





ABACUSHOD: a highly efficient extended multitracer HOD framework and its application to BOSS and eBOSS data

Sihan Yuan ^{1,2}★, Lehman H. Garrison ³, Boryana Hadzhiyska ¹, Sownak Bose ^{1,4} and Daniel J. Eisenstein¹

¹Center for Astrophysics | Harvard & Smithsonian, 60 Garden St, Cambridge, MA 02138, USA

²Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, Stanford, CA 94305, USA

³Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA

⁴Institute for Computational Cosmology, Department of Physics, Durham University, Durham DH1 3LE, UK

Accepted 2021 November 16. Received 2021 October 25; in original form 2021 August 23

ABSTRACT

We introduce the ABACUSHOD model and present two applications of ABACUSHOD and the ABACUSSUMMIT simulations to observations. ABACUSHOD is a Halo Occupation Distribution (HOD) framework written in PYTHON that is particle-based, multitracer, highly generalized, and highly efficient. It is designed specifically with multitracer/cosmology analyses for next-generation large-scale structure surveys in mind, and takes advantage of the volume and precision offered by the new state-of-the-art ABACUSSUMMIT cosmological simulations. The model is also highly customizable and should be broadly applicable to any upcoming surveys and a diverse range of cosmological analyses. In this paper, we demonstrate the capabilities of the ABACUSHOD framework through two example applications. The first example demonstrates the high efficiency and the large HOD extension feature set through an analysis of full-shape redshift-space clustering of BOSS galaxies at intermediate to small scales ($< 30 h^{-1}$ Mpc), assessing the necessity of introducing secondary galaxy biases (assembly bias). We find strong evidence for using halo environment instead of concentration to trace secondary galaxy bias, a result which also leads to a moderate reduction in the ‘lensing is low’ tension. The second example demonstrates the multitracer capabilities of the ABACUSHOD package through an analysis of the extended Baryon Oscillation Spectroscopic Survey cross-correlation measurements between three different galaxy tracers: luminous red galaxies, emission-line galaxies, and quasi-stellar objects. We expect the ABACUSHOD framework, in combination with the ABACUSSUMMIT simulation suite, to play an important role in a simulation-based analysis of the upcoming Dark Energy Spectroscopic Instrument data sets.

Key words: gravitational lensing: weak – methods: numerical – methods: statistical – galaxies: haloes – dark matter – large-scale structure of Universe.

1 INTRODUCTION

In the standard framework of structure formation in a Λ cold dark matter (Λ CDM) universe, galaxies are predicted to form and evolve in dark matter haloes (White & Rees 1978). To extract cosmological information and understand galaxy formation from observed galaxy clustering statistics, it is critical to correctly model the connection between galaxies and their underlying dark matter haloes. The most popular and efficient model of the galaxy-halo connection for cosmological studies is the Halo Occupation Distribution model (HOD; e.g. Peacock & Smith 2000; Scoccimarro et al. 2001; Berlind & Weinberg 2002; Zheng et al. 2005; Zheng, Coil & Zehavi 2007). The HOD is an empirical model that makes the assumption that all galaxies live inside dark matter haloes, and links galaxy occupation to specific halo properties. The most popular form of the HOD assumes that galaxy occupation is determined *solely* by halo mass, an assumption that rests on the long-standing and widely accepted theoretical prediction

that halo mass is the attribute that most strongly correlates with the halo abundance and halo clustering as well as the properties of the galaxies residing in it (White & Rees 1978; Blumenthal et al. 1984).

However, there is mounting evidence that the mass-only HOD model is insufficient in accurately capturing the observed galaxy clustering on moderate to small scales (around and less than $10 h^{-1}$ Mpc). A series of studies employing hydrodynamical simulations and semi-analytic models have found clear evidence that galaxy occupation correlates with secondary halo properties beyond just halo mass (e.g. Zhu et al. 2006; Artale et al. 2018; Zehavi et al. 2018; Bose et al. 2019; Contreras et al. 2019; Hadzhiyska et al. 2020; Xu, Zehavi & Contreras 2021). This phenomenon is commonly known as *Galaxy Assembly Bias*, or just assembly bias hereafter. In Yuan et al. (2021), we present observational evidence for galaxy assembly bias by analysing the full-shape redshift-space clustering of BOSS CMASS galaxies (Eisenstein et al. 2011; Dawson et al. 2013). We found that by including assembly bias and a secondary dependence on the local environment, the HOD model makes a significantly better prediction of the observed redshift-space clustering. We also found that the predicted galaxy–galaxy lensing signal also becomes significantly

* E-mail: sihan.yuan@cfa.harvard.edu

more consistent with data, thus potentially alleviating the ‘lensing is low’ tension, where the observed lensing signal was consistently lower than model predictions by 30–40 per cent (Leauthaud et al. 2017; Lange et al. 2019a). All these studies combine to refute the mass-only ansatz of the basic HOD model, demanding a set of robust physically motivated extensions to improve the HOD’s predictive power.

As shown in Yuan et al. (2021), a key challenge with extended HODs is computational efficiency. The secondary dependencies significantly increase the complexity of the model. Additionally, in order to produce more physical galaxy distributions, we adopt a particle-based approach, where we directly link galaxies to the dark matter particles, as opposed to estimating galaxy positions according to analytical models. However, due to the large number of particles (>100 times that of the haloes) in a high-resolution simulation box, a particle-based approach also significantly increases the computational cost. The combination of using particles and introducing secondary dependencies can make the HOD code too slow for comprehensive parameter space explorations. Indeed, a shortcoming of the Yuan et al. (2021) analysis is that instead of sampling the full extended HOD posterior space, we opted for a much cheaper global optimization routine, which could potentially miss interesting regions of parameter space, particularly in high dimension spaces. Moreover, future cosmological applications of the extended HODs will likely require HOD sampling at a range of different cosmologies, further increasing the computational cost. Thus, performance is of great priority for a robust extended HOD code.

With the advent of a new generation of cosmological surveys with much greater depth, such as the extended Baryon Oscillation Spectroscopic Survey (eBOSS; Dawson et al. 2016) and the Dark Energy Spectroscopic Instrument (DESI; DESI Collaboration 2016), there arises a new opportunity to simultaneously utilize multiple galaxy tracer types to probe structure, the so-called multitracers analysis (e.g. Alam et al. 2020; Zhao et al. 2021). There are three types of galaxies that are most relevant in current and upcoming cosmological surveys: luminous red galaxies (LRGs), which tend to be massive, spheroidal, and quenched; emission-line galaxies (ELGs), which tend to be less massive, disc-like, and star-forming; and quasi-stellar objects (QSOs), whose emissions are dominated by their active galactic nuclei (AGNs). Multitracer studies can not only bring additional statistical power to cosmology studies, but also leverage the potential difference in the clustering of different galaxy types to constrain the physics of galaxy formation. To enable such multitracers analyses, it is extremely helpful to devise a multitracers HOD model, where we simultaneously assign multiple types of galaxies to each halo.

In this paper, we introduce the ABACUSHOD framework and the accompanying PYTHON code module. This HOD framework is extended, efficient, and multitracers. It is developed in conjunction with the state-of-the-art ABACUSSUMMIT simulations (Maksimova et al. 2021) and designed specifically with DESI in mind. However, the model framework is also broadly applicable to other simulations and surveys. This paper also presents two applications of the ABACUSHOD framework and the ABACUSSUMMIT simulations to observational data, demonstrating its effectiveness in modelling observed galaxy clustering down to highly non-linear scales.

We should also mention that there have been several other important HOD analysis frameworks aimed at efficiently deriving HOD fits. One is the so-called emulator approach (e.g. DeRose et al. 2019; Wibking et al. 2019, 2020; Zhai et al. 2019), where a surrogate model is used to approximate the complex clustering predictions.

This approach saves computation by training the surrogate model on a modest number of HOD realizations. However, the success of an emulator model relies delicately on the choice of the surrogate model, while the accuracy often falls off drastically outside the training range. Another approach that avoids emulation is often referred to as tabulation (first described in Zheng & Guo 2016, a popular example is TABCORR¹), which focuses on minimizing the cost of recomputing clustering measurements. These approaches pre-tabulate the halo pair counts and then compute each HOD evaluation as a re-weighted sum of the tabulated halo pair counts. While tabulation can make HOD evaluations quite fast, it is limited by the type of statistics and the binning beforehand. HOD extensions and introducing particles also significantly increase the complexity in tabulated approaches. The ABACUSHOD framework aims to bypass the limits of emulators and tabulation by directly optimizing HOD evaluation and clustering calculation, ensuring maximum flexibility in the HOD model itself.

The paper is outlined as follows: In Section 2, we describe the theory behind the ABACUSHOD framework. In Section 3, we briefly describe ABACUSSUMMIT simulations that the ABACUSHOD is currently built on. In Section 4, we present the core algorithm and its many optimizations. In Section 5, we showcase the first example application of ABACUSHOD, specifically to model the full-shape redshift-space clustering of CMASS galaxies. In Section 6, we showcase the second example application, where we model clustering of the multitracers eBOSS sample. In Section 7, we discuss some interesting aspects of our analyses, especially with regard to the ‘lensing is low’ issue, and compare to previous analyses. Finally, we conclude in Section 8.

2 THE EXTENDED HOD FRAMEWORK

In this section, we introduce the extended multitracers HOD framework, starting with the baseline HOD model for the three dark-time tracers expected for DESI: LRG, ELG, and QSO. Then we describe extensions to the baseline model, including satellite profile variations, velocity bias, assembly bias, and environment-based secondary bias.

2.1 The baseline HOD model

The baseline HOD for LRGs comes from the five-parameter model described in Zheng et al. (2007), which gives the mean expected number of central and satellite galaxies per halo given halo mass:

$$\bar{n}_{\text{cent}}^{\text{LRG}}(M) = \frac{1}{2} \operatorname{erfc} \left[\frac{\log_{10}(M_{\text{cut}}/M)}{\sqrt{2}\sigma} \right], \quad (1)$$

$$\bar{n}_{\text{sat}}^{\text{LRG}}(M) = \left[\frac{M - \kappa M_{\text{cut}}}{M_1} \right]^{\alpha} \bar{n}_{\text{cent}}^{\text{LRG}}(M), \quad (2)$$

where the five parameters characterizing the model are M_{cut} , M_1 , σ , α and κ . M_{cut} characterizes the minimum halo mass to host a central galaxy. M_1 characterizes the typical halo mass that hosts one satellite galaxy. σ describes the steepness of the transition from 0 to 1 in the number of central galaxies. α is the power-law index on the number of satellite galaxies. κM_{cut} gives the minimum halo mass to host a satellite galaxy. We have added a modulation term $\bar{n}_{\text{cent}}^{\text{LRG}}(M)$ to the satellite occupation function to remove satellites from haloes without centrals.

In the baseline implementation, the actual number of central galaxy per halo is drawn from a Bernoulli distribution with mean equal to $\bar{n}_{\text{cent}}^{\text{LRG}}(M)$, and the actual number of satellite galaxies is drawn from

¹<https://github.com/johannesulf/TabCorr>

a Poisson distribution with mean equal to $\bar{n}_{\text{sat}}^{\text{LRG}}(M)$. The central galaxy is assigned to the centre of mass of the largest subhalo, with the velocity vector also set to that of the centre of mass of the largest subhalo. Satellite galaxies are assigned to particles of the halo with equal weights.

Similarly, the baseline HOD for ELGs is based on the parametric model described in Avila et al. (2020) and Alam et al. (2020):

$$\bar{n}_{\text{cent}}^{\text{ELG}}(M) = 2A\phi(M)\Phi(\gamma M) + \frac{1}{2Q} \left[1 + \text{erf} \left(\frac{\log_{10} M_h - \log_{10} M_{\text{cut}}}{0.01} \right) \right], \quad (3)$$

where

$$\phi(x) = \mathcal{N}(\log_{10} M_{\text{cut}}, \sigma_M), \quad (4)$$

$$\Phi(x) = \int_{-\infty}^x \phi(t) dt = \frac{1}{2} \left[1 + \text{erf} \left(\frac{x}{\sqrt{2}} \right) \right], \quad (5)$$

$$A = p_{\text{max}} - 1/Q, \quad (6)$$

whereas the satellite occupation continues to adopt the power-law form of equation 2. Compared to equations 9–12 in Alam et al. (2020), we modified the definition of A by skipping the denominator for ease of computation. We do not notice any significant change to the functional form of the central occupation. This baseline HOD form for ELGs is also confirmed in simulation and semi-analytic model approaches by studies such as Hadzhiyska et al. (2021b) and Gonzalez-Perez et al. (2020).

The baseline QSO HOD adopts the same functional form as the LRGs, referring back to equations 1–2. We also implement an overall amplitude parameter for each tracer to account for incompleteness.

In the current version of ABACUSHOD, we enforce that each halo can only host at most one central galaxy, which can be of any tracer type. Similarly, each particle can also host at most one satellite galaxy, which can be of any tracer type. We do not enforce any type of central-satellite conformity or 2-halo conformity.

In the following subsections, we introduce our framework of extending the baseline HODs with physical generalizations.

2.2 The satellite profile generalizations

In the baseline implementation, we assume that the 1-halo distribution of satellites tracks the halo density profile, where we assign satellites to halo particles with equal weight. Bose et al. (2019) used hydrodynamical simulations to show that this is a reasonable assumption for mass-selected galaxy samples. In this section, we relax that assumption by introducing several physically motivated generalizations that allow satellite profile to deviate from that of the dark matter halo. Our generalizations are based on re-weighting existing particles in the halo instead of simply moving galaxies, thus preserving Newtonian physics. These generalizations correspond to parameters s , s_p , and s_v , which were previously introduced in section 3.2 of Yuan, Eisenstein & Garrison (2018) and section 2.2 of Yuan et al. (2021). We summarize the key ideas here.

For instance, the $s \in [-1, 1]$ parameter biases the satellite profile radially, where $s > 0$ corresponds to the distribution of satellites being puffier than the halo profile, and $s < 0$ corresponding to satellite distribution being more concentrated. The s parameter works by first ranking all particles within a halo by their radial distance to halo centre, and then assigning to each particle a weight that linearly depends on rank. Specifically, the probability for the i -th ranked

particle to host a satellite galaxy is given by

$$p_i = \frac{\bar{n}_{\text{sat}}}{N_p} \left[1 + s \left(1 - \frac{2r_i}{N_p - 1} \right) \right], \quad (i = 0, 1, 2, \dots, N_p - 1), \quad (7)$$

where N_p is the number of particles in the halo, and r_i is the rank of the i -th particle.

