

**Forecasting Constraints on Primordial  
Non-Gaussianity: A Tomographic Study with  
SPHEREx**

*Submitted in partial fulfillment of the requirements  
for the Physics Honors Program*

by

**Nathan Olson**

*Department of Physics  
University of California, Berkeley*

**May 2025**

**Thesis Advisor Signature:**



**Date:** 5/12/25

# 1 Background

## 1.1 Primordial Non-Gaussianity

The measurement of primordial non-Gaussianity (PNG) is one of the key targets for upcoming cosmological surveys as detecting it would provide invaluable information about the physics of the early universe. Cosmic inflation is the leading theory for the rapid exponential expansion of the cosmos in its first tiny fraction of a second, generating nearly Gaussian fluctuations in the matter density field that grew through gravitational instability into the Large-Scale Structure (LSS) that we now observe [1]. Deviations from Gaussianity in the primordial matter density field are dependent on the potential of the inflaton field [2], which can be parameterized and measured both in the statistical distribution of fluctuations in the Cosmic Microwave Background (CMB) and in the galaxy density field. While the CMB has provided the most stringent constraints so far [3], its ability to further improve is fundamentally limited by cosmic variance (statistical uncertainty stemming from the reality that we can only observe a single realization of possible universes). Galaxy survey methods of probing PNG include the bispectrum, which is the Fourier-space three-point correlation function of the galaxy density field, as well as the scale-dependent galaxy bias that PNG induces in the clustering of galaxies. The latter can be parameterized via the parameter  $f_{\text{NL}}^{\text{loc}}$ , which quantifies the amplitude of local-type non-Gaussianity. In the presence of PNG, the linear galaxy bias  $b_1(z)$ , which across redshift space linearly relates the observed galaxy overdensity to the underlying dark matter overdensity that drives the growth of structure, acquires a scale-dependent correction  $\Delta b(k, z)$  [4] [1]:

$$\Delta b(k, z) = \frac{3(b_1(z) - 1)f_{\text{NL}}^{\text{loc}}\Omega_m H_0^2 \delta_c}{D(z)k^2 T_\delta(k)} \quad (1)$$

Here the quantity  $\Omega_m$  is the present-day matter density fraction,  $H_0$  is the Hubble constant, and  $\delta_c \simeq 1.686$  is the critical overdensity for spherical collapse. The function  $D(z)$  is the linear growth factor, which encodes the redshift evolution of density perturbations, and the matter transfer function  $T_\delta(k)$  accounts for the growth of the matter density perturbations through epochs of respective radiation and matter domination. For the purposes of computing Eq. [1], we make the approximation  $T_\delta(k) = 1$  across all  $k$  in our analysis, which is a great approximation at low  $k$  where the bias correction is most prominent.

This  $\Delta b(k, z)$  correction introduces a distinct  $k^{-2}$  dependence at large scales (small  $k$ ), making large-scale measurements particularly sensitive to  $f_{\text{NL}}^{\text{loc}}$ . The signal is strongest at high redshift where  $D(z)$  is small and  $b_1(z)$  is larger, enhancing the PNG signature in wide, deep galaxy surveys.

Simplest models (single-field, slow-roll inflation) predict  $|f_{\text{NL}}^{\text{loc}}| < 1$ . Achieving  $\sigma(f_{\text{NL}}^{\text{loc}}) \sim \mathcal{O}(1)$  will enable differentiation between current prospective models of inflation [6], including the number of fields and their interactions [7]. The SPHEREx satellite survey [8], which recently launched in March 2025, has the explicit purpose of measuring  $f_{\text{NL}}^{\text{loc}}$  through the construction of maps of LSS over 75% of the sky and at low redshift. Forecasts for  $\sigma(f_{\text{NL}}^{\text{loc}})$  from SPHEREx use the combined 3D power spectrum and bispectrum and predict sufficient constraints for detecting PNG. The modeling choices are arguably optimistic (e.g. assuming perfectly Gaussian photo- $z$  uncertainties for each galaxy, and minimal impact from variations in the galaxy selection function, i.e. LSS systematics, which are some of the most significant systematics for  $f_{\text{NL}}^{\text{loc}}$ ). Detecting PNG with a 3D survey using the bispectrum is a primary

---

<sup>1</sup>In reality there is a  $b_\phi$  term in this equation as well [5]. We assume for the current analysis that  $b_\phi = 1$ , or alternatively one can interpret our results for  $f_{\text{NL}}^{\text{loc}}$  to actually be for the parameter combination  $f_{\text{NL}}^{\text{loc}} b_\phi$ .

science goal of SPHEREx. In our work we instead do a tomographic study of a SPHEREx-like experiment to forecast constraints on  $f_{\text{NL}}^{\text{loc}}$  through the impacts of the scale-dependent galaxy bias on the angular power spectrum.

## 1.2 Galaxy Clustering & the Angular Power Spectrum

In wide-area surveys, we observe galaxies projected onto the celestial sphere, often with redshift estimates that are too imprecise to reconstruct the full three-dimensional density field. Instead, we work with the angular power spectrum, denoted  $C_\ell$ , which quantifies the statistical correlations of galaxy number counts across the sky at different angular scales.

The observed fluctuation in the number density of galaxies (that traces the underlying dark matter density) in a direction  $\hat{\mathbf{n}}$  on the sky can be expanded in spherical harmonics:

$$\delta_g(\hat{\mathbf{n}}) = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}). \quad (2)$$

The angular power spectrum is then defined as the variance of the harmonic coefficients: for two tracers denoted by A and B (for autocorrelations,  $A = B$ ) [9]

$$\langle a_{\ell m}^A a_{\ell' m'}^{B*} \rangle = C_\ell^{AB} \delta_{\ell\ell'} \delta_{mm'}. \quad (3)$$

Similarly to how wavenumber  $k$  corresponds to variations on characteristic scales,  $\ell$  relates to angular scales of  $\theta \simeq \frac{180^\circ}{\ell}$ , and  $C_\ell$  encodes the strength of fluctuations at that scale.

To connect these observables to the underlying matter distribution, we integrate over wavenumber  $k$ , with each mode weighed by the primordial power spectrum and transfer functions that encode how 3D fluctuations are projected onto the sky: [9]

$$C_\ell^{AB} = 4\pi \int_0^\infty \frac{dk}{k} \mathcal{P}_\Phi(k) \Delta_\ell^A(k) \Delta_\ell^B(k), \quad (4)$$

where  $\mathcal{P}_\Phi(k)$  is the dimensionless power spectrum of primordial curvature perturbations [10] and  $\Delta_\ell^A(k)$  &  $\Delta_\ell^B(k)$  are number count transfer functions for each respective tracer, which include contributions [11] from galaxy density (which is proportional to dark matter density by the linear galaxy bias, and incorporates the redshift distribution of galaxies), from systematic redshift-space distortions (RSDs) [12, 13] that flatten the shape of mass overdensities due to systematic infall toward the center and extend the shape in the radial direction due to the peculiar velocities of galaxies bound in galaxy clusters, and finally a magnification bias contribution that accounts for the effect of gravitational lensing on observed galaxy counts, which can increase the number of detected galaxies by brightening ones that were previously too faint to see or decrease the number by stretching the sky and spreading out their apparent density [14]. This effect over redshift is incorporated by the function  $s(z)$  that is baked into the transfer function, which formally is dependent on the faint-end slope of the luminosity function [15].

In tomographic analyses, the galaxy sample is divided into redshift bins, with each bin treated as a separate tracer. In this work, we consider only the angular auto-power spectra within each bin when forecasting constraints on  $f_{\text{NL}}^{\text{loc}}$ ; cross-power spectra between bins enter only through the covariance matrix.

Primordial non-Gaussianity affects the large-scale clustering of galaxies through its impact on the bias, as described earlier. This imprint manifests in the angular power spectrum as an enhancement of power at low multipoles (large angular scales), as can be seen in Figure 1. This is particularly impactful in auto-power spectra of high-redshift, high-bias galaxy

populations. Accurately modeling and extracting this signal requires careful treatment of survey systematics and of the redshift distribution.

