



Toward Accurate Modeling of Galaxy Clustering on Small Scales: Halo Model Extensions and Lingering Tension

Gillian D. Beltz-Mohrmann^{1,2} , Adam O. Szewciw¹ , Andreas A. Berlind^{1,3} , and Manodeep Sinha^{1,4,5} ¹ Department of Physics and Astronomy, Vanderbilt University, 2201 West End Avenue, Nashville, TN 37235, USA; gbeltzmohrmann@vanderbilt.edu² High Energy Physics Division, Argonne National Laboratory, 9700 South Cass Avenue, Lemont, IL 60439, USA³ National Science Foundation, Division of Astronomical Sciences, Alexandria, VA 22314, USA⁴ SA 118, Center for Astrophysics & Supercomputing, Swinburne University of Technology, 1 Alfred St., Hawthorn, VIC 3122, Australia⁵ ARC Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), Australia

Received 2023 January 3; revised 2023 March 15; accepted 2023 March 16; published 2023 May 11

Abstract

This paper represents an effort to provide robust constraints on the galaxy–halo connection and simultaneously test the Planck Λ CDM cosmology using a fully numerical model of small-scale galaxy clustering. We explore two extensions to the standard Halo Occupation Distribution model: assembly bias, whereby halo occupation depends on both halo mass and the larger environment, and velocity bias, whereby galaxy velocities do not perfectly trace the velocity of the dark matter within the halo. Moreover, we incorporate halo mass corrections to account for the impact of baryonic physics on the halo population. We identify an optimal set of clustering measurements to constrain this “decorated” HOD model for both low- and high-luminosity galaxies in SDSS DR7. We find that, for low-luminosity galaxies, a model with both assembly bias and velocity bias provides the best fit to the clustering measurements, with no tension remaining in the fit. In this model, we find evidence for both central and satellite galaxy assembly bias at the 99% and 95% confidence levels, respectively. In addition, we find evidence for satellite galaxy velocity bias at the 99.9% confidence level. For high-luminosity galaxies, we find no evidence for either assembly bias or velocity bias, but our model exhibits significant tension with SDSS measurements. We find that all of these conclusions still stand when we include the effects of baryonic physics on the halo mass function, suggesting that the tension we find for high-luminosity galaxies may be due to a problem with our assumed cosmological model.

Unified Astronomy Thesaurus concepts: Large-scale structure of the universe (902); Galaxy dark matter halos (1880); Galaxy groups (597); Clustering (1908); Redshift surveys (1378)

1. Introduction

Small-scale galaxy clustering contains a wealth of cosmological information. However, harnessing the constraining power of small scales requires highly accurate models of both dark matter structure formation and the galaxy–halo connection. Halo models are motivated by our understanding that galaxies form and reside in gravitationally bound, virialized regions of dark matter known as halos (e.g., Neyman & Scott 1952; Peebles 1974; McClelland & Silk 1977; Scherrer & Bertshinger 1991; Kauffmann et al. 1997; Jing et al. 1998; Baugh et al. 1999; Kauffmann et al. 1999; Benson et al. 2000; Ma & Fry 2000; Peacock & Smith 2000; Seljak 2000; Scoccimarro et al. 2001; Sheth et al. 2001; White et al. 2001; Cooray & Sheth 2002). These models assume that the clustering of galaxies can be fully described by (i) the clustering of their host halos and (ii) the way in which galaxies occupy these halos.

A key ingredient of the halo model is the Halo Occupation Distribution (HOD), which specifies via a few parameters the probability that a halo of mass M contains N galaxies (above some luminosity threshold) as well as how the galaxies are distributed within their halo (Berlind & Weinberg 2002; Berlind et al. 2003). The standard form of the HOD (Zheng et al. 2005) contains at most five free parameters that specify

the mean occupation number of central and satellite galaxies. This HOD formulation assumes that central galaxies reside at the center of their halo and move at the mean halo velocity, while satellite galaxies trace the spatial and velocity distribution of dark matter inside halos. Constraining these parameters when fitting to observational data provides a useful empirical measurement against which we can test competing theories of galaxy formation and evolution. Moreover, this framework provides a useful avenue for testing cosmology on small scales, assuming that the HOD model is sufficiently flexible to marginalize over the uncertainty of galaxy formation. Compared to other methods like subhalo abundance matching, the HOD model is more flexible and does not rely on having well-resolved subhalos from simulations, making it a more robust framework for probing cosmology.

Many works have used the standard HOD to model the clustering of galaxies in recent redshift surveys (e.g., Zehavi et al. 2011; Guo et al. 2016). Several of these studies yield fits that would rule out the Λ CDM + HOD model if taken at face value. However, these studies typically rely on analytic approximations for calculating clustering statistics, which can introduce unknown systematic uncertainties. Additionally, the errors and covariances used in these studies are typically derived via the jackknife method, which has been shown to produce biased results (Norberg et al. 2009).

Sinha et al. (2018) (S18 hereafter) developed a fully numerical mock-based forward modeling procedure, whereby the standard HOD model is applied to halo catalogs from cosmological N -body simulations and observational survey



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

systematics are added to create realistic mock catalogs. These mocks are then used both for model parameter exploration and for calculating covariance matrices. This significantly improved the accuracy of the HOD modeling framework and allowed for the use of arbitrary clustering statistics that could be *directly* measured on mocks (as opposed to calculated analytically). Using the projected correlation function, group multiplicity function, and galaxy number density, S18 compared the clustering of galaxies in the Sloan Digital Sky Survey (SDSS, York et al. 2000) to a Λ CDM cosmology (Planck Collaboration et al. 2014) + standard HOD model. Carefully controlling for systematic errors allowed them to reliably interpret the goodness of fit of their model. They found that their best-fit HOD model was unable to jointly fit the clustering statistics, revealing moderate tension with SDSS. Because this tension did not exist when they considered only measurements of the projected correlation function (as is done in many studies), S18 demonstrated the value of adding additional statistics in small-scale clustering analyses.

Szewciw et al. (2022) (S22 hereafter) enhanced the procedure used in S18 in order to maximize the return from spectroscopic surveys. They included the same clustering statistics used in S18 (galaxy number density, projected correlation function, and group multiplicity function) as well as four additional clustering statistics: redshift–space correlation function, group velocity dispersion, mark correlation function, and counts-in-cells statistics. Additionally, they developed an algorithm to identify an optimal set of clustering measurements at a variety of different scales in order to maximize constraining power and minimize noise. With these optimal observables, as well as several other improvements to the modeling procedure, they were able to significantly tighten their HOD parameter constraints, as well as dramatically increase the tension found in S18.

The tension found in S22 may be indicative of an issue with the cosmological model used, but it also may be the case that the standard HOD model is simply not flexible enough to accurately encompass the galaxy–halo connection. For example, the standard HOD model assigns galaxies to halos based solely on the halo’s mass, but it is possible that halo occupation depends on additional (secondary) features of the halo (e.g., concentration) that correlate with the halo’s larger-scale environment, a phenomenon known as assembly bias (e.g., Sheth & Tormen 2004; Gao et al. 2005; Wechsler et al. 2006; Croton et al. 2007; Gao & White 2007; Salcedo et al. 2018; Wechsler & Tinker 2018). Evidence for galaxy assembly bias has been found in multiple hydrodynamic simulations (e.g., Artale et al. 2018; Bose et al. 2019; Beltz-Mohrmann et al. 2020; Hadzhiyska et al. 2020; Xu & Zheng 2020; Hadzhiyska et al. 2021a, 2021b, 2021c; Contreras et al. 2021), indicating that it is an expected feature in a Λ CDM universe. Additionally, the standard HOD model assumes that galaxies trace the positions and velocities of the dark matter distribution within their host halo; however, as has been seen in hydrodynamic simulations, this may not be the case (e.g., Berlind et al. 2003; Beltz-Mohrmann et al. 2020). Finally, the typical halo modeling framework relies on dark matter only simulations for creating halo catalogs. These simulations lack baryonic physics, which has been shown to have a significant impact on the halo distribution itself (e.g., Beltz-Mohrmann & Berlind 2021). It is possible that failing to account for the impact of baryonic physics on the halo population is

contributing to the tension between the halo model and the clustering of SDSS galaxies.

Several works have examined the potential for the presence of assembly bias to affect constraints on the galaxy–halo connection and cosmology. For example, Zentner et al. (2014) examined the potential for assembly bias to induce systematic errors in inferred halo occupation statistics. They built mock galaxy catalogs that exhibited assembly bias as well as companion mock catalogs with identical HODs but no assembly bias. They fit HOD models to the galaxy clustering in each catalog, and found that the inferred HODs described the true HODs well in the mocks without assembly bias. However, in the mocks with assembly bias, the inferred HODs exhibited significant systematic errors.

In a later study, McCarthy et al. (2019) used a mock galaxy catalog with assembly bias to study how assembly bias might affect cosmological constraints. Specifically, they used the large-scale redshift–space distortions to probe $f\sigma_8$. They found that on small scales (a few to tens of h^{-1} Mpc) galaxy assembly bias can introduce systematic uncertainties in cosmological constraints if unaccounted for. They concluded that galaxy assembly bias can only be ignored when modeling scales above $8 h^{-1}$ Mpc, where clustering is determined purely by the large-scale bias. Similarly, Lange et al. (2019) explored how galaxy assembly bias affects cosmological inference and found a degeneracy between assembly bias and $f\sigma_8$. Ultimately, they found that not including galaxy assembly bias in the model leads to a small shift in the posterior of $f\sigma_8$, indicating that it is important to account for galaxy assembly bias to obtain unbiased cosmological constraints.

Several recent works have attempted to constrain the galaxy–halo connection and/or cosmology in observational surveys using an extended HOD model that includes assembly bias (e.g., Vakili & Hahn 2019; Zentner et al. 2019; Salcedo et al. 2022; Wang et al. 2022). Many of these works use the Hearin et al. (2016) “decorated” HOD model, which adds two free parameters to the standard HOD model to control the strength of central and satellite occupation on a secondary property. Other works have extended the HOD model to include galaxy velocity bias (Guo et al. 2015a, 2015b), while a few recent works have utilized an extended HOD that includes both assembly bias *and* velocity bias to constrain the galaxy–halo connection and/or cosmology (e.g., Lange et al. 2022; McCarthy et al. 2022; Zhai et al. 2022).

In this work, we build on the procedure established in Sinha et al. (2018) and Szewciw et al. (2022), extending the HOD model to include both assembly bias *and* velocity bias. We explore two different halo properties for implementing assembly bias, and identify an optimal set of clustering measurements to constrain our model. We also implement corrections to our halo masses to account for the impact of baryonic physics on the halo mass function. We use this framework to model the small-scale clustering of both low- and high-luminosity galaxies in SDSS. By using an optimal set of statistics, adding flexibility to our HOD model, and accounting for the potential impact of baryonic physics, our goal is to make the most robust test to date of our assumed Λ CDM cosmological model using small-scale galaxy clustering.

In Section 2, we describe our data, and in Section 3, we describe our simulations and halo catalogs. In Section 4, we describe our halo model, and in Section 5, we describe our full modeling procedure (including our mock galaxy catalogs,

Table 1
SDSS Volume-limited Sample Parameters

M_r^{lim}	z_{min}	z_{max}	z_{median}	V_{eff} ($h^{-3} \text{ Mpc}^3$)	n_g ($h^3 \text{ Mpc}^{-3}$)
-19	0.02	0.07	0.0562	6,087,119	0.01453
-21	0.02	0.158	0.1285	67,174,396	0.00123

Notes. The columns list (from left to right): the absolute magnitude threshold of each sample at $z = 0.1$; the minimum, maximum, and median redshifts; the effective volume; and the galaxy number density of each sample. The volumes and number densities of the samples are corrected for survey incompleteness.

covariance matrices, clustering measurements, and MCMC framework). In Section 6, we describe our selection of optimal observables for constraining our HOD model, and in Section 7, we describe our results. We summarize our findings in Section 8.

2. Observational Data

In this work, we use the same observational data set as that used in S22. We utilize the large-scale structure samples from the NYU Value Added Galaxy Catalog (NYU-VAGC; Blanton et al. 2005) from the seventh data release (DR7; Abazajian et al. 2009) of the Sloan Digital Sky Survey (SDSS; York et al. 2000). The absolute magnitudes of the galaxies in this sample have been k-corrected to rest-frame magnitudes at redshift $z = 0.1$ but have not been corrected for passive luminosity evolution.

From this sample, we construct two volume-limited subsamples, each complete down to a specified r -band absolute magnitude threshold ($M_r < -19$ and $M_r < -21$). We refer to these samples as the -19 and -21 samples throughout this paper. The luminosity thresholds, redshift limits, median redshifts, effective volumes, and number densities of our samples are listed in Table 1. The comoving distances of the SDSS galaxies in our samples are determined using a flat Λ CDM cosmological model with $\Omega_m = 0.302$ and $h = 1$. Our distances are in units of $h^{-1} \text{ Mpc}$, and our absolute magnitudes are actually $M_r + 5 \log h$.⁶

Fiber collisions are handled in the same way as in S22. Briefly, we first adopt the nearest-neighbor correction, and then, informed by SDSS plate overlap regions, we apply additional corrections to our galaxy clustering measurements in order to account for errors in the nearest-neighbor correction. This method was applied in S22, and was recently further validated using the Uchuu-SDSS galaxy lightcones (Dong-Páez et al. 2022). For more details on our observational data and our treatment of fiber collisions, see S22.

3. Simulations and Halo Catalogs

In our modeling procedure, we make use of the same dark matter only cosmological N -body simulations as those used in S22. These simulations are from the Large Suite of Dark Matter Simulations project (LasDamas; McBride et al. 2009) and were run on the Texas Advanced Computing Center’s Stampede supercomputer using the public code GADGET-2 (Springel 2005). Power spectra were generated with CMBFAST (Seljak & Zaldarriaga 1996; Zaldarriaga et al. 1998; Zaldarriaga & Seljak 2000), and initial conditions were generated

with 2LPTIC (Scoccimarro 1998; Crocce et al. 2006, 2012). All simulations were run with the following cosmological parameters, based on results from the Planck experiment (Planck Collaboration et al. 2014): $\Omega_m = 0.302$, $\Omega_\Lambda = 0.698$, $\Omega_b = 0.048$, $h = 0.681$, $\sigma_8 = 0.828$, and $n_s = 0.96$. For each observational sample of interest (i.e., -19 and -21), we run two sets of simulations: 100 low-resolution simulations to build a covariance matrix, and two high-resolution simulations for model parameter exploration. The details of these simulations are given in Table 2. We identify halos with a spherical overdensity algorithm (SO; Lacey & Cole 1994) using the ROCKSTAR phase-space temporal halo finder (Behroozi et al. 2013). We adopt a virial mass definition with a density threshold given by Bryan & Norman (1998). Finally, for computational purposes, we randomly downsample to keep only 5% of the dark matter particles in each halo, with no loss of accuracy (see S22).

