Optimizing Small-scale Redshift Space Distortion Measurements in eBOSS for Cosmological Inference

by

Michael J. Chapman

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Physics

Waterloo, Ontario, Canada, 2023

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner:	Baujui Li Professor, Dept. of Physics, Durham University
Supervisor(s):	Will Percival Professor, Dept. of Physics and Astronomy, University of Waterloo
Internal Member:	Mike Hudson Professor, Dept. of Physics and Astronomy, University of Waterloo
Internal-External Member:	Holger Kleinke Professor, Dept. of Chemistry, University of Waterloo
Other Member(s):	James Taylor Associate Professor, Dept. of Physics and Astronomy, University of Waterloo

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Modern cosmological datasets have grown substantially in size and the precision of their measurements. While the improvement has had a beneficial impact on our understanding of the cosmological model, it requires equal improvements in our analysis methods and the treatment of systematic biases to achieve optimal results. The model that best fits current observations is a spatially flat model with a cold dark matter (CDM) component that dominates the matter density and a cosmological constant (Λ) that dominates the energy density (ΛCDM). The objective of most cosmological datasets is to precisely measure the parameters of the model, discover an extension, or identify a tension with the expectations from another probe, with the eventual goal of discovering new physics. A probe of particular interest for this thesis is measurements of Redshift Space Distortions (RSD), which constrain the growth of structure through the parameter combination $f\sigma_8$, consisting of the logarithmic growth rate of density perturbations, f, and the amplitude of density fluctuations normalized using the standard deviation of fluctuations in a sphere of radius $8\,h^{-1}$Mpc, defined as σ_8. Not only do these measurements constrain core parameters of the ΛCDM model, they are also particularly interesting because they come from the velocity field rather than the density field directly. This makes them complementary to many other large-scale probes, and particularly useful for constraining theories of modified gravity.

One of the new datasets is the extended Baryon Oscillation Spectroscopic Survey (eBOSS), which spectroscopically observed over 1 million galaxies between 2014-2019 as part of the Sloan Digital Sky Survey (SDSS). I present the full eBOSS pipeline, from survey design to the cosmological analysis of the final data release (DR 16), with a focus on my contributions to its development. A key element is the treatment of observational systematics, which must be removed from the data to obtain reliable cosmological results. One of the most significant systematics for small-scale measurements is fibre collisions, where an observational limitation prevents observing close pairs of targets, producing a biased clustering measurement. I present the work of myself and my collaborators within eBOSS to generate Pairwise-Inverse-Probability (PIP) weights and combine them with Angular Upweighting (ANG) to fully remove the effect of fibre collisions, obtaining unbiased clustering measurements on all scales. I also describe my work to correct an observational systematic in the eBOSS Emission Line Galaxy (ELG) sample, caused by inconsistent calibration in the surveys used to identify targets for eBOSS, using a weight-based correction that does not require discarding already observed data. From the final cosmological analyses I present measurements of the Baryon Acoustic Oscillation (BAO) scale and RSD signal from each eBOSS sample. These measurements constrain the expansion history and

growth history over the range $0.6 < z < 2.2$, finding good agreement with the expectation from the 2018 *Planck* Cosmic Microwave Background (CMB) data for a flat ΛCDM model. When combined with other SDSS BAO measurements, as well as CMB and supernovae observations, we obtain precise measurements of the curvature of the Universe and the equation-of-state of the dark matter component, two of the simplest possible extensions to the cosmological model, and find both to be in agreement with flat ΛCDM to high precision.

Constraints on $f\sigma_8$ using small-scale RSD measurements have a significant statistical advantage over those made on large scales. My collaborators and I measure the small-scale clustering of the DR 16 eBOSS Luminous Red Galaxy (LRG) sample, using the PIP+ANG weights to correct for fibre collisions. We fit to the monopole and quadrupole moments of the 3D correlation function and to the projected correlation function over the separation range $7-60\,h^{-1}$Mpc with a model based on the AEMULUS cosmological emulator to measure $f\sigma_8$. 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 *Planck* 2018 measurements 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 modelled by the AEMULUS cosmological emulator, $0.1-60\,h^{-1}$Mpc, 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. 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.

The restriction of constraining $f\sigma_8$ using only the measurement scales $7 - 60\,h^{-1}$Mpc was motivated by the minimum scale at which the velocity scaling parameter used in the emulator to replicate changes in the growth rate still matched the expectation for a change in $f\sigma_8$. This issue highlights an important concern for small-scale RSD measurements: the need to carefully disentangle the linear and non-linear information when interpreting RSD in terms of $f\sigma_8$. It is particularly important to do this given the significant deviation from the expectation based on the *Planck*+ΛCDM model derived using the full range of scales modelled by the emulator in the previous analysis. We construct a new emulator-based model for small-scale galaxy clustering with scaling parameters for both the linear and non-linear velocities of galaxies, allowing us to isolate the linear growth rate. We train the emulator using simulations from the AbacusCosmos suite, estimating the linear velocity of galaxies by evolving the velocities of the simulations' Zel'dovich approximation initial conditions using linear growth. We apply a tophat smoothing kernel of radius $5\,h^{-1}$Mpc to the field to remove the remaining small-scale velocity dispersion, finding good agreement

between the behaviour of our linear velocity scaling parameter and the expectation for a change in $f\sigma_8$ on all scales. We apply the new emulator to the eBOSS LRG sample, obtaining a value of $f\sigma_8(z = 0.737) = 0.368 \pm 0.041$, in 2.3-$\sigma$ tension with the *Planck*+ΛCDM expectation. We also find less dependence on the minimum measurement scale than the previous analysis, validating our improved emulator.

The small- and large-scale eBOSS results provide a precise test of ΛCDM from both the expansion and growth history. While consistent with ΛCDM, these measurements give interesting insight into the current H_0 and S_8 tensions between various cosmological probes, and give some evidence for a third tension between the $f\sigma_8$ measurements of small-scale RSD analyses and the *Planck* 2018+ΛCDM expectation. The observations and analysis of the eBOSS samples, particularly the treatment of observational systematics, pave the way for the next generation of surveys, such as those currently being done by the Dark Energy Spectroscopic Instrument (DESI) and *Euclid* space mission. Applying the small-scale RSD analysis method to these surveys will be critical to achieving optimal constraints, which have the potential to revolutionize the ΛCDM model and our understanding of the Universe.

Acknowledgements

I would like to thank my supervisor, Prof. Will Percival, as well as my co-supervisors, Dr. Faizan Mohammad and Prof. Zhongxu Zhai, for their guidance, support, and patience when the code broke. I have learned more than I can list from each of you, and so much of this achievement would have been impossible without your efforts.

I would also like to thank my wife, Victoria, for sticking by me every step of this long (, long, long, LONG) journey. Your support and motivation got me to the end. I thank my parents, Jan and John, for their unwavering support, and guiding me through all the choices that have led up to this point. I wouldn't be here without you. I would also like to thank my other parents, Emma and Sal, for making sure I didn't starve over the course of all this studying. Your home and your company have always been a refuge for me when times were hard.

I'd like to thank my Advisory Committee, Prof. Mike Hudson, Prof. Niayesh Afshordi, and Prof. James Taylor, for their guidance and direction throughout my degree. I thank Dr. Meg Ward for giving me so many opportunities to teach and develop my skills while also putting food on the table. I thank my friends in Science Computing, especially Mirko Vucicevich, for teaching me that computers are only scary if you let them be. I would like to thank the many collaborators and reviewers that made my published research a reality, especially Prof. Jeremy Tinker, Dr. Dustin Lang, Prof. Etienne Burtin, Prof. Gong-Bo Zhao, Dr. Ashley Ross, Dr. Anand Raichoor, and Dr. Arnaud de Mattia. My thanks go out to all of my instructors, and the faculty and postdoctoral fellows of the Waterloo Centre for Astrophysics, for teaching me so many wonders of astrophysics. It has been the experience of a lifetime.

And last, but never least, I would like to thank my fellow grad students, from the office, class, Cookie Time, PhysGSA, and everywhere else. No one understands the ups and downs of this process like you do, and I appreciate you always being there for advice, a joke, or a game of Hanabi to help me through it.

Dedication

This thesis is dedicated to my grandfather, Dr. John Herbert Chapman, who inspired an early love of outer space, and to Victoria, who gives me a new reason every day to love life on Earth.

Table of Contents

List of Figures

List of Tables

List of Abbreviations

2PCF Two-point Correlation Function 23, 26, 68, 70, 71, 74, 77, 90

6dFGS 6-degree Field Galaxy Survey 17, 34, 35

ACT Atacama Cosmology Telescope 16

ANG Angular Upweighting 17, 45, 48, 65, 66, 68, 70–73, 75–77, 100, 102, 104, 109, 178, 179, 186, 192

AP Alcock-Paczynski 115, 132, 138

BAO Baryon Acoustic Oscillation 15–17, 19, 47–50, 53, 64, 90–101, 104, 115, 117, 169, 179, 180, 183, 185

BBN Big Bang Nucleosynthesis 92, 97, 100, 180, 183

BGS Bright Galaxy Survey 185

BOSS Baryon Oscillation Spectroscopic Survey 17, 33–35, 47–50, 53, 55–58, 65–67, 91, 98, 100, 102–104, 106, 107, 124, 143–146, 177, 181, 185, 193

C-HOD Constrained Halo Occupation Distribution 181

CAMB Code for Anisotropies in the Microwave Background 93, 119

CCD Charge-Coupled Device 59

CCHP Carnegie-Chicago Hubble Program 14

CDM Cold Dark Matter 166, 181

CEM Cosmological Evidence Modelling 181

CFHT LS Canada-France-Hawaii Telescope Legacy Survey 52–54

CLPT Convolutional Lagrangian Perturbation Theory 27, 28, 30, 31, 33, 90, 95

CMB Cosmic Microwave Background 2, 6, 11, 14–16, 19, 21, 22, 47, 92, 97, 99, 100, 119, 143, 169, 179, 180, 183, 194, 197

CP Close Pair 66, 71, 73, 75–77, 138

Dec Declination 82, 109, 190

DECaLS Dark Energy Camera Legacy Survey 51, 54, 60, 75, 78–88, 186

DES Dark Energy Survey 19, 51, 75, 78–88, 146, 147

DESI Dark Energy Spectroscopic Instrument 17, 50, 51, 146, 185, 186

DR Data Release 46, 48, 50, 51, 53, 54, 58, 63, 64, 73, 75–83, 85–87, 95–97, 105, 138, 139

eBOSS Extended Baryon Oscillation Spectroscopic Survey 12, 17, 21, 24, 27, 28, 31–35, 45–59, 62, 63, 65, 66, 68, 70, 71, 73, 75–81, 89, 91, 95–98, 100–102, 104–107, 110, 111, 118–128, 130–132, 135, 138–141, 144–146, 148–151, 159–161, 163–168, 171, 175, 177–181, 183–186, 189–193, 197, 198

EFT Effective Field Theory 27, 33

ELG Emission Line Galaxy 45–51, 53–57, 60, 62, 64, 71, 73, 75, 78, 80–82, 85–91, 93–96, 98, 100, 101, 144, 177, 179, 184–186, 191

EZmock Effective Zel'dovich Approximation mock 63, 71–73, 76, 77, 96, 111, 113, 114, 137, 179, 191

FoF Friends-of-Friends 38

FoG Finger-of-God 28, 29, 31, 32, 93

GAMA Galaxy and Mass Assembly 17, 34, 35

GLAM GaLAxy Mocks 64, 78, 79

GR General Relativity 3, 11, 108, 128, 182, 183

GS Gaussian Streaming 27, 30, 31, 90, 95

HACC Hardware/Hybrid/Accelerated Cosmology Code 64

HEALPix Hierarchical Equal Area isoLatitude Pixelation 59, 81, 82, 84–86, 90

HOD Halo Occupation Distribution 39, 40, 42, 64, 105–107, 111, 117–120, 122–126, 128, 131, 138, 143, 145, 146, 148, 152, 154, 155, 159–161, 163, 164, 166, 181, 182, 184, 186, 193, 197

HSC Hyper Suprime-Cam 19

IIP Individual-Inverse-Probability 69, 70

KiDS Kilo-Degree Survey 19

LRG Luminous Red Galaxy 32, 40, 44, 45, 47–53, 55–62, 64, 71–73, 76, 77, 91, 93–98, 100, 102–107, 110, 111, 118, 120, 122, 124, 126–128, 130, 135, 138, 140, 144, 146, 148–151, 159, 160, 163, 165–168, 171, 175, 177, 179–181, 184–186, 193, 198

MCMC Markov chain Monte Carlo 117, 118, 128, 161, 164, 168, 198

MG Modified Gravity 182, 183, 185

MGS Main Galaxy Sample 17, 34, 35, 98, 144, 177

NFW Navarro-Frenk-White 38, 42, 107

NGC North Galactic Cap 51, 53–55, 60, 61, 72, 76, 77, 80, 110

NN Nearest-Neighbour 66

NOAO National Optical Astronomy Observatory 81

NOIRLab National Optical-Infrared Astronomy Research Laboratory 81

PCA Principal Component Analysis 57

PIP Pairwise-Inverse-Probability 17, 24, 45, 46, 48, 65, 66, 68–73, 75–77, 100, 102, 104, 109, 138, 148, 178, 179, 186, 189, 191, 192

PTF Palomar Transient Factory 53

QPM Quick Particle-Mesh 64, 78, 79

QSO Quasi-Stellar Object 47, 49, 64, 71, 73

RA Right Ascension 82, 109, 190

RSD Redshift Space Distortion 1–3, 6, 10, 11, 17, 21, 25, 27–29, 31–35, 44, 45, 47, 48, 50, 63, 64, 90, 91, 95–98, 101–105, 108, 119, 137, 140, 149–151, 166, 176–186, 193

SDSS Sloan Digital Sky Survey 17, 33–35, 47–53, 55, 58–60, 64, 73, 91, 96–100, 104, 144, 176, 177, 179, 180

SGC South Galactic Cap 51, 53–55, 60, 61, 72, 76, 77, 80, 110

SHAM Subhalo Abundance Matching 39, 102, 126, 127, 163, 198, 199

SPIDERS SPectroscopic IDentification of ERosita Sources 56

SPT South Pole Telescope 16

TDSS Time-Domain Spectroscopic Survey 56

TNS Taruya-Nishimichi-Saito 27, 31, 33, 95

TRGB Tip of the Red Giant Branch 13, 14

VIPERS VIMOS Public Extragalactic Redshift Survey 17, 34, 35, 54, 68

WISE Wide-field Infrared Survey Explorer 50, 52

Chapter 1

Introduction

This chapter provides an introduction to the key concepts used in the remainder of this thesis. The majority of this chapter was written by me, with certain sections modified from the papers presented later in the thesis. Sec. 1.3.2 was adapted from a subsection originally written by Faizan Mohammad as part of [57]; a paper for which I lead the analysis of the data and the preparation of the paper (see Ch. 2). Sec. 1.4.2 and Sec. 1.4.3 were adapted from [23] (presented in Sec. 2.7), an analysis for which I was an author based on my contributions to the data catalogues and that I participated in proofreading. In both cases the content has been reformatted and rewritten by me for this thesis. All non-original plots contain a citation at the end of their captions detailing the source.

The modern field of cosmology is often described as the "era of precision cosmology". While past "eras" were focused on discriminating between very different models of the Universe, in the modern era we believe we have a firm understanding of much of the theory underlying the cosmological model of our Universe. The focus is instead on refining that model and identifying possible extensions through increasingly precise observations. Improved precision necessitates improved analysis tools and lower tolerances for systematic uncertainty, which in the context of Redshift Space Distortion (RSD) measurements form the basis of my thesis work.

In this chapter I will introduce the theoretical concepts applied and investigated in future chapters, including the currently leading cosmological model and the measurements that have defined it. I will also present the key analysis tools that are fundamental to later chapters.

I begin by providing an overview of our current standard model of cosmology in Sec. 1.1. In particular I introduce the key parameters that will be measured in later chapters, and

provide a detailed description of the growth of structure in the Universe, which is the aspect of the standard model most directly tested throughout my thesis. I also briefly introduce extensions to the standard model, with particular attention to those tested by RSD measurements.

Moving from theory to observations, in Sec. 1.2 I describe the key observables that have contributed to our cosmological knowledge. I discuss the methods and constraints from the local distance ladder, Cosmic Microwave Background (CMB), and weak lensing observations. I also introduce galaxy spectroscopic surveys and the ways in which they can be used to constrain the cosmological model, which will be central to later work. I end this section with a discussion of the current tensions between these measurements, which provide promising avenues for expanding our understanding of the cosmological model.

In Sec. 1.3 I expand on the subject of galaxy spectroscopic surveys by describing the measurement statistics that are used in their analysis. In particular I discuss two-point functions, which are the type of statistic measured and analyzed throughout this thesis, and give an overview of how they are measured in both configuration- and Fourier-space.

In Sec. 1.4 I focus on Redshift Space Distortion measurements, which are the key source of cosmological information from galaxy spectroscopic surveys used in much of my research. I discuss the connection between the growth of structure and RSD, and the theory behind RSD measurements. I then present the models commonly used for modern RSD analyses, and the current state of RSD measurements.

I expand on the particular type of RSD model used in my most significant analyses, emulator-based models, in Sec. 1.5. I first describe the N-body simulations that form the basis of this model. I then introduce the halo model, which is a framework for describing how the continuous density field of the simulation can be analyzed in terms of discrete tracers, and the methods of connecting the galaxies we observe to those halos. Lastly, I give an overview of cosmological emulators, which combine these tools to efficiently predict the clustering of galaxies. These emulators are one of the core pillars of the small-scale RSD analyses that make up the later parts of this thesis.

After introducing these key concepts, in Sec. 1.6 I provide an overview of the remaining chapters. I introduce the data and analyses that have resulted in this completed thesis, and the implications for our understanding of the cosmological model.

1.1 Cosmological theory

The standard model of cosmology has found great success in matching a variety of measurements using a small set of well understood parameters. In this section I will give an overview of this cosmological model, with a focus on the aspects most closely related to RSD. In Sec. 1.1.1 I introduce the model, the key parameters, and the current constraints on their values. In Sec. 1.1.2 I go into more detail about the growth of structure in this model, which is the progenitor of the RSD effect. Then in Sec. 1.1.3 I describe a few of the possible extensions to the standard model, with particular attention to those that are relevant for RSD analyses.

1.1.1 Standard model of cosmology

The standard model of cosmology is rooted in Albert Einstein's theory of General Relativity (GR), which provides a framework for the evolution of the Universe. The key element from GR for cosmology is Einstein's field equation:

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R - g_{\mu\nu}\Lambda = \frac{8\pi G}{c^4}T_{\mu\nu}\,, \tag{1.1}$$

where $R_{\mu\nu}$ is the Ricci tensor that describes the local curvature of space-time, R is the scalar curvature, $g_{\mu\nu}$ is the metric, Λ is the cosmological constant, and $T^{\mu\nu}$ is the energy-momentum tensor of the contents of the Universe. In essence, this equation relates the curvature of space-time to the energy content of the Universe.

Einstein's field equation can be solved to give a description of how the size of the Universe changes with time, one of the core fields of study for cosmology. Solutions are generally calculated under the assumptions that the Universe is homogeneous and isotropic on sufficiently large scales ($\gtrsim 100$ Mpc [113]). Together these assumptions are referred to as the cosmological principle, and have so far been supported by extensive observations of the Universe (see Sec. 1.2).

Typically the size of the Universe is characterized by a scale factor, $a(t) = l(t)/l(t_0)$, that measures how the distance, l, between two points fixed with respect to the cosmological background changes in time. A related quantity is the redshift, z, of light emitted at time t that we observe at time t_0. The expansion also affects the wavelength of light, with the change from the wavelength emitted at t, $\lambda_{\rm emit}$, to the wavelength we observe at t_0, $\lambda_{\rm obs}$, given by $1 + z = \lambda_{\rm obs}/\lambda_{\rm emit}$. The redshift can then be related to the scale factor using:

$$\frac{a}{a_0} = \frac{1}{1+z}\,, \tag{1.2}$$

where a_0 is the scale factor at the present day, commonly normalized to 1. Redshift is a common method of measuring the distance to objects in the Universe, and so will be used throughout this thesis as a distance scale and to define the epoch of various cosmological events.

Alexander Friedmann was the first to publish the solution to Einstein's field equation for an expanding Universe, now called the Friedmann equation:

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{Kc^2}{a^2} + \frac{\Lambda c^2}{3}\,, \tag{1.3}$$

where $H = \dot{a}/a$ measures the expansion rate of the universe at scale factor a, called the Hubble parameter, and ρc^2 is the combined energy density of the components of the Universe.

ρ can be separated into the contribution from each component, which also allows us to examine how the density of each component evolves with the scale factor. The matter component will dilute as the inverse change in volume of the universe, so $\rho_M \propto a^{-3}$. The number density of photons will also scale as a^{-3}, while the energy of each photon will scale as a^{-1}, meaning the overall scaling of the radiation component is $\rho_R \propto a^{-4}$. These scalings are often given in terms of an equation-of-state parameter, w, that relates the density and pressure of a component as $P = w\rho c^2$. The scale factor dependence of a component goes as $\rho \propto a^{-3(1+w)}$, meaning that the equation-of-state parameters of the various components are $w = 0$ for matter, consistent with a pressureless fluid; $w = 1/3$ for radiation, consistent with an ultra-relativistic fluid; and Λ behaves like a component with $w = -1$, consistent with a negative pressure fluid and such that the energy density does not change with scale factor.

In Eq. 1.3 K is the curvature signature of the Universe that has values:

$$K = \begin{cases} +1 & \text{Closed} \\ 0 & \text{Flat} \\ -1 & \text{Open} \end{cases} \tag{1.4}$$

In the absence of Λ, K determines the dynamics of the universe at very late times because the K term scales as a^{-2}, while the matter and radiation density terms scale as a^{-3} and a^{-4} respectively. If $K = +1$ then the right hand side of Eq. 1.3 will become negative, reversing the expansion of the universe and leading to collapse. We refer to this as a closed geometry of the universe, and occurs when the energy densities are sufficiently large enough that gravity counteracts the expansion of the universe, slowing and ultimately reversing it.

Alternatively, if $K = -1$ then the right side of the equation will remain positive, meaning the expansion will continue infinitely and the universe possesses an open geometry. This can be understood as the energy densities being too small to effectively counteract the expansion through gravitation, because as the scale factor increases the gravitational force between distant parts of the universe will diminish, resulting in a universe that expands without end. The final case is when $K = 0$, which is the critical boundary between the two cases and describes a universe with a flat geometry. In this case the Universe will also continue to expand forever as the right side of Eq. 1.3 approaches zero.

In order to determine the geometry of the universe it is useful to define the critical density,

$$\rho_{\text{crit}} \equiv \frac{3H^2}{8\pi G}\,, \tag{1.5}$$

and describe the densities of the various components in units of the critical density:

$$\Omega = \frac{\rho}{\rho_{\text{crit}}}\,, \tag{1.6}$$

Using these definitions, Eq. 1.3 can be re-arranged to give:

$$K = \left(\frac{Ha}{c}\right)^2 (\Omega - 1)\,, \tag{1.7}$$

where Λ has been included in Ω as a component with density $\rho_\Lambda = c^2\Lambda/(8\pi G)$. From Eq. 1.7 it is clear the comparison of ρ to ρ_{crit} determines the sign of K, and thus the geometry of the Universe. Current cosmological measurements (see Sec. 1.2) favour $\rho = \rho_{\text{crit}}$ to high precision, implying the geometry of the Universe is flat [213, 11].

The densities of each of the major components of the Universe have been measured to high precision. The component most familiar to humans is baryonic matter (Ω_b; protons, electrons, neutrons, etc.), and has been determined to be only a small fraction of the energy density of the Universe. Making up a considerably larger portion is a mysterious dark matter (Ω_c for dark matter only, Ω_m for total matter density when combined with baryonic matter). Dark matter is known to interact gravitationally [284, 198, 89, 236], but only weakly, if at all, using other mechanisms, making it currently impossible to observe through radiation. Observations point to the majority of this dark matter having non-relativistic velocities, so it is referred to as "cold" dark matter (CDM) [71]. The dominant component of the Universe at the current time is the even more mysterious dark energy (Ω_Λ) that accelerates the expansion of the Universe [228, 205]. Current observations find it to be consistent with a cosmological constant, Λ [213]. Together these two dark components

give the abbreviation for the standard model, ΛCDM. The final component of the standard model is radiation (Ω_r), which was dominant in the early Universe and is still observable today, but is now small compared to the other components. The measurements of these densities at the present time, denoted by a subscript "0", from the Cosmic Microwave Background [213] for a flat-ΛCDM universe are:

$$\Omega_{m,0} = 0.311 \pm 0.0056\,, \quad \Omega_{b,0}h^2 = 0.02242 \pm 0.00014\,, \quad \Omega_{\Lambda,0} = 0.6889 \pm 0.0056. \quad (1.8)$$

The density of the baryon component is reported as $\Omega_{b,0}h^2$, where $h \equiv H_0/(100\,\mathrm{km\,s^{-1}\,Mpc^{-1}})$ is a variable defined for convenience (see Sec. 1.2.1), because this is the physical density most directly measured by CMB observations (see Sec. 1.2.2). The radiation density can be calculated from the blackbody temperature of the CMB, $T_{\rm CMB} = 2.72548 \pm 0.00057$ [100], and the effective number of neutrino species, $N_{\rm eff} = 2.99 \pm 0.17$ [213], which are the largest radiation-like components of the Universe. They give a value of $\Omega_{r,0} = 9.2 \times 10^{-5}$ [187]. Throughout this thesis measurements of the density of each component, as well as other cosmological parameters, will typically be made using their values at the present time, so in most cases I will drop the subscript "0", and specify the redshift of the measurement only if it occurs at an earlier time (e.g. $\Omega_m(z = 0.7)$ or $\Omega_m(0.7)$). The only exceptions are measurements of $f\sigma_8$ from Redshift Space Distortion (RSD) (see Sec. 1.4), which are always presented at the effective redshift of the sample they are measured from.

In addition to these parameters, in order to specify the expansion history we need to know the rate of expansion at the present day, given by H_0. There is some tension in current measurements of H_0 (see Sec. 1.2.5), with the best CMB measurement giving $H_0 = 67.66 \pm 0.42\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ [213] and the best local measurements giving $H_0 = 73.04 \pm 1.04\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ [229].

In addition to measuring the expansion of the Universe, the field of cosmology is also occupied with investigating the distribution of matter and structure within the Universe. The origin of this structure is believed to be quantum fluctuations in the very early Universe. These fluctuations introduce perturbations to the density field that closely follow a Gaussian random field. The initial conditions of this field can be predicted by theory, and the resulting power spectrum is well matched by a power law, parameterized by two values. The first is the exponent of the initial power law, n_s, and the second is the amplitude, A_s. It is worth noting that for measurements of the late-Universe it is common to use the standard deviation of density fluctuations in a sphere of $8\,h^{-1}$Mpc, σ_8, to normalize the power spectrum of density perturbations instead of A_s. The best CMB estimates for these parameters, again for a flat-ΛCDM universe, are [213]:

$$n_s = 0.9665 \pm 0.0038\,, \quad A_s = 2.105 \times 10^9 \pm 0.030\,. \quad (1.9)$$

Together these parameters for the densities of the components of the Universe, the expansion rate, and the power law of density perturbations define the key characteristics of the standard cosmological model, and can be used to make accurate predictions for the expansion of the universe and the distribution of matter within it. Within the last several decades cosmological observations have improved to the point where all basic components of the model are well understood and the parameters precisely constrained, defining the era of "precision cosmology". The focus of new cosmological measurements, such as the ones presented in Ch. 2, 3, and 4, is to investigate deviations from and extensions to the standard model through increasingly precise measurements of the key parameters.

1.1.2 Growth of structure

While fluctuations in the initial density field are small due to their quantum origins, they do not remain that way due to the influence of gravity. Overdense areas have stronger gravitational attraction than the surrounding area, and so will become even denser, while the opposite is true for underdense regions. It is this growth that leads to the structure of the Universe we see today, including clusters, groups, and individual galaxies. While this growth would be of interest purely for the structures it produces, it can also be used as a method of measuring the parameters of the cosmological model and testing the ability of that model to match observations. This aspect of the growth of structure is a major focus of this thesis, and is the fundamental aspect of the cosmological model investigated in Ch. 3 and 4.

When working with the density field it is useful to define the density contrast, $\delta = (\rho - \bar{\rho})/\bar{\rho}$, where $\bar{\rho}$ is the mean density of the universe, so that $\delta > 0$ corresponds to an overdensity and $\delta < 0$ corresponds to an underdensity. Rather than working directly with δ it is convenient to define the Fourier expansion of the density contrast at wavenumber $\boldsymbol{k}$:

$$\delta_{\boldsymbol{k}}(t) = \frac{1}{\sqrt{V}} \int \delta(\boldsymbol{r}, t) e^{-i\boldsymbol{k}\cdot\boldsymbol{r}} d^3\boldsymbol{r}, \tag{1.10}$$

where V is the volume being integrated. Working in terms of $\delta_{\boldsymbol{k}}$, the density contrast of the mode with Fourier wavelength $\boldsymbol{k}$, leads to several simplifications. One of the most useful is that perturbations at different $\boldsymbol{k}$ are independent in a Gaussian random field.

We can derive a differential equation for how these overdensities change in time by combining the continuity, Euler, and Poisson equations, under the assumption of small fluctuations, $|\delta| \ll 1$, yielding:

$$\frac{\partial^2 \delta_{\boldsymbol{k}}}{\partial t^2} + 2\frac{\dot{a}}{a}\frac{\partial \delta_{\boldsymbol{k}}}{\partial t} = \left(4\pi G \rho_M - \frac{k^2 c_s^2}{a^2}\right) \delta_{\boldsymbol{k}} - \frac{2}{3}\frac{T}{a^2} k^2 S_{\boldsymbol{k}}, \tag{1.11}$$

where ρ_M is the mean matter density, c_s is the speed of sound, T is temperature, and all derivatives are with respect to cosmic time t. $S_{\boldsymbol{k}}$ is the Fourier expansion of entropy perturbations at wavenumber $\boldsymbol{k}$, equivalent to the Fourier expansion of perturbations in the density field defined in Eq. 1.10.

Eq. 1.11 can be simplified under a series of conditions. One is in the form of the initial perturbations, of which there are two usual types. The first is isentropic initial conditions, where the fluctuations of all components of the density field (matter, radiation, etc.) are proportional to each other, meaning the entropy is constant across the field ($S_{\boldsymbol{k}} \propto \nabla S = 0$) and the curvature varies. The second is isocurvature initial conditions, where the curvature is constant across the field, meaning the fluctuations in the various components are not proportional and the entropy varies. Isentropic initial conditions are predicted by currently favoured models and a better match to data [215], so we will assume that they are the correct form of initial conditions for the Universe, which eliminates the last term of Eq. 1.11.

The other key condition is the scales being investigated. It is helpful to refer to the Jeans length, $k_{\rm Jeans} = \sqrt{4\pi G \rho_M}(a/c_s)$, above which (smaller wavenumber in Fourier space) gravitational instability overcomes pressure and overdensities can grow through collapse. Defining a sphere with a diameter equal to the Jeans length gives a characteristic mass for which collapse is possible. Prior to recombination the Jeans mass is approximately $\sim 10^{19}\, M_\odot$, meaning perturbations on the scales of galaxies or clusters are unable to grow. However, shortly after recombination the Jeans mass drops significantly to $\sim 10^6\, M_\odot$, well below galactic scales [187]. Therefore, the second condition is to restrict the analysis to scales larger (smaller Fourier wavenumber) than the Jeans scale, effectively eliminating the term proportional to k^2 in Eq. 1.11. Using these two conditions, the differential equation for the growth of perturbations becomes:

$$\ddot{\delta} + 2H\dot{\delta} - 4\pi G \rho_M \delta = 0. \tag{1.12}$$

While Eq. 1.12 can be solved numerically for an arbitrary cosmological model, it is informative to examine the solutions for a flat universe dominated by a single component, matching the different epochs in the cosmic history of our Universe. For a radiation dominated universe $a \propto t^{1/2} \rightarrow H \propto 1/(2t)$ and $4\pi G\rho_M \ll H^2$, so that the solution to Eq. 1.12 is:

$$\delta(t) = A_1 + A_2 \ln(t)\,, \tag{1.13}$$

where A_1 and A_2 are constants. So in the radiation dominated era perturbations grew logarithmically in time, i.e. very slowly.

For a matter dominated universe $a \propto t^{2/3} \rightarrow H \propto 2/(3t)$ and $4\pi G\rho_M = (3/2)H^2$. The

solution is therefore:

$$\delta(t) = B_1 t^{2/3} + B_2 t^{-1}\,, \tag{1.14}$$

where B_1 and B_2 are again constants. These two terms are referred to as the growing and decaying modes respectively. The growing mode increases proportional to $t^{2/3} \propto a$, meaning perturbations grow significantly during this epoch. Perturbations in the decaying mode will rapidly decrease as t^{-1}.

The final epoch is that of a dark energy dominated universe, which is expected to occur in the future of our Universe. In that case $a \propto e^{Ht} \rightarrow H = \text{const}$ and $4\pi G\rho_M \ll H^2$, so that the solution to Eq. 1.12 is:

$$\delta(t) = C_1 + C_2 e^{-2Ht} \simeq \text{const}\,, \tag{1.15}$$

where once again C_1 and C_2 are constants. So once dark energy becomes dominant perturbations will cease to grow at all.

When investigating the growth of structure it is useful to define a couple of quantities. The first is the growth factor,

$$D(a) = \frac{\delta(a)}{\delta(1)}\,, \tag{1.16}$$

which characterizes the factor by which perturbations have grown between scale factor a and the present time. Within the matter dominated era D can be calculated as a function of redshift using the approximation $D(z) \propto g(z)/(1+z)$, where [55]

$$g(z) \approx \frac{5}{2}\Omega_m(z)\left\{\Omega_m^{4/7}(z) - \Omega_\Lambda(z) + \left[1 + \frac{\Omega_m(z)}{2}\right]\left[1 + \frac{\Omega_\Lambda(z)}{70}\right]\right\}^{-1}\,. \tag{1.17}$$

The second quantity is the growth rate,

$$f(a) = \frac{d\ln D}{d\ln a}\,, \tag{1.18}$$

which is the rate of change of the growth factor, and is particularly useful for analyses of the velocity field, as discussed in Sec. 1.4. Like D, f is not a fundamental parameter in the ΛCDM cosmological model, and can instead be calculated from the values of the other parameters, most importantly Ω_m. In our Universe a good approximation is $f(\Omega_m) \approx \Omega_m^{0.55}$ [201].

Eq. 1.11 and subsequent derivations were constructed under the assumption of small fluctuations, so that calculations could be limited to linear order. However, as perturbations grow to $\delta \sim 1$ these approximations are no longer valid. Instead, we need a method

for calculating the non-linear evolution of a density perturbation. One of the simplest models is to assume spherical symmetry for the overdensity, and that that symmetry is maintained throughout the collapse. This model, referred to as spherical collapse [116], relies on modelling the overdensity as a series of spherical shells with no mass mixing, i.e. the order of the shells is preserved. The result is that the collapse of the spherical overdensity can be calculated down to a radius of zero, when the collapse is determined to be complete. At this point the assumption that there is no mass mixing has been violated, called "shell-crossing", and rather than collapsing to a point the mass shells will form an extended and quasi-static virialized structure that we refer to as a halo (see Sec. 1.5.2).

A key consequence of this model is the critical density for collapse, $\delta_c = 1.69$, which is the overdensity a perturbation that finishes collapsing today would have at the present day if it had continued to grow linearly. It can be used to identify initial perturbations that will have finished collapse today if they exceed that threshold. This is in turn useful for determining the transition between the linear and non-linear regimes. In terms of separations of pairs of galaxies, s, which is the key independent variable of the clustering measurements described in Sec. 1.3, $s \gtrsim 40h^{-1}$Mpc corresponds to the linear regime, $20 \lesssim s \lesssim 40h^{-1}$Mpc corresponds to the quasi-linear regime where linear theory with corrections can still be successfully applied, $5 \lesssim s \lesssim 20h^{-1}$Mpc corresponds to the mildly non-linear regime, and $s \lesssim 5h^{-1}$Mpc corresponds to the strongly non-linear regime where structures begin to virialize and reach an equilibrium. These different regimes can be identified using the pairwise velocity between objects, as shown in Fig. C.1 of Appendix C.

The growth of structure provides several important methods for measuring and testing the cosmological model. Both the amplitude of perturbations and the rate at which they grow can be measured in observations (see Sec. 1.2). These measurements can be used either to constrain fundamental cosmological parameters, such as A_s and Ω_m, or to test for extensions to the standard model. Part of the strength of these constraints is that they are complementary to direct probes of the cosmological parameters, so can be used to validate other observations or falsify extensions that have been tuned to match a different set of observations. It is this aspect that makes RSD measurements (see Sec. 1.4) unique, and motivates their use in Ch. 3 and 4.

1.1.3 Extensions to the standard model

While the ΛCDM model has had great success predicting the growth of structure and matching a variety of observations (Sec. 1.2), there remain some tensions (see Sec. 1.2.5) and unexplained phenomena that motivate extensions to the model. One of the most

impactful for cosmology is the nature of dark energy. Observations of the expansion history of the universe are largely consistent with dark energy as a cosmological constant [213, 11], however there is not a compelling theoretical basis for the origin of a cosmological constant of this magnitude [270]. There exist many alternative models [14], such as a scalar field that behaves similarly to a cosmological constant, but has some deviation from an equation-of-state of exactly $w = -1$. Such models are generally parameterized as having an equation-of-state that has a constant term that differs from a cosmological constant, $w_0 \neq -1$, that varies with scale factor, $w_a(a)$, or both. More complicated models feature dark energy with multiple sub-components arising from different scalar fields, or interactions within the dark sector between dark energy and dark matter [82].

An alternative to the set of dark energy models is to construct a modified gravity model that differs from GR, and is able to produce accelerated expansion without invoking dark energy (see reviews in [252, 98]). These models are typically constructed to match the expansion history while leaving galaxy dynamics unaffected through a screening mechanism in order to be consistent with observations at both these scales. These models are therefore inherently difficult to test through observation.

This difficulty is a great motivator for measurements of the growth of structure, because that growth is very sensitive to the model of gravity. These measurements can also be used to constrain dark energy models through their dependence on the expansion rate, which opposes the formation of structure. RSD measurements in particular are very valuable because they probe the velocity field, making them complementary to other probes of cosmology.

1.2 Cosmological observations

The greatest success of the ΛCDM cosmological model is the ability to match a wide variety of observations with a small set of parameters. In this section I will introduce some of the key observations that have led to the acceptance of the ΛCDM model and provided our most important constraints on its parameters.

In Sec. 1.2.1 I discuss one of the earliest sources of cosmological information, and still one of the most impactful today: using a chain of calibrated distance measurements, called the distance ladder, to measure the local expansion rate. Then in Sec. 1.2.2 I discuss measurements of the Cosmic Microwave Background (CMB), which provide some of our most complete and precise measurements of the ΛCDM cosmological parameters. In Sec. 1.2.3 I give an overview of spectroscopic galaxy surveys, such as the Extended Baryon

Oscillation Spectroscopic Survey (eBOSS) presented in Ch. 2, and the unique constraints they provide on both the expansion and growth history of the Universe. In Sec. 1.2.4 I discuss weak lensing surveys, and the complementary information they provide on structure growth and the matter content of the Universe. Finally, in Sec. 1.2.5 I discuss tensions between the results of these different probes, and the opportunities they provide to expand our knowledge of the cosmological model.

1.2.1 Local distance ladder

One could argue that Edwin Hubble's measurements of the recession velocities of nearby galaxies through the local distance ladder [146] were the first cosmological measurements. The principle of the measurement is simple. In an expanding universe distant galaxies will appear to be receding from us, with more distant galaxies receding more quickly. Locally, the relationship between recession velocity and distance is linear, and goes as:

$$v_r = H_0 r + v_{\mathrm{pec},r} \tag{1.19}$$

where v_r is the radial velocity of the galaxy (that is, along the line-of-sight), H_0 is the Hubble constant at present time, and $v_{\mathrm{pec},r}$ is the radial velocity of the galaxy caused by the growth of structure that is in addition to the recession velocity, termed the peculiar velocity. This relationship can be used to measure H_0, one of the fundamental parameters of the ΛCDM model (see Sec. 1.1.1). However, in order to obtain an accurate and precise measurement we require $v_{\mathrm{pec},r}$ to be small compared to the recession velocity. Typical peculiar velocities are on the order of hundreds of km/s while H_0 is now known to have a value of $\sim 70\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ [213, 229], so objects with distances on the order of tens of Mpc or greater are required.

Determining accurate and precise measurements to such distant objects is non-trivial because of the difficulty in determining the intrinsic size or brightness of astrophysical objects. The basis of the distance ladder technique is to begin with a distance measurement that is known absolutely, such as through geometric means, and use that to calibrate the distance to nearby objects that are known to have a standard observational quantity, such as the luminosity. Those objects can then be used to calibrate the distance to more distant objects, and so on, until accurate distances are obtained for objects sufficiently far away to measure H_0.

Modern distance ladder measurements use parallax measurements of stars in the Milky Way and the Magellanic clouds, which are purely geometric measurements and do not require knowledge of the intrinsic qualities of the objects. Measurements from the *GAIA*

space telescope [107] provide the most accurate parallax measurements to date [106]. There are multiple options for further rungs of the ladder, the most well-known being Cepheid variable stars, which are what Hubble originally used to measure the distance to nearby galaxies. Cepheids vary in brightness cyclically with a fixed period [172], and the period of that cycle is highly correlated with the total luminosity of the star [43]. This allows them to be used as a standard candle: an object whose brightness is intrinsically known, and as such the distance to the object can be determined from the flux received from the object. Cepheids are bright enough to observe at a distance on the order of tens of Mpc, so that they can be used for a basic distance ladder measurement and to calibrate more distant standard objects.

The final rung of the distance ladder is type-Ia supernovae, which are supernovae from white dwarf stars that through accretion have exceeded the maximum mass for which electron degeneracy pressure can prevent collapse, $1.44\,M_\odot$ [144, 16, 63]. Because this event is related to a fixed mass the peak luminosities of type-Ia supernovae have little dispersion, and that dispersion is correlated with the rate at which the luminosity decays, allowing them to be calibrated as an accurate standard candle [127]. Type-Ia supernovae are also extremely bright, with peak luminosities on the order of $\sim 10^{10}\,L_\odot$ [62]. This allows distance ladder measurements to be extended to high redshift where the matter dominated epoch transitioned to the current dark energy dominated epoch, making them sensitive to the densities of those components, Ω_m and Ω_Λ. Type-Ia supernovae observations provided the first evidence for the accelerated expansion of the universe [228, 205], unveiling the role of dark energy in the evolution of our Universe. Currently, type-Ia supernovae observations calibrated using Cepheid variables from the SH0ES collaboration provide the tightest constraints on the expansion rate from the distance ladder, with a most recent result of $H_0 = 73.04 \pm 1.04\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ [229].

An alternative for Cepheids as the intermediate rung of the ladder is the Tip of the Red Giant Branch (TRGB) [105]. As helium builds up in the core of main sequence stars hydrogen fusion will continue in a shell around the core, causing the star's internal temperature to increase and the star to expand, transitioning to a red giant. This process will continue until the core temperature reaches the threshold required to start helium fusion, at which point the star will undergo a rapid increase in temperature and decrease in luminosity, leaving a sharp cutoff in colour-space. Like the trigger of type-Ia supernovae, this phenomenon is linked to a universal mechanism, the temperature required to fuse helium, so will occur in the same way for all stars and in all galaxies, with only weak dependencies on mass and metallicity. The observations of the TRGB in the halos, where observational systematics are easier to correct for, of supernovae type-Ia host galaxies have provided a strong intermediate rung on the distance ladder. A recent result from the

Carnegie-Chicago Hubble Program (CCHP) measured $H_0 = 69.8 \pm 2.4\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ using TRGB calibrations for type-Ia supernovae distances [105].

The value of H_0 has now been determined to high precision using a variety of probes, although there remains some tension over the true value (see Sec. 1.2.5). In the second half of the 20th century this tension was even greater and generated significant debate amongst astronomers, with estimates ranging from $50 - 100\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ [240]. This was particularly concerning because it is standard in cosmology to use comoving distances, which are distances corrected for the expansion of the Universe such that two points that are stationary with respect to the background expansion will maintain the same comoving separation throughout cosmic time. Unless otherwise specified, all distances described in this thesis are in comoving units. The issue in using comoving units to describe the distance to objects is that one must integrate over the expansion history up to the redshift of the object, which introduces a factor of $1/H_0$. Changing the value of H_0 would therefore require updating all comoving distance measurements. To avoid this problem it has become standard to calculate distances assuming a value of $H_0 = 100\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ and introduce an additional parameter, $h = H_0^T/(100\,\mathrm{km\,s^{-1}\,Mpc^{-1}})$, where H_0^T is the true value for the Universe. Comoving distance measurements can then be described in units of Mpc/h or $h^{-1}\,\mathrm{Mpc}$. In this way comoving distances can be calculated consistently regardless of values of H_0, and comoving distances for a particular cosmology can be easily calculated by dividing by the value of h for a given value of H_0^T. Throughout this thesis many distances will be described in units of $h^{-1}\,\mathrm{Mpc}$.

1.2.2 Cosmic Microwave Background

While the distance ladder provided some of the first cosmological measurements, many of the most precise and comprehensive tests of the ΛCDM model have come from observations of the Cosmic Microwave Background (CMB). By tracing the expansion of the Universe backwards we arrive at a very early Universe that was much denser and hotter than today, such that all baryonic matter was in the form of an ionized plasma that was coupled to the photons making up the radiation component. As expansion continued and temperatures decreased the ions combined to form neutral elements, a time referred to as recombination, and decoupled from the photons. With only small interaction cross-sections the photons began free streaming, and are observable today moving through the Universe in all directions, and following a blackbody spectrum with effective temperature $T_{\mathrm{CMB}} = 2.72548 \pm 0.00057$ [100], which we refer to as the Cosmic Microwave Background.

The CMB provides a wealth of cosmological information. Since it can be observed in all directions it presents a rigorous test of homogeneity and isotropy, which is passed with

high precision [216]. CMB radiation also contains information from a variety of interactions with matter post-recombination, such as gravitational lensing by foreground matter [39, 145, 175, 214]; fluctuations in the gravitational potential along the path of the photons, called the Sachs-Wolfe effect [237]; and scattering by moving charge particles, called the Sunyaev-Zel'dovich effect [254, 255, 256]. Together, the variations in the temperature of the CMB caused by these effects are referred to as secondary CMB anisotropies, because they occur after the photons were last scattered by the baryonic component, called the last scattering surface. The various secondary anisotropies contain significant information about the matter content and growth of structure in the Universe, and can also be combined with a variety of other probes for joint analyses.

While secondary anisotropies can be used to test many aspects of the cosmological model, the most powerful constraints come from the primary CMB anisotropies, arising from perturbations imprinted in the photon temperature before last scattering. Prior to decoupling perturbations in the photon temperature were correlated with perturbations in the matter density field. This correlation means that the temperature power spectrum of the CMB can be used to measure the initial amplitude, A_s, and slope, n_s, of the matter power spectrum.

As well as the initial power spectrum of density perturbations, the CMB temperature power spectrum also includes the imprint of the matter-radiation equality scale through the growth of perturbations. As explained in Sec. 1.1.2, prior to recombination density perturbations below super-cluster scales are unable to collapse. Instead, standing waves are produced in the baryonic density field driven by the balance between gravitational attraction and the photon pressure of the radiation component. The maximum scale of these Baryon Acoustic Oscillation (BAO)s is determined by the sound horizon, r_d: the maximum distance the waves are able to propagate before last scattering. The scale of this first peak can be used as a standard ruler in the CMB power spectrum, constraining H_0, Ω_m, and Ω_Λ through the expansion history. The position and relative heights of later peaks are determined by the physical densities of dark matter and baryons, $w_c = \Omega_c h^2$ and $w_b = \Omega_b h^2$, which combine to give $w_m = w_c + w_b = \Omega_m h^2$, providing additional constraints on these parameters. Together, the CMB primary and secondary anisotropies provide tight constraints on all of the key ΛCDM parameters [213].

Observations of the CMB are difficult to complete from the Earth due to atmospheric absorption by water vapour. Some of the most successful missions have been completed from space, including COBE [250], WMAP [28], and *Planck* [211]. *Planck* provides the most precise measurements of many cosmological parameters to date and is commonly used as the default parameter set for ΛCDM. As such it will commonly be referred to throughout this thesis, especially in terms of the results from the final major data release

in 2018. Additional observations of the CMB have been completed from the ground, especially focused on using larger diameter telescopes to exceed the angular resolution possible for space-based mission. The best resolution ground-based telescopes currently operating are the Atacama Cosmology Telescope (ACT) [104] and South Pole Telescope (SPT) [53].

1.2.3 Spectroscopic galaxy surveys

The distance ladder and CMB provide constraints from the very late and very early Universe respectively. We can obtain complementary observations by examining the matter field at intermediate redshifts through surveys of galaxies. Because distant galaxies have recession velocities that far exceed their peculiar velocities, we can use spectroscopic observations to determine the distance to the galaxies and build a 3D map of the Universe. While the positions of individual galaxies may have significant uncertainty because of their peculiar velocity and the limited exposure time that can be dedicated to each object, by observing large numbers of galaxies we can use their clustering to make statistically reliable measurements.

There are two main features targeted by these surveys for cosmological measurements. The first is the BAO scale discussed in Sec. 1.2.2, that remains imprinted in the matter density field after recombination and can be observed in the distribution of galaxies [94, 61]. Because the BAO is a fixed length scale, assuming the sound horizon is independently known it can be used as a standard ruler to constrain the expansion history [203]. These constraints are particularly useful because they occur in a different redshift regime than the distance ladder measurements, and can be observed back through cosmic history to before the transition from the matter dominated era to the dark energy dominated era. As such, the measurements of BAO from spectroscopic galaxy surveys provide an independent constraint of the expansion history with complementary degeneracies in the various parameters with distance ladder and CMB measurements (see Sec. 2.7.3 or e.g. [11]).

In addition to using the BAO scale as a standard ruler, galaxy surveys also contain a wealth of information about the growth history of structure from the velocities of the galaxies. While the recession velocity due to the expansion of the universe dominates our relative velocity with distant galaxies, those galaxies also possess an additional velocity arising from the growth of structure, termed the 'peculiar velocity'. When the redshift of a galaxy is used to infer the radial distance to the galaxy assuming that all of the velocity comes from the expansion of the Universe, as is the case for spectroscopic galaxy surveys, the peculiar velocities cause the radial positions of the galaxies to shift, called Redshift

Space Distortion (RSD) [152]. These RSD produce changes in the observed clustering of the galaxies in redshift space, and because of the link to structure growth are another source of significant cosmological information. A full description of RSD and their use as a cosmological observable is given in Sec. 1.4.

Many spectroscopic galaxy surveys have been used to constrain the expansion and growth histories of the Universe over the last several decades. Early surveys include the first observations made as part of the Sloan Digital Sky Survey (SDSS) [274], categorized as SDSS-I and -II and collected in the Main Galaxy Sample (MGS) [1]. There have also been several major surveys outside of the SDSS, such as the 6-degree Field Galaxy Survey (6dFGS) [151], Galaxy and Mass Assembly (GAMA) survey [84], WiggleZ Dark Energy Survey [83], and VIMOS Public Extragalactic Redshift Survey (VIPERS) [208]. The Baryon Oscillation Spectroscopic Survey (BOSS) [73], part of SDSS-III [93], observed the largest galaxy sample to date and provides the most precise measurements from both BAO and RSD. The Extended Baryon Oscillation Spectroscopic Survey (eBOSS) [72] is the successor to BOSS from SDSS-IV, expanding to higher redshift and introducing new galaxy samples. This thesis focuses on the analysis of eBOSS data. The Dark Energy Spectroscopic Instrument (DESI) [79, 80] is an on-going spectroscopic survey expected to increase the number of observed galaxies by an order of magnitude, and will be closely followed by the *Euclid* space mission [168].

There are several methods for constructing a multi-object spectrograph, with the choice of the SDSS and many other large surveys being a fibre-fed spectrograph. While fibres provide a versatile method of targeting many objects across a field-of-view, there are some challenges in this method. Fibre-fed surveys are limited in the radius on the sky for which two fibres can be simultaneously targeted due to the physical size of the fibres and their housing, referred to as "fibre collision". This limitation causes close objects to be systematically missed, which leads to a biased galaxy catalogue because galaxies that have correlated angular positions also have correlated radial positions. This bias becomes stronger for measurements made using smaller galaxy separation ranges, which are a focus of this thesis [226]. For that reason, an effective method is required to correct for fibre collisions. For the eBOSS analyses detailed in Ch. 2, 3, and 4, I use the combined method of Pairwise-Inverse-Probability (PIP) weighting [36] with Angular Upweighting (ANG) [202], which provides an unbiased correction to fibre collisions on all scales provided the probability that each pair of galaxies could have been observed in the survey is known, and no pairs have zero probability of being observed. Generating and applying PIP+ANG weights for eBOSS is described in Sec. 2.5.

1.2.4 Galaxy lensing surveys

In addition to spectroscopic surveys to measure galaxy clustering, purely photometric surveys of galaxies also provide significant cosmological information. The most common application of such surveys is measuring the gravitational lensing of background galaxies by foreground matter. When the source galaxy and lens matter are not closely aligned the change in the observed source galaxy is small, and the phenomenon is referred to as "weak lensing" [115, 87]. Because the intrinsic shape of galaxies varies widely it is not possible to obtain tight constraints from any single object. However, in the same way as the 3D clustering of galaxies in spectroscopic surveys can be used to produce informative clustering measurements, a statistical analysis of the correlations between galaxy shapes and foreground galaxies provides significant cosmological information [185, 40, 153].

Modern weak lensing surveys typically extract information from a set of three, two-point statistics (see e.g. [3, 136]). Two-point statistics refer to measurements that correlate the positions of two objects (see Sec. 1.3.1). In weak lensing analyses the first statistic used are the correlations between distortions in the shapes of galaxies, referred to as the galaxy shear [153, 21]. This measurement is therefore referred to as the shear-shear correlation, or the shear auto-correlation. The second is the correlation between the galaxy shear and the position of galaxies: the galaxy-shear cross-correlation. The third is the correlation between the positions of the galaxies themselves, the galaxy-galaxy correlation or galaxy auto-correlation, which provides complementary information to the shear statistics.

Because weak lensing is sourced from the foreground distribution of matter it is most sensitive to two parameters. The first is the total amount of matter, parameterized by Ω_m. The second is how clumped together the matter field is, which is described on various scales by the amplitude of the matter power spectrum, $P(k)$ (see Sec. 1.3.3). Because the shape of the initial power spectrum is known from the initial conditions (see Sec. 1.1.1), and in the linear regime the growth of structure is scale independent (see Eq. 1.12 in Sec. 1.1.2), the amplitude of the power spectrum only needs to be measured at an arbitrary scale to specify its behaviour. Typically, the standard deviation of the amplitude of mass fluctuations within a sphere of radius R, σ_R, is used for this normalization, and can be calculated from the power spectrum using:

$$\sigma_R^2(z) = \int_0^\infty P(k) \left(\frac{3}{(kR)^2} [\sin(kR) - kR\cos(kR)] \right)^2 k^2 dk \,. \tag{1.20}$$

For historical reasons the radius of the sphere is typically taken to be $R = 8\,h^{-1}$Mpc. Weak lensing measurements provide a degenerate measurement of Ω_m and σ_8, which is

often re-parameterized as

$$S_8 \equiv \sigma_8 \sqrt{\frac{\Omega_m}{0.3}}\,, \tag{1.21}$$

because S_8 is perpendicular to the degeneracy direction of the weak lensing parameter measurements, and so is the best constrained quantity.

There have been many weak lensing measurements since the first successful detection in the year 2000 [155, 264, 271]. The current leading constraints come from the Dark Energy Survey (DES) [70], Kilo-Degree Survey (KiDS) [163], and Hyper Suprime-Cam (HSC) [5] collaborations. The best constraints from these three surveys at the time of writing are $S_8 = 0.776 \pm 0.017$ [3] from DES, $S_8 = 0.766^{+0.020}_{-0.014}$ [136] from KiDS, and $S_8 = 0.804^{+0.032}_{-0.029}$ [126] from HSC, which are all mutually consistent.

1.2.5 Tensions between observations

While the ΛCDM model has been successful in matching observations from all these probes with a very small set of parameters, there do exist some tensions in the values of some parameters between probes. The most significant is between the values of H_0 determined by the distance ladder (e.g. [229]) and other probes of the late time expansion rate (e.g. [272, 206]), and the values determined from the CMB (e.g. [213, 6]) and BAO scale within spectroscopic galaxy surveys (e.g. [11]), which are calibrated by early-Universe physics. The measurements from each of these camps are generally in agreement with each other and in mild tension with the measurements of the other camp, producing an overall tension that is on the order of $4-6\sigma$ depending on analysis choices and the treatment of systematic uncertainties [4]. Fig. 1.1 shows a collection of recent H_0 measurements, where this division is clear, although exceptions do exist (e.g. [105]).

It is very difficult to reconcile these measurements through a modification to the late-Universe expansion rate because the BAO measurements from spectroscopic galaxy surveys span a wide range of redshifts and are individually consistent with a ΛCDM model, and because complementary measurements from other probes constrain many of the simplest modifications to the model [160]. Thorough investigations of possible sources of systematic error and their treatments have also been conducted for each probe without uncovering a source that can convincingly explain the difference, particularly because each camp consists of multiple probes containing measurements conducted by a variety of groups using a variety of instruments. Efforts continue to find an explanation using a modification to early-Universe physics, such as introducing early-Dark Energy to change the sound horizon and thus shift the BAO scale (e.g. [137]), but a perfect solution has yet to be discovered. As

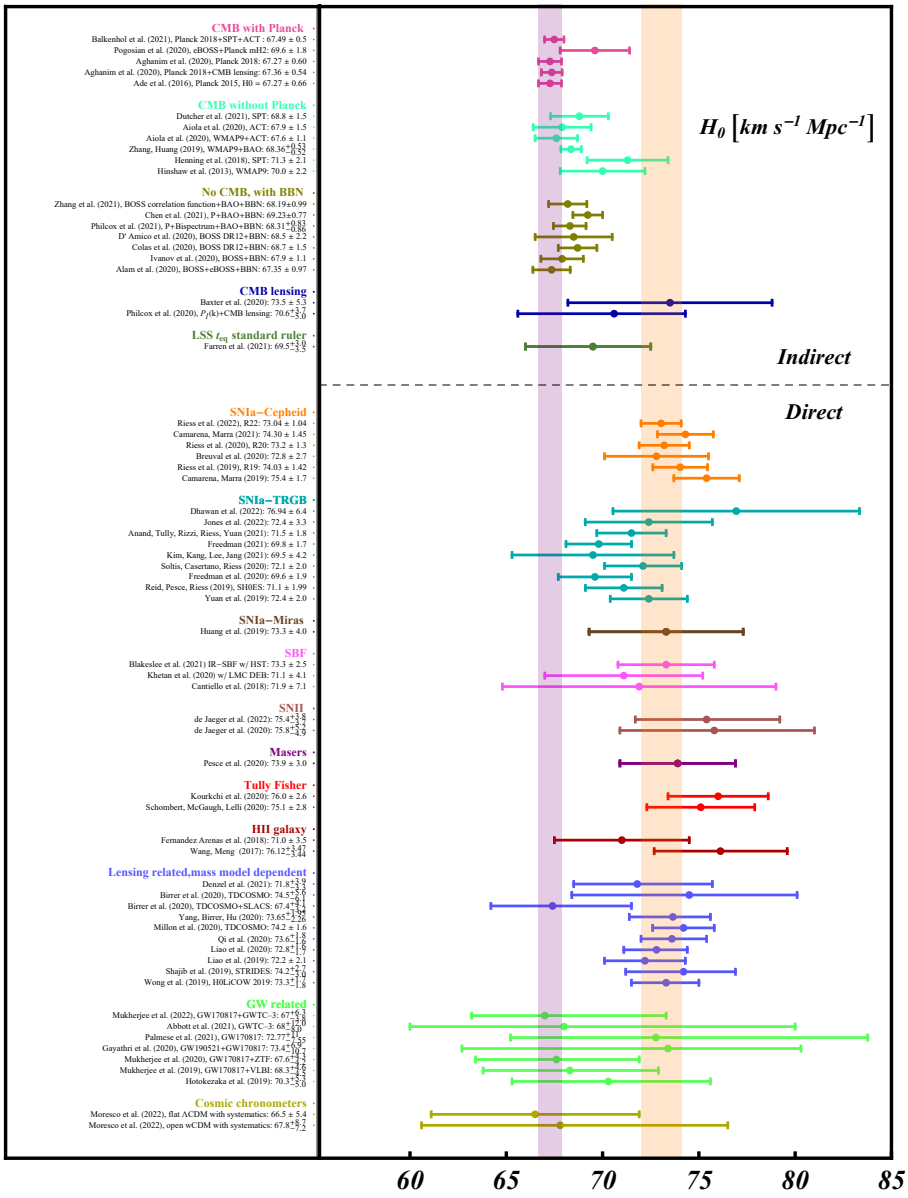

Figure 1.1: Measurements of H_0 from a variety of probes, divided into direct probes of the local expansion rate and indirect probes that assume a particular cosmological model. The purple band shows the *Planck* 2018 constraint [213], while the orange band shows a recent measurement from the SH0ES collaboration [229]. LSS refers to large-scale structure, SBF to surface brightness fluctuations, and GW to gravitational waves. Source [4].

a result, the H_0 tension remains one of the most significant unsolved problems of modern cosmology, and one of the most promising clues to discover new physics in the cosmological model.

The other most famous tension is between the measurements of S_8 inferred from observations of the CMB (e.g. [213, 6]) and the value found by weak lensing surveys (e.g. [3, 136, 126]) and other probes of the Universe at late times (e.g. [60]), shown in Fig. 1.2. While less significant than the H_0 tension, typically quoted as a $2-3\sigma$ tension depending on analysis choices, it is nonetheless concerning because it is found in all major weak lensing measurements, which are also found to agree with several other cosmological probes sensitive to the growth history of the Universe. This issue is further complicated because many theoretical methods of raising H_0 for CMB measurements also increase S_8 [148], making a simultaneous resolution for both tensions difficult to achieve. The S_8 tension remains another major unsolved issue within modern cosmology, and an important area of future research.

Each of the observations discussed above, in isolation, find a good fit to a ΛCDM cosmological model, so these tensions provide the most promising avenues for deepening our understanding of Universe and extending that model. For that reason new cosmological measurements, particularly ones that provide complementary constraints or make use of a different observational probe, are very valuable. This is particularly true of the eBOSS survey presented in Ch. 2, which extended spectroscopic galaxy surveys into a new redshift regime through observations of several novel galaxy populations to place greater constraints on H_0 and other parameters, and the small-scale RSD measurements presented in Ch. 3 and 4, which used new modelling techniques to improve constraints on structure growth, providing a complementary constraint to those presented for the S_8 tension.

1.3 Clustering measurements

In this section I introduce the two most common statistics used in galaxy clustering measurements. I begin by describing the larger class of two-point functions in Sec. 1.3.1, and the reasons that they are primary tools used for the cosmological analysis of galaxy clustering. Next, in Sec. 1.3.2, I describe the correlation function, which is a two-point function of galaxy clustering in configuration space. I provide a description of the correlation function, the methods used to calculate it, and the variations used for clustering analyses in the remainder of this thesis. Then, in Sec. 1.3.3, I provide the same details for the Fourier transform of the correlation function: the power spectrum.

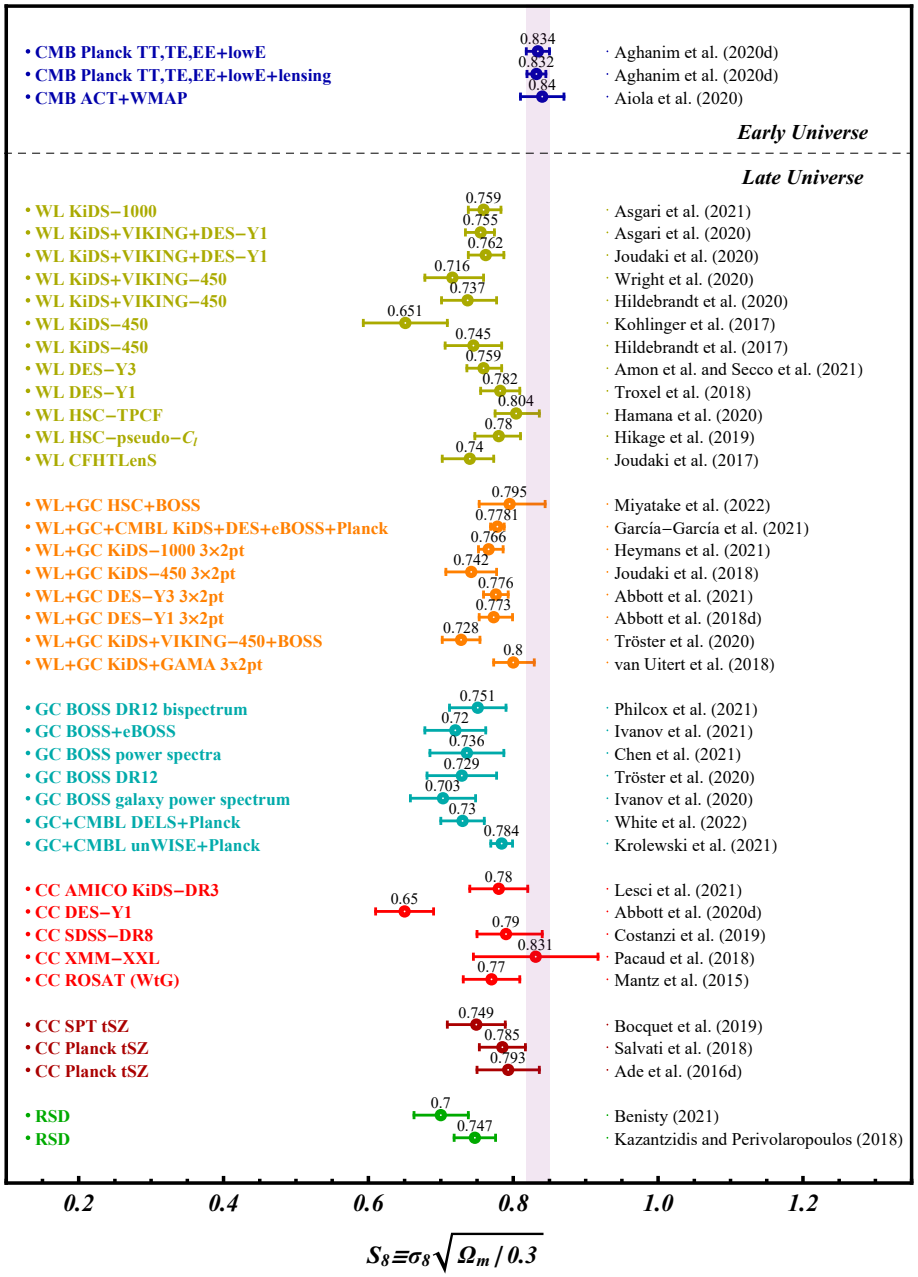

Figure 1.2: Recent measurements of S_8, divided into early Universe measurements of the CMB and late Universe measurements from a variety of probes. The purple band shows the *Planck* 2018 constraint [213]. WL refers to weak lensing, GC to galaxy clustering, and CC to cluster counts. Source [4].

1.3.1 Two-point functions

N-point functions in general describe a class of statistics that use correlations between N different points in the data catalogue to summarize the information contained within. There exist many possible statistics for a range of possible values of N, but here I will focus on two-point statistics due to their wide application in the field. The reason for the popularity of two-point functions is that the statistics of a Gaussian random field, which is a good approximation for the initial density field, are completely specified by the field's Two-point Correlation Function (2PCF) or its power spectrum. This compression of information is beneficial because for large datasets it is often expensive in terms of modelling complexity and computing resources to extract information from each individual object in a catalogue. Instead, by using a two-point function we are able to access the same information, but reduced to a smaller and more manageable set of values.

Of course, under the influence of gravity the initial field will undergo some non-linear evolution that will introduce more pronounced non-Gaussianities. Two-point statistics are thus not a complete description of the late-Universe density field, and higher order statistics are required to extract complete information [239, 200, 182]. Nevertheless, because higher order statistics require significantly more computation time, are higher-dimensional and thus require larger data vectors, and are noisier, in most analyses only two-point statistics are used because they are still able to access the majority of the information from the field.

1.3.2 Correlation function

The first two-point statistic I will discuss is the correlation function, $\xi(\boldsymbol{s})$. It describes the excess probability of finding another object in the catalogue at a given separation, $\boldsymbol{s}$, compared to a Poisson distribution. In galaxy clustering analyses, such as those described in this thesis, it is typically calculated using pair counting, where counts of each pair of objects in the data catalogue are binned in terms of the pair separation. Due to the complex nature of galaxy surveys it is often necessary to apply weights to the objects in the catalogue, either to remove observational systematics or to optimize the clustering signal. In that case a pair consisting of object i and object j will contribute a value of $w_{ij}^{\mathrm{Tot}} = w_i w_j w_{ij}$, where w_i and w_j are the weights of objects i and j respectively, and w_{ij} is any additional weight applied to the specific pair formed by i and j.

In order to assess the expected clustering for a Poisson distribution a catalogue of random points is constructed over the same angular and radial footprint, and pair counting is completed for the random catalogue as well. It is typical to use a random catalogue with

many times more points than the data catalogue to reduce the impact of shot-noise on the estimated correlation function. For example, in the eBOSS analyses described in Ch. 2 a factor of 50 times more randoms compared to the data catalogue were used [231].

The simplest estimator for the correlation function, called the natural estimator, can then be calculated as:

$$\xi(\boldsymbol{s}) = \frac{DD\boldsymbol{s})}{RR(\boldsymbol{s})} - 1, \tag{1.22}$$

where DD and RR are the weighted data-data and random-random pair counts respectively. However, while the natural estimator is simple, there exist better alternatives. The most commonly used is the Landy-Szalay estimator, which was shown to be least-bias and least-variance in [164]. The Landy-Szalay estimator is calculated as:

$$\xi\left(\boldsymbol{s}\right) = \frac{DD\left(\boldsymbol{s}\right) - 2DR\left(\boldsymbol{s}\right)}{RR\left(\boldsymbol{s}\right)} + 1, \tag{1.23}$$

where DR are weighted pair counts, with one object taken from the data catalogue and the other from the random catalogue.

For any estimator it is important to account for the difference in the number of galaxies and randoms by normalising the pair counts by the total number of distinct pairs. For the data-data pair counts this normalization is

$$(\mathrm{Norm}_{DD})^{-1} = \frac{N_D(N_D-1)}{2}, \tag{1.24}$$

where N_D is the weighted number of objects in the data catalog. The normalization for the random-random pair counts is equivalent after substituting N_R, the weighted number of randoms, for N_D, and the normalization of the data-random pair counts is simply $(\mathrm{Norm}_{DR})^{-1} = N_D N_R$. It should be noted that this calculation is valid when only individual object weights w_i and w_j are used, and not pair-based weights w_{ij}. In the case of pair-based weights, such as the Pairwise-Inverse-Probability (PIP) weights applied in later chapters, the normalization must be calculated by summing each weighted pair.

The pair separation, $\boldsymbol{s}$, is a vector quantity, and as such measuring the full 3D correlation function requires a large number of bins to measure the correlations along different directions. This is an issue for clustering measurements that empirically estimate the covariance between data points, a necessary ingredient in fitting a model to data. In order to obtain a stable estimate of the covariance matrix the number of realizations used to estimate the covariance typically must exceed the number of data points. Two of the most common methods of empirical covariance matrix estimation are from simulations and by

resampling the data, either using jackknife or bootstrap. In simulation-based estimates each simulation requires significant computing resources, while additional resampling realizations require reducing the size of the subsamples used in the resampling, which must be balanced against maintaining sufficient size that the subsamples remain an accurate representation of all the clustering information used in the analysis. As such, both methods face limitations in the number of realizations they can accommodate, so most analyses elect to compress the information of the full 3D correlation function to more manageable statistics. If the clustering of objects is isotropic it would be possible to reduce the correlation function to one-dimension by measuring in terms of the distance between pairs, $s = |\boldsymbol{s}|$. While this is the case for the real-space clustering of galaxies, in redshift-space the peculiar velocities of galaxies change their apparent position along the line-of-sight, as described in Sec. 1.4. In that case it is typical to also bin pairs in $\mu = r_\parallel/s$, the cosine of the angle between the line-of-sight direction and the pair separation vector.

The most common compression is to measure the multipoles of the redshift-space correlation function, $\xi(s, \mu)$. Multipole moments ξ_ℓ are defined as,

$$\xi_\ell(s) = (2\ell + 1) \int_0^1 \xi^s\,(s, \mu)\, L_\ell(\mu) d\mu, \tag{1.25}$$

where L_ℓ is the ℓ-order Legendre polynomial. Only the even multipole moments are non-zero, with progressively larger moments becoming noisier. For that reason the most common multipoles to use in clustering analyses are the monopole, ξ_0; quadrupole, ξ_2; and hexadecapole, ξ_4.

An alternative compression mechanism is to integrate along the line-of-sight to measure the projected correlation function, $w_p(r_\perp)$. This has the added benefit of removing the effect of RSD on the clustering. The projected correlation function $w_p(r_\perp)$ is estimated through,

$$w_p\,(r_\perp) = 2 \int_0^{r_{\parallel,\mathrm{max}}} \xi^s(r_\perp, r_\parallel) dr_\parallel, \tag{1.26}$$

where $r_\perp$ and $r_\parallel$ are the normal and parallel to the line-of-sight components of the pair separation, $\boldsymbol{s}$. $r_{\parallel,\mathrm{max}}$ is the maximum line-of-sight separation included in the integral. Integrating over the entire line-of-sight of the catalogue would completely remove RSD effects, but would require a much larger data vector. As such $r_{\parallel,\mathrm{max}}$ is usually taken to be a sufficiently large distance that the effects of RSD become minimal, on the order of $\sim 100h^{-1}$Mpc.

1.3.3 Power spectrum

The Fourier transform of the 2PCF is the power spectrum, and while it contains the same information about the density field it also has one important advantage for large-scale analyses that target the linear regime. The Fourier modes of a Gaussian random field are independent (see Sec. 1.1.2), so the large-scale modes that evolve linearly and whose growth can be predicted analytically have little contamination from small-scale non-linear modes whose evolution is much more difficult to predict. This is not the case for the correlation function, where all modes contribute to the amplitude on all scales, meaning small-scale corrections must be applied even in large-scale analyses.

Rather than pair counting, the power spectrum is estimated from the Fourier transform of the density field. The geometry of the data catalogue must still be accounted for, particularly in galaxy surveys where the geometry is highly irregular due to angular masks and a complicated radial selection function, so a random catalogue is again used to replicate the boundaries of the survey, called the survey window function. In the case of the power spectrum this is typically accomplished by dividing the survey volume into a grid and constructing a weighted over-density field using [97, 273]:

$$F(\boldsymbol{r}) = n_D(\boldsymbol{r}) - \alpha n_R(\boldsymbol{r})\,, \tag{1.27}$$

where $n_g(\boldsymbol{r})$ and $n_R(\boldsymbol{r})$ are the density of data objects and randoms at position $\boldsymbol{r}$. The scaling α corrects for the different effective numbers of data and randoms, including contributions from weights, and is calculated as:

$$\alpha = \frac{\sum_{i=1}^{N_D} w_{D,i}}{\sum_{i=1}^{N_R} w_{R,i}} \tag{1.28}$$

where $w_{D,i}$ and $w_{R,i}$ are the weights for the i-th data and random objects respectively.

As was the case with the correlation function, power spectrum analyses are usually compressed to analyze the information in the power spectrum multipoles, $P_l(k)$. The power spectrum multipoles can be estimated from the Fourier transform of $F(\boldsymbol{r})$ using [35]:

$$P_l(k) = \frac{2l+1}{I} \int \frac{d\Omega_k}{4\pi} F_0(\boldsymbol{k}) F_l(-\boldsymbol{k}) - P_l^{\text{noise}}(k)\,, \tag{1.29}$$

where I is a normalization term calculated as

$$I = \alpha \sum_{i=1}^{N_R} w_{R,i} n_{D,i}\,, \tag{1.30}$$

and $P_l^{\rm noise}(k)$ is a scale-independent shot noise component given by

$$P_0^{\rm noise} = \frac{1}{I}\left[\sum_{i=1}^{N_D} w_{D,i}^2 + \alpha^2 \sum_{i=1}^{N_R} w_{R,i}^2\right]. \tag{1.31}$$

$F_l(\boldsymbol{k})$ is the l-th multipole of the Fourier transform of $F(\boldsymbol{r})$:

$$F_l(\boldsymbol{k}) = \int d^3r F(\boldsymbol{r}) L_l(\hat{\boldsymbol{k}} \cdot \hat{\eta}) e^{i\boldsymbol{k}\cdot\boldsymbol{r}}\,, \tag{1.32}$$

where $\hat{\eta}$ is the line-of-sight direction.

