



JCAP ANNIVERSARY  
SPECIAL ISSUE

## Constraining cosmology with the *Gaia-unWISE* Quasar Catalog and CMB lensing: structure growth

David Alonso,<sup>a,\*</sup> Giulio Fabbian,<sup>b,c,\*</sup> Kate Storey-Fisher,<sup>d,\*</sup> Anna-Christina Eilers,<sup>e</sup>  
Carlos García-García,<sup>a</sup> David W. Hogg<sup>b,d,f</sup> and Hans-Walter Rix<sup>f</sup>

<sup>a</sup>Department of Physics, University of Oxford, Denys Wilkinson Building,  
Keble Road, Oxford OX1 3RH, United Kingdom

<sup>b</sup>Center for Computational Astrophysics, Flatiron Institute,  
162 Fifth Avenue, New York, NY, 10010, U.S.A.

<sup>c</sup>School of Physics and Astronomy, Cardiff University,  
The Parade, Cardiff, Wales CF24 3AA, United Kingdom

<sup>d</sup>Center for Cosmology and Particle Physics, Department of Physics, New York University,  
726 Broadway, New York, NY 10003, U.S.A.

<sup>e</sup>MIT Kavli Institute for Astrophysics and Space Research,  
77 Massachusetts Avenue, Cambridge, MA 02139, U.S.A.

<sup>f</sup>Max Planck Institute for Astronomy,  
Königstuhl 17, 69117 Heidelberg, Germany

E-mail: [david.alonso@physics.ox.ac.uk](mailto:david.alonso@physics.ox.ac.uk), [FabbianG@cardiff.ac.uk](mailto:FabbianG@cardiff.ac.uk),  
[k.sf@nyu.edu](mailto:k.sf@nyu.edu), [eilers@mit.edu](mailto:eilers@mit.edu), [carlos.garcia-garcia@physics.ox.ac.uk](mailto:carlos.garcia-garcia@physics.ox.ac.uk),  
[david.hogg@nyu.edu](mailto:david.hogg@nyu.edu), [rix@mpia.de](mailto:rix@mpia.de)

ABSTRACT: We study the angular clustering of Quasars, a *Gaia*- and *unWISE*-based catalog of over a million quasars with an exceptionally well-defined selection function. With it, we derive cosmology constraints from the amplitude and growth of structure across cosmic time. We divide the sample into two redshift bins, centered at  $z = 1.0$  and  $z = 2.1$ , and measure both overdensity auto-correlations and cross-correlations with maps of the Cosmic Microwave Background convergence measured by *Planck*. From these data, and including a prior from measurements of the baryon acoustic oscillations scale, we place constraints on the amplitude of the matter power spectrum  $\sigma_8 = 0.766 \pm 0.034$ , and on the matter density parameter  $\Omega_m = 0.343^{+0.017}_{-0.019}$ . These measurements are in reasonable agreement with *Planck*

\*Corresponding author.

at the  $\sim 1.4\sigma$  level, and are found to be robust with respect to observational and theoretical uncertainties. We find that our slightly lower value of  $\sigma_8$  is driven by the higher-redshift sample, which favours a low amplitude of matter fluctuations. We present plausible arguments showing that this could be driven by contamination of the CMB lensing map by high-redshift extragalactic foregrounds, which should also affect other cross-correlations with tracers of large-scale structure beyond  $z \sim 1.5$ . Our constraints are competitive with those from state-of-the-art  $3\times 2$ -point analyses, but arise from a range of scales and redshifts that is highly complementary to those covered by cosmic shear data and most galaxy clustering samples. This, coupled with the unprecedented combination of volume and redshift precision achieved by Quiaia, allows us to break the usual degeneracy between  $\Omega_m$  and  $\sigma_8$ .

KEYWORDS: cosmological parameters from LSS, galaxy clustering, weak gravitational lensing

ARXIV EPRINT: [2306.17748](https://arxiv.org/abs/2306.17748)

---

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Data</b>	<b>2</b>
2.1	The <i>Gaia-unWISE</i> Quasar Catalog	2
2.2	<i>Planck</i> CMB lensing	3
<b>3</b>	<b>Theory</b>	<b>3</b>
3.1	Projected galaxy clustering	3
3.2	CMB weak lensing	5
3.3	Statistics of projected tracers	5
3.4	The importance of high-redshift data	6
<b>4</b>	<b>Analysis methods</b>	<b>9</b>
4.1	Quasar overdensity maps	9
4.1.1	From catalog to maps	9
4.1.2	Magnification bias	10
4.1.3	The redshift distribution	12
4.1.4	Bias evolution	13
4.2	The CMB $\kappa$ map	14
4.3	Angular power spectra and covariances	14
4.4	Likelihood and priors	16
<b>5</b>	<b>Results</b>	<b>18</b>
5.1	Fiducial cosmological constraints	18
5.2	Robustness tests	21
5.2.1	Internal consistency	21
5.2.2	Bias model	24
5.2.3	Redshift uncertainties	26
5.2.4	Sky contamination	28
5.2.5	Magnification bias	30
5.2.6	Foreground contamination in CMB lensing	31
5.2.7	Impact of priors	33
<b>6</b>	<b>Conclusions</b>	<b>34</b>
<b>A</b>	<b>Simulation-based assessment of CMB lensing contamination</b>	<b>36</b>

---

## 1 Introduction

Much of the progress in constraining the physical parameters governing the initial conditions and evolution of our Universe is currently driven by the analysis of tracers of the large-scale structure. In these analyses, we study the spatial distribution and time evolution of various tracers of the matter density fluctuations, which allows us to constrain the Universe’s geometry, as well as the growth of structure within it. Two of the most powerful large-scale structure tracers are weak gravitational lensing and galaxy clustering [1–3]. The former provides largely unbiased maps of the matter fluctuations integrated along the line of sight from the source redshift, while the latter is a biased tracer of these fluctuations at the redshifts of the galaxies being observed.

This complementarity (unbiased vs. biased, cumulative in redshift vs. local) motivates the combined analysis of weak lensing and galaxy clustering data in a technique commonly known as *lensing tomography* (often also labelled “2×2-point” analysis) [4–7]. In this methodology one measures the two-point auto-correlations of a set of samples of galaxies that are reasonably localised in redshift, together with their cross-correlations with lensing convergence or shear maps from a background source. The cross-correlations effectively “slice” the lensing map into its contributions at different redshifts and, when combined with the auto-correlations, enables us to break the degeneracy between galaxy bias and the amplitude and growth of structure. The same technique can be applied to any projected tracer of the large-scale structure other than weak lensing, and can be used to e.g. measure the evolution in the mean gas pressure [8–12] and star-formation density [13, 14], or to reconstruct the redshift distribution of different probes [15, 16].

The lensing of the Cosmic Microwave Background (CMB) has become a treasure trove for lensing tomography. At  $z \sim 1100$ , the CMB is a lensing source that sits behind all tracers of the large-scale structure, and thus can be used to reconstruct the growth history at all redshifts. Various analyses have made use of lensing tomography to constrain cosmology, using both photometric [17–22] and spectroscopic samples [7, 23, 24]. Similar studies can be carried out in cross-correlation with galaxy weak lensing (cosmic shear), as well as combining all three tracers of structure [25, 26]. Constraints can be strengthened significantly by adding the auto-correlation of the weak lensing data, in what’s commonly called a “3 × 2-point” analysis (sometimes called 5 × 2-point when combining galaxy clustering, cosmic shear, and CMB lensing) [2, 3, 22, 25–31]. It is often not possible to fully break the degeneracy between growth and geometry with these datasets, and thus results are often reported in terms of the combination  $S_8 \equiv \sigma_8(\Omega_m/0.3)^{0.5}$ , where  $\sigma_8$  is defined as the standard deviation of the linear matter overdensity on spheres of radius 8 Mpc/ $h$ , and  $\Omega_m$  is the fractional energy density of non-relativistic matter. A growing number of these analyses have consistently recovered measurements of the amplitude of matter fluctuations that, qualitatively, seem to be in tension with the value inferred by CMB experiments such as *Planck* [32], or the Atacama Cosmology Telescope (ACT) [33, 34] (the so-called “ $S_8$  tension”). There are hints that the source of this tension may lie in the impact of non-linear baryonic effects, particularly in the case of cosmic shear data, which is more sensitive to small-scale structures ( $k \sim 0.1\text{--}1 \text{ Mpc}^{-1}$ ) [35–37]. Nevertheless, similar levels of tension have been found from pure CMB lensing tomography analyses [19, 21], and thus the origin and actual significance of this potential tension remains unclear.

In this context, quasars constitute an interesting probe of the large-scale structure, allowing us to trace the matter fluctuations at significantly larger redshifts than standard optical galaxy surveys, covering commensurately large volumes. Their study has been carried out through three-dimensional analysis of their auto-correlation [38–40], as well as their cross-correlation with other tracers, such as the Lyman- $\alpha$  forest [41], with the aim of studying both cosmology and the astrophysics of quasars. Since its first detection [42], the cross-correlation between quasars and CMB lensing maps has increasingly gained interest, particularly using quasar samples selected from the Sloan Digital Sky Survey [43–45], and as a way to improve our understanding of quasar properties. The use of this cross-correlation for cosmology, however, has so far been limited, although some works have exploited it to constrain the growth history at high redshifts [25], as well the level of primordial non-Gaussianity [46]. The availability of a full-sky quasar sample with a well-understood selection function would facilitate this kind of studies, as well as increase the sensitivity of these measurements by increasing the area overlap with existing CMB lensing maps from experiments such as *Planck* [47, 48].

In this paper we will study CMB lensing tomography making use of Quaia [49], a catalog of high-redshift quasars constructed based on the third data release of the *Gaia* satellite [50] as well as *unWISE* infrared observations [51–53]. Covering the redshift range  $z \lesssim 4$  with good redshift accuracy ( $|\Delta z/(1+z)| < 0.01$  (0.1) for 62% (83%) of the sources), and spanning the full sky, Quaia allows us to constrain the growth of structure over a range of cosmic times that is highly complementary to those accessed by previous analyses, as well as giving us access to significantly larger scales, where the impact of non-linearities in the matter power spectrum is less important. This complementarity is highly relevant to ascertain whether the source of the  $S_8$  tension lies in small-scale physics at low redshifts, an overall lack of power on all scales throughout the Universe’s evolution, or a combination of statistical and systematic uncertainties. Furthermore, as we will show, the large-scale reach of the sample will allow us to measure  $\sigma_8$  and  $\Omega_m$  independently (instead of through the  $S_8$  combination).

The paper is structured as follows. The Quaia sample and the *Planck* CMB lensing map used in the analysis are described in section 2. Section 3 presents the theoretical framework underpinning CMB lensing tomography. The data analysis methods used are outlined in section 4. Section 5 presents the main results of this analysis, as well as a large battery of tests aimed at confirming their robustness. Our conclusions are then summarised in section 6.

## 2 Data

### 2.1 The *Gaia-unWISE* Quasar Catalog

We use the *Gaia-unWISE* Quasar Catalog, Quaia [49], as our quasar sample. The catalog is all-sky and highly homogeneous, with 1,295,502 sources brighter than its magnitude limit of  $G < 20.5$  and a median redshift of  $z = 1.47$ . It is derived from the quasar candidates identified in the third data release of the space-based *Gaia* mission, which classified  $\sim 6.6$  million sources as potential quasars and estimated their redshifts *Gaia*’s low-resolution BP/RP spectra. Quaia combines *Gaia* data with mid-infrared information from the *unWISE* reprocessing of the Wide-field Infrared Survey Explorer (WISE) observations [51–53] to improve the sample selection and redshift estimation. The redshift estimation uses a  $k$ -Nearest Neighbors model

with *Gaia* and *unWISE* photometry and the *Gaia*-estimated spectral redshift as features, and cross-matched SDSS DR16Q quasars with high-precision spectroscopic redshifts [54] as labels. The Quaia redshift estimate,  $z_{\text{Quaia}}$ , is taken to be the median of the spectroscopic redshifts of the  $K = 27$  nearest neighbors, and the uncertainty  $\sigma_z$  is their standard deviation. Of the Quaia redshifts, an estimated 62% (83%) have  $|\Delta z/(1+z)| < 0.01$  (0.1) with respect to the SDSS redshifts, and 91% having  $|\Delta z/(1+z)| < 0.2$ .

In this analysis we use a modified version of the Quaia-modelled selection function, re-fit for each of the two redshift bins used for our fiducial results (see section 4.1). While Quaia already has relatively minimal selection effects, partly thanks to its space-based observations, the selection function model includes the effects of dust extinction, stellar crowding, and the *Gaia* scanning law (as detailed in [49]). In section 5.2.1 we will also explore a brighter sample with  $G < 20.0$ , for which the selection function was also refit.

## 2.2 *Planck* CMB lensing

Our baseline analysis uses the CMB lensing map reconstructed with the Planck PR4 data release based on the NPIPE processing pipeline [55]. The release improved over various assumptions of the PR3 2018 analysis [47] that made the analysis sub-optimal in terms of noise treatment and exploited the improved low-level data processing of NPIPE (encompassing additional data and more accurate simulations of the end-to-end data processing) to reach a  $\sim 20\%$  improvement in signal-to-noise at all scales. The details of the analysis are summarized in [48, 56]. The major pipeline improvements included the adoption of a more optimal generalized minimum-variance estimator (GMV) that performs a joint inverse-variance Wiener filtering of the temperature and polarization CMB maps at the same time instead of treating both observables separately [57], as well as an additional post-processing Wiener filtering of the lensing maps according to the local reconstruction noise level on the sky that improves the sensitivity of signal-dominated scales. The PR4 lensing release also includes an estimate of the lensing potential based on the 2018 pipeline and using CMB temperature-only and polarization-only data [58]. We use these CMB lensing maps to assess the impact of foregrounds in our results, given that the extragalactic foregrounds such as thermal Sunyaev-Zeldovich (tSZ) effect and the Cosmic Infrared Background (CIB) are largely unpolarized at this level of sensitivity. We used the PR3 and PR4 associated lensing simulations to evaluate the consistency between the measurements carried out with the different lensing maps, and to evaluate mask-dependent correction to the lensing map normalization. This is discussed in section 4.2, and in the PR4 lensing documentation.<sup>1</sup> All the data products are publicly available through the Planck legacy archive.<sup>2</sup>

## 3 Theory

### 3.1 Projected galaxy clustering

The first probe we will use in our analysis is the projected distribution of quasars. The core observable in this case is the angular quasar overdensity<sup>3</sup>  $\delta_g(\hat{\mathbf{n}}) \equiv n_g(\hat{\mathbf{n}})/\bar{n}_g - 1$ , where

<sup>1</sup>See [https://github.com/carronj/planck\\_pr4\\_lensing](https://github.com/carronj/planck_pr4_lensing).

<sup>2</sup><http://pla.esac.esa.int/>.

<sup>3</sup>Note that, although normally reserved to label “galaxies”, we will use the subscript  $g$  to denote any quantity related to the quasar samples used.

$n_g(\hat{\mathbf{n}})$  is the surface density of quasars in the sky direction defined by unit vector  $\hat{\mathbf{n}}$ , and  $\bar{n}_g$  is its ensemble average. The angular overdensity is related to the three-dimensional overdensity  $\Delta_g(\mathbf{x}, t)$  (defined analogously to the angular quasar overdensity), where  $\mathbf{x}$  is the tracer position, through the line-of-sight (LOS) projection

$$\delta_g(\hat{\mathbf{n}}) = \int d\chi H(z) p(z) \Delta_g(\chi \hat{\mathbf{n}}, t(z)), \quad (3.1)$$

where  $p(z)$  is the redshift distribution of quasars in the sample,  $z$  is the redshift to radial comoving distance  $\chi$ , and  $H(z)$  is the expansion rate at redshift  $z$ .<sup>4</sup>

The relation between the quasar overdensity and the underlying matter overdensity  $\Delta_m$  is, in general, non-linear, non-local, and stochastic. On sufficiently large-scales ( $r \gtrsim 10$  Mpc), such as those used in this work, a linear bias relation can be used, so that

$$\Delta_g = b_g \Delta_m + \epsilon, \quad (3.2)$$

where  $b_g$  is the linear quasar bias and  $\epsilon$  is an uncorrelated stochastic field representing the non-deterministic component of the  $\Delta_g$ - $\Delta_m$  relation. In our fiducial analysis, we will assume that this component is dominated by Poisson shot noise, in which case the power spectrum of  $\epsilon$  is simply  $P_\epsilon(k) = 1/\bar{N}_g$ , where  $\bar{N}_g$  is the mean 3-dimensional quasar density. In general,  $\epsilon$  may receive additional contributions due to physics on scales below the Lagrangian size of the protohalo hosting the galaxies composing our samples (e.g. halo exclusion effects). The low density of sources in the quasar sample used here makes it unlikely for these contributions to dominate over the pure Poisson term, but we will study the potential impact of stochastic bias in section 5.2.2.

The observed projected overdensity of sources is also affected by modifications to the properties of the light they emit caused by the large-scale structure. These contributions may be local (e.g. additional red- or blue-shifting caused by peculiar velocities or gravitational potentials), or LOS-integrated (e.g. lensing, integrated Sachs-Wolfe effect, Shapiro time delay) [59, 60]. Of these, the most relevant effect for distant sources with broad radial projection kernels, such as the samples used here, is the contribution from lensing magnification [61, 62]. Gravitational lensing causes a coherent distortion in the observed positions of background sources *away* from foreground overdensities, as well as modifies their observed flux, which can add sources to or remove sources from the sample (depending on the slope of the flux distribution, and on the sign of the lensing convergence). The combination of both effects leads to an extra additive contribution to  $\delta_g$  in eq. (3.1) of the form

$$\delta_g^{\text{mag}}(\hat{\mathbf{n}}) = \int d\chi (5s - 2) q_{\text{mag}}(\chi) \nabla_\perp^2 \nabla^{-2} \Delta_m(\chi \hat{\mathbf{n}}, t(z)), \quad (3.3)$$

where the lensing magnification kernel is

$$q_{\text{mag}}(\chi) \equiv \frac{3H_0^2 \Omega_m}{2} (1+z) \chi \int_{z(\chi)}^\infty dz' p(z') \frac{\chi(z') - \chi}{\chi(z')}, \quad (3.4)$$

and  $s(z)$  is the logarithmic slope of the cumulative number counts of sources at the magnitude limit of the sample:

$$s \equiv \left. \frac{d \log_{10} N(< m)}{dm} \right|_{m_{\text{lim}}}, \quad (3.5)$$

---

<sup>4</sup>Note that we use natural units with  $c = 1$  throughout.

where  $m$  is the source apparent magnitude, and  $m_{\text{lim}}$  is the sample's limiting magnitude. Section 4.1.2 describes the methods used to calculate  $s$  in practice. Finally,  $\nabla_{\perp}^2$  and  $\nabla^{-2}$  are the transverse Laplacian and the inverse Laplacian respectively. In the Limber approximation (see eq. (3.11)), they partially cancel each other up to a factor

$$K_{\ell} \equiv \frac{\ell(\ell+1)}{(\ell+1/2)^2} = 1 - \frac{1}{(2\ell+1)^2}. \quad (3.6)$$

which can be safely ignored on scales  $\ell > 10$  (although our theoretical calculations will account for this factor). We will therefore ignore the combination  $\nabla_{\perp}^2 \nabla^{-2}$  in what follows for simplicity.

Finally, it is worth noting that eq. (3.3) assumes that the matter overdensity is related to the Newtonian potential  $\Phi$  via Poisson's equation, and that the two scalar metric potentials in the Newtonian gauge,  $\Phi$  and  $\Psi$  are equal (which is valid for general relativity in the absence of anisotropic stress).

### 3.2 CMB weak lensing

Gravitational lensing perturbs the trajectories of CMB photons, thus distorting the observed pattern of CMB anisotropies, and inducing a statistical coupling between different harmonic scales. This distortion can be used to reconstruct the lensing displacement that causes it, and thus produce maps of the lensing convergence  $\kappa(\hat{\mathbf{n}})$ , defined as

$$\kappa(\hat{\mathbf{n}}) = \frac{3H_0^2 \Omega_m}{2} \int_0^{\chi_{\text{LSS}}} d\chi (1+z) \chi \frac{\chi_{\text{LSS}} - \chi}{\chi_{\text{LSS}}} \nabla_{\perp}^2 \nabla^{-2} \Delta_m(\chi \hat{\mathbf{n}}, t(z)). \quad (3.7)$$

Here,  $\chi_{\text{LSS}}$  is the comoving distance to the last-scattering surface (LSS). As in the case of lensing magnification, CMB lensing is subject to the same  $K_{\ell}$  correction factor, which can be neglected on the scales used here.

