

Systematics mitigation for catalogue-based angular power spectra

Thomas Cornish^{1,2}  ^{1,2}★ David Alonso,¹ Boris Leistedt² and Kevin Wolz¹

¹*Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK*

²*Department of Physics, Imperial College London, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, UK*

Accepted 2026 February 18. Received 2026 February 18; in original form 2025 November 20

ABSTRACT

Recent work has developed a formalism for computing angular power spectra directly from catalogues containing field values at discrete positions on the sky, thereby circumventing the need to create pixelized maps of the fields, as well as avoiding aliasing and finite-resolution effects. We adapt this formalism to incorporate template deprojection for mitigating systematic biases in the measured angular power spectra. We also introduce an alternative method of mitigating the ‘deprojection bias’ – the loss of modes induced by deprojection – employing simple simulations to compute a transfer function. We find that this approach performs at least as well as existing methods, and is relatively insensitive to how well one can guess the true power spectrum of the observed field, except at the largest scales ($\ell \lesssim 3$). Additionally, we develop exact expressions for the bias introduced by deprojection in the shot-noise component, which further improves the accuracy of this approach. We test our formalism on simulated data sets, demonstrating its applicability both to discretely sampled fields, and to the special case of galaxy clustering, with the survey selection function defined in terms of a random catalogue or as a continuous sky map. After removing the bias in the shot noise and correcting for the remaining mode loss using a transfer function, our formalism produces unbiased measurements of the angular power spectrum in all scenarios tested here. Finally, we apply our formalism to real data and show it produces results consistent with the standard map-based pseudo- C_ℓ formalism. We implement our method in the public code `NaMaster`.

Key words: methods: data analysis – methods: numerical.

1 INTRODUCTION

The accurate measurement of projected fields on the sky has become a cornerstone of modern cosmology, with perhaps the most famous example being the cosmic microwave background (CMB), for which measurements of anisotropies in its temperature and polarization enable constraints on models of inflation and the primordial density perturbations, as well as on the natures of neutrinos and dark energy (e.g. E. Camphuis et al. 2025; T. Louis et al. 2025; Planck Collaboration VI 2020). Our knowledge of the large-scale structure (LSS) in the Universe is also heavily dependent on our ability to measure projected fields on the sky: tomographic measurements of galaxy clustering exploit the fact that galaxies are biased tracers of the underlying dark matter distribution while leveraging photometric redshifts as a means of estimating their radial positions, thereby enabling indirect measurements of the matter density field in ‘slices’ across cosmic time without the need for spectroscopic observations, which are time-consuming and costly across large areas (e.g. A. Nicola et al. 2020; T. M. C. Abbott et al. 2022; M. Rodríguez-Monroy et al. 2022; S. Sugiyama et al. 2023; M. Vakili et al. 2023; Z. Yan et al. 2025). Furthermore, the effects of weak gravitational lensing can be expressed in terms of projected fields, such as

the lensing shear and convergence, both of which are sensitive probes of the intervening dark matter distribution (C. Heymans et al. 2021; T. M. C. Abbott et al. 2022; H. Miyatake et al. 2023; L. Faga et al. 2025; B. Stözlner et al. 2025). These are only two examples among myriad other projected probes of the LSS, which include other CMB secondary anisotropies, the Cosmic Infrared Background, etc. (e.g. R. Stompor & G. Efstathiou 1999; R. Adam et al. 2017; S. Bocquet et al. 2019; A. Maniyar et al. 2019; Í. Zúbel-dia & A. Challinor 2019; L. E. Bleem et al. 2022; J. Carron, M. Mirmelstein & A. Lewis 2022; B. Jogo et al. 2023; W. Coulton et al. 2024; X. Fang et al. 2024; F. J. Qu et al. 2024; S. Shaikh et al. 2024; F. Ge et al. 2025).

It is no surprise, then, that the angular power spectrum (denoted C_ℓ), which quantifies the covariance between angular fluctuations in two fields as a function of angular multipole ℓ (and encodes all information in the case of Gaussian fields), has become an essential summary statistic in cosmology. Unbiased estimators of the angular power spectrum have been extensively explored, and optimal quadratic minimum variance (QMV) estimators have been developed which are capable of recovering angular power spectra with minimal loss of information (M. Tegmark 1997). However, these estimators require the inversion of a pixel-space covariance matrix in order to optimally weight the data, an operation which can be computationally expensive when working with high-resolution maps. The prohibitive computational demand of these estimators often necessitates an al-

* E-mail: thomas.cornish@physics.ox.ac.uk

ternative, an example of which can be found in the pseudo- C_ℓ (PCL) algorithm (E. Hivon et al. 2002; D. Alonso et al. 2019). The PCL approach simplifies things by assuming a diagonal data covariance matrix – an assumption which holds if the data in question are spatially uncorrelated and/or dominated by white noise (B. Leistedt et al. 2013). This simplification significantly improves the speed of computation, making the PCL algorithm an often attractive approach to C_ℓ estimation.

The characterization and mitigation of systematic effects is becoming a problem of increasing relevance in cosmological surveys: as the hoards of data being collected continually increase in size, the statistical uncertainties have the potential to become comparable to those arising from systematics (e.g. S. Ho et al. 2012, 2015; D. Huterer, C. E. Cunha & W. Fang 2013; A. R. Pullen & C. M. Hirata 2013; N. Agarwal et al. 2014; D. L. Shafer & D. Huterer 2015; N. Weaverdyck, J. Muir & D. Huterer 2018). A plethora of techniques have been developed to tackle this issue by attempting to quantify the dependence of the desired field on maps of the survey properties or other contaminants thought to be contributing to the observed signal. These range from techniques interpretable through a linear regression framework (N. Weaverdyck & D. Huterer 2021) such as mode deprojection (also called template deprojection; G. B. Rybicki & W. H. Press 1992; A. Slosar, U. Seljak & A. Makarov 2004; F. Elsner, B. Leistedt & H. V. Peiris 2016, 2017), template subtraction (A. J. Ross et al. 2011; S. Ho et al. 2012), and the iterative systematics decontamination technique employed by the Dark Energy Survey (J. Elvin-Poole et al. 2018; M. Rodríguez-Monroy et al. 2022), to more sophisticated methods employing machine learning tools such as artificial neural networks (M. Rezaie et al. 2020) or self-organizing maps (H. Johnston et al. 2021).

The approaches discussed so far typically involve data in the form of pixelized maps that quantify the value of a projected field at a given position on the sky. However, it is often the case in cosmology that the fields whose angular power spectra we would like to measure are sampled at discrete positions on the sky, corresponding to the positions of sources in a catalogue. The typical course of action to obtain power spectrum measurements in this case is to bin these discrete points into pixels and produce a map of the mean field value in each pixel. However, recent work by A. Baleato Lizancos & M. White (2024); K. Wolz, D. Alonso & A. Nicola (2025) allows one to circumvent this step, through an update to the PCL formalism which facilitates the computation of the PCL directly from catalogue data. This approach naturally avoids any bias introduced by the process of pixelizing discretely sampled observables, which would otherwise influence measurements of the power spectrum at scales comparable to or smaller than the pixel scale. This extension to the PCL formalism was implemented in the publicly available code `NaMaster`¹ (D. Alonso et al. 2019), and its validity subsequently demonstrated on various discretely-sampled astrophysical fields, including cosmic shear, proper motions, dispersion measure of fast radio bursts, and galaxy positions. However, these validation tests were performed using idealized mock data with no contamination due to systematics; indeed, the catalogue-based formalism as it is described in K. Wolz et al. (2025) includes no means of mitigating systematic effects.

The aim of this paper is thus to extend the catalogue-based PCL formalism of K. Wolz et al. (2025) in order to incorporate

systematics mitigation, and to implement this in `NaMaster`. In particular, we adapt the template deprojection approach for applicability to catalogue data, as this is already implemented in `NaMaster` for map-based fields.

The structure of the paper is as follows. In Section 2, we outline the existing PCL formalism and its extension to catalogue-based data, before introducing our integration of template deprojection into the latter. In Section 3, we validate our approach using simulated data with systematic effects factored in. Section 4 contains our conclusions.

2 METHODS

2.1 Overview of pseudo- C_ℓ methods

This section contains a brief review of the pseudo- C_ℓ power spectrum estimator, including the concept of template deprojection, and the techniques used to extract power spectra directly from catalogues. This provides the background needed to present the extension of deprojection techniques to catalogue-based spectra described in Section 2.2. Readers are referred to E. Hivon et al. (2002) and D. Alonso et al. (2019) for a more detailed account of these techniques.

We will focus our discussion on the calculation of auto-correlations of a given field, since many of the new results presented in this paper are most relevant in this scenario, and all others can be trivially generalized to the case of cross-correlations.

2.1.1 Standard pseudo- C_ℓ estimator

The pseudo- C_ℓ (PCL) or MASTER algorithm for power spectrum estimation (E. Hivon et al. 2002; D. Alonso et al. 2019) is a simplified version of the optimal QMV estimator. The optimal inverse-variance weighting of the data map, prescribed by the QMV estimator, is a computationally expensive operation, effectively requiring the inversion of a large pixel-space covariance matrix. In the PCL approach, this matrix is assumed to be diagonal, turning the inverse variance weighting step into a simple multiplication of two maps: the data map containing the field measurements, and the ‘mask’, which plays the role of the local inverse variance of these measurements. Not only is this a significantly simpler and faster operation, but the statistical coupling between different angular multipoles induced by this weighting can be estimated analytically using efficient methods. This makes the PCL estimator easy to deploy on a large variety of data sets to extract fast power spectrum measurements. This is particularly beneficial on small scales, where the QMV estimator becomes prohibitively expensive.

The steps involved in the PCL estimator can be summarized as follows (see D. Alonso et al. (2019) for a more complete description of the algorithm):

(i) Let $\mathbf{a}(\hat{\mathbf{n}})$ be a spin- s field defined on the sphere, and let $v(\hat{\mathbf{n}})$ be its associated mask. Construct the masked field, given by the product $\mathbf{a}^v(\hat{\mathbf{n}}) \equiv v(\hat{\mathbf{n}}) \mathbf{a}(\hat{\mathbf{n}})$. Note that we denote spin- s quantities as a vector $\mathbf{a}(\hat{\mathbf{n}}) \equiv (a_1(\hat{\mathbf{n}}), a_2(\hat{\mathbf{n}}))$, where a_1 and a_2 are the two components of the field (e.g. the Q and U Stokes parameters for CMB polarization maps, or the γ_1 and γ_2 components of the cosmic shear field).

(ii) Compute the spherical harmonic transform (SHT) of the masked field, $\mathbf{a}'_{\ell m}$, ignoring the presence of a mask.

¹<https://github.com/LSSTDESC/NaMaster>

(iii) Calculate the pseudo- C_ℓ of \mathbf{a}^v , $\text{PCL}_\ell(\mathbf{a}^v, \mathbf{a}^v)$, defined as the naïve estimator of the field’s power spectrum, averaging over all values of m for a given ℓ , ignoring the fact that \mathbf{a}^v is in fact a masked field. In general, for two masked maps \mathbf{a}^v and \mathbf{b}^w , the PCL^2 is defined as

$$\text{PCL}_\ell(\mathbf{a}^v, \mathbf{b}^w) \equiv \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} \mathbf{a}_{\ell m}^v \mathbf{b}_{\ell m}^{w\dagger}. \quad (1)$$

Since the presence of a sky mask leads to the statistical coupling of different angular multipoles, the PCL is also often called the ‘coupled’ C_ℓ .

