

Cross-correlating radio continuum surveys and CMB lensing: constraining redshift distributions, galaxy bias, and cosmology

David Alonso,¹★ Emilio Bellini,¹ Catherine Hale,² Matt J. Jarvis,^{1,3} and Dominik J. Schwarz⁴

¹*Astrophysics, University of Oxford, DWB, Keble Road, Oxford OX1 3RH, UK*

²*CSIRO Astronomy and Space Science, P.O. Box 1130, Bentley WA 6102, Australia*

³*Department of Physics, University of the Western Cape, Bellville 7535, South Africa*

⁴*Fakultät für Physik, Universität Bielefeld, Postfach 100131, D-33501 Bielefeld, Germany*

Accepted 2021 January 6. Received 2021 January 4; in original form 2020 September 3

ABSTRACT

We measure the harmonic-space auto-power spectrum of the galaxy overdensity in the LOFAR Two-metre Sky Survey (LoTSS) first data release and its cross-correlation with the map of the lensing convergence of the cosmic microwave background (CMB) from the *Planck* collaboration. We report a $\sim 5\sigma$ detection of the cross-correlation. We show that the combination of the clustering power spectrum and CMB lensing cross-correlation allows us to place constraints on the high-redshift tail of the redshift distribution, one of the largest sources of uncertainty in the use of continuum surveys for cosmology. Our analysis shows a preference for a broader redshift tail than that predicted by the photometric redshifts contained in the LoTSS value-added catalogue, as expected, and more compatible with predictions from simulations and spectroscopic data. Although the ability of CMB lensing to constrain the width and tail of the redshift distribution could also be valuable for the analysis of current and future photometric weak lensing surveys, we show that its performance relies strongly on the redshift evolution of the galaxy bias. Assuming the redshift distribution predicted by the Square Kilometre Array Design simulations, we use our measurements to place constraints on the linear bias of radio galaxies and the amplitude of matter inhomogeneities σ_8 , finding $\sigma_8 = 0.69^{+0.14}_{-0.21}$ assuming the galaxy bias scales with the inverse of the linear growth factor, and $\sigma_8 = 0.79^{+0.17}_{-0.32}$ assuming a constant bias.

Key words: methods: data analysis – large-scale structure of Universe.

1 INTRODUCTION

Radio continuum surveys have long attracted the attention of the cosmological community as deep tracers of the large-scale structure, able to probe enormous swathes of the Universe. Several forecasting studies have been carried out to explore the cosmological potential of these probes with the Square Kilometre Array (SKA, Raccanelli et al. 2012; Jarvis et al. 2015). In particular, it has been argued that, using different radio galaxy populations as individual tracers, continuum surveys could be useful in detecting ultra-large-scale effects, such as relativistic light-cone effects and the impact of primordial non-Gaussianity (Ferramacho et al. 2014; Alonso & Ferreira 2015; Bengaly et al. 2019; Gomes et al. 2020).

Extragalactic radio continuum surveys produce a very different view of the Universe compared to those at shorter wavelengths (e.g. in the optical and IR). At ~ 0.1 –10 GHz, the dominant mechanism for radio continuum emission is from synchrotron radiation (Condon 1992), which originates from relativistic electrons spiralling in the magnetic fields. Due to this, radio observations are typically dominated by two main galaxy populations: star-forming galaxies (SFGs) and active galactic nuclei (AGNs). SFGs are often classed into normal star-forming galaxies (with star formation rates, $\text{SFR} \lesssim 100 \text{ M}_\odot \text{ yr}^{-1}$) as well as starburst galaxies, where intense periods

of star formation are present ($\text{SFR} \gtrsim 100 \text{ M}_\odot \text{ yr}^{-1}$). For AGN, there is a diverse polychotomy of sources as well as classification schemes. This includes classifying sources based on orientation to the observer (Antonucci 1993; Urry & Padovani 1995), morphology (e.g. Fanaroff–Riley Type I and II sources, Fanaroff & Riley 1974), or the accretion mechanism of material on to the central supermassive black hole (high/low excitation radio galaxies, e.g. Best & Heckman 2012; Heckman & Best 2014). A key advantage in observing both AGNs and SFGs at these frequencies is that attenuation of emission by dust, in the intergalactic medium and along the line of sight, is negligible. This is especially useful for the study of SFGs, where the radio emission can be used as an unbiased measurement of SFR (see e.g. Bell 2003; Davies et al. 2017; Gürkan et al. 2018).

Moreover, at these low frequencies, the instantaneous field of view from radio surveys is often large. For example, the LOW Frequency ARray (LOFAR; van Haarlem et al. 2013) has a field of view of $\sim 30 \text{ deg}^2$ at 150 MHz, the Meer Karoo Array Telescope (MeerKAT; Jonas 2009) has a field of view of $\sim 1 \text{ deg}^2$ at 1.2 GHz, and the Australian Square Kilometre Array Pathfinder (Johnston et al. 2007) has an instantaneous field of view of $\sim 30 \text{ deg}^2$, due to its use of phased array feeds. These large fields of view for radio interferometers are crucial for producing large, contiguous observations of the celestial sphere, which is especially advantageous for studying large-scale structure at large angular separations. As such, there exist many large sky surveys at radio frequencies, including: the NRAO VLA Sky Survey (NVSS; Condon et al. 1998, at

★ E-mail: David.Alonso@physics.ox.ac.uk

1.4 GHz), TIFR GMRT Sky Survey (Intema et al. 2017, at 150 MHz), and Sydney University Molongolo Sky Survey (Mauch et al. 2003). Further surveys such as the Evolutionary Map of the Universe (EMU; Norris et al. 2011) and the LOFAR Two-metre Sky Survey (LoTSS; Shimwell et al. 2019), will soon provide a huge leap by combining both depth and sky coverage (e.g. see fig. 1 of Shimwell et al. 2019).

Already though, the first data release (DR1) of LoTSS (Shimwell et al. 2019) has generated a low-frequency radio continuum survey, covering 424 deg² at 144 MHz to a depth of $\sim 70 \mu\text{Jy beam}^{-1}$. This is one of the deepest surveys available at this frequency that covers a large area of sky, ideal for studies of large-scale structure (see Siewert et al. 2019).

However, the main complication in using continuum data for cosmology is the absence of reliable redshifts for a large fraction of sources. This leads to large uncertainties in the redshift distribution of different samples, and in general on the redshift evolution of any of their properties, such as the galaxy bias or host halo masses. The large-scale structure analysis of continuum samples with optical matches, has been a useful method to address these issues and improve our understanding of the clustering properties of radio galaxies (Lindsay et al. 2014a; Hale et al. 2018; Siewert et al. 2019).

Further information can be gained through cross-correlations with other tracers of the large-scale structure (e.g. Lindsay, Jarvis & McAlpine 2014b). In particular, since the gravitational lensing of the cosmic microwave background (CMB) is sensitive to the distribution of matter inhomogeneities to very high redshifts, it constitutes an ideal probe to cross-correlate with deep radio continuum data (e.g. Planck Collaboration XVII 2014; Allison et al. 2015), including in the context of de-lensing (Namikawa et al. 2016). In this paper, we will explore the combination of the auto-correlation of the LoTSS survey and its cross-correlation with CMB lensing data from the *Planck* satellite. In particular, we will show how the different dependence of both correlations on the galaxy bias and redshift distribution of the sample allows us to use this combination to simultaneously constrain the bias and the high-redshift tail of the redshift distribution. This possibility is also relevant in the context of weak lensing analyses with current and future deep imaging surveys (Wright et al. 2020; Hildebrandt et al. 2020), for which the robustness of their cosmological constraints will rely heavily on their ability to calibrate the redshift distributions of their samples.

This paper is structured as follows: in Section 2, we present the theoretical background describing the auto- and cross-correlation signals. Section 3 presents the different data sets used in our analysis. The methods used to extract the auto- and cross-correlations, and to analyse them, are described in Section 4. The main results are then presented and discussed in Section 5, and we conclude in Section 6.

2 THEORY

2.1 Galaxy overdensity and CMB lensing

The two probes studied here are the projected overdensity of galaxies in the LoTSS sample, δ_g , and the gravitational lensing convergence of the CMB, κ .

The projected overdensity quantifies the overabundance of galaxies in a given sky position \hat{n} , $N_g(\hat{n})$ with respect to the sky average \bar{N}_g

$$\delta_g(\hat{n}) \equiv \frac{N_g(\hat{n}) - \bar{N}_g}{\bar{N}_g}, \quad (1)$$

and is related to the 3D galaxy overdensity, $\Delta_g(\mathbf{x}, z)$ (defined in a similar way as the fluctuation in the comoving number density of

galaxies at redshift z) through

$$\delta_g(\hat{n}) = \int dz \frac{dp}{dz} \Delta_g(\chi(z)\hat{n}, z), \quad (2)$$

where dp/dz is the redshift distribution of the galaxy sample normalized to 1 when integrated over z , and χ is the comoving radial distance.

