

Fermilab

Map-level baryonification: unified treatment of weak lensing
two-point and higher-order statistics

FERMILAB-PUB-25-0351-PPD

arXiv:2505.07949

This manuscript has been authored by Fermi Forward Discovery Group, LLC
under Contract No. 89243024CSC000002 with the U.S. Department of Energy,
Office of Science, Office of High Energy Physics.

Map-level baryonification: unified treatment of weak lensing two-point and higher-order statistics

Alan Junzhe Zhou^{1,3} Marco Gatti^{2,3} Dhayaa Anbajagane^{2,3}
Scott Dodelson^{2,3,4} Matthieu Schaller^{5,6} Joop Schaye⁵

¹Department of Physics, University of Chicago, Chicago, IL 60637, USA

²Department of Astronomy and Astrophysics, University of Chicago, Chicago, IL 60637, USA

³Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA

⁴Fermi National Accelerator Laboratory, P.O. Box 500, Batavia, IL 60510, USA

⁵Leiden Observatory, Leiden University, PO Box 9513, 2300 RA Leiden, the Netherlands

⁶Lorentz Institute for Theoretical Physics, Leiden University, PO box 9506, NL-2300 RA Leiden, the Netherlands

E-mail: ajzhou@uchicago.edu

Abstract. Precision cosmology benefits from extracting maximal information from cosmic structures, motivating the use of higher-order statistics (HOS) at small spatial scales. However, predicting how baryonic processes modify matter statistics at these scales has been challenging. The baryonic correction model (BCM) addresses this by modifying dark-matter-only simulations to mimic baryonic effects, providing a flexible, simulation-based framework for predicting both two-point and HOS. We show that a 3-parameter version of the BCM can jointly fit weak lensing maps' two-point statistics, wavelet phase harmonics coefficients, scattering coefficients, and the third and fourth moments to within 2% accuracy across all scales $\ell < 2000$ and tomographic bins for a DES-Y3-like redshift distribution ($z \lesssim 2$), using the FLAMINGO simulations. These results demonstrate the viability of BCM-assisted, simulation-based weak lensing inference of two-point and HOS, paving the way for robust cosmological constraints that fully exploit non-Gaussian information on small spatial scales.

Contents

1	Introduction	1
2	Baryonification of simulations with BARYONFORGE	3
3	Simulations	6
3.1	FLAMINGO simulations	6
3.2	DES-like weak lensing convergence maps	7
4	Statistics	8
4.1	Power Spectrum	8
4.2	Scattering Transform Coefficients (ST)	8
4.3	Wavelet Phase Harmonics Coefficients (WPH)	10
4.4	Higher Order Moments	11
4.5	The statistics of FLAMINGO and baryonified convergence maps	12
5	Method	12
5.1	Emulator	14
5.2	Sampling	14
6	Results	15
6.1	The 10-parameter baryon correction model	15
6.2	Baryon correction model with restricted degrees of freedoms	17
6.3	Baryon parameters	19
7	Conclusion	20
A	Scale cut	27
B	Statistics as a function of baryon parameters	29
C	Emulator	29
D	Best-fit baryon parameters	30
E	Baryonification error for strong and weak feedback variants	32
F	Cosmology dependence of baryonic suppression	32

1 Introduction

Weak gravitational lensing directly probes the total matter distribution in the Universe, making it a powerful cosmological tool. Unlike galaxy clustering, which is biased by

galaxy-matter relationships, weak lensing is sensitive to all matter, both dark and baryonic. Measurements of two-point weak lensing statistics have provided some of the tightest constraints on cosmological parameters, particularly $S_8 = \sigma_8 \sqrt{\Omega_m/0.3}$, which captures the degeneracy between the amplitude of matter fluctuations (σ_8) and matter density (Ω_m) (Amon et al. 2022a; Secco et al. 2022; Li et al. 2023; Dalal et al. 2023; Dark Energy Survey and Kilo-Degree Survey Collaboration et al. 2023; Anbajagane et al. 2025b; Wright et al. 2025b).

There is growing interest in extracting non-Gaussian features from the lensing map, both at the field level (Alsing et al. 2017; Loureiro et al. 2022; Boruah & Rozo 2023; Zhou et al. 2024a; Dai & Seljak 2024; Zhou et al. 2024b) and via higher-order statistics (HOS). In particular, many HOS have been successfully applied to data have shown significant gains in constraining power (e.g., Fluri et al. 2019; Gatti et al. 2022; Jeffrey et al. 2024; Marques et al. 2024; Novaes et al. 2024; Gatti et al. 2024; Cheng et al. 2025). They are also expected to improve our constraining power on extended cosmological models, such as those with primordial non-Gaussianities (Anbajagane et al. 2024b) or modified theories of gravity (Liu et al. 2016; Davies et al. 2024). Unlike two-point statistics, however, most HOS lack analytical predictions and rely on realistic weak lensing simulations for modeling.

Weak lensing observables – both standard two-point and HOS – probe relatively low redshifts and small physical scales, making them particularly sensitive to non-linear structure formation and astrophysical processes such as supernova and active galactic nucleus feedback, collectively referred to as baryonic feedback. These astrophysical processes alter both the distribution and thermodynamics of baryons, and affect the dark matter distribution through gravitational coupling (e.g., Gnedin et al. 2004; Duffy et al. 2010; Van Daalen et al. 2011; Semboloni et al. 2013; Anbajagane et al. 2022a; Shao et al. 2022; McCarthy et al. 2023; Broxterman et al. 2024; Schaller et al. 2025; Schaller & Schaye 2025). These effects are important on small and intermediate scales, extending into the quasi-linear regime ($\lesssim 10$ Mpc), where they contribute significantly to the lensing signal. If unaccounted for, baryonic effects can bias cosmological inferences (Semboloni et al. 2011; Yoon et al. 2019; Gatti et al. 2022; Amon & Efstathiou 2022; Anbajagane et al. 2023). For example, the observed discrepancy in S_8 between weak lensing and CMB measurements could arise from such mis-modeling rather than a fundamental deviation from the standard cosmological model. While a very conservative approach is to remove scales affected by baryons or to implement projection schemes that mitigate their impact (Amon et al. 2022a; Secco et al. 2022; Piccirilli et al. 2025; Truttero et al. 2025), accurate modeling remains the preferred solution.

On the modeling side, the baryonic imprints on the matter density field are typically studied using high-resolution cosmological hydrodynamical simulations (see Vogelsberger et al. 2020, for a review). However, there is significant freedom in defining the “subgrid” physics in these simulations – processes like star formation and collisional cooling that occur below the simulation’s resolution scale. Different but *equally well-motivated* subgrid

prescriptions result in a wide range of predictions for the distribution of galaxy properties (e.g., Velliscig et al. 2014; Van Daalen & Schaye 2015; Debackere et al. 2020; Anbajagane et al. 2020; Lim et al. 2021; Lee et al. 2022; Stiskalek et al. 2022; Anbajagane et al. 2022a,b; Shao et al. 2022; Gebhardt et al. 2023), and consequently for the baryonic imprints on the non-linear matter power spectrum (e.g., Van Daalen et al. 2011; Chisari et al. 2018; van Daalen et al. 2020; Amon & Efstathiou 2022; Schaller & Schaye 2025). While it is computationally unfeasible to explore broad cosmological and baryonic parameter spaces with these simulations, they are nonetheless essential for validating and calibrating simpler, more practical parametric models of baryonic effects.

Several parametric models have been developed to flexibly describe the impact of baryons on weak lensing observables. A well-known example is HMCode2020 (Mead et al. 2020), an augmented halo model designed to accurately predict the non-linear matter power spectrum across a wide range of cosmologies. To account for baryonic feedback, HMCode2020 modifies the concentration–mass relation through a parametric model that accounts for mass and redshift-dependence and is fitted to the matter power spectra of the BAHAMAS simulations (McCarthy et al. 2017). More recent examples include the power spectrum emulator trained on the FLAMINGO simulations (Schaller et al. 2025) and a wavenumber-dependent rescaling of the difference between the non-linear and linear gravity-only power spectra (Amon & Efstathiou 2022; Schaller & Schaye 2025). Another successful approach is the baryonic correction model (BCM), a phenomenological method that mimics baryonic effects by perturbing the density field of dark matter-only (DMO) simulations via post-processing (Schneider & Teyssier 2015; Schneider et al. 2019; Anbajagane et al. 2024a). These modified outputs can then be used to build emulators of the matter power spectrum that include such baryonic imprints. The BCM has demonstrated sufficient flexibility to reproduce a wide range of feedback scenarios seen in hydrodynamical simulations (Schneider et al. 2019; Aricò et al. 2020; Giri & Schneider 2021), and has already been applied to analyze two-point statistics from wide-field surveys (Chen et al. 2023; Aricò et al. 2023; Bigwood et al. 2024b).

Most of the recipes described above have been extensively tested in the context of two-point statistics. In principle, the BCM approach – via its use of N-body simulations – allows the construction of emulators for any observable, including HOS. However, the robustness of the model predictions for beyond two-point functions, in the context of analyzing current and future weak lensing datasets, remains largely unproven. Some studies have begun to explore this direction: Semboloni et al. (2013) proposed a halo-based model which corrects the baryonic effects on the two and three-point functions. Aricò et al. (2021) showed that the BCM can simultaneously fit the matter power spectrum and bispectrum of hydrodynamical simulations for snapshots at different redshifts. Lee et al. (2023) compared weak lensing peak counts from the full hydrodynamical simulation IllustrisTNG to those from a baryon-corrected version of its dark matter-only (DMO) counterpart, IllustrisTNG-Dark. They found that the BCM provided an insufficient fit when the noise levels match those expected from future surveys like LSST and Euclid. However, their analysis relied on BCM parameters calibrated solely to the power spectrum,

whereas a joint fit that included peak counts could have improved performance. Moreover, they did not consider other types of HOS. More recently, [Anbajagane et al. \(2024a\)](#) performed a joint fit of second-, third-, and fourth-moments of the matter density and gas pressure fields measured in IllustrisTNG, finding statistical agreement with the model on scales above ~ 1 Mpc and across multiple redshifts. However, this test was also limited to projected slices at fixed redshifts and did not validate against realistic weak lensing observables.

We show that the BCM is sufficiently flexible for simulation-based analyses by applying it to the FLAMINGO simulations, a state-of-the-art suite that features multiple baryonic feedback models, matching DMO runs, and is calibrated against observed $z = 0$ galaxy stellar mass function and the galaxy cluster gas fractions ([Schaye et al. 2023](#); [Kugel et al. 2023](#)). We baryonify the DMO simulations, generate realistic weak lensing maps, and build an emulator for both two-point and higher-order statistics (HOS) — including scattering transform coefficients ([Cheng et al. 2020](#)), wavelet phase harmonics ([Allys et al. 2020](#)), and third- and fourth-order moments. The emulator accurately recovers the summary statistics of the hydrodynamical runs, and we show that appropriate BCM parameter choices can jointly fit all statistics across the FLAMINGO feedback models down to $\ell \approx 2000$.

The paper is organized as follows. [Sec. 2](#) discusses the BCM model. [Sec. 3](#) describes the FLAMINGO simulations and the DES Y3-like tomographic weak lensing convergence maps. [Sec. 4](#) defines the relevant HOS and how they are measured in simulations. [Sec. 5](#) describes the procedure for fitting the BCM model to the FLAMINGO feedback variants, and [Sec. 6](#) presents the results. We conclude in [Sec. 7](#). Supporting figures and tables are provided in the Appendix.

2 Baryonification of simulations with BARYONFORGE

The BCM introduced in [Schneider & Teyssier \(2015\)](#) is a phenomenological framework that does not require specifying any small-scale subgrid physics. The method is parameterized by pairs of halo density profiles: one profile representing a halo in a DMO simulation, and another representing the halo in a hydrodynamical simulation (DMB). The latter includes contributions from dark matter (DM) and baryons (B), in the form of stars, gas, satellite galaxies, and dark matter. There exist different models for the exact form and parameterization of these profile components ([Schneider & Teyssier 2015](#); [Schneider et al. 2019](#); [Aricò et al. 2020](#); [Debackere et al. 2020](#); [Pandey et al. 2024](#)). Here we use the model of [Giri & Schneider \(2021\)](#), which is a slightly modified version of the model in [Schneider et al. \(2019\)](#). We refer the reader to those works for detailed discussions. Briefly, the total matter in the DMO case follows a Navarro–Frenk–White (NFW) profile ([Navarro et al. 1996, 1997](#)) with a power-law cut-off at large scales, $r > 4R_{200c}$. The stellar density follows an exponential profile, while the gas density follows a modified generalized NFW (gNFW) profile. The collisionless matter (dark matter and satellite galaxies) is modeled as a modified NFW profile that accounts for adiabatic relaxation due to the presence

of the star and gas components. In our implementation, each halo is assigned a NFW concentration using the semi-analytic model of [Diemer & Kravtsov \(2015\)](#).

