

PAPER • OPEN ACCESS

Catalog-based pseudo- C_ℓ s

To cite this article: Kevin Wolz *et al* JCAP01(2025)028

View the [article online](#) for updates and enhancements.

You may also like

- [CIBER Fourth Flight Fluctuation Analysis: Pseudopower Spectrum Formalism, Improved Source Masking, and Validation on Mocks](#)
Richard M. Feder, James J. Bock, Yun-Ting Cheng et al.
- [The Atacama Cosmology Telescope: Cosmology from Cross-correlations of unWISE Galaxies and ACT DR6 CMB Lensing](#)
Gerrit S. Farren, Alex Krolewski, Niall MacCrann et al.
- [Physical Properties of Large-scale Galactic Filaments](#)
Catherine Zucker, Cara Battersby and Alyssa Goodman

Catalog-based pseudo- C_ℓ s

Kevin Wolz ^a, David Alonso ^a and Andrina Nicola ^b

^aDepartment of Physics, University of Oxford,
Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, U.K.

^bArgelander Institut für Astronomie, Universität Bonn,
Auf dem Hügel 71, 53121 Bonn, Germany

E-mail: kevin.wolz@physics.ox.ac.uk, david.alonso@physics.ox.ac.uk,
anicola@uni-bonn.de

ABSTRACT: We present a formalism to extract the angular power spectrum of fields sampled at a finite number of points with arbitrary positions — a common situation for several catalog-based astrophysical probes — through a simple extension of the standard pseudo- C_ℓ algorithm. A key complication in this case is the need to handle the shot noise component of the associated discrete angular mask which, for sparse catalogs, can lead to strong coupling between very different angular scales. We show that this problem can be solved easily by estimating this contribution analytically and subtracting it. The resulting estimator is immune to small-scale pixelization effects and aliasing, and, most notably, unbiased against the contribution from measurement noise uncorrelated between different sources. We demonstrate the validity of the method in the context of cosmic shear datasets, and showcase its usage in the case of other spin-0 and spin-1 astrophysical fields of interest. We incorporate the method in the public [NaMaster](#) code.

KEYWORDS: power spectrum, galaxy clusters, gravitational lensing

ARXIV EPRINT: [2407.21013](https://arxiv.org/abs/2407.21013)

Contents

1	Introduction	1
2	Methods	3
2.1	Pseudo- C_ℓ s	3
2.2	The case of discretely-sampled fields	4
2.3	Spin- s fields	10
2.4	Clustering	12
3	Validation	14
3.1	Validation setup	14
3.2	Validation results	15
4	Example applications	20
4.1	Spin-2 example: cosmic shear from DES-Y3	20
4.2	Spin-1 example: proper motion surveys	22
4.3	Spin-0 example: FRB-based dispersion measure maps	23
4.4	Clustering example	26
5	Conclusion	28
A	Generalities of spherical harmonic functions	29
B	Pixelization effects in catalog-based observables	30

1 Introduction

In cosmology and astrophysics, a plethora of observable quantities can be described by projected fields on the sky. These fields may be observed and mapped in a continuous manner, as is the case of Cosmic Microwave Background (CMB) observations, or may be observed only at the discrete positions of a finite number of sources. Depending on the physical nature of the quantity at hand, these quantities will transform differently under rotations of our angular coordinates, leading to spin-0 fields (i.e., scalars, such as the CMB temperature anisotropies, the lensing convergence, or the galaxy overdensity), spin-1 fields (i.e., vectors such as transverse peculiar motions or lensing displacements), spin-2 fields (i.e., tensors such as the CMB polarization anisotropies or weak lensing shear), and higher-spin fields (e.g. the spin-4 polarization of gravitational waves). The angular power spectrum, corresponding to the covariance between angular fluctuations in two different fields at a given multipole scale ℓ , is one of the most used summary statistics in cosmology and astrophysics and, in particular, encodes all available information in the case of Gaussian fields. For this reason, being able to accurately estimate the angular power spectrum is a key requirement for many cosmological probes. In the case of galaxy weak lensing, for instance, a precise

and accurate reconstruction of the cosmic shear power spectrum on sub-arcminute scales is of vital importance to obtain meaningful constraints on the gravity-dominated, nonlinear regime of cosmic structure formation, scrutinizing the Universe’s matter content and nature of Dark Energy [1, 2]. This is a key goal for current and next-generation galaxy imaging surveys, such as the Kilo Degree Survey (KiDS),¹ the Dark Energy Survey (DES),² the Hyper Suprime-Cam,³ Euclid,⁴ and the Rubin Observatory Legacy Survey of Space and Time (LSST).⁵ Valuable complementary insight may be gained from the angular power spectrum of other discretely sampled tracers of the matter density, such as the overdensity of galaxies [3–5], fluctuations in their sizes [6, 7], fluctuations in the luminosity distance of standard sirens [8], transverse peculiar motions [9], or the dispersion measure of extragalactic Fast Radio Bursts (FRBs) [10–13].

There is a large literature of unbiased power spectrum estimators that are optimal, have well-understood statistical properties, and robustly meet common observational challenges such as partial sky coverage, linear sky contamination, and inhomogeneous noise [14–20]. For small angular scales, optimal quadratic estimators are usually computationally prohibitive, and the most common alternative is their simpler incarnation, the so-called “pseudo- C_ℓ ” estimator [21–23]. The pseudo- C_ℓ estimator effectively assumes a diagonal covariance matrix for the data in pixel space, and is therefore close to optimal for spatially uncorrelated or white-noise-dominated data. As a consequence of this assumption, the mode-coupling induced by the spatial variation of the resulting inverse-variance weights (i.e., the “mask”) can be estimated using fast, analytical methods. Variations of the method allow it to account for the impact of linear contaminant deprojection, and for the effects of E - and B -mode purification in non-scalar fields [23, 24].

In its standard implementation, the pseudo- C_ℓ method takes, as input, pixelized maps of the fields in question. When the original observations are in the form of discrete measurements at arbitrary sky positions, as is the case for any catalog-based observable, this results in two fundamental problems. First, generating a pixelized map from the discrete sources induces a bias at scales comparable to, and smaller than, the pixel size. The form of this bias depends on a combination of pixel smoothing and aliasing of small-scale power, and is non-trivial to characterize in general [25]. Second, when the number density of point sources constituting the mask is low, the mask is dominated by Poisson noise, also known as shot noise, which, as we will see, leads to statistical couplings between distant angular multipoles. If not treated carefully, this can lead to strong numerical instabilities and catastrophic biases in the estimator. The issue of small-scale power in the mask has been traditionally handled via apodization and smoothing of the mask edges [21, 23, 26], but this approach is highly impractical in the context of catalog-based observations for sparse samples.

This work addresses the issue of how to define an unbiased power spectrum estimator for fields of arbitrary spin discretely sampled on the sphere. Building on the recent work on this topic by [5], we forgo the pixelization step completely and directly estimate spherical

¹<http://kids.strw.leidenuniv.nl>.

²<https://www.darkenergysurvey.org>.

³<https://www.naoj.org/Projects/HSC/>.

⁴<https://www.euclid-ec.org>.

⁵<https://www.lsst.org>.

harmonic transforms of the field and its measurement weights at the discrete source positions. This allows us to build a pixel-independent version of the pseudo- C_ℓ estimator, addressing the first problem described above. We address the second problem (shot noise instabilities) by introducing a simple modification to the estimator that accounts for uncorrelated noise, both in the mask and the field, and simultaneously eliminates both the associated noise bias and the numerical instabilities of the mask-induced mode coupling. More specifically, the contribution to mode coupling from the mask shot noise is estimated analytically and subtracted from the measured pseudo- C_ℓ , together with a naive, data-based estimate of the noise bias. As we will show, the resulting estimator is then immune to any bias from uncorrelated noise, and benefits from a numerically stable mode-coupling matrix, independently of the number density of sources in the catalog. We implement this methodology in the publicly available `NaMaster` package.⁶ As we were preparing this manuscript, we learned of related work presented in Tessore et al. 2024 [27]. The results there agree with those presented here.

This paper is structured as follows. In section 2, we lay out the formalism that defines our unbiased estimator. In section 3, we validate it on simulated cosmic shear catalogs with a realistic source distribution, based on real data from the Kilo Degree Survey, also demonstrating its validity on significantly lower number densities. In section 4, we present several example applications of the methodology, including state-of-the-art cosmic shear data from the Dark Energy Survey, transverse proper motion surveys, and FRB-based dispersion measure maps. We conclude in section 5.

2 Methods

This section presents the formalism to estimate the angular power spectrum from sparse observations of a given field at a discrete set of points on the sphere using the pseudo- C_ℓ formalism. Section 2.1 reviews the main concepts behind the standard pseudo- C_ℓ method, and the application to discrete tracers is described in section 2.2. We first present the formalism in the simpler case of scalar fields, and then, in section 2.3 present the generalization to spin- s quantities.

2.1 Pseudo- C_ℓ s

Let $f(\hat{\mathbf{n}})$ be a scalar field on the sphere, and let $w(\hat{\mathbf{n}})$ be a sky mask (often also called weights maps, or sky window function). The mask may be as simple as a binary map (0 for unobserved regions, 1 for observed regions) but, more usefully, it should represent the precision with which f has been measured at different points (e.g. an inverse-variance map will be optimal for noise-dominated observations). Because of this, $w(\hat{\mathbf{n}})$ is often referred to as a “window function” or “weights map”, but we will stick to the simpler “mask” in what follows.

The field is related to its harmonic coefficients $f_{\ell m}$ through a spherical harmonic transform (and its inverse):

$$f(\hat{\mathbf{n}}) = \sum_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}) f_{\ell m} \leftrightarrow f_{\ell m} = \int d\hat{\mathbf{n}} f(\hat{\mathbf{n}}) Y_{\ell m}^*(\hat{\mathbf{n}}), \quad (2.1)$$

⁶<https://github.com/LSSTDESC/NaMaster/>.

where $Y_{\ell m}(\hat{\mathbf{n}})$ are the spherical harmonic functions. The spherical harmonic coefficients of the mask, or any other sky map, are defined in the same way.

Let $f^w(\hat{\mathbf{n}}) \equiv f(\hat{\mathbf{n}})w(\hat{\mathbf{n}})$ be the masked field. The ‘‘pseudo- C_ℓ ’’ of f is simply the naive estimate of the masked field’s power spectrum ignoring the fact that it’s masked:

$$\tilde{C}_\ell^f \equiv \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} f_{\ell m}^w f_{\ell m}^{w*}. \quad (2.2)$$

We commonly assume the field to receive signal and noise contributions: $f = s + n$, where the signal s is a statistically isotropic field, such that

$$\langle s_{\ell m} s_{\ell' m'}^* \rangle = \delta_{\ell\ell'} \delta_{mm'} S_\ell, \quad (2.3)$$

here S_ℓ is the signal power spectrum, and is the quantity we are ultimately after.

The key result of the pseudo- C_ℓ estimator is the relation between the ensemble average of \tilde{C}_ℓ^f and S_ℓ :

$$\langle \tilde{C}_\ell^f \rangle = \sum_{\ell'} M_{\ell\ell'}^w S_{\ell'} + \tilde{N}_\ell, \quad (2.4)$$

where $M_{\ell\ell'}^w$ is the mode-coupling matrix:

$$M_{\ell\ell'}^w = \frac{2\ell' + 1}{4\pi} \sum_{\ell''} \tilde{C}_{\ell''}^{rw} (2\ell'' + 1) \begin{pmatrix} \ell & \ell' & \ell'' \\ 0 & 0 & 0 \end{pmatrix}^2, \quad (2.5)$$

here $\tilde{C}_\ell^{rw} \equiv \sum_m |w_{\ell m}|^2 / (2\ell + 1)$ is the pseudo- C_ℓ of the mask. \tilde{N}_ℓ is the ‘‘noise bias’’, given by the ensemble average of the pseudo- C_ℓ of the masked noise:

$$\tilde{N}_\ell \equiv \left\langle \frac{1}{2\ell + 1} \sum_m n_{\ell m}^w n_{\ell m}^{w*} \right\rangle. \quad (2.6)$$

Depending on the type of observations, and our understanding of the noise properties, \tilde{N}_ℓ may be estimated analytically, via Monte-Carlo simulations, or avoided by using cross-correlations between observations with the same signal but different noise realizations (e.g. CMB maps constructed from disjoint sets of timestreams). As we will see, in the case of discretely-sampled fields, the noise bias can be entirely avoided by construction under the assumption that the noise contributions from different points are uncorrelated. The expressions above may be generalized to the case of cross-correlations between different fields, by simply replacing \tilde{C}_ℓ^{rw} with the cross-spectrum between the masks of both fields.

2.2 The case of discretely-sampled fields

Consider now observations of the field f only at a discrete set of N points on the sphere. Most commonly, these will be the sky positions of sources in a catalog. Let w_i be the weight associated to the i -th point, with sky position $\hat{\mathbf{n}}_i$, and $f_i \equiv f(\hat{\mathbf{n}}_i)$. In the case where the observable field is the point density itself, such as galaxy clustering, one can define the mask using a denser sampling than what is afforded by data itself. We will discuss the clustering case in section 2.4. To ease notation, in what follows, the subscript i applied to any function

defined on the sphere will be shorthand for the value of that function at $\hat{\mathbf{n}}_i$. The mask and masked field can, in this case, be written as a sum over delta functions:

$$w(\hat{\mathbf{n}}) = \sum_i w_i \delta^D(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i), \quad f^w(\hat{\mathbf{n}}) \equiv w(\hat{\mathbf{n}}) f(\hat{\mathbf{n}}) = \sum_i w_i f_i \delta^D(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i), \quad (2.7)$$

where $\delta^D(\hat{\mathbf{n}}_1, \hat{\mathbf{n}}_2)$ is the Dirac delta function on the sphere.⁷ The harmonic coefficients of w and f^w are then sums over the harmonic functions sampled at the positions of the discrete sources

$$w_{\ell m} = \sum_i w_i Y_{\ell m, i}^*, \quad f_{\ell m}^w = \sum_i w_i f_i Y_{\ell m, i}^*. \quad (2.8)$$

An efficient evaluation of such discrete harmonic transform over a non-regular set of points is, in general, computationally more expensive than standard approaches making use of regular pixelization schemes [28]. Nevertheless, approaches for dealing with this problem have been recently developed [5, 29, 30], and we will use the implementation of such “general” spherical harmonic transforms built in the `ducc` library,⁸ as described in [29]. This paper will not discuss the details of the numerical implementation of these methods and instead will deal with the implementational details of the pseudo- C_ℓ algorithm itself, once these discrete transforms are available. We refer readers to [29] further details.

The pseudo- C_ℓ of both mask and masked field is given by:

$$\tilde{C}_\ell^w = \sum_{ij} w_i w_j \frac{1}{2\ell + 1} \sum_m Y_{\ell m, i}^* Y_{\ell m, j}, \quad (2.9)$$

$$\tilde{C}_\ell^f = \sum_{ij} w_i w_j f_i f_j \frac{1}{2\ell + 1} \sum_m Y_{\ell m, i}^* Y_{\ell m, j}. \quad (2.10)$$

These equations may be further simplified making use of the addition theorem for spherical harmonics (eq. (A.7)). However, it will be more useful to leave them as is for now.

