V_\sigma}{V} \right)^2. \quad (3.12b)$$

Labelling evaluation at the initial time by ‘*’, the initial conditions are chosen to set up a significant kinetic-energy imbalance $\epsilon_\sigma^* \ll \epsilon_\phi^*$, and therefore $\delta_* \ll 1$. This kinetic energy imbalance implies that ϕ dominates ζ . For example, one can verify that the gauge transformation coefficients N_α defined in (2.27) satisfy $N_\sigma^* \sim \sqrt{\delta_*} N_\phi^*$, and therefore $N_\sigma^* \ll N_\phi^*$.

Eventually the trajectory will begin to depart from the ridge and the kinetic energy imbalance will begin to equalize. During this process the contribution of $\delta\sigma$ to ζ at a fixed scale becomes both more significant and more nonlinear, while the $\delta\phi$ contribution is largely unchanged. This causes the ζ bispectrum to grow. The peak amplitude is inversely proportional to the original imbalance, of order $|\eta_\sigma/\delta_*|$. It is typically reached at the point where both $\delta\phi$ and $\delta\sigma$ make comparable contributions to the ζ two-point function. This occurs while $\delta < 1$, and therefore before the point of equal kinetic energy which we describe as ‘the turn’ in Fig. 2. By the time the trajectory turns the bispectrum amplitude is already decaying. After the turn the evolution depends on the precise form of the potential at large distances from the ridge, labelled ‘ejection’ in Fig. 2.

This process is a close analogue of the curvaton mechanism, which relies on equalizing a large initial *energy density* imbalance between ϕ and σ rather than a kinetic energy imbalance. The equalization phase typically happens after inflation because the σ potential need not be of inflationary type. Here also the final amplitude is typically inversely proportional to the initial imbalance, but in a curvaton model the $\delta\sigma$ fluctuation would normally come to dominate both the ζ two- and three-point functions. If the $\delta\sigma$ spectrum is too far from scale invariance this will lead to an unacceptable spectral index. Therefore, in either case, we have to contrive an exit which occurs when $\delta\sigma$ contributes to the ζ three-point function but does not dominate the two-point function. We focus on the ridge case on the assumption that it will be less easy to realize a working curvaton scenario.

Evolution of σ .—We now consider the evolution of σ while the bispectrum is beginning to grow. Recall that, in this section, we are still temporarily imposing the slow-roll approximation in order to simplify our formulae.

While the evolution of σ is described by a constant η -parameter, it will grow like

$$\sigma \approx \sigma_* e^{-\eta_\sigma N}. \quad (3.13)$$

We have assumed that the ridge lies at $\sigma = 0$ and that the initial conditions displaced σ to slightly positive values, so that σ will roll in this direction at later times. During this era the reduced bispectrum amplitude on the equilateral configuration which left the horizon at the initial time will grow approximately like [57, 58]

$$\frac{6}{5} f_{\text{NL}} \approx -\frac{\eta_\sigma}{\delta_*} R^3, \quad (3.14)$$

where R is the quantity defined in (2.25). It measures the relative contribution of the σ and ϕ fluctuations to the ζ two-point function. Eq. (3.14) can be interpreted as a constraint on the initial kinetic energy imbalance given a value for η_σ and a desired amplitude f_{NL} . But we cannot make the initial imbalance too extreme without positioning σ_* very close to the top of the ridge where its evolution is dominated by quantum diffusion rather than classical evolution. To stay outside the diffusion regime we require the classical motion in a single e-fold $\sim d\sigma/dN$ to dominate the quantum motion $\delta\sigma/\delta N \sim H/2\pi$. That gives $2M_{\text{P}}^2 \epsilon_\sigma^* \gtrsim (H_*/2\pi)^2$, or equivalently

$$\delta_*^2 = \frac{\epsilon_\sigma^*}{\epsilon_\phi^*} \gtrsim \frac{1}{2} \frac{H_*^2}{4\pi^2} \frac{1}{M_{\text{P}}^2 \epsilon_\phi^*} \approx \frac{\mathcal{P}}{2}. \quad (3.15)$$

In the final step we have used the approximation that the ϕ contribution to \mathcal{P} barely evolves while f_{NL} is growing. If necessary it is possible to make a more precise estimate, but the outcome is hardly altered. Therefore we conclude

$$\frac{6}{5} |f_{\text{NL}}| \ll \left| \frac{\sqrt{2}\eta_\sigma R^3}{\mathcal{P}^{1/2}} \right|. \quad (3.16)$$

We have written ‘ \ll ’ because to obtain an acceptable outcome it is necessary to take δ_* substantially larger than the lower limit $\sim \mathcal{P}^{1/2}$. This is because initial conditions which

are too close to the diffusion regime would lead to very large fluctuations on the largest observable scales. For $\eta_\sigma \sim 0.2$ and $R \sim 0.1$ Eq. (3.16) gives $|f_{\text{NL}}| \ll 5$, which implies that it will hardly be possible to achieve even $f_{\text{NL}} \sim 1$. If we require more stringent bounds on R then the situation is worse. The conclusion is that, if one seeks to synthesize a scale-dependent bispectrum during inflation using the ridge mechanism with constant η_σ , and insists that the σ two-point function remains subdominant in the ζ power spectrum, it will not be possible to start with δ_* sufficiently small to yield a large amplitude.

Alternative strategies.—This argument invokes relatively strong assumptions, such as a constant η_σ , and does not preclude the possibility that a working model can be found. Instead it should be regarded as a guide, suggesting that in a large- η_σ model *either* the σ contribution to the ζ two-point function cannot remain always subdominant *or* that the σ potential cannot be featureless. These suggestions highlight generic difficulties. First, the σ power spectrum will typically be strongly scale-dependent if η_σ is large enough to generate a scale-dependent bispectrum. Second, features in the σ potential will typically imply fine-tuning. In the remainder of §§3–4 we develop these difficulties in more detail. A related discussion has been given by Kanno et al. [22].

Specifically, attempts to construct a working model using a large value for η_σ normally encounter at least one of the following difficulties.

- If the σ potential is featureless and we terminate the inflationary phase before the scale-dependent $\delta\sigma$ two-point function contributes significantly to the ζ two-point function, in order to protect near scale-invariance, then we will not normally generate a sufficiently large bispectrum to produce the desired response.
- To allow the development of large amplitudes one can add features to the σ power spectrum. This will be the approach taken in §4 below. Models of this kind may be successful but typically require several tunings, including at least the details of the feature and the exit from inflation. The exit point must be fine-tuned so that the bispectrum amplitude is sufficiently large but R is still sufficiently small. Even if this can be done there may be issues with the amplitude of the trispectrum, to be described below, and the required initial conditions may still be uncomfortably close to the diffusion regime.
- Alternatively one can abandon the idea of keeping R always small, allowing (3.14) to yield a larger amplitude. This could be done by retaining the ridge mechanism but allowing f_{NL} to pass through the point of maximum amplitude, or by extending the model to include a subsequent curvaton era.

In the ridge case one must still tune the exit from inflation to occur at the correct amplitude, while also finding some mechanism to suppress R later. We expect this requires further tuning. In the curvaton case there is no need to tune the inflationary exit, but it seems even more difficult to suppress R once the curvaton has come to dominate the energy density.

A different strategy would be to abandon the large- η_σ method for obtaining scaling in the bispectrum, instead attempting to find a suitable ξ_σ while keeping η_σ very small. This makes it easier to keep the $\delta\sigma$ two-point function nearly scale-invariant, but one would have to contrive a suitable ξ_σ evolution which gave a power-law response *and* did not induce a large η_σ at any point during the evolution. Although we have not proved that this cannot be done, it is far from easy to do so.

Trispectrum constraints.—The simplest possibility is to keep $R < 1$ and tune the time of exit from inflation. However, as we now explain, this typically leads to a large amplitude in the τ_{NL} mode of the trispectrum.

In a model which develops a large bispectrum amplitude through superhorizon evolution there will typically be two contributions to the trispectrum: a g_{NL} -mode, which is small for models without large cubic self-interactions [59–62], and a τ_{NL} -mode (discussed in §2.2 above) which obeys the Suyama–Yamaguchi relation [63],

$$\tau_{\text{NL}} \geq \left(\frac{6}{5} f_{\text{NL}}\right)^2. \quad (3.17)$$

In a model with significant scaling the precise momentum dependence associated with these trispectrum contributions will be modified in comparison with the standard templates, but we expect their basic character to be preserved.

Equality occurs for single-source models. In multiple-source models the amplitude of the τ_{NL} shape may be enhanced. Under the same assumptions which led to the estimates (2.24) and (2.26) for f_{NL} and τ_{NL} in terms of $A(k)$ we can infer a relationship between f_{NL} evaluated on squeezed configurations and τ_{NL} evaluated on collapsed configurations of a similar scale,

$$\tau_{\text{NL}} \simeq \frac{1}{R^2} \left(\frac{6}{5} f_{\text{NL}}\right)^2. \quad (3.18)$$

Therefore the trispectrum amplitude is enhanced by a factor $1/R^2$. This can be understood if we interpret (2.24) to mean that the price paid to obtain a strongly scale-dependent bispectrum is suppression of f_{NL} below its natural value by the factor $R^{1/2} < 1$. However, τ_{NL} is not suppressed in the same way and is therefore substantially larger than the lower bound provided by the Suyama–Yamaguchi relation [64].

If we choose $R \ll 1$ to protect the near scale-invariance of the power spectrum then the enhancement can be considerable. The same estimate (3.18) can be obtained using slow-roll results for the amplitude of τ_{NL} on tetrahedral configurations.⁶ These additionally enable us to estimate the scale-dependence of τ_{NL} on such configurations,

$$n_{\tau_{\text{NL}}} \simeq 6(\eta_\sigma - \eta_\phi) \simeq \frac{3}{2} n_{f_{\text{NL}}}, \quad (3.19a)$$

where $n_{f_{\text{NL}}}$, $n_{\tau_{\text{NL}}}$ refer to the spectral index of the reduced bispectrum and the amplitude of the τ_{NL} -shape with variations of k_i .

