

Consistently constraining f_{NL} with the squeezed lensing bispectrum using consistency relations

Samuel Goldstein^{1,*}, Oliver H. E. Philcox^{2,1}, J. Colin Hill¹, Angelo Esposito^{3,4} and Lam Hui¹

¹*Center for Theoretical Physics, Department of Physics,
Columbia University, New York, New York 10027, USA*

²*Simons Society of Fellows, Simons Foundation, New York, New York 10010, USA*

³*Dipartimento di Fisica, Sapienza Università di Roma, Piazzale Aldo Moro 2, I-00185 Rome, Italy*

⁴*INFN Sezione di Roma, Piazzale Aldo Moro 2, I-00185 Rome, Italy*



(Received 20 October 2023; accepted 26 January 2024; published 9 February 2024)

We introduce a nonperturbative method to constrain the amplitude of local-type primordial non-Gaussianity (f_{NL}) using squeezed configurations of the CMB lensing convergence and cosmic shear bispectra. First, we use cosmological consistency relations to derive a model for the squeezed limit of angular auto- and cross-bispectra of lensing convergence fields in the presence of f_{NL} . Using this model, we perform a Fisher forecast with specifications expected for upcoming CMB lensing measurements from the Simons Observatory and CMB-S4, as well as cosmic shear measurements from a Rubin LSST/*Euclid*-like experiment. Assuming a minimum multipole $\ell_{\text{min}} = 10$ and maximum multipole $\ell_{\text{max}} = 1400$, we forecast $\sigma_{f_{\text{NL}}} = 175$ (95) for Simons Observatory (CMB-S4). Our forecasts improve considerably for an LSST/*Euclid*-like cosmic shear experiment with three tomographic bins and $\ell_{\text{min}} = 10$ and $\ell_{\text{max}} = 1400$ (5000) with $\sigma_{f_{\text{NL}}} = 31$ (16). A joint analysis of CMB-S4 lensing and LSST/*Euclid*-like shear yields little gain over the shear-only forecasts; however, we show that a joint analysis could be useful if the CMB lensing convergence can be reliably reconstructed at larger angular scales than the shear field. The method presented in this work is a novel and robust technique to constrain local primordial non-Gaussianity from upcoming large-scale structure surveys that is completely independent of the galaxy field (and therefore any nuisance parameters such as b_{ϕ}), thus complementing existing techniques to constrain f_{NL} using the scale-dependent halo bias.

DOI: [10.1103/PhysRevD.109.043515](https://doi.org/10.1103/PhysRevD.109.043515)

I. INTRODUCTION

One of the primary goals of ongoing and upcoming galaxy surveys such as the Dark Energy Spectroscopic Instrument (DESI) [1], *Euclid* [2], *SPHEREx* [3], and the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST) [4] is to reveal information about the physics behind the primordial perturbations that evolved into present-day cosmic structures. In the standard cosmological model, these perturbations are produced during an inflationary epoch in which the Universe underwent a period of rapid accelerated expansion. The simplest single-field models of inflation predict initial conditions that are almost perfectly Gaussian and adiabatic [5–8]; however, a wealth of more complex models exist that predict departures from Gaussianity [9,10]. As such, searches for primordial non-Gaussianity (PNG) can powerfully probe the physics of the early Universe.

Currently, the tightest constraints on a wide range of parameters characterizing the amplitude of PNG in various

shapes come from analysis of the cosmic microwave background (CMB) [11]; nevertheless, large-scale structure (LSS) observations provide a complementary approach to competitively constrain PNG [12–18]. For local-type PNG, the subject of this work, LSS constraints are typically derived by taking advantage of its distinct imprint on halo clustering via the “scale-dependent bias.” This effect manifests as an enhancement in the amplitude of the large-scale power spectrum of biased tracers relative to the expectation from Gaussian initial conditions [19–22]. Constraining the amplitude of local PNG, f_{NL} ,¹ using the scale-dependent bias is a key goal of upcoming surveys. The potential of this technique was recently demonstrated in an analysis of DESI photometric clustering data, which provided the most precise LSS constraint on f_{NL} to date [18]. However, the impact of foreground and systematic effects on this result remains unclear.

A potential limitation of using the scale-dependent bias to constrain local PNG is that the derived constraints on

*sjg2215@columbia.edu

¹Throughout this paper, we use $f_{\text{NL}} \equiv f_{\text{NL}}^{\text{loc}}$.

f_{NL} require precise knowledge of the impact of local PNG on galaxy formation, due to the perfect degeneracy between f_{NL} and the non-Gaussian halo bias parameter b_ϕ [23–28]. This situation motivates the development of alternative methods to constrain local PNG using LSS that do not rely on the scale-dependent bias. A promising option is to instead use the weak lensing bispectrum, since this is sensitive to the (unbiased) total matter distribution.

The prospect of constraining local PNG using the weak lensing bispectrum was first discussed in Takada and Jain [29] (see also [30–32]), where the authors found that weak lensing convergence bispectrum tomography does not yield competitive constraints on f_{NL} . The goal of our paper is to revisit the feasibility of this approach and investigate whether the situation has improved given the significant advancements in both modeling and observations over the past two decades. Our work builds upon Ref. [29] in several ways. First, while Ref. [29] used a tree-level perturbation theory model for the local PNG contribution to the lensing bispectrum, we utilize a recently developed nonperturbative model for the matter bispectrum based on the LSS consistency relations [33,34] that has been validated deep into the nonlinear regime using N -body simulations [35,36]. This model enables us to include nonlinear modes in our analysis, leveraging the unprecedented depth of ongoing and upcoming weak lensing experiments. Second, we use realistic galaxy source distributions and number densities expected for an LSST/*Euclid*-like survey to directly forecast the constraining power of Stage-IV shear experiments. Finally, we forecast the constraining power of CMB lensing bispectrum measurements using lensing reconstruction noise properties expected from the imminent Simons Observatory [37] and future CMB-S4 [38] experiments. To our knowledge, this is the first forecast for constraining PNG using the CMB lensing convergence bispectrum, which has previously been shown to be a promising probe of cosmology [39].

The remainder of the paper is organized as follows. In Sec. II, we provide theoretical background on local PNG and weak lensing and present expressions for the convergence power spectra and bispectra, as well as their covariances. In Sec. III, we discuss the forecast setup, before presenting the corresponding results in Sec. IV. Section V summarizes our conclusions and the Appendix contains a discussion of the impact of nonlinear effects on our forecasts.

Conventions: Throughout this paper we work in natural units, $c = 1$. We assume a fiducial spatially flat Λ CDM cosmology based on the *Planck* 2018 results [40] with $\Omega_m = 0.311$, $\Omega_b = 0.0493$, $h = 0.674$, $n_s = 0.965$, $\sigma_8 = 0.811$, and $\tau = 0.054$. We assume three species of massless neutrinos.

II. THEORETICAL BACKGROUND

To forecast how well the squeezed lensing bispectrum can constrain f_{NL} , we require a model for the lensing

convergence bispectrum (and its covariance) in the presence of local PNG. We derive these results in this section. We first derive a nonperturbative expression for the unequal-time squeezed 3D matter bispectrum in the presence of f_{NL} based on Ref. [36]. We then introduce the weak lensing convergence field and compute expressions for the convergence power spectrum and bispectrum and their covariances in terms of the 3D matter power spectrum and bispectrum. Some of the results of this section are standard in the lensing literature [41–44]; nevertheless, we include them here both for the sake of completeness and to highlight the assumptions underlying our forecasts.

A. Conventions

We first establish some notation. For an overdensity field $\delta(\mathbf{q})$, the 3D power spectrum is defined by

$$\langle \delta(\mathbf{k})\delta(\mathbf{k}') \rangle \equiv (2\pi)^3 \delta^D(\mathbf{k} + \mathbf{k}') P(k, z_k), \quad (1)$$

where we have explicitly included the time-dependence of $\delta(\mathbf{k})$. Similarly, the 3D bispectrum is defined by

$$\langle \delta(\mathbf{q})\delta(\mathbf{k})\delta(\mathbf{p}) \rangle \equiv (2\pi)^3 \delta^D(\mathbf{q} + \mathbf{k} + \mathbf{p}) B(\mathbf{q}, \mathbf{k}, z_q, z_k, z_p). \quad (2)$$

Here, we are interested in the squeezed limit of the bispectrum, where one of the modes is much smaller than the other two, i.e., $q \ll k \simeq p$. In this limit, the contributions to the lensing bispectrum satisfy $z_k \simeq z_p$; therefore, we use the notation $B(\mathbf{q}, \mathbf{k}, z_q, z_k)$ to indicate the unequal-time bispectrum (fixing $z_p = z_k$). Finally, when working with correlators of the lensing convergence, we will often specify the time-dependence implicitly in terms of the comoving distance, χ , to redshift z .

B. Squeezed matter bispectrum in the presence of local primordial non-Gaussianity

To derive an expression for squeezed configurations of the lensing bispectrum, we first require a model for the unequal-time squeezed matter bispectrum in the presence of local PNG. Our derivation follows that of [36] (see also [45,46]), but is generalized to include unequal-time correlations between the long and short modes.

