

## Constraints on primordial non-Gaussianity from *Quaia*

Giulio Fabbian <sup>a,b,c</sup> David Alonso <sup>d</sup> Kate Storey-Fisher <sup>e</sup> and Thomas Cornish <sup>d</sup>

<sup>a</sup> *Université Paris-Saclay, CNRS, Institut d'Astrophysique Spatiale, 91405, Orsay, France*

<sup>b</sup> *School of Physics and Astronomy, Cardiff University, The Parade, Cardiff, Wales CF24 3AA, U.K.*

<sup>c</sup> *Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge CB3 0HA, U.K.*

<sup>d</sup> *Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, U.K.*

<sup>e</sup> *Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, 452 Lomita Mall, Stanford, CA 94305, U.S.A.*

*E-mail:* [giulio.fabbian@universite-paris-saclay.fr](mailto:giulio.fabbian@universite-paris-saclay.fr),  
[David.Alonso@physics.ox.ac.uk](mailto:David.Alonso@physics.ox.ac.uk), [kstoreyf@stanford.edu](mailto:kstoreyf@stanford.edu),  
[thomas.cornish@physics.ox.ac.uk](mailto:thomas.cornish@physics.ox.ac.uk)

**ABSTRACT:** We analyse the large-scale angular clustering of quasars in the *Gaia-unWISE* quasar catalog, *Quaia*, and their cross-correlation with maps of the lensing convergence of the Cosmic Microwave Background (CMB), to constrain the level of primordial non-Gaussianity (PNG). Specifically, we target the scale-dependent bias that would be induced by PNG on biased tracers of the matter inhomogeneities on large scales. The *Quaia* sample is particularly well suited for this analysis, given the large effective volume covered, and our ability to map out the main potential sources of systematic contamination and mitigate their impact. Using the universality relation to characterise the response of the quasar overdensity to PNG ( $p_\phi = 1$ ), we report constraints on the local-type PNG parameter  $f_{\text{NL}}$  of  $f_{\text{NL}} = -20.5^{+19.0}_{-18.1}$  (68% C.L.) by combining the quasar auto-correlation and its cross-correlation with CMB lensing in two tomographic redshift bins (or  $f_{\text{NL}} = -28.7^{+26.1}_{-24.6}$  if assuming a lower response for quasars,  $p_\phi = 1.6$ ). The error on  $f_{\text{NL}}$  can be further improved if the cross-correlation between the tomographic redshift bins is included. Using the CMB lensing cross-correlations alone, we find  $f_{\text{NL}} = -13.8^{+26.7}_{-25.0}$  and  $f_{\text{NL}} = -15.6^{+42.3}_{-34.8}$  for  $p_\phi = 1$  and  $p_\phi = 1.6$  respectively. These are the tightest constraints on  $f_{\text{NL}}$  to date from angular clustering statistics and cross-correlations with CMB lensing.

**KEYWORDS:** cosmological parameters from CMBR, cosmological parameters from LSS, non-gaussianity, gravitational lensing

ARXIV EPRINT: [2504.20992](https://arxiv.org/abs/2504.20992)

---

## Contents

|          |  |           |
|----------|--|-----------|
| <b>1</b> | <b>Introduction</b>                                  | <b>1</b>  |
| <b>2</b> | <b>Data</b>  | <b>3</b>  |
| 2.1      | Quaia  | 3         |
| 2.2      | Planck   | 5         |
| <b>3</b> | <b>Methods</b>                                       | <b>5</b>  |
| 3.1      | Power spectrum measurements                          | 5         |
| 3.2      | Covariances  | 12        |
| 3.3      | Low- $\ell$ scale cuts                               | 13        |
| 3.4      | Primordial non-Gaussianity and projected fields      | 14        |
| 3.5      | Likelihood   | 19        |
| <b>4</b> | <b>Results</b>                                       | <b>21</b> |
| 4.1      | Power spectrum measurements                          | 21        |
| 4.2      | Constraints on $f_{\text{NL}}$                       | 23        |
| <b>5</b> | <b>Discussion and conclusions</b>                    | <b>29</b> |
| <b>A</b> | <b>QML results and consistency with PCL pipeline</b> | <b>31</b> |

---

## 1 Introduction

One of the most general predictions of inflation, our most successful theory describing the early stages of our Universe, is that the primordial metric fluctuations should be very close to Gaussian-distributed [1, 2]. This prediction is borne out by the statistics of the CMB temperature and polarization fluctuations, for which no significant departures from Gaussianity have been detected [3, 4]. In detail, however, different inflationary models predict such departures at different levels, and improving current constraints on the level of primordial non-Gaussianity (PNG) is vital to advance our understanding of the early Universe.

Of particular relevance is the case of “local” PNG. In this case, the level of non-Gaussianity in the primordial gravitational potential  $\Phi(\mathbf{x})$  is parametrised as

$$\Phi(\mathbf{x}) = \phi_G(\mathbf{x}) + f_{\text{NL}} \left( \phi_G^2(\mathbf{x}) - \langle \phi_G^2 \rangle \right), \quad (1.1)$$

where  $\phi_G(\mathbf{x})$  is a Gaussian random field,  $\langle \dots \rangle$  denotes an ensemble average, and  $f_{\text{NL}}$  quantifies the amplitude of the non-Gaussian contributions, caused by the non-linear  $\phi_G^2$  term. While single-field models predict a negligibly small amplitude for local PNG [5], multi-field scenarios generally can produce PNG amplitudes of order  $f_{\text{NL}} \sim 1$  [6]. A constraint on  $f_{\text{NL}}$  at this level would therefore represent a powerful probe for a wider class of inflationary models including, e.g., curvaton, modulated reheating, models with an inhomogenous end to inflation, hybrid

inflation or N-flation (see e.g. [7] and references therein). Although current constraints on  $f_{\text{NL}}$  are driven by the study of the CMB three-point function (with an uncertainty  $\sigma(f_{\text{NL}}) = 5$  [4]), significant further progress in this direction is difficult due to cosmic variance. Instead, three-dimensional tracers of the large-scale structure, such as galaxies, offer a promising alternative to reach the desired  $\sigma(f_{\text{NL}}) \lesssim 1$  threshold. The response of the small-scale power to long-wavelength fluctuations caused by the initial non-linearity in  $\Phi$  modifies the relation between the gravitational potential (and its derivatives) and the overdensity of any biased LSS tracer. In particular, local PNG produces an additional, scale-dependent, contribution to the linear bias relation, which is most relevant on large scales [8, 9]. PNG thus manifests itself in the clustering statistics of biased tracers already at the level of two-point functions, with additional information encoded in higher-order correlators. This makes PNG one of the most relevant science cases for LSS studies.

The main challenges for local PNG searches with LSS are related to the  $\sim 1/k^2$  behaviour of the scale-dependent bias contribution (where  $k$  is the Fourier-space wavenumber) [8]. Firstly, since the  $f_{\text{NL}}$  signature peaks on large scales, maximising the effective volume covered by the LSS sample under study is of critical importance. This makes high-redshift quasars strong candidates to constrain  $f_{\text{NL}}$  [10–12], although other tracers have been used and proposed, including the Cosmic Infrared Background [13], radio continuum surveys [14], and Luminous Red Galaxies (LRGs) [15, 16]. Secondly, the depth and completeness of any galaxy sample is never perfectly homogeneous, and multiple systematic effects may lead to spurious fluctuations in the observed distribution of galaxies. Since many of these sources of contamination are associated with Galactic foregrounds or variations in observing conditions, their impact is often dominant on large angular scales and can lead to a significant bias on the inferred value of  $f_{\text{NL}}$ . Typical examples of systematics manifesting as a spurious modulation of power at the large scales, and that can mimic effects of primordial non-Gaussianities, are dust extinction, given the large apparent angular size of the Milky Way, or the expected variation of calibration or observing conditions between subsequent exposures of a telescope carrying out a survey on large sky fractions [17]. Therefore, a robust measurement of  $f_{\text{NL}}$  requires a tight control over the potential sources of systematic contamination in the data, or the use of data combinations that are resilient against them.

In this paper we present constraints on  $f_{\text{NL}}$  obtained by combining the clustering of quasars in the *Gaia-unWISE* quasar sample, *Quaia* [18], with the gravitational lensing of the CMB as reconstructed from the *Planck* temperature and polarization fluctuations [19]. Specifically, we will employ a tomographic approach, measuring and analysing the angular power spectrum of the *Quaia* sources in bins of redshift, as well as their large-scale cross-correlation with CMB lensing. Although this approach discards information from large-scale radial modes in the quasar distribution, it allows for a completely self-consistent analysis while avoiding potential systematics in the spectro-photometric redshifts of the sample. Crucially, the data combination used here addresses the challenges described in the previous paragraph: the *Quaia* sample covers one of the largest effective volumes mapped by current LSS datasets. Since it is mostly based on data from space-borne experiments, the selection function of the sample and its associated spatial inhomogeneity can be accurately characterised, and the impact of residual contamination may be further mitigated through the use of the CMB lensing cross-correlations. It is also similar to other recent  $f_{\text{NL}}$  analyses that exploited the projected

clustering of photometric sources (quasars and LRGs) and CMB lensing cross-correlations, allowing us to compare our results with those using other state-of-the-art datasets.

This paper is structured as follows: section 2 presents the datasets used (*Quaia* and *Planck* CMB lensing). Section 3 describes the methods used to estimate clustering statistics and their uncertainties, and to quantify the presence of sky contamination. It also presents the model and likelihood used to interpret our data. Our results are then described in section 4, including the battery of tests carried out to quantify their robustness. We summarise these results and conclude in section 5.

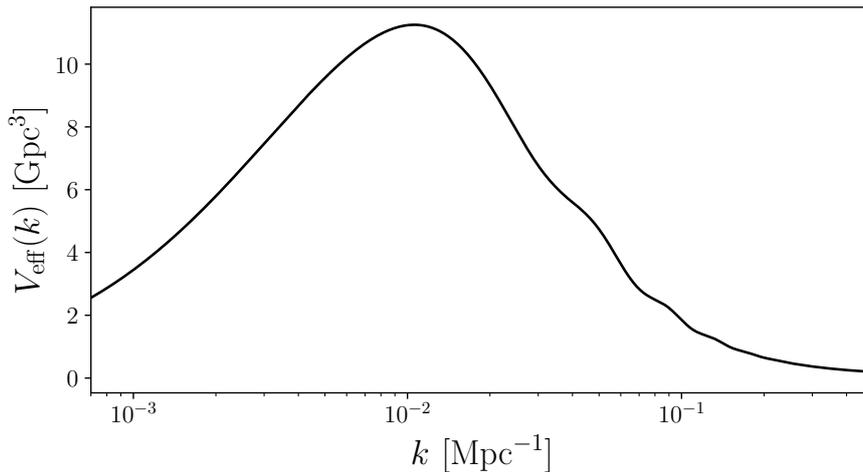
## 2 Data

### 2.1 Quaia

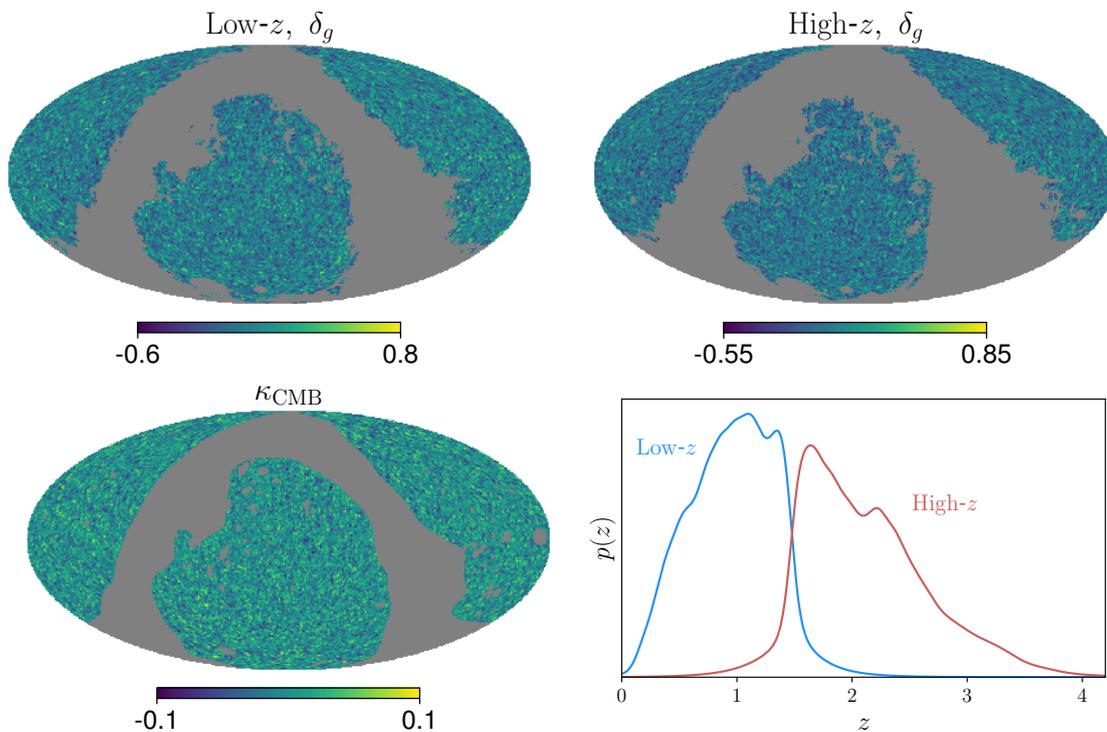
We use data from the *Gaia-unWISE* quasar sample, *Quaia* [18]. The catalog covers the entire sky and contains  $\sim 1.3$  million sources with magnitudes  $G < 20.5$ . The catalog was constructed by cross-matching the quasar candidates in the third *Gaia* data release [21] against photometry from the *unWISE* reprocessing of the Wide-Field Infrared Survey Explorer (WISE) [22], improving the level of star contamination and the quality of the estimated spectro-photometric redshifts.

Crucial to this analysis, the *Quaia* sample traces one of the largest volumes covered by existing samples, several times larger than previous datasets, such as the eBOSS quasar sample [23] (see table 1 in [18] and figure 1). According to the definition of spanning and effective volume adopted in [18], which covered a specific redshift range for the sake of comparison with other data sets, contemporary data sets such as DESI have a larger effective volume than *Quaia* (by about  $\sim 30\%$ ), thanks to their 2.5 times higher number density. However, the DESI spanning volume is still  $\sim 2.5$  times smaller than the one of *Quaia*, as it covers a reduced redshift range and a smaller sky fraction with effectively a comparable number of objects. Other samples derived from photometric full-sky surveys such as WISE (AGNs or unWISE galaxies [24, 25]) span smaller volumes and have lower redshift accuracy. This allows us to access the very large scales over which the PNG signal peaks. As highly-biased tracers of the matter density, quasars are also particularly sensitive to PNG. Furthermore, since *Quaia* is only based on data from space-borne experiments, the list of potential contaminants affecting the observed number of quasars is significantly shorter than that of ground-based surveys. This allows for a tighter control over the survey selection function, as well as the potential sources of large-scale systematics.

For our fiducial analysis we split the *Quaia* sample into two redshift bins corresponding to sources with spectro-photometric redshifts above and below the median redshift of the sample  $z_p = 1.47$ . These are the same bins used in the analyses of [26, 27], and we will refer to them as the “Low- $z$ ” and “High- $z$ ” redshift bins respectively. Splitting the catalog in these two redshift bin represents a good compromise to introduce tomographic capabilities while keeping shot noise in each redshift bin under control. In [28] we found in fact that splitting the catalog in more than two redshift bins led us to a minor loss in the precision of derived cosmological constraints. The redshift distributions for both bins were estimated by stacking the Gaussian probability density functions of each source, which was found to



**Figure 1.** Left: effective volume  $V_{\text{eff}}$  covered by *Quaia* as a function of Fourier wavenumber  $k$  (see [20] for the definition of  $V_{\text{eff}}(k)$ ).



**Figure 2.** Top: maps of the QSO overdensity of *Quaia* used in this work. Bottom: the reference *Planck* PR4 CMB lensing map used in our analysis is shown on the left while the redshift distribution of the two tomographic redshift bins  $p(z)$  is shown on the right. For visualization purpose, both the QSO and CMB lensing maps have been smoothed with a Gaussian kernel of 1 deg FWHM.

be a good approximation in [26]. The final redshift distribution  $p(z)$  for both bins is shown in figure 2. The survey selection function for each bin, which characterises the fraction of sources missed by the catalog across the sky, was estimated as described in [18]. The

catalog and survey selection function were then used to construct quasar overdensity maps as described in [26]. Since we are only interested in the clustering signal on relatively large scales, these maps were generated using the `HEALPix`<sup>1</sup> pixelisation scheme with resolution parameter  $N_{\text{side}} = 256$ , corresponding to a pixel size of  $\delta\theta_{\text{pix}} \sim 0.23^\circ$ , and they are shown in figure 2. As part of our robustness tests, we will also consider a single redshift bin containing all *Quaia* sources with  $0 < z_p \lesssim 4.5$ .

To quantify the presence of residual contamination in these maps, we consider template maps for 10 different potential systematics. These include: i) a dust extinction map constructed by [29], ii) two stellar contamination maps built by selecting random sources from *Gaia* and *unWISE*, iii) templates tracking the scanning patterns of *Gaia* and *unWISE*, iv) stellar density maps in the regions around the Large and Small Magellanic Clouds, and v) three dipole components, corresponding to the spherical harmonic functions with  $\ell = 1$ ,<sup>2</sup> ( $Y_{1,0}, Y_{1,\pm 1}$ ). These systematic templates are shown in figure 3. Note that the *Quaia* selection function was constructed as a non-linear combination of the systematic templates i)–iv) listed above. As we describe in section 3.1, these templates (as well as the dipole components) are linearly deprojected from the galaxy overdensity maps as part of the power spectrum estimation pipeline, to further remove residual contamination.

## 2.2 Planck

We use maps of the CMB lensing convergence from the *Planck* experiment. Specifically, we use the lensing maps based on data from the *Planck* “PR4” release,<sup>3</sup> generated with the `NPIPE` pipeline [19] (see figure 2). In our fiducial analysis we use the Generalised Minimum Variance (GMV) map, constructed from a joint inverse-variance Wiener filtering of the CMB temperature and polarization maps accounting for inhomogeneous noise. In addition to this map, we use the so-called “No-TT” map, first introduced in [28], and constructed from a quadratic estimator that uses both temperature and polarization but downweights *TT* correlations. This map is thus largely immune to the presence of unpolarized extragalactic foregrounds, such as the Cosmic Infrared Background (CIB) and the thermal Sunyaev-Zel’dovich effect. Contamination from the CIB is of particular relevance for the high-redshifts probed by *Quaia*, although no conclusive evidence of such contamination was found in [28].

These maps were preprocessed as described in [28]: the convergence harmonic coefficients and the analysis mask are rotated to Equatorial coordinates, and downgraded to our base resolution (including only multipoles up to  $\ell_{\text{max}} = 3N_{\text{side}} - 1$  to avoid aliasing effects).

