Statistical field theory for weak gravitational lensing

Zheng Zhang,[1,*] Philip Bull,[1,2] Chris Clarkson,[3,2] and Andrina Nicola[1]

[1]*Jodrell Bank Centre for Astrophysics, University of Manchester, Manchester, M13 9PL, United Kingdom*
[2]*Department of Physics and Astronomy, University of the Western Cape, Cape Town 7535, South Africa*
[3]*Astronomy Unit, School of Physical and Chemical Sciences, Queen Mary University of London, London E1 4NS, United Kingdom*

(Received 23 June 2026; accepted 8 October 2026; published 18 November 2026)

Standard weak-lensing calculations treat lensing as a linear remapping of the matter field along the line of sight. We instead formulate lensing as a stochastic field theory for the Sachs optical scalars, driven by random Ricci-focusing and Weyl-shearing fields. With convergence and shear expressed linearly in the Sachs-scalar fluctuations, the resulting path integral generates a diagrammatic expansion for their n-point correlation functions, organized into linear response, nonlinear propagation, and driving-field cumulants. The conventional calculation emerges as the lowest-order, linear-propagation limit. Beyond it, nonlinear Sachs evolution couples to driving-field non-Gaussianity, mixing the matter cumulant hierarchy into the lensing hierarchy. A selection rule governs the couplings: an n-point observable receives a direct contribution from the n-point driving-field cumulant, and its leading hierarchy-mixing correction from the $(n+1)$-point cumulant via one nonlinear Sachs interaction, with higher cumulants entering only at higher order. The two-point function, for instance, receives corrections from three-point cumulants of Ricci focusing and Weyl shearing through nonlinear Sachs evolution. Rather than a restrictive approximation scheme, the formalism is a paradigm shift: a unified framework naturally accommodating path corrections, higher-order matter statistics, stochasticity, and small-scale effects.

DOI: 10.1103/w3tj-slzm

I. INTRODUCTION

Weak gravitational lensing converts the tidal history of light rays into cosmological observables. Current cosmic-shear analyses already constrain large-scale structure at subpercent statistical precision [1,2], and at Stage-IV precision [3–5] the accuracy of the light-propagation model becomes a limiting systematic in its own right. The standard calculation treats lensing as a projection: matter correlations are filtered along the line of sight and read out as convergence and shear statistics [6–8]. In the Born limit this gives the familiar hierarchy in which the convergence two-point function is derived from the matter power spectrum, the convergence bispectrum inherits the matter bispectrum, and so on.

This projection picture is useful, but it packages together approximations that must be refined for next-generation wide, deep, high-precision weak-lensing surveys such as LSST, Euclid, and Roman [3–5]. Post-Born calculations include corrections from evaluating lensing fields along perturbed rays and from coupling successive lensing distortions (lens-lens coupling) [9–15]; higher-order perturbation theory and phenomenological matter models account for non-Gaussian structure formation [11,16–19]; beyond-Limber treatments restore long line-of-sight correlations [20–22]. These ingredients are usually developed as separate corrections to the same baseline projection calculation, often quantified one assumption at a time [23]. High-fidelity ray tracing can combine several of them numerically [14,24–26], but its accuracy is ultimately bounded by finite resolution and volume, sampling variance, and lens-plane artifacts, at substantial computational cost, and it does not by itself identify which physical mechanism controls a given residual. Our aim is to determine how nonlinear optical evolution combines with non-Gaussian gravitational-field statistics to generate lensing correlations, while keeping their contributions explicit.

We build such a framework from the Sachs optical equations [27], the exact relativistic equations that govern how the cross-section of a thin bundle of neighbouring light rays is focused and sheared as it propagates through an arbitrary spacetime. We treat the beam as a test congruence: spacetime curvature drives its optical evolution, with no feedback on matter or geometry. For a given metric, the optical expansion and shear evolve deterministically.

*Contact author: zheng.zhang@manchester.ac.uk

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

We prescribe the Ricci focusing scalar Φ_{00} and the Weyl shearing scalar Ψ_0 as random driving fields along the rays, making the optical problem stochastic throughout propagation. We promote the Sachs trajectory to a path integral [28–30], which encodes the nonlinear Sachs evolution together with the statistics of the driving fields in a single probability functional, so that in principle the statistics of observables reconstructed from the Sachs trajectory follow from it. The screen-space Jacobi map carries a small separation of observed directions to the transverse separation of neighboring rays at the source. Its distortion relative to the reference beam defines convergence and shear. We use linearized relations between these observables and the Sachs scalars (Sec. IV C), while retaining nonlinear Sachs evolution in the path integral. The resulting diagrammatic expansion resolves every n-point correlation function (nPCF) of convergence and shear into three ingredients: the linear response, nonlinear propagation, and the connected cumulants of the driving fields. For later reference we label these last two mechanisms F (nonlinear propagation) and $K^{(n)}$ (an n-point cumulant of the driving fields); in field-theory language they play the role of interaction vertices.

The statistics of these driving fields depend on the matter model. Standard perturbation theory predicts matter correlations from gravitational dynamics [31]. Effective descriptions account for unresolved small-scale physics through effective stresses, including stochastic contributions [32]. Vlasov perturbation theory incorporates velocity dispersion and higher cumulants of the local velocity distribution generated by orbit crossing [33]. These cumulants describe particle velocities at a given position. These matter models provide correlations of matter and gravitational fields from which the Ricci and Weyl driving-field statistics along light paths can be constructed.

Related diagrammatic and path-integral calculations of large-scale structure predict matter correlations by evolving random initial density and velocity fields under deterministic gravitational dynamics [34,35]. These examples use Gaussian initial statistics, with extensions to primordial non-Gaussianity [36]. Our optical problem has fixed observer boundary conditions for the Sachs fields and random Ricci and Weyl driving fields acting throughout propagation. For each driving-field realization, Sachs evolution is deterministic; averaging over their joint histories predicts correlations of convergence and shear. The driving fields may correlate different observed directions and affine parameters and need not be Gaussian. Reference [37] develops the general prescription adopted here to calculate how such driving statistics are transformed by nonlinear evolution into the statistics of the driven fields. It also automates the resulting diagrammatic calculations. The perturbative order of a matter model and the order of our optical diagrams thus refer to different expansions.

At leading order, this optical expansion recovers the conventional lensing-efficiency projection. The key identity is that propagating an n-point driving-field cumulant through the free Sachs response, with no nonlinear Sachs (F) insertion, is equivalent to weighting that cumulant by n copies of the lensing-efficiency window, the familiar geometric weight that sets how strongly structure at a given distance lenses the source. Thus the leading branch of the n-point observable is the usual line-of-sight projection of the n-point driving-field cumulant: the free two-point sector reproduces the standard two-point kernel, the leading three-point sector projects the driving-field three-point cumulant, and so forth.

With convergence and shear expressed as linear integrals of Sachs-scalar fluctuations, nonlinear Sachs evolution and gravitational non-Gaussianity combine to mix the statistical hierarchies, so that the order of a lensing observable and the order of the driving-field cumulant that sources it need not coincide: their interplay is governed by a selection rule. For an n-point lensing observable, the direct cumulant contribution comes from the n-point driving-field cumulant; the first hierarchy-mixing correction is sourced by the $(n+1)$-point cumulant through one nonlinear Sachs correction, and higher cumulants are absent at that order. Equivalently, a driving-field three-point cumulant contributes directly to a three-point lensing observable, but enters a two-point observable through one nonlinear interaction vertex, F. The observable hierarchy and the Ricci/Weyl driving-field hierarchy are therefore mixed, rather than tied by a strict one-to-one projection, with each step between them mediated by a nonlinear Sachs correction. This nonlinear-Sachs hierarchy mixing is present regardless of the Born limit or the perturbative matter setup.

As a demonstration, we work out this structure for the two-point functions of lensing observables. At Order-0 the formalism reproduces the conventional weak-lensing kernels for the four two-point functions considered. At the next order each of them acquires two physically distinct corrections: FF, from nonlinear optical propagation, and FK, in which Ricci/Weyl three-point cumulants enter the two-point statistic through one nonlinear Sachs correction. Under the Limber approximation adopted in the demonstration, the three-point cumulant enters in a collapsed configuration with two coincident legs, coupling large separations to small-scale driving-field fluctuations. Its amplitude depends on the matter model, source redshift, angular scale, and multipole cutoff (Sec. V D). For the linear spectra and tree-level bispectrum used here, at $z_s = 5$ and $\ell_{\max} = 15360$ the FK contribution raises the convergence correlation by 1.5%–1.6% over the sampled separations within $2'$–$12'$, the band considered in the discussion of current cosmic-shear analyses [1,2]. This increases from 0.86–0.94% at $z_s = 1$ (Fig. 15).

The demonstration scenario adopted in this paper is deliberately narrow: scalar perturbations with $\Phi = \Psi$, unperturbed path, a linear matter power spectrum, and a three-point cumulant evaluated from the tree-level bispectrum. The matter model determines the curvature statistics. Relaxing the unperturbed-path approximation changes both where the curvature is evaluated and how it is projected onto the ray and screen. At the level of the Sachs equations, these effects are encoded in the joint statistics of the resulting driving fields as functions of observed direction and affine parameter. The example neglects effects including shell crossing [33] and nonlinear baryonic feedback [38], so its numerical corrections illustrate the mechanism without providing calibrated small-scale predictions. We assume that each source is seen as a single, weakly distorted image. Caustics and multiple imaging invalidate this single-image description and cannot be accounted for by changing the driving-field cumulants alone.

In particular, the non-Gaussianity leakage into the 2PCF depends on the collapsed three-point cumulant, where the tree-level bispectrum underestimates the nonlinear small-scale clustering; a nonlinear evaluation, however, requires more than substituting a nonlinear bispectrum. It calls for N-body simulations with baryonic corrections, or for a separate-universe treatment of the squeezed response [39–41], and it calls for a scale below which the driving field is treated as unresolved, together with the counterterms that smoothing the field generates in the nonlinear propagation. We leave both to future work. Within the single-image regime, perturbed-path corrections [12,24], nonlinear matter statistics [17,31], primordial non-Gaussianity [18], and small-scale physics [38] modify the driving-field statistics entering the Sachs evolution. For perturbed paths, these must be the statistics of curvature projected along the corrected rays and onto their screen bases. Changes to the mapping from Sachs fields to measured observables must be treated separately.

Because the starting point is the Sachs system rather than a cosmology-specific projection operator, the same construction can also be adapted to other light-cone observables, including the lensing of the stochastic gravitational-wave background in the geometric-optics limit [42].

This paper is organized as follows. Section II formulates the Sachs evolution as a centred stochastic dynamical system. Section III builds the path integral for the Sachs equations and derives the diagrammatic rules and their selection rule. Section IV specializes the driving fields and the free Sachs response to a perturbed Friedmann-Lemaître-Robertson-Walker (FLRW) spacetime in Poisson gauge and establishes the lensing-kernel correspondence. Section V presents the Order-0 and FF/FK two-point results. Section VI summarizes the worked example and its three structural insights. The Appendixes collect conventions, Jacobi-map relations, driving-field three-point cumulants, and numerical validations.

II. NONLINEAR EVOLUTION OF NULL GEODESICS CONGRUENCE

We describe the observed light beam as a two-dimensional congruence of null geodesics [27,43,44], labeled by an affine parameter λ and an observed sky direction $\hat{\Omega}$. At each point on a ray we choose a Newman-Penrose null tetrad [45]

$$\{k^\mu, e^\mu, Z^\mu, \bar{Z}^\mu\}, \tag{1}$$

where k^μ is tangent to the ray, e^μ is the auxiliary null direction, and Z^μ, $\bar{Z}^\mu$ span the screen plane with $k^\mu e_\mu = -1$ and $Z^\mu \bar{Z}_\mu = 1$. The complete construction and the spin-weight conventions are collected in Appendix A; here we keep only the ingredients used below. Figure 1 summarizes the local geometry and the optical degrees of freedom used in this section.

The two screen-projected contractions of $\nabla_\nu k_\mu$ are the Newman-Penrose spin coefficients

$$\rho = -\nabla_\nu k_\mu Z^\mu \bar{Z}^\nu, \qquad \sigma = \nabla_\nu k_\mu Z^\mu Z^\nu. \tag{2}$$

For an affinely parametrized congruence with a parallel-transported screen basis, their propagation equations are [27]

$$\frac{\mathrm{d}\rho}{\mathrm{d}\lambda} = \rho^2 + \sigma\bar{\sigma} - \Phi_{00}, \tag{3}$$

$$\frac{\mathrm{d}\sigma}{\mathrm{d}\lambda} = (\rho + \bar{\rho})\sigma + \Psi_0. \tag{4}$$

The two source fields are the Ricci focusing scalar and the Weyl shear scalar,

$$\Phi_{00} = -\frac{1}{2} R_{\mu\nu} k^\mu k^\nu, \tag{5}$$

$$\Psi_0 = -C_{\alpha\beta\gamma\delta} k^\alpha Z^\beta k^\gamma Z^\delta = -R_{\alpha\beta\gamma\delta} k^\alpha Z^\beta k^\gamma Z^\delta. \tag{6}$$

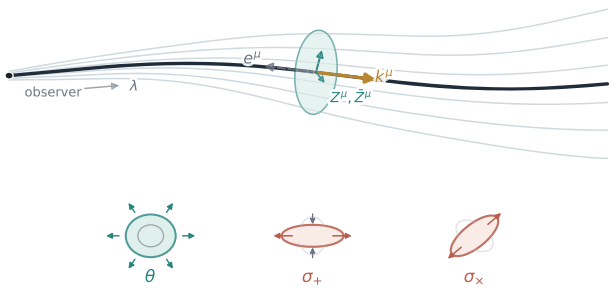


FIG. 1. Schematic of a null geodesic congruence evolved along a past-directed affine parameter, the local Newman-Penrose tetrad, and the Sachs scalar degrees of freedom $(\theta, \sigma_+, \sigma_\times)$ on the screen.

Here $C_{\alpha\beta\gamma\delta}$ is the Weyl tensor. The second equality in Eq. (6) follows because the Ricci and trace parts of the Riemann tensor vanish under the contraction $k^\alpha Z^\beta k^\gamma Z^\delta$. In natural units, the Einstein equations also give

$$R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu} = 8\pi G T_{\mu\nu} \Rightarrow \Phi_{00} = -4\pi G T_{\mu\nu} k^\mu k^\nu, \quad (7)$$

where the $R g_{\mu\nu}$ contraction drops out because k^μ is null. These equations provide the geometric-to-statistical bridge used throughout the paper: Φ_{00} drives focusing, while the real and imaginary parts of Ψ_0 control the two shear polarizations. Section IV evaluates these source fields explicitly in Poisson gauge.

A. Field equations of Sachs scalars

The complex scalars ρ and σ are in one-to-one correspondence with the more familiar real-valued optical scalars through the mapping

$$\rho = -(\theta + \mathrm{i}\omega), \qquad \sigma = \sigma_+ + \mathrm{i}\sigma_\times \quad (8)$$

where $\{\theta, \sigma_+, \sigma_\times, \omega\}$ are the Sachs (or optical) scalars. Geometrically, θ measures the expansion of the congruence, while the two real shear components describe the elliptic distortion of the beam cross-section: σ_+ stretches and compresses along the screen-basis axes $\mathbf{x}$ and $\mathbf{y}$, and $\sigma_\times$ gives the corresponding distortion along axes rotated by 45°. The remaining scalar ω is the optical twist of the null-ray congruence, the antisymmetric part of its local screen-projected deformation rate. Together, these scalars form the basis of the standard weak-lensing formalism.

In terms of the optical scalars, the Newman-Penrose field equations reduce to the familiar Sachs equations. Concretely, we view the congruence as a two-dimensional field of null rays, each labeled by its sky direction $\hat{\Omega}$ and evolving along its affine parameter λ. On the smooth part of the observer's past light cone, the rays are normal to wave fronts and the optical twist vanishes by Frobenius' theorem, $\omega = 0$ [46]. The Sachs (or Raychaudhuri) system then governs the evolution of the three nontrivial components,

$$\dot{\theta} = -\theta^2 - \sigma_+^2 - \sigma_\times^2 + \Phi_{00}, \quad (9)$$

$$\dot{\sigma_+} = -2\theta\sigma_+ + \mathcal{W}_1, \quad (10)$$

$$\dot{\sigma_\times} = -2\theta\sigma_\times + \mathcal{W}_2, \quad (11)$$

where the dot denotes differentiation with respect to λ, Φ_{00} is the Ricci focusing scalar introduced above, and $\mathcal{W}_{1,2}$ denote the real and imaginary parts of Ψ_0.

The Sachs equations combine quadratic optical couplings with the Ricci and Weyl driving fields, whose statistics encode the matter and curvature fluctuations along and across the light paths. For convenience, we collect the relevant fields into the vectors

$$X \equiv (\theta, \sigma_+, \sigma_\times), \qquad \boldsymbol{\phi} \equiv (\Phi_{00}, \mathcal{W}_1, \mathcal{W}_2). \quad (12)$$

In this notation, the evolution equations take the compact form

$$\dot{X}_a(\hat{\Omega}, \lambda) = F_{abc} X_b(\hat{\Omega}, \lambda) X_c(\hat{\Omega}, \lambda) + \boldsymbol{\phi}_a(\hat{\Omega}, \lambda) \quad (13)$$

where F_{abc} is a constant coupling array whose nonzero entries are listed in Table I. Throughout the remainder of this paper we adopt the convention that repeated Latin indices labeling the three field components are implicitly summed over $\{1, 2, 3\}$.

TABLE I. Entries of F_{abc} that are not zero.

abc	111	122	133	212	313
F_{abc}	−1	−1	−1	−2	−2

To connect this local optical evolution to image distortions, we first distinguish the different notions of rotation.

1. Optical twist and other notions of rotation

Optical twist, image rotation, matter vorticity, and screen-basis rotation describe different quantities despite their similar terminology. As noted above, optical twist vanishes on the observer's smooth light cone; nonzero optical twist requires a congruence that is not normal to a family of wave fronts.

Image rotation is defined by the antisymmetric part of the accumulated screen-space Jacobi map relative to parallel-transported screen bases. Zero optical twist does not require this map to be symmetric. The first-order Jacobi map is symmetric even for general linear perturbations [47], but scalar lens-lens coupling can generate image rotation at second order in gravitational perturbations [9]. This order counts powers of the gravitational fields, whereas our Order-0 and Order-2 labels count F and K insertions in the two-point functions.

Matter vorticity describes the rotation of a matter velocity field. For collisionless dark matter, a coarse-grained mean velocity field can develop vorticity after shell crossing [48]. This is distinct from optical twist: matter evolution changes the Ricci and Weyl driving fields, while the observer's smooth light-cone congruence remains normal to wave fronts.

Screen-basis rotation concerns the orientation of the screen frame rather than the deformation of the beam. The standard Sachs and screen-space Jacobi equations do not by themselves determine the basis orientation relative to a chosen reference frame; this requires tracking the basis transport separately [47]. We do not evaluate it in this work.

2. Nonlinear Sachs evolution and conversion to observables

The distinction between local deformation rates and accumulated image distortions also helps separate two sources of nonlinearity: the quadratic couplings in the Sachs equations and the relations that convert optical deformation rates into convergence and shear. Here nonlinear propagation refers to the former. Our demonstration retains these Sachs couplings but linearizes the conversion to observables. For convergence,

$$\kappa(\hat{\Omega},\lambda_s) = -\int_0^{\lambda_s} \delta\theta(\hat{\Omega},\lambda')d\lambda' + \mathcal{O}(\varepsilon^2),$$

where $\delta\theta = \theta - \theta^{(\mathrm{sa})}$, and ε counts powers of deformation-rate fluctuations in the observable map. Quadratic and higher-order contributions to these observable relations are omitted in this demonstration. They include both symmetric convergence/shear corrections and antisymmetric image rotation. The path-ordered reconstruction in Sec. IV C displays the quadratic terms explicitly. Our example therefore retains nonlinear Sachs evolution with a linearized observable map and does not provide a complete second-order weak-lensing calculation. The reference expansion $\theta^{(\mathrm{sa})}$ is defined below; Sec. IV C gives the linear convergence and shear relations in detail.

B. Mean field expansion

To organize a perturbative treatment of Eq. (13), we first isolate the evolution driven by the mean source. We solve the nonlinear Sachs equations with the ensemble-mean driving field $\bar{\boldsymbol{\phi}}(\lambda)$ and the same observer boundary conditions as the stochastic Sachs fields,

$$\dot{X}_a^{(\mathrm{sa})}(\lambda) = F_{abc}X_b^{(\mathrm{sa})}(\lambda)X_c^{(\mathrm{sa})}(\lambda) + \bar{\boldsymbol{\phi}}_a(\lambda). \tag{14}$$

The superscript (sa) indicates that the source is averaged before solving the nonlinear equations, so $X^{(\mathrm{sa})}$ need not equal the ensemble mean of X. Statistical isotropy and the observer boundary conditions make this solution independent of sky direction. In the flat FLRW application below, it is the homogeneous optical solution with zero shear and background distance $\bar{D} = a\chi$ (Appendix B).

We define fluctuations by subtracting this solution and the mean driving field,

$$\mathcal{X}(\hat{\Omega},\lambda) = X(\hat{\Omega},\lambda) - X^{(\mathrm{sa})}(\lambda), \tag{15}$$

$$\varphi(\hat{\Omega},\lambda) = \boldsymbol{\phi}(\hat{\Omega},\lambda) - \bar{\boldsymbol{\phi}}(\lambda). \tag{16}$$

In the flat FLRW application, the first subtraction removes the nonzero homogeneous optical expansion, while the second removes the mean Ricci-focusing source. Subtracting Eq. (14) from Eq. (13) yields the evolution equation for $\mathcal{X}$,

$$\dot{\mathcal{X}}_a = A_{ab}\mathcal{X}_b + F_{abc}\mathcal{X}_b\mathcal{X}_c + \varphi_a \tag{17}$$

where

$$A_{ab}(\lambda) = (F_{abc} + F_{acb})X_c^{(\mathrm{sa})}(\lambda). \tag{18}$$

The Sachs fluctuation vanishes at the observer, and the driving-field fluctuation has zero mean. The linear evolution set by A, together with the source two-point cumulant, defines the Gaussian reference used in Sec. III. The quadratic Sachs term and higher source cumulants are treated as interactions.

For an isotropic cosmological mean source, the Weyl components vanish and $\bar{\boldsymbol{\phi}}(\lambda) = (\bar{\Phi}_{00}(\lambda), 0, 0)$. The observer conditions then give zero shear for the source-averaged solution, and the linear coefficient takes the diagonal form

$$A_{ab}(\lambda) = -2\delta_{ab}\theta^{(\mathrm{sa})}(\lambda), \tag{19}$$

where δ_{ab} is the Kronecker delta. Appendix B gives the Jacobi equation and vertex conditions underlying $\theta^{(\mathrm{sa})}$.

III. PATH INTEGRAL OF SACHS TRAJECTORY

We now specialize the formulation of Ref. [37] to the centered Sachs equation [see Eq. (17)]. The response-field construction has also been used in path-integral calculations of matter correlations [35,36]. Here, averaging over the prescribed driving fields gives the probability functional of the Sachs fields. The linear evolution and two-point driving cumulant define the Gaussian reference, whose correlation and response propagators supply the diagram lines. Nonlinear Sachs evolution and three-point and higher driving-field cumulants supply the interaction vertices.

A. Probability functional for stochastic light propagation

Equation (17) describes the evolution of the Sachs fields along the affine parameter, with sky direction labeling the rays. For a given realization of the Ricci and Weyl driving fields and fixed observer boundary conditions, this evolution is deterministic. Its stochastic character comes from the ensemble of externally prescribed driving fields, whose values can be correlated across rays and affine parameters. We enforce this deterministic evolution in the path integral through a functional δ constraint,

$$P[\mathcal{X}|\varphi] = \prod_a \delta[\dot{\mathcal{X}}_a - A_{ab}\mathcal{X}_b - F_{abc}\mathcal{X}_b\mathcal{X}_c - \varphi_a]. \tag{20}$$

For brevity, we suppress overall normalization factors, which cancel against the partition function in any expectation value.

