University of Groningen

Density reconstruction from biased tracers and its application to primordial non-Gaussianity

Darwish, Omar; Foreman, Simon; Abidi, Muntazir M.; Baldauf, Tobias; Sherwin, Blake D.; Meerburg, P. Daniel

Published in:
Physical Review D

DOI:
[10.1103/PhysRevD.104.123520](https://doi.org/10.1103/PhysRevD.104.123520)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2021

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Darwish, O., Foreman, S., Abidi, M. M., Baldauf, T., Sherwin, B. D., & Meerburg, P. D. (2021). Density reconstruction from biased tracers and its application to primordial non-Gaussianity. *Physical Review D*, 104(12), [123520]. <https://doi.org/10.1103/PhysRevD.104.123520>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Density reconstruction from biased tracers and its application to primordial non-Gaussianity

Omar Darwish¹,[✉] Simon Foreman^{2,3},[✉] Muntazir M. Abidi,¹ Tobias Baldauf¹,[✉]
Blake D. Sherwin,^{1,4} and P. Daniel Meerburg⁵

¹*Department of Applied Mathematics and Theoretical Physics, University of Cambridge,
Wilberforce Road, Cambridge CB3 0WA, United Kingdom*

²*Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, Ontario N2L 2Y5, Canada*

³*Dominion Radio Astrophysical Observatory, Herzberg Astronomy & Astrophysics Research Centre,
National Research Council Canada, P.O. Box 248, Penticton, British Columbia V2A 6J9, Canada*

⁴*Kavli Institute for Cosmology, University of Cambridge, Cambridge CB3 0HA, United Kingdom*

⁵*Van Swinderen Institute for Particle Physics and Gravity, University of Groningen,
Nijenborgh 4, 9747 AG Groningen, Netherlands*

(Received 15 September 2020; accepted 4 October 2021; published 14 December 2021)

Large-scale Fourier modes of the cosmic density field are of great value for learning about cosmology because of their well-understood relationship to fluctuations in the early universe. However, cosmic variance generally limits the statistical precision that can be achieved when constraining model parameters using these modes as measured in galaxy surveys, and moreover, these modes are sometimes inaccessible due to observational systematics or foregrounds. For some applications, both limitations can be circumvented by reconstructing large-scale modes using the correlations they induce between smaller-scale modes of an observed tracer (such as galaxy positions). In this paper, we further develop a formalism for this reconstruction, using a quadratic estimator similar to the one used for lensing of the cosmic microwave background. We incorporate nonlinearities from gravity, nonlinear biasing, and local-type primordial non-Gaussianity, and verify that the estimator gives the expected results when applied to N -body simulations. We then carry out forecasts for several upcoming surveys, demonstrating that, when reconstructed modes are included alongside directly observed tracer density modes, constraints on local primordial non-Gaussianity are generically tightened by tens of percents compared to standard single-tracer analyses. In certain cases, these improvements arise from cosmic variance cancellation, with reconstructed modes taking the place of modes of a separate tracer, thus enabling an effective “multitracer” approach with single-tracer observations.

DOI: [10.1103/PhysRevD.104.123520](https://doi.org/10.1103/PhysRevD.104.123520)

I. INTRODUCTION

Our understanding of the Universe has benefited tremendously from measurements of the cosmic microwave background (CMB), primarily because of the linear relationship between fluctuations in the CMB and fluctuations generated in the very early universe. This relationship allows us to connect CMB measurements to the statistics of the initial fluctuations and their time evolution, and has led to the establishment of the current cosmological model. Extraction of similar information from the large-scale structure (LSS) of the universe is limited by nonlinear clustering at smaller distances and lower redshifts, requiring more elaborate modeling to interpret observations. This modelling burden is greatly reduced at the largest distances we can resolve with galaxy surveys, but this regime is in turn obscured by both statistical and systematic errors. In this paper, we explore a method to access these large scales while bypassing both types of errors: quadratic density reconstruction.

This idea of density reconstruction relies on the fact that a fixed long-wavelength density fluctuation correlates two different small-scale modes due to nonlinear evolution and higher-order biasing, with the amount of correlation proportional to the long-wavelength mode. This can be understood as arising from a violation of statistical homogeneity if the long-wavelength mode is considered as fixed and the shorter-wavelength modes are averaged over in an ensemble. Writing down a quadratic estimator that probes this induced correlation between two different modes, we can estimate the long-wavelength modes from the statistical properties of the smaller-scale modes.¹

¹In fact, these statements are independent of the relative wavelengths of the modes, and the formalism we present in this paper is not restricted to the so-called “squeezed limit” of the three modes involved. However, for our applications, the modes we are seeking to reconstruct have longer wavelengths than the two modes whose correlations are used for the reconstruction, so we focus on that situation in this paper.

There is a close analogy between this procedure and the common method of CMB lensing reconstruction, in which a quadratic estimator, making use of the lensing-induced correlation between two different CMB temperature modes, is used to reconstruct the lensing field (e.g., [1]). It is using this analogy that many of the methods for density reconstruction were derived. The idea of using a standard quadratic estimator in the CMB lensing form to perform this reconstruction was first proposed by [2], building on earlier work ([3–5], albeit with a somewhat modified estimator). Significant further work in this area has been presented by [6–10]; see further discussion in Sec. V.

The work in this paper broadly divides in two parts. In the first part, we present a considerable expansion of current technology for density reconstruction. We discuss the application of density reconstruction to biased tracers, including, for the first time, a full nonlinear bias model in such a formalism. We further validate our method on a suite of realistic N -body simulations, demonstrating that our methods perform just as expected from theoretical calculations for both the reconstruction and its noise level.

We note that this reconstruction has a wealth of applications. One simple application is the following: LSS surveys are often plagued by observational systematics that manifest at large scales, impeding the direct observation of low- k modes. Galaxy and quasar surveys are affected, for example, by variations in the density of foreground stars, seeing, and galactic dust extinction (e.g., [11–13]), while 21 cm surveys cannot access modes with low line-of-sight wave numbers that are dominated by galactic foregrounds, and imperfect knowledge of the instrument can spread this contamination throughout a wider region of Fourier space (e.g., [14–16]). A method of reconstructing these inaccessible modes using correlations between smaller-scale modes will improve the constraining power of a given survey for large-scale signals such as local non-Gaussianity, and allow cross-correlations involving 21 cm surveys that would otherwise be impossible (e.g., [6]). In this paper, we parametrize large-scale systematics with a wave number $K_{\min}$ below which the tracer modes are assumed to be inaccessible, and explore the precision with which modes with $K < K_{\min}$ can be recovered by our estimator. We note that this assumes that the relevant systematics can be parametrized as a large-scale additive component, rather than a possible modulation that might also significantly affect small scales; while there is evidence that this is a reasonable assumption for some of the currently known systematics (e.g., [13]), it may not hold in all cases.

In contrast to this general application, the second goal of our paper is to explore, in detail, a much more subtle application of density reconstruction: improving constraints on local-type primordial non-Gaussianity. We will briefly motivate the measurement of primordial

non-Gaussianity and the utility of density reconstruction for improving these constraints in the following paragraphs.

The CMB has taught us that the statistics of the primordial fluctuations can be accurately described by a red-tilted power law. If the initial conditions are completely described by this power law, they have to be Gaussian distributed, with statistics determined by only two degrees of freedom: the amplitude (A_s) and tilt (n_s) of the power law. If this is the case, however, it will be difficult to reach beyond our current understanding of the early Universe. The most widely accepted theory is known as cosmic inflation, which postulates a short early period of accelerated cosmic expansion. Effectively, the constraints we derive from the CMB tell us that inflation can be very well described by a scalar field slowly rolling down a potential (“single-field slow roll,” or SFSR), with only (weak) gravitational interactions. While such a model is certainly possible (it was the first to be considered [17–19]), it will not provide us with simple opportunities to understand the physics of inflation. If a proposed model of the early Universe has to comply with Gaussian initial conditions, effectively the model will observationally resemble SFSR. Any further distinction could be extracted from the details of the scale dependence of the primordial power spectrum [20], but so far, observations do not reveal any obvious deviations from a single parameter power law [21,22].

A much more powerful model discriminator would be available if the initial conditions showed a (small) deviation from Gaussianity. In the presence of non-Gaussianity, all moments beyond the power spectrum will generically be excited (starting with the 3-point function or bispectrum). Technically, these higher-point spectra probe the dynamics of the field(s) driving inflation. As a result, a measurement of non-Gaussianity would reveal details of inflation that can be directly related to the underlying fundamental physics. For example, non-Gaussianity could reveal the presence of more fields relevant during inflation, or could provide clues to how strongly coupled the inflation field is (see e.g., [23] and references therein). These powerful constraints cannot be exposed through any other measurement, making non-Gaussianities a unique probe of the early Universe.

To lowest order, primordial non-Gaussianities modulate the gravitational potential Φ via

$$\Phi(\mathbf{k}) = \varphi_G(\mathbf{k}) + f_{\text{NL}}^X \int \frac{d^3q}{(2\pi)^3} G_{\text{NL}}^X(\mathbf{q}, \mathbf{k} - \mathbf{q}) \varphi_G(\mathbf{q}) \varphi_G(\mathbf{k} - \mathbf{q}), \quad (1)$$

where φ_G is the Gaussian potential and G_{NL}^X is a kernel that describes how the potential is modulated. In this paper, we are interested in local non-Gaussianities for which $G_{\text{NL}}^{\text{local}} = 1$, i.e.,

$$\Phi(\mathbf{x}) = \varphi_G(\mathbf{x}) + f_{\text{NL}}(\varphi_G^2(\mathbf{x}) - \langle \varphi_G^2 \rangle), \quad (2)$$

where we have subtracted the mean to yield zero expectation value for the fluctuations and have renamed $f_{\text{NL}}^{\text{local}}$ to f_{NL} . Current constraints set $\sigma(f_{\text{NL}}) \sim \mathcal{O}(5)$ [22], while $f_{\text{NL}} \sim 1$ has been identified as a compelling theoretical threshold [24] which provides a strong motivation to go beyond current limits: if a measurement is made showing f_{NL} above this limit, it would effectively rule out SFSR inflation as a viable scenario. Future ground-based CMB experiments [25,26] may be able reach $\sigma(f_{\text{NL}}) \sim \mathcal{O}(1)$, but poor scaling and galactic and cosmological foregrounds will likely prevent the CMB from reaching (far) beyond this limit. Fortunately, the large scale structure (LSS) in the universe provides access to many more modes, for which $\sigma(f_{\text{NL}}) \propto (k_{\text{max}}^3 \log k_{\text{max}}/k_{\text{min}})^{-1/2}$ [27]. While increased dimensionality will help to improve constraints, the use of LSS will introduce many complications. For one, the scaling argument breaks down when k_{max} exceeds the nonlinear scale k_{NL} , which is of order $0.2h \text{ Mpc}^{-1}$ for current galaxy surveys [28,29]. Furthermore, line-of-sight information, which will be crucial in obtaining a sufficient number of modes, will require a careful treatment mainly due to redshift space effects [30]. Obtaining cosmological constraints from a measurement of the full LSS bispectrum will therefore be challenging, not least because of non-Gaussian covariance [27,31,32] which will likely require (a large number of) simulations to estimate [33]. Some of these difficulties can be overcome by simplifying the full bispectrum into more compressed statistics [34–41]. The advantage of these statistics is that they should capture nearly all the information [38], but are computationally and observationally less challenging.

Unlike in the CMB, in LSS local primordial non-Gaussianity can also significantly affect the power spectrum of biased tracers, such as galaxies. Specifically, it has been shown [42–46] that tracer bias will be affected by the primordial non-Gaussianity, with the bias acquiring a unique $1/k^2$ contribution, which is hard to produce otherwise. This signature has been used to place constraints on f_{NL} with current surveys [47–49]. Unfortunately, although the signal should be distinguishable from other effects, on large scales, the precision with which we can measure the power spectrum is ultimately limited by cosmic variance from the number of available modes. However, it was shown that this cosmic variance can be mitigated [50–54] by using multiple tracers of the same underlying density field (with different biases), which essentially allows a measurement of scale-dependent bias via a mode-by-mode comparison of the different tracers. A combination of two (or more, e.g., [53,55]) tracers will allow for cosmic variance cancellation, limiting a measurement of the scale-dependent bias from local primordial non-Gaussianity only by the number density of these tracers. Forecasts show that these techniques enable constraints to reach $\sigma(f_{\text{NL}}) \sim 1$ this decade [25,53,55,56].

In this paper, we show that this cosmic variance cancellation can also be achieved, to some extent, using only a single tracer. In order to do this, we compare our reconstructed density field (which provides information from higher-point functions) with a directly-measured tracer field. In the end, the constraints on f_{NL} will depend on the autocorrelation of the tracer field P_{gg} ,² the cross correlation of the tracer and the reconstructed field P_{gr} , and the autocorrelation of the reconstructed field P_{rr} . This idea is related to [40], where similar ideas are used to simplify a forecast of the combined information in the power spectrum, bispectrum and trispectrum. However, unlike in [40], we examine the reconstruction approach as a possible analysis tool rather than a method for more easily computing complex forecasts. In addition, whereas [40] relies on an extension of position-dependent power spectra [57–60], which draw information only from the squeezed limit, here we use a quadratic estimator formalism for the reconstructed field without imposing a squeezed-limit constraint.

Let us briefly summarize our most important results:

- (i) The modes of the tracer overdensity will be coupled due to nonlinearities from gravity, nonlinear bias, and primordial non-Gaussianity. The amplitudes (parameterized with bias coefficients) of several of these mode-couplings are unknown *a priori*. We incorporate this in our characterization of the quadratic estimator for long-wavelength modes, and marginalize over the unknown coefficients in our forecasts. We also highlight the important contribution of tracer shot noise to the noise on the reconstructed modes.
- (ii) We demonstrate density reconstruction using dark matter halos in N -body simulations, verifying that the performance agrees well with that predicted from analytical formulas. Though additional work using simulations will be required for a practical analysis, our results indicate that our forecasts are realistic.
- (iii) We show that the quadratic estimator is able to reconstruct long-wavelength modes at high signal-to-noise for a wide range of upcoming surveys (see Fig. 8).
- (iv) The addition of the reconstructed field to forecasts using the large-scale biased tracer field can improve constraints on f_{NL} by tens of percents depending on the survey configuration. The improvement arises from a combination of two sources: sample variance cancellation of signal in the large-scale tracer field, and additional scale-dependent signal in the reconstructed field on scales where the tracer field may be obscured by observational systematics.

²Since we will focus on the use of galaxies as tracers in this paper, we will use the subscript g to refer to these tracers, although the method we describe is equally applicable to quasars, line intensity maps, or other tracers.

TABLE I. Notation used for important quantities in this paper.

Quantity	Symbol	Defined in
Dirac delta function in 3d	$\delta_D(\mathbf{k})$	—
Wave numbers of modes used in reconstruction	$\mathbf{k}, \mathbf{q}$, etc.	—
Wave numbers of modes used for f_{NL} constraints	$\mathbf{K}, \mathbf{K}'$, etc.	—
Amplitude of local primordial non-Gaussianity	f_{NL}	Eq. (2)
Factor relating primordial potential and δ_1	$M(k, z)$	Eqs. (18)–(19)
Linear matter overdensity	$\delta_1(\mathbf{k}, z)$	—
Linear matter power spectrum	$P_{\text{lin}}(k, z)$	—
Tracer overdensity	$\delta_g(\mathbf{k}, z)$	Eq. (3) [generic]; Eq. (36) [second-order bias model]
Second-order mode-coupling	$F_\alpha(\mathbf{k}_1, \mathbf{k}_2)$	Eq. (3) [generic]; Eq. (36) [second-order bias model]
Second-order response of small-scale power spectrum to long mode	$f_\alpha(\mathbf{k}_1, \mathbf{k}_2, z)$	Eq. (9)
Coefficient of F_α in second-order bias model for δ_g	c_α	Eq. (36)
Linear tracer bias	$b_1 \equiv b_{10}^E$	Eq. (36)
Quadratic tracer bias	$b_2 \equiv b_{20}^E$	Eq. (36)
Other second-order bias parameters	$b_{s^2}^E, b_{01}^E, \dots$	Sec. II B
Quadratic estimator for mode with wave number $\mathbf{K}$	$\hat{\Delta}_\alpha(\mathbf{K})$	Eqs. (10), (16)
Weight function in $\hat{\Delta}_\alpha(\mathbf{K})$	$g_\alpha(\mathbf{k}_1, \mathbf{k}_2)$	Eq. (14)
Normalization and Gaussian noise of $\hat{\Delta}_\alpha(\mathbf{K})$	$N_{\alpha\beta}(\mathbf{K})$	Eq. (15)
Mode reconstructed with growth-coupling estimator $\hat{\Delta}_G(\mathbf{K}, z)$	$\delta_r(\mathbf{k}, z)$	—
Power spectrum of δ_g , ignoring shot noise contribution	P_{gg}	—
Sum of P_{gg} and shot noise contribution	P_{tot}	Sec. II A
Cross power spectrum between δ_g and δ_r , ignoring shot noise contribution	P_{gr}	—
Power spectrum of δ_r , ignoring shot noise contribution	P_{rr}	—
Shot noise contribution to δ_g power spectrum	$P_{\text{gg,shot}}$	Eq. (49)
Shot noise contribution to δ_r power spectrum	$P_{\text{rr,shot}}$	Eqs. (B15)–(B16)
Shot noise contribution to $\delta_g - \delta_r$ cross power spectrum	$P_{\text{gr,shot}}$	Eqs. (B20)–(B21)
Lowest wave number within survey volume	K_f	Sec. IV C 1
Wave number below which we assume δ_g cannot be measured	K_{min}	Sec. IV C 1
Maximum wave number used for f_{NL} constraints	K_{max}	Sec. IV C 1
Maximum wave number used in quadratic estimator for reconstructed modes	k_{max}	Sec. IV C 1

The additional information in the reconstructed modes can be viewed as a signature of non-Gaussian signal in the three and four-point functions, and our approach can be viewed as a simple method to obtain combined information from the three- and four-point functions and the power spectrum.

- (v) The performance of this approach to constraining f_{NL} is limited by a combination of tracer number density and maximum wave number of modes that can be used for reconstruction, with the details again depending on the survey configuration. Potential improvements using response function approaches [61,62] could be explored to extend the reconstruction wave number and gain signal-to-noise.

The outline of our paper is as follows. In Sec. II, we describe our methodology for density reconstruction, including the quadratic estimator formalism and bias expansion we use. In Sec. III, we apply this method to halos in N -body simulations. In Sec. IV, we present our forecasts for the expected precision on reconstructed

modes, as well as constraints on local non-Gaussianity. We compare this reconstruction formalism to other work involving higher-point statistics in Sec. V. Finally, we conclude in Sec. VI. Several derivations and technical details are included in the Appendixes, and a summary of our notation can be found in Table I. Except for in Sec. III, we use cosmological parameters from the Planck 2015 results, given in the “TT, TE, EE + lowP + lensing + ext” column of Table 4 of [63].

II. DENSITY RECONSTRUCTION

A. Quadratic estimator: General case

In this section, we will develop the general formalism for reconstructing large-scale³ density modes using observations of a biased tracer. This is largely based on the treatment in [2], but we have adapted their expressions

³We again remind the reader that the our formalism is generally applicable, without any strong assumptions on the wavelengths of the modes.

to 3D wave numbers rather than a separate treatment of line-of-sight and transverse components of $\mathbf{k}$.

Suppose that the overdensity field of the tracer, δ_g , is well-described by a linear bias with respect to the linear matter overdensity δ_1 , plus a set of quadratic terms that couple modes of δ_1 with kernels F_α and amplitudes c_α :

$$\delta_g(\mathbf{k}, z) \approx b_1(z) \delta_1(\mathbf{k}, z) + \sum_\alpha c_\alpha(z) \int_{\mathbf{q}} F_\alpha(\mathbf{q}, \mathbf{k} - \mathbf{q}; z) \delta_1(\mathbf{q}, z) \delta_1(\mathbf{k} - \mathbf{q}, z), \quad (3)$$

where $\int_{\mathbf{q}} \equiv (2\pi)^{-3} \int d^3\mathbf{q}$. For example, if we took δ_g to be the matter overdensity rather than a biased tracer, we would have $b_1 = 1$ and the sum would run over the second-order mode-couplings induced by gravitational evolution, which take the form (e.g., [64])

$$F_G(\mathbf{k}_1, \mathbf{k}_2; z) \equiv \frac{17}{21}, \quad F_S(\mathbf{k}_1, \mathbf{k}_2; z) \equiv \frac{1}{2} \left(\frac{1}{k_1^2} + \frac{1}{k_2^2} \right) \mathbf{k}_1 \cdot \mathbf{k}_2, \\ F_T(\mathbf{k}_1, \mathbf{k}_2; z) \equiv \frac{2}{7} \left[\frac{(\mathbf{k}_1 \cdot \mathbf{k}_2)^2}{k_1^2 k_2^2} - \frac{1}{3} \right], \quad (4)$$

with $c_G = c_S = c_T = 1$ and the subscripts indicating that these functions arise from isotropic **G**rowth, a large-scale coordinate **S**hift, and a **T**idel coupling. For a biased tracer, nonlinear biasing will lead to $c_\alpha \neq 1$ for the above couplings, and primordial non-Gaussianity will introduce additional mode-couplings. In Sec. II B, we will introduce the full set of mode-couplings that must be considered, but we note here that many of the corresponding c_α coefficients will not be known *a priori*, and this must be accounted for in the density reconstruction procedure. Henceforth, we will drop the z -dependence from the quantities defined above.

Now, we would like to use the mode-couplings in Eq. (3) to construct a quadratic estimator for a given mode of δ_1 .

We will present the logic in some detail, for readers who may not be familiar with the relevant arguments, but a reader who is comfortable with peak-background-split arguments or the CMB lensing formalism may wish to skip to the final result in Eqs. (8)–(9).

The analogous procedure for CMB lensing is to first consider an ensemble average over CMB fluctuations while keeping fluctuations in the lower-redshift matter density fixed. In this case, the fixed modes of the lensing potential ϕ (which is a line-of-sight projection of the lower-redshift density field—see e.g., [1]) break the statistical isotropy of the CMB fluctuations, inducing correlations between CMB fluctuation modes with different wave numbers: for temperature modes on the flat sky, the specific effect is given by

$$\langle T(\boldsymbol{\ell}) T(\mathbf{L} - \boldsymbol{\ell}) \rangle_{\phi \text{ fixed}} = (2\pi)^2 \delta_D(\mathbf{L}) C_L + f_\phi(\boldsymbol{\ell}, \mathbf{L} - \boldsymbol{\ell}) \phi(\mathbf{L}). \quad (5)$$

When analyzing CMB simulations or data, the temperature two-point function is estimated by a (weighted) sum over $\boldsymbol{\ell}$ within a given CMB realization, and this in fact approximates the ensemble average above, with ϕ modes effectively fixed because they do not explicitly enter the sum. Equation (5) is an efficient starting point for deriving quadratic estimators for a specific mode of ϕ , and we would like to find the analogous starting point for density reconstruction.

To proceed, we consider an ensemble average over all modes of δ_1 except those with wave numbers in a small neighborhood around $\mathbf{K}$, with $\delta_1(\mathbf{K})$ being the mode we will eventually want to reconstruct. (We must consider a neighborhood around $\mathbf{K}$ because we are working in the continuum limit, where we have integrals instead of discrete sums over wave numbers; we will return to this point below.) In this ensemble average, which we denote by “ $\sim \mathbf{K}$ fixed”, and using Eq. (3), the two-point function of δ_g is at next-to-leading-order in δ_1 is

$$\langle \delta_g(\mathbf{k}) \delta_g(\mathbf{K} - \mathbf{k}) \rangle_{\sim \mathbf{K} \text{ fixed}} = b_1^2 \langle \delta_1(\mathbf{k}) \delta_1(\mathbf{K} - \mathbf{k}) \rangle + b_1 \int_{\mathbf{q}} \sum_\alpha c_\alpha F_\alpha(\mathbf{q}, \mathbf{k} - \mathbf{q}) \langle \delta_1(\mathbf{q}) \delta_1(\mathbf{k} - \mathbf{q}) \delta_1(\mathbf{K} - \mathbf{k}) \rangle_{\sim \mathbf{K} \text{ fixed}} + [\mathbf{k} \leftrightarrow \mathbf{K} - \mathbf{k}]. \quad (6)$$

In the first line, we have assumed that $\mathbf{k}$ is not within the chosen neighborhood of $\mathbf{K}$ or the equivalent neighborhood of 0, so there is no difference between our special ensemble average and the standard one. In the second line, the integrand evaluates to zero if $\mathbf{q}$ and $\mathbf{K} - \mathbf{q}$ are not within the neighborhood of $\mathbf{K}$, since in that case, all three δ_1 modes are

averaged over, and the three-point function is zero for $\mathbf{K} \neq 0$. When $\mathbf{q} \sim \mathbf{K}$ or $\mathbf{K} - \mathbf{q} \sim \mathbf{K}$, where we use “ $\sim \mathbf{K}$ ” to indicate a vector falling within the neighborhood of $\mathbf{K}$, then $\delta_1(\mathbf{q})$ or $\delta_1(\mathbf{K} - \mathbf{q})$ factor out of the ensemble average because they are held fixed, and the remaining two modes are averaged over:

$$\begin{aligned} \langle \delta_g(\mathbf{k}) \delta_g(\mathbf{K}-\mathbf{k}) \rangle_{\sim \mathbf{K} \text{ fixed}} &= b_1^2 \langle \delta_1(\mathbf{k}) \delta_1(\mathbf{K}-\mathbf{k}) \rangle + b_1 \int_{\mathbf{q} \sim \mathbf{K}} \sum_{\alpha} c_{\alpha} F_{\alpha}(\mathbf{q}, \mathbf{k}-\mathbf{q}) \delta_1(\mathbf{q}) \langle \delta_1(\mathbf{k}-\mathbf{q}) \delta_1(\mathbf{K}-\mathbf{k}) \rangle + [\mathbf{k} \leftrightarrow \mathbf{K}-\mathbf{k}] \\ &+ b_1 \int_{\mathbf{K}-\mathbf{q} \sim \mathbf{K}} \sum_{\alpha} c_{\alpha} F_{\alpha}(\mathbf{q}, \mathbf{k}-\mathbf{q}) \delta_1(\mathbf{K}-\mathbf{q}) \langle \delta_1(\mathbf{q}) \delta_1(\mathbf{K}-\mathbf{k}) \rangle + [\mathbf{k} \leftrightarrow \mathbf{K}-\mathbf{k}]. \end{aligned} \quad (7)$$

From here, we simply evaluate the two-point correlators and use the resulting Dirac delta functions to collapse the $\mathbf{q}$ integrals.⁴ The final result is

$$\begin{aligned} \langle \delta_g(\mathbf{k}) \delta_g(\mathbf{K}-\mathbf{k}) \rangle_{\sim \mathbf{K} \text{ fixed}} \\ = (2\pi)^3 \delta_D(\mathbf{K}) b_1^2 P_{\text{lin}}(k) + b_1 \sum_{\alpha} c_{\alpha} f_{\alpha}(\mathbf{k}, \mathbf{K}-\mathbf{k}) \delta_1(\mathbf{K}), \end{aligned} \quad (8)$$

where

$$f_{\alpha}(\mathbf{k}_1, \mathbf{k}_2) \equiv 2[F_{\alpha}(\mathbf{k}_1 + \mathbf{k}_2, -\mathbf{k}_1) P_{\text{lin}}(\mathbf{k}_1, z) + 1 \leftrightarrow 2]. \quad (9)$$

In Eq. (8), we find the same structure as in the CMB lensing case in Eq. (5): the standard power spectrum term, plus a term from off-diagonal correlations induced by the fixed background mode.

