

The completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: measurement of the growth rate of structure from the small-scale clustering of the luminous red galaxy sample

Michael J. Chapman,^{1,2*} Faizan G. Mohammad,^{1,2} Zhongxu Zhai,³ Will J. Percival^{1,2,4}, Jeremy L. Tinker^{1,5}, Julian E. Bautista^{1,6,7}, Joel R. Brownstein^{1,8}, Etienne Burtin,⁹, Kyle S. Dawson,⁸ Héctor Gil-Marín^{10,11}, Axel de la Macorra,¹² Ashley J. Ross^{1,13}, Graziano Rossi,¹⁴ Donald P. Schneider^{15,16} and Gong-Bo Zhao¹⁷

¹Waterloo Centre for Astrophysics, University of Waterloo, Waterloo, ON N2L 3G1, Canada

²Department of Physics and Astronomy, University of Waterloo, Waterloo, ON N2L 3G1, Canada

³IPAC, California Institute of Technology, Mail Code 314-6, 1200 E. California Blvd., Pasadena, CA 91125, USA

⁴Perimeter Institute for Theoretical Physics, 31 Caroline St. North, Waterloo, ON N2L 2Y5, Canada

⁵Center for Cosmology and Particle Physics, Department of Physics, New York University, 726 Broadway, New York, NY 10003, USA

⁶Institute of Cosmology and Gravitation, Dennis Sciama Building, University of Portsmouth, Portsmouth PO1 3FX, UK

⁷Aix Marseille University, CNRS/IN2P3, CPPM, Marseille, France

⁸Department of Physics and Astronomy, University of Utah, 115 S. 1400 E., Salt Lake City, UT 84112, USA

⁹IRFU, CEA, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France

¹⁰Institut de Ciències del Cosmos, Universitat de Barcelona, ICCUB, Martí i Franquès 1, E-08028 Barcelona, Spain

¹¹Institut d'Estudis Espacials de Catalunya (IEEC), E-08034 Barcelona, Spain

¹²Instituto de Física, Universidad Nacional Autónoma de México, Apdo. Postal 20-364, México

¹³Department of Astronomy, The Ohio State University, 140 W. 18th Ave., Columbus, OH 43210, USA

¹⁴Department of Physics and Astronomy, Sejong University, Seoul 143-747, Korea

¹⁵Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA

¹⁶Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA

¹⁷National Astronomical Observatories of China, Chinese Academy of Sciences, 20A Datun Road, Chaoyang District, Beijing 100012, China

Accepted 2022 June 29. Received 2022 June 20; in original form 2021 June 28

ABSTRACT

We measure the small-scale clustering of the Data Release 16 extended Baryon Oscillation Spectroscopic Survey Luminous Red Galaxy sample, corrected for fibre-collisions using Pairwise Inverse Probability weights, which give unbiased clustering measurements on all scales. We fit to the monopole and quadrupole moments and to the projected correlation function over the separation range $7 - 60 h^{-1} \text{Mpc}$ with a model based on the AEMULUS cosmological emulator to measure the growth rate of cosmic structure, parametrized by $f\sigma_8$. We obtain a measurement of $f\sigma_8(z = 0.737) = 0.408 \pm 0.038$, which is 1.4σ lower than the value expected from 2018 Planck data for a flat ΛCDM model, and is more consistent with recent weak-lensing measurements. The level of precision achieved is 1.7 times better than more standard measurements made using only the large-scale modes of the same sample. We also fit to the data using the full range of scales $0.1 - 60 h^{-1} \text{Mpc}$ modelled by the AEMULUS cosmological emulator and find a 4.5σ tension in the amplitude of the halo velocity field with the Planck + ΛCDM model, driven by a mismatch on the non-linear scales. This may not be cosmological in origin, and could be due to a breakdown in the Halo Occupation Distribution model used in the emulator. Finally, we perform a robust analysis of possible sources of systematics, including the effects of redshift uncertainty and incompleteness due to target selection that were not included in previous analyses fitting to clustering measurements on small scales.

Key words: galaxies: distances and redshifts – cosmological parameters – cosmology: observations – large-scale structure of Universe.

1 INTRODUCTION

Understanding the accelerating expansion of the Universe is one of the primary goals for modern physics experiments. Many of these

experiments aim to accomplish this through measuring the observed positions of galaxies in the Universe, which depend on the cosmological model in a number of ways. The intrinsic distribution of galaxies results from the growth of initial matter perturbations through gravity, giving a window to the early Universe. However, the fundamental observables are the angular positions and redshifts of galaxies, while the intrinsic pattern is in comoving distances, so surveys are also

* E-mail: mj3chapman@uwaterloo.ca

sensitive to the link between these two coordinates. This link depends on the relationship between separations in angles and redshifts and distances across and along the line of sight (los; Alcock & Paczynski 1979), as well as on redshift-space distortions (Kaiser 1987). Because these depend on both cosmological expansion and the build-up of structure within the Universe, large galaxy surveys offer a unique opportunity to solve the question of the origin of the late acceleration of the expansion (Weinberg et al. 2013; Ferreira 2019).

The growth of structure most clearly manifests on the observed galaxy distribution through Redshift Space Distortions (RSD; Kaiser 1987). These are a consequence of the velocities of galaxies in a comoving frame distorting the los cosmological distances based on observed redshifts, and are sensitive to the growth rate of structure, which in turn depends on the strength of gravity. The strength of the RSD measurements depend on the parameter $f\sigma_8$, which is commonly used to quantify the amplitude of the velocity power spectrum and provides a strong test of modifications to gravity (Guzzo et al. 2008; Song & Percival 2009). The development of large galaxy surveys driven by advances in multi-object spectrographs has resulted in recent renewed interest in RSD including measurements from the WiggleZ (Blake et al. 2011), 6dFGS (Beutler et al. 2012), SDSS-II (Samushia, Percival & Raccanelli 2012), SDSS-MGS (Howlett et al. 2015), FastSound (Okumura et al. 2016), and VIPERS (Pezzotta et al. 2017) galaxy surveys.

The best precision measurements to date come from the Baryon Oscillation Spectroscopic Survey (BOSS; Dawson et al. 2013), part of the third generation of the Sloan Digital Sky Survey (SDSS; Eisenstein et al. 2011). Using large-scale modes, BOSS has achieved the best precision of ~ 6 per cent on the parameter combination $f\sigma_8$ (Beutler et al. 2017; Grieb et al. 2017; Sánchez et al. 2017; Satpathy et al. 2017). Note that these studies all measured RSD in the linear or quasi-linear regime, where proportionately small levels of non-linear modelling were required.

In contrast, Reid et al. (2014) made a measurement of the amplitude of the RSD signal from an early BOSS galaxy sample, fitting to the monopole and quadrupole moments of the correlation function over scales 0.8 to $32 h^{-1}\text{Mpc}$, obtaining a 2.5-per cent measurement of $f\sigma_8(z=0.57) = 0.450 \pm 0.011$. This demonstrates the increased precision available if RSD in the data can be accurately measured and modelled to small scales. The most accurate method to model small-scale clustering is to use N-body simulations, and this was the route taken by Reid et al. (2014). However, without a simulation for each model to be tested (Reid et al. 2014 used three simulation sets at three very similar cosmologies), one has to extrapolate solutions to different cosmologies, which needs care. The most pernicious problem faced in the Reid et al. (2014) analysis was correcting the small-scale clustering in the data, which suffers from fibre-collisions, where hardware limitations mean that some galaxies are excluded from the catalogue due to having close neighbours. A similar method was recently applied to the BOSS LOWZ galaxies (Lange et al. 2022), and a study is in preparation for the CMASS sample (Zhai et al. 2022).

The extended Baryon Oscillation Spectroscopic Survey (eBOSS; Dawson et al. 2016), part of the SDSS-IV experiment (Blanton et al. 2017) is the latest in a line of galaxy surveys done using the Sloan Telescope. This experiment was designed to make Baryon Acoustic Oscillations (BAO) and RSD measurements using three classes of galaxies used to directly trace the density field, together with a high redshift quasar sample (du Mas des Bourboux et al. 2020) that allows Lyman- α forest measurements at redshifts $z > 2.1$. We use the Luminous Red Galaxy (LRG) sample from Data Release 16 (Ahumada et al. 2020) to make RSD measurements

at $z \sim 0.7$ including small-scale information. Standard BAO and RSD measurements made with this sample on larger scales only are presented in Bautista et al. (2021), Gil-Marín et al. (2020), together with a test of their methodology using mock catalogues in Rossi et al. (2021). At intermediate redshifts, eBOSS probes the Universe using samples of emission line galaxies (Tamone et al. 2020; de Mattia et al. 2021; Raichoor et al. 2021) and quasars (Lyke et al. 2020; Neveux et al. 2020; Ross et al. 2020; Smith et al. 2020; Hou et al. 2021) as direct tracers of the density field lower redshifts. We do not analyse these data, focusing instead on the easier to model LRG sample. The cosmological interpretation of the BAO and RSD results from all eBOSS samples was presented in Alam et al. (2021).

Pushing the modelling to include small scales in our analysis is made possible by two key advances in methodology since the Reid et al. (2014) analysis. First, we use the AEMULUS emulator (Zhai et al. 2019) to create accurate models of the redshift-space correlation function moments to small scales (see Section 3.3). To correct for fibre-collisions, we use the Pairwise Inverse Probability (PIP) method (Bianchi & Percival 2017; Percival & Bianchi 2017), as described in Section 3.2. Together, these advances mean that we can now both make and model accurate clustering measurements from the eBOSS LRG sample, fitting the correlation function to small scales.

Our paper is structured as follows: the eBOSS LRG sample is described in Section 2, and the method for measuring and fitting the correlation functions in Section 3. In Section 4, we perform various tests of the method using mock catalogues. We present our results in Section 5, and discuss their significance in Section 6. Finally, we summarize our results in Section 7.

2 EBOSS LRG SAMPLE

The eBOSS LRG target sample was selected (Prakash et al. 2016) from SDSS DR13 photometry (Albareti et al. 2017), together with infrared observations from the *WISE* satellite (Lang, Hogg & Schlegel 2016). LRG targets were selected over 7500 deg^2 , and observed using the BOSS spectrographs (Smee et al. 2013) mounted on the 2.5-m Sloan telescope (Gunn et al. 2006).

In order to measure clustering we quantify the sample mask, detailing where we could observe galaxies, using the random catalogue with 50 times more points than galaxies as described in Ross et al. (2020). Regions with bad photometric properties, that are close to higher priority targets, or near the centrepiece region of the plates are masked, removing 17 per cent of the initial footprint. Redshifts for the randoms were sampled from those of the galaxies.

Redshifts were measured from the resulting spectra using the REDROCK algorithm.¹ REDROCK fits the data with templates derived from principal component analysis of SDSS data, followed by a redshift refinement procedure that uses stellar population models. We are unable to obtain a reliable redshift estimate from many spectra (3.5 per cent on average across the survey), with a failure fraction with systematic angular variations. We therefore apply a weight w_{noz} as described in Ross et al. (2020) to galaxies to remove these variations, calculated as a function of position of the fibre on the detector and the signal to noise of that set of observations.

Systematic variations in the density of galaxies caused by variations in the photometric data used for target selection are mitigated by applying weights w_{sys} to the galaxies. These were computed using a multilinear regression on the observed relations between the angular overdensities of galaxies versus stellar density, seeing and galactic

¹ Available at github.com/desihub/redrock

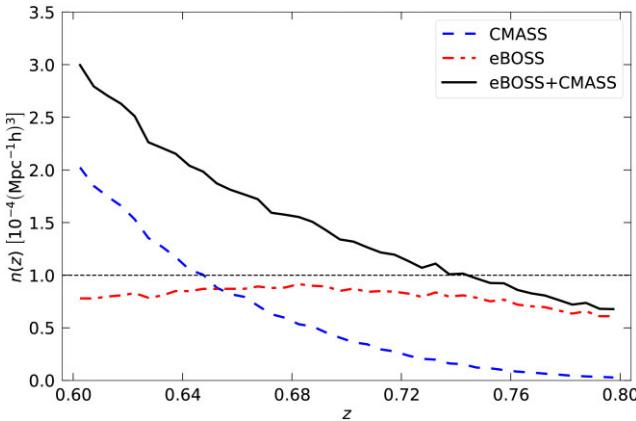


Figure 1. Redshift distribution of the eBOSS DR16 (red dash-dotted line), CMASS DR12 (blue dashed line), and the joint eBOSS + CMASS sample (black thick line, see Section 5.7 for details), optimized using w_{FKP} weights.

extinction. As we are interested primarily in small-scales, the exact correction is not important. Additional weights w_{FKP} that optimize the signal, which varies because the density varies across the sample (Feldman, Kaiser & Peacock 1994), are also included (Fig. 1).

A fibre could not be placed on 4 per cent of the LRG targets due to fibre-collisions: when a group of two or more galaxies are closer than 62 arcsec, they cannot all receive a fibre because of hardware limitations. We use PIP weights w^{PIP} together with angular upweighting (Bianchi & Percival 2017; Percival & Bianchi 2017) to correct for this effect, as described in Mohammad et al. (2020), and Section 3.2. The final combined weight applied to the galaxies is defined as $w^{\text{tot}} = w^{\text{noz}} w^{\text{sys}} w^{\text{FKP}}$, and we also use w^{PIP} applied to pairs.

The eBOSS sample of LRGs overlaps in area and redshift range with the high-redshift tail of the BOSS CMASS sample. Unlike many other eBOSS analyses including the large-scale measurements of BAO and RSD (Gil-Marín et al. 2020; Bautista et al. 2021), we do not combine the eBOSS LRG sample with all the $z > 0.6$ BOSS CMASS galaxies. We focus on the eBOSS sample to simplify the correction of the small-scale fibre assignment: fibre assignment was performed separately for BOSS and eBOSS using different configurations of the SDSS tiling code.

We define the effective redshift of our sample as the weighted mean redshift of galaxy pairs

$$z_{\text{eff}} = \frac{\sum_{m>n} w_{mn}^{\text{PIP}} w_n^{\text{tot}} w_m^{\text{tot}} (z_m + z_n) / 2}{\sum_{m>n} w_{mn}^{\text{PIP}} w_m^{\text{tot}} w_n^{\text{tot}}}, \quad (1)$$

where the indices m, n are over the objects in the data catalogue, and the description of the weights is given in Section 3.2. Additionally, we only include galaxy pairs which have a separation between $0.1-60 h^{-1}\text{Mpc}$, the scales used in our measurement. The effective redshift we obtain for our sample is $z = 0.737$, and an effective comoving volume of 1.28 Gpc^3 (Ross et al. 2020).

3 METHODS

3.1 Measurements

We measure and model the observed galaxy clustering in redshift space using the two-point correlation function as calculated using the least-bias and least-variance Landy–Szalay estimator (Landy &

Szalay 1993)

$$\xi(s) = \frac{DD(s) - 2DR(s)}{RR(s)} + 1, \quad (2)$$

with DD , DR , and RR being the data–data, data–random, and random–random pair counts at a given separation s . To reduce the impact of shot noise on the measured ξ from the random catalogue, we use a number of random points $N = 50$ times the number of galaxies in the DR16 sample. The difference in the number of galaxies and randoms is accounted for by normalizing the pair counts in equation (2) to the total number of distinct pairs.

The modelling of the 3D correlation function in equation (2) is complicated by the large number of separation bins. Indeed, this requires a very large number of survey realizations to estimate the data covariance matrix. We follow the standard technique of compressing the information contained in the full 3D correlation function $\xi(s)$. In particular, we fit our model to the projected correlation function $w_p(r_p)$ and the first two even multipole moments ξ_ℓ of the redshift space correlation function.

The halo-occupation properties of a given sample affect its intrinsic clustering. Classically, this effect is modelled using the projected correlation function $w_p(r_p)$ that is expected to be free of the apparent RSD effects. The projected correlation function $w_p(r_p)$ is estimated through

$$w_p(r_p) = 2 \int_0^{\pi_{\text{max}}} \xi^s(r_p, \pi) d\pi, \quad (3)$$

where r_p and π are the normal and parallel to the los components of the pair separation s . We limit the integral in equation (3) to a maximum los separation of $\pi_{\text{max}} = 80 h^{-1}\text{Mpc}$, matching the definition in the model to be fitted to these data (Zhai et al. 2019).