4. Halo Model

4.1. The Standard HOD Model

The Halo Occupation Distribution framework governs the number, positions, and velocities of galaxies within dark matter halos. The standard HOD model assigns galaxies to halos based on five free parameters, which depend only on the halo’s mass (Zheng et al. 2007). Galaxies are split into centrals and satellites within their halos (Kravtsov et al. 2004; Zheng et al. 2005). In this model, the mean number of central galaxies in a halo of mass M is described by

$$\langle N_{\text{cen}} \rangle = \frac{1}{2} \left[1 + \text{erf} \left(\frac{\log M - \log M_{\text{min}}}{\sigma_{\log M}} \right) \right], \quad (1)$$

where M_{min} is the mass at which half of halos host a central galaxy, $\sigma_{\log M}$ is the scatter around this halo mass, and $\text{erf}(x)$ is the error function, $\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x \exp(-y^2) dy$. For a specific halo of mass M , we draw a random number R from a uniform distribution on the interval $[0, 1)$. If $R < \langle N_{\text{cen}} \rangle$, then a central galaxy is assigned to the halo. The central galaxy is always placed at the center of the halo and given the mean velocity of the halo.

The number of satellite galaxies in a given halo is drawn from a Poisson distribution with mean

$$\langle N_{\text{sat}} \rangle = \langle N_{\text{cen}} \rangle \times \left(\frac{M - M_0}{M_1} \right)^\alpha, \quad (2)$$

where M_0 is the halo mass below which there are no satellite galaxies, M_1 is the mass where halos contain one satellite galaxy on average, and α is the slope of the power-law occupation function at high masses. Each satellite galaxy is given the position and velocity of a randomly selected dark matter particle within the halo.

4.2. Assembly Bias

One way in which we can extend the standard HOD model is to relax the assumption that halo occupation depends solely on the halo’s mass. In other words, we can allow for halo occupation to depend on both halo mass *and* a secondary halo property, a phenomenon known as assembly bias (Gao et al. 2005; Croton et al. 2007). To implement assembly bias, we use

⁶ Throughout this paper, \log refers to \log_{10} .

Table 2
Simulation Parameters

Use	Sample	Simulation	Seeds	L_{box} (h^{-1} Mpc)	N_{part}	m_{part} ($h^{-1}M_{\odot}$)	ϵ (h^{-1} kpc)	Number
Covariance matrix	−19	Consuelo	4001–4100	420	1400 ³	2.26×10^9	8	100
Covariance matrix	−21	Carmen	2001–2100	1000	1120 ³	5.97×10^{10}	25	100
MCMC	−19	ConsueloHD	4002, 4022	420	2240 ³	5.53×10^8	5	2
MCMC	−21	CarmenHD	2007, 2023	1000	2240 ³	7.46×10^9	12	2

Note. The columns list (from left to right): what each simulation is used for, the absolute magnitude threshold of the corresponding SDSS sample, the name of the simulation, the seeds used, the (comoving) box size, number of particles, mass resolution, (comoving) force softening, and the number of simulations.

the decorated HOD (dHOD) model of Hearin et al. (2016). In order to apply this decorated HOD model, we first split halos by mass into bins of width 0.05 dex. Then, within each mass bin, we split halos into two groups based on the median value of the secondary property s in each bin. We then assign galaxies to halos based on the following conditional relations:

$$\langle N_{\text{cen}}|M, s_{\text{high}} \rangle = \langle N_{\text{cen}}|M \rangle + \delta N_{\text{cen}}, \quad (3)$$

$$\langle N_{\text{cen}}|M, s_{\text{low}} \rangle = \langle N_{\text{cen}}|M \rangle - \delta N_{\text{cen}}, \quad (4)$$

$$\langle N_{\text{sat}}|M, s_{\text{high}} \rangle = \langle N_{\text{sat}}|M \rangle + \delta N_{\text{sat}}, \quad (5)$$

and

$$\langle N_{\text{sat}}|M, s_{\text{low}} \rangle = \langle N_{\text{sat}}|M \rangle - \delta N_{\text{sat}}, \quad (6)$$

where

$$\delta N_{\text{cen}} = A_{\text{cen}} \text{MIN}[\langle N_{\text{cen}}|M \rangle, 1 - \langle N_{\text{cen}}|M \rangle] \quad (7)$$

for central galaxies and

$$\delta N_{\text{sat}} = A_{\text{sat}} \langle N_{\text{sat}}|M \rangle, \quad (8)$$

for satellite galaxies. A_{cen} and A_{sat} have values between -1 and 1 ; values of 0 indicate no assembly bias. A key feature of this dHOD model is that, regardless of the strength of the assembly bias, $\langle N_{\text{cen}} \rangle$ and $\langle N_{\text{sat}} \rangle$ are preserved for a given halo mass. In other words, at fixed mass, for the same five-parameter standard HOD model, the decorated HOD has the same halo occupation distribution when averaged over all halos.

Several works have explored the variety of different halo properties that can be used to model assembly bias. Salcedo et al. (2018) explored halo assembly bias in the LasDamas simulations and found that a clustering bias exists if halos are binned by mass or by any other halo property, indicating that no single halo property encompasses all the spatial clustering information of the halo population. They also found that the mean values of some halo properties depend on their halo’s distance to a more massive neighbor, and they concluded that this “neighbor bias” largely accounts for the secondary bias seen in halos binned by mass and split by concentration or age. However, they also found that halos binned by other mass-like properties still show a secondary bias even when the neighbor bias is removed.

Meanwhile, Mao et al. (2018) presented a summary of secondary halo biases of high-mass halos due to various halo properties (e.g., concentration, spin, several proxies of assembly history, and subhalo properties) in the MultiDark Planck 2 simulation. They found that, while concentration, spin, and the abundance and radial distribution of subhalos exhibit significant secondary biases, properties that directly quantify halo assembly history do not.

Finally, Behroozi et al. (2021) examined the correlation of different properties of dark matter halos (e.g., growth rate, spin, and concentration) with environment in the Small MultiDark Planck simulation, and they demonstrated that these halo properties imprint distinct signatures in the galaxy two-point correlation function and in the distribution of distances to galaxies’ k th nearest neighbors. They also demonstrated that the agreement with observed clustering can be improved with a simple empirical model in which galaxy size correlates with halo growth.

In this work, the first secondary halo property that we use to model assembly bias is halo concentration, c , defined as the ratio of the virial radius R_{vir} of the halo to the scale radius R_s (Navarro et al. 1997). The dependence of the galaxy–halo connection on concentration or circular velocity has been explored in a number of previous works (e.g., Lehmann et al. 2017; Xu & Zheng 2020). For a given halo, the concentration can be found using the relationship between virial mass, maximum circular velocity, and concentration at $z=0$:

$$v_{\text{circ}}(M_{\text{vir}}) = \frac{6.72 \times 10^{-3} M_{\text{vir}}^{1/3} \sqrt{c}}{\sqrt{\ln(1+c) - c/(1+c)}}, \quad (9)$$

where M_{vir} is the virial mass of the halo in units of $h^{-1}M_{\odot}$ and v_{circ} is the maximum circular velocity of the halo in units of km/s (Klypin et al. 2011). In our case, when implementing halo concentration as our secondary bias property, we determined that the normalization is irrelevant and it is only the halo ranking that matters; thus, we use $v_{\text{circ}}/M_{\text{vir}}^{1/3}$ as a proxy for concentration. We refer to this assembly bias model using concentration as “ABcon.”

Another halo property that can be used to model assembly bias is its larger-scale environment. The reason that conditioning the galaxy occupation on concentration has an impact on clustering statistics is that concentration is correlated with a halo’s larger-scale environment at fixed mass. Because we do not know a priori what secondary halo property to use in modeling assembly bias, it makes sense to skip this intermediate step and condition galaxy occupation directly on environment. Several works have explored the dependence of the galaxy–halo connection on environment (e.g., Pujol et al. 2017; Shi & Sheth 2018; Hadzhiyska et al. 2020; Yuan et al. 2021). Motivated by these works, we choose to also explore the effects of using local halo environment to model assembly bias. We define local environment as the total mass of halos within a $5 h^{-1}$ Mpc spheres centered on the halo of interest (excluding the mass of the halo of interest itself). We do not impose any lower mass limit on the halos included in this sum. We refer to this assembly bias model using environment as “ABenv.”

4.3. Velocity Bias

Another way in which we can extend the HOD model is to relax the assumption that satellite galaxies trace the velocities of the dark matter particles within their host halo. In other words, we can introduce satellite velocity bias (“VB”) into our model. We do this by introducing a new parameter, B_{vel} , to the model. B_{vel} is defined as the ratio between the velocities of satellite galaxies and dark matter in the halo frame of reference:

$$B_{\text{vel}} = \frac{v_g - v_h}{v_m - v_h}, \quad (10)$$

where v_g is the velocity of the satellite galaxy, v_h is the velocity of the halo, and v_m is the velocity of the randomly chosen dark matter particle on which the satellite galaxy is placed. A value of B_{vel} less than 1 indicates that satellite galaxies are moving with slower velocities than the dark matter, while a value of B_{vel} greater than 1 indicates that satellite galaxies are moving faster than the dark matter, and a value of B_{vel} equal to 1 indicates no velocity bias.

In this study, we only model satellite velocity bias and not central velocity bias. In other words, we stick with the standard assumption that central galaxies inherit the same velocity as their host halo. In principle, central galaxies can move relative to their halo, as predicted by some hydrodynamic simulations (Berlind et al. 2003) and suggested for SDSS galaxies (e.g., Guo et al. 2015a, 2015b). However, when comparing HOD modeling to hydrodynamic simulations Beltz-Mohrmann et al. (2020) found that the presence of central velocity bias is likely to have a negligible effect on the galaxy clustering statistics that we use in our analysis—unlike satellite velocity bias, which is likely important for low-luminosity galaxies.

4.4. Accounting for Baryonic Effects

While not strictly part of the HOD model, another way in which we can extend the standard halo modeling framework is to account for the effect of baryonic physics on the halo mass function. The HOD model is typically applied to a halo catalog generated from a dark matter only (DMO) simulation, which does not account for the impact of baryonic physics on halo mass. Beltz-Mohrmann & Berlind (2021) investigated the differences in halo mass functions between matched DMO and hydrodynamic simulations in EAGLE, Illustris, and IllustrisTNG, and found that, for halos at $z = 0$, stellar feedback generally reduces the masses of low-mass halos ($\lesssim 10^{11} h^{-1} M_\odot$), while AGN feedback generally reduces the masses of high-mass halos (between 10^{12} and $10^{13} h^{-1} M_\odot$) compared to their DMO counterparts. However, the exact effect that feedback has on the halo masses differs dramatically from one hydrodynamic simulation to the next. By matching halos according to mass between dark matter and hydrodynamic simulations, Beltz-Mohrmann & Berlind (2021) produced formulae that can be used to adjust the halo masses in a DMO simulation in order to reproduce the halo mass function from a given hydrodynamic simulation. Additionally, they produced fits based on matching halos between dark matter and hydrodynamic simulations based on both mass and local halo environment. By taking halo environment into account, these fits can be used to adjust halo masses in a DMO simulation to not only reproduce the global halo mass function from a hydrodynamic simulation but also to reproduce the conditional mass function, which then also reproduces the halo correlation

function. In Section 7.4, we apply several of the halo mass corrections from Beltz-Mohrmann & Berlind (2021) to our halo catalogs, in order to assess the robustness of our results to changes in the mass function due to baryonic physics. It should be noted that these mass corrections do not modify the halo profile, nor do they alter the velocity dispersion of dark matter within the halo; they adjust only the mass of each halo.

4.5. Summary

In this work, we extend the standard HOD model in several ways. We first explore the effects of extending the standard HOD model to include concentration-based assembly bias. We then explore the effects of instead using environment-based assembly bias. Next, we extend the model to include both assembly bias *and* satellite velocity bias. Finally, we implement halo mass corrections to account for the impact of baryonic physics, and we investigate the effects this has on the results from our most complete halo model (i.e., the model with both assembly bias and velocity bias).

5. Modeling Procedure

5.1. Building Mock Galaxy Catalogs

We build mock galaxy catalogs to use as our model by populating the two high-resolution simulations for each sample (ConsueloHD and CarmenHD) with galaxies. Once we populate our dark matter halos with galaxies, we build realistic mock galaxy catalogs that resemble our SDSS samples of interest. To do this, we transpose the mock galaxies from Cartesian to spherical coordinates by positioning a virtual observer at the center of the box and converting the positions of the galaxies into R.A., decl., and comoving distances. We then carve out four independent mock galaxy catalogs from each simulation box and incorporate the same systematic effects that plague our observational data set, such as redshift–space distortions, sample geometry, and incompleteness. For more information on the forward modeling details, see S22.

5.2. Covariance Matrices

If we wish to take advantage of the information present at small scales to constrain the galaxy–halo connection, it is essential that we construct accurate covariance matrices for our clustering measurements. To do this, we run 100 low-resolution simulations for each sample (Consuelo and Carmen) that differ in the phases of the density modes of the power spectrum, which is controlled by a seed supplied to 2LPTIC. We populate these low-resolution simulations with galaxies using the same HOD parameters⁷ as were used to build the matrices in S22. These parameters are listed in Table 3.

We then build 400 independent mock galaxy catalogs for each sample, from which we can construct a covariance matrix to represent cosmic variance. The elements of the covariance matrix are given by

$$C_{ij} = \frac{1}{N-1} \sum_1^N (y_i - \bar{y}_i)(y_j - \bar{y}_j), \quad (11)$$

where the sum is taken over the $N = 400$ mocks. The values y_i and y_j are the i th and j th observables measured on each mock,

⁷ The covariance matrices are built using a model that does not include assembly bias or velocity bias.

