# Constraining Single-Field Inflation with MegaMapper

Giovanni Cabass,[1, *] Mikhail M. Ivanov,[1, 2, †] Oliver H. E. Philcox,[3, 4, ‡] Marko Simonović,[5] and Matias Zaldarriaga[1]

[1]*School of Natural Sciences, Institute for Advanced Study, 1 Einstein Drive, Princeton, NJ 08540, USA*
[2]*NASA Hubble Fellowship Program Einstein Postdoctoral Fellow*
[3]*Center for Theoretical Physics, Department of Physics,*
*Columbia University, New York, NY 10027, USA*
[4]*Simons Society of Fellows, Simons Foundation, New York, NY 10010, USA*
[5]*Theoretical Physics Department, CERN, 1 Esplanade des Particules, Geneva 23, CH-1211, Switzerland*

We forecast the constraints on single-field inflation from the bispectrum of future high-redshift surveys such as MegaMapper. Considering non-local primordial non-Gaussianity (NLPNG), we find that current methods will yield constraints of order $\sigma(f_{\rm NL}^{\rm eq}) \approx 23$, $\sigma(f_{\rm NL}^{\rm orth}) \approx 12$ in a joint power-spectrum and bispectrum analysis, varying both nuisance parameters and cosmology, including a conservative range of scales. Fixing cosmological parameters and quadratic bias parameter relations, the limits tighten significantly to $\sigma(f_{\rm NL}^{\rm eq}) \approx 17$, $\sigma(f_{\rm NL}^{\rm orth}) \approx 8$. These compare favorably with the forecasted bounds from CMB-S4: $\sigma(f_{\rm NL}^{\rm eq}) \approx 21$, $\sigma(f_{\rm NL}^{\rm orth}) \approx 9$, with a combined constraint of $\sigma(f_{\rm NL}^{\rm eq}) \approx 14$, $\sigma(f_{\rm NL}^{\rm orth}) \approx 7$; this weakens only slightly if one instead combines with data from the Simons Observatory. We additionally perform a range of Fisher analyses for the error, forecasting the dependence on nuisance parameter marginalization, scale cuts, and survey strategy. Lack of knowledge of bias and counterterm parameters is found to significantly limit the information content; this could be ameliorated by tight simulation-based priors on the nuisance parameters. The error-bars decrease significantly as the number of observed galaxies and survey depth is increased: as expected, deep dense surveys are the most constraining, though it will be difficult to reach $\sigma(f_{\rm NL}) \approx 1$ with current methods. The NLPNG constraints will tighten further with improved theoretical models (incorporating higher-loop corrections and improved understanding of nuisance parameters), as well as the inclusion of additional higher-order statistics.

*Motivation* The next generation of large-scale structure surveys will yield unprecedented measurements of the $z > 2$ Universe. The combination of huge volumes and high-redshifts will enable proposed surveys, such as MegaMapper [1, 2], MSE [3], GAUSS [4], SpecTel [5, 6] and Rubin [7], to constrain primordial physics by measuring a vast array of linear modes, improving on existing surveys by several orders of magnitude. A crucial question is the following: what can we hope to learn from this tranche of new data?

With precise data comes high-resolution measurements of the galaxy power spectrum, stretching to comparatively small (but yet still linear) scales. As with existing surveys, this can be used to place strong bounds on the cosmological model ($\nu\Lambda$CDM), via full-shape analyses [e.g., 8, 9], constraining parameters such as the matter density, primordial power spectrum amplitude, and Hubble constant. Such quantities are already relatively tightly constrained by the cosmic microwave background (CMB, [10]), however, thus it is interesting to shift our attention to non-standard parameters, in particular those set by early Universe physics. As shown in [11], by analyzing the high-redshift spectrum, we can hope to obtain strong constraints on particle physics, such as the mass of the neutrino and a number of proposed particles, such as axions, as well as energy deposition in the early Universe

(via the $N_{\rm eff}$ parameter) [e.g., 12–14]. Furthermore, we can directly probe inflation by considering the spectral tilt, $n_s$, and its running, as well as the scale-dependent bias induced by local primordial non-Gaussianity ($f_{\rm NL}^{\rm loc}$), which is a key signature of multi-field inflation [e.g., 15, 16].

By looking beyond the galaxy power spectrum, we can constrain a variety of other inflationary features. In particular, interactions in inflation and non-standard vacua can give rise to non-local primordial non-Gaussianity (NLPNG), whose primordial bispectra can be well described by the 'equilateral' and 'orthogonal' templates with amplitudes $f_{\rm NL}^{\rm eq}$ and $f_{\rm NL}^{\rm orth}$ [e.g., 17]. Careful analysis of the galaxy bispectra can yield constraints on these parameters, as demonstrated in [18–20] for current data. Of course, the rich landscape of inflation is not limited to two templates: we may utilize the galaxy bispectrum to constrain a wealth of models, including massive spinning particles, and yet more can be learnt from the galaxy trispectrum [e.g., 21–23]. For now, the constraining power on such parameters is dominated by the CMB [e.g., 24, 25]; however, with the advent of stage-five spectrosopic surveys, the attention will shift to large scale structure (LSS). In this work, we forecast how well proposed surveys such as MegaMapper can hope to constrain inflationary signals, via the NLPNG parameters, optionally in conjunction with CMB observations.

