Number counts and non-Gaussianity

Sarah Shandera,[1,*] Adrienne L. Erickcek,[2,3,†] Pat Scott,[4,‡] and Jhon Yana Galarza[3,5,§]

[1]*Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, Pennsylvania 16802, USA*

[2]*Canadian Institute for Theoretical Astrophysics, University of Toronto, 60 Street George Street, Toronto, Ontario M5S 3H8, Canada*

[3]*Perimeter Institute for Theoretical Physics, 31 Caroline Street N, Waterloo, Ontario N2L 2Y5, Canada*

[4]*Department of Physics, McGill University, Montréal QC H2W2L8, Canada*

[5]*Seminario Permanente de Astronomía y Ciencias Espaciales, Facultad de Ciencias Físicas,*
Universidad Nacional Mayor de San Marcos, Avenida Venezuela s/n, Lima 15081, Peru

We describe a general procedure for using number counts of any object to constrain the probability distribution of the primordial fluctuations, allowing for generic weak non-Gaussianity. We apply this procedure to use limits on the abundance of primordial black holes and dark matter ultracompact minihalos to characterize the allowed statistics of primordial fluctuations on very small scales. We present constraints on the power spectrum and the amplitude of the skewness for two different families of non-Gaussian distributions, distinguished by the relative importance of higher moments. Although primordial black holes probe the smallest scales, ultracompact minihalos provide significantly stronger constraints on the power spectrum and so are more likely to eventually provide small-scale constraints on non-Gaussianity.

PACS numbers: 98.80.Cq, 95.35.+d, 98.80.−k

I. INTRODUCTION

The power spectrum of primordial fluctuations is well measured on cosmological scales using the cosmic microwave background (CMB) and large-scale structure [1–6]. These observations provide compelling evidence that the fluctuations originated in an era of inflation. However, there is a great deal of new information waiting to be accessed in small-scale fluctuations and in higher-order statistics (non-Gaussianity). In this paper we explore both new regimes by examining how object number counts constrain the probability distribution of primordial fluctuations. By looking at very small objects (primordial black holes and ultracompact minihalos), we constrain the power in fluctuations on smaller scales than the CMB and large-scale structure currently probe, $k \gtrsim 3$ Mpc^{-1}. Since these objects are also very rare, we constrain the level of non-Gaussianity by limiting the abundance of extreme primordial overdensities.

Counts of rare objects are a useful probe of the primordial inhomogeneities in the gravitational field and their evolution: different objects probe different scales, different cosmological eras, and different particle physics and astrophysics. Primordial black holes (PBHs) have been used as a probe of small-scale power for many years [7]. Their abundance limits the allowed fluctuation power on very small scales, but the constraint is weak compared to the numbers expected from extrapolating a scale-invariant spectrum down from CMB scales [8,9]. More recently, a

*shandera@gravity.psu.edu
†erickcek@cita.utoronto.ca
‡patscott@physics.mcgill.ca
§yanajho@gmail.com

much stronger constraint on small-scale power was obtained by considering ultracompact minihalos (UCMHs) of dark matter [10–12]. The only drawback of this approach is that the strongest limits [11] require dark matter to annihilate into Standard Model (SM) particles, which can then be sought using standard indirect dark matter detection techniques [13]. Furthermore, UCMH constraints on the small-scale power obtainable by microlensing [12] are still stronger than PBH constraints, and apply even for nonannihilating dark matter.

The abundance of PBHs was also considered early on as a probe of non-Gaussianity [14]. In a non-Gaussian distribution, the number of rare objects is different from the Gaussian expectation, and rarer objects are typically sensitive to higher moments of the distribution. Number counts pick up *any* deviation from Gaussianity in a model-independent way and provide complementary constraints to those on the individual correlation functions, such as the shape of the three-point function in momentum space (the bispectrum). However, there are important limitations on what can be learned from number counts on both the analytic and observational sides. First, in general, one only knows the non-Gaussian probability density function (PDF) approximately, usually in terms of a few of the lowest-order moments. This limits how far out onto its tail the distribution is known, and so limits the utility of looking at very rare objects to constrain particular models. Furthermore, the best controlled approximations to weakly non-Gaussian PDFs are asymptotic expansions, so for a given level of non-Gaussianity, one trusts the expansion only so far out onto the tail. The greater the level of non-Gaussianity, the more limited the range of utility of the asymptotic expansion. On the observational side, one must know the amplitude of fluctuations (the variance of the distribution) very precisely

in order to find sharp constraints on the level of non-Gaussianity. The relationship between mass and the observable signature of the objects must be also be known accurately and precisely. Finally, extremely rare objects are not necessarily expected to be present in any finite sample (e.g. in any single survey or even our Universe), which makes drawing conclusions from their absence, or from very small number statistics, difficult. Nevertheless, constraints on the abundance of rare objects currently provide the only probe of the primordial fluctuations on very small scales, $k \gtrsim 10^4$ Mpc^{-1}. On larger scales, 50 Mpc$^{-1} \lesssim k \lesssim 10^4$ Mpc^{-1}, μ-type spectral distortions of the CMB also provide constraints [15–17]. We will return to the complementarity of these probes in the Conclusions.

Here we treat the analytic difficulties mentioned above carefully and develop a general prescription to use number counts to constrain weakly non-Gaussian primordial fluctuations. We apply this technique to derive bounds on non-Gaussianities with UCMHs and PBHs. Previous authors have used PBHs to constrain particular cases of the non-Gaussian PDF or particular inflation scenarios. Bullock and Primack [14], followed by Green and Liddle [18] and Ivanov [19], considered potentials with a localized feature that can both boost power dramatically on some scales and generate strong non-Gaussianity because of mode coupling near the sharp feature. Pina Avelino [20] worked out constraints on a χ^2 distribution for the fluctuations. Hidalgo [21] considered a non-Gaussian PDF from only the first term in the Edgeworth expansion [p_1 from Eq. (13)]. Klimai and Bugaev [22] constrained two-field inflation models. Most recently, Byrnes *et al.* [23] considered constraints on weak or strong non-Gaussianity of the "local" type [24], where the primordial curvature is $\mathcal{R}(x) = \mathcal{R}_G(x) + \frac{3}{5} f_{NL} \mathcal{R}_G^2(x)$, where $\mathcal{R}_G(x)$ is a Gaussian random field. Young and Byrnes [25] and Kohri *et al.* [26] went on to also consider non-Gaussianity in the curvaton scenario.

Our results generalize all previous work on weakly non-Gaussian models, demonstrating how constraints can be obtained from any object abundance and for any weakly non-Gaussian distribution. We use an extended Press-Schechter [27] approach where the abundance of an object is determined by the overdensity required for its formation ($\delta_c \equiv \delta \rho_c / \bar{\rho}$), the mass variance smoothed over the region that forms the object (σ_R), and similarly smoothed higher-order moments of the density fluctuations. These variables can be straightforwardly computed for *any* object: one need only compute the collapse threshold and smoothed statistics (by propagating the primordial fluctuations forward to the appropriate time of formation using transfer functions and a growth factor). This method is a straightforward extension of previous work on an analytic, weakly non-Gaussian mass function for galaxy clusters [28]. However, here we provide more precise criteria for where the analytic expansions are reliable in terms of the rareness of the object (the threshold $\nu_c \equiv \delta_c / \sigma_R$) and the level of

non-Gaussianity. In addition, we consider a wider range of non-Gaussian scenarios by parametrizing deviations from Gaussianity with both the value of the dimensionless skewness ($\mathcal{M}_3$) and a rule for how the higher moments scale with respect to the skewness. Most previous work assumes one particular choice, but physical mechanisms generating the primordial fluctuations do not all follow that pattern, and the distinction is important for interpreting constraints. By demonstrating the computational limits of the Edgeworth expansion, we also clarify why strongly non-Gaussian models must be evaluated on a case-by-case basis. Although we focus on small-scale constraints from PBHs and UCMHS in this work, the techniques we develop are equally applicable to limits on the abundances of high-redshift clusters, which provide strong constraints on non-Gaussianity on larger scales [29].

Section II develops our formalism for describing weakly non-Gaussian distributions and determining the abundance of objects when the primordial fluctuations are described by a non-Gaussian distribution. Sections III and IV apply this technique to PBHs and UCMHs, respectively, culminating with our final constraint plots in Fig. 5 (which a casual reader can skip to directly). We conclude by comparing current and future constraints on small-scale power and non-Gaussianity on a range of scales from a variety of existing and proposed techniques.

II. CALCULATING NUMBER COUNTS USING WEAKLY NON-GAUSSIAN DISTRIBUTIONS

A. Non-Gaussian probability distributions

Assuming that objects form from regions of the density field above some threshold, we need a probability distribution P of perturbation amplitudes δ_R and a threshold δ_c above which to count the objects. Here $\delta_{c,X}$ gives the minimum density contrast required for a perturbation to later collapse into an object of type $X \in \{\text{PBHs, UCMHs}\}$; $\delta_{c,\text{PBH}} \sim 1/3$ [7,30], while $\delta_{c,\text{UCMH}}$ is a scale-dependent quantity [11], although it is typically $\mathcal{O}(10^{-2}\text{–}10^{-3})$ [11,31]. Assuming a flat universe, the fractional density of matter in objects of type X is

$$\beta(\delta_{c,X}) = 2A \int_{\delta_{c,X}}^{\infty} P(\delta_R) d\delta_R, \quad (1)$$

where P is a function of quantities smoothed on the scale R associated with objects of mass M. For a generic PDF, the factor A can be determined by ensuring that counting all objects recovers the total mass density. Here we will just use the Press-Schechter factor of 2 ($A = 1$) as a good enough approximation, since we are only concerned with a subset of objects.[1] In addition, we have approximated the

---

[1]This factor would be more important if one had a particular model for how the primordial statistics varied as a function of scale and combined constraints on that model from objects on several scales.

integral as extending to infinity, even though for some rare objects (e.g. UCMHs) that are not quite as rare as others (e.g. PBHs), there will be a finite upper limit associated with δ_c of the "next rarer" object. Extending the integral to infinity is a very good approximation as long as δ_c is dissimilar for different rare objects, as P is always a steeply falling function of δ_R. If the probability distribution is Gaussian, we have

$$\beta_{X,\text{Gaussian}}(\nu_{c,X}) = \text{erfc}\left[\frac{\nu_{c,X}}{\sqrt{2}}\right], \qquad (2)$$

where $\nu_{c,X} \equiv \delta_{c,X}/\sigma_R$, and σ_R is calculated in the cosmological era relevant for the formation of objects of type X. In general we will suppress the subscript X on β, δ_c and ν_c, except where talking about a specific type of object. Our convention follows Ref. [11]: R refers to the comoving radius of the region that forms the object, $k = 1/R$ is the corresponding wave number, and both $\delta_{c,X}$ and σ_R are computed when the region enters the Hubble horizon. This is an analytic technique, which should be adjusted to agree with numerical simulations before it is trusted to give precise predictions of object abundance.