Equation (8) suggests we can multiply two different modes of the measured tracer field and then simply “divide” by the coupling strength $b_1 \sum_{\alpha} c_{\alpha} f_{\alpha}(\mathbf{k}, \mathbf{K}-\mathbf{k})$ to obtain an estimate of the linear field δ_1 at large scales. Unfortunately, in general we do not know the bias coefficients b_1 or c_{α} *a priori*, so the best we can do is to use the galaxy mode couplings to estimate the product $b_1 c_{\alpha} \delta_1$ for a chosen α . To reduce variance on the estimate, we will sum over all the mode couplings that involve the same large-scale mode. This can be achieved by writing the following general quadratic estimator

$$\hat{\Delta}_{\alpha}(\mathbf{K}) \equiv b_1 \widehat{c_{\alpha} \delta_1}(\mathbf{K}) = \int_{\mathbf{q}} g_{\alpha}(\mathbf{q}, \mathbf{K}-\mathbf{q}) \delta_g(\mathbf{q}) \delta_g(\mathbf{K}-\mathbf{q}), \quad (10)$$

with weights g_{α} , similar to what is done for CMB lensing [1] or “clustering fossils” from primordial gravitational waves [65–67]. For an alternative derivation of this estimator, based on optimizing the cross-correlation of a quadratic combination of measured modes with the true linear mode to be reconstructed, see Appendix A.

⁴One must integrate in a neighborhood around the argument of a Dirac delta function for this collapse to take place, and this is why we considered a neighborhood around $\mathbf{K}$ in the first place. In the discrete case, where we have sums instead of integrals over wave numbers, we could define our ensemble average to keep a single mode $\delta_1(\mathbf{K})$ fixed, since we would then have Kronecker deltas instead of Dirac delta functions.

The covariance between two such estimators α and β of the biased matter density field on large scales can be split into a Gaussian part, coming from all disconnected contributions, and a non-Gaussian part that includes all connected contributions:

$$\begin{aligned} \langle \hat{\Delta}_{\alpha}(\mathbf{K}) \hat{\Delta}_{\beta}^*(\mathbf{K}') \rangle - \langle \hat{\Delta}_{\alpha}(\mathbf{K}) \rangle \langle \hat{\Delta}_{\beta}^*(\mathbf{K}') \rangle \\ = (2\pi)^3 \delta_D(\mathbf{K}-\mathbf{K}') [\text{Cov}_G(\hat{\Delta}_{\alpha}(\mathbf{K}), \hat{\Delta}_{\beta}^*(\mathbf{K}')) \\ + \text{Cov}_{\text{NG}}(\hat{\Delta}_{\alpha}(\mathbf{K}), \hat{\Delta}_{\beta}^*(\mathbf{K}'))]. \end{aligned} \quad (11)$$

We constrain the weights to provide an estimator that is optimal in the sense of minimizing the Gaussian contributions to its variance,

$$\text{Var}_G[\hat{\Delta}_{\alpha}](\mathbf{K}) \equiv \text{Cov}_G(\hat{\Delta}_{\alpha}(\mathbf{K}), \hat{\Delta}_{\alpha}^*(\mathbf{K}')), \quad (12)$$

while requiring that it be unbiased if there were only a single mode-coupling, i.e.,

$$\int_{\mathbf{q}} g_{\alpha}(\mathbf{q}, \mathbf{K}-\mathbf{q}) f_{\alpha}(\mathbf{q}, \mathbf{K}-\mathbf{q}) = 1. \quad (13)$$

These criteria lead to the familiar quadratic estimator weights:

$$g_{\alpha}(\mathbf{k}_1, \mathbf{k}_2) = N_{\alpha\alpha}(\mathbf{k}_1 + \mathbf{k}_2) \frac{f_{\alpha}(\mathbf{k}_1, \mathbf{k}_2)}{2P_{\text{tot}}(k_1)P_{\text{tot}}(k_2)}, \quad (14)$$

where P_{tot} is the sum of the clustering and shot noise contributions to the tracer power spectrum. The normalization is given by

$$N_{\alpha\beta}(\mathbf{K}) = \left(\int_{\mathbf{q}} \frac{f_{\alpha}(\mathbf{q}, \mathbf{K}-\mathbf{q}) f_{\beta}(\mathbf{q}, \mathbf{K}-\mathbf{q})}{2P_{\text{tot}}(q)P_{\text{tot}}(|\mathbf{K}-\mathbf{q}|)} \right)^{-1} \quad (15)$$

which guarantees that $N_{\alpha\alpha}$ is equal to the Gaussian part of the variance of $\hat{\Delta}_{\alpha}$. We will refer to $N_{\alpha\alpha}$ as the reconstruction noise, which incorporates cosmic variance in the reconstruction and the disconnected contribution from shot noise of the tracer field. (It should be noted that $N_{\alpha\beta}$ is not equal to the noise when $\alpha \neq \beta$.) With the weights in Eq. (14), the estimator in Eq. (10) becomes

$$\hat{\Delta}_\alpha(\mathbf{K}) = N_{\alpha\alpha}(\mathbf{K}) \int_{\mathbf{q}} \frac{f_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q})}{2P_{\text{tot}}(q)P_{\text{tot}}(|\mathbf{K} - \mathbf{q}|)} \delta_g(\mathbf{q}) \delta_g(\mathbf{K} - \mathbf{q}). \quad (16)$$

The non-Gaussian part of the variance includes a trispectrum contribution from clustering of the tracers, and further contributions from tracer shot noise. We neglect the former, because it is subdominant to the latter; our comparisons with simulations in Sec. III show that this is a valid approximation. Importantly, the shot noise contributions can dominate over the Gaussian reconstruction noise in many cases, because these contributions couple to large-scale modes with large variance, while the Gaussian contribution only involves small-scale modes, which have smaller variance due to the shape of the matter power spectrum. We derive the full expressions for these contributions and discuss their hierarchy further in Appendix B.

The expectation value of the estimator in Eq. (16), for a given realization of the linear field at wave vector $\mathbf{K}$, is

$$\langle \hat{\Delta}_\alpha(\mathbf{K}) \rangle_{\delta_1(\mathbf{K}) \text{ fixed}} = b_1 \left[c_\alpha + \sum_{\beta \neq \alpha} c_\beta \frac{N_{\alpha\alpha}(\mathbf{K})}{N_{\alpha\beta}(\mathbf{K})} \right] \delta_1(\mathbf{K}). \quad (17)$$

We clearly see that there is a contamination of the estimator with respect to the case of only a single mode-coupling, given by the product of the (Gaussian) noise for the estimator α and a sum of bias terms divided by the cross normalization between estimators α and β .⁵ If the goal is to just reconstruct the linear mode of interest, then it is important to account for this contribution. One can attempt to construct a so-called “bias-hardened” estimator by forming a linear combination of the original estimators that is free of this contamination at leading order (e.g., [2,68,69]). However, for the specific mode-couplings relevant in this situation, the high degree of correlation between the original estimators implies that the noise on the new estimator will be so high that it is no longer useful; see Appendix C for details.

We claim that, for extracting non-Gaussianity, this contamination can actually be useful. As we will see later, some of these contaminating terms induce scale-dependence that reproduces the $1/K^2$ scaling created by primordial non-Gaussianity. Depending of the signs of these terms, they can either raise or lower the signal to noise on f_{NL} from the reconstructed field. We will discuss this further in Sec. II C.

⁵It can be seen from Eq. (15) that as the overlap integral of the two mode-couplings goes to zero, $N_{\alpha\beta}$ becomes very large and the contamination vanishes.

B. Non-Gaussianity and bias expansion

As we discussed in the introduction, primordial non-Gaussianity of the local type introduces a quadratic contribution to the metric perturbation. The metric perturbation (gravitational potential) φ is related to the linear matter overdensity through the usual Poisson equation (dropping the subscript G)

$$\varphi(\mathbf{k}, z) = \frac{\delta_1(\mathbf{k}, z)}{M(\mathbf{k}, z)}, \quad (18)$$

where the Poisson factor $M(k, z)$ is given by

$$M(k, z) = \frac{2c^2}{3H_0^2 \Omega_m} D(z) k^2 T(k). \quad (19)$$

Here, the growth factor $D(z)$ is normalized to agree with the scale factor $1/(1+z)$ during matter domination.

Galaxies and 21 cm fluctuations of the density field are biased tracers of the underlying, dynamically dominant matter distribution. In the presence of local primordial non-Gaussianity, the coupling of the long and short modes leads to an additional modulation of the abundance of collapsed objects by the long wavelength potential fluctuations φ . To describe biased tracers we thus follow [70,71] in performing a double expansion of the Eulerian galaxy (or tracer) density field in the nonlinear density and linear potential⁶

$$\begin{aligned} \delta_g^E(\mathbf{x}) = & b_{10}^E \delta(\mathbf{x}) + b_{01}^E \varphi(\mathbf{x}_L[\mathbf{x}]) + b_{20}^E \delta^2(\mathbf{x}) + b_{11}^E \delta(\mathbf{x}) \varphi(\mathbf{x}_L[\mathbf{x}]) \\ & + b_{02}^E \varphi^2(\mathbf{x}) + b_{s2}^E s_{ij}(\mathbf{x}) s^{ij}(\mathbf{x}) + \varepsilon(\mathbf{x}) \\ & + \varepsilon_\delta(\mathbf{x}) \delta(\mathbf{x}) + \varepsilon_\varphi(\mathbf{x}) \varphi(\mathbf{x}_L[\mathbf{x}]) + \dots \end{aligned} \quad (20)$$

Here the b_{ij}^E are the Eulerian bias parameters,⁷ s_{ij} is the tidal tensor

$$s_{ij}(\mathbf{x}) = \left[\frac{\nabla_i \nabla_j}{\nabla^2} - \frac{1}{3} \delta_{ij}^{(K)} \right] \delta(\mathbf{x}), \quad (21)$$

and ε is the stochasticity, which correlates with itself but not with the linear density field. In the simplest case where galaxies are a Poisson sample of the underlying matter field, the stochasticity leads to the fiducial $1/\bar{n}$ power spectrum. The higher order stochasticity contributions $\varepsilon_\delta \delta$ and $\varepsilon_\varphi \varphi$ lead to stochasticity contributions in the bispectrum [72], as we review in Appendix B. In simple

⁶In the peak-background split formalism, the abundance of collapsed objects is given by $e^{-\nu^2/2}$ where $\nu = \delta_c/\sigma$ with δ_c the collapse threshold and σ the variance. The long wavelength density modulates the collapse threshold as $\delta_c \rightarrow \delta_c - \delta$, whereas the metric perturbation φ modulates the variance $\sigma \rightarrow \sigma(1 + 2f_{\text{NL}}\varphi)$.

⁷In the introductory discussion in Sec. II A, we employed the notation $b_1 \equiv b_{10}^E$ for the sake of simplicity.

local-Lagrangian bias models, the tidal tensor bias $b_{s_2}^E$ can be related to the linear density bias as $b_{s_2}^E = -2/7(b_{10}^E - 1)$ [73]. Employing realistic simplifying assumptions, we will see that all of the bias parameters b_{ij}^E can be expressed in terms of b_{10}^E and b_{20}^E . We are truncating the above expansion at second order, since we will only consider tree level power spectra and bispectra as well as the Gaussian disconnected trispectrum in our derivations. We can thus also neglect higher derivative contributions, such as $k^2\delta_1(\mathbf{k})$, as they are equivalent to cubic contributions to the matter and galaxy density fields. Note that all of the δ terms in Eq. (20) refer to the underlying nonlinear matter density field including its quadratic couplings. The potential φ , in turn, is linear, as the dependence of the halo abundance on long wavelength potential fluctuations is set up in the early Universe.

There is, however, a nonlinearity in the potential terms that arises from the fact that the abundance of galaxies in the peak-background split is set up in Lagrangian space with coordinates $\mathbf{x}_L$. These Lagrangian positions are related to the Eulerian coordinates by $\mathbf{x}_L[\mathbf{x}] = \mathbf{x} - \mathbf{\Psi}(\mathbf{x})$ at leading order. The potential is thus advected by long wavelength displacements as [74]

$$\varphi(\mathbf{x}_L[\mathbf{x}]) = \varphi(\mathbf{x}) - \mathbf{\Psi}(\mathbf{x}) \cdot \nabla\varphi(\mathbf{x}) + \dots \quad (22)$$

The Fourier transform of the linear displacement field $\mathbf{\Psi}(\mathbf{x})$ is related to the linear matter overdensity by $\mathbf{\Psi}(\mathbf{k}) = i(\mathbf{k}/k^2)\delta(\mathbf{k})$.

At second order, the matter density field picks up a new quadratic contribution from primordial non-Gaussianity according to Eq. (2):

$$\delta(\mathbf{k}) = \delta_1(\mathbf{k}) + \int_{\mathbf{q}} \left[\sum_{\alpha=G,S,T} F_{\alpha}(\mathbf{q}, \mathbf{k}-\mathbf{q}) \right] \delta_1(\mathbf{q}) \delta_1(\mathbf{k}-\mathbf{q}) + f_{\text{NL}} M(k) \int_{\mathbf{q}} \varphi(\mathbf{q}) \varphi(\mathbf{k}-\mathbf{q}) + \dots, \quad (23)$$

where the growth, shift and tidal components of the gravitational coupling kernel are given by Eq. (4). For biased tracers, this expression gets multiplied by b_{10}^E . We can rewrite the last term in terms of the density field using the Poisson equation, resulting in a new quadratic coupling

$$F_{\varphi\varphi}(\mathbf{k}_1, \mathbf{k}_2) = \frac{M(|\mathbf{k}_1 + \mathbf{k}_2|)}{M(k_1)M(k_2)}, \quad (24)$$

such that

$$\delta(\mathbf{k}) = \delta_1(\mathbf{k}) + \int_{\mathbf{q}} \left[\sum_{\alpha=G,S,T,\varphi\varphi} c_{\alpha} f_{\text{NL}}^{p_{\alpha}} F_{\alpha}(\mathbf{q}, \mathbf{k}-\mathbf{q}) \right] \delta_1(\mathbf{q}) \delta_1(\mathbf{k}-\mathbf{q}) + \dots, \quad (25)$$

where now $c_{\alpha} = \{1, 1, 1, 1\}$ and $p_{\alpha} = \{0, 0, 0, 1\}$.

Combining this result with the Fourier transform of the other second order bias terms in Eq. (20) yields

$$\begin{aligned} \delta_g^E(\mathbf{k}) = & \left[b_{10}^E + \frac{b_{01}^E}{M(k)} \right] \delta_1(\mathbf{k}) + b_{01}^E \int_{\mathbf{q}} \frac{1}{2} \left[\frac{\mathbf{q} \cdot (\mathbf{k}-\mathbf{q})}{q^2 M(|\mathbf{k}-\mathbf{q}|)} + \frac{\mathbf{q} \cdot (\mathbf{k}-\mathbf{q})}{|\mathbf{k}-\mathbf{q}|^2 M(q)} \right] \delta_1(\mathbf{q}) \delta_1(\mathbf{k}-\mathbf{q}) \\ & + b_{10}^E \int_{\mathbf{q}} \left[\sum_{\alpha=G,S,T,\varphi\varphi} F_{\alpha}(\mathbf{q}, \mathbf{k}-\mathbf{q}) \right] \delta_1(\mathbf{q}) \delta_1(\mathbf{k}-\mathbf{q}) \\ & + f_{\text{NL}} b_{10}^E \int_{\mathbf{q}} \frac{M(k)}{M(q)M(|\mathbf{k}-\mathbf{q}|)} \delta_1(\mathbf{q}) \delta_1(\mathbf{k}-\mathbf{q}) + b_{20}^E \int_{\mathbf{q}} \delta_1(\mathbf{q}) \delta_1(\mathbf{k}-\mathbf{q}) \\ & + b_{11}^E \int_{\mathbf{q}} \frac{1}{2} \left(\frac{1}{M(q)} + \frac{1}{M(|\mathbf{k}-\mathbf{q}|)} \right) \delta_1(\mathbf{q}) \delta_1(\mathbf{k}-\mathbf{q}) + b_{02}^E \int_{\mathbf{q}} \frac{1}{M(q)M(|\mathbf{k}-\mathbf{q}|)} \delta_1(\mathbf{q}) \delta_1(\mathbf{k}-\mathbf{q}) \\ & + b_{s_2}^E \int_{\mathbf{q}} \left[\frac{[\mathbf{q} \cdot (\mathbf{k}-\mathbf{q})]^2}{q^2 |\mathbf{k}-\mathbf{q}|^2} - \frac{1}{3} \right] \delta_1(\mathbf{q}) \delta_1(\mathbf{k}-\mathbf{q}). \end{aligned} \quad (26)$$

The additional terms arising from the non-Gaussian bias can be encoded by the new quadratic coupling kernels

$$F_{01} = \frac{1}{2} \mathbf{k}_1 \cdot \mathbf{k}_2 \left(\frac{1}{k_2^2 M(k_1)} + \frac{1}{k_1^2 M(k_2)} \right), \quad F_{11} = \frac{1}{2} \left(\frac{1}{M(k_1)} + \frac{1}{M(k_2)} \right), \quad F_{02} = \frac{1}{M(k_1)M(k_2)}. \quad (27)$$

The Eulerian bias parameters can be related to their Lagrangian counterparts through a spherical collapse calculation [70,71]:

TABLE II. Mode couplings, f_{NL} exponents, bias parameters and coupling kernels of the quadratic interactions for Eq. (36).

Mode coupling (α)	p_α	c_α	$F_\alpha(\mathbf{k}_1, \mathbf{k}_2)$
G	0	$b_1 + \frac{21}{17}b_2$	$\frac{17}{21}$
S	0	b_1	$\frac{1}{2} \left[\frac{1}{k_1^2} + \frac{1}{k_2^2} \right] (\mathbf{k}_1 \cdot \mathbf{k}_2)$
T	0	$b_1 + \frac{7}{2}b_{s^2}$	$\frac{2}{7} \left[\frac{(\mathbf{k}_1 \cdot \mathbf{k}_2)^2}{k_1^2 k_2^2} - \frac{1}{3} \right]$
$\varphi\varphi$	1	b_1	$\frac{M(\mathbf{k}_1 + \mathbf{k}_2 , z)}{M(k_1)M(k_2)}$
01	1	$2\delta_c(b_1 - 1)$	$\frac{1}{2} \mathbf{k}_1 \cdot \mathbf{k}_2 \left(\frac{1}{k_1^2 M(k_1)} + \frac{1}{k_2^2 M(k_2)} \right)$
11	1	$2(\delta_c \left[\frac{b_2 - 2(a_1 + a_2)(b_1 - 1)}{a_1} \right] - a_1[b_1 - 1]) + 2\delta_c(b_1 - 1)$	$\frac{1}{2} \left(\frac{1}{M(k_1)} + \frac{1}{M(k_2)} \right)$
02	2	$4\delta_c \left[\delta_c \left[\frac{b_2 - 2(a_1 + a_2)(b_1 - 1)}{a_1^2} \right] - 2[b_1 - 1] \right]$	$\frac{1}{M(k_1)M(k_2)}$

$$b_{10}^E = b_{10}^L + 1, \quad (28)$$

$$b_{20}^E = 2(a_1 + a_2)b_{10}^L + a_1^2 b_{20}^L, \quad (29)$$

$$b_{01}^E = b_{01}^L, \quad (30)$$

$$b_{11}^E = a_1 b_{11}^L + b_{01}^L, \quad (31)$$

$$b_{02}^E = b_{02}^L, \quad (32)$$

where $a_1 = 1$ and $a_2 = -17/21$ are spherical collapse expansion factors. The non-Gaussian Lagrangian bias parameters can be obtained using the peak background split. They are given as the derivatives of the mass function with respect to the long wavelength potential fluctuations. Assuming a universal mass function, the derivatives with respect to the potential can be related to the derivatives with respect to the long wavelength density, and consequently the bias parameters of the potential terms can be related to the bias parameters of the density terms:

$$b_{01}^L = 2f_{\text{NL}}\delta_c(b_{10}^E - 1), \quad (33)$$

$$b_{11}^L = 2f_{\text{NL}} \left(\delta_c \left[\frac{b_{20}^E - 2(a_1 + a_2)(b_{10}^E - 1)}{a_1^2} \right] - [b_{10}^E - 1] \right), \quad (34)$$

$$b_{02}^L = 4f_{\text{NL}}^2 \delta_c \left(\delta_c \left[\frac{b_{20}^E - 2(a_1 + a_2)(b_{10}^E - 1)}{a_1^2} \right] - 2[b_{10}^E - 1] \right), \quad (35)$$

where δ_c is the spherical collapse threshold. Note that small deviations from this simple scaling of non-Gaussian bias b_{01}^L with Gaussian bias b_{10}^E have been found in simulations [75] and seem to depend on the way halos are identified.

In summary, we can write for the galaxy density field up to second order in the presence of local type primordial non-Gaussianity:

$$\delta_g(\mathbf{k}) = \left[b_{10}^E + f_{\text{NL}} \frac{c_{01}}{M(k)} \right] \delta_1(\mathbf{k}) + \int_q \left[\sum_\alpha c_\alpha f_{\text{NL}}^{p_\alpha} F_\alpha(\mathbf{q}, \mathbf{k} - \mathbf{q}) \right] \delta_1(\mathbf{q}) \delta_1(\mathbf{k} - \mathbf{q}), \quad (36)$$

where α now runs over $\{G, S, T, \varphi\varphi, 01, 11, 02\}$ with the couplings given in Table II. In this table, Eq. (36), and throughout the rest of the paper, we have simplified the notation to $b_1 \equiv b_{10}^E$, $b_2 \equiv b_{20}^E$, and $b_{s^2} \equiv b_{s^2}^E$. Note that we have not included mode-couplings due to lensing, which are expected to be a subdominant contribution that is somewhat degenerate with the S term [2], nor have we incorporated redshift space distortions or anisotropic selection effects (see Sec. IV C 5 for discussion).

C. Reconstruction noise and contamination

With this formalism in place, we can now examine the noise of the reconstructed modes, and the contamination arising from the presence of multiple mode-couplings in the tracer field used for reconstruction.⁸ We will show these quantities for a DESI-like survey (with specifications given in Sec. IV C), but we have checked that the conclusions we draw from this case also apply to the other surveys we consider.

Figure 1 shows the reconstruction noise power spectrum corresponding to estimators that use each of the quadratic mode-couplings discussed in Sec. II B. We see that the “growth” estimator has the lowest noise by far. We compare the predicted noise for the G, S, and T estimators with results from N -body simulations in Sec. III (among other tests), finding good agreement. Thus, we use the growth estimator in our forecasts for reconstruction,⁹ henceforth referring to reconstructed modes as $\delta_r(\mathbf{K})$ instead of $\hat{\Delta}_G(\mathbf{K})$. However, as we discussed in Sec. II A, the output of the G estimator (or any other single estimator) will be contaminated by the other mode-couplings, with the

⁸For producing matter power spectra for forecasts, we relied on the nbodyskit code (<https://github.com/bccp/nbodykit>).

⁹Out of the G, S, and T estimators, the G estimator yields both the lowest noise and the highest signal to noise on reconstructed modes. However, some of the other estimators (e.g., $\alpha = \varphi\varphi$) also have signal to noise approaching that of the G estimator, since the contaminating terms in Eq. (17) act as “signal” in a signal to noise computation. This indicates that a more optimal choice of estimator weights may be possible, although we leave this to future work.

FIG. 1. Reconstruction noise power spectra for estimators that use each of the quadratic mode-couplings discussed in Sec. II B. We omit curves for the c_{01} and c_{02} estimators, which are greater than the upper limit of the plot. The G (“growth”) estimator has the lowest noise by far. These curves are computed for a DESI-like survey, but the hierarchy between them is unchanged for the other surveys we consider. The signal to noise on reconstructed modes (not shown) is likewise much higher for the G estimator than for S or T, justifying our use of the G estimator for our main results.

specific contamination given by Eq. (17), and we must incorporate this contamination into our forecasts.

We show this contamination in the left panel of Fig. 2, in the form of each term $c_\beta N_{GG}/N_{G\beta}$ in the square brackets of Eq. (17). These curves each represent separate multiplicative biases on the amplitude of a reconstructed mode. Those

arising from late-time gravitational evolution (S, T) or from advection of the primordial potential (c_{01}) are white in K . In contrast, those arising from couplings between δ and φ (c_{11}) or φ and itself ($\varphi\varphi$, c_{02}) scale like $M(K)^{-1} \propto K^{-2}$. We derive these scalings analytically in the large-scale limit in Appendix D. Importantly, all terms that scale like K^{-2} involve f_{NL} , such that, as for δ_g , low- K scale-dependent bias in the reconstructed modes can be used as a probe of local primordial non-Gaussianity. The right panel of Fig. 2 shows that the relative size of this scale-dependent bias is comparable for δ_g and δ_r , reaching $\mathcal{O}(10\%)$ at $K \sim 0.001 h \text{ Mpc}^{-1}$, assuming $f_{\text{NL}} = 1$.