The DMO and DMB profiles can be expressed via the quantity $M(r)$, the mass enclosed within a radius r , which is a monotonically increasing function. The DMO and DMB enclosed mass profiles are then combined to compute a displacement function,

$$\Delta d(r) = M_{\text{DMB}}^{-1}(M_{\text{DMO}}(r)) - r. \quad (2.1)$$

where the functions above are implicitly also functions of the halo mass, M_{200c} . Here, the function M_{DMB}^{-1} takes a mass as its argument, and returns the radius that encloses that mass under the DMB model. This displacement function represents the radial offset required to transform a DMO matter distribution into a DMB one.

The BCM proceeds as follows. We start with a DMO simulation with a halo catalog. For each halo above a the mass threshold $10^{13} M_{\odot}/h^1$, we identify all (dark matter) particles within a predetermined radius. Then, given the halo mass M_{200c} and the radial separation of the particles from the halo, we evaluate the displacement $\Delta d(r)$. Next, we radially² offset the particles by $\Delta d(r)$ from the halo center. The displacements are generally positive on the scale of the halo radius, as the dark matter particles must be moved outwards to mimic the impact of extended gas distributions, and are generally negative on scales a hundredth of the halo radius, where the dark matter particles must coalesce to mimic the centrally peaked stellar distribution. The final particle distribution represents a density field that mimics the impact of baryonic processes.

One can then vary the parameters controlling the mass profiles and thereby obtain different corrections to the simulated density field. The profiles specified by a given model can then also be used to predict other quantities, such as the integrated stellar/mass as a function of halo mass. The BCM approach can simultaneously fit variations in the matter power spectrum alongside other quantities such as the scaling relations of gas and star fractions with halo mass, the halo gas profiles, etc. ([Schneider et al. 2019](#); [Giri & Schneider 2021](#); [To et al. 2024](#); [Bigwood et al. 2024a](#)).

Our discussions thus far describe the application of BCM to three-dimensional particle snapshots. However, for the purpose of weak lensing, it is computationally expensive to apply the BCM to each individual snapshot. To produce accurate lensing convergence maps, we require fine redshift resolution, e.g., $\mathcal{O}(100)$ snapshots, which increases the computational load of this task. The snapshot-based approach is also particularly expensive for the purpose of simulation-based inference, where of order $\mathcal{O}(10^4)$ simulations are needed to infer cosmological constraints. Therefore, we use the method outlined in [Anbajagane et al. \(2024a\)](#) which circumvents the computational cost by

¹We chose a threshold of $10^{13} M_{\odot}/h$ for computational reasons: the halo mass function is dominated by low-mass halos, which would exponentially increase the pipeline’s runtime. The impact of excluding these small halos is discussed in the results section.

²[Anbajagane et al. \(2024a\)](#) found the impact of anisotropic (elliptical) offsets on the moments of the density field to be sub-percent (see their Figure 7).

applying the BCM directly on 2D density shells on the full-sky.³ This is done by rewriting the model in terms of *projected* profiles and distances, rather than three-dimensional ones. Instead of particles, our units are the 2D pixels in the full-sky maps. For each pixel, we apply displacements from all halos — where the given pixel is contained within an aperture of $20R_{200c}$ around each selected halo — and accumulate the contributions from all halos to compute the total displacement of the pixel. This extension allows for efficient modeling of higher-order correlations in weak lensing observables by making the BCM more scalable, and therefore makes the method more easily applicable to large datasets from current and upcoming cosmological surveys.

All baryonification in this work is carried out using the publicly available BARYONFORGE⁴ package (Anbajagane et al. 2024a). We direct readers to Anbajagane et al. (2024a) for more details on the model and method choices. For reference, generating 12 different baryonified simulations, each simulation consisting of 60 tomographic density field shells with $N_{\text{side}} = 2,048$, takes approximately 4 hours on an AMD EPYC 7763 CPU node with 128 cores.

The BCM parameters are listed in Table 1, along with their ranges and descriptions. As we will see, the three most important parameters – the ones that can reproduce most of the variation across the hydrodynamic variants – are:

- M_c : a mass scale at which the characteristic slope of the gas profile is $\beta = 1.5$. Halos above/below this mass-scale have more/less cored gas profiles.
- θ_{ej} : a physical scale quantifying how extended the gas distribution is, expressed in units of R_{200c} of the halo. We will see that this ratio is of order 5, so the gas has quite an extended distribution.
- η : The power-law scaling of the stellar mass fraction with halo mass, $M_{\text{star}}/M_{200c} \propto M_{200c}^\eta$.

3 Simulations

3.1 FLAMINGO simulations

Our analysis relies on the FLAMINGO suite of cosmological hydrodynamical simulations, all of which share the initial conditions of a DMO simulation. We provide a brief summary here; for more details, see Schaye et al. (2023); Kugel et al. (2023). The simulations were run using the SPHENIX smoothed particle hydrodynamics implementation (Borrow et al. 2022) within SWIFT (Schaller et al. 2024). FLAMINGO includes radiative cooling and heating (Ploekinger & Schaye 2020), star formation (Schaye & Dalla Vecchia 2008), time-dependent stellar mass loss (Wiersma et al. 2009), massive neutrinos (Elbers et al. 2021), and kinetic supernova and stellar feedback (Dalla Vecchia & Schaye 2008; Chaikin

³An earlier version of this approach is also discussed in Fluri et al. (2022), but does not contain a number of modeling elements introduced in Anbajagane et al. (2024a).

⁴<https://github.com/DhayaaAnbajagane/BaryonForge>

Param	Prior	Description
ν_{M_c}	[-5, 5]	The redshift scaling of the gas mass scale.
$\nu_{\theta_{ej}}$	[-5, 5]	The redshift scaling of the ejection radius.
γ	[0.1, 8]	The slope of the generalized NFW profile at $r \sim R_{ej}$.
$\log_{10} \theta_{co}$	[-4.0, -0.1]	The logarithm of the radial scale of the collapsed gas, $R_{ej} = \theta_{co} R_{200c}$.
$\log_{10} \eta \delta$	[-3.0, 0.0]	The power-law scaling for the central galaxy stellar fraction–halo mass relation.
δ	[2, 20]	The slope of the generalized NFW profile at $r \gg R_{ej}$.
μ_β	[-3, 3]	The mass-dependent scaling of β , which controls the slope between $R_{co} < r < R_{ej}$.
$\log_{10} \frac{M_c}{M_\odot}$	[12.0, 16.0]	The logarithm of the mass-scale where the characteristic gas slope transitions.
$\log_{10} \eta$	[-3.0, 0.0]	The logarithm of the power-law scaling of stellar fraction–halo mass relation.
θ_{ej}	[0.8, 20]	The radial scale of the ejected gas, $R_{ej} = \theta_{ej} R_{200c}$.

Table 1: The input parameters to the profiles of the BCM. For each parameter, we list its prior range and a description of the parameter’s physical meaning. See [Anbajagane et al. \(2024a\)](#) for the specific equations that these parameters govern. The prior ranges are similar to those in [Anbajagane et al. \(2024a\)](#); [Pandey et al. \(2024\)](#) and are used in Sec. 5.1 to construct samples of the baryon-corrected maps and to train the BCM emulator.

[et al. 2023](#)). The simulations are calibrated to match the galaxy stellar mass function at $z = 0$ and the gas fraction of low- z clusters and groups of clusters inferred from X-Ray and weak lensing observations ([Kugel et al. 2023](#)).

We focus on AGN feedback, implemented in two forms: thermal ([Booth & Schaye 2009](#)), where particles in the AGN vicinity are heated, and kinetic jet-like, where particles are given momentum kicks along the black hole (BH) spin axis ([Huško et al. 2022](#)). We consider nine variations in AGN feedback strength from the models available in the FLAMINGO simulations, all in 1 Gpc^3 volumes:

- **Fiducial:** This model features thermal, AGN feedback without jets. Outflows follow paths of least resistance, forming buoyant high-entropy gas bubbles in clusters. Energy is injected into the nearest gas particle, reducing feedback at large distances compared to mass-weighted schemes. The AGN heating temperature, ΔT_{AGN} , is calibrated to match observations such as the cluster gas fractions. This model is labeled Fiducial (L1_m9) per FLAMINGO convention.
- **Thermal AGN – cluster gas and galaxy stellar mass function variations:** We include **6 variations** of the fiducial model where AGN feedback remains thermal, but the simulations feature different cluster gas fractions (Δf_{gas}) and/or galaxy stellar mass functions ($\Delta \Phi_*$), shifted by a specified number of observational standard deviations

(σ). These variations lead to different AGN feedback strengths. The models are labeled $\text{fgas}+2\sigma$, $\text{fgas}-2\sigma$, $\text{fgas}-4\sigma$, $\text{fgas}-8\sigma$, $\text{M}-\sigma$, and $\text{M}-\sigma_{\text{fgas}-4\sigma}$.

- Jet-like AGN Feedback: The **two additional models**, Jet and Jet_ $\text{fgas}-4\sigma$, replace the Fiducial thermal AGN feedback with kinetic, jet-like feedback. Compared to the Fiducial model, Jet has slightly weaker stellar feedback, while Jet_ $\text{fgas}-4\sigma$ features higher jet velocity.

We apply the baryonification procedure described in Sec. 2 to the full-sky spherical overdensity shells of the DMO simulation to generate baryonified realizations. All simulations – both hydrodynamical and baryonified – adopt the best-fit Λ CDM cosmology from DES Y3 “ $3\times 2\text{pt} + \text{All Ext.}$ ” (Abbott et al. 2022).

Finally, to investigate the cosmology dependence of baryonic feedback in Sec. 6, we also consider simulations run at the “LS8” cosmology (Amon et al. 2022b; Schaye et al. 2023), which has a lower amplitude of the power spectrum: $S_8 = 0.766$ compared to 0.815 for the fiducial cosmology. The simulations include DMO (LS8_DMO), the fiducial AGN (LS8), and the strong AGN (LS8_ $\text{fgas}-8\sigma$) variants. We did not apply the BCM to the LS8_DMO simulation and used these three simulations exclusively for the cosmology dependence analysis in Sec. 6.

3.2 DES-like weak lensing convergence maps

We post-process the FLAMINGO simulations to generate weak lensing convergence maps, κ . Each FLAMINGO simulation provides lightcone outputs in the form of HEALPix maps (Górski et al. 2005) as a function of redshift. The native HEALPix resolution of $N_{\text{side}} = 16,384$ is degraded to $N_{\text{side}} = 2,048$ for this project (corresponding to a pixel scale of approximately 1.4 arcminutes), with shell thickness $\Delta z = 0.05$. We produce noiseless, full-sky convergence maps assuming the DES Y3 weak lensing redshift distributions in Fig.1 from Gatti et al. 2020; Myles et al. 2021. Mimicking the DES Y3 tomographic sample, we generate 4 convergence maps by summing the shells weighted by the appropriate redshift distribution, enabling us to study baryonic effects as a function of redshift. Fig.1 shows the resulting 2D DMO power spectra in the four bins compared to theoretical predictions from the Euclid Emulator (Euclid Collaboration et al. 2019), with agreement at the 1–2% level, consistent with the emulator’s accuracy.

Since our focus is on the relative impact of baryonic feedback compared to the DMO simulation, we do not forward-model the DES Y3 noise or mask, and work exclusively with full-sky, noiseless convergence maps. In total, we use three sets of tomographic maps:

- κ_{dmo} : one set from the DMO simulation,
- κ_{flamingo} : 9 sets from the nine FLAMINGO feedback variants,
- κ_{bary} : $\sim 5,000$ sets, generated by baryonifying κ_{dmo} following the BCM prescription in Sec. 2, discussed further in Sec. 5.1.

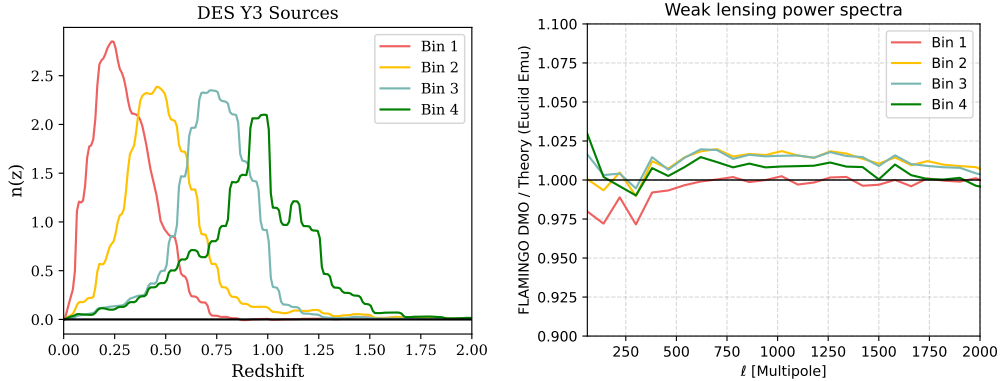


Figure 1: *Right:* redshift distribution of the four tomographic bins of the DES Y3 wak lensing sample considered here. *Left:* Power spectra of the 4 noiseless, full-sky convergence maps produced from the FLAMINGO DMO simulation, using the DES Y3 redshift distributions, compared to predictions from EuclidEmu (Euclid Collaboration et al. 2019).

