Title

The impact of anisotropic redshift distributions on angular clustering

Journal

Journal of Cosmology and Astroparticle Physics, 2023(07)

ISSN

1475-7516

Authors

Lizancos, Antón Baleato

White, Martin

Publication Date

2023-07-01

DOI

10.1088/1475-7516/2023/07/044

Prepared for submission to JCAP

The Impact of Anisotropic Redshift Distributions on Angular Clustering

Antón Baleato Lizancos and Martin White

Berkeley Center for Cosmological Physics, UC Berkeley, CA 94720, USA
Department of Physics, University of California, Berkeley, CA 94720, USA
Lawrence Berkeley National Laboratory, One Cyclotron Road, Berkeley, CA 94720, USA

E-mail: a.baleatolizancos@berkeley.edu, mwhite@berkeley.edu

Abstract. A leading way to constrain physical theories from cosmological observations is to test their predictions for the angular clustering statistics of matter tracers, a technique that is set to become ever more central with the next generation of large imaging surveys. Interpretation of this clustering requires knowledge of the projection kernel, or the redshift distribution of the sources, and the typical assumption is an isotropic redshift distribution for the objects. However, variations in the kernel are expected across the survey footprint due to photometric variations and residual observational systematic effects. We develop the formalism for anisotropic projection and present several limiting cases that elucidate the key aspects. We quantify the impact of anisotropies in the redshift distribution on a general class of angular two-point statistics. In particular, we identify a mode-coupling effect that can add power to auto-correlations, including galaxy clustering and cosmic shear, and remove it from certain cross-correlations. If the projection anisotropy is primarily at large scales, the mode-coupling depends upon its variance as a function of redshift; furthermore, it is often of similar shape to the signal. In contrast, the cross-correlation of a field whose selection function is anisotropic with another one featuring no such variations – such as CMB lensing – is immune to these effects. We discuss explicitly several special cases of the general formalism including galaxy clustering, galaxy-galaxy lensing, cosmic shear and cross-correlations with CMB lensing, and publicly release a code to compute the biases.

Contents

1 Introduction

The measurement of angular clustering of projected fields holds a central place in cosmology, and such measurements frequently allow us to push the cosmological frontier in sky coverage, source density and redshift [1–6]. The next generation of large imaging surveys, such as Vera Rubin LSST [7–9] and Euclid [10, 11] will ensure such analyses remain central to cosmological inference for the next decade and beyond. With some assumptions about evolution and statistical isotropy, observations along the past light cone can be used to infer the underlying 3D clustering if the radial distribution of the signal is known. This is typically related to the redshift distribution of a set of sources, and the general assumption is that this distribution is independent of location on the sky.

However, observational systematic effects are expected to introduce anisotropy in the selection of sources. For instance, it is well known that the observed number density of galaxies is a modulation of the true population by variations in detector sensitivity across time and location on the focal plane, changes in the observing conditions, completeness near bright stars, extinction due to Galactic dust, deblending, or the separation of galaxies from stars; this modulation can in turn bias cosmological constraints [12–14]. The same systematics will also compromise characterizations of the redshift distribution of the objects, presenting

a challenge to the interpretation of the observed angular pattern in terms of an underlying 3D distribution. In this paper, we investigate how the analysis needs to be changed if the assumption of isotropy in the redshift distribution is relaxed, illustrating the general formalism with several examples.

The outline of the paper is as follows. In §2 we introduce the general formalism for treating anisotropic projection kernels for 2D fields. We demonstrate the impact of such anisotropic projection in a simple "flat sky" calculation (§2.1) and then show that the main structure of the calculation – extra additive and mode-coupling contributions to the clustering – carries across to the full-sky calculation (§2.2). We give several examples and special cases in §3, showing what the general formalism implies for galaxy clustering, cross-correlations, cosmic shear and galaxy-galaxy lensing. Then, in §4, we provide a quantitative exploration of the additional contributions for fluctuations characteristic of present-day surveys. Our conclusions are presented in §5 while some technical details are relegated to appendices.

2 Formalism

We are interested in considering a field, such as the galaxy density, that has been projected along the line-of-sight direction, such that e.g.

$$1 + \delta^{(2D)}(\hat{\boldsymbol{n}}) = \int d\chi\ \phi(\chi, \hat{\boldsymbol{n}})\ \left[1 + \delta^{(3D)}(\chi\,\hat{\boldsymbol{n}})\right] \tag{2.1}$$

where χ is the comoving radial coordinate and $\hat{\boldsymbol{n}}$ is a unit vector on the sphere. The projection kernel, ϕ, is inferred from an assumed or observed redshift distribution, e.g. for the projected galaxy density example above $\phi \propto H\,(dN/dz)$, where H is the Hubble parameter and dN/dz is the redshift distribution of the sources; we shall discuss other kernels later. The standard assumption is that ϕ is a function only of χ, with no angular dependence. In this paper we are interested in the case where dN/dz varies with position on the sky, and hence $\phi = \phi(\chi, \hat{\boldsymbol{n}})$.

Throughout we shall assume that $\phi(\chi, \hat{\boldsymbol{n}})$ is well characterized so that we can ignore the problem of redshift-distribution errors and focus instead on the impact of a spatially-varying redshift distribution. The issue of properly inferring the redshift distribution of a population of objects, and its uncertainties, is a complex one with a large literature. Recent reviews of the current state-of-the-art can be found in refs. [15, 16]. We assume that this process has resulted in an estimate of the number of galaxies in redshift bins, including its variation across the survey, and hence of $\phi(\chi, \hat{\boldsymbol{n}})$. We denote the average of ϕ across the sky (i.e. survey) as $\bar{\phi}(\chi)$, which defines the residual through

$$\phi(\chi, \hat{\boldsymbol{n}}) = \bar{\phi}(\chi) + \Delta\phi(\chi, \hat{\boldsymbol{n}}) \tag{2.2}$$

Note that by definition $\Delta\phi$ averages to zero and $\bar{\phi}$ integrates to unity[1]. We shall also expand the angular dependence of our fields in spherical harmonics, e.g.

$$\delta^{(2D)}(\hat{\boldsymbol{n}}) = \sum_{\ell m} \delta^{(2D)}_{\ell m} Y_{\ell m}(\hat{\boldsymbol{n}})\,, \quad \Delta\phi(\chi, \hat{\boldsymbol{n}}) = \sum_{\ell m} \Delta\phi_{\ell m}(\chi) Y_{\ell m}(\hat{\boldsymbol{n}})\,. \tag{2.3}$$

We are concerned with the implications of non-zero $\Delta\phi$ on the observed auto- and cross-clustering of the projected field, $\delta^{(2D)}$.

[1]The latter condition, which holds automatically for densities, need not hold for more general fields, e.g. cosmic shear. We shall relax this condition when appropriate.

2.1 Flat sky

Before presenting our full calculation (§2.2), let us consider an algebraically simpler model that nonetheless illustrates many of the features of the full model. Specifically imagine a flat sky, with a redshift depth[2] small enough that we may directly convert angles and redshifts into Cartesian coordinates: $(\mathbf{x}_\perp, \chi)$. We split the selection function, $\phi(\mathbf{x})$, into an average piece, $\bar{\phi}(\chi)$, plus a fluctuation $\Delta\phi(\mathbf{x}_\perp, \chi)$. Then

$$\delta_{2D}(\mathbf{x}_\perp) = \int d\chi\ \Delta\phi(\mathbf{x}_\perp, \chi) + \int d\chi\ \left[\bar{\phi}(\chi) + \Delta\phi(\mathbf{x}_\perp, \chi)\right] \delta_{3D}(\mathbf{x}_\perp, \chi)\,. \tag{2.4}$$

Conservation of the mean density requires $\int \bar{\phi}\, d\chi = 1$ and $\int \Delta\phi\, d^2x_\perp = 0$ or equivalently $\Delta\phi(\mathbf{k}_\perp = 0, k_\parallel) = 0$. In moving to Fourier space two of the terms are straightforward Fourier transforms while the $\Delta\phi\,\delta_{3D}$ term is a convolution. Let us denote $X(\mathbf{x}) = \Delta\phi\,\delta_{3D}$. Then

$$X(\mathbf{k}) = \left[\Delta\phi\,\delta_{3D}\right](\mathbf{k}) = \int \frac{d^3k_1\, d^3k_2}{(2\pi)^6} (2\pi)^3 \delta^{(D)}(\mathbf{k} - \mathbf{k}_1 - \mathbf{k}_2) \Delta\phi(\mathbf{k}_1) \delta_{3D}(\mathbf{k}_2) \tag{2.5}$$

and if the 3D power spectrum of δ_{3D} is $P(\mathbf{k})$ then

$$\langle X(\mathbf{k}_1) X^\star(\mathbf{k}_2)\rangle = \int \frac{d^3k}{(2\pi)^3}\ \Delta\phi(\mathbf{k}_1 - \mathbf{k}) \Delta\phi^\star(\mathbf{k}_2 - \mathbf{k}) P(\mathbf{k})\,. \tag{2.6}$$

Note that the product in configuration space has led to a convolution in Fourier space that samples a range of $\mathbf{k}$-modes around $\mathbf{k}_1$ or $\mathbf{k}_2$.

The projection over χ in equation 2.4 implies that $\delta_{2D}(\mathbf{k}_\perp)$, is simply the Fourier transform of each contribution evaluated at $k_\parallel = 0$. The angular power spectrum thus becomes

$$C_K = \int \frac{dk_\parallel}{2\pi} P(\mathbf{k}_\perp, k_\parallel) \left|\bar{\phi}(k_\parallel)\right|^2 + C_K^{\Delta\phi} + \int \frac{d^3k}{(2\pi)^3} P(\mathbf{k}) \left|\Delta\phi(\mathbf{K} - \mathbf{k})\right|^2\,, \tag{2.7}$$

where $K = |\mathbf{k}_\perp|$ and in the last integral the $\mathbf{K} = (\mathbf{k}_\perp, 0)$ is interpreted as a 3-vector with zero line-of-sight component. The first term is the usual expression resulting from a fixed projection kernel, $\bar{\phi}$. It has the form of a power spectrum multiplied by a line-of-sight window function. The 'additive' term,

$$C_K^{\Delta\phi} = \int d\chi_1\, d\chi_2\ C_K^{\Delta\phi}(\chi_1, \chi_2) = \int d\chi_1\, d\chi_2 \int d^2x_\perp e^{-i\mathbf{K}\cdot\mathbf{x}_\perp} \Delta\phi(\mathbf{x}_\perp, \chi_1) \Delta\phi(0, \chi_2) \quad , \tag{2.8}$$

is just the angular power spectra of the projected fluctuation, $\Delta\phi$. It comes from modulation of the mean density by the varying projection kernel. The final term is a 'mode-coupling' term that samples from $P(\mathbf{k}_\perp, k_\parallel)$ over a range of $\mathbf{k}_\perp$ around $\mathbf{K}$. The three contributions to C come from each of the contributions to δ_{2D} squared. The cross terms vanish because of the constraint from the mean density.

This simple, flat-sky model thus leads us to expect that we will see three contributions to the measured clustering: the signal as for a uniform dN/dz, an additive contribution equal to the auto-correlation of the dN/dz fluctuations and a mode-coupling term that couples power at an observed scale, K, to that of nearby scales over a range defined by the Fourier transform of $\Delta\phi$. All three of these contributions and their behaviors carry across to the

[2]This analysis might be appropriate for medium- or narrow-band selected samples of emission line galaxies over small fields, for example Lyα emitter surveys (LAEs; see [17] for a review).

full calculation. This simpler model also suggests that the three contributions come from the auto-spectra of the three terms in equation (2.4). This will also have an analog in the full calculation. Finally, these results imply that if a projected field is cross-correlated with one that does not have any uncertainty in dN/dz (e.g. CMB lensing) the corrections vanish, or if such a field is correlated with one having mean zero (e.g. the shear field) the additive correction vanishes. These implications also hold in the full case, and are discussed further in what follows.

2.2 Beyond flat sky

Having built intuition for the physics underpinning the calculation, let us now repeat it in the more general spherical-sky formalism, using χ to parametrize comoving distances along a line-of-sight direction specified by the unit vector $\hat{\boldsymbol{n}}$. The general, position-dependent selection function is still given by Eq. (2.2). Projecting the 3D density field with this selection function, as in equation (2.1), and imposing the integral constraints $\int \bar{\phi}\, d\chi = 1$ and $\int \Delta\phi\, d^2\hat{\boldsymbol{n}} = 0$, we get

$$\delta^{(2D)}(\hat{\boldsymbol{n}}) = \int d\chi\, \left[\bar{\phi}(\chi) + \Delta\phi(\chi, \hat{\boldsymbol{n}})\right] \delta(\chi, \hat{\boldsymbol{n}}) + \int d\chi\, \Delta\phi(\chi, \hat{\boldsymbol{n}})\,, \tag{2.9}$$

with spherical harmonic coefficients

$$\delta^{(2D)}_{\ell m} = \int d\hat{\boldsymbol{n}}\, Y^*_{\ell m}(\hat{\boldsymbol{n}}) \int d\chi\, \left[\bar{\phi}(\chi)\delta(\chi, \hat{\boldsymbol{n}}) + \Delta\phi(\chi, \hat{\boldsymbol{n}})\delta(\chi, \hat{\boldsymbol{n}}) + \Delta\phi(\chi, \hat{\boldsymbol{n}})\right]\,. \tag{2.10}$$

The first term would be the only contribution if the redshift distribution were perfectly isotropic. When it is not, the other two terms give rise to a multiplicative and an additive contribution, respectively. Let us now carefully unpack the second term, leaving the other ones to follow by analogy. It gives[3]:

$$\begin{aligned}
\delta^{(2D)}_{\ell m} &\supset \int d\hat{\boldsymbol{n}}\, Y^*_{\ell m}(\hat{\boldsymbol{n}}) \int d\chi\, \Delta\phi(\chi, \hat{\boldsymbol{n}})\delta(\chi, \hat{\boldsymbol{n}}) \\
&= (4\pi)^2 \int d\chi \int \frac{\mathrm{d}^3\boldsymbol{k}_1}{(2\pi)^3} \frac{\mathrm{d}^3\boldsymbol{k}_2}{(2\pi)^3}\, \Delta\phi(\boldsymbol{k}_1)\delta\left(\boldsymbol{k}_2, z\right) \\
&\qquad \times (-1)^m \sum_{\ell_1 m_1} \sum_{\ell_2 m_2} i^{\ell_1+\ell_2} G^{\ell\ell_1\ell_2}_{-mm_1m_2} j_{\ell_1}(k_1\chi) j_{\ell_2}(k_2\chi) Y^*_{\ell_1 m_1}(\hat{\boldsymbol{k}}_1) Y^*_{\ell_2 m_2}(\hat{\boldsymbol{k}}_2)\,,
\end{aligned} \tag{2.12}$$

where, in going to the last line, we have used Rayleigh's plane wave expansion

$$e^{i\boldsymbol{k}\cdot\boldsymbol{x}} = 4\pi \sum_{\ell m} i^\ell j_\ell(k\chi) Y^*_{\ell m}(\hat{\boldsymbol{k}}) Y_{\ell m}(\hat{\boldsymbol{n}})\,, \tag{2.13}$$