The unconditional probability functional for the $\mathcal{X}$ trajectory then follows by marginalizing over the driving field φ,

$$P[\mathcal{X}] = \int P[\mathcal{X}|\varphi]P[\varphi]\mathrm{D}\varphi. \tag{21}$$

The δ functional in $P[\mathcal{X}|\varphi]$ is made tractable by trading it for an auxiliary functional integral through its Fourier representation. Equation (21) thereby takes the form

$$P[\mathcal{X}] = \int \mathrm{D}\varphi\mathrm{D}\tilde{\mathcal{X}}P[\varphi]\exp\Big[-\int \mathrm{i}\tilde{\mathcal{X}}_a \times (\dot{\mathcal{X}}_a - A_{ab}\mathcal{X}_b - F_{abc}\mathcal{X}_b\mathcal{X}_c - \varphi_a)d\hat{\Omega}d\lambda\Big], \tag{22}$$

where $\tilde{\mathcal{X}}_a$ is the so-called response field, conjugate to the full argument of the δ functional, namely $(\dot{\mathcal{X}}_a - A_{ab}\mathcal{X}_b - F_{abc}\mathcal{X}_b\mathcal{X}_c - \varphi_a)$. The remaining functional integral over the stochastic driving field φ is precisely the cumulant generating function (CGF) of φ evaluated at $\mathrm{i}\tilde{\mathcal{X}}$, so that the probability functional reduces into

$$P[\mathcal{X}] = \int \mathrm{D}\tilde{\mathcal{X}}e^{-S[\mathcal{X},\tilde{\mathcal{X}}]} \tag{23}$$

where S is the "action" given by

$$S[\mathcal{X},\tilde{\mathcal{X}}] = \int \mathrm{i}\tilde{\mathcal{X}}_a(\dot{\mathcal{X}}_a - A_{ab}\mathcal{X}_b - F_{abc}\mathcal{X}_b\mathcal{X}_c)d\hat{\Omega}d\lambda - W_\varphi[\mathrm{i}\tilde{\mathcal{X}}], \tag{24}$$

and W_φ is the CGF of the stochastic driving field,

$$W_\varphi[\mathrm{i}\tilde{\mathcal{X}}] = \sum_{n=2}^{\infty} W_\varphi^{(n)}[\mathrm{i}\tilde{\mathcal{X}}] \tag{25}$$

where the sum starts at $n = 2$ because the first cumulant (the mean) vanishes by construction, and each term is given by

$$W_\varphi^{(n)}[\mathrm{i}\tilde{\mathcal{X}}] = \frac{\mathrm{i}^n}{n!}\sum_{\{a_k\}}\int\prod_{m=1}^{n} dz_m \times \tilde{\mathcal{X}}_{a_1}(z_1)\ldots\tilde{\mathcal{X}}_{a_n}(z_n)\mathcal{K}^{(n)}_{a_1\ldots a_n}(z_1;\cdots;z_n) \tag{26}$$

where $a_k \in \{1,2,3\}$ for $k = 1,\ldots,n$, the shorthand $z_k \equiv \{\hat{\Omega}_k, \lambda_k\}$ collects the kth affine and directional coordinates, and

$$\mathcal{K}^{(n)}_{a_1\ldots a_n}(\hat{\Omega}_1,\lambda_1;\ldots;\hat{\Omega}_n,\lambda_n) \equiv \langle\varphi_{a_1}(\hat{\Omega}_1,\lambda_1)\ldots\varphi_{a_n}(\hat{\Omega}_n,\lambda_n)\rangle_c \tag{27}$$

are the connected n-point functions of φ. The full statistics of the stochastic sources are thus encoded in the cumulant hierarchy. This representation is particularly convenient: the first cumulant vanishes by centring, and for purely Gaussian sources all cumulants of order $n \geq 3$ vanish as well. The hierarchy can therefore be truncated at an order dictated by the assumed statistical model.

This hierarchy can be recast as the formal split used in the perturbative expansion below. After centering fixes the one-point source cumulant to zero, the effective linear Sachs operator and the connected two-point source cumulant define the Gaussian-reference theory: the linear operator generates the response propagator, while the two-point source cumulant supplies the noise kernel entering the correlation propagator. The remaining nonquadratic pieces, namely the nonlinear Sachs evolution and the connected source cumulants with $n \geq 3$, are treated as interaction vertices. At this stage no choice of driving-field channels or cumulant truncation is assumed.

B. Free propagators and interaction vertices

A direct evaluation of the path integral is, in general, intractable. Following standard practice in field theory, we therefore adopt a perturbative approach: the action is split into a free (quadratic) part and an interacting (higher-order) part, and observables are computed as convergent or asymptotic series in the interaction.

Concretely, we decompose the action S in Eq. (24) into a quadratic reference action S_0 and an interaction action S_{int} that collects all nonquadratic contributions,

$$S = S_0 + S_{\mathrm{int}}, \tag{28}$$

$$S_0 = \mathrm{i}\int \tilde{\mathcal{X}}_a(\dot{\mathcal{X}}_a - A_{ab}\mathcal{X}_b)d\hat{\Omega}d\lambda - W_\varphi^{(2)}[\mathrm{i}\tilde{\mathcal{X}}], \tag{29}$$

$$S_{\mathrm{int}} = -\mathrm{i}\int F_{abc}\tilde{\mathcal{X}}_a\mathcal{X}_b\mathcal{X}_c d\hat{\Omega}d\lambda - \sum_{k=3}^{\infty} W_\varphi^{(k)}[\mathrm{i}\tilde{\mathcal{X}}], \tag{30}$$

where $W_\varphi^{(k)}$ is defined in Eq. (26). The reference action S_0 encodes the linear dynamics together with the Gaussian statistics of the driving field, whereas S_{int} collects both the nonlinear dynamics and all non-Gaussian contributions of the driving field.

For any observable $\mathcal{O}$ of the Sachs fields, the expectation value with respect to the full action is

$$\langle\mathcal{O}\rangle_S \equiv \int \mathrm{D}\mathcal{X}\mathrm{D}\tilde{\mathcal{X}}\mathcal{O}e^{-S[\mathcal{X},\tilde{\mathcal{X}}]}. \tag{31}$$

Expanding the factor $e^{-S_{\rm int}}$ order by order, the expectation value reduces to a series of Gaussian (reference-theory) averages,

$$\langle \mathcal{O} \rangle_S = \sum_{n=0}^{N} \frac{(-1)^n}{n!} \langle \mathcal{O} S_{\rm int}^n \rangle_{S_0}. \tag{32}$$

Physical observables are typically expressed as moments (correlation functions) of the Sachs fields, and $S_{\rm int}$ itself is a polynomial in the physical field $\mathcal{X}$ and the response field $\tilde{\mathcal{X}}$. Every term in Eq. (32) is therefore a moment evaluated in the Gaussian reference theory defined by S_0, for which Wick's theorem applies: odd-order moments vanish, while even-order moments factorize into products of two-point functions. Two independent two-point functions suffice,

$$\langle \mathcal{X}_a(\hat{\Omega}, \lambda) \mathcal{X}_b(\hat{\Omega}', \lambda') \rangle_{S_0} \equiv \mathcal{C}_{ab}(\hat{\Omega}, \lambda; \hat{\Omega}', \lambda'), \tag{33}$$

$$\langle \tilde{\mathcal{X}}_a(\hat{\Omega}, \lambda) \mathcal{X}_b(\hat{\Omega}', \lambda') \rangle_{S_0} \equiv -i \mathcal{R}_{ab}(\hat{\Omega}, \lambda; \hat{\Omega}', \lambda') \tag{34}$$

where $\mathcal{C}$ is the correlation propagator of the physical fields, and $\mathcal{R}$ is the response propagator, which measures how $\mathcal{X}$ reacts to an infinitesimal source conjugate to $\tilde{\mathcal{X}}$. The autocorrelation of the response field itself vanishes identically,

$$\langle \tilde{\mathcal{X}}_a \tilde{\mathcal{X}}_b \rangle_{S_0} = 0. \tag{35}$$

More explicitly, the response propagator is defined as

$$\begin{aligned} \mathcal{R}_{ab}(\hat{\Omega}, \lambda; \hat{\Omega}', \lambda') &\equiv \delta(\hat{\Omega}' - \hat{\Omega}) \Theta(\lambda' - \lambda) \\ &\quad \times \left\{ \mathcal{T} \exp\left(\int_{\lambda}^{\lambda'} \mathbf{A}(\tau) d\tau \right) \right\}_{ab} \\ &= \delta_{ab} \delta(\hat{\Omega}' - \hat{\Omega}) \Theta(\lambda' - \lambda) e^{-2 \int_{\lambda}^{\lambda'} \theta^{(\rm sa)}(\tau) d\tau} \end{aligned} \tag{36}$$

where Θ is the Heaviside step function and $\mathcal{T}$ exp denotes the τ-ordered exponential along the congruence (propagators evaluated at larger τ are placed to the left), and the second equality uses the explicit form of $\mathbf{A}$ from Eq. (19). The correlation propagator is given by

$$\mathcal{C}_{ab}(\hat{\Omega}_1, \lambda_1; \hat{\Omega}_2, \lambda_2) = \int d\hat{\Omega}' d\hat{\Omega}'' d\lambda' d\lambda'' \mathcal{R}_{ad}(\hat{\Omega}', \lambda'; \hat{\Omega}_1, \lambda_1) \mathcal{K}^{(2)}_{dc}(\hat{\Omega}', \lambda'; \hat{\Omega}'', \lambda'') \mathcal{R}_{bc}(\hat{\Omega}'', \lambda''; \hat{\Omega}_2, \lambda_2). \tag{37}$$

Substituting Eq. (36) into Eq. (37) yields

$$\begin{aligned} \mathcal{C}_{ab}(\hat{\Omega}_1, \lambda_1; \hat{\Omega}_2, \lambda_2) = \int_0^{\lambda_1} d\lambda' \int_0^{\lambda_2} d\lambda'' \\ \mathcal{K}^{(2)}_{ab}(\hat{\Omega}_1, \lambda'; \hat{\Omega}_2, \lambda'') e^{-2\left(\int_{\lambda'}^{\lambda_1} + \int_{\lambda''}^{\lambda_2} \right) \theta^{(\rm sa)}(\tau) d\tau}. \end{aligned} \tag{38}$$

For statistics that are spatially homogeneous and isotropic on the celestial sphere, the two-point cumulant of the driving field depends on the two directions only through their relative angle and admits the Legendre expansion

$$\mathcal{K}^{(2)}_{ab}(\hat{r}_1, \lambda'; \hat{r}_2, \lambda'') = \sum_{\ell} \frac{2\ell + 1}{4\pi} C^{ab}_{\ell}(\lambda', \lambda'') P_{\ell}(\hat{r}_1 \cdot \hat{r}_2), \tag{39}$$

where P_ℓ is the Legendre polynomial of order ℓ and $C^{ab}_{\ell}(\lambda', \lambda'')$ collects the Legendre coefficients of the corresponding correlator between the driving-field components a and b at affine parameters λ' and λ'', taken in the pair-aligned screen basis fixed in Sec. IV. Each such correlator depends only on the angle between the two directions, so the expansion is exact for the spin-2 components as well as the scalar one; the spin weight is carried by the basis, not by the expansion. Substituting this multipole decomposition into Eq. (38) yields the angular expansion of the correlation propagator,

$$\begin{aligned} &\mathcal{C}_{ab}(\hat{\Omega}_1, \lambda_1; \hat{\Omega}_2, \lambda_2) \\ &= \sum_{\ell} \frac{2\ell + 1}{4\pi} P_{\ell}(\hat{\Omega}_1 \cdot \hat{\Omega}_2) \\ &\quad \int_0^{\lambda_1} d\lambda' \int_0^{\lambda_2} d\lambda'' C^{ab}_{\ell}(\lambda', \lambda'') e^{-2\left(\int_{\lambda'}^{\lambda_1} + \int_{\lambda''}^{\lambda_2} \right) \theta^{(\rm sa)}(\tau) d\tau}. \end{aligned} \tag{40}$$

Equation (40) thus expresses the correlation propagator entirely in terms of the angular power spectra of the stochastic driving field, $C^{ab}_{\ell}(\lambda', \lambda'')$, which serve as the only statistical input to the Gaussian reference theory.

In summary, we have (i) split the action into a Gaussian reference part S_0 and a non-Gaussian interaction part $S_{\rm int}$, (ii) identified the two independent free propagators of the reference theory, $\mathcal{R}_{ab}$ and $\mathcal{C}_{ab}$, and (iii) expressed them in closed form in terms of the source-averaged expansion scalar $\theta^{(\rm sa)}$ and the angular power spectra C^{ab}_{ℓ} of the driving field. With these ingredients in hand, the package SFT-WICK [37] automates the perturbative expansion of a given observable into Feynman diagrams built from $\mathcal{R}$, $\mathcal{C}$, and evaluates the resulting integrals numerically.

C. Diagrammatic expansion of correlation functions

The perturbative expansion in Eq. (32) can be organized graphically as Feynman diagrams. Each diagram represents an integral contribution, with its propagator factors, vertex coefficients, and combinatorial weight fixed by the expansion. For an observable $\mathcal{O} = \mathcal{X}_{a_1}(z_1)\cdots\mathcal{X}_{a_p}(z_p)$, the external vertices of the diagram denote the Sachs-scalar insertions. Each edge is evaluated in the Gaussian reference theory using the free propagators derived in Eqs. (36) and (37). Diagrammatically, we draw the correlation propagator $\mathcal{C}$ as an undirected solid line, while a response contraction is drawn as a directed dashed line from a response-field endpoint to a Sachs-field endpoint and carries the factor $-\mathrm{i}\mathcal{R}$. There is no line joining two response fields, since $\langle\tilde{\mathcal{X}}_a\tilde{\mathcal{X}}_b\rangle_{S_0} = 0$. Closed directed response cycles therefore vanish by causality.

The interaction action supplies the internal vertices. Each internal vertex stands for an integration over the spacetime point at which its attached legs meet, weighted by the corresponding interaction coefficient. The deterministic Sachs nonlinearity gives the local cubic vertex

$$\mathcal{V}_F = -\mathrm{i}\int F_{abc}\tilde{\mathcal{X}}_a(z)\mathcal{X}_b(z)\mathcal{X}_c(z)dz, \tag{41}$$

so a square vertex carries one response leg and two physical legs,[1] together with the component tensor F_{abc} and an integration over the internal spacetime point. Non-Gaussianity in the driving fields produces additional cumulant vertices through the terms $-W_\varphi^{(n)}[\mathrm{i}\tilde{\mathcal{X}}]$ in Eq. (30). Equivalently, for $n \geq 3$,

$$\mathcal{V}_{\mathcal{K}_n} = -\frac{\mathrm{i}^n}{n!}\sum_{\{a_r\}}\int\prod_{r=1}^{n}dz_r \times\tilde{\mathcal{X}}_{a_1}(z_1)\cdots\tilde{\mathcal{X}}_{a_n}(z_n)\mathcal{K}^{(n)}_{a_1\cdots a_n}(z_1;\cdots;z_n), \tag{42}$$

where each index a_r labels a component of the driving field. In this notation, the interaction action of Eq. (30) collapses to the compact form $S_{\mathrm{int}} = \mathcal{V}_F + \sum_{n\geq 3}\mathcal{V}_{\mathcal{K}_n}$.

The two vertex families differ structurally. The cumulant vertex $\mathcal{V}_{\mathcal{K}_n}$ is nonlocal: its n response legs are weighted by $\mathcal{K}^{(n)}_{a_1\cdots a_n}(z_1;\ldots;z_n)$ and integrated over n independent affine-direction arguments $z_r = (\hat{\Omega}_r, \lambda_r)$. By contrast, the Sachs vertex $\mathcal{V}_F$, drawn as an F node, is local: its three legs share a single $z = (\hat{\Omega}, \lambda)$.

The diagrammatic rule is then just Wick's theorem applied to the expanded path integral. To compute a connected p-point function to a fixed order, one inserts the required number of $\mathcal{V}_F$ and $\mathcal{V}_{\mathcal{K}_n}$ vertices, contracts all physical (Sachs) and response fields with the free propagators above, drops any subgraph that is disconnected from the external points, and integrates over every internal argument. Wick contractions that have the same connectivity are collected into a single topology with the corresponding symmetry and multiplicity factors. This is the bookkeeping automated by SFT-WICK: the algebraic expression for a given topology is obtained by multiplying all propagators, tensors, cumulants, and vertex signs associated with that graph, and then summing over all internal Sachs indices.

1. Example: Two-point correlation

As a first example, consider the two-point correlation function. The zeroth-order contribution is the single free correlation line $\mathcal{C}_{ab}(z_1, z_2)$. At first order, no single vertex insertion can contribute to a connected two-point function: the response-pairing rule of the reference action, $\langle\tilde{\mathcal{X}}\,\tilde{\mathcal{X}}\rangle_{S_0} = 0$, prevents any contraction in which the response legs ($\tilde{\mathcal{X}}$) of the inserted vertex outnumber the available Sachs partners ($\mathcal{X}$), and in every other Wick channel an odd number of fields is left over. We refer the reader to [37] for a systematic catalogue of these vanishing conditions. The leading interaction correction therefore appears at second order; concretely, the second-order expansion prior to Wick contraction reads

$$\frac{1}{2}\langle\mathcal{X}_a(\hat{\Omega},\lambda)\mathcal{X}_b(\hat{\Omega}',\lambda')S_{\mathrm{int}}^2\rangle_{S_0} = \frac{1}{2}\langle\mathcal{X}_a(\hat{\Omega},\lambda)\mathcal{X}_b(\hat{\Omega}',\lambda')(\mathcal{V}_F^2 + 2\mathcal{V}_{\mathcal{K}_3}\mathcal{V}_F)\rangle_{S_0}. \tag{43}$$

On the right-hand side we have already dropped every quadratic combination of S_{int} whose Wick contractions would force a response-response pair: any product $\mathcal{V}_F\mathcal{V}_{\mathcal{K}_n}$ with $n \geq 4$ or $\mathcal{V}_{\mathcal{K}_n}\mathcal{V}_{\mathcal{K}_{n'}}$ with $n, n' \geq 3$ carries more response legs than the available Sachs partners can absorb, and so collapses through $\langle\tilde{\mathcal{X}}\,\tilde{\mathcal{X}}\rangle_{S_0} = 0$. The corresponding Wick diagrams are collected in Fig. 2: each diagram stands for the corresponding internal integrations and the sum over Sachs component labels, and entries marked "disc" are topologies whose Wick graph is disconnected, with the two external points sitting in different connected components. For brevity, hereafter we use the shorthands $F \equiv \mathcal{V}_F$ and $K \equiv \mathcal{V}_{\mathcal{K}_3}$ for the local cubic Sachs vertex and the leading nonlocal source-cumulant vertex, matching the node labels in the figures. The two surviving topologies are therefore FF (two local cubic vertices) and FK (one local F paired with one nonlocal K).

2. Example: Three-point correlation

For the three-point correlation function, the zeroth-order contribution vanishes identically. At first order, both the

[1] All legs share the same integration variable, hence the term local vertex.

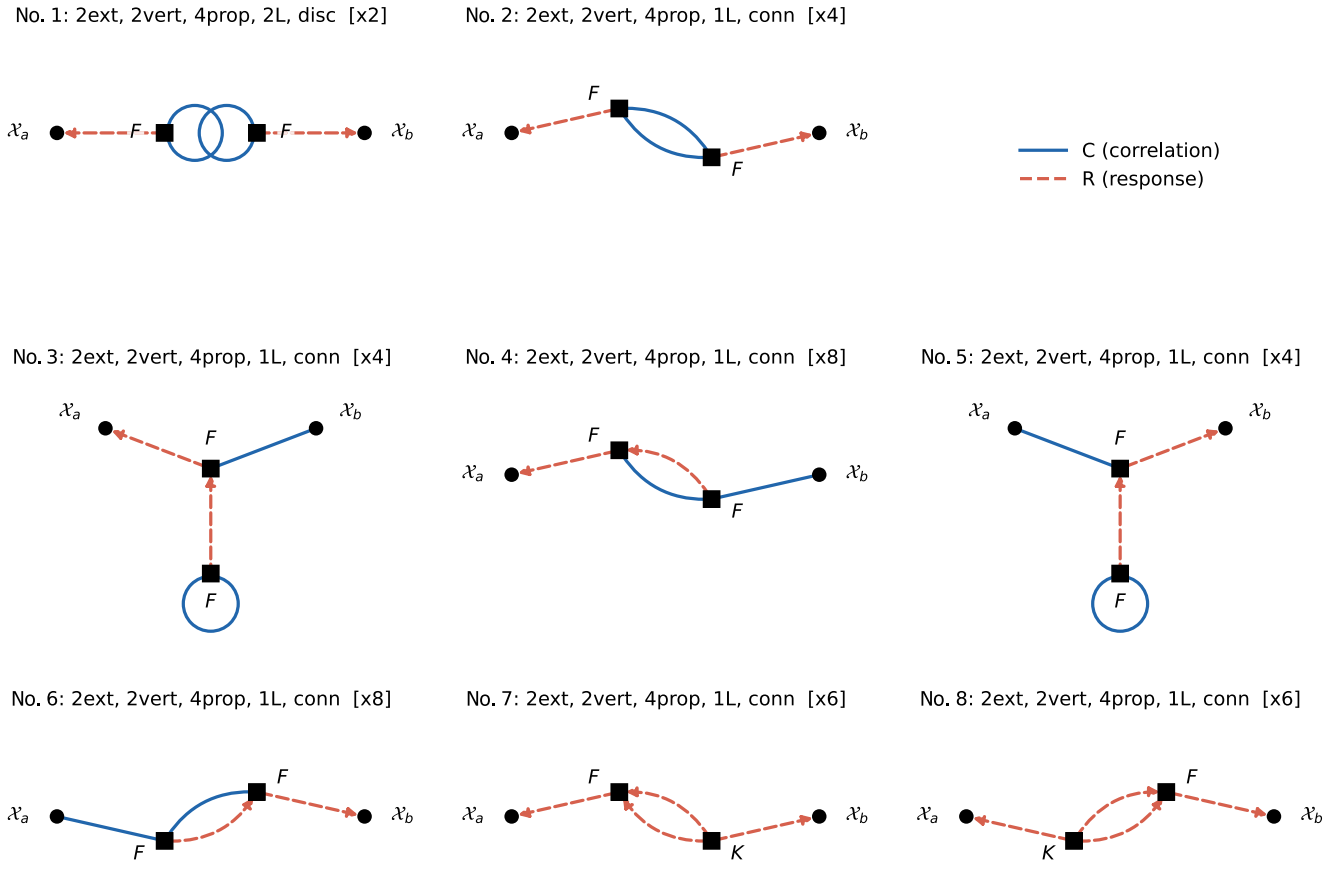


FIG. 2. Second-order diagrams for the Sachs two-point correlation function, each containing two interaction vertices. Circular external nodes denote Sachs fields. Square nodes labeled F represent the local interaction generated by nonlinear Sachs evolution [Eq. (41)]. Nodes labeled K represent the nonlocal vertex associated with the connected three-point cumulant of the Ricci and Weyl driving fields [Eq. (42) with $n = 3$]. Blue solid lines carry $\mathcal{C}$, the Sachs-field two-point correlation in the Gaussian reference theory [Eq. (37)]. Red directed dashed lines carry $-\mathrm{i}\mathcal{R}$, where $\mathcal{R}$ is the linear response propagator [Eq. (36)]. The label "disc" marks diagrams in which the two external points are not connected.

local cubic vertex $\mathcal{V}_F$ and the nonlocal cumulant vertex $\mathcal{V}_{\mathcal{K}_3}$ contribute,

$$-\langle \mathcal{X}_a(\hat{\Omega},\lambda)\mathcal{X}_b(\hat{\Omega}',\lambda')\mathcal{X}_c(\hat{\Omega}'',\lambda'')S_{\rm int}\rangle_{S_0} = -\langle \mathcal{X}_a(\hat{\Omega},\lambda)\mathcal{X}_b(\hat{\Omega}',\lambda')\mathcal{X}_c(\hat{\Omega}'',\lambda'')(\mathcal{V}_F+\mathcal{V}_{\mathcal{K}_3})\rangle_{S_0}, \quad (44)$$

whereas the higher driving field cumulant vertices $\mathcal{V}_{\mathcal{K}_n}$ with $n \geq 4$ are excluded by the same vanishing conditions invoked above. The local F vertex gives propagation-induced non-Gaussianity: the three external Sachs fields are contracted with the two Sachs legs and the response leg of the cubic vertex through the free propagators. The nonlocal K vertex instead represents the intrinsic three-point cumulant of the driving fields, propagated to the Sachs scalars by three response lines. The resulting topologies are displayed in Fig. 3.

The diagrams in this subsection are diagrams for correlation functions of the Sachs scalars themselves. The lensing observables studied in Sec. V are obtained by applying the linearized line-of-sight relations to the external Sachs insertions; the internal propagators and interaction vertices are left unchanged.