(iv) Calculate the ‘mode-coupling matrix’ (MCM), the matrix encoding the statistical coupling between different ℓ modes induced by the sky mask, using the analytical expressions presented in e.g. D. Alonso et al. (2019). The MCM depends solely on the PCL of the mask. If binning the power spectrum into bandpowers, calculate also the binned MCM. Invert the MCM and multiply the PCL calculated in the previous step by it.

(v) In most cases, the measured field receives both signal and noise contributions: $\mathbf{a} = \mathbf{s}_a + \mathbf{n}_a$, and more often than not we are interested in the power spectrum of the signal. Thus, an estimate of the noise power spectrum must be constructed and subtracted from the measured power spectrum. In certain cases (e.g. white noise) this may be done analytically. In other cases one must rely on noise simulations. Finally, noise bias may be avoided using only cross-correlations between different map ‘splits’ (e.g. maps constructed from data taken during disjoint time periods), containing the same signal but different noise realizations (T. Louis et al. 2025). The specific case of uncorrelated noise from discretely-sampled sources is discussed in Section 2.1.3.

(vi) The power spectrum must be corrected for finite-resolution effects. These may be caused by a finite telescope beam, or by the finite size of the sky pixels over which the map has been measured (assuming that the pixel values correspond to an average of the field value over the pixel area). This can be done exactly if a sufficiently accurate model of the smoothing kernel can be constructed (e.g. R. Datta et al. 2024).

It is worth noting that, under the usual interpretation of the QMV estimator, the mask $v(\hat{\mathbf{n}})$ in step 1 above should in principle correspond to the local inverse variance of the measured field.³ However, the estimator remains unbiased regardless of what $v(\hat{\mathbf{n}})$ is (as long as it is zero where the field is not measured), although its variance may become highly sub-optimal if $v(\hat{\mathbf{n}})$ deviates significantly from this prescription.

2.1.2 Template deprojection

Often the map under study is affected by contamination from non-cosmological sources. Although the level of contamination is usually not known a priori, it is often possible to construct a comprehensive library of templates representing the potential sources of contamination. If the level of contamination is relatively mild, we can then construct a linear model of the form

$$\mathbf{a}^{\text{obs}} = \mathbf{a}^v(\hat{\mathbf{n}}) + A_p \mathbf{f}^p(\hat{\mathbf{n}}), \quad (2)$$

²Note that we use ‘PCL’ as shorthand for both the ‘pseudo- C_ℓ ’ formalism and for the cut-sky estimator of the power spectrum defined here.

³We assume scalar masks, although this can be generalized to the case of fully anisotropic weighting—see D. Alonso (2025).

where \mathbf{a}^{obs} is the observed, contaminated field, \mathbf{a}^v is the masked and uncontaminated field, $\mathbf{f}^p(\hat{\mathbf{n}})$ are the different contaminant template maps, and A_p are unknown amplitudes that control the level of contamination from each source. Note that we implicitly sum over the repeated index p (and we will continue to do so throughout the paper).

A self-consistent procedure to account for this contamination is then to calculate the power spectrum of \mathbf{a}^{obs} after marginalizing over all possible values of A_p (G. B. Rybicki & W. H. Press 1992). As shown in F. Elsner, B. Leistedt & H. V. Peiris (2017), in the context of the PCL framework, this can be achieved by calculating the best-fitting value of A_p , and subtracting the resulting contamination from the original map. Mathematically, this is also equivalent to projecting the observed map onto the subspace that is orthogonal to all the contaminant templates, effectively ‘deprojecting’ these templates from the data. The resulting deprojected map is thus

$$\mathbf{a}^c(\hat{\mathbf{n}}) = \mathbf{a}^v(\hat{\mathbf{n}}) - \mathbf{f}^p(\hat{\mathbf{n}}) F_{pq} \langle \mathbf{f}^q, \mathbf{a}^v \rangle, \quad (3)$$

where the matrix \mathbf{F} is defined as the pseudo-inverse⁴ of \mathbf{F}^{-1} , where

$$(\mathbf{F}^{-1})_{pq} \equiv \langle \mathbf{f}^p, \mathbf{f}^q \rangle. \quad (4)$$

In these equations, we have defined the map-level dot product as

$$\langle \mathbf{a}, \mathbf{b} \rangle \equiv \int d\hat{\mathbf{n}} \mathbf{a}^\dagger(\hat{\mathbf{n}}) \mathbf{b}(\hat{\mathbf{n}}). \quad (5)$$

Expanding the fields in terms of their harmonic coefficients, this dot product may also be written in harmonic space as

$$\langle \mathbf{a}, \mathbf{b} \rangle = \sum_{\ell m} \mathbf{a}_{\ell m}^\dagger \mathbf{b}_{\ell m} = \sum_{\ell} (2\ell + 1) \text{Tr} [\text{PCL}_\ell(\mathbf{a}, \mathbf{b})]. \quad (6)$$

Since this will be important for our discussion in Section 2.2, note that the combination

$$\hat{A}_p \equiv F_{pq} \langle \mathbf{f}^q, \mathbf{a}^v \rangle \quad (7)$$

is the best-fitting value of the amplitude coefficient A_p in equation (2) estimated from the data. Equation (3) may thus be expressed directly in harmonic space:

$$\mathbf{a}_{\ell m}^c = \mathbf{a}_{\ell m}^v - \hat{A}_p \mathbf{f}_{\ell m}^p, \quad (8)$$

with \hat{A}_p calculated in real or harmonic space. Finally, note that all the template maps \mathbf{f}^p entering the equations above are implicitly masked using $v(\hat{\mathbf{n}})$ (the mask applied to the data).

Since contaminant deprojection effectively removes modes from the map (any mode that lies in the subspace of contaminant templates), it causes a bias in the estimated power spectrum that must be corrected for (F. Elsner et al. 2017). We will explore alternative approaches to correct for this bias in Section 2.3.

As first introduced in section 2.4.2 of D. Alonso et al. (2019), in the case of uncorrelated noise, it is possible to calculate the correction to the noise power spectrum caused by contaminant deprojection exactly, even if the noise is inhomogeneous. The additional contribution to the noise PCL is:

$$\begin{aligned} \Delta N_\ell = & -2F_{pq} \text{PCL}_\ell(\mathbf{f}^p, v^2 \sigma_N^2 \mathbf{f}^q) \\ & + F_{pq} F_{rs} \left[\int d\hat{\mathbf{n}} v^2 \sigma_N^2 \mathbf{f}^{q\dagger} \mathbf{f}^r \right] \text{PCL}_\ell(\mathbf{f}^p, \mathbf{f}^s), \end{aligned} \quad (9)$$

⁴The pseudo-inverse is used here to account for potential internal linear dependences in the library of contaminant templates.

where $\sigma_N^2(\hat{\mathbf{n}})$ is a map of the local field variance in a unit solid angle.

2.1.3 Catalogue-based pseudo- C_ℓ s

We are often interested in fields that are not mapped over continuous areas, but instead are recorded at the positions of discrete sources. A prime example of this is cosmic shear, where the lensing shear field is effectively measured at the positions of background sources. The power spectrum of such fields can still be calculated using the standard PCL approach, generating continuous maps by binning the sources into pixels, and using the density of these sources as a proxy for the sky mask. Some subtleties about this approach can be found in A. Nicola et al. (2021).

Instead, a more accurate treatment of discretely-sampled fields has been introduced recently (A. Baleato Lizancos & M. White 2024; K. Wolz et al. 2025), enabled by new algorithms to estimate the SHTs of irregularly-sampled fields (M. Reinecke, S. Belkner & J. Carron 2023). In this case, the field and its mask can be expressed as a sum of delta functions centred at the source coordinates:

$$v(\hat{\mathbf{n}}) = \sum_{i=1}^{N_d} w_i \delta^D(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i), \quad a^v(\hat{\mathbf{n}}) = \sum_{i=1}^{N_d} w_i a_i \delta^D(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i), \quad (10)$$

where the index i runs over all sources in the sample, w_i is the weight of the i th source, $\hat{\mathbf{n}}_i$ is its position on the sky, a_i is the value of field measured at that position, and δ^D is the Dirac delta function on the sphere. For simplicity we consider only a scalar field here, and the generalization to spin- s can be found in K. Wolz et al. (2025). The spherical harmonic coefficients of v and a^v are then simply

$$v_{\ell m} = \sum_{i=1}^{N_d} w_i Y_{\ell m, i}^*, \quad a_{\ell m}^v = \sum_{i=1}^{N_d} w_i a_i Y_{\ell m, i}^*, \quad (11)$$

which can be calculated using the irregular SHT methods implemented in e.g. M. Reinecke et al. (2023). Once computed, these harmonic coefficients can be used to calculate the PCL of the field and the associated MCM, as described in Section 2.1.1.

The most immediately obvious advantage of the catalogue-based approach is the absence of finite pixel resolution effects, allowing the resulting estimator to recover power spectra up to arbitrarily small scales (as long as the corresponding harmonic coefficients can be calculated and stored). In addition to this, K. Wolz et al. (2025) showed that the catalogue-based approach naturally allows for a more numerically stable treatment of uncorrelated noise bias. In particular, it was shown that all contributions from any uncorrelated noise sources can be removed exactly by subtracting the zero-lag contributions to the PCLs of the masked field and of the mask itself prior to estimating the MCM (i.e. the contributions to the PCLs coming from the product of v and a^v with themselves evaluated at the same source). These are given by

$$\tilde{N}^v \equiv \frac{1}{4\pi} \sum_{i=1}^{N_d} w_i^2, \quad \tilde{N}^a \equiv \frac{1}{4\pi} \sum_{i=1}^{N_d} w_i^2 a_i^2. \quad (12)$$

Doing so has the added advantage that, by removing the asymptotic value of the mask PCL at high ℓ (given by \tilde{N}^v), the associated MCM is significantly more numerically stable, and can be safely inverted.

Another application of the catalogue-based approach is the case of galaxy clustering, in which the density of sources is itself the field of interest. In the simplest scenario, a continuous map with the expected mean density of sources $\bar{n}(\hat{\mathbf{n}})$ can be constructed. The overdensity field is then given by

$$a(\hat{\mathbf{n}}) \equiv \delta_g(\hat{\mathbf{n}}) = \frac{1}{\bar{n}(\hat{\mathbf{n}})} \left[\sum_{i=1}^{N_d} w_i \delta^D(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i) - \bar{n}(\hat{\mathbf{n}}) \right]. \quad (13)$$

Assuming the mean density is itself used as a mask for this field (since $\bar{n}(\hat{\mathbf{n}})$ is indeed proportional to the inverse variance of δ_g in the Poisson limit), the masked field is simply

$$a^v(\hat{\mathbf{n}}) = \sum_{i=1}^{N_d} w_i \delta^D(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i) - \bar{n}(\hat{\mathbf{n}}). \quad (14)$$

Since in this case the mask is continuous, its shot noise component is simply $\tilde{N}^v = 0$, while the field noise is $\tilde{N}^a = \sum_{i=1}^{N_d} w_i^2 / (4\pi)$.