Table 3
Fiducial HOD Parameters for Covariance Matrices

M_r^{lim}	$\log M_{\text{min}}$	$\sigma_{\log M}$	$\log M_0$	$\log M_1$	α
−19	11.54	0.22	12.01	12.74	0.92
−21	12.72	0.46	7.87	13.95	1.17

Note. The HOD parameters used to construct the covariance matrices in our analysis. The matrices were constructed assuming zero assembly bias and velocity bias.

while \bar{y}_i and \bar{y}_j are the mean values of the i th and j th observables, respectively. Each diagonal element, C_{ii} , of the matrix is the variance across 400 mocks for observable i , and $\sqrt{C_{ii}}$ is the cosmic variance uncertainty of observable i . For an arbitrary observable, we refer to this uncertainty as σ_{obs} . S22 showed that the noise from using a finite number (400) of mock catalogs does not significantly affect our clustering analysis. Additionally, S22 examined the impact of resolution on the accuracy of our covariance matrices and determined that the lower resolution of the Carmen and Consuelo simulations causes us to overestimate the error on the smallest scales of the correlation function by 10%–20%. However, not only is this a small effect, but larger cosmic variance uncertainties result in lower chi-square measurements and in general make it more difficult to rule out incorrect models. We would rather have slightly broader constraints than artificially tight constraints.

5.3. Clustering Statistics

Several works have demonstrated the power of using a variety of different clustering statistics to constrain both the galaxy–halo connection as well as cosmology (Berlind & Weinberg 2002; Sinha et al. 2018; Hadzhiyska et al. 2021a; Storey-Fisher et al. 2022; Szewciw et al. 2022). In our analysis, we employ the following clustering statistics: the projected correlation function $w_p(r_p)$ (e.g., Zehavi et al. 2002, 2004; Zheng 2004; Zehavi et al. 2005; Zheng et al. 2007; Zehavi et al. 2011; Leauthaud et al. 2012; Zentner et al. 2014; Coupon et al. 2015), the redshift–space correlation function $\xi(s)$ (e.g., Tinker et al. 2006b; Parejko et al. 2013; Guo et al. 2015b; Padilla et al. 2019; Beltz-Mohrmann et al. 2020; Tonegawa et al. 2020), the group multiplicity function $n(N)$ (e.g., Berlind et al. 2006; Zheng & Weinberg 2007; Sinha et al. 2018; Beltz-Mohrmann et al. 2020), the average group velocity dispersion $\sigma_v(N)$, the mark correlation function $\text{mcf}(s)$ (e.g., Zu & Mandelbaum 2018; Storey-Fisher et al. 2022), and two special cases of counts in cells $P_N(R)$: the void probability function P_0 (VPF(R)) and the singular probability function P_0 (SPF(R)) (e.g., Tinker et al. 2006a, 2008; McCullagh et al. 2017; Walsh & Tinker 2019; Wang et al. 2019; Beltz-Mohrmann et al. 2020; Perez et al. 2021). A detailed description of each of these clustering statistics is given in S22. To calculate $w_p(r_p)$, $\xi(s)$, $\text{mcf}(s)$, VPF(R), and SPF(R), we make use of the publicly available code CORRFUNC (Sinha & Garrison 2019, 2020). In our modeling procedure, we measure each clustering statistic in the same way (i.e., either on the full box(es) or on the mock galaxy catalogs) as was done in S22. It is important to note that our clustering statistics range in scale from about 0.1 to 20 h^{-1} Mpc for both samples; thus, our analysis extends from the highly nonlinear regime all way out to the “quasi-linear” regime of clustering.

One of the main motivations for including so many higher-order statistics in our analysis is to ultimately obtain constraining power for both our model of the galaxy–halo connection and our cosmological model. For example, the redshift–space correlation function contains information about galaxy peculiar velocities due to redshift–space distortions that change the apparent positions of galaxies along the line of sight, which can help us constrain cosmic structure growth. Additionally, Storey-Fisher et al. (2022) found that statistics beyond the standard galaxy clustering statistics (e.g., $w_p(r_p)$) significantly increase the constraining power on cosmological parameters of interest. Specifically, they found that including counts in cells statistics and the mark correlation function improves the precision of constraints on σ_8 by 33%, Ω_m by 28%, and the growth of structure parameter, $f\sigma_8$, by 18%, compared to standard statistics. While we do not vary our cosmological model in this work, and thus cannot comment on the specific ability of each clustering measurement to constrain cosmological parameters, we include such a wide variety of clustering statistics in this work with the goal of ultimately constraining both HOD and cosmological parameters.

5.4. MCMC

We explore the HOD parameter space with a Markov Chain Monte Carlo (MCMC) algorithm, using a privately developed C-implementation of the popular affine-invariant sampler EMCEE (Foreman-Mackey et al. 2013), which we call EMCEE_IN_C.⁸ We impose flat priors on the same parameter ranges given in S18, as well as flat priors of [−1.0, 1.0] on A_{cen} and A_{sat} , and a flat prior of [0.5, 1.5] on B_{vel} .

At each point in the chain, we evaluate the likelihood that a particular HOD model could have generated a data set with the same clustering as SDSS. This likelihood is given by

$$\mathcal{L}(\mathbf{D}|\mathbf{M}) = \frac{\exp\left(-\frac{1}{2}(\mathbf{D} - \mathbf{M})\mathbf{C}^{-1}(\mathbf{D} - \mathbf{M})^T\right)}{\sqrt{(2\pi)^K \det(\mathbf{C})}}, \quad (12)$$

where \mathbf{D} is the K -dimensional vector of observables measured on the SDSS data set, \mathbf{M} is the corresponding vector of observables measured on the HOD model, and \mathbf{C} is the K -dimensional covariance matrix of these observables representing cosmic variance (see Equation (11)). (The term within the exponential is essentially χ^2 , multiplied by a factor of $-1/2$).

It is worth noting that our likelihood calculation assumes that all of our observables are Gaussian. However, Hahn et al. (2019) found that assuming a Gaussian likelihood in a group multiplicity function analysis slightly underestimates the uncertainties and biases HOD parameter constraints by $\sim 0.5\sigma$. We have examined all of our observables and determined that the vast majority of them (including the group multiplicity function) appear to follow Gaussian distributions and pass typical tests of Gaussianity, so we have proceeded with assuming a Gaussian likelihood.

In the HOD framework, the process of populating halos with galaxies is stochastic, and is controlled with a “population seed.” For a fixed HOD model, changes in this population seed can lead to significant differences in clustering statistics. To minimize the noise in our results due to this random variation,

⁸ https://github.com/aszewciw/emcee_in_c

at each point in the chain, we populate halos four times, using four fixed population seeds. Thus, the clustering measurements for a given point in HOD parameter space are the average measurements over these four population seeds. (See S22 for details).

6. Choosing Optimal Observables

In order to constrain the dHOD when fit to SDSS, we must first choose a set of observables to use in our MCMC. We cannot arbitrarily continue to increase the number of observables we use, because doing so increases the noise in our modeling and degrades our final constraints. Noise is introduced into the covariance matrix due to the fact that we are constructing it from only 400 mocks. This noise propagates into the likelihood function and ultimately into our posterior results. Thus, we need to choose our observables wisely. We seek a subset of observables that produce the tightest constraints on our HOD parameters, at the cost of little noise.

To choose an “optimal” set of high-information, low-noise observables, we employ the importance sampling algorithm described in S22. In this algorithm, we first create four mock SDSS catalogs for which we will determine optimal statistics. We use four mocks instead of the actual SDSS data, in order to minimize the impact of cosmic variance on the selection of optimal statistics. We build these four mocks using a fiducial dHOD (with concentration) model with parameters that we obtain by fitting this model to the S22 set of clustering statistics (listed in Table 4 under “sHOD”) measured on the SDSS.

We run an initial MCMC on each of the four mock galaxy catalogs, fitting the dHOD (with concentration) model to only two observables: n_{gal} and $w_p(r_p \sim 0.3 h^{-1} \text{Mpc})$. This results in a fairly broad MCMC nonuniform grid of points in parameter space for each mock. We then use importance sampling on these grids to explore the constraining power of different combinations of clustering statistics. The algorithm chooses observables one by one, each time selecting the observable that, when combined with all previously chosen observables, produces the tightest projected constraints on all HOD parameters of interest. When choosing an observable, we consider how it performs on across all four grids, minimizing any bias due to cosmic variance. Thus, at the end of running this algorithm, we have a list of observables (ordered in terms of cumulative constraining power) and a corresponding list of cumulative projected constraints for each sample. (We refer the reader to S22 for a more complete description of this procedure).

There are two key differences in our implementation of this algorithm compared to S22. First, when choosing the third observable for each sample, we only attempt to constrain A_{cen} and A_{sat} . This is because these parameters are entirely unconstrained when using only n_{gal} and w_p ($r_p \sim 0.3 h^{-1} \text{Mpc}$), which causes the MCMC to explore unrealistic HOD models; thus, it is essential to choose an observable early on that provides constraining power for these parameters. After the third observable is chosen, we make all successive choices by attempting to jointly constrain all HOD parameters (excluding $\log M_0$ for the -21 sample). Second, in the S22 algorithm, new grids are created (by running new MCMCs using the already chosen observables) whenever the old grids become insufficiently dense for importance sampling. S22 creates these new grids after choosing five observables for each sample, and again for the -19 sample after choosing eight

Table 4
Optimal Observable Order

Index	-19 sHOD	-19 dHOD	-21 sHOD	-21 dHOD
1	n_{gal}	n_{gal}	n_{gal}	n_{gal}
2	$w_p(r_p)$ 2	$w_p(r_p)$ 2	$w_p(r_p)$ 2	$w_p(r_p)$ 2
3	$w_p(r_p)$ 4	$\sigma_v(N)$ 3	$\xi(s)$ 8	$\sigma_v(N)$ 1
4	VPF(R) 3	$\xi(s)$ 8	$w_p(r_p)$ 4	$\xi(s)$ 9
5	$w_p(r_p)$ 8	$n(N)$ 3	mcf(s) 9	$\xi(s)$ 3
6	$\xi(s)$ 1	SPF(R) 1	$w_p(r_p)$ 1	mcf(s) 10
7	$n(N)$ 3	$w_p(r_p)$ 3	$\xi(s)$ 9	$w_p(r_p)$ 5
8	$\xi(s)$ 5	$n(N)$ 2	mcf(s) 7	$n(N)$ 1
9	$n(N)$ 2	$w_p(r_p)$ 8	$\xi(s)$ 4	$\sigma_v(N)$ 3
10	$n(N)$ 4	$\xi(s)$ 1	$\xi(s)$ 7	mcf(s) 3
11	$n(N)$ 1	$w_p(r_p)$ 4	mcf(s) 10	$\xi(s)$ 1
12	SPF(R) 4	VPF(R) 2	$\xi(s)$ 1	$\xi(s)$ 8
13	$\xi(s)$ 13	mcf(s) 1	$w_p(r_p)$ 14	$\xi(s)$ 5
14	mcf(s) 14	$\xi(s)$ 10	$n(N)$ 1	$w_p(r_p)$ 1
15	$\xi(s)$ 6	SPF(R) 2	SPF(R) 4	$n(N)$ 2
16	$n(N)$ 5	$\xi(s)$ 4	mcf(s) 3	SPF(R) 4
17	$\xi(s)$ 2	$n(N)$ 1	$\xi(s)$ 6	mcf(s) 5
18	SPF(R) 2	$n(N)$ 5	$\sigma_v(N)$ 4	$\sigma_v(N)$ 4
19	$\xi(s)$ 10	$w_p(r_p)$ 1	$\xi(s)$ 5	mcf(s) 14
20	mcf(s) 2	SPF(R) 4	$\xi(s)$ 3	SPF(R) 3
21	mcf(s) 3	mcf(s) 7	$n(N)$ 4	$w_p(r_p)$ 3
22	$\sigma_v(N)$ 1	mcf(s) 11	$w_p(r_p)$ 7	$\sigma_v(N)$ 5
23	$\sigma_v(N)$ 3	$\sigma_v(N)$ 5	$w_p(r_p)$ 3	$\sigma_v(N)$ 2
24	$\xi(s)$ 9	SPF(R) 3	mcf(s) 8	$\xi(s)$ 7
25	$\sigma_v(N)$ 4	$\xi(s)$ 3	VPF(R) 3	$n(N)$ 3
26	mcf(s) 1	$n(N)$ 4	$\xi(s)$ 2	$n(N)$ 4
27	$\sigma_v(N)$ 2	mcf(s) 2	$n(N)$ 5	$\xi(s)$ 4
28	$n(N)$ 6	$\sigma_v(N)$ 2	$n(N)$ 2	$\xi(s)$ 2
29	VPF(R) 1	VPF(R) 4	$\xi(s)$ 11	$n(N)$ 5
30	$w_p(r_p)$ 1	mcf(s) 8	$\sigma_v(N)$ 3	$w_p(r_p)$ 8
31	$w_p(r_p)$ 6	$w_p(r_p)$ 6	$\sigma_v(N)$ 2	$w_p(r_p)$ 4
32	$w_p(r_p)$ 5	$\xi(s)$ 9	SPF(R) 2	$\xi(s)$ 6
33	$\sigma_v(N)$ 5	$n(N)$ 6	mcf(s) 5	mcf(s) 8
34	$w_p(r_p)$ 3	mcf(s) 14	SPF(R) 3	$\xi(s)$ 10
35	$\sigma_v(N)$ 7	VPF(R) 5	mcf(s) 4	$\xi(s)$ 11
36	$n(N)$ 7	mcf(s) 12	SPF(R) 1	mcf(s) 7
37	mcf(s) 12
38	$\xi(s)$ 14
39	VPF(R) 5
40	VPF(R) 2
41	mcf(s) 1

Notes. The type of clustering statistic and the bin number (1-indexing) for the observables chosen (in order) for each sample. “sHOD” refers to the observables chosen for each sample using the standard HOD model in S22. “dHOD” refers to the observables chosen in this work using the decorated HOD model. The observables chosen in this work that were *not* chosen in S22 are shown in bold.

observables. In our case, because we are trying to constrain two additional parameters, our grids become insufficiently dense more quickly, and so we ultimately build denser grids after choosing 3, 5, 10, and 20 observables for each sample.