We can describe weakly non-Gaussian distributions as an expansion in moments about a Gaussian. A generic expression, generalizing the Edgeworth expansion, was written down by Petrov [32,33]:

$$P(\nu)\mathrm{d}\nu = \frac{\mathrm{d}\nu}{\sqrt{2\pi}} e^{-\nu^2/2}\bigg\{1 + \sum_{s=1}^{\infty} \sigma_R^s \sum_{\{k_m\}} H_{s+2r}(\nu) \prod_{m=1}^{s} \frac{1}{k_m!}$$
$$\times \left(\frac{S_{m+2,R}}{(m+2)!}\right)^{k_m}\bigg\}. \qquad (3)$$

In this expression, the second summation is over all sets of integers $\{k_m\}$ that satisfy

$$k_1 + 2k_2 + \cdots + sk_s = s; \qquad (4)$$

$H_n(\nu)$ are Hermite polynomials,[2]

$$H_n(\nu) = (-1)^n e^{\nu^2/2} \frac{\mathrm{d}^n}{\mathrm{d}\nu^n} e^{-\nu^2/2}; \qquad (5)$$

and $r = k_1 + k_2 + \cdots + k_m$. The "reduced cumulants" are defined as

$$S_{n,R} \equiv \frac{\langle\delta_R^n\rangle_c}{\langle\delta_R^2\rangle_c^{n-1}}. \qquad (6)$$

To decide how to organize the terms in this series and truncate it somewhere, we need additional input about the relative importance of higher-order moments for the physical case under consideration. The choices are more naturally phrased in terms of the connected part of the dimensionless moments (the cumulants) of the real space density contrast, smoothed on scale R:

$$\mathcal{M}_{n,R} = \frac{\langle\delta_R^n\rangle_c}{\langle\delta_R^2\rangle^{n/2}}. \qquad (7)$$

The moments $\mathcal{M}_{n,R}$ can equally well be thought of as functions $\mathcal{M}_n(k)$ of the wave number k corresponding to the smoothing scale R. We will use these two notations interchangeably, and sometimes we do not write out the scale dependence explicitly at all.

Two choices motivated by particle physics [34] are hierarchical scaling,

$$\mathcal{M}_n^{\text{h}} \propto (I\mathcal{P}_{\mathcal{R}}^{\frac{1}{2}})^{n-2}, \qquad (8)$$

and feeder scaling,

$$\mathcal{M}_n^{\text{f}} \propto I^n, \quad n \geq 3, \qquad (9)$$

where $\mathcal{P}_{\mathcal{R}}$ is the amplitude of fluctuations in the primordial curvature, $\langle\mathcal{R}^2(\mathbf{x})\rangle = \int \frac{dk}{k} \mathcal{P}_{\mathcal{R}}(k)$, and I is a parameter indicating the strength of the interaction that sourced the non-Gaussianity. In specific models, one can work out the coefficients relating $I \propto f_{\text{NL}}$, where f_{NL} is the parameter for which CMB experiments quote constraints on non-Gaussianity. Using simple local models with each scaling to determine representative combinatorics,[3] we can express the typical size of all higher-order moments ($n \geq 3$) in terms of the third moment $\mathcal{M}_3$ by

$$\text{Hierarchical: } \mathcal{M}_n^{\text{h}} = n!2^{n-3}\left(\frac{\mathcal{M}_3^{\text{h}}}{6}\right)^{n-2}, \qquad (10)$$

$$\text{Feeder: } \mathcal{M}_n^{\text{f}} = (n-1)!2^{n-1}\left(\frac{\mathcal{M}_3^{\text{f}}}{8}\right)^{n/3}. \qquad (11)$$

These are, of course, just two examples of how the moments could scale. For currently developed inflation models, they are representative of the full range of well-motivated possibilities (which we discuss in more detail in the next section). Some models (quasisingle field inflation [35]) are a hybrid between the two scalings, or have additional numerical coefficients (e.g. two-field local models). Any skewness in $P(\nu)$ must be positive (assuming that higher moments are subdominant) for non-Gaussianities to increase the production of rare objects like PBHs, as the positive tail of the distribution must be enhanced relative to the core.

Once the scaling of the moments has been established, Eq. (3) can be written in terms of the dimensionless smoothed moments $\mathcal{M}_{n,R}$, and the series can be organized in a form appropriate for any scaling. Making use of Eq. (4), we find

---

[2]Note that there are two conventions for naming the Hermite polynomials, and these are sometimes denoted He_n.

[3]We use the local ansatz $\mathcal{R}(x) = \mathcal{R}_G(x) + \frac{3}{5}f_{\text{NL}}[\mathcal{R}_G(x)^2 - \langle\mathcal{R}_G(x)^2\rangle]$ for the hierarchical scaling (in that case, $I = f_{\text{NL}}$). We use a two-field extension, $\mathcal{R}(x) = \phi_G + \sigma_G + \frac{3}{5}\tilde{f}_{\text{NL}}[\sigma_G(x)^2 - \langle\sigma_G(x)^2\rangle]$ with $\tilde{f}_{\text{NL}}\mathcal{P}_\sigma^{1/2} \gg 1$, for the feeder scaling (then $I = \tilde{f}_{\text{NL}}^{\text{eff}} = \tilde{f}_{\text{NL}}\mathcal{P}_\sigma/\mathcal{P}_{\mathcal{R}}^{1/2}$).

$$P(\nu)\mathrm{d}\nu = \frac{d\nu}{\sqrt{2\pi}}e^{-\nu^2/2}\left\{1 + \sum_{s=1}^{\infty}\sum_{\{k_m\}} H_{s+2r}(\nu)\prod_{m=1}^{s}\frac{1}{k_m!}\left(\frac{\mathcal{M}_{m+2,R}}{(m+2)!}\right)^{k_m}\right\} \equiv \frac{d\nu}{\sqrt{2\pi}}e^{-\nu^2/2}(1 + p_1(\nu,R) + p_2(\nu,R) + \cdots). \quad (12)$$

For the two scalings above, we have

$$p_1^{(h)}(\nu,R) = \frac{\mathcal{M}_{3,R}^h}{3!}H_3(\nu), \qquad p_2^{(h)}(\nu,R) = \frac{\mathcal{M}_{4,R}^h}{4!}H_4(\nu) + \frac{1}{2}\left(\frac{\mathcal{M}_{3,R}^h}{3!}\right)^2 H_6(\nu),$$

$$p_3^{(h)}(\nu,R) = \frac{\mathcal{M}_{5,R}^h}{5!}H_5(\nu) + \frac{\mathcal{M}_{3,R}^h\mathcal{M}_{4,R}^h}{3!4!}H_7(\nu) + \frac{1}{3!}\left(\frac{\mathcal{M}_{3,R}^h}{3!}\right)^3 H_9(\nu)$$

(13)

for the hierarchical scaling (this is the usual Edgeworth expansion) and

$$p_1^{(f)}(\nu,R) = \frac{\mathcal{M}_{3,R}^f}{3!}H_3(\nu), \qquad p_2^{(f)}(\nu,R) = \frac{\mathcal{M}_{4,R}^f}{4!}H_4(\nu), \qquad p_3^{(f)}(\nu,R) = \frac{\mathcal{M}_{5,R}^f}{5!}H_5(\nu),$$

$$p_4^{(f)}(\nu,R) = \left(\frac{\mathcal{M}_{6,R}^f}{6!} + \frac{1}{2}\left(\frac{\mathcal{M}_{3,R}^f}{3!}\right)^2\right)H_6(\nu), \qquad p_5^{(f)}(\nu,R) = \left(\frac{\mathcal{M}_{3,R}^f\mathcal{M}_{4,R}^f}{3!4!} + \frac{\mathcal{M}_{7,R}^f}{7!}\right)H_7(\nu)$$

(14)

for the feeder scaling. With this organization, the series in Eq. (12) is an asymptotic expansion for PDFs with moments of either scaling. Organizing the expansion in this way also exposes the key difference between the hierarchical and feeder scalings: $p_i^{(h)} \propto \mathcal{M}_3^i$, while $p_i^{(f)} \propto \mathcal{M}_3^{i/3}$, which implies that higher-order terms in the expansion make a larger contribution to feeder distributions than they do to hierarchical distributions.

Then, arranging the sums according to each scaling and performing the integral in Eq. (1), we find that for objects of type X,

$$\beta^{(h)}(\nu_{c,X}) = \mathrm{erfc}\left(\frac{\nu_{c,X}}{\sqrt{2}}\right) + 2\frac{e^{-\nu_{c,X}^2/2}}{\sqrt{2\pi}}\sum_{s=1}^{\infty}\sum_{\{k_m\}_h} H_{s+2r-1}(\nu_{c,X})\prod_{m=1}^{s}\frac{1}{k_m!}\left(\frac{\mathcal{M}_{m+2,R}}{(m+2)!}\right)^{k_m},$$

$$\beta^{(f)}(\nu_{c,X}) = \mathrm{erfc}\left(\frac{\nu_{c,X}}{\sqrt{2}}\right) + 2\frac{e^{-\nu_{c,X}^2/2}}{\sqrt{2\pi}}\sum_{s=1}^{\infty}\sum_{\{k_m\}_f} H_{s+1}(\nu_{c,X})\prod_{m=1}^{s}\frac{1}{k_m!}\left(\frac{\mathcal{M}_{m+2,R}}{(m+2)!}\right)^{k_m},$$

(15)

where the sets $\{k_m\}_h$ are again non-negative integer solutions to $k_1 + 2k_2 + \cdots + sk_s = s$ and $r = k_1 + k_2 + \cdots + k_m$, but the sets $\{k_m\}_f$ are solutions to $3k_1 + 4k_2 + \cdots + (s + 2)k_s = s + 2$. Now, for either scaling, truncating the series at some finite s in the sums above keeps all terms up to the same order in $\mathcal{M}_3$: $\mathcal{M}_3^s$ for hierarchical scalings and $\mathcal{M}_3^{s/3}$ for feeder scalings.