⁶A ‘tetrahedral’ configuration is a special case of an equilateral configuration for which $k_i = |\mathbf{k}_i - \mathbf{k}_j| = k$ for some k and all $i \neq j$. This is the appropriate trispectrum generalization of an equilateral bispectrum configuration, in the sense that it involves only a single scale. Therefore large effects due to a hierarchy of scales cannot enter [65].

Models which achieve strong scaling through a large $\xi_\sigma = M_{\text{p}}^3 V_{\sigma\sigma\sigma}/V$ are instead likely to generate large g_{NL} . Ref. [62] studied a concrete example based on a self-interacting curvaton, showing that it brought the model into conflict with observational constraints both for g_{NL} and the quadrupolar asymmetry of the power spectrum. To estimate how large a ξ_σ might be required to generate a relevant scale-dependence we use the slow-roll results given in Byrnes et al. [43], which yield roughly $|\xi_\sigma| \gtrsim 100$. (Notice that this does not imply a breakdown of the slow-roll conditions for σ , although as explained above we would expect a large ξ_σ to generate a large η_σ which in turn might conflict with scale-invariance of the ζ power spectrum.) In §5 we comment on further challenges arising in large ξ_σ models.

Non-inflationary scenarios.—Finally one might consider whether scenarios exist which do not rely on a large-amplitude bispectrum to couple long- and short-scale perturbations. Examples include: a domain wall [66]; ‘thawing’ cosmic strings [67]; parity violating fluctuations in the initial state [68]; a modulation of the long wavelength mode, coupling to isocurvature perturbations which persist until at least decoupling [69]; or spatially-modulated dissipation of the inflaton during inflation [70].

Some of these models are attractive because they naturally explain the scale dependence of the asymmetry, whereas in the inflationary case it requires a model-building choice to realize a large η_σ or ξ_σ . However, this does not imply that the resulting scale-dependence will match observation. For example, the simplest realization of the domain wall model predicts $A(k) \propto 1/k$, which is too steep to match the data. But more importantly (to the best of our knowledge) none of these models have been shown explicitly to match all the observational constraints, including those of the bi- and trispectrum, the low- ℓ multipoles of the CMB, *and* a quadrupolar modulation of the power spectrum. In many cases the shape and amplitude of the non-Gaussianity has not been computed or compared to observation—although this criticism applies equally to the inflationary case, which we will take up in §4 below.

4 A working but contrived model

In this section we illustrate the difficulties highlighted in §3 by exhibiting a concrete model in which the response has suitable amplitude and scale-dependence. The model is compatible with current constraints on the two- and three-point functions of ζ , and as we explain below it may also be compatible with constraints on the four-point function. However, its construction involves a number of arbitrary choices. We describe the model in §4.1 before summarizing the fine-tunings which are required.

A major advantage of working with a concrete model is that we can compute its bispectrum in detail, allowing the resulting shape and amplitude to be compared with precision CMB constraints. In §4.2 we discuss compatibility with constraints on the low CMB multipoles C_ℓ , and in §4.3 we estimate the bispectrum amplitudes $\hat{f}_{\text{NL}}^{\text{local}}$, $\hat{f}_{\text{NL}}^{\text{equi}}$ and $\hat{f}_{\text{NL}}^{\text{ortho}}$ which would be measured by a Planck-like experiment in this model.

4.1 A step in η

To avoid difficulties with the ζ spectral index we focus on models in which a large bispectrum is generated during inflation, with no subsequent curvaton era, as explained in §3.2. The discussion there showed that we can not expect to obtain a large amplitude if η_σ is constant, requiring the introduction of some feature which allows η_σ to evolve. The next simplest possibility is to allow a step which interpolates between two different constants. To realize this we adopt the potential

$$V = V_0 \left(1 + \frac{\eta_\phi}{2} \frac{\phi^2}{M_{\text{P}}^2} \right) \left(1 + \frac{1}{2} \frac{\sigma^2}{M_{\text{P}}^2} \left[\frac{\eta_2 - \eta_1}{2} \tanh \frac{\sigma - \sigma_c}{\sigma_{\text{step}}} + \frac{\eta_1 + \eta_2}{2} \right] - \frac{1}{2} \frac{\sigma_c^2}{M_{\text{P}}^2} \frac{\eta_2 - \eta_1}{2} \left[1 + \tanh \frac{\sigma - \sigma_c}{\sigma_{\text{step}}} \right] \right). \quad (4.1)$$

The step is centred at $\sigma = \sigma_c$ and has characteristic width σ_{step} in field units. By making σ_{step} small we can achieve a rapid transition. Prior to the transition the effective η_σ parameter is $\eta_\sigma \approx \eta_1$. After the transition we have $\eta_\sigma \approx \eta_2$. The scale V_0 should be chosen to match the normalization of the ζ power spectrum. In our numerical calculations we take $V_0 \approx 10^{-14} M_{\text{P}}^4$.

If we wish to work in a regime where $\delta\phi$ dominates the ζ two-point function then we should choose $\eta_\phi \approx -0.02$ to give an acceptable spectral index. We fix $\eta_1 = -0.25$ and $\eta_2 = -0.08$, and take initial conditions $\phi_* = 0.01 M_{\text{P}}$ and $\sigma_* = 8.94427 \times 10^{-8} M_{\text{P}}$ at time $N = 0$. The diffusion regime exists for roughly $|\sigma| < 4 \times 10^{-8} M_{\text{P}}$, making our initial conditions safe by ~ 3 e-folds. This is rather closer than one would like in a realistic scenario, but our interest is purely illustrative. In any case it exemplifies the difficulty of staying outside the diffusion regime—even when features are introduced in the σ potential. Finally, taking the transition to occur at $\sigma_c = 3.445 \times 10^{-6} M_{\text{P}}$ with width $\sigma_{\text{step}} = 10^{-10} M_{\text{P}}$ gives an acceptable phenomenology in which the step occurs at a time $N_{\text{step}} \sim 15$; see Fig. 3, which clearly demonstrates the need to tune the time at which the inflationary phase exits.

In this paper we do not discuss an exit mechanism, instead assuming that a choice can be found which brings inflation to an end around $N = 50$ while preserving the statistical properties of the field fluctuations. Accurately accounting for effects associated with the end of inflation and subsequent reheating [71–73] would be a challenge for any attempt to construct a realistic model describing the asymmetry.

Computing the response.—The large η_σ prior to the step and the rapidity of the transition imply that we should not trust analytic approximations for the correlation functions. Instead we calculate estimates for the response functions ρ_μ by combining Eq. (2.11) with numerical computations for the two- and three-point functions of the model. Our numerical method makes no use of the slow-roll approximation, instead treating the dynamics exactly. It also accounts for all quantum effects, including deformation of the wavefunctions due to mass terms, mixing with the metric, interference effects due to mode-coupling, and off-diagonal correlations present around the time of horizon exit. However, in practice, we find that these quantum effects are not very significant for the model (4.1). The details of these simulations will be described in forthcoming publications [74].




Figure 3: Time evolution of the reduced bispectrum $f_{\text{NL}}(k_1, k_2, k_3)$ in squeezed isosceles and equilateral configurations at the same scale k_t , chosen to be the value for which the average side $k_t/3$ roughly corresponds to $\ell = 1$. The spike at $N_{\text{step}} \sim 15$ is the effect of the step in σ . The growth from $N \gtrsim 35$ represents the growth of the bispectrum amplitude predicted by (3.14) as the initial kinetic energy imbalance is equalized. (Compare Fig. 2.) We have truncated the evolution at $N = 50$.

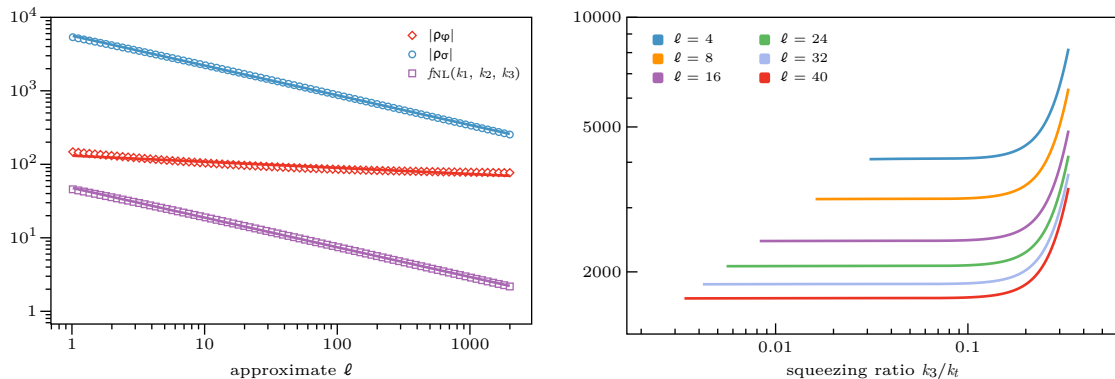


Figure 4: Left panel: Absolute value of response functions ρ_ϕ and ρ_σ for the step model (4.1). The responses are estimated using Eq. (2.11) on squeezed isosceles configurations where the long mode exits at time $N = 0$ and the approximate multipole $\ell = 14,000k$ Mpc corresponding to the short mode is plotted on the horizontal axis. We restrict to configurations for which the response can be estimated with $k_3/k_t < 0.1$ in order to ensure that the response has become adequately independent of the long mode; see right panel. Also plotted is the reduced bispectrum amplitude $f_{\text{NL}}(k, k, k_3)$ measured on the same isosceles configurations used to estimate the response functions. The data points represent values extracted from our numerical method and the solid lines are power-law fits.

Right panel: Variation in response functions with squeezing fraction k_3/k_t for different scales, measured by the approximate multipole ℓ . The response functions become independent of k_3 for squeezings $k_3/k_t < 0.1$.

