

A measurement of the integrated Sachs–Wolfe effect with the Rapid ASKAP Continuum Survey

Benedict Bahr-Kalus ¹★, David Parkinson ^{1,2}★, Jacobo Asorey ³, Stefano Camera ^{4,5,6},
Catherine Hale ⁷ and Fei Qin (秦斐) ¹

¹*Korea Astronomy and Space Science Institute, Yuseong-gu, Daedeok-daero 776, Daejeon 34055, Korea*

²*University of Science and Technology, Daejeon 34113, Korea*

³*Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Av. Complutense, 40, D-28040 Madrid, Spain*

⁴*Dipartimento di Fisica, Università degli Studi di Torino, Via P. Giuria 1, I-10125 Torino, Italy*

⁵*INFN – Istituto Nazionale di Fisica Nucleare, Sezione di Torino, Via P. Giuria 1, I-10125 Torino, Italy*

⁶*Department of Physics and Astronomy, University of the Western Cape, 7535 Cape Town, South Africa*

⁷*School of Physics and Astronomy, University of Edinburgh, Royal Observatory, Blackford Hill, Edinburgh EH9 3HJ, UK*

Accepted 2022 July 19. Received 2022 July 15; in original form 2022 May 3

ABSTRACT

The evolution of the gravitational potentials on large scales due to the accelerated expansion of the Universe is an important and independent probe of dark energy, known as the integrated Sachs–Wolfe (ISW) effect. We measure this ISW effect through cross-correlating the cosmic microwave background maps from the *Planck* satellite with a radio continuum galaxy distribution map from the recent Rapid ASKAP Continuum Survey (RACS). We detect a positive cross-correlation at $\sim 2.8\sigma$ relative to the null hypothesis of no correlation. We parametrize the strength of the ISW effect through an amplitude parameter and find the constraints to be $A_{\text{ISW}} = 0.94^{+0.42}_{-0.41}$, which is consistent with the prediction of an accelerating universe within the current concordance cosmological model, Λ CDM. The credible interval on this parameter is independent of the different bias models and redshift distributions that were considered when marginalizing over the nuisance parameters. We also detect a power excess in the galaxy autocorrelation angular power spectrum on large scales ($\ell \leq 40$), and investigate possible systematic causes.

Key words: dark energy – large-scale structure of Universe – radio continuum: galaxies.

1 INTRODUCTION

The mysterious acceleration of the expansion of the Universe, generated by the so-named dark energy, is now an established part of the concordance cosmological model, Λ CDM. The observational evidence comes not only from standard-candle and standard-ruler measurements of the expansion history but also from observations of the large-scale structure of matter and the distribution of the gravitational potential.

An accelerating expansion will act against gravitational in-fall, slowing the accretion rate and decreasing the growth rate of cosmic structures. These structures, and their evolution in time, are observed through tracer particles. For high-redshift observations, the tracers are the photons emitted at the surface of last scattering, which form the cosmic microwave background (CMB) and trace the density fluctuations at recombination through the anisotropies in the intensity (i.e. temperature) and polarization maps. In the CMB temperature power spectrum, the large-scale anisotropy is generated by the Sachs–Wolfe effect (SW; Sachs & Wolfe 1967) at last scattering, a gravitational redshift effect from photons climbing (or

falling) out of the gravitational potential to enter the homogeneous universe.

There is also a secondary effect generated long after recombination, caused by further evolution of the gravitational potentials, which is known as the integrated Sachs–Wolfe (ISW) effect. This late-time evolution of the potentials is driven by the accelerating universe, as the redshifting and blueshifting of photons moving into and out of density fields no longer exactly balances, but leaves some energy imprint in the photon frequencies. This process is an independent probe of the dark energy but it is difficult to see the effect on the CMB power spectrum alone. However, since the photon energies become correlated with the matter distribution at late times, the effect can be seen in the correlation between these two tracer fields (Crittenden & Turok 1996).

The ISW effect was first detected in cross-correlation using NVSS 1.4-GHz radio catalogue (Condon et al. 1998) and the HEAO1 A2 full-sky hard X-ray map (Boldt 1987) for large-scale structure tracers, and all-sky CMB map from the Wilkinson Microwave Anisotropy Probe (*WMAP*; Bennett et al. 2003), with a combined detection significance of 2.5σ (Boughn & Crittenden 2004, 2005). The statistical significance of this ISW detection with the NVSS sample was reassessed by Raccanelli et al. (2008), which examined the consistency of the modelled bias-weighted redshift distribution with the data, giving an adjusted 3σ detection.

* E-mail: benedictbahrkalus@kasi.re.kr (BBK); davidparkinson@kasi.re.kr (DP)

The ISW effect has also been detected using optical and infrared galaxies, cross-correlating *WMAP* with galaxy samples extracted from the Automated Plate Measurement survey (APM; Fosalba, Gaztanaga & Castander 2003; Fosalba & Gaztanaga 2004), the Sloan Digital Sky Survey (SDSS; Cabre et al. 2006), the 2MASS sample (Dupé et al. 2011), and *WISE* galaxies (Ferraro, Sherwin & Spergel 2015). Recent work has updated the CMB maps from the *Planck* mission (Planck Collaboration XXI 2016), and detected the ISW in cross-correlation at 4σ , again using the NVSS catalogue, as well as optical galaxies from the SDSS, infrared galaxies from the *WISE* survey, and the *Planck* 2015 convergence lensing map, as the low-redshift mass tracers. Most recently, it has been detected using the DR8 galaxy catalogue of the DESI Legacy imaging surveys (Dong et al. 2021), using a ‘low-density position’ filter, with a significance of 3.2σ .

However, all of these detections are at a relatively low significance and have not added so much to the total constraining power of a cosmological data compilation. The next generation of surveys, like those proposed for the Australian Square Kilometre Array Pathfinder (ASKAP; Johnston & Wall 2008; Hotan et al. 2021) and the SKA Observatory,¹ will detect objects down to a lower surface brightness, and this increase in number counts should in turn increase the significance of the ISW detection, as well as the utility of the measurement. As the number count is increased, the sample can be split into redshift bins, which would make such a sample more sensitive to the long-wavelength radial power that generates the signal, and allow it to be used for more than a simple detection of the dark energy (see Camera et al. 2012; Ballardini & Maartens 2019). In Raccanelli et al. (2015), the authors forecast the effectiveness of such a future sample in determining the amplitude of the non-Gaussian contribution to the primordial density fluctuation. They found it to increase the effectiveness of *Planck* for such, and be competitive with an all-sky optical survey such as that proposed for the *Euclid* satellite (see also Alonso & Ferreira 2015; Camera et al. 2015). Similar forecasts have been made for the effectiveness of measuring primordial non-Gaussianity using the multitracer technique (Yamauchi, Takahashi & Oguri 2014; Fonseca, Maartens & Santos 2017; Gomes et al. 2020), showing a predicted improvement over the constraints from *Planck* alone.

In this work, we present our analysis of the cross-correlation of the CMB maps from the *Planck* mission with a new radio continuum data set from the ‘band 1’ sample (Hale et al. 2021) of the first data release of the Rapid ASKAP Continuum Survey (RACS; McConnell et al. 2020). RACS is a large-area radio continuum survey, covering the sky south of $+41^\circ$ Dec. It is comparable to NVSS in depth, size of catalogue and area covered. It is different from NVSS in two key aspects. Firstly, it covers southern regions un-surveyed by NVSS. Secondly, whilst observations and re-observations were taken over 2019–2020, the total on-source time was only a few weeks (see McConnell et al. 2020). RACS demonstrates the impressive survey power of ASKAP and provides an opportunity to test the cosmological analysis methods for the Evolutionary Map of the Universe (EMU) survey (Norris et al. 2011, 2021).

In Section 2, we review the theoretical basis for the ISW effect. In Section 3, we describe our data sample and the methods and tools we use to analyse it. In Section 4, we give our results, and in Section 5, we summarize our findings.

¹<https://www.skatelescope.org>.

2 THEORY

The angular power spectrum of a set of tracers X (e.g. galaxies, or photons) can be measured from the overdensity field $\delta_X(\boldsymbol{\theta})$ (where $\boldsymbol{\theta}$ is a particular direction on the sky)

$$a_{\ell m}^X = \int d^2\boldsymbol{\theta} Y_{\ell m}^* \delta_X(\boldsymbol{\theta}). \quad (1)$$

Note that this is valid for a continuous density field. For a discrete density field, the integral is replaced with a sum.

Assuming an isotropic universe, we get the power spectrum from the autocorrelation

$$\langle a_{\ell m}^X a_{\ell' m'}^{*X} \rangle = \delta_{\ell\ell'}^K \delta_{mm'}^K C_\ell^{XX}, \quad (2)$$

where δ^K is the Kronecker symbol, and C_ℓ^{XX} is the angular autopower spectrum of a map of tracer X with itself. We can make a prediction for the angular power spectrum of a particular tracer using the three-dimensional power spectrum $P(k)$. Here the power needs to be averaged or ‘smoothed’ in the radial direction, and the theoretical prediction is given by

$$C_\ell^{XX} = \frac{2}{\pi} \int dk k^2 P(k) [W_\ell^X(k)]^2, \quad (3)$$

where $W_\ell(k)$ is the window function for the tracer X .

The galaxy window function (at linear order) is given by (e.g. Giannantonio et al. 2008; Raccanelli et al. 2008)

$$W_\ell^g(k) = \int dz n(z) b(z) D(z) j_\ell[kr(z)], \quad (4)$$

where $n(z) dz$ is the source distribution per steradian with redshift z within dz (brighter than some survey magnitude or flux limit), $b(z)$ is the linear bias factor relating tracer overdensity to matter overdensity, $D(z)$ is the growth factor of density perturbations, j_ℓ is the spherical Bessel function of order ℓ , and $r(z) = \eta_0 - \eta(z)$ is the radial comoving distance to redshift z , with $\eta(z)$ being the conformal time coordinate at redshift z .

The cross-correlation power spectrum between a density field of large-scale structure tracers at low redshift and the CMB temperature fluctuations is given by

$$C_\ell^{gT} := \langle a_{\ell m}^g a_{\ell m}^{*T} \rangle \quad (5)$$

$$= \frac{2}{\pi} \int dk k^2 P(k) W_\ell^g(k) W_\ell^T(k), \quad (6)$$

where we now have two different window functions: $W_\ell^g(k)$ for the large-scale structure tracer at low redshift, and $W_\ell^T(k)$ for the CMB photons. The window function for the CMB photons has a different structure to equation (4), as it is the power that is induced in the CMB temperature from the ISW effect, given by the equation

$$\left(\frac{\Delta T}{T}\right)_{\text{ISW}}(\boldsymbol{x}_0, \boldsymbol{\theta}) = 2 \int_{\eta_{\text{dec}}}^{\eta_0} d\eta \dot{\Upsilon}[\boldsymbol{x}_0 - \boldsymbol{\theta}(\eta - \eta_0), \eta], \quad (7)$$

where $\dot{\Upsilon}$ is the time-derivative of the lensing potential (i.e. the Weyl potential) $\Upsilon = (\Phi + \Psi)/2$, with respect to conformal time η . Here, \boldsymbol{x}_0 is the observer’s position (the photon position at time η_0), and $\boldsymbol{\theta}$ is the photon position at some general time.

Assuming no anisotropic growth, i.e. $\Phi = \Psi = \Upsilon$, the lensing potential obeys the field equation

$$\ddot{\Upsilon} + 3\mathcal{H}\dot{\Upsilon} + (2\dot{\mathcal{H}} + \mathcal{H}^2) = 4\pi G a^2 (\delta p), \quad (8)$$

where \mathcal{H} is the conformal-time Hubble–Lemaître rate and (δp) denotes the first-order perturbation on top of homogeneous and

isotropic pressure. Solving the Friedmann equations for a matter-dominated universe, one gets $\mathcal{H} = 2/\eta$ and, thus, $2\dot{\mathcal{H}} + \mathcal{H}^2 = 0$. As, on cosmic scales, matter is a pressureless fluid, i.e. $(\delta p) = 0$, equation (8) simplifies to

$$\ddot{\Upsilon} + \frac{6}{\eta} \dot{\Upsilon} = 0. \quad (9)$$

The solution of this equation has the form

$$\Upsilon = \aleph + \beth \eta^{-5}. \quad (10)$$

Now, unless \aleph is fine-tuned to be vanishingly small, we already have $\aleph \gg \beth \eta^{-5}$ at the epoch when the CMB photons are released and Υ is effectively constant during matter domination. Hence, as with equation (7), the CMB photons retain the integrated history of the gravitational evolution of the Universe, a non-vanishing $(\Delta T/T)_{\text{ISW}}$ proves that the Universe has undergone epochs where the cosmic fluid was not primarily composed of baryonic or dark matter ($\Omega_m \neq 1$). In the concordance model of cosmology, these epochs are the radiation-dominated epoch at early times when the CMB was released, and at late times, our current epoch, which is dominated by dark energy. As we are studying the ISW effect in cross-correlations between CMB anisotropies and the matter density field at relatively low redshifts $z \lesssim 5.2$, our analysis will establish evidence for or against the existence of dark energy.

To decrease the noise of the measured power spectra, we bin in multipole bins of width $\Delta\ell = 20$. We obtain the binned model power spectrum

$$C_\ell^{XY, \text{binned}} = \frac{\sum_{\ell' \in \ell\text{-bin}} \ell'(\ell' + 1) C_{\ell'}^{XY}}{\sum_{\ell' \in \ell\text{-bin}} \ell'(\ell' + 1)} \quad (11)$$

as the weighted average of the unbinned C_ℓ^{XY} , where the $\ell'(\ell' + 1)$ -weights are proportional to the variance, in turn, minimizing the variance on $C_\ell^{XY, \text{binned}}$ compared to a $(2\ell' + 1)$ -weighting scheme that corresponds to the number of modes entering each multipole ℓ . Since we apply the same weights to both the data and the models used to infer covariances and the significance of our findings, our conclusions are unaffected by the choice of weighting scheme.

Finally, a computation of the theoretical power can be increased in speed by making the Limber approximation (Limber 1953),

$$j_\ell[kr(z)] \xrightarrow{\ell \gg 1} \sqrt{\frac{\pi}{2\ell + 1}} \delta^D\left(\ell + \frac{1}{2} - kr(z)\right), \quad (12)$$

which approximates the full window function calculation and convolution to a simple distance integral, with δ^D the Dirac distribution. This approximation breaks down when we integrate over more angular than radial modes. Hence, applying the Limber approximation at multipoles ℓ below some ℓ_{min} can lead to catastrophic biases in the cosmological parameters of interest, e.g. as illustrated by Bernal et al. (2020) and proven by Martinelli et al. (2022) with a realistic analysis of a synthetic data set. However, in this instance, we are saved by not being able to locate radio continuum galaxies in redshift, thus, radial modes dominate even low multipoles and, hence, ℓ_{min} becomes a function of the width of the redshift bin. Tanidis, Camera & Parkinson (2020) have estimated $\ell_{\text{min}} = 2$ for a one-redshift-bin EMU-like survey. We are going to confirm the validity of the Limber approximation for our purposes in the following section before using it in cosmological analyses.

3 DATA ANALYSIS

In this section, we describe the input data catalogues that we use, as well as the angular selection functions and the estimators that we employ to measure the angular power spectrum.

3.1 Radio Data

The radio data used in this work is from RACS (McConnell et al. 2020; Hale et al. 2021), an ASKAP survey that aims to observe the entire Southern sky (Dec. $\lesssim +41^\circ$) using a rapid survey strategy in three frequency bands over the 700–1800 MHz range. Each frequency band will use a bandwidth of 288 MHz for the observations. The first such data release, McConnell et al. (2020), comprises images covering the Southern sky at Dec. $\lesssim +41^\circ$ and centred at a frequency of 888 MHz, using 15-min on-source observations. This is the lowest frequency band that will be observed with RACS. As part of the associated data release with McConnell et al. (2020), images and catalogues were released covering 903 pointings, each with varying angular resolution across the sky.