## 3 Methods

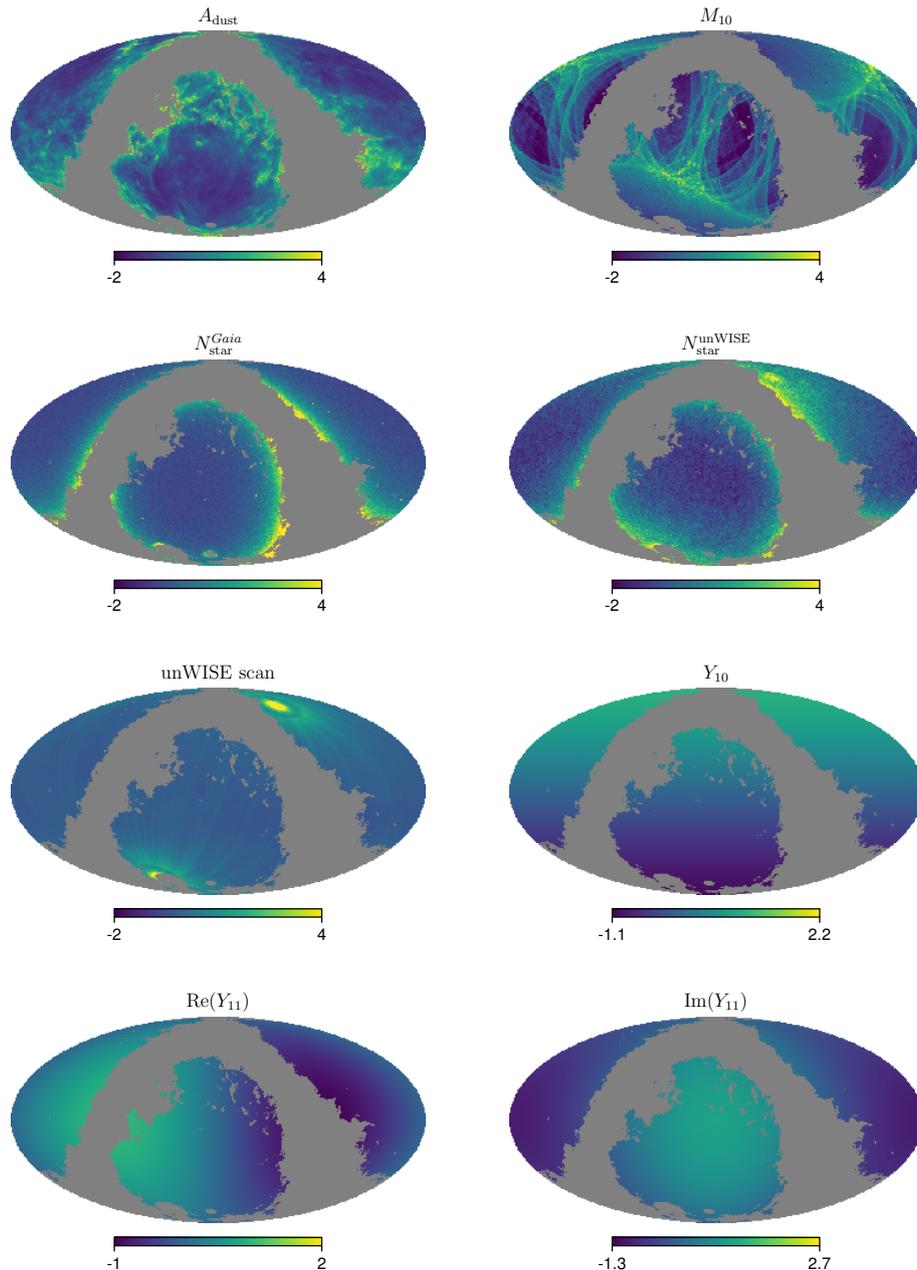
### 3.1 Power spectrum measurements

The effects of PNG on 2D projected statistics are most prominent on the largest angular scales. The estimate of 2-point statistics on those scales is complicated not only by the presence of observational systematics but also by partial sky coverage and survey geometry effects.

<sup>1</sup><https://healpix.sourceforge.io>.

<sup>2</sup>The dipole in the *Quaia* sample is being investigated in a separate work [30].

<sup>3</sup>[https://github.com/carronj/planck\\_PR4\\_lensing](https://github.com/carronj/planck_PR4_lensing).



**Figure 3.** Systematic templates used to construct the *Quia* selection function, and to quantify the impact of large-scale contamination. In all cases, we show the local deviation with respect to the mean as a fraction of the root-mean squared (rms) in each quantity, with mean and rms defined within the shown footprint. I.e. for systematic template  $S(\hat{\mathbf{n}})$ , we show  $(S(\hat{\mathbf{n}}) - \bar{S})/\sigma_S$ , with  $\bar{S}$  and  $\sigma_S$  the mean and rms deviation within the footprint. In addition to these 8 templates, we use two additional maps of the source density in *Gaia* and *unWISE* in the region around the Magellanic clouds, not shown here.

As such we investigated the use of two different pipelines that make different assumptions and simplifications when accounting for instrumental systematics and masking effects on the measured power spectrum.

### 3.1.1 Pseudo- $C_\ell$ pipeline

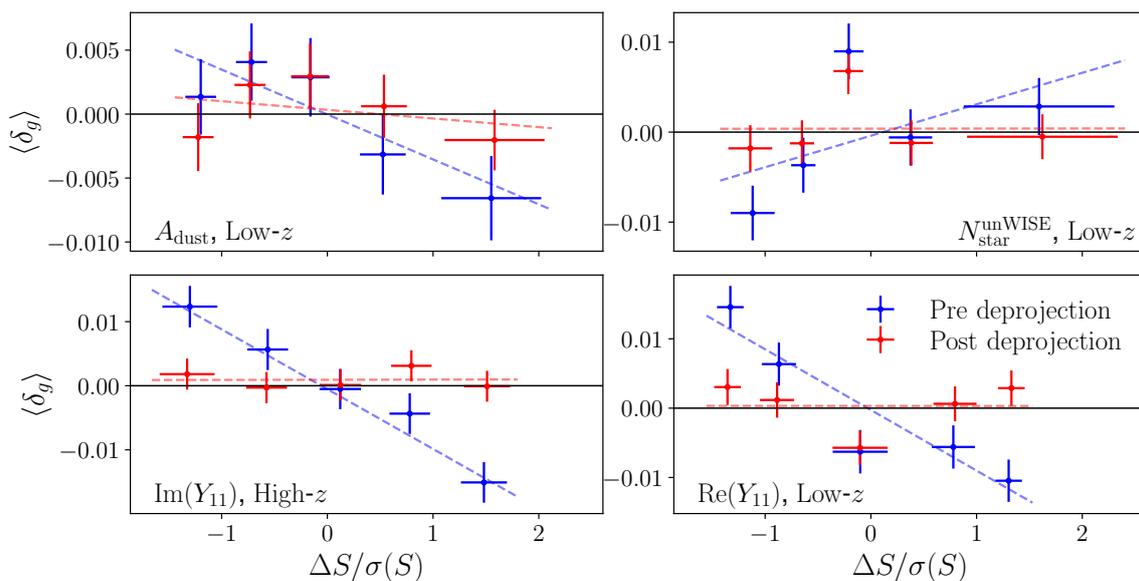
The first pipeline estimates the angular power spectrum with a state-of-the-art pipeline based on a pseudo- $C_\ell$  (PCL) power spectrum estimation method [31] as implemented in the publicly available code `NaMaster` [32]. The PCL estimator is in principle suboptimal, particularly on large angular scales, for data with steep power spectra or complex sky footprints [33]. However, in [27] we showed that this approximation is in fact quite accurate for *Quia*, since the sample is shot-noise dominated across most scales.

To suppress the presence of contamination from systematics, and to determine the angular scales over which our measurements may be safely used, we linearly deproject the angular fluctuations of the 10 templates described in section 2.1 from the QSO overdensity maps using the approach described in [32, 34]. Such mode-deprojection removes modes in the map that are proportional to any of the systematic templates deprojected or their linear combination without distinguishing between overlapping signal modes that behave in a similar way. As such, the resulting power spectrum is biased due to this mode loss. While for typical cases when the number of systematics templates is low and the resulting mode loss is negligible on small angular scales, for large scale analyses this bias needs to be accounted for; otherwise, this mode loss could bias the inferred value of  $f_{\text{NL}}$ .

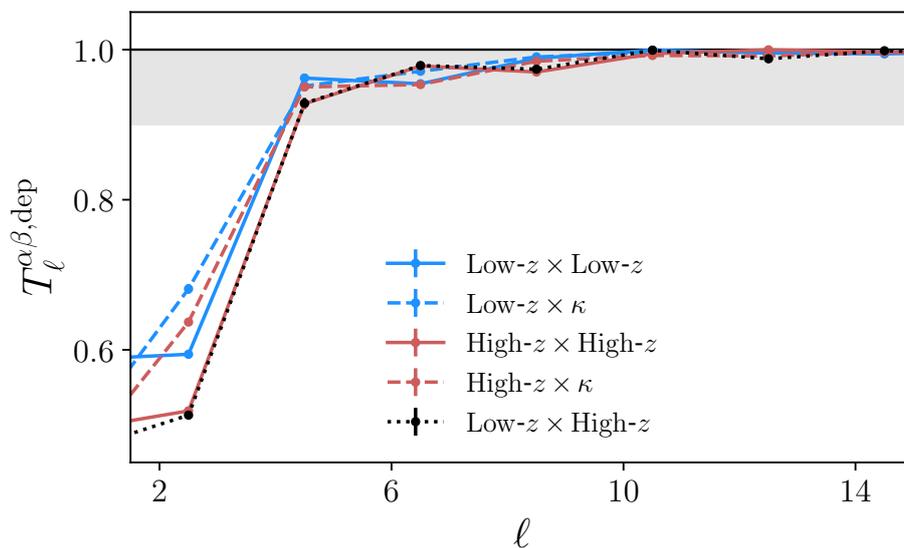
While the deprojection bias term can be corrected analytically if the true underlying power spectrum is known [32, 34], we followed the approach of [27] and estimated the deprojection bias as an effective multiplicative transfer function from simulated mock data. This is similar to the procedure used in CMB analyses, when the final maps have been subjected to non-trivial filtering operations (either in the time, pixel, or Fourier domains) to mitigate the impact of specific systematics (e.g. correlated noise, scan-synchronous signals, etc.) [35–37]. This approach has the advantage of being more robust against the assumed “true” power spectrum used when estimating the deprojection bias analytically, which in the case of studies involving large, cosmic-variance-dominated scales (as is our case) may not be known in detail. This method is discussed in detail in [38]. As discussed in [27], we found that the mode loss is important only on scales  $\ell \lesssim 6$  and is generally smaller for the cross-correlations with CMB lensing than for the quasar auto-correlations. The effect is, nevertheless, small, in particular on scales used in our cosmological analysis. In fact, the mode loss induced by deprojection affects mainly multipoles below  $\ell \simeq 4$  at  $\sim 40\%$  level, and at better than 5% level for  $\ell \gtrsim 6$ , with the large loss on largest scales induced by the deprojection of the  $Y_{1,0}$  and  $Y_{1,1}$  templates needed to account for the *Quia* dipole. In figure 5 we show the transfer functions for all the relevant spectra used in the following.

In section 4 we will quantify the impact of the choices made regarding contaminant deprojection on our final constraints. While in section 3.1.2 we describe an alternative analysis pipeline that is immune to transfer function effects and that was used to further validate the robustness of our analysis. We also verified that deprojecting mean-subtracted quadratic versions of the systematics templates did not change our conclusions.

To further illustrate the impact of deprojection in this analysis, figure 4 shows the correlation between the galaxy overdensity field and the systematic templates before and after deprojection (blue and red points, respectively). We show particular combinations of redshift bin and systematic templates that best show the effects of deprojection. In each



**Figure 4.** Map-level empirical correlation between a selection of contaminant templates and the galaxy overdensity  $\langle \delta_g \rangle$  in the *Quia* data. Results are shown for overdensity maps without any contaminant deprojection (blue) and after linear deprojection (red). Deprojection eliminates any residual linear correlation between the data and these particular contaminant templates. To illustrate this point, the best-fit linear trend, found by fitting a linear model to the data in each individual plot, is shown as dashed lines (although note that linear deprojection applies this to the full linear subspace spanned by all the systematic templates). The lower panels show that we find a strong dipole component, which is also removed via deprojection.



**Figure 5.** Systematics deprojection transfer function for the different auto and cross-correlations relevant for this work.

panel, the points with error bars show the mean value of  $\delta_g$  (and its standard deviation) across the map after binning the map pixels by the local value of a given systematic template  $S$ . The results are shown as a function of the fluctuation in the systematic map from its mean across the footprint as a fraction of its standard deviation. The dashed lines in the figure show the best-fit linear trend describing the data. By construction, deprojection eliminates any linear correlation between  $\delta_g$  and any map described by a linear combination of the systematic templates. We can also see that the residual linear correlation with the templates that were used to construct the *Quaia* selection function are relatively small (top panels). In turn, the data shows a clear linear trend against the dipole components. This correlation may be due to the local kinematic dipole in the observed density of distant quasars [30, 39], or a consequence of residual large-scale systematic contamination not captured by any of the other templates. Regardless of its origin, we choose to remove this dipole component altogether in order to suppress its contribution to the power spectrum uncertainties at higher multipoles via mode-coupling (see section 3.1.2).

We mask out pixels in the *Quaia* maps for which the corresponding *Quaia* selection function is  $< 0.5$  and use the selection function as pixel weights, and used the *Planck* lensing mask as mask for  $\kappa$ . For reference, our pixel cut based on the selection function reduces the observable sky area from  $\sim 87\%$  to  $61\%$  and  $52\%$  for the Low- $z$  and High- $z$  redshift bins respectively. We binned all power spectra into bin of width  $\Delta\ell = 2$  in the range  $\ell \leq 40$ ,  $\Delta\ell = 5$  in the range  $40 \leq \ell < 60$ , and  $\Delta\ell = 30$  in the range  $\ell \geq 60$ .

As a last notable ingredient of this pipeline, we accounted for the Monte Carlo correction of the normalization of the  $\kappa$  field that arises from the presence of inhomogeneous noise of the *Planck* data in our reference footprint. For this purpose we made use of the lensing reconstruction simulations made available with the PR4 lensing maps and followed the procedure discussed in [11, 26]. Such normalization becomes very large for the first two bandpowers but is smaller than 10% for smaller angular scales.

### 3.1.2 Quadratic maximum likelihood pipeline

As an alternative, we developed a second pipeline based on a quadratic maximum likelihood (QML) power spectrum estimation [40]. This pipeline allows us to test many of the assumptions inherent to the first pipeline. Firstly, QML power spectrum estimation methods effectively act as quadratic functions of the pixels and operates in pixel space. As such the estimate of the mode-coupling coefficients caused by the presence of a mask is not subject to the potential impact of finite band-limits in harmonic space.

Secondly, if the covariance matrix of the pixels is known, and if the field under study is Gaussian-distributed, the resulting estimator is manifestly optimal: all information available in the data is encoded in the estimated spectrum. By construction, the PCL estimator assumes a diagonal pixel covariance, and may be suboptimal when recovering spectra with a red tilt, exhibiting significant correlations across distant pixels. This is often particularly important on the largest angular scales. As in our analysis we are dealing with auto and cross-power spectra estimates we used the QML power spectrum estimators of [41] as implemented in the publicly available code `xQML`.<sup>4</sup>

<sup>4</sup><https://gitlab.in2p3.fr/xQML/xQML>.

This implementation writes down the quadratic estimate of the power spectrum  $\hat{C}_\ell^{AB,\text{QML}}$  of two pixelized maps of scalar fields  $\mathbf{d}_A, \mathbf{d}_B$  assuming that the cross-covariance between both fields is negligible when compared with their auto-covariances. In this case, the estimators for all auto- and cross-correlations are decoupled and may be calculated independently as

$$\hat{C}_\ell^{AB,\text{QML}} \equiv \sum [W^{-1}]_{\ell\ell'} \hat{y}_{\ell'}^{AB}. \quad (3.1)$$

The elements of eq. (3.1) are:

$$\hat{y}_{\ell'}^{AB} = \mathbf{d}_A^T \mathbf{E}_\ell \mathbf{d}_B - b_\ell^{AB}, \quad (3.2)$$

$$\mathbf{E}_\ell = \frac{1}{2} (\mathbf{C}^{AA})^{-1} \mathbf{P}_\ell (\mathbf{C}^{BB})^{-1}, \quad (3.3)$$

$$W_{\ell\ell'} = \frac{1}{2} \text{Tr} \left[ (\mathbf{C}^{AA})^{-1} \mathbf{P}_\ell (\mathbf{C}^{BB})^{-1} \mathbf{P}_{\ell'} \right], \quad (3.4)$$

where  $\mathbf{P}_\ell$  is a matrix involving the  $\ell$ -th Legendre polynomial:

$$\mathbf{P}_{\ell,ij} = \frac{2\ell+1}{4\pi} P_\ell(\cos \theta_{ij}), \quad (3.5)$$

with  $\theta_{ij}$  being the angular separation between pixels  $i$  and  $j$ . In the previous equations,  $\mathbf{C}^{XY}$  is the pixel covariance between fields  $X$  and  $Y$ , which can be decomposed into signal and noise components

$$\mathbf{C}^{AB} \equiv \langle \mathbf{d}_A \mathbf{d}_B^T \rangle = \mathbf{S}^{AB} + \mathbf{N}^{AB}. \quad (3.6)$$

$b_\ell^{AB} = \text{Tr} [\mathbf{E}_\ell \mathbf{N}^{AB}]$  is the noise bias estimate. In the case of cross-spectra  $A = \delta_g, B = \kappa$  the noise is uncorrelated between maps and as such the cross noise covariance matrix  $\mathbf{N}^{AB} = 0$  and  $b_\ell^{AB} = 0$ . In turn,  $C_\ell^{gg}$  is affected by shot-noise. We estimated the shot noise bias computing the average auto spectrum of 100 jackknife null maps of QSO overdensities derived from random equal splits of the *Quaia* catalog.<sup>5</sup> This allows us to minimize the residual noise biases arising from mismatches between the assumed data covariance and the true data one.

The computational time of the QML estimator roughly scales as the number of the observed pixel to the third power, driven by the complexity of the brute force matrix inversion involved in the algorithm. As such, it would be very hard to carry out the analysis at the same resolution of the PCL pipeline. We therefore reprocessed the *Quaia* catalog to produce maps of the QSO overdensity on a  $N_{\text{side}} = 16$  HEALPix grid. The *Planck* CMB  $\kappa$  estimates are given in terms of harmonic coefficients. To have a low resolution version of the CMB  $\kappa$  map, we resampled on the same grid only the Fourier modes  $\ell \leq 3N_{\text{side}} = 48$  (the maximum multipole supported by the grid) convolving them with the corresponding pixel window function. We reprocessed the simulated  $\kappa$  map of the *Planck* data release and use those to compute  $\mathbf{N}^{\kappa\kappa}$ . We then used the publicly available random realization of the *Quaia* catalog (which includes unclustered objects drawn from a uniform Poisson sampling of sky locations according to the selection function) to compute the expected number of objects in each pixel  $p$ ,  $\bar{n}_p$ , and assumed a diagonal QSO noise covariance matrix  $\mathbf{N}_{pp'}^{gg} = \delta_{pp'} / \bar{n}_p$

<sup>5</sup>We also verified that the noise-debiased  $C_\ell^{gg}$  agrees with the average of the cross-spectra of the aforementioned maps derived of the data splits which is inherently noise-bias free.

so that pixels where the selection function averaged in the pixel is lower are downweighted. We generated the signal covariance  $\mathbf{S}$  from the fiducial  $C_{\ell s}$  of *Quaia* sources and  $\kappa$  adopting  $f_{\text{NL}} = 0$ , the fiducial cosmology of the *Planck* FFP10 simulations' cosmology, and the best-fit bias model and redshift distributions of [26].