4 Statistics

We now describe the theory, measurement, and scale cuts of the statistics used here: the power spectra ($C(\ell)$), the scattering transform coefficients (ST), the wavelet phase harmonic coefficients (WPH), and the third- (m_3) and the fourth-moments (m_4) of the κ maps. As an example, the statistics of κ_{flamingo} are shown in Fig. 2 in colored lines. We also summarize all the statistics in Tab. 2 at the end of the section.

4.1 Power Spectrum

The power spectrum of a κ map is defined as

$$C(\ell) = \frac{1}{2\ell + 1} \sum_m |\hat{\kappa}_{\ell m}|^2, \quad (4.1)$$

where $\hat{\kappa}_{\ell m}$ is the spherical harmonic transform of κ that is downgraded to $N_{\text{side}} = 1,024$. We consider the auto-power spectra of the 4 DES redshift bins, linearly binned in $\ell \in [20, 2020]$.

A main objective is to validate BCM’s weak lensing statistics to spatial scales smaller than the previous DMO analyses. For example, the DES Y3 weak lensing simulation-based inference with HOS was limited to $\ell_{\text{max}} = 1,000$ (Gatti et al. 2023, 2024), because the analysis assumed a DMO model and could not model smaller scales contaminated by baryonic feedback. Here, we aim to validate the BCM to $\ell_{\text{max}} = 2000$. We do not attempt to model $\ell > 2,000$ due to 1) increased systematic effects and a low signal-to-noise ratio in real data, and 2) the computational challenge of producing enough simulations with low shot noise at high ℓ in the context of simulation-based inference.




Figure 2: Summary statistics of the FLAMINGO κ maps (Sec. 3). Rows correspond to different statistics, and columns to redshift bins (modeled after DES Y3 data). The y-axis shows the FLAMINGO statistics relative to the statistics of the DMO simulation of the same initial condition. On the x-axis, spatial scales decreases to the right. ST and WPH are ordered by the wavelet size (j) as a proxy for their characteristic scale. In FLAMINGO, baryon feedback physics suppresses all statistics on all scales, and is the strongest in low redshift bins.

4.2 Scattering Transform Coefficients (ST)

Wavelets are localized oscillatory functions that act as band-pass filters in both real and Fourier space. Unlike Fourier analysis which is a global transformation, wavelets provide spatial localization while simultaneously capturing frequency information. A wavelet family consists of wavelets of different scales and orientations, and can be used to extract multi-scale patterns in images while preserving directional information (Mallat 2012).

Given a template wavelet $\psi(\mathbf{r})$ ⁵, we can generate new wavelets of different sizes and orientations via

$$\psi_{j,l}(\mathbf{r}) \propto \psi\left(\frac{1}{2^j}R_{l\pi/L}^{-1}\mathbf{r}\right), \quad j \in \{\mathbb{Z} \mid 0 \leq j < J\}, \quad l \in \{\mathbb{Z} \mid 0 \leq l < L\}. \quad (4.2)$$

Here j and l label the scale and the orientation of the derived wavelets. Each wavelet covers $\sim 2^j$ image pixels, and larger j corresponds to larger scales. The operator R^{-1} in the argument is the rotation matrix with the angle of rotation given by the subscript, with the wavelets covering L directions uniformly distributed over $[0, \pi)$.

The scattering transformed κ maps are constructed by iteratively applying the wavelet convolution and the modulus operator (Mallat 2012; Cheng et al. 2020). The first and second-order scattering transformed κ maps are

$$\kappa_{j,l}^{\text{ST},(1)} = |\kappa \star \psi_{j,l}|, \quad (4.3)$$

$$\kappa_{j,l,j',l_\Delta}^{\text{ST},(2)} = \left| |\kappa \star \psi_{j,l}| \star \psi_{j',l-l_\Delta} \right|, \quad (4.4)$$

where \star denotes convolution. The modulus preserves the first-order field statistics while introducing the necessary non-linearity to extract information beyond the mean without the instability to outliers that higher-order moments would cause. The higher-order convolution extracts information about spatial distributions $\kappa^{\text{ST},(1)}$, and captures HOS by measuring the ‘‘clustering of clustered structures,’’ with each order n potentially containing information up to 2^n -point functions (Carron 2011; Cheng et al. 2020). For the second-order field, we typically require $j \leq j'$ to ensure that the second wavelet extracts information at scales equal to or larger than the first. Since κ is a homogeneous and isotropic field, we compress κ^{ST} through spatial and angular averaging (Allys et al. 2019; Cheng et al. 2020). More specifically, the ST coefficients are given by

$$s_j^{(1)} = \langle \kappa_{j,l}^{\text{ST},(1)} \rangle_{p,l}, \quad (4.5)$$

$$s_{j,j',l_\Delta}^{(2)} = \langle \kappa_{j,l,j',l_\Delta}^{\text{ST},(2)} \rangle_{p,l}, \quad (4.6)$$

where $\langle \cdot \rangle$ denotes the expectation value, and the subscripts indicate averaging over pixels (p) and orientations (l). These coefficients compactly and robustly summarize the information in the hierarchical scales of the clustering signal beyond the power spectrum.

To measure ST, we downgrade the κ maps to $N_{\text{side}} = 1,024$, and divide them into $N_{\text{side}} = 8$ HEALPix pixels, each gnomonically projected to a 256×256 Cartesian patch

⁵For scattering transform, we use the Morlet wavelet under the convention of Cheng et al. (2020).

with $3.44'$ resolution. Each patch is centered on an $N_{\text{side}} = 8$ pixel with surrounding HEALPix pixels masked to prevent double counting. We measure ST on $1/8$ of the sky using the `scattering_transform` package⁶ with parameters $J = 4$ and $L = 3$, leading to 4 $s^{(1)}$'s and 30 $s^{(2)}$'s per redshift bin. The final ST is calculated independently for each tomographic bin and averaged across all Cartesian patches.

Since ST does not have a sharp scale cutoff like $C(\ell)$, we empirically determine the characteristic scale of each ST, and perform consistent scale cut in App. A. All 34 ST coefficients per redshift bin are found to predominantly probe scales below $\ell_{\text{max}} = 2000$.

4.3 Wavelet Phase Harmonics Coefficients (WPH)

WPH is another class of wavelet-based summary statistics. Again, we use $\psi_{j,l}(\mathbf{r})$ to denote wavelets of different sizes (j) and orientations (l), but with a different waveform – the bump-steerable wavelets (Mallat et al. 2019). WPH statistics are obtained by first convolving κ with $\psi_{j,l}(\mathbf{r})$, then taking its phase harmonics, and finally computing the moments of the resulting fields (Mallat et al. 2019; Allys et al. 2020). Since κ is isotropic, we also average over the orientations of the wavelets. The equations for the first- and second-order WPH reduce to

$$S_j^{(1,1)} = \langle \text{Cov}(\kappa \star \psi_{j,l}, \kappa \star \psi_{j,l}) \rangle_l, \quad (4.7)$$

$$S_j^{(0,0)} = \langle \text{Cov}(|\kappa \star \psi_{j,l}|, |\kappa \star \psi_{j,l}|) \rangle_l, \quad (4.8)$$

$$S_j^{(0,1)} = \langle \text{Cov}(|\kappa \star \psi_{j,l}|, \kappa \star \psi_{j,l}) \rangle_l. \quad (4.9)$$

$S^{(0,0)}$ and $S^{(0,1)}$ contain both Gaussian and non-Gaussian information due to the nonlinearity of the phase harmonics (in this case, the modulus) operation. $S^{(1,1)}$ contains only Gaussian information, albeit with a different smoothing scheme compared to the power spectrum.

The measurement of WPH is similar to that of ST. We again work with Cartesian patches that are projected from the κ maps. On each patch, we compute WPH using the `PyWPH` package⁷. We consider wavelets with $J = 4$ and $L = 3$, leading to 4 of each $S^{(1,1)}$, $S^{(0,0)}$, and $S^{(0,1)}$ per redshift bin. We then average the WPH from each patch across $1/8$ of the sky to reduce cosmic variance.

The scale cut for WPH is empirically determined in App. A and is consistent with the $\ell_{\text{max}} = 2000$ scale cut for the power spectra. The three $j = 0$ WPH are removed from subsequent analyses.

4.4 Higher Order Moments

The third and fourth moments of the convergence field are defined as

$$m_3(\theta) = \langle (\kappa \star G(\theta))^3 \rangle_p, \quad (4.10)$$

$$m_4(\theta) = \langle (\kappa \star G(\theta))^4 \rangle_p, \quad (4.11)$$

⁶github.com/SihaoCheng/scattering_transform

⁷github.com/bregaldo/pywph

where $G(\theta)$ is a Gaussian kernel of width θ and p denotes pixels. m_3 measures skewness of κ . It contains only non-Gaussian information and is strictly zero for a Gaussian random field. m_4 measures the kurtosis of κ and has both a Gaussian (derivable from the power spectrum) and a non-Gaussian contribution following Wick’s theorem. We set the minimum scale at $4'$ and use $\theta \in [4', 48']$ in increments of $4'$. This is consistent with the scale cuts of other statistics since a $4'$ Gaussian-smoothed κ has negligible power beyond $\ell = 2,000$.

The modeling and the real data measurement of m_3 and m_4 are more sensitive to noise than two-point statistics. This is because 1) by the central limit theorem, for a sample with size N , while $\langle \kappa \rangle$ has an error $\propto N^{-1/2}$, for $\langle \kappa^n \rangle$, the effective standard error scales more like $\alpha N^{-1/2}$ where α grows superlinearly with n , and 2) the tail of outliers of the distribution contributes disproportionately, amplifying the fluctuations. Perturbation theory provides a similar perspective, where empirically $\langle \kappa^n \rangle \propto \langle \kappa^2 \rangle^{n-1}$ (Bernardeau et al. 2002). Therefore an x uncertainty in $\langle \kappa^2 \rangle$ propagates to $(n-1)x$ in $\langle \kappa^n \rangle$.

4.5 The statistics of FLAMINGO and baryonified convergence maps

In total, we consider 87 (348) data points per redshift bin (and in total), as summarized in Tab. 2. The statistics for κ_{flamingo} , relative to those of κ_{dmo} , are shown in Fig. 2. In all FLAMINGO feedback variants, baryons suppress all statistics at all scales. The suppression is stronger at smaller scales and lower redshifts, consistent with previous works (e.g., Van Daalen et al. 2011; Chisari et al. 2018; Huang et al. 2019; Amon & Efstathiou 2022; Grandón et al. 2024; Schaller & Schaye 2025). Relative to the fiducial model, the fgas+2 σ variant shows weaker baryonic signatures, whereas the fgas−2 σ and Jet_fgas−4 σ variants induce stronger suppression in the statistics.

Statistics	Equation	N per bin	Notes	
ST	$s_j^{(1)}$	Eq. 4.5	4	$0 \leq j < 4$
	$s_{j,j',l_\Delta}^{(2)}$	Eq. 4.6	30	$0 \leq j \leq j' < 4, 0 \leq l_\Delta < 3$
WPH	$S^{(1,1)}$	Eq. 4.7	3	$1 \leq j < 4$
	$S^{(0,0)}$	Eq. 4.8	3	$1 \leq j < 4$
	$S^{(0,1)}$	Eq. 4.9	3	$1 \leq j < 4$
$C(\ell)$	Eq. 4.1	20	linear bin, $\ell \in [20, 2020]$	
$m_3(\theta)$	Eq. 4.10	12	smoothing scale $\theta \in [4', 48']$	
$m_4(\theta)$	Eq. 4.11	12	smoothing scale $\theta \in [4', 48']$	

Table 2: Summary of statistics used in the analysis, restricted to scales with $\ell \lesssim 2000$.