### 3.3 Statistics of projected tracers

Consider a projected tracer  $u(\hat{\mathbf{n}})$  of a three-dimensional quantity  $U(\mathbf{x}, z)$ :

$$u(\hat{\mathbf{n}}) \equiv \int d\chi q_u(\chi) U(\chi \hat{\mathbf{n}}, z(\chi)), \quad (3.8)$$

where  $q_u(\chi)$  is the associated projection kernel. Assuming statistical homogeneity and isotropy, let us define the angular power spectrum  $C_{\ell}^{uv}$ , and the three-dimensional power spectrum  $P_{UV}(k, z)$  between two pairs of such quantities ( $(u, U)$  and  $(v, V)$ ) as:

$$\langle u_{\ell m} v_{\ell' m'}^* \rangle \equiv \delta_{\ell\ell'}^K \delta_{mm'}^K C_{\ell}^{uv}, \quad \langle U(\mathbf{k}, z) V(\mathbf{k}', z) \rangle \equiv (2\pi)^3 \delta^D(\mathbf{k} + \mathbf{k}') P_{UV}(k, z). \quad (3.9)$$

Here  $\langle \dots \rangle$  denotes ensemble averaging,  $\delta^K$  is the Kronecker delta,  $\delta^D$  is the Dirac delta, and  $u_{\ell m}$  and  $U(\mathbf{k}, z)$  are the spherical harmonic transform and the Fourier transform, respectively, of the corresponding configuration-space fields:

$$u_{\ell m} \equiv \int d\hat{\mathbf{n}} Y_{\ell m}^*(\hat{\mathbf{n}}) u(\hat{\mathbf{n}}), \quad U(\mathbf{k}, z) \equiv \int d^3x e^{i\mathbf{k}\cdot\mathbf{x}} U(\mathbf{x}, z). \quad (3.10)$$

The angular power spectrum can be connected with the three-dimensional  $P_{UV}(k)$  via

$$C_{\ell}^{uv} = \int \frac{d\chi}{\chi^2} q_u(\chi) q_v(\chi) P_{UV}\left(\frac{\ell+1/2}{\chi}, z(\chi)\right). \quad (3.11)$$



This equation is valid in the Limber approximation [63, 64], which is valid when the kernels under study overlap over a range of  $\chi$  significantly larger than the correlation length of  $U$  and  $V$ . This is an excellent approximation for the broad CMB lensing kernel, as well as the redshift distribution of the quasar samples we will study.

Within the linear bias model we use here, both the projected quasar overdensity  $\delta_g$  and the lensing convergence  $\kappa$  can be treated as tracers of the 3D matter overdensities  $\Delta_m$ , and therefore  $P_{UV}(k, z)$  above is simply the matter power spectrum for all correlations studied. The radial kernels for both tracers are:

$$q_g(\chi) = H(z)p(z)b_g + (5s - 2)q_{\text{mag}}(\chi), \quad q_\kappa(\chi) = \frac{3H_0^2\Omega_m}{2}(1+z)\chi\frac{\chi_{\text{LSS}} - \chi}{\chi_{\text{LSS}}}, \quad (3.12)$$

with the magnification kernel  $q_{\text{mag}}$  given in eq. (3.4).


### 3.4 The importance of high-redshift data

With the theoretical model described above, we can now explore the benefits of studying the growth of structure at high redshifts, as enabled by the Quiaa quasar sample we use here.

One of the most compelling reasons to study the history of structure growth, is the ongoing debate surrounding the so-called “ $S_8$  tension”: the fact that a number of low-redshift probes of structure recover an amplitude of the matter fluctuations that is in slight tension with the value inferred from CMB anisotropies. The evidence of this tension from individual experiments, and from some combinations of them, is mild, although not negligible (between  $2\sigma$  and  $3\sigma$ ), and the fact that several different probes have consistently found low values of the  $S_8$  parameter has motivated a widespread search for both further evidence of this tension from other datasets, as well as potential explanations for it in terms of both astrophysical effects and potential new physics.

The  $S_8$  tension is largely dominated by cosmic shear measurements [25, 65–69]. The weak lensing kernel for sources at  $z \lesssim 1$  peaks at  $z \sim 0.2$ – $0.5$ , and extends down to  $z \sim 0$  [70]. Thus, cosmic shear is mostly sensitive to the matter inhomogeneities at low redshifts and, due to projection effects, on relatively small scales  $k \sim 0.1$ – $1$  Mpc. On the other hand, probes such as the CMB lensing convergence auto-correlation, which access significantly higher redshifts (the CMB lensing kernel peaks at  $z \sim 2$ ), and are sensitive to larger, more linear scales, find no evidence of a low  $S_8$  [34]. This has motivated the exploration of the mis-modelling of the small-scale matter power spectrum due to non-linearities and baryonic effects as a plausible solution to the  $S_8$  tension [35, 37]. Additional, complementary analyses of structure formation at high redshifts and on linear scales are important in order to investigate this hypothesis; the clustering of high-redshift quasars offers such an opportunity.

Regardless of its potential ability to shed light on the ongoing  $S_8$  tension, the study of large-scale structure at redshifts  $z \sim 2$ – $4$  is interesting for cosmology in its own right. Although we expect the Universe’s expansion to be dominated by non-relativistic matter beyond  $z \sim 1$ , direct observations of this epoch are few and far between. The matter fluctuations at  $z > 2$  could therefore bear the imprint of new early-Universe physics that we are not sensitive to at lower redshift. At the same time, the comoving volume available at high- $z$  gives us access to significantly larger scales than those probed by low-redshift tracers. Figure 1 illustrates this



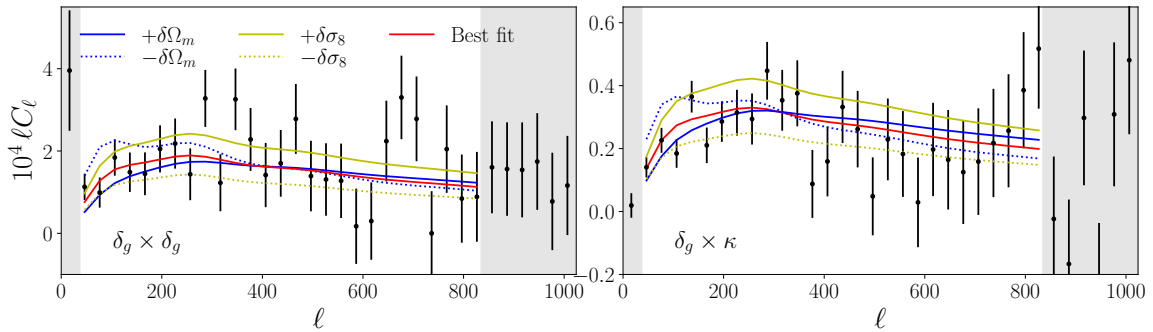
**Figure 1.** The matter power spectrum as observed by a hypothetical probe sensitive to different redshifts, as a function of the corresponding angular scale  $\ell \equiv k\chi - 1/2$ , where  $\chi$  is the comoving distance to that redshift. The central white band roughly corresponds to the range of scales used in cosmic shear analyses. Within these scales, at low redshifts we are mostly sensitive to the monotonic power-law-like part of the power spectrum, which is also affected by non-linear effects (the scale of non-linearities  $k_{\text{NL}}$ , defined in eq. (3.13), is shown by the black solid line). At higher redshifts we are able to better reconstruct the overall shape of the power spectrum at the accessible scales, and to use scales that are less affected by non-linearities. For reference, the plot also shows the position of  $k = 0.1 \text{ Mpc}^{-1}$  and  $k = 0.01 \text{ Mpc}^{-1}$ , roughly corresponding to the location of the BAO wiggles and the horizon scale at matter-radiation equality, respectively.

point. The figure shows the matter power spectrum as it would be observed by an imaginary probe of structure at different redshifts,<sup>5</sup> as a function of the angular multipole, related to the comoving wavenumber via  $k\chi \equiv \ell + 1/2$ , where  $\chi$  is the distance to the redshift being probed. The coloured lines show the power spectrum at different redshifts in the range  $0.3 < z < 4$ . The solid black line shows the angular multipole corresponding to the non-linear scale  $k_{\text{NL}}$ , defined as the comoving scale at which the overdensity variance reaches unity:

$$\sigma^2(< k_{\text{NL}}(z)) \equiv \int_0^{k_{\text{NL}}(z)} \frac{dk k^2}{2\pi^2} P_{mm}(k, z) \equiv 1. \quad (3.13)$$

The power spectra are shown in transparent colours above this scale, which marks the regime in which non-linear and baryonic effects become relevant. The dashed black line shows the position of the wavenumber  $k = 0.1 \text{ Mpc}^{-1}$ , roughly corresponding to the range of scales where baryon acoustic oscillations are most relevant. The dotted black line shows the wavenumber  $k = 0.01 \text{ Mpc}^{-1}$ , which corresponds approximately to the horizon scale at the matter-radiation equality, where the power spectrum reaches its maximum. Finally, the shaded gray bands mask out the scales  $\ell \lesssim 30$ , and  $\ell \gtrsim 1000$  (corresponding to  $6^\circ$  and  $10'$  respectively), usually discarded in cosmic shear analyses. As we can see, at low redshifts

<sup>5</sup>We multiply the power spectrum by  $1/(1+z)^3$  in order to better distinguish the shape measured at each redshift.



**Figure 2.** Quasar auto-spectrum (left panel) and cross-spectrum with the CMB lensing convergence (right panel) for the high-redshift bin used in our analysis, centered at redshift  $z = 2.1$  (see figure 4 and section 4.1). The measurements are shown as black points with error bars, with the best-fit model shown in red. The blue and yellow lines show the theoretical predictions after changing  $\Omega_m$  and  $\sigma_8$  by 0.1 in both directions, respectively. On the redshifts and scales our data are sensitive to, the effect of both parameters is markedly different, allowing us to break the degeneracy between them.

$z \lesssim 0.5$ , which dominate current cosmic shear data, we mostly have access to the range of scales where the power spectrum is approximately a monotonic power law (projection effects in weak lensing wash out the BAO wiggles), which also receives significant contribution from non-linear scales. In turn, at higher redshifts ( $z \gtrsim 1$ ) we gain access to larger scales, and are able to better resolve the shape of the power spectrum, reducing also the contribution from non-linear scales. Measurements at these high redshifts thus can allow us to break degeneracies between different cosmological parameters, which modify the power spectrum in complementary ways.

To illustrate this further, let us briefly jump ahead and examine some of the data we will analyse later on. The left panel of figure 2 shows the angular auto-correlation of the high-redshift bin we use in our analysis, centered at  $z \simeq 2$  (see figure 4 and section 4.1), and the right panel shows the cross-correlation of the same bin with the *Planck* CMB lensing map. The measured power spectra are shown as black points with error bars, and the red line shows the best-fit model describing these data. The blue solid and dotted lines show the result of changing the cosmic matter fraction  $\Omega_m$  by  $\delta\Omega_m = \pm 0.1$  with respect to its best fit value, while the analogous yellow lines show the result of perturbing  $\sigma_8$  by the same amount. We can see that, although on small scales ( $\ell \gtrsim 400$ ) both parameters modify the measured power spectra in a similar manner (scaling it up and down), when including the full range of scales used in our analysis (marked by the gray dashed bands), their impact is markedly different. While  $\sigma_8$  modifies the overall normalisation of the power spectrum,  $\Omega_m$  changes its spectral tilt.<sup>6</sup> As we will see, this will allow us to break the degeneracy between these parameters, which affects cosmic shear and lower-redshift data sets in general, and measure them independently of one another.

<sup>6</sup>It also changes the amplitude of the BAO peaks through the baryon fraction  $f_b \equiv \Omega_b/\Omega_m$ , as we keep  $\Omega_b$  constant, although we are not sensitive to the BAO wiggles with these data.

## 4 Analysis methods

### 4.1 Quasar overdensity maps

Our fiducial quasar sample will be the full  $G < 20.5$  Quiaia catalog, which we divide into two redshift bins ( $z_{\text{Quaia}} < 1.47$  and  $z_{\text{Quaia}} > 1.47$ ) each containing approximately the same number of sources (647,749 and 647,753 respectively). Splitting the sample into two bins allows us to explore any trends with redshift (both astrophysical and due to systematics), but we determined that further dividing the sample did not lead to significantly tighter cosmological constraints. In what follows, we will refer to these two bins as “Low- $z$ ” and “High- $z$ ” respectively. The following subsections describe the methodology used to produce quasar overdensity maps, calibrate their contamination from observational systematics, estimate the slope of the cumulative flux distribution to quantify the impact of magnification bias, and calibrate the redshift distribution of each redshift bin.

#### 4.1.1 From catalog to maps

We construct a map of the projected quasar overdensity in each redshift bin using the **HEALPix** pixelization scheme [71], with resolution parameter  $N_{\text{side}} = 512$  (corresponding to an angular resolution  $\delta\theta \sim 0.11^\circ$ ). In each pixel  $p$  we calculate the quasar overdensity as

$$\delta_{g,p} = \frac{N_p}{\bar{N}w_p} - 1. \quad (4.1)$$

Here  $N_p$  is the number of quasars in  $p$ ,  $\bar{N}$  is the mean number density of quasars in the sample, and  $w_p$  is the selection function value in  $p$ . The selection function is described below in more detail, and can be interpreted as the mean fraction of observed objects (out of those that should have been observed). The mean number density is calculated correcting for the selection function as:

$$\bar{N} = \frac{\sum_p N_p}{\sum_p w_p}. \quad (4.2)$$

The method used to construct the selection function is described in detail in [49]. In short, a non-linear relation, modelled as a Gaussian process, is found between the number of sources observed in each pixel and the value of a number of systematics templates in those pixels. This relation was found using pixels of resolution  $N_{\text{side}} = 64$  ( $\delta\theta \simeq 1^\circ$ ), and considering 4 systematic templates: Milky Way extinction as measured by [72], a star map constructed from *Gaia*, the specific contribution to the star map from the LMC and SMC, and a measure of depth in *Gaia* (the so-called  $M_{10}$  quantity [73]) that maps the *Gaia* scanning law and source completeness. It is worth noting that the method intrinsically assumes that the underlying galaxy distribution is statistically isotropic.

To avoid instabilities when the selection function is close to zero, we mask all pixels where more than half of the sources are missed on average ( $w_p < 0.5$ ), and only the resulting unmasked pixels are used in eq. (4.2). Elsewhere, the quasar overdensity is set to zero, and the impact of this masking is accounted for via the pseudo- $C_\ell$  estimator (see section 4.3). The resulting footprint covers 74% and 67% of the sky in the Low- $z$  and High- $z$  bins, respectively.

The results presented here were found to be stable against changes in the selection function threshold. In order to downweight regions of the map where a smaller fraction of objects were observed, we again use the selection (truncated at  $w_p = 0.5$  as above) as a weighting map/mask when estimating power spectra involving  $\delta_g$ .

To further remove any residual correlation between the quasar overdensity map and the sky systematics listed above, we make use of linear deprojection, as described in [74, 75]. Effectively, we fit for, and project out, any linear relation between  $\delta_{g,p}$  and the fluctuations of the different systematic templates about their mean.<sup>7</sup> Since deprojecting only 4 templates leads to the loss of only a small number of modes, we do not correct for the associated deprojection bias when computing angular power spectra.

The use of linear deprojection after having calibrated the selection function using a Gaussian process (which to some extent should also reconstruct linear trends in the data), may seem redundant. However, it is worth emphasizing that, in the presence of small levels of contamination, where a linear expansion is sufficiently accurate, linear deprojection is exact, non-degenerate with other potential trends in the data, and further guarantees no correlation between the cleaned maps and the different systematic templates. We thus employ it to further reduce the impact of any residual contamination. As we will see in section 5.2.4, although we find that deprojection is indeed able to remove the presence of residual systematics in the data, within the scale cuts used in our analysis, their impact is negligible.

#### 4.1.2 Magnification bias

In order to quantify the impact of magnification on the measured power spectra, we need to estimate the slope of the cumulative source number counts with limiting magnitude  $s$  (see eq. (3.5)).

The most relevant cuts in the Quaia sample are in  $G$ -band magnitude, and in colour space. Magnification leads to a coherent shift in magnitude in all bands, and thus the  $G$ -band cut is sensitive to it, but not the colour-space cuts. Therefore, as a first method to estimate  $s$ , we treat Quaia as a purely magnitude-limited sample, and simply calculate the slope of the number counts at the limiting magnitude.<sup>8</sup> The resulting value of  $s$  for the two redshift bins, and for a magnitude limit  $G < 20.5$  is

$$s(z < 1.47) = 0.388 \pm 0.004, \quad s(z \geq 1.47) = 0.420 \pm 0.005. \quad (4.3)$$


The statistical uncertainties were estimated via bootstrap resampling. The result for a brighter sample with  $G < 20$  (which we will explore in section 5.2.1) is

$$s(z < 1.47, G < 20) = 0.468 \pm 0.005, \quad s(z \geq 1.47, G < 20) = 0.550 \pm 0.010. \quad (4.4)$$

More in detail, the Quaia sample is not exactly magnitude-limited, as its definition also involves a cut in the joint space of magnitude and proper motion. Thus, as an alternative estimate of  $s$ , we generate two additional catalogs by first perturbing all magnitudes in the raw

<sup>7</sup>Note that it is crucially important to subtract the mean from all systematic templates before deprojection since, by design, the overall sky mean is removed when constructing  $\delta_{g,p}$ .

<sup>8</sup>The raw Quaia sample extends beyond the magnitude cut  $G < 20.5$ , which allowed us to estimate  $s$  via finite differences, comparing the number of sources with  $G < 20.45$  and  $G < 20.55$ .



**Figure 3.** Number counts slope (see eq. (3.5)) governing the impact of lensing magnification, for our fiducial  $G < 20.5$  sample (black) and for the brighter  $G < 20$  sample (red), as a function of redshift. The dotted lines show the value of  $s$  found for the full sample (i.e. not separated in redshift). The value of  $s$  in our fiducial sample,  $s \sim 0.4$ , implies that the impact of lensing magnification on our analysis is negligible.


catalog by  $\delta G = \pm 0.05$ , and then imposing the cuts defining our sample.  $s$  is then estimated from the difference in the number of objects between these two catalogues. The result is:

$$s(z < 1.47) = 0.389 \pm 0.003, \quad s(z \geq 1.47) = 0.416 \pm 0.004, \quad (4.5)$$

in excellent agreement with the more naïve estimate above.

So far we have treated  $s$  as a constant in each redshift bin, whereas in reality it likely evolves with redshift. To explore this, we have divided the sample into 6 redshift bins covering the range  $z_{\text{Quaia}} \in [0, 4]$ , and calculated  $s$  from the number counts slope in each of them. The resulting  $s(z)$  for the  $G < 20.5$  sample is shown in figure 3. We see that  $s$  only departs from its global value ( $s = 0.404 \pm 0.004$  and  $s = 0.515 \pm 0.006$  for the  $G < 20.5$  and  $G < 20$  samples respectively) significantly at low redshifts, where the contribution from magnification is negligible. We thus assume  $s$  to be constant in our fiducial analysis.

Since, for our sample, the counts slope is close to  $s \sim 0.4$ , the impact of lensing magnification is small, and likely negligible (although we account for it in our theory prediction). This is in contrast with studies carried out with previous optical quasar samples, such as [76, 77], where lower values were found (e.g.  $s = 0.2764$  for the DESI-selected quasar sample [46]). This may be due to differences in the selection of the various samples, or to the fact that Quaia targets brighter quasars on average, and hence is sensitive to a different part of the quasar luminosity function. A value of  $s < 0.4$  corresponds to a negative contribution from magnification, which mostly affects the CMB lensing cross-correlation. To make up for this deficit, we would then infer a larger value of  $\sigma_8$ , hence the importance of accurately accounting for lensing magnification. In our case, however, we find that this effect can be largely neglected.



**Figure 4.** Redshift distribution of the quasar samples studied in this work. The points with error bars show the DIR estimate of the redshift distributions for the two redshift bins used in our fiducial analysis (red and blue respectively). The solid lines show the same redshift distributions estimated via PDF stacking, as described in the main text. The PDF-stacked redshift distribution of the full sample is shown in dashed black (as used in section 5.2.3).

In section 5.2.5, we will present an alternative test, using cross-correlations with a low-redshift sample, to validate our inferred value of  $s$  in comparison with the value found for other quasar samples.