In Figure 1, we show our estimated constraint for each HOD parameter (excluding $\log M_0$) as we choose successive observables. The results for the -19 sample are shown in blue, and the results for the -21 sample are shown in red. The solid lines show the average constraint across the four mocks used in the algorithm described above. In Table 4, we list the observables chosen (in order) that we use for each sample (labeled “dHOD”). We also list the observables chosen in S22 using a standard HOD model (i.e., no assembly bias, labeled




Figure 1. Constraints on each HOD parameter as we increase the number of observables, for the -19 sample (blue) and the -21 sample (red). The solid line in each panel shows the average mock constraint (1σ) across four mocks, and the shaded region is an estimate of the uncertainty (inner 68%) in these constraints due to the noise present in the covariance matrix. The dot indicates the optimal number of observables for each sample, and the dashed line indicates the corresponding constraining power for each parameter.

“sHOD”). The observables chosen in this work that were *not* chosen in S22 are shown in bold.

After ordering the observables from greatest to least constraining power, we need to choose the total number of observables to use in our analysis to maximize our constraining power and minimize the noise in our procedure due to building our covariance matrix from a finite number of mocks (i.e., 400). To do this, we employ the same procedure as S22. Briefly, we estimate an uncertainty associated with each projected constraint (for a given number of observables K) by resampling the covariance matrix 100 times and then importance sampling the chain with each of these resampled matrices. Doing so lets us approximate the uncertainty in our constraints due to noise in the covariance matrix for each combination of observables that we consider. The shaded regions in each panel of Figure 1 show this uncertainty for each HOD parameter as we increase K . We choose the lowest value of K such that the constraint at this value is within one standard error of the constraint at all higher values of K . We require that this condition is met for all HOD parameters (except $\log M_0$ for the -21 sample). The optimal number of observables for each sample is indicated with a dot in each panel, and the corresponding constraining power is shown with a dashed line. For the -19 sample, the optimal number of observables is 36. For the -21 sample, the optimal number of observables is 41. Thus, the size of our data vector for each sample is 36 and 41, respectively.

Using these observables, we confirm that we can recover the truth when running chains on mocks created with different HOD parameters (i.e., different amounts of assembly bias) for each sample. In all of our validation tests with mocks, all

parameters for the -21 sample are always recovered within 1σ , and in the -19 sample all parameters are always recovered within or just outside the 1σ region. Additionally, for both samples, the best-fit result is always a good fit (i.e., it always has a p -value greater than 0.85).

Looking at the optimal observables in Table 4, it is noteworthy that, for both the -19 and -21 samples, the third observable (chosen to constrain only A_{cen} and A_{sat}) is a small bin of the average group velocity dispersion function ($\sigma_v(N)$ 3 for -19 and $\sigma_v(N)$ 1 for -21). It is also noteworthy that, for both samples, the majority of the first 20 observables chosen in this analysis (16/20 or 17/20) were also chosen in S22 to constrain an HOD model without assembly bias. Meanwhile, about half of the observables chosen beyond the initial twenty (8/16 or 9/21) in this analysis are unique to the model with assembly bias (i.e., they were not chosen in S22). This possibly indicates that the initial observables are chosen for their ability to constrain the standard HOD parameters, while the later observables are selected for their ability to constrain the assembly bias parameters. This may also indicate that it is difficult to constrain assembly bias until the standard HOD parameters are constrained, or that assembly bias is a smaller signal on top of the global clustering signal.

For the -19 sample, the unique observables chosen for this analysis include a large and small scale of $\xi(s)$, five scales of $P_N(R)$, and four large scales of $\text{mcf}(s)$. For the -21 sample, the unique observables chosen for this analysis include two bins of $\sigma_v(N)$, two intermediate scales of $w_p(r_p)$, one small scale and two large scales of $\text{mcf}(s)$, one intermediate bin of $n(N)$, two large scales of $\xi(s)$, and two bins of $\text{VPF}(R)$. It is worth

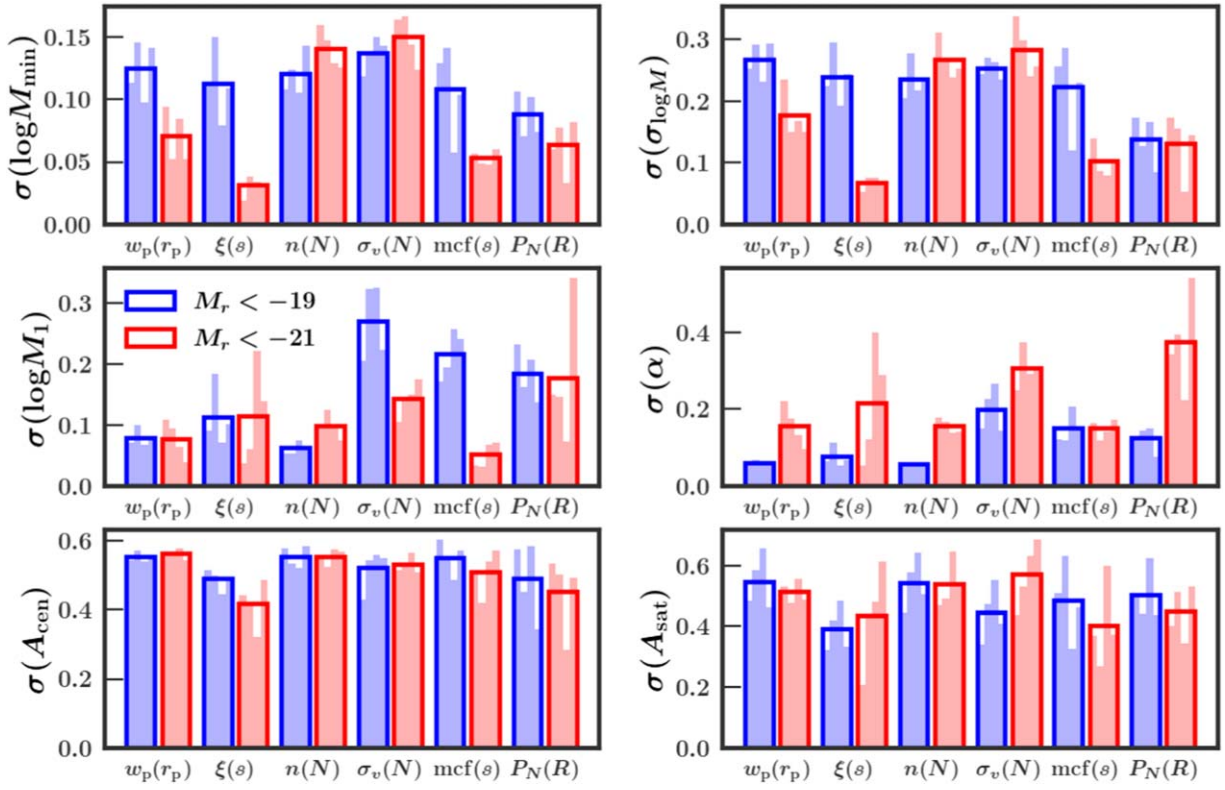


Figure 2. Projected constraints (1σ) of each clustering statistic (combined with n_{gal}) for each HOD parameter. The constraints for the -19 and -21 mocks are shown in blue and red, respectively. The height of each smaller vertical bar shows the projected constraints on one mock, while the larger open bar shows the average constraint across four mocks.

mentioning that, for the -19 sample, it is difficult to accurately constrain the decorated HOD model until the parameter $\log M_0$ is constrained. This occurs by about 15 observables, particularly after $\xi(s)$ 1 and $w_p(r_p)$ 4 are included. In the -21 sample, the parameter $\log M_0$ remains unconstrained. This is consistent with the results of S22, who found that constraining $\log M_0$ is important for obtaining accurate results in the -19 sample but not in the -21 sample.

Given the MCMC grids that we obtained from the first three observables in our optimal selection algorithm, we can use importance sampling to estimate the constraining power we would achieve for each HOD parameter had we run a chain using only one clustering statistic (e.g., $w_p(r_p)$) plus n_{gal} . We display the results of this exercise in Figure 2. In each panel, the y-axis shows the projected constraint (1σ) for a particular HOD parameter as we use different clustering statistics. The constraints for the -19 and -21 mocks are shown in blue and red, respectively. The height of each smaller vertical bar shows the projected constraints on one mock, while the larger open bar shows the average constraint across four mocks.

For the central and satellite parameters, our results are similar (though not identical) to the results from S22. For the assembly bias parameters, it is interesting to note that, for both samples, no single clustering statistic provides significant constraining power for either A_{cen} or A_{sat} . $\xi(s)$ seems to have the most constraining power for both A_{cen} and A_{sat} in both samples, but it performs only slightly better than the other clustering statistics.

Due to the nature of importance sampling, these results should be interpreted as estimates, purely for visual purposes. However, this figure illustrates that, while no single clustering

statistic provides significant constraining power for assembly bias, the combination of different scales of each clustering statistic is able to produce tighter constraints on the assembly bias parameters than any one statistic.

7. Results

7.1. Concentration-based Assembly Bias

Here, we present the results from using the optimal observables identified in the previous section to constrain the galaxy-halo connection in SDSS using a decorated HOD model with concentration-based assembly bias. The results for the -19 sample are shown in Figure 3, while the results for the -21 sample are shown in Figure 4. Dark and light regions depict the 1σ and 2σ regions, respectively. The best-fit parameters are listed in Table 5, along with their corresponding p -values, as well as the results from the previous S22 analysis for comparison. The constraints for each parameter are listed in Table 6.

For the -19 sample, our best-fit results suggest positive central galaxy assembly bias ($A_{\text{cen}} = 0.793$) and negative satellite galaxy assembly bias ($A_{\text{sat}} = -0.368$). In other words, central galaxies preferentially reside in halos with higher concentrations, while satellite galaxies preferentially reside in halos with lower concentrations, at fixed mass. This is consistent with previous results (e.g., Lange et al. 2022; Wang et al. 2022), which also found positive central galaxy assembly bias and negative satellite galaxy assembly bias when using concentration as the secondary halo property. Additionally, this best-fit model yields a significant decrease in tension compared to the results of S22 (2.0σ compared to 4.5σ). Unfortunately,




Figure 3. Parameter constraints for the SDSS -19 sample, using concentration as the secondary halo property and the “dHOD” optimal observables (listed in Table 4). The crosshairs in the third panel indicate $A_{\text{cen}} = A_{\text{sat}} = 0$ (i.e., no assembly bias).




Figure 4. Parameter constraints for the SDSS -21 sample, using concentration as the secondary halo property and the 41 “dHOD” optimal observables (listed in Table 4). The crosshairs in the third panel indicate $A_{\text{cen}} = A_{\text{sat}} = 0$ (i.e., no assembly bias).

Table 5
SDSS Best-fit Results for Different Halo Models

M_r^{lim}	Model	Obs.	$\log M_{\text{min}}$	$\sigma_{\log M}$	$\log M_0$	$\log M_1$	α	A_{cen}	A_{sat}	B_{vel}	p -value	AIC
-19	Standard HOD	S22	11.445	0.099	11.651	12.703	0.958	$6.8 \cdot 10^{-6}$	87.77
	ABcon	BM22	11.455	0.141	11.757	12.685	0.925	0.793	-0.368	...	0.047	56.83
	ABcon + VB	BM22	11.474	0.132	11.877	12.715	0.950	0.825	-0.251	0.898	0.155	51.54
	ABenv + VB	BM22	11.490	0.125	11.855	12.783	0.985	0.533	-0.224	0.826	0.364	45.99
-21	Standard HOD	S22	12.728	0.467	9.015	13.929	1.112	$3.5 \cdot 10^{-5}$	
	ABcon	BM22	12.774	0.554	9.447	13.926	1.067	-0.090	-0.240	...	$2.6 \cdot 10^{-6}$	99.38
	ABcon + VB	BM22	12.756	0.525	9.804	13.915	1.108	0.144	-0.198	0.976	$1.8 \cdot 10^{-6}$	100.96
	ABenv + VB	BM22	12.740	0.495	9.984	13.917	1.079	-0.025	0.165	1.011	$4.1 \cdot 10^{-6}$	98.43

Notes. Best-fit HOD parameters for each SDSS sample using four different models: the standard five-parameter model, a model with concentration-based assembly bias (“ABcon”), a model with concentration-based assembly bias plus satellite velocity bias (“ABcon + VB”), and a model with environment-based assembly bias plus satellite velocity bias (“ABenv + VB”). The Standard HOD results are taken from S22 and thus use the S22 observables, while the chains using extended HOD models use the optimal observables identified in this work (listed in Table 4). We indicate the goodness of fit of each parameter combination with a p -value, in addition to assessing the success of the model using the AIC.

even for our optimal combination of observables, it is difficult to tightly constrain central galaxy assembly bias for this sample (see the third panel in Figure 3). Despite the lack of tight constraints on A_{cen} , we are able to rule out a model with zero assembly bias (i.e., the standard HOD model).

For the -21 sample, we obtain slightly tighter constraints on A_{cen} than we are able to achieve in the -19 sample. However,

our best-fit results are consistent with zero assembly bias. Additionally, this model does not result in any decrease in tension compared to the results from S22.⁹ This finding is

⁹ In fact, the tension actually increased slightly compared to the previous analysis, from 4.1σ up to 4.7σ . This slight increase in tension can be attributed to the change in observables between this work and the previous work.