To build intuition about how each BCM parameter affects the statistics, we vary each parameter independently over its prior range defined in Tab. 1, and plot the resulting statistics of κ_{bary} relative to those of κ_{dmo} in Fig. 3. Our final results in Sec. 6 suggest

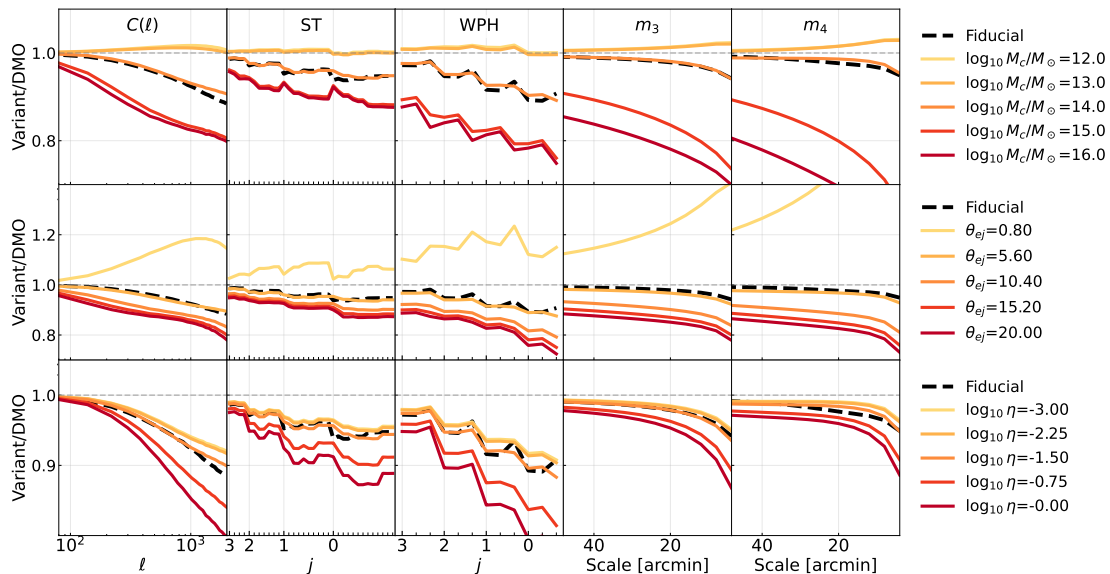


Figure 3: Summary statistics of the baryonified (colored) and the FLAMINGO fiducial variant’s (black, dashed) convergence map, focusing on the first redshift bin. For the baryonified maps, a single BCM parameter is varied at a time. Rows correspond to different BCM parameters, and columns to different statistics. The y-axis shows the statistics of both the baryonified and FLAMINGO maps, relative to those of the DMO simulation with the same initial condition. On the x-axis, spatial scale decreases to the right. ST and WPH statistics are ordered by wavelet size (j), used here as a proxy for characteristic scale. The BCM flexibly generates a range of baryonic signatures in the statistics, clustering around the FLAMINGO results. The variation of other parameters in Tab. 1 is shown in Fig. 8 in App. B.

that the three most important parameters to match and capture variations among different FLAMINGO feedback models are M_c , θ_{ej} , and η , which we highlight in Fig. 3. The effects of the remaining parameters are shown in Fig. 8 in App. B. We find that the prior ranges induce a wide range of changes in the statistics—including both suppression and, in some cases, enhancement—relative to the fiducial variant’s measurement.

5 Method

We want to demonstrate that the BCM can flexibly and accurately reproduce the wide range of baryonic feedback models in FLAMINGO, in terms of two-point and HOS and investigate what are the baryon parameter values that best fit the FLAMINGO simulations and how they vary as a function of hydrodynamical feedback strengths.

In principle, we could sample the best-fit baryon parameters in one step using a Markov Chain Monte Carlo (MCMC) – given the baryon parameters, we could generate baryonified κ maps using the DMO simulation, measure their statistics, and compare them

to those from FLAMINGO. In practice, generating new baryonified maps and computing their statistics at each MCMC step is too time-consuming. Therefore, we proceed in three steps:

1. Build an emulator that predicts the statistics for the given baryon parameters
2. Use the emulator in an MCMC on each of the FLAMINGO simulations to find the best-fit baryon parameters
3. Once the MCMC chains have converged, we generate baryonified κ maps at the MCMC best fit, measure the suppression of the statistics, and use them for our final considerations.

Therefore, we use the emulator as a fast and efficient tool to refine our search in parameter space, rather than brute-forcing the problem by baryonifying the maps within the MCMC, which would be computationally infeasible. This means that possible discrepancies between the suppression of the statistics in the FLAMINGO simulations and the one we obtain at the end of our three-step procedure can be attributed to two factors: (1) the accuracy of the emulator in reproducing the statistics compared to the full baryonification process, and (2) the intrinsic limitations of the baryonification method itself, even assuming a perfect emulator.

5.1 Emulator

We will refer to the collection of all summary statistics as \mathcal{S} . To accelerate the MCMC inference, we construct an emulator r , which predicts the ratio of the BCM-corrected summary statistics to its DMO counterpart given baryon parameters \mathbf{b} , i.e.,

$$r(\mathbf{b}) = \mathcal{S}_{\text{bary}}(\mathbf{b})/\mathcal{S}_{\text{dmo}}. \quad (5.1)$$

The numerator $\mathcal{S}_{\text{bary}}(\mathbf{b})$ is computed by first baryonifying κ_{dmo} using the BCM (Sec. 2) and then performing two-point and HOS measurements (Sec. 4). Since κ_{bary} and κ_{dmo} share the same initial condition, the ratio eliminates shot noise and cosmic variance.

The training set includes 5,468 κ_{bary} maps sampled in the ten-dimensional baryon parameter space. The prior ranges for the parameters are detailed in Tab. 1 and are similar to those used in [Schneider et al. \(2019\)](#); [Anbajagane et al. \(2024a\)](#). In Figs. 3 and 8, we have shown how varying a single BCM parameter can produce suppression similar to those observed in the fiducial model. We also explicitly verify that, across the training samples, the baryonified maps exhibit baryonic signatures spanning the full range of variations seen in the different FLAMINGO feedback variants. For each statistic, we construct an ensemble of 10 neural networks via cross-validation. We use the mean and 1σ scatter of the ensemble predictions as the point estimate and uncertainty of $r(\mathbf{b})$. The emulator accuracy is shown in Fig. 4 as a function of scale and redshift, and generally decreases at smaller scales and lower redshifts, where baryonic feedback is more nonlinear and varied. The emulator achieves a mean absolute error below 1% for all statistics except the low-redshift moments, which are intrinsically noisier, as discussed in Sec. 4.4. The

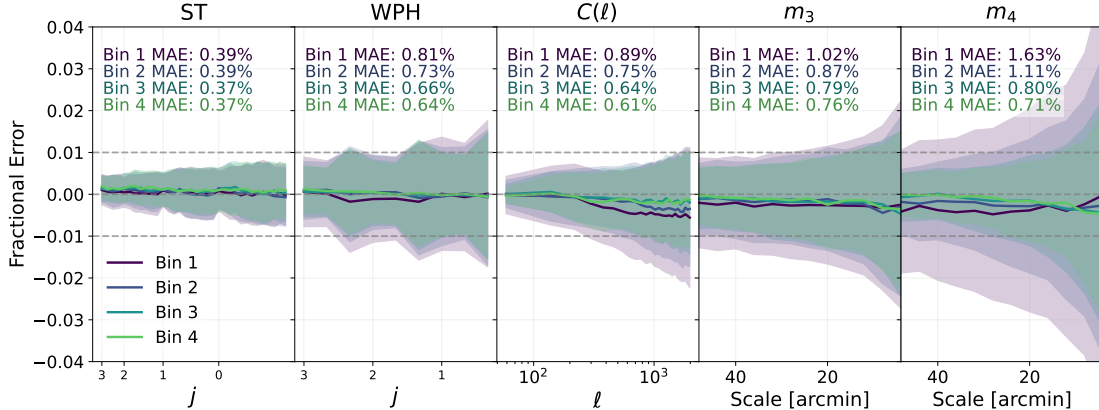


Figure 4: The emulator mean relative error, grouped by statistics and by redshift bins. The mean absolute (relative) error (MAE) is shown in text. The error is evaluated by averaging across the test set across the baryon parameter space. The emulator achieves sub-percent accuracy for all statistics besides the low redshift moments. The higher-order moment statistics are intrinsically noisier since it involves multiplications of random variables. We remind the reader that the emulator is used to efficiently explore the BCM parameter space when comparing to FLAMINGO statistics. However, the final assessment of how well the BCM reproduces a given FLAMINGO statistic is based on running the full baryonification pipeline at the emulator-derived best-fit parameters.

emulator is implemented in JAX and supports automatic differentiation (Bradbury et al. 2018). Details of the training set and emulator are provided in App. C.

5.2 Sampling

For each FLAMINGO feedback variant with statistics $\mathcal{S}_{\text{hydro}}$, we seek the best-fit baryon parameters by sampling the posterior \mathcal{P}

$$\mathcal{P}(\mathbf{b} \mid \mathcal{S}_{\text{hydro}}) = \mathcal{G} \left(\frac{\mathcal{S}_{\text{hydro}}}{\mathcal{S}_{\text{dmo}}} - r(\mathbf{b}), \sigma^2 + \epsilon(\mathbf{b})^2 \right) \quad (5.2)$$

where \mathcal{G} is a Gaussian distribution parametrized by mean and variance, σ is a constant uncertainty per statistic and ϵ is the emulator uncertainty. We assume a flat prior on \mathbf{b} . As discussed above, the simulations are noiseless and the ratio statistics suppress the cosmic variance and shot noise. Therefore, in theory, the emulator uncertainty ϵ is the only source of error in $r(\mathbf{b})$. However, since ϵ is a function of \mathbf{b} , the parameter-dependent covariance makes the sampling very difficult in high-dimensional space. Therefore, we introduce a constant noise σ to promote the smoothness of the posterior surface to facilitate sampling. In particular, σ does not reflect real observational uncertainty. We find $\sigma = 0.04$ for ST, WPH, and $C(\ell)$ and $\sigma = 0.08$ for m_3 and m_4 make the chains converge robustly. Although we are putting less weight on the moments, we shall see that the final best-fit baryon parameters reproduce all statistics equally well. We leverage

the gradient information provided by the emulator and use the `numpyro` implementation (Bingham et al. 2018; Phan et al. 2019) of the Hamiltonian Monte Carlo algorithm for efficient sampling (Neal 2011; Betancourt 2018).

6 Results

6.1 The 10-parameter baryon correction model

We sample the posterior $\mathcal{P}(\mathbf{b} \mid \mathcal{S}_{\text{hydro}})$ across the full 10-dimensional baryon parameter space for each of the FLAMINGO feedback variants. Since the marginal posterior distributions are approximately Gaussian, we use the medians – rather than the maximum a posteriori (MAP) values, which are less robust against sample variance – as our best-fit parameters $\hat{\mathbf{b}}$. The estimated $\hat{\mathbf{b}}$ values and their 1σ uncertainties are listed in Tab. 4 in App. D.

We then investigate how well the best-fit BCM reproduces the statistics in κ_{flamingo} . Ideally, we would directly compare the emulator prediction $r(\hat{\mathbf{b}})$ with FLAMINGO. However, the emulator includes bias from both the neural network and the BCM itself. To isolate the latter, we directly construct $\kappa_{\text{bary}}(\hat{\mathbf{b}})$, measure its statistics, and compare them to the FLAMINGO results. The emulator results agree with direct baryonification within expected emulator errors. For example, for FLAMINGO’s fiducial variant, the relative error of the 10-parameter best-fit BCM is shown in Fig. 5 in dark red. The BCM accurately reproduces ST, WPH, $C(\ell)$, and moment statistics at all scales and redshifts to within 2%. The residual error is the largest at small scales and low redshifts, where baryonic feedback is strongest.

We analyze all 8 other FLAMINGO variants in the same way and show the residuals for the strongest ($\text{fgas}-8\sigma$) and weakest ($\text{fgas}+2\sigma$) FLAMINGO feedback variants in Fig. 9 and Fig. 10 in App. E. The BCM can reproduce the statistics of different FLAMINGO variants equally well. For the strongest AGN feedback case ($\text{fgas}-8\sigma$), the BCM shows elevated errors of 3–4% in the first redshift bin of m_4 for $\theta < 20$ arcminutes, but otherwise maintains 2% accuracy. However, as discussed in Sec. 4.4, higher-order moments are intrinsically noisier in both modeling and measurement. The observed m_4 discrepancy remains well below the statistical uncertainty of a DES-like survey (Gatti et al. 2024) and is only detected at low significance in current DES Y3 data (Anbajagane et al. 2023).

We have verified that a significant fraction of the remaining 2% residual at small scales can be alleviated by lowering the halo mass threshold to $10^{12.5} M_{\odot}/h$. This test was performed using the best-fit of the baryonification emulator on the Fiducial simulation, and should be considered an approximate check. A formal evaluation would require regenerating the maps, retraining the emulator, and refitting the model using $10^{12.5} M_{\odot}/h$ as the threshold — a process we do not pursue due to its substantial computational cost. Nonetheless, we note that adopting a threshold of $10^{13} M_{\odot}/h$ still yields sub-2% precision.

Thus far, we have investigated the BCM at the fiducial cosmology. While both $\mathcal{S}_{\text{hydro}}$ and \mathcal{S}_{dmo} depend on cosmology and astrophysics, these dependencies are approximately




Figure 5: The residuals error of the BCM when fitted to the FLAMINGO fiducial variant. Rows correspond to different statistics, and columns to redshift bins. The y-axis shows the ratio between the statistics of the best-fit baryonified map to that of the fiducial variant. On the x-axis, spatial scales decreases to the right. ST and WPH are ordered by the wavelet size (j) as a proxy for their characteristic scale. Colors indicate the number of free parameters in the BCM fit. The shaded region marks $\pm 2\%$ for visual reference. For all n -parameter models, the BCM flexibly reproduces the baryonic features in the FLAMINGO fiducial variant simulation to percent level. The prediction is robust even with only 3 degrees of freedoms. Differences in accuracy between the models stem from how we estimate the best-fit parameter, as discussed in Sec. 6.2. The residuals for the strongest ($f_{\text{gas}} - 8\sigma$) and weakest ($f_{\text{gas}} + 2\sigma$) feedback variants are shown in Fig. 9 and Fig. 10 in App. E.

factorizable (Van Daalen et al. 2011; Elbers et al. 2025). As a result, the cosmology dependence largely cancels in the ratio $\mathcal{S}_{\text{hydro}}/\mathcal{S}_{\text{dmo}}$, which the BCM models. Using the low- S_8 variants (LS8, LS8_fgas- 8σ) of the FLAMINGO suite, we show in App. F that this ratio varies by less than 1% compared to the fiducial cosmology even for the strong feedback model (fgas- 8σ). Since the BCM (and thus the emulator) predicts the suppression ratio solely from the baryonic parameters, the recovered LS8 BCM parameters are consistent with those from the fiducial cosmology. The baryonic suppression has a stronger dependence on the baryon fraction, Ω_b/Ω_m (e.g. Schneider et al. 2019) but we leave a detailed analysis of this dependence to future work due to the computational cost of training separate emulators for each cosmology.