Similarly, we introduce the s_v parameter, which biases the satellite profile based on particle peculiar velocity, and the s_p parameter, which is based on particle perihelion distance. A detailed description of these two parameters, including how to estimate perihelion distance, can be found in sections 3.3 and 3.4 of Yuan et al. (2018). There are several motivations in including these satellite profile generalization parameters. Baryonic processes can bias the concentration of baryons within the dark matter potential well (e.g. Abadi et al. 2010; Duffy et al. 2010; Chua et al. 2017; Peirani et al. 2017; Amodeo et al. 2021). Splashback and infall can introduce biases for satellites with eccentric orbits (Adhikari, Dalal & Chamberlain 2014; Behroozi et al. 2014; Diemer & Kravtsov 2014; More et al. 2015a). In our previous analyses, we have also used the s_v parameter as an alternative but more physical model for satellite velocity bias. We discuss this approach further in the following subsection.

2.3 Velocity bias

While velocity measurements are not relevant for the study of galaxy positions in real space, the velocities do become entangled with the line-of-sight (LOS) positions in redshift-space due to redshift-space distortions (Kaiser 1987). Thus, to model the redshift-space clustering of galaxies with high fidelity, we need to introduce a more flexible model of galaxy velocities.

In the baseline implementation, we assume the velocity of the central galaxy to be the bulk velocity of the largest subhalo, and the velocity of the satellite galaxies to be the same as their host particles. Following observational evidence presented in Reid et al. (2014) and Guo et al. (2015), we introduce a velocity bias model that allows for deviations in central and satellite velocities from the underlying dark matter. First, we add an additional Gaussian scatter to the LOS component of the central galaxy velocity, with width equal to the halo particle velocity dispersion. The central galaxy velocity along the LOS is thus given by

$$v_{\text{cent},z} = v_{L2,z} + \alpha_c \delta v(\sigma_{\text{LOS}}), \quad (8)$$

where $v_{L2,z}$ denotes the LOS component of the subhalo velocity, $\delta v(\sigma_{\text{LOS}})$ denotes the Gaussian scatter, and α_c is the central velocity bias parameter, which modulates the amplitude of the central velocity bias effect. By definition, α_c is non-negative, and $\alpha_c = 0$ corresponds to no central velocity bias.

For the satellite galaxies, we allow for their peculiar velocities to be systematically higher or lower than their host particle velocities. This satellite velocity bias effect is given by

$$v_{\text{sat},z} = v_{L2,z} + \alpha_s (v_{p,z} - v_{L2,z}), \quad (9)$$

where $v_{p,z}$ denotes the LOS component of particle velocity, and α_s is the satellite velocity bias parameter. $\alpha_s = 1$ corresponds to no satellite velocity bias.

While this (α_c, α_s) model presented here is the most common implementation of velocity bias, it does break Newtonian physics by modifying satellite velocity without modifying its position. In Yuan et al. (2021), we used a different implementation of velocity bias, where we replaced the α_s parameter with the s_v parameter. The s_v parameter simultaneously modulates the radial distribution and peculiar velocity of the satellite by preferentially assigning

satellites to particles with higher or lower peculiar velocities, thus ensuring Newtonian orbits. However, the observed velocity bias is not necessarily exclusively due to halo physics, but could also be due to redshift systematics, which decouples the satellite velocities from satellite position. In this analysis, we use the more common (α_c , α_s) prescription to marginalize over such observation systematics.

2.4 Secondary biases

Following lessons learned in Yuan et al. (2021), we extend the standard HOD with two secondary dependencies, one on halo concentration (assembly bias), and one on the local environment. The concentration dependency describes the classical galaxy assembly bias effect, where the HOD model depends on the assembly history of the halo (encoded by halo concentration) in addition to halo mass. The local environment dependency is more novel, but it was found to be a necessary tracer of secondary biases in both simulations (Hadzhiyska et al. 2020; Xu et al. 2021) and observations (Yuan et al. 2021).


There are multiple frameworks for introducing these secondary dependencies. Hearin et al. (2016) and Yuan et al. (2021) both adopt a ‘galaxy swap’ approach, where galaxies are swapped between haloes of different secondary properties. This approach naturally conserves the total number density of galaxies, but it tends to be computationally expensive. In the ABACUSHOD package, we adopt a different but computationally cheaper approach, where we analytically mix the secondary property with halo mass. This approach was used previously both in Walsh & Tinker (2019) and Xu et al. (2021). We specifically follow the analytic prescription of Xu et al. (2021), where the secondary halo property is directly tied to the two mass parameters in the baseline HOD, M_{cut} and M_1 :

$$\log_{10} M_{\text{cut}}^{\text{mod}} = \log_{10} M_{\text{cut}} + A_{\text{cent}}(c^{\text{rank}} - 0.5) + B_{\text{cent}}(\delta^{\text{rank}} - 0.5) \quad (10)$$


$$\log_{10} M_1^{\text{mod}} = \log_{10} M_1 + A_{\text{sat}}(c^{\text{rank}} - 0.5) + B_{\text{sat}}(\delta^{\text{rank}} - 0.5), \quad (11)$$

where c and δ are the halo concentration and local overdensity, respectively. These secondary properties are ranked within narrow halo mass bins, and the resulting ranks c^{rank} and δ^{rank} are normalized to range from 0 to 1, with 0.5 corresponding to the median. For example, $c^{\text{rank}} > 0.5$ corresponds to a halo with above median concentration, $c^{\text{rank}} < 0.5$ corresponds to a halo with below median concentration, and the same logic applies for the environment rank δ^{rank} . The tetrad (A_{cent} , B_{cent} , A_{sat} , B_{sat}) form the four parameters describing the amplitude of secondary biases in our HOD model. No secondary bias corresponds to all four parameters being equal to zero. We should also point out that in this prescription, the sign of the secondary bias parameters goes in the opposite direction of the secondary bias parameters in Yuan et al. (2021). For example, in this new prescription, a positive A_{cent} would increase $\log M_{\text{cut}}$ for the more concentrated ($c^{\text{rank}} > 0.5$) haloes, which reduces the number of galaxies in more concentrated halo at fixed halo mass (refer to equations 1 and 3), whereas in the model used in Yuan et al. (2021), a positive A increases the number of galaxies in more concentrated haloes. The same logic applies to the environment-based bias, where we switch A with B . So to summarize so far, positive A (B) means high-concentration (environment) haloes have fewer galaxies and low-concentration (environment) haloes have more galaxies at fixed mass.

The concentration c is defined as the ratio $c = r_{90}/r_{25}$, where r_x refers to the radius that encloses x per cent of the total halo mass.



(a) Derivative against assembly bias parameters A_{cent} and A_{sat}



(b) Derivative against environment-based bias parameters B_{cent} and B_{sat}

Figure 1. Derivatives of the projected galaxy correlation function w_p (equation 15) against the secondary bias parameters. This is to help the readers gain intuition on how the four secondary bias parameters impact the predicted galaxy clustering. The top panel shows the derivative against the assembly bias parameters A_{cent} and A_{sat} , whereas the bottom panel shows the derivatives against the environment-based secondary bias parameters B_{cent} and B_{sat} .

The local overdensity δ is calculated in a very similar fashion to Yuan et al. (2021). First, for each halo, we compute the total enclosed mass of the neighbouring haloes, where a neighbour halo is defined to be within $5 h^{-1} \text{Mpc}$ but beyond the halo radius r_{98} . Then we divide the enclosed mass by the average enclosed mass to obtain the local overdensity. Mathematically, we express this definition as

$$\delta = \frac{M(r_{98} < r < 5 h^{-1} \text{Mpc})}{\langle M(r_{98} < r < 5 h^{-1} \text{Mpc}) \rangle} - 1, \quad (12)$$

where M denotes the enclosed mass of the neighbouring haloes.

To gain intuition on how the secondary biases impact galaxy clustering, we show the derivatives of the galaxy projected clustering function w_p (see equation 15) against each of the four secondary bias parameters in Fig. 1. The top panel shows the derivatives against

assembly bias parameters, whereas the bottom panel shows the derivatives against the environment-based secondary bias parameters. For the assembly bias parameters, we see that they are more important at very small scales ($r_p < 1 h^{-1}$ Mpc) while their effect diminishes at larger scales. At $r_p < 1 h^{-1}$ Mpc, the clustering is dominated by the 1-halo term, i.e. central-satellite clustering and satellite-satellite clustering. It makes sense that both the central and satellite derivatives are positive in this regime. Specifically, for a positive assembly bias parameter, the more concentrated haloes ($c_{\text{rank}} > 0.5$) correspond to a higher $M_{\text{cut}}^{\text{mod}}$ and M_1^{mod} , and higher $M_{\text{cut}}^{\text{mod}}$ and M_1^{mod} mean fewer centrals and satellites for those haloes. Thus, positive assembly bias parameters disfavour more concentrated haloes and favour less concentrated haloes. By putting more centrals and satellite into less concentrated haloes, positive assembly bias boosts the central-satellite and satellite-satellite pair counts at the halo-size scale, which is typically $0.1 - 1 h^{-1}$ Mpc, thus boosting the clustering at those scales.

The bottom panel of Fig. 1 shows that both environment-based bias parameters have negative derivatives in the projected galaxy clustering. This makes sense because for positive environment-based bias parameters, the $M_{\text{cut}}^{\text{mod}}$ and M_1^{mod} parameters for haloes in denser environments are increased. Thus, positive environment bias parameters favour haloes in less dense environments to host galaxies, which leads to lower galaxy pair counts, thus the lower clustering amplitude. We see that the effect is strongest at around $r_p \sim 3 h^{-1}$ Mpc, which is the characteristic scale that our environment is defined at. Compared to the concentration-based assembly bias parameters, it is clear that while assembly bias mostly impacts clustering in the 1-halo term, environment-based secondary bias affects mostly the 2-halo term and extends out to much larger scales.

The clustering signature of these secondary biases is ultimately the combined effect of occupational biases such as the ones modelled in this section, and of how halo clustering depends on secondary properties, an effect known as halo assembly bias (e.g. Gao, Springel & White 2005; Croton, Gao & White 2007). Specifically, at fixed halo mass, high-environment and high-concentration haloes tend to be more clustered, so when one varies the galaxy occupation as a function of these secondary parameters, one also changes the galaxy clustering by favouring more or less clustered haloes. We refer the readers to Hadzhiyska et al. (2021a) for a detailed presentation of the interaction between galaxy occupational variation and halo assembly bias.

3 THE ABACUSSUMMIT SIMULATIONS

In principle, our extended model is not simulation specific, as long as the simulation outputs a halo and particle catalogue. Currently, the ABACUSHOD code package is specifically set up for the ABACUSSUMMIT simulation suite, which is a set of large, high-accuracy cosmological N -body simulations using the ABACUS N -body code (Garrison, Eisenstein & Pinto 2019; Garrison et al. 2021), designed to meet the Cosmological Simulation Requirements of the DESI survey (Levi et al. 2013). ABACUSSUMMIT consists of over 150 simulations, containing approximately 60 trillion particles at 97 different cosmologies. A typical base simulation box contains 6912^3 particles within a $(2 h^{-1} \text{ Gpc})^3$ volume, which yields a particle mass of $2.1 \times 10^9 h^{-1} M_{\odot}$.²

²For more details, see https://abacussummit.readthedocs.io/en/latest/abacus_summit.html

The set of example fits presented in this paper are done primarily using the $z = 0.5$ slice of the AbacusSummit_base_c000_ph000 box, which is $(2 h^{-1} \text{ Gpc})^3$ in volume and adopts the Planck 2018 Λ CDM cosmology ($\Omega_c h^2 = 0.1200$, $\Omega_b h^2 = 0.02237$, $\sigma_8 = 0.811355$, $n_s = 0.9649$, $h = 0.6736$, $w_0 = -1$, and $w_a = 0$).

The COMPASO halo finder is a highly efficient on-the-fly group finder specifically designed for the ABACUSSUMMIT simulations (Hadzhiyska et al. 2022). COMPASO builds on the existing spherical overdensity (SO) algorithm by taking into consideration the tidal radius around a smaller halo before competitively assigning halo membership to the particles in an effort to more effectively deblend haloes. Among other features, the COMPASO finder also allows for the formation of new haloes on the outskirts of growing haloes, which alleviates a known issue of configuration-space halo finders of failing to identify haloes close to the centres of larger haloes.

We also run a post-processing ‘cleaning’ procedure that leverages the halo merger trees to ‘remerge’ a subset of haloes. This is done both to remove overdeblended haloes in the spherical overdensity finder, and to intentionally merge physically associated haloes that have merged and then physically separated. An example of such dissociation is what is known as splashback (e.g. Diemer & Kravtsov 2014; More, Diemer & Kravtsov 2015b; More et al. 2016), where haloes that were once part of a larger haloes have since exited following at least one orbital passage within their former hosts. In Bose et al. (2021), we find that remerging such haloes significantly improves the fidelity of the halo catalogue, and the resulting ‘cleaned’ halo catalogue achieves significantly better fits on data in an HOD analysis. The fits presented in later sections of this paper are carried out with the cleaned halo catalogues.

4 ABACUSHOD: CORE ALGORITHM AND OPTIMIZATIONS

The ABACUSHOD module loads the halo and particle catalogues from the ABACUSSUMMIT simulations and outputs multitracer mock galaxy catalogues. This code is designed particularly for efficient HOD parameter searches, in which many HOD parameter sets will be requested in quick succession. In this section, we describe the core algorithm and some key optimizations implemented to maximize efficiency.

The mock galaxy generation is divided into two stages: a preparation stage and an HOD stage. The preparation stage needs to be run first and serves to process the raw halo and particle files, front-loading all the expensive I/O and optimizes the simulation data for the second much faster HOD evaluation.

4.1 The preparation stage

One key objective of the preparation stage is to downsample the haloes and particles from the simulation box. This is because the speed of evaluating an HOD scales roughly linearly with the number of haloes and particles passed to the HOD code, barring a small amount of overheads. Thus, by optimally downsampling the haloes and particles in the preparation stage, we can substantially increase the efficiency of each evaluation of the HOD stage.