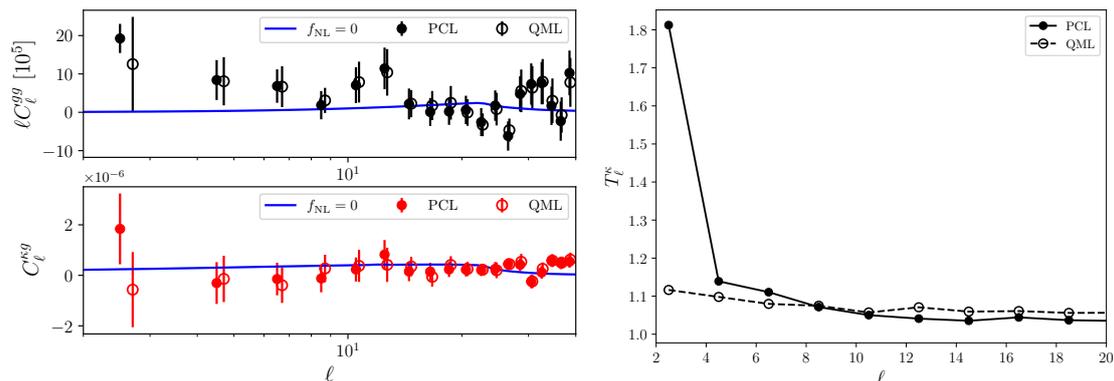
In the QML pipeline we treated the presence of systematics using mode deprojection via template marginalization [33, 42]. In this context the map-level systematics are removed by assigning infinite variance to contaminating templates, making therefore the power spectrum estimate insensitive to the templates and is equivalent to marginalizing over the contamination amplitude of each of them. Mathematically, this is equivalent to the linear deprojection approach implemented in the PCL estimator [34]. For a list of  $k$  systematics templates  $\mathbf{t}_k$  this involves modifying the covariance matrices entering equations (3.3), (3.4) as:

$$\mathbf{C} \rightarrow \mathbf{C} + \sum_k \beta_k \mathbf{t}_k \mathbf{t}_k^T, \quad (3.7)$$

where the sum is done over the number of templates and  $\beta_k$  are arbitrarily large numbers. We included the same templates employed in the PCL case downgraded to the same resolution of the data. This includes the three dipole templates, which were incorporated by adding a term proportional to the  $\ell = 1$  Legendre polynomial  $\Delta C_{ij} \propto P_1(\cos \theta_{ij})$  with a large prefactor. We also considered an additional template made by the *Planck* Zodiacal light map but we found this template to have no effect for the systematic marginalization.

Contrary to the PCL implementation of linear deprojection, the marginalisation approach does not bias the estimator, since the mode loss inherent to deprojection is consistently incorporated in the mode-coupling matrix (eq. (3.4)). We verified this empirically by applying the estimator, including mode-deprojection, to simulated data that do not contain systematic effects. In the QML pipeline we adopted the same binning for overlapping multipoles with the PCL pipeline to minimize the computational cost which becomes quickly important, while keeping enough resolution in the power spectrum estimate. However, the final power spectra combined the QML  $C_\ell$  estimates up to  $\ell \leq 16$  and use the estimates derived on a pseudo- $C_\ell$  power spectrum estimation on smaller scales, similarly to what is done for the PCL pipeline. The  $\ell = 16$  transition was chosen as we observed that the error of the QML and PCL estimates became effectively the same for the mock data on small angular scales when the latter did not deproject the dipole moments (see discussion section 3.2). We also did not include any deprojection transfer function for the PCL data points as it is negligible on those scales. For the computation of the power spectra across all scales we used a single galactic mask given by the product of the two selection function and lensing mask employed in the PCL pipeline.

Similarly to the PCL pipeline, we accounted for the Monte Carlo correction of the normalization of the CMB  $\kappa$  field (denoted  $T_\ell^\kappa$  hereafter for consistency with [27]) using the same procedure but processing the *Planck* PR4 lensing simulations in the same way we did with the data. As the evaluation of the transfer function requires the computation of the auto-spectrum of the input signal of the  $\kappa$  maps as well as their cross correlation with the corresponding reconstructed field  $\kappa_{\text{rec}}$  for the same realization, we set the  $\mathbf{N}^{\kappa\kappa} = 0$  and computed  $\mathbf{S}^{\kappa\kappa}$  from the theoretical power spectrum  $C_\ell^{\kappa\kappa}$  while we computed the  $\mathbf{C}^{\kappa_{\text{rec}}\kappa_{\text{rec}}}$  from the PR4 simulations directly. As the PR4 simulation include a large dipolar component due to the kinematic motion of our local frame with respect to the CMB, we included a



**Figure 6.** Left: comparison between PCL and QML pipeline used for this work for the largest angular scales where the treatment of systematics, estimation and covariance differs the most. We show the auto- (top) and cross-correlation with CMB lensing (bottom) for a single tomographic bin for simplicity. The shaded area identifies scales not included in the cosmological analysis. Right: comparison of the CMB lensing Monte Carlo correction for the QML and PCL pipelines. The QML estimate here includes a marginalization over the large  $\ell = 1$  mode in the CMB lensing field that better controls its leakage to higher multipoles.

marginalisation of the  $P_1$  template for computing both of these cross-spectra. Contrary to the PCL case, the CMB lensing transfer function is always well below 10% at all angular scales because the QML estimator minimizes the impact of coupling of the dipolar modulations of the lensing field connected to the local kinematic dipole. We show for comparison the different transfer functions in figure 6.

Because the QML pipeline is computationally expensive and less flexible, in the following we report the main results obtained for the PCL pipeline. However, we carried out extensive consistency tests in particular for the case involving one single tomographic bin that we discuss in appendix A. In figure 6 we show a comparison of the estimated large scale angular power spectrum for both pipelines. Despite the radical differences in terms of systematics treatment and assumptions implemented in the estimators and in the covariance modelling, the agreement of the measurements at all scales is very good, as it can be seen in the figure, and as are the final constraints.

### 3.2 Covariances

For the two fiducial power spectrum estimation pipelines, we adopt different fiducial models for the data covariances. For the PCL pipeline we take a similar approach to [27] and estimate the covariance matrix of the data from the power spectra of 1000 Gaussian simulations of the quasar overdensity and CMB convergence maps, based on the measured power spectra of the data as input to such simulations. This automatically incorporates the additional variance caused by the presence of large scale systematics in the *Quia* maps. This approach neglects in particular the noise inhomogeneity in both  $\kappa$  and *Quia* maps, as well as the fact that the input power spectra are noisy, non-smooth realizations of the underlying true power spectra. As in [27], we verified that both these effects have a negligible impact on the power spectrum uncertainties and on the final constraints. As a consistency test we also used an analytical

Gaussian covariance based on the Narrow Kernel Approximation [43], as implemented in `NaMaster`, which is less noisy but also in principle less accurate for the largest angular scales.

For the QML pipeline, we also adopted a simulation-based approach but revised some of the simplifications of the PCL implementation. We used the official *Planck* PR4 simulated reconstructed  $\kappa$  maps, which contain the effect of correlated inhomogeneous noise of the *Planck* satellite measurements, and generated correlated Gaussian realizations of the *Quaia* maps starting from the signal realizations of the PR4 simulations and a fiducial cross-correlation model assuming the same *Planck* cosmology of the PR4 simulations, the eBOSS QSO bias model of [44] and the measured redshift distribution of the *Quaia* sources. We then added to this signal map a shot-noise realization based on the random *Quaia* catalogues that are publicly available,<sup>6</sup> and a dipolar modulation with the same amplitude and direction observed in the *Quaia* maps. As the QML pipeline adopts a mix of pseudo- $C_\ell$  and QML estimates for the largest scales, we resampled each realizations on the native `HEALPix` grid resolution of the PCL and QML pipelines and processed these maps as we did with the data. The resulting simulated data vectors were then used to compute the data power spectrum covariance. The resulting covariance matrix incorporates more accurately the correlation between bandpowers and any other non-Gaussian contribution to the covariance. However they do not include the extra variance due to the excess power induced by the *Quaia* systematics. To overcome this issue in particular for  $\ell \leq 30$  we assumed a Gaussian approximation for these multipoles but calibrated the diagonal of the covariance measuring the effective number of degrees of freedom  $\nu_b$  for each multipole bin  $b$  using the mock data as [31]

$$\text{Var} \left( C_b^{AB, \text{QML}} \right)_s = \frac{1}{\nu_b} \left( \langle C_b^{AA, \text{QML}} \rangle_s \langle C_b^{BB, \text{QML}} \rangle_s + \langle C_b^{AB, \text{QML}} \rangle_s \right) \quad (3.8)$$

where the  $s$  subscript denotes quantities averaged over the number of realizations of our simulated data set and  $A, B \in \{g, \kappa\}$ . We then used the estimated  $\nu_b$  to rescale the covariance by the power spectrum measured on data as

$$\text{Var} \left( C_b^{AB, \text{QML}} \right) = \frac{1}{\nu_b} \left( C_b^{AA, \text{QML}} C_b^{BB, \text{QML}} + C_b^{AB, \text{QML}} \right), \quad (3.9)$$

so that potential non-Gaussian bandpower correlation is preserved and the extra variance due to observational systematics is accounted for. We note that for scales where the systematics in the auto-correlation do not clearly dominate, we adopted the largest between the analytical and the Monte Carlo error bar. For both the PCL and QML pipeline we applied a correction factor to the inverse covariance matrix to account for the finite number of simulations used to calculate the covariance, as described in [45]. As we start from the PR4 reconstructed  $\kappa$  realizations we are limited to 400 realizations for the QML pipeline (compared to the 1000 realisations used for the PCL estimator).

### 3.3 Low- $\ell$ scale cuts

The set of simulations produced for the QML pipeline covariance estimation was also used to define the large-scale cuts adopted in our analysis. For each of the simulated maps we

---

<sup>6</sup>The public random catalogues contain 10 times more objects than the *Quaia* sample, but we only draw a number equal to the size of the data sample in each realisation.

analysed the realizations with both the PCL and QML pipelines and applied systematics subtraction (via deprojection or marginalization) as we implemented for the real data. We then calculated the distribution of differences in each power spectrum multipole bin before and after systematics deprojection, for both  $C_\ell^{gg}$  and  $C_\ell^{kg}$ , and for both analysis pipelines. We then compared these simulated distributions to the differences observed in the real data after systematics deprojection. Since our simulations do not include systematic effects, any observed change in power due to deprojection is random — arising from chance correlations and mode loss. We therefore decided to exclude power spectrum bins where the observed shift in the data had a probability to exceed (PTE) below 5% compared to the simulation-based distribution. We computed the PTE using a two-sided distribution to account for the sign of the shift, given that systematics deprojection preferentially reduces power.

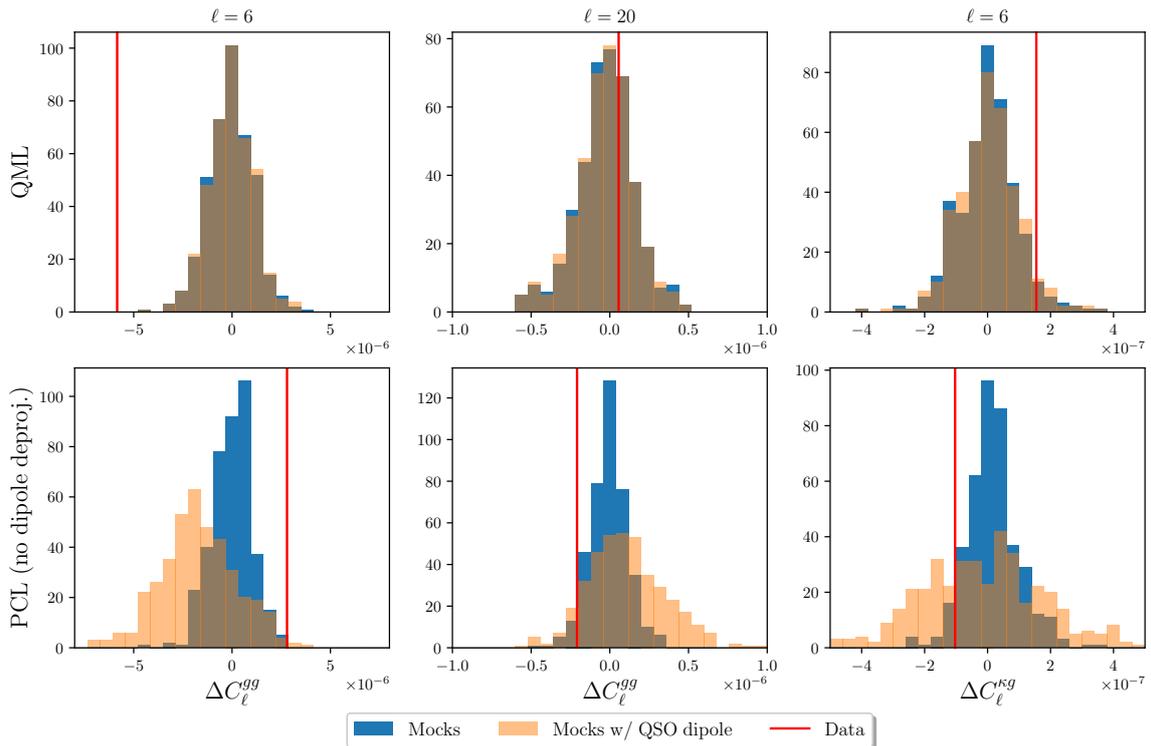
For both the QML and PCL pipeline we found that scales  $\ell \leq 15$  shift in an anomalous way for  $C_\ell^{gg}$  and as such are affected by systematics. We therefore adopted this conservative scale cut in the analysis of the QSO auto correlation as baseline. On the other hand, we found that the shift in all the  $C_\ell^{kg}$  bandpowers were consistent with random fluctuations. As such we decided to include all the angular scales down to  $\ell_{\min} = 2$  in the analysis of the CMB lensing cross-correlation. We show a summary of these results in figure 7.

This scale cut definition checks for anomalous shifts while removing part of the sample variance, as opposed to an assessment simply based on the amplitude of the power shifts induced by systematics subtraction in relation to the statistical uncertainties. Nevertheless, as we show in section 4, systematic deprojection shifts the bandpowers by an amount that is subdominant with respect of the statistical error on the scales used in this analysis. We note that we defined this set of scale cuts looking at the overall QSO clustering based on one single redshift bin. However, in the final cosmological constraints we assessed the impact of the scale cuts modifying the  $\ell_{\min}$  included in the analysis and show that we recover consistent results and thus the scale cuts applied are robust.

In figure 7 we also show the importance of properly treating the large dipole of the QSO distribution in our pipelines. For this purpose we show results obtained using mock data including a dipole in the QSO distribution as well as results based on dipole-free simulations. As it can be seen in the bottom panel for the case of the PCL pipeline, an untreated dipole generates additional variance and biases in the observed distributions of power spectra shifts induced by systematics deprojection compared to the case where mock data are dipole free. Including the  $\ell = 1$  spherical harmonics templates in the systematic deprojection removes biases and extra variance associated to the QSO dipole with the price of having to model this through a transfer function. In the case of the QML pipeline shown in the top panel conversely, the dipole marginalization effectively makes the presence of the QSO dipole irrelevant. We will test the robustness to all these aspects of the analysis in the following.

### 3.4 Primordial non-Gaussianity and projected fields

We focus our analysis on the impact of local-type PNG on the clustering of biased tracers. In the presence of local PNG, the primordial gravitational potential can be characterised as in eq. (1.1). As initially shown by [8], local PNG modifies the clustering of galaxies on large scales. Focusing only on large scales, we can model the three-dimensional quasar



**Figure 7.** Examples of the difference in the measured power spectra induced by systematics deprojection for different multipole bins and spectra. In solid red we show the measurements on data, while the histograms show the expected distribution of such shifts estimated from simulations with no systematics. Orange histograms show distributions evaluated adding a dipole in the simulated QSO maps while the blue histograms are derived from simulations with no QSO dipole. The two rows show the results for the QML (top) and PCL (bottom) pipeline. For  $\ell = 6$  the probability of data and simulations being consistent is  $< 1\%$ , hinting for the presence of systematics in the data (left). For  $\ell = 20$  (middle) shifts are consistent with random fluctuations.  $\ell = 15$  is the largest scales for which the observed shifts are consistent with random fluctuations at the 95% C.L. and we therefore select this as our scale cut. For  $C_\ell^{kg}$  (right) no sign of systematics can be observed even for scales contaminated in the QSO auto-correlation. For illustrative purposes, we did not remove the QSO dipole for the PCL panels. As can be seen, accounting for the QSO dipole is crucial to avoid biases and extra variance induced by the leakage of the dipole through the observational masks. We did so in our fiducial PCL analysis deprojecting the  $Y_{10}, Y_{\pm 1}$  templates.

overdensity in Fourier space at the tree level as

$$\Delta_g(\mathbf{k}, z) = \left[ b(z) + \frac{2\delta_c}{\alpha(k, z)} (b(z) - p_\phi) f_{\text{NL}} \right] \Delta_m(\mathbf{k}, z), \quad (3.10)$$

where  $\delta_c \equiv 1.686$  is the collapse threshold [46], and  $b(z)$  is the galaxy bias, which we explicitly show is redshift-dependent. The scale-dependent factor  $\alpha(k)$  allows us to connect the 3D matter overdensity  $\Delta_m$  with the primordial gravitational potential as

$$\alpha(k, z) = \frac{2c^2 k^2 T(k) D(z)}{3\Omega_m H_0^2}, \quad (3.11)$$

where  $T(k)$  is the matter transfer function, normalised to  $T(k \rightarrow 0) \rightarrow 1$ , and  $D(z)$  is the linear growth factor for matter perturbations, normalised to  $(1+z)D(z) \rightarrow 1$  at high redshifts, during the matter-dominated era.  $\Omega_m$  is the fractional amount of non-relativistic matter at  $z = 0$ , and  $H_0$  is the expansion rate today. Finally,  $p_\phi$  characterises the modification to the response function of the galaxy overdensity to a large-wavelength fluctuation caused by  $f_{\text{NL}}$ .  $p_\phi$  is assumed to range from  $p_\phi = 1$ , corresponding to the spherical collapse prediction for a Universal mass function [47], to  $p_\phi \sim 1.6$  for recently-merged haloes (as found in  $N$ -body simulations) [47–49]. The latter value was argued to be potentially more appropriate for quasars [47].