For this work, it is essential to have a single catalogue across the sky without any duplication. Therefore we used the catalogue released within the second RACS paper (Hale et al. 2021), which we shall briefly discuss. In Hale et al. (2021), the images of McConnell et al. (2020) were convolved to a common resolution of 25 arcsec and mosaicked together to produce a contiguous image across the majority of the sky covered by RACS. Convolution of the image to a common resolution was essential to retain flux scale across the images before mosaicking. This resulted in 799 pointings that had sufficient resolution to be convolved to 25 arcsec and hence mosaicked together. The missing regions compared to McConnell et al. (2020) were concentrated in the Dec. $=+30^\circ$ to $+40^\circ$ and -90° to -80° regimes. After mosaicking, sources were detected by running the source extraction software PYBDSF (Mohan & Rafferty 2015) over each of the 799 tiles using a 5σ criterion. The catalogues from the 799 tiles were then combined to avoid duplication, and to remove the Galactic plane, namely Galactic latitude between -5° and $+5^\circ$. The raw RACS overdensity field from Hale et al. (2021) is mapped in Fig. 1.

3.1.1 Radio data weighting function

Despite the radio data catalogue from Hale et al. (2021) having uniform resolution across the sky, it is not uniformly sensitive across the images. This is due to a variety of factors: bright sources in the field affecting the neighbouring image, hour angle coverage differing with observations and the amount of overlap in mosaicking with neighbouring tiles. We, therefore, use the completeness simulations from Hale et al. (2021, using resolved sources) to determine the detection fraction of sources within each HEALPIX bin.

The simulations from Hale et al. (2021) use simulated sources from Wilman et al. (2008), Wilman et al. (2010) and inject sources into the residual images and re-extract the sources using PYBDSF. These simulations use 5 million random sources across Dec. $=-85^\circ$ to $+30^\circ$ and each simulation is repeated 10 times. We combine all the recovered sources (which have a output ‘measured’ flux that would have resulted in a 5σ detection) within a HEALPIX bin and compare this to the number of sources within the HEALPIX bin that were injected to determine the weight within a given HEALPIX bin. A map of the radio data weights is shown in the top panel of Fig. 1.

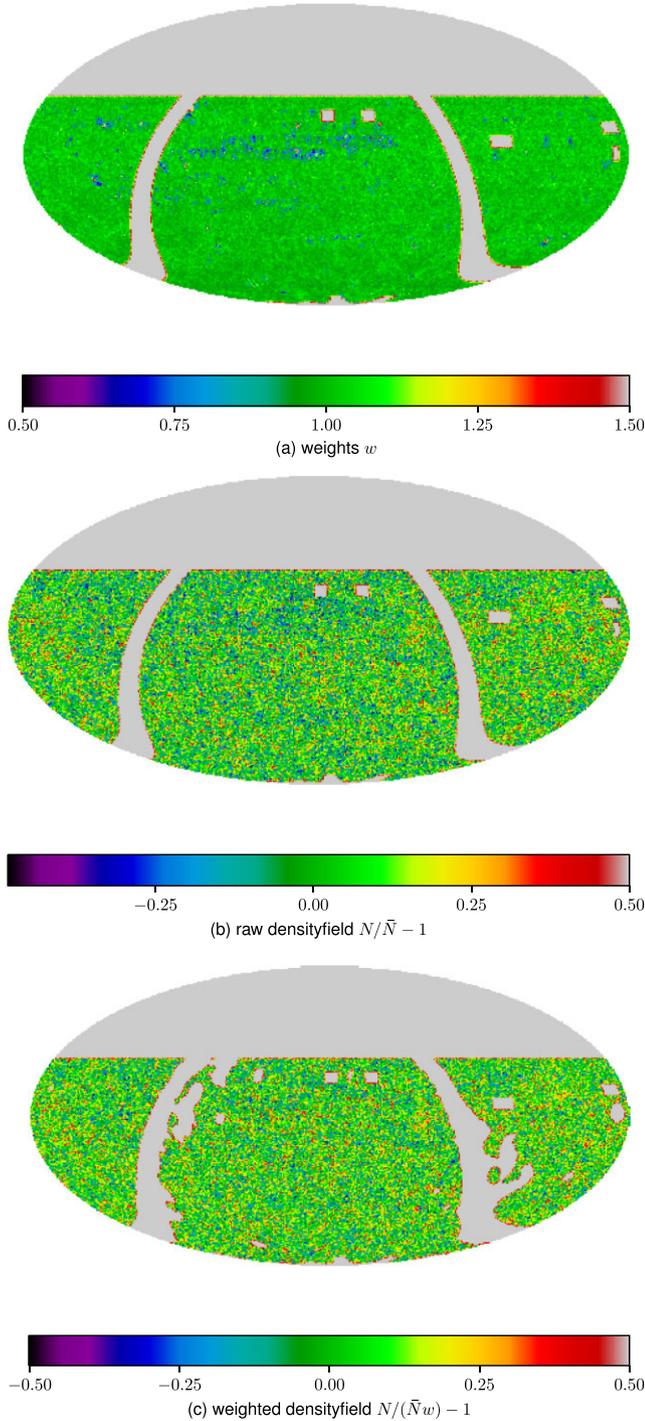


Figure 1. The distribution of the weights (top panel), normalized number counts of objects above a flux density threshold of 4 mJy (middle panel), and weighted overdensity field (bottom panel) of RACS radio continuum galaxies (still with a 4-mJy flux density limit) on the sky after masking. All maps are in Mollweide projection and equatorial coordinates in astronomical orientation, i.e. showing east on the left-hand side.

3.1.2 The radio dipole

The angular two-point statistics of the NVSS catalogue initially showed an excess at large scales that could have been interpreted as the signal due to the scale-dependent bias effect due to a

non-Gaussian distribution of the primordial density field (Xia et al. 2010). This excess signal has disappeared after Chen & Schwarz (2016) re-analysed the NVSS catalogue using a new mask taking sidelobe effects of bright sources, the Galactic foreground and the radio dipole signal into account. The radio dipole is believed to result from our peculiar motion whose velocity vector (in natural units) is written as \mathbf{v}_{pec} .

While the first two issues are addressed by the weighting function described in Section 3.1.1, the latter modulates the observed density field (Ellis & Baldwin 1984). Assuming that the flux density S at a given frequency ν and the number count $\tilde{N}(> S)$ are both given by power laws with, respectively, spectral indices α and x , i.e.

$$S \propto \nu^{-\alpha}, \quad (13)$$

$$\tilde{N}(> S) \propto S^{-x}, \quad (14)$$

the observed density field $\delta_{\text{g}}^{\text{obs}}$ in direction $\boldsymbol{\theta}$ is given by (e.g. Bengaly et al. 2019)

$$\delta_{\text{g}}^{\text{obs}} = \delta_{\text{g}}^{\text{rest}} + [2 + x(1 + \alpha)] \boldsymbol{\theta} \cdot \mathbf{v}_{\text{pec}}, \quad (15)$$

where $\delta_{\text{g}}^{\text{rest}}$ represents the overdensity field in the rest frame where the galaxy distribution is statistically isotropic.

We coincidentally estimate $\alpha = x = 0.76$ from SKADS (which we adopt from here on), as well as $\alpha = 0.82$ and $x = 0.90$ from T-RECS. While these simulation-based estimate might not provide us with the most accurate measurement of α and x (e.g. SKADS underestimates source counts at faint flux densities; see e.g. Smolčić et al. 2017; Norris et al. 2021; Gürkan et al. 2022) they are consistent with observations at higher flux densities and with that of Hale et al. (2021) above ~ 2 mJy. Measurements of α are commonly measured from radio surveys and assumed in studies within the literature to be ~ 0.7 – 0.8 (see e.g. Smolčić et al. 2017; De Gasperin, Intema & Fraix 2018; Norris et al. 2021), though we note that Hale et al. (2021) found slightly larger/smaller values dependent on the frequency being compared to. The scatter between the SKADS and T-RECS results is also dwarfed by the scatter among different measurements of the amplitude of \mathbf{v}_{pec} from radio surveys, such as Blake & Wall (2002), Singal (2011), Gibelyou & Huterer (2012), Rubart & Schwarz (2013), Tiwari et al. (2014), Tiwari & Jain (2015),

Tiwari & Nusser (2016), and Siewert, Schmidt-Rubart & Schwarz (2021), even find an apparent frequency dependence of the radio dipole amplitude. Given this uncertainty in the amplitude and the fact that all of them agree in direction with the CMB dipole, a more natural assumption of \mathbf{v}_{pec} when subtracting the second term of equation (15) is the CMB dipole measured by *Planck* (Aghanim et al. 2020a). By doing so, we are also consistent with the CMB data that we describe in the next subsection, from which the CMB dipole has been subtracted.

3.2 CMB data

We make use of the third release SMICA *Planck* Legacy Map (European Space Agency 2018). SMICA (Delabrouille, Cardoso & Patanchon 2003; Cardoso et al. 2008) stands for Spectral Matching Independent Component Analysis and is one of the four component separation methods used by the Planck Collaboration. The SMICA data model

$$\mathbf{R}_{\ell} = \mathbf{a} \mathbf{a}^{\dagger} C_{\ell}^{\text{TT}} + \mathbf{A} \mathbf{P}_{\ell} \mathbf{A}^{\dagger} + N_{\ell} \quad (16)$$

is a superposition of the true CMB signal (expressed in terms of the matrix \mathbf{a} composed of $a_{\ell m}^{\text{T}}$ in each frequency band and their

frequency-independent autopower spectrum C_ℓ^{TT}), the noise spectrum N_ℓ , and foreground signals $\mathbf{A} \mathbf{P}_\ell \mathbf{A}^\dagger$. The foreground signals are expressed in terms of a small number of templates with arbitrary frequency spectra, power spectra, and component correlations. These are fitted to the auto- and cross-power spectra of *Planck* maps $\mathbf{x}_{\ell m}$ in its nine frequency channels. The final SMICA map,

$$\hat{\mathbf{s}}_{\ell m} = \mathbf{w}_\ell^\dagger \mathbf{x}_{\ell m}, \quad (17)$$

is then obtained by fitting weights (note that these are unrelated to the weights in equation 11)

$$\mathbf{w}_\ell = \frac{\mathbf{R}_\ell^{-1} \mathbf{a}}{\mathbf{a}^\dagger \mathbf{R}_\ell^{-1} \mathbf{a}} \quad (18)$$

that minimize the discrepancy between the frequency channel map auto- and cross-power spectra, i.e.

$$\hat{\mathbf{w}}_\ell = \arg \min_{\mathbf{w}_\ell} \sum_\ell \left(\sum_m \mathbf{x}_{\ell m} \mathbf{x}_{\ell m}^\dagger \mathbf{R}_\ell + (2\ell + 1) \ln \det \mathbf{R}_\ell \right). \quad (19)$$

The fit is done in three steps:

- (i) Only the CMB power spectrum C_ℓ and \mathbf{a} are fitted on a clean patch of the sky;
- (ii) All other parameters are fitted on a large patch of the sky while keeping \mathbf{a} fixed at the best-fitting value of the previous step;
- (iii) \mathbf{a} and \mathbf{A} are fixed to their previously found values while the power spectra C_ℓ and \mathbf{P}_ℓ are fitted.

SMICA is the foreground component separation method that has performed best in a *Planck* foreground-cleaning mock challenge (Ade et al. 2014). However, we have found that the choice of component separation method has no significant impact on the galaxy–temperature cross-correlation and, thus, on the ISW signal.

The temperature map can be retrieved as the `I_STOKES` column from the FITS file downloadable from the digital object identifier given in the reference of European Space Agency (2018). We rotate and downgrade the resolution of the *Planck* map from its initial $N_{\text{side}} = 2048$ in galactic coordinates to match RACS’s $N_{\text{side}} = 128$ in equatorial coordinates. We perform the same transformations to the temperature confidence mask given in the `TMASK` column and we cut out pixels from the RACS map where the value of the temperature confidence is less than 0.5. Equally, we mask out CMB pixels that are also masked out by the RACS mask. We show the binary mask outlining the quality cuts imposed on the RACS and *Planck* data in grey in the bottom panel of Fig. 1.

3.3 Estimating the angular power spectra

The estimation of the spherical harmonic amplitudes, and the angular power spectrum, as given in equations (1) and (2), assumes that the full sky is available. For a cut-sky, as we have with both the CMB and radio continuum data, we need to apply an angular selection function (as described in Sections 3.1.1 and 3.2) and estimate from only those regions that are visible. This leads to measured amplitudes $\tilde{a}_{\ell m}$ s that are different from the true values, and a pseudo angular power spectrum \tilde{C}_ℓ , as computed by the MASTER algorithm (Hivon et al. 2002). The advantage of the MASTER algorithm is that the measured \tilde{C}_ℓ can then be directly compared to the theoretical prediction. In this work, we use the python implementation of the algorithm, `NAMASTER` (Alonso et al. 2019).

Following the approach of Alonso et al. (2021), we first generate a map of the radio continuum overdensity field, which we do by

Table 1. Fiducial cosmological parameters assumed throughout this paper.

Parameter	Symbol	Value/relationship
Hubble–Lemaître constant	H_0	$67.5 \text{ km s}^{-1} \text{ Mpc}^{-1}$
Reduced Hubble–Lemaître constant	h	$H_0/(100 \text{ km s}^{-1} \text{ Mpc}^{-1})$
Physical baryon density parameter	Ω_b	$0.022/h^2$
Cold dark matter density parameter	Ω_{cdm}	$0.12/h^2$
Total matter abundance	Ω_m	$\Omega_{\text{cdm}} + \Omega_b$
Dark energy density parameter	Ω_Λ	$1 - \Omega_m$
Reionization optical depth	τ	0.0544
Amplitude of scalar fluctuations	A_s	2×10^{-9}
Scalar spectral index	n_s	0.965

combining the galaxy number count map $N(\boldsymbol{\theta})$ with the radio data weighting function map $w(\boldsymbol{\theta})$ from Section 3.1, using the equation

$$\delta_g(\boldsymbol{\theta}) = \frac{N(\boldsymbol{\theta})}{\bar{N} w(\boldsymbol{\theta})} - 1, \quad (20)$$

where $\boldsymbol{\theta}$ is a particular direction (or `HEALPIX` pixel) on the sky and \bar{N} is the average weighted number of galaxies per `HEALPIX` cell. To construct the overdensity map, we cut all those pixels $\boldsymbol{\theta}$ that have weights $w(\boldsymbol{\theta}) < 0.5$, to prevent a bias. However, these are only a very small number that still lie inside the region selected in the angular window. The overdensity field is shown in Fig. 1.

As galaxies are discrete objects sampling the continuous density field, the pseudo galaxy autopower spectrum $\tilde{C}_\ell^{\text{gg}}$ will disagree with the model power spectrum C_ℓ^{gg} by a constant shot-noise term N_{shot} . Naïvely, one can think of the galaxies being drawn from the matter field as a Poisson point process. In spite of that, some galaxies appear as multiple sources in a radio catalogue, whereas in other instances, multiple sources may not be identified as such by the source finder. Consequently, the shot-noise level can deviate from its Poisson prediction. The source finding can be approximated as a supplementary Poisson sampling from the already Poisson sampled galaxy number count, resulting in a so-called compound Poisson distributed sample (Siewert et al. 2020). In any case, the compound Poisson distribution also predicts a scale-independent shot-noise power spectrum and, instead of modelling it, we fit a constant \hat{N}_{shot} that minimizes $(\tilde{C}_\ell^{\text{gg}} - \hat{N}_{\text{shot}} - C_\ell^{\text{gg, fid}})^2$ for the hereinafter defined fiducial power spectrum $C_\ell^{\text{gg, fid}}$.

3.4 Theoretical predictions and modelling

3.4.1 Cosmological parameters

To model theoretically the power spectra that we want to compare our data against, we assume a flat, homogeneous and isotropic universe where the laws of gravity are expressed by the theory of general relativity. As we cannot faithfully measure all cosmological parameters from RACS alone, we fix the parameters listed in Table 1 at the reported values. These are for the most part the default values of the ‘Code for Anisotropies in the Microwave Background’ (CAMB; Lewis, Challinor & Lasenby 2000; Howlett et al. 2012), with the exception of n_s and τ that we take from Aghanim et al. (2020b) for consistency with the *Planck* 2018 maps (European Space Agency 2018). Note that parameters that are not matched to *Planck* 2018 are within 1σ from the *Planck* 2018 best-fitting values.