Figure 2 also shows that the contamination from other mode-couplings is subdominant to the intrinsic bias on the reconstructed field (i.e., the c_G term in Eq. (17)). Thus, using $c_G = b_1 + (21/17)b_2$ from Table II, we can derive the rough dependence of P_π and P_{gr} on b_1 and b_2 :

$$P_\pi \propto b_1^2(b_1 + b_2)^2, \quad P_{\text{gr}} \propto b_1^2(b_1 + b_2). \quad (37)$$

If galaxy shot noise is negligible compared to P_{gg} , then the reconstruction noise N_{GG} satisfies $N_{\text{GG}} \propto b_1^4$, implying that $P_\pi/N_{\text{GG}} \propto (1 + b_2/b_1)^2$ in this regime. This scaling will be useful to help understand the behavior of our forecasts when we change the fiducial value of b_2 .

III. SIMULATIONS

To validate the quadratic estimator framework presented in Sec. II, we use a suite of 15 realizations of a

FIG. 2. Left: contamination in the expectation value of G estimator, corresponding to separate multiplicative biases on the amplitude of a reconstructed mode, computed for a DESI-like survey. Blue solid line is the estimator growth bias shown for comparison. Dashed lines indicate negative values. Several of these curves inherit the k^{-2} scaling of the scale-dependent bias in δ_g arising from nonzero f_{NL} , implying that reconstructed modes can be used to constrain f_{NL} in the same way. Right: ratio of scale-dependent bias from f_{NL} (for a fiducial value of $f_{\text{NL}} = 1$) to total bias for δ_g (solid) and δ_r (dot-dashed). Local primordial non-Gaussianity has roughly the same relative contribution to the bias of δ_g or reconstructed modes.

TABLE III. Properties of the halo mass bin employed in this study: the mean mass of the sample, the number density of halos $\bar{n}$, the linear bias b_1 , and the three relevant c_α parameters defined in Table II. The measured bias parameters are taken from [78], which is based on the same simulations and mass bin we use here.

Mass bin	Mean halo mass [$10^{13} h^{-1} M_\odot$]	$\bar{n}$ [$10^{-6} h^3 \text{ Mpc}^{-3}$]	b_1	c_g	c_t	c_s
I	0.77	627	1.07	0.62	1.14	1.07

cosmological N -body simulation. The initial conditions are generated with the second order Lagrangian perturbation theory (2-LPT) code [76] at the initial redshift $z_i = 99$ and are subsequently evolved using Gadget-2 [77]. The simulations are performed with $N_p = 1024^3$ dark matter particles in a cubic box of length $L = 1500 h^{-1} \text{ Mpc}$ with periodic boundary conditions. We assume a flat Λ CDM cosmology with the cosmological parameters $\Omega_m = 0.272$, $\Omega_\Lambda = 0.728$, $h = 0.704$, $n_s = 0.967$, $\sigma_8 = 0.81$.

We restrict our simulation comparisons to $z = 0$ because this is where nonlinearities are the strongest and our quadratic bias approximation breaks down at the lowest wave number; analyzing a field at $z = 0$ is therefore, arguably, the most rigorous test of our methodology. Since the purpose of the simulation comparisons is to provide a first, simple validation of our reconstruction procedure, rather than to fully explore its simulated performance as a function of redshift in detail, we defer any detailed explorations beyond our $z = 0$ validation to future work.

Dark matter halos in the final density field are identified using a friends-of-friends (FoF) algorithm with linking length $l = 0.2$ times the mean interparticle distance. The halos are binned in mass, with each bin spanning a factor of three in mass. We have checked the viability of our reconstruction method for a range of masses, finding qualitatively similar results in all cases; however, for simplicity, we present only the results for the lowest mass bin, the properties of which are given in Table III. Particles and halos are assigned to a regular grid using the cloud-in-cell (CIC) scheme. We Fourier transform the matter and halo density fields using the publicly available FFTW library.¹⁰

A. Generation of quadratic estimators

We generate quadratic estimators from the halo density field δ_g in N -body simulations using the convolution theorem. This means that we use a sequence of multiplications with powers of wave numbers in Fourier space, Fourier transforms, and subsequent multiplication of the weighted fields in configuration space. We generate three

quadratic estimators corresponding to the growth term δ^2 , shift term $\Psi \cdot \nabla \delta$, and the tidal term s^2 , with associated Fourier-space kernels given in Eq. (4). The first step in our procedure is to remove very small scale modes by applying a cut-off k_{max} in Fourier space through multiplication of the Fourier space density field with a filtering function. While the exact form of the cutoff is not important, we adopt a Gaussian filter $W(R\mathbf{k}) = \exp(-k^2 R^2/2)$ for numerical stability. We define the smoothed density field by $\delta_g^R(\mathbf{k})$. We choose three external smoothing scales: $R = 20 h^{-1} \text{ Mpc}$, $R = 10 h^{-1} \text{ Mpc}$, and $R = 4 h^{-1} \text{ Mpc}$, corresponding to maximum wave numbers $k_{\text{max}} \approx 0.05 h \text{ Mpc}^{-1}$, $k_{\text{max}} \approx 0.1 h \text{ Mpc}^{-1}$, and $k_{\text{max}} \approx 0.25 h \text{ Mpc}^{-1}$ respectively. The smoothing scale removes all wave numbers $k > k_{\text{max}}$, such that we reconstruct long wavelength modes using modes $k < k_{\text{max}}$ for three different cases.

The mode coupling functions $g_\alpha(\mathbf{q}, \mathbf{k} - \mathbf{q})$ defined by Eq. (14) contain a Wiener filter, which we implement by first generating the linear power spectrum on the simulation grid, and then defining two fields:

$$\delta_A(\mathbf{k}) = \frac{\delta_g^R(\mathbf{k})}{b_1^2 P_{\text{lin}}(\mathbf{k}) + \bar{n}^{-1}} \quad \text{and} \quad \delta_B(\mathbf{k}) = \frac{\delta_g^R(\mathbf{k}) P_{\text{lin}}(\mathbf{k})}{b_1^2 P_{\text{lin}}(\mathbf{k}) + \bar{n}^{-1}}, \quad (38)$$

where b_1 and $\bar{n}$ are the linear bias and halo number density corresponding to the halo mass bin defined in Table III. Using δ_A and δ_B we generate growth, shift and tidal estimators using multiplications of powers of wave numbers in Fourier space, Fourier transforms, and multiplication of fields in configuration space. For example, we generate the growth estimator as follows. First, we inverse Fourier transform both fields defined in Eq. (38) to obtain $\delta_A(\mathbf{x})$ and $\delta_B(\mathbf{x})$. Next, in configuration space, we multiply the product of both fields by 17/21 (Table II) and finally Fourier transform back to obtain the growth estimator in Fourier space. We generate shift and tidal estimators with a similar procedure.

Note that the main computational cost in generating the quadratic estimators comes from performing the Fourier transforms. The auto- and cross-spectrum analysis of quadratic estimators only requires the computational cost of a power spectrum analysis, which is quite efficient. In all our figures in this section, we estimate the error bars of our measurements using the standard deviation of 15 simulation realizations.

B. Cross-correlation of quadratic estimators with the initial linear field

In this section, we describe our results for the cross-correlations of three quadratic estimators $\hat{\Delta}_\alpha(\mathbf{k})$ with the initial linear field $\delta_1(\mathbf{k})$, and compare the theory predictions with simulations. The prediction is given by

¹⁰<http://www.fftw.org>

FIG. 3. Cross correlations of estimators $\hat{\Delta}_\alpha$ corresponding to the growth, shift, and tidal mode-couplings with the linear density field δ_1 . We compare theory predictions (lines) with simulations (points) for three different smoothing scales, $R = 20h^{-1}$ Mpc, $R = 10h^{-1}$ Mpc and $R = 4h^{-1}$ Mpc, corresponding to maximum wave numbers $k_{\max} = 0.05h$ Mpc $^{-1}$, $0.1h$ Mpc $^{-1}$, and $0.25h$ Mpc $^{-1}$ respectively. In this figure, we plot $\langle \hat{\Delta}_\alpha \delta_1 \rangle / N_{\alpha\alpha}$ (in contrast to what is defined in Eq. (16), in simulations we define the estimators $\hat{\Delta}_\alpha$ without a prefactor $N_{\alpha\alpha}$). We find very good agreement for the growth estimator for all smoothing scales, and also reasonably good agreement for the other estimators.

$$\begin{aligned}
 \langle \hat{\Delta}_\alpha(\mathbf{k}) \delta_1(\mathbf{k}') \rangle' &= b_1 N_{\alpha\alpha}(\mathbf{k}) \sum_{\beta \in \{G, S, T\}} c_\beta P_{\text{lin}}(\mathbf{k}) \\
 &\times \int_q \frac{f_\alpha(\mathbf{q}, \mathbf{k} - \mathbf{q}) f_\beta(\mathbf{q}, \mathbf{k} - \mathbf{q})}{2P_{\text{tot}}(\mathbf{q}) P_{\text{tot}}(\mathbf{k} - \mathbf{q})} \\
 &\times W(R\mathbf{q}) W(R(\mathbf{k} - \mathbf{q})) + P_{\alpha, \text{shot}}(\mathbf{k}) \\
 &= b_1 P_{\text{lin}}(\mathbf{k}) \sum_{\beta \in \{G, S, T\}} c_\beta \frac{N_{\alpha\alpha}(\mathbf{k})}{N_{\alpha\beta}(\mathbf{k})} + P_{\alpha, \text{shot}}(\mathbf{k}),
 \end{aligned} \tag{39}$$

where the prime on the left-hand side denotes that the factor of $(2\pi)^3 \delta_D(\mathbf{k} + \mathbf{k}')$ has been omitted, and c_β are bias parameters corresponding to the growth, shift and tidal terms and can be measured from either simulations or data. In our analysis we use the bias parameters from Table III, measured in simulations in [78]. In Eq. (39), $P_{\alpha, \text{shot}}$ is the bispectrum shot noise term. Since one field is the linear field, all contribution to this shot noise comes from the stochastic bias terms in the two galaxy fields δ_g in the quadratic estimator, such as ε and $\varepsilon_\delta \delta$ (see Appendix B or [72] for more discussion about stochastic bias terms). The expression for this shot noise contribution in this case can also be derived from Eq. (B20) and it takes the form

$$\begin{aligned}
 P_{\alpha, \text{shot}}(\mathbf{k}) &= \frac{b_1}{\bar{n}} P_{\text{lin}}(\mathbf{k}) N_{\alpha\alpha}(\mathbf{k}) \int_q \frac{f_\alpha(\mathbf{q}, \mathbf{k} - \mathbf{q})}{2P_{\text{tot}}(\mathbf{q}) P_{\text{tot}}(\mathbf{k} - \mathbf{q})} \\
 &\times W(R\mathbf{q}) W(R(\mathbf{k} - \mathbf{q})).
 \end{aligned} \tag{40}$$

In Fig. 3, we compare theory with simulations for three different values of $k_{\max}$. Although for the Fisher analysis in this work, we only use the growth estimator, here we also compare results in simulations for the shift and the tidal estimators. For the growth estimator, we find that the theory predictions agree very well with simulation results for up to

$k_{\max} = 0.25h$ Mpc $^{-1}$ at redshift $z = 0$. For the other estimators, we also find reasonably good agreement; however, upon close inspection we can see small disagreements which might arise from higher-order terms ignored in our theory predictions.

Interestingly, for $k_{\max} = 0.25h$ Mpc $^{-1}$, we can see in Fig. 3 that the shape of the cross-correlation of growth estimators with the density field is very similar to the linear power spectrum on large scales. The scale-dependent bias factor in Eq. (39) is flat on large scales, indicating that the reconstruction works very well for large $k_{\max}$.

C. Auto- and cross-correlations of quadratic estimators

In this section we discuss our results for the auto- and cross-correlations of three quadratic estimators from simulations and compare the results with our linear order theoretical prediction, given by

$$\begin{aligned}
 \langle \hat{\Delta}_\alpha(\mathbf{k}) \hat{\Delta}_\beta(\mathbf{k}') \rangle' &= b_1^4 N_{\alpha\alpha}(\mathbf{k}) N_{\beta\beta}(\mathbf{k}) \int_q \frac{f_\alpha(\mathbf{q}, \mathbf{k} - \mathbf{q}) f_\beta(\mathbf{q}, \mathbf{k} - \mathbf{q})}{[2P_{\text{tot}}(\mathbf{q}) P_{\text{tot}}(\mathbf{k} - \mathbf{q})]^2} \\
 &\times W(R\mathbf{q})^2 W(R(\mathbf{k} - \mathbf{q}))^2 P_{\text{lin}}(\mathbf{q}) P_{\text{lin}}(\mathbf{k} - \mathbf{q}) \\
 &+ b_1^2 P_{\text{lin}}(\mathbf{k}) \sum_{i,j} c_i c_j \frac{N_{\alpha\alpha}(\mathbf{k}) N_{\beta\beta}(\mathbf{k})}{N_{\alpha i}(\mathbf{k}) N_{\beta j}(\mathbf{k})} \\
 &+ P_{\alpha\beta, \text{shot}}(\mathbf{k}).
 \end{aligned} \tag{41}$$

The first term is of order $\mathcal{O}(\delta_1^4)$, while the second and third are of order $\mathcal{O}(\delta_1^6)$. The third term, $P_{\alpha\beta, \text{shot}}$, is the contribution arising from halo shot noise, and is given in Appendix B.

In Fig. 4 we compare cross-correlation results from simulations with theory, for the growth, shift, and tidal estimators, using the same three smoothing scales as above. The simulations and theory agree very well up to $k_{\max} = 0.1h$ Mpc $^{-1}$ at $z = 0$. For larger $k_{\max}$ we see good

FIG. 4. Autocorrelations of the quadratic estimators $\hat{\Delta}_a$, for the same smoothing scales shown in Fig. 3. The predictions for the growth estimator agree with simulations for all smoothing scales. However, for other estimators predictions agree with simulations for large smoothing scales but for the low smoothing scales, the predictions slightly disagree with simulation results as higher-order terms become more important.

agreement for the growth estimator and reasonable agreement for the tidal and shift estimators. The small disagreement of linear predictions for the tidal and shift estimators with simulations for the higher $k_{\max}$ show that higher-order

terms become important for these estimators. The detailed impact of these higher order corrections from biasing or scale dependent stochasticity will be subject of future inquiry. Although we appear to have excellent agreement

FIG. 5. Comparison of the auto power spectrum of the growth estimator $\hat{\Delta}_G$, normalized by N_{GG} computed from theory (in red), with $(\langle \hat{\Delta}_G \delta_1 \rangle)^2 / P_{\text{lin}}$ (in blue). We compare simulation results (points) with theory predictions (lines) for the same smoothing scales as Figs. 3 and 4. We again find excellent agreement between simulations and theory. In the bottom right panel, we plot the cross-correlation coefficient $r_{G\delta_1}$ between the growth estimator and the linear density field for three smoothing scales. We see that $r_{G\delta_1} > 0.9$ for $R = 4 h^{-1} \text{ Mpc}$ which is why in the bottom left panel, $\langle \hat{\Delta}_G \hat{\Delta}_G \rangle$ is signal dominated.

FIG. 6. 2D slices of the 3D linear density field (left panel) and the growth estimator $\hat{\Delta}_G$ (right panel). For the growth estimator we used $R = 4h^{-1}$ Mpc smoothing which corresponds to $k_{\max} = 0.25h$ Mpc $^{-1}$. We apply an external smoothing of $R = 20h^{-1}$ Mpc to both the linear and reconstructed fields. As expected, we find that the reconstruction reproduces many of the large-scale features in the linear density field.

for the growth term at higher $k_{\max}$, to be conservative, we still set the scale $k_{\max} = 0.1h$ Mpc $^{-1}$ at redshift $z = 0$ in our forecasts in Sec. IV. We scale this to other redshifts by making use of the fact that perturbation theory and the bias expansion at a given order will be valid at higher k for higher redshifts.

In Fig. 5, we plot the auto spectra of the growth estimator, normalized with N_{GG} (unlike in the previous plots), in order to compare them to an approximation of the signal power spectrum, given by the second term in Eq. (41) (the first and third terms represent noise). Since the contribution of the cross-shot noise is small, the signal part can be approximated by cross-correlating the growth estimator with the linear density field and dividing it by the linear power spectrum to ensure the correct normalization, i.e., $(\langle \hat{\Delta}_G \delta_1 \rangle)^2 / P_{\text{lin}}$; we show this in blue in Fig. 5. For the two larger smoothing scales, the spectra of the estimator are dominated by the noise contribution (which is white at low k). The excellent agreement between theory (red solid lines) and simulations (red points) for all smoothing scales serves as an additional verification that the reconstruction procedure is working as expected for reasonable values of $k_{\max}$. In addition to the auto spectra, to check how well the reconstruction is working, we plot the cross-correlation coefficients between the growth estimator and the linear density field in the bottom right panel of Fig. 5 for three different $k_{\max}$. The cross-correlation coefficient for low $k_{\max}$ is very low, $r_{G\delta_1} < 0.4$, explaining why the auto spectra in the top left panel are noise dominated. However, for highest $k_{\max}$ we consider, $0.25h$ Mpc $^{-1}$, the cross-correlation coefficient is $r_{G\delta_1} > 0.9$, which explains why the reconstruction works very well and the auto spectra for high $k_{\max}$ are signal dominated.

D. Visualization of reconstructed field

To visualize how well we are reconstructing the linear density field on large scales in simulations, we compare 2D slices of thickness $6h^{-1}$ Mpc of the linear density field and the reconstructed field in Fig. 6. We perform the reconstruction using $k_{\max} = 0.25h$ Mpc $^{-1}$, i.e., smoothing at a scale of $R = 4h^{-1}$ Mpc. In the visualization, we apply an external smoothing of $R = 20h^{-1}$ Mpc to both the linear field and the reconstructed field, which removes all modes with $k > 0.05h$ Mpc $^{-1}$. Our comparison of the linear and reconstructed fields in Fig. 6 shows that the reconstruction indeed recovers most of the large scale features in the linear density field. In Fig. 7 we show histograms, probing the one-point probability distribution functions, of the linear density field and the reconstructed field. We see that the reconstructed field is nearly Gaussian, partially justifying our approximation of a Gaussian likelihood in the next section.

IV. FORECASTS

A. Fisher matrix setup

To perform a Fisher forecast, we make the usual assumption that the measured tracer overdensity δ_g and the reconstructed field δ_r obey a Gaussian likelihood. For the matter and galaxy field, this approximation is partially justified by the fact that we are analyzing very large scales; for the reconstruction noise, this is partially justified by the fact that the reconstruction sums over a large number of mode pairs, so that to some extent the central limit theorem applies (although the pairs may not all be independent). Figure 7 supplies additional evidence that a simple Fisher forecast is sufficient, in that the PDFs of the density field

FIG. 7. Probability distribution functions (histograms) of the linear density field and the reconstructed field from the halo density field of mass bin I. As in Fig. 6, we use $k_{\max} = 0.25h \text{ Mpc}^{-1}$ for the reconstruction and apply an external smoothing scale of $R = 20h^{-1} \text{ Mpc}$ to both the linear field and the reconstructed field. The PDFs of the reconstructed field are scaled to have the same variance as the linear field, and shifted to have mean 0. We find that the PDF of the reconstructed field is very close to Gaussian. Note that here we have applied a low- k cutoff to the modes used for reconstruction of $k_{\min} = 0.05h \text{ Mpc}^{-1}$ in order to match the approach in our forecast section below.

(smoothed to correspond with our analysis range) do not greatly deviate from a Gaussian. This indicates that the influence of higher moments of the density field and noise is comparatively small for the purposes of a forecast.

Making this approximation and including the fact that δ_g has zero statistical mean, the Fisher matrix per mode $\mathbf{K}$ and redshift z is given by (e.g., [79])

$$\tilde{F}_{ab}(\mathbf{K}, z) = \frac{1}{2} \text{Tr}[\partial_a C(\mathbf{K}, z) C^{-1}(\mathbf{K}, z) \partial_b C(\mathbf{K}, z) C^{-1}(\mathbf{K}, z)], \quad (42)$$

where C is the total (signal plus noise) covariance matrix for our data vector $\mathbf{d}(\mathbf{K}) = (\delta_g(\mathbf{K}), \delta_r(\mathbf{K}))^T$, Tr is the trace matrix operator, $\partial_b C(\mathbf{K}, z) \equiv \frac{\partial}{\partial b} C(\mathbf{K}, z)$, and a, b are the parameters on which our quantities depend (in this case, f_{NL} and bias parameters). If the data vector is drawn from a Gaussian distribution and nothing is known about the parameters, then the inverse of the Fisher matrix gives the covariance matrix of the parameters, and the square root of the diagonal elements of F^{-1} give the error bars on the parameters and represent the minimum error achievable. Our goal is to calculate this minimum error, as it will determine our best ability to constrain parameters.

In reality, we do not just measure a single mode, but we measure several modes whose information can combined together in an integrated Fisher matrix for a specific redshift bin, i.e.,

$$F_{ab}(z) = \frac{V}{(2\pi)^2} \int_{K_{\min}}^{K_{\max}} dK \int_{-1}^1 d\mu K^2 \tilde{F}_{ab}(K, \mu, z). \quad (43)$$

Here V is the survey volume, $K_{\min}$ and $K_{\max}$ are the minimum and maximum moduli of the modes probed, and we already integrated over the azimuthal direction, supposing no dependence from it in the integrand.

For our specific case, the original field δ_g and the reconstructed field δ_r will give the total covariance matrix (which only depends on the magnitude of $\mathbf{K}$)

$$C(K, z) = \begin{bmatrix} C^{\text{gg}}(K, z) & C^{\text{gr}}(K, z) \\ C^{\text{gr}}(K, z) & C^{\text{rr}}(K, z) \end{bmatrix}, \quad (44)$$

with elements

$$C^{\text{gg}}(K, z) = \left(b_1(z) + \frac{c_{01} f_{\text{NL}}}{M(K, z)} D(z) \right)^2 P_{\text{lin}}(K, z) + P_{\text{gg,shot}}(K, z), \quad (45)$$

$$C^{\text{gr}}(K, z) = \left(b_1(z) + \frac{c_{01} f_{\text{NL}}}{M(K, z)} D(z) \right) b_r(K, z) P_{\text{lin}}(K, z) + P_{\text{gr,shot}}(K, z), \quad (46)$$

$$C^{\text{rr}}(K, z) = b_r(K, z)^2 P_{\text{lin}}(K, z) + N_{\text{GG}}(K, z) + P_{\text{rr,shot}}(K, z), \quad (47)$$

where

$$b_r \equiv b_1 \left(c_G + \sum_{\beta \neq G} c_\beta \frac{N_{\text{GG}}}{N_{G\beta}} \right), \quad (48)$$

and the sum runs over the mode-couplings found in Table II. We do not include redshift space distortions in these expressions; see Sec. IV C 5 for discussion.

The tracer shot noise is simply

$$P_{\text{gg,shot}}(K, z) = \frac{1}{\bar{n}(z)}, \quad (49)$$

where $\bar{n}$ is the comoving number density of observed tracers, while $P_{\text{rr,shot}}$ and $P_{\text{gr,shot}}$ are given in Eqs. (B15)–(B16) and (B20)–(B21), respectively. We will neglect the dependence of the reconstruction shot noise on f_{NL} . This is because in general these shot noise terms include the small scale tracer power spectrum, whose response to a change of f_{NL} is negligible compared to the response experienced by the large scale power spectrum. Moreover, even when the large scale tracer power spectrum enters the reconstruction shot noise, as in $P_{\text{rr,shot}}$ where there is a coupling between large and small scales as we explain in Appendix B, a small change from $f_{\text{NL}} = 0$, our fiducial value, is barely detectable. In principle, it may be possible extract additional information from the f_{NL} -dependence of the shot noise

contributions, but this will likely be difficult in practice, and therefore we conservatively choose not to consider these contributions as observables.