In this work we analyse the *projected* clustering of quasars. Within the linear bias model presented above, the observed 2D overdensity  $\delta_g(\hat{\mathbf{n}})$  in the sky position  $\hat{\mathbf{n}}$  is given by [11]

$$\delta_g(\hat{\mathbf{n}}) = \delta_g^{\text{int}}(\hat{\mathbf{n}}) + \delta_g^{\text{PNG}}(\hat{\mathbf{n}}) + \delta_g^{\text{RSD}}(\hat{\mathbf{n}}) + \delta_g^{\text{mag}}(\hat{\mathbf{n}}). \quad (3.12)$$

Here,  $\delta_g^{\text{int}}(\hat{\mathbf{n}})$  is the contribution from the intrinsic clustering of galaxies in the absence of PNG:

$$\delta_g^{\text{int}}(\hat{\mathbf{n}}) \equiv \int dz p(z) b(z) \Delta_m(\chi \hat{\mathbf{n}}, z), \quad (3.13)$$

where  $p(z)$  is the redshift distribution of the sample and  $\chi$  is the comoving radial distance.  $\delta_g^{\text{PNG}}(\hat{\mathbf{n}})$  is the scale-dependent contribution due to PNG

$$\delta_g^{\text{PNG}}(\hat{\mathbf{n}}) \equiv \int dz p(z) f_{\text{NL}} (b(z) - p_\phi) \Delta_g^{\text{PNG}}(\chi \hat{\mathbf{n}}, z), \quad (3.14)$$

with

$$\Delta_g^{\text{PNG}}(\mathbf{k}, z) \equiv \frac{2\delta_c}{\alpha(k, z)} \Delta_m(\mathbf{k}, z). \quad (3.15)$$

Finally,  $\delta_g^{\text{RSD}}(\hat{\mathbf{n}})$  and  $\delta_g^{\text{mag}}(\hat{\mathbf{n}})$  are the contributions due to redshift-space distortions (RSD) and magnification bias:

$$\delta_g^{\text{RSD}}(\hat{\mathbf{n}}) = - \int \frac{dz}{H(z)} p(z) \partial_{\parallel} v_{\parallel}, \quad \delta_g^{\text{mag}} = \int d\chi W_{\text{mag}}(\chi) \left[ \frac{1}{2} \nabla_{\perp}^2 (\phi + \psi) \right] (\chi \hat{\mathbf{n}}, z(\chi)), \quad (3.16)$$

where  $v_{\parallel}$  is the peculiar velocity field along the line of sight,  $\partial_{\parallel}$  is the line-of-sight gradient,  $\psi$  and  $\phi$  are the metric potential in the Newtonian gauge,  $\nabla_{\perp}^2$  is the Laplacian operator on the plane perpendicular to the line of sight, and

$$W_{\text{mag}}(\chi) \equiv \chi \int_{z(\chi)}^{\infty} dz' p(z') (5s(z') - 2) \left( 1 - \frac{\chi}{\chi(z')} \right). \quad (3.17)$$

Here  $s(z)$  is the logarithmic slope of the cumulative number of sources as a function of magnitude for the sample under study. RSD and magnification are the leading contributions to the fluctuations in the observed number counts of sources caused by perturbations to the comoving volume elements and source fluxes in the presence of inhomogeneities [50–52]. The remaining so-called “relativistic effects” are subdominant and can lead to a small increment or decrement of power on large scales, commensurate with a PNG signal of  $|f_{\text{NL}}| \sim 1$ . While this may become important with future, more sensitive searches, the effect of these terms can be safely neglected in this analysis [52–54].

We also study the cross-correlation of the *Quia* sample with maps of the CMB lensing convergence  $\kappa(\hat{\mathbf{n}})$ , given by

$$\kappa(\hat{\mathbf{n}}) = \int_0^{\chi_{\text{LSS}}} d\chi \chi \left(1 - \frac{\chi}{\chi_{\text{LSS}}}\right) \left[\frac{1}{2} \nabla_{\perp}^2(\phi + \psi)\right](\chi \hat{\mathbf{n}}, z(\chi)), \quad (3.18)$$

where  $\chi_{\text{LSS}}$  is the distance to the last-scattering surface, at  $z = 1089$ .

The theoretical prediction for the cross-spectrum between any of these two projected fields can be written as

$$C_{\ell}^{\alpha\beta} = \frac{2}{\pi} \iiint d\chi_1 d\chi_2 dk k^2 P(k, z(\chi_1), z(\chi_2)) \Delta_{\ell}^{\alpha}(k, \chi_1) \Delta_{\ell}^{\beta}(k, \chi_2), \quad (3.19)$$

where  $P(k, z_1, z_2)$  is the unequal-time matter power spectrum, and  $\Delta_{\ell}^{\alpha}(k, \chi)$  is the line-of-sight source function for field  $\alpha$ . We approximate the unequal-time power spectrum assuming perfect correlation between different times

$$P(k, z_1, z_2) = \sqrt{P(k, z_1) P(k, z_2)}, \quad (3.20)$$

which is sufficiently accurate [55, 56]. Within the linear bias model used here, and assuming a flat  $\Lambda$ CDM model to relate  $v_r$  and  $\phi + \psi$  with  $\Delta_m$ , the source function for the observed quasar overdensity is [9, 51]

$$\Delta_{\ell}^g(k, \chi) = \Delta_{\ell}^{\text{int}}(k, \chi) + \Delta_{\ell}^{\text{PNG}}(k, \chi) + \Delta_{\ell}^{\text{RSD}}(k, \chi) + \Delta_{\ell}^{\text{mag}}(k, \chi), \quad (3.21)$$

$$\Delta_{\ell}^{\text{int}}(k, \chi) \equiv b(z) H(z) p(z) j_{\ell}(k\chi), \quad (3.22)$$

$$\Delta_{\ell}^{\text{PNG}}(k, \chi) \equiv (b(z) - p) f_{\text{NL}} \frac{2\delta_c}{\alpha(k, z)} H(z) p(z) j_{\ell}(k\chi), \quad (3.23)$$

$$\Delta_{\ell}^{\text{RSD}}(k, \chi) \equiv -H(z) p(z) f(z) j_{\ell}''(k\chi), \quad (3.24)$$

$$\Delta_{\ell}^{\text{mag}}(k, \chi) \equiv \frac{\ell(\ell+1)}{(k\chi)^2} \frac{3}{2} H_0^2 \Omega_m (1+z) W_{\text{mag}}(\chi) j_{\ell}(k\chi), \quad (3.25)$$

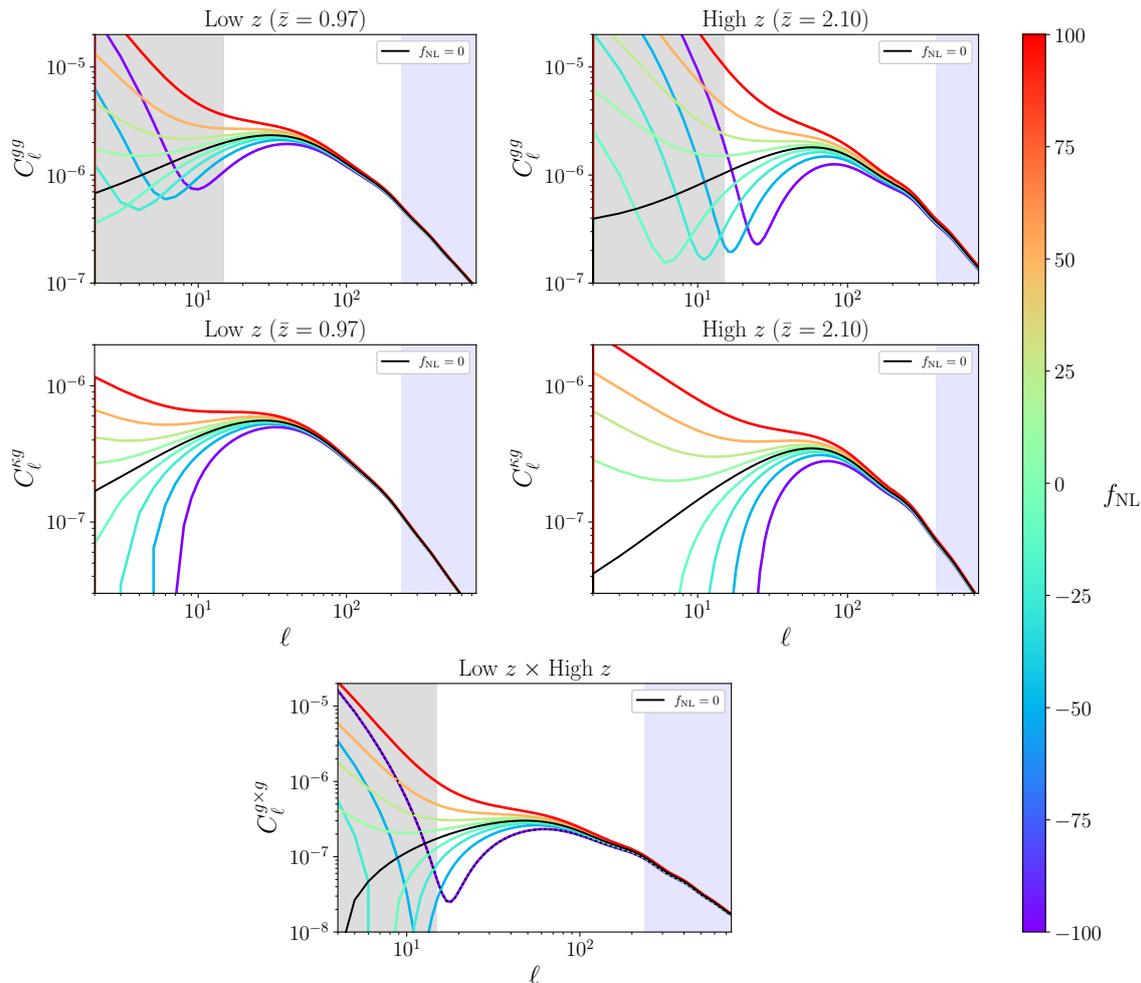
and the CMB lensing source function is

$$\Delta_{\ell}^{\kappa}(k, \chi) = \frac{\ell(\ell+1)}{(k\chi)^2} \frac{3}{2} H_0^2 \Omega_m (1+z) \chi \left(1 - \frac{\chi}{\chi_{\text{LSS}}}\right) \Theta(\chi < \chi_{\text{LSS}}). \quad (3.26)$$

Here,  $j_{\ell}(x)$  is the spherical Bessel function of order  $\ell$ ,  $\Theta(< x)$  is the Heaviside function,  $f(z)$  is the growth rate, and  $H(z)$  is the expansion rate at redshift  $z$ . In all cases above,  $z$  implicitly denotes the redshift at the comoving distance  $\chi$  in the background.

As eq. (3.19) shows, calculating the angular power spectrum involves, in principle, a triple integral over  $\chi_1$ ,  $\chi_2$ , and  $k$ . This can be significantly sped up using the Limber approximation [57, 58]. This approximation becomes less accurate on large scales, where most of the sensitivity to  $f_{\text{NL}} \neq 0$  lies, and therefore we only adopt it at  $\ell > 200$ , where it is very accurate [58, 59]. For smaller values of  $\ell$  we use the non-Limber calculation as implemented in the Core Cosmology Library (CCL) [60], specifically using the FKEM technique [61], first presented in [62]. Nevertheless, the impact of the Limber approximation was found to be highly subdominant for the *Quia* sample on the scales explored here in (see figure 3 in [27]).

Figure 8 shows the impact of local PNG on the angular power spectra used in this analysis for different values of  $f_{\text{NL}}$ . It is interesting to note the distinct effect of large negative



**Figure 8.** Impact of local primordial non-Gaussianity on the *Quiaia* auto- and cross-correlations used in this analysis. All predictions shown here were calculated assuming the fiducial *Planck* cosmology, the *Quiaia* redshift distribution and a PNG response parameter  $p_\phi = 1$ . Results are shown for the for the quasar auto-correlations, the cross-correlations with CMB lensing (middle), and for the cross-bin power spectrum (bottom). The shaded bands identify the low- $\ell$  (grey) and high- $\ell$  (blue) scale cuts adopted in our fiducial analysis setup.

values of  $f_{\text{NL}}$  on the auto-correlations and the CMB lensing cross-correlations. Since the scale-dependent bias appears squared in the auto-correlation, negative values of  $f_{\text{NL}}$  lead to a power *enhancement* on the largest scales, and a suppression on intermediate scales. In turn, negative  $f_{\text{NL}}$  always suppresses the cross-correlation. The latter therefore allows us to better distinguish between positive and negative values of  $f_{\text{NL}}$ , thanks to its monotonic trend with  $f_{\text{NL}}$  and the possibility to access the largest scales with minimal systematic contamination. In order to better disentangle the sign of  $f_{\text{NL}}$  from the auto-correlation, it is important to observe the characteristic signature of a potential negative  $f_{\text{NL}}$  value at the smallest possible angular scales, which have a lower sample variance and less systematic contamination compared to its largest scales. This can only be ensured by tomographic

analyses like the one carried out in this work using high redshift samples, as it can be seen comparing the top left and right panel of figure 8.

### 3.5 Likelihood

To derive constraints on  $f_{\text{NL}}$  we use a Gaussian power spectrum-level likelihood of the form

$$-2 \log p(\mathbf{d}|\vec{\theta}) = (\mathbf{d} - \mathbf{t}(\vec{\theta}))^T \mathbf{C}^{-1} (\mathbf{d} - \mathbf{t}(\vec{\theta})) + K. \quad (3.27)$$

Here,  $\mathbf{d}$  is the data vector that contains our power spectrum measurements.  $\mathbf{t}$  is the theoretical prediction for  $\mathbf{d}$ , dependent on a set of model parameters  $\vec{\theta}$ .  $\mathbf{C}$  is the covariance matrix of  $\mathbf{d}$ , and  $K$  is an irrelevant normalization constant.

The data vector  $\mathbf{d}$  contains our measurements of the quasar auto-correlation, and the quasar-lensing cross-correlation for each redshift bin. Each power spectrum is used over a different range of scales  $(\ell_{\text{min}}, \ell_{\text{max}})$ , defined as follows: the large-scale cut is chosen, as described in section 3.1, to ensure that our measurements are free from systematic contamination in the quasar density maps. Specifically, in our fiducial analysis we use  $\ell_{\text{min}}$  for the quasar-lensing cross-correlations, and  $\ell_{\text{min}} = 15$  in the quasar auto-correlations. The small-scale cut  $\ell_{\text{max}}$  is chosen to ensure the validity of the linear bias model adopted in this analysis. Specifically, we define a different scale cut for each redshift bin as  $\ell_{\text{max}} = k_{\text{max}} \bar{\chi}$ , where  $k_{\text{max}} = 0.07 \text{ Mpc}^{-1}$ , and  $\bar{\chi}$  is the comoving distance to the mean redshift of each bin. This corresponds to  $\ell_{\text{max}} = 236$  and  $390$  in the low- and high-redshift bins respectively.

Modelling the power spectrum likelihood as a multivariate Gaussian is only an approximation: even in the simplest case of Gaussian maps, the associated power spectra would follow a Wishart distribution [63]. As we demonstrated in [27], the distribution of the quasar-lensing cross-correlation measurements is well described by a Gaussian at all scales. This is also the case for the quasar auto-correlation on the smaller scales ( $\ell \geq 15$ ) over which we use it to avoid clustering systematics, due to the central limit theorem.

To construct  $\mathbf{t}(\vec{\theta})$  we follow the model described in section 3.4. The redshift distribution of each bin was constructed as described in section 2.1. The redshift-dependent bias in the  $i$ -th redshift bin is modelled as

$$b(z) = A_i b_*(z), \quad (3.28)$$

where  $b_*(z)$  is the fitting function for the bias of quasars in the sample studied by [44]:

$$b_*(z) = 0.278((1+z)^2 - 6.565) + 2.393, \quad (3.29)$$

and  $A_i$  is a free parameter of the model (one for each redshift bin). To quantify the impact of this choice of redshift evolution, we will also repeat our analysis for the alternative model presented in [28] (P24 hereafter), which provides a better fit for the *Quaia* sample:

$$b_*(z) = 1.26 \frac{D(0)}{D(z)}, \quad (3.30)$$

where  $D(z)$  is the linear growth factor for the best-fit  $\Lambda$ CDM cosmology found by *Planck* [64]. Note that the cosmological analysis of the *Quaia* sample presented in [26] found cosmological parameters that are slightly different from (but consistent with) the best-fit *Planck* values.

Our fiducial analysis will use the auto- and cross-correlations from the two redshift bins described in section 2.1. Thus, our model is described by three free parameters:  $\vec{\theta} \equiv \{A_1, A_2, f_{\text{NL}}\}$ , where  $A_1$  and  $A_2$  are the bias amplitudes in the low- and high-redshift bins, respectively. We sample this likelihood using the affine-invariant Markov chain Monte Carlo (MCMC) algorithm implemented in `emcee` [65].

In order to significantly speed up the evaluation of the theoretical prediction described above, we approximate the dependence of the angular power spectra on these parameters as follows: first, we write the source function for the quasar intrinsic clustering (eq. (3.22)) in the  $i$ -th redshift bin as:

$$\Delta_\ell^{\text{int},i}(k, \chi) = A_i \bar{b}_{*,i} \frac{b_*(z)}{\bar{b}_{*,i}} H(z) p_i(z) j_\ell(k\chi(z)) \equiv b_i \bar{\Delta}_\ell^{\text{int},i}(k, \chi), \quad (3.31)$$

which defines the parameter-independent source function  $\bar{\Delta}_\ell^{\text{int},i}(k, \chi)$ , and where  $\bar{b}_{*,i}$  is the mean value of the fiducial bias function  $b_*(z)$  over the redshift distribution of the  $i$ -th bin

$$\bar{b}_{*,i} \equiv \int dz p_i(z) b_*(z). \quad (3.32)$$

We have also defined  $b_i \equiv A_i \bar{b}_{*,i}$ . Then, we approximate the PNG source function (eq. (3.23)) as

$$\begin{aligned} \Delta_\ell^{\text{PNG},i}(k, \chi) &= (b(z) - p) f_{\text{NL}} \frac{2\delta_c}{\alpha(k, z)} H(z) p_i(z) j_\ell(k\chi) \\ &\simeq (b_i - p) f_{\text{NL}} \frac{b_*(z) - p}{\bar{b}_{*,i} - p} \frac{2\delta_c}{\alpha(k, z)} H(z) p_i(z) j_\ell(k\chi) \\ &\equiv (b_i - p) f_{\text{NL}} \bar{\Delta}_\ell^{\text{PNG},i}(k, \chi), \end{aligned} \quad (3.33)$$

which defines the parameter-independent PNG source function  $\bar{\Delta}_\ell^{\text{PNG},i}(k, \chi)$ . Evidently, this approximation is only appropriate if the bias amplitude parameters  $A_i$  are close to 1. Since in all the cases explored here the best-fit values of  $A_i$  never deviate from this value by more than  $\sim 10\%$ , we find that this approximation is indeed sufficiently accurate. Specifically, for a 10% deviation in  $A_i$ , this approximation incurs a very small error that leads to a difference in the  $\chi^2$  of the associated model of  $\Delta\chi^2 \simeq 0.01$ , which is negligible.