*Set-Up* To forecast the efficacy of high-redshift in constraining single-field inflationary parameters, we follow a similar procedure to [11], concentrating on the proposed MegaMapper experiment outlined in [26]. Unlike previous work, we include a full treatment of the bispectrum, mod-

* gcabass@ias.edu
† ivanov@ias.edu
‡ Corresponding Author; ohep2@cantab.ac.uk

eled in conjunction with the power spectrum to minimize parameter degeneracies. The fiducial experiment includes $\approx 40$ million galaxies in a redshift range $2 < z < 5$, which we divide into four contiguous bins with number density and redshift defined by Tab. 1 of [26]. In each redshift bin, we compute a fiducial power spectrum and bispectrum using the parameters of [11], in particular tidal biases set by the coevolution model [27] and an (optimistic) fingers-of-God (FoG) dispersion $\sigma_v = 100\,\mathrm{km\,s^{-1}}$. For the quadratic bias, we use the fitting formula of [28], which is significantly more accurate than the coevolution prediction for highly biased samples.

We utilize the one-loop power spectrum and the tree-level bispectrum model summarized in [29–31], depending on the following nuisance parameters:

$$\{b_1, b_2, b_{\mathcal{G}_2}, b_{\Gamma_3}, c_0, c_2, c_4, \tilde{c}, c_1, P_{\mathrm{shot}}, B_{\mathrm{shot}}, a_0, a_2, b_\phi\} \quad (1)$$

describing linear, quadratic and tidal bias, five counterterms, four stochasticity parameters, and non-Gaussian bias. These are subject to wide Gaussian priors following [18, 32]. We do not include one-loop corrections to the bispectrum; whilst these are available for $\Lambda$CDM [33, 34], the loops involving $f_{\mathrm{NL}}$ have yet to be self-consistently included. We include power spectrum multipoles up to $\ell = 4$ and the bispectrum monopole, noting that constraints may tighten somewhat if we include higher-order bispectrum multipoles [35]. We additionally include the $Q_0$ statistic of [36] (see also [37, 38]), which is a proxy for the real-space power spectrum, and allows extraction of information beyond the usual fingers-of-God limits. We vary the following cosmological parameters:

$$\{h, \omega_{\mathrm{cdm}}, \log 10^{10} A_s, f_{\mathrm{NL}}^{\mathrm{eq}}, f_{\mathrm{NL}}^{\mathrm{orth}}\}, \quad (2)$$

with wide flat priors, where the last two parameters control the single-field inflation model. We do not include the multi-field parameter, $f_{\mathrm{NL}}^{\mathrm{loc}}$, since the information content on this is dominated by the power spectrum [19] and is not strongly degenerate with the NLPNG amplitudes. We additionally fix the spectral tilt, $n_s$, since it will be precisely constrained by future CMB surveys. If one instead marginalizes over $n_s$, the $f_{\mathrm{NL}}$ constraints degrade by $< 5\%$, as verified in MCMC analyses (similar to the conclusions of [18] for the BOSS survey).

Following [11], we fix the minimum power spectrum wavenumber to $k_{\mathrm{min}} = \max[0.003h\,\mathrm{Mpc^{-1}}, 2\pi/V^{1/3}]$, and use bins of width $\Delta k = 0.005h\,\mathrm{Mpc^{-1}}$. We fix the maximum scale by asserting that the relative size of the FoG term (which usually dominates the theoretical error) is the same as for the BOSS analysis with $k_{\mathrm{max}}^P \approx 0.17h\,\mathrm{Mpc^{-1}}$ [18]. This is a relatively conservative choice and is equivalent to demanding that the FoG contributions are at most 10% of the tree-level theory. The $Q_0$ statistic is not affected by FoG, thus we consider modes up to the Zel'dovich velocity dispersion scale, fixing $k_{\mathrm{max}}^Q = \left[\int dq P_{\mathrm{lin}}(q)/(6\pi^2)\right]^{-1/2}$, again fol-

| Experiment | $\sigma(f_{\mathrm{NL}}^{\mathrm{eq}})$ | $\sigma(f_{\mathrm{NL}}^{\mathrm{orth}})$ |
|---|---|---|
| MegaMapper - A | 23 | 10 |
| MegaMapper - B | 22 | 10 |
| MegaMapper - C | 17 | 8 |
| Planck 2018 | 47 | 24 |
| Simons Observatory (SO) | 27 | 14 |
| CMB-S4 | 21 | 9 |
| MegaMapper + SO | 16 | 8 |
| MegaMapper + CMB-S4 | 14 | 7 |

TABLE I. 68% Constraints on the non-local primordial non-Gaussianity amplitudes shown in Fig. 1 & 2. The three MegaMapper analyses are (A) free cosmology and bias parameters, (B) fixed cosmology and free bias parameters, (C) fixed cosmology and fixed quadratic bias relations. We additionally quote the *Planck* 2018 constraints, as well as the forecasts for the Simons Observatory and CMB-S4. The final entry gives the joint constraints from MegaMapper and future CMB experiments, assuming that the latter datasets fix the cosmological parameters.