Redshift-space distortions change the apparent positions of targets in the radial direction with respect to those in real-space. RSD are classically measured and modelled in the multipole moments ξ_ℓ of the redshift-space correlation function $\xi^s(r_p, \pi)$. Multipole moments ξ_ℓ are defined as

$$\xi_\ell(s) = (2\ell + 1) \int_0^1 \xi^s(s, \mu) L_\ell(\mu) d\mu, \quad (4)$$

with $s = |s|$ and $\mu = \pi/s$ is the cosine of the angle between the los direction and the pair separation vector s and L_ℓ is the ℓ -order Legendre polynomial.

We bin r_p and s in nine logarithmically spaced bins between $0.1-60 h^{-1}\text{Mpc}$, matching the output of AEMULUS predictions for $w_p(r_p)$ and ξ_ℓ , while the los separation π and μ are binned using linear bins of width $\Delta\pi = 1 h^{-1}\text{Mpc}$ and $\Delta\mu = 0.1$. Given the discrete binning of different variables, we estimate the integrals in equations (3) and (4) as Riemann sums.

3.2 PIP correction

In eBOSS spectroscopic observations, fibre-collisions occur whenever two targets are closer than $\theta^{(\text{fc})} = 62''$ on the sky. While a fraction of such collisions are resolved thanks to multiple passes of the instruments in small chunks of the survey, fibre-collisions in single passes remain unresolved and correlate with the underlying target density. If not properly corrected, missed targets due to fibre-collisions can systematically bias the measured two-point correlation function on small scales. In the large-scale analysis of the eBOSS LRG sample (Bautista et al. 2021) fibre-collisions are accounted for by means of the nearest-neighbour (NN) weighting that is quantified through the weight w_{cp} .

In this work, we replace the standard NN correction for fibre-collisions with a more rigorous Pairwise-Inverse-Probability (PIP) weighting (see Bianchi & Percival 2017, for a discussion about inverse-probability estimators). The PIP weights are assigned to pairs of objects in the targeted sample and quantify the probability, for any pair, of being targeted in a random realization of the survey targeting. Under the assumption that no pair has zero probability of being observed, applying the PIP weighting provides statistically unbiased estimates of the two-point correlation function. The selection probabilities are characteristic of the particular fibre assignment algorithm used to select targets from a parent photometric sample for the spectroscopic follow-up. Therefore, these probabilities are extremely difficult to model analytically except for some simple targeting strategies. We infer the selection probabilities by generating multiple replicas of the survey target selection. Details on how these survey realizations are built are provided in Mohammad et al. (2020). Given a set of survey realizations, the inverse probability, or equivalently the PIP weight w_{mn} , is simply the number of realizations in which a given pair could have been targeted divided by the number of times it was targeted. The individual-inverse-probability (IIP) w_m are the single-object counterparts of the PIP weights, i.e. the inverse-probability for a given object m of being targeted in a random survey realization.

PIP weighting assumes that all pairs have a non-zero chance of being observed. However, in eBOSS pairs with separation smaller than the fibre-collision scale $\theta^{(\text{fc})}$ are missed in single-pass areas in all survey realizations. These pairs produce a systematic underestimation in the measured two-point correlation function. For the eBOSS LRG sample, the systematic bias is confined at transverse scales smaller than $r_p^{(\text{fc})} \sim 0.7 h^{-1} \text{Mpc}$ in $w_p(r_p)$ while it spreads to larger separations s in the multipole moments ξ_ℓ . Truncated multipoles $\hat{\xi}^{(\ell)}$ were proposed in Reid et al. (2014) and Mohammad et al. (2016) to remove transverse scales $r_p < r_p^{(\text{fc})}$ from the measured multipole moments, resulting in a loss of information at scales smaller than $r_p^{(\text{fc})}$. Alternatively, the angular up-weighting outlined in Percival & Bianchi (2017) can be used to de-bias the measurements at smaller scales. The angular up-weighting relies on the assumption that pairs missed due to fibre-collisions in the single pass of the instrument are statistically equivalent to those targeted in the multiple-pass areas. The combined PIP and angular up-weighting (PIP + ANG) is

$$\begin{aligned} DD(\vec{s}) &= \sum_{\substack{\vec{x}_m \cdot \vec{x}_n \approx \vec{s} \\ \vec{u}_m \cdot \vec{u}_n \approx \cos \theta}} w_{mn}^{\text{PIP}} w_m^{\text{tot}} w_n^{\text{tot}} \times \frac{DD_{\text{par}}(\theta)}{DD_{\text{fib}}^{\text{PIP}}(\theta)}, \\ DR(\vec{s}) &= \sum_{\substack{\vec{x}_m \cdot \vec{y}_n \approx \vec{s} \\ \vec{u}_m \cdot \vec{v}_n \approx \cos \theta}} w_m^{\text{IIP}} w_m^{\text{tot}} w_n^{\text{tot}} \times \frac{DR_{\text{par}}(\theta)}{DR_{\text{fib}}^{\text{IIP}}(\theta)}, \end{aligned} \quad (5)$$

where $w^{\text{tot}} = w^{\text{sys}} w^{\text{noz}} w^{\text{FKP}}$, and w_{mn}^{PIP} and w_m^{IIP} are PIP and IIP weights, respectively. The fractions on the right-hand side in equation (5) are the angular weights for DD and DR pair counts. An extensive description of different terms in equation (5) is given in Mohammad et al. (2020).

Mohammad et al. (2020) extensively tested the effectiveness of the method of PIP + ANG weighting using a sample of 100 Effective Zel'dovich mocks (EZmocks, Zhao et al. 2021) designed to match the eBOSS LRG sample. The mean of the corrected measurements was compared to the mean of the true clustering of the mocks for ξ_0 , ξ_2 , and w_p over a separation range of 0.1 – $100 h^{-1} \text{Mpc}$ (see figs 9 and 12 of Mohammad et al. 2020). The PIP + ANG correction was able to recover the clustering of the parent sample to within 1σ of the error on the mean at all measurement scales for ξ_0 , and ξ_2 and all

scales of w_p except for the fibre-collision scale, where the corrected measurements recovered the true clustering to within the error on a single mock. We can therefore be confident that the PIP + ANG correction to the eBOSS LRG sample produces unbiased results to within the statistical uncertainty of our sample on all scales.

3.3 AEMULUS cosmological emulator

We compare our measurements to the AEMULUS cosmological emulator (Zhai et al. 2019) predictions for ξ_0 , ξ_2 , and w_p for a galaxy sample in a universe with variable cosmological and galaxy-halo connection parameters. The AEMULUS emulator applies Gaussian process based machine learning to a training set of 40 N -body simulations and that use a latin hypercube to optimally sample a wCDM parameter space spanning the approximate 4σ range of the Planck (Planck Collaboration 2020b) or WMAP (Hinshaw et al. 2013) results (DeRose et al. 2019). A halo occupation distribution model (HOD) is used to connect a galaxy sample to the dark matter haloes. Unlike some galaxy clustering analyses, our emulator does not model ξ_4 , since it is considerably noisier than ξ_0 and ξ_2 . The emulator prediction would likely be noise dominated for ξ_4 , and would require adding more training complexity without a commensurate increase in cosmological information. In their measurement of $f\sigma_8$ from small-scale clustering within the BOSS LOWZ sample, Lange et al. (2022) found that excluding ξ_4 from their analysis of ξ_0 and ξ_2 did not produce a significant change in the best-fitting value or uncertainty.

AEMULUS allows for a flat wCDM described by seven parameters: Ω_M , Ω_b , σ_8 , h , n_s , w , and N_{eff} . For our analysis, we limit the cosmological parameter space by fixing $N_{\text{eff}} = 3.046$ and $w = -1$, since these parameters are not well constrained by our measurements but have been well measured by other probes, resulting in a five parameter flat Λ CDM cosmology. The HOD model used by the AEMULUS allocates a Poisson sampling of $N(M)$ galaxies to haloes of mass M , split into central galaxies and satellites following

$$\langle N(M) \rangle = N_{\text{cen}}(M) + N_{\text{sat}}(M), \quad (6)$$

$$N_{\text{cen}}(M) = \frac{f_{\text{max}}}{2} \left[1 + \text{erf} \left(\frac{\log_{10} M - \log_{10} M_{\text{min}}}{\sigma_{\log M}} \right) \right], \quad (7)$$

$$N_{\text{sat}}(M) = \left(\frac{M}{M_{\text{sat}}} \right)^\alpha \exp \left(-\frac{M_{\text{cut}}}{M} \right) \frac{N_{\text{cen}}(M)}{f_{\text{max}}}, \quad (8)$$

where the free parameters fit by the emulator are f_{max} , $\sigma_{\log M}$, $\log M_{\text{sat}}$, α , $\log M_{\text{cut}}$. Briefly, $\sigma_{\log M}$ defines the width of the transition from a mean occupation of 0–1 for centrals, M_{sat} is the typical mass for haloes to host one satellite, α is the power-law index for the mass dependence of the satellite occupation, and M_{cut} gives an exponential cutoff to the satellite occupation at low mass. M_{min} sets the transition point of the central occupation, and is fixed in the emulator to match the number density of the sample. By matching the number density, we ensure the correct linear bias, thus reducing the degeneracy between the HOD parameters and the growth rate in the correlation function measurements. Because of this choice, we do not use the number density as an observable in our analysis. f_{max} is a new parameter that we add to AEMULUS to address a possible inconsistency between the model and data. eBOSS was targeted using colour and magnitude cuts (Prakash et al. 2016) so it is not a complete sample, whereas the HOD model assumes that all galaxies are included in the sample. This is especially concerning for eBOSS since targets were selected using a lower magnitude limit in the i band to avoid overlap with the CMASS LRG sample (see fig. 1 of Zhai et al. 2017). f_{max} controls the fraction of centrals that are included in the sample, i.e. a value of $f_{\text{max}} < 1$ means that the very massive

haloes do not necessarily host a eBOSS LRG at the centre. While these targeting cuts would be expected to affect the completeness of both centrals and satellites, for satellites f_{\max} is completely degenerate with $M_{\text{sat}}^{-\alpha}$ (see similar discussion in Lange et al. 2022). Since these HOD parameters are primarily nuisance parameters in our constraint of the growth rate, we do not apply f_{\max} to the satellites. In Section 4.2, we perform a series of tests to determine the effect of excluding f_{\max} on the measured $f\sigma_8$.

The emulator also allows three additional parameters that control how galaxies are distributed in their host haloes: c_{vir} , v_{bc} , and v_{bs} (labelled η_{con} , η_{vc} , and η_{versus} in Zhai et al. 2019). c_{vir} is the ratio between the concentration parameters of the satellites to the host halo where the halo is assumed to have a Navarro–Frenk–White (NFW) profile (Navarro, Frenk & White 1996). v_{bc} and v_{bs} are the velocity biases of centrals and satellites respectively, where $\sigma_{\text{gal}} = v_{\text{gal}}\sigma_{\text{halo}}$ and σ_{halo} is the velocity dispersion of the halo calculated from its mass. Finally, the AEMULUS emulator uses a 15th parameter, γ_f , which rescales all halo bulk velocities in the simulation. The galaxy velocity can therefore be thought of as the sum of two components: a component equal to the bulk motion of the host halo scaled by γ_f , and a randomly directed component that depends on the halo mass through the velocity dispersion and that is scaled by either v_{bc} or v_{bs} for centrals and satellites, respectively. For a detailed description of the AEMULUS correlation function parameters see Zhai et al. (2019). See Section 3.7 for a description of how we treat these parameters in our fit.

The original AEMULUS emulator was trained to match a BOSS CMASS-like sample at $z = 0.57$ and space density $n = 4.2 \times 10^{-4} [h^{-1}\text{Mpc}]^{-3}$. However, our eBOSS sample is at an effective redshift of $z = 0.737$ and peak number density of $n = 9 \times 10^{-5}$. The difference in number density is particularly worrying, since a less dense sample will preferentially fill more massive haloes. The result will be a sample with a larger linear bias, which is degenerate with the growth rate in clustering measurements. In order to ensure an unbiased result, we rebuild the emulator from the original simulations, but using the $z = 0.7$ simulation time-slice and adjusting HOD parameters, especially M_{\min} , to match the eBOSS number density. The training ranges for the new emulator are given in Table 1.

3.4 Interpreting growth rate measurements

As shown in Reid et al. (2014), which used a similar parametrization to measure RSD from their simulations, in the linear regime a fractional change in γ_f is proportional to a fractional change in f , such that $f = \gamma_f f_{\Lambda\text{CDM}}$, where $f_{\Lambda\text{CDM}}$ is the linear growth rate for a flat ΛCDM cosmology specified by the model parameters. However, the link between the linear velocity power spectrum amplitude and the non-linear regime is possibly scale dependent i.e. a linear response on large scales might not necessarily lead to a linear response on small scales. γ_f is introduced in the simulations as a scaling of all velocities by the same amount and so γ_f also scales the non-linear velocities of haloes. In this case, γ_f still provides a consistency test with the amplitude of the velocity field expected in a ΛCDM universe with the model cosmology, where $\gamma_f = 1$ indicates agreement, but it no longer necessarily gives a pure rescaling of the linear growth rate. For models that do have such a linear response, then the measurement of γ_f over the full range of scales can be used to constrain the linear growth rate. However, as this is model dependent, we conservatively separate the contributions of the linear and non-linear regime in presenting our results (as described in Section 4.1).

Table 1. All model parameters divided into cosmological and HOD parameters, with the training range used by the AEMULUS emulator and the prior range used in the MCMC fit. Prior ranges were chosen to be slightly larger than the original training ranges, except where excluded by the physical meaning of the parameter, in order to be able to identify if the fit converges outside of the training range. The purpose of this extended range is only to more easily identify a prior dominated fit, since the emulator is not expected to produce accurate clustering outside of the training range. Instead, it would regress to the mean prediction. The exception is $\log M_{\text{cut}}$, where the prior excludes the lower part of training range since $\log M_{\text{cut}}$ ceases to have any impact on the halo occupation if it is below $\log M_{\min}$. This is the case for the eBOSS LRG sample, so $\log M_{\text{cut}}$ is poorly constrained. However, we found the chains tended to pile up at the lower end of the training range, which gave the misleading impression that the data strongly preferred the lowest possible value, although it had no effect on the cosmological constraints. For that reason, we set a more reasonable lower limit on $\log M_{\text{sat}}$ for our sample.

Parameter	Training range	Prior range
Ω_m	[0.255, 0.353]	[0.225, 0.375]
$\Omega_b h^2$	[0.039, 0.062]	[0.005, 0.1]
σ_8	[0.575, 0.964]	[0.5, 1]
h	[0.612, 0.748]	[0.58, 0.78]
n_s	[0.928, 0.997]	[0.8, 1.2]
N_{eff}	[2.62, 4.28]	3.046
w	[-1.40, -0.57]	-1
$\log M_{\text{sat}}$	[14.0, 16.0]	[13.8, 16.2]
α	[0.2, 2.0]	[0.1, 2.2]
$\log M_{\text{cut}}$	[10.0, 13.7]	[11.5, 14]
$\sigma_{\log M}$	[0.1, 1.6]	[0.08, 1.7]
v_{bc}	[0, 0.7]	[0, 0.85]
v_{bs}	[0.2, 2.0]	[0.1, 2.2]
c_{vir}	[0.2, 2.0]	[0.1, 2.2]
γ_f	[0.5, 1.5]	[0.25, 1.75]
f_{\max}	[0.1, 1]	[0.1, 1]

Although the AEMULUS code uses γ_f to adjust the RSD amplitude in the model, the RSD are sensitive to the parameter combination $f\sigma_8$. We therefore present our large-scale results in terms of $f\sigma_8 = \gamma_f f_{\Lambda\text{CDM}}\sigma_8$, which is used in the remainder of the paper and the abstract. It is also important to note that we calculate $f_{\Lambda\text{CDM}}\sigma_8$ from the model cosmology according to linear theory, rather than the value that would be obtained from the power spectrum on scales corresponding to $0.1\text{--}60 h^{-1}\text{Mpc}$. Thus, the value of $f\sigma_8$ we present is the value expected from linear theory for our model, and is directly comparable to measurements made on larger scales. However, care should be taken when using the resulting measurements of $f\sigma_8$ to constrain models where the other parameters deviate significantly from flat ΛCDM and general relativity ($\Lambda\text{CDM} + \text{GR}$, hereafter used interchangeably with ΛCDM). A problem inherent in many cosmological measurements and all previous RSD measurements is that one assumes various features of a particular model, here flat ΛCDM , in order to make the measurements. To test a different model, one should strictly have to perform a new fit including all properties of that model. This does not affect the validity of our measurement as a test of consistency with ΛCDM within the parameter space of the emulator, or as an indication of how the RSD measurements compare to those from other surveys.