1.4 Redshift space distortions

Redshift Space Distortion (RSD) are a powerful cosmological probe, and provide a unique measurement of structure growth in the Universe. I give an overview of RSD in Sec. 1.4.1, discussing the origin of RSD, how they can be used for cosmological constraints, and their effects on galaxy clustering. I then introduce the two RSD models used for the analysis of the eBOSS data in Ch. 2: Convolutional Lagrangian Perturbation Theory (CLPT) with Gaussian Streaming (GS) (Sec. 1.4.2) and the Taruya-Nishimichi-Saito (TNS) model (Sec. 1.4.3). Next, in Sec. 1.4.4 I introduce two classes of models used to extend RSD analyses towards non-linear scales: Effective Field Theory (EFT) models and simulation-based approaches. Lastly, in Sec. 1.4.5 I present some of the current constraints from measurements of RSD within galaxy clustering.

1.4.1 Theoretical overview

RSD are produced when the radial positions of galaxies are inferred from their redshifts assuming the recession velocity is due entirely to the expansion of the Universe. Galaxies possess an additional velocity component, the peculiar velocity, that is caused by the growth of structure and is sourced from gravity. The peculiar velocity shifts the position of the galaxies in redshift-space compared to real-space, distorting the clustering measured from the galaxy positions. While a nuisance in some analyses, RSD provide a unique probe of the growth of structure in the Universe and can be used to constrain key cosmological parameters. Additional constraints are vital to advancing our knowledge of the cosmological model and unravelling tensions between different probes. RSD measurements are

particularly interesting because they probe the velocity field of the galaxies, rather than directly measuring the density field like many other cosmological probes. This makes RSD measurements uniquely valuable for constraining theories of modified gravity, which produce significant changes in the expected growth of structure but are tuned to match other cosmological probes, as well as various dark energy models.

RSD have two main effects on galaxy clustering. The first is on small scales ($\lesssim 10\,h^{-1}$Mpc), where galaxies are located in virialized structures, and have highly non-linear velocities that are uncorrelated with their positions. These randomly directed velocities cause the collapsed structure to spread out along on the radial direction in redshift-space, in the so-called Finger-of-God (FoG) effect. The second occurs on large scales ($\sim 40h^{-1}$Mpc) as galaxies coherently infall to form larger structure [152]. The result is that galaxies located behind an overdensity are moved closer in redshift-space and objects located in front of the overdensity move farther away, compressing the structure along the line-of-sight, and known as the Kaiser effect. An equivalent process occurs in reverse for underdensities. The effect of RSD on the 2D redshift-space galaxy correlation function is demonstrated in Fig. 1.3. In real-space the clustering of galaxies is expected to be isotropic, but in redshift-space there is a clear stretching along the line-of-sight on small scales and compression on large scales, as expected.

The result is that in the linear regime the clustering of a biased tracer is enhanced by a factor of $(1 + \beta\mu^2)$, where $\beta = f/b$ is the ratio of the logarithmic growth rate (Eq. 1.18) and the linear bias, b [152]. While clustering is enhanced by a factor depending on β, 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 [119, 251].

1.4.2 Convolutional Lagrangian Perturbation Theory

Many models have been developed for modelling the effects of RSD on clustering measurements. Here I will focus on the Convolutional Lagrangian Perturbation Theory (CLPT) model (jointly developed in [227, 52, 267]), which is one of the most popular models in recent analyses (e.g. [143, 241]). In particular it was used in several of the eBOSS clustering analyses [23, 258] described in Sec. 2.7.2, so I will focus on the implementation used there.

CLPT is a non-perturbative resummation of Lagrangian perturbation to the two-point statistic of biased tracers in configuration space. For a given tracer the Lagrangian coordinates, $\boldsymbol{q}$, can be related to the Eulerian coordinates, $\boldsymbol{x}$ by:

$$\boldsymbol{x}(\boldsymbol{q}, t) = \boldsymbol{q} + \boldsymbol{\Psi}(\boldsymbol{q}, t) \tag{1.33}$$

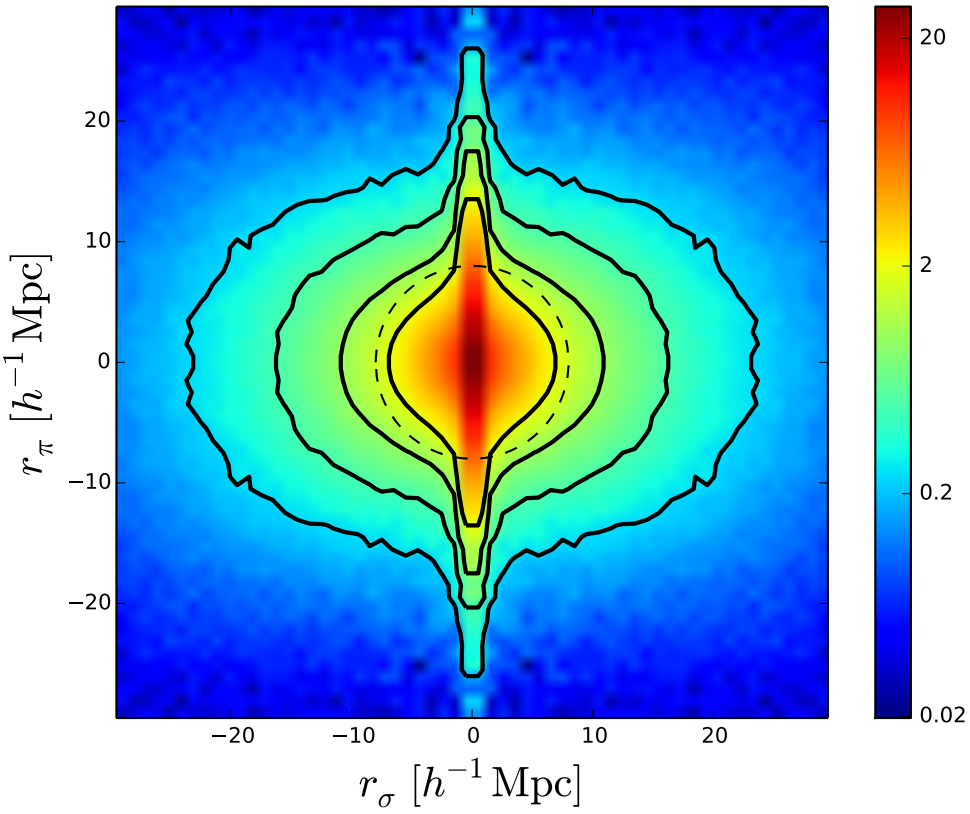

Figure 1.3: The 2D correlation function, $\xi(r_\sigma, r_\pi)$, of the SDSS-III CMASS galaxies from [226]. $r_\sigma \equiv r_\perp$ is the separation perpendicular to the line-of-sight, and $r_\pi \equiv r_\parallel$ is the separation parallel to the line-of-sight. For details of the correlation function see Sec. 1.3.2, but for the purpose of this demonstration a larger correlation function (the red end of the colourbar) indicates stronger clustering. The dashed circle shows the separation scale ($\sim 8h^{-1}$Mpc) at which the RSD transition from the large-scale behaviour (Kaiser effect) to the small-scale behaviour (Finger-of-God effect). The solid black curves show the contours at constant values of ξ=[2, 1, 0.5, 0.25] to demonstrate the deviation from circular symmetry. [226]

where $\boldsymbol{\Psi}(\boldsymbol{q}, t)$ is the displacement field evaluated at $\boldsymbol{q}$ and time t. The tracer used in the clustering analysis, denoted by subscript X, is assumed to be locally biased with respect to the matter overdensity, $\delta(\boldsymbol{q})$. The two-point correlation function is then expanded over different orders of the Lagrangian bias function, $F[\delta(\boldsymbol{q})]$, which is defined as:

$$1 + \delta_X(\boldsymbol{q}, t) = F[\delta(\boldsymbol{q})] \tag{1.34}$$

The Eulerian density contrast field is computed by convolving with the displacement:

$$1 + \delta_X(\boldsymbol{x}) = \int d^3q F[\delta(\boldsymbol{q})] \int \frac{d^3k}{(2\pi)^3} e^{i\boldsymbol{k}(\boldsymbol{x}-\boldsymbol{q}-\boldsymbol{\psi}(\boldsymbol{q}))} \tag{1.35}$$

The local Lagrangian bias function is approximated by a non-local expansion using its first and second derivative, where the n^{th} derivative is given by:

$$\langle F^n \rangle = \int \frac{d\delta}{\sqrt{2\pi}\sigma} e^{-\delta^2/2\sigma^2} \frac{d^n F}{d\delta^n}. \tag{1.36}$$

The two-point correlation function can then be obtained by evaluating the expression $\xi_X(\boldsymbol{r}) \langle \delta_X(\boldsymbol{x}) \delta_X(\boldsymbol{x} + \boldsymbol{r}) \rangle$, which can then be simplified (see Eq. 46 of [52]):

$$1 + \xi_X(\boldsymbol{r}) = \int d^3q M(\boldsymbol{r}, \boldsymbol{q}), \tag{1.37}$$

where $M(\boldsymbol{r}, \boldsymbol{q})$ is the kernel of convolution taking into account the displacement and bias expansion up to its second derivative term.

To this point the model has only included the effects of bias. The pairwise mean velocity $v_{12}(r)$ and the pairwise velocity dispersion $\sigma_{12}(r)$, where r is the real-space pair separation, can be computed using the method developed in [267]. This method is similar to the one described above, but with a modified kernel that takes into account the velocity rather than the density:

$$v_{12}(r) = (1 + \xi(\boldsymbol{r}))^{-1} \int M_1(\boldsymbol{r}, \boldsymbol{q}) d^3q, \tag{1.38}$$

and

$$\sigma_{12}(r) = (1 + \xi(\boldsymbol{r}))^{-1} \int M_2(\boldsymbol{r}, \boldsymbol{q}) d^3q. \tag{1.39}$$

The kernels $M_{1,2}(\boldsymbol{r}, \boldsymbol{q})$ also depend on the first two non-local derivatives of the Lagrangian bias $\langle F' \rangle$ and $\langle F'' \rangle$, which are free parameters in the model.

The small-scale modelling is further improved by mapping the real-space CLPT model of the two-point statistics into redshift space using the Gaussian Streaming (GS) model

[227]. In this model the pairwise velocity distribution of tracers is assumed to have a Gaussian distribution that depends on both the separation, r, and the cosine of the angle between the separation vector and the line-of-sight, μ.

What follows is the [267] implementation that uses CLPT results as input for the GS model. The redshift-space correlation function can finally be computed as:

$$1+\xi_{\rm X}(r_\perp,r_\parallel)=\int\frac{1}{\sqrt{2\pi\left[\sigma_{12}^2(r)+\sigma_{\rm FoG}^2\right]}}[1+\xi_{\rm X}(r)]\exp\left(-\frac{[r_\parallel-y-\mu v_{12}(r)]^2}{2\left[\sigma_{12}^2(r)+\sigma_{\rm FoG}^2\right]}\right)dy, \quad (1.40)$$

where $\xi(r)$, $v_{12}(r)$, and $\sigma_{12}(r)$ are obtained from CLPT. The exponential factor in the integral takes into account the scale-dependent halo-halo pairwise velocity, and introduces an additional parameter, $\sigma_{\rm FoG}$, that corrects for the FoG effect. In summary, given a fiducial cosmology, this RSD model can predict the redshift-space correlation function for a biased tracer of the density field using four free parameters, $[f, F', F'', \sigma_{\rm FoG}]$.

1.4.3 Taruya-Nishimichi-Saito model

The other RSD model used in the eBOSS clustering analyses [23, 111, 76, 141, 196] is the Taruya-Nishimichi-Saito (TNS) model developed in [259]. This model is based on the conservation of number density in real- and redshift-space [152], and produces a model for the power spectrum of biased tracers in redshift space. As such it is most often used for Fourier-space analyses, unlike the configuration-space CLPT model, although it can also be used to model the redshift-space two-point correlation function by taking the Hankel transform of the final model power spectrum multipoles.

To begin with, we can use the general form of the anisotropic power spectrum for unbiased tracers [244]:

$$P^s(k,\mu)=\int\frac{d^3\boldsymbol{r}}{(2\pi)^3}e^{-i\boldsymbol{k}\cdot\boldsymbol{r}}\left\langle e^{-ikf\mu\Delta u_\parallel}[\delta(\boldsymbol{x})+f\partial_\parallel u_\parallel(\boldsymbol{x})][\delta(\boldsymbol{x}')+f\partial_\parallel u_\parallel(\boldsymbol{x}')]\right\rangle \quad (1.41)$$

where $\mu=k_\parallel/k$ is the cosine of the wavevector with the line-of-sight and δ is the matter density field. Furthermore, it is convenient to define $u_\parallel(\boldsymbol{r})=-v_\parallel(\boldsymbol{r})/(faH(a))$ from $v_\parallel(\boldsymbol{r})$, the line-of-sight component of the peculiar velocity, and $\Delta u_\parallel=u_\parallel(\boldsymbol{x})-u_\parallel(\boldsymbol{x}')$ and $\boldsymbol{r}=\boldsymbol{x}-\boldsymbol{x}'$. In the TNS model, Eq. 1.41 can be rewritten as

$$P^s(k,\mu)=D(k,\mu,\sigma_v)\left[P_{\delta\delta}(k)+2\mu^2fP_{\delta\theta}(k)+\mu^4f^2P_{\theta\theta}(k)+C_A(k,\mu,f)+C_B(k,\mu,f)\right], \quad (1.42)$$

where θ is the divergence of the velocity field, defined as $\theta = -\nabla\cdot\mathbf{v}/(aHf)$. $P_{\delta\delta}$, $P_{\theta\theta}$ and $P_{\delta\theta}$ are the non-linear matter density, velocity divergence, and density-velocity divergence power spectra respectively. $C_A(k,\mu,f)$ and $C_B(k,\mu,f)$ are two correction terms that reduce to integrals of the matter power spectrum, and are given in [259]. $D(k,\mu,\sigma_v)$ is a phenomenological function that parameterizes damping of the power spectrum from multiple sources, including the FoG effect. This function usually takes the form of a Gaussian or Lorentzian, such as the Lorentzian damping function used in the configuration space analysis of the eBOSS LRG sample [23],

$$D(k,\mu,\sigma_v) = (1 + k^2\mu^2\sigma_v^2)^{-1}, \tag{1.43}$$

where σ_v represents an effective pairwise velocity dispersion.

So far we have included the effects of RSD for the analysis of an unbiased tracer. The model can be generalized to the case of biased tracers and rewritten as

$$P_{\rm g}^s(k,\mu) = D(k\mu\sigma_v)\left[P_{\rm gg}(k) + 2\mu^2 f P_{{\rm g}\theta} + \mu^4 f^2 P_{\theta\theta}(k) + C_A(k,\mu,f,b_1) + C_B(k,\mu,f,b_1)\right] \tag{1.44}$$

where b_1 is the galaxy linear bias. Explicit expressions for $C_A(k,\mu,f,b_1)$ and $C_B(k,\mu,f,b_1)$ can be found in [74].

$P_{\rm gg}$ is calculated by assuming a particular bias model (such as that presented in [19]) and applying it to analytically calculated non-linear power spectrum $P_{\delta\delta}$. $P_{{\rm g}\theta}$ is calculated by applying the bias model to the $P_{\delta\theta}$ power spectrum, which along with the $P_{\theta\theta}$ power spectrum can be calculated using the universal fitting functions of [27]:

$$\begin{aligned} P_{\theta\theta}(k) &= P_{\rm lin}(k) e^{-k\left(a_1 + a_2 k + a_3 k^2\right)}, \\ P_{\delta\theta}(k) &= \left[P_{\delta\delta}(k) P_{\rm lin}(k)\right]^{\frac{1}{2}} e^{-\frac{k}{k_\delta} - bk^6}. \end{aligned} \tag{1.45}$$

where $P_{\rm lin}(k)$ is an analytically calculated linear power spectrum.

The overall degree of nonlinear evolution is encoded by the amplitude of the matter fluctuations, σ_8, at the effective redshift of the sample under consideration. The explicit dependence of the fitting function coefficients on σ_8 is given by

$$\begin{aligned} a_1 &= -0.817 + 3.198\sigma_8 \\ a_2 &= 0.877 - 4.191\sigma_8 \\ a_3 &= -1.199 + 4.629\sigma_8 \\ 1/k_\delta &= -0.017 + 1.496\sigma_8^2 \\ b &= 0.091 + 0.702\sigma_8^2. \end{aligned} \tag{1.46}$$

In total, the TNS model has a minimum of three free parameters, $[f, b_1, \sigma_v]$, with more free parameters if the chosen bias model has additional degrees of freedom. Like the CLPT model it is able to accurately predict the redshift space clustering of a biased tracer within the linear regime. Because the TNS model is most naturally applied in Fourier-space analyses, while the CLPT model is well-suited for configuration space analyses, both models have been frequently applied in recent analyses, as was the case for eBOSS (see Sec. 2.7.2).

1.4.4 Extending RSD modelling towards non-linear scales

While the perturbation theory-based models discussed in Sec. 1.4.2 and 1.4.3 are accurate for linear and quasi-linear scales ($\gtrsim 20\,h^{-1}$Mpc), they break down at the shell-crossing scale. This failure requires analyses based on these models to implement hard scale-cuts to prevent contamination from non-linear effects that are not well-represented in the models. While most RSD analyses to date have abided by this limitation, other analyses that have included smaller scale clustering measurements using methods described below have found significant improvements in precision [166, 57, 277, 275, 58], such as the factor of 2.5 improvement found by [226] in their analysis of the SDSS-III BOSS CMASS sample compared to the large-scale analysis of the same sample [238].

The first method I will discuss is the class of Effective Field Theory (EFT) models [22, 54, 217]. These models make use of the relatively weak link between the small-scale non-linear structure of galaxy formation and the typical separation of galaxies in large-scale structure surveys [22, 54]. By integrating out short-wavelength perturbations (on scales $< 5\,h^{-1}$Mpc) it becomes possible to solve the resulting smoothed field with a high degree of accuracy into the quasi-linear regime by extending the perturbation theory calculations to arbitrarily high-order [69, 149, 60]. While these methods are successful at modelling the distribution of matter in the linear and quasi-linear regimes (typically limited to $> 30\,h^{-1}$Mpc, e.g. [60, 69]), they can not provide an analytic basis for the formation of galaxies or the non-linear motion of virialized structures. These effects are instead included as additional correction terms whose functional form can be predicted from perturbation, but with unknown amplitudes that must either be calibrated from simulations or fit from the data [51].

An alternative method is to attempt to fully model the formation of non-linear structure using N-body simulations. Ref. [226] used an N-body simulation at a single fixed cosmology to model the clustering of galaxies in the analysis cited above. This method has been expanded through the use of machine learning emulators to allow for varying cosmology

without needing to run additional N-body simulations for each new point in parameter space, finding similar improvements in precision over perturbation theory approaches [57, 277, 275, 161, 58]. Sec. 1.5 contains a more detailed overview of N-body simulations, emulators, and the other tools required to apply this method.

1.4.5 Growth constraints from RSD measurements

The models detailed above and their predecessors have been applied to spectroscopic galaxy surveys in order to produce constraints on $f\sigma_8$. A selection of these constraints for several of the largest galaxy surveys are shown in Fig. 1.4. Only results published prior to the year 2018 have been included in order to present a picture of the leading constraints prior to the publication of eBOSS and small-scale analyses that access the non-linear regime, presented in the remainder of this thesis. Results have been included from the 6dFGS [151], SDSS MGS [1], GAMA survey [84], WiggleZ Dark Energy Survey[83], BOSS [73], and VIPERS [208], which were the largest spectroscopic galaxy surveys prior to the release of the eBOSS data.

Comparing the results presented in Fig. 1.4 to the black line, which shows the expectation for a ΛCDM universe with *Planck* 2018 cosmological parameters [213], the measurements are very consistent. However, it is interesting to note that the majority of measurements, particularly the most precise ones, lie below the line. When considered in tandem with the S_8 tension introduced in Sec. 1.2.5, which finds that late time measurements generally prefer a lower value of σ_8 compared to CMB measurements, it raises an interesting question for what the result would be if we were able to improve the uncertainty of these constraints. This is precisely the objective of small-scale RSD measurements, which achieve higher precision and find some mild tension with the *Planck*+ΛCDM expectation (see Ch. 3 and Ch. 4). It is also worth noting that the measurements from these surveys are mostly concentrated in the redshift range $z \leq 0.6$. One of the primary advances of eBOSS was extending spectroscopic galaxy clustering measurements to a higher redshift range, with three populations of galaxies observed for clustering measurements between $0.6 \leq z \leq 2.2$ (see Ch. 2).

The models and measurements presented in this section are effective tools for deepening our understanding of the Universe. RSD measurements provide a unique and powerful constraint on the growth of structure, which is key to confirming or rejecting extensions to the ΛCDM model (Sec. 1.1.3) and unravelling the mystery of tensions between data sets (Sec. 1.2.5). Data from the eBOSS (Ch. 2) and small-scale clustering analyses (Ch. 3 and 4) have enabled significant advances in this field, and provided us with some of the most advanced tests of structure growth through RSD to date.

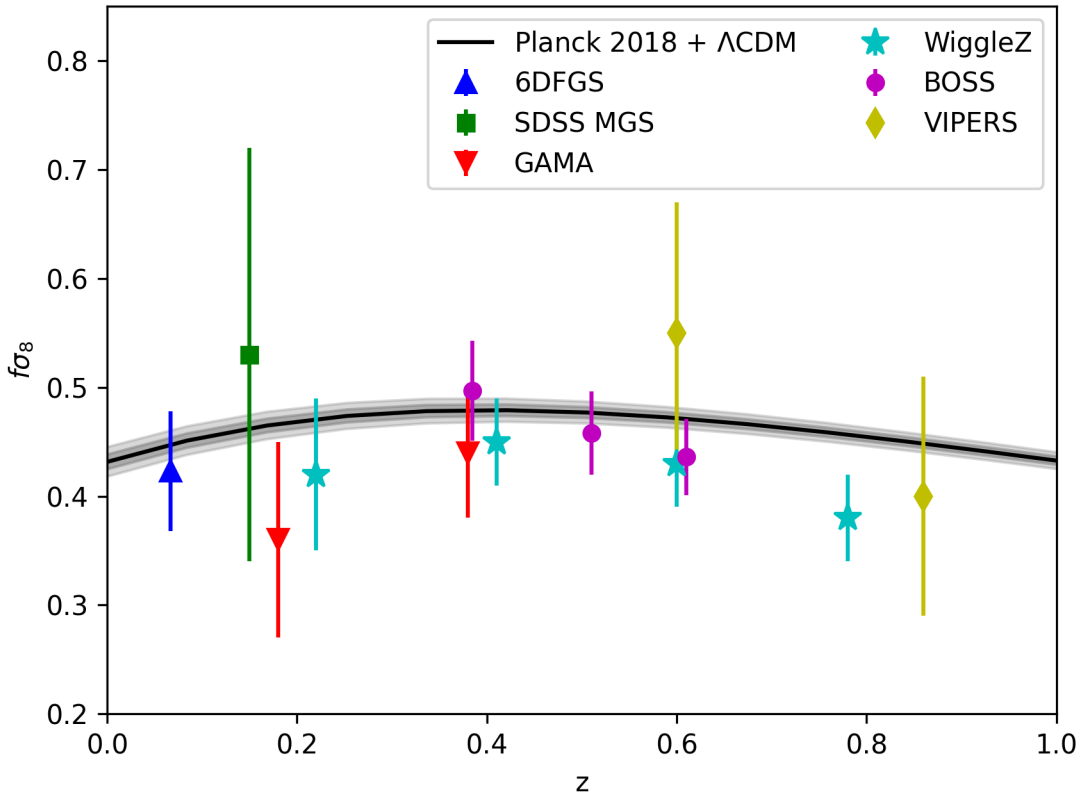

Figure 1.4: $f\sigma_8$ constraints from RSD measurements published prior to 2018, showing many of the leading constraints prior to the publication of eBOSS data and small-scale analyses. Results are shown for the 6dFGS (blue, [32]), SDSS MGS (green, [143]), GAMA survey (red, [37]), WiggleZ Dark Energy Survey (cyan, [38]), BOSS (magenta, [9]), and VIPERS (yellow, [208]). The black line shows the expected value of $f\sigma_8$ for a flat ΛCDM universe with best fit *Planck* 2018 cosmology [213], with the shaded regions showing the 1 and 2σ confidence regions.

1.5 Modelling small-scale galaxy clustering with simulations

In the strongly non-linear regime perturbation based approaches break down, and it becomes difficult to construct an analytic model for galaxy clustering. An alternative method that allows modelling clustering statistics to very small scales is to model the non-linear structure growth using simulations that include full gravitational interactions. I begin in Sec. 1.5.1 by describing the dark matter simulations used in this modelling. Then in Sec. 1.5.2 I provide an overview of dark matter halos, which are overdense regions of dark matter within which galaxies form. In order to compare simulations of dark matter to observations of galaxies we require a model for how galaxies occupy the underlying dark matter halos, which I describe in Sec. 1.5.3. Finally, in Sec. 1.5.4 I discuss how the large computational cost of simulations can be mitigated through the use of machine learning emulators of galaxy clustering, which is the approach used in Ch. 3 and 4 to obtain higher precision measurements of $f\sigma_8$ from small-scale clustering.

1.5.1 N-body simulations

Perturbation theory is able to provide accurate predictions down to the shell-crossing scale ($r \gtrsim 20\,h^{-1}\mathrm{Mpc}$, $k \lesssim 0.3\,h\mathrm{Mpc}^{-1}$), but beyond this scale purely analytic models become inaccurate [204, 56, 260]. However, because the matter density field of the Universe is dominated by dark matter, whose only detected interaction is through gravity, simulating the evolution of the density field can be relatively straight-forward computationally. The simplest of these simulations represent the evolution of the density field through the gravitational evolution of a set of N dark matter particles, and so are referred to as N-body simulations.

The key characteristics of an N-body simulation are the box size and the number of particles. Simulations are usually run using a box with periodic boundary conditions, so the side length of the box determines the total volume simulated. Simulations must be of sufficient volume to match the volume of the data set they are constructed to model, because only perturbations on scales smaller than the box size will be accurately simulated, although there exist methods to account for the variance of perturbations on scales larger than the box size [177]. The number of particles is important because simulations are run for a fixed set of cosmological parameters, so that total matter density in the box is set by Ω_m. As such, the number of particles determines the mass of each particle, which in turn determines the smallest mass scale that can be resolved in the simulation.

Typically these simulations are begun from initial conditions set early enough in the Universe that the density field evolution is entirely linear, and so can be modelled analytically. The most common method is to begin with a uniform distribution of particles, then displace them to match an analytically determined initial power spectrum using 1-loop Lagrangian perturbation theory, known as the Zel'dovich approximation (see e.g. [109]). Once the linear initial conditions are set the simulation can be evolved in time steps to the present day, $z = 0$. Rather than storing the state of the simulation at each time step, most simulations save only a subset of time steps that span the redshift range of interest, referred to as redshift slices.

These simulations have been shown to produce an accurate representation of the overall matter density field through dark matter [133, 243], however dark matter-only simulations are unable to reproduce baryonic structure such as galaxies, or reproduce their more complicated effects on the density field on small scales (see e.g. [140] and references therein). There exist a variety of methods to address this shortcoming. The least computationally expensive, which is discussed below in Sec. 1.5.3 in more detail, is to run the N-body simulation, then use an empirical model for how galaxies occupy the dark matter density field to populate the simulation with galaxies [30, 281, 25, 269]. The next is to approximate the evolution of the density field by including both dark matter and baryonic elements on a grid, and using empirically motivated functions to describe how the grid will evolve [29]. These are referred to as semi-analytic models, and are computationally cheap to produce, but depend strongly on the choices for the sub-grid physics. The last is to run a simulation that includes physically motivated evolution of both dark matter and baryon particles through hydrodynamical simulations [140]. While these simulations include the most fundamental description for the evolution of the baryon field they are also the most expensive to run, requiring reduced volume and mass resolution. And while they are the most physically motivated in the evolution of the baryon field they do still require applying empirically motivated sub-grid physics, because resolving individual star formation and evolution remains far beyond the capabilities of our cosmological simulations.

1.5.2 Dark matter halos

While dark matter-only N-body simulations provide an accurate representation of the dark matter density field for a given cosmology, we require statistical tools that can be used to relate the observed tracers of the density field, galaxies, to the simulations. To that end it is useful to introduce the halo model, which is a framework for describing the underlying density field and its tracers [66, 262]. The core of the model is that galaxies form within dark matter halos, which are collapsed regions where the density of dark matter is on the

order of hundreds of times denser than the background field. On average these halos are believed to be spherically symmetric, and their properties are determined entirely by the mass of the halo.

The key ingredients to the halo model are the halo bias, that describes how halos trace the underlying density field; a halo mass function, that describes the number of halos per unit mass; and the halo profile, that describes how the density of the halo changes as a function of radius. In analytic applications of the halo model a specific functional form must be assumed for each of these ingredients, and the parameters of those functions tuned from simulations.

For clustering models based off of simulations two of these three ingredients, the halo bias and halo mass function, can be determined directly from the simulation by using a halo finder to identify the halos within the simulation volume. These halo finders work by identifying sets of particles within the simulation that are tightly associated with one another, and identifying each set above a mass resolution threshold as being a dark matter halo. There are a variety of halo finder algorithms available. The simplest are referred to as Friends-of-Friends (FoF) algorithms. The key parameter in this method is a linking length, which is a spatial scale within which it is assumed two particles are members of the same halo. Halos are defined by finding all particles that are continuously connected to at least one other member by the linking length.

A more advanced method relies on matching particles that are closely associated both in position and velocity-space, and is applied using the ROCKSTAR[1][26] code. ROCKSTAR functions in a similar way to more basic FoF algorithms, but linking in the six-dimensional phase space defined by the particle positions and velocities, as well as in time. While this method has been shown to be more accurate in recovering halo properties than basic FoF algorithms [159], it is also more computationally expensive. ROCKSTAR is the halo finder applied throughout this thesis, unless otherwise specified.

The last halo property, the halo density profile, has been extensively studied using detailed simulations. [194] found that all halos could be fit using a universal density profile,

$$\rho(r) = \frac{\rho_0}{\frac{r}{R_s}\left(1+\frac{r}{R_s}\right)^2}\,, \tag{1.47}$$

where ρ_0 and R_s are parameters that vary from halo to halo. This profile is referred to as the Navarro-Frenk-White (NFW) profile, after the authors of the original study.

[1] https://github.com/yt-project/rockstar

A common threshold to define a dark matter halo is the region where the average density is 200 times the mean density of the Universe, with radius given by R_{200}. A final characteristic quantity of the profile is the concentration, $c = R_{200}/R_s$, which relates the scale radius, R_s, to the radius of the nominal edge of the halo.

1.5.3 Galaxy-halo connection

Once the halo catalogue has been assembled from the simulation, we require a model to populate those halos with galaxies. Unlike large-scale bias models, which provide an analytic formalism under a set of assumptions such as that galaxies are linearly biased relative to the overall mass distribution [186], on small scales the non-linear formation of halos and galaxies requires a more nuanced model. Because these conditions can not be predicted analytically the models used to connect galaxies to halos are phenomenological, in that they are motivated by some combination of theoretical and empirical expectation from simulations and data.

A relatively simple galaxy-halo connection model is abundance matching, also known as Subhalo Abundance Matching (SHAM) when extended to smaller halos that have been accreted within larger host halos. The primary motivation for these models is that observations have shown that halo mass and galaxy mass are tightly correlated [25]. In abundance matching galaxies are placed within halos monotonically based off of the halo mass, with the most massive halos receiving galaxies until the desired number density is achieved, matching what is observed in actual data [192]. In reflection of the fact that there is variance within the actual relationship between galaxy occupation and halo properties [261] it is also typical to include a parameter controlling variance in the primary halo property used in abundance matching. This parameter can be thought of as accounting for scatter in the observed relationship between galaxy mass and halo mass, and is determined empirically. While halo mass is the most natural property to use in abundance matching, it has been found that the peak circular velocity of the halo, $V_{\rm peak}$, has lower variance in the occupation of galaxies, so it is often used as the primary halo property in place of mass [269].

An alternative to abundance matching is the class of Halo Occupation Distribution models [30, 281]. Rather than the monotonic relationship of abundance matching, HOD models are based on the principle that each halo has a probability of hosting a galaxy determined by the mass of the halo. The occupation distribution as a function of mass is controlled by a set of parameters that provide an empirically motivated functional form, such as the observation that more massive halos are more likely to host galaxies, while

giving the model sufficient freedom to span the range of possible halo and galaxy formation processes. This occupation distribution is typically separated into the occupancy of central and satellite galaxies [281],

$$\langle N(M) \rangle = N_{\rm cen}(M) + N_{\rm sat}(M), \tag{1.48}$$

where each halo can be occupied by up to one central galaxy with additional satellite galaxies.

There are several possible parameterizations for $N_{\rm cen}$ and $N_{\rm sat}$ based on the perceived behaviour of the galaxy population. For populations of high mass galaxies commonly found in groups and clusters of galaxies, such as the Luminous Red Galaxy (LRG) sample analyzed in Ch. 3 and Ch. 4, a typical parameterization is for the central occupancy to transition from 0 to some maximum occupation, and for satellite galaxies to follow a Poisson distribution with first moment given by a power law for high mass halos. These choices ensure that most galaxies will be found in large halos hosting groups and clusters, with larger halos hosting more galaxies. The particular parameterization adopted by the analyses presented in this thesis, originally based on the model presented in [282], is [278, 57, 58]:

$$N_{\rm cen}(M) = \frac{f_{\rm max}}{2}\left[1 + {\rm erf}\left(\frac{\log_{10} M - \log_{10} M_{\rm min}}{\sigma_{\log M}}\right)\right], \tag{1.49}$$

$$N_{\rm sat}(M) = \left(\frac{M}{M_{\rm sat}}\right)^{\alpha} \exp\left(-\frac{M_{\rm cut}}{M}\right) \frac{N_{\rm cen}(M)}{f_{\rm max}}. \tag{1.50}$$

The free parameters of this model are $f_{\rm max}$, $\log M_{\rm min}$, $\sigma_{\log M}$, $\log M_{\rm sat}$, α, and $\log M_{\rm cut}$.$f_{\rm max}$ defines the maximum fraction of high mass halos that will contain a central galaxy. For a complete galaxy sample $f_{\rm max} = 1$, with $f_{\rm max} < 1$ indicating that some central galaxies are missing even from the most massive halos. $\log M_{\rm min}$ is the halo mass where the occupancy of centrals transitions from 0 to $f_{\rm max}$, effectively setting the minimum halo mass to host a central, and $\sigma_{\log M}$ sets the width of that transition. For samples dominated by centrals, such as LRGs, these three parameters have a strong impact on the number density of the sample. For the satellite occupancy, $M_{\rm sat}$ is the typical mass for halos to host one satellite, α is the power-law index for the mass dependence of the satellite occupation, and $M_{\rm cut}$ gives an exponential cutoff to the satellite occupation at low mass. Together these parameters determine the shape of the occupancy distributions, as shown by Fig. 1.5, and must be tuned to match the formation and bias of a particular galaxy population.

The six parameters listed above determine the number of galaxies that occupy each halo, however HOD models must also adopt a method of assigning positions and velocities to the galaxies. Here I will again follow the methods applied in Ch. 3 and Ch. 4, as they

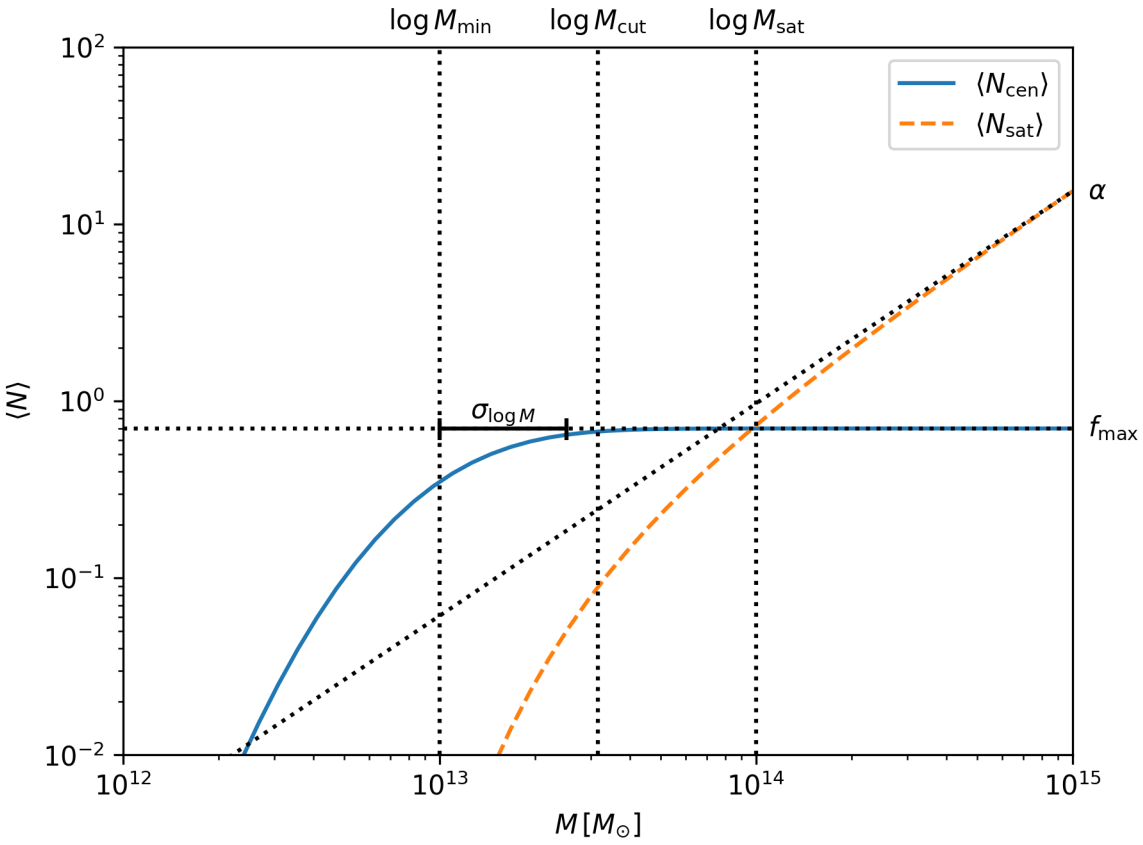

Figure 1.5: Example of a halo occupation distribution with labelled parameters. The blue line shows the expected occupation of central galaxies as a function of halo mass, while the dashed orange line shows the same for satellite galaxies. Labelled dotted lines show the values of the model parameters.

are the most relevant to this work and representative of typical methods used in the field. Central galaxies are assigned a position at the centre of their host halo, which in turn is determined by the halo finding algorithm from the particle positions. Satellite galaxies are assumed to follow the same distribution as the dark matter particles, so are randomly distributed with spherical symmetry about the halo centre with a probability distribution that follows a NFW profile. The concentration of the NFW profile is calculated from the mass of the halo using [46]

$$c(M) = \frac{c_0}{1+z}\frac{M}{M_*}^{\beta}, \tag{1.51}$$

where $c_0 = 11$, $\beta = -0.13$, and M_* is the nonlinear mass at $z = 0$, which depends on the fiducial cosmology used in the analysis. An additional parameter is added to the model, $c_{\rm vir} = c_{\rm sat}/c_{\rm halo}$, to allow the satellite concentration, $c_{\rm sat}$, to differ from the concentration of the host halo NFW, $c_{\rm halo}$.

Both centrals and satellites are assigned the velocity of their host halo, plus a random velocity dispersion controlled by separate velocity bias parameters: $v_{\rm bc}$ for centrals and $v_{\rm bs}$ for satellites. The velocity bias is calculated as a scaling of the halo velocity dispersion, given by $\sigma_{\rm gal} = v_{\rm gal}\sigma_{\rm halo}$. In Ch. 3 and Ch. 4 the halo velocity dispersion, $\sigma_{\rm halo}$, is calculated from the circular velocity of the halo at the virial radius, $\sigma_{\rm halo} = \sqrt{GM(< R_{\rm vir})/R_{\rm vir}}$. This method allows the velocity dispersion of all halos to be quickly and simply calculated, instead of alternative methods such as estimating the halo velocity dispersion from the dispersion of the constituent particles, which is more computationally expensive and noisy for low mass halos that contain fewer particles. The default velocity bias parameter for centrals is $v_{\rm bc} = 0$, meaning no velocity bias, and the default velocity bias parameter for satellites is $v_{\rm bs} = 1$, meaning they approximately follow the internal velocity dispersion of the host halo.

The general purpose of galaxy-halo connection models is to provide an empirically motivated method of placing galaxies within halos because the formation mechanisms of halos and galaxies can not be determined analytically. The benefit of the HOD model in particular is that the model parameters, as well as the choice of parameterization in general, gives the model the freedom to match a variety of galaxy populations exhibiting different bias relations, and can be marginalized over to account for deviations between the simulations or HOD parameterization and the true halo and galaxy formation mechanisms of the Universe.

1.5.4 Emulators

While N-body simulations provide an effective model for the non-linear clustering of halos and galaxies, it is computationally expensive to simulate large enough volumes (on the order of $\sim$ (Gpc)^3) to match modern galaxy surveys while maintaining the necessary particle mass to resolve the individual halos hosting galaxies (on the order of $10^{12}\,M_\odot$). This is an issue for analyses that must explore a multi-dimensional parameter-space, because it is not feasible to run a new simulation for each point in parameter-space. This was the challenge faced in the early small-scale RSD analysis performed by [226] (see discussion in Sec. 1.4.4), so that they were ultimately forced to use a single-fixed cosmology simulation to generate their clustering model. This reliance on a fixed cosmology leaves unanswered questions as to the reliability of the final measurements, and the size of the inferred uncertainty.

While it is generally not possible to generate a new simulation for each point in parameter-space, given a set of simulations it is possible to use machine learning to produce an emulator that is able to make predictions for any point in parameter-space within the span of the initial set of simulations [133, 135, 170, 134]. Once trained an emulator can typically produce a prediction on the order of O(1s) or faster, allowing them to be effectively used to explore the posterior in a Bayesian analysis. There are several possible machine learning methods that can be used to construct an emulator, with the most common being a Gaussian process (see e.g. [278]) or a neural network (see e.g. [235]). The strength of neural networks is that they are able to reproduce any arbitrary functional form, allowing them to effectively reproduce the desired behaviour from the training data regardless of complexity. The major drawback is that they can be susceptible to overfitting, making it difficult to ensure that the information being obtained is truly cosmological in origin. On the other hand Gaussian processes are a non-parametric method, meaning no specific functional form needs to be assumed, and are very effective at interpolating within a sparse parameter space, which is why Gaussian processes are used as the basis for the emulators applied in this work.

A Gaussian process is a collection of random variables such that each point in parameter space follows a multivariate normal distribution. For a simulation-based inference of galaxy clustering, the parameter-space consists of the combined cosmological and galaxy-halo connection model parameters. Initially the multivariate normal distribution has an arbitrary mean and variance, but the mean and variance at a particular point in parameter-space can be constrained by measuring the desired clustering statistics from N-body simulations at that point. The key to the Gaussian process method is the way that this constraint influences the area of parameter space around it, through the choice of a particular kernel function. For example, using an exponential kernel the influence of the measurement

on surrounding points in parameter space decays exponentially with the distance in that parameter. The decay rate is specified by a scale length in each parameter, which become hyperparameters to be trained. As more data is added the mean and variance become well constrained throughout the parameter-space. The choice of kernel functions is the key decision in constructing the emulator, and the machine learning process is tuning the hyperparameters of the kernel functions to achieve the best predictions when compared to test data.

The major limitation of an emulator-based model is that the emulator is only able to make accurate predictions for the region of parameter-space that is well sampled by its training data. For a Gaussian process emulator, outside of this region the predictions will regress back to the mean value, rendering them unreliable for cosmological inference. This limitation can be considered in terms of the parameter ranges of the simulations used to construct the emulator. The range must be large enough to encompass the entire region of parameter-space being investigated, while densely sampled enough to produce predictions with an uncertainty that is subdominant to the measurement uncertainty of the data.

Several emulators have been constructed to predict the small-scale clustering of galaxies to use in RSD analyses [278, 275, 161, 68]. The emulator of [278] is particularly important for this work, because it was constructed to match the clustering of a Luminous Red Galaxy (LRG) sample, and is the origin of the emulators used in the analyses presented in Ch. 3 and Ch. 4. It is a Gaussian process-based emulator that makes predictions for ξ_0, ξ_2, and w_p, each in 9 logarithmically spaced separation bins between $0.1 - 60\,h^{-1}$Mpc. Each of the 27 measurement bins is emulated independently using a sum of four kernels: a constant kernel, an exponential squared kernel multiplied by a constant kernel, a Matern-3/2 kernel, and a kernel that controls the logarithm of the white noise added to the diagonal of the covariance matrix during training. For more information see [278] and the `george`[2] Python package that it is constructed from.

1.6 Structure of this thesis

This chapter has provided an introduction to the key concepts applied in later chapters. I have introduced the ΛCDM cosmological model, as well as the key observables that we use to constrain its parameters and test for extensions to the model. In particular, I have presented the details of the measurement statistics and models used to constrain RSD from

[2] https://george.readthedocs.io/en/latest/

galaxy clustering surveys, which provide a unique and powerful test of the cosmological model from the growth of structure.

In Ch. 2 I present the final data release of the Extended Baryon Oscillation Spectroscopic Survey (eBOSS), which at the time of writing is the highest redshift large galaxy spectroscopic survey and the most recent to release data. I describe eBOSS from the selection of the different clustering samples through to the final cosmological measurements. I also present the corrections of two observational systematics: using Pairwise-Inverse-Probability (PIP) weights with Angular Upweighting (ANG) to correct fibre collisions and a weight-based method to correct density fluctuations in the Emission Line Galaxy (ELG) sample caused by variable depth in the targeting survey. Careful correction of these systematics is essential to optimally extract information from the data, and ensure the final results are robust.

I continue the analysis of eBOSS data in Ch. 3 by presenting my small-scale RSD analysis of the eBOSS Luminous Red Galaxy (LRG) sample. This analysis leveraged the unbiased small-scale clustering provided by the PIP+ANG weights and a simulation-based emulator for the clustering of LRGs to measure RSD from clustering in the separation range $0.1 - 60\,h^{-1}$Mpc, finding a factor of 1.7 improvement in precision over the large-scale analysis of the sample. This analysis also investigated several potential sources of systematic bias, and highlights the question of what information from non-linear scales can be correctly interpreted as constraining the linear growth rate. Separating the contributions of linear and non-linear information had not been addressed in contemporary small-scale analyses, and presents a major source of potential systematic bias. Mitigating this potential systematic required limiting the scales included in the measurement of $f\sigma_8$ to $7 - 60\,h^{-1}$Mpc, excluding a significant portion of the emulator range.

I address the issue of isolating the linear signal within small-scale analyses in Ch. 4, which presents a re-analysis of the small-scale eBOSS LRG clustering with an improved emulator that distinguishes between linear and non-linear information. The new emulator permits all scales to be included in the constraint on $f\sigma_8$ while reducing tension between different measurement scales. This chapter provides a detailed description of the new emulator and validation tests, as well as the results of the analysis of the eBOSS LRG clustering.

Finally, in Ch. 5 I present the conclusions drawn from the small- and large-scale analyses of the eBOSS data. I discuss the implications for the cosmological model, areas for future work, and the outlook for upcoming surveys.

Chapter 2

The Extended Baryon Oscillation Spectroscopic Survey

This chapter covers the content of the main Extended Baryon Oscillation Spectroscopic Survey (eBOSS) Data Release (DR)16 papers. I was a member of the eBOSS collaboration during this data release, and wam an author on many of the papers due to my contributions to generating Pairwise-Inverse-Probability (PIP) weights and resolving an observational systematic related to variable photometric depth in the Emission Line Galaxy (ELG) sample. Here I outline the structure of this chapter, the source of the content, and my contributions. Sec. 2.1 provides an introduction to eBOSS, and Sec. 2.2 provides a brief summary of the eBOSS survey design and target selection that gives context for later sections. Sec. 2.3 presents key details of the final large-scale structure catalogue papers [231, 223, 179]. My most significant contribution to those papers is to the ELG catalogue paper [223], and is covered later in the chapter. The mock papers [280, 178] and mock challenges [11, 234, 248] are introduced in Sec. 2.4. Sec. 2.5 presents the contents of the PIP correction paper [190], for which I am included in the first grouping of authors with major contributions due to my work using rotation as a means of removing zero-probability pairs, presented in Appendix A. I lead that work, writing and validating the necessary scripts, and consulted on the remaining work generating PIP weights as well as the writing of the paper. Sec. 2.6 covers a particular observational systematic of the ELG sample that I investigated using a weight-based correction. I led the work for this weight-based correction and presented my results to the eBOSS Collaboration, where they were used to inform the final correction applied in the analysis of the ELG sample. The majority of this section is my original work, with final treatment of the issue in the data analysis adapted from [223, 76, 258]. Sec. 2.7 contains the main cosmological results of DR16, summarized

in [10], and based on the analyses of the Luminous Red Galaxy (LRG) [23, 111], ELG [223, 258, 76], Quasi-Stellar Object (QSO) [141, 196], and Lyα forest [85] samples. In all cases the content has been reformatted and rewritten by me for this thesis. Throughout this chapter I use 'we' to refer to the work of the eBOSS collaboration that I was involved with but did not lead, and 'I' to highlight my primary contributions. All non-original plots contain a citation at the end of their captions detailing the source.

2.1 Introduction

In the era of precision cosmology, novel cosmological measurements are motivated by two factors: providing more precise measurements of key cosmological parameters, and presenting complimentary constraints that build a consensus model from independent cosmological probes. Spectroscopic galaxy surveys fulfill both of these roles. Modern surveys use multi-object spectrographs [247, 80, 168] to observe large samples of galaxies, improving the precision of measurements from previous generations. Those improved measurements are used to extract the Baryon Acoustic Oscillation (BAO) and Redshift Space Distortion (RSD) signals from clustering to constrain the expansion and growth history of the universe, and contain complementary information to local universe and Cosmic Microwave Background (CMB) observations. Successive generations of surveys have also pushed to higher redshifts [73, 208], constraining the evolution of cosmological parameters and providing insight into potential deviations from our standard cosmological model.

The Sloan Digital Sky Survey (SDSS) [41], which has operated the dedicated Sloan Foundation 2.5-meter telescope [117] at the Apache Point Observatory in south-eastern New Mexico since the year 2000, currently provides the largest publicly available catalogue of spectroscopically observed galaxies. The SDSS has been a major driver in our understanding of the Universe from its inception. The most stringent constraints come from the Baryon Oscillation Spectroscopic Survey (BOSS) [73], which observed over 1.5 million galaxies in the redshift range $0.15 < z < 0.7$ between 2009-2014. The final measurements of the BAO scale in BOSS clustering data provided 1% precision standard ruler measurements of cosmological distances up to redshift $z = 0.75$ [9], and a $\sim 6\%$ constraint on the parameter combination $f\sigma_8$ [33, 114, 257, 241] by measuring RSD.

The Extended Baryon Oscillation Spectroscopic Survey (eBOSS) [72] is part of stage IV of the SDSS, and was conceived as a successor to BOSS. eBOSS made use of the BOSS spectrographs [247], capable of simultaneously observing 1000 objects, to observe higher redshift Luminous Red Galaxy (LRG), Emission Line Galaxy (ELG), and Quasi-Stellar Object (QSO) samples for clustering, and greatly increased the sample of $z > 2.1$ quasars

observed for Lyα forest measurements in BOSS. This additional data not only provides tighter constraints on the cosmological model and extends our observations into a new redshift regime, but also investigates new tracers and techniques that will be essential for upcoming surveys [79, 168] to push our understanding even further.

The final eBOSS Data Release (DR), DR16, was accompanied by a coordinated release of papers measuring BAO and RSD in the clustering of the LRG [23, 111], ELG [223, 258, 76], quasar [141, 196], and Lyα forest [85] samples, making use of both configuration space and Fourier space measurements. Accompanying these analyses were papers detailing the construction of the large-scale structure catalogues [231, 179, 223], mock catalogues for testing systematics and estimating the covariance of the measurements [280, 178], and mock challenges to validate the analysis pipelines and assess systematic uncertainties [11, 234, 248]. After release, the final cosmological constraints were summarized and combined with previous SDSS measurements to determine the cosmological implications from two decades of observations [10].

This chapter covers the entirety of the eBOSS DR16 results, from the original survey design to the final cosmological results. In Sec. 2.2 I describe the survey plan of eBOSS and the samples targeted for observation. Sec. 2.3 presents the data reduction and final statistics of the large-scale structure catalogues used for cosmological analysis. In Sec. 2.4 I present the mock catalogues used to determine the covariance of the eBOSS clustering measurements, as well as the mock challenges performed to validate the analysis pipelines and assess the systematic uncertainty. I delve more deeply into two particular observational systematics in Sec. 2.5 and Sec. 2.6. Sec. 2.5 describes the generation, testing, and application of Pairwise-Inverse-Probability (PIP) weights and Angular Upweighting (ANG) to correct the issue of fibre collisions in the spectroscopic sample, providing the first unbiased correction to this important observational systematic for fibre-fed spectroscopic surveys. In Sec. 2.6 I discuss an observational systematic in the eBOSS ELG sample caused by inconsistent photometric calibration, and the weight-based correction I implement to remove the bias from the clustering. After addressing these systematic errors, I present the combined cosmological results from all three eBOSS samples in Sec. 2.7. The eBOSS analyses provide precise measurements of the parameters of the ΛCDM cosmological model, while placing stringent constraints on extensions to that model.

2.2 Survey design

Effective survey design is integral to making efficient observations and obtaining reliable results. In this section I provide a brief overview of the key design choices of eBOSS. I

begin by introducing the galaxy samples targeted for observation in Sec. 2.2.1, including two samples not previously observed in the SDSS. Then in Sec. 2.2.2 I describe the target selection algorithms used to construct the parent catalogues for each sample to be observed. Finally, in Sec. 2.2.3 I present the overall goals of eBOSS, and the observing strategy developed to meet them.

2.2.1 The eBOSS galaxy samples

eBOSS observations consist of three main galaxy samples: Luminous Red Galaxy (LRG), Emission Line Galaxy (ELG), and Quasi-Stellar Object (QSO), the last also commonly known as quasars. LRGs are a class of galaxies that are among the most massive, most luminous, and reddest galaxies, and are commonly found in clusters and groups. Because of their bright intrinsic luminosity and strong clustering they make an effective tracer of the large-scale structure, and have been used in spectroscopic galaxy surveys since the first detection of BAO [94] in the first SDSS spectroscopic sample [90]. LRGs have been used throughout the SDSS, and formed the main galaxy sample for BOSS, so they are well understood both as a tracer of the large-scale structure and as a target for observation.

BOSS also investigated a new probe of the large-scale structure of the Universe by spectroscopically observing quasars at redshifts $z > 2.1$ to measure the Lyα absorption of intervening neutral hydrogen. The frequency of these absorption lines is shifted according to the redshift of the gas, so this forest of absorption lines creates a map of the intervening neutral hydrogen clouds. As with other tracers of the matter density field, these clouds are imprinted with the BAO signal, so these Lyα forest quasars provide an independent BAO measurement compared to typical clustering samples [184, 50, 246, 102]. In fact, in BOSS the marginal value of each Lyα quasar in providing a BAO constraint exceeded that of the LRG sample, so Lyα quasars are designated as a high-priority target in eBOSS [72].

Following the example of BOSS, eBOSS also investigates two new galaxy samples. The first is a sample of quasars at redshifts lower than the Lyα forest sample, $0.9 < z < 2.2$, that are used as direct tracers of the underlying matter density. Quasars are ideal high-redshift tracers in that they are highly biased and intrinsically very bright, permitting them to be observed beyond the redshift range of other tracers [67].

The final galaxy sample observed in eBOSS is ELGs, which are galaxies containing active star-formation that provide strong emission lines. These strong emission lines allow for effective redshift determinations, and ELGs are much more common than LRGs at the redshifts targeted by eBOSS [65]. The large numbers of ELGs permits the creation of a high-density ELG catalogue that is observed separately from the LRG and QSO samples.

The ELG sample is also complementary to the LRG sample because they cover very similar redshift ranges, but are targeted using different selection algorithms and photometric surveys (see Sec. 2.2.2), observed independently, and use different redshift determination algorithms. Thus cross-correlation between the samples is a powerful tool for mitigating systematic effects in one sample. The eBOSS ELG sample is also an important test for the next generation of surveys, including the Dark Energy Spectroscopic Instrument (DESI) [79, 80] and *Euclid* [168], that plan to use ELGs as part of their core clustering samples. eBOSS provides a road map for these surveys, as well as an opportunity to address any difficulties in the observation or analysis of a large ELG sample.

2.2.2 Target selection algorithms

While the BAO signal is relatively insensitive to the homogeneity of the sample and large-scale observational systematics, because it depends on the observation of a relatively sharp feature that does not depend on tracer bias, this is not true for all of the cosmological measurements that use eBOSS data [72], such as the RSD measurements. For this reason it is important when constructing the eBOSS targeting catalogues to build samples that are homogeneous on large scales, including selecting the same population of galaxies throughout the survey window. It is also important to construct the target selection algorithm to ensure that targets will lie in the correct redshift range to prevent wasted observations.

The eBOSS LRG target sample is selected [219] from SDSS DR13 photometry [12]. SDSS photometry was used to target the LRG sample for BOSS, and provides sufficient depth and sky coverage in the *ugriz* bands for the eBOSS LRG sample. Compared to previous data releases, the SDSS DR13 photometry used for eBOSS uses the "ubercalibration" method [199] applied to imaging from the *Pan-STARRS* survey [154], providing an improved photometric calibration [242]. Typically, LRGs are targeted using the 4000 Å break caused by metals in the atmospheres of old late-type stars' atmospheres, however at the higher redshifts targeted by eBOSS this break shifts into the near infrared, and the colours of LRGs overlap more with M-type stars. These difficulties motivate including infrared observations from the Wide-field Infrared Survey Explorer (WISE) satellite [165], which observed the full sky in four infrared bands. The key advantage of including infrared observations is the 1.6 micron 'bump' in the spectral energy distribution of objects with old stellar populations, such as LRGs, caused by a minimum in the opacity of the H^- ion. At $z \sim 1$ this bump is shifted to around 3.4 microns, the centre of the first band of WISE observations, denoted W1. This band was therefore used in the LRG target selection to distinguish between LRGs and stars.

The target selection criteria applied for the eBOSS LRG sample is as follows. First, a series of flux limits (detailed in [219]) are applied to the sample to ensure good detections in all the necessary bands. Then three colour selections ware applied to the remaining objects to isolate LRGs in the desired redshift range:

$$\begin{aligned} r - i &> 0.98 \ , \\ r - W1 &> 2.0 \times (r - i) \ , \\ i - z &> 0.625 \ . \end{aligned} \tag{2.1}$$

The first and third conditions serve to separate high redshift LRGs from low redshift LRGs. The second condition separates galaxies from stars, as shown in Fig. 2.1[1].

Like the LRG sample, the ELG sample is also selected using colour cuts from photometric catalogues [221]. However, because ELGs tend to be lower mass and less luminous than LRGs, rather than using the SDSS photometry, the ELG target selection uses data from the deeper Dark Energy Camera Legacy Survey (DECaLS) [81] DR3 photometry. DECaLS was a grz imaging survey using the DECam camera [101] mounted on the Victor M. Blanco 4m telescope on Cerro Tololo, near La Serena, Chile. The primary purpose of DECaLS is to provide photometry to be used in targeting for DESI. As well as dedicated DECaLS observations, the program also incorporates all publicly available DECam observations taken over the DESI footprint, which primarily includes data from the Dark Energy Survey (DES) [2]. However, these DES observations only cover one of the two main observation areas of eBOSS, the South Galactic Cap (SGC). The North Galactic Cap (NGC) instead uses the slightly shallower DECaLS imaging.

The first step in the ELG target selection is to apply a series of cuts to ensure clean photometry, specified in [221]. For the SGC the following three selection cuts are then applied:

$$\begin{aligned} 21.825 < g &< 22.825 \, , \\ -0.068 \times (r - z) + 0.457 < g - r &< 0.112 \times (r - z) + 0.773 \, , \\ -0.218 \times (g - r) + 0.571 < r - z &< 0.555 \times (g - r) + 1.901 \, , \end{aligned} \tag{2.2}$$

while a slightly modified set of selection cuts are applied in the NGC:

$$\begin{aligned} 21.825 < g &< 22.9 \, , \\ -0.068 \times (r - z) + 0.457 < g - r &< 0.112 \times (r - z) + 0.773 \, , \\ -0.637 \times (g - r) + 0.399 < r - z &< 0.555 \times (g - r) + 1.901 \, . \end{aligned} \tag{2.3}$$

[1] http://www.cadc-ccda.hia-iha.nrc-cnrc.gc.ca/en/megapipe/docs/filt.html

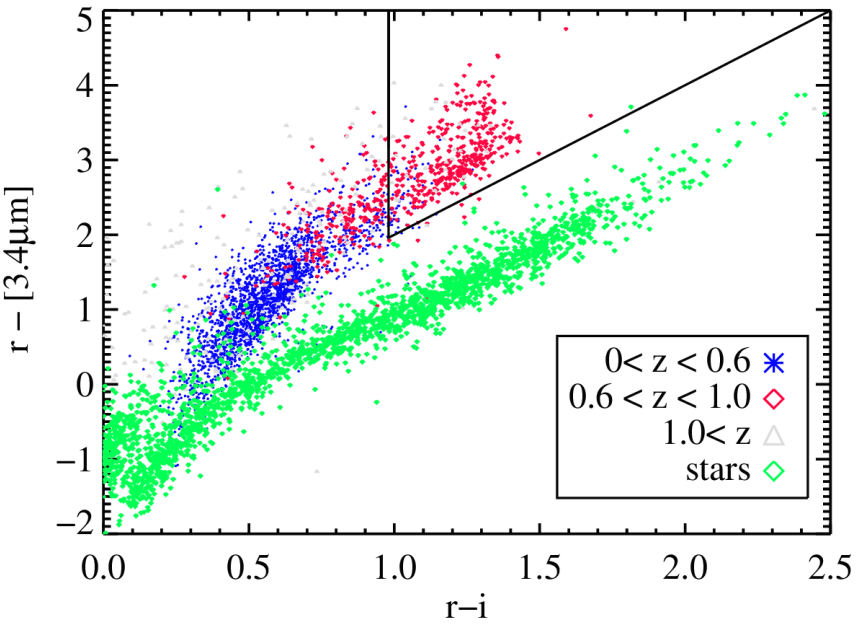

Figure 2.1: Colour-colour plot of galaxies observed by WISE and Canada-France-Hawaii Telescope Legacy Survey (CFHT LS). Photometric redshifts are taken from the COSMOS photo-z catalogue [147]. Blue stars show galaxies with photometric redshifts of $z < 0.6$, red diamonds show galaxies between $0.6 < z < 1.0$, and cyan triangles show galaxies at $z > 1.0$. Stars are shown as green diamonds. The triangular area shows the relevant eBOSS LRG selection cuts for these colours. Optical photometry is taken from the catalog of [120], transformed to SDSS passbands using the conversion relation found at the CFHT LS webpage (footnote in text). [219]

The g-band magnitude cut is used to select [OII] emitters. The reason it is extended for the NGC is because a coding error in the target selection script (see [221]) effectively introduced an additional angular mask in the NGC that caused $\sim$ 15% objects to be improperly rejected. The additional depth in the g-band raises the overall density of targets to match the survey goals. The latter two conditions for each cap define a quadrilateral in colour-space used to exclude stars and lower-redshift galaxies (see Fig. 2.2). The SGC selection window is extended in the $r - z$ colour direction because of the deeper DES imaging in the SGC. The deeper imaging reduces scatter, which permits an extended selection while keeping contamination from lower-redshift galaxies low.

The quasar selection is more complicated, and relies on two complementary methods. The first is the *XDQSOz* [44] method of classifying quasars in flux-space for a specfied range of redshift. The photometric data used for this classification is the same *ugriz* SDSS DR13 imaging used for the LRG sample. A test spectroscopic survey was performed using CFHT LS W3 data to verify this method, detailed in [7]. A secondary method based on variability in multi-epoch imaging from the Palomar Transient Factory (PTF) [224, 169] is used to increase the density of the sample. PTF is a wide-field photometric survey that repeated imaging of 20,000 deg^2 in the Northern Hemisphere to track optical transients. The details of both of these methods of quasar selection are given in [193].

These three target selection methods for the LRG, ELG, and quasar samples provide pure and homogeneous parent catalogues for the main spectroscopic observations of eBOSS. These requirements are essential for obtaining reliable results from the data analysis, and fulfilling the goals laid out in the survey plan.

2.2.3 Observing strategy

Following from the success of BOSS in achieving a 0.9% precision measurement of the BAO scale at an effective redshift $z = 0.57$ from LRG clustering and a 2% precision measurement at an effective redshift $z = 2.34$ from quasar Lyα forest, the goals set out for eBOSS for the LRG, quasar, and ELG clustering samples are to achieve 1%, 2%, and 2% precision measurements of the BAO scale, respectively, and to observe as many $z > 2.1$ quasars for Lyα forest measurements as possible because of their larger marginal value [72]. These goals, together with the roughly 5400 hours allocated for the eBOSS observing program, set the survey plan for eBOSS.

The typical BOSS spectroscopic plate was observed in five 15 minute exposures, with an additional 20 minutes of overhead per plate [42]. A pilot study using dedicated plates in 2013 and 2014 showed that the same observation strategy is sufficient to obtain redshifts

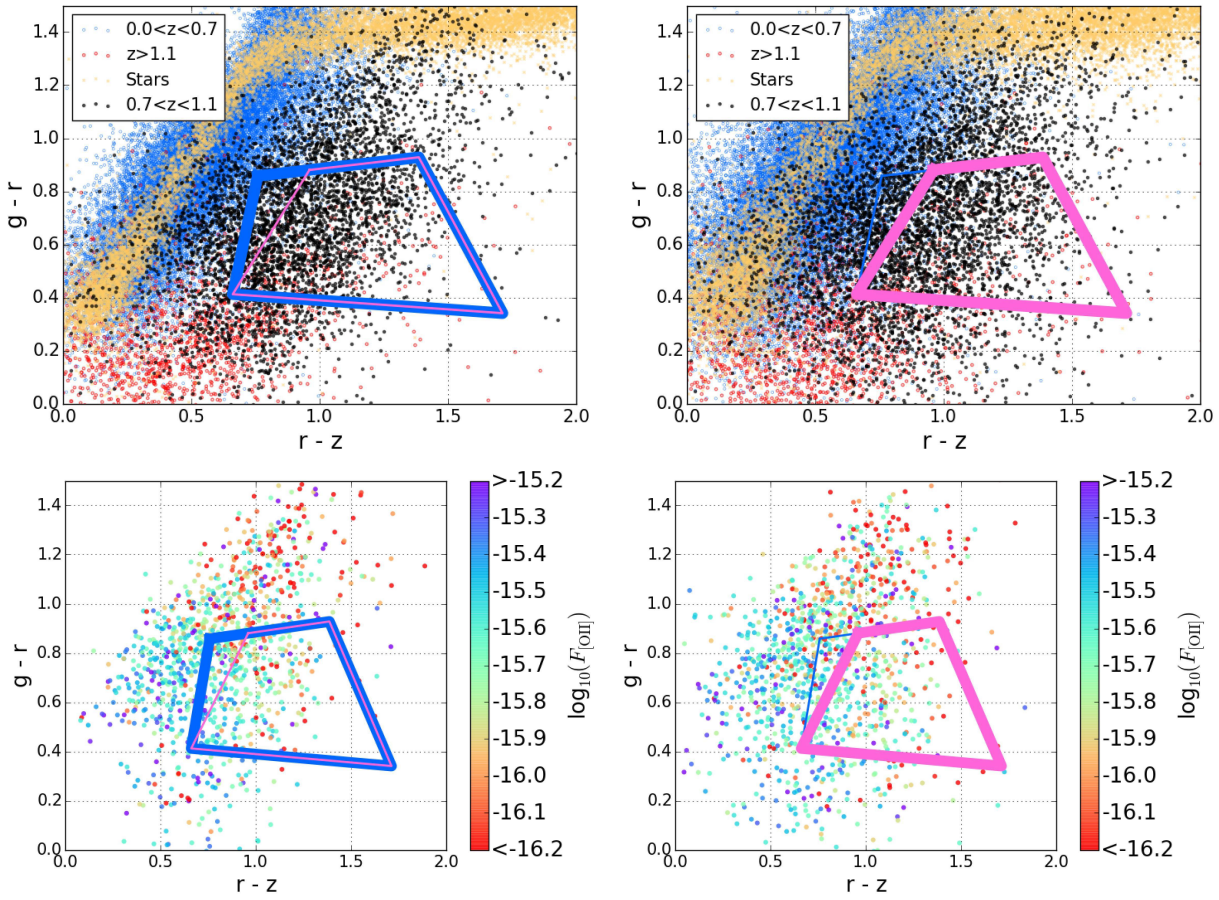

Figure 2.2: eBOSS ELG target selection in the grz-colour space with objects from the CFHT LS W4 field [120]. The SGC target selection is shown as the blue quadrilateral, and the NGC target selection is shown as the pink quadrilateral. All objects are observed using photometry from DECaLS DR3 over the CFHT LS W4 footprint, degraded down to DECaLS DR3 depths. The left column shows these objects over the SGC footprint, while the right column shows these objects over the NGC footprint. Photometric redshifts are determined from the CFHT LS survey. Stars are shown as beige crosses, $0 < z_{\rm phot} < 0.7$ objects as blue circles, $0.7 < z_{\rm phot} < 1.1$ objects as black circles, and $1.1 < z_{\rm phot}$ objects as red circles. Bottom Panels: objects in common with the VIPERS survey with $0.6 < z_{\rm spec} < 1.1$. The colour bar shows the [OII] fluxes measured from the VIPERS survey spectra. [221]

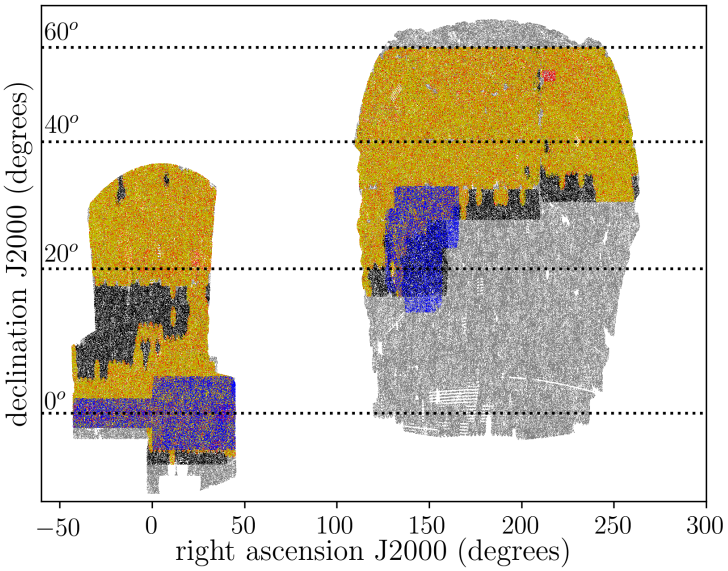

Figure 2.3: The footprint of all eBOSS targets before veto masks. Yellow points show quasars that are observed by eBOSS. Red points show observed LRGs, which overlap almost perfectly with the yellow points. Blue points are a random 20% subsample of observed ELGs. Black points show quasar and LRG targets that were tiled but did not ultimately receive observations, while grey points are BOSS CMASS galaxies. [231]

for the eBOSS samples as well [72]. As a result, it was planned to observe 1800 plates in the 5400 hours allocated to eBOSS, assuming the same 50% weather efficiency found during the SDSS-III observations. eBOSS is operated using a similar survey plan as BOSS, with overlapping pointings providing a mixture of single- and multi-pass observation areas, with up to a maximum of three observations of a given patch. This survey strategy gives varying completeness and allows for the observation of close pairs in overlap regions that could not be observed in a single pass due to fibre collisions, while also covering a large area of the sky. Each plate covers an area of 7 deg^2, so the field centres were aligned uniformly over the survey area with an average of 5 deg^2 per field centre to ensure overlap between neighbouring plates. The field centres were then perturbed to maximize the number of observed redshifts by correlating overlap regions with high target density areas. The final area planned for the survey is 9000 deg^2, shown in Fig. 2.3. Due to the large foregrounds, observations are restricted to be away from the galactic plane, forming two distinct regions labelled the North Galactic Cap (NGC) and South Galactic Cap (SGC).

The total eBOSS footprint is divided into smaller regions on the sky, called 'chunks', and tiled separately for observation. 300 plates are allocated almost entirely to the ELG sample because of the high density of targets, while the remaining 1500 plates, covering 7500 $\deg^2$, are used to simultaneously observe the LRG and quasar samples, as well as ancillary targets for the Time-Domain Spectroscopic Survey (TDSS) [191] and SPectroscopic IDentification of ERosita Sources (SPIDERS) [64] sub-programs. Once target catalogues were selected from the photometric surveys and telescope pointings were determined, a target selection algorithm is run to maximize the number of spectra obtained. An important observational limit of eBOSS is that two objects within $62''$ on the sky can not be simultaneously observed because of the finite size of the fibre casings (see Sec. 2.5.1 for more details). Objects within this radius are referred to as "collided". For the quasar and LRG chunks the priority in resolving collisions between targets, from highest to lowest, is: SPIDERS, TDSS, quasars, white dwarfs that can be used as calibration standards, and finally LRGs. The tiling algorithm is designed to ensure 100% completeness in the decollided set of non-LRG objects. The LRG target selection is tuned to ensure a sufficient density of possible LRG targets that there is an object available for all fibres, meaning the observed LRG sample is not fully complete by construction. For ELG chunks there is only a single sample, so all targets are assigned with an equal probability.

Once targeting was established a circular metal plate is drilled for each pointing, with fibres positioned in the holes to observe the targets of that plate. Fibres are hand plugged into each plate, and connected to the BOSS spectrographs for observation. The output of those observations and the resulting large-scale structure catalogues are described in Sec. 2.3.

2.3 Large-scale structure catalogues

Following from the survey design and observations described in Sec. 2.2, in this section I present the data reduction applied to the raw observations and the resultant large-scale structure catalogues. I begin by describing the process used to extract redshifts from spectra for the various galaxy samples (Sec. 2.3.1). Then, in Sec. 2.3.2 I explain how the data is corrected for observational effects and systematics, and present the final statistics of the large-scale structure catalogues used in the cosmological analyses.

2.3.1 Redshift estimation

Measuring redshifts is the key observation that allows galaxy surveys to transform a 2D distribution of objects on the sky to a 3D map of the large-scale structure of the Universe. Initially, the same redshift algorithms used for BOSS were planned to be applied in eBOSS (see [72] for discussion). However, once observations had begun it became apparent that these algorithms are insufficient for the fainter, higher redshift eBOSS LRG sample. Instead, redshifts for the eBOSS LRG sample are extracted from the spectra using the REDROCK algorithm[2]. REDROCK fits the data with a linear combination of redshift-varying spectral templates characterizing the spectral diversity of stars, galaxies, and quasars. The redshift and spectral class that results in the lowest χ^2 value is determined to be the best description of the spectra. Fits are only deemed to be reliable if the redshift and spectral class that produced the second best fit is separated by a sufficiently large difference in χ^2, specified by the parameter $\delta\chi^2$. The value of this difference was tuned using a sample of multi-epoch spectra, consisting of objects that were purposefully targeted by multiple plates. The final value of $\delta\chi^2$ was chosen to maximize the number of good redshifts without exceeding purity limits [231]. After making an initial match, the REDROCK algorithm applies a second step where a subset of the original spectral templates are used as archetype models, and a secondary match is performed, superseding the earlier match. The reason for performing this second match is to exclude non-physical combinations of the spectral templates that can lead to incorrect redshift detections. The galaxy spectral templates used for eBOSS are derived from a Principal Component Analysis (PCA) of 20 000 theoretical galaxy spectra spanning stellar age, metallicity, and star formation. The stellar templates are derived from a separate PCA analysis of 30 000 theoretical stellar spectral spanning a range of stellar types [231].

After applying REDROCK to the LRG spectra a total of 88% of observations result in a good LRG redshift, with 9% of original spectra ultimately attributed to stellar contaminants, and 1% to quasar contaminants. The final rate of catastrophic failures, where a redshift that deviates from the true value by more than 1000 km/s is confidently assigned to the spectra, was estimated to be less than 1% from multi-epoch observations. REDROCK is also applied to the ELG sample, with additional reliability criteria applied to ensure the redshift estimates match the purity conditions. These additional criteria, designed to ensure a sufficient signal-to-noise ratio of the [OII] line and enough continuum and emission line information from the spectra, are detailed in [223]. Reliable redshifts are obtained for 91% of ELG observations, with a catastrophic failure rate less than 1%.