Local PNG is parametrized by a primordial gravitational potential on subhorizon scales given by [e.g., [47,48]]

$$\Phi(\mathbf{x}) = \Phi_G(\mathbf{x}) + f_{\text{NL}}(\Phi_G^2(\mathbf{x}) - \langle \Phi_G^2 \rangle), \quad (3)$$

where $\Phi_G(\mathbf{x})$ is a Gaussian random field. To determine the squeezed bispectrum we evaluate the correlator $\langle \delta(\mathbf{q}) P(k, \chi_k | \Phi_L) \rangle$ where $\delta(\mathbf{q})$ is the soft mode density field and $P(k, \chi_k | \Phi_L)$ is the locally measured small-scale power spectrum in the presence of a background long-wavelength potential Φ_L [45]. This can be expanded as

$$P(k, \chi_k | \Phi_L) = P(k, \chi_k | 0) + \int d\mathbf{q}' \frac{\partial P(k, \chi_k)}{\partial \Phi_L(\mathbf{q}')} \Phi_L(\mathbf{q}') + \dots, \quad (4)$$

where we are assuming that $q' \equiv |\mathbf{q}'| \ll k$, such that we can treat the long mode as a background, in the presence of which the power spectrum is evaluated. This induces a coupling between the hard mode power spectrum, $P(k, \chi_k)$, and the soft mode density field, $\delta(\mathbf{q})$. The resulting squeezed limit bispectrum is given by

$$\begin{aligned} B(\mathbf{q}, \mathbf{k}, \chi_q, \chi_k) &= \int d\mathbf{q}' \frac{\partial P(k, \chi_k)}{\partial \Phi_L(\mathbf{q}')} \langle \delta(\mathbf{q}) \Phi_L(\mathbf{q}') \rangle \\ &= \frac{\partial P(k, \chi_k)}{\partial \Phi_L(\mathbf{q})} \frac{P(q, \chi_q)}{\alpha(q, \chi_q)}. \end{aligned} \quad (5)$$

Here, we have used Poisson's equation to relate the long-wavelength density and potential fields, $\delta(\mathbf{q}) = \alpha(q, \chi_q) \Phi_L(\mathbf{q})$, with

$$\alpha(q, \chi_q) \equiv \frac{2 q^2 T(q) D_{\text{md}}(\chi_q)}{3 \Omega_m H_0^2}, \quad (6)$$

where Ω_m and H_0 are the matter density and expansion rate today, $T(q)$ is the transfer function (normalized to unity for $q \rightarrow 0$), and $D_{\text{md}}(z)$ is the growth factor, normalized to $a(z) = 1/(1+z)$ in the matter-dominated era.

We evaluate the potential derivative using the separate Universe formalism leading to the following expression for the primordial contribution to the late-time matter bispectrum in the squeezed limit [36,46]:

$$\begin{aligned} B_{\text{prim}}(\mathbf{q}, \mathbf{k}, \chi_q, \chi_k) &= \frac{6 f_{\text{NL}} \Omega_m H_0^2}{D_{\text{md}}(\chi_q)} \frac{\partial P(k, \chi_k)}{\partial \log \sigma_8^2} \\ &\times \frac{P(q, \chi_q)}{q^2 T(q)} + \mathcal{O}(f_{\text{NL}}^2). \end{aligned} \quad (7)$$

Notice that the soft-mode dependence of the primordial contribution to the squeezed matter bispectrum scales as $P(q, \chi_q)/q^2$; however, the LSS consistency relations [33,34] ensure that, in the absence of local PNG and equivalence-principle-violating physics, the ratio $B(\mathbf{q}, \mathbf{k}, \chi_q, \chi_k)/P(q, \chi_q)$ has no $1/q$ and $1/q^2$ poles in the squeezed limit $q/k \rightarrow 0$ (see also [35]).² It is worth stressing that this statement about the lack of poles, in the absence of local PNG and equivalence-principle-violation, is robust: it holds even if the high momentum (k) modes are in the nonlinear regime, and even if they are affected by baryonic feedback

²In the unequal-time limit in the hard modes (i.e., $\chi_{k_1} \neq \chi_{k_2}$), the squeezed bispectrum has a term proportional to $(D_{\text{md}}(\chi_{k_1}) - D_{\text{md}}(\chi_{k_2}))P(q, \chi_q)/q$; however, this contribution vanishes in the limit $\chi_{k_2} \rightarrow \chi_{k_3}$, which is assumed here.

processes [49,50]. As such, we can split the late-time squeezed matter bispectrum into a primordial contribution and a gravitational contribution $B = B_{\text{prim}} + B_{\text{grav}}$, where the primordial contribution is, up to $\mathcal{O}(f_{\text{NL}}^2)$, described by Eq. (7).

A potential challenge in calculating the primordial contribution to the squeezed matter bispectrum using Eq. (7) is that the response function, $\partial \log(P(k, z))/\partial \log(\sigma_8^2)$, needs to be computed in the nonlinear regime. Ref. [36] estimated the response function using HALOFIT [51] and the QUIJOTE simulations [52], and found that both methods provide unbiased constraints on f_{NL} from the 3D matter bispectrum up to $k_{\text{max}} = 0.7 \text{ h/Mpc}$ for redshifts $z = 0$ and $z = 0.97$. Therefore, in this work, we compute the logarithmic derivative using HALOFIT. A more comprehensive analysis should include the impact of baryons on the response function; nevertheless, previous studies have found that these corrections are relatively small [53]. If, instead, we wanted to model the galaxy bispectrum, then the nonlinear response function would likely be much more difficult to model.

The gravitational term B_{grav} is difficult to model beyond perturbative scales; nevertheless, as shown in [35], the squeezed bispectrum is well-described by a power series

$$\begin{aligned} B_{\text{grav}}(\mathbf{q}, \mathbf{k}, \chi_q, \chi_k) &= \sum_{n=0}^{\infty} a_n(k, \theta, \chi_k) \left(\frac{q}{k}\right)^n \\ &\times P(q, \chi_q) P(k, \chi_k), \end{aligned} \quad (8)$$

where the a_i 's are coefficients characterizing the response of small-scale matter clustering to a long-wavelength mode and θ is the angle between $\mathbf{q}$ and $\mathbf{k}$. As described in Refs. [35,36], by angular averaging over all available short modes that satisfy the triangle inequality, the odd-order coefficients vanish, and the remaining coefficients become independent of the angle θ . In practice, one truncates the series at some finite n . By comparing to N -body simulations, Ref. [36] found that truncating at $n = 2$ is sufficient for a wide range of soft modes. The gravitational contribution to the matter bispectrum is then

$$\begin{aligned} B_{\text{grav}}(\mathbf{q}, \mathbf{k}, z_q, z_k) &= a_0(k, \chi_k) P(k) P(q) \\ &+ a_2(k, \chi_k) \frac{q^2}{k^2} P(k) P(q), \end{aligned} \quad (9)$$

which carries two scale- and redshift-dependent “nuisance” parameters, $a_0(k, \chi_k)$ and $a_2(k, \chi_k)$. In this work, we will ignore the $a_2(k, \chi_k)$ contribution because it is subdominant and largely uncorrelated with f_{NL} [36].³ Furthermore, for our fiducial forecasts, we will assume perfect knowledge of

³As shown in Appendix B of [36], $a_2(k, \chi_k) = 0$ is a valid assumption for a wide range of scale cuts.

$a_0(k, \chi_k)$ via the angular averaged bispectrum consistency condition from [54,55],

$$a_0(k, \chi_k) = 1 + \frac{13}{21} \frac{\partial \log P(k, \chi_k)}{\partial \log D_{\text{md}}(\chi_k)} - \frac{1}{3} \frac{\partial \log P(k, \chi_k)}{\partial \log k}. \quad (10)$$

Although Eq. (10) is nonperturbative and can therefore be applied to nonlinear scales, it is expected to break down at small scales and low redshifts due to baryonic effects and departures from an Einstein–de Sitter universe [54,55]. Furthermore, even if Eq. (10) is valid over the scales and redshifts considered in this work, modeling the derivatives in Eq. (10) can be challenging. Consequently, we also consider forecasts for a more pessimistic scenario in which $a_0(k, \chi_k)$ is a free amplitude that we marginalize over, as was done in Ref. [36].

C. Weak gravitational lensing

In this section, we introduce the weak lensing convergence field and derive theoretical predictions for the lensing convergence power spectrum and bispectrum and their associated covariances. For a detailed treatment of weak lensing see, e.g., [41–44].

Assuming the Born approximation and working at linear order in the matter density fluctuation, the convergence field $\kappa^{(i)}(\hat{n})$ is a weighted projection of the matter density field:

$$\kappa^{(i)}(\hat{n}) = \int_0^{\chi_s} d\chi W^{(i)}(\chi) \delta_m(\chi \hat{n}, \chi), \quad (11)$$

where $\hat{n}$ is a unit vector, χ_s is the comoving distance to the photon source, and $W^{(i)}(\chi)$ is the projection kernel. The exact form of $W^{(i)}(\chi)$ is determined by the specifics of the lensing source. For CMB lensing, the kernel is

$$W^{\text{CMB}}(\chi) \equiv \frac{3H_0^2 \Omega_m \chi}{2a(\chi)} \left(\frac{\chi_* - \chi}{\chi_*} \right), \quad (12)$$

where χ_* is the comoving distance to the last-scattering surface at $z_* \simeq 1090$, and $a(\chi)$ is the scale factor. For cosmic shear, the kernel is

$$W^{\text{gs},(i)}(\chi) = \frac{3H_0^2 \Omega_m \chi}{2a(\chi)} \int_\chi^\infty d\chi' p_s^{(i)}(\chi') \left(\frac{\chi' - \chi}{\chi'} \right), \quad (13)$$

where $p_s^{(i)}(\chi)$ is the redshift distribution of source galaxies in the i th tomographic bin, satisfying the normalization condition $\int d\chi' p_s^{(i)}(\chi') = 1$.

It is convenient to expand the convergence field in spherical harmonics,

$$\kappa^{(i)}(\hat{n}) = \sum_{\ell, m} \kappa_{\ell m}^{(i)} Y_{\ell m}(\hat{n}). \quad (14)$$

In the following sections, we will derive expressions for the auto- and cross-power spectra and bispectra of arbitrary convergence fields $\kappa_{\ell m}^{(i)}$ in terms of the 3D matter power spectrum and bispectrum.