Often, however, the mean source density is not provided as a continuous map, but as a catalogue of ‘random’ objects (or simply ‘randoms’). Catalogues of randoms have long been a critical component of configuration-space galaxy clustering analyses (e.g. H. Totsuji & T. Kihara 1969; S. D. Landy & A. S. Szalay 1993). These are point processes distributed according to the survey selection function and without any intrinsic clustering. They are typically generated by uniformly sampling across the survey footprint to create a catalogue of positions with a density significantly higher than that of the data, such that correlations between the randoms and observed galaxies can be used to mitigate systematics caused by masking or edge effects. It is worth noting that recent efforts have also been made to construct catalogues of ‘organized randoms’ which additionally account for spatial correlations induced by more complex survey properties such as inhomogeneous depth or seeing (H. Johnston et al. 2021; Z. Yan et al. 2025). However, since the aim of this work is to employ template deprojection to mitigate these more complex effects, we adopt the former approach to generating randoms when applying our formalism to simulations (Section 3.2) and observations (Section 3.5).

If randoms are used in place of a continuous map, the mean density map (and hence the mask) entering the equations above is

$$\bar{n}(\hat{\mathbf{n}}) = \alpha \sum_{i=1}^{N_r} w_i^r \delta^D(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i), \quad (15)$$

where w_i^r is the weight of the i th random object, and $\alpha \equiv \sum_{i=1}^{N_d} w_i / \sum_{i=1}^{N_r} w_i^r$ corrects for the different number of sources in the data and random catalogues. To simplify the notation, in what follows, we will omit the factor α , and absorb it into the random weights.⁵ In this case the mask has a shot noise component, and the randoms also contribute to the field noise:

$$\tilde{N}^w = \frac{1}{4\pi} \sum_{i=1}^{N_r} (w_i^r)^2, \quad \tilde{N}^a = \frac{1}{4\pi} \sum_{i=1}^{N_d} w_i^2 + \tilde{N}^w. \quad (16)$$

In summary, the treatment of catalogue-based fields is straightforward to extend to the case of galaxy clustering, by simply replacing all field values with 1, and including the contribution

⁵In other words, we will implicitly assume that the random weights are normalized such that $\sum_{i=1}^{N_r} w_i^r = \sum_{i=1}^{N_d} w_i$.

from random objects in the calculation of the field spherical harmonic coefficients and in the shot-noise components of field and mask (if randoms are used).

2.2 Template deprojection for catalogue-based C_ℓ s

We now present the extension of the template deprojection technique, outlined in Section 2.1.2 for continuous fields, to the case of fields sampled at the positions of discrete sources. Specifically, our aim is to extend equations (8) and (9), describing the harmonic coefficients of the contaminant-deprojected field and the modification to the white noise level brought about by deprojection, respectively. We will do so for both sampled fields and for galaxy clustering.

2.2.1 Sampled fields

Consider first a discretely-sampled field $a^p(\hat{\mathbf{n}})$ as given in equation (10). The masked contaminant templates in this case are

$$\tilde{f}^p(\hat{\mathbf{n}}) = \sum_i w_i f_i^p \delta^D(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i), \quad (17)$$

where f_i^p is the value of the template at the position of the i th source in the catalogue. Substituting this in equation (3), we find that the deprojected field is simply

$$a^c(\hat{\mathbf{n}}) = \sum_i w_i a_i^c \delta^D(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i), \quad (18)$$

where we have defined the deprojected sampled field values a_i^c as

$$a_i^c \equiv a_i - \hat{A}_p f_i^p, \quad (19)$$

and the estimated linear amplitude is

$$\hat{A}_p \equiv \tilde{F}_{pq} \sum_j w_j^2 f_j^q a_j, \quad (\tilde{\mathbf{F}}^{-1})_{pq} \equiv \sum_j w_j^2 f_j^p f_j^q. \quad (20)$$

Thus, we see that deprojection simply affects the values of the field sampled at the source positions, and it only depends on the values of the contaminant templates sampled at the same positions. The dot product introduced in equation (5) is replaced by a sum over all sources of the product of two weighted fields. Deprojection can therefore be achieved easily, without requiring any contaminant maps explicitly, and simply modifying the sampled field values as above. Likewise, the white noise amplitude, accounting for deprojection, can be calculated simply using equation (12), replacing a_i with a_i^c as given above.

The white-noise power spectrum receives a further scale-dependent correction from deprojection due to mode loss. In the case of continuous fields, this is given by equation (9). To calculate this correction in the case of a discretely-sampled field, we start by writing the PCL of a^c :

$$\text{PCL}_\ell(a^c, a^c) = \sum_{ij} (w_i w_j a_i^c a_j^c) \sum_{m=-\ell}^{\ell} \frac{|Y_{\ell m i}^* Y_{\ell m j}|}{2\ell + 1}. \quad (21)$$

We then expand a_i^c in terms of a_i and the templates as in equation (19), and then take the ensemble average of the result, assuming a_i to consist solely of uncorrelated noise: $\langle a_i a_j \rangle = \sigma_i^2 \delta_{ij}$, where σ_i^2 is the noise variance of the field at the i th discrete source. Writing the result as

$$\langle \text{PCL}_\ell(a^c, a^c) \rangle = \tilde{N}^a + \Delta N_\ell^a, \quad (22)$$

the correction is

$$\Delta N_\ell^a = -2\tilde{F}_{pq} \text{PCL}_\ell(\tilde{f}^p, \tilde{g}^q) + \tilde{F}_{pq} \tilde{F}_{rs} \left[\sum_j w_j^2 \sigma_j^2 (w_j f_j^q) (w_j f_j^s) \right] \text{PCL}_\ell(\tilde{f}^p, \tilde{f}^r), \quad (23)$$

where

$$\tilde{f}_{\ell m}^p \equiv \sum_i w_i f_i^p Y_{\ell m i}^*, \quad (24)$$

and where we have defined the noise-weighted field

$$\tilde{g}_{\ell m}^p \equiv \sum_i w_i^2 \sigma_i^2 (w_i f_i^p) Y_{\ell m i}^*. \quad (25)$$

As before, this contribution to the noise bias due to deprojection depends only on the values of the contaminant templates evaluated at the source positions, as well as an estimate of the noise variance at each source σ_i^2 . For noise-dominated fields, this may be estimated directly as $\sigma_i^2 \sim a_i^2$, but care must be exercised when this is not the case. For modest numbers of contaminant templates, this white noise correction is relatively small, and may be neglected. However, as we will see in Section 3.1, when many ($O(100)$) templates are deprojected, the correction may become significant, particularly on large angular scales.

2.2.2 Galaxy clustering

We turn now to the case of galaxy clustering. We will start with the more complex scenario in which the mask is provided as a discrete set of random points, and then turn to the simpler case of a continuous mask at the end of this section.

Two characteristics of this type of data separate it from the standard sampled fields in a way that impacts how contaminant deprojection must be carried out: first, the source positions and weights do not constitute the field's mask, the random catalogue does instead. Secondly, this discretely-sampled mask features as an additive contribution to the masked field itself, quantifying the mean density subtracted from the data (the second term in equation 14). Note that although this latter choice is indeed the most common one when analysing galaxy data, one could also use two different random catalogues – one to construct the subtracted mean density, and a different one to define the mask – and the resulting estimator for the galaxy overdensity would still be unbiased.⁶ Nevertheless, we will adopt the standard approach of using a single random catalogue for both purposes here.

This immediately leads to a problem if we simply apply equation (3) to the field defined in equation (14), with \bar{n} given by equation (15) and, hence, with the masked templates given by⁷

$$\tilde{f}^p(\hat{\mathbf{n}}) = \sum_{i=1}^{N_r} w_i^r f_i^p \delta^D(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i), \quad \tilde{f}_{\ell m}^p = \sum_{i=1}^{N_r} w_i^r f_i^p Y_{\ell m i}^*.$$

⁶Assuming that both random catalogues represent the same underlying mean density field.

⁷Recall that we assume the random weights to be normalized so that $\alpha = 1$.

Substituting this in the expression for the estimated linear coefficients \hat{A}_p (equation 7) we find:

$$\begin{aligned}\hat{A}_p &= F_{pq} \int d\hat{\mathbf{n}} a^v(\hat{\mathbf{n}}) \tilde{f}^q(\hat{\mathbf{n}}) \\ &= F_{pq} \left(\sum_{i=1}^{N_d} \sum_{j=1}^{N_r} w_i w_j^r f_j^q \int d\hat{\mathbf{n}} \delta^D(\hat{\mathbf{n}}_i, \hat{\mathbf{n}}) \delta^D(\hat{\mathbf{n}}_j, \hat{\mathbf{n}}) \right. \\ &\quad \left. - \sum_{i,j=1}^{N_r} w_i^r w_j^r f_j^q \int d\hat{\mathbf{n}} \delta^D(\hat{\mathbf{n}}_i, \hat{\mathbf{n}}) \delta^D(\hat{\mathbf{n}}_j, \hat{\mathbf{n}}) \right) \\ &\propto -F_{pq} \sum_{i=1}^{N_r} (w_i^r)^2 f_i^p,\end{aligned}$$

since we assume that no two objects from either the source or random catalogues occupy the exact same position. Likewise, the elements of the matrix \mathbf{F} satisfy

$$(\mathbf{F}^{-1})_{pq} \propto \sum_{i=1}^{N_r} (w_i^r)^2 f_i^p f_i^q. \quad (26)$$

This is problematic because, first, the estimated \hat{A}_p does not depend on the data catalogue, only on the randoms, and hence it does not capture any real correlations between the galaxy overdensity and the contaminants. Secondly, had we chosen to use two different catalogues for the subtracted mean density and the mask, the quantities above would both be exactly zero.

The cause of this problem lies in the key assumption underlying the pseudo- C_ℓ power spectrum estimator: the covariance matrices of all fields involved are diagonal in real space, and the inverse diagonals are represented by the field masks (in this case, the random catalogue positions). We know that, in practice, this is not the case: LSS fields exhibit correlations over a range of real-space separations. Furthermore, we actually expect the observed galaxy overdensity to be affected by systematic contamination (and hence to correlate with the contaminant templates) on large scales. We thus propose to use these large-scale correlations to redefine the linear deprojection coefficients. In particular, we can use the harmonic-space expression for the dot product of two fields (equation 6), truncating it at a given maximum multipole ℓ_{\max} , which characterizes the scale over which we want to capture correlations between data and contaminants. The linear deprojection coefficients, in the case of galaxy clustering, are therefore calculated via

$$\begin{aligned}\hat{A}_p &\equiv \tilde{F}_{pq} \sum_{\ell=0}^{\ell_{\max}} (2\ell+1) \left[\text{PCL}_\ell(a^v, \tilde{f}^q) + \sum_{i=1}^{N_r} \frac{(w_i^r)^2 f_i^q}{4\pi} \right], \\ (\mathbf{F}^{-1})_{pq} &= \sum_{\ell=0}^{\ell_{\max}} (2\ell+1) \left[\text{PCL}_\ell(\tilde{f}^p, \tilde{f}^q) - \sum_{i=1}^{N_r} \frac{(w_i^r)^2 f_i^p f_i^q}{4\pi} \right]. \quad (27)\end{aligned}$$

Note that we have subtracted the shot-noise contribution to the pseudo- C_ℓ s entering these equations, caused by coincident pairs of random points in the subtracted mean density and the mask. This contribution is completely accidental (e.g. it would be zero if we had decided to use two different random catalogues for both purposes), and must therefore be discarded.