To this end, we implement a mass-dependent downsampling of the haloes and particles. Specifically, we use a sigmoid function to aggressively downsample low-mass haloes while preserving most haloes at high mass, where the turn-off mass and the turn-off rate depends on the tracer type. For the particles, we apply a uniform downsampling for all masses in addition to the sigmoid turn-off. The goal is to reduce the number of particles until the number of particles

per halo is only 10–100 times the expected number of satellites. By default, we implement two sets of downsampling filters, one designed for CMASS LRGs, and the other designed for ELGs and QSOs. The second filter goes to substantially lower halo mass and thus contains a significantly larger number of haloes, resulting in lower performance in the HOD stage. The user should use these downsampling filters as a point of reference and customize the downsampling function as needed. The default filter for LRGs is shown in equation 13.

Another key objective of the preparation stage is to pre-compute all the halo and particle properties necessary for the HOD model and concatenate them into a large contiguous array. Along with relevant halo and particle properties, the code also marks each halo and particle with a random number. The random numbers are for drawing from the central Bernoulli distribution and the satellite Poisson distribution. Pre-generating random numbers for all haloes and particles not only reduces the computational cost when running HOD chains, but also carries the additional benefits of removing realization noise and making the mocks reproducible. The removal of realization noise also allows for calculation of more accurate derivatives of summary statistics against HOD parameters.

4.2 The HOD stage

The HOD stage centres around the ABACUSHOD class object, which loads the downsampled halo and particle catalogues from the preparation stage on to memory when initialized, and then takes input HOD parameters and returns galaxy mock catalogues. In each HOD call, the centrals and satellites are generated separately and then concatenated into one unified output dictionary.

To efficiently generate mocks given an input HOD, we adopt a multithreaded two-pass memory-in-place algorithm accelerated with NUMBA. To maximize the efficiency of multithreading, the first pass serves to exactly determine the galaxy generating workload, and evenly partition the workload across all threads and pre-allocate the amount of memory needed for each thread. It does so by looping through the haloes (particles) and calculating the number of centrals (satellites) to be generated for each halo by comparing the mean number of centrals (satellites) of that halo to its corresponding pre-generated random numbers. Then, it allocates an empty array for all galaxies to be generated, including their properties (position, velocity, and etc.). The galaxy array is then evenly partitioned by the number of threads and each partition is assigned to a thread. Finally on the second pass, each thread loops through its assigned haloes and particles and fills out the galaxy array. A two-pass approach achieves significantly better performance than a typical brute-force approach by storing haloes/particles data in memory, avoiding the costly operation of copy-pasting the entire galaxy array every time new galaxies need to be added.

We further accelerate the halo and particle for-loops with NUMBA no-python compiler to compile the slower Python code into faster machine code, and we take advantage of all available processor cores with by multithreading over the halo and particle loops.

4.3 Utility functions

As part of the ABACUSHOD package, we also provide additional utility functions that are commonly needed for HOD analyses. These include several two-point correlation function (2PCF) calculators, a galaxy–galaxy lensing calculator, and sampling scripts.

The provided 2PCF calculators are based on the highly optimized grid-based CORRFUNC code (Sinha & Garrison 2020). We further optimize the performance of CORRFUNC to match that of the HOD

code. The galaxy–galaxy lensing calculator is highly optimized compared to the HALOTOOLS lensing calculator (Hearin et al. 2017). It takes advantage of the fact that the g-g lensing measurement $\Delta\Sigma$ is a linearly combination of the $\Delta\Sigma$ at each galaxy position. It pre-computes and saves the $\Delta\Sigma$ at each halo and particle position. For each HOD evaluation, it simply conducts a weighted sum of the haloes and particles with galaxy weights given by the HOD. This lensing calculator is suited for fitting lensing measurements, not for making a single lensing prediction due to the high cost of pre-computing halo and particle $\Delta\Sigma$.

We provide two popular methods of HOD sampling with ABACUSHOD. The first one is an MCMC-based script with EMCEE (Foreman-Mackey et al. 2013), and the second one is a nested sampling script with DYNESTY (Speagle & Barbary 2018). We recommend the nested sampling script for much higher sampling efficiency and the natural calculation of model evidence, which is an essential metric for model comparisons. These scripts can be found at `abacusutils/scripts/hod/`. We provide a case study using the nested sampler in later sections.

4.4 Performance

A key characteristic of the ABACUSHOD code is its high efficiency. In this subsection, we offer a performance benchmark of the code. Our test system consists of two Intel Xeon Gold 5218 CPUs clocked at 2.30 GHz, for a total of 32 cores a single node, and 256GB of DDR4 RAM. We use PYTHON version 3.7.9, NUMPY version 1.19.2, and NUMBA version 0.51.2. For the test runs, we process a single ABACUSSUMMIT base simulation box at Planck cosmology, specifically the `AbacusSummit_base_c000_ph000` box. We load the cleaned COMPASO haloes at the $z = 0.5$ snapshot and downsample the haloes and particles using the default filters provided in the package, as shown in equation 13:

$$\begin{aligned} f_{\text{haloes}} &= \frac{1}{1 + 0.1 \exp(-(\log M_h - 13.3) \times 25)}, \\ f_{\text{particles}} &= \frac{4}{200 + \exp(-(\log M_h - 13.7) \times 8)}. \end{aligned} \quad (13)$$

The downsampling reduces the total number of haloes from 3.99×10^8 down to 6.18×10^6 , and the total number of subsample particles from 3.15×10^8 down to 1.77×10^7 . We note that, given the small number of satellites produced, it is likely possible to further downsample the particles. However, we do not further optimize the particle sample for this analysis since generating satellites is not the performance bottleneck in our tests.

We pick a fiducial baseline HOD prescription $\log_{10} M_{\text{cut}} = 12.8$, $\log_{10} M_1 = 14.0$, $\sigma = 0.5$, $\alpha = 1.0$, and $\kappa = 0.5$, roughly resembling a CMASS-like sample. The construction of the ABACUSHOD object, i.e. loading of halo and particle subsamples on to memory, takes approximately 10 s. Then, we run the HOD code once to compile the code with NUMBA, which takes around 10 s. Then we repeat each HOD run 20 times, and take the average run time. Here, we showcase how the run time of the mock generation and the 2PCF calculator depends on the number of threads, and the number density of galaxies. Specifically for the 2PCF calculator, we compute $\xi(r_p, \pi)$ with eight logarithmic bins in r_p between 0.169 and $30 h^{-1}$ Mpc, and six linear bins in π between 0 and $30 h^{-1}$ Mpc.

Fig. 2 shows how the timing of an HOD evaluation and the 2PCF calculation scales with the number of threads, where the galaxy number density is fixed at BOSS CMASS average density, $3 \times 10^{-4} h^3 \text{ Mpc}^{-3}$. We see that both calculations are highly scalable at below $N_{\text{thread}} < 32$, above which we start to lose per thread efficiency




Figure 2. Timing of the HOD evaluation and the 2PCF calculator as a function of the number of threads when running on a 32 core node, at fixed galaxy number density (CMASS average density). Both calculations are scalable. The dashed line shows the minimum total timing, at just below 0.4 s per HOD call. The timing plateaus above 32 threads, where running multiple threads per core (hyper-threading) provides little to no performance gain.




Figure 3. Timing of the HOD code and the 2PCF calculator as a function of galaxy number density at $N_{\text{thread}} = 32$. The timing of the HOD code is largely independent of number density, whereas the 2PCF timing scales roughly linearly with number density.

due to hyper-threading, i.e. running more than one thread per core provides little to no gain. The best-case timing for the HOD code (mock generation) is 0.17 s with 32 threads. For the 2PCF calculator, the best timing is 0.18 s with 64 threads. Fig. 3 shows how the timing scales with the galaxy number density, at fixed number of threads ($N_{\text{thread}} = 32$). The dashed line denotes the CMASS average density. We see that the HOD code is largely independent of the galaxy number density, since the HOD code timing depends on the number of haloes and particles considered, not the number of galaxies produced by that sweep. The 2PCF code scales roughly linearly with

the number density. This suggests that the 2PCF timing is likely dominated by linear overheads, such as gridding. We have optimized the 2PCF calculator by choosing the optimal grid size. The user may find further tuning of the grid size necessary depending on the simulation size and sample density.

The performance of ABACUSHOD in a real world application is highly dependent on the hardware, simulations, and the summary statistics. Our tests are conducted on a single node system with generous memory and relatively fast processors. On a cluster system like cori at NERSC, the user might also benefit from chain-level parallelization instead of focusing on the timing of a single HOD evaluation. We also did not present the timing for lower mass tracers such as ELGs, but internal testings show that evaluating an eBOSS-like ELG HOD is approximately two to three times slower than evaluating a CMASS-like LRG HOD. An important limiting factor in the performance is the summary statistics calculator. While we provide fast calculators for 2PCF and galaxy–galaxy lensing, anything beyond these statistics remains the responsibility of the user for now.

Compared to other existing HOD implementations, ABACUSHOD is ~ 100 times faster than other particle-based algorithms, including the GRAND-HOD code we developed in Yuan et al. (2018) and the HALOTOOLS implementation in Behroozi, Wechsler & Wu (2013). The tabulated HOD (e.g. TABCORR³) approach can be of similar or better performance by pre-computing all the halo clustering and then convolving halo occupation with the pre-computing clustering (Zheng & Guo 2016). The tabulated HOD approach achieves good performance by sacrificing model flexibility, where extending a tabulated HOD models with additional features could add significant complexities to the pre-tabulation and the computation of the convolution.

5 APPLICATION TO BOSS LRG CLUSTERING

In this section, we apply the ABACUSHOD package to fitting the BOSS CMASS and LOWZ LRG clustering. Besides serving as an example application for the ABACUSHOD package, we also compare the constraining power of the projected clustering and the redshift-space clustering. We test the necessity of various HOD extensions for this galaxy sample, which leads to implications in the galaxy–galaxy lensing tension (Leauthaud et al. 2017) (discussed in Section 7.2). For brevity, we lead with a detailed analysis on CMASS but only summarize the key results of the LOWZ analysis.

5.1 BOSS galaxy sample

The Baryon Oscillation Spectroscopic Survey (BOSS; Bolton et al. 2012; Dawson et al. 2013) is part of the SDSS-III programme (Eisenstein et al. 2011). BOSS Data Release 12 (DR12) provides redshifts for 1.5 million galaxies in an effective area of 9329 square degrees divided into two samples: LOWZ and CMASS. The LOWZ galaxies are selected to be the brightest and reddest of the low-redshift galaxy population at $z < 0.4$, whereas the CMASS sample is designed to isolate galaxies of approximately constant mass at higher redshift ($z > 0.4$), most of them being also LRGs (Reid et al. 2016; Rodríguez-Torres et al. 2016). The survey footprint is divided into chunks which are covered in overlapping plates of radius ~ 1.49 deg. Each plate can house up to 1000 fibres, but due to the finite size of

³<https://github.com/johannesulf/TabCorr>

the fibre housing, no two fibres can be placed closer than 62 arcsec, referred to as the fibre collision scale (Guo, Zehavi & Zheng 2012).

For the CMASS analysis, we limit our measurements to the galaxy sample in redshift range $0.46 < z < 0.6$ in DR12, whereas for LOWZ, we adopt redshift range $0.15 < z < 0.4$. We choose these moderate redshift ranges for completeness and to minimize the systematics due to redshift evolution. Applying this redshift range to both the north and south galactic caps gives a total of approximately 600 000 galaxies in our CMASS sample, and just under 400 000 galaxies in our LOWZ sample. The average galaxy number density is given by $n_{\text{CMASS}} = (3.01 \pm 0.03) \times 10^{-4} h^3 \text{Mpc}^{-3}$ for CMASS and $n_{\text{LOWZ}} = (3.26 \pm 0.03) \times 10^{-4} h^3 \text{Mpc}^{-3}$ for LOWZ.

We consider two key two-point statistics on the data. The first is the redshift-space 2PCF $\xi(r_p, \pi)$, which can be computed using the Landy & Szalay (1993) estimator:

$$\xi(r_p, \pi) = \frac{DD - 2DR + RR}{RR}, \quad (14)$$

where DD , DR , and RR are the normalized numbers of data–data, data–random, and random–random pair counts in each bin of (r_p, π) , and r_p and π are transverse and LOS separations in comoving units. For this paper, we choose a coarse binning to ensure reasonable accuracy on the covariance matrix, with eight logarithmically spaced bins between 0.169 and $30 h^{-1} \text{Mpc}$ in the transverse direction, and six linearly spaced bins between 0 and $30 h^{-1} \text{Mpc}$ bins along the LOS direction. The same binning is used for both the CMASS and LOWZ samples.

The second statistic is the projected galaxy 2PCF, commonly referred to as w_p . It is simply defined as the LOS integral of the redshift-space $\xi(r_p, \pi)$,

$$w_p(r_p) = 2 \int_0^{\pi_{\text{max}}} \xi(r_p, \pi) d\pi, \quad (15)$$

where $\pi_{\text{max}} = 30 h^{-1} \text{Mpc}$. We use a finer binning for w_p , with a total of 18 bins between 0.169 and $30 h^{-1} \text{Mpc}$.

We have corrected the data for fibre collision effects following the method of Guo et al. (2012), by separating galaxies into collided and decollided populations and assuming those collided galaxies with measured redshifts in the plate-overlap regions are representative of the overall collided population. The final corrected correlation function can be obtained by summing up the contributions from the two populations. However, scales below $0.5 h^{-1} \text{Mpc}$ likely still suffer from systematics even after the correction, and they show a turn-off that is qualitatively inconsistent with theoretical expectations and simulations. Thus, we remove the first three bins in w_p , and the first column of $\xi(r_p, \pi)$ from the fit. In fact, in our tests, we find that the removal of these bins yield a significantly better fit in terms of χ^2/dof . The covariance matrix is calculated from jackknife subsamples and is described in detail in section 3.1 of Yuan et al. (2021).

5.2 CMASS w_p fit

To fit the observed CMASS projected galaxy 2PCF w_p , we start with our ABACUSSUMMIT base box at Planck cosmology. For this analysis, we use the cleaned COMPASO haloes and their corresponding particles at the $z = 0.5$ snapshot.