In practice, galaxy weak lensing surveys measure cosmic shear instead of convergence; however, the convergence field can be reconstructed from shear measurements [56–62]. The reconstructed convergence field is a (somewhat) biased estimate of the underlying convergence field; therefore, in a real analysis, one would likely directly model the shear bispectrum instead of the convergence bispectrum. For CMB lensing, the convergence field itself can be directly reconstructed from the observed CMB temperature and polarization anisotropies (e.g., [63]).

1. Convergence power spectrum

The angular power spectrum between two convergence fields, $C_{\ell}^{(ij)}$, is defined by $\langle \kappa_{\ell m}^{(i)} \kappa_{\ell' m'}^{*(j)} \rangle = \delta_{\ell \ell'}^K \delta_{mm'}^K C_{\ell}^{(ij)}$, assuming statistical isotropy. Using Eqs. (11) and (14), we can express $C_{\ell}^{(ij)}$ as

$$C_{\ell}^{(ij)} = \frac{2}{\pi} \int d\chi_1 W^{(i)}(\chi_1) \int d\chi_2 W^{(j)}(\chi_2) \times \int dk k^2 P(k, \chi_1, \chi_2) j_{\ell}(k\chi_1) j_{\ell}(k\chi_2). \quad (15)$$

At high ℓ , we can use the Limber approximation [64–66] to replace the highly oscillatory spherical Bessel function by a Dirac delta function,

$$j_{\ell}(k\chi) \simeq \sqrt{\frac{\pi}{2\ell+1}} \delta^D\left(\ell + \frac{1}{2} - k\chi\right), \quad (16)$$

leading to

$$C_{\ell \gg 1}^{(ij)} \simeq \int \frac{d\chi}{\chi^2} W^{(i)}(\chi) W^{(j)}(\chi) P\left(k = \frac{\ell + \frac{1}{2}}{\chi}, \chi\right). \quad (17)$$

In this work, we compute the angular power spectrum using the exact expression via the FFTLog [67–70] algorithm for $\ell < 50$ and employ the Limber approximation for $\ell \geq 50$.

Assuming a fractional sky coverage f_{sky} , the covariance between $C_{\ell}^{(ij)}$ and $C_{\ell'}^{(mn)}$ is [29]

$$\text{Cov}[C_{\ell}^{(ij)}, C_{\ell'}^{(mn)}] = \frac{\delta_{\ell \ell'}^K}{f_{\text{sky}}(2\ell+1)} (\tilde{C}_{\ell}^{(im)} \tilde{C}_{\ell'}^{(jn)} + \tilde{C}_{\ell}^{(in)} \tilde{C}_{\ell'}^{(jm)}), \quad (18)$$

where we have neglected the connected non-Gaussian contribution and the super-sample covariance [71–76]. The connected non-Gaussian contribution is expected to be subdominant in Stage-IV convergence power spectra due to the suppression of non-Gaussianities in lensing,

which projects quantities along the line-of-sight [71,76]. Conversely, as shown in [76], the super-sample covariance can have a significant impact on the convergence power spectra for Stage-IV shear surveys. Nevertheless, we ignore the non-Gaussian covariance so that we can determine the most optimistic forecasts for a given survey (and thus obtain an upper bound on the utility of our method). Note that in Eq. (18), we write $\tilde{C}_\ell^{(ij)}$ to emphasize that this is the *observed* angular power spectrum, including the noise contribution as discussed in Sec. III A.

$$\langle \kappa_{\ell_1 m_1}^{(1)} \kappa_{\ell_2 m_2}^{(2)} \kappa_{\ell_3 m_3}^{(3)} \rangle = \frac{8}{\pi^3} \mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} \int dr r^2 \left[\prod_{i=1}^3 d\chi_i dk_i k_i^2 W^{(i)}(\chi_i) j_{\ell_i}(k_i \chi_i) j_{\ell_i}(k_i r) \right] B(k_1, k_2, k_3, \chi_1, \chi_2, \chi_3), \quad (20)$$

where we have introduced the Gaunt factor, $\mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} \equiv \sqrt{\frac{(2\ell_1+1)(2\ell_2+1)(2\ell_3+1)}{4\pi}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix}$. Again asserting isotropy, it is convenient to define the reduced bispectrum, $\langle \kappa_{\ell_1 m_1}^{(i)} \kappa_{\ell_2 m_2}^{(j)} \kappa_{\ell_3 m_3}^{(k)} \rangle \equiv \mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} b_{\ell_1 \ell_2 \ell_3}^{(ijk)}$, such that the reduced bispectrum and angular bispectrum are related by

$$B_{\ell_1 \ell_2 \ell_3}^{(ijk)} = \sqrt{\frac{(2\ell_1+1)(2\ell_2+1)(2\ell_3+1)}{4\pi}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix} b_{\ell_1 \ell_2 \ell_3}^{(ijk)} \equiv h_{\ell_1 \ell_2 \ell_3} b_{\ell_1 \ell_2 \ell_3}^{(ijk)}. \quad (21)$$

As in Eq. (17), in the squeezed limit, which corresponds to $\ell_1 \ll \ell_2 \simeq \ell_3$, we can use the Limber approximation for the spherical Bessel functions $j_{\ell_2}(k_2 \chi_2)$ and $j_{\ell_3}(k_3 \chi_3)$. This allows the reduced bispectrum to be written as

$$b_{\ell_1 \ell_2 \ell_3}^{(ijk)} = \frac{2}{\pi} \int d\chi_1 d\chi_2 \left(\frac{W^{(i)}(\chi_1) W^{(j)}(\chi_2) W^{(k)}(\chi_2)}{\chi_2^2} \right) \times \int dq q^2 j_{\ell_1}(q\chi_1) j_{\ell_1}(q\chi_2) \times B\left(q, \frac{\ell_2 + \frac{1}{2}}{\chi_2}, \frac{\ell_3 + \frac{1}{2}}{\chi_2}, \chi_1, \chi_2, \chi_2\right). \quad (22)$$

Notice that the assumption that $\chi_2 \simeq \chi_3$ is exact in the Limber approximation. If we also assume the Limber approximation for ℓ_1 , then

$$b_{\ell_1 \ell_2 \ell_3}^{(ijk)} \simeq \int d\chi \frac{W^{(i)}(\chi) W^{(j)}(\chi) W^{(k)}(\chi)}{\chi^4} \times B\left(\frac{\ell_1 + \frac{1}{2}}{\chi}, \frac{\ell_2 + \frac{1}{2}}{\chi}, \frac{\ell_3 + \frac{1}{2}}{\chi}, \chi\right). \quad (23)$$

Here, we compute the squeezed angular bispectrum using Eq. (22) via the FFTLog algorithm for $\ell_1 < 50$ and use the Limber approximation in Eq. (23) for $\ell_1 \geq 50$.

2. Convergence bispectrum

The angular bispectrum $B_{\ell_1 \ell_2 \ell_3}^{(ijk)}$ is defined by [e.g., [77]]

$$\langle \kappa_{\ell_1 m_1}^{(i)} \kappa_{\ell_2 m_2}^{(j)} \kappa_{\ell_3 m_3}^{(k)} \rangle \equiv \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix} B_{\ell_1 \ell_2 \ell_3}^{(ijk)}, \quad (19)$$

where $\begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix}$ is the Wigner-3j symbol. Using Eqs. (11) and (14), the convergence three-point function can be expressed as

Finally, assuming $\ell_1 \leq \ell_2 \leq \ell_3$, the Gaussian covariance between the reduced bispectra $b_{\ell_1 \ell_2 \ell_3}^{(abc)}$ and $b_{\ell'_1 \ell'_2 \ell'_3}^{(ijk)}$ is

$$\text{Cov}[b_{\ell_1 \ell_2 \ell_3}^{(abc)}, b_{\ell'_1 \ell'_2 \ell'_3}^{(ijk)}] \simeq h_{\ell_1 \ell_2 \ell_3}^{-1} h_{\ell'_1 \ell'_2 \ell'_3}^{-1} \frac{\Delta(\ell_1, \ell_2, \ell_3)}{f_{\text{sky}}} \times \tilde{C}_{\ell_1}^{(ai)} \tilde{C}_{\ell_2}^{(bj)} \tilde{C}_{\ell_3}^{(ck)} \delta_{\ell_1 \ell'_1}^K \delta_{\ell_2 \ell'_2}^K \delta_{\ell_3 \ell'_3}^K, \quad (24)$$

where $\Delta(\ell_1, \ell_2, \ell_3)$ is a symmetry factor that is equal to six for equilateral triangles, two for isosceles triangles, and one otherwise. As in Eq. (18), Eq. (24) includes only the Gaussian contribution to the covariance. Although non-Gaussian terms can have a significant impact on the lensing convergence bispectrum covariance [73,78–82], especially for squeezed configurations, we neglect them in order to forecast the most optimistic possible constraints on f_{NL} achievable with the presented method.