To check the validity of the Limber approximation (cf. Section 2), we evaluate equation (3) twice, once with and once without making use of the Limber approximation, for the same fiducial cosmology, bias and redshift distribution models. Binning the result in multipole

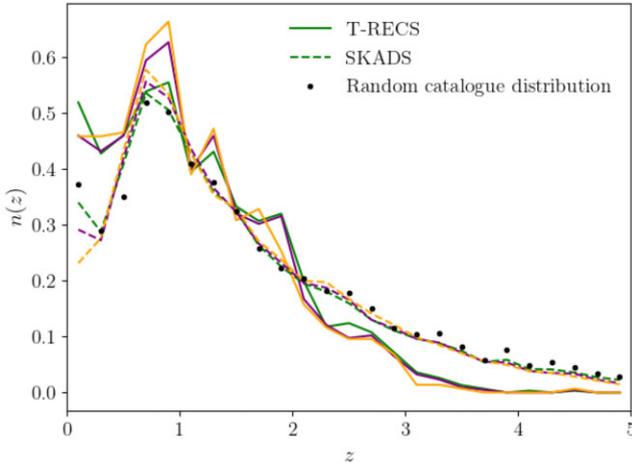


Figure 2. The distribution of radio continuum galaxies with redshift, as predicted by the SKADS (dashed) and T-RECS (solid line) mock radio continuum galaxy catalogues. The different colours correspond to 2- (green), 3- (purple), and 4-mJy (orange) flux density limits. The black dots are the distribution used in the generation of random catalogues, as described in Section 3.1.1.

bins with width $\Delta\ell = 20$ (as we are going to do in our analyses), we find a bias of ~ 1 per cent in the C_ℓ^{gg} prediction for the lowest ℓ -bin and much smaller biases at smaller scales. We revisit this assumption later when we evaluate the likelihood of our data.

As discussed in Section 2, the existence of dark energy causes a correlation between the CMB temperature map and the distribution of matter due to the late-time ISW effect. On the other hand, in a Universe without significant dark energy that would be dominated by matter until our present epoch, we should not measure a notable cross-correlation between the two fields. We, therefore, introduce a phenomenological parameter A_{ISW} , such that

$$C_\ell^{\text{gT}} = A_{\text{ISW}} C_\ell^{\text{gT, fid}}, \quad (21)$$

where $C_\ell^{\text{gT, fid}}$ is the galaxy–temperature cross-power spectrum computed for the fiducial parameters listed in Table 1. In this way, if we measure an A_{ISW} that is consistent with zero, we have not detected the ISW effect and, thus, we have found no evidence for dark energy. Should A_{ISW} , however, be consistent with unity, then our Λ CDM-based model of the galaxy–temperature cross-power spectrum is consistent with the data. If $A_{\text{ISW}} > 0$ but inconsistent with one, then we still have detected dark energy but we have to revisit our modelling assumptions.

3.4.2 Number count model

To make accurate predictions for the angular power spectrum of a galaxy sample, the window function needs to be computed using some well-motivated estimate for the redshift distribution of galaxy number per steradian $n(z)$ and bias $b(z)$. For our sample of radio continuum galaxies (being observed at ~ 1 GHz), we make use of simulations to inform this redshift distribution. Two of the major existing extra-galactic radio simulations that are available to use are the European SKA Design Study (SKADS) Simulated Skies (Wilman et al. 2008) and the Tiered Radio Extra-galactic Continuum Simulation (T-RECS; Bonaldi et al. 2019). In Fig. 2, we show the predicted $n(z)$ distribution for several different flux cuts from both the T-RECS and SKADS simulated catalogues, and the distribution used

in the generation of random catalogues, as described in Section 3.1.1. Although the predictions are very similar for 2, 3, and 4 mJy, we assume a value of 4 mJy for all theoretical predictions for the rest of this paper. This 4 mJy reflects a region where, above this flux density limit, we believe the random weight maps appropriately account for incompleteness within the survey, as can be seen in the source counts corrections of Hale et al. (2021).

We see that both simulations make roughly similar predictions for the redshift distribution, peaking at around $z = 1$ and slowly falling off at higher redshifts. However, the SKADS prediction has a larger high-redshift tail, with 90 per cent of galaxies lying below $z < 3.6$. In contrast, the T-RECS galaxies are more localized to $z \sim 1$, with 90 per cent lying below $z < 3.1$. This will affect the power spectrum predictions, as the window function given in equation (4) will average the radial fluctuations out over a larger range of k -values for SKADS than T-RECS, diluting the power and so leading to a lower amplitude for the same cosmology. We consider both $n(z)$ models in our analysis.

3.4.3 Bias model

Radio surveys are known to trace two galaxy populations: active galactic nuclei (AGN) and star-forming galaxies (SFGs). The peak-background split model (Bardeen et al. 1986; Cole & Kaiser 1989) predicts the simple relationship

$$\delta_{\text{g}} = b \delta_{\text{m}} \quad (22)$$

between the galaxy overdensity field δ_{g} and the matter overdensity field δ_{m} . This in turn means

$$C_\ell^{\text{gg}} = b^2 C_\ell^{\text{mm}}, \quad (23)$$

$$C_\ell^{\text{gT}} = b C_\ell^{\text{mT}}, \quad (24)$$

for the power spectra, under the assumption of a constant bias across redshift, i.e. $b(z) \equiv b$. As time progresses, more galaxies have the chance to form within haloes and evolve; thus, the galaxy bias is, in general, a redshift-dependent quantity.

For a combined sample, where the individual species of galaxies are not separated when the clustering is measured, the angular correlation function and power spectra are only sensitive to the total bias. For this total bias, we must combine the biases by weighting them with the individual number counts $n_i(z)$ of each galaxy-type population as in Ferramacho et al. (2014) (see also Bernal et al. 2019; Gomes et al. 2020; Asorey & Parkinson 2021), namely

$$b(z) = \frac{\sum_i b_i(z) n_i(z)}{n_{\text{all}}(z)}, \quad (25)$$

where i corresponds to the different populations and $n_{\text{all}}(z)$ is the whole sample redshift distribution. Then, we need some prescription for the biases of the individual populations.

At low redshifts, we have some good measurements of the bias values of each population (e.g. Magliocchetti et al. 2017; Hale et al. 2018; Dolfi et al. 2019). However, at higher redshifts the bias is a large source of uncertainty, amplified by our ignorance of what ratio of the observed population is composed of what type of radio source. For RACS, we estimate that SFGs make a considerable fraction of objects only at $z \sim 0.2$ (cf. Fig. 3).² At higher redshifts, the RACS catalogue is dominated by AGNs. Up until redshift $z \sim 1.8$, most

²Note that SKADS further subdivides this population into starburst galaxies (SBs) and true SFGs.

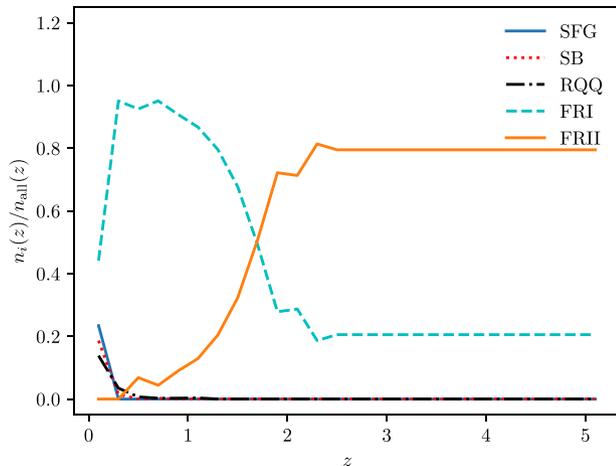


Figure 3. The predicted fraction of the observed number of SFGs, SBs, radio-quiet quasars (RQQs), AGN of type FR I or FR II in the Fanaroff–Riley classification, and the total number of observed objects estimated from SKADS (Wilman et al. 2008). This prediction assumes a 4-mJy flux cut.

of them fall into the first Fanaroff–Riley class (FR I). Above that redshift, FR IIs are the most important radio source.

The fitting that was done as part of the Wilman et al. (2008) analysis gave a parametrized form of this $b_i(z)$, and these bias models have been used extensively in forecasting the potential that radio continuum surveys have to probe cosmology (see e.g. Camera et al. 2012; Ferramacho et al. 2014; Raccanelli et al. 2015; Bernal et al. 2019; Asorey & Parkinson 2021), and are described in detail there. In these models each population has a bias that evolves exponentially with redshift. Wilman et al. (2008) argue that this leads to excessively strong clustering at high redshifts and, therefore, propose a constant bias above a certain cut-off redshift.

Instead of using theoretical models for the bias, that are based on N -body simulations, we can parameterize our ignorance, and attempt to measure the bias directly from the data. Here we consider the following effective $b(z)$ parametrizations (which we shall compare with the fiducial bias from SKADS and T-RECS in Section 4.4):

(i) As can be seen in Fig. 3, the RACS catalogue is expected to be composed mostly of AGNs. For both AGN FR subtypes, the SKADS bias model plateaus above $z > 1.5$. We therefore consider an exponential bias :

$$b(z) = \begin{cases} b(0) \exp\left(\frac{d \ln b}{d \ln z}\right) & \text{for } z < 1.5 \\ b(0) \exp\left(\frac{d \ln b}{d \ln z}\right)_{z=1.5} & \text{for } z \geq 1.5 \end{cases},$$

with an arbitrary redshift cap at $z = 1.5$, motivated by Wilman et al. (2008).

(ii) Since we find the redshift cut-off somewhat arbitrary, we also study an exponential bias $b(z) = b(0) \exp(d \ln b / d \ln z)$ that is still well motivated at the redshift range where we expect the bulk of our observed objects.

(iii) A linear bias $b(z) = b(0) + db/d \ln z$ that allows for redshift evolution without excessive clustering at the high-redshift tail.

(iv) Lastly, a constant bias b that has been assumed in forecasts at high redshift.

These models are plotted in Fig. 4. As galaxies only form in high-density regions, $b(z)$ has to be positive. We therefore impose hard

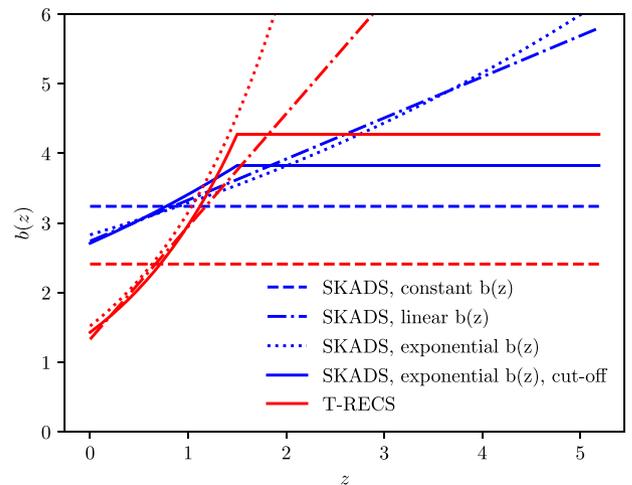


Figure 4. The bias parametrizations considered in this work evaluated at the best-fitting parameters for both SKADS- and T-RECS $n(z)$ distributions listed in Table 2.

priors $b(0) > 0$ in (i)–(iii), $db/dz > b(0)/z_{\max}$ (with $z_{\max} = 5.2$ the assumed maximum redshift attainable by the survey) in (iii), as well as $b > 0$ in (iv). Note that, since we keep $n(z)$ fixed in our analyses and $n(z)$ is degenerate with $b(z)$, our uncertainty on the bias $b(z)$ also effectively incorporates our uncertainty on $n(z)$.

3.5 Mock catalogues

We generate mock overdensity fields to test our analysis pipeline as well as to estimate the statistical significance of our measurements. To do so, we use the publicly available Full-sky Lognormal Astrofields Simulation Kit (FLASK; Xavier, Abdalla & Joachimi 2016) to draw Gaussian random fields for a given set of angular model power spectra C_ℓ^S and C_ℓ^T . We also have the option to further condition each pair of mock- galaxy and CMB maps by defining a model C_ℓ^{gT} . For the CMB maps, we just mask out regions that are not observed in the data from the Gaussian realizations of the temperature maps. For the mock RACS source maps, we first let FLASK Poisson sample the number of mock sources

$$n(\theta) = \text{Poisson} \{ \bar{n} w(\theta) [1 + \delta_g(\theta)] \} \quad (26)$$

from the Gaussian density field realizations $\delta_g(\theta)$, where we choose the average number density \bar{n} such that the total number of sources matches the number of sources in the data catalogue. We also apply the same completeness weights $w(\theta)$ and mask as for the data. Finally, the mock source count and CMB maps are saved in the HEALPIX format with $N_{\text{side}} = 128$.

3.6 Covariance matrices

We explore four different ways to estimate the covariances $\mathbf{S}_{\ell\ell'}^{WXYZ}$ between multipoles ℓ and ℓ' and fields W, X, Y , and Z . In the general case where all fields can be different, we use an analytic estimate based on a fiducial power spectrum and a mixing matrix encompassing the effect of the survey mask. We use one internal method, i.e. estimating the covariance by resampling the data, as well as two external methods where we estimate the covariance matrices from mock realizations of the data. We can use the covariance

matrices obtained in these different ways to validate them against each other.

3.6.1 Analytic covariance

Given two maps $X(\boldsymbol{\theta})$ and $Y(\boldsymbol{\theta})$, measurements \tilde{C}_ℓ^{XY} and $\tilde{C}_{\ell'}^{XY}$ of their harmonic-space cross-power spectrum at two different multipoles have covariance defined by

$$\mathbf{S}_{\ell\ell'} := \text{Cov} [\tilde{C}_\ell^{XY}, \tilde{C}_{\ell'}^{XY}] = \langle a_{\ell m}^X a_{\ell m}^{*Y} a_{\ell' m'}^Y a_{\ell' m'}^{*X} \rangle - \langle a_{\ell m}^X a_{\ell m}^{*Y} \rangle \langle a_{\ell' m'}^Y a_{\ell' m'}^{*X} \rangle. \quad (27)$$

Under the hypothesis of Gaussianity, and using Wick's theorem to break up the four-point correlator into products of two-point correlators, we find

$$\mathbf{S}_{\ell\ell'} = \frac{\tilde{C}_\ell^{XX} \tilde{C}_{\ell'}^{YY} + (\tilde{C}_\ell^{XY})^2}{(2\ell + 1) \Delta\ell} \delta_{\ell\ell'}^K. \quad (28)$$

In the case of partial sky coverage, a common approximation is to perform the rescaling $\mathbf{S}_{\ell\ell'} \rightarrow \mathbf{S}_{\ell\ell'} / f_{\text{sky}}$, where f_{sky} is the fraction of the sky observed. For $f_{\text{sky}} \lesssim 1$, this approximation performs well and has the advantage of correctly accounting for the increase in the (co)variance of the measurements due to a more limited number of available modes. However, if f_{sky} is significantly smaller than unity, or if the survey mask is highly non-trivial, or if coverage and depths change across the sky, more refined methods are needed. As mentioned in Section 3.3, one of such methods is represented by pseudo- C_ℓ s, where the coupling between different multipoles induced by the partial sky coverage is encoded in the so-called coupling matrix – in turn, related to the power spectrum of the mask/weight map. Once this quantity is given, the `NAMASTER` code allows for the evaluation of the masked covariance matrix.

3.6.2 Jackknife resampling

Internal covariance matrix estimation methods have the advantage that they are independent of any cosmological model, the survey selection is naturally accounted for, and the contribution of hidden or unforeseen systematic errors is inherent in the uncertainties estimated by internal methods. On the other hand, they rely on the assumption that the data are an accurate representation of the distribution of measurements. Sampling fluctuations known in the cosmology literature as *cosmic variance* are therefore not included in internal covariance matrix evaluations (see e.g. Norberg et al. 2009).

We make use of the ‘delete one’-jackknife method proposed by Shao (1986). We draw N_{sub} subsamples of non-adjacent non-zero `HEALPIX` cells without replacement, i.e. each `HEALPIX` cell (that is not excluded by the survey mask) is a member of exactly one subsample. We proceed by computing the angular power spectra omitting one subsample at a time. Calling the angular power spectrum obtained by omitting the i th subsample $\{C_\ell^{XY}\}_i$, we can estimate the covariance matrix as (e.g. Norberg et al. 2009)

$$\widehat{\mathbf{S}}_{\ell\ell'}^{WXYZ} = \frac{N_{\text{sub}} - 1}{N_{\text{sub}}} \sum_{i=1}^{N_{\text{sub}}} (\{C_\ell^{WX}\}_i - \bar{C}_\ell^{WX}) (\{C_{\ell'}^{YZ}\}_i - \bar{C}_{\ell'}^{YZ}), \quad (29)$$

where

$$\bar{C}_\ell^{XY} = \frac{1}{N_{\text{sub}}} \sum_{i=1}^{N_{\text{sub}}} \{C_\ell^{XY}\}_i \quad (30)$$

is the mean of the angular power spectrum over all subsamples, and the prefactor in equation (29) comes from the fact that $N_{\text{sub}} - 2$ pixel groups are the same between each pair of subsamples, thus, one has to correct the covariance matrix estimate for the correlation between each pair of $\{C_\ell^{XY}\}_i$.