#### 4.1.3 The redshift distribution

The redshift distribution  $p(z)$  of the samples under study is a crucial element of the theory prediction, and we must both estimate it and propagate the uncertainties of that estimate to the final inferred cosmological parameters. In this work we make use of two different estimates of the redshift distribution.

Our fiducial estimate of  $p(z)$  is based on the direct calibration (DIR) method of [78]. In this approach, a cross-matched spectroscopic sample is re-weighted to reproduce the magnitude/colour distribution of the target sample. The  $p(z)$  is then simply estimated as a weighted histogram of the spectroscopic redshifts. In our case, we make use of the cross-matched SDSS spectroscopic sample used to determine the Quaia redshift estimate  $z_{\text{Quaia}}$ , and re-weight it in the 5-dimensional space of magnitudes in the  $(G, BP, RP, W1, W2)$  bands using a  $k$ -nearest-neighbours algorithm. This spectroscopic sample contains a total of 243,206 sources. The weight of each spectroscopic quasar is proportional to the number of objects in Quaia within a sphere in magnitude space containing the 10 nearest spectroscopic neighbors (for a more detailed description of the method see e.g. section 2.2 of [79]). We verified that the re-weighted spectroscopic catalog follows the same magnitude distribution as Quaia in all four bands. We estimate the uncertainties in the resulting redshift distribution using jackknife resampling with 100 samples. The estimated  $p(z)$ s are shown in figure 4

as points with error bars for the two redshift bins used in our analysis. We propagate the uncertainties in the DIR redshift distribution making use of the approximate analytical marginalisation method proposed in [80, 81] (see also [82]). To account for any potential mis-calibration of the  $p(z)$  uncertainties, we verified that our final results did not change significantly when multiplying the jackknife errors by up to a factor of 4.

One potential drawback of the DIR-calibrated redshift distributions is the presence of stochastic noise associated with the specifics of the spectroscopic sample used (sample variance, shot noise, selection effects). If this is not correctly captured by the calibrated uncertainties (estimated via jackknife resampling above), this noise may lead to spurious shifts in the final inferred parameters. As an alternative to the DIR  $p(z)$ , we estimate the redshift distributions of both bins via *PDF stacking*. The Quiaia photometric redshifts  $z_{\text{Quaia}}$  carry an associated uncertainty  $\sigma_z$ . Assuming a Gaussian error model, we estimate the sample redshift distribution by adding one Gaussian probability distribution function (PDF)  $\mathcal{N}(z_{\text{Quaia}}, \sigma_z^2)$  for each quasar in the sample. The resulting redshift distribution is shown as solid lines in figure 4. We will only use this alternative estimate of the  $p(z)$  to quantify the sensitivity of our final results to mis-modelling of the true underlying redshift distribution (which, we will show, is small), and therefore, unlike in our fiducial case, we will not associate any uncertainties to it or marginalise over them.

As figure 4 shows, the redshift distributions exhibit non-negligible tails beyond the threshold  $z_{\text{Quaia}}$ , although both estimates of the  $p(z)$  are in qualitative agreement as to their extent. The final cosmological constraints are sensitive to the width of these tails, as they also control the amplitude of the  $p(z)$ , and hence it is important to ensure that they are correctly modelled. In section 5.2.3 we will make use of the cross-correlation between quasars in both bins as a sensibility check for the presence and extent of these tails.

#### 4.1.4 Bias evolution

Although calibrating the true redshift distribution is important in order to understand the samples under study, the power spectra ( $C_\ell^{gg}, C_\ell^{g\kappa}$ ) are in fact only sensitive to the combination  $b(z)p(z)$ . In other words, the final cosmological constraints depend not only on the redshift distribution, but also on the evolution of the sample’s bias with redshift.

The redshift evolution of the quasar clustering amplitude has been the subject of several works in the literature, including spectroscopic quasars in 2dF [83–85], SDSS [38, 43–45, 86–90] and DESI [46, 91], and photometrically-selected quasars [92–95]. Although there is significant scatter, especially at high redshifts, the general consensus is that quasar bias evolves strongly with redshift. Thus, although assuming a constant bias within each of the redshift bins studied here would potentially be a more agnostic approach, ignoring the potential evolution within each bin could lead to significant shifts in the final parameter constraints (although in section 5.2.2 we will show that our final constraints are rather resilient to this effect). Thus, in our fiducial analysis we will assume that the quasar bias within each redshift bin  $i$  takes the form

$$b(z) = b_g^i [0.278((1+z)^2 - 6.565) + 2.393], \quad (4.6)$$

with  $b_g^i$  a free parameter (one in each bin).  $b_g^i = 1$  would recover the best-fit bias model found by [38] from the eBOSS quasar sample. To quantify the impact of the assumed level



of bias evolution encoded in this model, in section 5.2.2 we will explore the constant-bias case, and a steeper evolution model, used in [93].

Another method to account for the redshift evolution of the sample bias is the so-called *clustering redshifts* approach. This technique [15, 96–98] is able to precisely constrain the combination  $b(z)p(z)$  from the cross-correlations of the target galaxy sample with a series of spectroscopic (or spectrophotometric) samples with well-determined redshifts. In section 5.2.3 we will make use of clustering redshifts to validate the robustness of our constraints to uncertainties in the combination  $b(z)p(z)$  in a model-independent way.

## 4.2 The CMB $\kappa$ map

We rotate the harmonic coefficients of the GMV CMB lensing convergence of the PR4 release into Equatorial coordinates, and truncate them to a maximum multipole  $\ell_{\max} = m_{\max} = 3N_{\text{side}} - 1$ , before transforming it into a HEALPix map of  $N_{\text{side}} = 512$  through an inverse spherical harmonic transform. We use the sky mask provided with the same release, rotate it in real space to Equatorial coordinates, apodize it with a C1 kernel (see section IIIB of [99]), and downgrade it to the same angular resolution.

Compared to previous releases, the PR4 lensing map has a non-trivial normalization that accounts, in a more optimal patch-dependent way, for the effect of anisotropic noise in the CMB maps, as well as its propagation in lensing reconstruction. However, simplifications in the response to the true sky signal and noise can introduce small biases to the normalization, in particular close to the mask boundaries. We compute the multiplicative correction to the estimator normalization (the so-called Monte Carlo (MC) correction) using the simulations associated with the PR4 release (which include realistic sky and noise modelling) measuring the ratio of the power spectrum of the input lensing map spectra to the cross-correlation between the input and the reconstructed lensing maps (see also discussion in [46, 100]). We compute the power spectra involved in this ratio on the Planck DR4 lensing mask with a binning of  $\delta\ell = 5$  and interpolate to the relevant effective multipoles of our chosen binning scheme. Such MC correction is of the order of  $\sim 5\%$  on small scales and only becomes significant on the largest scales  $\ell \lesssim 10$ , which have a minor statistical weight in our analysis. A mismatch in the normalization might also affect the accuracy of the subtraction of the mean-field anisotropy induced by the mask. This mean field becomes important in the CMB lensing autospectrum analysis, and dominates on the largest scales. However, it only contributes to the covariance of the spectrum in the cross-correlations with external tracers and does not introduce any bias. The use of 600 end-to-end Monte-Carlo simulations released with the NPIPE maps to estimate the mean field of the PR4 lensing estimator allowed us to reduce the uncertainty connected to this effect compared to the PR3 release. As such we do not propagate further uncertainties connected to the mean field.

## 4.3 Angular power spectra and covariances

To calculate all auto- and cross-power spectra we make use of the pseudo- $C_\ell$  approach as implemented in NaMaster<sup>9</sup> [75]. The estimator is described in detail in [75, 101], and we only provide a brief description here. The pseudo- $C_\ell$  estimator starts by assuming that all observed

<sup>9</sup><https://github.com/LSSTDESC/NaMaster>.

maps  $\tilde{a}$  are simply given by a locally-weighted version of the true underlying signal maps  $a$

$$\tilde{a}(\hat{\mathbf{n}}) = w_a(\hat{\mathbf{n}}) a(\hat{\mathbf{n}}), \quad (4.7)$$

where the weights map  $w_a$  is commonly called the “mask”. In the simplest case, we may consider a binary mask, where  $w_a = 1$  for observed pixels, and  $w_a = 0$  for unobserved ones. In practice, the mask can be used as an inverse-variance weight, which down-weights regions of the sky where the uncertainty on the true signal is high. This motivates us using the selection function itself as the mask for the quasar overdensity map, since it tracks the observed number density of the sample which, in the Poisson limit, corresponds to the inverse local variance of  $\delta_g$ .

A product in real space (eq. (4.7)) corresponds to a convolution in harmonic space, and the harmonic coefficients of  $\tilde{a}$  and  $a$  are related to one another via

$$\tilde{a}_{\ell m} = \sum_{\ell' m'} W_{\ell m, \ell' m'}^a a_{\ell' m'}, \quad W_{\ell m, \ell' m'}^a \equiv \int d\hat{\mathbf{n}} w_a(\hat{\mathbf{n}}) Y_{\ell m}^*(\hat{\mathbf{n}}) Y_{\ell' m'}(\hat{\mathbf{n}}). \quad (4.8)$$

The pseudo- $C_\ell$  of two observed maps  $(a, b)$  is defined as

$$\tilde{C}_\ell^{ab} \equiv \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} \tilde{a}_{\ell m} \tilde{b}_{\ell m}^*. \quad (4.9)$$

It is then straightforward to show that the ensemble average of  $\tilde{C}_\ell^{ab}$  is related to the true underlying power spectrum via

$$\langle \tilde{C}_\ell^{ab} \rangle = \sum_{\ell'} M_{\ell\ell'}^{ab} C_{\ell'}^{ab} + \tilde{N}_\ell^{ab}, \quad (4.10)$$

where  $\tilde{N}_\ell^{ab}$  is the pseudo- $C_\ell$  of the noise components of both fields.  $M_{\ell\ell'}^{ab}$  is the so-called *mode-coupling matrix*, and encodes the statistical coupling between different harmonic modes caused by the presence of a mask.  $M_{\ell\ell'}^{ab}$  depends solely on the pseudo- $C_\ell$  of the masks of both fields, and can be calculated analytically through an  $\mathcal{O}(\ell_{\max}^3)$  operation.

An unbiased estimator for the true power spectrum can thus be constructed by estimating and subtracting the noise bias from  $\tilde{C}_\ell^{ab}$ , and multiplying the result by the inverse of the mode-coupling matrix. To improve the condition of  $M_{\ell\ell'}^{ab}$  before inversion, and to reduce the size of the resulting data vector, it is often common to bin the pseudo- $C_\ell$  into finite bands of  $\ell$ , commonly called *bandpowers*. The estimated bandpowers  $\hat{C}_q^{ab}$  (where  $q$  labels each bandpower), can then be related to the true underlying power spectrum via

$$\hat{C}_q^{ab} = \sum_{\ell} \mathcal{F}_{q\ell}^{ab} C_\ell^{ab}, \quad (4.11)$$

where the *bandpower window functions*  $\mathcal{F}_{q\ell}^{ab}$  encode the effects of mode-coupling, binning, and inversion of the binned mode-coupling matrix. We bin all power spectra used here into bandpowers of width  $\Delta\ell=30$  between  $\ell = 2$  and  $\ell = 3N_{\text{side}} - 1$ .

Out of all power spectra used in our analysis, only the quasar auto-correlations have a non-zero noise bias due to shot noise. Assuming Poisson statistics, this bias can be estimated analytically as

$$\tilde{N}_\ell^{gg} = \frac{\langle w \rangle_P}{\bar{N}_\Omega}, \quad (4.12)$$

where  $\langle w \rangle_p$  is the sky average of the selection function, and  $\bar{N}_\Omega \equiv \bar{N}/\Omega_{\text{pix}}$  is the angular number density (in units of  $\text{srad}^{-1}$ ), with the mean number of objects per pixel given by eq. (4.2), and  $\Omega_{\text{pix}}$  the solid angle of each HEALPix pixel. In section 5.2.2 we will explore the impact of residual non-Poissonian shot noise in the data.

All power spectra are corrected for the effects of the pixel window function as described in [75], although we find the impact of this correction to be negligible on the scales used here. Besides this, as described in the previous section, the quasar- $\kappa$  cross-correlation is corrected for the CMB lensing transfer function. Finally, as discussed in section 4.1.1, we make use of linear deprojection as implemented in NaMaster, ignoring the negligible deprojection bias it incurs in our case.

We calculate the covariance matrix of the estimated power spectra analytically, following the methodology described in [102]. The method assumes that all fields involved are Gaussian, which is sufficiently accurate in our case, since all fields under study are largely noise-dominated. The approach described in [102] has been thoroughly validated, and accurately accounts for the effects of mode-coupling caused by the mask. To achieve this, the method only assumes that the true input power spectra are smooth functions of  $\ell$  that do not vary significantly compared to the width of the mode-coupling kernel caused by the masks involved (the so-called *Narrow-Kernel Approximation*, NKA). Given the broad sky coverage of the probes we study, this approximation is only mildly inaccurate on the largest scales, which we discard in our analyses in any case.

#### 4.4 Likelihood and priors

To infer free parameters of our model  $\vec{\theta}$ , given our data  $\mathbf{d}$ , we make use of a Gaussian (multivariate normal) likelihood:

$$-2 \log p(\mathbf{d}|\vec{\theta}) = (\mathbf{d} - \mathbf{t}(\vec{\theta}))^T \mathbf{C}^{-1} (\mathbf{d} - \mathbf{t}(\vec{\theta})) + K, \quad (4.13)$$

where  $\mathbf{t}(\vec{\theta})$  is the theoretical prediction for  $\mathbf{d}$ ,  $\mathbf{C}$  is the covariance matrix of  $\mathbf{d}$ , and  $K$  is an arbitrary normalisation constant that does not depend on  $\vec{\theta}$ . The *posterior* distribution  $p(\vec{\theta}|\mathbf{d})$  is calculated from the likelihood above using Bayes' theorem

$$\log p(\vec{\theta}|\mathbf{d}) = \log p(\mathbf{d}|\vec{\theta}) + \log p(\vec{\theta}) + K', \quad (4.14)$$

where  $p(\vec{\theta})$  is the *prior*.

In our case,  $\mathbf{d}$  is a collection of power spectra, comprising the quasar auto-correlations of the redshift bins under study, and their cross-correlations with the CMB lensing convergence.  $\mathbf{C}$  is the covariance matrix of these measurement, calculated as described in the previous section, and the model used to calculate the theoretical prediction  $\mathbf{t}$  was presented in section 3. We only include bandpowers in  $\mathbf{d}$  satisfying the following scale cuts:

- We discard the first bandpower ( $\ell < 32$ ) in all power spectra, where the impact of systematic contamination was found to be most relevant (see section 5.2.4).
- We impose a high- $\ell$  cut at  $\ell_{\text{max}} = k_{\text{max}}\chi(\bar{z})$ , where  $\bar{z}$  is the mean redshift of the bin under study, and  $k_{\text{max}}$  is the 3D wavenumber parametrizing our small-scale cut. In our

Parameter	Prior	Parameter	Prior
$\sigma_8$	$U(0.5, 1.2)$	$b_g^i$	$U(0.1, 3.0)$
$\Omega_m$	$U(0.05, 0.7) + \text{BAO}$	$A_{\text{SN}}^i(^*)$	$N(1, 0.1^2)$
$h$	$U(0.4, 1.0) + \text{BAO}$	$\Omega_m h^3(^{\dagger})$	$N(0.09633, 0.00029^2)$

(\*) The shot noise amplitudes  $A_{\text{SN}}^i$  are fixed to their Poisson level in our fiducial analysis, and we vary them with a 10% prior to test the dependence of our results on this assumption.

( $\dagger$ ) In section 5.2.7 we explore the impact of imposing a prior on the horizon scale at recombination from the positions of the acoustic peaks in the CMB power spectrum. Effectively, this puts a tight constraint on the combination  $\Omega_m h^3$  as shown on the table.

**Table 1.** Free parameters and priors used in our model. We use BAO data from BOSS and eBOSS [103, 104] to put a joint prior on  $\Omega_m$  and  $h$ .

fiducial results we choose  $k_{\text{max}} = 0.15 \text{ Mpc}^{-1}$ , corresponding to  $\ell_{\text{max}} = (497, 827)$  in the two redshift bins under study. The main motivation for this cut is ensuring the validity of the linear bias model used in our theoretical predictions. We will study the impact of this choice in section 5.2.1).

With these scale cuts, we are left with 16 and 27 bandpowers for power spectra involving the Low- $z$  and High- $z$  quasar redshift bins, respectively, and thus the fiducial size of the data vector combining all auto- and cross-correlations is  $N_d = 86$ . It is worth noting that the Gaussian likelihood of eq. (4.13), with a parameter-independent covariance, has been found to be an excellent approximation to the true likelihood of power spectra on scales  $\ell \gtrsim 10$  [105, 106].

Within the  $\Lambda\text{CDM}$  model, the main cosmological parameters that affect the power spectra considered here are the total non-relativistic matter density parameter  $\Omega_m$ , the amplitude of linear matter fluctuations  $\sigma_8$ , and the expansion rate  $h$ , and we will vary all three of them in our likelihood. We will fix the baryon density parameter, and the scalar spectral index to the best-fit values found by *Planck*:  $\Omega_b h^2 = 0.0224$ ,  $n_s = 0.9665$  [32]. Our data is not able to fully break the degeneracy between  $\Omega_m$  and  $h$  and hence we will impose a prior on both background parameters using low-redshift baryon acoustic oscillations (BAO) data. In particular, we include three measurements of the angular and radial BAO scales from Luminous Red Galaxies (LRGs) at redshifts  $z = 0.38, 0.51, 0.698$  obtained by the BOSS and eBOSS surveys in the DR12 and DR16 of the Sloan Digital Sky Survey [103, 104]. In section 5.2.7 we will explore the sensitivity of our final constraints to this choice of prior.

Besides the cosmological parameters, we consider the two linear bias parameters  $\{b_g^1, b_g^2\}$  with flat priors in the range  $b_g^i \in [0.1, 3.0]$ . We assume that the quasar bias evolves as in eq. (4.6) in our fiducial analysis, and explore the impact of this assumption on our final results in section 5.2.2. We will also quantify the impact of non-Poissonian shot noise on the final parameter constraints by freeing-up the amplitude of the shot-noise contribution to the quasar auto-correlations. This will add two new free parameters ( $A_{\text{SN}}^1, A_{\text{SN}}^2$ ), one for each redshift bin, with a 10% Gaussian prior centered around 1. Since both parameters enter linearly in the theoretical prediction, we will marginalise over them analytically. Table 1 lists all the parameters of the model, and their priors.

As described in section 4.1.3 we marginalise over the uncertainties in the calibrated redshift distribution. We do so following the analytical approach of [80], which results in a modification of the power spectrum covariance matrix, without adding any new free parameters to the model.

All theoretical calculations were carried out using the Core Cosmology Library<sup>10</sup> (CCL [107]). The linear matter power spectrum was computed with CLASS [108], and we used the HALOFit prescription of [109] to connect it with the non-linear power spectrum. We sample the posterior distribution using a the Metropolis-Hastings Monte-Carlo Markov Chain (MCMC) algorithm as implemented in Cobaya [110]. We determine chains to be converged when the Gelman-Rubin statistic reaches  $R < 0.01$ . We will present parameter constraints by quoting the peak of the marginalised 1D distribution, together with its 68% asymmetric confidence-level interval.

## 5 Results

### 5.1 Fiducial cosmological constraints

In our fiducial setup, we consider the  $G < 20.5$  Quia sample split into two redshift bins (see figure 4), with DIR-calibrated redshift distributions, fixed magnification bias slope parameters as in eq. (4.3), assuming a linear bias evolving as in eq. (4.6), and purely Poisson shot noise. We use only scales with  $k < k_{\max} = 0.15 \text{ Mpc}^{-1}$ , and impose a BAO prior on  $(\Omega_m, h)$ .

Figure 5 shows the angular power spectra measured in both bins ( $C_\ell^{gg}$  in black,  $C_\ell^{g\kappa}$  in red), together with the best-fit theoretical prediction (blue lines). The shaded bands show the angular scales discarded in our analysis. The bottom panels of both figures show the residuals with respect to the best-fit prediction normalised by the  $1\sigma$  uncertainties. We find that the best-fit model provides a reasonable fit to the data, with a  $\chi^2 = 96.9$ . With  $N_d = 86$  data points and  $N_p = 5$  free parameters, the associated probability to exceed (PTE) is  $p = 0.11$ , which corresponds to an adequate fit.