We assume Gaussian likelihood and express the log-likelihood in terms of the chi-squared, χ^2 . The χ^2 is given in two parts, corresponding to errors on the projected 2PCF and errors on the galaxy number density:

$$\chi^2 = \chi_{w_p}^2 + \chi_{n_g}^2, \quad (16)$$

Table 1. The assumed priors, the maximum-likelihood values, and posterior medians of the baseline HOD model, when constrained on w_p . We choose the priors to be Gaussians with broad non-informative width.

Parameter name	μ_{prior}	σ_{prior}	Best fit	Posterior median
$\log_{10}(M_{\text{cut}}/h^{-1} M_{\odot})$	13.3	0.5	12.9	13.1
$\log_{10}(M_1/h^{-1} M_{\odot})$	14.3	0.5	14.2	14.3
σ	0.5	0.2	2.7×10^{-3}	0.26
α	1.0	0.3	1.2	1.0
κ	0.5	0.2	0.08	0.45

where

$$\chi_{w_p}^2 = (w_{p,\text{mock}} - w_{p,\text{data}})^T \mathbf{C}^{-1} (w_{p,\text{mock}} - w_{p,\text{data}}), \quad (17)$$

and

$$\chi_{n_g}^2 = \begin{cases} \left(\frac{n_{\text{mock}} - n_{\text{data}}}{\sigma_n} \right)^2 & (n_{\text{mock}} < n_{\text{data}}) \\ 0 & (n_{\text{mock}} \geq n_{\text{data}}), \end{cases} \quad (18)$$

where \mathbf{C} as the jackknife covariance matrix on ξ , and σ_n is the uncertainty of the galaxy number density. The $\chi_{n_g}^2$ is an asymmetric normal around the observed number density n_{data} . When the mock number density is less than the data number density ($n_{\text{mock}} < n_{\text{data}}$), we give a Gaussian-type penalty on the difference between n_{mock} and n_{data} . When the mock number density is higher than data number density ($n_{\text{mock}} \geq n_{\text{data}}$), we invoke the incompleteness fraction f_{ic} that uniformly downsamples the mock galaxies to match the data number density. In this case, we impose no penalty. This definition of $\chi_{n_g}^2$ allows for incompleteness in the observed galaxy sample while penalizing HOD models that produce insufficient galaxy number density or too many galaxies. For the rest of this paper, we set $n_{\text{data}} = 3.0 \times 10^{-4} h^3 \text{Mpc}^{-3}$ and a rather lenient $\sigma_n = 4.0 \times 10^{-5} h^3 \text{Mpc}^{-3}$.

We sample the baseline HOD parameter space, without any extensions, using the DYNESTY nested sampler (Speagle & Barbary 2018; Speagle 2019). While being able to sample the posterior space more efficiently than an Markov Chain Monte Carlo sampler, nested sampling codes such as DYNESTY also compute the Bayesian evidence,

$$\mathcal{Z} = P(D|M) = \int_{\Omega_{\Theta}} P(D|\Theta, M) P(\Theta|M) d\Theta, \quad (19)$$

where M represents the model, D represents the data, and Θ represents the model parameters. The evidence can simply be interpreted as the marginal likelihood of the data given the model, and serves as an important metric in Bayesian model comparisons. Our tests show that, with sufficiently high number of live points, the nested samplings runs also are able to accurately identify the maximum likelihood point in high-dimensional spaces. In our DYNESTY runs, we use 500 live points and a uniform sampler. The stopping criterion is set to $d \log \mathcal{Z} > 0.01$. We assume broad Gaussian priors for all five baseline HOD parameters, as summarized in Table 1.

The best-fitting $\chi^2 = 11$ (dof = 10), and the best-fitting HOD parameters are summarized in Table 1. Fig. 4 shows the 1σ , 2σ , 3σ posterior constraints. The best fit corresponds to a galaxy number density of $n_{\text{fit}} = 5.0 \times 10^{-4} h^3 \text{Mpc}^{-3}$ and a satellite fraction of 9.6 per cent. The best-fitting parameters are largely within the expected range. The small σ value corresponds to a sharp mass cut-off for the central galaxies, which is reasonable given the constant mass selection cuts of CMASS galaxies. However, referring to the 2D marginalized posteriors, the typical value of σ in the fit is closer to 0.3, and the maximum-likelihood mode might be a relative outlier.

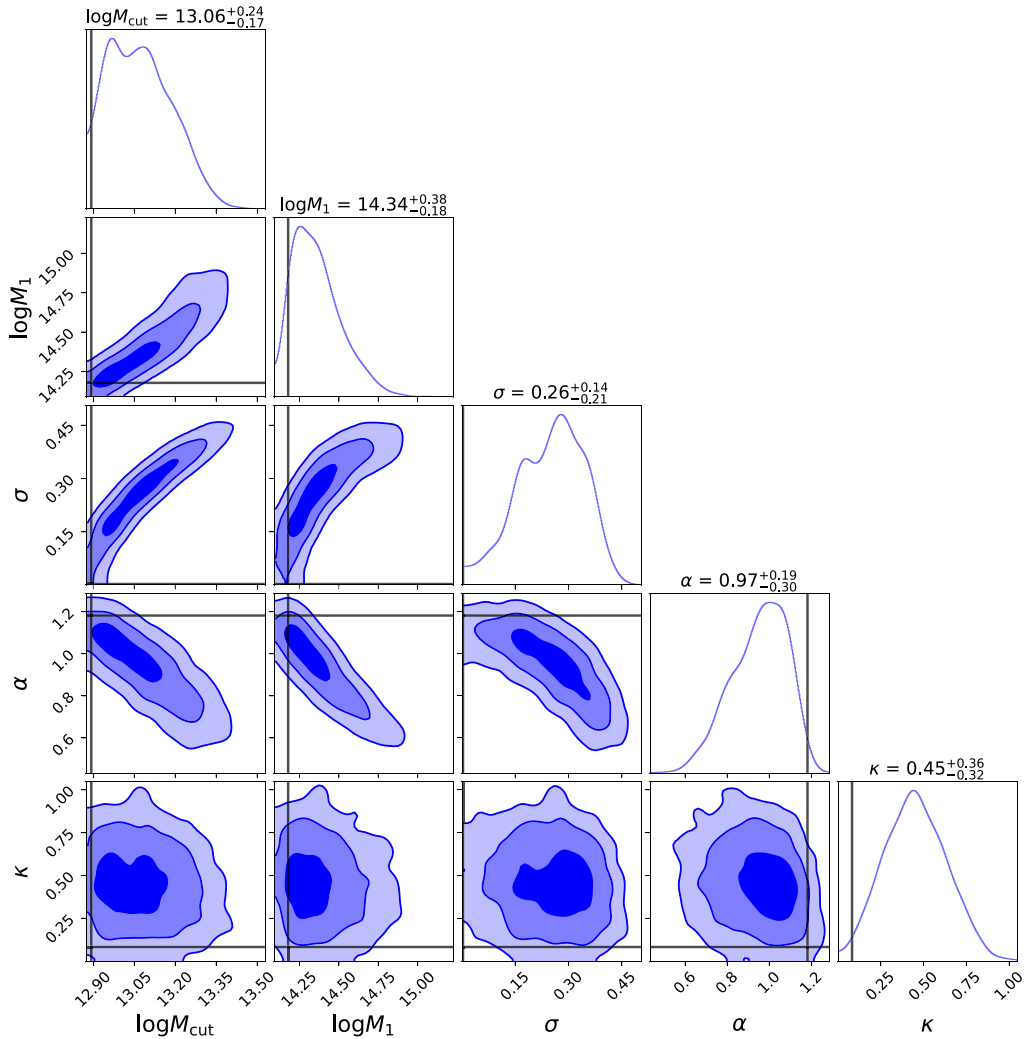


Figure 4. The 1D and 2D marginalized posterior constraints on the baseline HOD parameters from the w_p fit. The contours shown correspond to 1σ , 2σ , and 3σ uncertainties. The vertical and horizontal lines show the maximum-likelihood values for reference. The values displayed above the 1D marginals are posterior medians with the upper/lower bounds associated with the 0.025 and 0.975 quantiles.

It is also apparent from the posterior constraints that the HOD parameters are degenerate and not well constrained. The positive correlation between $\log M_{\text{cut}}$ and $\log M_1$ and the negative correlation between $\log M_{\text{cut}}$ and α suggest a well-constrained satellite fraction. The positive correlation between $\log M_{\text{cut}}$ and σ suggest a well-constrained average bias on the centrals. It is possible that using average bias and satellite fraction would result in a more orthogonal HOD parameter basis.

5.3 CMASS $\xi(r_p, \pi)$ fit

In Yuan et al. (2021), we found that the redshift-space 2PCF, specifically in the form of $\xi(r_p, \pi)$, offers significantly more constraining power on the HOD and assembly bias than the projected 2PCF. In this section, we present ABACUSHOD fits to the BOSS CMASS $\xi(r_p, \pi)$, and discuss evidence for various HOD extensions in the ABACUSHOD framework.

We follow the same routine as outlined for the projected 2PCF w_p , using the same simulation box and the same redshift snapshot ($z = 0.5$). We also assume Gaussian likelihood, where the covariance

matrix is computed from 400 jack knife samples of the data. We additionally apply corrections to the covariance matrix due to the limited simulation volume and the Hartlap correction following Hartlap, Simon & Schneider (2007). Sampling was performed with DYNESTY, with the same settings as before.

In the bare minimum, we need to extend the baseline HOD model by including velocity bias. In Yuan et al. (2021), we used a novel physically motivated implementation of satellite velocity bias, encoded by the parameter s_p . Here, we use the more canonical (α_s, α_c) model of velocity bias to also account for observational systematics. Fig. 5 showcases the posterior constraints, and the best-fitting values are summarized in the third column of Table 2, tagged ‘Baseline’ at the top.

With the $\xi(r_p, \pi)$ fit, we recover reasonable HOD parameter values. Unlike the w_p fit, we do not see strong degeneracies in the marginalized posteriors. As a result, the fit yields much tighter constraints on the HOD parameters $\log M_{\text{cut}}$, $\log M_1$, and α . For example, the 1σ interval on $\log M_{\text{cut}}$ is approximately 15 times tighter than when constrained just on w_p . Similarly, the 1σ interval is five times tighter in $\log M_1$ and three times tighter in α . This




Figure 5. The 1D and 2D marginalized posterior constraints on the baseline HOD parameters and velocity bias parameters from the $\xi(r_p, \pi)$ fit. The contours shown correspond to 1σ , 2σ , and 3σ uncertainties. The vertical and horizontal lines show the maximum-likelihood values for reference. The values displayed above the 1D marginals are posterior medians with the upper/lower bounds associated with the 0.025 and 0.975 quantiles, or approximately the 2σ interval. Compared to the projected clustering w_p constraints shown in Fig. 4, the full-shape redshift-space clustering gives much tighter constraints on parameters and breaks multiple parameter degeneracies.

showcases the extra information provided in the redshift-space 2PCF. In terms of the velocity bias, we find a non-zero central velocity bias with 5σ significance, and a satellite velocity bias consistent with 1. This measurement is consistent with the results of Guo et al. (2015), where the authors also found $\alpha_c \approx 0.2$ and $\alpha_s \approx 1$, albeit with lower signal-to-noise ratio. In Yuan et al. (2021), we found a similar $\alpha_c \approx 0.2$, but we found a significantly non-zero satellite velocity bias $s_v = 0.5\text{--}0.8$, which would suggest the satellite velocity dispersion to be larger than that of the dark matter. However, while the model

difference can be partly responsible for this discrepancy, we also find the fit to be simulation dependent. Yuan et al. (2021) used the ABACUSCOSMOS simulations, which are lower resolution and had a slightly different cosmology. The two simulations also use different halo finders, as detailed in Hadzhiyska et al. (2022) and Bose et al. (2021). We find that if we fit $\xi(r_p, \pi)$ with the same HOD + (α_c, s_v) model as in Yuan et al. (2021), but using the new ABACUSSUMMIT simulations, we recover a much smaller best-fitting value $s_v \approx 0.2 \pm 0.4$, statistically consistent with no satellite velocity bias.

Table 2. Summary of the key HOD fits on CMASS redshift-space 2PCF $\xi(r_p, \pi)$. The first column lists the HOD parameters, incompleteness factor f_{ic} , satellite fraction f_{sat} , the final χ^2 , degree of freedom (dof), and the marginalized Bayesian evidence. The second and third columns show the prior constraints. The third column summarizes the best-fitting parameter values of the baseline CMASS redshift-space 2PCF fit with baseline HOD + velocity bias (α_c, α_s). The following columns list the best-fitting parameters when we introduce additional parameters as shown in the top row. The errors shown are 1σ marginalized errors. The fourth column shows that the addition of the environment-based secondary bias parameters substantially improves the fit and their inclusion is strongly preferred by the data. The negative B_{cent} and B_{sat} values suggest that galaxies preferentially occupy less massive galaxies in denser environments. Fig. 6 shows the 2D posteriors of B_{cent} and B_{sat} , showcasing a $>3\sigma$ detection. The next column shows that the satellite profile parameter s further improves the fit, preferring a less concentrated satellite profile relative to the halo. The last column shows that the concentration-based assembly bias parameters moderately improve the fit. Fig. 8 shows the 2D posteriors of the assembly bias parameters, showing a weaker detection at just above 2σ . The best-fitting values suggest that centrals preferentially occupy more concentrated (older) haloes, whereas satellites occupy less concentrated (younger) haloes, which aligns with theory intuition.