lowing [11].[1] For the bispectrum, we consider broader bins of $\Delta k = 0.01h\,\mathrm{Mpc^{-1}}$ (since the desired signal is relatively smooth), and fix $k_{\mathrm{min}}^B = 0.01h\,\mathrm{Mpc^{-1}}$, and the maximum by asserting that the FoG contribution has similar contributions to those in the BOSS analyis, which here requires it to be less than 2% of the tree-level theory, yielding $k_{\mathrm{max}}^B \approx 0.20h\,\mathrm{Mpc^{-1}}$ at $z = 3$. This is lower than for the power spectrum, since we include only tree-level terms in the bispectrum model (though could extend beyond this via the approaches of [33, 34]).[2] Finally, we perform the forecast via an MCMC analysis of the fiducial $f_{\mathrm{NL}}^{\mathrm{eq}} = f_{\mathrm{NL}}^{\mathrm{orth}} = 0$ spectra, using the CLASS-PT code [42], the MONTEPYTHON sampler [43] and the public likelihoods described in [32],[3] assuming a Gaussian likelihood.

In this *Letter*, we assume a diagonal covariance matrix for the power spectrum multipoles, as in [44]. In principle, late-time non-Gaussianity leads to correlations both within and between correlators [45], which could somewhat reduce the detection significances (as shown in line-intensity-mapping contexts in [46]). For spectro-

---

[1] This is a conservative choice, since $\left[\int dq P_{\mathrm{lin}}(q)/(6\pi^2)\right]^{-1/2}$ in question is the characteristic scale of the displacement field, the bulk of which is accounted for via infrared resummation [39–41]. The true $k_{\mathrm{max}}$ relevant for $Q_0$ can be significantly larger than this (up to the non-linear scale) as discussed in [e.g., 36]. Since $Q_0$ does not contribute strongly to NLPNG constraints, this choice is not strongly relevant for our analysis.

[2] Imposing stronger bounds on the FoG contribution reduces the detection significance somewhat: we find a $\approx 30\%$ increase in $\sigma(f_{\mathrm{NL}})$ when restricting the terms to be 5% and 1% of the power spectrum and bispectrum respectively. This is equivalent to increasing $\sigma_v$ by a factor of $\sqrt{2}$. Further discussion of the dependence on $k_{\mathrm{max}}$ can be found below.
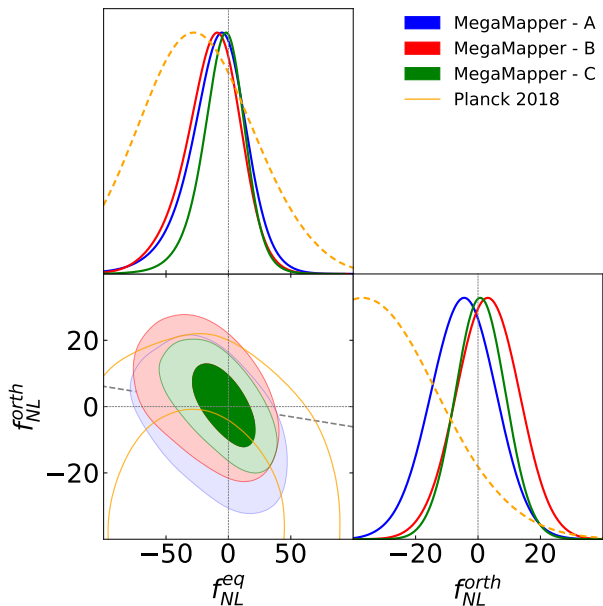
[3] Available at github.com/oliverphilcox/full_shape_likelihoods.

FIG. 1. Forecasted constraints on the non-local primordial non-Gaussianity amplitudes, $f_{\rm NL}^{\rm eq}$, $f_{\rm NL}^{\rm orth}$ from the proposed MegaMapper high-redshift experiment [2], alongside measurements from *Planck* [25]. The MegaMapper forecasts are computed via MCMC using the parameters described in the text, including both the one-loop power spectrum and tree-level bispectrum. We consider three types of forecast: A (blue), varying both cosmological and nuisance parameters; B (red) fixing cosmological parameters; C (green) fixing also quadratic bias relations. Numerical constraints are given in Tab. I, and we show the linear relationship between $f_{\rm NL}^{\rm orth}$ and $f_{\rm NL}^{\rm eq}$ relevant for DBI inflation as a dashed line [18, 48].

scopic surveys, the full impacts of these on $f_{\rm NL}$ constraints have yet to be explored; however, given our conservative choices of $k_{\rm max}$ and the high shot-noise of MegaMapper-like surveys, we expect this to be small, given the lack of importance of non-Gaussian covariance in power spectrum [47] and large-scale bispectrum [31–33] $\Lambda$CDM analyses.

*Fiducial Constraints* We consider three characteristic cases for the inflationary forecasts, similar to [18]. Firstly, we assume no knowledge of cosmology or bias, varying all parameters in (1) & (2). Secondly, we assume that the background $\nu\Lambda$CDM cosmology is known (for example from the power spectrum, or from external data e.g., CMB-S4), and vary only biases. Thirdly, we additionally fix the quadratic bias parameter relations ($b_2(b_1)$ and $b_{\mathcal{G}_2}(b_1)$); these generate shapes with significant correlations with NLPNG [e.g., 49], and could be potentially fixed via tight simulation-derived priors.[4]

FIG. 2. Comparison of MegaMapper (blue) constraints on $f_{\rm NL}^{\rm eq}$ and $f_{\rm NL}^{\rm orth}$ with those from CMB-S4 [53] (red). The joint constraint (green) is obtained by combining the fixed cosmology MegaMapper contour with the CMB-S4 forecasted posterior.