3.5 Covariance matrix

Clustering measurements in different separation bins are correlated, and we need an estimate of the covariance matrix when fitting a model to the observations. Mock surveys, either based on the output

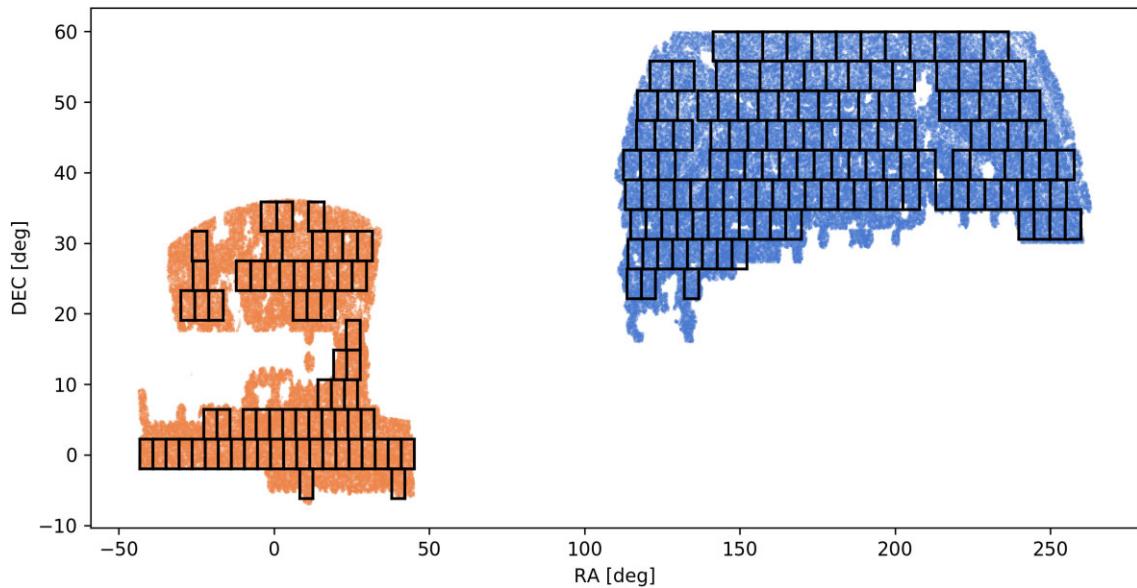


Figure 2. The footprint of the eBOSS LRG clustering catalogue with our jackknife regions. The blue points show the North Galactic Cap (NGC) observations, while the orange points show the South Galactic Cap (SGC) observations. It should be noted that the square jackknife regions all have approximately equal area on the sky, however due to the distortion of projecting a sphere on to a plane, the regions at larger declination appear wider in this plot.

of N -body simulations or approximate methods, have been widely used to estimate the data covariance matrix. However, in order to work on small scales, we would need a large number of simulations that accurately reproduce the small-scale clustering – a difficult task. In order to generate a covariance matrix that reflects the small-scale clustering of our sample, we instead use jackknife sampling. We split our survey footprint into equal area squares on the sky using right ascension (RA) and declination (Dec.) cuts. This method relies on the clustering of the sample being uncorrelated with position in the survey. Furthermore, because we expect the covariance to follow a simple volume scaling, we remove the squares with the smallest occupation as determined from the random catalogue over the survey footprint, so that each region included will contribute approximately the same statistical weight to the sampling (Fig. 2). Since the measurements from each sample are normalized, it is not necessary that they contain identical numbers of objects, however selecting regions in this way reduces variance from regions at the edge of the survey which are only partially filled or have peculiar geometries. The missing area is included in the final calculation by means of a volume-weighted correction.

For the objects in our data and random catalogues that are located within one of the 200 accepted regions, we store a region identification number. We then recalculate the monopole and quadrupole of the 3D correlation function and projected correlation function for this reduced sample 200 times, excluding one region from the calculation each time. We include the full PIP+ANG weighting scheme in these calculations, so that the variance in the PIP + ANG weights is included in the jackknife estimation. The covariance matrix is then estimated from this jackknife sampling using

$$C_{i,j} = \frac{n-1}{n} \sum_k (\xi_{i,k} - \bar{\xi}_i) (\xi_{j,k} - \bar{\xi}_j), \quad (9)$$

where the i, j indices are over the elements of the data vector, $n = 200$ is the number of jackknife regions, and k is an index over the jackknife realizations.

In order to more easily visualize the correlations between bins, we calculate the correlation matrix by

$$R_{i,j} = \frac{C_{i,j}}{(C_{i,i} C_{j,j})^{1/2}}. \quad (10)$$

The correlation matrix is highly diagonal, which is expected since we have a small number of widely separated bins, which are only expected to be weakly correlated. In order to reduce the noise in the off-diagonal terms, we smooth the correlation matrix using diagonally adjacent bins. Each off-diagonal element is assigned the average of itself and the two adjacent diagonal elements, excluding bins from other measurements. The result of this diagonal smoothing is shown in Fig. 3.

In addition to the data error, we include the emulator error in the covariance matrix. The emulator error is calculated as a fractional error on each correlation function bin using a sample of test HOD parameter sets which are selected from the same parameter ranges as the training sample, but were not used in the training (Zhai et al. 2019). The fractional error is converted to an absolute error, σ_E , by multiplying by the correlation function measurements from the data. The total variance for each measurement bin is then calculated from $\sigma_T^2 = \sigma_D^2 + \sigma_E^2$. In order to preserve the structure of the jackknife covariance matrix, we convert the smoothed correlation matrix back to the covariance matrix using $C_{i,i} = \sigma_{T,i}^2$. The contributions of the data and emulator errors to the total error are shown in Fig. 4. The data error is dominant in the region $s < 5 h^{-1}$ Mpc for the monopole and projected correlation function, while the emulator error is comparable for $s > 5 h^{-1}$ Mpc and across the full separation range of the quadrupole.

We also correct the inverse covariance matrix according to Hartlap, Simon & Schneider (2007), using

$$\hat{C}^{-1} = \frac{n-p-2}{n-1} C^{-1}, \quad (11)$$

where $n = 200$ is the number of jackknife regions, and $p = 27$ is the number of combined bins in our three measurements. Although n should properly be the number of completely independent

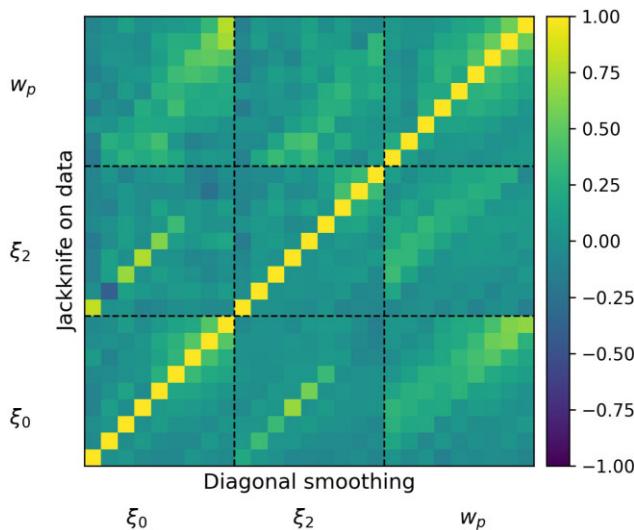


Figure 3. Comparison of the unsmoothed and smoothed correlation matrices. The upper diagonal elements correspond to the unsmoothed jackknife correlation matrix, while the lower diagonal elements show the result of our diagonal smoothing method.

measurements (Eifler, Kilbinger & Schneider 2008; Krause et al. 2013), we follow Reid et al. (2014) in using the number of regions, noting that this correction may therefore underestimate the true size of the effect. However, this factor has very little effect on our final fit, as well as not changing the best-fitting value.

3.6 AP scaling

Although we fit the AEMULUS correlation function predictions directly to our measurements from the data, our results are still affected by the Alcock–Paczynski (AP) effect (Alcock & Paczynski 1979), because we convert the data redshift to distance assuming a fixed fiducial cosmological model. We therefore need to scale the separations between model and data to account for the difference in comoving distance between our fiducial cosmology and the cosmology of the model. We apply the standard AP scaling from Bautista et al. (2021) to each model, first defining the perpendicular and parallel dilation factors

$$\alpha_{\perp} = \frac{D_M(z_{\text{eff}})}{D_M^{\text{fid}}(z_{\text{eff}})}, \quad \alpha_{\parallel} = \frac{D_H(z_{\text{eff}})}{D_H^{\text{fid}}(z_{\text{eff}})}, \quad (12)$$

where D_M is the comoving angular diameter distance, D_H is Hubble distance. We then scale the multipole moments of the correlation function as follows

$$\xi_0^{\text{fid}}(r^{\text{fid}}) = \xi_0(\alpha r) + \frac{2}{5}\epsilon \left[3\xi_2(\alpha r) + \frac{d\xi_2(\alpha r)}{d\ln(r)} \right], \quad (13)$$

$$\xi_2^{\text{fid}}(r^{\text{fid}}) = \left(1 + \frac{6}{7}\epsilon\right)\xi_2(\alpha r) + 2\epsilon \frac{d\xi_0(\alpha r)}{d\ln(r)} + \frac{4}{7}\epsilon \frac{d\xi_2(\alpha r)}{d\ln(r)}. \quad (14)$$

where $\alpha = \alpha_{\parallel}^{1/3} \alpha_{\perp}^{2/3}$ and $\epsilon = (\alpha_{\parallel}/\alpha_{\perp})^{1/3} - 1$. Once we have shifted the model, we used a cubic spline interpolation to recover the model values at the fiducial separations used to calculate the data values.

The projected correlation function was calculated similarly using the scaling

$$w_p^{\text{fid}}(r_p^{\text{fid}}) = w_p(\alpha_{\perp} r_p). \quad (15)$$

The accuracy of this method depends, in part, on the width of the bins used due to the calculation of the derivative and the interpolation

between points. In order to assess the importance of these factors, we perform an additional fit to the data without the AP correction. See Section 5.6 for details.

3.7 Exploring the likelihood

We assume our correlation function measurements are drawn from a multivariate Gaussian distribution, and use uniform priors for all model parameters, given in Table 1. We explore the posterior surface for the fit between data and the AEMULUS correlation function predictions using a Markov chain Monte Carlo (MCMC) sampler within the Cobaya² framework (Torrado & Lewis 2021). We include the full AEMULUS HOD parameter space in our fit, however, we limit the wCDM cosmological parameter space by fixing $N_{\text{eff}} = 3.046$ and $w = -1$, since these parameters are not well constrained by our measurements but have been well measured by other probes.

A concern for our small-scale analysis is that the separation range we use lacks a distinctive feature with a known scale to constrain the cosmological parameters, such as the BAO bump in large-scale analyses. Consequently, we consider a number of additional cosmological priors in order to set an accurate cosmology for our analyses. To begin with, we apply a uniform prior on the cosmological parameters based on the distance in 7D cosmological parameter space between the chain point and the cosmologies of the AEMULUS simulations used to train the emulator. If the distance is above a certain threshold the proposed step is forbidden, thus restricting the parameter space to the region which is well sampled by the training data, rather than the full uniform prior range given in Table 1. In practice, the main impact of the training prior is to add the restriction $\sigma_8 > 0.65$, since there is only one training cosmology with σ_8 below that range.

We also consider jointly fitting our data with the Planck 2018 TT,TE,EE and lensing likelihoods (Planck Collaboration 2020a, c) using the CAMB cosmological Boltzmann code (Lewis, Challinor & Lasenby 2000; Howlett et al. 2012), which constrain the cosmological parameters that control the shape of the power spectrum. It is important to note that γ_f is treated as a free parameter in addition to the standard cosmological parameters, and is only constrained by RSD as measured from the eBOSS data. In effect, it represents a consistency check between the large-scale structure and CMB data: if these are consistent, we expect that γ_f is close to one. We further consider three cases of the joint eBOSS and Planck fit. The first is a simple joint fit, where all of the cosmological parameters, including σ_8 , and jointly fit by both the eBOSS clustering measurements through AEMULUS and the Planck likelihoods, while the HOD parameters and γ_f are fit solely by the clustering measurements. The second is similar, except we explicitly account for the slight redshift offset between the emulator ($z = 0.7$) and the data ($z = 0.737$). The emulator takes all cosmological parameters at $z = 0$, so the shape of the linear power spectrum will be identical between the cosmology described by the Planck likelihoods and the emulator, however there will be a difference in amplitude due to the slight redshift offset.

²Cobaya, a code for *bayesian analysis in cosmology*, is the PYTHON successor to CosmoMC. Users are able to use the same MCMC sampler as CosmoMC (Lewis & Bridle 2002; Lewis 2013) in a PYTHON framework, while allowing access to likelihoods from many major cosmological datasets. The sampler is tailored for parameter spaces with a speed hierarchy and implements the ‘fast dragging’ procedure described in Neal (2005). See <https://cobaya.readthedocs.io> for details.

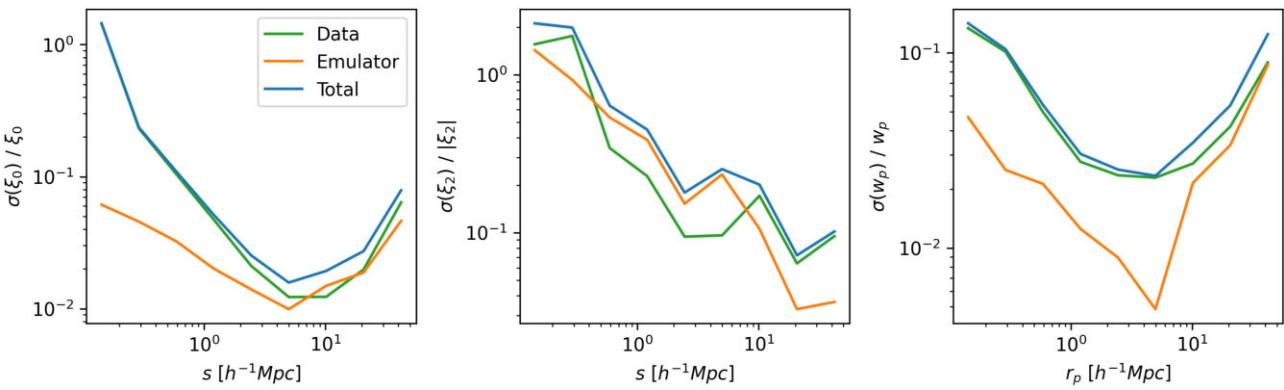


Figure 4. The contributions of the data error calculated through jackknife sampling (green), the emulator error (orange), and total error (blue), for the monopole, quadrupole, and projected correlation function (left- to right-hand panel).