Numerical results.—In the left panel of Fig. 4 we plot the response functions $\rho_\phi(k)$ and $\rho_\sigma(k)$ as a function of the approximate multipole $\ell = 14000k$ Mpc corresponding to the wavenumber

k . The ϕ response function is reasonably close to scale invariance, but substantially smaller than the σ response function which exhibits strong scale-dependence. It can be fit by the power law

$$\rho_\sigma(k) \approx 5600 \left(\frac{k}{k_{\ell=1}} \right)^{-0.405} \quad (4.2)$$

where the scale $k_{\ell=1} = 1/14,000 \text{ Mpc}^{-1}$ corresponds to the multipole $\ell \approx 1$. Eq. (4.2) is roughly acceptable as a description of the scale dependence of the modulation amplitude $A(k)$.⁷ The power-law for $f_{\text{NL}}(k, k, k_3)$ evaluated on the same configurations (represented by the purple line in the left panel of Fig. 4) is very close, with spectral index -0.404 . Both of these are close to the estimate $\sim k^{2\eta_\sigma - 2\eta_\phi} \sim k^{-0.36}$ obtained in Eq. (3.11), where we have approximated $n_s - 1 \approx 2\eta_\phi$, and reproduce the conclusion of Eq. (2.24) that $f_{\text{NL}}(k, k, k_3)$ and $\rho_\sigma(k)$ should scale similarly even if their amplitudes are different. A very small amount of running is visible in both ρ_σ and $f_{\text{NL}}(k, k, k_3)$, and fitting only to the region $1 \leq \ell \leq 60$ changes the spectral indices to -0.393 and -0.392 respectively.

We have checked that one can obtain steeper power laws if desired, by changing η_1 and modifying the initial conditions appropriately. However because we are not seriously advocating this model as an explanation of the anomaly, but using it only to illustrate general properties, we are content with the scaling (4.2) which allows simple numerical values for η_1 and η_2 .

These response functions are extracted from squeezed configurations, and we have verified that the level of squeezing is sufficient for each ρ_μ to become independent of the long mode. This is demonstrated in the right panel of Fig. 4 which shows ρ_σ as a function of the squeezing parameter k_3/k_t . In the left panel of Fig. 4 we have included response functions only for values of k which can be measured from our numerics on configurations with squeezing $k_3/k_t < 0.1$.

The power spectra for this model are shown in Fig. 5. On CMB scales the ζ power spectrum is dominated by ϕ , but at small ℓ it begins to receive contributions from σ which generate running. For the scales contributing to $\ell < 2000$ it can be fit by an approximate power law of the form $k^{-0.06}$, which would correspond to $n_s \sim 0.94$. Fitting instead to the region $500 < \ell < 1500$ gives $n_s \approx 0.96$. The σ power spectrum satisfies approximately

$$\mathcal{P}_\sigma(k) \approx 1.7 \times 10^{-12} \left(\frac{k}{k_{\ell=1}} \right)^{-0.464}. \quad (4.3)$$

Assuming the larger-scale modes were generated during the same period with large $|\eta_\sigma|$ we can estimate the required exceptionality E as a function of the quantity α , introduced in Eq. (2.15) to parametrize the physical scale of the modulating mode \mathbf{k}_L ,

$$E(\alpha) \approx \frac{A(k_{\ell=1})}{\pi\alpha} \frac{1}{\rho_\sigma(k_2)\mathcal{P}_\sigma^{1/2}(\alpha k_{\ell=1})} \approx 8.96\alpha^{-0.768} = 8.96 \left(\frac{k_L}{k_{\ell=1}} \right)^{-0.768}. \quad (4.4)$$

In the final numerical estimate we have taken $A(k_{\ell=1}) = 0.2$, as suggested by the scale-dependent analysis by Aiola et al. [12]. It follows that we can achieve exceptionalities in the

⁷We have verified that the response from the momentum perturbation $d\delta\sigma/dN$ is substantially smaller and can be neglected.




Figure 5: Left panel: power spectra in the step model (4.1), as a function of the approximate corresponding multipole $\ell \approx 14,000k$ Mpc. Right panel: Zoomed-in region showing only the ζ power spectrum. The small curvature means that a global power-law, represented by the solid red line, is not a perfect fit; instead, some running is required. Fitting to only the region $500 < \ell < 1500$ (represented by the green line) gives a shallower power law.

desired range provided k_L is not too much smaller than $k_{\ell=1}$. As explained in §2.2, this is anyway required to control the term $C(k)$.

4.2 The Grischuk–Zel’dovich effect

Having verified that a modulation with the correct amplitude and scale dependence can be synthesized in this model, there are two observational checks on its viability. The first comes from the requirement that the amplitude of the modulating mode is not so large that it would generate unacceptable contributions to the low CMB multipoles. This is the Grischuk–Zel’dovich effect. The second constraint is that, for exceptionalities which pass the Grischuk–Zel’dovich test, the required bispectrum amplitude is compatible with Planck measurements which are principally sensitive to squeezed configurations associated with CMB scales. In this section we pursue the Grischuk–Zel’dovich constraint, leaving the bispectrum amplitude to §4.3.

To compute the Grischuk–Zel’dovich effect, we note that the modulating mode will make a direct contribution to ζ which can be estimated from (2.27) after inverting the Fourier transform,

$$\zeta(\mathbf{x}) \supseteq N_\sigma \delta\sigma(\mathbf{x}) + \frac{1}{2} N_{\sigma\sigma} \delta\sigma(\mathbf{x})^2 + \dots \quad (4.5)$$

As in Eq. (2.27) the fluctuations on both sides of this expression are to be evaluated at the same time, in contrast to Eq. (3.1).

The contribution to CMB multipoles can be obtained by combining (4.5) with the formula $\Theta(\hat{\mathbf{n}}) \approx \zeta(x_{\text{ls}} \hat{\mathbf{n}})/5$ for the temperature anisotropy, after using (2.14) and expressing the result as a spherical harmonic transform. To do so we use Rayleigh’s formula for plane waves,

$$\exp(i\mathbf{k} \cdot \mathbf{x}) = 4\pi \sum_{\ell m} i^\ell j_\ell(kx) Y_{\ell m}^*(\hat{\mathbf{k}}) Y_{\ell m}(\hat{\mathbf{x}}), \quad (4.6)$$

where $j_\ell(x)$ is the spherical Bessel function of order ℓ and $Y_{\ell m}(\hat{\mathbf{n}})$ are spherical harmonics oriented with respect to the polar axis $\hat{\mathbf{n}}$. Using Rayleigh's formula to express the trigonometric functions as spherical harmonic transforms, we conclude⁸

$$a_{10} \supseteq -\frac{2\sqrt{3}\pi}{5} E(\alpha) \mathcal{P}_\sigma^{1/2}(k_L) \left(j_1(2\pi\alpha) N_\sigma \sin \vartheta + \frac{1}{4} E(\alpha) \mathcal{P}_\sigma^{1/2}(k_L) j_1(4\pi\alpha) N_{\sigma\sigma} \sin 2\vartheta \right) \quad (4.7a)$$

$$a_{20} \supseteq -\frac{2\sqrt{5}\pi}{5} E(\alpha) \mathcal{P}_\sigma^{1/2}(k_L) \left(j_2(2\pi\alpha) N_\sigma \cos \vartheta + \frac{1}{4} E(\alpha) \mathcal{P}_\sigma^{1/2}(k_L) j_2(4\pi\alpha) N_{\sigma\sigma} \cos 2\vartheta \right) \quad (4.7b)$$

$$a_{30} \supseteq \frac{2\sqrt{7}\pi}{5} E(\alpha) \mathcal{P}_\sigma^{1/2}(k_L) \left(j_3(2\pi\alpha) N_\sigma \sin \vartheta + \frac{1}{4} E(\alpha) \mathcal{P}_\sigma^{1/2}(k_L) j_3(4\pi\alpha) N_{\sigma\sigma} \sin 2\vartheta \right) \quad (4.7c)$$

One can suppress some contributions to Eqs. (4.7a)–(4.7c) by tuning the phase ϑ , but this is unattractive in a model which already requires significant tunings. As in §2.2 we assume that the Earth lies at a typical point where $\sin \vartheta \sim \cos \vartheta \sim \sin 2\vartheta \sim \cos 2\vartheta \sim \mathcal{O}(1)$. In the step model (4.1), N_σ settles down to a scale-independent value $N_\sigma \approx -14.2568$ long after the transition, and likewise for $N_{\sigma\sigma} \approx -8.07831$. The conclusion is that the Grischuk–Zel'dovich contributions to a_{10} , a_{20} and a_{30} can be written

$$a_{10} \supseteq 2.04 \times 10^{-4} \frac{j_1(2\pi\alpha)}{\alpha} + 3.38 \times 10^{-10} \frac{j_1(4\pi\alpha)}{\alpha^2} \quad (4.8a)$$

$$a_{20} \supseteq 2.64 \times 10^{-4} \frac{j_2(2\pi\alpha)}{\alpha} + 4.37 \times 10^{-10} \frac{j_2(4\pi\alpha)}{\alpha^2} \quad (4.8b)$$

$$a_{30} \supseteq 3.13 \times 10^{-4} \frac{j_3(2\pi\alpha)}{\alpha} + 5.17 \times 10^{-10} \frac{j_3(4\pi\alpha)}{\alpha^2}. \quad (4.8c)$$

Numerical values for a_{10} , a_{20} and a_{30} , together with the corresponding exceptionality $E(\alpha)$ and monopolar amplitude modulation $|C(k_{\ell=1})|$, are shown in Table 2 for $\alpha = 0.1, 0.01$ and 0.001 . We also display observational limits obtained from measurements of the low C_ℓ , where

$$C_\ell = \frac{1}{2\ell + 1} \sum_m |a_{\ell m}|^2. \quad (4.9)$$

We should therefore expect $a_{20} < C_2^{1/2}$ and $a_{30} < C_3^{1/2}$. The limits used here match those of Erickcek et al. [16]. Kanno et al. and Lyth used instead $C_2^{1/2} < 6.5 \times 10^{-6}$ [18, 22]

⁸These estimates differ in detail compared with results previously reported in the literature [16–18, 22, 23], which replace the spherical Bessel functions with other numerical factors. The difference arises because a_{10} and a_{30} receive contributions from all odd powers of $\mathbf{k}_L \cdot \mathbf{x}$ whereas a_{20} receives contributions from all even powers. Therefore any estimate requires an assumption about these higher order terms. The estimates made in Refs. [16–18, 22, 23] assumed a power series expansion in $\mathbf{k}_L \cdot \mathbf{x}$ which terminated at the quadratic or cubic level, so all higher-order terms were absent. In Eqs. (4.7a)–(4.7c) we have assumed the higher-order terms come from expansion of (2.14). The difference between these estimates is already negligible for $\alpha \lesssim 0.1$.