The lensing convergence $\kappa(\hat{n})$ quantifies the distortion in the trajectories of the CMB photons caused by the gravitational potential of the intervening matter structures, and is defined to be proportional to the divergence of the deflection in the photon arrival angle α : $\kappa \equiv -\nabla \cdot \alpha/2$. As such, κ is an unbiased tracer of the matter density fluctuations $\Delta_m(\mathbf{x}, z)$, and is related to them through:

$$\kappa(\hat{n}) = \int_0^{\chi_{\text{LSS}}} d\chi \frac{3H_0^2 \Omega_m}{2a} \chi \frac{\chi_{\text{LSS}} - \chi}{\chi_{\text{LSS}}} \Delta_m(\chi\hat{n}, z(\chi)) \quad (3)$$

where H_0 is the Hubble constant, Ω_m is the fractional matter density, $a = 1/(1+z)$ is the scale factor, and χ_{LSS} is the comoving distance to the surface of last scattering.

We will use the harmonic-space correlation between κ and δ_g to study the connection between Δ_g and Δ_m .

2.2 Angular power spectra

Both κ and δ_g can be generically described as a 2D field $u(\hat{n})$ related to an underlying 3D quantity $U(\mathbf{x}, z)$ through a projection on to the 2D sphere with a given radial kernel $W_u(\chi)$:

$$u(\hat{n}) = \int d\chi W_u(\chi) U(\chi\hat{n}, z(\chi)). \quad (4)$$

Any such projected quantity can be decomposed in terms of its spherical harmonic coefficients $u_{\ell m}$, the covariance of which is the so-called angular power spectrum:

$$\langle u_{\ell m} u_{\ell' m'}^* \rangle = \delta_{\ell\ell'} \delta_{mm'} C_\ell^{uv} \quad (5)$$

The angular power spectrum can be related to the power spectrum of the 3D fields P_{UV} through

$$C_\ell^{uv} = \int \frac{d\chi}{\chi^2} W_u(\chi) W_v(\chi) P_{UV} \left(k = \frac{\ell + 1/2}{\chi}, z(\chi) \right), \quad (6)$$

where $P_{UV}(k, z)$ is the variance of the Fourier coefficients of U and V :

$$\langle U(\mathbf{k}, z) V^*(\mathbf{k}', z) \rangle \equiv (2\pi)^3 \delta(\mathbf{k} - \mathbf{k}') P_{UV}(k, z). \quad (7)$$

We use the following definition for the Fourier transform:

$$U(\mathbf{x}, z) = \int \frac{d\mathbf{k}^3}{(2\pi)^2} e^{i\mathbf{k} \cdot \mathbf{x}} U(\mathbf{k}, z). \quad (8)$$

For the two fields under consideration, the radial kernels are given by

$$W_g(\chi) = \frac{H(z)}{c} \frac{dp}{dz}, \quad (9)$$

$$W_\kappa(\chi) = f_\ell \frac{3H_0^2 \Omega_m}{2a} \chi \frac{\chi_{\text{LSS}} - \chi}{\chi_{\text{LSS}}} \Theta(\chi_{\text{LSS}} - \chi) \quad (10)$$

where $H(z)$ is the expansion rate, and $\Theta(x)$ is the Heaviside function. Note that we have included a scale-dependent prefactor f_ℓ in the lensing kernel, given by

$$f_\ell = \frac{\ell(\ell+1)}{(\ell+1/2)^2} \simeq 1. \quad (11)$$

The presence of this factor is only relevant for $\ell \lesssim 10$, and accounts for the fact that κ is related to Δ_m through the angular Laplacian of

the gravitational potential Φ . We have also neglected the effects of lensing magnification, which are much smaller than our statistical uncertainties (Allison et al. 2015).¹

It is worth noting that equation (6) is strictly valid only in the Limber approximation, where the spherical Bessel functions are approximated by

$$j_\ell(x) \simeq \sqrt{\frac{\pi}{2\ell+1}} \delta(x - \ell - 1/2). \quad (12)$$

This is accurate when the radial kernels W_u are broader than the typical correlation length of the density inhomogeneities, as is the case for both δ_g and κ .

The final ingredient of the model needed to describe the measured angular power spectra is the auto- and cross-spectra between the matter and galaxy overdensity fields, $P_{gg}(k, z)$, $P_{gm}(k, z)$, and $P_{mm}(k, z)$. On the scales $k \lesssim 0.2 h \text{ Mpc}^{-1}$ we probe here, we assume a linear, scale-independent bias model between both fields: $\Delta_g(\mathbf{k}, z) = b_g(z)\Delta_\kappa(\mathbf{k}, z)$, such that

$$P_{gm}(k, z) = b_g(z) P_{mm}(k, z), \quad (13)$$

$$P_{gg}(k, z) = b_g^2(z) P_{mm}(k, z). \quad (14)$$

We explore two different models for the redshift dependence of the bias:

(i) A *constant bias* model, in which b_g does not evolve, and the clustering of galaxies follows the same growth as a function of time as the matter inhomogeneities.

(ii) A *constant amplitude* model, in which b_g evolves inversely with the linear growth factor $D(z)$

$$b_g(z) = b_g/D(z), \quad (15)$$

where the growth factor $D(z) \equiv \Delta_m(\mathbf{k}, z)/\Delta_m(\mathbf{k}, z=0)$ describes the evolution of the matter fluctuations on linear scales ($k \rightarrow 0$). In this scenario, the amplitude of Δ_g does not vary as a function of time on linear scales. This would correspond to a galaxy distribution that is fixed at some early time and preserves its large-scale properties unchanged.

These two models represent the two extremes of the expected redshift evolution of $b_g(z)$. Flux-density limited samples such as the one studied here are expected to have b_g grow with z , since more distant sources are intrinsically brighter and are therefore likely to reside in more massive haloes. At the same time, the LoTSS sample contains a wide variety of galaxy types, with different biases and different relative abundances as a function of redshift, and therefore the growth of $b_g(z)$ could be slower than the $1/D(z)$ model.

Finally, we model the matter power spectrum $P_{mm}(k, z)$ using the HALOFIT parametrization of Takahashi et al. (2012) as implemented in the CAMB Boltzmann code (Lewis, Challinor & Lasenby 2000). We use the Core Cosmology Library (CCL; Chisari et al. 2019) for all theoretical calculations.

3 DATA

3.1 The LOFAR Two-metre Sky Survey data release 1

The radio continuum sample used here is contained in the DR1 of the LoTSS. Whilst the full LoTSS survey will eventually cover the entire

northern sky, this DR1 covers a contiguous patch of 424 deg^2 in the Northern hemisphere over the HETDEX spring field, and contains a total of 325 694 sources. The DR1 is described in detail in Shimwell et al. (2019), and the clustering properties of the sample were studied for the first time in Siewert et al. (2019).

DR1 consists of 58 individual pointings, each covering an area approximately the shape of a circle of radius 1.7° at 6 arcsec resolution. Each of the 58 pointings was observed for a duration of 8 h, resulting in an rms noise of $\sim 70 \mu\text{Jy beam}^{-1}$. Siewert et al. (2019) found five of these pointings to result in a poor completeness (at the level of ~ 90 per cent at flux $S \simeq 1 \text{ mJy}$), we discard these pointings in our analysis as well. The average point-source completeness of the remaining area is approximately 99 per cent level for $S > 0.8 \text{ mJy}$.

Through a combination of likelihood-ratio matching and visual identification, a value-added catalogue (VAC) was provided as part of the DR1, which we use for our analysis (Williams et al. 2019). The VAC contains 318 520 sources, of which 73 per cent have additional optical and/or infrared information, primarily from cross-matches with Panoramic Survey Telescope and Rapid Response System (Pan-STARRS) and/or Wide-field Infrared Survey Explorer (WISE). For approximately half of the sample, redshifts were also assigned to sources within the VAC (Duncan et al. 2019). The majority of these are photometric redshifts, whilst a minority of sources have associated spectroscopic redshifts. Of those sources included in the VAC, we use a fiducial flux limited sample defined by a cut in total flux of $S > 2 \text{ mJy}$. We additionally discard all sources with a signal-to-noise ratio $S/N < 5$. The remaining sample contains 57 928 sources.

3.2 The Planck lensing maps

We use the publicly available CMB lensing convergence map provided by the Planck Collaboration VIII (2018b). The data are provided in the form of spherical harmonic coefficients $\kappa_{\ell m}$, which we transform into a HEALPIX map (Górski et al. 2005) with resolution parameter $N_{\text{side}} = 2048$, corresponding to an angular size of ~ 2 arcmin. We verified that this choice of resolution has a negligible impact on our results by re-calculating the δ_g - κ cross-correlation presented in Section 5 for $N_{\text{side}} = 256$. When estimating power spectra, we mask this convergence map using the mask provided by Planck, covering around 67 per cent of the sky. This mask was designed to remove regions dominated by galactic foregrounds as well as strong Sunyaev-Zel'dovich sources. This mask removes around ~ 2 per cent of the LoTSS footprint.

3.3 Redshift distributions

The redshift distribution of the LoTSS sample used here is a key ingredient in order to obtain a theoretical prediction for the clustering auto-spectrum and the CMB lensing cross-spectrum (dp/dz in equation 2).