Substituting Eq. (44) into Eq. (42), we can derive an explicit formula for the Fisher matrix per mode for our case, which can then be inserted into Eq. (43):

$$\begin{aligned} \tilde{F}_{ab} = & \frac{1}{2} \left(\frac{1}{C^{\text{rr}} C^{\text{gg}} (1 - r_{cc}^2)} \right)^2 [C^{\text{gg}} \{ \partial_b C^{\text{gr}} (-C^{\text{gr}} \partial_a C^{\text{rr}} + C^{\text{rr}} \partial_a C^{\text{gr}}) + \partial_b C^{\text{rr}} (C^{\text{gg}} \partial_a C^{\text{rr}} - C^{\text{gr}} \partial_a C^{\text{gr}}) \} \\ & - C^{\text{gr}} \{ \partial_b C^{\text{gg}} (-C^{\text{gr}} \partial_a C^{\text{rr}} + C^{\text{rr}} \partial_a C^{\text{gr}}) + \partial_b C^{\text{rr}} (C^{\text{gg}} \partial_a C^{\text{rr}} - C^{\text{gr}} \partial_a C^{\text{gr}}) \} \\ & - C^{\text{gr}} \{ \partial_b C^{\text{gr}} (-C^{\text{gr}} \partial_a C^{\text{gr}} + C^{\text{rr}} \partial_a C^{\text{gg}}) + \partial_b C^{\text{rr}} (C^{\text{gg}} \partial_a C^{\text{gr}} - C^{\text{gr}} \partial_a C^{\text{gg}}) \} \\ & + C^{\text{rr}} \{ \partial_b C^{\text{gg}} (-C^{\text{gr}} \partial_a C^{\text{gr}} + C^{\text{rr}} \partial_b C^{\text{gg}}) + \partial_b C^{\text{gr}} (C^{\text{gg}} \partial_a C^{\text{gr}} - C^{\text{gr}} \partial_a C^{\text{gg}}) \}], \end{aligned} \quad (50)$$

where r_{cc} is the g-r cross correlation coefficient:

$$r_{cc} \equiv \frac{C^{\text{gr}}}{\sqrt{C^{\text{gg}} C^{\text{rr}}}}. \quad (51)$$

For $a = b$, we obtain

$$\tilde{F}_{aa} = \frac{1}{2(1 - r_{cc}^2)^2} \left[\left(\frac{\partial_a C^{\text{gg}}}{C^{\text{gg}}} - 2r_{cc}^2 \frac{\partial_a C^{\text{gr}}}{C^{\text{gr}}} \right)^2 + 2r_{cc}^2 (1 - r_{cc}^2) \left(\frac{\partial_a C^{\text{gr}}}{C^{\text{gr}}} \right)^2 + 2r_{cc}^2 \frac{\partial_a C^{\text{rr}}}{C^{\text{rr}}} \left(\frac{\partial_a C^{\text{gg}}}{C^{\text{gg}}} - 2 \frac{\partial_a C^{\text{gr}}}{C^{\text{gr}}} \right) + \left(\frac{\partial_a C^{\text{rr}}}{C^{\text{rr}}} \right)^2 \right]. \quad (52)$$

On the other hand, if we only use δ_g , we get

$$\tilde{F}_{aa}^{(\text{g only})} = \frac{1}{2} \left(\frac{\partial_a C^{\text{gg}}}{C^{\text{gg}}} \right)^2. \quad (53)$$

B. Analytical derivation of cosmic variance cancellation

Cosmic variance cancellation will occur in the limit of low noise on the measured fields—that is, low reconstruction noise on the quadratic estimator, and low galaxy shot noise. To investigate this case analytically, let us work in the limit of very low shot noise, so that

$$\begin{aligned} C^{\text{gg}}(K) &= b_g(K)^2 P_{\text{lin}}(K), \\ C^{\text{gr}}(K) &= b_g(K) b_r(K) P_{\text{lin}}(K), \\ C^{\text{rr}}(K) &= b_r(K)^2 P_{\text{lin}}(K) + N_{\text{GG}}(K). \end{aligned} \quad (54)$$

Further, let us assume that f_{NL} is the only unknown parameter. If we define

$$\begin{aligned} x(K) &\equiv \frac{N_{\text{GG}}(K)}{b_r(K)^2 P_{\text{lin}}(K)}, \\ R_p(K) &\equiv \left(\frac{\partial_{f_{\text{NL}}} b_r(K)}{b_r(K)} \right) \left(\frac{\partial_{f_{\text{NL}}} b_g(K)}{b_g(K)} \right)^{-1}, \end{aligned} \quad (55)$$

where $x(K)$ is the inverse of the signal to noise per mode of the reconstructed field and $R_p(K)$ is a measure of similarity between the response of the bias of the reconstructed field and the one of the original tracer field, then a short

calculation gives the unmarginalized error bar on f_{NL} per K -mode:

$$\begin{aligned} \sigma_{f_{\text{NL}}}^2(K) &= \sigma_{f_{\text{NL}}, \text{g only}}^2(K) \frac{2x(K)}{(R_p(K) - 1)^2} \frac{1}{1 + 2(R_p(K) - 1)^{-2} x(K)}, \end{aligned} \quad (56)$$

where $\sigma_{f_{\text{NL}}, \text{g only}} = [F_{aa}^{(\text{g only})}]^{-1/2}$.

Let us investigate the general behavior of this equation in some limiting cases. If $(R_p - 1)^{-2} x$ is small, the $R_p < 0$ case (when $\partial_{f_{\text{NL}}} b_r$ and b_r have opposite signs) will result in smaller error bars than the $R_p > 0$ case, because the signatures of f_{NL} in b_r and b_g will be more distinguishable in that case. Expanding Eq. (56) in limit of small $(R_p - 1)^{-2} x$ gives

$$\begin{aligned} \sigma_{f_{\text{NL}}}^2(K) &= \sigma_{f_{\text{NL}}, \text{g only}}^2(K) \frac{2x(K)}{(R_p(K) - 1)^2} \\ &\times \sum_{n=0}^{\infty} [-2(R_p(K) - 1)^{-2} x(K)]^n. \end{aligned} \quad (57)$$

As we will see in Appendix D, $N_{\text{GG}} \propto k_{\text{max}}^{-3}$ in the low- K limit, so that we arrive at

$$\lim_{x \rightarrow 0} \sigma_{f_{\text{NL}}}^2(K) \propto 2\sigma_{f_{\text{NL}}, \text{g only}}^2(K) [k_{\text{max}}^{-3} + \mathcal{O}(k_{\text{max}}^{-6})], \quad (58)$$

TABLE IV. Survey characteristics used for our main forecasts. The DESI-like survey is based on the expected DESI emission-line galaxy sample, the MegaMapper-like survey is a next-generation survey targeting high-redshift “dropout” galaxies, and the PUMA-like survey represents a future 21 cm intensity mapping effort over half the sky. We marginalize over b_1 , b_2 , and b_{s^2} in our forecasts, and determine b_{11}^E and b_{02}^E using the relationships in Sec. II B. For the PUMA-like forecast, the main $\bar{n}$ values represent effective number densities that reproduce the same noise level as the sum of shot and instrumental noise power at $k = k_{\max}$, while the expected physical number densities are shown in parentheses. For this forecast, we also consider the effects of the so-called “foreground wedge” that will prevent direct measurement of certain modes. See main text for details.

	DESI-like	MegaMapper-like		PUMA-like	
	$0.6 < z < 1.6$	$2 < z < 2.5$	$4.5 < z < 5$	$2 < z < 3$	$5 < z < 6$
<i>Survey parameters</i>					
Survey volume (Gpc ³)	100	80	66	266	203
Mean galaxy density $\bar{n}$ (Mpc ⁻³)	10^{-4}	6×10^{-4}	2×10^{-5}	2×10^{-3} (6×10^{-3})	1×10^{-3} (2×10^{-2})
$K_{\max}$ for f_{NL} constraint ($h \text{ Mpc}^{-1}$)	0.05	0.08	0.14	0.09	0.15
$k_{\max}$ for reconstruction ($h \text{ Mpc}^{-1}$)	0.15	0.24	0.4	0.26	0.47
<i>Fiducial bias parameters</i>					
b_1	1.6	2.9	7.0	2.1	3.7
b_2	-0.30	1.1	17	0.041	2.8
b_{s^2}	-0.17	-0.54	-1.7	-0.31	-0.77
b_{11}^E	-3.0	-2.5	37	-3.5	0.58
b_{02}^E	-14	-21	85	-19	-16

where we assume that $R_p - 1$ varies slowly with K . This demonstrates that constraints on f_{NL} that use both reconstructed modes and modes of the original tracer will improve on a tracer-only analysis in a way that is only limited by the noise on the reconstructed modes (if shot noise is negligible).

Cosmic variance cancellation clearly requires that δ_r and δ_g are measured at the same wave number and in the same volume. To verify this, we can repeat the derivation above with $C^{\text{gr}} = 0$, corresponding to δ_r and δ_g being measured in different volumes. In this case, Eq. (56) becomes

$$\sigma_{f_{\text{NL}}}^2(K) = \sigma_{f_{\text{NL}}, \text{g-only}}^2(K) \frac{[1 + x(K)]^2}{R_p(K)^2 + [1 + x(K)]^2}, \quad (59)$$

which approaches a finite limit as $x \rightarrow 0$; thus, the improvement realized in Eq. (58) is only possible if δ_r and δ_g can be compared mode-by-mode in the same volume.

C. Assumptions and experimental configurations

1. Scales

In each forecast, for measuring f_{NL} , we use δ_g modes and reconstructed modes with wave number K satisfying $K_{\min} < K < K_{\max}$, and we also use reconstructed modes with $K_f < K < K_{\min}$, where $K_f \approx 0.002h \text{ Mpc}^{-1}$ is the lowest measurable wave number within each survey volume. In this way, $K_{\min}$ accounts for possible systematic effects that can prevent direct measurements of δ_g on large scales, but that do not impede reconstruction of these large-scale modes using smaller-scale correlations; an example is foreground contamination for intensity mapping

experiments, which has been a primary motivator for other work on reconstruction methods [2,4–8]. As input to the density-field reconstruction, we use modes with wave number k satisfying $K_{\max} < k < k_{\max}$. We consider a range of possible $K_{\min}$ values in our forecasts, while $k_{\max}$ and $K_{\max}$ are fixed for each survey, as described below.

2. Surveys

In our main forecasts, we consider three galaxy surveys, with properties summarized in Table IV. The first is similar to the emission-line galaxy sample expected from DESI [80]. For this survey, following [56], we consider 14000 deg² of sky area over $0.6 < z < 1.6$, which translates into a total comoving volume of roughly 100 Gpc³ and a mean redshift of $\bar{z} \approx 1$. We use a mean galaxy number density of $\bar{n} = 10^{-4} \text{ Mpc}^{-3}$, obtained by dividing the expected total number of redshifts in the DESI ELG sample (1.7×10^7 , from [80]) by the survey volume, and assume a mean linear galaxy bias of $b_1 = 1.6$. We take $K_{\max} = 0.05h \text{ Mpc}^{-1}$, since linear bias is expected to be an acceptable approximation for $K < K_{\max}$ at $z = 1$, and $k_{\max} = 0.15h \text{ Mpc}^{-1}$, since our quadratic bias expansion is valid for $k < k_{\max}$ at $z = 1$ (see Sec. III for justification based on simulations).

The second survey, which we call “MegaMapper-like,” is modeled on proposals for a next-generation spectroscopic survey targeting high-redshift “dropout” galaxies in the southern hemisphere [81–83]. For this, we assume a 14000 deg² survey, and separately consider two redshift bins, at $2 < z < 2.5$ and $4.5 < z < 5$, which have volumes of 80 Gpc³ and 66 Gpc³ respectively. The mean number density and linear bias in each bin are obtained from

averages of the values at the bin edges, taken from Table 1 of [82]; this yields $\bar{n} = 6 \times 10^{-4} \text{ Mpc}^{-3}$ and $b_1 = 2.9$ for the lower-redshift bin, and $\bar{n} = 2 \times 10^{-5} \text{ Mpc}^{-3}$ and $b_1 = 7.0$ for the higher-redshift bin. For K_{max} and k_{max} , we scale the DESI values using the ratio of linear growth factors between the mean redshifts of each redshift bin, to account for the increased range of validity of our perturbative expressions at higher redshift.¹¹

The third survey is based on specifications for PUMA, an envisioned radio interferometer designed for 21 cm intensity mapping [85,86]. We assume a survey over half the sky, and again consider two redshift bins, this time at $2 < z < 3$ and $5 < z < 6$, with volumes 266 Gpc^3 and 203 Gpc^3 respectively. For simplicity, we treat this survey as observing galaxy positions directly, rather than brightness temperature (which is just a rescaled biased tracer of the matter density). To do so, we set the noise contribution to the tracer power spectrum P_{gg} to equal the sum of the shot noise and instrumental noise power spectra computed using the PUMA noise calculator,¹² evaluated at $k = k_{\text{max}}$ in each redshift bin. We quote an effective number density that would result in the same noise level in Table IV. When computing the shot noise contributions to P_{gr} and P_{tr} , we use the expected number densities of 21 cm emitters, also taken from the PUMA noise calculator and quoted in parentheses in Table IV. For the linear bias in each bin, we use values from Fig. 33 of [85], evaluated at the mean redshifts. As for the MegaMapper-like survey, we scale K_{max} and k_{max} from DESI by the appropriate ratios of linear growth factors.

One could also consider applying our reconstruction procedure to data from the SPHEREx satellite [87], which aims to map large-scale structure by using low-resolution spectroscopy to determine the redshifts of over 500 million galaxies. The total SPHEREx galaxy sample will have higher number density than DESI ELGs by at least a factor of 5, but many of these galaxies will have significant redshift uncertainties that will suppress the measured clustering power in a scale-dependent way, and this would need to be included in our formalism. On the other hand, angular variations in the mean redshift uncertainty may also imprint spurious clustering signals on the data, and long-mode reconstruction is a promising avenue to probe or

mitigate this. We intend to explore these issues in future work.

In our derivation of stochastic contributions to the noise of the estimator and the cross-correlation between estimator and galaxy fields in App. B, we assume that the noise is Poissonian, i.e., that $\langle \epsilon \epsilon \rangle = (2\pi)^3 \delta_{\text{D}}(\mathbf{K} + \mathbf{K}')/\bar{n}$. There is evidence for halo stochasticity being sub-Poissonian for high mass haloes and super-Poissonian for low mass haloes [88,89]. Since the stochasticity corrections arise from small-scale exclusion and higher order biases, the actual shot noise levels cannot be theoretically predicted, implying that it may be advisable to marginalize over the stochasticity parameter(s). This approach is indeed adopted by some for the f_{NL} forecasting literature (e.g., [90]) but certainly not all of it (e.g., [53,56]). Here we decide to fix the stochasticity parameters to their fiducial Poissonian values and defer a more detailed investigation of the impact of noise corrections on the reconstructed fields to future work. We do note however, that we expect the impact of shot noise marginalization to be rather small, since we do not include the additional non-Gaussian signal arising in combination with stochastic terms in Eqs. (45)–(47).

3. 21 cm foregrounds

An additional consideration for 21 cm intensity mapping is the presence of foreground radiation, dominantly synchrotron from our own galaxy, that is brighter than the cosmological signal by several orders of magnitude. These foregrounds are extremely smooth in frequency, which implies that they mainly populate Fourier modes with low line-of-sight wave number $k_{\parallel}$; these modes will therefore likely not be usable for cosmology. Furthermore, the chromatic properties of interferometers generically spread foreground power from the low- $k_{\parallel}$ modes into a wedge-shaped region in the $k_{\parallel} - k_{\perp}$ plane (e.g., [14–16]), although this contamination can be removed with sufficiently precise instrumental calibration (e.g., [91,92]).

For constraining f_{NL} , the wedge will have two effects: it will reduce the number of short-wavelength modes available for the quadratic estimator, therefore increasing the noise N_{GG} on the reconstructed modes, and it will also reduce the number of long-wavelength δ_{g} modes available for measuring the scale-dependent bias induced by primordial non-Gaussianity.

We account for both effects in our forecasts for the PUMA-like survey, assuming a foreground wedge defined by 3 times the primary beam width, following [85]; see Appendix E 1 for details of how this is implemented in our computations. In addition, we perform forecasts that ignore the wedge, to represent the case when it can be completely eliminated via calibration. We account for lost low- $K_{\parallel}$ modes in two ways: either by restricting δ_{g} to have $K > K_{\parallel, \text{min}}$, or by approximating $K_{\parallel, \text{min}}$ as an isotropic K_{min} , matching our procedure for DESI and MegaMapper.

¹¹The argument for this k_{max} scaling is as follows. Nonlinear displacements are one of the physical effects that limit the accuracy of perturbation theory, and it is natural to assume that the redshift-dependence of k_{max} will scale with the rms displacement, which we compute in the Zel'dovich approximation. Under these assumptions, $k_{\text{max}}(z) \propto D^{-1}(z)$. In reality, the scaling of k_{max} with redshift is more complicated, involving the power spectrum tilt at the relevant wave numbers (e.g., [84]), but the simple growth factor scaling we use here should at least be roughly indicative of the useful scales for our forecasts.

¹²<https://github.com/slosar/PUMANoise>

The former approach is more realistic, while the latter is easier to compare with the other surveys, so we present the latter in the main text, and the former in Appendix E 2.

4. Bias parameters

For every survey, to perform forecasts, we assume a fiducial value of the quadratic bias parameter b_2 derived from the fitting formula of [93], which was fit to halo bias in separate-universe simulations over the range $1 \lesssim b_1 \lesssim 10$:

$$b_2(b_1) = \frac{1}{2}(0.412 - 2.143b_1 + 0.929b_1^2 + 0.008b_1^3), \quad (60)$$

where the extra factor of 2 arises from our different definition of b_2 compared to [93]. The fiducial value of the tidal bias b_{s^2} is found from

$$b_{s^2} = -\frac{2}{7}(b_1 - 1), \quad (61)$$

which assumes that the tidal bias in Lagrangian space is zero. In our forecasts, b_1 , b_2 , and b_{s^2} are allowed to vary independently (i.e., are marginalized over when we estimate uncertainties on f_{NL}), while b_{11} and b_{02} are assumed to obey the relationships in Eqs. (30)–(31) and (33)–(35). We take wide, flat priors on b_1 , b_2 , and b_{s^2} ; we have also implemented 10% Gaussian priors on b_2 and b_{s^2} , but these have a negligible effect on our baseline results.

5. Redshift space distortions

The line-of-sight component of a galaxy's position is observationally inferred from the galaxy's redshift, and the associated “redshift-space distortions” of δ_g should be included in a full treatment of the observed galaxy clustering. The leading-order effect is to add a $f\mu^2$ term to the linear bias of δ_g , such that Eq. (36) is modified to

$$\delta_g(\mathbf{k}) = \left[b_1 + f_{\text{NL}} \frac{c_{01}}{M(k)} + f\mu^2 \right] \delta_1(\mathbf{k}) + \dots, \quad (62)$$

where $f \equiv d \log D / d \log a$, $\mu \equiv k_{\parallel} / k$, and D is the linear growth factor [94]. Higher-order effects will create additional mode-couplings that can be described in perturbation theory (e.g., [95,96]). In a real tracer catalogue, there will also be line-of-sight-dependent selection effects that can be treated perturbatively [97].

We do not include any of these effects in our baseline forecasts, leaving them for future work. However, as a first step in this direction, we have checked the impact of including the Kaiser term. This raises the reconstruction noise N_{GG} by increasing $P_{\text{gg,tot}}$ in the denominator of Eq. (16), while also increasing the amplitude of P_{gg} and P_{gr} , thereby increasing the signal to noise on those quantities. For all surveys we consider, the former effect

overcomes the latter, with the result that $\sigma(f_{\text{NL}})$ increases by roughly 10%, and the improvement in $\sigma(f_{\text{NL}})$ from including reconstructed modes decreases by no more than the same amount. Additional mode-couplings from nonlinear redshift-space effects will likely dominate over this change, and a detailed analysis will be worthwhile to pursue, especially since some of these mode-couplings could potentially carry additional information about f_{NL} [90].

D. Expected precision on reconstructed modes

Aside from primordial local non-Gaussianity, there are many other applications of reconstructing large-scale modes, including more general constraints on cosmology, tests of predictions for the power spectrum on the largest scales, calibration of photometric redshifts [7], cross-correlations with other tracers (such as kSZ fluctuations in the CMB, e.g., [6]), and removing contamination from measurements of lensing of 21 cm fluctuations [2]. To represent the general utility of reconstructed modes from different surveys, in Fig. 8 we show the expected precision on the auto power spectrum of the reconstructed modes (plotted using the fiducial bias parameters from Table IV), computed in wave number bins with $\Delta K = 0.002h \text{ Mpc}^{-1}$. While these error bars are substantial for $K \lesssim 0.01h \text{ Mpc}^{-1}$ in DESI and the high- z bin of MegaMapper, the precision is expected to be much better for MegaMapper at low z and across the entire redshift range of PUMA, with most error bars approaching the cosmic variance limit. This will enhance many scientific applications of these surveys, particularly for PUMA, where large-scale modes can be reconstructed at high precision even in the presence of the foreground wedge.

Reference [85] also estimates the total signal to noise in reconstructed modes from PUMA over $1 < z < 6$, following the methodology of [2], finding $\mathcal{O}(1300)$ in the no-wedge case and $\mathcal{O}(500)$ for the same wedge model we use here. For comparison, we find a total S/N of 135 (108) for $2 < z < 3$ and 161 (134) for $5 < z < 6$ in the no-wedge (wedge) case. A direct comparison between the two sets of forecasts is difficult, because they use several distinct approximations: [85] treats the 21 cm brightness temperature as a linearly biased tracer of the matter density, while we have incorporated second-order biasing; [85] neglects the shot noise contribution to the reconstructed mode power spectrum, while we include it; [85] bias-harden their results against mode-couplings from gravitational lensing, while we do not; and, most importantly, [85] only consider reconstruction of modes that are purely transverse to the line of sight ($k_{\parallel} = 0$), while we use a 3d reconstruction formalism. Nevertheless, both forecasts reach the same broad conclusion that PUMA will be able to reconstruct long-wavelength density modes with total signal to noise of several hundred, which is strong motivation for continued

FIG. 8. Expected error bars on the reconstructed power spectrum P_π for the surveys and redshift bins we consider (blue), along with cosmic-variance-limited error bars (orange), computed for bandpowers with $\Delta K = 0.002 h \text{ Mpc}^{-1}$. Downward arrows indicate error bars whose lower limits fall outside of the y axis range. High-precision measurements of the power spectrum of reconstructed modes will be possible in several cases, even in the presence of a 21 cm foreground wedge for PUMA.

studies of the density reconstruction method we have presented in this paper.

E. Results: Constraints on non-Gaussianity

1. DESI

Figure 9 shows the results of our forecasts for the DESI-like survey. The left panel shows the various power spectra of interest, of linear matter density, galaxy number density, and reconstructed matter density modes, along with the cross spectrum between galaxies and reconstructed modes. This panel also shows the shot noise on P_{gg} , P_{gr} , and P_π , as well as the statistical noise (N_{GG}) on reconstructed modes. For DESI, the galaxy power spectrum is well above the shot

noise, while the reconstructed power spectrum is about an order of magnitude lower than the reconstruction noise. Despite the fact that galaxy shot noise is below P_{gg} , the shot noise contributions for both P_{gr} , and P_π are above the signal power spectra. As explained in Appendix B, this is due to coupling between galaxy shot noise and clustering at large scales, where the variance is larger than at small scales and therefore these shot noise spectra are significantly boosted compared to the $\bar{n}^{-1}$ contribution.

The middle panel of Fig. 9 shows the expected constraints on f_{NL} when only δ_g is used, or when reconstructed modes are also incorporated. The right panel shows the ratio of $\sigma(f_{\text{NL}})$ in these two cases. The improvement in $\sigma(f_{\text{NL}})$ is negligible at the lowest K_{min} we consider, which

FIG. 9. Forecasts for a DESI-like survey. Left: signal and noise power spectra involved in the forecast. The galaxy auto spectrum is well above the shot noise, while the auto spectrum of reconstructed modes (P_π) is roughly an order of magnitude below both the reconstruction noise (N_{GG}) in the quadratic estimator and the shot noise contribution to the estimator variance. Center: expected constraints on f_{NL} when only δ_g is used (solid), or when δ_r is also used (dotted). We assume that $\delta_g(\mathbf{K})$ cannot be directly measured for $K < K_{\text{min}}$, and marginalize over the b_1 , b_2 , and b_{s^2} bias parameters. Right: ratio of $\delta_g + \delta_r$ and δ_g -only cases from the center panel. We only notice an improvement for higher values of K_{min} , corresponding to using δ_r but not δ_g at $K < K_{\text{min}}$.

FIG. 10. The analog of the right panel of Fig. 9, with a variety of (mostly artificial) modifications to the forecasts. There is no improvement in $\sigma(f_{\text{NL}})$ when δ_r is neglected at $K < K_{\text{min}}$, indicating that the inclusion of δ_r at $K < K_{\text{min}}$ drives the improvement. Greater improvements are achieved for higher galaxy number density or if k_{max} can be increased by a factor of 2, with milder changes if the fiducial b_2 value is set to zero or shot noise on P_{rr} and P_{gr} is neglected.

corresponds to δ_g being measured on all scales resolvable within the survey volume ($K_{\text{min}} = K_f$). However, a larger improvement is seen when K_{min} is assumed to be higher: for $K_{\text{min}} = 0.02h \text{ Mpc}^{-1}$, for example, $\sigma(f_{\text{NL}})$ improves by around 15% when reconstructed modes are used.

To determine the origin of this behavior, we show several modifications of this forecast in Fig. 10. In particular, when reconstructed modes with $K < K_{\text{min}}$ are not included, there is no improvement of $\sigma(f_{\text{NL}})$, indicating that these modes are entirely responsible for the improvement. Therefore, DESI is not powerful enough to allow for cosmic variance cancellation between δ_g and δ_r at the same scales; rather, the primary use of reconstruction is to access scales ($K < K_{\text{min}}$) where δ_g cannot be directly measured. This naturally explains why the improvement of $\sigma(f_{\text{NL}})$ grows for higher K_{min} . While the absolute values of $\sigma(f_{\text{NL}})$ are not impressive at such high K_{min} —at $K_{\text{min}} = 0.02h \text{ Mpc}^{-1}$, for example, $\sigma(f_{\text{NL}}) \approx 50$ without reconstruction and 40 with reconstruction—the improvement comes “for free,” without requiring any other datasets.

The other curves in Fig. 10 illuminate other aspects of this forecast. Increasing $\bar{n}$ to an unrealistically high value of 10^2 Mpc^{-3} improves the δ_g -only forecast by roughly 10% (not shown), and also increases the improvement on $\sigma(f_{\text{NL}})$ from including δ_r , indicating that shot noise is a limiting factor in this improvement. Simply neglecting $P_{\text{rr,shot}}$ and $P_{\text{gr,shot}}$ has a similar effect, clarifying that shot noise in the galaxy power spectrum itself is comparatively less important than in these other spectra.

Also, we see the same type of change if we alter the fiducial value of b_2 . As mentioned at the end of Sec. II C, if $P_{\text{gg}} \gg P_{\text{gg,shot}}$ (as it is here), then $P_{\text{rr}}/N_{\text{GG}} \propto (1 + b_2/b_1)^2$,

so increasing b_2 from -0.3 to 0 boosts the signal to noise on the reconstructed modes. This would lead to a larger improvement if not for the large contribution of $P_{\text{rr,shot}}$. Boosting k_{max} by a factor of 2 leads to a better $\sigma(f_{\text{NL}})$ improvement at low K_{min} . This change lowers the Gaussian reconstruction noise N_{GG} , but also raises $P_{\text{rr,shot}}$ and $P_{\text{gr,shot}}$ by different amounts, and the combination of these changes ends up slightly boosting the constraining power of δ_r .

It may seem counterintuitive that δ_r adds anything at all to our forecasts, since the reconstruction noise and shot noise on P_{rr} are much larger than P_{rr} itself: one would expect such large noise to lead to a low cross-correlation coefficient between δ_r and δ_g , and also make it difficult to extract information from the auto spectrum of δ_r . However, the presence of a cross shot noise contribution to P_{gr} alters this picture, contributing to the $\delta_r - \delta_g$ cross-correlation coefficient and altering the structure of the covariance matrix. While it is not trivial to see in the Fisher matrix expression in Eq. (52), the net effect is to enhance the information content of δ_r with respect to f_{NL} . [54] reached a similar conclusion when examining cosmic variance cancellation between different line intensity maps, noticing that lowering the cross shot noise contribution led to worsened constraints on f_{NL} .