Our formula for the response was accurate only to $\mathcal{O}(k_L^2)$, but this does not prevent us from using terms such as $(\mathbf{k}_L \cdot \mathbf{x})^2$ or higher which combine \mathbf{k}_L with the same number of powers of \mathbf{x} . The $\mathcal{O}(k_L^2)$ corrections would not involve \mathbf{x} and are therefore subleading compared to the angular terms retained in these estimates.

Finally, the ansatz (2.14) takes the long-wavelength perturbation to be a pure cosine at the time of interest, whereas in Ref. [23] it was taken to be a pure cosine at the time of horizon exit. Nonlinear evolution between horizon exit and the time of interest would then generate \cos^2 and higher contributions sourced by Γ_{ab}^α and higher derivatives. These are not important for determining the asymmetry but would contribute to Eqs. (4.7a)–(4.7c). In both cases the pure cosine is just an ansatz and there will be further corrections which are being neglected, so we do not view this difference as material.

	scale of modulating mode			observational limit
	$\alpha = 0.1$	$\alpha = 0.01$	$\alpha = 0.001$	
a_{10}	4.12×10^{-4}	4.28×10^{-4}	4.30×10^{-4}	$C_1^{1/2} \lesssim 10^{-3}$
a_{20}	6.76×10^{-5}	6.95×10^{-6}	6.99×10^{-7}	$C_2^{1/2} \lesssim 5.9 \times 10^{-6}$
a_{30}	7.22×10^{-6}	7.39×10^{-8}	7.48×10^{-10}	$C_3^{1/2} \lesssim 9.1 \times 10^{-6}$
E	50.5	308	1800	
$ C(k_{\ell=1}) $	0.32	3.2	32	

Table 2: Grischuk–Zel’dovich contributions to a_{10} , a_{20} and a_{30} for the step model (4.1) at different values of α , together with the exceptionality E and the monopolar amplitude modulation $|C(k_{\ell=1})|$. Limits on the observed quadrupole and octupole C_2 and C_3 have been given by Efstathiou [75], who reported $\Delta T_2^2 \lesssim 250(\mu K)^2$ and $\Delta T_3^2 \lesssim 1183(\mu K)^2$ with the real value expected to lie near these limits. These correspond to $C_2^{1/2} \lesssim 5.9 \times 10^{-6}$ and $C_3^{1/2} \lesssim 9.1 \times 10^{-6}$. The dipole is comparatively unconstrained owing to uncertainties from the Earth’s motion relative to the CMB rest frame, but the values $\sim 10^{-4}$ given here are acceptable because they are smaller than the observed value $\sim 10^{-3}$.

and Kobayashi et al. used $C_2^{1/2} < 1.0 \times 10^{-5}$ [23]. In addition, Refs. [17, 18] used the method described above to estimate the contribution proportional to $N_{\sigma\sigma}$ (described as the ‘Erickcek–Kamionkowski–Carroll’ effect), but a different method to estimate the contribution proportional to N_σ (described as the ‘Grischuk–Zel’dovich’ effect). These different numerical choices led to $O(1)$ discrepancies between the constraints on f_{NL} reported in Refs. [17, 18] and Ref. [23]. We follow Kobayashi et al. in treating the contributions proportional to N_σ and $N_{\sigma\sigma}$ in the same way.

From Table 2 we conclude that, in the specific model (4.1), a value of α just a little smaller than $\alpha = 0.01$ should be acceptable, requiring an enhancement factor a little larger than $E = 300$ and perhaps a 1% to 10% tuning of $C(k)$. The Grischuk–Zel’dovich contributions to a_{20} and a_{30} can be suppressed comfortably below the observed constraints if we go as far as $\alpha = 0.001$ at the expense of a much larger exceptionality $E \sim 2000$ and substantially more tuning in $C(k)$. These values for E —but *not* the tuning in $C(k)$, cf. Eq. (2.18)—could be reduced by generating a bispectrum of larger amplitude, provided it remains compatible with the observational constraints discussed in §4.3.

Comparison with earlier literature.—If desired, Eqs. (4.7a)–(4.7c) can be rewritten in terms of A and f_{NL} . Focusing on the part of f_{NL} generated by the superhorizon mode and neglecting scale- and shape-dependence of each quantity, Refs. [17, 18, 22, 23] found a relation of the form

$$|a_{20}| \geq 6.9 \times 10^{-6} \times \frac{60}{|f_{\text{NL}}|} \left(\frac{A}{0.07} \right)^2 |\beta(\alpha, k_L)| \quad (4.10)$$

where β is defined by

$$\beta(\alpha, k_L) \approx \cos 2\vartheta + \frac{N_\sigma}{N_{\sigma\sigma}} \frac{1}{E(\alpha) \mathcal{P}_\sigma^{1/2}(k_L)} \cos \vartheta \quad (4.11)$$

for $\alpha \ll 1$. We have verified that a similar relation holds with scale- and shape-information retained, although because it is no more informative than (4.10) we do not write it explicitly.

The quantity β was introduced by Kobayashi et al. [23] and measures the relative contribution of the N_σ and $N_{\sigma\sigma}$ contributions in (4.7b). In β these translate to the terms proportional to $\cos \vartheta$ and $\cos 2\vartheta$, respectively. For the purposes of numerical estimates Kobayashi et al. assumed $|\beta| = \mathcal{O}(1)$ which implies the $N_{\sigma\sigma}$ term is dominant. In practice the N_σ term is often more important because it is enhanced by $\mathcal{P}_\sigma^{-1/2}$. Without tuning ϑ this makes values in the range 10 to 10^3 reasonable.

Ignoring scale- and shape-dependence, Eq. (4.10) suggests that, without unexpected cancellation between Eq. (4.10) and other contributions to a_{20} , we should expect $f_{\text{NL}} \gtrsim 60$ even in the optimistic case $\beta \sim 1$. This led to a discussion in the literature regarding compatibility of the model with observation, since the amplitude of local-type contributions to the bispectrum is now constrained to be substantially less than 60. In the next section, we will explicitly show that accounting for the scale- and shape-dependence of the bispectrum allows us to make a_{20} sufficiently small and A sufficiently large without demanding an unacceptable amplitude, even when we allow β to be substantially larger than unity.

4.3 The shape and amplitude of the bispectrum

Finally we must check the amplitude of three-point correlations. We have already observed that the reduced bispectrum $f_{\text{NL}}(k_1, k_2, k_3)$ will run as a function of scale and squeezing, and therefore will not match the ‘local’ template used to obtain the Planck2015 constraint $f_{\text{NL}}^{\text{local}} = 0.8 \pm 5.0$ [76]. In this section we study the shape of the bispectrum generated by (4.1) in more detail.

Variation with scale.—In Fig. 6 we show the dependence of the bispectrum on k_t for fixed shape, and its dependence on squeezing k_1/k_t at fixed scale. The variation with k_t at fixed shape can be fairly well fit by a constant power law. On equilateral triangles the amplitude is roughly

$$f_{\text{NL}} = 24 \left(\frac{k_t/3}{k_{\ell=1}} \right)^{-0.789}, \quad (4.12)$$

and on squeezed triangles with $k_1/k_t = 0.0025$ we have

$$f_{\text{NL}} = 94 \left(\frac{k_t/3}{k_{\ell=1}} \right)^{-0.716}. \quad (4.13)$$

Bearing in mind the size of η_σ , the scaling in (4.12) is a reasonable match for our lowest-order slow-roll prediction $k_t^{4\eta_\sigma - 4\eta_\phi} \sim k_t^{-0.92}$ obtained in (3.9)—and an even better match if the only large next-order term is included to give $k_t^{4\eta_\sigma + 4\eta_\sigma^2/3 - 4\eta_\phi} \sim k_t^{-0.84}$ [43, 52], within 6% of the measured result. (We have again approximated $n_s - 1 \approx 2\eta_\phi$ as in §4.1.) The accuracy of the slow-roll prediction is rather striking, and we have verified that similar accuracy persists even for larger values of $|\eta_1|$.

Comparison of Eqs. (4.12) and (4.13) shows that the k_t dependence varies with shape, in contradiction with the slow-roll prediction (3.9) which depends only on the Γ -matrices and is shape-independent. We interpret the shape dependence as a dominant contribution from (3.9) corrected by a smaller shape-dependent contribution from the n -point functions $\langle \delta\phi^a \delta\phi^b \rangle$, $\langle \delta\phi^a \delta\phi^b \delta\phi^c \rangle$ which were neglected in (3.9). In this model the variation in these n -point functions need not be small because of the large η_σ when relevant scales were leaving the horizon.

Collecting Eqs. (4.2), (4.3) and (4.12) yields the simple relations

$$A(k) \sim f_{\text{NL}}(k, k, k_3) \sim f_{\text{NL}}^{1/2}(k, k, k) \sim \frac{\mathcal{P}_\sigma(k)}{\mathcal{P}(k)}, \quad (4.14)$$

which should be interpreted as statements about the scaling behaviour of each quantity as a function of k , with all other quantities such as k_3 held fixed. The first follows from Eq. (2.24) and applies to any multi-source scenario in which a single source generates the bispectrum. The second follows from our assumption that the scale-dependence is generated by a large η_σ and follows from Eqs. (3.11) and (4.12). Finally, the third relation is another consequence of our assumption that σ dominates the bispectrum. The same scalings (4.14) will apply to any scenario which satisfies these criteria. Notice that the asymmetry scales like $f_{\text{NL}}^{1/2}(k, k, k)$, in contrast to the single-source case for which the asymmetry is independent of the scaling of the σ power spectrum and we instead have

$$A(k) \sim f_{\text{NL}}(k, k, k) \quad (\text{single source}). \quad (4.15)$$

Variation with shape.—The right-hand panel of Fig. 6 shows the variation with squeezing k_3/k_t on isosceles triangles at fixed scale. There is an approximate fit to a constant power law, but also some evidence for a change in the slope between large and small k_3/k_t . For the configuration whose average side $k_t/3$ left the horizon at time $N = 2$ we have, approximately,

$$f_{\text{NL}} \approx 58 \left(\frac{k_3/k_t}{0.01} \right)^{-0.255}, \quad (4.16)$$

and for the configuration whose average side left the horizon at time $N = 8$ we have, approximately,

$$f_{\text{NL}} \approx 0.73 \left(\frac{k_3/k_t}{0.01} \right)^{-0.365}. \quad (4.17)$$

Eqs. (4.16) and (4.17) show reasonable agreement with the slow-roll prediction (3.10). The discrepancy is presumably accounted for as above by corrections from larger-than-slow-roll scaling of $\langle \delta\phi^a \delta\phi^b \rangle$, $\langle \delta\phi^a \delta\phi^b \delta\phi^c \rangle$.