[3]We work in the asymmetric Fourier transform convention where

$$f(\boldsymbol{k}) = \int d^3\boldsymbol{x}\, f(\boldsymbol{x}) e^{-i\boldsymbol{k}\cdot\boldsymbol{x}} \quad \text{and} \quad f(\boldsymbol{x}) = \int \frac{\mathrm{d}^3\boldsymbol{k}}{(2\pi)^3}\, f(\boldsymbol{k}) e^{i\boldsymbol{k}\cdot\boldsymbol{x}}\,. \tag{2.11}$$

Further, here and in subsequent integrals, we simplify notation by letting the dependence on comoving distance appear implicitly via the redshift, $\delta(\boldsymbol{k}, z) \equiv \delta(\boldsymbol{k}, z(\chi))$. On the other hand, $\Delta\phi$ is only defined on the past lightcone, so its Fourier transform is fully specified by a wavevector.

as well as the definition of the Gaunt integral

$$G^{\ell_1\ell_2\ell_3}_{m_1m_2m_3} \equiv \int d\hat{\boldsymbol{n}}\, Y_{\ell_1 m_1}(\hat{\boldsymbol{n}}) Y_{\ell_2 m_2}(\hat{\boldsymbol{n}}) Y_{\ell_3 m_3}(\hat{\boldsymbol{n}}) \\ = \sqrt{\frac{(2\ell_1+1)(2\ell_2+1)(2\ell_3+1)}{4\pi}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix} . \tag{2.14}$$

At this point, it will be useful to extract the spherical harmonics of each radial slice. To do this, note that an arbitrary 3D field g can be expressed as

$$g_{\ell m}(\chi) \equiv \int d\hat{\boldsymbol{n}}\, Y^*_{\ell m}(\hat{\boldsymbol{n}}) g(\chi, \hat{\boldsymbol{n}}) . \tag{2.15}$$

In appendix B, we link this to the spherical Fourier-Bessel basis and use that to glean insights into the structure of the perturbations. If we have two statistically-isotropic random fields, δ^a and δ^b, with 3D cross-spectrum

$$\langle \delta^a(\boldsymbol{k}_1, z_1)\, \delta^{b,*}(\boldsymbol{k}_2, z_2)\rangle = (2\pi)^3\, \delta^{(3)}_{\rm D}(\boldsymbol{k}_1 - \boldsymbol{k}_2)\, P_{ab}(k; z_1, z_2) , \tag{2.16}$$

the angular cross-spectrum of two radial slices is given by

$$\langle \delta^a_{\ell_1 m_1}(\chi_1) \delta^{b,*}_{\ell_2 m_2}(\chi_2)\rangle = \delta_{\ell_1 l_2}\delta_{m_1 m_2} \frac{2}{\pi} \int dk\, k^2 j_{\ell_1}(k\chi_1) j_{\ell_1}(k\chi_2) P_{ab}(k; z(\chi_1), z(\chi_2)) \\ \equiv \delta_{\ell_1 l_2}\delta_{m_1 m_2} C^{ab}_{\ell_1}(\chi_1, \chi_2) . \tag{2.17}$$

In the literature, this sometimes goes by the name of multi-frequency angular power spectrum (MAPS [18]; see also refs. [19–21]).

The formalism is slightly different for the $\Delta\phi$'s, since these are fixed by whatever systematic effects are driving the variations in the dN/dz's and are therefore deterministic rather than stochastic. Despite commuting with ensemble averaging, we can still define a notion of their C_ℓ's as

$$C^{\Delta\phi^a \Delta\phi^b}_\ell(\chi_1, \chi_2) \equiv \frac{1}{2\ell+1} \sum_m \Delta\phi^a_{\ell m}(\chi_1) \Delta\phi^{b,*}_{\ell m}(\chi_2) . \tag{2.18}$$

We can now continue to simplify equation (2.12). With the toolkit we have developed, we can write

$$\delta^{(2D)}_{\ell m} \supset \int d\chi\, (-1)^m \sum_{\ell_1 m_1} \sum_{\ell_2 m_2} G^{\ell\ell_1\ell_2}_{-m m_1 m_2} \Delta\phi_{\ell_1 m_1}(\chi) \delta_{\ell_2 m_2}(\chi) . \tag{2.19}$$

Then, defining

$$\{\Delta\phi\, \delta\}_{\ell m}(\chi) \equiv (-1)^m \sum_{\ell_1 m_1} \sum_{\ell_2 m_2} G^{\ell\ell_1\ell_2}_{-m m_1 m_2} \Delta\phi_{\ell_1 m_1}(\chi) \delta_{\ell_2 m_2}(\chi) , \tag{2.20}$$

and proceeding similarly for the other terms in equation 2.10, we find

$$\delta^{(2D)}_{\ell m} = \int d\chi\, \left[\bar{\phi}(\chi)\delta_{\ell m}(\chi) + \Delta\phi_{\ell m}(\chi) + \{\Delta\phi\, \delta\}_{\ell m}(\chi)\right] . \tag{2.21}$$

This integral can be regarded as a sum of contributions from spherical shells positioned at increasing distance from the observer. The first two terms are responsible for projecting the 3D anisotropy in δ and $\Delta\phi$ onto the shells, while the last term is associated with the coupling of the angular momenta of the two fields.

Consider, in turn, the angular cross-correlation of two projected overdensity fields, $\delta_a^{(2D)}$ and $\delta_b^{(2D)}$, each with its own radial selection function and associated perturbation. On the full sky, the total angular power spectrum is

$$T_\ell^{ab} = \frac{1}{2\ell+1}\sum_m \langle \delta^{(2D)}_{a,\ell m}\delta^{(2D),*}_{b,\ell m}\rangle\,. \tag{2.22}$$

All the cross-terms in this expression vanish, some because the δ's have mean zero by definition, others because they entail a coupling of $\bar{\phi}$ with $\Delta\phi$ – since the former is isotropic, there can only be a contribution from $\Delta\phi_{00}$, and this is zero by construction (though see appendix A for a generalization to the case where it is not) – and we are left with the sum of three auto-spectra:

$$T_\ell^{ab} = \int d\chi_1\, d\chi_2 \;\left[U_\ell^{ab}(\chi_1,\chi_2) + A_\ell^{ab}(\chi_1,\chi_2) + R_\ell^{ab}(\chi_1,\chi_2)\right] \tag{2.23}$$

where the 'uniform' and 'additive' contributions are

$$U_\ell^{ab} = \bar{\phi}^a(\chi_1)\bar{\phi}^b(\chi_2)C_\ell^{ab}(\chi_1,\chi_2) \quad \text{and} \quad A_\ell^{ab} = C_\ell^{\Delta\phi^a\Delta\phi^b}(\chi_1,\chi_2)\,, \tag{2.24}$$

while the mode-coupling contribution is

$$\begin{aligned} R_\ell^{ab} &= (2\ell+1)^{-1}\sum_m \left\langle \{\Delta\phi^a\,\delta^a\}_{\ell m}(\chi_1)\{\Delta\phi^b\,\delta^b\}^\star_{\ell m}(\chi_2)\right\rangle && (2.25)\\ &= (2\ell+1)^{-1}\sum_m\sum_{123} G^{\ell\ell_1\ell_2}_{-mm_1m_2}G^{\ell\ell_1\ell_3}_{-mm_1m_3}\Delta\phi^a_{\ell_2m_2}(\chi_1)\Delta\phi^{b,\star}_{\ell_3m_3}(\chi_2)C^{ab}_{\ell_1}(\chi_1,\chi_2) && (2.26)\\ &= \sum_L M_{\ell L}^{\Delta\phi^a\Delta\phi^b}(\chi_1,\chi_2)C_L^{ab}(\chi_1,\chi_2)\,. && (2.27)\end{aligned}$$

In going to the last line, we used the definitions of the C_ℓ's in (2.17) and (2.18), and harnessed the analogy with the mode-coupling induced by a mask (see, e.g. ref. [22]) by defining

$$M_{\ell L}^{\Delta\phi^a\Delta\phi^b}(\chi_1,\chi_2) \equiv \frac{2L+1}{4\pi}\sum_{\ell'}(2\ell'+1)\begin{pmatrix}\ell & L & \ell'\\ 0&0&0\end{pmatrix}^2 C_{\ell'}^{\Delta\phi^a\Delta\phi^b}(\chi_1,\chi_2)\,, \tag{2.28}$$

using the identity

$$\sum_{mm_1} G^{\ell\ell_1\ell_2}_{-mm_1m_2}G^{\ell\ell_1\ell_3}_{-mm_2m_3} = \delta_{\ell_2\ell_3}\delta_{m_2m_3}\frac{(2\ell+1)(2\ell_1+1)}{4\pi}\begin{pmatrix}\ell & \ell_1 & \ell_2\\ 0&0&0\end{pmatrix}^2\,. \tag{2.29}$$

In the next section, we will use an analytically-tractable toy model to understand the structure of this mode coupling and derive a simple and accurate approximation valid in most cases of interest. In addition, in appendix C, we explain how to evaluate the integrals in equation (2.23) efficiently.

Our formalism reveals two important insights. First, we learn that, as long as the mean is correctly characterized, the cross-correlation of a field with uncertain dN/dz with another

one with no such uncertainty is completely unbiased[4]. Second, despite $\Delta\phi$ being in general anisotropic, the biases depend only on the diagonal component of the perturbation's angular spectrum, namely $C_\ell^{\Delta\phi^a\Delta\phi^b}(\chi_1,\chi_2)$. This suggests that an estimator for the diagonal elements of $C_\ell^{\Delta\phi^a\Delta\phi^b}$ could potentially be built from measurements of the off-diagonal elements of the angular power spectrum, $T_{\ell\ell'}^{ab}$, and used to mitigate the effects we have described. We defer a more detailed exploration of this approach to future work.

2.3 The shape of the mode-coupling integral

In order to better understand the implications of the formalism above, let us consider a very simplified example where $\Delta\phi$ contains only a single, very large angular scale mode. By working through this case we will be able to see how the shape of the mode-coupling contribution arises and under what conditions it mirrors that of the cosmological signal. Moreover, this simplified calculation will pave the way to a simple yet accurate analytic approximation that will be valid much more generally, in any situation where clustering is measured on smaller (angular) scales than the anisotropy of ϕ.

Assume the shift in the mean redshift of the distribution is small compared to the width, so we can approximate

$$\phi(\chi,\hat{\boldsymbol{n}}) = \frac{H(\chi)}{c}\frac{dn_g}{dz}(\chi,\hat{\boldsymbol{n}}) = \frac{H(\chi)}{c}\frac{d\bar{n}_g}{dz}\left(\chi-\chi_{\rm shift}(\hat{\boldsymbol{n}})\right) \tag{2.30}$$

$$\approx \frac{H(\chi)}{c}\left[\frac{d\bar{n}_g}{dz}(\chi)-\chi_{\rm shift}(\hat{\boldsymbol{n}})\frac{d^2\bar{n}}{d\chi dz}(\chi)\right]. \tag{2.31}$$

Identifying the first term with the fiducial selection function, we can isolate the perturbation

$$\Delta\phi(\chi,\hat{\boldsymbol{n}}) = -\chi_{\rm shift}(\hat{\boldsymbol{n}})\Psi(\chi) \tag{2.32}$$

where we have defined $\Psi = (H/c)\,d^2\bar{n}/(d\chi\,dz)$ having dimensions of inverse length squared. For pedagogical purposes, consider

$$\chi_{\rm shift}(\hat{\boldsymbol{n}}) = \epsilon\,\Re\{Y_{11}(\hat{\boldsymbol{n}})\}\,, \tag{2.33}$$

where ϵ is a small constant with units of distance. It follows that

$$\Delta\phi_{\ell m}(\chi) = -\frac{1}{2}\delta_{\ell 1}\left(\delta_{m1}-\delta_{m-1}\right)\epsilon\Psi(\chi)\,, \tag{2.34}$$

so our approximation will be valid whenever $\epsilon H(\chi)/c \ll 1$. We can use this to calculate

$$C_\ell^{\Delta\phi^a\Delta\phi^b}(\chi_1,\chi_2) = \frac{1}{6}\delta_{\ell 1}\frac{\epsilon_a\epsilon_b}{c^2}\Psi^a(\chi_1)\Psi^b(\chi_2)\,. \tag{2.35}$$

The additive bias is therefore only present at $\ell = 1$, the mode where we seeded the perturbation:

$$T_\ell^{ab} \supset \int d\chi_1\,d\chi_2 C_\ell^{\Delta\phi^a\Delta\phi^b}(\chi_1,\chi_2) = \frac{1}{6}\delta_{\ell 1}\frac{\epsilon_a\epsilon_b}{c^2}\int d\chi_1\,\Psi^a(\chi_1)\int d\chi_2\,\Psi^b(\chi_2)\,. \tag{2.36}$$

[4]Contrast this with the cross-correlation of a masked field on the sphere with another one covering the full sky. In that case, the result needs to be corrected by a factor of $f_{\rm sky}$. The difference is due to the fact that coupling $\Delta\phi$ with an isotropic δ can only depend on the monopole of $\Delta\phi$, which is zero by construction. On the other hand, the sky-mean of a survey mask is precisely equal to $f_{\rm sky}$. The analogy becomes more explicit in equation (A.6), where we allow $\Delta\phi$ to have a non-zero monopole.

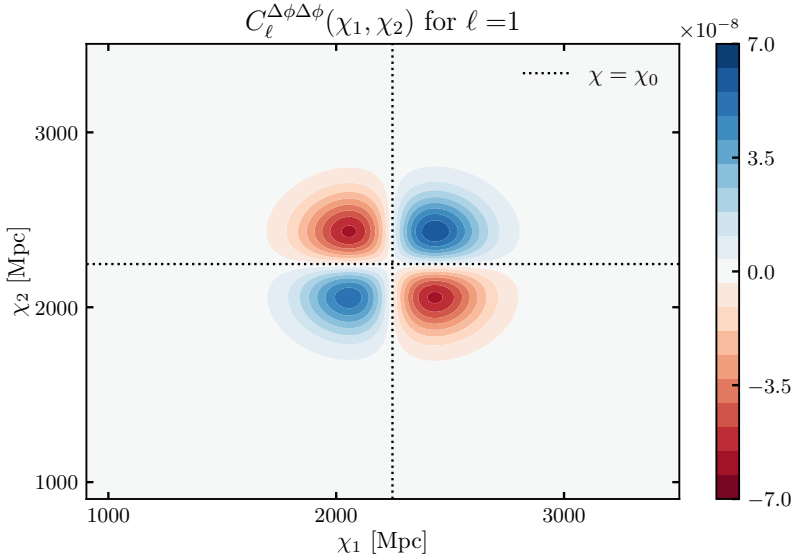

Figure 1: Multi-frequency angular power spectrum (MAPS) for our idealized example, where perturbations in the mean redshift of the distribution are seeded at $\ell = 1$. In this example, $\ell = 1$ is the only scale containing any structure at all. Dotted lines mark the comoving distance to the center of the galaxy distribution, χ_0. Because the sign of $\Delta\phi$ changes either side of χ_0, so does the sign of the MAPS. And since the additive bias comes from an integral of the MAPS over all of χ_1 and χ_2, cancellations make this contribution negligibly small.