This generalization of the standard contaminant deprojection approach is perfectly legitimate from a mathematical point of view: it is equivalent to a standard deprojection operation in harmonic space in which some data points (multipoles with $\ell, |m| > \ell_{\max}$) are assigned zero weight (or, equivalently, a very large variance).

Having defined the deprojection procedure, we must now derive the correction to the shot noise power spectrum caused by

deprojection. We can do so following the same steps outlined in the previous section for sampled fields, with two caveats:

(i) The noise variance in this case is sourced by Poisson noise in the data and random catalogues. Furthermore, if different random catalogues are used to quantify the subtracted mean density and to define the mask, only the randoms in the former (i.e. those entering $a_{\ell m}^v$) contribute to this variance, and not those in the latter (i.e. those entering $\tilde{f}_{\ell m}$).

(ii) We must take care to consistently apply the truncation scale ℓ_{\max} used to find the deprojection coefficients.

The resulting expression for the shot noise deprojection bias is

$$\begin{aligned}\Delta N_\ell &= -2\tilde{F}_{pq} \text{PCL}_\ell(\tilde{f}^p, \tilde{g}^q) \\ &\quad + \tilde{F}_{pq} \tilde{F}_{rs} \left[\sum_{i=1}^{N_d} w_i^2 \tilde{f}_{Fi}^q \tilde{f}_{Fi}^s \right. \\ &\quad \left. + \sum_{i=1}^{N_r} (w_i^r)^2 \tilde{f}_{Fi}^q \tilde{f}_{Fi}^s \right] \text{PCL}_\ell(\tilde{f}^p, \tilde{f}^q), \quad (28)\end{aligned}$$

where we have defined the filtered-weighted templates \tilde{f}_F^p as a truncated inverse SHT of the masked templates

$$\tilde{f}_{Fi}^p \equiv \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} \tilde{f}_{\ell m}^p Y_{\ell m, i}, \quad (29)$$

and where

$$\tilde{g}_{\ell m}^p \equiv \sum_{i=1}^{N_d} w_i^2 \tilde{f}_{Fi}^p Y_{\ell m, i}^* + \sum_{i=1}^{N_r} (w_i^r)^2 \tilde{f}_{Fi}^p Y_{\ell m, i}^*. \quad (30)$$

Note that, as in the case of sampled fields, we only require the values of the contaminant templates at the discrete positions of the points defining the field's mask (i.e. the random catalogue), and not explicitly in the form of maps.

Now, in the simpler case in which the mask (and hence the mean number density) is provided as a continuous map, deprojection can be carried out following the steps outlined above, with the following simplifications:

(i) The templates are also provided as continuous maps, and the masked templates are simply $\tilde{f}^p(\hat{\mathbf{n}}) = \bar{n}(\hat{\mathbf{n}}) f^p(\hat{\mathbf{n}})$, where $\bar{n}(\hat{\mathbf{n}})$ is the mean number density map, which also serves as our mask.

(ii) No shot-noise contributions must be subtracted when estimating the deprojection coefficients (i.e. the rightmost terms in both expressions in equation 27).

(iii) Only the data weights contribute to the shot noise variance in equations (28) and (30).

2.3 Correcting for mode loss

The process of template deprojection indiscriminately removes all power from modes proportional to any of the deprojected templates, introducing a bias in the final power spectrum estimate. As discussed in F. Elsner et al. (2017), this *deprojection bias* can be estimated analytically and subsequently removed: for a scalar field a with corresponding mask v and contaminants f^p , it is given by

$$\begin{aligned}\Delta C_\ell &= -2F_{pq} \text{PCL}_\ell(\tilde{f}_{\ell m}^q, f_{\ell m}^p) \\ &\quad + F_{pq} F_{rs} \left[\int d\hat{\mathbf{n}} f^{q\dagger}(\hat{\mathbf{n}}) \tilde{f}^r(\hat{\mathbf{n}}) \right] \text{PCL}_\ell(f^p, f^s), \quad (31)\end{aligned}$$

where F_{pq} is as defined via equation (4), and

$$\tilde{f}_{\ell m}^p \equiv \mathcal{S} [v(\hat{\mathbf{n}})\mathcal{S}^{-1} [C_{\ell'}^{\text{true}} \mathcal{S} [v f_{\ell' m'}^p]_{\hat{\mathbf{n}}}^{s_a}]_{\ell m}^{s_a}]. \quad (32)$$

with \mathcal{S} denoting a SHT and \mathcal{S}^{-1} its inverse.⁸ Note that we are again only considering autocorrelations of a single scalar field here; a generalization to cross-correlations between two fields with arbitrary spin can be found in D. Alonso et al. (2019).

This method of correcting for the loss of modes induced by deprojection has two drawbacks. First, it requires some knowledge of the true angular power spectrum C_ℓ^{true} , for which a reliable estimate is not always achievable. Secondly, while the above prescription works for map-based fields, an equivalent procedure for catalogue-based data is non-trivial to implement, owing to the inverse SHTs involved.

We therefore explore in this work an alternative approach, in which we use simulations to estimate a *transfer function* by which one can divide the measured power spectra to correct for mode loss following template deprojection. The process involves randomly generating many Gaussian realizations of some angular power spectrum C_ℓ^{input} , and applying to each realization the same mask as is being used for the data in question. For each realization, the angular power spectrum is then estimated both before and after deprojecting the set of contaminants believed to be influencing the data; we will denote these power spectra as C_ℓ^{before} and C_ℓ^{after} , respectively. The transfer function T_ℓ is then defined as

$$T_\ell = \frac{\langle C_\ell^{\text{after}} \rangle_{\text{sim}}}{\langle C_\ell^{\text{before}} \rangle_{\text{sim}}}, \quad (33)$$

where $\langle \cdot \rangle_{\text{sim}}$ denotes the average across Gaussian realizations. Dividing the power spectrum measured from the data by T_ℓ should then correct it for mode loss (see e.g. F. Elsner et al. 2016, in which this approach is used to validate analytical expressions for the deprojection bias).

In Section 3.3, we test this approach and verify it as a viable alternative to calculating the deprojection bias, which can in turn be applied to power spectra measured from catalogue-based data.

3 RESULTS

We now validate the catalogue-based adaptation of template deprojection described above using both simulated and real data. Sections 3.1 and 3.2 do this for the cases of sampled fields and galaxy clustering, respectively. Next, in Section 3.3 we examine the relative performances of the mode loss correction methods described in Section 2.3. Finally in Section 3.5 we apply this approach to real data from the *Gaia-unWISE* quasar catalogue (Quaia; K. Storey-Fisher et al. 2024).

3.1 Contaminant deprojection in sampled fields

We start by validating the formalism described in the previous section in the case of sampled fields. To do so, we generate a number of simulated catalogue-based observations including the presence of sky contamination, and show that we are able to recover unbiased power spectra by applying deprojection techniques and accounting for the impact of deprojection bias in both the signal and uncorrelated noise components.

⁸Defined as $S[a(\hat{\mathbf{n}})]_{\ell m} \equiv \int d\hat{\mathbf{n}} Y_{\ell m}^*(\hat{\mathbf{n}}) a(\hat{\mathbf{n}})$ and $S^{-1}[a_{\ell m}]_{\hat{\mathbf{n}}} \equiv \sum_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}) a_{\ell m}$ for a scalar field $a(\hat{\mathbf{n}})$.




Figure 1. Input power spectrum used to validate the approach to contaminant deprojection in the case of sampled fields (black line). The blue line shows the power spectrum of the contaminant maps used, while the horizontal dashed line shows the simulated white noise component. The circular points show the mean power spectrum measured from all contaminated simulations.

Specifically, in our simulations we use the positions of sources in the *Quaia* catalogue with spectro-photometric redshifts in the range $z_p < 1.47$, corresponding to a sample of 644 786 quasar candidates. This ensures that the simulated data have a realistic source distribution, including large-scale sky coverage and small-scale clustering. Using a HEALPIX⁹ (K. M. Górski et al. 2005; A. Zonca et al. 2019) pixelization scheme with a resolution parameter of $N_{\text{side}} = 256$ (corresponding to a pixel area of $\sim 0.05 \text{ deg}^2$), we then generate signal maps as Gaussian realizations with an angular power spectrum of the form

$$C_\ell^{\text{input}} = \frac{1}{\ell + 10}, \quad (34)$$

and sample them at the positions of the *Quaia* sources. We generate 1000 such realizations. Likewise, we generate 100 Gaussian maps from a steeper power spectrum of the form

$$C_\ell^{\text{syst}} = \frac{A_{\text{syst}}}{(\ell + 10)^3}, \quad (35)$$

which we sample at the same source positions and use as contaminant templates to deproject. We then construct a specific linear combination of these contaminant templates to contaminate the signal assigned to each source. For simplicity, we use the direct sum of all templates, scaling the proportionality constant A_{syst} to ensure that the total variance due to systematic contamination in each source is 30 per cent of the signal. Finally, we generate uncorrelated Gaussian numbers for each source to serve as a noise component, with a standard deviation equal to 10 times that of the signal. Fig. 1 shows the input and contaminant power spectra, as well as the expected noise level, together with the mean power spectrum of all simulations. We note that the amplitudes of the contaminant and noise components have been chosen to exacerbate the impact of systematic deprojection on the simulations: the measurements are systematics-dominated on scales $\ell \simeq 100$. Furthermore, the data are noise dominated on a source-by-source basis, as well as on all scales. This, together with the large number

⁹<http://healpix.sf.net/>




Figure 2. Power spectrum of noise-only simulations with and without contaminant deprojection (black and red points, respectively). The loss of power due to deprojection leads to a negative power spectrum (note that the catalogue-based estimator automatically removes the white noise component), which is accurately predicted by equation (23) (solid blue line).

of contaminant templates being deprojected, increases the amplitude of the correction to the noise bias due to deprojection (ΔN_ℓ^a in equation 23).

To first validate our analytical expression used to calculate ΔN_ℓ^a , we study the power spectrum of noise-only simulations, in which we only include the contribution from uncorrelated noise to the field value for each source. Fig. 2 shows the mean power spectrum of all such simulations before and after deprojecting all 100 contaminant templates (black and red points, respectively). We remind the reader that the white noise component is avoided by the catalogue-based estimator by construction. The loss of modes due to deprojection leads to a negative power spectrum at low ℓ , which is accurately recovered by our analytical expression (blue line).