III. FORECAST SETUP

A. Survey specifications

1. CMB lensing experiments

For the CMB lensing analysis, we consider two CMB experiments: Simons Observatory [37] and CMB-S4 [38,83]. We assume a fractional sky coverage $f_{\text{sky}} = 0.45$ for both surveys, for simplicity. The observed angular power spectrum of the CMB lensing convergence, used to compute the bispectrum covariance in Eq. (24), is

$$\tilde{C}_\ell = C_\ell^{\text{CMB}} + N_\ell, \quad (25)$$

where N_ℓ is the CMB lensing reconstruction noise. Here, we assume that the CMB lensing convergence is reconstructed using an iterative estimator [84,85] and model the reconstruction noise using the iterative noise curves




FIG. 1. Left: angular auto-power spectra of the CMB lensing convergence and cosmic shear convergence. The shaded region denotes the 1σ error bar estimated from the Gaussian covariance with $f_{\text{sky}} = 0.45$, as described in Sec. III. The CMB lensing results (orange and red) include lensing reconstruction noise for Simons Observatory and CMB-S4, while the shear results (blue) include shape noise and assume a single tomographic bin; the noise power spectra are shown in dash-dotted. The black dotted curves indicate linear-theory predictions. The cosmic shear power spectrum is significantly more sensitive to nonlinear structure formation than the CMB lensing power spectrum, but also has much higher signal-to-noise at small scales (large ℓ). The power spectra are computed using FFTLog for $\ell < 50$ and using the Limber approximation for $\ell \geq 50$. Right: assumed redshift source distribution for the cosmic shear forecasts presented in this work based on an LSST/Euclid-like survey. The full sample (solid) is normalized to unit integral and the three tomographic bins (dashed) are normalized to $1/3$ each. For comparison, we include the CMB lensing kernel $W^{\text{KCB}}(z)/H(z)$ computed using Eq. (12) with an arbitrary normalization.

from [86]. The convergence power spectrum C_ℓ^{KCB} is computed using Eqs. (15) and (17) where we compute the matter power spectrum using HALOFIT to model contributions from nonlinear structure formation.⁴

The left panel of Fig. 1 shows the CMB lensing convergence power spectrum and its covariance, including the cosmic variance and reconstruction noise expected for Simons Observatory and CMB-S4. The CMB lensing convergence power spectrum is cosmic-variance-limited (on the observed sky fraction) up to $\ell \simeq 350$ (700) for Simons Observatory (CMB-S4).

2. Cosmic shear experiments

For the cosmic shear analysis, we consider a generic Stage IV photometric galaxy survey with specifications similar to those expected of LSST [4] and Euclid [87]. As for the CMB, we assume that the cosmic shear surveys have a fractional sky coverage $f_{\text{sky}} = 0.45$. When cross-correlating cosmic shear and CMB lensing, we assume the surveys fully overlap.

We model the *true* source galaxy redshift distribution as

$$p_s^{\text{true}}(z) \propto z^2 \exp[-(z/z_0)^\alpha], \quad (26)$$

⁴We use the linear power spectrum for the soft mode when computing the non-Limber integrals since the FFTLog algorithm requires that the time dependence of the integrand factorizes. This has negligible impact on our results because these scales are well described by linear theory.

with $z_0 = 0.28$ and $\alpha = 0.9$ as specified in the LSST Dark Energy Science Collaboration Science Requirements Document [88]. To account for photometric-redshift uncertainties, we assume that the probability distribution for the observed photometric redshift, z_{ph} , given the true galaxy redshift, z , follows a Gaussian distribution,

$$\mathcal{P}(z_{\text{ph}}|z) = \frac{1}{\sqrt{2\pi\sigma^2(z)}} \exp\left[-\frac{1}{2}\left(\frac{z - z_{\text{ph}}}{\sigma(z)}\right)^2\right], \quad (27)$$

with uncertainty $\sigma(z) = 0.05(1+z)$ [88]. The source distribution in a tomographic bin $z_{\text{min}}^{(i)} \leq z \leq z_{\text{max}}^{(i)}$ is then

$$p_s^{(i)}(z) = p_s^{\text{true}}(z) \int_{z_{\text{min}}^{(i)}}^{z_{\text{max}}^{(i)}} dz' \mathcal{P}(z'|z). \quad (28)$$

For our fiducial forecast, we divide the source distribution into three tomographic bins with an equal number of galaxies in each bin. We investigate the impact of varying the number of tomographic bins on our parameter constraints in Sec. IV B.

The observed angular power spectrum of the cosmic shear convergence used to compute the bispectrum covariance in Eq. (24) is

$$\tilde{C}_\ell^{(ij)} = C_\ell^{g,(ij)} + \delta_{ij}^K \frac{\sigma_e^2}{\bar{n}_g^{(i)}}, \quad (29)$$

where we have included the shape noise contribution arising from the intrinsic ellipticities of galaxies. We set

the effective number density of galaxies to $\bar{n}_g = 31 \text{ arcmin}^{-2}$ and assume an intrinsic rms ellipticity $\sigma_e = 0.26$ [89]. When dividing the galaxy source sample into N_{tomo} tomographic bins, the effective number density is $\bar{n}_g^{(i)} = \bar{n}_g / N_{\text{tomo}}$.

The left panel of Fig. 1 includes the cosmic shear convergence power spectrum assuming $N_{\text{tomo}} = 1$ as well as the noise contribution due to shape noise. The cosmic shear power spectrum is cosmic-variance-limited (on the observed sky fraction) up to $\ell \simeq 1800$ assuming $N_{\text{tomo}} = 1$ and the survey specifications used in this work. The right panel of Fig. 1 shows the assumed galaxy redshift source distribution in a single tomographic bin and in three tomographic bins, as well as the CMB lensing kernel.

B. Fisher matrix

To estimate how well lensing convergence bispectra can constrain f_{NL} , we adopt the Fisher matrix formalism. For our fiducial forecasts, we assume that the auto- and cross-bispectra of the convergence fields are the only observables and that f_{NL} is the only free parameter.⁵ We also assume a Gaussian likelihood for the observed convergence bispectra. The Fisher matrix is then given by

$$\mathcal{F} = \sum_{\substack{ijk \\ abc}} \sum_{\ell_1 \ell_2 \ell_3} \frac{\partial b_{\ell_1 \ell_2 \ell_3}^{(ijk)}}{\partial f_{\text{NL}}} \text{Cov}^{-1} \left[b_{\ell_1 \ell_2 \ell_3}^{(ijk)}, b_{\ell_1 \ell_2 \ell_3}^{(abc)} \right] \frac{\partial b_{\ell_1 \ell_2 \ell_3}^{(abc)}}{\partial f_{\text{NL}}}, \quad (30)$$

where $\partial b_{\ell_1 \ell_2 \ell_3}^{(ijk)} / \partial f_{\text{NL}}$ is computed using Eqs. (7), (22), and (23), and the covariance is computed using Eq. (24). The Cramér–Rao bound guarantees that the minimum variance of an unbiased estimator of a given parameter is equal to the inverse Fisher matrix element of the associated parameter, hence $\sigma_{f_{\text{NL}}} \geq 1/\sqrt{\mathcal{F}}$.

The sum over i, j, k and a, b, c in Eq. (30) runs over all tomographic bins included in the analysis. For example, when considering CMB lensing and cosmic shear cross-correlations, the i, j, k , and a, b, c range from 1 to $N_{\text{tomo}} + 1$. The sum over ℓ_1, ℓ_2 , and ℓ_3 runs over all possible multipoles satisfying the following criteria:

- (i) $\ell_{\text{min}}^{\text{soft}} \leq \ell_1 \leq \ell_{\text{max}}^{\text{soft}}$;
- (ii) $\ell_{\text{min}}^{\text{hard}} \leq \ell_2 \leq \ell_{\text{max}}^{\text{hard}}$;
- (iii) $|\ell_i - \ell_j| \leq \ell_k \leq \ell_i + \ell_j$;
- (iv) $\ell_1 + \ell_2 + \ell_3 = \text{even}$.

The first two criteria ensure that the triangles are sufficiently squeezed such that the bispectrum model in Eq. (7) applies.⁶ The last two criteria arise from momentum conservation and parity, respectively. For our forecasts, we fix the maximum soft multipole $\ell_{\text{max}}^{\text{soft}} = 100$ and the minimum hard multipole $\ell_{\text{min}}^{\text{hard}} = 150$.⁷ In theory, the minimum soft multipole $\ell_{\text{min}}^{\text{soft}}$ is determined by the largest scales observed in the survey; in practice, however, the largest scales can be plagued by observational and theoretical systematics, including foreground contamination (note that high- ℓ CMB temperature and polarization foregrounds lead to low- ℓ biases in the reconstructed lensing convergence field) [90–93], relativistic corrections [94], and non-Gaussianity of the likelihood [95,96]. Consequently, we will consider $\ell_1^{\text{min}} = 2, 10$, and 20. We vary the maximum hard multipole $\ell_{\text{max}}^{\text{hard}}$ to determine the scales at which our constraints saturate due to noise.

In addition to the forecasts assuming f_{NL} is the only free parameter, we present forecasts where we marginalize over the leading-order gravitational contribution to the matter bispectrum, $a_0(k, \chi_k)$ [see Eq. (9)]. Directly marginalizing over $a_0(k, \chi_k)$ is challenging because $a_0(k, \chi_k)$ depends on scale and redshift, both of which enter the bispectrum evaluations in Eqs. (22) and (23). Therefore, we assume that $a_0(k, \chi_k)$ varies slowly over the integration volume and can thus be approximated by a single coefficient $\bar{a}_0^{(i)}$, where the index (i) is used to indicate that this parameter depends on the redshift kernel of the bispectrum considered.⁸ For simplicity, we present these marginalized forecasts using only a single convergence field, hence the only free parameters are $\bar{a}_0$ and f_{NL} . The 2×2 Fisher matrix is

⁶In principle, different scale cuts could be imposed for each tomographic bin, as well as for CMB lensing versus cosmic shear, because these measurements probe different physical scales and redshifts and are sensitive to different systematics. We do not account for this in our forecasts because quantifying the exact range of multipoles for which our bispectrum model applies for a given convergence field would require simulations. Nevertheless, we note that this approach of varying scale cuts could be useful in practice.

⁷This choice of multipoles is somewhat optimistic and a precise verification of the range of scales for which our model is valid would require analyzing simulations, similar to what was done for the 3D matter bispectrum in [36]. Nevertheless, since most of the f_{NL} information is coming from the lowest ℓ_1 bins, our results are largely insensitive to the precise choice of $\ell_{\text{max}}^{\text{soft}}$.