2. MegaMapper

Our results for the MegaMapper-like survey are shown in Fig. 11. For the low- z bin, the signal to reconstruction noise on the reconstructed modes is higher than for DESI, thanks to a combination of higher $\bar{n}$, higher k_{max} , and higher bias, and the signal to shot noise ratio is also correspondingly smaller. This leads to a greater improvement in $\sigma(f_{\text{NL}})$ when reconstructed modes are included. The left panel of Fig. 12 shows that, like DESI, this improvement comes not from cosmic variance cancellation, but from reconstructed modes with $K < K_{\text{min}}$, where we assume that δ_g cannot be directly measured. We see large changes if $\bar{n}$ is boosted or $P_{\text{gr,shot}}$ and $P_{\text{rr,shot}}$ are neglected, indicating that shot noise is a limiting factor in this bin. Changing the fiducial b_2 from 1.1 to 0 reduces the usefulness of the reconstructed modes for the same reason that changing b_2 increased their usefulness for DESI.

We see rather different behavior in the high- z bin. There, we find that the reconstruction noise is of the same order as P_{rr} while the shot noise contribution is much greater than the signal, and the shot noise contribution to P_{gr} is also greater than the signal. Despite this, the improvement in $\sigma(f_{\text{NL}})$ is larger than for the low- z bin, reaching 50% at $K_{\text{min}} = K_f$. The right panel of Fig. 12 shows that the improvement is the same whether or not we include modes of δ_r with $K < K_{\text{min}}$, and therefore, cosmic variance cancellation between δ_g and δ_r is solely responsible for the change in $\sigma(f_{\text{NL}})$.

FIG. 11. As Fig. 9, for low-redshift (top panels) and high-redshift (bottom panels) bins of a MegaMapper-like survey. The former has greater signal to noise on reconstructed modes than DESI, leading to a greater improvement in $\sigma(f_{\text{NL}})$ when these modes are included in the forecast. For the latter, the shot noise contributions to P_{rr} and P_{gr} are comparatively much larger, leading to a different scale-dependence for the improvement in $\sigma(f_{\text{NL}})$.

We also see from Fig. 12 that the low number density ($\bar{n} = 2 \times 10^{-5} \text{ Mpc}^{-3}$) in the high- z bin is not a huge limiting factor, with only a modest change if we use a much larger number density. This is because the reconstruction noise remains comparable to P_{rr} even for a much denser

survey, while further improvements are possible for a higher k_{max} but the same number density. If $P_{\text{rr,shot}}$ and $P_{\text{gr,shot}}$ are ignored, the results revert to the same situation as the low- z bin, with only slight gains in $\sigma(f_{\text{NL}})$ possible for low K_{min} . Finally, if b_2 is changed from 17 to 0, there is

FIG. 12. Modifications to the base MegaMapper forecasts. For the low-redshift bin (left panel), as for DESI, the improvement in $\sigma(f_{\text{NL}})$ is driven mostly by modes of δ_r with $K < K_{\text{min}}$, with further improvement possible for higher galaxy number density. For the high-redshift bin (right panel), neglecting δ_r at $K < K_{\text{min}}$ makes no difference, indicating that for lower K_{min} values, cosmic variance cancellation between δ_g and δ_r at the same K is driving the improvement in $\sigma(f_{\text{NL}})$. There are several ways to obtain greater improvements, as discussed in the main text.

FIG. 13. As Fig. 9, for low-redshift (top panels) and high-redshift (bottom panels) bins of a PUMA-like survey, treating the 21 cm brightness temperature in the same way as δ_g in our other forecasts, and translating thermal noise on the brightness temperature into an effective tracer number density for computing shot noise. We show $\sigma(f_{NL})$ either neglecting or incorporating the effects of the 21 cm foreground wedge; at high z , the benefit to $\sigma(f_{NL})$ from including reconstructed modes is greater in the presence of the wedge, since there are fewer δ_g modes that can be directly measured in that case. The results for the low-redshift bin are similar to those for MegaMapper, while larger improvements in $\sigma(f_{NL})$ are possible at higher redshift.

significantly more improvement in $\sigma(f_{NL})$: the amplitudes of P_{rr} and P_{gr} are reduced, but the relative uncertainty on f_{NL} from marginalizing over b_2 is also reduced, and the latter effect wins.

3. PUMA

We show results for the PUMA-like survey in Fig. 13, either neglecting or including the effects of the foreground wedge. Note that the left panels in Fig. 13 only show noise curves corresponding to the no-wedge case. As for the other surveys, we assume an isotropic K_{min} for δ_g in Fig. 13; we show results for a cutoff on $K_{||}$, which are qualitatively similar to those in Fig. 13, in Appendix E 2.

For both redshift bins, the shot noise in C^{gg} , C^{rr} , and C^{gr} is below the signal. However, the reconstruction noise is high enough in the low-redshift bin that the effect of reconstructed modes on $\sigma(f_{NL})$ is similar to DESI and the low- z MegaMapper bin, with the vast majority of the extra constraining power coming from reconstructed modes with $K < K_{min}$ (see the left panel of Fig. 14). The impacts of taking a higher k_{max} or tracer number density (equivalent to thermal noise in the interferometer) would only be mild.

Meanwhile, in the high- z bin, the improvement in $\sigma(f_{NL})$ arises from a combination of low- K reconstructed

modes and cosmic variance cancellation between δ_g and δ_r . There is greater improvement in the presence of the wedge, as reconstruction helps to recover modes that would otherwise be lost. This improvement is around 20% at the lowest K_{min} , and increases as more δ_g modes are lost, implying that reconstruction will be extremely useful for single-tracer constraints on f_{NL} from PUMA or other high- z intensity mapping. The right panel of Fig. 14 shows that lower thermal noise would lead to further improvements, while a lower value of b_2 would worsen the results due to a lowering of the signal to noise on P_{rr} .

4. ...and beyond

To demonstrate how the constraints on f_{NL} scale for surveys with extremely low shot noise and reconstruction noise, we also examine forecasts for the PUMA high-redshift bin where k_{max} is artificially increased, assuming that our quadratic bias model is valid to arbitrarily high k .¹³

¹³In practice the quadratic bias model will break down at sufficiently high k , but a theoretical framework such as the response function formalism (e.g., [61,62]) may allow the use of higher k_{max} , with suitable modifications of the reconstruction procedure. We leave this topic to future work.

FIG. 14. Modifications to the base PUMA forecasts, neglecting the foreground wedge. In the low- z bin, modes of δ_r with $K < K_{\min}$ are entirely responsible for the improvement in $\sigma(f_{\text{NL}})$, with most other modifications having little effect. In the high- z bin, the blue dotted curve demonstrates that the $\sigma(f_{\text{NL}})$ improvement comes from a combination of low- K modes of δ_r and cosmic variance cancellation at higher K . The improvement would get better if the thermal noise could be reduced (which maps onto a higher $\bar{n}$ in these forecasts).

We take the galaxy number density to infinity in these forecasts, to prevent shot noise from becoming the limiting factor. In this case, we expect the uncertainty on f_{NL} to scale like the inverse of the signal to noise on the reconstructed modes (see Sec. IV B). In turn, in this limit, the signal to noise scales like $k_{\text{max}}^{3/2}$ because the reconstruction noise spectrum N_{GG} becomes proportional to the number of modes with $k_{\min} < k < k_{\text{max}}$ (see Eq. (D13)).

In Fig. 15, we show the ratio of $\sigma(f_{\text{NL}})$ from $g + r$ or g -only forecasts as a function of k_{max} for two representative

values of $K_{\min}$. We indeed find that as the signal to noise on reconstructed modes is increased, the improvement on $\sigma(f_{\text{NL}})$ also increases, with the unmarginalized forecasts quickly satisfying the expected scaling. (Marginalization over bias parameters causes small deviations from this scaling.) This demonstrates the huge increases in constraining power that are possible in principle for a survey with high galaxy number density and many small-scale modes whose correlations can be used in reconstruction. We have also numerically verified that the $\sigma(f_{\text{NL}})$ ratio stays flat with increasing k_{max} if the noise in P_{gg} is taken

FIG. 15. The ratio of $\sigma(f_{\text{NL}})$ for the $\delta_g + \delta_r$ and δ_g -only forecasts for the high- z PUMA bin, where $\bar{n}$ is taken to infinity and k_{max} is artificially increased, assuming that our quadratic bias model is valid to arbitrarily high k . The left and right panels correspond to two representative values of $K_{\min}$. We show forecasts after marginalizing over bias parameters (black points) and without any marginalization (red points), along with the expected $k_{\text{max}}^{-3/2}$ scaling (dashed lines, each normalized to the highest- k_{max} point plotted). We find the unmarginalized curves quickly approach the ideal scaling, while the marginalized forecasts show small deviations from this scaling. This shows that large increases in constraining power are possible in principle for surveys with very high number density and a large allowed value of k_{max} .

very high, or if zero cross-correlation between δ_g and δ_r is assumed, further demonstrating that the scaling seen in Fig. 15 arises from the joint constraining power between δ_g and δ_r measured from the same volume.

V. DISCUSSION

The results presented here can be compared to other methods either utilizing reconstruction and/or combining a tower of n -point correlation functions. Compared to most methods proposed in the literature, this work presents an optimal quadratic estimator to reconstruct the large scale mode. As explained in Sec. IV D, in principle this reconstructed mode can be used for several (cosmological) applications, and here we only explored f_{NL} as an application of interest. When comparing this work with previous works, the main question is if the amount of information captured in the statistics of the tracer field is fully exploited. While it will be hard to compare methods directly, here we propose some heuristic arguments where we think our methods overlap and where they differ.

As mentioned in the introduction, some publications have aimed to simplify the search for primordial non-Gaussianities by proposing more compressed versions of the full bispectrum [34–41]. Common to these works is the fact that the information accessed is captured by the 2- and 3-point functions. In this work, besides the 3-point function, the 4-point function is also used and is important in obtaining cosmic variance cancellation. In other words, as shown in Fig. 15, significant improvements are possible when some conditions are met that would not be possible when considering the compressed statistics proposed in these earlier works.

Even if cosmic variance cancellation is not achieved, we generally observe improvements between 20%–50%. These numbers are similar to those projected in for example [37] for compressed statistics, but direct comparisons between our method and others are generally difficult. In the method presented in this paper, the improvement can roughly be attributed to coming from access to larger scales through the reconstruction, or, when both the linear and reconstructed mode are combined, cosmic variance cancellation. The projected improvement on the amplitude f_{NL} from compressed statistics is the result of adding the bispectrum information on top of the power spectrum. For a detailed comparison we would need to carefully associate every mode with improved signal-to-noise side by side for the two different methods. Although this would be interesting by itself, as it would help us understand to what extent these methods are overlapping or how they complement one another, we will leave this to future work.

The paper which our work has most in common with is [40], which discusses the information content of a joint analysis of the two point function and squeezed three- and

four-point functions. This work has several commonalities with our analysis. To perform forecasts, [40] uses the squeezed-limit position-dependent power spectrum as a field, in an approach that is quite similar to our long wavelength mode reconstruction. The author also makes similar arguments for how sample variance cancellation can significantly influence and improve constraints.

However, there are also many important differences to our approach. Most importantly, the specific squeezed-limit power spectrum picture in [40] is discussed as a tool to enable better forecasting of joint 2, 3 and 4-point analyses of local non-Gaussianity, rather than a practical data analysis method. In contrast, our method has been proposed as an analysis method and estimator to rapidly jointly analyze 2, 3 and 4- point functions, that is not only computationally tractable, but has been tested (to some extent) on simulations.

There are also significant differences in the details of the methodology. Our reconstruction quadratic estimator can infer the long wavelength mode from mode-pairs that are not much smaller than the mode to be reconstructed; in contrast, [40] always operates in the squeezed limit when analyzing the position-dependent power spectrum. While it is expected that the majority of information about local non-Gaussianity in the 3 and 4-point functions is contained in very squeezed shapes, it is not clear that nonsqueezed shapes do not contribute to long-wavelength mode reconstruction and hence sample variance cancellation. On the other hand, we note that in our analysis method we combine all quadratic estimator mode pairs into one long-wavelength mode estimate; in contrast, [40] shows that additional sample variance cancellation can be obtained when treating each mode pair (or position-dependent power spectrum bin) as a separate tracer. Although this suggests that further improvements to our method might be possible, the results of [40] suggest that this would only give significant improvements for very high k and very low noise, beyond the capabilities of next-generation surveys.

Finally, shortly before the completion of this work, in a follow-up to [9,10] presented results relating to reconstruction of large-scale density modes using biased tracers, although without discussing the application to constraining non-Gaussianity. While the core of this work is similar (using a quadratic estimator as proposed by [2]), here we explicitly account for the mode-coupling from higher-order biasing (which is non-negligible) in our estimator and compare theoretical estimates of the reconstruction noise, including bi- and trispectrum shot noise with additional contributions from primordial non-Gaussianity, to simulations. [10] include observations on the light-cone in their formalism, and also include the effect of redshift space distortions up to second order in the linear density, which we neglect in this work (although see Sec. IV C 5 for a discussion of the impact of the Kaiser term).

VI. CONCLUSIONS

In this paper, we have further developed a method for reconstructing modes of the cosmic density field using a quadratic estimator. This estimator extracts information about (typically) large-scale modes from correlations between smaller-scale modes, similar to standard methods for CMB lensing reconstruction. We have improved upon the estimator introduced in [2] by incorporating nonlinear biasing and local-type primordial non-Gaussianity, up to second order in the linear density field. At this order, there are several distinct sources of couplings between small-scale modes of the tracer density field, with amplitudes (i.e., bias coefficients) that are unknown *a priori*. We have found that an estimator based on the mode-coupling due to isotropic growth of the perturbations results in the lowest noise on reconstructed modes, and have enumerated the various multiplicative biases that will accompany the output of this estimator. We have also applied this estimator (along with those based on large-scale bulk flows and tidal interactions) to halos in N -body simulations, verifying that the results agree with analytical predictions.

In the course of this study, we have identified that it is crucial to include the shot noise contribution to the covariance between directly-observed tracer modes and reconstructed modes when performing an analysis. The shot noise not only adds a white noise contribution to the tracer power spectrum itself (in the case that the tracers Poisson-sample the density field, which we assume here), but also adds noise to the reconstructed-mode power spectrum and the cross-spectrum with the tracer modes. For sufficiently low tracer number density, this contribution can actually overwhelm the reconstruction noise from the quadratic estimator, and the cross spectrum alters the correlation coefficient between the tracer and reconstructed modes. We self-consistently include these features in our forecasts.

We have carried out forecasts that apply this formalism to several upcoming large-scale structure surveys: the emission-line galaxy survey from DESI [80], the high- z dropout survey envisioned in the MegaMapper proposal [82,83], and the 21 cm line-intensity survey from the PUMA proposal [85,86], treated like a galaxy survey with effective number density derived from PUMA's thermal noise model. Examining the expected error bars on the power spectrum of the reconstructed modes for $K < 0.02h \text{ Mpc}^{-1}$, we find that these error bars are several times larger than the signal for DESI and a high-redshift bin of MegaMapper. The latter is limited by the low number density of tracers, leading to a high shot noise contribution to the reconstruction noise, while the former's high reconstruction noise is sourced both by shot noise and a low number of modes used in the reconstruction (i.e., low k_{max}). In the other forecasts, we find that high-S/N reconstructions of the large-scale density power spectrum can be obtained, with the caveat that this spectrum comes

with multiplicative biases with known shapes but unknown amplitudes.

We have also computed the expected improvement in constraints on the amplitude of local-type primordial non-Gaussianity, f_{NL} , arising from analyzing reconstructed modes along with directly-observed tracer modes. For DESI and low- z bins of MegaMapper and PUMA, the improvement arises solely from being able to access reconstructed modes with $K < K_{\text{min}}$, where K_{min} denotes the minimum wave number at which we assume tracer modes can be measured (for systematics obscuring tracer modes with $K < K_{\text{min}}$ but not affecting the modes used for reconstruction). On the other hand, for a high- z bin of MegaMapper, the improvement in f_{NL} constraints arises solely from cosmic variance cancellation between tracer and reconstructed modes at the same wave numbers, similar to what can happen with different tracer populations or tracer-lensing cross-correlations [50,51,53,54]. For a high- z bin of PUMA, the $\sigma(f_{\text{NL}})$ improvement comes from a combination of cosmic variance cancellation and reconstructed modes alone. Generally, cosmic variance cancellation depends on having a sufficiently high cross-correlation coefficient between the tracer and reconstructed modes, but this depends on shot noise in a somewhat complicated way, due to the aforementioned cross shot noise contribution.

The improvement in $\sigma(f_{\text{NL}})$ also depends on the assumed value of K_{min} , so we have plotted the expected constraints as a function of K_{min} . In general, reconstructed modes improve $\sigma(f_{\text{NL}})$ by tens of percents: for example, at $K_{\text{min}} = 0.01h \text{ Mpc}^{-1}$, $\sigma(f_{\text{NL}})$ improves by a few percent for DESI, 15% and 40% for the low- z and high- z MegaMapper bins we consider, and at least 20% for both z -bins of PUMA, depending on what is assumed for the 21 cm foreground wedge. We have also shown that in the limit of zero shot noise, and if our quadratic bias model were valid to arbitrarily high k , $\sigma(f_{\text{NL}})$ scales like $k_{\text{max}}^{-3/2}$, reflective of the number of small-scale modes used for reconstruction.

There are several possible ways that this work could be extended. For example, we have neglected redshift-space distortions, but they should clearly be incorporated in advance of applying this technique to data. One could also consider applying reconstruction to photometric surveys, after an assessment of the impact of photometric redshift errors on the results. It would be interesting to see how things change if one were to consider the bias model from [98], based on shifted versions of bias operators designed to more fully incorporate large-scale displacements. Given that shot noise can be a limiting factor, the application of weighting schemes (such as [99]) could potentially further lower the shot noise level and improve the reconstruction performance, although these schemes can be difficult to implement in practice; along similar lines, one could apply a different weighting to better extract

the scale dependent signal given by f_{NL} (see, e.g., [40]). Finally, one could consider investigating nonlinear response functions [61,62] as a way to increase the number of small-scale modes that could be used in the quadratic estimator. Overall, we expect there to be many applications for reconstructed modes beyond constraints on local-type non-Gaussianity, and we therefore advocate for this reconstruction procedure as a useful tool to increase the scientific returns of upcoming large-scale structure surveys.

ACKNOWLEDGMENTS

We thank Emanuele Castorina, Azadeh Moradinezhad Dizgah, Simone Ferraro, Mat Madhavacheril, Moritz Münchmeyer, Will Percival, Matias Zaldarriaga, and Hong-Ming Zhu for useful conversations. We also thank Emanuele Castorina, Azadeh Moradinezhad Dizgah, Emmanuel Schaan and Marcel Schmittfull for thoughtful comments on a draft of this paper. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Colleges and Universities. S. F. thanks the Kavli Institute for Cosmology, Cambridge for hospitality while part of this work was carried out. The numerical part of this work was performed using the DiRAC COSMOS supercomputer and greatly benefited from the support of K. Korneet. M. A. acknowledges support from the Cambridge Commonwealth Trust, the Higher Education Commission, Pakistan, and the Cambridge Centre for Theoretical Cosmology. T. B. acknowledges support from the Cambridge Center for Theoretical Cosmology through the Stephen Hawking Advanced Fellowship. P. D. M. acknowledges support of the Netherlands organization for scientific research (NWO) VIDI grant (dossier 639.042.730). B. D. S. acknowledges support from an Isaac Newton Trust Early Career Grant, from a European Research Council (ERC) Starting Grant under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 851274), and from an STFC Ernest Rutherford Fellowship. O. D. is funded by the STFC CDT in Data Intensive Science.

APPENDIX A: DERIVATION OF DENSITY RECONSTRUCTION FROM THE BISPECTRUM

We will here consider how we can reconstruct an unknown field X given the knowledge of (i) its bispectrum B with two other fields Y, Z , and (ii) a measurement of these two other fields. While we write our argument exploiting the connection between bispectra and quadratic estimators in a form that is generally valid, for this paper we will assume that $X = \delta_1$, the linear density field, and both Y, Z are the observed nonlinear density field $Y, Z = \delta_g$; given that we can calculate the $\delta_1 \delta_g \delta_g$ bispectrum, we can

easily write an estimator for δ_1 given an observed δ_g . We will assume statistical homogeneity and isotropy of the fields and the bispectrum, which is a good approximation for large-scale structure surveys (although it may be broken for other applications).

We begin by defining the bispectrum of the unknown field with two observed fields:

$$\begin{aligned} \langle X(\mathbf{k}_1)Y(\mathbf{k}_2)Z(\mathbf{k}_3) \rangle \\ = (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B_{XYZ}(k_1, k_2, k_3). \end{aligned} \quad (\text{A1})$$

We will now write an ansatz for recovering the unknown field X from a quadratic estimator involving Y, Z :

$$\hat{X}(\mathbf{K}) = \int_{\mathbf{q}} g(\mathbf{q}, \mathbf{K} - \mathbf{q}) Y(\mathbf{q}) Z(\mathbf{K} - \mathbf{q}), \quad (\text{A2})$$

where we have introduced a function g which weights these pairs of modes. As we will see, the arguments of Y, Z assumed here are a consequence of the delta function momentum constraint in the bispectrum.

We now derive this function g . The function must obviously give an unbiased estimator. In a situation where one may not wish to average over Y, Z at fixed X (e.g., because they are the same fields), we will define unbiasedness by the condition that

$$\langle X(\mathbf{K}') \hat{X}(\mathbf{K}) \rangle = (2\pi)^3 \delta_D(\mathbf{K}' + \mathbf{K}) P_{XX}(\mathbf{K}). \quad (\text{A3})$$

For the estimator, this implies that since

$$\begin{aligned} \langle X(\mathbf{K}') \hat{X}(\mathbf{K}) \rangle &= \int_{\mathbf{q}} g(\mathbf{q}, \mathbf{K} - \mathbf{q}) \langle X(\mathbf{K}') Y(\mathbf{q}) Z(\mathbf{K} - \mathbf{q}) \rangle \\ &= (2\pi)^3 \delta_D(\mathbf{K}' + \mathbf{K}) \\ &\quad \times \int_{\mathbf{q}} g(\mathbf{q}, \mathbf{K} - \mathbf{q}) B_{XYZ}(\mathbf{K}, \mathbf{q}, |\mathbf{K} - \mathbf{q}|), \end{aligned} \quad (\text{A4})$$

we have a normalization condition on g

$$I[g] \equiv \int_{\mathbf{q}} g(\mathbf{q}, \mathbf{K} - \mathbf{q}) \frac{B_{XYZ}(\mathbf{K}, \mathbf{q}, |\mathbf{K} - \mathbf{q}|)}{P_{XX}(\mathbf{K})} = 1. \quad (\text{A5})$$

We would also like the estimator $\hat{X}$ to have as little variance per mode as possible. We will assume that, for the purposes of variance calculation, the fields can be approximated as Gaussian and statistically isotropic. Under these assumptions, the variance $V[f](\mathbf{K})$ is given by:

$$\begin{aligned}
\langle \hat{X}(\mathbf{K}) \hat{X}(\mathbf{K}') \rangle &= (2\pi)^3 V[g](\mathbf{K}) \delta_{\mathbf{D}}(\mathbf{K} + \mathbf{K}') \\
&= \int_{\mathbf{q}, \mathbf{q}'} g(\mathbf{q}, \mathbf{K} - \mathbf{q}) g(\mathbf{q}', \mathbf{K}' - \mathbf{q}') \langle Y(\mathbf{q}) Z(\mathbf{K} - \mathbf{q}) Y(\mathbf{q}') Z(\mathbf{K}' - \mathbf{q}') \rangle \\
&= \int_{\mathbf{q}, \mathbf{q}'} g(\mathbf{q}, \mathbf{K} - \mathbf{q}) g(\mathbf{q}', \mathbf{K}' - \mathbf{q}') [\langle Y(\mathbf{q}) Y(\mathbf{q}') \rangle \langle Z(\mathbf{K} - \mathbf{q}) Z(\mathbf{K}' - \mathbf{q}') \rangle + \langle Y(\mathbf{q}) Z(\mathbf{K}' - \mathbf{q}') \rangle \langle Z(\mathbf{K} - \mathbf{q}) Y(\mathbf{q}') \rangle] \\
&= \int_{\mathbf{q}, \mathbf{q}'} g(\mathbf{q}, \mathbf{K} - \mathbf{q}) g(\mathbf{q}', \mathbf{K}' - \mathbf{q}') [(2\pi)^3 P_{YY}(q) \delta_{\mathbf{D}}(\mathbf{q} + \mathbf{q}') (2\pi)^3 P_{ZZ}(|\mathbf{K} - \mathbf{q}|) \delta_{\mathbf{D}}(\mathbf{K} + \mathbf{K}' - \mathbf{q} - \mathbf{q}') \quad (A6)
\end{aligned}$$

$$+ (2\pi)^3 P_{YZ}(q) \delta_{\mathbf{D}}(\mathbf{q} + \mathbf{K}' - \mathbf{q}') (2\pi)^3 P_{YZ}(|\mathbf{K} - \mathbf{q}|) \delta_{\mathbf{D}}(\mathbf{K} - \mathbf{q} + \mathbf{q}')] \quad (A7)$$

$$\begin{aligned}
&= (2\pi)^3 \delta_{\mathbf{D}}(\mathbf{K} + \mathbf{K}') \int_{\mathbf{q}} [g(\mathbf{q}, \mathbf{K} - \mathbf{q}) (-\mathbf{q}, -\mathbf{K} + \mathbf{q}) P_{YY}(q) P_{ZZ}(|\mathbf{K} - \mathbf{q}|) \\
&\quad + g(\mathbf{q}, \mathbf{K} - \mathbf{q}) g(\mathbf{q} - \mathbf{K}, -\mathbf{q}) P_{YZ}(q) P_{YZ}(|\mathbf{K} - \mathbf{q}|)], \quad (A8)
\end{aligned}$$

where we have used Wick's theorem. In the following we will specialize to the case of $Y = Z$ (which is the relevant case for our application) and can thus consider g to be symmetric under exchange of its arguments. With the requirement for the reconstructed field to be real, this implies the following expression for the variance as a functional of g :

$$V[g](\mathbf{K}) = 2 \int_{\mathbf{q}} g^2(\mathbf{q}, \mathbf{K} - \mathbf{q}) P_{YY}(q) P_{YY}(|\mathbf{K} - \mathbf{q}|). \quad (A9)$$

We can thus solve for g by minimizing the variance $V[g](\mathbf{K})$ subject to the constraint $I[g] = 1$. We can

do this by introducing a Lagrange multiplier $N_{\alpha\alpha}$ and minimizing

$$V[g] - N_{\alpha\alpha} \times I[g], \quad (A10)$$

with respect to g . Minimizing this expression, we obtain

$$\begin{aligned}
g(\mathbf{q}, \mathbf{K} - \mathbf{q}) &= N_{\alpha\alpha}(K) \frac{1}{4P_{YY}(q)P_{YY}(|\mathbf{K} - \mathbf{q}|)} \frac{B_{XY}(K, q, |\mathbf{K} - \mathbf{q}|)}{P_{XX}(K)}, \quad (A11)
\end{aligned}$$

where from the constraint equation $I = 1$ we find

$$N_{\alpha\alpha}(K) = \left[\int_{\mathbf{q}} \frac{1}{4P_{YY}(q)P_{YY}(|\mathbf{K} - \mathbf{q}|)} \left(\frac{B_{XY}(K, q, |\mathbf{K} - \mathbf{q}|)}{P_{XX}(K)} \right)^2 \right]^{-1}. \quad (A12)$$

Applying this to our choice of fields, i.e., evaluating the $\langle \delta_1 \delta_g \delta_g \rangle$ bispectrum, we note that we recover a function g which gives the same expression for the quadratic estimator as used in the main part of our paper. With Eq. (14) we have for the bispectrum of a linear mode $X = \delta_1$ and two galaxy modes (ignoring biases) $Y = Z = \delta_g$

$$B_{\delta_1, \delta_g, \delta_g}(k_1, k_2, k_3) = 2[F_{\alpha}(\mathbf{k}_1, \mathbf{k}_3) P_{\text{lin}}(k_1) P_{\text{lin}}(k_3) + F_{\alpha}(\mathbf{k}_1, \mathbf{k}_2) P_{\text{lin}}(k_1) P_{\text{lin}}(k_2)]. \quad (A13)$$

Plugging this into Eq. (A12) we get

$$g(\mathbf{q}, \mathbf{K} - \mathbf{q}) = N_{\alpha\alpha}(K) \frac{F_{\alpha}(\mathbf{K}, -\mathbf{q}) P_{\text{lin}}(q) + F_{\alpha}(\mathbf{K}, -\mathbf{K} + \mathbf{q}) P_{\text{lin}}(|\mathbf{K} - \mathbf{q}|)}{2P_{\text{tot}}(q) P_{\text{tot}}(|\mathbf{K} - \mathbf{q}|)}, \quad (A14)$$

which agrees with Eq. (14) in the main text.