Therefore, we adjust the value of σ_8 given to AEMULUS as follows

$$\sigma_{8,Aem} = \sigma_8(z=0) \times \frac{D(z=0.737)}{D(z=0)} \times \frac{D(z=0)}{D(z=0.7)}. \quad (16)$$

This makes sure that the normalization of the AEMULUS output matches that expected at $z=0.737$ in the cosmology being tested: the first ratio corrects from $z=0$ to 0.737 in the cosmology being tested, and the second ratio corrects from $z=0.7$ to 0 , where the normalization is defined by AEMULUS. Thirdly, we consider a joint fit where the Planck likelihoods are used to constrain all of the cosmological parameters except for σ_8 , which is fit solely by the clustering data. We test the robustness of our results to the inclusion of the training prior and the Planck likelihoods through these three methods in Section 5.3. Based on the results of these tests we use the training prior but not the Planck likelihoods for our fiducial analysis.

4 ROBUSTNESS AND SYSTEMATIC ERROR CHECKS

In this section, we explore the robustness of our model in general and to several possible sources of systematic error in particular. We begin by assessing the impact of non-linear velocities on our measurements, and what information is included from different scales. We then perform a general check of our method by fitting to measurements made on a mock catalogue. Finally, we check the impact of the two possible discrepancies between our model and the data, the effects of galaxy selection on the completeness of the HOD model, and redshift uncertainty.

4.1 Contribution of non-linear velocities

In Section 3.4, we introduced the key parameter of our measurement, γ_f , and described its significance on linear and non-linear scales. In order to identify the transition between these regimes, we examine how the emulator prediction changes for various values of γ_f , shown in Fig. 5. For the three largest bins, varying γ_f produces an almost constant relative change in the monopole, with a larger growth rate giving a larger clustering amplitude, as expected from linear theory. In the middle three bins, the effect on the monopole changes signs as the quasi-linear regime transitions to the non-linear regime, where the random virial motions of the haloes begin to dominate and increasing γ_f , which rescales all halo velocities, begins to damp the clustering. In the three smallest bins, the effect of γ_f on the monopole begins to decrease as the one-halo term begins to dominate. Because γ_f

affects only the halo velocities, and in our HOD formalism, we do not assign galaxies based on subhalos, varying γ_f has no effect on the one-halo term. Motivated by this result, we divide our nine measurement bins into three groups of three bins, with individual ranges of $0.1-0.8 h^{-1}\text{Mpc}$, $0.8-7 h^{-1}\text{Mpc}$, and $7-60 h^{-1}\text{Mpc}$. These three ranges correspond roughly to the strongly non-linear regime where the one-halo term is dominant, the transition between the non-linear and quasi-linear regimes, and the quasi-linear regime. We therefore restrict our measurement of $f\sigma_8$ to the quasi-linear regime, where γ_f can be interpreted as a rescaling of the linear growth rate. For measurements performed over the full separation range, we instead use γ_f as a test of ΛCDM , where a deviation from $\gamma_f=1$ indicates that the velocity field of the data as parametrized by our emulator model is in disagreement with the expectation from ΛCDM .

4.2 Galaxy selection and the HOD model

As described in Section 3.3, we add an additional parameter f_{\max} to the emulator compared to previous uses that controls the maximum occupation fraction of central galaxies in the HOD framework, in order to address the incompleteness of the eBOSS LRG sample due to target selection. We test the necessity of this addition and the effect on the clustering using a series of HOD mock galaxy catalogues. We constructed these mocks from the Uchuu³ simulation. Briefly, Uchuu is a $(2000 h^{-1}\text{Mpc})^3$, 12800^3 particle simulation using the Planck2015 cosmology and a mass resolution of $m_p = 3.27 \times 10^8 h^{-1} M_\odot$. We construct the mocks from the $z=0.7$ slice, using the halotoools⁴ (Hearin et al. 2017) PYTHON package and an HOD parametrization identical to that outlined in Section 3.3. We constructed mocks using $\sigma_{\log M}$, $\log M_{\text{sat}}$, α , and $\log M_{\text{cut}}$ from five randomly selected test HOD parameter sets in AEMULUS, with $\log M_{\min}$ tuned to give $n = 1 \times 10^{-4}$. The AEMULUS test HOD sets are themselves randomly selected from the uniform training range given in Table 1, but were not used in training the emulator. In all mocks, we kept the additional parameters $v_{\text{bc}} = 0$, $v_{\text{bs}} = 1$, $c_{\text{vir}} = 1$, and $\gamma_f = 1$ fixed to their simplest, no scaling values. For each of the five HOD parameter sets, we then constructed five mocks with $f_{\max} = [0.2, 0.4, 0.6, 0.8, 1.0]$, for a total of 25 mocks.

We fit these 25 HOD mocks using two emulators: one matching the original AEMULUS HOD model that is equivalent to fixing $f_{\max} =$

³<http://skiesanduniverses.org/Simulations/Uchuu>

⁴<https://halotools.readthedocs.io/en/latest/>

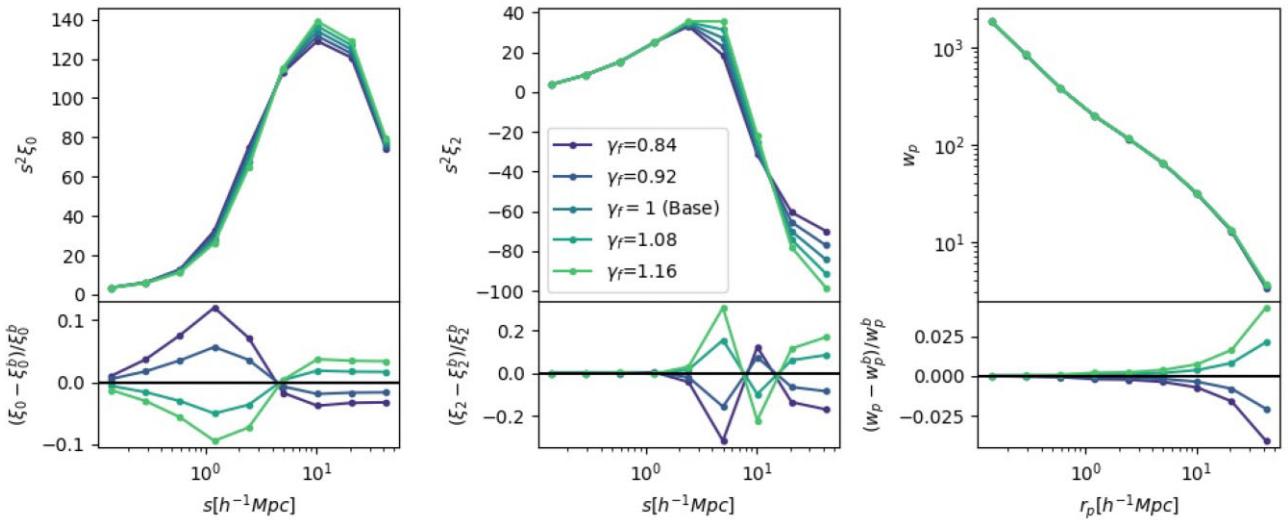


Figure 5. The effect on the emulator prediction of varying γ_f for the monopole (left-hand panel), quadrupole (centre hand panel), and projected correlation function (right-hand panel). All other parameters are kept fixed at reasonable values for the baseline eBOSS fit. Upper panels: Direct comparison of the predictions, ranging from low γ_f (blue) to high γ_f (red). Lower panels: Relative difference to the $\gamma_f = 1$ prediction.

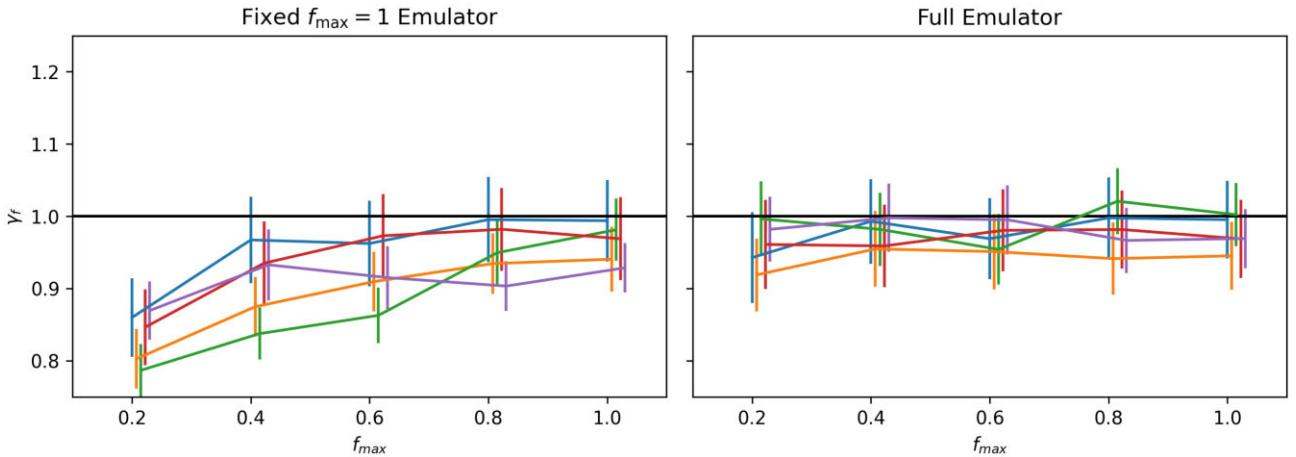


Figure 6. Performance of emulators with fixed or variable f_{\max} on HOD mocks constructed with varying f_{\max} . The left-hand panel shows the results from an emulator built with the original AEMULUS parameter set, which is equivalent to $f_{\max} = 1$. The right-hand panel shows the results from the emulator used in our analysis with variable f_{\max} . Both emulators were built to match the eBOSS redshift and number density. The horizontal line shows the expected value of γ_f used to construct the mocks. Points are shifted slightly along the x -axis to avoid overlap.

1, and the full emulator with variable f_{\max} . Both emulators were built to match the eBOSS redshift and number density, as described in Section 3.3. The γ_f constraints on the HOD mocks from both emulators are shown in Fig. 6, where the expected value is $\gamma_f = 1$ by the construction of the mocks. It should be noted that all of the mocks were constructed using the same halo catalogue from a single simulation box at a particular cosmology, so it is unsurprising that the constraints do not scatter evenly above and below $\gamma_f = 1$, since they are not fully independent. The key points to notice are that the variable f_{\max} emulator is able to recover the expected value of γ_f within the uncertainty over the full f_{\max} range, and shows no trend in f_{\max} . Conversely, the fixed f_{\max} emulator shows a clear bias in γ_f for $f_{\max} \leq 0.6$. This result matches what we would theoretically expect for model which overestimates the f_{\max} value of the sample. If the mismatch is small, there is not a significant change in the galaxy bias of the sample, however if f_{\max} is significantly overestimated then the model prediction has a larger galaxy bias, b , than the sample, which

is compensated by a lower growth rate since the amplitude of the linear clustering scales as fb^2 .

4.3 Redshift uncertainty

Another area of concern where the emulation based model may not accurately reflect the data is the effect of redshift uncertainties. As shown in Fig. 2 of Ross et al. (2020), the eBOSS LRG sample has a redshift uncertainty that is well approximated by a Gaussian with mean $\mu = 1.3 \text{ km s}^{-1}$ and standard deviation $\sigma = 91.8 \text{ km s}^{-1}$. On average, this means that each redshift is wrong by an absolute offset of 65.6 km s^{-1} . To first order this gives a Gaussian random velocity shift for all targets, which acts to damp the clustering of the multipoles on small scales. The parameters v_{bc} and v_{bs} , which control the velocity dispersion of centrals and satellites, respectively, should be able to mimic much of this effect in the model without affecting the constraints on other parameters. However, since γ_f scales all

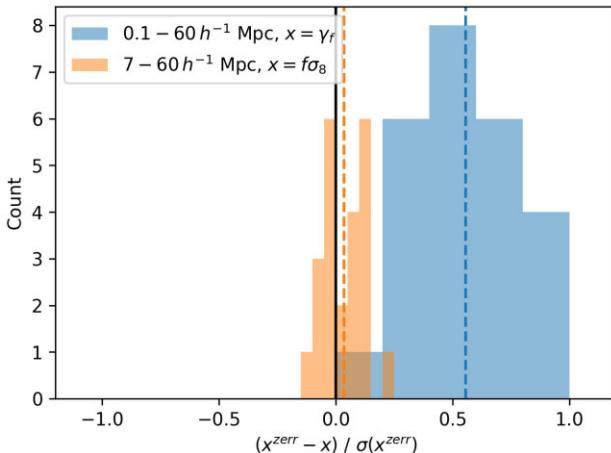


Figure 7. A histogram of the shifts in the measured cosmological parameters for 25 HOD mocks with and without a random velocity dispersion matching the eBOSS redshift uncertainty. Blue bars show the shift in γ_f measured over the full separation range, and orange bars show the shift in $f\sigma_8$ measured from the quasi-linear scales only. The x -axis shows the difference between the value measured for the mock with a random velocity dispersion ($zerr$) and the value measured from the same mock without the additional velocity dispersion, divided by the uncertainty of the measurement from the $zerr$ mock. Coloured dashed lines show the mean shift for each case. For the fit over $0.1-60 h^{-1} \text{Mpc}$, including a random velocity dispersion not represented in the model increased the measured value of γ_f for all 25 mocks, with a mean shift slightly larger than half of the statistical uncertainty. Conversely, for the fit over $7-60 h^{-1} \text{Mpc}$, the shifts from including a random velocity dispersion scatter around 0, with a mean shift that is negligible compared to the statistical error.

halo velocities in the simulation, on non-linear scales where the halo velocities are virialized, γ_f has a similar effect on the clustering as the redshift uncertainty, v_{bc} and v_{bs} . In addition, v_{bc} and v_{bs} are both calculated by scaling the virial dispersion of the host halo, so the galaxy velocities derived in the model have a mass dependence which is not reflected in the redshift uncertainty. The result is that the redshift uncertainty may bias the recovered value of γ_f on non-linear scales, with an unmodelled redshift uncertainty giving a larger than expected value of γ_f .

We test the effect of the redshift uncertainty on the γ_f and $f\sigma_8$ constraints using a second set of HOD mocks, constructed in the same way as those described in Section 4.2. We selected 25 new AEMULUS test HOD parameter sets and generated HOD catalogues using `halotools`. We then calculated the clustering with and without a random velocity shift along the los drawn from a Gaussian with mean $\mu = 1.3 \text{ km s}^{-1}$ and standard deviation $\sigma = 91.8 \text{ km s}^{-1}$. The change in the measured values of γ_f from the full separation range and $f\sigma_8$ from the quasi-linear scales only (matching the method used for our baseline results) due to the inclusion of the random velocity shift are shown in Fig. 7. For all 25 mocks, including a random velocity shift increased the value of γ_f measured from the full separation range, with an average shift slightly greater than half the statistical uncertainty. The larger value of γ_f measured due to the random velocity shift matches our theoretical expectation for the degeneracy between γ_f and the redshift uncertainty on non-linear scales, and the magnitude of the shift indicates that the redshift uncertainty is a significant concern when fitting to the non-linear scales. On the other hand, the shifts in the measured value of $f\sigma_8$ scatter around 0, with a mean shift over an order of magnitude smaller than the statistical uncertainty. This result also agrees with

what is expected for our model, since on quasi-linear scales the redshift uncertainty is not degenerate with a change in γ_f , and instead will change only v_{bc} and v_{bs} . Therefore, the redshift uncertainty is not a concern for our value of $f\sigma_8$ measured from the quasi-linear scales.

There are several barriers to including a correction for the redshift uncertainty in the model. Most significantly, the redshift uncertainty grows with redshift (see fig. 6 of Bolton et al. 2012 for BOSS redshift evolution), while the emulator is constructed from catalogues at a single redshift slice. The evolution with redshift is also important because the eBOSS LRG targeting cuts were made using the apparent magnitudes of the targets, so properties of the sample such as the mean mass will also evolve weakly with redshift and correlate with the growth of the redshift uncertainty. The result is that including the redshift uncertainty in the model may not be as simple as drawing from a uniform velocity shift, and would require more detailed testing and corrections. The effect of redshift uncertainty could instead be included as an additional systematic error or shift in our measured values. However, it is important to note that for every mock tested, the inclusion of redshift uncertainty (without it being present in the model) increased the measured value of γ_f , because on the non-linear scales where the redshift uncertainty is the most significant, it is degenerate with the larger random motions of the haloes provided by a larger value of γ_f . In Section 5, we consistently measure values of γ_f that are below the value expected from $\Lambda\text{CDM} + \text{Planck}2018$, so the presence of redshift uncertainty is actually expected to increase this tension rather than lowering it. We therefore take the conservative approach of excluding a shift in our measurements due to the redshift uncertainty, even though it would be expected to increase the tension shown by our measurements, and leave a complete treatment of the redshift uncertainty to future work.