The main advantage of this approximation is that it allows us to write the quasar auto-correlation in bin  $i$  as:

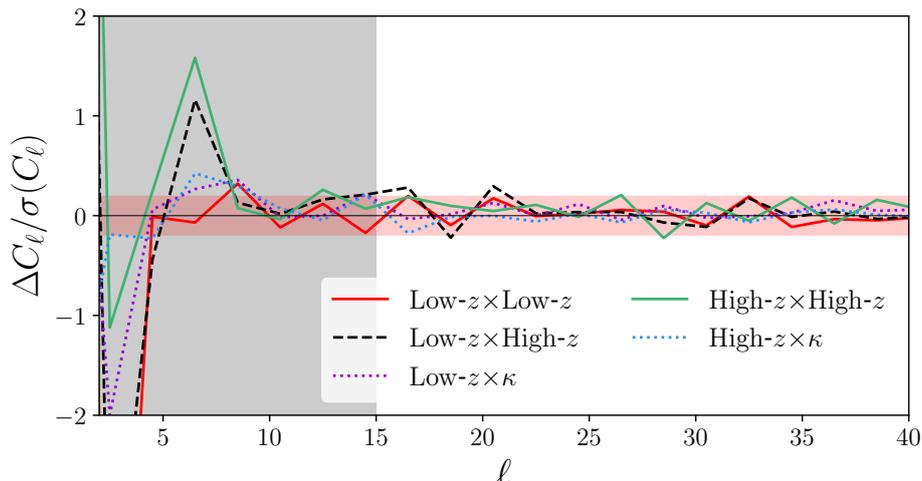
$$\begin{aligned} C_\ell^{gg} &= b_i^2 \bar{C}_\ell^{\text{int,int}} + 2b_i(b_i - p) f_{\text{NL}} \bar{C}_\ell^{\text{int,PNG}} + 2b_i \bar{C}_\ell^{\text{int,(RSD+mag)}} \\ &\quad + (b_i - p)^2 f_{\text{NL}}^2 \bar{C}_\ell^{\text{PNG,PNG}} + 2(b_i - p) f_{\text{NL}} \bar{C}_\ell^{\text{PNG,(RSD+mag)}} \end{aligned} \quad (3.34)$$

$$+ \bar{C}_\ell^{\text{(RSD+mag),(RSD+mag)}} \quad (3.35)$$

where all barred  $C_\ell$ s (e.g.  $\bar{C}_\ell^{\text{int,int}}$ ) denote power spectra involving the parameter-independent source functions. Likewise, the quasar-lensing cross-correlation is

$$C_\ell^{g\kappa} = b_i \bar{C}_\ell^{\text{int},\kappa} + (b_i - p) f_{\text{NL}} \bar{C}_\ell^{\text{PNG},\kappa} + \bar{C}_\ell^{\text{(RSD+mag),\kappa}}.$$

Although the barred power spectra are slow to calculate, particularly without the Limber approximation, they need only be computed once before sampling the likelihood (since they do not depend on the model parameters). Once this is done, the likelihood is a simple polynomial function of the parameters, which can be evaluated in  $\sim 0.3$  milliseconds on a standard laptop. This allows us to run Monte Carlo chains with tens of thousands of samples at a negligible computational cost.



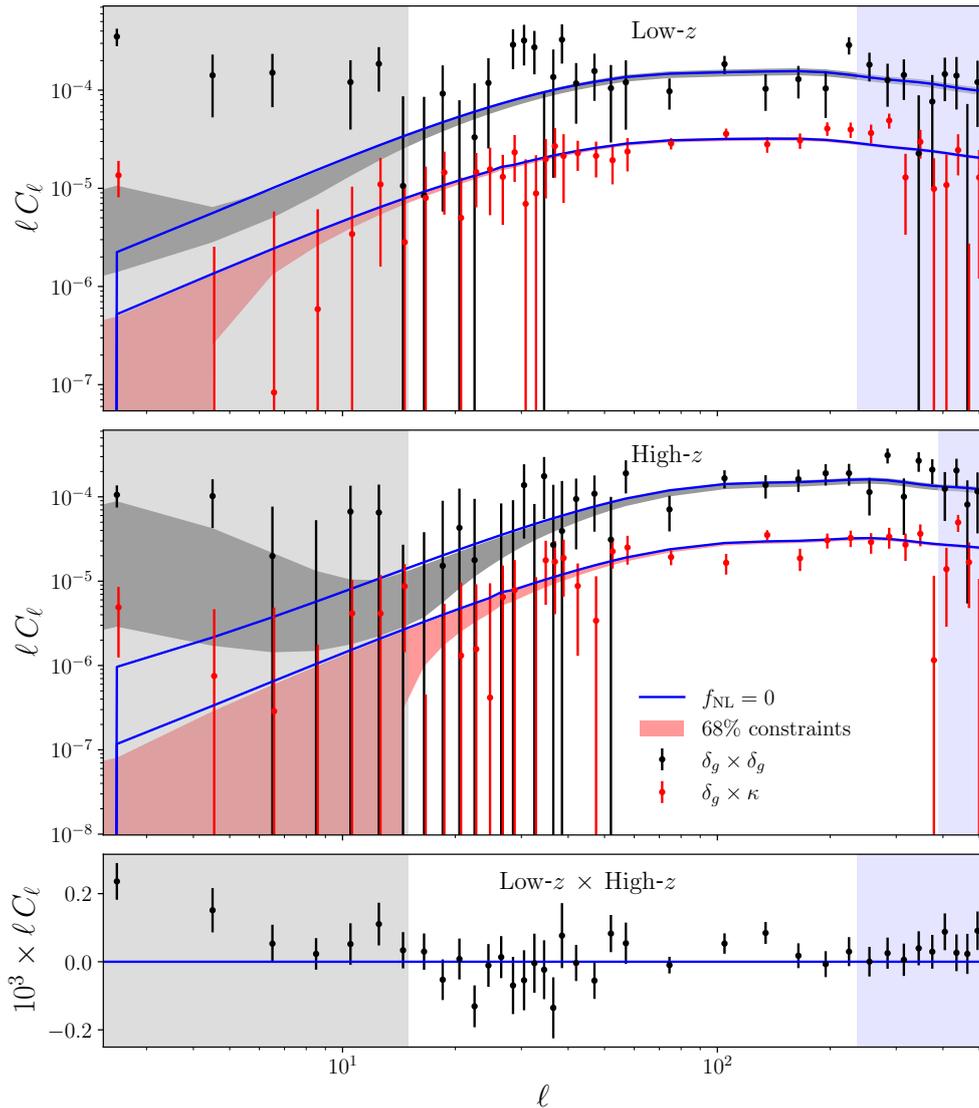
**Figure 9.** Difference between power spectra measured with and without linear contaminant deprojection as a fraction of the statistical uncertainties for  $\ell < 40$ . Results are shown for the quasar auto-correlations (solid lines), including the cross-bin correlation (black dashed lines — not used in the fiducial analysis), as well as the quasar — CMB lensing cross-correlations (dotted lines). On scales  $\ell \geq 15$  (gray shaded area), corresponding to the fiducial scale cut applied to the quasar auto-correlations, the differences remain below 20% of the measurement errors (shown as the red-shaded horizontal band).

## 4 Results

### 4.1 Power spectrum measurements

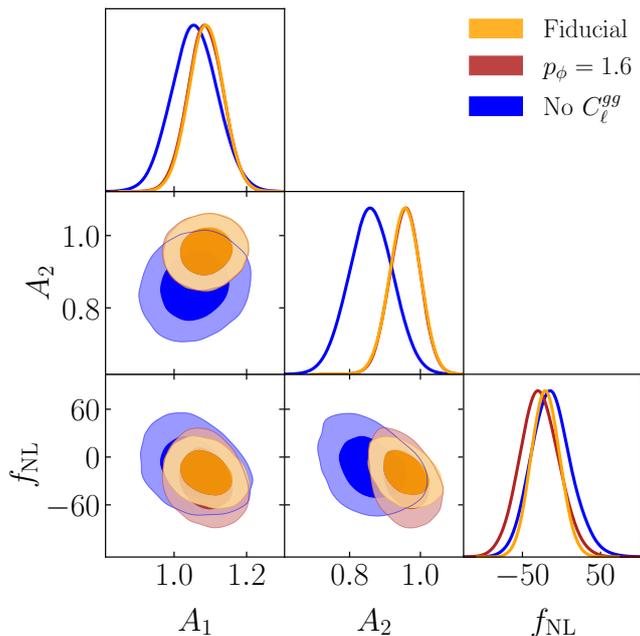
We measure the auto-correlation of the quasar overdensity in the two redshift bins, as well as the cross-bin correlation, and their cross-correlation with the CMB lensing map using the pseudo- $C_\ell$  approach described in section 3.1. Crucially for this analysis, we mitigate the impact of observational systematics on large scales by linearly deprojecting the ten contaminant templates presented in section 2.1. These trace the scanning patterns of both *Gaia* and *unWISE*, stellar abundance, and dust extinction, as well as three dipole moments. As discussed in 3.1, we find that the *Quaia* overdensity maps are most strongly correlated with the dipole, likely due to a combination of the kinematic dipole and residual large-scale contamination. Deprojecting the dipole components has a beneficial effect in the power spectrum measurements, reducing the variance caused by mode coupling on the smallest multipoles.

We correct for the loss of power caused by deprojection through the transfer function approach discussed in section 3.1. To quantify the potential presence of residual contamination in the data after this procedure, we compared the power spectrum measurements with and without contaminant deprojection. The difference between both measurements on large scales is shown in figure 9 as a fraction of the statistical uncertainties. We find that, after  $\ell = 15$ , the differences are subdominant (less than 20% of the statistical uncertainties), further justifying our choice of large-scale cut for the quasar auto-correlations. It is interesting to note that the cross-correlation between both bins (not used in the fiducial analysis) exhibits deviations between the cases with and without deprojection above  $0.2\sigma$  extending to smaller scales.



**Figure 10.** Measurements, after contaminant deprojection, of the quasar auto-correlation (black points) and their CMB lensing cross-correlation (red points) for the Low- $z$  and High- $z$  redshift bins (top and middle panels, respectively), and the cross-correlation between the redshift bins (bottom panel). The shaded red and black bands show the 68% confidence level (C.L.) bounds on the power spectra within our 3-parameter model, including  $f_{\text{NL}}$  and the linear bias in both redshift bins. On large scales, this is dominated by the uncertainty on  $f_{\text{NL}}$ , corresponding to  $f_{\text{NL}} = -20.2 \pm 18.4$  (constraints found when including data from both redshift bins). The best-fit prediction is shown as solid lines. The blue lines show the theoretical prediction for the best-fit bias parameters but fixing  $f_{\text{NL}} = 0$ . The blue shaded vertical band marks the high- $\ell$  scale cut applied for both  $C_\ell^{g\kappa}$  and  $C_\ell^{g\kappa}$ , while the gray vertical band shows the large-scale cut imposed on the quasar auto-correlations.

Our measurements of the quasar auto-correlations and cross-correlations with CMB lensing are shown in figure 10. The shaded bands show the  $1\sigma$  bounds for the theoretical prediction within our 3-parameter model ( $\{A_1, A_2, f_{\text{NL}}\}$ , see section 3.5). The vertical blue band marks the small scales excluded from our analysis of both  $C_\ell^{g\kappa}$  and  $C_\ell^{g\kappa}$ , while the gray



**Figure 11.** Constraints on the 3 parameters of our model: the bias amplitudes in each redshift bin  $A_1$  and  $A_2$ , and the PNG amplitude  $f_{\text{NL}}$ . Constraints are shown for our fiducial analysis, using both auto- and cross-correlations and assuming a universality relation for the bias response to PNG, with  $p_\phi = 1$  (yellow contours), for the case of a merger-driven population ( $p_\phi = 1.6$ , red contours), and for an analysis using only the CMB lensing cross-correlations (blue contours).

vertical band shows the large-scale cut used for  $C_\ell^{gg}$ . Within this range of scales, the model is able to provide a good fit to the data, with a best-fit  $\chi^2 = 117.9$  for a total of 109 degrees of freedom, corresponding to a probability-to-exceed (PTE) of 0.26.

## 4.2 Constraints on $f_{\text{NL}}$

The constraints on  $f_{\text{NL}}$  and the two bias amplitude parameters  $A_1$  and  $A_2$ , using our fiducial analysis choices, are shown in figure 11 (yellow contours). We find a constraint on the PNG amplitude parameter of

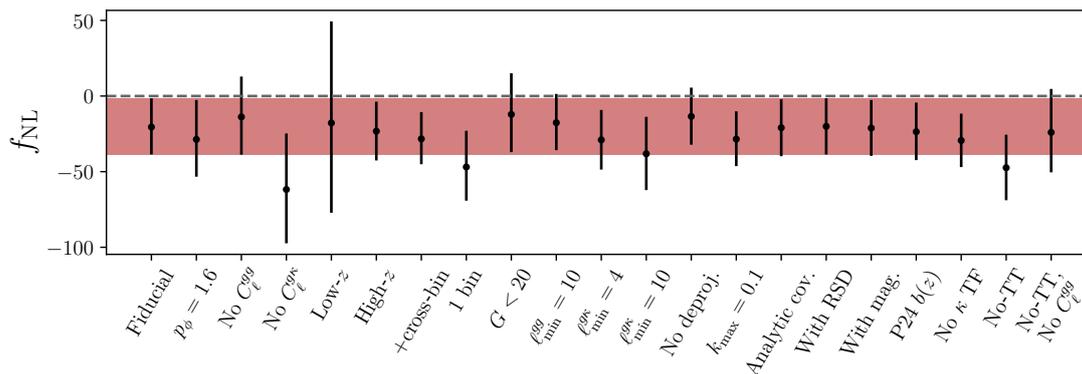
$$f_{\text{NL}} = -20.5_{-18.1}^{+19.0} \quad (p_\phi = 1, \text{ 68\% C.L.}), \quad (4.1)$$

compatible with  $f_{\text{NL}} = 0$  at the  $1.1\sigma$  level. The galaxy bias parameters are also in relatively good agreement (within  $\sim 2\sigma$  of unity) with the fiducial bias model. As mentioned in the previous section, the resulting best-fit model is a good fit to the data.

Figure 11 also shows the constraints found assuming a PNG response corresponding to a population dominated by recent mergers as may be the case for quasar populations, with  $p_\phi = 1.6$ , as discussed in 3.4. In this case, the constraints are degraded slightly, as expected, with

$$f_{\text{NL}} = -28.7_{-24.6}^{+26.1}, \quad (p_\phi = 1.6, \text{ 68\% C.L.}), \quad (4.2)$$

in similar agreement with  $f_{\text{NL}} = 0$ . Finally, the figure shows the constraints found using only the quasar-CMB lensing cross-correlations, which are more robust to the presence of residual



**Figure 12.** Constraints on  $f_{\text{NL}}$  obtained in our fiducial analysis (shaded horizontal band), and adopting alternative analysis strategies that test the robustness of our results. The numerical values shown here are listed in table 1, and the different tests are described in the main text.

contamination in the galaxy overdensity maps. The resulting constraints, assuming  $p_\phi = 1$ , are

$$f_{\text{NL}} = -13.8^{+26.7}_{-25.0}, \quad (p_\phi = 1, \text{ no } C_\ell^{gg}, \text{ 68\% C.L.}), \quad (4.3)$$

We observe a marginal shift towards  $f_{\text{NL}} = 0$  with respect to our fiducial result, combined with a  $\sim 40\%$  increase in the statistical uncertainties. Repeating this analysis assuming  $p_\phi = 1.6$  we obtain

$$f_{\text{NL}} = -15.6^{+42.3}_{-34.8}, \quad (p_\phi = 1.6, \text{ no } C_\ell^{gg}, \text{ 68\% C.L.}), \quad (4.4)$$

with a degradation in the constraints comparable to the one obtained for the  $p_\phi = 1$  case. In all cases, we find the posterior distribution on all parameters to be mostly symmetric and close to Gaussian. Nevertheless, in all cases we report the median and two-sided 68% confidence interval around it.

These constraints are robust to our specific analysis choices and against several potential sources of systematic errors. We carry out a number of tests to quantify this, repeating our analysis for variations in the data vector and theory model. The resulting constraints obtained in these cases are listed in table 1 and the impact on the  $f_{\text{NL}}$  measurements is shown in figure 12.

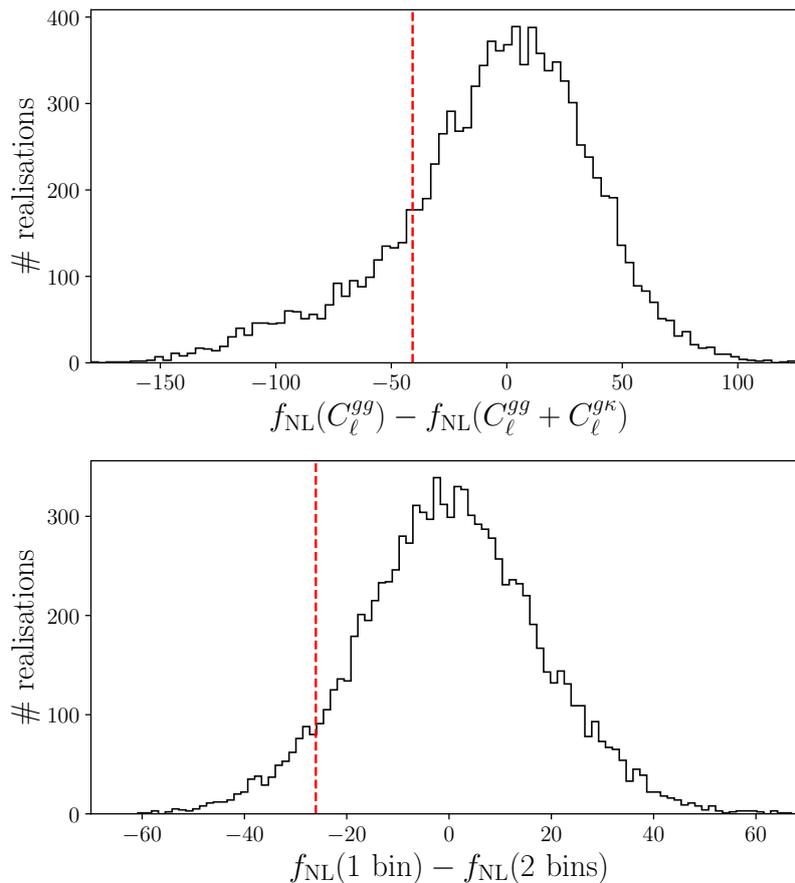
First, we repeat our analysis focusing on different subsets of the data to quantify the level of internal consistency of our constraints. The result of considering only the galaxy-lensing cross-correlation was reported above. Removing these cross-correlations and using only the quasar auto-correlations (row 4 in table 1) results in a lower central value for  $f_{\text{NL}}$ , accompanied by a 90% increase in the statistical uncertainties. This is one of the largest shifts we observe across all our tests, and is mostly driven by the marginally higher bias value favoured by the quasar auto-correlation in the High- $z$  bin.