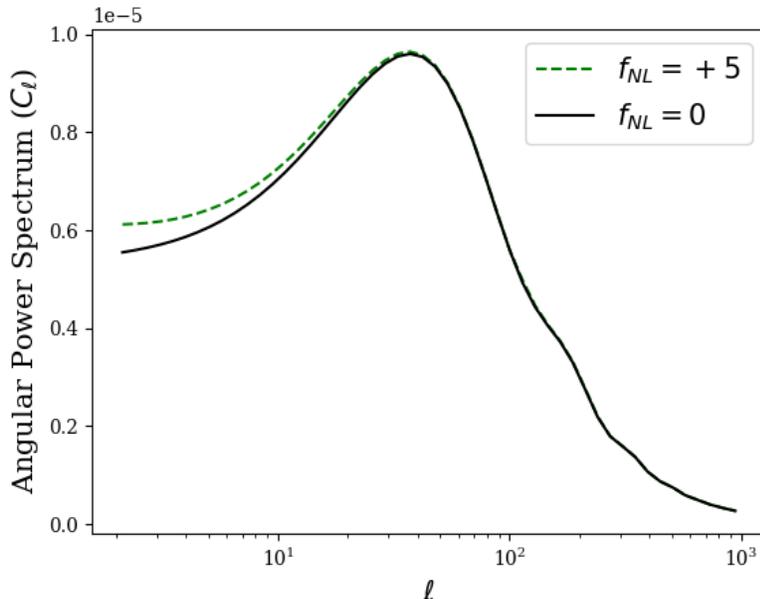


Figure 1: Imprint of  $f_{\text{NL}}^{\text{loc}}$  on the Angular Power Spectrum

## 2 Methods

### 2.1 SPHEREx-Like Experiment

We design a SPHEREx-like galaxy clustering survey using the Core Cosmology Library (CCL) [9] and public access SPHEREx statistical data [16], which provide simulated galaxy number densities and linear galaxy bias values for a specific redshift binning scheme. The galaxy sample is divided into five subsamples based on photometric redshift accuracy [8], each labeled by its maximum redshift uncertainty  $\tilde{\sigma}_z \equiv \sigma(z)/(1+z)$ . The five subsamples correspond to  $\tilde{\sigma}_z = \{0.003, 0.01, 0.03, 0.1, 0.2\}$ . For each subsample we perform a cubic interpolation of its simulated bias as can be seen in Figure 2, and fit its galaxy number counts to

$$N(z) = A \frac{z^{d-1}}{z_0^d} \exp\left\{-\left(z/z_0\right)^\beta\right\}, \quad (5)$$

where we fit for  $A$ ,  $d$ ,  $z_0$ , and  $\beta$ .

The SPHEREx simulated galaxy number densities are given in units of  $[h/\text{Mpc}]^3$ . In order to convert into units of  $[\text{galaxies}/\text{ster}]$  for fitting to a functional  $N(z)$  form, we first map the bin edges in redshift to comoving radial distance in units [Mpc] and calculate the volume of the shell defined by that bin. If  $a$  and  $b$  are the radial bin edges in comoving radial coordinates:

$$N(z) [\text{ster}^{-1}] = h^3 \frac{(b^3 - a^3)}{3} \times N(z) \left[ \frac{h^3}{\text{Mpc}^3} \right] \quad (6)$$

This requires assuming a specific cosmology but the resulting number densities are sufficiently insensitive to variations in cosmology for the purposes of being fit to a parent distribution.

With functions for galaxy count and bias in redshift for each subsample, we are able to vary the number and size of redshift bins. We incorporate redshift uncertainty in the binning by convolving over the bin slices of the galaxy count parent distributions with a Gaussian

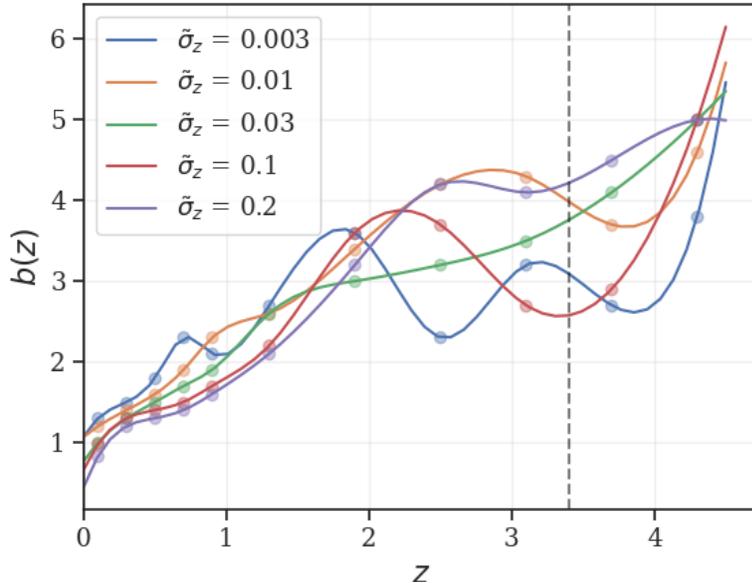


Figure 2: Interpolated bias for all five subsamples. The dashed line indicates the maximum redshift bin edge we use.

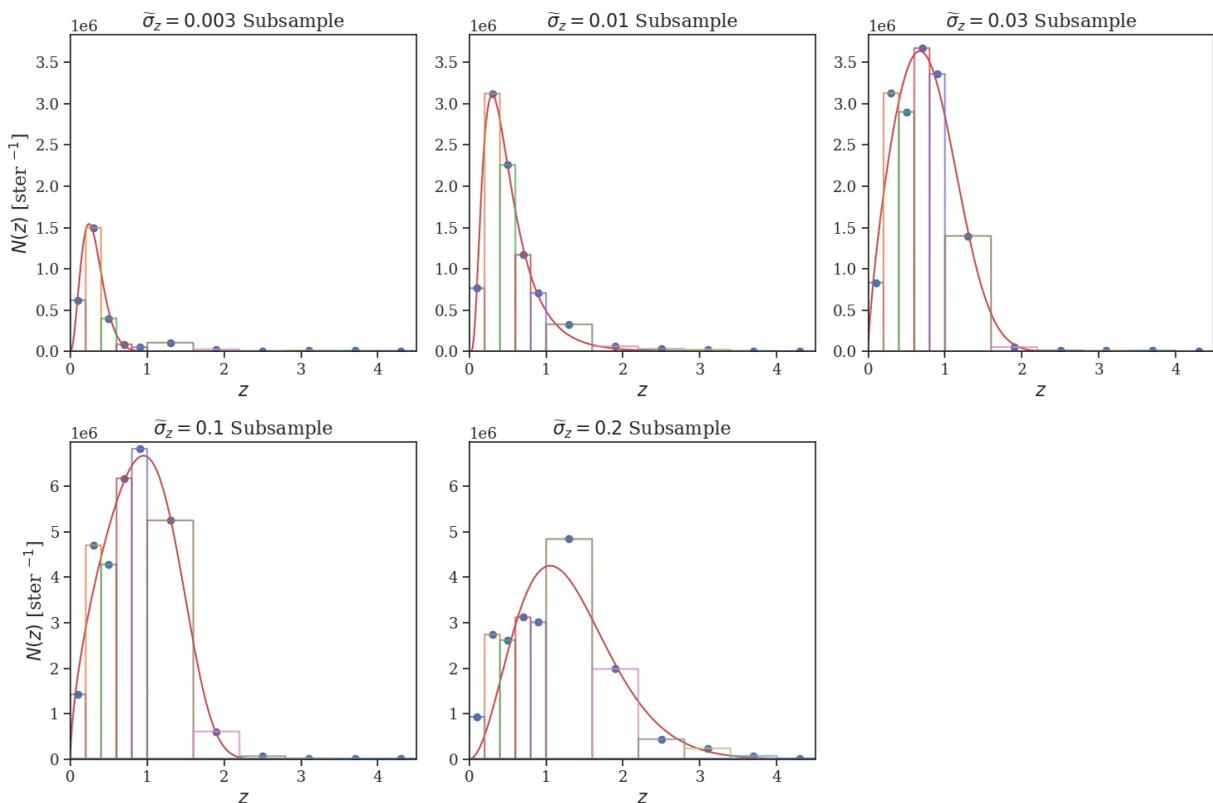


Figure 3: Parametrized Galaxy Number Counts

filter of standard deviation  $\tilde{\sigma}_z(1 + z^*)$ , where  $z^*$  is the mean redshift of the bin. We then perform a cosine window apodization at low  $z$  to ensure that the galaxy number counts go to zero at  $z = 0$ , and additionally perform an apodization at high redshift for the highest-uncertainty subsample because the wide convolution Gaussian bloats the distribution far beyond the parent distribution.