Observational constraints.—Fig. 6 shows that the bispectrum amplitude is large on some configurations but small on others. The constraints reported by the Planck collaboration are limits on the amplitude of scale-independent templates [76, 77] averaged over many configurations, and therefore as explained in §1 none of these can be related directly to a bispectrum which runs significantly with scale.

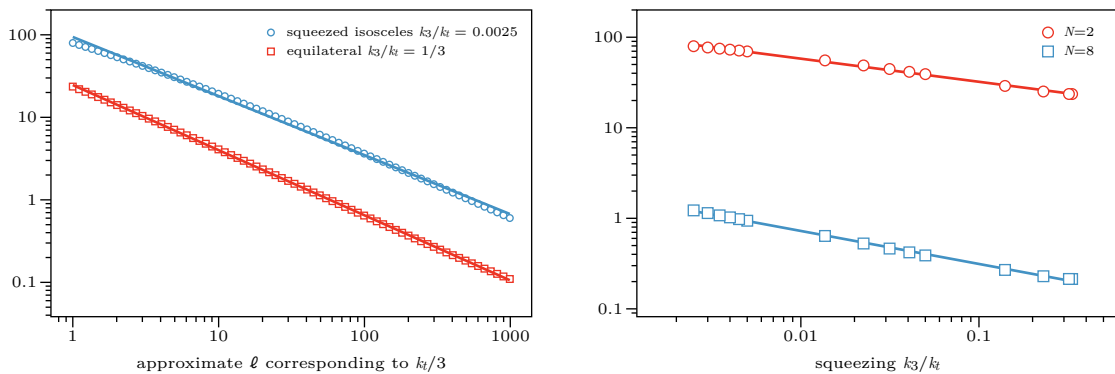


Figure 6: Dependence of the reduced bispectrum $f_{\text{NL}}(k_1, k_2, k_3)$ on scale and shape. The marked points are samples from our numerical results.

Left panel: variation of $f_{\text{NL}}(k_1, k_2, k_3)$ with k_t (represented by the corresponding approximate multipole ℓ) at fixed shape for two different shapes, (i) equilateral triangles, and (ii) squeezed isosceles triangles with $k_3/k_t = 0.0025$.

Right panel: variation of $f_{\text{NL}}(k_1, k_2, k_3)$ with k_3/k_t at fixed scale.

To determine how the estimators for the local, equilateral and orthogonal amplitudes would respond to the bispectrum produced by (4.1) we construct a Fisher matrix estimate. We numerically compute $\sim 5 \times 10^6$ bispectrum configurations for (4.1) covering the range from $\ell \sim 1$ to $\ell \sim 7000$ and use these to predict the observed angular bispectrum $b_{\ell_1 \ell_2 \ell_3}$ up to $\ell \sim 2000$ using the method of Refs. [78–81] and realistic estimates for the Planck beam and noise [82].

We find that the primordial bispectrum $B(k_1, k_2, k_3)$ is roughly 60% correlated with the local template in \mathbf{k} -space. Although the shape produced by (4.1) is ‘local-like’ in the sense of (2.3), the decorrelation can be attributed to the strong running with shape and scale. Despite the modest \mathbf{k} -space correlation, the angular bispectrum $b_{\ell_1 \ell_2 \ell_3}$ is 95% correlated with the local template in ℓ -space. The difference is caused by redistribution of power in the mapping from \mathbf{k} to ℓ , which reroutes structure towards modestly squeezed ℓ -configurations by mixing contributions from the low- k regime. This increases overlap with the local shape, although there is still marginally enhanced power on large scales; see Fig. 7.

In ℓ -space our bispectrum correlates at 98% with a local template normalized as in Eqs. (4.12)–(4.13), but with amplitude scaling as the power-law $k_t^{-0.7}$ and no dependence on the squeezing k_i/k_t . This template is a good match for our bispectrum in ℓ -space, and a better approximation than the pure local template.

For scale-independent bispectrum shapes it usually happens that the \mathbf{k} -space correlation is a good predictor for the ℓ -space correlation. In our example this is not true and the \mathbf{k} -space correlation would produce a misleading result. We believe this to be a fairly general feature of scale-dependent shapes, and if so it will imply that a running bispectrum shape intended to explain the asymmetry must be projected into ℓ -space before robust conclusions can be extracted regarding its observational viability.

We find that the amplitudes which would be measured for a bispectrum generated

by (4.1) are order unity. We obtain

$$\hat{f}_{\text{NL}}^{\text{local}} = 0.25, \quad \hat{f}_{\text{NL}}^{\text{equi}} = 0.6, \quad \hat{f}_{\text{NL}}^{\text{ortho}} = -1.0. \quad (4.18)$$

These amplitudes can be regarded as weighted averages of $f_{\text{NL}}(k_1, k_2, k_3)$ over many configurations. Since Fig. 6 shows that some configurations reach amplitudes of order $O(100)$, the $O(1)$ values reported in (4.18) imply that, on the configurations which contribute most signal-to-noise to the estimators $\hat{f}_{\text{NL}}^{\text{local}}$, $\hat{f}_{\text{NL}}^{\text{equi}}$ and $\hat{f}_{\text{NL}}^{\text{ortho}}$, the amplitude $f_{\text{NL}}(k_1, k_2, k_3)$ has already run to small values; in fact, the amplitude of $\hat{f}_{\text{NL}}^{\text{local}}$ agrees well with the value which would be inferred from (4.12) evaluated at the Planck pivot scale $0.05h \text{ Mpc}^{-1}$ (corresponding approximately to $\ell \sim 700$). These numbers should be compared to the Planck2013 temperature-only⁹ constraints [77]

$$\hat{f}_{\text{NL}}^{\text{local}} = 2.5 \pm 5.7, \quad \hat{f}_{\text{NL}}^{\text{equi}} = -16 \pm 70, \quad \hat{f}_{\text{NL}}^{\text{ortho}} = -34 \pm 33. \quad (4.19)$$

We conclude that the bispectrum amplitude generated by this model is well within the observational limits.

Eq. (4.18) appears to contrast with the much larger estimates appearing in Refs. [13, 17, 18, 23]. However, because our model still requires $f_{\text{NL}}(k_1, k_2, k_3)$ to be large on the configurations responsible for determining $A(k)$, our analysis does not disagree with the qualitative conclusions in these papers. Eqs. (4.18) are predictions for what would be observed in a realistic experiment, and this prediction is possible only because we have a concrete model, enabling the bispectrum to be accurately computed.

The numerical bispectrum we have used is strictly valid only for the model of Eq. (4.1), but in practice we believe it will be a good proxy for the bispectrum generated in any model which uses a large η_σ . If so, then after suitable rescaling Eq. (4.18) may be used to estimate the amplitudes produced in any model designed to generate the asymmetry by this mechanism.

Because our estimates (4.18) are so small there is enough headroom, if desired, to increase the amplitude of the bispectrum and decrease the exceptionality E . Based on the 95% correlation with the local template in ℓ -space we expect it is possible to increase the bispectrum amplitude by a factor of roughly ~ 50 while staying within the 2σ error bar. If this were done it would allow $f_{\text{NL}}(k_1, k_2, k_3)$ to have an amplitude as large as $O(10^3)$ on the configurations which determine $A(k)$ while still satisfying observational constraints.

Trispectrum.—We do not study the trispectrum in detail, because to do so accurately would require numerical calculations of the four-point functions which have not yet been developed. Instead we use Eq. (3.18) to estimate the amplitude of τ_{NL} , which should give fair results for both tetrahedral and collapsed configurations. However, because the scale- and shape-dependence of the τ_{NL} shape is rather strong, comparison with the upper bound $\tau_{\text{NL}} < 2800$ reported by Planck at 95%-confidence [76] is uncertain. A trustworthy estimate should take into account which configurations contribute the largest signal-to-noise. We provide numerical values for qualitative guidance only.

⁹We quote the temperature-only constraints because our analysis does not include polarization data.




Figure 7: Left: ℓ -space bispectrum generated by the model (4.1).

Right: ℓ -space bispectrum generated by the local template.

The shapes are 95% correlated even though the underlying k -space bispectra exhibit only 60% correlation. The ℓ -space shapes are normalised as described in [76]. Red regions represent positive values and blue regions representing negative values. By comparison with the local template, the red tilt can be seen qualitatively to enhance power at small ℓ .

We find that the amplitude is strongly scale-dependent, becoming very large on long scales but running to smaller values on short scales. We estimate that it has a rough scale-dependence $\propto k^{-0.7}$, with amplitude running from $\tau_{\text{NL}} \sim 30,000$ on the scale $\ell \approx 1$ to $\tau_{\text{NL}} \sim 400$ on the scale $\ell \approx 500$. Bearing in mind that, as for the bispectrum, the signal-to-noise for the estimator $\hat{\tau}_{\text{NL}}$ may receive its largest contribution from configurations at modest or high ℓ , these numbers suggest that τ_{NL} is not so large that the model is obviously unacceptable. However, this should be confirmed by a more accurate analysis.

5 Conclusions

In this section we collect our conclusions.

Technical results.—Our principal theoretical results are Eqs. (2.11) for the response function ρ_μ in a ‘local-like’ model, and (2.29) for the response to a long-wavelength ζ perturbation in the special case of a single-source model. These results extend the analyses given in Refs. [17, 18, 23, 24] which assumed a single-source model for which slow-roll was a good approximation while relevant scales were leaving the horizon, and used the separate universe approximation to estimate biasing of the short-wavelength power spectrum. Our key tools are the operator product expansion and linear response theory, of the kind used widely in applications of field theory to condensed matter. Our method does not invoke the slow-roll

approximation and applies to ‘local-like’ models where the long-wavelength response of each two-point function is dominated by the operator $\delta\phi^\mu$. It could be easily extended to obtain the response of the two-point function to any local operator of interest.