The LoTSS VAC provides photometric redshift (photo- z) information for about 51 per cent of the total sample (and ~ 48 per cent of our flux-limited sample), which can be used to estimate this redshift distribution. There are however significant uncertainties associated with these photometric redshifts. Of particular concern is the level to which the subset of the catalogue with measured redshifts is representative of the full sample, especially at high redshifts. Although the distribution of radio fluxes in the two samples (the full sample and the subset with measured redshifts), shows no significant differences, this is no guarantee that other selection

¹This may not be the case for the final LoTSS sample with smaller uncertainties.




Figure 1. Estimated redshift distribution of the flux-limited sample used in this work. The red histogram shows the distribution estimated from the photometric redshifts contained in the LoTSS VAC. The green histogram shows the estimate from the spectroscopic VLA-COSMOS sample at 3 GHz extrapolated to the LOFAR band. The blue histogram shows the distribution obtained from the SKADS simulation, based on existing measurements of the radio luminosity function. The dashed lines are empirical fits to the three distributions using the expression in equation (16) with a single free parameter z_{tail} . The black solid line shows the CMB lensing kernel normalized to make it viewable on the same scale. The distributions shown here are plotted as probability distributions dp/dz .

effects (e.g. cuts in the optical photometry) do not induce significant differences in their redshift distributions. This may be especially important in the tail of high-redshift AGN which is typically found in radio surveys. It is likely that the $N(z)$ distribution from Duncan et al. (2019) may underestimate the number of high-redshift sources, especially given radio emission does not suffer dust attenuation whilst the cross-matched catalogues is affected by this, especially for distant objects. Another potential source of error would be misidentification of high-redshift AGNs as low-redshift SFGs. Also, as noted in Siewert et al. (2019), the low- and high-redshift differential source count distributions show markedly different shapes. Therefore other estimates of dp/dz are needed in order to assess the reliability of our results. The redshift distribution estimated from the LoTSS VAC is shown in red in Fig. 1 as a solid line.

We obtain a second estimate of the redshift distribution from the Very Large Array Cosmic Evolution Survey (VLA-COSMOS) 3 GHz catalogue (Smolčić et al. 2017), which comprises $\sim 10\,000$ radio sources in a 2 sq-deg patch covering the COSMOS field in combination with optical and infrared data. Due to the vast wealth of deep multiwavelength data in the COSMOS field, 8995 of these sources have measured redshifts. This is likely to therefore provide a much better estimate of the high-redshift tail of radio sources. To obtain an estimate of the redshift distribution of the LoTSS sample, we extrapolate the radio fluxes of these objects to the LOFAR 144 MHz band assuming a spectral index $\alpha = -0.7^2$ and impose the same flux cut of used for our sample (see Section 3.1). Due to this relatively bright flux cut, the resulting sample contains only 378 galaxies. Although the small size of this sample prevents us from obtaining a well-resolved measurement of the redshift distribution, it is enough to quantify the amplitude of the high-redshift tail

in comparison with the LoTSS VAC measurement. The resulting redshift distribution, estimated by binning these galaxies into 20 linear redshift bins in the range $0 < z < 5$, is shown in green in Fig. 1. Although the VLA-COSMOS catalogue potentially underestimates the nearest extended objects due to its baseline sensitivity to extended emission, it displays a significantly larger high-redshift tail than the LoTSS VAC. As we will see, this high-redshift tail plays a significant role in determining the relative amplitudes of C_{ℓ}^{gg} and C_{ℓ}^{gk} .

Finally, we use the Square Kilometre Array Design Study Simulated Skies (SKADS), in particular the semi-empirical simulation (Wilman et al. 2008), to obtain a third estimate of the redshift distribution of our sample. This simulation contains a realistic distribution of radio galaxies, including AGN and SFGs, down to flux densities of 10 nJy based on estimates of the radio luminosity function for different source types, on a square 100 deg^2 patch. Each source has measured flux densities at several frequencies in the range 151 MHz to 18 GHz, which we use to infer their flux in the LOFAR band. The redshift distribution is then computed by binning all sources in the simulation with fluxes above our cut of 2 mJy. The resulting distribution is shown in blue in Fig. 1. Much like the distribution inferred from the VLA-COSMOS catalogue, the SKADS estimate displays a long high-redshift tail when compared with the LoTSS dp/dz . There is evidence though, that SKADS may underestimate the number of SFGs below flux densities of $S_{1.4\text{GHz}} < 0.1\text{ mJy}$ (Smolčić et al. 2017), however as this is below the survey limits considered in this work, the effect of this is likely to be minimal. Estimates of dp/dz will be improved in the future with deeper observations over well-studied multiwavelength fields, such as through the MIGHTEE (MeerKAT International Giga-Hertz Tiered Extragalactic Exploration, Jarvis et al. 2016) or deep-field LOFAR observations.

This large uncertainty in the true redshift distribution for radio continuum surveys motivates the use of the cross-correlation with CMB lensing as a means to mitigate it. To aid in this study, we fit these three redshift distributions to the same analytic expression

$$\frac{dp}{dz} \propto \frac{(z/z_0)^2}{1 + (z/z_0)^2} \frac{1}{1 + (z/z_{\text{tail}})^{\gamma}}. \quad (16)$$

The normalizing prefactor is determined by the requirement that the integral over dp/dz be unity. We find that $z_0 = 0.1$ and $\gamma = 3.5$ provide a good visual fit to the three redshift distributions, and we verified that varying the chosen z_0 in the range $[0, 0.25]$ had a negligible impact on the main results presented here. The free variable z_{tail} parametrizes the extent of the high-redshift tail of the distribution, and therefore its width. We find that this expression is able to provide a rough fit to the three redshift distributions with $z_{\text{tail}} = 0.8, 1.5$, and 2.0 for the LoTSS, VLA-COSMOS, and SKADS $N(z)$ s, respectively. These analytic fits are shown as dashed curves in Fig. 1. The figure also shows, as a black solid line, the shape of the CMB lensing kernel, which peaks at $z \sim 2$ but extends to significantly higher redshifts.

4 METHODS

4.1 Mean density, sky mask, and systematics

A crucial aspect of any analysis involving the fluctuations in galaxy number counts δ_g , is estimating the expected mean number of objects in each sky pixel. This involves accurately characterizing the geometry of the observed footprint, as well as the expected fluctuations in the number of observed sources due to variations in image noise. We will label these two the *mask* and *mean density* map respectively.

² α is defined as $S_{\nu} \propto \nu^{\alpha}$.

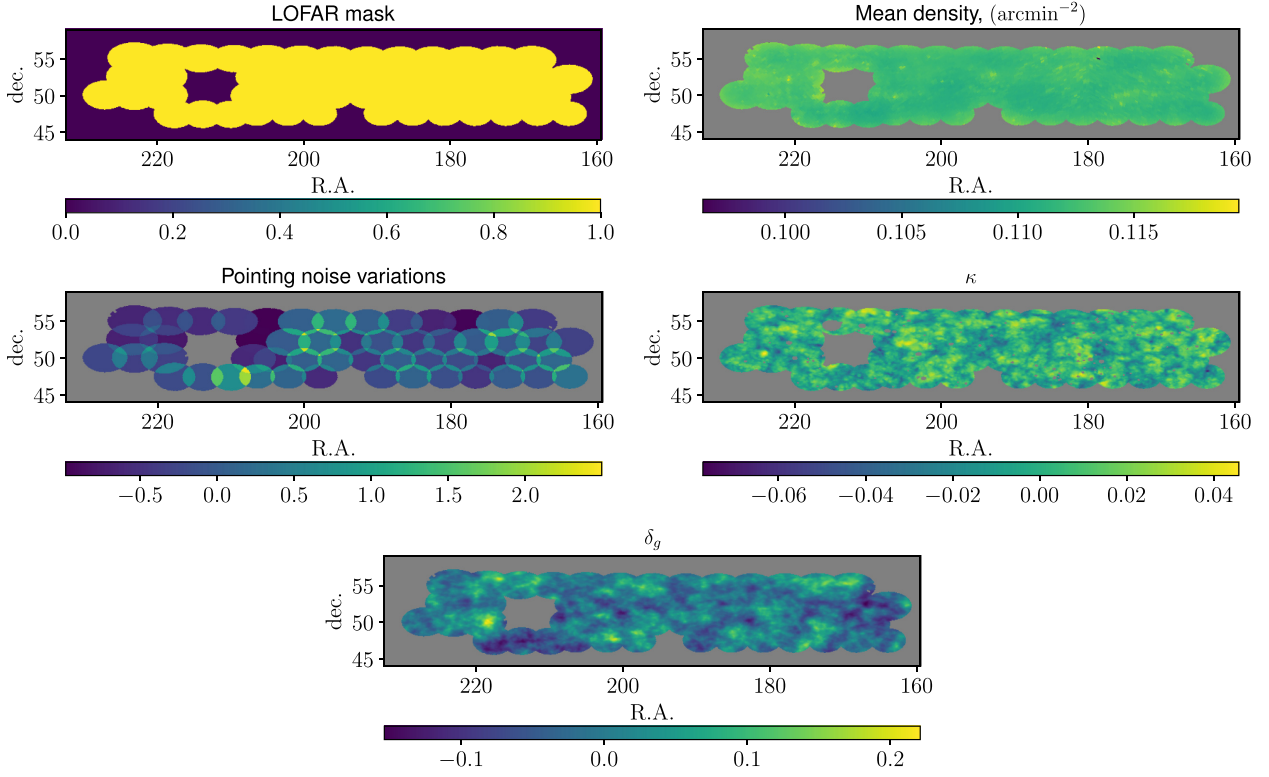


Figure 2. Various maps used in the analysis. Top left: binary mask for the LoTSS DR1 sample. Top right: $S_{\text{cut}} = 2$ mJy mean density map. Middle left: relative fluctuations in the noise rms. Middle right: Wiener-filtered CMB lensing convergence map. Bottom: galaxy overdensity map smoothed with the same filter.