D. The selection rule

The vanishing arguments invoked in the two examples above can be collected into a single selection rule, summarized in Fig. 4. A connected diagram contributing to a p-point Sachs correlator, built from k_F insertions of $\mathcal{V}_F$ and k_n insertions of each $\mathcal{V}_{\mathcal{K}_n}$, vanishes unless

$$\sum_{n\geq 3} n k_n \leq p + k_F \quad \text{and} \quad p + k_F + \sum_{n\geq 3} n k_n \in 2\mathbb{Z}. \quad (45)$$

The first inequality says that the response legs supplied by the cumulant vertices cannot exceed the Sachs legs available to absorb them through $\mathcal{R}$; the parity condition reflects the fact that Wick's theorem can only pair fields. Both restrictions descend from the single identity $\langle \tilde{\mathcal{X}}_a \tilde{\mathcal{X}}_b \rangle_{S_0} = 0$: every response leg must find a Sachs leg. The same identity

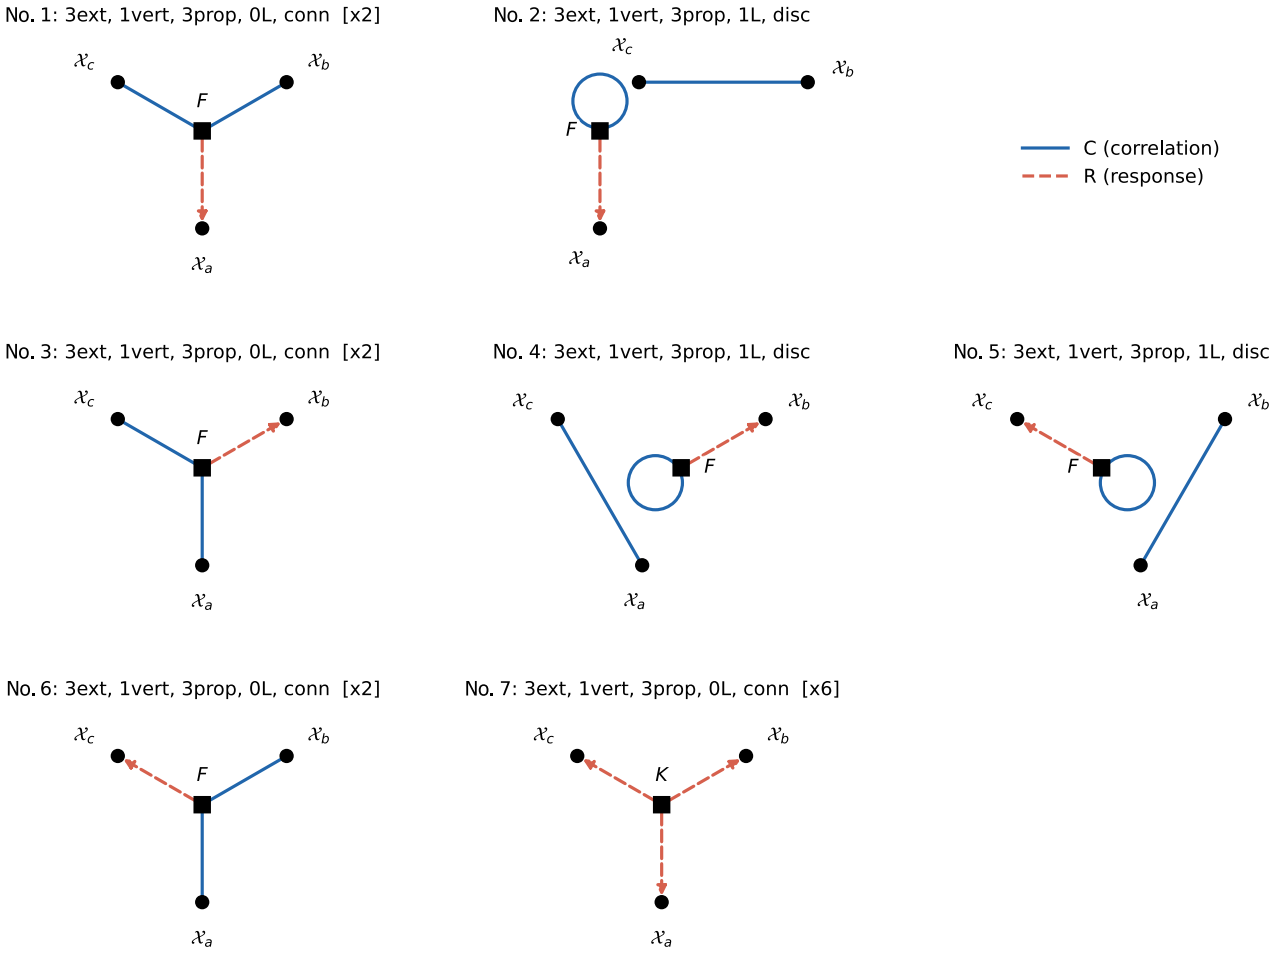


FIG. 3. First-order diagrams for the Sachs three-point correlation function generated by either one local cubic vertex $\mathcal{V}_F$ or one nonlocal source-cumulant vertex $\mathcal{V}_K \equiv \mathcal{V}_{\mathcal{K}_3}$. The diagrammatic conventions are the same as in Fig. 2.

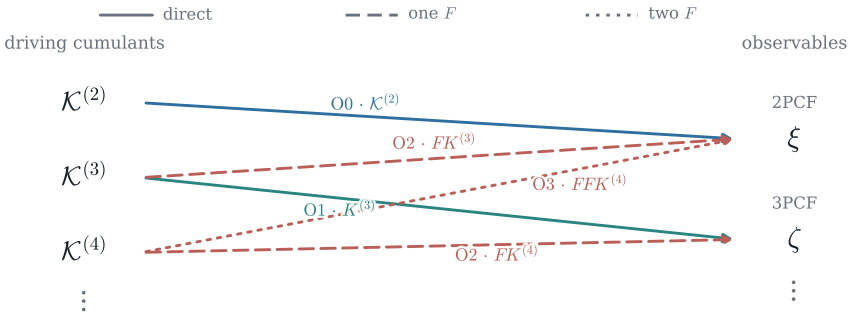


FIG. 4. Selection rule of Eq. (45) at a glance: which driving-field cumulant $\mathcal{K}^{(n)}$ feeds which observable, and at what leading order. Line style counts the nonlinear-propagation vertices F.

forbids any closed response cycle, consistent with the causality remark below Eq. (36). One can verify that Eq. (45) reproduces every vanishing case discussed above, in particular the collapse of $S^2_{\rm int}$ to $\mathcal{V}_F^2 + 2\mathcal{V}_{\mathcal{K}_3}\mathcal{V}_F$ at second order in the two-point function, and the exclusion of any first-order $\mathcal{V}_{\mathcal{K}_{n\geq 4}}$ contribution to the three-point function. We refer the reader to Zhang [37] for the generalization to correlators with external response insertions and for the enumeration of allowed topologies at arbitrary order.

IV. COSMOLOGICAL WEAK LENSING

The formalism of Secs. II–III is geometry-agnostic: the Sachs-scalar dynamics couple to an arbitrary stochastic spacetime through the Ricci focusing scalar Φ_{00} and the

Weyl shear scalar Ψ_0, and the path-integral expansion holds with no commitment to any particular background. We now specialize to a perturbed FLRW universe in Poisson gauge, a coordinate choice that extends conformal Newtonian gauge to retain vector and tensor perturbations [49]. The metric is

$$\mathrm{d}s^2 = a^2(\tau)[-(1+2\Psi)\mathrm{d}\tau^2 - 2B_i\mathrm{d}\tau\mathrm{d}x^i + ((1-2\Phi)\delta_{ij} + h_{ij})\mathrm{d}x^i\mathrm{d}x^j], \tag{46}$$

where $a(\tau)$ is the scale factor, τ the conformal time, Ψ the Newtonian gravitational potential perturbation, Φ the spatial curvature perturbation, B_i the transverse shift vector, and h_{ij} the transverse-traceless tensor perturbation. Poisson gauge removes the scalar part of the shift and the scalar and vector parts of the traceless spatial metric perturbation. The remaining shift B_i is transverse, $\partial^i B_i = 0$, and h_{ij} is transverse and traceless, $\partial^i h_{ij} = 0$ and $\delta^{ij}h_{ij} = 0$, with spatial indices raised using δ^{ij}. At first order in the metric perturbations, this form lets us compute separately the scalar, vector, and tensor contributions to the Ricci and Weyl fields that drive the Sachs equations.

Throughout this paper we adopt the flat ΛCDM fiducial cosmology used in the numerical pipeline, matching the Planck-2015 baseline [50]: $\Omega_c h^2 = 0.12029$, $\Omega_b h^2 = 0.02207$, $h = 0.6711$, $n_s = 0.97$, and $\sigma_8 = 0.81$, giving $\Omega_m = 0.31609$. These numerical choices enter the distance-redshift map, power spectra, and bispectra used in the figures and comparisons; the formal Sachs construction itself remains independent of this fiducial model. We use geometric units with $c = 1$, so that length and time are measured in a common unit and the speed of light never appears explicitly; the potentials Φ and Ψ are then dimensionless (in physical units, the usual Newtonian potentials divided by c^2), and the Hubble rate H_0 is correspondingly measured as an inverse comoving length. All four channels (Ψ, Φ, B_i, h_{ij}) are retained throughout this section; the simplifying assumptions used by the numerical analyses of Sec. V (scalar-only, $\Phi = \Psi$, Born-limit path) are not yet imposed. The treatment proceeds in three pieces.

Light-cone coordinate map (Sec. IV A). We relate the affine, line-of-sight (LOS) parametrization $(\lambda, \hat{\Omega})$ used by the Sachs dynamics to the cosmological coordinates $(\chi, \hat{r})$ and observed redshift z. The Born-limit map (i.e., the unperturbed path) collapses to an isotropic one-to-one correspondence and is used throughout the rest of the paper; the generic perturbed-path coordinate Jacobian is recorded in Appendix C for completeness.

Driving fields (Sec. IV B). Given the Poisson-gauge potentials, the Sachs sources Φ_{00} [Eq. (5)] and Ψ_0 [Eq. (6)] are evaluated to first order on the metric (46), across the scalar, vector (shift), and tensor sectors. Their connected two- and three-point cumulants enter the present work as inputs to the later analysis; the technical details of their construction and numerical implementation are presented in Ref. [51].

Free Sachs propagators (Sec. IV C). With the coordinate map and the driving-field cumulants in hand, the free Sachs response and correlation operators on the FLRW background reduce to closed forms in $\bar{D}(\lambda) = a(\lambda)\chi(\lambda)$, the background angular-diameter distance. These forms feed directly into the path-integral diagrams of Sec. V; as a consistency check, integrating the correlation operator along λ recovers the standard textbook weak-lensing window.

A. Light-cone coordinate map

The Sachs dynamics of Sec. II describe each ray by $(\lambda, \hat{\Omega})$, where λ is the affine parameter and $\hat{\Omega} \in S^2$ denotes the LOS direction in the observer rest frame. Cosmological theory and N-body simulations, however, are typically formulated in comoving coordinates $(\chi, \hat{r})$, where χ is the comoving radial distance and $\hat{r}$ the spatial direction from the observer. In practice, χ is often replaced by the corresponding background redshift. Observations, on the other hand, label points along the past light cone and therefore inherit $\hat{\Omega}$, but use the observed redshift z rather than the comoving distance χ. A complete cosmological interpretation therefore also requires the redshift evolution along each ray. In the unperturbed or Born limit, both the coordinate map and the redshift evolution reduce to isotropic one-to-one correspondences.

1. Coordinate-map Jacobian

Given the ray's spatial trajectory $\vec{x}(\lambda, \hat{\Omega})$ labeled by the affine parameter λ and the observer-today sky direction $\hat{\Omega}$, they are defined by

$$\chi(\lambda, \hat{\Omega}) \equiv |\vec{x}(\lambda, \hat{\Omega})|, \qquad \hat{r}^i(\lambda, \hat{\Omega}) \equiv \frac{x^i(\lambda, \hat{\Omega})}{\chi(\lambda, \hat{\Omega})}. \tag{47}$$

To construct the ray $\vec{x}(\lambda, \hat{\Omega})$ one needs to integrate the null geodesic equation

$$k^\mu = \frac{\mathrm{d}x^\mu}{d\lambda}, \qquad k^\nu \nabla_\nu k^\mu = 0, \tag{48}$$

from the observer with initial data fixed by $\hat{\Omega}$. Each past-directed photon that reaches the observer admits, both at the observer and along the ray, the decomposition

$$k^\mu = E(-u^\mu + n^\mu), \quad n^\mu u_\mu = 0, \quad n^\mu n_\mu = 1, \quad E = u_\mu k^\mu, \tag{49}$$

where u^μ is the four-velocity of the observer, with respect to which n^μ is the perceived four-vector of the outward line-of-sight direction, i.e., the unit spacelike vector in that observer's rest frame pointing from the observer along the

photon's spatial trajectory (with $n^\mu|_{\lambda=0}$ fixed by $\hat{\Omega}$), and E is the photon energy in observer's frame.[2]

To formulate the cosmological coordinate mapping, we identify the observer with the comoving observer at the ray event, namely the fundamental Hubble-flow observer, whose four-velocity is

$$u^\mu = (u^0, 0, 0, 0), \qquad u^\mu u_\mu = -1 \Rightarrow u^0 = \frac{1}{a\sqrt{1+2\Psi}}. \tag{50}$$

For the physically observed redshift evolution, however, the relevant observer is instead the local peculiar-flow frame intersected by the ray, allowing the tracer dynamics to manifest directly in observational cosmology.

The initial condition $k^\mu(\lambda = 0, \hat{\Omega})$ in Eqs. (49) and (50) is fixed by the observer, and integrating the null geodesic equation then determines the ray trajectory and redshift evolution. Rather than tracking the ray itself, we instead work with the 3×3 Jacobian of the coordinate map $(\lambda, \hat{\Omega}) \mapsto (\chi, \hat{r})$, which encodes the local deformation of the coordinate mapping,

$$J(\lambda, \hat{\Omega}) = \begin{pmatrix} \partial_\lambda \chi & \partial_{\hat{\Omega}} \chi \\ \partial_\lambda \hat{r} & \partial_{\hat{\Omega}} \hat{r} \end{pmatrix}. \tag{51}$$

At first order, $J = J_{(0)} + \delta J$ acquires a radial-rate correction $\delta J^\chi{}_\lambda = \partial_\lambda \delta\chi$, an angular-deflection rate $\delta J^{\hat{r}}{}_\lambda = \partial_\lambda \delta\hat{r}$, and a transverse 2×2 block $\delta J^{\hat{r}}{}_{\hat{\Omega}} = \partial_{\hat{\Omega}} \delta\hat{r}$; the explicit line-of-sight integral form of every block, and the scalar-sector reduction to the standard weak-lensing deflection and amplification matrix, are derived in Appendix C. The transverse block coincides, up to an observer-sphere $\leftrightarrow$ screen-frame rotation, with the standard Sachs Jacobi matrix $\mathcal{J}^A{}_B(\lambda)$ that maps initial observation angles to screen-space deviation vectors, and controls the familiar cosmic-shear convergence and shear. With $\hat{r} = \hat{\Omega}$ at background, the first-order corrections $\delta\chi(\lambda, \hat{\Omega})$ and $\delta\hat{r}^i(\lambda, \hat{\Omega})$ are driven by line-of-sight integrals of transverse potential gradients [schematically $\int \partial_\perp(\Phi + \Psi) d\chi'$ for $\delta\hat{r}$, i.e., the weak-lensing deflection] and take the form of the standard Sachs geodesic-deviation integrals.

We note that, here and after, any quantity X decomposes as $X = X^{(0)} + \varepsilon X^{(1)} + \mathcal{O}(\varepsilon^2)$, with the parenthesized index (0) [(1)] denoting the background [first-order] piece in the metric perturbations $\{\Psi, \Phi, B_i, h_{ij}\}$; the index may appear as a sub- or superscript depending on typographic convenience.

[2]The sign in $E = u_\mu k^\mu$ differs from the common future-directed convention. Here λ grows into the past, so $k^\tau = d\tau/d\lambda < 0$.

2. Hubble-flow redshift evolution

Redshift is an observer-dependent quantity; throughout this subsection it is defined relative to the fundamental Hubble-flow observer [four-velocity u^μ of Eq. (50)] at each ray event. The redshift of the local comoving frame is defined through

$$1 + z(\lambda) \equiv \frac{E(\lambda)}{E_0} \tag{52}$$

with $E_0 = E(\lambda = 0)$ the observer-today energy. For convenience we set $E_0 = 1$ throughout; this choice merely fixes the rescaling freedom of the affine parameter on the null geodesic and leaves every physical observable unchanged. Differentiating $E = u_\mu k^\mu$ along the ray

$$\frac{dE}{d\lambda} = k^\nu \nabla_\nu (u_\mu k^\mu) = k^\mu k^\nu \nabla_\nu u_\mu, \tag{53}$$

where the second equality uses $k^\nu \nabla_\nu k^\mu = 0$, one obtains the redshift evolution

$$\frac{dz}{d\lambda} = \frac{1}{E_0} k^\mu k^\nu \nabla_\nu u_\mu. \tag{54}$$

Inserting the Poisson-gauge metric (46), the Hubble-flow u^μ (50), and k^μ (49) into Eq. (54) and expanding to first order, the symbolic derivation yields

$$\left.\frac{dz}{d\lambda}\right|^{(1)}_{\hat{\Omega}} = \frac{E_0(1+z)^2}{a}\Big\{\mathcal{H} + \mathcal{H}[-\Psi + \hat{n}^i B_i] - \partial_\tau \Phi - \hat{n}^i \partial_i \Psi + \partial_\tau(\hat{n}^i B_i) + \frac{1}{2}\hat{n}^i \hat{n}^j \partial_\tau h_{ij}\Big\}, \tag{55}$$

where $\mathcal{H} \equiv a'/a$ is the conformal Hubble rate ($' \equiv \partial_\tau$) and all fields are evaluated along the ray of $\hat{\Omega}$ at affine parameter λ. The first term on the right-hand side is the background rate; the remaining terms collect the standard Sachs-Wolfe, Doppler-gradient, and integrated-Sachs-Wolfe contributions familiar from cosmological-perturbation theory; see, e.g., Clarkson *et al.* [52].

3. Born-limit coordinates and background redshift

In the following we adopt the Born approximation, in which first-order metric perturbations are evaluated along the unperturbed (background) null geodesic. This delivers the light-cone coordinate parametrization used throughout the rest of this section.

Coordinate-map Jacobian. A symbolic evaluation on the Poisson-gauge background yields a block-diagonal Jacobian,

$$J_{(0)} = \begin{pmatrix} E/a & 0 \\ 0 & \delta^B_A \end{pmatrix}, \qquad \det J_{(0)} = \frac{E}{a(\tau)} > 0, \tag{56}$$

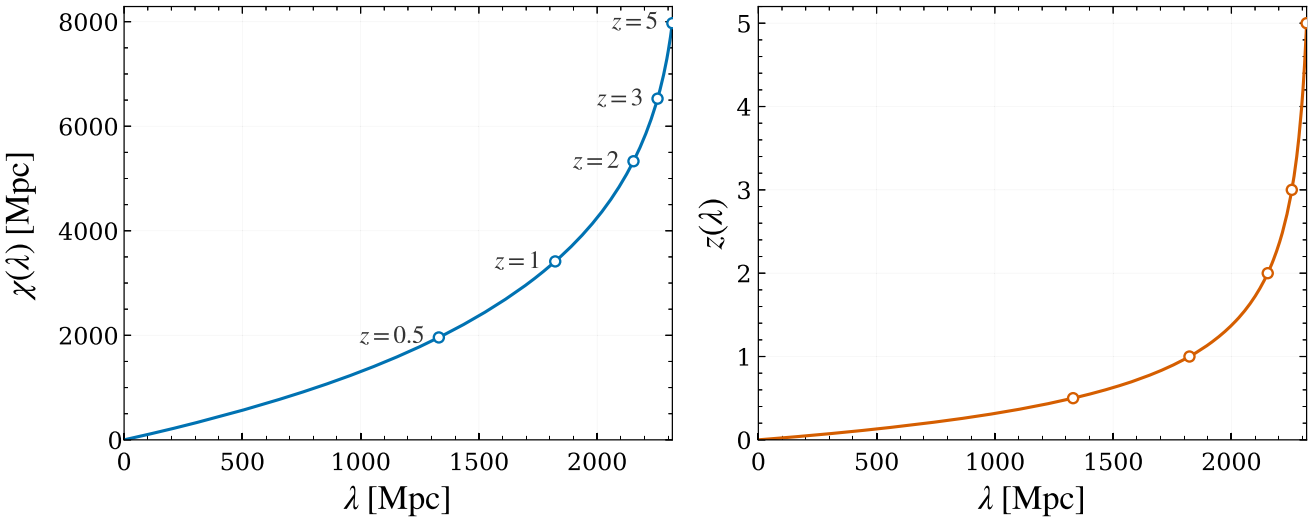


FIG. 5. Two scalar branches of the one-to-one light-cone coordinate map. Left: comoving distance $\chi(\lambda)$ obtained from Eq. (59) composed with Eq. (60). Right: redshift $z(\lambda)$ along the past light cone, the inverse of Eq. (60). Both panels share the affine-parameter λ axis; open circles mark $z \in \{0.5, 1, 2, 3, 5\}$ as orientation references. Curves use the fiducial cosmology ($\Omega_m = 0.31609$, $h = 0.6711$) adopted throughout the paper and the observer-frame normalization $E_0 = 1$ (so that λ has units of length); they reflect the exact radial grid the analyses of Sec. V consume.

whose nontrivial zeroth-order entries read

$$\left.\frac{d\chi}{d\lambda}\right|_{(0)} = \frac{E}{a(\tau)}, \qquad \hat{r}|_{(0)} = \hat{\Omega}. \tag{57}$$

Since $\det J = \det J_{(0)}(1 + \mathcal{O}(\varepsilon))$ is nowhere zero as long as the perturbations stay small, the coordinate map is a smooth diffeomorphism throughout the weak-lensing regime; caustics ($\det J \to 0$) signal the onset of strong lensing and lie outside the scope of this work. The first-order Born corrections $\delta\chi$ and $\delta\hat{r}$ are the line-of-sight integrals of transverse potential gradients derived in Appendix C.

Background redshift. For the observed redshift we adopt the background (zeroth-order) evolution, matching the convention of the general-relativistic N-body code GEVOLUTION used for crossvalidation. Dropping the first-order corrections in Eq. (55) gives

$$\left.\frac{dz}{d\lambda}\right|_{(0)} = E_0 H_c(z)(1+z)^2, \tag{58}$$

with $H_c \equiv (da/dt)/a = \mathcal{H}/a$ the cosmic-time Hubble rate (t the cosmic time, related to conformal time by $dt = ad\tau$). Dividing Eq. (57) by Eq. (58) gives the standard conformal-distance–redshift relation

$$\frac{d\chi}{dz} = \frac{1}{H_c(z)}, \qquad \chi(z) = \int_0^z \frac{dz'}{H_c(z')}, \tag{59}$$

and, equivalently, the affine parameter along the past light cone as a function of redshift,

$$\lambda(z) = \frac{1}{E_0}\int_0^z \frac{dz'}{H_c(z')(1+z')^2}. \tag{60}$$

Figure 5 shows the resulting $\chi(\lambda)$ and $z(\lambda)$ for the fiducial cosmology. The full Born-level first-order correction δz, obtained as a line-of-sight integral of Eq. (55)'s first-order right-hand side along the background ray, is deferred to future work.

B. Ricci and Weyl driving fields in Poisson gauge

Before entering the Poisson-gauge transfer functions, it is useful to fix the real component basis used for the two-point statistics. Figure 6 complements Fig. 1: the latter introduced the local screen and the optical scalars, while here we only fix the pair-aligned component convention. With $\hat{x}$ chosen along the separation of the two directions, the screen-projected optical tidal matrix $\mathcal{R}_{ab}$, $a, b \in \{\hat{x}, \hat{y}\}$, splits into three real pieces: Φ_{00} is the trace part, $\Psi_+ \equiv \mathrm{Re}\Psi_0$ is the plus-type trace-free part, and $\Psi_\times \equiv \mathrm{Im}\,\Psi_0$ is the cross-type trace-free part. Figure 6 matches these components to the source-driven rates $(\dot{\theta}, \dot{\sigma}_+, \dot{\sigma}_\times)$ that appear in the Sachs equations and to the line-segment convention used for the real and imaginary parts of Ψ_0. A reflection across the separation axis leaves the trace and plus pieces unchanged but flips the cross piece. With the convention of Eq. (5), a positive matter overdensity gives a negative Ricci-focusing source, $\Phi_{00} < 0$, which lowers $\dot{\theta}$ relative to the no-overdensity case.

This subsection records the explicit first-order form of the Ricci focusing Φ_{00} and the Weyl shear Ψ_0 on the Poisson-gauge metric (46), decomposed into scalar, vector (shift), and tensor sectors. The connected two- and three-point cumulants of $\{\Phi_{00}, \Psi_0\}$ that source the path-integral expansion of Sec. V are then obtained, by linearity, by transferring the corresponding statistics of the underlying Poisson-gauge perturbation potentials $\{\Psi, \Phi, B_i, h_{ij}\}$

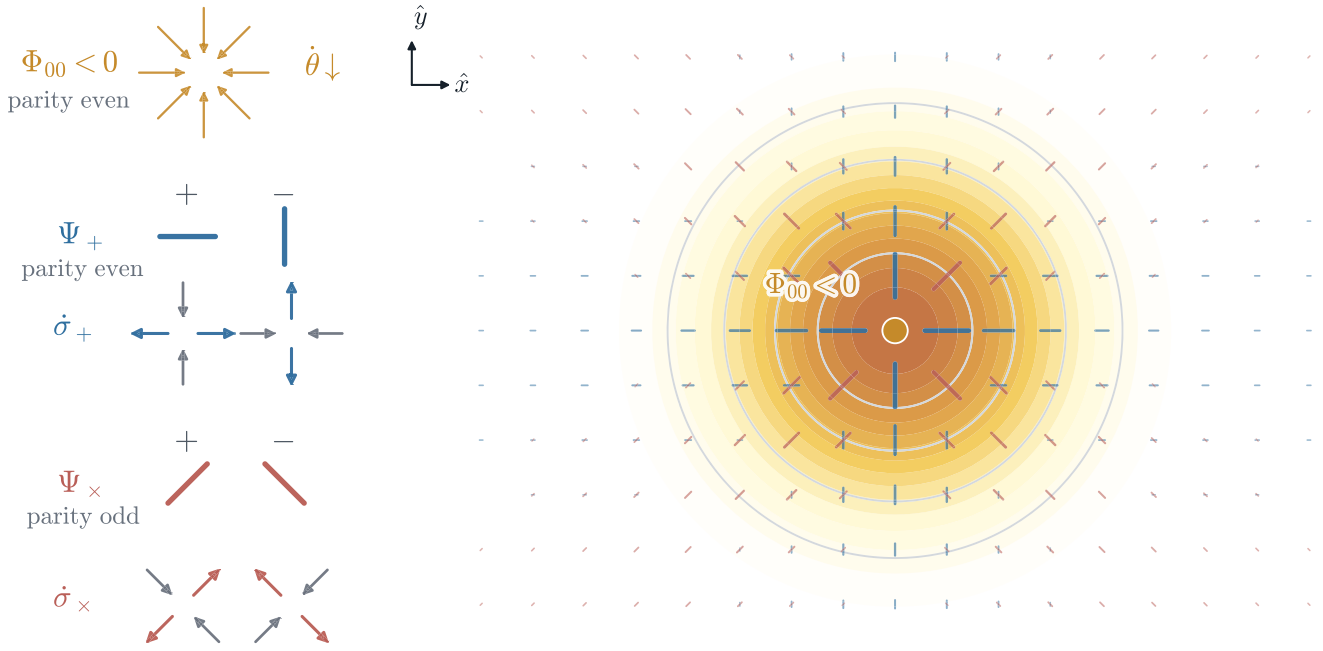


FIG. 6. Pair-aligned component dictionary and field pattern. The inset fixes the shared screen basis. The left dictionary relates the driving components to the corresponding Sachs-equation rates $(\dot{\theta}, \dot{\sigma}_+, \dot{\sigma}_\times)$; the $+/-$ marks define the line-sample sign convention; the component tags show the parity class under reflection about $\hat{x}$. Blue horizontal (vertical) segments encode positive (negative) source for $\dot{\sigma}_+$, while red $+45°$ ($-45°$) segments encode positive (negative) source for $\dot{\sigma}_\times$. The right field pattern applies this convention to an isolated focusing patch with $\Phi_{00} < 0$; line length indicates component amplitude. The dictionary arrows show rate trends, not accumulated shear shapes.

through per-multipole linear operators acting on their multipole amplitudes. Given these transfer formulas, the numerical construction of the cumulant functions of $\{\Phi_{00}, \Psi_0\}$ from the statistics of the Poisson-gauge perturbation potentials is carried out in Ref. [51]. The two-point cumulants follow directly from the power spectra of the perturbation potentials, whereas the three-point cumulants, sourced by non-Gaussianity, require additional treatment. As a concrete example, the analyses of Sec. V use a three-point cumulant built from the tree-level matter bispectrum; this model is summarized in Appendix D.

