

Self-consistent secondary cosmic microwave background anisotropies and extragalactic foregrounds in the FLAMINGO simulations

Tianyi Yang,^{1*} Ian G. McCarthy,^{1†} Fiona McCarthy,^{2,3} Boris Bolliet,^{3,4} Jens Chluba,⁵ William Coulton,^{2,3} John C. Helly,⁶ Matthieu Schaller,^{7,8} Joop Schaye⁷

¹*Astrophysics Research Institute, Liverpool John Moores University, Liverpool, L3 5RF, UK*

²*DAMTP, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, UK*

³*Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge CB3 0HA, UK*

⁴*Cavendish Astrophysics, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK*

⁵*Jodrell Bank Centre for Astrophysics, Alan Turing Building, University of Manchester, Manchester M13 9PL*

⁶*Institute for Computational Cosmology, Department of Physics, Durham University, South Road, Durham, DH1 3LE, United Kingdom*

⁷*Leiden Observatory, Leiden University, PO Box 9513, 2300 RA Leiden, the Netherlands*

⁸*Lorentz Institute for Theoretical Physics, Leiden University, PO box 9506, NL-2300 RA Leiden, the Netherlands*

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

Secondary anisotropies in the cosmic microwave background (CMB) contain information that can be used to test both cosmological models and models of galaxy formation. Starting from lightcone-based HEALPIX maps and catalogues, we present a new set of mock CMB maps constructed in a self-consistent manner from the FLAMINGO suite of cosmological hydrodynamical simulations, including CMB lensing, thermal and kinetic Sunyaev-Zel’dovich effects, cosmic infrared background, radio point source and anisotropic screening maps. We show that these simulations reproduce a wide range of observational constraints. We also compare our simulations with previous predictions based on dark matter-only simulations which generally model the secondary anisotropies independently from one another, concluding that our hydrodynamical simulation mocks perform at least as well as previous mocks in matching the observations whilst retaining self-consistency in the predictions of the different components. Using the model variations in FLAMINGO, we further explore how the signals depend on cosmology and feedback modelling, and we predict cross-correlations between some of the signals that differ significantly from those in previous mocks. The mock CMB maps should provide a valuable resource for exploring correlations between different secondary anisotropies and other large-scale structure tracers, and can be applied to forecasts for upcoming surveys.

Key words: cosmology: cosmic background radiation – cosmology: large-scale structure of Universe – galaxies: general – methods: numerical

1 INTRODUCTION

The cosmic microwave background (CMB) is a relic radiation field from the early Universe. During the epoch of recombination, photons decoupled from matter and began to propagate freely, leaving imprints of temperature fluctuations at the surface of last scattering. These fluctuations give rise to the primary anisotropies of the CMB, which contain information about the initial conditions and key fundamental parameters of the Universe. Observations of these anisotropies provide strong support for the standard cosmologi-

cal model, characterised by cold dark matter and a cosmological constant—the so-called Λ CDM model (e.g. [Planck Collaboration et al. 2020c](#); [Louis et al. 2025](#)).

As CMB photons travel through space following decoupling, they interact with intervening large-scale structure (LSS), which causes a range of secondary anisotropies in CMB observations. These include (but are not limited to): gravitational lensing of the CMB photons, which is sensitive to the total mass distribution along the line of sight; the Sunyaev–Zel’dovich (SZ) effects that probe the thermal energy density (thermal SZ, or tSZ) and the momentum (kinetic SZ, or kSZ) of the ionised gas; and so-called late-time anisotropic (or ‘patchy’) screening due to Thomson scattering off ionised electrons in and/or around haloes. In addition to CMB pho-

* E-mail: t.yang@ljamu.ac.uk

† E-mail: i.g.mccarthy@ljamu.ac.uk

tons, other sources of emission are present at radio and microwave wavelengths that can act as contaminants for CMB measurements. These include the so-called Cosmic Infrared Background (CIB), which probes the cosmic evolution of dusty star-forming galaxies, and radio point sources that reflect compact non-thermal emission from active galactic nuclei and other extragalactic populations. As the CIB and radio sources are associated with galaxies and therefore also trace LSS, they are expected to (and indeed are observed to) spatially correlate with CMB secondary anisotropies. Together, the primary CMB anisotropies and foregrounds, their secondary effects, and their cross-correlations with late-time tracers such as galaxies and weak gravitational lensing, provide powerful and complementary tools of cosmology and the astrophysical processes related to structure formation (e.g. [Planck Collaboration et al. 2014, 2016b,c](#); [Hojjati et al. 2017](#); [Planck Collaboration et al. 2018, 2020d](#); [Chiang et al. 2020](#); [McCarthy et al. 2023a](#); [McCarthy & Hill 2024b](#)).

Extracting a component of interest from CMB observations (e.g., the tSZ effect) is a challenging task owing to the finite sensitivity, resolution, and wavelength coverage of the observations coupled with uncertainties in the spectral characteristics of the CIB and radio sources (and their potential dependence on galaxy properties and their evolution) as well as uncertainties in the predicted spatial clustering of both these and the secondary CMB anisotropies and their evolution. To disentangle these contributions, several component separation algorithms have been developed, such as the COMMANDER method ([Eriksen et al. 2006](#)), which relies on a physical model of the sky, requiring both spatial and frequency-dependent templates for each component, and other more empirical methods such as the ILC, NILC, GILC, SEVEM, SMICA, and MILCA algorithms (see e.g. [Hurier et al. 2013](#); [Planck Collaboration et al. 2020b](#), and references therein), which rely primarily on just the spectral dependence of the target signal. Unsurprisingly, inaccurate assumptions employed in the component separation algorithms can result in biases in the extracted signals. For instance, [Efstathiou & McCarthy \(2025\)](#) demonstrated that the level of foreground leakage in the recovered tSZ power spectrum can vary significantly depending on the deprojection strategy used in the NILC algorithm.

To interpret observations, realistic multi-component CMB simulations are essential for predicting how different astrophysical and cosmological signals are correlated. Such mocks are also an invaluable tool for testing how cleanly the signals can be separated from one another in realistic observing conditions, as well for testing routines for source detection (e.g., cluster finding).

A number of previous studies have produced CMB secondary predictions based on the halo model formalism (e.g., `CLASS_SZ`¹; [Bolliet et al. 2024](#)) or dark matter-only cosmological simulations, which typically use parametric prescriptions for ‘painting on’ observable components such as the Sunyaev-Zel’dovich effect. We will mostly focus our comparison on the latter class of predictions, as they are more similar in spirit to our aims of producing full-sky maps that can be analysed in a way that is faithful to the observational analyses.

In a pioneering study producing the first CMB mocks based on dark matter-only simulations, [Sehgal et al. \(2010\)](#) started with a moderate resolution $1000 h^{-1}$ Mpc volume dark matter particle field, with haloes identified using a friends-of-friends (FoF) algorithm. The free parameters that govern the gas physics and CIB source counts were calibrated against existing external observations, including X-ray gas fractions and far-infrared measurements

available at that time. Updated models using this framework have been used for Simons Observatory forecasting (see [Ade et al. 2019](#)).

More recently, [Stein et al. \(2020\)](#) produced the WebSky CMB simulations² (see also [Li et al. 2022](#) for associated radio catalogues construction) that use the mass-Peak Patch formalism ([Bond & Myers 1996](#); [Stein et al. 2019](#)) to predict the growth of structure and the location of haloes. While this formalism approximates non-linear evolution (as opposed to full cosmological simulations, which directly evaluate it), a major strength is its speed and ability to follow large simulation volumes that better sample rare, massive haloes. Baryons within haloes are assigned using analytical prescriptions calibrated against hydrodynamical simulations and observations, while the dynamics and matter distribution of the diffuse field component are modelled using second-order Lagrangian perturbation theory (2LPT, [Bouchet et al. 1995](#)).

More similar to the approach of [Sehgal et al. \(2010\)](#), the AGORA CMB simulations³ ([Omori 2024](#)) are subarcminute-resolution CMB sky maps based on the higher-resolution MultiDark *Planck* 2 simulation (MDPL2, [Klypin et al. 2016](#)). Extragalactic observables are assigned using observationally constrained analytical models, and the resulting mock signals have been validated against recent CMB observations from *Planck* ([The Planck Collaboration 2006](#)), SPT-SZ ([Carlstrom et al. 2011](#)), and SPTpol ([Bleem et al. 2012](#)).

While dark matter-only and halo model-based simulations have proven valuable for constructing CMB secondary anisotropy maps, they rely on simplified or empirically tuned models to assign baryonic properties and in general there is no physical association between the components (e.g., the CIB and tSZ effect) apart from the fact that their clustering is spatially correlated with haloes. In contrast, hydrodynamical simulations offer a self-consistent treatment of gas physics, naturally capturing the interplay between complex processes such as radiative cooling, star formation, AGN feedback, diffuse gas flows, as well as their gravitational back reaction on the dark matter. Hydrodynamical simulations also have important uncertainties, since they cannot resolve all relevant scales or the predict from first principles the efficiencies of important processes such as feedback, thus requiring so-called subgrid models (see the discussion [Schaye et al. 2015](#)). Nevertheless, for a given feedback implementation the baryons (hot and cold gas and stars), dark matter, black holes, and neutrinos (if the latter are included) are followed in a self-consistent manner, which in general will result in correlations between the components. For example, relatively weak feedback might be expected to lead to higher gas fractions in groups and clusters but also to elevated star formation rates, potentially producing a much stronger cross-correlation between the CIB and the hot gas compared to a model where the feedback is considerably stronger. Furthermore, one can study the sensitivity of the results to the subgrid modelling by systematically varying the key parameters of the models, as we do here (see also, e.g., [Schaye et al. 2010](#); [Terrazas et al. 2020](#); [Ni et al. 2023](#); [Salcido et al. 2023](#); [Sorini et al. 2024](#)). In addition, environmental processes such as ram pressure stripping of gas from satellites are also naturally captured, as are effects which are challenging to paint on dark matter-only simulations, including the effects of substructure, triaxiality, mis-centring,

² WebSky maps are available at: https://lambda.gsfc.nasa.gov/simulation/mocks_data.html

³ AGORA maps are available at: <https://yomori.github.io/agora/index.html>

¹ <https://github.com/CLASS-SZ>

etc. on observable baryonic quantities (e.g., effects of triaxiality on the tSZ effect).

Individual secondary observables have been studied within the setups of several hydrodynamical simulation suites, including the SZ effect (e.g. Battaglia et al. 2012; Vogelsberger et al. 2014; McCarthy et al. 2014; Dolag et al. 2016; McCarthy et al. 2017; Yang et al. 2022) and the infrared source reconstruction (e.g. Hayward et al. 2013; Béthermin et al. 2017; Lovell et al. 2021; Cochrane et al. 2023). However, to our knowledge, none of these simulations have been used for producing full-sky, multi-component CMB maps, which in some cases is due to their limited box sizes, which restricts the ability to sample the large-scale modes and rare massive objects that are important for reproducing the statistical properties of full-sky signals such as the tSZ effect.

In this study, we present a set of full-sky CMB secondary anisotropy mock maps based on the FLAMINGO suite of cosmological hydrodynamical simulations (Schaye et al. 2023; Kugel et al. 2023), where the signals are directly derived from the physical properties of matter, gas, and accreting black holes within the simulation volume. FLAMINGO features a large fiducial run with a comoving box size of 2.8 Gpc. In addition, the suite includes multiple 1 Gpc runs with systematically varied subgrid feedback models and cosmological parameters, allowing us to further explore how different baryonic and cosmological assumptions affect the predicted signals and their correlations. This setup provides, for the first time, a powerful framework to study the signal correlations and their model dependencies using large-volume hydrodynamical simulations, which could bridge the gap between theoretical predictions and forthcoming high-resolution CMB observations.

This work presents a description of our methodology, simulated maps, and corresponding power spectra, along with initial comparisons to observational data and to previous mock studies based on dark matter-only simulations. In future work, we will extend this framework to evaluate the performance of component-separation algorithms and to conduct a wider range of comparisons with observations from current and forthcoming CMB experiments.

This paper is organised as follows. In Section 2, we briefly summarise the setup of the FLAMINGO fiducial run and its model variants used for map construction. In Section 3, we detail the procedures and models adopted to construct the maps of different CMB secondaries. Section 4 presents the auto- and cross-power spectra of various quantities from our mock maps, along with their comparisons to other multi-component CMB simulations and observations. In Section 5, we investigate the feedback and cosmology dependencies of these power spectra results. Finally, we conclude in Section 6.

2 FLAMINGO SIMULATIONS

We provide here a brief summary of the FLAMINGO simulations. A detailed description of the simulation is presented in Schaye et al. (2023).

FLAMINGO is a suite of large-scale cosmological hydrodynamical simulations designed to study cosmology and LSS physics. The suite includes three different baryonic gas particle mass resolutions: a high-resolution run with $m_{\text{gas}} = 1.3 \times 10^8 M_{\odot}$ (referred to as m8), an intermediate-resolution run with $m_{\text{gas}} = 1.1 \times 10^9 M_{\odot}$ (m9), and a low-resolution run with $8.6 \times 10^9 M_{\odot}$ (m10). The flagship runs are the $(1 \text{ Gpc})^3$ high-resolution run and the $(2.8 \text{ Gpc})^3$ intermediate-resolution run (denoted as L1_m8 and L2p8_m9 respectively). The latter follows 2.8×10^{11} particles, making it the

largest hydrodynamical simulation evolved to $z = 0$ at the time it was run⁴. Most runs adopt a Dark Energy Survey Year Three (DES Y3; Abbott et al. 2022) $3 \times 2\text{pt} + \text{All Ext. } \Lambda\text{CDM}$ cosmology (see Table 1 for the list of parameters). This cosmology assumes a spatially flat universe and is based on a combination of constraints from three DES Y3 two-point correlation functions—cosmic shear, galaxy clustering, and galaxy–galaxy lensing—along with external data from baryon acoustic oscillations (BAO), redshift-space distortions, Type Ia supernovae, *Planck* observations of the CMB (including CMB lensing), Big Bang nucleosynthesis, and local measurements of the Hubble constant. See Abbott et al. (2022) for details.

The simulations were performed using SWIFT (Schaller et al. 2024), a fully open-source coupled cosmology, gravity, hydrodynamics, and galaxy formation code. Short- and long-range gravitational forces are computed using a 4th-order fast multipole method (Greengard & Rokhlin 1987; Cheng et al. 1999; Dehnen 2014) and a particle-mesh method solved in Fourier space, respectively, following the force-splitting approach of Bagla & Ray (2003). The hydrodynamic equations are solved using the smoothed particle hydrodynamics (SPH) method (for a review, see Price 2012), in particular the SPHENIX flavour of SPH (Borrow et al. 2022), which was designed specifically for simulations of galaxy formation. The initial conditions are obtained from a modified version of monofonIC (Hahn et al. 2021; Elbers et al. 2022), and neutrinos are implemented with the δf method (Elbers et al. 2021).

Unresolved physical processes are modelled using subgrid physics with various parameter choices. The simulation includes a wide range of physical processes, including radiative cooling and heating (Ploekinger & Schaye 2020), star formation and evolution (Schaye & Dalla Vecchia 2008; Wiersma et al. 2009), black hole growth (Booth & Schaye 2009; Bahé et al. 2022), feedback from young stars and supernovae (Chaikin et al. 2022, 2023), and AGN feedback (Booth & Schaye 2009; Huško et al. 2022). A detailed description of the subgrid physics implementation is provided in Schaye et al. (2023).

Four subgrid parameters—two related to stellar feedback, one to black hole growth, and one to AGN feedback—are calibrated to match the observed present-day galaxy stellar mass function (SMF, Driver et al. 2022) and the low-redshift gas mass fraction within R_{500c} for galaxy groups and clusters (Akino et al. 2022; Kugel et al. 2023). Machine learning-based emulators were used in the calibration (Kugel et al. 2023). Note that the emulators were used not only to calibrate to the observations as in the fiducial model, but also to produce different model variations as discussed below.

The goal of our study is to provide multicomponent full-sky CMB mock maps predicted from hydrodynamical simulations. We use the $(2.8 \text{ Gpc})^3$ intermediate-resolution run as our primary model, referred to as the fiducial run. To further explore the feedback dependence of the correlation between different secondary anisotropies, we also make use of feedback model variants from a series of $(1 \text{ Gpc})^3$ runs that vary the strength of stellar and/or AGN feedback. These runs share the same particle resolution as the large-box fiducial model, but apply shifts to the observed galaxy stellar masses (for M_*) or cluster gas fractions (for f_{gas}) during calibration. In the strong stellar feedback model (denoted as $M_* - \sigma$), the observed stellar mass function (SMF) was shifted by -0.14 dex lower in stellar mass. The AGN feedback variants (denoted as $f_{\text{gas}} \pm N\sigma$) were calibrated to the observed gas mass fractions shifted by $+2$,

⁴ See the latest, larger Frontier-E simulation (Frontiere et al. 2025)

-2, -4, -8 times the measurement error, respectively. An additional jets and strong jets model are also included (denoted as Jet and Jet_ $f_{\text{gas}} - 4\sigma$, respectively), where AGN feedback is implemented via kinetic kicks to gas particles within the black hole’s SPH kernel, rather than through isotropic thermal energy injection. The jet model is calibrated against the same set of observational data as the fiducial thermal AGN feedback model. All feedback-variation runs are evolved under the fiducial DES Y3 cosmology.

The FLAMINGO simulations have previously been compared with a wide range of observations beyond the measurements used for calibration, including scaling relations between galaxy stellar mass and black hole mass, halo mass, quenched fractions, the cosmic star formation rate vs. redshift (Schaye et al. 2023), cluster X-ray and tSZ scaling relations and profiles (Schaye et al. 2023; Braspenning et al. 2024), the stacked kSZ effect profiles of massive galaxies (McCarthy et al. 2025; Siegel et al. 2026; Bigwood et al. 2026), the auto and cross-power spectra of CMB lensing, cosmic shear, and the tSZ effect (McCarthy et al. 2023b), the X-ray–cosmic shear cross-power spectra (McDonald et al. 2026), among others.

To study the cosmological dependence of those signals, we also include two cosmology variants from the $(1 \text{ Gpc})^3$ runs (see Table 1 for the list of parameters). One uses the Planck Collaboration et al. (2020c) maximum likelihood cosmology, denoted as Planck, while the other takes the lensing cosmology from Amon et al. (2023, denoted as LS8). This LS8 model has a lower amplitude of the power spectrum, $S_8 = 0.766$, compared with 0.815 and 0.833 for the fiducial and Planck runs, and provides a better match to some previous cosmic shear measurements, including from DES Y3 (Abbott et al. 2022), HSC Y1 (Aihara et al. 2018), and KiDS-1000 (Kuijken et al. 2019). All cosmology-variation runs are evolved using the same calibrated feedback model as the fiducial run.