2.2.1 Noise bias from uncorrelated measurements

Let us now consider our measurements of the field $f = s + n$ to contain some common signal s and noise component n . We assume that signal and noise are uncorrelated, and that noise is uncorrelated among different sources.⁹ Both assumptions are idealizations that may break down in the presence of catalog-specific systematic effects, such as source crowding and environmental effects. Under these assumptions, the noise contribution to \tilde{C}_ℓ^f reads

$$\langle n_i n_j \rangle = \delta_{ij} \sigma_{N, i}^2, \quad (2.11)$$

where $\sigma_{N, i}^2$ is the noise variance in the i -th point. In this case, the noise bias only receives contributions from the variance at the same source (i.e., the zero-lag correlation function):

$$\tilde{N}_\ell = \sum_{ij} w_i w_j \langle n_i n_j \rangle \frac{1}{2\ell + 1} \sum_m Y_{\ell m, i}^* Y_{\ell m, j} = \frac{1}{4\pi} \sum_i w_i^2 \sigma_{N, i}^2, \quad (2.12)$$

⁷That is, $\delta^D(\hat{\mathbf{n}}_1, \hat{\mathbf{n}}_2) \equiv \delta(\varphi_1 - \varphi_2) \delta(\cos \theta_1 - \cos \theta_2)$, where (θ_n, φ_n) are the spherical coordinates of the unit vector $\hat{\mathbf{n}}_n$.

⁸<https://mtr.pages.mpcdf.de/ducc>.

⁹In general, correlated noise between different sources leads to a scale-dependent noise power spectrum. Common methods to avoid this noise bias include cross-correlating fields and direct noise modeling.

where, in the last equality, we have made use of the addition theorem for spherical harmonics (eq. (A.7)) for the special case $i = j$. The noise bias is therefore “white” (i.e., scale-independent), and non-zero at all ℓ s.

Since the standard pseudo- C_ℓ estimator requires us to estimate \tilde{N}_ℓ and subtract it from \tilde{C}_ℓ^f , in order to obtain an unbiased estimator of the signal power spectrum S_ℓ , it is important for us to consider practical ways to estimate \tilde{N}_ℓ . Assuming our measurements of f to be dominated by noise at any single point, one potentially simple way to do this would be to estimate \tilde{N}_ℓ directly from the variance of the data itself. This methodology has been commonly applied, in the past, when estimating cosmic shear power spectra, where measurements are indeed dominated by noise [31–33]. In this situation, and looking at eq. (2.12), we can thus propose the following estimator for \tilde{N}_ℓ :

$$\tilde{N}^f \equiv \frac{1}{4\pi} \sum_i w_i^2 f_i^2. \quad (2.13)$$

Expanding the field measurements into signal and noise components, we can thus see that \tilde{N}^f is a biased estimate of \tilde{N}_ℓ due to the presence of signal in our measurements:

$$\langle \tilde{N}^f \rangle = \tilde{N}_\ell + \frac{1}{4\pi} \sum_i w_i^2 \langle s_i^2 \rangle. \quad (2.14)$$

For a statistically isotropic field, such as s , the field variance is related to its power spectrum via

$$\sigma_S^2 \equiv \langle s_i^2 \rangle = \sum_\ell \frac{2\ell + 1}{4\pi} S_\ell, \quad (2.15)$$

and thus we can write the bias of \tilde{N}^f as:

$$\langle \tilde{N}^f \rangle = \tilde{N}_\ell + \tilde{N}^w \sigma_S^2, \quad (2.16)$$

where we have defined the combination

$$\tilde{N}^w \equiv \frac{1}{4\pi} \sum_i w_i^2. \quad (2.17)$$

2.2.2 Mask noise

Drawing intuition from our discussion of uncorrelated noise in the field, we can see that the pseudo- C_ℓ of the mask \tilde{C}_ℓ^w will also receive a white-noise-like contribution from the terms $i = j$ in eq. (2.9). By inspection, we can see that this is none other than the quantity \tilde{N}^w we just calculated. We will thus refer to \tilde{N}^w as the “mask shot noise”. In terms of it, we can thus write the mask power spectrum as

$$\tilde{C}_\ell^w = \tilde{S}_\ell^w + \tilde{N}^w, \quad (2.18)$$

where the “signal” component is, mathematically:

$$\tilde{S}_\ell^w = \frac{1}{4\pi} \sum_{i \neq j} w_i w_j P_\ell(\mu_{ij}), \quad (2.19)$$

with $\mu_{ij} \equiv \hat{\mathbf{n}}_i \cdot \hat{\mathbf{n}}_j$, and $P_\ell(\mu)$ the Legendre polynomials. In practice, the fastest way to estimate \tilde{S}_ℓ^w is not to compute the double sum above, but to compute \tilde{C}_ℓ^w from the harmonic transform of $w(\hat{\mathbf{n}})$ (eq. (2.8)), and then subtract \tilde{N}^w from it, calculated using eq. (2.17).¹⁰

The presence of \tilde{N}^w is a feature exclusive of the discrete origin of our observations. It sets a floor below which \tilde{C}_ℓ^w cannot deviate significantly and, for most common point distributions of interest in astrophysics, \tilde{C}_ℓ^w will, in fact, tend to \tilde{N}^w at high ℓ . This is contrary to the usual scenario for power spectrum estimation over incomplete skies. In more common cases, in which the mask is a smooth function, and has support over a significant fraction of the sky (compared to the angular scales of interest), \tilde{C}_ℓ^w decays with ℓ . This causes the sum over ℓ'' in eq. (2.5) to converge quickly, and the resulting mode-coupling matrix to be highly concentrated along the diagonal. It is therefore important to understand what the contribution from \tilde{N}^w would be to $M_{\ell\ell'}^w$. Calling this contribution $M_{\ell\ell'}^{w,N}$, we can calculate it, replacing \tilde{C}_ℓ^w by \tilde{N}^w in eq. (2.5), and obtain

$$M_{\ell\ell'}^{w,N} = \frac{2\ell' + 1}{4\pi} \tilde{N}^w, \quad (2.20)$$

where we have made use of one of the orthogonality relations of the Wigner-3j coefficients (eq. (A.10)). We can thus see that the mask shot noise leads to a contribution that couples all multipole ℓ s with all other multipoles ℓ' (in principle up to $\ell' = \infty$). Although this does not necessarily imply that the mode-coupling matrix is singular, it does mean that an accurate implementation of the conventional pseudo- C_ℓ estimator might require (depending on the size of \tilde{N}^w ,) estimating $M_{\ell\ell'}^w$ up to very large values of both ℓ and ℓ' in order to account for mode coupling. This is problematic, since estimating $M_{\ell\ell'}^w$ is the most expensive step of the pseudo- C_ℓ estimator.

Instead, in this work we opt for a different approach that avoids this problem entirely, namely subtracting the mask-shot-noise contribution from the coupled mask pseudo- C_ℓ estimator before computing the mode coupling matrix. Employing our ansatz for the noise estimator \tilde{N}^f , we can combine eqs. (2.4) and (2.14) to compute the expectation value of the signal pseudo- C_ℓ estimator

$$\langle \tilde{C}_\ell^f - \tilde{N}^f \rangle = \sum_{\ell'} M_{\ell\ell'}^w S_{\ell'} + \tilde{N}_\ell - (\tilde{N}_\ell + \tilde{N}^w \sigma_S^2) = \sum_{\ell'} (M_{\ell\ell'}^w - M_{\ell\ell'}^{w,N}) S_{\ell'}, \quad (2.21)$$

where, in the last equality, we used that

$$\sum_{\ell'} M_{\ell\ell'}^{w,N} S_{\ell'} = \tilde{N}^w \sigma_S^2, \quad (2.22)$$

which follows directly from eq. (2.20). Thus, we find that the ‘‘spurious’’ signal contribution to our naive estimate of the noise bias precisely cancels the undesired mask-shot-noise contribution in the mode coupling matrix. This result allows us to construct a fully unbiased estimator for the power spectrum of a discretely-sampled field with uncorrelated noise, which we describe below.

¹⁰The idea of subtracting shot noise from the mask power spectrum was first mentioned, although not mathematically laid out, in [5].

2.2.3 An unbiased estimator

Gathering our results from the last two sections, the algorithm presented in this paper is as follows:

1. Estimate the spherical harmonic coefficients of the masked field, and of the mask (eq. (2.8)).
2. From these, compute the pseudo- C_ℓ of the mask (\tilde{C}_ℓ^w) and of the masked field (\tilde{C}_ℓ^f).
3. Estimate the “noise” components of both power spectra, \tilde{N}^w , \tilde{N}^f directly from the measurements, using eqs. (2.17) and (2.13), respectively.
4. Subtract both noise contributions to form the “signal” pseudo- C_ℓ s of the mask and the masked field¹¹ $\tilde{S}_\ell^w \equiv \tilde{C}_\ell^w - \tilde{N}^w$, $\tilde{S}_\ell^f \equiv \tilde{C}_\ell^f - \tilde{N}^f$.
5. Apply the standard pseudo- C_ℓ estimator to these “signal-only” spectra. This means, compute the mode-coupling matrix using \tilde{S}_ℓ^w instead of \tilde{C}_ℓ^w , and apply its inverse¹² to \tilde{S}_ℓ^f to obtain an unbiased estimate of the signal power spectrum.

To reiterate, subtracting the noise-like contribution from the mask pseudo- C_ℓ has two desirable consequences: first, it automatically allows us to account for the bias induced by correcting for a noise power spectrum computed from our own data. Secondly, the resulting mode-coupling matrix is now computed from the noise-debiased mask spectrum which, by construction, will decay with ℓ (at least for the usual point distributions that make up most cosmological observations). Hence, the sum over ℓ'' in (2.5) will have better convergence properties, and the mode-coupling matrix will be less numerically unstable.

2.2.4 Noise hardening

The result presented in section 2.2.2 may seem slightly disconcerting: by correcting for the noise-like component of the mask pseudo- C_ℓ , which we can estimate entirely from the data, we are able to *exactly* correct for the bias induced by the contribution to our estimate of the noise power spectrum coming from the signal variance. However, this contribution in principle depends on the true signal power spectrum, which we do not know a priori. For this to be possible, it must be that our correction of the mask power spectrum automatically makes the estimator immune to *any* noise-like (i.e., constant additive) contribution to the power spectrum. In other words, the corrected mode-coupling matrix $M_{\ell\ell'}^{w,S} \equiv M_{\ell\ell'}^w - M_{\ell\ell'}^{w,N}$ must satisfy that $\sum_{\ell'} M_{\ell\ell'}^{w,S} S_{\ell'} = 0$ if $S_{\ell'} = \text{constant}$ for all ℓ . In linear algebra terms, $M_{\ell\ell'}^{w,S}$

¹¹As we discussed, \tilde{S}_ℓ^f is actually a biased estimator of the signal component of \tilde{C}_ℓ^f , since we are subtracting a contribution sourced by the signal variance. Nevertheless, having corrected \tilde{C}_ℓ^w in the same way, the resulting estimator of S_ℓ will be unbiased.

¹²The usual caveats about power spectrum binning and invertible mode-coupling matrices should be borne in mind, however. See [23].

has a zero eigenvalue with eigenvector $v_\ell = 1 \forall \ell$. We can show this explicitly:

$$\begin{aligned}
\sum_{\ell'} M_{\ell\ell'}^{w,S} &= \frac{1}{4\pi} \sum_{\ell'\ell''} (2\ell' + 1)(2\ell'' + 1) \begin{pmatrix} \ell & \ell' & \ell'' \\ 0 & 0 & 0 \end{pmatrix}^2 [\tilde{C}_{\ell''}^{rw} - \tilde{N}^w], \\
&= \sum_{\ell''} \frac{2\ell'' + 1}{4\pi} [\tilde{C}_{\ell''}^{rw} - \tilde{N}^w], \\
&= \sum_{i \neq j} \frac{w_i w_j}{4\pi} \sum_{\ell'' m''} Y_{\ell'' m''}^* Y_{\ell'' m''}, \\
&= \sum_{i \neq j} \frac{w_i w_j}{4\pi} \delta^D(\hat{\mathbf{n}}_i, \hat{\mathbf{n}}_j) = 0,
\end{aligned}$$

where, in the second line, we have used the orthogonality relation eq. (A.10), in the third line we have expanded \tilde{C}_ℓ^{rw} and \tilde{N}^w using eqs. (2.9) and (2.17), and in the last line we have used the completeness of the spherical harmonic functions (eq. (A.8)).

In practice, what this means is that this power spectrum estimator is not able to recover any completely flat component of the true power spectrum or, in other words, that it has infinite variance for these flat modes. This is not a strong shortcoming, since the amplitude of such modes can be reconstructed from the zero-lag correlation function (i.e., the variance of the field across all sources), assuming a model for the noise. This also means that, to avoid numerical instabilities, it will be useful to employ singular value decomposition methods before inverting the noise-corrected mode-coupling matrix. The crucial advantage is the fact that we do not need to know which fraction of the total field variance is attributed to uncorrelated noise and which fraction to the signal autocorrelation: changes in either of them will only affect, and are exactly captured by, \tilde{N}^f . This fundamentally establishes the unbiasedness of our estimator.

2.2.5 Cross-correlations

Cross-correlations between two different fields measured on two disjoint sets of points, or cross-correlations between standard maps and discretely-sampled fields, can be estimated by making use of the standard pseudo- C_ℓ algorithm, taking advantage of discrete harmonic transform methods in the case of the sampled fields. No additional care must be taken in this case about subtracting noise components in the mask or field power spectra, since these are zero by construction. Care must be taken, however, if the two different fields to be correlated have been sampled at the same discrete positions. A good example is the cross-correlation between the cosmic shear signal of a set of sources and the PSF ellipticity at their position, as a diagnostic for potential systematics. In the most general case, different weights may be applied to the sources when mapping either field. In this case, the mask and masked-field pseudo- C_ℓ s must be corrected for noise bias.

Let e and f be the two fields being correlated, and let v_i and w_i , respectively, be the weights applied to each source when mapping either field. The noise-like contributions to the mask pseudo- C_ℓ and to the masked-field pseudo- C_ℓ are simple generalizations of those presented in the case of auto-correlations:

$$\tilde{N}^{vw} \equiv \frac{1}{4\pi} \sum_i v_i w_i, \quad \tilde{N}^{ef} \equiv \frac{1}{4\pi} \sum_i v_i w_i e_i f_i. \quad (2.23)$$

2.3 Spin- s fields

Let $f^a(\hat{\mathbf{n}})$, with $a \in \{1, 2\}$ be a spin- s_f field, and let $f_{\ell m}^\alpha$ be its spherical harmonic coefficients (with $\alpha \in \{E, B\}$, corresponding to E - and B -modes). The two are related via a spin- s_f spherical harmonic transform (and its inverse):

$$f^a(\hat{\mathbf{n}}) = \sum_{\ell m} {}_{s_f} Y_{\ell m}^{a\alpha}(\hat{\mathbf{n}}) f_{\ell m}^\alpha, \quad \leftrightarrow f_{\ell m}^\alpha = \int d\hat{\mathbf{n}} f^a(\hat{\mathbf{n}}) {}_{s_f} Y_{\ell m}^{a\alpha*}(\hat{\mathbf{n}}), \quad (2.24)$$

where the generalized spin- s spherical harmonic functions ${}_s Y_{\ell m}^{a\alpha}(\hat{\mathbf{n}})$ are defined in appendix A, and we implicitly sum over all repeated indices (e.g. a and α above).