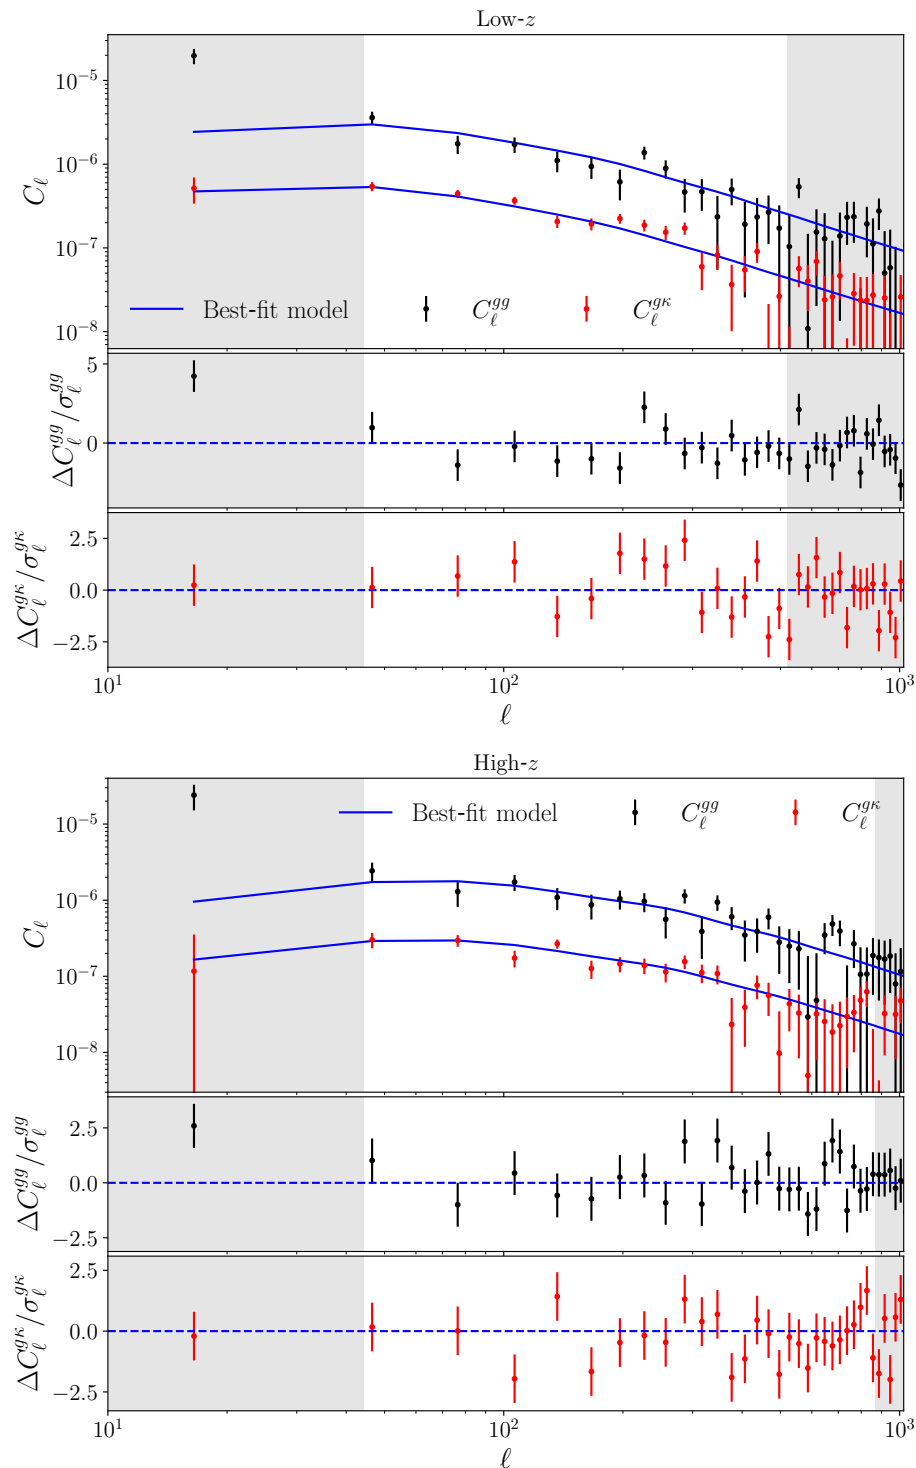
The cross-correlation between Quia and the *Planck* CMB lensing map is detected, within the scales used in this analysis, with a signal-to-noise ratio of 16.9 and 13.6 in the Low- $z$  and High- $z$  bins respectively, for a total detection significance of  $21.7\sigma$ . This is actually stronger than the significance of the clustering auto-correlation signal itself ( $9\sigma$  and  $8.4\sigma$  respectively), due to the impact of shot noise. Hence the cross-correlation carries a larger statistical weight in our final cosmological constraints.

The constraints on the cosmological parameters  $(\sigma_8, \Omega_m, h, S_8)$  found with this fiducial setup are shown in blue in figure 6. As explained in section 3.4, we are able to break the degeneracy between  $\sigma_8$  and  $\Omega_m$ , measuring them independently:

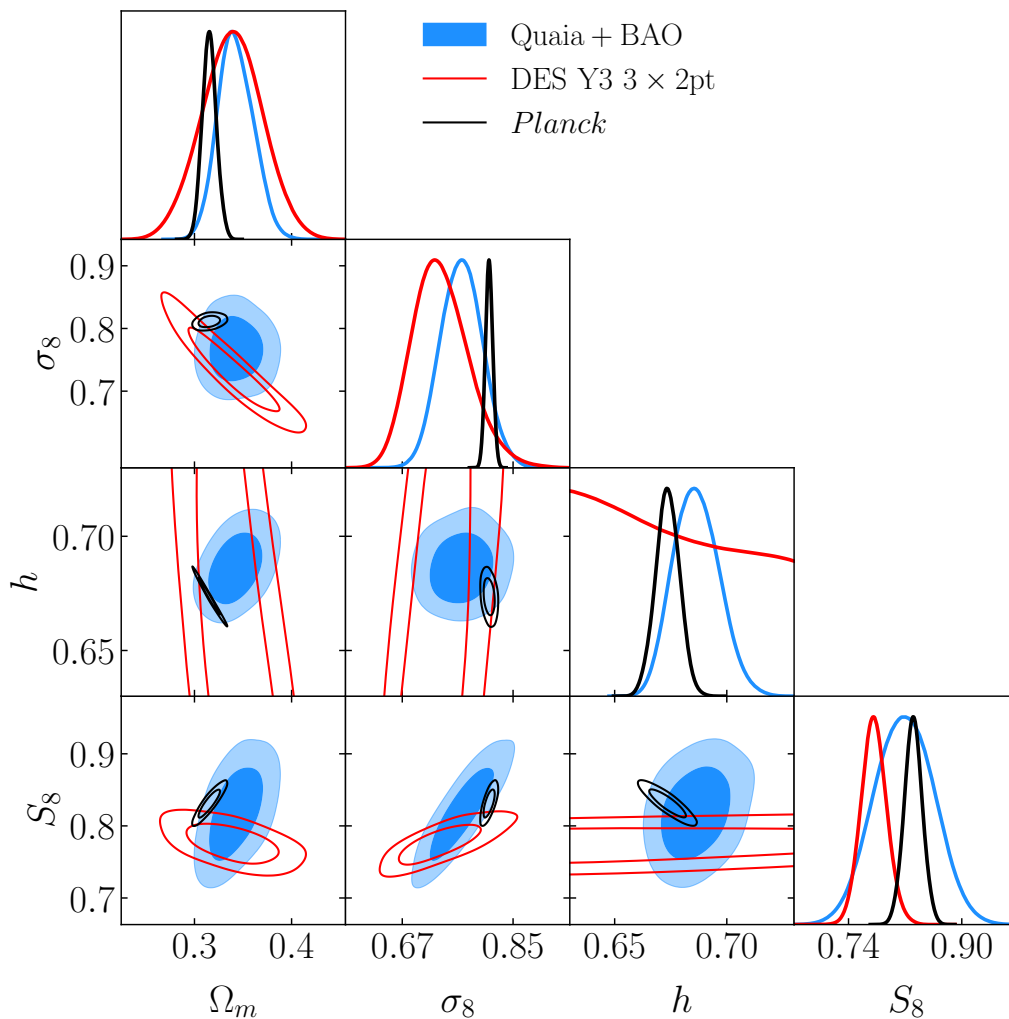
$$\sigma_8 = 0.766 \pm 0.034, \quad \Omega_m = 0.343_{-0.019}^{+0.017}. \quad (5.1)$$

The figure also shows, in red and black, the constraints on the same parameters found by the DES Y3  $3 \times 2$ -point analysis [3], combining cosmic shear and galaxy clustering from the MagLim sample, and by *Planck* [32], respectively. Our results are in reasonable agreement

<sup>10</sup><https://github.com/LSSTDESC/CCL>.



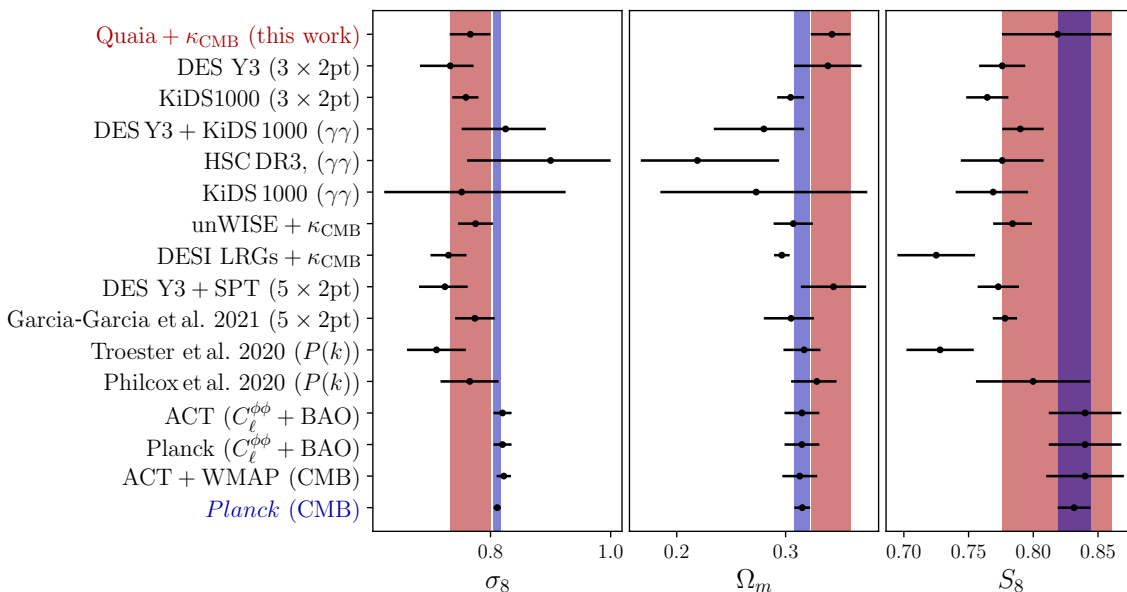
**Figure 5.** Angular auto-power spectra of the quasar overdensity field (black) and cross-correlation with the *Planck* CMB lensing convergence map (red) in the two Quiaia redshift bins (top and bottom panels for the Low- $z$  and High- $z$  bins respectively). The blue lines show the best-fit theoretical prediction found in our fiducial configuration. The bottom sub-plots in each panel show the residuals with respect to the best-fit prediction normalised by the  $1\sigma$  uncertainties. The central white band shows the angular scales used in our analysis for each redshift bin.



**Figure 6.** Cosmological constraints from our fiducial analysis of Quiaia with *Planck* CMB lensing, including a BAO prior from SDSS (filled blue contours), compared with those found by *Planck* (black), and the Y3 DES  $3 \times 2$ pt analysis (red).

with the *Planck* measurements, with only a mild tension at the  $1.3\sigma$  and  $1.4\sigma$ -level for  $\sigma_8$  and  $\Omega_m$  respectively. The constraints are also fully compatible with those found by DES Y3.

It is worth placing our constraints in the context of measurements obtained by other experiments. Figure 7 shows our constraints on  $\sigma_8$ ,  $\Omega_m$ , and  $S_8$ , together with those found in some of the latest analyses of galaxy clustering, cosmic shear, and CMB data. Our constraints are in broad agreement with those found by most of these experiments. We find that we are able to place a competitive constraint on  $\sigma_8$ , unmatched by cosmic-shear experiments due to its degeneracy with  $\Omega_m$ , and comparable to that found by  $3 \times 2$ pt,  $5 \times 2$ pt, and lensing tomography analyses, although with a significantly smaller sample size in the case of Quiaia. Our constraint on  $S_8$ , in turn, is less stringent than weak lensing and  $3 \times 2$ pt analyses, but comparable with those found by full-shape power spectrum analyses in spectroscopic galaxy surveys [111, 112].



**Figure 7.** Visual representation of the 68% confidence level constraints on  $\sigma_8$ ,  $\Omega_m$ ,  $S_8$  found in this analysis (salmon band), and found in other datasets. The points shown correspond to cosmic-shear analyses (marked  $\gamma\gamma$  [66–68]),  $3 \times 2$ pt and  $5 \times 2$ pt data [2, 3, 25, 26], lensing tomography [19, 21], full-shape power spectrum analyses [111, 112], the CMB lensing power spectrum + BAO data [34, 48], and CMB data [32, 33]. The *Planck* constraints are shown as a blue vertical band.

As highlighted in section 3.4, this is one of the highest-redshift constraints on the amplitude of matter fluctuations. That it agrees reasonably with probes at both higher and lower redshifts is therefore a non-trivial result that confirms the ability of the simple  $\Lambda$ CDM model to describe the large-scale properties of the Universe across most of its history.

## 5.2 Robustness tests

In this section we study the robustness of these results against the various analysis choices made, and potential sources of systematic uncertainty. Table 2 lists the final parameter constraints found for the different alternatives explored here, together with our fiducial measurements. The constraints on  $(\sigma_8, \Omega_m, S_8)$  from these cases are also shown in figure 8. In all these cases we find a reasonably good fit to the data with  $\chi^2$  PTE values ranging between 0.04 and 0.5.

### 5.2.1 Internal consistency

We begin by exploring whether we obtain consistent results from different sub-sections of our data:

- We explore the impact of our *small-scale cut* ( $k_{\max} = 0.15 \text{ Mpc}^{-1}$ ) on the final constraints by repeating our analysis imposing both more and less conservative cuts. Rows 2 and 3 in table 2 and figure 8 show the constraints found for  $k_{\max} = 0.1 \text{ Mpc}^{-1}$  and  $k_{\max} = 0.2 \text{ Mpc}^{-1}$ . In both cases we recover constraints that are compatible with our fiducial measurements within  $\sim 0.5\sigma$ , with a  $\sim 6\%$  increase and reduction in the final

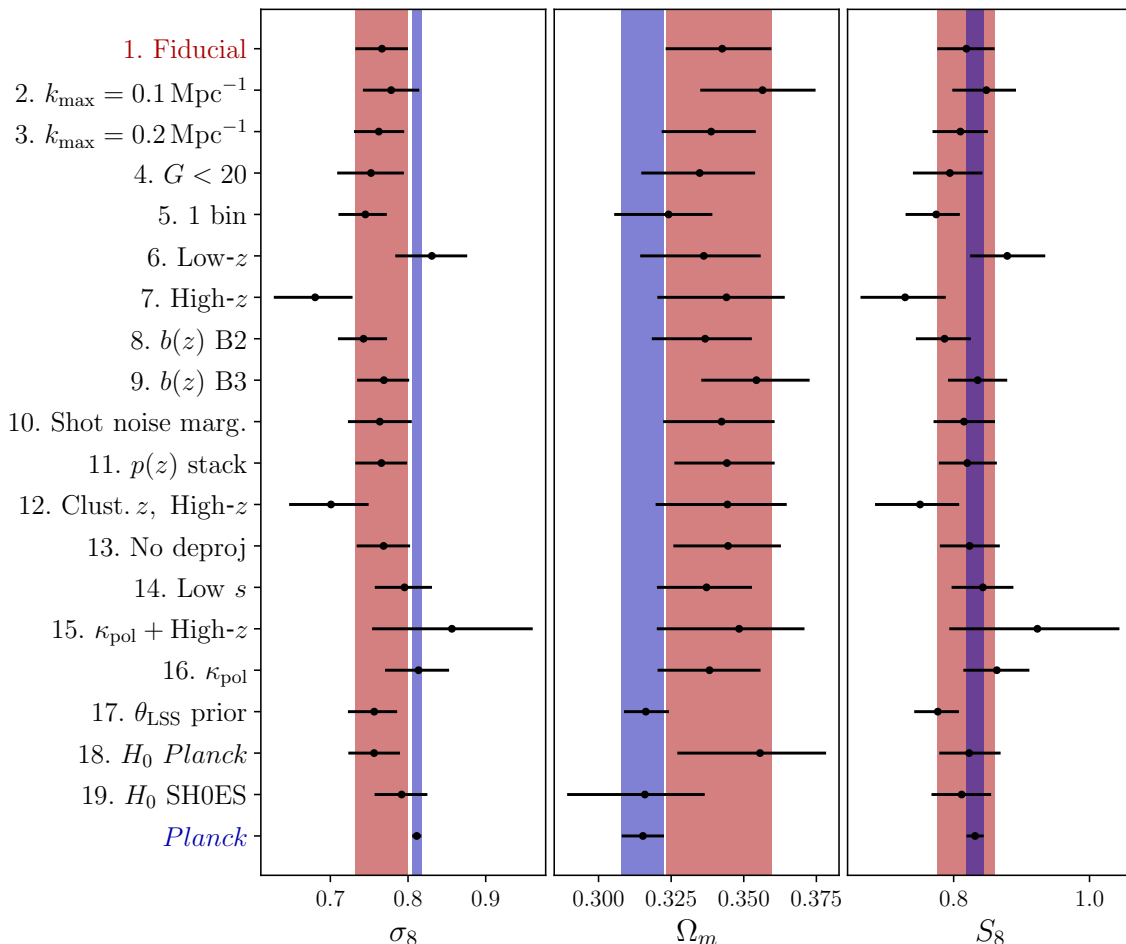


Case	$\sigma_8$	$\Omega_m$	$S_8$
1. <b>Fiducial</b>	$0.766 \pm 0.034$	$0.343^{+0.017}_{-0.019}$	$0.819 \pm 0.042$
2. $k_{\max} = 0.1 \text{ Mpc}^{-1}$	$0.778 \pm 0.036$	$0.356^{+0.018}_{-0.021}$	$0.848^{+0.044}_{-0.050}$
3. $k_{\max} = 0.2 \text{ Mpc}^{-1}$	$0.762 \pm 0.032$	$0.339 \pm 0.017$	$0.810 \pm 0.041$
4. $G < 20$	$0.752 \pm 0.043$	$0.335 \pm 0.020$	$0.795^{+0.048}_{-0.054}$
5. 1 bin	$0.745^{+0.028}_{-0.035}$	$0.324^{+0.015}_{-0.019}$	$0.774^{+0.035}_{-0.045}$
6. Low- $z$	$0.831 \pm 0.046$	$0.336^{+0.020}_{-0.022}$	$0.879 \pm 0.055$
7. High- $z$	$0.680 \pm 0.051$	$0.344^{+0.020}_{-0.024}$	$0.729 \pm 0.063$
8. $b(z)$ B2	$0.743 \pm 0.032$	$0.337^{+0.016}_{-0.018}$	$0.787 \pm 0.040$
9. $b(z)$ B3	$0.769 \pm 0.034$	$0.354 \pm 0.019$	$0.835 \pm 0.043$
10. Shot noise marg.	$0.764 \pm 0.041$	$0.342 \pm 0.019$	$0.815 \pm 0.045$
11. $p(z)$ stack	$0.766 \pm 0.034$	$0.344 \pm 0.017$	$0.820 \pm 0.042$
12. Clust. $z$ , High- $z$	$0.701 \pm 0.051$	$0.344^{+0.020}_{-0.025}$	$0.751^{+0.058}_{-0.066}$
13. No deproj	$0.769 \pm 0.035$	$0.345 \pm 0.018$	$0.823 \pm 0.045$
14. Low $s$	$0.795 \pm 0.036$	$0.337 \pm 0.016$	$0.843 \pm 0.045$
15. $\kappa_{\text{pol}} + \text{High-}z$	$0.86 \pm 0.11$	$0.348^{+0.023}_{-0.028}$	$0.92 \pm 0.13$
16. $\kappa_{\text{pol}}$	$0.814 \pm 0.042$	$0.338 \pm 0.018$	$0.864 \pm 0.050$
17. $\theta_{\text{LSS}}$ prior	$0.756^{+0.030}_{-0.034}$	$0.3163 \pm 0.0079$	$0.777 \pm 0.033$
18. $H_0$ <i>Planck</i>	$0.756 \pm 0.033$	$0.356^{+0.023}_{-0.028}$	$0.823 \pm 0.045$
19. $H_0$ SH0ES	$0.792 \pm 0.036$	$0.316^{+0.021}_{-0.027}$	$0.812 \pm 0.046$

**Table 2.** Constraints (marginalised mean and 68% confidence level interval) on  $(\sigma_8, \Omega_m, S_8)$  for the different cases studied here. Our fiducial result, described in section 5.1, is shown in the first row. All other rows show the results found in the various robustness tests described in section 5.2.

uncertainties for the more and less conservative cuts respectively. Our results are thus robust to this choice. This also reinforces the validity of the simple linear bias model used in our analysis, since the presence of large non-linear corrections would lead to significant changes in the final constraints when adding or removing smaller scales.

