Non-Gaussianity constrains hybrid inflation

David Mulryne[1,2]**, David Seery**[1] **and Daniel Wesley**[1,3]

[1] Centre for Theoretical Cosmology
Department of Applied Mathematics and Theoretical Physics
Wilberforce Road, Cambridge, CB3 0WA, United Kingdom
[2] Theoretical Physics Group,
Imperial College, London, SW7 2BZ, United Kingdom
[3] Center for Particle Cosmology
David Rittenhouse Laboratory, University of Pennsylvania
209 South 33rd Street, Philadelphia, PA 19104 USA

E-mail: `D.Mulryne@damtp.cam.ac.uk`, `djs61@cam.ac.uk`, `dwes@sas.upenn.edu`

Abstract. In hybrid inflationary models, inflation ends by a sudden instability associated with a steep ridge in the potential. Here we argue that this feature can generate a large contribution to the curvature perturbation on observable scales. This contribution is almost scale-invariant but highly non-Gaussian. The degree of non-Gaussianity can exceed current observational bounds, unless the inflationary scale is extremely low or the hybrid potential contains very large coupling constants. Non-linear effects on small scales may quench the non-Gaussian signal, and while we find no compelling evidence that this occurs, full lattice simulations are required to definitively address this issue.

Note added: We now believe that nonlinear effects will invalidate the original computation in this paper essentially instantaneously after the short-wavelength modes reach the minimum of their potential. This means that the mechanism described in this paper will not lead to appreciable curvature perturbations on long wavelengths, and no useful constraints on hybrid inflation will result. We have inserted a brief calculation on p2 of this manuscript to explain this fact, but have otherwise left the manuscript unchanged.

Keywords: Inflation, cosmological perturbation theory, physics of the early universe, quantum field theory in curved spacetime.

Version 2 of this paper withdraws the claim of large f_{NL}

After further review, we have come to the conclusion that nonlinear effects cut off the growth of long-wavelength modes nearly instantaneously after the short-wavelength modes reach the minimum of their potential. This means that even the small difference in timescale explored in Section 5 is not sufficient to guarantee the growth of appreciable long-wavelength fluctuations. Hence we believe that the mechanism originally described in our paper does not lead to the growth of curvature perturbations on large scales. Below, we present an argument which illustrates this fact.

We isolate the waterfall field χ and assume it has a potential of the form:

$$V(\chi) = \frac{\lambda}{4}\left(\chi^2 - v^2\right) = V_0 - \frac{\mu^2}{2}\chi^2 + \frac{\lambda}{4}\chi^4 \tag{1}$$

where $\mu^2 = \lambda v^2$. (Note μ has a different meaning in this note than in the rest of the paper). The equation of motion for χ is

$$\ddot{\chi} - \nabla^2\chi - \mu^2\chi + \lambda\chi^3 = 0, \tag{2}$$

and holds pointwise. We are really interested in the equation of motion for a related quantity, namely χ "smoothed" on some length scale L. Let us denote this "smoothed" version of $\chi(x)$ by $s_L[\chi](x)$. There are a few different ways to define this smoothed variable, such as convolution with a transfer function or via a cutoff in Fourier space, but all we really need to proceed with the argument is that the smoothing is linear, so

$$s_L[c_1\chi_1 + c_2\chi_2](x) = c_1 s_L[\chi_1](x) + c_2 s_L[\chi_2](x) \tag{3}$$

for all constants c_1 and c_2, and that the smoothing is a projection‡, so

$$s_L[s_L[\chi]] = s_L[\chi] \tag{4}$$

Since the smoothing process is a projection, we can use it to decompose χ into a "smoothed," long-wavelength piece χ_L, and a short-wavelength piece χ_S as follows

$$\chi = \chi_L + \chi_S \tag{5}$$

where $\chi_L(x) = s_L[\chi](x)$ and $\chi_S = \chi - \chi_L$.

We can now study how the short-wavelength dynamics affects the long-wavelength ones. Applying the smoothing process s_L to the equation of motion (2), one finds

$$\ddot{\chi}_L - \mu^2\chi_L + \lambda\chi_L^3 + \lambda\Delta = 0 \tag{6}$$

where

$$\Delta = s_L[\chi^3] - \chi_L^3 \tag{7}$$

We have dropped the gradient term, since L is much larger than the other scales in the problem. Equation (6) means that the smoothed field χ has the same equation of

‡ This property holds only approximately for convolutions, but the approximation will be sufficently accurate for large L.

motion as the unsmoothed field, up to an additional effective term $\lambda\Delta$. To estimate this term, we use the decomposition (5) and the properties (3) and (4) to find

$$\Delta = 3\chi_L^2 s_L[\chi_S] + 3\chi_L s_L[\chi_S^2] + s_L[\chi_S^3] \tag{8}$$

The first term must vanish because of the definition of the decomposition (5). The last term may be nonzero, but it seems likely that it should also vanish because of the $\chi \to -\chi$ symmetry in the problem. This leaves the middle term. Once the short-wavelength modes have reached the minimum of the potential, $\chi_S^2 = v^2$ Hence, once the short scale-wavelength modes have reached the minimum we have

$$\Delta = 3v^2\chi_L \tag{9}$$

Plugging this in to (6) results in

$$\ddot{\chi}_L + 2\mu^2\chi_L + \lambda\chi_L^3 = 0 \tag{10}$$

The short-wavelengthe modes have caused the appearance of an effective mass term in the equation of motion for χ_L. This term cancels the tachyonic mass term present in the action, and stabilizes the long-wavelength modes. This means that by the time the short-wavelength modes have reached the minimum of the potential, the long-wavelength modes should cease evolving essentially instantaneously.

It is perhaps interesting that if only quadratic terms were present in the action (1), then the equations of motion for χ would be linear and the tachyonic mass term would be preserved in the χ_L equation of motion. However, for inflation to end the potential must have a minimum, and the nonlienarities required to introduce the minimum also instantaneously cut off the growth of long-wavelength modes once the short-wavelength modes reach the minimum of the potential.

The correct calculation

The δN formalism can still be used to calculate the contribution to ζ from the hybrid transition, once we account correctly for the short scales modes. When we smooth the universe on large scales, we must remember that although the waterfall field χ averaged on these scales is extremely small, the kinetic energy density of the χ field averaged on these scales has a contribution from modes of all scales, and indeed is dominated by the shortest scale modes present. The time taken for the universe smoothed on large scales to transition from a flat initial hypersurface to a kinetically dominated comoving final hyperspace is therefore not given by equation (46), as we had calculated, but equation (46) with χ_*, the initial field smoothed on large scales, replaced by σ_χ^*, the RMS value of the initial χ field, with all modes included. This quantity is dominated by the shortest scale modes present in the problem, and can be read off from equation (43) using (44) with k given by k_{short}. From the point of view of the large scale modes σ_χ^* is almost a constant, but has a small cosmic variance given by χ_*. The induced variance in N and therefore ζ, due to this cosmic variance in σ_χ^* is therefore given by the expression

$$\delta N = \zeta = \frac{\partial N}{\partial \sigma_{\chi^*}}\chi_* , \tag{11}$$

where N is given by equation (46) multiplied by H and with χ_* replaced by σ_χ^* as discussed above. Therefore, the curvature perturbation on a particular scale is proportional to the field perturbation averaged on that scale, and is therefore exponentially suppressed on the largest scales.

1. Introduction

Recently, a rapid accumulation of data has made clear that the orthodox scenario of structure formation—in which small fluctuations oscillate under the influence of gravity within a almost-uniform primeval plasma—is in excellent agreement with observation. The initial conditions for these fluctuations are determined by the primordial curvature perturbation, ζ, usually supposed to have been synthesized during an earlier cosmic epoch, perhaps during inflation or ekpyrosis. Many versions of inflation are characterized by slow, smooth evolution and yield very Gaussian fluctuations. In other versions evolution is a dramatic process, allowing larger non-Gaussianities.

In this paper, we show that significant power and non-Gaussianity can be generated when a set of inflationary trajectories fall from a ridge in field space. Such a process is a crucial element of the hybrid inflationary scenario, introduced by Linde, in which inflation ends through a sudden instability [1, 2, 3]. Quantum fluctuations lead to slight differences in field values at the top of the ridge, causing different regions of the universe to follow one of a narrow bundle of trajectories. Trajectories near the edge of the bundle are ejected from the ridge early in their evolution, whereas those in the core remain on the arête longer. The dispersion generated in this way induces a variation in expansion history ("δN") from trajectory to trajectory within the bundle, resulting in a nearly scale-invariant but highly non-Gaussian spectrum of density perturbations. A careful analysis of this process reveals the surprising fact that the final dispersion in the bundle depends only weakly on the initial dispersion. Hence falling from the ridge can have a large effect on the final curvature perturbation, even if the bundle of trajectories is initially very tightly focused.

Ultimately these contributions to the power spectrum can be traced to the conversion of isocurvature modes into adiabatic ones during descent from the ridge. In the early days of the inflationary paradigm, it was assumed that the curvature perturbation was primarily determined by quantum fluctuations in the inflaton field "freezing out" as successive physical scales left the cosmological horizon. Later, when inflationary models containing two or more fields were constructed, it was understood that the curvature perturbation could grow or decay as neighbouring horizon volumes diverged along adjacent but inequivalent field-space trajectories. The evolution is caused by light isocurvature degrees of freedom which play no dynamical role while cosmic microwave background (CMB) scales are leaving the horizon, but later come to influence the expansion history of the universe. Although models with isocurvature effects are more complicated, they admit a richer phenomenology. In models with canonically normalized fields the reprocessing of isocurvature modes can lead to a significant non-Gaussian component of ζ [4, 5, 6]. In the present case, the isocurvature modes are massive while CMB scales leave the horizon and are suppressed by cosmological expansion. Nevertheless, falling from the ridge enormously amplifies these perturbations and can produce a significant effect.

Our analysis applies to the original hybrid scenario and many related models. The

qualitative conclusions may apply during a phase of two-field ekpyrosis [7, 8, 9, 10, 11, 12], which also depends on dispersion near a ridge. The inflationary production of non-Gaussianity in dispersing models has been source of interest in recent years. Alabidi determined the non-Gaussian yield in a hybrid model, beginning from a range of locations near but not coincident with the ridge [13]. Byrnes *et al.* performed a similar analysis, showing that the initial conditions could be tuned to produce a large effect [14, 15]. Related studies have appeared in the literature [16, 17].