⁸We find that approximating $a_0(k, \chi_k)$ as a constant evaluated at the peak of the lensing kernel can bias the bispectrum by at most $\simeq 10\%$ compared to full numerical integration using Eq. (10).

⁵Note that this differs from the analysis in Ref. [36], which used a joint likelihood in the 3D matter power spectrum and bispectrum to take advantage of the sample variance cancellation associated with the significant correlation between the squeezed bispectrum $B(\mathbf{q}, \mathbf{k})$ and the soft mode power spectrum $P(q)$. It would be interesting to consider whether a similar cancellation applies to the angular power spectrum C_{ℓ_1} and the squeezed angular bispectrum $b_{\ell_1 \ell_2 \ell_3}$; however, we leave this to future work because it would require including the non-Gaussian contributions to the power spectrum and bispectrum covariances, since the power spectrum and bispectrum are uncorrelated in the Gaussian limit. Additionally, $b_{\ell_1 \ell_2 \ell_3}$ and C_{ℓ_1} have different projection integrals, which likely reduces their correlation.




FIG. 2. Forecasted error on f_{NL} from the nonlinear squeezed bispectrum of CMB lensing (top) and cosmic shear (bottom) using a nonperturbative bispectrum model based on the cosmological consistency relations that is independent of the non-Gaussian bias, b_ϕ . The error is shown a function of the maximum multipole ℓ_{max} for several values of the minimum multipole $\ell_{\text{min}} = 2, 10$, and 20 . The forecasts are sensitive to the minimum multipole ℓ_{min} . Whereas the CMB lensing forecasts saturate around $\ell_{\text{max}} = 1000$ due to lensing reconstruction noise, the shear forecasts begin to plateau around $\ell_{\text{max}} = 3000$ due to shape noise. The cosmic shear forecasts improve if we divide the source galaxy sample into two tomographic bins, but there is little to gain with $N_{\text{tomo}} \geq 3$. The cosmic variance error is shown by the gray bands. The three gray bands in the shear forecasts correspond to the number of tomographic bins. All forecasts assume fractional sky coverage $f_{\text{sky}} = 0.45$.

$$\mathcal{F}_{\alpha\beta} = \sum_{\ell_1\ell_2\ell_3} \frac{\partial b_{\ell_1\ell_2\ell_3}}{\partial \theta_\alpha} \text{Cov}^{-1}[b_{\ell_1\ell_2\ell_3}, b_{\ell_1\ell_2\ell_3}] \frac{\partial b_{\ell_1\ell_2\ell_3}}{\partial \theta_\beta}, \quad (31)$$

where $\theta_\alpha \in \{f_{\text{NL}}, \bar{a}_0\}$. The Fisher error on f_{NL} is then $\sigma_{f_{\text{NL}}} = \sqrt{\mathcal{F}_{f_{\text{NL}}f_{\text{NL}}}^{-1}}$.

Finally, since the scale cuts used in this work typically contain $\mathcal{O}(10^6)$ triangles, we adopt a binning strategy based on Ref. [29]. In particular, we bin the first two multipoles with bin widths $\Delta\ell_1 = 5$ and $\Delta\ell_2 = 15$ and rescale the covariance by a factor of $(\Delta\ell_1\Delta\ell_2)^{-1}$. We let ℓ_3 range over all available multipoles subject to the triangle inequality, the parity constraint, and the ordering $\ell_2 \leq \ell_3$. We have verified that this binning approximation has negligible impact on our forecasts.

IV. RESULTS

A. CMB lensing

The top panels of Fig. 2 show the expected error on f_{NL} from the CMB lensing bispectrum as a function of the maximum hard multipole for three different choices of the minimum soft multipole. The shaded region denotes

the cosmic variance error, which assumes perfect knowledge of the lensing potential (i.e., $N_\ell = 0$). The forecasts for Simons Observatory (CMB-S4) saturate by $\ell_{\text{max}} \simeq 800$ (1200). For the fiducial Simons Observatory analysis with $\ell_{\text{min}} = 10$ and $\ell_{\text{max}} = 1400$, the error on f_{NL} is $\sigma_{f_{\text{NL}}} = 175$; with these scale cuts, CMB-S4 can improve the constraint by almost a factor of two, yielding $\sigma_{f_{\text{NL}}} = 95$. Notably, the forecasted error on f_{NL} is very sensitive to the lowest multipole. Assuming the CMB lensing bispectrum can be reliably reconstructed down to $\ell_{\text{min}} = 2$, then the method presented here constrains f_{NL} with precision $\sigma_{f_{\text{NL}}} = 105$ (56) for Simons Observatory (CMB-S4).

The findings of this section indicate that forthcoming measurements of the CMB lensing bispectrum from Simons Observatory and CMB-S4 may not offer competitive constraints on f_{NL} compared with those obtained from the primary CMB or from LSS constraints based on the scale-dependent bias, though the constraints are on different characteristic scales. Nevertheless, the forecasted error on f_{NL} is still better than the tightest current b_ϕ -independent LSS constraints on f_{NL} [14]. The main limitation of these forecasts is the CMB lensing reconstruction noise, which severely restricts our ability




FIG. 3. As Fig. 2, but showing forecasts for a joint analysis of the CMB lensing and cosmic shear auto- and cross-bispectra assuming three tomographic redshift bins. At low maximum multipoles ($\ell_{\text{max}} \approx 400$), including CMB lensing convergence measurements expected from Simons Observatory (CMB-S4) improves the forecasted error on f_{NL} by $\simeq 10\%$ (20%). At larger maximum multipoles ($\ell_{\text{max}} \approx 2000$), the constraint on f_{NL} is dominated by the shear bispectrum with improvements of only $\simeq 1\%$ (5%) when including the Simons Observatory (CMB-S4) lensing convergence. The two gray bands correspond to the shear and shear $\times$ CMB lensing cosmic variance errors.

to precisely measure highly squeezed configurations of the CMB lensing bispectrum. In the more long-term future, the method presented here could prove to be quite powerful for a CMB experiment that is cosmic-variance-limited up to much smaller scales, such as CMB-HD [97], although one should properly include the non-Gaussian covariance to properly assess how much information is practically available at smaller scales.

B. Cosmic shear

The bottom panels of Fig. 2 show the expected error on f_{NL} from cosmic shear bispectra as a function of ℓ_{max} . Interestingly, even for a single tomographic bin and fixed scale cuts, the cosmic variance error on f_{NL} from the shear bispectrum is approximately 20%–40% smaller than that from the CMB lensing bispectrum. This is because the squeezed limit convergence bispectrum is proportional to an integral over $\partial P(k)/\partial \log(\sigma_8^2)$, whereas the noise is proportional to $P(k)$. As a result, the signal-to-noise is enhanced by the logarithmic derivative $\partial \log P(k)/\partial \log(\sigma_8^2)$, which is more pronounced at low redshifts and small scales.⁹ This is discussed in more detail in the Appendix.

The forecasted constraints on f_{NL} from the shear bispectrum measured from an LSST/*Euclid*-like survey saturate at much higher ℓ_{max} than the forecasts from the CMB lensing bispectrum measured by Simons Observatory or CMB-S4. Assuming a single tomographic bin with $\ell_{\text{min}} = 10$ and $\ell_{\text{max}} = 5000$, the shear

bispectrum constrains f_{NL} with precision $\sigma_{f_{\text{NL}}} = 22$. The forecasted error improves to $\sigma_{f_{\text{NL}}} = 16$ using three tomographic bins, indicating that tomographic information of the source galaxies can significantly improve constraints on f_{NL} . Similar to the findings of Ref. [29], our forecasts do not improve significantly if we include more than three tomographic bins. Taken at face value, the results of this section suggest that upcoming measurements of tomographic lensing bispectra from Stage-IV shear surveys provide a more promising avenue toward constraining f_{NL} than upcoming measurements of the CMB lensing bispectrum, and one that could be vital in confirming detections from the scale-dependent bias method, such as [18]. However, it is important to note that, since the shear bispectrum probes smaller scales and lower redshifts than the CMB lensing bispectrum, the shear forecasts would likely be more sensitive to the non-Gaussian covariance, which we have ignored.

C. CMB lensing $\times$ cosmic shear

Figure 3 shows the forecasted error on f_{NL} for a joint analysis of the CMB lensing convergence and the cosmic shear auto- and cross-bispectra. These forecasts assume that the shear field is measured in three tomographic redshift bins. Assuming $\ell_{\text{max}} = 400$ and $\ell_{\text{min}} = 10$, including CMB lensing measurements from Simons Observatory (CMB-S4) reduces the error on f_{NL} from 94 to 86 (78). The improvement from a joint analysis is much less significant at smaller scales, where the signal-to-noise of the shear bispectrum is significantly larger than that of the CMB lensing convergence bispectrum. Indeed, assuming $\ell_{\text{max}} = 2000$ and $\ell_{\text{min}} = 10$, a joint analysis of CMB-S4 and an LSST/*Euclid*-like shear experiment leads to a less than 5%

⁹For the scales and redshifts considered here, the logarithmic derivative is greater than one; however, the situation can reverse at very small scales and low redshifts (see Fig. 5 in the Appendix).

improvement on the constraint on f_{NL} compared to a shear-only analysis.