We will work out constraints for the single parameter $\mathcal{M}_3$ characterizing the two classes of distribution (hierarchical and feeder) that we have motivated using their particular Petrov expansions shown above. First, we illustrate how the Petrov-type expansions work in Fig. 1, which shows a class of examples of the Edgeworth expansion compared to the corresponding complete distribution. The panels show three χ^2 distributions with varying levels of non-Gaussianity and three truncations of the Edgeworth expansion for each case. As the figures show, the Edgeworth series is accurate out to larger values of ν when the distribution is closer to Gaussian. But, for a given level of non-Gaussianity, there is a value of ν beyond which the series is not an accurate fit, and adding more terms does not improve the fit. A study comparing a variety of expansions of non-Gaussian distributions, with additional examples, can be found in Ref. [36].

B. The approximate PDF and inflation models

The parametrization of non-Gaussianity in terms of the dimensionless moments and their relative importance is the relevant organization for any observables sensitive to the PDF. There is a straightforward mapping between the dimensionless skewness and the power spectrum and bispectrum of any particular model (and choice of cosmological parameters). For example, for local and equilateral bispectral templates, the dimensionless skewness in the dark matter density perturbations at horizon crossing for modes that enter the horizon during radiation domination is

$$\mathcal{M}_{3,R} = 3.13 f_{\mathrm{NL}}^{\mathrm{local}}\mathcal{P}_{\mathcal{R}}^{1/2}(k \approx 2R^{-1}) \quad (16)$$

$$= 1.22 f_{\mathrm{NL}}^{\mathrm{equil}}\mathcal{P}_{\mathcal{R}}^{1/2}(k \approx 2R^{-1}). \quad (17)$$

We have integrated over the bispectra, using a top-hat window function and the transfer function described in the Appendix of Ref. [12]. We have assumed that both the power spectrum and f_{NL} parameters are constant over the range where the integrals peak, $k \approx 2/R$, which is not too restrictive, since only about a decade in k contributes significantly to the integrals.
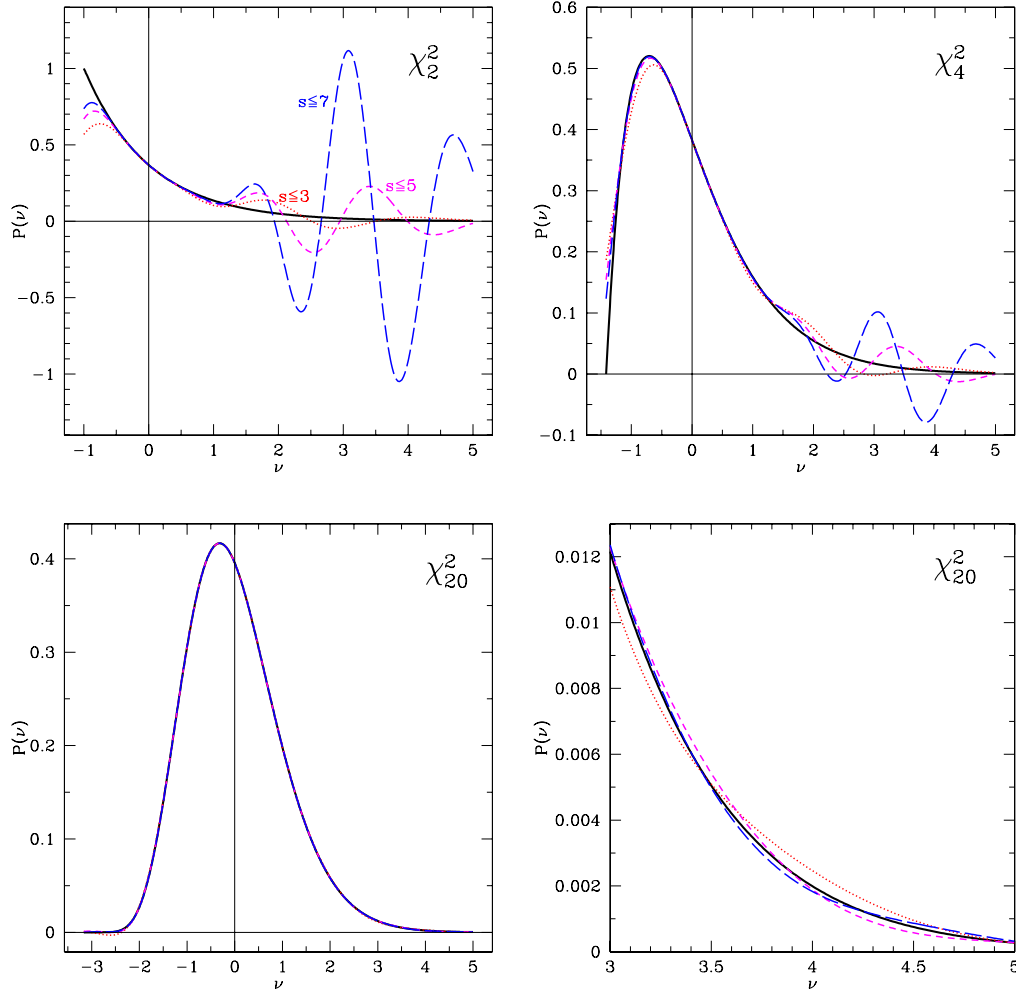
FIG. 1 (color online).   A comparison of the Edgeworth expansion to the exact PDF for the χ_k^2 distribution with mean 0 and variance 1, $P(\nu) = \sqrt{2k}(\sqrt{2k}\nu + k)^{(k/2-1)}\mathrm{Exp}[-(\sqrt{2k}\nu + k)/2]/(2^{k/2}\Gamma(k/2))$. The upper-left panel shows the full distribution (solid black line) for $k = 2$, as well as the Edgeworth expansion truncated after p_3 (dotted red), p_5 (short-dashed pink), and p_7 (long-dashed blue), using hierarchical ordering. The upper-right panel shows the same curves for $k = 4$, while the lower panels are for $k = 20$. The lower-right panel is a zoomed-in view of the tail of the $k = 20$ distribution.

The dimensionless moments are straightforward to calculate for any inflation model, and the hierarchical and feeder scalings we use are representative of the full range of primordial non-Gaussianity that has its origin in the particle physics of inflation. Since the moments are determined from a perturbation-theory expansion of some Lagrangian, they do indeed scale in a simple way depending on a single parameter that determines the overall level of interaction together with order-1 coefficients. Any scenario where the non-Gaussian interactions can be modeled by a single source, including inflaton self-interactions and curvaton models, gives rise to the hierarchical scaling. Scenarios where non-Gaussian fields couple to the primary source of curvature perturbation to generate the non-Gaussianity have the feeder scaling or a hybrid of the two types. Since a pure feeder scaling is the most non-Gaussian

scaling (for a given skewness) known to be generated from inflation, the purely hierarchical and purely feeder scalings nicely bracket the behavior of all known inflation models.

Finally, although we are dealing with much smaller scales than what is seen in the CMB, the behavior of the power spectrum and higher-order correlations on all scales can be computed in any given inflation model. So, the parameters constrained by probes of very small scales provide complementary constraints on the models. In addition, models where the non-Gaussianity changes on small scales typically have correlated changes in the amplitude of the power spectrum. The two effects are not arbitrary, as both come from the dynamics of the fields sourcing inflation. A full study of the relationship between the evolution of the power spectrum and amplitude of non-Gaussianity has not yet been done for a very
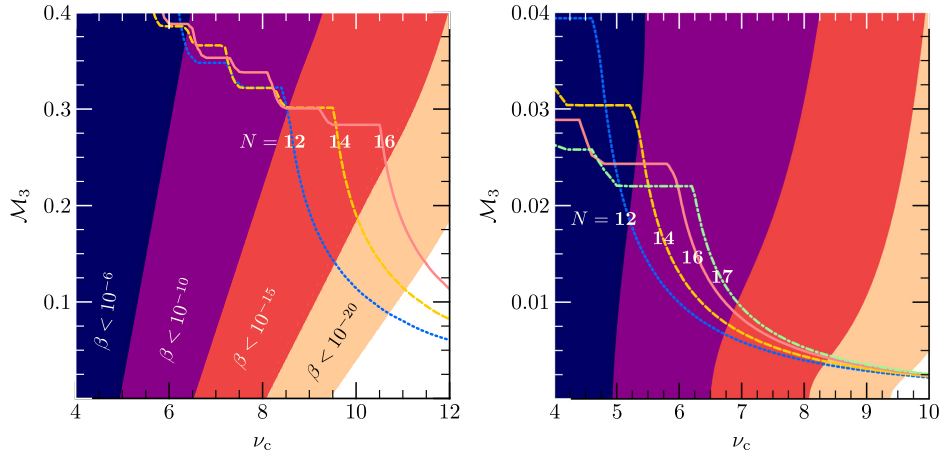
FIG. 2 (color online). Computational and observational constraints for hierarchical models (left) and feeder models (right). The curves show the values of $\mathcal{M}_3$ above which the errors in the PDF truncated at N terms exceed 20% within the range of ν values that contribute to $\beta(\nu_{c,X})$. The N for each curve is indicated directly on the curves. The shaded regions are excluded by observational constraints on β, with $N = 16$ for hierarchical models and $N = 17$ for feeder models, assuming $\beta < 10^{-6}$, 10^{-10}, 10^{-15} or 10^{-20}. Only the parts of the shaded regions falling below the colored curves are reliably excluded.

wide variety of models, but see Figs. 1 and 2 in Ref. [37] for an example.[4] Our joint constraints on the power spectrum and non-Gaussianity are particularly valuable for these models. We will return to this point when we present our results.