The null tetrad $\{k^\mu, e^\nu, Z^\mu, \bar{Z}^\mu\}$ and its real orthonormal underpinnings $\{\mathbf{x}^\mu, \mathbf{y}^\mu\}$ were introduced in Sec. II on a generic curved spacetime, where they are conventionally parallel-transported along the null geodesic, e.g., $k^\nu \nabla_\nu Z^\mu = 0$. However, for the first-order cosmological evaluation below, we do not need to solve the parallel-transport equation explicitly. It is enough that the tetrad is *first-order orthonormal* in the Poisson-gauge metric [Eq. (46)] around the Hubble-flow observer u^μ [Eq. (50)]: any two such bases differ by a first-order rotation $\mathbf{x}^\mu \to \mathbf{x}^\mu + \varepsilon\beta \mathbf{y}^\mu$, $\mathbf{y}^\mu \to \mathbf{y}^\mu - \varepsilon\beta \mathbf{x}^\mu$, as a result $Z^\mu \to Z^\mu e^{-\mathrm{i}\varepsilon\beta}$ and $\Psi_0 \to e^{-2\mathrm{i}\varepsilon\beta}\Psi_0$; the induced first-order shift $-2\mathrm{i}\varepsilon\beta\Psi_0^{(0)}$ vanishes because $\Psi_0^{(0)} = 0$ on the (conformally flat) FLRW background. The parallel-transported tetrad and the locally orthonormalized tetrad used below therefore give the same first-order Ψ_0. By contrast, Φ_{00} has no spin and is immune to the screen-basis rotation. First-order changes in the screen angle do not affect the first-order Ricci and Weyl driving fields used here. We do not evaluate screen-basis rotation in this demonstration.

Let $\{\hat{x}^i, \hat{y}^i\}$ be a pair of Cartesian coordinate unit three-vectors orthonormal with respect to the Euclidean metric δ_{ij} and transverse to the line of sight, i.e., $\delta_{ij}\hat{x}^i\hat{x}^j = \delta_{ij}\hat{y}^i\hat{y}^j = 1$ and $\delta_{ij}\hat{x}^i\hat{y}^j = 0$. Promoting them to four-vectors that are orthonormal in the full perturbed metric to first order in the Poisson potentials, namely enforcing $\mathbf{x}^\mu u_\mu = \mathbf{y}^\mu u_\mu = 0$, $\mathbf{x}^\mu \mathbf{x}_\mu = \mathbf{y}^\mu \mathbf{y}_\mu = 1$, and $\mathbf{x}^\mu \mathbf{y}_\mu = 0$, the screen four-vectors read

$$\mathbf{x}^\mu_{(1)} = \left(-\frac{1}{a}\hat{x}^i B_i, \frac{1}{a}\hat{x}^i + \frac{1}{a}\left(\Phi\hat{x}^i - \frac{1}{2}h^i{}_j\hat{x}^j\right)\right), \tag{61}$$

$$\mathbf{y}^\mu_{(1)} = \left(-\frac{1}{a}\hat{y}^i B_i, \frac{1}{a}\hat{y}^i + \frac{1}{a}\left(\Phi\hat{y}^i - \frac{1}{2}h^i{}_j\hat{y}^j\right)\right), \tag{62}$$

and $Z^\mu_{(1)} = (\mathbf{x}^\mu_{(1)} + \mathrm{i}\mathbf{y}^\mu_{(1)})/\sqrt{2}$ inherits the corrections through its real and imaginary parts. These corrections are not cosmetic: dropping them [i.e., using $\mathbf{x}^\mu = (0, \hat{x}^i/a)$

in the perturbed metric] would break the normalization $\mathbf{x}^\mu \mathbf{x}_\mu = 1$ and cause a spurious shift in $\Psi_0^{(1)}$, because the screen-plane rotation invariance of the previous paragraph holds only within the class of first-order-orthonormal bases.

With k^μ, u^μ, and Z^μ in hand, the driving fields defined in Eqs. (5) and (6) reduce to polynomials in the metric perturbations once $R_{\mu\nu}$ and $R_{\alpha\beta\gamma\delta}$ are expanded to first order. For later convenience, we now introduce the line-of-sight directional derivative $\hat{n}^i\partial_i$, the transverse Laplacian $\nabla_\perp^2 = (\delta^{ij} - \hat{n}^i\hat{n}^j)\partial_i\partial_j$, and the null-like differential operator

$$D \equiv \partial_\tau - \hat{n}^i\partial_i, \tag{63}$$

related to the affine-parameter derivative along the background null geodesic by $k^\mu\partial_\mu = -(E/a)D$. Tensor perturbation components projected onto the flat triad $\{\hat{n}^i, \hat{x}^i, \hat{y}^i\}$ are denoted

$$\begin{aligned} {}_0h &\equiv \hat{n}^i\hat{n}^j h_{ij}, \qquad {}_1h \equiv \hat{Z}^i\hat{n}^j h_{ij}, \\ {}_2h &\equiv \hat{Z}^i\hat{Z}^j h_{ij}, \qquad \hat{Z}^i \equiv (\hat{x}^i + \mathrm{i}\hat{y}^i)/\sqrt{2}, \end{aligned} \tag{64}$$

where the left-subscript $s = 0, 1, 2$ labels the spin weight of each projection under rotations of the screen basis $\{\hat{x}^i, \hat{y}^i\}$. Projecting onto the screen basis makes ${}_sh_{\ell m}$ coincide directly with the standard cosmological tensor-mode E/B amplitudes, facilitating polarization analyses in weak lensing. In the expressions below we use the exact relation $E = E_0(1+z)$; in the first-order prefactors, $(1+z)$ is evaluated at its background value.

Background. On the unperturbed FLRW spacetime the Weyl tensor vanishes and the Ricci tensor reduces to $R^{(0)}_{\tau\tau} = -3\mathcal{H}'$, $R^{(0)}_{ij} = (\mathcal{H}' + 2\mathcal{H}^2)\delta_{ij}$. Contracting with $k^\mu_{(0)} = (E/a)(-1, \hat{n}^i)$ gives

$$\Phi^{(0)}_{00} = -\frac{E_0^2(1+z)^2}{a^2}(\mathcal{H}^2 - \mathcal{H}'), \qquad \Psi^{(0)}_0 = 0, \tag{65}$$

so the Weyl shear is intrinsically a linear-order quantity.

Scalar sector. Keeping only Φ and Ψ, the first-order contributions to the driving fields are

$$\begin{aligned} \Phi^{(1,s)}_{00} = -\frac{E_0^2(1+z)^2}{a^2}\Big[&2(\mathcal{H}' - \mathcal{H}^2)\Psi + \mathcal{H}(\partial_\tau\Psi - \partial_\tau\Phi) \\ &- 2\mathcal{H}\hat{n}^i\partial_i\Psi + \frac{1}{2}\nabla_\perp^2(\Phi+\Psi) + D^2\Phi\Big], \end{aligned} \tag{66}$$

$$\Psi^{(1,s)}_0 = -\frac{E_0^2(1+z)^2}{a^2}\hat{Z}^i\hat{Z}^j\partial_i\partial_j(\Phi+\Psi). \tag{67}$$

Expanding $D^2\Phi = \partial_\tau^2\Phi - 2\hat{n}^i\partial_i\partial_\tau\Phi + \hat{n}^i\hat{n}^j\partial_i\partial_j\Phi$ and $\nabla_\perp^2 f = \nabla^2 f - \hat{n}^i\hat{n}^j\partial_i\partial_j f$, Eq. (66) can be rearranged into a form that isolates the well-known lensing potential $\Phi + \Psi$ and the line-of-sight projection,

$$\begin{aligned} \frac{a^2}{E_0^2(1+z)^2}\Phi^{(1,s)}_{00} = &-\frac{1}{2}\nabla^2(\Phi+\Psi) - \frac{1}{2}\hat{n}^i\hat{n}^j\partial_i\partial_j(\Phi-\Psi) \\ &- \partial_\tau^2\Phi + 2\hat{n}^i\partial_i\partial_\tau\Phi - \mathcal{H}(\partial_\tau\Psi - \partial_\tau\Phi) \\ &+ 2\mathcal{H}\hat{n}^i\partial_i\Psi - 2(\mathcal{H}' - \mathcal{H}^2)\Psi. \end{aligned} \tag{68}$$

The leading term $\frac{1}{2}\nabla^2(\Phi+\Psi)$ is the familiar Born-approximation subhorizon lensing source; the remaining contributions are subdominant on sub-Hubble scales but become relevant for integrated Sachs-Wolfe–like effects and for full-sky calculations. Equation (67) is the complex spin-2 Laplacian of the lensing potential on the screen, recovering the standard flat-sky expression $\Psi_0 = \frac{1}{2}(\mathcal{R}_{\hat{x}\hat{x}} - \mathcal{R}_{\hat{y}\hat{y}}) + \mathrm{i}\mathcal{R}_{\hat{x}\hat{y}}$ once $\hat{Z}^i = (\hat{x}^i + \mathrm{i}\hat{y}^i)/\sqrt{2}$ is expanded.

Figure 7 illustrates this scalar-sector chain at the map level. A matter overdensity first sources the scalar perturbation potential through Poisson's equation; the Ricci focusing Φ_{00} is the trace part of the resulting optical tidal field, while the Weyl scalar Ψ_0 is its spin-2 trace-free part. The same figure also shows the scale separation that underlies the weak-lensing expansion. In this snapshot, the potential has $\Phi \sim 10^{-5}$, and the displayed rescaled Ricci and Weyl fields, with the common local E^2/a^2 prefactor removed, are of order 10^{-8} Mpc^{-2}. Thus even regions that are comparatively dense in the cosmological sense still induce weak curvature driving. Away from strong-field environments such as the immediate vicinity of black holes, the dimensionless combinations built from Φ_{00} and Ψ_0 are therefore far smaller than unity, making the perturbative treatment well controlled.

Vector (shift, B_i) sector. The transverse shift vector B_i, with $\partial^i B_i = 0$, contributes to both driving fields via

$$\begin{aligned} \Phi^{(1,B)}_{00} = -\frac{E_0^2(1+z)^2}{a^2}\Big[&2(\mathcal{H}^2 - \mathcal{H}')\hat{n}^i B_i + \mathcal{H}\hat{n}^i\hat{n}^j\partial_i B_j \\ &- \frac{1}{2}\nabla^2(\hat{n}^k B_k) + \frac{1}{2}\partial_\tau(\hat{n}^i\hat{n}^j\partial_i B_j)\Big], \end{aligned} \tag{69}$$

$$\Psi^{(1,B)}_0 = \frac{E_0^2(1+z)^2}{a^2}[\hat{Z}^i\hat{Z}^j\partial_i\partial_j(\hat{n}^k B_k) + D\hat{Z}^i\partial_i(\hat{Z}^j B_j)]. \tag{70}$$

Here $D \equiv \partial_\tau - \hat{n}^i\partial_i$. The gauge condition $\partial^i B_i = 0$ removes the longitudinal part of the shift, not B_i itself, so these terms are not switched off by the gauge choice: they are the frame-dragging vector mode. Transversality does simplify the focusing, trading the transverse divergence for $-\hat{n}^i\hat{n}^j\partial_i B_j$ and leaving $\Phi^{(1,B)}_{00}$ dependent on B_i only through $\hat{n}^k B_k$, as is also true of the redshift evolution (55). The shear is not reduced this way: beside the spin-2 screen Hessian of $\hat{n}^k B_k$ it carries $D\hat{Z}^i\partial_i(\hat{Z}^j B_j)$, built from the spin-1 projection $\hat{Z}^j B_j$, which a longitudinal shift

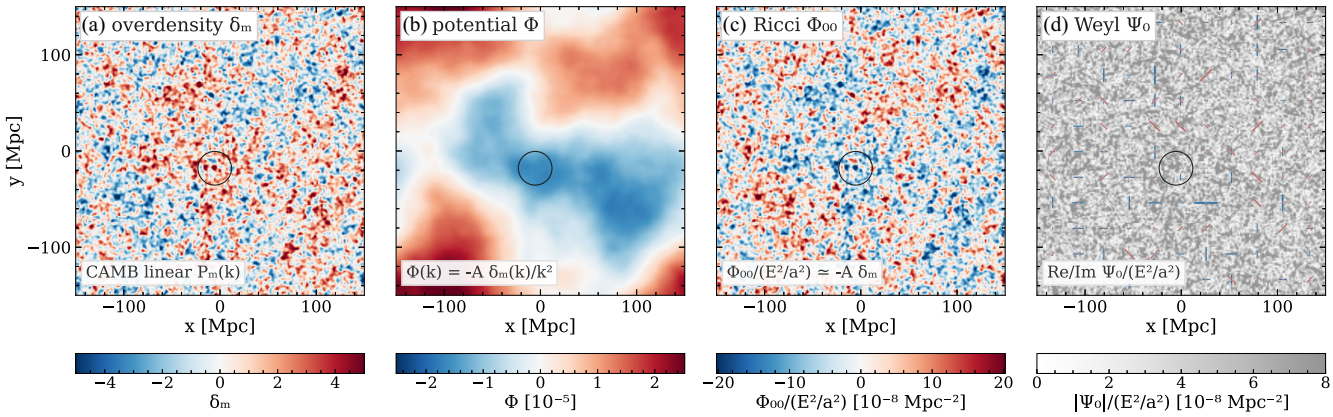


FIG. 7. Map-level relation between matter fluctuations and the Sachs driving fields in the scalar sector. Panel (a) shows a slice through a synthetic overdensity field δ_m drawn from a CAMB linear matter power spectrum at $z = 0.7$. Panel (b) shows the corresponding Poisson-gauge scalar potential $\Phi = \Psi$ for this scalar-only snapshot, obtained from $\Phi(k) = -\mathcal{A}\delta_m(k)/k^2$ with $\mathcal{A} = \frac{3}{2}\Omega_m H_0^2/a$; in the geometric units adopted here ($c = 1$) the Hubble rate is an inverse comoving length, $H_0 \simeq 2.24 \times 10^{-4}$ Mpc^{-1}, so that $\mathcal{A}$ carries dimensions of Mpc^{-2} and the dimensionless Φ follows from the dimensionless δ_m. Panel (c) shows the rescaled Ricci-focusing trace, with the common local E^2/a^2 prefactor removed, approximately $-\mathcal{A}\delta_m$; panel (d) shows the correspondingly rescaled Ψ_+ and $\Psi_\times$ from the trace-free screen Hessian of the scalar lensing potential $\Phi + \Psi$; the short blue and red line segments use the component convention illustrated in Fig. 6, with length tracking the corresponding local component amplitude. The colour bars are scaled to expose the relevant orders of magnitude in this snapshot: Φ is at the 10^{-5} level, while the displayed rescaled Ricci and Weyl fields are at the 10^{-8} Mpc^{-2} level. The black circle, identical in every panel, tracks the peak overdensity of panel (a) through the chain.

cannot supply on its own. Structurally this is Eq. (72) with the spin-2 slot absent, as a vector source requires. Vector modes have no growing linear mode in ΛCDM, so Sec. V keeps the scalar sector alone; the channel is retained because relativistic N-body output supplies it slot for slot [53].

Tensor sector. Keeping only the transverse-traceless tensor mode h_{ij} and imposing $h^i{}_i = 0$, $\partial^i h_{ij} = 0$, the contribution of gravitational waves reduces to the compact forms

$$\Phi_{00}^{(1,t)} = -\frac{E_0^2(1+z)^2}{4a^2}[\partial_{\tau_0}^2 h + 2\mathcal{H}\partial_{\tau 0} h - \nabla^2{}_0 h], \quad (71)$$

$$\Psi_0^{(1,t)} = \frac{E_0^2(1+z)^2}{2a^2}[D^2{}_2 h + 2D(\hat{Z}^i\partial_i)_1 h + (\hat{Z}^i\partial_i)^2{}_0 h]. \quad (72)$$

Equation (71) is proportional to the FLRW wave operator $\Box \equiv \partial_\tau^2 + 2\mathcal{H}\partial_\tau - \nabla^2$ acting on the line-of-sight projection of h_{ij}: for free gravitational waves obeying $\Box h_{ij} = 0$, the Ricci focusing contribution of tensor modes vanishes *"on shell."* This is the expected behavior, since gravitational waves source Weyl rather than Ricci curvature. The Weyl contribution in Eq. (72) couples the tensor polarizations ${}_2h$, ${}_1h$, ${}_0h$ through powers of the transverse complex screen gradient, $\hat{Z}^i\partial_i$. The mapping onto the standard E-mode/B-mode decomposition (equivalently, the real-space $+/\times$ polarization basis) and the on-shell reduction for a wave whose wave vector is aligned with the line of sight are recorded in [51].

Full first-order expression. Summing the three sectors, the complete first-order driving fields in the Poisson gauge are

$$\Phi_{00}^{(1)} = \Phi_{00}^{(1,s)} + \Phi_{00}^{(1,B)} + \Phi_{00}^{(1,t)}, \quad (73)$$

$$\Psi_0^{(1)} = \Psi_0^{(1,s)} + \Psi_0^{(1,B)} + \Psi_0^{(1,t)}. \quad (74)$$

Equations (65)–(74) are the explicit realizations of Eqs. (5) and (6) on the metric (46), and serve as the input to the per-multipole transfer operators that connect the angular power spectra of the Poisson-gauge perturbation fields $\{\Psi, \Phi, B_i, h_{ij}\}$ to those of the Ricci focusing Φ_{00} and the Weyl shear Ψ_0. Technical details of the implementations of those transfer operators and of the resulting two- and three-point driving-field cumulants are presented in Ref. [51].

C. Sachs propagators in cosmology and the standard weak-lensing kernel

In this section we first record the explicit form taken by the free Sachs response and correlation propagators on the flat-FLRW background. We then use that explicit form to show, from first principles, that the direct contribution of a connected n-point driving-field cumulant to linearized weak-lensing observables is its standard lensing-efficiency projection: the 2PCFs and angular power spectra inherit this projection already at Order-0 (free theory), while the single-K contribution to the 3PCFs and bispectra inherits it at Order-1. These results connect the path-integral formalism of Secs. II–III to the comoving-distance kernel

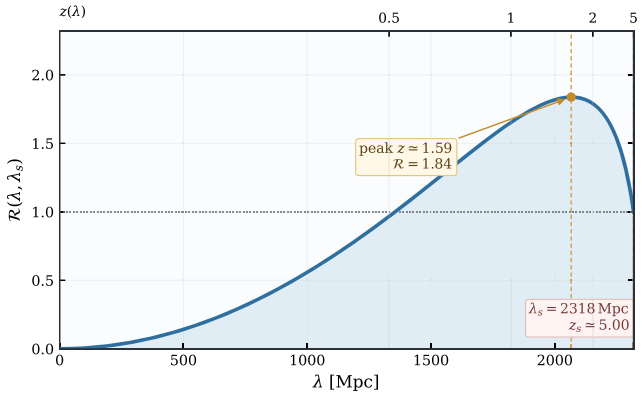


FIG. 8. FLRW response propagator [Eq. (75)] for $\lambda_s = 2318$ Mpc ($z_s \simeq 5$).

formulation for the same driving fields, without invoking the Limber approximation. Recovering the usual density-based lensing statistic also requires specializing the Ricci source to the standard transverse lensing source.

1. Response and correlation propagators

On the unperturbed FLRW background, $\Phi_{00}^{(0)}$ [Eq. (65)] sources the Jacobi equation for the background angular-diameter distance, whose closed-form solution is $\bar{D}(\lambda) = a(\lambda)\chi(\lambda)$ (Appendix B). The free response operator $\mathcal{R}$ and correlation operator $\mathcal{C}$ of Sec. II then collapse to compact forms. The response propagator is

$$\mathcal{R}_{ab}(\hat{\Omega},\lambda;\hat{\Omega}',\lambda') = \begin{cases} \delta_{ab}\delta(\hat{\Omega}-\hat{\Omega}')\left(\frac{\bar{D}(\lambda)}{\bar{D}(\lambda')}\right)^2, & \lambda < \lambda', \\ 0, & \text{otherwise.} \end{cases} \tag{75}$$

A fixed-source cut is shown in Fig. 8. The correlation propagator is given by

$$\begin{aligned} &\mathcal{C}_{ab}(\hat{\Omega}_1,\lambda_1;\hat{\Omega}_2,\lambda_2) \\ &= \int_0^{\lambda_1} d\lambda' \int_0^{\lambda_2} d\lambda'' \\ &\quad \times \left(\frac{\bar{D}(\lambda')}{\bar{D}(\lambda_1)}\right)^2 \mathcal{K}_{ab}^{(2)}(\hat{\Omega}_1,\lambda';\hat{\Omega}_2,\lambda'')\left(\frac{\bar{D}(\lambda'')}{\bar{D}(\lambda_2)}\right)^2. \end{aligned} \tag{76}$$

Representative slices of different components of this correlation propagator are shown in Fig. 9. The slices use the pair-aligned screen basis: $\hat{x}$ lies along the angular separation of the two lines of sight, while $\hat{y}$ is perpendicular to the ray-pair plane. A reflection through that plane leaves Φ_{00} and Ψ_+ even, but sends $\Psi_\times$ to minus itself, as summarized in Fig. 6. For the parity-even scalar statistics used here, components with a single $\Psi_\times$ insertion vanish in the pair-aligned full-sky continuum; any nonzero numerical residual is negligible compared with the displayed parity-even components. Equation (75) is the response propagator by which the Sachs fluctuations propagate from a source vertex at $(\hat{\Omega},\lambda)$ to a target vertex at $(\hat{\Omega}',\lambda')$; Eq. (76) is the double-λ "convolution" of the driving-field two-point cumulant through two such response kernels.

2. Convergence and shear as Sachs-scalar line-of-sight integrals

For a null geodesic labeled by $(\hat{\Omega},\lambda)$, the convergence is defined at linear order by the perturbation of the angular-diameter distance about the background trajectory,

$$D(\hat{\Omega},\lambda) = \bar{D}(\lambda)[1-\kappa(\hat{\Omega},\lambda)] + \mathcal{O}(\kappa^2), \tag{77}$$

with the sign convention $\kappa > 0 \Leftrightarrow D < \bar{D}$ of the standard weak-lensing literature. Using $\theta = \dot{D}/D$ and linearizing Eq. (77), the Sachs expansion fluctuation $\mathcal{X}_1 \equiv \theta - X_1^{(\mathrm{sa})}$ reduces to

$$\mathcal{X}_1(\hat{\Omega},\lambda) = \frac{d}{d\lambda}\left(\frac{\delta D}{\bar{D}}\right) = -\dot{\kappa}(\hat{\Omega},\lambda) + \mathcal{O}(\kappa^2), \tag{78}$$

so that integrating from the observer vertex $\lambda = 0$, where κ vanishes geometrically, gives

$$\kappa(\hat{\Omega},\lambda) = -\int_0^{\lambda} \mathcal{X}_1(\hat{\Omega},\lambda')d\lambda'. \tag{79}$$

The source-averaged reference component $X_1^{(\mathrm{sa})}$ is absorbed into $\bar{D}(\lambda)$ and drops out by construction, so κ vanishes on the reference trajectory. Its statistics are inherited entirely from those of $\mathcal{X}_1$; its ensemble mean need not vanish.

The same LOS integration applies to the spin-2 Sachs fluctuations $\mathcal{X}_{2,3} = (\sigma_+,\sigma_\times)$, the shear-rate components of the Sachs system that vanish on the background and whose integrals along the past null cone are the two real shear observables on the sky,

$$\gamma_+(\hat{\Omega},\lambda) \pm \mathrm{i}\gamma_\times(\hat{\Omega},\lambda) = -\int_0^{\lambda} [\mathcal{X}_2 \pm \mathrm{i}\mathcal{X}_3](\hat{\Omega},\lambda')d\lambda'. \tag{80}$$

Equations (79) and (80) are the linearized relations used in Sec. V to attach external κ and $\gamma_\pm$ legs to the path-integral diagrams: every linear weak-lensing observable on the sky inherits its statistics from the spin-aware Sachs three-vector $\mathcal{X}_a$ propagated through the Sachs equations.