For definiteness, let us consider the specific example of a Gaussian redshift distribution with width $\sigma_0 = 0.06$ centered at $z_0 = 0.59$, and variations about the mean redshift with amplitude $\epsilon = 0.3 \times 10^{-2}$. As we will see in greater detail in §4.1, our choices are inspired by DES's REDMAGIC sample [23]. We also adopt a galaxy power spectrum appropriate for this sample and described in that same section. Translated to distance units (in the Planck 2018 [24, 25] cosmology), these parameter values correspond to $\chi_0 = 2241\,\text{Mpc}$, $\sigma_0 = 187\,\text{Mpc}$ and $\epsilon = 9.5\,\text{Mpc}$.

The integrand of equation (2.36) can be visualized in figure 1. Note the extensive cancellations between regions where the integrand takes on opposite signs, leading us to expect that the additive bias will ultimately be very small. This can be seen more explicitly if we integrate by parts:

$$\int_{\chi=0}^{\infty} d\chi \left\{ H \frac{d^2\bar{n}}{d\chi dz} \right\} (\chi) = \left[H \frac{d\bar{n}}{dz} \right]_{\chi=0}^{\infty} - \int_{\chi=0}^{\infty} d\chi \left\{ \frac{dH}{d\chi} \frac{d\bar{n}}{dz} \right\} (\chi) \,. \tag{2.37}$$

The first term on the right vanishes because the dN/dz is zero at the boundary, and the second term is very small because H varies slowly over a typical redshift distribution – in the limit that H is constant over this range, the integrals are exactly zero. The narrower the mean dN/dz, the smaller this contribution will be. For the scenario at hand, the additive bias is consistent with zero up to numerical error.

Meanwhile, the mode-coupling matrix also takes a transparent form

$$M_{\ell L}^{\Delta\phi^a \Delta\phi^b}(\chi_1, \chi_2) = (2L+1) \begin{pmatrix} \ell & L & 1 \\ 0 & 0 & 0 \end{pmatrix}^2 \frac{\epsilon_a \epsilon_b}{8\pi c^2} \Psi^a(\chi_1) \Psi^b(\chi_2) \,, \tag{2.38}$$

with the triangle conditions of the $3j$ symbol imposing the simplest of off-diagonal couplings,

$$\begin{pmatrix} \ell & L & 1 \\ 0 & 0 & 0 \end{pmatrix}^2 = \delta_{L\,(\ell-1)} \begin{pmatrix} \ell & \ell-1 & 1 \\ 0 & 0 & 0 \end{pmatrix}^2 + \delta_{L\,(\ell+1)} \begin{pmatrix} \ell & \ell+1 & 1 \\ 0 & 0 & 0 \end{pmatrix}^2 \tag{2.39}$$

$$= \delta_{L\,(\ell-1)} \frac{\ell}{(2\ell+1)(2\ell-1)} + \delta_{L\,(\ell+1)} \frac{\ell+1}{(2\ell+3)(2\ell+1)} \,. \tag{2.40}$$

The range of ℓs over which $C_\ell^{\Delta\phi^a\Delta\phi^b}(\chi_1,\chi_2)$ has support determines the width of the convolution kernel to which P_{ab} is subjected. Had $\Delta\phi$ had structure on smaller angular scales, the mode-coupling would have involved a wider range of L's. With the mode-coupling matrix above, the integrand of the mode-coupling bias becomes

$$\begin{aligned} R_\ell^{ab} &= \sum_L M_{\ell L}^{\Delta\phi^a\Delta\phi^b}(\chi_1,\chi_2) C_L^{ab}(\chi_1,\chi_2) \\ &= \frac{\epsilon_a\epsilon_b}{8\pi c^2}\Psi^a\Psi^b \left[(2\ell-1) C_{\ell-1}^{ab} \begin{pmatrix} \ell & \ell-1 & 1 \\ 0 & 0 & 0 \end{pmatrix}^2 + (2\ell+3) C_{\ell+1}^{ab} \begin{pmatrix} \ell & \ell+1 & 1 \\ 0 & 0 & 0 \end{pmatrix}^2 \right] , \end{aligned} \tag{2.41}$$

giving

$$\begin{aligned} T_\ell^{ab} &\supset \int d\chi_1\, d\chi_2\; R_\ell^{ab}(\chi_1,\chi_2) \\ &= \frac{\epsilon_a\epsilon_b}{8\pi c^2}(2\ell-1) \begin{pmatrix} \ell & \ell-1 & 1 \\ 0 & 0 & 0 \end{pmatrix}^2 \int d\chi_1\, \Psi^a(\chi_1) \int d\chi_2\, \Psi^b(\chi_2)\; C_{\ell-1}^{ab}(\chi_1,\chi_2) \\ &+ \frac{\epsilon_a\epsilon_b}{8\pi c^2}(2\ell+3) \begin{pmatrix} \ell & \ell+1 & 1 \\ 0 & 0 & 0 \end{pmatrix}^2 \int d\chi_1\, \Psi^a(\chi_1) \int d\chi_2\, \Psi^b(\chi_2)\; C_{\ell+1}^{ab}(\chi_1,\chi_2) \\ &\approx \frac{\epsilon_a\epsilon_b}{8\pi c^2} \int \frac{d\chi}{\chi^2} \Psi^a(\chi)\Psi^b(\chi) \\ &\qquad \times \left\{ (2\ell-1) \begin{pmatrix} \ell & \ell-1 & 1 \\ 0 & 0 & 0 \end{pmatrix}^2 P_{ab}\left(\frac{\ell-1/2}{\chi}; z\right) \right. \\ &\qquad\quad \left. + (2\ell+3) \begin{pmatrix} \ell & \ell+1 & 1 \\ 0 & 0 & 0 \end{pmatrix}^2 P_{ab}\left(\frac{\ell+3/2}{\chi}; z\right) \right\} . \end{aligned} \tag{2.42}$$

In the second equality, we have used equation (C.2), the Limber approximation for C_ℓ^{ab}. When we evaluate this expression, we find it to be in very good agreement with a full calculation of equation (2.25) using simulated, perturbed distributions, underestimating it by only 1% below $\ell \lesssim 2000$.

The advantage of this analytic route, however, is that it sheds light on how the mode-coupling bias depends on the characteristics of the problem. Except on the very largest angular scales we can approximate $P_{ab}\left(\frac{\ell+3/2}{\chi}; z\right) \approx P_{ab}\left(\frac{\ell-1/2}{\chi}; z\right) \approx P_{ab}\left(\frac{\ell+1/2}{\chi}; z\right)$ with an error of $\mathcal{O}(\ell^{-1})$, so

$$\begin{aligned} \Delta T_{\ell\gg1}^{ab} \approx & \frac{\epsilon_a\epsilon_b}{8\pi c^2} \left\{ (2\ell-1) \begin{pmatrix} \ell & \ell-1 & 1 \\ 0 & 0 & 0 \end{pmatrix}^2 + (2\ell+3) \begin{pmatrix} \ell & \ell+1 & 1 \\ 0 & 0 & 0 \end{pmatrix}^2 \right\} \\ & \times \int \frac{d\chi}{\chi^2}\; \Psi^a(\chi)\Psi^b(\chi)\; P_{ab}\left(\frac{\ell+1/2}{\chi}; z\right) . \end{aligned} \tag{2.43}$$

Then, using the identity [26],

$$\sum_{\ell_3=|\ell_1-\ell_2|}^{\ell_1+\ell_2}(2\ell_3+1)\begin{pmatrix}\ell_1 & \ell_2 & \ell_3\\ 0 & 0 & 0\end{pmatrix}^2=1\,, \tag{2.44}$$

the term in $\{\cdots\}$ above becomes 1 and we obtain

$$\begin{aligned}\Delta T^{ab}_{\ell\gg1}\approx&\frac{\epsilon_a\epsilon_b}{8\pi c^2}\int\frac{d\chi}{\chi^2}\;\Psi^a(\chi)\Psi^b(\chi)\;P_{ab}\left(\frac{\ell+1/2}{\chi};z\right)\\ \equiv&\int d\chi\left[\frac{\bar\phi^a(\chi)\bar\phi^b(\chi)}{\chi^2}\right]f_a(\chi)f_b(\chi)P_{ab}\left(\frac{\ell+1/2}{\chi};z\right)\,.\end{aligned} \tag{2.45}$$

The functions we have defined in the last line are there to facilitate comparison with the cosmological signal, equation (C.3). Compared to the signal, the bias integrand is suppressed by two factors of

$$f_x(\chi)\equiv(8\pi)^{-1/2}(\chi-\chi^x_0)\frac{\epsilon_x}{(\sigma^x_0)^2}\,. \tag{2.46}$$

This tells us that the mode-coupling bias to the power spectrum scales as $\Delta T_\ell\propto\epsilon^2/\sigma_0^4$, where ϵ parametrizes the amplitude of the $\phi(\chi)$ variations, and σ_0 is the width of the mean distribution. Samples with broad redshift distributions are therefore more robust against bias (though remember the signal amplitude also depends upon σ_0).

Moreover, the integration kernel in equation (2.45) is independent of ℓ and probes very similar effective scales and redshifts of the 3D power spectrum, $P_{ab}(k\sim\ell/\chi;z)$, as the standard kernel in equation (C.3) – see the left panel of figure 2. We thus expect that, in this regime where $\ell\gg1$, the mode-coupling bias will have the same shape as the unbiased angular power spectrum, only differing from it by an ℓ-independent amplitude factor.

This qualitative behaviour should also hold in more realistic scenarios as long as $C^{\Delta\phi}_{\ell'}$ only has support at relatively low ℓ', and we are looking at much smaller scales in the angular power spectrum. When this is the case, $\ell\gg\ell'$, and the triangle condition imposes $\ell\approx L\gg\ell'$. Taking as our starting point the Limber-approximated expression in equation (C.6), we can write

$$\Delta T^{ab}_{\ell\gg1}\approx\int\frac{d\chi}{\chi^2}\sum_L M^{\Delta\phi^a\Delta\phi^b}_{\ell L}(\chi,\chi)P_{ab}\left(\frac{L+1/2}{\chi};z\right) \tag{2.47}$$

$$\begin{aligned}&\approx\int\frac{d\chi}{\chi^2}P_{ab}\left(\frac{\ell+1/2}{\chi};z\right)\sum_L M^{\Delta\phi^a\Delta\phi^b}_{\ell L}(\chi,\chi)\\ &\approx\int\frac{d\chi}{\chi^2}P_{ab}\left(\frac{\ell+1/2}{\chi};z\right)\sum_{\ell'}\frac{(2\ell'+1)}{4\pi}C^{\Delta\phi^a\Delta\phi^b}_{\ell'}(\chi)\sum_L(2L+1)\begin{pmatrix}\ell & L & \ell'\\ 0 & 0 & 0\end{pmatrix}^2\,.\end{aligned} \tag{2.48}$$

In going to the last line, we just used the definition of the mode-coupling matrix, equation (2.28). To simplify this expression further, notice that the covariance of the perturbations across a comoving distance slice is

$$\mathrm{Cov}[\phi^a,\phi^b](\chi)=\mathrm{Cov}[\Delta\phi^a,\Delta\phi^b](\chi)=\sum_\ell\frac{(2\ell+1)}{4\pi}C^{\Delta\phi^a\Delta\phi^b}_\ell(\chi)\,. \tag{2.49}$$

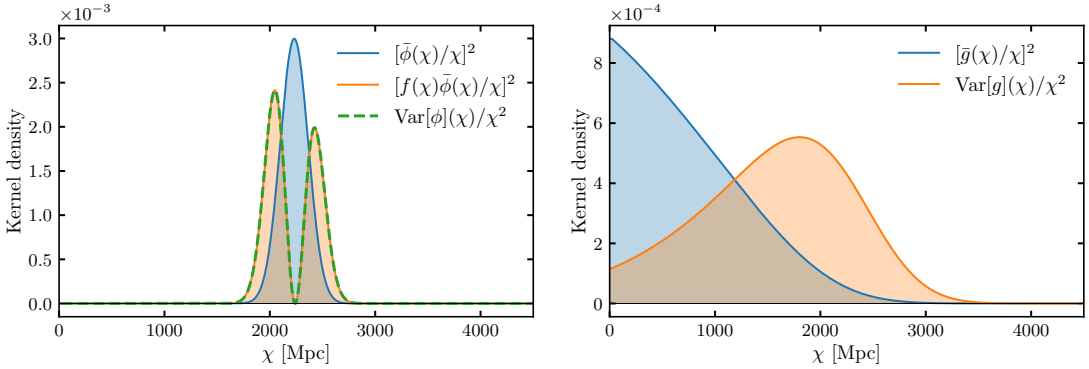

Figure 2: Kernel densities involved in the standard angular power spectrum (blue) and the mode-coupling term in the $\ell \gg 1$ limit. The left panel pertains to our toy example of galaxy clustering, so we show the bias kernel in two different approximations (orange and green, dashed) which are practically indistinguishable from each other. Integrating the blue or orange kernels against $P^{gg}(k \sim \ell/\chi; z)$ yields very similar results, because they are probing similar effective scales and redshifts. Consequently, at large ℓ, the bias to galaxy clustering auto-spectra has very similar shape to the signal. On the other hand, in the case of cosmic shear (the right panel being a typical such example) the effective scales and redshifts probed by the signal and bias kernels differ significantly, so we expect the shape of the two integrals to deviate accordingly.

Finally, using identity (2.44) to do the sum over L, we arrive at the very simple expression

$$\Delta T^{ab}_{\ell\gg1} \approx \int d\chi \, \frac{\mathrm{Cov}[\phi^a, \phi^b](\chi)}{\chi^2} P_{ab}\left(\frac{\ell+1/2}{\chi}; z\right) . \tag{2.50}$$

In hindsight, this motivates equation (2.46): $[f(\chi)\bar{\phi}(\chi)]^2$ is a remarkably good approximation to $\mathrm{Var}[\phi](\chi)$ – see the left panel of figure 2 – which is why it produced such a good approximation to the full result in the context of our toy model[5]. By the same token as above, we expect the shape of the mode-coupling contribution to the galaxy clustering power spectrum to track the signal, though the statement is now more general. In §4.1, figure 5, we verify this insight and show that the expression is indeed an excellent approximation to the full calculation on scales smaller than those where there is significant projection anisotropy. The expression above also suggests that the shape and amplitude of this term relative to the galaxy clustering signal is independent of the underlying power spectrum of the objects and their distance to the observer.

Equation (2.50) is in fact very general. It applies whenever the separation of angular scales is respected, even in cases where the shape of $\mathrm{Cov}[\Phi^a, \Phi^b](\chi)$ is very different from that of $\bar{\Phi}^a(\chi)\bar{\Phi}^b(\chi)$ (where Φ is now some general selection function) as is the case for example with cosmic shear or galaxy-galaxy lensing. In these cases, the various ℓ/χ are weighted differently in the two integrals, and ultimately the shape of the mode-coupling term deviates from that of the signal. The difference in kernels is illustrated in the right panel of figure 2 for cosmic shear, and discussed further in the coming sections.