Table 6
 SDSS Constraints

M_r^{lim}	Model	$\log M_{\text{min}}$	$\sigma_{\log M}$	$\log M_0$	$\log M_1$	α	A_{cen}	A_{sat}	B_{vel}
-19	sHOD	$11.442^{+0.016}_{-0.015}$	$0.106^{+0.074}_{-0.065}$	$11.674^{+0.089}_{-0.094}$	$12.691^{+0.028}_{-0.029}$	$0.954^{+0.019}_{-0.019}$
	ABcon	$11.469^{+0.019}_{-0.017}$	$0.159^{+0.074}_{-0.077}$	$11.750^{+0.093}_{-0.095}$	$12.685^{+0.025}_{-0.031}$	$0.930^{+0.025}_{-0.028}$	$0.673^{+0.245}_{-0.529}$	$-0.361^{+0.107}_{-0.103}$...
	ABcon+VB	$11.484^{+0.022}_{-0.021}$	$0.121^{+0.072}_{-0.067}$	$11.864^{+0.100}_{-0.099}$	$12.724^{+0.037}_{-0.037}$	$0.947^{+0.024}_{-0.028}$	$0.338^{+0.472}_{-0.720}$	$-0.194^{+0.145}_{-0.134}$	$0.876^{+0.042}_{-0.041}$
	ABenv+VB	$11.481^{+0.021}_{-0.021}$	$0.132^{+0.073}_{-0.052}$	$11.786^{+0.119}_{-0.140}$	$12.776^{+0.041}_{-0.040}$	$0.986^{+0.027}_{-0.027}$	$0.444^{+0.265}_{-0.190}$	$-0.173^{+0.108}_{-0.115}$	$0.857^{+0.041}_{-0.042}$
-21	sHOD	$12.748^{+0.015}_{-0.015}$	$0.517^{+0.029}_{-0.029}$	$9.015^{+2.017}_{-2.036}$	$13.919^{+0.014}_{-0.014}$	$1.088^{+0.031}_{-0.033}$
	ABcon	$12.737^{+0.019}_{-0.020}$	$0.494^{+0.038}_{-0.040}$	$8.980^{+2.019}_{-2.013}$	$13.914^{+0.015}_{-0.015}$	$1.110^{+0.035}_{-0.039}$	$-0.236^{+0.290}_{-0.297}$	$-0.148^{+0.181}_{-0.156}$...
	ABcon+VB	$12.747^{+0.020}_{-0.020}$	$0.505^{+0.038}_{-0.040}$	$9.525^{+1.676}_{-1.879}$	$13.923^{+0.018}_{-0.017}$	$1.111^{+0.043}_{-0.050}$	$0.154^{+0.305}_{-0.283}$	$-0.101^{+0.215}_{-0.208}$	$0.971^{+0.047}_{-0.041}$
	ABenv+VB	$12.773^{+0.039}_{-0.043}$	$0.546^{+0.067}_{-0.081}$	$9.984^{+1.161}_{-1.456}$	$13.925^{+0.022}_{-0.022}$	$1.059^{+0.052}_{-0.062}$	$0.014^{+0.038}_{-0.045}$	$0.115^{+0.108}_{-0.118}$	$1.000^{+0.050}_{-0.044}$

Notes. Marginalized constraints on SDSS for both samples using four different models: the standard HOD model from S22 (using the optimal observables from S22), an HOD model with concentration-based assembly bias, a model with concentration-based assembly bias plus satellite velocity bias, and a model with environment-based assembly bias plus satellite velocity bias, using the optimal observables identified in this work. We present the median parameter values along with upper and lower limits corresponding to the 84th and 16th percentiles respectively. All of these chains were run using the optimal observables identified in this work (listed in Table 4).




Figure 5. Residuals between the best-fit model and the SDSS measurements for the -19 (top) and -21 (bottom) samples. For each sample, the model includes concentration-based assembly bias. We show residuals for all observables, but the model was constrained using the “dHOD” optimal observables for each sample (listed in Table 4), which are displayed with larger points.

consistent with the results of Beltz-Mohrmann et al. (2020), which found assembly bias to be present in hydrodynamic simulations for lower-luminosity galaxies but not significant for higher-luminosity galaxies. It is thus to be expected that, for the -21 sample, the addition of assembly bias parameters to the model did not result in any relief of tension. Furthermore, the constraints on the standard HOD parameters in the -21 sample do not change considerably compared to what they were in S22, indicating that the addition of assembly bias has very little affect on the outcome of the model.

In Figure 5, we show the deviation between each observable as measured on SDSS (D) and on our best-fit model (M) for each sample. This deviation is shown as a factor of the cosmic variance uncertainty, σ_{obs} , for each observable. This quantity is

shown for all observables, where each point is a different scale or bin of a given clustering statistic. Each clustering statistic is plotted in a different color and is labeled on the x -axis. The specific observables actually used in our analysis are plotted as larger bold points. The results for the -19 sample are shown in the top panel, and the results for the -21 panel are shown in the bottom panel. In the -19 sample, much of the remaining tension seems to be coming from several scales of $\xi(s)$, and to a lesser extent, n_{gal} , $n(N)$, $\text{VPF}(R)$, and $\text{SPF}(R)$. In the -21 sample, most of the clustering statistics exhibit a high degree of tension at various scales. Overall, the -19 sample exhibits a greater improvement in the observable residuals compared to the S22 results than the -21 sample does, which explains the greater overall reduction in tension seen in this sample.

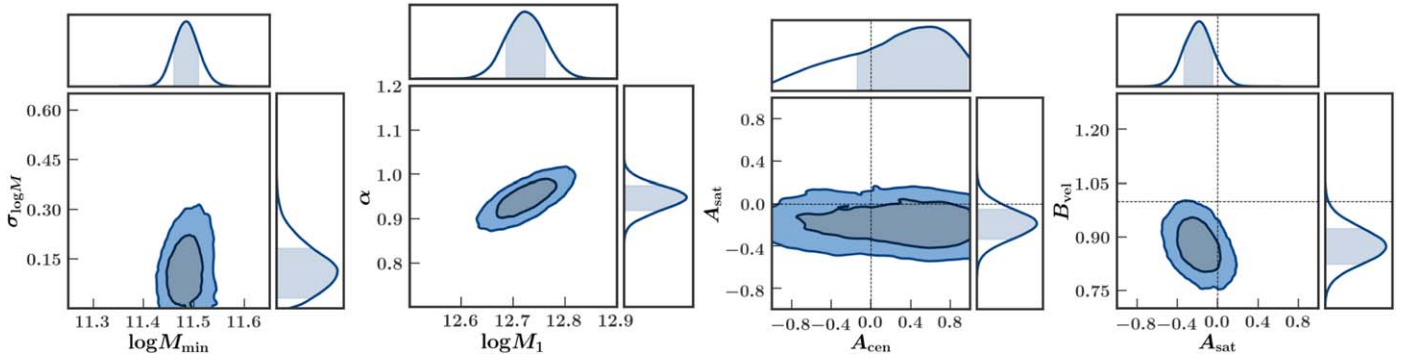


Figure 6. Parameter constraints for the SDSS –19 sample, from a model with assembly bias (with concentration as the secondary halo property) and satellite velocity bias, using the “dHOD” optimal observables (listed in Table 4). The crosshairs in the third and fourth panels indicate no assembly bias and no velocity bias ($A_{\text{cen}} = A_{\text{sat}} = B_{\text{vel}} = 0$).

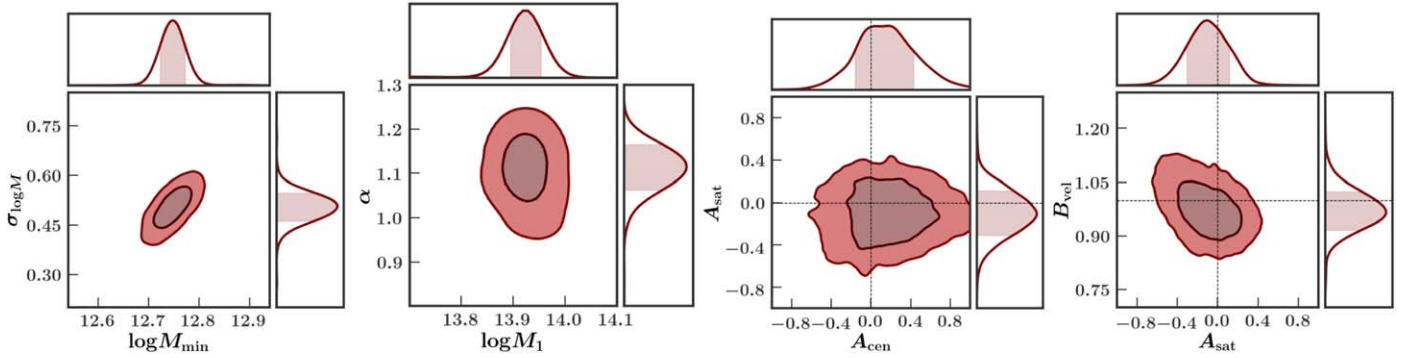


Figure 7. Parameter constraints for the SDSS –21 sample, from a model with assembly bias (with concentration as the secondary halo property) and satellite velocity bias, using the “dHOD” optimal observables (listed in Table 4). The crosshairs in the third and fourth panels indicate no assembly bias and no velocity bias ($A_{\text{cen}} = A_{\text{sat}} = B_{\text{vel}} = 0$).

The remaining tension found for both the –19 and –21 samples could indicate that the HOD model needs to be made even more flexible with the inclusion of spatial and velocity bias parameters (e.g., Beltz-Mohrmann et al. 2020). Additionally, it is possible that a different secondary halo property other than concentration could be more strongly correlated with galaxy occupation and is thus a more appropriate choice for our assembly bias model. It is also possible that accounting for the impact of baryonic physics on the halo mass function could relieve some of the remaining tension. Finally, these results are for a fixed cosmology sample; it is possible that a slight change in cosmological parameter values could also result in a further relief of this tension. We explore some of these possibilities in the remaining sections.

7.2. Satellite Velocity Bias

Here, we investigate whether adding satellite velocity bias to our HOD model results in better agreement with SDSS. Using the same set of optimal observables for each sample listed in Table 4, we run chains on our SDSS samples using an HOD model with both concentration-based central and satellite galaxy assembly bias (i.e., A_{cen} and A_{sat}) and additionally with satellite galaxy velocity bias (B_{vel}). The results for the –19 sample are shown in Figure 6, while the results for the –21 sample are shown in Figure 7. The best-fit parameter values and constraints are listed in Tables 5 and 6. Additionally, in Figure 8 we show the deviation between each observable as measured on SDSS (D) and on our best-fit model (M) for each sample, with the same layout as in Figure 5.

We do not identify a new set of optimal observables for constraining this new model, but rather run chains for each of our SDSS samples using the same optimal observables listed in Table 4 (“dHOD”). While choosing a new set of optimal observables could potentially lead to tighter constraints on B_{vel} , we emphasize that our goal is not to optimally constrain satellite velocity bias, but rather to allow B_{vel} to vary, with the hope of alleviating the lingering tension with SDSS. Additionally, we note that our optimal set of observables already includes many measurements that are sensitive to galaxy velocities (e.g., $\xi(s)$, $n(N)$, $\sigma_v(N)$) and so should contain constraining power for B_{vel} .

For the –19 sample, our best-fit results for this model indicate moderate satellite galaxy velocity bias, with satellite galaxies moving slightly slower than the dark matter distribution ($B_{\text{vel}} = 0.898$). However, when velocity bias is included in the model, the strength of the assembly bias signal is significantly reduced. This is in part due to the anticorrelation between B_{vel} and (concentration-based) A_{sat} , which can be seen in the fourth panel of Figure 6: when $B_{\text{vel}} = 1$, lower values of A_{sat} are preferred by the model, but when B_{vel} is allowed to be less than 1, A_{sat} increases. While the best-fit parameter values still suggest positive central assembly bias and negative satellite assembly bias ($A_{\text{cen}} = 0.825$ and $A_{\text{sat}} = -0.251$), the constraints on A_{cen} and A_{sat} are much weaker, and we can no longer rule out a model with zero assembly bias (though we can rule out a model with zero velocity bias).

The fact that adding velocity bias to our model reduces the strength of the assembly bias signature is strong evidence in favor of having a sufficiently flexible HOD model, without




Figure 8. Residuals between the best-fit model and the SDSS measurements for the -19 (top) and -21 (bottom) samples. For each sample, the model includes assembly bias (with concentration as the secondary halo property) as well as velocity bias. We show residuals for all observables, but the model was constrained using the “dHOD” optimal observables for each sample (listed in Table 4), which are displayed with larger points.

which we cannot claim to have made a robust detection of assembly bias, nor can we hope to reliably constrain cosmology. It is likely that our analysis is sensitive to B_{vel} because of our use of clustering measurements that are particularly affected by galaxy velocities (e.g., $\xi(s)$). It is possible that an analysis that does not include these statistics would have less constraining power for B_{vel} and therefore might not affirm the presence of velocity bias.

After including B_{vel} as well as A_{cen} and A_{sat} in our HOD model, the best-fit result for the -19 sample exhibits a substantial decrease in tension, down from 2.0σ to 1.4σ . The $\chi^2/\text{d.o.f.}$ also decreased, from 1.53 to 1.27. Thus, when we include both assembly bias and velocity bias in our model, our clustering results are in agreement with SDSS. Looking at the top panel of Figure 8 and comparing it to Figure 5, we can see that this relief in tension comes from the improvement in observables across the board, but particularly in n_{gal} , $w_p(r_p)$,¹⁰ $\xi(s)$, and $n(N)$.

For the -21 sample, our best-fit results for this model yield minimal satellite galaxy velocity bias ($B_{\text{vel}} = 0.976$), as well as minimal central and satellite galaxy assembly bias ($A_{\text{cen}} = 0.144$ and $A_{\text{sat}} = -0.198$). For this sample, the constraints on A_{cen} and A_{sat} do not significantly degrade after adding velocity bias to the model, and the constraints on B_{vel} are comparable to what they are in the -19 sample. Despite this, we cannot claim a significant detection of assembly or satellite galaxy velocity bias in the SDSS -21 sample. Once again, the constraints on the standard HOD parameters remain roughly the same, suggesting that the further addition of velocity bias to the model has negligible impact. Additionally,

¹⁰ We note that the improvement in n_{gal} and $w_p(r_p)$ comes not from their relationship with satellite velocity, but rather from the changes in other HOD parameters as a result of including velocity bias in the model.

we are still unable to relieve the tension present in the -21 sample even after adding this new flexibility to the HOD model.

Looking at the lower panel of Figure 8, we see very little improvement in our residuals compared to Figure 5, which illustrates why the addition of velocity bias to the model results in no relief in tension for this sample. It is particularly noteworthy that even the statistics that contain velocity information (like $\xi(s)$ and $\sigma_v(N)$) do not show any substantial improvement after adding velocity bias to the model. While it is possible that a different velocity bias prescription could lead to more improvement, it is also possible that these dynamical clustering measurements are sensitive to an issue that exists not within our HOD model but within our cosmological model.