The pseudo- C_ℓ of the masked field is defined in the same way as in the scalar case, except now we have to keep track of E/B -mode indices:

$$\tilde{C}_\ell^{f,\alpha\beta} \equiv \frac{1}{2\ell+1} \sum_m f_{\ell m}^\alpha f_{\ell m}^{\beta*}. \quad (2.25)$$

The generalization of the scalar pseudo- C_ℓ estimator from eq. (2.4) to spin- s fields is:

$$\langle \tilde{C}_\ell^{f,\alpha\beta} \rangle = \sum_{\ell'} (\mathbf{M}_{\ell\ell'}^w)^{\alpha\beta} S_{\ell'}^{\alpha'\beta'} + \tilde{N}_\ell^{\alpha\beta}, \quad (2.26)$$

where the mode-coupling matrix is [22, 23]:

$$\mathbf{M}_{\ell\ell'}^w = \frac{2\ell'+1}{4\pi} \sum_{\ell''} \tilde{C}_{\ell''}^w (2\ell''+1) \begin{pmatrix} \ell & \ell' & \ell'' \\ s_f & -s_f & 0 \end{pmatrix}^2 \mathbf{D}_{\ell+\ell'+\ell''}, \quad (2.27)$$

and the matrix \mathbf{D}_n is:

$$(\mathbf{D}_n)_{EE}^{EE} = (\mathbf{D}_n)_{EB}^{EB} = (\mathbf{D}_n)_{BE}^{BE} = (\mathbf{D}_n)_{BB}^{BB} = \frac{1+(-1)^n}{2}, \quad (2.28)$$

$$(\mathbf{D}_n)_{BB}^{EE} = -(\mathbf{D}_n)_{BE}^{EB} = -(\mathbf{D}_n)_{EB}^{BE} = (\mathbf{D}_n)_{EE}^{BB} = \frac{1-(-1)^n}{2}. \quad (2.29)$$

Following the same logic used in section 2.2, in the case of discretely-sampled fields, the spherical harmonic transform of a sampled and masked spin- s_f field is

$$f_{\ell m}^{w,\alpha} = \sum_i w_i f_i^a {}_s Y_{\ell m,i}^{a\alpha*}, \quad (2.30)$$

and its pseudo- C_ℓ is

$$\tilde{C}_\ell^{f,\alpha\beta} = \sum_{ij} w_i w_j f_i^a f_j^b \frac{1}{2\ell+1} \sum_m {}_s Y_{\ell m,i}^{a\alpha*} {}_s Y_{\ell m,j}^{b\beta}. \quad (2.31)$$

Now, let us separate the field into signal and noise components $f^a = s^a + n^a$, where the signal is statistically isotropic ($\langle s_{\ell m}^\alpha s_{\ell' m'}^{\beta*} \rangle = \delta_{\ell\ell'} \delta_{mm'} S_\ell^{\alpha\beta}$), and the noise is uncorrelated between different sampling points, which in this case means $\langle n_i^a n_j^b \rangle = \delta^{ab} \delta_{ij} \sigma_{N,i}^2$. Inserting this in the previous equation, and using eq. (A.6), we obtain the noise bias:

$$\tilde{N}_\ell^{\alpha\beta} = \sum_i w_i^2 \sigma_{N,i}^2 \frac{1}{2\ell+1} \sum_m {}_s Y_{\ell m,i}^{a\alpha*} {}_s Y_{\ell m,i}^{a\beta} = \frac{\delta^{\alpha\beta}}{4\pi} \sum_i w_i^2 \sigma_{N,i}^2. \quad (2.32)$$

Based on this result, a potentially useful estimate of the noise bias, using only our measurements of f , would be:

$$\tilde{N}^{f,\alpha\beta} \equiv \frac{\delta^{\alpha\beta}}{4\pi} \sum_i w_i^2 \frac{1}{2} f_i^a f_i^b \delta_{ab}. \quad (2.33)$$

As in the spin-0 case, this estimator will be biased due to the contribution from the signal component of f . To calculate this bias, we take the expectation value of the previous equation, and find:

$$\begin{aligned} \langle \tilde{N}^{f,\alpha\beta} \rangle &= \tilde{N}_\ell^{\alpha\beta} + \frac{\delta^{\alpha\beta}}{4\pi} \sum_i w_i^2 \frac{\delta_{ab}}{2} \langle s_i^a s_i^b \rangle \\ &= \tilde{N}_\ell^{\alpha\beta} + \frac{\delta^{\alpha\beta}}{4\pi} \sum_i w_i^2 \frac{1}{2} \sum_\ell S_\ell^{\alpha'\beta'} \sum_m s Y_{\ell m, i}^{a\alpha'} s Y_{\ell m, i}^{b\beta'} \\ &= \tilde{N}_\ell^{\alpha\beta} + \delta^{\alpha\beta} \tilde{N}^w \sum_\ell \frac{2\ell + 1}{4\pi} \frac{S_\ell^{EE} + S_\ell^{BB}}{2} \\ &\equiv \tilde{N}_\ell^{\alpha\beta} + \delta^{\alpha\beta} \tilde{N}^w \sigma_S^2, \end{aligned} \quad (2.34)$$

in analogy with eq. (2.16), after generalizing the expression for the field variance σ_S^2 in eq. (2.15). In the second line above, we have simply expanded s_i^a in its spherical harmonic coefficients, and used their statistical isotropy. We have then made use of eq. (A.6) in the third line.

The final step in our derivation shows that the contribution from \tilde{N}^w to eq. (2.26) cancels out the term $\propto \sigma_S^2$ in the last equation, thus yielding an unbiased power spectrum estimator. To begin with, let us replace \tilde{C}_ℓ^w with \tilde{N}^w in eq. (2.27) to calculate the mask noise contribution to the mode-coupling matrix, $M_{\ell\ell'}^{w,N}$:

$$M_{\ell\ell'}^{w,N} = \tilde{N}^w \frac{2\ell' + 1}{4\pi} \sum_{\ell''} (2\ell'' + 1) \begin{pmatrix} \ell & \ell' & \ell'' \\ s_f & -s_f & 0 \end{pmatrix}^2 D_{\ell+\ell'+\ell''}. \quad (2.35)$$

To simplify this further, note that all elements of $D_{\ell+\ell'+\ell''}$ either are zero, or take the form $[1 \pm (-1)^{\ell+\ell'+\ell''}]/2$ (see eq. (2.28)). We can use the term $\propto (-1)^{\ell+\ell'+\ell''}$ to swap the sign of the lower spin indices of one of the two Wigner $3j$ symbols, using eq. (A.11). The resulting sum over ℓ'' for this term is then zero by virtue of eq. (A.10), and we are left only with the contribution from the first term in the D matrix elements, which is always 1/2 (or zero). Again, using eq. (A.10) on this term, we obtain the simplified result:

$$M_{\ell\ell'}^{w,N} = \tilde{N}^w \frac{2\ell' + 1}{4\pi} \frac{d}{2}, \quad (2.36)$$

where the matrix d has elements:

$$(d)_{EE}^{EE} = (d)_{EB}^{EB} = (d)_{BE}^{BE} = (d)_{BB}^{BB} = (d)_{BB}^{EE} = (d)_{EE}^{BB} = 1, \quad (d)_{BE}^{EB} = (d)_{EB}^{BE} = -1, \quad (2.37)$$

and zero otherwise. Substituting this in eq. (2.26), and after a bit of algebra,¹³ we find that the resulting contribution to $\langle \tilde{C}_\ell^{f,\alpha\beta} \rangle$ is

$$(\Delta \tilde{C}_\ell^{f,\alpha\beta})_N = \delta^{\alpha\beta} \tilde{N}^w \sigma_S^2. \quad (2.38)$$

¹³This can be done by writing out d as a 4×4 matrix, and multiplying it by the vector $(S_\ell^{EE}, S_\ell^{EB}, S_\ell^{BE}, S_\ell^{BB})$, with $S_\ell^{BE} = S_\ell^{EB}$.

Thus, the recipe for an unbiased estimator described in section 2.2.3 can be directly applied to spin- s fields, simply replacing the scalar noise bias \tilde{N}^f by its spin generalization $\tilde{N}^{f,\alpha\beta}$ in eq. (2.33).

2.3.1 Different spins

One important generalization of the formalism we just described is the calculation of cross-power spectra between two fields of different spin sampled at the same source positions. A good example of this would be measurements of both the shape (spin-2) and size (spin-0) of galaxies, as proxies for weak lensing shear and convergence. The formalism we just described, in which the mode-coupling matrix is estimated from the noise-debiased mask pseudo- C_ℓ , can be applied directly to these situations without any modifications. In fact, the orthogonality relation eq. (A.10) guarantees that, for fields of different spin, the contribution from \tilde{N}^w to the mode-coupling matrix vanishes exactly. Likewise, by virtue of the addition theorem for spin-weighted spherical harmonics (eq. (A.6)), the noise bias to the pseudo- C_ℓ of the masked fields also vanishes in the limit of uncorrelated noise. Thus, in these cases, one may even forego the calculation of \tilde{N}^w and \tilde{N}^f altogether.

2.4 Clustering

A conceptually different type of discretely sampled field is the case of galaxy clustering, in which the overdensity of points in the discrete source catalog is itself the field. This case has been discussed in detail in [5], but we repeat part of the discussion here for completeness. In this case, the number density of objects is

$$n_d(\hat{\mathbf{n}}) = \sum_{i=1}^N w_i^d \delta^D(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i), \quad (2.39)$$

where w_i^d is the weight of the i -th galaxy (which may be used to account for observational conditions, instrumental effects, such as fiber collisions, and other systematics [34]).

Critical in the study of galaxy clustering is our ability to account for spatial variations in the expected mean number density of sources $\bar{n}(\hat{\mathbf{n}})$, caused by depth variations across the footprint, contamination by galactic or instrumental systematics, and the complex geometry of the observed footprint itself. This has been commonly done through the use of “random catalogs”: synthetic catalogs containing discrete points that trace the underlying expected density of sources, but that are otherwise not clustered. The mean number density of sources in this case is then

$$\bar{n}(\hat{\mathbf{n}}) = \alpha \sum_{i=1}^{N_r} w_i^r \delta^D(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i), \quad (2.40)$$

where w_i^r is the weight of the i -th random point, N_r is the total number of random points, and the normalization factor α accounts for the potentially different number of objects in the data and random catalogs:

$$\alpha \equiv \frac{\sum_{i=1}^N w_i^d}{\sum_{i=1}^{N_r} w_i^r}. \quad (2.41)$$

Ideally, the field whose power spectrum we wish to estimate is the galaxy overdensity $f(\hat{\mathbf{n}}) \equiv \delta_g(\hat{\mathbf{n}}) \equiv n_d(\hat{\mathbf{n}})/\bar{n}(\hat{\mathbf{n}}) - 1$, although, obviously, the discretized nature of $\bar{n}(\hat{\mathbf{n}})$ complicates the construction of such a map. We can proceed, however, under the following consideration: a nearly optimal choice of mask for the standard pseudo- C_ℓ method is one that is proportional to the inverse noise variance of the field under study. Assuming galaxy positions are known exactly, the measurement noise of the overdensity field is purely Poissonian with local mean and variance both given by $1/\bar{n}(\hat{\mathbf{n}})$, respectively. A reasonable choice of mask is therefore $w(\hat{\mathbf{n}}) = \bar{n}(\hat{\mathbf{n}})$. With this choice, the masked field $f^w(\hat{\mathbf{n}}) = \delta_g(\hat{\mathbf{n}})\bar{n}(\hat{\mathbf{n}})$ and the associated mask are simply

$$f^w(\hat{\mathbf{n}}) = n_d(\hat{\mathbf{n}}) - \bar{n}(\hat{\mathbf{n}}) = \sum_{i=1}^N w_i^d \delta^D(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i) - \alpha \sum_{i=1}^{N_r} w_i^r \delta^D(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i), \quad (2.42)$$

$$w(\hat{\mathbf{n}}) = \alpha \sum_{i=1}^{N_r} w_i^r \delta^D(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i). \quad (2.43)$$

The methodology described in section 2.2.3 can therefore be extended directly to the case of galaxy clustering with the following changes:

- The SHTs of the masked field and the mask are

$$f_{\ell m}^w = \sum_{i=1}^N w_i^d Y_{\ell m, i}^* - w_{\ell m}, \quad w_{\ell m} = \alpha \sum_{i=1}^{N_r} w_i^r Y_{\ell m, i}^*, \quad (2.44)$$

- The shot noise components to field and mask are:

$$\tilde{N}^f = \frac{1}{4\pi} \sum_{i=1}^N (w_i^d)^2 + \tilde{N}^w, \quad \tilde{N}^w = \frac{1}{4\pi} \sum_{i=1}^{N_r} (\alpha w_i^r)^2. \quad (2.45)$$

Although the use of randoms to characterize survey homogeneity has been common, particularly in the analysis of spectroscopic galaxy surveys, it is also possible to map spatial variations in survey homogeneity at every point in the celestial sphere. It is therefore important to also consider the case in which a ‘‘completeness map’’ W is provided, such that $W(\hat{\mathbf{n}}) \propto \bar{n}(\hat{\mathbf{n}})$. In this case, the method must be adapted as follows:

- The sky mask is generated by normalizing the completeness map, to make it correspond to the expected number density of galaxies:

$$w(\hat{\mathbf{n}}) = \alpha W(\hat{\mathbf{n}}), \quad \alpha = \frac{\sum_{i=1}^N w_i^d}{\int d\hat{\mathbf{n}} W(\hat{\mathbf{n}})}. \quad (2.46)$$

- The SHT of the masked field is still given by the first equality in eq. (2.44), with the mask SHT $w_{\ell m}$ estimated from the continuous $w(\hat{\mathbf{n}})$ as usual.
- Since the mask is now a map, rather than a collection of sources, its shot noise is $\tilde{N}^w = 0$. The shot-noise component of the field’s pseudo- C_ℓ , in turn, is $\tilde{N}^f = \sum_{i=1}^N (w_i^d)^2 / (4\pi)$.

3 Validation

Having presented the mathematical formalism of an unbiased power spectrum estimator for discretely sampled fields, we now proceed to validate it against simulations. We focus on what is likely the most relevant probe for which this formalism is useful: cosmic shear catalogs. The results presented here will therefore correspond to spin-2 fields, but other spin cases do not change the problem conceptually.¹⁴ We showcase examples of the spin-1 and spin-0 cases in sections 4.2 and 4.3, respectively.

Cosmic shear measures the apparent distortion of a galaxy’s shape due to the gravitational lensing effect caused by the foreground matter distribution integrated along the line of sight. This effect, while weak and indistinguishable from the galaxy’s intrinsic ellipticity when considering single objects, represents a powerful probe of the cosmic large-scale structure when studied in a statistical way on large catalogs of galaxies. In the weak lensing regime, the shear $\gamma_{a,i}$ that the i -th galaxy’s image has undergone, provides a linear contribution to its measured ellipticity $e_{a,i}$ (estimated from the traceless component of the projected inertia tensor, see e.g. [25]):

$$e_{a,i} = \gamma_{a,i} + n_{a,i}, \quad (3.1)$$

where $a = 1, 2$ labels the two components of ellipticity and shear, and the noise component $n_{a,i}$ includes both the intrinsic (i.e., unlensed) galaxy ellipticity and any additional measurement noise. We work under the common assumption that intrinsic noise dominates over correlated measurement noise, hence we consider $n_{a,i}$ to be uncorrelated between different sources.