4.4 SHAM mocks

We test the robustness of our model and analysis pipeline using a subhalo abundance matching (SHAM) mock generated from the Uchuu simulation. By using a SHAM mock rather than a HOD mock, we remove the dependence on the specific galaxy-halo connection model used in our analysis, providing the best approximation to a model independent test. If our analysis is able to correctly recover the expected value of $\gamma_f = 1$ for the SHAM mock, then we can be confident it will be able to match the data, even if there are deviations from the specific functional form of the galaxy-halo connection model described in Section 3.3. We use the $z = 0.7$ slice of the simulation to construct a SHAM mock using the peak halo velocity, V_{peak} , with a scatter of 0.2 dex, and a number density of $n = 1 \times 10^{-4}$ in order to match the eBOSS LRG number density and redshift.

The result of our fit to the SHAM mock is shown in Fig. 8. The primary purpose of the Uchuu SHAM mock test is to assess the robustness of the cosmological parameter recovery using our HOD based emulator, so we have only included the parameters which have the greatest impact on the γ_f constraint. The constraints on all of the cosmological parameters are in good agreement with the known values from the simulation, and the 1D marginalized constraint on γ_f is $\gamma_f = 0.964 \pm 0.049$, which agrees to within $1-\sigma$ with the known value of $\gamma_f = 1$ for the mock. All well constrained HOD parameters converge within the training parameter space indicating that the emulator is able to accurately model the clustering of the mock, despite the mock being constructed using a different galaxy-halo connection. This result shows that our analysis pipeline and model provide robust constraints on the growth rate.

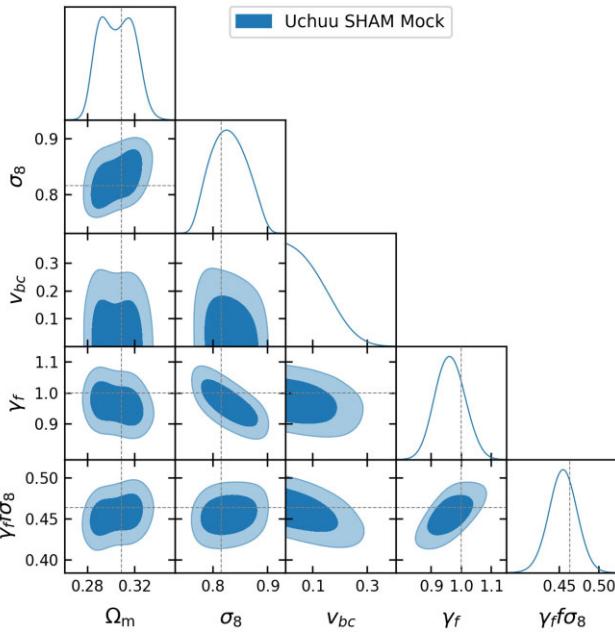


Figure 8. Two dimensional and one dimensional marginalized constraints of the key parameters from the fit to an Uchuu SHAM mock matching the eBOSS LRG number density and redshift. Dotted lines show the values of the cosmological parameters from the simulation.

5 RESULTS

In this section, we present the results of our fit to the small-scale LRG clustering. We also investigate the robustness of our results by testing the inclusion of additional constraints on the cosmological parameters, examining how the constraints change depending on which scales and measurements are included in the analysis, the effect of covariance matrix smoothing on the measured parameters, and consistency with the constraints from a combined CMASS + eBOSS sample.

5.1 Headline results

We fit the eBOSS LRG monopole, quadrupole, and projected correlation function over scales $0.1 < r < 60 h^{-1}$ Mpc using the Cobaya MCMC sampler. We restrict the cosmological parameter space using the AEMULUS training prior described in Section 3.7, but do not include any external data. We obtain a value of $\gamma_f = 0.767 \pm 0.052$, 4.5σ below what would be expected in a Λ CDM + GR universe. The 1D and 2D likelihood contours of the full parameter set are shown in Fig. 9. All well constrained parameters are within the prior ranges described in Table 1, and the parameters that are most impactful for our results, Ω_m , σ_8 , v_{bc} , and γ_f , all show roughly Gaussian constraints. The best-fitting values of the cosmological parameters other than γ_f are consistent with recent measurements from the Planck Collaboration (Planck Collaboration 2020b). The best-fit model prediction is plotted relative to the data in Fig. 10, showing reasonable agreement within the measurement uncertainty on all scales. The best-fitting prediction has $\chi^2 = 14.1$, with 14 degrees of freedom and 27 data points, indicating a good fit. In addition, we consider a fit over only the quasi-linear scales of our measurements, $7-60 h^{-1}$ Mpc as described in Section 4.1, from which we obtain a value of $f\sigma_8(z = 0.737) = 0.408 \pm 0.038$. This value is 1.4σ below what is expected from the 2018 Planck data for a flat Λ CDM universe, and is a factor of 1.7 improvement in statistical

error over the more standard large-scale analysis of the same data set. See Section 5.4 for more details.

5.2 Testing the quasi-linear scales for overfitting

One concern for our fit to the quasi-linear scales is that by reducing the separation range to $7-60 h^{-1}$ Mpc, we are fitting nine data points with a 14 free parameter model. However, it is important to note that many of the HOD parameters have a negligible effect on these scales. In particular, the three parameters that control the satellite occupation ($\log M_{\text{sat}}$, α , $\log M_{\text{cut}}$) and the three parameters that control the positions of galaxies in the haloes (v_{bc} , v_{bs} , c_{vir}) have very little impact and are almost entirely constrained by the $0.1-7 h^{-1}$ Mpc bins. Therefore, while there are 14 free parameters in the model, only eight are significant when fitting to the nine bins of the quasi-linear scales. While this provides a theoretical explanation for why the quasi-linear scales will not be overfit, our fit over the scales $7-60 h^{-1}$ Mpc has a minimum $\chi^2 = 0.36$ (Table 2), indicating that the small scale HOD parameters may still be causing some overfitting.

To test if this overfitting affects our results, we perform additional fits over the $7-60 h^{-1}$ Mpc separation range with the predominantly small scale HOD parameters fixed to their best-fitting values from the fit over the full $0.1-60 h^{-1}$ Mpc separation range. In the first additional fit we keep the six parameters listed above fixed, leaving eight parameters (Ω_m , $\Omega_b h^2$, σ_8 , h , n_s , $\sigma_{\log M}$, γ_f , f_{max}) free. In the second fit, we also keep $\sigma_{\log M}$ and f_{max} fixed to their best-fitting values from the full fit, allowing only the six cosmological parameters to vary. The γ_f constraints from these fits are shown in Table 2 and Fig. 11. The results of both fits show that reducing the parameter space increases the precision of the γ_f constraint without significantly shifting the central value, while increasing the minimum χ^2 . We conclude that allowing the small scale HOD parameters to be free does lead to the quasi-linear scales being overfit, however, it does not bias our cosmological constraints and instead only increases the uncertainty. Fixing these HOD parameters would increase the precision of our measurement from the quasi-linear scales, but it would also introduce an indirect dependence on the non-linear scales. We therefore take the conservative choice of using the measurement with all 14 parameters free as our baseline result. However, this test does show the value of including the non-linear scales in a measurement of the linear growth rate.

5.3 Testing the impact of the cosmological priors

We consider a number of prior constraints on the cosmological parameters, as described in Section 3.7. The three most significant cases are a uniform prior as described in Table 1, a uniform prior that restricts the cosmological parameters to be within the volume that is well sampled by the training simulations, and a joint fit with Planck2018 likelihoods with a scaled value of σ_8 to account for the redshift difference between the data and the model. The constraints on the key parameters for these three prior choices are shown in Fig. 12. The parameter that is most significantly impacted by the prior choice is σ_8 , with all three methods giving consistent values but with large differences in precision. However, the constraint on $f\sigma_8$ is almost unchanged for all prior choices. This result clearly shows the robustness of the $f\sigma_8$ fit from the data, and demonstrates the freedom of the model where changes in σ_8 can be balanced by γ_f . It is also important to note that because the uncertainty on $f\sigma_8$ is dominated by the uncertainty of γ_f that the training prior and the joint fit with Planck achieve almost the same precision on $f\sigma_8$, despite

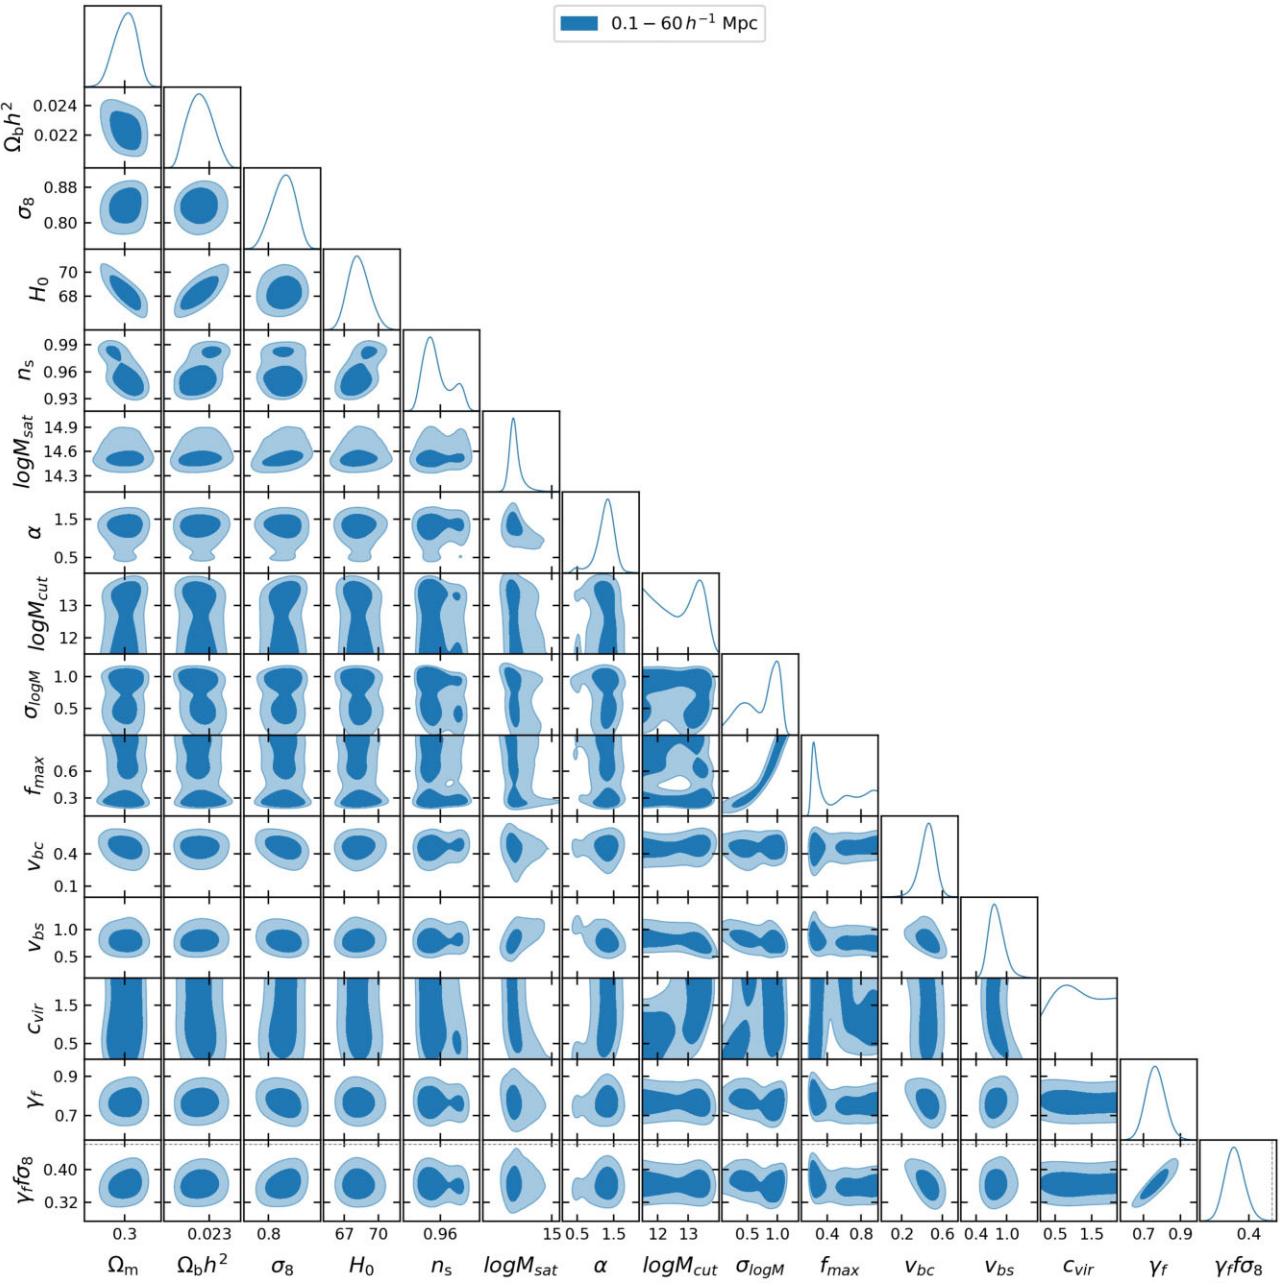


Figure 9. One dimensional and two dimensional contours of the parameters used in our baseline fit, as well as the derived constraints on $\gamma_f \sigma_8$.

having comparable constraints on γ_f but a significant difference in precision on σ_8 .

The effect of the three treatments of σ_8 for the joint Planck fit described in Section 3.7 can be found in Table 2. Using the same value of σ_8 for the Planck chains and model, scaling to account for the redshift offset, or excluding the Planck constraints on σ_8 all give consistent values for the growth rate, again demonstrating the robustness of the fit.