The corresponding constraints on $f_{\rm NL}$ are given in Fig. 1 and Tab. I. In the conservative case, we find $\sigma(f_{\rm NL}^{\rm eq}) \approx 23$ and $\sigma(f_{\rm NL}^{\rm orth}) \approx 10$; neither constraints are found to tighten significantly when we fix the fiducial cosmology (and thus the NLPNG templates), but we see a modest ($\approx 30\%$) improvement by fixing the quadratic bias relations, and thus reducing the relevant degeneracies. In the most optimistic case, we find $\sigma(f_{\rm NL}^{\rm eq}) \approx 17$, $\sigma(f_{\rm NL}^{\rm orth}) \approx 8$, though we note that the relevant bias parameter relation priors may be difficult to obtain in practice (since the data already constrain the tidal bias to $\approx 10\%$, for example). Note that the impact of bias parameter priors is weaker here than in [18] due to the larger $k$ range, and thus greater internal degeneracy breaking.

To place our results in context, we may compare them to current limits from the CMB and LSS. The BOSS data constrain $f_{\rm NL}^{\rm eq} = 260 \pm 300$, $f_{\rm NL}^{\rm orth} = -23 \pm 120$ [18] (with fixed biases), which are considerably weaker than the CMB bounds of $f_{\rm NL}^{\rm eq} = -26 \pm 47$, $f_{\rm NL}^{\rm orth} = -38 \pm 24$ from *Planck* 2018 [25]. The next generation of CMB data is expected to tighten these bounds somewhat (cf. Fig. 2), with $\sigma(f_{\rm NL}^{\rm eq}) = 27$, $\sigma(f_{\rm NL}^{\rm orth}) = 14$ expected from the Simons Observatory (SO) [54], and $\sigma(f_{\rm NL}^{\rm eq}) = 21$, $\sigma(f_{\rm NL}^{\rm orth}) = 9$ from CMB-S4 [53, 55, 56]. This improvement is relatively small given that *Planck* is already cosmic variance limited for

---

[4] One may also place priors on the non-Gaussian biases, such as $b_\phi$: this has limited effect for NLPNG (but is crucial for $f_{\rm NL}^{\rm loc}$, see [50, 51]), as seen by the similar constraints from MCMC and Fisher forecasts in the below, noting that the latter require fixed
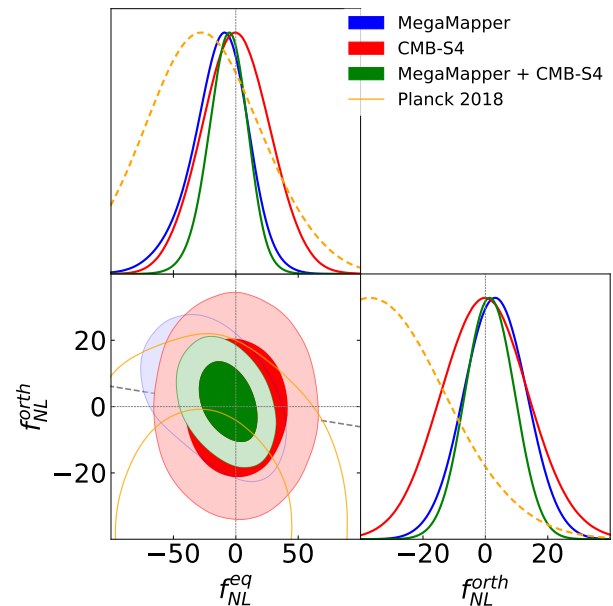
$b_\phi$ (as it appears only proportional to $f_{\rm NL}$). This is consistent with the forecast of [52], which found $\sigma(f_{\rm NL}^{\rm equil}) \approx 500$ from the power spectrum scale-dependent bias alone, for a MegaMapper-type survey.

the large-scale temperature modes, and there is little gain from small scales due to the numerous secondary contributions.

The forecasts above improve upon LSS constraints by around an order of magnitude, and, even in the conservative case, where we vary cosmology and all bias parameters, are significantly tighter than those of *Planck*. This matches expectations, since the survey volume of MegaMapper is roughly $30\times$ that of BOSS, and we work at higher redshift, facilitating larger $k_{\max}$. If the $\nu\Lambda\mathrm{CDM}$ parameters are known (or at least highly constrained), MegaMapper can provide competitive constraints to CMB-S4, and somewhat tighter than the SO baseline. In combination, the two yield stronger constraints still, as shown in Fig. 2, with bounds of $\sigma(f_{\mathrm{NL}}^{\mathrm{eq}}) \approx 14$, $\sigma(f_{\mathrm{NL}}) \approx 7$ expected from CMB-S4, broadening to $\sigma(f_{\mathrm{NL}}^{\mathrm{eq}}) \approx 16$, $\sigma(f_{\mathrm{NL}}) \approx 8$ with SO. We may also compare the results to previous simplified forecasts, in particular [26], based on [57]. This obtained $\sigma(f_{\mathrm{NL}}^{\mathrm{eq}}) = 40$, $\sigma(f_{\mathrm{NL}}^{\mathrm{orth}}) = 9$, varying all relevant parameters: these agree with our fiducial analysis ('A') to $\approx 40\%$; a factor certainly appropriate for Fisher forecasts. Of course, these results depend on the various modeling choices shown above, in particular the $k$ ranges and fiducial bias and FoG parameters. However, our wavenumber limits may be regarded as conservative, and the modelling can likely be improved further by the addition of one-loop terms, though it is uncertain whether this will significantly impact parameter constraints [33].