3 METHODOLOGY

In this section, we describe our methodology for producing maps of the various CMB secondary anisotropies, based on the lightcone maps from the FLAMINGO simulations. Section 3.1 provides a brief summary of the construction of lightcone maps and particle lightcone data in the simulation. In Section 3.2, we summarise the modelling of the stacked CMB lensing convergence field. Section 3.3 discusses the thermal and kinetic SZ effects, including the modelling of their relativistic corrections. Section 3.4 provides a brief introduction to the optical depth map construction, which is directly related to the patchy screening effect. Sections 3.5 and 3.6 describe the construction of the CIB map and mock catalogues for radio point sources. In these sections, we also highlight some key differences between our models and those used in other mock CMB simulations, such as AGORA and WebSky. Finally, in Section 3.7, we outline how lensing effects are added to each simulated quantity. Relevant maps and their power spectra are discussed in Section 4.

3.1 HEALPix lightcone maps and halo lightcone catalogues

Here we provide a brief summary of the lightcone data generation. A detailed description of the algorithm can be found in the appendix of Schaye et al. (2023).

To produce the HEALPix (Górski et al. 2005) maps for each quantity, an observer’s past lightcone is split into a set of concentric spherical shells, with a redshift interval of $\Delta z = 0.05$ from $z = 0$ to 3 and $\Delta z = 0.25$ for higher redshifts. Whenever a particle is found to have crossed the lightcone, we determine which shell it lies in at

the time of crossing and accumulate the particle’s contributions to the HEALPix maps for that shell. In this study, we mainly focused on the full-sky maps of total matter, Compton y , Doppler B (for the kSZ), optical depth (for anisotropic screening), star formation rate (for the CIB), and black holes (for radio point sources).

To produce the halo lightcone data, structure finding is performed in post-processing on the snapshot particle data using HBTHERONS, which is a modified version of the HBT algorithm (Han et al. 2012, 2018; Forouhar Moreno et al. 2025). We first read in the snapshot subhalo catalogue corresponding to a given snapshot, then load a spherical shell from the black hole particle lightcone that spans the redshift range halfway to the previous snapshot and halfway to the next. For each halo in the snapshot, we identify a black hole particle to use as a tracer of that halo. Whenever the tracer particle ID appears among the particles read from the black hole particle lightcone, we place a copy of the halo at that position. This process is then repeated for every snapshot to build the full halo lightcone catalogue.

In our study, the mock maps are integrated up to either $z = 4.5$ or $z = 3.0$, depending on the simulation box size. Such a redshift range is sufficient for CIB-related statistics (see Figure 9), which are the main focus of our study here. For the default fiducial $(2.8 \text{ Gpc})^3$ run, we use full-sky particle lightcone HEALPix maps per shell with $N_{\text{side}}=4096$, extending out to $z = 4.5$. For the feedback model and cosmology variants in $(1 \text{ Gpc})^3$ boxes, however, lightcone map outputs are only available up to $z = 3.0$. To account for any missing signal beyond $z = 3.0$ for the $(1 \text{ Gpc})^3$ runs, we rescale their power spectra using a redshift evolution correction with the ratio $C_{\ell, L2p8_m9 \ z<4.5} / C_{\ell, L2p8_m9 \ z<3.0}$. This ensures a fair comparison with the results obtained from the fiducial $(2.8 \text{ Gpc})^3$ run and that our summary statistics are fully converged in terms of redshift integration. Note that, as only the fiducial feedback model in the fiducial cosmology was run in the large volume, the correction factor that we apply is assumed to be independent of cosmology and feedback. In principle the correction factor ought to depend on both feedback and cosmology, but in practice we expect these dependencies to be small for the summary statistics that we examine. For radio sources, we integrate the halo lightcone data up to $z = 2.5$ only, which corresponds to the maximum redshift for which the adopted radio luminosity function is available.

We use the NAMASTER⁵ (Alonso et al. 2019) package to compute the auto- and cross-power spectra for various statistics. We initially use a multipole moment binning (bandpower) of $\Delta \ell = 7$, and then smooth the resulting curves using a Savitzky–Golay filter of order 3 with a window size of 15, which is only for visualization purposes. We deconvolve the $N_{\text{side}} = 4096$ pixel window function from the computed spectra using the *pixwin* function in the HEALPix package. To avoid repeating structures along the line of sight, the shells in the lightcone are rotated randomly every box-length interval, which preserves the line-of-sight correlations over the scales captured by the simulations. To preserve the correlations between the different signals (e.g., CIB and tSZ), the same rotations are applied to all quantities considered in this study.

3.2 CMB lensing maps

CMB lensing is caused by the deflection of CMB photons when they travel through the LSS of the Universe. As these photons prop-

⁵ <https://namaster.readthedocs.io/en/latest/source/installation.html>

Table 1. Cosmological parameters adopted in different FLAMINGO simulations. Only the models used in this work are listed here (for the full table, please see Schaye et al. 2023). These include the dimensionless Hubble constant, h ; the total matter density parameter, Ω_m ; the dark energy density parameter, Ω_Λ ; the baryonic matter density parameter, Ω_b ; the sum of the particle masses of the neutrino species, $\Sigma m_\nu c^2$; the amplitude of the primordial matter power spectrum, A_s ; the power-law index of the primordial matter power spectrum, n_s ; the amplitude of the initial power spectrum parametrised as the r.m.s. mass density fluctuation in spheres of radius $8 h^{-1}$ Mpc extrapolated to $z = 0$ using linear theory, σ_8 ; the amplitude of the initial power spectrum parametrised as $S_8 \equiv \sigma_8 \sqrt{\Omega_m/0.3}$; the neutrino matter density parameter, $\Omega_\nu \equiv \Sigma m_\nu c^2 / (93.14 h^2 \text{ eV})$. Note that the values of the Hubble and density parameters are given at $z = 0$.

Model name	h	Ω_m	Ω_Λ	Ω_b	$\Sigma m_\nu c^2$	A_s	n_s	σ_8	S_8	Ω_ν
Fiducial	0.681	0.306	0.694	0.0486	0.06 eV	2.099×10^{-9}	0.967	0.807	0.815	1.39×10^{-3}
Planck	0.673	0.316	0.684	0.0494	0.06 eV	2.101×10^{-9}	0.966	0.812	0.833	1.42×10^{-3}
LS8	0.682	0.305	0.695	0.0473	0.06 eV	1.836×10^{-9}	0.965	0.760	0.766	1.39×10^{-3}

agate from the last scattering surface to the observer, their paths are deflected by the intervening matter distribution, leading to small distortions in the observed CMB temperature and polarization images. This effect is characterised by a projected convergence field, κ , which can be used to infer the (CMB kernel-weighted) integrated mass distribution along the line of sight (see Lewis & Challinor 2006, for a review).

To produce the projected convergence map, we first compute the two-dimensional projected overdensity field $\delta(\chi, \theta)$ using the total mass HEALPIX output lightcone maps from the FLAMINGO simulation. Here, χ denotes the comoving radial distance along the line of sight, and θ represents the angular position on the sky. Therefore, each map corresponds to the projected overdensity within a thin shell at χ . We then integrate these overdensity maps (shells) along the line of sight up to $z = 4.5$, weighting them by the CMB lensing kernel to yield the total CMB lensing convergence map. This is expressed as:

$$\kappa_{\text{CMB}}(\theta) = \int_0^{\chi(z_{\text{max}})} W^{\kappa_{\text{CMB}}}(\chi) \delta(\chi, \theta) d\chi. \quad (1)$$

In constructing the convergence field we adopt the Born approximation, i.e. the lensing kernel is integrated along unperturbed photon paths. Previous studies have shown that, for CMB lensing and other weak-lensing-dominated probes, corrections from full ray-tracing are typically at the percent level and therefore subdominant for the statistics considered in this work (e.g. Giocoli et al. 2016; Fabbian et al. 2019; Hilbert et al. 2020; Broxterman et al. 2024). The CMB lensing kernel is written as:

$$W^{\kappa_{\text{CMB}}}(\chi) = \frac{3}{2} \left(\frac{H_0}{c} \right)^2 \Omega_m \frac{\chi}{a(\chi)} \left(1 - \frac{\chi}{\chi_{\text{CMB}}} \right), \quad (2)$$

with χ_{CMB} as the comoving distance to the last scattering surface (assumed to be at $z_{\text{CMB}} = 1100$), and $a(\chi)$ the scale factor at comoving distance χ . Although a non-negligible fraction of the κ_{CMB} signal might be expected from $z > 4.5$, given the broadness of the CMB lensing kernel, integration up to this redshift is sufficient for studying cross-correlations between the convergence field and the other fields we consider in this study (e.g., CIB, tSZ effect), which are the main summary statistics of interest in this study.

A detailed discussion of the κ_{CMB} auto power spectrum and the κ_{CMB} -cosmic shear power spectra can be found in McCarthy et al. (2023c). In summary, the κ_{CMB} statistics predicted by the FLAMINGO simulation show good convergence with respect to variations in box size, resolution and feedback models. A Planck-like cosmology yields excellent agreement between the predicted κ_{CMB} auto-power spectrum and measurements from Planck (Planck Collaboration et al. 2020d), SPTpol (Wu et al. 2019), and ACT DR6 (Qu et al. 2024). Under a Planck-like cosmology, the predicted amplitude of the κ_{CMB} -cosmic shear cross-power spectrum

is marginally higher than that measured between KiDS-1000 and Planck and ACT CMB lensing (Robertson et al. 2021), though this tension is not statistically significant. As we have previously considered the cross correlation between CMB lensing and cosmic shear and CMB lensing and Compton y (see Schaye et al. 2023), we focus our discussion on lensing-CIB cross-correlation, as outlined in the following sections.

3.3 Sunyaev-Zel'dovich effect maps

The Sunyaev-Zel'dovich (SZ) effect arises from the interaction between CMB photons, emitted at the last scattering surface, and intervening free electrons (Sunyaev & Zeldovich 1972, 1980; Mroczkowski et al. 2019, for a review). This effect can be primarily classified as the thermal SZ (tSZ) effect and the kinetic SZ (kSZ) effect.

The thermal SZ effect results from inverse Compton scattering of CMB photons off hot electrons around and within galaxy groups and clusters. Through this scattering process, photons gain energy, causing a distortion of the energy distribution compared to the original CMB radiation spectrum. This causes a decrease of CMB intensity at frequencies $\lesssim 217$ GHz and an increase at higher frequencies. The distortion is expressed as

$$\frac{\Delta T_{\text{tSZ}}}{T_{\text{CMB}}} = y f(\nu), \quad (3)$$

where the frequency dependence $f(\nu)$ is given by $x \coth(x/2) - 4$ with $x = h\nu/k_B T_{\text{CMB}}$. T_{CMB} is the CMB temperature and k_B is the Boltzmann constant. The magnitude of the distortion is characterised by the Compton parameter y , a dimensionless quantity that is related to the integral of gas pressure along the line of sight to the surface of last scattering:

$$y = \frac{\sigma_T}{m_e c^2} \int P_e dl \propto \int n_e T_e dl, \quad (4)$$

where σ_T is the cross-section for Thomson scattering and $m_e c^2$ is the electron rest energy. P_e is the electron pressure calculated as $n_e k_B T_e$, with n_e and T_e being the electron number density and temperature respectively.

The tSZ effect is linearly proportional to the product of gas density and temperature, and it is independent of redshift. These properties make it a powerful probe of the thermal state of the ionised gas over a wide range of spatial and mass scales, including not only massive clusters but also galaxy groups and diffuse components like the circumgalactic and intergalactic medium. Furthermore, the tSZ signal can be cross-correlated with other LSS tracers to provide joint constraints on the thermodynamic and baryonic evolution of the Universe (e.g., Vikram et al. 2017; Tröster et al. 2022; Battaglia et al. 2015; Koukoufilippas et al. 2020).

The kSZ effect, on the other hand, originates from the line-of-sight bulk motion of the ionised gas in the Universe, inducing a Doppler shift in the thermodynamic temperature of the CMB photons. Unlike the tSZ effect, the kSZ effect does not produce a spectral distortion but instead causes a frequency-independent shift in the CMB temperature. This shift is described by

$$\frac{\Delta T_{\text{kSZ}}}{T_{\text{CMB}}} = -\frac{\sigma_{\text{T}}}{c} \int_{\text{los}} \frac{d\chi}{1+z} n_e(\chi, z) (\mathbf{v}_{\text{e,p}} \cdot \hat{\mathbf{n}}), \quad (5)$$

where χ is the comoving distance to redshift z , $\mathbf{v}_{\text{e,p}} \cdot \hat{\mathbf{n}}$ is line-of-sight part of the peculiar velocity for the free electron field, and where $\hat{\mathbf{n}}$ is the line-of-sight unit vector pointing away from the observer.

All the above equations assume the SZ effect is induced by non-relativistic electrons. However, in galaxy clusters, where electron temperatures can exceed $T_e \gtrsim 0.1$ keV, relativistic effects become relevant and should be considered (e.g. [Challinor & Lasenby 1998](#)). As a result, the SZ spectral distortion becomes temperature, and velocity-dependent, and the standard frequency kernel $f(\nu)$ must be replaced with a more general relativistic kernel. In the case of the thermal SZ effect⁶, the relativistically corrected fluctuation is given by

$$\Delta I_{\text{rSZ}}(\hat{\mathbf{n}}, \nu, T_e(\hat{\mathbf{n}})) = f^{\text{rel}}(\nu, T_e(\hat{\mathbf{n}})) y(\hat{\mathbf{n}}), \quad (6)$$

where $f^{\text{rel}}(\nu, T_e(\hat{\mathbf{n}}))$ is the relativistic frequency response function, which depends both on frequency and the local electron temperature. This correction introduces additional angular and spectral anisotropies in SZ maps, and has become an important direction of recent observational and theoretical studies aiming to better constrain the thermal properties of massive clusters (see e.g. [Coulton et al. 2024b](#); [Kay et al. 2024](#); [Remazeilles & Chluba 2020, 2025](#)).

In our simulations, Compton y parameter maps of the tSZ effect are constructed for each lightcone. When a gas particle crosses the lightcone, we accumulate the following dimensionless quantity

$$y = \frac{\sigma_{\text{T}} k_{\text{B}}}{m_e c^2} \frac{m_{\text{g}} n_e T_e}{\Omega_{\text{pix}} d_{\text{A}}^2 \rho} \quad (7)$$

to the map, where m_{g} and ρ are the mass and mass density of the gas particle. Ω_{pix} is the solid angle of a HEALPIX pixel and d_{A} is the angular diameter distance to the observer. Since the gas particles have associated smoothing lengths, quantities derived from the gas are smoothed onto the HEALPIX maps. A detailed description of the smoothing scheme can be found in [Schaye et al. \(2023\)](#).

Similarly, smoothed lightcone maps of the dimensionless Doppler parameter b for the kSZ effect are also provided. This is computed as

$$b = \frac{n_e m_{\text{g}} \sigma_{\text{T}} v_r}{\Omega_{\text{pix}} d_{\text{A}}^2 \rho c}, \quad (8)$$

where v_r is the particle's radial velocity relative to the observer. Under this definition, the mapping between the Doppler parameter is simply given by $\Delta T_{\text{kSZ}}/T_{\text{CMB}} = -b$, where T_{CMB} is taken as 2.73 K. The stacked SZ map is then produced by summing all HEALPIX lightcone maps from individual shells up to $z = 4.5$.

To account for relativistic corrections to the tSZ map, we construct the y -weighted electron temperature map, $\langle T_{\text{e,y}}(\hat{\mathbf{n}}) \rangle$, by dividing the stacked gas-temperature–Compton y product map⁷ by

⁶ Here, we focus only on the relativistic corrections to the tSZ effect. However, non-negligible corrections are also expected from the kSZ and the tSZ–kSZ cross terms (see e.g. [Kuhn et al. 2025](#)). A full discussion of these corrections is left to future work.

⁷ This map is not available in the public repository and is regenerated using

the stacked y map (e.g., [Coulton et al. 2024b](#); [Remazeilles & Chluba 2020, 2025](#)). The relativistic frequency response function, $f^{\text{rel}}(\nu, \langle T_{\text{e,y}}(\hat{\mathbf{n}}) \rangle)$, is computed using the COMBO integration mode of the SZPACK⁸ ([Chluba et al. 2012, 2013](#)). This code allows us to obtain precise SZ response functions for all electron temperatures.

3.4 Anisotropic screening (optical depth τ) maps

The anisotropic, or ‘patchy’, screening effect originates from Thomson scattering of CMB photons by free electrons along the line of sight between the last scattering surface and the observer. This effect provides a novel probe of ionized gas in and around haloes, as well as during the epoch of reionization. Patchy screening effectively acts as a spatially varying screen on the observed CMB sky, damping the primary CMB anisotropies and generating new polarization patterns. Its amplitude is generally much smaller than that of the tSZ and kSZ effects and it scales as

$$\left(\frac{\Delta T}{T_{\text{CMB}}} \right)_{\text{screening}} \approx \tau(\hat{\mathbf{n}}) \left(\frac{\Delta T}{T_{\text{CMB}}} \right)_{\text{primary}}, \quad (9)$$

where $\tau(\hat{\mathbf{n}})$ is the line-of-sight optical depth, determined by the integrated free electron density in the direction $\hat{\mathbf{n}}$ on the sky ($\tau = \int an_e \sigma_{\text{T}} d\chi$).

When combined with the tSZ and kSZ effects, the patchy screening effect has the potential to provide a more complete picture of the electron distribution in and around haloes. Joint analyses can also help disentangle degeneracies between electron density, gas temperature, and galaxy velocities. Several observational attempts have been made to study patchy screening, often through cross-correlations of current high-resolution CMB surveys such as ACT with LSS surveys including unWISE, BOSS, and DESI (e.g. [Coulton et al. 2024a](#); [Hadzhiyska et al. 2025](#)). However, current measurements remain highly challenging, as they are significantly contaminated by other sources of anisotropy in the CMB. Future CMB surveys, such as the Simons Observatory (SO) and CMB-HD, with greater sensitivity and resolution, are expected to reduce these contaminants, which would open the possibility of robust detections of the patchy screening effect and its use as a new probe of gas thermodynamics (e.g. [Schutt et al. 2024](#); [Kramer et al. 2025](#)).

The screening effect is proportional to the line-of-sight integrated optical depth $\tau(\hat{\mathbf{n}})$, which is simply as $\tau = \sum_i^{N_{\text{shell}}} \sigma_{\text{T}} \Delta\tau(\hat{\mathbf{n}})$. We provide a smoothed map of the optical depth, defined as the accumulated contribution per pixel of

$$\Delta\tau = \frac{n_e m_{\text{g}}}{\Omega_{\text{pix}} d_{\text{A}}^2 \rho}. \quad (10)$$

We generate stacked τ maps for different model variations in our repository. Meanwhile, a detailed analysis, including cross-correlations with other LSS tracers and the study of relevant statistics, is left to future work (Conley et al., in prep).

3.5 Cosmic infrared background maps

The cosmic infrared background (CIB) is primarily sourced by star-forming regions. Light emitted by young, massive stars in star-forming galaxies is absorbed by surrounding dust grains and then gets re-emitted at infrared wavelengths. Such emission builds up over time to form a generally unresolved background, which traces

the Lightcone I/O code, which is available online at <https://github.com/jchelly/LightconeIO>

⁸ SZPACK is available at: <https://github.com/CMBSPEC/SZpack>.

the cumulative infrared output of dusty galaxies across a wide range of redshift, $0 < z \lesssim 6$.