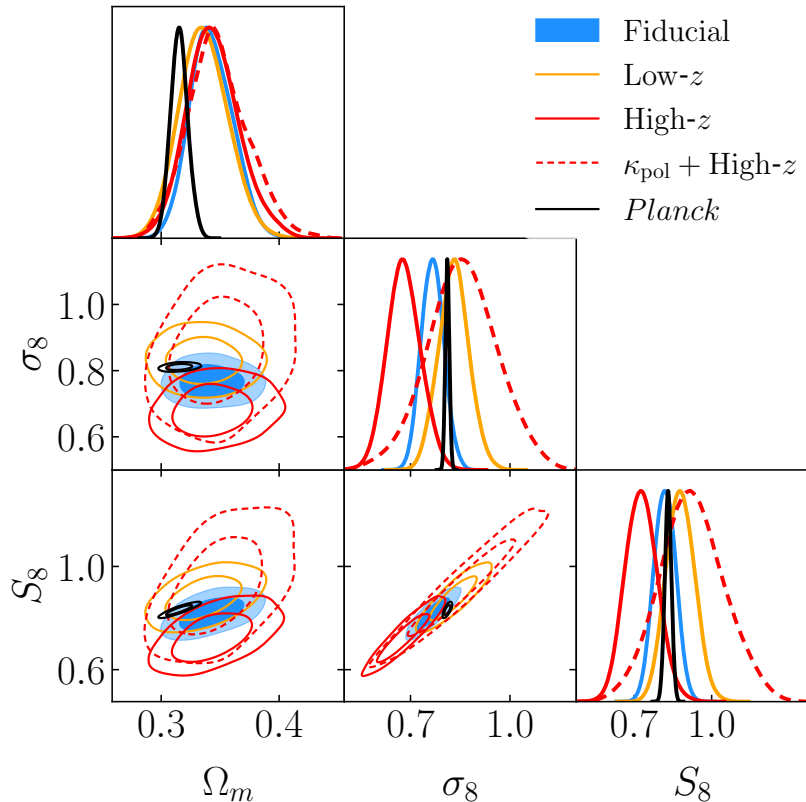
- We repeat the analysis for a *brighter quasar sample*, cutting at  $G < 20$ . This brighter sample is potentially less sensitive to sky contamination and photo- $z$  mis-estimation than the fiducial catalog used in our analysis. For this sample, the magnification bias slope is given by eq. (4.4). As shown in row 4 of table 2 and figure 8, the resulting parameter constraints are fully compatible with the fiducial case, with a  $\sim 30\%$  increase in the final uncertainties, caused by the higher shot noise of this sample.
- We obtain cosmological constraints from all Quia sources combined into a *single redshift bin* (with the redshift distribution shown in black in figure 4). This test allows us to check for the impact of any potential mis-estimation of the intermediate-redshift tails of the redshift distributions in the 2-bin case, as well as increased sky contamination



**Figure 8.** Visual representation of the results shown in table 2. Together with the results found in the different robustness described in section 5.2, we show the constraints found in our fiducial analysis (first row, and salmon band), and those found by *Planck* (last row and blue band).

caused by the additional cut in redshift. The results, shown in row 5 of table 2 and figure 8, are largely compatible with our fiducial constraints, with a shift of less than  $1\sigma$  in both  $\sigma_8$  and  $\Omega_m$ , and barely any change in the final uncertainties.

- Finally, we repeat our analysis for each of the *two redshift bins separately*. The constraints are shown in rows 6 and 7 of table 2 and figure 8, and the corresponding 2D contours are also shown separately in figure 9. The cosmological parameters inferred from each redshift bin are compatible with one another, although the High- $z$  bin favours a value of  $\sigma_8$  that is  $2.2\sigma$  lower than the Low- $z$  bin,  $1.5\sigma$  lower than our fiducial constraints, and  $2.5\sigma$  lower than *Planck*. We thus find that the slightly lower value of  $\sigma_8$  found in our analysis in comparison with *Planck* is mostly driven by the higher redshift bin. We will therefore scrutinise the results obtained from this bin further in the following sections.



**Figure 9.** Constraints on  $\Omega_m$ ,  $\sigma_8$ , and  $S_8$  found in our fiducial analysis (blue contours), in comparison with those found using only each of the two Quiaa redshift bins independently (solid orange and red contours for the Low- $z$  and High- $z$  bins respectively). The *Planck* constraints are shown in black. The higher redshift bin recovers a value of  $\sigma_8$  that is smaller than that found by *Planck* at the  $2.5\sigma$  level, and is hence the main driver of the slightly lower value of  $\sigma_8$  found in our fiducial analysis. The dashed red contours show the constraints found with the High- $z$  when cross-correlating with the polarization-only CMB lensing map (which is more reliable against extragalactic foreground contamination in the CMB map), as presented in section 5.2.6. The  $\sim 1.5\sigma$  upwards shift in  $\sigma_8$  observed is a hint that the low  $\sigma_8$  found with the fiducial CMB  $\kappa$  map may be caused by extragalactic foreground contamination in that map.

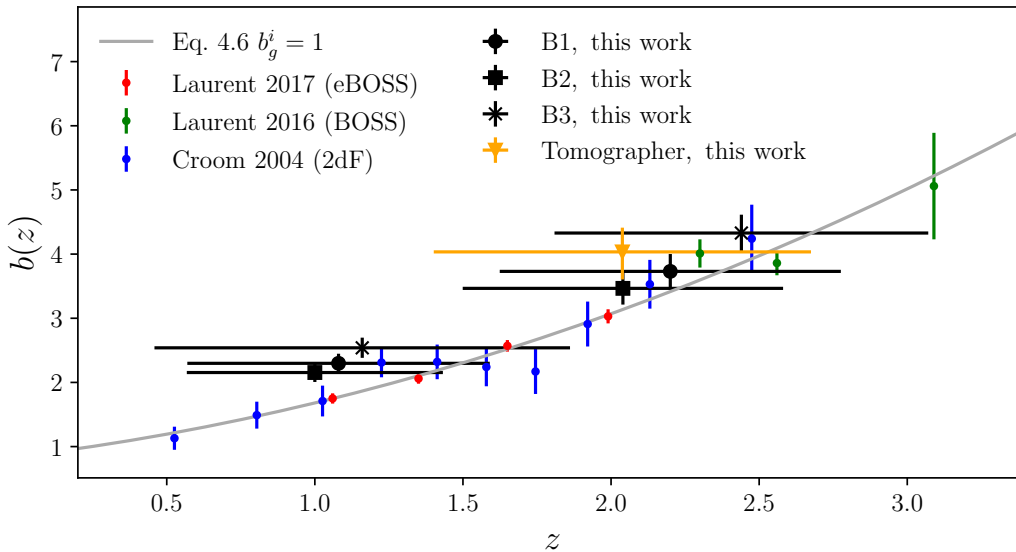
### 5.2.2 Bias model

Our fiducial constraints assume the specific redshift dependence of the quasar bias in eq. (4.6). For concreteness, we will label this **Model B1** in this section. To quantify the dependence of our results on the assumed evolution of the quasar bias, we repeat our analysis for two other models:

- **Model B2:** the model of [93], which is significantly steeper, and likely ruled out by current data:

$$b(z) = b_g^i \left[ 1 + \left( \frac{1+z}{2.5} \right)^5 \right], \quad (5.2)$$

- **Model B3:** a constant bias model ( $b(z) = b_g^i$ ), also ruled out by current data if applied at all redshifts.



**Figure 10.** Measurements of the bias  $b(z)$  for various quasar samples: the 2dF QSO Redshift Survey [83] (blue), BOSS [89] (green), and eBOSS [38] (red). The grey line shows the fiducial bias evolution model used in our analysis, with the amplitude parameter  $b_g$  fixed to 1, to reproduce the best-fit of [38] for eBOSS. The bias values measured in this analysis for the Quiaia sample are shown in black for the three different bias evolution models considered (models B1, B2, and B3, described in section 5.2.2). The value found when using Tomographer to constrain the redshift dependence of the product  $b(z)p(z)$  is shown in orange. Our constraints are in rough agreement with previous measurements. The systematically higher value of the quasar bias found for Quiaia is consistent with it being slightly brighter than the spectroscopic samples shown in the figure.

The constraints found with these two alternative models are shown in rows 8 and 9 of table 2 and figure 8. We observe that models with a steeper bias evolution lead to lower values of both  $\sigma_8$ , and  $\Omega_m$ . Nevertheless, for both models, B2 and B3, which encompass the degree of evolution exhibited by existing measurements, the parameter shifts are mild (less than  $1\sigma$ ). Hence, our results do not depend strongly on this choice.

It is useful to place our constraints on quasar bias in the context of previous measurements. For the three bias evolution models studied, we calculate the effective bias in each of the two redshift bins as

$$b_{\text{eff}} = \frac{\int dz p(z) b(z)}{\int dz p(z)}, \quad (5.3)$$

and assign to it an effective redshift given by the mean  $z$  in each bin weighted by the product  $b(z)p(z)$ . The results are shown in figure 10 for all three models (black symbols), together with the measurements of [38, 83, 89] (blue, green, and red points), and our fiducial model (eq. (4.6), with  $b_g^i = 1$ ). Reassuringly, our measurements are in rough agreement with other quasar samples. In particular, we find no strong evidence of the lower quasar bias at high redshifts ( $z \gtrsim 2$ ) reported by previous analyses using CMB lensing cross-correlations [43, 46, 90, 113]. As discussed in section 4.1.2, the impact of magnification bias, an effect that the CMB lensing cross-correlation is particularly sensitive to, is likely negligible in our sample. Finally, it is worth noting that, although our bias constraints lie somewhat above most previous

measurements in both redshift bins, this is not entirely surprising, as our sample is different (brighter, in fact) from the spectroscopic samples targeted by previous measurements.

Our fiducial bias model is purely linear, and ignores contributions from higher-order terms. This assumption is likely valid given the relatively conservative scale cuts used [114–116], and we are further reassured by the robustness of our constraints to changes in these scale cuts (as shown in section 5.2.1). However, we have assumed that the stochastic term in eq. (3.2) is completely described by Poisson statistics, and can be subtracted analytically. Since a stochastic term would affect the quasar auto-correlation on all harmonic scales, we test this assumption by repeating our analysis varying the amplitude of the shot-noise contribution to  $C_\ell^{gg}$  with a 10% prior. This also allows us to test for any misestimation of the purely-Poisson noise level. The result is shown in row 10 of table 2 and figure 8. The mean value of all parameters remains virtually unchanged, the only effect being a  $\sim 10$ –20% increase in the final uncertainties. The data thus shows no preference for additional flat contribution to  $C_\ell^{gg}$ , and its inclusion causes no shift on the final parameter constraints.

### 5.2.3 Redshift uncertainties


To test the robustness of our constraints to a mis-calibration of the quasar redshift distributions, we repeat our analysis assuming the  $p(z)$  calculated via PDF stacking, as described in section 4.1.3. The results are shown in row 11 of table 2 and figure 8. The constraints are virtually unchanged, which shows that our results are insensitive to mild changes in the redshift distributions (such as the height of the bump<sup>11</sup> at  $z \sim 2.25$  shown in figure 4).

Another important systematic related to the  $p(z)$  is a potential mis-estimate of the inter-bin tails of both samples. Driven by photometric redshift outliers, they control the amplitude of the redshift distribution, and hence the amplitude of the quasar auto-correlation. Fortunately, the cross-correlation between both bins is particularly sensitive to the details of the redshift overlap between them [20, 117]. This cross-correlation is shown in figure 11, together with the theoretical prediction for our fiducial best-fit model. On scales  $30 \leq \ell \leq 800$ , the cross-correlation is detected at the  $3.2\sigma$  level, and our best-fit model provides a good fit to it (PTE = 0.51), without any additional parameters. This shows that the combination of best-fit bias parameters found in both bins, and the intermediate-redshift tails of their redshift distributions, are compatible with a complementary piece of the data vector that was not used in the likelihood.

As shown in section 5.2.1, the lower value of  $\sigma_8$  we recover compared to *Planck* is largely driven by the auto- and cross-correlation of the High- $z$  quasar redshift bin. Since both the shape of the redshift distribution, and the evolution of the linear bias may depart from the fiducial models used here, particularly at high redshifts. (e.g. due to photo- $z$  outliers or a steeper-than-expected  $b(z)$  relation), we perform one further check to test the robustness of the low  $\sigma_8$  value measured for that sample. We made use of the clustering redshifts technique, as implemented in *Tomographer*,<sup>12</sup> to measure the product  $b(z)p(z)$  up to an unknown normalization factor. *Tomographer* uses a spectroscopic reference sample from SDSS comprising

<sup>11</sup>The origin of this bump is discussed in more detail in [49], and is in fact robust to the choice of spectroscopic training sample.

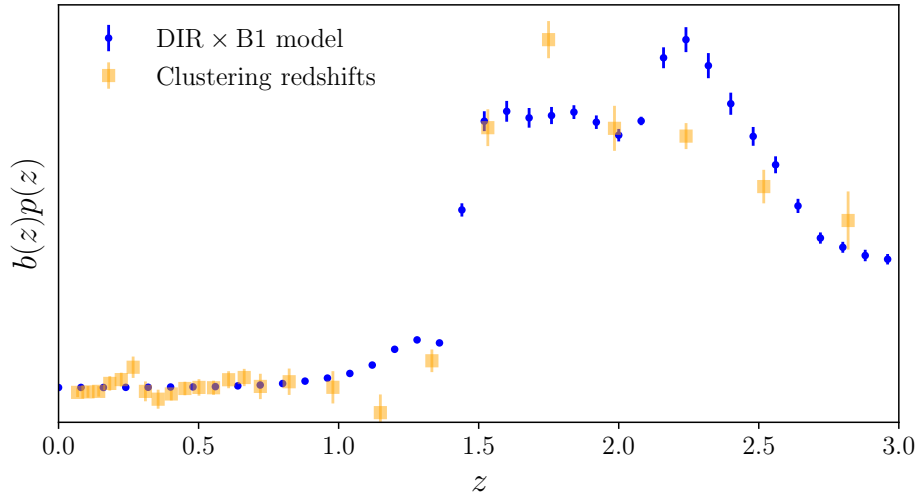
<sup>12</sup><http://tomographer.org>.



**Figure 11.** Cross-correlation between the two Quiaia redshift bins (black points with error bars), and prediction from our best-fit model. This cross-correlation is sensitive to the overlap between the redshift distributions of both bins, shown in figure 4. Our best-fit model provides a good fit to this cross-correlation, which was not used in that analysis.

2 million sources at redshifts  $z \lesssim 3$ . Through its online interface, the code transforms an input catalog into a HEALPix map at resolution  $N_{\text{side}} = 2048$ , imposes a high-pass filter removing scales larger than  $\sim 2^\circ$  to mitigate the impact of sky systematics, and calculates the correlations between the map and the reference sample in thin redshift bins on scales  $r \in [2, 8]$  Mpc. In combination with the reference sample auto-correlations, this is then used to provide constraints (mean and standard deviation) of  $b(z)p(z)$  in each spectroscopic redshift bin.

We use **Tomographer** to estimate  $b(z)p(z)$  for the High- $z$  bin in our sample. The redshift coverage of the reference spectroscopic sample in **Tomographer** prevents a reliable determination of  $b(z)p(z)$  beyond  $z \sim 3$ , and thus, for this test, we impose an additional cut in the catalog at  $z_{\text{Quiaia}} < 3$ . As with the DIR redshift distribution, we marginalise over the uncertainties in  $b(z)p(z)$  obtained by **Tomographer** using the analytical marginalisation method of [80]. The cosmological constraints obtained from this test are shown in line 12 of table 2 and figure 8. Compared with the constraints found with the High- $z$  bin using the DIR redshift distribution and our fiducial bias evolution template, we observe only a small upwards shift in  $\sigma_8$  ( $\delta\sigma_8 \simeq 0.02$ , less than  $0.5\sigma$ ). Hence, the data still favours a lower value of  $\sigma_8$  compared to *Planck*, although at a lower significance of  $2.1\sigma$ . This slightly better agreement with *Planck* could be simply a statistical fluctuation, it could be caused by the slightly lower effective redshift of the sample, or could be evidence of a departure from the quasar bias evolution model assumed here. The effective bias obtained from our MCMC chains for this bin using clustering redshifts is shown in orange in figure 10, and agrees well with the values found for the three other bias evolution models explored. Figure 12 shows the values of  $b(z)p(z)$  recovered from **Tomographer** in orange, together with the DIR-calibrated  $p(z)$  multiplied by the B1 bias evolution model (eq. (4.6)). Both predictions for the product

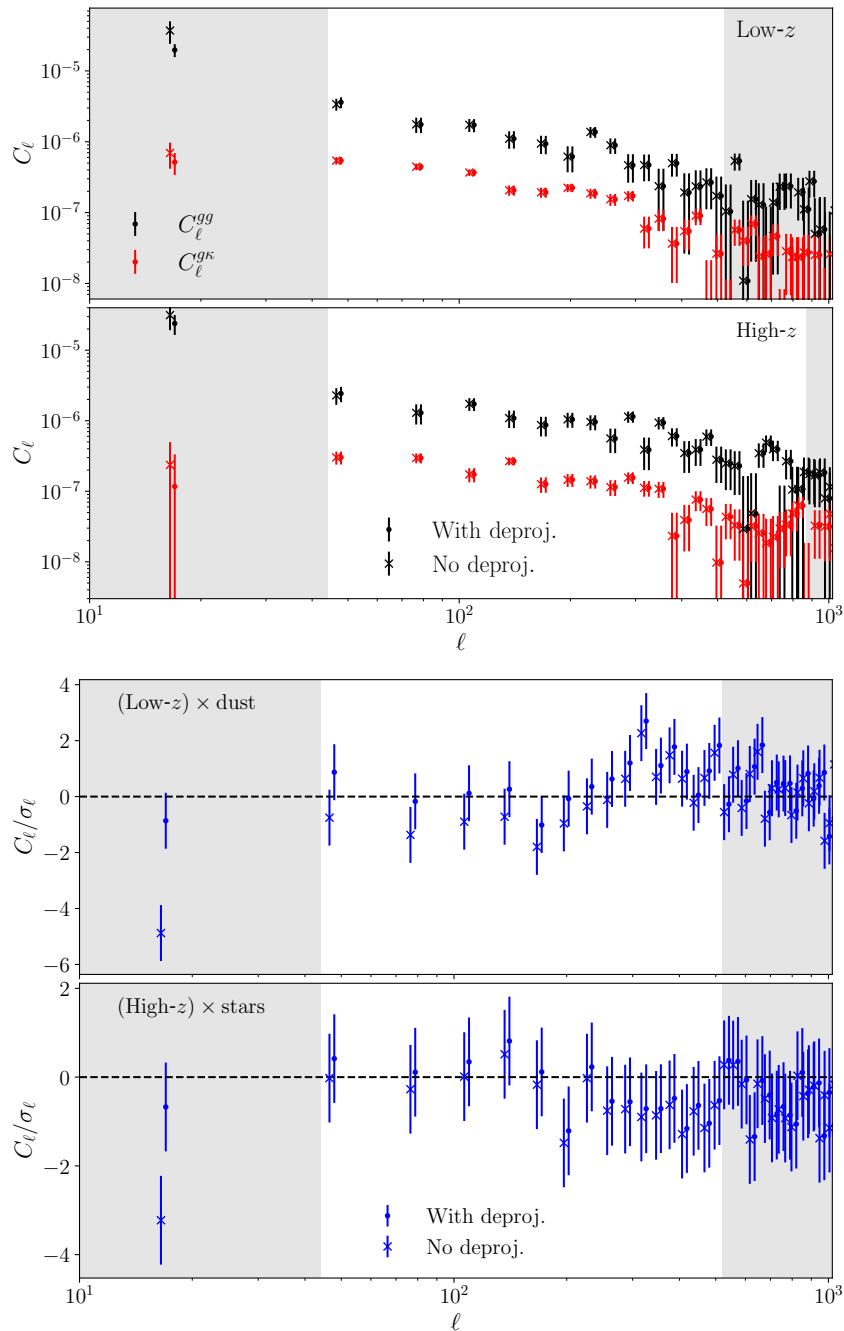


**Figure 12.** Product of the redshift distribution and linear bias of the High- $z$  Quiaia redshift bin. The blue points show the DIR estimate of the redshift distribution multiplied by our fiducial bias evolution model (eq. (4.6)). The orange points show the measurement of  $b(z)p(z)$  made using the clustering redshifts approach of Tomographer.