Let $\mathcal{D}$ be the screen-space Jacobi matrix, with local Sachs matrix $\mathsf{S} = \dot{\mathcal{D}}\mathcal{D}^{-1}$ in parallel-transported screen bases [[54], Eqs. (2.13)–(2.14)]. Define $\delta\mathsf{S} = \mathsf{S} - \theta^{(\mathrm{sa})}\mathsf{I}$ and $\delta\theta = \theta - \theta^{(\mathrm{sa})}$, where $\theta^{(\mathrm{sa})} = \dot{\bar{D}}/\bar{D}$. In our shear convention,

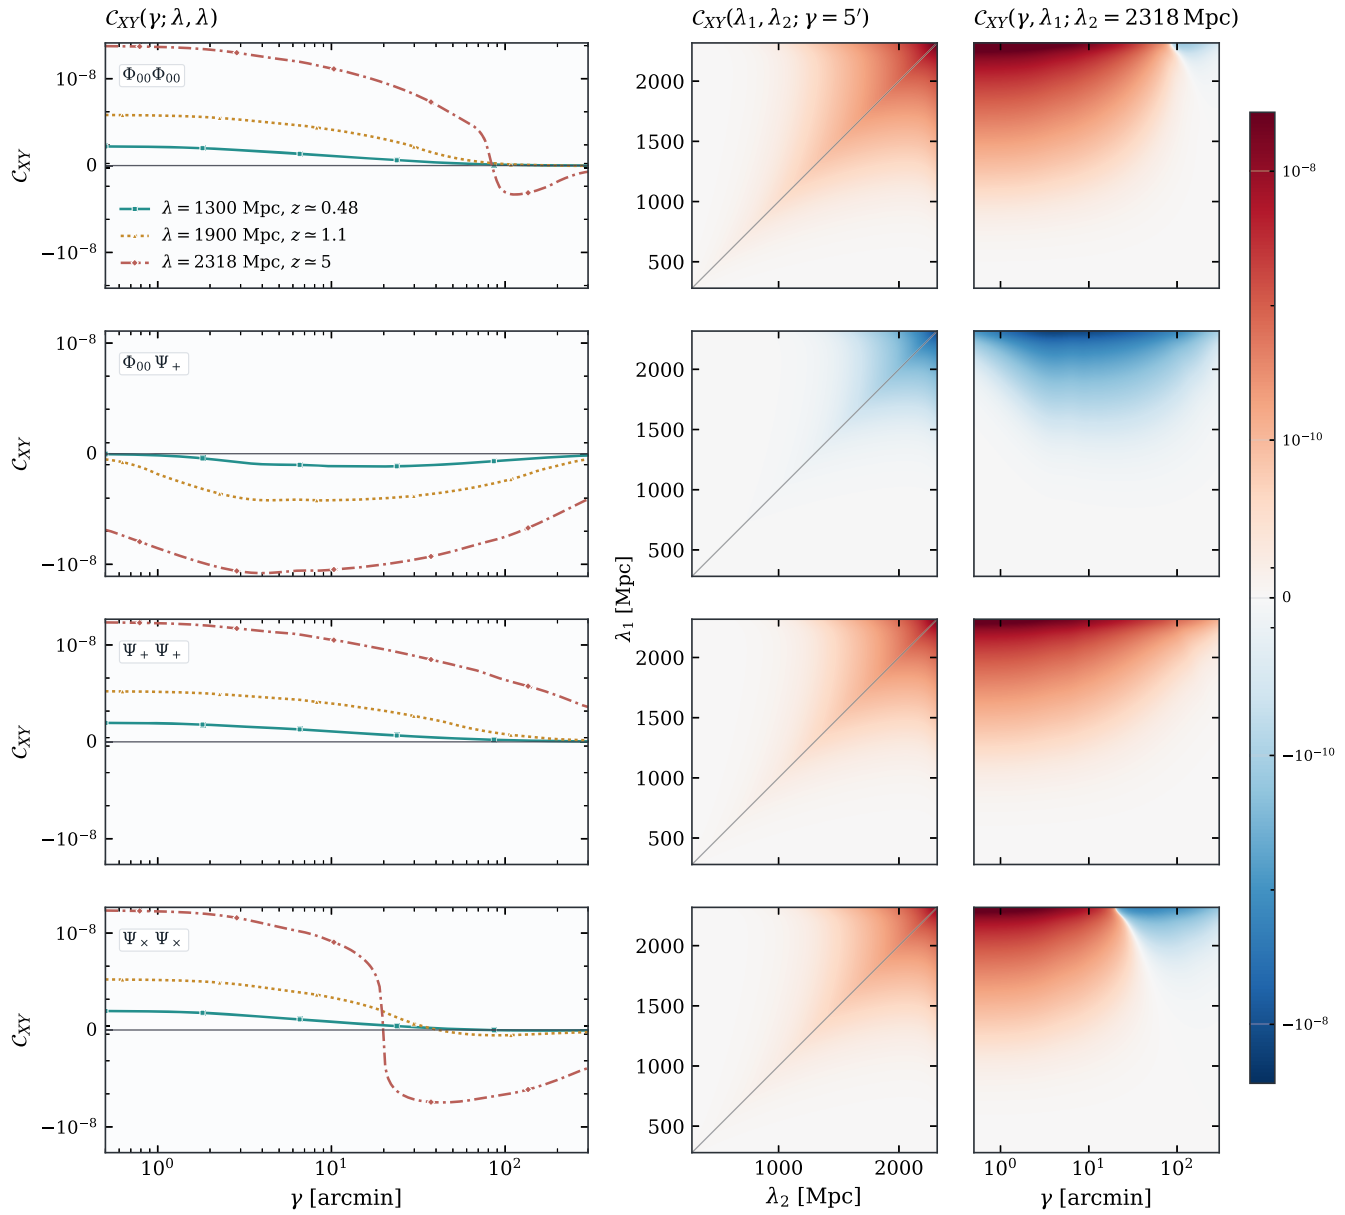


FIG. 9. Slices of different components of the correlation propagator: $\mathcal{C}_{XY}(\gamma;\lambda,\lambda)$, $\mathcal{C}_{XY}(\lambda_1,\lambda_2;\gamma=5')$, and $\mathcal{C}_{XY}(\gamma,\lambda_1;\lambda_2=\lambda_s)$. The omitted $\Phi_{00}\Psi_\times$ and $\Psi_+\Psi_\times$ components each contain one parity-odd $\Psi_\times$ insertion; any nonzero numerical residual is negligible compared with the displayed components.

$$\delta\mathsf{S}=\begin{pmatrix}\delta\theta+\sigma_+ & \sigma_\times\\ \sigma_\times & \delta\theta-\sigma_+\end{pmatrix},$$
$$\mathcal{M}=\frac{\mathcal{D}}{\bar{D}}=\begin{pmatrix}1-\kappa-\gamma_+ & -\gamma_\times-\varpi\\ -\gamma_\times+\varpi & 1-\kappa+\gamma_+\end{pmatrix}. \quad (81)$$

$$\mathcal{M}(\lambda_s)=\mathcal{P}\exp\left[\varepsilon\int_0^{\lambda_s}\delta\mathsf{S}_1 d\lambda_1\right]$$
$$=\mathsf{I}+\varepsilon\int_0^{\lambda_s}\delta\mathsf{S}_1 d\lambda_1$$
$$+\varepsilon^2\int_0^{\lambda_s}d\lambda_1\int_0^{\lambda_1}d\lambda_2\delta\mathsf{S}_1\delta\mathsf{S}_2+\mathcal{O}(\varepsilon^3). \quad (82)$$

Here ϖ denotes image rotation, distinct from optical twist ω. Writing $\delta\mathsf{S}_i=\delta\mathsf{S}(\lambda_i)$, the solution with $\lim_{\lambda\to0}\mathcal{M}=\mathsf{I}$ is

Here ε counts factors of Sachs fluctuations and is set to 1; $\mathcal{P}$ places the later matrix on the left. Although each local matrix is symmetric, noncommuting successive shears can

give the accumulated Jacobi map an antisymmetric part, hence image rotation. This is the multiple-lens effect [9]; a single scalar lens plane, or commuting shears, gives no such rotation. Our numerical example omits these quadratic and higher-order corrections to the observable map, including image rotation and symmetric convergence/shear terms, and is therefore not a complete second-order weak-lensing calculation.

The four two-point correlation functions used throughout this paper, all evaluated as functions of the angular separation $\gamma = \arccos(\hat{\Omega}_1 \cdot \hat{\Omega}_2)$ between two lines of sight and of a pair of source affine parameters (λ_s, λ_s'), are

$$\xi_\kappa(\gamma;\lambda_s,\lambda_s') \equiv \langle \kappa(\hat{\Omega}_1,\lambda_s)\kappa(\hat{\Omega}_2,\lambda_s')\rangle, \tag{83}$$

$$\xi_\pm(\gamma;\lambda_s,\lambda_s') \equiv \langle \gamma_t(\hat{\Omega}_1,\lambda_s)\gamma_t(\hat{\Omega}_2,\lambda_s') \pm \gamma_\times(\hat{\Omega}_1,\lambda_s)\gamma_\times(\hat{\Omega}_2,\lambda_s')\rangle, \tag{84}$$

$$\xi_{\kappa\gamma_t}(\gamma;\lambda_s,\lambda_s') \equiv \langle \kappa(\hat{\Omega}_1,\lambda_s)\gamma_t(\hat{\Omega}_2,\lambda_s')\rangle, \tag{85}$$

where $(\gamma_t, \gamma_\times)$ are the tangential and cross components of the shear obtained by rotating $(\gamma_+, \gamma_\times)$ of Eq. (80) into the basis aligned with the great circle connecting $\hat{\Omega}_1$ and $\hat{\Omega}_2$ [6,7]. In this pair-aligned basis we use $\gamma_t = -\gamma_+$.

3. From the correlation operator to the standard weak-lensing kernel

By construction, the correlation propagator $\mathcal{C}_{ab}$ of Eq. (76) is the free-theory two-point function of the Sachs scalars $\mathcal{X}_a$ on the past light cone. Integrating the Sachs scalars along the line of sight via Eqs. (79) and (80) then promotes that propagator to the free-reference two-point function of the weak-lensing observables themselves. As we show below, the resulting expression coincides exactly with the textbook weak-lensing efficiency kernel projection of the driving-field two-point cumulant functions.

As a concrete example, we work through the convergence 2PCF ξ_κ in detail. The first step substitutes Eq. (79) into ξ_κ, which lays the two convergence legs onto the past light cone as a pair of LOS integrals of the Sachs scalar $\mathcal{X}_1$. Averaging under the Gaussian reference action contracts the two $\mathcal{X}_1$ legs through the correlation operator of Eq. (76),

$$\xi_\kappa(\gamma;\lambda_s,\lambda_s') = \int_0^{\lambda_s} d\lambda_1 \int_0^{\lambda_s'} d\lambda_2 \mathcal{C}_{11}(\hat{\Omega}_1,\lambda_1;\hat{\Omega}_2,\lambda_2). \tag{86}$$

The second step substitutes the explicit form of $\mathcal{C}_{11}$ from Eq. (76), which exposes the real-space driving-field two-point cumulant $\mathcal{K}_{11}^{(2)}$ that ultimately gets convolved with the response propagator,

$$\xi_\kappa(\gamma;\lambda_s,\lambda_s') = \int_0^{\lambda_s} d\lambda_1 \int_0^{\lambda_s'} d\lambda_2 \int_0^{\lambda_1} d\lambda' \int_0^{\lambda_2} d\lambda'' \times \left[\frac{\bar{D}(\lambda')\bar{D}(\lambda'')}{\bar{D}(\lambda_1)\bar{D}(\lambda_2)}\right]^2 \mathcal{K}_{11}^{(2)}(\gamma;\lambda',\lambda''), \tag{87}$$

where $\mathcal{K}_{11}^{(2)}(\gamma;\lambda',\lambda'')$ is the two-point cumulant (or connected two-point correlation) of the Ricci focusing field.

The two outer integrals, over λ_1 and λ_2, act only on the response factors $\bar{D}^{-2}(\lambda_1)$ and $\bar{D}^{-2}(\lambda_2)$. Swapping the order of integration via

$$\int_0^{\lambda_s} d\lambda_1 \int_0^{\lambda_1} d\lambda' = \int_0^{\lambda_s} d\lambda' \int_{\lambda'}^{\lambda_s} d\lambda_1, \tag{88}$$

and folding the inner λ_1- and λ_2-integrals into effective weights, Eq. (87) reduces to the two-fold form

$$\xi_\kappa(\gamma;\lambda_s,\lambda_s') = \int_0^{\lambda_s} d\lambda' \int_0^{\lambda_s'} d\lambda'' \times K(\lambda',\lambda_s)K(\lambda'',\lambda_s')\mathcal{K}_{11}^{(2)}(\gamma;\lambda',\lambda''), \tag{89}$$

with effective affine-parameter kernel

$$K(\lambda,\lambda_s) \equiv \bar{D}^2(\lambda) \int_\lambda^{\lambda_s} \frac{d\lambda_1}{\bar{D}^2(\lambda_1)}. \tag{90}$$

Equation (90) is the affine-parameter form of the weak-lensing window. Substituting $\bar{D}(\lambda) = a(\lambda)\chi(\lambda)$ and the affine measure $d\lambda = a^2(\chi)d\chi$ into Eq. (90), the a^2 factors cancel inside the integrand and the affine kernel collapses to the closed form

$$K(\lambda,\lambda_s) = a^2(\chi)\chi^2 \int_\chi^{\chi_s} \frac{d\chi_1}{\chi_1^2} = a^2(\chi)\frac{\chi(\chi_s-\chi)}{\chi_s}. \tag{91}$$

The geometric factor $\chi(\chi_s-\chi)/\chi_s$ is the textbook lensing efficiency [31]; the prefactor $a^2(\chi)$ is the Jacobian that converts the affine-parameter density of the source field to a comoving-distance density. Because $\mathcal{K}_{11}^{(2)}(\lambda,\lambda')$ is essentially a density in λ, the same change of radial measure $d\lambda = a^2(\chi)d\chi$ acts on it as Jacobian factors,

$$\mathcal{K}_{11}^{(2)}(\gamma;\chi,\chi') = a^2(\chi)a^2(\chi')\mathcal{K}_{11}^{(2)}(\gamma;\lambda(\chi),\lambda(\chi')). \tag{92}$$

The same symbol on the two sides denotes the same physical correlation expressed through two different radial measures; the field Φ_{00} itself is not rescaled. Substituting into Eq. (89) and changing the integration variable from λ to χ then gives

$$\xi_\kappa(\gamma;\chi_s,\chi_s') = \int_0^{\chi_s} d\chi \int_0^{\chi_s'} d\chi' \times \frac{\chi(\chi_s-\chi)}{\chi_s}\frac{\chi'(\chi_s'-\chi')}{\chi_s'}\mathcal{K}_{11}^{(2)}(\gamma;\chi,\chi'). \quad (93)$$

This is the standard kernel projection: two lensing-efficiency windows acting on the same Ricci-focusing two-point cumulant, read in comoving-distance coordinates. For continuity with the more familiar harmonic-space formula, the Legendre transform [see Eq. (39)] of Eq. (93) gives

$$C_\ell^{\kappa\kappa}(\chi_s,\chi_s') = \int_0^{\chi_s} d\chi \int_0^{\chi_s'} d\chi' \times \frac{\chi(\chi_s-\chi)}{\chi_s}\frac{\chi'(\chi_s'-\chi')}{\chi_s'} C_\ell^{\Phi_{00}\Phi_{00}}(\chi,\chi'), \quad (94)$$

where $C_\ell^{\Phi_{00}\Phi_{00}}(\chi,\chi')$ is the harmonic projection of the same driving-field two-point cumulant on each radial slice. Equation (94) is the harmonic form of the identity proved above. For a fixed driving-field cumulant, the equivalence is exact: no Limber approximation, no equal-shell collapse, and no flat-sky simplification have been invoked. The same chain runs for every other linear weak-lensing 2PCF: ξ_+, ξ_-, $\xi_{\kappa\gamma_t}$ take the form of Eq. (89) with the appropriate spin-block of $\mathcal{K}_{ab}^{(2)}$ on the right-hand side (now including the Ψ_0 contribution that drives the shear-rate channels).

A unified view. The derivation above is not specific to the two-point case. It is the consequence of the generic ingredients: the FLRW response operator [Eq. (75)], which acts on driving field cumulants. For any n-point correlation function of the weak-lensing observables, the direct cumulant contribution in the path-integral expansion folds the connected n-point cumulant of the driving fields through n external response propagators, and the change of radial measure to χ applies on each external propagator. The general statement is therefore the following:

> The conventional lensing-efficiency-kernel formulation corresponds to the leading connected diagram associated with the corresponding driving-field cumulant.

In comoving radial coordinates the response operator $\mathcal{R}$ collapses to the textbook lensing-efficiency window $\chi(\chi_s-\chi)/\chi_s$, with one such collapse per external propagator giving the n-window projection of the same driving-field cumulant.

Two concrete cases bracket the general statement. For the weak-lensing 2PCFs, this result is already the free-theory (order-0) prediction, since the response operator covers the two-point cumulant case exactly. For the connected 3PCFs ($\langle\kappa\kappa\kappa\rangle$ and friends), the same result appears as the single-K term at Order 1 in the path-integral expansion, while the Order-0 contribution vanishes identically: a centred Gaussian free theory has no three-point function. A single-F contribution also occurs at Order 1 and generates propagation-induced non-Gaussianity, as shown in Fig. 3.

The remaining diagrams of the path-integral expansion describe nonlinear Sachs propagation; Sec. V presents illustrative examples.

V. OBSERVABLE CONSEQUENCES FOR WEAK LENSING

Expressing convergence and shear as linear integrals of Sachs-scalar fluctuations, we use the framework of Secs. II–IV to write weak-lensing correlation functions as sums of diagrams, organized along two physical axes: how nonlinearly the ray bundle propagates, carried by the cubic Sachs vertex (the F insertions), and how non-Gaussian the driving fields are, carried by their connected cumulants (the K insertions). Each diagram is labeled by how many insertions of each kind it carries, (n_F, n_K), with the total order $n = n_F + n_K$. We use the two-point correlation functions (2PCFs) as a worked example. Their lowest order, Order 0, is the free theory, with no insertion of either kind; the leading corrections to it appear at Order 2 through two channels, FF (2, 0) and FK (1, 1) (Fig. 2). Throughout we follow the four 2PCFs ξ_κ, ξ_+, ξ_- and $\xi_{\kappa\gamma_t}$ on a single source plane at $z_s = 5$, and take the driving-field statistics, built from the linear matter power spectrum and the tree-level matter bispectrum projected onto the light cone, from Ref. [51] as tabulated inputs; the three-point cumulant that seeds the FK channel is constructed in Appendix D. The matter sector is kept linear and tree-level; nonlinear-matter and perturbed-path refinements would enter by the same construction, and we leave them to future work. We take Order 0 and the two Order-2 channels in turn.

A. A unified view

Order-0 of any 2PCF is the free theory: the two external Sachs legs are paired through the Gaussian reference action S_0, with no vertex inserted. Section IV C showed analytically that, after the change of radial variable to comoving distance, this contraction reduces to the comoving double integral [Eq. (93) for the convergence].[3] This is the lensing-efficiency projection of the same driving-field cumulant, with no Limber, equal-shell, or flat-sky step. Specializing the focusing field to the standard transverse lensing source and applying the Limber approximation then gives the familiar density-based Born-Limber statistic,

$$\xi_\kappa(\gamma) = \frac{1}{2\pi}\int d\ell\,\ell P_\ell(\cos\gamma) C_\ell^{\kappa\kappa,\mathrm{BL}}, \quad (95)$$

$$C_\ell^{\kappa\kappa,\mathrm{BL}} \propto \int_0^{\chi_s} d\chi \frac{W^2(\chi)}{\chi^2} P_m((\ell+1/2)/\chi,\chi), \quad (96)$$

[3]The other three 2PCFs follow with the matching spin-0 or spin-2 two-point cumulant in place of $\mathcal{K}_{11}^{(2)}$ in Eq. (93).

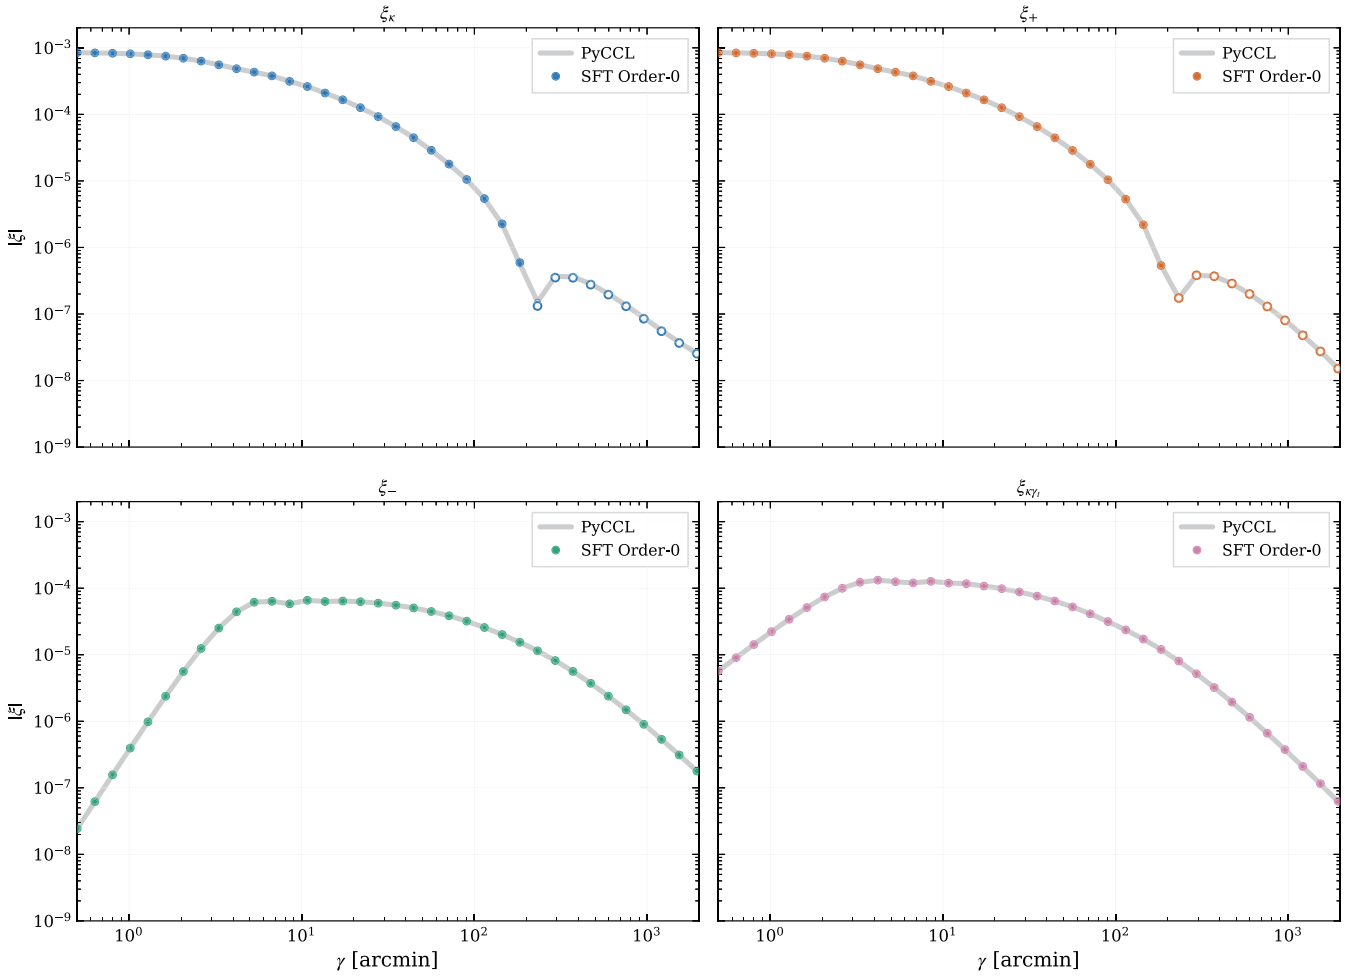


FIG. 10. Order-0 prediction compared with an independent PyCCL calculation at $z_s = 5$, using matched matter statistics, background evolution, radial support, and multipole range. The non-Limber calculation extends through $\ell = 800$, with Limber integration at higher multipoles. The panels show ξ_κ (top left), ξ_+ (top right), ξ_- (bottom left), and $\xi_{\kappa\gamma_t}$ (bottom right). Markers are filled for positive values and hollow for negative values. Grey curves show PyCCL. The Ricci-source distinction and numerical checks are described in the text.

with $W(\chi) = \frac{3}{2}\Omega_m H_0^2(1+z)\chi(\chi_s - \chi)/\chi_s$ the density-lensing window. The source specialization and the equal-shell collapse used in this expression do not change its optical Order 0. Adding F or K insertions gives the departures studied below.

Figure 10 provides an independent numerical cross-check against PyCCL [55] at $z_s = 5$. We match the linear matter power spectrum, background evolution, radial support, and multipole range. PyCCL evaluates the angular power spectra with the method of Fang, Krause, Eifler, and MacCrann (FKEM) [21] through $\ell = 800$ and uses Limber integration at higher multipoles. FKEM retains correlations between different radial distances. This choice of integration method introduces no F or K insertions and is independent of the optical expansion order.

At low multipoles, our Ricci source retains the radial and time-derivative terms in Eq. (68), whereas the standard PyCCL lensing kernel uses the transverse lensing projection. PyCCL includes the background expansion but does not retain all these additional Ricci terms. Above $\ell = 800$, our Limber implementation omits radial-derivative Ricci terms. Over the displayed angular range, the largest absolute difference is below 0.013% of the corresponding PyCCL curve's peak amplitude. Tested numerical refinements change the predicted correlations by up to 0.017% of a channel's peak. These diagnostics do not isolate the effect of the additional Ricci terms. The analytic identity in Eq. (93) concerns the projection of a fixed driving-field cumulant and holds independently of the Limber approximation.