We use CCL to create galaxy number count tracers from the bins (with the scale-dependent bias included), cross-correlate the tracers to get angular power spectra, and then forecast constraining power on  $f_{\text{NL}}^{\text{loc}}$  through the Fisher matrix formalism.

We set  $k_{\text{max}} = 0.2h$  as we are in the quasi-linear regime and can only model down to certain scales. This choice does not affect constraints on  $f_{\text{NL}}^{\text{loc}}$  much, as constraining power is concentrated at low  $k$ . We convert from  $k_{\text{max}}$  to  $\ell_{\text{max}}^{(i)}$  for the  $i$ th redshift bin by the relation  $\ell_{\text{max}}^{(i)} \simeq k * \chi(z_*^{(i)})$  where  $\chi(z_*^{(i)})$  is the co-moving radial distance evaluated at the

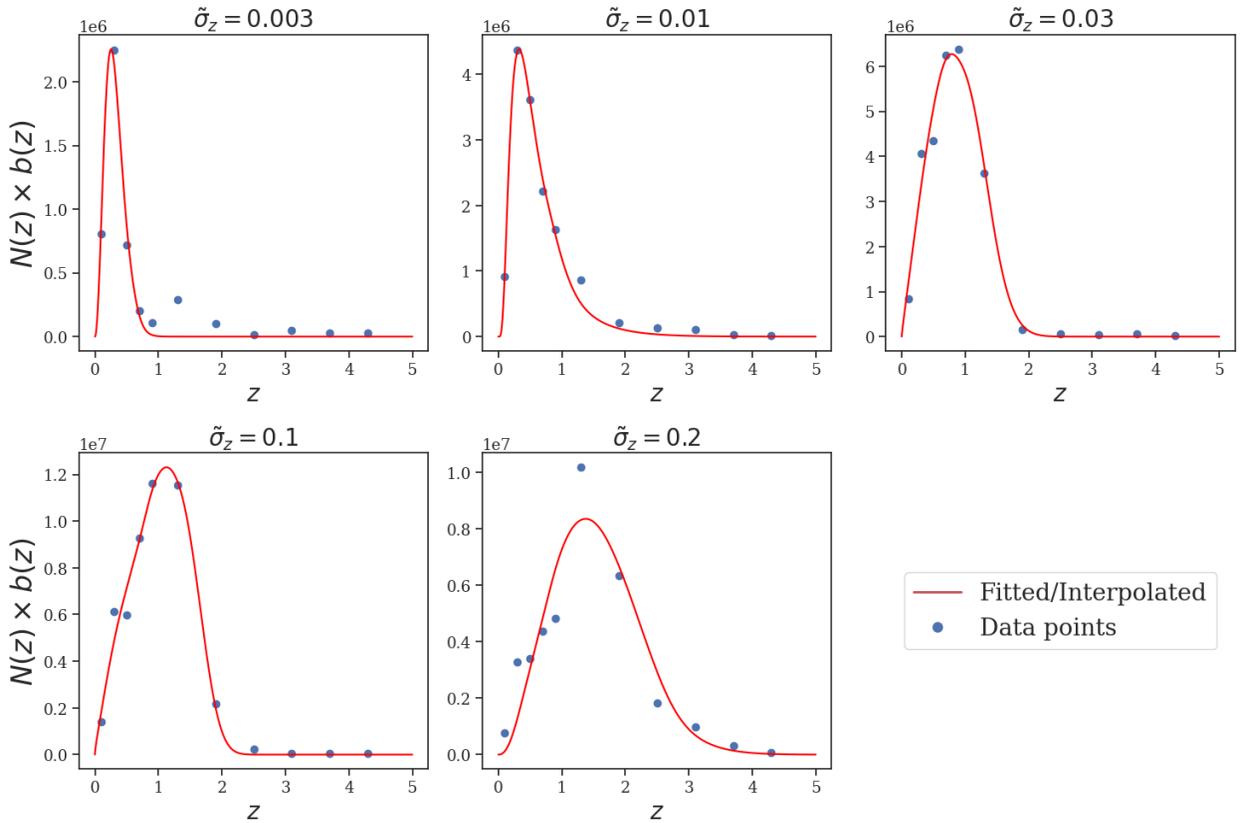


Figure 4: Galaxy counts times bias, showing how closely the fitted and interpolated values match the simulated data provided by SPHEREx [17]

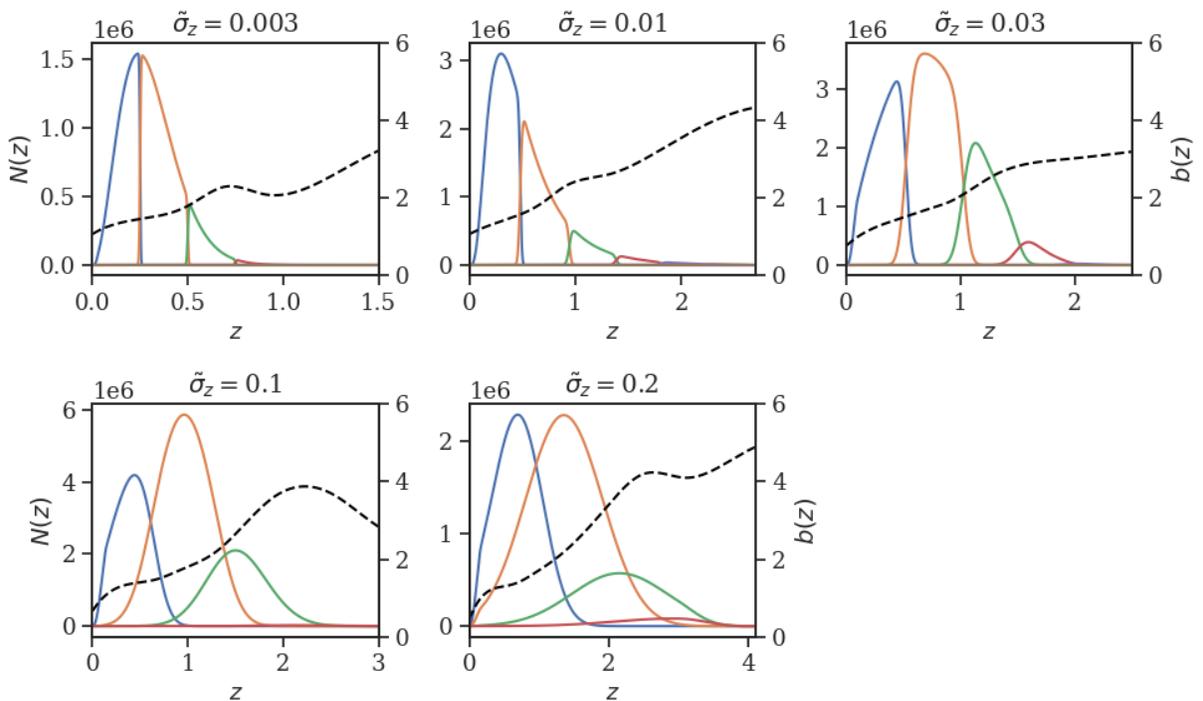


Figure 5: Final redshift bins (after convolution, colored) and linear galaxy bias (black dashed) for each subsample of SPHEREx galaxies.

mean redshift of the  $i$ th bin. For each bin we evaluate  $C_\ell$  from  $\ell = 3$  to  $\ell_{\max}^{(i)}$ . We bin in  $\ell$  with  $\Delta\ell = 2$  from  $\ell = 2$  to 30 and  $\Delta\ell = 20$  from  $\ell = 30$  to  $\ell_{\max}^{(i)}$ .

CCL allows RSD and magnification bias additions to its tracers. We incorporate these effects, picking a fiducial constant  $s(z) = 0.6$  which we will vary and marginalize over in our analysis. With magnification incorporated, we find that the limber approximation implemented in CCL fails even on small angular scales, and so we rely fully on the FKEM non-limber integration method [18] across all  $\ell$ .