Specializing to the single-source, slow-roll case we reproduce the formulae of Lyth [17, 18], Namjoo et al [19–21] and Kobayashi, Cortês & Liddle [23]. We verify the statement made in Ref. [18], that in these models it is the reduced bispectrum on equilateral configurations which controls the response of the two-point function to biasing. This is a special case of the more general result (2.24) that if a single source dominates the bispectrum then the asymmetry scales with k like $f_{\text{NL}}(k, k, k_3)$ at fixed $k_3 \ll k$. The special feature of truly single-source models is that $f_{\text{NL}}(k, k, k_3)$ is independent of k_3 , making the amplitude the same as the equilateral configuration $f_{\text{NL}}(k, k, k)$.

If more than one field contributes to the bispectrum then Eq. (2.11) shows that it is a combination of suitable response functions ρ_μ rather than the reduced bispectrum which will determine the response. These response functions are determined from squeezed isosceles configurations of the mixed three-point function $\langle\delta\phi^\alpha\zeta\zeta\rangle$. However, because the linear response calculation for local-like models predicts that each n -point function responds to all long-wavelength modes in the same way, the response does not depend on the squeezing ratio k_3/k_t . Nevertheless this does not mean it is related in any simple way to the bispectrum amplitude on equilateral configurations, although in some models that will be the case.

In §3.1 we have developed a formalism to compute the scale- and shape-dependence of each n -point function without invoking a perturbative expansion of the ‘separate universe coefficients’ $\Gamma_a^\alpha, \Gamma_{ab}^\alpha$.

Model building constraints.—Even if we know how to compute the response it is still necessary to construct a model. In §3 we have described in general terms why this is difficult. If one generates scale-dependence using a large η -parameter associated with an isocurvature field σ then the potential cannot be featureless, and it appears unavoidable that this introduces fine-tuning. Also, the generic result is contamination of the ζ spectral index if σ contributes to the ζ two-point function, or an enhanced trispectrum amplitude if it does not. In either case it may also be necessary to tune the time and mechanism by which inflation ends. If one instead generates scale-dependence using a large ξ -parameter while keeping η small then the time dependence of ξ must be tuned to give an approximate power-law. We have not succeeded in constructing an example model of this type because typically a large ξ sources a large η within a few e-folds, and it is not clear whether this problem can be overcome.¹⁰

To exemplify these difficulties we have constructed an explicit model giving an acceptable fit to present constraints on the ζ two-, three- and four-point functions, and avoiding (if only marginally) pathologies such as a quantum diffusion regime. This model has a large η parameter for the isocurvature field σ , and therefore the slow-roll approximation is not obviously acceptable. Instead, to obtain accurate predictions, we have used a numerical method

¹⁰It can be overcome if the effective ξ is oscillatory, as in the model of Enqvist et al. [83]. Then η remains small due to cancellation of opposite-sign contributions from ξ . However the oscillations lead to oscillatory effects in the bispectrum, meaning it is still not easy to manufacture an approximate power-law over a sufficient number of e-folds.

to estimate the two- and three-point functions and the response functions. We find (perhaps surprisingly) that the slow-roll estimates continue to apply.

To determine whether the bispectrum amplitude is compatible with recent constraints from Planck we compute the angular bispectrum $b_{\ell_1\ell_2\ell_3}$ up to $\ell \sim 2000$ and obtain the response of the local, equilateral and orthogonal estimators $\hat{f}_{\text{NL}}^{\text{local}}$, $\hat{f}_{\text{NL}}^{\text{equi}}$ and $\hat{f}_{\text{NL}}^{\text{ortho}}$. We conclude that these are all order unity. Despite the large reduced bispectrum $f_{\text{NL}}(k_1, k_2, k_3)$ on those configurations responsible for the asymmetry $A(k)$, this shows that the model is comfortably compatible with present-day constraints; indeed, it is even possible to increase the bispectrum amplitude if we wish to decrease the required enhancement factor E .

Phenomenology.—Our numerical computations use the precise bispectrum generated by the step-like model (4.1). However the details of the bispectrum do not strongly depend on the model; the close correspondence between the generic estimates obtained in §3.1 and those measured from our numerics show that the bispectrum is mostly determined by the large value for η_σ . Neither the scale- or shape-dependence is influenced by the tanh step, which is only important to obtain a suitable amplitude. Therefore the results reported in §4.1—especially the estimates for $\hat{f}_{\text{NL}}^{\text{local}}$, $\hat{f}_{\text{NL}}^{\text{equi}}$ and $\hat{f}_{\text{NL}}^{\text{ortho}}$ —have a wider significance for models which attempt to explain the asymmetry using a bispectrum of this type. In particular, our results could be used to estimate $\hat{f}_{\text{NL}}^{\text{local}}$, $\hat{f}_{\text{NL}}^{\text{equi}}$ and $\hat{f}_{\text{NL}}^{\text{ortho}}$ for such models by suitable rescaling.

In addition, we have clarified the scaling properties of $A(k)$ in different scenarios (§4.3). In truly single-source scenarios it is already known that $A(k) \sim f_{\text{NL}}(k, k, k)$ [17, 18, 23]. In a more general class of scenarios where a single source dominates the bispectrum (but not necessarily the two-point function) we have shown that $A(k) \sim f_{\text{NL}}(k, k, k_3)$ at fixed k_3 . If the field responsible for sourcing the bispectrum does not dominate $\mathcal{P}(k)$ then when rewritten in terms of the equilateral amplitude $f_{\text{NL}}(k, k, k)$ the dependence of $f_{\text{NL}}(k, k, k_3)$ on k_3 changes the scaling law to $A(k) \sim f_{\text{NL}}^{1/2}(k, k, k)$. In this sense the two scenarios are strikingly different, rather than one being a perturbative refinement of the other.

Using the operator product expansion we provide the formulae (2.24) and (2.26) which relate the asymmetry amplitude $A(k)$ to $f_{\text{NL}}(k_1, k_2, k_3)$ on squeezed configurations, and $\tau_{\text{NL}}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4)$ on collapsed configurations. These relations are model independent, assuming only that a single field dominates the bispectrum and trispectrum respectively, and apply for any bispectrum shape controlled by the local part of the OPE. Together they predict an enhanced τ_{NL} amplitude whenever the σ field does not contribute to the ζ two-point function.

Discussion.—In our opinion none of the early-universe scenarios which have been proposed to date are compelling. Whether we are forced to take them seriously may become clearer once polarization data become available, which will provide an independent probe of the amplitude of long-wavelength modes.

If the asymmetry is not a statistical accident, and is to be explained by using a bispectrum to couple long- and short-scale modes, then it has been understood for a long time that a single-field inflationary model is not viable [15]. In this paper we additionally argue that although multiple-field models satisfying the required observational criteria may exist, they typically require multiple independent fine-tunings including at least some of the following.

1. The existence of a long-wavelength mode enhanced by an exceptionality $E \gtrsim 10$ compared to the naïve estimate from power-law scaling. In the absence of new physics to explain its amplitude this would correspond to a 10σ fluctuation or more, and is a poor explanation of a 3σ anomaly. If there is new physics which naturally makes the exceptionality large—for example, in the scenarios of Refs. [35, 36]—then care must be taken to include the effect of many long-wavelength modes with similar amplitude [13, 31]. This may perhaps lead to a large quadrupolar modulation of the power spectrum.¹¹
2. A tuning of the Taylor coefficient of order $(\mathbf{k}_L \cdot \mathbf{x})^0$ with respect to that of order $(\mathbf{k}_L \cdot \mathbf{x})^1$ for the long-wavelength mode. For small α this is required to prevent $C(k) \propto A(k)/\alpha$ growing too large, as described below Eq. (2.18), and generating an unwanted scale-dependent, monopolar modulation of power. This tuning is independent of the exceptionality E and the amplitude of the bispectrum.
3. A large effective mass, corresponding to a large η_σ , or large self-interaction of the scalar field σ which generates the asymmetry. In both cases the initial value of σ requires significant fine-tuning to generate a bispectrum of sufficient amplitude.

- **Large η_σ .**—In this case the power spectrum must predominantly be generated by the inflaton field ϕ in order to preserve the near scale-invariance of the ζ power spectrum. We have shown that, if this is achieved by keeping $R \equiv \mathcal{P}_{\zeta_\sigma}/\mathcal{P}_\zeta \lesssim 0.1$ throughout the evolution, then η_σ cannot be constant without trespassing on the diffusion region near a hilltop of the potential.

Although one can construct models which avoid this constraint by introducing features, such as the step-like tanh model in §4, the location of the feature represents another fine-tuning. In addition it adds complexity to a model which is already complicated. Even if this can be done successfully the amplitude of the τ_{NL} trispectrum shape is typically enhanced above its single-source value $\sim f_{\text{NL}}^2$. Because of the strong scale-dependence of the τ_{NL} amplitude we are not yet able to determine conclusively whether the model of §4 is ruled out by observation.

- **Large self-interaction.**—In this case the large self-interaction will typically lead to a large effective mass as σ evolves, potentially spoiling scale-invariance of the ζ power spectrum. It also generates a large contribution to the amplitude of the g_{NL} trispectrum shape. If ξ is constant then this amplitude may be close to scale invariant (unlike f_{NL} or τ_{NL}) and therefore in conflict with observation; alternatively, if ξ evolves in such a way that it produces an approximate power-law bispectrum the situation is less clear. An example of this case was studied in Ref. [64], where it was demonstrated that the quadrupolar modulation of the power spectrum is typically much too large.

¹¹Quadrupolar modulation would correspond to the next-order term in (2.16), giving an angular dependence of the form $(\hat{\mathbf{p}} \cdot \hat{\mathbf{n}})^2$.

4. The long-wavelength mode must respect observational constraints on C_ℓ for low ℓ . Although we have not done so here, other authors have invoked fine-tuning of our position on the long wavelength mode to help evade these constraints [23]. Strategies of this sort are possible but unattractive, and indeed it is not clear whether they remain viable if enough ℓ -modes are considered. Alternatively, the Grischuk–Zel’dovich contributions to C_2 and C_3 can be suppressed by increasing the wavelength of the large-scale mode, but that simultaneously increases the tuning required in $|C(k)|$.