B. Nonlinear propagation (FF)

The first correction beyond the free theory comes from nonlinear propagation alone. Switching off driving-field non-Gaussianity, so that the driving fields are Gaussian, and keeping the cubic Sachs vertex leaves a single Order-2 diagram in the 2PCF, FF $(n_F = 2)$: two propagation insertions tied together by the Gaussian two-point statistics of the driving fields. This is the result of the nonlinear

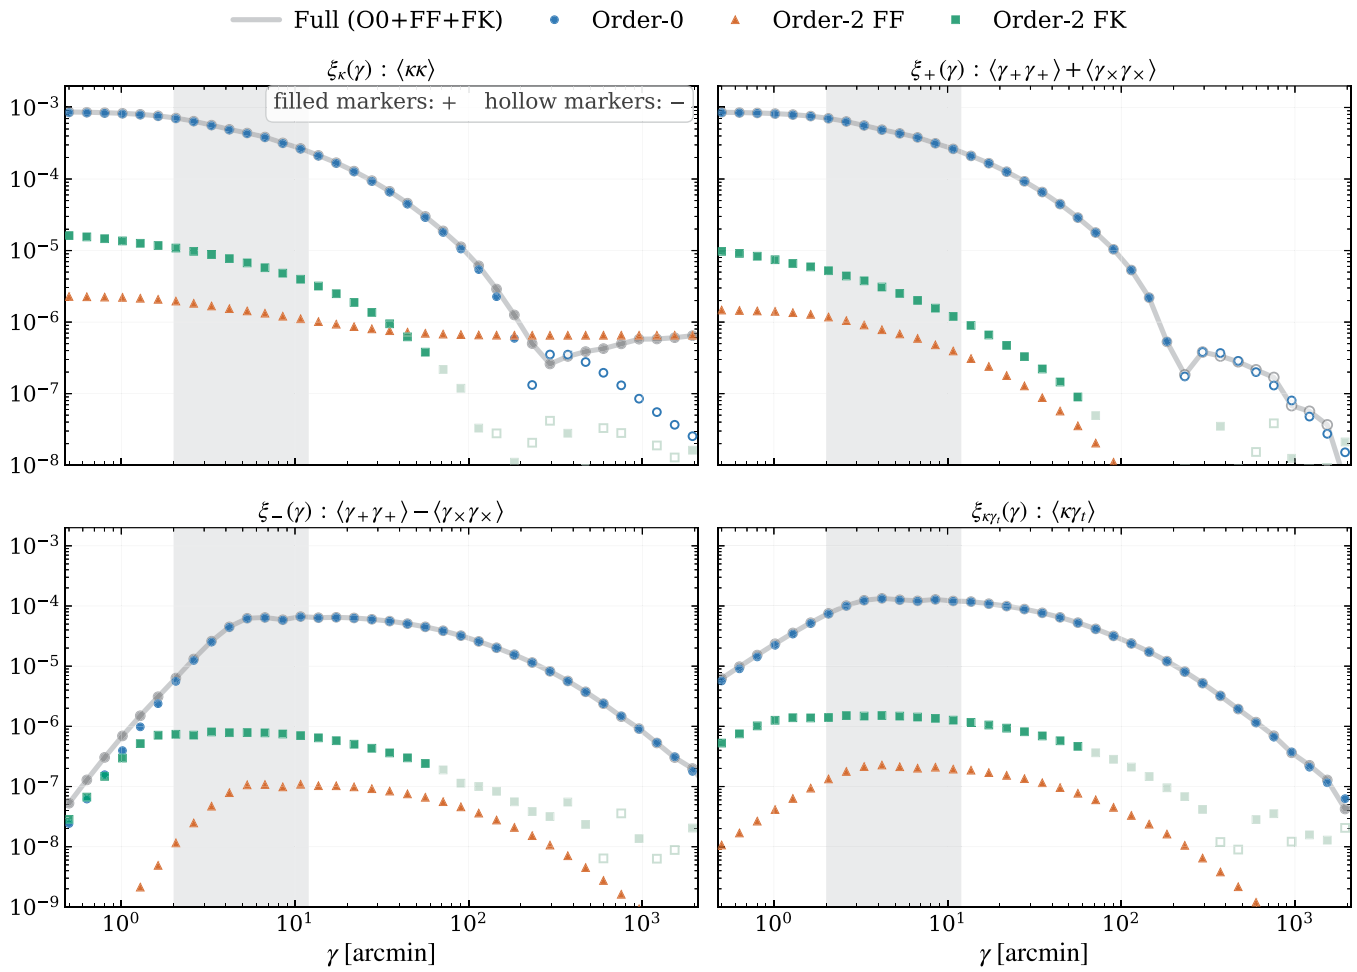


FIG. 11. Order-0 and Order-2 decomposition of the four 2PCFs at $z_s = 5$. Order-0 (blue) reproduces Fig. 10; the Order-2 correction splits into FF (orange, nonlinear propagation) and FK (green, driving-field non-Gaussianity, Sec. V C), and the grey curve is their sum. The band marks 2′–12′, the range over which FK is quoted in the text. FK increases all four correlations at the sampled points in this band; its fractional size depends on the observable. As Order-0 decays, FF supplies the large-separation plateau of the convergence moment. The FK series is drawn faint beyond $\gamma \simeq 1°$, where values are excluded from the quantitative discussion (Sec. V D). Markers are filled where positive, hollow where negative.

propagation sourced by the nonlinear terms in Sachs equations. It is a pure propagation effect, and survives even when the driving fields are perfectly Gaussian.

On the convergence (top-left of Fig. 11) FF reaches 2.28×10^{-6} at $\gamma = 0.5'$, about 0.27% of Order 0 there. It decreases toward a large-separation plateau of approximately 6.55×10^{-7}, so as the Order-0 signal falls steeply with angle the FF correction grows in relative importance toward large separations. In the sampled curve, it overtakes FK at 44′ and the absolute Order-0 signal at 183′. Figure 12 shows the same decomposition as angular power spectra, from a curved-sky transform of the four 2PCFs.[4]

[4]The Sachs vertex gives the convergence a small mean value at second order. Because ξ is defined as a moment rather than a covariance, the product of that mean along the two lines of sight is not subtracted off, and unlike a connected correlation it does not fall away with separation.

C. Leakage of driving-field non-Gaussianity (FK)

We now look at the contributions from non-Gaussian statistics of the driving fields. At second order in the Feynman diagrammatic expansion, the component with two cumulant vertices (KK) vanishes (see Fig. 2), but one nonlinear-propagation vertex F allows another vertex from the driving-field three-point cumulant $\mathcal{K}^{(3)}$. This contribution can be viewed as the leakage of driving-field non-Gaussianity into the lensing 2PCF as a result of nonlinear Sachs evolution. More generally, for this linearized observable map, the hierarchy mixing effect follows the selection rule as stated in Sec. III D: an n-point observable couples directly to the n-point driving-field cumulant, while its first hierarchy-mixing correction is sourced by the $(n+1)$-point cumulant through one nonlinear-propagation insertion (a single F vertex), and no higher cumulant can reach that order. For the 2PCF this selection rule singles out the

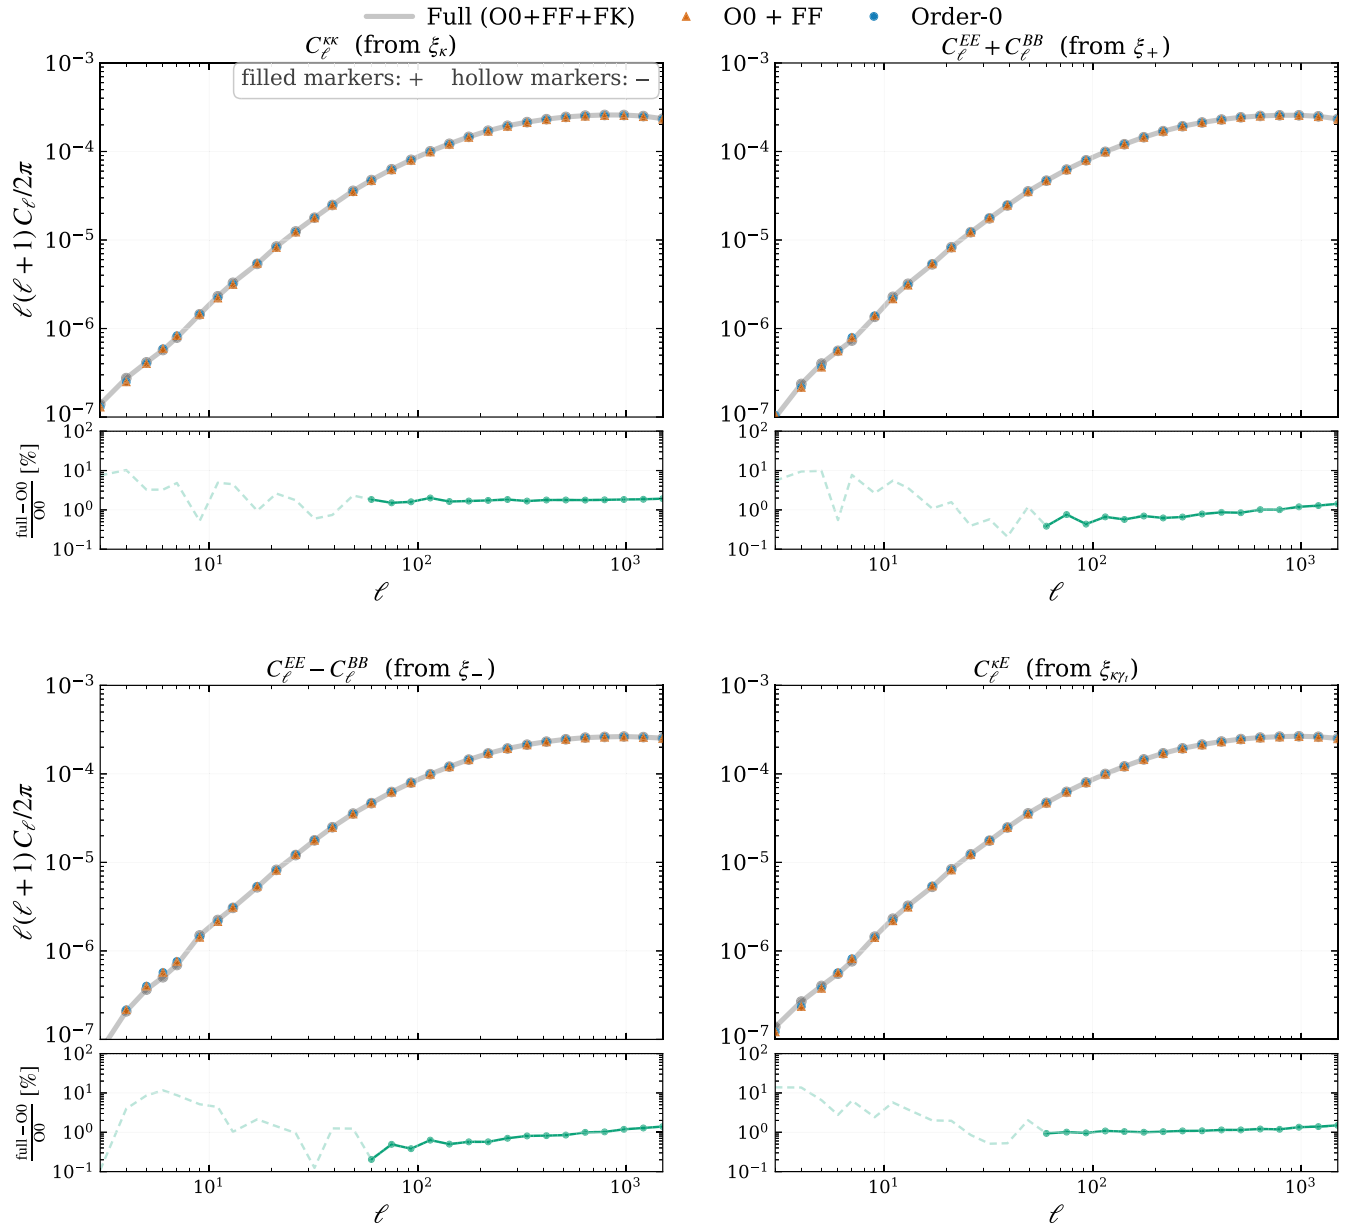


FIG. 12. Angular power spectra of the four observables of Fig. 11: the full prediction (Order-0 + FF + FK, grey), Order-0 + FF (orange), and Order 0 (blue), as the band power $\ell(\ell+1)C_\ell/2\pi$ at $z_s = 5$, from a curved-sky (Wigner-d) transform of the real-space 2PCFs; the strip below each panel shows the total Order-2 correction as a percentage of Order 0. Before transforming ξ_κ, its large-separation plateau is estimated from the largest sampled separation and subtracted. No constant is subtracted from the spin correlations. Finite angular coverage and sampling introduce transform residuals, including apparent E/B mixing, assessed in Fig. 14. Markers filled/hollow for positive/negative. Points below $\ell \sim 50$ are drawn faint and ratios are quoted for $\ell \gtrsim 50$.

three-point cumulant $\mathcal{K}^{(3)}$: the four-point cumulant and beyond are forbidden from the leading correction by the vertex counting itself, not removed by hand. In other words, which cumulants can appear is fixed by the structure of the expansion, not by any truncation.

In the numerical example, the three-point cumulant that seeds FK is represented by the equal-shell reduction $\zeta_{XYZ}(\gamma,\lambda)$ of the full cross-shell driving-field cumulant in the Limber approximation, evaluated on the collapsed (or "squeezed" if beyond the Limber approximation) configuration $(1, \cos\gamma, \cos\gamma)$ selected by the 2PCF FK diagram.[5] It is built from the tree-level matter bispectrum

[5]The propagators in the FK diagrams are all given by $\mathcal{R}$, which carries a directional δ function. Consequently, the three legs attached to a K vertex must share the directions of the two external points. In the Limber approximation, two legs sharing the same direction correspond to the same point.

[31], converted to a bispectrum of the gravitational potential by the Poisson relation and then transferred to Ricci focusing and Weyl shearing potential statistics, which are then projected onto the convergence (spin-0) and shear (spin-2) patterns seen on the sky; Appendix D gives the construction and the equal-shell reduction it relies on. Because the configuration is collapsed (or squeezed once the Limber approximation is relaxed), a long-wavelength separation couples to two local, small-scale modes, so the FK amplitude depends on the small-scale matter bispectrum entering this cumulant (Fig. 16).

How large FK is depends on each observable's spin structure. The shear is a spin-2 field, so a pair of shear legs can combine in two ways. The conjugate pairing in ξ_+ has net spin 0, while the pairing in ξ_- has spin 4. At $z_s = 5$ and $\ell_{\max} = 15360$, FK raises the convergence ξ_κ by 1.5%–1.6% and the parity-even shear ξ_+ by 0.5%–0.8% over the sampled separations within $2'$–$12'$. These correct the propagation; the mapping from (κ, γ) to a measured galaxy shape carries its own reduced-shear correction [10,12]. In the convergence-shear cross $\xi_{\kappa\gamma_t}$ it raises the signal by 1.2%–2.3% across that range. In the shear difference ξ_- the ratio is not flat: 1.1%–1.3% at the sampled points above $5'$, it climbs to 13.2% near the $2'$ edge of the band. These ratios also depend on the size of the corresponding Order-0 statistic. A difference between ξ_+ and ξ_- does not by itself imply B-mode power, since the two correlations use different angular kernels even for a pure E field.

D. Convergence in the multipole cutoff

The FK term depends on the driving-field three-point cumulant in its collapsed configuration, with two legs at one point. A coincident-point moment depends on the smallest scales retained, so the multipole cutoff $\ell_{\max}$ of the vertex table regulates the sum and its convergence has to be demonstrated rather than assumed. With the tree-level bispectrum (Fig. 13), the FK contribution to ξ_κ at $\gamma = 0.5'$, the smallest separation computed, rises from 0.36% to 1.91% of Order 0 as $\ell_{\max}$ runs from 960 to 15360; the final doubling adds 17% here, and 8.5%–10.8% over the sampled separations within $2'$–$12'$. These changes measure finite-cutoff sensitivity; they do not establish an infinite-cutoff value or error bound. Apart from this cutoff comparison, every FK result in this paper is evaluated at $\ell_{\max} = 15360$.

The angular dependence also matters when testing the cutoff. The Legendre factors in Eq. (D7) oscillate at finite separation, so cancellations can make the residual sensitive to the multipole range. The cutoff test here samples separations below $14'$. It does not establish convergence of the wide-angle tail.

With a nonlinear matter bispectrum the same sum is more sensitive to small-scale power: the BiHalofit [56] curve in Fig. 13 increases by 106% at $\gamma = 0.5'$ when $\ell_{\max}$ increases from 7680 to 15360, compared with 17% for the tree-level result. A physical nonlinear prediction depends on the assumed small-scale matter statistics and any smoothing applied to the driving fields, which is why we quote no calibrated nonlinear amplitude.

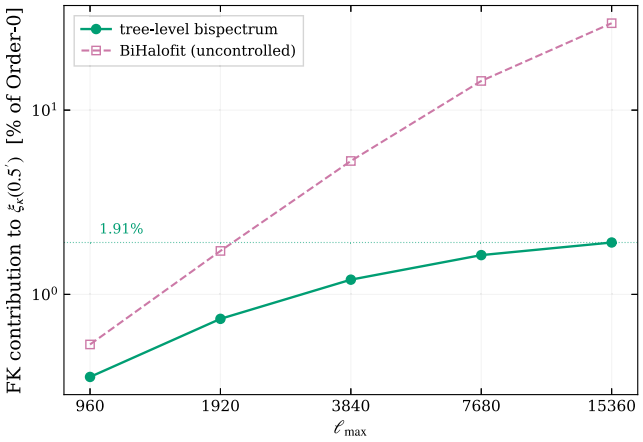


FIG. 13. FK contribution to ξ_κ at $\gamma = 0.5'$, as a fraction of Order 0, against the multipole cutoff of the vertex table. With the tree-level bispectrum, the final doubling from 7680 to 15360 adds 17%; with the BiHalofit bispectrum it adds 106%. These finite-cutoff changes do not establish an infinite-cutoff limit.

Validation. A Monte-Carlo integration of the stochastic Sachs equation tests both Order-2 channels with matched inputs (Fig. 17). The two channels need different estimators; Appendix E describes both tests and gives the numbers. The Monte Carlo is independent of the expansion, solving the Sachs system directly rather than folding it. FK agrees within 1% at four tested separations. FF agrees within 0.7% with its continuum Order-F^2 reference for the local covariance used by the simulation; this does not test the nonlocal cosmological FF prediction. A second check contracts the tabulated cumulant with the FK kernel by plain quadrature, agreeing within 1%. The expansion and the package that evaluates it are validated more broadly in Zhang [37].

Across the four 2PCFs the worked example shows a clear pattern. Order-0 reproduces the conventional kernels for matched driving-field cumulants; the leading corrections then arrive at Order-2 in two physically distinct channels, nonlinear propagation (FF) and the leakage of driving-field non-Gaussianity (FK). The four correlations respond differently to these terms, so keeping them separate preserves the origin of each correction when relating the two-point statistics to the matter model. The same structure points beyond the 2PCF, to higher-point lensing statistics, which we take up in the conclusion.

VI. CONCLUSIONS AND DISCUSSIONS

We have reformulated weak gravitational lensing as a stochastic field theory for the Sachs optical scalars of a ray bundle, driven by the random Ricci-focusing and

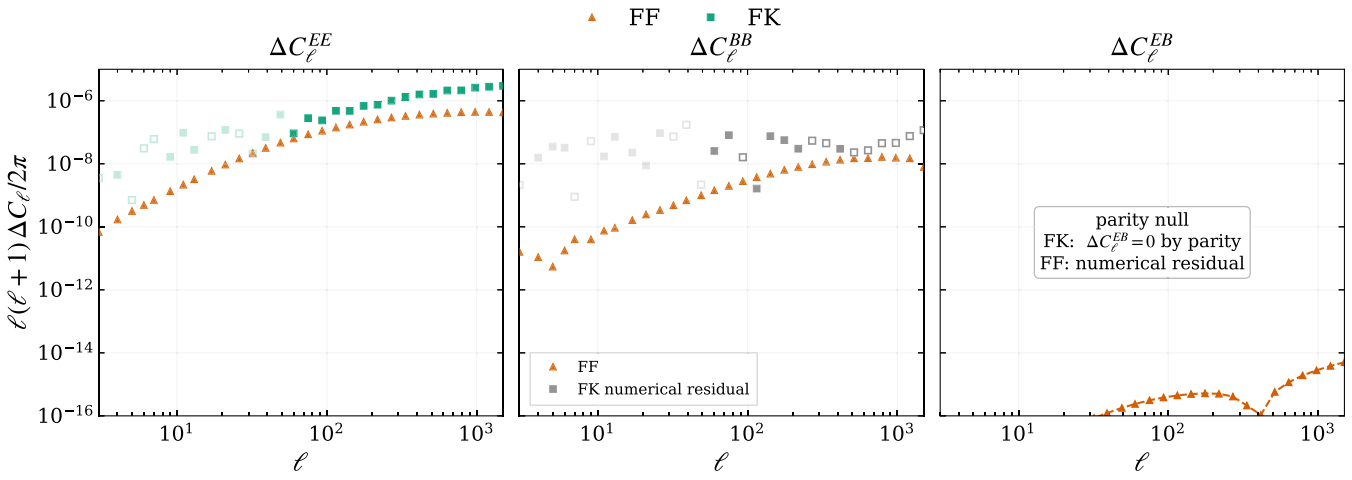


FIG. 14. Polarization decomposition of the Order-2 corrections at $z_s = 5$, shown as band powers $\ell(\ell+1)\Delta C_\ell/2\pi$. For this scalar example, FK contributes to EE while its physical BB and EB contributions vanish. The grey FK BB curve is a numerical residual. Sparse sampling over a finite angular range can mix E and B components, as confirmed by a pure-E control on the same angular grid. FF can contribute to both EE and BB, while its EB component vanishes by parity apart from a numerical residual. Fainter points below $\ell = 50$ are not used for quantitative statements. Markers are filled for positive values and hollow for negative values.

Weyl-shearing fields of the intervening matter. With the linearized observable map used here, promoting the Sachs trajectory to a path integral turns n-point functions of convergence and shear into diagrammatic expansions, ordered by how nonlinearly the bundle propagates and how non-Gaussian the driving fields are. Working the two-point correlation functions as a worked example brings out three structural insights.

First, the direct contribution of a specified driving-field cumulant projects it through the usual lensing-efficiency windows. For the two-point function this is the free theory, with no F or K insertion, and the projection identity holds independently of the Limber approximation. The standard density-based statistic additionally uses the transverse lensing source and, in its Born-Limber form, an equal-shell collapse. Optical Order 0 therefore remains the baseline whether the radial projection is evaluated with or without the Limber approximation.

Second, for this linearized observable map, a selection rule decides which driving-field cumulants can reach a given observable: only the next cumulant up the hierarchy enters the leading correction, and it does so through a single nonlinear-propagation insertion. For the two-point function this singles out the three-point cumulant, with the four-point cumulant and beyond shut out. The cutoff is structural, fixed by the diagram counting rather than imposed by hand.

Third, this mixing gives small-scale structure a route into large-angle lensing statistics. The three-point cumulant enters the two-point function in a near-degenerate configuration, two short-wavelength legs sharing a line of sight while the third follows the separation, so the contribution is sensitive to small-scale modes of the matter bispectrum at every separation. At $z_s = 5$ and $\ell_{\max} = 15360$, FK raises the convergence correlation by 1.5%–1.6% and ξ_+ by 0.5%–0.8% over the sampled separations within $2'$–$12'$, the band used in the discussion of current cosmic-shear cuts [1,2]. The other two correlations respond differently, including a positive convergence–tangential-shear correction. For convergence, the fractional FK correction increases with source redshift and exceeds FF over these sampled angles (Fig. 15). Reduced shear is a separate correction to the shape observable, rather than to propagation, and can be comparable in magnitude [10,12]. The nonlinear-propagation and non-Gaussianity corrections are physically distinct and show up in different combinations of the data, a separation that later analyses can exploit.

In this scalar example, the polarization decomposition distinguishes the two Order-2 channels. For scalar perturbations and the linearized observable map used here, the linear shear has no B component. FK pairs this linear shear with a shear containing one nonlinear Sachs insertion, so its contribution to C_ℓ^{BB} vanishes. Parity also removes C_ℓ^{EB}. FF can contribute BB at the same diagrammatic order because both external shear fields can contain a nonlinear insertion. This contribution arises from the nonlinear Sachs couplings with the linearized observable map used here. The quadratic observable reconstruction, including scalar lens-lens image rotation and the corresponding symmetric convergence/shear corrections, is not included in this calculation. Our example therefore isolates nonlinear Sachs evolution with a linearized observable map and does not provide a complete second-order weak-lensing calculation. Figure 14 shows the calculated polarization components and the numerical FK BB residual.