Section 3.1 below describes the setup used to validate the catalog-based methodology described in the previous section, comparing it against the results of the standard map-based approach. The results of this validation, as a function of the number density of catalog sources, are then presented in section 3.2.

3.1 Validation setup

To validate our estimator, we generate 500 synthetic cosmic shear catalogs. Each catalog is created by first generating a sky realization of the spin-2 cosmic shear field as a Gaussian random field from a theoretical E -mode power spectrum. This realization is generated using the HEALPix pixelization scheme with a resolution parameter $N_{\text{side}}^{\text{hires}} = 4096$, corresponding to a pixel size of $\sim 0.86'$. The native resolution is chosen to be significantly higher than the angular scales we will try to extract using either the map-based or catalog-based methods, to ensure that the finite-resolution effects inherent to these ground-truth maps are negligible compared to the typical statistical error. From this map, we then produce a mock observed shear catalog, by sampling it at the positions of a discrete set of sources. Specifically, to each source we assign the value of the shear field corresponding to the HEALPix pixel it falls in.¹⁵ We do not add any shape noise to the catalog, in order to reduce the statistical uncertainties of the recovered power spectra, thereby maximizing our sensitivity to any potential bias.

Once a mock shear catalog has been generated, we estimate its angular power spectrum using the catalog-based method described above, and the standard map-based method. For

¹⁴Except perhaps in the case $s = 0$, where the separation between E - and B -modes is meaningless.

¹⁵Unfortunately, more sophisticated interpolation schemes are highly computationally intensive for HEALPix.

the map-based case, we use a `HEALPix` resolution of $N_{\text{side}} = 1024$ to keep the computational cost at a manageable level. We note that this is a lower resolution than $N_{\text{side}} = 4096$, used by most state-of-the-art pixel-based cosmic shear analyses [25, 33, 35], but the same resolution as the one used by the official DES-Y3 cosmic shear analysis [31]. We follow the procedure outlined in [25]: firstly, we construct a map where each pixel is assigned the average shear value. In each pixel, the local shear field is calculated as the sum of source ellipticities in it, each one multiplied by the individual measurement weights, and divided by the sum of the weights of all sources in that pixel. As our simulated catalogs do not contain any additional shape noise, we do not subtract any data-based noise bias in the standard map-based approach (although we do subtract \tilde{N}^f and \tilde{N}^w in the catalog-based method, as described in section 2). The `HEALPix` pixelization suppresses the small-scale information in the shear map, reducing the map RMS by a pixel window function, which needs to be accounted for when computing pseudo- $C_{\ell s}$ [36]. Assuming that the pixelization is coarse enough so that every pixel contains at least one galaxy, we can approximate the pixel window function in real space by a Gaussian smoothing kernel with an FWHM of $\sim 41.7 \text{ deg}/N_{\text{side}}$ size [37]. We do not correct for this bias in the simulated data, but account for it when comparing with the bandpower-convolved theory power spectrum.

In order to produce mock catalogs with realistic source distributions, we use the angular positions and shape weights of real data from the public KiDS-1000 shear catalog [38]. In particular, we use all sources from the fourth photometric bin, which contains 4,544,395 sources, observed in two different fields, covering an area of approximately 800 deg^2 after masking. The ground-truth maps were generated using the theoretical prediction for this redshift bin based on the best-fit ΛCDM parameters from *Planck* [39], assuming no intrinsic alignments. We note that, by choosing to use real KiDS-1000 sources, we ensure that the method is validated on catalogs with not only a realistic number density, but also a spatial distribution that reproduces both the expected clustering of sources, and their inhomogeneity caused by depth variations across the observed footprint. In order to quantify the performance of the catalog- and map-based methods as a function of source number density, we repeat our analysis on down-sampled versions of these catalogs, in which we select a fraction of the KiDS-1000 sources. Specifically, we will study the results of selecting only 1% and 0.01% of the sources. We choose a multipole binning of constant width $\Delta\ell = 30$ up to a maximum of $\ell_{\text{max}} = 2N_{\text{side}} = 2048$.

3.2 Validation results

Figure 1 shows the main results of our validation exercise. The upper right panel shows the results in the fiducial setup, in which all the KiDS-1000 sources are used to measure the power spectrum. The power spectrum averaged over all simulations is shown as orange and blue circles for the catalog- and map-based approaches, respectively. The theoretical predictions, estimated by convolving the input power spectrum with the bandpower window functions corresponding to the mode-coupling matrix derived with either method, are shown as dashed red and solid grey lines (the latter barely visible behind the red dashed lines). We multiply the theory spectra by the squared `HEALPix` pixel window function ($N_{\text{side}}^{\text{ hires}} = 4096$ for the catalog-based estimator, $N_{\text{side}} = 1024$ for the map-based one). The narrow subpanel below

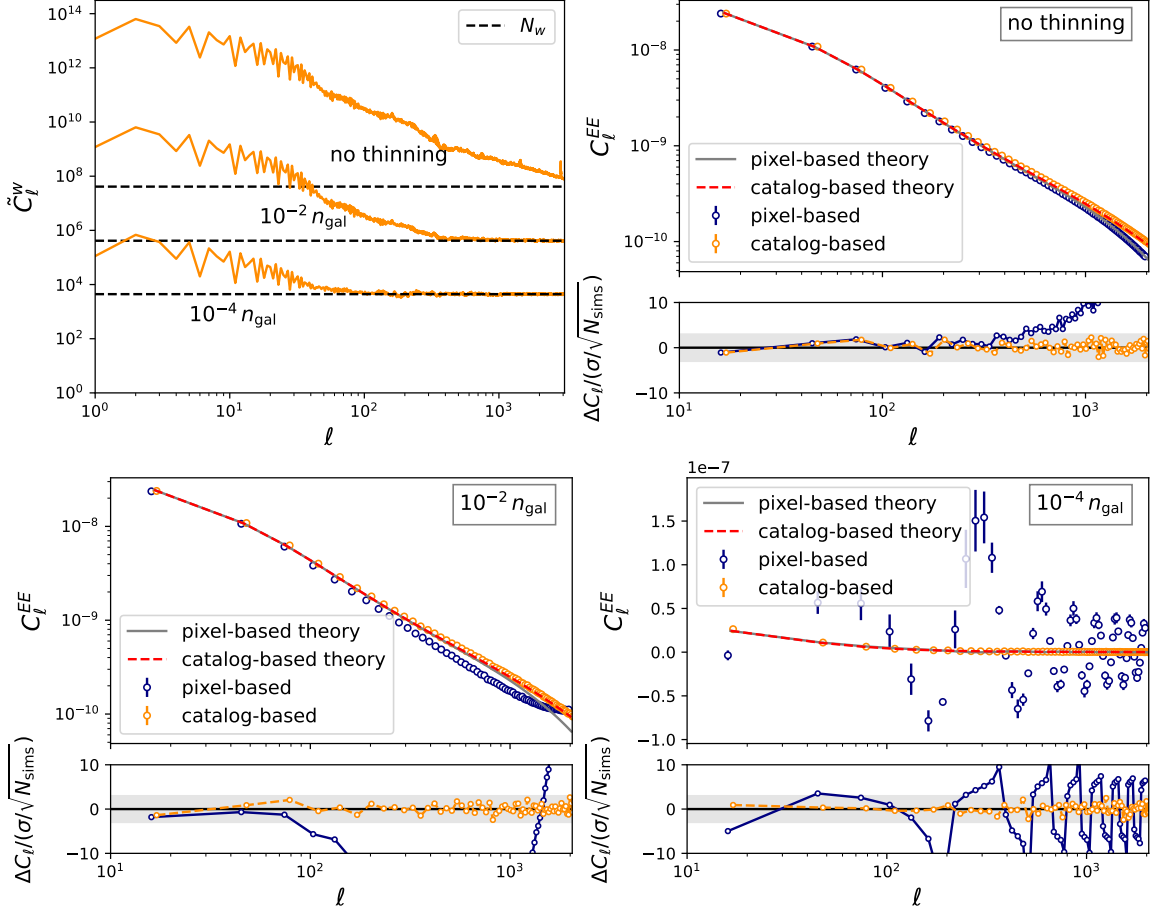


Figure 1. Validation shear EE power spectra from 500 KiDS-1000-like mock catalogs with different thinning factors. *Left upper panel:* Coupled pseudo power spectrum of the sky mask using the catalog-based approach (orange lines). Black horizontal dashed lines mark the mask shot noise level, which we subtract before computing the mode coupling matrix. The three sets of lines from top to bottom are the KiDS-1000 catalog without thinning, with a thinning factor of 100, and a thinning factor of 10,000. *Right upper panel:* Mean decoupled power spectrum from simulations, evaluated with the pixel-based (blue circles) and the catalog-based approach (orange circles, including a horizontal offset for clarity), compared to theory bandpowers computed using either the pixel-based or the catalog-based mode coupling (grey solid lines and red dashed lines, respectively). The narrow subpanel below shows residuals between estimates and theory (corrected for the corresponding HEALPix pixel window), measured in units of the expected error on the mean. The grey horizontal band shows the $\pm\sigma$ error range. *Lower panels:* results from 500 simulated catalogs that contain 100 times and 10^4 times less sources than the original catalog.

shows the residual of each method with respect to its corresponding prediction normalized by the error on the mean of the 500 simulations (i.e., the scatter across simulations divided by $\sqrt{500}$). The horizontal grey bands show $3\text{-}\sigma$ confidence intervals.

The catalog-based method is able to produce an unbiased estimate of the power spectrum, while at $\ell \gtrsim 500$, the map-based approach deviates significantly due to pixelization effects at the map level and shot-noise-related bias in the estimator itself. Note that, in a realistic scenario, the statistical uncertainties would be at least $\sqrt{500} \sim 22$ times larger, not accounting for the impact of shape noise, and hence these deviations would likely be subdominant. We discuss the various sources of this bias in appendix B. For the catalog-based estimator, we must account for the pixelization of our ground truth shear map, and can do so accurately by means of the HEALPix window function.¹⁶ Ideally, we could choose an even higher $N_{\text{side}}^{\text{hires}}$ to avoid this effect, but given the higher numerical cost, we prefer to correct the theory accordingly. With the map-based approach, the estimate is biased both for the uncorrected case (not shown in the figure) and the case where the theory is convolved with the pixel window function (shown). As explained in appendix B, the two main reasons are the impact of aliasing, and the numerical instability of the pixel-based mode-coupling matrix due to mask shot noise. Note that this coupling matrix is truncated to multipoles $\ell < 3N_{\text{side}}$, and thus fails to capture the full correlation between scales induced by the mask for the sparsest catalogs.

The bottom panels of figure 1 show the results of repeating this analysis on catalogs containing 100 times and 10,000 times fewer sources (left and right panels, respectively). We find that, in spite of the substantial drop in number density, the catalog-based method is able to recover unbiased power spectrum measurements in both cases within the range of scales of interest. The map-based approach, however, recovers significantly biased measurements, which cannot be corrected for by our analytical pixel window function. This bias becomes gradually larger on smaller scales for the 1% catalog, reaching very significant levels after $\ell \sim 100$. For the 0.01% catalog, in turn, the map-based method fails catastrophically, due to significant shot noise in the field pseudo- C_ℓ estimator, as well as numerical instabilities in the estimated mode-coupling matrix. These lead to the method misestimating the measured power spectrum at all scales (bottom right panel). In contrast, the catalog-based approach remains unbiased even for this extremely low number density (the shear field is sampled at the positions of just 454 sources). For the thinned catalogs, the catalog-based estimator is again fully unbiased even at scales $\ell > 1024$ after accounting for pixelization-related bias mentioned above, which becomes less important for smaller source densities due to higher estimator noise ($\sim 2\sigma$ and $\lesssim 1\sigma$ at $\ell = 1024$ for 1% and 0.01% thinning, respectively).

The upper left panel of figure 1 shows the mask pseudo- C_ℓ in the catalog-based case, \tilde{C}_ℓ^w , for the fiducial catalog, and the 1% and 0.01% catalogs (from top to bottom). In each case, the mask shot noise level \tilde{N}^w is marked by horizontal black dashed lines, and the raw mask pseudo- C_ℓ s are shown as the solid lines. We see that, for the original KiDS source density, the mask shot noise is low enough to not affect the mode-coupling matrix significantly, except above $\ell \sim 2000$. In turn, the 1% catalog becomes shot-noise dominated at $\ell \sim 200$ (compared

¹⁶The simulated catalog contains field values that are, by design, constant within every pixel of the ground truth map. This makes the catalog-based field effectively band-limited, with no power at scales below the pixel scale. We can account for this by means of the analytic pixel window function, as shown in appendix B.

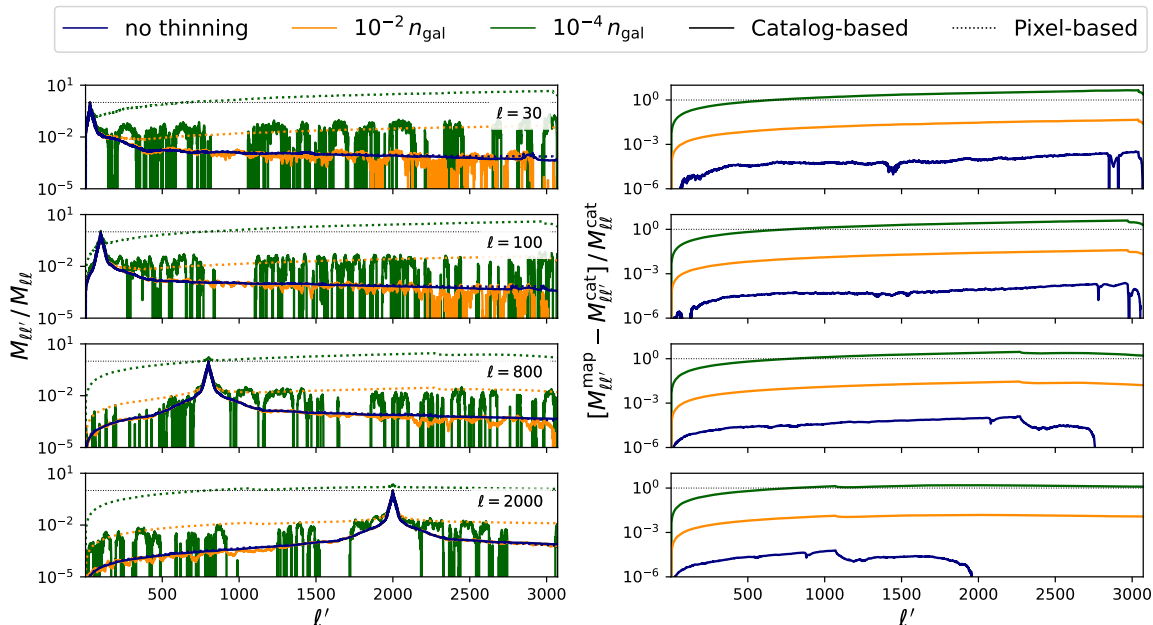


Figure 2. Rows of catalog-based mode coupling matrix and comparison between pixel-based and catalog-based approach, displayed for KiDS-1000-like mock catalogs of three different thinning factors. Blue, orange, and green are the cases of no thinning, thinning factor 100, and thinning factor 10,000. From top bottom, we show different rows of the mode coupling matrix: $\ell = 30, 100, 800, 2000$. *Left panels:* Mode coupling matrix for catalog- (solid lines) and pixel-based $C_{\ell s}$ (dotted lines), normalized to the value of the $\ell = \ell'$ element. *Right panels:* Difference between the rows of the mode coupling matrix in the map-based and the catalog-based case, normalized by the catalog-based element at $\ell = \ell'$.

to $\ell \sim 100$ where the decoupled power spectrum estimate starts deviating significantly from the input). Finally, the 0.01% catalog is shot-noise dominated at almost all scales, leading to severe numerical instabilities in the standard pixel-based algorithm, which does not account for the impact of mask shot noise.