To summarize, for $\ell \leq 2000$, the 10-parameter BCM is flexible enough to reproduce the range of baryonic signatures in two-point and HOS.

6.2 Baryon correction model with restricted degrees of freedoms

Variant	$\log_{10} M_c/M_\odot$	$\log_{10} \eta$	θ_{ej}
Fiducial(L1_m9)	13.9 ± 0.19	-1.51 ± 0.39	4.80 ± 0.65
Jet(3)	13.7 ± 0.16	-1.87 ± 0.39	$6.23^{+0.89}_{-0.75}$
Jet_fgas- 4σ (3)	14.0 ± 0.16	$-1.02^{+0.21}_{-0.25}$	5.94 ± 0.61
fgas+ 2σ	13.7 ± 0.16	-1.95 ± 0.38	$5.05^{+0.68}_{-0.59}$
fgas- 2σ	14.0 ± 0.22	$-1.20^{+0.34}_{-0.35}$	4.85 ± 0.59
fgas- 4σ	14.0 ± 0.19	$-0.983^{+0.23}_{-0.29}$	5.05 ± 0.54
fgas- 8σ	14.0 ± 0.14	$-0.885^{+0.17}_{-0.18}$	5.40 ± 0.52
M*- σ	14.0 ± 0.21	$-1.17^{+0.33}_{-0.37}$	4.69 ± 0.55
M*- σ _fgas- 4σ	14.0 ± 0.15	-0.852 ± 0.2	5.05 ± 0.5

Table 3: The best-fit BCM parameters obtained from the joint fitting of all HOS to different FLAMINGO variants. These parameter combinations yield HOS that agree with FLAMINGO measurements to 2%. Other BCM parameters are fixed to $\nu_{M_c} = 0$, $\nu_{\theta_{\text{ej}}} = 0$, $\gamma = 2.60$, $\log_{10} \theta_{c\sigma} = -1.00$, $\log_{10} \eta_\delta = -0.28$, $\delta = 14.0$, and $\mu_\beta = 1.60$. The uncertainty is the 16% and 84% credible interval of our fitting method. See Tab. 4 in App. D for the best-fit values for all n-parameter models.

After validating that the 10-parameter BCM accurately reproduces the statistics, we ask whether similar success can be achieved with significantly fewer degrees of freedom. This simplification has three advantages. First, in real survey analyses, using fewer baryon parameters reduce the impact on the constraining power of key cosmological parameters. Second, in simulation-based inference, it lowers the number of simulations required to train the neural posterior estimator, thus reducing computational cost. Third, the 10-dimensional \mathbf{b} space contains significant degeneracies; fewer parameters improves the interpretability how \mathbf{b} vary across FLAMINGO feedback variants.

We consider 4 restricted versions of the BCM, with 8, 6, 4, and 3 parameters. In

general, at each step, we fix parameters whose best-fit values are consistent across all feedback variants. The 8-parameter model fixes the redshift scalings ν_{M_c} and $\nu_{\theta_{ej}}$ to zero, consistent with their posteriors in the 10-parameter fits. The 6-parameter model further fixes $\gamma = 2.6$ and $\log_{10} \theta_{co} = -1$. For the 4-parameter model, we also fix $\log_{\delta} \eta_{\delta} = 0.28$ and $\delta = 14$. Finally, for the 3-parameter model, we additionally fix $\mu_{\beta} = 1.6$. The remaining free parameters in the 3-parameter model are $\log_{10} M_c$, $\log_{10} \eta$, and θ_{ej} . For each n-parameter model and feedback variant, we tabulate $\hat{\mathbf{b}}$ in Tab. 4 in App. D, with fixed parameters shown in red. We also summarize $\hat{\mathbf{b}}$ values for the 3-parameter model in Tab. 3. As with the 10-parameter case, we validate the best-fit BCMs by directly constructing κ_{bary} and measuring their statistics. Residuals for the fiducial, strong feedback (fgas -8σ), and weak feedback (fgas $+2\sigma$) variants are shown in Figs. 5, 9, and 10, respectively.

Although we have dramatically reduced the number of parameters, all n-parameter models faithfully reproduce all statistics across scales, redshifts, and feedback variants. Notably, the 3-parameter model often fits the statistics better than the 10-parameter model. This is because our best estimate, $\hat{\mathbf{b}}$, is defined as the median of the marginal posterior. In high-dimensional spaces, $\hat{\mathbf{b}}$ can deviate from the MAP due to projection over non-Gaussian volumes. As the number of parameters decreases, the posterior becomes more Gaussian-like, and $\hat{\mathbf{b}}$ approaches the MAP, leading to improved fits. In other words, were we to have a less noisy emulator, we could potentially use the MAP as $\hat{\mathbf{b}}$ and obtain even higher-quality fits for the 10-parameter model.

6.3 Baryon parameters

Having demonstrated that the BCM can flexibly reproduce the baryonic signatures across a range of FLAMINGO feedback variants, we now examine the best-fit baryon parameters themselves. When comparing with other results in the literature, we emphasize that our modeling choices — such as the specific implementation of the BCM and the set of free parameters — are in some aspects different from those adopted in other studies, and that our results are entirely based on simulations. In this context, consistency at the parameter level with earlier work is not required for our approach to be effective, as success is defined by how well it reproduces the suppression of the lensing statistics in the FLAMINGO simulations, rather than by agreement in individual parameter values. Nonetheless, the inferred $\hat{\mathbf{b}}$ is broadly consistent with earlier results based on both simulations and observations (Schneider et al. 2019; Giri & Schneider 2021; Bigwood et al. 2024b). This qualitative agreement reinforces that our model is not overfitting to a narrow region of parameter space and is instead capturing physically meaningful baryonic effects.

Focusing on the 10-parameter model, all variants are consistent with $\nu_{M_c} = 0$ and $\nu_{\theta_{ej}} = 0$, suggesting minimal sensitivity of our measurements to the mass dependence of the gas ejection radius and characteristic slope. This lack of sensitivity is consistent with the forecasted constraints from weak lensing presented in Anbajagane et al. (2024a). Interestingly, all feedback variants prefer a high δ , which corresponds to a sharp drop off in the gas density profile past $r > \theta_{ej} R_{200c}$ (Anbajagane et al. 2024a; Schneider et al. 2019;




Figure 6: The inferred baryon parameter $\hat{\mathbf{b}}$ from different FLAMINGO feedback variants. Markers represent inference using different degrees of freedom. The color map labels the feedback variants and inherits from Fig. 2. In the legend, the feedback strength approximately increases from top to bottom, with the strongest (weakest) feedback model suppression being the $fgas-8\sigma$ ($fgas+2\sigma$) model. For each feedback variant, the clustering of the points demonstrate the BCM is robust against removing many nuisance parameter dependencies.

Bigwood et al. 2024b). To investigate if different values of δ affect the best-fit statistics, we rerun the 4- and 3-parameter models, setting $\delta = 7$ and $\mu_\beta = 2$ while keeping the other fixed parameters the same as in Sec. 6.2. The BCM still achieves a 2% fit across all statistics and feedback variants, although the best-fit θ_{ej} lowers from 5 to 3. This is expected as δ and θ_{ej} both control the radial extent of the gas profile and are therefore quite negatively correlated.

Next, we investigate how much the estimated $\hat{\mathbf{b}}$ varies across the n-parameter models. Fig. 6 shows $\hat{\mathbf{b}}$ as a function of feedback variant for the 6-, 4-, and 3-parameter models. The constraints on the remaining free parameters do not vary significantly as we reduce the number of parameters. This demonstrates that the BCM’s weak lensing predictions are robust against removing many nuisance parameters and the BCM allows us to track how the phenomenological baryon parameters change as a function of FLAMINGO feedback models. The best-fit parameters of the different models show some degree of clustering

in the 3D parameter space, indicating that, if sufficient statistical power is achieved by observational probes, it may be possible to distinguish between feedback models using real data. Concerning the physical interpretation of the position of the models in the parameter space, we know that $\log_{10} M_c$, $\log_{10} \eta$, and θ_{ej} all scale positively with feedback strength. Indeed, we find that the strongest feedback model ($f_{\text{gas}}-8\sigma$) lies in the upper range of each of the three parameters, while the weakest feedback model ($f_{\text{gas}}+2\sigma$ weak) occupies the lowest values in each dimension. Beyond a qualitative ranking of suppression strength, the interplay between the three parameters and the resulting impact on the lensing statistics is non-trivial. Each parameter modulates the suppression differently and in a scale-dependent way, leading to best-fit values that are scattered across parameter space in a complex pattern.

7 Conclusion

Data from current and upcoming weak lensing surveys are probing cosmological structures at increasingly small spatial scales. To extract both cosmological and astrophysical information in this regime robustly and optimally, we require models that can accurately capture the full range of baryonic feedback effects allowed by theory and observed in state-of-the-art hydrodynamical simulations.

In this work, we investigate the map-level baryonification approach from [Anbajagane et al. \(2024a\)](#), which efficiently and flexibly models baryonic effects in weak lensing convergence maps. By calibrating the baryon parameters against FLAMINGO hydrodynamical simulations, we demonstrate that this baryon correction model (BCM) accurately reproduces baryonic impacts on various convergence field statistics, including the two-point power spectrum, scattering transform coefficients (ST), wavelet phase harmonics (WPH), and third- and fourth-order moments. This agreement holds across scales $\ell < 2000$ and throughout the redshift range of DES-like surveys. We construct an emulator to predict BCM-generated statistics efficiently, then sample baryon parameters by comparing emulator predictions with measurements from different FLAMINGO feedback variants. Testing BCM versions with 10, 8, 6, 4, and 3 free parameters, we find all models can reproduce statistics within $< 2\%$ accuracy with accuracy improving at larger scales and higher redshifts. A minor exception is the fourth moment, where, for models with more than 3 free parameters, the residual error approaches 3% around $10'$, partially due to the increased uncertainty in our fitting procedure.

These results have several implications for future weak lensing analyses. First, this shows that the BCM is robust for a wide range of weak lensing statistics and can be embedded in simulation-based inference frameworks to marginalize over baryonic uncertainties or directly constrain baryon parameters. Reduced-parameter versions of the BCM further improve efficiency by lowering the number of simulation needed and sampling costs. Second, our work can be extended to test the robustness of other tracers and probes, such as cross-correlations with large-scale structure observables like galaxy clustering or thermal Sunyaev-Zel'dovich maps. Joint analyses with these additional

datasets may offer tighter constraints on baryonic physics.

Although our results are based on the DES Y3 redshift distributions, they are broadly applicable to other current weak lensing surveys due to similar source redshift distributions (Rau et al. 2023; Anbajagane et al. 2025a; Wright et al. 2025a). They are also relevant for upcoming surveys like LSST, Euclid, and Roman, whose higher-redshift distributions may allow our models to extend to even smaller angular scales due to weaker baryonic feedback at early times. The ability to model baryonic physics in this regime opens new opportunities for extracting both cosmological and astrophysical information from small-scale data. Future work could extend this framework to alternative summary statistics, such as voids, peaks, or CNN-based approaches (e.g., Jain & Waerbeke 2000; Jeffrey et al. 2024). Our results highlight the promise of map-level baryonification—without relying on 3D particle snapshots—as a computationally efficient and physically grounded approach for incorporating baryonic feedback in weak lensing analyses. As upcoming surveys push to smaller scales, such methods will be essential for maximizing constraining power while remaining robust to astrophysical systematics.

Acknowledgments

AZ is supported by Jane Street through the Jane Street Graduate Research Fellowship. MG is supported by funds provided by the Kavli Institute for Cosmological Physics at the University of Chicago through an endowment from the Kavli Foundation. DA is supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE 1746045. This work was supported by FermiForward Discovery Group, LLC under Contract No. 89243024CSC000002 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.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.

Data Availability

The BARYONFORGE pipeline used in this work is publicly available at <https://github.com/DhayaaAnbajagane/BaryonForge>. The FLAMINGO simulations are available upon request by contacting Joop Schaye at schaye@strw.leidenuniv.nl. The total matter power spectra from the FLAMINGO simulations have been publicly released and are available at <https://flamingo.strw.leidenuniv.nl>.