[5]Equation (2.50) leads to slightly better agreement with simulations – residuals are below 0.1% in this case on scales $\ell \leq 2000$ – because measuring the variance directly on each χ-slice can capture the effect of simulation artefacts like leakage of power between angular scales.

3 Special cases

The formalism above is general enough that it encompasses both auto- and cross-correlations between two projected fields. To elucidate the implications, in the following sections we tackle some special cases of particular observational interest. The full results simplify in these cases.

3.1 Cross-correlations

By far the easiest scenario is when at least one of the fields has a precisely isotropic redshift distribution or weight, such as occurs for example with CMB lensing. As we noted above, the three terms in equation (2.23) come from the auto-correlations of the $\Delta\phi$, $\bar{\phi}\,\delta$ and $\Delta\phi\,\delta$ contributions to $\delta^{(2D)}$.[6] This means that in cross correlation, if one of the fields has no $\Delta\phi$ then only the 'cosmological' signal remains. Both the additive and mode-coupling corrections vanish.

On the other hand, when the two fields being cross-correlated have projection anisotropy, a mode-coupling contribution is possible. A particularly likely and interesting scenario is when this projection anisotropy is correlated across the two fields. In §3.4, we explain that this contribution will be negative (or positive if the anisotropy is anti-correlated across fields). Then, in §4.1, we quantify its expected impact on current surveys.

3.2 Cosmic shear

A second case, where at least one of the fields has mean zero, retains the cosmological signal and the mode-coupling term. However, in this case the additive contribution, $C_\ell^{\Delta\phi}$, vanishes. Let us now show this explicitly for the case of cosmic shear.

Absent variations in the redshift distribution of the source galaxies the lensing convergence is given by[7]

$$\kappa(\hat{\boldsymbol{n}}) = \int d\chi\; \bar{g}(\chi)\,\delta(\chi,\hat{\boldsymbol{n}})\,, \tag{3.1}$$

where the lens efficiency kernel is defined as

$$\bar{g}(\chi) \equiv \frac{3}{2}\Omega_{m,0}\frac{H_0^2}{c^2}\frac{\chi}{a(\chi)}\int_\chi^\infty d\chi_s \frac{H(\chi_s)}{c}\frac{(\chi_s-\chi)}{\chi_s}\frac{d\bar{n}_g}{dz}(\chi_s)\,. \tag{3.2}$$

As with galaxy clustering, the redshift distribution of source galaxies is typically normalized such that $\int dz(d\bar{n}_g/dz) = 1$. If required, convergence can then be related to shear using the Kaiser-Squires method [27].

Suppose now that the photometry varies across the sky, so that we have fluctuations around the mean redshift distribution of the source galaxy sample in different sky locations. We can introduce $\hat{\boldsymbol{n}}$ dependence in $d\bar{n}_g/dz$ and hence g in equation (3.2), such that

$$g(\chi,\hat{\boldsymbol{n}}) \equiv \bar{g}(\chi) + \Delta g(\chi,\hat{\boldsymbol{n}})\,. \tag{3.3}$$

As in previous sections, we assume that $\bar{g}$ accurately captures the sky mean, so the monopole of Δg vanishes. In presence of these perturbations, the convergence becomes

$$\kappa(\hat{\boldsymbol{n}}) = \int d\chi\,[\bar{g}(\chi) + \Delta g(\chi,\hat{\boldsymbol{n}})]\,\delta(\chi,\hat{\boldsymbol{n}})\,. \tag{3.4}$$

[6] This assumes that the fiducial dN/dz is set to the footprint mean, as is most often done. When this is not the case, a multiplicative bias is possible, as explained in appendix A.

[7] We assume a spatially-flat Universe and work in the Born approximation throughout.

Notice that the perturbed convergence has no additive contribution from just Δg. In the case of galaxy clustering, this contribution appeared because any non-cosmological variation in the number of sampled galaxies across the footprint can be mistaken for the signal of interest. By contrast, in the context of cosmic shear, a variation in the number of source galaxies only affects the number of measurements available to extract the shear signal (thus imprinting inhomogeneity in the shape noise across the sky) but does not add spurious lensing signal.

Proceeding by analogy with §2.2, we have

$$\kappa_{\ell m} = \int d\chi \, \left[\bar{g}(\chi)\delta_{\ell m}(\chi) + \{\Delta g \delta\}_{\ell m}(\chi)\right] , \tag{3.5}$$

with the difference that δ now denotes the matter instead of the galaxy overdensity. Indeed, the only extra effect is a coupling of the cosmological anisotropy of δ with the newly-induced anisotropy in the lens efficiency kernel.

The impact on the cross-correlation between various tomographic shear bins can be obtained by exact analogy with §2.2, simply replacing $\phi \to g$. On the full sky, the total angular cross-spectrum between bins i and j is

$$\begin{aligned} T_\ell^{ij} =& \frac{1}{2\ell+1}\sum_m \langle \kappa^i_{\ell m} \kappa^{j,*}_{\ell m}\rangle \\ =& \int d\chi_1 \, d\chi_2 \left[\bar{g}^i(\chi_1)\bar{g}^j(\chi_2) C_\ell^{ij}(\chi_1,\chi_2) + \sum_L M_{\ell L}^{\Delta g^i \Delta g^j}(\chi_1,\chi_2) C_L^{ij}(\chi_1,\chi_2)\right] . \end{aligned} \tag{3.6}$$

As before, the absence of cross-terms is due to Δg having a vanishing monopole. The first term is the usual expression, while the second is a multiplicative bias. Limber approximating following appendix C, we find

$$T_\ell^{ij} \approx \int \frac{d\chi}{\chi^2} \left[\bar{g}^i(\chi)\bar{g}^j(\chi) P_{ij}\left(\frac{\ell+1/2}{\chi}; z\right) + \sum_L M_{\ell L}^{\Delta g^i \Delta g^j}(\chi,\chi) P_{ij}\left(\frac{L+1/2}{\chi}; z\right)\right] . \tag{3.7}$$

In section §3.4, we explain that this new term is expected to be positive for both auto- and cross-correlation of tomographic cosmic shear measurements. On scales smaller than those on which there is significant projection anisotropy, the second term above can be approximated with exquisite accuracy by equation (2.50), just replacing ϕ with g (see, e.g., figure 9).

3.3 Galaxy-galaxy lensing

Similarly, cross-correlations of the cosmic shear signal with a sample of lens galaxies are susceptible only to the multiplicative, mode-coupling bias. Given the machinery we have developed, it is easy to show that variations in the redshift distributions of lens and source galaxy samples lead to a measured angular spectrum of the form

$$\begin{aligned} T_\ell^{\kappa g} =& \frac{1}{2\ell+1}\sum_m \langle \kappa_{\ell m} \delta^{(2D),*}_{\ell m}\rangle \\ =& \int d\chi_1 \, d\chi_2 \left[\bar{g}(\chi_1)\bar{\phi}(\chi_2) C_\ell(\chi_1,\chi_2) + \sum_L M_{\ell L}^{\Delta g \Delta \phi}(\chi_1,\chi_2) C_L(\chi_1,\chi_2)\right] , \end{aligned} \tag{3.8}$$

where $M_{\ell L}^{\Delta g \Delta \phi}$ is defined by replacing $\{\Delta\phi^a, \Delta\phi^b\} \to \{\Delta g, \Delta\phi\}$ in equations (2.28) and (2.18). Likewise, if projection anisotropy is confined to large angular scales, the mode-coupling term

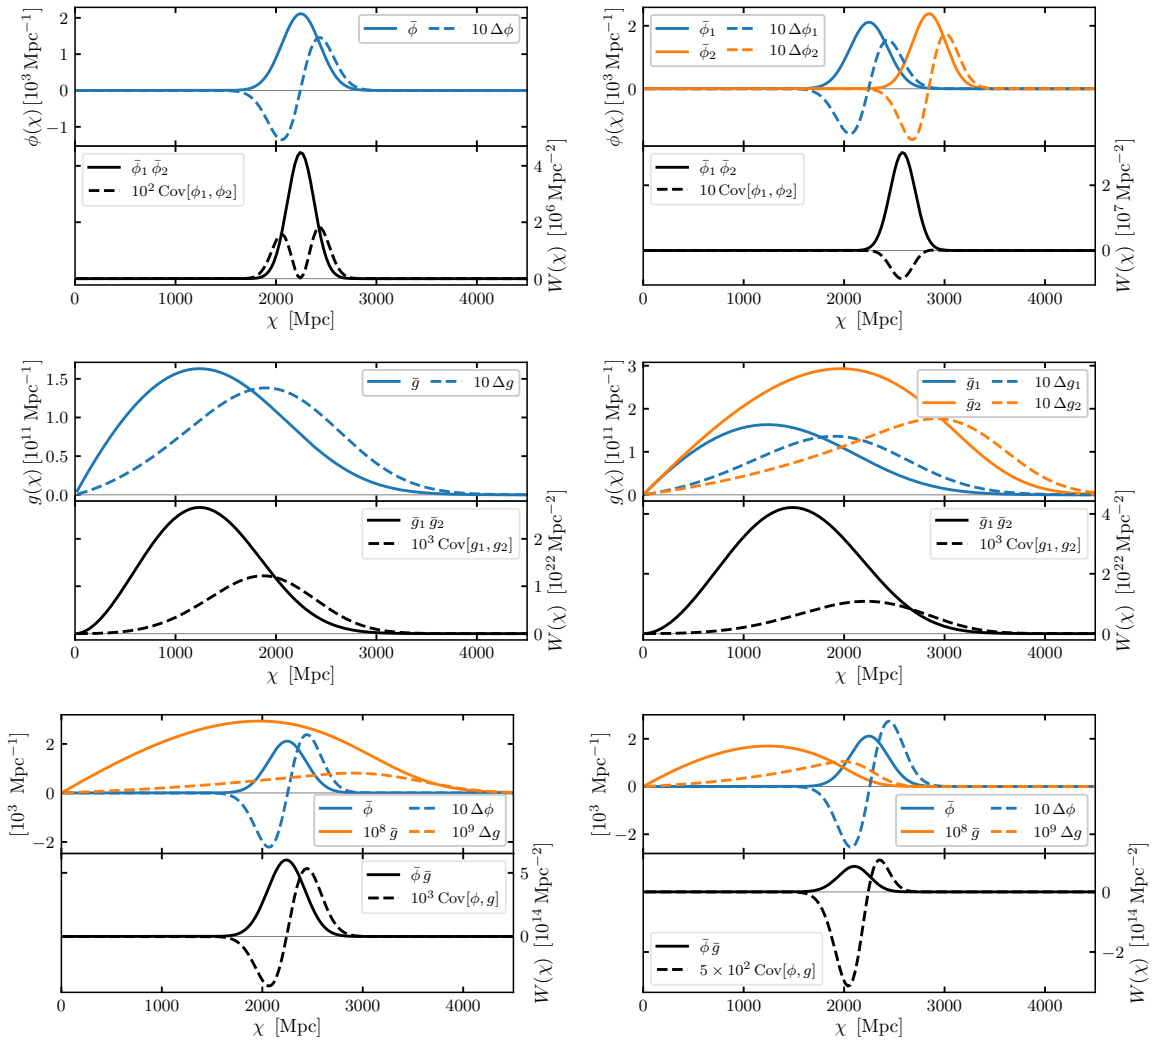

Figure 3: The signal and mode-coupling kernels for certain special scenarios discussed in §3. These include galaxy clustering auto- (top left) and cross-correlations (top right), the auto- (middle left) and cross-correlation (middle right) of cosmic shear tomographic bins, and galaxy-galaxy lensing analyses where the source galaxies are far behind (bottom left) or close to (bottom right) the lenses. Within each panel, the top plot shows the unperturbed selection function (solid) along with a random perturbation to it (dashed) associated with a shift of the distribution to higher redshift. The bottom panel then shows the signal kernel (black, solid) and an approximation to the mode-coupling kernel (black, dashed).

can be approximated on smaller scales by substituting $\{\phi^a, \phi^b\} \to \{g, \phi\}$ in equation (2.50). As we will now see, the magnitude and sign of this new contribution depends subtly on the relative arrangement of lens and source galaxy distributions.

3.4 Implications of the kernels

The kernels for the signal and the mode-coupling term for the special cases above are illustrated in figure 3, from which we can understand much of the resulting phenomenology. To

shift the central redshift of a galaxy bin, $\Delta\phi$ must take the form of a wiggle. The auto-covariance of this wiggle is a function with two peaks on either side of the peak of $\bar{\phi}$. We showed in figure 1 that this structure was an excellent approximation to the full result and led to a bias that had the same shape as the signal. On the other hand, a shift of the source galaxy distribution to higher/lower redshift is associated with a Δg that is consistently positive/negative across χ.

This has interesting implications if the redshift anisotropy is correlated across both legs of the two-point function: while the mode-coupling contribution to galaxy clustering auto-spectra (figure 3; left, top) and cosmic shear auto- (left, middle) and cross-spectra (right, middle) are always positive, the contribution to the cross-spectrum of different galaxy density bins (right, top) is expected to be negative. The case of galaxy-galaxy lensing is more nuanced: when the source galaxies are far behind the lenses (left, bottom), the mode-coupling kernel is a wiggle that should produce negligible contributions; however, when source and lens galaxies are sufficiently close to each other (right, bottom), Δg overlaps only with the lower-χ part of the $\Delta\phi$ wiggle, and since the two are of opposite sign, the mode-coupling kernel is primarily negative and can potentially lead to significant bias. We shall illustrate these general points with some numerical examples in the following section.

4 Examples

It is worth quantifying the effects we have been describing in the context of some concrete examples. To do so, we will come up with a fiducial redshift distribution and perturb it. We take the galaxy redshift distribution to be Gaussian in comoving distance, albeit a different one in every pixel of a HEALPIX [28] pixelization:

$$\frac{dN_g}{dz}(\chi,\hat{\boldsymbol{n}}) = \frac{1}{\sqrt{2\pi}\,\sigma_0} e^{-[\chi-\chi_0'(\hat{\boldsymbol{n}})]^2/2\sigma_0^2} \quad , \quad \chi_0'(\hat{\boldsymbol{n}}) \equiv \chi_0 + \chi_{\rm shift}(\hat{\boldsymbol{n}})\,, \tag{4.1}$$

with χ_0 and σ_0 some fiducial central distance and standard deviation of the distributions appropriate for the sample at hand, and $\chi_{\rm shift}(\hat{\boldsymbol{n}})$ a random perturbation to the mean.

We generate HEALPIX `nside`=128 templates of these perturbations by drawing harmonic coefficients from independent, zero-mean Gaussian distributions with power-law angular power spectra, $C_\ell \propto \ell^\alpha$, up to some cutoff scale $\ell'_{\rm max}$, and adjusting the normalization so that the resulting template map has the desired variance (the quantity that can more easily be quantified observationally). Note that our code, CARDIAC[8], can also take in user-defined templates of $\chi_{\rm shift}(\hat{\boldsymbol{n}})$ and/or a shift in the width $[\sigma_{\rm shift}(\hat{\boldsymbol{n}})]$ and from them calculate the expected contributions.