A further strength of the construction is its indifference to the detailed statistics of the driving fields. Strongly

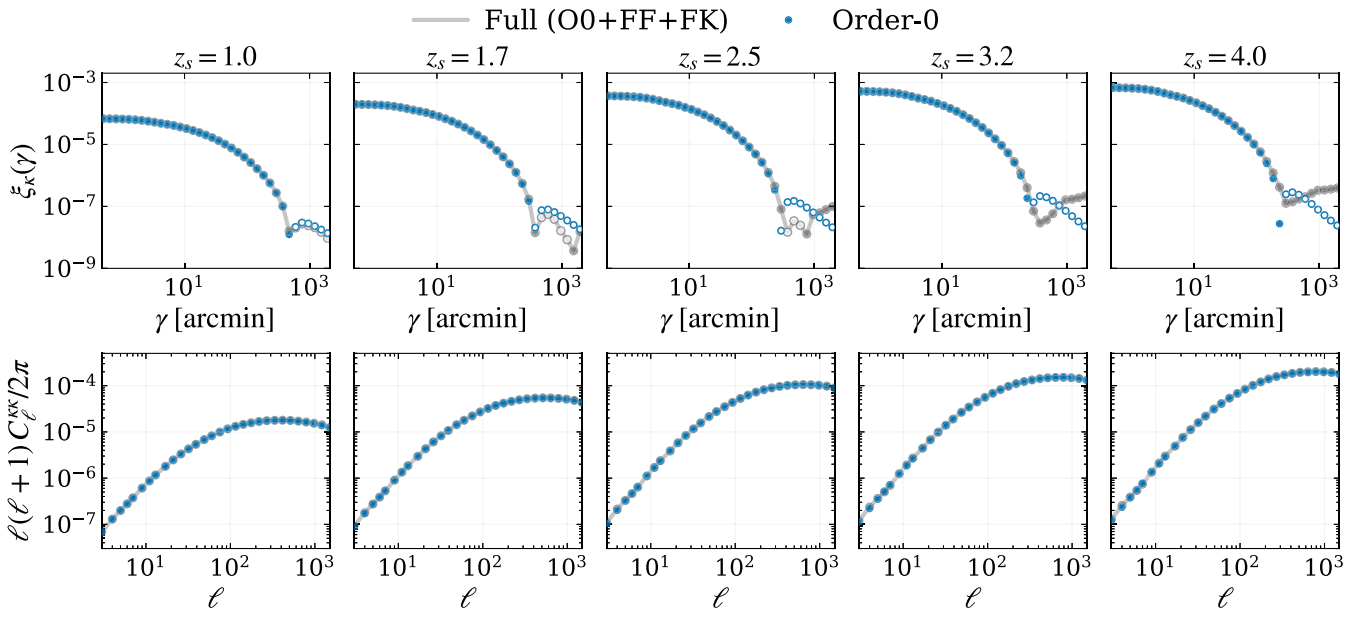


FIG. 15. Source-redshift dependence of the convergence two-point statistics. The top row shows $\xi_\kappa(\gamma)$ and the bottom row shows its curved-sky transform $C_\ell^{\kappa\kappa}$ for $z_s = 1$, 1.7, 2.5, 3.2, 4.0. Each panel compares Order-0 + FF + FK (grey) with Order-0 (blue), at $\ell_{\max} = 15360$ for FK. The fractional FK contribution increases with source redshift over the sampled separations within 2′–12′. FF supplies a large-separation plateau in the convergence moment. Before the harmonic transform, this plateau is estimated from the largest sampled separation and subtracted. Markers are filled for positive values and hollow for negative values.

correlated focusing and shearing fields and discrete, white-noise-like sources are accommodated within the same diagrammatic language, the only difference being the form of the connected cumulants fed into the vertices. For white-noise sources the Dirac-δ correlations collapse the corresponding affine and angular integrals, reducing the diagrams to closed forms with little additional effort. In this sense the formalism naturally encompasses the line-of-sight stochastic-lensing description of Fleury *et al.* [54,57], while generalizing it in two respects: it does not assume Gaussian driving fields, and it is formulated over the full observational domain rather than along a single line of sight.

Direct stochastic Sachs integrations check the two Order-2 channels at $z_s = 5$. FK uses the corrected tabulated three-point cumulant and agrees with the diagrammatic prediction within 1% at four tested angles. FF agrees within 0.7% with a continuum reference using the same local covariance as the simulation, a test distinct from the nonlocal cosmological FF prediction (Fig. 17). Independent quadrature of the FK cumulant also agrees within 1%. The sampling and numerical diagnostics are reported in Appendix E. A realistic input cumulant would call for N-body simulations with baryonic corrections, or a separate-universe treatment of the squeezed response. The simulation route is demanding precisely because of the kinematics that generate the effect: the FK leakage originates in the collapsed configuration of the cumulant, where the leg at the observed separation couples to two short-wavelength modes. Capturing it calls for a simulation box large enough to sample the separation leg together with a resolution fine enough to resolve the small-scale modes that dominate the collapsed cumulant.

These results were obtained for a deliberately narrow demonstration scenario (scalar perturbations, an unperturbed ray, a linear power spectrum, and a tree-level bispectrum), but that slice is a bookkeeping choice rather than a limitation of the formalism. In particular, the Limber reduction makes FK depend on a cumulant with two coincident legs, whose value depends on the small-scale input. A physical nonlinear prediction therefore depends on the assumed small-scale matter statistics and any smoothing applied to the driving fields. We quote no calibrated nonlinear amplitude here (Sec. V D). A further caveat concerns near-horizon scales: at the lowest multipoles a manifestly gauge-invariant observable would also need the relativistic Poisson source and the observed-redshift correction [52], fixed here at subhorizon, background level. In the opposite, wide-angle limit the collapsed cumulant is genuinely squeezed, one soft leg against two hard ones. The natural tool there is the separate-universe response of the small-scale power to a long-wavelength background [41], rather than the multipole sum used here, whose high-multipole terms alternate in sign and largely cancel at those separations (Sec. V D); that route would also carry over to the nonlinear regime, though it would still need a named scale below which the driving field is treated as unresolved. These near-horizon refinements, and a more accurate nonlinear three-point cumulant, are left to future work; the toy-model driving field already suffices to illustrate the effect. Perturbed-path remapping, nonlinear matter statistics, and

primordial non-Gaussianity all enter through the driving-field cumulants inside the same path integral, the perturbed path through the cumulants evaluated along the corrected light cone. The selection rule also points to a natural next step. Because an n-point observable in this map receives the n-point cumulant as the leading-order term, higher-point lensing statistics are cleaner windows on driving-field non-Gaussianity than the two-point function; the lensing three-point function and the one-point convergence distribution are the obvious targets. Finally, because the construction starts from the Sachs system rather than a lensing-specific projection, the same machinery extends to other light-cone observables, including the lensing of the stochastic gravitational-wave background in the geometric-optics limit [42].

ACKNOWLEDGMENTS

We are grateful to Ruth Durrer, Roy Maartens, Yan-Chuan Cai, Julien Larena, and Robert Reischke for valuable advice and discussions. Throughout the course of this work, Z. Z. has been deeply grateful for and fondly remembers the late Nick Kaiser, who brilliantly guided his PhD journey, particularly in shaping his passion for GR and field theory. Z. Z. would also like to express sincere gratitude for all the encouragement and feedback received during the "GR Effects in ΛSS 2026" workshop. This result is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant Agreement No. 948764; Z. Z., P. B.). Z. Z. also acknowledges support from the RadioForegroundsPlus Project No. HORIZON-CL4-2023-SPACE-01, GA 101135036.

DATA AVAILABILITY

The full Python toolbox and analysis pipeline (including scripts to reproduce all analyses) are publicly available as repositories [58]. The *Mathematica* scripts used for the symbolic derivations presented in this paper, based on the xAct [59] and xPerm [60] tensor-algebra packages, are also included there.

APPENDIX A: THE NEWMAN-PENROSE FORMALISM

In general relativity, the dynamics of a geodesic congruence is conventionally described either through the geodesic-deviation equation or through a tensorial decomposition of the covariant derivative of the tangent vector field. For null congruences specifically, several equivalent formalisms have been developed: one class works with the projected screen-space metric, for example the "hatting" operation in Wald's textbook, while another more prevalent in the weak-lensing literature is rendered as a second-order equation for the Jacobi map.

For the path-integral formulation pursued in this work, we require instead a system of first-order evolution equations along each null geodesic. It is also natural to express the Sachs-scalar evolution directly in the parallel-transported screen-space basis, rather than projecting from a global coordinate chart. These two requirements motivate a coordinate-free starting point: the Newman-Penrose formalism, which replaces the use of the metric with a complex null tetrad,

$$\{k^\mu, e^\mu, Z^\mu, \bar{Z}^\mu\}. \tag{A1}$$

Here k^μ is a real null vector, identified with the tangent to the photon congruence; e^μ is an auxiliary real null vector pointing in the direction "opposite" to k^μ; and Z^μ, $\bar{Z}^\mu$ are complex-conjugate vectors spanning the screen plane perpendicular to the light ray. Given two real, orthogonal, unit spacelike vectors $\mathbf{x}$ and $\mathbf{y}$[6] in the screen plane, the complex vectors are defined by

$$Z^\mu = \frac{1}{\sqrt{2}}(\mathbf{x}^\mu + \mathrm{i}\mathbf{y}^\mu), \tag{A2}$$

$$\bar{Z}^\mu = \frac{1}{\sqrt{2}}(\mathbf{x}^\mu - \mathrm{i}\mathbf{y}^\mu). \tag{A3}$$

By construction, the tetrad vectors satisfy the normalization and orthogonality conditions

$$k^\mu e_\mu = -1, \qquad Z^\mu \bar{Z}_\mu = 1, \tag{A4}$$

with all other inner products vanishing. Besides being coordinate-free, this basis makes the spin weight of a field manifest: each factor of Z raises the spin by one unit, while each $\bar{Z}$ lowers it by 1.

The dynamics of a null geodesic congruence are encoded in the covariant derivative of its tangent field, $\nabla_\nu k^\mu$. Observationally, however, only the projection of this tensor onto the screen plane is of interest. It is therefore natural to work with its contractions against the screen basis, which define the complex scalars

$$\rho = -\nabla_\nu k_\mu Z^\mu \bar{Z}^\nu, \tag{A5}$$

$$\sigma = \nabla_\nu k_\mu Z^\mu Z^\nu, \tag{A6}$$

known as the Newman–Penrose spin coefficients, with ρ and σ of spin weight 0 and 2, respectively. Their evolution along the congruence is governed by the remarkably concise system

$$\frac{d\rho}{d\lambda} = \rho^2 + \sigma\bar{\sigma} - \Phi_{00}, \tag{A7}$$

$$\frac{d\sigma}{d\lambda} = (\rho + \bar{\rho})\sigma + \Psi_0. \tag{A8}$$

[6]Typically one takes $\mathbf{x} = \hat{\mathbf{e}}_\theta$ and $\mathbf{y} = \hat{\mathbf{e}}_\phi$.

Here λ is the affine parameter along the geodesic. The right-hand sides introduce two driving fields. The first is the Ricci focusing scalar

$$\Phi_{00} = -\frac{1}{2} R_{\mu\nu} k^\mu k^\nu, \tag{A9}$$

a real-valued field built from the Ricci tensor $R_{\mu\nu}$, which drives the convergence of the beam. The second is the complex Weyl scalar

$$\Psi_0 = -C_{\alpha\beta\gamma\delta} k^\alpha Z^\beta k^\gamma Z^\delta, \tag{A10}$$

which sources the tidal shearing induced by spacetime curvature. Here $C_{\alpha\beta\gamma\delta}$ is the Weyl tensor, defined as the trace-free part of the Riemann tensor,

$$C_{\alpha\beta\gamma\delta} = R_{\alpha\beta\gamma\delta} - (g_{\alpha[\gamma} R_{\delta]\beta} - g_{\beta[\gamma} R_{\delta]\alpha}) + \frac{1}{3} R g_{\alpha[\gamma} g_{\delta]\beta}, \tag{A11}$$

where square brackets denote antisymmetrization of the enclosed indices, e.g., $A_{[\alpha\beta]} = \frac{1}{2}(A_{\alpha\beta} - A_{\beta\alpha})$.

Throughout this work we adopt natural units in which $c = 1$. Via the Einstein equations, Φ_{00} can equivalently be expressed in terms of the stress-energy tensor,

$$R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu} = 8\pi G T_{\mu\nu} \Rightarrow \Phi_{00} = -4\pi G T_{\mu\nu} k^\mu k^\nu, \tag{A12}$$

where the $R g_{\mu\nu}$ contraction vanishes because k^μ is null. In the same spirit, contracting Eq. (A11) with $k^\alpha Z^\beta k^\gamma Z^\delta$ causes the Ricci and trace terms to drop out, so that Ψ_0 reduces to its Riemann-tensor form,

$$\Psi_0 = -C_{\alpha\beta\gamma\delta} k^\alpha Z^\beta k^\gamma Z^\delta = -R_{\alpha\beta\gamma\delta} k^\alpha Z^\beta k^\gamma Z^\delta. \tag{A13}$$

Together, Eqs. (A13) and (A12) [or Eq. (A9) when direct contraction of the Ricci tensor is more convenient] provide general expressions for the driving fields that determine the evolution of weak-lensing observables. The main text quotes these definitions in Eqs. (5)–(6) before turning to their explicit cosmological realization in Sec. IV.

APPENDIX B: CLOSED-FORM JACOBI SOLUTION ON FLAT FLRW

On a flat FLRW background the mean Weyl source vanishes, and the isotropic observer conditions set the source-averaged shear to zero. The remaining expansion obeys the Riccati equation

$$\dot{\theta}^{(\mathrm{sa})}(\lambda) = -[\theta^{(\mathrm{sa})}(\lambda)]^2 + \bar{\Phi}_{00}(\lambda). \tag{B1}$$

At the observer vertex, $\theta^{(\mathrm{sa})}(\lambda) \to 1/\lambda$ as $\lambda \to 0^+$. Writing $\theta^{(\mathrm{sa})} = \dot{\bar{D}}/\bar{D}$ gives the Jacobi amplitude equation with regular vertex conditions,

$$\ddot{\bar{D}}(\lambda) = \bar{\Phi}_{00}(\lambda)\bar{D}(\lambda), \qquad \bar{D}(0) = 0, \qquad \dot{\bar{D}}(0) = 1. \tag{B2}$$

This equation admits the closed-form solution (see [61] for the expression of $\bar{D}$ as the angular diameter distance)

$$\bar{D}(\lambda) = a(\lambda)\chi(\lambda), \tag{B3}$$

where a is the scale factor and χ the comoving radial coordinate along the past null geodesic. Here the dot denotes $d/d\lambda$ (consistent with Sec. II) and $\bar{\Phi}_{00} = -(\mathcal{H}^2 - \mathcal{H}')/a^4$ is the FLRW limit of the Sachs scalar [Eq. (65)]; $\mathcal{H} \equiv a'/a$ with the prime denoting ∂_τ (conformal time).

We verify Eq. (B3) by direct substitution. With $E_0 = 1$ by construction, the affine parameter satisfies $d\tau/d\lambda = -1/a^2$ and $d\chi/d\lambda = +1/a^2$ along the radial null geodesic, so that $da/d\lambda = -\mathcal{H}/a$ and $d^2a/d\lambda^2 = -(\mathcal{H}^2 - \mathcal{H}')/a^3$ after using $a'' = a(\mathcal{H}^2 + \mathcal{H}')$. Differentiating $\bar{D} = a\chi$ twice in λ and substituting these expressions yields

$$\ddot{\bar{D}}(\lambda) = -\frac{\mathcal{H}^2 - \mathcal{H}'}{a^4}(a\chi) = \bar{\Phi}_{00}(\lambda)\bar{D}(\lambda), \tag{B4}$$

identically, with no use of the Friedmann equations: the proof is purely geometric and holds for any flat-FLRW $a(\tau)$. The vertex conditions $\bar{D}(0) = a(0)\chi(0) = 0$ and $\dot{\bar{D}}(0) = a(0)(d\chi/d\lambda)(0) = 1$ follow directly from $a(0) = 1$ and $d\chi/d\lambda(0) = 1/a^2(0) = 1$.

Two consequences enter the analysis pipeline. First, since $\theta^{(\mathrm{sa})} = d\ln\bar{D}/d\lambda$, any line-of-sight integral of this expansion scalar reduces to a difference of logs,

$$\int_{\lambda'}^{\lambda} \theta^{(\mathrm{sa})}(\tau')d\tau' = \ln\bar{D}(\lambda) - \ln\bar{D}(\lambda'). \tag{B5}$$

Second, the response propagator [Eq. (36)] reduces to the closed form

$$\mathcal{R}(\lambda, \lambda') = \Theta(\lambda' - \lambda)\left(\frac{\bar{D}(\lambda)}{\bar{D}(\lambda')}\right)^2, \tag{B6}$$

which is finite at the vertex by construction [since $\bar{D}(0) = 0$ exactly and $\bar{D}(\lambda) \sim \lambda$ as $\lambda \to 0$]. Working with $\bar{D}$ directly avoids the $1/\lambda$ integrand singularity of $\theta^{(\mathrm{sa})}$ at the observer, eliminating any quadrature error in $\mathcal{R}$ when $\bar{D}$ is tabulated from the numerical solution of the ordinary differential equation, Eq. (B2).

APPENDIX C: DETAILED FIRST-ORDER FORM OF THE 3 × 3 JACOBI MAP

This appendix records the first-order Poisson-gauge form of the past-light-cone Jacobian $J^I{}_\alpha = \partial y^I/\partial q^\alpha$, with $y^I = (\chi, \hat{r}^B)$ and $q^\alpha = (\lambda, \hat{\Omega}^A)$, used schematically in Subsec. IV A. This form keeps the three-dimensional

map explicit while showing how its transverse block reduces to the usual weak-lensing Jacobi map.

We start from the perturbed ray. Writing $k^\mu = k^{(0)\mu} + \varepsilon\delta k^\mu$ and similarly $x^\mu = x^{(0)\mu} + \varepsilon\delta x^\mu$, the null-geodesic equation $k^\nu\nabla_\nu k^\mu = 0$ splits into the background geodesic and the first-order equation

$$\frac{d\delta k^\mu}{d\lambda} = S^\mu(\tau,\vec{x}) - 2\Gamma^{(0)\mu}{}_{\nu\rho}k^{(0)\nu}\delta k^\rho,$$
$$S^\mu \equiv -\Gamma^{(1)\mu}{}_{\nu\rho}k^{(0)\nu}k^{(0)\rho}. \tag{C1}$$

The explicit Poisson-gauge form of S^μ, including the scalar potentials $\{\Psi,\Phi\}$, the transverse shift vector B_i, and the transverse-traceless mode h_{ij}, follows by substituting the linearized Christoffels $\Gamma^{(1)\mu}{}_{\nu\rho}$ and the background null tangent into Eq. (C1). Its formal solution along the background ray,

$$\delta k^\mu(\lambda) = \int_0^\lambda d\lambda' \mathcal{G}^\mu{}_\nu(\lambda,\lambda')S^\nu(\tau^{(0)}(\lambda'),\vec{x}^{(0)}(\lambda')),$$
$$\delta x^\mu(\lambda) = \int_0^\lambda d\lambda'\delta k^\mu(\lambda'), \tag{C2}$$

uses the retarded propagator $\mathcal{G}$ for the homogeneous part of Eq. (C1). In flat FLRW, $\mathcal{G}$ contributes a single factor of $1/a^2$ along the ray, and the two nested integrals in λ become the standard light-cone kernels displayed below.

The spatial perturbation then separates cleanly into radial and angular pieces. Using the background projector $P^{ij} \equiv \delta^{ij} - \hat{\Omega}^i\hat{\Omega}^j$, the spatial perturbation decomposes into

$$\delta\chi = \hat{\Omega}_i\delta x^i, \quad \delta x^i_\perp = P^i{}_j\delta x^j, \quad \delta\hat{r}^i = \frac{\delta x^i_\perp}{\chi^{(0)}} + \mathcal{O}(\varepsilon^2), \tag{C3}$$

and the four Jacobian blocks follow by differentiating these decompositions in λ (along the ray) or in $\hat{\Omega}^A$ (across observer directions). With the same convention as Eq. (51) (upper index for the output coordinate, lower index for the input coordinate), they are

$$\delta J^\chi{}_\lambda = \partial_\lambda\delta\chi = \hat{\Omega}_i\delta k^i, \tag{C4}$$

$$\delta J^A{}_\lambda = \partial_\lambda\delta\hat{r}^A = \frac{1}{\chi^{(0)}}P^A{}_i\delta k^i, \tag{C5}$$

$$\delta J^\chi{}_A = \partial_{\hat{\Omega}^A}\delta\chi = \hat{\Omega}_i\partial_{\hat{\Omega}^A}\delta x^i, \tag{C6}$$

$$\delta J^B{}_A = \partial_{\hat{\Omega}^A}\delta\hat{r}^B = \frac{1}{\chi^{(0)}}(P^B{}_i\partial_{\hat{\Omega}^A}\delta x^i - \delta^B_A\delta\chi). \tag{C7}$$

The 2×2 transverse block (C7) coincides, up to the rotation from the observer-sky basis to the screen basis, with the standard Sachs Jacobi matrix $\mathcal{J}^A{}_B(\lambda)$. This is the block that carries the usual weak-lensing convergence and shear.

For the scalar part, substituting the Poisson-gauge Christoffels and setting the shift and tensor modes ($B_i = 0$, $h_{ij} = 0$), the source S^μ reduces to

$$S^\tau|_{\rm scalar} = \frac{E^2}{a^2}[2\mathcal{H}(\Phi+\Psi) + 2\hat{\Omega}^i\partial_i\Psi + \partial_\tau(\Phi-\Psi)], \tag{C8}$$

$$\hat{\Omega}_iS^i|_{\rm scalar} = \frac{E^2}{a^2}[\hat{\Omega}^i\partial_i(\Phi-\Psi) - 2\partial_\tau\Phi], \tag{C9}$$

$$P^A{}_iS^i|_{\rm scalar} = -\frac{E^2}{a^2}\partial^A_\perp(\Phi+\Psi). \tag{C10}$$

Equation (C8) is the Sachs-Wolfe plus integrated-Sachs-Wolfe integrand that drives $\delta\tau$, the radial term (C9) gives the gauge-consistent Shapiro shift of $\delta\chi$, and the transverse term (C10) is the familiar gradient of the lensing potential $\Phi+\Psi$ that sources the weak-lensing deflection.

The weak-lensing limit follows from the transverse source. Integrating Eqs. (C2) and (C10) twice along the background ray, converting from λ to χ via Eq. (57), and dividing by $\chi_s = \chi^{(0)}(\lambda_s)$, one recovers the standard transverse deflection and its sky Jacobian,

$$\delta\hat{r}^A(\hat{\Omega},\chi_s) = -\int_0^{\chi_s}\frac{\chi_s-\chi}{\chi_s}[\partial^A_\perp(\Phi+\Psi)](\tau^{(0)}(\chi),\chi\hat{\Omega})d\chi, \tag{C11}$$

$$\delta J^B{}_A(\hat{\Omega},\chi_s) = -\int_0^{\chi_s}\frac{(\chi_s-\chi)\chi}{\chi_s} \times [\partial^B_\perp\partial^\perp_A(\Phi+\Psi)](\tau^{(0)}(\chi),\chi\hat{\Omega})d\chi. \tag{C12}$$

From Eq. (C12) the convergence and shear of the beam bundle are

$$\kappa = -\frac{1}{2}\delta J^A{}_A, \qquad \gamma_1 \pm i\gamma_2 = -\frac{1}{2}(\delta J^1{}_1 - \delta J^2{}_2) \mp i\delta J^1{}_2, \tag{C13}$$

matching the standard CMB/large-scale-structure weak-lensing literature and, by construction, the standard Sachs Jacobi map.

The same construction extends to the nonscalar perturbations. The shift-B_i and tensor-h_{ij} contributions enter through S^μ, while Eqs. (C4)–(C7) remain unchanged as block identities. Equivalently, only the source terms in Eqs. (C8)–(C10) gain additional B_i and h_{ij} terms.

APPENDIX D: THE DRIVING-FIELD THREE-POINT CUMULANT

This appendix builds the driving-field three-point cumulant that sources the FK term, in three steps. We first reduce

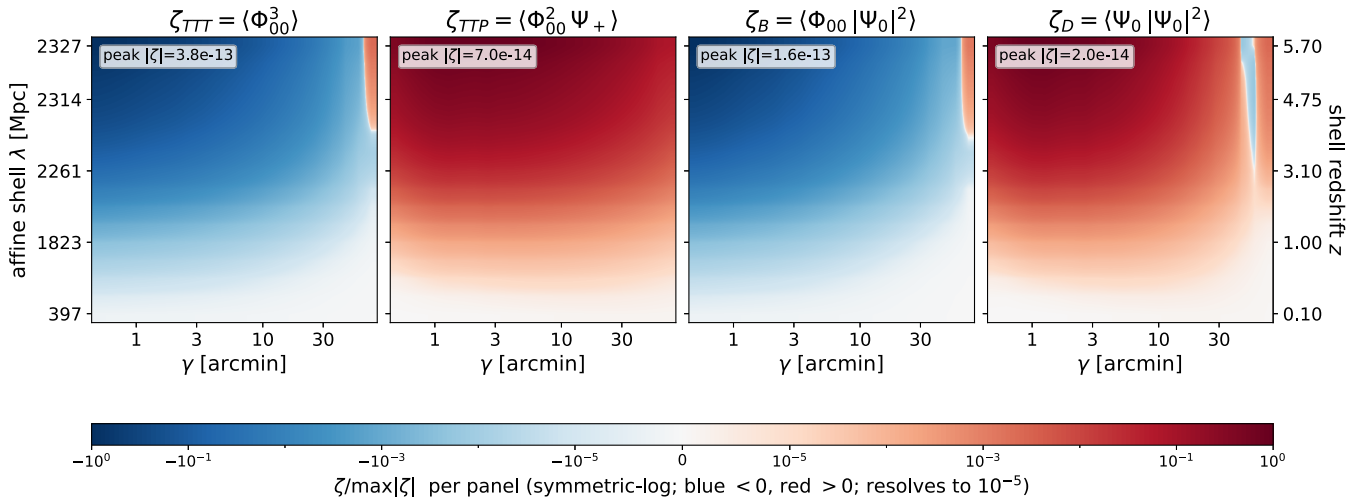


FIG. 16. Equal-shell driving-field three-point cumulants in the collapsed configuration $(1, \cos\gamma, \cos\gamma)$. The panels show ζ_{TTT}, ζ_{TTP}, and the modulus channels $\zeta_B = \langle\Phi_{00}|\Psi_0|^2\rangle$ and $\zeta_D = \langle\Psi_0|\Psi_0|^2\rangle$ defined in Eq. (D10). The channels ζ_{TPP} and ζ_{PPP} are not shown. Each panel is normalized to its own peak absolute amplitude and uses the same symmetric-logarithmic color scale. Rows correspond to the 16 driving shells, labeled by affine distance and shell redshift. In the higher-redshift shells, TTT and B change sign between 56.27′ and 71.26′, while D changes sign between 44.43′ and 56.27′. TTP has no sign change over the displayed range. The input table uses $\ell_{\rm max} = 15360$.

the full object to the single quantity the two-point expansion actually consumes, then construct that quantity from the present-day matter power spectrum, and finally resolve the one spin-2 subtlety the reduction exposes. We introduce the equal-shell approximation below.