## 2.2 Forecasting Constraints with Fisher Formalism

The Fisher formalism is an efficient way of estimating model parameter uncertainties. By assuming the likelihood distribution  $\mathcal{L}$  in the model parameters  $p_i$  is a multivariate Gaussian about its peak, one can calculate the Fisher information matrix [19]:

$$F_{ij} = \left\langle -\frac{\partial^2 \ln \mathcal{L}}{\partial p_i \partial p_j} \right\rangle \quad (7)$$

The matrix measures how sharply the likelihood function is curved about its (assumed) maximum point. A steeper curvature means the parameters can be estimated more precisely, indicating that the data contains more information about the cosmological parameters. For measurements of the angular power spectrum, the Fisher matrix can instead be expressed as

$$F_{ij} = \sum_{\ell} \frac{\partial \mathbf{C}}{\partial p_i} \mathbf{Cov}^{-1} \frac{\partial \mathbf{C}}{\partial p_j} \quad (8)$$

Here  $\mathbf{C}$  is a vector of angular autopower spectra for each pair of redshift bins evaluated at  $\ell$ , and  $\mathbf{Cov}^{-1}$  is the inverse covariance matrix between angular power spectra taking Gaussian covariance on the  $C_{\ell}$ s, and whose matrix element corresponding to bins A and B is given by [20]

$$\mathbf{Cov} [C_{\ell}^{AA}, C_{\ell}^{BB}] = \frac{2}{(2\ell + 1)f_{\text{sky}}\Delta\ell} (\tilde{C}_{\ell}^{AB})^2. \quad (9)$$

$f_{\text{sky}}$  is the fraction of the sky covered by the experiment (which for SPHEREx is 0.75),  $\Delta\ell$  is the width of the  $\ell$  bin, and the auto power spectra (when  $A = B$ ) includes shot noise contributions  $1/\bar{n}^{(i)}$ , where  $\bar{n}^{(i)}$  is the angular number density of galaxies (with units of  $[\text{ster}^{-1}]$ ) in redshift bin  $i$ :

$$\tilde{C}_{\ell}^{AB} = C_{\ell}^{AB} + \frac{\delta_{AB}}{\bar{n}^{(i)}}. \quad (10)$$

$p_i$  runs over the cosmological parameters  $[\Omega_c, \Omega_b, h, \sigma_8, n_s, w_0, w_a, f_{\text{NL}}^{\text{loc}}]$ . We additionally incorporate uncertainty in the linear galaxy bias with multiplicative bias parameters  $b_z$  for each bin, uncertainty in photometric redshift systematics with parameters  $\text{shift}_z$  that shift redshift bins translationally as well as parameters  $\text{stretch}_z$  that stretch/squeeze the redshift bin as is done in [21], and lastly uncertainty in the magnification bias  $s(z)$  with parameters  $s_z$  for each bin, to which we add a  $+/- 0.02$  Gaussian prior [22]. We numerically evaluate the derivatives of  $C_{\ell}$  by generating power spectra with small perturbations in the parameters around their fiducial values. Fiducial parameter values and step sizes for use in calculating the numerical derivatives are given in Table 1. With this formalism we are easily able to combine constraints from the different subsamples by adding their respective Fisher matrices after marginalizing over nuisance parameters.

The resulting Fisher matrix yields the best possible constraints on model parameters given a set of data, as calculated by the Cramer-Rao inequality which gives the lower bound for error on a parameter:

$$\sigma(p_i) \geq \begin{cases} \sqrt{(F^{-1})_{ii}} & (\text{marginalized}) \\ 1/\sqrt{F_{ii}} & (\text{unmarginalized}) \end{cases} \quad (11)$$

Marginalized constraints account for correlations between parameters by letting them communicate through the matrix inverse, and are therefore more conservative and realistic than unmarginalized ones, which assume all other parameters are known perfectly [23]. Henceforth, all reported parameter errors are assumed to be marginalized over the rest of the parameters unless otherwise specified.

Parameter ( $p_i$ )	Fiducial Value	Step Size ( $\Delta p_i$ )
$\Omega_c$	0.25	$\pm 1\%$
$\Omega_b$	0.05	$\pm 1\%$
$h$	0.67	$\pm 1\%$
$\sigma_8$	0.81	$\pm 1\%$
$n_s$	0.96	$\pm 1\%$
$w_0$	-1.0	$\pm 0.1$
$w_a$	0	$\pm 0.1$
$f_{\text{NL}}^{\text{loc}}$	0	$\pm 1$
$b_z$	1	$\pm 0.05$
shift $_z$	0	$\pm 0.01$
stretch $_z$	1	$\pm 0.05$
$s_z$	0.6	$\pm 0.02$

Table 1: Fiducial values and step sizes for cosmological and nuisance parameters.

### 3 Results

With the binning architecture and Fisher matrix formalism established in the previous section, we now examine how different choices of redshift binning and subsample configurations affect constraints on  $f_{\text{NL}}^{\text{loc}}$ . We begin by evaluating the constraining power of individual bins and subsamples, then explore the effects of binning choices more broadly, and finally perform some robustness tests to see what modeling assumptions are most impactful.

#### 3.1 Constraints Per Bin

We begin by examining how constraining power on  $f_{\text{NL}}^{\text{loc}}$  varies across redshift and subsample. Picking the same bin edges that SPHEREx used in their simulations [16], we investigate the amount of unmarginalized constraining power per redshift bin and per photometric-redshift-accuracy subsample (defined by maximum values of  $\tilde{\sigma}_z$ ). Figure 6 shows the result, making it clear that constraints are most stringent at high redshift and thus at high uncertainty subsamples that go deeper in redshift. At the highest redshifts, uncertainties increase significantly in several subsamples due to large shot noise from the low number density of galaxies in these bins. The  $\tilde{\sigma}_z = 0.01$  subsample retains good constraining power even at high redshift despite being a low-uncertainty subsample. This apparent strength is partly a result of how the parent distribution for that subsample was fitted; this effect is mitigated when marginalizing over nuisance parameters that affect bin shape and positioning.

However, bins in a given subsample do not independently contribute to a subsample’s constraining power, due to off-diagonal terms in the covariance matrix in equation 8 that come from cross-power between bins. Therefore we cannot simply add the constraints on  $f_{\text{NL}}^{\text{loc}}$  from each bin together and must instead consider constraints from an entire subsample. One way to still examine constraining power per bin on  $f_{\text{NL}}^{\text{loc}}$  while accounting for correlated  $C_{\ell s}$  (and marginalization over all other parameters) is to see how constraints change when removing that bin completely from the subsample. The results of this analysis are shown in figure 7 and are visualized with and without normalization to the full subsample’s constraint with no bins removed. The normalized plot most clearly shows which bins are most important in each subsample. High-redshift bins still provide the strongest constraints as is expected, while the bins that contain almost no power at all are either at very low redshift where the PNG signal is weak, or at very high redshift in low-uncertainty subsamples where shot noise dominates, or for certain bins in the highest-uncertainty subsample where the bins are so