4.1.1 Sky mask

We produce masks for each of the 58 pointings from their associated low-resolution images at resolution $N_{\text{side}} = 2048$. Each mask is a binary sky map containing 0s or 1s depending on whether the centre of a given pixel lies inside the pointing. The final mask is then the logical addition of the 58 pointing masks.

When computing power spectra with the lower resolution maps, the associated mask is computed by averaging down the $N_{\text{side}} = 2048$ masks. The resulting low-resolution mask contains, in each pixel, the fraction of its area that was observed. This mask is shown in the top left panel of Fig. 2.

4.1.2 Mean density

The expected mean number of observed objects in a given pixel varies across the sky as a result of spatial variations in observing conditions, resulting on fluctuations in flux noise. If not accounted for (i.e. if interpreted as intrinsic fluctuations in the galaxy number density), they will bias the estimated two-point functions and, through them, the final parameter constraints. We used a semi-analytic method to reconstruct the expected fluctuations in the mean density for a given sample that uses the same noise model used by Siewert et al. (2019).

Let S_{obs} be the observed flux of a given object with true flux S_{true} . The probability to detect objects with observed fluxes above a given threshold S_{thr} is:

$$P(> S_{\text{thr}}) = \int_0^\infty dS_{\text{true}} p(S_{\text{true}}) \Pi(S_{\text{true}}, S_{\text{thr}})$$

where $p(S_{\text{true}})$ is the distribution of true fluxes, and $\Pi(S_{\text{true}}, S_{\text{thr}})$ is the probability for a source with true flux S_{true} to have an observed flux

above S_{thr} , given in terms of the conditional distribution of observed fluxes $p(S_{\text{obs}}|S_{\text{true}})$ as

$$\Pi(S_{\text{true}}, S_{\text{thr}}) \equiv \int_{S_{\text{thr}}}^\infty dS_{\text{obs}} p(S_{\text{obs}}|S_{\text{true}}). \quad (17)$$

Assuming the noise in the measured fluxes has a Gaussian distribution with standard deviation σ_S (as done in Siewert et al. 2019), this is simply

$$\Pi(S_{\text{true}}, S_{\text{thr}}) = \frac{1}{2} \text{erfc} \left[\frac{S_{\text{true}} - S_{\text{thr}}}{\sqrt{2}\sigma_N} \right], \quad (18)$$

where erfc is the complementary error function.

We therefore need two ingredients in order to compute $P(> S_{\text{thr}})$:

(i) An estimate of the noise rms in each pixel $\sigma_N(\hat{\mathbf{n}})$. As done in Siewert et al. (2019), we make histograms of the rms noise for all sources in a given pixel, and use them to produce maps of the mean and median rms noise per pixel. The rms noise of each source is given as the averaged background rms of the corresponding island. Note that, in order to have enough galaxies in each pixel on average, we produce these maps with resolution $N_{\text{side}} = 256$.

(ii) An estimate of the distribution of true fluxes $p(S_{\text{true}})$. We obtain this from the SKADS semi-empirical simulation (Wilman et al. 2008). We queried the SKADS data base to retrieve the flux distribution at 151 MHz, which we scaled to 144 MHz assuming a spectral index $\alpha = -0.7$.

We use this method to produce a mean density map for our sample. Since we only consider sources detected above 5σ , the threshold flux is $S_{\text{thr}} = \max(5\sigma_N, S_{\text{cut}})$, where S_{cut} is the flux cut defining the sample. Fig. 2 shows, in the top right panel, the resulting map for our fiducial $S_{\text{cut}} = 2$ mJy. The sample is highly homogeneous, with a

mean density $\bar{n} = 0.12 \text{ arcmin}^{-2}$ that varies by ~ 1.4 per cent across the footprint (1 standard deviation).

4.1.3 Noise variations

The number of detected sources is likely to depend strongly on the depth of a given sky region. As a tracer for those depth variations, we map the co-added inverse background noise variance. For each pointing, we estimate its noise rms σ_N as the standard deviation of the low-resolution residual map made available with the DR1. The final inverse noise variance map is generated by adding the values of σ_N^{-2} for all pointings overlapping on each pixel. It is worth noting that the procedure used to generate this map does not reflect exactly the mosaicing process used in LoTSS. Additionally, there are probably fluctuations of the flux density scale between pointings that have not been quantified. It is expected that these shortcomings will improve with future DRs. This map is shown in middle left panel of Fig. 2, and will be used in the next section to study the impact of this systematic on the power spectrum.

4.2 Maps and power spectra

Before computing the power spectra, we first generate a map of the overdensity field $\delta_g(\hat{n})$ as:

$$\delta_g(\hat{n}) = \frac{N(\hat{n})}{\bar{N}w_g(\hat{n})} - 1, \quad (19)$$

where $N(\hat{n})$ is the number of galaxies observed in the pixel with position \hat{n} , and $w_g(\hat{n})$ is a weight map given by the product of the sky mask described in Section 4.1.1, and the mean density map described in Section 4.1.2. $w_g(\hat{n})$ therefore quantifies the fraction of galaxies passing our cuts at a given pixel that should have been observed. The quantity \bar{N} is the mean number of objects per pixel in the sample, estimated as $\bar{N} = \langle N(\hat{n}) \rangle_\theta / \langle w_g(\hat{n}) \rangle_\theta$, where $\langle \cdot \rangle_\theta$ denotes an average over all pixels in the map. Finally, to avoid using heavily masked pixels, we set to zero all pixels in $\delta_g(\hat{n})$ and $w_g(\hat{n})$ where $w_g(\hat{n}) < 0.5$.

Once we have the two maps $\delta_g(\hat{n})$ and $\kappa(\hat{n})$ (both shown in Fig. 2, and their associated masks $w_g(\hat{n})$ and $w_\kappa(\hat{n})$ (the latter described in Section 3.2), we compute the auto-correlation of δ_g (C_ℓ^{gg}), and its cross-correlation with κ ($C_\ell^{g\kappa}$) using the so-called *pseudo- C_ℓ* estimator (Peebles 1973; Hivon et al. 2002) using the NAMASTER code³ (Alonso et al. 2019). The pseudo- C_ℓ algorithm uses analytical methods to calculate the coupling of different multipoles ℓ due to the incomplete sky coverage. Although the method is in principle less optimal than minimum variance quadratic estimators, it is able to achieve almost equivalent uncertainties for sufficiently flat power spectra (as is the case here) with a much higher computational speed. We bin both power spectra in bands of $\Delta\ell = 50$ starting at $\ell = 2$.

The estimated overdensity map contains shot noise due to the discrete nature of galaxy number counts. The associated contribution to the angular power spectrum (commonly called the ‘noise bias’), can be estimated analytically as described in Alonso et al. (2019). We will assess the validity of this calculation in Section 5.1.

We also use NAMASTER to estimate the covariance matrix of the power spectra following García-García, Alonso & Bellini (2019). This is calculated as the so-called *Gaussian covariance*, which approximates both κ and δ_g as Gaussian random fields. The covariance matrix estimator implemented in NAMASTER accurately accounts for

the correlation between different ℓ bins induced by the incomplete sky coverage. Since the estimator requires a best-guess estimate of the underlying power spectra (C_ℓ^{gg} , $C_\ell^{g\kappa}$, and $C_\ell^{\kappa\kappa}$), the covariance is estimated in two steps: first, we compute theoretical power spectra for cosmological parameters fixed to the best-fitting values found by *Planck* (Planck Collaboration VI 2018a) and a constant galaxy bias $b_g = 1.3$ assuming the LoTSS redshift distribution, which provides a good visual fit to the data. The resulting covariance is then used in the likelihood described in Section 4.3 to find the best-fitting parameters. These are used to estimate new theory power spectra that are then used by NAMASTER to estimate our final covariance matrix.⁴ Note that the auto-spectra C_ℓ^{gg} and $C_\ell^{\kappa\kappa}$ should contain both the signal and noise contributions. For C_ℓ^{gg} , we use the shot noise component described above, while for $C_\ell^{\kappa\kappa}$ we use the noise curves provided in the *Planck* DR.

It is worth noting that, although the Gaussian covariance described above assumes Gaussian statistics for both δ_g and κ , which is known to be inaccurate due to the non-linear growth of structure, the covariance estimated using this method accounts for the largest fraction of the statistical uncertainty (Barreira, Krause & Schmidt 2018; Nicola et al. 2020), and is an excellent approximation (better than 10 per cent) in the range of scales studied here (Nicola et al. 2020).