Although these results suggest that there is little to gain from a joint analysis of CMB lensing and cosmic shear in comparison to a shear-only analysis, the situation may differ considerably in practice. For instance, a realistic analysis would likely impose different scale cuts for the CMB lensing measurements and the cosmic shear measurements to account for their distinct systematics (with CMB lensing potentially providing easier access to low ℓ , modulo foreground biases or other reconstruction systematics). If one uses only the CMB lensing convergence to measure large-scale modes, then the joint analysis is restricted to 16 of the 64 total bispectra combinations, assuming $N_{\text{tomo}} = 3$. Under these conditions, with $\ell_{\text{min}} = 10$ and $\ell_{\text{max}} = 2000$, the resulting error on f_{NL} is $\sigma_{f_{\text{NL}}} = 32$ for CMB-S4. Whereas this value is worse than the forecasted $\sigma_{f_{\text{NL}}} = 23$ from a joint analysis of all 64 bispectra combinations, it still represents a significant improvement over the $\sigma_{f_{\text{NL}}} = 95$ value obtained for a CMB-S4 lensing convergence bispectrum-only analysis with these scale cuts.

An alternative approach for constraining f_{NL} using CMB lensing and cosmic shear measurements would be to directly remove low redshift contributions to the CMB lensing kernel through cross-correlation with the shear maps, following the procedure described in Ref. [98]. By isolating the high redshift contributions to the CMB lensing kernel, this approach has the potential to alleviate several challenges encountered when constraining f_{NL} using projected lensing maps. First, by removing the low redshift contribution, the effective projection kernel becomes more localized in redshift, protecting the PNG signal from the projection effects encountered when integrating over a broad kernel [31]. Second, since the effective weight function for the bispectrum, $W^3(\chi)/\chi^3$, peaks at small scales [31], constraints on local PNG improve at higher redshifts, which include contributions from larger scales. Finally, the high redshift CMB lensing signal is less sensitive to the non-Gaussian covariance. We leave an investigation of this method to a future work.

D. Marginalizing over gravitational non-Gaussianity

Finally, we consider how our forecasts would change if we no longer assume perfect knowledge of the leading-order gravitational non-Gaussianity contribution, $a_0(k, \chi_k)$ (see Eq. (9)). The left panel of Fig. 4 compares the forecasted error on f_{NL} from the CMB lensing bispectrum with and without marginalization over the leading-order gravitational non-Gaussianity parameter $a_0(k, \chi_k)$ using the marginalization procedure described in Sec. III B. The red curves assume CMB-S4 lensing reconstruction noise and the gray shaded regions denote the cosmic variance error. Marginalizing over gravitational non-Gaussianity degrades the constraining power quite significantly, with the




FIG. 4. Forecasted error on f_{NL} from the CMB lensing bispectrum measured by CMB-S4 with and without marginalization over gravitational non-Gaussianity. The solid result assumes perfect knowledge of the leading-order gravitational non-Gaussianity parameter, $a_0(k, \chi_k)$. The dashed line marginalizes over $a_0(k, \chi_k)$ using the procedure described in Sec. III B. Marginalization over $a_0(k, \chi_k)$ weakens the constraint on f_{NL} by approximately 50%. The two gray shaded regions correspond to the cosmic-variance-limited error on f_{NL} from the CMB lensing bispectrum with and without marginalizing over $a_0(k, \chi_k)$.

forecasted error on f_{NL} increasing from $\sigma_{f_{\text{NL}}} = 95$ to 150, assuming $\ell_{\text{max}} = 1400$. This loss in constraining power arises from the large correlation between the primordial contribution to the matter bispectrum, Eq. (7), and the leading-order gravitational contribution, Eq. (9), which, at equal times, is only broken by the $1/(q^2 T(q))$ term.¹⁰ These findings are in agreement with Ref. [36], which found a sizeable cross-correlation ($r \simeq 0.7$ at $z = 0.97$) between $\bar{a}_0$ and f_{NL} .

Based on these results, it is clear that marginalizing over the gravitational contribution to the squeezed matter bispectrum significantly degrades constraints on f_{NL} using the method presented here. Nevertheless, there are several possible ways to extend our analysis that could improve constraints on $a_0(k, \chi_k)$ and, hence, f_{NL} . First of all, since $a_0(k, \chi_k) \simeq B(q, k)/P(q)P(k)$, a joint analysis of the convergence bispectrum and convergence power spectrum could improve constraints on $a_0(k, \chi_k)$. Reference [36] exploited this sample variance cancellation in 3D and found that it significantly improved constraints on f_{NL} . A precise determination of the impact of sample variance cancellation is beyond the scope of this work because it requires non-Gaussian covariances. Secondly, cross-correlation of CMB lensing and cosmic shear maps could help constrain

¹⁰The k -dependence of the $a_0(k, \chi_k)$ coefficient also breaks this correlation; however, this k -dependence is small and is neglected in our forecasts.

$a_0(k, \chi_k)$, since both the CMB lensing bispectrum and the cosmic shear bispectrum depend on differently weighted integrals over $a_0(k, \chi_k)$.

V. CONCLUSIONS

Constraining the amplitude of local primordial non-Gaussianity, f_{NL} , is a key goal of upcoming cosmological surveys. Here, we have quantified how well squeezed configurations of the CMB lensing convergence and cosmic shear bispectra, as measured by forthcoming experiments, can constrain f_{NL} . Our method has utilized a nonperturbative model for the squeezed bispectrum based on the cosmological consistency relations, which is, in principle, applicable down to very small scales (high ℓ_{max}). In practice, we have found that, even with a joint analysis of the auto- and cross-bispectra of convergence measurements from an LSST/*Euclid*-like experiment and CMB-S4, it will be difficult to constrain $\sigma_{f_{\text{NL}}} \lesssim 10$.

In the modern world, with an abundance of forecasts predicting $\sigma_{f_{\text{NL}}} \lesssim 1$ using Stage-IV surveys (with varying degrees of optimism in their assumptions), it is natural to question the utility of the approach presented here. We emphasize that the primary advantage of using the squeezed lensing bispectrum to constrain f_{NL} is that it actually constrains f_{NL} . This stands in contrast to most existing methods for constraining f_{NL} using LSS, which rely on the scale-dependent bias, and hence require accurate knowledge of galaxy formation physics to obtain a direct constraint on f_{NL} .¹¹ Moreover, the squeezed lensing bispectrum is sensitive to different large-scale systematics compared to alternative methods to constrain f_{NL} (e.g., the galaxy power spectrum); thus, our method can be used as a valuable cross-check for future constraints on f_{NL} using LSS data. Finally, our approach could be extended to include galaxy clustering information. Combining galaxy clustering with CMB lensing and/or cosmic shear has already shown great promise as a probe of f_{NL} using two-point statistics [99]; therefore, a joint analysis of galaxy clustering, cosmic shear, and CMB lensing two-point and three-point functions could offer highly informative results. Although these constraints would no longer be independent of b_ϕ , the potential for sample variance cancellations and forming cross-bispectra with the galaxy density field as the long mode (similar to the discussion in Sec. IV C for CMB lensing) makes this approach particularly promising from a modeling perspective. It is also worth noting that aside from constraining f_{NL} , checking the consistency relation with data constitutes a test of equivalence principle as well. For that purpose, a measurement of the bispectrum involving different galaxy populations would be needed [100].

¹¹In a not entirely unrealistic, but still unrealistic, scenario in which cosmologists find themselves debating $b_\phi f_{\text{NL}}$ vs. f_{NL} , the method outlined here could potentially quell the controversy.

As with all forecasts, our results depend on several modeling assumptions. First of all, we have ignored various systematics that could impact the CMB lensing and/or cosmic shear bispectra, such as post-Born corrections [101,102], reconstruction noise biases [103], intrinsic alignments [104], source clustering [105], and relativistic effects [94]. Furthermore, we omitted baryonic effects in our analysis, despite their significant impact on the correlators of weak lensing convergence across various scales [106]. One may worry that baryons pose a significant challenge to the methodology presented here, which uses measurements from extremely small scales. However, it is important to note that, since the cosmological consistency relations are still satisfied with baryons, our model that marginalizes over gravitational non-Gaussianity remains valid as long as baryonic corrections are included when computing the response function, $\partial \log(P(k, z)) / \partial \log(\sigma_8^2)$. Previous studies have shown that these corrections are relatively small [53].¹² Finally, to obtain an upper bound on our method's utility, we have neglected non-Gaussian contributions to the bispectrum covariance. These contributions could significantly degrade the cosmic shear forecasts at high ℓ_{max} , but we leave this to a future work.

There are several ways to build upon our analysis. An immediate follow-up would be to use simulations to explicitly verify the range of scale cuts for which our bispectrum model is valid and to quantify the impact of the non-Gaussian covariance. Such an analysis could be compared with the results of Ref. [107], which assessed the information content of primordial non-Gaussianity in the lensing convergence field at nonlinear scales. Additionally, it would be valuable to generalize the method presented here to include galaxy clustering statistics and information from the scale-dependent bias. Finally, the method presented in this work can be readily extended to test alternative nonstandard cosmological scenarios, such as quasisingle field inflation [108–111] or equivalence-principle-violating physics [100], all of which generate analogous poles in the squeezed bispectrum. An observational test of the LSS consistency relations using weak lensing bispectra would also complement the recent test of the cosmological consistency relations using the anisotropic three-point correlation function [112].

ACKNOWLEDGMENTS

We acknowledge computing resources from Columbia University's Shared Research Computing Facility project,

¹²It is worth reiterating that gravitational non-Gaussianity cannot produce poles in the squeezed bispectrum to soft power spectrum ratio, regardless of the complexity of the baryonic feedback effects [35,49]. Thus observing such poles immediately points to PNG (beyond that expected from single-field slow roll inflation) or equivalence principle violation. To infer parameters such as f_{NL} from the residue of such poles though, does require incorporating baryonic corrections in the response function.

which is supported by NIH Research Facility Improvement Grant No. 1G20RR030893-01, and associated funds from the New York State Empire State Development, Division of Science Technology and Innovation (NYSTAR) Contract No. C090171, both awarded April 15, 2010. We acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing HPC resources that have contributed to the research results reported within this paper. O. H. E. P is a Junior Fellow of the Simons Society of Fellows and edited this draft at a speed inspired by the Costa Rican sloth population. J. C. H. acknowledges support from NSF grant No. AST-2108536, NASA grants No. 21-ATP21-0129 and No. 22-ADAP22-0145, the Sloan Foundation, and the Simons Foundation. L. H. acknowledges support by the Grant No. DOE DE-SC011941 and a Simons Fellowship in Theoretical Physics.