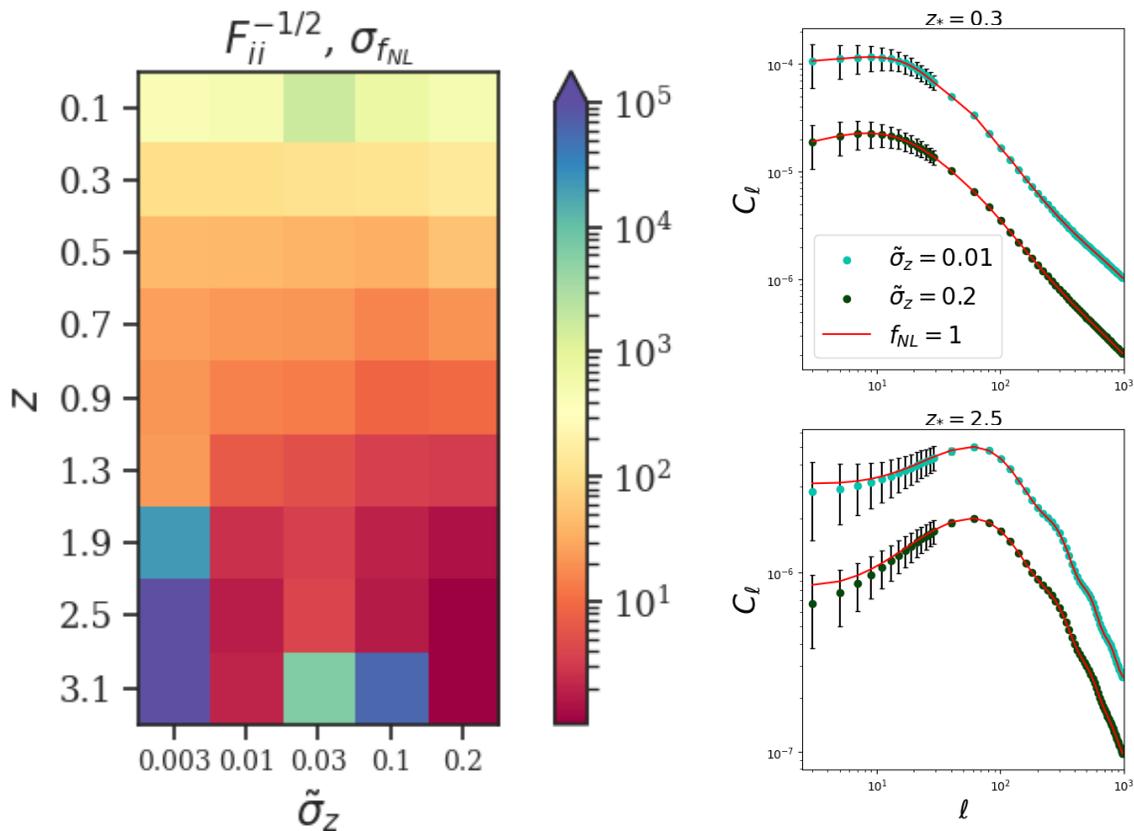


Figure 6: Left: Unmarginalized constraints on  $f_{\text{NL}}^{\text{loc}}$  from each bin in each subsample, taken individually. Note that bins are not all independent. Right: Auto power spectra and the imprint of  $f_{\text{NL}}^{\text{loc}}$  for a low and a high redshift bin, and for a low and high uncertainty subsample.

overlapped that the information becomes saturated.

### 3.2 Optimal Binning

The choice of redshift binning plays a crucial role in determining the constraining power of a tomographic analysis. Binning too coarsely risks diluting cosmological signals by averaging over redshift evolution, while overly fine binning increases shot noise and makes the analysis more sensitive to systematic effects. An optimal binning scheme must strike a balance between maximizing sensitivity to scale-dependent bias and maintaining robustness against systematics.

To explore this tradeoff, we vary the number of redshift bins for each photometric-accuracy subsample and assess the resulting constraints on  $f_{\text{NL}}^{\text{loc}}$ . We also explore the impact of non-uniform binning schemes by introducing a skew factor, defining redshift bin edges that modify the width of each bin linearly with redshift:

$$\Delta z_i = \Delta z_{\text{base}} [1 + (\text{skew} - 1)z_i], \quad (12)$$

where  $\Delta z_{\text{base}}$  is the average bin width in the unskewed (evenly spaced) case, and  $z_i$  is the lower edge of the  $i$ -th bin. The bin edges are then rescaled to exactly span the desired redshift interval. When  $\text{skew} = 1$ , bins are evenly spaced in redshift;  $\text{skew} > 1$  stretches bin widths toward higher redshift (wider bins at high  $z$ ), while  $\text{skew} < 1$  compresses bins toward high redshift (wider bins at low  $z$ ). This allows us to test whether concentrating redshift resolution in specific ranges improves constraints over uniformly spaced binning.

Figure 8 summarizes the dependence of marginalized uncertainties in  $f_{\text{NL}}^{\text{loc}}$ , combining constraints from all subsamples, on both the number of evenly spaced bins and the skew parameter. The skew variation analysis is performed at a fixed bin count of six. We find that beyond two to three bins, increasing the number of bins yields only modest improvements in

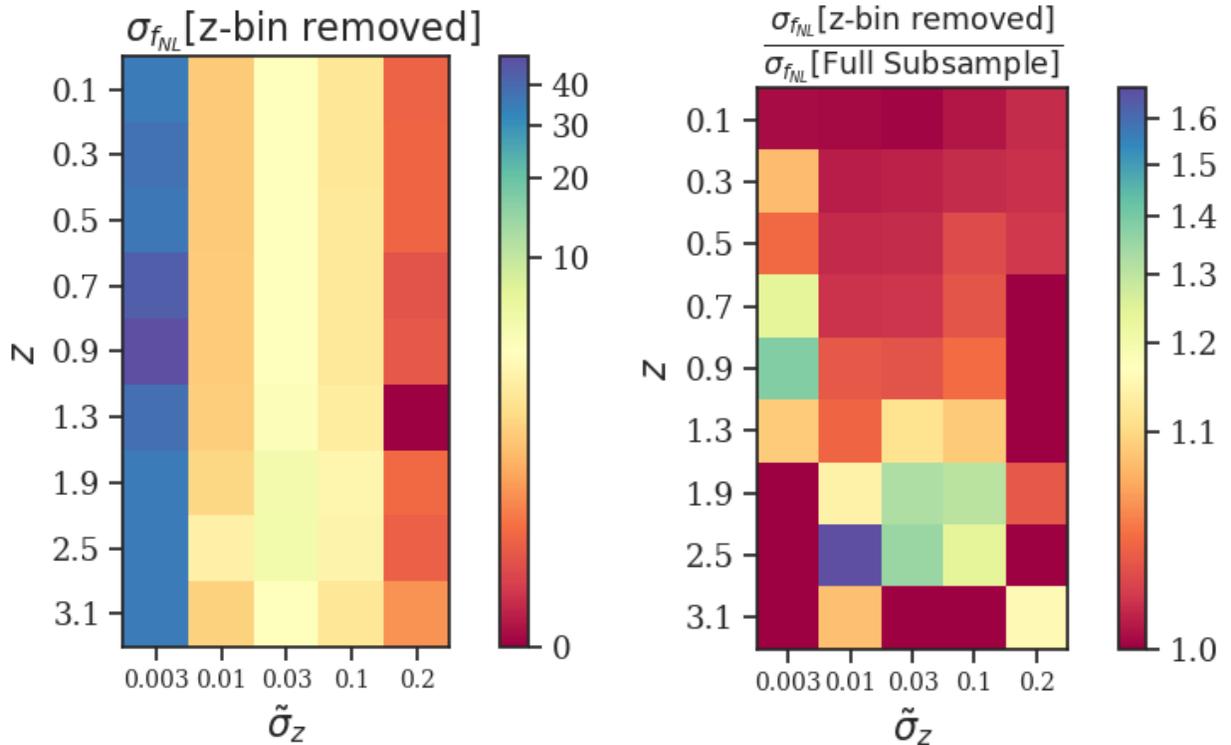


Figure 7: Marginalized uncertainties on  $f_{\text{NL}}^{\text{loc}}$  when individual redshift bins are removed from their respective subsamples (columns). Each entry corresponds to the uncertainty obtained after excluding the bin centered at the specified redshift in the specified subsample. The left panel shows the raw marginalized uncertainties, indicating which bins are most critical to constraining power. The right panel shows the same values normalized by the uncertainty from the full subsample (with all bins included), highlighting the relative importance of each bin within its subsample.

constraints, and changes in the skew factor have some effect but make little overall difference unless extreme. These results suggest that when subsamples are combined, the precise choices in binning configuration have limited impact on the overall constraining power, which are primarily driven by sample depth in redshift. Therefore, maintaining control over systematics (see e.g. [24]) should be prioritized over fine-tuning bin placement or count.