The CIB is a dominant component at high frequencies ($\nu \gtrsim 200$ GHz), where it has been measured in detail by surveys such as *IRIS* (Miville-Deschênes & Lagache 2005), *Herschel* (Berta et al. 2011; Viero et al. 2013a) and *Planck* (Planck Collaboration et al. 2011, 2014). At these frequencies, the CIB is often treated as a contaminant that must be carefully removed in order to isolate other signals of interest (e.g. Hill & Spergel 2014; McCarthy & Hill 2024a; Efstathiou & McCarthy 2025). At lower frequencies, however, the CIB becomes subdominant and is therefore more poorly constrained. Nevertheless, CMB experiments detect and must include models of the emission properties of infrared sources (e.g. Dunkley et al. 2013; George et al. 2015; Reichardt et al. 2021; Louis et al. 2025). Beyond being a foreground contaminant to CMB studies, the CIB is an astrophysically rich signal in its own right. It directly traces the cosmic star formation rate density and is sensitive to the emission properties of dusty infrared galaxies. Because the CIB traces faint, unresolved galaxies, it offers a unique tool to study the star formation activities in low-mass or high-redshift systems (e.g. Viero et al. 2013c; Maniyar et al. 2018; Yan et al. 2022; Chiang et al. 2025). In addition, because it traces galaxies and galaxies trace LSS, the clustering of the CIB (and its cross-correlation with other tracers of LSS) will also be sensitive to the underlying cosmological model.

However, modelling of the CIB signal is challenging because of significant uncertainties in how both the SFR and spectral energy distribution (SED) of dusty sources potentially vary with galaxy/halo mass, redshift, local environment, and so on. Previous studies based on dark matter-only simulations adopted the halo model formalism to describe the link between SFR (or sometimes dust mass) and galaxy/halo mass and its evolution (e.g., Sehgal et al. 2010; Stein et al. 2020; Omori 2024), though the specific choices differed in detail. Different studies also approached the SED modelling in different ways (e.g. Magdis et al. 2012; Béthermin et al. 2012; Planck Collaboration et al. 2014, 2016a; Maniyar et al. 2021). On the observational side, accurate separation of the CIB from Galactic dust foregrounds is non-trivial. In particular, the Galactic cirrus emission, which is often traced by low- z HI surveys, must be subtracted carefully to avoid large-scale biases in the measured CIB signal (e.g. see Lenz et al. 2019; McCarthy 2024). In this regard, a detailed understanding of the CIB is crucial not only for studying galaxy evolution but also for precise CMB cosmology.

3.5.1 Three-parameter model

Since the CIB is strongly correlated with the distribution of star-forming galaxies, for this work, we use the star formation rate (SFR) lightcone output maps from the simulations and convert them into mock CIB maps. We start with the conversion between the SFR and the bolometric luminosity of infrared sources, $L_{\text{bol,IR}}$, by assuming that the SFR is proportional to $L_{\text{bol,IR}}$ for a dusty galaxy (Kennicutt 1998). For a Chabrier (2003) stellar initial mass function, as adopted by the FLAMINGO simulation suite, the conversion between $L_{\text{bol,IR}}$ and SFR is given by (also see e.g. Viero et al. 2013c; Figueira et al. 2022)

$$\frac{L_{\text{bol,IR}}}{1 \times 10^{10} L_{\odot}} = \frac{\text{SFR}}{1 M_{\odot} \text{yr}^{-1}}. \quad (11)$$

As discussed earlier, what is observed is not the bolometric luminosity, but rather the luminosity within a specific bandpass. To convert $L_{\text{bol,IR}}$ into the IR luminosity at a given frequency, one needs to include the SED of infrared sources, $\Phi(\nu, T_{\text{dust}})$. We follow

Planck Collaboration et al. (2016a) and express the SED as

$$\Phi(\nu, T_{\text{dust}}, z) = \begin{cases} \left[\exp\left(\frac{h\nu}{k_{\text{B}} T_{\text{dust}}(z)}\right) - 1 \right]^{-1} \nu^{\beta_d + 3} & (\nu \leq \nu') \\ \left[\exp\left(\frac{h\nu'}{k_{\text{B}} T_{\text{dust}}(z)}\right) - 1 \right]^{-1} \nu'^{\beta_d + 3} \left(\frac{\nu}{\nu'}\right)^{-\alpha_d} & (\nu > \nu'), \end{cases} \quad (12)$$

which is modelled as a greybody radiation at low frequencies, with a power law transition at high frequencies. β_d is a free parameter in the model that correlates with dust properties, and we model the dust temperature as a simple power law in redshift: $T_{\text{dust}} = T_0(1+z)^{\alpha}$ (as modelled and used in e.g. Viero et al. 2013b; Stein et al. 2020; Chiang et al. 2025; Bolliet et al. 2025), since dust properties are not explicitly modelled in the FLAMINGO simulations. T_0 and α are also free parameters. In the regime considered in this work, $z \lesssim 4.5$ and $\nu < 1000$ GHz, the SED is expected to remain a greybody and the power-law cutoff component is not considered in the following calculations (e.g. the α_d defined in Equation 12 is not used in Equation 13, with the latter is simply an integral over a greybody). The luminosity at a specific frequency is then calculated as

$$L_{\nu, \text{IR}} = L_{\text{bol,IR}}(\text{SFR}) \frac{\Phi(\nu, T_{\text{dust}})}{\int d\nu \Phi(\nu, T_{\text{dust}})}. \quad (13)$$

Since we will fit our model to data from the *Planck* survey, we further convolve the SED with a detector bandpass function in the modelling⁹, which is defined as the normalised spectral transmission profile of a particular *Planck* channel. The flux is then calculated based on the source's luminosity, radial comoving distance, and redshift:

$$S_{\nu, \text{IR}} = \frac{L_{\nu(1+z), \text{IR}}}{4\pi\chi^2(1+z)}, \quad (14)$$

with the luminosity is calculated in the rest frame. As shown in Equation 13, we shift both the SED and the bandpass function to the rest frame.

In summary, the free parameters in our CIB model are: β_d , T_0 , and the redshift dependence of the dust temperature, α . We fit the model to the auto-power spectra measured by Lenz et al. (2019) at 353, 545, and 857 GHz. In principle, while varying different sets of parameters to determine the best fit (which we do via Monte Carlo Markov chain optimisation), we need to recalculate the power spectrum for a new set of CIB HEALPix lightcone maps at each redshift, which can be computationally expensive. To reduce this cost during the fitting procedure, we take advantage of the fact that the SFR maps are independent of the modelled CIB parameters. As a result, any CIB model-dependent terms can be factored out as constant normalisation terms in the final power spectrum calculation. Therefore, for each lightcone shell, we first compute the SFR power spectrum independently of the fitting routine. The modelled stacked $C_{\ell, \text{CIB}}$ can then be expressed as:

$$C_{\ell, \text{CIB}}^{\text{stacked}} = \sum_{i=1}^N A_i^2 C_{\ell, \text{SFR}}^{i,i} + 2\sum_{i>j} A_i A_j C_{\ell, \text{SFR}}^{i,j}, \quad (15)$$

where the stacking of lightcone shells is performed up to $z = 4.5$. Here, $C_{\ell, \text{SFR}}^{i,i}$ denotes the SFR auto-power spectrum for each lightcone shell. Because the redshift intervals of our output maps are relatively narrow at low z ($\Delta z = 0.05$), while diffuse star-forming structures can span a broader redshift range, we also include the

⁹ The HFI-RIMO-R3.00.fits file available at https://irsa.ipac.caltech.edu/data/Planck/release_3/ancillary-data/HFI_Products.html

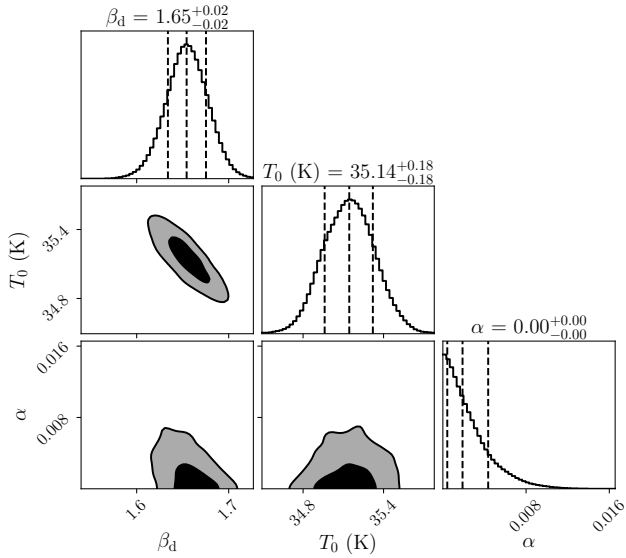


Figure 1. Constraints on the β_d , T_0 , and α parameters of the SED of the CIB model, obtained by fitting to the 353/545/857 GHz auto-power spectrum measurements from [Lenz et al. \(2019\)](#).

cross-shell term $C_{\ell, \text{SFR}}^{i,j}$ of the SFR power spectrum in the calculation. The factor A_i encapsulates all the free parameters in our CIB model and is defined from Equations 11 to 14.

In the fitting, a simple Gaussian likelihood is assumed, with

$$\ln \mathcal{L}(\vec{d}|\vec{\theta}) = -\frac{1}{2} [\vec{d} - \vec{m}(\vec{\theta})]^T \Sigma^{-1} [\vec{d} - \vec{m}(\vec{\theta})]^T, \quad (16)$$

where $\vec{\theta} = [\beta_d, T_0, \alpha]$. \vec{d} is the data vector and \vec{m} are the modelled powers at each frequency band. The covariance matrix is estimated using a simple Gaussian approach. In the absence of sky masks, this is given by

$$\Sigma = \delta_{\ell\ell'} \frac{1}{(2\ell+1)} \left[C_{\alpha\gamma}(\ell) C_{\beta\delta}(\ell) + C_{\alpha\delta}(\ell) C_{\beta\gamma}(\ell) \right], \quad (17)$$

where a Kronecker delta $\delta_{\ell\ell'}$ indicates that the covariance is diagonal in this simple Gaussian limit. The subscripts $[\alpha, \gamma, \beta, \delta]$ correspond to one of the frequencies 353, 545, or 857 GHz. The $C_{\ell s}$ in the covariance matrix are taken as the auto-spectra from the [Lenz et al. \(2019\)](#) CIB maps. We note that the power spectra in [Lenz et al. \(2019\)](#) are measured from CIB maps where individual bright IR point sources have been masked. In our standard fitting methodology we do not include any masking of the simulations, however we show in Appendix C that applying a realistic flux cut on the simulations has only a very small impact on the predicted power spectra and does not change any of our main results or conclusions.

Figure 1 shows the best-fitting values of β_d , T_0 , and α , which are used to generate the mock maps and compute the power spectrum. We find $\beta_d = 1.65 \pm 0.02$ and $T_0 = 35.14 \pm 0.18$ K for the fiducial FLAMINGO model (we discuss how these parameters may vary as a function of feedback and cosmology in Section 5.1.2), where the values are quite close to those adopted in the [Planck Collaboration et al. \(2016a\)](#) study. As indicated by the α parameter, our best-fitting model prefers weak to no redshift evolution of the dust temperature.

3.5.2 Extended models

In the three-parameter model discussed above, the bolometric IR luminosity is specified by the SFR only and the mapping is precisely specified using the [Kennicutt \(1998\)](#) relation. However, in reality this relation has uncertainties and may depend on other properties such as dust mass and galaxy stellar mass. To examine how sensitive our results are to these possible effects, we explore two alternative cases. For the first case, we allow the constant of proportionality in the [Kennicutt \(1998\)](#) relation to vary:

$$\frac{L_{\text{bol,IR}}}{1 \times 10^{10+N_s} L_{\odot}} = \frac{\text{SFR}}{1 \text{ M}_{\odot} \text{ yr}^{-1}}, \quad (18)$$

In the second case, we adopt an alternative parametric form for modelling the infrared flux density, motivated by e.g. [Hayward et al. \(2013\)](#); [Lovell et al. \(2021\)](#); [Cochrane et al. \(2023\)](#); [Kumar et al. \(2025\)](#). Analogous to those studies, we introduce a non-linear dependence on both SFR and dust mass in the computation of the infrared luminosity:

$$L_{\text{bol,IR}} = 10^{10+a} L_{\odot} \left(\frac{\text{SFR}}{100 \text{ M}_{\odot} \text{ yr}^{-1}} \right)^b \left(\frac{M_{\text{dust}}}{10^8 \text{ M}_{\odot}} \right)^c. \quad (19)$$

Due to strong parameter degeneracies in this model, we fix the indices b and c to 0.5, since they do not vary significantly across different simulation setups (e.g., see Table 2 in [Kumar et al. \(2025\)](#)). For the M_{dust} calculation, we assume a fixed dust-to-metal mass ratio (DTM) of 0.4 (e.g. see measurements from [De Cia et al. 2013, 2016](#); [Wiseman et al. 2017](#) and predictions from [Li et al. 2019](#))¹⁰. Since metal mass lightcone maps are not available for the fiducial (2.8 Gpc)³ run, we construct our own HEALPIX maps using halo lightcone outputs from the fiducial (1 Gpc)³ run up to $z = 3.0$. We then apply a normalisation to account for the redshift evolution of the curve to $z = 4.5$, as described in Section 3.1. While computing the metal mass per halo, we consider either only the cold star-forming gas, or all gas within a 3D aperture of 30 kpc. Note that for both of these models, we keep the rest of the CIB modelling, such as the SED of infrared galaxies, unchanged.

In the following discussion, the CIB model with $[\beta_d, T_0, \alpha]$ as free parameters is treated as our default case (denoted the *three-parameter* fit case). The corresponding best-fitting CIB power spectra are shown later in Figure 8. The version with a varied proportionality is referred to as the *four-parameter* fit case, and the model with both SFR and M_{dust} dependencies in $L_{\text{bol,IR}}$ is denoted as the *extended* model (with three free parameters: β_d , T_0 and the amplitude a as defined in Equation 19). The results of these alternative models are provided and discussed in Appendix A. In general, the resulting curves are similar in all three cases, although the best-fitting parameters vary across models.

It is worth noting that more complexity can also be introduced into the SED template, for example when modelling the dust temperature and its redshift evolution. In this work, we adopt a simple power law relation; however, one could in principle trace the dust properties directly in hydrodynamical simulations if explicit dust modelling is included (e.g. in SIMBA and COLIBRE, see [Davé et al. 2019](#); [Schaye et al. 2025](#)), or model them as functions of galaxy properties which could be fitted to the data. Since our simplified model already provides a good fit to the data we compare to, we leave such analyses to future work.

¹⁰ Note that the exact choice of DTM is inconsequential, since the uncertainty of this constant factor is absorbed into the free amplitude in Equation 19.

3.6 Radio point sources

Radio point sources make a non-negligible contribution to the CMB anisotropy signal at low frequencies, typically for $\nu \lesssim 150$ GHz. Individually detected bright sources are often masked out prior to component separation. If such sources are strongly correlated with the signal of interest, however, there is a risk that simply masking these radio sources will also mask the main contribution to the signal of interest. For example, if most massive local clusters have bright radio sources that are masked, the contribution of these clusters to the measured tSZ power spectrum may be underestimated, depending on the angular size of the mask and what compensation (if any) is made for the excluded area. Aside from individually detected bright sources, we also expect there to be many more undetected radio sources that could contribute significantly to the measurements and this contribution therefore needs to be modelled.

In the simulations, black holes are characterised solely by their masses and accretion rates (and positions and velocities). As there is no explicit modelling of the accretion disk or of the scales near it, the simulations cannot predict which black holes ought to correspond to radio-loud AGN (as opposed to radiatively efficient, radio-quiet quasars, for example). We therefore must identify potential radio AGN based on properties that are followed and then assign each radio AGN an appropriate luminosity.

To identify potential radio AGN sources, we use the ratio of the bolometric luminosity to the derived Eddington luminosity for each black hole. This is motivated based on both theoretical and observational studies (e.g. see Best & Heckman 2012; Fabian 2012; Heckman & Best 2014), which show that radiatively-inefficient radio-loud AGN are much more likely to have low Eddington ratios, typically below a few percent. In more detail, for each black hole with a subgrid mass greater than $10^7 M_\odot$, we compute its bolometric luminosity using $L_{\text{bol}} = \epsilon_r (1 - \epsilon_f) \dot{M}_{\text{BH}} c^2$, where \dot{M}_{BH} is the accretion rate of each black hole particle, ϵ_r is the fraction of accreted rest-mass energy converted into radiation and $(1 - \epsilon_f)$ accounts for the fraction of radiated energy not coupled to feedback¹¹. The FLAMINGO simulations set ϵ_r and ϵ_f as 0.1 and 0.15, respectively. We apply the aforementioned $10^7 M_\odot$ mass cut to include only physically meaningful black holes in the analysis, which is far from the black hole seed mass, i.e. the mass scale at which it reaches the black hole–halo mass relation (see Fig. 1 in Costello et al. 2025). We explore the sensitivity of our results to the specific Eddington ratio cut applied. Specifically, we consider three cuts in Eddington ratio, $\lambda_{\text{Edd}} \equiv L_{\text{bol}}/L_{\text{Edd}}$: $\lambda_{\text{Edd}} < 10^{-2}$, 10^{-3} and 10^{-6} .

After applying an Eddington ratio cut, we proceed to assign radio luminosities to each selected black hole. As radio point source statistics are poorly constrained at high frequency bands in CMB surveys due to their low number density, our approach is to assign radio fluxes to each selected black hole based on deep, well-measured luminosity functions at low frequencies observed with LOFAR. These fluxes are then extrapolated to higher frequencies using a simple power law spectral model (as expected for synchrotron radiation), with free parameters calibrated against the source number count curves from other CMB surveys.

More specifically, for each of the λ_{Edd} cases, we rank order the selected black holes based on their L_{bol} values¹². We then re-assign the radio luminosity of each black hole by abundance matching to the radio luminosity function (RLF) of low-excitation radio galaxies (LERG) measured at 150 MHz from the LOFAR survey (Kondapally et al. 2022). We select these measurements because they are less uncertain (when compared to the RLF of high-excitation radio galaxies), and the parametric fits as a function of redshift (up to $z = 2.5$) are well calibrated against observations¹³. Abundance matching is done by computing the cumulative number density as a function of luminosity from the simulation and comparing it to the observed cumulative number density curve at 150 MHz, ranked by luminosity. An example of the reconstructed RLF from our simulation is shown in Figure 2. The fluxes are then computed as

$$S_{150 \text{ MHz}} = \frac{L_{150 \text{ MHz}}}{4\pi D_L^2(z)} (1+z)^{1+\alpha}, \quad (20)$$

with α assumed to be -0.7 (as adopted in e.g. Kondapally et al. 2022, and references therein).