The quasar redshift estimation follows a multi-step process. First an automated decision-

[2] Available at github.com/desihub/redrock

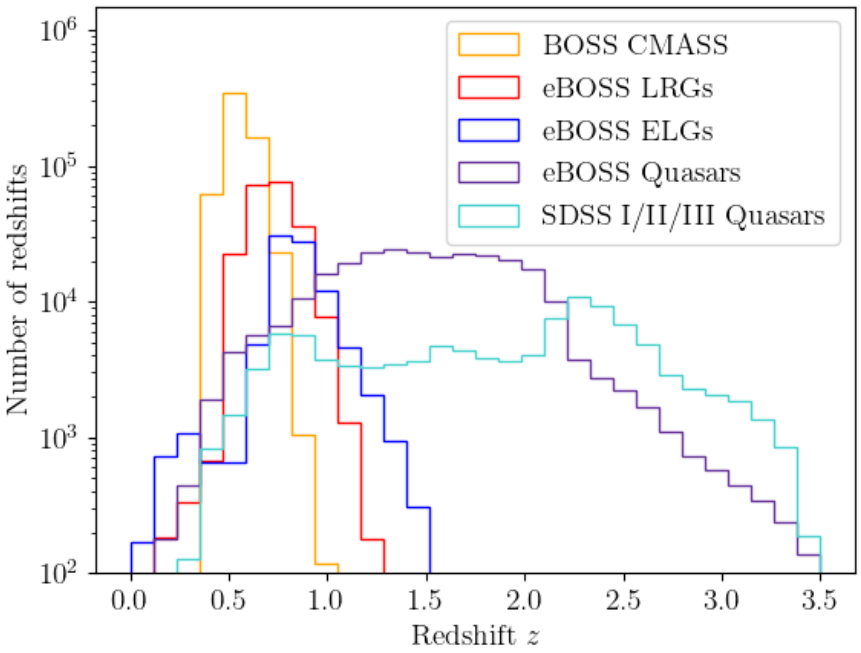

Figure 2.4: Histograms of the redshift distributions of the eBOSS DR16 large-scale structure catalogues. BOSS CMASS LRGs and quasars from previous generations of SDSS are also shown because they were jointly analysed with the relevant eBOSS samples. [231]

tree is applied to the data to ensure proper classification as a quasar. Then, the REDVSBLUE[3] algorithm, described in [179], is used to estimate the redshift. In total, 86% of quasars are classified as receiving a good redshift.

The redshift estimation of all three clustering samples results in a sufficiently high success rate to match the survey design. The final redshift distributions of all the eBOSS samples are shown in Fig. 2.4. This figure shows the large number of objects and wide redshift range observed by eBOSS, which are key to providing the wealth of cosmological results presented in Sec. 2.7

[3]https://github.com/londumas/redvsblue

2.3.2 Correcting for observational systematics

The eBOSS observations resulted in large spectroscopic samples over a wide range in redshift, ideal for providing complementary constraints to earlier galaxy surveys. However, in order to produce robust results the data must be carefully analyzed and corrected for any signals imprinted by the observational methods, particularly systematic variations in density that are non-cosmological in origin.

One such issue is caused by redshift failures, which are not completely random across the survey. The origin of this issue is that the signal-to-noise ratio of spectra is not uniform across the spectrograph Charge-Coupled Device (CCD), and the signal-to-noise is naturally correlated with the rate of redshift failure. Because individual fibres are consistently connected to the same portion of the CCD, and are also consistently connected to the same regions of the spectroscopic plate, the correlation of redshift failure with CCD position is imprinted in the angular distribution of the clustering sample. These additional density perturbations effectively introduce angular clustering modes that are not cosmological in origin, potentially biasing the cosmological inference. The variations are corrected by calculating the probability of obtaining a good redshift for each fibre, and weighting each observation by the inverse of this probability for the fibre used to observe that object, statistically accounting for the missing objects. These weights are stored in the large-scale structure catalogues as $w_{\rm noz}$.

Redshift failures are not the only observational effect that can impact the angular clustering of the sample. A variety of factors cause variations in the photometric data used for targeting, which carry forward to cause systematic variations in the density of the final spectroscopic observations. These variations are corrected using a multi-variate linear regression to fit the trend in density with each factor, and then inversely weighting by the fitted trends to remove them from the final analysis. These weights, $w_{\rm sys}$, are also stored in the final catalogues.

To determine how each factor impacts the density in angular patches the sample is divided into equal area pixels on the sky using Hierarchical Equal Area isoLatitude Pixelation (HEALPix)[4], with $N_{\rm side} = 512$. The density of target samples is then estimated by taking the ratio of the number of targets to the number of randoms in each pixel, normalized by the average ratio of targets to randoms across the survey.

For the eBOSS LRGs the factors included in the multi-linear regression are the stellar density ($N_{\rm star}/{\rm deg}^2$), the Galactic extinction (E[B-V]), the sky background in the i-band, and the i-band seeing maps of the SDSS photometry. The reason for only including the

[4]https://healpix.sourceforge.io/

i-band maps in the regression is because the SDSS photometry was observed using a drift scanning technique where all five bands were observed almost simultaneously, meaning the backgrounds and seeing are highly correlated between the bands. The regressions are performed separately for the NGC and SGC. The raw and corrected trends in target density with each of these factors, as well as the airmass and g-band depth, are shown in Fig. 2.5. The regression is able to remove strong trends in the galactic extinction and stellar density, and lead to an improvement in all of the parameters considered [231].

The same process is applied to the quasar sample, but the stellar density is exchanged for the g-band depth in the regression because the raw quasar sample display a strong trend with g-band depth, but is only weakly correlated with stellar density. The quasar correction achieves a similar result to the LRG correction, removing significant density trends in the corrected sample, and is described in [231].

A similar method is used for the ELG sample, but includes more potential sources of systematics in the regression. The ELG correction regresses against DECaLS imaging depth and seeing in the grz-bands, the stellar density, Galactic extinction, dust temperature, and HI column density. The full results of this regression and equivalent plots to Fig. 2.5 are shown in [223].

As well as the weights to remove systematic variations in angular density, a set of weights are also applied that optimize the signal for the BAO analysis by reducing variations in the redshift distribution of the data, $n(z)$. These weights, denoted as $w_{\rm FKP}$, are calculated using the method of [97]:

$$w_{\rm FKP} = 1/[1 + n(z)P_0]\,, \tag{2.4}$$

where P_0 is an estimate for the magnitude of the power spectrum. The values used for each sample are:

$$P_0 = \begin{cases} 6000\,({\rm Mpc}/h)^3 & \text{quasars} \\ 10000\,({\rm Mpc}/h)^3 & \text{LRGs} \\ 4000\,({\rm Mpc}/h)^3 & \text{ELGs} \end{cases} \tag{2.5}$$

In each case P_0 matches the amplitude of the power spectrum for that tracer at $k \sim 0.15h{\rm Mpc}^{-1}$, which is the optimal choice for the BAO analysis [103].

A final set of weights is applied to the data to correct for systematic variations in the data due to fibre collisions. The fibre-collision issue and the weights used to correct it are discussed in detail in Sec. 2.5.

The statistics of the large-scale structure catalogues for the LRG, quasar, and ELG clustering samples are shown in Table 2.1. The number of objects in each sample are

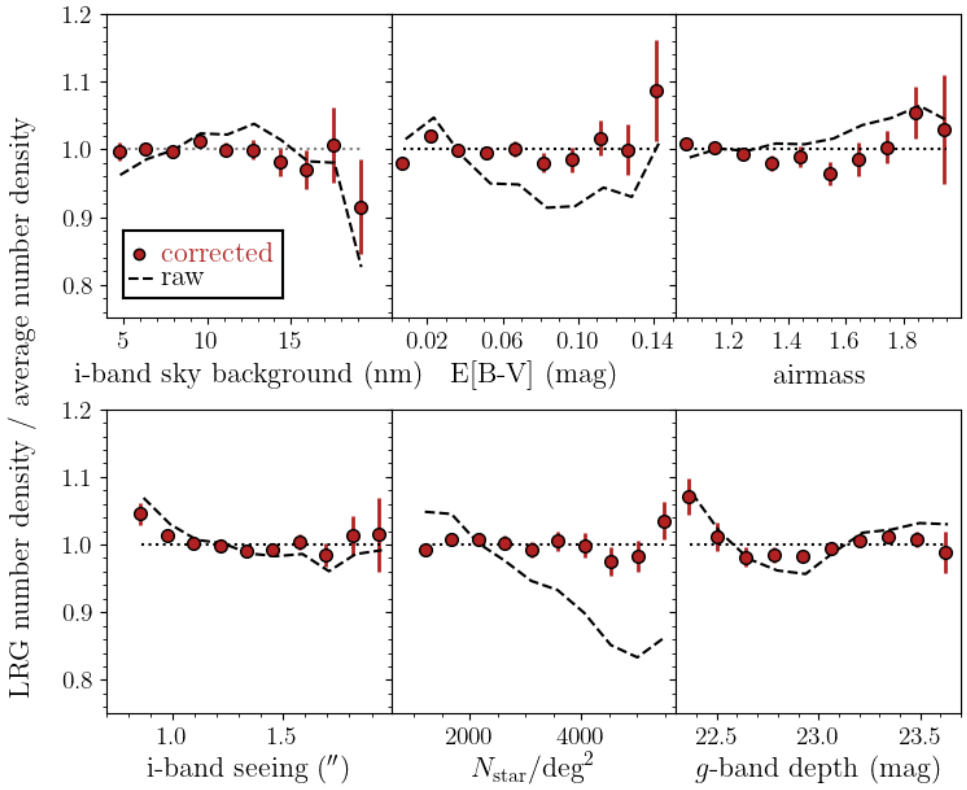

Figure 2.5: Trends in angular LRG target density with various observational systematics. The dashed black lines show the raw trends before any correction, and the red points show the results after the data has been weighted to remove the trends using a multi-linear regression. The regression is applied separately for the NGC and SGC, and then the two regions are combined here. *Top row from left to right:* density of targets relative to the random catalogue plotted against sky background in the i-band, galactic extinction, and airmass. *Bottom row from left to right:* density plotted against i-band seeing, stellar density, and depth in the g-band of the targeting photometry. [231]

	LRG	Quasars	ELG
$N_{\rm par}$	377,633	703,521	269,178
$N_{\rm targ}$	311,848	655,521	229,694
$N_{\rm cp}$	8,386	11,710	10,602
$N_{\rm gal}$	209,894	28,041	195,085
$N_{\rm QSO}$	2,889	454,452	-
$N_{\rm star}$	20,766	7,085	1,704
$N_{\rm zfail}$	5,224	18,778	20,202
$N_{\rm LSS}$	174,816	343,708	173,736
Tiled Area ($\deg^2$)	6,309	6,309	1170.2
Unvetoed Area ($\deg^2$)	5,223	5,858	733.8
Effective area ($\deg^2$)	4,103	4,699	727.0

Table 2.1: Statistics of the eBOSS LRG, quasar, and ELG clustering samples. $N_{\rm par}$ is the total number of objects selected for the parent catalogue of each sample. $N_{\rm targ}$ is the number of eligible targets after applying veto masks. $N_{\rm cp}$ is the number of objects in close pairs, that are not observed because of fibre collisions. $N_{\rm gal}$, $N_{\rm QSO}$, and $N_{\rm star}$ are the numbers of observed spectra with a successful redshift estimation that are classified as galaxies, quasars, and stars respectively. $N_{\rm zfail}$ is the number of observed spectra for which redshift estimation failed. $N_{\rm LSS}$ is the final number of objects with successful redshifts of the correct type for that sample, after applying completeness and redshift cuts. The redshift ranges selected for the LRG, quasar, and ELG clustering samples are $0.6 < z < 1.0$, $0.8 < z < 2.2$, and $0.6 < z < 1.1$ respectively. 'Tiled Area' is the total area tiled for each sample. 'Unvetoed Area' is the remaining area after veto masks are applied. 'Effective Area' is the area after applying completeness cuts, and multiplying remaining sectors by their completeness. [231, 223]

shown from the initial parent catalogues selected according to the targeting algorithms (see Sec. 2.2.2), through to the final large-scale structure catalogues used in the cosmological analyses (see Sec. 2.7). There is a significant reduction in the LRG and quasar samples, apparent in the large change between $N_{\rm targ}$ and the number of successful redshifts of the correct type for each sample, and in the difference between the unvetoed and effective area of each sample. This difference is due to the removal of sectors with completeness below 50% (see [231] for details) that includes the regions of the LRG and quasar parent catalogues that were tiled but never observed, shown in black in Fig. 2.3. The full details of the LRG and quasar samples are given in [231], and the details of the ELG sample are presented in [223].

2.4 Mocks

Before performing the analysis of the large-scale structure catalogues there remain two very important steps. The first is determining the covariance of the data vector, which is essential to calculating the agreement between model and data. The other is validating the analysis pipeline to ensure robust results. We use synthetic mock galaxy catalogues to perform both these tasks. In Sec. 2.4.1 I describe the main mock catalogues used to determine the covariance matrices for the final analyses. Then in Sec. 2.4.2 I describe the mock challenges performed on each data sample to validate the analysis pipelines and determine the systematic error budget.

2.4.1 Mock catalogue ensembles used for cosmological analysis

When analyzing clustering information it is important to note that clustering measurements in different separation bins are correlated, necessitating the covariance of the data to be taken into account in the analysis. For the eBOSS analyses presented in Sec. 2.7, this covariance is estimated using mock galaxy catalogues that accurately reproduce the expected clustering of the various galaxy samples. It is necessary to invert the covariance matrix during the analysis, so in order to produce a numerically stable result it is desirable to estimate the covariance matrix from a significantly larger number of samples than there are data points in the analysis. This presents a challenge for a large volume survey such as eBOSS, since it is not computationally tractable to produce enough equal volume N-body simulations.

Instead, the covariance of the eBOSS clustering measurements are estimated using a set of 1000 approximate mock catalogues for each sample generated using the Effective Zel'dovich Approximation mock (EZmock) method. The EZmocks are based on a Gaussian random field in a $5\,h^{-3}\mathrm{Gpc}^3$ box with an initial power spectrum and cosmology of a flat ΛCDM universe with parameters $\Omega_m = 0.307115$, $\Omega_b = 0.048206$, $h = 0.6777$, $\sigma_8 = 0.8225$, and $n_s = 0.9611$. Matter particles were then displaced from their initial to final positions using the Zel'dovich approximation. The tracer bias relation is calibrated to match the observed clustering of the target sample in the eBOSS DR16 data. The linear component of the RSD is imprinted by the Zel'dovich approximation, while the non-linear term is modelled through an isotropic Gaussian motion. The mocks are then trimmed to match the geometry and radial selection function of the relevant eBOSS sample catalogues. Observational effects are also introduced to the final EZmock catalogues based on their impact on the actual data, including fibre collisions, redshift failures, and angular photometric systematics. Full details of the EZmock generation and tests are reported in [280].

An additional suite of 2000 approximate simulations are generated to test for systematic errors in the ELG sample using the GaLAxy Mocks (GLAM)-Quick Particle-Mesh (QPM) method. These mocks are generated in four steps. First, 2000 GLAM [157] simulations are run with box size $3000\,h^{-1}$Mpc. Second, the QPM code is used to assign dark matter halos to the GLAM density fields. Third, the halos are populated with an HOD model to replicate the small-scale clustering of the ELG sample on scales around $\sim 10h^{-1}$Mpc. Finally, the mocks are cut to match the ELG geometry. These mocks, as well as validation of their large-scale clustering compared to the ELG sample, are presented in [178].

2.4.2 Mock challenges

In addition to the mock catalogues used to assess the covariance of the sample, mock challenges are also performed for the LRG [234], ELG [11], and QSO [248] samples to validate and assess the experimental uncertainty of the models used for the cosmological analyses. These mock challenges are based around the Outer Rim N-body simulation [132], a large and high resolution gravity-only N-body simulation run using the Hardware/Hybrid/Accelerated Cosmology Code (HACC) [121]. The simulation consists of a periodic box with side length $3\,h^{-1}$Gpc and 10240^3 particles with mass $m_p = 1.85 \times 10^9\,h^{-1}M_\odot$, making it one of the largest and highest resolution N-body simulations at the time that DR16 was being produced. The LRG challenge also used the NSERIES mocks that were previously used in the SDSS DR12 analysis [9], and the ELG challenge includes additional tests using the MultiDark *Planck* simulation [158].

Each mock challenge began by investigating a variety of HOD models in order to create an accurate representation of their respective samples from the Outer Rim mocks. For each sample, multiple catalogues are produced using separate HOD models to reduce the dependence of the results on any particular parameterization. The analysis pipelines are then applied to the HOD catalogues to determine if the BAO and RSD models are able to faithfully recover the known cosmological parameters of the simulation, to identify and assess the magnitude of any systematic errors, and to make adjustments where necessary. In the case of the ELG and QSO mock challenges additional tests are made using blinded mocks, where the cosmological parameters of interest have been shifted from the base values of the simulation and are unknown to the experimenters, in order to ensure that the analysis methods are not tuned to produce an expected result. For all three samples, the mock challenges showed that the analysis pipelines are able to recover the true values to within statistical uncertainty. For full details of the results, as well as systematic error estimations, see the mock challenge paper for each sample (LRG [234], ELG [11], and QSO [248]).

2.5 Correcting fibre collisions using PIP+ANG weighting

An important observational systematic in targeted spectroscopic surveys is the fibre-collision issue, which biases the clustering measured from the observed sample. Fibre collisions have typically been corrected using approximate methods that give accurate large-scale results, but are unable to completely remove the bias on small scales. Pairwise-Inverse-Probability (PIP) [36, 202] weights provide the first theoretically unbiased correction to the fibre-collision issue, and in [190] we apply this method to a fibre-fed spectroscopic survey for the first time.

In this section I begin by describing the fibre-collision issue, and discuss the correction used for large-scale analyses in eBOSS (Sec. 2.5.1). I then provide a theoretical background for the PIP method in Sec. 2.5.2, as well as the Angular Upweighting (ANG) method that we apply in tandem. Then in Sec. 2.5.3 I describe the method we use to generate and apply PIP+ANG weights for all three eBOSS samples. Next, I present the results of tests performed by applying PIP+ANG weights to mock catalogues to validate the technique (Sec. 2.5.4). Lastly, I show the final PIP+ANG corrected clustering measurements for eBOSS and compare them to measurements corrected using an approximate method (Sec. 2.5.5).

2.5.1 Fibre collisions in spectroscopic surveys

Fibre collisions are an observational limitation of fibre-fed spectroscopic surveys, where the physical size of the fibre casing prevents simultaneously targeting two targets within a certain radius on the sky, called the fibre-collision radius. For BOSS and eBOSS, which both use the BOSS spectrographs, the fibre-collision radius is $\theta^{(\mathrm{fc})} = 62''$. In regions of the survey where multiple passes of the instrument overlap it may be possible to observe multiple collided objects, but fibre collisions in single pass regions will always remain unresolved. These missing objects present an issue because they are correlated with higher density regions, introducing a systematic bias in the 3D clustering of the sample. If uncorrected this bias affects all clustering scales, but is most impactful on small scales, where it is the leading source of systematic error [226].

BOSS and eBOSS adopt the same selection method for resolving fibre collisions. First, an algorithm is applied to maximize the number of observed spectra, then any remaining conflicts are resolved randomly. This selection can be illustrated using a simple example of three objects, where object A collides with both B and C, but B and C are not collided.

In a single pass region if A is observed then neither B nor C can be observed, so the selection algorithm will always target both B and C, and leave A unselected, to maximize the number of observed spectra. In a double pass region any pair (A and B, A and C, or B and C) can be observed, but it is not possible to observe all three objects due to the double collision of A, so the observed pair is determined randomly. In a triple pass region each object can be observed on a different plate, so it is possible to observe all objects. An example of fibre collisions in the BOSS CMASS sample is shown in Fig. 2.6. Fibre-collided close pairs are mostly resolved in multi-pass regions, with few exceptions, however all close pairs in single-pass regions are unresolved. eBOSS uses the same instrument and similar tiling, fibre-assignment, and observational strategies as BOSS, so the effect is very similar.

A common correction for fibre collisions is the application of Nearest-Neighbour (NN) weighting, where the weight of missing observations is transferred to its nearest neighbour. This method is able to recover accurate clustering on large scales, but performs worse at scales approaching the fibre-collision radius, and loses all information below the fibre-collision radius. In eBOSS a variation of NN weights is produced where the weight of the missing object is equally distributed to all good observations in the collision group. This method, denoted as Close Pair (CP) weights and given the symbol $w_{\rm cp}$, improves the accuracy of the NN correction, but faces the same issues on small scales (see [231]). CP weights are used in most large-scale eBOSS analyses, detailed in Sec. 2.7, but for the small-scale analyses presented in Ch. 3 and Ch. 4, we require a correction that is unbiased on all scales, so we use the more accurate PIP+ANG correction.

2.5.2 Pairwise-inverse-probability weights and angular upweighting

Unlike NN weights, the PIP weighting proposed in [36] provides a theoretically unbiased correction to the issue of fibre collisions on all scales. The PIP weights are assigned to pairs of objects in the targeted sample and quantify the probability, for any pair, of being selected in a random realisation of the survey targeting. Each pair is then inversely weighted by that observation probability. An intuitive example is if there was a particular configuration of potential targets that is only observed 10% of the time, then each such pair would be weighted by a factor of $w_{\rm PIP} = 1/0.1 = 10$, perfectly correcting for the missing 90%.

In order to correctly apply PIP weights two conditions must be met. The first is that the targeting probability for each pair is known, which is non-trivial. 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. These probabilities are

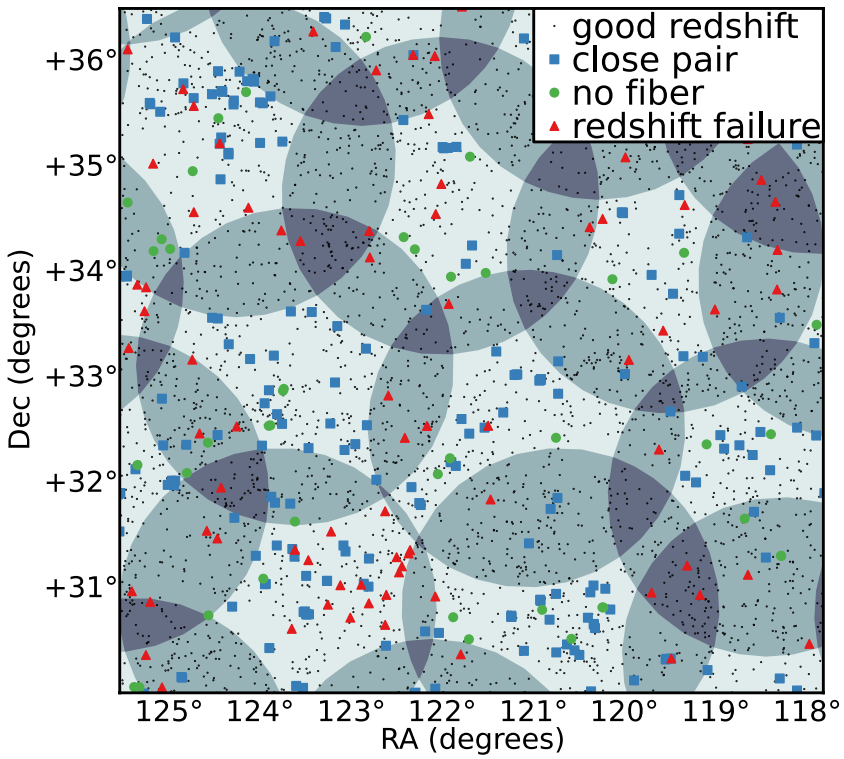

Figure 2.6: Distribution and outcome of targets for a region of the CMASS sample within BOSS. Large circles show the placement of observing tiles, with darker shading indicating overlap areas. Small black dots show the positions of targets for which a good redshift was obtained. Blue squares show targets that were not allocated a fibre because they were part of a close pair, and thus could not be observed due to the fibre-collision issue. Green circles show targets that were not allocated a fibre but were not part of a close pair. Red triangles show targets that were allocated a fibre but a good redshift was not obtained [232].

extremely difficult to model analytically, except in the case of some simple targeting strategies. Our method for measuring these probabilities in eBOSS is presented in Sec. 2.5.3. The second condition is that no pair has zero probability of being observed. Trivially, in this case the PIP weights would be undefined, and because the pair is never observed it is impossible to account for its effect on the clustering. In the absence of zero probability pairs PIP weights provide statistically unbiased corrections to the 2PCF. The theory of PIP weights and proof that they provide an unbiased correction are presented in [36]. Prior to eBOSS, PIP weights had also been successfully applied to the slit-based VIMOS Public Extragalactic Redshift Survey [188].

We combine PIP weights with the Angular Upweighting (ANG) method presented in [202]. Angular upweighting is predicated on the assumption that the subsample of the parent catalogue that is observed spectroscopically is statistically identical to the full parent catalogue. In that case the angular clustering of the parent catalogue can be used to reduce the variance of the observed subsample. The angular weights are defined simply as the ratio of the measured angular correlation function of the parent catalogue to the observed subsample. The combination of PIP+ANG weights therefore gives an unbiased and minimum variance estimate of the clustering. Particularly for eBOSS, including angular upweighting is an important tool for addressing the issue of collided pairs in single-pass regions, as described in Sec. 2.5.3.

2.5.3 Generating PIP+ANG weights for eBOSS

In order to generate PIP weights we must determine the selection probability for each possible pair in the survey. We accomplish this by generating multiple replicas of the survey target selection. The eBOSS tiling placement and fibre assignment strategy are detailed in Sec. 2.2.3. In particular, as discussed in Sec. 2.5.1, cases of fibre collision are resolved randomly using a predetermined "random seed", except where one of the possible targeting configurations results in more observable objects. This deviation from purely randomly selected targets and its effect on the PIP weights is explored in Sec. 2.5.4.

To generate replicas of the target selection we rerun the eBOSS target selection, changing the random seed to vary which members of collided pairs receive fibres. We perform a total of $N_{\rm runs} = 1860$ fibre assignment runs for the eBOSS samples, which provide a large enough sample to ensure the inferred selection probabilities would not be noise dominated. The first run used the actual random seed of the eBOSS target selection, so that the true run is included in our ensemble. This step is necessary to ensure that each actually targeted pair will have a non-zero selection probability in our finite number of survey realizations.

For the mock catalogues analyzed in Sec. 2.5.4 we perform $N_{\rm runs} = 310$ fibre assignments for each mock, balancing the number of fibre assignments against the number of mock catalogues for computational efficiency. In order to run the fibre assignment algorithm such a large number of times we modify it to remove unnecessary output steps and to permit parallel computation of the algorithm on a supercomputing cluster, while leaving the actual selection method unchanged. Details of these target selection replicas and modifications to the fibre assignment code are described in [190].

Given a set of survey realisations, the inverse probability, or equivalently the PIP weight w_{mn}, is simply the number of realisations in which a given pair could have been targeted divided by the number of times it was targeted. For efficiency, rather than storing pairwise weights for all $O(N^2)$ pairs, we store bitwise weights $w_i^{(b)}$ for each target, following the method established in [36]. These bitwise weights are simply binary arrays of length $N_{\rm runs}$, where each bit (either one or zero) represents the outcome of the corresponding fibre assignment run for target i (either this target is, or is not, included in run b). Bitwise weights can then be combined efficiently on the fly to compute the pairwise weights for targets m and n using:

$$w_{mn} = \frac{N_{\rm runs}}{\mathrm{popcnt}\left[w_m^{(b)} \& w_n^{(b)}\right]}, \tag{2.6}$$

where popcnt and & are standard bitwise operators. popcnt is the 'population count' operator that, given an array, returns the number of elements different than 0. & is the bitwise 'and' that, given two arrays of equal length, performs the logical 'AND' operation on each pair of the corresponding bits and returns the result as an array with length equal to that of the input arrays. The Individual-Inverse-Probability (IIP) weights, 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 realisation. w_m can be calculated simply by replacing $m = n$ in Eq. 2.6.

The weights applied to the pair counts when using the PIP and IIP weights are then:

$$\begin{aligned} DD(\vec{s}) &= \sum_{\vec{x}_m - \vec{x}_n \approx \vec{s}} w_{mn} w'_{\rm tot,m} w'_{\rm tot,n} \ , \\ DR(\vec{s}) &= \sum_{\vec{x}_m - \vec{y}_n \approx \vec{s}} w_m w'_{\rm tot,m} w_{\rm tot,n} \ , \end{aligned} \tag{2.7}$$

In Eq. 2.7, $w'_{\rm tot} = w_{\rm sys} \times w_{\rm noz} \times w_{\rm FKP}$ are the total weights described in Sec. 2.3.2, while $w_{\rm tot} = w_{\rm sys} \times w_{\rm noz} \times w_{\rm FKP} \times w_{\rm CP}$ are applied to the random catalogue, which is valid for all fibre assignment runs.

Because the eBOSS tile positions are fixed, collided pairs in single pass regions can never be resolved, meaning there are zero probability pairs in the sample. These pairs produce a systematic underestimation in the measured 2PCF. Initially, we investigated resolving this issue by rotating the parent catalogue on the sky in random directions and magnitudes, giving a non-zero probability for any pair to lie in a multi-pass region. Unfortunately, due to the observing choices made for eBOSS, this method is not able to faithfully reproduce the baseline of the parent catalogue. The method, as well as the reasons for rejecting it, are described in Appendix A.

Alternatively, the Angular Upweighting outlined in [202] can be used to de-bias the measurements at smaller scales. The parent catalogue contains all collided pairs, and gives the correct angular correlation function below the fibre-collision scale. The spectroscopic catalogue with PIP weighting misses all collided pairs in single-pass regions, but in multi-pass regions collided pairs can be resolved and, after applying the PIP weights, give an unbiased measurement of the clustering within their region. Since collided pairs in multi-pass regions are the only angular pairs with separations less than the fibre-collision scale observed in the survey, angular upweighting increases their weight to match the clustering of the parent catalogue, accounting for all the missing pairs in single-pass regions. This correction relies on the assumption that pairs missed due to fibre collisions in the single-pass regions are statistically equivalent to those targeted in the multiple-pass areas. It is worth noting that this is not necessarily true for eBOSS given the perturbation of the tile placement to correlate multi-pass regions with higher density regions (Sec. 2.2.3). However this effect is likely small enough to be subdominant to the statistical uncertainty, as shown by the recovery tests with mock catalogues presented in Sec. 2.5.4.

The angular weights used to upweight DD and DR are defined as,

$$\begin{aligned} w_{\rm ang}^{\rm DD}(\theta) &= \frac{DD^{\rm par}\left(\theta\right)}{DD^{\rm fib}_{\rm PIP}\left(\theta\right)}\,, \\ w_{\rm ang}^{\rm DR}(\theta) &= \frac{DR^{\rm par}\left(\theta\right)}{DR^{\rm fib}_{\rm IIP}\left(\theta\right)}\,, \end{aligned} \tag{2.8}$$

where θ is the angular separation of the pair. The superscripts -par and -fib in Eq. 2.8 denote pairs of targets from the reference parent sample and pairs of targets that receive fibres, respectively. The subscript PIP and IIP denote the fact that the pair counts are upweighted using the PIP weights for DD pairs, and IIP weights for the DR pairs. We measure the angular correlation function in 34 bins, with upper limits logarithmically spaced between $0.01 - 3^\circ$. It is worth noting that weights only significantly deviate from 1 below the fibre-collision scale, $\theta^{\rm (fc)} = 62''$.

Combining both of the corrections, the complete PIP+ANG weightings are:

$$
\begin{aligned}
DD(\vec{s}) &= \sum_{\substack{\vec{x}_m - \vec{x}_n \approx \vec{s} \\ \vec{u}_m \cdot \vec{u}_n \approx \cos\theta}} w_{mn}^{\mathrm{PIP}} w_m^{\mathrm{tot}} w_n^{\mathrm{tot}} \times \frac{DD_{\mathrm{par}}(\theta)}{DD_{\mathrm{fib}}^{\mathrm{PIP}}(\theta)} \,, \\
DR(\vec{s}) &= \sum_{\substack{\vec{x}_m - \vec{y}_n \approx \vec{s} \\ \vec{u}_m \cdot \vec{v}_n \approx \cos\theta}} w_m^{\mathrm{IIP}} w_m^{\mathrm{tot}} w_n^{\mathrm{tot}} \times \frac{DR_{\mathrm{par}}(\theta)}{DR_{\mathrm{fib}}^{\mathrm{IIP}}(\theta)} \,,
\end{aligned}
\tag{2.9}
$$

2.5.4 Validation using mock catalogues

We test the effectiveness of the PIP+ANG weights using a sample of 100 EZmocks for the LRG, ELG, and QSO samples (see Sec. 2.4.1). By default the EZmocks contain an approximation of fibre collisions, where objects that did not receive a fibre are removed and corrected using CP weights. In order to be able to accurately test the PIP+ANG correction we instead use the parent EZmock catalogues, which are already cut to match the eBOSS footprint, and apply the full eBOSS fibre assignment algorithm to realistically introduce fibre collisions, as they appear in the actual eBOSS samples. The full steps we take to prepare the 100 EZmock catalogues and ensure they are an accurate representation of the eBOSS samples are described in [190].

For each eBOSS sample (LRG, ELG, and QSO), we measure the clustering of the parent catalogue, the catalogues affected by fibre collisions and corrected using PIP weighting, and the catalogues affected by fibre collisions with the full PIP+ANG correction for all 100 EZmocks. In each case we measure the projected correlation function, w_p, as well as the first three even multipoles of the redshift-space 2PCF, ξ_0, ξ_2, and ξ_4, over a separation range of $0.1 - 100h^{-1}$Mpc. We then compare the mean of the corrected measurements to the mean of the true clustering of the parent mocks.

Fig. 2.7 shows the result of this comparison for the LRG EZmock projected correlation function. I choose to display the LRG measurement here because it is the most clustered sample, and therefore displays the largest differences caused by fibre collisions, and because it is the sample that is analyzed in Ch. 3 and Ch. 4. Equivalent plots for the ELG and QSO samples are shown in [190], although the same trends, features, and performance of the PIP and PIP+ANG corrections are observed for all three samples.

On scales above those corresponding to the fibre-collision radius the PIP and PIP+ANG corrections give nearly identical results, both of which are in agreement with the true clustering of the parent catalogue. For comparison, the lower panel shows the difference

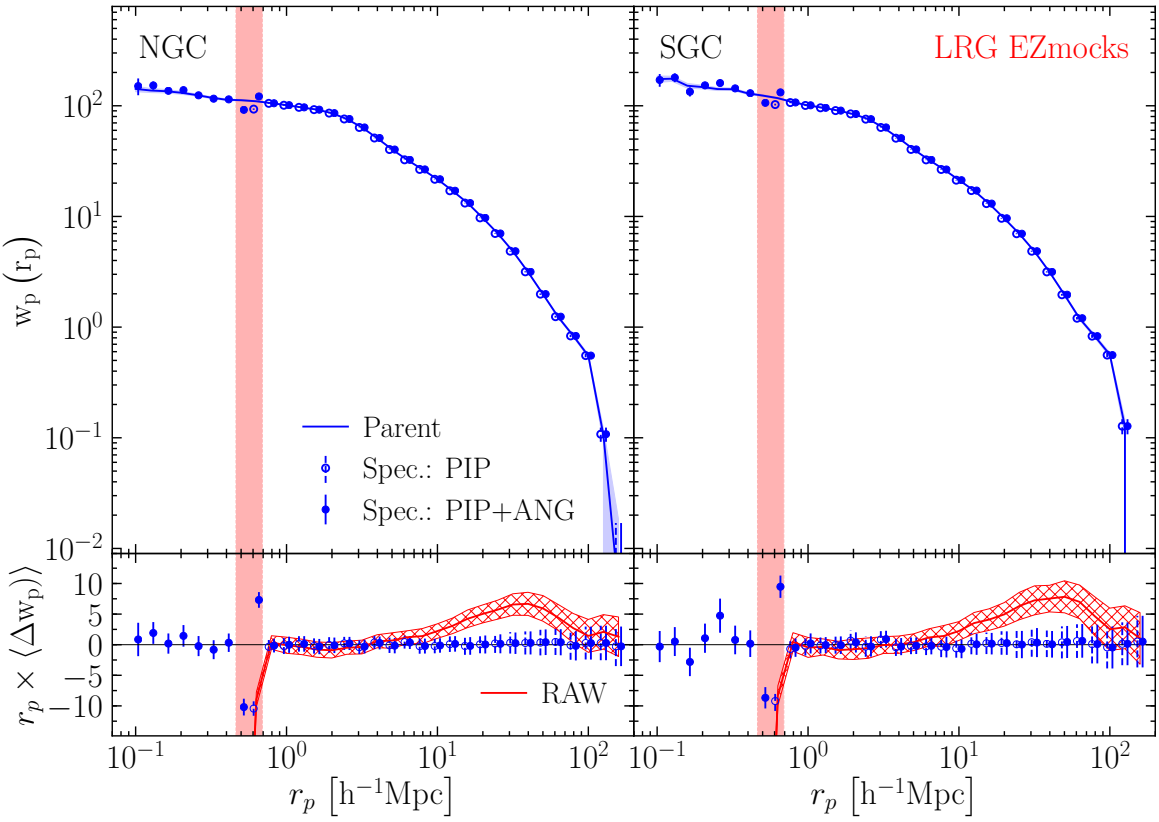

Figure 2.7: Mean projected correlation function measurements of 100 LRG EZmocks in the NGC (left column) and SGC (right column). *Top:* measurements of the parent mocks (solid lines with shaded bands), from catalogues affected by fibre collisions and corrected using PIP weights (empty markers with dashed error-bars), and corrected using PIP+ANG (filled points with continuous error-bars). The blue shaded bands and error-bars show the error on the mean. The vertical red shaded bands show the transverse scales corresponding to the fibre-collision angle between the minimum and maximum redshifts of the sample. *Bottom:* mean of the differences between the corrected measurements from mocks affected by fibre collisions and the corresponding parent mock. To reduce the range of variation, each quantity in the bottom panel is multiplied by r_p. For comparison, the red continuous lines and hatch regions (RAW) show the mean measurements where no fibre-collision correction is applied. Note: empty markers at scales smaller than the fibre-collision scale are not visible in the plot because they are well below the minimum limit set on the y axis.

with the uncorrected measurements of the catalogues affected by fibre collisions, which deviate significantly from the true clustering for measurements with separations between $10 - 100h^{-1}$Mpc. Below the fibre-collision scale the PIP+ANG weights continue to accurately recover the true clustering, but the PIP corrected measurements are much lower than the true clustering due to the large number of missing pairs in single-pass regions.

The only scales at which the PIP+ANG measurements deviate from the clustering of the parent catalogue is at the fibre-collision scale, shown as the vertical red band in the plot. The reason for this deviation is likely the zero-probability pairs introduced by the choice of maximizing the number of observable objects in the target selection algorithm for collided groups, discussed in Sec. 2.5.1. This deviation shows why it is preferable to choose an observing strategy that treats all objects equally (in this case, applying a fully random selection where necessary), rather than tuning the strategy to produce slightly more observations. In the case of all objects being statistically equivalent, PIP weights can recover the true clustering without bias. But in this case, where the observing strategy is tuned, there is an unrecoverable loss of information. However, it is important to note that this deviation is only visible because we have displayed the uncertainty on the mean of the 100 EZmock measurements in Fig. 2.7. This deviation is significantly smaller than the statistical uncertainty on a single mock measurement, and so is not a concern for the corrected eBOSS measurements.

Fig. 2.8 shows an equivalent plot for the multipoles of the LRG EZmocks. For the multipoles, the PIP+ANG correction is still able to recover the correct clustering on all scales, but the PIP correction is only able to recover the true clustering above a scale of $\sim 20h^{-1}$Mpc, below which it significantly underestimates the true clustering. This failure is also caused by missing fibre-collided pairs in single-pass regions, but unlike for w_p where these pairs are restricted to a particular transverse scale, in the 3D clustering they are spread over a range of separation. This can also be seen in the uncorrected measurements, which deviate significantly from the true clustering on all scales.

2.5.5 Applying PIP+ANG weights to eBOSS data

We generate PIP+ANG weights for the eBOSS DR16 LRG, ELG, and QSO samples, as described in Sec. 2.5.3. These measurements were made publicly available as part of SDSS DR16[5].

Fig. 2.9 shows the projected correlation function of the eBOSS DR16 LRG sample corrected by CP, PIP, and PIP+ANG weights for comparison. As was seen in Fig. 2.7,

[5]https://www.sdss4.org/dr16/

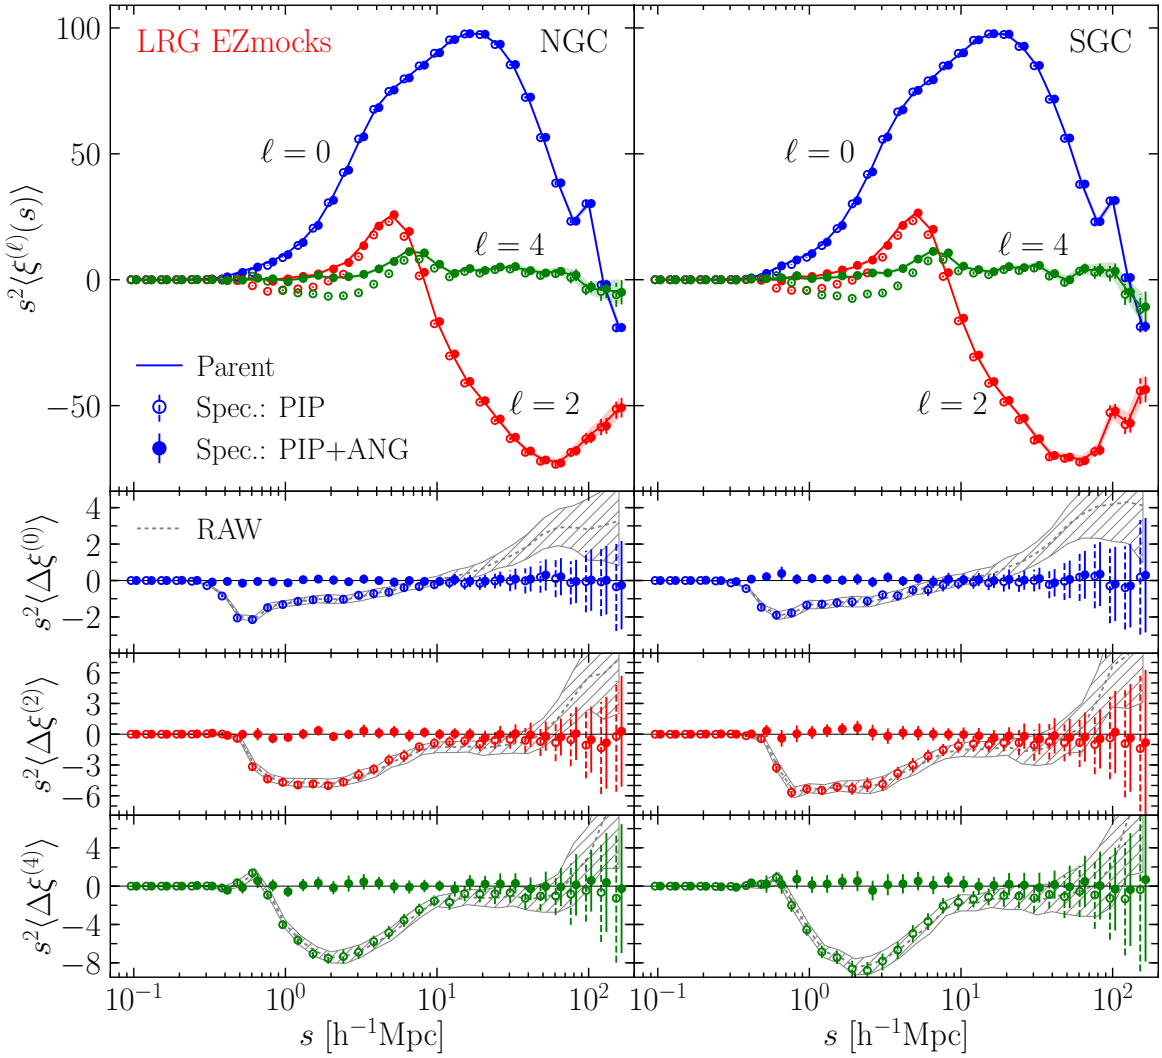

Figure 2.8: Same as Fig. 2.7, but for the monopole (blue), quadrupole (red), and hexadecapole (green) of the redshift-space 2PCF. In the lower panels, showing the differences between the corrected measurements and the measurements of the parent mocks for each multipole, the uncorrected measurements are now shown as grey continuous lines and hatch regions (RAW). All measurements are multiplied by the square of the separation to better simultaneously display the small- and large-scale behaviour.

above the fibre-collision scale the PIP and PIP+ANG corrections are in good agreement, however below the fibre-collision scale the PIP corrected measurements drop significantly, showing the importance of including angular upweighting. The CP and PIP corrections are generally in agreement above the fibre-collision scale. This agreement shows that the selection probabilities are highly uncorrelated on these scales, and so can be well approximated using an empirical prescription such as the CP method. This is to be expected given that the eBOSS fibre assignment algorithm used a random seed to resolve most instances of fibre collision.

Fig. 2.10 shows a similar plot for the multipoles. As was found in Fig. 2.8, changing to the three dimensional clustering spreads out the scales affected by the missing observations in single-pass regions. The result is that the PIP and CP corrections are only in agreement with the PIP+ANG corrected measurements above $\sim 20h^{-1}$Mpc. These measurements highlight the importance of properly correcting fibre collisions for small-scale analyses, like those presented in Ch. 3 and Ch. 4.

2.6 ELG targeting depth systematic

After completing observations of the eBOSS Emission Line Galaxy (ELG) sample we discovered that there was an observational systematic affecting the clustering that could bias the cosmological analysis of the sample. After some investigation, we linked this systematic to the photometric calibration of the DECaLS data used to target the ELGs. In this section I describe this systematic, and the work done by myself and others to correct it. I begin by giving an overview of the discovery and evidence for this observational systematic, and the failure of simple techniques, such as additional masking, to resolve the issue, despite removing significant amounts of observed data (Sec. 2.6.1). I investigate correcting the issue without discarding any data by applying a weight-based correction, which I introduce in Sec. 2.6.2. The key requirement of this method is a measurement that is highly correlated with the variations in density caused by the uneven calibration. In Sec. 2.6.3 I describe my work using the magnitude differences of stars between the DECaLS DR7 and DES DR1 observations as a proxy for variations in the photometric calibration. I then discuss the results of this weight-based correction on the most strongly affected observation chunk, and how it informs our understanding of the systematic issue (Sec. 2.6.4). Finally, in Sec. 2.6.5 I describe the corrections applied to the ELG sample when using the data for cosmological inference.

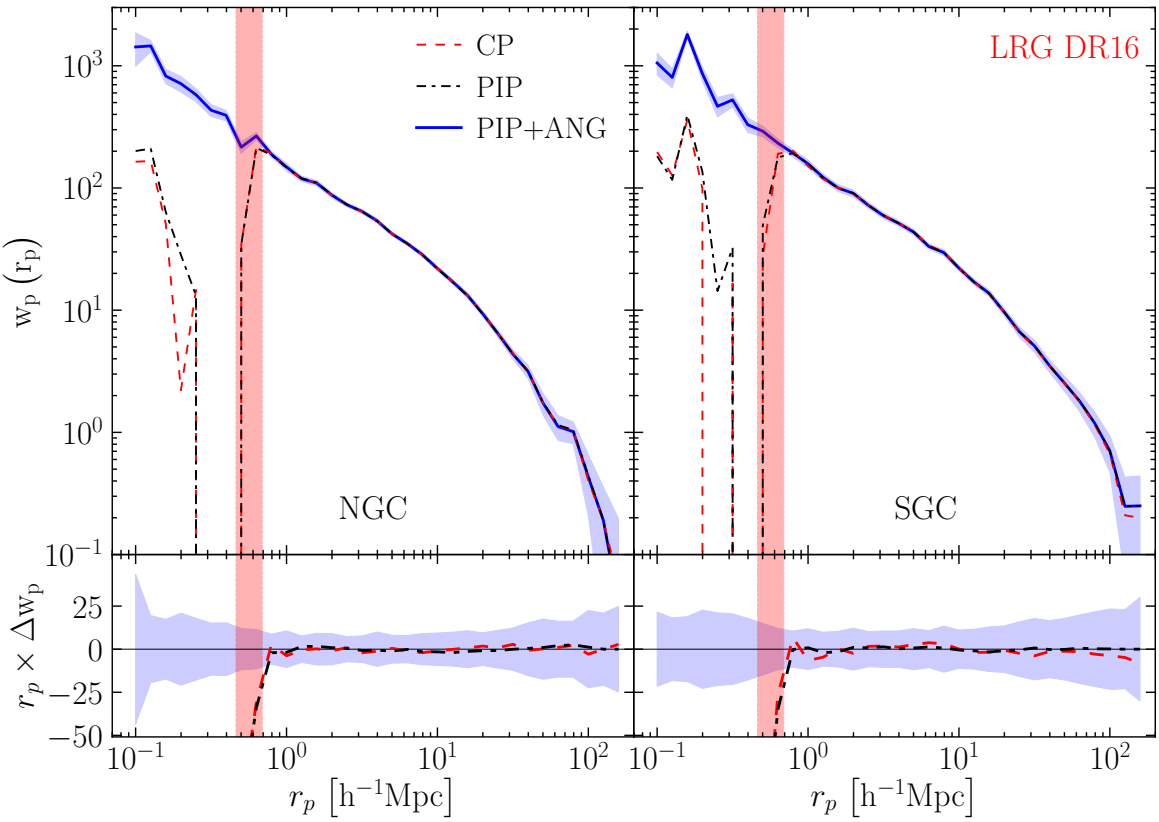

Figure 2.9: *Top:* projected correlation function measurements of the eBOSS DR16 LRG sample using different fibre-collision correction schemes: CP weights (dashed red), PIP weights (black dashed-dotted), and PIP+ANG weights (blue). The shaded region around the PIP+ANG correction shows the 1σ uncertainty estimated from 100 EZmocks. The vertical red shaded bands show the transverse scales corresponding to the fibre-collision angle between the minimum and maximum redshifts of the sample. *Bottom:* differences between the CP and PIP corrected measurements and the PIP+ANG correction. To reduce the range of variation, each quantity in the bottom panel is multiplied by r_p. The left column shows the result for the NGC, and the right column shows the SGC.

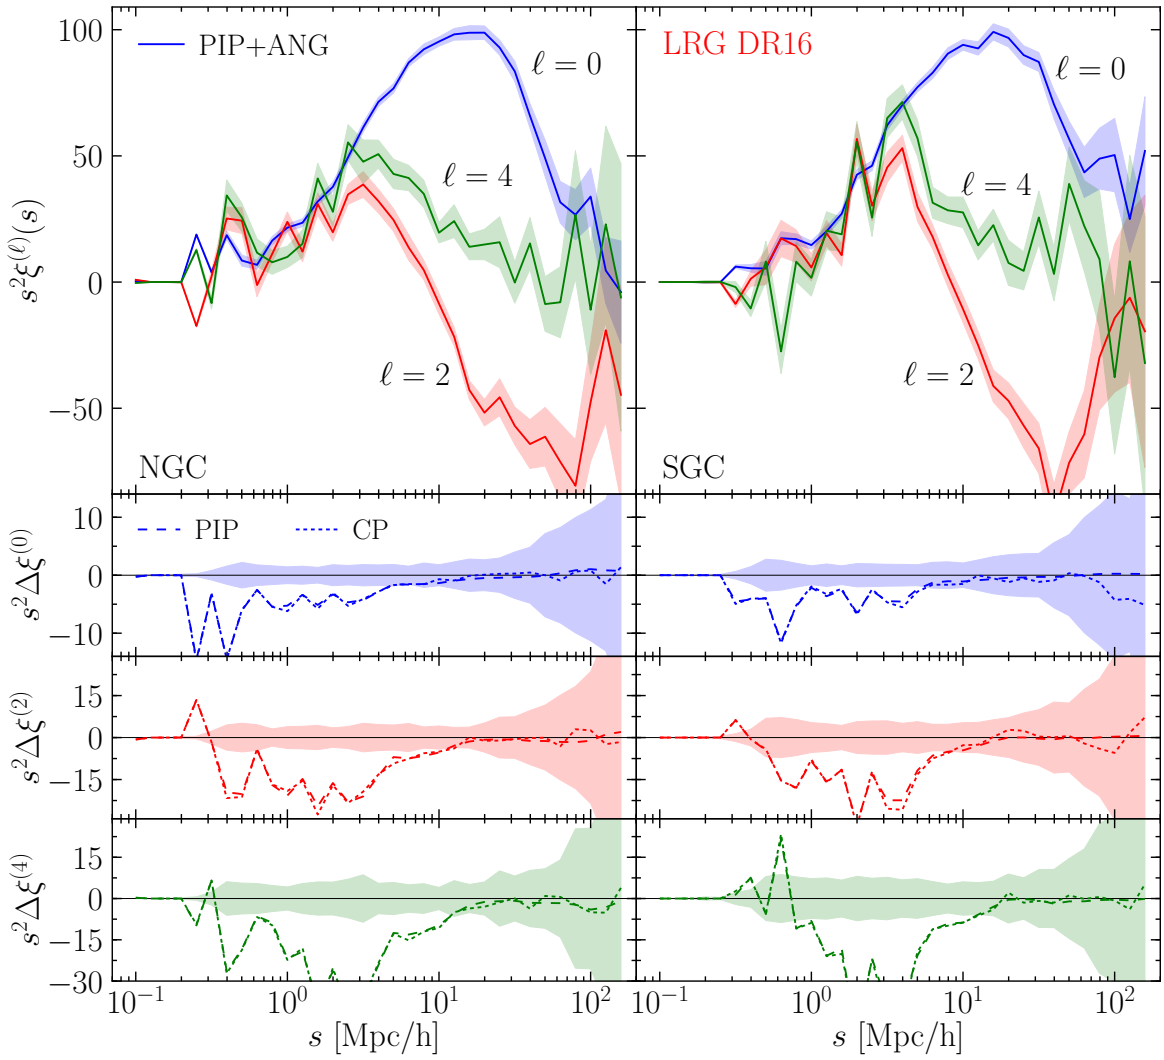

Figure 2.10: Redshift-space 2PCF multipole measurements of the eBOSS DR16 LRG sample in the NGC (left column) and SGC (right column). *Top:* the monopole (blue), quadrupole (red), and hexadecapole (green) of the redshift-space 2PCF corrected by PIP+ANG weighting. The shaded bands show 1σ uncertainties estimated from 100 EZ-mocks. *Bottom:* differences between the CP and PIP corrected measurements and the PIP+ANG correction. All measurements are multiplied by the square of the separation to better simultaneously display the small and large-scale behaviour.

2.6.1 Identifying the observational systematic

The eBOSS ELG targets are selected using photometric data from DECaLS DR3 and DR5 (see Sec. 2.2.2 for more details). These DECaLS data releases are a collection of all public grz-band DECam imaging taken over the DECaLS footprint, and were analyzed using the Tractor[6] image processing software. While primarily composed of DES and dedicated DECaLS observations, observations from other programs with different calibration prescriptions are also included.

The danger of using inconsistent calibration techniques is variable depth across the survey, effectively shifting the magnitude cuts used for target selection (see Sec. 2.2.2) between different regions. Because galaxies are not uniformly distributed in colour-space (see e.g. Fig. 2.2), shifting the target window can lead to a greater flux of galaxies entering the target window than leaving, or vice-versa. Changing the number of selected targets in patches introduces density variations across the survey, causing additional angular correlations in the data that are not cosmological in origin. While not initially a concern for the eBOSS ELGs, after observations were complete it was discovered that the ELG sample displayed excess angular clustering and a larger quadrupole compared to mock galaxy catalogues. This issue is also inconsistent between the observation chunks, with the **eboss22** ELGs showing both the largest angular correlation function and quadrupole. **eboss22** is the largest ELG chunk, covering 445 $\deg^2$ and containing 69 071 reliably observed ELG spectra, making up 35% of the entire eBOSS ELG sample [223].

In order to assess whether this issue is related to the photometric calibration we compared the angular clustering of DECaLS DR7, which contains the original data used to target the eBOSS ELGs, including the additional observation programs, and the angular clustering of the DES DR1 [2] sample within the **eboss22** footprint. DES is an additional photometric survey performed using DECam that overlaps with DECaLS, but using different observing parameters and SExtractor[7][31] for imaging processing. The results of this comparison are shown in Fig. 2.11. The DECaLS data shows significant excess angular clustering compared to DES on the scale of a few degrees, matching what was observed for the **eboss22** ELG sample. The DES data is in agreement with the prediction from the GLAM-QPM mocks, which are constructed to match the clustering of the ELG sample (see Sec. 2.4.1). The disagreement between the DECaLS and DES observations over the same footprint and with the same instrument implies that the source of the issue is the DECaLS photometric calibration and image processing, a conclusion supported by the agreement between the DES and GLAM-QPM mocks.

[6] https://github.com/dstndstn/tractor
[7] https://www.astromatic.net/software/sextractor/

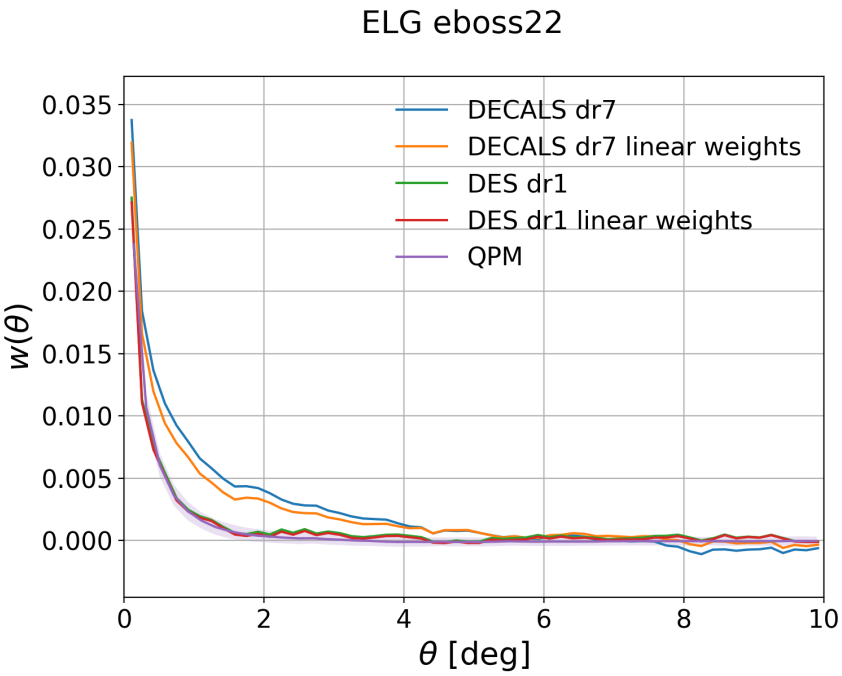

Figure 2.11: Comparison of angular correlation functions over the **eboss22** footprint. Blue and orange lines show the clustering of DECaLS DR7 while green and red show the clustering of DES DR1, in both cases with and without a linear correction to the clustering. The purple line shows the clustering of the GLAM-QPM mocks, with the shaded region showing the 1σ uncertainty. This plot was produced by Arnaud de Mattia for the eBOSS Collaboration.

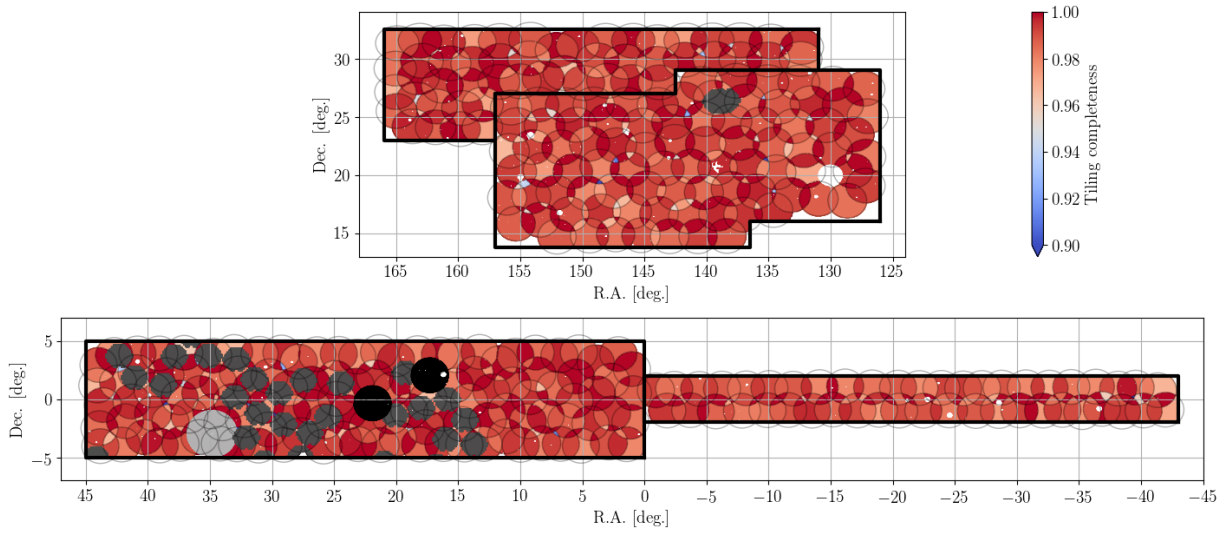

Figure 2.12: The footprint of the eBOSS ELG program. The upper panel shows the NGC chunks, where **eboss25** is at higher declination and **eboss23** is at lower declination. The lower panel shows the SGC chunks, with **eboss22** at positive right ascension (the left of this plot) and **eboss21** at negative right ascension. The blue-red colourbar shows the tiling completeness in each sector. The open black circles show the positions of the eBOSS plates. Filled circles show additional angular veto masks that are added after observation: the star Mira (light gray), DECam pointings with inconsistent photometric calibration in DECaLS DR3 and DR5 (dark gray), and eBOSS spectroscopic plates with abnormally high failure rates that are removed from the analysis (black) [223]

Unlike DECaLS DR3 and DR5 that are used for the ELG targeting, or DR7 that is analyzed in Fig. 2.11, DECaLS DR8 restricted the data included in the catalogue to mostly DES and DECaLS observations. DR8 also uses a significantly improved photometric calibration procedure to address the issues discovered in DR7. By comparing DECaLS DR8 to the data releases used for targeting we are able to identify the exposures used in targeting with the most extreme variations in calibration. These exposures are shown in Fig. 2.12, and are primarily found in **eboss22**, matching the excess angular clustering and quadrupole.

We attempted to correct this systematic by adding the most strongly affected DECaLS exposures, shown in Fig. 2.12, to the final angular veto mask. However, this additional masking only slightly improved the agreement of the angular correlation function and monopole with the mock measurements, and has little effect on the quadrupole and hex-

adecapole. The additional masks also remove tens of thousands of observed objects from the ELG sample, with even more aggressive masking resulting in even more lost data.

2.6.2 Weight-based correction

Rather than masking the miscalibrated exposures and reducing the size of the sample, I investigate correcting this systematic using a weight-based approach. This method has been successfully applied to correct other artificial density trends in the observed data, as described in Sec. 2.3.2.

The basis of this method is to divide the data into equal area angular pixels on the sky, using HEALPix, and compare the density in each pixel to a measurement that is correlated with the spurious density pattern. To be effective this method requires knowledge of the cause of the variations in density, in order to produce a strong trend between the measurement and the density variations. Weights can then be defined from the inverse of this trend, which when applied to the data will remove the density variations and the additional angular clustering. As long as the cause of the density variation is not correlated with the cosmological clustering, as is typically the case for observational systematics and foreground effects, the true signal will be unaffected and no data needs to be discarded.

The key to this method is identifying a statistic that is highly correlated with the false density variations. In this case the variation in density is correlated with the DECaLS photometric calibration, as shown in Fig. 2.12, but is not present in the DES observations over the same footprint, as shown in Fig. 2.11. As a proxy for variations in the DECaLS calibration, I investigate differences in the apparent magnitudes of stars between DECaLS DR7 and DES DR1 in the g, r, z-bands used to target the eBOSS ELGs. DES provides a consistent calibration over the same footprint and observed with the same instrument, so the magnitude differences are expected to be highly correlated with the variations in calibration of the DECaLS exposures. Stars are an ideal probe of this difference because they are point sources, which allows for a simple cross-match between the two surveys and a consistent estimate of the total flux. Stars are also numerous, giving a large signal-to-noise ratio in each pixel.

2.6.3 DES-DECaLS cross correlation

Using the National Optical-Infrared Astronomy Research Laboratory (NOIRLab) Astro Data Lab[8] (previously the National Optical Astronomy Observatory (NOAO) Datalab) I

[8]https://datalab.noirlab.edu/

cut the DECaLS DR7 main Tractor catalog in Right Ascension (RA) and Declination (Dec) to match the **eboss22** footprint ($0 <$ RA < 45, $-5 <$ Dec < 5). I select objects classified as stars by the Tractor image processor with positive, non-zero flux in the g, r, z-bands used for ELG targeting. I then perform a nearest-neighbour cross-match with the DES DR1 catalog to find the closest object within $1''$, removing any object with negative or zero flux in the three bands of interest.

I set an initial maximum cross-match distance of $5''$. In Fig. 2.13 I examine the number of objects matched as a function of cross-match distance. The number of objects decreases rapidly between $0''$ and $1''$, before increasing again. This behaviour is consistent with the majority of good matches occurring within $1''$, beyond which a false match is more likely to occur. To confirm this hypothesis I examined the relationship between the observed magnitude in each survey for the matched objects. I found that for objects matched within $1''$ the magnitudes in both surveys were highly correlated, as expected for correct matches given the differences in calibration between the two surveys. Objects matched between $1''$ and $5''$ tended to have uncorrelated magnitudes, with the DECaLS magnitude often greater than 25. These objects are most likely artifacts identified by the image processing software, rather than true stars. The likelihood of matching these artifacts increases with larger cross-match distance, but is greatly reduced when the maximum cross-match distance is limited to $1''$. I therefore adopt $1''$ as an upper limit to the cross-match distance, and find a total of $N_{\rm match} = 2\ 152\ 681$ matched objects within the **eboss22** footprint.

To generate the weights I divide the sky into equal area HEALPix pixels with $N_{\rm side} = 128$. Within the HEALPix method the total number of equal-area pixels is $N_{\rm Tot} = 12 \times N_{\rm side}^2$, so the value of $N_{\rm side}$ determines the size of each pixel. I choose $N_{\rm side} = 128$ so that the size of each pixel approximately matches the field-of-view of the DECam exposures used for targeting, although $N_{\rm side} = 256$ and $N_{\rm side} = 512$ were also tested with similar results. After dividing the footprint I find the pixel-by-pixel average magnitude difference between the two surveys for the cross-matched stars, in each of the three bands. Fig. 2.14 shows the average magnitude difference in each pixel of **eboss22** for each of the three bands. The r-band map shows the largest and most consistent differences, with most pixels showing a positive $m_{r,\rm DES} - m_{r,\rm DECaLS}$ difference. The largest shifts, around 0.05 dex in magnitude, are concentrated towards the larger RA end of the chunk, matching the area most heavily affected by poor exposures in Fig. 2.12. This miscalibration shifts the boundaries of the colour cuts based on the r-band magnitude used to target the ELG sample (see Eq. 2.3, Eq. 2.2, and Fig. 2.2 in Sec. 2.2.2) since the same shifts are not present in the g and z band maps, which show smaller shifts more evenly distributed around 0.

After determining the mean magnitude difference in each pixel I count the number of galaxies in the full **eboss22** ELG catalog and the number of objects from the random

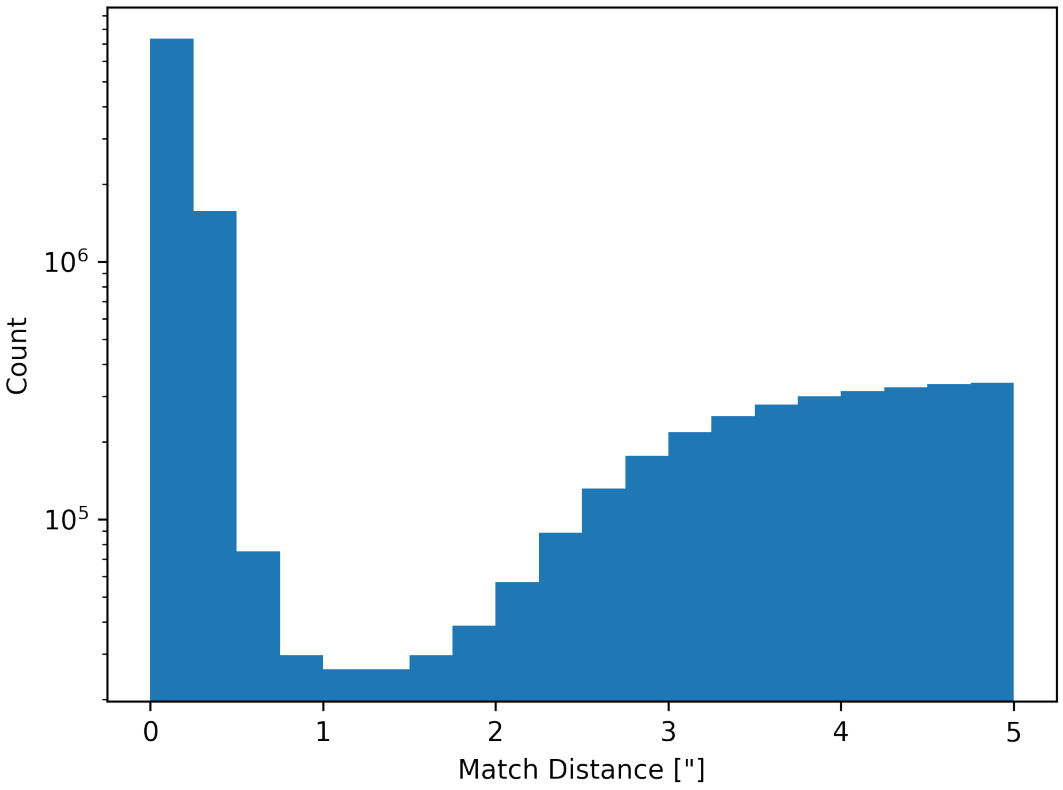

Figure 2.13: Histogram of the angular separation of stars matched in DES DR1 and DECaLS DR7.

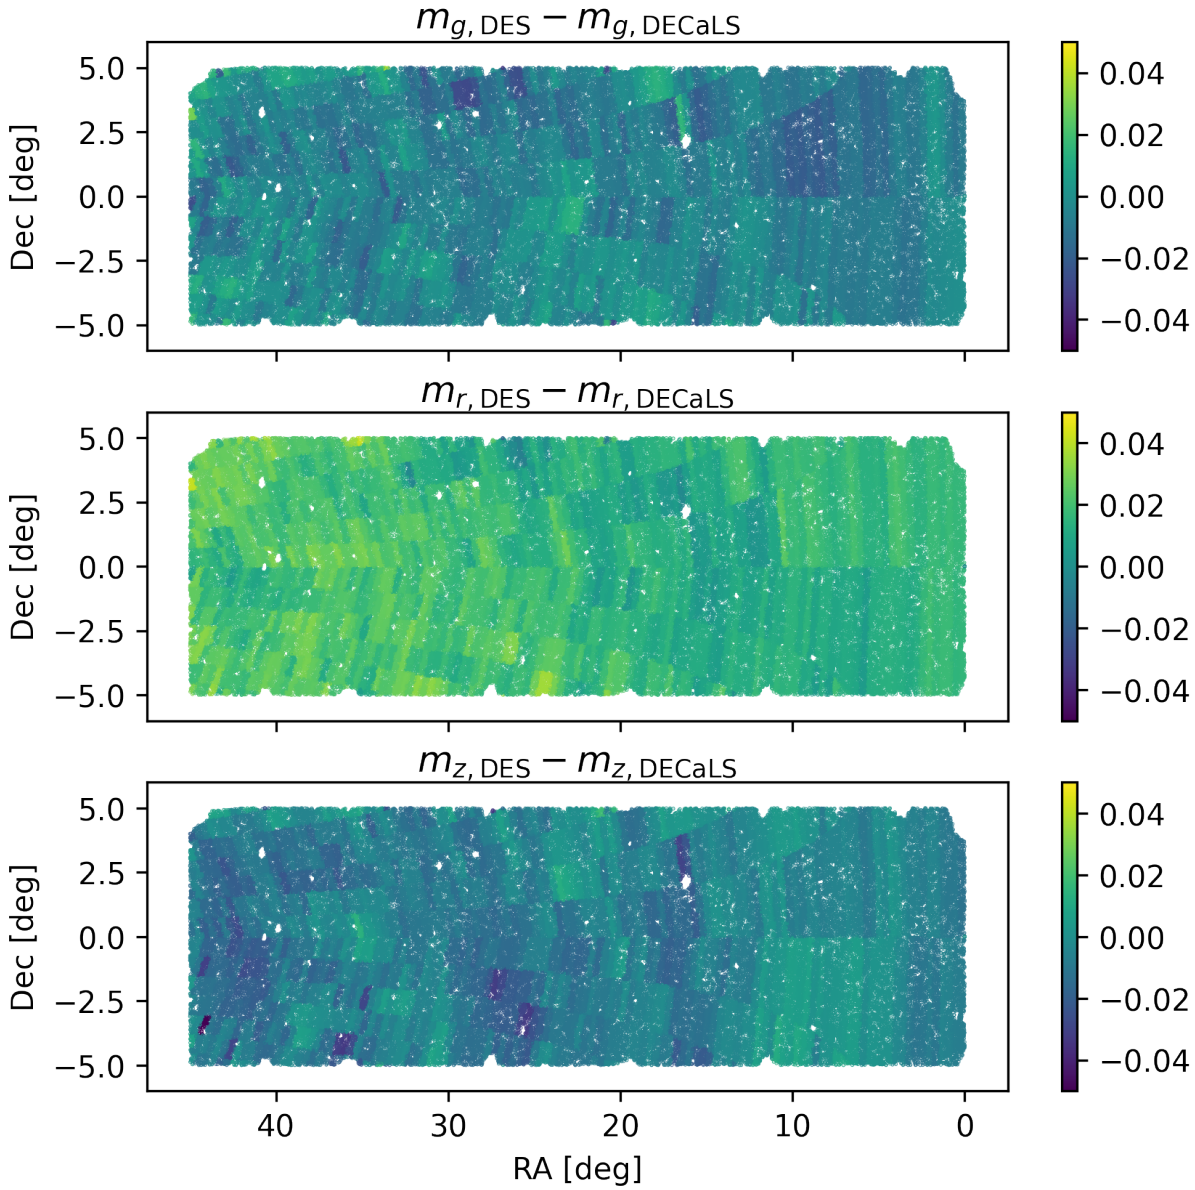

Figure 2.14: Mean DES-DECaLS stellar magnitude differences in HEALPix pixels for the g-, r-, and z-bands within **eboss22**.

catalog in each pixel. To find the target density of each pixel I calculate the ratio of target galaxies to randoms, normalized by the average across the chunk. I then correlate the density with the magnitude difference, divided into 10 equal width magnitude bins. Fig. 2.15 shows the density of **eboss22** ELG targets in HEALPix pixels, plotted against the average stellar magnitude difference between DES and DECaLS in that pixel. A clear trend in density is observed in the r-band for the unweighted sample, with a slighter trend in the g-band, and no clear trend in the z-band. Inversely weighting by the density in each bin removes all trends, and gives a sample unbiased in terms of the DES-DECaLS magnitude difference.

2.6.4 Performance of the weight-based correction

I inversely weight each galaxy by the density in its pixel's bin to remove any trends, and store the weights for all three bands. To test the success of the method I calculate the angular correlation function, as well as the 3D monopole, quadrupole, and hexadecapole, using the Landy-Szalay estimator for both the corrected and uncorrected samples (see Sec. 1.3.2).

These measurements are shown in Fig. 2.16. The weighted correction removes some excess angular clustering, shifting away from the measurement of the DECaLS DR7, which contains the miscalibrated observations, and towards the correctly calibrated DES DR1, but does not fully eliminate the issue. The correction also has a slight effect on the monopole, but no significant effect on the quadrupole or hexadecapole. Thus the quadrupole remains positive on large scales, in disagreement with measurements from the mock catalogues. These changes are similar to what is observed when the miscalibrated exposures are masked, but without any data being removed. The similarity of the results implies that the weight-based method can successfully remove the angular component of the photometric systematic without the need to remove data, but neither method fully corrects the issue.

This failure is likely caused by a redshift dependence in the density variations, which cannot be corrected with weights determined on a purely angular basis. Some redshift dependence is expected in this case because of the colour cuts used to select targets. These cuts are the same for all galaxies, but because they are applied using apparent magnitudes they will occupy a different region of colour-space in the rest frame of low redshift galaxies compared to high redshift galaxies. As the window shifts in colour-space the flux of possible targets across the boundaries will change, impacting the way in which the variance in calibration affects the density of targets.