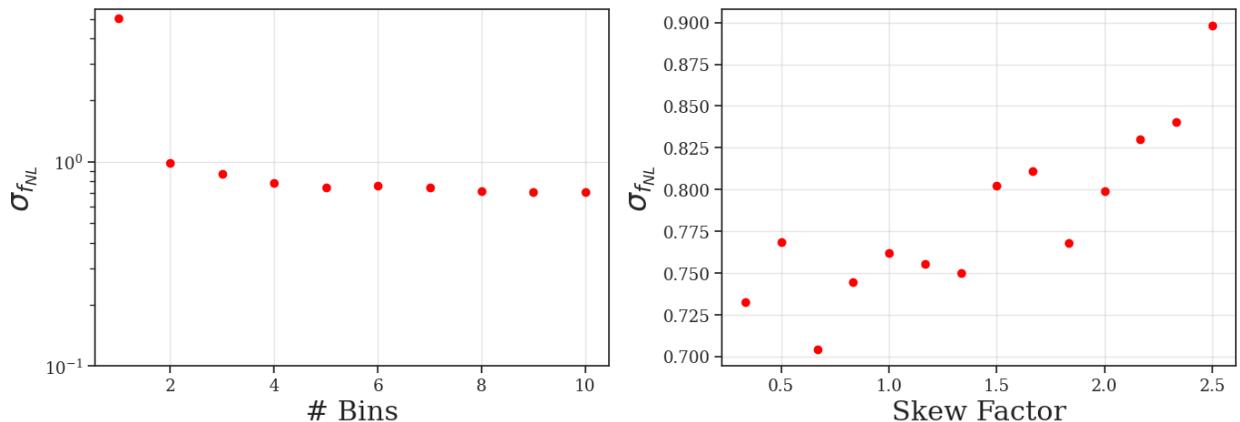


Figure 8: Marginalized  $f_{\text{NL}}^{\text{loc}}$  constraints combining all subsamples, as a function of binning choices. **Left:** Constraints versus the number of evenly spaced redshift bins. Beyond two to three bins, improvements quickly plateau as information saturates. **Right:** Constraints versus the skew factor at a fixed bin count of six. Skew = 1 corresponds to even binning, while skew > 1 and < 1 shift resolution toward high and low redshift, respectively. The weak dependence on binning configuration indicates that maximizing redshift depth and controlling systematics are more impactful than fine-tuning bin structure.

While these results demonstrate that binning configuration has limited impact when subsamples are combined, we next determine optimal binning strategies *within* each individual subsample. Because the five photometric-accuracy subsamples reach different redshift

depths, we assign a separate maximum bin edge for each:  $\{1.5, 2.6, 2.8, 2.8, 3.4\}$  from lowest to highest redshift uncertainty. For each subsample, we first vary the number of evenly spaced bins up to its redshift limit and identify the smallest number of bins that yields sufficient constraints on  $f_{\text{NL}}^{\text{loc}}$  without saturating information with too many bins. With the bin count picked for each subsample and fixed, we then vary the skew parameter and pick optimal values. This procedure balances computational efficiency with robustness and serves as the basis for our final binning choices used in the combined forecasts. The results of this optimization procedure are shown in Figure 9 and the optimal redshift bins picked from this analysis are shown in Figure 5. We find that 4-6 bins are sufficient bin numbers and that near-even spacing in the bins is preferred.

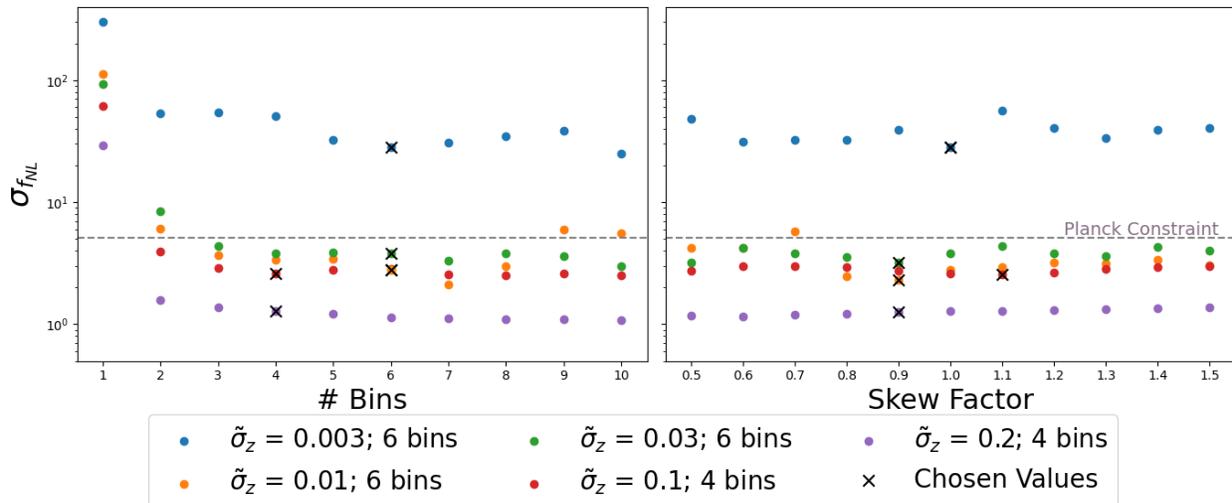


Figure 9: Optimization of binning strategy for each subsample. We use a different maximum redshift for each subsample:  $\{1.5, 2.6, 2.8, 2.8, 3.4\}$  for  $\tilde{\sigma}_z = \{0.003, 0.01, 0.03, 0.1, 0.2\}$ , respectively. The dashed line denotes the Planck 2018 constraint on  $f_{\text{NL}}^{\text{loc}}$  for reference [3]. Black crosses mark the chosen bin count and skew for each subsample, selected to balance performance and robustness. **Left:** Marginalized constraints on  $f_{\text{NL}}^{\text{loc}}$  versus the number of evenly spaced redshift bins. The chosen number of bins for each subsample are marked. **Right:** Constraints versus skew parameter at fixed bin count (selected per subsample based on the left panel).

Subsample $\tilde{\sigma}_z$	# Bins	Skew Factor	$\sigma(f_{\text{NL}}^{\text{loc}})$
0.003	6	1.0	<b>28.4</b>
0.01	6	0.9	<b>2.31</b>
0.03	6	0.9	<b>3.22</b>
0.1	4	1.1	<b>2.56</b>
0.2	4	0.9	<b>1.25</b>
Combined	–	–	<b>0.823</b>

Table 2: Marginalized constraints on  $f_{\text{NL}}^{\text{loc}}$  from chosen optimal binning for each individual subsample and then with all subsamples combined.

Table 2 shows the resulting constraints on  $f_{\text{NL}}^{\text{loc}}$  from the chosen redshift binning, achieving a projected uncertainty of  $\sigma(f_{\text{NL}}^{\text{loc}}) = 0.823$  when combining all subsamples. This result is close to the Fisher forecast projected constraints from the 3D power spectrum done in [8] using different redshift binning and different assumptions (e.g. uncorrelated power spectra between bins).

To better understand how constraints on  $f_{\text{NL}}^{\text{loc}}$  arise from the Fisher formalism, we examine the correlation matrices between all model parameters, including both cosmological and

nuisance parameters. These are derived by normalizing the inverse Fisher matrix and thus reflect the degree to which uncertainties in different parameters are degenerate with one another.

In particular, correlations with  $f_{\text{NL}}^{\text{loc}}$  are important because they directly impact the marginalized uncertainty: strong degeneracies lead to degraded constraints, while weak or no correlation implies that the constraint is more robust and driven by the unique features of the  $f_{\text{NL}}^{\text{loc}}$  signal. We find that  $f_{\text{NL}}^{\text{loc}}$  is generally weakly correlated with other parameters — both cosmological and nuisance — indicating that the PNG signal extracted from scale-dependent galaxy bias is relatively distinct and well-isolated in the parameter space.

The full correlation matrices for our chosen binnings are shown in Figure 10 and then zoomed in on the  $\tilde{\sigma}_z = 0.2$  subsample in Figure 11. These plots provide insight into which parameters are most degenerate and help validate the stability of the resulting marginalized forecasts.