$b(z)p(z)$  are qualitatively similar in terms of the redshift support and high-redshift tails. The small difference observed in the final constraints might be caused by the enhancement in the bump at  $z \sim 2.25$  in the DIR distribution caused by the strong bias evolution at that redshift. Overall, however, we find that both predictions yield highly compatible constraints for the High- $z$  bin, and hence, when combined with the results presented in this section and section 5.2.2, it is unlikely for the lower value of  $\sigma_8$  found in this bin to be caused by mis-modelling of the sample’s redshift distribution or its bias evolution.

#### 5.2.4 Sky contamination

The presence of residual sky contaminants in the quasar overdensity maps could cause a bias in the final cosmological constraints. For instance, if purely additive, contamination would raise the amplitude of the quasar auto-correlation, leaving the cross-correlation intact. When fitting for bias and  $\sigma_8$ , this would then lead to an upwards shift in the quasar bias (to accommodate the higher auto-correlation) and, correspondingly, a downwards shift in  $\sigma_8$  (to compensate for the larger bias in the cross-correlation). To quantify the robustness of our results to the presence of residual contamination, we repeat the analysis skipping the final linear deprojection step when computing all power spectra. As a reminder, this step was meant to ensure that any small residual contamination from the known sky systematics was removed from the maps. The top panel of figure 13 shows the four angular power spectra used in our analysis with (points) and without (crosses) performing the final linear deprojection step. We see that linear deprojection has a significant effect on the first bandpower, which is discarded in our analysis, and less of an effect on smaller scales. The constraints found without linear deprojection are shown in row 13 of table 2 and figure 8. We observe only negligible shifts in the best-fit parameters ( $\lesssim 0.1\sigma$ ), and thus conclude that our results are likely robust against residual contamination from the known sky systematics.



**Figure 13.** *Top:* quasar-quasar (black) and quasar- $\kappa$  (red) angular power spectra for the two redshift bins used in this analysis (top and bottom subpanels). The measurements made using linear deprojection to remove residual sky contamination are shown as points with error bars, while the crosses show the results found without deprojection. The impact of residual contamination is mostly relevant for the first bandpower of the auto-correlation, which we discard in our analysis. *Bottom:* example cross-correlations between the quasar overdensity maps and the most relevant sky contaminants for each of the Quiaia redshift bins. Results are shown with and without linear deprojection (points and crosses respectively), divided by their statistical uncertainties to better highlight significant departures from zero correlation. Evidence of contamination is mostly present in the first bandpower, in agreement with the results of the top panel.



To further understand this result, we compute the cross-correlation between the quasar overdensity maps and the 4 contaminant maps described in section 4.1.1. The detection of a non-zero correlation would constitute evidence of residual contamination in the maps. We find that, without linear deprojection, the quasar overdensity maps exhibit large-scale residual contamination, particularly from dust extinction and stars. This contamination disappears after deprojection (unsurprisingly, since linear deprojection is precisely intended to cancel these correlations at the pixel level). An example of this is shown in the lower panel of figure 13 for the two Quaia bin-systematics combinations exhibiting the most contamination. Although linear deprojection modifies some of the cross-correlations on all scales, they are only significantly different from zero in the first bandpower without deprojection, in agreement with the results found above.

### 5.2.5 Magnification bias

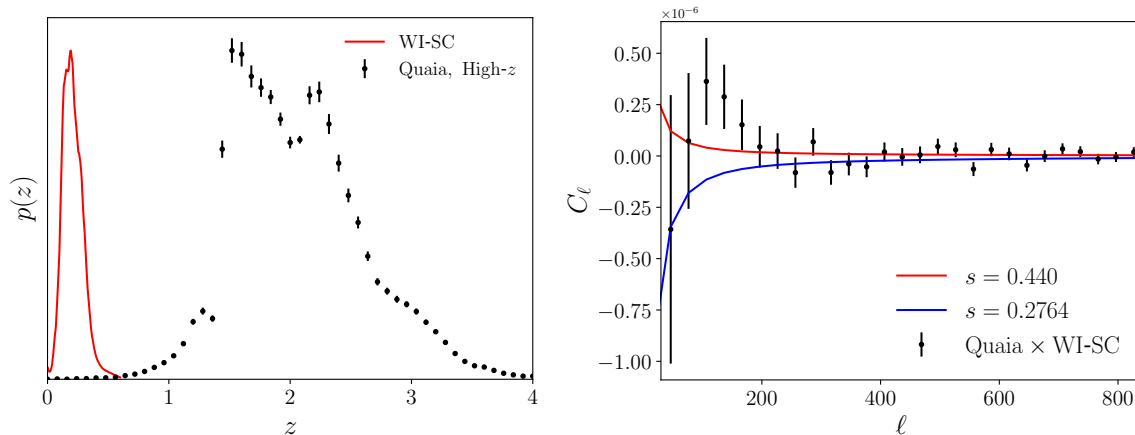
As discussed in section 4.1.2, the magnification bias slope of our fiducial quasar sample, measured with two different methods, is close to  $s \simeq 0.4$ , and hence our measurements are largely unaffected by magnification bias. To quantify the sensitivity of our final constraints to a potential mis-estimation of this quantity, we repeat our analysis fixing the number counts slope to the value found for the DESI quasar sample,  $s = 0.2764$  [46]. This value is in rough agreement with previous estimates for similar spectroscopic quasar samples (e.g. in BOSS and eBOSS). The result from this exercise, shown in row 14 of table 2 and figure 8, leads to a  $< 1\sigma$  upwards shift in  $\sigma_8$ , and almost no change in the  $\Omega_m$  constraint. It is worth noting that this value of  $s$  is  $\gtrsim 3\sigma$  away from our estimate for Quaia (see section 4.1.2), and is therefore not supported by the data. Nevertheless, the impact on our constraints of such a significant misestimate of  $s$  is small.

As an additional test to validate our estimate of  $s$  for the High- $z$  bin (which shows the largest tension with *Planck*), we measure the cross-correlation between this bin and a low-redshift sample of galaxies selected from the WISE-SuperCOSMOS catalog (WI-SC, [118]). In particular, we select 14.7 million galaxies with photometric redshifts  $z_p \in [0.1, 0.3]$ . The methods used to analyse this sample are described in detail in [9], and we will not repeat them here for brevity. As shown in the left panel of figure 14, both samples have negligible overlap in redshift, and thus, any detection of the cross-correlation between them would be largely driven by magnification bias with  $s \neq 0.4$ .

Fixing cosmological parameters to the best-fit values found by *Planck* [32], we determine the linear bias of the WI-SC sample from its auto-correlation on large scales to be  $b_g = 1.1$ , in good agreement with [5]. Fixing the galaxy bias to this value, and the quasar bias to the value found for the High- $z$  bin in our fiducial analysis, we use the cross-correlation between both samples to constrain the value of  $s$ . The right panel of figure 14 shows this cross-correlation, together with the prediction for the best-fit value of the number counts slope (see eq. (5.4)) in red, and that for the value of  $s$  found by [46] for DESI in blue. The value of  $s$  obtained from this cross-correlation is

$$s(z > 1.47) = 0.440 \pm 0.047, \quad (5.4)$$

and provides a good fit to the data ( $\chi^2 = 26.7$  for 26 degrees of freedom). This value is in good agreement with that found with two different methods in section 4.1.2, whereas the



**Figure 14.** *Left:* redshift distribution of the High- $z$  Quaia bin (points with error bars), and of the WISE-SuperCOSMOS sample we cross-correlate it with to validate our estimate of the magnification bias slope. *Right:* cross-correlation between the High- $z$  Quaia bin and the low-redshift WISE-SuperCOSMOS sample. The solid lines show the predictions for the best-fit value of the magnification bias slope  $s$  (red), and for the value found by [46] in the DESI quasar sample (blue). The cross-correlation is able to rule out a value of  $s$  as low as that found in the DESI sample at  $\sim 3.5\sigma$ .


DESI value lies  $\sim 3.5\sigma$  away from it. It is thus unlikely for magnification bias to be the cause of the low  $\sigma_8$  value found from the higher-redshift quasar sample.

### 5.2.6 Foreground contamination in CMB lensing

So far we have mostly explored the impact of various systematics affecting the Quaia sample. One of the most important sources of systematic uncertainty for cross-correlation studies, however, is contamination in the CMB lensing map, particularly from extragalactic foregrounds, tracing the same large-scale structure as the probe being cross-correlated.

To quantify the robustness of our results against this potential systematic, we computed the cross-correlation of the Quaia sample with an alternative map of the CMB convergence constructed using only polarization information (and therefore largely devoid of any extragalactic foreground contamination that is mainly unpolarized) and a PR3-like pipeline run on NPIPE maps. We note that lensing reconstruction based on polarization can be more prone to biases generated by Galactic foregrounds, such as Galactic dust emission, which might in turn correlate with uncorrected dust extinction in the Quaia sample. However, [119] showed that dust biases are mitigated to a negligible level if lensing reconstruction is performed on a CMB map cleaned of foregrounds through component separation (such as those used for PR3 and PR4 lensing analyses), even for data much deeper than *Planck*. The systematic deprojection performed on the Quaia sample would reduce this effect further. Using the methods described below we find no indication of contamination in the cross-correlation with the Low- $z$  bin, and therefore we will limit our discussion to the High- $z$  bin, which displays the largest deviation in  $\sigma_8$  with respect to *Planck* (see section 5.2.1).

Figure 15 shows the cross-correlation of the High- $z$  bin with our fiducial generalised minimum variance (GMV)  $\kappa$  map in red, and with the polarization-only map  $\kappa_{\text{Pol}}$  in orange. Although, in each bandpower, both power spectra are compatible within  $1\sigma$ , we find that the



**Figure 15.** Cross-correlation between the High- $z$  Quiaia redshift bin, and maps of the CMB convergence. Results are shown for the fiducial generalised minimum variance lensing map (red points), and for the polarization-only map (orange triangles). When the power spectrum is negative, we show its absolute value as empty circles with dashed error bars. The polarization-only data is systematically above the GMV estimator on scales  $\ell \lesssim 300$ , although with larger uncertainties, with all deviations at the level of  $\sim 1\sigma$ . The blue solid line shows the best-fit theoretical prediction found in our fiducial analysis. In turn, the dashed blue line shows the theoretical prediction resulting from multiplying the fiducial best-fit curve by a free amplitude and fitting it to the difference between both power spectra. The best-fit curve found this way is  $\sim 30\%$  higher than our fiducial best fit, which is significant at the  $\sim 2\sigma$  level. This translates into a  $\sim 1.5\sigma$  upwards shift in  $\sigma_8$  for this bin when using the polarization-only map, making it fully compatible with the *Planck* value.

polarization-only cross-correlation is systematically  $\sim 30\%$  higher on large scales ( $\ell \lesssim 300$ ). At the *Planck* sensitivity, the polarization-only reconstructed  $\kappa$  map is significantly noisier than the GMV estimator (the errors in the cross-correlation are between  $\sim 2$  and  $\sim 4$  times larger on the scales explored here), and hence we find only mild evidence that this shift is indeed caused by contamination in the  $\kappa$  map. To quantify this in a model-independent way, we first calculate the  $\chi^2$  of the difference between both power spectra with respect to a null model.<sup>13</sup> The PTE of the resulting  $\chi^2$  value is 0.37, and hence, by this metric, both power spectra are compatible with each other within  $1\sigma$ . To quantify the evidence for a coherent deviation between both spectra, we instead take the theoretical prediction for the best-fit parameters found in our fiducial analysis for this cross-correlations, and multiply it by a free amplitude  $A_{\text{Pol}}$ . We fit this 1-parameter model to the difference between both power spectra on scales  $30 \leq \ell \leq 600$ , finding:

$$1 + A_{\text{pol}} = 1.296 \pm 0.125. \quad (5.5)$$

We thus find weak evidence, at the  $2.3\sigma$  level, that the power spectrum computed with the polarization-only  $\kappa$  map is  $\sim 30\%$  higher than that for the minimum-variance map. A more

<sup>13</sup>Note that the covariance of the difference between two variables  $\Delta \mathbf{x}_{12} \equiv \mathbf{x}_1 - \mathbf{x}_2$  is  $\mathbf{C}_{11} + \mathbf{C}_{22} - 2\mathbf{C}_{12}$ , where  $\mathbf{C}_{ij} \equiv \text{Cov}(\mathbf{x}_i, \mathbf{x}_j)$ .

in-depth investigation of this issue, making use of the *Planck* CMB lensing simulations in the PR3 release, is described in appendix A.

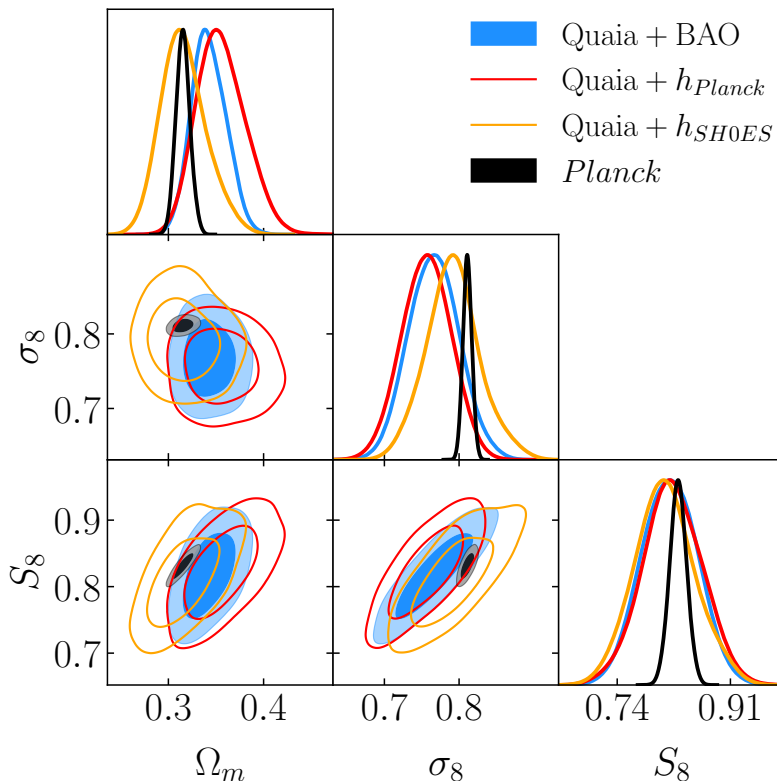
Finally, we propagate the impact of this potential systematics to our final constraints. For the High- $z$  Quaia bin, using the polarization-only  $\kappa$  map, we find the cosmological constraints shown in row 15 of table 2 and figure 8. These constraints are also shown as dashed blue lines in figure 9. We observe a  $\sim 1.4\sigma$  upwards shift in  $\sigma_8$ , which is now fully compatible with *Planck* within  $1\sigma$ . Although the evidence is relatively weak (at the  $\sim 2\sigma$  level), this provides an attractive explanation for the low amplitude of our fiducial cross-correlation between CMB lensing and the High- $z$  bin with respect to the *Planck* prediction. The most likely source of contamination behind this effect would be the Cosmic Infrared Background (CIB), given its relevance at redshifts  $\gtrsim 2$ , and the fact that the resulting bias to the CMB lensing map and power spectrum is expected to be negative (see figure 23 of [47]). The cosmological constraints found using the GMV  $\kappa$  map with the Low- $z$  bin, and the polarization-only map with the High- $z$  bin, are shown in row 16 of table 2 and figure 8. The main effect is an upwards shift of  $\sigma_8$  by  $\sim 1\sigma$ , and a  $\sim 30\%$  increase in the final uncertainties caused by the significantly higher noise of the polarization-only map.

### 5.2.7 Impact of priors

As explained in section 4.4, our data is not sufficient to break the degeneracy between  $\Omega_m$  and  $h$ , and we have relied on low-redshift BAO data to pin down  $\Omega_m$ . To explore the dependence of our results on this choice, we have repeated our analysis with three alternative prior choices.

First, we consider the option of adding, to our fiducial BAO prior, a constraint from the angular sound horizon scale at recombination,  $\theta_{\text{LSS}}$ , from the positions of the peaks in the CMB power spectrum. This is one of the cleanest CMB observables, purely geometrical by nature, and is remarkably well measured ( $\theta_{\text{LSS}} = 1.04109 \pm 0.00030$  [32]). Following [19], we include this constraint by imposing a Gaussian prior on the combination  $\Omega_m h^3 = 0.09633 \pm 0.00029$  (and varying this quantity instead of  $h$  in our MCMC chains). The result is shown in row 17 of table 2 and figure 8. The addition of this prior leads to a significant tightening of the constraint on  $\Omega_m$ , which now agrees fully with *Planck* with almost matching uncertainties. The constraint on  $\sigma_8$ , on the other hand, remains almost unchanged, with only a  $0.3\sigma$  downwards shift (away from *Planck*), and a 5% tightening of its uncertainty. It is worth noting that, in spite of the relatively large change in the posterior distribution of  $\Omega_m$ , the best-fit model in this case is still a good fit to the data (the  $\chi^2$  only increases by  $\sim 2$  with respect to our fiducial best fit).

We then study the impact of replacing the BAO prior by a Gaussian prior directly on  $h$ . We consider two possibilities: a *Planck* prior of the form  $h = 0.674 \pm 0.005$ , and a supernova prior from the SH0ES team [120]:  $h = 0.7304 \pm 0.0104$ . The results are shown in rows 18 and 19 of table 2 and figure 8. The 2D constraints in the  $(\sigma_8, \Omega_m, S_8)$  plane are additionally shown in figure 16. Using the *Planck* prior does not lead to significant changes in the final constraints beyond a broadening of the  $\Omega_m$  uncertainty. Interestingly, using the SH0ES prior instead, shifts  $\Omega_m$  and  $\sigma_8$  closer to the values preferred by *Planck* (by  $\sim 1.3\sigma$  in the case of  $\Omega_m$ ). The origin of the  $H_0$  tension thus has an impact on the interpretation of our constraints (as would be the case for any data set displaying a degeneracy with  $h$ ).



**Figure 16.** Constraints on  $(\sigma_8, \Omega_m, S_8)$  from our analysis using a BAO prior (our fiducial case, in blue), and imposing a prior on the expansion rate today, from *Planck* (red) and from the SH0ES collaboration (orange). Although the constraints on  $\sigma_8$  are only mildly affected by the choice of prior, the effect on  $\Omega_m$  is larger, shifting it by up to  $1.3\sigma$ .

## 6 Conclusions

In this work, we have presented a study of the tomographic clustering of quasars in the Quaia sample, a quasar catalog selected from the third *Gaia* data release in combination with *unWISE* [49]. The analysis is based on the angular auto-power spectrum of Quaia in two redshift bins, centered at  $z = 1.0$  and  $z = 2.1$ , as well as the cross-correlation with CMB lensing data from *Planck*.

We find that, thanks to the large range of scales made available by Quaia’s large cosmological volume, we are able to break the degeneracy between  $\Omega_m$  and  $\sigma_8$  that is usually present in tomographic large-scale structure analyses with lower-redshift samples. We are able to constrain  $\sigma_8$  at a level that is competitive with (and in some cases better than) the precision obtained by state-of-the-art weak lensing and galaxy clustering experiments (see figure 7). Furthermore, our measurement is highly complementary to those datasets, since it makes use of significantly higher-redshift data, and mostly linear scales, where the modelling of astrophysical systematics is, arguably, cleaner. Our constraints on the amplitude of matter fluctuations,  $\sigma_8 = 0.766 \pm 0.034$ , and on the fractional abundance of non-relativistic matter,  $\Omega_m = 0.343^{+0.017}_{-0.019}$ , are compatible with *Planck* at the  $1.4\sigma$  level. We also find no evidence of the so-called “ $S_8$  tension”, although our data is less sensitive to that parameter combination

than weak lensing experiments. We have shown that our measurement is robust against a large suite of potential systematic effects and analysis choices, including the assumed scale cuts, the assumed redshift evolution of the quasar bias, the impact of non-Poissonian shot noise, uncertainties in the redshift distribution of our samples, contamination from residual sky systematics in the galaxy overdensity map, a mis-estimation of the magnification bias slope, and the choice of priors on other cosmological parameters. Furthermore, we find that our measurements of the quasar bias is in reasonable agreement with previous studies.