Start from the most general object the formalism produces, the cross-shell three-point cumulant $\mathcal{K}^{(3)}_{abc}(\hat{\Omega}_1, \hat{\Omega}_2, \hat{\Omega}_3; \lambda_1, \lambda_2, \lambda_3)$, a function of three directions and three affine labels. Its real driving-field components $a, b, c \in \{\Phi_{00}, \Psi_+, \Psi_\times\}$ are the spin-0 Ricci focusing Φ_{00} and the real and imaginary parts of the spin-2 Weyl shear Ψ_0 (Sec. IV B). What enters the *two-point* expansion is far simpler, and the FK diagram (Sec. V C) shows why. There the nonlocal K-vertex is folded with response propagators, and each propagator carries a directional δ-function that pins a vertex leg to an external line [Eq. (36)]. Two of the three vertex directions are thereby tied to the two observed directions, which differ by a single angle γ. The object we actually need is therefore the one-angle, equal-shell reduction $\zeta_{abc}(\gamma, \lambda)$: two directions coincident, the third offset by γ, all on a common source shell λ. Although the statistic we ultimately need is the one carried in the affine-parameter (λ) measure that the path integral folds over, the same quantity is most naturally computed first in comoving space, as a function of χ, and only then converted into the corresponding λ-measure form by applying the correct Jacobian.

Collapsing the three affine labels onto one shell is the three-point counterpart of the Limber approximation. As in two-point lensing projections [20–22], it is motivated by lensing response kernels that are broad compared with the radial coherence length of the matter correlators. Explicitly, the collapse replaces the cross-shell cumulant by its diagonal,

$$\mathcal{K}^{(3)}_{abc}(\gamma; \lambda_1, \lambda_2, \lambda_3) \approx \zeta_{abc}(\gamma, \lambda_1)\delta(\lambda_1 - \lambda_2)\delta(\lambda_1 - \lambda_3), \quad \text{(D1)}$$

turning the three radial integrals of the cross-shell vertex into one and making the three-point cumulant analysis far easier to implement with the CANOES code [51].

From the matter spectrum to ζ. The cumulant follows from a single time-independent input, the present-day matter power spectrum $P_\delta(k)$, through the chain $P_\delta \to B_\delta \to B_\Phi \to (\Phi_{00}, \Psi_0) \to \zeta$. At tree level the matter bispectrum is the symmetrized second-order coupling of standard perturbation theory (e.g., [31]),

$$B_\delta(k_1, k_2, k_3) = 2F_2^{(s)}(\mathbf{k}_1, \mathbf{k}_2)P_\delta(k_1)P_\delta(k_2) + 2\,\text{cyclic}, \quad \text{(D2)}$$

with $F_2^{(s)}$ the usual mode-coupling kernel; the tree-level redshift evolution is carried entirely by the growth factor,

$$B_\delta(z) = D(z)^4 B_\delta(0). \quad \text{(D3)}$$

The Poisson relation turns each matter leg into a potential leg,

$$\Phi(\mathbf{k}) = \frac{A(a)}{k^2}\delta_m(\mathbf{k}), \qquad A(a) = -\frac{3}{2}\Omega_m H_0^2/a, \quad \text{(D4)}$$

so the potential bispectrum

$$B_\Phi(k_1,k_2,k_3) = \left[\prod_{i=1}^{3}\frac{A(a)}{k_i^2}\right] B_\delta(k_1,k_2,k_3) \quad \text{(D5)}$$

inherits the negative Poisson cube $A(a)^3 < 0$. In the transverse scalar approximation used for this equal-shell cumulant, the driving fields are obtained from the trace and trace-free parts of the screen-space Hessian of the lightcone potential [Eqs. (66) and (67)]. Transformed to a spherical-harmonic basis, that Hessian reduces per multipole to a simple angular multiplier on each leg,

$$\begin{aligned} &\Phi_{00}\colon\ +\frac{L^2}{\chi^2} \text{ on } Y_{\ell m},\\ &\Psi_0\colon\ -\frac{\sqrt{L^2(L^2-2)}}{\chi^2} \text{ on } {}_2Y_{\ell m},\\ &L^2 \equiv \ell(\ell+1). \end{aligned} \quad \text{(D6)}$$

the Laplacian L^2/χ^2 being the spin-0 trace [Ricci focusing, the transverse-Laplacian term of Eq. (66)] and $-\sqrt{L^2(L^2-2)}/\chi^2$ the spin-2 $\eth^2$ shear (the trace-free screen Hessian of Eq. (67)); the shared $1/\chi^2$ converts the two transverse screen derivatives at distance χ, and the two moduli agree at leading order for large ℓ. It is convenient to compute ζ_{abc} in the field basis: labeling the spin-0 leg $T = \Phi_{00}$ and each spin-2 leg $P = \Psi_0$, the four projections $XYZ \in \{TTT, TTP, TPP, PPP\}$ are the natural computational *channels*. Contracting the projected legs and summing over the parity-even multipole triangles ($\ell_1+\ell_2+\ell_3$ even, selected by the scalar Wigner-3j) gives each channel,

$$\zeta_{XYZ}(\gamma,\chi) = \sum_{\ell_1\ell_2\ell_3} \mathcal{G}_{\ell_1\ell_2\ell_3} P_{\ell_3}(\cos\gamma) B^{XYZ}_{\ell_1\ell_2\ell_3}(\chi), \quad \text{(D7)}$$

with $\mathcal{G}$ the parity-even Gaunt weight and each reduced bispectrum B^{XYZ} a single radial integral over P_δ [Eqs. (D2)–(D6)], evaluated in closed form by an FFTlog double-Bessel method, with an exact Wigner-3j sum at low multipoles ($\ell \le 60$) and the Limber approximation above.[7]

The high-ℓ branch makes B^{XYZ} explicit and is the most transparent. With $\chi = \chi(\lambda)$ the comoving distance to the shell and $k_i = \ell_i/\chi$ the comoving wave number that an angular scale ℓ_i probes at that distance,

$$B^{XYZ}_{\ell_1\ell_2\ell_3}(\chi) = \frac{D(z)^4}{\chi^4}\mathcal{K}^{XYZ}_{\ell_1\ell_2\ell_3}(z) B_\delta\left(\frac{\ell_1}{\chi},\frac{\ell_2}{\chi},\frac{\ell_3}{\chi}\right). \quad \text{(D8)}$$

Reading the factors one by one: B_δ is the matter bispectrum of Eq. (D2), the three-point clustering of the density field, evaluated at the three wave numbers that the multipole triangle (ℓ_1,ℓ_2,ℓ_3) picks out at distance χ. The growth factor $D(z)$ measures how much structure has assembled by redshift z; it enters at the fourth power because the tree-level bispectrum is built from two power spectra, each growing as D^2. The kernel $\mathcal{K}^{XYZ}_{\ell_1\ell_2\ell_3}(z)$ converts each matter leg into the field that does the lensing, combining the gravitational map from density to potential [$A(a) = -\frac{3}{2}\Omega_m H_0^2/a$] with the on-sky derivatives of Eq. (D6) that isolate the spin-0 focusing (Φ_{00}) or spin-2 shear (Ψ_0) part of each leg. The $1/\chi^4$ is purely geometric. Because the matter correlations fade over a radial distance much shorter than the depth of the source shell, the three points are effectively pinned to a common distance; collapsing the two now-redundant radial integrals leaves a single integral along the line of sight, and the surviving $1/\chi^4$ is the three-point analog of the $1/\chi^2$ in the familiar two-point Limber projection.

The spin-2 modulus channels. The path integral contracts the K-vertex as the local cumulant $\zeta_{abc} = \langle\varphi_a\varphi_b\varphi_c\rangle$ with the three real driving fields $\varphi \in \{\Phi_{00}, \Psi_+, \Psi_\times\}$, so the full cumulant needs every component of this triple. Parity removes most of them: reflection about the separation axis makes $\Psi_\times$ parity odd, so any component with an odd number of $\Psi_\times$ legs vanishes.

The field-basis channels of Eq. (D7), with $T = \Phi_{00}$ and $P = \Psi_0$, pair the two spin-2 legs *without conjugation* ($\Psi_0\Psi_0$), so they deliver ζ_{TTT}, ζ_{TTP}, and the parity-even *differences*,

$$\zeta_{TPP} = \langle\Phi_{00}(\Psi_+^2 - \Psi_\times^2)\rangle, \quad \zeta_{PPP} = \langle\Psi_+^3 - 3\Psi_+\Psi_\times^2\rangle. \quad \text{(D9)}$$

Resolving the individual $\langle\Phi_{00}\Psi_+^2\rangle$ and $\langle\Phi_{00}\Psi_\times^2\rangle$ (and likewise the Ψ_+^3, $\Psi_+\Psi_\times^2$ pieces) needs the complementary *sums*, supplied by two modulus channels in which one spin-2 leg is conjugated,

$$\zeta_B = \langle\Phi_{00}|\Psi_0|^2\rangle, \qquad \zeta_D = \langle\Psi_0|\Psi_0|^2\rangle, \quad \text{(D10)}$$

with $|\Psi_0|^2 = \Psi_+^2 + \Psi_\times^2$. Difference and sum reconstruct each component,

$$\langle\Phi_{00}\Psi_+^2\rangle = \frac{1}{2}(\zeta_{TPP} + \zeta_B), \quad \langle\Phi_{00}\Psi_\times^2\rangle = \frac{1}{2}(\zeta_B - \zeta_{TPP}), \quad \text{(D11)}$$

and likewise $\langle\Psi_+^3\rangle$, $\langle\Psi_+\Psi_\times^2\rangle$ from ζ_{PPP}, ζ_D.

[7]The collapsed FK geometry weights the near-degenerate triangles $\ell_1 \simeq \ell_2$ most heavily, but the sum runs over all parity-allowed triples. For the spin-2 channels the scalar $w_{3j}(\ell_1,\ell_2,\ell_3;0,0,0)$ enters as a triangle/parity selector; the spin weight is carried separately by the spin-weighted-harmonic contraction.

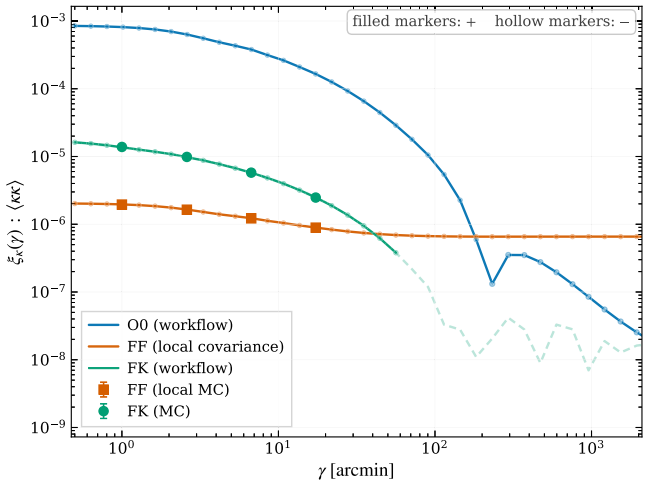


FIG. 17. *Workflow validation.* The convergence two-point function $\xi_\kappa(\gamma)$ split into its Order-0 (O0), nonlinear-propagation (FF) and three-point cumulant (FK) channels: the analytic workflow (lines) against a direct Monte Carlo of the stochastic Sachs equation with matched input statistics (markers). The FF line uses the simulation's local covariance, distinct from the nonlocal cosmological FF prediction. The Order-0 curve is shown for reference. FK markers are extrapolated to vanishing colored-noise correlation length, with error bars showing the standard error from independent realization blocks. Marker angles are $1', 2.6', 6.7', 17.3'$. Filled and hollow markers denote positive and negative values.

APPENDIX E: WORKFLOW VALIDATION

We validate that the workflow correctly propagates matched input driving-field statistics into lensing observables. Both Order-2 channels are checked against direct Monte-Carlo integration of the stochastic Sachs equation (Fig. 17); FK agrees within 1% and FF within 0.7% at $\gamma = 1', 2.6', 6.7', 17.3'$. FK uses the corrected tabulated cumulant at $\ell_{\max} = 15360$. The FF test uses the local covariance actually sampled by the simulation, including its positive-semidefinite adjustment, and a matched continuum Order-F^2 reference. It is separate from the nonlocal cosmological FF prediction of Sec. V B. Sampling standard errors are 0.36%–2.10% for FK and 0.27%–0.30% for FF, relative to their respective references. As a second, independent check of FK, we contract the tabulated cumulant ζ_{abc} with the FK kernel by direct quadrature, independently of the workflow code. This reproduces the workflow result within 1% at the same four separations. The tests use the fiducial flat-ΛCDM cosmology ($\Omega_m = 0.316$, $h = 0.671$, $n_s = 0.97$) with a single source plane at $z_s = 5$.

The two channels require different estimators because their stochastic noise enters differently. For FF, the nonlinear Sachs coupling can be switched on and off for the same random draw, and differencing the resulting matched runs cancels the common leading fluctuation. The FK signal, by contrast, is a property of the input statistics: it is the response of $\langle\kappa\kappa\rangle$ to the three-point cumulant of the driving field. A naive estimate of a small third moment is affected by Gaussian fluctuations, motivating a variance-controlled FK estimator. For the colored-noise regularization used here, the pointwise skewness scales as $\sigma_\lambda^{-1/2}$ when the integrated cumulants are held fixed. The FK estimate is therefore obtained at two finite correlation lengths along the ray and linearly extrapolated to zero. Error bars come from independent realization blocks. The FF estimator also retains higher propagation orders, which we check by varying the Sachs coupling. Integration-step, cumulant-injection, finite-correlation and source-grid checks assess numerical sensitivities separately from the sampling errors.

[1] M. Asgari, C.-A. Lin, B. Joachimi, B. Giblin, C. Heymans, H. Hildebrandt, A. Kannawadi, B. Stölzner, T. Tröster, J. L. van den Busch *et al.*, Kids-1000 cosmology: Cosmic shear constraints and comparison between two point statistics, Astron. Astrophys. **645**, A104 (2021).

[2] A. Amon, D. Gruen, M. A. Troxel, N. MacCrann, S. Dodelson, A. Choi, C. Doux, L. F. Secco, S. Samuroff, E. Krause *et al.*, Dark energy survey year 3 results: Cosmology from cosmic shear and robustness to data calibration, Phys. Rev. D **105**, 023514 (2022).

[3] Ž. Ivezić, S. M. Kahn, J. A. Tyson, B. Abel, E. Acosta, R. Allsman, D. Alonso, Y. AlSayyad, S. F. Anderson, J. Andrew *et al.*, Lsst: From science drivers to reference design and anticipated data products, Astrophys. J. **873**, 111 (2019).

[4] R. Laureijs, J. Amiaux, S. Arduini, J.-L. Augueres, J. Brinchmann, R. Cole, M. Cropper, C. Dabin, L. Duvet, A. Ealet *et al.*, Euclid definition study report, arXiv: 1110.3193.

[5] D. Spergel, N. Gehrels, C. Baltay, D. Bennett, J. Breckinridge, M. Donahue, A. Dressler, B. S. Gaudi, T. Greene, O. Guyon *et al.*, Wide-field InfrarRed survey telescope-astrophysics focused telescope assets WFIRST-AFTA 2015 report, arXiv:1503.03757.

[6] M. Bartelmann and P. Schneider, Weak gravitational lensing, Phys. Rep. **340**, 291 (2001).

[7] M. Kilbinger, Cosmology with cosmic shear observations: A review, Rep. Prog. Phys. **78**, 086901 (2015).

[8] R. Mandelbaum, Weak lensing for precision cosmology, Annu. Rev. Astron. Astrophys. **56**, 393 (2018).

[9] A. Cooray and W. Hu, Second-order corrections to weak lensing by large-scale structure, Astrophys. J. **574**, 19 (2002).
[10] S. Dodelson, C. Shapiro, and M. White, Reduced shear power spectrum, Phys. Rev. D **73**, 023009 (2006).
[11] F. Bernardeau, C. Bonvin, and F. Vernizzi, Full-sky lensing shear at second order, Phys. Rev. D **81**, 083002 (2010).
[12] E. Krause and C. M. Hirata, Weak lensing power spectra for precision cosmology-multiple-deflection, reduced shear, and lensing bias corrections, Astron. Astrophys. **523**, A28 (2010).
[13] G. Pratten and A. Lewis, Impact of post-born lensing on the CMB, J. Cosmol. Astropart. Phys. 08 (2016) 047.
[14] G. Fabbian, M. Calabrese, and C. Carbone, CMB weak-lensing beyond the born approximation: A numerical approach, J. Cosmol. Astropart. Phys. 02 (2018) 050.
[15] G. Marozzi, G. Fanizza, E. Di Dio, and R. Durrer, Cmb lensing beyond the leading order: Temperature and polarization anisotropies, Phys. Rev. D **98**, 023535 (2018).
[16] J. Fry, Gravity, bias, and the galaxy three-point correlation function, Phys. Rev. Lett. **73**, 215 (1994).
[17] R. Scoccimarro, H. A. Feldman, J. N. Fry, and J. A. Frieman, The bispectrum of IRAS redshift catalogs, Astrophys. J. **546**, 652 (2001).
[18] D. Jeong and E. Komatsu, Primordial non-gaussianity, scale-dependent bias, and the bispectrum of galaxies, Astrophys. J. **703**, 1230 (2009).
[19] A. Halder, O. Friedrich, S. Seitz, and T. N. Varga, The integrated three-point correlation function of cosmic shear, Mon. Not. R. Astron. Soc. **506**, 2780 (2021).
[20] P. Lemos, A. Challinor, and G. Efstathiou, The effect of Limber and flat-sky approximations on galaxy weak lensing, J. Cosmol. Astropart. Phys. 05 (2017) 014.
[21] X. Fang, E. Krause, T. Eifler, and N. MacCrann, Beyond Limber: Efficient computation of angular power spectra for galaxy clustering and weak lensing, J. Cosmol. Astropart. Phys. 05 (2020) 010.
[22] T. D. Kitching, J. Alsing, A. F. Heavens, R. Jimenez, J. D. McEwen, and L. Verde, The limits of cosmic shear, Mon. Not. R. Astron. Soc. **469**, 2737 (2017).
[23] C. Cuesta-Lazaro, A. Quera-Bofarull, R. Reischke, and B. M. Schäfer, Gravitational corrections to light propagation in a perturbed FLRW universe and corresponding weak-lensing spectra, Mon. Not. R. Astron. Soc. **477**, 741 (2018).
[24] S. Hilbert, J. Hartlap, S. D. M. White, and P. Schneider, Ray-tracing through the millennium simulation: Born corrections and lens-lens coupling in cosmic shear and galaxy-galaxy lensing, Astron. Astophys. **499**, 31 (2009).
[25] A. Petri, Z. Haiman, L. Hui, M. May, and J. M. Kratochvil, Cosmology with Minkowski functionals and moments of the weak lensing convergence field, Phys. Rev. D **88**, 123002 (2013).
[26] R. Takahashi, T. Hamana, M. Shirasaki, T. Namikawa, T. Nishimichi, K. Osato, and K. Shiroyama, Full-sky gravitational lensing simulation for large-area galaxy surveys and cosmic microwave background experiments, Astrophys. J. **850**, 24 (2017).
[27] R. Sachs, Gravitational waves in general relativity. VI. The outgoing radiation condition, Proc. R. Soc. A **264**, 309 (1961).
[28] P. C. Martin, E. D. Siggia, and H. A. Rose, Statistical dynamics of classical systems, Phys. Rev. A **8**, 423 (1973).
[29] H.-K. Janssen, On a Lagrangean for classical field dynamics and renormalization group calculations of dynamical critical properties, Z. Phys. B Condens. Matter **23**, 377 (1976).
[30] J. A. Krommes, Fundamental statistical descriptions of plasma turbulence in magnetic fields, Phys. Rep. **360**, 1 (2002).
[31] F. Bernardeau, S. Colombi, E. Gaztañaga, and R. Scoccimarro, Large-scale structure of the universe and cosmological perturbation theory, Phys. Rep. **367**, 1 (2002).
[32] D. Baumann, A. Nicolis, L. Senatore, and M. Zaldarriaga, Cosmological non-linearities as an effective fluid, J. Cosmol. Astropart. Phys. 07 (2012) 051.
[33] M. Garny, D. Laxhuber, and R. Scoccimarro, Perturbation theory with dispersion and higher cumulants: Framework and linear theory, Phys. Rev. D **107**, 063539 (2023).
[34] M. Crocce and R. Scoccimarro, Renormalized cosmological perturbation theory, Phys. Rev. D **73**, 063519 (2006).
[35] P. Valageas, Expansion schemes for gravitational clustering: Computing two-point and three-point functions, Aston. Astrophys. **484**, 79 (2008).
[36] K. Izumi and J. Soda, Renormalized Newtonian cosmic evolution with primordial non-Gaussianity, Phys. Rev. D **76**, 083517 (2007).
[37] Z. Zhang, SFT-WICK: A formalism and package for Feynman-diagram expansion and evaluation in stochastic field theories, arXiv:2606.19480.
[38] M. P. van Daalen, J. Schaye, C. M. Booth, and C. Dalla Vecchia, The effects of galaxy formation on the matter power spectrum: A challenge for precision cosmology, Mon. Not. R. Astron. Soc. **415**, 3649 (2011).
[39] C. Wagner, F. Schmidt, C.-T. Chiang, and E. Komatsu, The angle-averaged squeezed limit of nonlinear matter N-point functions, J. Cosmol. Astropart. Phys. 08 (2015) 042.
[40] Y. Li, W. Hu, and M. Takada, Super-sample covariance in simulations, Phys. Rev. D **89**, 083519 (2014).
[41] A. Barreira and F. Schmidt, Responses in large-scale structure, J. Cosmol. Astropart. Phys. 06 (2017) 053.
[42] J. M. Diego, T. Broadhurst, and G. F. Smoot, Evidence for lensing of gravitational waves from LIGO-Virgo data, Phys. Rev. D **104**, 103529 (2021).
[43] J. E. Gunn, On the Propagation of Light in Inhomogeneous Cosmologies. I. Mean Effects, Astrophys. J. **150**, 737 (1967).
[44] G. F. Ellis, R. Maartens, and M. A. MacCallum, *Relativistic Cosmology* (Cambridge University Press, Cambridge, England, 2012).
[45] E. Newman and R. Penrose, An approach to gravitational radiation by a method of spin coefficients, J. Math. Phys. (N. Y.) **3**, 566 (1962).
[46] R. M. Wald, *General Relativity* (University of Chicago Press, Chicago, 1984).
[47] G. Fanizza, E. Di Dio, R. Durrer, and G. Marozzi, The gauge invariant cosmological Jacobi map from weak lensing at leading order, J. Cosmol. Astropart. Phys. 08 (2022) 052.

[48] S. Pueblas and R. Scoccimarro, Generation of vorticity and velocity dispersion by orbit crossing, Phys. Rev. D **80**, 043504 (2009).
[49] E. Bertschinger, Cosmological dynamics, in *Cosmology and Large Scale Structure, Proceedings of the Les Houches Summer School, Session LX*, edited by R. Schaeffer, J. Silk, M. Spiro, and J. Zinn-Justin (Elsevier, Amsterdam, 1996), pp. 273–348.
[50] P. A. R. Ade *et al.* (Planck Collaboration), Planck 2015 results. XIII. Cosmological parameters, Astron. Astrophys. **594**, A13 (2016).
[51] Z. Zhang, P. Bull, and A. Nicola, CANOES: Cosmological Angular Observables and Estimator Suite (to be published).
[52] C. Clarkson, G. F. Ellis, A. Faltenbacher, R. Maartens, O. Umeh, and J.-P. Uzan, (Mis) interpreting supernovae observations in a lumpy universe, Mon. Not. R. Astron. Soc. **426**, 1121 (2012).
[53] J. Adamek, D. Daverio, R. Durrer, and M. Kunz, General relativity and cosmic structure formation, Nat. Phys. **12**, 346 (2016).
[54] P. Fleury, J. Larena, and J.-P. Uzan, The theory of stochastic cosmological lensing, J. Cosmol. Astropart. Phys. 11 (2015) 022.
[55] N. E. Chisari, D. Alonso, E. Krause, C. D. Leonard, P. Bull, J. Neveu, A. Villarreal, S. Singh, T. McClintock, J. Ellison *et al.*, Core cosmology library: Precision cosmological predictions for lsst, Astrophys. J. Suppl. Ser. **242**, 2 (2019).
[56] R. Takahashi, T. Nishimichi, T. Namikawa, A. Taruya, I. Kayo, K. Osato, Y. Kobayashi, and M. Shirasaki, Fitting the nonlinear matter bispectrum by the Halofit approach, Astrophys. J. **895**, 113 (2020).
[57] P. Fleury, J. Larena, and J.-P. Uzan, Weak gravitational lensing of finite beams, Phys. Rev. Lett. **119**, 191101 (2017).
[58] https://github.com/StatFieldTheory.
[59] J. M. Martin-Garcia, A. García-Parrado, A. Stecchina, B. Wardell, C. Pitrou, D. Brizuela *et al.*, xAct: Efficient tensor computer algebra for the wolfram language (2020), http://www.xact.es.
[60] J. M. Martín-García, xPerm: Fast index canonicalization for tensor computer algebra, Comput. Phys. Commun. **179**, 597 (2008).
[61] D. W. Hogg, Distance measures in cosmology, arXiv:astro-ph/9905116.