4.3 Likelihood

In order to extract information from the measured C_ℓ^{gg} and $C_\ell^{g\kappa}$ we use a Gaussian likelihood of the form:

$$\chi^2 \equiv -2 \log p(\mathbf{d}|\mathbf{q}) = (\mathbf{d} - \mathbf{t}(\mathbf{q}))^T \text{Cov}^{-1}(\mathbf{d} - \mathbf{t}(\mathbf{q})), \quad (20)$$

where the data vector \mathbf{d} denotes all measured power spectra and $\mathbf{t}(\mathbf{q})$ is the theoretical prediction for \mathbf{d} given a set of parameters \mathbf{q} .

We will present results for different choices of data vector and parameters. Specifically, we will present constraints based on C_ℓ^{gg} and $C_\ell^{g\kappa}$ alone, as well as from the combination of both. We will also explore different combinations of three free parameters:

- (i) The galaxy bias b_g (within the two redshift evolution models described in Section 2.2).
- (ii) z_{tail} , which quantifies the extent of the redshift distribution tail when parametrized according to equation (16).
- (iii) The amplitude of matter fluctuations as parametrized by σ_8 .

We fix all cosmological parameters (including σ_8 when not used as a free parameter) to the best-fitting values found by *Planck* (Planck Collaboration VI 2018a).

We will only use the multipoles ℓ smaller than $\ell_{\text{max}} = 500$, corresponding to a wavenumber $k_{\text{max}} \simeq 0.15 \text{ Mpc}^{-1}$ at $z \simeq 1$. Thus we make sure that we only use modes where a linear, scale-dependent bias relation is a good approximation, and where the non-Gaussian contributions to the covariance matrix can be neglected.

5 RESULTS

5.1 Power spectra and covariances

The power spectra C_ℓ^{gg} and $C_\ell^{g\kappa}$, measured using the methods described in Section 4.2, are shown as black points with error bars in the upper and lower panels of Fig. 3. The correlation matrix

⁴The differences between the two covariance matrices were very small, and had a negligible effect on the final results presented here.

³<https://github.com/LSSTDESC/NaMaster>




Figure 3. Galaxy auto-correlation C_ℓ^{gg} (top panel) and its cross-correlation with the CMB lensing convergence $C_\ell^{g\kappa}$ (bottom panel). The measured power spectra are shown as black points with error bars. The solid red curve shows the best-fitting theory prediction assuming the SKADS redshift distribution, the *Planck* cosmological parameters, and a galaxy bias that grows inversely with the linear growth factor (equation 15), with an amplitude $b_g = 2.1$. The noise bias due to shot noise in the auto-correlation is shown as a grey line in the top panel. The grey points and error bars in the bottom panel show a null cross-correlation calculated by correlating the galaxy overdensity map with a randomly rotated CMB convergence map.




Figure 4. Correlation matrix associated with the covariance matrix used in this analysis. The covariance was calculated analytically using the methods in García-García et al. (2019). The auto- and cross-correlation are ~ 50 per cent correlated with each other on the largest scales.

$r_{ij} = \text{Cov}_{ij} / \sqrt{\text{Cov}_{ii} \text{Cov}_{jj}}$ associated with the covariance matrix of the full data vector is shown in Fig. 4. The lowest multipoles of C_ℓ^{gg} and $C_\ell^{g\kappa}$ are ~ 40 per cent correlated.

Before using these measurements to extract parameter constraints, we perform a number of sanity checks and null tests. The main




Figure 5. Difference between the fiducial auto- and cross-spectra (top and bottom panels, respectively) and two alternative measurements as a fraction of the 1σ uncertainties. The red points show the effect of using coarser pixels ($N_{\text{side}} = 256$). The green points show the result of deprojecting the fluctuations in the mean density map and the noise variations map (described in Sections 4.1.2 and 4.1.3) from the galaxy overdensity map, as proxies for potential contaminants of the observed δ_g .

concern regarding C_ℓ^{gg} is the presence of residual systematics in the galaxy overdensity producing extra power on large scales (Ross et al. 2012; Leistedt et al. 2013, 2016). In our case, the most likely source of such systematics is fluctuations in survey completeness caused by variations in survey properties. As a proxy for these residual systematics, we use the maps of pointing noise variation and the fluctuations in the mean density map itself. We then compare the power spectrum \hat{C}_ℓ^{gg} estimated from the original overdensity map and from a ‘systematics-deprojected’ map, in which these systematic templates are projected out at the map level (see Alonso et al. 2019, for details). We also verify that the power spectra are stable against the choice of pixelization by recomputing them with a resolution parameter $N_{\text{side}} = 256$. The result of these two tests is displayed in Fig. 5. The figure shows the relative difference between the fiducial and alternative power spectra as a fraction of the 1σ uncertainties. The power spectra are robust to the choice of pixelization and to the potential contamination from variations in survey completeness within the range of scales explored here.

In order to validate the cross-correlation $\hat{C}_\ell^{g\kappa}$, and to test for a potential mis-estimate of the statistical uncertainties, we perform one further null test by cross correlating the overdensity map δ_g with a random convergence map generated by arbitrarily rotating the original *Planck* map. The measured cross correlation is shown by the grey points in the lower panel of Fig. 3. Fitting a constant amplitude to these measurements, we find the best-fitting value $A_{\text{null}} = (-2.9 \pm 4.1) \times 10^{-8}$, consistent with zero.

The galaxy auto-correlation \hat{C}_ℓ^{gg} could still be affected by unknown systematics that are not well described by the two template maps used above. These systematics would typically affect the largest scales, comparable with the size of a pointing ($\ell \sim \pi/\delta\theta \sim 50$), and therefore we can test for the relevance of unknown contaminants by removing those scales from the analysis. Another source of systematic uncertainty that could bias our results is a mis-estimate of the noise bias in \hat{C}_ℓ^{gg} . This noise bias is shown as a grey solid line in the top panel of Fig. 3, and was calculated analytically. An error in this estimate would mostly affect the smallest scales, which are

dominated by shot noise, potentially biasing the inferred parameters. To study the impact of these two systematics, we constrain the galaxy bias parameter b_g from the auto-correlation \hat{C}_ℓ^{gg} alone in three cases:

- (i) A fiducial case, using all the data and the fiducial Gaussian covariance matrix. We find a bias value $b_g = 2.10 \pm 0.10$.
- (ii) We remove the first bandpower, covering scales $\ell < 50$. The inferred bias is $b_g = 1.98 \pm 0.12$.
- (iii) Instead of analytically computing the noise bias and subtracting it from the data, we marginalize over a constant offset to \hat{C}_ℓ^{gg} with a free amplitude parameter. This is done analytically by simply modifying the inverse covariance as (Rybicki & Press 1992):

$$\text{Cov}^{-1} \rightarrow \text{Cov}^{-1} - \frac{(\text{Cov}^{-1}\mathbf{n})(\text{Cov}^{-1}\mathbf{n})^T}{\mathbf{n}^T \text{Cov}^{-1}\mathbf{n}}, \quad (21)$$

where \mathbf{n} is our fiducial estimate of the noise bias. The corresponding constraint is $b_g = 2.23 \pm 0.16$.

In all cases, we parametrize the redshift dependence of $b_g(z)$ using the ‘constant amplitude’ model, and we assume the SKADS redshift distribution. Given the restricted range of scales used in this analysis, marginalizing over the noise bias leads to a 60 per cent increase in the final uncertainties. Nevertheless, the inferred values of b_g are compatible with each other in all cases, and therefore we conclude that the measured spectra are robust to unknown sky contaminants and uncertainties in the shot noise level.

5.2 Cross-correlation significance

The measured cross-correlation \hat{C}_ℓ^{gk} is shown in the lower panel of Fig. 3. We can quantify the significance of this detection in two different ways. First, the probability to exceed associated with the χ^2 of the measurement of \hat{C}_ℓ^{gk} with respect to the null hypothesis ($\mathbf{t} = 0$) is $p(> \chi^2) = 8 \times 10^{-5}$, corresponding to a $\sim 4\sigma$ detection.

Secondly, we fit the measured cross-correlation to a single parameter model of the form

$$\hat{C}_\ell^{gk} = b_g C_\ell^{gk}|_{b=1}, \quad (22)$$

where $C_\ell^{gk}|_{b=1}$ is the predicted signal assuming the *Planck* best-fitting cosmological parameters, the SKADS redshift distribution and a unit galaxy bias parameter $b_g = 1$ in the constant-amplitude redshift dependent model. The constraint on the free amplitude b_g is $b_g = 1.59 \pm 0.31$, corresponding to a 5.2σ measurement of the cross-correlation. Note that this is mathematically equivalent to quantifying the significance as the square root of the difference in χ^2 between the null hypothesis and this best-fitting model $\sqrt{\Delta\chi^2}$.