The generation and evolution of fluctuations in hybrid inflation was studied by Randall, Soljačić & Guth [18], and later García-Bellido, Linde & Wands [19, 20]. More detailed calculations were undertaken by Copeland, Pascoli & Rajantie [21]. The field whose instability ultimately ends inflation is termed the *waterfall field* and at early times its potential is designed so that it is very heavy, causing its perturbations to decay. Only fluctuations orthogonal to the waterfall make an appreciable contribution to the number of e-folds, N, required to reach a subsequent surface of uniform energy density. The result is a spectrum of density perturbations which can be evaluated by ignoring the waterfall field. In the simplest scenario, where only one canonically normalized field ϕ is relevant, the single-field formula $\delta N \simeq H \delta\phi / |\dot{\phi}|$ applies. There is an added complication, because the waterfall field rolls very slowly at the epoch of ejection. Therefore fluctuations in N may be large on scales leaving the horizon at that time, and one must verify that there is not an unacceptable synthesis of topological defects or primordial black holes [20]. The original hybrid model produced a blue spectrum of perturbations, but a larger class of so-called "hill-top" models [22, 23] yield red spectra compatible with CMB constraints [24].

In this paper, we point out that the traditional calculation of curvature perturbations in this model should be augmented with a new contribution sourced by the waterfall transition. While the trajectories remain on the ridge, fluctuations in δN are well-described by conventional perturbation theory applied to the transverse directions, as explained above. At the epoch of ejection, however, fluctuations in the waterfall field become important, *no matter how small they are.* To see that this can be so, note that δN measured between the central trajectory (which does not leave the ridge) and any other trajectory (which falls into a non-inflating minimum) tends to infinity. Thus, even exponentially suppressed fluctuations can be amplified. It follows that a tiny fluctuation between neighbouring horizon volumes may ultimately give rise to a large density fluctuation.

The basic mechanism by which large non-Gaussianity is generated can be illustrated using an inverted simple harmonic oscillator. Suppose we consider a particle of unit mass and position x, for which the Lagrangian

$$L = \frac{1}{2}\dot{x}^2 + \frac{1}{2}\omega^2 x^2 \tag{12}$$

describes an inverted harmonic oscillator of natural frequency ω. We wish to compute the time t required for the particle to roll down to some specified final position, x_F, as a function of its initial position x_0. We assume that $\omega x_F \gg 1$, $x_0 \ll x_F$, and that the

particle has zero initial velocity. At late times $x(t) \sim (x_0/2)e^{\omega t}$, so

$$t(x_0) = \frac{1}{\omega} \ln \frac{2x_F}{|x_0|} \tag{13}$$

Next we study the statistics of t for an ensemble of particles in which x_0 is Gaussian distributed with variance σ^2, assuming $\omega\sigma \ll 1$. The mean rolling time is

$$\langle t \rangle = \int_{-\infty}^{\infty} \frac{e^{-x_0^2/2\sigma^2}}{\sqrt{2\pi}\sigma} \frac{1}{\omega} \ln \frac{2x_F}{|x_0|} \, \mathrm{d}x_0 = \frac{1}{\omega} \left[\frac{\gamma + 3\ln 2}{2} + \ln\left(\frac{x_F}{\sigma}\right) \right] \tag{14}$$

where γ is the Euler–Mascheroni constant. The average time $\langle t \rangle$ depends only weakly on the initial distribution width, σ, with narrower distributions taking longer to roll off the arête. More surprising results obtain if we compute higher statistics of t, such as its variance

$$\sigma_t^2 = \langle (t - \langle t \rangle)^2 \rangle = \frac{\pi^2}{8\omega^2} \tag{15}$$

third moment

$$\mu_3 = \langle (t - \langle t \rangle)^3 \rangle = \frac{7\zeta(3)}{4\omega^3} \tag{16}$$

and fourth moment

$$\mu_4 = \langle (t - \langle t \rangle)^4 \rangle = \frac{7\pi^4}{64\omega^4} \tag{17}$$

where $\zeta(z)$ is Euler's zeta function and $\zeta(3) \sim 1.202$ is Apéry's constant. Unlike $\langle t \rangle$, these moments are completely independent of the initial variance σ^2, and depend only on the curvature of the potential. Therefore, no matter how small the variance associated with x_0, the variance in arrival times is the same. Moreover, one typically measures departures from Gaussianity by computing quantities such as the skew μ_3/σ_t^3, and excess kurtosis $\mu_4/\sigma_t^4 - 3$. In our case these are pure numbers, of order unity, which are completely independent of ω. Therefore, our purely Gaussian initial distribution always becomes highly non-Gaussian through the process of rolling off the hill.

This toy example can be translated into the inflationary context using a simple dictionary. Roughly speaking, the coordinate x represents the waterfall field which mediates the end of hybrid inflation, the end point x_F represents the location of the reheating minimum, and the time t represents the number of e-foldings required to reach reheating, at which point inflation ends. The initial Gaussian distribution of release points models the quantum flucutations in the waterfall field just before the hybrid transition. Fluctuations in the e-folding history are related to the curvature perturbation, ζ. Therefore, the moments of t represent the moments of ζ. The fact that they are large suggests that large non-Gaussianity can be generated by the hybrid waterfall. Of course, the full inflationary story is much more complicated, and we will support our conclusions with semi-analytic arguments and numerical calculations using realistic forms of the inflationary potential. However, the basic mechanism is sufficiently general that its essence is captured by the inverted harmonic oscillator.

This mechanism generates power on all scales. In the traditional picture of hybrid inflation, effects occuring near the waterfall transition could only generate significant

power on short scales. For example, Randall *et al.* and García-Bellido *et al.* identified large fluctuations which collapse to black holes or topological defects. These are generated on the horizon scale at the epoch of ejection, yielding a spike in the power spectrum at small wavelengths. In contrast, a careful analysis of the process described above indicates that its power spectrum is nearly scale-invariant. This is possible because the final variance between expansion histories depends very weakly on the initial dispersion of release points, as illustrated by the toy model above. During inflation, long-wavelength fluctuations in the waterfall field are inevitably produced, although these are strongly suppressed by cosmic expansion. In the picture we employ, these tiny fluctuations are encoded in the width of the bundle of field-space trajectories followed by cosmological-scale patches of the universe. After the hybrid transition, the tilt of the power spectrum can be determined by comparing the dispersion of the bundle of trajectories followed by large-scale patches with that of horizon-scale patches. Since the final dispersion of trajectories is very weakly dependent on the initial dispersion, the final dispersion in each bundle is roughly the same. This is the hallmark of a scale-invariant process.

Could this new component in the curvature perturbation become observationally detectable? If the growth of perturbations during the waterfall is not curtailed it will be amplified indefinitely, yielding a bispectrum in conflict with observation. However, any hybrid model automatically contains a cut-off on the amplification which can be achieved. While traversing the ridge, small fluctuations push horizon-scale patches to one side or the other, which roll down and reheat into radiation. This so-called tachyonic preheating [25, 26, 27] is non-perturbative, and does not admit a description in terms of a homogeneous coarse-grained background field with small fluctuations. Beyond some limiting time, of order that required for the field to reach the minimum, the evolution of large-scale modes becomes dominated by radiation produced during the reheating phase rather than the inflationary potential. Because the ridge in field space no longer drives trajectories to diverge, the growth of isocurvature modes from the waterfall ceases.

Ultimately, numerical simulations may be required to determine whether a significant effect can be generated before perturbation theory breaks down. One might have expected that a very long time would be required before isocurvature fluctuations from the waterfall could compete with the adiabatic mode. In fact, our analysis suggests that the field smoothed on large scales typically reaches the reheating minimum only a fraction of an e-fold after the field smoothed on the horizon scale. Whether a given model can generate acceptably Gaussian fluctuations depends on a subtle competition between the perturbative growth of isocurvature modes and a non-perturbative cut-off from tachyonic preheating. In our opinion, it is not clear that the outcome can be decided on the basis of perturbative calculations. Our aim is to show that there is serious interest in obtaining the answer. We hope that the quantitative intricacies of the competition can eventually be settled using lattice simulations.

The structure of this paper is as follows. In §2, we fix our notation by reviewing familiar material concerning inflationary perturbations and the δN formalism. In §3 we

give a phenomenological description of evolving field fluctuations near a sharp ridge, and use this to derive some approximate expressions for the spectrum and bispectrum sourced in the curvature perturbation. In §4 we discuss the validity of our calculation. We work in natural units where $c = \hbar = 1$, and the Planck mass associated with Newton's gravitational constant is set to unity, $M_\mathrm{P} = (8\pi G)^{-1/2} = 1$, although M_P will occasionally be restored to clarify the relative magnitude of terms.

2. The δN formalism

Experience has shown that an especially convenient accounting of isocurvature modes can be obtained using the "δN formalism," introduced by Starobinsky [28] and Stewart & Sasaki [29], and extended beyond linear order by Lyth & Rodríguez [5]. The δN formula is an example of the "separate universe picture," according to which distant Hubble volumes evolve like separate unperturbed universes.

The comoving curvature perturbation, ζ, measures fluctuations in the expansion history between spatially disjoint regions of the universe. According to the separate universe picture it can be evaluated using the δN formula,

$$\zeta(t^c, \mathbf{x}) = N(\rho^c, t^*, \mathbf{x}) - N(\rho^c, t^*) \equiv \delta N, \tag{18}$$

where t^* labels an initial spatially flat hypersurface, and t^c labels a subsequent uniform energy-density hypersurface on which the energy takes a prescribed value ρ^c.

Eq. (18) is the lowest term of an expansion in powers of the dimensionless gradient k/aH but is otherwise non-perturbative. Analytic calculations can frequently be simplified by constructing a Taylor expansion of $N(\rho^c, t^*, \mathbf{x})$ around an arbitrary fiducial trajectory. When inserted in expectation values, the high order terms in this expansion generate contributions which grow with volume, describing a delicately shifting pattern of correlations in the large-volume limit. In analogy with the growing ultra-violet ("loop") contributions which govern correlations in quantum field theory as one samples fluctuations of increasing energy, these terms have been interpreted as "classical loops" or "c-loops" [30, 31], and it has become popular to frame δN calculations in this language. In this paper we do not make use of a loop expansion, but apply the non-perturbative definition, Eq. (18), directly. We believe this choice has important advantages. First, issues associated with convergence are avoided. Second, the calculation need not be artificially truncated at a low order in the expansion. Third, it is unnecessary to introduce an infra-red regulator or "factorization" scale (without physical significance) to make intermediate steps of the calculation finite. In quantum field theory the loop expansion is unavoidable owing to our ignorance of ultra-violet physics. In contrast, in many models the infra-red physics associated with Eq. (18) is under reasonable control and in these cases we believe the loop expansion is, at best, an avoidable complication. In unfavourable cases it may be rather misleading.

In the following sections, we supply expressions which allow N to be calculated in a hybrid model. Eq. (18) relates the statistical properties of N to those of the observable

quantity, ζ, by the rules

$$\langle \zeta \zeta \rangle = \langle N^2 \rangle - \langle N \rangle^2, \tag{19}$$

$$\langle \zeta \zeta \zeta \rangle = \langle N^3 \rangle - 3\langle N \rangle \langle N^2 \rangle + 2\langle N \rangle^3. \tag{20}$$

Similar formulae can be written for higher-order correlation functions.