Parameter name	$\mu_{prior} \pm \sigma_{prior}$	Baseline	B_{cent}, B_{sat}	s, B_{cent}, B_{sat}	A_{cent}, A_{sat}
$\log_{10} M_{cut}$	13.3 ± 0.5	$12.86^{+0.01}_{-0.01}$	$12.80^{+0.02}_{-0.02}$	$12.78^{+0.02}_{-0.02}$	$12.88^{+0.02}_{-0.01}$
$\log_{10} M_1$	14.3 ± 0.5	$14.10^{+0.02}_{-0.02}$	$14.00^{+0.04}_{-0.04}$	$13.88^{+0.07}_{-0.05}$	$14.17^{+0.03}_{-0.02}$
$\log_{10} \sigma$	-1 ± 1	$-2.8^{+0.4}_{-0.7}$	$-2.9^{+0.4}_{-0.7}$	$-2.9^{+0.4}_{-0.7}$	$-2.2^{+0.4}_{-0.7}$
α	1.0 ± 0.3	$1.12^{+0.04}_{-0.04}$	$1.03^{+0.04}_{-0.04}$	$1.05^{+0.04}_{-0.04}$	$1.09^{+0.04}_{-0.04}$
κ	0.5 ± 0.2	$0.2^{+0.2}_{-0.1}$	$0.3^{+0.2}_{-0.2}$	$0.5^{+0.2}_{-0.2}$	$0.15^{+0.17}_{-0.15}$
α_c	0.3 ± 0.2	$0.22^{+0.02}_{-0.02}$	$0.18^{+0.03}_{-0.04}$	$0.10^{+0.04}_{-0.05}$	$0.22^{+0.02}_{-0.02}$
α_s	1.0 ± 0.3	$0.98^{+0.03}_{-0.04}$	$1.00^{+0.03}_{-0.03}$	$0.84^{+0.07}_{-0.05}$	$0.98^{+0.03}_{-0.03}$
s	0.0 ± 0.3	–	–	$-0.63^{+0.2}_{-0.1}$	–
A_{cent}	0.0 ± 0.3	–	–	–	$-0.40^{+0.09}_{-0.17}$
A_{sat}	0.0 ± 0.3	–	–	–	$0.2^{+0.2}_{-0.3}$
B_{cent}	0.0 ± 0.3	–	$-0.04^{+0.02}_{-0.02}$	$-0.04^{+0.03}_{-0.03}$	–
B_{sat}	0.0 ± 0.3	–	$-0.17^{+0.11}_{-0.12}$	$-0.15^{+0.09}_{-0.10}$	–
f_{ic}	–	0.58	0.46	0.43	0.58
f_{sat}	–	0.11	0.13	0.15	0.11
χ^2 (dof)	–	60 (35)	42 (33)	33 (32)	54 (33)
$\log \mathcal{Z}$	–	–62	–52	–51	–59

5.4 Introducing secondary biases

In this section, we further extend the baseline + velocity bias model with the environment-based secondary bias and the concentration-based assembly bias. First, we extend the baseline + velocity bias model with the following parameters one at a time: $s, s_v, s_p, (A_{cent}, A_{sat})$, and (B_{cent}, B_{sat}) , where each pair of secondary biases is considered one ‘parameter’. We find that $s, (A_{cent}, A_{sat})$, and (B_{cent}, B_{sat}) significantly improve the best-fitting χ^2 , with $\Delta\chi^2 = -7.1, \Delta\chi^2 = -6.1$, and $\Delta\chi^2 = -18.2$ respectively. The other parameters do not significantly improve the fit. The results that assembly bias and environment-based secondary bias improve the fit on the redshift-space 2PCF are qualitatively consistent with our findings in Yuan et al. (2021).

Since the environment-based secondary bias brings the largest improvement to the fit, we first introduce B_{cent} and B_{sat} and use DYNesty to compute the model evidence and sample the posterior space. This model thus includes the five baseline HOD parameters, velocity bias (α_c, α_s), and environment-based secondary bias (B_{cent}, B_{sat}). The fourth column of Table 2 summarizes the resulting best-fitting parameter values. The model yields the best-fitting $\chi^2 = 42$, a significant improvement over the baseline + velocity bias model. The marginalized evidence also significantly improves, suggesting that the observed redshift-space 2PCF significantly prefers the inclusion of environment-based secondary bias in the model. The blue contours in Fig. 6 show the posterior constraints on B_{cent} , and B_{sat} from this fit. We see that while the marginalized 1D posteriors do not show significant detections, the 2D posterior shows that the preference for non-zero B_{cent} and B_{sat} is quite significant. The

negative values of B_{cent} and B_{sat} are consistent with the positive A_e value reported in Yuan et al. (2021) due to the definitional differences. In both analyses, we find that the data preferentially put galaxies in haloes in denser environments. It is also worth noting that we find the environment-based bias for the satellite galaxies to be stronger than that for the central galaxies. This shows the need for separate secondary bias prescriptions for the centrals and satellites, as opposed to the unified prescription used in Yuan et al. (2021).

Revisiting the other parameters in this fit, we see that we continue to find strong evidence for central velocity bias, but no evidence for satellite velocity bias. Interestingly, the baseline HOD parameter fits seem to be sensitive to the inclusion of environment-based secondary bias. Specifically, we see a decrease in M_{cut}, M_1 , and α compared to the fit without the secondary bias. These decreases translate to moving both centrals and satellites to lower mass haloes. This is the same preference we saw in Yuan et al. (2021), which in turns decreases the predicted weak lensing signal. We revisit the lensing discussion in the following section.

The fifth column of Table 2 summarizes the best-fitting parameters when we also include the satellite profile parameter s in addition to the environment-based secondary bias parameters (B_{cent}, B_{sat}). We see a further improvement to the best-fitting χ^2 , down to 33 with 32 degrees of freedom. The model evidence also sees a further improvement. The introduction of s does not significantly bias the best-fitting values of the environment-based secondary bias parameters, but does affect the best-fitting values of the baseline parameters and the velocity bias parameters. Specifically, we see a further decrease in

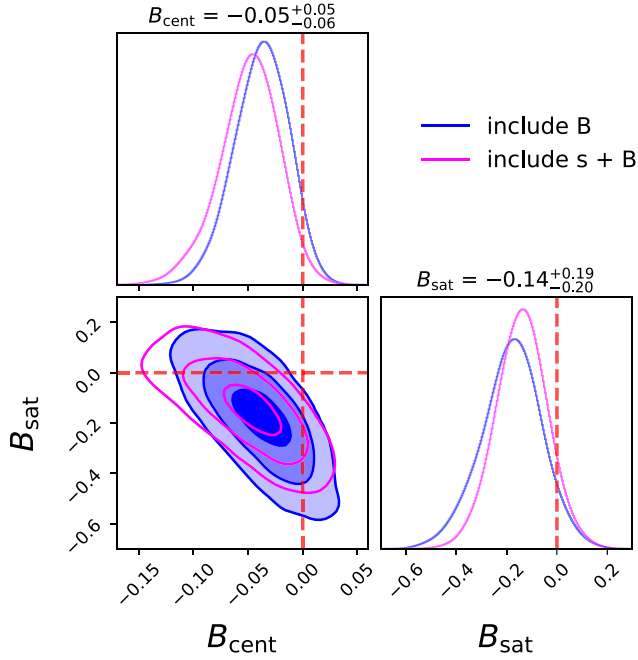


Figure 6. The 1D and 2D marginalized posterior constraints on the environment-based secondary bias parameters B_{cent} , B_{sat} from the $\xi(r_p, \pi)$ fit. The contours shown correspond to 1σ , 2σ , and 3σ uncertainties. The vertical and horizontal lines mark the zeros. The values displayed above the 1D marginals are posterior medians with the upper/lower bounds associated with the 0.025 and 0.975 quantiles, or approximately the 2σ interval. The blue contours show the constraints when we include B_{cent} and B_{sat} in addition to the baseline HOD + velocity bias model. The magenta contours show the constraints when we also include the satellite profile parameter s . We see that while the marginalized 1D posteriors do not show significant detections, the 2D posterior shows that the preference for non-zero B_{cent} and B_{sat} is strong.

the halo mass of central galaxies and satellite galaxies. The decrease in M_1 while α remains the same results in an increase in the inferred satellite fraction. The negative s itself implies a less concentrated satellite galaxy distribution relative to the halo profile, preferring the outer regions of the halo over the halo core. The magenta contours in Fig. 6 show that the inclusion of s does not alter the posterior constraints on environment-based secondary bias parameters.

Fig. 7 showcases our best fit with environment-based secondary bias, corresponding to the fifth column of Table 2. Specifically, the left-hand side shows the target data vector of our analysis, i.e. the CMASS redshift-space $\xi(r_p, \pi)$ measurement. The right-hand side shows the difference between our best fit and the data vector, normalized by data error bars, which we compute from the diagonal of the data covariance matrix. We achieve good fit on most bins, within $1 - 2\sigma$, with the exception being a few bins at $3-5 h^{-1}$ Mpc transverse and large π . However, note that these bins at larger r_p and π tend to be covariant, so the diagonal errors quoted underestimate the true level of uncertainty in the data, and the discrepancy between the data and the best fit is less statistically significant. There is, however, a trend the transverse direction, where the model tends to overestimate at small r_p but underestimate at larger r_p . This suggests there is still a small residual signal that our model has not accounted for fully.

Next, we test the baseline + velocity bias (α_c, α_s) + assembly bias ($A_{\text{cent}}, A_{\text{sat}}$) model. Again, we use DYNesty to compute the model evidence and sample the posterior space. The sixth column of Table 2 summarizes the best-fitting parameter values when we include parameter A_{cent} and A_{sat} . The inclusion of the assembly bias

parameters moderately reduces the χ^2 per degree of freedom, and improves the marginalized Bayesian evidence. The blue contours of Fig. 8 show the posterior constraints on the assembly bias parameters. The 2D constraints show that the detection of central assembly bias is approximately 2σ confidence, whereas that of satellite assembly bias is less than 1σ . The magenta contours show the constraints when we also add the satellite profile parameter s . We recover the same constraints for the central assembly bias but no evidence for the satellite assembly bias. Overall, the evidence for assembly bias is less significant than that of environment-based secondary bias.

The central and satellite assembly bias also seem to exhibit opposite behaviours. The best-fitting values suggest that the centrals tend to live in more concentrated haloes, while the satellites prefer to live in less concentrated haloes. This is consistent with the fact that more concentrated haloes tend to be older, where more satellites have already merged with the central. The discrepant assembly bias signature between centrals and satellites has also recently been found in the BOSS LOWZ sample (Lange et al. 2022), though they found the assembly bias signature to be dependent on redshift. It also appears, based on the 2D constraints, that the central assembly bias and the satellite assembly bias are seemingly uncorrelated.

5.5 LOWZ fits

We repeat our analysis for the BOSS LOWZ sample in redshift range $0.15 < z < 0.40$. We continue to find that the baseline five-parameter HOD model provides a good fit on the projected galaxy 2PCF w_p , yielding best-fitting $\chi^2 = 10$ with a dof = 10. We also find consistent results in the full-shape $\xi(r_p, \pi)$ fit. Most notably, the environment-based secondary bias continues to enable a significantly better fit on the observed redshift-space clustering than the concentration-based assembly bias. Specifically, without any secondary biases, the baseline HOD plus velocity bias model achieves a best-fitting $\chi^2/\text{dof} = 1.42$. With assembly bias parameters A_{cent} and A_{sat} , we get $\chi^2/\text{dof} = 1.36$, but with environment-based bias parameters B_{cent} and B_{sat} , we get a much improved $\chi^2/\text{dof} = 0.97$. This is similar to the behaviour we see for the CMASS fits in Table 2. We skip the detailed figures and tables for brevity.

6 APPLICATION TO EBOSS MULTITRACER CLUSTERING

In this section, we apply the ABACUSHOD framework to fitting the multitracer cross-correlation measurements of the eBOSS galaxy samples. This serves as a scientifically interesting application that showcases the multitracer capabilities of the ABACUSHOD framework.

6.1 The eBOSS sample

The data set comes from the extended Baryon Oscillation Spectroscopic (eBOSS) survey (Dawson et al. 2016). The eBOSS project is one of the programmes within the wider 5-yr Sloan Digital Sky Survey-IV (SDSS-IV; Blanton et al. 2017). The eBOSS sample consists of four different types of tracers, namely LRGs, ELGs, QSOs, and Lyman Alpha Forest. For this analysis, we are using a subset of the eBOSS samples that covers the redshift range from 0.7 to 1.1, where all of the three tracers of interest, namely LRGs, ELGs, and QSOs, overlap. The overlap region can be used to study these tracers with cross-correlations and this results in dense enough galaxy samples to probe the underlying dark matter distribution through the combined samples. We use intermediate versions of the data




Figure 7. The CMASS $\xi(r_p, \pi)$ that we fit to in this section (left-hand panel) and our best fit with environment-based secondary biases (right-hand panel). The right-hand side specifically shows the difference between the best fit and the data, normalized by the data error bars, computed from the diagonal of the jackknife covariance matrix. We achieve good fit on most bins, with the exceptions being the a few bins at a few megaparsecs transverse and large π .

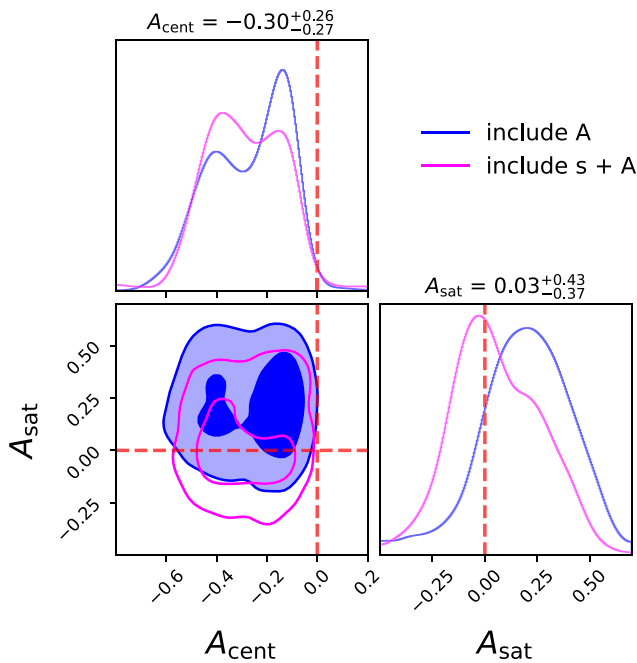


Figure 8. The 1D and 2D marginalized posterior constraints on the assembly bias parameters A_{cent} , A_{sat} from the $\xi(r_p, \pi)$ fit. The contours shown correspond to 1σ and 2σ uncertainties. The 3σ contour is less constrained, and we omit it for better visualization. The vertical and horizontal lines show the zeros. The values displayed above the 1D marginals are posterior medians with the upper/lower bounds associated with the 0.025 and 0.975 quantiles, or approximately the 2σ interval. The blue contours correspond to the constraints when only assembly bias parameters are added to the baseline HOD + velocity bias model, whereas the magenta corresponds to when we also add satellite profile parameter s to the model. We find a 2σ detection of central assembly bias. The detection of satellite assembly bias is weak, especially when we include satellite profile parameter s .

release 16 (DR16) catalogues produced by the eBOSS collaboration (Raichoor et al. 2021). Any changes between the version we have used and the final versions are expected to be minor and mainly affect the results at large scales. The details of target selection and the cross-correlation measurements are found in Alam et al. (2020). The mean number density per tracer is $n_{\text{LRG}} = 1 \times 10^{-4} h^3 \text{Mpc}^{-3}$, $n_{\text{ELG}} = 4 \times 10^{-4} h^3 \text{Mpc}^{-3}$, and $n_{\text{QSO}} = 2 \times 10^{-5} h^3 \text{Mpc}^{-3}$. The exact $n(z)$ distribution can be found in fig. 1 of Alam et al. (2020).