Finally, Fig. 3 shows the result of our validation of the formalism outlined here. The purple points show the average power spectrum for contaminated simulations after contaminant deprojection bias, before accounting for any sources of deprojection bias. On large scales, the power spectrum shows a very significant lack of power, even becoming negative at the lowest ℓ s. Much of this power is recovered by correcting for the white noise deprojection bias calculated via equation (23), as shown by the orange points. The remaining power loss can be recovered by correcting for the deprojection bias of the signal component. As described in Section 2.3, this may be done analytically if the true signal power spectrum is known, or through a simulation-based transfer function approach. The red points show the result of the latter method,¹⁰ which is able to recover an unbiased power spectrum at all scales.

3.2 Contaminant deprojection in galaxy clustering

Here, we validate the galaxy clustering formalism described in Section 2.2.2 by again using simulated catalogue-based observations, with contributions from sky contaminants imposed on the

¹⁰It is worth noting that the true signal power spectrum is used here to generate the simulations required for deriving the transfer function (see Section 2.3). However, as we will show in Section 3.3, this approach is relatively insensitive to this choice of power spectrum.




Figure 3. Mean and standard deviation of the power spectra from contaminated simulations after contaminant deprojection (purple points, dashed line). The orange points (solid line) show the result of accounting for the impact of deprojection in the white noise component (via equation 23), while the red points (dotted line) show the power spectra after accounting for deprojection bias in both the noise and signal components (the latter via the transfer function approach). The result is an unbiased estimate of the input power spectrum (solid black line). The bottom panel shows the residuals with respect to the input spectrum normalized by the error in the mean over all simulations.

distribution of the sources. We assess the two possible treatments of the mask: using the positions and weights of randoms to define it, and providing it as a continuous map. We will hereafter refer to these as the randoms-based and mask-based approaches, respectively.

For each simulation, we begin by generating a Gaussian realization of a map, $\delta_{g,i}^{\text{in}}$, from some input angular power spectrum, C_ℓ^{input} . For this test, we use the best-fitting angular power spectrum for the *Quiaia* ‘low- z ’ ($0 < z < 1.47$) sample, described in D. Alonso et al. (2025) and G. Fabbian et al. (2026), as the input for our simulations. A second map, $\delta_{g,i}^{\text{cont}}$, is also generated from C_ℓ^{input} and contaminated with templates f^p using known linear coefficients A_p such that the value in pixel i is given by

$$\delta_{g,i}^{\text{cont}} = \delta_{g,i}^{\text{in}} + \sum_{p=1}^{N_{\text{sys}}} A_p f_i^p. \quad (36)$$

We use the same 10 templates as used in D. Alonso et al. (2025) and G. Fabbian et al. (2026), which the authors deproject from *Quiaia* overdensity maps using the standard formalism as outlined in Section 2.2. The use of these templates here is motivated by the desire to test our catalogue-based formalism in a realistic scenario for galaxy clustering. Details of these templates can be found in Appendix A, including Mollweide projections of each (Fig. A1).

We then use a mask defining the *Quiaia* survey geometry to convert each of these maps into a probability map, P , whose value in a pixel i is given by:

$$P_i^X = (1 + \delta_{g,i}^X) w^i, \quad (37)$$

where w is a map describing the survey mask, and $X \in \{\text{in}, \text{cont}\}$. As in Section 3.1, all maps are created using a HEALPIX pixelization scheme with a resolution parameter of $N_{\text{side}} = 256$. These maps are then normalized such that $\max(P^X) = 1$, and subsequently Poisson-sampled to create realizations of the catalogue with and without contamination from systematics. We note that it is possible, when generating these probability maps, for a pixel

to end up with a value less than zero, and cannot therefore be Poisson-sampled. To test whether this is likely to be an issue, we use the prescription outlined above to generate two sets of 10^5 probability maps, one with contamination and one without, and count how many pixels in each have values below zero. We find that there are no such pixels in any of the maps, and thus conclude that they are unlikely to contribute significantly to our analysis.

As highlighted in Section 2.2.2, it is necessary in the case of galaxy clustering systematics to truncate at some maximum multipole ℓ_{\max} the harmonic-space expression for the dot product of two fields in order to avoid losing information about real correlations between the galaxy overdensity and the contaminants (see equation 27). The choice of ℓ_{\max} depends on the contaminants in question. To determine a suitable value for the templates used here, we rerun the suite of simulations using several different values of ℓ_{\max} ranging from $\ell_{\max} = 10$ to $\ell_{\max} = 3N_{\text{side}} - 1 = 767$.¹¹

The results of this test are shown in Fig. 4, wherein each panel shows the means and standard deviations in the deprojection coefficient \hat{A}_p for a given contaminant as measured using the randoms- and mask-based approaches (teal circles and red squares, respectively), along with the input value used when generating the simulated data. Both the randoms- and mask-based deprojection methods produce biased estimates of the deprojection coefficients for several of the contaminants evaluated here if one uses an ℓ_{\max} that is either too low or too high. We find that the most stable and least biased estimates tend to occur at $\ell_{\max} \sim 100$, and as such adopt this value going forward with our simulations.

The reason for this bias in the estimated \hat{A}_p is related to the pixel window function, and is a by-product of the fact that our contaminant templates take the form of maps. While in the map-based formalism it is straightforward to correct for the pixel window function, we can see from equation (27) that such a correction is not applicable here, owing to calculation of $\text{PCL}_\ell(a^v, \hat{f}^q)$ required to obtain \hat{A}_p . We reiterate that a^v is the (masked) catalogue-based field being observed, and the \hat{f}^q is a given template (also masked), the nature of which depends on which approach is being used. In the mask-based approach, the template maps are used directly, resulting in a PCL cross-correlating a catalogue-based field with a map-based field, for which a correction for the pixel window function is not viable. One might think this is mitigated by using the randoms-based approach, in which \hat{f}^q is computed by sampling the template at the positions of the randoms, but since these values are still being sampled from a map of finite resolution, the pixel window function will still have an effect that is not easy to remove. While this effect is inherent to our simulations, it is also likely in a realistic scenario that one will be working with pixelized contaminant templates, in which case this will still be an issue. Since systematics tend to affect clustering measurements at larger scales, a relatively conservative ℓ_{\max} of ~ 100 should be sufficient to capture the contribution from these systematics in most use cases, while significantly mitigating the effects of the pixel window function.

Fig. 5 shows the result of applying the template deprojection

procedure outlined in Section 2.2.2 to the set of contaminated catalogues from our simulations, using both the randoms-based (teal circles; top panel) and mask-based (red squares; middle panel) approaches. As was seen for sampled fields in Section 3.1, there is a significant loss of power at large scales, which is almost entirely mitigated by accounting for the noise deprojection bias via equation (28), since these data are noise-dominated in this regime. Consequently, attempts to apply the transfer function approach outlined in Section 2.3 appear to exacerbate the deviations from the true power spectrum, although these shifts are marginal and the resulting power spectra are still consistent within the errorbars. Overall, the final C_ℓ estimates are consistent with the input power spectrum to within $\sim 3\sigma$ at all scales. Since our simulated catalogues have been sampled from maps of finite resolution, aliasing effects are expected to arise beyond $\ell \approx 2N_{\text{side}}$; this can be seen for the mask-based case in the shaded grey region of the bottom panel in Fig. 5, while the randoms-based approach appears relatively unaffected. We therefore caution users against using the mask-based approach to probe scales of $\ell > 2N_{\text{side}}$. Importantly, we also note that the uncertainties in Fig. 5 correspond to the mean over 1000 simulated data sets, and are therefore $\sqrt{1000} \sim 30$ times smaller than the actual errors for a real-world data set. The small deviations shown in the figure should therefore be completely negligible in a realistic scenario.

Given that the systematics prevalent in LSS analyses tend to dominate at large angular scales, it is worth discussing here the efficacy of our method in the context of analyses for which these scales are the most relevant. For example, tomographic clustering-based constraints on the amplitude of local primordial non-Gaussianity, f_{NL} , are particularly reliant on having unbiased measurements of the large-scale angular power spectrum (e.g. N. Dalal et al. 2008; A. Slosar et al. 2008; J.-Q. Xia et al. 2011; A. R. Pullen & C. M. Hirata 2013; T. Giannantonio et al. 2014; B. Leistedt & H. V. Peiris 2014; B. Leistedt, H. V. Peiris & N. Roth 2014; S. Ho et al. 2015; E. Castorina et al. 2019; A. Krolewski et al. 2024; S. Chiarenza et al. 2025; G. Fabbian et al. 2026). We thus now explore the results of this section from the perspective of constraining f_{NL} .

The ‘low- z ’ *Quiaia* sample, used as the basis for the simulations used in this work, has also been used in G. Fabbian et al. (2026) for exactly this purpose. A simple yet conservative test of the implications of our systematics mitigation approach for constraining f_{NL} can be performed as follows: using the redshift distribution for this QSO sample (obtained from the data products of D. Alonso et al. 2023) and assuming the Planck 2018 cosmology (Planck Collaboration VI 2020) with a PNG response parameter $p_\phi = 1$, we generate theoretical predictions of the quasar clustering C_ℓ s for $f_{\text{NL}} \in [-100, 100]$. We then evaluate each of these theoretical C_ℓ s at the effective ℓ of the lowest bandpower considered here ($\ell_{\text{eff}} = 2.5$), and use this to estimate the dependence of f_{NL} on the measured value of C_ℓ at this scale. Comparing the corresponding values of f_{NL} for the final corrected C_ℓ and the input C_ℓ used to generate the simulations, we find that for the randoms-based approach this difference amounts to a shift of $\Delta f_{\text{NL}} \approx 6$. This is larger than the precision of f_{NL} measurements required to distinguish between single- and multifield models of inflation is $\Delta f_{\text{NL}} \lesssim 1$ (e.g. P. Creminelli & M. Zaldarriaga 2004), however we also cannot rule out that this shift is driven by the statistical uncertainty in our validation process. Indeed, we find that the 1σ uncertainty on the final corrected C_ℓ at $\ell_{\text{eff}} = 2.5$ corresponds to

¹¹HEALPIX maps generated with a resolution parameter of N_{side} are approximately band-limited to multipoles with $\ell \lesssim 3N_{\text{side}} - 1$.