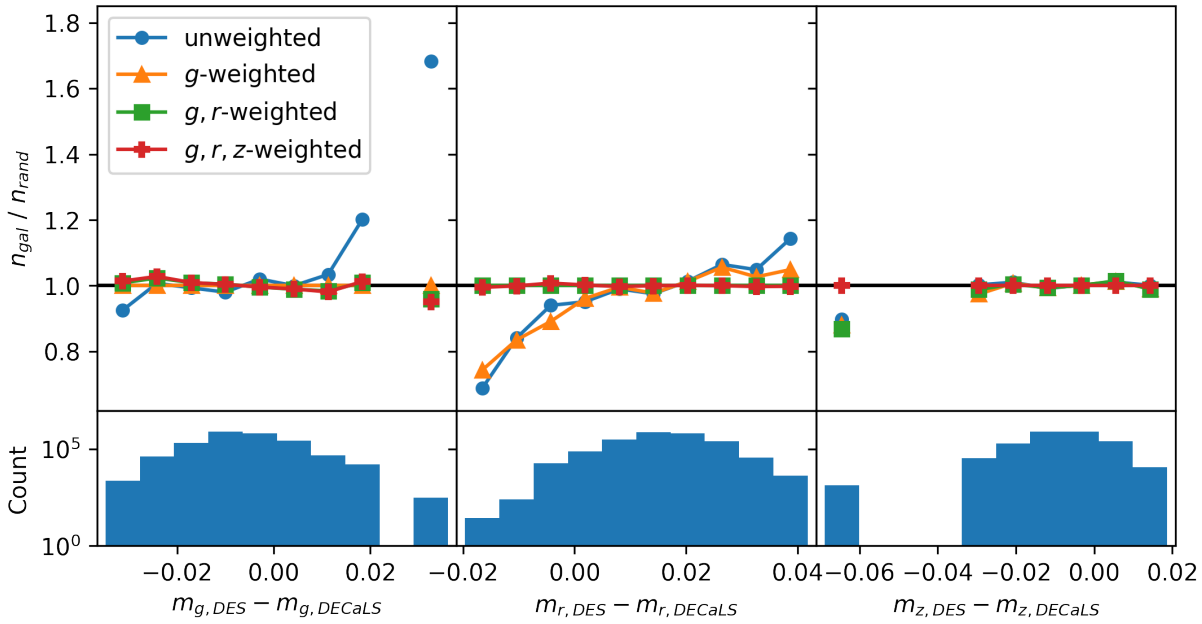

Figure 2.15: The density of **eboss22** ELG galaxies in HEALPix pixels on the sky, plotted against the average magnitude difference between stars in DES DR1 and DECaLS DR7 in that pixel. Density is calculated from the ratio of targeting catalog galaxies to random catalog points in that pixel, and normalized by the average galaxy-random ratio of the chunk. Blue circles show the original, unweighted distribution, orange triangles show the distribution weighted so that the density as a function of g-band magnitude difference is uniform, green squares show the distribution with both g- and r-band weights, and red crosses show the distribution with weights applied for all three bands. *Left:* g-magnitude difference. *Centre:* r-magnitude difference. *Right:* z-magnitude difference. The lower panels show the number of stars used to calculate the average magnitude difference in each bin.

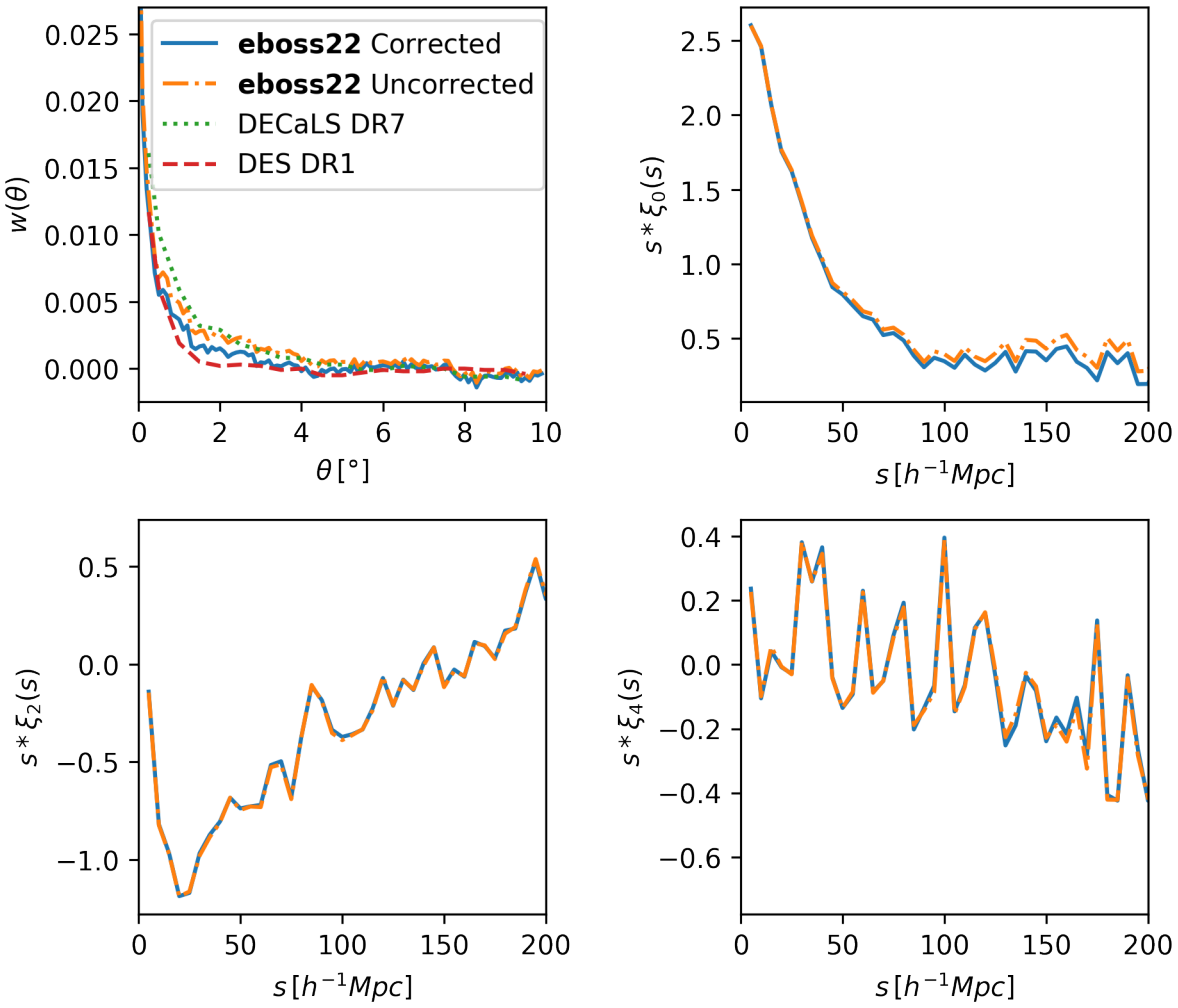

Figure 2.16: Effect of weight-based correction on **eboss22** ELG correlation functions. Uncorrected (orange) and corrected (blue) measurements are shown for the angular correlation function (top left), and the monopole (top right), quadrupole (bottom left), and hexadecapole (bottom right) of the 3D correlation function. Angular correlation functions from the DECaLS DR7 (green) and DES DR1 (red) are also shown for reference. Multipole measurements are multiplied by the pair separation to better display the large-scale behaviour.

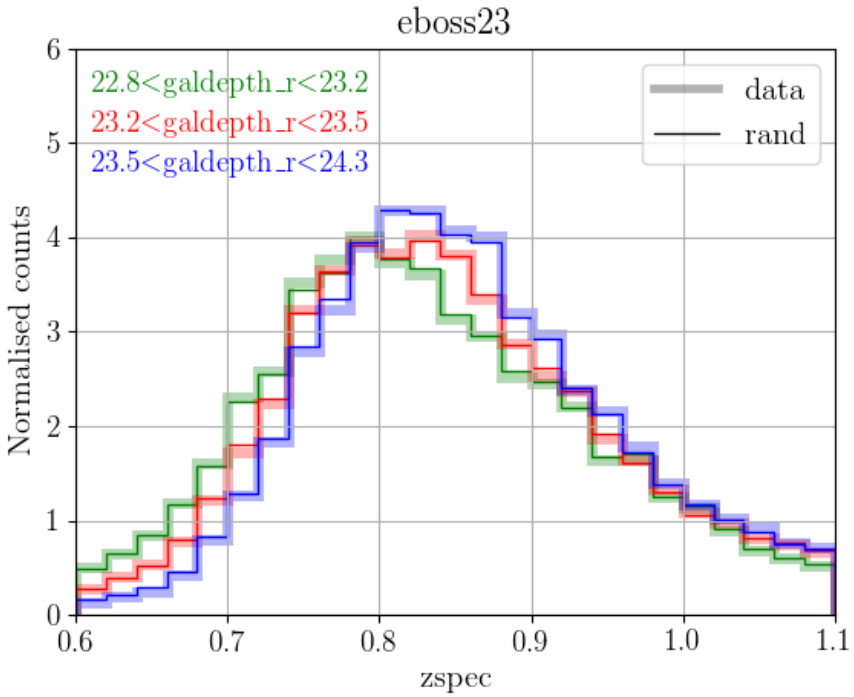

Figure 2.17: Redshift distribution of **eboss23** ELGs divided into three bins of target imaging depth in the r-band. Faint, thick lines show the distribution of the data in each bin, while thin lines show the distributions of the random catalogues. [223]

In Fig. 2.17 we demonstrate this relationship between imaging depth and redshift distribution for the **eboss23** chunk. The chunk is divided into three bins by imaging depth of the targeting photometry in the r-band. **eboss23** was chosen to display because it exhibits the clearest separation between the three bins. r-band imaging depth is used because it shows the strongest effect, in agreement with my results from the weight-based correction (Fig. 2.15). Fig. 2.17 shows that regions with shallower imaging also skew to lower redshifts, because lower apparent magnitude galaxies at high redshift are excluded from the colour selection. In order to fully correct this issue the redshift dependence would need to be included in the weights, requiring a more complicated statistic than the magnitude differences between DECaLS and DES.

2.6.5 Treatment for cosmological inference

My results show that it is not possible to correct the variations in density caused by inconsistent photometric calibration without a loss of information. We therefore adopt two corrections for the cosmological analysis of the eBOSS ELGs: calibrating the radial distributions of the randoms separately in bins of equal targeting depth, and nulling angular modes in the data.

In eBOSS the redshifts of the random catalogue are assigned using the *shuffled* method, where the redshifts of the randoms are randomly drawn from the redshifts of the data catalogue, ensuring that the randoms exactly capture the radial window function of the survey. Due to the variations in redshift distribution with imaging depth shown in Fig. 2.17, for the ELG data we split each chunk into subregions with approximately equal imaging depth and draw the redshifts of the randoms separately for each subregion. The imaging depth at each point in the survey is defined as a linear combination of the 5σ flux detection limits of the imaging survey in the g-, r-, and z-bands. For complete details, see Sec. 3.7 of [223]. The resulting $n(z)$ for the randoms is shown as the faint, thick lines in Fig. 2.17, closely matching the distribution of the data in each depth bin.

This method ensured the randoms capture the correct radial window function across the survey, although determining the window function from the data in this way introduces an integral constraint that we must address in the analysis. The purpose of the window function is to give an expected density of observed galaxies in any redshift-space position within the survey in the absence of clustering. If the window function has been accurately determined, then taking the difference of the window function from the data gives a fair estimate of the density fluctuations, which contains the cosmological information of interest. However, the window function of spectroscopic surveys is complex in nature due to the many observational challenges and their various angular and redshift dependencies. For this reason it is typically estimated from the data, as described above for the *shuffled* method of determining the radial window function. These methods fix the window function used in the analysis to match the data, which is not necessarily the case for the true window function due to the presence of large-scale clustering modes beyond the size of the survey. While not directly analogous, this issue could be compared to the more common problem of estimating the standard deviation of a sample. If the population mean is not known and the mean must be estimated from the sample itself, the normalization used in the calculation must be adjusted to account for the loss of freedom. Integral constraints provide the correction to the window function in our clustering analysis. For the ELG sample the radial integral constraint is particularly important to include because we tune the $n(z)$ in subregions of the survey, so we follow the method laid out in [75] to account

for it.

The second correction is to completely null all angular modes in the survey. This ensures that the angular variations in density will not bias the cosmological inference, at the expense of discarding all angular information. This is done separately in the Fourier space analysis [76] and the configuration space analysis [258]. In Fourier space the randoms are weighted within HEALPix pixels on the sky to exactly match the density of the data within that pixel. This is equivalent to introducing an angular integral constraint to correct the observational systematic, which is described in [75] and applied in [76].

The configuration space analysis uses a modified 2PCF that nulls the angular modes of the clustering, based on the shuffled 2PCF method presented in [49]. In this method the angular positions of randoms are also selected from the data, analogous to the *shuffled* method for determining the radial positions, to form a shuffled catalogue. The shuffled 2PCF is then estimated using a modification of the Landy-Szalay estimator:

$$\hat{\xi}^{\text{shuff}}(r_\perp, r_\parallel) = \frac{DD(r_\perp, r_\parallel) - 2DS(r_\perp, r_\parallel) + SS(r_\perp, r_\parallel)}{RR(r_\perp, r_\parallel)}, \tag{2.10}$$

where S denotes the shuffled random catalogue. Since the angular positions of the shuffled random catalogue are drawn from the angular positions of the data catalogue, DD, DS, and SS will all possess the same angular clustering, which will cancel out in the final estimator, nulling all angular modes.

[49] presented a model for the shuffled 2PCF, calculated by subtracting terms integrated over the line of sight. However, the shuffled 2PCF method was constructed with a focus on BAO modelling, whereas the ELG analysis includes both BAO and RSD. For this reason the modelling of the shuffled 2PCF is poorly suited to the ELG analysis, because the small radial scales included in the line-of-sight integrals are not trusted for the CLPT-GS model used for the RSD. Instead, [258] introduces a novel modified 2PCF based on the shuffled 2PCF model, but excluding small radial scales by controlling the bounds of the integration. Full details of the model as well as its validation and performance are included in [258]. The results of both the configuration-space and Fourier-space analyses, with the targeting depth systematic removed, are presented in Sec. 2.7.

2.7 Cosmological constraints from clustering measurements

The final eBOSS data release was accompanied by a coordinated release of papers measuring BAO and RSD in the clustering of the LRG [23, 111], ELG [223, 258, 76], and quasar [141, 196] samples, making use of both configuration space and Fourier space measurements. In general, two measurements are performed for each clustering sample: a BAO-only measurement made by fitting a template to the relevant clustering measurements, and a full-shape analysis that simultaneously measures BAO and RSD. The Lyα forest sample is analyzed separately, and is presented in [85]. In this section I present the methods and models used to make the BAO-only template measurements (Sec. 2.7.1) and full-shape BAO+RSD measurements (Sec. 2.7.2) for each clustering sample, as well as the final results of those analyses. In Sec. 2.7.3 I describe the impact of these measurements, together with the Lyα forest results and the results of previous SDSS measurements, on our understanding of the cosmological model.

2.7.1 BAO-only measurements

The eBOSS BAO analyses uses two approaches to measuring the BAO scale. The first is to use a template-fitting approach, which is effective at isolating the measurement of the BAO scale without contamination from redshift space distortions, but at the expense of excluding some broadband shape information. The second is to allow the shape of the model to vary in order to extract information from all data points, termed a full-shape fit, which simultaneously measures the BAO scale and RSD signal. In this section I will focus only on the template fitting approach, and discuss the RSD constraints from the full shape analysis in Sec. 2.7.2.

In order to extract 3D clustering information from the large-scale structure catalogues a fiducial cosmological model must be assumed to convert from angular positions and redshifts to 3D positions. The fiducial cosmology used in the eBOSS analyses is the same as that used in the final BOSS analysis [9]. It consists of the following values for the major cosmological parameters at redshift $z = 0$: $\Omega_m = 0.310$, $h = 0.676$, $\Omega_b h^2 = 0.022$, $\Omega_c = 0.260$, $\sigma_8 = 0.800$, $n_s = 0.970$, and $r_{\rm drag} = 147.78\,{\rm Mpc}$. Any deviations from this fiducial cosmology or additional parameters are noted in the relevant analysis paper. Once the fiducial cosmology is set and redshifts have been converted to distances, clustering measurements can then be extracted from the data.

These clustering measurements display a peak from BAO, however the scale at which the peak occurs depends on the fiducial cosmology. In order to measure the BAO scale independent of the fiducial cosmology, a template for the clustering is generated at the desired point in cosmological parameter space, and then the distances used in the template are scaled by a set of dilation parameters that account for the difference between the cosmology being investigated and the fiducial cosmology. These dilation parameters provide a measurement of the BAO scale that is independent of the fiducial cosmology assumed in the analysis. The two dilation parameters, $\alpha_\perp$ and $\alpha_\parallel$, scale separations in the transverse and radial directions respectively. They are closely related to the comoving angular diameter distance, $D_M = (1+z)D_A(z)$, and the Hubble distance, $D_H = c/H(z)$, and defined as:

$$\begin{aligned} \alpha_\perp &= \frac{D_M(z_{\rm eff})/r_d}{D_M^{\rm fid}(z_{\rm eff})/r_d^{\rm fid}}\,, \\ \alpha_\parallel &= \frac{D_H(z_{\rm eff})/r_d}{D_H^{\rm fid}(z_{\rm eff})/r_d^{\rm fid}}\,. \end{aligned} \tag{2.11}$$

Here r_d is the sound horizon, the maximum distance that sound waves could travel prior to the decoupling of baryons and photons after recombination, which sets the scale of the BAO. Given the value of r_d, measurements of the BAO scale from galaxy clustering can be used to provide tight constraints on Ω_m, Ω_Λ, and H_0 through the dependence on D_M and D_H.

r_d can be calculated from Ω_m, Ω_b, Ω_r, and H_0, but it can not be measured directly from spectroscopic galaxy surveys because galaxy clustering does not constrain Ω_b. This is why we measure $\alpha_\perp$ and $\alpha_\parallel$ rather than D_M and D_H directly. However, since Ω_r is well known from the CMB temperature and the effective number of neutrino species (together the dominant radiation-like components of the Universe), r_d can be calibrated given an independent measurement of Ω_b. The two sources typically used are CMB observations or Big Bang Nucleosynthesis (BBN) measurements. BBN refers to the period in the very early Universe, prior to recombination, when the nuclei of elements heavier than hydrogen were first formed. Each element can only be formed once the Universe has cooled to below their respective binding energy, which in turn depends on the baryon-to-photon ratio, η, of the Universe. The present day temperature of the CMB is well known through observation, so η can be easily converted to a constraint on Ω_b. The final abundance of each element is limited by the proton-to-neutron ratio, which can be accurately calculated for any point in the early Universe using particle physics knowledge, so that the element abundances can be predicted by integrating forwards through the reaction network ([266]; or see [187] for a detailed description). The observed primordial abundance of light elements such as ^{4}He, ^{3}He, and, in particular, deuterium [266], can then be used to place strong constraints on

η, and by extension Ω_b [207].

Measuring the BAO scale in galaxy clustering is further complicated by the non-linear evolution of the density field, which shifts the positions of galaxies relative to the initially sharply defined BAO scale of the early Universe, thus smearing out the BAO scale in the clustering data. The velocities produced by the non-linear evolution also impact the positions of the galaxies in redshift space, further smearing the BAO scale and making detection and precise measurement more difficult. The impact on the BAO measurement can be reduced by shifting the position of galaxies to undo the non-linear evolution, a method known as reconstruction [92]. The fiducial BAO measurements from the LRG and ELG samples use clustering from catalogues with reconstruction applied, using the method developed in [48, 47, 24]. Reconstruction is not used in the fiducial analysis of the quasars, because the low density of the sample means the matter field is not sampled with the accuracy required to successfully apply reconstruction.

The basic BAO template is estimated from the linear matter power spectrum, adjusted to match the expectations for the redshift-space clustering of the biased tracers. It is given by:

$$P(k,\mu) = \frac{B[1+\beta(1-S(k))\mu^2]^2}{(1+k^2\mu^2\Sigma_s^2/2)}\left[P_{\text{no peak}}(k) + P_{\text{peak}}(k)e^{-k^2\Sigma_{\text{nl}}^2/2}\right], \tag{2.12}$$

where k is the modulus of the Fourier wave-vector and μ is the cosine of the angle between the wave-vector and the line-of-sight. In the numerator of the expression, B is a free parameter controlling the amplitude of the clustering that is closely linked to the linear bias, b, and $\beta = f/b$ is a factor introduced by the redshift-space distortions. $S(k) = e^{-k^2\Sigma_r^2/2}$ accounts for the smoothing used in the reconstruction technique, where Σ_r is the smoothing radius used. For analyses that do not apply reconstruction, such as the fiducial measurements of the quasar sample, $S(k) = 0$. The denominator includes the contribution from non-linear random motions on small scales, the Finger-of-God effect, modelled by a Lorentzian distribution and parameterized by Σ_s. The final factor is divided into a P_{peak} term and a $P_{\text{no peak}}$ term, which are the BAO peak component and the remaining components of the power spectrum respectively. The reason for splitting those terms is to allow for the non-linear broadening of the BAO peak by a Gaussian distribution, parameterized by $\Sigma_{\text{nl}}^2(\mu) = \Sigma_\parallel^2\mu^2 + \Sigma_\perp^2(1-\mu^2)$. The three parameters controlling the damping terms, $\Sigma_\parallel$, $\Sigma_\perp$, and Σ_s, are held fixed in the analysis, and are set from measurements of the relevant mock catalogues (see Sec. 2.4.1).

To compute P_{peak} and $P_{\text{no peak}}$, we begin by calculating the linear power spectrum, P_{lin}, using the Code for Anisotropies in the Microwave Background (CAMB)[9][176] package. For

[9] camb.info

both ELG analyses and the configuration space analysis of the quasars we calculate $P_{\text{no peak}}$ by applying the fitting formula described in [91] to remove the BAO wiggles from P_{lin}. For the LRG analyses and the Fourier space analysis of the quasars we instead calculate $P_{\text{no peak}}$ using the method established in [156]. In this method, P_{lin} is Fourier transformed to give the correlation function, and then the peak is removed from the correlation function by replacing the peak region with a polynomial function fitted using the regions around the peak ($50 < r < 80\,h^{-1}$Mpc and $160 < r < 190\,h^{-1}$Mpc). The no-peak correlation function is then Fourier transformed back to give $P_{\text{no peak}}$. The peak component is then obtained from $P_{\text{peak}} = P_{\text{lin}} - P_{\text{no peak}}$. The details of the method used in each analysis, as well as comparisons of the two methods and checks for systematic biases, can be found in the relevant analysis papers.

Rather than fitting a model to the two-dimensional power spectrum, $P(k, \mu)$, we compress the information by calculating the multipoles of the power spectrum:

$$P_l(k) = \frac{2l+1}{2} \int_{-1}^{1} P(k, \mu) L_l(\mu) d\mu \,, \tag{2.13}$$

where L_l are the Legendre polynomials. Each analysis also includes a smooth polynomial function of separation that adds several degrees of freedom to the fit. The reason for adding these additional degrees of freedom is to prevent any unknown large-scale systematics, which would introduce additional correlations in the measurements, from biasing the cosmological inference. For example, the factor in the model used to account for the reconstruction is known to be insufficient, and there does not exist an accurate analytic model that is able to capture the full effects of reconstruction on the multipoles, so these additional large-scale degrees of freedom provide the freedom to capture these effects. The final power spectrum multipole templates are therefore:

$$P_l(k) = \frac{2l+1}{2} \int_{-1}^{1} P(k, \mu) L_l(\mu) d\mu + \sum_{i=i_{\text{min}}}^{i_{\text{max}}} a_{l,i} k^i \,, \tag{2.14}$$

where $a_{l,i}$ is the i-th order free parameter in the fit to the l-th multipole. The number of additional free parameters, set by i_{min} and i_{max}, varies between the analyses of the different samples, and is described in the relevant papers.

For the configuration space analysis, the power spectrum multipoles without additional terms (Eq. 2.13) are converted to correlation function multipoles using the Hankel transform, and then the additional terms are added:

$$\xi_l(r) = \frac{i^l}{2\pi^2} \int_0^{\infty} k^2 j_l(kr) P_l(k) dk + \sum_{i=i_{\text{min}}}^{i_{\text{max}}} a_{l,i} r^i \,, \tag{2.15}$$

Paper	Tracer	Type	Measurements	Fit Method
[23]	LRG	Config.	ξ_0, ξ_2	Quasi-Newton
[111]	LRG	Fourier	P_0, P_2	MCMC
[223]	ELG	Config.	ξ_0	Grid search
[76]	ELG	Fourier	P_0	Quasi-Newton
[141]	Quasar	Config.	ξ_0, ξ_2	MCMC
[196]	Quasar	Fourier	P_0, P_2, P_4	Quasi-Newton

Table 2.2: Characteristics of the eBOSS DR16 BAO-only analyses.

where j_l is the l-th spherical Bessel functions.

The only fiducial analysis that does not use a close variation of Eq. 2.12 is the Fourier space measurement of the ELG sample. Instead of using an anisotropic BAO template, which measures the two-point clustering both in terms of the separation of the points and μ, we use an isotropic BAO template developed in [112]. The isotropic model simplifies the template, particularly in that it removes the dependence on μ. The advantages are that it is faster to evaluate and does not need to be integrated along the line-of-sight, and is more suitable for using an analytical solver to find the maximum likelihood. The disadvantage is that the damping of the BAO from non-linear effects is not as accurately described. The isotropic BAO measurement is controlled by a single combined dilation parameter, $\alpha = \alpha_{\parallel}^{1/3}\alpha_{\perp}^{2/3}$. While the configuration space analysis of the ELG sample is based on the anisotropic BAO template in Eq. 2.12, only the monopole (which is insensitive to μ) is included in the fit, so it also produces an isotropic BAO fit.

The key characteristics of the BAO analyses from the different eBOSS samples are displayed in Table 2.2. In each analysis the covariance matrix is estimated from the 1000 EZmocks for the relevant sample (see Sec. 2.4.1). Additional details, such as the parameter values of $[\Sigma_r, \Sigma_{\parallel}, \Sigma_{\perp}, \Sigma_s, i_{\rm min}, i_{\rm max}]$, can be found in the individual analysis papers. The results of each measurement are given in Table 2.3.

2.7.2 Full-shape BAO+RSD measurements

The eBOSS full-shape analyses use a RSD model to predict the relevant clustering statistics from the quasi-linear regime to beyond the BAO scale. Each analysis makes use of one of two RSD models: Convolutional Lagrangian Perturbation Theory (CLPT) with Gaussian Streaming (GS), or the Taruya-Nishimichi-Saito (TNS) model. These models are described in Sec. 1.4.2 and Sec. 1.4.3 respectively. The LRG configuration space analysis makes

Paper	Tracer	(D_M/r_d)	(D_H/r_d)	(D_V/r_d)	χ^2/dof
[23]	LRG C.	17.86 ± 0.33	19.34 ± 0.54	-	$39/(40-9)$
[111]	LRG F.	17.86 ± 0.37	19.30 ± 0.56	-	$108/(112-17)$
[223]	ELG C.	-	-	18.23 ± 0.58	44.4/31
[76]	ELG F.	-	-	18.33 ± 0.60	$42.8/(54-13)$
[141]	Quasar C.	30.82 ± 0.82	13.22 ± 0.56	-	34.1/30
[196]	Quasar F.	30.60 ± 0.90	13.34 ± 0.60	-	$87.63/(126-22)$

Table 2.3: Final constraints of the eBOSS DR16 BAO-only analyses.

Paper	Tracer	Type	Model	Measurements	Fit Method
[23]	LRG	Config	CLPT-GS, TNS	ξ_0, ξ_2, ξ_4	MCMC
[111]	LRG	Fourier	TNS	P_0, P_2, P_4	MCMC
[258]	ELG	Config.	CLPT-GS	ξ_0, ξ_2, ξ_4	Nested Sampler
[76]	ELG	Fourier	TNS	P_0, P_2, P_4	MCMC
[141]	Quasar	Config.	TNS	ξ_0, ξ_2, ξ_4	MCMC
[196]	Quasar	Fourier	TNS	P_0, P_2, P_4	MCMC

Table 2.4: Characteristics of the eBOSS DR16 full-shape BAO+RSD analyses.

measurements using each of these models, and presents a final combined constraint as the best full-shape measurement [23]. It is important to note that the full-shape analyses are performed without reconstruction, and the dilation parameters (Eq. 2.11) are still applied to the model clustering measurements.

The key characteristics of the full-shape BAO+RSD analyses, using the same 1000 EZmock covariance matrices as the BAO-only fits, are given in Table 2.4, while the results are listed in Table 2.5.

2.7.3 Cosmological implications

The BAO and RSD measurements presented in Sec. 2.7.1 and Sec. 2.7.2 provide a wealth of cosmological information. Fig. 2.18 shows the final measurement of each sample, where the configuration-space and Fourier-space measurements have been combined to give a best measurement for each tracer. Also shown are the constraints from previous SDSS measurements, demonstrating the constraining power of 20 years of observations using the SDSS telescope, from redshift $z = 0.07$ galaxies through to $z > 2.1$ Lyα forest quasars. Over this entire range the SDSS samples, including eBOSS, are in agreement with the

Paper	(D_M/r_d)	(D_H/r_d)	$f\sigma_8$	χ^2/dof
[23]	17.42 ± 0.40	20.46 ± 0.70	0.460 ± 0.050	$83.7/(63-6)$, $85.2/(65-7)$
[111]	17.49 ± 0.52	20.18 ± 0.78	0.454 ± 0.046	$77/(78-11)$
[258]	19.9 ± 1.0	19.1 ± 2.0	0.35 ± 0.10	-
[76]	19.17 ± 0.99	20.0 ± 2.3	0.289 ± 0.091	$141/(146-21)$
[141]	30.66 ± 0.88	13.11 ± 0.52	0.439 ± 0.048	-
[196]	30.68 ± 0.90	13.52 ± 0.51	0.476 ± 0.047	-

Table 2.5: Final constraints of the eBOSS DR16 full-shape BAO+RSD analyses. The two χ^2 values for the LRG configuration space measurement are for the fits to the CLPT-GS and TNS models respectively, while the parameter constraints are from the combination of the two measurements. Some analysis papers did not quote a final χ^2 value, so those cells are left blank.

expectation for a ΛCDM universe with the *Planck* 2018 cosmological parameters [213], providing a strong test of the cosmological model.

The measurements from eBOSS and other SDSS samples also place valuable constraints on extensions to the ΛCDM model. Fig. 2.19 shows the constraints on an oΛCDM universe (non-zero curvature) from the combined SDSS BAO measurements [10], the Pantheon type-Ia supernovae, and the *Planck* 2018 temperature and polarization data [213]. The SDSS BAO provides the tightest constraints on Ω_m of the three probes and a competitive constraint on Ω_Λ. The greatest benefit is that the BAO measurements provide a complementary degeneracy direction in these parameters, leading to the very tight constraint on Ω_k from the combined measurements.

As well as curvature, the SDSS BAO measurements provide an important constraint on models with a dark energy equation-of-state $w \neq -1$ (wCDM). Fig. 2.20 shows the constraints from the *Planck* CMB, Pantheon supernovae, and SDSS BAO measurements on w and Ω_m. The BAO measurements again provide a complementary constraint due to the difference in degeneracy between the parameters, leaving only a small portion of parameter-space that is consistent with all three probes.

Finally, the SDSS BAO constraints give new information on tensions between different probes within a ΛCDM framework, particularly the measurements of H_0 between the CMB and the local distance ladder. Fig. 2.21 shows the constraints on H_0 and r_d from the BAO, BAO with sound horizon set by constraints from BBN, the SH0ES distance ladder measurements, and the *Planck* CMB constraints. While the BAO constraints on their own have a strong degeneracy between H_0 and r_d, making them consistent with both the distance ladder and CMB measurements, when combined with BBN constraints on r_d they

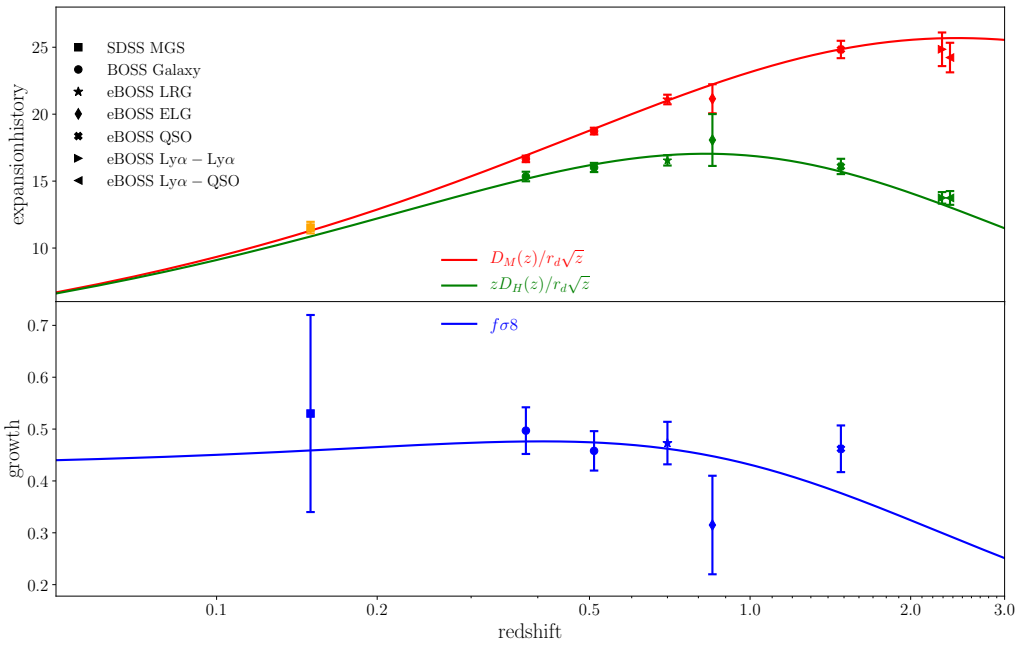

Figure 2.18: BAO and RSD measurements from all SDSS samples. The top panel shows the BAO measurements, with red corresponding to the transverse BAO and green to the radial BAO. The orange point shows the isotropic BAO measurement from the MGS sample, translated to a D_M constraint assuming a ΛCDM universe. The bottom panel shows $f\sigma_8$ constraints. In both cases the solid lines show the expectation for a ΛCDM universe with *Planck* 2018 best fit cosmological parameters [213]. The SDSS constraints come from the SDSS MGS [233, 143], BOSS galaxies [9], eBOSS LRGs [23, 111], eBOSS ELGs [223, 258, 76], eBOSS quasars [141, 196], BOSS+eBOSS Lyα auto-correlation, and BOSS+eBOSS Lyα-quasar cross-correlation [85]. [10]

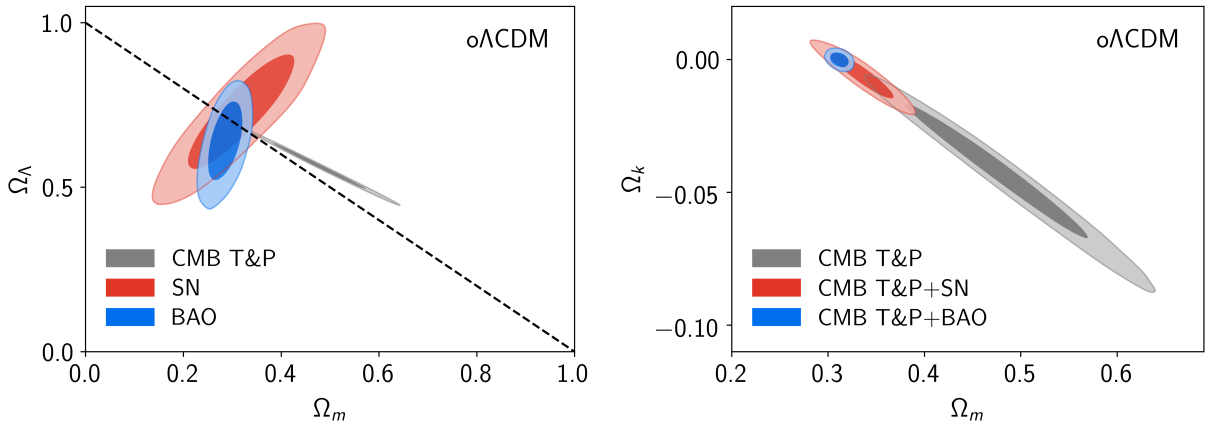

Figure 2.19: Curvature constraints from the combined SDSS BAO measurements (blue), the Pantheon type-Ia supernovae sample (red), and *Planck* temperature and polarization measurements (grey) for a ΛCDM universe with variable curvature (oΛCDM). The left panel shows the constraints on Ω_m and Ω_Λ, with the dark shaded regions showing the 68% credible intervals and the lighter shaded regions show the 95% credible intervals for each measurement. The dashed line shows the values for a flat universe (zero curvature). The right panel shows the combined constraints on Ω_m and Ω_k from *Planck* (grey), *Planck*+Pantheon supernovae (red), and *Planck*+SDSS BAO (blue). [10]

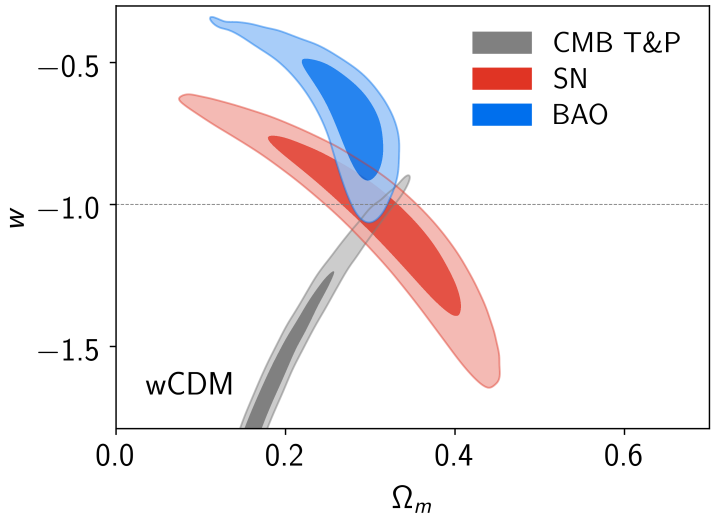

Figure 2.20: Constraints on the dark energy equation-of-state from SDSS BAO (blue), Pantheon supernovae (red), and *Planck* CMB measurements (grey). [10]

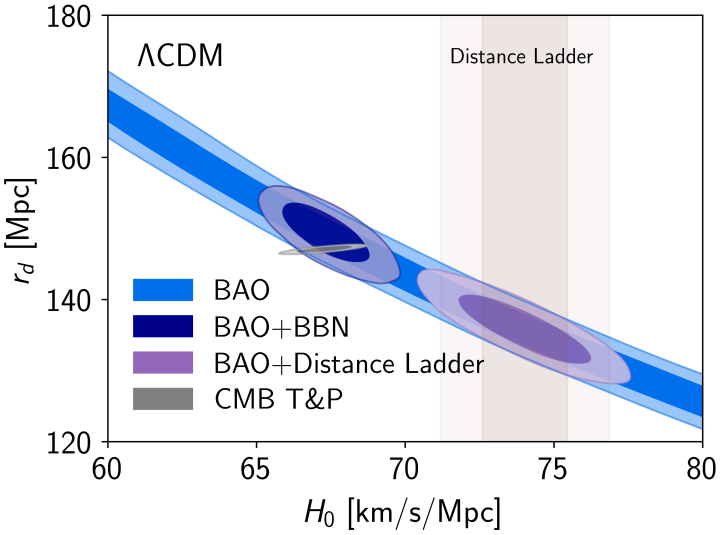

Figure 2.21: Constraints on H_0 and r_d in a ΛCDM model from SDSS BAO (blue), combination BAO and BBN (dark blue), combination BAO and SH0ES distance ladder (purple), and from *Planck* CMB (grey). The shaded bands show the constraint on H_0 from the distance ladder. [10]

are in almost perfect agreement with the CMB measurements, and in significant tension with the distance ladder. This independent constraint provides additional evidence that new physics is required to reconcile these early and late time distance measurements. The BAO also provide a consistency test across a wide range of redshifts, as shown by Fig. 2.18, severely limiting the space of possible models to resolve the tension.

2.8 Conclusions

eBOSS was conceived as a continuation of the successful BOSS program, intended to extend the LRG sample to a higher redshift range, greatly expand the catalogue of high-efficiency Lyα forest quasars, and add new samples of ELGs and clustering quasars. These goals have been met through careful target selection, observation, and data reduction.

A key aspect of undertaking such a large survey is the careful treatment of systematic biases. As part of eBOSS we provide the first unbiased correction to the fibre-collision issue in a fibre-fed spectroscopic survey, through the application of PIP weights with Angular Upweighting. We also investigate a troubling observational systematic in the ELG sample

caused by variable depth in the imaging surveys used to target the ELGs. The impact of this systematic ultimately requires removing significant information from the sample in the cosmological analysis, highlighting the importance of correcting such issues before observation in future surveys.

After carefully validating our large-scale structure catalogues we extract important cosmological constraints through a BAO-only template fitting model and a full-shape BAO+RSD model. These measurements push our constraints on key parameters of the ΛCDM model to a new redshift range, and provide unique and powerful restrictions on extensions to the ΛCDM model, as well as insight into tensions between different cosmological probes.

In addition to these core cosmological analyses, the eBOSS data facilitates many more analyses that extract additional information from the eBOSS large-scale structure catalogues (Ch. 3 and Ch. 4). Together, these analyses show the value of eBOSS both for our current understanding of cosmology, and for our future discoveries.

Chapter 3

Measuring small-scale RSD with eBOSS LRGs

This chapter is adapted from [57]. I was the principal investigator of this work and performed all steps of the data analysis, with the exceptions of building, training, and estimating the uncertainty of the emulator, performed by Prof. Zhongxu Zhai; constructing the SHAM mock, performed by Prof. Jeremy Tinker; and measuring the CMASS+eBOSS correlation functions, performed by Dr. Faizan Mohammad. A central aspect of this work is the PIP+ANG weights produced for eBOSS by Dr. Mohammad in [190], which I applied in the correlation function measurements of this work. I lead the writing of the paper for publication, with input from my co-authors. Specifically, Prof. Will Percival wrote the original Sec. 1: Introduction, and Dr. Mohammad wrote the original Sec. 2: eBOSS LRG Sample and Sec. 3.1 Measurements, all of which were edited by myself and other co-authors. Throughout this chapter I used 'we' to describe all contributions, including my own. The content of this work has been adapted and reformatted by me for this thesis.

3.1 Introduction

The first measurement of the amplitude of the RSD signal from small-scale clustering was made by [226] using an early galaxy sample from the Baryon Oscillation Spectroscopic Survey (BOSS). [226] fit to the monopole and quadrupole moments of the correlation function over scales 0.8 to $32\,h^{-1}$Mpc, obtaining a 2.5% measurement of $f\sigma_8(z = 0.57) = 0.450 \pm 0.011$, compared to the leading $\sim 6\%$ measurement from the large-scale analysis of

the same sample [33, 114, 257, 241]. This demonstrated the increased precision available *if* RSD in the data can be accurately measured and modeled 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 [226]. However, without a simulation for each model to be tested ([226] used three simulation sets at three very similar cosmologies), one has to extrapolate solutions to different cosmologies, which needs care. In addition to fixed cosmological parameters, the [226] analysis also faced the pernicious problem of correcting fibre collisions in the small-scale clustering in the data, which they identified as the leading observational systematic for small-scale clustering. Since the [226] analysis, similar methods have been applied to the BOSS LOWZ galaxies [166], CMASS sample [275], and combined LOWZ+CMASS [277].

RSD depend on the velocity field in the universe along the line of sight. In practice, in order to measure this we need to specify the cosmological frame within which the velocities are defined. Thus the goal of most analyses is to test whether the amplitudes of the velocities match those expected by the cosmological model used to define the frame within which the velocities are measured. For most RSD analyses this is accomplished by setting a constraint on the parameter combination $f\sigma_8$, which can then be compared to the ΛCDM expectation for a given set of cosmological parameters. [226] faced the restriction that their small-scale clustering model was constructed from a single N-body simulation at a fixed cosmology, so they accomplished this goal using a novel technique: a velocity scaling parameter, $\gamma_{\rm HV}$, that multiplied all halo velocities in the simulation. They also assumed that the cosmological model used for the framework did not change the RSD measurement. If the amplitude of the velocity field is directly proportional to $f\sigma_8$, as predicted in linear theory, then $\gamma_{\rm HV}$ is simply a scaling parameter measuring the multiplicative offset between the measured $f\sigma_8$ and that in the framework cosmology. This parameterization has the added benefit that, because a change in $\gamma_{\rm HV}$ is directly proportional to a change in f (in linear theory), the [226] results can be directly compared to other RSD analyses by measuring the derived parameter $f\sigma_8 = \gamma_{\rm HV} f_{\Lambda{\rm CDM}}\sigma_8$, where $f_{\Lambda{\rm CDM}}\sigma_8$ are calculated from the simulation cosmology.

The Gaussian process-based emulator constructed by [278] to match the BOSS LRG samples used the same velocity scaling technique, shifting velocities within the simulations using an extra parameter. Without this scaling, the emulator would force the velocities (and $f\sigma_8$) to match that predicted within the modeled framework cosmology (with $f\sigma_8$ depending on Ω_m and σ_8). Including a velocity scaling parameter allows for finer sampling of the amplitude of the velocity field in the training data of the emulator, which is otherwise restricted by the number of N-body simulations that can be produced with variable values of Ω_m and σ_8. Most importantly, fitting the velocity scaling parameter specifically assesses

deviations in the expected growth rate within a ΛCDM framework, which allows for testing of models that do not follow the same relation between the amplitude of the velocity field and the cosmological parameters as ΛCDM, such as modified gravity models.

In this chapter we perform a measurement of $f\sigma_8$ from the small-scale clustering of the eBOSS LRG sample. Pushing the modelling to include small scales in our analysis is made possible by two key advances in methodology since the [226] analysis. First, we use the AEMULUS emulator [278] to create accurate models of the redshift-space correlation function moments to small scales (see Sec. 3.3.2). To correct for fibre collisions in eBOSS, we use the Pairwise-Inverse-Probability (PIP) method with Angular Upweighting (ANG) [36, 202], as described in Sec. 2.5. 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.

This chapter is structured as follows: the eBOSS LRG sample is described briefly in Section 3.2, and the method for measuring the data and fitting to the model is described in Section 3.3. In Section 3.4 we perform various tests of the method using mock catalogues. We present our results in Section 3.5, and discuss their significance in Section 3.6. Finally, we summarize our results in Section 3.7.

3.2 eBOSS LRG sample

For our analysis we use the eBOSS LRG sample, described in Ch. 2. We apply the PIP+ANG weights described in Sec. 2.5 to correct fibre collisions, allowing us to make unbiased clustering measurements down to very small scales.

It is worth noting that the eBOSS sample of LRGs overlaps in area and redshift range with the high-redshift tail of the BOSS CMASS sample (see Fig. 2.4). Unlike many other eBOSS analyses, including the large-scale measurements of BAO and RSD described in Sec. 2.7 [23, 111], 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,

$$z_{\rm eff} = \frac{\Sigma_{m>n} w_{mn}^{\rm PIP} w_m^{\rm tot} w_n^{\rm tot} (z_m + z_n)/2}{\Sigma_{m>n} w_{mn}^{\rm PIP} w_m^{\rm tot} w_n^{\rm tot}}, \tag{3.1}$$

where the indices m, n are over the objects in the data catalogue, and the description of the weights is given in Sec. 2.5. Additionally, we only include galaxy pairs which have a separation between $0.1 - 60\,h^{-1}$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 $1.28\,\mathrm{Gpc}^3$ [231].

3.3 Methods

3.3.1 Measurements

We measure the monopole, $\xi_0(s)$, and quadrupole, $\xi_2(s)$, of the two-point correlation function, as well as the projected correlation function, $w_p(r_\perp)$, of the eBOSS LRG sample using the Landy-Szalay estimator, as detailed in Sec. 1.3.2. We include $w_p(r_\perp)$ because the halo-occupation properties of a given sample affect its intrinsic clustering, which can be accurately modelled in $w_p(r_\perp)$ free of the apparent RSD effects. We bin $r_\perp$ and s in 9 logarithmically spaced bins between $0.1 - 60\,h^{-1}$Mpc, matching the output of AEMULUS predictions for $w_p(r_\perp)$ and ξ_ℓ, while the line-of-sight separation $r_\parallel$ and μ are binned using linear bins of width $\Delta r_\parallel = 1\,h^{-1}$Mpc and $\Delta\mu = 0.1$. 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 eBOSS DR 16 LRG sample. Given the discrete binning of different variables, we estimate the integrals in Eq. 1.26 and Eq. 1.25 as Riemann sums. We limit the integral in Eq. 1.26 to a maximum line-of-sight separation of $r_{\parallel,\mathrm{max}} = 80\,h^{-1}$Mpc, matching the definition in the model to be fitted to these data [278].

3.3.2 AEMULUS cosmological emulator

In order to access RSD information on small scales we need to model the clustering of galaxies into the non-linear regime. The solution we choose is to use an emulator for the small-scale clustering, trained and validated using N-body simulations. Constructing the emulator applies machine learning with a Gaussian process to emulate the correlation function measurements in each separation bin, as a function of the set of parameters specifying the cosmology and HOD model. First, a set of *training* data is used to specify the value of the emulator at a series of points in parameter space. These are the means of the Gaussian distributions. Next, a different set of *test* data is used to optimise the width and shape of an "interpolation kernel", such that the final model, given a set of model

parameters, is the linear sum of the means coming from the training data, weighted by this kernel. The training data is generated from N-body simulations, where a Halo Occupation Distribution (HOD) is used to connect galaxies to halos. While the training data consists of only a limited number of possible values in our parameter space, the trained emulator is very effective at interpolating within this parameter space to produce accurate clustering measurements.

We compare our measurements to the AEMULUS cosmological emulator [278] 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 training data of the AEMULUS emulator is a set of 40 N-body simulations that use a latin hypercube to optimally sample a flat wCDM parameter space (i.e. a cosmological model equivalent to a ΛCDM model, but where the dark energy equation-of-state parameter is constant but not restricted to $w = -1$) spanning the approximate 4σ range of the *Planck* [213] or WMAP [138] results [77]. Details of the kernel and optimisation are available in [278]. 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, [166] found that excluding ξ_4 from their analysis of ξ_0 and ξ_2 did not produce a significant change in best fit value or uncertainty.

AEMULUS allows for a flat wCDM universe described by 7 parameters: Ω_M, Ω_b, σ_8, h, n_s, w, and $N_{\rm eff}$. For our analysis we limit the cosmological parameter space by fixing $N_{\rm eff} = 3.046$ and $w = -1$, since these parameters are not well constrained by our measurements but have been well measured by other probes, resulting in a 5 parameter flat ΛCDM cosmology.

The HOD model used by AEMULUS allocates a Poisson sampling of $N(M)$ galaxies to halos of mass M, split into central galaxies and satellites following the model detailed in Sec. 1.5.3. The free parameters of this HOD model are $f_{\rm max}$, $\sigma_{\log M}$, $\log M_{\rm sat}$, α, $\log M_{\rm cut}$. $\log M_{\rm min}$ is an additional parameter in the model that 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. $f_{\rm max}$ 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 (see Sec. 2.2.2) 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 [279]). $f_{\rm max}$ controls the fraction of centrals that are included in the sample, i.e. a value of $f_{\rm max} < 1$ means that the very massive halos do not necessarily host an eBOSS LRG at the center. While these targeting cuts would be expected to affect the completeness of both centrals and satellites, for satellites $f_{\rm max}$ is completely degenerate with $M_{\rm sat}^{-\alpha}$ (see similar discussion in [166]). Since these HOD parameters are primarily nuisance parameters in our constraint of the growth rate, we do not apply $f_{\rm max}$ to the satellites. In Sec. 3.4.2 we perform a series of tests to determine the effect of excluding $f_{\rm max}$ on the measured $f\sigma_8$.

The emulator also allows three additional parameters that control how galaxies are distributed in their host halos: $c_{\rm vir}$, $v_{\rm bc}$, and $v_{\rm bs}$ (labelled $\eta_{\rm con}$, $\eta_{\rm vc}$, and $\eta_{\rm vs}$ in [278]). $c_{\rm vir} = c_{\rm sat}/c_{\rm halo}$ is the ratio of the concentration of the satellite distribution, $c_{\rm sat}$, to the concentration of the host halo, $c_{\rm halo}$, assuming the halo follows a NFW profile [194]. $v_{\rm bc}$ and $v_{\rm bs}$ are the velocity biases of centrals and satellites respectively, where $\sigma_{\rm gal} = v_{\rm gal}\sigma_{\rm halo}$ and $\sigma_{\rm halo}$ 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. Throughout this chapter we define 'bulk velocities' to mean the velocity of the halo as a single unit, rather than the velocity of the individual particles making up the halo or the internal velocity dispersion of the halo. 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 $v_{\rm bc}$ or $v_{\rm bs}$ for centrals and satellites respectively. For a detailed description of the AEMULUS correlation function parameters see [278]. See Sec. 3.3.6 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 $n = 4.2 \times 10^{-4}[h^{-1}{\rm Mpc}]^{-3}$. 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 halos. The result will be a sample with a larger linear bias, which is degenerate with the growth rate in clustering measurements. The effect of these differences, as well as methods to correct them, are investigated in Appendix B. In order to ensure an unbiased result for our analysis we decided to rebuild the emulator from the original AEMULUS simulations, but using the $z = 0.7$ simulation time-slice and adjusting HOD parameters, especially $M_{\rm min}$, to match the eBOSS number density. The training ranges for the new emulator are given in Table 3.1.

3.3.3 Interpreting growth rate measurements

As shown in [226], which used a similar parameterization 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 $f = \gamma_f f_{\Lambda\mathrm{CDM}}$, where $f_{\Lambda\mathrm{CDM}}$ 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, so γ_f scales both the linear and non-linear velocities of halos. 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 Sec. 3.4.1).

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 $f\sigma_8 = \gamma_f f_{\Lambda\mathrm{CDM}}\sigma_8$, which is used in the remainder of this chapter and other sections where these results are referenced. It is also important to note that we calculate $f_{\Lambda\mathrm{CDM}}\sigma_8$ from the model cosmology according to linear theory, rather than the value that would be obtained from the power spectrum on scales corresponding to $0.1-60\,h^{-1}$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 (ΛCDM+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.

3.3.4 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 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. 3.1). 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

$$C_{i,j} = \frac{n-1}{n}\sum_{k}^{n}(\xi_{i,k} - \bar{\xi}_i)(\xi_{j,k} - \bar{\xi}_j), \tag{3.2}$$

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 realisations.

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

$$R_{i,j} = \frac{C_{i,j}}{(C_{i,i}C_{j,j})^{1/2}}\,. \tag{3.3}$$

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

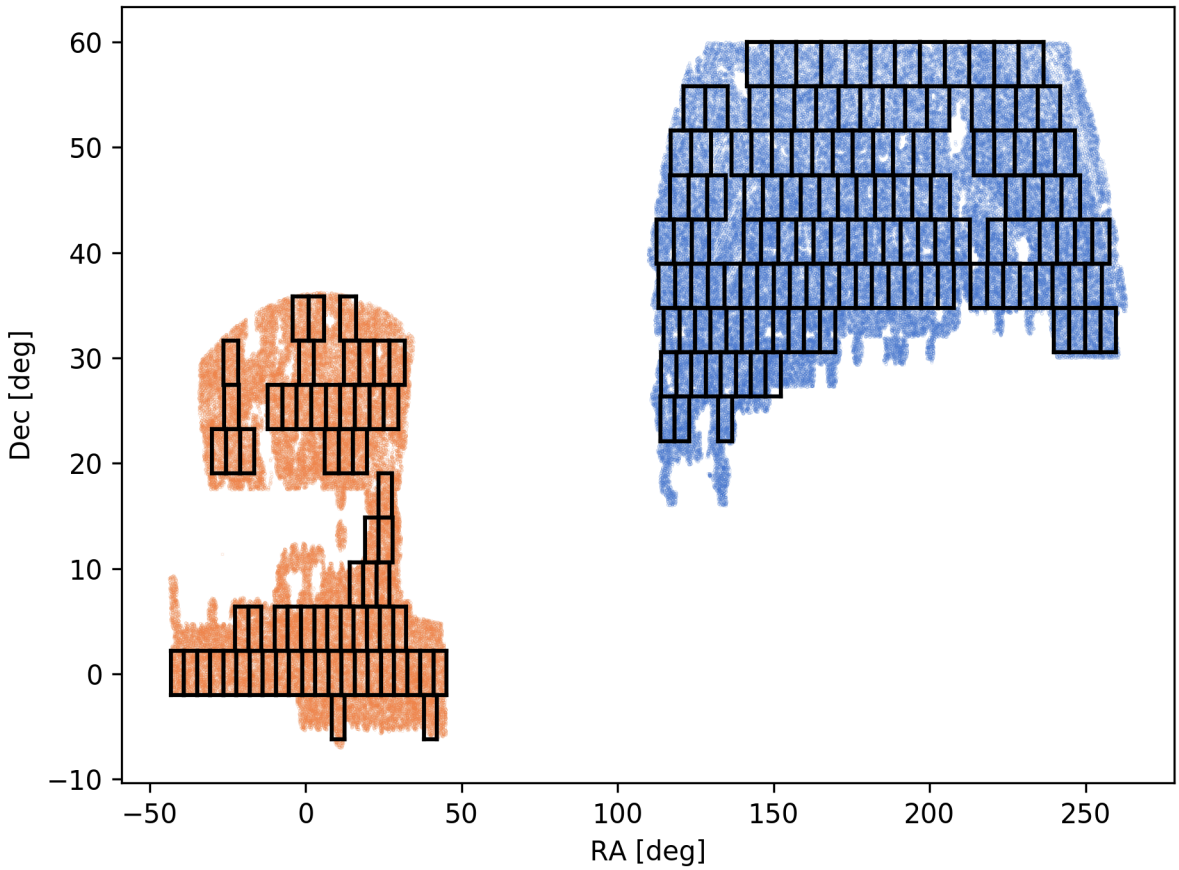

Figure 3.1: The footprint of the eBOSS LRG clustering catalogue with our jackknife regions. The blue points show the North Galactic Cap (NGC) observations, while the orange points show the South Galactic Cap (SGC) observations. It should be noted that the square jackknife regions all have approximately equal area on the sky, however due to the distortion of projecting a sphere onto a plane the regions at larger declination appear wider in this plot.

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. 3.2.

We test the robustness of our jackknife covariance matrix by comparing to a covariance matrix estimated from 1000 EZmocks constructed to match the eBOSS LRG sample (see Sec. 2.4.1). The variance from the jackknife on the data and 1000 EZmocks are compared in Fig. 3.3. Because the correlation function measured for the EZmocks is significantly smaller than the correlation function of the data on small scales we compare the ratio of the square root of the variance to the correlation function measurement in each separation bin. The correlation function measurements from the data are used in the denominators for the jackknife on the data, and the mean of the EZmocks measurement is used for the denominator of the 1000 EZmocks. We find that both methods are in agreement at all scales for the three measurements used in the analysis. In Fig. 3.3 there is an apparent deviation for the small scales of the quadrupole, but this is actually caused by the mean quadrupole measurement of the 1000 EZmocks. Individual quadrupole measurements have a large variance on these scales but a mean that is close to 0, increasing the ratio plotted here for the 1000 EZmocks. For comparison we have applied our jackknife method to 10 EZmocks, also shown in Fig. 3.3. These measurements are consistent with the jackknife on the data, demonstrating that both covariance estimation methods are in agreement.

We also examine the correlation between bins in both methods by comparing their correlation matrices (Fig. 3.4). In general, the two correlation matrices agree closely after diagonal smoothing has been applied to the jackknife correlation matrix, although the jackknife correlation matrix shows an additional feature compared to the mock correlation matrix where the small-scale ξ_0 and ξ_2 are correlated in the same separation bin. This additional feature, along with the difference in clustering, highlights the importance of using jackknife to accurately measure the small-scale covariance.

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 that are selected from the same parameter ranges as the training sample, but were not used in the training [278]. 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. 3.5. The data error is dominant in the region $s < 5\,h^{-1}$Mpc for the monopole and projected correlation function,

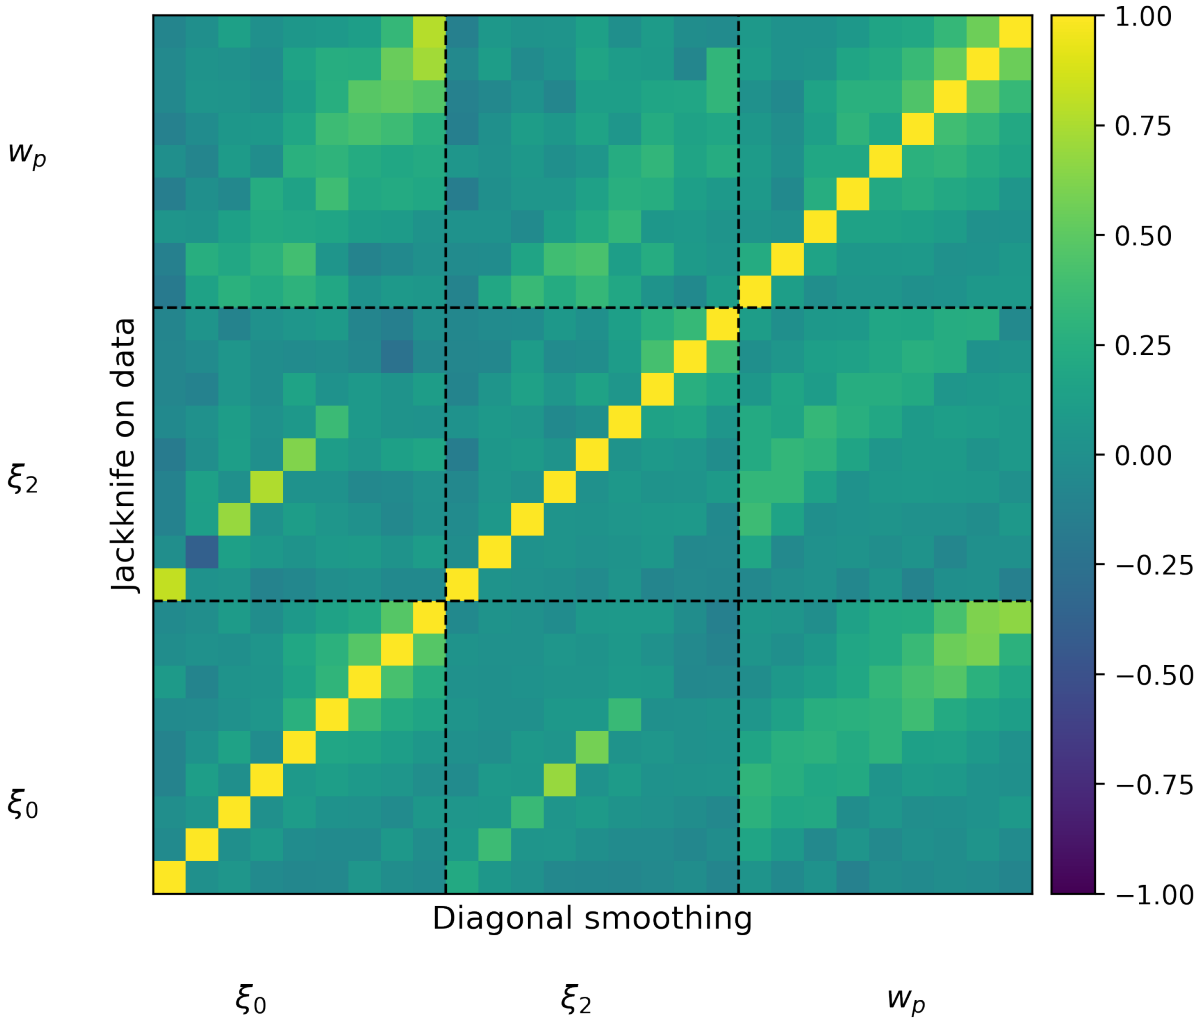

Figure 3.2: Comparison of the unsmoothed and smoothed correlation matrices. The upper diagonal elements correspond to the unsmoothed jackknife correlation matrix, while the lower diagonal elements show the result of our diagonal smoothing method.

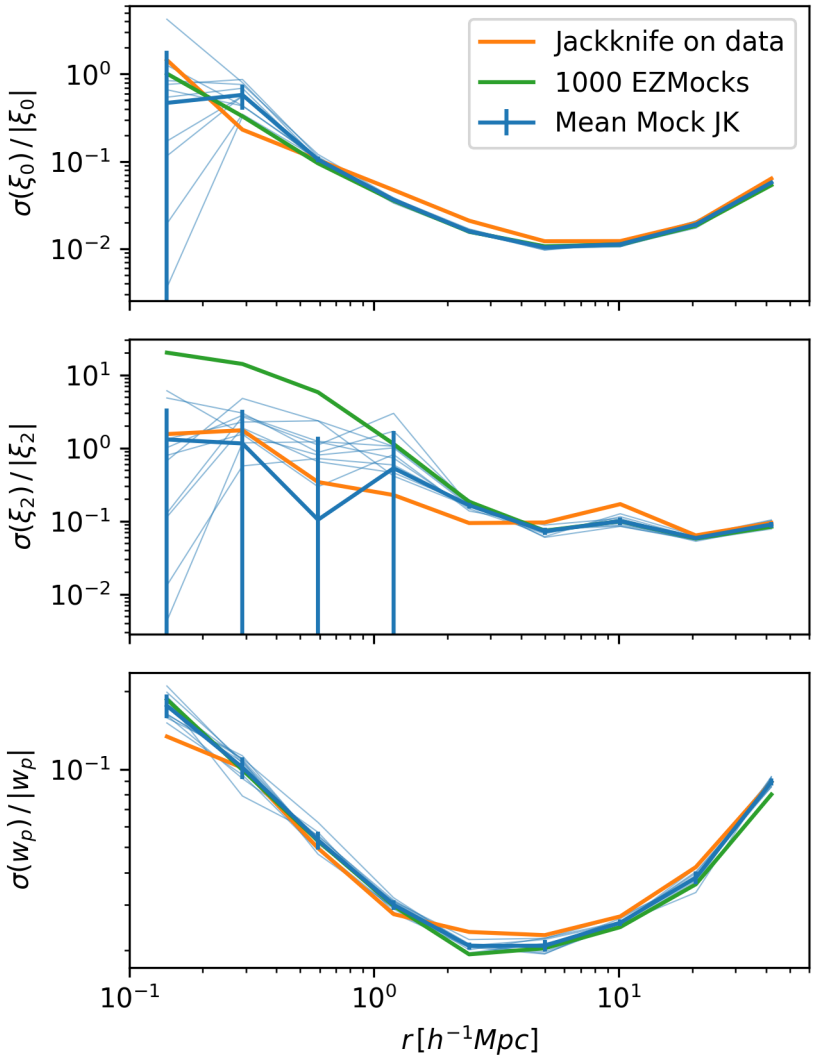

Figure 3.3: Comparison of the variances estimated from different covariance matrix methods. In each case the y-axis shows the square root of the variance divided by the correlation function measurement for that separation bin. The orange line shows the result from the jackknife covariance matrix, while the green line shows the result from estimating the covariance using 1000 EZmocks. The thin blue lines show the result of applying the jackknife method to 10 randomly selected EZmocks. The thick blue line shows the mean of the 10 thin blue lines, with error bars showing their variance. Upper, middle, and lower panels show the monopole, quadrupole, and projected correlation functions respectively.

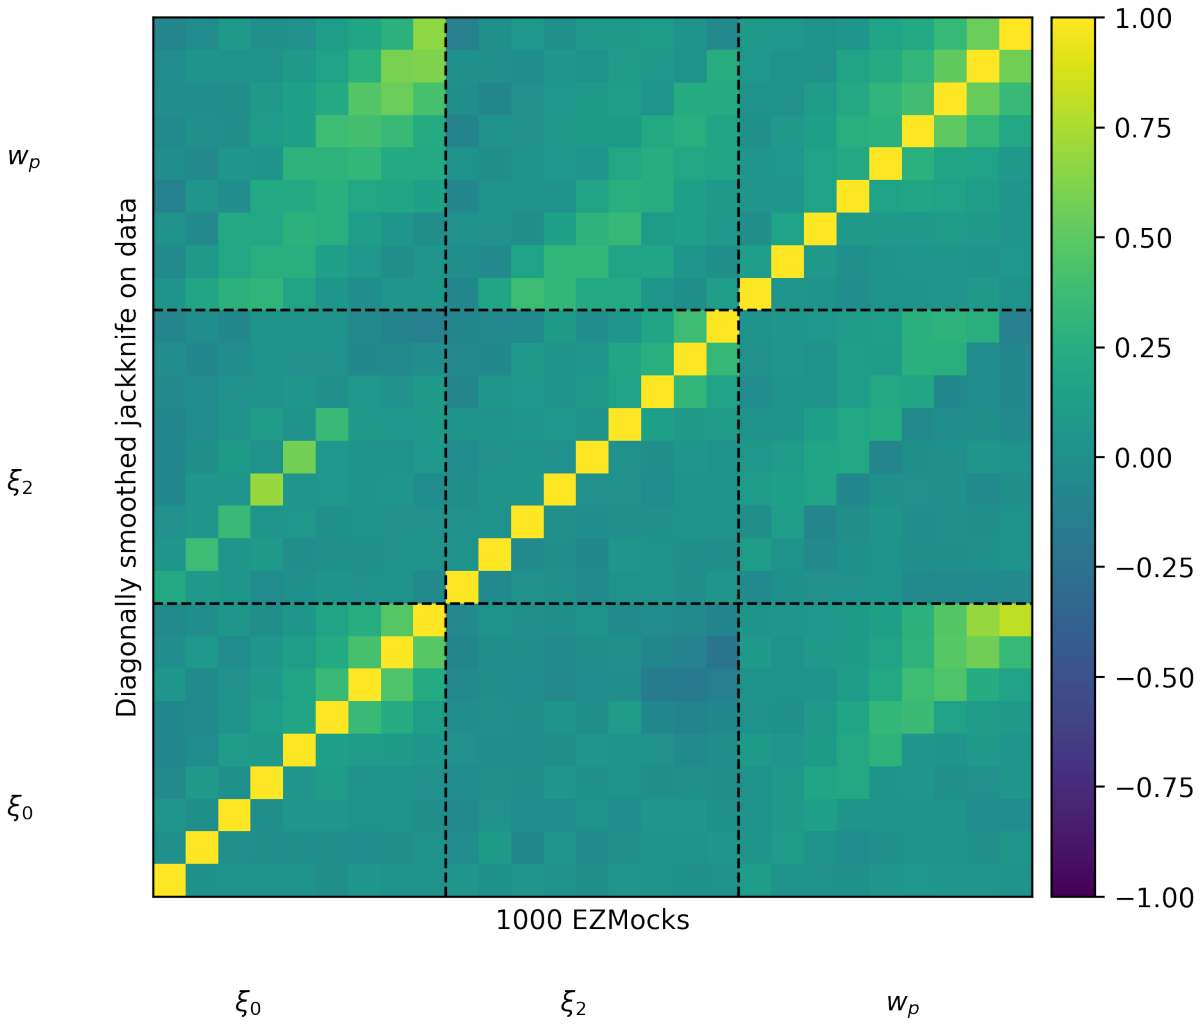

Figure 3.4: Comparison of correlation matrices estimated from jackknife on the data and 1000 EZmocks. The upper diagonal elements correspond to the correlation matrix estimated from applying jackknife to the data with diagonal smoothing, while the lower diagonal elements show the correlation matrix estimated from the 1000 EZmocks.

while the emulator error is comparable to the data error for $s > 5\,h^{-1}$Mpc and across the full separation range of the quadrupole.

We also correct the inverse covariance matrix to account for the finite number of samples used to estimate the covariance using [130]

$$\hat{\boldsymbol{C}}^{-1} = \frac{n-p-2}{n-1}\boldsymbol{C}^{-1}, \tag{3.4}$$

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 measurements [162, 88], we follow [226] 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 fit value.

3.3.5 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 [13], because we convert the data redshift to distance assuming a fixed fiducial cosmological model. The AP effect is a geometrical distortion in the 3D clustering caused by converting redshifts to distances using a cosmological model that differs from the true 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 [23] to each model (similar to the scaling parameters used for the BAO-only template fit in Sec. 2.7.1), first defining the perpendicular and parallel dilation factors

$$\alpha_\perp = \frac{D_M(z_{\rm eff})}{D_M^{\rm fid}(z_{\rm eff})}\,, \quad \alpha_\parallel = \frac{D_H(z_{\rm eff})}{D_H^{\rm fid}(z_{\rm eff})}\,, \tag{3.5}$$

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

$$\xi_0^{\rm fid}(r^{\rm fid}) = \xi_0(\alpha r) + \frac{2}{5}\epsilon\left[3\xi_2(\alpha r) + \frac{d\xi_2(\alpha r)}{d\ln(r)}\right], \tag{3.6}$$

$$\xi_2^{\rm fid}(r^{\rm fid}) = (1+\frac{6}{7}\epsilon)\xi_2(\alpha r) + 2\epsilon\frac{d\xi_0(\alpha r)}{d\ln(r)} + \frac{4}{7}\epsilon\frac{d\xi_2(\alpha r)}{d\ln(r)}. \tag{3.7}$$

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.

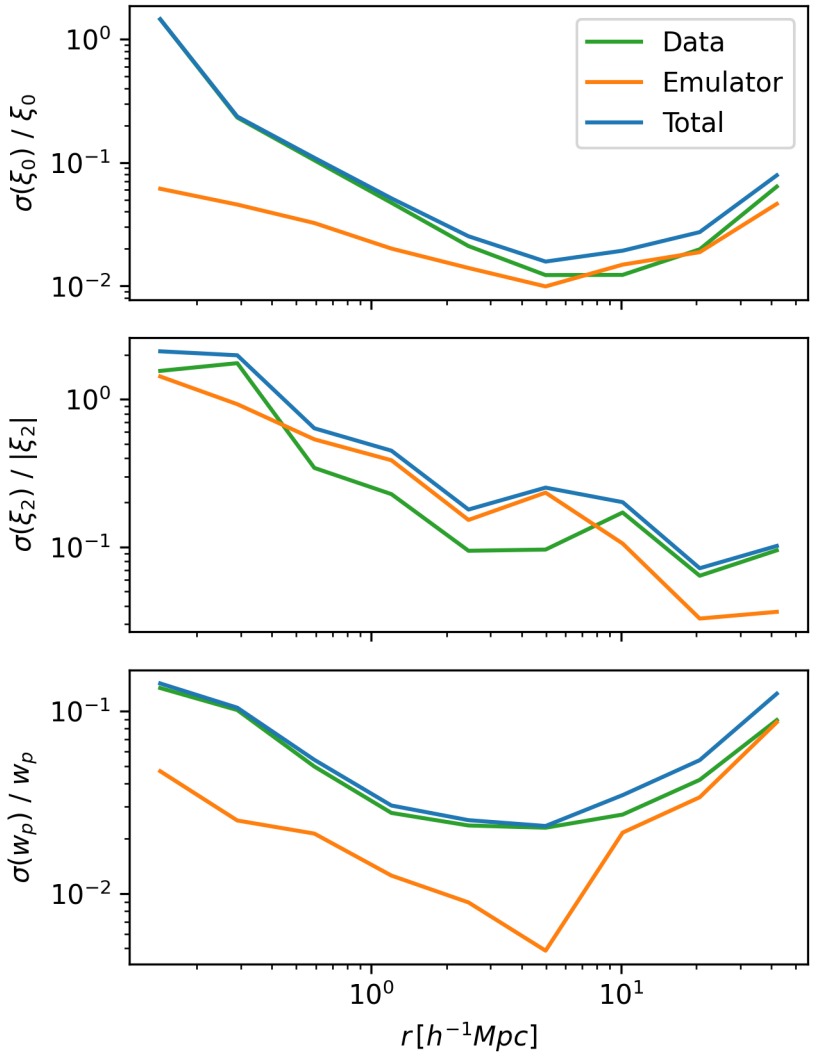

Figure 3.5: The contributions of the data error calculated through jackknife sampling (green), the emulator error (orange), and total error (blue), for the monopole, quadrupole, and projected correlation function (top to bottom).

The projected correlation function was calculated similarly using the scaling

$$w_p^{\rm fid}(r_\perp^{\rm fid}) = w_p(\alpha_\perp r_\perp)\,. \tag{3.8}$$

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 Sec. 3.5.6).

3.3.6 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 3.1. 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 Cobaya [1] framework [263]. We include the full AEMULUS HOD parameter space in our fit, however we limit the wCDM cosmological parameter space by fixing $N_{\rm eff} = 3.046$ 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, in order to set an accurate cosmology for our analyses we consider a number of cosmological priors beyond the simplest uniform priors given in Table 3.1. 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 3.1. 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.

[1] Cobaya, a **co**de for **bay**esian **a**nalysis in cosmology, is the Python successor to CosmoMC. Users are able to use the same MCMC sampler as CosmoMC [174, 173] 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 [195]. See https://cobaya.readthedocs.io for details.