C. Truncation and error evaluation

To calculate the abundance of rare objects, we must truncate the β expansions given by Eq. (15), and we must determine the accuracy of the truncated series. For a given level of non-Gaussianity, there will be some region in the high-ν tail of the PDF for which the truncated expansion of the distribution is not reliable. Furthermore, for a given value of ν, there is also a level of non-Gaussianity beyond which the expansion is not reliable. This maximum accessible value of $\mathcal{M}_3$ can be increased by including higher-order terms, but only up to a point. Eventually the inclusion of more terms will stop improving the accuracy of the truncated sum, and adding additional terms makes the expansion less reliable (e.g., see the blue long-dashed curves in Fig. 1). In this section, we determine the optimal number of terms to include in the expressions for the PDF and β. We also describe how we evaluate the accuracy of the resulting expressions and find the maximum value of $\mathcal{M}_3$ that can be reliably constrained by a given upper bound on β. Some previous work on the validity of truncating this series for the hierarchical case only and applying it to cluster number counts can be found in Ref. [38]. A recent study of observational constraints on the feeder and hierarchical scalings from x-ray detected

clusters demonstrates that current data analysis is sensitive to the choice of scaling [29].

To calculate the number of objects of a given type, we must be able to trust the PDF for values of $\nu > \nu_{c,X}$ that contribute to the β integral given by Eq. (1). Fortunately, only a limited range of ν values make a significant contribution to β, because the PDF decreases rapidly as ν increases ($\sim \exp[-\nu^2/2]$). Therefore, we only need to trust the PDF for $\nu_{c,X} < \nu < \nu_{\max}$, where $\nu_{\max}$ is chosen so that truncating the β integral at $\nu_{\max}$ has a minimal impact on β. More explicitly, we demand that

$$\frac{\beta(\nu_{\max})}{\beta(\nu_{c,X})} < 0.02 \qquad (18)$$

so that the PDF for values of $\nu > \nu_{\max}$ contributes less than 2% of the β integral. If the observational upper bound on β is greater than 5×10^{-24}, then Eq. (18) is satisfied for all relevant $\nu_{c,X}$ values and all $\mathcal{M}_3$ values for which we can trust the PDF for $\nu_{c,X} < \nu < \nu_{\max}$ if we set $\nu_{\max} = K\nu_{c,X}^{0.7}$, where $K = 2.1$ for feeder models and $K = 2.2$ for hierarchical models. Since all the bounds on UCMHs and most of the bounds on PBHs are less restrictive than $\beta < 5 \times 10^{-24}$, we adopt these expressions for $\nu_{\max}$ in our analysis. Using stricter bounds on β to constrain non-Gaussianity would require a larger value for $\nu_{\max}$, which could prevent these improved constraints on β from providing improved constraints on $\mathcal{M}_3$ using asymptotic expansions of the PDF.

Next, we evaluate the error in the PDF over the range $\nu_{c,X} < \nu < \nu_{\max}$ as a function of $\mathcal{M}_3$ and $\nu_{c,X}$. Since the series expansion of the PDF is an asymptotic expansion, the error introduced by truncating the series at N terms is of

---

[4]We thank D. Seery and R. Ribeiro for correspondence on this point.

the same order as the $N + 1$ term in the expansion.[5] Given $\mathcal{M}_3$ and $\nu_{c,X}$, we compute the ratio of the $N + 1$ term in the PDF to the sum of the first N terms and find its maximum value in the range $\nu_{c,X} < \nu < \nu_{max}$: this ratio estimates the maximum fractional error in the truncated PDF over the range of ν values that contribute significantly to $\beta(\nu_{c,X})$.

In our analysis, we restrict ourselves to values of $\mathcal{M}_3$ for which this maximum fractional error is less than 20%. We use this criterion to define a "computationally accessible" region in the $(\nu_{c,X}, \mathcal{M}_3)$ parameter space. Given that the PDF is accurate to within 20% for all ν values that significantly contribute to $\beta(\nu_{c,X})$, we know that the fractional error in $\beta(\nu_{c,X})$ is at most 20% for these values of $\mathcal{M}_3$ and $\nu_{c,X}$, and we consider the truncated expansion to be trustworthy in this region. We expect that the error in β is actually lower than 20%, because the errors in the PDF are less than 20% for most values of ν between $\nu_{c,X}$ and ν_{max}. Furthermore, these errors may be positive or negative and therefore may partially cancel when the PDF is integrated to obtain β.

This expectation was confirmed when we applied our procedure to the χ_k^2 probability distributions and compared the values of β obtained by integrating the truncated PDF expansion over all $\nu > \nu_c$ to the values obtained by integrating the true PDF (given in the caption to Fig. 1). We tested series that were truncated at $s = 7, 8, 9,$ and 10, and we considered the series trustworthy only if the estimated PDF errors were less than 20% within the range $\nu_{c,X} < \nu < \nu_{max}$. We estimated the errors in the PDF using the first term not included in the expansion; we did not compare the truncated PDF to the true PDF. For all truncations, the β values obtained from the truncated series with ν_c and k values in that truncation's trustworthy region differed from the true β values by less than 6%.

Figure 2 shows the region of the $\nu_{c,X}$-$\mathcal{M}_3$ plane that is computationally accessible for $N = 12, 14,$ and 16 for both hierarchical and feeder models, and $N = 17$ for feeder models; the trustworthy regions lie below the thick curves. For hierarchical models, we see that adding terms with $N > 12$ does not significantly change the upper bound on computationally accessible $\mathcal{M}_3$ values for $\nu_{c,X} < 8$. For larger values of $\nu_{c,X}$, adding more terms to the PDF brings larger values of $\mathcal{M}_3$ into the trustworthy region. For feeder models, adding more terms decreases the range of accessible $\mathcal{M}_3$ values for small values of $\nu_{c,X}$ but improves the range for larger values of $\nu_{c,X}$. However, for $\nu_{c,X} > 8$, adding more terms does not significantly change the maximum trustworthy value of $\mathcal{M}_3$.

To determine the optimal number of terms to consider when computing the PDF, we must know what $\nu_{c,X}$ values we are interested in. The shaded regions in Fig. 2 show the regions that are excluded by different observational upper bounds on β. To maximize the power of a given upper bound on β, we want to choose N to maximize the portion of the corresponding excluded region that lies within the computationally accessible region (i.e. beneath the thick curves in Fig. 2). For example, the darkest two shaded areas show the values excluded by the bound $\beta < 10^{-10}$. For hierarchical models, we should choose $N \gtrsim 13$ when evaluating this constraint. For feeder models, $N = 17$ is the best choice, because it maximizes the excluded range of $\nu_{c,X}$ values, but we could reach higher $\mathcal{M}_3$ by also considering PDFs with fewer terms. In practice, however, changing the number of terms in the expansion for β only changes the excluded regions for values of $\nu_{c,X}$ near the excluded region's boundary and for untrustworthy $\mathcal{M}_3$ values that are far above all the thick curves in Fig. 2. Therefore, we should choose N such that the computationally accessible excluded region reaches the highest possible values of $\nu_{c,X}$. We can also trust this excluded region for all $\mathcal{M}_3$ that are computationally accessible for any smaller value of N. Our default value for N is 16 for hierarchical models and 17 for feeder models. We chose this value of N due to computational limitations, but from Fig. 2 we see that including more terms offers no benefits for hierarchical models with upper bounds on β greater than 10^{-15}. For feeder models, including more terms could marginally extend the computationally accessible region excluded by $\beta < 10^{-10}$, but it would not have any impact for lower or higher bounds on β.

Figure 3 compares several Gaussian and non-Gaussian PDFs (left panel) and their positive tails (right panel). The distributions have parameters found using the procedure above and the bound on the abundance of UCMHs at $k = 2 \times 10^7$ Mpc^{-1}, for which $\delta_c \simeq 0.00156$ and $\beta \lesssim 3.0 \times 10^{-6}$. The thick solid black curve is the Gaussian distribution that saturates the bound on β, with $\sigma_R = 0.000335$. The blue solid curve is the non-Gaussian distribution with hierarchical scaling and the maximum trustworthy value of $\mathcal{M}_3$ that gives $\beta = 3.0 \times 10^{-6}$ ($\mathcal{M}_{3,max} = 0.4$ for $N = 12$). This case also defines a minimum variance $\sigma_R = 0.000259$ for a hierarchical distribution that saturates the bound on β; the dotted blue curve is the Gaussian distribution with that same value of σ. In the right hand panel, this Gaussian distribution has very little area under the positive tail and lies nearly on top of the x axis. The long-dashed red curve is the non-Gaussian distribution with feeder scaling and the maximum trustworthy value of $\mathcal{M}_3$ that gives $\beta = 3.0 \times 10^{-6}$ ($\mathcal{M}_{3,max} = 0.03$ for $N = 13$). This curve has $\sigma_R = 0.000308$, and the dotted red curve in the right panel shows the tail of a Gaussian distribution with the same value of σ. In the right panel, this curve would be nearly indistinguishable from the red

---

[5]For both the hierarchical and the feeder models, we define the Nth term in the expansion according to the power of $\mathcal{M}_3$ in that term. For hierarchical models, the Nth term is proportional to $\mathcal{M}_3^N$ and corresponds to the $s = N$ term in Eq. (12). For feeder models, the Nth term is proportional to $\mathcal{M}_3^{N/3}$ and receives contributions from all terms with $s + 2r = N$ in Eq. (12).
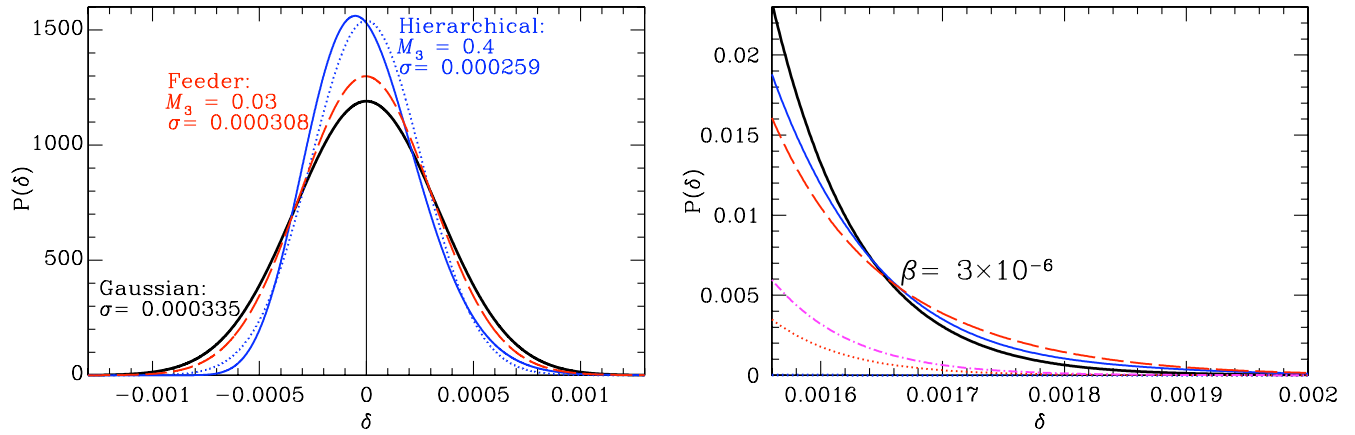
FIG. 3 (color online). Comparing Gaussian and non-Gaussian PDFs for UCMHs at $k = 2.0 \times 10^7$ Mpc^{-1}; the right panel shows $\delta \gtrsim \delta_c \simeq 1.56 \times 10^{-3}$. In both panels, the thick solid black curve is the Gaussian PDF that saturates the observational bound on the UCMH abundance ($\beta \leq 3.0 \times 10^{-6}$). Both panels also show the maximally non-Gaussian hierarchical PDF (solid blue) and feeder PDF (dashed red) that saturate this bound. (That is, with the maximum value of $\mathcal{M}_3$ for which the expansion is controlled out to the largest value of ν that contributes significantly to the integral.) In the right panel, we indicate that all three of these curves have the same value of β. The dotted curves show Gaussian PDFs that have the same values of σ as these maximally non-Gaussian PDFs (blue for hierarchical and red for feeder). In addition, the dot-dashed curve in the right-hand panel shows a hierarchical PDF with the same values of σ and $\mathcal{M}_3$ as the maximally non-Gaussian feeder PDF that saturates the β bound (the red dashed curve).