APPENDIX B: NOISE EXPRESSIONS FOR QUADRATIC ESTIMATOR

In this Appendix, we derive expressions for the noise power spectrum corresponding to the auto-correlation of the reconstructed field $\hat{\Delta}_\alpha(\mathbf{K})$ and its cross-correlation with the original input tracer field $\delta_g(\mathbf{K})$. We will see that the noise comes from a combination of shot noise, due to discrete sampling of the underlying matter field, and cosmic variance.

1. Noise for the autocorrelation of the reconstructed field

In this Appendix, we calculate the covariance of our quadratic estimators, which is defined as

$$\begin{aligned} & \langle \hat{\Delta}_\alpha(\mathbf{K}) \hat{\Delta}_\beta(\mathbf{K}') \rangle - \langle \hat{\Delta}_\alpha(\mathbf{K}) \rangle \langle \hat{\Delta}_\beta(\mathbf{K}') \rangle \\ &= \int_{\mathbf{q}} \int_{\mathbf{q}'} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) g_\beta(\mathbf{q}', \mathbf{K}' - \mathbf{q}') \\ & \quad \times (\langle \delta_g(\mathbf{q}) \delta_g(\mathbf{K} - \mathbf{q}) \delta_g(\mathbf{q}') \delta_g(\mathbf{K}' - \mathbf{q}') \rangle \\ & \quad - \langle \delta_g(\mathbf{q}) \delta_g(\mathbf{K} - \mathbf{q}) \rangle \langle \delta_g(\mathbf{q}') \delta_g(\mathbf{K}' - \mathbf{q}') \rangle). \end{aligned} \quad (\text{B1})$$

To compute this expression we have to first derive the four point function for the input tracer field. It is possible to derive the shot noise formulae directly in Fourier space, with a discretized version of the tracer field. Alternative derivations of bispectrum shot noise can be found in [100] following [101]. We cross-checked our results with [33,102].

Let us start by rederiving the stochasticity contributions to the power spectrum, bispectrum and trispectrum. Let us consider a finite number N of pointlike tracers, such as galaxies, at positions $\mathbf{x}_i$ in a finite volume V .¹⁴ Their Fourier space density field is then given as a sum of plane waves

$$\delta_g(\mathbf{k}) = \frac{1}{\bar{n}} \sum_i \exp[i\mathbf{k}\mathbf{x}_i], \quad (\text{B2})$$

where $\bar{n} = N/V$. The power spectrum of the discrete tracers in the finite volume can then be computed as

$$\begin{aligned} P_g(k) &= \frac{1}{V} \langle \delta_g(\mathbf{k}) \delta_g(-\mathbf{k}) \rangle \\ &= \frac{V}{N^2} \sum_{i=j} \langle \exp[i\mathbf{k}(\mathbf{x}_i - \mathbf{x}_j)] \rangle + \frac{V}{N^2} \sum_{i \neq j} \langle \exp[i\mathbf{k}(\mathbf{x}_i - \mathbf{x}_j)] \rangle \\ &= \frac{1}{\bar{n}} + P_{g,\text{cont}}(k). \end{aligned} \quad (\text{B3})$$

¹⁴It is useful to recall that all wave vectors in a finite volume are integer multiples of the fundamental wave number. The Dirac delta distribution thus becomes a Kronecker delta $(2\pi)^2 \delta_D(\mathbf{k}_1 + \mathbf{k}_2) \rightarrow V \delta_K(\mathbf{k}_1, \mathbf{k}_2)$.

Here, the constant $1/\bar{n}$ is denoted the shot noise term and we have identified the nonzero lag expectation value with the continuous part of the discrete tracer power spectrum $P_{g,\text{cont}}(k)$. In the local bias model at linear order we have $P_{g,\text{cont}}(k) = b_1^2 P_{\text{lin}}(k)$, which becomes $P_{g,\text{cont}}(k) = [b_{10} + b_{01}/M(k)]^2 P_{\text{lin}}(k)$ in the presence of primordial non-Gaussianities of the local kind. Let us now consider the bispectrum. Following the same steps that led to the power spectrum above, we have to consider the case where all three positions coincide, the case where two positions coincide but are different from the third, and finally the case where all three positions are distinct:

$$\begin{aligned} B_g(\mathbf{k}_1, \mathbf{k}_2) &= \frac{1}{V} \langle \delta_g(\mathbf{k}_1) \delta_g(\mathbf{k}_2) \delta_g(-\mathbf{k}_1 - \mathbf{k}_2) \rangle \\ &= \frac{V^2}{N^3} \sum_{i=j=l} \langle \exp[i\mathbf{k}_1(\mathbf{x}_i - \mathbf{x}_l) + i\mathbf{k}_2(\mathbf{x}_j - \mathbf{x}_l)] \rangle \\ & \quad + \frac{V^2}{N^3} \sum_{i=l \neq j} \langle \exp[i\mathbf{k}_1(\mathbf{x}_i - \mathbf{x}_l) + i\mathbf{k}_2(\mathbf{x}_j - \mathbf{x}_l)] \rangle \\ & \quad + 2\text{perm.} \\ & \quad + \frac{V^2}{N^3} \sum_{i \neq j, j \neq l, i \neq l} \langle \exp[i\mathbf{k}_1(\mathbf{x}_i - \mathbf{x}_l) + i\mathbf{k}_2(\mathbf{x}_j - \mathbf{x}_l)] \rangle \\ &= \frac{1}{\bar{n}^2} + \frac{1}{\bar{n}} [P_{g,\text{cont}}(k_1) + 2\text{perm.}] + B_{g,\text{cont}}(\mathbf{k}_1, \mathbf{k}_2). \end{aligned} \quad (\text{B4})$$

Again, the nonzero lag correlators are identified with the continuous power spectrum and bispectrum of the tracer field. We see that two different stochasticity corrections arise: a $1/\bar{n}^2$ constant shot noise term and a product of the shot noise and the continuous power spectrum. As above for the power spectrum, in the presence of primordial non-Gaussianity, both of these continuous statistics contain the respective non-Gaussian bias corrections. Note that there is now a coupling between stochasticity and clustering which is enhanced with respect to the pure noise term on large scales.

Let us connect this result to the noise terms introduced in Eq. (20), where the relevant contributions are given by

$$\delta_g \supset \epsilon + \epsilon_\delta \delta + \epsilon_\varphi \varphi. \quad (\text{B5})$$

The three-point correlator of the noise fields ϵ can be associated with the white-noise term in Eq. (B4)

$$\langle \epsilon(\mathbf{k}_1) \epsilon(\mathbf{k}_2) \epsilon(\mathbf{k}_3) \rangle = (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) \frac{1}{\bar{n}^2}. \quad (\text{B6})$$

The contributions from $\epsilon_\delta \delta$ and $\epsilon_\varphi \varphi$ arise by correlating with the linear δ and φ fields and the linear noise term ϵ

$$\begin{aligned} \langle \epsilon(\mathbf{k}_1)([\epsilon_\delta \star \delta] + [\epsilon_\varphi \star \varphi])(\mathbf{k}_2)(b_{10}\delta + b_{01}\varphi)(\mathbf{k}_3) \rangle = \int_{\mathbf{q}} [\langle \epsilon(\mathbf{k}_1)\epsilon_\delta(\mathbf{q}) \rangle (b_{10}\langle \delta(\mathbf{k}_2 - \mathbf{q})\delta(\mathbf{k}_3) \rangle + b_{01}\langle \varphi(\mathbf{k}_2 - \mathbf{q})\delta(\mathbf{k}_3) \rangle) \\ + \langle \epsilon(\mathbf{k}_1)\epsilon_\varphi(\mathbf{q}) \rangle (b_{10}\langle \varphi(\mathbf{k}_2 - \mathbf{q})\delta(\mathbf{k}_3) \rangle + b_{01}\langle \varphi(\mathbf{k}_2 - \mathbf{q})\varphi(\mathbf{k}_3) \rangle)]. \end{aligned} \quad (\text{B7})$$

Using [72,103] we have for the noise correlators

$$\langle \epsilon\epsilon_\delta \rangle = (2\pi)^3 \delta_D(\mathbf{k} + \mathbf{k}') \frac{b_{10}}{\bar{n}}, \quad \langle \epsilon\epsilon_\varphi \rangle = \frac{b_{01}}{b_{10}} \langle \epsilon\epsilon_\delta \rangle = (2\pi)^3 \delta_D(\mathbf{k} + \mathbf{k}') \frac{b_{01}}{\bar{n}}. \quad (\text{B8})$$

We finally obtain for the mixed contribution to the three-point correlator

$$\langle \epsilon(\epsilon_\delta \delta + \epsilon_\varphi \varphi)(b_{10}\delta + b_{01}\varphi) \rangle = (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) \frac{P_{\text{lin}}(k_3)}{\bar{n}} \left[b_{10} \left(b_{10} + \frac{b_{01}}{M(k_3)} \right) + \frac{b_{01}}{M(k_3)} \left(b_{10} + \frac{b_{01}}{M(k_3)} \right) \right]. \quad (\text{B9})$$

In summary, we have

$$B_g(k_1, k_2, k_3) = \frac{1}{\bar{n}^2} + \frac{1}{\bar{n}} \left[\left(b_{10} + \frac{b_{01}}{M(k_1)} \right)^2 P_{\text{lin}}(k_1) + 2 \text{perm.} \right] + B_{g,\text{cont}}(k_1, k_2, k_3). \quad (\text{B10})$$

This is equivalent to Eq. (B4), as long as the galaxy power spectrum $P_{g,\text{cont}}$ in that equation is taken to be the one with the scale dependent non-Gaussian bias $b_{10} + b_{01}/M(k)$.

For the connected trispectrum, we have four positions, which allow for five different configurations: all positions equal, three positions equal but different from the fourth one, two pairs of positions equal but different from the other pair, one pair of positions equal but different from all other positions and finally, all four positions distinct. The trispectrum can then be written as

$$\begin{aligned} T_{g,\text{conn}}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) &= \frac{1}{V} \langle \delta_g(\mathbf{k}_1) \delta_g(\mathbf{k}_2) \delta_g(\mathbf{k}_3) \delta_g(-\mathbf{k}_1 - \mathbf{k}_2 - \mathbf{k}_3) \rangle = \frac{V^3}{N^3 N} \sum_{i,j,s,t} \langle \exp[i\mathbf{k}_1(\mathbf{x}_i - \mathbf{x}_t) + i\mathbf{k}_2(\mathbf{x}_j - \mathbf{x}_t) + i\mathbf{k}_3(\mathbf{x}_s - \mathbf{x}_t)] \rangle \\ &= \frac{V^3}{N^3 N} \sum_{i=j=s=t} \langle \exp[i\mathbf{k}_1(\mathbf{x}_i - \mathbf{x}_t) + i\mathbf{k}_2(\mathbf{x}_j - \mathbf{x}_t) + i\mathbf{k}_3(\mathbf{x}_s - \mathbf{x}_t)] \rangle + \frac{V^2}{N^2 N^2} \sum_{i=j=s \neq t} \langle \exp[-i\mathbf{k}_4(\mathbf{x}_i - \mathbf{x}_t)] \rangle + 3 \text{perm.} \\ &\quad + \frac{V^2}{N^2 N^2} \sum_{i=j \neq s=t} \langle \exp[i(\mathbf{k}_1 + \mathbf{k}_2)(\mathbf{x}_i - \mathbf{x}_t)] \rangle + 2 \text{perm.} \\ &\quad + \frac{V}{N N^3} \sum_{i=j \neq s \neq t, j \neq t} \langle \exp[i(\mathbf{k}_1 + \mathbf{k}_2)(\mathbf{x}_i - \mathbf{x}_t) + i\mathbf{k}_3(\mathbf{x}_s - \mathbf{x}_t)] \rangle + 5 \text{perm.} \\ &\quad + \frac{V^3}{N^4} \sum_{i \neq j \neq s \neq t, i \neq t, i \neq s, j \neq t} \langle \exp[i\mathbf{k}_1(\mathbf{x}_i - \mathbf{x}_t) + i\mathbf{k}_2(\mathbf{x}_j - \mathbf{x}_t) + i\mathbf{k}_3(\mathbf{x}_s - \mathbf{x}_t)] \rangle \\ &= \frac{1}{\bar{n}^3} + \frac{1}{\bar{n}^2} [P_{g,\text{cont}}(\mathbf{k}_4) + 3 \text{perm.}] + \frac{1}{\bar{n}^2} [P_{g,\text{cont}}(\mathbf{k}_1 + \mathbf{k}_2) + 2 \text{perm.}] \\ &\quad + \frac{1}{\bar{n}} [B_{g,\text{cont}}(\mathbf{k}_1 + \mathbf{k}_2, \mathbf{k}_3) + 5 \text{perm.}] + T_{g,\text{cont}}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3). \end{aligned} \quad (\text{B11})$$

Furthermore, there is a disconnected cosmic variance contribution for counteraligned pairs of momenta

$$\begin{aligned} T_{g,\text{disconn}}(\mathbf{k}_1, -\mathbf{k}_1, \mathbf{k}_3) &= \frac{1}{V} \langle \delta_g(\mathbf{k}_1) \delta_g(-\mathbf{k}_1) \delta_g(\mathbf{k}_3) \delta_g(-\mathbf{k}_3) \rangle = \frac{V^3}{N^3 N} \sum_{i,j,s,t} \langle \exp[i\mathbf{k}_1(\mathbf{x}_i - \mathbf{x}_j) + i\mathbf{k}_3(\mathbf{x}_s - \mathbf{x}_t)] \rangle \\ &= V \frac{V}{N^2} \sum_{i,j} \langle \exp[i\mathbf{k}_1(\mathbf{x}_i - \mathbf{x}_j)] \rangle \frac{V}{N^2} \sum_{s,t} \langle \exp[i\mathbf{k}_3(\mathbf{x}_s - \mathbf{x}_t)] \rangle. \end{aligned} \quad (\text{B12})$$

In the continuous case this disconnected contribution to the four-point function becomes

$$\langle \delta_g(\mathbf{k}_1) \delta_g(\mathbf{k}_2) \delta_g(\mathbf{k}_3) \delta_g(\mathbf{k}_4) \rangle \supset (2\pi)^6 \delta_D(\mathbf{k}_1 + \mathbf{k}_2) \delta_D(\mathbf{k}_3 + \mathbf{k}_4) \left[P_{g,\text{cont}}(k_1) + \frac{1}{\bar{n}} \right] \left[P_{g,\text{cont}}(k_3) + \frac{1}{\bar{n}} \right] + 2 \text{ perm.} \quad (\text{B13})$$

This will give the Gaussian contribution to the covariance of the estimator.

Thus, the total covariance for the auto-spectrum of the reconstructed field is

$$\begin{aligned} \langle \hat{\Delta}_\alpha(\mathbf{K}) \hat{\Delta}_\beta(\mathbf{K}') \rangle - \langle \hat{\Delta}_\alpha(\mathbf{K}) \rangle \langle \hat{\Delta}_\beta(\mathbf{K}') \rangle &= \int_{\mathbf{q}} \int_{\mathbf{q}'} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) g_\beta(\mathbf{q}', \mathbf{K}' - \mathbf{q}') [\langle \delta_g(\mathbf{q}) \delta_g(\mathbf{K} - \mathbf{q}) \delta_g(\mathbf{q}') \delta_g(\mathbf{K}' - \mathbf{q}') \rangle \\ &\quad - \langle \delta_g(\mathbf{q}) \delta_g(\mathbf{K} - \mathbf{q}) \rangle \langle \delta_g(\mathbf{q}') \delta_g(\mathbf{K}' - \mathbf{q}') \rangle] \\ &= (2\pi)^3 \delta_D(\mathbf{K} + \mathbf{K}') \left[\frac{N_{\alpha\alpha}(\mathbf{K}) N_{\beta\beta}(\mathbf{K})}{N_{\alpha\beta}(\mathbf{K})} + N_{\alpha\beta,\text{shot}}(\mathbf{K}) \right], \end{aligned} \quad (\text{B14})$$

where

$$\begin{aligned} N_{\alpha\beta,\text{shot}}^T(\mathbf{K}) &\equiv \int_{\mathbf{q}} \int_{\mathbf{q}'} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) g_\beta(\mathbf{q}', -\mathbf{K} - \mathbf{q}') T_{g,\text{conn}}(\mathbf{q}, \mathbf{K} - \mathbf{q}, \mathbf{q}', -\mathbf{K} - \mathbf{q}') \\ &= \int_{\mathbf{q}} \int_{\mathbf{q}'} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) g_\beta(\mathbf{q}', -\mathbf{K} - \mathbf{q}') \left\{ \frac{1}{\bar{n}^3} + \frac{1}{\bar{n}^2} [P_{g,\text{cont}}(\mathbf{q}) + P_{g,\text{cont}}(\mathbf{K} - \mathbf{q}) + P_{g,\text{cont}}(\mathbf{q}') + P_{g,\text{cont}}(\mathbf{K} + \mathbf{q}')] \right. \\ &\quad + \frac{1}{\bar{n}^2} [P_{g,\text{cont}}(\mathbf{K}) + P_{g,\text{cont}}(\mathbf{q} - \mathbf{K} - \mathbf{q}') + P_{g,\text{cont}}(\mathbf{q} + \mathbf{q}')] \\ &\quad + \frac{1}{\bar{n}} [B_{g,\text{cont}}(\mathbf{K}, \mathbf{q}', -\mathbf{K} - \mathbf{q}') + B_{g,\text{cont}}(\mathbf{q} + \mathbf{q}', \mathbf{K} - \mathbf{q}, -\mathbf{K} - \mathbf{q}')] \\ &\quad + B_{g,\text{cont}}(\mathbf{q} - \mathbf{K} - \mathbf{q}', \mathbf{K} - \mathbf{q}, \mathbf{q}') + B_{g,\text{cont}}(\mathbf{K} - \mathbf{q} + \mathbf{q}', \mathbf{q}, -\mathbf{K} - \mathbf{q}') \\ &\quad \left. + B_{g,\text{cont}}(\mathbf{K} - \mathbf{q} - \mathbf{K} - \mathbf{q}', \mathbf{q}, \mathbf{q}') + B_{g,\text{cont}}(\mathbf{q}' - \mathbf{K} - \mathbf{q}', \mathbf{q}, \mathbf{K} - \mathbf{q}) \right\}. \end{aligned} \quad (\text{B15})$$

For our forecasts, we use the “growth” estimator, and therefore the conversion to the notation of the main text is

$$P_{\text{rr},\text{shot}}(K) = N_{\text{GG},\text{shot}}^T(K). \quad (\text{B16})$$

Note an important feature of Eq. (B15): the shot noise contribution to the quadratic estimator’s noise power spectrum depends on the tracer power spectrum and bispectrum at the same scale K as the mode being reconstructed. This is to be contrasted with the Gaussian estimator noise in Eq. (15), which is mainly determined by the tracer power spectrum at the smallest scale k_{max} used in the estimator. Since the tracer power spectrum and bispectrum both increase at smaller wave numbers (down to the matter-radiation equality scale), Eq. (B15)’s sensitivity to large scales can cause it to dominate over the Gaussian

estimator noise if $\bar{n}$ is sufficiently low. In our forecasts in the main text, this condition is met for DESI (Fig. 9) and MegaMapper (Fig. 11).

We can simplify the above expression by noticing that changes of variables can make some terms of the integrand equivalent. For example, we can simplify the final 3 lines of Eq. (B15) into

$$\begin{aligned} \int_{\mathbf{q}} \int_{\mathbf{q}'} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) g_\beta(\mathbf{q}', -\mathbf{K} - \mathbf{q}') \frac{1}{\bar{n}} (B_{g,\text{cont}}(\mathbf{K}, \mathbf{q}', -\mathbf{K} - \mathbf{q}') \\ + 4B_{g,\text{cont}}(-\mathbf{q} - \mathbf{q}', \mathbf{q}, \mathbf{q}') + B_{g,\text{cont}}(-\mathbf{K}, \mathbf{q}, \mathbf{K} - \mathbf{q})). \end{aligned} \quad (\text{B17})$$

In our calculations, we take the tree-level expression for the bispectrum, obtainable from Eq. (36) as (see also [71,74])

$$B_{g,\text{cont}}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) = 2 \left(b_{10} + f_{\text{NL}} \frac{c_{01}}{M(\mathbf{k}_1)} \right) \left(b_{10} + f_{\text{NL}} \frac{c_{01}}{M(\mathbf{k}_2)} \right) \sum_{\alpha} c_{\alpha} F_{\alpha}(\mathbf{k}_1, \mathbf{k}_2) P_{\text{lin}}(\mathbf{k}_1) P_{\text{lin}}(\mathbf{k}_2) + 2 \text{ perms.} \quad (\text{B18})$$

2. Noise for the cross correlation of the reconstructed field with the tracer field

The noise calculation for the cross-correlation of the reconstructed field with the tracer field is very similar to the one above. The variance of the cross-correlation of the reconstructed field with the tracer field depends on the bispectrum of the tracer field derived above in Eq. (B4) (see also [100]):

$$\begin{aligned} & \langle \hat{\Delta}_\alpha(\mathbf{K}) \delta_g(\mathbf{K}') \rangle - \langle \hat{\Delta}_\alpha(\mathbf{K}) \rangle \langle \delta_g(\mathbf{K}') \rangle \\ &= \langle \hat{\Delta}_\alpha(\mathbf{K}) \delta_g(\mathbf{K}') \rangle = \int_{\mathbf{q}} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) \langle \delta_g(\mathbf{q}) \delta_g(\mathbf{K} - \mathbf{q}) \delta_g(\mathbf{K}') \rangle \\ &= (2\pi)^3 \delta_D(\mathbf{K} + \mathbf{K}') \int_{\mathbf{q}} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) \left[B_{g,\text{cont}}(\mathbf{q}, \mathbf{K} - \mathbf{q}, -\mathbf{K}) \right. \\ & \quad \left. + \frac{1}{\bar{n}} (P_{g,\text{cont}}(\mathbf{q}) + P_{g,\text{cont}}(\mathbf{K} - \mathbf{q}) + P_{g,\text{cont}}(-\mathbf{K})) + \frac{1}{\bar{n}^2} \right]. \end{aligned} \quad (\text{B19})$$

Then, the shot noise power is given by

$$\begin{aligned} N_{\alpha,\text{shot}}^B(\mathbf{K}) &\equiv \frac{1}{\bar{n}} \int_{\mathbf{q}} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) \left[P_{g,\text{cont}}(\mathbf{q}) + P_{g,\text{cont}}(\mathbf{K} - \mathbf{q}) \right. \\ & \quad \left. + P_{g,\text{cont}}(-\mathbf{K}) + \frac{1}{\bar{n}} \right]. \end{aligned} \quad (\text{B20})$$

For the growth estimator we use in our forecasts we thus have

$$P_{\text{gr,shot}}(K) = N_{G,\text{shot}}^B(K). \quad (\text{B21})$$

APPENDIX C: BIAS-HARDENING

In Sec. II A, we saw that a quadratic estimator designed to have unit response to a specific form of mode-coupling will generically acquire a mean-field contamination from other forms of mode coupling that are not incorporated in the estimator's weights [see Eq. (17)]. Our main approach in this paper is to include that contamination in our model for the estimator's output, marginalizing over the associated free (bias) parameters where necessary. Alternatively, one can attempt to define an estimator that is orthogonal to those extra mode-couplings; such a “bias-hardening” procedure has been applied to weak lensing of the CMB (e.g., [68,69]) and line intensity maps [2]. In this Appendix, we explore this approach and explain why we did not find it to be useful for this study.