3.6.3 Sample covariance of mock realizations

The first external covariance estimator is the most straight forward and most used one. Having generated N_{mock} mock realizations of the data as described in Section 3.5, one can simply compute the sample covariance as

$$\widehat{\mathbf{S}}_{\ell\ell'}^{WXYZ} = \frac{1}{N_{\text{mock}} - 1} \sum_{i=1}^{N_{\text{mock}}} (\{C_\ell^{WX}\}_i - \bar{C}_\ell^{WX}) (\{C_{\ell'}^{YZ}\}_i - \bar{C}_{\ell'}^{YZ}). \quad (31)$$

Here, one has to be aware of the fact that even though equation (31) is an unbiased estimator of the covariance matrix; this is not true for its inverse, the precision matrix $\mathbf{K}_{\ell\ell'}^{WXYZ} \equiv (\widehat{\mathbf{S}}^{WXYZ})_{\ell\ell'}^{-1}$, which is actually the crucial quantity for inference purposes. An unbiased estimator of the $p \times p$ precision matrix is given by (Kaufman 1967; Hartlap, Simon & Schneider 2007)

$$\widehat{\mathbf{K}}_{\ell\ell'}^{WXYZ} = \frac{N_{\text{mock}} - p - 2}{N_{\text{mock}} - 1} (\widehat{\mathbf{S}}^{WXYZ})_{\ell\ell'}^{-1}. \quad (32)$$

3.6.4 Covariance from mock realizations using the graphical lasso

As we are primarily interested in the precision matrix, we can also apply an estimator designed to directly find sparse precision matrices from realizations of the data. Such an estimator is the graphical lasso (Friedman, Hastie & Tibshirani 2007). The algorithm works by finding the non-negative definite matrix $\widehat{\mathbf{K}}_{\ell\ell'}^{XYXY}$ that minimizes the log-likelihood of the mock realizations. The strength of the graphical lasso is recovering the graphical structure from correlations in the data. This works better for the inverse correlation matrix $\widehat{\mathbf{R}}_{\ell\ell'}^{XYXY} \equiv \widehat{\mathbf{R}}_{\ell\ell'}^{XYXY}$ than for the precision matrix $\widehat{\mathbf{K}}_{\ell\ell'}^{XYXY}$. We get the precision matrix as $\widehat{\mathbf{K}}_{\ell\ell'}^{XY} = \widehat{\mathbf{R}}_{\ell\ell'}^{XY} / (\sigma_\ell \sigma_{\ell'})$, where $\sigma_\ell \equiv \sqrt{\langle (\{C_\ell^{XY}\}_i - \bar{C}_\ell^{XY})^2 \rangle}$ is the standard deviation of the angular power spectra estimated from the mocks. As covariance, precision and correlation matrices are usually sparse, there is also a penalty term on off-diagonal terms. The full cost function with the penalty term reads

$$-\ln \det \widehat{\mathbf{R}}^{XY} + \sum_{\ell\ell'} \left[\sum_{i=1}^{N_{\text{mock}}} S_\ell^i \widehat{\mathbf{R}}_{\ell\ell'}^{XY} S_{\ell'}^i + \lambda \left| \widehat{\mathbf{R}}_{\ell\ell'}^{XY} \right| (1 - \delta_{\ell\ell'}^K) \right], \quad (33)$$

where

$$S_\ell^i \equiv \frac{\{C_\ell^{XY}\}_i - \bar{C}_\ell^{XY}}{\sigma_\ell} \quad (34)$$

are the standardized angular power spectra and λ is a hyperparameter that describes the assumed noisiness of the off-diagonal terms. In the limit of $\lambda = 0$, thus assuming the off-diagonal terms of the sample covariance to be noise-free, one can show that equation (32) minimizes equation (33). We use the graphical lasso implementation of the `SCIKIT-LEARN` PYTHON package (Pedregosa et al. 2011), which also includes a cross-validation method to automatically choose the value for λ .

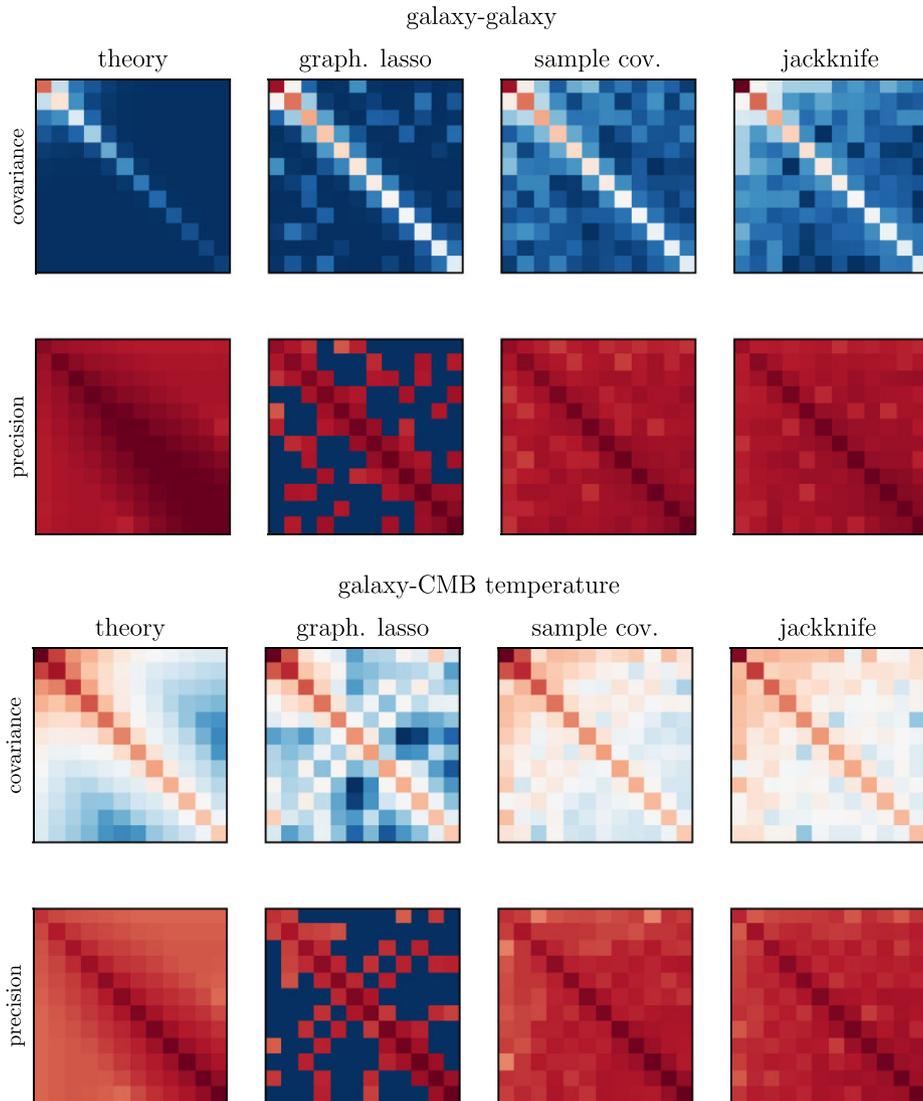


Figure 5. Comparison of the absolute values of the covariance (top line) and precision (bottom line) matrices for C_{ℓ}^{gg} and C_{ℓ}^{gT} obtained analytically/theoretically, from mock data using the graphical lasso algorithm and by computing their sample covariances, and from jack-knife resampling. The variance in the lowest ℓ -bin is shown in the top left-hand side of each panel, while ℓ increases towards the right and bottom, with the bins matching those of the measured power spectrum. The colour scaling is logarithmic.

3.6.5 Comparison of covariance matrices

We plot the covariance and precision matrices obtained with the above-mentioned estimators in Fig. 5. There is reasonable agreement among all of them, though one can spot some significant differences:

(i) The analytic galaxy–galaxy covariance shows smaller values on the diagonal at small scales as those obtained using numerical methods.

(ii) The graphical lasso variances agree well with the sample variance and the jack-knife variance. The off-diagonal values are smaller, which is expected as the method is set up to find sparse matrices. For the ISW covariance, the off-diagonal terms are smaller than the analytic prediction though, which hints at a too large value of the hyperparameter λ . However, increasing λ would also increase the suspicious lines of increased covariance perpendicular to the diagonal that are also prominent in the precision matrix.

(iii) The sample covariance matrix agrees on the diagonal well with the graphical lasso estimates, whereas the off-diagonal entries look like the analytic covariance matrix with added noise, as expected.

(iv) The jack-knife resampling method slightly underestimates the galaxy–temperature covariance at large scales, which is expected as the method is inherently blind to cosmic variance. However, for the galaxy–galaxy covariance, jack-knife resampling yields larger estimates of the covariance at large scales, which is because all other methods make use of a model, whereas the data show a large-scale power offset compared to our fiducial model that we further discuss in the following section. At smaller scales, the jack-knife covariance agrees remarkably well with the sample covariance.

As we shall later justify, ignoring the gg power spectrum at large scales, we use the sample covariance of our mocks to attain the main results of this article because it absorbs effects from the survey

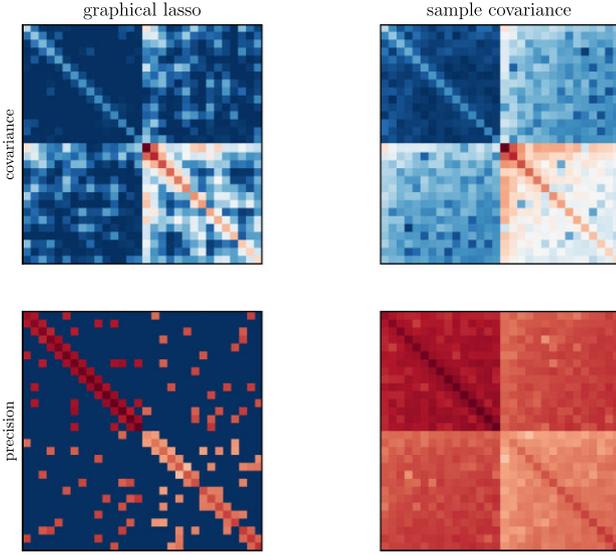


Figure 6. Covariance and precision matrices estimated from data vectors containing both galaxy autopower spectra and galaxy–temperature power spectra. The top left-hand block is the galaxy–galaxy covariance/precision, the bottom right-hand block the galaxy–temperature submatrix, and the top right- and bottom left-hand blocks show the galaxy–galaxy–galaxy–temperature covariance/precision.

window, does not rely on any hyperparameters and embodies cosmic variance in the large-scale gT power spectrum.

When performing a joint analysis of the galaxy–galaxy and galaxy–temperature power spectra, we generally have to take the galaxy–galaxy–galaxy–temperature covariance into account. We have estimated $\hat{\mathbf{S}}_{\ell\ell'}^{\text{gggT}}$ from mock realizations only because we do not have a reliable analytic model for it and jack-knife realizations have little advantage here, as cross-correlations are mostly unaffected by observational systematic errors such as foregrounds. Our estimated galaxy–galaxy–galaxy–temperature covariance and precision matrices are visualized in Fig. 6. By eye, we cannot identify any particular features in the $\hat{\mathbf{S}}_{\ell\ell'}^{\text{gggT}}$ estimated using the graphical lasso method. In the sample covariance, one can make out a slight increase on the diagonal at large scales, but above the first five ℓ -bins, we do not see any difference between diagonal and off-diagonal terms, raising the suspicion that these matrices are dominated by noise rather than actual correlations. Fortunately, using the full matrix shown in Fig. 6 provides an equivalent value of χ^2 as when dropping gggT correlations in the χ^2 computation. We henceforth set $\hat{\mathbf{S}}_{\ell\ell'}^{\text{gggT}} = 0$ for all ℓ and ℓ' .

3.7 Markov chain Monte Carlo sampling

Even though the angular power spectrum is not normally distributed at large scales (e.g. Verde et al. 2003; Percival & Brown 2006), the distribution of angular power spectra measured from our mock catalogues is approximately Gaussian when binning in relatively wide bins of width $\Delta\ell = 20$ due to the central limit theorem. We, therefore, conjecture the likelihood of the data \tilde{C}_ℓ , given the model C_ℓ as (ignoring the constant normalization term)

$$-2 \ln \mathcal{P} \left(\tilde{C}_\ell | C_\ell, \hat{\mathbf{K}}_{\ell\ell'} \right) = \sum_{X \in \{\text{gg, gT}\}} \sum_{\ell\ell'} \Delta C_\ell^X \hat{\mathbf{K}}_{\ell\ell'}^{XX} \Delta C_\ell^X, \quad (35)$$

with $\Delta C_\ell^X = \tilde{C}_\ell^X - C_\ell^X$. Note that we assume here that the shot noise has already been subtracted from \tilde{C}_ℓ (cf. Section 3.3).

Before using this likelihood in any Markov chain Monte Carlo (MCMC) sampling, we evaluate it for our fiducial model with and without putting the Limber approximation into service. A difference occurs only at the third significant digit that justifies our reliance on the Limber approximation to avoid our MCMC sampling being considerably more expensive.

We sample the posterior distribution of the parameters of interest using *Ensemble Slice Sampling* (Karamanis & Beutler 2020) implemented in the `ZEUS` code (Karamanis, Beutler & Peacock 2021). For one parameter, given a starting point x_0 and calling the probability density function to be sampled $f(x)$, *Slice Sampling* works by iterating over the following steps (Neal 2003):

- (i) Draw a uniformly distributed height y_i from the interval $[0, f(x_i)]$;
- (ii) Define the slice $S = \{x: y_i < f(x)\}$;
- (iii) Uniformly draw a new point x_{i+1} from S .

The advantages of this sampler compared to many other MCMC samplers are that one does not have to define any proposal distribution for efficient application (it is a so-called black box) and that its acceptance rate is 1. On the downside, the Slice Sampler has to evaluate $f(x)$ multiple times per step to numerically approximate the slice interval S .

For more than one parameter, each slice S has as many dimensions as parameters, thus, one has to define a direction along which the next point \mathbf{x}_{i+1} is chosen. `ZEUS` runs an ensemble of Slice Samplers in parallel, and, by default, the new point $\mathbf{x}_{i+1}^{\mathcal{T}}$ of the \mathcal{T} th walker is chosen along the vector

$$\boldsymbol{\eta}_{\mathcal{T}} = \mu \left(\mathbf{x}_i^{\text{L}} - \mathbf{x}_i^{\text{R}} \right), \quad (36)$$

where \mathbf{x}_i^{L} and \mathbf{x}_i^{R} are the current position of two walkers other than \mathcal{T} drawn uniformly and without replacement, and μ is a length-scale that, as the sampling progresses, is tuned to reduce the number of $f(x)$ evaluations needed to find the slice interval. As the distribution of walkers, after a burn-in period, resembles the target distribution, equation (36) naturally prefers directions of correlated parameters (Karamanis & Beutler 2020).

We employ `CHAINCONSUMER` (Hinton, Adams & Badger 2020) to analyse our chains.

4 RESULTS

4.1 The galaxy–galaxy autopower spectrum

In Fig. 7, we show the measured angular galaxy autopower spectrum $\tilde{C}_\ell^{\text{gg}}$ in ℓ -bins with width $\Delta\ell = 20$ for a flux limit of 4 mJy. We also plot the fiducial power spectrum that we use to set up `FLASK` along with percentile regions estimated from 3000 `FLASK` realizations. We see a good agreement of the fiducial model with the data at $\ell > 40$. At larger scales, however, we see more power than expected. We suspect that this offset is due to hitherto unidentified systematic effects and discuss this further below and in Appendix A. In Alonso et al. (2021), the angular clustering data from LOFAR on scales larger than the size of a pointing was removed due to systematic effects.