When extracting constraints from each of the two redshift bins of the catalog independently, we find that the higher redshift sample favours a value of  $\sigma_8$  that is lower than the *Planck* measurement at the  $\sim 2.5\sigma$  level. This is caused by the relatively low amplitude of the quasar-lensing cross-correlation in that bin, and drives the lower value of  $\sigma_8$  found in our fiducial analysis. This result is robust against all sources of uncertainty in the modelling of the Quia sample we have explored (quasar bias evolution, redshift distribution calibration, magnification bias). The explanation for this low amplitude thus most likely lies in the CMB lensing map itself. In particular, by reanalysing the data using an alternative CMB  $\kappa$  map, constructed using only polarization information, we have found compelling (although not definitive) evidence of potential contamination from extragalactic foregrounds, which would lower the amplitude of this particular cross-correlation. The most likely foreground source responsible for this effect would be the Cosmic Infrared Background, given its relevance at  $z \gtrsim 2$ . We find that, using the polarization-only  $\kappa$  map, the amplitude of the quasar- $\kappa$  cross-correlation can grow by up to  $\sim 20\text{--}30\%$ , although the significantly higher noise level of this map prevents us from detecting this shift at more than  $\sim 2\sigma$ . This correction is substantially larger than the nominal level of contamination due to CIB leakage in the CMB lensing power spectrum determined by *Planck* (at the level of  $\sim 1\%$ ), but it is able to bring our results into complete agreement with the *Planck* constraints.

Fully understanding the presence and origin of this potential contamination would be highly relevant, since it could have a substantial effect on other cross-correlation studies involving high-redshift datasets (such as [43, 46, 90, 121, 122]). A better characterisation of the effect could be achieved using more sensitive CMB lensing estimators that are still robust against CIB contamination, such as including  $TE$  information, or using hardening methods [123]. The use of high-resolution ground-based CMB lensing data [7, 34] will also be able to shed light on this effect in the near future.

Another caveat of our analysis is the degeneracy with the value of the Hubble constant  $h$ , which we mitigated by making use of low-redshift BAO data. In the future, other spectroscopic quasar catalogs, such as the sample targeted by DESI [124], may be able to break this degeneracy independently, making use of their own measurements of the redshift-distance relation. We also showed that our final constraints, especially in the case of  $\Omega_m$ , are sensitive to the choice of prior on  $h$  (if one is used), and thus may be affected by the ongoing tension between high redshift CMB and large-scale structure data, and direct measurements from supernovae.

Our analysis has demonstrated the significant and complementary constraining power of high-redshift large-scale structure data sets covering large volumes. The characteristics of the Quia catalog (comoving volume, redshift precision, homogeneity, reduced observational

systematic effects enabled by space-quality data), make it ideal to carry out a number of other cosmological studies. These include constraining the level of large-scale homogeneity [125] and isotropy [126], and detecting ultra-large scale effects, such as primordial non-Gaussianity [127]. The main limiting factor of this sample is its relatively low density, which significantly reduces the amount of information that can be extracted from small (but still only mildly non-linear) scales. This will likely be improved upon by future quasar catalogs from next-generation experiments such as DESI [124] or the Vera C. Rubin Observatory [128], which will be able to improve on the results obtained here. Future *Gaia* data releases may also enable improvements to the catalog in terms of redshift precision, quasar identification, and decontamination, which could allow for improved measurements.

**Data access and software.** The per-bin selection functions used in this analysis, and the complete set of power spectra, covariance, and associated metadata used in our fiducial analysis are publicly available at <https://doi.org/10.5281/zenodo.8098635>. The public Quiaia catalog can be found at <https://doi.org/10.5281/zenodo.8060755>.

This work made extensive use of Astropy [129–131], cobaya [110], healpy [71, 132], matplotlib [133], NaMaster [75], numpy [134], and scipy [135].

## Acknowledgments

We would like to thank Nestor Arsenov, Michael Blanton, Julien Carron, Gerrit Farren, Pedro Ferreira, Andras Kovacs, Anthony Lewis, Blake Sherwin, Lyuba Slavcheva-Mihova, Anže Slosar, and Abby Williams for useful comments and discussions. DA acknowledges support from the Beecroft Trust, and from the John O’Connor Research Fund, at St. Peter’s College, Oxford. GF acknowledges the support of the European Research Council under the Marie Skłodowska Curie actions through the Individual Global Fellowship No. 892401 PiCOGAMBAS and of the Simons Foundation. K.S.F. is supported by the NASA FINESST program under award number 80NSSC20K1545. CGG acknowledges support from the European Research Council Grant No: 693024 and the Beecroft Trust. We made extensive use of computational resources at the University of Oxford Department of Physics, funded by the John Fell Oxford University Press Research Fund, and at the Flatiron Institute.

## A Simulation-based assessment of CMB lensing contamination

To further investigate the potential evidence of contamination described in section 5.2.6, we performed a consistency test of our finding by cross-correlating the Quiaia sample with different CMB lensing maps of the PR3 release. For the PR3 release, in fact, we have access to a complete set of paired signal only ( $\kappa^s$ ) and reconstructed  $\kappa$  maps obtained running MV, polarization-only or temperature-only lensing reconstruction ( $\kappa_{MV}^s, \kappa_{Pol}^s, \kappa_{TT}^s$  respectively) on simulated CMB maps that contain a full end-to-end treatment of the whole Planck data processing chain [136]. As such, those simulated reconstructed  $\kappa$  maps represent statistical realizations of our best knowledge of the Planck lensing data (including instrumental systematics, realistic correlated noise, and foreground residuals after component separation) that can be used to model more accurately the expected difference in bandpowers amplitude

in the cross-correlation between Quaia quasars and different versions of the CMB lensing map. The PR4 has a similar data set only for the GMV and MV lensing reconstruction and, as such, we could only perform consistency tests using the data-driven null model discussed in section 5.2.6. Our pipeline operated as follows:

1. For each  $\kappa^s$  realization, we generated a constrained correlated Gaussian realization of the clustered component of the Quaia sample. For this we assumed the theoretical correlation between the Quaia quasars and CMB lensing described in section 3, for the same cosmology used to obtain the  $\kappa^s$  realizations.
2. For each of these signal-only realizations of the Quaia quasar overdensity, we added a shot-noise component modeled with random realizations of the catalog<sup>14</sup> to produce realistic Quaia-like quasar overdensity maps  $g_{\text{Quaia}}^s$ .
3. For each reconstructed CMB lensing map and its corresponding correlated  $g_{\text{Quaia}}^s$  map, we computed  $C_\ell^{\kappa_{\text{MV}}^s g_{\text{Quaia}}^s}$ ,  $C_\ell^{\kappa_{\text{Pol}}^s g_{\text{Quaia}}^s}$  and their difference for all the different CMB lensing reconstruction methods.<sup>15</sup>
4. We evaluated the covariance of the shifts in bandpowers  $\text{Cov}(C_\ell^{\kappa_{\text{MV}}^s g_{\text{Quaia}}^s} - C_\ell^{\kappa_{\text{Pol}}^s g_{\text{Quaia}}^s})$  averaging over all available realizations.
5. Using this covariance, we computed the  $\chi^2$  for each simulation realization as well as its distribution. We repeated the same procedure for the  $\chi^2$  distribution of each bandpower.

We then assessed the significance of the observed shifts in the data bandpowers  $C_\ell^{g_{\kappa_{\text{MV}}}} - C_\ell^{g_{\kappa_{\text{Pol}}}}$  comparing the value obtained for the  $\chi^2$  performed on data with the  $\chi^2$  distribution obtained from simulations using only the multipoles in the  $30 \leq \ell \leq 600$  range. We adopted the same approach for the  $\chi^2$  test for each bandpower observed for the data. We found a PTE of 6% for the full spectrum, supporting the conclusions derived with our data-based model for PR4, although with worse PTE approaching a  $2\sigma$ -level tension as we found fitting the  $A_{\text{pol}}$  parameter. When considering the  $\chi^2$  test per bandpower, we found that almost all the observed shifts in the data are consistent with simulations (PTEs above 5%). However, the bin at  $\ell \approx 100$  and the one at  $\ell \approx 370$  have a PTE of about 2%, hinting to some mild evidence of inconsistency with the hypothesis of bandpower shifts only due to statistical fluctuations. The feature at  $\ell \approx 100$  in particular is the one driving the low (yet acceptable) PTE for the full  $\chi^2$  test. We also note that lowering the covariance estimated from our Monte Carlo approach by a factor of 20% to mimic the gain in noise obtained in PR4 GMV analysis compared to PR3 did not change these conclusions. We also repeated a similar test comparing the MV with TT-only reconstruction, the MV with the TT lensing reconstruction obtained explicitly nulling the tSZ contribution when performing component separation, and the TT with the TT-only tSZ-deprojected reconstruction. In all these cases we found that the observed shifts in the  $g_\kappa$  bandpowers are consistent with the differences induced by statistical fluctuations, reinforcing the conclusion that, among the extragalactic foregrounds, the CIB emission rather than tSZ is more likely to drive the low amplitude of the  $g_\kappa$  signal at high redshift.

<sup>14</sup>More details on these simulations are given in [127].

<sup>15</sup>We corrected each estimate of cross-correlation power spectra by the corresponding lensing normalization correction for PR3 following the approach used in section 4.2.



## References

- [1] EBOSS collaboration, *Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological implications from two decades of spectroscopic surveys at the Apache Point Observatory*, *Phys. Rev. D* **103** (2021) 083533 [[arXiv:2007.08991](#)] [[INSPIRE](#)].
- [2] C. Heymans et al., *KiDS-1000 Cosmology: Multi-probe weak gravitational lensing and spectroscopic galaxy clustering constraints*, *Astron. Astrophys.* **646** (2021) A140 [[arXiv:2007.15632](#)] [[INSPIRE](#)].
- [3] DES collaboration, *Dark Energy Survey Year 3 results: Cosmological constraints from galaxy clustering and weak lensing*, *Phys. Rev. D* **105** (2022) 023520 [[arXiv:2105.13549](#)] [[INSPIRE](#)].
- [4] DES collaboration, *CMB lensing tomography with the DES Science Verification galaxies*, *Mon. Not. Roy. Astron. Soc.* **456** (2016) 3213 [[arXiv:1507.05551](#)] [[INSPIRE](#)].
- [5] J.A. Peacock and M. Bilicki, *Wide-area tomography of CMB lensing and the growth of cosmological density fluctuations*, *Mon. Not. Roy. Astron. Soc.* **481** (2018) 1133 [[arXiv:1805.11525](#)] [[INSPIRE](#)].
- [6] A. Krolewski, S. Ferraro, E.F. Schlafly and M. White, *unWISE tomography of Planck CMB lensing*, *JCAP* **05** (2020) 047 [[arXiv:1909.07412](#)] [[INSPIRE](#)].
- [7] O. Darwish et al., *The Atacama Cosmology Telescope: A CMB lensing mass map over 2100 square degrees of sky and its cross-correlation with BOSS-CMASS galaxies*, *Mon. Not. Roy. Astron. Soc.* **500** (2020) 2250 [[arXiv:2004.01139](#)] [[INSPIRE](#)].
- [8] DES collaboration, *Constraints on the redshift evolution of astrophysical feedback with Sunyaev-Zel'dovich effect cross-correlations*, *Phys. Rev. D* **100** (2019) 063519 [[arXiv:1904.13347](#)] [[INSPIRE](#)].
- [9] N. Koukoufilippas, D. Alonso, M. Bilicki and J.A. Peacock, *Tomographic measurement of the intergalactic gas pressure through galaxy-tSZ cross-correlations*, *Mon. Not. Roy. Astron. Soc.* **491** (2020) 5464 [[arXiv:1909.09102](#)] [[INSPIRE](#)].
- [10] Y.-K. Chiang, R. Makiya, B. Ménard and E. Komatsu, *The Cosmic Thermal History Probed by Sunyaev-Zeldovich Effect Tomography*, *Astrophys. J.* **902** (2020) 56 [[arXiv:2006.14650](#)] [[INSPIRE](#)].
- [11] Z. Yan et al., *Probing galaxy bias and intergalactic gas pressure with KiDS Galaxies-tSZ-CMB lensing cross-correlations*, *Astron. Astrophys.* **651** (2021) A76 [[arXiv:2102.07701](#)] [[INSPIRE](#)].
- [12] Z. Yan et al., *Cosmic star formation history with tomographic cosmic infrared background-galaxy cross-correlation*, *Astron. Astrophys.* **665** (2022) A52 [[arXiv:2204.01649](#)] [[INSPIRE](#)].
- [13] S.J. Schmidt et al., *Inferring the Redshift Distribution of the Cosmic Infrared Background*, *Mon. Not. Roy. Astron. Soc.* **446** (2015) 2696 [[arXiv:1407.0031](#)] [[INSPIRE](#)].
- [14] B. Jogo et al., *The star-formation history in the last 10 billion years from CIB cross-correlations*, *Mon. Not. Roy. Astron. Soc.* **520** (2023) 1895 [[arXiv:2206.15394](#)] [[INSPIRE](#)].
- [15] M. Rahman et al., *Clustering-based Redshift Estimation: Comparison to Spectroscopic Redshifts*, *Mon. Not. Roy. Astron. Soc.* **447** (2015) 3500 [[arXiv:1407.7860](#)] [[INSPIRE](#)].
- [16] Y.-K. Chiang, B. Ménard and D. Schiminovich, *Broadband Intensity Tomography: Spectral Tagging of the Cosmic UV Background*, *Astrophys. J.* **877** (2019) 150 [[arXiv:1810.00885](#)] [[INSPIRE](#)].

- [17] L.E. Bleem et al., *A Measurement of the Correlation of Galaxy Surveys with CMB Lensing Convergence Maps from the South Pole Telescope*, *Astrophys. J. Lett.* **753** (2012) L9 [[arXiv:1203.4808](#)] [[INSPIRE](#)].
- [18] HERSCHEL ATLAS collaboration, *Cross-correlation between the CMB lensing potential measured by Planck and high- $z$  sub-mm galaxies detected by the Herschel-ATLAS survey*, *Astrophys. J.* **802** (2015) 64 [[arXiv:1410.4502](#)] [[INSPIRE](#)].
- [19] A. Krolewski, S. Ferraro and M. White, *Cosmological constraints from unWISE and Planck CMB lensing tomography*, *JCAP* **12** (2021) 028 [[arXiv:2105.03421](#)] [[INSPIRE](#)].
- [20] Q. Hang, S. Alam, J.A. Peacock and Y.-C. Cai, *Galaxy clustering in the DESI Legacy Survey and its imprint on the CMB*, *Mon. Not. Roy. Astron. Soc.* **501** (2021) 1481 [[arXiv:2010.00466](#)] [[INSPIRE](#)].
- [21] M. White et al., *Cosmological constraints from the tomographic cross-correlation of DESI Luminous Red Galaxies and Planck CMB lensing*, *JCAP* **02** (2022) 007 [[arXiv:2111.09898](#)] [[INSPIRE](#)].
- [22] DES and SPT collaborations, *Joint analysis of Dark Energy Survey Year 3 data and CMB lensing from SPT and Planck. II. Cross-correlation measurements and cosmological constraints*, *Phys. Rev. D* **107** (2023) 023530 [[arXiv:2203.12440](#)] [[INSPIRE](#)].
- [23] G.A. Marques and A. Bernui, *Tomographic analyses of the CMB lensing and galaxy clustering to probe the linear structure growth*, *JCAP* **05** (2020) 052 [[arXiv:1908.04854](#)] [[INSPIRE](#)].
- [24] S.-F. Chen, M. White, J. DeRose and N. Kokron, *Cosmological analysis of three-dimensional BOSS galaxy clustering and Planck CMB lensing cross correlations via Lagrangian perturbation theory*, *JCAP* **07** (2022) 041 [[arXiv:2204.10392](#)] [[INSPIRE](#)].
- [25] C. García-García et al., *The growth of density perturbations in the last  $\sim 10$  billion years from tomographic large-scale structure data*, *JCAP* **10** (2021) 030 [[arXiv:2105.12108](#)] [[INSPIRE](#)].
- [26] DES and SPT collaborations, *Joint analysis of Dark Energy Survey Year 3 data and CMB lensing from SPT and Planck. III. Combined cosmological constraints*, *Phys. Rev. D* **107** (2023) 023531 [[arXiv:2206.10824](#)] [[INSPIRE](#)].
- [27] A. Nicola, A. Refregier and A. Amara, *Integrated approach to cosmology: Combining CMB, large-scale structure and weak lensing*, *Phys. Rev. D* **94** (2016) 083517 [[arXiv:1607.01014](#)] [[INSPIRE](#)].
- [28] R. Sgier et al., *Combined  $13 \times 2$ -point analysis of the Cosmic Microwave Background and Large-Scale Structure: implications for the  $S_8$ -tension and neutrino mass constraints*, [arXiv:2110.03815](#) [[INSPIRE](#)].
- [29] A. Amon et al., *Consistent lensing and clustering in a low- $S_8$  Universe with BOSS, DES Year 3, HSC Year 1 and KiDS-1000*, *Mon. Not. Roy. Astron. Soc.* **518** (2023) 477 [[arXiv:2202.07440](#)] [[INSPIRE](#)].
- [30] S. More et al., *Hyper Suprime-Cam Year 3 Results: Measurements of Clustering of SDSS-BOSS Galaxies, Galaxy-Galaxy Lensing and Cosmic Shear*, [arXiv:2304.00703](#) [[INSPIRE](#)].
- [31] H. Miyatake et al., *Hyper Suprime-Cam Year 3 Results: Cosmology from Galaxy Clustering and Weak Lensing with HSC and SDSS using the Emulator Based Halo Model*, [arXiv:2304.00704](#) [[INSPIRE](#)].
- [32] PLANCK collaboration, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6 [*Erratum ibid.* **652** (2021) C4] [[arXiv:1807.06209](#)] [[INSPIRE](#)].

- [33] ACT collaboration, *The Atacama Cosmology Telescope: DR4 Maps and Cosmological Parameters*, *JCAP* **12** (2020) 047 [[arXiv:2007.07288](#)] [[INSPIRE](#)].
- [34] ACT collaboration, *The Atacama Cosmology Telescope: DR6 Gravitational Lensing Map and Cosmological Parameters*, [arXiv:2304.05203](#) [[INSPIRE](#)].
- [35] A. Amon and G. Efstathiou, *A non-linear solution to the  $S_8$  tension?*, [arXiv:2206.11794](#) [[DOI:10.1093/mnras/stac2429](#)] [[INSPIRE](#)].
- [36] S. Contreras et al., *Consistent and simultaneous modelling of galaxy clustering and galaxy-galaxy lensing with subhalo abundance matching*, *Mon. Not. Roy. Astron. Soc.* **520** (2023) 489 [[arXiv:2211.11745](#)] [[INSPIRE](#)].
- [37] G. Aricò et al., *DES Y3 cosmic shear down to small scales: constraints on cosmology and baryons*, [arXiv:2303.05537](#) [[INSPIRE](#)].
- [38] P. Laurent et al., *Clustering of quasars in SDSS-IV eBOSS: study of potential systematics and bias determination*, *JCAP* **07** (2017) 017 [[arXiv:1705.04718](#)] [[INSPIRE](#)].
- [39] E. Castorina et al., *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](#)].
- [40] R. Neveux et al., *Combined full shape analysis of BOSS galaxies and eBOSS quasars using an iterative emulator*, *Mon. Not. Roy. Astron. Soc.* **516** (2022) 1910 [[arXiv:2201.04679](#)] [[INSPIRE](#)].
- [41] BOSS collaboration, *Quasar-Lyman  $\alpha$  Forest Cross-Correlation from BOSS DR11: Baryon Acoustic Oscillations*, *JCAP* **05** (2014) 027 [[arXiv:1311.1767](#)] [[INSPIRE](#)].
- [42] B.D. Sherwin et al., *The Atacama Cosmology Telescope: Cross-Correlation of CMB Lensing and Quasars*, *Phys. Rev. D* **86** (2012) 083006 [[arXiv:1207.4543](#)] [[INSPIRE](#)].
- [43] D. Alonso, J. Colosimo, A. Font-Ribera and A. Slosar, *Bias of damped Lyman- $\alpha$  systems from their cross-correlation with CMB lensing*, *JCAP* **04** (2018) 053 [[arXiv:1712.02738](#)] [[INSPIRE](#)].
- [44] J. Han, S. Ferraro, E. Giusarma and S. Ho, *Probing Gravitational Lensing of the CMB with SDSS-IV Quasars*, *Mon. Not. Roy. Astron. Soc.* **485** (2019) 1720 [[arXiv:1809.04196](#)] [[INSPIRE](#)].
- [45] J.E. Geach et al., *The Halo Mass of Optically Luminous Quasars at  $z \approx 1-2$  Measured via Gravitational Deflection of the Cosmic Microwave Background*, *Astrophys. J.* **874** (2019) 85 [[arXiv:1902.06955](#)] [[INSPIRE](#)].
- [46] DESI collaboration, *Constraining primordial non-Gaussianity from DESI quasar targets and Planck CMB lensing*, [arXiv:2305.07650](#) [[INSPIRE](#)].
- [47] PLANCK collaboration, *Planck 2018 results. VIII. Gravitational lensing*, *Astron. Astrophys.* **641** (2020) A8 [[arXiv:1807.06210](#)] [[INSPIRE](#)].
- [48] J. Carron, M. Mirmelstein and A. Lewis, *CMB lensing from Planck PR4 maps*, *JCAP* **09** (2022) 039 [[arXiv:2206.07773](#)] [[INSPIRE](#)].
- [49] K. Storey-Fisher et al., *Quaia, the Gaia-unWISE Quasar Catalog: An All-Sky Spectroscopic Quasar Sample*, submitted to *Astrophys. J.* (2023).
- [50] GAIA collaboration, *Gaia Data Release 3 — The extragalactic content*, *Astron. Astrophys.* **674** (2023) A41 [[arXiv:2206.05681](#)] [[INSPIRE](#)].
- [51] 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](#)].