1. General derivation

First, we derive a form of bias-hardening that is a light generalization of the standard form (e.g., [68]). Recall that a quadratic estimator with weights g_α is given by

$$\hat{\Delta}_\alpha(\mathbf{K}) = \int_{\mathbf{q}} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) \delta_g(\mathbf{q}) \delta_g(\mathbf{K} - \mathbf{q}), \quad (\text{C1})$$

with expectation value

$$\begin{aligned} & \langle \hat{\Delta}_\alpha(\mathbf{K}) \rangle_{\delta_1(\mathbf{K})\text{fixed}} \\ &= \sum_{\beta} \left[c_\beta \int_{\mathbf{q}} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) f_\beta(\mathbf{q}, \mathbf{K} - \mathbf{q}) \right] b_1 \delta_1(\mathbf{K}). \end{aligned} \quad (\text{C2})$$

The estimator will be unbiased with respect to the α mode-coupling if

$$\int_{\mathbf{q}} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) f_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) = 1. \quad (\text{C3})$$

The $\beta \neq \alpha$ terms in the sum in Eq. (C2) could be subtracted if we knew the values of the c_β coefficients ahead of time, but this will generally not be true. Instead, we can attempt to set the weights g_α such that

$$\int_{\mathbf{q}} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) f_\beta(\mathbf{q}, \mathbf{K} - \mathbf{q}) = 0, \quad (\text{C4})$$

for all $\beta \neq \alpha$.

We begin by assuming that there is only one additional mode-coupling β that we are concerned with, and requiring that the estimator's response to it is not necessarily zero, but a chosen constant c instead:

$$\int_{\mathbf{q}} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) f_\beta(\mathbf{q}, \mathbf{K} - \mathbf{q}) = c, \quad (\text{C5})$$

while also imposing Eq. (C3) and minimizing the Gaussian contribution to the variance of $\hat{\Delta}_\alpha(\mathbf{K})$, given by

$$\begin{aligned} & \text{Var}_G[\hat{\Delta}_\alpha(\mathbf{K})] \\ &= 2 \int_{\mathbf{q}} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) g_\alpha^*(\mathbf{q}, \mathbf{K} - \mathbf{q}) P_{\text{tot}}(\mathbf{q}) P_{\text{tot}}(|\mathbf{K} - \mathbf{q}|). \end{aligned} \quad (\text{C6})$$

We find g_α that satisfies these conditions by the method of Lagrange multipliers, starting with the following function:

$$\begin{aligned} L[g_\alpha, \lambda, \lambda_*] &= 2 \int_{\mathbf{q}} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) g_\alpha^*(\mathbf{q}, \mathbf{K} - \mathbf{q}) P_{\text{tot}}(\mathbf{q}) P_{\text{tot}}(\mathbf{K} - \mathbf{q}) \\ & \quad - \lambda \left[\int_{\mathbf{q}} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) f_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) - 1 \right] \\ & \quad - \lambda_* \left[\int_{\mathbf{q}} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) f_{\beta_*}(\mathbf{q}, \mathbf{K} - \mathbf{q}) - c \right]. \end{aligned} \quad (\text{C7})$$

If we wanted to control the effect of other mode-couplings in Eq. (C2), we would simply add other terms with similar constraints to this equation. We demand that the functional derivative of L with respect to $g_\alpha(\mathbf{q}', \mathbf{K} - \mathbf{q}')$ vanishes:

$$\begin{aligned}
0 \stackrel{!}{=} \frac{\delta L[g_\alpha, \lambda, \lambda_*]}{\delta g_\alpha(\mathbf{q}', \mathbf{K} - \mathbf{q}')} &= 2 \int_{\mathbf{q}} [\delta_D(\mathbf{q} - \mathbf{q}') g_\alpha^*(\mathbf{q}, \mathbf{K} - \mathbf{q}) + g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) \delta_D(\mathbf{q} - \mathbf{q}')] P_{\text{tot}}(\mathbf{q}) P_{\text{tot}}(\mathbf{K} - \mathbf{q}) \\
&\quad - \lambda \int_{\mathbf{q}} \delta_D(\mathbf{q} - \mathbf{q}') f_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) - \lambda_* \int_{\mathbf{q}} \delta_D(\mathbf{q} - \mathbf{q}') f_\beta(\mathbf{q}, \mathbf{K} - \mathbf{q}) \\
&= 4g_\alpha(\mathbf{q}', \mathbf{K} - \mathbf{q}') P_{\text{tot}}(\mathbf{q}') P_{\text{tot}}(\mathbf{K} - \mathbf{q}') - \lambda f_\alpha(\mathbf{q}', \mathbf{K} - \mathbf{q}') - \lambda_* f_\beta(\mathbf{q}', \mathbf{K} - \mathbf{q}'), \tag{C8}
\end{aligned}$$

where in the last line, we took g_α to be real. This implies that (relabeling $\mathbf{q}' \rightarrow \mathbf{q}$)

$$g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) = \frac{\lambda f_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) + \lambda_* f_\beta(\mathbf{q}, \mathbf{K} - \mathbf{q})}{4P_{\text{tot}}(\mathbf{q})P_{\text{tot}}(\mathbf{K} - \mathbf{q})}, \tag{C9}$$

and plugging this into Eq. (C3) gives

$$\lambda = 2N_{\alpha\alpha} - \lambda_* \frac{N_{\alpha\alpha}}{N_{\alpha\beta}}, \tag{C10}$$

using the definition of $N_{\alpha\beta}$ from Eq. (15). (Note that we obtain the original filter if $\lambda_* = 0$.) Inserting Eqs. (C9) and (C10) into (C5), we get¹⁵

$$\lambda_* = 2 \frac{1}{1 - r_{\alpha\beta}^2} N_{\beta\beta}(\mathbf{K}) \left[c(\mathbf{K}) - \frac{N_{\alpha\alpha}(\mathbf{K})}{N_{\alpha\beta}(\mathbf{K})} \right], \tag{C11}$$

where

$$r_{\alpha\beta}^2 \equiv \frac{N_{\alpha\alpha} N_{\beta\beta}}{N_{\alpha\beta}^2}. \tag{C12}$$

Thus, the final form of the weight function is

$$g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) = \frac{N_{\alpha\alpha}(\mathbf{K}) f_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q})}{2P_{\text{tot}}(\mathbf{q})P_{\text{tot}}(\mathbf{K} - \mathbf{q})} + \frac{f_{\beta_*}(\mathbf{q}, \mathbf{K} - \mathbf{q}) N_{\beta\beta}(\mathbf{K}) - \frac{N_{\alpha\alpha}(\mathbf{K}) N_{\beta\beta}(\mathbf{K})}{N_{\alpha\beta}(\mathbf{K})} f_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q})}{2P_{\text{tot}}(\mathbf{q})P_{\text{tot}}(\mathbf{K} - \mathbf{q})} \frac{1}{1 - r_{\alpha\beta}^2} \left[c(\mathbf{K}) - \frac{N_{\alpha\alpha}(\mathbf{K})}{N_{\alpha\beta}(\mathbf{K})} \right]. \tag{C13}$$

This is the standard filter for mode-coupling α , plus some additional terms related to the response to mode-coupling β . For the case $c = N_{\alpha\alpha} N_{\alpha\beta}^{-1}$, one obtains the standard estimator and contamination term.

2. Application to long-mode reconstruction

To relate the bias-hardened estimator derived above to the specific application we consider in this paper, let us examine the variance of the estimator:

$$\text{Var}[\hat{\Delta}_\alpha^H(\mathbf{K})] = N_{\alpha\alpha}(\mathbf{K}) + N_{\beta\beta}(\mathbf{K}) \frac{[c(\mathbf{K}) - \frac{N_{\alpha\alpha}(\mathbf{K})}{N_{\alpha\beta}(\mathbf{K})}]^2}{1 - r_{\alpha\beta}^2}. \tag{C14}$$

If $|c(\mathbf{K})| < |N_{\alpha\alpha}(\mathbf{K}) N_{\alpha\beta}(\mathbf{K})^{-1}|$, then the increase in the estimator's variance scales with $(1 - r_{\alpha\beta}^2)^{-1}$, where $r_{\alpha\beta}$ is

¹⁵To generalize to several mode-couplings, the coefficients λ_{i*} are obtained from $\lambda_* = -2\mathbf{A}^{-1}\mathbf{c}$, where $\mathbf{A}$ has elements $A_{ij} = (N_{\beta_i\beta_j}^{-1} - N_{\alpha\alpha} N_{\alpha\beta_j}^{-1} N_{\alpha\beta_i}^{-1})$ and $\mathbf{c}$ has elements $N_{\alpha\alpha} N_{\alpha\beta_i}^{-1}$.

the correlation coefficient between the unhardened estimator $\hat{\Delta}_\alpha$ and the analogous estimator for the other mode-coupling, $\hat{\Delta}_\beta$. In Fig. 16, we show $r_{\alpha\beta}(\mathbf{K})$ for the growth, shift, and tidal mode-couplings introduced in Sec. II A. It is clear that the corresponding estimators are highly correlated, so that any bias-hardened estimator will have a much larger variance than without bias-hardening. In our numerical tests (with $c = 0$), when one of these three estimators was bias-hardened with respect to the other two, we found that the variance increased enough to eliminate any advantages of removing the mean-field contamination, and therefore we did not implement any bias-hardening in our final forecasts.

In the course of this investigation, we derived a compact form for a bias-hardened quadratic estimator in the case of three mode-couplings, and we reproduce this derivation here in case it may be useful in other contexts. Considering only G, S, and T, the expectation values of the corresponding quadratic estimators [see Eq. (17)] can be written in matrix form:

FIG. 16. Cross-correlation coefficients between G, S, and T estimators, as given by Eq. (C12), for our DESI-like survey forecast (the other surveys give similar results) We see that all three estimators are highly correlated, implying, through Eq. (C14), that applying bias-hardening will strongly increase the noise of the resulting estimator.

$$\begin{bmatrix} \langle \hat{\Delta}_G \rangle \\ \langle \hat{\Delta}_S \rangle \\ \langle \hat{\Delta}_T \rangle \end{bmatrix} = b_1 \begin{bmatrix} 1 & N_{GG}N_{GS}^{-1} & N_{GG}N_{GT}^{-1} \\ N_{SS}N_{SG}^{-1} & 1 & N_{SS}N_{TS}^{-1} \\ N_{TT}N_{TG}^{-1} & N_{TT}N_{TS}^{-1} & 1 \end{bmatrix} \begin{bmatrix} c_G \\ c_S \\ c_T \end{bmatrix} \delta_1. \quad (\text{C15})$$

We can derive bias-hardened estimators (in the $c(\mathbf{K}) = 0$ case) by inverting this system, solving for each $b_1 c_\alpha \delta_1$. After some lengthy algebra, the results can be written in terms of the original variances plus certain combinations of the original cross-correlation coefficients:

$$\begin{aligned} \text{Var}[\hat{\Delta}_G^H] &= N_{GG} \times \frac{1 - r_{ST}^2}{\det M}, \\ \text{Var}[\hat{\Delta}_S^H] &= N_{SS} \times \frac{1 - r_{GT}^2}{\det M}, \\ \text{Var}[\hat{\Delta}_T^H] &= N_{TT} \times \frac{1 - r_{GS}^2}{\det M}, \end{aligned} \quad (\text{C16})$$

where M denotes the matrix in Eq. (C15), and

$$\det M = 1 - r_{GS}^2 - r_{GT}^2 - r_{ST}^2 + 2r_{GS}r_{GT}r_{ST}. \quad (\text{C17})$$

There can be nontrivial cancellations within the above determinant, and in our case, these lead to large increases in the variances of the bias-hardened estimators.

Finally, we mention a few other possible solutions to the problem of mean-field contamination. Instead of fixing the contamination to some value, as in Eq. (C5), one could require it to be smaller than some fixed value, or one could minimise some total function that depends on the contamination and the (Gaussian) variance. For example it is possible to define a “bouncing estimator” by solving the following minimization problem:

$$\begin{aligned} L[g_\alpha, \lambda, \mathbf{K}] &= 2 \int_{\mathbf{q}} |g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q})|^2 P_{\text{tot}}(\mathbf{q}) P_{\text{tot}}(\mathbf{K} - \mathbf{q}) \\ &+ A \left[\int_{\mathbf{q}} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) f_\beta(\mathbf{q}, \mathbf{K} - \mathbf{q}) \right]^2 \\ &- \lambda \left[\int_{\mathbf{q}} g_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) f_\alpha(\mathbf{q}, \mathbf{K} - \mathbf{q}) - 1 \right]. \end{aligned} \quad (\text{C18})$$

The intuition behind this is that we want to minimize the variance of the α estimator, trying also to take into account the contamination from the other mode-coupling. The square is to ensure that the modulus of the contamination is minimized in the combination that makes L the smallest. If we go to the standard minimum variance solution, the solution here will “bounce” from it, because it would increase L , if we take into account the square of the contamination. We also want to decrease the contamination, but without it taking a large negative value. The last term enforces the standard unbiasedness condition.

We reiterate, however, that for some applications, such as the f_{NL} constraints we consider in this paper, a mean-field contamination can actually be advantageous. We leave the problem of finding a fully optimal estimator for long-wavelength reconstruction to future work.

APPENDIX D: CONTAMINATION OF QUADRATIC ESTIMATOR BY f_{NL} TERMS: ANALYTICAL EXPRESSIONS

The overdensity of a biased tracer has second order contributions in the linear field coming from the presence of primordial non-Gaussianity. In this Appendix, we show that when reconstructing the new field on large scales, we get an f_{NL} term proportional to $\frac{1}{M(\mathbf{K})}$. We will show this for the low- $|\mathbf{K}|$ limit, which is the relevant regime for the reconstructed modes we are concerned with.

1. Expansion of basic quantities

We need to expand a few quantities first. Expanding the linear power spectrum around $\mathbf{q}$ gives

$$\begin{aligned} \lim_{|\mathbf{K}| \rightarrow 0} P_{\text{lin}}(|\mathbf{K} - \mathbf{q}|) &= P_{\text{lin}}(|\mathbf{q}|) - \nabla_{\mathbf{q}} P_{\text{lin}} \cdot \mathbf{K} + O(|\mathbf{K}|^2) \\ &= P_{\text{lin}}(|\mathbf{q}|) - \frac{\partial P_{\text{lin}}}{\partial |\mathbf{q}|} \nabla_{\mathbf{q}} |\mathbf{q}| \cdot \mathbf{K} + O(|\mathbf{K}|^2) \\ &= P_{\text{lin}}(\mathbf{q}) \left(1 - \frac{|\mathbf{K}| \mu}{|\mathbf{q}|} \frac{\partial \ln P_{\text{lin}}}{\partial \ln |\mathbf{q}|} \right) + O(|\mathbf{K}|^2) \\ &\equiv P_{\text{lin}}(\mathbf{q}) (1 + \Delta_P) + O(|\mathbf{K}|^2), \end{aligned} \quad (\text{D1})$$

where the last expression is useful because we can see the expansion in powers of $\frac{|\mathbf{K}|}{|\mathbf{q}|}$, and where μ is the cosine of the angle between $\mathbf{K}$ and $\mathbf{q}$. In the same way, we can also expand for

$$\begin{aligned} \lim_{|\mathbf{K}| \rightarrow 0} P_{\text{NL}}(|\mathbf{K} - \mathbf{q}|) &= P_{\text{NL}}(\mathbf{q}) \left(1 - \frac{|\mathbf{K}| \mu}{|\mathbf{q}|} \frac{\partial \ln P_{\text{NL}}}{\partial \ln |\mathbf{q}|} \right) + O(|\mathbf{K}|^2) \\ &\equiv P_{\text{NL}}(\mathbf{q})(1 + \Delta_{\text{NL}}) + O(|\mathbf{K}|^2), \end{aligned} \quad (\text{D2})$$

and

$$\begin{aligned} \lim_{|\mathbf{K}| \rightarrow 0} M|\mathbf{K} - \mathbf{q}| &= M(\mathbf{q}) \left(1 - \frac{|\mathbf{K}| \mu}{|\mathbf{q}|} \frac{\partial \ln M}{\partial \ln |\mathbf{q}|} \right) + O(|\mathbf{K}|^2) \\ &\equiv M(\mathbf{q})(1 + \Delta_M). \end{aligned} \quad (\text{D3})$$

We can also write

$$\begin{aligned} \lim_{|\mathbf{K}| \rightarrow 0} |\mathbf{K} - \mathbf{q}| &= \lim_{|\mathbf{K}| \rightarrow 0} \sqrt{|\mathbf{K}|^2 + |\mathbf{q}|^2 - 2|\mathbf{K}||\mathbf{q}|\mu} \\ &\approx |\mathbf{q}| \left(1 - \frac{|\mathbf{K}|}{|\mathbf{q}|} \mu \right), \end{aligned} \quad (\text{D4})$$

and the expression for the cosine of the angle between $\mathbf{K}$ and $\mathbf{K} - \mathbf{q}$ as

$$\begin{aligned} \mu' &= \lim_{|\mathbf{K}| \rightarrow 0} \approx -\mu + (1 - \mu^2) \frac{|\mathbf{K}|}{|\mathbf{q}|} + \frac{|\mathbf{K}|^2}{|\mathbf{q}|^2} \mu \\ &\approx -\mu + (1 - \mu^2) \frac{|\mathbf{K}|}{|\mathbf{q}|}. \end{aligned} \quad (\text{D5})$$

2. Expansion of mode coupling expressions

As $|\mathbf{K}| \rightarrow 0$ for the next calculations we will assume $\frac{|\mathbf{K}|}{|\mathbf{q}|} \ll 1$ so that we keep terms linear in this variable. If the linear terms completely cancel, we include terms at the next order.

Recall the definition of $f_\beta(\mathbf{q}, \mathbf{K} - \mathbf{q})$ from Eq. (9):

$$f_\beta(\mathbf{q}, \mathbf{K} - \mathbf{q}) = 2[F_\alpha(\mathbf{K}, -\mathbf{q})P_{\text{lin}}(|\mathbf{q}|) + F_\alpha(\mathbf{K}, -\mathbf{K} + \mathbf{q})P_{\text{lin}}(|\mathbf{K} - \mathbf{q}|)]. \quad (\text{D6})$$

Starting from Eq. (D6), we have that for the $\varphi\varphi$ term

$$\begin{aligned} f_{\varphi\varphi}(\mathbf{q}, \mathbf{K} - \mathbf{q}) &= \frac{2}{M(\mathbf{K})} \left[\frac{M(\mathbf{K} - \mathbf{q})}{M(-\mathbf{q})} P_{\text{lin}}(\mathbf{q}) + \mathbf{q} \rightarrow \mathbf{K} - \mathbf{q} \right] \\ &= \frac{2P_{\text{lin}}(\mathbf{q})}{M(\mathbf{K})} \left[\frac{M(\mathbf{q})(1 + \Delta_M(\mathbf{q}))}{M(\mathbf{q})} + \frac{M(\mathbf{q})}{M(\mathbf{q})(1 + \Delta_M(\mathbf{q}))} (1 + \Delta_P(\mathbf{q})) \right] \approx \frac{2P_{\text{lin}}(\mathbf{q})}{M(\mathbf{K})} [2 + \Delta_P(\mathbf{q})], \end{aligned} \quad (\text{D7})$$

and for the b_{01} term

$$\begin{aligned} f_{01}(\mathbf{q}, \mathbf{K} - \mathbf{q}) &= \frac{2}{M(\mathbf{K})} \frac{1}{2} \left[(\mathbf{K} \cdot (-\mathbf{q})) \left(\frac{1}{|\mathbf{q}|^2} + \frac{M(\mathbf{K})}{|\mathbf{K}|^2 M(-\mathbf{q})} \right) P_{\text{lin}}(-\mathbf{q}) + \mathbf{q} \rightarrow \mathbf{K} - \mathbf{q} \right] \\ &= \frac{1}{2} \frac{2P_{\text{lin}}(\mathbf{q})}{M(\mathbf{K})} \left[\mathbf{K} \cdot (-\mathbf{q}) \left(\frac{1}{|\mathbf{q}|^2} + \frac{M(\mathbf{K})}{|\mathbf{K}|^2 M(\mathbf{q})} \right) + \mathbf{K} \cdot (-1)(\mathbf{K} - \mathbf{q}) \right. \\ &\quad \times \left. \left(\frac{1}{|\mathbf{q}|^2} \left(1 + \frac{|\mathbf{K}|}{|\mathbf{q}|} \mu \right) + \frac{M(\mathbf{K})}{K^2 M(\mathbf{q})} (1 - \Delta_M(\mathbf{q})) \right) (1 + \Delta_P(\mathbf{q})) \right] \\ &\approx \frac{1}{2} \frac{2P_{\text{lin}}(\mathbf{q})}{M(\mathbf{K})} \left[-\frac{|\mathbf{K}|^2}{|\mathbf{q}|^2} - \frac{|\mathbf{K}|^3}{|\mathbf{q}|^3} \mu - \frac{M(\mathbf{K})}{M(\mathbf{q})} (1 + \Delta_P(\mathbf{q}) - \Delta_M(\mathbf{q})) + \frac{|\mathbf{K}|^2}{|\mathbf{q}|^2} \mu^2 + \frac{q\mu M(\mathbf{K})}{|\mathbf{K}| M(\mathbf{q})} (\Delta_P(\mathbf{q}) - \Delta_M(\mathbf{q})) \right] \\ &\approx \frac{1}{2} \frac{2P_{\text{lin}}(\mathbf{q})}{M(\mathbf{q})} \left[-1 - \Delta_P(\mathbf{q}) - \Delta_M(\mathbf{q}) - (1 - \mu^2) \frac{|\mathbf{K}|^2}{|\mathbf{q}|^2} \right], \end{aligned} \quad (\text{D8})$$

where we remember that $M(\mathbf{q}) \propto |\mathbf{q}|^{-2}$. For the b_{11} term,

$$\begin{aligned} f_{11}(\mathbf{q}, \mathbf{K} - \mathbf{q}) &= \frac{2}{M(\mathbf{K})} \frac{1}{2} \left[\left(1 + \frac{M(\mathbf{K})}{M(-\mathbf{q})} \right) P_{\text{lin}}(-\mathbf{q}) + \mathbf{q} \rightarrow \mathbf{K} - \mathbf{q} \right] \\ &= \frac{P_{\text{lin}}(\mathbf{q})}{M(\mathbf{K})} \left[1 + \frac{M(\mathbf{K})}{M(\mathbf{q})} + 1 + \frac{M(\mathbf{K})}{M(\mathbf{q})} (1 - \Delta_M(\mathbf{q})) (1 + \Delta_P(\mathbf{q})) \right] \\ &\approx \frac{P_{\text{lin}}(\mathbf{q})}{M(\mathbf{K})} \left[2 + \frac{M(\mathbf{K})}{M(\mathbf{q})} (1 + \Delta_P(\mathbf{q}) - \Delta_M(\mathbf{q})) \right], \end{aligned} \quad (\text{D9})$$

and finally for the b_{02} term

$$\begin{aligned}
f_{02}(\mathbf{q}, \mathbf{K}-\mathbf{q}) &= \frac{2}{M(\mathbf{K})} \left[\frac{P_{\text{lin}}(-\mathbf{q})}{M(-\mathbf{q})} + \mathbf{q} \rightarrow \mathbf{K}-\mathbf{q} \right] \\
&= \frac{2P_{\text{lin}}(\mathbf{q})}{M(\mathbf{K})} \left[\frac{1}{M(\mathbf{q})} + \frac{1}{M(\mathbf{q})} (1 - \Delta_M(\mathbf{q}) + \Delta_P(\mathbf{q})) \right] \\
&\approx \frac{2P_{\text{lin}}(\mathbf{q})}{M(\mathbf{K})} \frac{1}{M(\mathbf{q})} [2 - \Delta_M(\mathbf{q}) + \Delta_P(\mathbf{q})]. \quad (\text{D10})
\end{aligned}
\quad
\begin{aligned}
f_G(\mathbf{q}, \mathbf{K}-\mathbf{q}) &= 2 \left[\frac{17}{21} P(\mathbf{q}) + \frac{17}{21} P_{\text{lin}}(\mathbf{K}-\mathbf{q}) \right] \\
&= 2 \frac{17}{21} P(\mathbf{q}) [2 + \Delta_P(\mathbf{q})]. \quad (\text{D11})
\end{aligned}$$

We wrote all of them in such a way that, when possible, we can factor out a $\frac{1}{M(\mathbf{K})}$. Finally the growth term can be written as

3. Writing the large scale contamination terms

Recall that the terms that contaminate the expectation value of the quadratic estimator are of the form $c_\alpha \frac{N_{\text{GG}}}{N_{\text{G}\alpha}}$ (see Eq. (17)). Therefore, for the large scale limit, we need to consider

$$\begin{aligned}
\lim_{|\mathbf{K}| \rightarrow 0} N_{\text{G}\alpha}^{-1}(\mathbf{K}) &= \lim_{|\mathbf{K}| \rightarrow 0} \frac{2\pi}{(2\pi)^3} \int_{-1}^1 d\mu \int_{q_{\min}}^{q_{\max}} dq q^2 \frac{f_\alpha(\mathbf{q}, \mathbf{K}-\mathbf{q})}{2P_{\text{tot}}(\mathbf{q})P_{\text{tot}}(\mathbf{K}-\mathbf{q})} f_g(\mathbf{q}, \mathbf{K}-\mathbf{q}) \\
&= \lim_{|\mathbf{K}| \rightarrow 0} \frac{1}{b_1^4} \frac{2\pi}{(2\pi)^3} \int_{-1}^1 d\mu \int_{q_{\min}}^{q_{\max}} dq q^2 \frac{f_\alpha(\mathbf{q}, \mathbf{K}-\mathbf{q})}{2P_{\text{NL}}^2(\mathbf{q})} 2 \frac{17}{21} P_{\text{lin}}(\mathbf{q}) [2 + \Delta_P(\mathbf{q})] (1 - \Delta_{\text{PN}}), \quad (\text{D12})
\end{aligned}$$

where we assume no shot noise in the total galaxy power spectrum. Dropping any nonzero power of $\frac{|\mathbf{K}|}{q}$, for small $\mathbf{K}$ (with respect to reconstruction modes $\mathbf{q}$), we obtain

$$\begin{aligned}
N_{\text{GG}}(\mathbf{K}) &\approx \left[\frac{1}{b_1^4} \frac{2\pi}{(2\pi)^3} \left(2 \frac{17}{21} \right)^2 4 \int_{q_{\min}}^{q_{\max}} dq q^2 \frac{P_{\text{lin}}^2(\mathbf{q})}{P_{\text{NL}}^2(\mathbf{q})} [1 + \Delta_P(\mathbf{q}) - \Delta_{\text{PN}}(\mathbf{q})] \right]^{-1} \\
&\approx \left[\frac{1}{b_1^4} \frac{4\pi}{(2\pi)^3} 8 \left(\frac{17}{21} \right)^2 \int_{q_{\min}}^{q_{\max}} dq q^2 \frac{P_{\text{lin}}^2(\mathbf{q})}{P_{\text{NL}}^2(\mathbf{q})} \right]^{-1}, \quad (\text{D13})
\end{aligned}$$

and we can see that if we approximate $P_{\text{lin}} \approx P_{\text{NL}}$ for the small scales of reconstruction, then the Gaussian reconstruction noise is roughly proportional to the volume shell between $q_{\min}$ and $q_{\max}$. Thus, we can take the noise as roughly proportional to $q_{\max}^3$, although in practice this relation is not exactly correct.