6.2 Fitting eBOSS auto/cross-correlations

The eBOSS cross-correlations are measured within the overlapping footprint of the three tracers, resulting in a set of six projected auto/cross-correlation measurements, as showcased by the orange data points in Fig. 9. The errorbars are estimated from the jackknife covariance for each measurement. We refer the readers to section 4 of Alam et al. (2020) for a detailed description of the data vector measurements and the associated systematics.

For this analysis, we only invoke the baseline HOD for each of the three tracers in ABACUSHOD (equations 1–6), for a total of 20 HOD parameters. Additionally, we account for incompleteness in each tracer. We define the χ^2 similar to equation 16, except the $\chi_{w_p}^2$ is now the summation of six individual χ^2 terms, one for each auto/cross-correlation measurement, and the $\chi_{n_g}^2$ term is the summation of three terms, one for each tracer. We calculate a jackknife covariance for each of the auto/cross-correlation measurement, but we do not account for the covariance between the six measurements in this analysis. To derive the best fit, we follow the methodology of Yuan et al. (2021) in using a global optimization technique known as the covariance matrix adaptation evolution strategy (CMAES; Hansen & Ostermeier 2001). We use an implementation that is part of the publicly available StochOPy (STOCHastic OPTimization for PYthon) package.⁴

⁴<https://github.com/keurfonluu/StochOPy>

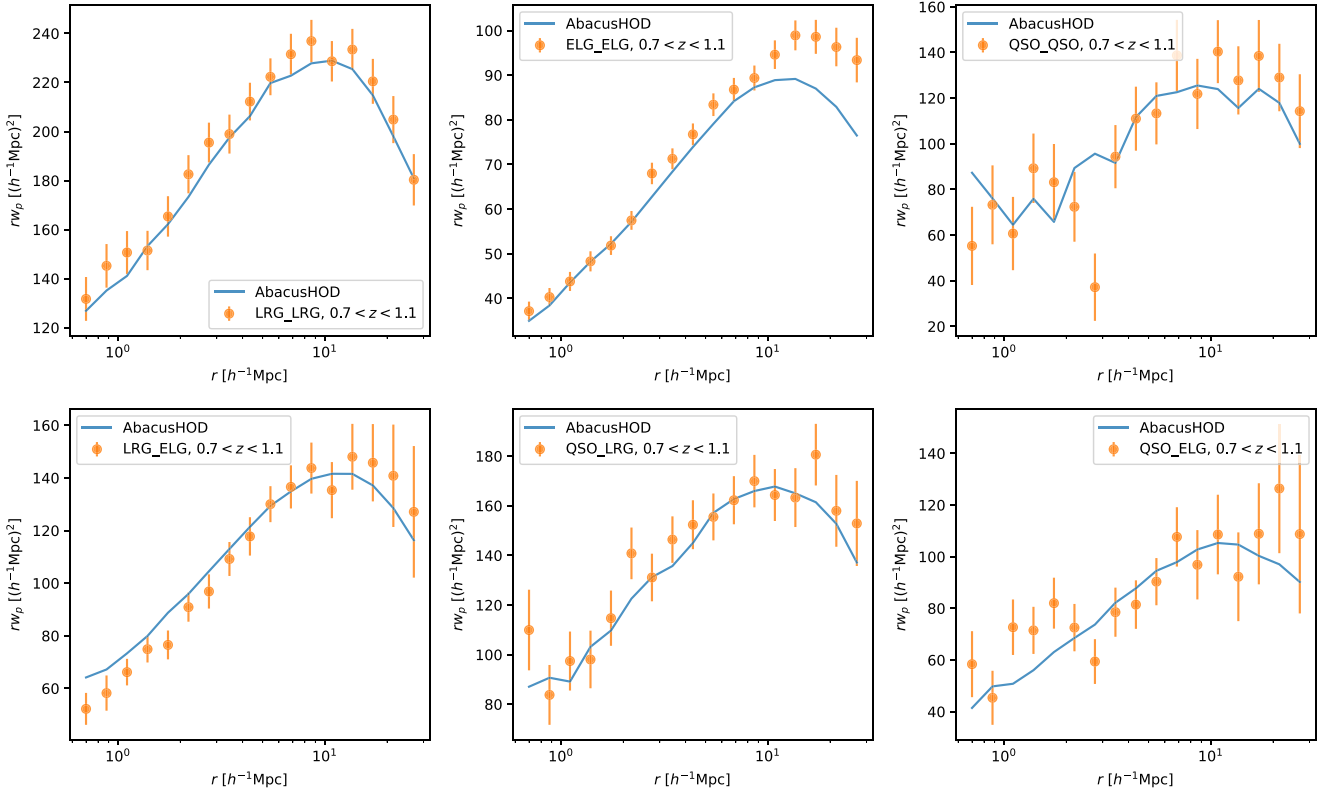


Figure 9. The eBOSS auto/cross-correlation measurements in orange, and the ABACUSHOD best fit in blue. We achieve a good fit, with the best-fitting $\chi^2 = 89$ and dof = 82. Visually, the ELG autocorrelation function appears to show large discrepancies between the data and the prediction in the largest r bins. However, those bins are highly covariant, and the errorbars shown significantly underestimate the true level of uncertainty.

Table 3. The best-fitting ABACUSHOD parameters for all three eBOSS tracers. This should be compared to table 1 of Alam et al. (2020). The ELG column here should be compared to the ELG (HMQ) column in Alam et al. (2020). We do not implement p_{\max} for LRGs as it is redundant with the incompleteness factor in our implementation.

Parameters	LRG	QSO	ELG
$\log_{10} M_{\text{cut}}$	12.8	12.2	11.83
$\log_{10} \sigma$	-1.0	-1.63	-0.24
γ	-	-	5.8
Q	-	-	19.0
$\log_{10}(M_1)$	14.0	14.0	14.0
κ	0.63	0.63	0.82
α	0.78	1.04	0.70
p_{\max}	1.0 (fixed)	0.85	0.68

We achieve a good fit on the data, with the best-fitting $\chi^2 = 89$ and dof = 82. We showcase the best fit with the blue curves in Fig. 9. Visually, the largest deviation comes from the large-scale bins of the ELG autocorrelation function. However, those large-scale bins are highly covariant, and the errorbars shown significantly underestimate the true level of uncertainties. These measurements will be dramatically improved with DESI. Compared to the best fit shown in fig. 5 of Alam et al. (2020), we see broad consistencies between the two fits. This is to be expected as the two fits use equivalent HOD models, though implemented on different simulations. Another difference is that the Alam et al. (2020) analysis only fits the three autocorrelation functions, whereas we fit all six measurements simultaneously.

The best-fitting HOD parameters are summarized in Table 3 and visualized in Fig. 10. Compared to table 1 of Alam et al. (2020),

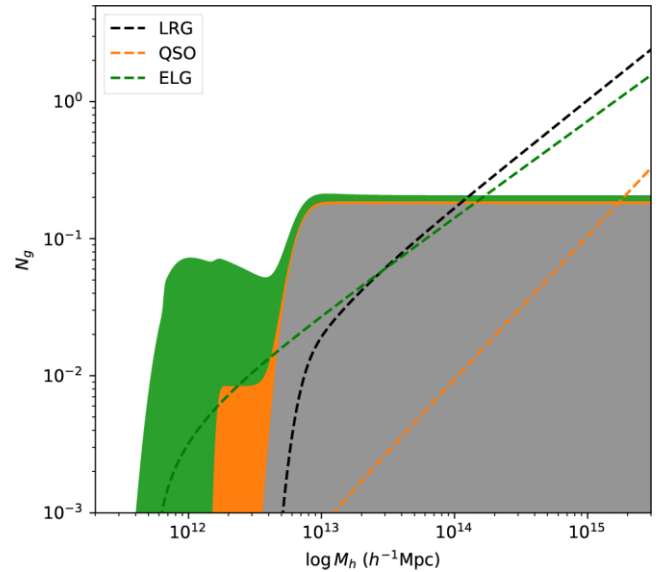


Figure 10. The best-fitting HOD model as a function of halo mass for all three eBOSS tracers. The corresponding parameter values are listed in Table 3. The shaded areas show the central occupation, stacked to show the total galaxy occupation as a function of halo mass. The dashed curves show the satellite distribution. This plot shows the different mass regimes of the three tracers, with LRGs occupying the most massive haloes, whereas ELGs occupation extends down to much lower mass.

there are some inconsistencies. However, these differences can be due to differences in HOD implementation and specifications of the simulations and halo finders. The more important point is that both models can model the auto/cross-correlation functions sufficiently well, and the best-fitting predictions are consistent between the two models. We can compute the typical halo mass from the best-fitting HOD parameters, yielding $M_h^{\text{LRG}} = 1.9 \times 10^{13} h^{-1} M_\odot$, $M_h^{\text{ELG}} = 3.0 \times 10^{12} h^{-1} M_\odot$, and $M_h^{\text{QSO}} = 6.8 \times 10^{12} h^{-1} M_\odot$. This is consistent with the findings of Alam et al. (2020), where the authors found a mean mass per tracer of $M_h^{\text{LRG}} = 1.9 \times 10^{13} h^{-1} M_\odot$, $M_h^{\text{ELG}} = 2.9 \times 10^{12} h^{-1} M_\odot$, and $M_h^{\text{QSO}} = 5 \times 10^{12} h^{-1} M_\odot$. While the mean halo mass of the LRGs and ELGs match exactly, our inferred QSO halo mass is somewhat larger than previous studies but within statistical uncertainty (also refer to Laurent et al. 2017; Rodríguez-Torres et al. 2017).

In terms of satellite fraction, we find the LRGs have a satellite fraction of 13 per cent, consistent with Alam et al. (2020) but slightly lower than Zhai et al. (2017). For QSOs, we find a satellite fraction of approximately 5 per cent, consistent with previous QSO studies (e.g. Rodríguez-Torres et al. 2017), but much lower than the outlier 30 per cent reported in Alam et al. (2020). However, we do find a different mode in the likelihood surface that offers comparable χ^2 and a much higher QSO satellite fraction (~ 34 per cent). We reject that mode for rather extreme parameter values in other HOD parameters. Nevertheless, this highlights the limitation of an optimization analysis, and calls for a comprehensive posterior analysis. For ELGs, we find a satellite fraction of 7 per cent, which is lower than previously reported values in the range of 12 per cent to 17 per cent. Compared to Alam et al. (2020), this is due to us finding both a higher M_1 and a lower α . We suspect this is partially due to differences in the ELG HOD implementation and differences in halo finder. The COMPASO haloes used in this analysis is known to deblend haloes more aggressively than Friends-of-friends halo finders and ROCKSTAR (Behroozi et al. 2013), resulting in more ELGs being identified as centrals.

7 DISCUSSIONS

7.1 Sensitivity to environment definition

The choice of $r_{\text{max}} = 5 h^{-1} \text{Mpc}$ in our environment definition is somewhat arbitrary. It is the recommended value from Hadzhiyska et al. (2020), in which the authors found $r_{\text{max}} = 5 h^{-1} \text{Mpc}$ to be best at capturing the secondary bias signature in hydrodynamical simulations. In Yuan et al. (2021), we also found $r_{\text{max}} = 4\text{--}5 h^{-1} \text{Mpc}$ to be the value that yields the best fit. However, now that we are using a different simulation set and different halo finder, we again test different values of r_{max} . We find no significant improvement to the fit as we vary r_{max} . The lensing prediction of the best-fitting HOD is also largely insensitive to r_{max} . We also test the definition $r_{\text{max}} = 5 \times r_{98}$, where we remove any fixed scale from the definition. Again we find no improvement to the fit compared to the $r_{\text{max}} = 5 h^{-1} \text{Mpc}$ case.

An alternative, computationally cheaper, definition of the local environment is to use a density grid. Specifically, one calculates the local density smoothed over a Gaussian kernel at fixed grid points spanning the entire simulation box. The local density of each halo can be approximately from the overdensities at its nearest grid points, avoiding explicit pair counts. We test this environment definition with a grid size of $N = 1024^3$ and a Gaussian smoothing scale of $3 h^{-1} \text{Mpc}$. Surprisingly, we find that this grid-based environmental secondary bias yields no significant improvement over the no-environment HOD in either the best-fitting χ^2 or the model evidence. While we still need to explore more grid-based variations before

declaring a clear preference, we highlight the high sensitivity of the fit to the details of the HOD model. For now, we continue to recommend the use of enclosed neighbour mass as the local environment definition.

7.2 Galaxy–galaxy lensing comparison

A well-known observational tension exists between galaxy clustering and galaxy–galaxy lensing (g-g lensing). Leauthaud et al. (2017) found discrepancies of 20–40 per cent between their measurements of g-g lensing for CMASS galaxies and a model predicted from mock galaxy catalogues generated at Planck cosmology that match the CMASS projected correlation function (Reid et al. 2014; Saito et al. 2016, see fig. 7 of Leauthaud et al. 2017). Lange et al. (2019b, 2021) extended this result by finding a similar ~ 25 per cent discrepancy between the projected clustering measurement and the g-g lensing measurement in the BOSS LOWZ sample. In Yuan, Eisenstein & Leauthaud (2020), we reaffirmed this tension by fitting simultaneously the projected galaxy clustering and g-g lensing with an extended HOD incorporating a concentration-based assembly bias prescription. However, in Yuan et al. (2021), we found that the inclusion of both the assembly bias and an environment-based secondary bias can significantly reduce ($\sim 10\text{--}20$ per cent) the predicted lensing signal when constrained on the redshift-space 2PCF. We concluded that assembly bias and secondary biases in the HOD could be part of the explanation for the lensing tension.