7.3. Environment-based Assembly Bias

The previous section shows that central galaxy occupation is, at most, only loosely tied to the concentration of the host halo. We next investigate whether modeling assembly bias with a different halo property can improve the goodness of fit of our model. Given that we do not know which halo property other than mass most strongly affects the presence of a central galaxy in a given halo, we choose to use local halo environment as our new assembly bias property. This allows us to investigate the general assumption that central galaxy occupation is tied to some halo property that is correlated with the local environment. For this reason, environment is a useful property for probing assembly bias and gives our model the flexibility that it needs to model galaxy clustering without knowing the “true” halo property that leads to assembly bias. We define local environment as the total mass in halos within a $5 h^{-1}$ Mpc radius.

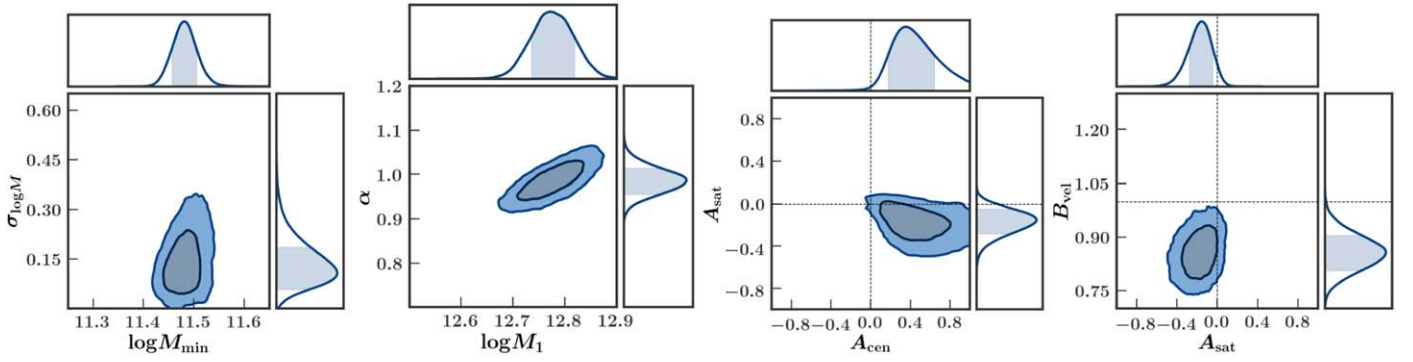


Figure 9. Parameter constraints for the SDSS -19 sample, from a model with assembly bias (with environment as the secondary halo property) and satellite velocity bias, using the “dHOD” optimal observables (listed in Table 4). The crosshairs in the third and fourth panels indicate no assembly bias and no velocity bias ($A_{\text{cen}} = A_{\text{sat}} = B_{\text{vel}} = 0$).

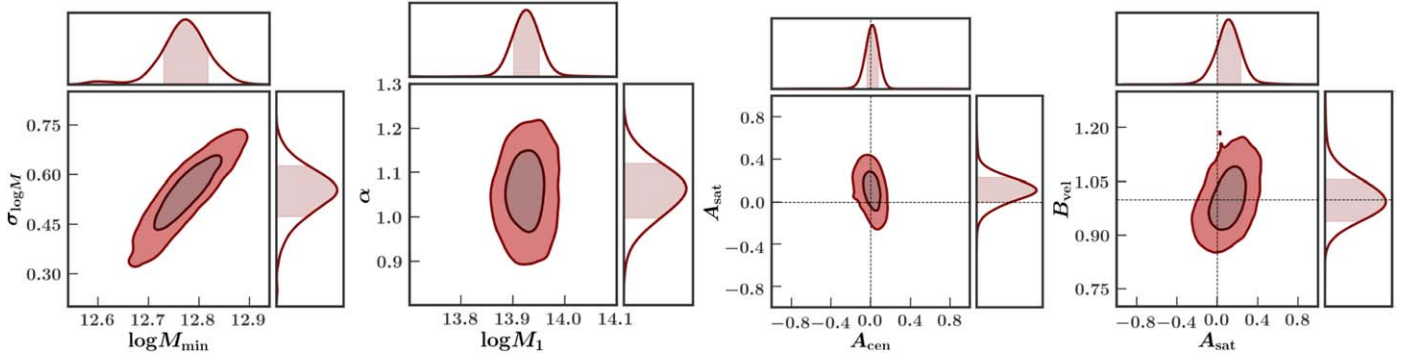


Figure 10. Parameter constraints for the SDSS -21 sample, from a model with assembly bias (with environment as the secondary halo property) and satellite velocity bias, using the “dHOD” optimal observables (listed in Table 4). The crosshairs in the third and fourth panels indicate no assembly bias and no velocity bias ($A_{\text{cen}} = A_{\text{sat}} = B_{\text{vel}} = 0$).

Once again, we do not identify a new set of optimal observables for constraining this environment-based assembly bias model, but rather run chains for each of our SDSS samples using the same optimal observables listed in Table 4 (“dHOD”). The results for the -19 sample are shown in Figure 9, while the results for the -21 sample are shown in Figure 10. The best-fit parameter values and constraints are listed in Tables 5 and 6.

For the -19 sample, our best-fit results once again indicate positive central galaxy assembly bias ($A_{\text{cen}} = 0.533$) and negative satellite galaxy assembly bias ($A_{\text{sat}} = -0.224$). In other words, low-luminosity central galaxies preferentially reside in halos with denser environments, while satellite galaxies preferentially reside in halos with less dense environments, at fixed mass. We can rule out a model with no central galaxy assembly bias at the 99% confidence level and a model with no satellite galaxy assembly bias at the 95% confidence level. We again find that satellite galaxies have velocities that are slightly slower than the dark matter distribution ($B_{\text{vel}} = 0.826$). We can rule out a model with no satellite velocity bias at the 99.9% confidence level. It is noteworthy that, when environment is used to model assembly bias instead of concentration, B_{vel} and A_{sat} are correlated rather than anticorrelated.

Additionally, the constraints on A_{cen} are much improved compared to the constraints when using concentration, and the constraints on A_{sat} are slightly improved. This improvement in constraints is seen despite the fact that the observables used were chosen to optimally constrain a concentration-based assembly bias model. We attribute this improvement to the fact

that environment is more directly associated with clustering than concentration.

Ultimately, this model is an even better fit to the data than the model with concentration-based assembly bias plus velocity bias (0.9σ compared to 2.0σ), and we can rule out a model with zero assembly bias and zero velocity bias. Looking at the top panel of Figure 11 and comparing it to Figure 8, we can see that switching the assembly bias property from concentration to environment leads to improvement in almost every observable that we measure, which explains the reduction in tension. In particular, n_{gal} , $\xi(s)$, $\text{mcf}(s)$, $\text{VPF}(R)$, and $\text{SPF}(R)$ see sizeable improvement.

For the -21 sample, our best-fit results indicate negligible central galaxy assembly bias ($A_{\text{cen}} = -0.025$) and minimal satellite galaxy assembly bias ($A_{\text{sat}} = 0.165$), as well as negligible velocity bias ($B_{\text{vel}} = 1.011$). A model with no assembly bias and no velocity bias is entirely consistent with the data. This means that, for high-luminosity galaxies, neither central nor satellite galaxies show any meaningful preference toward local halo environment, and satellite galaxies move with velocities similar to the dark matter within the halo.

Like in the -19 sample, the constraints on A_{cen} and A_{sat} are substantially improved for the -21 sample when using environment as the secondary halo property as opposed to concentration. However, the constraints on M_{min} and $\sigma_{\log M}$ are actually degraded when environment-based assembly bias is used. Once again, we see no improvement in tension between our model and SDSS (4.6σ). Looking at the lower panel of Figure 11, we see little to no improvement in our residuals compared to Figure 8, which illustrates why switching the




Figure 11. Residuals between the best-fit model and the SDSS measurements for the -19 (top) and -21 (bottom) samples. For each sample, the model includes assembly bias (with environment as the secondary halo property) as well as velocity bias. We show residuals for all observables, but the model was constrained using the “dHOD” optimal observables for each sample (listed in Table 4), which are displayed with larger points.

assembly bias parameter from concentration to environment fails to reduce the tension for this sample.

These results demonstrate that low-luminosity galaxies exhibit an assembly bias signature that is present in some capacity regardless of the secondary halo property used, although the exact strength of the central and satellite assembly bias may differ for different secondary properties. Meanwhile, high-luminosity galaxies do not display any assembly bias for either secondary halo property. Moreover, the tension found between our model and SDSS is not easily alleviated with a change in secondary halo property. Investigating many different secondary halo properties for modeling assembly bias is beyond the scope of this paper; however, given our results using concentration and environment, we do not anticipate that some other secondary halo property would alleviate all of the tension that we find in the -21 sample.

7.4. Baryonic Effects

In this section, we present the results from applying the halo mass corrections from Beltz-Mohrmann & Berlind (2021) to our halo catalogs and then repeating our analysis using an HOD model with both assembly bias and velocity bias. Specifically, we utilize the mass corrections for M_{vir} halos at $z = 0$ according to the IllustrisTNG and EAGLE simulations, as well as the environment-dependent mass correction from IllustrisTNG. (The EAGLE mass correction shows very little environmental dependence, so we do not employ it here). The best-fit model parameters for these analyses are listed in Table 7. The constraints on the model parameters remain similar in size after the different mass corrections, and thus are not listed separately; we refer the reader to Table 6 for the constraints on the assembly bias + velocity bias (ABe+VB) model.

In Figure 12, we show the results of applying these halo mass corrections to the -19 halo catalogs and performing our analysis. The panels show the model parameters, using the same layout as in Figure 6. The original results (i.e., with no mass correction) are depicted in blue. The results from the EAGLE mass correction are shown in yellow, the results from the TNG mass correction are shown in green, and the results from the environment-dependent TNG mass correction (“TNG, env”) are shown in purple. In Figure 13, we show the results of applying these halo mass corrections to the -21 sample. The original results (i.e., with no mass correction) are depicted in red, and the mass-corrected results are shown with the same colors as in Figure 12.

For both samples, the mass corrections produce minimal changes to our HOD parameter constraints (with the exception of $\log M_{\text{min}}$, which does experience significant shifts in each sample). Additionally, our conclusions about the presence of assembly bias and velocity bias remain the same for both samples after the mass corrections: the -19 sample exhibits significant positive central assembly bias and negative satellite assembly bias, as well as significant velocity bias, while the -21 sample exhibits no such biases. Furthermore, the goodness of fit of the model remains roughly the same after each of the mass corrections: an HOD model with environment-based assembly bias and satellite velocity bias produces good agreement with the clustering of low-luminosity SDSS galaxies, while the same model yields significant tension with the clustering of high-luminosity SDSS galaxies. None of the mass corrections are able to alleviate this tension.

It is unsurprising that the mass corrections mainly affect the best-fit value of $\log M_{\text{min}}$ in each sample, have a slight impact on the other standard HOD parameters, and have a negligible affect on the assembly bias and velocity bias parameters. This is because the mass corrections shift the masses of our halos

Table 7
SDSS Best-fit Results with Different Mass Corrections

M_r^{lim}	Model	Mass Correction	$\log M_{\text{min}}$	$\sigma_{\log M}$	$\log M_0$	$\log M_1$	α	A_{cen}	A_{sat}	B_{vel}	p -value
−19	ABenv + VB	...	11.490	0.125	11.855	12.783	0.985	0.533	−0.224	0.826	0.364
	ABenv + VB	TNG	11.478	0.140	11.780	12.760	0.963	0.606	−0.286	0.844	0.250
	ABenv + VB	EAGLE	11.412	0.146	11.758	12.766	0.963	0.575	−0.273	0.825	0.311
	ABenv + VB	TNG, env.	11.481	0.122	11.777	12.777	0.971	0.506	−0.228	0.816	0.301
−21	ABenv + VB	...	12.740	0.495	9.984	13.917	1.079	−0.025	0.165	1.011	$4.1 \cdot 10^{-6}$
	ABenv + VB	TNG	12.671	0.455	9.258	13.924	1.122	−0.035	0.145	1.001	$1.1 \cdot 10^{-6}$
	ABenv + VB	EAGLE	12.720	0.539	11.028	13.925	1.059	−0.026	0.177	0.970	$1.5 \cdot 10^{-6}$
	ABenv + VB	TNG, env.	12.745	0.575	8.594	13.914	1.053	0.023	0.074	0.998	$8.5 \cdot 10^{-7}$

Notes. Best-fit HOD parameters for each SDSS sample using different HOD models as well as different mass corrections from Beltz-Mohrmann & Berlind (2021). The model includes concentration-based assembly bias plus velocity bias. All of these chains were run using the optimal observables identified in this work (listed in Table 4). We list the best-fit values of each parameter and indicate the goodness of fit of each parameter combination with a p -value.




Figure 12. HOD parameter constraints for the SDSS −19 sample, from a model with assembly bias (with environment as the secondary halo property) and velocity bias, after applying three different mass corrections. The crosshairs in the third and fourth panels indicate no assembly bias and no velocity bias ($A_{\text{cen}} = A_{\text{sat}} = B_{\text{vel}} = 0$).




Figure 13. HOD parameter constraints for the SDSS −21 sample, from a model with assembly bias (with environment as the secondary halo property) and velocity bias, after applying three different mass corrections. The crosshairs in the third and fourth panels indicate no assembly bias and no velocity bias ($A_{\text{cen}} = A_{\text{sat}} = B_{\text{vel}} = 0$).

(albeit in a nontrivial way), and so the parameter that governs the minimum halo mass that can host a galaxy shifts to compensate. To a lesser extent, the parameter that governs the scatter in this minimum halo mass ($\sigma_{\log M}$) and the parameters that determine the number of satellite galaxies in a halo of a given mass ($\log M_1$ and α) also shift to compensate for the halo mass correction. Meanwhile, the parameters that govern the dependence of halo occupation on a halo property *other* than mass (A_{cen} and A_{sat}) and the parameter that governs the relative velocities of the satellite galaxies to the dark matter (B_{vel}) are unaffected by changes to the halo mass function.

Overall, it is difficult to distinguish between any of the mass corrections based on their agreement with the clustering of

SDSS, nor would we rule out any of these models of baryonic physics based on our results. It is possible that, with different clustering statistics, we could tighten some of our constraints and thus differentiate between the results of the different mass corrections. It is also possible that, with a better fitting model for the −21 sample, the effect of different mass corrections on the overall tension could be seen. Ideally, we would be able to vary HOD parameters, cosmology parameters, and mass correction prescriptions simultaneously and use our results to rule out certain baryonic physics models; however, this is a challenge left for future work.