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]
$N_{\rm eff}$	[2.62, 4.28]	3.046
w	[-1.40, -0.57]	-1
$\log M_{\rm sat}$	[14.0, 16.0]	[13.8, 16.2]
α	[0.2, 2.0]	[0.1, 2.2]
$\log M_{\rm cut}$	[10.0, 13.7]	[11.5, 14]
$\sigma_{\log M}$	[0.1, 1.6]	[0.08, 1.7]
$v_{\rm bc}$	[0, 0.7]	[0, 0.85]
$v_{\rm bs}$	[0.2, 2.0]	[0.1, 2.2]
$c_{\rm vir}$	[0.2, 2.0]	[0.1, 2.2]
γ_f	[0.5, 1.5]	[0.25, 1.75]
$f_{\rm max}$	[0.1, 1]	[0.1, 1]

Table 3.1: 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 $\log M_{\rm cut}$, where the prior excludes the lower part of training range since $\log M_{\rm cut}$ ceases to have any impact on the halo occupation if it is below $\log M_{\rm min}$. This is the case for the eBOSS LRG sample, so $\log M_{\rm cut}$ 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 $\log M_{\rm cut}$ for our sample.

We also consider jointly fitting our data with the *Planck* 2018 TT,TE,EE and lensing likelihoods [212, 214] using the CAMB cosmological Boltzmann code [176, 142], 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, are 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. Therefore, we adjust the value of σ_8 given to AEMULUS as follows

$$\sigma_{8,Aem} = \sigma_8(z = 0) \times \frac{D(z = 0.737)}{D(z = 0)} \times \frac{D(z = 0)}{D(z = 0.7)} \,. \tag{3.9}$$

This makes sure that the normalisation of the AEMULUS output matches that expected at $z = 0.737$ in the cosmology being tested: the first ratio corrects from $z = 0$ to $z = 0.737$ in the cosmology being tested, and the second ratio corrects from $z = 0.7$ to $z = 0$, where the normalisation 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 Sec. 3.5.3. Based on the results of these tests we use the training prior but not the *Planck* likelihoods for all measurements, including our headline results, unless otherwise specified.

3.4 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 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. Finally, we perform a general check of our method by fitting to measurements made on a mock catalogue.

3.4.1 Contribution of non-linear velocities

In Sec. 3.3.3 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. 3.6. 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 halos 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 9 measurement bins into three groups of three bins, with individual ranges of $0.1 - 0.8\,h^{-1}$Mpc, $0.8 - 7\,h^{-1}$Mpc, and $7 - 60\,h^{-1}$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 parameterized by our emulator model is in disagreement with the expectation from ΛCDM.

3.4.2 Galaxy selection and the HOD model

As described in Sec. 3.3.2, we add an additional parameter $f_{\rm max}$ to the emulator 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 Uchuu[2] simulation. Briefly, Uchuu is a $(2000\,h^{-1}\text{Mpc})^3$, 12800^3 particle simulation using the *Planck* 2015 cosmology

[2]http://skiesanduniverses.org/Simulations/Uchuu/

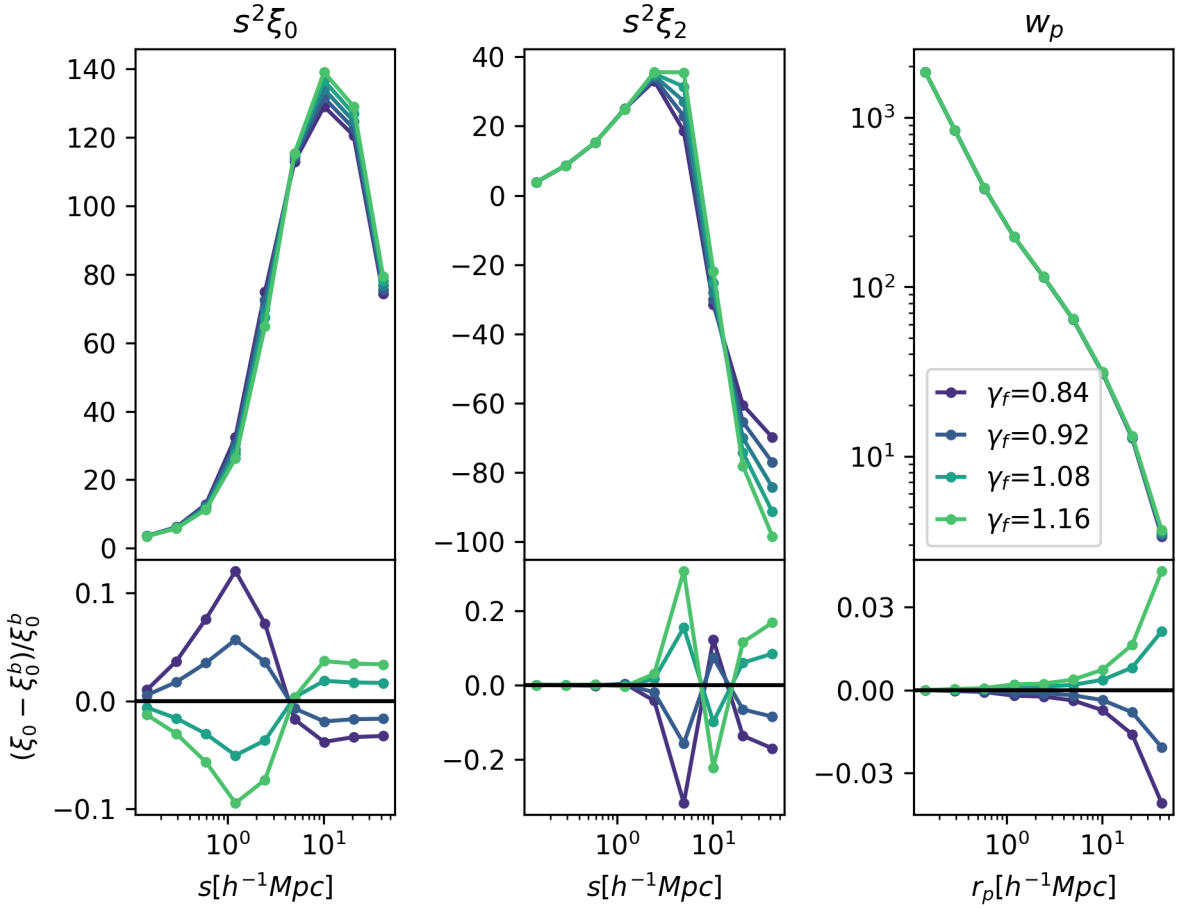

Figure 3.6: The effect on the emulator prediction of varying γ_f for the monopole (*left*), quadrupole (*centre*), and projected correlation function (*right*). All other parameters are kept fixed at reasonable values for the baseline eBOSS fit. *Upper panels:* Direct comparison of the predictions, ranging from low γ_f (purple) to high γ_f (green). *Lower panels:* Relative difference to the $\gamma_f = 1$ prediction.

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 `halotools`[3] [131] Python package and a HOD parameterization identical to that outlined in Sec. 1.5.3 and Sec. 3.3.2. We constructed mocks using $\sigma_{\log M}$, $\log M_{\rm sat}$, α, and $\log M_{\rm cut}$ from five randomly selected test HOD parameter sets in AEMULUS, with $\log M_{\rm min}$ 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 3.1, but were not used in training the emulator. In all mocks we kept the additional parameters $v_{\rm bc} = 0$, $v_{\rm bs} = 1$, $c_{\rm vir} = 1$, 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 $f_{\rm max} = [0.2, 0.4, 0.6, 0.8, 1.0]$, 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 $f_{\rm max} = 1$, and the full emulator with variable $f_{\rm max}$. Both emulators were built to match the eBOSS redshift and number density, as described in Sec. 3.3.2. The γ_f constraints on the HOD mocks from both emulators are shown in Fig. 3.7, 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 catalog 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 $f_{\rm max}$ emulator is able to recover the expected value of γ_f within the uncertainty over the full $f_{\rm max}$ range, and shows no trend in $f_{\rm max}$. Conversely, the fixed $f_{\rm max}$ emulator shows a clear bias in γ_f for $f_{\rm max} \leq 0.6$. This result matches what we would theoretically expect for a model that overestimates the $f_{\rm max}$ value of the sample. If the mismatch is small there is not a significant change in the galaxy bias of the sample, however if $f_{\rm max}$ is significantly overestimated then the model prediction has a larger galaxy bias, b, than the sample. In the fit this is compensated by a lower value of γ_f, since the amplitudes of the correlation function multipoles on linear scales are proportional to both the galaxy bias, b (controlled by $f_{\rm max}$), and the logarithmic growth rate, f (controlled by γ_f).

3.4.3 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 [231], the eBOSS LRG sample has a redshift uncertainty that is well approximated by a Gaussian with mean $\mu = 1.3$ km s^{-1} and standard deviation $\sigma = 91.8$ km s^{-1}. On average, this means that each

[3] https://halotools.readthedocs.io/en/latest/

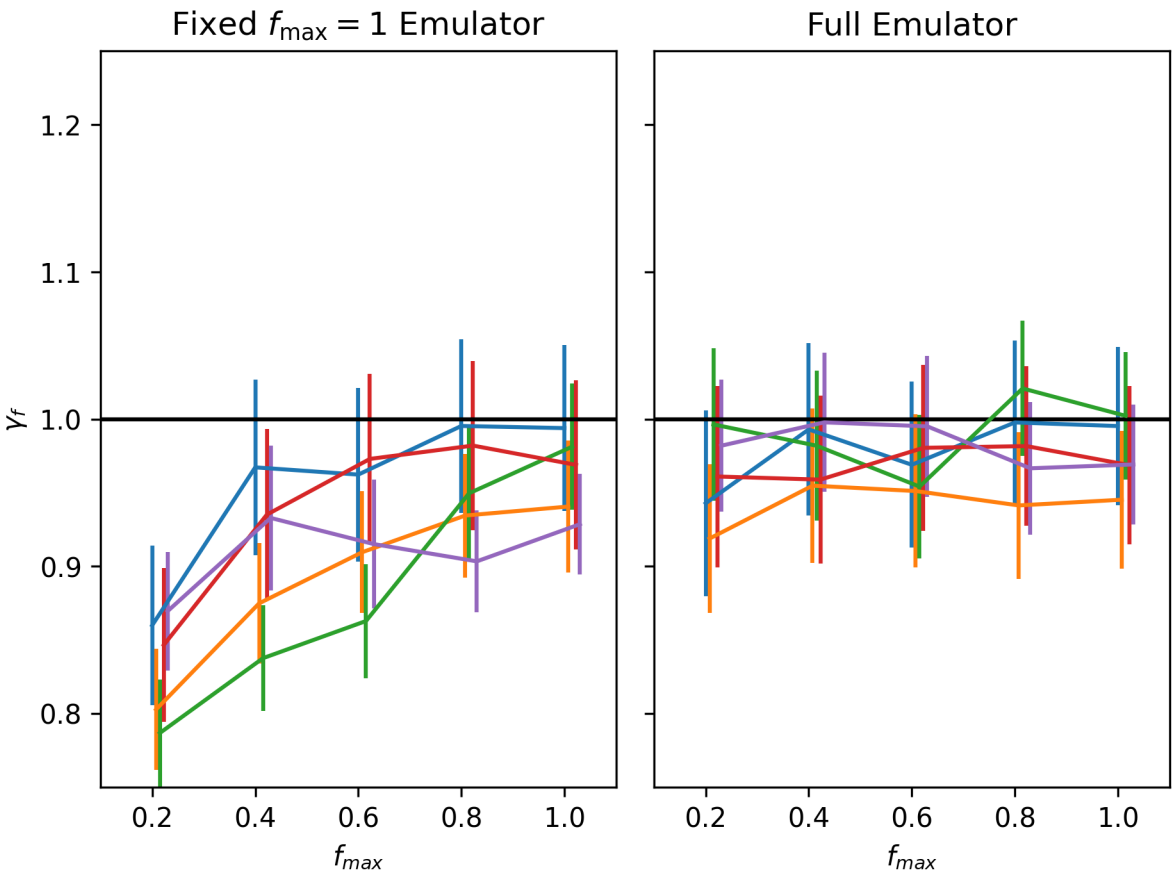

Figure 3.7: Performance of emulators with fixed or variable $f_{\rm max}$ on HOD mocks constructed with varying $f_{\rm max}$. The left panel shows the results from an emulator built with the original AEMULUS parameter set, which is equivalent to $f_{\rm max} = 1$. The right panel shows the results from the emulator used in our analysis with variable $f_{\rm max}$. 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.

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 $v_{\rm bc}$ and $v_{\rm bs}$, 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 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, $v_{\rm bc}$, and $v_{\rm bs}$. In addition, $v_{\rm bc}$ and $v_{\rm bs}$ 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 Sec. 3.4.2. 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 line of sight drawn from a Gaussian with mean $\mu = 1.3$ km s^{-1} and standard deviation $\sigma = 91.8$ 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. 3.8. 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 $v_{\rm bc}$ and $v_{\rm bs}$. 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 [42] 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

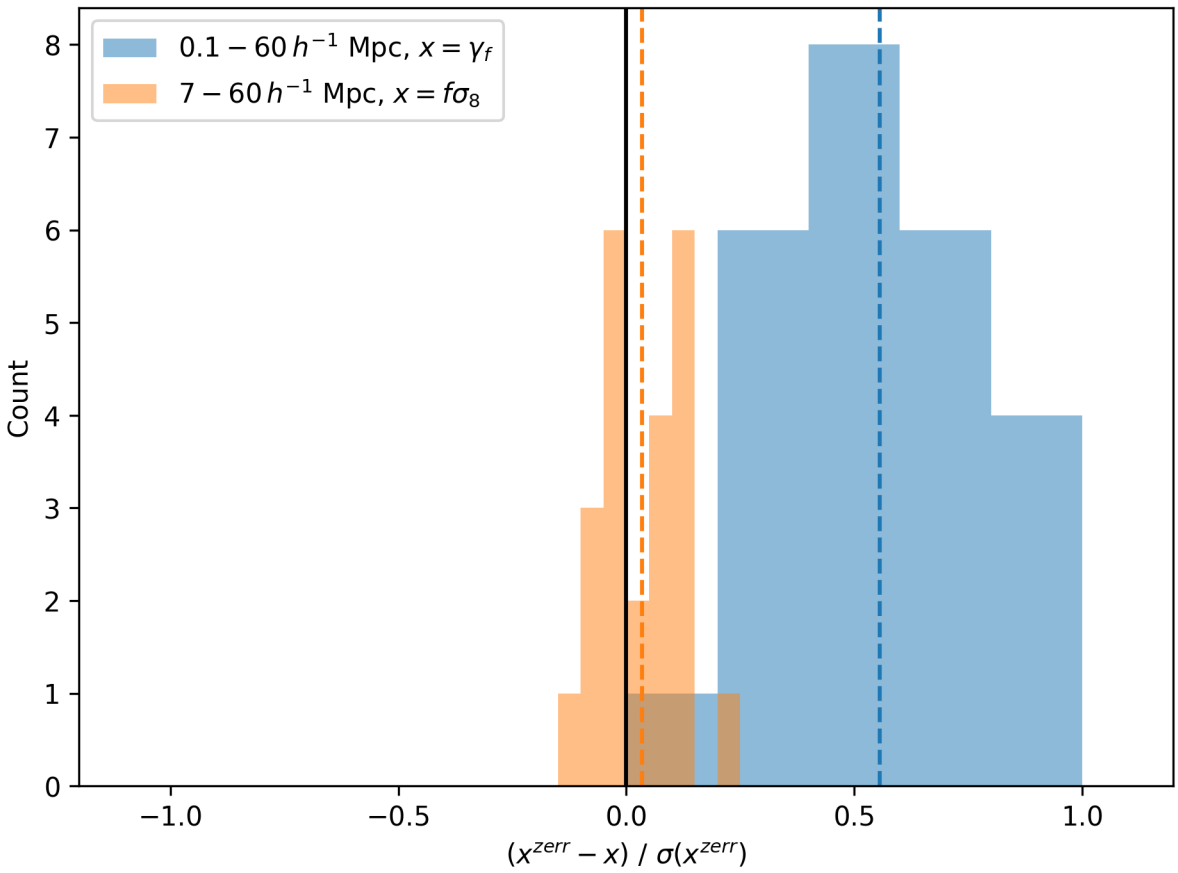

Figure 3.8: A histogram of the shifts in the measured cosmological parameters for 25 HOD mocks with and without a random velocity dispersion matching the eBOSS redshift uncertainty. Blue bars show the shift in γ_f measured over the full separation range, and orange bars show the shift in $f\sigma_8$ 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 $0.1 - 60\,h^{-1}$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 $7 - 60\,h^{-1}$Mpc, the shifts from including a random velocity dispersion scatter around 0, with a mean shift that is negligible compared to the statistical error.

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 halos provided by a larger value of γ_f. In Sec. 3.5 we consistently measure values of γ_f that are below the value expected from ΛCDM+*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.

3.4.4 SHAM mock

We test the robustness of our model and analysis pipeline using a 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 used in the construction of the emulator. We use the $z = 0.7$ slice of the simulation to construct a SHAM mock using the peak halo maximum circular velocity, $V_{\rm peak}$, 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 over the full range of the emulator ($0.1 - 60\,h^{-1}$Mpc) is shown in Fig. 3.9. 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 1σ 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.

3.5 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.

3.5.1 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 Sec. 3.3.6, but do not include any external data. We obtain a value of $\gamma_f = 0.767 \pm 0.052$, 4.5σ below the expectation of $\gamma_f = 1$ in a ΛCDM+GR universe. The 1D and 2D likelihood contours of the full parameter set are shown in Fig. 3.10. All well constrained parameters are within the prior ranges described in Table 3.1, and the parameters that are most impactful for our results, Ω_m, σ_8, $v_{\rm bc}$, and γ_f, all show roughly Gaussian constraints. The best-fit values of the cosmological parameters other than γ_f are consistent with recent measurements from the *Planck* Collaboration [213]. The best fit model prediction is plotted relative to the data in Fig. 3.11, showing reasonable agreement within the measurement uncertainty on all scales. The best fit 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 Sec. 3.4.1, from which we obtain a value of $f\sigma_8(z = 0.737) = 0.408 \pm 0.038$ (calculated from $f\sigma_8 = \gamma_f f_{\Lambda{\rm CDM}}\sigma_8$ as specified in Sec. 3.3.3). This value is 1.4σ below what is expected from the 2018 *Planck* data for a flat ΛCDM 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 Sec. 3.5.4 for more details.

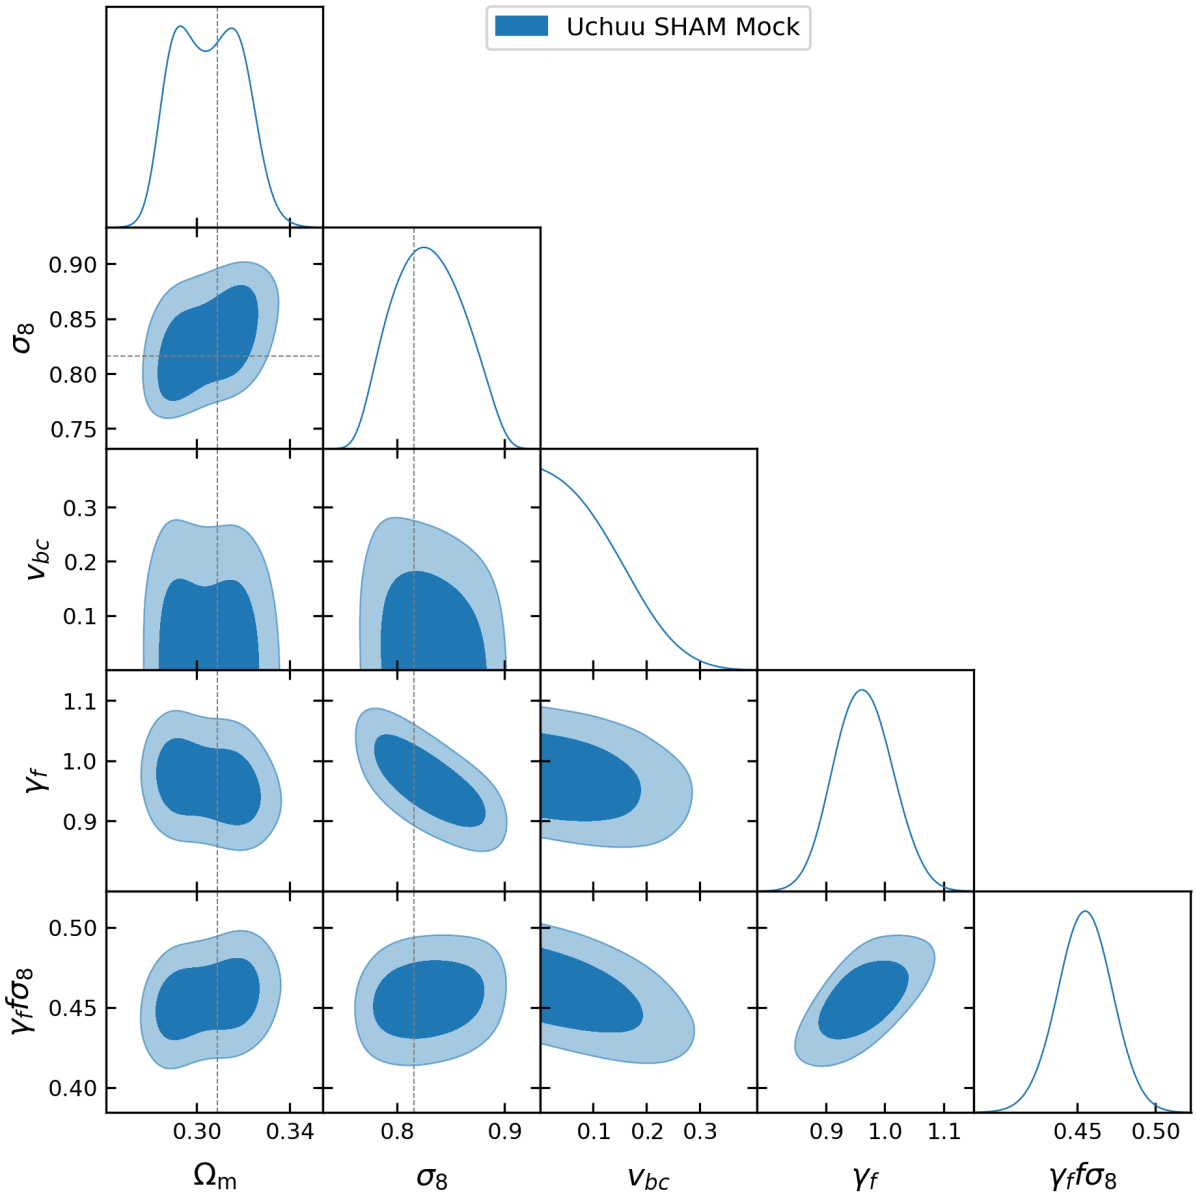

Figure 3.9: 2D and 1D marginalized constraints of the key parameters from the fit to an Uchuu SHAM mock matching the eBOSS LRG number density and redshift. Dotted lines show the values of the cosmological parameters from the simulation.

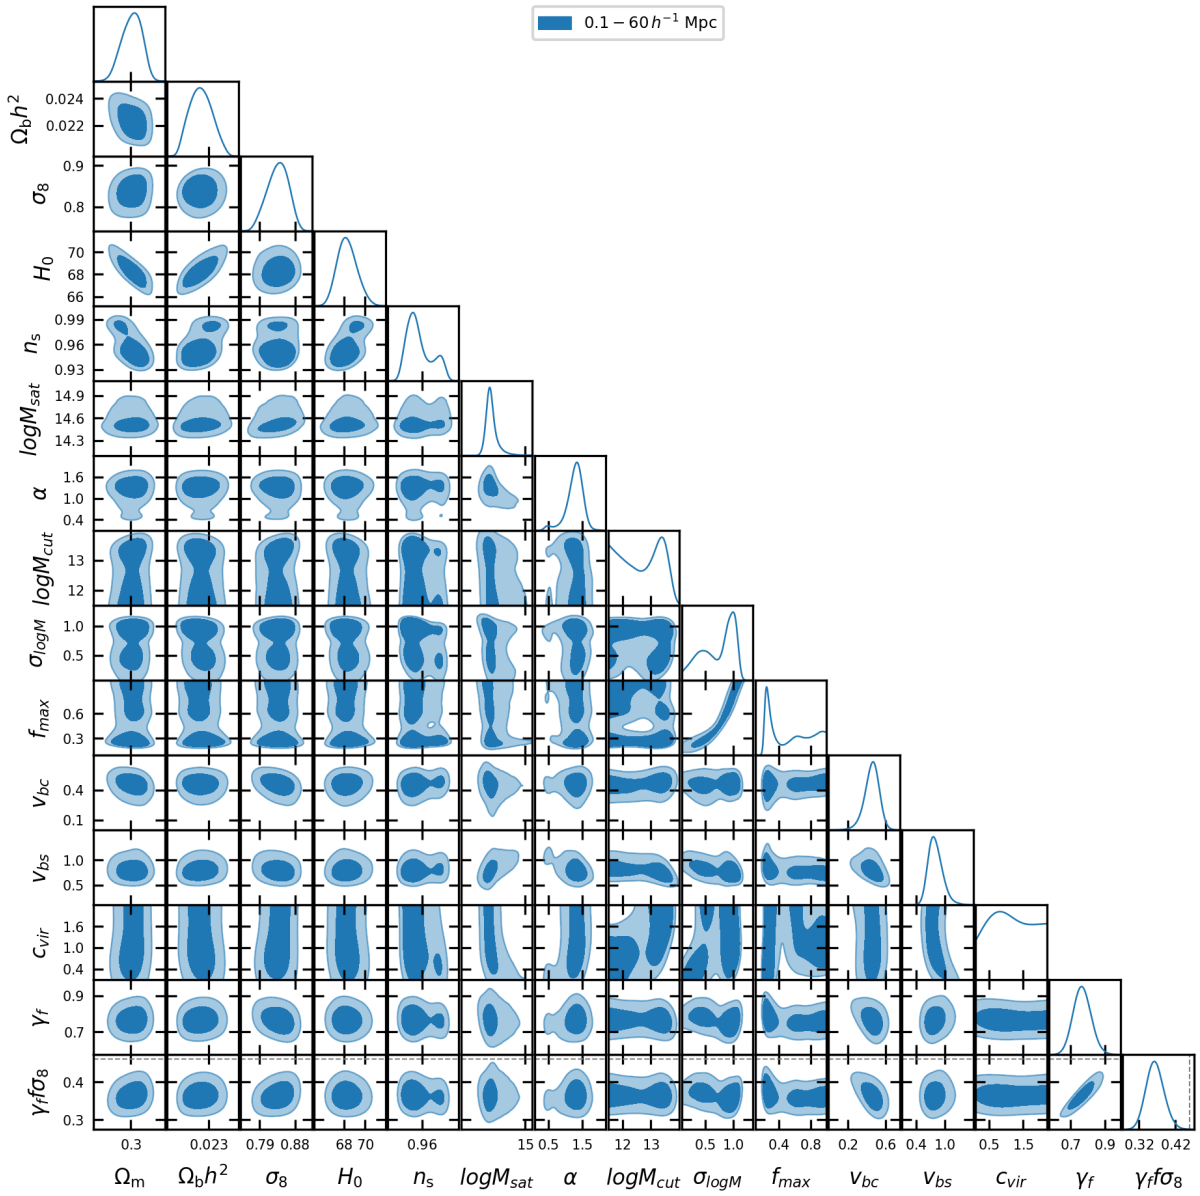

Figure 3.10: 1D and 2D contours of the parameters used in our baseline fit, as well as the derived constraints on $\gamma_f f\sigma_8$.

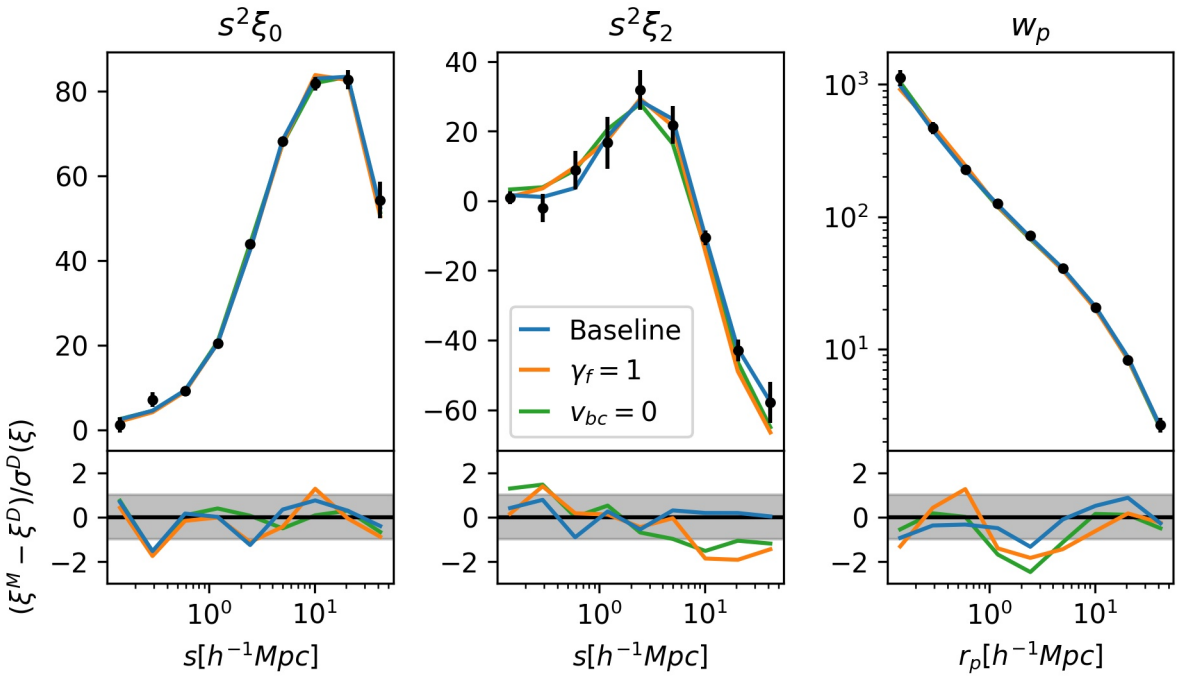

Figure 3.11: Comparison of the best fit model predictions to the data for several fits to the eBOSS LRG sample for the monopole (*left*), quadrupole (*centre*) and projected correlation function (*right*). *Upper panels:* The baseline fit (blue), fixed $\gamma_f = 1$ fit (orange), and $v_{\rm bc} = 0$ fit (green), with the data and measurement uncertainty (black). *Lower panels:* The difference between the best fit models and the data divided by the measurement uncertainty. The 1σ region is shown in grey.

3.5.2 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 ($\log M_{\rm sat}$, α, $\log M_{\rm cut}$) and the three parameters that control the position and velocity bias of galaxies within halos ($v_{\rm bc}$, $v_{\rm bs}$, $c_{\rm vir}$) 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 3.2), 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 fit 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, $f_{\rm max}$) free. In the second fit we also keep $\sigma_{\log M}$ and $f_{\rm max}$ fixed to their best fit values from the full fit, allowing only the six cosmological parameters to vary. The γ_f constraints from these fits are shown in Table 3.2 and Fig. 3.12. 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.

3.5.3 Testing the impact of the cosmological priors

We consider a number of prior constraints on the cosmological parameters, as described in Sec. 3.3.6. The three most significant cases are a uniform prior as described in Table 3.1, 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 *Planck* 2018 likelihoods with

Run	γ_f	N_P	N_D	χ^2
$0.1 - 60\,h^{-1}$Mpc	0.767 ± 0.052	14	27	14.1
$0.1 - 7\,h^{-1}$Mpc	0.71 ± 0.14	14	18	7.8
$0.8 - 60\,h^{-1}$Mpc	0.783 ± 0.066	14	18	4.2
$7 - 60\,h^{-1}$Mpc	0.854 ± 0.083	14	9	0.36
$7 - 60\,h^{-1}$Mpc, 8 parameters	0.821 ± 0.064	8	9	0.74
$7 - 60\,h^{-1}$Mpc, 6 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
$v_{\rm bc} = 0$	0.958 ± 0.088	13	27	22.5
$f_{\rm max} = 1$	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+*Planck*18	0.784 ± 0.048	14*	27	18.5
eBOSS+*Planck*18 scaled σ_8	0.798 ± 0.047	14*	27	19.1
eBOSS+*Planck*18 free σ_8	0.766 ± 0.053	14*	27	18.0
No AP scaling	0.772 ± 0.053	14	27	14.5

Table 3.2: 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. Unless specified in the table or text, the analysis methods for each fit match the headline, i.e. all data over the measurement range $0.1 - 60\,h^{-1}$Mpc is included, the smoothed jackknife covariance matrix is used, the training prior is applied, and AP scaling is included. *The eBOSS+*Planck*18 runs jointly fit 5 of the 14 parameters with *Planck*, so they are not fully independent.

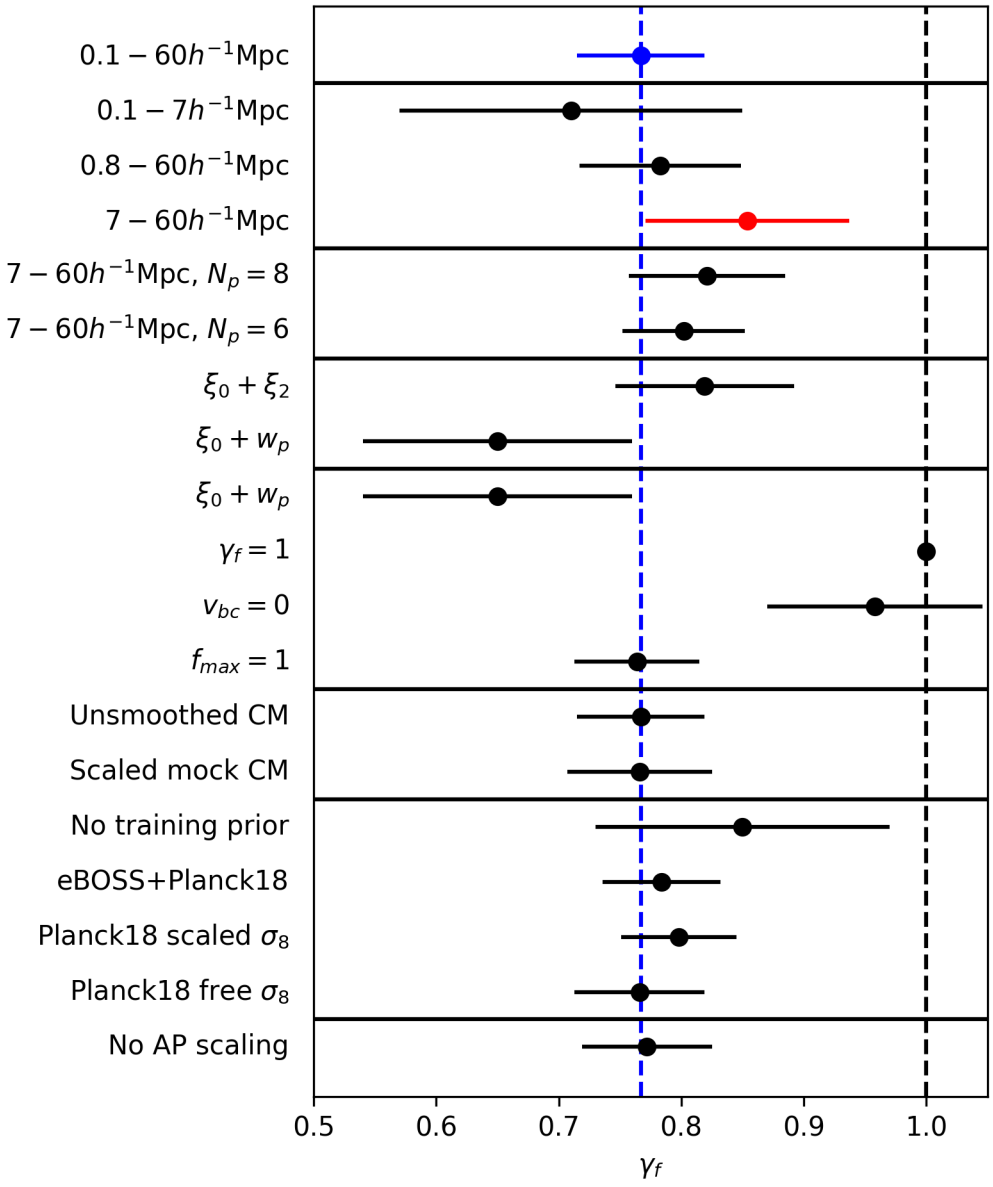

Figure 3.12: γ_f constraints from all the runs listed in Table 3.2. The blue point shows the baseline fit to the full separation range, extended by the blue dashed line for comparison to other points. The red point shows the fit to the quasi-linear scales only. The black dashed line shows $\gamma_f = 1$ for comparison, the value expected if the amplitude of the halo velocity field matches the ΛCDM expectation.

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 from fits to the eBOSS data over the full emulator range ($0.1-60\,h^{-1}$Mpc) are shown in Fig. 3.13. 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 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 Sec. 3.3.6 can be found in Table 3.2. 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.

3.5.4 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 Sec. 3.4.1 we fit to the full non-linear regime ($0.1-7\,h^{-1}$Mpc), the weakly non-linear and quasi-linear regimes ($0.8-60\,h^{-1}$Mpc), and the quasi-linear regime only ($7-60\,h^{-1}$Mpc). 1D and 2D contours in the $v_{\rm bc}-\gamma_f$ parameter space for these three fits are shown in the left panel of Fig. 3.14. There is little variation in the other parameters between these fits to different scales, however some important insight is gained from examining the $v_{\rm bc}-\gamma_f$ 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 $v_{\rm bc}$, while obtaining smaller and less precise constraints on γ_f. The full fit to all scales is located at the intersection in $v_{\rm bc}-\gamma_f$ 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 than 1σ. 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, with these results showing that there is a strong tension that is most significant in the non-linear regime.

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

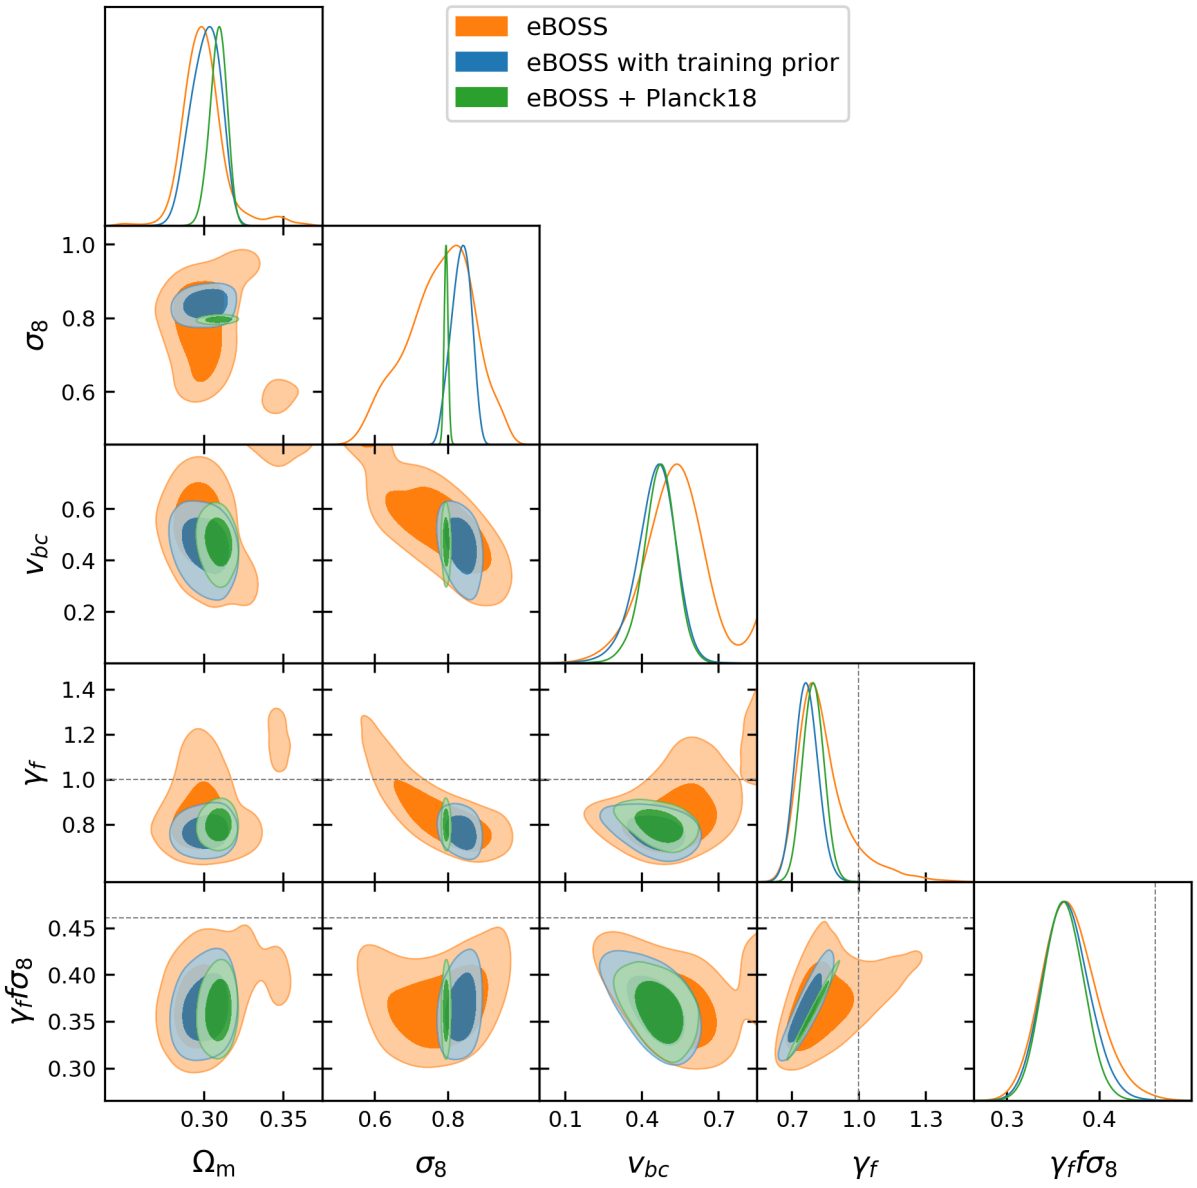

Figure 3.13: 1D and 2D contours of the key fit parameters for the fit to the eBOSS LRG sample with no additional cosmological constraints (orange), restricted by the Aemulus training prior (blue), and jointly fit with the *Planck* 2018 likelihoods (green).

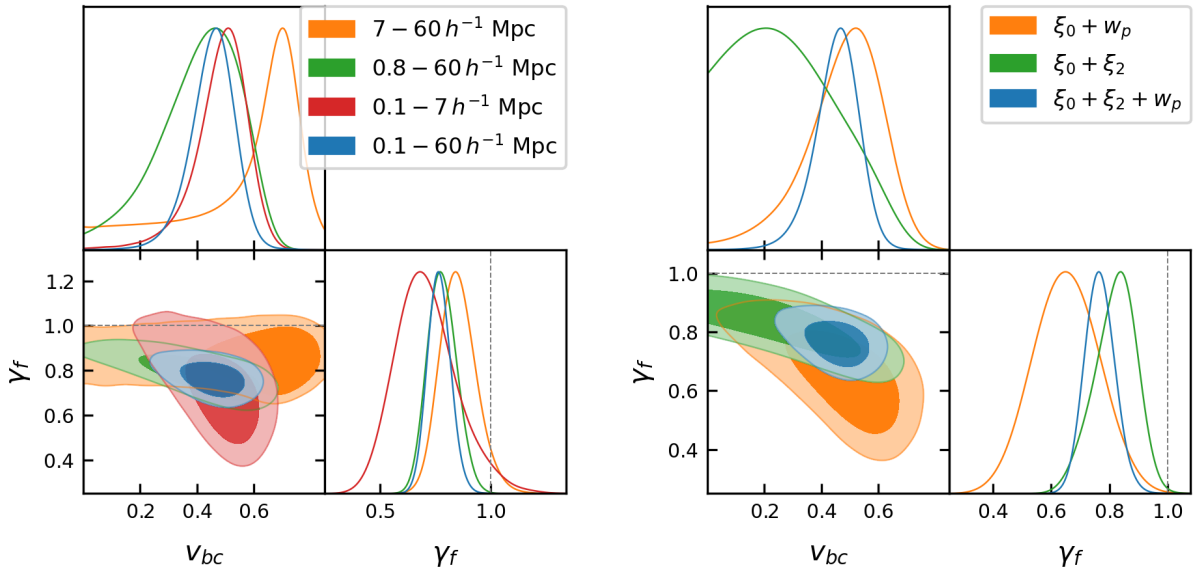

Figure 3.14: 2D and 1D marginalized constraints on $v_{\rm bc}$ and γ_f for fits to different scales and measurements. *Left:* Constraints from the three largest separation bins (orange), six largest separation bins (green), six smallest separation bins (red), and all nine separation bins (blue) for 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 ΛCDM. *Right:* Constraints from the joint fit to the monopole and projected correlation function (orange), monopole and quadrupole (green), and all three measurements (blue).

with any value of $v_{\rm bc}$ 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 panel of Fig. 3.14 we show the constraints in $v_{\rm bc} - \gamma_f$ parameter space from the joint fit to only the monopole and projected correlation function, and the joint fit to the multipoles only. The multipole only fit is less sensitive to the degeneracy between $v_{\rm bc}$ and γ_f, but prefers a smaller value of $v_{\rm bc}$ 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 $v_{\rm bc}$ 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.

3.5.5 Testing the dependence on the covariance matrix

In Sec. 3.3.4 we compared our jackknife covariance matrix to a covariance matrix estimated from 1000 EZmocks. We found good agreement in both the variance and correlation structure between the two methods, so we performed additional fits to the data using the mock covariance matrix and the unsmoothed jackknife covariance matrix to assess the robustness of our constraints with respect to these choices.

It should be noted that the clustering of the mocks on scales below $\sim 1\,h^{-1}$Mpc is significantly lower than the data, so that 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

$$\sigma_{i,i}^{M,s} = \sigma_{i,i}^{M} \frac{\xi_i^D}{\bar{\xi}_i^M}, \tag{3.10}$$

where ξ^D is the correlation function from the data and $\bar{\xi}^M$ 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 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 3.2. 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.

3.5.6 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 3.2 and Fig. 3.12. Excluding the AP correction has a negligible effect on the constraint on γ_f and slightly increases the best fit 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.

3.5.7 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 $9376\,\mathrm{deg}^2$ over a redshift range of $0.4 < z < 0.8$ [225] with a target density of $99.5\,\mathrm{deg}^{-2}$. 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 [225] 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. 3.15.

The advantage of this sample is that it is more complete due to the complementary 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.

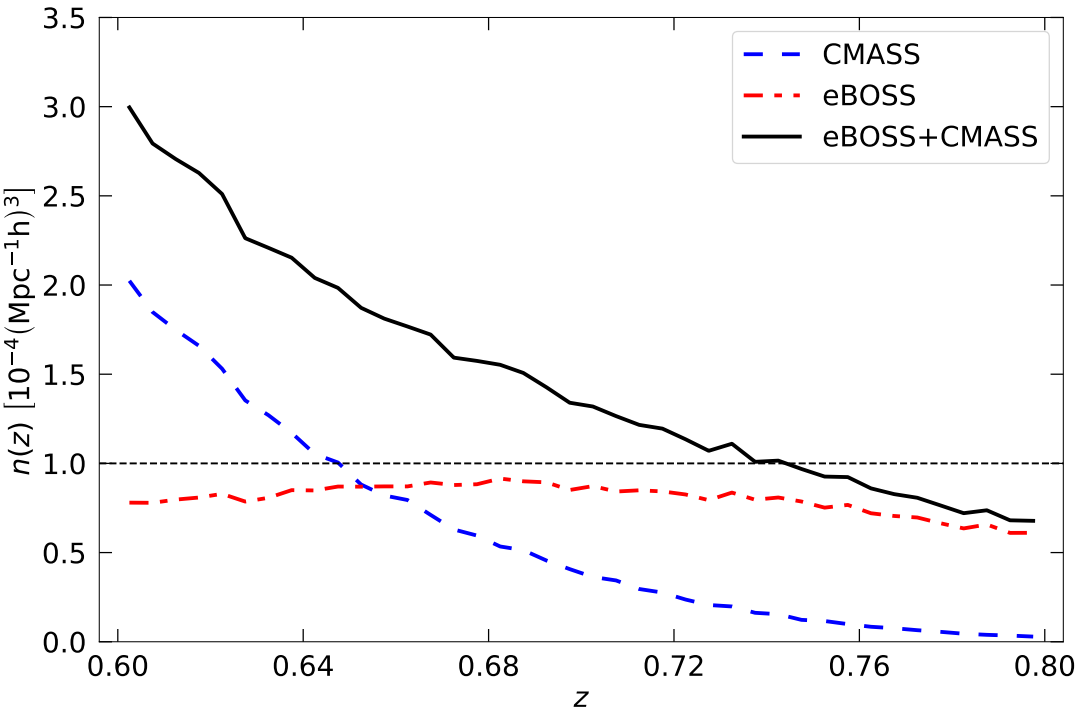

Figure 3.15: Redshift distribution of the eBOSS DR16 (red dash-dotted line), CMASS DR12 (blue dashed line) and the joint eBOSS+CMASS sample (black thick line, see Sec. 3.5.7 for details), optimized using $w_{\rm FKP}$ weights.

To correct fibre collisions in the CMASS sample we use the Close Pair (CP) weights described in Sec. 2.5.1, and the standard angular upweighting method described in Sec. 2.5.3. 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 Fig. 2.9 and Fig. 2.10). 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. 3.16 shows the result of our fit compared to the eBOSS only fit in the most important parameters of our analysis for the full emulator range, while Fig. 3.17 shows the result of fits to 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 $f_{\rm max}$ values, as expected. It is interesting to note that in the fit over the full emulator range (Fig. 3.16) 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. Firstly, 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 complementary constraining of different scales and probes on parameter combinations such as $v_{\rm bc}$ and γ_f (Fig. 3.14). 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 (Fig. 3.17), 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.

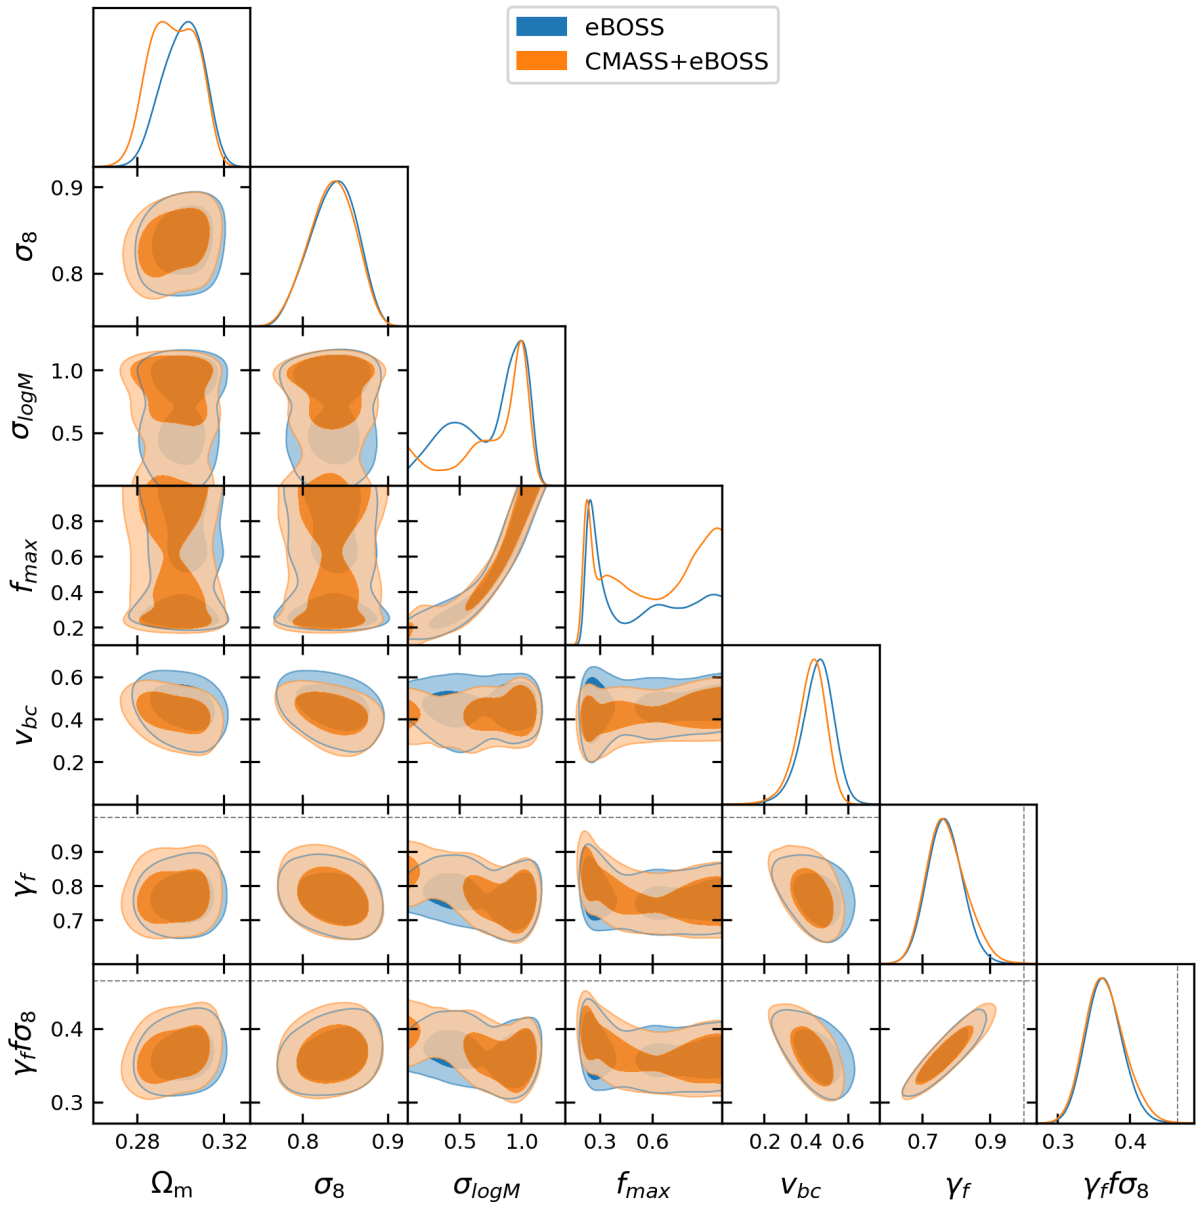

Figure 3.16: 2D and 1D marginalized constraints of the key parameters of our fit for our fiducial eBOSS measurement (blue) and combined CMASS+eBOSS sample (orange) over the full emulator range, $0.1 - 60\,h^{-1}$Mpc.

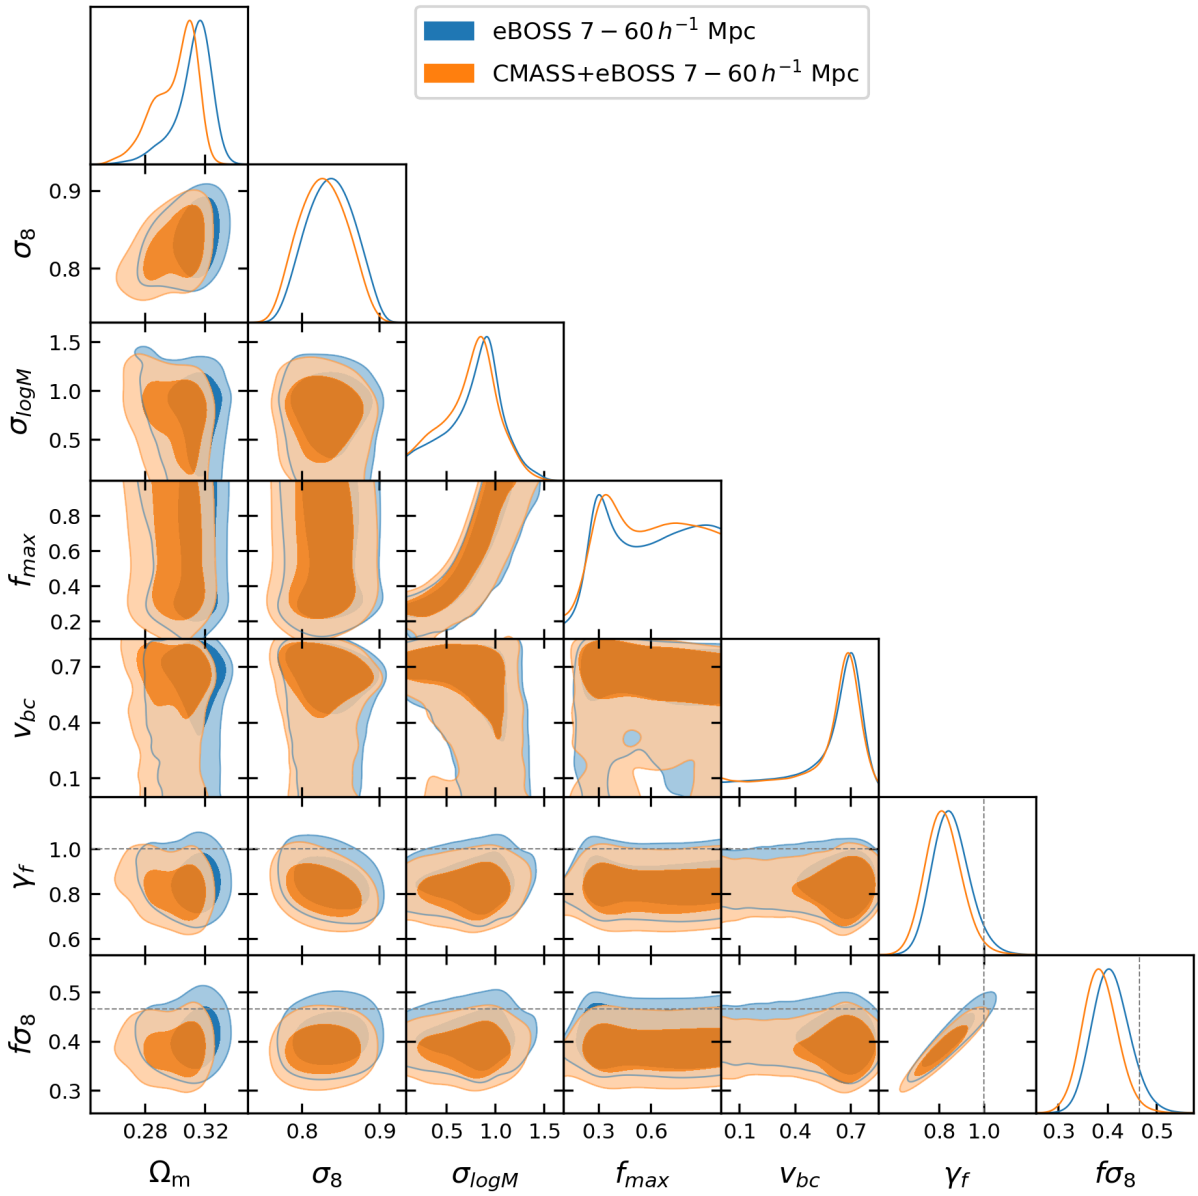

Figure 3.17: Same as Fig. 3.16, but fit over the quasi-linear scales only, $7-60\,h^{-1}$Mpc.

3.6 Discussion

3.6.1 Comparison to other measurements

We compare our result to other measurements of $f\sigma_8$ from galaxy clustering surveys in Fig. 3.18. Taken as a whole sample, there is clearly good consistency with the ΛCDM prediction. For the eBOSS LRGs, [23] analyzed pairs with separations between $25 - 130\,h^{-1}$Mpc, and obtained measurements of $f\sigma_8 = 0.446 \pm 0.066$ and $f\sigma_8 = 0.420 \pm 0.065$ depending on the RSD model used in the analysis (see Table B1 of [23]). Our measurement is consistent with these results at around the $\sim 1\sigma$ 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. 3.18 we also compare our results to other attempts to measure $f\sigma_8$ on small scales. [226] used a similar parameterization 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 [226] used a fixed cosmology, which has been shown by [278] to significantly reduce the uncertainty. Conversely, [166] 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 an equivalent of our γ_f parameter that allows the linear growth rate to change independently of the ΛCDM 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 Sec. 3.5.1 and Sec. 3.5.4, using the full separation range significantly increases the tension with the result expect for ΛCDM, 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 [166] found a similar dependence on the measurement scales, with smaller scales preferring a smaller value of $f\sigma_8$. [166] also found that adding the projected correlation function to their fiducial measurement of the monopole, quadrupole, and hexadecapole reduced the best fit value of their lower redshift sample by $\sim 1\sigma$, 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

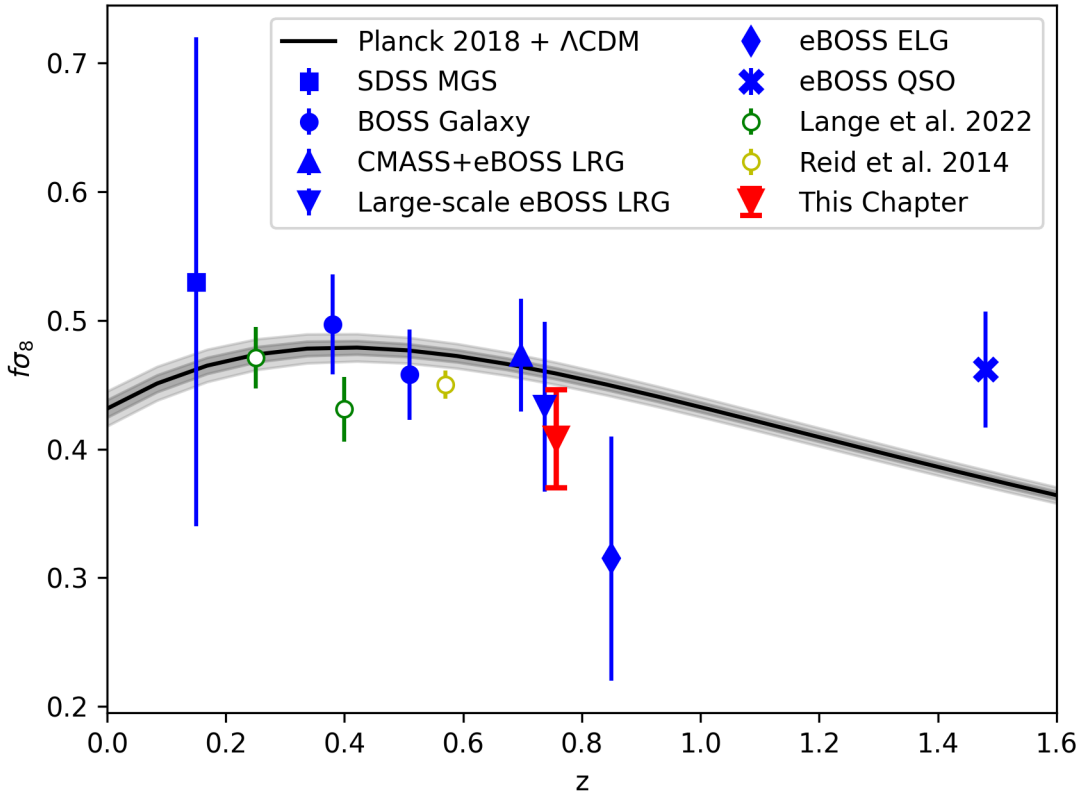

Figure 3.18: $f\sigma_8$ measurements from various SDSS samples. The blue points show the results of the more conventional large-scale analyses from the SDSS MGS [143], BOSS galaxies [9], CMASS+eBOSS LRGs and eBOSS LRGs [23], eBOSS ELGs [76], and eBOSS quasars [196]. Our small-scale analysis of the eBOSS 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 ([166], green) and BOSS CMASS sample ([226], yellow) that included non-linear scales in the analysis. The black line shows the expected value of $f\sigma_8$ for a flat ΛCDM universe with best fit *Planck* 2018 cosmology, where the shaded regions show the 1 and 2σ confidence regions. The large-scale eBOSS LRG result is shifted in the x-axis to avoid overlap with the small-scale result from this chapter.

impact of the different measurements and scales between our results. This is particularly true since [166] do not include a parameter comparable to our γ_f, given the importance of w_p in breaking the $v_{\rm bc} - \gamma_f$ degeneracy in our analysis.

3.6.2 Galaxy-halo connection parameters

The parameter found to be most degenerate with our γ_f constraint is $v_{\rm bc}$, the scaling of the velocity dispersion of centrals in the HOD framework (Fig. 3.10). A lower value of $v_{\rm bc}$ corresponds to a larger γ_f, as expected in the non-linear regime since both parameters increase the observed velocity dispersion of galaxies (see Sec. 3.4.3). Our fit over the full $0.1 - 60\,h^{-1}$Mpc separation range strongly prefers a non-zero $v_{\rm bc}$ and low γ_f. However, our fit to the quasi-linear regime finds no discernible degeneracy between $v_{\rm bc}$ and γ_f and recovers both a relatively large value of γ_f and non-zero value of $v_{\rm bc}$, although the constraint on $v_{\rm bc}$ is weak due to the small impact it has on those scales (Fig. 3.14). This result indicates that the degeneracy between $v_{\rm bc}$ and $f\sigma_8$ may illustrate the degree to which the non-linear scales affect the overall constraint. [166] 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 $v_{\rm bc} > 0$ and low $f\sigma_8$ compared to the ΛCDM prediction. [226] elected to fix the velocity of centrals to match that of the host halo, and find closer agreement with the ΛCDM expectation, which we also find when using a fixed $v_{\rm bc} = 0$. $v_{\rm bc} > 0$ 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 $f_{\rm max}$ parameter. Due to the strong degeneracy between $\sigma_{\log M}$ and $f_{\rm max}$, our fit to the data is broadly consistent with a wide range of values for $f_{\rm max}$ between 0.2 and 1, however there is a large peak at $f_{\rm max} = 0.25$. A low value of $f_{\rm max}$ is not surprising for the eBOSS sample given the magnitude and color 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 $f_{\rm max}$ is not expected to bias our measurement.

Numerical simulations have shown that the clustering of dark matter halos can depend on properties other than halo mass, a.k.a halo assembly bias ([245, 108, 128, 268, 197]). This bias can propagate into the distribution of galaxies that live in these halos and thus introduce additional bias in the clustering measurement. In the analysis of BOSS galaxies over a wider redshift range using a similar emulator model [277], the basic HOD approach

used here was enhanced with an assembly bias model depending on the environment of dark matter halos. 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 modeling of assembly bias in this analysis.

3.6.3 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 S_8-tension between *Planck* and weak lensing surveys (see Sec. 1.2.5) and the low amplitude of the galaxy-galaxy lensing amplitude measured using the BOSS CMASS sample by [171], 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 [2], *Planck* 2018 [213], and our results (Fig. 3.19). The left panel shows our measurement using the full separation range, while the right 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 or velocity field of the simulations, rather than a purely cosmological tension.

3.6.4 Emulator robustness and potential improvements

We have performed rigorous tests of the emulator performance (see Sec. 3.4), 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 $f_{\rm max}$ to the emulator, and verify that the full emulator gives an unbiased measurement for

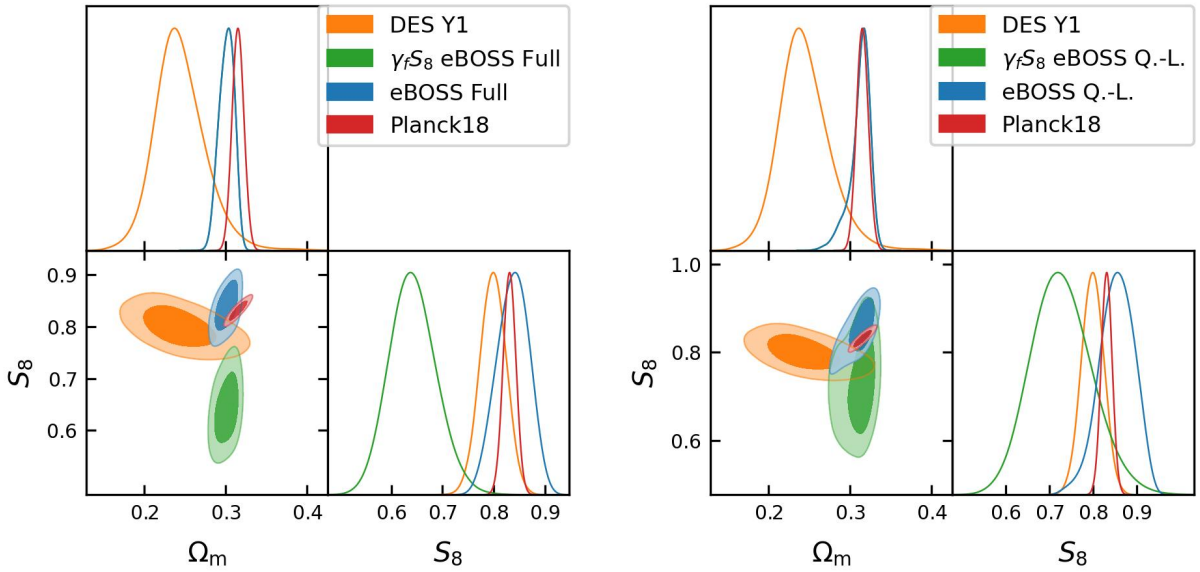

Figure 3.19: 2D and 1D marginalized constraints on Ω_m and S_8 from our analysis (blue), the Dark Energy Survey (DES) year 1 results [2] (orange) and *Planck* 2018 results [212, 214] (red). Since our low value of $f\sigma_8$ mostly comes from γ_f, we also plot $\gamma_f S_8$ (green) for our fit, which shows the constraint we would have if the low value of $f\sigma_8$ came entirely from the σ_8 value. *Left:* results of our fit over the full emulator range, $0.1 - 60\,h^{-1}$Mpc. *Right:* results from only the quasi-linear scales, $7 - 60\,h^{-1}$Mpc, used to constrain the linear growth rate.

$0.2 < f_{\rm max} < 1.0$. 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 [79, 80] and *Euclid* [168].

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. 3.5), 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 Sec. 3.5.2). 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.

3.7 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 ΛCDM 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 ΛCDM 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 $f_{\rm max} \leq 0.6$. 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 halos 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 ΛCDM 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.

Chapter 4

Isolating the linear signal when making RSD measurements

This chapter is adapted from [58]. I was the principal investigator of that work and performed all steps of the data analysis, with the exception of training the emulator and estimating the emulator uncertainty, both performed by Prof. Zhongxu Zhai. I lead the writing of the paper for publication, with input from my co-authors. Throughout this chapter I use 'we' to describe all contributions, including my own. The content of that work has been adapted and reformatted by me for this thesis.

4.1 Introduction

Our first measurement of RSD from small-scale clustering within the eBOSS LRG sample, presented in Ch. 3, successfully found a significant improvement in precision over the large-scale analysis of the same sample when the full separation range of the emulator was used [57]. However, that improvement was reduced because we were forced to restrict our analysis when quoting a constraint on $f\sigma_8$ with a minimum scale cut to match the scale where changing γ_f no longer directly matched the expectation for a change in $f\sigma_8$. While the small-scale, non-linear velocities are certainly affected by a change in the growth rate, it is no longer necessary that that change be directly proportional, so there is a potential for a systematic bias in applying a linear velocity scaling to non-linear velocities.

This highlights a larger issue in the area of small-scale RSD measurements; how to measure a linear quantity in the non-linear regime without allowing the non-linear velocity

evolution to bias the results. This is the primary motivation for this chapter. We build on the previous model by splitting the velocity scaling parameter γ_f into two parameters: γ_l to scale the linear component of the velocity, and γ_n to scale the non-linear component. This new parameterization allows us to interpret a change in γ_l as a change in the amplitude of the linear velocity field consistent with a change in $f\sigma_8$ within a ΛCDM framework, while γ_n allows enough freedom for the non-linear velocity to vary without directly matching the scaling of the linear velocity.

This chapter is structured as follows. In Sec. 4.2 we expand on the model of Ch. 3 to isolate the linear signal in the non-linear regime using our new velocity scaling parameters. Then we refit the eBOSS LRG data using the new emulator, and present the results in Sec. 4.3. Finally, in Sec. 4.4 we discuss the significance of our new results and compare to the work of the previous emulator and other related measurements.

4.2 Modelling RSD including velocity scaling

4.2.1 Measuring the growth of structure with velocity scaling

Following the work of Ch. 3, in this chapter we use velocity scaling to allow our model to test for RSD that differ from the model used to fit to other aspects of the clustering measurements. The velocity scaling parameter, γ_f, is implemented as a rescaling of all halo bulk velocities in a simulation, where we use the same definition of 'bulk velocities' as Ch. 3, meaning the velocity of the halo as a single unit, rather than the velocity of the individual particles making up the halo or the internal velocity dispersion of the halo. In the linear regime the amplitude of the velocity field is directly proportional to $f\sigma_8$, so a scaling of the velocity field has the same effect as scaling the logarithmic growth rate, f [226]. In this way γ_f can be used to measure an effective value of $(f\sigma_8)^{\rm meas} = \gamma_f f_{\Lambda{\rm CDM}}\sigma_8$. It is important to note that in the ΛCDM cosmological model f is a derived parameter that is determined entirely by the value of more fundamental cosmological parameters, most notably Ω_m, so here $f_{\Lambda{\rm CDM}}$ is used to represent the value of f calculated for a particular set of cosmological parameters within the ΛCDM framework. Throughout the remainder of this chapter we will drop the superscript "meas", and refer to results calculated as $f\sigma_8 = \gamma_f f_{\Lambda{\rm CDM}}\sigma_8$ when a velocity scaling parameter is included in the model.

A value of $\gamma_f \neq 1$ indicates disagreement between the data and the ΛCDM cosmological model and provides a test of ΛCDM. While that is an accurate interpretation, we can also present our main results in terms of a constraint on $f\sigma_8$, as described above. It is useful

to present our results in this way in order to make them more easily comparable to results from other RSD measurements, including those from the eBOSS LRG configuration-space clustering on larger scales using more conventional RSD modelling presented in Sec. 2.7.2 (see [23]). This comparison is key to assessing the information content of the non-linear and quasi-linear regimes compared to linear scales.

It is also worth considering whether including a parameter such as γ_f is still necessary for an emulator-based model, such as the one applied in Ch. 3 and the one that will be described in Sec. 4.2.2. An emulator is capable of making predictions for galaxy clustering over a wide region of cosmological parameter-space, meaning the model is able to accommodate different values of $f\sigma_8$ by changing Ω_m and σ_8. However, without γ_f such a model is only able to produce a good fit to data where the velocities match the geometry and clustering of a ΛCDM cosmological model, as determined by the combination of simulations and other model elements such as the galaxy-halo connection. Fitting to data that is not well-matched by this framework will result in a poor agreement between data and model, characterized by the χ^2 statistic, and could bias the inferred parameter constraints. It is also difficult to separate the constraints from the RSD from those from fitting the overall clustering pattern.

By including a velocity scaling parameter the model is given the freedom to fit differences in the amplitude of the velocity field directly, unlike constraints on Ω_m and σ_8 that affect other elements of the clustering in addition to their influence on the velocity field (see Sec. 4.2.5). Splitting the parameters that contribute to the $f\sigma_8$ constraint in this way also provides more information about where the constraint comes from, since the parameters γ_f, Ω_m, and σ_8 are not completely degenerate. For example, in Ch. 3 we found a value of $f\sigma_8$ lower than the expectation for a ΛCDM universe with *Planck* 2018 [213] cosmology when including a velocity scaling parameter, with a reasonable χ^2 value. However, fixing the scaling parameter to $\gamma_f = 1$ and allowing $f\sigma_8$ to vary only through Ω_m and σ_8 resulted in a much worse fit to the data, and a shifted constraint on $f\sigma_8$. This result indicated that the discrepancy between the data and model was specifically in the amplitude of the velocity field, which is information that can only be gained by including additional degrees of freedom.

Related to these results from Ch. 3, it is worth considering an issue highlighted in that analysis: what velocities can be considered as linear for the purposes of the growth rate? While a change in the growth rate will affect all components of the velocity, the relation between the amplitude of the non-linear velocity field and f may not be directly proportional. In Ch. 3 we investigated the effect of varying γ_f on the correlation function and identified a scale of $\sim 7\,h^{-1}\,$Mpc as the transition between the quasi-linear and non-linear regimes, so we restricted our measurement of $f\sigma_8$ to between $7 - 60\,h^{-1}\,$Mpc to

isolate the linear signal when using a single scaling parameter. The primary purpose of this chapter is to improve on the model used in Ch. 3 to be able to include all scales in the constraint on $f\sigma_8$ without bias from misinterpreting non-linear information. That process will be outlined in the following sections.