References

- Abbott T. M. C., et al., 2022, *Phys. Rev. D*, **105**, 023520
- Allys E., Levrier F., Zhang S., Colling C., Blancard B. R.-S., Boulanger F., Hennebelle P., Mallat S., 2019, *Astronomy & Astrophysics*, **629**, A115
- Allys E., Marchand T., Cardoso J.-F., Villaescusa-Navarro F., Ho S., Mallat S., 2020, *Physical Review D*, **102**, 103506
- Alsing J., Heavens A. F., Jaffe A. H., 2017, *Monthly Notices of the Royal Astronomical Society*, **466**, 3272
- Amon A., Efstathiou G., 2022, *Monthly Notices of the Royal Astronomical Society*, **516**, 5355
- Amon A., et al., 2022a, *Phys. Rev. D*, **105**, 023514
- Amon A., et al., 2022b, *Monthly Notices of the Royal Astronomical Society*, **518**, 477
- Anbajagane D., Evrard A. E., Farahi A., Barnes D. J., Dolag K., McCarthy I. G., Nelson D., Pillepich A., 2020, *MNRAS*, **495**, 686
- Anbajagane D., Evrard A. E., Farahi A., 2022a, *MNRAS*, **509**, 3441
- Anbajagane D., et al., 2022b, *MNRAS*, **510**, 2980
- Anbajagane D., et al., 2023, *MNRAS*, **526**, 5530
- Anbajagane D., Pandey S., Chang C., 2024a ([arXiv:2409.03822](https://arxiv.org/abs/2409.03822))
- Anbajagane D., Chang C., Lee H., Gatti M., 2024b, *J. Cosmology Astropart. Phys.*, **2024**, 062
- Anbajagane D., et al., 2025a, *arXiv e-prints*, p. [arXiv:2502.17675](https://arxiv.org/abs/2502.17675)
- Anbajagane D., et al., 2025b, *arXiv e-prints*, p. [arXiv:2502.17677](https://arxiv.org/abs/2502.17677)
- Aricò G., Angulo R. E., Hernández-Monteagudo C., Contreras S., Zennaro M., Pellejero-Ibañez M., Rosas-Guevara Y., 2020, *Monthly Notices of the Royal Astronomical Society*, **495**, 4800
- Aricò G., Angulo R. E., Hernández-Monteagudo C., Contreras S., Zennaro M., 2021, *MNRAS*, **503**, 3596
- Aricò G., Angulo R. E., Zennaro M., Contreras S., Chen A., Hernández-Monteagudo C., 2023, *A&A*, **678**, A109
- Bernardeau F., Colombi S., Gaztanaga E., Scoccimarro R., 2002, *Physics Reports*, **367**, 1
- Betancourt M., 2018 ([arXiv:1701.02434](https://arxiv.org/abs/1701.02434))
- Bigwood L., et al., 2024a, *arXiv e-prints*, p. [arXiv:2404.06098](https://arxiv.org/abs/2404.06098)
- Bigwood L., et al., 2024b ([arXiv:2404.06098](https://arxiv.org/abs/2404.06098)), [doi:10.48550/arXiv.2404.06098](https://doi.org/10.48550/arXiv.2404.06098)
- Bingham E., et al., 2018 ([arXiv:1810.09538](https://arxiv.org/abs/1810.09538))
- Booth C. M., Schaye J., 2009, [doi:10.1063/1.3293041](https://doi.org/10.1063/1.3293041)
- Borrow J., Schaller M., Bower R. G., Schaye J., 2022, *MNRAS*, **511**, 2367
- Boruah S. S., Rozo E., 2023 ([arXiv:2307.00070](https://arxiv.org/abs/2307.00070))
- Bradbury J., et al., 2018
- Broxterman J. C., et al., 2024, *Monthly Notices of the Royal Astronomical Society*, **529**, 2309
- Carron J., 2011, *The Astrophysical Journal*, **738**, 86

- Chaikin E., Schaye J., Schaller M., Benítez-Llambay A., Nobels F. S. J., Ploeckinger S., 2023, [MNRAS](#), **523**, 3709
- Chen A., et al., 2023, [MNRAS](#), **518**, 5340
- Cheng S., Ting Y.-S., Ménard B., Bruna J., 2020, [Monthly Notices of the Royal Astronomical Society](#), **499**, 5902
- Cheng S., Marques G. A., Grandón D., Thiele L., Shirasaki M., Ménard B., Liu J., 2025, [J. Cosmology Astropart. Phys.](#), **2025**, 006
- Chisari N. E., et al., 2018, [MNRAS](#), **480**, 3962
- Dai B., Seljak U., 2024 ([arXiv:2306.04689](#))
- Dalal R., et al., 2023, [Phys. Rev. D](#), **108**, 123519
- Dalla Vecchia C., Schaye J., 2008, [MNRAS](#), **387**, 1431
- Dark Energy Survey and Kilo-Degree Survey Collaboration et al., 2023, [The Open Journal of Astrophysics](#), **6**, 36
- Davies C. T., Harnois-Déraps J., Li B., Giblin B., Hernández-Aguayo C., Paillas E., 2024, [MNRAS](#), **533**, 3546
- Debackere S. N. B., Schaye J., Hoekstra H., 2020, [Monthly Notices of the Royal Astronomical Society](#), **492**, 2285
- Diemer B., Kravtsov A. V., 2015, [ApJ](#), **799**, 108
- Duffy A. R., Schaye J., Kay S. T., Dalla Vecchia C., Battye R. A., Booth C. M., 2010, [MNRAS](#), **405**, 2161
- Elbers W., Frenk C. S., Jenkins A., Li B., Pascoli S., 2021, [MNRAS](#), **507**, 2614
- Elbers W., et al., 2025, [Monthly Notices of the Royal Astronomical Society](#), **537**, 2160
- Euclid Collaboration et al., 2019, [MNRAS](#), **484**, 5509
- Fluri J., Kacprzak T., Lucchi A., Refregier A., Amara A., Hofmann T., Schneider A., 2019, [Phys. Rev. D](#), **100**, 063514
- Fluri J., Kacprzak T., Lucchi A., Schneider A., Refregier A., Hofmann T., 2022, [Phys. Rev. D](#), **105**, 083518
- Gatti M., Giannini G., et al., 2020, arXiv e-prints, p. [arXiv:2012.08569](#)
- Gatti M., et al., 2022, [Phys. Rev. D](#), **106**, 083509
- Gatti M., et al., 2023 ([arXiv:2310.17557](#))
- Gatti M., et al., 2024 ([arXiv:2405.10881](#)), [doi:10.48550/arXiv.2405.10881](#)
- Gebhardt M., et al., 2023, arXiv e-prints, p. [arXiv:2307.11832](#)
- Giri S. K., Schneider A., 2021, [Journal of Cosmology and Astroparticle Physics](#), **2021**, 046
- Gnedin O. Y., Kravtsov A. V., Klypin A. A., Nagai D., 2004, [ApJ](#), **616**, 16
- Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, [ApJ](#), **622**, 759
- Grandón D., Marques G. A., Thiele L., Cheng S., Shirasaki M., Liu J., 2024, [Physical Review D](#), **110**, 103539
- Huang H.-J., Eifler T., Mandelbaum R., Dodelson S., 2019, [Monthly Notices of the Royal Astronomical Society](#), **488**, 1652

Huško F., Lacey C. G., Schaye J., Schaller M., Nobels F. S. J., 2022, *MNRAS*, 516, 3750

Jain B., Waerbeke L. V., 2000, *The Astrophysical Journal*, 530, L1

Jeffrey N., et al., 2024, *Monthly Notices of the Royal Astronomical Society*, 536, 1303

Kingma D. P., Ba J., 2017 ([arXiv:1412.6980](https://arxiv.org/abs/1412.6980)), doi:10.48550/arXiv.1412.6980

Kugel R., et al., 2023, *MNRAS*, 526, 6103

Lee E., et al., 2022, *MNRAS*, 517, 5303

Lee M. E., Lu T., Haiman Z., Liu J., Osato K., 2023, *MNRAS*, 519, 573

Li X., et al., 2023, *Phys. Rev. D*, 108, 123518

Lim S. H., Barnes D., Vogelsberger M., Mo H. J., Nelson D., Pillepich A., Dolag K., Marinacci F., 2021, *MNRAS*, 504, 5131

Liu X., et al., 2016, *Phys. Rev. Lett.*, 117, 051101

Loureiro A., Whiteway L., Sellentin E., Lafaurie J. S., Jaffe A. H., Heavens A. F., 2022 ([arXiv:2210.13260](https://arxiv.org/abs/2210.13260))

Mallat S., 2012, *Communications on Pure and Applied Mathematics*, 65, 1331

Mallat S., Zhang S., Rochette G., 2019 ([arXiv:1810.12136](https://arxiv.org/abs/1810.12136)), doi:10.48550/arXiv.1810.12136

Marques G. A., et al., 2024, *MNRAS*, 528, 4513

McCarthy I. G., Schaye J., Bird S., Le Brun A. M. C., 2017, *Monthly Notices of the Royal Astronomical Society*, 465, 2936

McCarthy I. G., et al., 2023, *Monthly Notices of the Royal Astronomical Society*, 526, 5494

Mead A. J., Tröster T., Heymans C., Van Waerbeke L., McCarthy I. G., 2020, *A&A*, 641, A130

Myles J., et al., 2021, *MNRAS*, 505, 4249

Navarro J. F., Frenk C. S., White S. D. M., 1996, *The Astrophysical Journal*, 462, 563

Navarro J. F., Frenk C. S., White S. D. M., 1997, *The Astrophysical Journal*, 490, 493

Neal R. M., 2011, MCMC Using Hamiltonian Dynamics. ([arXiv:1206.1901](https://arxiv.org/abs/1206.1901)), doi:10.1201/b10905

Novaes C. P., et al., 2024, *arXiv e-prints*, p. [arXiv:2409.01301](https://arxiv.org/abs/2409.01301)

Pandey S., Salcido J., To C.-H., Hill J. C., Anbajagane D., Baxter E. J., McCarthy I. G., 2024, *arXiv e-prints*, p. [arXiv:2401.18072](https://arxiv.org/abs/2401.18072)

Phan D., Pradhan N., Jankowiak M., 2019 ([arXiv:1912.11554](https://arxiv.org/abs/1912.11554))

Piccirilli G., Zennaro M., García-García C., Alonso D., 2025, *arXiv e-prints*, p. [arXiv:2502.17339](https://arxiv.org/abs/2502.17339)

Ploekinger S., Schaye J., 2020, *MNRAS*, 497, 4857

Rau M. M., et al., 2023, *MNRAS*, 524, 5109

Schaller M., Schaye J., 2025 ([arXiv:2504.15633](https://arxiv.org/abs/2504.15633)), doi:10.48550/arXiv.2504.15633

Schaller M., et al., 2024, *MNRAS*, 530, 2378

Schaller M., Schaye J., Kugel R., Broxterman J. C., van Daalen M. P., 2025, *Monthly Notices of the Royal Astronomical Society*, 539, 1337

Schaye J., Dalla Vecchia C., 2008, *MNRAS*, 383, 1210

- Schaye J., et al., 2023, *MNRAS*, 526, 4978
- Schneider A., Teyssier R., 2015, *Journal of Cosmology and Astroparticle Physics*, 2015, 049
- Schneider A., Teyssier R., Stadel J., Chisari N. E., Brun A. M. C. L., Amara A., Refregier A., 2019, *Journal of Cosmology and Astroparticle Physics*, 2019, 020
- Secco L. F., et al., 2022, *Phys. Rev. D*, 105, 023515
- Semboloni E., Hoekstra H., Schaye J., Van Daalen M. P., McCarthy I. G., 2011, *Monthly Notices of the Royal Astronomical Society*, 417, 2020
- Semboloni E., Hoekstra H., Schaye J., 2013, *Monthly Notices of the Royal Astronomical Society*, 434, 148
- Shao M., Anbajagane D., Chang C., 2022, *arXiv e-prints*, p. [arXiv:2212.05964](https://arxiv.org/abs/2212.05964)
- Srivastava N., Hinton G., Krizhevsky A., Sutskever I., Salakhutdinov R., 2014, *Journal of Machine Learning Research*, 15, 1929
- Stiskalek R., Bartlett D. J., Desmond H., Anbajagane D., 2022, *MNRAS*, 514, 4026
- To C.-H., Pandey S., Krause E., Dalal N., Anbajagane D., Weinberg D. H., 2024, *arXiv e-prints*, p. [arXiv:2402.00110](https://arxiv.org/abs/2402.00110)
- Truttero O., Zuntz J., Pourtsidou A., Robertson N., 2025, *The Open Journal of Astrophysics*, 8, 19
- Van Daalen M. P., Schaye J., 2015, *Monthly Notices of the Royal Astronomical Society*, 452, 2247
- Van Daalen M. P., Schaye J., Booth C. M., Dalla Vecchia C., 2011, *Monthly Notices of the Royal Astronomical Society*, 415, 3649
- Velliscig M., van Daalen M. P., Schaye J., McCarthy I. G., Cacciato M., Brun A. M. C. L., Vecchia C. D., 2014, *Monthly Notices of the Royal Astronomical Society*, 442, 2641
- Vogelsberger M., Marinacci F., Torrey P., Puchwein E., 2020, *Nature Reviews Physics*, 2, 42
- Wiersma R. P. C., Schaye J., Theuns T., Dalla Vecchia C., Tornatore L., 2009, *MNRAS*, 399, 574
- Wright A. H., et al., 2025a, *arXiv e-prints*, p. [arXiv:2503.19440](https://arxiv.org/abs/2503.19440)
- Wright A. H., et al., 2025b, *arXiv e-prints*, p. [arXiv:2503.19441](https://arxiv.org/abs/2503.19441)
- Xu B., Wang N., Chen T., Li M., 2015 ([arXiv:1505.00853](https://arxiv.org/abs/1505.00853)), [doi:10.48550/arXiv.1505.00853](https://doi.org/10.48550/arXiv.1505.00853)
- Yoon M., Jee M. J., Tyson J. A., Schmidt S., Wittman D., Choi A., 2019, *ApJ*, 870, 111
- Zhou A. J., Li X., Dodelson S., Mandelbaum R., 2024a, *Physical Review D*, 110, 023539
- Zhou A. J., Li Y., Dodelson S., Mandelbaum R., Zhang Y., Li X., Fabbian G., 2024b, *Journal of Cosmology and Astroparticle Physics*, 2024, 069
- van Daalen M. P., McCarthy I. G., Schaye J., 2020, *Monthly Notices of the Royal Astronomical Society*, 491, 2424