Figure 2 shows the mode coupling matrices computed from mock catalogs, and compares the result to the map-based analog. We show this for the full KiDS-1000 catalog (blue) and the two thinned versions (orange, green) but focus on the former for now. The left panels show specific columns of the $\frac{EE}{EE}$ component of the catalog-based corrected mode coupling matrix $M_{\ell\ell'}^{w,S}$ (solid lines, see section 2.2.4) and the standard pixel-based mode coupling (dashed lines), normalized to one at $\ell = \ell'$. The right panels show, in the same units, the relative difference between the columns of the map-based mode coupling matrix (2.27) and the catalog-based one.¹⁷ From top to bottom, we display columns with $\ell = (30, 100, 800, 2000)$. We find that, in the case of the full catalog, both matrices differ by a factor of less than 10^{-4} , indicating that the mask shot noise correction is subdominant. The matrix columns peak at $\ell = \ell'$ as expected, rapidly decline to percent level at $|\ell - \ell'| \sim 200$, and converge to a nearly constant value at the 0.1 percent level.

¹⁷For a direct comparison between the catalog-based and the map-based mode coupling matrix elements, we need to divide the latter by the squared pixel area, $(4\pi/N_{\text{pix}})^2$.

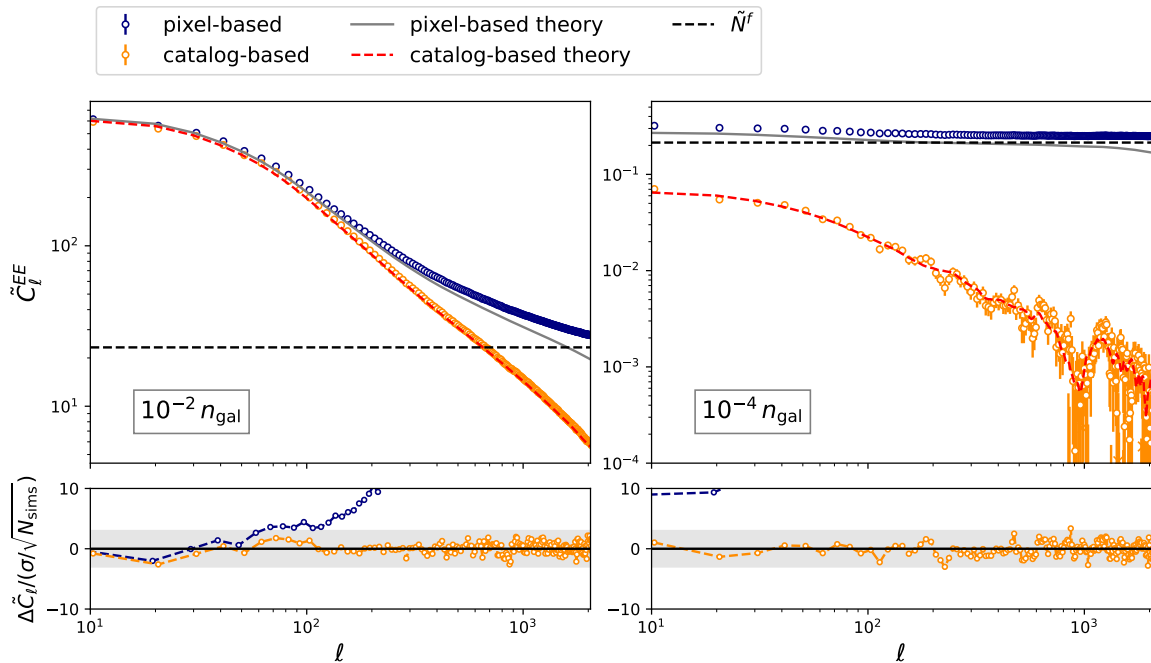


Figure 3. Coupled pseudo power spectra for KiDS1000-like validation catalogs, thinned by factors 100 and 10,000. The blue and orange error bars are centered at the mean from 500 simulations and their length corresponds to the standard deviations across 500 simulations divided by $\sqrt{500}$ in the map-based and catalog-based approach, respectively. The lower panels show the deviation between data and coupled theory, normalized by the error. The difference between both estimates corresponds to the total field noise \tilde{N}_f (black dashed line, see eq. (2.16)), which in this example is dominated by shot noise.

For the 1%-thinned catalog, both methods start differing at the percent level. The elements of the catalog-based matrix match on average those for the full catalog, but are much noisier for couplings between distant multipoles, while the pixel-based coupling matrix includes a now very significant contribution due to mask shot noise, with the expected scaling $\propto (2\ell + 1)$, see eq. (2.20). For the highly shot-noise dominated case with 454 sources, the catalog-based mode coupling becomes more noisy, with fluctuations reaching percent level, but its mean stays compatible with the one for the full catalog. On the other hand, the pixel-based coupling matrix is almost completely dominated by the shot noise component, and differs from the catalog-based version by more than 100% at $\ell \gtrsim 500$. This shows that even with an extremely low source number density, the catalog-based approach recovers a stable mode coupling matrix that allows us to estimate unbiased power spectra from sparsely sampled sources on the sky.

To understand what causes the bias in the pixel-based estimator, we compare both approaches in terms of the coupled pseudo- C_ℓ estimators. Figure 3 shows, in the same color scheme as in figure 1, the coupled shear EE pseudo- C_ℓ s (see eq. (2.26)) computed using the catalog-based and pixel-based approach, and compared to the input theory, multiplied by the square of the respective HEALPix pixel window, and coupled with the respective mode coupling matrix. The field noise level N_f is shown as a horizontal black dashed line

for comparison. The left and the right panel show the cases of 1% and 0.01% thinning, respectively, and the lower panels show the bias with respect to the mode-coupled theory prediction. The pixel- and catalog-based coupled estimators differ at similar scales as their decoupled versions shown in figure 1, but without being affected by the numerical instability of the inverse binned mode-coupling matrix.

In contrast, the coupled theory bandpowers do depend on the mode coupling matrix. We see that the catalog-based theory prediction becomes noisy for small scales, which is inherited from the noisy mode-coupling matrix seen in figure 2. For both thinned catalogs, the catalog-based estimator is unbiased, while the map-based estimator receives a bias at intermediate and small scales. In summary, figure 3 illustrates that for low source number density, the pixel-based coupled pseudo power spectrum is shot-noise dominated. Analogously, the mask shot noise dominates the pixel-based mode coupling at small scales and leads to a severe misestimation of the mode-coupled theory prediction. The decoupled pixel-based estimator inherits this field noise bias, and suffers from numerical instabilities from applying the inverse mode coupling matrix. In contrast, the catalog-based estimator does not deviate significantly from theory predictions, even when considering the statistical power of 500 noiseless simulations combined.

4 Example applications

4.1 Spin-2 example: cosmic shear from DES-Y3

As a first real-data example, we apply our estimator to the highest redshift bin of the galaxy shear catalog from DES-Y3 [40]. This catalog contains $\mathcal{O}(10)$ more galaxies in a four times larger survey footprint than the fourth redshift bin of KiDS-1000 used in section 3. Cosmic shear is a common example of a spin-2 field sampled at discrete positions, and one of the most powerful probes of the cosmic LSS. It measures the lensing-induced distortion of galaxy shapes due to the foreground matter distribution along the line of sight. In the Limber approximation [41], the cosmic shear angular power spectrum is related to the 3D matter power spectrum $P(k, z(\chi))$, and can be modeled as

$$C_\ell = G_\ell^2 \int \frac{d\chi}{\chi^2} q^2(\chi) P\left(k = \frac{\ell + 1/2}{\chi}, z(\chi)\right), \quad (4.1)$$

where k is the modulus of the Fourier wave vector, z is the redshift, and χ is the comoving radial distance. The lensing kernel $q(\chi)$ quantifies the radial contribution of foreground matter to the angular galaxy lensing power spectrum, and depends on the redshift distribution $p(z)$ of galaxies that may act as lensed images. In the case of a flat universe in which general relativity is valid and $c = 1$, the lensing kernel is given by

$$q(\chi) = \frac{3}{2} H_0^2 \Omega_m \frac{\chi}{a(\chi)} \int_{z(\chi)}^\infty dz' p(z') \frac{\chi(z') - \chi}{\chi(z')}, \quad (4.2)$$

where H_0 is the Hubble constant, Ω_m is the matter density parameter, and $a = 1/(1+z)$ is the scale factor. Transitioning from the 3D Laplacian of the gravitational potential to the angular Hessian of the associated lensing potential requires the correction factor [42]

$$G_\ell \equiv \sqrt{\frac{(\ell+2)!}{(\ell-2)! (\ell+1/2)^2}} \cdot 1. \quad (4.3)$$

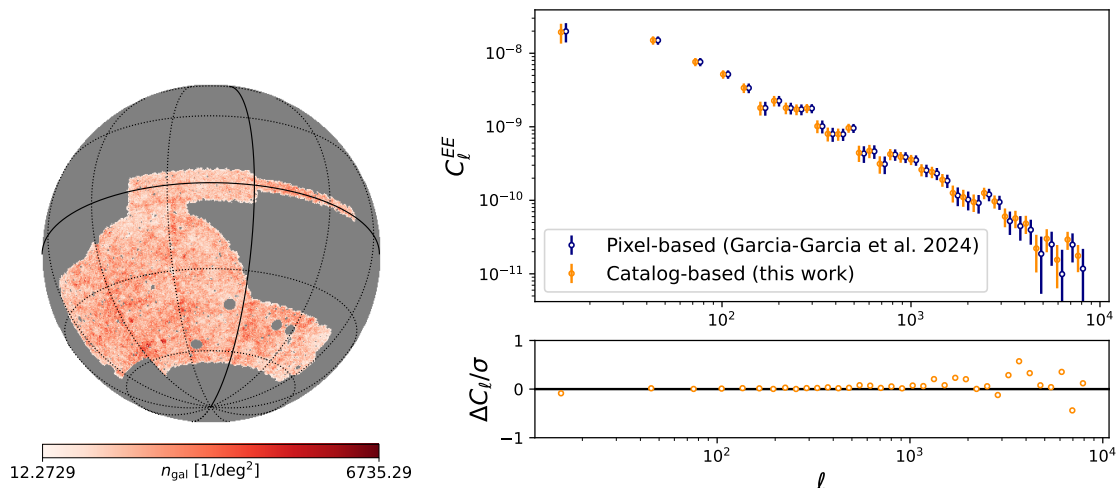


Figure 4. *Left panel:* Galaxy source catalog of the fourth redshift bin of DESY3 data used in this work, displayed in celestial coordinates. *Right panel:* decoupled C_{ℓ}^{EE} using the catalog-based approach (orange), and comparing with literature results from [33] (blue, including a horizontal offset for clarity). We consider error bars from Gaussian and super-sample covariance terms taken from [33]. The lower panel shows residuals of the catalog-based method with respect to the literature in units of the 1-sigma error.

Galaxy shear data are notoriously biased by galaxy shape noise, which has to be corrected for explicitly in the pixel-based approach [25], but is removed at the catalog level using our estimator.

We use the fourth redshift bin of the DES-Y3 data, containing 25,091,297 catalog galaxies between redshifts ~ 0.4 and ~ 1.5 . We correct the catalog galaxy ellipticities for linear and multiplicative bias as described in [33]. The survey footprint covers an effective sky area of 4143 deg² as shown in the left panel of figure 4. The weighted source number density of DES-Y3 is $n_{\text{eff}} = 5.59$ gal/arcmin², compared to 6.17 gal/arcmin² for KiDS-1000, yielding a similar mask shot noise level. Note that since we use the sparsest DES-Y3 redshift bin, our results on catalog-based power spectra represent an upper limit on the significance of pixelization results presented in [33].

Figure 4 in the right panel shows the results in terms of the decoupled power spectrum, comparing our catalog-based approach with recent pixel-based measurements using the pixel-based approach published in [33]. Both estimators use the same set of multipole bins up to $\ell = 8192$, corresponding to $2N_{\text{side}}$ of the HEALPix map used with the pixel-based estimator, or an angular scale of ~ 26 arcsec. We corrected the map-based estimate by the analytical pixel window function. Both results use 1-sigma error bars from [33], derived from an empirical power spectrum covariance that includes a Gaussian and a non-Gaussian super-sample contribution. The lower panel shows the relative difference between both methods in units of sigma. We find that both methods agree within 1 sigma within the shown multipole range.




Figure 5. *Left:* spatial distribution of the *Quiaia* quasar catalog used in this exercise, color-coded by its number density on the sphere. *Right:* distribution of spectro-photometric redshift estimates for the same sample, with the vertical band indicating the redshift range of selected sources used in this work.

4.2 Spin-1 example: proper motion surveys

Although most quantities of cosmological interest are either scalars or spin-2 objects (e.g. cosmic shear, CMB polarization), there are examples of particularly useful spin-1 fields, especially in terms of catalog-based observations. Examples of these are lensing displacements, galaxy angular momenta, and transverse peculiar velocities [9, 43]. Astrometric surveys, such as the *Gaia* satellite, are able to track the angular position of sources over time with high precision. Although the main use of such observations is reconstructing the dynamics of stars in the Milky Way, it has been argued that they may be useful for cosmology when applied to extra-galactic sources, placing constraints on the local rate of expansion [44, 45], detecting the Solar System’s proper motion dipole [46], and even constraining structure growth [9].

As an example application, we generate simulated observations of quasar proper motions for a *Gaia*-like sample. In particular, we use the positions of quasars in the public *Quiaia* quasar sample [47]. We select 342,971 quasars with spectro-photometric redshifts in the range $1 < z_{\text{sp}} < 1.5$. The spatial and redshift distribution of this sample is shown in figure 5. Next, we generate the theoretical prediction for the angular power spectrum of the expected cosmological proper motion signal. Following the formalism in [9], the projected proper motion map $\dot{\mathbf{A}}$ is a spin-1 field given by

$$\dot{\mathbf{A}}(\hat{\mathbf{n}}) = \int \frac{dz}{(1+z)\chi} p(z) (1 - \hat{\mathbf{n}}^T \hat{\mathbf{n}}) \cdot \mathbf{v}(\chi \hat{\mathbf{n}}), \quad (4.4)$$

where $p(z)$ is the redshift distribution of the sample, \mathbf{v} is the peculiar velocity, projected on the plane transverse to the line of sight via the projector $1 - \hat{\mathbf{n}}^T \hat{\mathbf{n}}$. At linear order in perturbation theory, $\dot{\mathbf{A}}$ is a pure E -mode field, since it can be written as the covariant derivative of a projected velocity potential [9]. Thus, in what follows, we study only the power spectrum of its E -mode component.