Eqs. (19)–(20) give the variance and skew of a collection of independent spacetime volumes smoothed on some characteristic scale, L. They are not spatially dependent. To obtain the real- and Fourier-space correlation functions which can be compared with observation, one makes a second use of the separate universe picture. This implies that any two spatially disconnected volumes will be uncorrelated. Therefore,

$$\langle \zeta_{\mathbf{x}} \zeta_{\mathbf{y}} \rangle = \sigma^2(L) L^3 \delta(\mathbf{x} - \mathbf{y}), \tag{21}$$

where we have written $\langle \zeta \zeta \rangle_L = \sigma^2(L)$. In Fourier space, the power spectrum $P(k)$ is defined by

$$\langle \zeta_{\mathbf{k}_1} \zeta_{\mathbf{k}_2} \rangle = P(k_1)(2\pi)^3 \delta(\mathbf{k}_1 + \mathbf{k}_2), \tag{22}$$

and after making the identification $k \sim L^{-1}$ we can conclude that $P(k) = \sigma^2(k)/k^3$. The bispectrum, $B(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)$, can be obtained by similar means. It satisfies

$$\langle \zeta_{\mathbf{k}_1} \zeta_{\mathbf{k}_2} \zeta_{\mathbf{k}_3} \rangle = B(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)(2\pi)^3 \delta(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3). \tag{23}$$

It is conventional to use momentum conservation to make B a function of the magnitudes $\{k_1, k_2, k_3\}$ alone. We write $\langle \zeta \zeta \zeta \rangle_L = \alpha(L)$. Then,

$$\langle \zeta_{\mathbf{x}} \zeta_{\mathbf{y}} \zeta_{\mathbf{z}} \rangle = \frac{\alpha(L)}{3}$$
$$\times \left\{ L_2^3 L_3^3 \delta(\mathbf{x} - \mathbf{y}) \delta(\mathbf{x} - \mathbf{z}) + L_1^3 L_3^3 \delta(\mathbf{y} - \mathbf{x}) \delta(\mathbf{y} - \mathbf{z}) + L_1^3 L_2^3 \delta(\mathbf{z} - \mathbf{x}) \delta(\mathbf{z} - \mathbf{y}) \right\}. \tag{24}$$

Comparing Eqs. (23) and (24), it follows that the bispectrum can be written

$$B = \frac{\alpha}{3} \frac{\sum_i k_i^3}{\prod_j k_j^3}, \tag{25}$$

which is the 'local' form typically generated by superhorizon evolution. This is a natural consequence of Eqs. (21)–(24), which require that correlations are local in real space.

To measure the amplitude of three-point correlations it is conventional to use the f_{NL} parameter, defined by

$$B(k_1, k_2, k_3) = \frac{6}{5} f_{\mathrm{NL}} \left\{ P(k_1)P(k_2) + P(k_1)P(k_3) + P(k_2)P(k_3) \right\}. \tag{26}$$

Eqs. (25) and (26) imply that f_{NL} satisfies

$$f_{\mathrm{NL}} = \frac{5}{18} \frac{\alpha(L)}{\sigma^4(L)} = \frac{5}{18} \frac{\langle \zeta \zeta \zeta \rangle_L}{\langle \zeta \zeta \rangle_L^2}. \tag{27}$$

This formula can be derived from many alternative constructions (see Ref. [32]).

3. Conversion of isocurvature fluctuations at the waterfall

Consider any region of field space described by a collection of light scalars ϕ^α and a waterfall field χ, evolving according to the action

$$S = -\frac{1}{2} \int \mathrm{d}^4 x \, \sqrt{-g} \Big\{ \partial \phi^\alpha \partial \phi_\alpha + (\partial \chi)^2 + 2V(\phi^\alpha, \chi) \Big\}, \tag{28}$$

where $(\partial \chi)^2 \equiv \partial^a \chi \partial_a \chi$ and lower-case Latin labels $\{a, b, \ldots\}$ index spacetime dimensions. The potential $V(\phi^\alpha, \chi)$ should be suitable for realising hybrid inflation but is otherwise arbitrary.

Near the point at which the waterfall field, χ, becomes tachyonic, many hybrid potentials can be well-approximated by

$$V = \hat{V}(\varphi, \ldots) - \lambda \varphi \chi^2 + \ldots. \tag{29}$$

where $\hat{V}$ involves only fields orthogonal to χ. The field φ is a single direction in field space which behaves as an order parameter controlling the onset of the waterfall, and λ is a coupling with the dimensions of mass. If the density perturbation is dominated by adiabatic fluctuations among these fields, then the spectral tilt, n_s, depends on $\hat{V}$ alone. It is given by

$$n_s - 1 = -6\epsilon + 2\eta, \tag{30}$$

where the slow-roll parameters are defined by

$$\epsilon \equiv \frac{M_{\mathrm{P}}^2}{2} \left(\frac{\hat{V}'}{\hat{V}} \right)^2 \qquad \text{and} \qquad \eta \equiv M_{\mathrm{P}}^2 \frac{\hat{V}''}{\hat{V}}. \tag{31}$$

A prime $'$ denotes a derivative with respect to the adiabatic direction φ. Linde's original formulation of hybrid inflation yielded a blue spectrum [1], which is now known to be incompatible with observation. In hybrid models ϵ is typically negligible, so acceptable $\hat{V}$ generally require η to be slightly negative, satisfying $\eta \approx -0.02$. Such models are sometimes described as "hill-top" potentials [23]. Potentials with a negative η complicate the higher-order terms necessary to guarantee a consistent model, but are conceptually no different to the original hybrid proposal. We will use an explicit example of such a potential in §3.2.

3.1. Analytic approximation

The dynamics of Eq. (29) can be complicated. Our analysis relies on certain simplifying approximations, and closely parallels that of Copeland *et al.* [21]. We ignore Hubble damping and assume that $\dot{\varphi}$ is practically constant. These will be acceptable during ejection from the ridge, provided that φ is rolling slowly and the trajectory is ejected over a timescale much shorter than a Hubble time. Ultimately, we will justify these approximations by checking our results numerically. For this purpose we use *both* the effective potential, Eq. (29), and global completions which describe the descent into a reheating minimum, to be discussed below.

It is convenient to work in Fourier space, where χ can be decomposed into a series of coupled oscillators $\chi_{\mathbf{k}}$. We may freely choose coordinates so that the effective potential for χ becomes tachyonic at $\varphi = 0$, and set $t = 0$ at that time. For $t > 0$ the evolution of each $\chi_{\mathbf{k}}$ is governed by a time-dependent equation of motion,

$$\ddot{\chi}_{\mathbf{k}} + (k^2 - \mu^3 t)\chi_{\mathbf{k}} = 0, \tag{32}$$

where the linear growth with t is a consequence of our assumption that $\dot{\varphi}$ is practically constant. The overdot indicates a derivative with respect to t, and μ has dimensions of mass,

$$\mu^3 \equiv 2\lambda\dot{\varphi}|_{t=0}. \tag{33}$$

The general solution is

$$\chi_{\mathbf{k}}(t) = A_{\mathbf{k}}\operatorname{Ai}(\mu t - k^2/\mu^2) + B_{\mathbf{k}}\operatorname{Bi}(\mu t - k^2/\mu^2), \tag{34}$$

where $A_{\mathbf{k}}$ and $B_{\mathbf{k}}$ are arbitrary coefficients, and $\operatorname{Ai}(x)$ and $\operatorname{Bi}(x)$ are Airy functions of the first and second kind, respectively. For $x < 0$ these functions are oscillatory, whereas for $x > 0$ the first Airy function decays and $\operatorname{Bi}(x)$ grows exponentially. This growth corresponds to an exponential rise in occupation number of the low-lying modes of χ. It describes a rapid reordering of the spatial field configuration, during which gradients may become large. We will return to this issue in §4.

For a mode of wavenumber k, exponential growth begins at the time $t_k \equiv k^2/\mu^3$. It follows that the phase transition proceeds by spinodal decomposition, beginning with the zero mode and extending to higher k-modes at successively later times. The coefficients $A_{\mathbf{k}}$ and $B_{\mathbf{k}}$ can be expressed in terms of $\chi_{\mathbf{k}}(t_k)$ and $\dot{\chi}_{\mathbf{k}}(t_k)$, giving

$$A_{\mathbf{k}} = \frac{3^{2/3}}{2}\Gamma(2/3)\chi_{\mathbf{k}}(t_k) - \frac{3^{1/3}}{2\mu}\Gamma(1/3)\dot{\chi}_{\mathbf{k}}(t_k) \tag{35}$$

$$B_{\mathbf{k}} = \frac{3^{1/6}}{2}\Gamma(2/3)\chi_{\mathbf{k}}(t_k) + \frac{3^{-1/6}}{2\mu}\Gamma(1/3)\dot{\chi}_{\mathbf{k}}(t_k). \tag{36}$$

Our interest lies in a description of the process by which spatially disjoint spacetime regions depart from the arête. It is at this point that our analysis departs from that of Copeland *et al.*, whose interest lay with the smallest possible scales. Consider an ensemble of spacetime volumes passing over the waterfall, smoothed on a lengthscale of order k^{-1}. The mutual scatter in field value between the members of this ensemble can be computed, for which it is a good approximation to include only the growing mode. We find

$$\sigma_\chi^2(k, t) = \frac{1}{2\pi^2}\int_0^k P_B(k')\operatorname{Bi}[\mu(t - t_{k'})]^2 k'^2 \, \mathrm{d}k', \tag{37}$$

where we have filtered out modes of wavelength shorter than $1/k$, and introduced the power spectrum of B_k,

$$\langle B_{\mathbf{k}}^* B_{\mathbf{k}}\rangle = P_B(k)(2\pi)^3\delta(\mathbf{k} - \mathbf{k}'). \tag{38}$$

Eq. (37) is valid when the relevant k-modes are already rolling, and hence applies provided $k \lesssim \sqrt{\mu^3 t}$. We have set an implicit infrared cutoff scale to zero. It transpires that $P_B(k)$ is a very blue function of k, so this is typically an excellent approximation.

Determining $P_B(k)$ requires information about the quantum fluctuations in χ.† Immediately prior to the tachyonic transition, χ carries a power spectrum imprinted by the preceeding epoch, during which its mass was large. It is a reasonable approximation to determine the fluctuations synthesized during this era by setting each $\mathbf{k}$-mode in a Bunch–Davies vacuum state, with large constant mass M. Well inside the horizon, each mode is normalized to the corresponding Minkowksi space oscillator. Once $k/a \lesssim M$ the mass term will dominate its evolution, and each mode decays as $\chi_{\mathbf{k}} \propto a^{-3/2}$ owing to Hubble damping. For super-horizon modes, where $k/a \ll H$, this leads to‡

$$\langle \chi_{\mathbf{k}}^* \chi_{\mathbf{k}'} \rangle = \frac{1}{2M} \left(\frac{H}{k_\star} \right)^3 (2\pi)^3 \delta(\mathbf{k} - \mathbf{k}')$$

$$\langle \dot{\chi}_{\mathbf{k}}^* \dot{\chi}_{\mathbf{k}'} \rangle = \frac{M}{2} \left(\frac{H}{k_\star} \right)^3 (2\pi)^3 \delta(\mathbf{k} - \mathbf{k}'), \tag{40}$$

where $k_\star$ is the comoving wavenumber of the mode leaving the horizon at the time of evaluation. These correlators scale like a^{-3}, because the wavenumber $k_\star$ scales like a as inflation proceeds. Employing (35) and (40) yields§

$$P_B(k) = \frac{H^3}{8Mk_\star^3} \left[3^{1/3}\Gamma(2/3)^2 + 3^{-1/3}\Gamma(1/3)^2 \left(\frac{M}{\mu} \right)^2 \right] = \frac{CH^3}{Mk_\star^3}, \tag{41}$$

where C aggregates the numerical constants. It is at least of order unity, and can be much larger if $\mu \ll M$.