All these challenges can be understood as manifestations of the difficulty in constructing a $\sim 20\%$ modulation of the power spectrum amplitude on large scales while maintaining consistency with observational constraints on the smallness of the bi- and trispectrum, and quadrupolar modulation of the power spectrum. Even worse, the largest CMB multipoles are actually observed to be suppressed whereas the existence of a large amplitude long-wavelength fluctuation would naturally be expected to enhance them. It is challenging to construct any model which seeks to explain these conflicting demands without emerging as more unlikely than the hemispherical asymmetry itself.

Acknowledgements

DS acknowledges support from the Science and Technology Facilities Council [grant number ST/L000652/1]. CTB is a Royal Society University Research Fellow. The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007–2013) / ERC Grant Agreement No. [308082]. This work was supported in part by National Science Foundation Grant No. PHYS-1066293 and the hospitality of the Aspen Center for Physics.

Some numerical work was undertaken on the COSMOS Shared Memory system at DAMTP, University of Cambridge, operated on behalf of the STFC DiRAC HPC Facility. This equipment is funded by BIS National E-infrastructure capital grant ST/J005673/1 and STFC grants ST/H008586/1, ST/K00333X/1.

We would like to thank Shaun Hotchkiss, Zac Kenton, Takeshi Kobayashi, Antony Lewis, David Mulryne, Marjorie Schillo and Sarah Shandera for helpful discussions.

Data availability statement.—Please contact the authors to obtain the bispectrum for the step model (4.1), which was used to estimate the responses (4.18).

References

- [1] D. J. Schwarz, C. J. Copi, D. Huterer, and G. D. Starkman, *CMB Anomalies after Planck*, 2015. [arXiv:1510.07929](https://arxiv.org/abs/1510.07929).
- [2] H. K. Eriksen, F. K. Hansen, A. J. Banday, K. M. Gorski, and P. B. Lilje, *Asymmetries in the Cosmic Microwave Background anisotropy field*, *Astrophys. J.* **605** (2004) 14–20, [[astro-ph/0307507](https://arxiv.org/abs/astro-ph/0307507)]. [Erratum: *Astrophys. J.*609,1198(2004)].
- [3] F. K. Hansen, A. J. Banday, and K. M. Gorski, *Testing the cosmological principle of isotropy: Local power spectrum estimates of the WMAP data*, *Mon. Not. Roy. Astron. Soc.* **354** (2004) 641–665, [[astro-ph/0404206](https://arxiv.org/abs/astro-ph/0404206)].

- [4] H. K. Eriksen, A. J. Banday, K. M. Gorski, F. K. Hansen, and P. B. Lilje, *Hemispherical power asymmetry in the three-year Wilkinson Microwave Anisotropy Probe sky maps*, *Astrophys. J.* **660** (2007) L81–L84, [[astro-ph/0701089](#)].
- [5] J. Hoftuft, H. K. Eriksen, A. J. Banday, K. M. Gorski, F. K. Hansen, and P. B. Lilje, *Increasing evidence for hemispherical power asymmetry in the five-year WMAP data*, *Astrophys. J.* **699** (2009) 985–989, [[arXiv:0903.1229](#)].
- [6] **Planck** Collaboration, P. A. R. Ade *et al.*, *Planck 2013 results. XXIII. Isotropy and statistics of the CMB*, *Astron. Astrophys.* **571** (2014) A23, [[arXiv:1303.5083](#)].
- [7] **Planck** Collaboration, P. A. R. Ade *et al.*, *Planck 2015 results. XVI. Isotropy and statistics of the CMB*, [[arXiv:1506.07135](#)].
- [8] F. Paci, A. Gruppuso, F. Finelli, A. De Rosa, N. Mandolesi, and P. Natoli, *Hemispherical power asymmetries in the WMAP 7-year low-resolution temperature and polarization maps*, *Mon. Not. Roy. Astron. Soc.* **434** (2013) 3071, [[arXiv:1301.5195](#)].
- [9] Y. Akrami, Y. Fantaye, A. Shafieloo, H. K. Eriksen, F. K. Hansen, A. J. Banday, and K. M. Górski, *Power asymmetry in WMAP and Planck temperature sky maps as measured by a local variance estimator*, *Astrophys. J.* **784** (2014) L42, [[arXiv:1402.0870](#)].
- [10] S. Flender and S. Hotchkiss, *The small scale power asymmetry in the cosmic microwave background*, *JCAP* **1309** (2013) 033, [[arXiv:1307.6069](#)].
- [11] C. M. Hirata, *Constraints on cosmic hemispherical power anomalies from quasars*, *JCAP* **0909** (2009) 011, [[arXiv:0907.0703](#)].
- [12] S. Aiola, B. Wang, A. Kosowsky, T. Kahniashvili, and H. Firouzjahi, *Microwave Background Correlations from Dipole Anisotropy Modulation*, *Phys. Rev.* **D92** (2015), no. 6 063008, [[arXiv:1506.04405](#)].
- [13] S. Adhikari, S. Shandera, and A. L. Erickcek, *Large-scale anomalies in the Cosmic Microwave Background as signatures of non-Gaussianity*, [[arXiv:1508.06489](#)].
- [14] S. Mukherjee and T. Souradeep, *Litmus Test for Cosmic Hemispherical Asymmetry in the CMB B-mode polarization*, [[arXiv:1509.06736](#)].
- [15] A. L. Erickcek, M. Kamionkowski, and S. M. Carroll, *A Hemispherical Power Asymmetry from Inflation*, *Phys. Rev.* **D78** (2008) 123520, [[arXiv:0806.0377](#)].
- [16] A. L. Erickcek, S. M. Carroll, and M. Kamionkowski, *Superhorizon Perturbations and the Cosmic Microwave Background*, *Phys. Rev.* **D78** (2008) 083012, [[arXiv:0808.1570](#)].
- [17] D. H. Lyth, *The CMB modulation from inflation*, *JCAP* **1308** (2013) 007, [[arXiv:1304.1270](#)].
- [18] D. H. Lyth, *Generating f_{NL} at $\ell \lesssim 60$* , *JCAP* **1504** (2015), no. 04 039, [[arXiv:1405.3562](#)].
- [19] M. H. Namjoo, S. Baghran, and H. Firouzjahi, *Hemispherical Asymmetry and Local non-Gaussianity: a Consistency Condition*, *Phys. Rev.* **D88** (2013) 083527, [[arXiv:1305.0813](#)].
- [20] A. A. Abolhasani, S. Baghran, H. Firouzjahi, and M. H. Namjoo, *Asymmetric Sky from the Long Mode Modulations*, *Phys. Rev.* **D89** (2014), no. 6 063511, [[arXiv:1306.6932](#)].
- [21] M. H. Namjoo, A. A. Abolhasani, S. Baghran, and H. Firouzjahi, *CMB Hemispherical Asymmetry: Long Mode Modulation and non-Gaussianity*, *JCAP* **1408** (2014) 002, [[arXiv:1405.7317](#)].
- [22] S. Kanno, M. Sasaki, and T. Tanaka, *A viable explanation of the CMB dipolar statistical*

- anisotropy*, *PTEP* **2013** (2013) 111E01, [[arXiv:1309.1350](#)].
- [23] T. Kobayashi, M. Cortês, and A. R. Liddle, *A separate universe view of the asymmetric sky*, *JCAP* **1505** (2015), no. 05 029, [[arXiv:1501.05864](#)].
- [24] Z. Kenton, D. J. Mulryne, and S. Thomas, *Generating the cosmic microwave background power asymmetry with g_{NL}* , *Phys. Rev.* **D92** (2015) 023505, [[arXiv:1504.05736](#)].
- [25] J. McDonald, *Hemispherical Power Asymmetry from Scale-Dependent Modulated Reheating*, *JCAP* **1311** (2013) 041, [[arXiv:1309.1122](#)].
- [26] P. Creminelli, G. D’Amico, M. Musso, and J. Norena, *The (not so) squeezed limit of the primordial 3-point function*, *JCAP* **1111** (2011) 038, [[arXiv:1106.1462](#)].
- [27] E. Pajer, F. Schmidt, and M. Zaldarriaga, *The Observed Squeezed Limit of Cosmological Three-Point Functions*, *Phys. Rev.* **D88** (2013), no. 8 083502, [[arXiv:1305.0824](#)].
- [28] K. Hinterbichler, L. Hui, and J. Khoury, *Conformal Symmetries of Adiabatic Modes in Cosmology*, *JCAP* **1208** (2012) 017, [[arXiv:1203.6351](#)].
- [29] P. Creminelli, J. Norena, and M. Simonovic, *Conformal consistency relations for single-field inflation*, *JCAP* **1207** (2012) 052, [[arXiv:1203.4595](#)].
- [30] M. Mirbabayi and M. Simonović, *Effective Theory of Squeezed Correlation Functions*, [arXiv:1507.04755](#).
- [31] F. Schmidt and L. Hui, *Cosmic Microwave Background Power Asymmetry from Non-Gaussian Modulation*, *Phys. Rev. Lett.* **110** (2013) 011301, [[arXiv:1210.2965](#)]. [Erratum: *Phys. Rev. Lett.*110,059902(2013)].
- [32] Z. Kenton and D. J. Mulryne, *The squeezed limit of the bispectrum in multi-field inflation*, [arXiv:1507.08629](#).
- [33] C.-T. Chiang, C. Wagner, F. Schmidt, and E. Komatsu, *Position-dependent power spectrum of the large-scale structure: a novel method to measure the squeezed-limit bispectrum*, *JCAP* **1405** (2014) 048, [[arXiv:1403.3411](#)].
- [34] F. Schmidt, N. E. Chisari, and C. Dvorkin, *Imprint of inflation on galaxy shape correlations*, *JCAP* **1510** (2015), no. 10 032, [[arXiv:1506.02671](#)].
- [35] A. R. Liddle and M. Cortês, *Cosmic microwave background anomalies in an open universe*, *Phys. Rev. Lett.* **111** (2013), no. 11 111302, [[arXiv:1306.5698](#)].
- [36] A. Mazumdar and L. Wang, *CMB dipole asymmetry from a fast roll phase*, *JCAP* **1310** (2013) 049, [[arXiv:1306.5736](#)].
- [37] J. McDonald, *Negative Running of the Spectral Index, Hemispherical Asymmetry and the Consistency of Planck with Large r* , *JCAP* **1411** (2014), no. 11 012, [[arXiv:1403.6650](#)].
- [38] C. R. Contaldi, M. Peloso, and L. Sorbo, *Suppressing the impact of a high tensor-to-scalar ratio on the temperature anisotropies*, *JCAP* **1407** (2014) 014, [[arXiv:1403.4596](#)].
- [39] A. H. Guth and S. Y. Pi, *Fluctuations in the New Inflationary Universe*, *Phys. Rev. Lett.* **49** (1982) 1110–1113.
- [40] D. H. Lyth, *Large Scale Energy Density Perturbations and Inflation*, *Phys. Rev.* **D31** (1985) 1792–1798.
- [41] M. Dias, J. Elliston, J. Frazer, D. Mulryne, and D. Seery, *The curvature perturbation at second order*, *JCAP* **1502** (2015), no. 02 040, [[arXiv:1410.3491](#)].