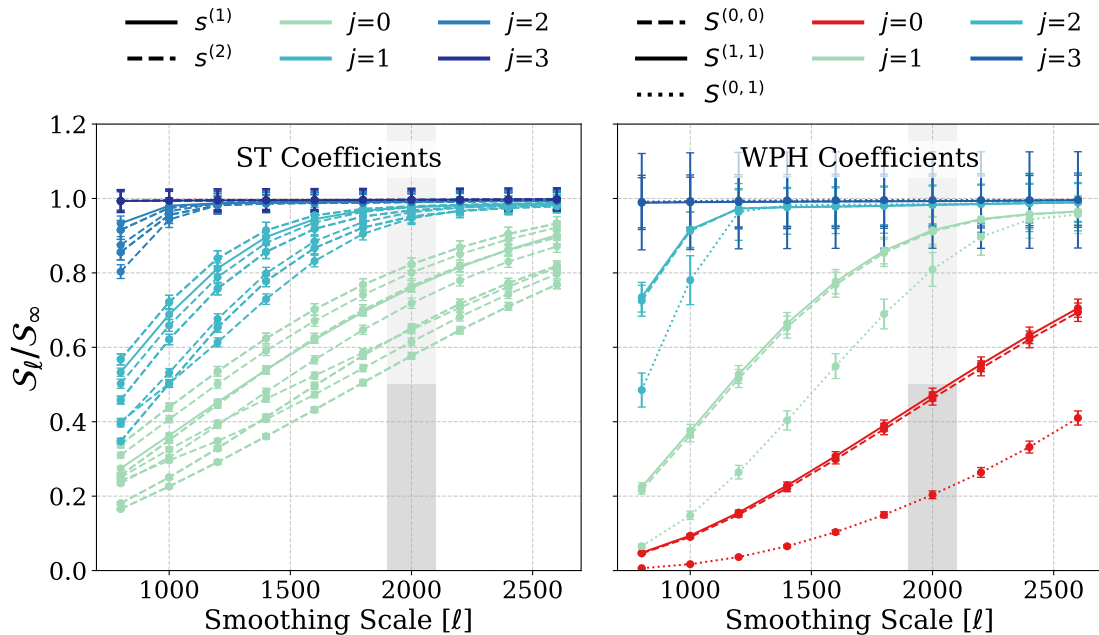


Figure 7: The effect of smoothing scale on the two kinds of wavelet coefficients, ST and WPH. Given a statistics \mathcal{S} , we define \mathcal{S}_ℓ as the statistics measured on κ maps that are low-pass filtered (with a step function) at different ℓ . The statistics measured on the smoothed map is called \mathcal{S}_∞ . We test this idea on κ_{dmo} , and plot the relative suppression $\mathcal{S}_\ell/\mathcal{S}_\infty$ as a function of smoothing scale ℓ . The line and the error bar represent the mean and the standard error of \mathcal{S} measured on different patches across the full sky. We further color code the lines using the characteristic scale of the wavelet (higher j equals larger wavelets). If $\mathcal{S}_\ell/\mathcal{S}_\infty \ll 1$, then \mathcal{S} contains information beyond ℓ . The $j = 0$ WPH coefficients (shown in red) are suppressed more than 50% at $\ell = 2000$ (shaded region), and hence removed from subsequent analyses.

A Scale cut

We empirically determine the characteristic scale of each ST and WPH statistic and enforce a consistent scale cut with respect to the power spectrum. For each statistic \mathcal{S} , we smooth the $N_{\text{side}} = 1,024$ κ_{dmo} maps with low-pass filters at different ℓ_{max} values and measure the statistics (\mathcal{S}_ℓ), normalized by the unsmoothed map's statistics (\mathcal{S}_∞). The results for ST and WPH are shown in the left and right panels of Fig. 7, respectively. For a given statistic, if $\mathcal{S}_\ell/\mathcal{S}_\infty \ll 1$, then it contains information beyond ℓ , and vice versa. We exclude statistics where $\mathcal{S}_\ell/\mathcal{S}_\infty < 0.5$ at $\ell_{\text{max}} = 2,000$. All 34 ST coefficients per redshift bin remain within our analysis scales. For WPH, the three $j = 0$ coefficients (shown in red) fail the scale cut criterion and are removed from the analysis.




Figure 8: Same as Fig. 3, but for the remaining BCM parameters not shown in the main text. Colored curves show the summary statistics of baryonified convergence maps with one parameter varied at a time, while black dashed curves show the corresponding statistics from the FLAMINGO simulations.

B Statistics as a function of baryon parameters

Here we show how the summary statistics vary with the remaining BCM parameters listed in Tab. 1, excluding M_c , θ_{ej} , and η , which are highlighted in the main text (Fig. 3). Each parameter is varied individually across its prior range, and the resulting baryonified statistics are compared to those from the corresponding DMO simulation. These results show the flexibility of the BCM in reproducing a wide range of baryonic effects around the hydrodynamical predictions.

C Emulator

We build a neural network-based emulator to map the baryon parameters to the tomographic ST, WPH, $C(\ell)$ and the m_3 and m_4 statistics. Each type of statistic is modeled independently. We first standardize the input parameters and then use three dense hidden layers following the architectures: (512, 256, 256) for ST, (512, 256, 128) for WPH, (512, 256, 128) for $C(\ell)$ statistics, and (512, 256, 128) for the concatenated m_3 and m_4 data vector. Each hidden layer is followed by a leaky ReLU activation function ($\alpha = 0.1$) (Xu et al. 2015) and a 3% dropout layer to prevent over-fitting (Srivastava et al. 2014). For robustness, we train 10 independent cross-validated models for each type of statistics, resulting in a total ensemble of 40 neural networks. We use 95% of the data for the 10-fold cross-validation and use the rest as the test set for the performance evaluation as shown in Fig. 4. During training, we use the mean absolute error (MAE) as the loss function and the Adam optimizer with weight decay for l_2 regularization (Kingma & Ba 2017). The emulator implements a dual early stopping strategy to prevent over-fitting: (1) a global early stopping mechanism with a patience of 80 iterations that monitors validation loss to terminate training when no improvement is observed, and (2) a learning rate scheduler that starts with a learning rate of 10^{-3} and reduces it by a factor of 0.5 after 5 epochs without improvement, with a minimum threshold of 10^{-6} . Since each statistic is predicted 10 times by the cross-validated models, our emulator not only provides a point estimate (given by the ensemble mean) but also uncertainty estimates (given by the ensemble variance). The entire implementation supports automatic differentiation through JAX, enabling gradient-based optimization and sampling, as well as sensitivity analysis of the statistics with respect to the baryon parameters (Bradbury et al. 2018).

We used the baryonification model to generate 6695 $\kappa_{\text{bary}}(\mathbf{b})$ maps and measured their statistics $\mathcal{S}_{\text{bary}}$, where the 10-dimensional \mathbf{b} is sampled using a Sobol sequence. Since some $\mathcal{S}_{\text{bary}}$ show an amplification rather than a suppression as observed in all FLAMINGO measurements, we remove $\mathcal{S}_{\text{bary}}$ where $C(\ell)$ is on average amplified by 10% across all linear bins. The final training set includes 5468 $\mathcal{S}_{\text{bary}}$. Since the training set is small, the emulator has high outlier rate near parameter boundaries. Therefore, we truncate parameter domain boundaries where the absolute error averaged across all statistics is above 1%. The most significant truncation occurs for $\gamma < 0.9$. We later verify these restrictions do not impact the posteriors of the MCMC chains.

D Best-fit baryon parameters

Table 4: The best-fit BCM parameters obtained from the joint fitting of all HOS to different FLAMINGO variants. These parameter combinations yield HOS that agree with FLAMINGO measurements to 2%. The integers following the feedback variant names are the number of free parameters. The fixed parameters are highlighted in red. The uncertainty is the 16% and 84% credible interval.