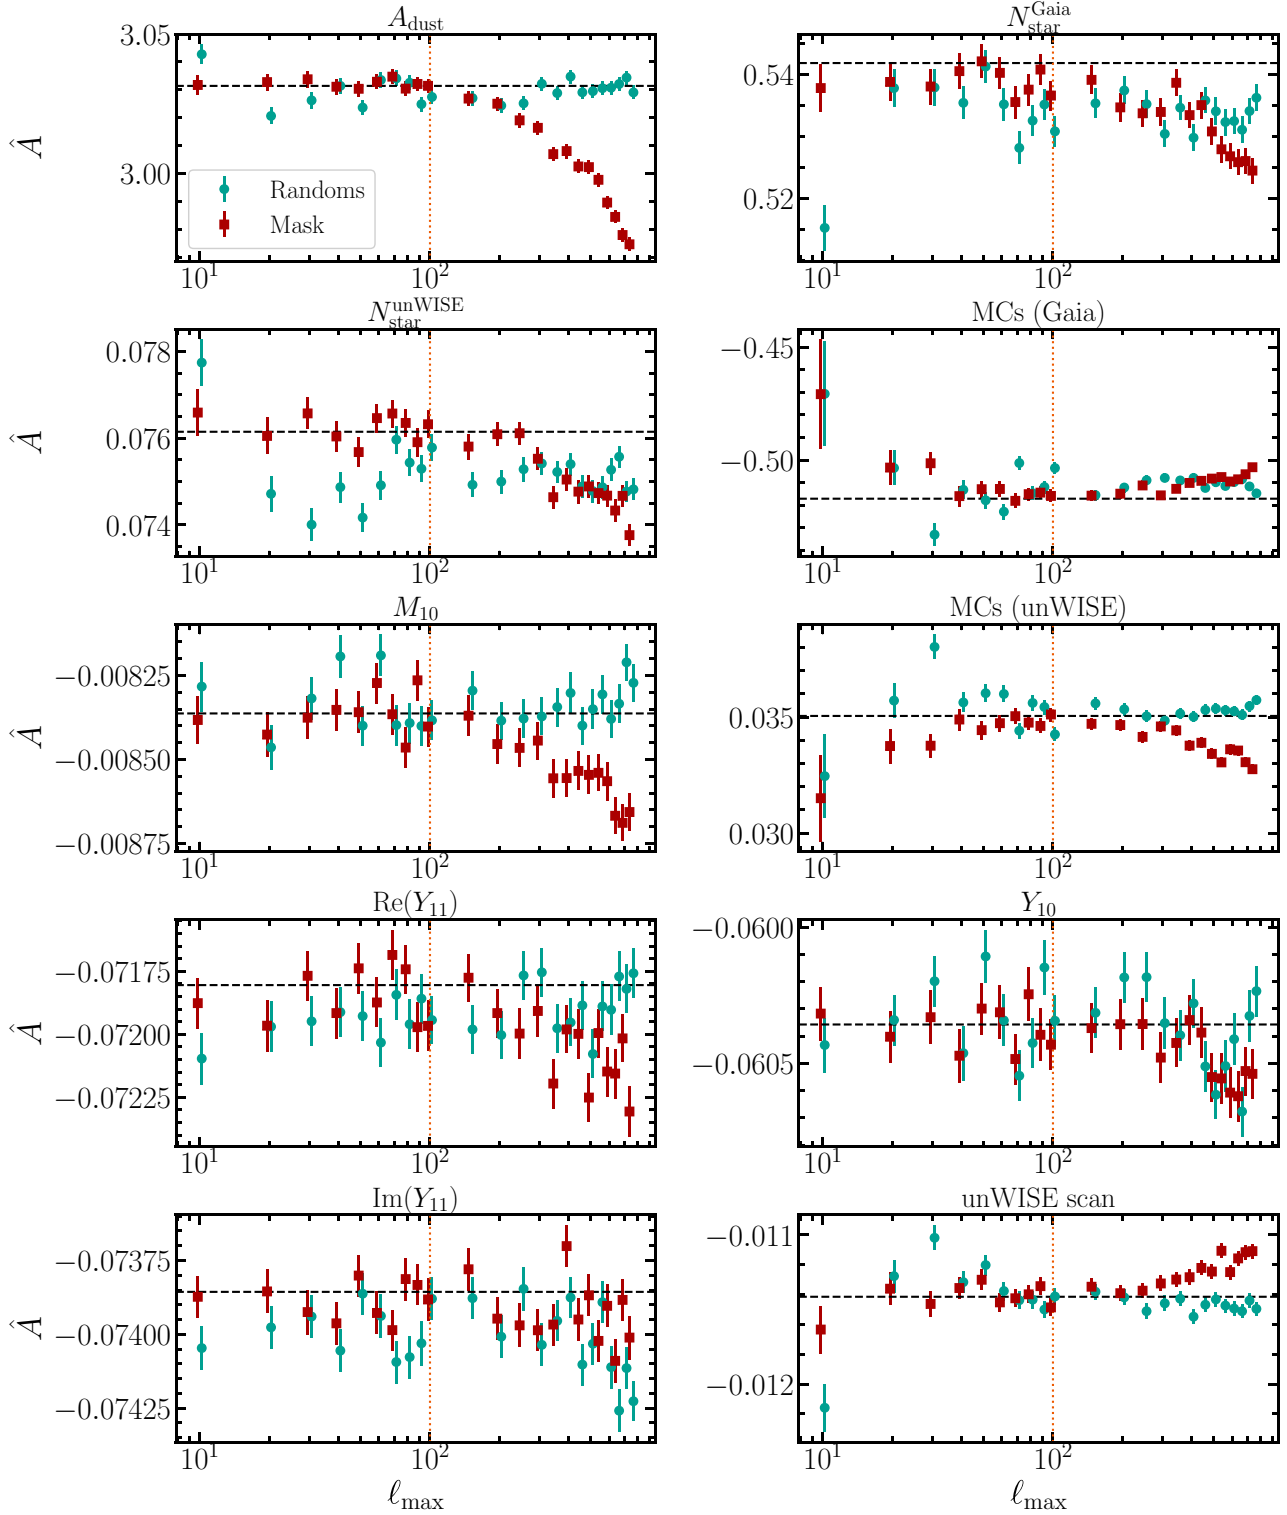


Figure 4. Linear deprojection coefficients measured as a function of ℓ_{\max} from simulations using the randoms-based (teal circles) and mask-based (red squares) approaches. The procedure is outlined in Section 3.2, and details of the templates (indicated at the top of each panel) can be found in Appendix A. The dashed black line in each panel shows the input deprojection coefficient used to contaminate the simulated data. Small horizontal offsets have been applied to the data for clarity. Both methods show significant variation in the deprojection coefficients for several contaminants, particularly as one moves to the lowest and highest multipoles. The dotted orange line indicates $\ell_{\max} = 100$, which we deem a suitable value for minimizing the bias in our measured deprojection coefficients for the majority of templates.




Figure 5. Means and standard deviations in angular power spectra from contaminated simulations measured at various stages of correction using the randoms-based (top panel) and mask-based (middle panel) approaches. Different marker sizes and line styles represent the different stages: after contaminant deprojection (small markers, dashed lines); after removing the noise deprojection bias (medium-sized markers, solid lines); after also correcting for deprojection bias in the signal component using the transfer function approach (large markers, dotted lines). The black solid line shows the angular power spectrum used to generate the (non-contaminated) data, and as such is indicative of the ‘true’ power spectrum. The shaded region begins at $\ell = 2N_{\text{side}}$, beyond which we expect aliasing effects to arise from the fact that we have sampled the catalogue from a map with finite resolution. The bottom panel shows the residuals of the fully corrected power spectra with respect to the input power spectrum, normalized by the error in the mean across the simulations. Both the randoms-based (teal circles) and mask-based (red squares) approaches are able to accurately reproduce the input power spectrum, with the randoms-based approach maintaining better accuracy than the mask-based at the higher multipoles. Exacerbated deviations from the input power spectrum after applying the transfer function (T_ℓ) at the largest scales are due to the data being noise-dominated in this regime (see text).

$\Delta f_{\text{NL}} \approx 4$. We therefore conclude that our systematics mitigation approach has a low risk of producing significantly biased measurements of f_{NL} .

3.3 Comparison between debiasing methods

We now use simulations to validate the transfer function approach described in Section 2.3 (see equation 33) as a means of correcting for mode loss when deprojecting templates, and compare with the deprojection bias approach normally employed for this purpose (see equation 31; G. B. Rybicki & W. H. Press 1992; F. Elsner et al. 2017). In particular, we wish to test how dependent the transfer function approach is on assumptions about the true power spectrum.

As in Section 3.2, we begin with 1000 Gaussian realizations of the angular power spectrum obtained from the *Quaia* survey, where we have contaminated each realization with those same

templates via equation (36) and applied the *Quaia* mask w . We will refer to this as the ‘main’ suite of simulations henceforth. Using the standard map-based PCL approach (Sections 2.1.1 and 2.1.2), we then deproject the templates to obtain biased estimates of the input power spectrum, \hat{C}_ℓ . The transfer function or the deprojection bias estimates are then used to debias these angular power spectra. We use the debiased power spectra as a means of assessing the efficacy of the transfer function approach.

To calculate the various transfer functions we wish to test, we generate several more suites of simulated maps, each producing 1000 Gaussian realizations of a different input C_ℓ . For each realization, we apply the same mask as for the main suite and use the standard PCL approach to measure the angular power spectrum of the masked field, C_ℓ^{before} . We then use the map-based approach to deproject the same templates used to contaminate the main suite and measure the angular power spectrum post-deprojection, C_ℓ^{after} . The transfer function is then computed using equation (33), averaging the relevant quantities across all simulations in the suite.

We use the following angular power spectra as inputs for each simulation suite:

- (i) The *Quaia* C_ℓ used to generate the the main suite, whose measured angular power spectra we aim to debias. This is the true power spectrum and therefore should, by definition, produce unbiased results after applying the transfer function.
- (ii) A guess at the true power spectrum, derived by measuring the PCL of one simulated *Quaia*-like data set and dividing by the covered sky fraction, estimated as $f_{\text{sky}} = \langle w^2 \rangle$.
- (iii) A power spectrum with an arbitrarily chosen functional form, $C_\ell = 1/(\ell + 10)$.
- (iv) A flat power spectrum with arbitrary amplitude, $C_\ell = K$.

In addition, we use the standard deprojection bias approach to debias the measured power spectra, using one of the first two C_ℓ s described above as the estimate of the true power spectrum (C_ℓ^{true} in equation 32).

Fig. 6 compares the deviation (normalized by the standard error of the mean) of the measured angular power spectra from the input power spectrum, after correcting for mode loss due to deprojection using one of the methods outlined above. We find that all methods perform similarly well at most scales, being consistent with each other within the errorbars. Ignoring scales at which $\ell \gtrsim 2N_{\text{side}}$ (demarcated by the grey shaded region in the figure) where aliasing effects are expected to arise from the pixelization of our maps, it is only at the largest scales ($\ell \lesssim 3$) that we see any difference between the various methods. Namely, the transfer function approach fails to produced unbiased estimates at these scales if an unrealistic guess at the true angular power spectrum is used (see orange stars and cyan diamonds in Fig. 6). We thus conclude that the transfer function approach is a viable alternative to calculating the deprojection bias, so long as one uses a reasonable guess at the true power spectrum when generating the required simulations. For a realistic scenario in which the truth is not known, one can derive a sufficient guess by dividing the measured PCL by the observed sky fraction f_{sky} , as has been done here (teal downward-pointing and purple right-pointing triangles in Fig. 6, for the deprojection bias and transfer function approaches, respectively).