In this section, we revisit this issue with the new ABACUSHOD fits. First, we reiterate the key differences compared to the Yuan et al. (2021) analysis. First of all, we are now using a different set of simulations, with higher spatial and force resolution, a fundamentally different halo finder, and a slight difference in cosmology. Then, we have updated the HOD model, with a different model for velocity bias, and a different model for secondary biases. Finally, while Yuan et al. (2021) used optimization routines to identify best fits, this analysis enables full posterior sampling, achieving more robust results.

Fig. 11 showcases the g-g lensing predictions of the best-fitting HODs in this analysis, specifically comparing the HOD models with and without environment-based secondary biases. Again, we find that relative to the baseline HOD constrained on w_p , the inclusion of the environment-based secondary bias reduces the lensing prediction by 10–15 per cent, towards better agreement with observation. This shows that the correct secondary bias models not only significantly improve the fit on the redshift-space 2PCF, but also serve an important role in resolving the g-g lensing tension. These two key results, in addition to the findings of Yuan et al. (2021), combine to provide strong evidence for the existence of environment-based secondary bias in the CMASS galaxy population. This detection is only present in fitting the redshift-space 2PCF, demonstrating the extra information contained in the velocity space structure of the full-shape 2PCF.

The blue ‘observed’ lensing data come from the Canada France Hawaii Telescope Lensing Survey (CFHTLenS, Heymans et al. 2012; Miller et al. 2013), and the Canada France Hawaii Telescope Stripe 82 Survey (CS82, Erben et al. 2013). We have also internally compared these measurements to updated measurements from Hyper-Suprime Cam (HSC) survey (Mandelbaum et al. 2018a, b), the Dark Energy Survey (DES, Dark Energy Survey Collaboration 2016; Drlica-Wagner et al. 2018) Y1, and the Kilo-Degree Survey (KiDS-1000, Kuijken et al. 2019; Wright et al. 2020; Giblin et al. 2021; Hildebrandt et al. 2021). We find that the updated measurement is largely consistent with the older CFHTLenS data. The detailed comparison




Figure 11. The lensing prediction of the best-fitting HOD models when including the environment-based secondary bias parameters. The top panel shows the r weighted surface overdensity profile, whereas the bottom panel shows the relative deviation of the predictions from the data. The red curve corresponds to the baseline HOD fit on w_p presented in Section 5.2. The orange curve corresponds to the baseline HOD + velocity bias fit presented in Section 5.3. The solid and dashed purple curves correspond to the environment-based secondary bias fits presented in Section 5.4, where the dashed line also includes the satellite profile parameter s . The shaded purple region shows the corresponding 1σ errorbars. These fits are summarized in Tables 1 and 2.

will be presented in upcoming papers (Amon et al. in preparation, Lange et al. in preparation).

Fig. 12 showcases the g-g lensing predictions when including the concentration-based assembly bias parameters instead of the environment-based parameters. It is clear that the inclusion of assembly bias parameter does not help resolve the lensing tension. This is contrary to the results of Yuan et al. (2021), where we found both assembly bias and environment-based secondary bias reduce the lensing tension. We attribute this to differences in halo finders and the assembly bias models. We discuss these differences more in Appendix A.

In LOWZ, we continue to find that the inclusion of the environment-based bias in the HOD model results in a more realistic lensing prediction, as we show in Fig. 13. The figure shows the relative change to the predicted lensing signal due to the inclusion of assembly bias parameters (dashed lines) and environment-based secondary bias parameters (solid lines). Specifically, $\delta\Delta\Sigma = (\Delta\Sigma_{\text{with bias}}/\Delta\Sigma_{\text{base}}) - 1$, where the baseline prediction comes from the w_p fit with the vanilla five-parameter HOD model. The CMASS prediction is plotted in orange, whereas the LOWZ prediction is plotted in purple. Clearly, on small scales, the 10 per cent reduction in the lensing prediction due to the environment-based bias persists over both samples. The effect of the assembly bias is also consistent across both samples, but remains small in amplitude. This lends further weight to the conclusion from the previous sections that the




Figure 12. The lensing prediction of the best-fitting HOD models when including the assembly bias parameters. The top panel shows the r weighted surface overdensity profile, whereas the bottom panel shows the relative deviation of the predictions from the data. The red curve corresponds to the baseline HOD fit on w_p presented in Section 5.2. The orange curve corresponds to the baseline HOD + velocity bias fit presented in Section 5.3. The solid and dashed purple curves correspond to the assembly bias fits presented in Section 5.4, where the dashed line also includes the satellite profile parameter s . The purple shaded region denotes the corresponding 1σ errorbars.




Figure 13. The correction to the baseline lensing prediction due to the inclusion of assembly bias (dashed lines) and environment-based bias (solid line). CMASS fits are shown in orange, whereas LOWZ fits are in purple. $\delta\Delta\Sigma = (\Delta\Sigma_{\text{with bias}}/\Delta\Sigma_{\text{base}}) - 1$, where the baseline prediction comes from the naive five-parameter HOD fit on w_p . We see environment-based bias consistently lowers the lensing prediction for both samples by 10 per cent on small scales, whereas assembly bias has less impact.

inclusion of environment-based biases in the HOD model not only achieves a good fit on the small-scale full-shape clustering, but also reduces the lensing tension on small scales.

We do caution that the lensing tension remains despite accounting for the environment-based secondary bias. A full resolution of the lensing tension likely requires a combination of secondary bias effects, as demonstrated here, improvements in baryonic effects modelling (Amodeo et al. 2021), and potentially a better accounting of observational systematics (Amon et al. in preparation). A statistically rigorous joint-analysis of the clustering and lensing measurements is required to determine whether the combination of these effects can indeed resolve the lensing tension. We reserve such analysis to future papers.

7.3 Synergies with other works

This work provides growing evidence that the local environment is an important tracer for secondary galaxy bias. Hadzhiyska et al. (2020) and Xu et al. (2021) both systematically tested the effectiveness of various secondary HOD dependencies in capturing the secondary bias signature, Hadzhiyska et al. (2020) through hydrodynamical simulations and Xu et al. (2021) through semi-analytical models. Both studies found the halo environment to be the best indicator of secondary bias. Yuan et al. (2021) and this work complement the previous works by providing the observational support for including environment-based secondary bias in HOD models.

This work also adds another piece to the ‘lensing is low’ puzzle, by showing that the environment-based secondary bias can account for 30 per cent of the lensing discrepancy. Another recent development in resolving this discrepancy comes from the kinetic Sunyaev–Zeldovich (kSZ) effect measurements from the Atacama Cosmological Telescope (ACT) collaboration (Amodeo et al. 2021; Schaan et al. 2021). These studies found that, due to baryonic feedback, the gas profile is significantly more extended in and around the host haloes. A first-order estimate by the ACT team shows that this extended gas profile can reduce the predicted lensing signal by approximately 50 per cent. This shows that a combination of baryonic effects and secondary biases, and potentially a more thorough accounting of data systematics, can reconcile the lensing tension, without the need to invoke any change to the underlying cosmology.

8 CONCLUSIONS

In this paper, we present ABACUSHOD, a new extended multitracer HOD framework built on top of the state-of-art ABACUSUMMIT simulation suite. This HOD framework is feature rich, incorporating flexible models for secondary biases and satellite properties. The code is highly optimized for fast HOD evaluation, enabling robust sampling of extended HOD parameter spaces. In the age of DESI, this code will be an important tool in multitracer analyses and cosmology inference on relatively small scales.

We present two examples applying ABACUSHOD and ABACUSUMMIT to BOSS and eBOSS data. First, we model the full-shape redshift-space 2PCF on small to intermediate scales. We find that the redshift-space 2PCF is significantly more constraining on the HOD parameters than the projected 2PCF, while also breaking parameter degeneracies between M_{cut} , M_1 , σ , and α . We tested various extensions to the baseline + velocity bias model. We find that the observed redshift-space 2PCF strongly prefers the inclusion of environment-based secondary bias, with greater than 3σ detection for the environment-based secondary bias parameters. We find weaker evidence for the canonical concentration-based assembly bias, with just over 2σ detection. This is consistent with several recent studies

that have found the local environment to be the far better indicator of galaxy secondary biases in the HOD.

In the second example, we showcase the multitracer capabilities of ABACUSHOD package by analysing the auto/cross-correlation functions of eBOSS LRG, ELG, and QSO samples. Our model achieves a good fit on the full data vector, yielding consistent predictions with previous analyses.

In the discussion section, we also highlight the fact that by including the environment-based secondary bias, the best-fitting g - g lensing prediction is decreased by approximately 10 per cent in magnitude, accounting for about one-third of the lensing tension. We also find that assembly bias does not significantly lower the lensing prediction. This result is reproduced in both the CMASS and LOWZ sample. In the general sense, this is consistent with the conclusion of Yuan et al. (2021), that secondary biases can potentially partially resolve the lensing tension. Combined with baryonic effects, which have recently been shown to account for up to 50 per cent of the lensing tension, and a more careful accounting of data systematics, we could potentially explain the full ‘lensing is low’ tension.

ACKNOWLEDGEMENTS

We would like to thank Shadab Alam, Johannes Lange, and Josh Speagle for technical guidance in the analyses. We would like to thank David Weinberg, Charlie Conroy, Lars Hernquist, and Douglas Finkbeiner for their helpful feedback.

This work was supported by U.S. Department of Energy grant DE-SC0013718, NASA ROSES grant 12-EUCLID12-0004, NSF PHY-2019786, and the Simons Foundation. SB is supported by the UK Research and Innovation (UKRI) Future Leaders Fellowship (grant number MR/V023381/1).

This work used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231. The ABACUSUMMIT simulations were conducted at the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725, through support from projects AST135 and AST145, the latter through the Department of Energy ALCC program.

DATA AVAILABILITY

The simulation data are available at <https://abacussummit.readthedocs.io/en/latest/>. The ABACUSHOD code package is publicly available as a part of the ABACUSUTILS package at <http://https://github.com/abacusorg/abacusutils>. Example usage can be found at <https://abacusutils.readthedocs.io/en/latest/hod.html>.

REFERENCES

- Abadi M. G., Navarro J. F., Fardal M., Babul A., Steinmetz M., 2010, *MNRAS*, 407, 435
- Adhikari S., Dalal N., Chamberlain R. T., 2014, *J. Cosmol. Astropart. Phys.*, 11, 019
- Alam S., Peacock J. A., Kraljic K., Ross A. J., Comparat J., 2020, *MNRAS*, 497, 581
- Amodeo S. et al., 2021, *Phys. Rev.*, 103, 063514
- Artale M. C., Zehavi I., Contreras S., Norberg P., 2018, *MNRAS*, 480, 3978
- Avila S. et al., 2020, *MNRAS*, 499, 5486
- Behroozi P. S., Wechsler R. H., Wu H.-Y., 2013, *ApJ*, 762, 109
- Behroozi P. S., Wechsler R. H., Lu Y., Hahn O., Busha M. T., Klypin A., Primack J. R., 2014, *ApJ*, 787, 156
- Berlind A. A., Weinberg D. H., 2002, *ApJ*, 575, 587