For large x, the growing Airy function $\mathrm{Bi}(x)$ takes the asymptotic form

$$\mathrm{Bi}(x) \simeq \frac{1}{\sqrt{\pi}} x^{-1/4} \exp\left(\frac{2}{3} x^{3/2} \right). \tag{42}$$

At late times, this implies that the scatter among field values is well-approximated by

$$\sigma_\chi^2(k,t) \approx \frac{CH^3}{2\pi^3 Mk_\star^3} \int_0^k \frac{1}{\sqrt{\mu(t - t_{k'})}} \exp\left(\frac{4}{3} \left[\mu(t - t_{k'}) \right]^{3/2} \right) k'^2 \, \mathrm{d}k'. \tag{43}$$

---

† In their analysis, Copeland *et al.* noted that modes are light at the tachyonic transition and assumed that each mode was in a massless Bunch–Davis vacuum state at time t_k. In this state, the correlation functions satisfy

$$\langle \chi_{\mathbf{k}}^* \chi_{\mathbf{k}'} \rangle = \frac{1}{2k}(2\pi)^3 \delta(\mathbf{k} - \mathbf{k}')$$

$$\langle \dot{\chi}_{\mathbf{k}}^* \dot{\chi}_{\mathbf{k}'} \rangle = \frac{k}{2}(2\pi)^3 \delta(\mathbf{k} - \mathbf{k}'). \tag{39}$$

This would be a good choice whenever the effective mass M of χ varies adiabatically, or whenever $\dot{M}/M^2 \ll 1$. Very close to the tachyonic transition, $M^2 \sim t$ and evolution will not be adiabatic.

‡ See Appendix A for details.

§ We have evaluated the correlators $\langle \chi_{\mathbf{k}}^*(t_k)\chi_{\mathbf{k}'}(t_k) \rangle$ and $\langle \dot{\chi}_{\mathbf{k}}^*(t_k)\dot{\chi}_{\mathbf{k}'}(t_k) \rangle$ at a common time, $t = 0$, rather than t_k. The error we commit in this approximation is negligible, since the e-foldings of expansion between $t = 0$ and $t = t_k$ is of order $(k/k_*)^2 (H/\mu)^3 \ll 1$.

When $t \gg t_k$ the k-dependence becomes trivial: all modes behave collectively, evolving coherently with the background field. Moreover, the k-integral can be performed at once. In this limit, we find

$$\sigma_\chi^2(k, t) = \frac{CH^3 k^3}{6\pi^3 M k_\star^3} \frac{1}{\sqrt{\mu t}} \exp\left(\frac{4}{3}[\mu t]^{3/2}\right). \tag{44}$$

Eq. (44) implies that the scatter within an ensemble of spacetime volumes grows according to the equation of motion for the background field in a single volume,

$$\chi = \frac{\chi_*}{(\mu t)^{1/4}} \exp\left(\frac{2}{3}[\mu t]^{3/2}\right). \tag{45}$$

This justifies our application of the separate universe picture, but the preceding Fourier space analysis is required to determine the effective release point, $\chi_\star$, as a function of scale. This can be obtained from inspection of Eq. (44).

3.2. The statistics of N

Once the waterfall potential has become tachyonic, χ rolls down the ridge. Its velocity increases until the universe becomes dominated by the kinetic energy $\dot{\chi}^2/2$. When this phase of "kinetic domination" is achieved, comoving hypersurfaces are practically determined by hypersurfaces of constant kinetic energy, and ζ ceases to evolve. According to Eq. (45), the distribution of initial values, χ_*, will lead to a spread in arrival times at the kinetically dominated epoch. We write the perturbation in expansion history associated with this spread as $\zeta_{\rm w}$. It is this fluctuation which we propose must be included when calculating the properties of perturbations generated in a typical hybrid model. The following perturbative analysis will be trustworthy only if kinetic domination is reached. In §4 we discuss the possibility that a non-perturbative effect, *tachyonic preheating*, can quench the dispersion of trajectories before kinetic domination is attained.

The transit time to kinetic domination, $t(\chi_*)$, can be estimated by requiring $\dot{\chi} \sim H$. Taking H to be constant, Eq. (45) implies that $t(\chi_*)$ must solve

$$\ln \frac{H}{\chi_*} = -\frac{1}{4} \ln \left[\mu t(\chi_*)\right] + \ln \left(\frac{\mu}{M_{\rm P}} \sqrt{\mu t(\chi_*)} - \frac{1}{4t(\chi_*)M_{\rm P}}\right) + \frac{2}{3}\left[\mu t(\chi_*)\right]^{3/2}, \tag{46}$$

in which the Planck mass, $M_{\rm P}$, has temporarily been restored to exhibit the relative magnitude of each term. The asymptotic regime of late times corresponds to $\mu t \gg 1$, and is approximately reached when $N \sim H/\mu$ e-folds have elapsed since the waterfall transition. In this region $\ln(\mu t)$ can be neglected in comparison with $(\mu t)^{3/2}$. Under the same assumptions, $(tM_{\rm P})^{-1}$ can be neglected in the middle logarithm, after which the remaining term is of order $\ln \mu M_{\rm P}^{-1}$. This may be large if μ is much smaller than the Planck scale. On the left-hand side, however, $\ln H\chi_*^{-1}$ is typically of order $1.5 \ln(k_*/k) + 0.5 \ln M H^{-1} \sim 100$. To avoid fatal problems with overproduction of black holes and topological defects, discussed by García-Bellido *et al.* [20], the entire waterfall phase should complete within an e-fold. Since the scale for roll-down is set by

μ, this requires $\mu \gtrsim H$. Hence μ cannot be arbitrarily small, and $|\ln(\mu M_{\rm P}^{-1})|$ is likely to be much smaller than 100, and can be discarded in comparison with $\ln H\chi_*^{-1}$. The remaining terms imply that the transit time is given to a fair approximation by

$$t(\chi_*) = \frac{1}{\mu}\left(\frac{3}{2}\ln\frac{H}{\chi_*}\right)^{2/3}. \tag{47}$$

Note that the typical smallness of χ_*, which at first sight might cause one to imagine that any isocurvature fluctuation is tiny, has played an important role in ascertaining the rough validity of Eq. (47).

In virtue of our assumption that H is practically constant, Eq. (47) can be converted to an estimate of the transit time in e-folds, $N(\chi_*)$. The moments of N satisfy

$$\langle N \rangle = \frac{H}{\mu}\int_0^\infty P(\chi_*, \sigma_\chi)\left(\frac{3}{2}\ln\frac{H}{\chi_*}\right)^{2/3}\mathrm{d}\chi_*, \tag{48}$$

$$\langle N^2 \rangle = \left(\frac{H}{\mu}\right)^2\int_0^\infty P(\chi_*, \sigma_\chi)\left(\frac{3}{2}\ln\frac{H}{\chi_*}\right)^{4/3}\mathrm{d}\chi_*, \tag{49}$$

$$\langle N^3 \rangle = \left(\frac{H}{\mu}\right)^3\int_0^\infty P(\chi_*, \sigma_\chi)\left(\frac{3}{2}\ln\frac{H}{\chi_*}\right)^2\mathrm{d}\chi_*, \tag{50}$$

where $P(\chi_*, \sigma_\chi)$ is the distribution of initial values, χ_*. It is determined by the quantum mechanics of a massive scalar field evolving under Hubble damping. Therefore, we expect P to be close to Gaussian, characterized by its variance but with negligible higher moments.

In analogy with the inverted simple harmonic oscillator considered in §1, these expressions lead to the surprising result that the variance and skew of $\zeta_{\rm w}$ depend only weakly on the initial variance, σ_χ^2. Furthermore, even if the initial distribution $P(\chi_*, \sigma_\chi)$, is exactly Gaussian, a skew of order unity is generated. This behaviour persists over an exponentially large range of σ_χ^2. We have evaluated Eqs. (48)–(50) numerically, and used Eqs. (19)–(20) to determine statistics of the observational quantity, $\zeta_{\rm w}$. Our results are depicted in Fig. 1. To a good approximation, it is clear that $\langle\zeta_{\rm w}\zeta_{\rm w}\rangle = \sigma_{\rm w}^2 \sim \gamma(H/\mu)^2$, where γ varies slowly within the range $0.01 < \gamma < 0.1$ over an exponentially large range of σ_χ^2. Similar behaviour occurs for the skewness, defined by

$$\mathrm{Skew}(\zeta_{\rm w}) = \frac{\langle\zeta_{\rm w}\zeta_{\rm w}\zeta_{\rm w}\rangle}{\langle\zeta_{\rm w}\zeta_{\rm w}\rangle^{3/2}} \tag{51}$$

and displayed in Fig. 2. Over a similar exponentially large range of σ_χ it is clear that $\mathrm{Skew}(\zeta_{\rm w})$ is practically constant, taking values close to 1.5 for an Gaussian initial distribution of χ_*.

We have verified these results using numerical Monte Carlo simulations, which give good agreement with the semi-analytic results given above. We have carried out simulations for a simple example of a potential which agrees with Eq. (29) in the transition region, but which is bounded from below and describes the descent into a stable reheating minimum at zero energy. The potential we use is

$$V = V_0\left[\left(1 + \frac{\eta\phi^2}{4M_{\rm P}^2}\right) - \frac{\lambda\left(\phi^2 - \phi_0^2\right)\chi^2}{4V_0\phi_0}\right]^2 \tag{52}$$

$$\sim V_0 + \frac{\eta V_0}{2M_{\mathrm{P}}^2}\varphi^2 - \lambda\varphi\chi^2 + O(\varphi^3, \ldots) \tag{53}$$

where $\varphi = \phi - \phi_0$. Near the hybrid transition, Eq. (53) indicates that (52) has the approximate form of Eq. (29). In using Eq. (53) we are free to end the simulation at any time after kinetic domination is reached. For convenience we choose the time at which the field is close to minimum for the first time. This occurs just after kinetic domination. We evaluate N as a function of the initial values ϕ_* and χ_* by solving the full cosmological field equations. Initially ϕ is fixed so that the system is on the cusp of the hybrid instability. For the full potential, Eq. (52), this corresponds to choosing $\phi_* = \phi_0$. The initial value, χ_*, is drawn from a Gaussian distribution with variance σ_χ^2. This process is repeated $\sim 10^4$ times, for which N is evaluated at a fixed value of H. This procedure accurately provides values for $\langle N \rangle$, $\langle N^2 \rangle$ and $\langle N^3 \rangle$ evaluated on a comoving hypersurface. Eqs. (19)–(20) are employed to determine the statistics of ζ. Physically, this procedure is equivalent to allowing an ensemble of trajectories, with initial conditions characterized by a distribution of values $\chi = \chi_*$, to evolve under the influence of the hybrid potential. No approximations are made during the dynamical evolution. We include the effect of cosmological expansion on the evolution of ϕ and χ, and retain all contributions to the Hubble rate.