From the 150 MHz fluxes we extend to other CMB frequencies using a simple power law, modelled as $S_\nu = S_{150 \text{ MHz}} (\nu/150 \text{ MHz})^{\alpha_{\text{radio}}}$, where α_{radio} is obtained by jointly fitting our curve to the measured differential source number counts from the SPT survey at 95, 150, and 220 GHz (Everett et al. 2020). Specifically, we fit to the range 10^{-2} – 10^{-1} Jy, since the shape of the bright end depends strongly on the particular λ_{Edd} cut and is also affected by large uncertainties in the measured data. Figure 3 shows a comparison between our recovered differential source number counts, $S^{2.5} dN/dS$, at the three frequencies and the SPT measurements. We also include the measurements from ACT and *Planck*, as well as the predictions from the AGORA simulation, in the same figure. We find good agreement between our recovered curves and the other measurements over the fitted range. We find that the best-fitting α_{radio} values are not very sensitive to the particular frequency we fit to. For example, the best-fitting values are -0.56 , -0.58 and -0.61 when fitting independently to the 95, 150, and 220 GHz SPT measurements, respectively, as opposed to our default method of jointly fitting all three data sets.

3.7 Lensing of observables

Large-scale structure not only acts as a lens for the primary CMB, it will also lens any CMB secondary sources more distant than the lens. To account for the lensing effects on each mock CMB secondary anisotropy map, we use PIXELL¹⁴ to deflect the background maps. Observables in the lightcone are lensed shell by shell, using the integrated κ -map up to the shell of interest. For radio point sources, we add the lensing effect to each object before projecting them onto 2D maps. The lensed flux density at the undeflected positions is

¹² We note that the FLAMINGO simulations do not reproduce the observed evolution of the quasar luminosity function particularly well (see Ding et al. 2025), suggesting that the bolometric luminosities of the simulated BHs may not be realistic. However, this should not be an issue for our purposes since only the luminosity *ranking* matters in the abundance matching when assigning radio luminosities to the simulated BH samples.

¹³ It is worth noting that it is now possible to explore the radio properties of AGNs at higher redshifts (e.g. at $z \sim 5 - 6$ from the James Webb Space Telescope (JWST) AGN samples, Mazzolari et al. 2024). We will investigate the cross-correlation between CMB secondary anisotropies and radio emission at higher redshifts in future work.

¹⁴ <https://github.com/simonsobs/pixell>

¹¹ Note that the inclusion of $(1 - \epsilon_f)$ factor implies that radiation that is coupled to the gas (i.e., as feedback) is not observable. Its inclusion may not be appropriate for comparisons to observations that have been explicitly corrected for intrinsic absorption. We include the factor, but its inclusion has no practical consequences for our results as it does not change the rank ordering of black hole luminosities, and our results are also insensitive to small variations in the Eddington ratio cut.

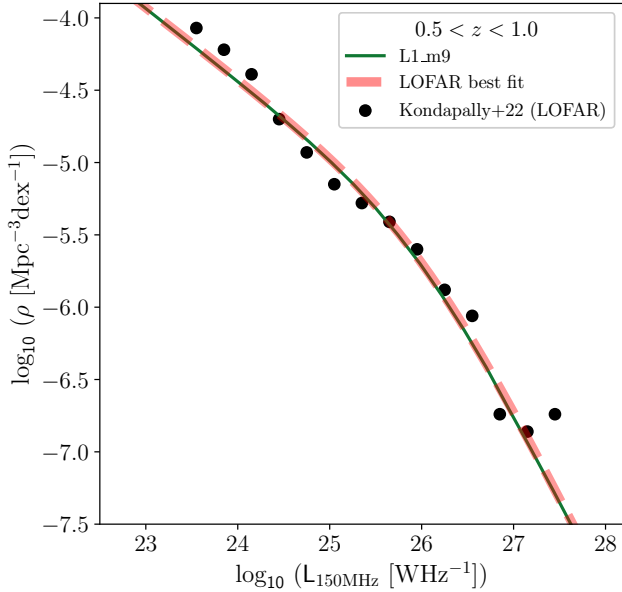


Figure 2. The 150 MHz radio luminosity function (RLF) reconstructed from the black hole particle lightcone in the fiducial $(1 \text{ Gpc})^3$ run. This is an example showing the case for a black hole selection of $\lambda_{\text{Edd}} < 10^{-2}$ over the redshift range $0.5 < z < 1.0$. The observed RLF from the LOFAR survey within the same redshift interval (black points), as well as its best-fitting parametric model (red thick dashed line), are overplotted for comparison (see [Kondapally et al. 2022](#)). Abundance matching is used to map from the bolometric luminosity function of black holes in the simulation to the observed RLF (see text).

given by $S^{\text{lens}} = |\mu|S^{\text{unlensed}}$, where μ is the magnification map. This map is constructed based on

$$\mu = \frac{1}{(1 - \kappa)^2 - |\gamma|^2}, \quad (21)$$

where we use the relationship (see e.g. [Hu 2000](#))

$$\kappa_\ell = \frac{\ell(\ell + 1)}{2} \Phi_\ell,$$

$$\gamma_\ell = \frac{1}{2} \sqrt{\frac{(\ell + 2)!}{(\ell - 2)!}} \Phi_\ell$$

to convert the convergence κ map into the shear γ map.

Figure 4 shows an example of the impact of this lensing effect on the resulting power spectrum analysis, where we show the tSZ, kSZ and 217 GHz CIB auto-power spectrum as an illustration. In general, the lensing-induced modifications are small, where a suppression of power by around 1-2% is expected at small scales. The lensing effect on the CIB map at 217 GHz is slightly stronger than that on the SZ effect, due to its broader redshift kernel which has more overlap with the CMB lensing kernel. Although this is a relatively small effect, it may become important for future high-precision cosmological analyses.

4 MAPS AND POWER SPECTRA

In this section, we present the mock maps and the resulting power spectra predicted by the FLAMINGO simulations, and we compare them with other multicomponent CMB simulations and observational data.

4.1 tSZ and kSZ auto-spectra

The left panel of Figure 5 shows an illustration of the mock, non-relativistic, lensed tSZ intensity map at 857 GHz. This map is created by stacking lightcone data from the fiducial $(2.8 \text{ Gpc})^3$ run up to redshift $z = 4.5$. The middle panel is the corresponding difference map between the relativistically corrected tSZ intensities and the non-relativistic anisotropy map. These relativistic corrections are computed using the SZPACK, as described in Section 3.3. The rightmost panel shows the y -weighted temperature map for comparison. It is visually evident that the relativistic corrections to the tSZ anisotropies are significantly enhanced in regions containing massive clusters (red regions in the rightmost panel) owing to their high temperatures. This suggests that neglecting relativistic effects in these regions could introduce appreciable biases, especially when extracting the tSZ auto-power spectrum from the observed temperature maps (see e.g. [Remazeilles et al. 2019](#)).

In Figure 6, we present the comparison of tSZ $(y - y)$ auto-power spectra computed from our mock maps and other simulations. For this comparison the power spectra are derived from unlensed maps without relativistic corrections, for consistency with the previous studies we compare to. Note that a detailed comparison between the FLAMINGO tSZ auto-power spectrum and observational data can be found in [McCarthy et al. \(2023c, 2025\)](#); here, we mainly focus on comparing different simulation results.

Among the FLAMINGO model variants, the LS8 run yields the lowest power amplitude, due to the strong σ_8 dependence of the tSZ power spectrum. And, as expected, the strongest AGN model (fgas- 8σ) shows the greatest suppression of power on small angular scales, due to gas expulsion from groups and low-mass clusters. The two BAHAMAS curves show a similar trend and amplitude to the FLAMINGO fiducial and fgas- 8σ curves, which is perhaps unsurprising given the similar approaches to feedback calibration of the two suites of simulations. Note that the differences between the FLAMINGO fiducial and BAHAMAS curves on large-scales ($\ell \lesssim 300$) are likely due to a combination of factors, including cosmic variance, differences in cosmology, and box size limitations for BAHAMAS (also see discussions in [Schaller et al. 2025](#)). The shaded region shows the spread in power spectra for the large FLAMINGO 2.8 Gpc using 8 independent lightcones.

The WebSky simulation shows a trend similar to that in the fiducial case but with a lower power amplitude, in spite of the WebSky cosmology being much more similar to the fiducial FLAMINGO cosmology rather than LS8. It is unclear what is driving this difference. An interesting comparison comes from the AGORA and BAHAMAS results. In principle, AGORA adopted the bound gas profile from the HMx code ([Mead et al. 2020](#)), which was calibrated on the BAHAMAS simulations. However, AGORA produces significantly less tSZ power than BAHAMAS at small scales. [Mead et al. \(2020\)](#) cautioned against using HMx for tSZ effect power spectrum analyses given the relatively poor fit of the model to the pressure power spectrum of BAHAMAS. In addition, AGORA does not model gas outside of haloes, apart from gas that was ejected from the haloes (the typical effective range of the feedback effect is $\approx 2 - 3 r_{200}$). The diffuse electron distribution in the intergalactic medium is therefore likely to be underestimated. Together, these factors plausibly explain the lower tSZ power predicted by AGORA compared to BAHAMAS and FLAMINGO.

In principle the tSZ monopole can also be used as a test of cosmology and astrophysics (e.g., [Hill et al. 2015](#); [Thiele et al. 2022](#)). However, measurements of the monopole are challenging. In fact, the only experiment to date that has been able to place a




Figure 3. Comparison of the differential source number counts between observations and simulations. The blue, orange, and green solid curves are the best-fitting source count models for the three black hole selections considered in this study (corresponding to different Eddington rate cuts), with the best-fitting α_{radio} values obtained by fitting to the measured SPT data at three frequencies (Everett et al. 2020) (black diamonds) in the flux range $S_\nu = 10^{-2}$ – 10^{-1} Jy. Other observations from the ACT (Vargas et al. 2023) and Planck (Planck Collaboration et al. 2018) are shown as violet triangles and light green crosses respectively. Results from the AGORA simulation are shown as a red thick dashed line for comparison. Overall, the simulations reproduce the number counts well.

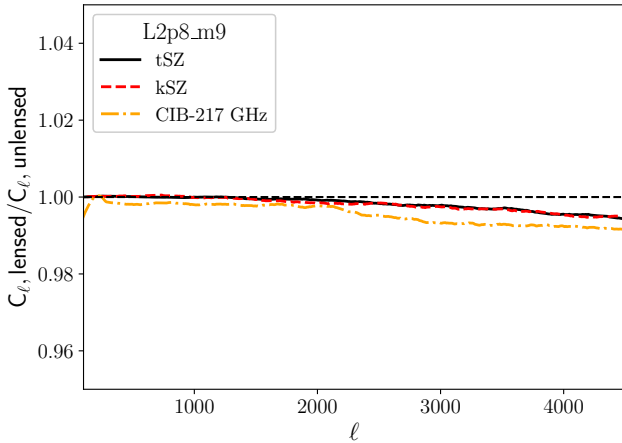


Figure 4. Comparison of the lensed and unlensed auto-power spectra of different LSS tracers from the fiducial $(2.8 \text{ Gpc})^3$ run. All curves are obtained by averaging over 8 independent lightcones. A 1-2% suppression of power at small scales is expected from foreground lensing effects.

direct constraint on $\langle y \rangle$, from spectral distortion measurements of the primary CMB, is COBE-FIRAS (Fixsen et al. 1996), which placed an upper limit on the mean y of the Universe. A very recent re-analysis of the COBE-FIRAS data using an improved astrophysical foreground cleaning technique has reduced this upper limit by approximately a factor of three (Fabbian et al. 2025). Future CMB spectrometers, such as BISO (Maffei et al. 2021), TMS (Rubio Martín et al. 2020), COSMO (Masi et al. 2021), and FOSSIL (FOSSIL Collaboration 2022), will be able to place much stronger direct constraints on the mean y -distortion. Note that the monopole from haloes (which is expected to dominate the overall y signal) can be indirectly inferred from experiments such as Planck, by cross-correlating Compton y maps with tracers of large-scale structure

Table 2. tSZ monopole measured from observations and from different FLAMINGO model variations. For each model, the predicted monopole is estimated by averaging the mock full-sky Compton y map integrated up to $z = 3.0$.

Observations	$\langle y \rangle$
COBE-FIRAS (Fixsen et al. 1996)	$< 15 \times 10^{-6}$
COBE-FIRAS re-analysis (Fabbian et al. 2025)	$< 5.2 \times 10^{-6}$
Planck (Chiang et al. 2020)	$1.22^{+0.23}_{-0.17} \times 10^{-6}$
FLAMINGO models	$\langle y \rangle$
Fiducial	$(1.52 \pm 0.01) \times 10^{-6}$
$M_* - \sigma$	1.59×10^{-6}
$f_{\text{gas}} + 2\sigma$	1.46×10^{-6}
$f_{\text{gas}} - 8\sigma$	1.75×10^{-6}
Jet	1.40×10^{-6}
Planck	1.57×10^{-6}
LS8	1.17×10^{-6}

(e.g., galaxies, quasars) in a tomographic way, allowing one to constrain a bias-weighted estimate of dy/dz . Modelling the bias of the tracers and integrating over redshift allows one to infer $\langle y \rangle$ (e.g., Chiang et al. 2020).

Table 2 compares the observed tSZ monopole constraints from COBE-FIRAS and Planck with predictions from the various FLAMINGO model variations. The $\langle y \rangle$ value for each model is obtained by averaging the Compton y map integrated up to $z = 3.0$. The uncertainty value quoted after the fiducial run result is obtained by averaging over 8 different observers from the $(2.8 \text{ Gpc})^3$ box. All runs are easily compatible with the 95% upper limit provided by COBE-FIRAS, including the more stringent re-analysis by Fabbian et al. (2025). For the fiducial run, the result is slightly higher than the model-dependent tomographic measurement from Planck. The strongest AGN model ($f_{\text{gas}} - 8\sigma$ run) yields the largest y -monopole due to the enhanced thermal energy of the expelled gas, yielding a $\langle y \rangle$ that is approximately 2.3σ larger than derived by Chiang et al.




Figure 5. Full-sky mock intensity maps of the lensed thermal Sunyaev–Zel’dovich (tSZ) effect and its relativistic correction. These maps are generated from the lightcone outputs of the fiducial $(2.8 \text{ Gpc})^3$ run integrated up to $z = 4.5$. *Left:* tSZ intensity map at 857 GHz, computed using Equation 3 (i.e., non-relativistic tSZ). *Middle:* Difference map between the relativistically corrected and non-relativistic tSZ intensity maps at the same frequency, with the relativistic corrections applied using SZPACK. *Right:* the corresponding y -weighted temperature map. Note the large differences in the color bar ranges across the three maps.

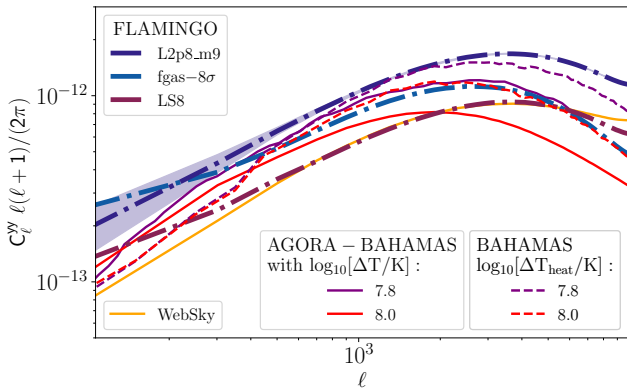


Figure 6. Comparison of the tSZ angular power spectrum from different simulations. Thick dash-dot lines in different colors show results from the FLAMINGO $(2.8 \text{ Gpc})^3$ fiducial run and a subset of its $(1 \text{ Gpc})^3$ model variants, integrated up to $z = 3.0$. The shaded region shows the spread obtained by averaging over 8 independent lightcones. The orange line shows the spectrum from the WebSky simulation. Solid lines are the spectra computed from the AGORA tSZ mock maps, calibrated against two BAHAMAS AGN models with heating temperatures of $10^{7.8}$ and $10^{8.0}$ K (the latter is the default AGORA gas model) using the HMX code of Mead et al. (2020). Note that AGORA does not model the full gas distribution outside of haloes. For comparison, dashed lines show the corresponding results from the BAHAMAS simulations using the same AGN models, where the BAHAMAS maps are constructed in a lightcone with a field of view of 5 square degrees.

(2020). The LS8 model predicts the lowest value and, interestingly, is in excellent agreement with the *Planck* mean value. This is consistent with the findings of McCarthy et al. (2023c), who shown that the measured tSZ power spectrum (Bolliet et al. 2018, but see Efstathiou & McCarthy 2025 for a discussion of the uncertainties in the power spectrum measurements) has a lower amplitude than expected for a *Planck* primary CMB cosmology, while a LS8 cosmology was in relatively good agreement with the measurements on large scales.

We now move on to the kSZ effect. Figure 7 shows the kSZ auto-power spectrum computed from the fiducial $(2.8 \text{ Gpc})^3$ run, compared with kSZ power spectra from previous simulations. Here we see that the fiducial FLAMINGO model yields results that are fairly consistent with the WebSky simulation. Although the AGORA result is integrated only up to $z = 3.0$, its amplitude is comparable to ours at intermediate scales ($\ell \approx 500\text{--}1000$). However, it shows a significant suppression of power at large scales ($\ell \lesssim 300$), which

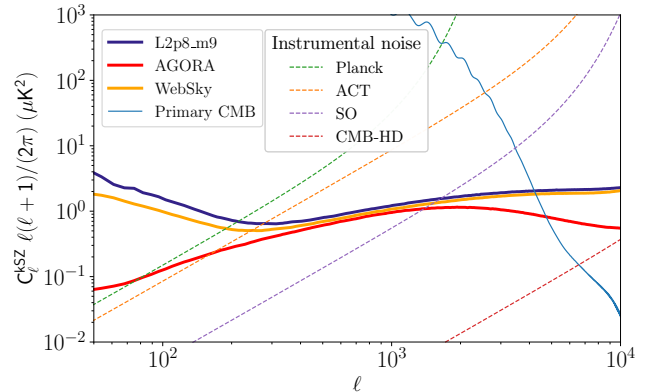


Figure 7. Comparison of the kSZ auto-power spectra from different simulations. The thick dark line shows the result from the FLAMINGO fiducial $(2.8 \text{ Gpc})^3$ run, integrated up to $z = 4.5$. The thick red line shows the AGORA result up to $z = 3.0$, and the thick orange line shows the WebSky result up to $z = 4.5$. For reference, the primary CMB power spectrum (thin blue line) and the CMB instrumental noise curves for some current and future CMB surveys (dashed lines) are over-plotted.

is plausibly due to the lack of a diffuse IGM component in their model. The difference at small scales ($\ell \gtrsim 2000$) is likely because their model is based on AGN 8.0 BAHAMAS, which has stronger feedback effects on the gas distribution than fiducial FLAMINGO. We comment on the cosmology and feedback dependence of the kSZ power spectra in Section 5.