5.4 Testing the dependence on the data fitted

In order to test the consistency of the constraint on γ_f from the different regimes described in Section 4.1, we fit to the full non-

linear regime ($0.1-7 h^{-1} \text{Mpc}$), the weakly non-linear and quasi-linear regimes ($0.8-60 h^{-1} \text{Mpc}$), and the quasi-linear regime only ($7-60 h^{-1} \text{Mpc}$). One dimensional and two dimensional contours in the $v_{bc} - \gamma_f$ parameter space for these three fits are shown in the left-hand panel of Fig. 13. There is little variation in the other parameters between these fits to different scales, however some important insight is gained from examining the $v_{bc} - \gamma_f$ degeneracy since both parameters have a similar effect on the clustering in the non-linear regime. The fits to smaller scales yield larger and more precise values of v_{bc} , while obtaining smaller and less precise constraints on γ_f . The full fit to all scales is located at the intersection in $v_{bc} - \gamma_f$ space of the small and larger scale fits. The result is that there is mild tension between the constraints on small and large scales, although the significance when considering the combined uncertainty is less

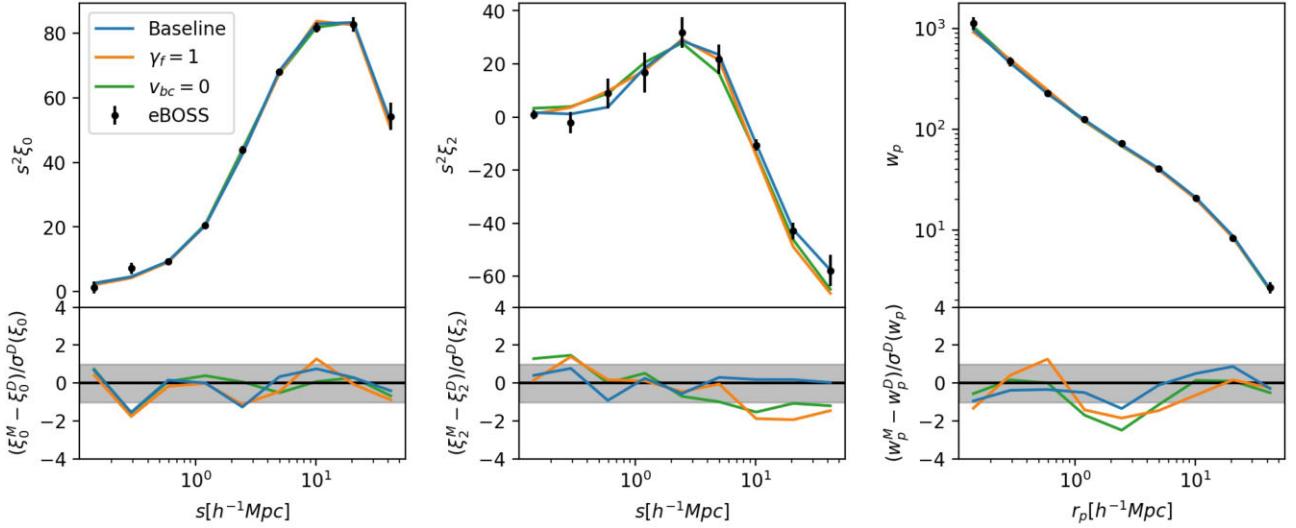


Figure 10. Comparison of the best-fitting model predictions to the data several fit to the eBOSS LRG sample for the monopole (left-hand panel), quadrupole (centre panel) and projected correlation function (right-hand panel). Upper panels: the baseline fit (blue), fixed $\gamma_f = 1$ fit (orange), and $v_{bc} = 0$ fit (green), with the data and measurement uncertainty (black). Lower panels: The difference between the best-fitting models and the data divided by the measurement uncertainty. The $1 - \sigma$ region is shown in grey.

Table 2. γ_f constraints with statistical errors calculated from the width of the 1D marginalized posterior and χ^2 values for the fits used in our analysis. N_P gives the number of free model parameters in the fit and N_D gives the number of data points. *The eBOSS + Planck18 runs jointly fit 5 of the 14 parameters with Planck, so they are not fully independent.

Run	γ_f	N_P	N_D	χ^2
0.1–60 $h^{-1}\text{Mpc}$	0.767 ± 0.052	14	27	14.1
0.1–7 $h^{-1}\text{Mpc}$	0.71 ± 0.14	14	18	7.8
0.8–60 $h^{-1}\text{Mpc}$	0.783 ± 0.066	14	18	4.2
7–60 $h^{-1}\text{Mpc}$	0.854 ± 0.083	14	9	0.36
7–60 $h^{-1}\text{Mpc}$, eight parameters	0.821 ± 0.064	8	9	0.74
7–60 $h^{-1}\text{Mpc}$, six parameters	0.802 ± 0.050	6	9	1.8
$\xi_0 + \xi_2$	0.819 ± 0.073	14	18	5.0
$\xi_0 + w_p$	0.65 ± 0.11	14	18	5.4
$\gamma_f = 1$	1	13	27	28.0
$v_{bc} = 0$	0.958 ± 0.088	13	27	22.5
$f_{\text{max}} = 1$	0.764 ± 0.051	13	27	16.6
Unsmoothed covariance matrix	0.767 ± 0.052	14	27	14.3
Scaled mock covariance matrix	0.766 ± 0.059	14	27	12.0
No training prior	0.85 ± 0.12	14	27	12.1
eBOSS + Planck18	0.784 ± 0.048	14*	27	18.5
eBOSS + Planck18 scaled σ_8	0.798 ± 0.047	14*	27	19.1
eBOSS + Planck18 free σ_8	0.766 ± 0.053	14*	27	18.0
No AP scaling	0.772 ± 0.053	14	27	14.5

than $1 - \sigma$. It is worth recalling that since γ_f rescales all halo velocities in the simulation, in the linear regime it can be used to derive a constraint on the linear growth rate $f\sigma_8$, in the non-linear it also enhances the effects of non-linear growth. So the fit to the small-scales is really a consistency check between the data and model with ΛCDM , and these results showing that there is a strong tension which is most significant in the non-linear regime.

The fit to the quasi-linear scales only does not show the same degeneracy between v_{bc} and γ_f since they no longer have the same effect on the clustering, and is broadly consistent with any value of

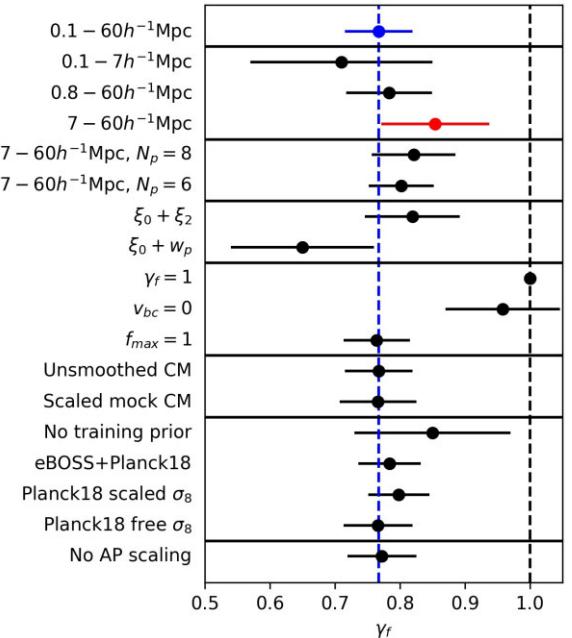


Figure 11. γ_f constraints from all the runs listed in Table 2. The blue point shows the baseline fit to the full separation range, extended by the blue dashed line for comparison to other points. The red point shows the fit to the quasi-linear scales only. The black dashed line shows $\gamma_f = 1$ for comparison, the value expected if the amplitude of the halo velocity field matches the ΛCDM expectation.

v_{bc} since it ceases to be impactful on such large scales. However, the large scale fit is still able to recover a relatively tight constraint on γ_f that can be compared directly to the linear growth rate, giving a measurement $f\sigma_8 = 0.408 \pm 0.038$, which is 1.4σ lower than the value expected from the 2018 Planck data for a flat ΛCDM model.

We also examine the effect of excluding certain measurements from the fit. In the right-hand panel of Fig. 13, we show the constraints in $v_{bc} - \gamma_f$ parameter space from the joint fit to only the monopole and

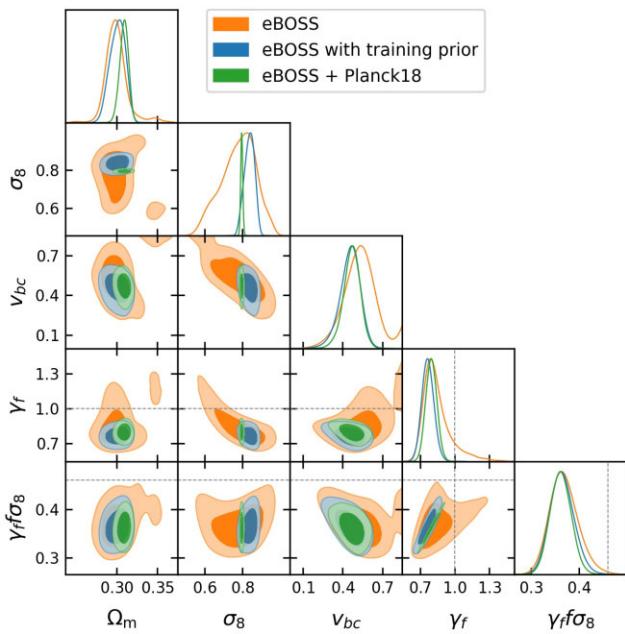


Figure 12. One dimensional and two dimensional contours of the key fit parameters for the fit to the eBOSS LRG sample with no additional cosmological constraints (orange), restricted by the AEMULUS training prior (blue), and jointly fit with the Planck2018 likelihoods (green).

projected correlation function, and the joint fit to the multipoles only. The multipole only fit is less sensitive to the degeneracy between v_{bc} and γ_f , but prefers a smaller value of v_{bc} and larger γ_f compared to the full fit. On the other hand, the joint fit of the monopole and projected correlation function, which contain similar clustering information but are sensitive and insensitive to the effects of RSD respectively, prefer a non-zero value of v_{bc} with much greater confidence, compensated by a low but less well constrained value of γ_f . As with the fits to different scales, the full fit lies in the overlap region produced by the different sensitivities of these measurements.

5.5 Testing the dependence on the covariance matrix

We test the robustness of our covariance matrix smoothing by fitting to the unsmoothed covariance matrix and a scaled version of the covariance matrix estimated from 1000 EZmocks. These mocks are constructed to match the clustering of the eBOSS DR16 samples on mildly non-linear and linear scales, but are not matched on small-scales, where the mocks exhibit very different clustering from the data. EZmocks are based on a Gaussian random field in a $5 h^{-3} \text{Gpc}^3$ box and an initial power spectrum and geometry of a flat ΛCDM cosmology with parameters $\Omega_m = 0.307115$, $\Omega_b = 0.048206$, $h = 0.6777$, $\sigma_8 = 0.8225$, $n_s = 0.9611$. Matter particles are displaced from their initial to final positions using the Zel'dovich approximation. Tracer bias relation is calibrated to match the observed clustering of the target sample in the eBOSS DR16 data. The linear component of the redshift-space distortions is imprinted using Zel'dovich approximation while the non-linear term is modelled through an isotropic Gaussian motion. Mocks are then trimmed to match the geometry and radial selection function of the eBOSS DR16 LRG catalogue. The unscaled mock covariance matrix displays a similar correlation structure to the covariance matrix calculated from applying jackknife to the data, however because the clustering of the mocks on scales below $\sim 1 h^{-1} \text{Mpc}$ is significantly

lower than the data, the mock covariance matrix underestimates the variance on those scales. To bring the mock covariance matrix into better agreement we calculate the correlation matrix from the mocks, and then convert the correlation matrix to the covariance matrix by scaling the original diagonal values of the mock covariance matrix according to

$$\sigma_{i,i}^{M,s} = \sigma_{i,i}^M \frac{\xi_i^D}{\xi_i^M}, \quad (17)$$

where ξ^D is the correlation function from the data and ξ^M is the mean correlation function from the 1000 EZmocks. This scaling preserves the original correlation structure and $\sigma(\xi)/\xi$ ratio of the mock covariance matrix, but adjusts for the higher clustering of the data. However, this method results in a very large variance for the quadrupole bins because the the mean quadrupole of the mocks goes to 0 on small scales. In order to prevent this artificial inflation of the quadrupole bins, we instead use $\sigma_{i,i}^{M,s} = \sigma_{i,i}^D$ for the quadrupole.

The results of the fits using this scaled mock covariance matrix and the original unsmoothed jackknife covariance are shown in Table 2. The constraints in both cases are nearly identical to our baseline fit using the smoothed jackknife covariance matrix, indicating that our analysis is robust to the choice of covariance matrix.

5.6 Testing the dependence on AP correction

We test the dependence of our result on the AP correction by running a full fit excluding the AP correction. The constraint on γ_f from this fit can be seen in Table 2 and Fig. 11. Excluding the AP correction has a negligible effect on the constraint on γ_f and slightly increases the best-fitting χ^2 . We therefore conclude that any uncertainty in the AP correction due to the large bin width and approximate calculation will not have a significant effect on our cosmological constraints.

5.7 Including the BOSS CMASS data

We test the reliability of our fit using a combined CMASS + eBOSS sample in the redshift range $0.6 \leq z \leq 0.8$. In particular, in our analysis we use the CMASS sample from the DR12 data release. The CMASS DR12 catalogue covers an area of 9376 deg^2 over a redshift range of $0.4 < z < 0.8$ (Reid et al. 2016) with a target density of 99.5 deg^{-2} . The target selection is calibrated to provide a sample of galaxies with approximately constant stellar mass over the spanned redshift range. We refer the reader to Reid et al. (2016) for a detailed description of the target selection and properties for CMASS sample. In order to perform a joint measurement of the two-point correlation function using the eBOSS and CMASS catalogues, we restrict the two samples (and the corresponding random catalogues) only to the area of the sky where they overlap and to the redshift range of $0.6 < z < 0.8$. The redshift distributions of the two samples as well as their joint distribution are shown in Fig. 1.

The advantage of this sample is that it is more complete due to the complimentary nature of the CMASS and eBOSS colour cuts. However, the inclusion of the additional CMASS objects skews the redshift distribution of the sample, which is not ideal for an HOD-based analysis where the galaxy–halo connection parameters are implicitly assumed to be the same across the full redshift range of the sample, and several are dependent on the density of galaxies. As such, we use our combined CMASS + eBOSS measurement to provide a consistency check with our fit, particularly our assumption that the target selection of eBOSS does not affect our measurement, but we continue to use the eBOSS only constraint as our fiducial measurement.

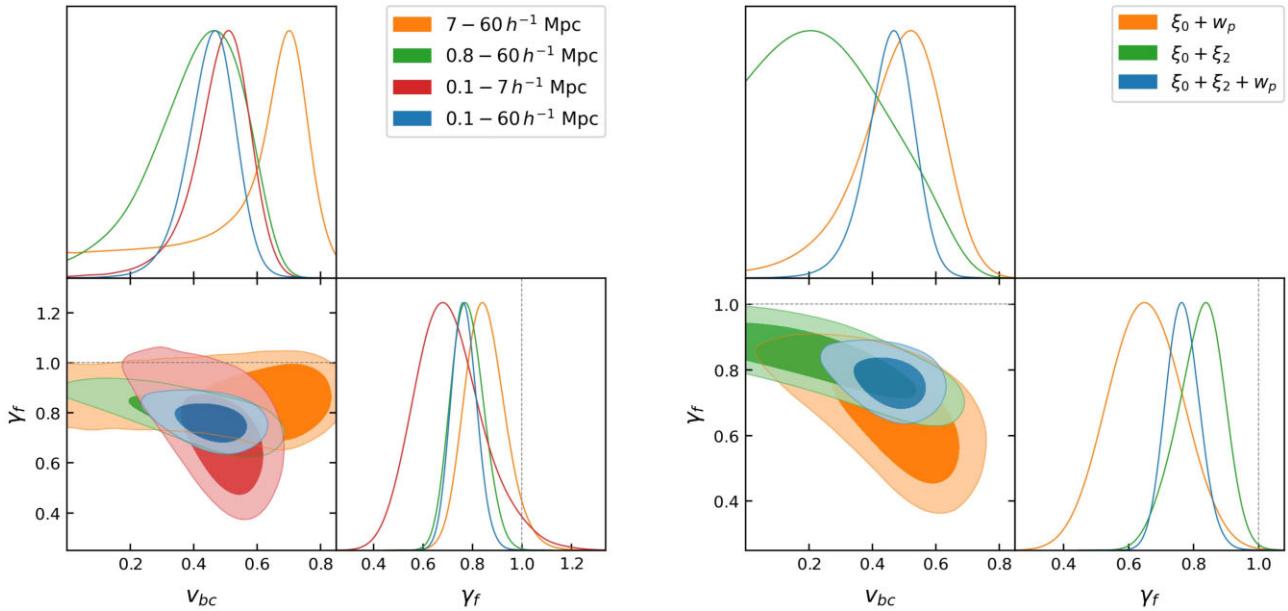


Figure 13. Two dimensional and one dimensional marginalized constraints on v_{bc} and γ_f for fits to different scales and measurements. Left-hand panel: constraints from the three largest separation bins (orange), six largest separation bins (green), six smallest separation bins (red), and all nine separation bins (blue) for all three measurements. The dotted line shows $\gamma_f = 1$, the value expected if the amplitude of the halo velocity field matches the expectation from Λ CDM. Right-hand panel: constraints from the joint fit to the monopole and projected correlation function (orange), monopole and quadrupole (green), and all three measurements (blue).