Variant	ν_{M_c}	$\nu_{\theta_{ej}}$	γ	$\log_{10} \theta_{co}$	$\log_{10} \eta_{\delta}$	δ	μ_{β}	$\log_{10} M_c/M_{\odot}$	$\log_{10} \eta$	θ_{ej}
Fiducial(L1_m9)(10)	$1.48^{+2.0}_{-2.4}$	$-0.727^{+1.1}_{-0.99}$	$2.61^{+1.7}_{-0.81}$	-1.06 ± 0.19	$-0.247^{+0.12}_{-0.16}$	13.7 ± 3.5	1.81 ± 0.53	13.9 ± 0.25	-1.57 ± 0.43	$5.36^{+2.4}_{-1.7}$
Fiducial(L1_m9)(8)	0	0	$2.64^{+1.8}_{-0.85}$	-1.04 ± 0.18	$-0.255^{+0.13}_{-0.18}$	$14.4^{+3.1}_{-3.7}$	1.77 ± 0.56	13.9 ± 0.18	$-1.49^{+0.48}_{-0.44}$	$4.85^{+1.8}_{-1.3}$
Fiducial(L1_m9)(6)	0	0	2.60	-1.00	$-0.282^{+0.13}_{-0.17}$	$15.1^{+2.8}_{-3.9}$	$1.87^{+0.56}_{-0.48}$	13.9 ± 0.18	-1.43 ± 0.46	$5.06^{+0.91}_{-1.1}$
Fiducial(L1_m9)(4)	0	0	2.60	-1.00	-0.28	14.0	$1.92^{+0.5}_{-0.42}$	14.0 ± 0.16	$-1.52^{+0.37}_{-0.41}$	5.00 ± 0.7
Fiducial(L1_m9)(3)	0	0	2.60	-1.00	-0.28	14.0	1.60	13.9 ± 0.19	-1.51 ± 0.39	4.80 ± 0.65
Jet(10)	$0.697^{+2.4}_{-2.7}$	$-0.227^{+1.1}_{-1.1}$	$2.86^{+2.0}_{-0.98}$	$-0.959^{+0.22}_{-0.21}$	-0.305 ± 0.15	$13.5^{+3.8}_{-3.9}$	$1.50^{+0.72}_{-0.55}$	13.6 ± 0.32	-1.89 ± 0.41	$5.25^{+2.6}_{-1.7}$
Jet(8)	0	0	$2.92^{+2.0}_{-1.0}$	-0.956 ± 0.2	-0.313 ± 0.15	$14.0^{+3.6}_{-4.2}$	$1.45^{+0.72}_{-0.58}$	13.7 ± 0.28	-1.86 ± 0.44	$5.01^{+2.1}_{-1.5}$
Jet(6)	0	0	2.60	-1.00	$-0.282^{+0.13}_{-0.13}$	$14.7^{+4.2}_{-4.1}$	$1.45^{+0.76}_{-0.54}$	13.6 ± 0.27	$-1.80^{+0.38}_{-0.41}$	5.98 ± 1.4
Jet(4)	0	0	2.60	-1.00	-0.28	14.0	$1.49^{+0.74}_{-0.5}$	13.6 ± 0.26	$-1.84^{+0.38}_{-0.42}$	6.15 ± 0.98
Jet(3)	0	0	2.60	-1.00	-0.28	14.0	1.60	13.7 ± 0.16	-1.87 ± 0.39	$6.23^{+0.89}_{-0.75}$
Jet_fgas-4 σ (10)	$1.79^{+1.9}_{-2.3}$	$-0.785^{+1.2}_{-1.0}$	$2.70^{+1.6}_{-0.77}$	-1.08 ± 0.15	$-0.249^{+0.13}_{-0.23}$	$13.5^{+3.4}_{-3.2}$	1.93 ± 0.39	14.0 ± 0.2	$-1.13^{+0.34}_{-0.37}$	$6.84^{+2.8}_{-2.0}$
Jet_fgas-4 σ (8)	0	0	$2.86^{+1.6}_{-0.84}$	-1.07 ± 0.13	$-0.295^{+0.16}_{-0.34}$	$14.3^{+3.2}_{-3.5}$	1.92 ± 0.37	14.1 ± 0.16	$-0.99^{+0.32}_{-0.35}$	$5.78^{+1.9}_{-1.3}$
Jet_fgas-4 σ (6)	0	0	2.60	-1.00	$-0.328^{+0.17}_{-0.27}$	$14.9^{+2.8}_{-3.5}$	$1.94^{+0.36}_{-0.32}$	14.1 ± 0.14	$-0.961^{+0.24}_{-0.31}$	6.14 ± 0.98
Jet_fgas-4 σ (4)	0	0	2.60	-1.00	-0.28	14.0	$1.96^{+0.35}_{-0.31}$	14.1 ± 0.13	$-1.05^{+0.21}_{-0.26}$	$6.12^{+0.69}_{-0.62}$
Jet_fgas-4 σ (3)	0	0	2.60	-1.00	-0.28	14.0	1.60	14.0 ± 0.16	$-1.02^{+0.21}_{-0.25}$	5.94 ± 0.61
fgas+2 σ (10)	$0.761^{+2.4}_{-2.6}$	$-0.328^{+1.1}_{-1.1}$	$2.77^{+2.0}_{-0.9}$	-1.07 ± 0.22	-0.291 ± 0.15	$13.7^{+3.6}_{-4.0}$	$1.58^{+0.71}_{-0.59}$	13.7 ± 0.29	-1.88 ± 0.43	$5.08^{+2.4}_{-1.7}$
fgas+2 σ (8)	0	0	$2.75^{+1.9}_{-0.93}$	$-1.05^{+0.2}_{-0.19}$	-0.302 ± 0.15	$14.1^{+3.4}_{-3.9}$	$1.57^{+0.72}_{-0.59}$	13.7 ± 0.24	-1.85 ± 0.43	$4.82^{+2.1}_{-1.4}$
fgas+2 σ (6)	0	0	2.60	-1.00	$-0.317^{+0.13}_{-0.13}$	$15.0^{+3.2}_{-3.9}$	$1.66^{+0.75}_{-0.64}$	13.7 ± 0.25	-1.80 ± 0.43	$5.19^{+1.1}_{-1.1}$
fgas+2 σ (4)	0	0	2.60	-1.00	-0.28	14.0	$1.66^{+0.79}_{-0.61}$	13.7 ± 0.25	-1.91 ± 0.4	5.10 ± 0.73
fgas+2 σ (3)	0	0	2.60	-1.00	-0.28	14.0	1.60	13.7 ± 0.16	-1.95 ± 0.38	$5.05^{+0.68}_{-0.59}$
fgas-2 σ (10)	$1.69^{+2.0}_{-2.2}$	$-0.802^{+1.0}_{-0.95}$	$2.59^{+1.5}_{-0.73}$	-1.06 ± 0.16	$-0.237^{+0.12}_{-0.19}$	$13.8^{+3.5}_{-3.3}$	1.84 ± 0.45	13.9 ± 0.22	$-1.31^{+0.43}_{-0.39}$	$5.53^{+2.3}_{-1.6}$
fgas-2 σ (8)	0	0	$2.73^{+1.4}_{-0.88}$	-1.03 ± 0.15	$-0.267^{+0.14}_{-0.29}$	$14.5^{+3.2}_{-3.5}$	$1.83^{+0.45}_{-0.43}$	14.0 ± 0.18	$-1.17^{+0.41}_{-0.43}$	$4.86^{+1.8}_{-1.2}$
fgas-2 σ (6)	0	0	2.60	-1.00	$-0.295^{+0.15}_{-0.23}$	$15.2^{+2.8}_{-3.7}$	$1.86^{+0.45}_{-0.38}$	14.0 ± 0.17	$-1.11^{+0.33}_{-0.42}$	$5.08^{+0.83}_{-0.93}$
fgas-2 σ (4)	0	0	2.60	-1.00	-0.28	14.0	$1.88^{+0.4}_{-0.36}$	14.0 ± 0.17	$-1.23^{+0.32}_{-0.34}$	4.99 ± 0.63
fgas-2 σ (3)	0	0	2.60	-1.00	-0.28	14.0	1.60	14.0 ± 0.22	$-1.20^{+0.34}_{-0.35}$	4.85 ± 0.59
fgas-4 σ (10)	$1.68^{+2.0}_{-2.1}$	$-0.822^{+1.0}_{-0.95}$	$2.64^{+1.4}_{-0.73}$	-1.04 ± 0.15	$-0.24^{+0.13}_{-0.24}$	$13.8^{+3.3}_{-3.1}$	1.89 ± 0.41	14.0 ± 0.2	$-1.12^{+0.35}_{-0.37}$	$5.64^{+2.1}_{-1.5}$
fgas-4 σ (8)	0	0	$2.81^{+1.4}_{-0.83}$	-1.02 ± 0.13	$-0.293^{+0.17}_{-0.34}$	$14.4^{+3.1}_{-3.4}$	$1.86^{+0.4}_{-0.37}$	14.0 ± 0.16	$-0.964^{+0.31}_{-0.36}$	$4.86^{+1.5}_{-1.0}$

Continued on next page

Table 4: The best-fit BCM parameters obtained from the joint fitting of all HOS to different FLAMINGO variants. These parameter combinations yield HOS that agree with FLAMINGO measurements to 2%. The integers following the feedback variant names are the number of free parameters. The fixed parameters are highlighted in red. The uncertainty is the 16% and 84% credible interval.

Variant	ν_{M_c}	$\nu_{\theta_{ej}}$	γ	$\log_{10} \theta_{co}$	$\log_{10} \eta_{\delta}$	δ	μ_{β}	$\log_{10} M_c/M_{\odot}$	$\log_{10} \eta$	θ_{ej}
fgas-4 σ (6)	0	0	2.60	-1.00	-0.309 ^{+0.16} _{-0.25}	15.1 ^{+2.8} _{-3.6}	1.85 ^{+0.39} _{-0.34}	14.0 \pm 0.15	-0.935 ^{+0.25} _{-0.32}	5.23 \pm 0.85
fgas-4 σ (4)	0	0	2.60	-1.00	-0.28	14.0	1.88 ^{+0.37} _{-0.31}	14.0 \pm 0.15	-1.02 ^{+0.22} _{-0.28}	5.16 \pm 0.58
fgas-4 σ (3)	0	0	2.60	-1.00	-0.28	14.0	1.60	14.0 \pm 0.19	-0.983 ^{+0.23} _{-0.29}	5.05 \pm 0.54
fgas-8 σ (10)	1.51 ^{+2.0} _{-2.1}	-0.753 ^{+1.1} _{-0.93}	2.74 ^{+1.4} _{-0.76}	-1.04 \pm 0.13	-0.275 ^{+0.16} _{-0.32}	14.0 \pm 3.2	1.92 \pm 0.38	14.0 \pm 0.19	-0.93 ^{+0.28} _{-0.3}	5.94 ^{+2.4} _{-1.6}
fgas-8 σ (8)	0	0	2.96 ^{+1.4} _{-0.85}	-1.03 \pm 0.12	-0.323 ^{+0.18} _{-0.36}	14.6 ^{+3.0} _{-3.4}	1.93 \pm 0.36	14.1 \pm 0.14	-0.857 ^{+0.26} _{-0.24}	5.12 ^{+1.5} _{-1.0}
fgas-8 σ (6)	0	0	2.60	-1.00	-0.322 ^{+0.17} _{-0.25}	15.0 ^{+2.7} _{-3.4}	1.91 ^{+0.35} _{-0.33}	14.1 \pm 0.12	-0.848 ^{+0.21} _{-0.2}	5.63 \pm 0.83
fgas-8 σ (4)	0	0	2.60	-1.00	-0.28	14.0	1.92 \pm 0.35	14.1 \pm 0.12	-0.893 ^{+0.16} _{-0.19}	5.57 \pm 0.56
fgas-8 σ (3)	0	0	2.60	-1.00	-0.28	14.0	1.60	14.0 \pm 0.14	-0.885 ^{+0.17} _{-0.18}	5.40 \pm 0.52
M*- σ (10)	1.70 ^{+2.0} _{-2.3}	-0.901 ^{+1.1} _{-0.94}	2.58 ^{+1.5} _{-0.74}	-1.06 \pm 0.16	-0.232 ^{+0.12} _{-0.22}	13.6 ^{+3.5} _{-3.3}	1.89 ^{+0.47} _{-0.44}	13.9 \pm 0.22	-1.28 ^{+0.43} _{-0.4}	5.48 ^{+2.2} _{-1.6}
M*- σ (8)	0	0	2.75 ^{+1.4} _{-0.87}	-1.04 \pm 0.15	-0.263 ^{+0.14} _{-0.29}	14.4 ^{+3.2} _{-3.4}	1.85 ^{+0.48} _{-0.42}	14.0 \pm 0.18	-1.15 ^{+0.4} _{-0.44}	4.73 ^{+1.6} _{-1.1}
M*- σ (6)	0	0	2.60	-1.00	-0.296 ^{+0.15} _{-0.25}	15.0 ^{+2.9} _{-3.7}	1.90 ^{+0.46} _{-0.39}	14.0 \pm 0.16	-1.09 ^{+0.35} _{-0.42}	4.94 ^{+0.81} _{-0.92}
M*- σ (4)	0	0	2.60	-1.00	-0.28	14.0	1.92 ^{+0.42} _{-0.36}	14.0 \pm 0.16	-1.20 ^{+0.32} _{-0.36}	4.85 \pm 0.57
M*- σ (3)	0	0	2.60	-1.00	-0.28	14.0	1.60	14.0 \pm 0.21	-1.17 ^{+0.33} _{-0.37}	4.69 \pm 0.55
M*- σ _fgas-4 σ (10)	1.60 ^{+2.0} _{-2.2}	-0.863 ^{+1.0} _{-0.9}	2.61 ^{+1.3} _{-0.73}	-1.02 \pm 0.13	-0.268 ^{+0.15} _{-0.34}	13.7 \pm 3.3	1.92 ^{+0.41} _{-0.36}	14.0 \pm 0.19	-0.906 ^{+0.29} _{-0.31}	5.64 ^{+2.2} _{-1.5}
M*- σ _fgas-4 σ (8)	0	0	2.89 ^{+1.3} _{-0.86}	-1.01 \pm 0.12	-0.329 ^{+0.19} _{-0.41}	14.4 ^{+3.0} _{-3.3}	1.91 ^{+0.39} _{-0.35}	14.0 \pm 0.14	-0.814 \pm 0.27	4.81 ^{+1.4} _{-0.97}
M*- σ _fgas-4 σ (6)	0	0	2.60	-1.00	-0.33 ^{+0.17} _{-0.27}	15.1 ^{+2.6} _{-3.4}	1.88 ^{+0.38} _{-0.34}	14.0 \pm 0.13	-0.809 \pm 0.23	5.28 ^{+0.72} _{-0.81}
M*- σ _fgas-4 σ (4)	0	0	2.60	-1.00	-0.28	14.0	1.91 ^{+0.36} _{-0.3}	14.0 \pm 0.12	-0.867 \pm 0.2	5.16 \pm 0.51
M*- σ _fgas-4 σ (3)	0	0	2.60	-1.00	-0.28	14.0	1.60	14.0 \pm 0.15	-0.852 \pm 0.2	5.05 \pm 0.5

E Baryonification error for strong and weak feedback variants

Following the discussion in Sec. 6, we show here the residual errors for the best-fit BCM models applied to two extreme feedback variants from the FLAMINGO suite: the strongest ($f_{\text{gas}}-8\sigma$) and the weakest ($f_{\text{gas}}+2\sigma$). These figures are analogous to Fig. 5, which showed results for the fiducial case. For the strong feedback case ($f_{\text{gas}}-8\sigma$), the BCM reproduces most statistics to within 2%, with a mild 3–4% excess in m_4 at low redshift and small angular scales (Fig. 9). This reflects both modeling limitations and the intrinsic noise in higher-order moments. For the weak feedback case ($f_{\text{gas}}+2\sigma$), residuals remain within 2% across all statistics and scales (Fig. 10).

F Cosmology dependence of baryonic suppression

To test how baryonic suppression of weak lensing summary statistics varies with cosmology, we compute the statistics for both the fiducial and LS8 cosmologies, and for each with the fiducial and strong-AGN ($f_{\text{gas}}-8\sigma$) feedback variants. Fig. 11 shows the relative difference in suppression between LS8 and fiducial cosmologies in the first redshift bin, which is sub-percent for all statistics and scales.




Figure 9: Similar to Fig. 5 but with BCM fitted to the strongest feedback variant in the FLAMINGO suite. The BCM error is still at the percent level. The elevated residual for low redshift m_4 is potentially related to both the baryon modelling error and/or m_4 's sensitivity to noise.




Figure 10: Similar to Fig. 5 but with BCM fitted to the weakest feedback variant in the FLAMINGO suite. The BCM error is still at the percent level.




Figure 11: The sensitivity of the baryonic effect to cosmology for the first redshift bin. The y-axis shows the fractional change between the baryonic suppression of summary statistics at the LS8 cosmology compared to the suppression at the fiducial cosmology, for two different subgrid models: the fiducial models (LS8 and L1_m9) and the strong AGN feedback variants (LS8_fgas-8 σ and fgas-8 σ). Across both feedback variants, the amount of baryonic suppression changes only at the sub-percent level, indicating weak dependence on σ_8 .