For comparison, we also include the primary CMB power spectrum and the instrumental noise power spectra of several current and upcoming CMB surveys in the same figure. The primary CMB spectrum is generated using the CAMB (Lewis et al. 2000), while the instrumental noise power spectrum is modelled as $C_{\text{noise}}^{\ell}/W_{\text{beam}}^2(\ell)$, where:

$$C_{\text{noise}}^{\ell} = \left(\frac{\Delta T}{T_{\text{CMB}}} \right)^2, \quad (22)$$

and

$$W_{\text{beam}}(\ell) = \exp(-\ell(\ell+1)/\ell_{\text{beam}}^2/2). \quad (23)$$

Here, ℓ_{beam} is defined as $\sqrt{8 \ln 2}/\theta_{\text{FWHM}}$, and ΔT is the white noise level. We adopted the values of $\theta_{\text{FWHM}} = 1.3'$ and $\Delta T = 25 \mu\text{K-arcmin}$ for the ACT survey at 150 GHz (Naess et al. 2020; Qu et al. 2024), $\theta_{\text{FWHM}} = 7.2'$ and $\Delta T = 33 \mu\text{K-arcmin}$ for the *Planck* survey

at 143 GHz (Planck Collaboration et al. 2020a), $\theta_{\text{FWHM}} = 0.25'$ and $\Delta_T = 0.5 \mu\text{K-arcmin}$ for the CMB-HD survey at 150 GHz (Sehgal et al. 2019), and $\theta_{\text{FWHM}} = 1.4'$ and $\Delta_T = 6.3 \mu\text{K-arcmin}$ for the SO-LAT survey at 145 GHz (Ade et al. 2019). In general, the kSZ power spectrum is subdominant to the combined signal from the primary CMB and instrumental noise. However, detection seems to be achievable at small angular scales ($\ell \gtrsim 5000$), particularly with next-generation experiments such as the CMB-HD. As we will show later, this is also the regime where the kSZ signal becomes most sensitive to feedback and cosmological effects. This highlights the importance of accurately modelling the kSZ signal for future CMB surveys.

4.2 CIB statistics: auto-spectra and cross correlations with other probes

The best-fitting CIB auto- and cross-power spectra are shown in Figure 8, where we compare the 353/545/857 GHz power spectra with observational measurements from Lenz et al. (2019) and Planck Collaboration et al. (2014). The model provides a qualitatively good fit to the observed data across a broad range of frequencies and multipoles. In Appendix B, we show the CIB power spectra as ℓC_ℓ and plotted on a linear y-axis scale. The curves are also qualitatively similar to the AGORA prediction, for which the free parameters of its adopted SED model were also constrained by fitting to the 353/545/857 GHz data. Note, however, that neither our model nor AGORA were fitted to the 217 GHz measurements, which is where we see the largest discrepancy between the observations and the AGORA model. Although we did not explicitly fit the 217 GHz data, our best-fitting SED model reproduces the observational measurements at this frequency reasonably well. This provides a reassuring check on the accuracy of the best-fitting SED.

We also find broad agreement between the FLAMINGO results and those from WebSky. Note that the WebSky model was not fitted to the Lenz et al. (2019) data, but was adjusted to approximately match the 545 GHz measurement at $\ell = 500$ from Planck Collaboration et al. (2014). For this comparison, we use the *masked* WebSky CIB maps, in which all pixels with flux density greater than 400 mJy are masked and refilled. As already noted, such bright-source masking has only a very minor effect on our model (as well as on AGORA), but it has a very large impact on the WebSky CIB power spectrum at high ℓ (see Appendix C for a detailed discussion). Therefore, we use the masked WebSky CIB maps for the following comparisons.

Another interesting property of the CIB map is its frequency decoherence, which quantifies how similar the spatial structures of CIB anisotropies are across different frequency channels. This is characterised by the decorrelation coefficient, defined as $C_\ell^{\nu\nu'} / \sqrt{C_\ell^{\nu\nu} C_\ell^{\nu'\nu'}}$. This quantity captures the degree of overlap in the redshift contributions of CIB signals across frequencies: adjacent frequency channels tend to be more correlated (with value closer to 1) because their redshift-dependent emission peaks overlap more. To estimate this, we compute the mean and standard deviation of this decorrelation coefficient within the multipole range $150 < \ell < 1000$, where the CIB clustering signal dominates. The results are summarised in Table 3. Compared with other simulations and observations, our predictions show good agreement with previous measurements, where the decorrelation level increases as the frequency separation between channels gets larger.

Similar to the case for the tSZ effect, the CIB monopole can be inferred via cross correlation between CIB maps with other

Table 3. Frequency decoherence of the CIB measured by averaging $C_\ell^{\nu\nu'} / \sqrt{C_\ell^{\nu\nu} C_\ell^{\nu'\nu'}}$ over the range $150 < \ell < 1000$. Error bars correspond to the standard deviation in this range. The results from Planck Collaboration et al. (2014), Lenz et al. (2019), the WebSky and AGORA simulations are also listed for comparison.

		Frequency (GHz)			
		857	545	353	217
857	<i>Planck</i>	1	0.949 ± 0.005	0.911 ± 0.003	0.85 ± 0.05
	Lenz et al.	1	0.96 ± 0.01	0.91 ± 0.01	-
	FLAMINGO	1	0.959 ± 0.004	0.895 ± 0.009	0.841 ± 0.012
	WebSky	1	0.933 ± 0.017	0.882 ± 0.021	0.838 ± 0.026
	AGORA	1	0.970 ± 0.003	0.907 ± 0.008	0.840 ± 0.013
545	<i>Planck</i>	...	1	0.983 ± 0.007	0.90 ± 0.05
	Lenz et al.	...	1	0.98 ± 0.01	-
	FLAMINGO	...	1	0.983 ± 0.001	0.956 ± 0.003
	WebSky	...	1	0.960 ± 0.014	0.935 ± 0.018
	AGORA	...	1	0.981 ± 0.002	0.941 ± 0.005
353	<i>Planck</i>	1	0.91 ± 0.05
	Lenz et al.	1	-
	FLAMINGO	1	0.993 ± 0.001
	WebSky	1	0.968 ± 0.014
	AGORA	1	0.989 ± 0.001

tracers of large-scale structure. Figure 9 shows the CIB monopole integrated up to different maximum redshifts, using our best-fitting CIB SED as defined in Equation 12. For each frequency

$$I_\nu = \int_0^{z_{\text{max}}} dz \frac{dI_\nu}{dz}, \quad (24)$$

with dI_ν/dz defined as

$$\frac{dI_\nu}{dz} = \frac{c}{4\pi} \frac{\epsilon(\nu_{\text{emitted}})}{H(z)(1+z)}, \quad (25)$$

where the emissivity $\epsilon(\nu_{\text{emitted}})$ is derived from our best-fitting CIB SED. From this figure, we see that the CIB signal converges by $z \approx 3$, and the power-law break happens at $\nu \gtrsim 1000$ GHz. We find that our predictions are in reasonably good agreement with the *Planck*-based measurements of Chiang et al. (2025) (we show their best-fitting parametric model to their measurements). It would also be interesting to compare to tomographic measurements of CIB intensity using similar observational techniques (e.g. galaxy–CIB cross-correlations as applied by Yan et al. 2022 and Chiang et al. 2025). We plan to explore this comparison in detail in future work.

We now explore the cross-power spectra between the CIB and other LSS tracers. Here, we present the results for the CIB–tSZ and CIB–CMB lensing (κ) cross-power spectra, as shown in Figure 10.

For the CIB–tSZ cross-correlation, our predicted curves are qualitatively consistent with those from the AGORA. The simulation uncertainties on the FLAMINGO curves, which are estimated by averaging the results from 8 different lightcones in the $(2.8 \text{ Gpc})^3$ fiducial run, are relatively large at large angular scales (small ℓ s). This is expected since the CIB–tSZ signal mainly traces warm–hot gas surrounding star-forming regions in intermediate-mass haloes. On large scales, the contribution is dominated by more massive and rarer structures, which introduces significant sample variance and hence larger uncertainties in the measurement. The WebSky results are higher at all three frequencies compared to our predicted curves (as well as AGORA’s; see Appendix B for the power spectra plotted with a linear y-axis scale, where the differences are more noticeable). This may be due to their adopted CIB modelling, which could result in higher spatial correlations between the CIB and tSZ signals for low-redshift haloes (see e.g. Zubeldia et al. 2023, 2025). The halo model-based prediction from Planck Collaboration et al.

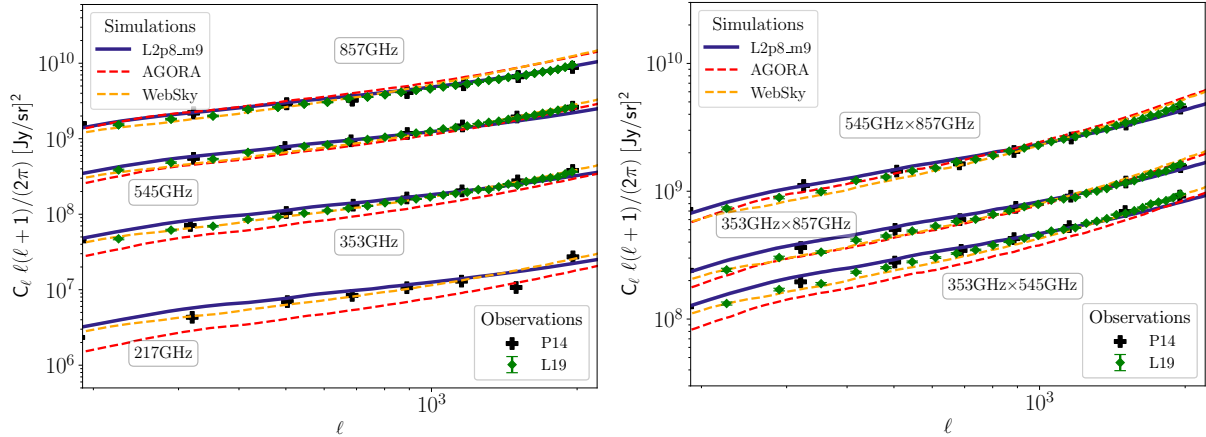


Figure 8. The auto- (left) and cross- (right) power spectra measured from the CIB maps generated from the FLAMINGO fiducial (2.8 Gpc)³ run at the *Planck* HFI frequencies. All curves are obtained by averaging over 8 independent lightcones. Results from the AGORA (red dashed) and the WebSky (orange dashed) simulations are displayed for comparison. Observational data points from [Planck Collaboration et al. \(2014\)](#) and [Lenz et al. \(2019\)](#) are shown as black crosses and green diamonds, respectively.




Figure 9. CIB monopole computed using the best-fitting CIB SED from the FLAMINGO fiducial (2.8 Gpc)³ run, integrated up to different maximum redshifts. For comparison, the measurements provided by [Chiang et al. \(2025\)](#) are overlotted as a black dashed line. For reference, the vertical dashed lines are the *Planck* HFI frequencies — 353 GHz, 545 GHz, and 857 GHz — which are the frequencies used in the SED fitting to the CIB power spectra.

(2016c) is systematically lower at all frequencies compared to the simulation-based results.

In practice, detecting the CIB–tSZ cross-correlation in observations can be challenging because the two signals trace quite distinct physical processes: the CIB is more sensitive to star-forming galaxies, while the tSZ effect traces hot electron pressure, preferentially in massive galaxy groups and clusters. Their overlap is limited due to environmental processes (e.g., ram pressure stripping), which tends to quench star formation in high density regions. This leads to a relatively weak correlation signal that is challenging to measure, particularly in the face of limited sensitivity and foreground contaminants. Due to the limited observational constraints for this statistic, we simply highlight the potential non-negligible contribution of this correlation and leave a more detailed discussion for future work.

The right panel of Figure 10 shows the CIB– κ cross-power spectrum, which is a more robustly measured quantity compared to the CIB–tSZ cross-correlation. Note that the uncertainties in this statistic are much smaller across all multipole scales. This is because κ traces the total matter distribution and has a broad redshift kernel similar to that of the CIB.

Overall, the shape and amplitude predicted by our models are consistent with those from other simulations and observations. However, caution must be taken when interpreting the observational measurements: the data from [Lenz et al. \(2019\)](#) are known to be biased at $\ell \lesssim 100$, due to a non-optimal method of Galactic dust foreground subtraction (see e.g. [McCarthy 2024](#)). Instead, the measurements from [McCarthy \(2024\)](#), which adopt an improved dust-cleaning approach, are considered more reliable on these large angular scales. At smaller scales, the simulations match the [Lenz et al. \(2019\)](#) data relatively well. Although we have a slightly higher amplitude at $\ell \sim 100$ –1000 when compared to the [McCarthy \(2024\)](#) measurements, the predictions are still reasonably close to the measurements. A small adjustment in cosmology (see Fig. 14 below) might completely remove this offset. On larger scales, all curves are consistent with the measurements of [Lenz et al. \(2019\)](#) and halo model predictions.

4.3 Radio point source statistics

The left panel of Figure 11 shows the auto-power spectrum of the radio point sources for three different λ_{Edd} cases, measured on our stacked HEALPIX maps integrated to $z = 2.5$. The power spectra for all cases generally follow a Poisson-like distribution, which is consistent with both model predictions and other simulation results. In particular, our predictions show that the population with the lowest accretion rate threshold ($\lambda_{\text{Edd}} < 10^{-6}$) has the highest power. By construction, we have the same space density of bright sources regardless of the applied Eddington cuts. However, in the low Eddington-cut case, we select sources that typically reside in massive clusters (as shown in Figure 12), which are more highly biased tracers with respect to the overall matter distribution, and therefore appear more strongly clustered in the resulting power spectrum.

Observational analyses often mask individually-detected

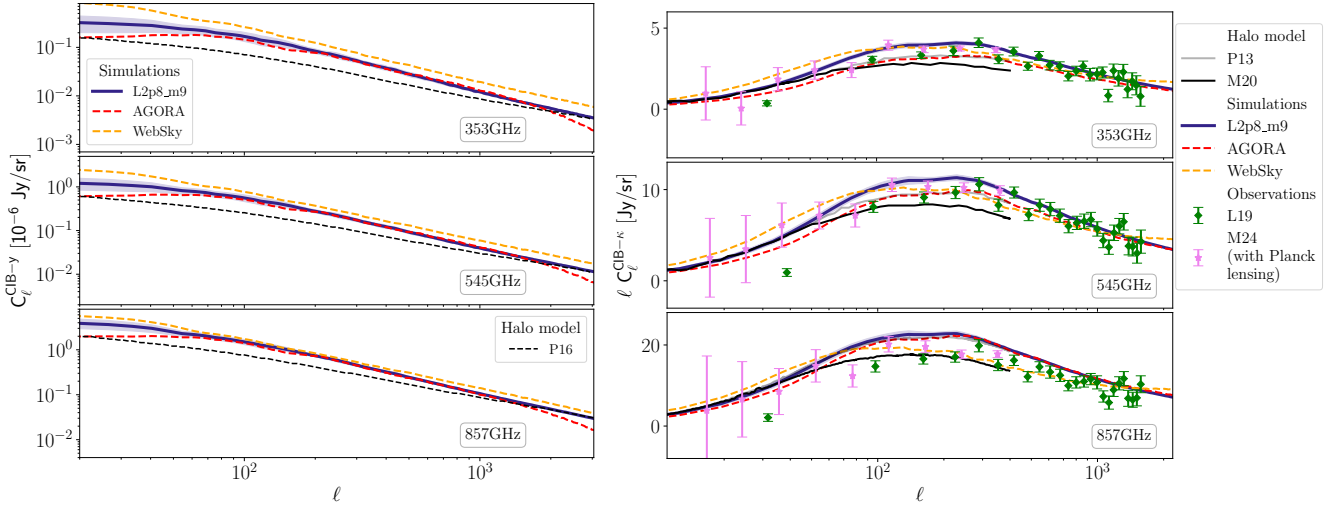


Figure 10. CIB-LSS tracers cross-power spectra at different *Planck* frequencies. *Left:* Comparison of the CIB-tSZ cross-power spectrum between simulations (dark solid: fiducial $(2.8 \text{ Gpc})^3$ FLAMINGO run, with shaded regions estimated by averaging the results from eight different lightcones; red dashed: AGORA; orange dashed: WebSky) and halo model predictions from [Planck Collaboration et al. \(2016c\)](#) (black dashed). *Right:* Comparison of the CIB-CMB lensing (κ) cross-power spectrum between simulations, halo model predictions (black solid: [Planck Collaboration et al. \(2014\)](#); gray solid: [Maniyar et al. \(2021\)](#)), and observations (green diamonds: [Lenz et al. \(2019\)](#), violet stars: [McCarthy \(2024\)](#), taken from the Galactic dust-cleaning case of $N_{\text{HI}} < 2.5 \times 10^{20} \text{ cm}^{-2}$ and real ILC on $N_{\text{side}} = 1$ super pixels).




Figure 11. Auto-power spectrum of radio point sources for the three λ_{Edd} cuts considered in our mock catalogue, where the radio luminosities are assigned by abundance matching to the observed radio luminosity function of low-excitation radio galaxies, measured at 150 MHz from the LOFAR survey (see Section 3.6 for details). *Left:* power spectrum computed from the full sample without masking bright sources. For comparison, the corresponding full power spectra at the same frequency from other simulations are shown as dashed lines. *Right:* power spectrum after removing bright sources with a 7 mJy flux cut at 148 GHz. Dashed lines are the curves from other simulations with the same flux cut applied.

bright sources. We demonstrate the impact this has by excluding sources with flux densities greater than 7 mJy at 148 GHz. This is the same flux cut as adopted by [Ade et al. \(2019\)](#) and in the WebSky simulation ([Li et al. 2022](#)). As done in [Ade et al. \(2019\)](#), we populate the remaining sources onto a HEALPix map and recompute the power spectrum. While in the WebSky simulation, an apodized mask was applied with a 15-arcminute C2 apodization around pixels with sources brighter than the flux cut. The results are shown in the right panel of Figure 11. After applying the flux cut, the recovered power spectra show excellent agreement among the three cases, as

well as with the predictions from [Ade et al. \(2019\)](#) (based on the [Sehgal et al. 2010](#) mocks) and WebSky. This is expected because all the cases considered in our study and other simulations are calibrated to reproduce the fainter end of the differential source count curves at CMB frequencies. Given this consistency, we then use the maps with bright sources removed to further study the cross-correlation between radio point sources and other secondary CMB anisotropies.

To test the influence of bright-source masking on the recovered statistics, we perform a test using the Compton $y - y$ power spec-




Figure 12. Cross-power spectrum between radio flux density (after bright source removal) and other LSS tracers, where the radio luminosities are assigned by abundance matching to the observed radio luminosity function of low-excitation radio galaxies, measured at 150 MHz from the LOFAR survey (see Section 3.6 for details). *Left:* cross-correlation with the Compton y field for three different λ_{Edd} cuts. *Right:* cross-correlation with the CIB map at 143 GHz. Predictions from Ade et al. (2019) (based on the Sehgal et al. 2010 mocks) and WebSky are shown as black and red dashed lines.

trum. Specifically, we replace pixels with radio source flux densities greater than 7 mJy by the average of their nearest neighbors (as used by e.g. Planck Collaboration et al. 2011; Stein et al. 2020). We repeat this test for maps of different resolutions, from $n_{\text{side}} = 512$ up to $n_{\text{side}} = 4096$. We find that for maps with a resolution higher than $n_{\text{side}} = 1024$ (corresponding to an angular resolution of 3.4 arcmin), the true $y - y$ power spectrum can be accurately recovered. We leave the investigation of the effects of different masking schemes on various statistics and their observational implications to a future study.