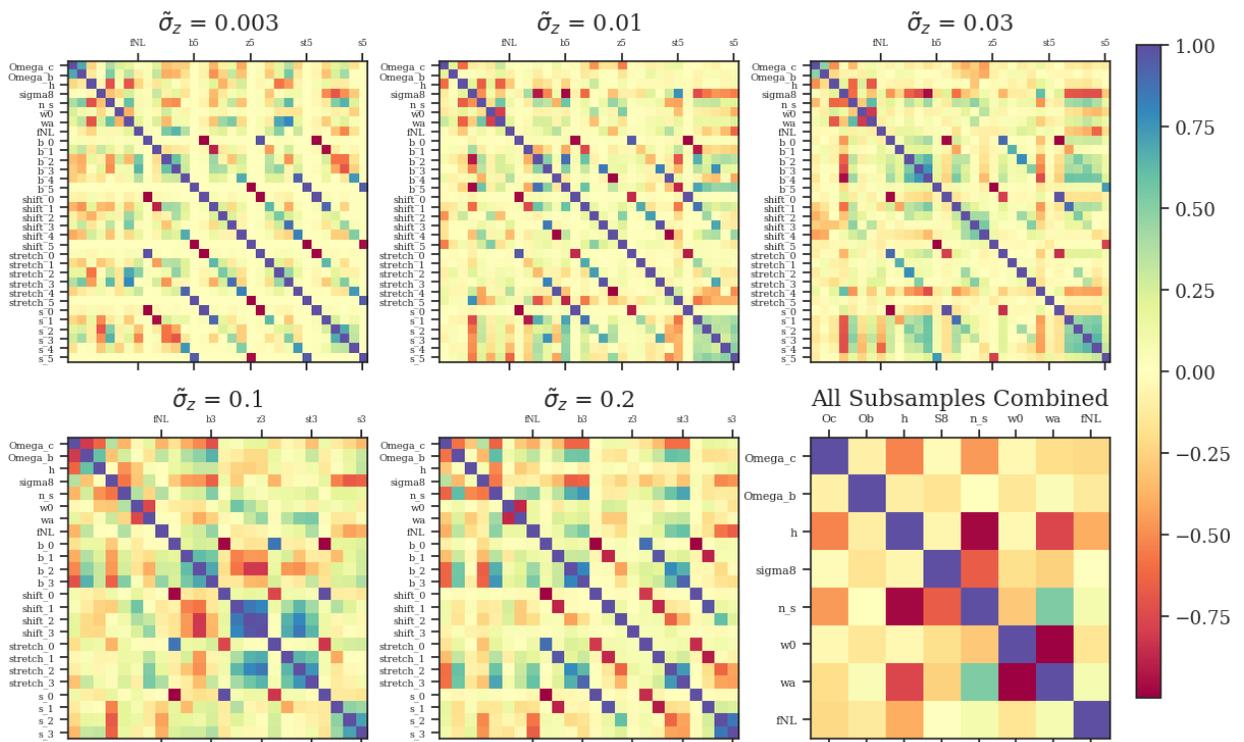


Figure 10: Fisher matrix correlation matrices for each subsample and the combined case (bottom right). Each matrix shows the normalized correlations between model parameters after applying the optimized binning and skew configuration for that subsample. The combined correlation matrix includes only the cosmological parameters, after marginalizing over all nuisance parameters.

### 3.3 Robustness Tests

To assess the sensitivity of our forecasted constraint on  $f_{\text{NL}}^{\text{loc}}$  to modeling choices, we perform a set of robustness tests which are designed to evaluate what assumptions are the most impactful.

In the left panel of Figure 12, we test the impact of various components of our model. Starting from the full model, we evaluate the change in  $\sigma(f_{\text{NL}}^{\text{loc}})$  when fixing specific nuisance parameters instead of marginalizing over them. We find that modeling uncertainty on the bias has the most impact on results, whereas nuisance parameters that model uncertainty on  $s(z)$  has very little effect. As expected, the bin shape parameters also have some impact but not as much as the bias parameters which directly affect the PNG signal through equation

1

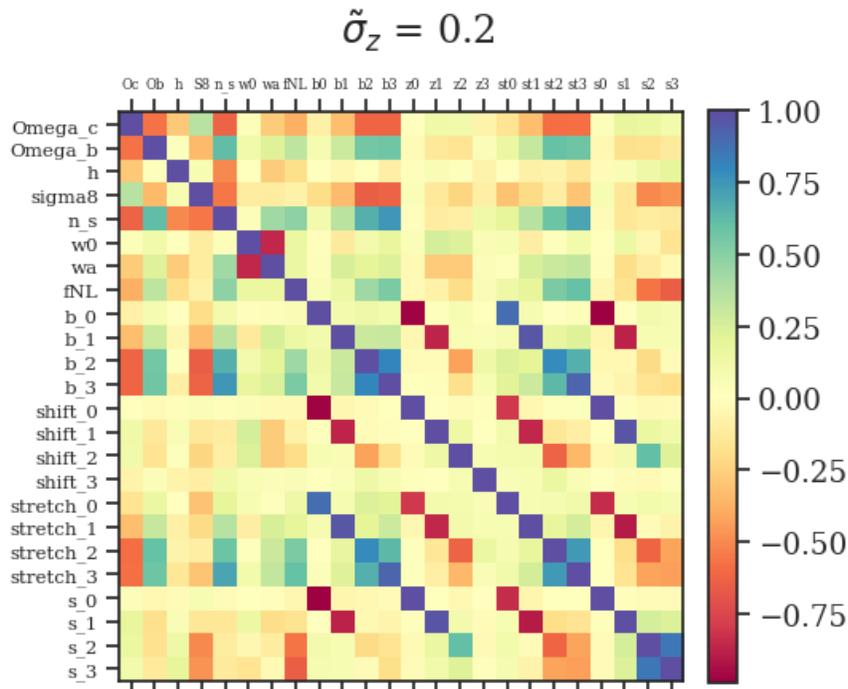


Figure 11: Correlation matrix for the  $\tilde{\sigma}_z = 0.2$  subsample

We also check the constraint on  $f_{\text{NL}}^{\text{loc}}$  without including contributions from the  $\tilde{\sigma}_z = 0.2$  subsample in case we can't trust such high redshift uncertainties, which could lead to systematics and dangerous dependence on the assumption of Gaussian error which drives constraints at the tails (which for the high- $z$  tail drives constraining power). We find in this case that  $\sigma(f_{\text{NL}}^{\text{loc}}) = 1.27$ , still beating down the uncertainty to the  $\mathcal{O}(1)$  level.

In the right panel, we show how the  $f_{\text{NL}}^{\text{loc}}$  constraint depends on the minimum multipole  $\ell_{\text{min}}$  included in the Fisher calculation. Since the scale-dependent bias signature from  $f_{\text{NL}}^{\text{loc}}$  is strongest on large angular scales, excluding low- $\ell$  modes quickly degrades the constraint. For example, increasing  $\ell_{\text{min}}$  from 2 to 10 more than doubles the uncertainty. This underscores the importance of robust large-scale systematics control and accurate modeling at low  $\ell$ , where cosmic variance, observational challenges, and potentially unmodeled contributions are greatest.

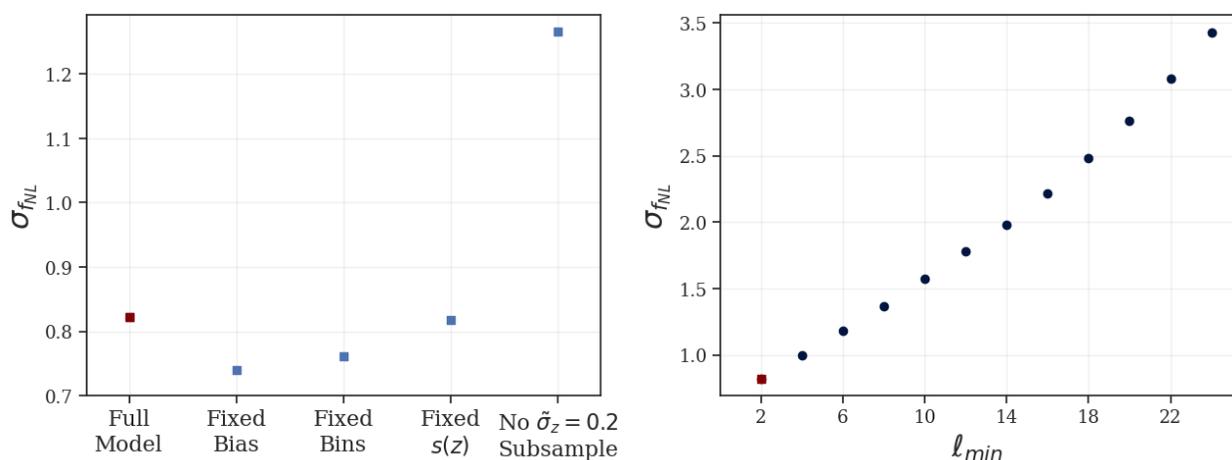


Figure 12: **Left:** How constraints on  $f_{\text{NL}}^{\text{loc}}$  change when fixing various model components (that we usually vary with parameters and marginalize over in the Fisher matrix), as well as the constraint without including the largest uncertainty subsample. **Right:** How constraints on  $f_{\text{NL}}^{\text{loc}}$  depend on the lowest value of  $\ell$  used in calculating angular power spectra.