The assumption that this large-scale power excess is due to systematic effects is further supported by the fact that when we measure the galaxy–galaxy autopower spectrum in stripes of constant Declination with a width of 6° , we see less power at the largest scales in stripes that are closer to the South Pole (cf. Fig. 8). Interestingly,

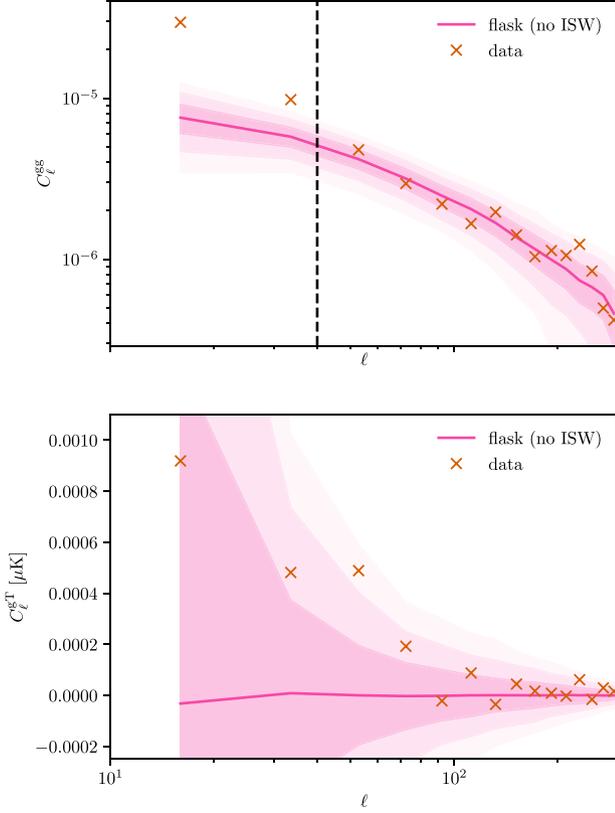


Figure 7. The angular autopower spectra $\tilde{C}_\ell^{\text{gg}}$ measured from the RACS island catalogue (crosses, top panel) and the RACS-Planck cross-power spectrum $\tilde{C}_\ell^{\text{gT}}$ (bottom). The magenta line shows the median C_ℓ of the FLASK realizations and the shaded regions show the 68-, 95-, and 99.75-percentile regions. In the top plot, we mark $\ell = 40$ as the upper bound of the distrusted multipole range that we do not include in our analyses of C_ℓ^{gg} .

our mock catalogues suggest that the error on $\tilde{C}_{\ell=24}^{\text{gg}}$ increases towards the equator as well, regardless of the increased area subtended by the Declination strip. Since the only direction-dependent information that enters the generation of the mock catalogues is the radio data weighting function $w(\theta)$, we suspect this unexpected behaviour to be due to an increased number of pixels where $w(\theta)$ is low as we go further north. This shall be studied in more detail in future work in preparation for the EMU survey.

Despite this behaviour that is correlated with Declination, we cannot simply ignore data on the fact that they do not match our expectations. We will therefore perform a first bias measurement both with and without considering large-scale (i.e. $\ell \leq 40$) galaxy clustering data. The measured bias parameters for both $n(z)$ models and bias models (i)–(iv) are tabulated in Table 2. These have been obtained by simple numerical optimization methods and thus are reported without errors that we deliver later (cf. Tables 3 and 4) after running MCMC jointly on the gg and gT power spectra. The aim here is to check how well our modelling assumptions can describe the data.

When we include multipoles at $\ell \leq 40$, the galaxy bias (for non-constant bias parametrizations) surprisingly decreases with redshift. Furthermore, the minimum χ^2 is from three to more than twelve times larger than the number of degrees of freedom, suggesting that our model is insufficient at large scales. We, therefore, make use of the galaxy autopower spectrum at $\ell > 40$ only (unless otherwise

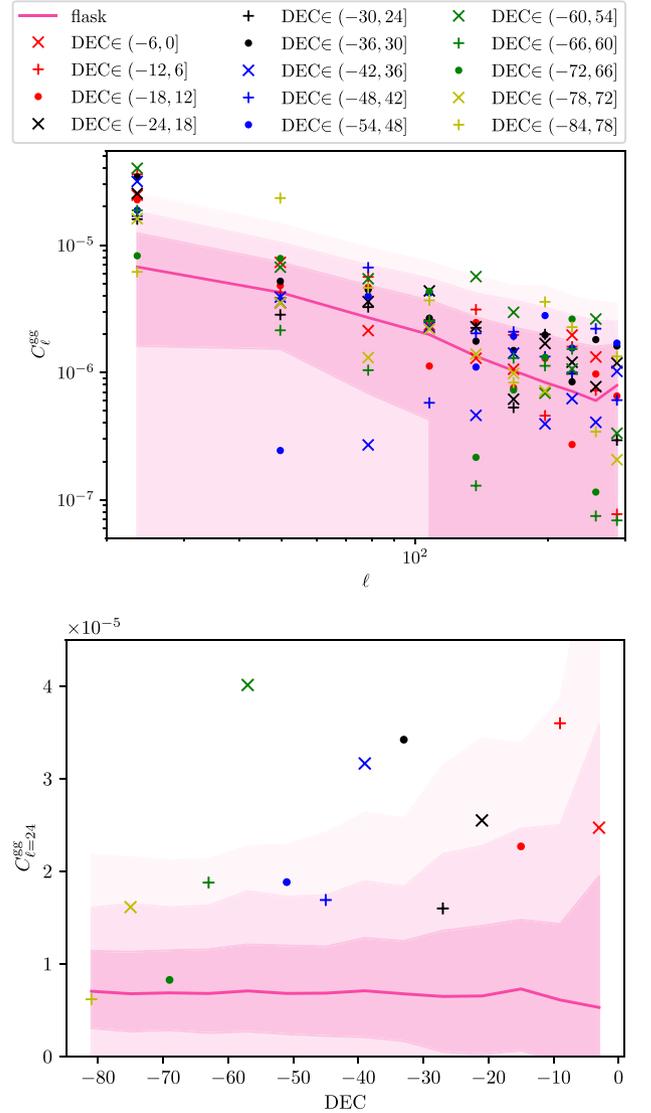


Figure 8. The top panel shows the $\tilde{C}_\ell^{\text{gg}}$ measured in Declination (Dec.) bands with widths of 6° . The bottom panel shows the $\tilde{C}_{\ell=24}^{\text{gg}}$ at $\ell = 24$ as a function of Dec. The marker colours and shapes coincide in both plots. The magenta line and shaded regions show the median and 68-, 95-, and 99.75-percentile regions of the FLASK realizations, as in Fig. 7.

stated) and leave it to be re-analysed in the future when either an extended model or a better understanding of systematic effects is at hand.

Omitting large-scale multipoles at $\ell \leq 40$, we find almost equal values of $\chi^2/\text{d.o.f.}$ for all bias parametrizations and both $n(z)$ models, with the exception of using the T-RECS $n(z)$ with a constant bias. This model stands out in Fig. 9 as the one where $n(z)b(z)$ drops quite sharply above $z > 1$, whereas other T-RECS models have a wider peak region that extends up to $z \sim 2$ and the $n(z)b(z)$ of best-fitting SKADS models have a peak similar to the constant-bias T-RECS model but have a plateau between $1.4 \lesssim z \lesssim 2.6$ such that, in this redshift range, the average $n(z)b(z)$ is the same as for the T-RECS models with bias evolution. Yet, even in the constant-bias T-RECS case, $\chi^2/\text{d.o.f.}$ is much lower than in any full ℓ -range case. In all other cases, $\chi^2/\text{d.o.f.}$ is only marginally greater than unity, implying that all of these models describe the data well. Instead of trying to choose

Table 2. Maximum posterior values of the galaxy bias b_0 at redshift $z = 0$ and the bias's redshift evolution from minimizing the C_ℓ^{gg} χ^2 .

Model	b_0	db/dz or $d\ln b/dz$	z_{eff}	$b(z_{\text{eff}})$	χ_{min}^2	$\chi_{\text{min}}^2/\text{d.o.f.}$
All ℓ						
SKADS						
const. $b(z)$	3.63	–	1.56	3.63	172.8	12.34
linear $b(z)$	4.25	–0.85	1.28	3.16	131.6	10.12
exp. $b(z)$	7.80	–1.79	0.65	2.44	45.27	3.482
- w/cut-off	7.79	–1.79	0.88	1.61	45.10	3.469
T-RECS						
const. $b(z)$	2.66	–	1.13	2.66	93.85	6.703
linear $b(z)$	2.95	–0.59	1.00	2.36	72.05	5.542
exp. $b(z)$	3.96	–1.11	0.72	1.78	45.83	3.525
- w/cut-off	3.96	–1.11	0.65	1.92	45.72	3.517
Only $\ell > 40$						
SKADS						
const. $b(z)$	3.24	–	1.56	3.24	12.70	1.058
linear $b(z)$	2.74	0.59	1.72	3.75	12.11	1.101
exp. $b(z)$	2.83	0.15	1.72	3.66	12.11	1.101
- w/cut-off	2.71	0.23	1.64	3.62	12.09	1.099
T-RECS						
const. $b(z)$	2.41	–	1.13	2.41	17.19	1.433
linear $b(z)$	1.33	1.62	1.40	3.60	11.78	1.071
exp. $b(z)$	1.52	0.61	1.51	3.82	11.71	1.065
- w/cut-off	1.43	0.73	1.34	3.80	11.67	1.061

Notes. The effective redshift z_{eff} has been obtained by integrating over $zb(z)n(z)$ for the best-fitting bias parameters. The bias $b(z)$ is included in the z_{eff} integral as it is degenerate with $n(z)$, and, therefore, our $b(z)$ measurement is, to some degree, also effectively accounts for potential deviations from the fiducial $n(z)$ distribution.

Table 3. Maximum posterior values of the galaxy bias b and A_{ISW} from jointly analysing $\tilde{C}_\ell^{\text{gg}}$ and $\tilde{C}_\ell^{\text{gT}}$ assuming a SKADS $n(z)$.

gg ℓ -range	b	A_{ISW}	χ_{min}^2	d.o.f.	$\chi_{\text{min}}^2/\text{d.o.f.}$
All ℓ	$3.613^{+0.087}_{-0.050}$	$0.68^{+0.32}_{-0.36}$	187	30	6.2
$\ell > 40$	$3.248^{+0.068}_{-0.094}$	$0.82^{+0.39}_{-0.33}$	27.1	28	0.97

Note. We use the full available ℓ range in the gT spectrum, but omit $\ell \leq 40$ when analysing the gg spectrum for the bottom line.

one particular model, we shall use the scatter of the results obtained with these different models to estimate the systematic uncertainty.

In any case, neither the best-fitting parameters nor the minimum χ^2 show much difference between the pure exponential bias parametrization (ii) and its variant (i) with a constant bias above $z > 1.5$. Considering this result and the fact that we regard the redshift cut as arbitrary, we do not pursue model (i) any further.

4.2 The galaxy–temperature cross-power spectrum

As the ground-based radio observations of galaxies are subject to very different systematic effects as CMB observations from space, we do not expect significant systematic contributions to the measured cross-power spectrum $\tilde{C}_\ell^{\text{gT}}$ shown in Fig. 7. We mark again the median and confidence regions estimated from 3000 FLASK realizations, however, this time, we initialize each simulation to have no intrinsic correlation between the galaxy and CMB map.

We expect most of the ISW signal at large scales. Thus, even though we ignore the first two multipole bins in C_ℓ^{gg} , these are

crucial in the C_ℓ^{gT} analysis. As we show in Appendix B, it is actually conservative to include large-scale C_ℓ^{gT} multipoles in the ISW analysis. Assuming that the observed $\tilde{a}_{\ell m}^{\text{g}} = a_{\ell m}^{\text{g}} + f_{\ell m}$ is the sum of the true cosmological $a_{\ell m}^{\text{g}}$ and some unknown systematic $f_{\ell m}$, we have the observed power spectra

$$\begin{aligned} \langle \tilde{C}_\ell^{\text{gg}} \rangle &= \langle \tilde{a}_{\ell m}^{\text{g}} \tilde{a}_{\ell m}^{\text{g}*} \rangle = C_\ell^{\text{gg}} + (\langle f_{\ell m} a_{\ell m}^{\text{g}*} \rangle + \text{c.c.}) + \langle f_{\ell m} f_{\ell m}^* \rangle, \text{ and} \\ \langle \tilde{C}_\ell^{\text{gT}} \rangle &= \langle \tilde{a}_{\ell m}^{\text{g}} \tilde{a}_{\ell m}^{\text{T}*} \rangle = C_\ell^{\text{gT}} + \langle f_{\ell m} a_{\ell m}^{\text{T}*} \rangle. \end{aligned} \quad (37)$$

If $f_{\ell m}$ is an observational systematic, e.g. a terrestrial or Galactic foreground, then it is uncorrelated with the true cosmological signal, i.e. $\langle f_{\ell m} a_{\ell m}^{\text{g}*} \rangle = \langle f_{\ell m} a_{\ell m}^{\text{T}*} \rangle = 0$. Hence, $\langle \tilde{C}_\ell^{\text{gT}} \rangle = C_\ell^{\text{gT}}$ is unaffected by the systematic, whereas $\tilde{C}_\ell^{\text{gg}}$ is biased by the autopower spectrum of $f_{\ell m}$. On the other hand, if the observed excess is due to a theoretical systematic, i.e. it is not predicted well by our modelling of the density field, we will see unexpected behaviour in the gT cross-power spectrum as well. In the latter case, we will see values of χ^2 that exceed the number of degrees of freedom by far. We therefore proceed including the full available multipole range in the gT analysis and will present a simple χ^2 test later to justify this.

Our first step in analysing the significance of the ISW signal in the gT cross-power spectrum is to compare the values of

$$\chi^2 = \sum_{\ell, \ell'} \left(\tilde{C}_\ell^{\text{gT}} - C_\ell^{\text{gT}} \right) \mathbf{K}_{\ell\ell'}^{\text{gTgT}} \left(\tilde{C}_{\ell'}^{\text{gT}} - C_{\ell'}^{\text{gT}} \right), \quad (38)$$

for the two hypotheses of existence and non-existence of gT cross-correlations due to the ISW effect. In the former case, C_ℓ^{gT} is as defined in equation (6), while in the latter, we just have $C_\ell^{\text{gT}} = 0$. Using the sample covariance matrix of 3000 mock catalogues, we obtain $\chi^2 = 17.7$ for the null hypothesis ($C_\ell^{\text{gT}} = 0$) and $\chi^2 = 10.9$ for the C_ℓ^{gT} -model given in equation (6). If we use instead a precision matrix estimated from the same set of mock catalogues using the graphical lasso method, we find $\chi^2 = 17.8$ for the null hypothesis and $\chi^2 = 11.0$ for ISW hypothesis. So in both cases, adopting an ISW model reduces χ^2 by 6.8. Using the theoretical precision matrix, we obtain $\chi^2 = 10.7$ for the null hypothesis and $\chi^2 = 7.4$ for the ISW hypothesis, underestimating the mode-coupling contribution of the survey mask, and, hence, the significance of the ISW detection. On the contrary, ignoring cosmic variance, the jack-knife resampling increases the significance with $\chi^2 = 21.3$ and 12.2 for the null and ISW hypotheses, respectively.

We can further describe the significance of this finding in terms of the signal-to-noise ratio (Becker et al. 2016):

$$S/N = \frac{\sum_{\ell, \ell'} \tilde{C}_\ell^{\text{gT}} \mathbf{K}_{\ell\ell'}^{\text{gTgT}} C_{\ell'}^{\text{gT}}}{\sqrt{\sum_{\ell, \ell'} C_\ell^{\text{gT}} \mathbf{K}_{\ell\ell'}^{\text{gTgT}} C_{\ell'}^{\text{gT}}}}. \quad (39)$$

We evaluate equation (39) again using both covariance matrices estimated from simulations and an ℓ -binning with $\Delta\ell = 20$, which yields

$$\frac{S}{N} = 2.8. \quad (40)$$

Alternatively, we attain $S/N = 1.9$ with the analytic and $S/N = 3.2$ with the jack-knife covariance matrices.

4.3 Parameter constraints

In the previous subsection, we have detected a positive cross-correlation between the galaxy and temperature maps at 2.8σ compared to the null hypothesis of no correlation. However, in the gT cross-power spectrum, the amplitude of the ISW signal A_{ISW}

Table 4. Maximum posterior values of the galaxy bias b_0 at redshift $z = 0$, its redshift evolution expressed as db/dz or $d\ln(b)/dz$, and A_{ISW} from jointly analysing $\tilde{C}_\ell^{\text{gg}}$ and $\tilde{C}_\ell^{\text{gT}}$.