We compute a theoretical prediction for the signal power spectrum using the redshift distribution of the sample and assuming standard Λ CDM cosmological parameters. We then generate Gaussian realizations of this field at resolution $N_{\text{side}} = 1024$, and sample them at the positions of the *Quiaia* quasars. Since the peculiar motion signal is most relevant at large scales,

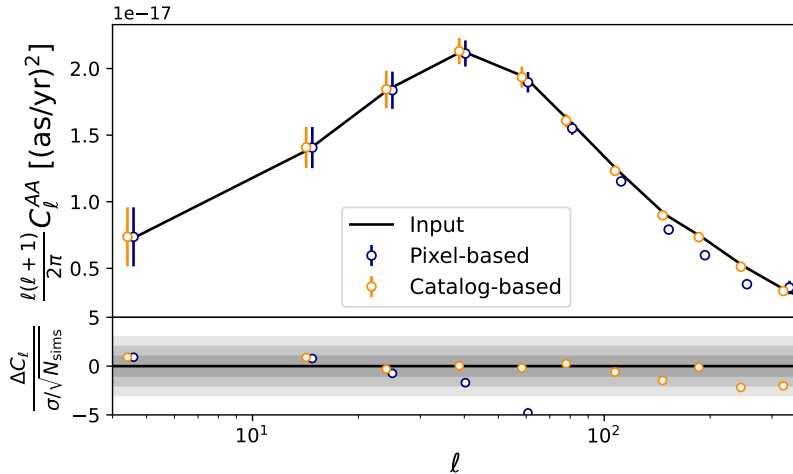


Figure 6. Angular power spectrum from 100 simulated *Quia*-like quasar proper motion catalogs, comparing the catalog-based (orange error bars) and the pixel-based approach (blue error bars, including a horizontal offset for clarity). The pixel-based result was obtained using a resolution parameter $N_{\text{side}} = 128$. The theoretical input is shown in black, and the relative difference with respect to it as a fraction of the statistical error on the mean of all 100 simulations is shown in the lower panel.

and given the relative sparsity of the sample, we will only attempt to reconstruct the power spectrum up to scales $\ell = 360$, and thus this resolution is sufficient. We generate 100 such simulated catalogs, estimate their power spectrum using the catalog-based method presented here and the standard pixel-based approach, using maps of resolution $N_{\text{side}} = 128$ in the latter case. This corresponds to pixels containing, on average, between one and two sources.

The result of this exercise is shown in figure 6. The upper panel shows the E -mode power spectrum bandpowers measured and averaged over 100 simulations, as well as the associated theoretical prediction, given by the input power spectrum convolved with the bandpower window functions. The lower panel shows the relative residuals between simulations and theory prediction in both cases. We find that, while the catalog-based approach is able to obtain an unbiased estimate of the power spectrum at all scales, the pixel-based method under-predicts it severely on small scales ($\ell \gtrsim 70$).

4.3 Spin-0 example: FRB-based dispersion measure maps

As an example of a scalar (spin-0) quantity sampled at discrete positions, we study the case of maps of the projected electron density distribution constructed from the dispersion measure of fast radio bursts (FRBs). FRBs are powerful millisecond-long radio pulses emitted by (most likely) extra-galactic sources of unknown nature. As they propagate through the interstellar and intergalactic plasma, the group velocity of these pulses changes with frequency. This leads to a time delay between pulses at different frequencies ν that scales as $\Delta t = t_0 \text{DM}/\nu^2$, where $t_0 \equiv 4.15 \times 10^{-3} \text{ s GHz}^2 \text{ cm}^3 \text{ pc}^{-1}$, and the dispersion measure DM is

$$\text{DM}(\hat{\mathbf{n}}) \equiv \int_0^{\chi_s} \frac{d\chi}{a} n_e(\chi \hat{\mathbf{n}}), \quad (4.5)$$




Figure 7. *Left:* spatial distribution of the FRB catalog used in this exercise, with the color of each source displaying its measured DM. *Right:* distribution of DM-based redshift estimates for the same sample, and fitting function used to generate the input angular power spectrum used in this study.

with χ_s the comoving radial distance to the source, a the scale factor, and $n_e(\mathbf{x})$ the comoving electron density. As in the case of cosmic shear and peculiar transverse velocities, one can thus make use of FRB catalogs to construct maps of the projected quantity of interest, in this case the electron density. For a catalog with a redshift distribution $p(z)$ (and assuming $c = 1$ as before), this is given by [48–50]

$$\text{DM}(\hat{\mathbf{n}}) = \int \frac{dz}{H(z)} (1+z) P(>z) n_e(\chi(z)\hat{\mathbf{n}}), \quad P(>z) \equiv \int_z^\infty dz' p(z'), \quad (4.6)$$

where $P(>z)$ is the cumulative redshift distribution.

Current FRB catalogs are significantly sparser than the transverse velocity and cosmic shear examples from previous sections, with existing FRB samples containing a few hundred sources (although this number will grow significantly in the future [51]). Another distinctive feature of FRB catalogs is the fact that DM measurements are highly accurate. Typically, the DM can be measured with a precision of $\sim 0.1 \text{ cm}^{-3} \text{ pc}$, while extragalactic DM values are orders of magnitude higher (e.g. $\text{DM}_{z \sim 0.8} \sim 1000 \text{ cm}^{-3} \text{ pc}$), and the contributions from the Milky Way are also comparatively small ($\text{DM}_{\text{MW}} \sim 50 \text{ cm}^{-3} \text{ pc}$). Typical host galaxy contributions are likely of the order $\text{DM}_{\text{host}} \sim 100 \text{ cm}^{-3} \text{ pc}$ (although the distribution of host DM may have significant tails [52]). Thus, in this case, the per-source measurements of the field are signal-dominated, unlike the previous examples.

To showcase the ability of the method presented here to recover accurate power spectrum measurements in this regime, we use the latest public catalog of FRBs made available by the CHIME collaboration [53]. After removing duplicates, the catalog contains 492 unique sources, with DM values in the range $\text{DM} \lesssim 2000$. The left panel of figure 7 shows the spatial distribution of this sample and their measured dispersion measures. We assign a nominal redshift to each source from their measured DM using the background relation between DM and redshift:¹⁸

$$\overline{\text{DM}}(z) - \overline{\text{DM}}_0 = \int_0^z \frac{dz'}{H(z')} (1+z') \bar{n}_e, \quad \bar{n}_e = \frac{3H_0^2 \Omega_b x_e (1+x_H)}{8\pi G m_p 2}, \quad (4.7)$$

¹⁸Estimating redshifts for FRBs without host identification is not possible. However, we can use the DM- z relation and the observed FRB dispersion measures to obtain a noisy estimate for an FRB redshift. For more details, see e.g. [54].

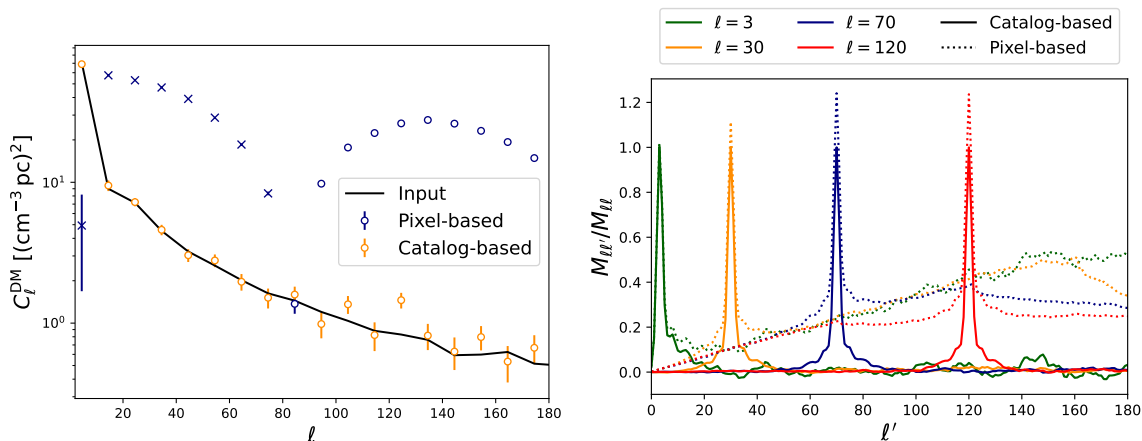


Figure 8. Results of the spin-0 FRB-based analysis. *Left:* power spectrum, estimated using the catalog-based approach, averaged over 1000 Gaussian realizations (orange points), with error bars corresponding to the standard error on the mean C_ℓ computed from simulations, and theoretical expectation (solid black line), corresponding to the input power spectrum convolved with the bandpower window functions. In blue, we show the result of the standard pixel-based approach, which fails catastrophically. *Right:* rows of the mode-coupling matrix $M_{\ell\ell'}$ for different values of ℓ , as a function of ℓ' . Results are shown for the catalog-based and pixel-based approaches (solid and dotted lines, respectively).

where m_p is the proton mass, H_0 is the expansion rate today, Ω_b is the fractional density of baryons (which we set to $\Omega_b = 0.05$), x_e is the ionization fraction (which we set to $x_e = 1$), and x_H is the hydrogen mass fraction (which we set to $x_H = 0.75$). We fix the offset $\overline{\text{DM}}_0 = 100 \text{ cm}^{-3} \text{ pc}$ to account for the typical host / Milky Way contribution (although our results do not depend on this choice). The corresponding redshift distribution of the sample is shown in the right panel of figure 7, and can be approximated as $p(z) \propto z^2 \exp(-z/0.15)$ (solid line in the same figure).

We use this redshift distribution to generate a theoretical prediction for the power spectrum of the DM field in eq. (4.6), assuming a linear bias for the electron overdensity of $b_e = 0.8$ [55]. We generate a suite of 1000 simulated Gaussian realizations of $\text{DM}(\hat{\mathbf{n}})$ at a HEALPix resolution of $N_{\text{side}} = 512$, which we then sample at the positions of the CHIME catalog sources. Given the sparsity of the sample, we only attempt to reconstruct the power spectrum up to scales $\ell = 180$, and hence this reduced resolution is sufficient. The left panel of figure 8 shows, in orange error bars, the reconstructed power spectrum averaged over all simulations and their standard error on the mean, computed across simulations, together with its theoretical prediction (i.e., the bandpower-convolved version of the power spectrum used to generate the simulated maps) in solid black lines. In blue error bars, we show also the result of using the standard, pixelized version of the estimator, using a resolution $N_{\text{side}} = 64$. In this case, the pixel size is much smaller than the typical distance between sources. The numerical instability of the mode-coupling matrix, and aliasing due to limited pixel resolution, causes the standard estimator to fail catastrophically in this case.

The right panel of the same figure shows rows of the normalized mode-coupling matrix $M_{\ell\ell'}$ for four different values of ℓ . Results are shown for the catalog-based and pixel-based

approaches as solid and dotted lines, respectively. In the catalog-based case, $M_{\ell\ell'}$ displays significant coupling between multipoles separated by up to $\Delta\ell \sim 20$, with a more non-trivial structure at larger multipole separations. In turn, the pixel-based approach, which does not correct for the effect of mask shot noise, shows significant correlation extending across the whole range of multipoles, growing like $(2\ell' + 1)$ (see eq. (2.20)). The mode-coupling matrix in this case is thus significantly more numerically unstable. Together with the impact of aliasing (see appendix B), this leads to the failure observed in the left panel.

4.4 Clustering example

We end this section showcasing the applicability of the discrete pseudo- C_ℓ approach to the case of source clustering, in which the density distribution of sources is itself the field under study. The formalism for this case was first presented in [5], and we described it within the general framework presented here in section 2.4. To validate our implementation of the formalism in NaMaster, we generate fast simulations of a galaxy clustering sample as follows:

- We calculate an initial angular power spectrum corresponding to the projected overdensity of sources with a linear galaxy bias $b = 1$ and a normal redshift distribution centered at $z = 0.6$ with a standard deviation $\sigma_z = 0.1$.
- We create a completeness map $c(\hat{\mathbf{n}})$ to simulate the observational survey inhomogeneities that the random catalogs described in section 2.4 aim to trace. To do this, we draw 200 random unit vectors on the sphere and select a spherical cap with a 10° radius centered at each of them. For each cap, we then draw a random number between 0.97 and 1 to represent the sky completeness in that patch (i.e., up to 3% incompleteness). This procedure is aimed to reproduce completeness variations between different exposures for a survey with relatively poor cadence (although the patch size used here is larger than the field of view of typical optical telescopes). We multiply the resulting map by a sky mask that removes regions of high dust extinction and star contamination, constructed as described in [56].
- We generate a Gaussian overdensity field, $\delta_G(\hat{\mathbf{n}})$ from this power spectrum, and perform a lognormal transformation to turn it into a positive-definite angular density field:

$$n_g^{\text{true}}(\hat{\mathbf{n}}) = \bar{n}_g \exp(\delta_G(\hat{\mathbf{n}}) - \sigma_G^2/2), \quad (4.8)$$

where σ_G is the standard deviation of δ_G , and \bar{n}_g is the average number of sources per steradian as observed by an “ideal” experiment free of survey inhomogeneities.¹⁹ The observed mean source density is then constructed by multiplying this field by the completeness map described above: $n_g^{\text{obs}}(\hat{\mathbf{n}}) = c(\hat{\mathbf{n}}) n_g^{\text{true}}(\hat{\mathbf{n}})$.

- We Poisson-sample this field, assigning each source a random position within its pixel, to generate the “data” catalog. The “random” catalog is generated by drawing random

¹⁹The $-\sigma_G^2/2$ counterterm ensures that the mean of the resulting field is indeed \bar{n}_g , assuming δ_G is Gaussian.

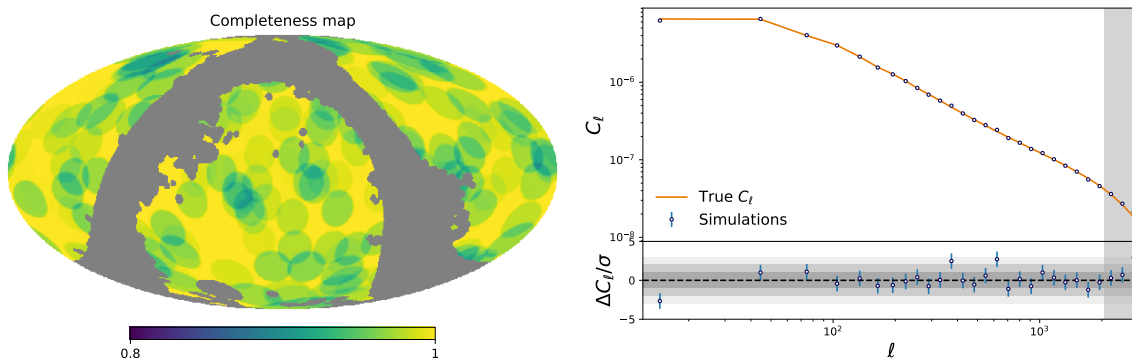


Figure 9. *Left:* completeness map used for the clustering study described in section 4.4. Fully masked pixels are shown in gray. *Right:* power spectrum averaged over 100 simulations (dark blue circles with error bars), true underlying spectrum (orange line), and residual relative to the statistical uncertainties in the mean (bottom panel).

positions on the sphere, and discarding them with a probability given by the local value of the completeness map. The random catalog thus constructed has 50 times more sources than the data.