To quantify the significance of this shift, and the potential presence of internal tension in the data, we carry out the following frequentist exercise. We generate simulated realisations of our full  $C_\ell^{gg} + C_\ell^{g\kappa}$  data vector by drawing from a multivariate normal distribution with a mean given by the theoretical prediction corresponding to our fiducial best-fit parameters, and

| Analysis setting                       | $f_{\text{NL}}$         | $A_1$           | $A_2$           | $\chi^2/N_{\text{dof}}$ | PTE |
|--|-------------------------|-----------------|-----------------|-------------------------|-----|
| 1. Fiducial                            | $-20.5^{+19.0}_{-18.1}$ | $1.09 \pm 0.05$ | $0.96 \pm 0.04$ | 117.9/109               | 0.3 |
| 2. $p_\phi = 1.6$                      | $-28.7^{+26.1}_{-24.6}$ | $1.09 \pm 0.05$ | $0.96 \pm 0.04$ | 118.0/109               | 0.3 |
| 3. No $C_\ell^{gg}$                    | $-13.8^{+26.7}_{-25.0}$ | $1.06 \pm 0.06$ | $0.86 \pm 0.06$ | 56.0/60                 | 0.6 |
| 4. No $C_\ell^{g\kappa}$               | $-61.8^{+37.0}_{-35.6}$ | $1.11 \pm 0.06$ | $1.05 \pm 0.06$ | 55.5/46                 | 0.2 |
| 5. Low- $z$                            | $-17.8^{+67.1}_{-59.3}$ | $1.07 \pm 0.06$ | N.A.            | 52.2/49                 | 0.4 |
| 6. High- $z$                           | $-23.2^{+19.5}_{-19.3}$ | N.A.            | $0.96 \pm 0.04$ | 60.9/59                 | 0.4 |
| 7. +cross-bin                          | $-28.3^{+17.7}_{-16.8}$ | $1.08 \pm 0.05$ | $0.97 \pm 0.04$ | 146.9/131               | 0.2 |
| 8. 1 bin                               | $-46.9^{+24.0}_{-22.3}$ | $1.03 \pm 0.04$ | N.A.            | 50.0/55                 | 0.7 |
| 9. $G < 20$                            | $-12.1^{+27.2}_{-25.0}$ | $1.18 \pm 0.07$ | $0.94 \pm 0.07$ | 110.9/109               | 0.4 |
| 10. $\ell_{\text{min}}^{gg} = 10$      | $-17.7^{+19.0}_{-18.1}$ | $1.09 \pm 0.05$ | $0.95 \pm 0.04$ | 126.6/115               | 0.2 |
| 11. $\ell_{\text{min}}^{g\kappa} = 4$  | $-29.0^{+19.8}_{-19.6}$ | $1.09 \pm 0.05$ | $0.96 \pm 0.04$ | 110.7/107               | 0.4 |
| 12. $\ell_{\text{min}}^{g\kappa} = 10$ | $-38.2^{+24.4}_{-24.0}$ | $1.10 \pm 0.05$ | $0.97 \pm 0.05$ | 109.2/101               | 0.3 |
| 13. No deproj.                         | $-13.5^{+19.0}_{-18.7}$ | $1.10 \pm 0.05$ | $0.95 \pm 0.04$ | 116.0/109               | 0.3 |
| 14. $k_{\text{max}} = 0.1$             | $-28.5^{+18.4}_{-17.9}$ | $1.12 \pm 0.04$ | $0.96 \pm 0.04$ | 142.3/127               | 0.2 |
| 15. Analytic cov.                      | $-21.0^{+18.8}_{-18.8}$ | $1.09 \pm 0.05$ | $0.94 \pm 0.04$ | 124.0/109               | 0.2 |
| 16. With RSD                           | $-20.1^{+18.6}_{-18.6}$ | $1.09 \pm 0.05$ | $0.96 \pm 0.04$ | 117.9/109               | 0.3 |
| 17. With mag.                          | $-21.2^{+18.5}_{-18.4}$ | $1.09 \pm 0.05$ | $0.95 \pm 0.04$ | 118.3/109               | 0.3 |
| 18. P24 $b(z)$                         | $-23.6^{+19.3}_{-18.8}$ | $0.88 \pm 0.04$ | $0.99 \pm 0.05$ | 118.7/109               | 0.2 |
| 19. No $\kappa$ TF                     | $-29.4^{+17.7}_{-17.6}$ | $1.07 \pm 0.05$ | $0.96 \pm 0.05$ | 121.3/109               | 0.2 |
| 20. No-TT                              | $-47.4^{+21.8}_{-21.5}$ | $1.06 \pm 0.05$ | $1.07 \pm 0.05$ | 89.8/109                | 0.9 |
| 21. No-TT, No $C_\ell^{gg}$            | $-24.0^{+28.7}_{-26.4}$ | $1.02 \pm 0.09$ | $1.07 \pm 0.10$ | 46.4/60                 | 0.9 |
| 22. $p_\phi = 1.6$ , No $C_\ell^{gg}$  | $-15.6^{+42.3}_{-34.8}$ | $1.06 \pm 0.06$ | $0.86 \pm 0.06$ | 56.1/63                 | 0.6 |

**Table 1.** Constraints on the model parameters: the PNG amplitude  $f_{\text{NL}}$  and the bias scaling amplitudes for the different redshift bins ( $A_1$  and  $A_2$ ). Results are shown for our fiducial analysis (first row), as well as all other variations used to explore the stability of our results to different analysis choices and potential systematics; detailed descriptions are in the text. The best-fit  $\chi^2$ , number of degrees of freedom, and associated PTE are listed in the last two columns.

the covariance matrix of the data. For each realisation, we then find the best-fit parameters of our model, consider either the full set of auto- and cross-correlations, or only the quasar auto-correlations. We then study the distribution of the difference between the best-fit values of  $f_{\text{NL}}$  found in either case. The distribution reconstructed from  $10^4$  such realisations is shown in the top panel of figure 13, with the  $f_{\text{NL}}$  difference found in the real data shown as a vertical line. We find that 18.6% of all realisations recover downward shifts in  $f_{\text{NL}}$  that are larger than that found in the data. The shift we recover is therefore not statistically significant.



**Figure 13.** *Top panel:* distribution of the difference between the best-fit values of  $f_{\text{NL}}$  found using only the galaxy auto-correlation and that found using both auto- and cross-correlations, obtained from a suite of  $10^4$  Gaussian simulations. The red dashed line shows the value obtained from the real data. *Bottom panel:* as the top panel for the difference between the best-fit values found using one or two redshift bins.

The large increase in statistical uncertainties when using only the quasar auto-correlations also shows that our constraints are dominated by the  $C_\ell^{g\kappa}$  cross-correlation.

We then consider the impact of restricting our analysis to only one of the two redshift bins (rows 5 and 6 for the Low- $z$  and High- $z$  bins, respectively). Discarding the high-redshift bin results in a dramatic  $\sim 3.5$ -fold increase in the statistical uncertainties, clearly showing that our constraints are driven by the high-redshift data. This is reasonable, as this sample covers a larger volume, and the low- $\ell$  angular power spectra are able to access smaller physical wavenumbers  $k$ , more sensitive to PNG. In either case, the resulting  $f_{\text{NL}}$  constraints are compatible with our fiducial measurement. Finally, we repeat our analysis including the cross-correlation between quasars in both redshift bins. The resulting constraints on  $f_{\text{NL}}$  (row 7 in table 1) shift downwards by  $\sim 0.4\sigma$ , while the errors shrink by approximately 10%. The impact of including this cross-correlation is interesting for several reasons. First, in principle this cross-correlation should allow us to recover some information contained in long-wavelength radial modes that is lost by considering only projected statistics [66]. At

the same time, cross-bin correlations are less robust to various forms of systematics. Besides genuine extra power on very large scales, the only other possible contributions to the signal in this spectrum come from the overlap between the redshift distributions of both bins, or from magnification bias. On the other hand, a non-zero cross-correlation could also be sourced by correlated residual sky contamination in both bins. Therefore, the fact that our fiducial results do not change significantly when including the cross-bin correlation, is further reassurance that these sources of systematic uncertainty (photometric redshift errors, uncertainties in the magnification bias slope, and residual sky contamination) are under control in our sample. The goodness of fit for this case also supports this conclusions (see also the discussion in [26]). However, in order to be conservative as we do not have a fully overlapping spectroscopic sample to accurately test or recalibrate our redshift distribution, we did not quote this as our baseline.

We also explored the impact our analysis varying our selection of the *Quiaia* sources. First, we repeated our analysis using a single redshift bin covering the range  $z_p < 5$  (row 8). The resulting constraints shift towards negative  $f_{\text{NL}}$  by approximately  $1\sigma$ , with uncertainties that grow by  $\sim 30\%$ . To quantify the significance of this shift, we repeat the analysis described above to study the shift observed when excluding the lensing cross-correlations. We generate Gaussian realisations of our power spectrum measurements for the single-bin and two-bin versions of the analysis, fully accounting for the covariance between both datasets at the level of both signal and shot noise (we estimate the cross-covariance between the single-bin and two-bin power spectra following the same approach used for our fiducial measurements). We then find the best-fit values of  $f_{\text{NL}}$  in both datasets, and recover the distribution of the difference between them. The result is shown in the bottom panel of figure 13, with the difference found in the real data marked as a vertical red dashed line. We find that 13% of all realisations recover larger shifts than those found in the data. This shift is therefore only mildly significant ( $1.5\sigma$ ), although there are several potential reasons for it. For instance, the single-bin case is potentially more strongly affected by the inclusion of small-scale structures at low redshifts within our scale cuts, as well as more sensitive to the assumed evolution of the quasar bias within the bin. A misestimation of the effective galaxy bias due to either of these effects would be compensated by a shift in  $f_{\text{NL}}$ . As a second test of our quasar selection, we obtained constraints after imposing a brighter magnitude cut on the *Quiaia* sample, using only sources with  $G > 20$ ; this is a cleaner sample with better redshift uncertainties and lower systematics. The resulting constraints (row 9) are compatible with our fiducial measurement, with a measurement error that grows by  $\sim 40\%$  due to the impact of shot noise.

We find that our constraints are broadly stable against variations in the large-scale cut imposed on the data vector, both for  $C_\ell^{gg}$  and  $C_\ell^{g\kappa}$  (rows 10–12). In particular, using a less conservative scale cut for  $C_\ell^{gg}$  (row 10) leads to a shift in  $f_{\text{NL}}$  towards zero. Since potential contamination from observational systematics would most likely manifest itself on the largest scales probed, this is indirect confirmation that our results are robust against this contamination. As further evidence of this robustness, repeating our analysis on galaxy overdensity maps on which no deprojection of contaminant templates is performed (row 13) does not change our results significantly. Changing the small-scale cut used from  $k_{\text{max}} = 0.07 \text{ Mpc}^{-1}$  to  $k_{\text{max}} = 0.1 \text{ Mpc}^{-1}$  also does not strongly affect our constraints (row 14).

To test the stability of our constraints against the procedure used to estimate the covariance matrix of our measurements, we repeat our analysis using the analytical Gaussian covariance matrix estimated by `NaMaster`. The narrow-kernel approximation used in this estimate is known to lose accuracy on large scales, where the PNG signal is concentrated. The methodology used here also ignores the impact of inhomogeneity in the associated maps (e.g. caused by the spatially-varying *Quia* selection function). In spite of this, the constraints found are in very good agreement with our fiducial measurement (row 15).

Our constraints are also robust against potential errors in the theoretical predictions. Our fiducial model ignores both magnification bias and redshift-space distortions (RSDs), and assumes a specific functional form for the redshift evolution of the linear galaxy bias. Accounting for RSDs leads to virtually no change in the constraints (row 16). The inclusion of magnification bias, using the values of  $s(z)$  measured for the *Quia* sample in [26] (see figure 3), also has a negligible effect on our results (row 17). Using the bias evolution model of P24 instead (see eq. (3.30)) also does not change our conclusions significantly (row 18).

Finally, we study the impact of potential systematics in the CMB lensing maps used in our analysis. First, we quantify the effect of correcting for the mode loss in lensing reconstruction by repeating our analysis without correcting for the associated transfer function described in section 2.2. Failing to include this transfer function results in a downward shift in  $f_{\text{NL}}$  by  $\sim 0.5\sigma$  (row 19). Thus, we do not expect any potential misestimation of the lensing reconstruction transfer function to affect our constraints significantly. Secondly, we test for the potential presence of residual contamination from extragalactic foregrounds in the CMB lensing map. Most importantly, the presence of Cosmic Infrared Background residuals could correlate strongly with the high-redshift quasars, leading to a biased cross-correlation measurement. This could then bias our estimate of  $f_{\text{NL}}$ , either by directly modifying the shape of the cross-correlation on large scales, or by affecting our estimate of the quasar bias. As was done in [26, 28], to test for this systematic, we make use of an alternative CMB lensing map constructed avoiding  $TT$  contributions in the quadratic estimator, described in section 2.2. Since only combinations involving polarization data contribute to this map, it should be largely free from extragalactic foreground contamination, at the cost of a degraded sensitivity. Repeating our analysis using this map, using both  $C_\ell^{gg}$  and  $C_\ell^{g\kappa}$ , leads to the constraints found in row 20 of table 1. The constraints are shifted by  $\sim -1\sigma$  with respect to our fiducial measurements. However, this shift is not caused by the change in CMB lensing map, but is instead due to the preference of  $C_\ell^{gg}$  for a more negative value of  $f_{\text{NL}}$ , discussed above (see row 4): the lower sensitivity of the No- $TT$  map downweights the contribution of  $C_\ell^{g\kappa}$ , which then leads to this shift. We can verify this by removing the quasar auto-correlations from the analysis. The resulting constraints on  $f_{\text{NL}}$  (row 21) are in good agreement with our fiducial results (row 1), and with the  $C_\ell^{g\kappa}$ -only analysis carried out with the GMV  $\kappa$  map (row 3).

In summary, none of these tests reveal a significant systematic in our fiducial  $f_{\text{NL}}$  constraints, or a lack of internal consistency in our measurements. The best-fit  $\chi^2$  and associated PTE for these alternative analyses are listed in the two last columns of table 1. In all cases we find that our best-fit model provides a reasonable description of the data, with PTEs ranging from 0.15 to 0.9.

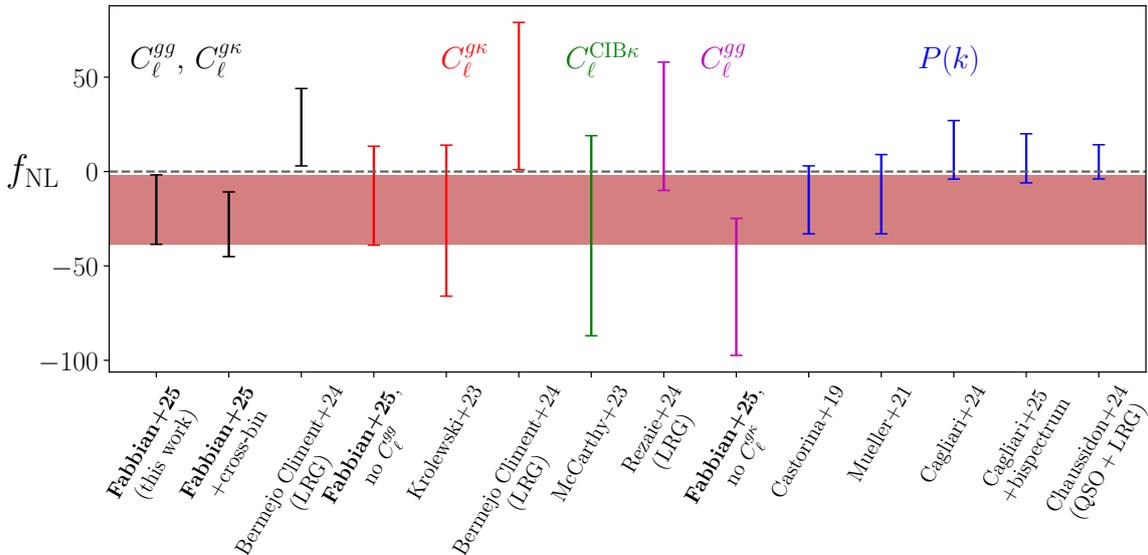
## 5 Discussion and conclusions

We have presented constraints on local-type primordial non-Gaussianity from the clustering of quasars in the *Quaia* sample, using both their projected auto-correlation and their cross-correlation with maps of the CMB lensing convergence by *Planck*. We validated our power spectrum estimation pipeline against an independent implementation of the quadratic minimum-variance estimator, and quantified the impact of systematic contamination in the quasar overdensity maps via linear deprojection techniques. We find the PNG amplitude  $f_{\text{NL}}$  to be compatible with zero, at the  $\sim 1.1\sigma$  level in our fiducial analysis, and at  $1.0\sigma$  with the *Planck* estimate of  $f_{\text{NL}}^{\text{local}} = -0.9 \pm 5.1$  [3]. The measurement uncertainties, assuming an  $f_{\text{NL}}$  response with  $p_\phi = 1$  or  $p_\phi = 1.6$ , are  $\sigma(f_{\text{NL}}) \simeq 18$  and  $\sigma(f_{\text{NL}}) \simeq 25$ , respectively (see eqs. (4.1) and (4.2)). These results are consistent with those found using only the quasar-CMB lensing cross-correlation, which is potentially more robust against systematic contamination (see eq. (4.3)). Exploiting the quasar cross-correlation between different redshift bins allows us to further reduce the statistical uncertainties by  $\sim 10\%$ .

Our results are robust against the most likely sources of systematic uncertainty, both observational and theoretical. This includes our choice of scale cuts, the presence of residual contamination in the galaxy overdensity and CMB lensing maps, the impact of magnification bias and redshift-space distortions, and assumptions about bias evolution. The results also display a good level of internal consistency. Although we observe a shift in  $f_{\text{NL}}$  towards more negative values when considering only the quasar auto-correlations, we showed that the shift is not statistically significant. Our constraints are driven by the large-scale cross-correlation with CMB lensing, and by the High- $z$  quasar sample.

It is interesting to compare our results with the most recent PNG constraints from the literature. Figure 14 summarises the current state of the art. The results presented here are consistent with those found from the projected clustering of quasars and luminous red galaxies (LRGs) in the DESI Legacy Survey [11, 15], both in auto- and cross-correlation, as well as the cross-correlation between CMB lensing and the Cosmic Infrared Background [13]. Our analysis gives the tightest constraints to date on  $f_{\text{NL}}$  from CMB lensing cross-correlations, and from projected datasets overall. Using the three-dimensional clustering of quasars allows for tighter constraints on  $f_{\text{NL}}$ , however. In particular, although the precision obtained in this analysis is similar to that achieved by early analyses of quasar clustering in the extended Baryon Oscillation Spectroscopic Survey (eBOSS) [10, 67], the uncertainties on  $f_{\text{NL}}$  obtained here are between 15% and 50% larger than those found using the latest eBOSS quasar data [12, 68] with optimized redshift-dependent weights, or the combination of LRG and quasar clustering from the first data release of the Dark Energy Spectroscopic Instrument [16]. Nevertheless, the CMB lensing cross-correlation analysis offers the key advantage of being less sensitive to the impact of residual large-scale sky contamination, and *Quaia* is an independent sample that complements the measurements from these data sets.