Model	b_0	A_{ISW}	db/dz	$d\ln(b)/dz$	z_{eff}	$b(z_{\text{eff}})$
SKADS						
const. $b(z)$	$3.248^{+0.068}_{-0.094}$	$0.82^{+0.39}_{-0.33}$	–	–	1.56	3.248
linear $b(z)$	$2.60^{+0.78}_{-0.59}$	$0.87^{+0.48}_{-0.31}$	$0.73^{+0.73}_{-0.78}$	–	1.78	3.900
exp. $b(z)$	$2.82^{+0.65}_{-0.59}$	$0.89^{+0.43}_{-0.34}$	–	$0.13^{+0.24}_{-0.22}$	1.69	3.513
T-RECS						
const. $b(z)$	$2.407^{+0.059}_{-0.061}$	$0.86^{+0.32}_{-0.39}$	–	–	1.13	2.407
linear $b(z)$	$1.10^{+0.70}_{-0.45}$	$0.96^{+0.41}_{-0.38}$	$1.87^{+0.66}_{-0.79}$	–	1.43	3.774
exp. $b(z)$	$1.56^{+0.32}_{-0.38}$	$0.99^{+0.33}_{-0.42}$	–	0.59 ± 0.24	1.50	3.852

Notes. The effective redshift z_{eff} is obtained as in Table 2. We omit $\ell \leq 40$ when analysing the gg spectrum.

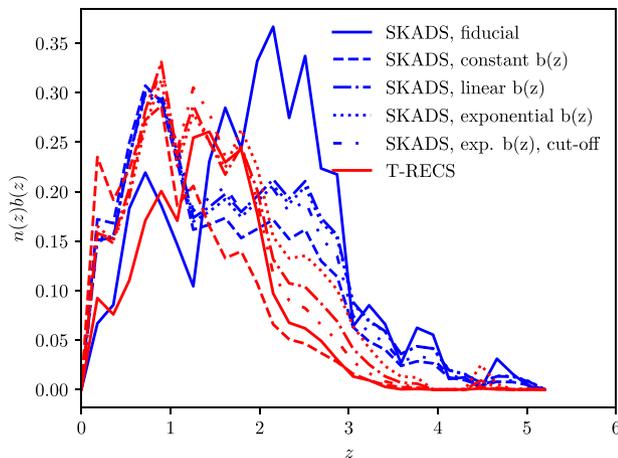


Figure 9. The product of the SKADS/T-RECS $n(z)$ distribution and the best-fitting galaxy biases $b(z)$ for when we use the fiducial bias parameters used in Raccanelli et al. (2015), Bernal et al. (2019), and Asorey & Parkinson (2021). The solid red line shows the estimate of $n(z)b(z)$ from T-RECS (cf. Fig. 13).

is degenerate with the galaxy bias $b(z)$, as well as the redshift distribution of radio continuum sources per steradian $n(z)$. In this section, we reevaluate the significance of our ISW detection taking our ignorance on $b(z)$ and $n(z)$ into account.

As the gg autopower spectrum depends only on $b^2(z)n^2(z)$, we can use it to anchor $b(z)$ and, thus, lift the $b(z)$ – A_{ISW} degeneracy. At the outset, we fix $n(z)$ to the one predicted by SKADS (Wilman et al. 2008). We ensemble slice sample a constant bias parameter b and the ISW signal amplitude A_{ISW} first using the full measured gg autopower spectrum, and then repeat the same analysis restricting the gg autopower spectrum to $\ell > 40$ only, while still taking the full gT cross-power spectrum. The resulting b – A_{ISW} posterior contours are plotted in Fig. 10. To fit the excess power at low multipoles with our two-parameter model, the galaxy bias b is required to be significantly larger than for the case where we ignore galaxy autocorrelations at $\ell \leq 40$. As we perform both analyses on the same multipole range of the gT cross-power spectrum whose amplitude is given by the product bA_{ISW} , using the full available range of scales favours smaller values of A_{ISW} . We are reassured by the fact that the marginalized posteriors on A_{ISW} are mostly consistent with each other. The significance of our ISW detection is thus largely unaffected by the large-scale power excess.

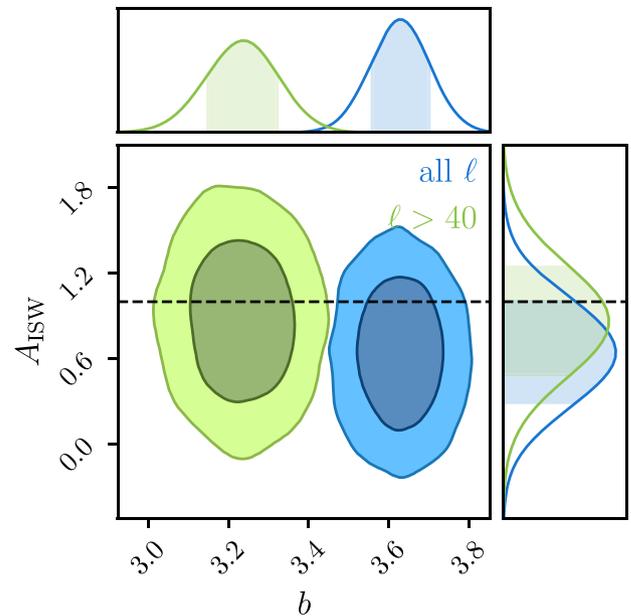


Figure 10. Posterior contours of the galaxy bias b and A_{ISW} from jointly analysing $\tilde{C}_\ell^{\text{gg}}$ and $\tilde{C}_\ell^{\text{gT}}$ assuming an SKADS $n(z)$. We use the full available ℓ -range in the gT spectrum, but we omit $\ell \leq 40$ for the green contours. The dashed lines indicate $A_{\text{ISW}} = 1$. The dark (light) shaded contours contain 68 (95) percent of the MCMC chain elements. The shaded regions in the histograms correspond to the 68 per cent credible interval.

The best-fitting values are given in Table 3 along with χ^2_{min} , the minimum value of χ^2 . When omitting large scales in the gg autopower spectrum, we obtain a reduced χ^2 of 0.97, indicating that our modelling works well to describe the data at these scales. When we include multipoles at $\ell \leq 40$, the reduced χ^2 increases by more than six times the previous value, suggesting that our model is insufficient at large scales. We, therefore, ignore the galaxy autopower spectrum at $\ell \leq 40$ in the following parts of this article and leave it to be re-analysed in the future when either an extended model or a better understanding of systematic effects is readily available.

Our next step is to allow the bias to evolve with redshift. To avoid unphysical results, we adopt an additional prior on combinations of $b(z = 0)$ and db/dz in the linear bias case, that is that $b(z) < 0$ is excluded at all redshifts z probed by the survey. This condition is always fulfilled by the exponential bias parametrization as long as

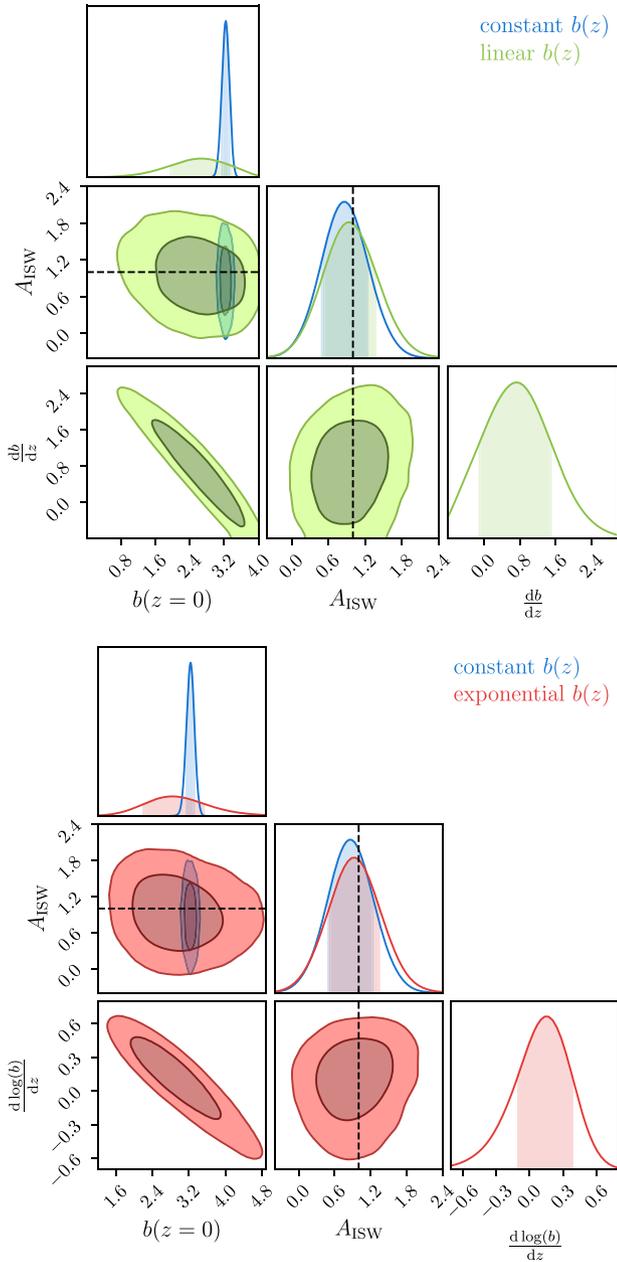


Figure 11. Posterior contours of the galaxy bias $b(z=0)$ at redshift $z=0$, its evolution and A_{ISW} . In the top plot, we compare the contours for constant and linear bias parametrizations. The bottom panel shows the contours obtained using an exponential bias parametrization alongside the constant bias contours. All contours shown in this figure have been made without the contributions of multipoles $\ell \leq 40$ to the galaxy–galaxy autopower spectrum \hat{C}_ℓ^{gg} . The shading signifies the same as in Fig. 10.

$b(z=0) > 0$. In the top (bottom) panel of Fig. 11, we compare the posterior contours of a constant bias with those resulting from using a parametrization where the bias evolves linearly (exponentially). In both cases, introducing more freedom to the bias model leads to a larger uncertainty in the bias, but the lower bounds on A_{ISW} are largely unaffected by the bias parametrization. However, we can also observe that if the bias evolves more strongly with redshift, slightly larger values of A_{ISW} are likely, and overall, the evolving bias parametrizations favour to some degree higher values of A_{ISW} ,

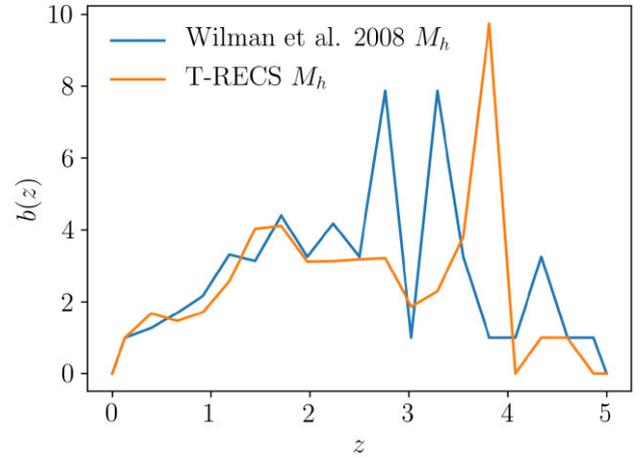


Figure 12. Ensemble bias for the T-RECS medium simulation for a flux cut of 4 mJy for an 888-MHz catalogue. In this case, we fix the redshift at which the bias is evaluated for high redshifts, depending on the population.

bringing its best-fitting value closer to unity (cf. Table 4). For the most part though, the marginalized posterior distribution of A_{ISW} is robust under different bias parametrizations.

Finally, we check for the impact of our ignorance on $n(z)$. We have fairly good knowledge of $n(z)$ for redshifts $z \lesssim 2$. The high-redshift tail of the distribution is one of the largest sources of uncertainty in the use of radio continuum surveys. This is reflected in the differences between the $n(z)$ estimated from SKADS and the one estimated from T-RECS, as plotted in Fig. 2. We, therefore, repeat all the analyses done so far using the $n(z)$ from T-RECS and also list their results in Table 4. Using the T-RECS $n(z)$, we obtain significantly lower values of the bias b at redshift $z=0$, but also significantly stronger redshift evolution, such that at the effective redshift z_{eff} , when allowing for redshift evolution, the bias is roughly the same as when using the SKADS $n(z)$ (cf. Table 3). As can be seen in Fig. 4, at redshift $z \sim 1$, also a constant bias with SKADS $n(z)$ agrees with the evolving parametrizations. Despite T-RECS’s preference for strong bias evolution, we can see in Fig. 9 that the product of $n(z)$ and the best-fitting $b(z)$ is generally unaffected by the choice of $n(z)$ and bias model. T-RECS favours a stronger localization of objects below $z < 2$ and suppresses the high-redshift tail present in SKADS. The T-RECS analysis provides us with 68 per cent-credible intervals that are almost equal to the ones from the SKADS analysis, as can be seen by comparing red and blue whiskers in the top panel of Fig. 14. Nevertheless, we obtain larger best-fitting values, with $A_{\text{ISW}} = 0.99^{+0.33}_{-0.42}$ measured by assuming the T-RECS $n(z)$ and an exponentially evolving bias being the closest to 1.

We are going to use the scatter among these different predictions to estimate the systematic uncertainty of our final A_{ISW} result.

4.4 Comparison of bias measurements with previous models

Before presenting a combined final result of our A_{ISW} , we compare briefly our phenomenological bias results with previous results.

Most studies consider a bias model based on N -body dark matter simulations (e.g. Wilman et al. 2008; Bonaldi et al. 2019), in which the bias for each population of radio-galaxies is defined as belonging to a given halo mass M_h . To check the robustness of this approach, we have used the halo masses of each species of radio galaxy (as specified in Bonaldi et al. 2019). For each galaxy type i in the T-

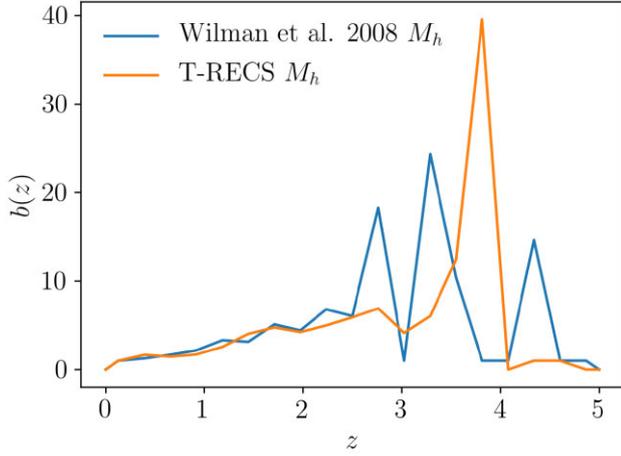


Figure 13. Same as Fig. 12 but without redshift cut when evaluating the halo bias in equation (41).

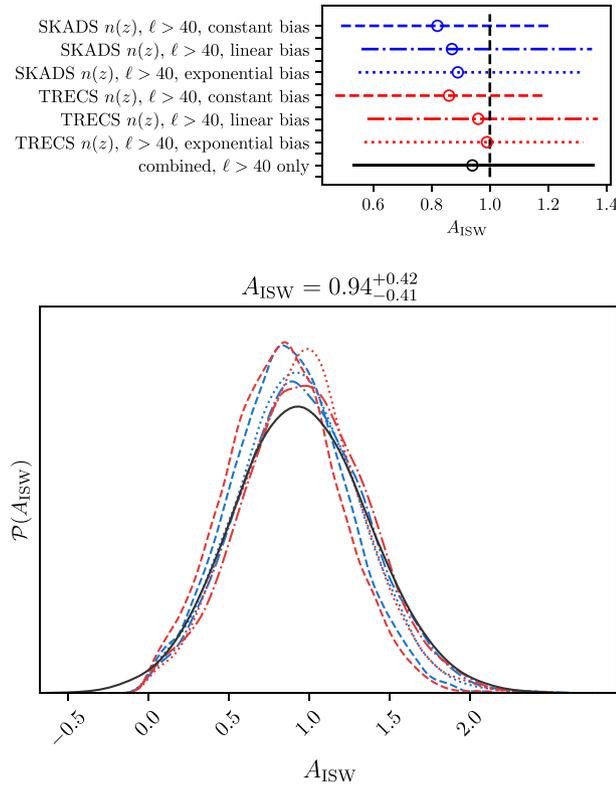


Figure 14. Top panel: boxplot summarizing the A_{ISW} results obtained using different $n(z)$ and bias parametrizations, leaving out $\tilde{C}_\ell^{\text{gg}}$ at $\ell \leq 40$. Potential systematic biases are lessened when the different $n(z)$ and bias models are combined using the BACCUS approach (black line). Bottom panel: the corresponding marginalized posterior distributions on A_{ISW} .