We obtain the fiducial redshift distribution as an average of (4.1) over the footprint, construct the fiducial selection function (after assuming a cosmology) and normalize it to satisfy $\int \bar{\phi}\, d\chi = 1$. We then use this same normalization to obtain dn_g/dz from the dN_g/dz in equation (4.1) and thus obtain a spatially-varying selection function as

$$\phi(\chi,\hat{\boldsymbol{n}}) \equiv \frac{H(\chi)}{c}\frac{dn_g}{dz}(\chi,\hat{\boldsymbol{n}})\,. \tag{4.2}$$

Finally, we isolate the perturbation as $\Delta\phi(\chi,\hat{\boldsymbol{n}}) = \phi(\chi,\hat{\boldsymbol{n}}) - \bar{\phi}(\chi)$.

[8]Code for Anisotropic Redshift Distributions in Angular Clustering: https://github.com/abaleato/CARDiAC.

Note that $\phi(\chi, \hat{\boldsymbol{n}})$ still respects the integral constraint

$$\int d\chi \int_{\mathcal{A}} d\hat{\boldsymbol{n}}\, \phi(\chi, \hat{\boldsymbol{n}}) / \int_{\mathcal{A}} d\hat{\boldsymbol{n}} = 1 \iff \int d\chi \int_{\mathcal{A}} d\hat{\boldsymbol{n}}\, \Delta\phi(\chi, \hat{\boldsymbol{n}}) = 0$$
$$\iff \sum_{\ell m} \int d\chi\, \Delta\phi_{\ell m}(\chi) \int_{\mathcal{A}} d\hat{\boldsymbol{n}}\, Y^*_{\ell m}(\hat{\boldsymbol{n}}) = 0\,; \tag{4.3}$$

as long as

$$\int d\chi\, \Delta\phi_{00}(\chi) = 0\,, \tag{4.4}$$

but this is satisfied trivially because $\Delta\phi_{00} = 0$ by construction[9]. This means that, besides the condition on its average over the analysis region, $\Delta\phi$ is unconstrained.

4.1 Galaxy clustering

As a somewhat realistic example, let us consider galaxy samples loosely inspired by the Dark Energy Survey's (DES) RedMaGiC and MagLim selections, which have been used in several cosmological analyses (e.g., ref. [23]), both for studies of angular clustering and as samples of lens galaxies in galaxy-galaxy lensing, and estimate the new contributions due to projection anisotropy.

Our calculations will involve a model for the galaxy power spectrum. We obtain this from the Anzu[10] code, which combines N-body simulations of the dark matter component with an analytic treatment of Lagrangian galaxy bias – we use the best-fit bias values measured by [29] from simulated samples of RedMaGiC-like galaxies (similar to those in [30]) at $z = 0.59$. We assume the fiducial dN/dz is a Gaussian centered at this redshift of $z = 0.59$, and with a standard deviation of 0.06 in redshift. This resembles the dN/dz of the third redshift bin of both the RedMaGiC and MagLim samples (see, e.g., figure 1 of [23]). Note that equation (2.50) suggests that two samples with similar mean redshift distribution and projection anisotropy will see mode-coupling terms with similar shape and amplitude relative to the signal, independent of the exact galaxy power spectrum (and the same holds for the additive term). Hence, we do not need to consider separately a MagLim-specific galaxy power spectrum.

We generate templates of variations in the mean redshift of the dN/dz, with characteristics summarized in table 1, and two such examples shown in figure 4. As a first approximation, the variance across the shifts template can be related to an uncertainty in the determination of the mean redshift of a given sample (though note that this is likely to underestimate the true variations present in the data). Some values in the literature can give us a sense of the scale of variations appropriate for DES: in their cosmological constraints, the width of their Gaussian prior on the mean redshift of the third bin is $\sigma(z_0) = 0.003$ for RedMaGiC, and $\sigma(z_0) = 0.006$ for MagLim [23] both more than an order of magnitude smaller than the bin width. The mean uncertainty on the redshift of a given RedMaGiC galaxy in this bin is much higher, at $\sigma_z = 0.027$ [31].

We then calculate the mode-coupling contribution in each case working in the Limber approximation; i.e., evaluating equation (C.6). We find that the shape of the mode-coupling

[9] On a cut sky, there are in principle additional conditions on certain other $\Delta\phi_{\ell m}$'s, those with ℓ so small that the corresponding $Y_{\ell m}$'s do not average to zero over the footprint. However, for this same reason, they will be indistinguishable from the monopole for all practical purposes, and will therefore get absorbed into $\bar{\phi}$.

[10] https://github.com/kokron/anzu

z_0	χ_0 [Mpc]	$\sigma_0 \times 10^2$	σ_0 [Mpc]
0.59	2241	3, 6, 9, 12	95, 187, 279, 368

(a) Parameters characterizing the fiducial dN/dz.

α	$\ell'_{\rm max}$	$\sigma(z_{\rm shift}) \times 10^2$	$\sigma(\chi_{\rm shift})$ [Mpc]
0, −1, −2	100	0.3, 0.6, 1.0, 2.7	9.5, 19, 32, 85

(b) Parameters characterizing the perturbations.

Table 1: Parameters we consider in our galaxy clustering example. We set $\Delta\phi^a = \Delta\phi^b$ in order to study the impact on the galaxy clustering power spectrum. Values given in comoving distance units are converted from redshifts assuming the Planck 2018 [24, 25] ΛCDM cosmology.

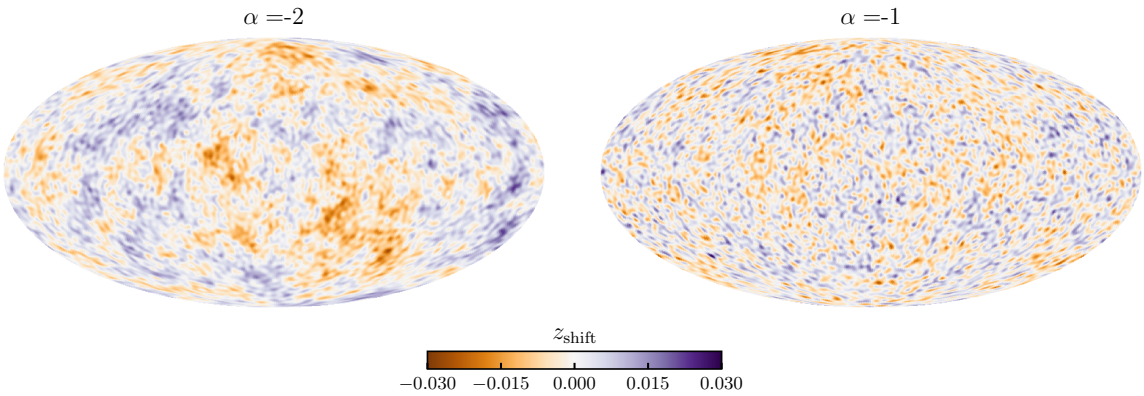

Figure 4: Example variations of $z_{\rm shift}$ across the sky consistent with the MagLim uncertainty of $\sigma(z_{\rm shift}) = 0.006$. We draw these from Gaussian distributions with power-law angular spectra $C_\ell \propto \ell^\alpha$, truncated at $\ell_{\rm max} = 100$.

term tracks the unperturbed signal very closely for $\ell \gg 1$ – i.e., $\Delta C^{gg}_{\ell\gg1}/C^{gg}_{\ell\gg1}$ is flat – as expected from the discussion around equation (2.46). Moreover, figure 5 demonstrates that the analytic approximation we developed in equation (2.50) is in excellent agreement with a full calculation of equation (C.6). This is a consequence of the projection anisotropy being confined to large angular scales – though see figure 6 for the response of the mode-coupling term to changes in the $\ell'_{\rm max}$ cutoff, which appears to be small.

Given the flatness of $\Delta C_\ell/C_\ell$, we can define a 'bias amplitude',

$$\int_{\ell=50}^{1500} d\ell \, (\Delta C_\ell/C_\ell) \, / \int_{\ell=50}^{1500} d\ell \,, \tag{4.5}$$

where the lower end of the range of integration is set by the validity of the Limber approximation. In figure 7 we plot this metric against σ_0, the width of the fiducial distribution, for various galaxy samples. At fixed σ_0, the bias amplitude scales as $\Delta C_\ell/C_\ell \propto \sigma^2(\chi_{\rm shift})$, as identified in equation (2.46). That same equation also predicted the scaling $\Delta C_\ell \propto \sigma_0^{-4}$; however, making the dN/dz narrower also increases the clustering signal. All in all, the bias

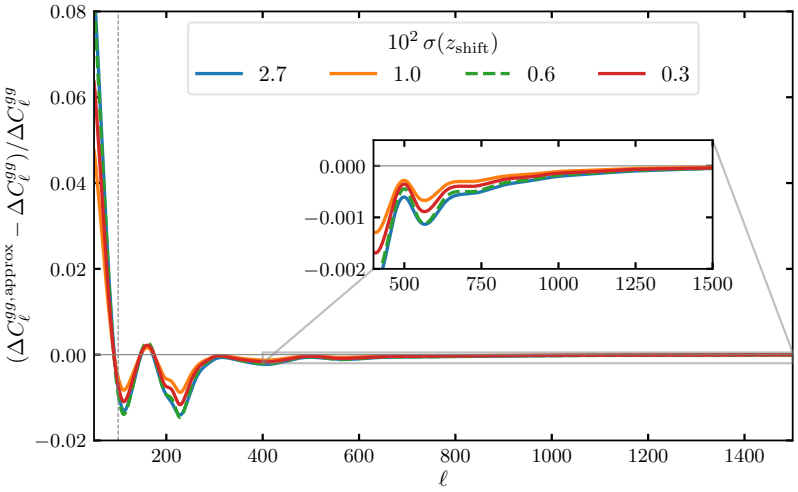

Figure 5: Accuracy of the analytic approximation to the mode-coupling contribution, equation (2.50), in the limit where projection anisotropy is restricted to $\ell' < \ell'_{\rm max} = 100$ (dashed, vertical line); in the cases shown, the anisotropy falls off as a power law below the cut, with spectral index $\alpha = -2$ though the result is little changed with $\alpha = -1, 0$. The approximation is excellent above scales where there is significant projection anisotropy, for reasons explained in §4.1.

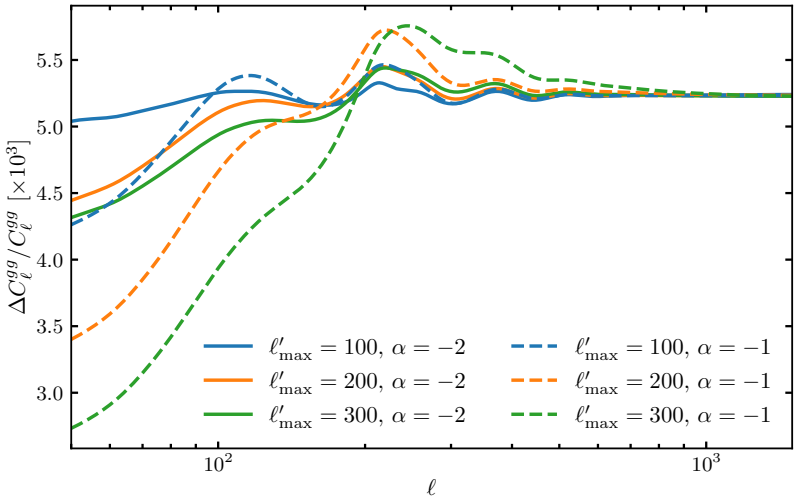

Figure 6: Response of the mode-coupling term to the maximum scale on which there is projection anisotropy. Underlying all curves is the same standard deviation for the mean redshift variations, which we fix to the MagLim value of $\sigma(z_{\rm shift}) = 0.6 \times 10^{-2}$, though results are qualitatively the same for other values of $\sigma(z_{\rm shift})$. The redder the anisotropy spectrum, the faster we converge to the $\ell \gg \ell'$ prediction of flatness from equation (2.50).

amplitude goes approximately as

$$\frac{\Delta C_\ell}{C_\ell} \propto \left[\frac{\sigma(\chi_{\rm shift})}{\sigma_0}\right]^2 \quad . \tag{4.6}$$

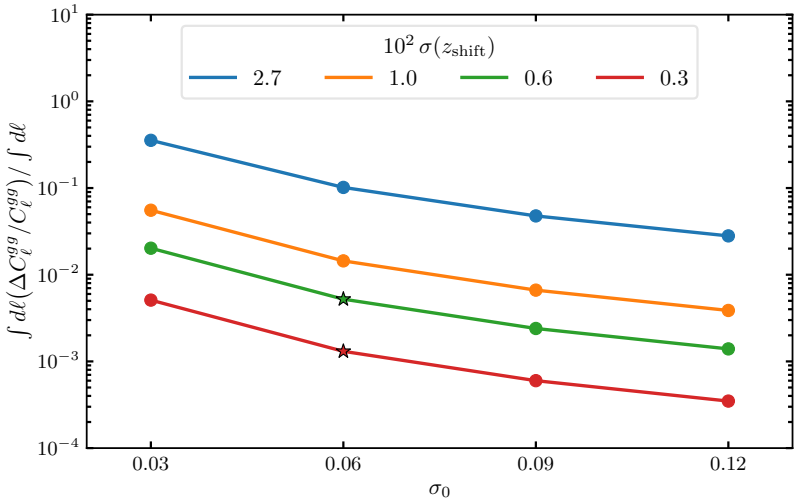

Figure 7: Fractional amplitude of the mode-coupling contribution to the galaxy clustering power spectrum for the samples described in table 1 and the text. The standard deviation of the injected shifts in central redshift, $\sigma(\chi_{\rm shift})$, can be compared to uncertainties on the mean redshift of actual samples: the green curve roughly corresponds to MagLim errors, and the red to RedMaGiC. We consider various values for the width of the fiducial distribution, σ_0, and denote with stars the approximate characteristics of the MagLim and RedMaGiC samples. In the cases shown, the anisotropy follows a power-law with spectral index $\alpha = -2$, truncated at $\ell'_{\rm max} = 100$, though the result is little changed with $\alpha = -1, 0$ (see figure 6).

The stars in figure 7 denote the rough characteristics of the DES lens galaxy samples. If their redshift uncertainties and distribution width are well characterized, our results suggest that the mode-coupling effect should be negligible for them: a 0.5% and 0.1% correction for the MagLim and RedMaGiC bins we have looked at, respectively. We also calculate the additive contribution and find it to be completely negligible.

So far, we have looked only at the auto-correlation of galaxy overdensity bins. However, in §3.4, we anticipated that the cross-correlation of two samples that are only partially overlapping should be especially affected by mode-coupling biases when the anisotropy is correlated across both samples. To put this on a more quantitative footing, we consider two bins inspired by the DES lens galaxy samples: one centered at $z_0 = 0.59$, the other at $z_0 = 0.79$, and both with width $\sigma_0 = 0.06$. We allow for a level of anisotropy consistent with the RedMaGiC $[\sigma(z_{\rm shift}) = 0.003]$ or MagLim errors $[\sigma(z_{\rm shift}) = 0.006]$, make the variations common to both fields, and propagate this through to biases on the angular cross-spectrum.