5.3 Galaxy bias and high-redshift tails

As described in Section 3.3, we have at our disposal at least three different estimates of the redshift distribution for our sample, which differ mostly in terms of the size of their high-redshift tails. It is therefore important to understand the level to which the galaxy bias constraints depend on the redshift distribution uncertainties. Assuming the *Planck* cosmological parameters and a constant-amplitude redshift-dependent model for the galaxy bias, the constraints on b_g assuming these three different redshift distributions, and using both \hat{C}_ℓ^{gg} and \hat{C}_ℓ^{gk} are:

$$\begin{aligned} b_g &= 1.12 \pm 0.05, \quad (\text{LoTSS-VAC}), \\ b_g &= 1.84 \pm 0.09, \quad (\text{VLA} - \text{COSMOS}), \\ b_g &= 2.10 \pm 0.10, \quad (\text{SKADS}). \end{aligned} \quad (23)$$

While the VLA-COSMOS and SKADS estimates are roughly compatible with each other, they differ from the LoTSS-VAC result by more than 7σ . The differences between the different redshift distribution estimates are therefore highly significant in this context. In order to investigate this further, we repeat the exercise using only the CMB lensing cross-correlation. In this case, the constraints are fully compatible with each other (albeit with larger error bars):

$$\begin{aligned} b_g &= 1.35 \pm 0.25, \quad (\text{LoTSS}), \\ b_g &= 1.46 \pm 0.28, \quad (\text{VLA} - \text{COSMOS}), \\ b_g &= 1.59 \pm 0.31, \quad (\text{SKADS}). \end{aligned} \quad (24)$$

The reason for the disagreement between the different inferred bias values with the full data vector is the well-known fact that both the galaxy bias and the width of the redshift distribution have a degenerate effect on the amplitude of the galaxy auto-correlation. Structure is washed out in samples with broader redshift distributions, which can be compensated with a larger b_g .

To understand the role played by both parameters in the auto- and cross-correlation, Fig. 6 shows the measured power spectra together with the theoretical predictions for varying values of the galaxy bias b_g (left-hand panel) and the high-redshift tail z_{tail} (right-hand panel). While the galaxy bias affects the amplitude of the auto- and cross-correlations (quadratically and linearly, respectively), an increasing high-redshift tail lowers the auto-correlation while leaving the cross-correlation almost constant. We can understand the latter result as follows: since the CMB lensing kernel W_k extends to very high redshifts, a variation in the width of the galaxy redshift distribution leaves the overlap between W_g and W_k almost unchanged. On the one hand, this makes the cross-correlation between CMB lensing and galaxy clustering robust to uncertainties in the redshift distribution width. On the other hand, the different response of the auto- and cross-correlation to z_{tail} should allow us to break its degeneracy with b_g , allowing us to simultaneously constrain both parameters by combining \hat{C}_ℓ^{gg} and \hat{C}_ℓ^{gk} .

To explore this, we parametrize the redshift distribution according to equation (16) and derive constraints on two free parameters, b_g and z_{tail} , with flat priors $b_g \in (0.6, 6)$ and $z_{\text{tail}} \in (0.1, 5)$. The resulting constraints are shown in Fig. 7 for the constant-amplitude (orange contours) and constant-bias (black-grey contours) redshift-dependent models of $b_g(z)$. Looking at the constant-amplitude constraints, the uncertainty on b_g ($\sigma(b_g) = 0.28$) degrades significantly with respect to the previous results where the redshift distribution was fixed to the SKADS estimate. The high-redshift tail, on the other hand, is constrained to be $z_{\text{tail}} = 1.30^{+0.27}_{-0.40}$. In the constant-bias case, a larger value of b_g is preferred, to compensate for the lack of growth as a function of redshift, which is accompanied by a preference for larger values of z_{tail} and overall larger uncertainties on both parameters. In both cases, however, the data show a preference for larger redshift tails than that implied by the photometric redshifts included in the LoTSS VAC. This can be seen explicitly in Fig. 8, which shows the 1σ bounds on the recovered redshift distribution in comparison with the estimates from the LoTSS VAC, VLA-COSMOS, and SKADS.

The corresponding constraints on $b_g(z)$ are shown in Fig. 9 in both cases. The true redshift evolution of the effective galaxy bias for the LoTSS sample analysed here is most likely not constant but potentially less steep than the constant-amplitude model, therefore lying somewhere between these two extremes. To illustrate this, the figure also shows the bias values measured for different radio populations by Hale et al. (2018), and the NVSS sample of Nusser & Tiwari (2015). The results above therefore imply that the true underlying

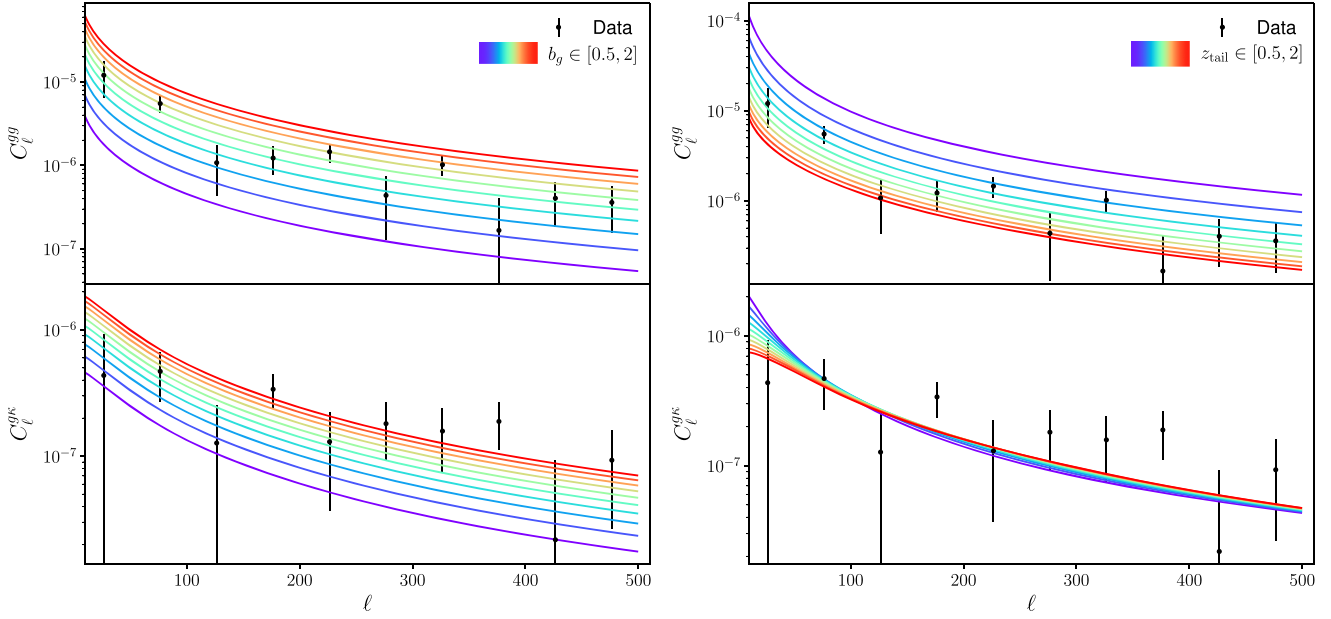


Figure 6. Measured auto- and cross-correlation (black dots with error bars in the top and bottom panels, respectively), together with the theory prediction for different values of the galaxy bias b_g and the high-redshift tail z_{tail} (left- and right-hand panels, respectively), both in the range (0.5, 2.0). A ‘constant amplitude’ model is assumed for the redshift evolution of the galaxy bias. b_g is fixed to 1.3 in the right-hand panel, while $z_{\text{tail}} = 1.1$ in the left one. While both b_g and z_{tail} affect the amplitude of the auto-correlation, the cross-correlation depends only mildly on the high-redshift tail, making it possible to break the degeneracy between both parameters by combining C_ℓ^{gg} and C_ℓ^{gk} .




Figure 7. Constraints on galaxy bias b_g and the high-redshift tail parameter z_{tail} from the combination of C_ℓ^{gg} and C_ℓ^{gk} . Results are shown for a constant galaxy bias as a function of redshift, the ‘constant bias’ model (black), and for a bias that scales with the inverse of the linear growth factor, the ‘constant amplitude’ model (orange).

distribution is probably more compatible with the SKADS or VLA-COSMOS estimates than the distribution of photometric redshifts in the LoTSS VAC.

5.4 Constraints on bias and σ_8

Under the assumption that the true underlying redshift distribution of the LoTSS sample is well described by the SKADS estimate, we can use the combination of \hat{C}_ℓ^{gg} and \hat{C}_ℓ^{gk} to break the degeneracy between b_g and the amplitude of matter fluctuations parametrized




Figure 8. 1σ constraints on the redshift distribution of the LoTSS flux-limited sample obtained from the combination of C_ℓ^{gg} and C_ℓ^{gk} . Results are shown for a constant bias model (grey region), and for a bias that grows with the inverse of the linear growth factor (orange region). The redshift distributions inferred from the photometric redshifts in the LoTSS VAC, from the 3 GHz VLA-COSMOS catalogue, and from the SKADS simulations are shown in red, green, and blue, respectively, for comparison. All redshift distributions are normalized to have the same maximum amplitude. The constraints show a preference for longer redshift tails than the one predicted by the LoTSS photo-zs.

by σ_8 . Given the existing redshift distribution uncertainties, and the degeneracy with other cosmological parameters, the resulting constraints would be neither robust nor competitive. The aim of this exercise is therefore twofold: a sanity check to verify that our data




Figure 9. 1σ constraints on the galaxy bias $b_g(z)$ from the combination of C_ℓ^{gg} and C_ℓ^{gk} . Results are shown for a constant bias model (grey region), and for a bias that grows with the inverse of the linear growth factor (orange region). The figure also shows measurements of the bias for different radio galaxy samples in Hale et al. (2018, black symbols with error bars), and for the NVSS sample of Nusser & Tiwari (2015).