dashed curve. The dot-dashed curve in the right panel is one way of visualizing the difference between the feeder and hierarchical scalings: it shows a hierarchical non-Gaussian PDF with values of σ and $\mathcal{M}_3$ that saturate β when used in a *feeder* PDF ($\mathcal{M}_{3,\mathrm{max}} = 0.03$ and $\sigma_R = 0.000308$, as shown by the red dashed curve). Clearly, for a particular value of $\mathcal{M}_3$, the feeder scaling gives a PDF that is overall substantially more non-Gaussian than the hierarchical scaling.

Now we can get some measure of the maximum shift in the PDF that can come from adding weak non-Gaussianity in a controlled way. Figure 4 illustrates this in two ways. The left-hand panel illustrates how a constraint on the variance σ derived assuming a Gaussian distribution can be shifted to a constraint on a non-Gaussian distribution with a new (smaller) variance. The vertical axis shows the relative shift in ν_c between the Gaussian distribution that gives the abundance fraction β and the maximally
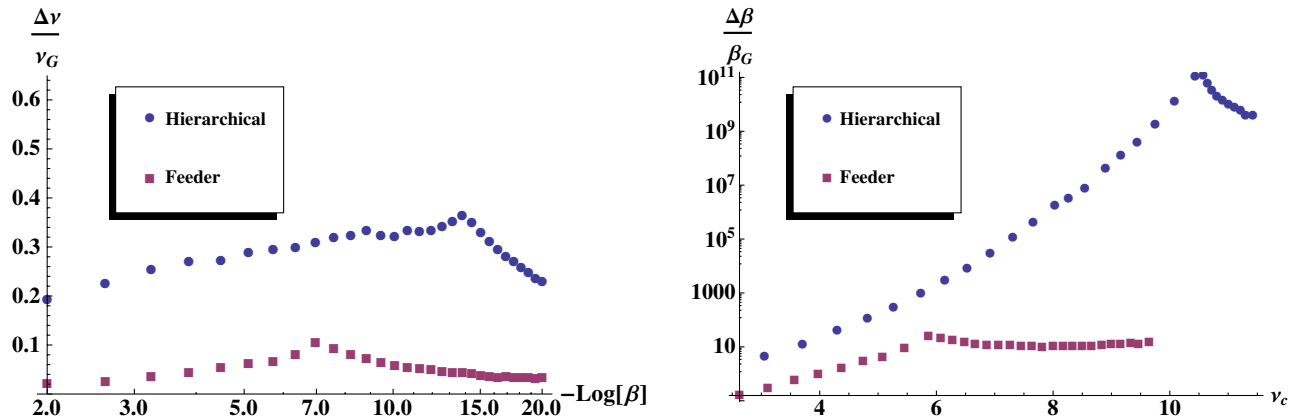

FIG. 4 (color online). *Left panel:* The ratio $[\nu_{c,\mathrm{NG}}(\mathcal{M}_{3,\mathrm{max}}) - \nu_{c,G}]/\nu_{c,G}$, where $\nu_{c,G}$ is the value of ν_c such that a Gaussian distribution gives the plotted value of β [see Eq. (2)], and $\nu_{c,\mathrm{NG}}(\mathcal{M}_{3,\mathrm{max}})$ is the value of ν_c such that a non-Gaussian distribution with $\mathcal{M}_3 = \mathcal{M}_{3,\mathrm{max}}$ gives the same value of β. Here $\mathcal{M}_{3,\mathrm{max}}$ is the largest amplitude of non-Gaussianity where the Petrov expansion, calculated out to $N = 16$ for hierarchical scaling and $N = 17$ for feeder scaling, is trustworthy, according to the criteria discussed below Eq. (18). *Right panel:* The ratio $[\beta(\mathcal{M}_{3,\mathrm{max}}, \nu_c) - \beta(0, \nu_c)]/\beta(0, \nu_c)$ plotted versus ν_c, assuming hierarchical (blue circles) or feeder (red squares) scaling. In both plots, the change in slope in the hierarchical points appears to be an artifact of our truncation to a finite number of terms—going beyond $N = 16$ would probably allow one to continue farther onto the tail. For the feeder scaling, however, it is less clear that adding more terms will help—the distribution is quite non-Gaussian, so the Petrov expansion breaks down irreparably at smaller ν and smaller $\mathcal{M}_3$. The scatter in the points about a smooth line comes from the approximate nature of our criteria for determining $\mathcal{M}_{3,\mathrm{max}}$.

non-Gaussian distribution, under control according to our criteria above, that gives the same β. To use this for any particular object, ν_c should be written in terms of the appropriate δ_c and σ_R for the object of interest. $\mathcal{M}_{3,\max}$ can likewise be converted to a non-Gaussian parameter in the primordial distribution. This plot gives a sense of how accurately the variance of fluctuations must be known in order for number counts to provide an independent constraint on non-Gaussianity using Petrov expansions. For instance, if $\Delta \nu / \nu_G = 0.3$ in the left panel of Fig. 4, then adding weak non-Gaussianity in a controlled way has the same effect on the abundance of rare objects as increasing σ_R by $\sim 30\%$. If the uncertainty in σ_R exceeds 30%, then number counts with the corresponding value β cannot independently constrain non-Gaussianity without knowing the full PDF of the density fluctuations.

The right-hand panel of Fig. 4 shows the maximum fractional change in the abundance β that results from adding controlled non-Gaussianity for a given object and fixed variance (so fixed ν_c): $[\beta(\mathcal{M}_{3,\max}, \nu_c) - \beta(0, \nu_c)]/\beta(0, \nu_c)$. Notice that the primary limit on how much β can change is our ability to accurately determine the PDF. Since the feeder scaling produces a distribution that is overall much more non-Gaussian than the hierarchical for a fixed value of $\mathcal{M}_3$, we cannot accurately calculate at as large a ν_c (rarer objects) or for as large a shift in β as for the hierarchical case. Again, the effects of more non-Gaussian distributions can be computed if one knows the full distribution rather than just a finite number of moments.

III. PRIMORDIAL BLACK HOLES

For a Gaussian spectrum of fluctuations, the present-day fractional density of black holes has typically been used to constrain the spectral index n_s under the assumption of a scale-free spectrum [14,18,39,40]. The resulting bound, $n_s \lesssim 1.3$ (based on e.g. VERITAS searches at $k = 3 \times 10^{15}$ Mpc^{-1} [41]), is easily satisfied by the current experimentally measured value $n_s = 0.9603 \pm 0.0073$ [42], with little or no running. In terms of generalized, scale-dependent Gaussian power spectra, PBHs have been used to place bounds on the amplitude of curvature perturbations at the level of $\mathcal{P}_\mathcal{R} \lesssim 10^{-2}$ [8,43], on scales corresponding to $10^{-2} < k \lesssim 10^{19}$ Mpc^{-1}. Stronger constraints than this can be found if the analysis is done using a top-hat window function rather than a Gaussian, which is arguably a more correct method [11,44].

Present observational bounds on PBHs [8,9] limit β_{PBH} to values ranging from $\beta_{\mathrm{PBH}} \lesssim 10^{-10}$ for solar-mass PBHs, to $\beta_{\mathrm{PBH}} \lesssim 10^{-28}$ for a small range of PBH masses around 10^{13} g, to $\beta_{\mathrm{PBH}} \lesssim 10^{-20}$ for even smaller masses. A PBH mass M_{PBH} corresponds to horizon mass $M_{\mathrm{H}} = 3^{3/2} M_{\mathrm{PBH}}$ [8], and therefore wave number

$$\frac{k}{k_{\mathrm{eq}}} = (2 \times 3^3)^{-1/4} \left(\frac{g_\star^R}{g_\star^{\mathrm{eq}}} \right)^{1/4} \left(\frac{g_{\star S}^{\mathrm{eq}}}{g_{\star S}^R} \right)^{1/3} \left(\frac{M_{\mathrm{H}}^{\mathrm{eq}}}{M_{\mathrm{PBH}}} \right)^{1/2}. \quad (19)$$

Here, $k_{\mathrm{eq}} = 0.072 \Omega_m h^2$ Mpc$^{-1} = 9.68 \times 10^{-3}$ Mpc^{-1} [1] and $M_{\mathrm{H}}^{\mathrm{eq}} = 3.5 \times 10^{17} M_\odot$ [13] are the wave number and horizon mass, respectively, at equality, and $g_{\star S}^{\mathrm{eq}} = 3.91$, $g_\star^{\mathrm{eq}} = 3.36$, $g_\star^R \geq 106.75$ [45] are the respective degrees of freedom at equality and the time of horizon entry.