We have performed this proceedure for a wide range of parameter values. The results are always in good agreement with our semi-analytic formulae. We find that the variance σ_{w}^2 depends only weakly on the initial variance, σ_χ^2. More importantly, the skew is of order unity, irrespective of the value of σ_χ^2. Our conclusion that ζ_{w} carries a strongly non-Gaussian distribution appears to be robust. In Figs. 1–2, we present results obtained using the full potential (52) and the specific choices

$$\eta = -0.02, \quad \lambda = 10^{-5}, \quad V_0 = 10^{-11}, \quad \phi_0 = 1. \tag{54}$$

In this model $H/\mu \simeq 0.02$.

One may have harboured concerns the approximation of Eq. (32), in which $\dot{\varphi}$ is taken to be constant, could be a good approximation near the top of the arête, but a poor description of the waterfall evolution near kinetic domination, where the dominant contribution to ζ_{w} is synthesized. The good agreement shown in Figs. 1–2 between our exact Monte Carlo simulations and the semi-analytic formulae shows that Eq. (32) provides a good approximation throughout the evolution of the waterfall field.

3.3. Power spectrum and f_{NL}

According to the argument of §§3.1–3.2, we have identified a new source of curvature perturbations in hybrid-like scenarios. However, we cannot yet conclude that this perturbation is present in typical hybrid models because the foregoing analysis depends on the validity of the separate universe picture. Our numerical experiments show that the dominant part of ζ_{w} is generated close to the era of kinetic domination. In the Introduction (§1), we discussed the role of tachyonic preheating in quenching the dispersion of trajectories. This is associated with a failure of the separate universe
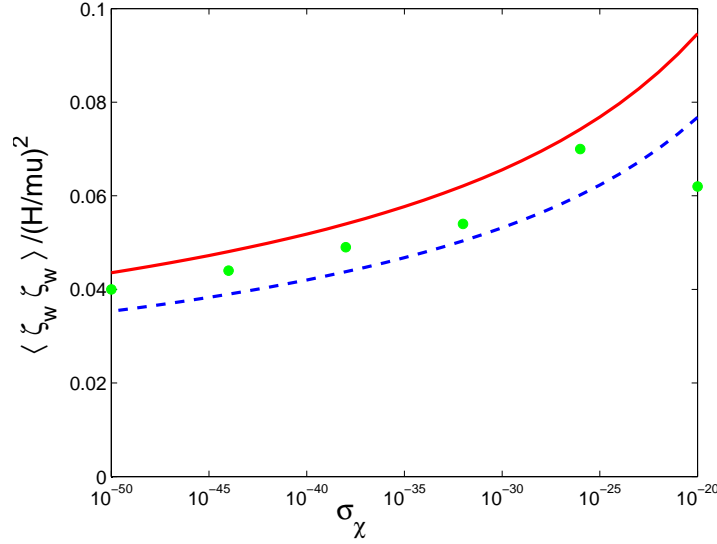
Figure 1. $\sigma_{\mathrm{w}}^2/(H/\mu)^2$ plotted as a function of the RMS initial fluctuation σ_χ in χ. The green dots are computed numerically for a Gaussian ensemble of trajectories, using the potential (52) and parameter values (54). The lines are semi-analytic estimates, obtained by numerically integrating (48-50). The solid red line uses a Gaussian distribution of initial conditions of variance σ_χ^2, while the dashed blue line employs a uniform distribution of varaince σ_χ^2. The final variance depends weakly on the form of the initial distribution, and varies extremely slowly over many decades of the initial variance.
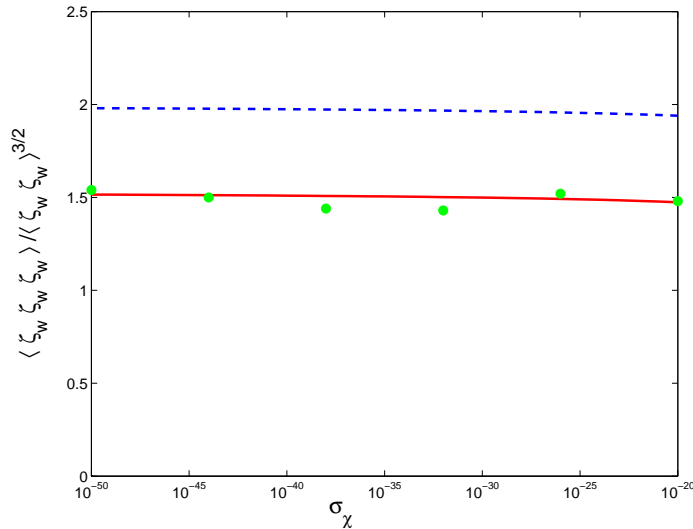

Figure 2. The skew of ζ_{w} plotted as a function of the RMS initial fluctuation σ_χ in χ. The green dots are computed numerically for a Gaussian ensemble of trajectories using the potential (52) and parameter values (54). The lines are semi-analytic estimates, obtained by numerically integrating (48-50). The solid red line uses a Gaussian distribution of initial conditions of variance σ_χ^2, while the dashed blue line employs a uniform distribution of variance σ_χ^2. The final skew depends weakly on the form of the initial distribution, and varies very little over many decades of variation in the initial variance.

picture. If tachyonic preheating completes before kinetic domination is attained on CMB scales, then the large-scale waterfall field acquires negligible dispersion. In this case the evolution of perturbations follows the traditional picture of Refs. [18, 20].‖ We study this possibility in detail in §4.

In the remainder of §3 we assume that the dynamics of some specific model are such that kinetic domination occurs before the onset of gradient instabilities associated with tachyonic preheating. In any such model ζ_{w} is uncorrelated with, and adds incoherently to, the curvature perturbation synthesized during the previous inflationary epoch. This is similar to the scenario studied by Boubekeur and Lyth [33], who considered the effect of an uncorrelated χ^2 perturbation which added incoherently to a scale-invariant inflationary contribution.

Does ζ_{w} unacceptably contaminate the primordial density perturbation in the hybrid scenario? The answer depends on the spectrum of ζ_{w}, and whether it makes a significant contribution on CMB scales. The spectrum of ζ_{w} is determined by the dependence of σ_{w}^2 on the moments which characterize the initial distribution of χ_*. The continuum power spectrum associated with ζ_{w} takes the form

$$\mathcal{P}_{\mathrm{w}}(k) = \frac{\mathrm{d}\sigma_{\mathrm{w}}^2}{\mathrm{d}\ln k} = \frac{\partial\sigma_{\mathrm{w}}^2}{\partial\sigma_{\chi}^2}\frac{\mathrm{d}\sigma_{\chi}^2}{\mathrm{d}\ln k} + \sum_{m\geqslant 3}\frac{\partial\sigma_{\mathrm{w}}^2}{\partial\mu_m}\frac{\mathrm{d}\mu_m}{\mathrm{d}\ln k}, \tag{55}$$

where μ_m is the mth moment of χ_*, σ_{χ}^2 is the initial variance of χ, and $\mathcal{P}$ denotes the dimensionless power spectrum, defined for any fluctuation by the rule $\mathcal{P}(k) = k^3 P(k)/2\pi^2$. Comparison of the Gaussian and uniform distributions in Fig. 1 (which have very different higher moments) gives reasonable evidence that $\mathrm{d}\sigma_{\mathrm{w}}^2/\mathrm{d}\sigma_{\chi}^2$ is roughly independent of the μ_m. There is no reason to believe that the $\mathrm{d}\sigma_{\mathrm{w}}^2/\mathrm{d}\mu_m$ are small in comparison, but if the initial distribution of χ is close to Gaussian then these moments can be neglected. As we argued above it seems reasonable to expect that $P(\chi_*)$ is approximately Gaussian, because it is determined by the fluctuations of very massive oscillators evolving under Hubble damping. Under these assumptions the power spectrum of ζ_{w} is approximately

$$\mathcal{P}_{\mathrm{w}}(k) \simeq \frac{\mathrm{d}\sigma_{\mathrm{w}}^2}{\mathrm{d}\sigma_{\chi}^2}\mathcal{P}_{\chi}(k). \tag{56}$$

In view of the steep, blue spectrum associated with χ_* one might worry that Eq. (56) implies $\mathcal{P}_{\mathrm{w}}$ is also very blue. This turns out not to be the case. A slowly-varying function such as $\sigma_{\mathrm{w}}^2(\sigma_{\chi}^2)$ can be well-approximated by a series in powers of $\ln(H/\sigma_{\chi})$, of the form

$$\sigma_{\mathrm{w}}^2 = A_0\left[\ln\left(\frac{H}{\sigma_{\chi}}\right)\right]^{-p} + A_1\left[\ln\left(\frac{H}{\sigma_{\chi}}\right)\right]^{-p-1} + \cdots + \mathrm{O}\left(\frac{\sigma_{\chi}}{H}\right). \tag{57}$$

The parameter p must be strictly positive, because σ_{w}^2 must be an increasing function of σ_{χ}^2. Subsequent terms in the series are subleading. One can show analytically that

‖ A significant non-Gaussian perturbation is always generated on the horizon scale, which may or may not be competitive with the perturbations studied by Randall *et al.* [18] and García-Bellido *et al.* [20]. We defer consideration of this effect until §5.

σ_{w}^2 takes precisely this form when the distribution for χ_* is uniform, with leading term equal to $(2/3)^{2/3}(H/\mu)^2 \left[\ln(H/\sigma_\chi)\right]^{-2/3}$. Focusing on the leading term, one obtains a power spectrum

$$\mathcal{P}_{\mathrm{w}}(k) = \frac{\mathrm{d}\ln\sigma_\chi^2}{\mathrm{d}\ln k}\frac{(pA_0)}{2}\left[\ln\left(\frac{H}{\sigma_\chi}\right)\right]^{-p-1} + \cdots . \tag{58}$$

Eq. (58) shows that $\mathcal{P}_{\mathrm{w}}(k)$ is proportional to the tilt of the power spectrum of $\mathcal{P}_\chi(k)$, rather then the input spectrum itself. Even if χ has an extremely steep power spectrum, provided its running is small, the power spectrum of ζ_{w} is approximately scale invariant. Had we retained subleading terms from (57), proportional to higher powers of $[\ln H/\sigma_\chi]^{-1}$, the same conclusion would have been obtained. Such terms may modify the normalization of $\mathcal{P}_{\mathrm{w}}$ but do not change its shape.