*Impact of Analysis Choices*   To better understand the above results, it is useful to consider the dependence of the NLPNG constraints on the maximum wavenumber used in the analysis, $k_{\max}^{P,B}$, and the assumed priors on nuisance parameters. To this end, we perform a Fisher analysis based on the above methodology, forecasting $\sigma(f_{\mathrm{NL}}^{\mathrm{eq}})$ and $\sigma(f_{\mathrm{NL}}^{\mathrm{orth}})$ for the fiducial MegaMapper survey, using the same galaxy distribution parameters parameters as before. For the fiducial MegaMapper set-up, the Fisher forecast matches the full MCMC constraints to $\approx 10\%$, at significantly lower computational cost.

Fisher analyses are performed for nine values of $k_{\max}^{B} \in [0.05, 0.5]h\,\mathrm{Mpc}^{-1}$, using the same scale-cut for each redshift bin and fixing $k_{\max}^{P} = 2k_{\max}^{B}$, motivated by the above physical limits. In each case, we consider two scenarios: one in which all nuisance parameters (encompassing bias, stochasticity, and counterterms) are varied, akin to method 'B' in the MCMC forecasts, and one in which only $f_{\mathrm{NL}}^{\mathrm{eq}}$ and $f_{\mathrm{NL}}^{\mathrm{orth}}$ are free, with all other parameters fixed to their fiducial values. Whilst the extent to which future analyses can strongly bound nuisance parameters is unclear, this provides a practical bound to the precision of NLPNG measurements, in the limit of perfect understanding of the UV physics of galaxy formation, as well as nonlinear gravitational collapse.

The results are shown in Fig. 3, and indicate considerable dependence of the error-bar on $k_{\max}$. This is unsurprising, since the signal-to-noise of the statistics increases considerably with $k$ (up to the shot-noise domi-
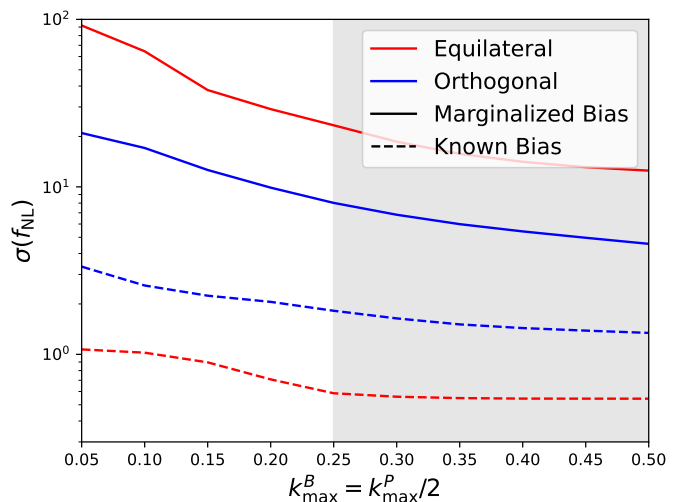


FIG. 3. Impact of scale cuts and nuisance parameter marginalization on the $f_{\mathrm{NL}}$ constraints obtained from an idealized MegaMapper-like survey. Fisher forecasts are shown for both equilateral (red) and orthogonal (blue) primordial non-Gaussianity, marginalizing over all NLPNG and nuisance parameters (full lines) or only the NLPNG parameters (dashed lines), jointly varying both the maximum power spectrum and bispectrum wavenumber across all redshift slices. We show the fiducial MegaMapper constraints as horizontal dotted lines (cf. Tab. I), noting that these use a different $k_{\max}$ for each redshift bin, and show in grey the rough range of scales for which higher-loop effects become important. The plot indicates that $f_{\mathrm{NL}}$ constraints are strongly limited by our lack of knowledge of nuisance parameters, and show considerable dependence on $k_{\max}$. Whilst the known-bias case represents an optimistic lower bound on the parameter error, we caution that higher-loop terms (both in hydrodynamic and gravitational physics) will be required to extend the modelling to large $k_{\max}$, and that precise knowledge of nuisance parameters will require accurate hydrodynamic simulations coupled with good understanding of observational effects.

nated regime), allowing breaking of parameter degeneracies. At large $k_{\max}$, we caution that additional loops (and thus nuisance parameters) will be required for accurate modelling, which will somewhat temper the reduction in $\sigma(f_{\mathrm{NL}})$; this idealized plot indicates the utility of extending the modelling further however. We also find a considerable improvement in constraining power (at least an order of magnitude) when the nuisance parameters are known, particularly for equilateral non-Gaussianity. This implies that the majority of the signal-to-noise goes into constraining parameters such as higher-order bias and shot-noise rather than NLPNG directly (particularly at low $k_{\max}$, whereupon information from the power spectrum dominates), and motivates further study into simulation-based priors on such effects. If one had perfect knowledge of the nuisance parameters and non-linear physics (which will likely be challenging, even in the far future), constraints of $\sigma(f_{\mathrm{NL}}) = \mathcal{O}(1)$ could be obtained from a MegaMapper-like survey.