4.2.2 Building an emulator

In this work we build on the emulator used in Ch. 3, originally based on [278]. The emulator used a 5-parameter cosmological model consisting of Ω_M, Ω_b, σ_8, h, and n_s, as well as an 8-parameter HOD model to connect galaxies to halos in the simulation, described by the parameters $f_{\rm max}$, $\sigma_{\log M}$, $\log M_{\rm sat}$, α, $\log M_{\rm cut}$, $c_{\rm vir}$, $v_{\rm bc}$, and $v_{\rm bs}$. The final parameter of the Ch. 3 emulator was the velocity scaling parameter, γ_f, described in Sec. 4.2.1.

We improve on the Ch. 3 emulator using the method described in Sec.4.2.3 to model the linear and nonlinear velocity components. In order to apply this new method we require access to the initial conditions of the simulation, which are not publicly available for the Aemulus suite of simulations [77] used by the Ch. 3 emulator. For our new emulator we use the AbacusCosmos suite of simulations [110], with available first-order initial conditions generated from the zeldovich-PLT[1] code [109]. AbacusCosmos consists of 40 variable cosmology, 1100 h^{-1} Mpc simulation boxes with 1440^3 particles that we use to train the emulator, as well as 20 simulation boxes at the *Planck* 2015 cosmology [210] that are used for testing. Since the AbacusCosmos and Aemulus suites are similar in terms of number of boxes, box size, and number of particles we use the same method to estimate the emulator uncertainty as [278], adapted to the boxes available in AbacusCosmos. We use the 20 AbacusCosmos boxes with *Planck* cosmology to estimate the sample variance, and assess the performance of the emulator throughout the cosmological parameter space by retraining the emulator with one variable cosmology box excluded at a time, and comparing emulator predictions to measurements from the excluded box.

4.2.3 Isolating the linear signal

In order to ensure that our results are not biased by the assumption that all components of the velocity will be scaled in the same way by a change in f, we split the velocity of halos into two components: a linear and a non-linear component. We scale each component by an independent parameter: γ_l for the linear component and γ_n for the non-linear component.

[1] https://github.com/abacusorg/zeldovich-PLT

If these parameters are constrained such that $\gamma_l = \gamma_n$, then all velocities are scaled by the same amount and the model reduces to the single scaling parameter, γ_f, used in Ch. 3. The split is performed on halo velocities rather than galaxy velocities because the velocity bias of galaxies is implemented by other independent parameters in the emulator. Galaxies are assigned the velocity of their host halo, with an additional velocity term calculated as $\sigma_{\rm gal} = v_{\rm gal}\sigma_{\rm halo}$, where $v_{\rm gal}$ is the velocity bias parameter for that galaxy type ($v_{\rm bc}$ for centrals and $v_{\rm bs}$ for satellites), and $\sigma_{\rm halo}$ is the velocity dispersion of the halo. Halo velocity dispersion is calculated from its mass using the circular velocity at the virial radius, which provides a physically motivated estimate that is fast and simple to calculate, and is not affected by noise from small numbers of particles for low mass halos in the simulation. The additional velocity term is calculated independently of the velocity scaling by γ_l and γ_n, so that it is controlled entirely by $v_{\rm bc}$ and $v_{\rm bs}$. This choice reduces the degeneracy between the velocity scaling and velocity bias parameters while still allowing for sufficient freedom in the model to address both a change in the growth rate and the presence of velocity bias [118].

The challenge of this new model is determining what component of the velocity is linear at late time. While this is difficult to do for the halo velocities, we can make use of the fact that the initial conditions of the emulator provide a method for calculating particle linear velocities, which can then be combined to provide an estimate of the linear velocity of the halo. The AbacusCosmos initial conditions were generated by calculating Zel'dovich approximation displacements for a grid of particles at $z = 49$ using the zeldovich-PLT code. The Zel'dovich approximation provides a first order calculation of the displacements and velocities of particles, so $z = 49$ is chosen as an arbitrarily large redshift where the motion of particles will very closely follow linear theory. We can use these initial particle linear velocities to predict the particle linear velocities at the $z = 0.7$ simulation slice by evolving them using the linear theory prediction for the amplitude of the velocity field,

$$\boldsymbol{v}_{\boldsymbol{k}} = \frac{i\boldsymbol{k}}{k^2} Ha\delta_{\boldsymbol{k}} f(\Omega_m). \tag{4.1}$$

The velocity scaling of the initial conditions is simply the ratio of Eq. 4.1 between the redshift of the initial conditions and the desired final redshift,

$$\boldsymbol{v}(z_2) = \frac{Haf\sigma_8(z_2)}{Haf\sigma_8(z_1)} \boldsymbol{v}(z_1). \tag{4.2}$$

We define the non-linear velocity as all components of the total velocity not included in the linear velocity, and calculate it by subtracting the linear velocity vector from the total velocity vector. By separately scaling the linear velocity by γ_l and the non-linear

velocity by γ_n, we allow for the non-linear velocity of the data to deviate from the ΛCDM expectation of the simulations without biasing the value of $f\sigma_8$ we infer from γ_l. While only γ_l affects our measurement of $f\sigma_8$, there are several benefits to also allowing γ_n to vary in our model. The amplitude of the non-linear velocity component is not necessarily directly proportional to $f\sigma_8$, however it is still reasonable to expect that a change in $f\sigma_8$ would affect the amplitude of the non-linear component. γ_n provides our model with the freedom to reproduce such changes, which are then marginalized over in the constraints of γ_l and $f\sigma_8$. γ_n also provides the model with freedom in the case that the prediction for the non-linear velocity component from the simulation+HOD model is a poor fit to the data. There is some evidence for such a disagreement, as the results of both Sec. 3.5.4 and [166] found that measurements from smaller separation ranges produced lower constraints on $f\sigma_8$, which is in conflict with the ΛCDM expectation that $f\sigma_8$ should be scale-independent. There are several explanations for a discrepancy in the amplitude of the non-linear velocity component between the model and the data. The simulations may fail to produce the correct non-linear velocity distribution, particularly due to the effects of baryonic physics. Another possibility is if the chosen galaxy-halo connection model is not able to reproduce the correct non-linear velocity distribution. A final benefit is that the value of γ_n itself provides a test of the agreement between the non-linear velocity component of the model and the data, with $\gamma_n \neq 1$ indicating a disagreement. This disagreement could either be caused by one of the factors listed above, or by a breakdown in the ΛCDM cosmological model.

It should also be noted that while γ_n is constructed to have a minimal impact on the constraint of $f\sigma_8$, in general γ_l and γ_n will have some correlation with each other. For example, this will be true for quasi-linear velocity evolution that happens along the direction of the linear velocity. Reducing this correlation would require a more detailed estimation of the linear and non-linear velocity components, which we leave to future work.

4.2.4 Smoothing the linear velocity field

Pairs of galaxies with small separation in collapsed objects have lost all dependence on the initial linear velocities. This approximately occurs at shell crossing and means that our split into linear and non-linear components is ineffective on such scales - a portion of the velocity ascribed to non-linear motion simply cancels out the linear one (see Appendix C). In an extreme situation, if two objects are located sufficiently close to each other along the line-of-sight and have a large enough infall velocity, the shift in position in redshift space reverses the orientation of the pair along the line-of-sight. In this situation scaling the velocity will increase the pair separation, leading to damping of the correlation function.

We therefore elect to smooth the particle linear velocity field around the shell-crossing scale, which from our analysis in Ch. 3 we know to occur at approximately $5\,h^{-1}$Mpc. This smoothing reduces the pairwise linear velocity of nearby objects, transferring the component of the velocity that provokes shell crossing to what we have termed the 'non-linear' component, since total velocity is still conserved. Meanwhile, the linear pairwise velocity of more distant objects is unaffected, preserving the signal we wish to extract with our linear velocity scaling parameter.

To illustrate the smoothing effect we use a projected $5\,h^{-1}$Mpc thick slice of the Abacus-Cosmos *Planck* 00-0 box to demonstrate the arrangement of the different particle velocity components in a high density region, shown in Fig. 4.1. The panels of the plot show the total particle velocity; the linear velocity calculated from the initial conditions; the smoothed linear velocity calculated using a tophat smoothing kernel with radius $5\,h^{-1}$Mpc; and the non-linear velocity component, calculated as the difference between the total velocity and the linear velocity. While the velocity of field particles is largely unchanged between total, linear, and smoothed linear velocities, the behaviour of particles in the cluster differs greatly. The unsmoothed linear velocity displays a distinct preferred direction when compared to the total velocity, however some scatter persists. The smoothed velocity is significantly more collimated, so that close particles will maintain their separation in redshift space, as intended. The non-linear velocities show the difference between the total velocity and smoothed linear velocities. As expected, the non-linear velocities are significantly larger in collapsed structures compared to the field, and do not show an obvious preferred direction. While we use particle velocities to in Fig. 4.1 to illustrate the various components, it is important to note that the velocities of galaxies in the final HOD catalogue are determined from the halo velocities, which are detailed below.

Our process of smoothing and assigning halo velocities is as follows. First, we construct a 3D grid with side length $1\,h^{-1}$Mpc over the simulation box, and assign to each grid cell a linear velocity equal to the mean linear velocity of the particles contained within the cell. Next, we smooth the grid using a 3D spherical tophat kernel of radius $5\,h^{-1}$Mpc, equally weighting each grid cell. Finally, halos are assigned the smoothed linear velocity of the cell they inhabit. The smoothing radius of $5\,h^{-1}$Mpc was chosen to match the approximate scale found in Sec. 3.4.1 where increasing the velocity scaling parameter, γ_f, transitioned from amplifying the monopole to damping the monopole. A tophat kernel was chosen because of the small width of this transition, and because it reduces the number of calculations required compared to other possible kernel choices, such as a Gaussian kernel. The grid spacing was chosen to balance the resolution of the grid and the memory requirements of the computation. Testing these choices is discussed below.

In Fig. 4.2, we investigate the effect of scaling the smoothed halo linear velocity on

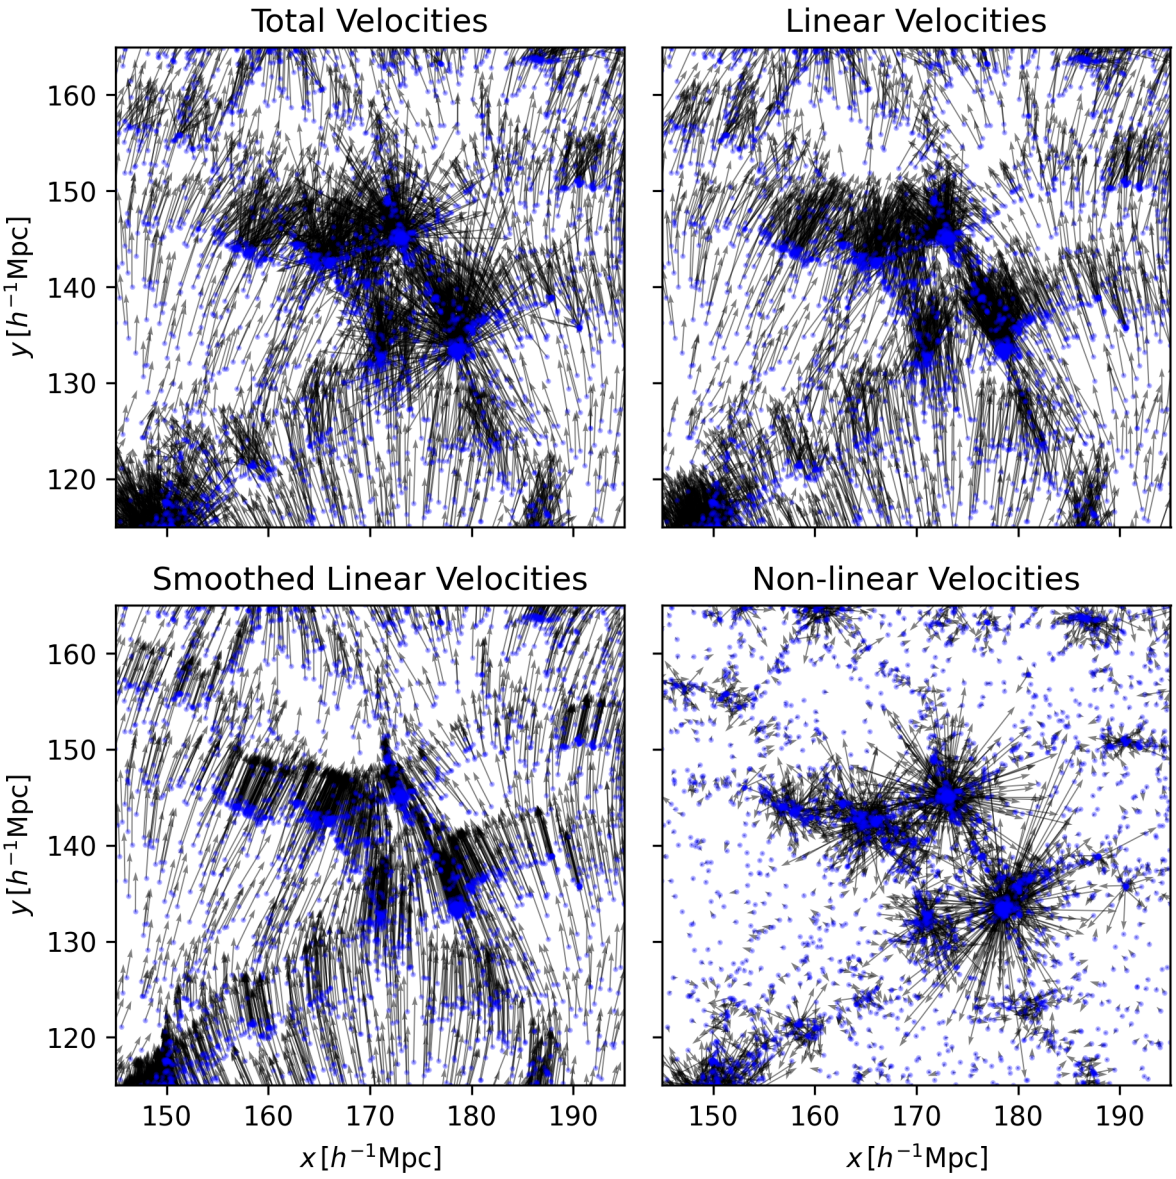

Figure 4.1: A slice of one of the AbacusCosmos *Planck* boxes showing the particle positions and velocities. Blue points show the position of particles from a uniform 10% down sampling, and black arrows show the velocities of the particles where the size of the arrow is proportional to the amplitude of the velocity.

the monopole of the halo correlation function and compare to the results of scaling the unsmoothed halo velocities. For the unsmoothed linear velocity field we define the linear halo velocity as the mean linear velocity of the constituent particles. Scaling both the smoothed and unsmoothed velocities has a nearly identical effect on the large scales of the monopole for both the linear velocity scaling parameter, γ_l, and the non-linear velocity scaling parameter, γ_n. This result is expected since the velocity smoothing primarily affects the pairwise velocity of small separation objects by construction, and desired because the large-scale behaviour follows the expectation from linear theory in that the amplitude of the monopole is proportional to f, and scaling up the velocities increases the amplitude of the correlation function. However, around $\sim 2\,h^{-1}$Mpc scaling up the unsmoothed linear velocities changes behaviour and damps the monopole due to the shell-crossing issue discussed above. Scaling up the smoothed linear velocity increases the amplitude at all scales, although the effect is reduced below the smoothing scale. This matches our desired behaviour for the linear velocity field, which was visualized in Fig. 4.1, that close pairs that have already collapsed maintain their separation as the linear growth rate is increased, rather than being spread apart. When scaling γ_n the effect is similar for both methods of calculating the velocity components, although the smoothed velocity field shows a greater change in amplitude. The quadrupole is not included in this plot because the change in sign makes these trends more difficult to see intuitively, but the same behaviour of the scaling parameters is seen in the quadrupole as displayed in the monopole. The projected correlation function is largely insensitive to the radial velocity by construction, and the difference between smoothed and unsmoothed velocities is insignificant.

Fig. 4.2 also shows the results of varying the parameters used to smooth the linear velocity field. Faint, coloured lines show the effects of smoothing using a tophat radius of $3.0\,h^{-1}$Mpc or $7.0\,h^{-1}$Mpc instead of the default $5.0\,h^{-1}$Mpc, using a grid of side length $2.0\,h^{-1}$Mpc or $0.8\,h^{-1}$Mpc instead of the default $1.0\,h^{-1}$Mpc, and of using a Gaussian kernel with standard deviation $2.0\,h^{-1}$Mpc. In all cases the effect is quite similar to our default choice of parameters at all scales and for both scaling parameters, indicating that our smoothing method is robust to varying these choices.

4.2.5 Visualizing parameter degeneracies

The model we have constructed contains multiple parameters affecting the linear (γ_l, Ω_m, σ_8) and non-linear (γ_n, $v_{\rm bc}$, $v_{\rm bs}$) components of the velocity. In applying the model it is therefore important to understand the degeneracies between these parameters. We can visualize the impact of the different parameters on the correlation functions by using the emulator to make predictions while varying each parameter individually. As a base model,

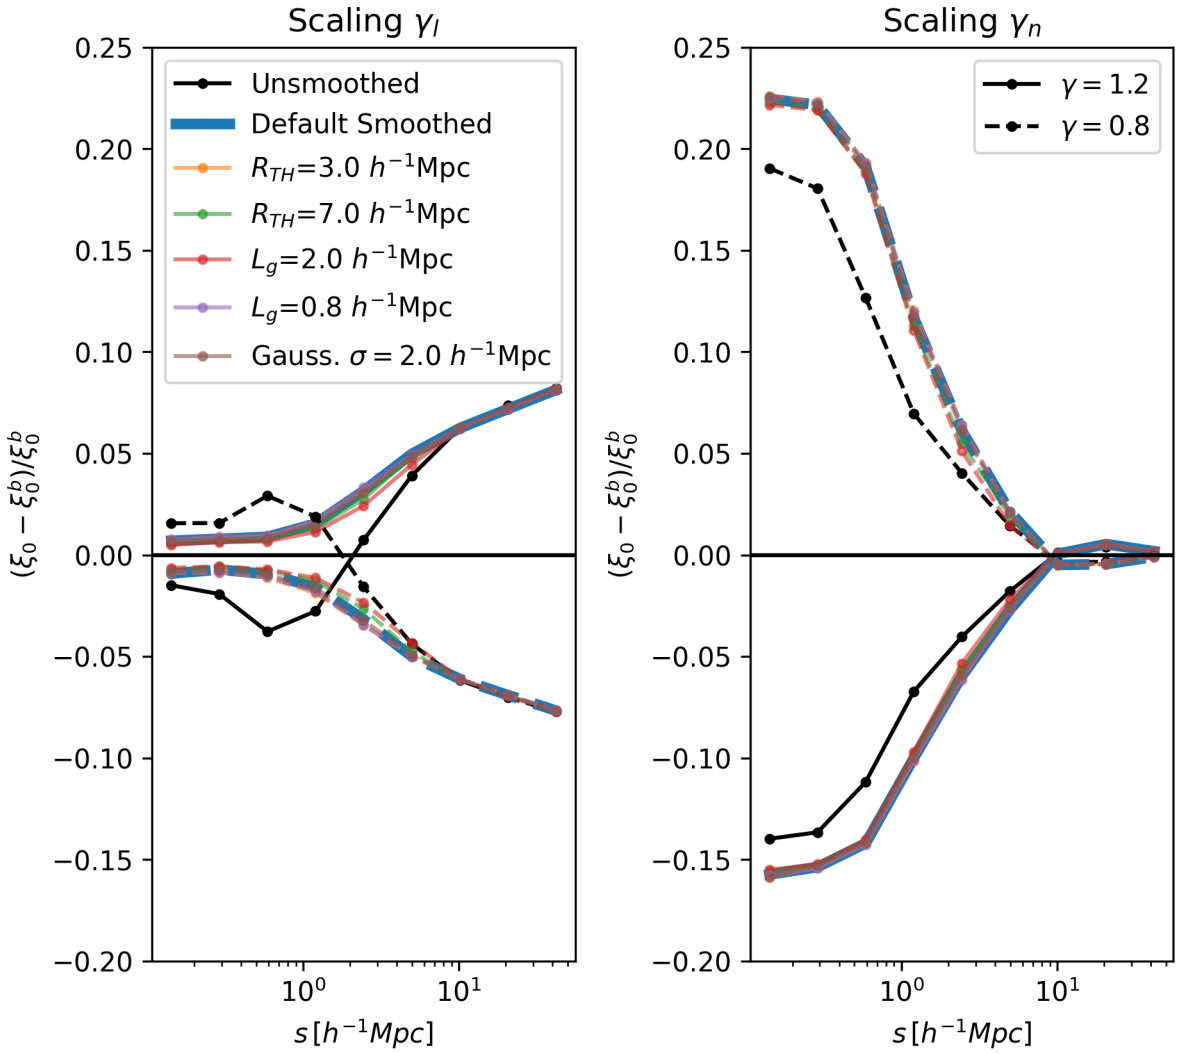

Figure 4.2: The mean change in the monopole of the halo correlation functions after velocity scaling from the 20 *Planck* cosmology boxes. The y-axis shows the relative difference between the correlation function with scaling and the base (superscript b) correlation function without scaling. Solid lines show scaling by $\gamma = 1.2$, while dashed lines show scaling by $\gamma = 0.8$. The black lines show the result using the unsmoothed linear velocity, and the thick blue line shows the result of our fiducial smoothing method; a tophat kernel with radius $5\,h^{-1}$Mpc on a grid of side length $1\,h^{-1}$Mpc. Faint coloured lines show the results of variations on the smoothing method, described in the text.

we use the *Planck* 2018 best fit cosmological parameters (Ω_m, Ω_b, H_0, σ_8, n_s; [213]), with the best fit HOD parameters ($\log M_{\rm sat}$, α, $\log M_{\rm cut}$, $\sigma_{\log M}$, $f_{\rm max}$) from our new analysis of the eBOSS LRG data (see Sec. 4.3.2), and default velocity parameters ($\gamma_l = 1$, $\gamma_n = 1$, $v_{\rm bc} = 0$, $v_{\rm bs} = 1$, $c_{\rm vir} = 1$). We use the core HOD parameters from our fit to the eBOSS data because aspects of the galaxy-halo connection, such as the satellite fraction, will change how the velocity parameters affect the correlation functions. Using the HOD parameters of our fit ensures that we are building the correct intuition for the analysis of the eBOSS LRG sample.

Fig. 4.3 shows the emulator predictions from varying the parameters that directly contribute to the constraint on $f\sigma_8$ in our analysis: γ_l, Ω_m, and σ_8. The amount of variation in each parameter is chosen to give a similar change in amplitude in the monopole on scales $\sim 10\,h^{-1}$Mpc, to make direct comparisons simpler. In general, increasing γ_l or σ_8 increases the amplitude of clustering, while increasing Ω_m decreases the amplitude. The behaviour of γ_l matches the expectations for increased linear growth (see Fig. 4.2 and discussion). While increasing Ω_m or σ_8 will also increase the amplitude of the linear velocity component, both parameters also affect the amplitude of density perturbations, and Ω_m affects the shape of the matter power spectrum. These additional effects can be seen in w_p, which is insensitive to the velocity field. This highlights the difference between γ_l, which affects only the amplitude of the linear velocity, and Ω_m and σ_8, which also affect other aspects of the clustering. In addition to the general behaviour, there are some specific ways the parameters deviate in their effects on ξ_0, ξ_2, and w_p, particularly across separation ranges. While varying γ_l or σ_8 both have a similar impact on the large scales of the monopole, the small-scale behaviour is very different, as are the effects on the quadrupole on all scales. These features highlight the importance of including all three parameters in the emulator.

We perform a similar analysis of the parameters that affect the non-linear velocity component, γ_l, $v_{\rm bc}$, and $v_{\rm bs}$, in Fig. 4.4. For this comparison the amount of variation is tuned to give a similar change in amplitude of the monopole on scales $\sim 1\,h^{-1}$Mpc, within the regime where the non-linear velocity component is strong (see Appendix C). The effect on w_p is omitted because all three parameters purely affect the velocities of galaxies, so any effect from changing the parameters is only an artifact of the finite $r_{\parallel,\rm max} = 80\,h^{-1}$Mpc used in calculating w_p (see Eq. 1.26 in Sec. 1.3.2). γ_n and $v_{\rm bc}$ show very similar effects on the correlation functions, with the exception of the small scales of the quadrupole, implying that they will be highly degenerate in a fit to data. However, it is worth noting that the "default" value of $v_{\rm bc}$ is $v_{\rm bc} = 0$, the case of no central galaxy velocity bias, and that $v_{\rm bc}$ can only be increased and will only ever produce a damping of the correlation function measurements. This is in contrast with values of $\gamma_n < 1$, which reduce the damping of the correlation function, so the degeneracy between γ_n and $v_{\rm bc}$ is only expected to hold

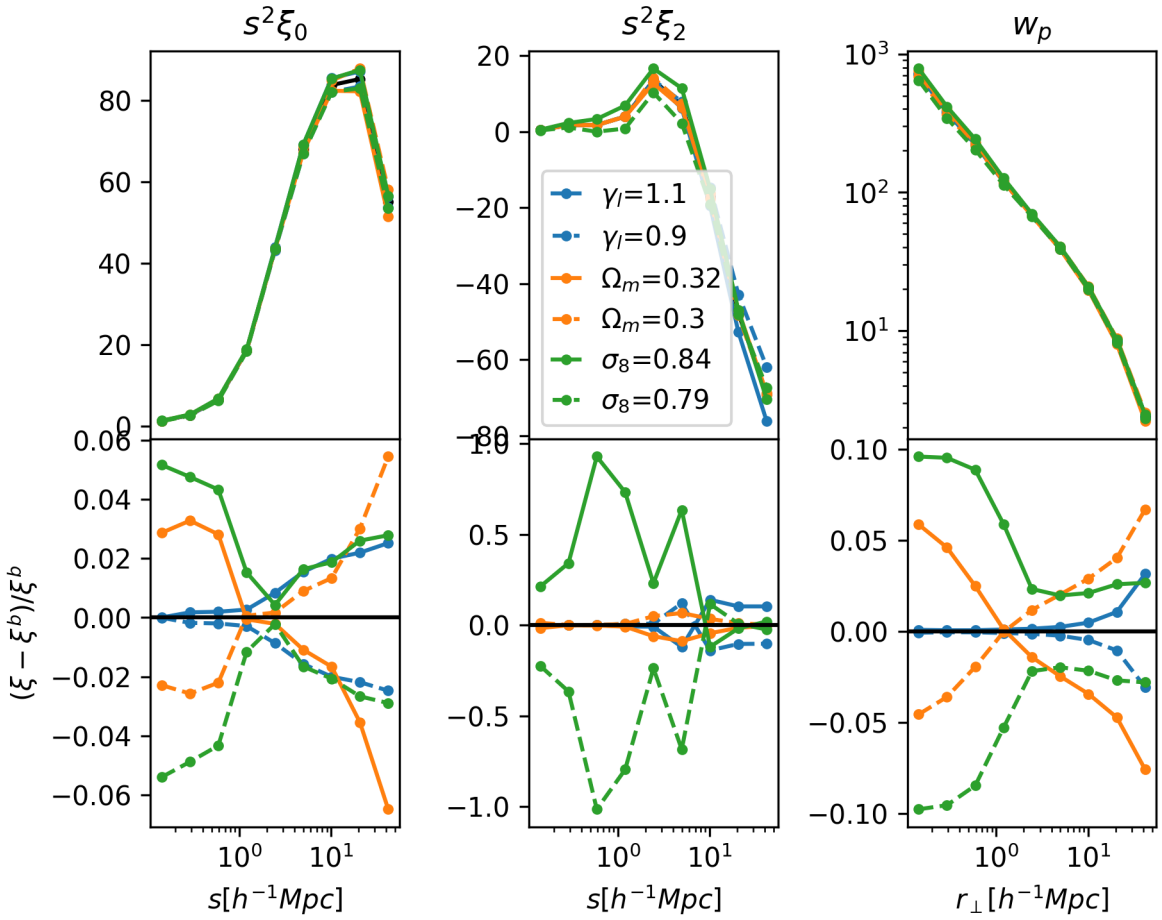

Figure 4.3: Emulator predictions for varying parameters that directly contribute to the constraint on $f\sigma_8$. The black line shows the prediction for the base parameter set consisting of the *Planck* 2018 cosmological parameters [213], best fit HOD parameters for the eBOSS LRG sample from Sec. 4.3.2, and default velocity parameters (see text for full listing). Blue lines show the emulator prediction when γ_l is varied with all other parameters kept fixed, with orange and green lines showing equivalent results for varying Ω_m and σ_8 respectively. The columns shows the predictions for the monopole (left), quadrupole (centre), and projected correlation function (right). Panels in the upper row directly compare the clustering predictions, while the lower row shows the relative difference between the predictions for the varied and base parameter sets.

for $\gamma_n > 1$. We would therefore expect $\gamma_n \neq 1$ to become important in a fit to data when the amplitude of the non-linear component of the velocity in the data is less than the expectation from the ΛCDM+HOD simulations, and $v_{\rm bc} > 0$ to become important if the data is consistent with a velocity bias for central galaxies, leading to an increased Finger-of-God effect in the small scales of the quadrupole. While $v_{\rm bs}$ also affects the damping of the correlation function multipoles through the amount of non-linear velocity, it affects only the velocity of satellite galaxies in the HOD model, and so has a stronger effect than the other parameters on scales $s < 1\,h^{-1}$Mpc, and a weaker affect above that. It is therefore expected to be only weakly degenerate with the other parameters.

4.2.6 Testing the improved emulator

We validate our emulator by performing a MCMC fit to a subsample of the measurements of the *Planck* 2015 boxes used for determining the emulator uncertainty. We randomly select 10 test HOD parameter sets and measure the redshift-space galaxy correlation function for all 20 simulation boxes with line-of-sight along each of the three axes, giving a total of 60 measurements. We average the results of these 60 measurements for each HOD parameter set and fit the data using our improved emulator. To estimate the covariance we use the same data covariance matrix calculated in Sec. 3.3.4, scaled along the diagonal to match the volume of the mock measurements without modifying the correlation structure. While the true effective volume of our measurement will be between 20-60 simulation boxes, because we use 20 independent boxes each measured along three independent lines-of-sight, we conservatively choose a volume of 20 simulation boxes as our fiducial amount.

For all 10 parameter sets we recover the known value of γ_l and the expected value of $f\sigma_8$ to within the 68% credible interval. This is expected given our conservative choices for the emulator uncertainty, which lead to slightly inflated credible intervals while ensuring that our parameter inference is not biased. Likewise, the known cosmological and HOD parameters are recovered for the majority of the parameter sets. The HOD parameters that are least often recovered are $\log M_{\rm cut}$, $\sigma_{\log M}$, and $f_{\rm max}$, however none are degenerate with our key cosmological parameters and there is no significant impact on the $f\sigma_8$ constraints, so there is no concern for our measurement of the eBOSS data.

We also investigate the scale dependence of the constraints from the 10 test HOD parameter sets. For each parameter set we perform a fit to the full separation range of the model, $0.1 - 60\,h^{-1}$Mpc, as well as four additional fits restricted to the separation ranges $0.1 - 7\,h^{-1}$Mpc, $0.8 - 7\,h^{-1}$Mpc, $0.8 - 60\,h^{-1}$Mpc, and $7 - 60\,h^{-1}$Mpc, matching the methodology used to test the data in Sec. 4.3.4. For each set of parameters we find

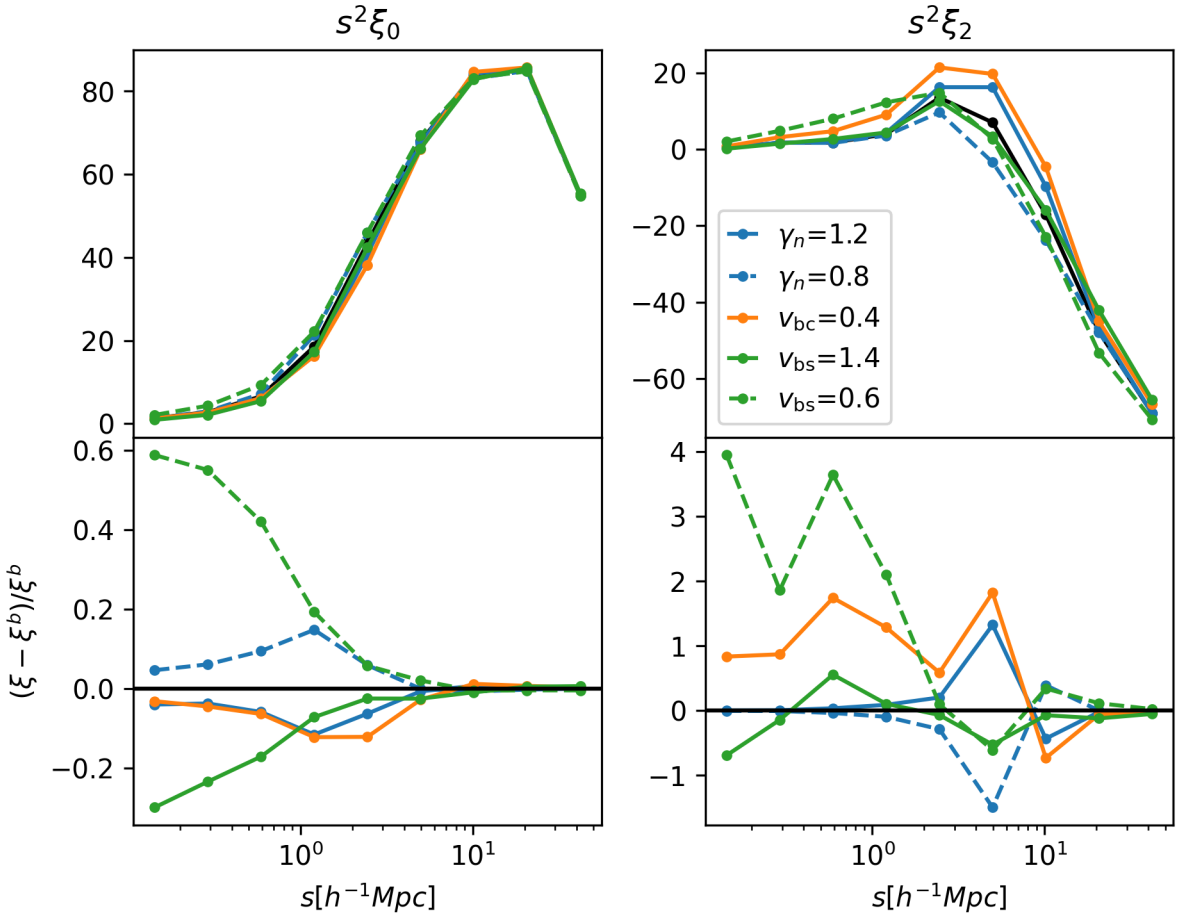

Figure 4.4: Emulator predictions for varying parameters that affect the non-linear velocity component. The black line shows the prediction for the same base parameter set used in Fig. 4.3. Blue, orange, and green lines show the emulator prediction when varying γ_n, $v_{\rm bc}$, and $v_{\rm bs}$ respectively, with all other parameters kept fixed. The monopole is shown on the left and the quadrupole on the right, with the upper row showing a direct comparison, and the lower row showing the relative difference with the prediction for the base parameter set.

all separation ranges give a mutually consistent value of $f\sigma_8$ at the 1σ level, with approximately half of the parameter sets showing a slight offset between the $0.1 - 7\,h^{-1}$Mpc and $0.8 - 7\,h^{-1}$Mpc results (which are in agreement for all parameter sets) and the remaining separation ranges. The offset is equally likely to occur to larger and smaller values and is within the measurement uncertainty, so it is not a concern for our cosmological inference.

Finally, we validate our entire pipeline using the same Uchuu SHAM mock used to test the previous emulator in Sec. 3.4.4. As was the case then, using a different galaxy halo connection model and simulation is a necessary test of the robustness of our model in order to be able to confidently apply it to the eBOSS data. Fitting the correlation function of the SHAM mock using our new emulator we are able to recover the known cosmological parameters within the 68% credible interval for all parameters, and find all well constrained HOD parameters to be within their respective prior ranges. We recover $\gamma_l = 1.00 \pm 0.08$ and $\gamma_n = 0.90 \pm 0.14$, both consistent with their expected values of 1 since the mock was constructed with a ΛCDM growth rate and no velocity scaling. This test and the test using the HOD parameter sets give us confidence that our updated emulator is functioning correctly, and is ready to be applied to real data.

4.3 Measuring the eBOSS LRG RSD

4.3.1 eBOSS LRG sample

We fit our new emulator model to the same eBOSS LRG data (see Ch. 2) analyzed in Ch. 3. Specifically, we use the same measurements of ξ_0, ξ_2, and w_p (see Sec. 1.3.2) binned in nine logarithmically spaced separation bins (see Sec. 3.3.1), and the same jackknife covariance matrix (see Sec. 3.3.4). The only change we make to the data is applying the $v_{\rm match}$ weighting scheme detailed in [189] to the estimation of the jackknife covariance matrix, which corrects the ratio of auto-pairs and cross-pairs removed by the jackknife sampling for a low density of galaxies. Applying these weights we find a minor reduction in the covariance of large separation bins, matching what was seen in [189], although there is no significant change to the correlation between different separation bins.

The reasoning behind using a consistent data set is to be able to give a direct comparison to the results of the original emulator used Ch. 3. In particular, using the same data set but an updated model is useful for identifying systematic biases originating from the modelling, such as the choice of simulations (see Sec. 4.4).

4.3.2 Headline results

We fit our new velocity-split emulator to the small-scale eBOSS correlation functions using a MCMC sampler from the Cobaya package [263]. We apply the same two priors detailed in Sec. 3.3.6, namely a uniform prior over the training range for the HOD parameters, and another uniform prior over the cosmological parameters that limits the parameter-space to be within a threshold distance of the cosmologies of the simulations used for training. The reason for these priors is to limit the parameter values to the range used in training the emulator, as outside of this range the emulator predictions are expected to regress to the mean of the untrained Gaussian distribution. We update the cosmological prior to match the AbacusCosmos cosmologies, rather than the original Aemulus cosmologies.

Our analysis of the eBOSS sample with the new velocity-split emulator yields a value of $f\sigma_8(z = 0.737) = 0.368 \pm 0.041$, representing the posterior mean and 68% credible interval, with a minimum $\chi^2 = 16.6$ at $f\sigma_8 = 0.334$ from 27 data points and 15 free parameters. This value is 2.3-σ below the expectation for a ΛCDM universe with the *Planck* 2018 cosmology, and an increase in tension from the 1.4-σ offset found in Ch. 3. The headline results from Ch. 3 were $f\sigma_8(z = 0.737) = 0.408 \pm 0.038$ when using a single velocity scaling parameter and limiting the measurement scales to $7-60\,h^{-1}$Mpc, so this increase in tension is caused by a shift to a lower value of $f\sigma_8$ rather than an increase in precision, although the two results are mutually consistent.

In Fig. 4.5 we compare the best fit models for various choices of scaling parameters to the eBOSS data. All models are able to accurately fit the data on all scales, although our baseline model of allowing both γ_l and γ_n to vary results in the lowest χ^2 value. The largest improvement in performance over the single scaling parameter model is in the intermediate scales of the monopole and quadrupole. This is likely caused by improved flexibility in simultaneously fitting the smallest and largest measurement scales by decoupling the scaling of the velocity terms, with the non-linear velocity dominating on the smallest scales and the linear velocity dominating on the largest scales. It should be noted that while the fixed $\gamma_l = 1$ model is restricted to match the linear velocity amplitude expected for a ΛCDM universe from the AbacusCosmos simulations, it does not indicate agreement with the value of $f\sigma_8$ expected from the *Planck* 2018 observations because Ω_m and σ_8 are still allowed to vary. That model results in values of $\sigma_8 = 0.792 \pm 0.026$ and $f\sigma_8 = 0.450 \pm 0.015$, with $\chi^2 = 19.8$.

Our fit to the data only weakly constrains the amplitude of the non-linear velocity field, giving a value of $\gamma_n = 0.692 \pm 0.29$, where a value of $\gamma_n = 1$ indicates agreement between the data and the expectation for a ΛCDM universe from the model. This constraint is slightly limited by the lower edge of the prior at $\gamma_n = 0.2$, but does show a clear

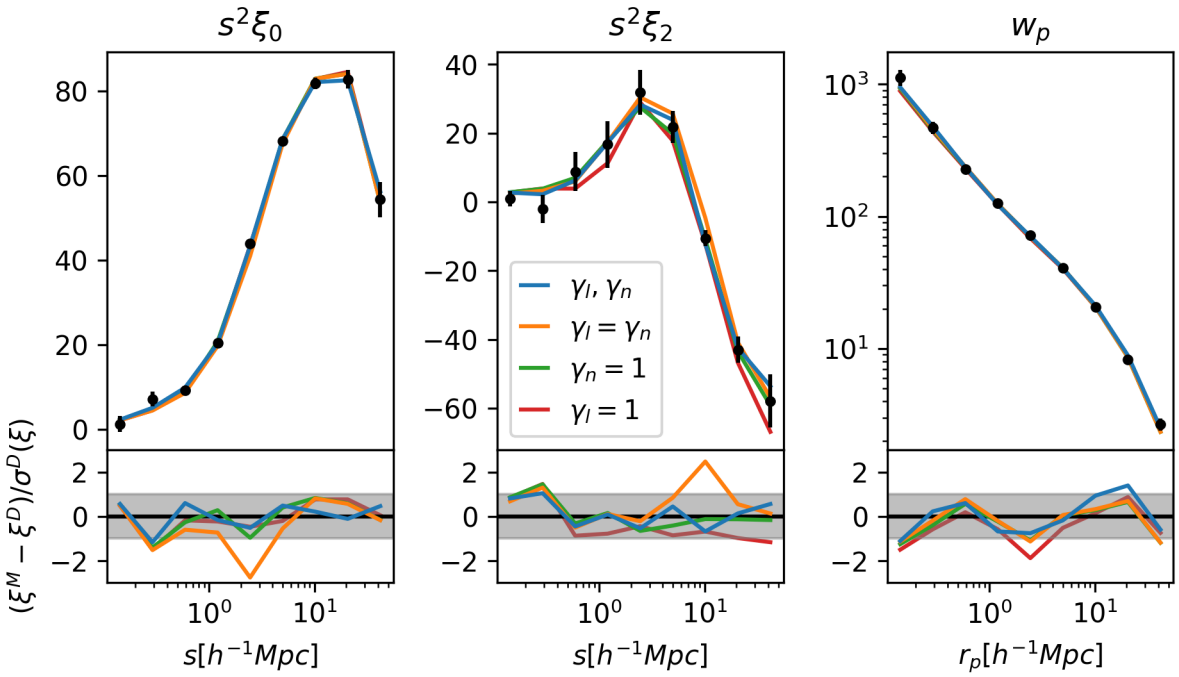

Figure 4.5: Best fit models compared to the eBOSS LRG measurement data for several choices of scaling parameters. Our baseline fit, allowing both γ_l and γ_n to vary, is shown in blue. A single velocity scaling parameter model, constrained so that $\gamma_l = \gamma_n$ and equivalent to the model used in Ch. 3, is shown in orange. The green line shows the result of allowing γ_l to vary while fixing $\gamma_n = 1$, and the red line shows the result of allowing γ_n to vary while fixing $\gamma_l = 1$. The left, centre, and right columns show the monopole, quadrupole, and projected correlation function respectively. The top row of panels directly compares the model to the data, while the lower row shows the difference between model and data in units of the data uncertainty, with the grey-shaded region indicating the 1σ region.

preference for $\gamma_n < 1$. The poor constraint is likely due to the lower magnitude of the non-linear velocity field compared to the linear velocity field (see Appendix C), as well as the degeneracy between γ_n and $v_{\rm bc}$ (Fig. 4.6). It may also indicate that our parameterization of γ_n needs further refinement in order to fully describe the behaviour of the actual non-linear velocity field. We have implemented γ_n as a uniform scaling for all components of the velocity that do not match the initial linear velocity. It is possible that there are multiple contributions to the non-linear velocity, requiring a more nuanced parameterization to capture the deviations between the data and the best fit ΛCDM+HOD model. Non-linear velocity scaling is also not necessarily uniform for all galaxies, and may be dependent on characteristics such as galaxy mass and environment. Investigating these alternatives is a possible avenue for future research. While this result is independent of our cosmological constraint by construction, it does indicate that non-linear velocities in the data are lower than those generated by combining our HOD model with a CDM-only simulation.

Fig. 4.6 shows the 2D and 1D marginalized constraints on all the free parameters of the model described in Sec. 4.2 for our fiducial fit to the eBOSS LRG sample. The parameters are well constrained within their training range and prior, with few exceptions. There is a strong degeneracy between $\sigma_{\log M}$ and $f_{\rm max}$ due to their effects on the occupation of central galaxies, with the posterior occupying a narrow region of the two-dimensional space, but each individual parameter is poorly constrained in one-dimension when the other is marginalized over. $\log M_{\rm cut}$ has a peak above 13, but a long tail towards the lower end of the prior range. This is because $\log M_{\rm cut}$ controls the exponential cutoff of the satellite occupation function, and has no impact when it is significantly smaller than the cutoff of the central occupation, parameterized by $\log M_{\rm min}$, because the satellite occupation is modulated by the central occupation (see Sec. 1.5.3). For the eBOSS LRG sample, $\log M_{\rm min} \approx 13.4$, so $\log M_{\rm cut} < 13$ is unconstrained. $c_{\rm vir}$ is poorly constrained because it is less impactful than other parameters, but does not show any significant degeneracies so is not a concern for our constraints on cosmological parameters. Related to the discussion of degeneracies, it is worth commenting on the derived parameter $f\sigma_8 = \gamma_l f_{\Lambda{\rm CDM}}\sigma_8$. Because $f\sigma_8$ is evaluated from γ_l, Ω_m, and σ_8 at each chain step the constraint on $f\sigma_8$ is correlated with all three parameters. However, because Ω_m and σ_8 affect multiple elements of the clustering, and γ_l is constructed to affect only the portion of the RSD effect caused by the amplitude of the linear velocity component, the strongest correlation is with γ_l, as shown by the $f\sigma_8$-γ_l panel of Fig. 4.6. That panel shows a narrow diagonal ellipse, with the width of the ellipse in the marginalized parameter constraint coming from variance in σ_8 and Ω_m at fixed γ_l.

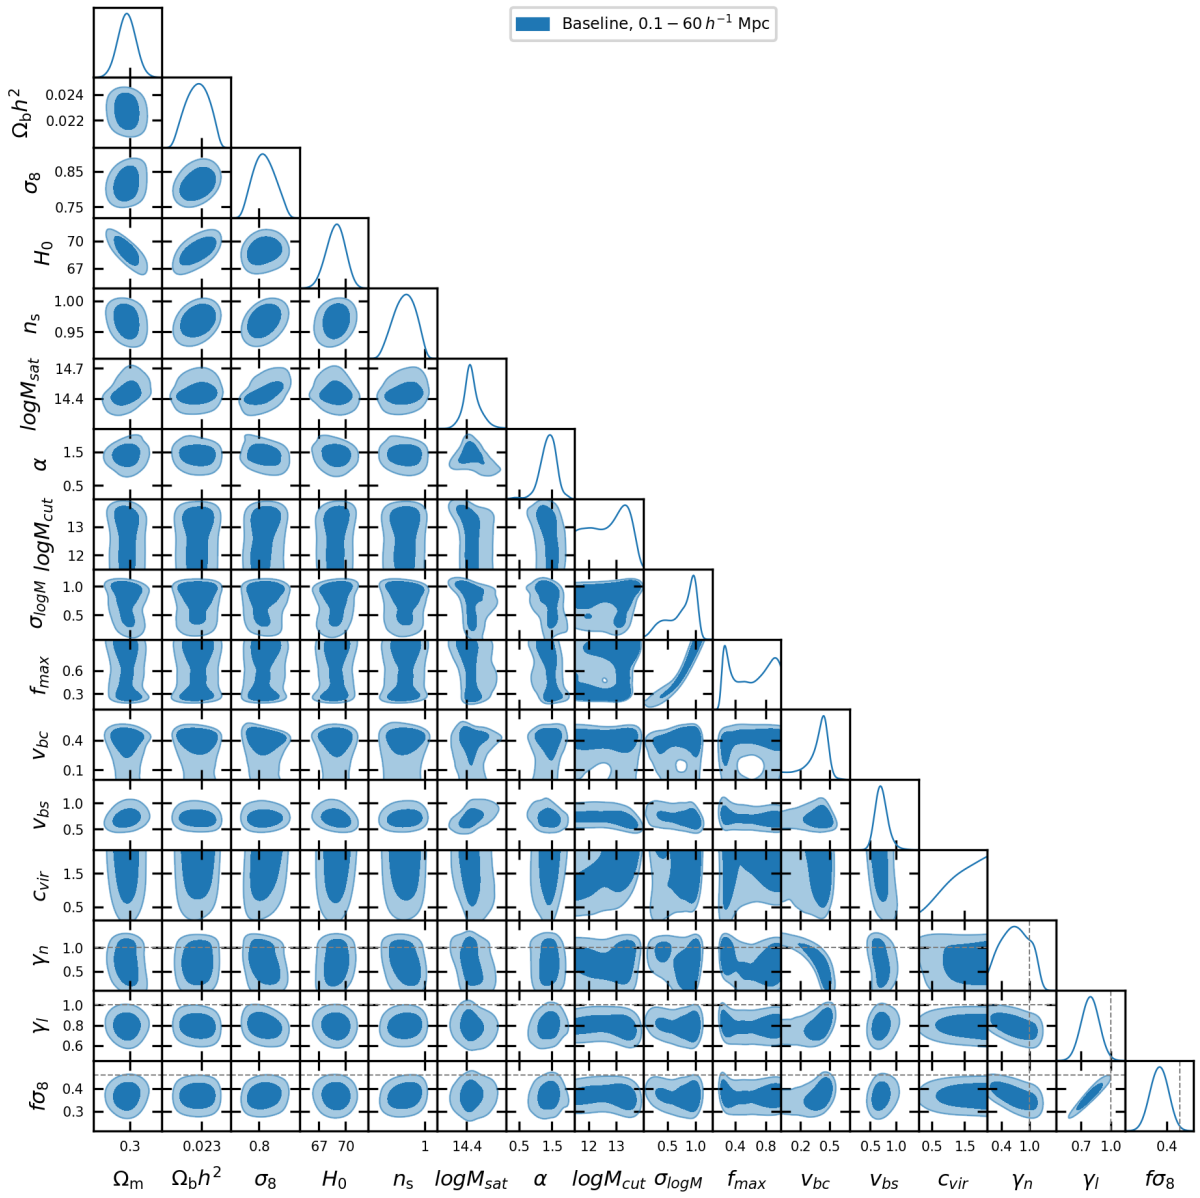

Figure 4.6: 1D and 2D contours of the parameters used in our baseline fit of the eBOSS LRG. The constraint on $f\sigma_8$ is calculated as $f\sigma_8 = \gamma_l f_{\Lambda\mathrm{CDM}}\sigma_8$, where for each chain point σ_8 and $f_{\Lambda\mathrm{CDM}}$ are calculated using the model parameters for that point (see Sec. 4.2.1). The dashed lines highlight $\gamma_l = 1$ and $\gamma_n = 1$, which would indicate that no velocity scaling is needed to match the data to the ΛCDM expectation of the emulator.

γ_l	Ω_m	σ_8	$f\sigma_8$	χ^2
Headline results				
0.795	0.297	0.815	0.368	17.0
2σ variations				
0.619	0.297	0.815	**0.286**	**20.2**
0.630	**0.253**	0.815	**0.286**	**32.3**
0.681	0.297	**0.740**	**0.286**	**21.3**
0.690	**0.276**	**0.736**	**0.286**	**25.0**
0.970	0.297	0.815	**0.449**	**19.1**
0.951	**0.361**	0.815	**0.449**	**44.2**
0.884	0.297	**0.894**	**0.449**	**20.4**
0.869	**0.329**	**0.899**	**0.449**	**24.6**
3σ variations				
0.531	0.297	0.815	**0.246**	**25.2**
1.058	0.297	0.815	**0.489**	**22.8**

Table 4.1: Best fit χ^2 values for models with various $f\sigma_8$ values, obtained by fixing the values of γ_l, Ω_m, and σ_8. The first section uses the marginalized values from the full fit to the eBOSS LRG sample described in Sec. 4.3.2. Later sections show the results for fits where the values of γ_l, Ω_m, and σ_8 have been fixed to give a particular value of $f\sigma_8$ that deviates from the headline results by varying numbers of standard deviations. Parameter values that are changed from the headline results are displayed in boldface.

4.3.3 Testing the credible intervals

The uncertainty assigned to the value of $f\sigma_8$ derived from our analysis is just as important as the value itself. For that reason we test our uncertainties, which are determined as credible intervals from the posterior of the fit to the data, by investigating the quality of fit we obtain when we restrict the model to a value of $f\sigma_8$ at multiples of our uncertainty away from the posterior mean value. For each value of $f\sigma_8$ we fix the values of γ_l, σ_8, and Ω_m, which collectively specify the value of $f\sigma_8$, and perform a new fit to the eBOSS LRG data allowing the remaining parameters of the model to be free. We then select the best fit model from that MCMC fit, and report the χ^2 between the emulator prediction and the data. The details of the additional fits we perform are given in Table 4.1.

As a baseline for comparison we fix γ_l, σ_8, and Ω_m individually to their posterior mean values when all other parameters are marginalized over (i.e., for the value of γ_l, σ_8 and Ω_m are also marginalized over). This gives the same value of $f\sigma_8$ reported in the headline

results. We then vary that value of $f\sigma_8$ higher and lower in increments of 2σ and 3σ, as determined by the uncertainty reported in our headline results.

We are limited in how we can vary the parameters by the range used in training the emulator, as outside of this range the emulator predictions are expected to regress to the mean of the untrained Gaussian distribution. This restriction is more severe for Ω_m and σ_8, where the range is determined by the cosmologies of the AbacusCosmos simulations that correspond to the 4σ constraints [110] of the *Planck* 2013 release [209]. To investigate variations in Ω_m and σ_8 we therefore find the minimum and maximum values of $f\sigma_8$ that can be obtained within the prior while varying Ω_m and keeping σ_8 fixed to the headline results, fixing Ω_m and varying σ_8, and varying both. It is worth noting that while these values represent the extremes of our prior range, they are all excluded at high confidence by other cosmological probes such as CMB [213], weak lensing [3, 18], and BAO measurements [10]. In all cases these variations are insufficient to reach a 2σ change in $f\sigma_8$, due to the relatively tight parameter constraints of *Planck*, and because the uncertainty in our measured value of $f\sigma$ is dominated by the uncertainty on γ_l. We therefore vary γ_l as well in each case to give the correct value of $f\sigma_8$. Due to the limits imposed by the training prior we investigate varying $\gamma_l + \Omega_m$, $\gamma_l + \sigma_8$, and $\gamma_l + \Omega_m + \sigma_8$ for the 2σ variations only, while we investigate varying γ_l alone for both 2σ and 3σ variations.

The results of the fits to these models are reported in Table 4.1. As expected, the minimum χ^2 is obtained for the headline results. It is worth noting that the χ^2 value reported here is higher than in Sec. 4.3.2. That is because the value of $f\sigma_8$ used here and reported in the results is the posterior mean, while the maximum likelihood value is at a slightly lower value of $f\sigma_8$, which is within the 68% credible interval. Table 4.1 shows that variations in $f\sigma_8$ by increasing multiples of the uncertainty lead to progressively larger values of χ^2. These χ^2 differences correspond roughly to the expectation for the 95% and 99.7% confidence intervals for a one parameter Gaussian distribution, $\Delta\chi^2_{95\%} = 2.71$ and $\Delta\chi^2_{99.7\%} = 9.00$ [265]. Varying only γ_l gives lower values of χ^2 than varying $\gamma_l + \sigma_8$, $\gamma_l + \Omega_m + \sigma_8$, or $\gamma_l + \Omega_m$, which increase in that order. This is reflective of the relative uncertainties of each parameter in our fit to the data, with Ω_m being the most strongly constrained and γ_l being the weakest. It is also worth noting that for a given set of parameters the χ^2 is generally smaller for variations to higher values of $f\sigma_8$ than it is for lower values of $f\sigma_8$. This difference matches the slight skewness in our constraint on $f\sigma_8$, visible in Fig. 4.6. However, because the 68% credible interval is symmetric around the posterior mean to within the number of significant figures included in our result, we report a single symmetric uncertainty on our final result.

In order to visualize the ability of the different parameter sets to fit the data, we plot the best fit models from the 2σ variations to lower values of $f\sigma_8$ in Fig. 4.7, and the 2σ

Measurement Scales	Ch. 3	This Chapter
$0.1 - 7\,h^{-1}$Mpc	0.334 ± 0.061	0.335 ± 0.105
$0.8 - 60\,h^{-1}$Mpc	0.373 ± 0.031	0.368 ± 0.041
$7 - 60\,h^{-1}$Mpc	0.408 ± 0.038	0.412 ± 0.048
$0.1 - 60\,h^{-1}$Mpc	0.365 ± 0.025	0.368 ± 0.041

Table 4.2: Comparison of $f\sigma_8$ constraints from different scales between the velocity-split emulator and a single velocity scaling parameter emulator. In each analysis the constraint is calculated as $f\sigma_8 = \gamma_{(f,l)} f_{\Lambda\text{CDM}} \sigma_8$ and marginalized over all other parameters, including marginalizing over γ_n.

variations to higher values of $f\sigma_8$ in Fig. 4.8. There are some common trends in both plots. The largest deviation between the headline results model and the varied $f\sigma_8$ models occurs at the large scales of the quadrupole, matching the results of Fig. 4.5 and reflecting the importance of that measurement to the $f\sigma_8$ constraint. The deviation is smaller for the high values of $f\sigma_8$ than the small values, which is a major source of the lower χ^2 for high values of $f\sigma_8$ than low values. Another common feature is that models that vary σ_8 and Ω_m provide a poorer fit to the projected correlation function on large scales. This result is consistent with the behaviour of Fig. 4.3, which showed that σ_8 and Ω_m have a significant impact on w_p, unlike γ_l. Overall, the results of these plots and the χ^2 values shown in Table 4.1 give us confidence that the mean posterior value and 68% credible interval reported in Sec. 4.3.2 are a reasonable representation of the constraint of the data on $f\sigma_8$.

4.3.4 Testing the dependence on the data fitted

A key motivating factor for constructing our new velocity-split model was the scale dependence observed when using a single velocity scaling parameter in Ch. 3. In that analysis we found that fitting to various measurement scales found lower values of $f\sigma_8$ at smaller scales, although all measurements were consistent with each other and below the expectation for a ΛCDM universe with *Planck* 2018 cosmology. Using our updated emulator we find the smallest measurement scales to be in better agreement with the larger scales of our analysis. A small offset still exists between the quasi-linear scales and transition scales, as shown in Fig. 4.9. A comparison of the constraints on $f\sigma_8$ using various measurement scales between the new emulator and the result of the single velocity scaling parameter emulator used in Ch. 3 is shown in Table 4.2.

This result follows our expectation for splitting the velocity parameters into γ_n and γ_l.

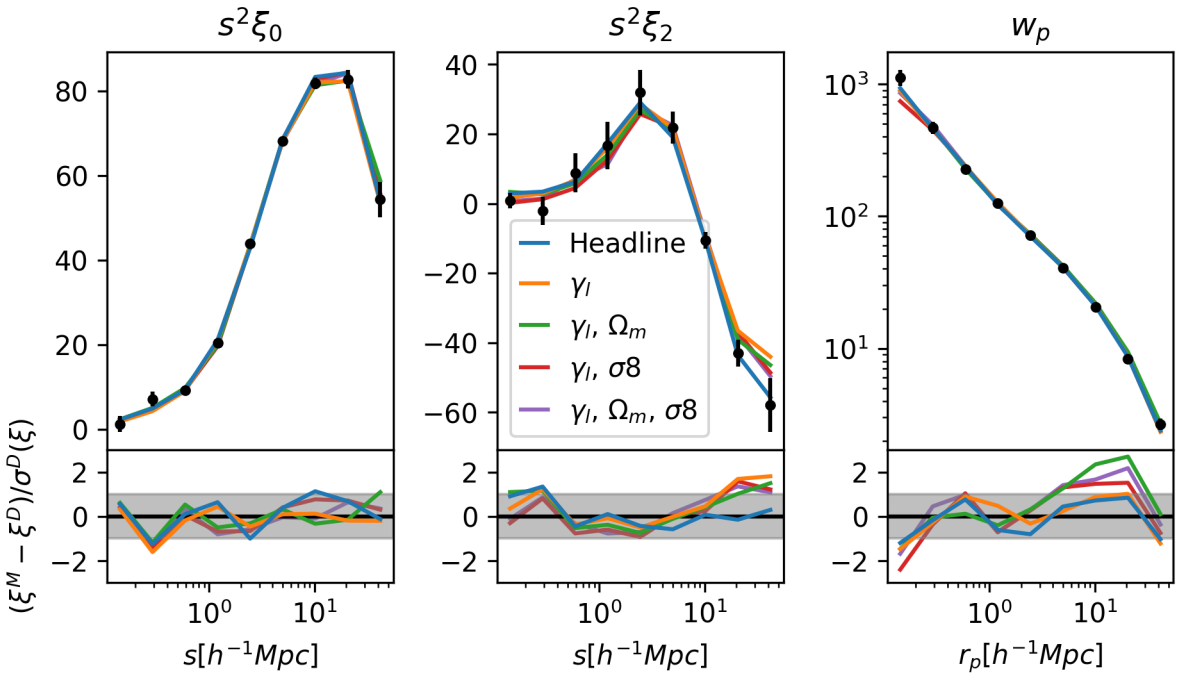

Figure 4.7: Best fit models compared to the eBOSS LRG measurement data for several choices of fixed $f\sigma_8$. In each case γ_l, Ω_m, and σ are fixed to give a specific value of $f\sigma_8$, as described in Table 4.1. The blue line shows the best fit model when $f\sigma_8$ is fixed to match our headline results, while the other coloured lines show the results when $f\sigma_8$ is shifted 2σ lower. The different colours indicate the parameters that have been changed to give the new value of $f\sigma_8$. The left, centre, and right columns show the monopole, quadrupole, and projected correlation function respectively. The top row of panels directly compares the model to the data, while the lower row shows the difference between model and data in units of the data uncertainty, with the grey-shaded region indicating the 1σ region.

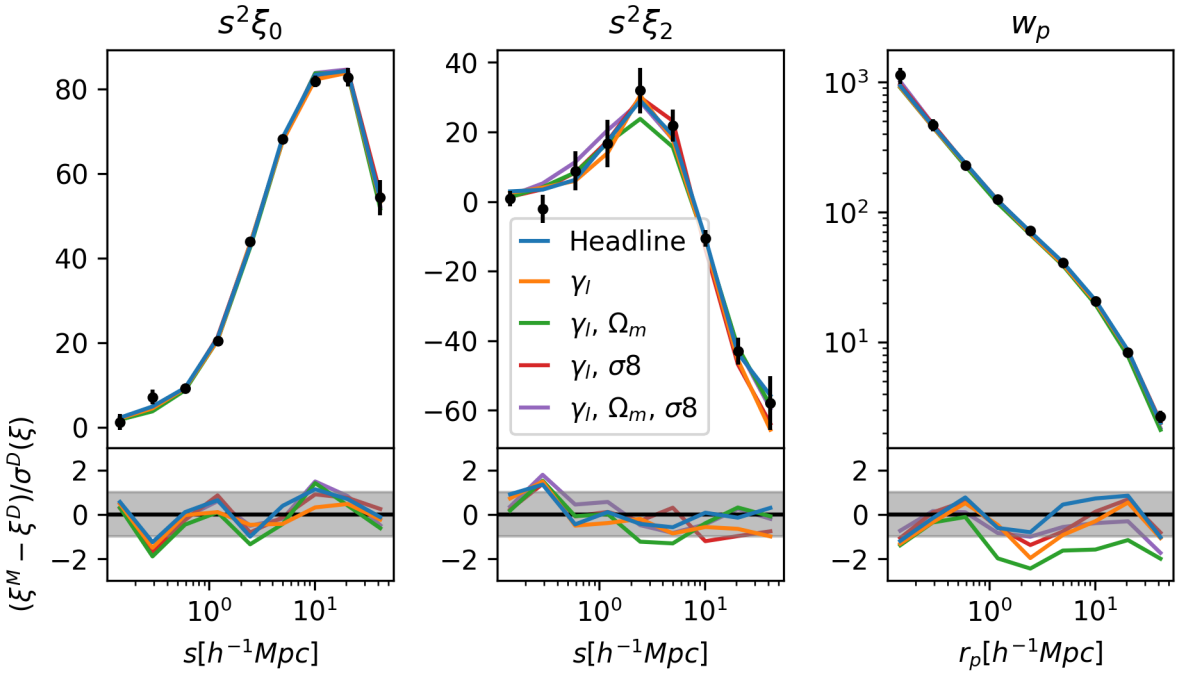

Figure 4.8: Same as Fig. 4.7, but when $f\sigma_8$ is shifted 2σ higher.

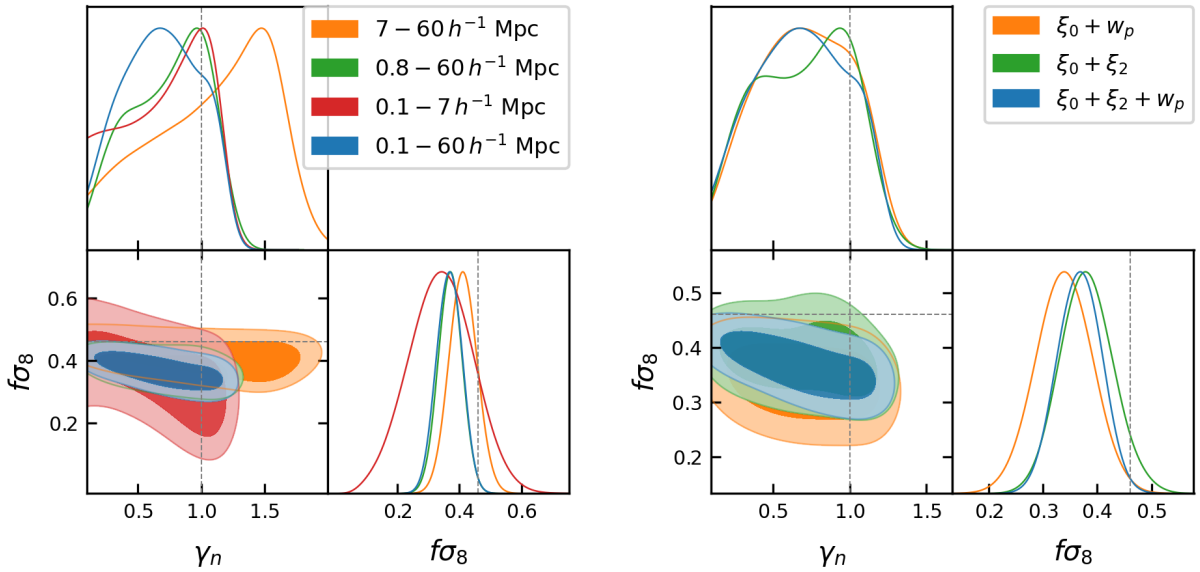

Figure 4.9: 2D and 1D marginalized constraints on γ_n and $f\sigma_8$ for fits to different scales and measurements. *Left:* Constraints from the three largest separation bins (orange), six largest separation bins (green), six smallest separation bins (red), and all nine separation bins (blue) for all three measurements. The dotted line shows the value of $f\sigma_8$ expected from the *Planck* 2018 results assuming a ΛCDM cosmological model. *Right:* Constraints from the joint fit to the monopole and projected correlation function (orange), monopole and quadrupole (green), and all three measurements (blue).

γ_n is more important on the smallest scales, which are fit to the lowest velocity amplitude. Introducing an additional degree of freedom for the non-linear velocities through γ_n reduces the tension in γ_l, but leaves the constraints from the scales around the transition from non-linear to quasi-linear ($\sim 0.8 - 7h^{-1}$Mpc) and quasi-linear scales largely unaffected. Our lower overall value of $f\sigma_8$ from the new analysis is caused by the inclusion of these transition scales, which also preferred a low value of $f\sigma_8$ in Ch. 3, but could not be definitively attributed to the linear signal until the introduction of our new model.

While including γ_n and γ_l has not fully removed the scale dependence of our measurement, it has significantly improved the agreement of different measurement probes, as shown in the right panel of Fig. 4.9. The results of fitting to the multipoles alone, the monopole with projected correlation function, and all three measurements are in close agreement. This result is a significant improvement over the result of Ch. 3, which found some tension between the different measurements due to the degeneracy between the combined velocity scaling parameter and $v_{\rm bc}$.

4.3.5 Comparison to previous emulator

In Fig. 4.10 we compare the constraints of key parameters between our new velocity-split model and the single velocity scaling parameter model used in Ch. 3. To ensure an accurate comparison between the two methods we produce a new fit using our current emulator by setting $\gamma_l = \gamma_n$, which is equivalent to scaling all velocities by a single value. We find that all parameters are consistent between the two methods, with the most significant differences occurring in the velocity parameters γ_n, γ_l and $v_{\rm bc}$, as expected. The new method slightly increases the uncertainty on γ_l, which follows through to the $f\sigma_8$ constraint, since splitting the velocity scaling parameters causes γ_l to have a smaller impact on the fit. This is particularly true on small scales. γ_n and $v_{\rm bc}$ show significant degeneracy since they both contribute dispersive components to the galaxy velocity, and there are lesser degeneracies between γ_l and γ_n, and γ_l and $v_{\rm bc}$.

4.4 Discussion and conclusions

Using an emulator-based model with individual scaling parameters for the linear velocity, γ_l, and non-linear velocity, γ_n, we measure $f\sigma_8(z = 0.737) = 0.368 \pm 0.041$ from clustering between $0.1 - 60\,h^{-1}$Mpc. In Ch. 3 we measured the same sample using an emulator with a single parameter scaling for the total velocity, but restricted our range of measurement

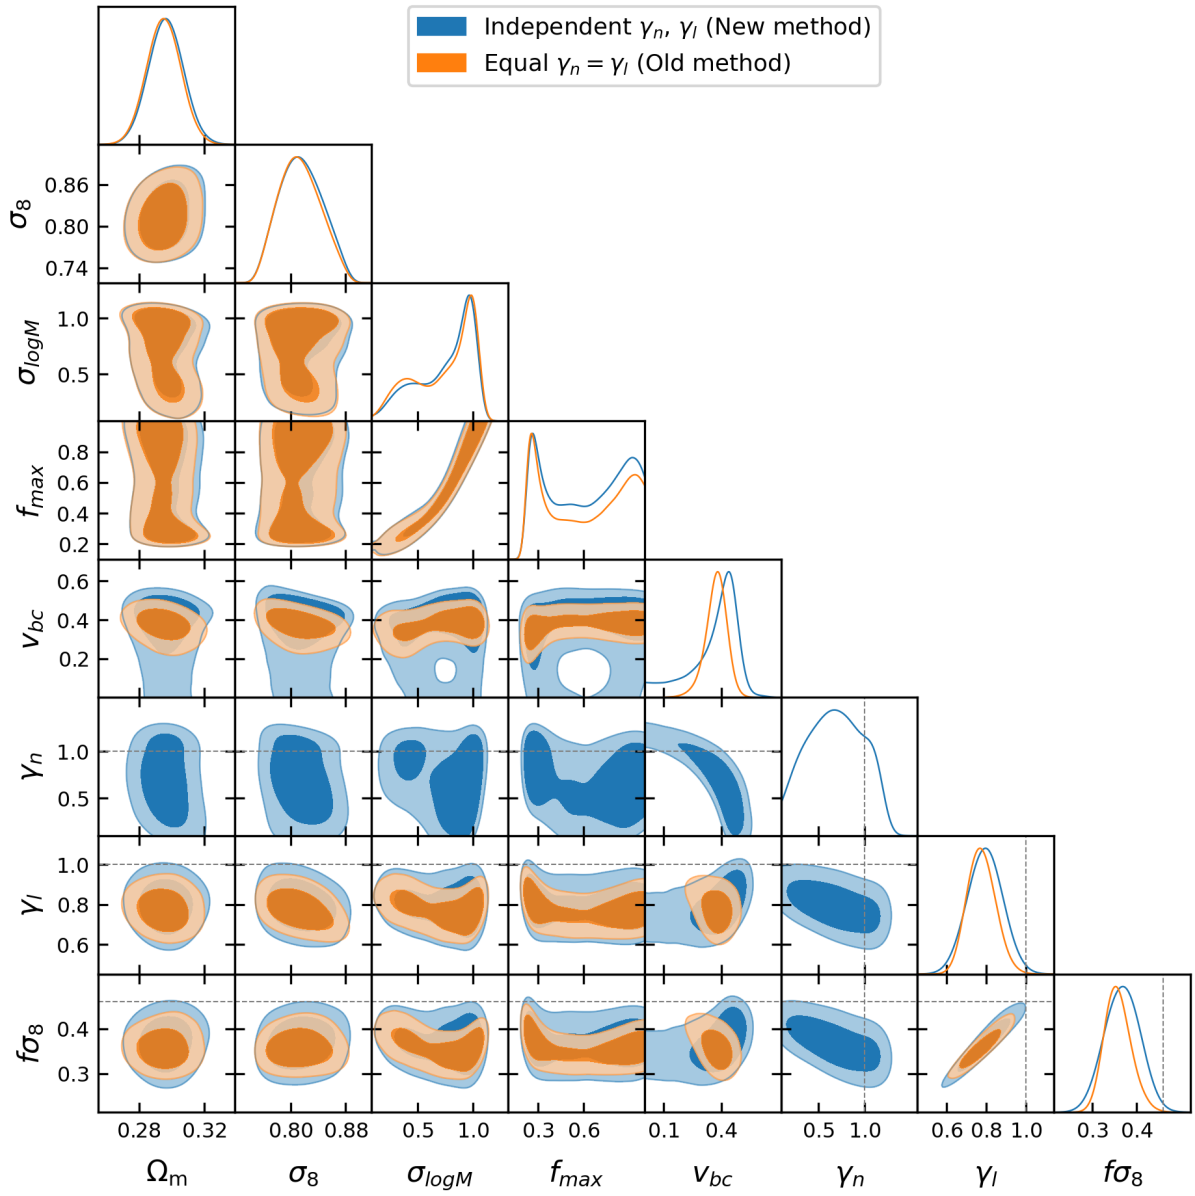

Figure 4.10: Marginalized constraints of several key parameters from the eBOSS LRG data using the independent velocity-split scaling parameters introduced in this chapter (blue) compared to the results from a single scaling parameter (orange). Both fits were made using the updated emulator described in Sec. 4.2. For the single scaling parameter fit we constrain $\gamma_n = \gamma_l$ to mimic the effect of a single scaling parameter.

to $0.7 - 60\,h^{-1}$Mpc in order to isolate the linear signal, and found $f\sigma_8(z = 0.737) = 0.408 \pm 0.038$. The shift to lower values in the updated emulator is caused by the inclusion of smaller scale clustering, and is very similar to the measurement from the same scales using the older emulator.

The consistency of the $f\sigma_8$ constraints between the two models gives confidence that our cosmological constraint is robust to the form of the velocity scaling. The advantage of the new model is that by isolating the linear signal we can now confidently extend our fitted data to small scales, which gives an increased tension with the expectation from *Planck*+ΛCDM. By splitting the velocity scaling parameter to isolate the linear signal we can identify where the information for our constraint comes from, and be sure that we are optimally extracting linear information from the small-scale RSD signal without contamination from non-linear structure growth. This theory is borne out by the consistency between the results of the two emulators given the difference in modelling choices, which indicates that the non-linear velocities are not significantly affecting the linear measurements. Therefore, the most significant advancement of the new emulator is removing non-linear contamination as a potential source of systematic. In addition to the change in parameters, the older emulator was trained on the Aemulus simulation suite while the updated emulator was trained on the AbacusCosmos suite. The consistency between two different simulation suites, generated using different codes, indicates the reliability of the training data. Combined, these factors place severe limits on potential systematic biases in the analysis that could produce the low value of $f\sigma_8$ found from the data.

The results of both emulators, as well as other measurements of $f\sigma_8$ from SDSS galaxy samples, are shown in Fig. 4.11. Our result is still consistent with the large-scale analysis of the same sample at the 1σ level, but is now in 2.3σ tension with the expectation for a ΛCDM universe with a *Planck* 2018 cosmology. There also remains a consistent trend in small-scale RSD measurements to lower values of $f\sigma_8$. This trend is now remarkable when considering the differences in modelling, data, and simulations between these analyses.