Our next step is to estimate f_{NL}. Our semi-analytic discussion and the numerical evidence shows that the skew of ζ_{w} is of order unity over a very wide range of σ_χ, and is essentially independent of k. It follows that

$$\langle\zeta_{\mathrm{w}}\zeta_{\mathrm{w}}\zeta_{\mathrm{w}}\rangle \sim \langle\zeta_{\mathrm{w}}\zeta_{\mathrm{w}}\rangle^{3/2}. \tag{59}$$

Assuming there are no other non-Gaussian contributions, f_{NL} is given by the ratio

$$f_{\mathrm{NL}} = \frac{\langle\zeta_{\mathrm{w}}\zeta_{\mathrm{w}}\zeta_{\mathrm{w}}\rangle}{\mathcal{P}_{\mathrm{total}}(k)^2} \sim \frac{\mathcal{P}_{\mathrm{w}}(k)^{3/2}}{\mathcal{P}_{\mathrm{total}}(k)^2}, \tag{60}$$

where $\mathcal{P}_{\mathrm{total}}$ is the total power spectrum of ζ, including the usual inflationary contributions as well as those from the waterfall process. The contributions from the waterfall need not dominate $\mathcal{P}_{\mathrm{total}}$. Fig. 1 shows that when σ_χ varies over ~ 30 orders of magnitude, σ_{w} varies only over a factor of roughly two. We will take a conservative point of view and estimate that the variance in ζ_{w} is given approximately by

$$\mathcal{P}_{\mathrm{w}} \sim \langle\zeta_{\mathrm{w}}\zeta_{\mathrm{w}}\rangle \sim \gamma\left(\frac{H}{\mu}\right)^2, \quad \text{with} \quad \gamma \sim 10^{-2}. \tag{61}$$

The dependence on H/μ follows from our semi-analytic arguments and is supported by our numerical evidence. The dimensionless factor $\gamma \sim 10^{-2}$ arises in the same way. From Fig. 1 it is evident that γ varies from about $0.03 - 0.1$ over many decades of σ_χ, so Eq. (61) is conservative. In practice, γ is slightly larger in all of the explicit examples we have studied. In what follows, we quote requirements on f_{NL} in terms of constraints on H/μ. At the end of this section we will relate these constraints to parameters of the inflationary model such as λ, V, and ϵ.

In models where tachyonic preheating does not quench the production of ζ_{w}, the power spectrum and f_{NL} put tight constraints on H/μ. Whatever the other predictions of a given model of hybrid inflation, the power spectrum of ζ must match observation. Therefore its value is fixed at $\mathcal{P}_{\mathrm{total}} \sim 3 \times 10^{-9}$. If $\mathcal{P}_{\mathrm{w}}$ makes a dominant contribution to $\mathcal{P}_{\mathrm{total}}$ the model is experimentally ruled out, because a correct normalization of the primordial power spectrum implies $f_{\mathrm{NL}} \sim \mathcal{P}_{\mathrm{total}}^{-1/2} \sim 2 \times 10^4$. This is in grave conflict with experiment [34], which requires $|f_{\mathrm{NL}}| \lesssim 100$. To avoid ζ_{w} dominating the power

spectrum, the estimate of Eq. (61) shows that we must have $\mathcal{P}_{\mathrm{w}} \lesssim \mathcal{P}_{\mathrm{total}}$, corresponding to $(H/\mu) \lesssim 5 \times 10^{-4}$.

In the regime where ζ_{w} does not dominate $\mathcal{P}_{\mathrm{total}}$, a slightly stronger constraint on H/μ arises from considering f_{NL}. Combining Eqs. (60) and (61), and neglecting the contribution of $\mathcal{P}_{\mathrm{w}}$ to $\mathcal{P}_{\mathrm{total}}$, we find

$$f_{\mathrm{NL}} \sim 10^{14} \left(\frac{H}{\mu} \right)^3 \tag{62}$$

Since current experimental constraints require $f_{\mathrm{NL}} \lesssim 10^2$, we obtain the constraint

$$\frac{H}{\mu} \lesssim 10^{-4}. \tag{63}$$

This is slightly stronger than the constraint from the power spectrum alone. If future constraints on f_{NL} improve to $|f_{\mathrm{NL}}| \lesssim 1$, and we take a less conservative estimate that $\mathcal{P}_{\mathrm{w}} \sim 10^{-1}(H/\mu)^2$, the resulting constraint on H/μ would be an order of magnitude tighter.

3.4. Constraints on hybrid model parameters

The constraints on H/μ turn out to be surprisingly restrictive when written in terms of the fundamental inflationary model parameters λ, V, and ϵ. Using the Friedmann equation, the slow-roll condition for ϕ, and the definitions (29) of λ and (33) of μ, we find

$$\frac{H}{\mu} = \left(\frac{V_{\mathrm{end}}}{6\sqrt{2} M_{\mathrm{P}}^3 \lambda \sqrt{\epsilon_{\mathrm{end}}}} \right)^{1/3} \tag{64}$$

Here, V_{end} and ϵ_{end} are the potential and slow-roll parameter evaluated at the end of inflation, which corresponds to the beginning of the hybrid transition. (Recall also that ϵ is defined using the derivative of V in the φ direction only, just before the hybrid transition). In the following, we will find it convenient to parametrize V_{end} as

$$V_{\mathrm{end}} = M_I^4 \tag{65}$$

where M_I is the inflationary mass scale, measured at the end of inflation. Since current cosmological scales have exited the horizon long before the end of inflation, there are no *a priori* constraints on M_I and ϵ_{end} from experiment.

We can get an idea of how strongly non-Gaussianity constrains the underlying hybrid inflationary model using a simple estimate, in which we make certain assumptions about the naturalness of the potential. On CMB scales the normalization of the scalar power spectrum constrains a combination of M_I and ϵ, specifically

$$\frac{V}{24\pi^2 M_{\mathrm{P}}^4 \epsilon} = \mathcal{P}_{\mathrm{total}} = 3 \times 10^{-9}. \tag{66}$$

Because H and ϵ are close to constant during a hybrid phase it would be unnatural for this ratio to change significantly by the end of inflation, although such behaviour could be introduced by a tuned or highly featured potential. Moreover, ϵ does not

approach unity towards the end of inflation, because there is no graceful exit in a hybrid model. Therefore, we generically expect (66) to hold, roughly, throughout the hybrid phase. (For example, it holds for the potential (52) for reasonable choices of parameters.) Accordingly, we will take $M_I \sim (3 \times 10^{-2}) \epsilon_{\text{end}}^{1/4} M_{\text{P}}$. This immediately implies that $f_{\text{NL}} \lesssim 100$ requires $M_I^2 \lesssim \lambda 10^{-8} M_{\text{P}}$. M_I sets the scale of the inflationary potential, and even if λ took a value close to the Planck mass this inequality would imply a rather low value for M_I. It is more reasonable to assume that λ would take a value at a similar scale to M_I — otherwise we would be forced to introduce a hierarchy between these mass parameters. Taking $\lambda \sim M_I$ implies the constraint

$$M_I \lesssim 10^{-8} M_{\text{P}}. \tag{67}$$

The f_{NL} parameter scales linearly with M_I, so if future experiments constrain $|f_{\text{NL}}| \lesssim 1$, this constraint tightens to $M_I \lesssim 10^{-10} M_{\text{P}}$. We conclude that unless the inflationary scale is very much lower than the Planck scale, unacceptable non-Gaussianity will be introduced.

The constraints we have derived here should be viewed as guidelines for reasonable hybrid inflationary models. If one has a particular example in mind, one should first compute the dimensionless factor γ appearing in (61) — although in practice we find this parameter varies very little between different potentials. One must then apply the definition of f_{NL} in (60) to find the corresponding constraint on H/μ. The formula (64) constrains the underlying inflationary parameters V, λ, and ϵ. The resuting constraints may be summarized by

$$\left(\frac{f_{\text{NL}}}{1.3 \times 10^4} \right) \sim \left(\frac{\gamma}{10^{-2}} \right)^{3/2} \left(\frac{M_I}{10^{-3} M_{\text{P}}} \right)^4 \left(\frac{10^{-2} M_{\text{P}}}{\lambda} \right) \left(\frac{10^{-2}}{\epsilon_{\text{end}}} \right)^{1/2}, \tag{68}$$

where we have assumed that H/μ is small enough that the waterfall makes a subdominant contribution to the ζ power spectrum. While there may be models for which Eq. (68) is effectively unconstraining — especially hybrid models with a very low inflationary scale and a correspondingly tiny ϵ — our analysis indicates that there are surprisingly strong constraints on a broad family of hybrid models. Moreover, these constraints can be expected to improve as future experimental results become available.

4. The validity of the separate universe picture

In this section we return to the question posed at the beginning of §3.3, namely whether the trajectories followed by adjacent cosmological patches can be prevented from dispersing sufficiently to generate an appreciable ζ_{w}. Felder *et al.* [25] and later Felder, Kofman & Linde [26] argued that a phenomenon known as *tachyonic preheating* typically occurs in hybrid models. Lattice simulations of the effect show a rapid destabilization of the background evolution owing to gradient instabilities formed as the universe reheats on small scales. This destabilization invalidates the separate universe assumption. To determine how much dispersion occurs, we must decide whether

CMB-scale regions of the universe reach kinetic domination before or after the onset of gradient instabilities.

Which are the relevant timescales? To avoid overproduction of topological defects and black holes it is already known that the hybrid transition should complete in of order one e-fold [20]. This constraint is independent of ζ_{w}. Accordingly the shortest wavelength modes, which leave the cosmological horizon immediately prior to the tachyonic transition, must reach kinetic domination after $\lesssim 1$ e-fold. Their effective release point $\chi_*^{(\mathrm{short})}$ can be found using (44) and (45) with $k = k_\star$, giving approximately

$$\chi_*^{(\mathrm{short})} = \sqrt{\frac{CH^3}{6\pi^3 M}}. \tag{69}$$

We recall that C is implicitly defined in Eq. (41), and is of the form $C \sim a + b(M/\mu)^2$, where a and b are roughly unity. Together with (47), Eq. (69) allows us to estimate the number of e-foldings, N_S, for the shortest wavelength modes to reach kinetic domination. It is

$$N_S \sim \frac{H}{\mu} \left(\ln \frac{6\pi^3 M}{HC} \right)^{2/3}. \tag{70}$$

The analysis of García–Bellido *et al.* [20] shows that the logarithm cannot be larger than $(H/\mu)^{-3/2}$. We also define N_L to be the number of e-foldings required for CMB scales to reach kinetic domination. These scales left the horizon N_* e-foldings before the end of inflation, and hence the CMB scale is roughly e^{N_*} larger than the horizon scale. In typical models N_* is of order $50 - 75$. The scaling of χ_* with $k/k_\star$ indicates that

$$\chi_*^{(\mathrm{long})} = e^{-3N_*/2} \chi_*^{(\mathrm{short})}. \tag{71}$$

Hence,

$$N_L \sim \frac{H}{\mu} \left(3N_* + \ln \frac{6\pi^3 M}{HC} \right)^{2/3}. \tag{72}$$

In typical models, the evolving field rolls into its reheating minimum almost immediately after kinetic domination. When it does so, Felder *et al.* [25, 26] argued that a rapid growth in occupation numbers of low-lying $\chi_{\mathbf{k}}$ modes would efficiently drain energy from the rolling field, so that preheating would complete within a single oscillation. On this basis, we should expect the separate universe picture to provide an accurate description up to $\sim$ few $\times N_S$ e-folds from the hybrid transition, but to fail at later times. A definitive determination of the range of validity of the separate universe picture will likely require a full lattice simulation of the waterfall phase. As a proxy for this unknown condition, we will assume that if $N_L/N_S \lesssim 10$ then there is a reasonable basis for belief that kinetic domination on CMB scales occurs sufficiently quickly for ζ_{w} to be synthesized.