We show our results in figure 8. As expected from the reasoning in §3.4 the biases are negative, have very similar shape to the signal, and are very well approximated by the analytic expression in equation (2.50) based on the anisotropy covariance. This effect could be present at the level of a couple percent for MagLim, or half a precent for RedMaGiC.

4.2 Cosmic shear

Next, let us explore a quantitative example in the realm of cosmic shear. For simplicity, we will consider the shear auto-spectrum of source galaxies in a single redshift bin. The more general case of cross-correlations between bins can be studied as required using our publicly-available

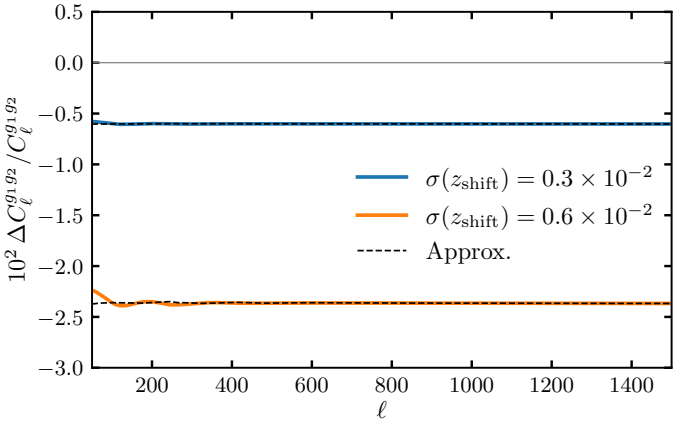

Figure 8: Fractional mode-coupling bias to the cross-correlation of two partially-overlapping galaxy density bins in the limit where the projection anisotropy is common to both bins. The cases we show roughly correspond to cross-correlating two adjacent bins of the REDMAGIC (blue) or MAGLIM (orange) samples. The dashed lines show an analytic approximation based on the anisotropy covariance, equation (2.50). Note the negative sign of this contribution.

z_0	χ_0 [Mpc]	$\sigma_0 \times 10$
0.2, 0.4, 0.7, 1	844, 1600, 2579, 3396	1, 2, 3, 4

(a) Parameters characterizing the fiducial dN/dz of the source galaxies.

α	$\ell'_{\rm max}$	$\sigma(z_{\rm shift}) \times 10^2$
-2	100	0.1, 0.6, 1.1, 1.5, 2.0

(b) Parameters characterizing the perturbations.

Table 2: Parameters we consider in our cosmic shear examples. We set $\Delta g^a = \Delta g^b$ in order to study the impact on the cosmic shear auto-spectrum. Values given in comoving distance units are converted from redshifts assuming the Planck 2018 [24, 25] ΛCDM cosmology. Since the conversion is different at each redshift, we quote the width of the fiducial and perturbation distributions in redshift units only.

code, CARDiAC. From the discussion in §3.4, we expect no major qualitative difference between the two scenarios.

As in the previous example, we generate templates of $z_{\rm shift}$ – a spatially-varying shift in the mean redshift of the source galaxy dN/dz – with variance $\sigma(z_{\rm shift})$, for several choices of the width of the fiducial distribution, parametrized by σ_0. For this example we consider also various values for the central redshift of the fiducial distribution; for galaxy clustering we did not do this because we expect $\Delta C_\ell^{gg}/C_\ell^{gg}$ to be independent of z_0, all other things being equal. The range of parameter values we explore is detailed in table 2. These ranges encompass values that roughly correspond to the DES source galaxy samples summarized in figure 1 and table 1 of ref. [23]; where relevant, we will identify them as such in our plots. The last ingredient we need to evaluate equation (3.7) is the non-linear matter power spectrum, for which we use the ref. [32] version of the HaloFit prescription [33] as implemented in CAMB [34].

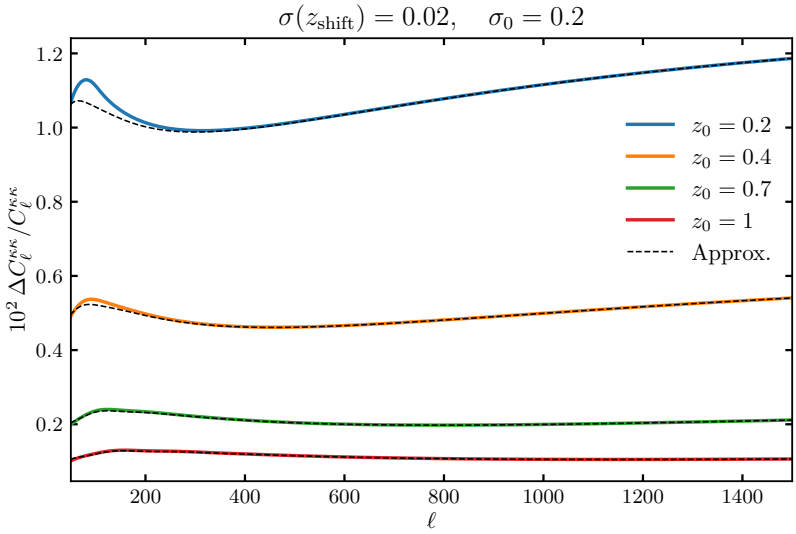

Figure 9: Fractional amplitude of the mode-coupling contribution to the shear power spectrum for a specific example where the source galaxy distribution is centered at z_0, has width $\sigma_0 = 0.2$, and variations about the central redshift with standard deviation $\sigma(z_{\rm shift}) = 0.02$ (similar to the DES source galaxy samples). The mode-coupling term has the shape of the signal when the sources are at high redshift, and begins to deviate from it when they are below $z_0 \lesssim 0.4$. In either regime, equation (2.50) provides an excellent approximation (dashed) to the full calculation (solid) at high ℓ. The behavior seen here extends to other values of σ_0 and $\sigma(z_{\rm shift})$.

Considering once again the limit where the projection anisotropy is confined to large angular scales ($\ell'_{\rm max} = 100$, $\alpha = -2$), we find that the analytic approximation in equation (2.50) – replacing ϕ with g – is in exquisite agreement with the full calculation on scales smaller than the injected anisotropy, as evidenced by how accurately the dashed curves (the approximation) track the solid ones (the full calculation) in the example scenario of figure 9. At high enough redshift, $z_0 \gtrsim 0.4$, the shape of the mode-coupling term resembles that of the signal. However, when the source galaxies are closer by, the integration kernels of bias and signal begin to differ significantly (see e.g. the right panel of figure 2) and thus $\Delta C_\ell^{\kappa\kappa}/C_\ell^{\kappa\kappa}$ begins to deviate from flatness, with the departure being greater for narrower source galaxy distributions (lower σ_0). But even when this is the case, we find our analytic expression to remain an excellent approximation.

Figure 10 shows the average fractional bias in $C_\ell^{\kappa\kappa}$ as a function of redshift width (σ_0) for four different mean redshifts (z_0). Different lines indicate the scale of variation of the mean redshift across the sky, $\sigma(z_{\rm shift})$. Black stars mark values approximately consistent with those of source galaxy samples in DES, from which we see that the bias is expected to be subdominant to the statistical error except at the lowest redshift bin, where it can be a percent-level effect.

4.3 Galaxy-galaxy lensing

Having looked at galaxy clustering and cosmic shear, let us study the intersection of the two: galaxy-galaxy lensing. In particular, let us address the case where there is correlated

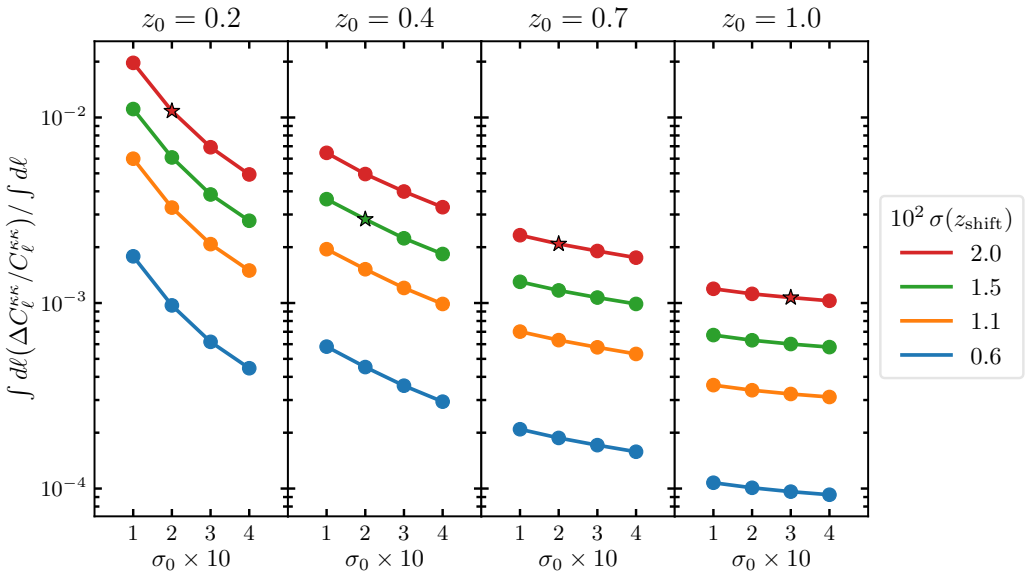

Figure 10: Mean fractional bias to the cosmic shear angular power spectrum for source galaxy redshift bins centered at z_0 and with width σ_0, as a function of $\sigma(z_{\rm shift})$, the scale of variation of the mean redshift across the sky. Stars denote the approximate characteristics of the DES samples.

projection anisotropy across the lens and source galaxy samples. In § 3.4, we flagged this scenario as being of special interest because when this is the case and both distributions are close together, the mode-coupling bias is expected to be amplified and negative.

To verify this, we consider a lens galaxy sample with width $\sigma_0^{\rm lens} = 0.06$ centered at $z_0 = 0.59$, and source samples centered at $z_0^{\rm source} = 0.59$ and $z_0^{\rm source} = 0.8$ with various possible redshift widths. We allow for a level of projection anisotropy in either sample consistent with DES estimates, but ensure that the anisotropy templates in both fields are scaled versions of each other.

The right panel of figure 11 shows that when the sources are far behind the lenses, the bias is positive and small. This can be understood from the bottom left panel of figure 3: the bias kernel is a wiggle across the integration domain, which leads to extensive cancellations.

On the other hand, when the sources are centered near the lenses, the situation changes dramatically. To understand why, it is useful to keep in mind the bottom right panel of figure 3. For a narrow enough distribution of sources, Δg overlaps only with the low-χ part of the $\Delta\phi$ wiggle, which has its opposite sign when the projection anisotropies in source and lens distributions are positively correlated. Equipped with this intuition, we can understand the left panel of figure 11, which quantifies the bias in the limiting case where both lenses and sources are centered at the same redshift[11]. Indeed, the biases are all negative, and grow rapidly with decreasing $\sigma_0^{\rm source}$. The intuition developed around figure 3 tells us that as $\sigma_0^{\rm source}$ drops, the bias kernel must be converging to a single peak near $\bar{\phi}\bar{g}$, which is why the shapes of signal and bias become more and more similar. All in all, these effects ought to be below percent-level for DES.

[11]Though we do not show it explicitly in the figures, we verify that in all cases equation (2.50) provides an excellent approximation to the full calculation.

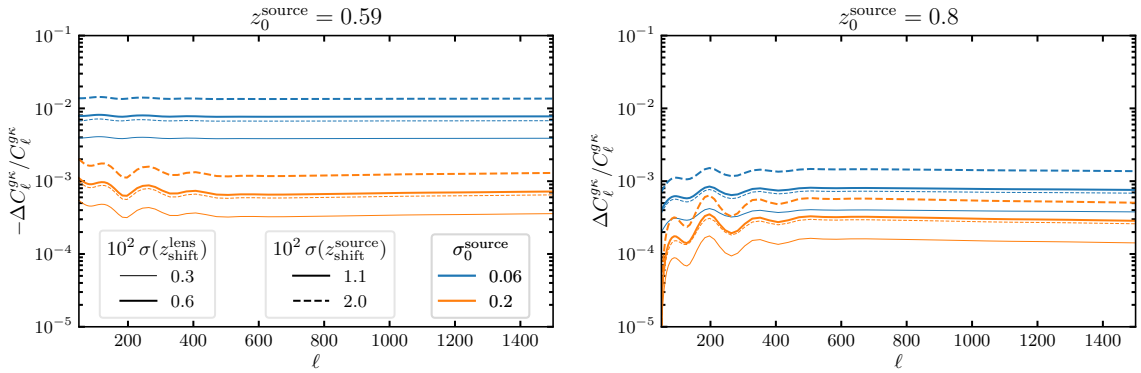

Figure 11: Fractional mode-coupling bias to the galaxy-galaxy lensing angular spectrum. Without loss of generality, we fix the lens galaxy distribution to have width $\sigma_0^{\rm lens} = 0.06$ and be centered at $z_0^{\rm lens} = 0.59$, and consider two values for each of $\sigma(z_{\rm shift}^{\rm lens})$, $\sigma(z_{\rm shift}^{\rm source})$ and $\sigma_0^{\rm source}$, denoting them via line width, line style and color, respectively. In the figure on the right, the sources are far behind the lenses and the biases are all positive; in the one on the left, the two distributions overlap, and the biases are all negative and amplified.

4.4 Multi-modal distributions

Finally, as a different application of our formalism, let us address the topic of multi-modal redshift distributions. Interlopers, or catastrophic outliers in photometric redshifts, are salient examples. The presence of such a poorly-characterized component in the distribution may signal challenges in the photometric selection, which in turn suggests that the fraction of galaxies in each of the modes might be uncertain and varying across the footprint.

To be more quantitative, let us suppose, following e.g., [35], that the redshift distribution is given by the sum of two components,

$$\frac{dn_g}{dz} = (1-f)\left(\frac{dn_g}{dz}\right)_1 + f\left(\frac{dn_g}{dz}\right)_2 , \tag{4.7}$$

where f is the fraction of sources that are misidentified. The top plot in the left panel of figure 12 shows a plausible such example where the bulk of the mean distribution is at $z = 0.59$ and has width $\sigma = 0.06$, similar to the DES lens galaxy samples, while $f = 0.1$ of the galaxies are interlopers at $z = 0.2$ with the same distribution width; we assume both components are Gaussian. We allow f to vary about its mean value across the sky following a red power law ($\alpha = -2$, $\ell'_{\rm max} = 100$) to produce a map-level standard deviation $\sigma(f) = 0.01$.

Once again, equation (2.50) will provide an excellent approximation to the mode-coupling contribution on angular scales smaller than those on which f varies significantly. In this limit, we can use the simplified kernel in the bottom plot of the left panel of figure 12 for insight. Since the bump at low-χ induces a different ℓ-to-k mapping than the signal kernel does, the bias kernel will no longer be as good an approximation to the signal kernel as it was in the discussion around equation (2.46), and we now expect the shape of the mode-coupling and signal spectra to differ. Indeed, in the right panel of the figure, we see that the bias has acquired a slight blue tilt. Nevertheless, the analytic approximation remains very accurate. As before, the additive contribution appears to be negligible.