Keeping the completeness map and associated random catalog fixed, we generate 100 data simulations at $N_{\text{side}} = 1024$ with a number density $\bar{n}_g = 0.01 \text{ arcmin}^{-2}$, and estimate the power spectrum of each simulation using the procedure described in section 2.4. The use of a lognormal transformation to generate the galaxy density field prevents us from using the initial angular power spectrum as a theoretical prediction to compare against. Although the effect of a lognormal transformation can be propagated to the power spectrum analytically (see [57]) we resort to a more direct approach that avoids potential numerical instabilities of Hankel transforms: for each simulation, we calculate the power spectrum of the respective realization of the lognormal density field ($\delta_{\text{LN}} = \exp(\delta_G - \sigma_G^2/2) - 1$) projected on a full-sky HEALPix map at $N_{\text{side}} = 1024$ (i.e., in the absence of completeness, masking, and discrete-source effects). The “true” power spectrum against which we compare the result of our simulations is then given by the average of this power spectrum over simulations, corrected for the impact of the HEALPix pixel window function used in these simulations. We do so because our findings in section 3 suggest that when sampling simulated catalogs from ground-truth HEALPix maps, our catalog-based estimator is unbiased with respect to the pixel-window-corrected theory.

The left panel of figure 9 shows the completeness map used in this exercise. The right panel of the same figure shows the mean power spectrum recovered from our simulation (blue points with error bars), the true C_ℓ (orange line), and the relative residual between them (bottom subpanel). The catalog method implemented here is able to recover the true underlying power spectrum at very high accuracy (the error bars shown represent the statistical uncertainties in the mean, corresponding to 1/10 of the error of a single realization).

5 Conclusion

In this work, we present a simple extension to the standard pseudo- C_ℓ estimator to tackle the problem of fields sampled at the positions of a discrete and finite catalog of sources. By treating the problem directly at the catalog level, without resorting to an intermediate pixelization of the sources onto a regular grid, we demonstrate that the method is able to tackle three prevalent sources of systematic bias in the standard pixel-based approach: resolution and aliasing effects due to the finite pixel area, shot noise in the field estimate due to the discrete nature of the field tracers, and numerically unstable mode-coupling matrices due to shot noise in the pixelized mask constructed from the sources. We interpret the catalog data as representing a masked field, with a mask given by a set of delta functions at the position of each source, and take advantage of new methods developed for the calculation of spherical harmonic transforms on irregular grids. We show that a simple correction of the shot-noise contributions to the power spectra of the masked field and of the mask itself, using fast and model-independent analytical expressions, is able to significantly improve the numerical stability of the mode-coupling matrix, while simultaneously rendering the estimator immune to any uncorrelated noise bias, irrespective of the size of the signal contribution. The method is a generalization of the technique presented in [5] in the context of galaxy clustering. The most evident use of this method is in the analysis of cosmic shear data, although it is broadly applicable to any datasets described by observed properties of discrete sources.

We validate our implementation of this method on a set 500 simulated KiDS-1000-like cosmic shear catalogs, with sources sampled at their real sky positions. We show that the resulting estimator is unbiased even for unrealistically small samples (100 and 10,000 times less dense than the real KiDS-1000 sample), where mode-coupling instabilities due to shot noise render standard pixel-based methods unusable. Even at realistic sample sizes, including the larger densities achieved by deeper samples, we show that pixel-based methods encounter difficulties to accurately characterize the effects of pixelization on the small-scale power spectrum. While these can be circumvented through the use of higher pixel resolutions, this is at a significantly higher cost in computational time and complexity, which the catalog-based method presented here avoids altogether.

To demonstrate the practical use of our estimator on other suitable astrophysical datasets, with varying sample sizes and signal-to-noise ratios, we apply it in four additional distinct scenarios. First, we demonstrate its feasibility in the analysis of significantly larger catalogs, reproducing the analysis of the 3-year DES cosmic shear sample, comparing the results to those presented in [33]. We then study the case of a spin-1 field: the peculiar transverse motion of extragalactic objects as traced by an astrometric survey, using the *Quaia* quasar sample. Besides the different spin, this type of dataset features a significantly lower source density and signal-to-noise ratio than cosmic shear. We also consider measurements of the spin-0 projected electron density in the form of Fast Radio Burst catalogs. These correspond to samples with extremely low number densities (10^2 - 10^3 objects) but a very high signal-to-noise ratio. Finally, we showcase the applicability of our formalism to the analysis of galaxy clustering, in which the density of catalog sources is itself the field of interest. Our results from all four examples confirm our previous findings that our estimator is unbiased and can overcome the numerical challenge of source sparseness when characterizing mode coupling.

We expect the method presented to be useful in the analysis of upcoming large-scale structure datasets, particularly in the context of cosmic shear, as well as novel probes such as FRBs, radial and transverse peculiar motion surveys, weak lensing of gravitational waves, and any other catalog-based observations. Our implementation is made available with the latest release of the public power spectrum package `NaMaster`.²⁰

Acknowledgments

When the work presented here was at an advanced stage, we learnt of similar work being developed by Tessore et al. 2024 [27]. We thank the authors of this work, particularly Nicolas Tessore and Benjamin Joachimi, for sharing their draft paper and for useful discussions. We also thank Anton Baleato, Elyas Farah, Robert Reischke, Martin White, and the anonymous referee for useful comments and feedback. KW acknowledges support from the Science and Technology Facilities Council (STFC) under grant ST/X006344/1. DA acknowledges support from the Beecroft trust. We made extensive use of computational resources at the University of Oxford Department of Physics, funded by the John Fell Oxford University Press Research Fund.

A Generalities of spherical harmonic functions

We use the following definition for the generalized spin- s spherical harmonic functions:²¹

$${}_s Y_{\ell m}^{1E} = -\frac{1}{2} [{}_s Y_{\ell m} + (-1)^s {}_{-s} Y_{\ell m}], \quad {}_s Y_{\ell m}^{1B} = -\frac{i}{2} [{}_s Y_{\ell m} - (-1)^s {}_{-s} Y_{\ell m}], \quad (\text{A.1})$$

$${}_s Y_{\ell m}^{2E} = \frac{i}{2} [{}_s Y_{\ell m} + (-1)^s {}_{-s} Y_{\ell m}], \quad {}_s Y_{\ell m}^{1B} = -\frac{1}{2} [{}_s Y_{\ell m} - (-1)^s {}_{-s} Y_{\ell m}], \quad (\text{A.2})$$

where ${}_s Y_{\ell m}(\hat{\mathbf{n}})$ are the spin-weighted spherical harmonic functions of spin s . These functions satisfy the following properties:

$$\int d\hat{\mathbf{n}} {}_s Y_{\ell m}^{a\alpha*}(\hat{\mathbf{n}}) {}_s Y_{\ell' m'}^{a\beta}(\hat{\mathbf{n}}) = \delta_{\alpha\beta} \delta_{\ell\ell'} \delta_{mm'}, \quad (\text{A.3})$$

$$\int d\hat{\mathbf{n}} {}_s Y_{\ell m}^{a\alpha*}(\hat{\mathbf{n}}) {}_s Y_{\ell' m'}^{a\beta}(\hat{\mathbf{n}}) Y_{\ell'' m''}(\hat{\mathbf{n}}) = (-1)^{s+m} \sqrt{\frac{(2\ell+1)(2\ell'+1)(2\ell''+1)}{4\pi}} \times \begin{pmatrix} \ell & \ell' & \ell'' \\ -m & m' & m'' \end{pmatrix} \begin{pmatrix} \ell & \ell' & \ell'' \\ s & -s & 0 \end{pmatrix} d_{\ell+\ell'+\ell''}^{\alpha\beta}, \quad (\text{A.4})$$

$$d_n^{EE} = d_n^{BB} = [1 + (-1)^n]/2, \quad d_n^{EB} = -d_n^{BE} = -i[1 - (-1)^n]/2. \quad (\text{A.5})$$

$$\sum_m {}_s Y_{\ell m}^{a\alpha*}(\hat{\mathbf{n}}) {}_{s'} Y_{\ell m}^{a\beta}(\hat{\mathbf{n}}) = \frac{2\ell+1}{4\pi} \delta_{\alpha\beta} \delta_{ss'}. \quad (\text{A.6})$$

The last relation can be derived from the addition theorem for spin-weighted spherical harmonics [58]. For scalar harmonics, it reads simply

$$\sum_m Y_{\ell m}^*(\hat{\mathbf{n}}_1) Y_{\ell m}(\hat{\mathbf{n}}_2) = \frac{2\ell+1}{4\pi} P_\ell(\hat{\mathbf{n}}_1 \cdot \hat{\mathbf{n}}_2), \quad (\text{A.7})$$

²⁰<https://github.com/LSSTDESC/NaMaster/>.

²¹The -1 sign preceding these definitions is a mostly inconvenient convention that is dropped for the case of spin-0 fields, so that ${}_0 Y_{\ell m}^{1E} = {}_0 Y_{\ell m}^{2B} = Y_{\ell m}$, and ${}_0 Y_{\ell m}^{2E} = {}_0 Y_{\ell m}^{1B} = 0$.

which, for $\hat{\mathbf{n}}_1 = \hat{\mathbf{n}}_2$, simply evaluates to $(2\ell + 1)/(4\pi)$. Furthermore, the spherical harmonic functions are a complete basis on the sphere, so that:

$$\sum_{\ell m} Y_{\ell m}^*(\hat{\mathbf{n}}_1) Y_{\ell m}(\hat{\mathbf{n}}_2) = \delta^D(\hat{\mathbf{n}}_1, \hat{\mathbf{n}}_2). \quad (\text{A.8})$$

We will also employ the following properties of the Wigner $3j$ symbols:

$$\sum_{mm'} \begin{pmatrix} \ell & \ell' & \ell'' \\ m & m' & m'' \end{pmatrix} \begin{pmatrix} \ell & \ell' & \ell''' \\ m & m' & m''' \end{pmatrix} = \frac{\delta_{\ell'' \ell'''} \delta_{m'' m'''}}{2\ell'' + 1}, \quad (\text{A.9})$$

$$\sum_{\ell''} (2\ell'' + 1) \begin{pmatrix} \ell & \ell' & \ell'' \\ s & -s & 0 \end{pmatrix} \begin{pmatrix} \ell & \ell' & \ell'' \\ s' & -s' & 0 \end{pmatrix} = \delta_{ss'}, \quad (\text{A.10})$$

$$\begin{pmatrix} \ell & \ell' & \ell'' \\ -m & -m' & -m'' \end{pmatrix} = (-1)^{\ell + \ell' + \ell''} \begin{pmatrix} \ell & \ell' & \ell'' \\ m & m' & m'' \end{pmatrix}. \quad (\text{A.11})$$

B Pixelization effects in catalog-based observables

The main text has emphasized the impact of the numerical instability of the mode-coupling matrix due to the presence of shot noise in the mask power spectrum for catalog-based observables that are first pixelized before estimating their angular power spectrum. This appendix explores the two other sources of potential bias that affect pixelized catalog-based data: the impact of pixel window functions, and aliasing.

For this exercise we generate two types of simulated realizations of a scalar Gaussian random field using a base pixelization of $N_{\text{side}} = 1024$:

- **Band-limited realizations** with a power spectrum of the form

$$C_\ell = \frac{\exp[-(\ell\Theta)^2]}{\ell + 50}. \quad (\text{B.1})$$

The exponential cutoff, with $\Theta = 0.3^\circ$ sharply reduces any structure on scales $\ell \gtrsim 200$.

- **Non-band-limited realizations** with a power spectrum of the form

$$C_\ell = \frac{1}{\ell + 50}, \quad (\text{B.2})$$

which is rather flat, decaying significantly more slowly at high ℓ than the previous case.

First, to isolate the effects of pixel averaging and aliasing from the practicalities of fields sampled at arbitrary discrete positions, we construct maps at lower resolution, $N_{\text{side}} = 256$ and $N_{\text{side}} = 64$ by simply averaging the $N_{\text{side}} = 1024$ maps onto larger pixels. We calculate the angular power spectrum of each of these maps, and estimate the effective pixel window function of the measured power spectrum as:

$$b_\ell^{\text{eff}} \equiv \sqrt{\frac{\hat{C}_\ell}{C_\ell^{\text{in}}}}, \quad (\text{B.3})$$

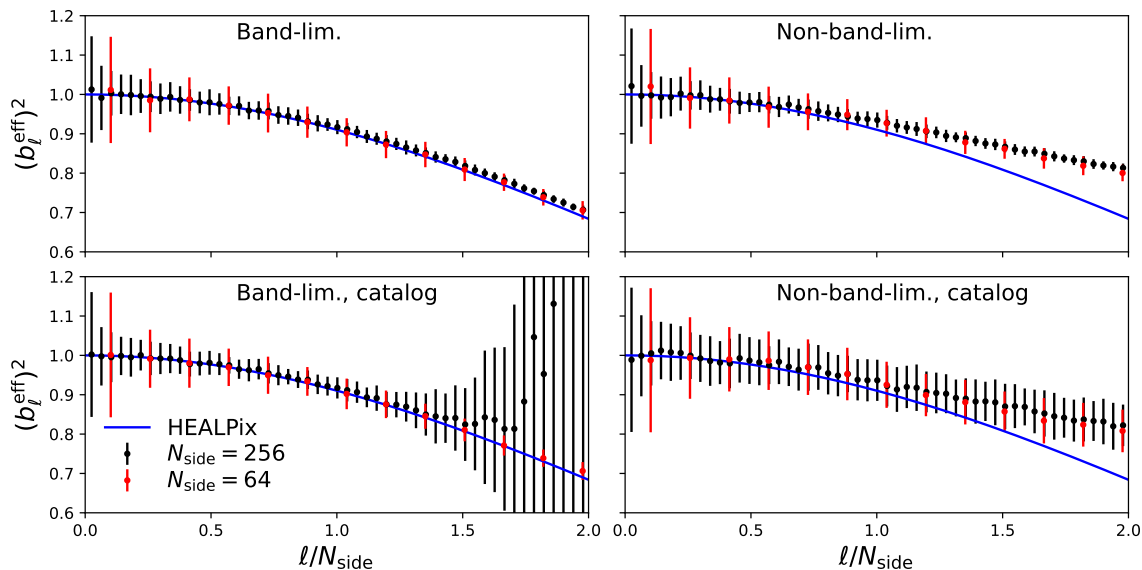


Figure 10. Effective pixel window function for band-limited simulations (left panels) and non-band-limited simulations (right panels). Results for medium- and low-resolution pixels ($N_{\text{side}} = 256$ and 64 , respectively) are shown as black and red points with error bars, compared against the HEALPix window function (blue). The top panels show results for simulations in which the base maps at $N_{\text{side}} = 1024$ are averaged down to the lower resolutions. The bottom panels show results for catalogs constructed from the base maps and then pixelized at lower resolution. Besides the numerical instability of the mode-coupling matrix for catalog-based measurements, which is factored out in this exercise, the most relevant residual resolution effect relevant for map-based analyses of catalog datasets is aliasing.

where \hat{C}_ℓ is the estimated power spectrum, and C_ℓ^{in} is the input power spectrum (given by eq. (B.1) or eq. (B.2)). We repeat this for 100 simulations and compute the mean and scatter of the effective pixel window function.

The top left and top right panels of figure 10 shows the result of this exercise for the band-limited and non-band-limited realizations, respectively. We find that, for the band-limited case, the impact of pixel averaging is captured by the HEALPix pixel window function with relatively good accuracy.²² In the case of non-band-limited simulations, however, we see very significant differences in the effective pixel window function. These differences are caused by aliasing of the small-scale power in the base maps onto the larger pixels via averaging. Interestingly, the resulting effective pixel window function has a similar shape both in the $N_{\text{side}} = 256$ and the $N_{\text{side}} = 64$ case.