There are two regimes in which the estimate of N_L/N_S is especially simple, depending on the relative size of the two parenthesized terms in (71). Although one might have expected that the exponential suppresion of CMB-scale waterfall fluctuations would imply $N_L \gg N_S$, we will find that in both regimes, there is good evidence that the

synthesis of $\zeta_{\rm w}$ can be successfully completed. The parity between N_L and N_S improves when H/μ is made small enough to satisfy observational constraints.

The first regime applies when $N_* \gtrsim \ln(6\pi^3 M/HC)$. This regime is relevant when no large hierarchies exist between M, H, and μ. In this case,

$$N_L - N_S \sim \frac{H}{\mu}\,(3N_*)^{2/3}, \qquad\qquad \text{for } N_* \gg \ln(6\pi^3 M/HC). \qquad (73)$$

Both N_L and N_S are much smaller than unity, since they are proportional to the very small parameter H/μ. In §3.3, we showed that consistency with constraints on $f_{\rm NL}$ requires $H/\mu \lesssim 10^{-4}$. For $N_* \sim 50 - 75$ this gives $N_L - N_S \lesssim 3 \times 10^{-3}$, showing that CMB-scale modes reach kinetic domination only a thousandth of an e-fold after short-scale modes. In this regime the ratio N_L/N_S is controlled by the hierarchy between N_* and the logarithmic factor, yielding

$$\frac{N_L}{N_S} \sim \left[1 + \frac{3N_*}{\ln(6\pi^3 M/HC)}\right]^{2/3} \qquad\qquad \text{for } N_* \gg \ln(6\pi^3 M/HC). \qquad (74)$$

Eq. (70) shows that the logarithm may be of order N_* without spoiling the requirement $N_S \lesssim 1$. In the typical models studied in §3 we have found that N_L/N_S is generally of order $5 - 7$. This analytic estimate is confirmed by our detailed Monte Carlo simulations. Pending detailed lattice calculations, there seems no reason to believe that perturbation theory should fail long before kinetic domination occurs on CMB scales. We also note that, while inflationary initial conditions imply that arguments based on causality give limited information, it may happen that the transition from an inflationary to a radiation-dominated equation of state takes some time to propagate from short to long scales. If so, it may be possible to tolerate larger values of N_L/N_S than might seem reasonable on the basis of Refs. [25, 26].

Even if a typical member of a bundle of CMB-scale trajectories does not reach kinetic domination, a small subset of trajectories in the bundle can be expected to do so. In this case the final curvature perturbation is likely to depend in detail on the tails of the distribution of χ_*, and cannot be calculated in perturbation theory. We leave the exact details of this regime for future work.

The second regime corresponds to $N_* \ll \ln(6\pi^3 M/HC)$, which requires large hierarchies between M, H, and μ. In this regime, N_L and N_S are very nearly equal, with

$$N_L \sim N_S \sim \frac{H}{\mu}\left(\ln\frac{6\pi^3 M}{HC}\right)^{2/3} \qquad\qquad \text{for } N_* \ll \ln(6\pi^3 M/HC). \qquad (75)$$

It follows that N_S must be exceedingly small, since

$$\frac{N_L - N_S}{N_S} \sim \frac{2N_\star}{\ln(6\pi^3 M/HC)} \qquad\qquad \text{for } N_* \ll \ln(6\pi^3 M/HC). \qquad (76)$$

Then, because $N_\star/\ln(6\pi^3 M/HC) \ll 1$ by assumption, the difference $N_L - N_S$ is itself much smaller than N_S. In this regime there seems a very good basis for belief that the separate universe assumption is not invalidated. Note that, however large the logarithm becomes, the requirement $N_S \lesssim 1$ can be satisfied by choosing a correspondingly small H/μ.

5. Conclusions

We have argued that the transition ending hybrid inflation can generate a large contribution to the curvature perturbation. This contribution arises from fluctuations in the "waterfall" field, χ, which are created as modes exit the inflationary horizon. These perturbations are subsequently suppressed by the cosmological expansion. Although these fluctuations are very small when inflation ends, they are amplified by the tachyonic instability in χ and can provide a significant contribution to the curvature perturbation. To calculate the enhancement we use a gradient expansion in which we treat the moments of the initial distribution non-perturbatively. This method is similar to the method of "moment transport" which has recently been used to calculate the curvature perturbation in inflationary models [32]. Here, the procedure is simplified owing to our ability to calculate an approximate analytic solution. This technique predicts that the fluctuations generated by the waterfall are highly non-Gaussian, and can easily exceed observational constraints. If present in a given model, these effects lead to new constraints on viable hybrid models together with a unique observational signature.

There are two phenomena which enable the tiny fluctuations in χ to generate significant effects. First, the variance in δN following completion of the waterfall depends extremely weakly on the variance in χ before the waterfall begins. While surprising, this phenomenon is sufficiently general that it can be illustrated using an inverted simple harmonic oscillator. Similar behaviour is seen in the semi-analytic and numerical examples given above. Therefore, even very tiny fluctuations in χ can lead to large fluctuations in δN. This means that the exponential suppression of χ-fluctuations during inflation does not guarantee that these fluctuations are harmless when hybrid inflation completes. In a single-field model, this kind of amplification would be forbidden because the curvature perturbation ζ would be conserved on super-horizon scales. Since hybrid inflation necessarily involves multiple fields, the presence of isocurvature modes means that ζ is no longer conserved. As we have argued above, the hybrid waterfall can cause it to grow significantly.

Second, the spectrum of fluctuations in δN, and hence the spectrum of the curvature perturbation ζ, is nearly scale-invariant. Naively, one might have expected that any process operating at the end of inflation could only affect modes of wavelength smaller than the horizon. This intuition is borne out for some processes: for example, the production of black holes or defects at the end of hybrid inflation results in a very blue power spectrum, which would be unimportant on cosmological scales. In the present case, it is true that χ has a strongly blue power spectrum immediately prior to the hybrid transition. This is because fluctuations in χ are suppressed during inflation. However, the waterfall process ensures that the final fluctution in δN is nearly independent of that in χ. Therefore, the *final* curvature fluctuation is nearly the same, independently of scale, even though it is seeded by fluctuations which, on long wavelengths, are much smaller. Thus, the hybrid transition can reprocess the blue spectrum of χ into a nearly scale-invariant one. The key difference in behaviour between large and small scales is

the transit time from the onset of the waterfall to kinetic domination.

The analysis we have presented here depends crucially on the validity of the separate universe picture while CMB scales are being ejected from the arête. On small scales, it is known that the hybrid transition creates large (nonlinear) gradients in the scalar fields. This rapid growth in gradients is responsible for the "tachyonic preheating" effect studied by Felder *et al.* [25, 26]. It is possible that short wavelength (horizon-size) perturbations reach this nonlinear regime very rapidly, before significant curvature perturbations are generated on longer wavelength (cosmological-size) scales. If so, this would prevent the dispersion of trajectories and lead to effects which are much weaker than those estimated here. An estimate of the onset of this gradient instability can be obtained by studying the time taken for different scales to reach the kinetic-dominated regime, at which time the curvature pertubation ζ ceases to evolve. Our estimates indicate that long and short wavelength perturbations reach the kinetic-dominated regime at nearly the same time, within a tiny fraction of an e-folding from the onset of the hybrid transition. The hierarchy between the transit time for long and short wavelengths is typically a number of order unity. In this regime a subtle interplay exists between effects on short and long length scales, and the condition for validity of the separate universe picture is not known. We expect that a definitive resolution of these questions requires numerical simulations of the hybrid transition. We conclude that perturbation theory provides no compelling reason to believe that waterfall contributions to the curvature perturbation should be strongly suppressed.

Even if tachyonic preheating prevents the generation of a significant curvature perturbation on CMB scales, it seems impossible to prevent the appearance of strong non-linearities on the shortest scales — those associated with the horizon scale at the transition. If these fluctuations are large, they will collapse to form black holes or topological defects. The analysis we have presented indicates that the formation of these non-perturbative objects should be taken to proceed from very non-Gaussian initial conditions. It is known that the final mass fraction contained in such collapsed objects can depend on the detailed profile of initial fluctuations [35, 36, 37], which potentially provides another means to constrain models containing a waterfall transition.

Acknowledgments

Appendix A. Evolution of a massive scalar field in de Sitter space

Here we give more details regarding the calculation of the power spectrum of χ before the tachyonic transition. This is used to determine $\sigma_\chi^2(k)$, the variance of χ smoothed on a real space length scale of $2\pi/k$.

We assume that χ has the effective Lagrangian

$$S = \int \left[-\frac{1}{2}(\partial\chi)^2 + \frac{1}{2}M^2\chi^2 \right] \sqrt{-g}\, \mathrm{d}^4 x \tag{A.1}$$

and take m to be constant. We work in conformal time and assume that the spacetime is flat de Sitter space, and hence that

$$\mathrm{d}s^2 = \frac{-\mathrm{d}\eta^2 + \mathrm{d}x_3^2}{(H\eta)^2} \tag{A.2}$$

where H is the Hubble parameter. As is conventional, we take η to increase along $(-\infty, 0)$ as proper time increases and the universe expands. We will find it convenient to use a rescaled field u, defined by

$$u(t, \vec{x}) = (-H\eta)\chi(t, \vec{x}) \tag{A.3}$$

Substituting this definition into the action for χ, and decomposing u into its Fourier modes, we find that each Fourier mode possesses the effective action

$$S = \int \frac{1}{2}(u_k')^2 - \frac{1}{2}\left(k^2 - \frac{2 - (M/H)^2}{\eta^2} \right) u_k^2 \, \mathrm{d}t \tag{A.4}$$

where $' = d/d\eta$. The rescaled field u has a canonically normalized kinetic term in this action. This makes it simple to impose the usual Minkowski space boundary conditions well inside the cosmological horizon. The subtlety is that well outside the horizon the mode functions are very different from their Minkowskian form.