*Survey Design* Finally, it is interesting to consider what types of survey could yield the best constraints on NLPNG parameters. This can also be probed using Fisher forecasts, here considering an idealized Lyman-Break galaxy survey (LBG, modelled using the $n(z)$ of [26]), and computing $\sigma(f_{\mathrm{NL}})$ as a function of the total number of observed galaxies, $N_g$, and the maximum redshift, $z_{\max}$ (using four redshift bins, as before).

The results are shown in Fig. 4, displaying 30 (fixed-cosmology) forecasts, with $z_{\max} \in [3, 5]$, $\log_{10} N_g \in [7, 9]$, and an additional cosmic-variance dominated sample (with the idealized limit of $N_g \to \infty$).[5] At small $N_g$, we find little gain in pushing to large volumes due to the prohibitively low sample density; however, as $N_g$ increases, we see significant reduction in $\sigma(f_{\mathrm{NL}})$ as the survey volume increases for a fixed number of targets. This occurs due to the trade-off between low cosmic variance (scaling as the root of the survey volume, $V^{1/2}$) and low shot-noise (depending on $N_g/V$). In contrast, increasing the total number of galaxies observed (at fixed $V$) yields much increased sensitivity at all $z_{\max}$, with the cosmic variance limit yielding $\sigma(f_{\mathrm{NL}}^{\mathrm{eq}}) \approx 10$, $\sigma(f_{\mathrm{NL}}^{\mathrm{orth}}) \approx 5$. Even with futuristic surveys such as MegaMapper, it seems that we are far from this limit, and are thus not shot-noise dominated on all scales of interest; approaching would require the inclusion of additional galaxy samples however.

All of our forecasts indicate that future high-redshift surveys will yield strong constraints on NLPNG. These

are likely to surpass those from the CMB, and are made possible by the analysis of higher-order statistics. This is particularly true for the non-local non-Gaussianity parameters: whilst similar forecasts to the above can be performed for $f_{\mathrm{NL}}^{\mathrm{loc}}$, we expect that the bulk of the signal-to-noise will come from the power spectrum, with the bispectrum adding only $\approx 30\%$ [19]. Our results, however, have strong dependence on the $k$ ranges assumed, which themselves depend on (poorly understood) MegaMapper sample parameters. To obtain a better understanding of the capabilities of future surveys, it will be vital to generate accurate mock catalogs appropriate for the high-redshift regime. We close by stressing that the three $f_{\mathrm{NL}}$ parameters do not paint a full picture of inflation: MegaMapper, or some similar experiment, will allow a wide variety of models to be probed, including ghosts, colliders, and axions.

---

[1] D. J. Schlegel *et al.*, (2019), arXiv:1907.11171 [astro-ph.IM].

[2] D. J. Schlegel *et al.*, (2022), arXiv:2209.04322 [astro-ph.IM].

[3] C. Babusiaux *et al.* (MSE Science Team), (2019), arXiv:1904.04907 [astro-ph.IM].

[4] A. Blanchard *et al.*, Exper. Astron. **51**, 1623 (2021), arXiv:2102.03931 [astro-ph.CO].

[5] R. Ellis *et al.*, (2019), arXiv:1907.06797 [astro-ph.IM].

[6] K. Dawson, J. Frieman, K. Heitmann, B. Jain, S. Kahn, R. Mandelbaum, S. Perlmutter, and A. Slosar, (2018), arXiv:1802.07216 [astro-ph.CO].

[7] D. Spergel *et al.*, (2015), arXiv:1503.03757 [astro-ph.IM].

[8] M. M. Ivanov, M. Simonović, and M. Zaldarriaga, JCAP **2020**, 042 (2020), arXiv:1909.05277 [astro-ph.CO].

[9] G. d'Amico, J. Gleyzes, N. Kokron, K. Markovic, L. Senatore, P. Zhang, F. Beutler, and H. Gil-Marín, JCAP **2020**, 005 (2020), arXiv:1909.05271 [astro-ph.CO].

[10] Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, *et al.*, A&A **641**, A6 (2020), arXiv:1807.06209 [astro-ph.CO].

[11] N. Sailer, E. Castorina, S. Ferraro, and M. White, JCAP **12**, 049 (2021), arXiv:2106.09713 [astro-ph.CO].

[12] D. Baumann and D. Green, Phys. Rev. D **85**, 103520 (2012), arXiv:1109.0292 [hep-th].

[13] D. Green, Y. Guo, and B. Wallisch, JCAP **02**, 019 (2022), arXiv:2109.12088 [astro-ph.CO].

[14] F. Beutler, M. Biagetti, D. Green, A. c. v. Slosar, and B. Wallisch, Phys. Rev. Research **1**, 033209 (2019).

[15] V. Desjacques, U. Seljak, and I. Iliev, Mon. Not. Roy. Astron. Soc. **396**, 85 (2009), arXiv:0811.2748 [astro-ph].

[16] D. Jeong and E. Komatsu, Astrophys. J. **703**, 1230 (2009), arXiv:0904.0497 [astro-ph.CO].

[17] L. Senatore, K. M. Smith, and M. Zaldarriaga, JCAP **01**, 028 (2010), arXiv:0905.3746 [astro-ph.CO].