APPENDIX: IMPACT OF NONLINEARITIES ON COSMIC SHEAR AND CMB LENSING FORECASTS

In this section, we discuss the impact of nonlinear structure formation on the forecast results in the main text. In the realm of 3D matter distributions, and ignoring gravitational nonlinearities, the redshift dependence of the bispectrum scales as $B \sim P^2/D(z) \sim D^3(z)$, whereas the redshift dependence of the covariance scales as $P^3 \approx D^6$. Consequently, the signal-to-noise should be roughly independent of redshift. Nevertheless, our forecast results show that the cosmic variance error on f_{NL} can differ by up to 50% between the cosmic shear forecasts and the CMB lensing forecasts for fixed scale cuts and

assuming a single tomographic bin. We investigate the source of this discrepancy in this section.

The right panel of Fig. 5 compares the ratio of the cosmic variance error on f_{NL} as a function of ℓ_{max} for a cosmic shear analysis with a single tomographic bin compared to that from a CMB lensing analysis for a variety of nonlinear modeling choices. These results assume the fiducial scale cuts in the main text, with $\ell_{\text{min}} = 10$. The red line corresponds to the analysis choices used in the main forecasts, where all power spectra and the logarithmic derivative $\partial \log(P(k))/\partial \log(\sigma_8^2)$ are computed using HALOFIT. In this case, the shear bispectrum provides $\sim 30\%$ tighter constraints on f_{NL} than the CMB lensing bispectrum. However, if we fix the logarithmic derivative to the linear theory prediction of unity (blue), then the improvement from cosmic shear diminishes considerably. This shows that the main source of discrepancy between the cosmic variance CMB lensing and cosmic shear forecasts is due to the nonperturbative enhancement of the squeezed matter bispectrum due to local PNG. It remains to be seen to what extent this also reduces our constraining power due to the associated non-Gaussian covariance. Finally, we can use linear theory to also compute the bispectrum Eq. (7) and its covariance (gray). In this case, the cosmic shear and CMB lensing forecasts are consistent to within $\simeq 4\%$, with the residual attributed to differences in the projection kernels.

The right panel of Fig. 5 shows the HALOFIT prediction for the logarithmic derivative $\partial \log(P(k))/\partial \log(\sigma_8^2)$ over redshifts $0 < z < 9$. At high redshifts and large wavenumbers the logarithmic derivative is consistent with unity,




FIG. 5. Left: ratio of cosmic variance error on f_{NL} as a function of ℓ_{max} from the cosmic shear bispectrum with $N_{\text{tomo}} = 1$ to that from the CMB lensing convergence bispectrum for a variety of nonlinear modeling choices. The results from the forecasts of the main text are shown in red where the shear bispectrum yields $\sim 30\%$ tighter constraints on f_{NL} than the CMB lensing bispectrum. If we neglect the nonlinear enhancement from the logarithmic derivative (blue), then the forecasted error on f_{NL} agrees within 10% between the two approaches. Finally, if we use linear theory, then the two constraints are consistent to within 4%. The slight discrepancy arises from differences in the projection kernels. Right: HALOFIT prediction for $\partial \log(P(k, z))/\partial \log(\sigma_8^2)$ over redshifts $0 < z < 9$. Vertical lines indicate the approximate maximum 3D wavenumber $k_{\text{max}} \simeq \ell_{\text{max}}/\chi(z)$ probed at a certain redshift assuming $\ell_{\text{max}} = 2000$. For $z \gtrsim 9$ the HALOFIT predictions can become unstable so we fix the logarithmic derivative to 1 (i.e., linear theory) which is accurate at all scales of interest.

as expected from linear theory. At smaller scales and lower redshifts, however, it can pick up a sizeable enhancement which improves the constraints from cosmic shear. The vertical dashed lines indicate the approximate maximum

wavenumber probed $k_{\text{max}} \simeq \ell_{\text{max}}/\chi$ assuming $\ell_{\text{max}} = 2000$. The impact of baryons on the logarithmic derivative could be a significant systematic [53], which should be explored in future studies.

-
- [1] A. Aghamousa *et al.* (DESI Collaboration), [arXiv:1611.00036](#).
 - [2] L. Amendola *et al.*, *Living Rev. Relativity* **21**, 2 (2018).
 - [3] O. Doré *et al.*, [arXiv:1412.4872](#).
 - [4] P. A. Abell *et al.* (LSST Science and LSST Project Collaborations), [arXiv:0912.0201](#).
 - [5] J. M. Maldacena, *J. High Energy Phys.* **05** (2003) 013.
 - [6] P. Creminelli and M. Zaldarriaga, *J. Cosmol. Astropart. Phys.* **10** (2004) 006.
 - [7] P. Creminelli, G. D'Amico, M. Musso, and J. Norena, *J. Cosmol. Astropart. Phys.* **11** (2011) 038.
 - [8] E. Pajer, F. Schmidt, and M. Zaldarriaga, *Phys. Rev. D* **88**, 083502 (2013).
 - [9] P. D. Meerburg *et al.*, *Bull. Am. Astron. Soc.* **51**, 107 (2019).
 - [10] A. Achúcarro *et al.*, [arXiv:2203.08128](#).
 - [11] Y. Akrami *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A9 (2020).
 - [12] G. D'Amico, M. Lewandowski, L. Senatore, and P. Zhang, [arXiv:2201.11518](#).
 - [13] G. Cabass, M. M. Ivanov, O. H. E. Philcox, M. Simonović, and M. Zaldarriaga, *Phys. Rev. Lett.* **129**, 021301 (2022).
 - [14] G. Cabass, M. M. Ivanov, O. H. E. Philcox, M. Simonović, and M. Zaldarriaga, *Phys. Rev. D* **106**, 043506 (2022).
 - [15] F. McCarthy, M. S. Madhavacheril, and A. S. Maniyar, *Phys. Rev. D* **108**, 083522 (2023).
 - [16] G. Cabass, M. M. Ivanov, O. H. E. Philcox, M. Simonovic, and M. Zaldarriaga, *Phys. Lett. B* **841**, 137912 (2023).
 - [17] A. Krolewski *et al.* (DESI Collaboration), [arXiv:2305.07650](#).
 - [18] M. Rezaie *et al.*, [arXiv:2307.01753](#).
 - [19] N. Dalal, O. Dore, D. Huterer, and A. Shirokov, *Phys. Rev. D* **77**, 123514 (2008).
 - [20] S. Matarrese and L. Verde, *Astrophys. J. Lett.* **677**, L77 (2008).
 - [21] A. Slosar, C. Hirata, U. Seljak, S. Ho, and N. Padmanabhan, *J. Cosmol. Astropart. Phys.* **08** (2008) 031.
 - [22] V. Desjacques, U. Seljak, and I. Iliev, *Mon. Not. R. Astron. Soc.* **396**, 85 (2009).
 - [23] B. Reid, L. Verde, K. Dolag, S. Matarrese, and L. Moscardini, *J. Cosmol. Astropart. Phys.* **07** (2010) 013.
 - [24] A. Barreira, G. Cabass, F. Schmidt, A. Pillepich, and D. Nelson, *J. Cosmol. Astropart. Phys.* **12** (2020) 013.
 - [25] A. Barreira, *J. Cosmol. Astropart. Phys.* **01** (2022) 033.
 - [26] T. Lazeyras, A. Barreira, F. Schmidt, and V. Desjacques, *J. Cosmol. Astropart. Phys.* **01** (2023) 023.
 - [27] A. Barreira, *J. Cosmol. Astropart. Phys.* **11** (2022) 013.
 - [28] A. Barreira and E. Krause, *J. Cosmol. Astropart. Phys.* **10** (2023) 044.
 - [29] M. Takada and B. Jain, *Mon. Not. R. Astron. Soc.* **348**, 897 (2004).
 - [30] B. M. Schäfer, A. Grassi, M. Gerstenlauer, and C. T. Byrnes, *Mon. Not. R. Astron. Soc.* **421**, 797 (2012).
 - [31] D. Jeong, F. Schmidt, and E. Sefusatti, *Phys. Rev. D* **83**, 123005 (2011).
 - [32] A. Grassi, L. Heisenberg, C. T. Byrnes, and B. M. Schäfer, *Mon. Not. R. Astron. Soc.* **442**, 1068 (2014).
 - [33] M. Peloso and M. Pietroni, *J. Cosmol. Astropart. Phys.* **05** (2013) 031.
 - [34] A. Kehagias and A. Riotto, *Nucl. Phys.* **B864**, 492 (2012).
 - [35] A. Esposito, L. Hui, and R. Scoccimarro, *Phys. Rev. D* **100**, 043536 (2019).
 - [36] S. Goldstein, A. Esposito, O. H. E. Philcox, L. Hui, J. C. Hill, R. Scoccimarro, and M. H. Abitbol, *Phys. Rev. D* **106**, 123525 (2022).
 - [37] P. Ade *et al.* (Simons Observatory Collaboration), *J. Cosmol. Astropart. Phys.* **02** (2019) 056.
 - [38] K. N. Abazajian *et al.* (CMB-S4 Collaboration), [arXiv:1610.02743](#).
 - [39] T. Namikawa, *Phys. Rev. D* **93**, 121301 (2016).
 - [40] N. Aghanim *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A6 (2020); **652**, C4(E) (2021).
 - [41] M. Bartelmann and P. Schneider, *Phys. Rep.* **340**, 291 (2001).
 - [42] A. Lewis and A. Challinor, *Phys. Rep.* **429**, 1 (2006).
 - [43] M. Bartelmann and M. Maturi, [arXiv:1612.06535](#).
 - [44] M. Kilbinger, *Rep. Prog. Phys.* **78**, 086901 (2015).
 - [45] U. Giri, M. Münchmeyer, and K. M. Smith, *Phys. Rev. D* **107**, L061301 (2023).
 - [46] U. Giri, M. Münchmeyer, and K. M. Smith, [arXiv:2305.03070](#).
 - [47] E. Komatsu and D. N. Spergel, *Phys. Rev. D* **63**, 063002 (2001).
 - [48] R. Scoccimarro, L. Hui, M. Manera, and K. C. Chan, *Phys. Rev. D* **85**, 083002 (2012).
 - [49] B. Horn, L. Hui, and X. Xiao, *J. Cosmol. Astropart. Phys.* **09** (2014) 044.
 - [50] B. Horn, L. Hui, and X. Xiao, *J. Cosmol. Astropart. Phys.* **09** (2015) 068.
 - [51] R. Takahashi, M. Sato, T. Nishimichi, A. Taruya, and M. Oguri, *Astrophys. J.* **761**, 152 (2012).
 - [52] F. Villaescusa-Navarro *et al.*, *Astrophys. J. Suppl. Ser.* **250**, 2 (2020).
 - [53] A. Barreira, D. Nelson, A. Pillepich, V. Springel, F. Schmidt, R. Pakmor, L. Hernquist, and M. Vogelsberger, *Mon. Not. R. Astron. Soc.* **488**, 2079 (2019).
 - [54] P. Valageas, *Phys. Rev. D* **89**, 123522 (2014).