The equation of motion resulting from (A.4) is

$$u'' + \left(k^2 - \frac{2 - (M/H)^2}{\eta^2} \right) = 0 \tag{A.5}$$

with solution

$$u_k(\eta) = c_1 \sqrt{-k\eta} H_\nu^{(1)}(-k\eta) + c_2 \sqrt{-k\eta} H_\nu^{(2)}(-k\eta) \tag{A.6}$$

where $H_\nu^{(1,2)}$ are Hankel functions and

$$\nu = \sqrt{\frac{9}{4} - \left(\frac{M}{H} \right)^2} \tag{A.7}$$

For applications to hybrid inflation, we require that the waterfall field is very massive, and in particular that $M \gg H$. Hence the index of the Hankel function is purely imaginary. We assume that we work in this limit throughout this derivation and hence

take $\nu = i(M/H)$. We now canonically normalize the mode functions by taking the limit $-k\eta \to 0$ and using the small argument expansion†

$$H_\nu^{(1,2)}(x) \to \sqrt{\frac{2}{\pi x}} \exp\left(\pm ix \mp \frac{i\pi\nu}{2} \mp \frac{\pi}{4}\right) \qquad x \gg |\nu| \gg 1 \qquad \text{(A.8)}$$

To canonically normalize the fields, we should have the positive frequency mode asymptote to $u_k(x) \to (2k)^{-1/2}e^{-ix}$ where $x = -k\eta$. Keeping in mind that μ is purely imaginary, this implies that the correctly normalized positive frequency mode is described by

$$c_1 = \frac{e^{-\pi|\nu|/2}}{2}\sqrt{\frac{\pi}{k}}, \qquad c_2 = 0 \qquad \text{(A.9)}$$

To find the super-horizon limit, we take $x \ll 1$ and use

$$H_\nu^{(1)}(x) = -\frac{i}{\pi}\left[\Gamma(-\nu)e^{-i\pi\nu}\left(\frac{x}{2}\right)^\nu + \Gamma(\nu)\left(\frac{x}{2}\right)^{-\nu}\right] \qquad x \ll 1 \qquad \text{(A.10)}$$

Starting with the Euler reflection formula

$$\Gamma(-\nu)\Gamma(\nu) = -\frac{\pi}{\nu\sin(\pi\nu)} \qquad \text{(A.11)}$$

and using the fact that $\Gamma(\nu)^* = \Gamma(\nu^*)$, the definition of χ in terms of u, and the assumption that $\nu \gg 1$, the super horizon limit of χ becomes

$$\chi_k \to \sqrt{\frac{H^3}{k^3 M}}(-k\eta)^{(3/2)+i(M/H)+i\theta} \qquad k\eta \gg 1, \quad M/H \gg 1. \qquad \text{(A.12)}$$

where θ is an irrelevant constant phase angle. Using this expression, one can easily derive the correlators (40) of χ_k on super-horizon scales.

References

[1] A. D. Linde, *Axions in inflationary cosmology*, Phys. Lett. **B259** (1991) 38–47.

[2] A. D. Linde, *Hybrid inflation*, Phys. Rev. **D49** (1994) 748–754, [arXiv:astro-ph/9307002].

[3] E. J. Copeland, A. R. Liddle, D. H. Lyth, E. D. Stewart, and D. Wands, *False vacuum inflation with Einstein gravity*, Phys. Rev. **D49** (1994) 6410–6433, [arXiv:astro-ph/9401011].

[4] J. M. Maldacena, *Non-Gaussian features of primordial fluctuations in single field inflationary models*, JHEP **05** (2003) 013, [arXiv:astro-ph/0210603].

[5] D. H. Lyth and Y. Rodríguez, *The inflationary prediction for primordial non- gaussianity*, Phys. Rev. Lett. **95** (2005) 121302, [arXiv:astro-ph/0504045].

[6] D. Seery and J. E. Lidsey, *Primordial non-gaussianities from multiple-field inflation*, JCAP **0509** (2005) 011, [arXiv:astro-ph/0506056].

[7] K. Koyama and D. Wands, *Ekpyrotic collapse with multiple fields*, JCAP **0704** (2007) 008, [arXiv:hep-th/0703040].

[8] K. Koyama, S. Mizuno, and D. Wands, *Curvature perturbations from ekpyrotic collapse with multiple fields*, Class. Quant. Grav. **24** (2007) 3919–3932, [arXiv:0704.1152].

† In cosmological applications, usually one takes the "sub-horizon limit" of these Hankel functions $H(x)$ to correspond to $x \gg 1$. In the present case of very massive fields with $M/H \gg 1$, the limit we require is actually that the proper field momentum is much larger than the inverse Compton wavelength, which corresponds to $x \gg M/H = |\nu|$. It is actually in the limit $x \gg |\nu|$ that the Hankel functions can be approximated by an expression such as (A.8).

[9] E. I. Buchbinder, J. Khoury, and B. A. Ovrut, *Non-Gaussianities in New Ekpyrotic Cosmology*, Phys. Rev. Lett. **100** (2008) 171302, [arXiv:0710.5172].

[10] K. Koyama, S. Mizuno, F. Vernizzi, and D. Wands, *Non-Gaussianities from ekpyrotic collapse with multiple fields*, JCAP **0711** (2007) 024, [arXiv:0708.4321].

[11] J.-L. Lehners and P. J. Steinhardt, *Non-Gaussian Density Fluctuations from Entropically Generated Curvature Perturbations in Ekpyrotic Models*, Phys. Rev. **D77** (2008) 063533, [arXiv:0712.3779].

[12] J.-L. Lehners and P. J. Steinhardt, *Intuitive understanding of non-gaussianity in ekpyrotic and cyclic models*, Phys. Rev. **D78** (2008) 023506, [arXiv:0804.1293].

[13] L. Alabidi, *Non-gaussianity for a two component hybrid model of inflation*, JCAP **0610** (2006) 015, [arXiv:astro-ph/0604611].

[14] C. T. Byrnes, K.-Y. Choi, and L. M. H. Hall, *Conditions for large non-Gaussianity in two-field slow- roll inflation*, JCAP **0810** (2008) 008, [arXiv:0807.1101].

[15] C. T. Byrnes, K.-Y. Choi, and L. M. H. Hall, *Large non-Gaussianity from two-component hybrid inflation*, JCAP **0902** (2009) 017, [arXiv:0812.0807].

[16] H. R. S. Cogollo, Y. Rodríguez, and C. A. Valenzuela-Toledo, *On the Issue of the ζ Series Convergence and Loop Corrections in the Generation of Observable Primordial Non-Gaussianity in Slow-Roll Inflation. Part I: the Bispectrum*, JCAP **0808** (2008) 029, [arXiv:0806.1546].

[17] Y. Rodríguez and C. A. Valenzuela-Toledo, *On the Issue of the ζ Series Convergence and Loop Corrections in the Generation of Observable Primordial Non- Gaussianity in Slow-Roll Inflation. Part II: the Trispectrum*, arXiv:0811.4092.

[18] L. Randall, M. Soljačić, and A. H. Guth, *Supernatural Inflation: Inflation from Supersymmetry with No (Very) Small Parameters*, Nucl. Phys. **B472** (1996) 377–408, [arXiv:hep-ph/9512439].

[19] J. García-Bellido and D. Wands, *The spectrum of curvature perturbations from hybrid inflation*, Phys. Rev. **D54** (1996) 7181–7185, [arXiv:astro-ph/9606047].

[20] J. García-Bellido, A. D. Linde, and D. Wands, *Density perturbations and black hole formation in hybrid inflation*, Phys. Rev. **D54** (1996) 6040–6058, [arXiv:astro-ph/9605094].

[21] E. J. Copeland, S. Pascoli, and A. Rajantie, *Dynamics of tachyonic preheating after hybrid inflation*, Phys. Rev. **D65** (2002) 103517, [arXiv:hep-ph/0202031].

[22] B. A. Ovrut and P. J. Steinhardt, *Inflationary cosmology and the mass hierarchy in locally supersymmetric theories*, Phys. Rev. Lett. **53** (1984) 732.

[23] L. Boubekeur and D. H. Lyth, *Hilltop inflation*, JCAP **0507** (2005) 010, [arXiv:hep-ph/0502047].

[24] **WMAP** Collaboration, E. Komatsu *et al.*, *Five-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation*, arXiv:0803.0547.

[25] G. N. Felder *et al.*, *Dynamics of symmetry breaking and tachyonic preheating*, Phys. Rev. Lett. **87** (2001) 011601, [arXiv:hep-ph/0012142].

[26] G. N. Felder, L. Kofman, and A. D. Linde, *Tachyonic instability and dynamics of spontaneous symmetry breaking*, Phys. Rev. **D64** (2001) 123517, [arXiv:hep-th/0106179].

[27] J. García-Bellido, M. Garcia Perez, and A. Gonzalez-Arroyo, *Symmetry breaking and false vacuum decay after hybrid inflation*, Phys. Rev. **D67** (2003) 103501, [arXiv:hep-ph/0208228].

[28] A. A. Starobinsky, *Multicomponent de Sitter (Inflationary) Stages and the Generation of Perturbations*, JETP Lett. **42** (1985) 152–155.

[29] M. Sasaki and E. D. Stewart, *A General analytic formula for the spectral index of the density perturbations produced during inflation*, Prog. Theor. Phys. **95** (1996) 71–78, [arXiv:astro-ph/9507001].

[30] I. Zaballa, Y. Rodríguez, and D. H. Lyth, *Higher order contributions to the primordial non-gaussianity*, JCAP **0606** (2006) 013, [arXiv:astro-ph/0603534].

[31] D. H. Lyth and D. Seery, *Classicality of the primordial perturbations*, Phys. Lett. **B662** (2008) 309–313, [arXiv:astro-ph/0607647].

[32] D. J. Mulryne, D. Seery, and D. Wesley, *Moment transport equations for non-Gaussianity,*

arXiv:0909.2256.

[33] L. Boubekeur and D. H. Lyth, *Detecting a small perturbation through its non-Gaussianity*, *Phys. Rev.* **D73** (2006) 021301, [arXiv:astro-ph/0504046].

[34] K. M. Smith, L. Senatore, and M. Zaldarriaga, *Optimal limits on $f_{\mathrm{NL}}^{\mathrm{local}}$ from WMAP 5-year data*, *JCAP* **0909** (2009) 006, [arXiv:0901.2572].

[35] J. S. Bullock and J. R. Primack, *Non-Gaussian fluctuations and primordial black holes from inflation*, *Phys. Rev.* **D55** (1997) 7423–7439, [arXiv:astro-ph/9611106].

[36] P. Ivanov, *Non-linear metric perturbations and production of primordial black holes*, *Phys. Rev.* **D57** (1998) 7145–7154, [arXiv:astro-ph/9708224].

[37] J. C. Hidalgo, *The effect of non-Gaussian curvature perturbations on the formation of primordial black holes*, arXiv:0708.3875.