We use the current limits on β_{PBH} from Ref. [9], along with Eq. (15) and the requirement that it represents a well-behaved expansion, to determine the range of values of $\mathcal{M}_3(k)$ that are excluded by nonobservation of PBHs. For this calculation, we draw upon the expressions for σ_R in a generalized Gaussian power spectrum, given in Appendix B of Ref. [11]; in particular, we employ the top-hat window function advocated in that paper. We give our limits in Fig. 5 as a function of the power spectrum of curvature perturbations at any given k:

$$\mathcal{P}_\mathcal{R}(k) = 1.10 \sigma_R^2. \quad (20)$$

We show the excluded areas as dark shaded regions, for both hierarchical and feeder scaling, and at three different example scales. The corresponding limits on β at these scales are $\beta < 2.2 \times 10^{-13}$ ($k = 2 \times 10^3$ Mpc^{-1}), $\beta < 9.4 \times 10^{-12}$ ($k = 2 \times 10^7$ Mpc^{-1}) and $\beta < 1.4 \times 10^{-23}$ ($k = 2 \times 10^{17}$ Mpc^{-1}). More positive non-Gaussianity (larger $\mathcal{M}_3$) results in stronger exclusions, up until the point at which our expansion in Eq. (15) begins to break down. Above these very large values of $\mathcal{M}_3$, the excluded regions shrink with additional non-Gaussianity, as more of the available parameter space is made inaccessible by the increasingly poor behavior of the expansion.

IV. ULTRACOMPACT MINIHALOS

UCMHs were initially discussed in early papers by Berezinsky and collaborators [46–49], and have been recently studied in the context of lensing signatures [12,31,50], gamma-ray and neutrino signals from dark matter annihilation or decay [10,11,13,51–56], and impacts on the CMB [55,57–59]. To date, the strongest limits come from gamma-ray observations [11], indicating that $n_s < 1.17$ for a scale-free spectrum, and $\mathcal{P}_\mathcal{R}(k) \lesssim 10^{-6}$–$10^{-7}$ for generalized spectra, on scales corresponding to 3 Mpc$^{-1} < k < 3 \times 10^7$ Mpc^{-1}. The corresponding direct limits on β at the two example scales of Fig. 5 accessible with UCMHs are $\beta < 1.8 \times 10^{-9}$ ($k = 2 \times 10^3$ Mpc^{-1}) and $\beta < 3.0 \times 10^{-6}$ ($k = 2 \times 10^7$ Mpc^{-1}).

To explore the impact of UCMHs on the allowed parameter space for non-Gaussianity, we have implemented Eq. (15) in the analysis framework of Ref. [11], using combined one-year Fermi-LAT limits on the abundance of UCMHs from Galactic and extragalactic source searches,[6] as well as the overall diffuse gamma-ray flux in the direction of the Galactic poles. While two-year

---

[6]See the Appendix for details on our slight improvement of the treatment of sources over Ref. [11].
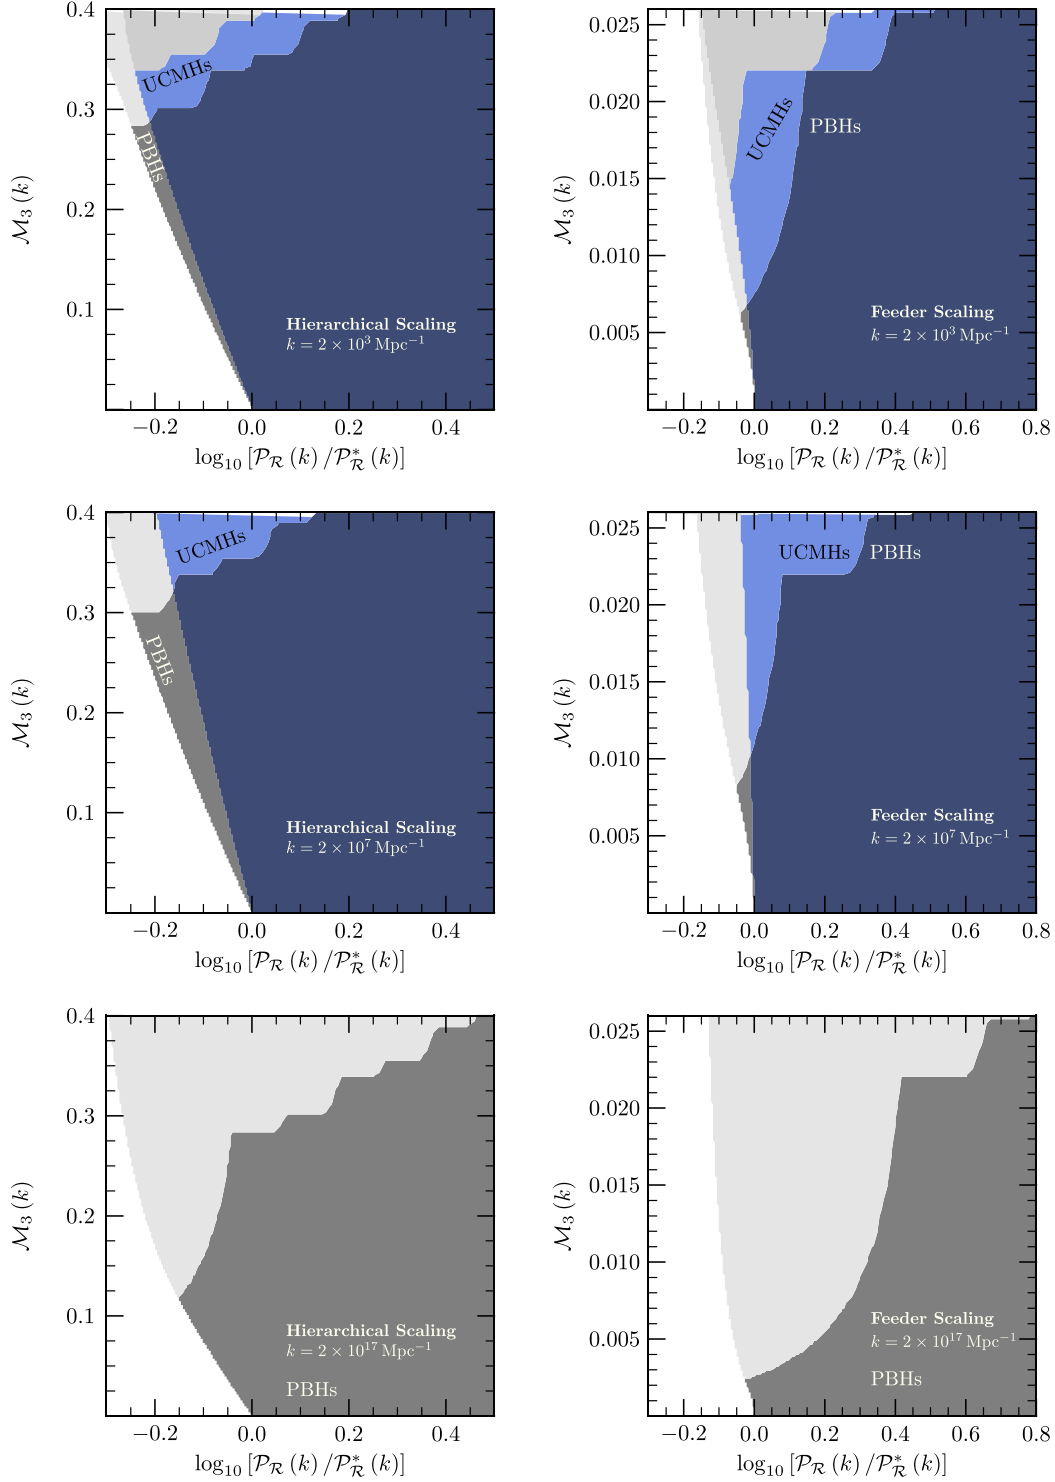
FIG. 5 (color online).   Bounds from UCMHs and PBHs on non-Gaussianity expressed as $\mathcal{M}_3(k)$ [the dimensionless skewness, defined in Eq. (7)], as a function of the amount of Gaussian power in curvature perturbations. Dark shaded regions are ruled out at a 95% C.L. by observations. Light shaded regions would have been ruled out by a naive application of the observational limits but correspond to areas where our analysis is invalid because our expansion in Eq. (15) breaks down. The limits are given at $k = 2 \times 10^3$ Mpc^{-1} (top), $k = 2 \times 10^7$ Mpc^{-1} (center) and $k = 2 \times 10^{17}$ Mpc^{-1} (bottom), for both the feeder and hierarchical scalings. Here the Gaussian power is expressed relative to the respective limits from PBHs and UCMHs on a Gaussian spectrum. In general, UCMHs probe non-Gaussianity at much lower overall amplitudes, but PBHs are sensitive to a broader range of scales. At $k = 2 \times 10^3, 2 \times 10^7, 2 \times 10^{17}$ Mpc^{-1}, for PBHs $\log_{10}(\mathcal{P}_{\mathcal{R}}^*) = -2.65, -2.58, -2.91$. For UCMHs with collapse redshift $z_c = 1000$, $\log_{10}(\mathcal{P}_{\mathcal{R}}^*) = -6.54, -6.87$ at $k = 2 \times 10^3, 2 \times 10^7$ Mpc^{-1}, whereas with $z_c = 200$, $\log_{10}(\mathcal{P}_{\mathcal{R}}^*) = -8.25, -8.47$.

source searches (e.g. Ref. [60]) and corresponding diffuse data are available, we conservatively adopt one-year sensitivities to be certain that no dark matter sources were found in the observing period we consider. The reader is referred to Ref. [11] for extensive discussion of this issue. Our gamma-ray limits assume the canonical thermal cross section for dark matter annihilation $(3 \times 10^{-26} \text{ cm}^3 \text{ s}^{-1})$ into b quarks and antiquarks, and a relatively conservative dark matter mass of 1 TeV (smaller masses would result in stronger limits).

In Fig. 5, we show the ability of UCMHs and PBHs to exclude different values of $\mathcal{M}_3$ for the two scaling scenarios, as a function of the departure of the underlying Gaussian power from the current Gaussian limit. Dark shaded regions are excluded at a 95% confidence level (C.L.). By expressing the underlying Gaussian power *relative to present limits on Gaussian power* from each class of compact object, it is possible to compare the limits on M_3 from UCMHs and PBHs directly. Although UCMHs provide a much stronger limit than PBHs on the Gaussian power (as known from Ref. [11]), PBHs provide a stronger lever arm for constraining non-Gaussianity below their Gaussian limit than UCMHs. On the other hand, our expansion is well behaved to larger $\mathcal{M}_3$ for UCMHS, allowing them to exclude stronger non-Gaussianities than PBHs.