[18] G. Cabass, M. M. Ivanov, O. H. E. Philcox, M. Simonović, and M. Zaldarriaga, Phys. Rev. Lett. **129**, 021301 (2022), arXiv:2201.07238 [astro-ph.CO].

[19] G. Cabass, M. M. Ivanov, O. H. E. Philcox, M. Simonović, and M. Zaldarriaga, Phys. Rev. D **106**, 043506 (2022), arXiv:2204.01781 [astro-ph.CO].

[20] G. D'Amico, M. Lewandowski, L. Senatore, and P. Zhang, (2022), arXiv:2201.11518 [astro-ph.CO].

[21] N. Arkani-Hamed and J. Maldacena, arXiv e-prints , arXiv:1503.08043 (2015), arXiv:1503.08043 [hep-th].

[22] A. Moradinezhad Dizgah, H. Lee, J. B. Muñoz, and C. Dvorkin, JCAP **05**, 013 (2018), arXiv:1801.07265 [astro-ph.CO].

---

[5] In practice, there are only a finite number of LBGs in the Universe, thus $\log_{10} N_g \lesssim 8$. Here, we allow arbitrarily large $N_g$ to emulate other galaxy populations one may wish to study.
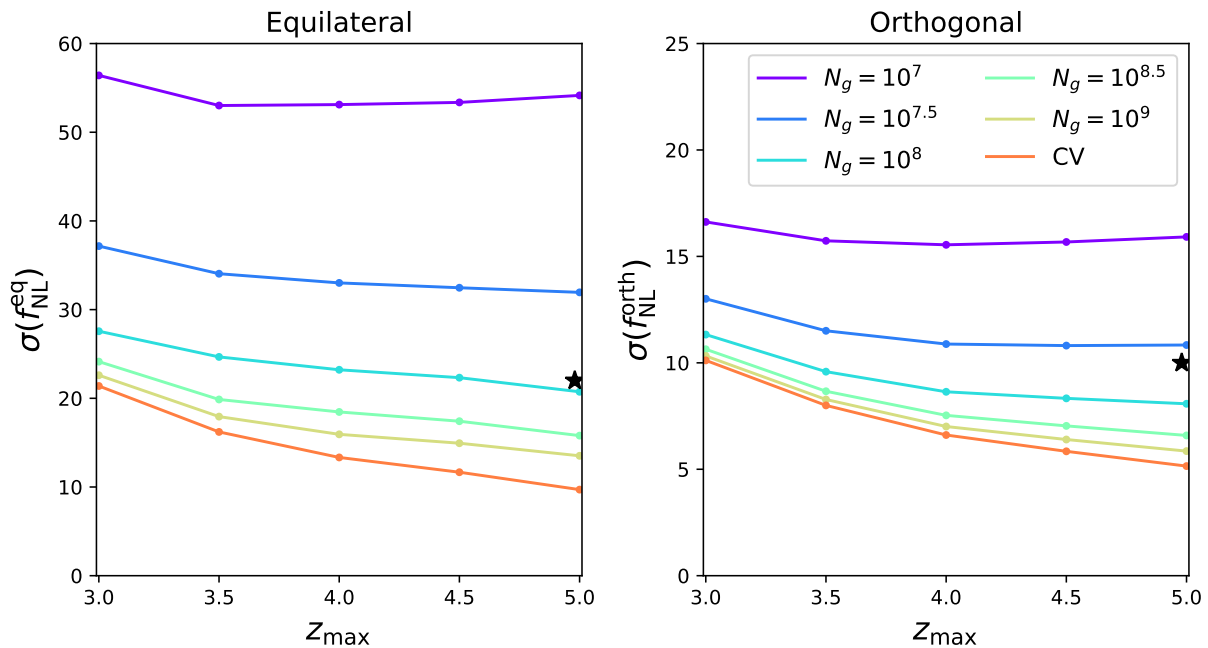
FIG. 4. Fisher forecast of the error-bar on NLPNG parameters as a function of the maximum survey redshift ($z_{\mathrm{max}}$) and the total number of galaxies ($N_g$). 30 results are shown, with the density indicated by the caption, where 'CV' indicates an idealized cosmic-variance limited sample. The MCMC forecasts appropriate for a MegaMapper-like experiment (Tab. I) are shown as black stars. In each case, we fix cosmology, but allow all bias and nuisance parameters to vary freely, in four redshift bins. Increasing the survey volume at fixed $N_g$ yields a slight increase in precision, with a much larger one seen by increasing the number of target galaxies.

[23] G. Cabass, M. M. Ivanov, and O. H. E. Philcox, (2022), arXiv:2210.16320 [astro-ph.CO].

[24] Planck Collaboration, Y. Akrami, F. Arroja, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, *et al.*, A&A **641**, A10 (2020), arXiv:1807.06211 [astro-ph.CO].

[25] Planck Collaboration, Y. Akrami, F. Arroja, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, *et al.*, A&A **641**, A9 (2020), arXiv:1905.05697 [astro-ph.CO].

[26] S. Ferraro *et al.*, (2019), arXiv:1903.09208 [astro-ph.CO].

[27] V. Desjacques, D. Jeong, and F. Schmidt, Phys. Rept. **733**, 1 (2018), arXiv:1611.09787 [astro-ph.CO].