RECS medium sample:

$$b_i(z) = \frac{\int dM n_i(M, z) b_h(M, z)}{\int dM n_i(M, z)}, \quad (41)$$

where $n_i(M, z)$ is the halo mass function for galaxies of population type i in the redshift bin $(z, z + dz)$ and $b_h(M, z)$ is the halo bias, which we estimate using Colossus suite (Diemer 2018) and the halo

model from Tinker et al. (2010). The total bias is then computed using equation (25).

In Fig. 12, we show the T-RECS total bias $b(z)$ redshift evolution when using a flux cut of 4 mJy at 888 MHz, using the approach above. As proposed in Wilman et al. (2008), for high redshifts, we evaluate equation (41) at a fixed redshift ($z = 1.5$ or $z = 3$ depending on the population). We compare this result with the Wilman et al. (2008) approach by using the SKADS simulation, in which we evaluate the halo bias at the corresponding halo mass value proposed for the SKADS simulation for each population type. We see that the estimated bias is similar for both approaches. We also see that if we mix information from both simulations, we obtain a wrong result. The scattered and noisy behaviour at larger redshifts is due to the small number of haloes that remain after all the applied cuts. In Fig. 13, we show the same $b(z)$ but without applying any redshift cut-off on the bias evaluation. We see that the bias grows, both for SKADS and T-RECS to extremely high values due to the FR II galaxy population.

As the T-RECS $n(z)$ model predicts only a few objects at high redshifts, we still see agreement between the N -body result and our measurements of the T-RECS $n(z)b(z)$ high-redshift tail (cf. Fig. 9). Surprisingly, the N -body based T-RECS model underpredicts $n(z)b(z)$ at redshifts $z \lesssim 1$. This might be due to the SFG abundance being underpredicted and needs further investigation in the preparation for EMU.

We also added the parametrized form of the bias as given as part of the Wilman et al. (2008) analysis to Fig. 9. Similar as in the T-RECS case, we measure a lower $n(z)b(z)$ at $z \lesssim 1$ with RACS as the N -body simulations suggest; thus, SKADS might also predict less SFGs as there are in reality. Furthermore, we see that $n(z)b(z)$ peaks at a much higher redshift than all best-fitting models. The peak is where we expect FR IIs to dominate the sample. We deduce from this that the abstruse FR II bias has so far been overestimated. Another possible explanation for the mismatch in Fig. 9 could be that FR IIs start to dominate at higher redshifts than previously thought. We leave a detailed examination of this issue for future work.

4.5 Combining different predictions

In Fig. 14, we show the best-fitting values and marginalized posterior distributions for A_{ISW} obtained using different bias models and $n(z)$ distributions (for an extended version of this plot including $\tilde{C}_\ell^{\text{gg}}$ data at $\ell \leq 40$, we refer to Appendix C). Given their scatter, we can estimate the overall uncertainty A_{ISW} including the uncertainty due to unknown systematics. To avoid expectation bias, we follow the BACCUS (BAYesian Conservative Constraints and Unknown Systematics) approach of Bernal & Peacock (2018), where we assume that unknown systematic effects have biased all of our measurements $A_{\text{ISW}}^{(i)}$ by an unknown $\Delta A_{\text{ISW}}^{(i)}$ and that they also degrade the variance σ_i^2/ζ_i by a factor of ζ_i . As all our A_{ISW} results are single-peaked but asymmetric, we summarize their marginalized posterior distributions in terms of variable Gaussians, i.e. as a normal distribution whose scale parameter $\sigma(A_{\text{ISW}})$ depends on the value of A_{ISW} in the exponential (while keeping the log-determinant fixed; Bartlett 1953):

$$-2 \ln \mathcal{P}(A_{\text{ISW}} | \sigma_i, A_{\text{ISW}}^{(i)}) = \left(\frac{A_{\text{ISW}} - A_{\text{ISW}}^{(i)}}{\sigma_i(A_{\text{ISW}})} \right)^2 + \text{const.} \quad (42)$$

Assuming a linear relationship and imposing that the asymmetric errors σ_i^+ and σ_i^- define the full width at half-maximum (FWHM),

one finds (Barlow 2004)

$$\sigma_i(A_{\text{ISW}}) = \frac{2\sigma_i^+\sigma_i^-}{\sigma_i^+ + \sigma_i^-} + \frac{\sigma_i^+ - \sigma_i^-}{\sigma_i^+ + \sigma_i^-} (A_{\text{ISW}} - A_{\text{ISW}}^{(i)}). \quad (43)$$

As A_{ISW} is also statistically independent from the galaxy bias parameters, introducing the systematic bias $\Delta A_{\text{ISW}}^{(i)}$ and variance degradation parameters ζ_i , we write the log-likelihood of each measurement (dropping constant terms) as

$$\begin{aligned} 2 \ln \mathcal{P} \left(A_{\text{ISW}}, \zeta_i, \Delta A_{\text{ISW}}^{(i)} \middle| \sigma_i, A_{\text{ISW}}^{(i)} \right) \\ = \ln(\zeta_i) - \zeta_i \left(\frac{A_{\text{ISW}} + \Delta A_{\text{ISW}}^{(i)} - A_{\text{ISW}}^{(i)}}{\sigma_i(A_{\text{ISW}})} \right)^2. \end{aligned} \quad (44)$$

Since measurements (i) come from the same data, we cannot just add up the individual log-likelihoods. Instead, we consider the average log-likelihood:

$$\begin{aligned} \ln \mathcal{P} (A_{\text{ISW}}, \zeta, \Delta A_{\text{ISW}} \middle| \sigma, A_{\text{ISW}}) \\ = \frac{1}{6} \sum_{i=0}^6 \ln \mathcal{P} \left(A_{\text{ISW}}, \zeta_i, \Delta A_{\text{ISW}}^{(i)} \middle| \sigma_i, A_{\text{ISW}}^{(i)} \right). \end{aligned} \quad (45)$$

We follow Bernal & Peacock (2018) in our choice of priors on the scaling parameters σ and the systematic bias shifts ΔA_{ISW} . Thus, we assume that we have estimated the size of our statistical errors correctly, and, as a consequence, that the prior on the scaling parameters is (Hobson, Bridle & Lahav 2002)

$$\mathcal{P}(\zeta_i) \propto \begin{cases} \exp(-\zeta_i) & \text{if } \zeta_i > 0, \\ 0 & \text{else.} \end{cases} \quad (46)$$

We also choose a zero-centred Gaussian prior on $\Delta A_{\text{ISW}}^{(i)}$ with width σ_i .

We display at the bottom of Fig. 14 the posterior on A_{ISW} after marginalizing over ζ and ΔA_{ISW} . We find a best-fitting value of $A_{\text{ISW}} = 0.94^{+0.42}_{-0.41}$, thus 2.3σ away from $A_{\text{ISW}} = 0$. Allowing for both statistical and systematic uncertainties in this conservative approach, we obtain a probability for a positive A_{ISW} of 98.9 per cent.

5 SUMMARY

(i) We have measured the angular power spectrum of the radio continuum sources detected above a 4-mJy flux-density limit by the Rapid ASKAP Continuum Survey at 888 MHz, in autocorrelation and also in cross-correlation with temperature maps of the cosmic microwave background from the *Planck* mission.

(ii) We constructed estimates of the variance of the angular power spectra, using the purely analytic prediction from theory, jack-knife resampling of the catalogue data, and two methods that use simulation of mock RACS catalogues using the Full-sky Lognormal Astrofields Simulation Kit (FLASK) (sample covariance of the mocks, and a graphical lasso estimator learning sparse covariance matrices from simulations). All of these gave roughly consistent results, with the sample-covariance and jack-knife approaches predicting more off-diagonal covariance compared to the others.

(iii) We have tested four different bias parametrizations, the goodness of fit is almost indistinguishable among them, making it impossible to pick only one of them. We have found that the product of the best-fitting biases $b(z)$ and the redshift distributions of sources per steradian $n(z)$ has lower values at $z \gtrsim 2$ than what we had predicted from SKADS and T-RECS simulations, hinting towards the assumed FR II bias value being too large.

(iv) We have found that the angular autopower spectrum of RACS galaxies is consistent with the prediction from Λ CDM, except on large scales, $\ell \leq 40$, where we detect an excess that we believe is due to systematics.

(v) We have split the RACS catalogue into different regions and measured the angular power spectrum, to test for systematic causes for the excess. We have found a tentative trend showing that the large-scale excess is more pronounced for regions with Declination closer to the equator. However, the error on the power spectrum estimated from mocks also increases towards the equator, thus, the measured power is consistent with the Λ CDM prediction in almost all Declination strips. In comparison, we see no such trend with Right Ascension. This strongly implies that the excess is due to a systematic effect associated with the noise or some other observational effect.

(vi) We have detected a cross-correlation between the galaxy distribution and the distribution of hot and cold spots in the cosmic microwave background. This cross-correlation has been measured through the angular power cross-spectrum C_ℓ^{gT} , and is significant at 2.8σ relative to the null hypothesis of no cross-correlation.

(vii) We have found that when fitting the data from both the auto- and cross-correlation, the fit is consistent with the Λ CDM prediction (when the $\ell \leq 40$ data are removed from the autopower spectrum).

(viii) We have parametrized the amplitude of the cross-correlation signal A_{ISW} . We find that when combining the angular auto- and cross-power spectra, and assuming an $n(z)$ from SKADS and a constant bias model, that $A_{\text{ISW}} = 0.82^{+0.39}_{-0.33}$. These constraints are not very sensitive to the choice of the number count model or bias model.

(ix) When using the BACCUS approach to marginalize over different assumptions on the bias and number count model, allowing for unknown systematic biases and for possible posterior widening due to unknown systematic effects, we have found $A_{\text{ISW}} = 0.94^{+0.42}_{-0.41}$, corresponding to a 2.3σ or 98.9 per cent detection of the ISW effect and, hence, of dark energy.

This analysis has demonstrated that a few weeks on-source time of ASKAP observations provide data for meaningful cosmological analyses, while identifying what points have to be addressed in the analysis pipeline to reap the full potential of the upcoming EMU survey. The cosmological utility of the clustering statistics of radio continuum galaxies can only improve through the pathfinder era, to reach maturity with the SKA Observatory.

ACKNOWLEDGEMENTS

BBK, DP, and FQ are supported by the project ‘Understanding Dark Universe Using Large Scale Structure of the Universe’, funded by the Ministry of Science of the Republic of Korea. SC acknowledges support from the ‘Departments of Excellence 2018-2022’ Grant (L. 232/2016) awarded by the Italian Ministry of University and Research (MUR). JA has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 776247 EWC. CLH is supported by a Leverhulme Trust Early Career Research Fellowship.

This work was supported by the high performance computing clusters Seondeok at the Korea Astronomy and Space Science Institute. This research made substantial use of `ASTROPY`,³ a community-developed core PYTHON package for astronomy (Astropy Collaboration et al. 2013; Price-Whelan et al. 2018), `CAMB` (Lewis et al. 2000;

³<http://www.astropy.org>.

Howlett et al. 2012), CHAINCONSUMER (Hinton et al. 2020), FLASK (Xavier et al. 2016), HEALPIX and HEALPY (Górski et al. 2005; Zonca et al. 2019), NAMASTER (Hivon et al. 2002; Alonso et al. 2019), NUMPY package (Oliphant 2006), SCIKIT-LEARN (Pedregosa et al. 2011), SCIPY (Virtanen et al. 2020), MATPLOTLIB (Hunter 2007), and ZEUS packages (Karamanis et al. 2021).

The Australian SKA Pathfinder is part of the Australia Telescope National Facility (<https://ror.org/05qajvd42>) that is managed by CSIRO. Operation of ASKAP is funded by the Australian Government with support from the National Collaborative Research Infrastructure Strategy. ASKAP uses the resources of the Pawsey Supercomputing Centre. Establishment of ASKAP, the Murchison Radio-astronomy Observatory and the Pawsey Supercomputing Centre are initiatives of the Australian Government, with support from the Government of Western Australia and the Science and Industry Endowment Fund. We acknowledge the Wajarri Yamatji people as the traditional owners of the Observatory site.

DATA AVAILABILITY

The RACS radio continuum Stokes I source catalogue used in this analysis was generated from data available from the CSIRO ASKAP Science Data Archive (CASDA). The angular power spectra and covariance matrices that were measured are made available via a GitHub repository, which can be found at <https://github.com/racs-cosmology/isw>. The MCMC chains generated for the analysis are available on request to the authors.