To correct fibre-collisions in the CMASS sample, we use a modified version of the NN upweighting with completeness correction, designated CP, as described in section 2.3 of Mohammad et al. (2020), and the standard angular upweighting method described in Section 3.2. For the eBOSS LRG sample, the CP correction was found to perform similarly to the PIP only result on all scales of w_p , ξ_0 , and ξ_2 (see figs 15 and 18 of Mohammad et al. 2020). Given the similarities in sample type and targeting between CMASS and eBOSS, it is reasonable to expect a similar result for CMASS. When combined with angular upweighting, any systematic bias is expected to be below the statistical uncertainty of the measurement. Since our primary goal in analyzing the combined CMASS + eBOSS sample is as a consistency check, this correction is sufficient for our purposes.

Fig. 14 shows the result of our fit compared to the eBOSS only fit in the most important parameters of our analysis for both the full emulator range and the quasi-linear scales only. The CMASS + eBOSS measurement is consistent with the eBOSS only measurement in all parameters, although there is a greater preference for larger f_{\max} values, as expected. It is interesting to note that in the fit over the full emulator range the inclusion of the CMASS data does not affect our γ_f constraint, including not reducing the 1D marginalized uncertainty. However, there are several reasons why including additional data may not reduce 1D marginalized constraints. First, the additional data may reduce the allowed parameter space in 14 dimensions without affecting the 1D constraints on a specific parameter. Additionally, the uncertainty in our measurement is limited by the emulator accuracy in several bins, notably the quadrupole and the large scale bins of the monopole and w_p , so a reduction of measurement uncertainty in these bins will not be reflected in the fit. Finally, the constraint on γ_f seems to rely on the complimentary constraining of different scales and probes on parameter combinations such as v_{bc} and γ_f (Fig. 13). The fit to CMASS + eBOSS has slightly less tension between the small and

large scales than the eBOSS only measurement, so the overlap region remains the same size even though the uncertainty from separated scales has been reduced. This can be seen in the fit to the quasi-linear scales, where the combined CMASS + eBOSS sample gives a constraint of $f\sigma_8 = 0.384 \pm 0.036$. This constraint is consistent with the eBOSS only measurement from the quasi-linear scales, but because it is slightly lower, it is in less tension with the fit over the full separation range.

6 DISCUSSION

6.1 Comparison to other measurements

We compare our result to other measurements of $f\sigma_8$ from galaxy clustering surveys in Fig. 15. Taken as a whole sample, there is clearly good consistency with the Λ CDM prediction. For the eBOSS LRGs, Bautista et al. (2021) analyzed pairs with separations between $25-130 h^{-1}$ Mpc, and obtained measurements of $f\sigma_8 = 0.446 \pm 0.066$ and 0.420 ± 0.065 depending on the RSD model used in the analysis (see Table B1 of Bautista et al. 2021). Our measurement is consistent with these results at around the $1 - \sigma$ level, but has a factor of 1.7 improvement in the statistical error. Our measurement also continues the trend of galaxy clustering measurements of $f\sigma_8$ falling slightly below the prediction from observations of the CMB.

In Fig. 15, we also compare our results to other attempts to measure $f\sigma_8$ on small scales. Reid et al. (2014) used a similar parametrization as our analysis to measure $f\sigma_8$ from the small-scale clustering of the BOSS CMASS sample, and achieved the highest precision to date. However, due to the difficulty of modelling the non-linear regime Reid et al. (2014) used a fixed cosmology, which has been shown by Zhai et al. (2019) to significantly reduce the uncertainty. Conversely, Lange et al. (2022) use a novel modelling method in their analysis of the BOSS LOWZ sample that does not require an emulator. It should also be noted that their model does not include

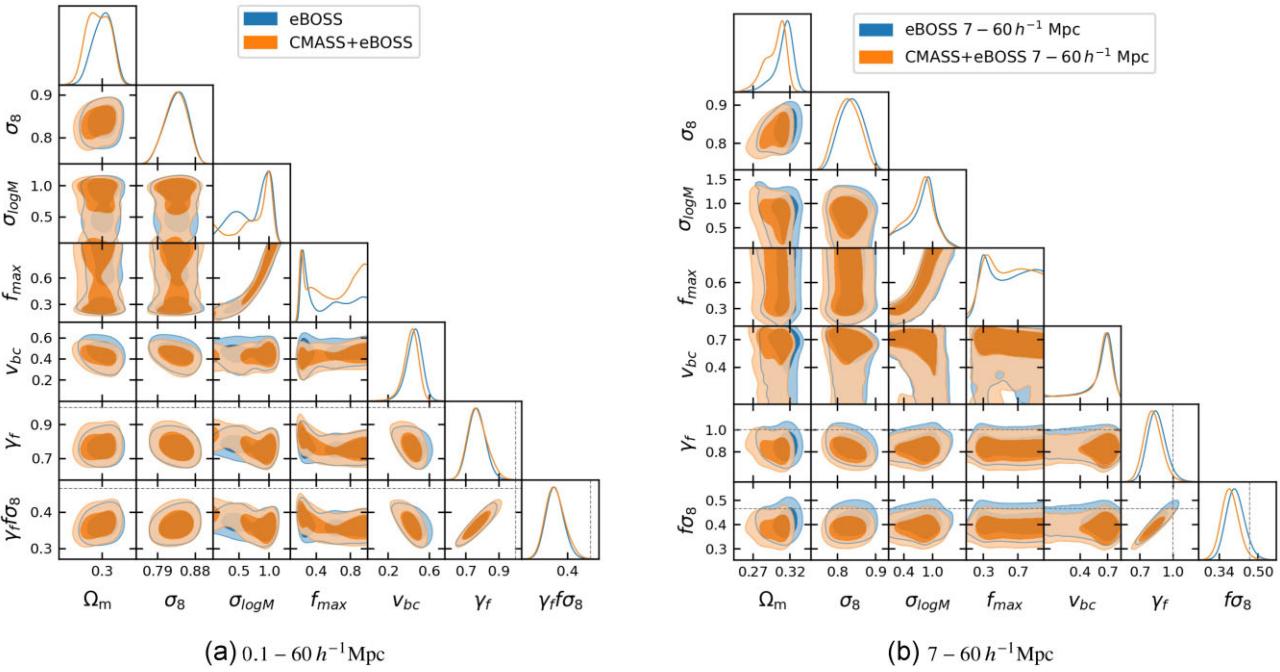


Figure 14. Two dimensional and one dimensional marginalized constraints of the key parameters of our fit for our fiducial eBOSS measurement (blue) and combined CMASS + eBOSS sample (orange). The two plots show (a) the fit over the full emulator range, and (b) the fit to the quasi-linear scales only.

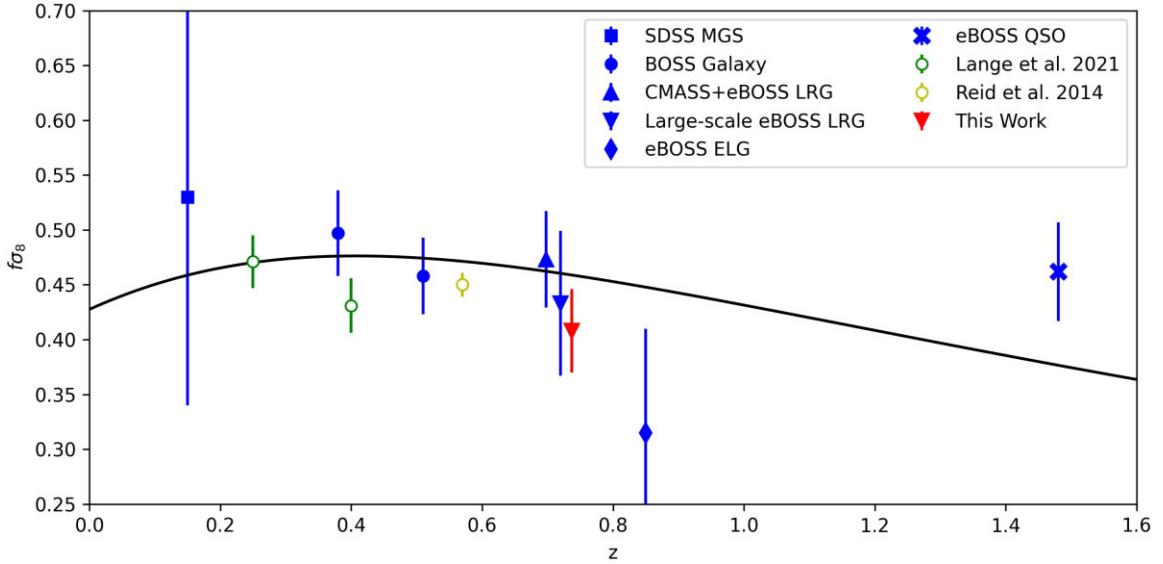


Figure 15. $f\sigma_8$ measurements from various SDSS samples. The blue points show the results of the standard large-scale analyses from the SDSS MGS (Howlett et al. 2015), BOSS galaxies (Alam et al. 2017), CMASS + eBOSS LRGs, eBOSS LRGs (Bautista et al. 2021), eBOSS ELGs (de Mattia et al. 2021), and eBOSS quasars (Neveux et al. 2020). Our small-scale analysis of the eBOSS LRGs using only the quasi-linear regimes is shown in red. Empty coloured points show the results of small-scale analyses from the BOSS LOWZ sample (Lange et al. 2022, green) and BOSS CMASS sample (Reid et al. 2014, yellow) that included non-linear scales in the analysis. The black line shows the expected value of $f\sigma_8$ for a flat Λ CDM universe with the best-fitting Planck2018 cosmology. The large-scale eBOSS LRG result is shifted in the x -axis to avoid overlap with the small-scale result from this work.

an equivalent of our γ_f parameter that allows the linear growth rate to change independently of the Λ CDM cosmology. Both of these analyses have split in linear and non-linear regimes differently than our analysis, which significantly affects the claimed uncertainty. By restricting our measurement of $f\sigma_8$ to only the quasi-linear scales, our uncertainty increases by a factor of ~ 1.5 compared to our fit over the full 0.1 – $60 h^{-1}$ Mpc separation range, however, we can be confident that what we are measuring is purely the linear growth

rate, and so can be directly compared to other more standard large-scale measurements. As shown in Sections 5.1 and 5.4, using the full separation range significantly increases the tension with the result expect for Λ CDM, with the non-linear scales in greater disagreement with the expected value than the quasi-linear scales, however it is no longer clear if this tension arises from a discrepancy in the linear growth rate or a difference in the non-linear velocity field measured in the data using the emulator model.

It is interesting to note that Lange et al. (2022) found a similar dependence on the measurement scales, with smaller scales preferring a smaller value of $f\sigma_8$. Lange et al. (2022) also found that adding the projected correlation function to their fiducial measurement of the monopole, quadrupole, and hexadecapole reduced the best-fitting value of their lower redshift sample by $1 - \sigma$, but did not significantly affect the measurement from their higher redshift sample. Differences between the two analysis methods mean it is expected that there would be some variation in the impact of the different measurements and scales between our results. This is particularly true since Lange et al. (2022) do not include a parameter comparable to our γ_f , given the importance of w_p in breaking the $v_{bc} - \gamma_f$ degeneracy in our analysis.

6.2 Galaxy–halo connection parameters

The parameter found to be most degenerate with our γ_f constraint is v_{bc} , the scaling of the velocity dispersion of centrals in the HOD framework (Fig. 9). A lower value of v_{bc} corresponds to a larger γ_f , as expected in the non-linear regime since both parameters increase the observed velocity dispersion of galaxies (see Section 4.3). Our fit over the full 0.1 – $60 h^{-1}$ Mpc separation range strongly prefers a non-zero v_{bc} and low γ_f . However, our fit to the quasi-linear regime finds no discernible degeneracy between v_{bc} and γ_f and recovers both a relatively large value of γ_f and non-zero value of v_{bc} , although the constraint on v_{bc} is weak to the small impact it has on those scales (Fig. 13). This result indicates that the degeneracy between v_{bc} and $f\sigma_8$ may illustrate the degree to which the non-linear scales affect the overall constraint. Lange et al. (2022) also find a strong degeneracy between the velocity scaling of central galaxies and their constraint on $f\sigma_8$, with their higher redshift sample yielding $v_{bc} > 0$ and low $f\sigma_8$ compared to the Λ CDM prediction. Reid et al. (2014) elected to fix the velocity of centrals to match that of the host halo, and find closer agreement with the Λ CDM expectation, which we also find when using a fixed $v_{bc} = 0$. $v_{bc} > 0$ indicates that a central galaxy is in motion relative to the centre of the host halo, either because the central galaxy is oscillating in the potential or because the system is not fully relaxed. Understanding the physical processes that would lead to this effect, especially if the process is redshift dependent, will be important for future analyses.

We also investigate the dependence of our measurement on the f_{\max} parameter. Due to the strong degeneracy between $\sigma_{\log M}$ and f_{\max} , our fit to the data is broadly consistent with a wide range of values for f_{\max} between 0.2 and 1, however there is a large peak at $f_{\max} = 0.25$. A low value of f_{\max} is not surprising for the eBOSS sample given the magnitude and colour cuts made when selecting the target sample, particularly since the highest magnitude objects were removed. We do not find a degeneracy with $f\sigma_8$, so the lack of constraint on $\sigma_{\log M}$ and f_{\max} is not expected to bias our measurement.

Numerical simulations have shown that the clustering of dark matter haloes can depend on properties other than halo mass, a.k.a halo assembly bias (Sheth & Tormen 2004; Gao, Springel & White 2005; Harker et al. 2006; Wechsler et al. 2006; Obuljen, Dalal & Percival 2019). This bias can propagate into the distribution of galaxies that live in these haloes and thus introduce additional bias in the clustering measurement. In the analysis of BOSS galaxies over a wider redshift range Zhai et al. (2022), we enhance the basic HOD approach used here with an assembly bias model depending on the environment of dark matter haloes. Although the results of that analysis imply the mild existence of assembly bias, there is a negligible impact on the cosmological constraint and measurement of structure growth rate. Therefore, we exclude explicit modelling of assembly bias in this paper.

6.3 Comparison to tension from lensing surveys

It is interesting to note that we obtain a lower value of $f\sigma_8$ than expected from Planck measurements, given the current σ_8 -tension between Planck and weak lensing surveys and the low amplitude of the galaxy–galaxy lensing amplitude measured using the BOSS CMASS sample by Leauthaud et al. (2017), since both tensions could be resolved by a lower value of σ_8 than that measured by Planck. To see approximately how our result might relate to this tension, we compare the constraints on $S_8 = \sigma_8(\Omega_M/0.3)^{0.5}$ for the DES Y1 results (Abbott et al. 2018), Planck 2018 (Planck Collaboration 2020b), and our results (Fig. 16). The left-hand panel shows our measurement using the full separation range, while the right-hand panel shows our measurement from the quasi-linear scales only. Our constraint, shown as the blue contour, is consistent with both the DES Y1 and Planck results in both cases. However, it is important to note that our low value of $f\sigma_8$ comes almost entirely from $\gamma_f < 1$, which reduces the magnitude of peculiar velocities in the simulation without affecting the amplitude of fluctuations, σ_8 . If the low value of $f\sigma_8$ we measure was due to the value of σ_8 instead then the constraint would shift down the S_8 axis, shown as a green contour. For our measurement from the quasi-linear scales this shift maintains consistency with both DES Y1 and Planck 2018, however for our fit to all scales this shift puts the green constraint in tension with the Planck results, and in more mild disagreement with the DES results. This result may indicate that the increased tension we find from the non-linear scales may be caused by an issue with the HOD model, rather than a purely cosmological tension.