We now repeat the exercise with a key modification: we sample the base maps at the discrete positions of 10^6 sources drawn at random on the sphere, and we then construct maps of the resulting field using the standard pixel-based approach at $N_{\text{side}} = 256$ and $N_{\text{side}} = 64$ (corresponding to ~ 1.3 and ~ 21 sources per pixel on average, respectively). When calculating the power spectra of these maps we use the density of said sources as their corresponding masks. Finally, to factor out the effects caused by shot noise in the

²²The slight difference for the $N_{\text{side}} = 256$ maps is due to the fact that averaging over only 4×4 sub-pixels is less representative of a true area averaging than the 16×16 pixels for $N_{\text{side}} = 64$.

mask, we subtract the analytical shot noise contribution to the mask and field pseudo- C_ℓ (\tilde{N}^w and \tilde{N}^f), mimicking the procedure described in the main text to avoid these effects. The bottom left and right panels of figure 10 show the results of this exercise for the band-limited and non-band-limited maps, respectively. The conclusions are analogous to those found in the previous exercise: regardless of the average number of sources per pixel, the most relevant effect is aliasing in the case of non-band-limited data. The effects of pixelization on band-limited data are well captured by the HEALPix pixel window function, suggesting that aliasing-related bias in the standard pseudo- C_ℓ approach can be reduced by band-limiting the signal (e.g. by smoothing the input map with an artificial beam window function).

References

- [1] A. Albrecht et al., *Report of the Dark Energy Task Force*, [astro-ph/0609591](#) [INSPIRE].
- [2] D.H. Weinberg et al., *Observational Probes of Cosmic Acceleration*, *Phys. Rept.* **530** (2013) 87 [[arXiv:1201.2434](#)] [INSPIRE].
- [3] LSST collaboration, *Tomographic galaxy clustering with the Subaru Hyper Suprime-Cam first year public data release*, *JCAP* **03** (2020) 044 [[arXiv:1912.08209](#)] [INSPIRE].
- [4] M. White et al., *Cosmological constraints from the tomographic cross-correlation of DESI Luminous Red Galaxies and Planck CMB lensing*, *JCAP* **02** (2022) 007 [[arXiv:2111.09898](#)] [INSPIRE].
- [5] A. Baleato Lizancos and M. White, *Harmonic analysis of discrete tracers of large-scale structure*, *JCAP* **05** (2024) 010 [[arXiv:2312.12285](#)] [INSPIRE].
- [6] G. Bertin and M. Lombardi, *Looking at the Fundamental Plane through a gravitational lens*, *Astrophys. J. Lett.* **648** (2006) L17 [[astro-ph/0606672](#)] [INSPIRE].
- [7] H. Johnston et al., *Intrinsic correlations of galaxy sizes in a hydrodynamical cosmological simulation*, *Mon. Not. Roy. Astron. Soc.* **520** (2023) 1541 [[arXiv:2209.11063](#)] [INSPIRE].
- [8] C.T. Mpetha, G. Congedo and A. Taylor, *Future prospects on testing extensions to Λ CDM through the weak lensing of gravitational waves*, *Phys. Rev. D* **107** (2023) 103518 [[arXiv:2208.05959](#)] [INSPIRE].
- [9] I. Duncan, D. Alonso, A. Slosar and K. Storey-Fisher, *Can we constrain structure growth from galaxy proper motions?*, *Open J. Astrophys.* **7** (2024) 9 [[arXiv:2305.15893](#)] [INSPIRE].
- [10] M.S. Madhavacheril, N. Battaglia, K.M. Smith and J.L. Sievers, *Cosmology with the kinematic Sunyaev-Zeldovich effect: Breaking the optical depth degeneracy with fast radio bursts*, *Phys. Rev. D* **100** (2019) 103532 [[arXiv:1901.02418](#)] [INSPIRE].
- [11] D. Alonso, *Linear anisotropies in dispersion-measure-based cosmological observables*, *Phys. Rev. D* **103** (2021) 123544 [[arXiv:2103.14016](#)] [INSPIRE].
- [12] A. Nicola et al., *Breaking baryon-cosmology degeneracy with the electron density power spectrum*, *JCAP* **04** (2022) 046 [[arXiv:2201.04142](#)] [INSPIRE].
- [13] R. Reischke and S. Hagstotz, *Consistent constraints on the equivalence principle from localized fast radio bursts*, *Mon. Not. Roy. Astron. Soc.* **523** (2023) 6264 [[arXiv:2302.10072](#)] [INSPIRE].
- [14] M. Tegmark, *How to measure CMB power spectra without losing information*, *Phys. Rev. D* **55** (1997) 5895 [[astro-ph/9611174](#)] [INSPIRE].

- [15] J.R. Bond, A.H. Jaffe and L. Knox, *Estimating the power spectrum of the cosmic microwave background*, *Phys. Rev. D* **57** (1998) 2117 [[astro-ph/9708203](#)] [[INSPIRE](#)].
- [16] J. Borrill, *Power spectrum estimators for large CMB datasets*, *Phys. Rev. D* **59** (1999) 027302 [[astro-ph/9712121](#)] [[INSPIRE](#)].
- [17] G. Efstathiou, *Myths and truths concerning estimation of power spectra: the case for a hybrid estimator*, *Mon. Not. Roy. Astron. Soc.* **349** (2004) 603 [[astro-ph/0307515](#)] [[INSPIRE](#)].
- [18] A.J.S. Hamilton, *Power Spectrum Estimation I. Basics*, *Lect. Notes Phys.* **665** (2008) 415 [[astro-ph/0503603](#)] [[INSPIRE](#)].
- [19] A.J.S. Hamilton, *Power Spectrum Estimation II. Linear Maximum Likelihood*, *Lect. Notes Phys.* **665** (2008) 433 [[astro-ph/0503604](#)] [[INSPIRE](#)].
- [20] M. Tegmark and A. de Oliveira-Costa, *How to measure CMB polarization power spectra without losing information*, *Phys. Rev. D* **64** (2001) 063001 [[astro-ph/0012120](#)] [[INSPIRE](#)].
- [21] E. Hivon et al., *Master of the cosmic microwave background anisotropy power spectrum: a fast method for statistical analysis of large and complex cosmic microwave background data sets*, *Astrophys. J.* **567** (2002) 2 [[astro-ph/0105302](#)] [[INSPIRE](#)].
- [22] M.L. Brown, P.G. Castro and A.N. Taylor, *Cosmic microwave background temperature and polarization pseudo- C_l estimators and covariances*, *Mon. Not. Roy. Astron. Soc.* **360** (2005) 1262 [[astro-ph/0410394](#)] [[INSPIRE](#)].
- [23] LSST DARK ENERGY SCIENCE collaboration, *A unified pseudo- C_l framework*, *Mon. Not. Roy. Astron. Soc.* **484** (2019) 4127 [[arXiv:1809.09603](#)] [[INSPIRE](#)].
- [24] F. Elsner, B. Leistedt and H.V. Peiris, *Unbiased pseudo- Cl power spectrum estimation with mode projection*, *Mon. Not. Roy. Astron. Soc.* **465** (2017) 1847 [[arXiv:1609.03577](#)] [[INSPIRE](#)].
- [25] A. Nicola et al., *Cosmic shear power spectra in practice*, *JCAP* **03** (2021) 067 [[arXiv:2010.09717](#)] [[INSPIRE](#)].
- [26] G. Chon et al., *Fast estimation of polarization power spectra using correlation functions*, *Mon. Not. Roy. Astron. Soc.* **350** (2004) 914 [[astro-ph/0303414](#)] [[INSPIRE](#)].
- [27] EUCLID collaboration, *Euclid preparation. LIX. Angular power spectra from discrete observations*, [arXiv:2408.16903](#) [[INSPIRE](#)].
- [28] M. Reinecke and D.S. Seljebotn, *Libsharp — spherical harmonic transforms revisited*, *Astron. Astrophys.* **554** (2013) A112 [[arXiv:1303.4945](#)].
- [29] M. Reinecke, S. Belkner and J. Carron, *Improved cosmic microwave background (de-)lensing using general spherical harmonic transforms*, *Astron. Astrophys.* **678** (2023) A165 [[arXiv:2304.10431](#)] [[INSPIRE](#)].
- [30] S. Belkner et al., *cuNuSHT: GPU Accelerated Spherical Harmonic Transforms on Arbitrary Pixelizations*, *RAS Tech. Instrum.* **3** (2024) 711 [[arXiv:2406.14542](#)] [[INSPIRE](#)].
- [31] DES collaboration, *Dark energy survey year 3 results: cosmological constraints from the analysis of cosmic shear in harmonic space*, *Mon. Not. Roy. Astron. Soc.* **515** (2022) 1942 [[arXiv:2203.07128](#)] [[INSPIRE](#)].
- [32] R. Dalal et al., *Hyper Suprime-Cam Year 3 results: Cosmology from cosmic shear power spectra*, *Phys. Rev. D* **108** (2023) 123519 [[arXiv:2304.00701](#)] [[INSPIRE](#)].
- [33] C. García-García et al., *Cosmic shear with small scales: DES-Y3, KiDS-1000 and HSC-DR1*, *JCAP* **08** (2024) 024 [[arXiv:2403.13794](#)] [[INSPIRE](#)].

- [34] DES collaboration, *Dark energy survey year 3 results: weak lensing shape catalogue*, *Mon. Not. Roy. Astron. Soc.* **504** (2021) 4312 [[arXiv:2011.03408](#)] [[INSPIRE](#)].
- [35] KiDS collaboration, *KiDS-1000 Cosmology: Cosmic shear constraints and comparison between two point statistics*, *Astron. Astrophys.* **645** (2021) A104 [[arXiv:2007.15633](#)] [[INSPIRE](#)].
- [36] C. Hikage, M. Takada, T. Hamana and D. Spergel, *Shear power spectrum reconstruction using the pseudo-spectrum method*, *Mon. Not. Roy. Astron. Soc.* **412** (2011) 65 [[arXiv:1004.3542](#)] [[INSPIRE](#)].
- [37] K.M. Górski et al., *HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere*, *Astrophys. J.* **622** (2005) 759 [[astro-ph/0409513](#)] [[INSPIRE](#)].
- [38] B. Giblin et al., *KiDS-1000 catalogue: Weak gravitational lensing shear measurements*, *Astron. Astrophys.* **645** (2021) A105 [[arXiv:2007.01845](#)] [[INSPIRE](#)].
- [39] PLANCK collaboration, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6 [*Erratum ibid.* **652** (2021) C4] [[arXiv:1807.06209](#)] [[INSPIRE](#)].
- [40] DES collaboration, *Dark Energy Survey Year 3 results: Cosmology from cosmic shear and robustness to data calibration*, *Phys. Rev. D* **105** (2022) 023514 [[arXiv:2105.13543](#)] [[INSPIRE](#)].
- [41] D.N. Limber, *The Analysis of Counts of the Extragalactic Nebulae in Terms of a Fluctuating Density Field*, *Astrophys. J.* **117** (1953) 134.
- [42] M. Kilbinger et al., *Precision calculations of the cosmic shear power spectrum projection*, *Mon. Not. Roy. Astron. Soc.* **472** (2017) 2126 [[arXiv:1702.05301](#)] [[INSPIRE](#)].
- [43] S. Mishra-Sharma, *Inferring dark matter substructure with astrometric lensing beyond the power spectrum*, *Mach. Learn. Sci. Tech.* **3** (2022) 01LT03 [[arXiv:2110.01620](#)] [[INSPIRE](#)].
- [44] J. Darling, *Objects Appear Smaller as They Recede: How Proper Motions Can Directly Reveal the Cosmic Expansion, Provide Geometric Distances, and Measure the Hubble Constant*, *Astrophys. J. Lett.* **777** (2013) L21 [[arXiv:1310.0500](#)] [[INSPIRE](#)].
- [45] J. Darling and A.E. Truebenbach, *All Transverse Motion is Peculiar: Connecting the Proper Motions of Galaxies to the Matter Power Spectrum*, *Astrophys. J.* **864** (2018) 37 [[arXiv:1807.06658](#)] [[INSPIRE](#)].
- [46] A. Hall, *Cosmology with extragalactic proper motions: harmonic formalism, estimators, and forecasts*, *Mon. Not. Roy. Astron. Soc.* **486** (2019) 145 [[arXiv:1811.05454](#)] [[INSPIRE](#)].
- [47] K. Storey-Fisher et al., *Quaia, the Gaia-unWISE Quasar Catalog: An All-sky Spectroscopic Quasar Sample*, *Astrophys. J.* **964** (2024) 69 [[arXiv:2306.17749](#)] [[INSPIRE](#)].
- [48] K. Ioka, *The Cosmic Dispersion Measure from Gamma-Ray Burst Afterglows: Probing the Reionization History and the Burst Environment*, *Astrophys. J. Lett.* **598** (2003) L79 [[astro-ph/0309200](#)] [[INSPIRE](#)].
- [49] S. Inoue, *Probing the cosmic reionization history and local environment of gamma-ray bursts through radio dispersion*, *Mon. Not. Roy. Astron. Soc.* **348** (2004) 999 [[astro-ph/0309364](#)] [[INSPIRE](#)].
- [50] W. Deng and B. Zhang, *Cosmological Implications of Fast Radio Burst/Gamma-Ray Burst Associations*, *Astrophys. J. Lett.* **783** (2014) L35 [[arXiv:1401.0059](#)] [[INSPIRE](#)].
- [51] E. Petroff, J.W.T. Hessels and D.R. Lorimer, *Fast radio bursts at the dawn of the 2020s*, *Astron. Astrophys. Rev.* **30** (2022) 2 [[arXiv:2107.10113](#)] [[INSPIRE](#)].

- [52] A. Theis, S. Hagstotz, R. Reischke and J. Weller, *Galaxy dispersion measured by Fast Radio Bursts as a probe of baryonic feedback models*, [arXiv:2403.08611](#) [INSPIRE].
- [53] CHIME/FRB collaboration, *The First CHIME/FRB Fast Radio Burst Catalog*, *Astrophys. J. Supp.* **257** (2021) 59 [[arXiv:2106.04352](#)] [INSPIRE].
- [54] R. Reischke, S. Hagstotz and R. Lilow, *Probing primordial non-Gaussianity with Fast Radio Bursts*, *Phys. Rev. D* **103** (2021) 023517 [[arXiv:2007.04054](#)] [INSPIRE].
- [55] K.M. Smith et al., *KSZ tomography and the bispectrum*, [arXiv:1810.13423](#) [INSPIRE].
- [56] N. Koukoufilippas, D. Alonso, M. Bilicki and J.A. Peacock, *Tomographic measurement of the intergalactic gas pressure through galaxy-tSZ cross-correlations*, *Mon. Not. Roy. Astron. Soc.* **491** (2020) 5464 [[arXiv:1909.09102](#)] [INSPIRE].
- [57] P. Coles and B. Jones, *A lognormal model for the cosmological mass distribution*, *Mon. Not. Roy. Astron. Soc.* **248** (1991) 1 [INSPIRE].
- [58] K. Seibert, *Spin-Weighted Spherical Harmonics and Their Application for the Construction of Tensor Slepian Functions on the Spherical Cap*, Ph.D. thesis, Universität Siegen, Germany (2018).