- [55] T. Nishimichi and P. Valageas, *Phys. Rev. D* **90**, 023546 (2014).
- [56] N. Kaiser and G. Squires, *Astrophys. J.* **404**, 441 (1993).
- [57] P. G. Castro, A. F. Heavens, and T. D. Kitching, *Phys. Rev. D* **72**, 023516 (2005).
- [58] B. Leistedt, J. D. McEwen, M. Büttner, and H. V. Peiris, *Mon. Not. R. Astron. Soc.* **466**, 3728 (2017).
- [59] C. G. R. Wallis, M. A. Price, J. D. McEwen, T. D. Kitching, B. Leistedt, and A. Plouviez, *Mon. Not. R. Astron. Soc.* **509**, 4480 (2021).
- [60] C. Chang *et al.* (DES Collaboration), *Mon. Not. R. Astron. Soc.* **475**, 3165 (2018).
- [61] M. Gatti *et al.* (DES Collaboration), *Mon. Not. R. Astron. Soc.* **498**, 4060 (2020).
- [62] A. Barthelemy, A. Halder, Z. Gong, and C. Uhlemann, [arXiv:2307.09468](#).
- [63] W. Hu and T. Okamoto, *Astrophys. J.* **574**, 566 (2002).
- [64] D. N. Limber, *Astrophys. J.* **117**, 134 (1953).
- [65] N. Kaiser, *Astrophys. J.* **388**, 272 (1992).
- [66] M. LoVerde and N. Afshordi, *Phys. Rev. D* **78**, 123506 (2008).
- [67] A. J. S. Hamilton, *Mon. Not. R. Astron. Soc.* **312**, 257 (2000).
- [68] H. S. Grasshorn Gebhardt and D. Jeong, *Phys. Rev. D* **97**, 023504 (2018).
- [69] V. Assassi, M. Simonović, and M. Zaldarriaga, *J. Cosmol. Astropart. Phys.* **11** (2017) 054.
- [70] X. Fang, E. Krause, T. Eifler, and N. MacCrann, *J. Cosmol. Astropart. Phys.* **05** (2020) 010.
- [71] M. Takada and B. Jain, *Mon. Not. R. Astron. Soc.* **395**, 2065 (2009).
- [72] M. Sato, T. Hamana, R. Takahashi, M. Takada, N. Yoshida, T. Matsubara, and N. Sugiyama, *Astrophys. J.* **701**, 945 (2009).
- [73] M. Sato and T. Nishimichi, *Phys. Rev. D* **87**, 123538 (2013).
- [74] M. Takada and D. N. Spergel, *Mon. Not. R. Astron. Soc.* **441**, 2456 (2014).
- [75] E. Krause and T. Eifler, *Mon. Not. R. Astron. Soc.* **470**, 2100 (2017).
- [76] A. Barreira, E. Krause, and F. Schmidt, *J. Cosmol. Astropart. Phys.* **10** (2018) 053.
- [77] W. Hu, *Phys. Rev. D* **62**, 043007 (2000).
- [78] I. Kayo, M. Takada, and B. Jain, *Mon. Not. R. Astron. Soc.* **429**, 344 (2012).
- [79] I. Kayo and M. Takada, [arXiv:1306.4684](#).
- [80] K. C. Chan, A. Moradinezhad Dizgah, and J. Noreña, *Phys. Rev. D* **97**, 043532 (2018).
- [81] T. Flöss, M. Biagetti, and P. D. Meerburg, *Phys. Rev. D* **107**, 023528 (2023).
- [82] A. E. Bayer, J. Liu, R. Terasawa, A. Barreira, Y. Zhong, and Y. Feng, *Phys. Rev. D* **108**, 043521 (2023).
- [83] K. Abazajian *et al.*, [arXiv:1907.04473](#).
- [84] C. M. Hirata and U. Seljak, *Phys. Rev. D* **68**, 083002 (2003).
- [85] J. Carron and A. Lewis, *Phys. Rev. D* **96**, 063510 (2017).
- [86] M. Robertson and A. Lewis, *J. Cosmol. Astropart. Phys.* **08** (2023) 048.
- [87] R. Laureijs *et al.* (EUCLID Collaboration), [arXiv:1110.3193](#).
- [88] R. Mandelbaum *et al.* (LSST Dark Energy Science Collaboration), [arXiv:1809.01669](#).
- [89] C. Chang, M. Jarvis, B. Jain, S. M. Kahn, D. Kirkby, A. Connolly, S. Krughoff, E. Peng, and J. R. Peterson, *Mon. Not. R. Astron. Soc.* **434**, 2121 (2013).
- [90] S. J. Osborne, D. Hanson, and O. Doré, *J. Cosmol. Astropart. Phys.* **03** (2014) 024.
- [91] A. van Engelen, S. Bhattacharya, N. Sehgal, G. P. Holder, O. Zahn, and D. Nagai, *Astrophys. J.* **786**, 13 (2014).
- [92] S. Ferraro and J. C. Hill, *Phys. Rev. D* **97**, 023512 (2018).
- [93] H. Cai, M. S. Madhavacheril, J. C. Hill, and A. Kosowsky, *Phys. Rev. D* **105**, 043516 (2022).
- [94] F. Bernardeau, C. Bonvin, N. Van de Rijt, and F. Vernizzi, *Phys. Rev. D* **86**, 023001 (2012).
- [95] M. S. Wang, W. J. Percival, S. Avila, R. Crittenden, and D. Bianchi, *Mon. Not. R. Astron. Soc.* **486**, 951 (2019).
- [96] B. Tucci and F. Schmidt, [arXiv:2310.03741](#).
- [97] S. Aiola *et al.* (CMB-HD Collaboration), [arXiv:2203.05728](#).
- [98] F. J. Qu, B. D. Sherwin, O. Darwish, T. Namikawa, and M. S. Madhavacheril, *Phys. Rev. D* **107**, 123540 (2023).
- [99] M. Schmittfull and U. Seljak, *Phys. Rev. D* **97**, 123540 (2018).
- [100] P. Creminelli, J. Gleyzes, L. Hui, M. Simonović, and F. Vernizzi, *J. Cosmol. Astropart. Phys.* **06** (2014) 009.
- [101] G. Pratten and A. Lewis, *J. Cosmol. Astropart. Phys.* **08** (2016) 047.
- [102] G. Fabbian, M. Calabrese, and C. Carbone, *J. Cosmol. Astropart. Phys.* **02** (2018) 050.
- [103] A. Kalaja, G. Orlando, A. Bowkis, A. Challinor, P. D. Meerburg, and T. Namikawa, *J. Cosmol. Astropart. Phys.* **04** (2023) 041.
- [104] M. A. Troxel and M. Ishak, *Phys. Rep.* **558**, 1 (2014).
- [105] M. Gatti *et al.* (DES Collaboration), *Mon. Not. R. Astron. Soc.* **527**, L115 (2023).
- [106] F. Ferlito *et al.*, *Mon. Not. R. Astron. Soc.* **524**, 5591 (2023).
- [107] D. Anbajagane, C. Chang, H. Lee, and M. Gatti, [arXiv:2310.02349](#).
- [108] V. Assassi, D. Baumann, and D. Green, *J. Cosmol. Astropart. Phys.* **11** (2012) 047.
- [109] X. Chen and Y. Wang, *J. Cosmol. Astropart. Phys.* **04** (2010) 027.
- [110] D. Baumann, A. Nicolis, L. Senatore, and M. Zaldarriaga, *J. Cosmol. Astropart. Phys.* **07** (2012) 051.
- [111] T. Noumi, M. Yamaguchi, and D. Yokoyama, *J. High Energy Phys.* **06** (2013) 051.
- [112] N. S. Sugiyama, D. Yamauchi, T. Kobayashi, T. Fujita, S. Arai, S. Hirano, S. Saito, F. Beutler, and H.-J. Seo, *Mon. Not. R. Astron. Soc.* **524**, 1651 (2023).