[28] T. Lazeyras, C. Wagner, T. Baldauf, and F. Schmidt, JCAP **02**, 018 (2016), arXiv:1511.01096 [astro-ph.CO].

[29] M. M. Ivanov, M. Simonović, and M. Zaldarriaga, JCAP **05**, 042 (2020), arXiv:1909.05277 [astro-ph.CO].

[30] A. Chudaykin, K. Dolgikh, and M. M. Ivanov, Phys. Rev. D **103**, 023507 (2021), arXiv:2009.10106 [astro-ph.CO].

[31] M. M. Ivanov, O. H. E. Philcox, T. Nishimichi, M. Simonović, M. Takada, and M. Zaldarriaga, Phys. Rev. D **105**, 063512 (2022), arXiv:2110.10161 [astro-ph.CO].

[32] O. H. E. Philcox and M. M. Ivanov, arXiv e-prints , arXiv:2112.04515 (2021), arXiv:2112.04515 [astro-ph.CO].

[33] O. H. E. Philcox, M. M. Ivanov, G. Cabass, M. Simonović, M. Zaldarriaga, and T. Nishimichi, Phys. Rev. D **106**, 043530 (2022), arXiv:2206.02800 [astro-ph.CO].

[34] G. D'Amico, Y. Donath, M. Lewandowski, L. Senatore, and P. Zhang, (2022), arXiv:2206.08327 [astro-ph.CO].

[35] M. M. Ivanov, O. H. E. Philcox, G. Cabass, T. Nishimichi, M. Simonović, and M. Zaldarriaga, (2023),

arXiv:2302.04414 [astro-ph.CO].

[36] M. M. Ivanov, O. H. E. Philcox, M. Simonović, M. Zaldarriaga, T. Nischimichi, and M. Takada, Phys. Rev. D **105**, 043531 (2022), arXiv:2110.00006 [astro-ph.CO].

[37] R. Scoccimarro, Phys. Rev. D **70**, 083007 (2004), arXiv:astro-ph/0407214.

[38] G. D'Amico, L. Senatore, P. Zhang, and T. Nishimichi, (2021), arXiv:2110.00016 [astro-ph.CO].

[39] D. Blas, M. Garny, M. M. Ivanov, and S. Sibiryakov, JCAP **07**, 028 (2016), arXiv:1605.02149 [astro-ph.CO].

[40] M. M. Ivanov and S. Sibiryakov, JCAP **07**, 053 (2018), arXiv:1804.05080 [astro-ph.CO].

[41] L. Senatore and M. Zaldarriaga, JCAP **02**, 013 (2015), arXiv:1404.5954 [astro-ph.CO].

[42] A. Chudaykin, M. M. Ivanov, O. H. E. Philcox, and M. Simonović, Phys. Rev. D **102**, 063533 (2020), arXiv:2004.10607 [astro-ph.CO].

[43] T. Brinckmann and J. Lesgourgues, Phys. Dark Univ. **24**, 100260 (2019), arXiv:1804.07261 [astro-ph.CO].

[44] A. Chudaykin and M. M. Ivanov, JCAP **2019**, 034 (2019), arXiv:1907.06666 [astro-ph.CO].

[45] N. S. Sugiyama, S. Saito, F. Beutler, and H.-J. Seo, Mon. Not. Roy. Astron. Soc. **497**, 1684 (2020), arXiv:1908.06234 [astro-ph.CO].

[46] T. Flöss, M. Biagetti, and P. D. Meerburg, (2022), arXiv:2206.10458 [astro-ph.CO].

[47] D. Wadekar, M. M. Ivanov, and R. Scoccimarro, Phys. Rev. D **102**, 123521 (2020), arXiv:2009.00622 [astro-ph.CO].

[48] M. Alishahiha, E. Silverstein, and D. Tong, Phys. Rev. D **70**, 123505 (2004), arXiv:hep-th/0404084.

[49] D. Baumann and D. Green, JCAP **08**, 061 (2022), arXiv:2112.14645 [astro-ph.CO].

[50] A. Barreira, JCAP **12**, 031 (2020), arXiv:2009.06622 [astro-ph.CO].

[51] A. Barreira, JCAP **11**, 013 (2022), arXiv:2205.05673 [astro-ph.CO].

[52] J. Gleyzes, R. de Putter, D. Green, and O. Doré, JCAP **04**, 002 (2017), arXiv:1612.06366 [astro-ph.CO].

[53] K. N. Abazajian *et al.* (CMB-S4), (2016), arXiv:1610.02743 [astro-ph.CO].

[54] Simons Observatory Collaboration, P. Ade, J. Aguirre, Z. Ahmed, S. Aiola, A. Ali, D. Alonso, M. A. Alvarez, K. Arnold, P. Ashton, *et al.*, JCAP **2019**, 056 (2019), arXiv:1808.07445 [astro-ph.CO].

[55] K. Abazajian *et al.*, Bull. Am. Astron. Soc. **51**, 209 (2019), arXiv:1908.01062 [astro-ph.IM].

[56] K. Abazajian *et al.* (CMB-S4), (2022), arXiv:2203.08024 [astro-ph.CO].

[57] D. Karagiannis, A. Lazanu, M. Liguori, A. Raccanelli, N. Bartolo, and L. Verde, Mon. Not. Roy. Astron. Soc. **478**, 1341 (2018), arXiv:1801.09280 [astro-ph.CO].