These insights carry over to anisotropy in multi-modal redshift distributions of source galaxy samples, and thus have implications for cosmic shear and galaxy-galaxy lensing. How-

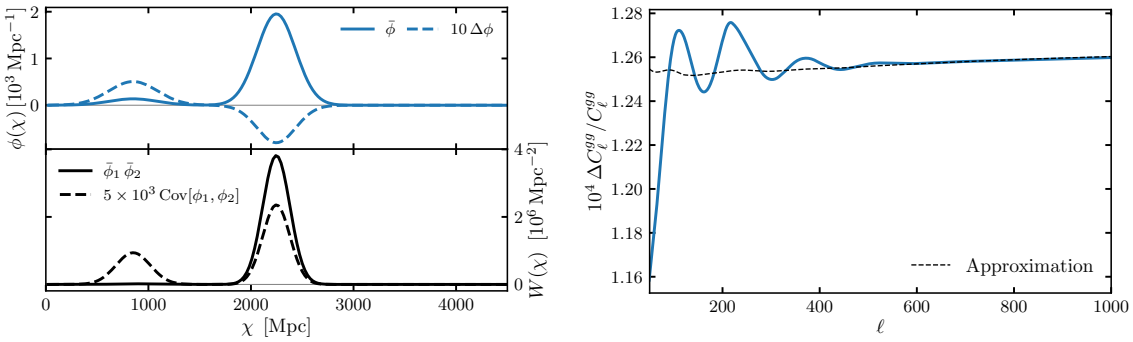

Figure 12: Impact of spatial variations of a multi-modal dN/dz on the galaxy clustering power spectrum. In the top plot of the left panel, we show the mean distribution (solid) and a typical perturbation to it (dashed); in the bottom plot, we show the signal kernel (solid) along with an approximation to the mode-coupling kernel (dashed). The right panel then shows the associated mode-coupling bias, as calculated using the full expression (solid) or in the approximation of equation (2.50) (dashed). Because of the interloper component, the ℓ-to-k mapping is slightly different in the mode-coupling and signal integrals, so $\Delta C_\ell^{gg}/C_\ell^{gg}$ is not exactly flat.

ever, we defer a rigorous treatment of all possible cases to the practitioners, noting that much of the infrastructure required for the task is present in our publicly-available code.

5 Conclusions

The angular auto- and cross-power spectra or correlation functions of fields that are projected onto the sky form one of the key observables from which we extract cosmological information. As long as the projection kernel, ϕ, is known, the assumption of statistical isotropy allows us to infer the 3D clustering from the observed, projected clustering. However, variations in photometry and observational non-idealities across a survey inevitably mean that the projection kernel is anisotropic. We have shown (§2) that such anisotropy leads, in general, to two additional contributions to the observed clustering: an additive term from the auto-correlation of $\Delta\phi$ and a mode-coupling term that arises from the interaction of $\Delta\phi$ with $\delta^{(3D)}$ and couples power at an observed angular wavenumber, ℓ, to that of nearby scales.

The signal and the two bias terms arise from auto-correlations of the three contributions to the projected density. This implies that in the special case of a cross-correlation with one field where ϕ is 'perfectly' known (e.g. CMB lensing) both bias contributions vanish. Since the sky-average of the perturbation to ϕ vanishes by definition, if at least one of the fields has mean zero – as is the case with cosmic shear or galaxy-galaxy lensing – then only the mode-coupling term survives. In the general case both can contribute, though the mode-coupling term frequently dominates.

In the limit that the variation, $\Delta\phi$, is primarily on scales much larger than those being probed by the clustering, the mode-coupling contribution to the angular spectrum is a projection of $\mathrm{Cov}[\Delta\phi, \Delta\phi]$ with a known kernel (equation 2.50). This represents our key result, and allows a simple and accurate estimation of the impact of varying projection across the sky in terms of the observable variance in e.g. mean redshift. Figure 3 illustrates the key

ingredients for translating the spatial variation in the kernel into the power bias, and can be used to understand multiple special cases described in the text.

We present the implications of our general formalism to several special cases – including galaxy auto- and cross-correlations, cosmic shear, galaxy-galaxy lensing and CMB lensing – in §3. Numerical examples for the impact of shifts in the mean redshift across the sky that follow a power-law power spectrum are presented in §4, along with a brief exploration of spatially-varying, multi-modal redshift distributions.

We show that for galaxy clustering the mode-coupling bias has a similar shape to the cosmological signal. For variations in the mean redshift with amplitude ϵ the bias to the power spectrum is positive and scales as ϵ^2/σ_0^4 with σ_0 the width of the sky-average dN/dz. Since on small scales the clustering signal scales as σ_0^{-2} the ratio of bias to signal scales as $(\epsilon/\sigma_0)^2$ (equation 4.6). For current-generation surveys, such as DES, this bias on the power spectrum is at worst percent level[12]. In contrast, cross-correlations of galaxy samples can be prone to a negative bias when projection anisotropy is correlated across both tracers; this can also be a percent-level effect for current surveys.

A similar story holds for cosmic shear (figures 9 and 10). Except for redshift distributions with low mean redshift, the quoted uncertainties in the mean redshift from DES would lead to sub-percent biases in the shear auto-spectra that are almost the same shape as the signal. (For the lowest-redshift bin of DES, the bias could be around percent-level.) Such biases would be subdominant to the statistical errors.

The same conclusion holds for galaxy-galaxy lensing (figure 11), though this case is interesting because the amplitude and sign of the bias depends on the distance between lens and source galaxy distributions – the bias is amplified when the two are close together, in which case it takes a negative sign. Nevertheless, it is a negligible effect for current surveys.

In closing we have presented a general formalism that allows one to assess the bias introduced on angular clustering measurements of 2D fields by anisotropic projection kernels. We make available a code, CARDiAC[13], that allows the user to do the same for any specific application. While we have illustrated the formalism with multiple examples, we leave a detailed exploration of specific scenarios to the groups analyzing particular observations. We have further assumed that the dN/dz, while anisotropic, is perfectly known, in which case our formalism explains how this information can be incorporated into analyses. We defer consideration of anisotropic and uncertain projection kernels to future work.

Acknowledgments

We are grateful to William Coulton, Emmanuel Schaan, Colin Hill, Giulio Fabbian, Simone Ferraro, Minas Karamanis, Xiao Fang and especially Noah Weaverdyck for useful conversations. M.W. is supported by the DOE. This research has made use of NASA's Astrophysics Data System, the arXiv preprint server, the Python programming language and packages NumPy, Matplotlib, SciPy, AstroPy, HealPy [36], Hankl [37] and FFTLog-and-Beyond [38]. This research is supported by the Director, Office of Science, Office of High Energy Physics of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231,

[12]In our comparisons to current data, we focus on anisotropy in the mean redshift of some dN/dz, and assume the extent of the variations is comparable to the uncertainty on the mean redshift quoted by DES. This might be overly optimistic, in which case our results should be regarded as a lower bound on the actual biases.

[13]Code for Anisotropic Redshift Distributions in Angular Clustering: https://github.com/abaleato/CARDiAC.

and by the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility under the same contract. This work was carried out on the territory of xučyun (Huchiun), the ancestral and unceded land of the Chochenyo speaking Ohlone people, the successors of the sovereign Verona Band of Alameda County.

A Multiplicative bias from a mischaracterized monopole

It is possible to conceive of situations where the footprint-average of the selection function's anisotropy does not vanish, i.e., $\Delta\phi_{00} \neq 0$. This will be the case, for example, when the fiducial dN/dz is calibrated on a patch that does not exactly match that on which the analysis is performed[14], or when photometric redshifts are systematically offset from their true value. In such a scenario, the general cross-correlation in Eq. (2.23) receives an additional multiplicative contribution so that

$$T_\ell^{ab} = \int d\chi_1\, d\chi_2 \,\left[U_\ell^{ab}(\chi_1,\chi_2) + A_\ell^{ab}(\chi_1,\chi_2) + R_\ell^{ab}(\chi_1,\chi_2) + Q_\ell^{ab}(\chi_1,\chi_2)\right] . \tag{A.1}$$

The new term, $Q_\ell^{ab}(\chi_1,\chi_2)$, comes from contractions of $\{\Delta\phi\delta\}$ and $\Delta\phi$ across two projected δ^{2D}s. Let us now derive it.

Statistical isotropy of the δs implies

$$Q_\ell^{ab}(\chi_1,\chi_2) = \frac{C_\ell^{ab}(\chi_1,\chi_2)}{2\ell+1} \sum_{\substack{m\\ \ell_1 m_1}} (-1)^m G^{\ell\ell_1\ell}_{-mm_1m} \left[\bar\phi(\chi_1)\Delta\phi_{\ell_1 m_1}(\chi_2) + \bar\phi(\chi_2)\Delta\phi_{\ell_1 m_1}(\chi_1)\right] . \tag{A.2}$$

We now make use of the selection rules for the Wigner-3j symbols comprising the Gaunt integral, specifically $m_1 + m_2 + m_3 = 0$, to obtain

$$G^{\ell\ell_1\ell}_{-mm_1m} = (2\ell+1)\sqrt{\frac{(2\ell_1+1)}{4\pi}} \begin{pmatrix} \ell & \ell & \ell_1 \\ m & -m & 0 \end{pmatrix} \begin{pmatrix} \ell & \ell & \ell_1 \\ 0 & 0 & 0 \end{pmatrix} \delta_{m_1 0} . \tag{A.3}$$

This, together with the identities

$$\sum_m (-1)^m \begin{pmatrix} \ell & \ell & \ell_1 \\ m & -m & 0 \end{pmatrix} = (-1)^\ell \sqrt{2\ell+1}\; \delta_{\ell_1 0} , \tag{A.4}$$

and

$$\begin{pmatrix} \ell & \ell & 0 \\ 0 & 0 & 0 \end{pmatrix} = \frac{(-1)^\ell}{\sqrt{2\ell+1}} , \tag{A.5}$$

allows us to simplify extensively, giving

$$Q_\ell^{ab}(\chi_1,\chi_2) = C_\ell^{ab}(\chi_1,\chi_2) \left[\bar\phi^a(\chi_1)\frac{\Delta\phi^b_{00}(\chi_2)}{\sqrt{4\pi}} + \bar\phi^b(\chi_2)\frac{\Delta\phi^a_{00}(\chi_1)}{\sqrt{4\pi}}\right] . \tag{A.6}$$

Naturally, only the monopole of $\Delta\phi$ can survive after coupling to an isotropic δ.

[14] One potential example of this is a cross-correlation analysis where the tracers are defined on somewhat different footprints. To avoid the bias described in this section, one could use different fiducial selection functions when projecting to the theoretical angular auto- and cross-spectra, ensuring that the fiducials are always representative of the patch where each measurement is made.

B Spherical Fourier-Bessel decomposition

Let us make contact with the spherical Fourier-Bessel decomposition (sFB) [20, 21, 39–47]. In this language, the arbitrary 3D field defined in equation (2.15) as

$$g_{\ell m}(\chi) \equiv \int d\hat{\boldsymbol{n}}\, Y^*_{\ell m}(\hat{\boldsymbol{n}}) g(\chi, \hat{\boldsymbol{n}})\,, \tag{B.1}$$

has conjugate

$$g_{\ell m}(k) = \sqrt{\frac{2}{\pi}} k \int d\chi\, \chi^2 j_\ell(k\chi) g_{\ell m}(\chi)\ . \tag{B.2}$$

It will be particularly interesting to work with the sFB decomposition of the perturbation field, $\Delta\phi_{\ell m}(k)$. In §2.2, we defined its C_ℓ's as

$$C_\ell^{\Delta\phi^a \Delta\phi^b}(\chi_1, \chi_2) \equiv \frac{1}{2\ell+1} \sum_m \Delta\phi^a_{\ell m}(\chi_1) \Delta\phi^{b,*}_{\ell m}(\chi_2)\,. \tag{B.3}$$

These can be related to the sFB cross-spectrum

$$C_\ell^{\Delta\phi^a \Delta\phi^b}(k_1, k_2) \equiv \frac{1}{2\ell+1} \sum_m \Delta\phi^a_{\ell m}(k_1) \Delta\phi^{b,*}_{\ell m}(k_2)\,, \tag{B.4}$$

via

$$C_\ell^{\Delta\phi^a \Delta\phi^b}(\chi_1, \chi_2) = \frac{2}{\pi} \int dk_1\, dk_2\, k_1 k_2 j_\ell(k_1\chi_1) j_\ell(k_2\chi_2) C_\ell^{\Delta\phi^a \Delta\phi^b}(k_1, k_2)\,. \tag{B.5}$$

These expressions can give us valuable insight into the mode-coupling kernels underlying the effects discussed in the main text. Let us rewrite the perturbation as

$$\Delta\phi_{\ell m}(k) = \sqrt{\frac{2}{\pi}} k \int d\chi\, \chi^2 j_\ell(k\chi) \int d\hat{\boldsymbol{n}}\, Y^*_{\ell m}(\hat{\boldsymbol{n}}) \Delta\phi(\chi, \hat{\boldsymbol{n}}) \tag{B.6}$$

$$= \int d\chi\, \chi J_{\ell+1/2}(k\chi) \int d\hat{\boldsymbol{n}}\, Y^*_{\ell m}(\hat{\boldsymbol{n}}) \Delta\phi(\chi, \hat{\boldsymbol{n}})\,. \tag{B.7}$$

where J_μ the μ-th order Bessel function of the first kind. This way of writing things makes explicit our approach to evaluating these quantities, which harnesses the computational speed of the FFTLog algorithm for Hankel transforms [48]. The procedure is as follows:

1. Define some HEALPix pixelization of the sky and generate templates of $z_{\rm shift}(\hat{\boldsymbol{n}}_i)$ and/or $\sigma_{\rm shift}(\hat{\boldsymbol{n}}_i)$ following §4.

2. Define a grid of values of χ_i that are distributed uniformly in logarithmic space, as required by FFTLog. At each χ_i, compute $\Delta\phi(\chi_i, \hat{\boldsymbol{n}}_i)$ from the $z_{\rm shift}(\hat{\boldsymbol{n}}_i)$ and $\sigma_{\rm shift}(\hat{\boldsymbol{n}}_i)$ we drew in the previous step, and take the spherical harmomic transform of $\Delta\phi(\chi_i, \hat{\boldsymbol{n}}_i)$ to obtain $\Delta\phi_{\ell m}(\chi_i)$.

3. Use the FFTLog algorithm to compute the $(\ell+1/2)$-th order Hankel tranform of every m mode of $\Delta\phi_{\ell m}(\chi_i)$, obtaining $\Delta\phi_{\ell m}(k_i)$.

For illustration, we show in figure 13 the sFB auto- and cross-spectrum of an example we have seen previously: the anisotropy produced by the mean-redshift shifts in figure 4.