6.4 Emulator robustness and potential improvements

We have performed rigorous tests of the emulator performance (see Section 4), and found that the model performs well when fit to an independent simulation and galaxy–halo connection prescription. We also find that a model that assumes all central galaxies are observed leads to a systematic bias in the recovered cosmological parameters if the actual fractional occupation of centrals is lower than 0.6. We correct this bias by adding the parameter f_{\max} to the emulator, and verify that the full emulator gives an unbiased measurement for $0.2 < f_{\max} < 1.0$. We also identify the redshift uncertainty as a source of systematic bias on non-linear scales, with a redshift uncertainty missing from the model leading to an offset in γ_f to larger values by more than half of the statistical error for the eBOSS sample. This is a significant concern for future small-scale analyses, and will require careful attention due to the difficulties in implementing a redshift dependent effect in a model constructed at a single redshift. The redshift uncertainty has also been found to scale with redshift, so it will be an even greater concern for future large surveys at high redshift such as DESI (DESI Collaboration 2016a, b) and Euclid (Laureijs et al. 2011).

Our measurement of the clustering within the eBOSS LRG sample also meets or exceeds the emulator precision in several of the measurement bins (see Fig. 4), showing the importance of improving the model precision for future surveys. This must be balanced against ensuring there are sufficient bins to yield a well defined fit, given the number of model parameters (see Section 5.2). Finally, careful attention must be given to the non-linear scales, and identifying what information can be used to constrain the linear growth rate. A key aspect includes ensuring the performance of the HOD model on these scales, and investigating the effect of baryonic physics.

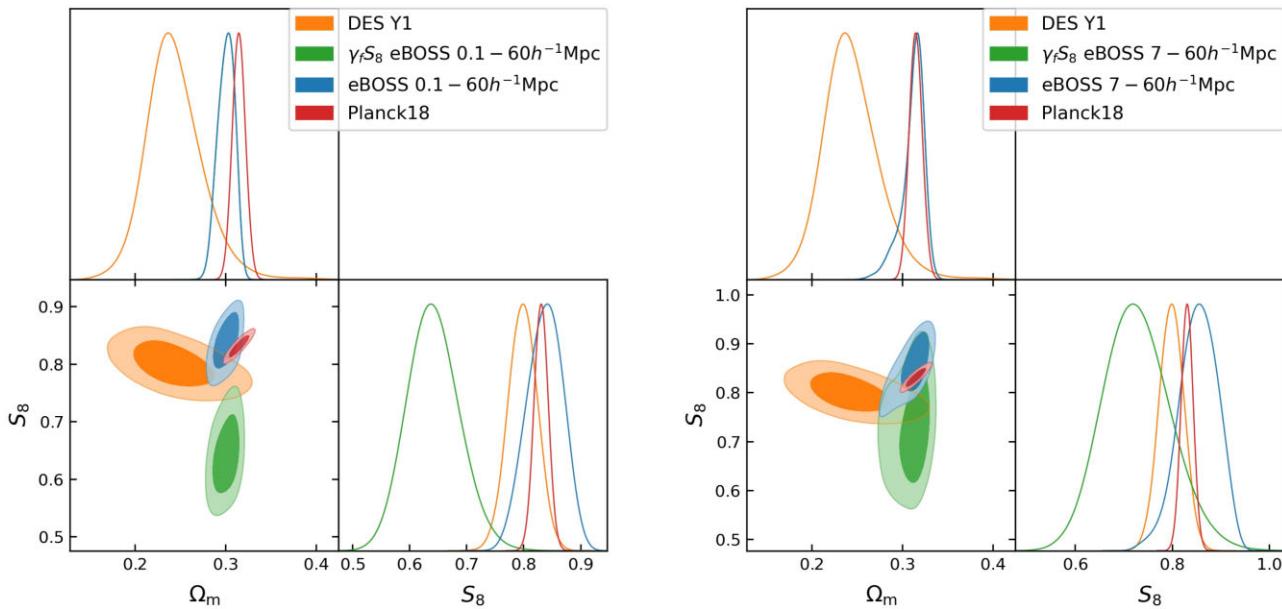


Figure 16. Two dimensional and one dimensional marginalized constraints on Ω_m and S_8 from our analysis (blue), the Dark Energy Survey (DES) year 1 results (Abbott et al. 2018) (orange) and Planck 2018 results (Planck Collaboration 2020a, c) (red). Since our low value of $f\sigma_8$ mostly comes from γ_f , we also plot $\gamma_f S_8$ (green) for our fit, which shows the constraint we would have if the low value of $f\sigma_8$ came entirely from the σ_8 value. Left-hand panel: results of our full fit to all scales. Right-hand panel: results from only the quasi-linear scales used to constrain the linear growth rate.

7 SUMMARY

We have measured the growth rate of structure from the small-scale clustering of the eBOSS LRG sample corrected by PIP weights and modelled using the AEMULUS cosmological emulator. Using the quasi-linear scales of our measurement range, we obtain a measurement of $f\sigma_8(z = 0.737) = 0.408 \pm 0.038$, which is 1.4σ lower than the value expected from 2018 Planck data for a flat Λ CDM model. Our measurement is a significant improvement over more standard measurements made using only the large-scale modes, achieving a level of precision that is 1.7 times better than the large-scale analysis of the same sample. Using the full separation range of our measurement, we find a 4.5σ tension in the amplitude of the halo velocity field with the expectation for a Λ CDM universe. This tension is driven by the non-linear scales of our analysis and so may not be well modelled by a change in the linear growth rate, but may instead reflect a breakdown in the HOD model used in the emulator.

We perform a robust check of possible sources of systematic error not included in previous analyses. We find that using a model that assumes all central galaxies are observed leads to a systematic bias if the actual occupation of centrals is lower; a fractional occupation of $f_{\max} \leq 0.6$. We also investigate the effect of redshift uncertainty, and find that the presence of a velocity shift from redshift uncertainty in the data that is not included in the model results in a higher measurement of γ_f with an offset of $\sim 0.5\sigma$, where σ is the typical statistical error. This effect is caused by the degeneracy between the increased velocity dispersion due to the redshift uncertainty and the random motions of the haloes in the non-linear regime. Lastly, we investigate the consistency between the non-linear and quasi-linear scales of our analysis. While we find them to be consistent within the statistical error, there is a trend to lower γ_f on non-linear scales, which increases the tension with the expectation from Λ CDM for the fit to all scales. This result highlights the importance of distinguishing between results obtained from the linear scales and thus directly

constraining the linear growth rate $f\sigma_8$, and those that include non-linear scales and may have a non-linear dependence on the linear growth rate together with a dependence on other factors.

While our results are consistent with the expectation from Planck 2018 parameter constraints, we are also consistent with recent weak lensing results giving a low value of S_8 , particularly if our low value of $f\sigma_8$ was driven by an adjustment to σ_8 . In light of these lensing results and the mild disagreement we find with Planck expectations, extending this type of analysis to future surveys including DESI and Euclid will be an important area of future research. With considerably larger samples and probing a different redshift range, the improvement in precision from moving to smaller scales will be key to achieving optimal constraints and identifying or rejecting a tension in the growth rate of cosmic structure.

ACKNOWLEDGEMENTS

ZZ is supported in part by NASA grant 15-WFIRST15-0008, Cosmology with the High Latitude Survey Roman Science Investigation Team (SIT). JLT acknowledges the support of NSF AAG grant 2009291. MJC, FGM, ZZ, and WJP acknowledge financial support from the Canadian Space Agency (CSA) and the Natural Sciences and Engineering Research Council of Canada (NSERC). GR acknowledges support from the National Research Foundation of Korea (NRF) through Grants Nos 2017R1E1A1A01077508 and 2020R1A2C1005655 funded by the Korean Ministry of Education, Science and Technology (MoEST).

Research at Perimeter Institute is supported, in part, by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Colleges and Universities.

This research was enabled, in part, by the support provided by Compute Ontario (www.computeontario.ca) and Compute Canada (www.computecanada.ca).

Funding for the Sloan Digital Sky Survey IV has been provided by the Alfred P. Sloan Foundation, the U.S. Department of Energy Office of Science, and the Participating Institutions. SDSS-IV acknowledges support and resources from the Center for High Performance Computing at the University of Utah. The SDSS website is www.sdss.org. SDSS-IV is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS Collaboration including the Brazilian Participation Group, the Carnegie Institution for Science, Carnegie Mellon University, Center for Astrophysics | Harvard & Smithsonian, the Chilean Participation Group, the French Participation Group, Instituto de Astrofísica de Canarias, The Johns Hopkins University, Kavli Institute for the Physics and Mathematics of the Universe (IPMU) / University of Tokyo, the Korean Participation Group, Lawrence Berkeley National Laboratory, Leibniz Institut für Astrophysik Potsdam (AIP), Max-Planck-Institut für Astronomie (MPIA Heidelberg), Max-Planck-Institut für Astrophysik (MPA Garching), Max-Planck-Institut für Extraterrestrische Physik (MPE), National Astronomical Observatories of China, New Mexico State University, New York University, University of Notre Dame, Observatório Nacional / MCTI, The Ohio State University, Pennsylvania State University, Shanghai Astronomical Observatory, United Kingdom Participation Group, Universidad Nacional Autónoma de México, University of Arizona, University of Colorado Boulder, University of Oxford, University of Portsmouth, University of Utah, University of Virginia, University of Washington, University of Wisconsin, Vanderbilt University, and Yale University.

DATA AVAILABILITY

The eBOSS galaxy and random catalogues are publicly available at: <https://data.sdss.org/sas/dr16/eboss/lss/catalogs/DR16/> with a description here: <https://www.sdss.org/dr16/spectro/lss/> We used the AEMULUS emulator, which is available here: <https://aemulusproject.github.io>, and the COBAYA package, which is available here: <https://github.com/CobayaSampler>.

REFERENCES

Abbott T. M. C. et al., 2018, *Phys. Rev. D*, 98, 043526
 Ahumada R. et al., 2020, *ApJS*, 249, 3
 Alam S. et al., 2017, *MNRAS*, 470, 2617
 Alam S. et al., 2021, *Phys. Rev. D*, 103, 083533
 Albareti F. D. et al., 2017, *ApJS*, 233, 25
 Alcock C., Paczynski B., 1979, *Nature*, 281, 358
 Bautista J. E. et al., 2021, *MNRAS*, 500, 736
 Beutler F. et al., 2012, *MNRAS*, 423, 3430
 Beutler F. et al., 2017, *MNRAS*, 466, 2242
 Bianchi D., Percival W. J., 2017, *MNRAS*, 472, 1106
 Blake C. et al., 2011, *MNRAS*, 415, 2876
 Blanton M. R. et al., 2017, *ApJ*, 154, 28
 Bolton A. S. et al., 2012, *AJ*, 144, 144
 Dawson K. S. et al., 2013, *ApJ*, 145, 10
 Dawson K. S. et al., 2016, *ApJ*, 151, 44
 de Mattia A. et al., 2021, *MNRAS*, 501, 5616
 DeRose J. et al., 2019, *ApJ*, 875, 69
 DESI Collaboration, 2016a, preprint ([arXiv:1611.00036](https://arxiv.org/abs/1611.00036))
 DESI Collaboration, 2016b, preprint ([arXiv:1611.00037](https://arxiv.org/abs/1611.00037))
 du Mas des Bourboux H. et al., 2020, *ApJ*, 901, 153
 Eifler T., Kilbinger M., Schneider P., 2008, *A&A*, 482, 9
 Eisenstein D. J. et al., 2011, *ApJ*, 142, 72
 Feldman H. A., Kaiser N., Peacock J. A., 1994, *ApJ*, 426, 23
 Ferreira P. G., 2019, *ARA&A*, 57, 335
 Gao L., Springel V., White S. D. M., 2005, *MNRAS*, 363, L66
 Gil-Marín H. et al., 2020, *MNRAS*, 498, 2492
 Grieb J. N. et al., 2017, *MNRAS*, 467, 2085
 Gunn J. E. et al., 2006, *AJ*, 131, 2332
 Guzzo L. et al., 2008, *Nature*, 451, 541
 Harker G., Cole S., Helly J., Frenk C., Jenkins A., 2006, *MNRAS*, 367, 1039
 Hartlap J., Simon P., Schneider P., 2007, *A&A*, 464, 399
 Hearin A. P. et al., 2017, *AJ*, 154, 190
 Hinshaw G. et al., 2013, *ApJS*, 208, 19
 Hou J. et al., 2021, *MNRAS*, 500, 1201
 Howlett C., Lewis A., Hall A., Challinor A., 2012, *JCAP*, 1204, 027
 Howlett C., Ross A. J., Samushia L., Percival W. J., Manera M., 2015, *MNRAS*, 449, 848
 Kaiser N., 1987, *MNRAS*, 227, 1
 Krause E., Hirata C. M., Martin C., Neill J. D., Wyder T. K., 2013, *MNRAS*, 428, 2548
 Landy S. D., Szalay A. S., 1993, *ApJ*, 412, 64
 Lang D., Hogg D. W., Schlegel D. J., 2016, *AJ*, 151, 36
 Lange J. U., Hearin A. P., Leauthaud A., van den Bosch F. C., Guo H., DeRose J., 2022, *MNRAS*, 509, 1779
 Laureijs R. et al., 2011, preprint ([arXiv:1110.3193](https://arxiv.org/abs/1110.3193))
 Leauthaud A. et al., 2017, *MNRAS*, 467, 3024
 Lewis A., 2013, *Phys. Rev.*, D87, 103529
 Lewis A., Bridle S., 2002, *Phys. Rev.*, D66, 103511
 Lewis A., Challinor A., Lasenby A., 2000, *ApJ*, 538, 473
 Lyke B. W. et al., 2020, *ApJS*, 250, 8
 Mohammad F. G., de la Torre S., Bianchi D., Guzzo L., Peacock J. A., 2016, *MNRAS*, 458, 1948
 Mohammad F. G. et al., 2020, *MNRAS*, 498, 128
 Navarro J. F., Frenk C. S., White S. D. M., 1996, *ApJ*, 462, 563
 Neal R. M., 2005, preprint ([arXiv:math/0502099](https://arxiv.org/abs/math/0502099))
 Neveux R. et al., 2020, *MNRAS*, 499, 210
 Obuljen A., Dalal N., Percival W. J., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 020
 Okumura T. et al., 2016, *PASJ*, 68, 38
 Percival W. J., Bianchi D., 2017, *MNRAS*, 472, L40
 Pezzotta A. et al., 2017, *A&A*, 604, A33
 Planck Collaboration V, 2020a, *A&A*, 641, A5
 Planck Collaboration VI, 2020b, *A&A*, 641, A6
 Planck Collaboration VIII, 2020c, *A&A*, 641, A8
 Prakash A. et al., 2016, *ApJS*, 224, 34
 Raichoor A. et al., 2021, *MNRAS*, 500, 3254
 Reid B. A., Seo H.-J., Leauthaud A., Tinker J. L., White M., 2014, *MNRAS*, 444, 476
 Reid B. et al., 2016, *MNRAS*, 455, 1553
 Rossi G. et al., 2021, *MNRAS*, 505, 377
 Ross A. J. et al., 2020, *MNRAS*, 498, 2354
 Samushia L., Percival W. J., Raccanelli A., 2012, *MNRAS*, 420, 2102
 Sánchez A. G. et al., 2017, *MNRAS*, 464, 1640
 Satpathy S. et al., 2017, *MNRAS*, 469, 1369
 Sheth R. K., Tormen G., 2004, *MNRAS*, 350, 1385
 Smee S. A. et al., 2013, *AJ*, 146, 32
 Smith A. et al., 2020, *MNRAS*, 499, 269
 Song Y.-S., Percival W. J., 2009, *J. Cosmol. Astropart. Phys.*, 2009, 004
 Tamone A. et al., 2020, *MNRAS*, 499, 5527
 Torrado J., Lewis A., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 057
 Wechsler R. H., Zentner A. R., Bullock J. S., Kravtsov A. V., Allgood B., 2006, *ApJ*, 652, 71
 Weinberg D. H., Mortenson M. J., Eisenstein D. J., Hirata C., Riess A. G., Rozo E., 2013, *Phys. Rep.*, 530, 87
 Zhai Z. et al., 2017, *ApJ*, 848, 76
 Zhai Z. et al., 2019, *ApJ*, 874, 95
 Zhai Z. et al., 2022, preprint ([arXiv:2203.08999](https://arxiv.org/abs/2203.08999))
 Zhao C. et al., 2021, *MNRAS*, 503, 1149