- Blanton M. R. et al., 2017, *AJ*, 154, 28
- Blumenthal G. R., Faber S. M., Primack J. R., Rees M. J., 1984, *Nature*, 311, 517
- Bolton A. S. et al., 2012, *AJ*, 144, 144
- Bose S., Eisenstein D. J., Hernquist L., Pillepich A., Nelson D., Marinacci F., Springel V., Vogelsberger M., 2019, *MNRAS*, 490, 5693
- Bose S., Eisenstein D. J., Hadzhiyska B., Garrison L. H., Yuan S., 2021, preprint (arXiv:2110.11409)
- Chua K. T. E., Pillepich A., Rodriguez-Gomez V., Vogelsberger M., Bird S., Hernquist L., 2017, *MNRAS*, 472, 4343
- Contreras S., Zehavi I., Padilla N., Baugh C. M., Jiménez E., Lacerna I., 2019, *MNRAS*, 484, 1133
- Croton D. J., Gao L., White S. D. M., 2007, *MNRAS*, 374, 1303
- Dark Energy Survey Collaboration, 2016, *MNRAS*, 460, 1270
- Dawson K. S. et al., 2013, *AJ*, 145, 10
- Dawson K. S. et al., 2016, *AJ*, 151, 44
- DeRose J. et al., 2019, *ApJ*, 875, 69
- DESI Collaboration, 2016, preprint (arXiv:1611.00036)
- Diemer B., Kravtsov A. V., 2014, *ApJ*, 789, 1
- Drlica-Wagner A. et al., 2018, *ApJS*, 235, 33
- Duffy A. R., Schaye J., Kay S. T., Dalla Vecchia C., Battye R. A., Booth C. M., 2010, *MNRAS*, 405, 2161
- Eisenstein D. J. et al., 2011, *AJ*, 142, 72
- Erben T. et al., 2013, *MNRAS*, 433, 2545
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
- Gao L., Springel V., White S. D. M., 2005, *MNRAS*, 363, L66
- Garrison L. H., Eisenstein D. J., Pinto P. A., 2019, *MNRAS*, 485, 3370
- Garrison L. H., Eisenstein D. J., Ferrer D., Maksimova N. A., Pinto P. A., 2021, *MNRAS*, 508, 575
- Giblin B. et al., 2021, *A&A*, 645, A105
- Gonzalez-Perez V. et al., 2020, *MNRAS*, 498, 1852
- Guo H., Zehavi I., Zheng Z., 2012, *ApJ*, 756, 127
- Guo H. et al., 2015, *MNRAS*, 446, 578
- Hadzhiyska B., Bose S., Eisenstein D., Hernquist L., Spergel D. N., 2020, *MNRAS*, 493, 5506
- Hadzhiyska B., Liu S., Somerville R. S., Gabrielpillai A., Bose S., Eisenstein D., Hernquist L., 2021a, *MNRAS*, 508, 698
- Hadzhiyska B., Tacchella S., Bose S., Eisenstein D. J., 2021b, *MNRAS*, 502, 3599
- Hadzhiyska B., Eisenstein D., Bose S., Garrison L. H., Maksimova N., 2022, *MNRAS*, 509, 501
- Hansen N., Ostermeier A., 2001, *Evol. Comput.*, 9, 159
- Hartlap J., Simon P., Schneider P., 2007, *A&A*, 464, 399
- Hearin A. P., Zentner A. R., van den Bosch F. C., Campbell D., Tollerud E., 2016, *MNRAS*, 460, 2552
- Hearin A. P. et al., 2017, *AJ*, 154, 190
- Heymans C. et al., 2012, *MNRAS*, 427, 146
- Hildebrandt H. et al., 2021, *A&A*, 647, A124
- Kaiser N., 1987, *MNRAS*, 227, 1
- Kuijken K. et al., 2019, *A&A*, 625, A2
- Landy S. D., Szalay A. S., 1993, *ApJ*, 412, 64
- Lange J. U., Yang X., Guo H., Luo W., van den Bosch F. C., 2019a, *MNRAS*, 488, 5771
- Lange J. U., van den Bosch F. C., Zentner A. R., Wang K., Hearin A. P., Guo H., 2019b, *MNRAS*, 490, 1870
- Lange J. U., Leauthaud A., Singh S., Guo H., Zhou R., Smith T. L., Cyr-Racine F.-Y., 2021, *MNRAS*, 502, 2074
- Lange J. U., Hearin A. P., Leauthaud A., van den Bosch F. C., Guo H., DeRose J., 2022, *MNRAS*, 509, 1779
- Laurent P. et al., 2017, *J. Cosmol. Astropart. Phys.*, 2017, 017
- Leauthaud A. et al., 2017, *MNRAS*, 467, 3024
- Levi M. et al., 2013, preprint (arXiv:1308.0847)
- Maksimova N. A., Garrison L. H., Eisenstein D. J., Hadzhiyska B., Bose S., Satterthwaite T. P., 2021, preprint (arXiv:2110.11398)
- Mandelbaum R. et al., 2018a, *PASJ*, 70, S25
- Mandelbaum R. et al., 2018b, *MNRAS*, 481, 3170
- Miller L. et al., 2013, *MNRAS*, 429, 2858
- More S., Miyatake H., Mandelbaum R., Takada M., Spergel D. N., Brownstein J. R., Schneider D. P., 2015a, *ApJ*, 806, 2
- More S., Diemer B., Kravtsov A. V., 2015b, *ApJ*, 810, 36
- More S. et al., 2016, *ApJ*, 825, 39
- Peacock J. A., Smith R. E., 2000, *MNRAS*, 318, 1144
- Peirani S. et al., 2017, *MNRAS*, 472, 2153
- Raichoor A. et al., 2021, *MNRAS*, 500, 3254
- Reid B. A., Seo H.-J., Leauthaud A., Tinker J. L., White M., 2014, *MNRAS*, 444, 476
- Reid B. et al., 2016, *MNRAS*, 455, 1553
- Rodríguez-Torres S. A. et al., 2016, *MNRAS*, 460, 1173
- Rodríguez-Torres S. A. et al., 2017, *MNRAS*, 468, 728
- Saito S. et al., 2016, *MNRAS*, 460, 1457
- Schaan E. et al., 2021, *Phys. Rev. D*, 103, 063513
- Scoccimarro R., Sheth R. K., Hui L., Jain B., 2001, *ApJ*, 546, 20
- Sinha M., Garrison L. H., 2020, *MNRAS*, 491, 3022
- Speagle J. S., 2020, *MNRAS*, 493, 3132
- Speagle J., Barbary K., 2018, Astrophysics Source Code Library, record ascl:1809.013
- Walsh K., Tinker J., 2019, *MNRAS*, 488, 470
- White S. D. M., Rees M. J., 1978, *MNRAS*, 183, 341
- Wibking B. D. et al., 2019, *MNRAS*, 484, 989
- Wibking B. D., Weinberg D. H., Salcedo A. N., Wu H.-Y., Singh S., Rodríguez-Torres S., Garrison L. H., Eisenstein D. J., 2020, *MNRAS*, 492, 2872
- Wright A. H., Hildebrandt H., van den Busch J. L., Heymans C., 2020, *A&A*, 637, A100
- Xu X., Zehavi I., Contreras S., 2021, *MNRAS*, 502, 3242
- Yuan S., Eisenstein D. J., Garrison L. H., 2018, *MNRAS*, 478, 2019
- Yuan S., Eisenstein D. J., Leauthaud A., 2020, *MNRAS*, 493, 5551
- Yuan S., Hadzhiyska B., Bose S., Eisenstein D. J., Guo H., 2021, *MNRAS*, 502, 3582
- Zehavi I., Contreras S., Padilla N., Smith N. J., Baugh C. M., Norberg P., 2018, *ApJ*, 853, 84
- Zhai Z. et al., 2017, *ApJ*, 848, 76
- Zhai Z. et al., 2019, *ApJ*, 874, 95
- Zhao C. et al., 2021, *MNRAS*, 503, 1149
- Zheng Z., Guo H., 2016, *MNRAS*, 458, 4015
- Zheng Z. et al., 2005, *ApJ*, 633, 791
- Zheng Z., Coil A. L., Zehavi I., 2007, *ApJ*, 667, 760
- Zhu G., Zheng Z., Lin W. P., Jing Y. P., Kang X., Gao L., 2006, *ApJ*, 639, L5

APPENDIX A: LENSING COMPARISON TO ABACUSCOSMOS

Comparing to the lensing predictions using the ABACUSCOSMOS simulations presented in Yuan et al. (2021), our updated lensing predictions are systematically higher, especially at inner halo scales ($r_p < 0.1 h^{-1}$ Mpc). The ABACUSCOSMOS predictions have a much steeper drop off at small scales. We find that this is at least partially due to the more concentrated core structure of the ABACUSSUMMIT haloes, which is a result of the higher force resolution in the simulation and potentially the COMPASO finder preferentially selecting more compact objects at fixed mass. Fig. A1 showcases the halo density profile of two samples of haloes abundance matched between ABACUSCOSMOS and ABACUSSUMMIT. It is clear that the ABACUSSUMMIT haloes have a significantly more concentrated core at $r < 0.1 h^{-1}$ Mpc.

While the profile difference largely disappears at above $0.1 h^{-1}$ Mpc, the g-g lensing measurement is a cumulative density measurement, meaning that the differences in the core propagates to larger radii. We showcase the volume overdensity $\Delta\rho$ of the same samples of abundance matched haloes in Fig. A2, where the volume overdensity is defined in a similar fashion to the lensing

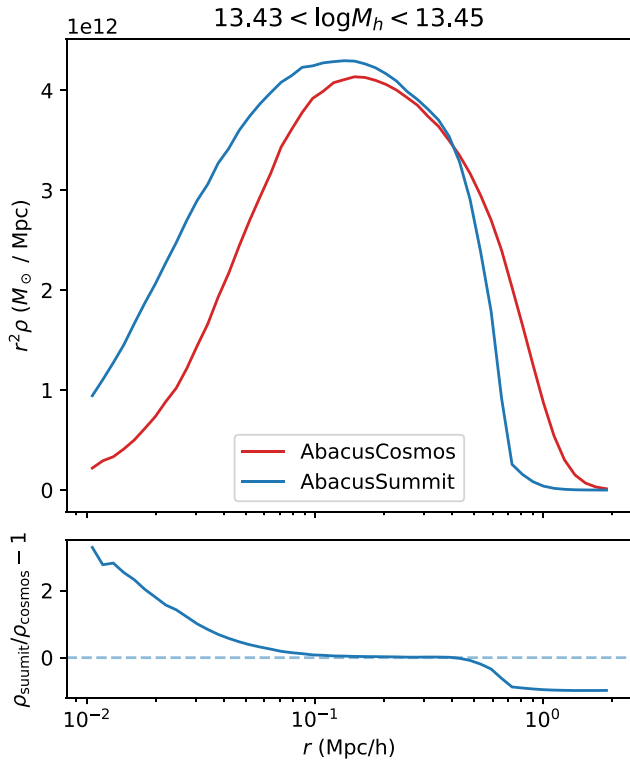


Figure A1. The radial density profile of two samples of haloes abundance matched between ABACUSCOSMOS and ABACUSSUMMIT. The ABACUSSUMMIT sample is selected to be in the mass range $13.43 < \log M_h < 13.45$, which is approximately the typical mass of galaxy hosting haloes. The bottom panel shows the relative difference between the two profiles, and it is apparent that the ABACUSSUMMIT haloes have a much more concentrated core at $r < 0.1 h^{-1}$ Mpc.

measurement, but in 3D:

$$\Delta\rho = \bar{\rho}(< r) - \bar{\rho}(r), \quad (\text{A1})$$

where $\bar{\rho}(< r)$ is the mean density within radius r , and $\bar{\rho}(r)$ is the mean density at r .

It is clear that the difference in core density at $r < 0.1 h^{-1}$ Mpc has a significant impact on the overdensity at $r > 0.1 h^{-1}$ Mpc, and this effect is significant all the way up to $0.8 h^{-1}$ Mpc, where the lensing signal peaks. This suggests that the sub-megaparsec $\Delta\Sigma$ measurement is likely sensitive to uncertainties on the simulated halo core structure.

To confirm this conjecture, Fig. A3 showcases the predicted lensing measurement of two abundance matched samples of halo centres, one in ABACUSCOSMOS and one in ABACUSSUMMIT, both matched to the mean CMASS galaxy number density. As expected, the lensing prediction of the ABACUSCOSMOS haloes is biased low compared to the ABACUSSUMMIT haloes, at scales up to $0.8 h^{-1}$ Mpc. The effect is much stronger at even smaller scales of $0.1 h^{-1}$ Mpc. This confirms that the smallest scale lensing measurements are highly sensitive to systematics in the halo core structure. Even at moderate scales of close to $1 h^{-1}$ Mpc, the halo core structure remains an important systematic. Different force resolution between the ABACUSCOSMOS and ABACUSSUMMIT simulations lead to very different core structure in equivalent haloes, and this difference propagates up to megaparsec scales, at least partially explaining why the ABACUSCOSMOS lensing predictions are biased low relative to the ABACUSSUMMIT prediction at these scales.

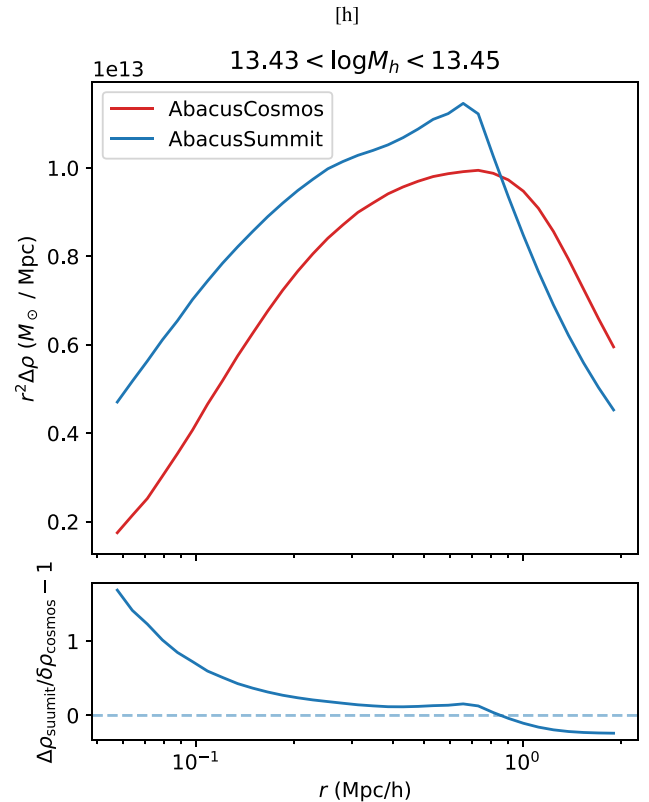


Figure A2. The radial overdensity profile of two samples of haloes abundance matched between ABACUSCOSMOS and ABACUSSUMMIT, the same two samples as shown in Fig. A1. Note the different scale ranges in r when comparing the two plots. It is clear that the difference in core structure at below $r < 0.1 h^{-1}$ Mpc propagates to larger scales in the overdensity. The sharp peak in the ABACUSSUMMIT overdensity profile at $r \approx 0.8 h^{-1}$ Mpc is due to deliberate choices of halo boundary in the COMPASO halo finder.

The other key difference with the Yuan et al. (2021) results is that we do not find significant evidence for the concentration-based assembly bias. Specifically, in the ABACUSCOSMOS simulations, we find that the data prefer to put galaxies in low-mass, low-concentration haloes, thus resulting in the strong assembly bias signature, but we do not reproduce this preference in ABACUSSUMMIT simulations. We believe that this is largely due to the different halo-finding algorithms and possibly different halo concentration definitions. We find that the low-mass, low-concentration haloes in ABACUSSUMMIT simulations tend to be ‘edge haloes’ that live right on the outskirts of larger haloes, resulting in a strong peak in the halo autocorrelation function at $r_p \sim 1 h^{-1}$ Mpc. The low-mass, low-concentration ABACUSCOSMOS haloes on the other hand do not show this peak in the autocorrelation function, suggesting that they are a different population of haloes that do not tend to live on the outskirts of large haloes. This difference is a result of deliberate choices made in the different halo finders. Specifically, as shown in Fig. A1, the COMPASO haloes in ABACUSSUMMIT have sharper boundaries and smaller radii than ROCKSTAR haloes, yielding the remaining mass to additional haloes. In other words, COMPASO deblends haloes more aggressively than ROCKSTAR. None of the halo-finding approaches are necessarily right or wrong, they choose to summarize the density field differently, leading to different results in assembly bias. The key is that both sets of haloes can describe the data reasonably well when combined with a flexible galaxy-halo connection model.




Figure A3. The predicted lensing measurement of halo centres abundance matched to the BOSS CMASS mean galaxy number density. The bottom panel shows the relative difference between the predicted lensing measurement and the observations. The small-scale discrepancy between the two simulations confirms the significant effect of halo core structure on the lensing signal.

This paper has been typeset from a \TeX/L\AA\TeX file prepared by the author.