Figure 10. Constraints on the galaxy bias parameter b_g and the amplitude of matter inhomogeneities σ_8 from the combination of C_ℓ^{gg} and C_ℓ^{gk} . Results are shown for a constant-bias model (black contours) and for a bias that evolves with the inverse of the linear growth factor (orange contours). The blue band shows the constraints on σ_8 found by Planck Collaboration VI (2018a). Our power spectrum measurements are in good agreement with the standard Λ CDM model as constrained by Planck.

is broadly compatible with the standard cosmological model, and a demonstration of the use of continuum surveys for cosmology.

The results are shown in Fig. 10 for the constant-bias and constant-amplitude redshift-dependent models. The constraints on σ_8 for both

bias models are

$$\sigma_8 = 0.69^{+0.14}_{-0.21}, \quad (\text{constant amplitude}),$$

$$\sigma_8 = 0.79^{+0.17}_{-0.32}, \quad (\text{constant bias}). \quad (25)$$

Both values are in agreement with each other as well as with those found by Planck ($\sigma_8 = 0.811 \pm 0.006$, shown in blue in the figure), with a preference for lower σ_8 values along the b_g - σ_8 degeneracy direction.

6 CONCLUSIONS

Radio continuum surveys are a valuable probe to study the properties and evolution of SFGs and AGNs. Unimpeded by dust attenuation, radio surveys cover a much broader range of redshifts, and correspondingly larger volumes, than their optical counterparts. Because of this, continuum surveys also constitute a potential avenue to reconstruct the growth of structure over the largest scales, and therefore have drawn the interest of the cosmology community.

However, the lack of redshift information from radio continuum data makes these surveys reliant on matching to deep optical catalogues to constrain the radial distribution of the sources, and to reconstruct the redshift evolution of their properties. The potential incompleteness of the optical cross-matches makes redshift calibration one of the largest sources of systematic uncertainty in the potential use of radio continuum surveys for cosmology. This is akin to the challenges faced by ongoing and future photometric weak lensing surveys when characterizing and propagating the uncertainties in the redshift distribution of galaxies (Sánchez et al. 2020; Hildebrandt et al. 2020; Schaaf, Ferraro & Seljak 2020).

In this paper, we have studied the clustering of galaxies on the flux-density limited LoTSS DR1, both through the harmonic-space auto-correlation, and through their cross-correlation with the lensing convergence of the CMB, as measured by Planck. The cross-correlation is detected at the 5σ level.

To illustrate the challenge posed by the uncertainties in the redshift distribution dp/dz , we have considered three estimates of this quantity: from the LoTSS VAC, from the VLA-COSMOS cross-matched sample (which includes both spectroscopic and photometric redshifts), and from the SKADS simulations. We have shown that the truncated high-redshift tail from the LoTSS-VAC photometric redshifts leads to radically different interpretations of the clustering amplitude when compared with the results from the VLA-COSMOS or SKADS redshift distributions. On the other hand, the CMB lensing cross-correlation alone is fairly insensitive to variations in dp/dz , and leads to consistent measurements of the galaxy bias.

The robustness of the CMB-lensing cross-correlation to redshift distribution uncertainties, allows us to break the degeneracy between the galaxy bias and the width of the dp/dz by combining it with the clustering auto-correlation. Through this joint analysis, we are able place constraints on the high-redshift tail of the distribution, showing that it is underestimated by the LoTSS-VAC photo-zs, and better represented by the VLA-COSMOS and SKADS estimates. To our knowledge, this is the first attempt in the literature at calibrating redshift distributions through cross-correlations with CMB lensing. This is akin to the ‘cross-correlation redshifts’ approach, based on using cross-correlations against samples with a known redshift distribution (Newman 2008; Alonso et al. 2017; Gatti et al. 2018, in this case, the CMB lensing kernel). Given the broad range of redshifts covered by the CMB lensing kernel, it is unlikely that this approach will be able to constrain a shift in the mean redshift of the target distribution (e.g. see Alonso et al. 2017), as is usually done in weak

lensing analyses (Hildebrandt et al. 2020), but it should be useful to calibrate the width of the distribution or the presence of photo- z outliers. This could provide a robust way to extract cosmological information from samples with poor spectroscopic coverage in future surveys, such as the Legacy Survey of Space and Time (LSST; The LSST Dark Energy Science Collaboration et al. 2018).

Finally, assuming that the redshift distribution is well described by the SKADS estimate (which is broadly consistent with the VLA-COSMOS redshift distribution), we use the combination of clustering and CMB lensing to measure the amplitude of matter inhomogeneities σ_8 . Although the resulting constraints are neither competitive nor robust, given the existing uncertainties on $d\mu/dz$ and on the redshift dependence of the galaxy bias, this exercise allows us to demonstrate the use of continuum surveys for cosmology. The measured value of $\sigma_8 = 0.79^{+0.17}_{-0.32}$, assuming a constant bias, agrees well with the constraints found by CMB and large-scale structure experiments (Planck Collaboration VI 2018a; Abbott et al. 2018; Hikage et al. 2019; Asgari et al. 2020).

This paper highlights the novel science that can be carried out by combining CMB and radio-continuum survey using the DR1 from LoTSS, which covers 424 deg². LoTSS itself should cover the vast majority of the northern sky when completed ($\sim 20\,000$ deg², allowing much stronger constraints on the various results presented here. Furthermore, the EMU Survey (Norris et al. 2011) will provide a complementary Southern hemisphere survey, down to a similar equivalent flux-density limit at GHz frequencies. However, as we have shown in this work, one of the key uncertainties in using radio continuum surveys for cosmology arises from the lack of a well-constrained redshift distribution. Fortunately, there are a wealth of surveys that will aid in pinning this down on the time-scale of these all-sky surveys. Future LoTSS DRs will be able to make use of deep fields with ~ 90 per cent photometric redshift coverage. Furthermore, the WEAVE-LOFAR Survey (Smith et al. 2016) will provide a huge number of spectroscopic redshifts for the LoTSS sources, combining deep spectroscopy in the LOFAR deep survey fields, with spectroscopy of the rarer brighter sources, through a tiered radio-continuum selected survey. However, spectroscopic surveys are rarely complete, and tend to be limited by the optical magnitude depth or the ability to detect emission lines (an advantage for radio continuum selected surveys). As such, further important information will be gleaned from measuring the redshift distribution of radio sources in the deep extragalactic survey fields, similar to how we have used the VLA-COSMOS survey here. However, a single deep field is inherently limited by sample variance, as such the MIGHTEE (Jarvis et al. 2016) Survey, which will cover ~ 20 deg² over four extragalactic deep fields, accessible from the Southern hemisphere, will help pin down the redshift distribution of radio continuum sources to very deep flux-density limits using both photometric (e.g. Jarvis et al. 2013; Adams et al. 2020) and spectroscopic (e.g. Davies et al. 2018; Cirasuolo et al. 2014; Driver et al. 2019) redshifts.

Taken together these advances will allow us to begin to divide the radio continuum source population into distinct subsamples, for example, SFGs and AGN. This would then allow us to introduce more realistic bias evolution models for these different populations, with different redshift distributions, which is a limitation on our current work due to the relatively small survey area of the LoTSS DR1.

Looking further in the future, the sky surveys planned for the SKA will lead to higher number density of objects across huge swathes of sky. Although this may not substantially increase the number of AGN at high redshift ($z > 2$), these surveys will increase the number of SFGs across all redshifts, and allow the possibility of a more robust decoupling of the SFG and AGN populations (e.g. Makhathini

et al. 2015; Muxlow et al. 2020), given the much higher angular resolution (< 1 arcsec). This could also be achieved by LOFAR using its international baselines.

ACKNOWLEDGEMENTS

We would like to thank Emmanuel Schaan for useful comments and discussions. DA acknowledges support from the Beecroft Trust, and from the Science and Technology Facilities Council through an Ernest Rutherford Fellowship, grant reference ST/P004474. EB is supported by the European Research Council grant no.: 693024 and the Beecroft Trust. MJJ acknowledges support from the UK Science and Technology Facilities Council [ST/N000919/1], the South African Radio Astronomy Observatory (SARAO; www.ska.ac.za) and the Oxford Hintze Centre for Astrophysical Surveys which is funded through generous support from the Hintze Family Charitable Foundation. DJS acknowledges the Research Training Group 1620 ‘Models of Gravity’, supported by Deutsche Forschungsgemeinschaft (DFG) and support by the German Federal Ministry for Science and Research BMBF-Verbundforschungsprojekt D-LOFAR IV (grant no. 05A17PBA). This publication arises from research funded by the John Fell OUP Research Fund. LOFAR data products were provided by the LOFAR Surveys Key Science project (LSKSP⁵) and were derived from observations with the International LOFAR Telescope (ILT). LOFAR (van Haarlem et al. 2013) is the Low Frequency Array designed and constructed by ASTRON. It has observing, data processing, and data storage facilities in several countries, that are owned by various parties (each with their own funding sources), and that are collectively operated by the ILT foundation under a joint scientific policy. The efforts of the LSKSP have benefited from funding from the European Research Council, NOVA, NWO, INSU, CNRS, the SURF Co-operative, the UK Science and Technology Funding Council, and the Jülich Supercomputing Centre. This work uses data based on observations obtained with *Planck* (<http://www.esa.int/Planck>), an ESA science mission with instruments and contributions directly funded by ESA Member States, NASA, and Canada. Contour plots were generated using the GETDIST package (Lewis 2019). Some of the results in this paper have been derived using the HEALPIX (Górski et al. 2005) package. We made extensive use of the SCIPY (Virtanen et al. 2020), HEALPY (Zonca et al. 2019), and MATPLOTLIB (Hunter 2007) PYTHON packages.