At this point, we can start listing the $N_{\text{G}\alpha}$ terms where α is a mode-coupling involving f_{NL} . We begin with $\alpha = \varphi\varphi$:

$$N_{\text{G}\varphi\varphi}^{-1}(\mathbf{K}) \approx \frac{1}{b_1^4} \frac{1}{2\pi^2} \left(\frac{136}{21} \right) \frac{1}{M(\mathbf{K})} \int_{q_{\min}}^{q_{\max}} dq q^2 \frac{P_{\text{lin}}^2(\mathbf{q})}{P_{\text{NL}}^2(\mathbf{q})}, \quad (\text{D14})$$

such that multiplying by Eq. (D13), we have a scaling for the $\alpha = \varphi\varphi$ term in the bias of the reconstructed field:

$$N_{\text{GG}}(\mathbf{K}) N_{\text{G}\varphi\varphi}^{-1}(\mathbf{K}) \approx \left(\frac{21}{17} \right) \frac{1}{M(\mathbf{K})}. \quad (\text{D15})$$

Similarly we can calculate an approximate expression for the $\alpha = 01$ term:

$$N_{\text{G}01}^{-1}(\mathbf{K}) \approx (-1) \frac{1}{b_1^4} \frac{1}{2\pi^2} \frac{1}{21} \frac{68}{21} \int_{q_{\min}}^{q_{\max}} dq q^2 \frac{P_{\text{lin}}^2(\mathbf{q})}{P_{\text{NL}}^2(\mathbf{q})} \frac{1}{M(\mathbf{q})}. \quad (\text{D16})$$

Multiplying this by the approximate N_{GG} expression, we obtain

$$N_{\text{GG}}(\mathbf{K}) N_{\text{G}01}^{-1}(\mathbf{K}) \propto -\frac{42}{17}, \quad (\text{D17})$$

implying that on large scales for the $\alpha = 01$ term we do not have a $\frac{1}{K^2}$ behavior, but a negative constant bias.

FIG. 17. We plot the contamination curves from numerical and analytic approximation. It can be seen that for the $q\phi\phi$, b_{11} and b_{02} terms we have a $\frac{1}{K^2}$ behavior. For the analytical calculation of the b_{02} term we use an approximation $(\int_{q_{\min}}^{q_{\max}} dq q^2 \frac{P_{\text{lin}}^2(q)}{P_{\text{NL}}^2(q)} \frac{1}{M(q)}) / (\int_{q_{\min}}^{q_{\max}} dq q^2 \frac{P_{\text{lin}}^2(q)}{P_{\text{NL}}^2(q)}) \approx \frac{1}{2M(q_{\min})}$. We also show the absolute value of the numerical b_{01} curve, with an approximate constant value on large scales.

Turning to the $\alpha = 02$ and 11 terms, we find

$$N_{G02}^{-1}(\mathbf{K}) \approx \frac{1}{b^4} \frac{1}{2\pi^2} \left(\frac{136}{21} \right) \frac{1}{M(\mathbf{K})} \int_{q_{\min}}^{q_{\max}} dq q^2 \frac{P_{\text{lin}}^2(q)}{P_{\text{NL}}^2(q)} \frac{1}{M(q)}, \quad (\text{D18})$$

and

$$N_{G11}^{-1}(\mathbf{K}) \approx \frac{1}{b^4} \frac{1}{2\pi^2} \left(\frac{68}{21} \right) \frac{1}{M(\mathbf{K})} \int_{q_{\min}}^{q_{\max}} dq q^2 \frac{P_{\text{lin}}^2(q)}{P_{\text{NL}}^2(q)} \times \left(1 + \frac{M(\mathbf{K})}{2M(q)} \right). \quad (\text{D19})$$

In both cases we end up with a $1/K^2$ behavior:

$$N_{\text{GG}}(\mathbf{K}) N_{G02}^{-1}(\mathbf{K}) = 21/17 \frac{1}{M(\mathbf{K})} \times \left(\int_{q_{\min}}^{q_{\max}} dq q^2 \frac{P_{\text{lin}}^2(q)}{P_{\text{NL}}^2(q)} \frac{1}{M(q)} \right) / \left(\int_{q_{\min}}^{q_{\max}} dq q^2 \frac{P_{\text{lin}}^2(q)}{P_{\text{NL}}^2(q)} \right), \quad (\text{D20})$$

$$N_{\text{GG}}(\mathbf{K}) N_{G11}^{-1}(\mathbf{K}) = \frac{121}{217} \frac{1}{M(\mathbf{K})}. \quad (\text{D21})$$

To summarise, we have found an induced contamination on the G estimator of the following form:

$$\sum_{\alpha \in \{q\phi\phi, 01, 11, 02\}} \int_{\mathbf{q}} c_{\alpha} g_{\text{G}}(\mathbf{q}, \mathbf{K} - \mathbf{q}) f_{\alpha}(\mathbf{q}, \mathbf{K} - \mathbf{q}) = \frac{1}{M(|\mathbf{K}|)} [f_{\text{NL}} A(\mathbf{K}) + f_{\text{NL}}^2 B(\mathbf{K})], \quad (\text{D22})$$

where A, B are some functions that can be calculated from the definitions or numerically. In Fig. 17, we show that the analytical approximations are in excellent agreement with the full numerical computations for the contamination curves.

APPENDIX E: FOREGROUNDS FOR 21 CM INTENSITY MAPPING

1. Implementation in forecasts

As discussed in Sec. IV C, the presence of foregrounds in 21 cm intensity mapping limits the modes of δ_{g} that can be directly observed. Specifically, foregrounds impose a minimum $k_{\parallel}$ value for these modes, and also obscure modes within a wedge-shaped region in the $k_{\parallel} - k_{\perp}$ plane. The modes within this wedge satisfy (e.g., [85])

$$k_{\parallel} < \beta(z) k_{\perp}, \quad (\text{E1})$$

where

$$\beta(z) \equiv \frac{\chi(z) H(z)}{c(1+z)} \sin(\theta_{\text{w}}), \quad (\text{E2})$$

$\chi(z)$ is the comoving distance to redshift z , and θ_{w} is the maximum angle from the beam center at which the power of a spectrally-smooth source will leak into other regions of Fourier space. The angle θ_{w} is typically related to the width of the primary beam; following [85], we take it to be 3 times the primary beam width of PUMA, or

FIG. 18. The fraction of small-scale modes lying outside the foreground wedge, for the low- z (solid) and high- z (dashed) PUMA redshift bins we use in our forecasts. To a good approximation, the reconstruction noise integrals $N_{\alpha\beta}$ will simply be scaled by the inverse of these fractions in the presence of the wedge.

FIG. 19. Forecasts for a PUMA-like survey, analogous to Fig. 13 except for using a cutoff on the line-of-sight component of accessible δ_g modes instead of an isotropic K cutoff. This is motivated by the fact that 21 cm foregrounds will preferentially obscure modes with low wave number components along the line of sight. The absolute values of $\sigma(f_{\text{NL}})$ are slightly higher when using $K_{\parallel,\text{min}}$ instead of K_{min} , since more modes are eliminated with a $K_{\parallel}$ cut, but the improvement in $\sigma(f_{\text{NL}})$ from including reconstructed modes is qualitatively similar to the case with K_{min} .

$\theta_w \approx 3 \times 1.22\lambda(z)/D_{\text{eff}}$, where $\lambda(z) = 21(1+z)$ cm and $D_{\text{eff}} \approx 5$ m is the effective dish diameter ($\eta_a^{1/2} \times 6$ m with aperture efficiency factor $\eta_a = 0.7$). Using the mean redshifts of the low- z and high- z bins we use in our forecasts, this yields $\beta \approx 0.38$ and 1.3 for each bin respectively.

The wedge will restrict the small-scale modes that can be used for reconstructing the longer modes via the quadratic estimator in Eq. (16), and we can account for this by restricting the reconstruction noise integral in Eq. (15) to modes outside the wedge. This means that, when reconstructing a mode with wave vector $\mathbf{K}$, the integration variable $\mathbf{q}$ must satisfy

$$|q_{\parallel}| > \beta q_{\perp}, \quad |K_{\parallel} - q_{\parallel}| > \beta(\mathbf{K} - \mathbf{q})_{\perp}. \quad (\text{E3})$$

Rather than implementing these restrictions directly in the integral for $N_{\alpha\beta}$, which would cause the result to depend on the full vector $\mathbf{K}$ instead of just the norm K , we use an approximation based on the fact that in the $q \gg K$ limit, $N_{\alpha\beta}$ scales like the inverse of the number of modes that contribute to the reconstruction.¹⁶ Thus, the effect of the

wedge is mostly to rescale $N_{\alpha\beta}$ by the inverse of the fraction of modes that are outside of the wedge, i.e., the fraction of the integration domain that satisfies Eq. (E3). This fraction will depend slightly on the direction of $\mathbf{K}$, and we account for this dependence by averaging the fraction over $\mu_K \equiv K_{\parallel}/K$, although the dependence is only mild.

We plot this angle-averaged fraction in Fig. 18 for both redshift bins we use for our PUMA forecast. In the $2 < z < 3$ bin, we find that around 60% of the small-scale modes are untouched by the wedge, while for the $5 < z < 6$ bin, only 20% of the modes remain, corresponding to a factor of 5 increase in the reconstruction noise compared to the no-wedge case.

For the shot noise contributions to the δ_r auto spectrum and $\delta_r - \delta_g$ cross spectrum (Eqs. (B15)–(B16) and (B20)–(B21), respectively), we directly implement the wedge in the angular limits of the integrals, but we find that it has a negligible effect, since these integrals are normalized with $N_{\alpha\beta}$ and the fractional change in the integrals and $N_{\alpha\beta}$ is very similar.

We also need to incorporate the wedge when we integrate the Fisher matrix in Eq. (42) over long-wavelength modes used for the f_{NL} constraint [Eq. (43)]. Within the wedge, we will not have access to δ_g , but we will have access to

¹⁶This is identical to what happens to the noise on the standard quadratic estimator for CMB lensing in the $\ell \gg L$ limit (e.g., [104]).

modes δ_r reconstructed with the quadratic estimator. Thus, as in our baseline forecasts with an isotropic $K_{\min}$, we sum the outside-wedge and inside-wedge Fisher matrices, each restricted to the appropriate integration domain, with the latter Fisher matrix determined solely by the covariance of the reconstructed modes.

In the next subsection, where we consider a $K_{\parallel,\min}$ instead of an isotropic $K_{\min}$, we likewise implement the restriction $K > K_{\parallel,\min}$ in Eq. (43), and add the contribution from reconstructed modes with $K < K_{\parallel,\min}$.

2. PUMA forecasts with $K_{\parallel,\min}$

In Sec. IV E 3, we showed forecasts assuming an isotropic $K_{\min}$ for δ_g . In Fig. 19, we repeat those forecasts, but with a $K_{\parallel,\min}$, assuming that all values of $K_{\perp}$ within the survey volume can be accessed. The absolute values we find for $\sigma(f_{\text{NL}})$ are slightly higher, due to the number of inaccessible modes being larger with a $K_{\parallel,\min}$ cutoff, but the results for the improvement in $\sigma(f_{\text{NL}})$ due to the inclusion of reconstructed modes are qualitatively similar to those in Sec. IV E 3.

-
- [1] W. Hu and T. Okamoto, *Astrophys. J.* **574**, 566 (2002).
 - [2] S. Foreman, P.D. Meerburg, A. van Engelen, and J. Meyers, *J. Cosmol. Astropart. Phys.* **07** (2018) 046.
 - [3] U.-L. Pen, R. Sheth, J. Harnois-Deraps, X. Chen, and Z. Li, [arXiv:1202.5804](#).
 - [4] H.-M. Zhu, U.-L. Pen, Y. Yu, X. Er, and X. Chen, *Phys. Rev. D* **93**, 103504 (2016).
 - [5] H.-M. Zhu, U.-L. Pen, Y. Yu, and X. Chen, *Phys. Rev. D* **98**, 043511 (2018).
 - [6] D. Li, H.-M. Zhu, and U.-L. Pen, *Phys. Rev. D* **100**, 023517 (2019).
 - [7] C. Modi, M. White, A. Slosar, and E. Castorina, *J. Cosmol. Astropart. Phys.* **11** (2019) 023.
 - [8] N.G. Karaçaylı and N. Padmanabhan, *Mon. Not. R. Astron. Soc.* **486**, 3864 (2019).
 - [9] P. Li, S. Dodelson, and R. A. Croft, *Phys. Rev. D* **101**, 083510 (2020).
 - [10] P. Li, R. A. Croft, and S. Dodelson, [arXiv:2007.00226](#).
 - [11] A. J. Ross *et al.*, *Mon. Not. R. Astron. Soc.* **428**, 1116 (2013).
 - [12] S. Ho *et al.*, *J. Cosmol. Astropart. Phys.* **05** (2015) 040.
 - [13] B. Kalus, W. Percival, D. Bacon, E. Mueller, L. Samushia, L. Verde, A. Ross, and J. Bernal, *Mon. Not. R. Astron. Soc.* **482**, 453 (2019).
 - [14] A. R. Parsons, J. C. Pober, J. E. Aguirre, C. L. Carilli, D. C. Jacobs, and D. F. Moore, *Astrophys. J.* **756**, 165 (2012).
 - [15] A. Liu, A. R. Parsons, and C. M. Trott, *Phys. Rev. D* **90**, 023018 (2014).
 - [16] A. Liu, A. R. Parsons, and C. M. Trott, *Phys. Rev. D* **90**, 023019 (2014).
 - [17] A. H. Guth, *Phys. Rev. D* **23**, 347 (1981).
 - [18] A. Linde, *Phys. Lett.* **108B**, 389 (1982).
 - [19] A. Linde, *Phys. Lett.* **116B**, 335 (1982).
 - [20] A. Slosar *et al.*, [arXiv:1903.09883](#).
 - [21] Y. Akrami *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A10 (2020).
 - [22] Y. Akrami *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A9 (2020).
 - [23] P. D. Meerburg *et al.*, [arXiv:1903.04409](#).
 - [24] M. Alvarez *et al.*, [arXiv:1412.4671](#).
 - [25] P. Ade *et al.* (Simons Observatory), *J. Cosmol. Astropart. Phys.* **02** (2019) 056.
 - [26] CMB-S4 Collaboration, [arXiv:1610.02743](#).
 - [27] R. Scoccimarro, E. Sefusatti, and M. Zaldarriaga, *Phys. Rev. D* **69**, 103513 (2004).
 - [28] G. D’Amico, J. Gleyzes, N. Kokron, D. Markovic, L. Senatore, P. Zhang, F. Beutler, and H. Gil-Marín, *J. Cosmol. Astropart. Phys.* **05** (2020) 005.
 - [29] M. M. Ivanov, M. Simonovic, and M. Zaldarriaga, *J. Cosmol. Astropart. Phys.* **05** (2020) 042.
 - [30] H. Gil-Marín, C. Wagner, J. Noreña, L. Verde, and W. Percival, *J. Cosmol. Astropart. Phys.* **12** (2014) 029.
 - [31] E. Sefusatti, M. Crocce, S. Pueblas, and R. Scoccimarro, *Phys. Rev. D* **74**, 023522 (2006).
 - [32] I. Kayo and M. Takada, [arXiv:1306.4684](#).
 - [33] K. C. Chan and L. Blot, *Phys. Rev. D* **96**, 023528 (2017).
 - [34] M. Schmittfull, T. Baldauf, and U. Seljak, *Phys. Rev. D* **91**, 043530 (2015).
 - [35] J. R. Fergusson, D. M. Regan, and E. P. S. Shellard, *Phys. Rev. D* **86**, 063511 (2012).
 - [36] J. Byun, A. Eggemeier, D. Regan, D. Seery, and R. E. Smith, *Mon. Not. R. Astron. Soc.* **471**, 1581 (2017).
 - [37] J.-P. Dai, L. Verde, and J.-Q. Xia, *J. Cosmol. Astropart. Phys.* **08** (2020) 007.
 - [38] A. Moradinezhad Dizgah, H. Lee, M. Schmittfull, and C. Dvorkin, *J. Cosmol. Astropart. Phys.* **04** (2020) 011.
 - [39] C.-T. Chiang, C. Wagner, F. Schmidt, and E. Komatsu, *J. Cosmol. Astropart. Phys.* **05** (2014) 048.
 - [40] R. de Putter, [arXiv:1802.06762](#).
 - [41] D. Gualdi, H. Gil-Marín, R. L. Schuhmann, M. Manera, B. Joachimi, and O. Lahav, *Mon. Not. R. Astron. Soc.* **484**, 3713 (2019).
 - [42] N. Dalal, O. Dore, D. Huterer, and A. Shirokov, *Phys. Rev. D* **77**, 123514 (2008).
 - [43] S. Matarrese and L. Verde, *Astrophys. J. Lett.* **677**, L77 (2008).
 - [44] A. Slosar, C. Hirata, U. Seljak, S. Ho, and N. Padmanabhan, *J. Cosmol. Astropart. Phys.* **08** (2008) 031.
 - [45] V. Desjacques and U. Seljak, *Classical Quantum Gravity* **27**, 124011 (2010).
 - [46] F. Schmidt and M. Kamionkowski, *Phys. Rev. D* **82**, 103002 (2010).
 - [47] T. Giannantonio, A. J. Ross, W. J. Percival, R. Crittenden, D. Bacher, M. Kilbinger, R. Nichol, and J. Weller, *Phys. Rev. D* **89**, 023511 (2014).

- [48] B. Leistedt, H. V. Peiris, and N. Roth, *Phys. Rev. Lett.* **113**, 221301 (2014).
- [49] E. Castorina *et al.*, *J. Cosmol. Astropart. Phys.* **09** (2019) 010.
- [50] U. Seljak, *Phys. Rev. Lett.* **102**, 021302 (2009).
- [51] P. McDonald and U. Seljak, *J. Cosmol. Astropart. Phys.* **10** (2009) 007.
- [52] N. Hamaus, U. Seljak, and V. Desjacques, *Phys. Rev. D* **84**, 083509 (2011).
- [53] M. Schmittfull and U. Seljak, *Phys. Rev. D* **97**, 123540 (2018).
- [54] R. H. Liu and P. C. Breysse, *Phys. Rev. D* **103**, 063520 (2021).
- [55] M. Ballardini, W. L. Matthewson, and R. Maartens, *Mon. Not. R. Astron. Soc.* **489**, 1950 (2019).
- [56] M. Münchmeyer, M. S. Madhavacheril, S. Ferraro, M. C. Johnson, and K. M. Smith, *Phys. Rev. D* **100**, 083508 (2019).
- [57] C.-T. Chiang, C. Wagner, F. Schmidt, and E. Komatsu, *J. Cosmol. Astropart. Phys.* **05** (2014) 048.
- [58] C.-T. Chiang, C. Wagner, A. G. Sánchez, F. Schmidt, and E. Komatsu, *J. Cosmol. Astropart. Phys.* **09** (2015) 028.
- [59] C.-T. Chiang, *Phys. Rev. D* **95**, 123517 (2017).
- [60] S. Adhikari, D. Jeong, and S. Shandera, *Phys. Rev. D* **94**, 083528 (2016).
- [61] A. Barreira and F. Schmidt, *J. Cosmol. Astropart. Phys.* **06** (2017) 053.
- [62] A. Barreira and F. Schmidt, *J. Cosmol. Astropart. Phys.* **11** (2017) 051.
- [63] P. Ade *et al.* (Planck Collaboration), *Astron. Astrophys.* **594**, A13 (2016).
- [64] B. D. Sherwin and M. Zaldarriaga, *Phys. Rev. D* **85**, 103523 (2012).
- [65] K. W. Masui and U.-L. Pen, *Phys. Rev. Lett.* **105**, 161302 (2010).
- [66] D. Jeong and M. Kamionkowski, *Phys. Rev. Lett.* **108**, 251301 (2012).
- [67] K. W. Masui, U.-L. Pen, and N. Turok, *Phys. Rev. Lett.* **118**, 221301 (2017).
- [68] T. Namikawa, D. Hanson, and R. Takahashi, *Mon. Not. R. Astron. Soc.* **431**, 609 (2013).
- [69] S. J. Osborne, D. Hanson, and O. Doré, *J. Cosmol. Astropart. Phys.* **03** (2014) 024.
- [70] T. Giannantonio and C. Porciani, *Phys. Rev. D* **81**, 063530 (2010).
- [71] T. Baldauf, U. Seljak, and L. Senatore, *J. Cosmol. Astropart. Phys.* **04** (2011) 006.
- [72] V. Desjacques, D. Jeong, and F. Schmidt, *Phys. Rep.* **733**, 1 (2018).
- [73] T. Baldauf, U. Seljak, V. Desjacques, and P. McDonald, *Phys. Rev. D* **86**, 083540 (2012).
- [74] M. Tellarini, A. J. Ross, G. Tasinato, and D. Wands, *J. Cosmol. Astropart. Phys.* **07** (2015) 004.
- [75] M. Biagetti, T. Lazeyras, T. Baldauf, V. Desjacques, and F. Schmidt, *Mon. Not. R. Astron. Soc.* **468**, 3277 (2017).
- [76] R. Scoccimarro, L. Hui, M. Manera, and K. C. Chan, *Phys. Rev. D* **85**, 083002 (2012).
- [77] V. Springel, *Mon. Not. R. Astron. Soc.* **364**, 1105 (2005).
- [78] M. M. Abidi and T. Baldauf, *J. Cosmol. Astropart. Phys.* **07** (2018) 029.
- [79] M. Tegmark, A. Taylor, and A. Heavens, *Astrophys. J.* **480**, 22 (1997).
- [80] A. Aghamousa *et al.* (DESI Collaboration), *arXiv:1611.00036*.
- [81] M. J. Wilson and M. White, *J. Cosmol. Astropart. Phys.* **10** (2019) 015.
- [82] S. Ferraro *et al.*, *arXiv:1903.09208*.
- [83] D. J. Schlegel *et al.*, *arXiv:1907.11171*.
- [84] J. J. M. Carrasco, S. Foreman, D. Green, and L. Senatore, *J. Cosmol. Astropart. Phys.* **07** (2014) 057.
- [85] R. Ansari *et al.* (Cosmic Visions 21 cm), *arXiv:1810.09572*.
- [86] A. Slosar *et al.* (PUMA Collaboration), *arXiv:1907.12559*.
- [87] O. Doré *et al.*, *arXiv:1412.4872*.
- [88] N. Hamaus, U. Seljak, V. Desjacques, R. E. Smith, and T. Baldauf, *Phys. Rev. D* **82**, 043515 (2010).
- [89] T. Baldauf, U. s. Seljak, R. E. Smith, N. Hamaus, and V. Desjacques, *Phys. Rev. D* **88**, 083507 (2013).
- [90] E. Castorina and A. M. Dizgah, *J. Cosmol. Astropart. Phys.* **10** (2020) 007.
- [91] J. R. Shaw, K. Sigurdson, M. Sitwell, A. Stebbins, and U.-L. Pen, *Phys. Rev. D* **91**, 083514 (2015).
- [92] A. Ghosh, F. Mertens, and L. V. E. Koopmans, *Mon. Not. R. Astron. Soc.* **474**, 4552 (2018).
- [93] T. Lazeyras, C. Wagner, T. Baldauf, and F. Schmidt, *J. Cosmol. Astropart. Phys.* **02** (2016) 018.
- [94] N. Kaiser, *Mon. Not. R. Astron. Soc.* **227**, 1 (1987).
- [95] A. Perko, L. Senatore, E. Jennings, and R. H. Wechsler, *arXiv:1610.09321*.
- [96] L. F. de la Bella, D. Regan, D. Seery, and S. Hotchkiss, *J. Cosmol. Astropart. Phys.* **11** (2017) 039.
- [97] V. Desjacques, D. Jeong, and F. Schmidt, *J. Cosmol. Astropart. Phys.* **12** (2018) 035.
- [98] M. Schmittfull, M. Simonović, V. Assassi, and M. Zaldarriaga, *Phys. Rev. D* **100**, 043514 (2019).
- [99] U. Seljak, N. Hamaus, and V. Desjacques, *Phys. Rev. Lett.* **103**, 091303 (2009).
- [100] D. Jeong, http://www.personal.psu.edu/duj13/dissertation/djeong_diss.pdf (2010).
- [101] H. A. Feldman, N. Kaiser, and J. A. Peacock, *Astrophys. J.* **426**, 23 (1994).
- [102] N. S. Sugiyama, S. Saito, F. Beutler, and H.-J. Seo, *Mon. Not. R. Astron. Soc.* **497**, 1684 (2020).
- [103] F. Schmidt, *Phys. Rev. D* **93**, 063512 (2016).
- [104] D. Hanson, A. Challinor, G. Efstathiou, and P. Bielewicz, *Phys. Rev. D* **83**, 043005 (2011).