REFERENCES

- Ade P. A. R. et al., 2014, *A&A*, 571, A12
- Aghanim N. et al., 2020a, *A&A*, 641, A1
- Aghanim N. et al., 2020b, *A&A*, 641, A6
- Alonso D., Ferreira P. G., 2015, *Phys. Rev. D*, 92, 063525
- Alonso D., Sanchez J., Slosar A., LSST Dark Energy Science Collaboration, 2019, *MNRAS*, 484, 4127
- Alonso D., Bellini E., Hale C., Jarvis M. J., Schwarz D. J., 2021, *MNRAS*, 502, 876
- Asorey J., Parkinson D., 2021, *MNRAS*, 506, 4121
- Astropy Collaboration et al., 2013, *A&A*, 558, A33
- Ballardini M., Maartens R., 2019, *MNRAS*, 485, 1339
- Bardeen J. M., Bond J. R., Kaiser N., Szalay A. S., 1986, *ApJ*, 304, 15
- Barlow R., 2004, *PHYSTAT* (2005): Statistical Problems in Particle Physics, Astrophysics and Cosmology. p. 56
- Bartlett M., 1953, *London Edinburgh Dublin Phil. Mag. J. Sci.*, 44, 249
- Becker M. R. et al., 2016, *Phys. Rev. D*, 94, 022002
- Bengaly C. A. P., Siewert T. M., Schwarz D. J., Maartens R., 2019, *MNRAS*, 486, 1350
- Bennett C. L. et al., 2003, *ApJS*, 148, 1
- Bernal J. L., Peacock J. A., 2018, *J. Cosmol. Astropart. Phys.*, 07, 002
- Bernal J. L., Raccanelli A., Kovetz E. D., Parkinson D., Norris R. P., Danforth G., Schmitt C., 2019, *J. Cosmol. Astropart. Phys.*, 02, 030
- Bernal J. L., Bellomo N., Raccanelli A., Verde L., 2020, *J. Cosmol. Astropart. Phys.*, 10, 017
- Blake C., Wall J., 2002, *Nature*, 416, 150
- Boldt E., 1987, *Phys. Rep.*, 146, 215
- Bonaldi A., Bonato M., Galluzzi V., Harrison I., Massardi M., Kay S., De Zotti G., Brown M. L., 2019, *MNRAS*, 482, 2
- Boughn S., Crittenden R., 2004, *Nature*, 427, 45
- Boughn S., Crittenden R., 2005, *New Astron. Rev.*, 49, 75
- Cabre A., Gaztanaga E., Manera M., Fosalba P., Castander F., 2006, *MNRAS*, 372, L23
- Camera S., Santos M. G., Bacon D. J., Jarvis M. J., McAlpine K., Norris R. P., Raccanelli A., Röttgering H., 2012, *MNRAS*, 427, 2079
- Camera S. et al., 2015, *PoS*, AASKA14, 025
- Cardoso J.-F., Le Jeune M., Delabrouille J., Betoule M., Patanchon G., 2008, *IEEE J. Sel. Top. Signal Process.*, 2, 735
- Chen S., Schwarz D. J., 2016, *A&A*, 591, A135
- Cole S., Kaiser N., 1989, *MNRAS*, 237, 1127
- 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
- Crittenden R. G., Turok N., 1996, *Phys. Rev. Lett.*, 76, 575
- de Gasperin F., Intema H. T., Frail D. A., 2018, *MNRAS*, 474, 5008
- Delabrouille J., Cardoso J. F., Patanchon G., 2003, *MNRAS*, 346, 1089
- Diemer B., 2018, *ApJS*, 239, 35
- Dolfi A., Branchini E., Bilicki M., Balaguera-Antolínez A., Prandoni I., Pandit R., 2019, *A&A*, 623, A148
- Dong F., Yu Y., Zhang J., Yang X., Zhang P., 2021, *MNRAS*, 500, 3838
- Dupé F. X., Rassat A., Starck J. L., Fadili M. J., 2011, *A&A*, 534, A51
- Ellis G. F. R., Baldwin J. E., 1984, *MNRAS*, 206, 377
- European Space Agency, 2018, PR3 Legacy CMB Maps, Version PR2 and PR3. ESA, Noordwijk
- Ferramacho L. D., Santos M. G., Jarvis M. J., Camera S., 2014, *MNRAS*, 442, 2511
- Ferraro S., Sherwin B. D., Spergel D. N., 2015, *Phys. Rev. D*, 91, 083533
- Fonseca J., Maartens R., Santos M. G., 2017, *MNRAS*, 466, 2780
- Fosalba P., Gaztanaga E., 2004, *MNRAS*, 350, L37
- Fosalba P., Gaztanaga E., Castander F., 2003, *ApJ*, 597, L89
- Friedman J., Hastie T., Tibshirani R., 2007, *Biostatistics*, 9, 432
- Giannantonio T., Scranton R., Crittenden R. G., Nichol R. C., Boughn S. P., Myers A. D., Richards G. T., 2008, *Phys. Rev. D*, 77, 123520
- Gibelyou C., Huterer D., 2012, *MNRAS*, 427, 1994
- 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., 2022, *MNRAS*, 512, 6104
- Hale C. L., Jarvis M. J., Delvecchio I., Hatfield P. W., Novak M., Smolcic V., Zamorani G., 2018, *MNRAS*, 474, 4133
- Hale C. L. et al., 2021, *Publ. Astron. Soc. Aust.*, 38, e058
- Hartlap J., Simon P., Schneider P., 2007, *A&A*, 464, 399
- Hinton S., Adams C., Badger C., 2020, Samreay/ChainConsumer v0.33.0. Zenodo. Available at: <https://doi.org/10.5281/zenodo.4280904>
- Hivon E., Górski K. M., Netterfield C. B., Crill B. P., Prunet S., Hansen F., 2002, *ApJ*, 567, 2
- Hobson M. P., Bridle S. L., Lahav O., 2002, *MNRAS*, 335, 377
- Hotan A. W. et al., 2021, *Publ. Astron. Soc. Aust.*, 38, e009
- Howlett C., Lewis A., Hall A., Challinor A., 2012, *J. Cosmol. Astropart. Phys.*, 1204, 027
- Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
- Johnston S., Wall J., 2008, *Exp. Astron.*, 22, 151
- Karamanis M., Beutler F., 2020, *Stat. Comput.*, 31, 61
- Karamanis M., Beutler F., Peacock J. A., 2021, *MNRAS*, 508, 3589
- Kaufman G. M., 1967, Center for Operations Research and Econometrics, Catholic University of Louvain, Heverlee, Report No. 6710
- Land K., Magueijo J., 2005, *Phys. Rev. Lett.*, 95, 071301
- Lewis A., Challinor A., Lasenby A., 2000, *ApJ*, 538, 473
- Limber D. N., 1953, *ApJ*, 117, 134
- McConnell D. et al., 2020, *Publ. Astron. Soc. Aust.*, 37, e048
- Magliocchetti M., Popesso P., Brusa M., Salvato M., Laigle C., McCracken H. J., Ilbert O., 2017, *MNRAS*, 464, 3271
- Martinelli M., Dalal R., Majidi F., Akrami Y., Camera S., Sellentin E., 2022, *MNRAS*, 510, 1964
- Mohan N., Rafferty D., 2015, Astrophysics Source Code Library, record ascl:1502.007
- Neal R. M., 2003, *Ann. Stat.*, 31, 705
- Norberg P., Baugh C. M., Gaztanaga E., Croton D. J., 2009, *MNRAS*, 396, 19
- Norris R. P. et al., 2011, *Publ. Astron. Soc. Aust.*, 28, 215
- Norris R. P. et al., 2021, *Publ. Astron. Soc. Aust.*, 38, e046
- Oliphant T., 2006, Guide to NumPy. Continuum Press, London and New York City

- Pedregosa F. et al., 2011, *J. Machine Learning Res.*, 12, 2825
 Percival W. J., Brown M. L., 2006, *MNRAS*, 372, 1104
 Planck Collaboration et al., 2016, *A&A*, 594, A21
 Price-Whelan A. M. et al., 2018, *AJ*, 156, 123
 Raccanelli A., Bonaldi A., Negrello M., Matarrese S., Tormen G., De Zotti G., 2008, *MNRAS*, 386, 2161
 Raccanelli A. et al., 2015, *J. Cosmol. Astropart. Phys.*, 01, 042
 Rubart M., Schwarz D. J., 2013, *A&A*, 555, A117
 Sachs R. K., Wolfe A. M., 1967, *ApJ*, 147, 73
 Shao J., 1986, *Ann. Stat.*, 14, 1322
 Siewert T. M. et al., 2020, *A&A*, 643, A100
 Siewert T. M., Schmidt-Rubart M., Schwarz D. J., 2021, *A&A*, 653, A9
 Singal A. K., 2011, *ApJ*, 742, L23
 Smolčić V. et al., 2017, *A&A*, 597, A4
 Tanidis K., Camera S., Parkinson D., 2020, *MNRAS*, 491, 4869
 Tinker J. L., Robertson B. E., Kravtsov A. V., Klypin A., Warren M. S., Yepes G., Gottlöber S., 2010, *ApJ*, 724, 878
 Tiwari P., Jain P., 2015, *MNRAS*, 447, 2658
 Tiwari P., Nusser A., 2016, *J. Cosmol. Astropart. Phys.*, 1603, 062
 Tiwari P., Kothari R., Naskar A., Nadkarni-Ghosh S., Jain P., 2014, *Astropart. Phys.*, 61, 1
 Verde L. et al., 2003, *ApJS*, 148, 195
 Virtanen P. et al., 2020, *Nat. Methods*, 17, 261
 Wilman R. J. et al., 2008, *MNRAS*, 388, 1335
 Wilman R. J., Jarvis M. J., Mauch T., Rawlings S., Hickey S., 2010, *MNRAS*, 405, 447
 Xavier H. S., Abdalla F. B., Joachimi B., 2016, *MNRAS*, 459, 3693
 Xia J.-Q., Viel M., Baccigalupi C., De Zotti G., Matarrese S., Verde L., 2010, *ApJ*, 717, L17
 Yamauchi D., Takahashi K., Oguri M., 2014, *Phys. Rev. D*, 90, 083520
 Zonca A., Singer L., Lenz D., Reinecke M., Rosset C., Hivon E., Gorski K., 2019, *J. Open Source Softw.*, 4, 1298

APPENDIX A: IS THERE AN ‘AXIS OF EVIL’?

We have omitted multipoles at $\ell \leq 40$ in our C_{ℓ}^{gg} analyses due to

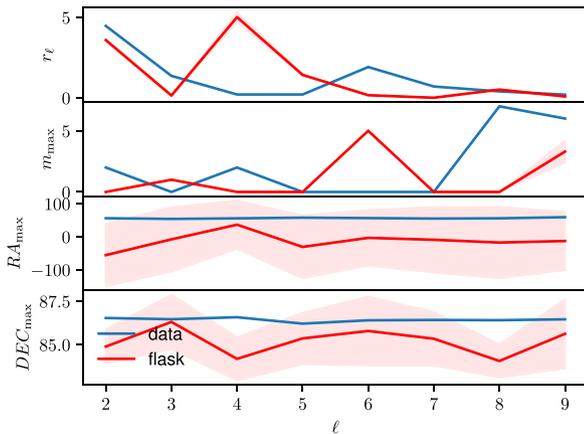


Figure A1. The ratio r_{ℓ} of power absorbed by the maximum mode with ‘shape’ m_{max} in direction $\mathbf{n}_{\text{max}} = (\text{Dec}_{\text{max}}, \text{RA}_{\text{max}})$. The blue line has been obtained from the data. The magenta line shows the mean of 10 flask realizations, whereas the shaded region displays the 1σ region estimated from the same 10 realizations.

an excess in the power spectrum that cannot be described by our model. Whilst we believe this is due to systematics, one possible cause of this excess power might also be a large anisotropy in the distribution of continuum galaxies. It would be possible to test this explanation in harmonic space, by conducting a similar analysis

to that which detected a preferred axis of the cosmic microwave background anisotropy (Land & Magueijo 2005). We consider the ratio

$$r_{\ell} = \max_{mn} \frac{C_{\ell m}}{(2\ell + 1)C_{\ell}} \quad (\text{A1})$$

of power absorbed by the maximum mode with ‘shape’ m_{max} in direction \mathbf{n}_{max} , where

$$C_{\ell m} = \begin{cases} |a_{\ell m}|^2, & \text{for } \ell = 0 \\ 2|a_{\ell m}|^2, & \text{else} \end{cases}, \quad (\text{A2})$$

and which we plot in Fig. A1. We compute r_{ℓ} for the data and for 10 flask realizations. Surprisingly, the direction \mathbf{n}_{ℓ} of the data is within the 1σ bounds of the flask directions, but the data and flask r_{ℓ} are discrepant. In any case, the data r_{ℓ} have the same magnitude as the ones we see in the flask realizations. The fact that the flask realizations show almost no scatter at most multipoles suggests that this is mostly driven by the mask. This is further supported by the ‘axis of evil’ pointing towards a direction close to the North Pole, around which the RACS mask is almost symmetric (cf. Fig. 1).

APPENDIX B: MEASURING ISW EXCLUDING LARGE-SCALE MULTIPOLES IN BOTH THE GALAXY-GALAXY AUTO POWER SPECTRUM AND GALAXY-TEMPERATURE CROSS POWER SPECTRUM

In Section 4.2, we have argued that, even though we ignore large-

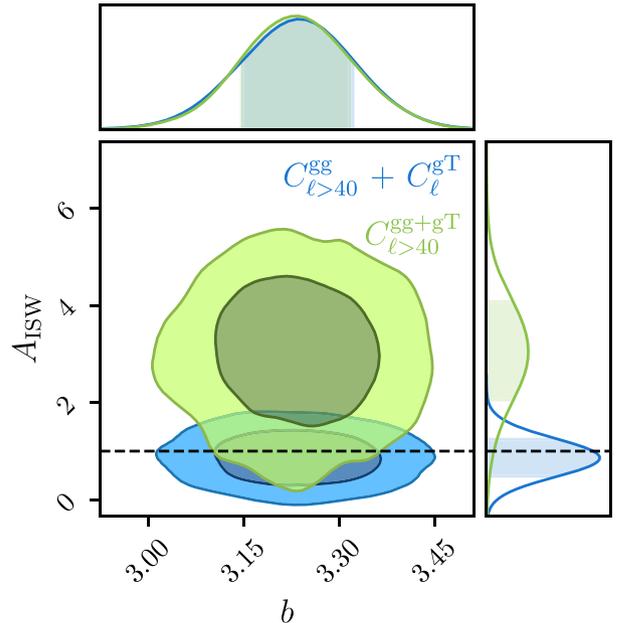


Figure B1. Comparison of the b - A_{ISW} contours from a combined C_{ℓ}^{gg} and C_{ℓ}^{gT} fit. In both cases, we omit large-scale multipoles ($\ell < 40$) in the galaxy-galaxy autopower spectrum. For the blue contour, we use the full available multipole range in the galaxy-temperature cross power spectrum, whereas for the green contour, we also omit $\ell < 40$ in the gT spectrum. The dashed line marks the expectation of $A_{\text{ISW}} = 1$.

scale C_{ℓ}^{gg} -multipoles, we can still trust large-scale C_{ℓ}^{gT} -multipoles. We present in Fig. B1 a combined measurement of a constant bias b

and A_{ISW} where we omit $\ell < 40$ in both C_ℓ^{gg} and C_ℓ^{gT} . Without the first two ℓ -bins, Fig. 7 shows that the ISW signal is only distinguishable from the null hypothesis of no galaxy–temperature correlation in the third and fourth ℓ -bin. We can also see there that the third ℓ -bin has a larger value of C_ℓ^{gT} than expected from neighbouring values. Without the first two ℓ -bins, we therefore see an increased value of $A_{\text{ISW}} = 3.09_{-1.03}^{+0.99}$. Thus, using the full ℓ -range is actually more conservative since cutting out large-scale multipoles pushes the significance of the ISW detection up to 3σ . As can also be seen in Fig. B1, the bias is unaffected by the large-scale gT power.

APPENDIX C: ISW CONSTRAINTS INCLUDING LARGE-SCALE MULTIPOLES

In this Appendix, we present the results using the full available

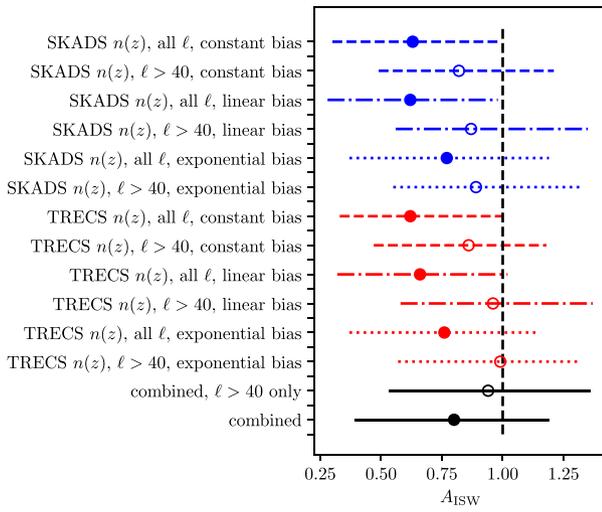


Figure C1. Boxplot summarizing the A_{ISW} results obtained using different $n(z)$, ℓ ranges, and bias parametrizations. Including all of the data, including the C_ℓ^{gg} for $\ell \leq 40$ (the solid circles) lowers the mean value of A_{ISW} by approximately 0.5σ , in comparison to estimates where these data are left out (empty circle). This potential systematic bias is lessened when the different models are combined using the BACCUS approach (black line).

multipole range also in C_ℓ^{gg} . In line with Fig. 10, for all bias and $n(z)$ models, larger values of the galaxy bias $b(z)$ are favoured when

including multipoles at $\ell \leq 40$ as the bias is the only parameter we vary in our C_ℓ^{gg} model. As the amplitude of $b(z)$ is degenerate with A_{ISW} in C_ℓ^{gT} , the full-range analysis hence supports lower values of A_{ISW} , which we present in Fig. C1.

As can be seen there, even when we consider the ℓ -range where our C_ℓ^{gg} model breaks down, there is no set of $b(z)$ and $n(z)$ model where the data are consistent with $A_{\text{ISW}} = 0$. This result is not unexpected, given that we believe the $\ell \leq 40$ autopower spectrum excess to be a systematic, and so uncorrelated with the CMB photon distribution.

We repeat the BACCUS-like analysis described in Section 4.5. The combined posterior distribution of A_{ISW} is graphed in Fig. C2. We measure $A_{\text{ISW}} = 0.75 \pm 0.43$; thus, even when we include the $\ell \leq 40$ -range, the data are in favour of the ISW effect with 1.7σ . It has to be noted, however, that the Variable Gaussian distribution does not summarize the individual A_{ISW} posteriors well and that the

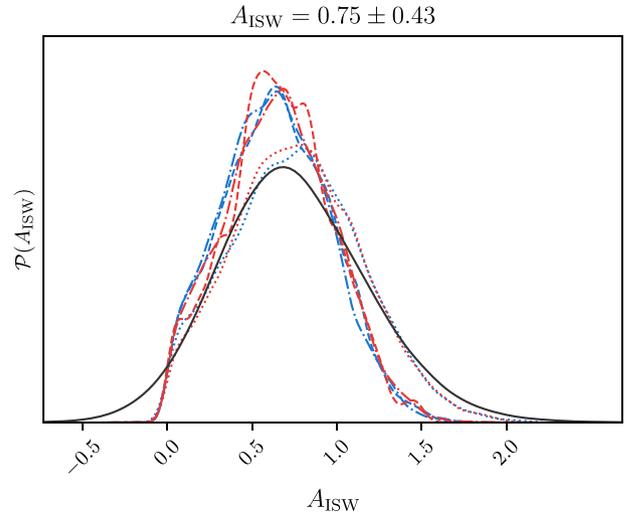


Figure C2. The posterior on A_{ISW} after combining the measurements presented in Fig. C1 in a BACCUS-like (Bernal & Peacock 2018) fashion.

true low- A_{ISW} tails are less pronounced than those of the Gaussian approximations (cf. Fig. C2). Our estimate of 1.7σ is therefore conservative.

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