The reference power $\mathcal{P}_{\mathcal{R}}^*$ that we have used for normalizing the curves shown in Fig. 5 is the observational limit on the Gaussian power at each scale. For PBHs at $k = 2 \times 10^3$, 2×10^7, 2×10^{17} Mpc^{-1}, this corresponds to $\log_{10}(\mathcal{P}_{\mathcal{R}}^*) = -2.65, \ -2.58, \ -2.91$. For UCMHs, the limit on the Gaussian power is much stronger. Precisely how much stronger depends on the latest redshift at which UCMHs can be assumed to undergo gravitational collapse. Assuming that only objects collapsing before a redshift of $z_\text{c} = 1000$ form UCMHs, the reference power is $\log_{10}(\mathcal{P}_{\mathcal{R}}^*) = -6.54$ at $k = 2 \times 10^3$ Mpc^{-1}, and $\log_{10}(\mathcal{P}_{\mathcal{R}}^*) = -6.87$ at 2×10^7 Mpc^{-1}.

Setting $z_\text{c} = 1000$ follows earlier work [11,31] and is a very conservative choice. The defining characteristic of a UCMH is that it forms in isolation from a spherical collapse; the processes of radial collapse and secondary infall give UCMHs extremely steep dark matter density profiles, thereby distinguishing them from halos formed at late times. Exactly when the approximation of radial infall breaks down is not yet known [11], as this depends sensitively on the velocities of dark matter particles and interactions between growing overdensities. More detailed simulation work is sorely needed to answer this question. In principle, UCMHs may form at lower redshifts, as long as the presence of different-sized overdensities does not spoil the assumption of spherical collapse. The smaller the redshift, the more unlikely it is that this requirement will be fulfilled. Later collapse redshifts allow smaller amplitude perturbations to form UCMHs, reducing $\delta_{\text{c,UCMH}}$ and increasing β_{UCMH} for a given combination of $\mathcal{M}_3$ and $\mathcal{P}_{\mathcal{R}}$. If
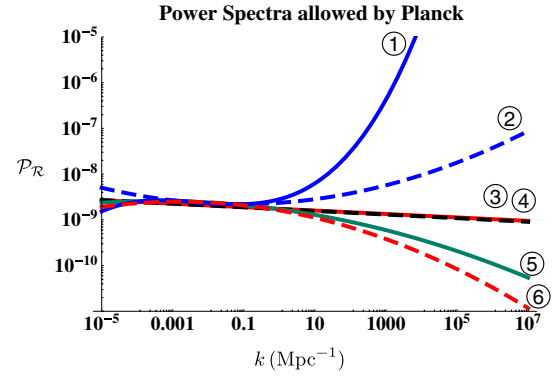

FIG. 6 (color online). The power spectra corresponding to the best-fit results of the Planck satellite for several choices of allowed parameters and combinations with other data sets [42]. The minimal six-parameter "ΛCDM" fit models the power spectrum amplitude (A_s) with a constant tilt ($n_s - 1$). If additional scale dependence of n_s and/or a nonzero amplitude of tensor modes (r) are allowed, the best-fit power spectrum can be different. From top to bottom at large k, the allowed parameters and data sets shown (where WP = WMAP polarization data) are as follows: (1) ΛCDM + $dn_s/d\ln k + d^2(n_s)/d\ln k^2$, Planck + WP (solid blue); (2) ΛCDM + $r + dn_s/d\ln k$, Planck + WP (dashed blue); (3) ΛCDM + r, Planck + WP (solid red); (4) ΛCDM, Planck + WP (dashed black); (5) ΛCDM + $dn_s/d\ln k$, Planck + WP (solid green); (6) ΛCDM + $r + dn_s/d\ln k$, Planck + WP + BAO (dashed red). Planck does not report the new best-fit value of A_s for the models with more than six parameters, but the 2σ uncertainty in the amplitude for the minimal model is within the thickness of the lines. The ΛCDM and ΛCDM + r models (lines 3 and 4) are reported at a pivot of $k_* = 0.002$ Mpc^{-1}, while the others have a pivot of $k_* = 0.05$ Mpc^{-1}.

UCMHs continue to be formed as late as $z_\text{c} = 200$, their nonobservation begins to constrain non-Gaussianities on small scales, even if the power spectrum is near the level observed on CMB scales: $\log_{10}(\mathcal{P}_{\mathcal{R}}^*) = -8.25$ at $k = 2 \times 10^3$ Mpc^{-1}, and $\log_{10}(\mathcal{P}_{\mathcal{R}}^*) = -8.47$ at 2×10^7 Mpc^{-1}. We stress again, however, that there exists *no observational evidence whatsoever* that the power on small scales is actually this low (see Fig. 6); even with a very conservative collapse redshift of $z_\text{c} = 1000$, UCMHs therefore place significant new constraints on non-Gaussianity at small scales.

V. DISCUSSION AND CONCLUSIONS

We have provided a general technique for using the abundance of any object to constrain weakly non-Gaussian statistics of the primordial fluctuations. We suggest precise criteria for the amplitude of non-Gaussianity that can be accurately constrained using Petrov-type expansions and emphasize that there is not a one-to-one mapping between the overall level of non-Gaussianity and the size of the skewness ($\mathcal{M}_3$). This is important when using the abundance of rare objects to constrain

non-Gaussianity: number counts are sensitive to *any* deviation from Gaussianity, but a measured value of β does not constrain a parameter like $f_{\rm NL}$ (typically interpreted as the amount of skewness) without additional information about the higher moments of the PDF.

In general, PBHs provide stringent constraints on non-Gaussianities only in the presence of fluctuations with relatively large variance, from $\mathcal{P}_{\mathcal{R}} \gtrsim 1\text{--}2 \times 10^{-3}$ at intermediate ($k = 2 \times 10^3$ Mpc^{-1}) and small ($k = 2 \times 10^7$ Mpc^{-1}) scales, to $\mathcal{P}_{\mathcal{R}} \gtrsim 7 \times 10^{-4}$ at very small scales. Notably though, at very small scales, PBHs provide the only available limit on the allowed amount of non-Gaussianity. UCMHs provide stronger constraints, and so may be more likely to eventually also constrain non-Gaussianity, but do not reach as small scales as PBHs.

For both UCMHs and PBHs, it is straightforward to map our constraints on $\mathcal{M}_3$ to constraints on a non-Gaussian parameter in the primordial fluctuations. For example, for the local ansatz $\mathcal{R}(x) = \mathcal{R}(x)_G + \frac{3}{5} f_{\rm NL}[\mathcal{R}_G(x)^2 - \langle \mathcal{R}_G(x)^2 \rangle]$, the smoothed, integrated, dimensionless skewness $\mathcal{M}_3 \sim \langle \mathcal{R}(x)^3 \rangle_R / (\langle \mathcal{R}(x)^2 \rangle_R)^{3/2}$ is very nearly equal to $f_{\rm NL} \mathcal{P}_{\mathcal{R}}^{1/2}$. For this case, if the variance is large enough, the abundance of UCMHs constrains $f_{\rm NL} \lesssim \mathcal{O}(100)$, and PBHs constrain $f_{\rm NL} \lesssim \mathcal{O}(10)$ on smaller scales. However, it is important to remember that this does not mean UCMHs are less powerful—quite the opposite. For the constraint from either object, the maximum value of $f_{\rm NL}$ corresponds to $\mathcal{M}_3 \sim \mathcal{O}(0.1)$, so those values of $f_{\rm NL}$ correspond to equally non-Gaussian distributions. The difference in constraints phrased as a value of $f_{\rm NL}$ comes from the difference in the variance of the fluctuations. For reference, recall that on CMB scales, where the amplitude of fluctuations is currently measured, $f_{\rm NL} \sim \mathcal{O}(10)$ corresponds to $\mathcal{M}_3 \sim \mathcal{O}(10^{-3})$.

To constrain non-Gaussianity, both PBHs and UCMHs require significantly more power than extrapolating the CMB spectrum from the standard six-parameter fit would suggest, but PBHs require even more power than UCMHs. So, given a prediction for the power on small scales, UCMHs are more likely to constrain non-Gaussianity. In addition, current observational constraints on cosmological models that allow more than six parameters show that additional power on small scales is hardly ruled out. Figure 6 shows the power spectra allowed by fits to the Planck satellite data [42].

The constraints we quote here assume that the dark matter particle is a standard thermal relic with detectable annihilation products [11]. Unfortunately, near-future astrometric microlensing searches for UCMHs can only constrain $\beta_{\rm UCMH} \lesssim 0.3$ [12]. For these large values of β, UCMHs are not sufficiently rare to be a useful probe of non-Gaussianity because $\nu_c \simeq 1$, which is smaller than the crossing point between Gaussian and skewed PDFs. Consequently, changing $\mathcal{M}_3$ does not significantly affect the abundance of relatively common objects ($\beta \gtrsim 0.1$).

Aside from the abundances of rare objects, the only other probe of the small-scale inhomogeneities is the spectral distortion of the CMB, specifically the μ distortion, which is sensitive to the fluctuations on scales on the order of $k = 50\text{--}10^4$ Mpc^{-1} [15–17]. As with PBHs and UCMHs, CMB μ distortion can also constrain non-Gaussianity, but the nature of the constraints on both characteristics is qualitatively different. First, CMB μ distortion is sensitive to the total power in fluctuations over a wider range of scales, while the objects are sensitive to the power only very near the scale of the object, $k \sim 1/R$. Second, μ distortion is sensitive to only the squeezed limit of the bispectrum rather than the total skewness. That is, this measurement can constrain the amplitude in the correlation of three momentum modes with $k_1 \ll k_2 \sim k_3$ [61,62]. If non-Gaussianity follows the local ansatz (which has a strong signal in the squeezed bispectrum), a futuristic CMB probe like PIXIE [63] could constrain $f_{\rm NL} \sim \mathcal{O}(10^3)$ for a scale-invariant power spectrum and $f_{\rm NL} \sim \mathcal{O}(10)$ for an enhanced small-scale power spectrum that saturates current bounds ($\mathcal{M}_3 \sim 0.04$ in both cases) [61]. As can be seen from the top two panels in Fig. 5, the μ-distortion bound would be about an order of magnitude better than the bounds from object counts at that scale if the non-Gaussianity is local and hierarchical. The bounds may be similar for non-Gaussianity with a stronger, feeder-type scaling and a local bispectrum. The μ-distortion probe is then quite complementary to the abundance of PBHs and UCMHs in terms of the range of scales probed and the relationship between the observables and the primordial fluctuations.