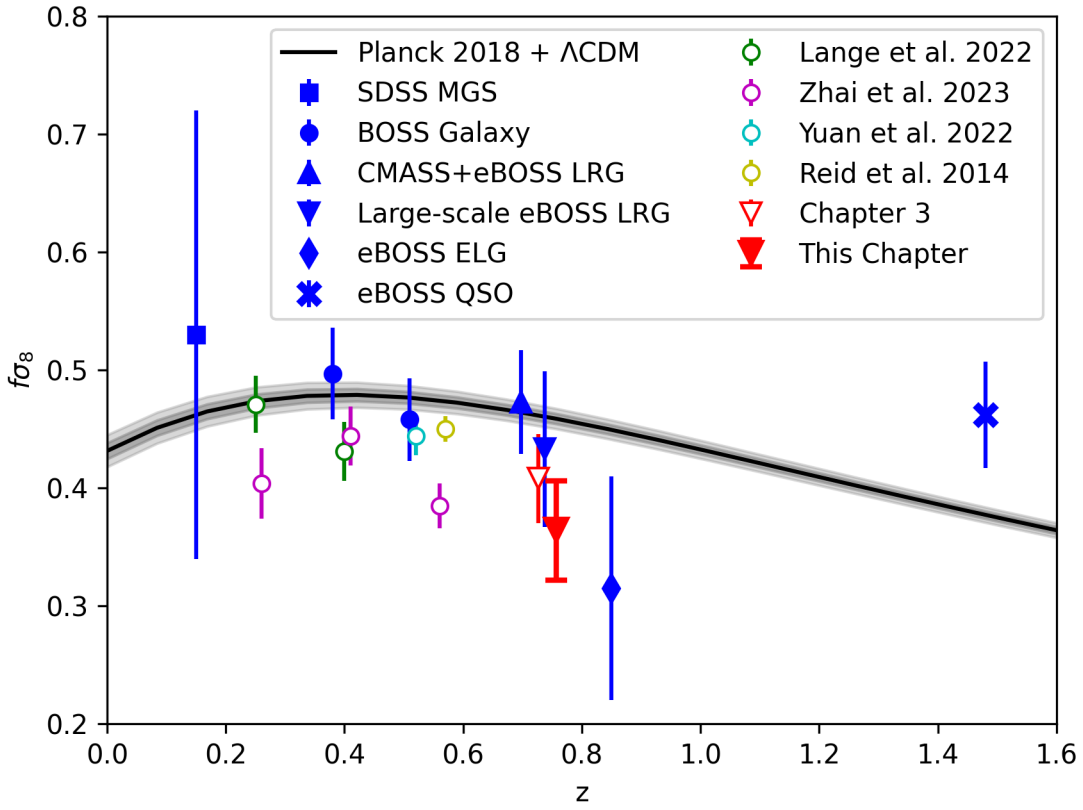

Figure 4.11: $f\sigma_8$ constraints from RSD measurements of SDSS samples. The blue points show the results of more conventional large-scale analyses from the SDSS MGS [143], BOSS galaxies [9], CMASS+eBOSS LRGs and eBOSS LRGs only [23], eBOSS ELGs [76], and eBOSS quasars [196]. The results of this chapter are shown as the red solid triangle, while the red empty triangle shows the results from Ch. 3 using only the separation range $7-60\,h^{-1}$Mpc. Other coloured points show the results from various small-scale analyses of the BOSS galaxy samples that do not distinguish between linear and non-linear information ([166] in green, [277] in magenta, [275] in cyan, and [226] in yellow). The black line shows the expected value of $f\sigma_8$ for a flat ΛCDM universe with best fit *Planck* 2018 cosmology, with the shaded regions showing the 1 and 2σ confidence regions. Measurements from the same galaxy sample are shifted slightly in the x-axis to avoid overlap.

Chapter 5

Conclusions

This chapter summarizes the results of this thesis and provides the main conclusions. The majority of this chapter was written by me for this thesis, with some passages adapted from the Abstracts, Discussions, and Conclusions of my two lead-author papers [57, 58]. In all cases the original content was written by me with input from my co-authors, and has been reformatted and rewritten by me for this thesis.

This thesis has explored a number of ways the Extended Baryon Oscillation Spectroscopic Survey (eBOSS) data has improved our understanding of the standard cosmological model, with particular focus on optimizing the information gained from Redshift Space Distortion (RSD) measurements and correcting systematic biases to give robust results. In this chapter I will summarize the work presented in this thesis, and discuss the future outlook of this research.

In Ch. 1 I introduced the standard model of cosmology, the observations used to identify and constrain that model, and the tensions between measurements that give evidence for extensions to the model. I presented the growth of structure, its effect on galaxy spectroscopic surveys through RSD, and the measurements and models that are used to analyze the effect. These tools are instrumental to the analyses presented in later chapters, which in turn produce significant insights into the nature of the cosmological model and our Universe.

Ch. 2 presented the Extended Baryon Oscillation Spectroscopic Survey from the initial survey design to the final cosmological measurements. An important aspect of this process is identifying and removing observational biases from the sample. I described our process for generating Pairwise-Inverse-Probability (PIP) weights and applying them in tandem with Angular Upweighting (ANG) to remove the bias caused by fibre collisions

in the large-scale structure catalogues. We test these weights using a sample of 1000 EZmocks, and are able to recover unbiased clustering on all scales for all of the eBOSS samples. I also described an observational systematic in the Emission Line Galaxy (ELG) sample caused by variable depth in the photometric catalogues used for targeting, and presented my method for correcting this systematic without removing any data by using a weight-based correction. This weight-based correction gives important insight into the effects of this systematic on the ELG sample, which informed the final approach of nulling the angular modes applied in the cosmological analysis of the sample. After addressing the observational systematics I presented the final cosmological analyses of the LRG, ELG, quasar, and Lyα forest samples. These samples provide stringent constraints on the expansion history of the Universe by measuring $D_M(z)/r_d$ and $D_H(z)/r_d$ using the BAO scale, and the growth history of the Universe by measuring $f\sigma_8$ from RSD. The three clustering samples collectively span the redshift range $0.6 < z < 2.2$, with the Lyα forest measurements providing additional higher redshift information from $z > 2.1$ quasars, opening a new redshift regime for galaxy spectroscopic surveys. When combined with other SDSS samples these constraints become even stronger, and severely limit extensions to the ΛCDM model. Combined with CMB and supernovae observations, the BAO measurements yield $\Omega_k = -0.0001 \pm 0.0018$ (for oΛCDM) and $w = -1.026 \pm 0.033$ (for wCDM), both consistent with a flat ΛCDM model at high precision [11].

In Ch. 3 I shifted from the large-scale analysis of eBOSS to the small scales, presenting our analysis measuring small-scale RSD within clustering of the eBOSS LRG sample [57]. We use the PIP+ANG correction to remove the effects of fibre collisions and obtain unbiased clustering measurements on all scales. We fit to the monopole and quadrupole moments of the 3D correlation function and to the projected correlation function over the separation range $7 - 60\,h^{-1}$Mpc with a model based on the AEMULUS cosmological emulator to measure $f\sigma_8$. We measure $f\sigma_8(z = 0.737) = 0.408 \pm 0.038$, which is 1.4σ lower than the value expected from 2018 *Planck* data for a flat ΛCDM model, with a factor of 1.7 improvement in precision of the large-scale analysis using more conventional techniques of the same sample. We also fit to the data using the full range of scales modelled by the AEMULUS emulator, $0.1 - 60\,h^{-1}$Mpc, 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.

In Ch. 4 I introduced our new model to isolate the linear signal when making small-scale RSD measurements [58]. Isolating the linear signal enables us to avoid bias from misinterpreting the effect of linear growth on non-linear scales while still accessing the full range of the emulator. We construct a new emulator with scaling parameters for both the linear and non-linear velocities of galaxies, allowing us to isolate the linear growth rate.

We train the emulator using simulations from the AbacusCosmos suite, determining the linear velocity of the dark matter halos from the evolved and smoothed velocity field of the initial conditions. We apply the new emulator to the same eBOSS LRG data, and obtain a value of $f\sigma_8(z = 0.737) = 0.368 \pm 0.041$, yielding an increased 2.3-σ tension with the *Planck* 2018 ΛCDM expectation. We also find less dependence on the minimum measurement scale than the analysis presented in Ch. 3, indicating that our model works as intended and that misinterpreting non-linear velocities is not the cause of the low value of $f\sigma_8$ obtained from these measurements.

Our small- and large-scale eBOSS analyses also give valuable insight into current tensions between observations within the ΛCDM framework. The combined SDSS BAO measurements are consistent on their own with the values of H_0 determined using both early-Universe observations, such as the CMB, and late-Universe observations, such as the local distance ladder, due to the freedom of r_d within the constraints on $D_M(z)/r_d$ and $D_H(z)/r_d$. However, when r_d is calibrated independently using Big Bang Nucleosynthesis (BBN) observations, the SDSS measurements are in close agreement with the CMB, and in significant tension with the late Universe measurements. This result gives strong evidence against a simple systematic bias being the root of the tension, and severely limits the available space of models to resolve it. In a similar way, our small-scale measurements of $f\sigma_8$ can be compared to the early and late time measurements of S_8, which are also in tension. Our results are consistent with both sets of measurements when considering only the constraints on Ω_m and σ_8 and excluding the effects of the velocity scaling parameter. However, if the low value we obtain for the velocity scaling parameter is instead caused by a low value of σ_8, we find increased tension with the *Planck* CMB expectations. This result provides motivation for extending this type of analysis to upcoming surveys with greater precision, in order to provide the same kind of independent constraint on the S_8 tension that the BAO+BBN does for the H_0 tension. In extending this analysis it is important to note that our emulators' ability to make predictions for low values of S_8 is limited by the range of values represented in the simulations, and implemented as the training prior. Future emulators must have simulations that extend this range to be able to make predictions for the full range of values represented by late time S_8 measurements.

Our two small-scale analyses begin to give evidence for a third cosmological tension in the value of $f\sigma_8$ measured by RSD within galaxy clustering and the expectation for a ΛCDM universe with *Planck* 2018 cosmological parameters. Our results, as well as all other small-scale RSD measurements, find consistently low values of $f\sigma_8$ when compared to the *Planck*+ΛCDM expectation. This tension is particularly concerning given the differences in modelling, data, and simulations used by each analysis, shown in Table 5.1. These differences in methodology and the consistency in the results leave limited options to

Analysis	Data	Simulations	Model	P18+ΛCDM
Chapter 3 [57]	eBOSS LRG	Aemulus	γ_f	-1.4σ
Chapter 4 [58]	eBOSS LRG	AbacusCosmos	γ_l, γ_n	-2.3σ
Lange 2022 [166]	BOSS LOWZ	Aemulus	CEM	+0.04σ, -1.8σ
Yuan 2022 [275]	BOSS CMASS	AbacusSummit	C-HOD	-1.8σ
Zhai 2023 [277]	LOWZ+CMASS	Aemulus	γ_f	-1.9σ, -0.3σ, -3.4σ

Table 5.1: Data, simulations, and models used by a variety of small-scale RSD analyses, and the difference between their $f\sigma_8(z_{\rm eff})$ constraints and the *Planck* 2018 expectation for a flat ΛCDM universe. The models used in Ch. 3, Ch. 4, and Zhai et al. 2023 are all emulator models based on [278], and listed by the velocity scaling parameters they use. Lange et al. 2022 is the only analysis that does not use an emulator for modelling, instead using the Cosmological Evidence Modelling (CEM) method [167] to infer the likelihood directly from the simulations. Yuan et al. 2022 introduced a new emulator-based model that used the ABACUSHOD [276] code to more efficiently sample the HOD parameter space. Differences from *Planck* 2018+ΛCDM expectations are given in units of the combined uncertainty between the $f\sigma_8(z_{\rm eff})$ constraint of the relevant analysis and the *Planck* prediction, added in quadrature. Negative values indicate the small-scale RSD analysis measured a lower value of $f\sigma_8(z_{\rm eff})$ than the *Planck* expectation. Multiple values indicate the offset in each redshift range for analyses that subdivide their sample.

explain the tension with the ΛCDM expectation without modifying the cosmological model. However, there are several common tools shared by all these analyses. All have models based on CDM-only simulations, with a HOD model to connect galaxies to halos, and all are used to analyse data composed mainly of LRGs observed using the BOSS spectrographs. A non-cosmological solution could take the form of an overlooked systematic related to one of these three shared tools that is able to affect all analyses. However, it should be stated that each analysis has attempted to test these factors, and none have given evidence of an unknown systematic.

The primary systematic concern for the simulations used to construct the models is that the amplitude of the velocity field is overestimated compared to the actual Universe, requiring a lower value of $f\sigma_8$ ($\gamma < 1$ for our scaling parameters) to match the observed data. If this overestimation is scale dependent it would also explain why the small-scale RSD measurements give slightly lower values than large-scale analyses of the same samples, and why some measurements (such as Ch. 3, Ch. 4, and [166]) find lower values of $f\sigma_8$ when using a lower minimum scale cut. The most likely reason for this discrepancy is suppression of growth by feedback mechanisms from baryon physics, which is expected to suppress the

amount of non-linear growth and possibly reduce the amplitude of the velocity field in a scale-dependent way, although the effects on the velocity field over the scales measured here have not been well-studied in recent analyses of hydrodynamical simulations [86, 183, 253]. Suppression of the matter power spectrum on non-linear scales due to baryon feedback has been suggested as a possible resolution to the S_8 tension [15, 220], although this would require suppression that is stronger than expected from hydrodynamical simulations. Nevertheless, it is interesting to investigate if a similar suppression could explain the smaller than expected amplitude of the velocity field, particularly on non-linear scales, found in Ch. 3 and 4. An important avenue for future research is therefore to test small-scale clustering emulators against simulations including baryonic physics.

Alternatively, systematically low measurements of $f\sigma_8$ could be caused by a breakdown in the HOD models used to connect galaxies to halos. If the distribution of galaxies within halos in the model differs significantly from the true distribution of the data, or if the velocities assigned to galaxies in the model differ significantly from what is found in the data, it could cause a systematic shift in the measured values of $f\sigma_8$. HOD models are empirically motivated and based on a relationship between halos and galaxies that is undoubtedly simpler than what is found in the real Universe [123, 122], raising the possibility of such a discrepancy. However, it should be noted that the analyses listed in Table 5.1 perform numerous tests and extensions to their HOD models, including introducing assembly bias parameters, without significantly changing their results. One of the strengths of the HOD method is that the number of free parameters theoretically allows a particular model to match a wide range of possible galaxy-halo connections, reducing the possibility of a deviation significant enough to bias the cosmological inference. While this freedom gives confidence for continuing to use HOD models in small-scale RSD analyses, continuing to test the HOD models and investigating additional parameters should remain a focus for future analyses.

If the low measurements of $f\sigma_8$ are found to be statistically significant and not caused by deficiencies in the modelling or data analysis, either for the small-scale RSD measurements or the *Planck* measurements that they are compared to, it indicates a deviation from the ΛCDM model and new physics in our understanding of the Universe. The most natural explanation would be a deviation from General Relativity (GR), of which there exist many possible Modified Gravity (MG) models (see Sec. 1.1.3). These models are expected to produce modifications to the growth history of the Universe that should be strongly tested by RSD measurements. Some of these models also allow for scale-dependent growth, which could explain the slight shift to lower values of $f\sigma_8$ seen between conventional large-scale RSD measurements and the small-scale measurements listed in Table 5.1, as well as the dependence of $f\sigma_8$ constraints on the minimum measurement scale seen in Ch. 3 and 4

(and others, e.g. [166]).

A MG model or other extension to ΛCDM+GR could also be used to resolve the S_8 tension (see Sec. 1.2.5) through a modification to late-time structure growth [98]. This tension is most clearly seen in the shear-shear measurements of weak lensing, which like the small-scale RSD measurements make use of significant small-scale, non-linear information (see e.g. [181, 150]). Similar to the RSD measurements of $f\sigma_8$, these S_8 measurements from weak lensing are not in significant tension with the *Planck*+ΛCDM expectations, but provide another avenue for exploring and constraining possible extensions. Such extensions are also restricted by complementary observations, such as the expansion history. While the value of H_0 measured by BAO+CMB or BAO+BBN and the value measured be the local distance ladder are in tension, each measurement is individually very consistent with a ΛCDM universe. Because of the large redshift range covered by BAO measurements, thanks in large part to eBOSS, this places strong constraints on deviations to the expansion history in the late Universe.

Current observations are still far from a statistically significant detection of the signature of a MG model, or even favouring one over ΛCDM. An important consideration is that all of the small-scale RSD measurements listed in Table 5.1 implicitly assume a ΛCDM+GR universe because they are constructed from simulations run using that framework. In order to obtain accurate measurements of $f\sigma_8$ or an equivalent set of parameters for a particular MG model, a new set of simulations would need to be run using that gravity framework. Due to the large space of allowed models and lack of a significant detection, it is still a reasonable approach to limit modelling complexity by measuring RSD using models constructed within a ΛCDM+GR framework, and if a significant deviation from the expected value of $f\sigma_8$ is detected then the most promising set of models could be investigated further. Alternatively, work has begun to construct an emulator from MG simulations directly, which could be applied to future data sets [17, 129, 235].

While our velocity-split emulator model provides precise results and shows promise in isolating the linear signal, there remain a number of possible areas of improvement. The smoothing of the linear velocity field, while empirically motivated and tested, could be directly connected to a physical phenomenon [139]. The optimal smoothing scale is likely related to some physical characteristic of the density field, such as the radius for shell crossing. The non-linear velocity parameter is poorly constrained in our analysis, and several significant degeneracies are found between parameters in the fit. Refining these parameters could lead to more informative and precise results. Finally, while the model successfully separates linear growth and random motions, quasi-linear growth along the direction of the linear velocity remains a point of degeneracy between γ_l and γ_n, and a potential bias in the model, that should be addressed using a more sophisticated velocity

split method.

Beyond the velocity scaling parameters, future emulators would do well to carefully consider the choice of HOD parameterization or other galaxy-halo connection model. Our first analysis of small-scale RSD from eBOSS in Ch. 3 highlights how incompleteness due to target selection can be degenerate with measurements of $f\sigma_8$ when not properly included in the model, motivating us to rebuild our emulator to include $f_{\rm max}$. Other small-scale analyses have included parameters to account for the influence of halo parameters other than the mass on galaxy occupation, referred to as halo assembly bias [166, 277]. While not necessary for the current generation of surveys [277], halo assembly bias may be impactful for future, more precise clustering measurements or samples other than LRGs, and should be included in future models. Extending this analysis to a non-LRG sample would also require careful consideration of the expected functional form of the HOD for that sample. For example, ELGs occupy less massive halos than LRGs, requiring an alternative to the simple transition from low to high occupation of centrals commonly used for LRG samples (see Eq. 1.48). Finding the most accurate parameterization remains an active area of research [20, 124, 230].

Another important finding of Ch. 3 is that the effects of redshift uncertainty can lead to larger inferred values of $f\sigma_8$ in small-scale analyses. Because this uncertainty scales with increasing target redshift this is particularly concerning for upcoming surveys. Correctly modelling redshift uncertainty would require a detailed model that takes into account any correlation between the increasing redshift uncertainty and the effects of using an apparent magnitude-based target selection over a wide redshift range. It should be noted that a model like our velocity-split emulator is insensitive to redshift uncertainty by construction, due to the distinction between the coherent linear velocities and the randomly directed non-linear velocities.

As well as refining the modelling choices, an emulator constructed for the next generation of galaxy surveys will require improved simulations. In our eBOSS analyses the intrinsic uncertainty of the emulators constructed using the Aemulus and AbacusCosmos suites is comparable to the data uncertainty on many scales (see Fig. 3.5). To improve the emulator uncertainty to match the expected improvements in future galaxy surveys we require a simulation suite that simulates larger volumes and more densely samples the cosmological parameter space. The suites used in this thesis each consist of 40 simulations, with side length $\sim 1000\,h^{-1}$Mpc and $(\sim 1400)^3$ particles. An ideal candidate is AbacusSummit [180], the successor to AbacusCosmos, which consists of simulations spanning 80 different sets of cosmological parameters, with side length $2000\,h^{-1}$Mpc and $(6912)^3$ particles. AbacusSummit provides both a significant improvement in the number of cosmologies sampled and the simulation volume, as well as improved mass resolution that would allow

extending this method to galaxy samples occupying less massive halos, such as ELGs.

Extending small-scale RSD analyses to the next generation of spectroscopic galaxy surveys will be an important area of future research. This next generation is being led by the Dark Energy Spectroscopic Instrument (DESI) [79, 80] and *Euclid* space mission [168]. DESI is mounted on the 4m Mayall Telescope at Kitt Peak National Observatory in New Mexico, USA, and began its main survey May 14th of 2021. The DESI focal plane contains 5020 robotic fibre positioners connected to a set of 10 identical spectrographs [78], which allows DESI to significantly improve the rate of redshift acquisition compared to BOSS and eBOSS, which each used a set of 1000 hand-plugged fibres. Over the course of its planned five year survey DESI is expected to cover 14 000 $\deg^2$ and observe 40 million galaxies and quasars. The main clustering samples of DESI are the Bright Galaxy Survey (BGS), consisting of >10 million galaxies at $z < 0.6$ and divided into a $r < 19.5$ magnitude limited sample, a fainter colour-selected sample, and low redshift quasars [125]; the DESI LRG sample, consisting of ∼8 million LRGs in the redshift range $0.4 < z < 1.0$ [283]; the DESI ELG sample, consisting of ∼13 million ELGs in the redshift range $0.6 < z < 1.6$ [222]; and the DESI quasar samples, consisting of ∼3 million quasars divided between a clustering sample in the range $0.9 < z < 2.1$ and a Lyα forest sample at $z > 2.1$ [59]. These samples are projected to provide precise expansion and growth history constraints across the redshift range by dividing these samples into 13 redshift bins of width $\delta z = 0.1$, with the first seven bins each providing $< 1\%$, $< 2\%$, and $\sim 1\%$ precision constraints on the transverse BAO scale, radial BAO scale, and $f\sigma_8$ respectively [79]. While the latter six bins will have increasing uncertainty, they will still provide the most precise measurements in that redshift range to date. Constraining MG models with DESI has also been investigated using simulations, with the results indicating that DESI will provide considerably improved constraints on many classes of models, particularly from small-scale RSD measurements [8].

Euclid is a 1.2 m space-based telescope located at the Earth-Sun Lagrange point 2 (on route at the time of writing), launched July 1st of 2023. Post-launch, *Euclid* will undergo 3 months commissioning and performance verification before beginning a six year survey. The *Euclid* wide survey will cover 15 000 $\deg^2$, consisting of both spectroscopic observations for 3D galaxy clustering measurements and photometric observations for weak lensing measurements [96]. The spectroscopic component of the wide survey will use two "red" grisms (1250-1850 nm) to perform slitless spectroscopy in the Y, I, and H bands, complete to an apparent magnitude limit of 24.0 in each band. The survey is expected to measure the redshifts of up to 30 million galaxies [218] at a maximum redshift of $z \approx 2$, and is optimized for observing Hα emitters in the redshift range $0.9 < z < 1.8$ [96]. Combining the spectroscopic galaxy clustering measurements with the weak lensing and

photometric galaxy clustering from the *Euclid* wide survey is expected to yield sub-percent level constraints on $\Omega_{m,0}$, σ_8, H_0, and n_s for a flat ΛCDM universe [95]. Data releases are planned approximately two years (covering 2 500 deg^2), four years (covering 7 500 deg^2), and seven years (full 15 000 deg^2) after the start of the survey. Between the forthcoming data and measurements from both DESI and *Euclid*, this is a very exciting and crucial time for spectroscopic galaxy surveys.

Our work with eBOSS provides a road-map for these surveys. We explored extending observations to the highest redshift range for spectroscopic galaxy clustering to date, and pioneered the use of ELGs and quasars for clustering measurements, paving the way for their inclusion in DESI. In particular, the work performed in eBOSS to correct observational systematics should serve as an important lesson when working with upcoming data. PIP+ANG weights remain the only fully unbiased method for removing the effects of fibre collisions, and our work in eBOSS highlights the importance of using a random target selection algorithm that can be easily rerun and instrument pointings that are uncorrelated with the underlying mass density in order to avoid an unrecoverable loss of information. This information is directly applicable, as PIP+ANG weights have been investigated for correcting fibre collisions in DESI [34, 249]. The ELG targeting systematic also provides a cautionary tale for using a consistent calibration scheme across the survey. Our results have already had an impact, as the DESI collaboration updated the photometric catalogue they used for target selection to a version of DECaLS that removes the poor exposures that caused issues for the eBOSS ELG sample [222].

Extending small-scale RSD analyses to the next generation of surveys will provide a key test of non-cosmological sources of discrepancy in the $f\sigma_8$ measurements. Performing a similar analysis as presented in Ch. 4 with DESI would eliminate many possible sources of observational bias in the data because DESI uses a different target selection and significantly improved instrument. Measuring the DESI ELG sample would be particularly interesting since ELGs are expected to have a different HOD from LRGs, allowing an independent test of the HOD model. Investigating these factors with next generation surveys could also provide insight into the low value of γ_f we obtain using the full emulator range in Ch. 3 and the low value of γ_n we obtain in Ch. 4, which indicate a discrepancy between the model and the data in both the linear and non-linear velocity fields.

Performing small-scale RSD analyses with the next generation of spectroscopic galaxy surveys will be essential for extracting optimal cosmological constraints. Small-scale analyses have consistently found a factor of ~ 2 improvement in precision over more conventional large-scale measurements, so will be key to identifying or rejecting a tension in the growth rate of cosmic structure. If current offsets from *Planck*+ΛCDM expectations are maintained, then the next generation of surveys will provide the first significant detection of a

tension in $f\sigma_8$ measurements, opening a new avenue for expanding our knowledge of the Universe.

References

[1] Abazajian, Kevork N. , Adelman-McCarthy, Jennifer K. , Agüeros, Marcel A. , Allam, Sahar S. , Allende Prieto, Carlos , An, Deokkeun , Anderson, Kurt S. J. , Anderson, Scott F. , et al. The Seventh Data Release of the Sloan Digital Sky Survey. *ApJS*, 182(2):543–558, June 2009. doi:10.1088/0067-0049/182/2/543.

[2] Abbott, T. M. C. , Abdalla, F. B. , Alarcon, A. , Aleksić, J. , Allam, S. , Allen, S. , Amara, A. , Annis, J. , et al. Dark Energy Survey year 1 results: Cosmological constraints from galaxy clustering and weak lensing. *Phys. Ref. D*, 98(4):043526, August 2018. doi:10.1103/PhysRevD.98.043526.

[3] Abbott, T. M. C. , Aguena, M. , Alarcon, A. , Allam, S. , Alves, O. , Amon, A. , Andrade-Oliveira, F. , Annis, J. , et al. Dark Energy Survey Year 3 results: Cosmological constraints from galaxy clustering and weak lensing. *Phys. Ref. D*, 105(2):023520, January 2022. doi:10.1103/PhysRevD.105.023520.

[4] Abdalla, Elcio , Abellán, Guillermo Franco , Aboubrahim, Amin , Agnello, Adriano , Akarsu, Özgür , Akrami, Yashar , Alestas, George , Aloni, Daniel , et al. Cosmology intertwined: A review of the particle physics, astrophysics, and cosmology associated with the cosmological tensions and anomalies. *Journal of High Energy Astrophysics*, 34:49–211, June 2022. doi:10.1016/j.jheap.2022.04.002.

[5] Aihara, Hiroaki , Arimoto, Nobuo , Armstrong, Robert , Arnouts, Stéphane , Bahcall, Neta A. , Bickerton, Steven , Bosch, James , Bundy, Kevin , et al. The Hyper Suprime-Cam SSP Survey: Overview and survey design. *Publications of the Astronomicall Society of Japan*, 70:S4, January 2018. doi:10.1093/pasj/psx066.

[6] Aiola, Simone , Calabrese, Erminia , Maurin, Loïc , Naess, Sigurd , Schmitt, Benjamin L. , Abitbol, Maximilian H. , Addison, Graeme E. , Ade, Peter A. R.

, et al. The Atacama Cosmology Telescope: DR4 maps and cosmological parameters. *J. Cosmology Astropart. Phys.*, 2020(12):047, December 2020. doi: 10.1088/1475-7516/2020/12/047.

[7] Alam, Shadab , Albareti, Franco D. , Allende Prieto, Carlos , Anders, F. , Anderson, Scott F. , Anderton, Timothy , Andrews, Brett H. , Armengaud, Eric , et al. The Eleventh and Twelfth Data Releases of the Sloan Digital Sky Survey: Final Data from SDSS-III. *ApJS*, 219(1):12, July 2015. doi:10.1088/0067-0049/219/1/12.

[8] Alam, Shadab , Arnold, Christian , Aviles, Alejandro , Bean, Rachel , Cai, Yan-Chuan , Cautun, Marius , Cervantes-Cota, Jorge L. , Cuesta-Lazaro, Carolina , et al. Towards testing the theory of gravity with DESI: summary statistics, model predictions and future simulation requirements. *J. Cosmology Astropart. Phys.*, 2021(11):050, November 2021. doi:10.1088/1475-7516/2021/11/050.

[9] Alam, Shadab , Ata, Metin , Bailey, Stephen , Beutler, Florian , Bizyaev, Dmitry , Blazek, Jonathan A. , Bolton, Adam S. , Brownstein, Joel R. , et al. The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample. *MNRAS*, 470(3):2617–2652, September 2017. doi:10.1093/mnras/stx721.

[10] Alam, Shadab , Aubert, Marie , Avila, Santiago , Balland, Christophe , Bautista, Julian E. , Bershady, Matthew A. , Bizyaev, Dmitry , Blanton, Michael R. , et al. Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological implications from two decades of spectroscopic surveys at the Apache Point Observatory. *Phys. Ref. D*, 103(8):083533, April 2021. doi:10.1103/PhysRevD.103.083533.

[11] Alam, Shadab , de Mattia, Arnaud , Tamone, Amélie , Ávila, S. , Peacock, John A. , Gonzalez-Perez, V. , Smith, Alex , Raichoor, Anand , et al. The completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: N-body mock challenge for the eBOSS emission line galaxy sample. *MNRAS*, 504(4):4667–4686, July 2021. doi: 10.1093/mnras/stab1150.

[12] Albareti, Franco D. , Allende Prieto, Carlos , Almeida, Andres , Anders, Friedrich , Anderson, Scott , Andrews, Brett H. , Aragón-Salamanca, Alfonso , Argudo-Fernández, Maria , et al. The 13th Data Release of the Sloan Digital Sky Survey: First Spectroscopic Data from the SDSS-IV Survey Mapping Nearby Galaxies at Apache Point Observatory. *ApJS*, 233(2):25, December 2017. doi:10.3847/1538-4365/aa8992.

[13] Alcock, C. and Paczynski, B. . An evolution free test for non-zero cosmological constant. *Nature*, 281:358, October 1979. doi:10.1038/281358a0.

[14] Amendola, Luca and Tsujikawa, Shinji . *Dark Energy: Theory and Observations*. Cambridge University Press, 2010.

[15] Amon, Alexandra and Efstathiou, George . A non-linear solution to the S_8 tension? *MNRAS*, 516(4):5355–5366, November 2022. doi:10.1093/mnras/stac2429.

[16] Arnett, W. David . A Possible Model of Supernovae: Detonation of ^{12}C. *Astrophys. Space Sci.*, 5(2):180–212, October 1969. doi:10.1007/BF00650291.

[17] Arnold, Christian , Li, Baojiu , Giblin, Benjamin , Harnois-Déraps, Joachim , and Cai, Yan-Chuan . FORGE: the f(R)-gravity cosmic emulator project - I. Introduction and matter power spectrum emulator. *MNRAS*, 515(3):4161–4175, September 2022. doi:10.1093/mnras/stac1091.

[18] Asgari, Marika , Lin, Chieh-An , Joachimi, Benjamin , Giblin, Benjamin , Heymans, Catherine , Hildebrandt, Hendrik , Kannawadi, Arun , Stölzner, Benjamin , et al. KiDS-1000 cosmology: Cosmic shear constraints and comparison between two point statistics. *A&A*, 645:A104, January 2021. doi:10.1051/0004-6361/202039070.

[19] Assassi, Valentin , Baumann, Daniel , Green, Daniel , and Zaldarriaga, Matias . Renormalized halo bias. *J. Cosmology Astropart. Phys.*, 2014(8):056–056, August 2014. doi:10.1088/1475-7516/2014/08/056.

[20] Avila, S. , Gonzalez-Perez, V. , Mohammad, F. G. , de Mattia, A. , Zhao, C. , Raichoor, A. , Tamone, A. , Alam, S. , et al. The Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: exploring the halo occupation distribution model for emission line galaxies. *MNRAS*, 499(4):5486–5507, December 2020. doi: 10.1093/mnras/staa2951.

[21] Bacon, D. J. , Goldberg, D. M. , Rowe, B. T. P. , and Taylor, A. N. . Weak gravitational flexion. *MNRAS*, 365(2):414–428, January 2006. doi:10.1111/j.1365-2966.2005.09624.x.

[22] Baumann, Daniel , Nicolis, Alberto , Senatore, Leonardo , and Zaldarriaga, Matias . Cosmological non-linearities as an effective fluid. *J. Cosmology Astropart. Phys.*, 2012(7):051, July 2012. doi:10.1088/1475-7516/2012/07/051.

[23] Bautista, Julian E. , Paviot, Romain , Vargas Magaña, Mariana , de la Torre, Sylvain , Fromenteau, Sebastien , Gil-Marín, Hector , Ross, Ashley J. , Burtin, Etienne , et al. The completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: measurement of the BAO and growth rate of structure of the luminous red

galaxy sample from the anisotropic correlation function between redshifts 0.6 and 1. *MNRAS*, 500(1):736–762, January 2021. doi:10.1093/mnras/staa2800.

[24] Bautista, Julian E. , Vargas-Magaña, Mariana , Dawson, Kyle S. , Percival, Will J. , Brinkmann, Jonathan , Brownstein, Joel , Camacho, Benjamin , Comparat, Johan , et al. The SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations at Redshift of 0.72 with the DR14 Luminous Red Galaxy Sample. *ApJ*, 863(1):110, August 2018. doi:10.3847/1538-4357/aacea5.

[25] Behroozi, Peter S. , Conroy, Charlie , and Wechsler, Risa H. . A Comprehensive Analysis of Uncertainties Affecting the Stellar Mass-Halo Mass Relation for $0 < z < 4$. *ApJ*, 717(1):379–403, July 2010. doi:10.1088/0004-637X/717/1/379.

[26] Behroozi, Peter S. , Wechsler, Risa H. , and Wu, Hao-Yi . The ROCKSTAR Phase-space Temporal Halo Finder and the Velocity Offsets of Cluster Cores. *ApJ*, 762(2):109, January 2013. doi:10.1088/0004-637X/762/2/109.

[27] Bel, J. , Pezzotta, A. , Carbone, C. , Sefusatti, E. , and Guzzo, L. . Accurate fitting functions for peculiar velocity spectra in standard and massive-neutrino cosmologies. *A&A*, 622:A109, February 2019. doi:10.1051/0004-6361/201834513.

[28] Bennett, C. L. , Larson, D. , Weiland, J. L. , Jarosik, N. , Hinshaw, G. , Odegard, N. , Smith, K. M. , Hill, R. S. , et al. Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results. *ApJS*, 208(2):20, October 2013. doi:10.1088/0067-0049/208/2/20.

[29] Benson, Andrew J. . Galaxy formation theory. *Phys. Rept.*, 495(2-3):33–86, October 2010. doi:10.1016/j.physrep.2010.06.001.

[30] Berlind, Andreas A. and Weinberg, David H. . The Halo Occupation Distribution: Toward an Empirical Determination of the Relation between Galaxies and Mass. *ApJ*, 575(2):587–616, August 2002. doi:10.1086/341469.

[31] Bertin, E. and Arnouts, S. . SExtractor: Software for source extraction. *A&AS*, 117:393–404, June 1996. doi:10.1051/aas:1996164.

[32] Beutler, Florian , Blake, Chris , Colless, Matthew , Jones, D. Heath , Staveley-Smith, Lister , Poole, Gregory B. , Campbell, Lachlan , Parker, Quentin , et al. The 6dF Galaxy Survey: $z \approx 0$ measurements of the growth rate and σ_8. *MNRAS*, 423(4):3430–3444, July 2012. doi:10.1111/j.1365-2966.2012.21136.x.

[33] Beutler, Florian , Seo, Hee-Jong , Saito, Shun , Chuang, Chia-Hsun , Cuesta, Antonio J. , Eisenstein, Daniel J. , Gil-Marín, Héctor , Grieb, Jan Niklas , et al. The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: anisotropic galaxy clustering in Fourier space. *Monthly Notices of the Royal Astronomical Society*, 466:2242–2260, April 2017. ISSN 0035-8711. doi: 10.1093/mnras/stw3298.

[34] Bianchi, Davide , Burden, Angela , Percival, Will J. , Brooks, David , Cahn, Robert N. , Forero-Romero, Jaime E. , Levi, Michael , Ross, Ashley J. , et al. Unbiased clustering estimates with the DESI fibre assignment. *MNRAS*, 481(2):2338–2348, December 2018. doi:10.1093/mnras/sty2377.

[35] Bianchi, Davide , Gil-Marín, Héctor , Ruggeri, Rossana , and Percival, Will J. . Measuring line-of-sight-dependent Fourier-space clustering using FFTs. *MNRAS*, 453(1):L11–L15, October 2015. doi:10.1093/mnrasl/slv090.

[36] Bianchi, Davide and Percival, Will J. . Unbiased clustering estimation in the presence of missing observations. *MNRAS*, 472(1):1106–1118, Nov 2017. doi:10.1093/mnras/stx2053.

[37] Blake, Chris , Baldry, I. K. , Bland-Hawthorn, J. , Christodoulou, L. , Colless, M. , Conselice, C. , Driver, S. P. , Hopkins, A. M. , et al. Galaxy And Mass Assembly (GAMA): improved cosmic growth measurements using multiple tracers of large-scale structure. *MNRAS*, 436(4):3089–3105, December 2013. doi:10.1093/mnras/stt1791.

[38] Blake, Chris , Brough, Sarah , Colless, Matthew , Contreras, Carlos , Couch, Warrick , Croom, Scott , Davis, Tamara , Drinkwater, Michael J. , et al. The WiggleZ Dark Energy Survey: the growth rate of cosmic structure since redshift z=0.9. *MNRAS*, 415(3):2876–2891, August 2011. doi:10.1111/j.1365-2966.2011.18903.x.

[39] Blanchard, A. and Schneider, J. . Gravitational lensing effect on the fluctuations of the cosmic background radiation. *A&A*, 184(1-2):1–6, October 1987.

[40] Blandford, R. D. , Saust, A. B. , Brainerd, T. G. , and Villumsen, J. V. . The distortion of distant galaxy images by large-scale structure. *MNRAS*, 251:600, August 1991. doi:10.1093/mnras/251.4.600.

[41] Blanton, Michael R. , Bershady, Matthew A. , Abolfathi, Bela , Albareti, Franco D. , Allende Prieto, Carlos , Almeida, Andres , Alonso-García, Javier , Anders, Friedrich , et al. Sloan Digital Sky Survey IV: Mapping the Milky Way, Nearby Galaxies,

and the Distant Universe. *The Astronomical Journal*, 154(1):28, July 2017. ISSN 0004-6256. doi:10.3847/1538-3881/aa7567.

[42] Bolton, Adam S. , Schlegel, David J. , Aubourg, Éric , Bailey, Stephen , Bhardwaj, Vaishali , Brownstein, Joel R. , Burles, Scott , Chen, Yan-Mei , et al. Spectral Classification and Redshift Measurement for the SDSS-III Baryon Oscillation Spectroscopic Survey. *AJ*, 144(5):144, November 2012. doi:10.1088/0004-6256/144/5/144.

[43] Bono, Giuseppe , Marconi, Marcella , and Stellingwerf, Robert F. . Classical Cepheid Pulsation Models. I. Physical Structure. *ApJS*, 122(1):167–205, May 1999. doi:10.1086/313207.

[44] Bovy, Jo , Myers, Adam D. , Hennawi, Joseph F. , Hogg, David W. , McMahon, Richard G. , Schiminovich, David , Sheldon, Erin S. , Brinkmann, Jon , et al. Photometric Redshifts and Quasar Probabilities from a Single, Data-driven Generative Model. *ApJ*, 749(1):41, April 2012. doi:10.1088/0004-637X/749/1/41.

[45] Bueno Belloso, Alicia , Pettinari, Guido W. , Meures, Nikolai , and Percival, Will J. . Using galaxy pairs as cosmological tracers. *Phys. Ref. D*, 86(2):023530, July 2012. doi:10.1103/PhysRevD.86.023530.

[46] Bullock, J. S. , Kolatt, T. S. , Sigad, Y. , Somerville, R. S. , Kravtsov, A. V. , Klypin, A. A. , Primack, J. R. , and Dekel, A. . Profiles of dark haloes: evolution, scatter and environment. *MNRAS*, 321(3):559–575, March 2001. doi:10.1046/j.1365-8711.2001.04068.x.

[47] Burden, A. , Percival, W. J. , and Howlett, C. . Reconstruction in Fourier space. *MNRAS*, 453(1):456–468, October 2015. doi:10.1093/mnras/stv1581.

[48] Burden, A. , Percival, W. J. , Manera, M. , Cuesta, Antonio J. , Vargas Magana, Mariana , and Ho, Shirley . Efficient reconstruction of linear baryon acoustic oscillations in galaxy surveys. *MNRAS*, 445(3):3152–3168, December 2014. doi:10.1093/mnras/stu1965.

[49] Burden, Angela , Padmanabhan, Nikhil , Cahn, Robert N. , White, Martin J. , and Samushia, Lado . Mitigating the impact of the DESI fiber assignment on galaxy clustering. *J. Cosmology Astropart. Phys.*, 2017(3):001, March 2017. doi:10.1088/1475-7516/2017/03/001.

[50] Busca, N. G. , Delubac, T. , Rich, J. , Bailey, S. , Font-Ribera, A. , Kirkby, D. , Le Goff, J. M. , Pieri, M. M. , et al. Baryon acoustic oscillations in the Lyα forest of BOSS quasars. *A&A*, 552:A96, April 2013. doi:10.1051/0004-6361/201220724.

[51] Cabass, Giovanni , Ivanov, Mikhail M. , Lewandowski, Matthew , Mirbabayi, Mehrdad , and Simonović, Marko . Snowmass White Paper: Effective Field Theories in Cosmology. *arXiv e-prints*, arXiv:2203.08232, March 2022. doi:10.48550/arXiv.2203.08232.

[52] Carlson, Jordan , Reid, Beth , and White, Martin . Convolution Lagrangian perturbation theory for biased tracers. *MNRAS*, 429(2):1674–1685, February 2013. doi:10.1093/mnras/sts457.

[53] Carlstrom, J. E. , Ade, P. A. R. , Aird, K. A. , Benson, B. A. , Bleem, L. E. , Busetti, S. , Chang, C. L. , Chauvin, E. , et al. The 10 Meter South Pole Telescope. *PASP*, 123(903):568, May 2011. doi:10.1086/659879.

[54] Carrasco, John Joseph M. , Hertzberg, Mark P. , and Senatore, Leonardo . The effective field theory of cosmological large scale structures. *Journal of High Energy Physics*, 2012:82, September 2012. doi:10.1007/JHEP09(2012)082.

[55] Carroll, Sean M. , Press, William H. , and Turner, Edwin L. . The cosmological constant. *ARA&A*, 30:499–542, January 1992. doi:10.1146/annurev.aa.30.090192.002435.

[56] Cataneo, Matteo , Foreman, Simon , and Senatore, Leonardo . Efficient exploration of cosmology dependence in the EFT of LSS. *J. Cosmology Astropart. Phys.*, 2017(4):026, April 2017. doi:10.1088/1475-7516/2017/04/026.

[57] Chapman, Michael J. , Mohammad, Faizan G. , Zhai, Zhongxu , Percival, Will J. , Tinker, Jeremy L. , Bautista, Julian E. , Brownstein, Joel R. , Burtin, Etienne , et al. 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. *MNRAS*, 516(1):617–635, October 2022. doi:10.1093/mnras/stac1923.

[58] Chapman, Michael J. , Zhai, Zhongxu , and Percival, Will J. . Isolating the linear signal when making redshift space distortion measurements. *Accepted in MNRAS*, arXiv:2302.11621, July 2023. doi:10.48550/arXiv.2302.11621.

[59] Chaussidon, Edmond , Yèche, Christophe , Palanque-Delabrouille, Nathalie , Alexander, David M. , Yang, Jinyi , Ahlen, Steven , Bailey, Stephen , Brooks, David , et al. Target Selection and Validation of DESI Quasars. *ApJ*, 944(1):107, February 2023. doi:10.3847/1538-4357/acb3c2.

[60] Chen, Shi-Fan , Vlah, Zvonimir , Castorina, Emanuele , and White, Martin . Redshift-space distortions in Lagrangian perturbation theory. *J. Cosmology Astropart. Phys.*, 2021(3):100, March 2021. doi:10.1088/1475-7516/2021/03/100.

[61] Cole, Shaun , Percival, Will J. , Peacock, John A. , Norberg, Peder , Baugh, Carlton M. , Frenk, Carlos S. , Baldry, Ivan , Bland-Hawthorn, Joss , et al. The 2dF Galaxy Redshift Survey: power-spectrum analysis of the final data set and cosmological implications. *MNRAS*, 362(2):505–534, September 2005. doi:10.1111/j.1365-2966.2005.09318.x.

[62] Colgate, S. A. . Supernovae as a standard candle for cosmology. *ApJ*, 232:404–408, September 1979. doi:10.1086/157300.

[63] Colgate, Stirling A. and McKee, Chester . Early Supernova Luminosity. *ApJ*, 157:623, August 1969. doi:10.1086/150102.

[64] Comparat, J. , Merloni, A. , Dwelly, T. , Salvato, M. , Schwope, A. , Coffey, D. , Wolf, J. , Arcodia, R. , et al. The final SDSS-IV/SPIDERS X-ray point source spectroscopic catalogue. *A&A*, 636:A97, April 2020. doi:10.1051/0004-6361/201937272.

[65] Comparat, Johan , Richard, Johan , Kneib, Jean-Paul , Ilbert, Olivier , Gonzalez-Perez, Violeta , Tresse, Laurence , Zoubian, Julien , Arnouts, Stephane , et al. The $0.1 < z < 1.65$ evolution of the bright end of the [O ii] luminosity function. *A&A*, 575:A40, March 2015. doi:10.1051/0004-6361/201424767.

[66] Cooray, Asantha and Sheth, Ravi . Halo models of large scale structure. *Phys. Rept.*, 372(1):1–129, December 2002. doi:10.1016/S0370-1573(02)00276-4.

[67] Croom, S. M. , Smith, R. J. , Boyle, B. J. , Shanks, T. , Miller, L. , Outram, P. J. , and Loaring, N. S. . The 2dF QSO Redshift Survey - XII. The spectroscopic catalogue and luminosity function. *MNRAS*, 349(4):1397–1418, April 2004. doi:10.1111/j.1365-2966.2004.07619.x.

[68] Cuesta-Lazaro, Carolina , Nishimichi, Takahiro , Kobayashi, Yosuke , Ruan, Cheng-Zong , Eggemeier, Alexander , Miyatake, Hironao , Takada, Masahiro , Yoshida,

Naoki , et al. Galaxy clustering from the bottom up: a streaming model emulator I. *MNRAS*, 523(3):3219–3238, August 2023. doi:10.1093/mnras/stad1207.

[69] d'Amico, Guido , Gleyzes, Jérôme , Kokron, Nickolas , Markovic, Katarina , Senatore, Leonardo , Zhang, Pierre , Beutler, Florian , and Gil-Marín, Héctor . The cosmological analysis of the SDSS/BOSS data from the Effective Field Theory of Large-Scale Structure. *J. Cosmology Astropart. Phys.*, 2020(5):005, May 2020. doi: 10.1088/1475-7516/2020/05/005.

[70] Dark Energy Survey Collaboration, Abbott, T. , Abdalla, F. B. , Aleksić, J. , Allam, S. , Amara, A. , Bacon, D. , Balbinot, E. , et al. The Dark Energy Survey: more than dark energy - an overview. *MNRAS*, 460(2):1270–1299, August 2016. doi: 10.1093/mnras/stw641.

[71] Davis, M. , Efstathiou, G. , Frenk, C. S. , and White, S. D. M. . The evolution of large-scale structure in a universe dominated by cold dark matter. *ApJ*, 292:371–394, May 1985. doi:10.1086/163168.

[72] Dawson, Kyle S. , Kneib, Jean-Paul , Percival, Will J. , Alam, Shadab , Albareti, Franco D. , Anderson, Scott F. , Armengaud, Eric , Aubourg, Éric , et al. The SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Overview and Early Data. *The Astronomical Journal*, 151:44, February 2016. ISSN 0004-6256. doi: 10.3847/0004-6256/151/2/44.

[73] Dawson, Kyle S. , Schlegel, David J. , Ahn, Christopher P. , Anderson, Scott F. , Aubourg, Éric , Bailey, Stephen , Barkhouser, Robert H. , Bautista, Julian E. , et al. The Baryon Oscillation Spectroscopic Survey of SDSS-III. *The Astronomical Journal*, 145:10, January 2013. ISSN 0004-6256. doi:10.1088/0004-6256/145/1/10.

[74] de la Torre, Sylvain and Guzzo, Luigi . Modelling non-linear redshift-space distortions in the galaxy clustering pattern: systematic errors on the growth rate parameter. *MNRAS*, 427(1):327–342, November 2012. doi:10.1111/j.1365-2966.2012.21824.x.

[75] de Mattia, Arnaud and Ruhlmann-Kleider, Vanina . Integral constraints in spectroscopic surveys. *J. Cosmology Astropart. Phys.*, 2019(8):036, August 2019. doi: 10.1088/1475-7516/2019/08/036.

[76] de Mattia, Arnaud , Ruhlmann-Kleider, Vanina , Raichoor, Anand , Ross, Ashley J. , Tamone, Amélie , Zhao, Cheng , Alam, Shadab , Avila, Santiago , et al. The completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: measurement of

the BAO and growth rate of structure of the emission line galaxy sample from the anisotropic power spectrum between redshift 0.6 and 1.1. *MNRAS*, 501(4):5616–5645, March 2021. doi:10.1093/mnras/staa3891.

[77] DeRose, Joseph , Wechsler, Risa H. , Tinker, Jeremy L. , Becker, Matthew R. , Mao, Yao-Yuan , McClintock, Thomas , McLaughlin, Sean , Rozo, Eduardo , et al. The AEMULUS Project. I. Numerical Simulations for Precision Cosmology. *ApJ*, 875(1):69, April 2019. doi:10.3847/1538-4357/ab1085.

[78] DESI Collaboration, Abareshi, B. , Aguilar, J. , Ahlen, S. , Alam, Shadab , Alexander, David M. , Alfarsy, R. , Allen, L. , et al. Overview of the Instrumentation for the Dark Energy Spectroscopic Instrument. *AJ*, 164(5):207, November 2022. doi:10.3847/1538-3881/ac882b.

[79] DESI Collaboration, Aghamousa, Amir , Aguilar, Jessica , Ahlen, Steve , Alam, Shadab , Allen, Lori E. , Allende Prieto, Carlos , Annis, James , et al. The DESI Experiment Part I: Science,Targeting, and Survey Design. *arXiv e-prints*, arXiv:1611.00036, October 2016.

[80] DESI Collaboration, Aghamousa, Amir , Aguilar, Jessica , Ahlen, Steve , Alam, Shadab , Allen, Lori E. , Allende Prieto, Carlos , Annis, James , et al. The DESI Experiment Part II: Instrument Design. *arXiv e-prints*, arXiv:1611.00037, October 2016.

[81] Dey, Arjun , Schlegel, David J. , Lang, Dustin , Blum, Robert , Burleigh, Kaylan , Fan, Xiaohui , Findlay, Joseph R. , Finkbeiner, Doug , et al. Overview of the DESI Legacy Imaging Surveys. *AJ*, 157(5):168, May 2019. doi:10.3847/1538-3881/ab089d.

[82] Di Valentino, Eleonora , Melchiorri, Alessandro , Mena, Olga , and Vagnozzi, Sunny . Nonminimal dark sector physics and cosmological tensions. *Phys. Ref. D*, 101(6):063502, March 2020. doi:10.1103/PhysRevD.101.063502.

[83] Drinkwater, Michael J. , Jurek, Russell J. , Blake, Chris , Woods, David , Pimbblet, Kevin A. , Glazebrook, Karl , Sharp, Rob , Pracy, Michael B. , et al. The WiggleZ Dark Energy Survey: survey design and first data release. *MNRAS*, 401(3):1429–1452, January 2010. doi:10.1111/j.1365-2966.2009.15754.x.

[84] Driver, S. P. , Hill, D. T. , Kelvin, L. S. , Robotham, A. S. G. , Liske, J. , Norberg, P. , Baldry, I. K. , Bamford, S. P. , et al. Galaxy and Mass Assembly (GAMA): survey diagnostics and core data release. *MNRAS*, 413(2):971–995, May 2011. doi:10.1111/j.1365-2966.2010.18188.x.

[85] du Mas des Bourboux, Hélion , Rich, James , Font-Ribera, Andreu , de Sainte Agathe, Victoria , Farr, James , Etourneau, Thomas , Le Goff, Jean-Marc , Cuceu, Andrei , et al. The Completed SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations with Lyα Forests. *ApJ*, 901(2):153, October 2020. doi:10.3847/1538-4357/abb085.

[86] Dubois, Y. , Pichon, C. , Welker, C. , Le Borgne, D. , Devriendt, J. , Laigle, C. , Codis, S. , Pogosyan, D. , et al. Dancing in the dark: galactic properties trace spin swings along the cosmic web. *MNRAS*, 444(2):1453–1468, October 2014. doi: 10.1093/mnras/stu1227.

[87] Dyer, C. C. and Roeder, R. C. . Observations in Locally Inhomogeneous Cosmological Models. *ApJ*, 189:167–176, April 1974. doi:10.1086/152784.

[88] Eifler, T. , Kilbinger, M. , and Schneider, P. . Comparing cosmic shear measures. Optimizing the information content of cosmic shear data vectors. *A&A*, 482(1):9–19, April 2008. doi:10.1051/0004-6361:20078573.

[89] Einasto, Jaan , Kaasik, Ants , and Saar, Enn . Dynamic evidence on massive coronas of galaxies. *Nature*, 250(5464):309–310, July 1974. doi:10.1038/250309a0.

[90] Eisenstein, Daniel J. , Annis, James , Gunn, James E. , Szalay, Alexander S. , Connolly, Andrew J. , Nichol, R. C. , Bahcall, Neta A. , Bernardi, Mariangela , et al. Spectroscopic Target Selection for the Sloan Digital Sky Survey: The Luminous Red Galaxy Sample. *AJ*, 122(5):2267–2280, November 2001. doi:10.1086/323717.

[91] Eisenstein, Daniel J. and Hu, Wayne . Baryonic Features in the Matter Transfer Function. *ApJ*, 496(2):605–614, March 1998. doi:10.1086/305424.

[92] Eisenstein, Daniel J. , Seo, Hee-Jong , Sirko, Edwin , and Spergel, David N. . Improving Cosmological Distance Measurements by Reconstruction of the Baryon Acoustic Peak. *ApJ*, 664(2):675–679, August 2007. doi:10.1086/518712.

[93] Eisenstein, Daniel J. , Weinberg, David H. , Agol, Eric , Aihara, Hiroaki , Allende Prieto, Carlos , Anderson, Scott F. , Arns, James A. , Aubourg, Éric , et al. SDSS-III: Massive Spectroscopic Surveys of the Distant Universe, the Milky Way, and Extra-Solar Planetary Systems. *The Astronomical Journal*, 142:72, September 2011. ISSN 0004-6256. doi:10.1088/0004-6256/142/3/72;.

[94] Eisenstein, Daniel J. , Zehavi, Idit , Hogg, David W. , Scoccimarro, Roman , Blanton, Michael R. , Nichol, Robert C. , Scranton, Ryan , Seo, Hee-Jong , et al. Detection of

the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies. *ApJ*, 633(2):560–574, November 2005. doi:10.1086/466512.

[95] Euclid Collaboration, Blanchard, A. , Camera, S. , Carbone, C. , Cardone, V. F. , Casas, S. , Clesse, S. , Ilić, S. , et al. Euclid preparation. VII. Forecast validation for Euclid cosmological probes. *A&A*, 642:A191, October 2020. doi:10.1051/0004-6361/202038071.

[96] Euclid Collaboration, Scaramella, R. , Amiaux, J. , Mellier, Y. , Burigana, C. , Carvalho, C. S. , Cuillandre, J. C. , Da Silva, A. , et al. Euclid preparation. I. The Euclid Wide Survey. *A&A*, 662:A112, June 2022. doi:10.1051/0004-6361/202141938.

[97] Feldman, H. A. , Kaiser, N. , and Peacock, J. A. . Power-spectrum analysis of three-dimensional redshift surveys. *ApJ*, 426:23–37, May 1994. doi:10.1086/174036.

[98] Ferreira, Pedro G. . Cosmological Tests of Gravity. *ARA&A*, 57:335–374, August 2019. doi:10.1146/annurev-astro-091918-104423.

[99] Fisher, Karl B. . On the Validity of the Streaming Model for the Redshift-Space Correlation Function in the Linear Regime. *ApJ*, 448:494, August 1995. doi:10.1086/175980.

[100] Fixsen, D. J. . The Temperature of the Cosmic Microwave Background. *ApJ*, 707(2):916–920, December 2009. doi:10.1088/0004-637X/707/2/916.

[101] Flaugher, B. , Diehl, H. T. , Honscheid, K. , Abbott, T. M. C. , Alvarez, O. , Angstadt, R. , Annis, J. T. , Antonik, M. , et al. The Dark Energy Camera. *AJ*, 150(5):150, November 2015. doi:10.1088/0004-6256/150/5/150.

[102] Font-Ribera, Andreu , Kirkby, David , Busca, Nicolas , Miralda-Escudé, Jordi , Ross, Nicholas P. , Slosar, Anže , Rich, James , Aubourg, Éric , et al. Quasar-Lyman α forest cross-correlation from BOSS DR11: Baryon Acoustic Oscillations. *J. Cosmology Astropart. Phys.*, 2014(5):027, May 2014. doi:10.1088/1475-7516/2014/05/027.

[103] Font-Ribera, Andreu , McDonald, Patrick , Mostek, Nick , Reid, Beth A. , Seo, Hee-Jong , and Slosar, Anže . DESI and other Dark Energy experiments in the era of neutrino mass measurements. *J. Cosmology Astropart. Phys.*, 2014(5):023, May 2014. doi:10.1088/1475-7516/2014/05/023.

[104] Fowler, J. W. , Niemack, M. D. , Dicker, S. R. , Aboobaker, A. M. , Ade, P. A. R. , Battistelli, E. S. , Devlin, M. J. , Fisher, R. P. , et al. Optical design of the Atacama Cosmology Telescope and the Millimeter Bolometric Array Camera. *Appl. Opt.*, 46(17):3444–3454, June 2007. doi:10.1364/AO.46.003444.

[105] Freedman, Wendy L. , Madore, Barry F. , Hatt, Dylan , Hoyt, Taylor J. , Jang, In Sung , Beaton, Rachael L. , Burns, Christopher R. , Lee, Myung Gyoon , et al. The Carnegie-Chicago Hubble Program. VIII. An Independent Determination of the Hubble Constant Based on the Tip of the Red Giant Branch. *ApJ*, 882(1):34, September 2019. doi:10.3847/1538-4357/ab2f73.

[106] Gaia Collaboration, Brown, A. G. A. , Vallenari, A. , Prusti, T. , de Bruijne, J. H. J. , Babusiaux, C. , Biermann, M. , Creevey, O. L. , et al. Gaia Early Data Release 3. Summary of the contents and survey properties. *A&A*, 649:A1, May 2021. doi:10.1051/0004-6361/202039657.

[107] Gaia Collaboration, Prusti, T. , de Bruijne, J. H. J. , Brown, A. G. A. , Vallenari, A. , Babusiaux, C. , Bailer-Jones, C. A. L. , Bastian, U. , et al. The Gaia mission. *A&A*, 595:A1, November 2016. doi:10.1051/0004-6361/201629272.

[108] Gao, L. , Springel, V. , and White, S. D. M. . The age dependence of halo clustering. *MNRAS*, 363:L66–L70, October 2005. doi:10.1111/j.1745-3933.2005.00084.x.

[109] Garrison, Lehman H. , Eisenstein, Daniel J. , Ferrer, Douglas , Metchnik, Marc V. , and Pinto, Philip A. . Improving initial conditions for cosmological N-body simulations. *MNRAS*, 461(4):4125–4145, October 2016. doi:10.1093/mnras/stw1594.

[110] Garrison, Lehman H. , Eisenstein, Daniel J. , Ferrer, Douglas , Tinker, Jeremy L. , Pinto, Philip A. , and Weinberg, David H. . The Abacus Cosmos: A Suite of Cosmological N-body Simulations. *ApJS*, 236(2):43, June 2018. doi:10.3847/1538-4365/aabfd3.

[111] Gil-Marín, Héctor , Bautista, Julián E. , Paviot, Romain , Vargas-Magaña, Mariana , de la Torre, Sylvain , Fromenteau, Sebastien , Alam, Shadab , Ávila, Santiago , et al. The Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: measurement of the BAO and growth rate of structure of the luminous red galaxy sample from the anisotropic power spectrum between redshifts 0.6 and 1.0. *MNRAS*, 498(2):2492–2531, October 2020. doi:10.1093/mnras/staa2455.

[112] Gil-Marín, Héctor , Percival, Will J. , Cuesta, Antonio J. , Brownstein, Joel R. , Chuang, Chia-Hsun , Ho, Shirley , Kitaura, Francisco-Shu , Maraston, Claudia ,

et al. The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: BAO measurement from the LOS-dependent power spectrum of DR12 BOSS galaxies. *MNRAS*, 460(4):4210–4219, August 2016. doi:10.1093/mnras/stw1264.

[113] Gonçalves, R. S. , Carvalho, G. C. , Bengaly, C. A. P. , Carvalho, J. C. , and Alcaniz, J. S. . Measuring the scale of cosmic homogeneity with SDSS-IV DR14 quasars. *MNRAS*, 481(4):5270–5274, December 2018. doi:10.1093/mnras/sty2670.

[114] Grieb, Jan Niklas , Sánchez, Ariel G. , Salazar-Albornoz, Salvador , Scoccimarro, Román , Crocce, Martín , Dalla Vecchia, Claudio , Montesano, Francesco , Gil-Marín, Héctor , et al. The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological implications of the Fourier space wedges of the final sample. *Monthly Notices of the Royal Astronomical Society*, 467(2):2085–2112, May 2017. ISSN 0035-8711. doi:10.1093/mnras/stw3384. Publisher: Oxford Academic.

[115] Gunn, James E. . A Fundamental Limitation on the Accuracy of Angular Measurements in Observational Cosmology. *ApJ*, 147:61, January 1967. doi:10.1086/148981.

[116] Gunn, James E. and Gott, III , J. Richard. On the Infall of Matter Into Clusters of Galaxies and Some Effects on Their Evolution. *ApJ*, 176:1, August 1972. doi:10.1086/151605.

[117] Gunn, James E. , Siegmund, Walter A. , Mannery, Edward J. , Owen, Russell E. , Hull, Charles L. , Leger, R. French , Carey, Larry N. , Knapp, Gillian R. , et al. The 2.5 m Telescope of the Sloan Digital Sky Survey. *AJ*, 131(4):2332–2359, April 2006. doi:10.1086/500975.

[118] Guo, Hong , Zheng, Zheng , Zehavi, Idit , Dawson, Kyle , Skibba, Ramin A. , Tinker, Jeremy L. , Weinberg, David H. , White, Martin , et al. Velocity bias from the small-scale clustering of SDSS-III BOSS galaxies. *MNRAS*, 446(1):578–594, January 2015. doi:10.1093/mnras/stu2120.

[119] Guzzo, L. , Pierleoni, M. , Meneux, B. , Branchini, E. , Fèvre, O. Le , Marinoni, C. , Garilli, B. , Blaizot, J. , et al. A test of the nature of cosmic acceleration using galaxy redshift distortions. *Nature*, 451(7178):541–544, January 2008. ISSN 1476-4687. doi:10.1038/nature06555.

[120] Gwyn, Stephen D. J. . The Canada-France-Hawaii Telescope Legacy Survey: Stacked Images and Catalogs. *AJ*, 143(2):38, February 2012. doi:10.1088/0004-6256/143/2/38.

[121] Habib, Salman , Pope, Adrian , Finkel, Hal , Frontiere, Nicholas , Heitmann, Katrin , Daniel, David , Fasel, Patricia , Morozov, Vitali , et al. HACC: Simulating sky surveys on state-of-the-art supercomputing architectures. *New Astron.*, 42:49–65, January 2016. doi:10.1016/j.newast.2015.06.003.

[122] Hadzhiyska, Boryana , Bose, Sownak , Eisenstein, Daniel , and Hernquist, Lars . Extensions to models of the galaxy-halo connection. *MNRAS*, 501(2):1603–1620, February 2021. doi:10.1093/mnras/staa3776.

[123] Hadzhiyska, Boryana , Bose, Sownak , Eisenstein, Daniel , Hernquist, Lars , and Spergel, David N. . Limitations to the 'basic' HOD model and beyond. *MNRAS*, 493(4):5506–5519, April 2020. doi:10.1093/mnras/staa623.

[124] Hadzhiyska, Boryana , Tacchella, Sandro , Bose, Sownak , and Eisenstein, Daniel J. . The galaxy-halo connection of emission-line galaxies in IllustrisTNG. *MNRAS*, 502(3):3599–3617, April 2021. doi:10.1093/mnras/stab243.

[125] Hahn, ChangHoon , Wilson, Michael J. , Ruiz-Macias, Omar , Cole, Shaun , Weinberg, David H. , Moustakas, John , Kremin, Anthony , Tinker, Jeremy L. , et al. The DESI Bright Galaxy Survey: Final Target Selection, Design, and Validation. *AJ*, 165(6):253, June 2023. doi:10.3847/1538-3881/accff8.

[126] Hamana, Takashi , Shirasaki, Masato , Miyazaki, Satoshi , Hikage, Chiaki , Oguri, Masamune , More, Surhud , Armstrong, Robert , Leauthaud, Alexie , et al. Cosmological constraints from cosmic shear two-point correlation functions with HSC survey first-year data. *Publications of the Astronomicall Society of Japan*, 72(1):16, February 2020. doi:10.1093/pasj/psz138.

[127] Hamuy, Mario , Phillips, M. M. , Suntzeff, Nicholas B. , Schommer, Robert A. , Maza, Jose , and Aviles, R. . The Absolute Luminosities of the Calan/Tololo Type IA Supernovae. *AJ*, 112:2391, December 1996. doi:10.1086/118190.

[128] Harker, G. , Cole, S. , Helly, J. , Frenk, C. , and Jenkins, A. . A marked correlation function analysis of halo formation times in the Millennium Simulation. *MNRAS*, 367:1039–1049, April 2006. doi:10.1111/j.1365-2966.2006.10022.x.

[129] Harnois-Déraps, Joachim , Hernandez-Aguayo, Cesar , Cuesta-Lazaro, Carolina , Arnold, Christian , Li, Baojiu , Davies, Christopher T. , and Cai, Yan-Chuan . MGLenS: Modified gravity weak lensing simulations for emulation-based cosmological inference. *arXiv e-prints*, arXiv:2211.05779, November 2022. doi:10.48550/arXiv.2211.05779.

[130] Hartlap, J. , Simon, P. , and Schneider, P. . Why your model parameter confidences might be too optimistic. Unbiased estimation of the inverse covariance matrix. *A&A*, 464(1):399–404, March 2007. doi:10.1051/0004-6361:20066170.

[131] Hearin, Andrew P. , Campbell, Duncan , Tollerud, Erik , Behroozi, Peter , Diemer, Benedikt , Goldbaum, Nathan J. , Jennings, Elise , Leauthaud, Alexie , et al. Forward Modeling of Large-scale Structure: An Open-source Approach with Halotools. *AJ*, 154(5):190, November 2017. doi:10.3847/1538-3881/aa859f.

[132] Heitmann, Katrin , Finkel, Hal , Pope, Adrian , Morozov, Vitali , Frontiere, Nicholas , Habib, Salman , Rangel, Esteban , Uram, Thomas , et al. The Outer Rim Simulation: A Path to Many-core Supercomputers. *ApJS*, 245(1):16, November 2019. doi:10.3847/1538-4365/ab4da1.

[133] Heitmann, Katrin , Higdon, David , White, Martin , Habib, Salman , Williams, Brian J. , Lawrence, Earl , and Wagner, Christian . The Coyote Universe. II. Cosmological Models and Precision Emulation of the Nonlinear Matter Power Spectrum. *ApJ*, 705(1):156–174, November 2009. doi:10.1088/0004-637X/705/1/156.

[134] Heitmann, Katrin , Lawrence, Earl , Kwan, Juliana , Habib, Salman , and Higdon, David . The Coyote Universe Extended: Precision Emulation of the Matter Power Spectrum. *ApJ*, 780(1):111, January 2014. doi:10.1088/0004-637X/780/1/111.

[135] Heitmann, Katrin , White, Martin , Wagner, Christian , Habib, Salman , and Higdon, David . The Coyote Universe. I. Precision Determination of the Nonlinear Matter Power Spectrum. *ApJ*, 715(1):104–121, May 2010. doi:10.1088/0004-637X/715/1/104.

[136] Heymans, Catherine , Tröster, Tilman , Asgari, Marika , Blake, Chris , Hildebrandt, Hendrik , Joachimi, Benjamin , Kuijken, Konrad , Lin, Chieh-An , et al. KiDS-1000 Cosmology: Multi-probe weak gravitational lensing and spectroscopic galaxy clustering constraints. *A&A*, 646:A140, February 2021. doi:10.1051/0004-6361/202039063.

[137] Hill, J. Colin , McDonough, Evan , Toomey, Michael W. , and Alexander, Stephon . Early dark energy does not restore cosmological concordance. *Phys. Ref. D*, 102(4):043507, August 2020. doi:10.1103/PhysRevD.102.043507.

[138] Hinshaw, G. , Larson, D. , Komatsu, E. , Spergel, D. N. , Bennett, C. L. , Dunkley, J. , Nolta, M. R. , Halpern, M. , et al. Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results. *ApJS*, 208(2):19, October 2013. doi:10.1088/0067-0049/208/2/19.

[139] Hollinger, Amber M. and Hudson, Michael J. . Assessing the accuracy of cosmological parameters estimated from velocity - density comparisons via simulations. *MNRAS*, 502(3):3723–3732, April 2021. doi:10.1093/mnras/staa4039.

[140] Hopkins, Philip F. , Kereš, Dušan , Oñorbe, José , Faucher-Giguère, Claude-André , Quataert, Eliot , Murray, Norman , and Bullock, James S. . Galaxies on FIRE (Feedback In Realistic Environments): stellar feedback explains cosmologically inefficient star formation. *MNRAS*, 445(1):581–603, November 2014. doi: 10.1093/mnras/stu1738.

[141] Hou, Jiamin , Sánchez, Ariel G. , Ross, Ashley J. , Smith, Alex , Neveux, Richard , Bautista, Julian , Burtin, Etienne , Zhao, Cheng , et al. The completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: BAO and RSD measurements from anisotropic clustering analysis of the quasar sample in configuration space between redshift 0.8 and 2.2. *MNRAS*, 500(1):1201–1221, January 2021. doi: 10.1093/mnras/staa3234.

[142] Howlett, Cullan , Lewis, Antony , Hall, Alex , and Challinor, Anthony . CMB power spectrum parameter degeneracies in the era of precision cosmology. *JCAP*, 1204:027, 2012. doi:10.1088/1475-7516/2012/04/027.

[143] Howlett, Cullan , Ross, Ashley J. , Samushia, Lado , Percival, Will J. , and Manera, Marc . The clustering of the SDSS main galaxy sample - II. Mock galaxy catalogues and a measurement of the growth of structure from redshift space distortions at z = 0.15. *Monthly Notices of the Royal Astronomical Society*, 449:848–866, May 2015. ISSN 0035-8711. doi:10.1093/mnras/stu2693.

[144] Hoyle, F. and Fowler, William A. . Nucleosynthesis in Supernovae. *ApJ*, 132:565, November 1960. doi:10.1086/146963.

[145] Hu, Wayne and Okamoto, Takemi . Mass Reconstruction with Cosmic Microwave Background Polarization. *ApJ*, 574(2):566–574, August 2002. doi:10.1086/341110.

[146] Hubble, Edwin . A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae. *Proceedings of the National Academy of Science*, 15(3):168–173, March 1929. doi:10.1073/pnas.15.3.168.

[147] Ilbert, O. , Salvato, M. , Capak, P. , Le Floc'h, E. , Aussel, H. , McCracken, H. J. , Arnouts, S. , Mobasher, B. , et al. Photometric Redshifts and Stellar Mass Assembly in the 2-deg$\hat{2}$ COSMOS Field. In Kodama, T. , Yamada, T. , and Aoki, K. , editors,

Panoramic Views of Galaxy Formation and Evolution, volume 399 of *Astronomical Society of the Pacific Conference Series*, page 169. October 2008.

[148] Ivanov, Mikhail M. , McDonough, Evan , Hill, J. Colin , Simonović, Marko , Toomey, Michael W. , Alexander, Stephon , and Zaldarriaga, Matias . Constraining early dark energy with large-scale structure. *Phys. Ref. D*, 102(10):103502, November 2020. doi: 10.1103/PhysRevD.102.103502.

[149] Ivanov, Mikhail M. , Simonović, Marko , and Zaldarriaga, Matias . Cosmological parameters from the BOSS galaxy power spectrum. *J. Cosmology Astropart. Phys.*, 2020(5):042, May 2020. doi:10.1088/1475-7516/2020/05/042.

[150] Joachimi, B. , Lin, C. A. , Asgari, M. , Tröster, T. , Heymans, C. , Hildebrandt, H. , Köhlinger, F. , Sánchez, A. G. , et al. KiDS-1000 methodology: Modelling and inference for joint weak gravitational lensing and spectroscopic galaxy clustering analysis. *A&A*, 646:A129, February 2021. doi:10.1051/0004-6361/202038831.

[151] Jones, D. Heath , Read, Mike A. , Saunders, Will , Colless, Matthew , Jarrett, Tom , Parker, Quentin A. , Fairall, Anthony P. , Mauch, Thomas , et al. The 6dF Galaxy Survey: final redshift release (DR3) and southern large-scale structures. *MNRAS*, 399(2):683–698, October 2009. doi:10.1111/j.1365-2966.2009.15338.x.

[152] Kaiser, N. . Clustering in real space and in redshift space. *MNRAS*, 227:1–21, July 1987. doi:10.1093/mnras/227.1.1.

[153] Kaiser, Nick . Weak Gravitational Lensing of Distant Galaxies. *ApJ*, 388:272, April 1992. doi:10.1086/171151.

[154] Kaiser, Nick , Burgett, William , Chambers, Ken , Denneau, Larry , Heasley, Jim , Jedicke, Robert , Magnier, Eugene , Morgan, Jeff , et al. The Pan-STARRS wide-field optical/NIR imaging survey. In Stepp, Larry M. , Gilmozzi, Roberto , and Hall, Helen J. , editors, *Ground-based and Airborne Telescopes III*, volume 7733 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 77330E. July 2010. doi:10.1117/12.859188.

[155] Kaiser, Nick , Wilson, Gillian , and Luppino, Gerard A. . Large-Scale Cosmic Shear Measurements. *arXiv e-prints*, astro-ph/0003338, March 2000. doi:10.48550/arXiv.astro-ph/0003338.

[156] Kirkby, David , Margala, Daniel , Slosar, Anže , Bailey, Stephen , Busca, Nicolás G. , Delubac, Timothée , Rich, James , Bautista, Julian E. , et al. Fitting methods for

baryon acoustic oscillations in the Lyman-α forest fluctuations in BOSS data release 9. *J. Cosmology Astropart. Phys.*, 2013(3):024, March 2013. doi:10.1088/1475-7516/2013/03/024.

[157] Klypin, Anatoly and Prada, Francisco . Dark matter statistics for large galaxy catalogues: power spectra and covariance matrices. *MNRAS*, 478(4):4602–4621, August 2018. doi:10.1093/mnras/sty1340.

[158] Klypin, Anatoly , Yepes, Gustavo , Gottlöber, Stefan , Prada, Francisco , and Heß, Steffen . MultiDark simulations: the story of dark matter halo concentrations and density profiles. *MNRAS*, 457(4):4340–4359, April 2016. doi:10.1093/mnras/stw248.

[159] Knebe, Alexander , Knollmann, Steffen R. , Muldrew, Stuart I. , Pearce, Frazer R. , Aragon-Calvo, Miguel Angel , Ascasibar, Yago , Behroozi, Peter S. , Ceverino, Daniel , et al. Haloes gone MAD: The Halo-Finder Comparison Project. *MNRAS*, 415(3):2293–2318, August 2011. doi:10.1111/j.1365-2966.2011.18858.x.

[160] Knox, L. and Millea, M. . Hubble constant hunter's guide. *Phys. Ref. D*, 101(4):043533, February 2020. doi:10.1103/PhysRevD.101.043533.

[161] Kobayashi, Yosuke , Nishimichi, Takahiro , Takada, Masahiro , and Miyatake, Hironao . Full-shape cosmology analysis of the SDSS-III BOSS galaxy power spectrum using an emulator-based halo model: A 5% determination of σ_8. *Phys. Ref. D*, 105(8):083517, April 2022. doi:10.1103/PhysRevD.105.083517.