Figure 6. Residuals of the angular power spectra measured from simulations with respect to the power spectrum used to generate the simulated data, having mitigated the mode loss due to deprojection using different approaches. Namely, either the standard deprojection bias (DB) approach or the transfer function (TF) approach is used, with the chosen guess at the input power spectrum indicated in the legend. Small horizontal shifts have been applied to the data to aid visualization, and the results have been partially split between two panels for ease of comparison between the most relevant methods. Arrows of a given colour denote scales at which the angular power spectrum deviates from the truth by more than 5σ after debiasing with the corresponding method. *Top panel:* The DB approach using the true angular power spectrum (blue upward-pointing triangles) and a data-driven guess at its form (teal downward-pointing triangles), along with the TF approach using the same angular power spectra (red left-pointing and purple right-pointing triangles, respectively). The DB and TF approaches produce similarly unbiased C_ℓ estimates in both cases. *Bottom panel:* The TF approach using the four different angular power spectra described in the text and indicated in the legend. All four perform similarly well overall, although the choice of guess C_ℓ matters more at low ℓ .

3.4 Catalogue-based transfer functions

While we find in Section 3.3 that the approach of using map-based simulations to debias catalogue-based data appears to work sufficiently well, one question that remains is whether one can instead use catalogue-based simulations such as the those used in Section 3.2 to estimate a transfer function that more fairly represents the data in question.

To test this, we again begin by generating 1000 Gaussian realizations of an input C_ℓ . Since we find that the transfer function derived using a data-driven guess for the true C_ℓ performs better than if some arbitrary C_ℓ is used (see Fig. 6), we opt for this as our input C_ℓ here. This also represents a more realistic scenario one would encounter when working with real data, for which the true power spectrum is unknown. From each of these Gaussian

realizations we then draw Poisson realizations of a catalogue with a mean of 644 786 sources per catalogue (i.e. the same number of sources as in the *Quaia* ‘low- z ’ sample), following the procedure outlined in Section 3.2. From these we compute the angular power spectra using the catalogue-based approach, both with and without deprojecting the contaminant templates. We do this for both the case where the mask is provided directly as a continuous map, and the case where randoms are used to define the mask, and compute the transfer function for each using equation (33).

In both cases, we find that the catalogue-based simulations produce significantly noisier transfer functions than map-based simulations using the same input C_ℓ , with the scatter across simulations being up to ~ 100 (1000) times greater for the randoms-based (mask-based) approach. This is due to the impact of shot noise on the measured spectra before and after deprojection. We therefore recommend that the map-based approach of Section 3.3 be used to derive transfer functions, particularly when working with shot noise-dominated data sets. Note that this also has the advantage of being less expensive computationally in most real-world applications.

3.5 Application to *Quaia*

Having validated our method using simulations, we now test its application to real data. We again turn to the *Quaia* catalogue (K. Storey-Fisher et al. 2024), focusing in particular on the ‘low- z ’ sample of 644 768 quasars with spectro-photometric redshifts in the range $z_p < 1.47$. After removing any quasars with redshift estimates outside of this range, we deproject the same templates used by D. Alonso et al. (2025) and G. Fabbian et al. (2026) (see Appendix A for details) to mitigate the contribution of systematics. We do this using three approaches for the sake of comparison: the standard map-based approach, the catalogue-based approach in the case of a continuous mask, and the catalogue-based approach using randoms to define the mask. In the latter two approaches, the shot noise is calculated with equation (12) (ignoring the contribution from randoms in the mask-based case) and subtracted from the measured power spectra. For the map-based approach, we assume that the shot noise is entirely Poissonian, in which case it can also be calculated analytically (D. Alonso et al. 2019; A. Nicola et al. 2020): the PCL of the noise is given by

$$\tilde{N}_\ell = \Omega_{\text{pix}} \frac{\langle w \rangle}{\langle n \rangle}, \quad (38)$$

where Ω_{pix} is the pixel area in steradians ($\approx 1.6 \times 10^{-5}$ sr for $N_{\text{side}} = 256$), $\langle w \rangle$ is the mean value of the *Quaia* mask across the map, and $\langle n \rangle$ is the mean effective number of sources per pixel. The power spectrum of the shot noise is then obtained by multiplying it by the inverse MCM calculated from the *Quaia* mask.

Fig. 7 shows the noise-subtracted angular power spectra pre- and post-deprojection as measured with the three different approaches. Uncertainties are estimated via the computation of Gaussian (or ‘disconnected’) power spectrum covariance matrices. Once the noise deprojection bias has been accounted for in the catalogue-based approaches, all three methods produce very similar measurements with respect to each other, both before and after template deprojection has been applied. The consistency of the deprojection methods themselves is exemplified in the bottom panel, in which we show the deviation in power induced by this step, normalized by the uncertainties: the deviations agree within




Figure 7. Application of three methods of deprojection to the low- z sample of quasars from *Quiaia*. The first three panels show (from top to bottom) the results of applying the traditional map-based approach to template deprojection, the catalogue-based approach in the case where randoms have been used to define the mask, and the catalogue-based approach where the mask has been provided as a continuous map. In each case, we show the angular power spectra pre- and post-deprojection (correcting for the noise deprojection bias in the catalogue-based approaches). The shot noise has already been subtracted in all cases. Open symbols denote negative power, for which we plot the absolute values. The bottom panel shows the difference between the pre- and post-deprojection angular power spectra, normalized by the corresponding uncertainties. All three methods are found to result in similar power deviations across all scales.

errorbars for all three methods at all angular scales. This test thus validates our implementation of template deprojection in the catalogue-based PCL framework when applied to real data.

4 CONCLUSIONS

In this work, we present an extension to the catalogue-based PCL formalism of K. Wolz et al. (2025) that allows one to apply template deprojection in order to mitigate the contribution from systematic effects. In doing so, we combine the advantages of the existing catalogue-based formalism – i.e. avoiding biases stemming from the pixelization of the observed fields – with systematics mitigation capabilities previously only available in the map-based formalism. We describe the application of this formalism to fields sampled at the positions of discrete sources, as well as to the particular case of galaxy clustering, in which the density of sources is the field of interest. Note that while the fields considered here are all scalar, our implementation is applicable to fields of arbitrary spin. A key part of this formalism lies in being able

to account for mode loss in the noise component, induced by the deprojection procedure, which we demonstrate can be accurately computed analytically. We also introduce an alternative method to account for deprojection bias in the signal component, via the computation of a transfer function, and demonstrate its efficacy using simulations.

For the case of discretely-sampled fields, we validate our implementation by simulating 1000 Gaussian realizations of an arbitrary angular power spectrum, which are then contaminated with 100 synthetic templates and subsequently Poisson-sampled at the positions of galaxies from a real data set (*Quiaia*) to create 1000 realizations of a catalogue with a realistic source distribution. We show that our implementation can be used to obtain unbiased estimates of the true power spectrum, after the mode loss induced by deprojection has been accounted for in both the noise and signal components.

We then simulate 1000 *Quiaia*-like catalogues and test the clustering-specific part of our formalism. We first demonstrate that care must be taken when choosing a maximum multipole, ℓ_{\max} , for the deprojection procedure, as choosing a scale that is too small or too large can result in biased estimates of the deprojection parameters (which in turn will bias the final angular power spectrum). Having selected a suitable multipole ($\ell_{\max} = 100$), we show that unbiased estimates of the true power spectrum can again be obtained after correcting for mode loss in the noise component. Since these simulated catalogues are shot noise-dominated, the effect of the transfer function is minimal in comparison, but we emphasize that for denser data sets this step will likely be necessary to achieve unbiased results. We find that all of this is true regardless of whether the mask is defined via a catalogue of randoms, or as a continuous map.

To assess the efficacy of the transfer function approach to mitigating mode loss caused by deprojection, we derive four transfer functions using different C_ℓ s as input for the simulations, and compare their performance with the standard deprojection bias approach by applying them each to a set of 1000 simulated *Quiaia*-like catalogues. All methods are found to produce similarly unbiased results, except at scales $\ell \lesssim 3$, where at least a reasonable approximation of the true angular power spectrum is required.

Finally, we apply our formalism to real data from *Quiaia* and compare it to the standard map-based approach. After correcting for the noise deprojection bias, both the randoms-based and mask-based approaches produce measurements consistent with the map-based approach, pre- and post-deprojection.

We have thus demonstrated that our formalism is a viable alternative to the traditional map-based approach when working with fields sampled at discrete positions. We anticipate that the method presented here will be useful in the analysis of large data sets incoming from Stage IV surveys such as the Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST), for which systematics mitigation will be a crucial step towards obtaining unbiased cosmological measurements, and the wealth of information contained at small scales will be better exploited by avoiding pixelization. Our implementation is made available with the latest release of the public PCL package `NaMaster`.

ACKNOWLEDGEMENTS

This work has been supported by Science and Technology Facilities Council (STFC) funding for UK participation in LSST, through grant ST/X00127X/1. DA and TC acknowledge support from the Becroft Trust. BL is supported by the Royal Society

through a University Research Fellowship. KW acknowledges support from the STFC under grant ST/X006344/1. Extensive use was made of computational resources at the University of Oxford Department of Physics, funded by the John Fell Oxford University Press Research Fund.

DATA AVAILABILITY

The quasar catalogue used in this work is publicly available at <https://zenodo.org/records/12805938>.