8. Conclusions

In this work, we have explored several extensions to the standard HOD model and employed an optimal set of galaxy clustering measurements to constrain this model for both high- and low-luminosity galaxies in SDSS. We first extended the standard HOD model to include parameters for central and satellite galaxy assembly bias, using halo concentration as the secondary halo property for implementing this assembly bias. We identified a set of observables to best constrain this model, using the algorithm laid out in Szezewicw et al. (2022). We then further extended our model to include an additional parameter for satellite galaxy velocity bias and repeated our analysis for both SDSS samples using this new model with both concentration-based assembly bias *and* satellite velocity bias. We then repeated this analysis using local halo environment as the assembly bias property instead of concentration. Last, we applied three different halo mass corrections to our dark matter halos to account for the impact of baryonic physics on the halo mass function; we repeated our analysis by applying our extended halo model (with environment-based assembly bias *and* satellite velocity bias) to these corrected halo masses. This is the first time that an extended HOD modeling framework, with assembly bias *and* velocity bias, a prescription to account for the impact of baryonic physics on the halo mass function, and a variety of galaxy clustering statistics measured on a wide range of scales, has been used to constrain the galaxy–halo connection in the SDSS -19 and -21 samples. Our conclusions are listed below:

1. Low-luminosity galaxies in SDSS exhibit both central and satellite galaxy assembly bias when fit with an HOD model that includes concentration-based assembly bias, with satellite galaxies displaying a negative dependence of occupation on concentration and central galaxies displaying a positive dependence on concentration (although central galaxy assembly bias is difficult to constrain). With this model, we find evidence for satellite assembly bias at the 99.8% confidence level. Additionally, this model is a substantially better fit to the clustering of low-luminosity galaxies than the standard HOD model (i.e., a model with no assembly bias).
2. When fitting the clustering of low-luminosity galaxies with an HOD model that includes both concentration-based assembly bias and satellite galaxy velocity bias, we find evidence for satellite velocity bias at the 99.8% confidence level, with satellite galaxies moving $\sim 10\%$ – 15% slower than the dark matter. The assembly bias is quite unconstrained, making it difficult to rule out a model with zero assembly bias. However, this model does further reduce the tension with SDSS.
3. When fit with an HOD model that instead uses environment-based assembly bias, low-luminosity galaxies exhibit significant negative satellite assembly bias and significant positive central assembly bias. Using environment also helps to tighten the constraints on the assembly bias parameters. This model also results in significant satellite velocity bias. We find evidence for satellite assembly bias, central assembly bias, and satellite velocity bias at the 95%, 99%, and 99.9% confidence levels, respectively. This model ultimately results in the tightest constraints on assembly bias and velocity bias, as

well as the best agreement with SDSS, with essentially no remaining tension.

4. High-luminosity galaxies exhibit negligible assembly bias when using either concentration or local environment as the assembly bias property (although the constraints are once again tighter using environment). They also exhibit negligible satellite velocity bias when fit with a model that includes both assembly bias and velocity bias. Additionally, none of these models yield good agreement with SDSS ($>4\sigma$ tension).
5. While each different treatment of baryonic physics leads to a slight change in best-fit HOD parameters, none of them significantly change our conclusions about the presence of assembly bias and velocity bias in each sample, nor do they change the goodness of fit of the HOD model used. Thus, we cannot draw any conclusions on the accuracy of our baryonic physics models based on this analysis, nor can we use baryonic physics to explain the tension we find in the -21 sample.

For low-luminosity galaxies, our results using either concentration or environment are consistent with recent results from semi-analytic models and hydrodynamic simulations (e.g., Artale et al. 2018; Zehavi et al. 2018; Bose et al. 2019). Additionally, the presence of assembly bias and velocity bias among low-luminosity galaxies but not among high-luminosity galaxies is consistent with recent findings from semi-analytic models and hydrodynamic simulations (e.g., Contreras et al. 2019; Beltz-Mohrmann et al. 2020; Contreras et al. 2021, 2022).

Our findings are also consistent with several recent observational studies. For example, Zentner et al. (2019) used concentration to model assembly bias in SDSS and found evidence for satellite assembly bias among faint galaxies ($M_r < -19$) but found no evidence for assembly bias in the $M_r < -21$ sample. Similarly, Vakili & Hahn (2019) used concentration to model assembly bias in SDSS and detected moderate central assembly bias among faint galaxies ($M_r < -20.5$, -20 , -19.5) but did not detect central galaxy assembly bias among bright galaxies ($M_r < -21.5$, -21). Meanwhile, Salcedo et al. (2022) instead used environment to model assembly bias in SDSS and similarly found no evidence for assembly bias among bright galaxies. Wang et al. (2022) also used concentration to model assembly bias in SDSS and detected positive central assembly bias for faint galaxies ($M_r < -20.5$, -20 , -19.5 , -19) and marginal negative satellite galaxy assembly bias in the $M_r < -20$ and $M_r < -19$ samples, but they did not detect assembly bias in the $M_r < -21$ sample.

The assembly bias signature among low-luminosity galaxies can be understood as follows: Early-forming halos ultimately contain fewer satellites, because they acquired their satellites earlier, and thus these satellites were subject to the destructive processes of the host halo (i.e., merging) for a longer period of time (Zentner et al. 2005). Thus, it is reasonable that satellite galaxies would preferentially reside in late-forming halos. Formation time is strongly correlated with halo concentration, with early-forming halos having higher concentrations (e.g., Wechsler et al. 2002; Zhao et al. 2003) and thus fewer satellite galaxies. This also explains why satellite galaxies preferentially reside in halos with low-density environments: satellites residing in high-density environments are more vulnerable to

mergers, and thus host halos in high-density environments will ultimately contain fewer satellites.

Meanwhile, among Milky Way–sized halos, galaxies residing in higher concentration halos tend to be more luminous (Zentner et al. 2019). This is because, at fixed mass, halos with higher concentrations have deeper potential wells, allowing gravity to more strongly bind the stellar and gas contents of these halos, possibly leading to more rapid star formation (or less vulnerability to processes that suppress star formation). Thus, we can understand why central galaxies residing in Milky Way–sized halos seem to preferentially reside in high-concentration halos: the deep potential well of the host halo ultimately leads to a more luminous central galaxy, which is more likely to pass our luminosity cut than a central galaxy living in a shallow potential well. The same logic can be used to understand why central galaxies prefer to reside in halos with high-density environments—such environments are more conducive to merging, leading to a more luminous central galaxy in the end.

The lack of assembly bias signature among high-luminosity galaxies has a few possible explanations. It has been found that, among Milky Way–sized halos, the number of subhalos in a given host halo depends heavily on host halo environment, but among cluster-mass halos, the abundance of subhalos exhibits no such environmental dependence (e.g., Zentner et al. 2019). This explains why, when using halo environment as our secondary halo property, we detect satellite assembly bias among low-luminosity galaxies but not among high-luminosity galaxies. Another possible explanation is that satellite galaxies in the -21 sample are all recent additions to the halo, and so they have not had enough time to be destroyed via mergers; thus, whether the host halo is high- or low-concentration (or lives in a high- or low-density environment) makes no difference for the presence of satellite galaxies in this sample. This could also explain why we detect satellite galaxy velocity bias among low-luminosity galaxies but not among high-luminosity galaxies: satellite galaxies in the -19 sample have likely been slowed down via dynamical friction over time, whereas satellite galaxies in the -21 sample are all more recent acquisitions to the halo and thus have not had time to be significantly affected by dynamical friction.

While several physical explanations are reasonable, the fact remains that, among low-luminosity galaxies, the model that includes both assembly bias and satellite velocity bias exhibits minimal tension; in other words, the model is in good agreement with the clustering of -19 SDSS galaxies. This is consistent with our expectation for low-luminosity galaxies and with the minimal tension found in previous studies of this nature for low-luminosity galaxies. By contrast, the model is *not* in good agreement with the clustering of -21 SDSS galaxies. This high degree of tension for high-luminosity galaxies is in contrast with these previous studies, which did not find any significant tension with SDSS. The larger constraining power of our results is likely, in part, due to the large set of optimal clustering statistics that we use.

This tension that we find among the higher-luminosity sample could be indicative of several things. For example, it is possible that central galaxies do indeed exhibit significant velocity bias (e.g., Guo et al. 2015a, 2015b), and including this in the model would lead to better agreement with SDSS. However, we have examined the impact on our observables after adding central velocity bias to our best-fit model at the

level found in previous works, and we found that this has a negligible effect on our clustering measurements. Thus, if central velocity bias is indeed present in SDSS, it is not currently detectable with our clustering measurements, given our uncertainty due to cosmic variance. It is also possible that galaxies do not trace the spatial distribution of dark matter within halos (i.e., there is spatial bias; see Watson et al. 2012; Piscionere et al. 2015). Additionally, the standard HOD model assumes that the number of satellite galaxies in each halo is governed by a Poisson distribution, but recent results indicate that this is probably not the case (Boylan-Kolchin et al. 2010; Mao et al. 2015; Jiménez et al. 2019).

It is also possible that a change in halo definition or halo finder could alter our results. S18 repeated their analysis twice, once using M_{200b} halos and again using M_{vir} halos, and found similar results, with M_{vir} halos producing slightly tighter constraints. For this reason, we have proceeded using only M_{vir} halos. Because a small change in halo definition simply leads to a small change in mass for all halos, we think it likely that the HOD parameters can compensate for any small change in halo definition; in this case, our best-fit parameter values would change slightly, but our overall conclusions about the presence of assembly bias or the goodness of fit of our model would not change. However, a significant change in halo definition or halo finder could potentially lead to changes in our conclusions. For example, several works have found that the proper treatment of splashback halos could lead to a reduction in the assembly bias signature for low-luminosity galaxies (Villarreal et al. 2017; Mansfield & Kravtsov 2020).





In future work, it is worth investigating whether accounting for these possibilities leads to improved agreement between our model and the observed clustering of high-luminosity galaxies. However, we think it unlikely that accounting for these affects would be enough to explain the amount of tension we are finding. In fact, in hydrodynamic simulations, the standard HOD model proved to be a good fit to the clustering of high-luminosity galaxies, provided that the model was applied to a DMO simulation with the same cosmological model as the hydrodynamic simulation in question (Beltz-Mohrmann et al. 2020). Thus, the fact that we find such significant tension in our analysis between our best-fit HOD model and the clustering of high-luminosity galaxies leads us to believe that there may be an issue with our cosmological model. It is possible that our clustering statistics are able to detect such an issue for our high-luminosity sample, but are not sensitive enough to pick up on a cosmological discrepancy among low-luminosity galaxies. Such a result would be consistent with the findings of several other recent analyses (e.g., Wibking et al. 2020; Chapman et al. 2022; Lange et al. 2022; Zhai et al. 2022). For example, Lange et al. (2022) used an HOD model with both assembly bias and velocity bias parameters to obtain cosmological constraints from the BOSS LOWZ sample. Using V_{max} as their assembly bias property, they did not find significant evidence of either central or satellite galaxy assembly bias, and they found only minimal evidence for central velocity bias and no evidence of satellite velocity bias. However, they found that their best cosmological constraints were slightly inconsistent with the Planck observations. Similarly, Zhai et al. (2022) used the Aemulus suite of cosmological N -body simulations to model the clustering of BOSS galaxies, using an HOD model with both assembly bias (based on environment) and velocity bias. They found some evidence for positive galaxy assembly bias

but no evidence for satellite galaxy velocity bias. Additionally, they found that their cosmological constraints exhibited some tension with the Planck observations.

In future work, we intend to explore whether a change in cosmological parameters could be the key to alleviating this tension that we are finding. It is worth noting that the Dark Energy Spectroscopic Instrument (DESI; DESI Collaboration et al. 2016) will have better precision than the SDSS, due to its larger volume, allowing it to potentially detect even smaller differences in clustering measurements. Applying our model to upcoming DESI data could allow us to gain better constraints on our halo model parameters, differentiate between different baryonic feedback implementations, and ultimately constrain cosmology using small-scale galaxy clustering.

We would like to sincerely thank our anonymous referee for helpful comments, which improved the quality of this paper. This project has been supported by the National Science Foundation (NSF) through Award (AST-1909631), and has made use of NASA’s Astrophysics Data System; MATPLOTLIB, a Python library for publication quality graphics (Hunter 2007); SCIPY (Virtanen et al. 2020); the IPYTHON package (Pérez & Granger 2007); ASTROPY, a community-developed core Python package for Astronomy (Astropy Collaboration et al. 2013, 2018); NUMPY (Harris et al. 2020); PANDAS (McKinney 2010, 2011); and CHAINCONSUMER (Hinton 2016). Funding for the SDSS and SDSS-II has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, the U.S. Department of Energy, the National Aeronautics and Space Administration, the Japanese Monbukagakusho, the Max Planck Society, and the Higher Education Funding Council for England. The SDSS website is <http://www.sdss.org/>. The SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions. The Participating Institutions are the American Museum of Natural History, Astrophysical Institute Potsdam, University of Basel, University of Cambridge, Case Western Reserve University, University of Chicago, Drexel University, Fermilab, the Institute for Advanced Study, the Japan Participation Group, Johns Hopkins University, the Joint Institute for Nuclear Astrophysics, the Kavli Institute for Particle Astrophysics and Cosmology, the Korean Scientist Group, the Chinese Academy of Sciences (LAMOST), Los Alamos National Laboratory, the Max-Planck-Institute for Astronomy (MPIA), the Max-Planck-Institute for Astrophysics (MPA), New Mexico State University, Ohio State University, University of Pittsburgh, University of Portsmouth, Princeton University, the United States Naval Observatory, and the University of Washington. The mock catalogs used in this paper were produced by the LasDamas project (<http://lss.phy.vanderbilt.edu/lasdamas/>); we thank NSF XSEDE for providing the computational resources for LasDamas. The MCMCs in this work were run on the Texas Advanced Computing Center’s Stampede2 supercomputer. Some of the computational facilities used in this project were provided by the Vanderbilt Advanced Computing Center for Research and Education (ACCRES). Parts of this research were conducted by the Australian Research Council Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), through project number CE170100013. These acknowledgements were compiled using the Astronomy Acknowledgement Generator (<http://astrofrog.github.io/acknowledgment-generator/>).