Figure 12 shows the cross-correlation between radio point sources (after bright source removal) and other LSS tracers. The left panel shows the cross-correlation between tSZ and radio flux density. The FLAMINGO predictions broadly agree in shape with those based on the maps from Ade et al. (2019) and Websky, although with a higher amplitude compared to the other curves. The differences among the three λ_{Edd} cuts are relatively small, which is likely due to the removal of bright and massive sources which typically trace massive clusters that also dominate the tSZ signal. The predicted correlation coefficient¹⁵ values at $\ell = 3000$ for radio-tSZ, $\rho_{\ell=3000}^{y-S143 \text{ GHz}}$, are 0.19, 0.24, 0.11 for the $\lambda_{\text{Edd}} < 10^{-2}$, 10^{-3} and 10^{-6} cases respectively, compared to 0.056 and 0.089 as predicted by Sehgal et al. (2010) and the WebSky simulations.

The right panel shows the cross-correlation between flux density and the CIB at 143 GHz (after bright source removal in the radio catalogues). Compared to the left panel, these curves are much smoother, as the CIB is more sensitive to more common, lower mass haloes, where the bright source cut has less of an impact. The lowest- λ_{Edd} case yields the lowest cross-correlation amplitude. This trend is likely driven by the fact that imposing a lower λ_{Edd} limit preferentially picks out relatively higher mass haloes which tend to have lower star formation activity due to environmental quenching and enhanced AGN feedback. correlation coefficient values at $\ell = 3000$ for radio-CIB, $\rho_{\ell=3000}^{\text{CIB}-S143 \text{ GHz}}$, are 0.056, 0.052, 0.023 for the $\lambda_{\text{Edd}} < 10^{-2}$, 10^{-3} and 10^{-6} cases respectively, compared to

0.08 and 0.074 as predicted by Sehgal et al. (2010) and the WebSky simulations.

5 DISCUSSION: FEEDBACK AND COSMOLOGY DEPENDENCIES

A major advantage of the FLAMINGO suite is that it has a number of variations in feedback modelling and background cosmology. This enables us to explore the dependencies of the key summary statistics on feedback processes and cosmology. In this section, we focus mainly on the CIB statistics (Section 5.1) and the kSZ auto-power spectrum (Section 5.2). For a discussion of such dependencies in other statistics (such as the tSZ effect and its cross correlations), see McCarthy et al. (2023c, 2025).

5.1 Variations in the CIB power spectrum

In previous studies modelling the CIB (e.g. in Planck Collaboration et al. 2014; Stein et al. 2020; Maniyar et al. 2021; Omori 2024), the analysis is typically performed under the assumption of a fixed cosmological model. In addition, the nature of the SED has no direct link to the physics of galaxy formation and feedback making the derived best-fitting parameters difficult to interpret. While this indeed simplifies the modelling, both the CIB power spectrum and the SED modelling encode information about feedback physics and cosmology, as we will demonstrate below.

We adopt two approaches to explore the dependencies of the CIB power spectrum on feedback and cosmological models. In the first approach, we fix the SED parameters to the best-fitting values obtained from the fiducial $(2.8 \text{ Gpc})^3$ run in Section 3.5.1 and use the same set of parameters to compute the CIB statistics for other model variants. As a result, any variations in the resulting CIB-CIB auto- and cross-power spectra are solely driven by differences in the spatial clustering of star-forming galaxies predicted by each FLAMINGO model. This is discussed in Section 5.1.1.

In the second approach, we refit the CIB SED model from each FLAMINGO variant to the measurements from Lenz et al.

¹⁵ See its definition in Section 5.1.2

(2019). This method ensures that the CIB auto- and cross-power spectra are well matched across all models by construction, but it results in a distinct set of best-fitting SED parameters for each case. We further examine how much variation remains in the CIB–LSS cross-correlations under consistent CIB power constraints. Results from this approach are presented in Section 5.1.2.

5.1.1 Fixed SED template

The left panel of Figure 13 shows the CIB auto-power spectrum for different feedback and cosmology models, adopting the best-fitting SED parameters derived from the fiducial run. Here, we show only the curves from the strong stellar feedback ($M_* - \sigma$), strongest AGN (fgas- 8σ), and LS8 models for illustration. In the bottom panel, we show the ratio between different models and the fiducial curve. Only the results at 857 GHz are shown here for illustration, as the other frequencies display the same general trends across models. This comparison clearly shows the sensitivity of the CIB power spectrum to feedback and cosmological variations: the $M_* - \sigma$ model predicts the lowest power amplitude ($\approx 30 - 40\%$ reduction at all scales) due to its strong suppression of star formation activity. The fgas- 8σ model yields a moderately higher power amplitude ($\approx 20\%$) compared to the fiducial model. Here we point out that, while the energy per AGN feedback event was increased to yield a lower gas fraction in galaxy groups in the fgas- 8σ model, the other subgrid feedback parameters (particularly controlling stellar feedback) were also varied to attempt to rematch the observed galaxy stellar mass function. From figure 8 of Schaye et al. (2023), we see that the fgas- 8σ has a slightly higher mass function compared to the fiducial model (see also the SFR density curve in the right panel of Figure 15), which is consistent with the difference between models in the CIB power spectrum. For the LS8 model, although it uses same subgrid physics as the fiducial model, it also predicts a lower power amplitude, primarily due to an overall suppression of matter clustering. However, this suppression is generally less significant than that seen in the strong stellar feedback model. We discuss the right panel of Figure 13 below.

Figure 14 shows the feedback and cosmology dependencies of the CIB–tSZ and CIB– κ cross-power spectra. On large scales, the CIB–tSZ cross spectrum (left panel) is noisy due to the limited number of massive structures that dominate at these scales¹⁶. The LS8 model predicts the lowest power amplitude, which is expected due to the strong σ_8 dependence of tSZ statistics. On small scales, the strongest AGN model shows a suppression of power due to gas depletion within the one-halo regime.

For the CIB– κ cross-power spectrum (right panel), the feedback dependencies follow similar trends as shown in the CIB auto-spectrum though with generally reduced offset amplitudes, which is expected given that the κ statistics are not strongly sensitive to feedback processes. For the LS8 model, however, a similar amplitude offset with respect to the fiducial run is seen in the CIB– κ cross spectrum and the CIB auto-spectrum, which is due to the suppressed overall matter clustering predicted by this low- σ_8 cosmology.

In summary, if the SED is fixed (or precisely known), the CIB auto-power spectrum and the cross-correlation of the CIB with other LSS tracers are sensitive probes of both cosmology and feedback. In reality, however, the SED is not precisely known, and in general,

one often constrains a model of the SED to match observations, an approach which we consider immediately below.

5.1.2 Varying the SED template

The right panel of Figure 13 presents the CIB auto-power spectra for different FLAMINGO feedback and cosmology models, where the SED parameters have now been refitted for each model to match observational data of Lenz et al. (2019). As an illustration, the ratio between each model and the fiducial run for 857 GHz is shown in the bottom panel. In general all the models match the CIB auto-power spectra similarly well. Green data points in the bottom panel are the ratios between the Lenz et al. (2019) measurements and the fiducial curves. The ratio plot suggests that the fiducial model provides a slightly worse fit to the data compared to other model variations. A better fit, however, can be achieved with the *four-parameter* model (see the right panel of Figure A2).

The best-fitting SED parameters for all considered model variations, specifically the present-day dust temperature (T_0) and the spectral index (β_d), are shown in the left panel of Figure 15. This plot highlights how these parameters vary across different physical and cosmological assumptions (see Equation 12 for the SED definition). For reference, the right panel of Figure 15 shows the cosmic SFR density as a function of redshift, computed from all gas cells in each lightcone shell for every model variant.

Interestingly, an anti-correlation is observed between T_0 and β_d , which is consistent with the parameter degeneracies revealed in the MCMC corner plot (Figure 1). Moreover, the $M_* - \sigma$ and LS8 models, which predict the lowest SFRs across redshifts, are also the most extreme outliers in the $T_0 - \beta_d$ parameter space, for models with thermal (non-jet) AGN feedback. These findings demonstrate that the SED fitting procedure is flexible enough to absorb most of the feedback and cosmology dependencies into the SED parameters. This implies that, if one constrains the SED parameters to match the CIB auto-power spectrum, the sensitivity to feedback and cosmology variations is mostly lost.

Importantly, however, this does *not* imply that cross-correlations between the CIB and others tracers of LSS will be insensitive to feedback and cosmology variations. To test this, Figure 16 shows the resulting CIB–tSZ and CIB– κ cross-power spectra for this analysis at 857 GHz. Compared to Figure 14, where SED parameters were fixed, we see there is still tangible (if somewhat reduced) sensitivity to variations in the feedback model and background cosmology. The LS8 model shows a $\approx 10\%$ suppression in the CIB– κ cross-spectrum and up to $\approx 20\%$ in the CIB–tSZ case. Also, the small-scale suppression in the CIB–tSZ cross spectrum is still significant for the fgas- 8σ run relative to the fiducial model.

Figure 17 shows the cross-correlation coefficient for the CIB–tSZ (left) and CIB– κ (right) cross correlations, computed as $\rho_\ell = C_\ell^{\text{CIB-XXX}} / \sqrt{C_\ell^{\text{CIB-CIB}} C_\ell^{\text{XXX-XXX}}}$, where XXX refers to either the tSZ or κ field. Note that the CIB and tSZ statistics are well converged up to redshift $z \approx 3$ (see Figure 9 and discussions in McCarthy et al. 2023c). However, the κ – κ auto-power spectrum still has non-negligible contributions from higher redshifts. Since our κ maps are integrated only up to $z = 4.5$, we correct for the missing high-redshift power using the Limber approximation (e.g. see Kaiser 1992; McCarthy et al. 2023c) and the matter power spectrum from different FLAMINGO model variants, evaluated at particle snapshots up to $z = 30$. For illustration, we only show results for the 545 GHz channel.

The predicted correlation coefficients $\rho_\ell^{\text{CIB-tSZ}}$ are qualita-

¹⁶ It is possible to explore these fluctuations further by, e.g., decomposing the CIB–tSZ cross-correlation by halo mass. We leave this for future work.

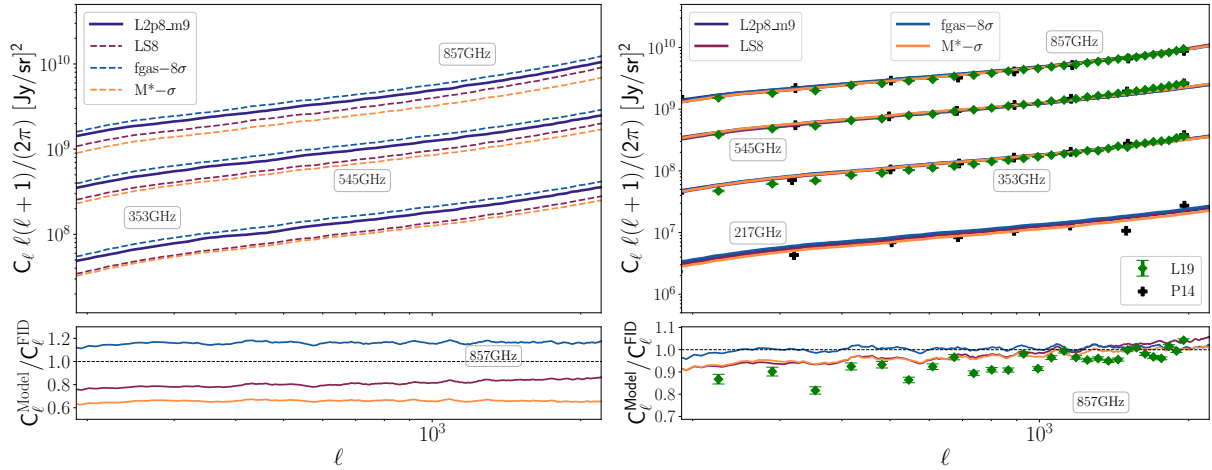


Figure 13. Feedback and cosmology dependencies of the CIB auto-power spectra for the case when SED parameters are fixed to the values from fitting the fiducial $(2.8 \text{ Gpc})^3$ run to the [Lenz et al. \(2019\)](#) data (see Section 3.5.1, *left*), and for the case when SEDs are refitted to the [Lenz et al. \(2019\)](#) data for different models (*right*). Green data points are the measurements from [Lenz et al. \(2019\)](#). The ratio curves between different models and the fiducial curves for 857 GHz are shown on the bottom, along with the ratios between the [Lenz et al. \(2019\)](#) measurements and the fiducial curves for comparison.

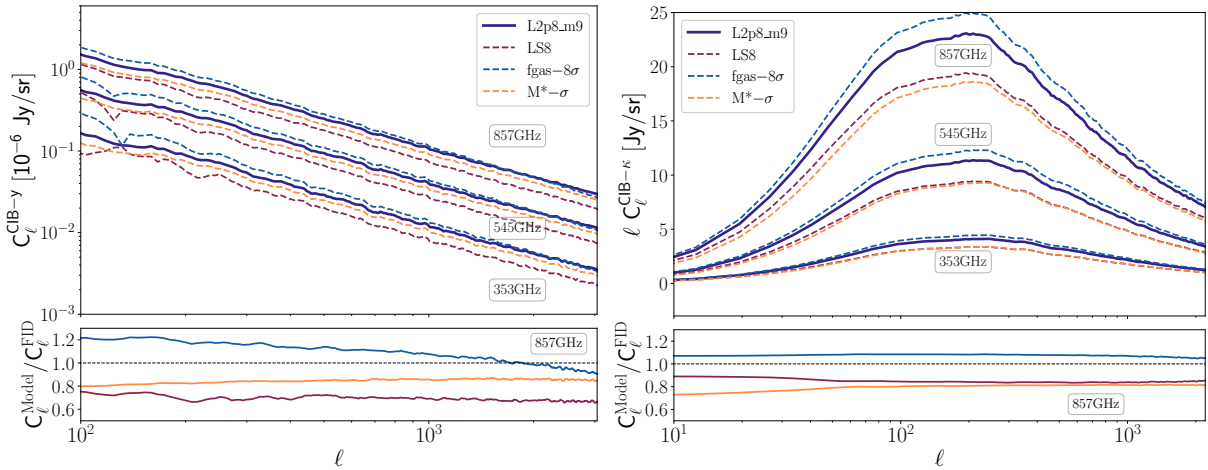


Figure 14. Feedback and cosmology dependencies of the CIB-tSZ (*left*) and CIB- κ (*right*) cross-power spectra for the case when CIB SED parameters are fixed to the values obtained from fitting the fiducial $(2.8 \text{ Gpc})^3$ run to the data from [Lenz et al. \(2019\)](#) (see Section 3.5.1). Ratios between different models and the fiducial curves for 857 GHz are shown on the bottom panel.

tively consistent with those from the AGORA simulations. The values from WebSky are however higher than those from FLAMINGO and AGORA, mainly due to WebSky's lower tSZ auto power (see Figure 6). The $\rho_\ell^{\text{CIB}-\kappa}$ predictions are quite consistent across different models.

Compared with Figure 16, once the model dependence on the auto-power spectra is divided out, the differences of the CIB statistics become weaker than in the cross-power spectra alone. This is expected since the auto-power spectra themselves also carry model-dependent information. Nevertheless, as seen in the $\rho_\ell^{\text{CIB-tSZ}}$ curves, differences of $\sim 10 - 20\%$ still remain among the various FLAMINGO models, and there are larger discrepancies among CMB simulations when different gas-physics prescriptions and cosmology are adopted.

These results suggest that, even after SED-based matching, CIB-LSS cross-correlations are still sensitive to the underlying baryonic physics and cosmology. Although statistics such as the

CIB-tSZ cross-correlation are still difficult to measure with current survey setups, here we simply note that future CIB observations, especially when they are combined with other LSS tracers, may be able to provide meaningful, independent constraints on stellar (CIB auto) and AGN (CIB-LSS cross) feedback models, as well as cosmological parameters.

Note that an alternative approach to constraining the SED of the CIB could be to instead match the CIB cross-correlations and predict the auto-correlations. The optimal strategy for calibration likely depends on the scientific goals (e.g., to best constrain cosmological parameters as opposed to test feedback models).

5.2 Variations in kSZ power spectrum

Figure 18 shows the impact of baryonic feedback and cosmology on the kSZ auto power spectrum. This summary statistic has several interesting trends at intermediate to small angular scales across




Figure 15. *Left:* The best-fitting rest-frame dust temperature (T_0) and spectral index (β_d) for all FLAMINGO model variants considered in this study, with best-fitting values obtained by refitting the modelled CIB power spectra to Lenz et al. (2019) data. *Right:* Cosmic star formation rate density as a function of redshift for the different model variants.

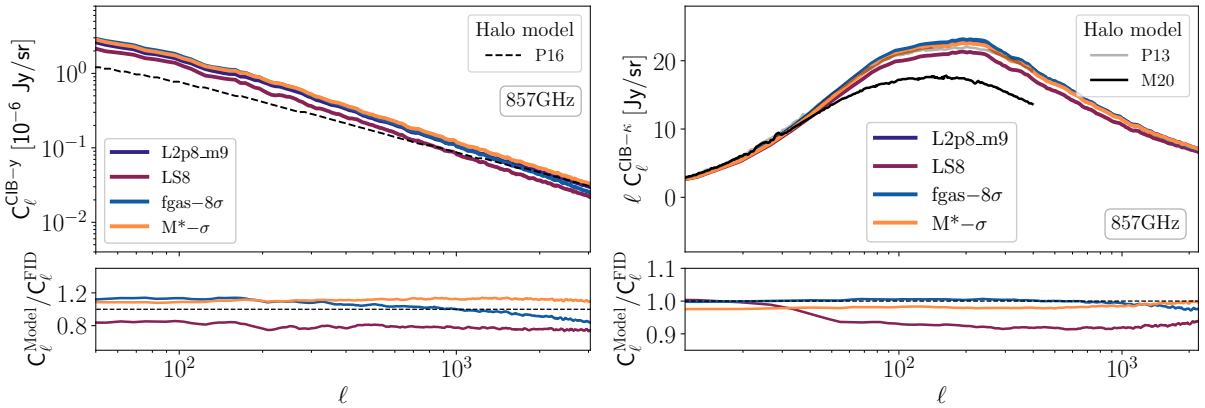


Figure 16. Feedback and cosmology dependencies of the CIB-LSS cross-power spectrum at 857 GHz for the case when SEDs are refitted to the Lenz et al. (2019) data for different models. Ratios between different models and the fiducial curves are shown in the bottom subpanels. The left panel shows the results for the CIB-TSZ cross-power spectrum, and the right panel shows the CIB- κ curves.

the different FLAMINGO models. Notably, the fiducial, $M_* - \sigma$, and jet models yield very similar predictions, particularly at high multipoles. This is not unexpected, since all three of these models have been calibrated to the same gas mass fraction data and the kSZ effect scales in proportion to the integrated gas mass. In contrast, models which systematically vary the gas mass fractions, such as fgas+2 σ and fgas-8 σ , lead to very different behaviours on small angular scales.