Together, these tests show that our forecasted constraint on  $f_{\text{NL}}^{\text{loc}}$  is relatively robust to reasonable assumptions about nuisance parameters, and highlight the critical role of the largest-scale modes and deep redshift coverage in driving sensitivity to primordial non-Gaussianity.

## 4 Discussion

In this work, we have forecasted constraints on  $f_{\text{NL}}^{\text{loc}}$  from a SPHEREx-like tomographic galaxy clustering survey, focusing on the imprint of primordial non-Gaussianity in the angular power spectrum via the scale-dependent galaxy bias. By incorporating Gaussian photometric redshift uncertainties, nuisance parameter modeling, and Fisher forecasting techniques, we have estimated the constraining power of different subsamples and identified binning strategies that balance precision and robustness.

We find that if angular systematics can be controlled, a tomographic analysis of SPHEREx may be able to constrain  $f_{\text{NL}}^{\text{loc}}$  down to the  $\mathcal{O}(1)$  level, which is critical for distinguishing between single-field and multi-field inflationary models [6]. Notably,  $f_{\text{NL}}^{\text{loc}}$  is only weakly correlated with other cosmological and nuisance parameters, reinforcing the idea that it can be cleanly extracted from large-scale clustering measurements.

Our results also show that while the choice of redshift binning does have some effect on the final constraints, the differences are modest. This suggests that we should favor binning schemes that minimize sensitivity to systematic errors and reduce dependence on uncertain modeling assumptions. Much of the constraining power comes from high-redshift bins, where galaxy bias and the bin tails driven by the assumption of Gaussian redshift error dominate — both of which introduce theoretical uncertainties that will need to be better quantified in real data analyses [24].

In addition, our forecasts do not explicitly model selection function errors — mismatches between the true galaxy distribution and the assumed survey response, such as spatial variations in detection efficiency or contamination from stars. These unmodeled effects can bias clustering measurements and will be an important target for mitigation in future observational analyses.

In conclusion, SPHEREx has the potential to significantly advance our understanding of the early universe, with a tomographic analysis alone pushing us close to the constraining power needed to detect PNG, and with further gains available from the bispectrum. However, achieving this will require meticulous control of systematics as well as careful modeling of redshift uncertainties and bias. This analysis helps indicate where we should focus our efforts to obtain clean samples with maximal leverage on  $f_{\text{NL}}^{\text{loc}}$ .

## References

- [1] D. Baumann, “Tasi lectures on inflation,” 2012.
- [2] M. Rezaie, A. J. Ross, H.-J. Seo, H. Kong, A. Porredon, L. Samushia, E. Chaussidon, A. Krolewski, A. de Mattia, F. Beutler, *et al.*, “Local primordial non-gaussianity from the large-scale clustering of photometric desi luminous red galaxies,” *Monthly Notices of the Royal Astronomical Society*, p. stae886, 2024.
- [3] Y. Akrami, F. Arroja, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. Banday, R. Barreiro, N. Bartolo, S. Basak, *et al.*, “Planck 2018 results-ix. constraints on primordial non-gaussianity,” *Astronomy & Astrophysics*, vol. 641, p. A9, 2020.
- [4] D. Jeong and E. Komatsu, “Primordial non-gaussianity, scale-dependent bias, and the bispectrum of galaxies,” *The Astrophysical Journal*, vol. 703, no. 2, p. 1230, 2009.
- [5] A. Barreira, “Can we actually constrain  $f_{NL}$  using the scale-dependent bias effect? an illustration of the impact of galaxy bias uncertainties using the boss dr12 galaxy power spectrum,” *Journal of Cosmology and Astroparticle Physics*, vol. 2022, no. 11, p. 013, 2022.
- [6] R. de Putter, J. Gleyzes, and O. Doré, “Next non-gaussianity frontier: What can a measurement with  $\sigma(f_{NL}) \leq 1$  tell us about multifield inflation?,” *Physical Review D*, vol. 95, no. 12, p. 123507, 2017.
- [7] P. D. Meerburg, D. Green, M. Abidi, M. A. Amin, P. Adshead, Z. Ahmed, D. Alonso, B. Ansarinejad, R. Armstrong, S. Avila, *et al.*, “Primordial non-gaussianity,” *arXiv preprint arXiv:1903.04409*, 2019.
- [8] O. Doré, J. Bock, M. Ashby, P. Capak, A. Cooray, R. de Putter, T. Eifler, N. Flagey, Y. Gong, S. Habib, *et al.*, “Cosmology with the spherex all-sky spectral survey,” *arXiv preprint arXiv:1412.4872*, 2014.
- [9] N. E. Chisari, D. Alonso, E. Krause, C. D. Leonard, P. Bull, J. Neveu, A. Villarreal, S. Singh, T. McClintock, J. Ellison, *et al.*, “Core cosmology library: Precision cosmological predictions for lsst,” *The Astrophysical Journal Supplement Series*, vol. 242, no. 1, p. 2, 2019.
- [10] A. S. Josan, A. M. Green, and K. A. Malik, “Generalized constraints on the curvature perturbation from primordial black holes,” *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, vol. 79, no. 10, p. 103520, 2009.
- [11] X. Fang, E. Krause, T. Eifler, and N. MacCrann, “Beyond limber: Efficient computation of angular power spectra for galaxy clustering and weak lensing,” *Journal of Cosmology and Astroparticle Physics*, vol. 2020, no. 05, p. 010, 2020.
- [12] N. Kaiser, “Clustering in real space and in redshift space,” *Monthly Notices of the Royal Astronomical Society*, vol. 227, no. 1, pp. 1–21, 1987.
- [13] A. Hamilton, “Linear redshift distortions: A review,” *The Evolving Universe: Selected Topics on Large-Scale Structure and on the Properties of Galaxies*, pp. 185–275, 1998.
- [14] M. Bartelmann and P. Schneider, “Weak gravitational lensing,” *Physics Reports*, vol. 340, no. 4-5, pp. 291–472, 2001.

- [15] N. Bellomo, J. L. Bernal, G. Scelfo, A. Raccanelli, and L. Verde, “Beware of commonly used approximations. part i. errors in forecasts,” *Journal of Cosmology and Astroparticle Physics*, vol. 2020, no. 10, p. 016, 2020.
- [16] <https://github.com/SPHEREx/Public-products>.
- [17] <https://github.com/LSSTDESC/CCL>.
- [18] C. D. Leonard, T. Ferreira, X. Fang, R. Reischke, N. Schoeneberg, T. Tröster, D. Alonso, J.-E. Campagne, F. Lanusse, A. Slosar, *et al.*, “The n5k challenge: Non-limber integration for lsst cosmology,” *arXiv preprint arXiv:2212.04291*, 2022.
- [19] A. Albrecht, L. Amendola, G. Bernstein, D. Clowe, D. Eisenstein, L. Guzzo, C. Hirata, D. Huterer, R. Kirshner, E. Kolb, *et al.*, “Findings of the joint dark energy mission figure of merit science working group,” *arXiv preprint arXiv:0901.0721*, 2009.
- [20] D. Huterer, *A Course in Cosmology*. Cambridge University Press, 3 2023.
- [21] A. Porredon, M. Crocce, J. Elvin-Poole, R. Cawthon, G. Giannini, J. De Vicente, A. Carnero Rosell, I. Ferrero, E. Krause, X. Fang, *et al.*, “Dark energy survey year 3 results: Cosmological constraints from galaxy clustering and galaxy-galaxy lensing using the maglim lens sample,” *Physical Review D*, vol. 106, no. 10, p. 103530, 2022.
- [22] L. Wenzl, S.-F. Chen, and R. Bean, “Magnification bias estimators for realistic surveys: an application to the boss survey,” *Monthly Notices of the Royal Astronomical Society*, vol. 527, no. 2, pp. 1760–1773, 2024.
- [23] D. Coe, “Fisher matrices and confidence ellipses: a quick-start guide and software,” *arXiv preprint arXiv:0906.4123*, 2009.
- [24] N. Weaverdyck and D. Huterer, “Mitigating contamination in lss surveys: a comparison of methods,” *Monthly Notices of the Royal Astronomical Society*, vol. 503, no. 4, pp. 5061–5084, 2021.