Although the large angular scales of *Quaia* are easier to model compared to ground-based data sets, modes  $\ell \lesssim 15$  in the projected QSO distribution are still affected by systematics. We carried out a forecast with our likelihood pipeline to assess the impact of systematics in our final statistical uncertainties. For this purpose, we replaced our measured data with theoretical predictions based on a cosmological model with  $f_{\text{NL}} = 0$ , and considered two



**Figure 14.** Constraints on  $f_{\text{NL}}$  at 68% C.L. found in this work in comparison with recent bounds from other groups in the literature [10–13, 15, 16, 67–69]. Black, red, and blue error bars show constraints found from the combination of  $C_\ell^{gg}$  and  $C_\ell^{g\kappa}$ , from  $C_\ell^{g\kappa}$  alone, and from the three-dimensional clustering of galaxies, respectively. The green error bars show the constraints found from the cross-correlation of CIB maps and CMB lensing [13] and the magenta one from  $C_\ell^{gg}$  alone. Unless stated the constraints used samples of QSOs.

cases: first, we used the same covariance of the data, and secondly, a modified covariance that includes only contributions of signal and shot noise at all scales. We found that given the assumed scale cuts, the derived marginalized error bar on  $f_{\text{NL}}$  is consistent in both cases. This shows that our constraints (which are driven by the sources at  $z \gtrsim 1.5$ ) are at present limited by the shot noise in *Quiaia*, the *Planck* lensing noise, and by the choice of the scale cuts rather than the additional error induced by systematics. Conversely, for the 1-bin case (and for the analysis of the low- $z$  bin) systematics limit more significantly the results. We also verified that improvements expected from the addition from the cross-correlation between the two tomographic redshift bins should be around 7%, in the ballpark of what we saw in the data. If *Quiaia* can be improved in future work or using future *Gaia* data releases, and therefore a lower  $\ell_{\text{min}}^{gg}$  can be included in the analysis, the variance of  $C_\ell^{g\kappa}$  would be consequently reduced and the constraining power if  $C_\ell^{gg}$  would increase. In this scenario,  $f_{\text{NL}}$  could be constrained with a precision  $\sigma(f_{\text{NL}}) \approx 9.5$  for a tomographic analysis with 2 redshift bins and  $\ell_{\text{min}}^{gg} = 2$  combining  $C_\ell^{g\kappa}$  and  $C_\ell^{gg}$ . Using  $C_\ell^{gg}$  alone would give  $\sigma(f_{\text{NL}}) \approx 11.5$  and would become the most constraining observable, with the *Planck* CMB lensing noise as limiting factor. Using CMB lensing from upcoming deep CMB polarization surveys covering large sky fractions such as Simons Observatory [70, 71] will allow to overcome this limitation, as shown recently for the case of ACT [72].

We also note that the analysis presented here could be improved in a number of ways, all of which could lead to a significant increase in the constraining power of the *Quiaia* dataset. Firstly, we have not attempted to derive and apply redshift-dependent weights for

the *Quaia* sources that would maximise the sensitivity to  $f_{\text{NL}}$ . The use of optimal weights has been shown to lead to very significant improvements in 3D analyses [10, 68], and similar approaches have been advanced for projected data [73, 74]. Secondly, although the precision of the spectro-photometric redshifts in *Quaia* should prevent a detailed measurement of the small-scale clustering along the line of sight, it should be sufficient to recover long-wavelength radial modes carrying additional information about PNG. More quantitatively, the mean normalised error in *Quaia* is approximately  $\sigma_z/(1+z) \simeq 0.06$ . Translating this into a radial wavenumber as  $k_{\parallel, \text{max}} \equiv \pi H(z)/\sigma_z$ ,  $k_{\parallel, \text{max}}$  ranges between  $\sim 0.01 \text{ Mpc}^{-1}$  and  $0.016 \text{ Mpc}^{-1}$  across the redshift range covered by *Quaia*. Although this is not enough to resolve the radial BAO, it should allow *Quaia* to constrain the clustering of quasars on horizon-sized radial scales over which PNG plays a significant role. Finally, we could further constrain PNG by including information from higher-order statistics, namely the galaxy bispectrum [68, 75–78]. Deploying efficient estimators of the projected bispectrum [79], targeting both the quasar auto-correlation and cross-correlations with CMB lensing, could therefore allow us to sharpen the constraints presented here using *Quaia*.

## Acknowledgments

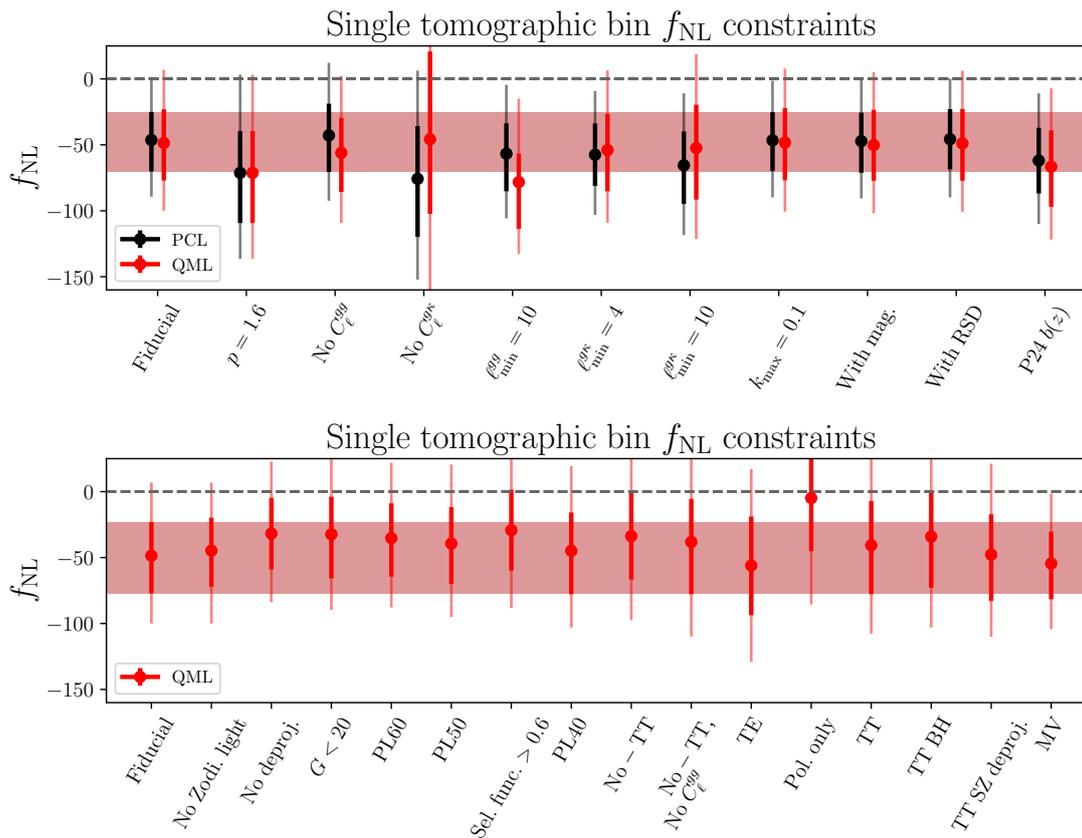
We would like to thank Nestor Arsenov, José Bermejo-Climent, Carlos Hernández-Monteagudo, David W. Hogg, Andras Kovacs, Anže Slosar, Abby Williams and the members of the *Quaia* team for useful comments and discussions. GF thanks Gabriella Contardo for fostering this collaboration. GF acknowledges the support of the STFC Ernest Rutherford fellowship, of the European Research Council under the Marie Skłodowska Curie actions through the Individual Global Fellowship No. 892401 PiCOGAMBAS, of the European Union’s Horizon 2020 research and innovation program (Grant agreement No. 851274), and of the Simons Foundation for the initial phase of this work. DA acknowledges support from the Beecroft Trust. KSF acknowledges the support of the Kavli Foundation.

**Software and Data availability** This work relied on open-source software packages, including NaMaster [32], CCL [60], SciPy [80], NumPy [81], Matplotlib [82], emcee [65], healpy [83], HEALPix [84] and xQML [85]. We make our key data products and code available at <https://github.com/gfabbian/quaia-fnl>; the rest of the analysis code will be made available upon request to the authors.

## A QML results and consistency with PCL pipeline

In figure 6 we show that the measurements of the angular power spectrum derived from the QML and PCL pipeline show a very good level of agreement despite implementing very different approaches to remove systematics and to estimate the data points and their covariance. As the PCL pipeline is numerically less costly and more easily adaptable, we use it to carry out the tomographic analysis reported in the main text. Before doing so, we performed several tests to assess the level of consistency on the estimated  $f_{\text{NL}}$  in the case of a single tomographic redshift bin across both pipeline.

In the upper panel of figure 15 we show a detailed comparison of the same analysis setups considered in figure 12 and table 1. As it can be seen both the QML and PCL pipeline give



**Figure 15.** Top:  $f_{\text{NL}}$  constraints from a single tomographic bin analysis of *Quia* derived with the QML pipeline compared to the results obtained with the PCL pipeline for common analysis setups of the  $G < 20.5$  sample of figure 12. The results show excellent agreement both in terms of central value and error bars. Bottom: additional consistency tests on  $f_{\text{NL}}$  constraints ran for the QML pipeline mostly focus on stability with respect to changes in analysis mask, CMB lensing map or catalog definition. For both top and bottom panel, the shared error bars indicate the 95% confidence level errors after marginalization over the QSO bias amplitude.

very consistent results. The differences in the estimated value of  $f_{\text{NL}}$  are negligible. The QML pipeline shows a slightly larger error bar of the order of 10% or less, consistent with the difference in the number of realizations employed to estimate the measurements' covariances in the two pipelines. As the QML pipeline also models the noise properties of the measurements using more realistic simulations, these results also prove that the data-based approach used to model the covariance in the PCL pipeline is sufficiently accurate to capture the error of the measurements. The only minor differences across pipeline can be seen when we relax the scale cuts used for  $C_{\ell}^{gg}$  and include larger angular scales where the systematics subtraction is very different. However, even in this case, the shifts in the estimated  $f_{\text{NL}}$  amounts to  $\lesssim 0.5\sigma$ .

The lower panel of figure 6 here shows the stability of the  $f_{\text{NL}}$  estimates of the QML with respect to the change of catalog (using only the brightest  $G < 20$  QSOs), systematics templates as well as CMB lensing reconstruction method. Specifically we tested that removing the Zodiacal light template from the list of systematics templates that are marginalized

over, does not lead to a noticeable difference in the results. For the mask variations we considered the official *Planck* galactic masks that retain the 60% (PL60) and 40% (PL40) of the sky available on the *Planck* legacy archive. To this we considered an extra mask obtained thresholding the brightest pixels of the *Planck* 353 GHz map to retain 50% of the sky (PL50) as well a mask including sky locations where the *Quaia* selection function exceeds 60%. This roughly retains a 40% sky fraction. Finally we also considered additional variations of the *Planck* PR4 CMB lensing estimation compared those in section 4. In particular, we included CMB lensing maps derived using the minimum variance combination (MV) that treats suboptimally the CMB noise inhomogeneity [86], using CMB polarization data only, or correlating CMB temperature and E-mode polarization only (TE). We also included CMB lensing maps reconstructed using CMB temperature data only (TT) with and without specific extragalactic foreground mitigation techniques such as bias hardening (BH) against point source emission, or deprojection of the SZ effect in the component separation step prior to lensing reconstruction. All these variation of lensing reconstruction and related simulation (that we used for the covariance estimation) are publicly available.<sup>7</sup> For all these cases we found shifts in the parameter within fractions of the error. The polarization only reconstruction experiences larger shifts, however the CMB lensing map is considerably noisier and the final error on  $f_{\text{NL}}$  is 50% larger. As such it is not statistically significant. Because we applied the different galactic masks on top of the mask used for the fiducial analysis (and therefore the data sets are effectively nested), we tested the consistency of the parameters using the criteria outlined in [87], and found no indication of anomalies. We also found that using data for which no systematics marginalization is done, led to minor changes in the results, showing the robustness of our scale cuts and lack of systematics in  $C_\ell^{g\kappa}$  that drives the constraints. For the cases considered here, the estimated value of  $f_{\text{NL}}$  is also consistent with 0 at  $\sim 2\sigma$  level (95% C.L.). All the results of the tests we considered are summarized in table 2.

---

<sup>7</sup>[https://data.cmb-s4.org/planck\\_pr4.html](https://data.cmb-s4.org/planck_pr4.html).

| Analysis setting                   | $f_{\text{NL}}$         | $A_1$           | $\chi^2/N_{\text{dof}}$ | PTE  |
|------------------------------------|-------------------------|-----------------|-------------------------|------|
| Fiducial                           | $-48.6^{+25.5}_{-28.5}$ | $1.02 \pm 0.04$ | 54.3/55                 | 0.50 |
| $p = 1.6$                          | $-71.3^{+31.7}_{-38.0}$ | $1.03 \pm 0.04$ | 50.1/55                 | 0.66 |
| No $C_\ell^{gg}$                   | $-56.0^{+26.5}_{-29.7}$ | $1.00 \pm 0.05$ | 24.6/30                 | 0.74 |
| No $C_\ell^{g\kappa}$              | $-45.8^{+66.6}_{-56.5}$ | $1.03 \pm 0.06$ | 23.2/23                 | 0.45 |
| $G < 20$                           | $-32.3^{+28.4}_{-33.5}$ | $1.05 \pm 0.06$ | 43.2/55                 | 0.87 |
| $\ell_{\text{min}}^{gg} = 10$      | $-78.2^{+21.6}_{-35.4}$ | $1.04 \pm 0.04$ | 61.6/58                 | 0.35 |
| $\ell_{\text{min}}^{g\kappa} = 4$  | $-53.9^{+27.3}_{-31.3}$ | $1.02 \pm 0.04$ | 54.3/54                 | 0.46 |
| $\ell_{\text{min}}^{g\kappa} = 10$ | $-52.4^{+32.9}_{-39.1}$ | $1.02 \pm 0.05$ | 53.8/51                 | 0.37 |
| No deproj.                         | $-31.9^{+27.1}_{-27.1}$ | $1.01 \pm 0.04$ | 58.2/55                 | 0.36 |
| $k_{\text{max}} = 0.1$             | $-48.2^{+26.1}_{-28.6}$ | $1.01 \pm 0.04$ | 69.0/65                 | 0.34 |
| With RSD                           | $-48.9^{+26.1}_{-28.3}$ | $1.02 \pm 0.04$ | 54.3/55                 | 0.50 |
| With mag.                          | $-50.2^{+26.8}_{-27.0}$ | $1.01 \pm 0.04$ | 54.4/55                 | 0.50 |
| P24 $b(z)$                         | $-66.6^{+27.6}_{-30.4}$ | $0.93 \pm 0.04$ | 55.5/55                 | 0.46 |
| No TT                              | $-33.7^{+32.6}_{-32.9}$ | $1.04 \pm 0.05$ | 47.3/55                 | 0.76 |
| No TT No $C_\ell^{gg}$             | $-38.1^{+32.6}_{-39.6}$ | $1.01 \pm 0.08$ | 19.5/30                 | 0.93 |
| TT                                 | $-40.7^{+33.5}_{-37.0}$ | $1.00 \pm 0.05$ | 49.5/55                 | 0.68 |
| TTBH                               | $-34.1^{+34.4}_{-38.8}$ | $0.99 \pm 0.05$ | 49.3/55                 | 0.69 |
| TT SZ deproj.                      | $-47.7^{+30.5}_{-35.2}$ | $1.00 \pm 0.05$ | 50.6/55                 | 0.64 |
| TE                                 | $-56.0^{+37.2}_{-37.7}$ | $1.03 \pm 0.05$ | 46.0/55                 | 0.80 |
| Pol. only                          | $-4.9^{+40.6}_{-40.3}$  | $1.02 \pm 0.05$ | 43.3/55                 | 0.87 |
| MV                                 | $-54.5^{+24.2}_{-27.0}$ | $1.02 \pm 0.04$ | 49.2/55                 | 0.69 |
| PL60                               | $-35.3^{+26.4}_{-29.2}$ | $1.00 \pm 0.04$ | 52.7/55                 | 0.56 |
| PL50                               | $-39.4^{+27.6}_{-30.6}$ | $0.98 \pm 0.04$ | 56.5/55                 | 0.42 |
| PL40                               | $-44.9^{+29.1}_{-32.9}$ | $0.99 \pm 0.05$ | 46.8/55                 | 0.78 |
| Sel. func. $> 0.6$                 | $-29.2^{+30.6}_{-30.7}$ | $0.99 \pm 0.04$ | 60.1/55                 | 0.30 |
| No Zodi. light                     | $-44.7^{+25.0}_{-27.4}$ | $1.02 \pm 0.04$ | 54.4/55                 | 0.50 |

**Table 2.** Constraints on the model parameters obtained with the QML pipeline for the benchmark case of a single tomographic redshift bin as in table 1. Results are shown for our fiducial analysis (first row), as well as all other variations used to explore the stability of our results to different analysis choices and potential systematics. Error bars are 68% C.L. values.