REFERENCES

- Abbott T. M. C. et al., 2022, *Phys. Rev. D*, 105, 023520
- Adam R. et al., 2017, *A&A*, 598, A115
- Agarwal N. et al., 2014, *J. Cosmol. Astropart. Phys.*, 2014, 007
- Alonso D., 2025, *Open J. Astrophys.*, 8, 53
- Alonso D., Sanchez J., Slosar A., *LSST Dark Energy Science Collaboration*, 2019, *MNRAS*, 484, 4127
- Alonso D., Fabbian G., Storey-Fisher K., Eilers A.-C., García-García C., Hogg D. W., Rix H.-W., 2023, *J. Cosmology Astropart. Phys.*, 2023, 043
- Alonso D., Hetmantsev O., Fabbian G., Slosar A., Storey-Fisher K., 2025, *Open J. Astrophys.*, 8, 42
- Baleato Lizancos A., White M., 2024, *J. Cosmol. Astropart. Phys.*, 2024, 010
- Bleem L. E. et al., 2022, *ApJS*, 258, 36
- Bocquet S. et al., 2019, *ApJ*, 878, 55
- Camphuis E. et al., 2025, preprint ([arXiv:2506.20707](https://arxiv.org/abs/2506.20707))
- Carron J., Mirmelstein M., Lewis A., 2022, *J. Cosmol. Astropart. Phys.*, 2022, 039
- Castorina E. et al., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 010
- Chiang Y.-K., 2023, *ApJ*, 958, 118
- Chiarenza S. et al., 2025, preprint ([arXiv:2512.17865](https://arxiv.org/abs/2512.17865))
- Coulton W. et al., 2024, *Phys. Rev. D*, 109, 063530
- Creminelli P., Zaldarriaga M., 2004, *Phys. Rev. D*, 70, 083532
- Dalal N., Doré O., Huterer D., Shirokov A., 2008, *Phys. Rev. D*, 77, 123514
- Datta R. et al., 2024, *ApJS*, 273, 26
- Elsner F., Leistedt B., Peiris H. V., 2016, *MNRAS*, 456, 2095
- Elsner F., Leistedt B., Peiris H. V., 2017, *MNRAS*, 465, 1847
- Elvin-Poole J. et al., 2018, *Phys. Rev. D*, 98, 042006
- Fabbian G., Alonso D., Storey-Fisher K., Cornish T., 2026, *J. Cosmol. Astropart. Phys.*, 2026, 40
- Faga L. et al., 2025, *MNRAS*, 536, 1586
- Fang X. et al., 2024, *MNRAS*, 527, 9581
- Ge F. et al., 2025, *Phys. Rev. D*, 111, 083534
- Giannantonio T., Ross A. J., Percival W. J., Crittenden R., Bacher D., Kilbinger M., Nichol R., Weller J., 2014, *Phys. Rev. D*, 89, 023511
- Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, 622, 759
- Heymans C. et al., 2021, *A&A*, 646, A140
- Hivon E., Górski K. M., Netterfield C. B., Crill B. P., Prunet S., Hansen F., 2002, *ApJ*, 567, 2
- Ho S. et al., 2012, *ApJ*, 761, 14
- Ho S. et al., 2015, *J. Cosmol. Astropart. Phys.*, 2015, 040
- Huterer D., Cunha C. E., Fang W., 2013, *MNRAS*, 432, 2945
- Jego B., Ruiz-Zapatero J., García-García C., Koukoufilippas N., Alonso D., 2023, *MNRAS*, 520, 1895
- Johnston H. et al., 2021, *A&A*, 648, A98
- Krolewski A. et al., 2024, *J. Cosmol. Astropart. Phys.*, 2024, 021
- Landy S. D., Szalay A. S., 1993, *ApJ*, 412, 64
- Leistedt B., Peiris H. V., 2014, *MNRAS*, 444, 2
- Leistedt B., Peiris H. V., Mortlock D. J., Benoit-Lévy A., Pontzen A., 2013, *MNRAS*, 435, 1857
- Leistedt B., Peiris H. V., Roth N., 2014, *Phys. Rev. Lett.*, 113, 221301
- Louis T. et al., 2025, *J. Cosmol. Astropart. Phys.*, 2025, 062
- Maniyar A., Lagache G., Béthermin M., Ilić S., 2019, *A&A*, 621, A32
- Miyatake H. et al., 2023, *Phys. Rev. D*, 108, 123517
- Nicola A. et al., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 044
- Nicola A., García-García C., Alonso D., Dunkley J., Ferreira P. G., Slosar A., Spergel D. N., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 067
- Planck Collaboration VI, 2020, *A&A*, 641, A6
- Pullen A. R., Hirata C. M., 2013, *PASP*, 125, 705
- Qu F. J. et al., 2024, *ApJ*, 962, 112
- Reinecke M., Belkner S., Carron J., 2023, *A&A*, 678, A165
- Rezaie M., Seo H.-J., Ross A. J., Bunescu R. C., 2020, *MNRAS*, 495, 1613
- Rodríguez-Monroy M. et al., 2022, *MNRAS*, 511, 2665
- Ross A. J. et al., 2011, *MNRAS*, 417, 1350
- Rybicki G. B., Press W. H., 1992, *ApJ*, 398, 169
- Shafer D. L., Huterer D., 2015, *MNRAS*, 447, 2961
- Shaikh S. et al., 2024, *MNRAS*, 528, 2112
- Slosar A., Seljak U., Makarov A., 2004, *Phys. Rev. D*, 69, 123003
- Slosar A., Hirata C., Seljak U., Ho S., Padmanabhan N., 2008, *J. Cosmol. Astropart. Phys.*, 2008, 031
- Stözlner B. et al., 2025, *A&A*, 702, A169
- Stompor R., Efstathiou G., 1999, *MNRAS*, 302, 735
- Storey-Fisher K., Hogg D. W., Rix H.-W., Eilers A.-C., Fabbian G., Blanton M. R., Alonso D., 2024, *ApJ*, 964, 69
- Sugiyama S. et al., 2023, *Phys. Rev. D*, 108, 123521
- Tegmark M., 1997, *Phys. Rev. D*, 55, 5895
- Totsuji H., Kihara T., 1969, *PASJ*, 21, 221
- Vakili M. et al., 2023, *A&A*, 675, A202
- Weaverdyck N., Huterer D., 2021, *MNRAS*, 503, 5061
- Weaverdyck N., Muir J., Huterer D., 2018, *Phys. Rev. D*, 97, 043515
- Wolz K., Alonso D., Nicola A., 2025, *J. Cosmol. Astropart. Phys.*, 2025, 028
- Xia J.-Q., Baccigalupi C., Matarrese S., Verde L., Viel M., 2011, *J. Cosmol. Astropart. Phys.*, 2011, 033
- Yan Z. et al., 2025, *A&A*, 694, A259
- Zonca A., Singer L., Lenz D., Reinecke M., Rosset C., Hivon E., Gorski K., 2019, *J. Open Source Softw.*, 4, 1298
- Zubeldia Í., Challinor A., 2019, *MNRAS*, 489, 401

APPENDIX A: QUAIA TEMPLATES

We include here brief details of the *Quaia* systematics templates used for the analysis carried out in Sections 3.2 and 3.3. These templates include: a Galactic dust extinction map (A_{dust}) from Y.-K. Chiang (2023); two maps of stellar contamination from *Gaia* and *unWISE* ($N_{\text{star}}^{\text{Gaia}}$ and $N_{\text{star}}^{\text{unWISE}}$, respectively); the scanning patterns of *Gaia* and *unWISE* (M_{10} and *unWISE scan*); stellar density maps in the regions around the Large and Small Magellanic Clouds measured with *Gaia* and *unWISE* (MCs (*Gaia*) and MCs (*unWISE*)); the three dipole moments associated with a multipole of $\ell = 1$ (Y_{10} , $\text{Re}(Y_{11})$, $\text{Im}(Y_{11})$).

We show Mollweide-projected maps of these templates in Fig. A1. In each case, we show the residual of a given pixel with respect to the mean value of that systematic across the footprint ($S(\hat{n}) - \bar{S}$), normalized by the rms deviation across the footprint (σ_S).

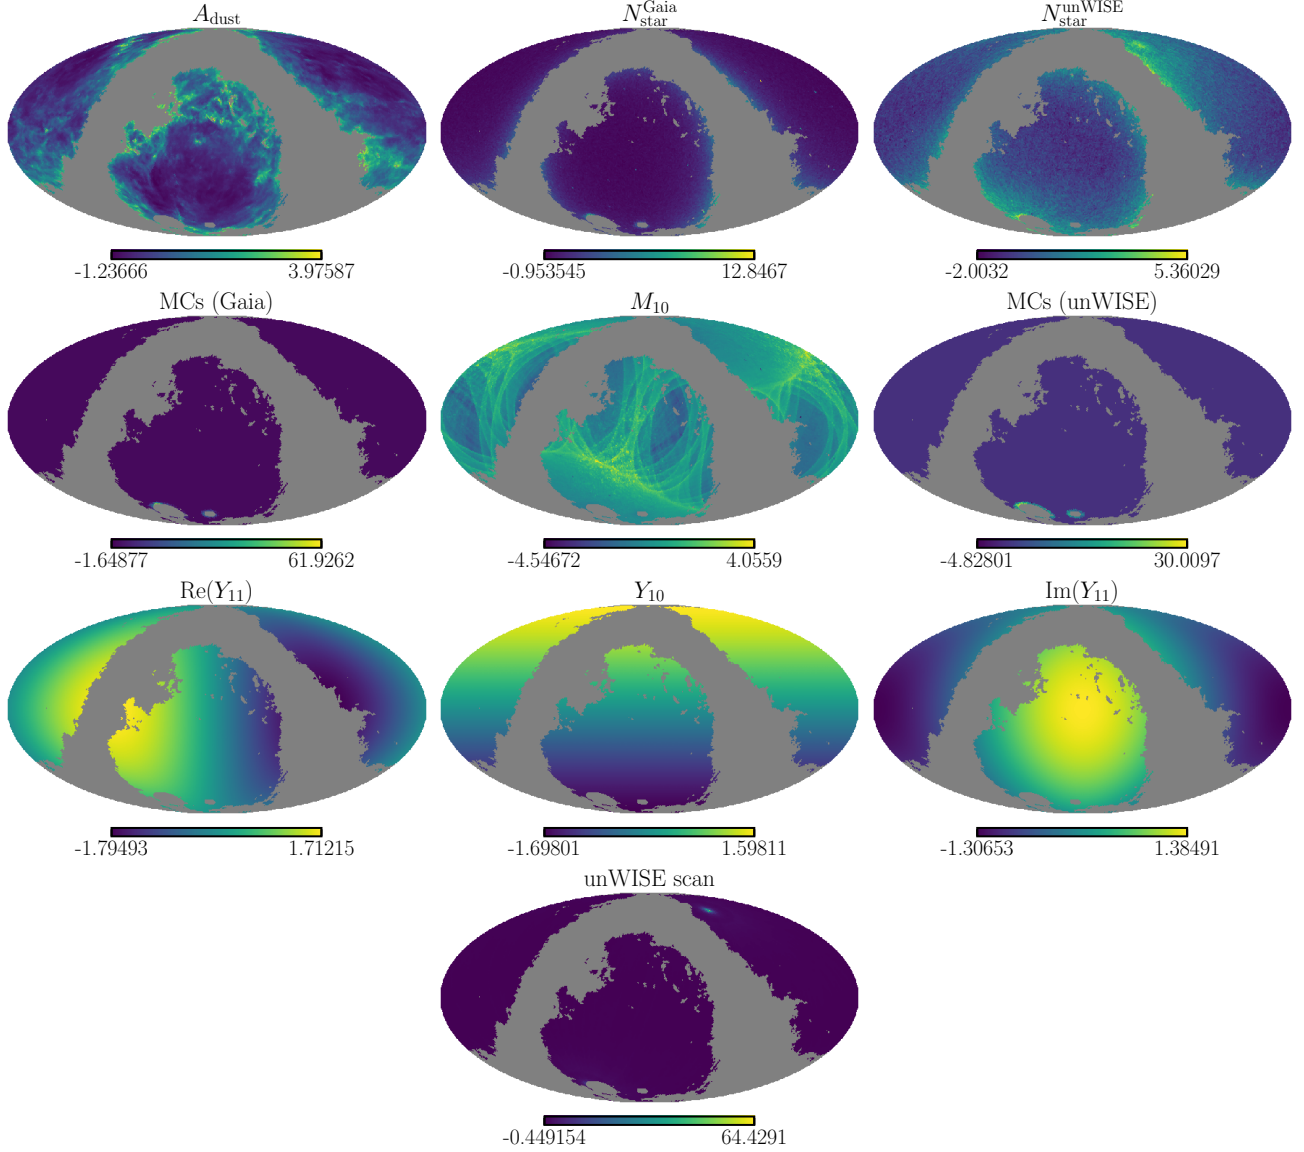


Figure A1. Systematic templates used when simulating clustering data in Sections 3.2 and 3.3, and when applying our catalogue-based deprojection formalism to *Quia* data in Section 3.5. As in G. Fabbian et al. (2026), we show for each systematic template $S(\hat{n})$ the quantity $(S(\hat{n}) - \bar{S})/\sigma_S$, where \bar{S} and σ_S are the mean and rms deviation within the footprint, respectively.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.