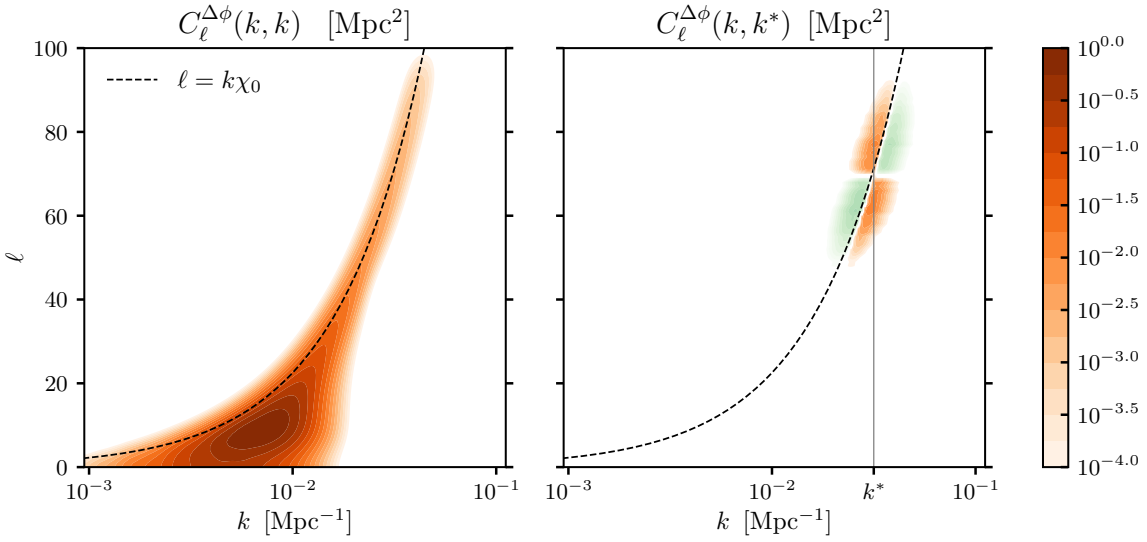

Figure 13: Spherical Fourier-Bessel auto- (left) and cross-spectrum (right, plotted in green where negative). The specific case examined here is produced by the variations in mean redshift shown in the left panel of figure 4, but some key features are general. The left panel shows that the anisotropy power peaks on large angular scales and is sourced primarily by modes transverse to the line-of-sight – for which $\ell \approx k\chi_0$, where χ_0 is the comoving distance to the center of the distribution – with this projection being sharper at high ℓ. This is a common feature of projections and can be related to the structure of the Bessel functions involved, as can be the fact that no anisotropy projects to $\ell \gtrsim k\chi_0$, while some does enter larger scales $\ell < k\chi_0$ via the oscillatory tails of the functions. The right panel shows that off-diagonal spectra are generally weaker, oscillatory, and constrained to a limited region of k-space near k^*.

C Evaluating the integrals

In this section we describe our approach to evaluating the various contributions to equation (2.23), the total angular cross-spectrum of two arbitrary fields in light of variations in their redshift distributions.

Consider first the unbiased contribution. Using equation (2.17) to relate $C_\ell^{ab}(\chi_1,\chi_2)$ to the 3D power spectrum of the tracers, we can write

$$T_\ell^{ab} \supset \int d\chi_1\, d\chi_2\; U_\ell^{ab}(\chi_1,\chi_2) = \int d\chi_1\, d\chi_2\, \bar{\phi}^a(\chi_1)\bar{\phi}^b(\chi_2) \times \frac{2}{\pi}\int dk\, k^2 j_\ell(k\chi_1) j_\ell(k\chi_2) P_{ab}\,(k;z_1,z_2)\;. \tag{C.1}$$

In the limit that P_{ab} varies slowly along the radial direction compared to the oscillations of the Bessel functions, it is a good approximation to assume that

$$\int dk\, k^2 j_\ell(k\chi_1) j_\ell(k\chi_2) P_{ab}\,(k;z_1,z_2) \approx \delta_{\rm D}(\chi_1-\chi_2)\;\frac{\pi}{2\chi_1^2} P_{ab}\left(\frac{\ell+1/2}{\chi_1};z_1\right)\;. \tag{C.2}$$

This is the leading-order version of the Limber approximation [49]. For typical tracers, it works well away from the largest angular scales. In this limit,

$$T_\ell^{ab} \supset \int d\chi_1\, d\chi_2\; U_\ell^{ab}(\chi_1,\chi_2) \approx \int d\chi\; \left[\frac{\bar{\phi}(\chi)}{\chi}\right]^2 P_{ab}\left(\frac{\ell+1/2}{\chi}; z\right) . \tag{C.3}$$

Let us now move on to the corrections. We evaluate the additive bias directly as

$$T_\ell^{ab} \supset \int d\chi_1\, d\chi_2\; A_\ell^{ab}(\chi_1,\chi_2) = \int d\chi_1\, d\chi_2\; C_\ell^{\Delta\phi^a\Delta\phi^b}(\chi_1,\chi_2)\,, \tag{C.4}$$

without resorting to Limber. On the other hand, the multiplicative bias

$$\begin{aligned} T_\ell^{ab} \supset \int d\chi_1\, d\chi_2\; R_\ell^{ab}(\chi_1,\chi_2) &= \int d\chi_1\, d\chi_2 \sum_L M_{\ell L}^{\Delta\phi^a\Delta\phi^b}(\chi_1,\chi_2) C_L^{ab}(\chi_1,\chi_2) \\ &= \frac{2}{\pi}\int d\chi_1\, d\chi_2 \sum_L M_{\ell L}^{\Delta\phi^a\Delta\phi^b}(\chi_1,\chi_2) \\ &\quad \times \int dk\, k^2 j_L(k\chi_1) j_L(k\chi_2) P_{ab}\left(k; z_1, z_2\right) , \end{aligned} \tag{C.5}$$

can be evaluated efficiently by applying Limber to the integral over k, independent of the smoothness of the $\Delta\phi$'s, giving

$$T_\ell^{ab} \supset \int d\chi_1\, d\chi_2\; R_\ell^{ab}(\chi_1,\chi_2) \approx \int d\chi\, \frac{1}{\chi^2} \sum_L M_{\ell L}^{\Delta\phi^a\Delta\phi^b}(\chi,\chi) P_{ab}\left(\frac{L+1/2}{\chi}; z\right) . \tag{C.6}$$

References

[1] P. J. E. Peebles, *The large-scale structure of the universe*. 1980.

[2] P. J. E. Peebles, *Principles of Physical Cosmology*. 1993.

[3] J. A. Peacock, *Cosmological Physics*. Jan., 1999.

[4] S. Dodelson, *Modern cosmology*. 2003.

[5] S. Dodelson and F. Schmidt, *Modern Cosmology*. Elsevier Science, 2020.

[6] D. Baumann, *Cosmology*. Cambridge University Press, 2022, 10.1017/9781108937092.

[7] LSST Dark Energy Science Collaboration, *Large Synoptic Survey Telescope: Dark Energy Science Collaboration*, *ArXiv e-prints* (2012) [1211.0310].

[8] LSST Science Collaboration, P. A. Abell, J. Allison, S. F. Anderson, J. R. Andrew, J. R. P. Angel et al., *LSST Science Book, Version 2.0*, *ArXiv e-prints* (2009) [0912.0201].

[9] Ž. Ivezić et al., *LSST: From Science Drivers to Reference Design and Anticipated Data Products*, *ApJ* **873** (2019) 111 [0805.2366].

[10] L. Amendola et al., *Cosmology and Fundamental Physics with the Euclid Satellite*, *Living Reviews in Relativity* **16** (2013) 6 [1206.1225].

[11] L. Amendola, S. Appleby, A. Avgoustidis, D. Bacon, T. Baker, M. Baldi et al., *Cosmology and fundamental physics with the Euclid satellite*, *Living Reviews in Relativity* **21** (2018) 2 [1606.00180].

[12] D. Huterer, C. E. Cunha and W. Fang, *Calibration errors unleashed: effects on cosmological parameters and requirements for large-scale structure surveys*, *Mon. Not. R. Astron. Soc.* **432** (2013) 2945 [1211.1015].

[13] D. L. Shafer and D. Huterer, *Multiplicative errors in the galaxy power spectrum: self-calibration of unknown photometric systematics for precision cosmology*, *Mon. Not. R. Astron. Soc.* **447** (2015) 2961 [1410.0035].

[14] N. Weaverdyck and D. Huterer, *Mitigating contamination in LSS surveys: a comparison of methods*, *Mon. Not. R. Astron. Soc.* **503** (2021) 5061 [2007.14499].

[15] M. Salvato, O. Ilbert and B. Hoyle, *The many flavours of photometric redshifts*, *Nature Astronomy* **3** (2019) 212 [1805.12574].

[16] J. A. Newman and D. Gruen, *Photometric Redshifts for Next-Generation Surveys*, *ARA&A* **60** (2022) 363 [2206.13633].

[17] M. Ouchi, Y. Ono and T. Shibuya, *Observations of the Lyman-α Universe*, *ARA&A* **58** (2020) 617 [2012.07960].

[18] K. K. Datta, T. R. Choudhury and S. Bharadwaj, *The multifrequency angular power spectrum of the epoch of reionization 21-cm signal*, *Mon. Not. R. Astron. Soc.* **378** (2007) 119 [astro-ph/0605546].

[19] J. R. Shaw, K. Sigurdson, U.-L. Pen, A. Stebbins and M. Sitwell, *All-sky Interferometry with Spherical Harmonic Transit Telescopes*, *ApJ* **781** (2014) 57 [1302.0327].

[20] E. Castorina and M. White, *Beyond the plane-parallel approximation for redshift surveys*, *Mon. Not. R. Astron. Soc.* **476** (2018) 4403 [1709.09730].

[21] E. Castorina and M. White, *The Zeldovich approximation and wide-angle redshift-space distortions*, *Mon. Not. R. Astron. Soc.* **479** (2018) 741 [1803.08185].

[22] E. Hivon, K. M. Górski, C. B. Netterfield, B. P. Crill, S. Prunet and F. Hansen, *MASTER of the Cosmic Microwave Background Anisotropy Power Spectrum: A Fast Method for Statistical Analysis of Large and Complex Cosmic Microwave Background Data Sets*, *ApJ* **567** (2002) 2 [astro-ph/0105302].

[23] DES Collaboration, *Dark Energy Survey Year 3 results: Cosmological constraints from galaxy clustering and weak lensing*, *Phys. Rev. D* **105** (2022) 023520 [2105.13549].

[24] Planck Collaboration, *Planck 2018 results. I. Overview and the cosmological legacy of Planck*, *arXiv e-prints* (2018) arXiv:1807.06205 [1807.06205].

[25] Planck Collaboration, *Planck 2018 results. VI. Cosmological parameters*, *A&A* **641** (2020) A6 [1807.06209].

[26] D. A. Varshalovich, A. N. Moskalev and V. K. Khersonskii, *Quantum Theory of Angular Momentum*. 1988, 10.1142/0270.

[27] N. Kaiser and G. Squires, *Mapping the Dark Matter with Weak Gravitational Lensing*, *ApJ* **404** (1993) 441.

[28] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke et al., *HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere*, *ApJ* **622** (2005) 759 [arXiv:astro-ph/0409513].

[29] N. Kokron, J. DeRose, S.-F. Chen, M. White and R. H. Wechsler, *The cosmology dependence of galaxy clustering and lensing from a hybrid N-body-perturbation theory model*, *Mon. Not. R. Astron. Soc.* **505** (2021) 1422 [2101.11014].

[30] J. DeRose et al., *The Buzzard Flock: Dark Energy Survey Synthetic Sky Catalogs*, *arXiv e-prints* (2019) arXiv:1901.02401 [1901.02401].

[31] J. Elvin-Poole and DES Collaboration, *Dark Energy Survey year 1 results: Galaxy clustering for combined probes*, *Phys. Rev. D* **98** (2018) 042006 [1708.01536].

[32] A. J. Mead, S. Brieden, T. Tröster and C. Heymans, *HMCODE-2020: improved modelling of non-linear cosmological power spectra with baryonic feedback*, *Mon. Not. R. Astron. Soc.* **502** (2021) 1401 [2009.01858].

[33] R. E. Smith et al., *Stable clustering, the halo model and non-linear cosmological power spectra*, *Mon. Not. R. Astron. Soc.* **341** (2003) 1311 [astro-ph/0207664].

[34] A. Lewis, A. Challinor and A. Lasenby, *Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models*, *ApJ* **538** (2000) 473.

[35] C. Modi, M. White and Z. Vlah, *Modeling CMB lensing cross correlations with CLEFT*, *J. Cosmol. Astropart. Phys.* **2017** (2017) 009 [1706.03173].

[36] A. Zonca, L. Singer, D. Lenz, M. Reinecke, C. Rosset, E. Hivon et al., *healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in python*, *Journal of Open Source Software* **4** (2019) 1298.

[37] M. Karamanis and F. Beutler, *hankl: A lightweight Python implementation of the FFTLog algorithm for Cosmology*, *arXiv e-prints* (2021) arXiv:2106.06331 [2106.06331].

[38] 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.* **2020** (2020) 010 [1911.11947].

[39] O. Lahav, K. B. Fisher, Y. Hoffman, C. A. Scharf and S. Zaroubi, *Wiener Reconstruction of All-Sky Galaxy Surveys in Spherical Harmonics*, *ApJL* **423** (1994) L93 [astro-ph/9311059].

[40] K. B. Fisher, C. A. Scharf and O. Lahav, *A spherical harmonic approach to redshift distortion and a measurement of Omega(0) from the 1.2-Jy IRAS Redshift Survey*, *Mon. Not. R. Astron. Soc.* **266** (1994) 219 [astro-ph/9309027].

[41] A. F. Heavens and A. N. Taylor, *A spherical harmonic analysis of redshift space*, *Mon. Not. R. Astron. Soc.* **275** (1995) 483 [astro-ph/9409027].

[42] W. J. Percival, D. Burkey, A. Heavens, A. Taylor, S. Cole, J. A. Peacock et al., *The 2dF Galaxy Redshift Survey: spherical harmonics analysis of fluctuations in the final catalogue*, *Mon. Not. R. Astron. Soc.* **353** (2004) 1201 [astro-ph/0406513].

[43] N. Padmanabhan, M. Tegmark and A. J. S. Hamilton, *The Power Spectrum of the CFA/SSRS UZC Galaxy Redshift Survey*, *ApJ* **550** (2001) 52 [astro-ph/9911421].

[44] G. Pratten and D. Munshi, *Effects of linear redshift space distortions and perturbation theory on BAOs: a 3D spherical analysis*, *Mon. Not. R. Astron. Soc.* **436** (2013) 3792 [1301.3673].

[45] L. Samushia, *Proper Fourier decomposition formalism for cosmological fields in spherical shells*, *arXiv e-prints* (2019) arXiv:1906.05866 [1906.05866].

[46] H. S. Grasshorn Gebhardt and O. Doré, *Fabulous code for spherical Fourier-Bessel decomposition*, *Phys. Rev. D* **104** (2021) 123548 [2102.10079].

[47] S. Passaglia, A. Manzotti and S. Dodelson, *Cross-correlating 2D and 3D galaxy surveys*, *Phys. Rev. D* **95** (2017) 123508 [1702.03004].

[48] A. J. S. Hamilton, *Uncorrelated modes of the non-linear power spectrum*, *Mon. Not. R. Astron. Soc.* **312** (2000) 257 [astro-ph/9905191].

[49] M. LoVerde and N. Afshordi, *Extended Limber approximation*, *Phys. Rev. D* **78** (2008) 123506 [0809.5112].