ACKNOWLEDGMENTS

APPENDIX: GENERALIZED SOURCE LIMITS ON UCMHS

Our analysis here actually includes a slightly improved form of the point source analysis of Ref. [11]. Following from Eq. (18) in Ref. [11], if there is a single UCMH

present somewhere in the Milky Way, we identify the probability P_1^{seen} of actually seeing that UCMH as

$$P_1^{\text{seen}} = x P_{d < d_{\text{obs}}, 1}^{\text{exists}} = x \frac{M_{d < d_{\text{obs}}}}{M_{\text{MW}}}. \qquad (\text{A1})$$

Here d is the true distance of the UCMH from Earth, d_{obs} is the maximum observable distance required for detection at C.L. x, $P_{d < d_{\text{obs}}, 1}^{\text{exists}}$ is the probability of the UCMH residing within distance d_{obs} of Earth, M_{MW} is the dark mass of the Milky Way and $M_{d < d_{\text{obs}}}$ is the dark mass within d_{obs}.

This leads to an improved version of Eq. (24) in Ref. [11], which gives the maximum allowed fraction of matter in UCMHs today:

$$f_{\text{max}} = \frac{f_\chi M_{\text{UCMH}}^0}{M_{\text{MW}}} \frac{\log(1 - y)}{\log\left(1 - x \frac{M_{d < d_{\text{obs}}}}{M_{\text{MW}}}\right)}. \qquad (\text{A2})$$

Here $f_\chi \equiv \rho_\chi / \rho_{\text{m}}$ is the fraction of matter in the Universe that is dark, M_{UCMH}^0 is the present-day UCMH mass, and y is the desired C.L. of the limit f_{max}. The relationship between β and f_{max} can be found in Eq. (21) of Ref. [11]. Equation (A2) is valid for all x, $y < 1$. For $1 - y \gg 1 - x$, it agrees with Eq. (24) in Ref. [11]; while both forms are valid for the values of x and y we use here ($1 - y = 5 \times 10^{-2}$, $1 - x = 6 \times 10^{-7}$, as in Ref. [11]), the latter breaks down as $y \to x$, so Eq. (A2) is more general.

---

[1] E. Komatsu *et al.*, Astrophys. J. Suppl. Ser. **192**, 18 (2011).

[2] J. Dunkley, R. Hlozek, J. Sievers, V. Acquaviva, P. Ade *et al.*, Astrophys. J. **739**, 52 (2011).

[3] R. Keisler, C. Reichardt, K. Aird, B. Benson, L. Bleem *et al.*, Astrophys. J. **743**, 28 (2011).

[4] B. A. Reid, W. J. Percival, D. J. Eisenstein, L. Verde, D. N. Spergel *et al.*, Mon. Not. R. Astron. Soc. **404**, 60 (2010).

[5] N. Sehgal, H. Trac, V. Acquaviva, P. A. Ade, P. Aguirre *et al.*, Astrophys. J. **732**, 44 (2011).

[6] P. McDonald *et al.* (SDSS Collaboration), Astrophys. J. Suppl. Ser. **163**, 80 (2006).

[7] B. J. Carr, Astrophys. J. **201**, 1 (1975).

[8] A. S. Josan, A. M. Green, and K. A. Malik, Phys. Rev. D **79**, 103520 (2009).

[9] B. Carr, K. Kohri, Y. Sendouda, and J. Yokoyama, Phys. Rev. D **81**, 104019 (2010).

[10] A. S. Josan and A. M. Green, Phys. Rev. D **82**, 083527 (2010).

[11] T. Bringmann, P. Scott, and Y. Akrami, Phys. Rev. D **85**, 125027 (2012).

[12] F. Li, A. L. Erickcek, and N. M. Law, Phys. Rev. D **86**, 043519 (2012).

[13] P. Scott and S. Sivertsson, Phys. Rev. Lett. **103**, 211301 (2009).

[14] J. S. Bullock and J. R. Primack, Phys. Rev. D **55**, 7423 (1997).

[15] W. Hu, D. Scott, and J. Silk, Astrophys. J. **430**, L5 (1994).

[16] J. Chluba and R. A. Sunyaev, Mon. Not. R. Astron. Soc. **419**, 1294 (2012).

[17] J. Chluba, A. L. Erickcek, and I. Ben-Dayan, Astrophys. J. **758**, 76 (2012).

[18] A. M. Green and A. R. Liddle, Phys. Rev. D **56**, 6166 (1997).

[19] P. Ivanov, Phys. Rev. D **57**, 7145 (1998).

[20] P. Pina Avelino, Phys. Rev. D **72**, 124004 (2005).

[21] J. C. Hidalgo, arXiv:0708.3875.

[22] P. Klimai and E. Bugaev, arXiv:1210.3262.

[23] C. T. Byrnes, E. J. Copeland, and A. M. Green, Phys. Rev. D **86**, 043512 (2012).

[24] D. S. Salopek and J. R. Bond, Phys. Rev. D **42**, 3936 (1990).

[25] S. Young and C. T. Byrnes, J. Cosmol. Astropart. Phys. 08 (2013) 052.

[26] K. Kohri, C.-M. Lin, and T. Matsuda, Phys. Rev. D **87**, 103527 (2013).

[27] W. H. Press and P. Schechter, Astrophys. J. **187**, 425 (1974).

[28] M. LoVerde, A. Miller, S. Shandera, and L. Verde, J. Cosmol. Astropart. Phys. 04 (2008) 014.

[29] S. Shandera, A. Mantz, D. Rapetti, and S. W. Allen, J. Cosmol. Astropart. Phys. 08 (2013) 004.

[30] J. C. Niemeyer and K. Jedamzik, Phys. Rev. D **59**, 124013 (1999).

[31] M. Ricotti and A. Gould, Astrophys. J. **707**, 979 (2009).

[32] V. V. Petrov, *Sums of Independent Random Variables* (Springer-Verlag, Berlin, 1975).

[33] V. V. Petrov, *Limit Theorems of Probability Theory: Sequences of Independent Random Variables*, Oxford Studies in Probability Vol. 4 (Clarendon, Oxford, New York, 1995).

[34] N. Barnaby and S. Shandera, J. Cosmol. Astropart. Phys. 01 (2012) 034.

[35] X. Chen and Y. Wang, J. Cosmol. Astropart. Phys. 04 (2010) 027.

[36] S. Blinnikov and R. Moessner, Astron. Astrophys. Suppl. Ser. **130**, 193 (1998).

[37] M. Dias and D. Seery, Phys. Rev. D **85**, 043519 (2012).

[38] S. Chongchitnan and J. Silk, Astrophys. J. **724**, 285 (2010).

[39] A. M. Green, A. R. Liddle, K. A. Malik, and M. Sasaki, Phys. Rev. D **70**, 041502 (2004).

[40] S. Chongchitnan and G. Efstathiou, J. Cosmol. Astropart. Phys. 01 (2007) 011.

[41] G. Tesic for the VERITAS Collaboration, J. Phys. Conf. Ser. **375**, 052024 (2012).

[42] P. A. R. Ade *et al.* (Planck Collaboration), arXiv:1303.5082.

[43] L. Alabidi, K. Kohri, M. Sasaki, and Y. Sendouda, J. Cosmol. Astropart. Phys. 09 (2012) 017.

[44] T. Bringmann, C. Kiefer, and D. Polarski, Phys. Rev. D **65**, 024008 (2001).

[45] E. W. Kolb and M. S. Turner, *The Early Universe*, Frontiers in Physics (Addison-Wesley, Reading, MA, 1990).

[46] V. Berezinsky, V. Dokuchaev, and Y. Eroshenko, Phys. Rev. D **68**, 103003 (2003).

[47] V. Berezinsky, V. Dokuchaev, and Y. Eroshenko, Phys. Rev. D **73**, 063504 (2006).

[48] V. Berezinsky, V. Dokuchaev, and Y. Eroshenko, J. Cosmol. Astropart. Phys. 07 (2007) 011.

[49] V. Berezinsky, V. Dokuchaev, and Y. Eroshenko, Phys. Rev. D **77**, 083519 (2008).

[50] E. Zackrisson, S. Asadi, K. Wiik, J. Jonsson, P. Scott, K. K. Datta, M. M. Friedrich, H. Jensen, J. Johansson, C.-E. Rydberg, and A. Sandberg, Mon. Not. R. Astron. Soc. **431**, 2172 (2013).

[51] B. C. Lacki and J. F. Beacom, Astrophys. J. **720**, L67 (2010).

[52] V. Berezinsky, V. Dokuchaev, Y. Eroshenko, M. Kachelriess, and M. A. Solberg, Phys. Rev. D **81**, 103529 (2010).

[53] V. Berezinsky, V. Dokuchaev, Y. Eroshenko, M. Kachelriess, and M. A. Solberg, Phys. Rev. D **81**, 103530 (2010).

[54] Y. Yang, L. Feng, X. Huang, X. Chen, T. Lu, and H. Zong, J. Cosmol. Astropart. Phys. 12 (2011) 020.

[55] Y. Yang, G. Yang, and H. Zong, Europhys. Lett. **101**, 69001 (2013).

[56] Y. Yang, G. Yang, and H. Zong, Phys. Rev. D **87**, 103525 (2013).

[57] D. Zhang, Mon. Not. R. Astron. Soc. **418**, 1850 (2011).

[58] Y. Yang, X. Huang, X. Chen, and H. Zong, Phys. Rev. D **84**, 043506 (2011).

[59] Y.-P. Yang, X. Chen, T. Lu, and H.-S. Zong, Eur. Phys. J. Plus **126**, 123 (2011).

[60] H.-S. Zechlin and D. Horns, J. Cosmol. Astropart. Phys. 11 (2012) 050.

[61] E. Pajer and M. Zaldarriaga, Phys. Rev. Lett. **109**, 021302 (2012).

[62] J. Ganc and E. Komatsu, Phys. Rev. D **86**, 023518 (2012).

[63] A. Kogut, D. Fixsen, D. Chuss, J. Dotson, E. Dwek *et al.*, J. Cosmol. Astropart. Phys. 07 (2011) 025.