- [52] D. Lang, *unWISE: Unblurred Coadds of the WISE Imaging*, *Astron. J.* **147** (2014) 108 [[arXiv:1405.0308](#)].
- [53] 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](#)].
- [54] B.W. Lyke et al., *The Sloan Digital Sky Survey Quasar Catalog: Sixteenth Data Release*, *Astrophys. J. Suppl.* **250** (2020) 8 [[arXiv:2007.09001](#)] [[INSPIRE](#)].
- [55] PLANCK collaboration, *Planck intermediate results. LVII. Joint Planck LFI and HFI data processing*, *Astron. Astrophys.* **643** (2020) A42 [[arXiv:2007.04997](#)] [[INSPIRE](#)].
- [56] J. Carron, A. Lewis and G. Fabbian, *Planck integrated Sachs-Wolfe-lensing likelihood and the CMB temperature*, *Phys. Rev. D* **106** (2022) 103507 [[arXiv:2209.07395](#)] [[INSPIRE](#)].
- [57] A.S. Maniyar et al., *Quadratic estimators for CMB weak lensing*, *Phys. Rev. D* **103** (2021) 083524 [[arXiv:2101.12193](#)] [[INSPIRE](#)].
- [58] M. Mirmelstein, J. Carron and A. Lewis, *Optimal filtering for CMB lensing reconstruction*, *Phys. Rev. D* **100** (2019) 123509 [[arXiv:1909.02653](#)] [[INSPIRE](#)].
- [59] C. Bonvin and R. Durrer, *What galaxy surveys really measure*, *Phys. Rev. D* **84** (2011) 063505 [[arXiv:1105.5280](#)] [[INSPIRE](#)].
- [60] A. Challinor and A. Lewis, *The linear power spectrum of observed source number counts*, *Phys. Rev. D* **84** (2011) 043516 [[arXiv:1105.5292](#)] [[INSPIRE](#)].
- [61] R. Narayan, *Gravitational lensing and quasar-galaxy correlations*, *Astrophys. J.* **339** (1989) L53.
- [62] P. Schneider, *The number excess of galaxies around high redshift quasars*, *Astron. Astrophys.* **221** (1989) 221.
- [63] D.N. Limber, *The Analysis of Counts of the Extragalactic Nebulae in Terms of a Fluctuating Density Field*, *Astrophys. J.* **117** (1953) 134.
- [64] N. Kaiser, *Weak gravitational lensing of distant galaxies*, *Astrophys. J.* **388** (1992) 272 [[INSPIRE](#)].
- [65] S. Joudaki et al., *KiDS+VIKING-450 and DES-Y1 combined: Cosmology with cosmic shear*, *Astron. Astrophys.* **638** (2020) L1 [[arXiv:1906.09262](#)] [[INSPIRE](#)].
- [66] KiDS collaboration, *KiDS-1000 Cosmology: Cosmic shear constraints and comparison between two point statistics*, *Astron. Astrophys.* **645** (2021) A104 [[arXiv:2007.15633](#)] [[INSPIRE](#)].
- [67] DES collaboration, *Dark Energy Survey Year 3 results: Cosmology from cosmic shear and robustness to data calibration*, *Phys. Rev. D* **105** (2022) 023514 [[arXiv:2105.13543](#)] [[INSPIRE](#)].
- [68] R. Dalal et al., *Hyper Suprime-Cam Year 3 Results: Cosmology from Cosmic Shear Power Spectra*, [arXiv:2304.00701](#) [[INSPIRE](#)].
- [69] S.-S. Li et al., *KiDS-1000: Cosmology with improved cosmic shear measurements*, [arXiv:2306.11124](#) [[INSPIRE](#)].
- [70] M. Bartelmann and P. Schneider, *Weak gravitational lensing*, *Phys. Rept.* **340** (2001) 291 [[astro-ph/9912508](#)] [[INSPIRE](#)].
- [71] 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](#)].
- [72] E.F. Schlafly and D.P. Finkbeiner, *Measuring Reddening with SDSS Stellar Spectra and Recalibrating SFD*, *Astrophys. J.* **737** (2011) 103 [[arXiv:1012.4804](#)] [[INSPIRE](#)].

- [73] T. Cantat-Gaudin et al., *An empirical model of the Gaia DR3 selection function*, *Astron. Astrophys.* **669** (2023) A55 [[arXiv:2208.09335](#)].
- [74] 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](#)].
- [75] LSST DARK ENERGY SCIENCE collaboration, *A unified pseudo- $C_\ell$  framework*, *Mon. Not. Roy. Astron. Soc.* **484** (2019) 4127 [[arXiv:1809.09603](#)] [[INSPIRE](#)].
- [76] A.D. Myers et al., *On statistical lensing and the anti-correlation between 2dF QSOs and foreground galaxies*, *Mon. Not. Roy. Astron. Soc.* **359** (2005) 741 [[astro-ph/0502481](#)] [[INSPIRE](#)].
- [77] SDSS collaboration, *Detection of cosmic magnification with the Sloan Digital Sky Survey*, *Astrophys. J.* **633** (2005) 589 [[astro-ph/0504510](#)] [[INSPIRE](#)].
- [78] M. Lima et al., *Estimating the Redshift Distribution of Faint Galaxy Samples*, *Mon. Not. Roy. Astron. Soc.* **390** (2008) 118 [[arXiv:0801.3822](#)] [[INSPIRE](#)].
- [79] C. García-García et al., *Combining cosmic shear data with correlated photo- $z$  uncertainties: constraints from DESY1 and HSC-DR1*, *JCAP* **01** (2023) 025 [[arXiv:2210.13434](#)] [[INSPIRE](#)].
- [80] B. Hadzhiyska, D. Alonso, A. Nicola and A. Slosar, *Analytic marginalization of  $N(z)$  uncertainties in tomographic galaxy surveys*, *JCAP* **10** (2020) 056 [[arXiv:2007.14989](#)] [[INSPIRE](#)].
- [81] J. Ruiz-Zapatero et al., *Analytical marginalization over photometric redshift uncertainties in cosmic shear analyses*, *Mon. Not. Roy. Astron. Soc.* **522** (2023) 5037 [[arXiv:2301.11978](#)] [[INSPIRE](#)].
- [82] B. Stölzner et al., *Self-calibration and robust propagation of photometric redshift distribution uncertainties in weak gravitational lensing*, *Astron. Astrophys.* **650** (2021) A148 [[arXiv:2012.07707](#)] [[INSPIRE](#)].
- [83] S.M. Croom et al., *The 2dF QSO Redshift Survey. 14. Structure and evolution from the two-point correlation function*, *Mon. Not. Roy. Astron. Soc.* **356** (2005) 415 [[astro-ph/0409314](#)] [[INSPIRE](#)].
- [84] C. Porciani and P. Norberg, *Luminosity and redshift dependent quasar clustering*, *Mon. Not. Roy. Astron. Soc.* **371** (2006) 1824 [[astro-ph/0607348](#)] [[INSPIRE](#)].
- [85] B. Chehade et al., *The 2QDES Pilot: The luminosity and redshift dependence of quasar clustering*, *Mon. Not. Roy. Astron. Soc.* **459** (2016) 1179 [[arXiv:1603.04849](#)] [[INSPIRE](#)].
- [86] Y. Shen et al., *Quasar Clustering from SDSS DR5: Dependences on Physical Properties*, *Astrophys. J.* **697** (2009) 1656 [[arXiv:0810.4144](#)] [[INSPIRE](#)].
- [87] N.P. Ross et al., *Clustering of Low-Redshift ( $z < 2.2$ ) Quasars from the Sloan Digital Sky Survey*, *Astrophys. J.* **697** (2009) 1634 [[arXiv:0903.3230](#)] [[INSPIRE](#)].
- [88] S. Eftekharzadeh et al., *Clustering of intermediate redshift quasars using the final SDSS III-BOSS sample*, *Mon. Not. Roy. Astron. Soc.* **453** (2015) 2779 [[arXiv:1507.08380](#)] [[INSPIRE](#)].
- [89] P. Laurent et al., *A  $14 h^{-3} \text{Gpc}^3$  study of cosmic homogeneity using BOSS DR12 quasar sample*, *JCAP* **11** (2016) 060 [[arXiv:1602.09010](#)] [[INSPIRE](#)].
- [90] X. Lin et al., *Constraining the Halo Mass of Damped Ly $\alpha$  Absorption Systems (DLAs) at  $z = 2-3.5$  Using the Quasar-CMB Lensing Cross-correlation*, *Astrophys. J.* **905** (2020) 176 [[arXiv:2011.01234](#)] [[INSPIRE](#)].
- [91] DESI collaboration, *First Detection of the BAO Signal from Early DESI Data*, [arXiv:2304.08427](#) [[INSPIRE](#)].

- [92] A.D. Myers et al., *First measurement of the clustering evolution of photometrically-classified quasars*, *Astrophys. J.* **638** (2006) 622 [[astro-ph/0510371](#)] [[INSPIRE](#)].
- [93] B. Leistedt, H.V. Peiris and N. Roth, *Constraints on Primordial Non-Gaussianity from 800 000 Photometric Quasars*, *Phys. Rev. Lett.* **113** (2014) 221301 [[arXiv:1405.4315](#)] [[INSPIRE](#)].
- [94] M.A. DiPompeo, R.C. Hickox, S. Eftekharzadeh and A.D. Myers, *The characteristic halo masses of half-a-million WISE-selected quasars*, *Mon. Not. Roy. Astron. Soc.* **469** (2017) 4630 [[arXiv:1705.05306](#)].
- [95] G.C. Petter et al., *Host Dark Matter Halos of SDSS Red and Blue Quasars: No Significant Difference in Large-scale Environment*, *Astrophys. J.* **927** (2022) 16 [[arXiv:2201.07803](#)] [[INSPIRE](#)].
- [96] J.A. Newman, *Calibrating Redshift Distributions Beyond Spectroscopic Limits with Cross-Correlations*, *Astrophys. J.* **684** (2008) 88 [[arXiv:0805.1409](#)] [[INSPIRE](#)].
- [97] B. Ménard et al., *Clustering-based redshift estimation: method and application to data*, [arXiv:1303.4722](#) [[INSPIRE](#)].
- [98] C.B. Morrison et al., *The-wiZZ: Clustering redshift estimation for everyone*, *Mon. Not. Roy. Astron. Soc.* **467** (2017) 3576 [[arXiv:1609.09085](#)] [[INSPIRE](#)].
- [99] J. Grain, M. Tristram and R. Stompor, *Polarized CMB spectrum estimation using the pure pseudo cross-spectrum approach*, *Phys. Rev. D* **79** (2009) 123515 [[arXiv:0903.2350](#)] [[INSPIRE](#)].
- [100] M. Lembo, G. Fabbian, J. Carron and A. Lewis, *CMB lensing reconstruction biases from masking extragalactic sources*, *Phys. Rev. D* **106** (2022) 023525 [[arXiv:2109.13911](#)] [[INSPIRE](#)].
- [101] 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](#)].
- [102] 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](#)].
- [103] BOSS collaboration, *The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample*, *Mon. Not. Roy. Astron. Soc.* **470** (2017) 2617 [[arXiv:1607.03155](#)] [[INSPIRE](#)].
- [104] J.E. Bautista et al., *The Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: measurement of the BAO and growth rate of structure of the luminous red galaxy sample from the anisotropic correlation function between redshifts 0.6 and 1*, *Mon. Not. Roy. Astron. Soc.* **500** (2020) 736 [[arXiv:2007.08993](#)] [[INSPIRE](#)].
- [105] S. Hamimeche and A. Lewis, *Likelihood Analysis of CMB Temperature and Polarization Power Spectra*, *Phys. Rev. D* **77** (2008) 103013 [[arXiv:0801.0554](#)] [[INSPIRE](#)].
- [106] D. Kodwani, D. Alonso and P. Ferreira, *The effect on cosmological parameter estimation of a parameter dependent covariance matrix*, *Open J. Astrophys.* **2** (2019) 3 [[arXiv:1811.11584](#)] [[INSPIRE](#)].
- [107] LSST DARK ENERGY SCIENCE collaboration, *Core Cosmology Library: Precision Cosmological Predictions for LSST*, *Astrophys. J. Suppl.* **242** (2019) 2 [[arXiv:1812.05995](#)] [[INSPIRE](#)].
- [108] D. Blas, J. Lesgourgues and T. Tram, *The Cosmic Linear Anisotropy Solving System (CLASS) II: Approximation schemes*, *JCAP* **07** (2011) 034 [[arXiv:1104.2933](#)] [[INSPIRE](#)].
- [109] R. Takahashi et al., *Revising the Halofit Model for the Nonlinear Matter Power Spectrum*, *Astrophys. J.* **761** (2012) 152 [[arXiv:1208.2701](#)] [[INSPIRE](#)].

- [110] J. Torrado and A. Lewis, *Cobaya: Code for Bayesian Analysis of hierarchical physical models*, *JCAP* **05** (2021) 057 [[arXiv:2005.05290](#)] [[INSPIRE](#)].
- [111] T. Tröster et al., *Cosmology from large-scale structure: Constraining  $\Lambda$ CDM with BOSS*, *Astron. Astrophys.* **633** (2020) L10 [[arXiv:1909.11006](#)] [[INSPIRE](#)].
- [112] O.H.E. Philcox, M.M. Ivanov, M. Simonović and M. Zaldarriaga, *Combining Full-Shape and BAO Analyses of Galaxy Power Spectra: A 1.6% CMB-independent constraint on  $H_0$* , *JCAP* **05** (2020) 032 [[arXiv:2002.04035](#)] [[INSPIRE](#)].
- [113] C. Doux et al., *Cosmological constraints from a joint analysis of cosmic microwave background and spectroscopic tracers of the large-scale structure*, *Mon. Not. Roy. Astron. Soc.* **480** (2018) 5386 [[arXiv:1706.04583](#)] [[INSPIRE](#)].
- [114] DES collaboration, *Dark Energy Survey Year 3 results: Cosmological constraints from galaxy clustering and galaxy-galaxy lensing using the MagLim lens sample*, *Phys. Rev. D* **106** (2022) 103530 [[arXiv:2105.13546](#)] [[INSPIRE](#)].
- [115] S. Sugiyama et al., *HSC Year 1 cosmology results with the minimal bias method: HSC $\times$ BOSS galaxy-galaxy weak lensing and BOSS galaxy clustering*, *Phys. Rev. D* **105** (2022) 123537 [[arXiv:2111.10966](#)] [[INSPIRE](#)].
- [116] S. Sugiyama et al., *Hyper Suprime-Cam Year 3 Results: Cosmology from Galaxy Clustering and Weak Lensing with HSC and SDSS using the Minimal Bias Model*, [arXiv:2304.00705](#) [[INSPIRE](#)].
- [117] LSST collaboration, *Tomographic galaxy clustering with the Subaru Hyper Suprime-Cam first year public data release*, *JCAP* **03** (2020) 044 [[arXiv:1912.08209](#)] [[INSPIRE](#)].
- [118] M. Bilicki et al., *WISE  $x$  SuperCOSMOS photometric redshift catalog: 20 million galaxies over 3pi steradians*, *Astrophys. J. Suppl.* **225** (2016) 5 [[arXiv:1607.01182](#)] [[INSPIRE](#)].
- [119] D. Beck, J. Errard and R. Stompor, *Impact of Polarized Galactic Foreground Emission on CMB Lensing Reconstruction and Delensing of B-Modes*, *JCAP* **06** (2020) 030 [[arXiv:2001.02641](#)] [[INSPIRE](#)].
- [120] A.G. Riess et al., *A Comprehensive Measurement of the Local Value of the Hubble Constant with 1 km s<sup>-1</sup> Mpc<sup>-1</sup> Uncertainty from the Hubble Space Telescope and the SH0ES Team*, *Astrophys. J. Lett.* **934** (2022) L7 [[arXiv:2112.04510](#)] [[INSPIRE](#)].
- [121] POLARBEAR collaboration, *Cross-correlation of POLARBEAR CMB Polarization Lensing with High- $z$  Sub-mm Herschel-ATLAS galaxies*, *Astrophys. J.* **886** (2019) 38 [[arXiv:1903.07046](#)] [[INSPIRE](#)].
- [122] M.J. Wilson and M. White, *Cosmology with dropout selection: straw-man surveys & CMB lensing*, *JCAP* **10** (2019) 015 [[arXiv:1904.13378](#)] [[INSPIRE](#)].
- [123] N. Sailer, E. Schaan and S. Ferraro, *Lower bias, lower noise CMB lensing with foreground-hardened estimators*, *Phys. Rev. D* **102** (2020) 063517 [[arXiv:2007.04325](#)] [[INSPIRE](#)].
- [124] E. Chaussidon et al., *Target Selection and Validation of DESI Quasars*, *Astrophys. J.* **944** (2023) 107 [[arXiv:2208.08511](#)] [[INSPIRE](#)].
- [125] D.W. Hogg et al., in preparation (2023).
- [126] A. Williams et al., in preparation (2023).
- [127] G. Fabbian et al., in preparation (2023).
- [128] Ž. Ivezić, *LSST survey: millions and millions of quasars*, *Proc. Int. Astron. Union* **12** (2016) 330.

- [129] ASTROPY collaboration, *The Astropy Project: Sustaining and Growing a Community-oriented Open-source Project and the Latest Major Release (v5.0) of the Core Package\**, *Astrophys. J.* **935** (2022) 167 [[arXiv:2206.14220](#)] [[INSPIRE](#)].
- [130] ASTROPY collaboration, *The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package*, *Astron. J.* **156** (2018) 123 [[arXiv:1801.02634](#)] [[INSPIRE](#)].
- [131] ASTROPY collaboration, *Astropy: A Community Python Package for Astronomy*, *Astron. Astrophys.* **558** (2013) A33 [[arXiv:1307.6212](#)] [[INSPIRE](#)].
- [132] 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](#)].
- [133] J.D. Hunter, *Matplotlib: A 2D Graphics Environment*, *Comput. Sci. Eng.* **9** (2007) 90 [[INSPIRE](#)].
- [134] S. van der Walt, S.C. Colbert and G. Varoquaux, *The NumPy Array: A Structure for Efficient Numerical Computation*, *Comput. Sci. Eng.* **13** (2011) 22 [[arXiv:1102.1523](#)] [[INSPIRE](#)].
- [135] P. Virtanen et al., *SciPy 1.0 — Fundamental Algorithms for Scientific Computing in Python*, *Nature Meth.* **17** (2020) 261 [[arXiv:1907.10121](#)] [[INSPIRE](#)].
- [136] PLANCK collaboration, *Planck 2018 results. III. High Frequency Instrument data processing and frequency maps*, *Astron. Astrophys.* **641** (2020) A3 [[arXiv:1807.06207](#)] [[INSPIRE](#)].