In the LS8 cosmology, fewer massive haloes form. This results in a globally lower abundance of dense, coherently moving ionised gas in this model, which is reflected as a significant suppression of the kSZ power spectrum across all scales. As shown in Figure 7, upcoming high-resolution CMB experiments may be sensitive to the kSZ signal at multipoles $l \gtrsim 5000$. Our calculations highlight the potential the kSZ effect power spectrum has as a probe of both feedback and cosmology.

6 CONCLUSION

Secondary anisotropies in the cosmic microwave background (CMB) contain a wealth of cosmological and astrophysical information. In this work, we have used the state-of-the-art FLAMINGO cosmological hydrodynamical simulations to generate a set of realistic mock CMB anisotropy maps, which are derived directly from the simulated properties of matter, gas, and accreting black holes. To our knowledge, this is the first time in which a broad range of anisotropies on full-sky lightcones has been self-consistently produced from cosmological hydrodynamical simulations.

Using the full-sky HEALPix lightcone maps, we have generated mock CMB lensing maps, Sunyaev-Zel'dovich (SZ) effect maps (with and without relativistic corrections), anisotropic screening maps (or equivalently, optical depth τ maps), and cosmic infrared background (CIB) maps, all integrated up to $z = 4.5$. In addition, we have produced radio point source maps from the black hole

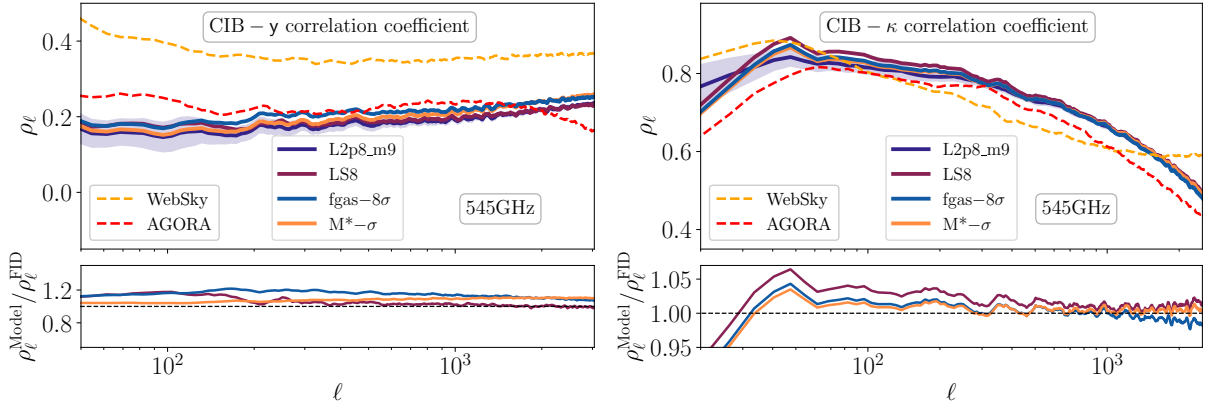


Figure 17. Feedback and cosmology dependencies of the CIB-LSS correlation coefficient at 545 GHz for the case when SEDs are refitted to the [Lenz et al. \(2019\)](#) data for different models, with shaded regions estimated by averaging the results from eight different lightcones. Ratios between different models and the fiducial curves are shown in the bottom subpanels. The left panel shows the results for the CIB-tSZ correlation coefficient, and the right panel shows the CIB- κ correlation coefficient.




Figure 18. The kSZ auto-power spectrum for different FLAMINGO model variants.

lightcone catalogues, integrated up to $z = 2.5$. Descriptions of how each component are modelled and derived from the lightcone-based catalogues are outlined in Section 3, and comparisons between the power spectra from our mock maps and observational data are presented in Section 4. To fully utilise the model variations within the FLAMINGO simulations, we further explored the impact of feedback effects and cosmology on the auto- and cross-correlations of various anisotropies, with a particular focus on CMB lensing, the SZ effect, CIB, and radio point sources. The dependence of different statistics on feedback and cosmology is discussed in Section 5. Our main findings are summarised as follows:

- From the $(2.8 \text{ Gpc})^3$ intermediate-resolution run (L2p8_m9), we recover the general shape of the thermal SZ and kinetic SZ auto-power spectra, which are consistent with the predictions from other hydrodynamical and CMB mock simulations. However, the amplitudes differ due to varying treatments of subgrid physics and intergalactic medium (IGM) modelling (Figures 6 and 7).
- Using a simplified *three-parameter* model, which includes a linear $\text{SFR}-L_{\text{bol,IR}}$ conversion law, a modified blackbody spectral energy distribution (SED) template for infrared sources, and a power

law redshift evolution of the dust temperature, we obtain good fits to the CIB auto- and cross-power spectra for the *Planck* HFI channels (Figure 8). From this fitting, we successfully constrain the SED and dust properties to values that are broadly consistent with those adopted in [Planck Collaboration et al. \(2016c\)](#), although our best-fitting model favours a weaker redshift evolution of the dust temperature (Figure 1).

- Additional complexity can be introduced into this simplified CIB model, such as varying the proportionality constant in the $\text{SFR}-L_{\text{bol,IR}}$ conversion law, or incorporating a dust-mass dependence with a more sophisticated SFR scaling. In these extended frameworks, we are still able to recover both the shape and amplitude of the measured curves (Figure A1). Other improvements in the SED, for example, treating the dust temperature as a function of various galaxy properties, are left for future work.

- The predicted CIB-Compton γ and CIB- κ cross-correlations are in good agreement with both observations and previous halo model predictions (Figure 10).

- We construct radio point-source catalogues from the black hole lightcone data, with luminosities assigned via abundance matching to the 150 MHz radio luminosity function from the LOFAR survey ([Kondapally et al. 2022](#)). A frequency scaling is then applied to convert $L_{150 \text{ MHz}}$ to CMB frequencies, with the scaling factor calibrated against the observed differential source counts (Figure 3). Based on this, we successfully reproduce the radio point-source auto-power spectrum at 143 GHz (Figure 11). Furthermore, we find that the cross-correlations between radio sources and other anisotropies depend on the properties of the radio population (characterised by the Eddington accretion ratio λ_{Edd} of black holes, Figure 12).

- The CIB power spectrum and its cross-correlations with other large-scale structure (LSS) tracers show strong dependence on feedback processes and cosmology. When keeping the CIB modelling parameters fixed across all model variations, we find that the strong stellar feedback model produces the largest suppression of both the CIB auto-correlation and the CIB- κ cross-correlation, which is consistent with the suppression of star formation. AGN feedback suppresses small-scale power in the CIB-Compton γ cross-correlation. Finally, the low- σ_8 (LS8) model consistently yields the lowest amplitudes across all CIB statistics (Figures 13 and 14).

- When refitting the CIB-predicted curves from each model variation to the same set of measured power spectra (and noting that

the various runs can reproduce the CIB power spectra similarly well when the SED is refitted), we find that the best-fitting SED parameters show strong model dependence (Figure 15). In other words, the flexibility/uncertainty in the SED can absorb the dependence of the CIB power spectrum on feedback and cosmology. However, cross-correlations between the CIB and other LSS tracers retain sensitivity to the feedback modelling and background cosmology (Figures 16 and 17) even after the SED parameters have been refitted to the CIB auto power spectra. An alternative approach to SED calibration could be to fit the CIB-LSS cross-correlations and predict the CIB auto power spectra.

- The kinetic SZ power spectrum also shows a strong dependence on baryonic feedback and cosmology, especially at small scales where feedback effects dominate. Increasing the strength of AGN feedback results in the greatest suppression of power on small scales, while the LS8 model yields the lowest amplitude across all scales (Figure 18).

- We also present predictions for the tSZ monopole (Table 2) and the CIB monopole (Figure 9). All predicted y monopoles from the different FLAMINGO models lie well below the direct COBE–FIRAS 95% upper limit. The fiducial model sits slightly above the indirect *Planck* tomographic estimate of Chiang et al. (2020). The strongest AGN model ($f_{\text{gas}} - 8\sigma$ run) yields the largest value, which is in mild tension with the tomographic measurements, while the LS8 cosmology (with fiducial feedback modelling) gives the lowest monopole and is in excellent agreement with the *Planck* measurement from Chiang et al. (2020). For the CIB monopole, using the best-fitting SED from the *three-parameter* model, we show that the CIB intensity converges by $z \approx 3$ with a spectral break above $\nu \gtrsim 1000$ GHz. The predicted CIB monopoles broadly match the *Planck*-based constraints of Chiang et al. (2025).

Our simulations reproduce a wide range of observational constraints with good accuracy and can be used to make forecasts for upcoming CMB surveys such as the Simons Observatory (Ade et al. 2019) and CMB-HD (Sehgal et al. 2019). In particular, we highlight the importance of CIB and radio point-source statistics, as well as their dependence on feedback models and cosmology. We expect these mock CMB maps to be valuable for testing cross-correlations between different secondary anisotropies, and for validating component-separation pipelines used to separate their individual contributions in real observations.

ACKNOWLEDGEMENTS

The authors thank the anonymous referee for helpful suggestions which improved the paper. This work was supported by the Science and Technology Facilities Council (grant number ST/Y002733/1). This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 769130). This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (<https://dirac.ac.uk/>). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, ST/R002371/1 and ST/S002502/1, Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure. JCH is supported by STFC consolidated grant ST/X001075/1.

DATA AVAILABILITY

The maps and catalogues produced as part of this study will be made publicly available on the FLAMINGO data release website upon acceptance of the paper. Additional derived data products can be provided upon reasonable request to the corresponding author.

REFERENCES

- Abbott T. M. C., et al., 2022, *Phys. Rev. D*, **105**, 023520
- Ade P., et al., 2019, *J. Cosmology Astropart. Phys.*, 2019, 056
- Aihara H., et al., 2018, *PASJ*, **70**, S8
- Akino D., et al., 2022, *PASJ*, **74**, 175
- Alonso D., Sanchez J., Slosar A., LSST Dark Energy Science Collaboration 2019, *MNRAS*, **484**, 4127
- Amon A., et al., 2023, *MNRAS*, **518**, 477
- Bagla J. S., Ray S., 2003, *New Astron.*, **8**, 665
- Bahé Y. M., et al., 2022, *MNRAS*, **516**, 167
- Battaglia N., Bond J. R., Pfrommer C., Sievers J. L., 2012, *ApJ*, **758**, 75
- Battaglia N., Hill J. C., Murray N., 2015, *APJ*, **812**, 154
- Berta S., et al., 2011, *A&A*, **532**, A49
- Best P. N., Heckman T. M., 2012, *MNRAS*, **421**, 1569
- Béthermin M., et al., 2012, *ApJ*, **757**, L23
- Béthermin M., et al., 2017, *A&A*, **607**, A89
- Bigwood L., et al., 2026, *arXiv e-prints*, p. arXiv:2510.15822
- Bleem L., et al., 2012, *Journal of Low Temperature Physics*, **167**, 859
- Bolliet B., Comis B., Komatsu E., Macías-Pérez J. F., 2018, *MNRAS*, **477**, 4957
- Bolliet B., et al., 2024, in *mm Universe 2023 - Observing the Universe at mm Wavelengths*. EDP, p. 00008, doi:10.1051/epjconf/202429300008
- Bolliet B., et al., 2025, *arXiv e-prints*, p. arXiv:2507.07346
- Bond J. R., Myers S. T., 1996, *ApJS*, **103**, 1
- Booth C. M., Schaye J., 2009, *MNRAS*, **398**, 53
- Borrow J., Schaller M., Bower R. G., Schaye J., 2022, *MNRAS*, **511**, 2367
- Bouchet F. R., Colombi S., Hivon E., Juszkiewicz R., 1995, *A&A*, **296**, 575
- Braspenning J., et al., 2024, *MNRAS*, **533**, 2656
- Broxterman J. C., et al., 2024, *MNRAS*, **529**, 2309
- Carlstrom J. E., et al., 2011, *PASP*, **123**, 568
- Chabrier G., 2003, *PASP*, **115**, 763
- Chaikin E., Schaye J., Schaller M., Bahé Y. M., Nobels F. S. J., Ploeckinger S., 2022, *MNRAS*, **514**, 249
- Chaikin E., Schaye J., Schaller M., Benítez-Llambay A., Nobels F. S. J., Ploeckinger S., 2023, *MNRAS*, **523**, 3709
- Challinor A., Lasenby A., 1998, *ApJ*, **499**, 1
- Cheng H., Greengard L., Rokhlin V., 1999, *Journal of Computational Physics*, **155**, 468
- Chiang Y.-K., Makiya R., Ménard B., Komatsu E., 2020, *ApJ*, **902**, 56
- Chiang Y.-K., Makiya R., Ménard B., 2025, *arXiv e-prints*, p. arXiv:2504.05384
- Chluba J., Nagai D., Sazonov S., Nelson K., 2012, *MNRAS*, **426**, 510
- Chluba J., Switzer E., Nelson K., Nagai D., 2013, *MNRAS*, **430**, 3054
- Cochrane R. K., Hayward C. C., Anglés-Alcázar D., Somerville R. S., 2023, *MNRAS*, **518**, 5522
- Costello E. E., McCarthy I. G., Salcido J., Helly J. C., McGibbon R. J., Schaller M., Schaye J., 2025, *arXiv e-prints*, p. arXiv:2510.17980
- Coulton W. R., et al., 2024a, *arXiv e-prints*, p. arXiv:2401.13033
- Coulton W. R., et al., 2024b, *arXiv e-prints*, p. arXiv:2410.19046
- Davé R., Anglés-Alcázar D., Narayanan D., Li Q., Rafieferantsoa M. H., Appleby S., 2019, *MNRAS*, **486**, 2827
- De Cia A., Ledoux C., Savaglio S., Schady P., Vreeswijk P. M., 2013, *A&A*, **560**, A88
- De Cia A., Ledoux C., Mattsson L., Petitjean P., Srianand R., Gavignaud I., Jenkins E. B., 2016, *A&A*, **596**, A97
- Dehnen W., 2014, *Computational Astrophysics and Cosmology*, **1**, 1
- Ding B., Pizzati E., Schaye J., Hennawi J. F., McDonald W., Schaller M., 2025, *arXiv e-prints*, p. arXiv:2510.24283
- Dolag K., Komatsu E., Sunyaev R., 2016, *MNRAS*, **463**, 1797