ORCID iDs

Gillian D. Beltz-Mohrmann  <https://orcid.org/0000-0002-4392-8920>
 Adam O. Szewciw  <https://orcid.org/0000-0001-9094-8433>
 Andreas A. Berlind  <https://orcid.org/0000-0002-1814-2002>
 Manodeep Sinha  <https://orcid.org/0000-0002-4845-1228>

References

- Abazajian, K. N., Adelman-McCarthy, J. K., Agüeros, M. A., et al. 2009, *ApJS*, **182**, 543
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
- Artale, M. C., Zehavi, I., Contreras, S., & Norberg, P. 2018, *MNRAS*, **480**, 3978
- Baugh, C. M., Benson, A. J., Cole, S., Frenk, C. S., & Lacey, C. G. 1999, *MNRAS*, **305**, L21
- Behroozi, P., Hearin, A., & Moster, B. P. 2022, *MNRAS*, **509**, 2800
- Behroozi, P. S., Wechsler, R. H., & Wu, H.-Y. 2013, *ApJ*, **762**, 109
- Beltz-Mohrmann, G. D., & Berlind, A. A. 2021, *ApJ*, **921**, 112
- Beltz-Mohrmann, G. D., Berlind, A. A., & Szewciw, A. O. 2020, *MNRAS*, **491**, 5771
- Benson, A. J., Cole, S., Frenk, C. S., Baugh, C. M., & Lacey, C. G. 2000, *MNRAS*, **311**, 793
- Berlind, A. A., Frieman, J., Weinberg, D. H., et al. 2006, *ApJS*, **167**, 1
- Berlind, A. A., & Weinberg, D. H. 2002, *ApJ*, **575**, 587
- Berlind, A. A., Weinberg, D. H., Benson, A. J., et al. 2003, *ApJ*, **593**, 1
- Blanton, M. R., Schlegel, D. J., Strauss, M. A., et al. 2005, *AJ*, **129**, 2562
- Bose, S., Eisenstein, D. J., Hernquist, L., et al. 2019, *MNRAS*, **490**, 5693
- Boylan-Kolchin, M., Springel, V., White, S. D. M., & Jenkins, A. 2010, *MNRAS*, **406**, 896
- Bryan, G. L., & Norman, M. L. 1998, *ApJ*, **495**, 80
- Chapman, M. J., Mohammad, F. G., Zhai, Z., et al. 2022, *MNRAS*, **516**, 617
- Contreras, S., Angulo, R. E., Chaves-Montero, J., White, S. D. M., & Aricò, G. 2023, *MNRAS*, **520**, 489
- Contreras, S., Angulo, R. E., & Zennaro, M. 2021, *MNRAS*, **504**, 5205
- Contreras, S., Zehavi, I., Padilla, N., et al. 2019, *MNRAS*, **484**, 1133
- Cooray, A., & Sheth, R. 2002, *PhR*, **372**, 1
- Coupon, J., Arnouts, S., van Waerbeke, L., et al. 2015, *MNRAS*, **449**, 1352
- Crocce, M., Pueblas, S., & Scoccimarro, R. 2006, *MNRAS*, **373**, 369
- Crocce, M., Pueblas, S., & Scoccimarro, R. 2012, 2LPTIC: 2nd-order Lagrangian Perturbation Theory Initial Conditions, Astrophysics Source Code Library, ascl:1201.005
- Croton, D. J., Gao, L., & White, S. D. M. 2007, *MNRAS*, **374**, 1303
- DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016, arXiv:1611.00036
- Dong-Páez, C. A., Smith, A., Szewciw, A. O., et al. 2022, arXiv:2208.00540
- 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
- Gao, L., & White, S. D. M. 2007, *MNRAS*, **377**, L5
- Guo, H., Zheng, Z., Behroozi, P. S., et al. 2016, *MNRAS*, **459**, 3040
- Guo, H., Zheng, Z., Zehavi, I., et al. 2015a, *MNRAS*, **446**, 578
- Guo, H., Zheng, Z., Zehavi, I., et al. 2015b, *MNRAS*, **453**, 4368
- Hadzhiyska, B., Bose, S., Eisenstein, D., & Hernquist, L. 2021a, *MNRAS*, **501**, 1603
- Hadzhiyska, B., Bose, S., Eisenstein, D., Hernquist, L., & Spergel, D. N. 2020, *MNRAS*, **493**, 5506
- Hadzhiyska, B., Liu, S., Somerville, R. S., et al. 2021b, *MNRAS*, **508**, 698
- Hadzhiyska, B., Tacchella, S., Bose, S., & Eisenstein, D. J. 2021c, *MNRAS*, **502**, 3599
- Hahn, C., Beutler, F., Sinha, M., et al. 2019, *MNRAS*, **485**, 2956
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
- Hearin, A. P., Zentner, A. R., van den Bosch, F. C., Campbell, D., & Tollerud, E. 2016, *MNRAS*, **460**, 2552
- Hinton, S. R. 2016, *JOSS*, **1**, 00045
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Jiménez, E., Contreras, S., Padilla, N., et al. 2019, *MNRAS*, **490**, 3532
- Jing, Y. P., Mo, H. J., & Börner, G. 1998, *ApJ*, **494**, 1
- Kauffmann, G., Colberg, J. M., Diaferio, A., & White, S. D. M. 1999, *MNRAS*, **303**, 188
- Kauffmann, G., Nusser, A., & Steinmetz, M. 1997, *MNRAS*, **286**, 795

- Klypin, A. A., Trujillo-Gomez, S., & Primack, J. 2011, *ApJ*, 740, 102
- Kravtsov, A. V., Berlind, A. A., Wechsler, R. H., et al. 2004, *ApJ*, 609, 35
- Lacey, C., & Cole, S. 1994, *MNRAS*, 271, 676
- Lange, J. U., Hearin, A. P., Leauthaud, A., et al. 2022, *MNRAS*, 509, 1779
- Lange, J. U., van den Bosch, F. C., Zentner, A. R., et al. 2019, *MNRAS*, 490, 1870
- Leauthaud, A., Tinker, J., Bundy, K., et al. 2012, *ApJ*, 744, 159
- Lehmann, B. V., Mao, Y.-Y., Becker, M. R., Skillman, S. W., & Wechsler, R. H. 2017, *ApJ*, 834, 37
- Ma, C.-P., & Fry, J. N. 2000, *ApJ*, 543, 503
- Mansfield, P., & Kravtsov, A. V. 2020, *MNRAS*, 493, 4763
- Mao, Y.-Y., Williamson, M., & Wechsler, R. H. 2015, *ApJ*, 810, 21
- Mao, Y.-Y., Zentner, A. R., & Wechsler, R. H. 2018, *MNRAS*, 474, 5143
- McBride, C., Berlind, A., Scoccimarro, R., et al. 2009, *BAAS*, 41, 253
- McCarthy, K. S., Zheng, Z., & Guo, H. 2019, *MNRAS*, 487, 2424
- McCarthy, K. S., Zheng, Z., Guo, H., Luo, W., & Lin, Y.-T. 2022, *MNRAS*, 509, 380
- McClelland, J., & Silk, J. 1977, *ApJ*, 217, 331
- McCullagh, N., Norberg, P., Cole, S., et al. 2017, arXiv:1705.01988
- McKinney, W. 2010, in Proc. the 9th Python in Science Conf., ed. S. van der Walt & J. Millman, 51
- McKinney, W. 2011, Python for High Performance and Scientific Computing, 14
- Navarro, J. F., Frenk, C. S., & White, S. D. M. 1997, *ApJ*, 490, 493
- Neyman, J., & Scott, E. L. 1952, *ApJ*, 116, 144
- Norberg, P., Baugh, C. M., Gaztañaga, E., & Croton, D. J. 2009, *MNRAS*, 396, 19
- Padilla, N., Contreras, S., Zehavi, I., Baugh, C. M., & Norberg, P. 2019, *MNRAS*, 486, 582
- Parejko, J. K., Sunayama, T., Padmanabhan, N., et al. 2013, *MNRAS*, 429, 98
- Peacock, J. A., & Smith, R. E. 2000, *MNRAS*, 318, 1144
- Peebles, P. J. E. 1974, *A&A*, 32, 197
- Pérez, F., & Granger, B. E. 2007, *CSE*, 9, 21
- Perez, L. A., Malhotra, S., Rhoads, J. E., & Tilvi, V. 2021, *ApJ*, 906, 58
- Piscionere, J. A., Berlind, A. A., McBride, C. K., & Scoccimarro, R. 2015, *ApJ*, 806, 125
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014, *A&A*, 571, A16
- Pujol, A., Hoffmann, K., Jiménez, N., & Gaztañaga, E. 2017, *A&A*, 598, A103
- Salcedo, A. N., Maller, A. H., Berlind, A. A., et al. 2018, *MNRAS*, 475, 4411
- Salcedo, A. N., Zu, Y., Zhang, Y., et al. 2022, *SCPMA*, 65, 109811
- Scherrer, R. J., & Bertschinger, E. 1991, *ApJ*, 381, 349
- Scoccimarro, R. 1998, *MNRAS*, 299, 1097
- Scoccimarro, R., Sheth, R. K., Hui, L., & Jain, B. 2001, *ApJ*, 546, 20
- Seljak, U. 2000, *MNRAS*, 318, 203
- Seljak, U., & Zaldarriaga, M. 1996, *ApJ*, 469, 437
- Sheth, R. K., Hui, L., Diaferio, A., & Scoccimarro, R. 2001, *MNRAS*, 325, 1288
- Sheth, R. K., & Tormen, G. 2004, *MNRAS*, 350, 1385
- Shi, J., & Sheth, R. K. 2018, *MNRAS*, 473, 2486
- Sinha, M., Berlind, A. A., McBride, C. K., et al. 2018, *MNRAS*, 478, 1042
- Sinha, M., & Garrison, L. H. 2019, in Software Challenges to Exascale Computing. Second Workshop, ed. A. Majumdar & R. Arora (Singapore: Springer), 3
- Sinha, M., & Garrison, L. H. 2020, *MNRAS*, 491, 3022
- Springel, V. 2005, *MNRAS*, 364, 1105
- Storey-Fisher, K., Tinker, J., Zhai, Z., et al. 2022, arXiv:2210.03203
- Szewciw, A. O., Beltz-Mohrmann, G. D., Berlind, A. A., & Sinha, M. 2022, *ApJ*, 926, 15
- Tinker, J. L., Conroy, C., Norberg, P., et al. 2008, *ApJ*, 686, 53
- Tinker, J. L., Weinberg, D. H., & Warren, M. S. 2006a, *ApJ*, 647, 737
- Tinker, J. L., Weinberg, D. H., & Zheng, Z. 2006b, *MNRAS*, 368, 85
- Tonegawa, M., Park, C., Zheng, Y., et al. 2020, *ApJ*, 897, 17
- Vakilii, M., & Hahn, C. 2019, *ApJ*, 872, 115
- Villarreal, A. S., Zentner, A. R., Mao, Y.-Y., et al. 2017, *MNRAS*, 472, 1088
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, 17, 261
- Walsh, K., & Tinker, J. 2019, *MNRAS*, 488, 470
- Wang, K., Mao, Y.-Y., Zentner, A. R., et al. 2019, *MNRAS*, 488, 3541
- Wang, K., Mao, Y.-Y., Zentner, A. R., et al. 2022, *MNRAS*, 516, 4003
- Watson, D. F., Berlind, A. A., McBride, C. K., Hogg, D. W., & Jiang, T. 2012, *ApJ*, 749, 83
- Wechsler, R. H., Bullock, J. S., Primack, J. R., Kravtsov, A. V., & Dekel, A. 2002, *ApJ*, 568, 52
- Wechsler, R. H., & Tinker, J. L. 2018, *ARA&A*, 56, 435
- Wechsler, R. H., Zentner, A. R., Bullock, J. S., Kravtsov, A. V., & Allgood, B. 2006, *ApJ*, 652, 71
- White, M., Hernquist, L., & Springel, V. 2001, *ApJL*, 550, L129
- Wibking, B. D., Weinberg, D. H., Salcedo, A. N., et al. 2020, *MNRAS*, 492, 2872
- Xu, X., & Zheng, Z. 2020, *MNRAS*, 492, 2739
- York, D. G., Adelman, J., Anderson, J. E., Jr., et al. 2000, *AJ*, 120, 1579
- Yuan, S., Garrison, L. H., Hadzhiyska, B., Bose, S., & Eisenstein, D. J. 2021, *MNRAS*, 510, 3301
- Zaldarriaga, M., & Seljak, U. 2000, *ApJS*, 129, 431
- Zaldarriaga, M., Seljak, U., & Bertschinger, E. 1998, *ApJ*, 494, 491
- Zehavi, I., Blanton, M. R., Frieman, J. A., et al. 2002, *ApJ*, 571, 172
- Zehavi, I., Contreras, S., Padilla, N., et al. 2018, *ApJ*, 853, 84
- Zehavi, I., Weinberg, D. H., Zheng, Z., et al. 2004, *ApJ*, 608, 16
- Zehavi, I., Zheng, Z., Weinberg, D. H., et al. 2005, *ApJ*, 630, 1
- Zehavi, I., Zheng, Z., Weinberg, D. H., et al. 2011, *ApJ*, 736, 59
- Zentner, A. R., Berlind, A. A., Bullock, J. S., Kravtsov, A. V., & Wechsler, R. H. 2005, *ApJ*, 624, 505
- Zentner, A. R., Hearin, A., van den Bosch, F. C., Lange, J. U., & Villarreal, A. 2019, *MNRAS*, 485, 1196
- Zentner, A. R., Hearin, A. P., & van den Bosch, F. C. 2014, *MNRAS*, 443, 3044
- Zhai, Z., Tinker, J. L., Banerjee, A., et al. 2022, arXiv:2203.08999
- Zhao, D. H., Mo, H. J., Jing, Y. P., & Börner, G. 2003, *MNRAS*, 339, 12
- Zheng, Z. 2004, *ApJ*, 610, 61
- Zheng, Z., Berlind, A. A., Weinberg, D. H., et al. 2005, *ApJ*, 633, 791
- Zheng, Z., Coil, A. L., & Zehavi, I. 2007, *ApJ*, 667, 760
- Zheng, Z., & Weinberg, D. H. 2007, *ApJ*, 659, 1
- Zu, Y., & Mandelbaum, R. 2018, *MNRAS*, 476, 1637