DATA AVAILABILITY STATEMENT

The data underlying this article are available in the *Planck* legacy archive <https://www.cosmos.esa.int/web/planck/pla>, the LOFAR surveys website <https://lofar-surveys.org/releases.html>, and the VLA-COSMOS 3 GHz repository at the Strasbourg astronomical Data Center <http://cdsarc.u-strasbg.fr/viz-bin/qcat?J/A+A/602/A2>.

REFERENCES

- Abbott T. M. C. et al., 2018, *Phys. Rev. D*, 98, 043526
- Adams N. J., Bowler R. A. A., Jarvis M. J., Häußer B., McLure R. J., Bunker A., Dunlop J. S., Verma A., 2020, *MNRAS*, 494, 1771
- Allison R. et al., 2015, *MNRAS*, 451, 849
- Alonso D., Ferreira P. G., 2015, *Phys. Rev. D*, 92, 063525
- Alonso D., Ferreira P. G., Jarvis M. J., Moodley K., 2017, *Phys. Rev. D*, 96, 043515

⁵<https://lofar-surveys.org/>

- Alonso D., Sanchez J., Slosar A., LSST Dark Energy Science Collaboration, 2019, *MNRAS*, 484, 4127
- Antonucci R., 1993, *ARA&A*, 31, 473
- Asgari M. et al., 2020, preprint ([arXiv:2007.15633](https://arxiv.org/abs/2007.15633))
- Barreira A., Krause E., Schmidt F., 2018, *J. Cosmol. Astropart. Phys.*, 2018, 053
- Bell E. F., 2003, *ApJ*, 586, 794
- Bengaly C. A. P., Maartens R., Randriamiarinarivo N. r., Baloyi A., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 025
- Best P. N., Heckman T. M., 2012, *MNRAS*, 421, 1569
- Chisari N. E. et al., 2019, *ApJS*, 242, 2
- Cirasuolo M. et al., 2014, in Ramsay S. K., McLean I. S., Takami H., eds, Proc. SPIE Conf. Ser. Vol. 9147, Ground-based and Airborne Instrumentation for Astronomy V, SPIE, Bellingham, p. 91470N
- Condon J. J., 1992, *ARA&A*, 30, 575
- Condon J. J., Cotton W. D., Greisen E. W., Yin Q. F., Perley R. A., Taylor G. B., Broderick J. J., 1998, *AJ*, 115, 1693
- Davies L. J. M. et al., 2017, *MNRAS*, 466, 2312
- Davies L. J. M. et al., 2018, *MNRAS*, 480, 768
- Driver S. P. et al., 2019, *The Messenger*, 175, 46
- Duncan K. J. et al., 2019, *A&A*, 622, A3
- Fanaroff B. L., Riley J. M., 1974, *MNRAS*, 167, 31P
- Ferramacho L. D., Santos M. G., Jarvis M. J., Camera S., 2014, *MNRAS*, 442, 2511
- García-García C., Alonso D., Bellini E., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 043
- Gatti M. et al., 2018, *MNRAS*, 477, 1664
- Gomes Z., Camera S., Jarvis M. J., Hale C., Fonseca J., 2020, *MNRAS*, 492, 1513
- Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, 622, 759
- Gürkan G. et al., 2018, *MNRAS*, 475, 3010
- Hale C., Jarvis M., Delvecchio I., Hatfield P., Novak M., Smolcic V., Zamorani G., 2018, *MNRAS*, 474, 4133
- Heckman T. M., Best P. N., 2014, *ARA&A*, 52, 589
- Hikage C. et al., 2019, *PASJ*, 71, 43
- Hildebrandt H. et al., 2020, preprint ([arXiv:2007.15635](https://arxiv.org/abs/2007.15635))
- Hivon E., Górski K. M., Netterfield C. B., Crill B. P., Prunet S., Hansen F., 2002, *ApJ*, 567, 2
- Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
- Intema H. T., Jagannathan P., Mooley K. P., Frail D. A., 2017, *A&A*, 598, A78
- Jarvis M. J. et al., 2013, *MNRAS*, 428, 1281
- Jarvis M. et al., 2016, MeerKAT Science: On the Pathway to the SKA, p. 6. preprint ([arXiv:1709.01901](https://arxiv.org/abs/1709.01901))
- Jarvis M., Bacon D., Blake C., Brown M., Lindsay S., Raccanelli A., Santos M., Schwarz D. J., 2015, Proc. Sci., Cosmology with SKA Radio Continuum Surveys, SISSA, Trieste. p. PoS(AASKA14)018, preprint ([arXiv:1501.03825](https://arxiv.org/abs/1501.03825))
- Johnston S. et al., 2007, *PASA*, 24, 174
- Jonas J. L., 2009, *IEEE Proc.*, 97, 1522
- Leistedt B. et al., 2016, *ApJS*, 226, 24
- Leistedt B., Peiris H. V., Mortlock D. J., Benoit-Lévy A., Pontzen A., 2013, *MNRAS*, 435, 1857
- Lewis A., 2019, preprint ([arXiv:1910.13970](https://arxiv.org/abs/1910.13970))
- Lewis A., Challinor A., Lasenby A., 2000, *ApJ*, 538, 473
- Lindsay S. N. et al., 2014a, *MNRAS*, 440, 1527
- Lindsay S. N., Jarvis M. J., McAlpine K., 2014b, *MNRAS*, 440, 2322
- Makhathini S., Jarvis M., Smirnov O., Heywood I., 2015, Proc. Sci., Morphological classification of radio sources for galaxy evolution and cosmology with the SKA, SISSA, Trieste. p. PoS(AASKA14)081, preprint ([arXiv:1412.5990](https://arxiv.org/abs/1412.5990))
- Mauch T., Murphy T., Buttery H. J., Curran J., Hunstead R. W., Piestrzynski B., Robertson J. G., Sadler E. M., 2003, *MNRAS*, 342, 1117
- Muxlow T. W. B. et al., 2020, *MNRAS*, 495, 1188
- Namikawa T., Yamauchi D., Sherwin B., Nagata R., 2016, *Phys. Rev. D*, 93, 043527
- Newman J. A., 2008, *ApJ*, 684, 88
- Nicola A. et al., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 044
- Norris R. P. et al., 2011, *PASA*, 28, 215
- Nusser A., Tiwari P., 2015, *ApJ*, 812, 85
- Peebles P. J. E., 1973, *ApJ*, 185, 413
- Planck CollaborationXVII, 2014, *A&A*, 571, A17
- Planck Collaboration et al., 2018a, *A&A*, 641, A6
- Planck Collaboration et al., 2018b, *A&A*, 641, A8
- Raccanelli A. et al., 2012, *MNRAS*, 424, 801
- Ross A. J. et al., 2012, *MNRAS*, 424, 564
- Rybicki G. B., Press W. H., 1992, *ApJ*, 398, 169
- Sánchez C., Raveri M., Alarcon A., Bernstein G. M., 2020, *MNRAS*, 498, 2984
- Schaan E., Ferraro S., Seljak U., 2020, *J. Cosmol. Astropart. Phys.*, 12, 001
- Shimwell T. W. et al., 2019, *A&A*, 622, A1
- Siewert T. M. et al., 2019, *A&A*, 643, A100
- Smith D. J. B. et al., 2016, in Reylé C., Richard J., Cambrésy L., Deleuil M., Pécontal E., Tresse L., Vauglin I., eds, SF2A-2016: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics. p. 271, preprint ([arXiv:1611.02706](https://arxiv.org/abs/1611.02706))
- Smolčić V. et al., 2017, *A&A*, 602, A1
- Takahashi R., Sato M., Nishimichi T., Taruya A., Oguri M., 2012, *ApJ*, 761, 152
- The LSST Dark Energy Science Collaboration et al., 2018, preprint ([arXiv:1809.01669](https://arxiv.org/abs/1809.01669))
- Urry C. M., Padovani P., 1995, *PASP*, 107, 803
- van Haarlem M. P. et al., 2013, *A&A*, 556, A2
- Virtanen P. et al., 2020, *Nat. Methods*, 17, 261
- Williams W. L. et al., 2019, *A&A*, 622, A2
- Wilman R. J. et al., 2008, *MNRAS*, 388, 1335
- Wright A. H., Hildebrandt H., van den Busch J. L., Heymans C., Joachimi B., Kannawadi A., Kuijken K., 2020, *A&A*, 640, L14
- Zonca A., Singer L., Lenz D., Reinecke M., Rosset C., Hivon E., Gorski K., 2019, *J. Open Source Softw.*, 4, 1298

This paper has been typeset from a \LaTeX file prepared by the author.