- Driver S. P., et al., 2022, *MNRAS*, 513, 439
- Dunkley J., et al., 2013, *J. Cosmology Astropart. Phys.*, 2013, 025
- Efstathiou G., McCarthy F., 2025, *MNRAS*, 540, 1055
- Elbers W., Frenk C. S., Jenkins A., Li B., Pascoli S., 2021, *MNRAS*, 507, 2614
- Elbers W., Frenk C. S., Jenkins A., Li B., Pascoli S., 2022, *MNRAS*, 516, 3821
- Eriksen H. K., et al., 2006, *ApJ*, 641, 665
- Everett W. B., et al., 2020, *ApJ*, 900, 55
- FOSSIL Collaboration 2022, <https://www.ias.u-psud.fr/en/content/fossil>
- Fabbian G., Lewis A., Beck D., 2019, *JCAP*, 2019, 057
- Fabbian G., Bianchini F., Sabyr A., Hill J. C., Lovell C. C., Thiele L., Spergel D. N., 2025, arXiv e-prints, p. arXiv:2512.03038
- Fabian A. C., 2012, *ARA&A*, 50, 455
- Figueira M., et al., 2022, *A&A*, 667, A29
- Fixsen D. J., Cheng E. S., Gales J. M., Mather J. C., Shafer R. A., Wright E. L., 1996, *ApJ*, 473, 576
- Forouhar Moreno V. J., Helly J., McGibbon R., Schaye J., Schaller M., Han J., Kugel R., 2025, arXiv e-prints, p. arXiv:2502.06932
- Frontiere N., et al., 2025, arXiv e-prints, p. arXiv:2510.03557
- George E. M., et al., 2015, *ApJ*, 799, 177
- Giocoli C., et al., 2016, *MNRAS*, 461, 209
- Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, 622, 759
- Greengard L., Rokhlin V., 1987, *Journal of Computational Physics*, 73, 325
- Hadzhiyska B., Sailer N., Ferraro S., 2025, arXiv e-prints, p. arXiv:2506.17379
- Hahn O., Rampf C., Uhlemann C., 2021, *MNRAS*, 503, 426
- Han J., Jing Y. P., Wang H., Wang W., 2012, *MNRAS*, 427, 2437
- Han J., Cole S., Frenk C. S., Benitez-Llambay A., Helly J., 2018, *MNRAS*, 474, 604
- Hayward C. C., Narayanan D., Kereš D., Jonsson P., Hopkins P. F., Cox T. J., Hernquist L., 2013, *MNRAS*, 428, 2529
- Heckman T. M., Best P. N., 2014, *ARA&A*, 52, 589
- Hilbert S., et al., 2020, *MNRAS*, 493, 305
- Hill J. C., Spergel D. N., 2014, *J. Cosmology Astropart. Phys.*, 2014, 030
- Hill J. C., Battaglia N., Chluba J., Ferraro S., Schaun E., Spergel D. N., 2015, *Phys. Rev. Lett.*, 115, 261301
- Hojjati A., et al., 2017, *MNRAS*, 471, 1565
- Hu W., 2000, *Phys. Rev. D*, 62, 043007
- Hurier G., Macías-Pérez J. F., Hildebrandt S., 2013, *A&A*, 558, A118
- Huško F., Lacey C. G., Schaye J., Schaller M., Nobels F. S. J., 2022, *MNRAS*, 516, 3750
- Kaiser N., 1992, *ApJ*, 388, 272
- Kay S. T., Braspenning J., Chluba J., Helly J. C., Kugel R., Schaller M., Schaye J., 2024, *MNRAS*, 534, 251
- Kennicutt Robert C. J., 1998, *ApJ*, 498, 541
- Klypin A., Yepes G., Gottlöber S., Prada F., Heß S., 2016, *MNRAS*, 457, 4340
- Kondapally R., et al., 2022, *MNRAS*, 513, 3742
- Koukoufilippas N., Alonso D., Bilicki M., Peacock J. A., 2020, *MNRAS*, 491, 5464
- Kramer D., van Engelen A., Cain C., MacCrann N., Trac H., Grayson S., Scannapieco E., Sherwin B., 2025, arXiv e-prints, p. arXiv:2501.07623
- Kugel R., et al., 2023, *MNRAS*, 526, 6103
- Kuhn L., Li Z., Coulton W. R., 2025, arXiv e-prints, p. arXiv:2504.18637
- Kuijken K., et al., 2019, *A&A*, 625, A2
- Kumar A., et al., 2025, arXiv e-prints, p. arXiv:2501.19327
- Lenz D., Doré O., Lagache G., 2019, *ApJ*, 883, 75
- Lewis A., Challinor A., 2006, *Phys. Rep.*, 429, 1
- Lewis A., Challinor A., Lasenby A., 2000, *ApJ*, 538, 473
- Li Q., Narayanan D., Davé R., 2019, *MNRAS*, 490, 1425
- Li Z., Puglisi G., Madhavacheril M. S., Alvarez M. A., 2022, *J. Cosmology Astropart. Phys.*, 2022, 029
- Louis T., et al., 2025, arXiv e-prints, p. arXiv:2503.14452
- Lovell C. C., Geach J. E., Davé R., Narayanan D., Li Q., 2021, *MNRAS*, 502, 772
- Maffei B., et al., 2021, arXiv e-prints, p. arXiv:2111.00246
- Magdis G. E., et al., 2012, *ApJ*, 760, 6
- Maniyar A. S., Béthermin M., Lagache G., 2018, *A&A*, 614, A39
- Maniyar A., Béthermin M., Lagache G., 2021, *A&A*, 645, A40
- Masi S., et al., 2021, arXiv e-prints, p. arXiv:2110.12254
- Mazzolari G., et al., 2024, arXiv e-prints, p. arXiv:2412.04224
- McCarthy F., 2024, arXiv e-prints, p. arXiv:2405.13470
- McCarthy F., Hill J. C., 2024a, *Phys. Rev. D*, 109, 023528
- McCarthy F., Hill J. C., 2024b, *Phys. Rev. D*, 109, 023529
- McCarthy I. G., Le Brun A. M. C., Schaye J., Holder G. P., 2014, *MNRAS*, 440, 3645
- McCarthy I. G., Schaye J., Bird S., Le Brun A. M. C., 2017, *MNRAS*, 465, 2936
- McCarthy F., Madhavacheril M. S., Maniyar A. S., 2023a, *Phys. Rev. D*, 108, 083522
- McCarthy I. G., et al., 2023b, *MNRAS*, 526, 5494
- McCarthy I. G., et al., 2023c, *MNRAS*, 526, 5494
- McCarthy I. G., et al., 2025, *MNRAS*, 540, 143
- McDonald W., Schaye J., Kuijken K., Helly J., Braspenning J., Schaller M., 2026, arXiv e-prints, p. arXiv:2602.02484
- Mead A. J., Tröster T., Heymans C., Van Waerbeke L., McCarthy I. G., 2020, *A&A*, 641, A130
- Miville-Deschênes M.-A., Lagache G., 2005, *ApJS*, 157, 302
- Mroczkowski T., et al., 2019, *SSR*, 215, 17
- Naess S., et al., 2020, *J. Cosmology Astropart. Phys.*, 2020, 046
- Ni Y., et al., 2023, *ApJ*, 959, 136
- Omori Y., 2024, *MNRAS*, 530, 5030
- Planck Collaboration et al., 2011, *A&A*, 536, A18
- Planck Collaboration et al., 2014, *A&A*, 571, A30
- Planck Collaboration et al., 2016a, *A&A*, 594, A15
- Planck Collaboration et al., 2016b, *A&A*, 594, A22
- Planck Collaboration et al., 2016c, *A&A*, 594, A23
- Planck Collaboration et al., 2018, *A&A*, 619, A94
- Planck Collaboration et al., 2020a, *A&A*, 641, A1
- Planck Collaboration et al., 2020b, *A&A*, 641, A4
- Planck Collaboration et al., 2020c, *A&A*, 641, A6
- Planck Collaboration et al., 2020d, *A&A*, 641, A8
- Ploeckinger S., Schaye J., 2020, *MNRAS*, 497, 4857
- Price D. J., 2012, *Journal of Computational Physics*, 231, 759
- Qu F. J., et al., 2024, *ApJ*, 962, 112
- Reichardt C. L., et al., 2021, *ApJ*, 908, 199
- Remazeilles M., Chluba J., 2020, *MNRAS*, 494, 5734
- Remazeilles M., Chluba J., 2025, *MNRAS*, 538, 1576
- Remazeilles M., Bolliet B., Rotti A., Chluba J., 2019, *MNRAS*, 483, 3459
- Robertson N. C., et al., 2021, *A&A*, 649, A146
- Rubiño Martín J. A., et al., 2020, in Zmuidzinas J., Gao J.-R., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 11453, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy X. p. 114530T, doi:10.1117/12.2561309
- Salcido J., McCarthy I. G., Kwan J., Upadhye A., Font A. S., 2023, *MNRAS*, 523, 2247
- Schaller M., et al., 2024, *MNRAS*, 530, 2378
- Schaller M., Schaye J., Kugel R., Broxterman J. C., van Daalen M. P., 2025, *MNRAS*, 539, 1337
- Schaye J., Dalla Vecchia C., 2008, *MNRAS*, 383, 1210
- Schaye J., et al., 2010, *MNRAS*, 402, 1536
- Schaye J., et al., 2015, *MNRAS*, 446, 521
- Schaye J., et al., 2023, *MNRAS*, 526, 4978
- Schaye J., et al., 2025, arXiv e-prints, p. arXiv:2508.21126
- Schutt T., Maniyar A. S., Schaun E., Coulton W. R., Mishra N., 2024, *Phys. Rev. D*, 109, 103539
- Sehgal N., Bode P., Das S., Hernandez-Monteagudo C., Huppenberger K., Lin Y.-T., Ostriker J. P., Trac H., 2010, *ApJ*, 709, 920
- Sehgal N., et al., 2019, in Bulletin of the American Astronomical Society. p. 6 (arXiv:1906.10134), doi:10.48550/arXiv.1906.10134
- Siegel J., et al., 2026, arXiv e-prints, p. arXiv:2509.10455
- Sorini D., Bose S., Davé R., Anglés-Alcázar D., 2024, *The Open Journal of Astrophysics*, 7, 115

- Stein G., Alvarez M. A., Bond J. R., 2019, *MNRAS*, **483**, 2236
- Stein G., Alvarez M. A., Bond J. R., van Engelen A., Battaglia N., 2020, *J. Cosmology Astropart. Phys.*, 2020, 012
- Sunyaev R. A., Zeldovich Y. B., 1972, *Comments on Astrophysics and Space Physics*, **4**, 173
- Sunyaev R. A., Zeldovich I. B., 1980, *ARAA*, **18**, 537
- Terrazas B. A., et al., 2020, *MNRAS*, **493**, 1888
- The Planck Collaboration 2006, *arXiv e-prints*, pp astro-ph/0604069
- Thiele L., et al., 2022, *Phys. Rev. D*, **105**, 083505
- Tröster T., et al., 2022, *AAP*, **660**, A27
- Vargas C., et al., 2023, *arXiv e-prints*, p. arXiv:2310.17535
- Viero M. P., et al., 2013a, *ApJ*, **772**, 77
- Viero M. P., et al., 2013b, *ApJ*, **772**, 77
- Viero M. P., et al., 2013c, *ApJ*, **779**, 32
- Vikram V., Lidz A., Jain B., 2017, *MNRAS*, **467**, 2315
- Vogelsberger M., et al., 2014, *MNRAS*, **444**, 1518
- Wiersma R. P. C., Schaye J., Theuns T., Dalla Vecchia C., Tornatore L., 2009, *MNRAS*, **399**, 574
- Wiseman P., Schady P., Bolmer J., Krühler T., Yates R. M., Greiner J., Fynbo J. P. U., 2017, *A&A*, **599**, A24
- Wu W. L. K., et al., 2019, *ApJ*, **884**, 70
- Yan Z., van Waerbeke L., Wright A. H., Bilicki M., Gu S., Hildebrandt H., Maniyar A. S., Tröster T., 2022, *A&A*, **665**, A52
- Yang T., Cai Y.-C., Cui W., Davé R., Peacock J. A., Sorini D., 2022, *MNRAS*, **516**, 4084
- Zubeldia Í., Chluba J., Battye R., 2023, *MNRAS*, **522**, 5123
- Zubeldia Í., Melin J.-B., Chluba J., Battye R., 2025, *MNRAS*, **539**, 185

APPENDIX A: POWER SPECTRUM FITS FOR THE CIB FOUR-PARAMETER MODEL AND THE EXTENDED MODEL

In this appendix, we present the power spectrum fitting results for the two alternative CIB models: the *four-parameter* model (including a free normalisation parameter in the Kennicutt relation) and a model that depends on both SFR and dust mass, as defined in equations 18 and 19, respectively. Figure A1 shows the CIB auto-power spectrum fits from these two models. The leftmost panel shows the results from the *four-parameter* case. The middle panel shows the results when only star-forming gas is used to compute the dust mass within each halo, while the rightmost panel shows the results when all gas is considered. For comparison, we also show the ratio between the best-fitting curves from different models and those obtained from the default *three-parameter* case at each frequency.

For each model, we are able to recover approximately the same shape and amplitude of the CIB power spectrum. The associated uncertainties can be largely absorbed into the SED modelling. For the *four-parameter* model, we have best-fitting SED parameters of $\beta_d = 2.65 \pm 0.11$, $T_0 = 21.00 \pm 0.95$ K, $\alpha = 0.07 \pm 0.01$, and $N_s = -0.26 \pm 0.03$. For the *extended* model, we have $\beta_d = 2.49 \pm 0.05$ (2.56 ± 0.04), $T_0 = 21.55 \pm 0.24$ K (20.35 ± 0.25 K), $a = 1.23 \pm 0.01$ (1.10 ± 0.02), and a weak T_{dust} redshift evolution constraint for the star-forming gas (all gas) case, respectively.

Figure A2 shows the feedback dependence of the CIB auto-power spectrum for the *four-parameter* case. By comparing this plot with Figure 13 (for the default *three-parameter* fit), we find that in general, introducing additional degrees of freedom into the CIB model does not lead to a significant improvement in the quality of the best-fitting power spectra when compared to the measured data, and the feedback dependencies of the CIB statistics are comparable to the results obtained from the *three-parameter* case.

APPENDIX B: LINEAR SCALE OF THE CIB STATISTICS

In this appendix, we show the ℓC_ℓ^{CIB} and $\ell C_\ell^{\text{CIB-y}}$ with a linear y-axis scale to highlight the model differences more clearly. While the FLAMINGO simulations provide a qualitatively good fit to the observational measurements that improves on previous models, it is clear from visual inspection of Fig. B1 that the best-fit model does not provide a formally good fit from a statistical point of view, given the statistical uncertainties on the measurements. Note, however, that no account has been taken for any possible sources of systematic uncertainty.

APPENDIX C: EFFECT OF FLUX DENSITY CUT ON THE CIB STATISTICS

In this appendix, we evaluate the impact of applying a simple flux-density cut on the CIB maps used in our analysis. The goal is to test whether bright infrared sources have a significant effect on the recovered CIB auto- and cross-power spectra for the three models considered in this work: FLAMINGO (blue), WebSky (orange), and AGORA (red), as shown in Figures C1 and C2. These figures demonstrate the comparison between the unmasked and masked cases for the CIB power spectrum, the CIB- κ cross-spectrum, and the CIB-tSZ cross-spectrum, respectively.

For simplicity, we adopt a single flux-density threshold of 400 mJy for all frequencies and all models (unlike the frequency-dependent cuts as used in the *Planck* analysis, see Planck Collaboration et al. 2014). Pixels with values exceeding the 400 mJy threshold are replaced with the mean of the remaining pixels below the cut, which is the same refilling strategy as implemented in WebSky. For a fair comparison, the masks and refilling procedure are applied consistently across all simulations. For both the FLAMINGO and AGORA models, we find that the flux-density cut only has minor changes in the recovered CIB power spectra. The masked and unmasked curves remain nearly identical across all frequencies and for all cross-correlations with κ and tSZ. This suggests that our model fitting is robust to such masking choices, and that flux-density cuts are not a major systematic uncertainty for the FLAMINGO-based CIB predictions.

However, the flux-density cut produces a substantial suppression in the WebSky CIB auto-spectrum, as well as noticeable changes in some of the cross-spectra. As discussed in Section 4.2, the CIB modelling adopted by WebSky appears to assign higher infrared signals to low-redshift haloes. When these bright resolved low- z structures are masked and refilled, the WebSky CIB power is significantly reduced, as shown in Figure C1.

Despite this strong effect on the CIB auto-/cross-spectrum, the WebSky CIB- κ cross-spectrum remains largely unaffected. This is expected because this statistic is more sensitive to the overall matter distribution and a wider range of redshifts, which is not strongly altered by removing bright, low-redshift sources. For the CIB-tSZ cross-spectrum, however, the most visible difference appears at low multipoles in 857 GHz. This is again consistent with the removal of bright, low-redshift structures, which contributes to both CIB and tSZ signals. Masking these objects removes correlated large-scale structure, leading to a slightly lower masked spectrum compared to the unmasked case.

Overall, this test shows that our FLAMINGO-based CIB modelling is stable under this simplistic flux-density cuts, whereas the WebSky predictions are more sensitive to masking due to their bright low- z population. These results indicate that our model fitting is not strongly impacted by point-source masking choices.




Figure A1. CIB auto-power spectrum fit using the *four-parameter* model (see Equation 18) and an extended model using both SFR and dust mass (see Equation 19). Black data points show the measurements from Lenz et al. (2019) and Planck Collaboration et al. (2014) for comparison. The left panel shows the best-fitting curves from the *four-parameter* model. The middle panel corresponds to the case where only star-forming gas within each halo (an aperture size of 30 kpc) is used when computing the dust mass, while the right panel shows the result when all gas is included. For reference, the bottom subpanels show the ratios between the predicted best-fitting curves at each frequency from the varied model and the default *three-parameter* curves.

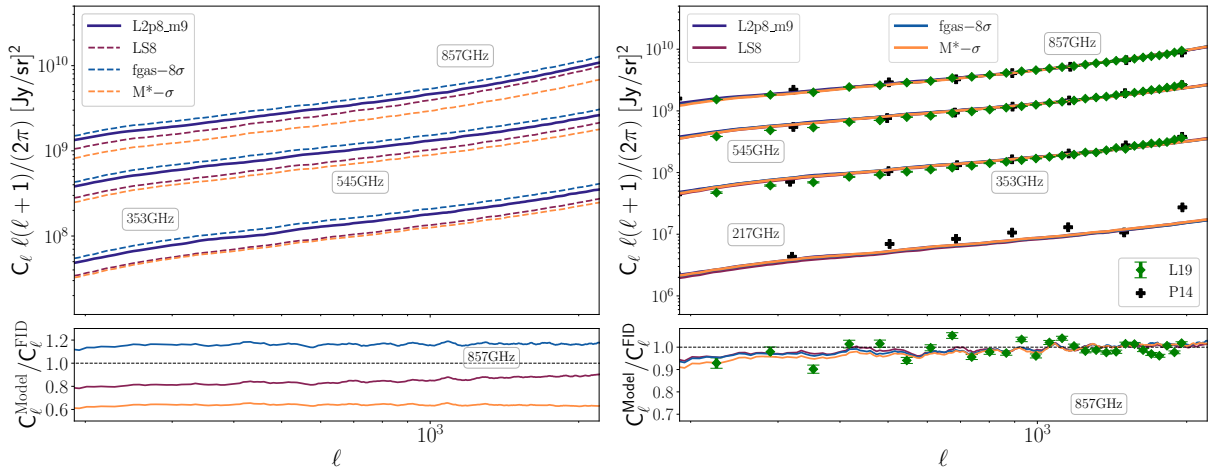


Figure A2. As Figure 13, but for the *four-parameter* case as defined in Equation 18.

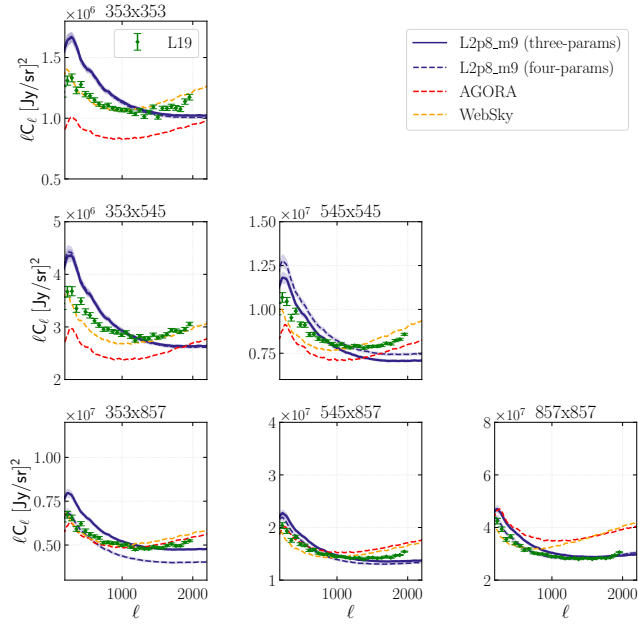


Figure B1. C_ℓ^{CIB} plotted on a linear scale (solid: best-fitting curves from the *three-parameter* model; dashed: best-fitting curves from the *four-parameter* model). For the logarithmic-scale versions, please see Figures 8. Shaded regions are the cosmic variance estimated by averaging the results from eight different lightcones. Green data with error bars are measurements from Lenz et al. (2019).




Figure B2. As Figure B1, but for $C_\ell^{\text{CIB-}\gamma}$. For the logarithmic-scale versions, please see Figures 10.

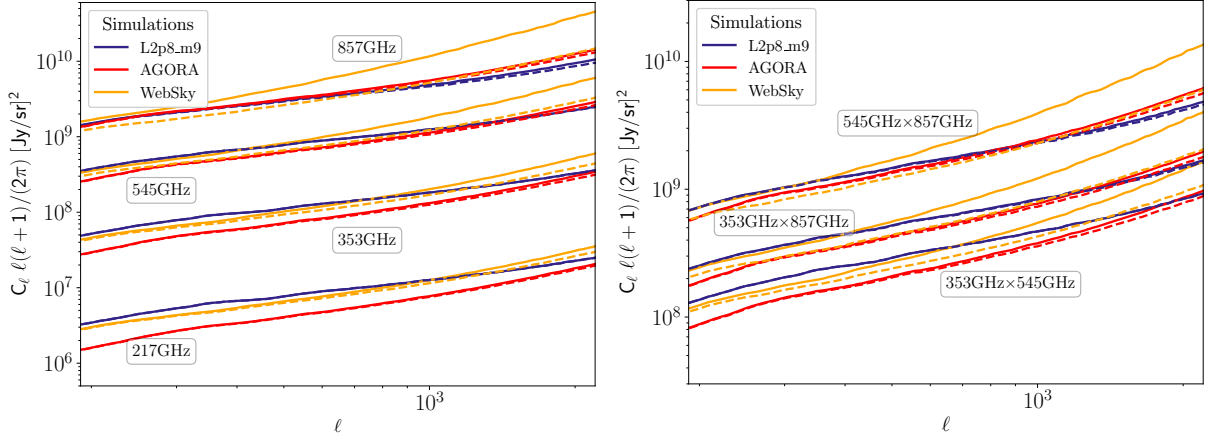


Figure C1. The CIB auto- (*left*) and cross- (*right*) power spectra measured from the masked (dashed) and unmasked (solid) CIB maps from the FLAMINGO fiducial $(2.8 \text{ Gpc})^3$ run, the AGORA (red) and the WebSky (orange) simulations.




Figure C2. CIB-tSZ (*left*) and CIB- κ (*right*) cross-power spectra at different frequencies, generated from the masked (dashed) and unmasked (solid) CIB maps.