- [42] M. Dias, R. H. Ribeiro, and D. Seery, *Scale-dependent bias from multiple-field inflation*, *Phys. Rev.* **D87** (2013), no. 10 107301, [[arXiv:1303.6000](#)].
- [43] C. T. Byrnes, M. Gerstenlauer, S. Nurmi, G. Tasinato, and D. Wands, *Scale-dependent non-Gaussianity probes inflationary physics*, *JCAP* **1010** (2010) 004, [[arXiv:1007.4277](#)].
- [44] C. T. Byrnes and J.-O. Gong, *General formula for the running of f_{NL}* , *Phys. Lett.* **B718** (2013) 718–721, [[arXiv:1210.1851](#)].
- [45] E. Tzavara and B. van Tent, *Momentum dependence of the bispectrum in two-field inflation*, *JCAP* **1306** (2013) 001, [[arXiv:1211.6325](#)].
- [46] D. H. Lyth and Y. Rodriguez, *The Inflationary prediction for primordial non-Gaussianity*, *Phys. Rev. Lett.* **95** (2005) 121302, [[astro-ph/0504045](#)].
- [47] D. Seery, D. J. Mulryne, J. Frazer, and R. H. Ribeiro, *Inflationary perturbation theory is geometrical optics in phase space*, *JCAP* **1209** (2012) 010, [[arXiv:1203.2635](#)].
- [48] S. Yokoyama, T. Suyama, and T. Tanaka, *Primordial Non-Gaussianity in Multi-Scalar Inflation*, *Phys. Rev.* **D77** (2008) 083511, [[arXiv:0711.2920](#)].
- [49] S. Yokoyama, T. Suyama, and T. Tanaka, *Primordial Non-Gaussianity in Multi-Scalar Slow-Roll Inflation*, *JCAP* **0707** (2007) 013, [[arXiv:0705.3178](#)].
- [50] S. Yokoyama, T. Suyama, and T. Tanaka, *Efficient diagrammatic computation method for higher order correlation functions of local type primordial curvature perturbations*, *JCAP* **0902** (2009) 012, [[arXiv:0810.3053](#)].
- [51] G. J. Anderson, D. J. Mulryne, and D. Seery, *Transport equations for the inflationary trispectrum*, *JCAP* **1210** (2012) 019, [[arXiv:1205.0024](#)].
- [52] C. T. Byrnes, S. Nurmi, G. Tasinato, and D. Wands, *Scale dependence of local f_{NL}* , *JCAP* **1002** (2010) 034, [[arXiv:0911.2780](#)].
- [53] D. Seery and J. E. Lidsey, *Primordial non-Gaussianities from multiple-field inflation*, *JCAP* **0509** (2005) 011, [[astro-ph/0506056](#)].
- [54] F. Vernizzi and D. Wands, *Non-gaussianities in two-field inflation*, *JCAP* **0605** (2006) 019, [[astro-ph/0603799](#)].
- [55] D. H. Lyth and I. Zaballa, *A Bound concerning primordial non-Gaussianity*, *JCAP* **0510** (2005) 005, [[astro-ph/0507608](#)].
- [56] **Planck** Collaboration, P. A. R. Ade *et al.*, *Planck 2015 results. XX. Constraints on inflation*, [[arXiv:1502.02114](#)].
- [57] J. Elliston, D. J. Mulryne, D. Seery, and R. Tavakol, *Evolution of f_{NL} to the adiabatic limit*, *JCAP* **1111** (2011) 005, [[arXiv:1106.2153](#)].
- [58] J. Elliston, D. Mulryne, D. Seery, and R. Tavakol, *Evolution of non-Gaussianity in multi-scalar field models*, *Int. J. Mod. Phys.* **A26** (2011) 3821–3832, [[arXiv:1107.2270](#)]. [*Int. J. Mod. Phys. Conf. Ser.*03,203(2011)].
- [59] D. Seery, J. E. Lidsey, and M. S. Sloth, *The inflationary trispectrum*, *JCAP* **0701** (2007) 027, [[astro-ph/0610210](#)].
- [60] D. Seery, M. S. Sloth, and F. Vernizzi, *Inflationary trispectrum from graviton exchange*, *JCAP* **0903** (2009) 018, [[arXiv:0811.3934](#)].
- [61] D. Seery and J. E. Lidsey, *Non-Gaussianity from the inflationary trispectrum*, *JCAP* **0701**

- (2007) 008, [[astro-ph/0611034](#)].
- [62] C. T. Byrnes, M. Sasaki, and D. Wands, *The primordial trispectrum from inflation*, *Phys. Rev. D* **74** (2006) 123519, [[astro-ph/0611075](#)].
- [63] T. Suyama and M. Yamaguchi, *Non-Gaussianity in the modulated reheating scenario*, *Phys. Rev. D* **77** (2008) 023505, [[arXiv:0709.2545](#)].
- [64] C. T. Byrnes and E. R. M. Tarrant, *Scale-dependent non-Gaussianity and the CMB Power Asymmetry*, *JCAP* **1507** (2015) 007, [[arXiv:1502.07339](#)].
- [65] M. Dias, R. H. Ribeiro, and D. Seery, *The δN formula is the dynamical renormalization group*, *JCAP* **1310** (2013) 062, [[arXiv:1210.7800](#)].
- [66] S. Jazayeri, Y. Akrami, H. Firouzjahi, A. R. Solomon, and Y. Wang, *Inflationary power asymmetry from primordial domain walls*, *JCAP* **1411** (2014) 044, [[arXiv:1408.3057](#)].
- [67] C. Ringeval, D. Yamauchi, J. Yokoyama, and F. R. Bouchet, *Large scale CMB anomalies from thawing cosmic strings*, [arXiv:1510.01916](#).
- [68] A. Ashoorioon and T. Koivisto, *Hemispherical Asymmetry from Parity-Violating Excited Initial States*, [arXiv:1507.03514](#).
- [69] H. Assadullahi, H. Firouzjahi, M. H. Namjoo, and D. Wands, *CMB hemispherical asymmetry from non-linear isocurvature perturbations*, *JCAP* **1504** (2015), no. 04 017, [[arXiv:1410.8036](#)].
- [70] G. D'Amico, R. Gobetti, M. Kleban, and M. Schillo, *Large-scale anomalies from primordial dissipation*, *JCAP* **1311** (2013) 013, [[arXiv:1306.6872](#)].
- [71] G. Leung, E. R. M. Tarrant, C. T. Byrnes, and E. J. Copeland, *Reheating, Multifield Inflation and the Fate of the Primordial Observables*, *JCAP* **1209** (2012) 008, [[arXiv:1206.5196](#)].
- [72] J. Meyers and E. R. M. Tarrant, *Perturbative Reheating After Multiple-Field Inflation: The Impact on Primordial Observables*, *Phys. Rev. D* **89** (2014), no. 6 063535, [[arXiv:1311.3972](#)].
- [73] J. Elliston, S. Orani, and D. J. Mulryne, *General analytic predictions of two-field inflation and perturbative reheating*, *Phys. Rev. D* **89** (2014), no. 10 103532, [[arXiv:1402.4800](#)].
- [74] M. Dias, J. Frazer, D. Mulryne, and D. Seery. To appear.
- [75] G. Efstathiou, *A Maximum likelihood analysis of the low CMB multipoles from WMAP*, *Mon. Not. Roy. Astron. Soc.* **348** (2004) 885, [[astro-ph/0310207](#)].
- [76] **Planck** Collaboration, P. A. R. Ade *et al.*, *Planck 2015 results. XVII. Constraints on primordial non-Gaussianity*, [arXiv:1502.01592](#).
- [77] **Planck** Collaboration, P. A. R. Ade *et al.*, *Planck 2013 Results. XXIV. Constraints on primordial non-Gaussianity*, *Astron. Astrophys.* **571** (2014) A24, [[arXiv:1303.5084](#)].
- [78] J. R. Fergusson, M. Liguori, and E. P. S. Shellard, *General CMB and Primordial Bispectrum Estimation I: Mode Expansion, Map-Making and Measures of f_{NL}* , *Phys. Rev. D* **82** (2010) 023502, [[arXiv:0912.5516](#)].
- [79] J. R. Fergusson, D. M. Regan, and E. P. S. Shellard, *Optimal Trispectrum Estimators and WMAP Constraints*, [arXiv:1012.6039](#).
- [80] D. Regan, P. Mukherjee, and D. Seery, *General CMB bispectrum analysis using wavelets and separable modes*, *Phys. Rev. D* **88** (2013), no. 4 043512, [[arXiv:1302.5631](#)].
- [81] D. Regan, M. Gosenca, and D. Seery, *Constraining the WMAP9 bispectrum and trispectrum with needlets*, *JCAP* **1501** (2015), no. 01 013, [[arXiv:1310.8617](#)].

- [82] **CMBPol Study Team** Collaboration, D. Baumann *et al.*, *CMBPol Mission Concept Study: Probing Inflation with CMB Polarization*, *AIP Conf. Proc.* **1141** (2009) 10–120, [[arXiv:0811.3919](#)].
- [83] K. Enqvist, D. J. Mulryne, and S. Nurmi, *Resolving primordial physics through correlated signatures*, *JCAP* **1505** (2015), no. 05 010, [[arXiv:1412.5973](#)].