[162] Krause, Elisabeth , Hirata, Christopher M. , Martin, Christopher , Neill, James D. , and Wyder, Ted K. . Halo occupation distribution modelling of green valley galaxies. *MNRAS*, 428(3):2548–2564, January 2013. doi:10.1093/mnras/sts221.

[163] Kuijken, Konrad , Heymans, Catherine , Hildebrandt, Hendrik , Nakajima, Reiko , Erben, Thomas , de Jong, Jelte T. A. , Viola, Massimo , Choi, Ami , et al. Gravitational lensing analysis of the Kilo-Degree Survey. *MNRAS*, 454(4):3500–3532, December 2015. doi:10.1093/mnras/stv2140.

[164] Landy, Stephen D. and Szalay, Alexander S. . Bias and Variance of Angular Correlation Functions. *ApJ*, 412:64, July 1993. doi:10.1086/172900.

[165] Lang, Dustin , Hogg, David W. , and Schlegel, David J. . WISE Photometry for 400 Million SDSS Sources. *AJ*, 151(2):36, February 2016. doi:10.3847/0004-6256/151/2/36.

[166] Lange, Johannes U. , Hearin, Andrew P. , Leauthaud, Alexie , van den Bosch, Frank C. , Guo, Hong , and DeRose, Joseph . Five per cent measurements of the growth rate from simulation-based modelling of redshift-space clustering in BOSS LOWZ. *MNRAS*, 509(2):1779–1804, January 2022. doi:10.1093/mnras/stab3111.

[167] Lange, Johannes U. , van den Bosch, Frank C. , Zentner, Andrew R. , Wang, Kuan , Hearin, Andrew P. , and Guo, Hong . Cosmological Evidence Modelling: a new simulation-based approach to constrain cosmology on non-linear scales. *MNRAS*, 490(2):1870–1878, December 2019. doi:10.1093/mnras/stz2664.

[168] Laureijs, R. , Amiaux, J. , Arduini, S. , Augueres, J. L. , Brinchmann, J. , Cole, R. , Cropper, M. , Dabin, C. , et al. Euclid Definition Study Report. *arXiv e-prints*, arXiv:1110.3193, October 2011.

[169] Law, Nicholas M. , Kulkarni, Shrinivas R. , Dekany, Richard G. , Ofek, Eran O. , Quimby, Robert M. , Nugent, Peter E. , Surace, Jason , Grillmair, Carl C. , et al. The Palomar Transient Factory: System Overview, Performance, and First Results. *PASP*, 121(886):1395, December 2009. doi:10.1086/648598.

[170] Lawrence, Earl , Heitmann, Katrin , White, Martin , Higdon, David , Wagner, Christian , Habib, Salman , and Williams, Brian . The Coyote Universe. III. Simulation Suite and Precision Emulator for the Nonlinear Matter Power Spectrum. *ApJ*, 713(2):1322–1331, April 2010. doi:10.1088/0004-637X/713/2/1322.

[171] Leauthaud, Alexie , Saito, Shun , Hilbert, Stefan , Barreira, Alexandre , More, Surhud , White, Martin , Alam, Shadab , Behroozi, Peter , et al. Lensing is low: cosmology, galaxy formation or new physics? *MNRAS*, 467(3):3024–3047, May 2017. doi:10.1093/mnras/stx258.

[172] Leavitt, Henrietta S. and Pickering, Edward C. . Periods of 25 Variable Stars in the Small Magellanic Cloud. *Harvard College Observatory Circular*, 173:1–3, March 1912.

[173] Lewis, Antony . Efficient sampling of fast and slow cosmological parameters. *Phys. Rev.*, D87(10):103529, 2013. doi:10.1103/PhysRevD.87.103529.

[174] Lewis, Antony and Bridle, Sarah . Cosmological parameters from CMB and other data: A Monte Carlo approach. *Phys. Rev.*, D66:103511, 2002. doi:10.1103/PhysRevD.66.103511.

[175] Lewis, Antony and Challinor, Anthony . Weak gravitational lensing of the CMB. *Phys. Rept.*, 429(1):1–65, June 2006. doi:10.1016/j.physrep.2006.03.002.

[176] Lewis, Antony , Challinor, Anthony , and Lasenby, Anthony . Efficient computation of CMB anisotropies in closed FRW models. *Astrophys. J.*, 538:473–476, 2000. doi: 10.1086/309179.

[177] Li, Yin , Hu, Wayne , and Takada, Masahiro . Super-sample covariance in simulations. *Phys. Ref. D*, 89(8):083519, April 2014. doi:10.1103/PhysRevD.89.083519.

[178] Lin, Sicheng , Tinker, Jeremy L. , Klypin, Anatoly , Prada, Francisco , Blanton, Michael R. , Comparat, Johan , Dawson, Kyle S. , de Mattia, Arnaud , et al. The completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: GLAM-QPM mock galaxy catalogues for the emission line galaxy sample. *MNRAS*, 498(4):5251–5262, November 2020. doi:10.1093/mnras/staa2571.

[179] Lyke, Brad W. , Higley, Alexandra N. , McLane, J. N. , Schurhammer, Danielle P. , Myers, Adam D. , Ross, Ashley J. , Dawson, Kyle , Chabanier, Solène , et al. The Sloan Digital Sky Survey Quasar Catalog: Sixteenth Data Release. *ApJS*, 250(1):8, September 2020. doi:10.3847/1538-4365/aba623.

[180] Maksimova, Nina A. , Garrison, Lehman H. , Eisenstein, Daniel J. , Hadzhiyska, Boryana , Bose, Sownak , and Satterthwaite, Thomas P. . ABACUSSUMMIT: a massive set of high-accuracy, high-resolution N-body simulations. *MNRAS*, 508(3):4017–4037, December 2021. doi:10.1093/mnras/stab2484.

[181] Mandelbaum, Rachel . Weak Lensing for Precision Cosmology. *ARA&A*, 56:393–433, September 2018. doi:10.1146/annurev-astro-081817-051928.

[182] Massara, Elena , Villaescusa-Navarro, Francisco , Hahn, ChangHoon , Abidi, Muntazir M. , Eickenberg, Michael , Ho, Shirley , Lemos, Pablo , Dizgah, Azadeh Moradinezhad , et al. Cosmological Information in the Marked Power Spectrum of the Galaxy Field. *ApJ*, 951(1):70, July 2023. doi:10.3847/1538-4357/acd44d.

[183] McCarthy, Ian G. , Schaye, Joop , Bird, Simeon , and Le Brun, Amandine M. C. . The BAHAMAS project: calibrated hydrodynamical simulations for large-scale structure cosmology. *MNRAS*, 465(3):2936–2965, March 2017. doi:10.1093/mnras/stw2792.

[184] McDonald, Patrick . Toward a Measurement of the Cosmological Geometry at z ~2: Predicting Lyα Forest Correlation in Three Dimensions and the Potential of Future Data Sets. *ApJ*, 585(1):34–51, March 2003. doi:10.1086/345945.

[185] Miralda-Escude, Jordi . The Correlation Function of Galaxy Ellipticities Produced by Gravitational Lensing. *ApJ*, 380:1, October 1991. doi:10.1086/170555.

[186] Mo, H. J. and White, S. D. M. . An analytic model for the spatial clustering of dark matter haloes. *MNRAS*, 282(2):347–361, September 1996. doi:10.1093/mnras/282.2.347.

[187] Mo, Houjun , van den Bosch, Frank C. , and White, Simon . *Galaxy Formation and Evolution*. Cambridge University Press, 2010.

[188] Mohammad, F. G. , Bianchi, D. , Percival, W. J. , de la Torre, S. , Guzzo, L. , Granett, B. R. , Branchini, E. , Bolzonella, M. , et al. The VIMOS Public Extragalactic Redshift Survey (VIPERS). Unbiased clustering estimate with VIPERS slit assignment. *A&A*, 619:A17, November 2018. doi:10.1051/0004-6361/201833853.

[189] Mohammad, Faizan G. and Percival, Will J. . Creating jackknife and bootstrap estimates of the covariance matrix for the two-point correlation function. *MNRAS*, 514(1):1289–1301, July 2022. doi:10.1093/mnras/stac1458.

[190] Mohammad, Faizan G. , Percival, Will J. , Seo, Hee-Jong , Chapman, Michael J. , Bianchi, D. , Ross, Ashley J. , Zhao, Cheng , Lang, Dustin , et al. The completed SDSS-IV extended baryon oscillation spectroscopic survey: pairwise-inverse probability and angular correction for fibre collisions in clustering measurements. *MNRAS*, 498(1):128–143, October 2020. doi:10.1093/mnras/staa2344.

[191] Morganson, Eric , Green, Paul J. , Anderson, Scott F. , Ruan, John J. , Myers, Adam D. , Eracleous, Michael , Kelly, Brandon , Badenes, Carlos , et al. The Time Domain Spectroscopic Survey: Variable Selection and Anticipated Results. *ApJ*, 806(2):244, June 2015. doi:10.1088/0004-637X/806/2/244.

[192] Moster, Benjamin P. , Naab, Thorsten , and White, Simon D. M. . Galactic star formation and accretion histories from matching galaxies to dark matter haloes. *MNRAS*, 428(4):3121–3138, February 2013. doi:10.1093/mnras/sts261.

[193] Myers, Adam D. , Palanque-Delabrouille, Nathalie , Prakash, Abhishek , Pâris, Isabelle , Yeche, Christophe , Dawson, Kyle S. , Bovy, Jo , Lang, Dustin , et al. The SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Quasar Target Selection. *ApJS*, 221(2):27, December 2015. doi:10.1088/0067-0049/221/2/27.

[194] Navarro, Julio F. , Frenk, Carlos S. , and White, Simon D. M. . The Structure of Cold Dark Matter Halos. *ApJ*, 462:563, May 1996. doi:10.1086/177173.

[195] Neal, R. M. . Taking Bigger Metropolis Steps by Dragging Fast Variables. *ArXiv Mathematics e-prints*, February 2005.

[196] Neveux, Richard , Burtin, Etienne , de Mattia, Arnaud , Smith, Alex , Ross, Ashley J. , Hou, Jiamin , Bautista, Julian , Brinkmann, Jonathan , et al. The completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: BAO and RSD measurements from the anisotropic power spectrum of the quasar sample between redshift 0.8 and 2.2. *MNRAS*, 499(1):210–229, November 2020. doi:10.1093/mnras/staa2780.

[197] Obuljen, Andrej , Dalal, Neal , and Percival, Will J. . Anisotropic halo assembly bias and redshift-space distortions. *J. Cosmology Astropart. Phys.*, 2019(10):020, October 2019. doi:10.1088/1475-7516/2019/10/020.

[198] Ostriker, J. P. , Peebles, P. J. E. , and Yahil, A. . The Size and Mass of Galaxies, and the Mass of the Universe. *ApJ*, 193:L1, October 1974. doi:10.1086/181617.

[199] Padmanabhan, Nikhil , Schlegel, David J. , Finkbeiner, Douglas P. , Barentine, J. C. , Blanton, Michael R. , Brewington, Howard J. , Gunn, James E. , Harvanek, Michael , et al. An Improved Photometric Calibration of the Sloan Digital Sky Survey Imaging Data. *ApJ*, 674(2):1217–1233, February 2008. doi:10.1086/524677.

[200] Paillas, Enrique , Cuesta-Lazaro, Carolina , Zarrouk, Pauline , Cai, Yan-Chuan , Percival, Will J. , Nadathur, Seshadri , Pinon, Mathilde , de Mattia, Arnaud , et al. Constraining $\nu\Lambda$CDM with density-split clustering. *MNRAS*, 522(1):606–625, June 2023. doi:10.1093/mnras/stad1017.

[201] Peebles, P. J. E. . *The large-scale structure of the universe*. Princeton University Press, 1980.

[202] Percival, Will J. and Bianchi, Davide . Using angular pair upweighting to improve 3D clustering measurements. *MNRAS*, 472(1):L40–L44, November 2017. doi:10.1093/mnrasl/slx135.

[203] Percival, Will J. , Cole, Shaun , Eisenstein, Daniel J. , Nichol, Robert C. , Peacock, John A. , Pope, Adrian C. , and Szalay, Alexander S. . Measuring the Baryon Acoustic Oscillation scale using the Sloan Digital Sky Survey and 2dF Galaxy Redshift Survey. *MNRAS*, 381(3):1053–1066, November 2007. doi:10.1111/j.1365-2966.2007.12268.x.

[204] Perko, Ashley , Senatore, Leonardo , Jennings, Elise , and Wechsler, Risa H. . Biased Tracers in Redshift Space in the EFT of Large-Scale Structure. *arXiv e-prints*, arXiv:1610.09321, October 2016. doi:10.48550/arXiv.1610.09321.

[205] Perlmutter, S. , Aldering, G. , Goldhaber, G. , Knop, R. A. , Nugent, P. , Castro, P. G. , Deustua, S. , Fabbro, S. , et al. Measurements of Ω and Λ from 42 High-Redshift Supernovae. *ApJ*, 517(2):565–586, June 1999. doi:10.1086/307221.

[206] Pesce, D. W. , Braatz, J. A. , Reid, M. J. , Riess, A. G. , Scolnic, D. , Condon, J. J. , Gao, F. , Henkel, C. , et al. The Megamaser Cosmology Project. XIII. Combined Hubble Constant Constraints. *ApJ*, 891(1):L1, March 2020. doi:10.3847/2041-8213/ab75f0.

[207] Pettini, Max , Zych, Berkeley J. , Murphy, Michael T. , Lewis, Antony , and Steidel, Charles C. . Deuterium abundance in the most metal-poor damped Lyman alpha system: converging on $\Omega_{b,0}h^2$. *MNRAS*, 391(4):1499–1510, December 2008. doi:10.1111/j.1365-2966.2008.13921.x.

[208] Pezzotta, A. , de la Torre, S. , Bel, J. , Granett, B. R. , Guzzo, L. , Peacock, J. A. , Garilli, B. , Scodeggio, M. , et al. The VIMOS Public Extragalactic Redshift Survey (VIPERS). The growth of structure at $0.5 < z < 1.2$ from redshift-space distortions in the clustering of the PDR-2 final sample. *A&A*, 604:A33, July 2017. doi:10.1051/0004-6361/201630295.

[209] Planck Collaboration, Ade, P. A. R. , Aghanim, N. , Armitage-Caplan, C. , Arnaud, M. , Ashdown, M. , Atrio-Barandela, F. , Aumont, J. , et al. Planck 2013 results. XVI. Cosmological parameters. *A&A*, 571:A16, November 2014. doi:10.1051/0004-6361/201321591.

[210] Planck Collaboration, Ade, P. A. R. , Aghanim, N. , Arnaud, M. , Ashdown, M. , Aumont, J. , Baccigalupi, C. , Banday, A. J. , et al. Planck 2015 results. XIII. Cosmological parameters. *A&A*, 594:A13, September 2016. doi:10.1051/0004-6361/201525830.

[211] Planck Collaboration, Aghanim, N. , Akrami, Y. , Arroja, F. , Ashdown, M. , Aumont, J. , Baccigalupi, C. , Ballardini, M. , et al. Planck 2018 results. I. Overview and the cosmological legacy of Planck. *A&A*, 641:A1, September 2020. doi:10.1051/0004-6361/201833880.

[212] Planck Collaboration, Aghanim, N. , Akrami, Y. , Ashdown, M. , Aumont, J. , Baccigalupi, C. , Ballardini, M. , Banday, A. J. , et al. Planck 2018 results. V. CMB power spectra and likelihoods. *A&A*, 641:A5, September 2020. doi:10.1051/0004-6361/201936386.

[213] Planck Collaboration, Aghanim, N. , Akrami, Y. , Ashdown, M. , Aumont, J. , Baccigalupi, C. , Ballardini, M. , Banday, A. J. , et al. Planck 2018 results. VI. Cosmological parameters. *A&A*, 641:A6, September 2020. doi:10.1051/0004-6361/201833910.

[214] Planck Collaboration, Aghanim, N. , Akrami, Y. , Ashdown, M. , Aumont, J. , Baccigalupi, C. , Ballardini, M. , Banday, A. J. , et al. Planck 2018 results. VIII. Gravitational lensing. *A&A*, 641:A8, September 2020. doi:10.1051/0004-6361/201833886.

[215] Planck Collaboration, Akrami, Y. , Arroja, F. , Ashdown, M. , Aumont, J. , Baccigalupi, C. , Ballardini, M. , Banday, A. J. , et al. Planck 2018 results. X. Constraints on inflation. *A&A*, 641:A10, September 2020. doi:10.1051/0004-6361/201833887.

[216] Planck Collaboration, Akrami, Y. , Ashdown, M. , Aumont, J. , Baccigalupi, C. , Ballardini, M. , Banday, A. J. , Barreiro, R. B. , et al. Planck 2018 results. VII. Isotropy and statistics of the CMB. *A&A*, 641:A7, September 2020. doi:10.1051/0004-6361/201935201.

[217] Porto, Rafael A. , Senatore, Leonardo , and Zaldarriaga, Matias . The Lagrangian-space Effective Field Theory of large scale structures. *J. Cosmology Astropart. Phys.*, 2014(5):022, May 2014. doi:10.1088/1475-7516/2014/05/022.

[218] Pozzetti, L. , Hirata, C. M. , Geach, J. E. , Cimatti, A. , Baugh, C. , Cucciati, O. , Merson, A. , Norberg, P. , et al. Modelling the number density of Hα emitters for future spectroscopic near-IR space missions. *A&A*, 590:A3, May 2016. doi:10.1051/0004-6361/201527081.

[219] Prakash, Abhishek , Licquia, Timothy C. , Newman, Jeffrey A. , Ross, Ashley J. , Myers, Adam D. , Dawson, Kyle S. , Kneib, Jean-Paul , Percival, Will J. , et al. The SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Luminous Red Galaxy Target Selection. *ApJS*, 224(2):34, June 2016. doi:10.3847/0067-0049/224/2/34.

[220] Preston, Calvin , Amon, Alexandra , and Efstathiou, George . A non-linear solution to the S_8 tension II: Analysis of DES Year 3 cosmic shear. *arXiv e-prints*, arXiv:2305.09827, May 2023. doi:10.48550/arXiv.2305.09827.

[221] Raichoor, A. , Comparat, J. , Delubac, T. , Kneib, J. P. , Yèche, Ch , Dawson, K. S. , Percival, W. J. , Dey, A. , et al. The SDSS-IV extended Baryon Oscillation Spectroscopic Survey: final emission line galaxy target selection. *MNRAS*, 471(4):3955–3973, November 2017. doi:10.1093/mnras/stx1790.

[222] Raichoor, A. , Moustakas, J. , Newman, Jeffrey A. , Karim, T. , Ahlen, S. , Alam, Shadab , Bailey, S. , Brooks, D. , et al. Target Selection and Validation of DESI Emission Line Galaxies. *AJ*, 165(3):126, March 2023. doi:10.3847/1538-3881/acb213.

[223] Raichoor, Anand , de Mattia, Arnaud , Ross, Ashley J. , Zhao, Cheng , Alam, Shadab , Avila, Santiago , Bautista, Julian , Brinkmann, Jonathan , et al. The completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: large-scale structure catalogues and measurement of the isotropic BAO between redshift 0.6 and 1.1 for the Emission Line Galaxy Sample. *MNRAS*, 500(3):3254–3274, January 2021. doi:10.1093/mnras/staa3336.

[224] Rau, Arne , Kulkarni, Shrinivas R. , Law, Nicholas M. , Bloom, Joshua S. , Ciardi, David , Djorgovski, George S. , Fox, Derek B. , Gal-Yam, Avishay , et al. Exploring the Optical Transient Sky with the Palomar Transient Factory. *PASP*, 121(886):1334, December 2009. doi:10.1086/605911.

[225] Reid, Beth , Ho, Shirley , Padmanabhan, Nikhil , Percival, Will J. , Tinker, Jeremy , Tojeiro, Rita , White, Martin , Eisenstein, Daniel J. , et al. SDSS-III Baryon Oscillation Spectroscopic Survey Data Release 12: galaxy target selection and large-scale structure catalogues. *MNRAS*, 455(2):1553–1573, January 2016. doi:10.1093/mnras/stv2382.

[226] Reid, Beth A. , Seo, Hee-Jong , Leauthaud, Alexie , Tinker, Jeremy L. , and White, Martin . A 2.5 per cent measurement of the growth rate from small-scale redshift space clustering of SDSS-III CMASS galaxies. *MNRAS*, 444(1):476–502, October 2014. doi:10.1093/mnras/stu1391.

[227] Reid, Beth A. and White, Martin . Towards an accurate model of the redshift-space clustering of haloes in the quasi-linear regime. *MNRAS*, 417(3):1913–1927, November 2011. doi:10.1111/j.1365-2966.2011.19379.x.

[228] Riess, Adam G. , Filippenko, Alexei V. , Challis, Peter , Clocchiatti, Alejandro , Diercks, Alan , Garnavich, Peter M. , Gilliland, Ron L. , Hogan, Craig J. , et al. Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. *AJ*, 116(3):1009–1038, September 1998. doi:10.1086/300499.

[229] Riess, Adam G. , Yuan, Wenlong , Macri, Lucas M. , Scolnic, Dan , Brout, Dillon , Casertano, Stefano , Jones, David O. , Murakami, Yukei , et al. A Comprehensive Measurement of the Local Value of the Hubble Constant with 1 km s^{-1} Mpc^{-1}

Uncertainty from the Hubble Space Telescope and the SH0ES Team. *ApJ*, 934(1):L7, July 2022. doi:10.3847/2041-8213/ac5c5b.

[230] Rocher, Antoine , Ruhlmann-Kleider, Vanina , Burtin, Etienne , Yuan, Sihan , de Mattia, Arnaud , Ross, Ashley J. , Aguilar, Jessica , Ahlen, Steven , et al. The DESI One-Percent survey: exploring the Halo Occupation Distribution of Emission Line Galaxies with AbacusSummit simulations. *arXiv e-prints*, arXiv:2306.06319, June 2023. doi:10.48550/arXiv.2306.06319.

[231] Ross, Ashley J. , Bautista, Julian , Tojeiro, Rita , Alam, Shadab , Bailey, Stephen , Burtin, Etienne , Comparat, Johan , Dawson, Kyle S. , et al. The Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Large-scale structure catalogues for cosmological analysis. *MNRAS*, 498(2):2354–2371, October 2020. doi: 10.1093/mnras/staa2416.

[232] Ross, Ashley J. , Percival, Will J. , Sánchez, Ariel G. , Samushia, Lado , Ho, Shirley , Kazin, Eyal , Manera, Marc , Reid, Beth , et al. The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: analysis of potential systematics. *MNRAS*, 424(1):564–590, July 2012. doi:10.1111/j.1365-2966.2012.21235.x.

[233] Ross, Ashley J. , Samushia, Lado , Howlett, Cullan , Percival, Will J. , Burden, Angela , and Manera, Marc . The clustering of the SDSS DR7 main Galaxy sample - I. A 4 per cent distance measure at z = 0.15. *MNRAS*, 449(1):835–847, May 2015. doi:10.1093/mnras/stv154.

[234] Rossi, Graziano , Choi, Peter D. , Moon, Jeongin , Bautista, Julian E. , Gil-Marín, Hector , Paviot, Romain , Vargas-Magaña, Mariana , de la Torre, Sylvain , et al. The Completed SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: N-body Mock Challenge for Galaxy Clustering Measurements. *MNRAS*, 505(1):377–407, July 2021. doi:10.1093/mnras/staa3955.

[235] Ruan, Cheng-Zong , Cuesta-Lazaro, Carolina , Eggemeier, Alexander , Li, Baojiu , Baugh, Carlton M. , Arnold, Christian , Bose, Sownak , Hernández-Aguayo, César , et al. An emulator-based halo model in modified gravity – I. The halo concentration-mass relation and density profile. *arXiv e-prints*, arXiv:2301.02970, January 2023. doi:10.48550/arXiv.2301.02970.

[236] Rubin, V. C. , Ford, Jr. , W. K., and Thonnard, N. . Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 (R=4kpc) to UGC 2885 (R=122kpc). *ApJ*, 238:471–487, June 1980. doi:10.1086/158003.

[237] Sachs, R. K. and Wolfe, A. M. . Perturbations of a Cosmological Model and Angular Variations of the Microwave Background. *ApJ*, 147:73, January 1967. doi:10.1086/148982.

[238] Samushia, Lado , Reid, Beth A. , White, Martin , Percival, Will J. , Cuesta, Antonio J. , Lombriser, Lucas , Manera, Marc , Nichol, Robert C. , et al. The clustering of galaxies in the SDSS-III DR9 Baryon Oscillation Spectroscopic Survey: testing deviations from Λ and general relativity using anisotropic clustering of galaxies. *MNRAS*, 429(2):1514–1528, February 2013. doi:10.1093/mnras/sts443.

[239] Samushia, Lado , Slepian, Zachary , and Villaescusa-Navarro, Francisco . Information content of higher order galaxy correlation functions. *MNRAS*, 505(1):628–641, July 2021. doi:10.1093/mnras/stab1199.

[240] Sandage, A. and Tammann, G. A. . Steps toward the Hubble constant. V. The Hubble constant from nearby galaxies and the regularity of the local velocity field. *ApJ*, 196:313–328, March 1975. doi:10.1086/153413.

[241] Satpathy, Siddharth , Alam, Shadab , Ho, Shirley , White, Martin , Bahcall, Neta A. , Beutler, Florian , Brownstein, Joel R. , Chuang, Chia-Hsun , et al. The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: on the measurement of growth rate using galaxy correlation functions. *Monthly Notices of the Royal Astronomical Society*, 469:1369–1382, August 2017. ISSN 0035-8711. doi:10.1093/mnras/stx883.

[242] Schlafly, E. F. , Finkbeiner, D. P. , Jurić, M. , Magnier, E. A. , Burgett, W. S. , Chambers, K. C. , Grav, T. , Hodapp, K. W. , et al. Photometric Calibration of the First 1.5 Years of the Pan-STARRS1 Survey. *ApJ*, 756(2):158, September 2012. doi:10.1088/0004-637X/756/2/158.

[243] Schneider, Aurel , Teyssier, Romain , Potter, Doug , Stadel, Joachim , Onions, Julian , Reed, Darren S. , Smith, Robert E. , Springel, Volker , et al. Matter power spectrum and the challenge of percent accuracy. *J. Cosmology Astropart. Phys.*, 2016(4):047, April 2016. doi:10.1088/1475-7516/2016/04/047.

[244] Scoccimarro, Román , Zaldarriaga, Matias , and Hui, Lam . Power Spectrum Correlations Induced by Nonlinear Clustering. *ApJ*, 527(1):1–15, December 1999. doi:10.1086/308059.

[245] Sheth, R. K. and Tormen, G. . On the environmental dependence of halo formation. *MNRAS*, 350:1385–1390, June 2004. doi:10.1111/j.1365-2966.2004.07733.x.

[246] Slosar, Anže , Iršič, Vid , Kirkby, David , Bailey, Stephen , Busca, Nicolás G. , Delubac, Timothée , Rich, James , Aubourg, Éric , et al. Measurement of baryon acoustic oscillations in the Lyman-α forest fluctuations in BOSS data release 9. *J. Cosmology Astropart. Phys.*, 2013(4):026, April 2013. doi:10.1088/1475-7516/2013/04/026.

[247] Smee, Stephen A. , Gunn, James E. , Uomoto, Alan , Roe, Natalie , Schlegel, David , Rockosi, Constance M. , Carr, Michael A. , Leger, French , et al. The Multi-object, Fiber-fed Spectrographs for the Sloan Digital Sky Survey and the Baryon Oscillation Spectroscopic Survey. *AJ*, 146(2):32, August 2013. doi:10.1088/0004-6256/146/2/32.

[248] Smith, Alex , Burtin, Etienne , Hou, Jiamin , Neveux, Richard , Ross, Ashley J. , Alam, Shadab , Brinkmann, Jonathan , Dawson, Kyle S. , et al. The completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: N-body mock challenge for the quasar sample. *MNRAS*, 499(1):269–291, November 2020. doi:10.1093/mnras/staa2825.

[249] Smith, Alex , He, Jian-hua , Cole, Shaun , Stothert, Lee , Norberg, Peder , Baugh, Carlton , Bianchi, Davide , Wilson, Michael J. , et al. Correcting for fibre assignment incompleteness in the DESI Bright Galaxy Survey. *MNRAS*, 484(1):1285–1300, March 2019. doi:10.1093/mnras/stz059.

[250] Smoot, G. F. , Bennett, C. L. , Kogut, A. , Wright, E. L. , Aymon, J. , Boggess, N. W. , Cheng, E. S. , de Amici, G. , et al. Structure in the COBE Differential Microwave Radiometer First-Year Maps. *ApJ*, 396:L1, September 1992. doi:10.1086/186504.

[251] Song, Yong-Seon and Percival, Will J. . Reconstructing the history of structure formation using redshift distortions. *Journal of Cosmology and Astroparticle Physics*, 2009(10):004–004, October 2009. ISSN 1475-7516. doi:10.1088/1475-7516/2009/10/004. Publisher: IOP Publishing.

[252] Sotiriou, Thomas P. and Faraoni, Valerio . f(R) theories of gravity. *Reviews of Modern Physics*, 82(1):451–497, January 2010. doi:10.1103/RevModPhys.82.451.

[253] Springel, Volker , Pakmor, Rüdiger , Pillepich, Annalisa , Weinberger, Rainer , Nelson, Dylan , Hernquist, Lars , Vogelsberger, Mark , Genel, Shy , et al. First results from the IllustrisTNG simulations: matter and galaxy clustering. *MNRAS*, 475(1):676–698, March 2018. doi:10.1093/mnras/stx3304.

[254] Sunyaev, R. A. and Zeldovich, Ya. B. . Small-Scale Fluctuations of Relic Radiation. *Astrophys. Space Sci.*, 7(1):3–19, April 1970. doi:10.1007/BF00653471.

[255] Sunyaev, R. A. and Zeldovich, Ya. B. . The Observations of Relic Radiation as a Test of the Nature of X-Ray Radiation from the Clusters of Galaxies. *Comments on Astrophysics and Space Physics*, 4:173, November 1972.

[256] Sunyaev, R. A. and Zeldovich, Ya. B. . The velocity of clusters of galaxies relative to the microwave background - The possibility of its measurement. *MNRAS*, 190:413–420, February 1980. doi:10.1093/mnras/190.3.413.

[257] Sánchez, Ariel G. , Scoccimarro, Román , Crocce, Martín , Grieb, Jan Niklas , Salazar-Albornoz, Salvador , Dalla Vecchia, Claudio , Lippich, Martha , Beutler, Florian , et al. The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: Cosmological implications of the configuration-space clustering wedges. *Monthly Notices of the Royal Astronomical Society*, 464:1640–1658, January 2017. ISSN 0035-8711. doi:10.1093/mnras/stw2443.

[258] Tamone, Amélie , Raichoor, Anand , Zhao, Cheng , de Mattia, Arnaud , Gorgoni, Claudio , Burtin, Etienne , Ruhlmann-Kleider, Vanina , Ross, Ashley J. , et al. The completed SDSS-IV extended baryon oscillation spectroscopic survey: growth rate of structure measurement from anisotropic clustering analysis in configuration space between redshift 0.6 and 1.1 for the emission-line galaxy sample. *MNRAS*, 499(4):5527–5546, December 2020. doi:10.1093/mnras/staa3050.

[259] Taruya, Atsushi , Nishimichi, Takahiro , and Saito, Shun . Baryon acoustic oscillations in 2D: Modeling redshift-space power spectrum from perturbation theory. *Phys. Ref. D*, 82(6):063522, September 2010. doi:10.1103/PhysRevD.82.063522.

[260] Tinker, Jeremy L. . Redshift-space distortions with the halo occupation distribution - II. Analytic model. *MNRAS*, 374(2):477–492, January 2007. doi:10.1111/j.1365-2966.2006.11157.x.

[261] Tinker, Jeremy L. , Brownstein, Joel R. , Guo, Hong , Leauthaud, Alexie , Maraston, Claudia , Masters, Karen , Montero-Dorta, Antonio D. , Thomas, Daniel , et al. The Correlation between Halo Mass and Stellar Mass for the Most Massive Galaxies in the Universe. *ApJ*, 839(2):121, April 2017. doi:10.3847/1538-4357/aa6845.

[262] Tinker, Jeremy L. , Robertson, Brant E. , Kravtsov, Andrey V. , Klypin, Anatoly , Warren, Michael S. , Yepes, Gustavo , and Gottlöber, Stefan . The Large-scale Bias of Dark Matter Halos: Numerical Calibration and Model Tests. *ApJ*, 724(2):878–886, December 2010. doi:10.1088/0004-637X/724/2/878.

[263] Torrado, Jesús and Lewis, Antony . Cobaya: code for Bayesian analysis of hierarchical physical models. *J. Cosmology Astropart. Phys.*, 2021(5):057, May 2021. doi:10.1088/1475-7516/2021/05/057.

[264] Van Waerbeke, L. , Mellier, Y. , Erben, T. , Cuillandre, J. C. , Bernardeau, F. , Maoli, R. , Bertin, E. , McCracken, H. J. , et al. Detection of correlated galaxy ellipticities from CFHT data: first evidence for gravitational lensing by large-scale structures. *A&A*, 358:30–44, June 2000. doi:10.48550/arXiv.astro-ph/0002500.

[265] Verde, L. . Statistical Methods in Cosmology. In Wolschin, Georg , editor, *Lectures on Cosmology: Accelerated Expansion of the Universe*, volume 800, pages 147–177. Springer Berlin Heidelberg, 2010. doi:10.1007/978-3-642-10598-2_4.

[266] Wagoner, Robert V. , Fowler, William A. , and Hoyle, F. . On the Synthesis of Elements at Very High Temperatures. *ApJ*, 148:3, April 1967. doi:10.1086/149126.

[267] Wang, Lile , Reid, Beth , and White, Martin . An analytic model for redshift-space distortions. *MNRAS*, 437(1):588–599, January 2014. doi:10.1093/mnras/stt1916.

[268] Wechsler, R. H. , Zentner, A. R. , Bullock, J. S. , Kravtsov, A. V. , and Allgood, B. . The Dependence of Halo Clustering on Halo Formation History, Concentration, and Occupation. *ApJ*, 652:71–84, November 2006. doi:10.1086/507120.

[269] Wechsler, Risa H. and Tinker, Jeremy L. . The Connection Between Galaxies and Their Dark Matter Halos. *ARA&A*, 56:435–487, September 2018. doi:10.1146/annurev-astro-081817-051756.

[270] Weinberg, Steven . The cosmological constant problem. *Reviews of Modern Physics*, 61(1):1–23, January 1989. doi:10.1103/RevModPhys.61.1.

[271] Wittman, David M. , Tyson, J. Anthony , Kirkman, David , Dell'Antonio, Ian , and Bernstein, Gary . Detection of weak gravitational lensing distortions of distant galaxies by cosmic dark matter at large scales. *Nature*, 405(6783):143–148, May 2000. doi:10.1038/35012001.

[272] Wong, Kenneth C. , Suyu, Sherry H. , Chen, Geoff C. F. , Rusu, Cristian E. , Millon, Martin , Sluse, Dominique , Bonvin, Vivien , Fassnacht, Christopher D. , et al. H0LiCOW - XIII. A 2.4 per cent measurement of H_0 from lensed quasars: 5.3σ tension between early- and late-Universe probes. *MNRAS*, 498(1):1420–1439, October 2020. doi:10.1093/mnras/stz3094.

[273] Yamamoto, Kazuhiro , Nakamichi, Masashi , Kamino, Akinari , Bassett, Bruce A. , and Nishioka, Hiroaki . A Measurement of the Quadrupole Power Spectrum in the Clustering of the 2dF QSO Survey. *Publications of the Astronomicall Society of Japan*, 58:93–102, February 2006. doi:10.1093/pasj/58.1.93.

[274] York, Donald G. , Adelman, J. , Anderson, Jr. , John E., Anderson, Scott F. , Annis, James , Bahcall, Neta A. , Bakken, J. A. , Barkhouser, Robert , et al. The Sloan Digital Sky Survey: Technical Summary. *AJ*, 120(3):1579–1587, September 2000. doi:10.1086/301513.

[275] Yuan, Sihan , Garrison, Lehman H. , Eisenstein, Daniel J. , and Wechsler, Risa H. . Stringent σ_8 constraints from small-scale galaxy clustering using a hybrid MCMC + emulator framework. *MNRAS*, 515(1):871–896, September 2022. doi:10.1093/mnras/stac1830.

[276] Yuan, Sihan , Garrison, Lehman H. , Hadzhiyska, Boryana , Bose, Sownak , and Eisenstein, Daniel J. . ABACUSHOD: a highly efficient extended multitracer HOD framework and its application to BOSS and eBOSS data. *MNRAS*, 510(3):3301–3320, March 2022. doi:10.1093/mnras/stab3355.

[277] Zhai, Zhongxu , Tinker, Jeremy L. , Banerjee, Arka , DeRose, Joseph , Guo, Hong , Mao, Yao-Yuan , McLaughlin, Sean , Storey-Fisher, Kate , et al. The Aemulus Project. V. Cosmological Constraint from Small-scale Clustering of BOSS Galaxies. *ApJ*, 948(2):99, May 2023. doi:10.3847/1538-4357/acc65b.

[278] Zhai, Zhongxu , Tinker, Jeremy L. , Becker, Matthew R. , DeRose, Joseph , Mao, Yao-Yuan , McClintock, Thomas , McLaughlin, Sean , Rozo, Eduardo , et al. The Aemulus Project. III. Emulation of the Galaxy Correlation Function. *ApJ*, 874(1):95, March 2019. doi:10.3847/1538-4357/ab0d7b.

[279] Zhai, Zhongxu , Tinker, Jeremy L. , Hahn, ChangHoon , Seo, Hee-Jong , Blanton, Michael R. , Tojeiro, Rita , Camacho, Hugo O. , Lima, Marcos , et al. The Clustering of Luminous Red Galaxies at z $\sim$ 0.7 from EBOSS and BOSS Data. *ApJ*, 848(2):76, October 2017. doi:10.3847/1538-4357/aa8eee.

[280] Zhao, Cheng , Chuang, Chia-Hsun , Bautista, Julian , de Mattia, Arnaud , Raichoor, Anand , Ross, Ashley J. , Hou, Jiamin , Neveux, Richard , et al. The completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: 1000 multi-tracer mock catalogues with redshift evolution and systematics for galaxies and quasars of the final data release. *MNRAS*, 503(1):1149–1173, May 2021. doi:10.1093/mnras/stab510.

[281] Zheng, Zheng , Berlind, Andreas A. , Weinberg, David H. , Benson, Andrew J. , Baugh, Carlton M. , Cole, Shaun , Davé, Romeel , Frenk, Carlos S. , et al. Theoretical Models of the Halo Occupation Distribution: Separating Central and Satellite Galaxies. *ApJ*, 633(2):791–809, November 2005. doi:10.1086/466510.

[282] Zheng, Zheng , Coil, Alison L. , and Zehavi, Idit . Galaxy Evolution from Halo Occupation Distribution Modeling of DEEP2 and SDSS Galaxy Clustering. *ApJ*, 667(2):760–779, October 2007. doi:10.1086/521074.

[283] Zhou, Rongpu , Dey, Biprateep , Newman, Jeffrey A. , Eisenstein, Daniel J. , Dawson, K. , Bailey, S. , Berti, A. , Guy, J. , et al. Target Selection and Validation of DESI Luminous Red Galaxies. *AJ*, 165(2):58, February 2023. doi:10.3847/1538-3881/aca5fb.

[284] Zwicky, F. . On the Masses of Nebulae and of Clusters of Nebulae. *ApJ*, 86:217, October 1937. doi:10.1086/143864.

APPENDICES

Appendix A

Generating PIP weights using catalogue rotation

Pairwise-Inverse-Probability (PIP) weights can be used to correct a set of pair counts to match a baseline state, where the weight for a given pair is the inverse of the selection probability of that pair relative to that baseline state. In the case of an observed sample affected by fibre collisions where there is a non-zero probability of each possible pair in the parent catalogue being selected for the observed sample, applying PIP weights to the observed sample using the parent catalogue selection probabilities will cause it to match the original parent catalogue, removing the effect of fibre collisions entirely.

This ideal scenario is difficult to achieve for eBOSS because the majority of the footprint is only covered by a single spectroscopic plate. In these regions it is not possible to observe both members of a collided pair, so those collided pairs have zero probability of being selected. It is still possible to apply PIP weights, but they will only correct the pair counts to a baseline where these single-pass collided pairs are missed. While it is possible to account for these missing pairs using angular upweighting, which is the final solution used in eBOSS and presented in Sec. 2.5.3, our first solution was to construct a baseline where each pair in the parent catalogue has some probability of being in a multi-pass region, and so has a non-zero selection probability.

Building such a picture is non-trivial, because the eBOSS instrument pointings are fixed and irregular due to shifts made to maximize the number of observed spectra. Rather than modifying the tile placement algorithm in a way that would make it deviate from the actual observations, equivalent again to choosing a different baseline, I investigate generating PIP weights by applying random rotations to the target catalogue, thus giving a non-zero

probability that any given pair could be found in a multi-pass region and have a non-zero probability of observation.

The field of view of each plate in the eBOSS survey is 3' in diameter, and plates are arranged so that all plates overlap with at least one neighbour (see Sec. 2.2.3 for details), meaning that the angular distance between any target and the nearest multi-pass region is $\theta < 3'$. I therefore construct my rotation scheme to choose a random rotation angle uniformly between $-3' \leq \theta_{\rm rot} < 3'$. To determine the angle of the rotation I first determine a unit vector to the centre of the catalogue being rotated, $\boldsymbol{v}_{\rm cen}$, using $\theta_{\rm cen} = (\theta_{\rm max} + \theta_{\rm min})/2$ for both the Right Ascension (RA) and Declination (Dec) limits of the catalogue. I then define a vector on the unit sphere perpendicular to $\boldsymbol{v}_{\rm cen}$. By rotating that vector by a random angle between $0-2\pi$, I obtain a random axis of rotation, $\boldsymbol{v}_{\rm rot}$, that is perpendicular to the centre of the catalogue. It is worth noting that it is not strictly necessary to use a vector perpendicular to the catalogue footprint as the rotation axis, since even small rotations can still be valid realizations of the selection ensemble. I choose to use a perpendicular rotation axis to ensure that each rotation is impactful, reducing the number of targeting runs required to generate signal-dominated weights.

After determining the rotation axis and angle I perform the rotation by converting the RA and Dec positions of the target catalogues to Cartesian coordinates on the unit sphere:

$$\begin{aligned} x &= \cos(\mathrm{RA})\cos(\mathrm{Dec})\,, \\ y &= \sin(\mathrm{RA})\cos(\mathrm{Dec})\,, \\ z &= \sin(\mathrm{Dec})\,, \end{aligned} \tag{A.1}$$

and rotating them by applying Rodrigues' rotation formula:

$$\boldsymbol{v}_i' = \boldsymbol{v}_i\cos(\theta_{\rm rot}) + (\boldsymbol{v}_{\rm rot}\times\boldsymbol{v}_i)\sin(\theta_{\rm rot}) + \boldsymbol{v}_{\rm rot}(\boldsymbol{v}_{\rm rot}\cdot\boldsymbol{v}_i)(1-\cos(\theta_{\rm rot})), \tag{A.2}$$

where $\boldsymbol{v}_i$ is the vector to the i-th object in the target catalogue, and $\boldsymbol{v}_i'$ is the rotated vector to that object. I then convert the rotated vector back to equatorial coordinates from Cartesian coordinates using:

$$\begin{aligned} R &= \sqrt{x^2+y^2+z^2}\,, \\ \mathrm{RA} &= \arctan\left(\frac{y}{x}\right)\,, \\ \mathrm{Dec} &= \arcsin\left(\frac{z}{R}\right)\,. \end{aligned} \tag{A.3}$$

An example rotation is shown in Fig. A.1.

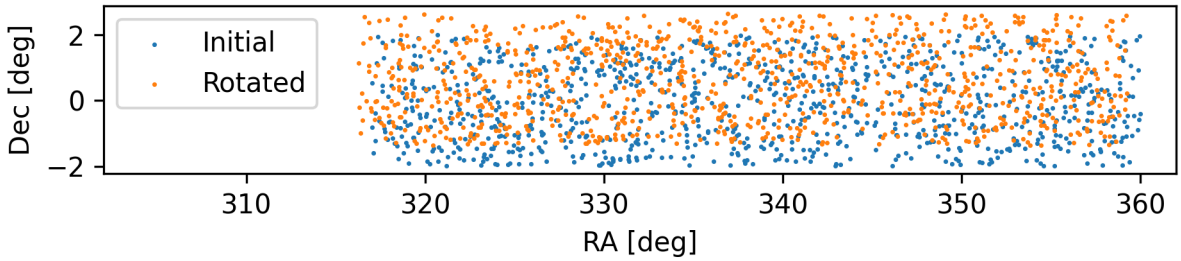

Figure A.1: Sample rotation for 1000 objects randomly selected from the **eboss21** ELG chunk. Blue points show the initial positions of the objects, and orange points show their rotated positions.

It it is important to note that objects on the edge of the survey footprint will be frequently rotated outside of the survey footprint, and thus be ineligible to receive a fibre. So that these realizations do not unrealistically lower the observation probability for objects close to masked regions, when using this method I modify Eq. 2.6 to be:

$$w_{mn} = \frac{\text{popcnt}\left[p_m^{(b)} \& p_n^{(b)}\right]}{\text{popcnt}\left[w_m^{(b)} \& w_n^{(b)}\right]}, \tag{A.4}$$

where $p_m^{(b)}$ and $p_n^{(b)}$ are bitwise records of when targets m and n were within the survey footprint and possible to observe. Like $w_m^{(b)}$ and $w_n^{(b)}$ they are binary arrays of length N_{runs}, where each bit (either one or zero) represents the possibility of assigning a fibre to that target (either this target is possible to observe, or not possible to observe in run b). $\text{popcnt}\left[p_m^{(b)} \& p_n^{(b)}\right]$ then calculates the number of runs where it is possible to observe both objects.

I generate a separate set of PIP weights for the 100 EZmock catalogues with realistic fibre collisions described in Sec. 2.5.4 by following the method outlined in Sec. 2.5.3, but applying a random rotation to the target catalogues before all but one of the target selection runs. For the remaining run I do not apply any rotation and use the actual random seed of the eBOSS targeting runs, ensuring that the realization used to select the actual eBOSS targets is part of the ensemble. Using this method I am able to generate non-zero observation probabilities for all pairs.

However, when comparing the PIP corrected clustering to the parent catalogue for the EZmocks I found a deviation from the expected clustering. Upon investigation, I

determined that targets around the edge of the survey have an excess probability of being observed compared to the completeness of the actual survey. The reason for this difference is that when a plate is only partially covered by targets the fibre assignment algorithm will assign all fibres to the covered area, increasing the completeness compared to a plate that has the same density of targets but is fully covered. While this situation occurs to a small degree for many edge plates in the actual eBOSS survey, it is much more common in my rotation scheme that frequently shifts targets beyond the edge of the footprint, leading to plates that have very low coverage by targets, and thus very high completeness. As discussed at the start of this section, this discrepancy does not mean that the PIP weights are invalid, but the baseline they have corrected the pair counts to is not the same baseline as the parent catalogue of the targeting sample, meaning the rotation PIP-weighted catalogue is not appropriate for the cosmological analysis, despite having fibre collisions removed.

This difference in baselines prevents us from using rotation to correct for the zero-probability pairs in single-pass regions. However, as was discussed in Sec. 2.5.3 and shown in Sec. 2.5.4, we are instead able to account for the missing collided pairs using angular upweighting. We generate PIP weights using the probability from the eBOSS target selection algorithm without rotation, which corrects the sample so that single-pass region pairs above the fibre-collision scale and multi-pass region pairs on all scales match the baseline of the target catalogue. Angular upweighting increases the weight of the multi-pass pairs below the fibre collision scale to account for the missing pairs in single pass regions, so under the assumption that the collided pairs in the multi-pass regions and single-pass regions are statistically similar, our method is able to match the baseline of the target catalogue. As a result, applying our PIP+ANG weights produces a catalogue with the effect of fibre collisions completely removed while matching the baseline of the parent catalogue.

Appendix B

Correcting redshift differences between an emulator and data

The original AEMULUS emulator adapted for the small-scale RSD analysis of the eBOSS LRG sample in Ch. 3 was first designed to match a BOSS CMASS-like sample at $z = 0.57$ and space density $n = 4.2 \times 10^{-4}[h^{-1}\mathrm{Mpc}]^{-3}$. These choices were built into the training of the emulator, which used the $z = 0.57$ redshift slices of the AEMULUS simulations, and fixed M_{min} across the training HOD parameter sets to give a constant number density of $n = 4.2 \times 10^{-4}[h^{-1}\mathrm{Mpc}]^{-3}$. However, our eBOSS sample is at an effective redshift of $z = 0.737$ and peak number density of $n = 9 \times 10^{-5}$. Ultimately, this difference required us to retrain the emulator using the $z = 0.7$ redshift slices and updated values of M_{min} to give the correct value of n (see Sec. 3.3.2). However, before making that choice we investigated if it was possible to analyze a galaxy sample using an emulator trained at a different redshift by altering the cosmological parameters given to the emulator.

The differences between the original AEMULUS emulator and the eBOSS LRG sample present two major concerns. The first is the difference in redshift, which means that the data and model have experienced different amounts of linear growth, changing the amplitude of perturbations at all scales by the ratio of the growth factor, $D(z)$, between the two redshifts, and different amounts of non-linear growth that will primarily affect the small-scale clustering. However, we can allow for this difference in redshift by adjusting the parameters provided to the emulator away from those in the model being tested and shifting the resulting correlation functions to correct for the adjustment.

First, we note that a slice of the Universe when observed at $z_{\mathrm{data}} = 0.737$ would appear to be at a redshift of $z_{\mathrm{emul}} = 0.57$ for an observer located at $z_{\mathrm{obs}} = 1.737/1.57 - 1$ due to

the ratio of the scale factors:

$$\begin{aligned} \frac{a_{\rm data}}{a_{\rm obs}} &= \frac{a_{\rm emul}}{a_0} \\ \frac{1+z_{\rm obs}}{1+z_{\rm data}} &= \frac{1}{1+z_{\rm emul}} \\ z_{\rm obs} &= \frac{1.737}{1.57} - 1\,. \end{aligned} \tag{B.1}$$

Therefore, we can input to the emulator the parameters of the Universe at $z_{\rm obs}$, denoted by the subscript "emul", to match the evolution of the two cosmologies at different redshifts. Specifically, this requires setting the input matter density to be $\Omega_{m,\rm emul} = \Omega_m(z_{\rm obs})$ and the physical matter density to be

$$\Omega_{m,\rm emul}h^2_{\rm emul} = \left(\frac{1.737}{1.57}\right)^3 \Omega_m h^2\,, \tag{B.2}$$

while keeping the baryon fraction constant,

$$\Omega_{b,\rm emul} = \Omega_{m,\rm emul}\left(\frac{\Omega_m}{\Omega_b}\right)\,, \tag{B.3}$$

and CMB temperature constant,

$$T_{\rm CMB,emul} = (1+z_{\rm obs})T_{\rm CMB,emul}\,. \tag{B.4}$$

The output from the emulator would now be appropriate for analyzing a data sample at $z = 0.737$ - with exactly the same cosmological evolution expected for the emulator training simulations run to $z = 0.57$ as in the model being tested to $z = 0.737$ (e.g. the age of the Universe at $z = 0.57$ for the cosmological model given to the emulator matches the age at $z = 0.737$ in the model being compared to the data).

Note that in making this adjustment the length scales need to be shifted to match the definition of comoving units for the data cosmology,

$$r'_{\rm emul} = r_{\rm emul} \times \frac{H_0}{H(z_{\rm obs})} \times \frac{1+z_{\rm data}}{1+z_{\rm emul}}\,, \tag{B.5}$$

$$k'_{\rm emul} = k_{\rm emul} \times \frac{H(z_{\rm obs})}{H_0} \times \frac{1+z_{\rm emul}}{1+z_{\rm data}}\,, \tag{B.6}$$

$$P'(k'_{\rm emul}) = P(k'_{\rm emul}) \times \left(\frac{k_{\rm emul}}{k'_{\rm emul}}\right)^3 . \tag{B.7}$$

Although the correlation function is dimensionless and hence we do not need to change the normalisation due to the change in scale, we do need to be careful with the value of $\sigma_8(z=0)$ given to the emulator as an input parameter. There are two factors to consider: the effect of the change in scale on the normalisation scale of $8\,h^{-1}$Mpc, and the different evolution to $z=0$ where $\sigma_8(z=0)$ is normalised. The emulator cosmology therefore requires the normalization value,

$$\sigma_{8,emul} = \sigma_R(z=0) \times \frac{D(z=0.737)}{D(z=0)} \times \frac{D_{\rm emul}(z=0)}{D_{\rm emul}(z=0.57)} . \tag{B.8}$$

Note that many codes calculate the linear growth factor $g = D/a$, so extra scale factors are required to use these calculations. The first ratio corrects from $z=0$ to $z=0.737$ in the *Planck* cosmology, and the second ratio corrects from $z=0.57$ to $z=0$, where the normalisation is defined in the emulator cosmology. Because of the shift in scales, we also have to shift the scale at which the linear power spectrum is normalised, assuming $R = 8 \times H_0/H(z_{\rm obs}) \times 1.737/1.57$. Using this conversion, the emulator results would be correlation functions of the expected amplitude for the true cosmology being tested at $z=0.737$.

We can simplify this equation by noting that the *Planck* and AEMULUS cosmologies are identical up to $z_{\rm obs}$, so we have that

$$\frac{D_{\rm Aem}(z=0)}{D_{\rm Aem}(z=0.57)} = \frac{D(z=z_{\rm obs})}{D(z=0.737)} . \tag{B.9}$$

Substituting this into Eq. B.8 gives that

$$\sigma_{8,Aem} = \sigma_R(z=0) \times \frac{D(z_{\rm obs})}{D(z=0)} . \tag{B.10}$$

Using this conversion, AEMULUS results will be correlation functions of the expected amplitude for the true cosmology being tested at $z=0.737$.

A comparison of the comoving linear power spectrum for the data and emulator cosmologies is shown in Fig. B.1, with corresponding cosmological parameters and ages given in Table B.1. To make the differences between lines more apparent for this example we use $z_{\rm data} = 1.0$ and $z_{\rm emul} = 0.5$. The adjusted cosmology matches the data cosmology power

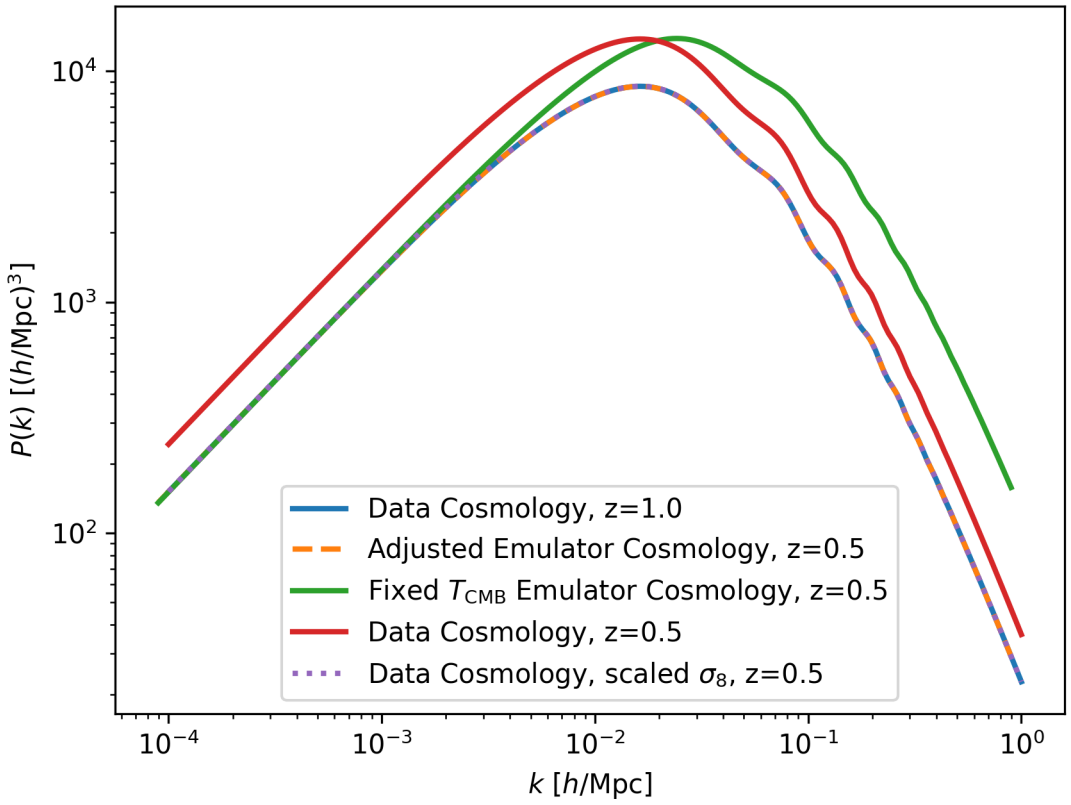

Figure B.1: Demonstration of adjusting the cosmological parameters to match the emulator comoving linear matter power spectrum and age to the data. The blue line shows the power spectrum for the data cosmology at the redshift of the data. The dashed orange line shows the power spectrum for the adjusted cosmology given to the emulator at the emulator redshift. The green line shows the power spectrum for the emulator if T_{CMB} can not be changed from the default value. The red line shows the data cosmology at the emulator redshift. The dotted purple line is the same data cosmology at the emulator redshift, but with σ_8 scaled to match the amplitude of the data power spectrum at the data redshift. The detailed cosmological parameters and universe age at the given redshift for each line is displayed in Table B.1.

Line	H_0 $[kms^{-1}Mpc^{-1}]$	Ω_m	Ω_b	$T_{\rm CMB}$	z	Age [Gyr]
Data Cosmology	67.0	0.310	0.496	2.7255	1.0	5.92
Adjusted Emulator Cos.	80.0	0.516	0.8256	3.634	0.5	5.92
Fixed $T_{\rm CMB}$ Emulator	80.0	0.516	0.8256	2.7255	0.5	5.92
Data Cos. at $z = 0.5$	67.0	0.310	0.496	2.7255	0.5	8.69

Table B.1: Detailed cosmological parameters and universe age at the given redshift for each line in Fig. B.1.

spectrum exactly despite the difference in redshift, and because both universes have the same age they will have experienced the exact same amount of non-linear growth.

While the adjustment detailed above would provide an exact match between the data and emulator at different redshifts, not all of the parameter changes are possible for the AEMULUS cosmological emulator. Since $T_{\rm CMB}$ has been measured to such high precision by *Planck* and other CMB probes it is fixed in the emulator, meaning we are not able to modify the model cosmology to exactly match between the data and model. The effect of keeping $T_{\rm CMB}$ fixed while changing the other parameters is also shown in Fig. B.1, with significant deviations at intermediate and large k.

Instead, we can provide the emulator with the $z = 0$ cosmological parameters of the cosmology being tested, but adjust the value of σ_8 as described above to give the correct normalisation. This correction is also shown in Fig. B.1, with the unscaled data cosmology at the emulator cosmology shown for reference. Using only the σ_8 scaling, the comoving linear power spectrum at the emulator redshift also matches the power spectrum at the data redshift exactly. However, unlike the full adjustment, the universe ages are no longer the same at the different redshifts, as shown in Table B.1, meaning the amount of non-linear growth will be different between the emulator cosmology and the cosmology being tested for the data.

Although the age and evolution to the output redshift emulated by AEMULUS does not match that of the cosmology being tested, the linear power spectrum does, which should still allow for this change to give robust results for the following reasons: for an Einstein-de Sitter model (i.e. a flat, matter only universe), where there is a symmetry between growth and evolution, the match would be perfect - the amplitude of the growth completely defines the cosmological evolution to that point. This is not the case for cosmological models where Λ is present because this symmetry breaks down. However, the redshifts of interest in the eBOSS analysis are close to the epoch of matter-Λ equality, and so the effect of Λ on the evolution to the redshifts of interest small. The halo mass function in standard theory depends only on the linear power spectrum, and the HOD model used in the analysis for

small-scale clustering is based on this halo mass function. Thus we expect that the model can accurately match the small-scale clustering, and so the only concern is small changes in the interpolations between large and small scales. Finally, we note that the value of f input to AEMULUS is the value in the true cosmological model at $z_{\rm eff} = 0.737$: this is a dimensionless parameter calculated at the redshift of interest, and has the property that it is the same for observers at different redshifts.

The second issue is the difference in number density. It is particularly worrying since a less dense sample will preferentially fill more massive halos. The result will be a sample with a larger linear bias, which is degenerate with the growth rate in clustering measurements (see Sec. 3.4.2 for how an incorrect linear bias can influence the inferred cosmological results). In order to assess the impact of the difference in number density, as well as test our cosmological scaling method, we use a set SHAM mocks generated from the Uchuu simulation. We generate a mock using the $z = 0.7$ redshift slice of the simulation and number density $n = 1 \times 10^{-4}$, matching the eBOSS redshift and number density. We then generate two additional mocks, one from the $z = 0.57$ redshift slice with $n = 1 \times 10^{-4}$ and one from the $z = 0.7$ slice with $n = 4 \times 10^{-4}$, to investigate the dependence on each characteristic separately. All mocks are constructed using the peak halo velocity, $V_{\rm peak}$, with a scatter of 0.2 dex.

To test the ability of our emulator to model these different conditions we fit the emulator predictions to the correlation function measurements from the various SHAM mocks using a MCMC sampler (see Sec. 3.3.6), and compare the $f\sigma_8$ constraints to the known values for the simulations. The results of these fits are shown in Fig. B.2. The fit to the $z = 0.7$, $n = 4 \times 10^{-4}$ mock is able to recover the expected $f\sigma_8$, showing that the emulator is robust to changes in redshift without biasing constraints of $f\sigma_8$. However, the $n = 1\times10^{-4}$ mocks for both $z = 0.57$ and $z = 0.7$ give significantly larger values of $f\sigma_8$ compared to the known value, as expected from the degeneracy between linear bias and $f\sigma_8$. This discrepancy means that while our method for analyzing a galaxy sample using an emulator trained at a different redshift can be applied for samples with the same number density, it can not be used for our analysis of the eBOSS LRG sample with the original AEMULUS emulator. As a result we elect to retrain the emulator to match the eBOSS LRG sample, as detailed in Sec. 3.3.2. After retraining the emulator we repeat the fit to the $z = 0.7$, $n = 1 \times 10^{-4}$ SHAM mock using the new emulator (see Sec. 3.4.4), and find robust constraints on all cosmological parameters.

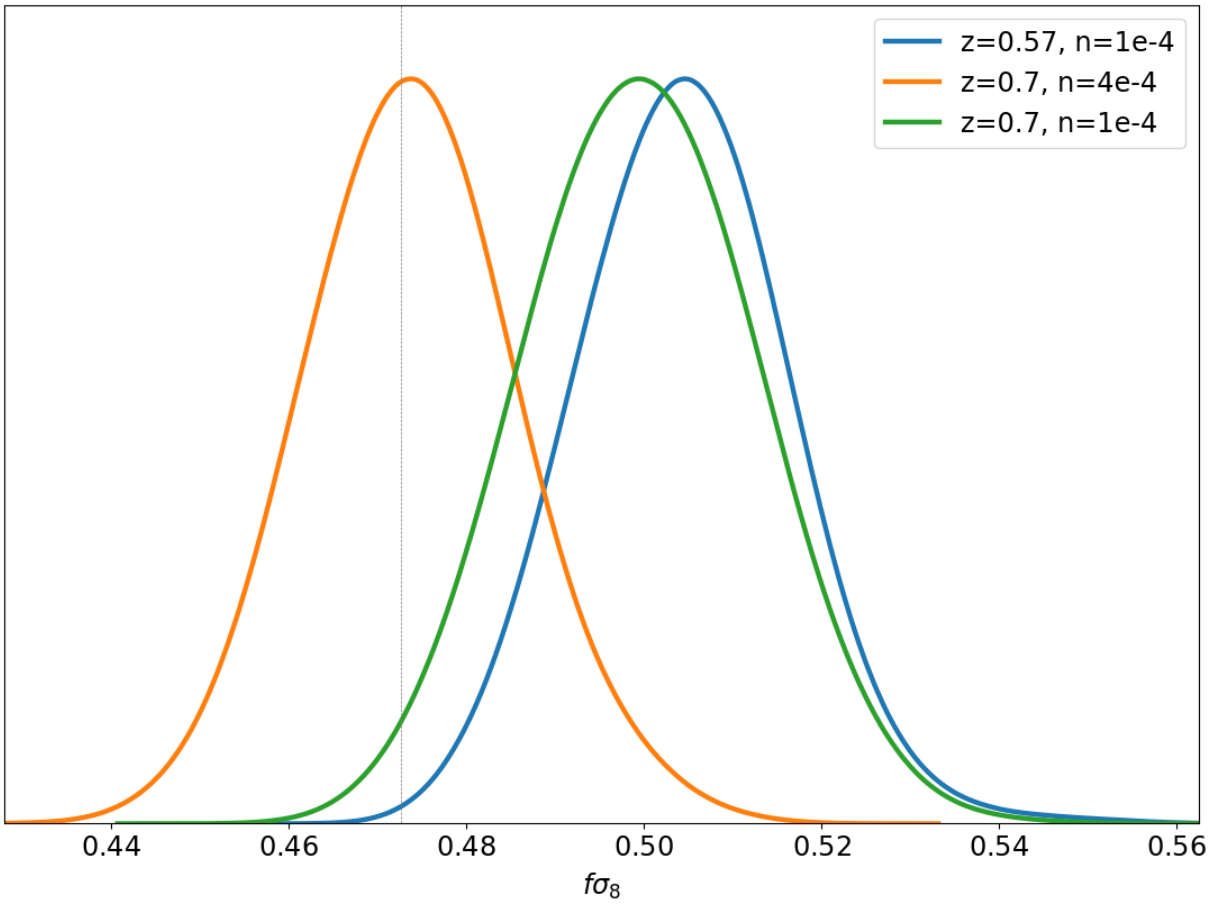

Figure B.2: $f\sigma_8$ constraints from fitting the AEMULUS cosmological emulator to SHAM mocks constructed from the Uchuu simulation at various redshifts and number densities. The vertical line shows the expected value of $f\sigma_8$ at the emulator redshift.

Appendix C

Linear and non-linear velocity distributions

In order to test our calculation of the linear velocity in Ch. 4 and to observe the relative impacts of the linear and non-linear components as a function of separation we examine the mean pairwise velocity of halos as a function of pair separation. We use the halos of the 20 boxes of the AbacusCosmos simulation suite with *Planck* 2015 cosmologies described in Sec. 4.2.2, at $z = 0.700$. For each simulation we divide the halos into mass bins of width 0.5 dex from $10^{12} - 10^{15} M_{\odot}/h$, and calculate the mean pairwise velocity for each mass interval in 80 separation bins of equal logarithmic width from $10^{-2} - 10^{2}\, h^{-1}$Mpc. We perform this calculation for both the unsmoothed and smoothed halo velocities (see Sec. 4.2.4). These pairwise velocities are shown in Fig. C.1. Also included in the plot is the static solution, which is the pairwise velocity required to maintain a constant proper separation in an expanding background, and the linear theory prediction for the pairwise velocity [99, 227, 45]. It should be noted that the linear theory prediction contains a bias factor that we have set to 1 in the calculation for all halo masses, leading to a difference in amplitude for the high mass halo bins.

As expected, at large separation the linear velocity calculated from the initial conditions is in good agreement with the total halo velocity, providing a good test of the linear velocities calculated at high redshift as well as the scaling to low redshift. The shape of the pairwise velocity is also a good match to the expectation from linear theory above $\sim 20\, h^{-1}$Mpc. Below that scale the linear velocity begins to deviate from the total velocity, and goes towards zero at small scales. Because the non-linear velocity is defined as the difference between the total velocity and the linear component this leads to a larger non-linear component, which peaks around $1\, h^{-1}$Mpc after crossing the static solution, before

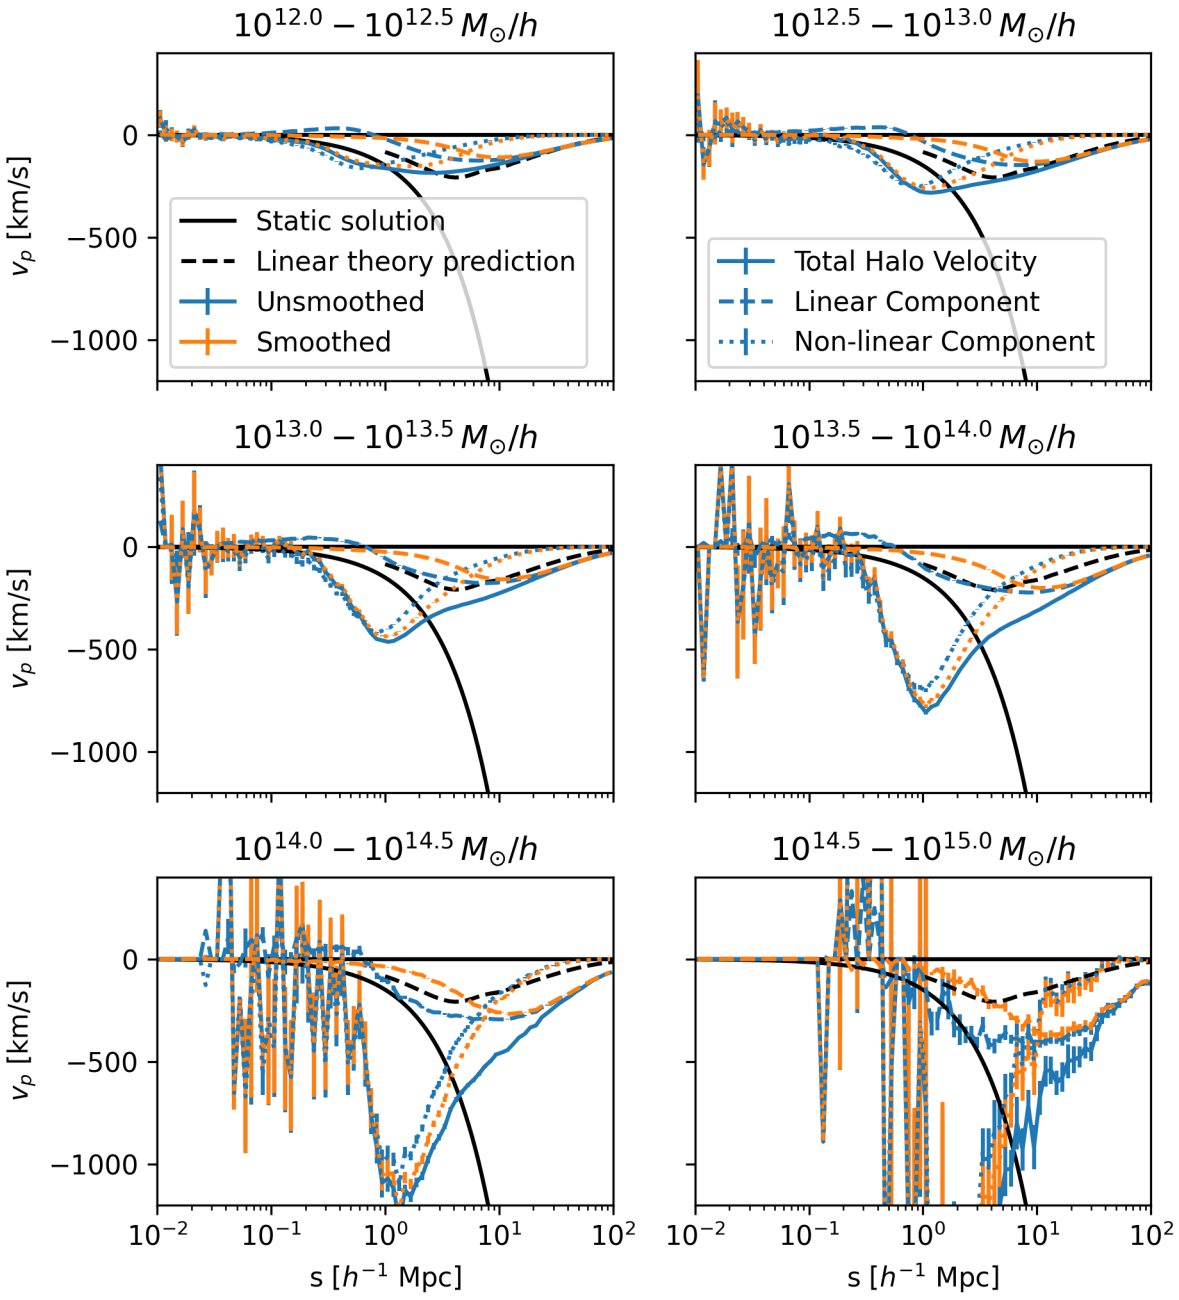

Figure C.1: Mean pairwise velocities as a function of separation for the halos of the 20 AbacusCosmos *Planck* simulation boxes, split into halo mass bins.

trending towards zero. The unsmoothed linear velocity becomes positive, indicating a pair increasing in separation, below $\sim 0.5\,h^{-1}$Mpc. This behaviour is counter intuitive for structure growth and accounts for the unexpected effect of the scaling parameter on the correlation function at small scales shown in Fig. 4.2. Smoothing the linear velocity field causes it to turn towards zero below the smoothing scaling, without crossing to positive velocities at small separations.

We also investigate the magnitudes of the total, linear, and non-linear halo velocities as a function of halo mass for the unsmoothed and smoothed velocity fields in Fig. C.2. Total velocity increases with halo mass, caused by an increase in non-linear velocity. The linear velocity dominates for small halo masses, but decreases slightly with halo mass. These trends are seen for both the unsmoothed and smoothed velocities, with the only significant difference being a slight decrease in the magnitude of the linear component from smoothing.

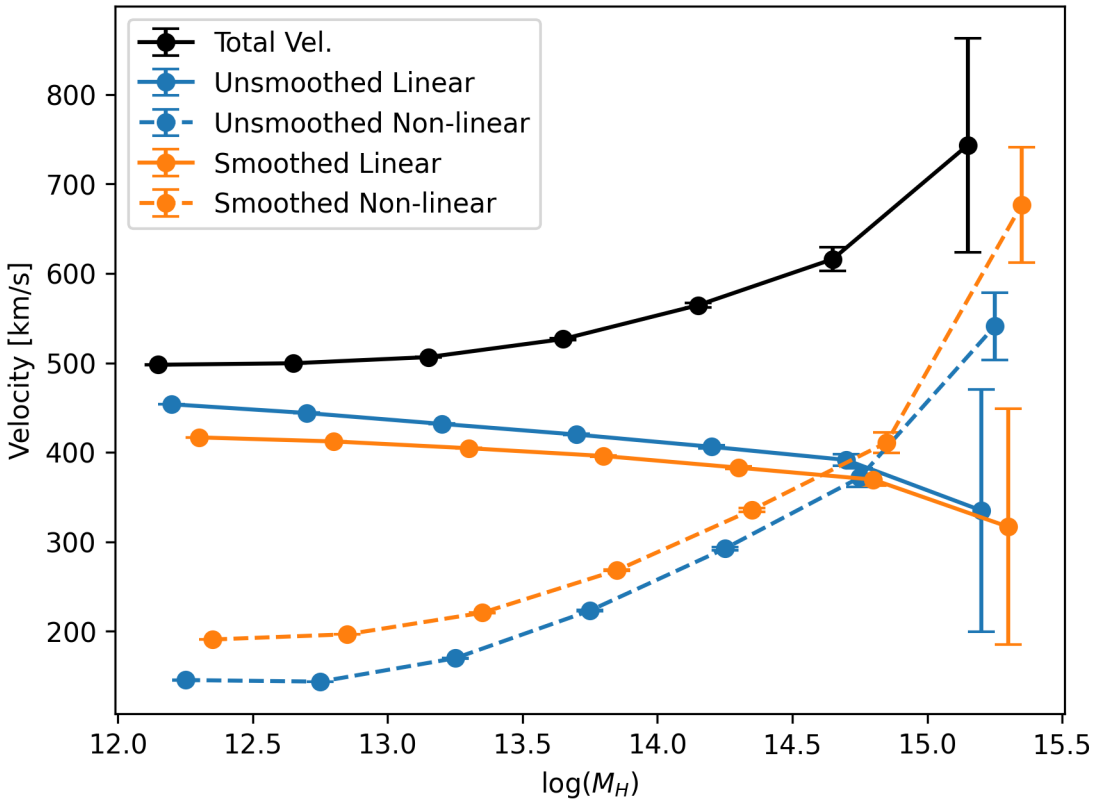

Figure C.2: Mean halo velocity magnitude in mass bins of width 0.5 dex. The black line shows the mean total halo velocity, while solid coloured lines show the mean linear velocity magnitude and dashed coloured lines show the mean non-linear velocity magnitude. Blue lines represent the result using the unsmoothed linear velocity, while orange lines show the result after our fiducial velocity smoothing. The position of points on the x-axis are offset slightly to prevent overlap.