## References

- [1] N. Bartolo, E. Komatsu, S. Matarrese and A. Riotto, *Non-Gaussianity from inflation: Theory and observations*, *Phys. Rept.* **402** (2004) 103 [[astro-ph/0406398](#)] [[INSPIRE](#)].
- [2] D. Baumann, *Inflation*, in the proceedings of the *Theoretical Advanced Study Institute in Elementary Particle Physics: Physics of the Large and the Small*, Boulder, U.S.A. (2011) [[DOI:10.1142/9789814327183\\_0010](#)] [[arXiv:0907.5424](#)] [[INSPIRE](#)].
- [3] PLANCK collaboration, *Planck 2018 results. IX. Constraints on primordial non-Gaussianity*, *Astron. Astrophys.* **641** (2020) A9 [[arXiv:1905.05697](#)] [[INSPIRE](#)].
- [4] G. Jung et al., *Constraints on primordial non-Gaussianity from Planck PR4 data*, *Astron. Astrophys.* **702** (2025) A204 [[arXiv:2504.00884](#)] [[INSPIRE](#)].
- [5] P. Creminelli and M. Zaldarriaga, *Single field consistency relation for the 3-point function*, *JCAP* **10** (2004) 006 [[astro-ph/0407059](#)] [[INSPIRE](#)].
- [6] C.T. Byrnes and K.-Y. Choi, *Review of local non-Gaussianity from multi-field inflation*, *Adv. Astron.* **2010** (2010) 724525 [[arXiv:1002.3110](#)] [[INSPIRE](#)].
- [7] R. de Putter, J. Gleyzes and O. Doré, *Next non-Gaussianity frontier: What can a measurement with  $\sigma(f_{NL}) \lesssim 1$  tell us about multifield inflation?*, *Phys. Rev. D* **95** (2017) 123507 [[arXiv:1612.05248](#)] [[INSPIRE](#)].
- [8] N. Dalal, O. Dore, D. Huterer and A. Shirokov, *The imprints of primordial non-gaussianities on large-scale structure: scale dependent bias and abundance of virialized objects*, *Phys. Rev. D* **77** (2008) 123514 [[arXiv:0710.4560](#)] [[INSPIRE](#)].
- [9] V. Desjacques and U. Seljak, *Primordial non-Gaussianity from the large scale structure*, *Class. Quant. Grav.* **27** (2010) 124011 [[arXiv:1003.5020](#)] [[INSPIRE](#)].
- [10] eBOSS collaboration, *Redshift-weighted constraints on primordial non-Gaussianity from the clustering of the eBOSS DR14 quasars in Fourier space*, *JCAP* **09** (2019) 010 [[arXiv:1904.08859](#)] [[INSPIRE](#)].
- [11] DESI collaboration, *Constraining primordial non-Gaussianity from DESI quasar targets and Planck CMB lensing*, *JCAP* **03** (2024) 021 [[arXiv:2305.07650](#)] [[INSPIRE](#)].
- [12] M.S. Cagliari, E. Castorina, M. Bonici and D. Bianchi, *Optimal constraints on Primordial non-Gaussianity with the eBOSS DR16 quasars in Fourier space*, *JCAP* **08** (2024) 036 [[arXiv:2309.15814](#)] [[INSPIRE](#)].
- [13] F. McCarthy, M.S. Madhavacheril and A.S. Maniyar, *Constraints on primordial non-Gaussianity from halo bias measured through CMB lensing cross-correlations*, *Phys. Rev. D* **108** (2023) 083522 [[arXiv:2210.01049](#)] [[INSPIRE](#)].
- [14] L.D. Ferramacho, M.G. Santos, M.J. Jarvis and S. Camera, *Radio galaxy populations and the multitracers technique: pushing the limits on primordial non-Gaussianity*, *Mon. Not. Roy. Astron. Soc.* **442** (2014) 2511 [[arXiv:1402.2290](#)] [[INSPIRE](#)].
- [15] J.R. Bermejo-Climent et al., *Constraints on primordial non-Gaussianity from the cross-correlation of DESI luminous red galaxies and Planck CMB lensing*, *Astron. Astrophys.* **698** (2025) A177 [[arXiv:2412.10279](#)] [[INSPIRE](#)].
- [16] E. Chaussidon et al., *Constraining primordial non-Gaussianity with DESI 2024 LRG and QSO samples*, *JCAP* **06** (2025) 029 [[arXiv:2411.17623](#)] [[INSPIRE](#)].
- [17] N. Weaverdyck and D. Huterer, *Mitigating contamination in LSS surveys: a comparison of methods*, *Mon. Not. Roy. Astron. Soc.* **503** (2021) 5061 [[arXiv:2007.14499](#)] [[INSPIRE](#)].

- [18] 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](#)].
- [19] J. Carron, M. Mirmelstein and A. Lewis, *CMB lensing from Planck PR4 maps*, *JCAP* **09** (2022) 039 [[arXiv:2206.07773](#)] [[INSPIRE](#)].
- [20] M. Tegmark, *Measuring cosmological parameters with galaxy surveys*, *Phys. Rev. Lett.* **79** (1997) 3806 [[astro-ph/9706198](#)] [[INSPIRE](#)].
- [21] A. Vallenari et al., *Gaia Data Release 3: Summary of the content and survey properties*, *Astron. Astrophys.* **674** (2023) A1 [[arXiv:2208.00211](#)].
- [22] A.M. Meisner, D. Lang, E.F. Schlafly and D.J. Schlegel, *unWISE Coadds: The Five-year Data Set*, *Publ. Astron. Soc. Pac.* **131** (2019) 124504 [[arXiv:1909.05444](#)] [[INSPIRE](#)].
- [23] EBOSS collaboration, *The Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Large-scale structure catalogues for cosmological analysis*, *Mon. Not. Roy. Astron. Soc.* **498** (2020) 2354 [[arXiv:2007.09000](#)] [[INSPIRE](#)].
- [24] E.L. Wright et al., *The Wide-field Infrared Survey Explorer (WISE): Mission Description and Initial On-orbit Performance*, *Astron. J.* **140** (2010) 1868 [[arXiv:1008.0031](#)] [[INSPIRE](#)].
- [25] E.F. Schlafly, A.M. Meisner and G.M. Green, *The unWISE Catalog: Two Billion Infrared Sources from Five Years of WISE Imaging*, *Astrophys. J. Suppl.* **240** (2019) 30 [[arXiv:1901.03337](#)].
- [26] D. Alonso et al., *Constraining cosmology with the Gaia-unWISE Quasar Catalog and CMB lensing: structure growth*, *JCAP* **11** (2023) 043 [[arXiv:2306.17748](#)] [[INSPIRE](#)].
- [27] D. Alonso et al., *Measurement of the power spectrum turnover scale from the cross-correlation between CMB lensing and Quaia*, [arXiv:2410.24134](#) [[DOI:10.33232/001c.136891](#)] [[INSPIRE](#)].
- [28] G. Piccirilli et al., *Growth history and quasar bias evolution at  $z < 3$  from Quaia*, *JCAP* **06** (2024) 012 [[arXiv:2402.05761](#)] [[INSPIRE](#)].
- [29] Y.-K. Chiang, *Corrected SFD: A More Accurate Galactic Dust Map with Minimal Extragalactic Contamination*, *Astrophys. J.* **958** (2023) 118 [[arXiv:2306.03926](#)] [[INSPIRE](#)].
- [30] A. Williams et al., in preparation (2025).
- [31] 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](#)].
- [32] LSST DARK ENERGY SCIENCE collaboration, *A unified pseudo- $C_\ell$  framework*, *Mon. Not. Roy. Astron. Soc.* **484** (2019) 4127 [[arXiv:1809.09603](#)] [[INSPIRE](#)].
- [33] B. Leistedt et al., *Estimating the large-scale angular power spectrum in the presence of systematics: a case study of Sloan Digital Sky Survey quasars*, *Mon. Not. Roy. Astron. Soc.* **435** (2013) 1857 [[arXiv:1306.0005](#)] [[INSPIRE](#)].
- [34] F. Elsner, B. Leistedt and H.V. Peiris, *Unbiased pseudo- $C_\ell$  power spectrum estimation with mode projection*, *Mon. Not. Roy. Astron. Soc.* **465** (2017) 1847 [[arXiv:1609.03577](#)] [[INSPIRE](#)].
- [35] POLARBEAR collaboration, *A Measurement of the Degree Scale CMB B-mode Angular Power Spectrum with POLARBEAR*, *Astrophys. J.* **897** (2020) 55 [[arXiv:1910.02608](#)] [[INSPIRE](#)].
- [36] C. Hervías-Caimapo et al., *The Simons Observatory: validation of reconstructed power spectra from simulated filtered maps for the small aperture telescope survey*, *JCAP* **06** (2025) 055 [[arXiv:2502.00946](#)] [[INSPIRE](#)].

- [37] ATACAMA COSMOLOGY TELESCOPE collaboration, *The Atacama Cosmology Telescope: DR6 power spectra, likelihoods and  $\Lambda$ CDM parameters*, *JCAP* **11** (2025) 062 [[arXiv:2503.14452](#)] [[INSPIRE](#)].
- [38] T. Cornish, D. Alonso, B. Leistedt and K. Wolz, *Systematics mitigation for catalogue-based angular power spectra*, [arXiv:2510.19912](#) [[INSPIRE](#)].
- [39] V. Mittal, O.T. Oyda and G.F. Lewis, *The Cosmic Dipole in the Quiaia Sample of Quasars: A Bayesian Analysis*, *Mon. Not. Roy. Astron. Soc.* **527** (2024) 8497 [[arXiv:2311.14938](#)] [[INSPIRE](#)].
- [40] M. Tegmark, *How to measure CMB power spectra without losing information*, *Phys. Rev. D* **55** (1997) 5895 [[astro-ph/9611174](#)] [[INSPIRE](#)].
- [41] S. Vanneste, S. Henrot-Versillé, T. Louis and M. Tristram, *Quadratic estimator for CMB cross-correlation*, *Phys. Rev. D* **98** (2018) 103526 [[arXiv:1807.02484](#)] [[INSPIRE](#)].
- [42] G.B. Rybicki and W.H. Press, *Interpolation, realization, and reconstruction of noisy, irregularly sampled data*, *Astrophys. J.* **398** (1992) 169 [[INSPIRE](#)].
- [43] C. García-García, D. Alonso and E. Bellini, *Disconnected pseudo- $C_\ell$  covariances for projected large-scale structure data*, *JCAP* **11** (2019) 043 [[arXiv:1906.11765](#)] [[INSPIRE](#)].
- [44] EBOSS collaboration, *Clustering of quasars in SDSS-IV eBOSS: study of potential systematics and bias determination*, *JCAP* **07** (2017) 017 [[arXiv:1705.04718](#)] [[INSPIRE](#)].
- [45] J. Hartlap, P. Simon and P. Schneider, *Why your model parameter confidences might be too optimistic: Unbiased estimation of the inverse covariance matrix*, *Astron. Astrophys.* **464** (2007) 399 [[astro-ph/0608064](#)] [[INSPIRE](#)].
- [46] W.H. Press and P. Schechter, *Formation of galaxies and clusters of galaxies by selfsimilar gravitational condensation*, *Astrophys. J.* **187** (1974) 425 [[INSPIRE](#)].
- [47] A. Slosar et al., *Constraints on local primordial non-Gaussianity from large scale structure*, *JCAP* **08** (2008) 031 [[arXiv:0805.3580](#)] [[INSPIRE](#)].
- [48] A. Barreira et al., *Galaxy bias and primordial non-Gaussianity: insights from galaxy formation simulations with IllustrisTNG*, *JCAP* **12** (2020) 013 [[arXiv:2006.09368](#)] [[INSPIRE](#)].
- [49] A. Barreira, *Predictions for local PNG bias in the galaxy power spectrum and bispectrum and the consequences for  $f_{NL}$  constraints*, *JCAP* **01** (2022) 033 [[arXiv:2107.06887](#)] [[INSPIRE](#)].
- [50] C. Bonvin and R. Durrer, *What galaxy surveys really measure*, *Phys. Rev. D* **84** (2011) 063505 [[arXiv:1105.5280](#)] [[INSPIRE](#)].
- [51] A. Challinor and A. Lewis, *The linear power spectrum of observed source number counts*, *Phys. Rev. D* **84** (2011) 043516 [[arXiv:1105.5292](#)] [[INSPIRE](#)].
- [52] D. Alonso et al., *Ultra large-scale cosmology in next-generation experiments with single tracers*, *Astrophys. J.* **814** (2015) 145 [[arXiv:1505.07596](#)] [[INSPIRE](#)].
- [53] D. Alonso and P.G. Ferreira, *Constraining ultralarge-scale cosmology with multiple tracers in optical and radio surveys*, *Phys. Rev. D* **92** (2015) 063525 [[arXiv:1507.03550](#)] [[INSPIRE](#)].
- [54] S.L. Guedezounme, S. Jolicoeur and R. Maartens, *Primordial non-Gaussianity — the effects of relativistic and wide-angle corrections to the power spectrum*, *JCAP* **07** (2025) 063 [[arXiv:2412.06553](#)] [[INSPIRE](#)].
- [55] T.D. Kitching and A.F. Heavens, *Unequal-Time Correlators for Cosmology*, *Phys. Rev. D* **95** (2017) 063522 [[arXiv:1612.00770](#)] [[INSPIRE](#)].

- [56] N.E. Chisari and A. Pontzen, *Unequal time correlators and the Zel'dovich approximation*, *Phys. Rev. D* **100** (2019) 023543 [[arXiv:1905.02078](#)] [[INSPIRE](#)].
- [57] D.N. Limber, *The Analysis of Counts of the Extragalactic Nebulae in Terms of a Fluctuating Density Field.*, *Astrophys. J.* **117** (1953) 134.
- [58] M. LoVerde and N. Afshordi, *Extended Limber Approximation*, *Phys. Rev. D* **78** (2008) 123506 [[arXiv:0809.5112](#)] [[INSPIRE](#)].
- [59] P. Simon, *How accurate is Limber's equation?*, *Astron. Astrophys.* **473** (2007) 711 [[astro-ph/0609165](#)] [[INSPIRE](#)].
- [60] LSST DARK ENERGY SCIENCE collaboration, *Core Cosmology Library: Precision Cosmological Predictions for LSST*, *Astrophys. J. Suppl.* **242** (2019) 2 [[arXiv:1812.05995](#)] [[INSPIRE](#)].
- [61] LSST DARK ENERGY SCIENCE collaboration, *The N5K Challenge: Non-Limber Integration for LSST Cosmology*, *Open J. Astrophys.* **6** (2023) 1 [[arXiv:2212.04291](#)] [[INSPIRE](#)].
- [62] X. Fang, E. Krause, T. Eifler and N. MacCrann, *Beyond Limber: Efficient computation of angular power spectra for galaxy clustering and weak lensing*, *JCAP* **05** (2020) 010 [[arXiv:1911.11947](#)] [[INSPIRE](#)].
- [63] S. Hamimeche and A. Lewis, *Likelihood Analysis of CMB Temperature and Polarization Power Spectra*, *Phys. Rev. D* **77** (2008) 103013 [[arXiv:0801.0554](#)] [[INSPIRE](#)].
- [64] PLANCK collaboration, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6 [*Erratum ibid.* **652** (2021) C4] [[arXiv:1807.06209](#)] [[INSPIRE](#)].
- [65] D. Foreman-Mackey, D.W. Hogg, D. Lang and J. Goodman, *emcee: The MCMC Hammer*, *Publ. Astron. Soc. Pac.* **125** (2013) 306 [[arXiv:1202.3665](#)] [[INSPIRE](#)].
- [66] J. Asorey, M. Crocce, E. Gaztanaga and A. Lewis, *Recovering 3D clustering information with angular correlations*, *Mon. Not. Roy. Astron. Soc.* **427** (2012) 1891 [[arXiv:1207.6487](#)] [[INSPIRE](#)].
- [67] EBOSS collaboration, *Primordial non-Gaussianity from the completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey II: measurements in Fourier space with optimal weights*, *Mon. Not. Roy. Astron. Soc.* **514** (2022) 3396 [[arXiv:2106.13725](#)] [[INSPIRE](#)].
- [68] M.S. Cagliari et al., *Bispectrum constraints on Primordial Non-Gaussianities with the eBOSS DR16 quasars*, *JCAP* **07** (2025) 043 [[arXiv:2502.14758](#)] [[INSPIRE](#)].
- [69] M. Rezaie et al., *Local primordial non-Gaussianity from the large-scale clustering of photometric DESI luminous red galaxies*, *Mon. Not. Roy. Astron. Soc.* **532** (2024) 1902 [[arXiv:2307.01753](#)] [[INSPIRE](#)].
- [70] SIMONS OBSERVATORY collaboration, *The Simons Observatory: Science goals and forecasts*, *JCAP* **02** (2019) 056 [[arXiv:1808.07445](#)] [[INSPIRE](#)].
- [71] SIMONS OBSERVATORY collaboration, *The Simons Observatory: science goals and forecasts for the enhanced Large Aperture Telescope*, *JCAP* **08** (2025) 034 [[arXiv:2503.00636](#)] [[INSPIRE](#)].
- [72] C.E. Villagra et al., *The Atacama Cosmology Telescope: high-redshift measurement of structure growth from the cross-correlation of Quiaa quasars and CMB lensing from ACT DR6 and Planck PR4*, *JCAP* **12** (2025) 033 [[arXiv:2507.08798](#)] [[INSPIRE](#)].
- [73] D. Alonso, *Science-driven 3D data compression*, *Mon. Not. Roy. Astron. Soc.* **473** (2018) 4306 [[arXiv:1707.08950](#)] [[INSPIRE](#)].
- [74] M. Schmittfull and U. Seljak, *Parameter constraints from cross-correlation of CMB lensing with galaxy clustering*, *Phys. Rev. D* **97** (2018) 123540 [[arXiv:1710.09465](#)] [[INSPIRE](#)].

- [75] E. Sefusatti and E. Komatsu, *The Bispectrum of Galaxies from High-Redshift Galaxy Surveys: Primordial Non-Gaussianity and Non-Linear Galaxy Bias*, *Phys. Rev. D* **76** (2007) 083004 [[arXiv:0705.0343](#)] [[INSPIRE](#)].
- [76] A. Moradinezhad Dizgah et al., *Primordial Non-Gaussianity from Biased Tracers: Likelihood Analysis of Real-Space Power Spectrum and Bispectrum*, *JCAP* **05** (2021) 015 [[arXiv:2010.14523](#)] [[INSPIRE](#)].
- [77] G. D’Amico, M. Lewandowski, L. Senatore and P. Zhang, *Limits on primordial non-Gaussianities from BOSS galaxy-clustering data*, *Phys. Rev. D* **111** (2025) 063514 [[arXiv:2201.11518](#)] [[INSPIRE](#)].
- [78] G. Cabass et al., *Constraints on multifield inflation from the BOSS galaxy survey*, *Phys. Rev. D* **106** (2022) 043506 [[arXiv:2204.01781](#)] [[INSPIRE](#)].
- [79] L. Harscouet et al., *Fast Projected Bispectra: the filter-square approach*, [arXiv:2409.07980](#) [[DOI:10.33232/001c.128309](#)] [[INSPIRE](#)].
- [80] P. Virtanen et al., *SciPy 1.0 — Fundamental Algorithms for Scientific Computing in Python*, *Nature Meth.* **17** (2020) 261 [[arXiv:1907.10121](#)] [[INSPIRE](#)].
- [81] C.R. Harris et al., *Array programming with NumPy*, *Nature* **585** (2020) 357 [[arXiv:2006.10256](#)] [[INSPIRE](#)].
- [82] J.D. Hunter, *Matplotlib: A 2D Graphics Environment*, *Comput. Sci. Eng.* **9** (2007) 90 [[INSPIRE](#)].
- [83] A. Zonca et al., *healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in Python*, *J. Open Source Softw.* **4** (2019) 1298 [[INSPIRE](#)].
- [84] 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](#)].
- [85] S. Vanneste, J.-C. Hamilton and J. Grain, *xqml: A pure python implementation of the qml method*, (2021).
- [86] PLANCK collaboration, *Planck 2018 results. VIII. Gravitational lensing*, *Astron. Astrophys.* **641** (2020) A8 [[arXiv:1807.06210](#)] [[INSPIRE](#)].
- [87] S. Gratton and A. Challinor, *Understanding parameter differences between analyses employing nested data subsets*, *Mon. Not. Roy. Astron. Soc.* **499** (2020) 3410 [[arXiv:1911.07754](#)] [[INSPIRE](#)].